CSC 59866-E: Senior Project I
AI Agents for Decision Making in the Real World

By Saptarashmi Bandyopadhyay
Email: sbandyopadhyay@ccny.cuny.edu, sbandyopadhyay@gc.cuny.edu
Assistant Professor of Computer Science
City College of New York and Graduate Center at the City University of New York

May 4, 2026 CSC 59866

Advanced Topics: Human-Agent and Agent-Agent Coordination and Competition

Logistics & The State of the Class

Recall Lecture 23: We dealt with *some* messy real-world problems including in finance, supply chain logistics, and robotics.

Even then, almost all of the environments we looked at involved either single-agent or purely cooperative multi-agent settings.

Today we will explore some of the theory behind how AI Agents interact not just with other AI Agents, but with *humans* as well.

Today's Agenda

1. **Agent-Agent Coordination & Scaling:** Topologies, Dec-POMDPs, and the limits of scale (Tran et al., 2025; Kim et al., 2026).
2. **Agent-Agent Competition & Mixed Motives:** Strategic deception and "Cooperate to Compete" (O'Neill et al., 2026).
3. **Human-Agent Collaboration:** Zero-shot adaptation, generative partner modeling, and goal abstractions (Liang et al., 2024; Long et al., 2026; Tankelevitch & Rintel, 2026).

Agent-Agent Coordination & Scaling

Types of Multia-Agent Architectures

There are a few different ways Multi-Agent Systems (MAS) can be organized depending on the task at hand:

In a **Centralized** approach, one *Manager* Agent coordinates with multiple worker agents, dictating tasks to them. This creates a highly aligned system that is nonetheless prone to single-point-of-failure issues and communication bottlenecks with the Manager.

In **Peer-to-Peer (Decentralized)** Agent networks, there is no single point of failure but now there's the possibility for infinite loops and failures during the consensus process.

A hybrid approach, **Hierarchical** Agents, has local leaders that form squads with other agents to reduce the overhead from communication while still effectively accomplishing tasks.

Recall: Dec-POMDPs

How do we formalize decentralized agents cooperating under uncertainty?

Recall: Decentralized Partially Observable Markov Decision Processes (Dec-POMDPs) are defined by the tuple $\langle I, S, A, P, R, \Omega, O, \gamma \rangle$

- I : A finite set of agents $i \in \{1, \ldots, N\}$.
- $A = \times_i A_i$: The joint action space.
- $\Omega = \times_i \Omega_i$: The joint observation space.
- $P(s'|s, \vec{a})$: State transition probabilities.
- $R(s, \vec{a})$: A *shared* global reward function.

The Challenge: Agents do not share O_i perfectly. They must learn a decentralized policy $\pi_i(a_i|h_i)$ based strictly on their local history h_i, making exact optimization NEXP-complete.

Towards a Science of Scaling Agent Systems

If one agent is smart, are ten agents ten times smarter?

The Scaling Law of Agents: Kim et al. (2026) developed a predictive model capturing how performance varies with coordination and model capability across benchmarks.

Their study evaluated 260 configurations, finding a robust R^2 = 0.373 correlation dictating system performance.

The Finding: Adding more agents reliably improves performance *only up to a critical threshold*, after which capability saturation occurs.

The Cost of Coordination

Why do multi-agent systems face diminishing returns?

Communication Overhead: In a fully connected peer-to-peer MAS, the number of communication channels scales at $O(N^2)$.

Context Dilution: As agents debate, the context window fills with redundant or argumentative tokens, burying the actual logical reasoning.

We must mathematically penalize unnecessary communication to optimize inference efficiency.

Memory Management: Context Compression

To mitigate the "coordination tax" during long ideation sessions, agents must actively prune their memory (Quan et al., 2025).

Persona-Based Compression: When context length exceeds a threshold T (e.g., 15 turns):

1. The k most recent turns are preserved verbatim.
2. The T-k older turns are compressed via a Multi-Chat Summary Prompt.
3. User/Human inputs are weighted higher and preserved, while verbose AI-to-AI debates are aggressively collapsed into <200 token summaries.

Agent-Agent Competition & Mixed Motives

The Mixed-Motive Dilemma

Most MARL research focuses on purely cooperative tasks (shared reward) or purely zero-sum games (Chess/Go).

Mixed-Motive Environments: Settings where agents must leverage short-term cooperation for long-term competitive goals (e.g., politics, corporate negotiation).

In these environments, alliances are temporary, and communication channels are frequently used for *deception* rather than coordination.

Cooperate-to-Compete

O'Neill et al. (2026) introduced the C2C environment: agents engage in private, non-binding negotiations while competing to achieve secret objectives.

Asymmetric Objectives: Agent A needs Resource X. Agent B needs Resource Y. They can trade, but helping B might accidentally give B the win.

Because negotiations are *non-binding*, agents can agree to a truce and then immediately backstab their partner on the next turn.

Humans vs. AI in Strategic Negotiation

How do LLM agents differ from humans in mixed-motive conquest? (O'Neill et al., 2026)

Human Profiles: Humans favor lower-complexity deals, are more aggressive, and are significantly *less reliable partners*. Humans accept deals without a counteroffer only 56.3% of the time.

AI Profiles: LLM agents (out-of-the-box) are highly compliant, naive, and overly cooperative, accepting bad deals 67.6% of the time.

Prompt Engineering for Cutthroats: By explicitly prompting agents with "Aggressive Negotiation" and "Deceitful" personas, their win rates in mixed-motive games jump from 22.2% to 32.7%.