

DOCTORAL CONSORTIUM PAPER SUBMISSION ■ AAMAS 2026

Adaptive Human-AI Shared Autonomy Frameworks for Dynamic Multi-Agent Decision Environments

Elena Rostova ■ Third-Year PhD Candidate

Primary Advisors: Prof. Marcus Vance & Dr. Sophia Al-Mansoor

[Institute of Advanced Technology](#) elena.rostova@iat-edu.org elena-rostova.research.org

Executive Summary & Research Abstract

High-stakes multi-agent environments—such as autonomous fleet management, robotic disaster response, and dynamic network defense—increasingly rely on hybrid human-AI teams. However, conventional shared autonomy frameworks struggle with two critical flaws: static allocations of authority that fail under non-stationary task demands, and cognitive overload induced when human operators must intervene during unexpected agent policy drifts. This dissertation proposal presents **SYNAPSE-LOOP**, a novel adaptive shared control paradigm that dynamically models human cognitive state, predicts intervention intent, and negotiates task authority in real time. By fusing bi-directional intent forecasting with multi-agent reinforcement learning (MARL), our framework maintains mutual alignment while minimizing unnecessary operator burden. Early empirical evaluations on simulated emergency logistics demonstrate a 34% reduction in human cognitive workload metrics alongside a 28% improvement in mission completion velocity compared to standard supervisory baselines.

Candidacy Snapshot

Institution: Inst. of Advanced Tech.**Department:** Computer Science & AI**PhD Start Date:** September 2023**Expected Graduation:** May 2027**Candidacy Status:** Proposal Approved**Primary Area:** Human-AI Collaboration, MARL, Shared Autonomy**Research Keywords:**

- Shared Autonomy
- Human-in-the-Loop
- MARL Policy Shift
- Cognitive Workload

1 Problem Statement & Research Motivation

Autonomous multi-agent systems are rapidly transitioning from controlled laboratory settings into dynamic, high-uncertainty operational environments. In scenarios such as robotic search-and-rescue, urban aerial mobility management, and dynamic supply chain dispatch, full agent autonomy remains unfeasible due to unpredictable environmental edge cases and complex ethical trade-offs. Consequently, system architects implement *Human-in-the-Loop* (HITL) architectures to retain human judgment for high-level strategic supervision.

Despite significant advances in single-agent shared control, existing multi-agent supervisory frameworks suffer from three systemic limitations:

- P1. Rigid Authority Boundaries:** Existing frameworks rely on pre-scripted state transitions (e.g., manual override toggles), which lack the granularity required to transfer partial authority across heterogeneous sub-tasks under fluctuating network latency.
- P2. Cognitive Context Disconnect:** When autonomous agents encounter out-of-distribution states and request human intervention, operators suffer severe situational awareness degradation (*out-of-the-loop latency*), frequently leading to suboptimal or contradictory control inputs.

P3. Asymmetric Mutual Modeling: While human operators develop internal mental models of agent behavior over time, autonomous agents typically treat human interventions as exogenous noise rather than intentional feedback derived from unobserved environmental context.

Addressing these challenges requires a paradigm shift from *passive supervisory control* toward *symmetrical dynamic co-adaptation*, where both human operators and autonomous agents continuously model each other's intent, confidence, and cognitive capacity.

2 ? Key Research Questions & Hypotheses

This dissertation investigates the formalization, algorithmic implementation, and human-subject validation of adaptive multi-agent shared control. The core research is guided by three primary research questions:

RQ1 (Bi-Directional Intent Estimation): How can an autonomous agent team infer human intervention intent and cognitive fatigue from implicit behavioural streams (e.g., eye gaze, control input jitter) before performance degradation occurs?

Hypothesis 1 (H1): A latent temporal belief model trained on multimodal interaction telemetry can predict intervention intent 2.5 seconds prior to action execution with $> 88\%$ precision.

RQ2 (Dynamic Authority Allocation): What game-theoretic framework optimally balances agent policy execution against human control authority while strictly bounding operator cognitive workload under threshold $\Omega_{\max}$?

Hypothesis 2 (H2): Bounding authority distribution via constrained Markov Decision Processes (CMDP) with dynamic slack variables achieves higher mission safety guarantees than static sliding-autonomy approaches.

RQ3 (Reciprocal Policy Alignment): How can human intervention trajectories be converted into real-time reward shaping signals without inducing reward hacking or catastrophic forgetting in underlying MARL policies?

Hypothesis 3 (H3): Online constrained preference optimization with bounded KL divergence prevents policy degradation while accelerating policy adaptation by at least 40%.

3 ⚙️ Proposed SYNAPSE-LOOP Architecture

To evaluate these hypotheses, we propose the **SYNAPSE-LOOP** (Symmetrical Network for Adaptive Policy Synchronization & Estimation) architecture. Figure 1 illustrates the closed-loop interactions between human operators, the cognitive monitoring suite, the authority arbitrator, and the multi-agent execution team.

The architecture comprises three core modules:

- **Cognitive State & Intent Estimator:** Processes continuous biometric data (pupillometry, heart rate variability) combined with control gesture dynamics to calculate operator cognitive load index $\Omega_t \in [0, 1]$ and forecast intervention intent probability $\hat{I}_t$.
- **Dynamic Authority Arbitrator:** Solves a constrained game-theoretic optimization problem at each decision epoch t , outputting an authority allocation vector $\alpha_t \in [0, 1]^N$ across N sub-task dimensions.




Figure 1: The SYNAPSE-LOOP framework connecting human cognitive state estimation with multi-agent policy optimization.

- **Multi-Agent Adaptation Engine:** Utilizes parameter-shared Soft Actor-Critic (SAC) augmented with an online human feedback loss function to adjust agent policy execution without causing instability.

4 Methodology & Preliminary Results

We evaluated a prototype version of SYNAPSE-LOOP in a simulated disaster logistics testbed involving 12 autonomous delivery drones (Vertex-UAV Platform) and a single human supervisor managing emergency supply dispatch under severe weather perturbations.

4.1 Experimental Setup & Baselines

We compared our approach against three standard operational paradigms:

- Full Autonomy (FA):** Agents operate independently without human intervention options.
- Fixed Supervisory Control (FSC):** Human operator must manually request control of specific drones.
- Threshold Sliding Autonomy (TSA):** Authority switches automatically based on fixed rule-based performance metrics.

4.2 Empirical Performance Comparison

Table 1 summarizes the performance metrics collected across 45 evaluation trials with 15 human subjects.

Table 1: Comparative evaluation across control paradigms (Mean $\pm$ Std. Dev.).

Control Paradigm	Task Completion (%)	Mean Workload (Ω_t)	Out-of-Loop Delay (s)	Safety Violations
Full Autonomy (FA)	62.4 $\pm$ 5.1	N/A	N/A	14.2 $\pm$ 2.3
Fixed Supervisory (FSC)	78.1 $\pm$ 3.8	0.78 $\pm$ 0.06	4.1 $\pm$ 0.8	6.5 $\pm$ 1.2
Threshold Sliding (TSA)	81.5 $\pm$ 3.2	0.64 $\pm$ 0.05	2.8 $\pm$ 0.5	4.1 $\pm$ 0.9
SYNAPSE-LOOP (Ours)	94.7 $\pm$ 1.9	0.41 $\pm$ 0.04	0.9 $\pm$ 0.2	0.8 $\pm$ 0.3

The empirical findings confirm H1 and H2: SYNAPSE-LOOP drastically reduces out-of-the-loop delay from 4.1 seconds down to 0.9 seconds while maintaining low operator workload scores ($\Omega_t = 0.41$).

5 📅 Dissertation Status & Timeline

As a third-year PhD candidate, I have completed all required coursework, passed the departmental comprehensive examinations, and formally defended the dissertation proposal. Table 2 outlines the completed milestones and target schedule through graduation.

Table 2: Dissertation research roadmap and milestones.

Phase / Date	Milestone / Objective	Status / Target
Phase I (2023–2024)	Literature survey, baseline testbed setup, initial cognitive monitoring pipeline.	✅ Completed
Phase II (2024–2025)	Formulated CMDP authority arbitration algorithm; published initial results at AAMAS.	✅ Completed
Phase III (2025–2026)	Human-subject trials with multi-agent physical UAV testbed; refine reciprocal learning.	🔄 In Progress
Phase IV (Late 2026)	Complete full multi-modal ablation studies; finalize dissertation draft.	🕒 Q4 2026
Phase V (Spring 2027)	Final dissertation defense and manuscript publication.	🎓 May 2027

6 🎓 Doctoral Consortium Goals & Expected Feedback

Participating in the AAMAS 2026 Doctoral Consortium will provide pivotal guidance at a critical stage of my PhD research. Specifically, I seek feedback from senior consortium mentors on the following aspects:

- **Methodological Validation for Human Trials:** Guidance on scaling human-subject user studies to ensure statistical power when evaluating complex multi-modal biometrics across diverse operator experience levels.
- **Theoretical Bounding of Safety in MARL:** Insights into formal verification techniques for bounding safety constraints in online preference optimization when human rewards fluctuate due to subjective operator biases.
- **Career Mentorship & Post-PhD Trajectory:** Advice on structuring a research agenda suited for academic faculty positions vs. industrial research laboratories focused on human-centric AI systems.

References

- [1] M. Vance, E. Rostova, and S. Al-Mansoor. Dynamic authority allocation in multi-agent shared control. *IEEE Transactions on Human-Machine Systems*, 54(3):310–322, 2024.
- [2] E. Rostova and M. Vance. Predicting human intervention intent in multi-uav supervisory tasks. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pages 412–420, 2025.
- [3] S. Al-Mansoor, E. Rostova, et al. Bio-behavioral metrics for cognitive load estimation in complex autonomous systems. *ACM Transactions on Interactive Intelligent Systems*, 15(2):145–162, 2025.