Symbolic Guidance for LLM Agents in Distributed Multiagent Coordination

Extended Abstract

Ben Rachmut
Washington University in St. Louis
Saint Louis, MO, United States
benr@wustl.edu

Ning Zhang
Washington University in St. Louis
Saint Louis, MO, United States
zhang.ning@wustl.edu

Yevgeniy Vorobeychik
Washington University in St. Louis
Saint Louis, MO, United States
yvorobeychik@wustl.edu

William Yeoh
Washington University in St. Louis
Saint Louis, MO, United States
wyeoh@wustl.edu

ABSTRACT

Large language models (LLMs) are increasingly deployed as autonomous agents in multi-agent systems, yet their ability to reliably execute distributed coordination protocols remains poorly understood. While AgentsNet, a benchmark framework for distributed coordination among LLM agents, enables such coordination, granting full reasoning autonomy often leads to inconsistent or degraded performance in complex domains.

We hypothesize that coordination can be improved by regulating agent autonomy through symbolic guidance derived from established algorithms. To investigate this, we introduce the *Symbolic Guidance Taxonomy (SGT)*, which characterizes a spectrum of autonomy ranging from open-ended natural language reasoning to fully prescribed algorithmic execution, with intermediate levels providing partial pseudocode guidance. Our results show that *intermediate autonomy levels* consistently *outperform* both unguided agents and fully prescriptive specifications. These findings identify autonomy regulation as a key design principle for LLM-based distributed coordination.

KEYWORDS

Large Language Models; Multiagent Systems; Distributed Coordination; Symbolic Guidance; Adjustable Autonomy

ACM Reference Format:
Ben Rachmut, Ning Zhang, Yevgeniy Vorobeychik, and William Yeoh. 2026. Symbolic Guidance for LLM Agents in Distributed Multiagent Coordination: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 3 pages. https://doi.org/10.65109/KEVK7310

1 INTRODUCTION

Coordinating multiple autonomous agents to achieve a common objective is a long-standing challenge in AI. Applications ranging from meeting scheduling and multi-robot coordination to satellite constellations require agents to make local decisions under limited information while collectively achieving coherent global behavior [5, 7, 11, 15]. These distributed coordination problems are inherently complex, as agents must reason about both their own actions and those of others in dynamic, decentralized environments [1, 9].

Classical symbolic approaches provide principled solutions through rigorously defined distributed algorithms with formal guarantees of correctness and convergence [6]. However, these methods rely on complete and well-specified problem formulations, making them highly dependent on precise specifications and thus unsuited to informal or underspecified settings. In contrast, recent work has explored LLMs as autonomous agents capable of interpreting instructions and interacting in natural language [8, 14]. While LLM agents offer flexibility and robustness to under-specification, they often struggle to coordinate reliably, particularly when global coherence must emerge from local interactions [2]. There is thus a significant opportunity in bridging the two paradigms [10].

To this end, we introduce the *Symbolic Guidance Taxonomy (SGT)*, a framework for characterizing how symbolic guidance modulates the autonomy of LLM-based agents in distributed coordination settings. Inspired by adjustable autonomy [12, 13], SGT organizes agents along a spectrum from fully open-ended reasoning with no symbolic structure to fully prescribed algorithmic execution, with intermediate levels providing partial, function-level guidance.

Using SGT, we evaluate how varying degrees of symbolic guidance affect coordination performance on standard distributed graph problems, building on the AgentsNet benchmarking framework [2]. Our results show that *intermediate autonomy levels are most effective*: agents provided with partial, function-level guidance consistently outperform both unguided agents and fully prescriptive specifications. These findings highlight symbolic guidance as a key mechanism for balancing autonomy and structure in LLM-based multiagent coordination.

2 SYMBOLIC GUIDANCE TAXONOMY

LLM agents can be guided in different ways when solving distributed coordination problems, ranging from being left entirely to

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). https://doi.org/10.65109/KEVK7310

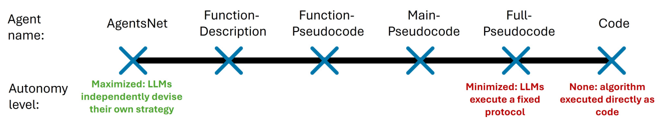

Figure 1: The Symbolic Guidance Taxonomy (SGT) autonomy spectrum, from maximal autonomy (AgentsNet) to no autonomy (Code).

their own devices to being constrained by precise symbolic instructions. The *Symbolic Guidance Taxonomy (SGT)* provides a unified framework for describing this design space by formalizing how symbolic structure regulates agent autonomy in LLM-based multiagent systems. Formally, let A denote a symbolic algorithm for a distributed coordination problem. A guidance strategy is defined as a mapping $G(S, E; A) \mapsto I$, where $S \subseteq A$ specifies which components of the algorithm are revealed to the agent, E specifies how these components are encoded (e.g., natural language or pseudocode), and I denotes the resulting guidance provided to the agent. Agent autonomy is determined by both the amount of algorithmic information exposed through S and the precision of its representation through E.

This formulation induces a spectrum of autonomy. At one extreme, agents receive no symbolic guidance ($S = \emptyset$), corresponding to maximal autonomy where coordination strategies are self-generated; at the other extreme, agents are provided with a complete algorithmic specification ($S = A$), corresponding to minimal autonomy where behavior is fully prescribed. *Intermediate-autonomy agents lie between these extremes: they are guided by partial algorithmic structure while retaining freedom in how this structure is instantiated during reasoning.*

SGT can be viewed as an operational instantiation of adjustable autonomy for LLM-based agents. Classical work on adjustable autonomy focuses on dynamically transferring control between agents and external decision-makers depending on context and uncertainty [12, 13]. In contrast, SGT regulates autonomy by varying the amount and representation of symbolic algorithmic information provided to the agent, offering a controllable mechanism for balancing flexibility and reliability in distributed coordination settings.

Figure 1 summarizes the resulting autonomy spectrum and the agent types evaluated in this work, and serves as the basis for our empirical analysis of how regulating autonomy through symbolic guidance affects coordination performance.

3 EXPERIMENTAL EVALUATION

We evaluate how different levels of symbolic guidance influence the performance of LLM agents on distributed coordination problems using extended AgentsNet [2], across multiple representative LLM models.[1] Specifically, we consider three canonical graph-based domains—graph coloring, matching, and vertex cover—and examine both *local variants*, where feasible solutions can be obtained using local heuristics, and *global variants*, where optimal solutions require stronger coordination. For these domains, we employ the *Maximal Independent Set (MIS)* heuristic as the underlying symbolic algorithm, a classical local distributed coordination method that resolves conflicts through priority-based decisions using only neighborhood information [3, 4]. We evaluate coordination quality using a unified error metric based on the sum of squared errors (SSE), which captures deviations from feasibility or optimality, with lower values indicating better coordination. For each domain, we instantiate multiple agent types corresponding to distinct points along the SGT autonomy spectrum shown in Figure 1. Across all domains and problem variants, our results reveal a consistent pattern: symbolic guidance profoundly shapes coordination outcomes. Structured scaffolding significantly improves coherence and stability, while excessive prescriptiveness can reduce adaptability. In particular, agents operating at intermediate autonomy levels consistently outperform both unguided AgentsNet agents and fully specified algorithmic agents. Function-level guidance, which exposes partial algorithmic structure while allowing flexibility in execution, yields the most robust performance across settings.

	Local Coordination (MIS)		
	(Δ+1)-Coloring	**Matching**	**Vertex Cover**
AgentsNet	1.99 ± 0.43	1.03 ± 0.34	2.74 ± 0.44
Function-Description	0.55 ± 0.26	**0.78 ± 0.31**	**2.21 ± 0.34**
Function-Pseudocode	**0.34 ± 0.12**	0.80 ± 0.30	3.03 ± 0.40
Main-Pseudocode	1.05 ± 0.33	1.85 ± 0.46	2.34 ± 0.38
Full-Pseudocode	0.76 ± 0.25	2.61 ± 0.49	3.48 ± 0.25
	Global Coordination (MIS)		
	3-Coloring	**Matching**	**Vertex Cover**
AgentsNet	2.35 ± 0.33	1.24 ± 0.11	2.90 ± 0.31
Function-Description	1.92 ± 0.30	1.11 ± 0.20	**2.35 ± 0.32**
Function-Pseudocode	**1.56 ± 0.24**	**1.01 ± 0.17**	3.29 ± 0.39
Main-Pseudocode	2.49 ± 0.33	1.37 ± 0.26	2.83 ± 0.26
Full-Pseudocode	1.89 ± 0.36	1.85 ± 0.25	3.68 ± 0.26

Table 1: Initial results (Gemini-2.5 Flash): mean $\ln(1 + \text{SSE}) \pm \text{SEM}$ **for local and global coordination using MIS.**

Table 1 summarizes these initial results for Gemini-2.5 Flash, illustrating the non-monotonic relationship between autonomy and coordination performance, with intermediate-autonomy agents achieving the strongest overall performance across both local and global variants.

4 CONCLUSION

We introduced the Symbolic Guidance Taxonomy (SGT), a principled framework for analyzing how symbolic structure regulates autonomy in LLM-based multiagent coordination. Through initial experiments on distributed graph problems, we showed that coordination performance does not improve monotonically with either increased autonomy or increased guidance. Instead, agents operating at intermediate autonomy levels—guided by partial, function-level symbolic structure—consistently achieve the strongest performance. These findings highlight symbolic guidance as an effective mechanism for balancing flexibility and reliability in distributed coordination. Future work will explore adaptive and heterogeneous guidance strategies that dynamically adjust agent autonomy.

[1] Code is publicly available at https://github.com/YODA-Lab/Symbolic-Guidance-Taxonomy.

ACKNOWLEDGMENTS

This work was partially supported by the Office of Naval Research (N00014-24-1-2663).

REFERENCES

[1] Ferdinando Fioretto, Enrico Pontelli, and William Yeoh. 2018. Distributed Constraint Optimization Problems and Applications: A Survey. *Journal of Artificial Intelligence Research* 61 (2018), 623–698.

[2] Florian Grötschla, Luis Müller, Jan Tönshoff, Mikhail Galkin, and Bryan Perozzi. 2025. AgentsNet: Coordination and Collaborative Reasoning in Multi-Agent LLMs. *arXiv preprint* (2025). arXiv:2507.08616

[3] Nathan Linial. 1992. Locality in distributed graph algorithms. *SIAM J. Comput.* 21, 1 (1992), 193–201.

[4] Michael Luby. 1985. A simple parallel algorithm for the maximal independent set problem. In *Proceedings of the Annual ACM Symposium on Theory of Computing (STOC)*. 1–10.

[5] Rajiv Maheswaran, Milind Tambe, Emma Bowring, Jonathan Pearce, and Pradeep Varakantham. 2004. Taking DCOP to the Real World: Efficient Complete Solutions for Distributed Event Scheduling. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 310–317.

[6] Pragnesh Jay Modi, Wei-Min Shen, Milind Tambe, and Makoto Yokoo. 2005. ADOPT: Asynchronous Distributed Constraint Optimization with Quality Guarantees. *Artificial Intelligence* 161, 1–2 (2005), 149–180.

[7] Duc Thien Nguyen, William Yeoh, Hoong Chuin Lau, and Roie Zivan. 2019. Distributed Gibbs: A Linear-Space Sampling-Based DCOP Algorithm. *Journal of Artificial Intelligence Research* 64 (2019), 705–748.

[8] Joon Sung Park, Joseph O'Brien, Carrie Jun Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein. 2023. Generative Agents: Interactive Simulacra of Human Behavior. In *Proceedings of the Annual ACM Symposium on User Interface Software and Technology (UIST)*. 1–22.

[9] David Peleg. 2000. *Distributed Computing: A Locality-Sensitive Approach.* SIAM.

[10] Gauthier Picard, William Yeoh, and Roie Zivan. 2026. Agentic LLMs and Distributed Constraint Reasoning: A Symbiotic Perspective for Neurosymbolic Multi-Agent Systems. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

[11] Pierre Rust, Gauthier Picard, and Fano Ramparany. 2022. Resilient Distributed Constraint Reasoning to Autonomously Configure and Adapt IoT Environments. *ACM Transactions on Internet Technology* 22, 4 (2022), 1–31.

[12] Paul Scerri, David Pynadath, and Milind Tambe. 2001. Adjustable Autonomy in Real-World Multi-Agent Environments. In *Proceedings of the International Conference on Autonomous Agents (ICA)*. 300–307.

[13] Paul Scerri, David Pynadath, and Milind Tambe. 2002. Towards adjustable autonomy for the real world. *Journal of Artificial Intelligence Research* 17 (2002), 171–228.

[14] Lei Wang, Chen Ma, Xueyang Feng, Zeyu Zhang, Hao Yang, Jingsen Zhang, Zhiyuan Chen, Jiakai Tang, Xu Chen, Yankai Lin, et al. 2024. A survey on large language model based autonomous agents. *Frontiers of Computer Science* 18, 6 (2024), 186345.

[15] William Yeoh, Ariel Felner, and Sven Koenig. 2010. BnB-ADOPT: An Asynchronous Branch-and-Bound DCOP Algorithm. *Journal of Artificial Intelligence Research* 38 (2010), 85–133.