Master Computer Science

Novelty-promoting behaviour in agents: Curiosity and Diversity

Name: Shreyansh Sharma
Student ID: S3772241

Date: 30/01/2025

Specialisation: Artificial Intelligence

1st supervisor: Rob Saunders
2nd supervisor: Tessa Verhoef

Master’s Thesis in Computer Science

Leiden Institute of Advanced Computer Science
Leiden University
Niels Bohrweg 1
2333 CA Leiden
The Netherlands

Abstract

Novelty is a key driver of innovation, acting as the catalyst that propels advancements across various domains, including technology, economics, and artificial intelligence. This paper explores how novelty influences innovation within Multi-Agent Systems (MAS) using Reinforcement Learning (RL) frameworks. By categorizing novelty into two components—Curiosity and Diversity—we examine their impact on exploration and role specialization. We propose an integrated methodology that combines Intrinsic Curiosity Modules (ICM) and Diversity Control (DiCo), providing insights into their interaction and contribution to innovation. Experimental results in two distinct environments highlight how these novelty components drive emergent behaviors, offering implications for learning efficiency and system robustness.

1. Introduction

Measuring novelty and its impact on innovation poses significant challenges. Novelty is a multifaceted concept that can manifest at individual and collective levels, influencing both personal exploration and group dynamics. By dissecting novelty into quantifiable components and tracking their influence over time, we can better understand its role in driving innovation within complex systems.

Multi-agent systems (MAS) have become increasingly important in artificial intelligence research, offering unique insights into complex interactions, emergent behaviors, and collective problem-solving. Recent advancements in multi-agent systems, language models, and reinforcement learning have opened new avenues for exploring creativity and innovation in artificial systems (Lazaridou et al., 2016; Eecke et al., 2023).

Agents operating within MAS can exhibit behaviors influenced by novelty, leading to emergent innovation that mirrors real-world processes (Leibo et al., 2019). Reinforcement Learning (RL) provides the tools for agents to learn and adapt based on interactions with their environment, making it an ideal framework for modeling and measuring the influence of novelty on innovation.

Innovation drives technological and societal progress, with novelty as its cornerstone. However, understanding and quantifying novelty's role in fostering innovation in artificial systems remains challenging. This research investigates novelty in Multi-Agent Systems (MAS), where agents interact and adapt, simulating complex systems. Using Reinforcement Learning (RL), we measure how curiosity and diversity—two key aspects of novelty—affect individual and collective behaviors. Specifically, this research aims to answer the question: How does novelty, quantified through curiosity and diversity, influence innovation in Multi-Agent Systems (MAS) using Reinforcement Learning (RL). The findings aim to bridge gaps in modeling innovation and provide tools for analyzing emergent intelligence.

2. Literature Review

2.1. The Role of Novelty in Innovation

Novelty search has been widely recognized as a key driver of innovation, stimulating exploration and the discovery of new knowledge or strategies. In the context of MAS and RL, novelty search algorithms encourage agents to explore the state-action space beyond the objective of reward maximization (Aydeniz et al., 2023). Recent studies have applied novelty search to address the exploration-exploitation dilemma in RL, demonstrating its effectiveness in avoiding local optima and promoting the discovery of diverse behaviors (Ecoffet et al., 2021).

Novelty is inherently linked to innovation, representing the introduction of new elements—ideas, products, or behaviors—that differ from existing ones. Studies in innovation diffusion within multi-agent systems have provided valuable insights into how new ideas and behaviors spread through agent interactions (Ting, 2006; Zhang & Vorobeychik, 2019). These studies underscore the potential of multi-agent systems to model and analyze the dynamics of innovation diffusion, contributing to our understanding of how innovative behaviors emerge and propagate in decentralized environments. This research focuses on the interplay between individual curiosity and group diversity, two facets of novelty, to model innovation diffusion in MAS.

2.2. Measuring Novelty: Curiosity and Diversity

Novelty can be characterised through two distinct components that reflect both individual and collective aspects of innovation: Curiosity and Diversity. These components were chosen based on their relevance to both biological and artificial systems. Curiosity, inspired by cognitive models of learning, drives individual agents to explore new and unfamiliar states, enhancing their ability to discover novel solutions (Berlyne, 2014). Techniques such as Intrinsic Curiosity Modules (ICM) (Pathak et al., 2017) and Random Network Distillation (RND) (Burda et al., 2018) provide methods for incorporating curiosity into RL agents.

Diversity, on the other hand, promotes a range of behaviors across agents, preventing homogenization and encouraging the exploration of varied strategies. Methods like promoting behavioral diversity (Parker-Holder et al., 2020) and sharing novel experiences among agents (Gerstgrasser et al., 2024) enhance group novelty within MAS. Measuring diversity can be achieved through metrics such as System Neural Diversity (SND) (Bettini et al., 2024b), which quantifies the behavioral differences among agent policies.

This combination allows us to study how individual exploration and group heterogeneity contribute to the emergence of innovative behaviors in multi-agent systems.

2.3. Multi-Agent Systems and Reinforcement Learning

Multi-agent systems consist of autonomous agents which observe their environment, reason about their actions, and interact to accomplish shared goals. These interactions can be competitive, collaborative, self-replicating, or mixed, leading to diverse emergent behaviors and solutions (Baker et al., 2020; y Arcas et al., 2024). Agents, termed adaptive units, learn and adapt their strategies to benefit themselves and their peers. These dynamics have been well documented across various domains, including biology, economics, sociology, and psychology.

Reinforcement Learning provides a framework for agents to learn optimal behaviors through environment interaction. In Multi-Agent Systems, RL enables agents to adapt not only to the environment but also to other agents' actions. This framework allows us to study how curiosity-driven exploration and diversity among agents contribute to innovation.

2.4. Motivation

Understanding the influence of novelty on innovation has several practical implications:

- **Enhancing Learning Efficiency:** By promoting novelty through curiosity and diversity, agents can discover innovative solutions more efficiently, leading to improved problem-solving capabilities.
- **Facilitating Cooperation and Role Specialization:** Diversity in agent behaviors can lead to better cooperation strategies and role specialization in MAS, essential for complex task completion (Bettini et al., 2024a).
- **Increasing Interpretability:** Tracking feature importance over time aids in understanding how agents prioritize different novelty components, enhancing transparency and explainability in AI systems.
- **Emergent Behavior Understanding:** Studying novelty in MAS provides insights into how complex behaviors emerge from simple interactions, contributing to our understanding of collective intelligence (y Arcas et al., 2024).

3. Methodology

3.1. MADDPG Architecture for Multi-Agent Reinforcement Learning

We adopt the Multi-Agent Deep Deterministic Policy Gradient architecture for training agents in cooperative environments (Lowe et al., 2020). MADDPG extends the DDPG algorithm to handle multiple agents, enabling centralized training with decentralized execution. Agents share a common critic network while maintaining separate actor networks, allowing them to learn individual policies that collectively optimize a shared objective.

3.2. Intrinsic Curiosity Modules (ICM)

Intrinsic Curiosity Modules introduce intrinsic rewards based on the prediction error of an agent's internal forward model (Pathak et al., 2017). The curiosity reward encourages agents to explore unfamiliar or unpredictable transitions in the environment.

The curiosity reward at time t is defined as:

$$r^{\text{curiosity}}(t) = \|\hat{s}_{t+1} - s_{t+1}\|^2 \tag{1}$$

where $\hat{s}_{t+1}$ is the predicted next state from the agent's internal model, and s_{t+1} is the actual next state observed.

3.3. Measuring Diversity: System Neural Diversity

System Neural Diversity (SND) is a metric used to quantify the diversity among agent policies in a multi-agent system. It measures the functional differences between agents' policies by computing the pairwise Wasserstein distance between their action distributions over a set of observations and then aggregating these distances to provide a system-level diversity measure.

1. Pairwise Diversity: For every pair of agents, compute the Wasserstein distance W_2 between their action distributions over a set of observations O:

$$d(\pi_i, \pi_j) = \frac{1}{|O|} \sum_{o \in O} W_2(\pi_i(o), \pi_j(o)) \tag{2}$$

2. Aggregate to System-Level Diversity:

$$SND(\{\pi_i\}_{i \in N}) = \frac{2}{N(N-1)} \sum_{i=1}^{N} \sum_{j=i+1}^{N} d(\pi_i, \pi_j) \tag{3}$$

where N is the number of agents.

3.4. Diversity Control (DiCo)

Diversity Control is a method designed to maintain desired levels of diversity among agent policies by scaling agent policies to achieve a desired diversity level quantified through SND (Bettini et al., 2024a). DiCo adjusts the balance between homogeneous (shared) policy components and heterogeneous (agent-specific) components, ensuring that diversity is controlled without altering the learning objective.

$$\pi_i(o) = \pi_h(o) + \frac{SND_{\text{des}}}{SND_{\text{current}}}\pi_{h,i}(o) \tag{4}$$

where $\pi_i(o)$ is the policy of agent i, $\pi_h(o)$ is the homogeneous policy component, $\pi_{h,i}(o)$ is the heterogeneous component, SND_{current} is the current System Neural Diversity, and SND_{des} is the desired diversity level.

3.5. Integrating Curiosity and Diversity

We propose an approach that integrates ICM and DiCo within the MARL framework to study their combined effect on innovation.

3.5.1. Integrated Reward Function

We define the overall reward for each agent as a combination of extrinsic rewards and curiosity-driven rewards:

$$R_{\text{total}} = R_{\text{extrinsic}} + \alpha R_{\text{curiosity}} \tag{5}$$

where:

- $R_{\text{extrinsic}}$ is the external reward from the environment. - $R_{\text{curiosity}}$ is the intrinsic curiosity reward as defined by ICM. - α is a scaling factor that adjusts the influence of curiosity on the agent's learning process.

Diversity on the other hand is maintained by directly scaling agent policies to achieve the desired diversity level as described in the Diversity Control (DiCo) method. Agents update their policies using the integrated reward, balancing exploration driven by curiosity with the maintenance of diversity among agent policies.

4. Experiments

This section describes the experimental setup designed to examine the effects of curiosity and diversity as potential novelty drivers in Multi-Agent Reinforcement Learning (MARL). Our objective is to understand their role individually and collectively in enhancing exploration, improving learning efficiency, and fostering innovative behavior in agents across two distinct environments.

4.1. Experimental Design

4.1.1. Selected Environments

To evaluate the impact of integrating curiosity and diversity in MARL agents, we selected two environments with distinct characteristics that highlight different aspects of agent coordination and exploration.

- **Balance Environment:** The Balance environment as seen in Figure 1 is designed as a cooperative task where agents must work together to balance an object (e.g., a plank or beam) and move it to a target location. Precise synchronization is required, emphasizing the need for homogeneous behavior among agents. This environment allows us to assess the effectiveness of diversity control mechanisms in situations where coordinated actions are critical and where curiosity-driven exploration may have limited impact due to a relatively small and well-defined state space.
- **Multiagent Navigation Environment:** The Multiagent Navigation environment as seen in Figure 2 presents a more complex, multi-objective task where agents navigate to dispersed target locations. Each agent has a unique target, requiring individual exploration to find the most efficient path while also requiring coordination to avoid collisions with other agents. This scenario necessitates higher levels of exploration and role specialization, making it an ideal test bed for evaluating both curiosity and diversity mechanisms.

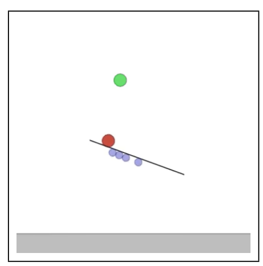

Figure 1. Balance Environment: Purple agents can be seen taking the platform with the red ball towards the green target.

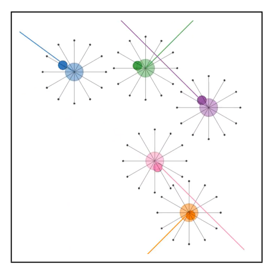

Figure 2. Multiagent Navigation Environment: Five agents (represented by different colors) have successfully navigated to their respective target locations.

4.2. Configurations and Methodology

We systematically tested the influences of curiosity and diversity by configuring the following agent setups in each environment:

- **MADDPG Baseline:** Serving as the control group, this configuration evaluates agents using the standard

MADDPG algorithm without any novelty incentives. Agents learn purely from extrinsic rewards provided by the environment.

- **Curiosity-Driven Agents:** Agents are attached with Intrinsic Curiosity Modules (ICM), receiving intrinsic rewards based on the prediction error of their internal forward model. This setup highlights the effect of intrinsic motivation on exploration potential. The α parameter is set to a fixed value of 0.1 for all experiments to measure a consistent level of curiosity influence.
- **Diverse Agents:** Agents use Diversity Control (DiCo) to maintain varying levels of System Neural Diversity (SND). We tested multiple desired diversity levels (SND_{des}) to assess how systemic diversity influences agent behavior and task performance.
- **Combined Strategy:** This configuration integrates both ICM and DiCo within the agents to determine the interactions between curiosity-driven exploration and diversity in enhancing overall performance.

5. Results

The following section discusses the experimental findings, providing insight into how curiosity and diversity influence MARL agents' behaviors, with implications for innovative capacity and learning dynamics.

5.1. Findings from the Balance Environment

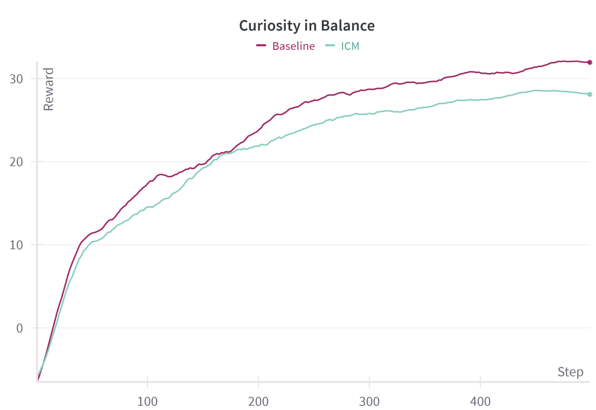

Figure 3. Comparison of Baseline vs. Curiosity-Driven Agents in the Balance Environment. The graph shows cumulative rewards over training episodes, highlighting the impact of intrinsic curiosity on exploration efficiency and task performance.

5.1.1. Curiosity Insights

Introducing ICM to the baseline configuration increased exploration in the balance environemnt. Curiosity-driven exploration initially slowed down the learning curve before reaching a plateau, which led to lower cumulative rewards compared to baseline agents, as shown in Figure 3.

5.1.2. Diversity Observations

Diversity control mechanisms exhibited nuanced effects on task performance. Agents with lower diversity levels ($SND_{\text{des}} = [0.2]$) achieved better cumulative rewards compared to curiosity-driven agents, indicating an optimal level of diversity can enhance learning efficiency without compromising coordination, as can be seen in Figure 4. Agents with $SND_{\text{des}} = [0.2]$ demonstrated faster learning in the early stages and achieved marginally higher cumulative rewards compared to the baseline. Where as agents with higher diversity level ($SND_{\text{des}} = [0.5]$) showed decreased performance in the early stages of training, possibly due to excessive behavioral divergence hindering the precise coordination required for the task, followed by a lower cumulative reward compared to the baseline.

From Figure 5, it can be observed that the SND value of the baseline agents, which do not use any scaling methods, remains close to 0.2 over time. This helps confirm that the 'balance' environment requires lower diversity levels to achieve optimal performance. By using DiCo to constrain diversity levels to 0.2, the agents were able to achieve a higher cumulative reward compared to the baseline agents.

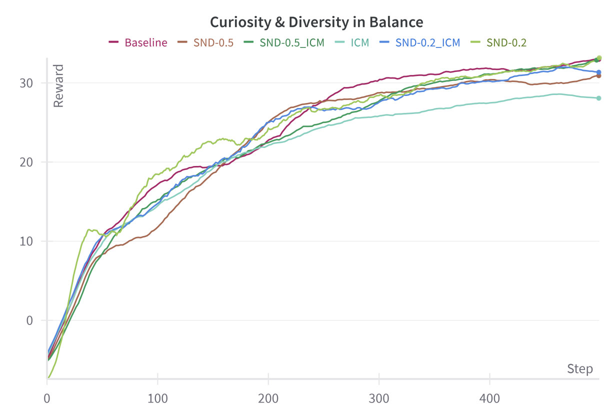

Figure 4. Performance Comparison of Curisos, Diverse, and Combined Strategy Agents in the Balance Environment. The graph illustrates cumulative rewards over training episodes.

5.1.3. Combined Effects

When combining curiosity and diversity mechanisms, agents with 0.2 SND slightly underperformed compared to those using pure diversity mechanisms, as shown in Figure 4. Conversely, agents with 0.5 SND and curiosity outperformed their counterparts, indicating a complex relationship be-

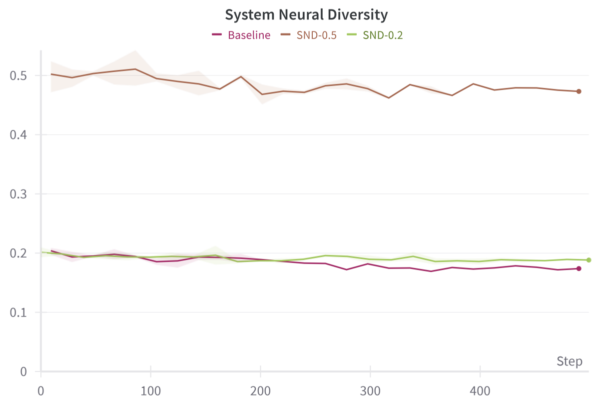

Figure 5. System Neural Diversity (SND) Levels Over Training Episodes in the Balance Environment. The graph shows the evolution of diversity among agent policies as measured by SND.

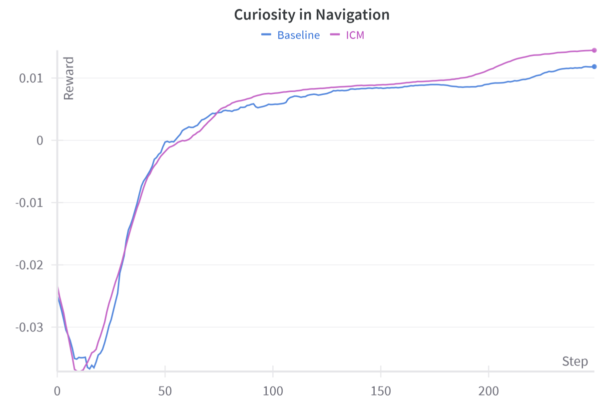

Figure 6. Comparison of Baseline vs. Curiosity-Driven Agents in the Multiagent Navigation Environment. The graph illustrates cumulative rewards over training episodes, highlighting the impact of intrinsic curiosity on exploration efficiency and task performance.

tween these novelty components.

5.2. Findings from the Multiagent Navigation Environment

5.2.1. Curiosity Implications

In the Multiagent Navigation environment, the implementation of ICM significantly enhanced exploration efficiency. Curiosity-driven agents exhibited increased state coverage, exploring a broader range of the environment and achieving higher cumulative rewards, as illustrated in Figure 6. This improved exploration led to faster convergence, evidenced by the reduced episode lengths shown in Figure 7. Notably, agents discovered more efficient paths to their targets and exhibited innovative behaviors, such as exploring less congested routes to avoid collisions.

5.2.2. Diversity Impacts

Diversity control mechanisms had pronounced effects on task performance and agent behaviors in the Navigation environment. Agents with desired SND levels of $SND_{\text{des}} = [0.2]$ and $SND_{\text{des}} = [0.5]$ achieved superior task performance compared to the baseline. This can be seen in Figure 7, where agents with lower diversity levels ($SND_{\text{des}} = [0.2]$) took the lead, converging faster to optimal policies. The faster learning is further evidenced by the cumulative rewards, indicating that constraining diversity to these levels can enhance learning efficiency.

5.2.3. Combined Influences

Combining curiosity and diversity mechanisms in the Navigation environment led to mixed results. Agents with $SND_{\text{des}} = [0.2]$ and ICM modules performed as well as the curiosity-driven agents, as shown in Figure 7. This indicates that curiosity-driven exploration can enhance the performance of agents even in scenarios where diversity is constrained. However, when combining curiosity and diversity with $SND_{\text{des}} = [0.5]$, agents performed worse than their counterparts with pure diversity. This suggests a more complex relationship between curiosity and diversity, where higher levels of diversity combined with curiosity-driven exploration may lead to excessive exploration and degradation in performance.

6. Discussion

Our experiments reveal that incorporating novelty components like curiosity and/or diversity can enhance MARL agent effectiveness, particularly in environments requiring exploration and role specialization. The Balance and Multiagent Navigation environments provided contrasting insights into the interplay between curiosity and diversity, highlighting the importance of task complexity and coordination requirements in shaping the benefits of novelty mechanisms.

6.1. Balance Environment Insights

In the Balance environment, the benefits of incorporating curiosity and diversity were limited.

- **Curiosity Limitations:** The small and well-defined state space meant that agents quickly learned optimal policies without needing extensive exploration. Excessive exploration driven by curiosity disrupted coordination without providing significant benefits.
- **Diversity Trade-offs:** While moderate diversity

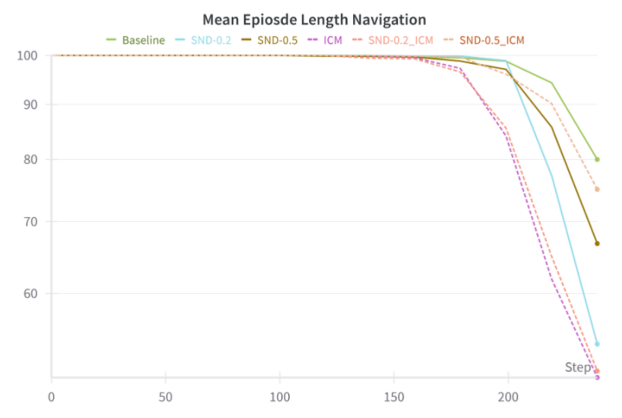

Figure 7. Comparison of Baseline vs. ICM vs. Curiosity-Driven Agents in the Multiagent Navigation Environment. The graph illustrates the mean episode length over training episodes, highlighting the impact of factors of novelty on exploration efficiency.

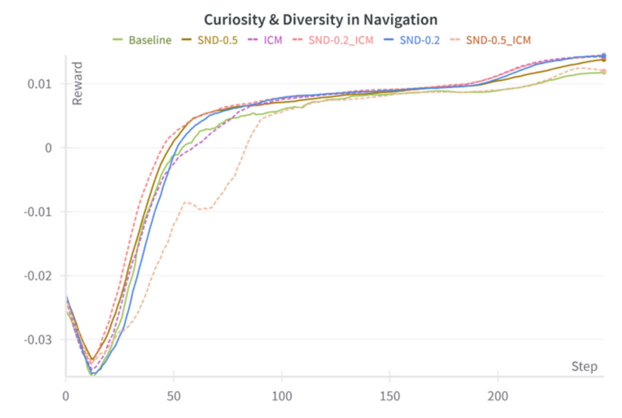

Figure 8. Performance Comparison of Baseline, Curiosity-Driven, Diverse, and Combined Strategy Agents in the Multiagent Navigation Environment. The graph illustrates cumulative rewards over training episodes.

slightly improved performance by preventing premature convergence, high diversity levels hindered the precise synchronization required, reducing overall performance.

- **Combined Mechanisms:** Integrating curiosity and diversity did not yield notable improvements. The task's cooperative nature favors homogeneous behaviors, suggesting that in environments demanding tight coordination, novelty mechanisms must be carefully calibrated.

6.2. Multiagent Navigation Environment Insights

In contrast, the Multiagent Navigation environment benefited significantly from the integration of curiosity and diversity.

- **Curiosity Benefits:** The larger and more complex state space allowed curiosity-driven agents to explore effectively, discovering efficient paths and improving task engagement.
- **Diversity Advantages:** Constraining diversity to 0.2 and 0.5 led to role specialization, enhancing agents' ability to navigate to their unique targets while avoiding collisions. Both levels of diversity showed improved performance over the baseline, with 0.2 achieving faster convergence and higher cumulative rewards.
- **Complementary Effects:** Curiosity-driven exploration complemented constrained diversity, resulting in agents performing on par with the best configuration, which were curiosity-driven agents. This indicates that maintaining a controlled level of diversity does not undermine the value of curiosity-driven exploration.
- **Innovation Indicators:** Faster convergence is a sign that exploration led to the discovery of less congested routes and the avoidance of collisions, indicating that integrating these novelty components can drive innovation and improve performance in complex tasks.

7. Limitations

- **Environment Constraints:** The environments used in the experiments were limited by the complexity of the tasks, potentially constraining the benefits of combining curiosity and diversity.
- **Parameter Sensitivity:** The performance of novelty mechanisms is sensitive to parameter settings (e.g., α, lr, SND_{des}). Finding optimal values may require extensive tuning. For example, for the α parameter, the value of 0.1 was used for all experiments, which may vary as the diversity levels change.
- **Computational Resources:** Increased computational demands due to additional modules like ICM and DiCo may limit scalability.

8. Future Work

These findings suggest that adaptive strategies that dynamically adjust the contributions of curiosity and diversity based on task demands could optimize innovation and learning stability. In tasks requiring tight coordination, mechanisms should favor homogeneity, possibly reducing the influence of curiosity and diversity. This task-specific calibration ensures that agents can synchronize their actions effectively without being disrupted by excessive exploration or behavioral divergence.

Implementing adaptive mechanisms that adjust the scaling factors α (curiosity) and SND (diversity) during training could allow agents to balance exploration and exploitation more effectively. By dynamically adjusting their novelty-seeking behaviors based on task requirements, agents can optimize their learning processes and improve overall performance.

Future studies could explore adaptive integration strategies and investigate the impact of these novelty components in other environments that facilitate more complex interactions and emergent behaviors. This research could provide deeper insights into how curiosity and diversity contribute to innovation in artificial systems, paving the way for more advanced and adaptive multi-agent learning frameworks.

Further research could explore additional novelty components, to gain a more comprehensive understanding of their interactions and influence on agent behaviors, leading to more sophisticated insights. This approach provides a foundation for innovation modeling in artificial systems and offers new perspectives on how novelty mechanisms can drive emergent behaviors and learning dynamics in multi-agent systems.

9. Conclusion

This thesis introduces a novel approach to studying innovation in multi-agent systems by examining how factors of novelty: curiosity and diversity contribute to agent behavior. By combining Intrinsic Curiosity Modules and Diversity Control within a MARL framework, we provide new tools for investigating how different aspects of novelty influence learning and behavior in artificial systems.

The key contribution of our work lies not in demonstrating dramatic performance improvements, but in establishing that curiosity and diversity can be effectively combined as components of novelty without mutual interference. This finding opens new avenues for studying innovation in artificial systems and suggests that the relationship between these novelty components is more complex than previously understood.

Our experiments across two distinct environments revealed that the effectiveness of novelty components in driving innovation depends strongly on environmental characteristics and task constraints. In the Balance environment, the benefits were limited by task constraints, while in the Navigation environment, the combination of curiosity and diversity supported the emergence of diverse behavioral strategies.

These findings lay the groundwork for future research into how multiple components of novelty can be integrated and dynamically balanced to promote innovation in artificial systems. The framework we've developed provides new tools for measuring and analyzing innovation-promoting factors, contributing to our understanding of how complex behaviors emerge from simple mechanisms in multi-agent systems.

References

Aydeniz, A. A., Loftin, R. T., and Tumer, K. Novelty seeking multiagent evolutionary reinforcement learning. *Proceedings of the Genetic and Evolutionary Computation Conference*, 2023. URL `https://api.semanticscholar.org/CorpusID:259833682`.

Baker, B., Kanitscheider, I., Markov, T., Wu, Y., Powell, G., McGrew, B., and Mordatch, I. Emergent tool use from multi-agent autocurricula, 2020. URL `https://arxiv.org/abs/1909.07528`.

Berlyne, D. E. Conflict, arousal, and curiosity. 2014. URL `https://api.semanticscholar.org/CorpusID:143207836`.

Bettini, M., Kortvelesy, R., and Prorok, A. Controlling behavioral diversity in multi-agent reinforcement learning, 2024a. URL `https://arxiv.org/abs/2405.15054`.

Bettini, M., Shankar, A., and Prorok, A. System neural diversity: Measuring behavioral heterogeneity in multi-agent learning, 2024b. URL `https://arxiv.org/abs/2305.02128`.

Burda, Y., Edwards, H., Storkey, A., and Klimov, O. Exploration by random network distillation, 2018. URL `https://arxiv.org/abs/1810.12894`.

Ecoffet, A., Huizinga, J., Lehman, J., Stanley, K. O., and Clune, J. Go-explore: a new approach for hard-exploration problems, 2021. URL `https://arxiv.org/abs/1901.10995`.

Eecke, P. V., Beuls, K., Ekila, J. B., and Rădulescu, R. Language games meet multi-agent reinforcement learning: A case study for the naming game. *Journal of Language Evolution*, 2023. URL `https://api.semanticscholar.org/CorpusID:258223656`.

Gerstgrasser, M., Danino, T., and Keren, S. Selectively sharing experiences improves multi-agent reinforcement learning, 2024. URL `https://arxiv.org/abs/2311.00865`.

. **AUTHORERR: Missing \icmlcorrespondingauthor.**

Lazaridou, A., Peysakhovich, A., and Baroni, M. Multi-agent cooperation and the emergence of (natural) language. *ArXiv*, abs/1612.07182, 2016. URL https://api.semanticscholar.org/CorpusID:14911774.

Leibo, J. Z., Hughes, E., Lanctot, M., and Graepel, T. Autocurricula and the emergence of innovation from social interaction: A manifesto for multi-agent intelligence research, 2019. URL https://arxiv.org/abs/1903.00742.

Lowe, R., Wu, Y., Tamar, A., Harb, J., Abbeel, P., and Mordatch, I. Multi-agent actor-critic for mixed cooperative-competitive environments, 2020. URL https://arxiv.org/abs/1706.02275.

Parker-Holder, J., Pacchiano, A., Choromanski, K., and Roberts, S. Effective diversity in population based reinforcement learning, 2020. URL https://arxiv.org/abs/2002.00632.

Pathak, D., Agrawal, P., Efros, A. A., and Darrell, T. Curiosity-driven exploration by self-supervised prediction, 2017. URL https://arxiv.org/abs/1705.05363.

Ting, Z. A review of innovation diffusion models based on cellular automata. *Systems Engineering*, 2006. URL https://api.semanticscholar.org/CorpusID:63882113.

y Arcas, B. A., Alakuijala, J., Evans, J., Laurie, B., Mordvintsev, A., Niklasson, E., Randazzo, E., and Versari, L. Computational life: How well-formed, self-replicating programs emerge from simple interaction, 2024. URL https://arxiv.org/abs/2406.19108.

Zhang, H. and Vorobeychik, Y. Empirically grounded agent-based models of innovation diffusion: A critical review. *Artificial Intelligence Review*, 52, 06 2019. doi: 10.1007/s10462-017-9577-z.