Agentic LLMs and Distributed Constraint Reasoning: A Symbiotic Perspective for Neurosymbolic Multi-Agent Systems

Blue Sky Ideas Track

Gauthier Picard
DTIS, ONERA, Université de Toulouse
Toulouse, France
gauthier.picard@onera.fr

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

Roie Zivan
Ben-Gurion University of the Negev
Beer-Sheva, Israel
zivanr@bgu.ac.il

ABSTRACT

Distributed Constraint Reasoning (DCR) has long provided a principled framework for modeling and solving multi-agent coordination and optimization problems. However, its practical adoption in real-world, human-centric domains has been hindered by the challenge of translating human intentions, preferences, and constraints into formal symbolic models. At the same time, recent advances in LLMs have enabled powerful agentic capabilities, including natural language understanding, flexible reasoning, and interactive problem solving, but these systems lack the formal rigor and guarantees needed for scalable multi-agent coordination. In this paper, we argue that the convergence of these two paradigms offers a timely and transformative opportunity. We articulate several synergistic research directions: leveraging LLMs for translating natural language into DCR specifications, eliciting and refining user preferences, and enhancing inter-agent communication; and conversely, applying DCR models and algorithms to improve coordination, structured reasoning, resource allocation, and communication sensitivity in Agentic LLM systems. Together, these threads point toward hybrid neurosymbolic systems that combine the adaptability of LLMs with the mathematical rigor of DCR.

KEYWORDS

LLMs, Agentic AI, Distributed Constraint Reasoning

ACM Reference Format:
Gauthier Picard, William Yeoh, and Roie Zivan. 2026. Agentic LLMs and Distributed Constraint Reasoning: A Symbiotic Perspective for Neurosymbolic Multi-Agent Systems: Blue Sky Ideas Track. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), Paphos, Cyprus, May 25 – 29, 2026*, IFAAMAS, 6 pages. https://doi.org/10.65109/YNZM5391

1 INTRODUCTION

Multi-agent systems (MAS) study scenarios involving multiple autonomous entities, or agents, interacting to achieve individual and shared goals. A common modeling framework for such interactions is *distributed constraint reasoning (DCR)* [61], where agents control variables, assign them values from predefined domains, and coordinate under constraints that capture the implications of their choices. Over the past three decades, extensive research has developed DCR models, algorithms, and metrics to enable cooperative decision making among agents [17]. DCR problems typically take two forms. *Distributed constraint satisfaction problems (DCSPs)* seek assignments satisfying all hard constraints [61]. *Distributed constraint optimization problems (DCOPs)* further account for preferences or costs, aiming for solutions that minimize violation penalties or maximize collective utility [33].

Historically, these models assumed that agents accurately represent human intentions, preferences, and constraints – an assumption often unrealistic in practice. Yet recent advances in *LLMs* and *generative AI* provide the missing link: LLMs can translate between human language and the formal specifications required by DCR, enabling richer human-AI interaction. Moreover, LLMs exhibit emergent reasoning abilities [3, 51] and can themselves act as reasoning components within DCR frameworks [47].

This position paper envisions a synergistic unification of DCR and Agentic LLM systems. LLMs can enhance DCR by translating natural language, eliciting preferences, and supporting flexible inter-agent communication, while DCR offers LLM agents formal tools for coordination, optimization, and structured problem solving. Together, they form the basis of hybrid neurosymbolic systems that combine LLMs' adaptability with DCR's rigor, paving the way for scalable, explainable, and human-aligned multi-agent intelligence.

2 BACKGROUND

Distributed Constraint Reasoning. A *distributed constraint reasoning (DCR)* framework models coordination among multiple autonomous agents, each responsible for one or more variables that must be assigned values from predefined domains under shared constraints. These constraints capture dependencies among agents' decisions, such as avoiding conflicts or satisfying resource limits. When all constraints must be satisfied, the problem is formulated as a *distributed constraint satisfaction problem (DCSP)*; when constraints have associated utilities representing preferences, it becomes a *distributed constraint optimization problem (DCOP)*. In both cases, each agent selects values for its variables and communicates with other agents to jointly determine a complete assignment that maximizes the number of constraints satisfied (in DCSPs) or the total constraint utility (in DCOPs). Communication is typically assumed to be asynchronous and reliable (i.e., messages may be delayed but are never lost) allowing agents to operate concurrently and cooperatively in decentralized environments.

Solving DCOPs optimally is NP-hard [33], and even approximate methods must balance communication, computation, and privacy

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/YNZM5391

trade-offs. Consequently, extensive research has produced a spectrum of algorithms, from inference-based [4, 7, 8, 16, 42, 45, 65], search-based [14, 21, 23, 31, 33, 36, 59, 60, 64], and hybrid inference- and search-based methods [19] to sampling [35, 38] and learning-based approximations [11–13, 34], enabling scalable decision making in dynamic, uncertain, or resource-constrained settings. These models have been widely applied in domains such as meeting scheduling, sensor networks, traffic control, and multi-robot coordination, establishing DCR as a foundational paradigm for distributed problem solving in multi-agent systems.

LLM Agents and Agentic LLMs. The integration of LLMs into autonomous systems has led to the emergence of *LLM agents* – architectures that extend the generative and reasoning capabilities of foundation models with *agentic* functionalities such as planning, memory, tool-use, and action execution. Unlike traditional LLMs that operate as passive, prompt-response systems, LLM agents are designed to autonomously perceive, reason, and act within complex task environments. Architecturally, these agents typically comprise a foundation model (e.g., GPT [25], Qwen [55], DeepSeek [10], and LLaMA [15]), a memory subsystem supporting both short- and long-term context retention, a planning module capable of decomposing goals into subtasks, and an action interface for interacting with external tools, APIs, or environments.

Recent frameworks such as AutoGPT [56], BabyAGI, and CrewAI operationalize these components in feedback-driven loops, enabling agents to iteratively reflect, plan, invoke tools, and revise actions. Such capabilities have made LLM agents applicable across domains including software development, healthcare, scientific research, and content generation. These agents also exhibit key properties associated with cognitive architectures, such as context-aware decision-making, multi-modal integration, and adaptability to dynamic inputs.

Agentic LLMs and Multi-Agent Problem Solving. Recent efforts have begun to apply LLMs to multi-agent decision-making. Sun et al. [47] provide a comprehensive analysis of this burgeoning field, moving beyond individual LLM agent behaviors to address the complexities of collective decision-making among multiple Agentic LLMs. Notably, they map the multi-agent decision-making problem as Dec-POMDPs, where Agentic LLMs are tasked with learning to coordinate and communicate in a fully end-to-end manner. However, the authors emphasize the inherent scalability limitations of this approach, typically restricted to a few agents, and the persistent lack of robust security guarantees. Aside from that, AgentsNet is a benchmark to measure the ability of Agentic LLM systems to self-organize, coordinate, and communicate effectively on fundamental distributed computing problems, such as graph coloring, leader election, consensus [22]. Finally, Agashe et al. [1] evaluate LLM agents in pure coordination games, finding that while they excel in scenarios relying on environmental variables, they face challenges when active consideration of partners' beliefs and intentions (akin to local constraints and their implications) is required.

3 MOTIVATING APPLICATIONS

As a motivating example, consider a multi-agent travel logistics problem in which researchers planning conference attendance are

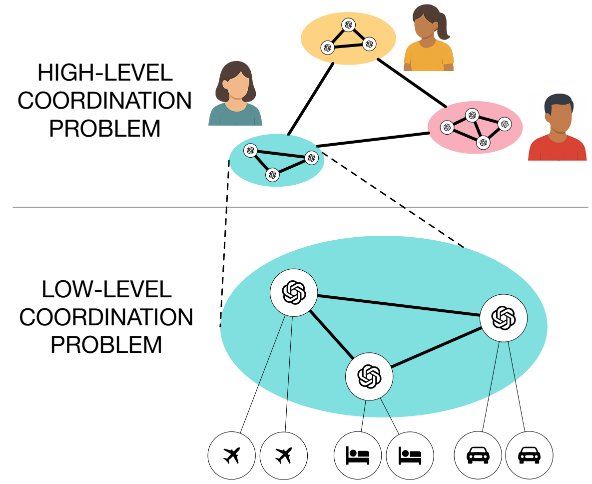

Figure 1: Illustration of our motivating travel logistics planning and coordination problem as a hierarchical problem.

each represented by an agent. High-level coordination is needed to satisfy interdependent preferences (e.g., carpooling to reduce costs or synchronizing arrival times for workshops) naturally giving rise to distributed coordination among agents with overlapping constraints (see Figure 1). Within each researcher's planning process lies another lower level of multi-agent coordination among specialized task agents – one handling flight bookings, another managing accommodations, and so on. These agents must coordinate locally to ensure consistency (e.g., flight and hotel dates match). This example represents a broader class of hierarchical multi-agent coordination problems, where high-level agents (individuals) solve their shared problem in a decentralized manner, preserving privacy over preferences and constraints, while lower-level agents (task agents) cooperate centrally within each individual's domain.

Hierarchical multi-agent coordination arises in many domains. In smart-grid management, high-level agents (households or micro-grids) negotiate consumption and pricing while preserving privacy, and lower-level agents control appliances, batteries and renewables to respect grid limits. In supply-chain logistics, firms coordinate delivery deadlines, inventory costs and carrier capacities, while internal agents handle order processing, warehousing and routing. Autonomous-vehicle fleets and collaborative manufacturing cells similarly negotiate routes, loads or service levels at the top level and rely on internal agents for routing, charging reservations, task assignment and quality control. Disaster-response, healthcare scheduling, and space-mission planning follow the same pattern: Agencies negotiate resource allocation, evacuation routes or launch windows without exposing sensitive data, while internal agents synchronize rescue schedules, appointments or trajectory designs. Multi-robot exploration and urban mobility-as-a-service also fit, with high-level negotiation of traffic, charging infrastructure or congestion pricing and local agents managing navigation, power management or bike-share redistribution.

This hierarchical mix of decentralized *negotiation* and centralized sub-task *coordination* preserves *privacy*, handles *overlapping constraints* and *scales* to complex real-world problems.

4 POTENTIAL SYNERGIES

Given this class of hierarchical multi-agent coordination problems, we discuss several potential synergistic research directions for integrating DCR models and algorithms with Agentic LLM capabilities.

4.1 LLMs for DCR

LLMs for Translating Natural Language into DCR Specifications. Translating natural language into formal symbolic representations (e.g., first-order logic, logic programs, and planning domains) is an active research area in neurosymbolic AI [24, 37, 41, 49, 57]. Despite encouraging results, the process remains challenging. A key difficulty lies in selecting suitable intermediate formalisms bridging natural and symbolic representations [2]. LLMs often produce syntactically valid but semantically inconsistent logical forms (e.g., misusing quantifiers, variable bindings, or scopes) and suffer from limited generalization due to a lack of diverse training data [26]. Approaches such as symbolic validation and human-in-the-loop correction have been explored to refine generated specifications [39]. Relatively little work has focused on translating natural language into DCR specifications [32].

Research Direction 1. *Extending existing neurosymbolic frameworks, such as Logic-LM [39], which supports translation into CSP and SAT formulations, offers a promising direction. Similar constraint-based approaches [24, 41] could be adapted for DCR, especially since symbolic validators already exist for constraint systems, enabling automatic verification and iterative correction of generated models.*

LLMs for User Preference Elicitation and Learning. Beyond translating given specifications, an important challenge is eliciting those specifications: Identifying user constraints and utilities through interaction. Recent work explores LLMs as interactive elicitors that infer user preferences via natural dialogue; e.g., GATE [29] enables LLMs to pose open-ended or hypothetical queries to uncover nuanced, implicit preferences, while MEDIQ [30] demonstrates adaptive questioning in medical decision-making, where the model requests missing information before committing to an answer. This capability aligns closely with Incomplete DCOPs (I-DCOPs) [54], in which some constraint utilities are unknown and must be elicited during problem solving. Prior work has developed strategies to minimize the number of queries while maintaining solution quality [48, 67].

Research Direction 2. *Integrating LLMs as the conversational interface for generating those targeted queries could yield practical, human-interactive DCR systems. Also, rather than eliciting all preferences from scratch, LLMs can learn and refine user models over time using methods such as reinforcement learning from human feedback [9], preference fine-tuning [63], and continual adaptation [62].*

LLMs for Inter-Agent Communication and Information Sharing. While the previous directions use LLMs for DCR specification and symbolic solvers for optimization, a complementary direction is to embed LLMs within the reasoning and communication processes.

Research Direction 3. *A DCR agent's functionality can be divided into two modules: (1) A symbolic optimization module that performs local computations, and (2) an LLM-based communication module that determines what information to request or share with other agents. This leverages LLMs' strengths in language-driven reasoning and negotiation while preserving symbolic rigor for optimization.*

Consider MGM [31], where agents iteratively exchange variable assignments and local utility gains until convergence. Here, the computation of local gains belongs to the symbolic optimization module, while decisions about which neighbors to query or inform could be handled by the LLM-based communication module. Although gains for simple algorithms like MGM may be limited, more complex methods, such as region-optimal approaches [28], may benefit from LLMs' flexibility in dynamically forming regions, proposing coordination heuristics, or adapting communication strategies in heterogeneous, human-in-the-loop environments.

LLMs for Intra-Agent Optimization and Reasoning. Finally, this last research direction is the most ambitious of all.

Research Direction 4. *Instead of using a symbolic module to perform the reasoning process in optimization module above, one can use LLMs instead. For example, many DCR algorithms rely on variable- and value-ordering heuristics that determine the order in which agents evaluate their different variables and values for the solution [17]. LLMs may offer an alternative ordering heuristic that is more tailored to the specific domain and exploits known domain characteristics or biases.*

Instead of using LLMs as subroutines within an optimization module, an alternative is to task an LLM agent with solving the entire optimization problem. Recent work shows that such LLM-based approaches can outperform stochastic DCR algorithms on some benchmarks [32], and that LLM agents benefit from structured symbolic guidance derived from established DCR algorithms [43]. This approach may be more practical in real-world applications, where LLMs can exploit implicitly encoded commonsense knowledge, and can be further improved through fine-tuning on domain-specific training data.

4.2 DCR for Agentic LLMs

DCR for Coordination and Collaboration in Agentic LLMs. When multiple LLM agents are used, coordinating their actions and ensuring they work towards a *common, consistent goal* without conflicts is challenging. Without explicit coordination mechanisms, they might engage in redundant work or generate conflicting outputs. DCR provides formal models for multi-agent coordination and decentralized decision-making. Each LLM could represent an agent responsible for a subset of variables or constraints. DCR algorithms can then guide their interactions, ensuring that their individual assignments or decisions are globally consistent or contribute to a global optimum. This formalizes the communication and negotiation protocols among LLM agents.

DCR is equipped to handle complex problems that necessitate coordination schemes extending beyond simple pairwise interactions, often to achieve a global optimum. For example, some DCR problems involve higher arity constraints, which demand coordinated decision-making among more than two agents to reach an optimal solution. The recent efforts to study the multi-agent problem solving capacity of Agentic LLMs [1, 22] fall within the region-optimal family of approaches [28, 40, 50]. Region-optimal algorithms constitute a class of incomplete (local search) DCOP algorithms that provide solutions with quality guarantees relative to the global optimum. Instead of aiming for a costly global optimum, these methods focus on identifying solutions that are optimal within specific "regions" or "neighborhoods" of agents. Specifically, the aforementioned efforts often correspond to the 1-optimal algorithm family, where solutions align with Nash equilibria [6].

Orthogonally, there is work on decomposing high-arity constraints to binary constraints [20], which Agentic LLMs can easily adapt in their frameworks. A classical example is the all-different constraint, which imposes all variables to take on different values, can be decomposed into pairwise binary not-equal constraints between all variables. Additionally, DCR researchers have also proposed methods to decompose a problem into subproblems, for example, which agent should be responsible for which variables in order to minimize coordination overhead or to speed up the overall resolution process [5, 18, 27]. These methods can also be adapted for Agentic LLM systems to determine how subproblems should be scoped out and solved by an individual LLM agent such that, collectively, they solve the overall problem.

Research Direction 5. *Future research should focus on (i) studying the ability of Agentic LLMs to effectively leverage and adapt to richer coordination schemes in DCR, moving beyond 1-optimal solutions to achieve more general region-optimal outcomes, and (ii) exploring the definition and implementation of advanced coordination schemes for Agentic LLMs that are capable of reaching global optima or providing provable guarantees on solution quality.*

DCR for Consistent and Structured Reasoning in Multi-Step Problem Solving. While LLMs show impressive *chain-of-thought* capabilities [52], their reasoning is often heuristic, implicit, and can break down on complex, multi-step problems requiring precise logical deduction or combinatorial search. They might struggle with complex interdependencies. DCR excels at breaking down complex problems into smaller, manageable subproblems, where agents (or LLMs acting as agents) work cooperatively to find a globally consistent or optimal solution. The explicit definition of variables, domains, and constraints in DCR provides a structured backbone for the LLM's reasoning process. The LLM can then focus on generating plausible values or identifying relevant constraints, with the DCR framework ensuring the overall logical coherence and completeness. DCR would enable Agentic LLMs to tackle problems that require formal, systematic, and exhaustive search, which is beyond their native capabilities, by embedding them within a rigorous reasoning framework.

Research Direction 6. *While* tree-of-thought *methods can enhance reasoning in complex question-answering tasks [58], a promising avenue is to leverage classical DCR coordination protocols as the underlying structure for a more sophisticated* network-of-thoughts. *In this paradigm, the graphical models that form the basis of most DCR solution methods (e.g., pseudo-trees, factor graphs) could serve as the network's organizational structure [17].*

DCR for Agentic LLM Resource Management and Optimization. LLMs themselves consume significant computational resources (inference cost, memory). In distributed scenarios, optimizing resource allocation among multiple LLM instances or other computational agents is critical. Wilkins et al. [53] provides a foundational understanding for sustainable LLM deployment and offers insights for designing energy-efficient AI infrastructure in the future. It highlights the importance of considering workload characteristics and system configurations for accurate energy estimation and effective optimization. DCOPs are well suited for distributed resource allocation and optimization where agents have limited local knowledge and must coordinate to optimize a global objective.

Research Direction 7. *LLMs could use DCOP principles to decide how to allocate computational budget, which knowledge sources to retrieve from, or which sub-tasks to take on, based on explicit constraints, as proposed by Rust et al. [46] in the context of the Internet-of-Things. This would lead to more efficient and cost-effective deployment of LLMs, especially in large-scale or real-time distributed applications.*

DCR for Communication-Sensitive Agentic LLMs. Like early work in DCR, existing Agentic LLM frameworks typically make naive assumptions on the communication between agents (e.g., they typically do not model communication costs and assume that communication is perfect). Recent work on communication-aware DCR has found that the performance and properties of popular DCR algorithms can vary significantly in the presence of communication latency and message loss [44, 66].

Research Direction 8. *Communication-aware DCR frameworks and designated algorithms can be used to inform the design of Agentic LLM systems to operate more robustly in communication-sensitive applications.*

5 CONCLUSIONS

DCR and Agentic LLMs represent complementary paradigms for multi-agent intelligence. DCR offers mathematically grounded frameworks for coordination, optimization, and decision-making, while Agentic LLMs provide powerful capabilities for natural language understanding, flexible reasoning, and human interaction. We argued that the time is ripe to bring these traditions together. Accordingly, we outlined several synergistic directions: Using LLMs to translate human intent into DCR specifications, to elicit and refine user preferences, and to support inter-agent communication; and, conversely, using DCR algorithms to provide Agentic LLMs with principled methods for coordination, structured reasoning, resource allocation, and communication-sensitive operation.

Taken together, these research threads highlight the potential for hybrid systems that combine the adaptability and expressiveness of LLMs with the rigor and guarantees of DCR. Beyond technical synergies, this integration addresses long-standing challenges in both communities: Grounding DCR models in real-world human input and augmenting LLM-based agents with formal mechanisms for consistency, efficiency, and scalability.

We therefore issue a call to the research community: To systematically investigate these intersections, to develop benchmarks and evaluation methodologies that capture both symbolic rigor and language-based adaptability, and to design architectures that are not only technically powerful but also explainable, robust, and aligned with human needs. By jointly advancing DCR and Agentic LLM research, we can lay the foundation for the next generation of multi-agent systems that is capable of operating in complex, distributed, and human-centric environments.

ACKNOWLEDGMENTS

This research is partially supported by the US-Israel Binational Science Foundation (BSF) under grant 2022189 and the US Office of Naval Research (ONR) under grant N00014-24-1-2663.

REFERENCES

[1] Saaket Agashe, Yue Fan, Anthony Reyna, and Xin Eric Wang. 2025. LLM-Coordination: Evaluating and Analyzing Multi-agent Coordination Abilities in Large Language Models. In *Findings of the Association for Computational Linguistics (NAACL)*.

[2] Alexander Beiser and David Penz. 2025. Making LLMs Reason? The Intermediate Language Problem in Neurosymbolic Approaches. *arXiv preprint arXiv:2502.17216* (2025).

[3] Maciej Besta, Julia Barth, Eric Schreiber, Ales Kubicek, Afonso Claudino Catarino, Robert Gerstenberger, Piotr Nyczyk, Patrick Iff, Yueling Li, Sam Houliston, Tomasz Sternal, Marcin Copik, Grzegorz Kwasniewski, Jürgen Müller, Lukasz Flis, Hannes Eberhard, Hubert Niewiadomski, and Torsten Hoefler. 2025. Reasoning Language Models: A Blueprint. *arXiv preprint arXiv:2501.11223* (2025).

[4] Ismel Brito and Pedro Meseguer. 2010. Improving DPOP with function filtering. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 141–148.

[5] David A. Burke and Kenneth N. Brown. 2006. Efficient Handling of Complex Local Problems in Distributed Constraint Optimization. In *Proceedings of the European Conference on Artificial Intelligence (ECAI)*. 701–702.

[6] A. C. Chapman, A. Rogers, N. R. Jennings, and D. S. Leslie. 2011. A unifying framework for iterative approximate best-response algorithms for distributed constraint optimization problems. *Knowledge Engineering Review* 26, 4 (2011), 411–444.

[7] Dingding Chen, Ziyu Chen, Yanchen Deng, Zhongshi He, and Lulu Wang. 2023. Inference-based complete algorithms for asymmetric distributed constraint optimization problems. *Artificial Intelligence Review* 56, 5 (2023), 4491–4534.

[8] Ziyu Chen, Yanchen Deng, Tengfei Wu, and Zhongshi He. 2018. A class of iterative refined Max-sum algorithms via non-consecutive value propagation strategies. *Autonomous Agents and Multi-Agent Systems* 32, 6 (2018), 822–860.

[9] Paul F. Christiano, Jan Leike, Tom B. Brown, Miljan Martic, Shane Legg, and Dario Amodei. 2017. Deep reinforcement learning from human preferences. In *Proceedings of the International Conference on Neural Information Processing Systems (NIPS)*. 4302–4310.

[10] DeepSeek-AI, Aixin Liu, Bei Feng, Bing Xue, et al. 2024. DeepSeek-V3 Technical Report. *arXiv preprint arXiv:2412.19437* (2024), 53.

[11] Yanchen Deng and Bo An. 2020. Speeding Up Incomplete GDL-based Algorithms for Multi-agent Optimization with Dense Local Utilities. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. 31–38.

[12] Yanchen Deng, Shufeng Kong, and Bo An. 2022. Pretrained Cost Model for Distributed Constraint Optimization Problems. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. 9331–9340.

[13] Yanchen Deng, Shufeng Kong, Caihua Liu, and Bo An. 2022. Deep Attentive Belief Propagation: Integrating Reasoning and Learning for Solving Constraint Optimization Problems. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*.

[14] Yanchen Deng, Xinrun Wang, and Bo An. 2025. GDBA Revisited: Unleashing the Power of Guided Local Search for Distributed Constraint Optimization. *arXiv preprint arXiv:2508.06899* (2025).

[15] Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, et al. 2024. The Llama 3 Herd of Models. *arXiv preprint arXiv:2407.21783* (2024).

[16] A. Farinelli, A. Rogers, A. Petcu, and N. Jennings. 2008. Decentralised Coordination of Low-Power Embedded Devices Using the Max-Sum Algorithm. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 639–646.

[17] 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.

[18] Ferdinando Fioretto, William Yeoh, and Enrico Pontelli. 2016. Multi-Variable Agents Decomposition for DCOPs. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. 2480–2486.

[19] Junsong Gao, Ziyu Chen, Dingding Chen, Wenxin Zhang, and Qiang Li. 2024. Toward fast belief propagation for distributed constraint optimization problems via heuristic search. *Autonomous Agents and Multi-Agent Systems* 38, 1 (2024), 15.

[20] Ian P. Gent, Kostas Stergiou, and Toby Walsh. 2000. Decomposable constraints. *Artificial Intelligence* 123, 1-2 (2000), 133–156.

[21] A. Gershman, A. Meisels, and R. Zivan. 2009. Asynchronous Forward-Bounding for Distributed COPs. *Journal of Artificial Intelligence Research* 34 (2009), 61–88.

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

[23] K. Hirayama and M. Yokoo. 2005. The Distributed Breakout Algorithms. *Artificial Intelligence* 161, 1–2 (2005), 89–115.

[24] Lothar Hotz, Christian Bähnisch, Sebastian Lubos, Alexander Felfernig, Albert Haag, and Johannes Twiefel. 2024. Exploiting Large Language Models for the Automated Generation of Constraint Satisfaction Problems. In *Proceedings of the International Workshop on Configuration (ConfWS)*. 91–100.

[25] Aaron Jaech, Adam Kalai, Adam Lerer, Adam Richardson, et al. 2024. OpenAI o1 System Card. *arXiv preprint arXiv:2412.16720* (2024), 52.

[26] Rushang Karia, Daniel Bramblett, Daksh Dobhal, and Siddharth Srivastava. 2025. Autonomous Evaluation of LLMs for Truth Maintenance and Reasoning Tasks. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

[27] Md. Mosaddek Khan, Long Tran-Thanh, William Yeoh, and Nicholas R. Jennings. 2018. A Near-Optimal Node-to-Agent Mapping Heuristic for GDL-Based DCOP Algorithms in Multi-Agent Systems. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1604–1612.

[28] C. Kiekintveld, Z. Yin, A. Kumar, and M. Tambe. 2010. Asynchronous Algorithms for Approximate Distributed Constraint Optimization with Quality Bounds. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 133–140.

[29] Belinda Z. Li, Alex Tamkin, Noah Goodman, and Jacob Andreas. 2025. Eliciting Human Preferences with Language Models. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

[30] Shuyue Stella Li, Vidhisha Balachandran, Shangbin Feng, Jonathan S. Ilgen, Emma Pierson, Pang Wei Koh, and Yulia Tsvetkov. 2024. MediQ: Question-Asking LLMs and a Benchmark for Reliable Interactive Clinical Reasoning. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*.

[31] R. Maheswaran, J. Pearce, and M. Tambe. 2004. Distributed Algorithms for DCOP: A Graphical Game-Based Approach. In *Proceedings of the International Conference on Parallel and Distributed Computing Systems (PDCS)*. 432–439.

[32] Saaduddin Mahmud, Dorian Benhamou Goldfajn, and Shlomo Zilberstein. 2025. Distributed Multi-Agent Coordination Using Multi-Modal Foundation Models. *arXiv preprint arXiv:2501.14189* (2025).

[33] P. J. Modi, W. Shen, M. Tambe, and M. Yokoo. 2005. ADOPT: Asynchronous Distributed Constraint Optimization with Quality Guarantees. *Artificial Intelligence* 161, 1–2 (2005), 149–180.

[34] Duc Thien Nguyen, William Yeoh, Hoong Chuin Lau, Shlomo Zilberstein, and Chongjie Zhang. 2014. Decentralized Multi-Agent Reinforcement Learning in Average-Reward Dynamic DCOPs. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. 1447–1455.

[35] 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.

[36] Steven Okamoto, Roie Zivan, and Aviv Nahon. 2016. Distributed Breakout: Beyond Satisfaction. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. 447–453.

[37] James Oswald, Kavitha Srinivas, Harsha Kokel, Junkyu Lee, Michael Katz, and Shirin Sohrabi. 2024. Large language models as planning domain generators. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*.

[38] Brammert Ottens, Christos Dimitrakakis, and Boi Faltings. 2017. DUCT: An Upper Confidence Bound Approach to Distributed Constraint Optimization Problems. *ACM Transactions on Intelligent Systems and Technology* 8, 5 (2017), 69:1–69:27.

[39] Liangming Pan, Alon Albalak, Xinyi Wang, and William Wang. 2023. Logic-LM: Empowering Large Language Models with Symbolic Solvers for Faithful Logical Reasoning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*. 3806–3824.

[40] Jonathan Pearce and Milind Tambe. 2007. Quality Guarantees on k-Optimal Solutions for Distributed Constraint Optimization Problems. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. 1446–1451.

[41] Roberto Penco, Damir Pintar, Mihaela Vranić, and Marko Šoštarić. 2025. Large Language Model-Driven Framework for Automated Constraint Model Generation in Configuration Problems. *Applied Sciences* 15, 12 (2025). https://doi.org/10.3390/app15126518

[42] A. Petcu and B. Faltings. 2005. A Scalable Method for Multiagent Constraint Optimization. In *Proceedings of the International Joint Conference on Artificial Intelligence (IJCAI)*. 1413–1420.

[43] Ben Rachmut, Ning Zhang, Yevgeniy Vorobeychik, and William Yeoh. 2026. Symbolic Guidance for LLM Agents in Distributed Multiagent Coordination. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.

[44] Ben Rachmut, Roie Zivan, and William Yeoh. 2022. Communication-Aware Local Search for Distributed Constraint Optimization. *Journal of Artificial Intelligence Research* 75 (2022), 637–675.

[45] A. Rogers, A. Farinelli, R. Stranders, and N. Jennings. 2011. Bounded Approximate Decentralised Coordination via the Max-Sum Algorithm. *Artificial Intelligence* 175, 2 (2011), 730–759.

[46] 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, Article 100 (2022), 31 pages.

[47] Chuanneng Sun, Songjun Huang, and Dario Pompili. 2025. LLM-Based Multi-Agent Decision-Making: Challenges and Future Directions. *IEEE Robotics and Automation Letters* 10, 6 (2025), 5681–5688.

[48] Atena M. Tabakhi, Yuanming Xiao, William Yeoh, and Roie Zivan. 2021. Branch-and-Bound Heuristics for Incomplete DCOPs. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1677–1679.

[49] Paul Tarau. 2024. On LLM-generated Logic Programs and their Inference Execution Methods. In *Proceedings of the International Conference on Logic Programming (ICLP)*. 1–14.

[50] M. Vinyals, E. Shieh, J. Cerquides, J. Rodriguez-Aguilar, Z. Yin, M. Tambe, and E. Bowring. 2011. Quality Guarantees for Region Optimal DCOP algorithms. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 133–140.

[51] Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. Finetuned Language Models are Zero-Shot Learners. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

[52] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. 2022. Chain-of-thought prompting elicits reasoning in large language models. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*.

[53] Grant Wilkins, Srinivasan Keshav, and Richard Mortier. 2025. Offline Energy-Optimal LLM Serving: Workload-Based Energy Models for LLM Inference on Heterogeneous Systems. *SIGENERGY Energy Informatics Review* 4, 5 (2025), 113–119.

[54] Yuanming Xiao, Atena M. Tabakhi, and William Yeoh. 2020. Embedding Preference Elicitation Within the Search for DCOP Solutions. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 2044–2046.

[55] Jin Xu, Zhifang Guo, Jinzheng He, Hangrui Hu, Ting He, Shuai Bai, Keqin Chen, Jialin Wang, Yang Fan, Kai Dang, Bin Zhang, Xiong Wang, Yunfei Chu, and Junyang Lin. 2025. Qwen2.5-Omni Technical Report. *arXiv preprint arXiv:2503.20215* (2025).

[56] Hui Yang, Sifu Yue, and Yunzhong He. 2023. Auto-GPT for Online Decision Making: Benchmarks and Additional Opinions. *arXiv preprint arXiv:2306.02224* (2023).

[57] Yuan Yang, Siheng Xiong, Ali Payani, Ehsan Shareghi, and Faramarz Fekri. 2024. Harnessing the Power of Large Language Models for Natural Language to First-Order Logic Translation. In *Proceedings of the Annual Meeting of the Association for Computational Linguistics (ACL)*. 6942–6959.

[58] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L. Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: deliberate problem solving with large language models. In *Proceedings of the Conference on Neural Information Processing Systems (NeurIPS)*.

[59] W. Yeoh, A. Felner, and S. Koenig. 2010. BnB-ADOPT: An Asynchronous Branch-and-Bound DCOP Algorithm. *Journal of Artificial Intelligence Research* 38 (2010), 85–133.

[60] W. Yeoh, P. Varakantham, X. Sun, and S. Koenig. 2011. Incremental DCOP Search Algorithms for Solving Dynamic DCOPs. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 1069–1070.

[61] Makoto Yokoo, Edmund H. Durfee, Toru Ishida, and Kazuhiro Kuwabara. 1998. The Distributed Constraint Satisfaction Problem: Formalization and Algorithms. *IEEE Transactions on Knowledge and Data Engineering* 10, 5 (1998), 673–685.

[62] Han Zhang, Yu Lei, Lin Gui, Min Yang, Yulan He, Hui Wang, and Ruifeng Xu. 2024. CPPO: Continual Learning for Reinforcement Learning with Human Feedback. In *Proceedings of the International Conference on Learning Representations (ICLR)*.

[63] Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul F. Christiano, and Geoffrey Irving. 2019. Fine-Tuning Language Models from Human Preferences. *arXiv preprint arXiv:1909.08593* (2019).

[64] R. Zivan. 2008. Anytime Local Search for Distributed Constraint Optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence (AAAI)*. 393–398.

[65] R. Zivan and H. Peled. 2012. Max/Min-Sum Distributed Constraint Optimization through Value Propagation on an Alternating DAG. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. 265–272.

[66] Roie Zivan, Ben Rachmut, Omer Perry, and William Yeoh. 2023. Effect of asynchronous execution and imperfect communication on max-sum belief propagation. *Autonomous Agents and Multi-Agent Systems* 37, 2 (2023), 40.

[67] Roie Zivan, Shiraz Regev, and William Yeoh. 2024. Ex-Ante Constraint Elicitation in Incomplete DCOPs. In *Proceedings of the International Conference on Principles and Practice of Constraint Programming (CP)*. 33:1–33:16.