Communication-as-Control: Intent-Aware Interaction for Scalable Multi-Agent Coordination

Mahdi Imani[1], **Tian Lan**[2]
[1]Northeastern University, Boston, MA 02215
[2]George Washington University, Washington, DC 20052
m.imani@northeastern.edu, tlan@gwu.edu

Abstract

Uncrewed Aerial Systems (UAS) increasingly operate as large, heterogeneous teams under partial observability. In such settings, coordination depends on how agents infer, represent, and act upon uncertainty about one another: belief divergence is a natural consequence of asymmetric sensing and latent intent, and naive communication mechanisms do not scale with team size. We argue for a reframing of interaction as *communication-as-control*, grounded in a small set of tightly coupled capabilities that link sensing to distributed execution under interaction constraints. Rather than prescribing a fixed pipeline, coordination in such teams relies on several interdependent capabilities. First, agents must perform *implicit intent inference without communication*, forming uncertainty-aware beliefs over teammate intent and progress from observable behavior. Second, such uncertainty can inform *selective explicit interaction*, with targeted, minimal exchanges invoked only when unresolved ambiguity becomes consequential for joint decision-making. Third, uncertainty arising from both implicit inference and explicit interaction can be leveraged to support belief-guided coordination and adaptation when alignment is partial, delayed, or unattainable. This makes selective interaction a scalability mechanism: as teams grow in size and heterogeneity, coordination remains feasible only if communication is sparse, recipient-specific, and tied to decision relevance.

Motivation: Coordination Under Latent Intent and Partial Observability

Coordination in autonomous teams hinges not only on sensing and control, but on how agents reason about one another under uncertainty. In multi-agent settings, each agent operates with partial and asymmetric observability, forming internal beliefs about the environment, the task, and the latent intent and progress of its teammates (Zhang, Imani, and Lan 2024; Saboia et al. 2022; Ravari et al. 2025). Because intent and task engagement are not directly observable (Lin et al. 2024), the same behavior may correspond to different goals, contingencies, or constraint-driven deviations. As a result, belief divergence is a natural consequence of heterogeneous sensing, delayed observations, and local context (Alali et al. 2025), particularly in large-scale teams where interactions are local and only a subset of neighbors is observable at any time.

Many existing autonomy frameworks implicitly assume that belief divergence can be managed through shared state, synchronized updates, or predefined communication protocols. While such mechanisms may reduce disagreement in small teams, they do not scale: sharing information, whether raw data, learned messages, or model updates, does not guarantee alignment of decision-relevant beliefs as team size, heterogeneity, and asynchrony increase. As a result, coordination failures emerge from misalignment in internal beliefs at moments when agents' decisions become interdependent.

Current systems lack principled mechanisms to reason about how intent and behavior should be inferred in the absence of communication, how uncertainty in such inferences should be represented, and how that uncertainty should influence interaction and downstream decision-making (Lin et al. 2025). Communication is typically treated either as a background channel for state dissemination or as a heuristic trigger, rather than as a deliberate decision variable that shapes the evolution of shared understanding (Hu et al. 2021; Li and Zhang 2024; Zhu, Dastani, and Wang 2024; Lin et al. 2023; Zhang et al. 2025).

This paper argues for a unifying vision of coordination in distributed autonomy, built around a small set of tightly coupled capabilities. These include the ability to implicitly infer teammate intent and behavior without communication (Lin et al. 2024), to reason explicitly about uncertainty as an indicator of when belief divergence becomes coordination-relevant (Ravari, Ghoreishi, and Imani 2024; Alali and Imani 2024), and to leverage both implicit and explicit information at a higher level to support belief-guided, distributed autonomy (Lin et al. 2025). Together, these capabilities enable scalable coordination under partial observability without assuming persistent communication or shared internal models.

Conceptual Setting and Tactical Constraints

We consider UAS operating in dynamic, adversarial environments with partial observability and heterogeneous teammates. Each platform has distinct sensing geometry, compute and energy limits, and mission-local objectives. Perception is multi-modal and uneven across the team, so uncertainty is not only about what is present, but also about what each teammate could plausibly observe, confirm, and act

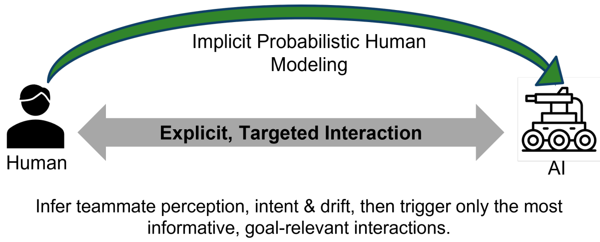

Figure 1: Conceptual framing of communication-as-control. Implicit inference maintains uncertainty over teammate intent and progress; an interaction broker invokes explicit, targeted interaction only when that uncertainty is expected to change a sensing-to-execution commitment (e.g., deconfliction, handoff, hazard response) under interaction cost and timing constraints.

upon within a limited horizon (Cao et al. 2023). Intent and progress are latent and non-stationary: the same observed behavior may correspond to different goals, contingencies, or constraint-driven deviations, and ambiguity grows under occlusion, degraded sensing, and rapidly shifting mission priorities.

Communication is constrained in multiple coupled ways and is often mission-risk bearing (Saboia et al. 2022; Heuer et al. 2024). Bandwidth and latency limits are compounded by intermittent connectivity, asymmetric reachability, contention with higher-priority traffic, and operational exposure risk under policies such as low probability of intercept/detection and burst-only relay windows (Ancel et al. 2024). Even when a link exists, transmission competes with flight-critical computation and timing budgets, and it can be unsafe if it delays actions that are difficult to revise once taken. Distributed autonomy, therefore, requires mechanisms that decide not only *what* to share but *also whether* sharing is worth its cost *now* (Lin et al. 2023; Xue, Wang, and Zhang 2022).

The design objective is not to maximize information transfer. It is maintaining sufficient belief alignment on intent, progress, and risk to enable safe, effective joint action from sensing to execution. We use *belief* to denote a locally maintained, uncertainty-aware coordination state with explicit confidence and staleness. A practical consequence is that perception and fusion are naturally expressed as compact coordination-belief summaries—e.g., decision-relevant ambiguity about hazards, feasibility, or task progress together with uncertainty/staleness—so that coordination decisions remain feasible on-board under edge compute and timing limits (Moss et al. 2024; Sun et al. 2023). Human teammates introduce qualitatively different uncertainty, including delayed responses, cognitive state variability, and authority asymmetries, making belief alignment both more critical and more costly than in purely autonomous teams.

Core Vision: Communication-as-Control via Intent-Aware Interaction

The resulting coordination loop described above couples three interdependent elements: (i) implicit maintenance of uncertainty-aware beliefs about teammate intent and progress from observable behavior, (ii) selective, recipient-specific interaction triggered by *consequential uncertainty*—uncertainty that can change a near-term coordination commitment—and (iii) belief-aware planning that consumes these uncertainty-tagged beliefs to adapt execution when alignment is partial or delayed, without prescribing a specific planning or optimization algorithm.

Implicit intent learning as silent coordination. In the absence of explicit interaction, agents maintain local beliefs over teammate intent and task progress by interpreting observable behavioral cues such as motion patterns, dwell times, formation changes, sensor orientation, route choices, and task artifacts. These cues are inherently ambiguous and context-dependent, and deviations may reflect external contingencies or local constraints rather than changes in goal. The objective of implicit intent learning is therefore calibrated uncertainty over plausible intent and progress hypotheses that supports anticipation, deconfliction, and conservative coordination when communication is unavailable.

Explicit interaction only for consequential uncertainty. Explicit communication is introduced only when unresolved uncertainty is expected to affect a near-term coordination decision. The trigger is not elapsed time since last contact nor a generic entropy threshold; such heuristics fail under asynchrony and heterogeneity, where stale information and differing local contexts decouple "latest" from "decision-relevant." Instead, interaction is driven by *consequential uncertainty*: ambiguity that can flip a commitment, alter a safety margin, invalidate an assumed task allocation, or disrupt a synchronized time window governing interdependent action.

This yields targeted, contextual, and often asymmetric interaction mediated by an *interaction broker* at the sensing-to-execution seam. The broker selects both the recipient and a minimal interaction primitive sufficient to resolve the decision-relevant belief gap, and withholds interaction when uncertainty is irrelevant or counterproductive. In this view, interaction is sparse by design and shaped by its expected impact on coordination rather than by the availability of bandwidth or the volume of state information.

Belief alignment, not telemetry volume, as the control objective. The purpose of interaction is to selectively align the subset of beliefs that govern safe interdependence. Teams do not require a fully synchronized global state; they require alignment on intent, commitments, constraints, and hazards at the points where misalignment would lead to unsafe or inefficient execution. Silence is therefore a first-class action: choosing not to communicate is correct when uncertainty is not coordination-critical or when interaction would not change downstream decisions.

Figure 1 illustrates the proposed communication-as-control loop that couples belief maintenance, interaction decisions, and execution. Table 1 provides illustrative interaction primitives that resolve decision-relevant belief gaps while maximizing coordination value under constrained interaction budgets.

Decision-relevant belief gap	Minimal semantic resolution	Decision consequence if unresolved
Uncertainty over task ownership	"Confirmed responsibility for a shared task over a specified interval."	Conflicting commitments and duplicated effort at execution time.
Uncertainty over task completion state	"Verified completion status within the current decision window."	Unsafe convergence or premature handoff decisions.
Uncertainty over assumed future actions	"Disconfirmed assumed action in the upcoming interval."	Irreversible actions taken under incorrect coordination assumptions.
Uncertainty over localized risk affecting plans	"Risk affecting planned actions elevated (confidence- and staleness-tagged)."	Misaligned risk posture leading to unsafe or overly conservative decisions.

Table 1. Decision-relevant belief gaps and minimal semantic resolutions for scalable coordination. The emphasis is on resolving only those belief misalignments whose resolution would change near-term coordination decisions, not on prescribing message formats or protocols.

Acknowledgements

The authors acknowledge the support of the National Science Foundation award IIS-2311969, Office of Naval Research award N00014-23-1-2850, Army Research Laboratory awards W911NF-23-2-0207 and W911NF-23-20175, and Army Research Office award W911NF-24-2-0166, and the Defense Advanced Research Projects Agency (DARPA) under grant HR00112420366.

References

Alali, M.; and Imani, M. 2024. Deep reinforcement learning data collection for Bayesian inference of hidden Markov models. *IEEE transactions on artificial intelligence*.

Alali, M.; Lan, T.; Ghoreishi, S. F.; and Imani, M. 2025. Decision-Theoretic Learning of Human Perception in Uncertain Environments. In *Proceedings of the 59th Annual Asilomar Conference on Signals, Systems, and Computers*.

Ancel, E.; Moore, A. J.; Young, S. D.; Dill, E. T.; Quach, C.; and Smalling, K. M. 2024. Testing a Run-Time Assurance Framework Coupled with Integrated Risk Mitigation Capabilities for Autonomous Urban UAS Flights. In *2024 AIAA DATC/IEEE 43rd Digital Avionics Systems Conference (DASC)*, 1–9. IEEE.

Cao, C.; Zhu, H.; Ren, Z.; Choset, H.; and Zhang, J. 2023. Representation granularity enables time-efficient autonomous exploration in large, complex worlds. *Science Robotics*, 8(80): eadf0970.

Heuer, L.; Palmieri, L.; Mannucci, A.; Koenig, S.; and Magnusson, M. 2024. Benchmarking Multi-robot coordination in realistic, unstructured human-shared environments. In *2024 IEEE International Conference on Robotics and Automation (ICRA)*, 14541–14547. IEEE.

Hu, G.; Zhu, Y.; Zhao, D.; Zhao, M.; and Hao, J. 2021. Event-triggered communication network with limited-bandwidth constraint for multi-agent reinforcement learning. *IEEE Transactions on Neural Networks and Learning Systems*, 34(8): 3966–3978.

Li, X.; and Zhang, J. 2024. Context-aware Communication for Multi-agent Reinforcement Learning. In *Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems*, 1156–1164.

Lin, Y.; Ghoreishi, S. F.; Lan, T.; and Imani, M. 2024. High-level human intention learning for cooperative decision-making. In *2024 IEEE Conference on Control Technology and Applications (CCTA)*, 209–216. IEEE.

Lin, Y.; Ghoreishi, S. F.; Lan, T.; and Imani, M. 2025. Reinforcement learning for human-AI collaboration via probabilistic intent inference. In *Reinforcement Learning Conference*.

Lin, Y.; Li, W.; Zha, H.; and Wang, B. 2023. Information design in multi-agent reinforcement learning. *Advances in Neural Information Processing Systems*, 36: 25584–25597.

Moss, R. J.; Jamgochian, A.; Fischer, J.; Corso, A.; and Kochenderfer, M. J. 2024. ConstrainedZero: chance-constrained POMDP planning using learned probabilistic failure surrogates and adaptive safety constraints. In *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence*, 6752–6760.

Ravari, A.; Ghoreishi, S. F.; and Imani, M. 2024. Optimal inference of hidden Markov models through expert-acquired data. *IEEE Transactions on Artificial Intelligence*, 5(8): 3985–4000.

Ravari, A.; Ghoreishi, S. F.; Lan, T.; Bastian, N. D.; and Imani, M. 2025. Hybrid Modeling of Heterogeneous Human Teams for Collaborative Decision Processes. *Proceedings of Machine Learning Research vol*, 283: 1–14.

Saboia, M.; Clark, L.; Thangavelu, V.; Edlund, J. A.; Otsu, K.; Correa, G. J.; Varadharajan, V. S.; Santamaria-Navarro, A.; Touma, T.; Bouman, A.; et al. 2022. Achord: Communication-aware multi-robot coordination with intermittent connectivity. *IEEE Robotics and Automation Letters*, 7(4): 10184–10191.

Sun, J.; Jiang, Y.; Qiu, J.; Nobel, P.; Kochenderfer, M. J.; and Schwager, M. 2023. Conformal prediction for uncertainty-aware planning with diffusion dynamics model. *Advances in Neural Information Processing Systems*, 36: 80324–80337.

Xue, Y.; Wang, J.; and Zhang, C. 2022. Efficient Multi-Agent Communication via Shapley Message Value. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*.

Zhang, Z.; Imani, M.; and Lan, T. 2024. Modeling other players with Bayesian beliefs for games with incomplete information. *arXiv preprint arXiv:2405.14122*.

Zhang, Z.; Zhou, H.; Imani, M.; Lee, T.; and Lan, T. 2025. Learning to Collaborate with Unknown Agents in the Absence of Reward. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 39, 14502–14511.

Zhu, C.; Dastani, M.; and Wang, S. 2024. A survey of multi-agent deep reinforcement learning with communication. *Autonomous Agents and Multi-Agent Systems*, 38(1): 4.