

The Impact of Information Flows on Multi-Agent Systems

Tim Cahoe
timothy.m.cahoe@nps.edu
Naval Postgraduate School
1 University Circle, Monterey, CA 93943

Kasey Miller
millerkc@uncw.edu
University of North Carolina Wilmington
Wilmington NC 28403

Tim Shives
timothy.shives@nps.edu
Naval Postgraduate School
1 University Circle, Monterey, CA 93943

Abstract

Multi-Agent Systems (MAS) are increasingly employed in autonomous and distributed applications, yet many continue to underperform in collaborative task execution. This study investigates whether information flow architecture is a key determinant of MAS performance across five commonly used agent interaction strategies. A Monte Carlo simulation consisting of 1,500 experimental runs was conducted in a Python-based environment, in which 10 autonomous agents collaboratively mapped a bounded area under three communication conditions: high, low, and no communication. Performance differences were evaluated using regression and Risk Simulator diagnostics to assess the effect of information flow on a composite of map coverage and steps to completion. Results demonstrate that communication architecture significantly influences system performance, although the effects vary by coordination strategy. Increased communication frequency substantially degraded performance in flocking and consensus-based models ($p < .001$) but significantly improved task completion in the Artificial Bee Colony framework. These findings suggest that communication intensity should be matched to the underlying coordination algorithm rather than uniformly increased. The results provide empirical guidance for configuring distributed autonomous systems under constrained communications and establish a quantitative foundation for future live-flight unmanned aerial vehicle (UAV) teaming experiments.

Keywords: Multi-Agent Systems, Information Flow, Monte Carlo Simulation, Autonomous Systems, UAV Teaming, Swarm Intelligence, Artificial Bee Colony, Distributed Coordination

The Impact of Information Flows on Multi-Agent Systems

Tim Cahoe, Kasey Miller and Tim Shives

1. INTRODUCTION

Hughes & Girrier (2018) and Kline (2017) note that naval warfare is transitioning from the missile age to an information and robotics age. Groups of small unmanned aerial vehicles (UAVs) and unmanned surface vessels (USVs) are part of the vision for a hybrid fleet that integrates autonomous vehicles with traditional naval assets. This hybrid operational concept, developed by Kline (2017), is currently being implemented in strategic naval initiatives. Autonomous machine groups interact similarly to general machine groups but introduce additional complexity through factors such as goal-seeking, reasoning, and machine learning, which can lead to extensive feedback loops and emergent behaviors. Russell (2021) observes that contemporary autonomous systems struggle to interpret complex human intent and often optimize assigned objective functions at the expense of secondary constraints or human values. Arquilla (2000) notes that information flows are among the key factors to be explored in understanding swarm technology, yet there is a critical lack of quantitative empirical studies examining whether increasing the frequency of information flow universally benefits or actively degrades Multi-Agent Systems (MAS) task performance across divergent algorithmic architectures. The intersection of translating human goals to MAS and the role of information flows in MAS performance offers a promising area of study, not only for the US Navy but also for companies like OpenAI or xAI pursuing LLM-based MAS.

Problem Statement

The US Navy plans to deploy autonomous groups of UAVs and USVs in communications-restricted environments, yet it lacks empirical frameworks to match algorithmic architectures to degraded network conditions risking systemic failure if bandwidth-dependent models are deployed under contested communications. This Navy problem is a direct reflection of a larger information systems problem that plagues modern multi-agent LLM research. The problem is we do not understand how multi-agent systems may be configured to accomplish human goals with minimal to no human intervention.

Purpose Statement/Study Proposal

This study examines how information flow affects MAS performance across several promising MAS methodologies. This empirical investigation was conducted using custom Python-based simulation models and Risk Simulator diagnostic analytics to evaluate agent performance across communication tiers, laying the groundwork for upcoming live MAS UAV field experiments. This study will provide value by demonstrating the impact of information flows on MAS performance across a variety of scenarios, yielding generalizable knowledge. The results from each experiment type will be compared using statistical analyses, including ANOVA, t-tests, and regression analysis.

2. LITERATURE REVIEW

Rizk et al. (2018) define Multi-Agent Systems (MAS) as networks of autonomous entities operating with shared or conflicting goals based on distributed sensor data. Key to MAS is the coordination and information flows that exist between agents (Wooldridge & Jennings, 1995). Tang (2021) notes over 80,000 articles on Google Scholar on this topic at the time of writing, with multiple theses and dissertations existing on the topic and a vast number of approaches to coordination and information flows. However, in terms of Navy applications Arquilla (2000) and Hughes & Girrier (2018) note a gap in how information flows impact MAS for Naval warfare and Hughes & Girrier (2018) and Kline (2017) note the importance of the use of UAVs in modern Naval warfare. Likewise, more generally in the field Cao et al. (2024) notes that adaptive cognitive networking is necessary in dynamic environments. Yet, Cattai et al. (2025) and Bayer & Faigl (2026) note that traditional approaches to UAV MAS rely on high bandwidth ideal environments for group functioning which doesn't align to reality where robots may be bandwidth constrained and unable to pass data they have been designed to rely on, such as complex 3D mapping. Further, Katsiuba et al. (2025) notes the need for standard operational procedures and algorithms that account for the dynamic needs of information flow protocols in dynamic drone communication. On a similar note Garzón et al. (2025) and Dietz et al. (2023) point out that human

designers show bias in selecting communication protocols for tasks rather than allowing optimization techniques to provide dynamic selection logic for best fit protocols to the situation. The value of dynamic information flow method shifts also extends to the advent of MAS LLM systems, notable recent applications include Allmendinger et al. (2026) 5 layer Multi-agent Artificial Intelligence (MAAI) workflow that dynamically selects appropriate agent communication for the task, and Kim et al. (2026) who present an approximate linear threshold of 45% for when cognitive load of an agent is benefitted and not hindered by multi-agent information flows.

Game theory and utility theory are classic fields dating back to at least Nash (1951) and Bentham (1789), and unique solutions have been proposed that leverage quantified utility and Nash equilibrium to develop optimal solutions for MAS coordination and information flow. Swarm technology is a subset of MAS (Brambilla et al., 2013). Swarm systems typically follow an emergent, nature-based model, such as modeling the behavior of birds (Reynolds, 1987), particles (Kennedy & Eberhart, 1995), bees (Karaboga, 2005), and ants (Dorigo, 1992). Swarm systems also have more human-like information-sharing and coordination algorithms that can be either highly decentralized or hierarchical (Arquilla & Ronfeldt, 2000). Finally, more recent MAS studies examine LLM-based MAS and unique ways natural language processing can be used to produce coordinated machine outcomes.

Utility Equations and Game Theory for MAS Coordination

Russell (2021) and others have noted that utility equations are a good angle for how to model complex goals and values in machines. The philosophical foundations of utility theory trace back to Bentham's (1789) ethical frameworks. In computational decision-making, modern utility theory originated with von Neumann and Morgenstern (1944) and has evolved significantly alongside algorithmic complexity. Utility theory seeks to value decisions mathematically to represent human goal seeking (Von Neumann & Morgenstern, 1944). Unique solutions have been developed for MAS utility applications. Koutsoupas and Papadimitriou (2009) suggest a framework for quantifying stability vs. anarchy (Price of Anarchy) in a decision process. Rădulescu et al. (2019) developed the Multiobjective multi-agent systems (MOMAS) to value compromises in complex multi-goal scenarios where agents may be assigned unique goals for the tasks they are assigned in the MAS. In review of a large range of MAS research Paccagnan et al. (2022) notes the need to design and optimize utility values for individual agents for system-wide benefit and Yang & Liu (2024) suggest viewing utility within a framework of an agentic needs hierarchy developing trust and an understanding of other agents in the MAS to better coordinate for common goals. Game theory covers complex interactions between autonomous agents and has been significantly utilized in computer science models for machine interaction. Initially theorized by Von Neumann (1944) and later modified by Nash (1951). Game theory offers a framework for making optimal decisions with multiple agents with limited information (Paccagnan et al., 2022), often involving reaching an equilibrium for minimizing the cost of a solution (Rădulescu et al., 2019).

Swarm Coordination Methodologies

Swarms are a subset of MAS that operate cooperatively under the direction of a single operator-in-charge through a common architecture (Price, 2025). Shahzad et al. (2023) note that a swarm requires at least 3 agents, limited human control, and that its members must communicate and work cooperatively. Many swarm methodologies draw inspiration from social animals (Brambilla et al., 2013; Tan & Zheng, 2013), including particles, fish, birds, bees, and ants (Tang et al., 2021). These methodologies solve complex problems such as pathfinding, reconnaissance, and search-and-rescue (Brambilla et al., 2013). Other methodologies include decentralized and centralized control approaches, in which one or more agents are responsible for the cooperative assignments and information-sharing procedures for sub-agents (Kinnison et al., 2022). Gulosh (2018) notes remarkable improvements in decentralized hierarchical drone performance in support of military operations, reducing casualties by up to 50% compared to emergent solutions. Despite operational benefits, Kinnison et al. (2022) note hierarchical solutions add operational complexity for human operators and recommend developing autonomous task allocation techniques to reduce the cognitive burden.

LLM Agentic MAS Coordination

Agentic AI using LLMs is a unique modern solution in MAS, where the use of natural language and agent identity enables novel agent communication and coordination (Schmidgall et al., 2024). Schmidgall et al. (2024) & Xi et al. (2024) note the need to train agents in their roles to optimize performance. Li et al. (2023) & Karten et al. (2025) suggest a voting framework in which the MAS follows a voting process, and Li et al. (2023) note that improvements in performance on complex tasks correlate with adding

more agents to the voting pool. These solutions are emerging and show promise in addressing issues experienced with MAS to date (Zhu, Du, Hong et al., 2025).

Hypothesis for this Study

HA: Information flow frequency has a statistically significant impact on task performance across differing MAS methodologies.

H0: Information flow frequency has no statistically significant impact on task performance across differing MAS methodologies.

Further, we hypothesize that the direction and magnitude of information flow impact on task performance are significantly dependent on the underlying MAS coordination methodology.

3. METHODOLOGY

For this analysis five methodologies were chosen based on common models utilized for MAS cooperation. The MAS methodologies were adapted to cover a two-dimensional box, shown in Figure 1 below, and to communicate with one another based on the information flow of each methodology. In figure 1 density maps are shown of representative runs from high and low communications runs for each of the MAS methodologies, comms represents the specific number of communications that took place as a function of the steps taken, the communication frequency set, and the methodology being used. Each methodology was tested at high communications, low communications, and no communications and a Monte Carlo of 100 runs was performed for each combination of communications and MAS methodology, resulting in 300 runs per MAS methodology and a total of 1500 runs.

The following five foundational swarm algorithms were implemented in Python, governed by their respective mathematical behavioral rules:

1. **Reynolds' Boids Model: One of the original MAS methodologies based on bird flocking patterns (Reynolds, 1987).**

$$\vec{v}_a(t+1) = \vec{v}_a(t) + w_s \vec{f}_{sep} + w_a \vec{f}_{ali} + w_c \vec{f}_{coh}$$

2. **Vicsek Model: Self-propelled models that move based on the average direction of their neighbors (Hernandez-Lopez, 2010).**

$$\theta_a(t+1) = \text{Arg} \left(\frac{1}{|N_a|} \sum_{b \in N_a} \vec{v}_b \right) + \Delta \theta$$

3. **Ant Colony Optimization (ACO): Ant pheromone trails are laid by leading ants, providing direction, ideal for path search applications (Tang et al., 2021).**

$$\tau_{xy}(t+1) = (1 - \rho) \tau_{xy}(t) + \sum_{a=1}^N \Delta \tau_{xy}^a$$

4. **Particle Swarm Optimization (PSO): Where particles change position based on personal and global best solutions (Tang et al., 2021).**

$$\mathbf{v}_i(t+1) = w \mathbf{v}_i(t) + c_1 r_1 (\mathbf{p}_i^{\text{best}} - \mathbf{x}_i(t)) + c_2 r_2 (\mathbf{g}^{\text{best}} - \mathbf{x}_i(t))$$

5. **Artificial Bee Colony (ABC) Algorithm: Where employed, onlooker, and scout bees use dancing to draw other bees to ideal exploration paths (Tang et al., 2021)**

$$\vec{p}_{\text{cand}} = \vec{p}_a(t) + \vec{\phi} \odot (\vec{p}_a(t) - \vec{p}_j(t))$$

For the Python code used in this experiment, the communication frequency is unique to each chosen MAS methodology. Tables 1a through 1e describe how information flow frequency (0.0 to 1.0) was operationalized across each algorithm. Low communication runs were set to 0.01, high communication runs were set to 0.8 and no communication runs were set to 0. The no-communication condition (comm_freq = 0) disables the coordination channel for that algorithm and reduces every agent to an independent random walk, and is included only as a no-coordination baseline against which low and high communication are compared. Communication frequency levels were selected to bracket operationally relevant extremes rather than to characterize a continuous dose-response relationship. A low-communications condition (0.01) representing sparse, intermittent coordination consistent with Navy Emissions Control (EMCON) or heavily degraded tactical networks, and a high-communications condition (0.8) representing near-continuous connectivity while preserving a residual margin of independent agent behavior to avoid conflating comms availability with the complete absence of exploratory movement.

1. Boids

comm_freq	Practical translation
0.8	Neighbor-awareness "refreshes" ~80% of moves; ~20% of the time agents have no awareness of separation/alignment/cohesion.
0.01	Flocking behavior only ~1 in 5 steps.
0.0	Random walk baseline, no coordination.

Table 1a Boids Information Flow Levels for Simulation

2. Vicsek

comm_freq	Practical translation
0.8	Heading updates happen on 80% of moves; agent continues on last heading (plus noise) the rest of the time.
0.01	Rare re-alignment; agents behave close to independent random walkers with occasional consensus correction.
0.0	Random walk baseline, no coordination.

Table 2b Vicsek Information Flow Levels for Simulation

3. Ant Colony Optimization (ACO)

comm_freq	Practical translation
0.8	Pheromone map is refreshed 80% of moves; evaporation and deposit both pause on the other 20%, so the map is "frozen" intermittently rather than degrading gradually.
0.01	Environment-mediated communication is rare; the pheromone map barely changes, so agents get little benefit from others' recent exploration.
0.0	Random walk baseline, no coordination.

Table 3c ACO Information Flow Levels for Simulation

4. Particle Swarm Optimization (PSO)

comm_freq	Practical translation
0.8	Global best is recomputed on 80% of moves; agents continue on last-known velocity vector 20% of the time.
0.01	Centralized reporting is sparse; agents rarely learn where the best-performing agent is, so exploration looks much more like independent search.
0.0	Random walk baseline, no coordination.

Table 4d PSO Information Flow Levels for Simulation

5. Artificial Bee Colony (ABC)

comm_freq	Practical translation
0.8	Both random-partner exploration and onlooker fitness-broadcast happen 80% of ticks; agents idle the rest independently.

comm_freq	Practical translation
0.01	Hive-level information sharing is sparse.
0.0	Random walk baseline, no coordination.

Table 5e ABC Information Flow Levels for Simulation

The following variables were used in this experiment:

Intervening variables (not changed)

- Number of Agents: 10
- Size of Map: 5 x 5 map
- Max Steps: 40000 (ACO) 100000 (PSO)


Runs ended when grid coverage reached 100 percent or when an algorithm-specific iteration bound was met. ACO was bounded at 40,000 iterations, a window within which pheromone repulsion still reaches full coverage. PSO was bounded at 100,000 iterations to test whether further steps would overcome stagnation from attraction to the global best; high-communication runs still ended at that bound with incomplete coverage. Boids, Vicsek, and ABC used the same coverage rule. Because each method used one bound across its no, low, and high-communication runs, the within-method slopes are not an artifact of the unequal ceilings. Of note the ceilings also do not produce the ABC high-communication advantage since the cells nearest that result finished on the coverage break, and the censored ACO and PSO high-communication cells remain well below it.

Independent Variables

- MAS Methodology
- Frequency of Communication: a number between 0 and 1 that represents a range of communications
- Communication instances: The number of communications that take place in a scenario as a function of the MAS methodology and the communication frequency set.

Dependent Variable

- Performance = ((mean steps taken for that methodology)/steps taken) x (% complete of map). Weighted performance is that score multiplied by a method-specific constant (1.7236 for ABC), not an average across methods.






Figure 1: Sample Run from Each Run of the 2D Map Covered by Each MAS Methodology. Cells are color-coded by Visit Count. Blue dots are the Final Locations of the Agents.

4. ANALYSIS AND DISCUSSION

Performance was regressed on the specific communications count in Risk Simulator, not on the assigned frequency. Each additional message was associated with lower performance for Boids, Vicsek, PSO, and ACO ($t = -46.61$ to -14.49) and higher performance for ABC ($t = 35.08$). All five coefficients and the accompanying ANOVA tests were significant at $p < .001$. R^2 ranged from 0.413 for PSO to 0.879 for ACO. On weighted performance, low communication differed little from none for Boids (1.335 versus 1.303) and Vicsek (1.314 versus 1.281). Table 2 regresses unweighted performance on the message

count. Table 3 and the ABC mean of 6.757 use weighted performance.

ACO declined from 1.300 with no communication to 1.209 at low communication and to 0.441 at high communication. PSO declined at both steps, from 1.341 to 0.281 to 0.019. A Mann-Whitney check on PSO confirmed that ordering: no communication (median 1.34) exceeded low communication (median 0.24), and low exceeded high (median 0.02), $U = 0$, $z = 12.22$, $p < .001$ in both contrasts. Results of this analysis are shown in Table 2. ABC, however, showed a significant increase in performance when going from no to high communication. In terms of performance, ABC outperformed all other methodologies and communication combinations when operating in high communication mode. Table 3 shows Games-Howell comparisons of ABC high-communication performance ($M = 6.757$, $SD = 1.057$, $N = 100$) against every other method-by-communication cell. A positive difference means ABC high comms scored higher. All 14 contrasts are significant at $p < .001$. Figure 2 graphically represents average performance score differences.

Boids and Vicsek held their performance when communication was removed. Low versus none was not significant ($p = .198$ and $p = .211$). High communication reduced performance. They did not beat the no-communication baseline. The ABC algorithm, however, exhibited a contrasting effect where high-frequency, accurate information exchange dramatically increased performance, resulting in statistically superior task completion compared to all other methodologies and communication combinations.

Methodology	Slope per message	Per 1,000 messages	t (df 298)	p	R ²	ANOVA F
Boids	-0.00001967	-0.020	-24.36	< .001	0.666	593.64
Vicsek	-0.00001603	-0.016	-15.89	< .001	0.459	252.45
PSO	-0.00001852	-0.019	-14.49	< .001	0.413	209.93
ACO	-0.00002073	-0.021	-46.61	< .001	0.879	2,172.07
ABC	0.001149	1.149	35.08	< .001	0.805	1,230.44

Table 2: Summary of Results

Model Input Comparison	Sample 1 Median	Sample 2 Median	Test Statistic (U)	z-score	p-value
PSO No Comms vs. PSO Low Comms*	1.34	0.24	0	12.22	<0.001
PSO Low Comms vs. PSO High Comms*	0.24	0.02	0	12.22	<0.001

Table 2a: Summary of Results, Mann-Whitney Follow Up for PSO




Figure 2: Comparison of Average Performance, ABC High Comms Significantly Performed Far Better than Any other Method Comms Combination

Comparison group	Group mean	Difference	t	df	P
ABC low comms	1.252	5.506	51.29	105.3	< .001
ABC no comms	1.322	5.435	50.76	104.3	< .001
ACO high comms	0.441	6.316	59.76	99.1	< .001
ACO low comms	1.209	5.549	51.97	103.1	< .001
ACO no comms	1.3	5.458	50.88	105	< .001
Boids high comms	0.915	5.843	54.42	105.3	< .001
Boids low comms	1.335	5.422	50.62	104.4	< .001
Boids no comms	1.303	5.454	50.94	104.2	< .001
PSO high comms	0.019	6.739	63.77	99	< .001
PSO low comms	0.281	6.477	61.08	100.3	< .001
PSO no comms	1.341	5.416	50.22	107.2	< .001
Vicsek high comms	1.005	5.752	53.82	103.5	< .001
Vicsek low comms	1.314	5.444	50.69	105.5	< .001
Vicsek no comms	1.281	5.477	51.06	105	< .001

Table 3: Summary of Results, Games-Howell

From these tests we can reject the null hypothesis that information-flow frequency has no statistically significant association with task performance. In the separate linear regressions (Table 2), each slope differs from zero at $p < .001$. The claim that the direction of that association depends on methodology is supported by the slope signs: positive for ABC ($t = 35.08$) and negative for Boids, Vicsek, ACO, and PSO (t from -14.49 to -46.61). Mann-Whitney tests (Table 2a) show the same PSO decline at both steps, from no communication to low and from low to high (both $p < .001$). PSO often reached the 100,000-step ceiling, the near-zero variance in the high-communication cell indicates this is likely because frequent global-best updates pulled agents onto the same reported location and collapsed coverage into repeated visits rather than new cells. Games-Howell contrasts (Table 3) corroborate that high-communication ABC exceeded every other method-by-communication cell by 5.42 to 6.74 points.

5. LIMITATIONS

Although the baseline algorithms were not highly optimized for spatial mapping, the use of standardized, unconstrained behavioral rules enhances internal validity by isolating the inherent explanatory power of communication frequency, free from the influence of algorithmic heuristics. However, this experiment has limitations, as it focuses on a very specific task, which may favor certain methodologies over others, thereby constraining the external validity of the findings. Additionally, the information flows in this simulation are relatively simplistic and could benefit from a more comprehensive literature review on best practices for each methodology.

6. CONCLUSIONS & FUTURE WORK

Within this 2D spatial-coverage simulation, information flow architecture impacts Multi-Agent System (MAS) performance and its effects depended heavily on the underlying coordination algorithm. Results from 1,500 Monte Carlo simulation runs indicate that increasing communication frequency does not universally improve task performance. Instead, excessive information exchange significantly degraded performance in flocking and consensus-based models, while substantially improving performance in the Artificial Bee Colony (ABC) framework. These findings suggest that communication strategies should be designed in conjunction with task allocation and coordination algorithms rather than treated as independent design parameters. Consequently, effective MAS design requires a systems perspective that aligns information flow, agent behavior, and mission objectives.

This work provides an empirical foundation for the next phase of research, which will extend simulation results to live-flight unmanned aerial vehicle (UAV) experiments involving human-machine teaming during search-and-rescue missions. Future studies will evaluate whether the relationships observed in simulation persist in operational environments by testing the following hypotheses:

- HA1: Information flow architecture significantly affects autonomous team performance.
- H01: Information flow architecture does not significantly affect autonomous team performance.
- HA2: Autonomous team performance significantly affects overall human-machine team performance during a man-overboard search-and-rescue operation.
- H02: Autonomous team performance does not significantly affect overall human-machine team performance during a man-overboard search-and-rescue operation.

Validating these hypotheses in real-world operational environments will provide evidence to inform the design of adaptive communication architectures that improve the effectiveness, robustness, and scalability of autonomous and human-machine teams.

7. REFERENCES

Allmendinger, S., Bonenberger, L., Endres, K., Fetzer, D., Gimpel, H., & Kühn, N. (2026). Multi-agent AI. *Electronic Markets*, 36, Article 18. <https://doi.org/10.1007/s12525-025-00862-z>

- Arquilla, J., & Ronfeldt, D. (2000). *Swarming and the future of conflict*. RAND Corporation. https://www.rand.org/pubs/documented_briefings/DB311.html
- Bayer, J., & Faigl, J. (2026). Decentralized multi-robot exploration under low-bandwidth communications. *Autonomous Robots*, 50(1), Article 7. <https://doi.org/10.1007/s10514-025-10234-3>
- Bentham, J. (1789). *An introduction to the principles of morals and legislation*. T. Payne and Son.
- Brambilla, M., Ferrante, E., Birattari, M., & Dorigo, M. (2013). Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence*, 7(1), 1–41. <https://doi.org/10.1007/s11721-012-0075-2>
- Cao, P., Lei, L., Cai, S., Shen, G., Liu, X., Wang, X., Zhang, L., Zhou, L., & Guizani, M. (2024). Computational intelligence algorithms for UAV swarm networking and collaboration: A comprehensive survey and future directions. *IEEE Communications Surveys & Tutorials*, 26(4), 2110–2150. <https://doi.org/10.1109/COMST.2024.3395358>
- Cattai, T., Frattolillo, F., Lacava, A., Raut, P., Simonjan, J., D'Oro, S., Melodia, T., Vinogradov, E., Natalizio, E., Colonnese, S., Cuomo, f., & Iocchi, L. (2025). Multi-UAV reinforcement learning with realistic communication models: Recent advances and challenges. *IEEE Open Journal of Vehicular Technology*, 6, 1200–1225. <https://doi.org/10.1109/OJVT.2025.3586774>
- Dietz, E., Davis, D., & Hale, B. (2023). Utilizing the Messaging Layer Security Protocol in a lossy communications aerial swarm. In *Proceedings of the 56th Hawaii International Conference on System Sciences (HICSS 2023)* (pp. 6591–6600). <https://hdl.handle.net/10125/103431>
- Dorigo, M. (1992). *Optimization, learning and natural algorithms* (Doctoral dissertation, Politecnico di Milano).
- Garzón Ramos, D., Medrano, J. B., Hauert, S., & Birattari, M. (2025). Direct communication or stigmergy? Selecting communication mechanisms via AutoMoDe. In *ALIFE 2025: Proceedings of the 2025 Artificial Life Conference* (Vol. 37, Article 915). MIT Press. https://doi.org/10.1162/isal_a_00915
- Gulosh, N. J. (2018). *Employment of intelligence, surveillance, and reconnaissance drone swarms to enhance ground combat operations* [Master's thesis, Naval Postgraduate School]. Defense Technical Information Center.
- He, Z., Sun, Y., & Feng, Z. (2023). Research on resource allocation of autonomous swarm robots based on game theory. *Electronics*, 12(20), 4370.
- Hernandez-Lopez, R. (2010). *Complex networks and collective behavior in nature*. Massachusetts Institute of Technology. https://web.mit.edu/8.334/www/grades/projects/projects10/Hernandez-Lopez-Rogelio/dynamics_2.html
- Hughes, W. P., Jr., & Girrier, R. P. (2018). *Fleet tactics and naval operations* (3rd ed.). Naval Institute Press.
- Karaboga, D. (2005). *An idea based on honey bee swarm for numerical optimization* (Technical Report TR06). Erciyes University, Engineering Faculty, Computer Engineering Department.
- Karten, S., Li, W., Ding, Z., Kleiner, S., Bai, Y., & Jin, C. (2025). LLM Economist: Large population models and mechanism design in multi-agent generative simulacra. *arXiv preprint arXiv:2507.15815*. <https://arxiv.org/abs/2507.15815>

- Katsiuba, D., Brander, S., Dolata, M., & Schwabe, G. (2025). From ad hoc pilots to repeatable patterns: Structuring drone collaboration in emergency services with DroneLets. In *Proceedings of the International Conference on Information Systems (ICIS 2025)*. Association for Information Systems.
- Kennedy, J., & Eberhart, R. (1995). Particle swarm optimization. In *Proceedings of ICNN'95 - International Conference on Neural Networks* (Vol. 4, pp. 1942–1948). IEEE. <https://doi.org/10.1109/ICNN.1995.488968>
- Kim, Y., Gu, K., Park, C., Park, C., Schmidgall, S., Heydari, A. A., Yan, Y., Zhang, Z., et al. (2026). Capable language models can outgrow the benefits of collaboration. *Nature Machine Intelligence*, 8, 1157–1172. <https://doi.org/10.1038/s42256-026-01268-y>
- Kinnison, M. C., Colter, J. T., Schlager, S., & Harbour, S. (2022). Hierarchical Architecture of Swarm Command and Control Operations (HASCCO). *2022 IEEE/AIAA 41st Digital Avionics Systems Conference (DASC)*, 1–5.
- Kline, J. (2017). Impacts of the robotics age on naval force design, planning, and operations. *Naval War College Review*, 70(3), 67–88.
- Koutsoupias, E., & Papadimitriou, C. (2009). Worst-case equilibria. *Computer Science Review*, 3, 65–69.
- Li, J., Zhang, Q., Yu, Y., Fu, Q., & Ye, D. (2024). More agents is all you need. *Transactions on Machine Learning Research (TMLR) / arXiv:2402.05120*. <https://arxiv.org/abs/2402.05120>
- Nash, J. (1951). Non-cooperative games. *Annals of Mathematics*, 54(2), 286–295.
- Paccagnan, D., Chandan, R., & Marden, J. R. (2022). Utility and mechanism design in multi-agent systems: An overview. *Annual Reviews in Control*, 53, 315–328.
- Price, R. E. (2025). Defining swarm: A critical step toward harnessing the power of autonomous systems. *Military Review*. <https://www.armyupress.army.mil/journals/military-review/online-exclusive/2025-ole/defining-swarm/>
- Rădulescu, R., Mannion, P., Roijers, D. M., & Nowé, A. (2019). Multi-objective multi-agent decision making: A utility-based analysis and survey. *Autonomous Agents and Multi-Agent Systems*, 34.
- Reynolds, C. W. (1987). Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '87)* (pp. 25–34). Association for Computing Machinery. <https://doi.org/10.1145/37401.37406>
- Rizk, Y., Awad, M., & Tunstel, E. W. (2018). Decision making in multiagent systems: A survey. *IEEE Transactions on Cognitive and Developmental Systems*, 10, 514–529.
- Russell, S. J., & Norvig, P. (2021). *Artificial intelligence: A modern approach* (4th ed.). Pearson.
- Schmidgall, S., et al. (2024). AgentClinic: A multimodal agent benchmark to evaluate AI in simulated clinical environments. *arXiv preprint*.
- Shahzad, M. M., Saeed, Z., Akhtar, A., Munawar, H., Yousaf, M. H., Baloach, N. K., & Hussain, F. (2023). A review of swarm robotics in a nutshell. *Drones*, 7(4), 269.
- Tan, Y., & Zheng, Z. Y. (2013). Research advance in swarm robotics. *Defence Technology*, 9(1), 18–39.

- Tang, J., Liu, G., & Pan, Q. (2021). A review on representative swarm intelligence algorithms for solving optimization problems: Applications and trends. *IEEE/CAA Journal of Automatica Sinica*, 8(10), 1627–1643. <https://doi.org/10.1109/JAS.2021.1004129>
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press.
- Wooldridge, M., & Jennings, N. R. (1995). Intelligent agents: Theory and practice. *The Knowledge Engineering Review*, 10(2), 115–152. <https://doi.org/10.1017/S0269888900008122>
- Xi, Z., et al. (2025). AGENTGYM: Evolving Large Language Model-based agents across diverse environments. *arXiv preprint arXiv:2412.14161*. <https://arxiv.org/abs/2412.14161>
- Yang, Q., & Liu, R. (2026). Understanding the application of utility theory in robotics and artificial intelligence: A survey. *arXiv preprint arXiv:2306.09445*. <https://arxiv.org/abs/2306.09445>
- Zhu, K., Du, H., Hong, Z., Yang, X., Guo, S., Wang, Z., Wang, Z., Qian, C., Tang, J., Ji, H., & You, J. (2025). MultiAgentBench: Evaluating the Collaboration and Competition of LLM agents. *arXiv preprint arXiv:2503.01935*. <https://doi.org/10.48550/arxiv.2503.01935>