Dynamic Hybrid Governance of Sociotechnical Systems

Balancing Top-Down Interventions and Bottom-Up Responses

Babak Heydari, Northeastern University

AI4SE 2024

MAGICS LAB

Home Research Team Lab Director Collaboration

MAGICS Lab

Multi-AGent Intelligent Complex Systems

Qingtao Cao (PhD student)

Research Interests: Network Science, Multi-Agent Systems, Sharing Economy, Computational Methods for Policy and Regulations, Spatiotemoral Data Analysis

Negin Maddah (PhD Student)

Research Interests: Complex Networks, Computational Social Science, Intelligent SocioTechnical Systems, Process Mining

Soumyakant Padhee (PhD Student)

Research Interests: Innovation in Engineering Design Teams, Causal Identification in SocioTechnical Systems, Experimental methods.

Qiliang Chen (PhD Student)

Research Interests: Artificial Intelligence, Meta Reinforcement Learning, Modularity, Multi-Agent Systems, Interpretable AI

Babak Heydari (Associate Professor & Lab Director)

(Link to more details)

Nunzio Lore (PhD Student)

Research Interests: Network Science, Game Theory, Agent-Based Modeling, Fairness in AI Platforms.

Gabriele Ansaldo (MSc. Student)

Research Interest: Human - Robot Teaming, Multi-Agent Systems, Trustable AI

Alireza (Sepehr) Ilami (PhD Student)

Research Interest: Computational Models for AI Policy and Governance

Hybrid Mechanisms of Sociotechnical Systems

- **Socio-Technical Systems (STS):** The Intersection of Organic (BU, Evolving) and Engineered (TD, Controlled) Systems.
- **Synergy or Discord:** Interplay Between Top-Down and Bottom-Up Forces in STS.
- **Embracing Inherent Forces:** Different from "human error"; Suppression is Not the Solution.
- **Governance Perspective:** Fostering Synergistic Mechanisms through Effective Bottom-Up Dynamics. Can't Force a desired Bottom-up Dynamic.

Synergistic Adaptive Governance (SAG)

- ❑ **Working definition:** A set of precautions throughout the lifecycle of the system that **enable effective bottom-up mechanisms** through **top-down design.**

- ❑ **Steering Emergence:** e.g. Resilience/cooperation/ trust as an emergent property.

- ❑ **Broad notion of TD design**: includes system's architecture, incentives, and regulations.

- ❑ Can be extended to more levels of hierarchy: Enabling effective **lower-level responses** through **higher-level design**.

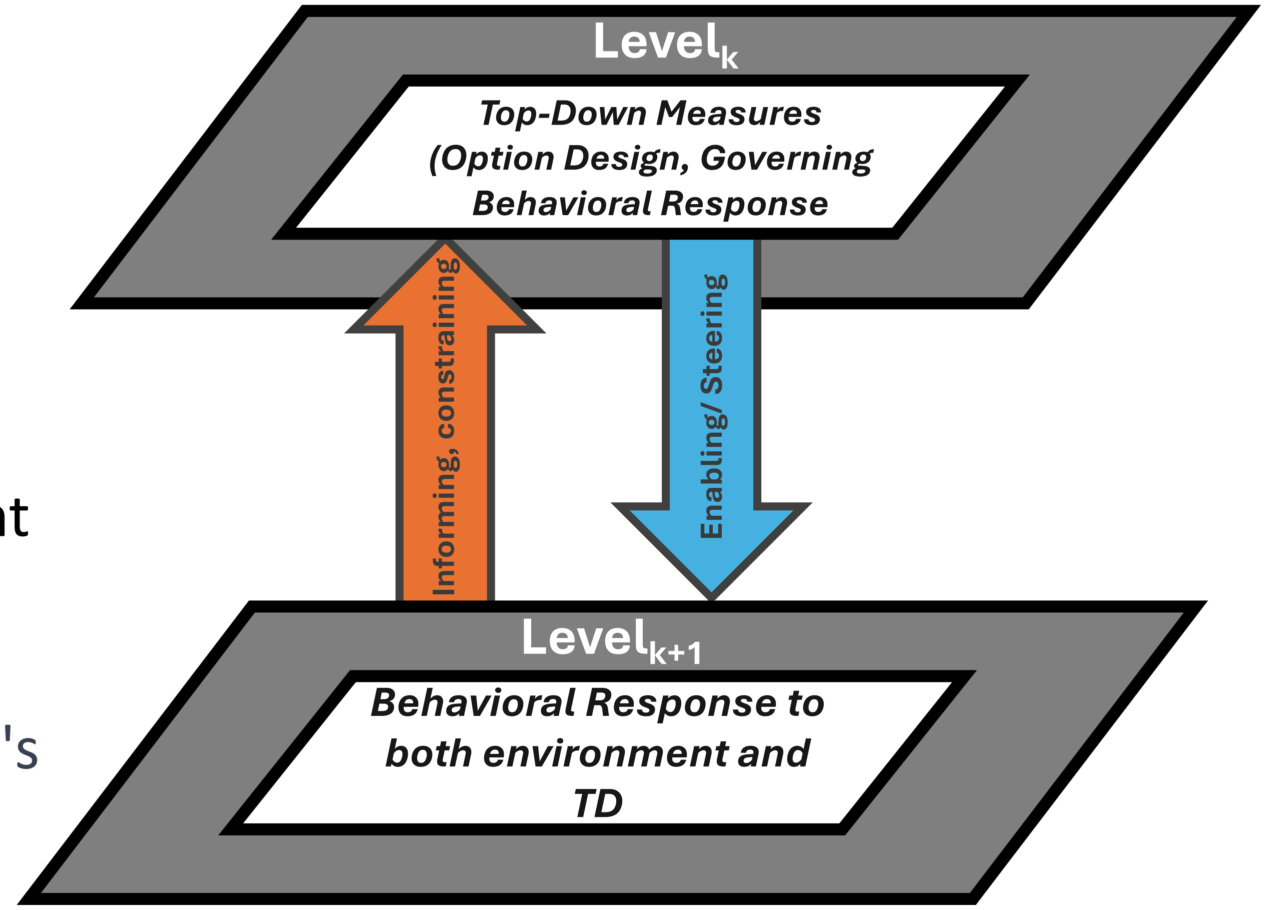

Two Pillars of Synergistic Adaptive Governance

- ❑ Enabling Effective recovery through **Option Design**

- ❑ Guiding the **Behavioral Response**

- ❑ There is an **overlap** because of how system structure/architecture **co-evolves** with agents' behavior.

- ❑ **AI can make this overlap more significant by:**
  - ❑ Making Option Design Adaptive
  - ❑ Expanding the types of options (architecture, resources, information)
  - ❑ Using adaptive options to steer the collective behavior

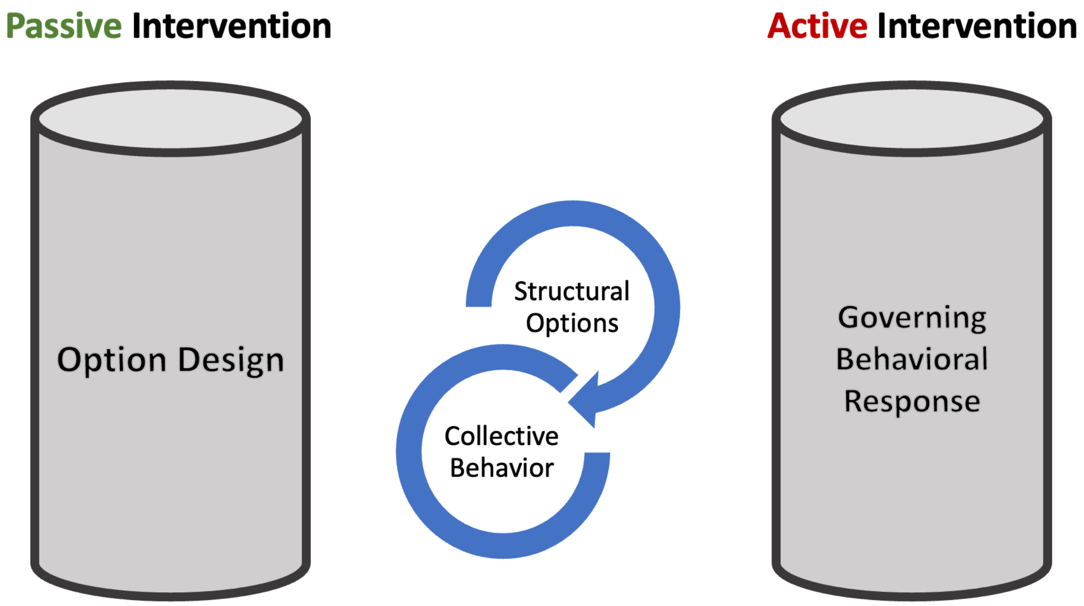

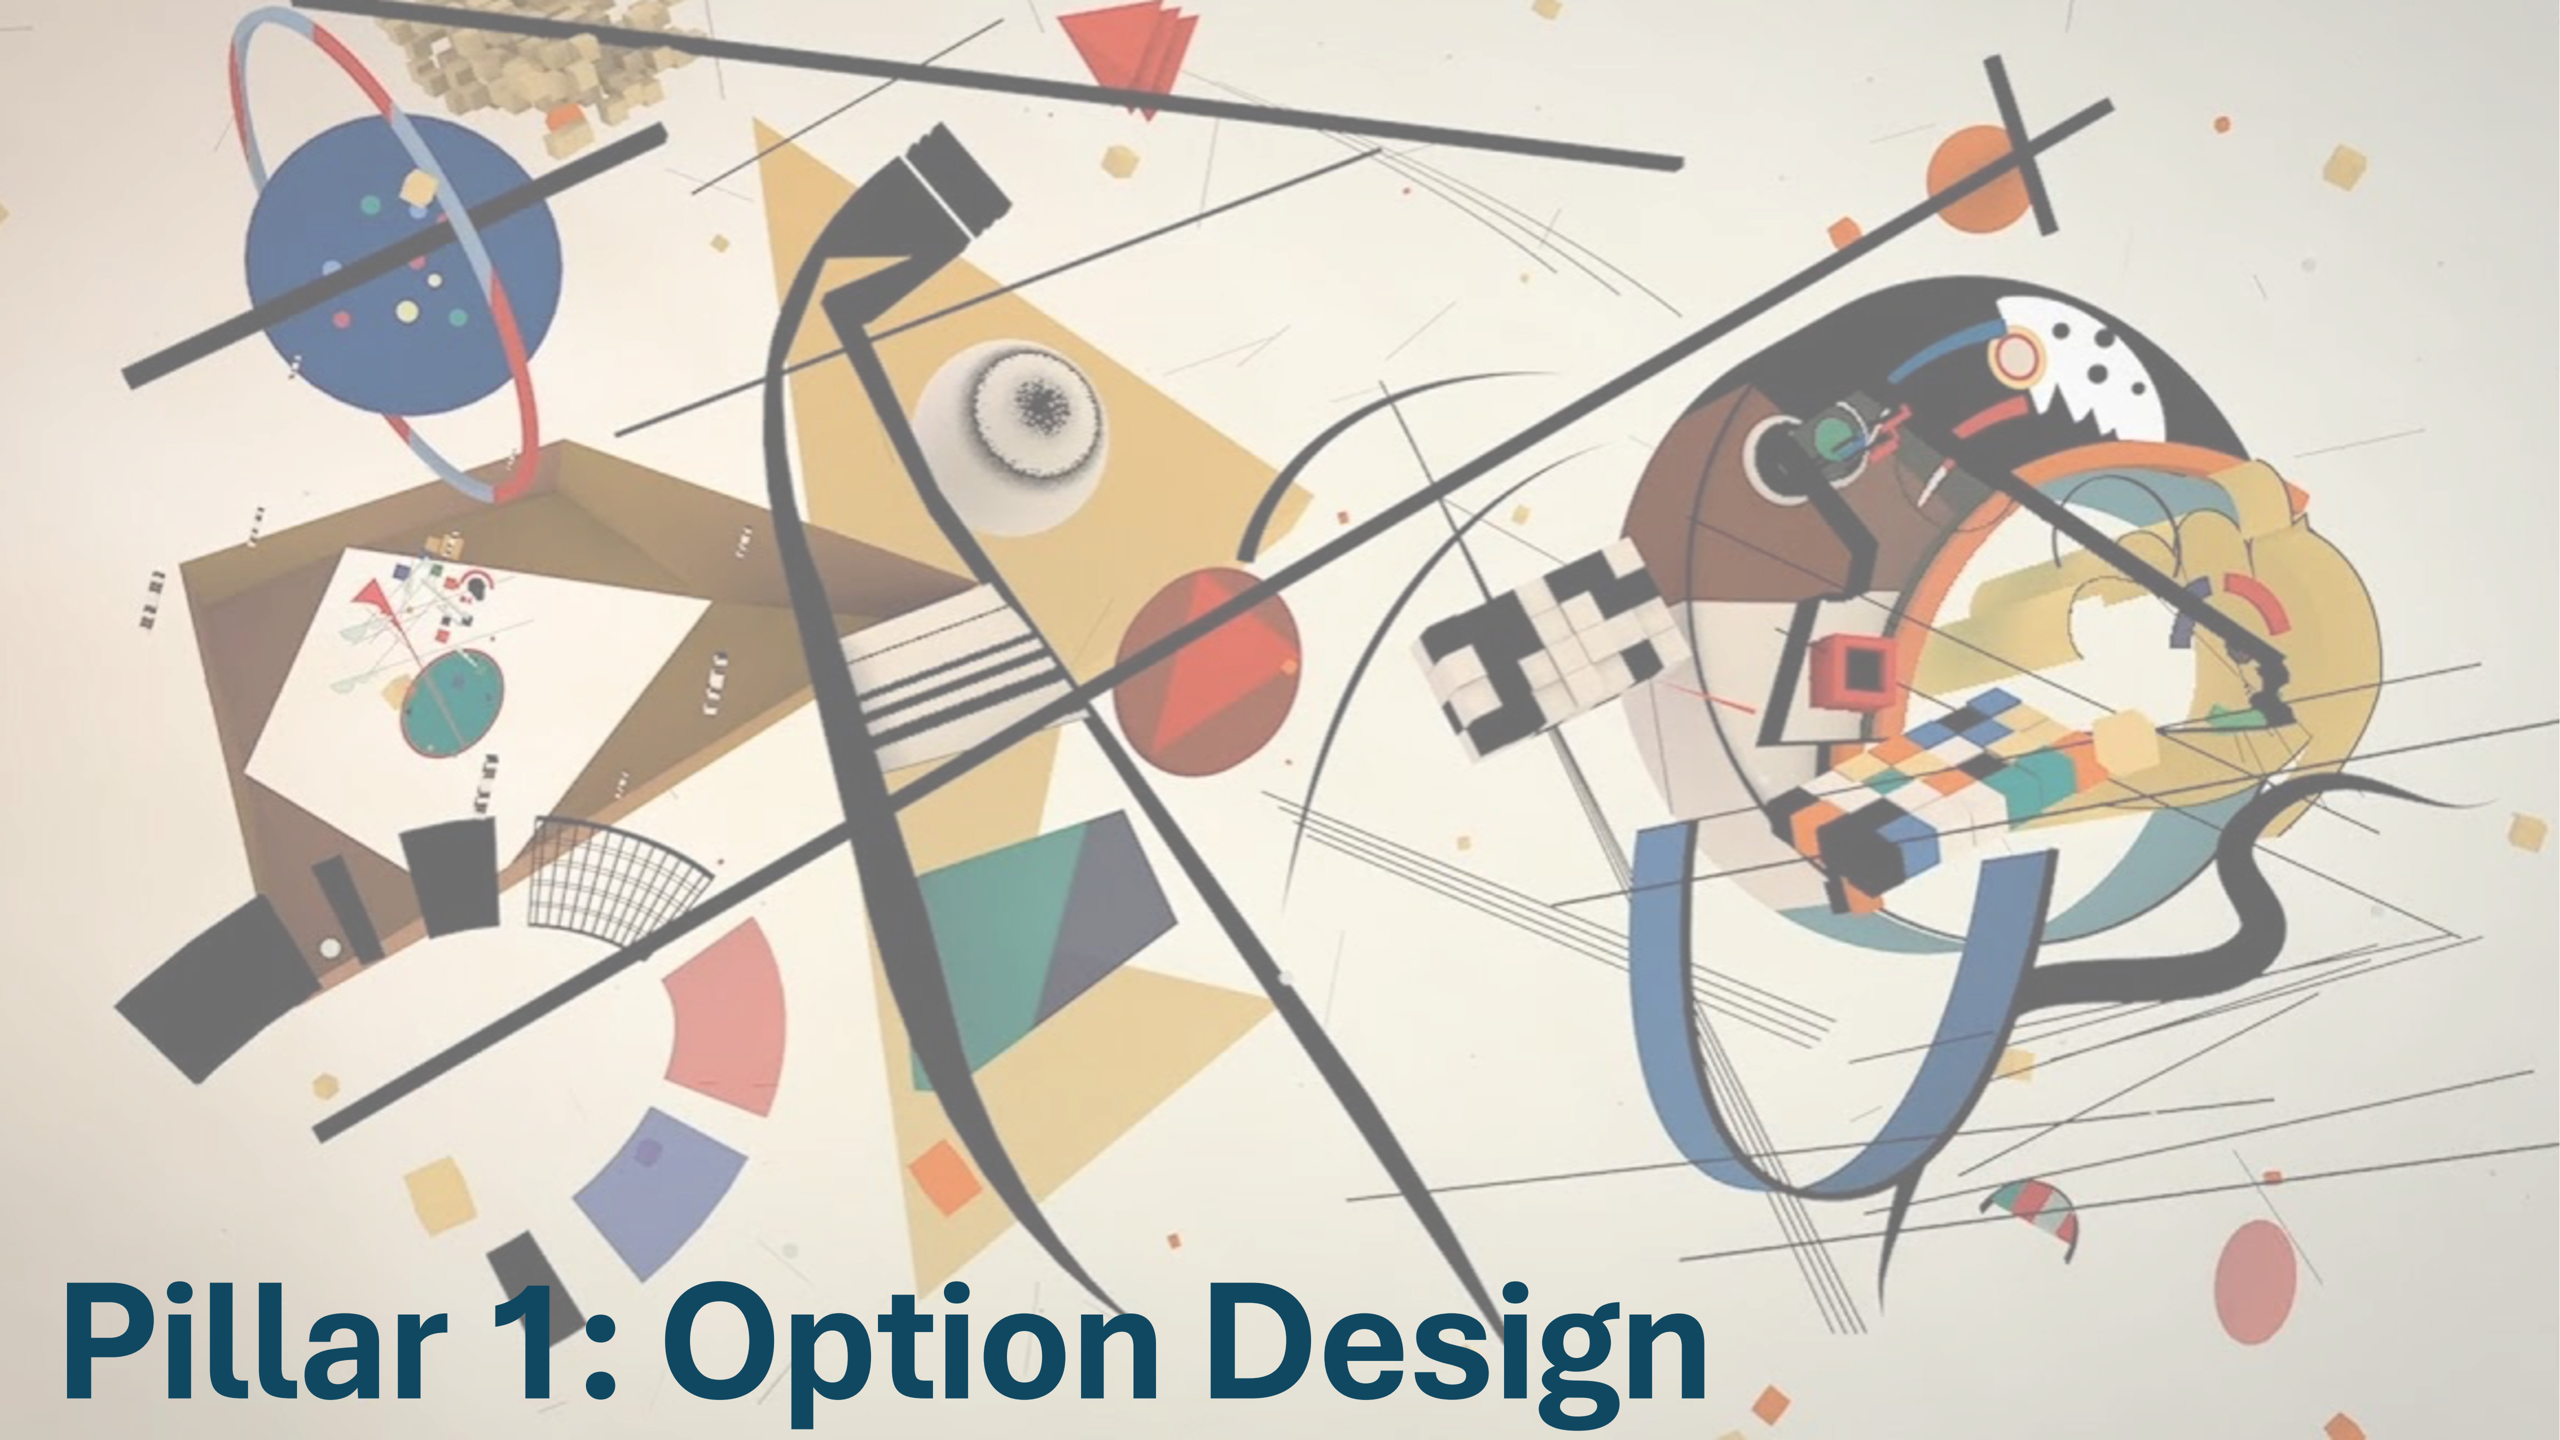

Pillar 1: Option Design

Governance through Option Design

- Options are rights, but not obligations, of executing a decision in the future (if it is beneficial).
- Embedding options at higher level enables effective behavioral response at lower levels.
- Many different forms; I categorize them into four clusters, in the order of complexity (gaining lower fixed and transaction-cost).

Redundancy

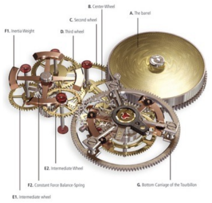

Modularity

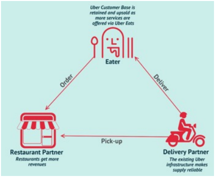

Multi-sided Platforms

Adaptive Options (AI-based)

Heydari, Babak, Mohsen Mosleh, and Kia Dalili. "From modular to distributed open architectures: A unified decision framework." *Systems Engineering* 19.3 (2016): 252-266.

Hierarchical SoS governance framework
Iterative Agent-Based Reinforcement Learning (IAB-RL)

- Hierarchy of autonomy (sub-systems, systems, SoS)
- Iterative Agent-Based Reinforcement Learning (IAB-RL)
- The higher-level RL learns to adjust the set of *options* available to each agent in the layer below.
- Can move the SoS away form inefficient Nash Equilibrium to more efficient outcomes, without compromising the autonomy of individual agents.
- Seeking interpretability in learned policies.

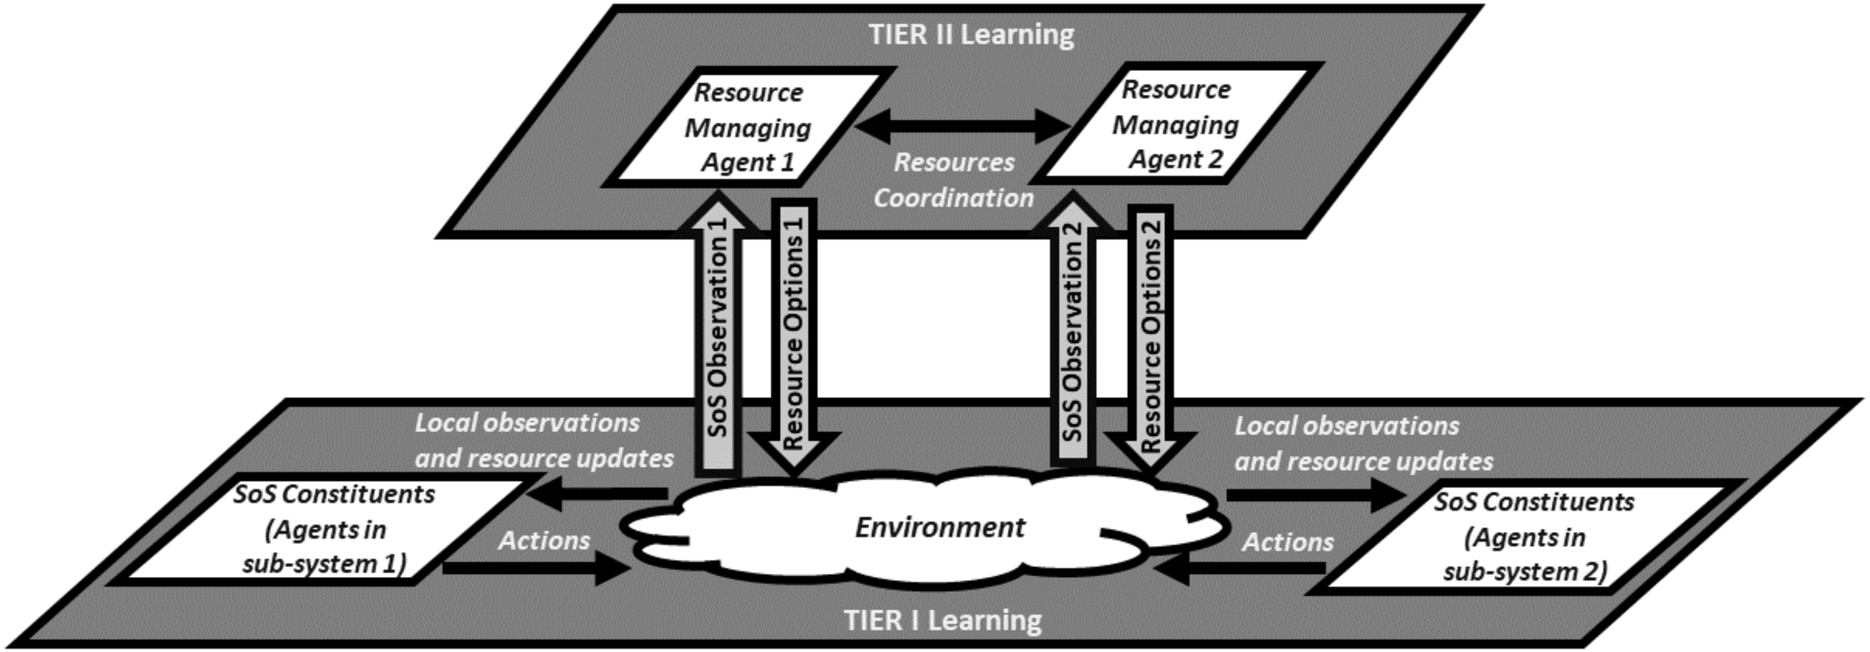

Chen, Qiliang, and Babak Heydari. "The SoS conductor: Orchestrating resources with iterative agent-based reinforcement learning." *Systems Engineering* (2024).

Two-Tier adaptive governance framework

- First implementation: A two-tier SoS
- A **multi-robot navigation problem** using modified OpenAI environment (agents needs to coordinate to spread on landmarks as soon as possible with less collision.
- Resource managing agents can assign two types of resources to agents: additional vision range and enabling communication between agents.
- The policies of agents and manager are represented using Deep neural networks.

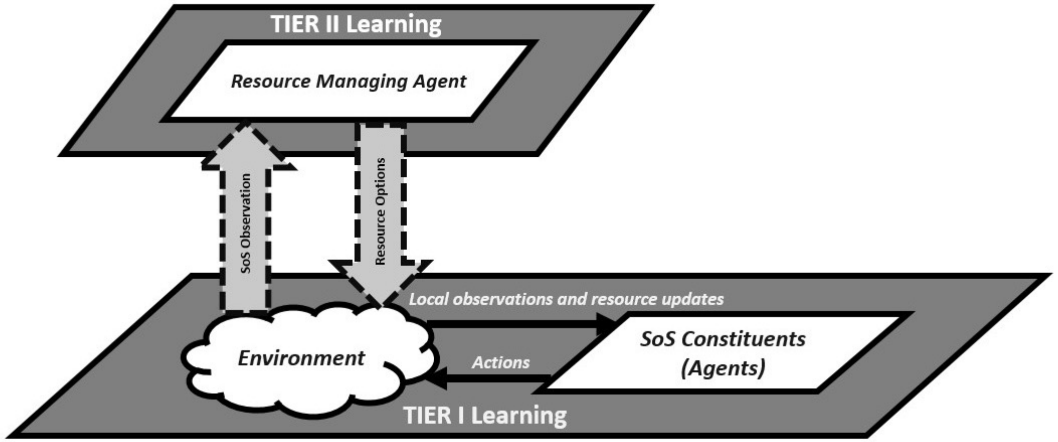

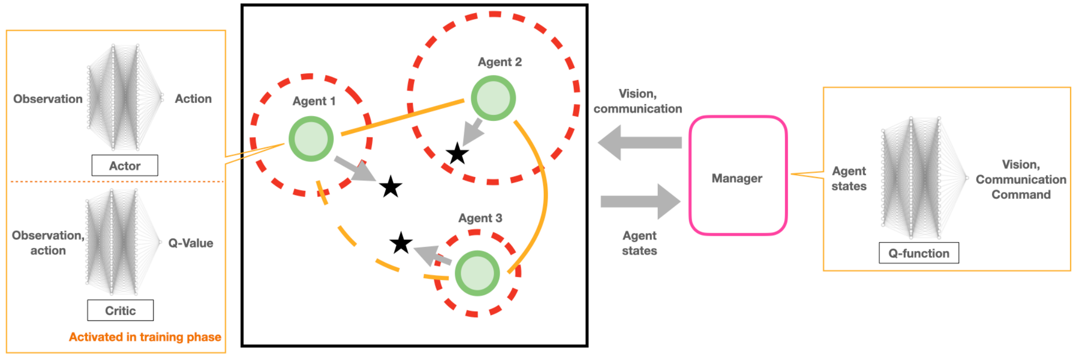

Chen, Qiliang, and Babak Heydari. "Dynamic Resource Allocation in Systems-of-Systems Using a Heuristic-Based Interpretable Deep Reinforcement Learning." *Journal of Mechanical Design* 144.9 (2022): 091711.

Two-Tier Results

- **Efficient Governance:** Performance comparison between the RL manager and a manager using a baseline policy shows that the **RL manager achieves Pareto optimality** across all methods in terms of performance and resource usage.
- **Interpretability:** Distribution of resource usage over time as dictated by the RL manager's learned policy. An interesting pattern emerges where it initially allocates expensive resources but shifts to cheaper resources once the agents start coordinating independently.

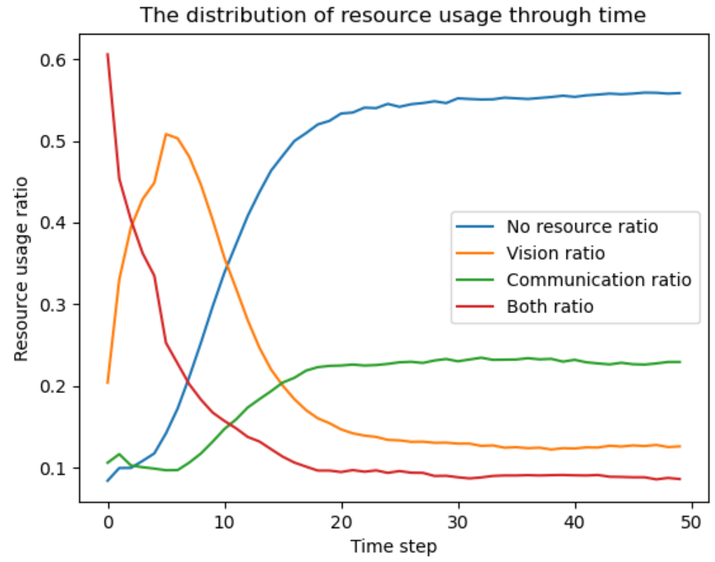

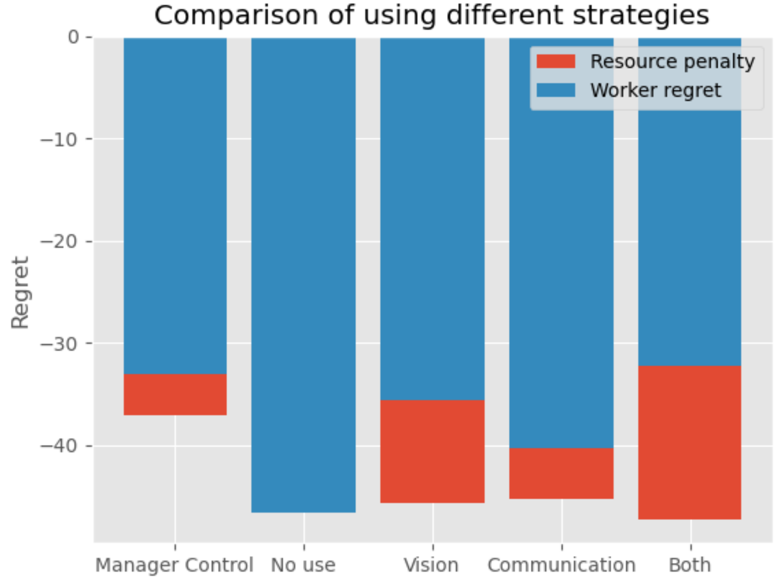

Chen, Qiliang, and Babak Heydari. "Dynamic Resource Allocation in Systems-of-Systems Using a Heuristic-Based Interpretable Deep Reinforcement Learning." *Journal of Mechanical Design* 144.9 (2022): 091711.

Experiment results – behavioral analysis

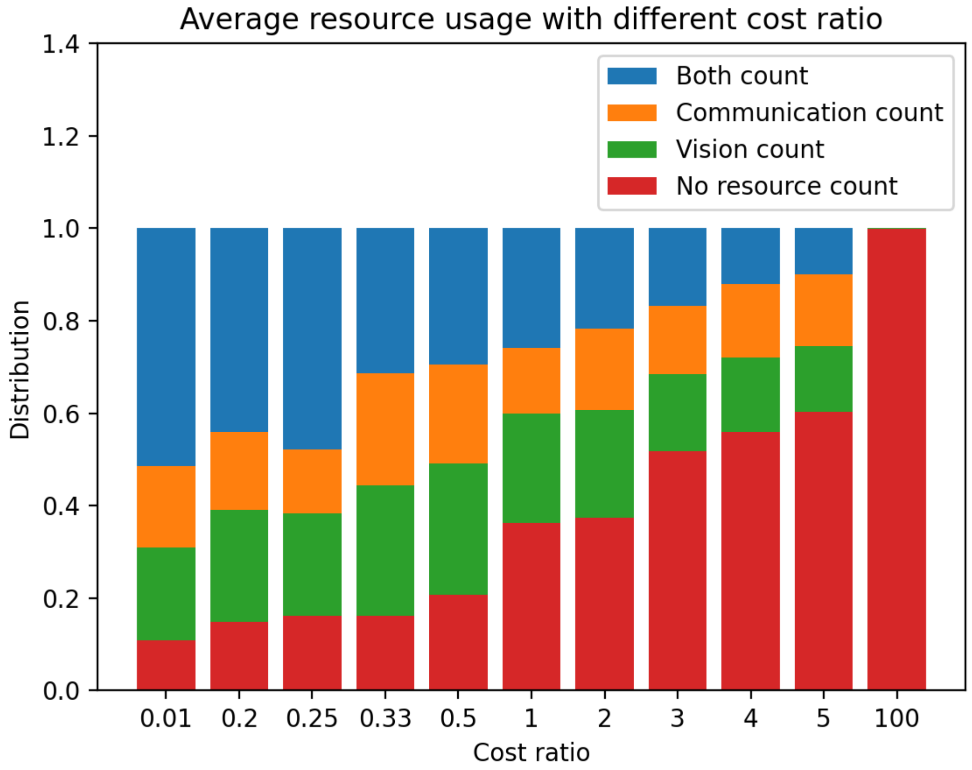

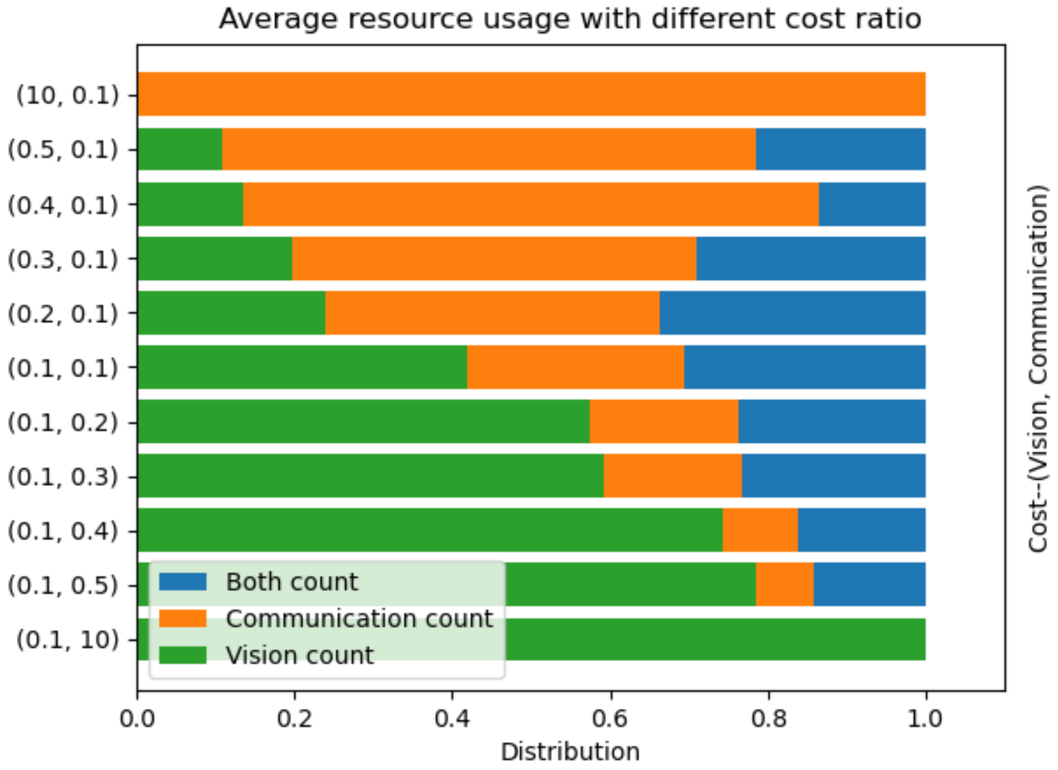

- The left figure displays the ratio of different resource usages by the RL manager, with varying importance weights between performance and resources. A clear pattern emerges: as resources become more expensive, the RL manager learns to use fewer resources.
- The right figure details the ratio of different resource usages by the RL manager, influenced by varying costs between two types of resources. It shows a distinct pattern: as the cost of one type of resource increases, the RL manager learns to use less of it and shifts towards utilizing the other type.

Generalized two-tier framework

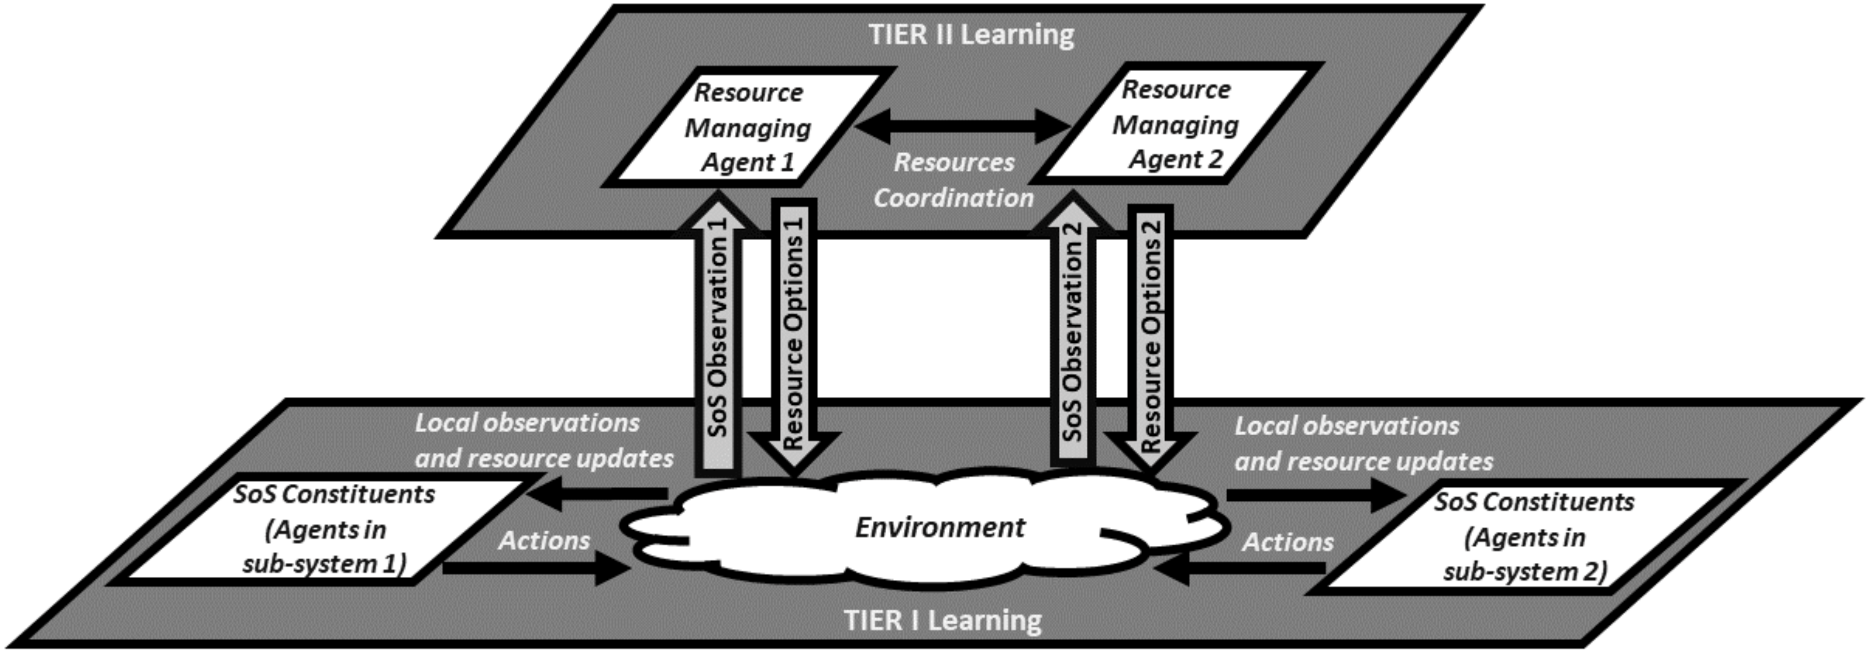

- A generalized version of the two-tier framework involves a system comprising multiple subsystems, each responsible for different tasks. In this model, not only do agents need to interact with each other, but resource managers also must coordinate with one another due to the generally limited availability of resources.
- The objectives of the resource managers extend beyond optimizing the performance of their own subsystems to also enhancing the performance of other subsystems. The weighting of these objectives depends on the level of coordination between the subsystems.

Experiment results

- Top Figure: The system performance across varying levels of coordination and resource budgets. With higher resource budgets, system performance improves.
- In all scenarios, the RL manager's policy consistently outperforms the baseline method. Additionally, the figure shows that with the same resource budget, a higher level of coordination decreases system performance.
- Bottom figure: Quantify the effect of coordination level. Small sacrifice in SoS goal pay-off can result in large gain in resource usage.

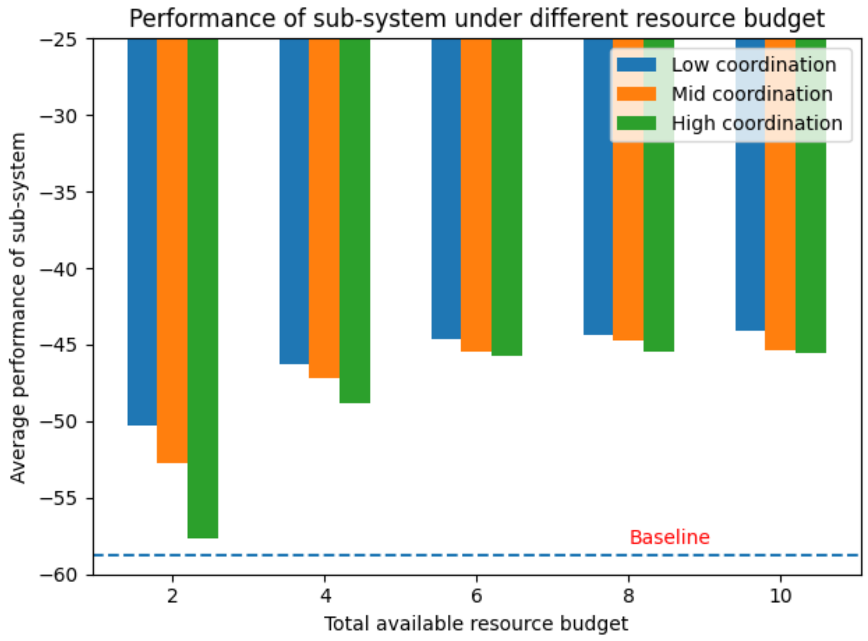

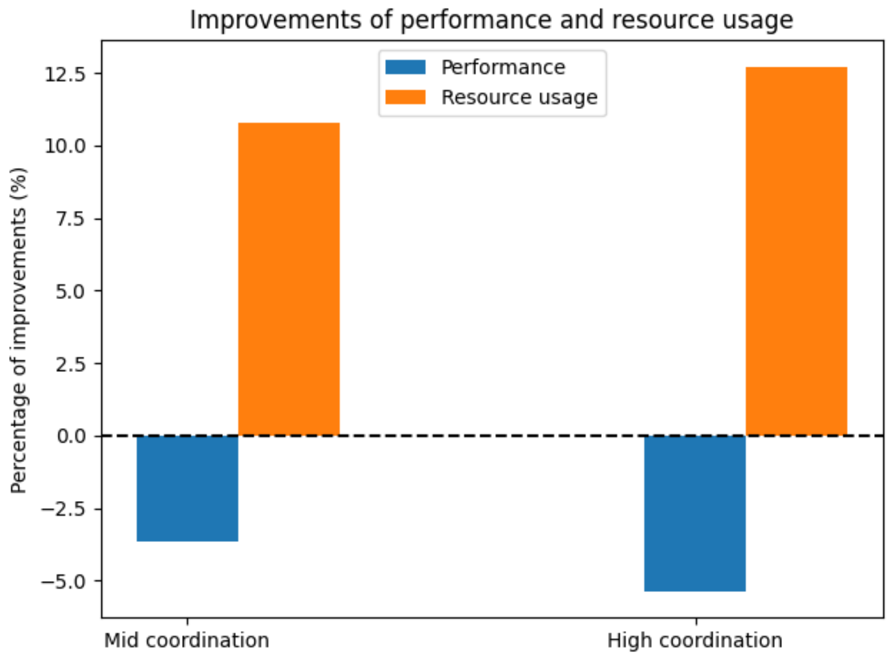

Adaptive network intervention for multi-agent systems

Variational Auto Encoder + Reinforcement Learning (VAE-RL) framework

- ❑ Dynamic Communication Access
- ❑ Heterogeneous Options → Dynamic Network Structure
- ❑ VAE to embed large action space to a few latent variables; then use continuous RL policies to control these latent spaces.

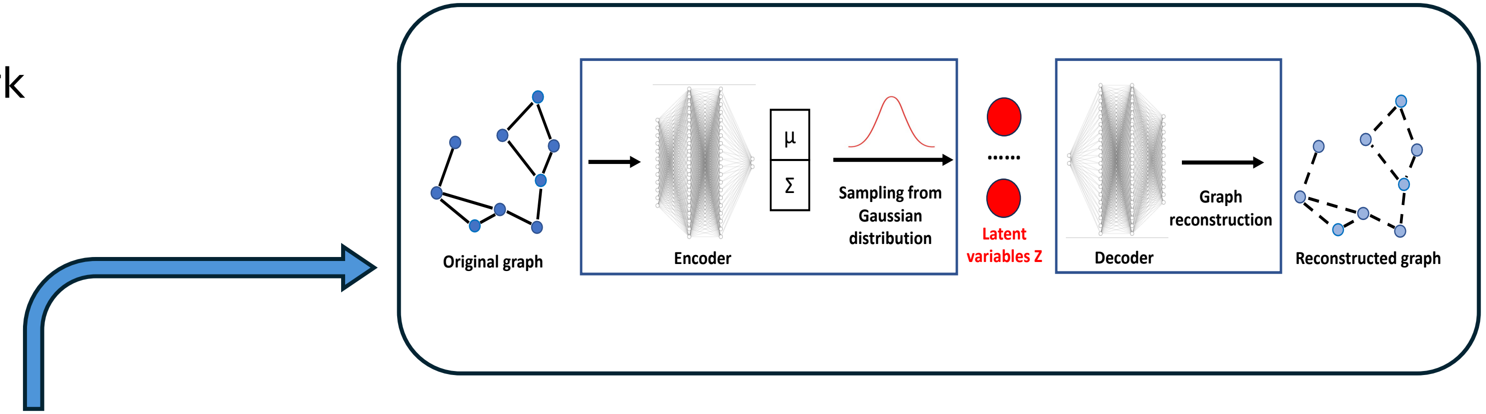

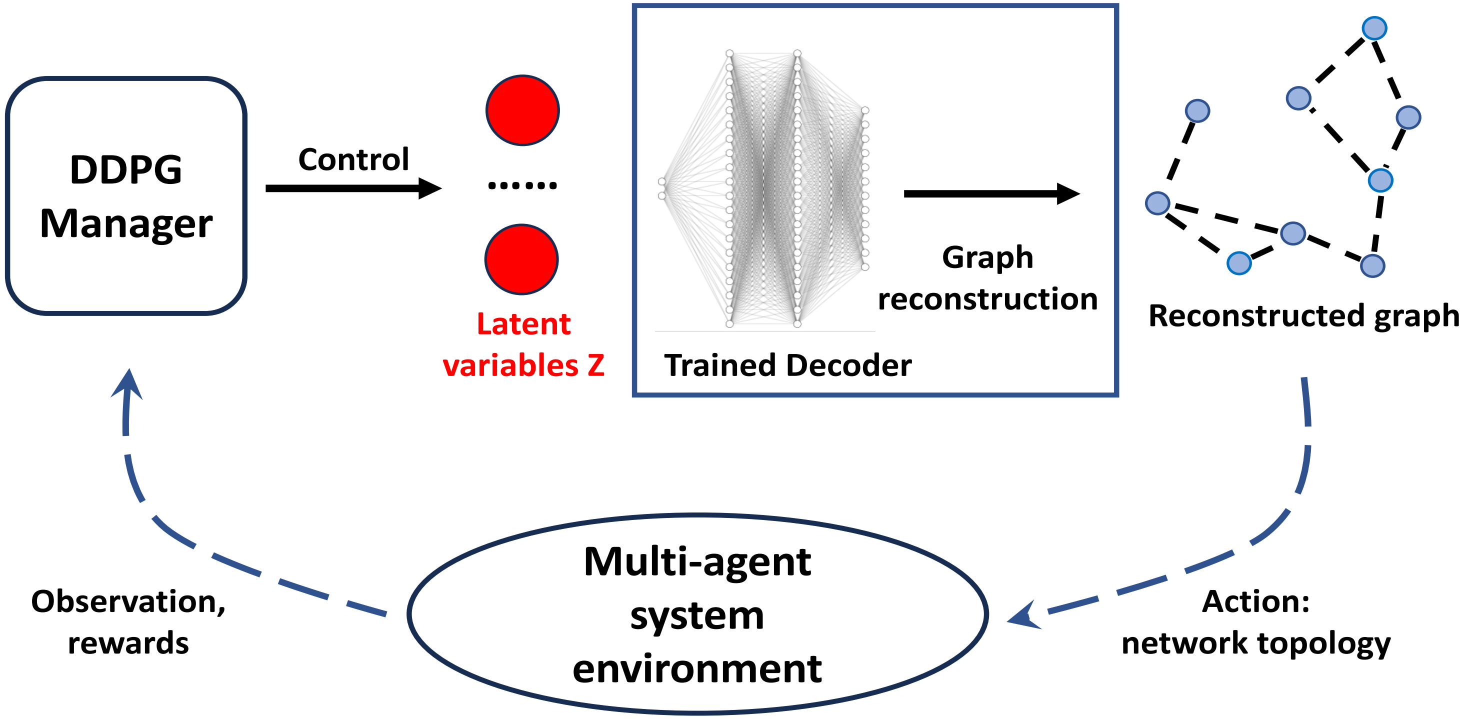

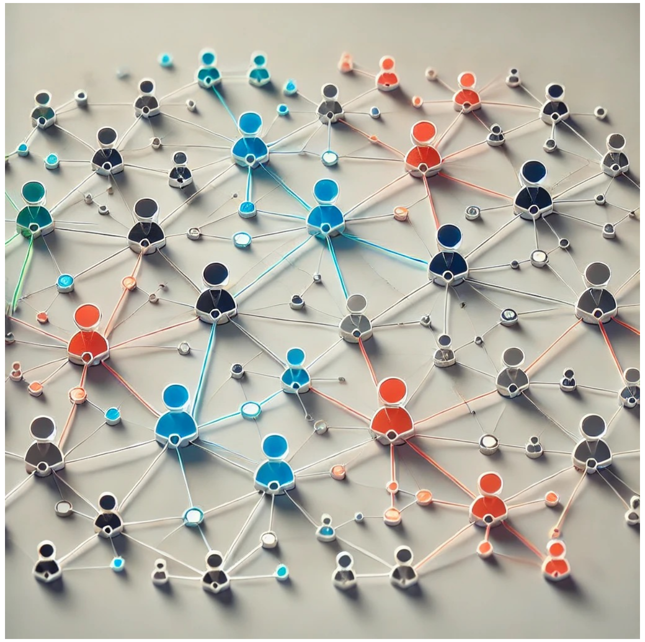

Experiment results

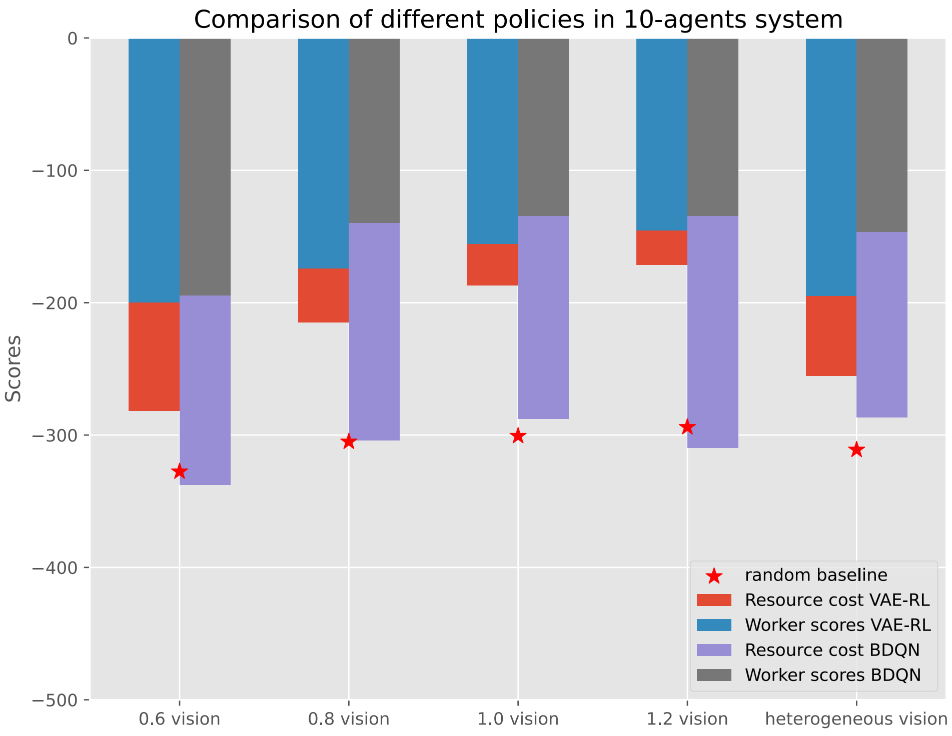

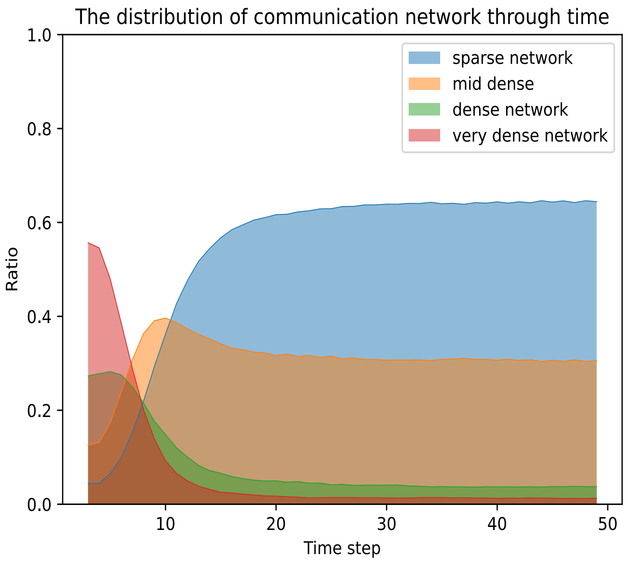

0.6 vision

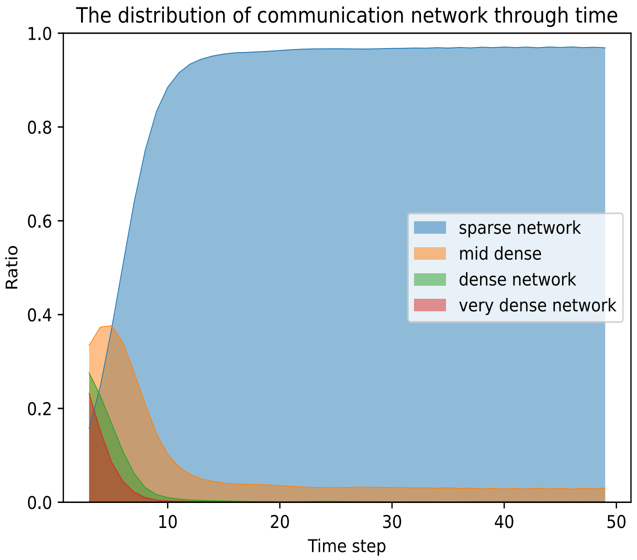

0.8 vision

Steering the behavioral response

Direct Modeling of Competition and social dilemmas among agents

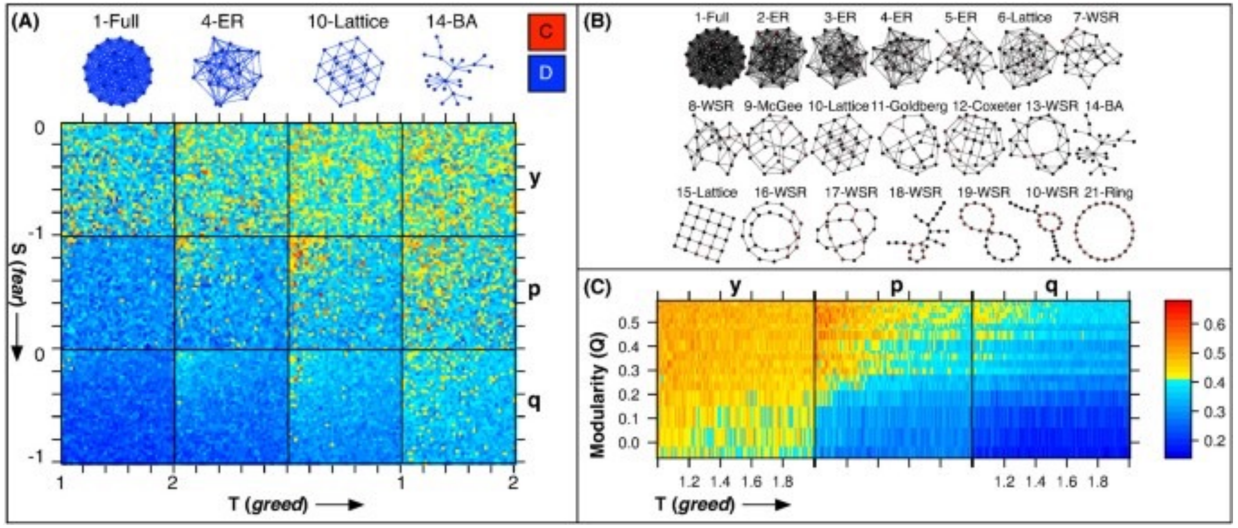

- ❑ Network Structure Plays a strong role in governing social dilemma and balancing cooperation/competition.
- ❑ What if we dynamically change the network structures?

Gianetto, David A., and Babak Heydari. "Network modularity is essential for evolution of cooperation under uncertainty." *Scientific reports* 5.1 (2015): 9340.

- ❑ Graph Neural Networks are good candidates, but…
- ❑ Introducing a new model (HGRL) based on the idea of link separability.
- ❑ Evolutionary Prisoners' Dilemma on Networks (Social Learning + Strategic Behavior)
- ❑ Learns to identify source and destination separately

Heydari, Babak, Mohsen Mosleh, and Kia Dalili. "Efficient network structures with separable heterogeneous connection costs." *Economics Letters* 134 (2015): 82-85.

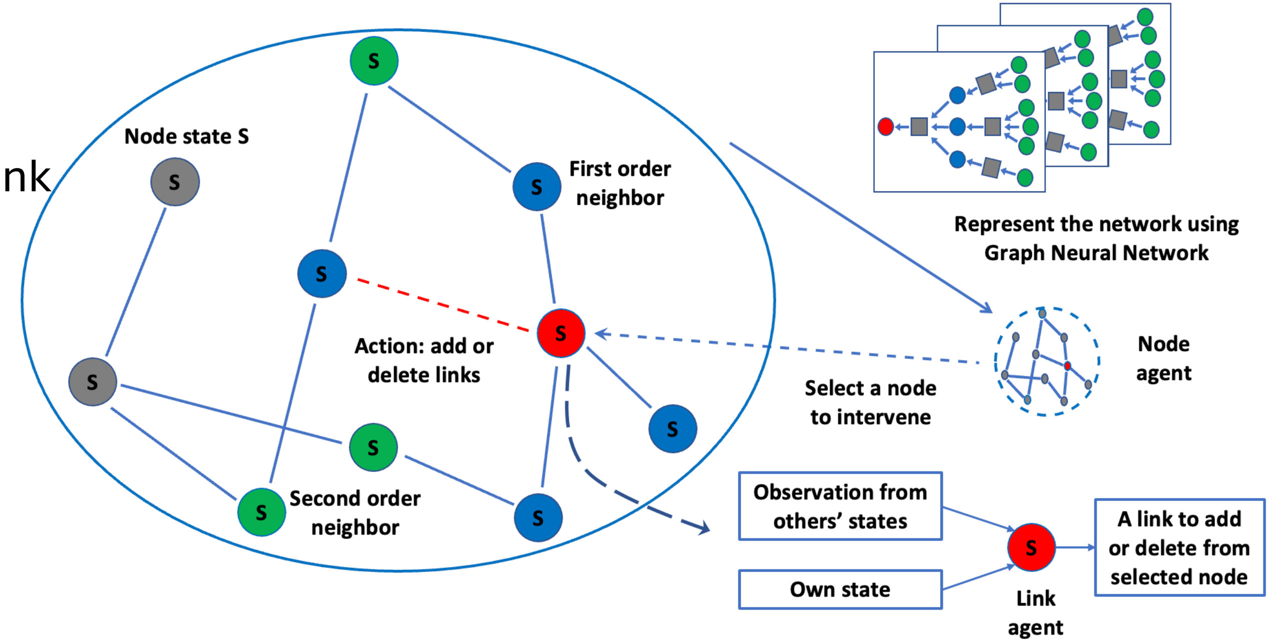

Experimental Results

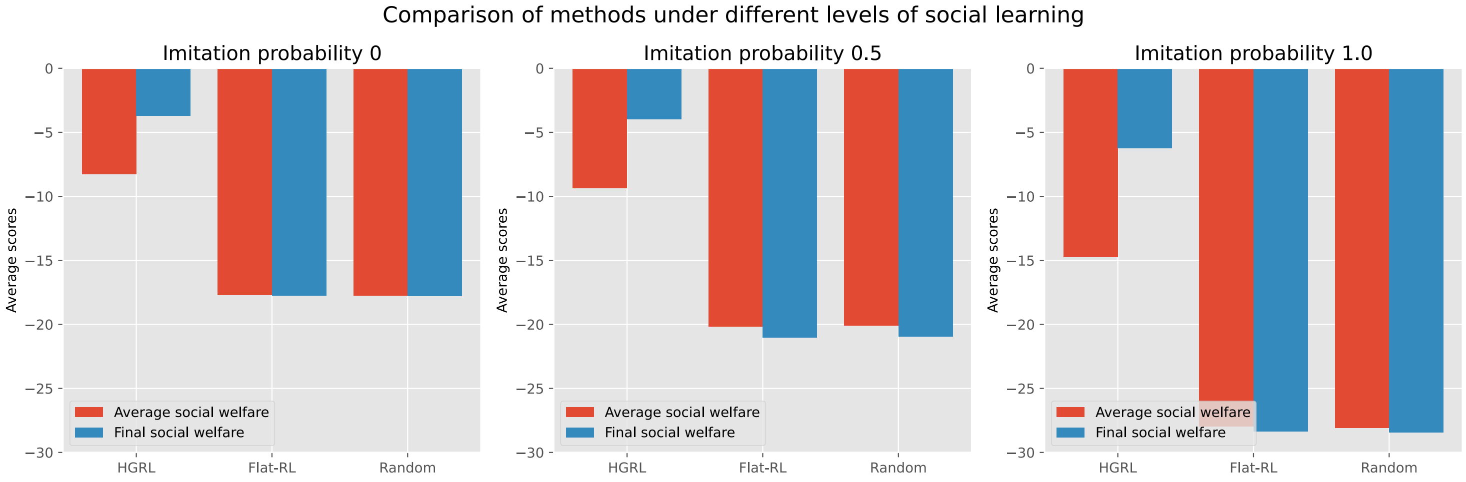

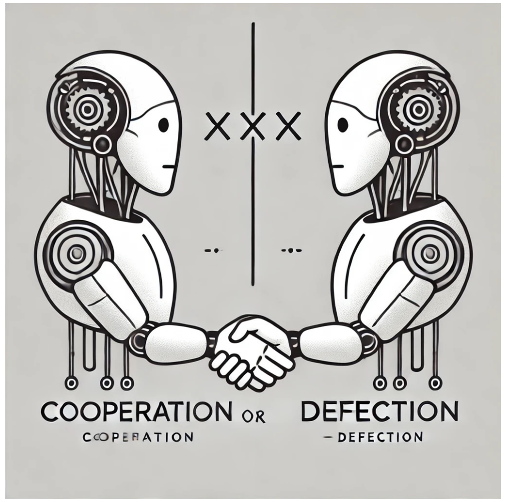

0.5 imitation probs

1.0 imitation probs

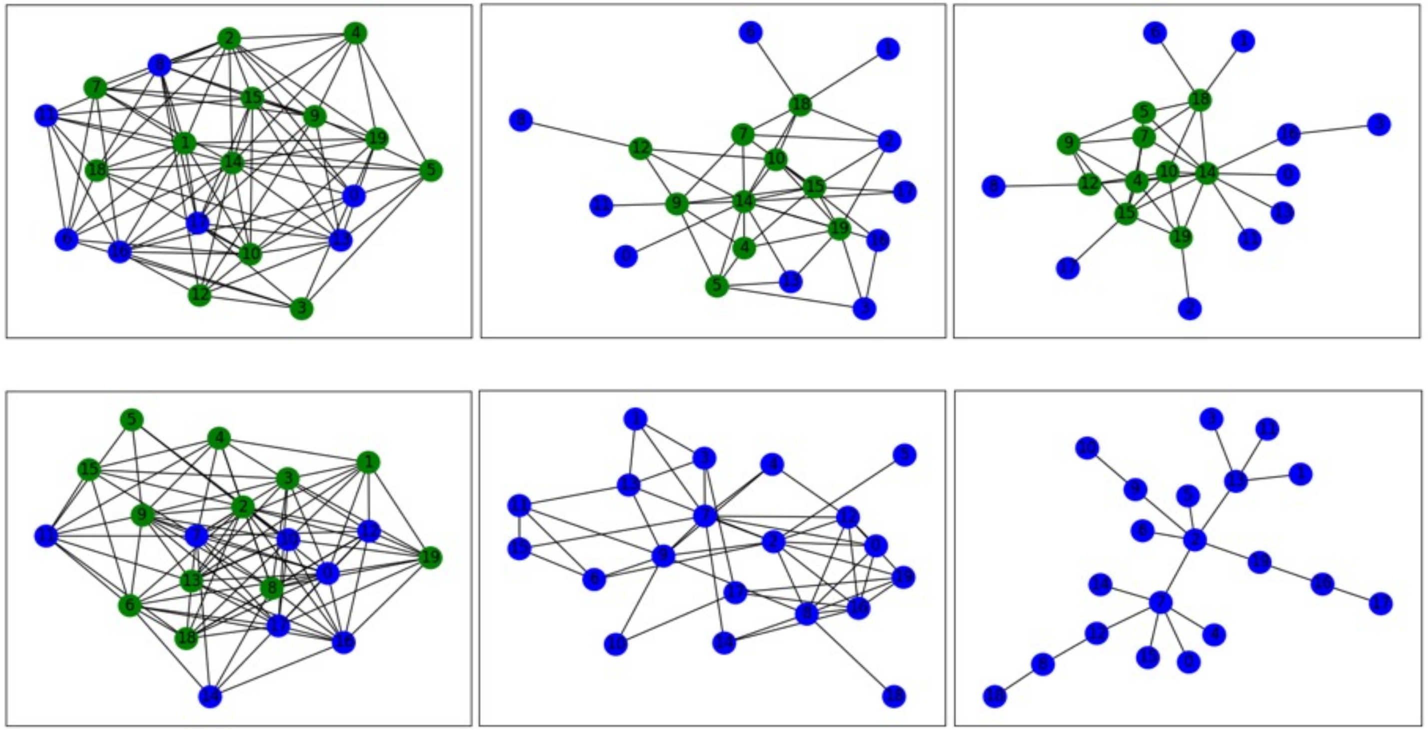

Early phase

Middle

Late phase

Governance through Network Design in multi-equilibria systems

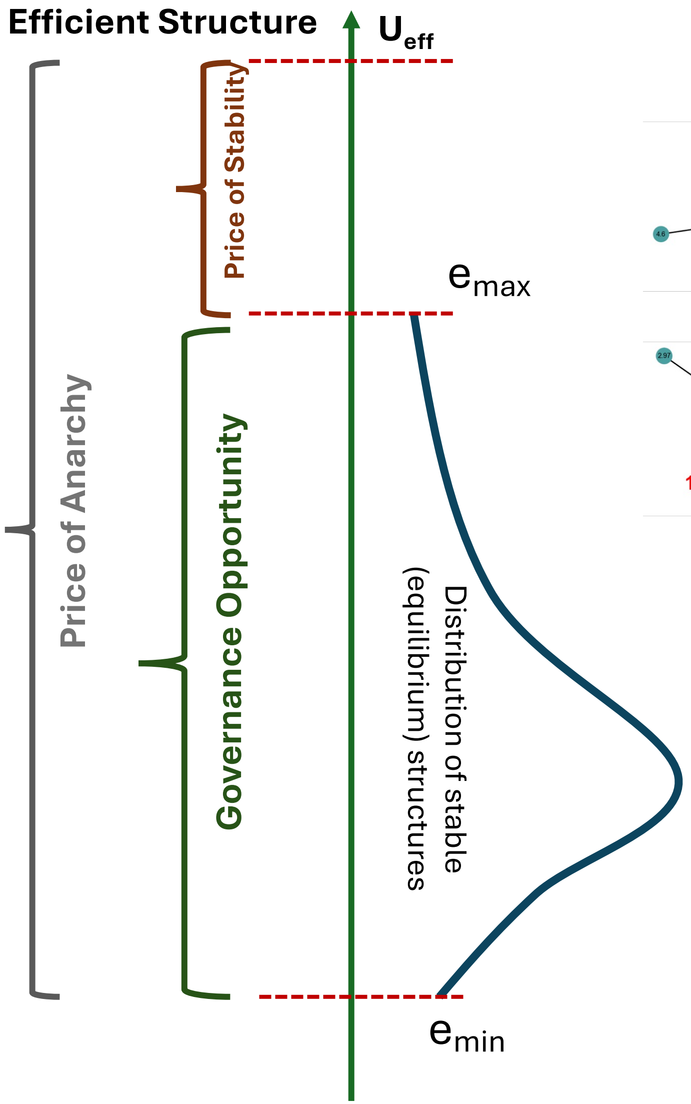

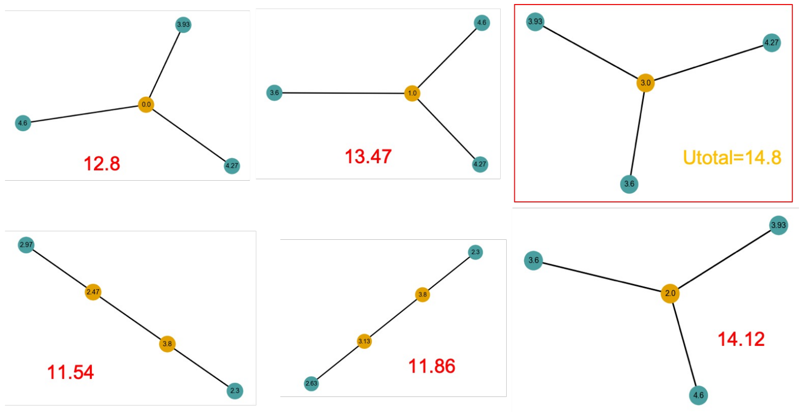

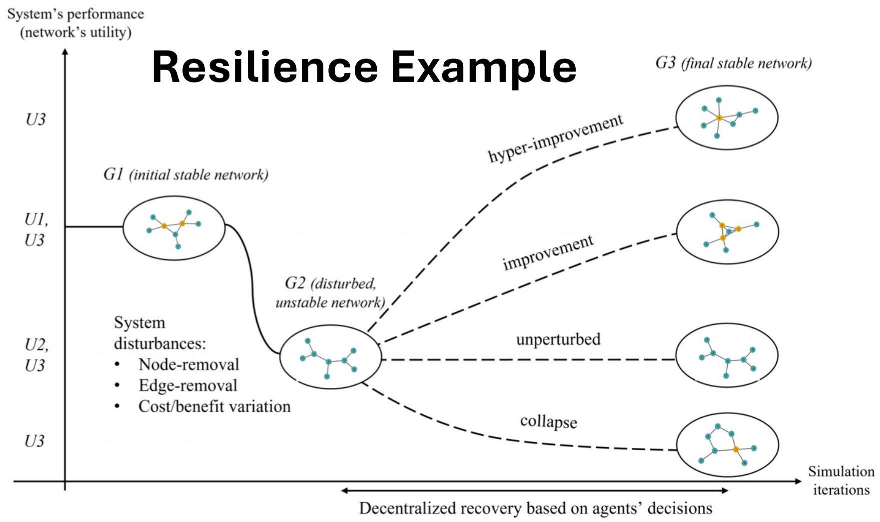

Figure 2: Different elements and scenarios for decentralized recovery analysis for networked systems: Three stages (G1-G3) and example corresponding network structures.

Let $\mathcal{E}$ be the set of all possible equilibrium states, e^* be the set of desired equilibria, e_0 be the initial state, $C(e_0, e)$ be the transition cost from state e_0 to state e, and $\mathcal{P}$ and $\mathcal{C}$ be sets of properties and costs, respectively.

We formally define $G'(\mathcal{S})$ as follows:

$$G'(\mathcal{S}) = \min_{e \in \mathcal{E}} \sum_i w_i C_i(e_0, e), \quad subjectto \quad e \in e^*, \quad \forall p \in \mathcal{P}.$$

Maddah, Negin, and Babak Heydari. "Building back better: Modeling decentralized recovery in sociotechnical systems using strategic network dynamics." Reliability Engineering & System Safety 246 (2024): 110085.

Future Directions: Digital Twins of Sociotechnical systems using (LLM+SAG)

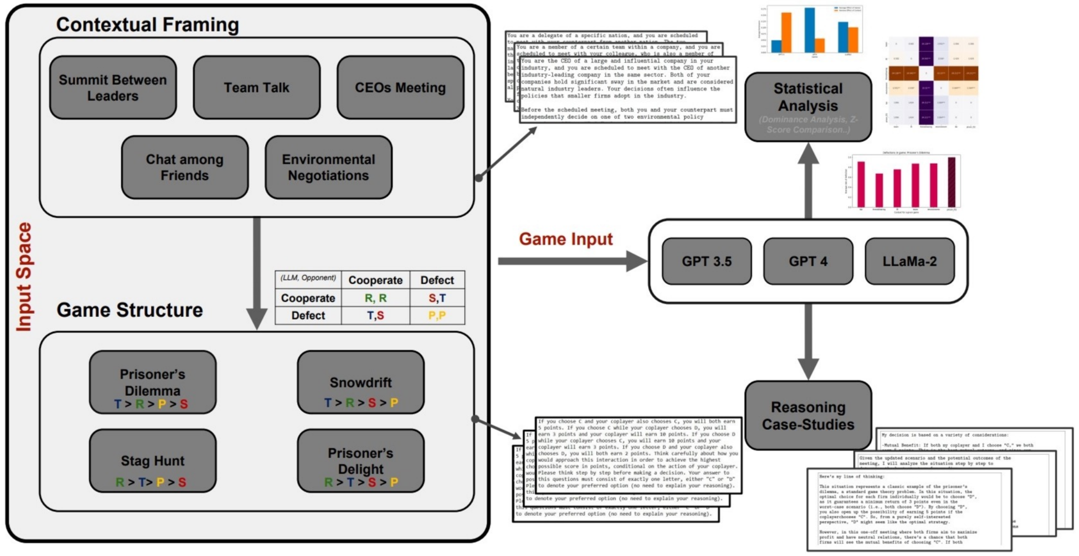

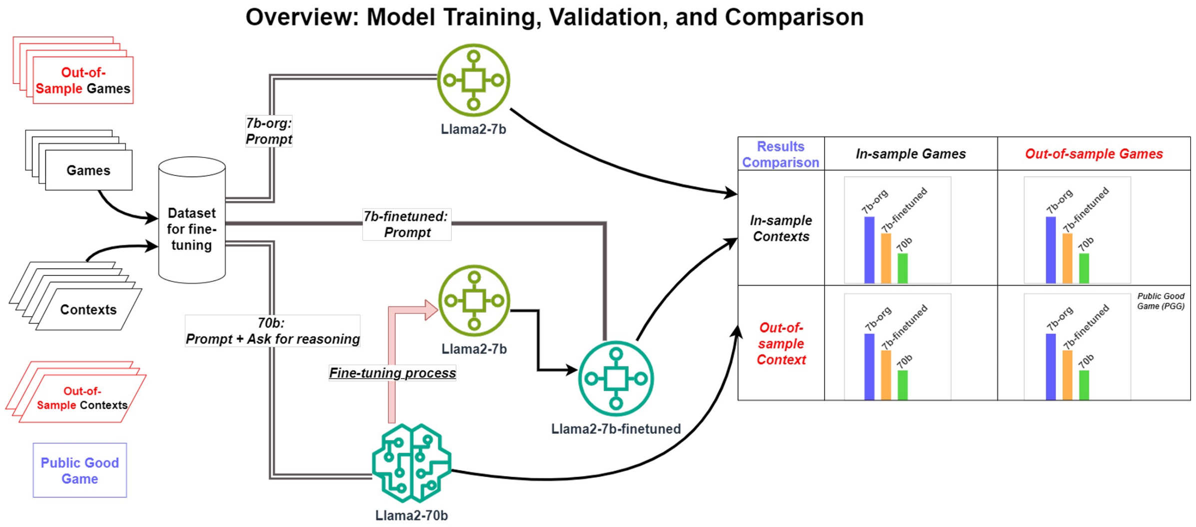

Lorè, Nunzio, and Babak Heydari. "Strategic behavior of large language models and the role of game structure versus contextual framing." *Scientific Reports* 14.1 (2024): 18490.

Lore, Nunzio, and Babak Heydari. "Large Model Strategic Thinking, Small Model Efficiency: Transferring Theory of Mind in Large Language Models." *arXiv preprint arXiv:2408.05241* (2024).

We can integrate interactive and strategic behavior of social agents into digital twins of complex systems to enable the next generation of SAG.

Some References

- Chen, Qiliang, and Babak Heydari. "The SoS conductor: Orchestrating resources with iterative agent-based reinforcement learning." *Systems Engineering* (2024).
- Chen, Qiliang, and Babak Heydari. "Dynamic Resource Allocation in Systems-of-Systems Using a Heuristic-Based Interpretable Deep Reinforcement Learning." *Journal of Mechanical Design* 144.9 (2022): 091711.
- *Maddah, Negin, and Babak Heydari. "Building back better: Modeling decentralized recovery in sociotechnical systems using strategic network dynamics." Reliability Engineering & System Safety 246 (2024): 110085.*
- Chen, Qiliang, and Babak Heydari. Adaptive Network Steering: Combining Variational Autoencoders and Reinforcement Learning for Multi-Agent Resource Allocation, Under Review at IEEE. T. Network Sci. and Eng.(2024)
- Heydari, Babak, and Michael J. Pennock. "Guiding the behavior of sociotechnical systems: The role of agent-based modeling." *Systems Engineering* 21.3 (2018): 210-226.
- Lorè, Nunzio, and Babak Heydari. "Strategic behavior of large language models and the role of game structure versus contextual framing." *Scientific Reports* 14.1 (2024): 18490.
- Lore, Nunzio, and Babak Heydari. "Large Model Strategic Thinking, Small Model Efficiency: Transferring Theory of Mind in Large Language Models." *arXiv preprint arXiv:2408.05241* (2024).

Integrated Perspective of Sociotechnical systems