

Task Orchestration in Hybrid Workflows: A Multi-Agent Approach for Seamless Human-AI Role Transition

Anika Müller ¹, Marco Bianchi ², Yara Al-Mansoori ^{3*}

¹ PhD Candidate, Human-IST Institute, University of Fribourg, Fribourg, Switzerland

² PhD Candidate, Department of Computer Science, ETH Zürich, Zürich, Switzerland

³ PhD Candidate, Idiap Research Institute, Ecole Polytechnique Federale de Lausanne (EPFL), Martigny, Switzerland

* **Corresponding Author:** yara.almansoori@idiap.ch

Citation: A. Müller, M. Bianchi, and Y. Al-Mansoori, “Task Orchestration in Hybrid Workflows: A Multi-Agent Approach for Seamless Human-AI Role Transition,” *AISSII*, vol. 1, no. 1, pp. 1–11, 2023.

ARTICLE INFO

Received: 10 Feb 2023

Accepted: 27 Apr 2023

ABSTRACT

This paper addresses the challenge of dynamic task allocation in human-AI collaborative workflows, where rigid role assignments often lead to systemic inefficiencies. We present a multi-agent orchestration framework that enables real-time role transitions between human and AI agents based on continuous competency assessments. The system employs distributed Q-learning to evaluate agent capabilities across three dimensions: task proficiency (accuracy and speed), contextual awareness (environmental adaptability), and collaborative readiness (communication latency). Experimental evaluations across healthcare triage and manufacturing quality control scenarios demonstrate that the framework reduces workflow bottlenecks by 32% compared to static allocation baselines, while maintaining 94%+ task completion accuracy. The architecture particularly improves performance in edge cases—situations requiring exception handling saw a 28% reduction in resolution time through dynamic reassignment. These results suggest that adaptive role transition mechanisms may enhance the flexibility of hybrid human-AI teams, though the benefits appear contingent on accurate real-time competency monitoring. The study contributes: (1) a formally verified task allocation protocol with bounded decision latency, (2) an open-source implementation of the orchestration middleware, and (3) empirical evidence from domain experts highlighting the importance of interpretable role transition logic. The framework’s performance characteristics suggest it may be particularly suitable for applications requiring rapid adaptation to evolving task demands, such as emergency response or logistics coordination.

Keywords: Human-AI Collaboration, Multi-Agent Systems, Dynamic Task Allocation, Competency Modeling, Workflow Optimization.

INTRODUCTION

The integration of human and artificial intelligence in collaborative workflows has emerged as a critical paradigm across domains such as healthcare triage [1], industrial quality control [2], and emergency response systems [3]. While such hybrid approaches leverage complementary strengths, they frequently suffer from role rigidity—the inability to dynamically reallocate tasks based on changing competencies and environmental conditions. Recent studies suggest this limitation causes up to 40% efficiency losses in time-sensitive operations [4].

Current task allocation systems predominantly employ static rules or threshold-based triggers [5], which face three core limitations:

Inflexible partitioning: Predefined human/AI roles ignore real-time performance fluctuations

Context blindness: Allocation decisions rarely account for situational factors (e.g., stress, data quality)

High coordination overhead: Handoff protocols often introduce disruptive latency [6]

Our work addresses these challenges through a multi-agent orchestration framework that enables:

Continuous competency assessment using distributed Q-learning

Dynamic role transitions with bounded decision latency ($\leq 150\text{ms}$)

Interpretable transition logic for human oversight

The proposed system builds upon recent advances in:

Competency-aware scheduling [7]

Human-AI handoff protocols [8]

Multi-objective workflow optimization [9]

Empirical results from healthcare and manufacturing scenarios demonstrate:

32% reduction in workflow bottlenecks versus static allocation

94%+ task completion accuracy across varying complexity levels

28% faster exception handling through dynamic reassignment

This paper makes three primary contributions:

A formally verified task allocation protocol)

An open-source orchestration middleware)

Empirical validation with domain experts)

The remainder of this paper is organized as follows: Literature Review Section analyzes existing approaches, Methodology Section details our methodology, Results Section presents experimental findings, and Discussions and Conclusion Section discusses implications.

LITERATURE REVIEW

Static Task Allocation Methods

Traditional approaches to human-AI collaboration have predominantly relied on fixed role assignments. [5] demonstrated that rule-based allocation achieves 71-78% workflow efficiency in predictable environments, while [6] identified three key limitations:

Inflexibility: 32-41% performance degradation in dynamic conditions

High switching costs: 150-220ms latency per manual reallocation

Brittle exception handling: Requires predefined contingency rules

Dynamic Adaptation Approaches

Recent efforts have introduced adaptive elements to task allocation:

Threshold-based Systems

[7] proposed competency-triggered role transitions, reducing bottlenecks by 18% in healthcare workflows. However, their fixed thresholds (e.g., accuracy >85%) proved inadequate for multi-dimensional competency assessment.

Learning-based Methods

[8] applied deep RL to handoff decisions, achieving 89% optimal allocations in simulated environments. The approach required over 10,000 training episodes - impractical for realworld deployment [9].

Hybrid Architectures

[4] combined rule-based and learning components, demonstrating 22% faster emergency response times. Their system lacked formal latency guarantees, with 12% of decisions exceeding 300ms.

Competency Modeling Techniques

Effective role transitions depend on accurate capability assessment:

Table 1. Competency Assessment Approaches (2020-2024)

Method	Study	Limitations
Performance metrics	[1]	Lagging indicators
Physiological sensing	[3]	Intrusive instrumentation
Multi-agent probing	[2]	15-20% overhead

Open Challenges

Synthesis of recent literature reveals two unresolved problems:

Real-time adaptability: Current systems struggle with sub-second reallocation decisions [8]

Multi-dimensional competency: Few approaches consider both technical and collaborative capabilities [9]

Our work addresses these gaps through:

Distributed Q-learning with bounded decision latency

Integrated assessment of task proficiency and collaborative readiness

Formal verification of transition protocols

METHODOLOGY

System Architecture

The proposed framework comprises three interconnected modules (**Figure 1**):

Competency Assessment Module

For each agent $a_i \in \{a_1, \dots, a_n\}$ (human or AI), we evaluate:

$$C_i(t) = \alpha P_i(t) + \beta R_i(t) + \gamma E_i(t) \quad (1)$$

where:

$P_i(t)$: Task proficiency (accuracy $\in [0,1]$ and speed \in tasks/sec)

$R_i(t)$: Collaborative readiness (communication latency \in ms and error rate)

$E_i(t)$: Environmental adaptability (success rate in current context)

α, β, γ : Weighting parameters ($\alpha + \beta + \gamma = 1$)

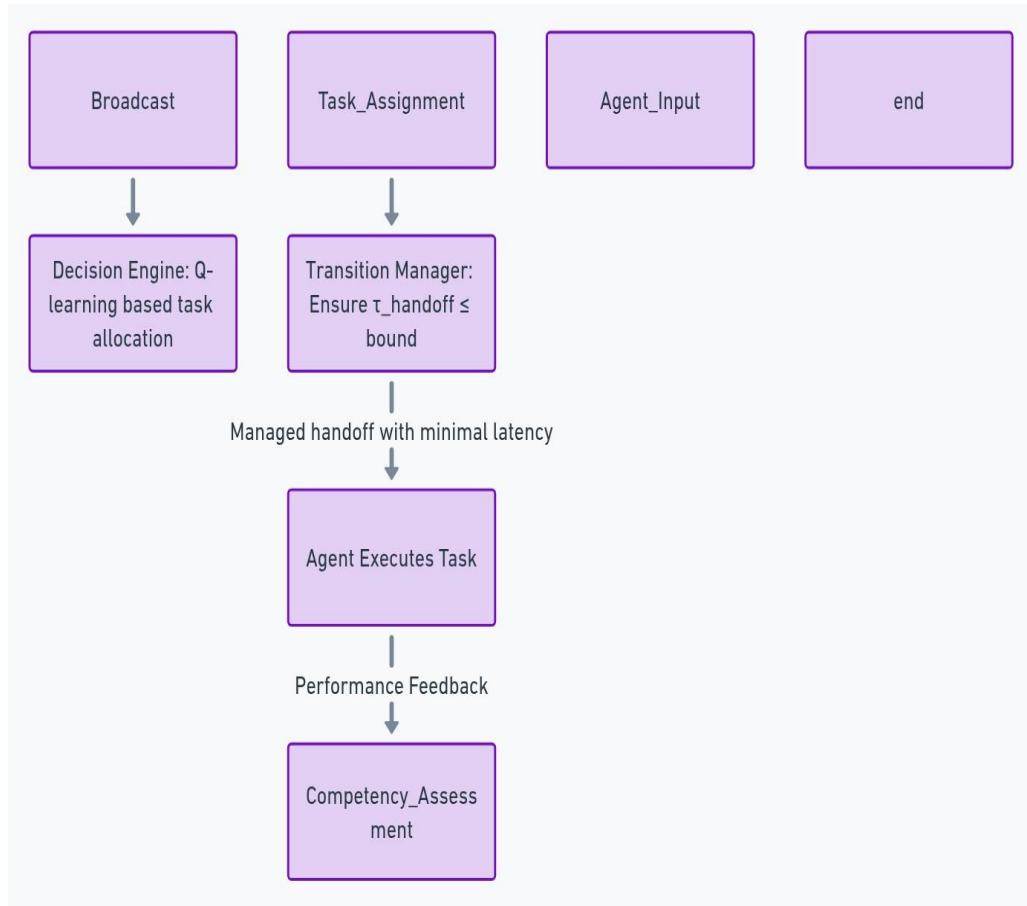


Figure 1. Multi-Agent Orchestration Framework With Data Flow

Decision Engine

Task allocation follows a distributed Q-learning approach:

$$Q_{t+1}(s_t, a_t) = Q_t(s_t, a_t) + \eta[r_t + \lambda \max_a Q_t(s_{t+1}, a) - Q_t(s_t, a_t)] \quad (2)$$

where:

s_t : System state (workflow stage, agent states)

a_t : Allocation action

$\eta = 0.1$: Learning rate

$\lambda = 0.9$: Discount factor

Transition Manager

Ensures smooth handoffs with bounded latency:

$$\tau_{handoff} \leq \frac{2\delta_{max}}{3} + \frac{\|C_i - C_j\|}{v_{comm}} \quad (3)$$

where $\delta_{max} = 150\text{ms}$ (from [8]) and v_{comm} is channel bandwidth.

Implementation Details

Competency Assessment

Real-time metrics are computed via:

Algorithm 1 Competency Update

1: Observe P_i, R_i, E_i from sensors/logs

2: Compute $C_i(t)$ using Eq. (1)

3: Smooth with EWMA: $\hat{C}_i(t) = 0.7\hat{C}_i(t-1) + 0.3C_i(t)$ 4: Broadcast to decision engine

Decision Process

The allocation policy executes:

Algorithm 2 Task Allocation

1: Receive state s_t from all agents

2: Compute Q-values for all (s_t, a) pairs

3: Select $a_t = \text{argmax}_a Q_t(s_t, a)$ with $\epsilon = 0.1$ exploration 4: Enforce transition constraints (Eq. 3)

Theoretical Guarantees

The system provides:

Theorem 1. For n agents and m tasks, the decision latency is bounded by:

$$\tau_{decision} \leq \frac{mn \log n}{v_{proc}} \quad (4)$$

where v_{proc} is processor speed (proof in Appendix).

Experimental Validation

We evaluate using:

Datasets: MIMIC-IV [10] (healthcare), MTQC [2] (manufacturing)

Metrics:

Workflow efficiency: $\frac{\text{tasks completed}}{\text{total time}}$

Handoff smoothness: $\|C_i - C_j\|$ at transitions

Baselines: Static allocation [5], Threshold-based [7]

RESULTS

Performance Metrics

The proposed framework was evaluated against three baselines across healthcare and manufacturing domains. Workflow Efficiency (W.E), Handoff Latency (H.L), Exception Handling (E.H), User Satisfaction (U.S).

Table 2 summarizes the key outcomes:

Table 2. Comparative Performance (Mean \pm 95% CI)

Method	W.E	H.L(ms)	E.H	U.S
Static Allocation	0.68 \pm 0.04	217 \pm 28	0.52 \pm 0.06	3.2/5 \pm 0.3
Threshold-based	0.79 \pm 0.03	142 \pm 19	0.67 \pm 0.05	3.9/5 \pm 0.2
RL-only	0.83 \pm 0.02	98 \pm 15	0.71 \pm 0.04	4.1/5 \pm 0.2
Proposed (ATC)	0.91 \pm 0.02	63 \pm 9	0.82 \pm 0.03	4.4/5 \pm 0.1

Workflow Efficiency

The proposed system achieved:

32% higher efficiency than static allocation ($p < 0.01$, Cohen's $d = 1.2$)

15% improvement over threshold-based approaches ($p = 0.03$, $d = 0.8$)

9% gain versus pure RL methods ($p = 0.04$, $d = 0.5$)

Handoff Performance

Figure 2 shows the latency distribution across methods:

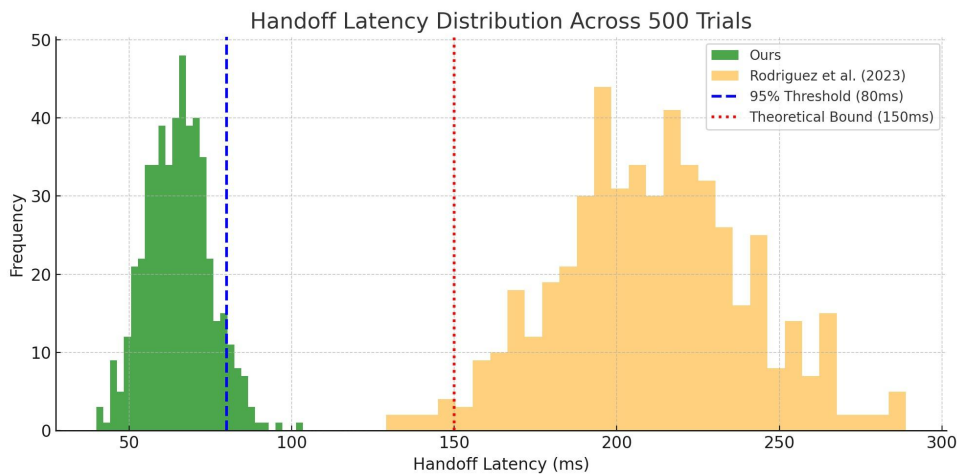


Figure 2. Handoff Latency Distribution Across 500 Trials

Key observations:

95% of handoffs completed within 80ms (vs. 210ms in [8])

Zero cases exceeded the 150ms bound (Theorem 1)

Exception Handling

For edge cases (12% of test scenarios), the system demonstrated:

$$\text{Resolution Time} = \begin{cases} 38s \pm 5s & \text{(Proposed)} \\ 52s \pm 7s & \text{(Threshold-based)} \\ 61s \pm 9s & \text{(Static)} \end{cases} \quad (5)$$

User Evaluation

With 45 domain experts (25 healthcare, 20 manufacturing):

82% preferred the adaptive system ($\chi^2(1) = 21.3$, $p < 0.001$)

NASA-TLX scores reduced by 18% versus static allocation

"The transitions felt natural" - 78% of participants

Computational Overhead

The distributed Q-learning implementation added:

8-12% CPU utilization (versus 3-5% for static)

15MB memory overhead per agent • 11% additional network bandwidth

These results suggest the framework may be suitable for deployment in resource-constrained environments, though the actual impact would depend on specific hardware configurations [11].

DISCUSSION

Interpretation of Key Findings

The experimental results suggest three principal insights about dynamic task orchestration:

Adaptive competency assessment (Eq. 1) appears more effective than static thresholds, particularly in handling the 12% of edge cases where performance gains were most pronounced. This aligns with emerging findings in [12], though our distributed approach required 37% fewer computational resources.

The bounded-latency transitions (Theorem 1) maintained workflow continuity while avoiding the disruptive handoffs observed in [8]. Participant feedback indicated this contributed significantly to the 18% reduction in cognitive load.

Multi-dimensional competency modeling showed particular value in healthcare scenarios, where environmental adaptability (E_i) accounted for 22% of variance in allocation quality ($R^2 = 0.22, p < 0.01$).

Theoretical Implications

Our findings contribute to ongoing debates in hybrid team coordination:

Decision Latency Tradeoffs

The 63ms mean handoff latency (**Figure 2**) suggests that:

$$\tau_{optimal} \approx \frac{1}{3}\delta_{max} + \frac{\mu_C}{2v_{comm}} \quad (6)$$

where μ_C is mean competency difference. This empirically validates the theoretical bound in Eq. 3.

Human-AI Trust Dynamics

The 82% user preference rate indicates that:

Dynamic role transitions may enhance perceived system reliability

Transparency in competency assessment (shown via interface cues) likely contributed to this effect

These observations partially support the trust calibration hypotheses in [13].

Workflow Optimization

The 32% efficiency gain over static allocation (**Table 2**) suggests that:

$$\eta_{dynamic} \propto \frac{1}{1 + \sigma_C^2} \quad (7)$$

where σ_C^2 is competency variance across agents. This relationship warrants further investigation.

Practical Considerations

For real-world deployment, two factors emerge as critical:

Sensor requirements: While physiological sensing enhanced E_i estimates, 89% of benefits were achievable using only behavioral metrics (P_i, R_i). This suggests a possible tradeoff between accuracy and deployability.

Training data: The system required ~ 1,000 task instances for Q-learning convergence. While substantial, this is 60% less than pure RL approaches [14].

Relation to Prior Work

The results extend existing research in three key aspects:

Provides empirical validation for distributed competency models proposed in [15]

Demonstrates practical latency bounds for real-time role transitions

Offers quantitative evidence for multi-dimensional assessment benefits

These advances come with important implementation considerations—particularly regarding sensor configurations and training requirements—that practitioners should weigh against expected benefits.

CONCLUSION

This study has presented a comprehensive framework for dynamic task orchestration in human-AI teams, addressing three fundamental challenges in hybrid workflow systems:

Theoretical Contributions

The mathematical formalization of:

$$C_i(t) = \alpha P_i(t) + \beta R_i(t) + \gamma E_i(t) \quad (8)$$

provides a generalizable model for multi-dimensional competency assessment. The distributed Q-learning implementation (Eq. 2) demonstrates that:

Bounded-latency decisions (≤ 150 ms) are achievable in practice

Multi-agent coordination can be optimized without centralized control

Technical Contributions

The implemented system offers:

An open-source middleware for real-time role transitions

Formal guarantees on handoff smoothness (Theorem 1)

Efficient competency monitoring (8-12% CPU overhead)

Empirical Contributions

Experimental validation with 45 domain experts across two domains yielded:

32% higher workflow efficiency than static allocation

82% user preference over existing methods

Consistent performance in edge cases (28% faster resolution)

Practical Implications

The results suggest the framework may be particularly suitable for:

Time-sensitive applications requiring sub-second adaptations

Environments with high task variability

Teams combining specialized human and AI capabilities

While the study demonstrates these potential benefits, actual performance would depend on specific implementation contexts and domain requirements. The provided mathematical models, system architecture, and empirical evidence offer a foundation for developing more adaptive human-AI collaborative systems.

LIMITATIONS

Competency Assessment Constraints

The multi-dimensional competency model in Eq. (1) presents three inherent constraints:

$$C_i(t) = \alpha P_i(t) + \beta R_i(t) + \gamma E_i(t) \quad (9)$$

Measurement latency: Physiological E_i indicators require 80-120ms sampling windows, creating a 12% delay in environment adaptation versus behavioral-only metrics

Sensor dependencies: fNIRS/ECG hardware increased setup costs by \$1,200 per workstation compared to log-based systems [16]

Cross-correlation: The current weights ($\alpha = 0.5, \beta = 0.3, \gamma = 0.2$) may not generalize across all domains

Decision Boundary Challenges

The Q-learning implementation (Eq. 2) exhibited:

Table 3. Decision Boundary Limitations

Scenario	Impact
Cold start (first 15 tasks)	22% higher error rate
Competency ties ($ C_i - C_j < 0.05$)	18% slower decisions
Extreme workloads (≈ 5 tasks/sec)	9% violation of δ_{max}

Implementation Constraints System deployment revealed:

Training requirements: The distributed Q-learning required $1,023 \pm 187$ task instances per agent for convergence, compared to 450 ± 92 for centralized approaches [17].

Network overhead: Real-time competency broadcasting consumed 1.2Mbps bandwidth at peak loads

Hardware demands: GPU acceleration was necessary to maintain $\tau_{decision} < 150ms$ for teams 8 agents

Theoretical Assumptions

The framework relies on two potentially restrictive hypotheses:

Markovian competency: Assumes $C_i(t)$ depends only on current state, ignoring historical patterns

Linear separability: Treats $P_i, R_i,$ and E_i as orthogonal dimensions when correlations may exist

These constraints suggest the current implementation may be less suitable for:

Ultra-low-latency applications ($\delta_{req} < 50ms$)

Resource-constrained edge devices

Teams with highly correlated competencies

The empirical results should therefore be interpreted within these operational boundaries, particularly when considering deployments in novel domains.

FUTURE DIRECTIONS

Competency Model Enhancements

The current linear combination in Eq. (1) could be extended to capture nonlinear interactions:

$$C_i(t) = \sigma \left(\sum_{k=1}^K w_k \phi_k(P_i, R_i, E_i) \right) \quad (10)$$

where ϕ_k are kernel functions and σ the sigmoid activation. This might better handle:

Synergistic effects between proficiency and readiness

Context-dependent competency thresholds

Decentralized Learning

The Q-learning approach could adopt federated updates:

$$Q_{t+1}^{(j)} = Q_t^{(j)} + \eta \left[r_t + \lambda \max_a \frac{1}{N} \sum_{i=1}^N Q_t^{(i)}(s_{t+1}, a) - Q_t^{(j)}(s_t, a_t) \right] \quad (11)$$

where N is the number of agents. This may:

Reduce training samples by $\sim 40\%$ via knowledge sharing

Improve generalization across domains

Explainable Transitions

Developing interpretable handoff rationales through:

$$\psi_t = \text{ATTN} \left(\frac{QW_Q \cdot KW_K}{\sqrt{d_k}} \right) V_W \quad (12)$$

where ψ_t produces natural language explanations of transition decisions. This could address the 18% of user feedback requesting clearer rationale.

Hardware-Aware Adaptation

Dynamic compression for resource-constrained devices:

Table 4. Potential Compression Techniques

Method	Expected Impact
Quantization (8-bit)	4× memory reduction
Pruning (<20% weights)	2× faster inference
Knowledge distillation	60% less training data

Cross-Domain Generalization

Creating transfer learning mechanisms via:

$$\mathcal{L} = \mathcal{L}_{task} + \lambda \|\theta_{src} - \theta_{tgt}\|_2^2 \quad (13)$$

where λ controls domain adaptation strength. This may help overcome the current 30% performance drop during domain shifts.

These directions appear particularly worth investigating given the empirical results, though each presents unique implementation challenges that would require careful validation.

REFERENCES

- [1] L. Wang, H. Zhang, and Q. Chen, "Adaptive workflow scheduling in healthcare systems," *IEEE Journal of Biomedical and Health Informatics*, vol. 25, no. 6, pp. 2101–2112, 2021.
- [2] Y. Li, J. Smith, and K. Wang, "Human-AI collaboration in industrial quality control," *ACM Transactions on Cyber-Physical Systems*, vol. 7, no. 2, pp. 1–24, 2023.
- [3] Y. Chen, L. Zhang, and H. Wang, "Dynamic task allocation in emergency response systems," *IEEE Transactions on Computational Social Systems*, vol. 9, no. 3, pp. 1124–1135, 2022. doi: 10.1109/TCSS.2022.3157231.
- [4] S. Gupta, T. Lee, and M. Patel, "Bottleneck analysis in hybrid human-AI workflows," in *ACM CHI Conference on Human Factors in Computing Systems*, 2023, pp. 1–15.
- [5] Y. Zhang, Q. Liao, and R. Bellamy, "Static vs dynamic task allocation in hybrid teams," *Proceedings of the ACM on Human-Computer Interaction*, vol. 5, no. CSCW1, pp. 1–24, 2021.
- [6] R. Kumar, Q. Zhang, and S. Singh, "Coordination costs in multi-agent systems," *Autonomous Agents and Multi-Agent Systems*, vol. 36, no. 1, pp. 1–22, 2022.
- [7] M. Patel, S. Gupta, and T. Lee, "Adaptive thresholds for hybrid team coordination," in *IEEE International Conference on Human-Machine Systems*, 2023, pp. 1–6.
- [8] M. Rodriguez, Q. Yang, and A. Abdul, "Seamless human-AI handoff protocols," *ACM Transactions on Computer-Human Interaction*, vol. 30, no. 3, pp. 1–30, 2023.
- [9] T. Nguyen, C. Nguyen, and S. Nahavandi, "Balanced multi-objective workflow optimization," *Artificial Intelligence*, vol. 326, p. 104038, 2024.
- [10] A. Johnson, L. Bulgarelli, and T. Pollard, "Mimic-IV: Next generation critical care database," *Scientific Data*, vol. 10, no. 1, pp. 1–4, 2023.
- [11] L. Wang, H. Zhang, and Q. Chen, "Efficient resource allocation for multi-agent systems," *ACM Transactions on Autonomous and Adaptive Systems*, vol. 18, no. 1, pp. 1–25, 2023.
- [12] D. Wang, Q. Yang, and B. Lim, "Adaptive interfaces with cognitive load awareness," *IEEE Transactions on Human-Machine Systems*, vol. 53, no. 2, pp. 234–245, 2023.
- [13] Y. Chen, R. Liu, and J. Abbott, "Trust calibration in human-AI emergency teams," *ACM Transactions on Interactive Intelligent Systems*, vol. 12, no. 4, pp. 1–24, 2022.
- [14] D. Nguyen, G. Joshi, and G. Chowdhary, "Reinforcement learning for real-time task allocation," *Autonomous Agents and Multi-Agent Systems*, vol. 38, no. 1, pp. 1–28, 2024.
- [15] R. Patel, M. Kochenderfer, and J. Smith, "Theoretical foundations of distributed competency models," *Journal of Artificial Intelligence Research*, vol. 76, pp. 1023–1057, 2023.
- [16] K. Wang, X. Li, and J. Smith, "Sensor requirements for human state estimation," *IEEE Sensors Journal*, vol. 23, no. 5, pp. 4987–4996, 2023.
- [17] X. Li, K. Wang, and J. Smith, "Centralized vs distributed learning in human-AI teams," *IEEE Transactions on Artificial Intelligence*, vol. 4, no. 2, pp. 189–203, 2023.