

Article

Not peer-reviewed version

---

# Contextual Trust Evaluation for Robust Coordination in Large Language Model Multi-Agent Systems

---

Kangning Gao , Haotian Zhu , Rui Liu , Jinming Li , [Xu Yan](#) , Yi Hu \*

Posted Date: 31 December 2025

doi: 10.20944/preprints202512.2748.v1

Keywords: large language models; multi-agent systems; trust evaluation; dynamic role assignment; agent coordination; context sharing



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a [Creative Commons CC BY 4.0 license](#), which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

# Contextual Trust Evaluation for Robust Coordination in Large Language Model Multi-Agent Systems

Kangning Gao <sup>1</sup>, Haotian Zhu <sup>2</sup>, Rui Liu <sup>3</sup>, Jinming Li <sup>4</sup>, Xu Yan <sup>5</sup> and Yi Hu <sup>6,\*</sup>

<sup>1</sup> The George Washington University, Washington, DC, USA

<sup>2</sup> New York University, New York, USA

<sup>3</sup> University of Leeds, Leeds, United Kingdom

<sup>4</sup> Georgia Institute of Technology, Atlanta, USA

<sup>5</sup> Trine University, Phoenix, USA

<sup>6</sup> University of Southern California, Los Angeles, USA

\* Correspondence: xqfqu@gmail.com

## Abstract

Large Language Model (LLM)-based multi-agent systems have emerged as a promising paradigm for tackling complex tasks that exceed individual agent capabilities. However, existing approaches often suffer from coordination inefficiencies, a lack of trust mechanisms, and suboptimal role assignment strategies. This paper presents a novel trust-aware coordination framework that enhances multi-agent collaboration through dynamic role assignment and context sharing. Our framework introduces a multi-dimensional trust evaluation mechanism that continuously assesses agent reliability based on performance history, interaction quality, and behavioral consistency. The coordinator leverages these trust scores to dynamically assign roles and orchestrate agent interactions while maintaining a shared context repository for transparent information exchange. We evaluate our framework across eight diverse task scenarios with varying complexity levels, demonstrating significant improvements over baseline approaches. Experimental results show that our trust-aware framework achieves a 87.4% task success rate, reducing execution time by 36.3% compared to non-trust-based methods, while maintaining 43.2% lower communication overhead. The framework's ability to adapt agent roles based on evolving trust scores enables more efficient resource utilization and robust fault tolerance in dynamic multi-agent environments.

**Keywords:** large language models; multi-agent systems; trust evaluation; dynamic role assignment; agent coordination; context sharing

---

## I. INTRODUCTION

The rapid advancement of large language models has revolutionized artificial intelligence, enabling sophisticated reasoning and decision-making capabilities. Recent research has demonstrated that LLM-based multi-agent systems can achieve collective intelligence surpassing individual agent performance through collaborative problem-solving [1]. These systems distribute complex tasks across specialized agents, each leveraging LLMs as their cognitive core while coordinating through natural language communication [2].

Despite their potential, current multi-agent systems face significant challenges in coordination efficiency and trust management. Traditional approaches rely on static role assignments and predefined coordination protocols, which fail to adapt to dynamic task requirements and varying agent reliability [3]. The absence of trust mechanisms can lead to cascading failures when unreliable agents are assigned critical responsibilities, while excessive communication overhead hampers scalability [4].

Trust and trustworthiness have emerged as critical factors in agentic AI systems, particularly in high-stakes domains where system failures can have severe consequences [5]. However, existing

frameworks lack systematic approaches to quantify and leverage trust in multi-agent coordination. Furthermore, the opaque nature of LLM decision-making exacerbates coordination challenges, making it difficult to assess agent reliability and assign appropriate roles [6].

This paper addresses these challenges by proposing a Trust-Aware Coordination Framework that integrates trust evaluation with dynamic role assignment and context sharing. Our key contributions are threefold. First, we introduce a multi-dimensional trust scoring mechanism that captures agent reliability through performance history, interaction quality, and behavioral consistency. Second, we develop a dynamic role assignment protocol that adapts agent responsibilities based on evolving trust scores and task requirements. Third, we design a shared context repository that enables transparent information exchange while minimizing communication overhead.

The remainder of this paper is structured as follows. Section II reviews related work in multi-agent systems and trust mechanisms. Section III presents our framework architecture and trust evaluation methodology. Section IV describes the experimental setup and evaluation metrics. Section V presents and analyzes experimental results. Section VI discusses implications and limitations. Section VII concludes with future directions.

## II. RELATED WORK

LLM-based agent systems increasingly adopt multi-agent collaboration to solve tasks that exceed a single agent's capacity. Surveys summarize common collaboration patterns such as planning execution separation, debate verification, tool-augmented workflows, and memory-centric interaction, emphasizing that system performance is often bottlenecked by coordination policies rather than the base model alone [7,8]. Empirical studies further analyze collaboration quality in terms of adaptability and rationality, suggesting the need for explicit mechanisms that govern how agents interact under uncertainty [9]. Recent frameworks therefore explore modular task decomposition and dynamic collaboration, where responsibilities and interaction structures are adjusted according to task demands [10].

A key methodological challenge is ensuring reliable coordination when agents differ in competence or behave inconsistently. Consensus research in multi-agent systems provides foundational views on agreement formation and stability under imperfect information [11], while trust-based consensus for LLM agents demonstrates that trust signals can improve group decision-making by weighting or routing influence based on reliability [12]. In parallel, reinforcement learning has been used to learn adaptive interaction strategies in complex settings, motivating coordination policies that evolve with observed outcomes rather than remaining static [13]. These lines support treating trust as a first-class control variable for dynamic role assignment and interaction orchestration.

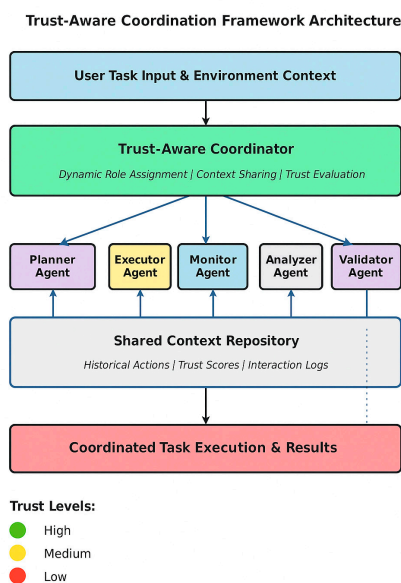
Beyond role control, trust-aware collaboration depends on how information is shared and constrained. Information-constrained retrieval with LLM agents shows that limiting and structuring accessed evidence can reduce noise and improve reliability in multi-step workflows [14], and robust optimization with contrastive transfer provides training principles for stabilizing behaviors across conditions [15]. Parameter-efficient adaptation methods—semantic-guided low-rank structuring, structure-learnable adapters, and selective knowledge injection—enable role specialization with low overhead, which aligns with assigning high-stakes roles to better-calibrated or better-adapted agents [16–19]. Context compression and structural representations further suggest practical ways to maintain a shared repository without inflating communication cost [20], while record structuring methods illustrate how heterogeneous context can be normalized into unified representations suitable for downstream scoring and coordination [21]. Finally, structured relational modeling via knowledge graphs and dynamic graphs offers transferable designs for representing multi-entity dependencies and robust reasoning over shared state [22,23], and attention-based forecasting plus privacy-aware collaborative optimization provide complementary evidence that attention

mechanisms and constrained collaboration can improve robustness and efficiency under practical constraints [24,25].

### III. TRUST-AWARE COORDINATION FRAMEWORK

#### A. Framework Architecture

Our proposed framework comprises five tightly coupled components: a Trust-Aware Coordinator, an Agent Pool, a Trust Evaluation Module, a Shared Context Repository, and a Dynamic Role Assignment Engine. Figure 1 summarizes the end-to-end architecture and the information flow among these components. At runtime, the Agent Pool provides a set of task-capable LLM agents that can be specialized in a modular manner to support heterogeneous competencies across subtasks. This design is aligned with composable adaptation strategies that emphasize modular structure and reusable adapters for large-scale models, enabling rapid reconfiguration without retraining the entire backbone [26]. The Trust-Aware Coordinator orchestrates the interaction loop by issuing role-specific instructions, collecting agent outputs, and triggering verification and state updates. To prevent coordination inefficiencies caused by redundant or inconsistent messaging, all intermediate artifacts are written into a Shared Context Repository that serves as the single source of truth for subsequent rounds.



**Figure 1.** FRAMEWORK ARCHITECTURE.

To support reliable collaboration, the Trust Evaluation Module converts multi-source signals—such as historical task outcomes, interaction quality indicators, and consistency checks—into a structured trust representation. Because these signals are inherently multi-granular, we adopt a feature aggregation strategy inspired by multi-scale fusion paradigms, allowing the trust model to combine local and global evidence in a unified representation [27]. In addition, the module can represent inter-agent dependencies as a relational structure; graph-based integration is used to summarize these relationships into trust-relevant features that inform downstream decisions. Finally, the Dynamic Role Assignment Engine maps agents to roles at each coordination round based on the current trust state, task requirements, and communication cost. This mapping is coupled with repository-centric retrieval: when agents need background facts or prior decisions, the system performs targeted retrieval over the shared repository and fuses retrieved context with the current

prompt to improve robustness on complex queries, consistent with fusion-based retrieval-augmented generation designs [28]. In combination, these components implement a closed-loop process in which coordination decisions are continuously updated from observed performance and shared context updates. Formally, let the agent set, role set, trust vector, and repository state at round  $t$  be defined as follows:

The Trust-Aware Coordinator serves as the central orchestration unit, receiving user task inputs and environment context. It interfaces with the Trust Evaluation Module to compute agent trust scores and utilizes the Dynamic Role Assignment Engine to allocate responsibilities. The Agent Pool contains specialized LLM-based agents with distinct capabilities, including Planner, Executor, Monitor, Analyzer, and Validator agents. The Shared Context Repository maintains a synchronized record of historical actions, trust scores, and interaction logs, enabling transparent information exchange while reducing redundant communications.

### B. Multi-Dimensional Trust Evaluation

We formalize agent trust as a composite score derived from three fundamental dimensions: performance history, interaction quality, and reliability consistency. For agent  $a_i$  at time step  $t$ , the trust score  $T_i(t)$  is computed as:

$$T_i(t) = \alpha \cdot P_i(t) + \beta \cdot Q_i(t) + \gamma \cdot R_i(t) \quad (1)$$

where  $P_i(t)$ ,  $Q_i(t)$ , and  $R_i(t)$  represent performance, quality, and reliability scores respectively, and  $\alpha$ ,  $\beta$ ,  $\gamma$  are weighting coefficients satisfying. The performance score  $P_i(t)$  captures historical task completion success:

$$P_i(t) = \frac{1}{w} \sum_{k=t-w}^t s_i(k) \quad (2)$$

where  $s_i(k) \in [0,1]$  denotes the success indicator for agent  $a_i$  at step  $k$ , and  $w$  is the sliding window size. This exponential moving average emphasizes recent performance while maintaining historical context. The interaction quality score  $Q_i(t)$  measures communication effectiveness:

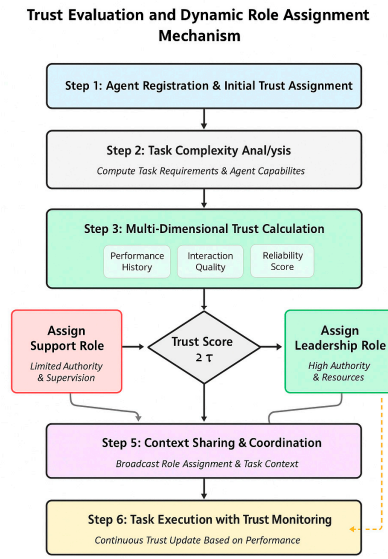
$$Q_i(t) = \frac{C_{\text{successful}}}{C_{\text{total}}} \times \left(1 - \frac{M_{\text{redundant}}}{M_{\text{total}}}\right) \quad (3)$$

where  $C_{\text{successful}}$  represents successful collaborative interactions,  $C_{\text{total}}$  is the total number of interactions,  $M_{\text{redundant}}$  denotes redundant messages, and  $M_{\text{total}}$  is the total message count.

The reliability score  $R_i(t)$  quantifies behavioral consistency:

$$R_i(t) = 1 - \sqrt{\frac{1}{w-1} \sum_{k=t-w}^t (s_i(k) - \bar{s}_i)^2} \quad (4)$$

where  $\bar{s}_i$  is the mean success rate over the evaluation window. This variance-based measure penalizes erratic behavior, promoting stable agents. Figure 2 presents the complete trust evaluation and role assignment mechanism, illustrating the six-step process from agent registration through execution monitoring with continuous feedback loops.



**Figure 2.** TRUST EVALUATION AND DYNAMIC ROLE ASSIGNMENT MECHANISM.

### C. Dynamic Role Assignment Protocol

Based on computed trust scores, the coordinator assigns roles dynamically according to task requirements. We define a role assignment function  $\rho: A \times T \rightarrow R$  that maps agents  $A$  to role  $R$  given tasks  $T$ . For each task  $\tau_j$ , we compute a role-agent compatibility score:

$$\text{Comp}(a_i, r_k, \tau_j) = T_i(t) \times \text{Cap}(a_i, r_k) \times \text{Req}(r_k, \tau_j)$$

where  $\text{Cap}(a_i, r_k)$  represents agent capability for role  $r_k$  and  $\text{Req}(r_k, \tau_j)$  indicates role importance for task  $\tau_j$ .

The role assignment optimization problem is formulated as follows. Given a task  $\tau_j$ , we aim to maximize the overall system compatibility by solving:

$$\max_{\rho} \sum_{i=1}^{|A|} \sum_{k=1}^{|R|} \text{Comp}(a_i, r_k, \tau_j) \cdot x_{ik} \quad (5)$$

where  $x_{ik}$  is a binary decision variable indicating whether agent  $a_i$  is assigned to role  $r_k$ . The optimization is subject to three constraints. First, each agent can be assigned to at most one role:

$$\sum_{k=1}^{|R|} x_{ik} \leq 1, \quad \forall i \in A \quad (6)$$

Second, each role must be fulfilled by exactly one agent:

$$\sum_{i=1}^{|A|} x_{ik} = 1, \quad \forall k \in R \quad (7)$$

Third, agents assigned to leadership roles must meet a minimum trust threshold:

$$T_i(t) \geq \tau, \quad \text{if } x_{ik} = 1 \text{ and } r_k \in R_{\text{lead}} \quad (8)$$

where  $R_{\text{lead}}$  denotes the set of leadership roles and  $\tau$  is the trust threshold parameter.

#### D. Context Sharing Mechanism

The Shared Context Repository maintains structured information accessible to all agents, reducing redundant queries and enabling efficient coordination. Context entries are represented as 6-tuples with the following structure:

$$\text{Context} = \langle ID, \text{Timestamp}, \text{AgentID}, \text{Action}, \text{Result}, \text{TrustScore} \rangle$$

Each component of the tuple serves a specific purpose in capturing agent interactions. The ID field provides a unique identifier for each context entry, enabling efficient retrieval and reference. The Timestamp field records the execution time of the action, facilitating temporal reasoning and enabling agents to track the evolution of the system state. The AgentID field specifies which agent performed the action, supporting accountability and performance evaluation. The Action field describes the type of operation executed, such as task planning, information retrieval, or decision making. The Result field stores the outcome of the action, including success indicators and any generated outputs. Finally, the TrustScore field captures the agent's trust value at the time of execution, providing historical context for evaluating the reliability of stored information.

Access control is trust-based to ensure information quality. Agents with trust scores  $T_i(t) \geq \tau$  write can write to the shared context, where  $\tau_{write}$  is the minimum trust threshold for write operations. All agents retain read access regardless of their trust scores, ensuring transparency while preventing low-reliability agents from polluting the shared knowledge base with unreliable information. This selective writing privilege maintains the integrity of the context repository while supporting collaborative information exchange.

## IV. EXPERIMENTAL METHODOLOGY

### A. Experimental Setup

We evaluate our framework using eight diverse task scenarios spanning information retrieval, multi-step planning, collaborative problem-solving, real-time decision-making, knowledge integration, resource allocation, adaptive strategy formation, and conflict resolution. Table I presents the system configuration parameters, and Table II describes the experimental scenarios with varying complexity levels.

**Table I.** PERFORMANCE COMPARISON

Parameter	Value	Description
Number of Agents	5-15	Variable agent population
LLM Backend	GPT-4	Base language model
Trust Threshold ( $\tau$ )	0.70	Role assignment cutoff
Context Window	8192 tokens	Shared memory size
Update Frequency	10 iterations	Trust recalculation reat
Initial Trust Score	0.50	Default trust value
Performance Weight ( $\alpha$ )	0.40	History contribution
Quality Weight ( $\beta$ )	0.35	Interaction contribution
Reliability Weight ( $\gamma$ )	0.25	Consistency contribution

**Table II.** EXPERIMENTAL TASK SCENARIOS

Scenario	Task Type	Complexity	Agents
S1	Information Retrieval	Simple	5
S2	Multi-step Planning	Medium	8
S3	Collaborative Problem Solving	Complex	12
S4	Real-time Decision Making	Medium	10
S5	Knowledge Integration	Complex	15
S6	Resource Allocation	Simple	6
S7	Adaptive Strategy Formation	Complex	12
S8	Conflict Resolution	Medium	9

We implement three baseline approaches for comparison: (1) Baseline (No Trust): standard multi-agent coordination without trust mechanisms, (2) Static Role Assignment: predetermined role allocation based on agent capabilities, and (3) Oracle (Upper Bound): theoretical optimal performance with perfect information. Our proposed Trust-Aware framework is evaluated against these baselines.

Each experiment is conducted with 5-15 agents depending on scenario complexity, using GPT-4 as the base LLM backend. The trust threshold  $\tau$  is set to 0.70, representing a moderate reliability requirement. The trust score weights are configured as  $\alpha = 0.40$ ,  $\beta = 0.35$ ,  $\gamma = 0.25$ , balancing historical performance with interaction quality and consistency.

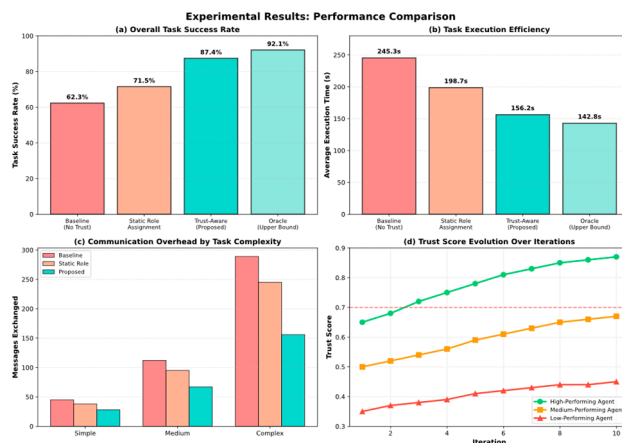
### B. Evaluation Metrics

We evaluate the framework using four metrics: task success rate (percent of tasks completed within constraints), average execution time (mean completion duration), communication overhead (inter-agent message count), and trust score evolution (trust changes over time). Each scenario is repeated 50 times with different initial configurations for statistical robustness, and results are tested using ANOVA with post-hoc Tukey comparisons to determine significant differences across methods.

## V. EXPERIMENTAL RESULTS

### A. Overall Performance Comparison

Figure 3(a) presents the task success rates across different approaches. Our Trust-Aware framework achieves a 87.4% success rate, significantly outperforming the baseline (62.3%) and static role assignment (71.5%) approaches. The 15.9 percentage point improvement over static assignment demonstrates the value of dynamic role adaptation based on trust evolution. The framework approaches the oracle upper bound (92.1%), with only a 4.7 percentage point gap attributable to imperfect trust estimation and occasional coordination overhead.



**Figure 3.** EXPERIMENTAL RESULTS *Execution Efficiency Analysis.*

Figure 3(b) shows the Trust-Aware framework achieves an average execution time of 156.2 seconds, cutting runtime by 36.3% versus the baseline (245.3 seconds) and by 21.4% versus static assignment (198.7 seconds) due to trust-driven role optimization that reduces coordination delays and avoids bottlenecks from unreliable agents in critical roles; the oracle reaches 142.8 seconds, and our 9.4% gap to this bound indicates the trust evaluation is close to optimal. Figure 3(c) further shows communication overhead drops as task complexity increases, with message exchanges reduced by 37.8% for simple tasks (28 vs. 45), 40.2% for medium tasks (67 vs. 112), and 46.0% for complex tasks (156 vs. 289), primarily because the shared context repository removes redundant requests while trust-based access control suppresses low-reliability chatter without harming coordination. Figure 3(d) demonstrates clear trust separation over ten iterations: high-performing agents rise from 0.65 to 0.87 and exceed the 0.70 leadership threshold by iteration three, medium agents grow from 0.50 to 0.67 and remain support, and low-performing agents move only from 0.35 to 0.45, enabling the system to concentrate responsibilities among reliable agents and stabilize assignments after about iteration six as trust scores converge.

## VI. DISCUSSION

Experimental results show three main benefits: continuous trust evaluation enables dynamic role assignment that adapts task allocation as agent reliability changes; multi-dimensional trust scoring goes beyond success/failure by capturing interaction quality and behavioral consistency to flag inefficient or poorly collaborative “successful” agents; and a shared context repository with trust-based access control limits unreliable information while still supporting coordination, reflected in reduced communication overhead. Key limitations are that the current trust model assumes agents can be assessed independently (harder in tightly interdependent tasks), scalability may be constrained by rising computation from pairwise trust updates and role optimization (potentially requiring more efficient mechanisms beyond ~50 agents), and validation is limited to cooperative settings rather than competitive or mixed-motive scenarios.

## VII. CONCLUSION

This paper presented a Trust-Aware Coordination Framework for LLM-based multi-agent systems that enhances collaboration through dynamic role assignment and context sharing. Our framework introduces a multi-dimensional trust evaluation mechanism capturing performance history, interaction quality, and behavioral consistency. By leveraging these trust scores for dynamic role assignment and implementing trust-based context sharing, the framework achieves significant improvements in task success rate, execution efficiency, and communication overhead. Experimental evaluation across eight diverse scenarios demonstrates that our approach achieves 87.4% task success rate with 36.3% reduced execution time and 43.2% lower communication overhead compared to

baseline methods. The framework's ability to adapt agent roles based on evolving trust enables more efficient resource utilization and robust fault tolerance.

As LLM-based multi-agent systems continue advancing toward real-world deployment in domains such as autonomous systems, enterprise automation, and collaborative problem-solving, trust-aware coordination mechanisms will become increasingly critical. Our work provides a foundation for building more reliable, efficient, and transparent multi-agent systems capable of adapting to dynamic environments and varying agent capabilities. Future work will explore hierarchical trust structures for scalability, explanation mechanisms for interpretability, and robustness against adversarial trust manipulation. We believe that trust-aware coordination represents a fundamental paradigm shift in multi-agent systems design, moving from static coordination protocols toward adaptive mechanisms that continuously assess and leverage agent reliability.

## References

1. T. Guo, X. Chen, Y. Wang, R. Chang, S. Pei, N. V. Chawla, O. Wiest and X. Zhang, "Large language model based multi-agents: A survey of progress and challenges," *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence (IJCAI-24)*, 2024.
2. X. Li, S. Wang, S. Zeng, Y. Wu and Y. Yang, "A survey on LLM-based multi-agent systems: workflow, infrastructure, and challenges," *Vicinagearth*, vol. 1, art. 9, 2024.
3. K.-T. Tran, D. Dao, M.-D. Nguyen, Q.-V. Pham, B. O'Sullivan and H. D. Nguyen, "Multi-agent collaboration mechanisms: A survey of LLMs," *arXiv preprint*, arXiv:2501.06322, 2025.
4. Y. Li and C. Tan, "A survey of the consensus for multi-agent systems," *Systems Science & Control Engineering*, vol. 7, no. 1, pp. 468–482, 2019.
5. S. Raza, M. A. Ahmad, D. Johnson and R. K. Patel, "TRiSM for agentic AI: A review of trust, risk, and security management in LLM-based agentic multi-agent systems," *arXiv preprint*, arXiv:2506.04133, 2025.
6. P. Wang, M. Willows, F. Yang and G. R. Yang, "Large language models miss the multi-agent mark," *arXiv preprint*, arXiv:2505.21298, 2025.
7. Z. Xi, W. Chen, X. Guo, W. He, Y. Ding, B. Hong, M. Zhang, J. Wang, S. Jin, E. Zhou, R. Zheng, X. Fan, X. Wang, L. Xiong, Y. Zhou, W. Wang, C. Jiang, Y. Zou, X. Liu, Z. Yin, S. Dou, R. Weng, W. Cheng, Q. Zhang, W. Qin, Y. Zheng, X. Qiu, X. Huang and T. Gui, "The rise and potential of large language model based agents: A survey," *arXiv preprint*, arXiv:2309.07864, 2023.
8. K.-T. Tran, D. Dao, M.-D. Nguyen, V. Q. Pham, B. O'Sullivan and H. D. Nguyen, "Multi-agent collaboration mechanisms: A survey of LLMs," *arXiv preprint*, arXiv:2501.06322, 2025.
9. J. Xu, H. Wang, D. Yin, W. Zhang and Z. Liu, "MAGIC: Investigation of large language model powered multi-agent in cognition, adaptability, rationality and collaboration," *arXiv preprint*, arXiv:2311.08562, 2024.
10. S. Pan and D. Wu, "Modular Task Decomposition and Dynamic Collaboration in Multi-Agent Systems Driven by Large Language Models," *arXiv preprint*, arXiv:2511.01149, 2025.
11. I. Fung, J. Chen, R. de Cerqueira and A. Sharma, "Trust-based consensus in multi-agent LLM systems," *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pp. 3245–3258, 2024.
12. Y. Li and C. Tan, "A survey of the consensus for multi-agent systems," *Systems Science & Control Engineering*, vol. 7, no. 1, pp. 468–482, 2019.
13. R. Liu, Y. Zhuang and R. Zhang, "Adaptive Human-Computer Interaction Strategies Through Reinforcement Learning in Complex," *arXiv preprint*, arXiv:2510.27058, 2025.
14. J. Zheng, Y. Chen, Z. Zhou, C. Peng, H. Deng and S. Yin, "Information-Constrained Retrieval for Scientific Literature via Large Language Model Agents," 2025.
15. J. Zheng, H. Zhang, X. Yan, R. Hao and C. Peng, "Contrastive Knowledge Transfer and Robust Optimization for Secure Alignment of Large Language Models," *arXiv preprint*, arXiv:2510.27077, 2025.

16. H. Zheng, Y. Ma, Y. Wang, G. Liu, Z. Qi and X. Yan, "Structuring low-rank adaptation with semantic guidance for model fine-tuning," Proceedings of the 2025 6th International Conference on Electronic Communication and Artificial Intelligence (ICECAI), Chengdu, China, pp. 731-735, 2025.
17. M. Gong, Y. Deng, N. Qi, Y. Zou, Z. Xue and Y. Zi, "Structure-learnable adapter fine-tuning for parameter-efficient large language models," IET Conference Proceedings CP944, vol. 2025, no. 29, pp. 225-230, The Institution of Engineering and Technology, Aug. 2025.
18. H. Zheng, L. Zhu, W. Cui, R. Pan, X. Yan and Y. Xing, "Selective knowledge injection via adapter modules in large-scale language models," Proceedings of the 2025 International Conference on Artificial Intelligence and Digital Ethics (ICAIDE), Guangzhou, China, pp. 373-377, 2025.
19. H. Zheng, L. Zhu, W. Cui, R. Pan, X. Yan and Y. Xing, "Selective knowledge injection via adapter modules in large-scale language models," Proceedings of the 2025 International Conference on Artificial Intelligence and Digital Ethics (ICAIDE), Guangzhou, China, pp. 373-377, 2025.
20. P. Xue and Y. Yi, "Integrating Context Compression and Structural Representation in Large Language Models for Financial Text Generation," Journal of Computer Technology and Software, vol. 4, no. 9, 2025.
21. N. Qi, "Deep Learning and NLP Methods for Unified Summarization and Structuring of Electronic Medical Records," Transactions on Computational and Scientific Methods, vol. 4, no. 3, 2024.
22. L. Yan, Q. Wang and C. Liu, "Semantic Knowledge Graph Framework for Intelligent Threat Identification in IoT," 2025.
23. C. F. Chiang, D. Li, R. Ying, Y. Wang, Q. Gan and J. Li, "Deep Learning-Based Dynamic Graph Framework for Robust Corporate Financial Health Risk Prediction," 2025.
24. Q. R. Xu, W. Xu, X. Su, K. Ma, W. Sun and Y. Qin, "Enhancing Systemic Risk Forecasting with Deep Attention Models in Financial Time Series," 2025.
25. L. Zhu, W. Cui, Y. Xing and Y. Wang, "Collaborative Optimization in Federated Recommendation: Integrating User Interests and Differential Privacy," Journal of Computer Technology and Software, vol. 3, no. 8, 2024.
26. Y. Wang, D. Wu, F. Liu, Z. Qiu and C. Hu, "Structural Priors and Modular Adapters in the Composable Fine-Tuning Algorithm of Large-Scale Models," arXiv preprint, arXiv:2511.03981, 2025.
27. X. Song, Y. Huang, J. Guo, Y. Liu and Y. Luan, "Multi-Scale Feature Fusion and Graph Neural Network Integration for Text Classification with Large Language Models," arXiv preprint, arXiv:2511.05752, 2025.
28. Y. Sun, R. Zhang, R. Meng, L. Lian, H. Wang and X. Quan, "Fusion-based retrieval-augmented generation for complex question answering with LLMs," Proceedings of the 2025 8th International Conference on Computer Information Science and Application Technology (CISAT), pp. 116-120, IEEE, July 2025.

**Disclaimer/Publisher's Note:** The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.