

# ROUTERHGC: Optimized Router for LLM-based Multi-Agent Systems via Heterogeneous Graph Contrastive Learning

Extended Abstract

Yitao Xiao  
Beijing University of Posts and  
Telecommunications  
Beijing, China  
yitoooom@bupt.edu.cn

Shaoyong Guo\*  
Beijing University of Posts and  
Telecommunications  
Beijing, China  
syguo@bupt.edu.cn

Guoming Yang  
Beijing University of Posts and  
Telecommunications  
Beijing, China  
gmyang@bupt.edu.cn

Qingnan Wang  
Beijing University of Posts and  
Telecommunications  
Beijing, China  
wqn@bupt.edu.cn

Yinlin Ren  
Beijing University of Posts and  
Telecommunications  
Beijing, China  
yinlinren@bupt.edu.cn

Xuesong Qiu  
Beijing University of Posts and  
Telecommunications  
Beijing, China  
xsqiu@bupt.edu.cn

Feng Qi  
Beijing University of Posts and  
Telecommunications  
Beijing, China  
qifeng@bupt.edu.cn

## ABSTRACT

While Large Language Model (LLM)-based Multi-Agent Systems (MAS) offer scalable solutions for complex tasks, dynamically configuring the optimal combination of collaboration modes, agent roles, and LLMs remains challenging. Existing routing methods typically resort to computationally expensive post-response optimization or focus on single-agent settings. We introduce ROUTERHGC, the first framework to harness heterogeneous graph contrastive learning for MAS routing. We formalize routing as a node selection problem on a heterogeneous graph, capturing high-order dependencies among queries, collaboration patterns, and roles. To optimize this, we design a novel global-local contrastive loss that aligns query graphs with high-performing configurations while distancing them from suboptimal or costly ones. Experiments on five datasets demonstrate that ROUTERHGC significantly improves reasoning accuracy while reducing costs.

## KEYWORDS

MAS; LLM; Agent Routing; Graph Contrastive Learning

### ACM Reference Format:

Yitao Xiao, Shaoyong Guo, Guoming Yang, Qingnan Wang, Yinlin Ren, Xuesong Qiu, and Feng Qi. 2026. ROUTERHGC: Optimized Router for LLM-based Multi-Agent Systems via Heterogeneous Graph Contrastive Learning. Extended Abstract. In *Proc. of the 25th International Conference on*

*Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 4 pages. <https://doi.org/10.65109/PGW17290>

## 1 INTRODUCTION

Large Language Models (LLMs) have evolved into autonomous agents capable of handling complex reasoning tasks [15, 20, 54] through independent decision-making [10, 39, 41]. As the Agentic Web continues to expand [44], Multi-Agent Systems (MAS) organize agents with distinct roles and collaboration modes (e.g., Chain [55], Tree [59], Debate [37]) to yield emergent intelligence that surpasses single models [11, 45]. However, the capabilities of different LLMs vary significantly [16], and various agent collaboration modes also exhibit marked differences in cost and task proficiency [28, 29, 38]. The core challenge for MAS providers is delivering query-driven, cost-sensitive, and domain-adaptive services [12, 25].

Existing approaches are categorized into post-response optimization and pre-response routing. Post-response optimization [21, 52] adjusts agent contributions over multiple rounds but incurs high computational cost and latency. Pre-response routing [4, 7, 33] uses lightweight routers to predict LLM suitability, reducing latency and cost. However, current routing research predominantly focuses on single-agent scenarios [1, 26, 31, 35], overlooking the determination of collaboration patterns [5] and role [27] assignments in MAS. The only existing MAS routing scheme [50] employs cascaded controllers that often fail to capture high-order contextual dependencies [51], leading to suboptimal decisions. Furthermore, we observe that the widely used KL-divergence reward mechanism [22, 34] lacks discriminative power due to Softmax uniformity. As a representative self-supervised method, graph contrastive learning centers on pulling similar samples closer and pushing dissimilar samples apart, encouraging the model to learn representations that effectively capture graph-structural features and

\*Corresponding author



This work is licensed under a Creative Commons Attribution International 4.0 License.

thereby improving the robustness and generalization of graph representations [58]. Thus far, graph contrastive learning has been applied in recommender systems [2, 13, 32, 43, 46, 48], social network analysis [23, 24, 36, 57], and genomics-related tasks [6, 9, 18, 42].

Inspired by this line of work, we propose ROUTERHGC, a heterogeneous graph-based framework that leverages a global-local contrastive loss to distinctively separate optimal from suboptimal MAS configurations in the representation space.

## 2 METHOD

**Heterogeneous Graph Modeling:** To capture the intricate dependencies within a MAS, we formalize the routing task as a hierarchical node selection problem on a heterogeneous graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ . The node set  $\mathcal{V}$  comprises four distinct entities: User Queries ( $V_q$ ), Collaboration Modes ( $V_p$ ) (e.g., CoT, Debate), Agent Roles ( $V_r$ ) (e.g., Coder, Reviewer), and LLM Backbones ( $V_m$ ). Instead of treating these components as independent variables, we employ a Heterogeneous Graph Neural Network (HGNN) to perform message passing, explicitly modeling their high-order interactions. The decision-making process is instantiated as a path traversal starting from  $V_q$ : the system first predicts the optimal collaboration mode via edges ( $V_q, V_p$ ), identifies necessary roles via ( $V_p, V_r$ ), and finally maps each role to a cost-effective LLM via ( $V_r, V_m$ ). This graph-structured formulation enables the router to construct an effective agent team by capturing essential structural and semantic relationships.

**Global-Local Contrastive Learning:** To address weak supervision signals in standard routing, we design a contrastive mechanism guided by a joint performance-cost reward function:

$$\text{Reward} = \alpha \cdot \text{Performance} - \beta \cdot \text{Cost}. \quad (1)$$

At the macro level, the training process utilizes an extended InfoNCE loss [30], we design a Global Graph Contrastive Loss to optimize the representation of the entire routing graph:

$$\mathcal{L}_C = -\frac{1}{|B|} \sum_{G \in B} \log \frac{\sum_{G_p \in G_P} \exp(\text{Sim}_p/\tau)}{\sum_{G_n \in G_N} \exp(\text{Sim}_n/\tau)}, \quad (2)$$

where  $B$  denotes the batch of graphs,  $G_P \in \{G_s, G_c\}$  is the positive sample,  $G_N$  represents negative samples,  $\text{Sim}_{n/p} = \text{sim}(z_G, z_{G_n/p})$  computes cosine similarity between graph embeddings, and  $\tau$  is the temperature parameter. Specifically, we define positive samples as the routing configurations that yield the highest joint rewards for the current query, as well as the optimal graphs derived from semantically similar queries in the embedding space. In contrast, negative samples are selected from configurations that exhibit poor reasoning accuracy or incur excessive computational costs relative to their performance. By minimizing the distance to these high-quality positives and maximizing the distance from the suboptimal negatives, the router learns to align the query representation with cost-effective and high-performing team structures. At the micro level, we incorporate a Local Edge-Weighted Loss to calibrate specific edge predictions, ensuring precise role assignment:

$$\mathcal{L}_e = \frac{1}{|\mathcal{E}|} \sum_{(u,v) \in \mathcal{E}} (p_{uv} - y_{uv})^2, \quad (3)$$

where  $p_{uv}$  represents the predicted probability of selecting edge  $(u, v)$  (e.g., connecting a specific role to an LLM) generated by the HGNN, and  $y_{uv} \in \{0, 1\}$  is the ground-truth label indicating

**Table 1: Performance comparison across multiple benchmarks. Best: bold, second-best: underlined.**

Method	CMMLU	GSM8K	MATH	HumanEval	HotpotQA
Vanilla	78.98	90.14	68.31	79.27	73.30
CoT [56]	79.80	89.76	68.52	79.62	73.83
Complex-CoT [8]	80.29	89.51	69.41	80.24	74.08
SC [40]	80.13	90.33	68.39	80.86	74.04
Debate [19]	81.40	94.34	61.48	89.02	75.07
AFlow [53]	79.40	94.88	65.84	85.50	76.79
DyLAN [21]	81.99	95.15	<u>71.75</u>	86.42	76.97
RouterDC [4]	82.39	92.62	68.70	80.75	70.68
MasRouter [50]	<b>83.18</b>	<u>95.26</u>	70.51	<u>89.68</u>	<u>78.02</u>
<b>ROUTERHGC</b>	<u>82.93</u>	<b>96.37</b>	<b>72.15</b>	<b>90.53</b>	<b>78.88</b>

whether the edge exists in the optimal routing path. The final objective  $\mathcal{L} = \lambda_1 \mathcal{L}_c + \lambda_2 \mathcal{L}_e$  jointly optimizes global structural discrimination and local selection precision.

## 3 EXPERIMENTS

As presented in Table 1, ROUTERHGC achieves the highest average accuracy (84.17%) across the CMMLU [17], GSM8K [49], MATH [14], HumanEval [3], and HotpotQA [47] datasets, consistently outperforming all baselines. Notably, our framework surpasses the state-of-the-art MasRouter, exemplified by a 1.64% improvement on the complex MATH dataset. Crucially, this performance advantage is accompanied by significant cost savings; compared to MasRouter, ROUTERHGC reduces inference costs on HotpotQA by 27.4%. These results validate that our global-local contrastive learning effectively optimizes the accuracy-budget trade-off. By strategically allocating cost-effective models, our approach excels in both effectiveness and economy, surpassing single-agent baselines, post-response optimization strategies, and mainstream pre-response routing methods.

## 4 CONCLUSION

This paper proposes ROUTERHGC, pioneering the application of graph contrastive learning to multi-agent routing. By leveraging a heterogeneous graph neural network, ROUTERHGC effectively captures the high-order dependencies among queries, collaboration patterns, roles, and LLMs. A custom-designed global-local contrastive loss function maximizes the distinctions between multi-agent routing options, thereby recommending the most performant and cost-effective multi-agent service configuration for each specific user query. Experiments across five public datasets demonstrate that ROUTERHGC surpasses existing baseline methods in both performance and cost. Future work will focus on exploring the scalability and adaptability of this multi-agent routing service in edge computing environments.

## ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (62322103) and Beijing-Tianjin-Hebei Natural Science Foundation Cooperation Project (Grant No. 25JJJC0049).

## REFERENCES

- [1] Lingjiao Chen, Matei Zaharia, and James Zou. 2023. Frugalgpt: How to use large language models while reducing cost and improving performance. *arXiv preprint arXiv:2305.05176* (2023).
- [2] Mengru Chen, Chao Huang, Lianghao Xia, Wei Wei, Yong Xu, and Ronghua Luo. 2023. Heterogeneous graph contrastive learning for recommendation. In *Proceedings of the sixteenth ACM international conference on web search and data mining*. 544–552.
- [3] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgren Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. Evaluating Large Language Models Trained on Code. (2021). arXiv:2107.03374 [cs.LG]
- [4] Shuhao Chen, Weisen Jiang, Baijiong Lin, James Kwok, and Yu Zhang. 2024. Routerdc: Query-based router by dual contrastive learning for assembling large language models. *Advances in Neural Information Processing Systems* 37 (2024), 66305–66328.
- [5] Dujan Ding, Ankur Mallick, Chi Wang, Robert Sim, Subhabrata Mukherjee, Victor Rühle, Laks VS Lakshmanan, and Ahmed Hassan Awadallah. 2024. Hybrid llm: Cost-efficient and quality-aware query routing. *arXiv preprint arXiv:2404.14618* (2024).
- [6] Yanhao Fan, Che Zhang, Xiaowen Hu, Zhijian Huang, Jiameng Xue, and Lei Deng. 2024. SGCLDGA: unveiling drug–gene associations through simple graph contrastive learning. *Briefings in Bioinformatics* 25, 3 (2024), bbae231.
- [7] Tao Feng, Yanzhen Shen, and Jiaxuan You. 2024. Graphrouter: A graph-based router for llm selections. *arXiv preprint arXiv:2410.03834* (2024).
- [8] Yao Fu, Hao Peng, Ashish Sabharwal, Peter Clark, and Tushar Khot. 2023. COMPLEXITY-BASED PROMPTING FOR MULTI-STEP REASONING. In *11th International Conference on Learning Representations, ICLR 2023*.
- [9] Zihao Gao, Huifang Ma, Xiaohui Zhang, Yike Wang, and Zheyu Wu. 2023. Similarity measures-based graph co-contrastive learning for drug–disease association prediction. *Bioinformatics* 39, 6 (2023), btad357.
- [10] Shivam Goel. 2024. Towards building Autonomous AI Agents and Robots for Open World Environments.. In *AAMAS*. 2743–2745.
- [11] Taicheng Guo, Xiuying Chen, Yaqi Wang, Ruidi Chang, Shichao Pei, Nitesh V Chawla, Olaf Wiest, and Xiangliang Zhang. 2024. Large language model based multi-agents: A survey of progress and challenges. *arXiv preprint arXiv:2402.01680* (2024).
- [12] Lewis Hammond, Alan Chan, Jesse Clifton, Jason Hoelscher-Obermaier, Akbir Khan, Euan McLean, Chandler Smith, Wolfram Barfuss, Jakob Foerster, Tomáš Gavenčák, et al. 2025. Multi-agent risks from advanced ai. *arXiv preprint arXiv:2502.14143* (2025).
- [13] Wei He, Guohao Sun, Jinhu Lu, and Xiu Susie Fang. 2023. Candidate-aware graph contrastive learning for recommendation. In *Proceedings of the 46th international ACM SIGIR conference on research and development in information retrieval*. 1670–1679.
- [14] Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring Mathematical Problem Solving With the MATH Dataset. *NeurIPS* (2021).
- [15] Bin Lei, Yi Zhang, Shan Zuo, Ali Payani, and Caiwen Ding. 2024. Macm: Utilizing a multi-agent system for condition mining in solving complex mathematical problems. *Advances in Neural Information Processing Systems* 37 (2024), 53418–53437.
- [16] Pierre Lepagnol, Thomas Gerald, Sahar Ghannay, Christophe Servan, and Sophie Rosset. 2024. Small language models are good too: An empirical study of zero-shot classification. *arXiv preprint arXiv:2404.11122* (2024).
- [17] Haonan Li, Yixuan Zhang, Fajri Koto, Yifei Yang, Hai Zhao, Yeyun Gong, Nan Duan, and Timothy Baldwin. 2023. CMMLU: Measuring massive multitask language understanding in Chinese. arXiv:2306.09212 [cs.CL]
- [18] Shuangli Li, Jingbo Zhou, Tong Xu, Dejing Dou, and Hui Xiong. 2022. Geomgl: Geometric graph contrastive learning for molecular property prediction. In *Proceedings of the AAAI conference on artificial intelligence*, Vol. 36. 4541–4549.
- [19] Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2024. Encouraging Divergent Thinking in Large Language Models through Multi-Agent Debate. In *EMNLP*.
- [20] Feng Lin, Dong Jae Kim, et al. 2024. Soen-101: Code generation by emulating software process models using large language model agents. *arXiv preprint arXiv:2403.15852* (2024).
- [21] Zijun Liu, Yanzhe Zhang, Peng Li, Yang Liu, and Diyi Yang. 2024. A dynamic llm-powered agent network for task-oriented agent collaboration. In *First Conference on Language Modeling*.
- [22] Keming Lu, Hongyi Yuan, Runji Lin, Junyang Lin, Zheng Yuan, Chang Zhou, and Jingren Zhou. 2023. Routing to the expert: Efficient reward-guided ensemble of large language models. *arXiv preprint arXiv:2311.08692* (2023).
- [23] Guanghui Ma, Chunming Hu, Ling Ge, Junfan Chen, Hong Zhang, and Richong Zhang. 2022. Towards robust false information detection on social networks with contrastive learning. In *Proceedings of the 31st ACM international conference on information & knowledge management*. 1441–1450.
- [24] Guanghui Ma, Chunming Hu, Ling Ge, and Hong Zhang. 2022. Open-topic false information detection on social networks with contrastive adversarial learning. In *Proceedings of the 2022 conference on empirical methods in natural language processing*. 2911–2923.
- [25] Kai Mei, Wujiang Xu, Shuhang Lin, and Yongfeng Zhang. 2025. OmniRouter: Budget and Performance Controllable Multi-LLM Routing. *arXiv preprint arXiv:2502.20576* (2025).
- [26] Alireza Mohammadshahi, Arshad Rafiq Shaikh, and Majid Yazdani. 2024. Routoo: Learning to route to large language models effectively. *arXiv preprint arXiv:2401.13979* (2024).
- [27] Isaac Ong, Amjad Almahairi, Vincent Wu, Wei-Lin Chiang, Tianhao Wu, Joseph E Gonzalez, M Waleed Kadous, and Ion Stoica. 2024. Routellm: Learning to route llms with preference data. *arXiv preprint arXiv:2406.18665* (2024).
- [28] Yinghui Pan, Fanke Chen, Biyang Ma, Yifeng Zeng, and Prashant Doshi. 2025. Model-based Individual Learning for Competitive Agents. (2025).
- [29] David Radke, Kate Larson, Tim Brecht, and Kyle Tilbury. 2023. Towards a better understanding of learning with multiagent teams. *arXiv preprint arXiv:2306.16205* (2023).
- [30] Evgenia Rusak, Patrik Reizinger, Attila Juhas, Oliver Bringmann, Roland S. Zimmermann, and Wieland Brendel. 2024. InfoNCE: Identifying the Gap Between Theory and Practice. In *ICML 2024 Workshop on Geometry-grounded Representation Learning and Generative Modeling*. <https://openreview.net/forum?id=sYoGm14cle>
- [31] Tal Shnitzer, Anthony Ou, Mírian Silva, Kate Soule, Yuekai Sun, Justin Solomon, Neil Thompson, and Mikhail Yurochkin. 2023. Large language model routing with benchmark datasets. *arXiv preprint arXiv:2309.15789* (2023).
- [32] Jie Shuai, Kun Zhang, Le Wu, Peijie Sun, Richang Hong, Meng Wang, and Yong Li. 2022. A review-aware graph contrastive learning framework for recommendation. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*. 1283–1293.
- [33] Wei Song, Zhenya Huang, Cheng Cheng, Weibo Gao, Bihan Xu, GuanHao Zhao, Fei Wang, and Runze Wu. 2025. IRT-Router: Effective and Interpretable Multi-LLM Routing via Item Response Theory. *arXiv preprint arXiv:2506.01048* (2025).
- [34] KV Srivatsa, Kaushal Kumar Maurya, and Ekaterina Kochmar. 2024. Harnessing the power of multiple minds: Lessons learned from LLM routing. *arXiv preprint arXiv:2405.00467* (2024).
- [35] Dimitris Stripelis, Zijian Hu, Jipeng Zhang, Zhaozhuo Xu, Alay Dilipbhai Shah, Han Jin, Yuhang Yao, Salman Avestimehr, and Chaoyang He. 2024. Tensor-opera router: A multi-model router for efficient llm inference. *arXiv preprint arXiv:2408.12320* (2024).
- [36] Tiensun Song, Zhong Qian, Sujun Dong, Peifeng Li, and Qiaoming Zhu. 2022. Rumor detection on social media with graph adversarial contrastive learning. In *Proceedings of the ACM web conference 2022*. 2789–2797.
- [37] Yiliu Sun, Zicheng Zhao, Sheng Wan, and Chen Gong. 2025. CortexDebate: Debating Sparsely and Equally for Multi-Agent Debate. *arXiv preprint arXiv:2507.03928* (2025).
- [38] Khanh-Tung Tran, Dung Dao, Minh-Duong Nguyen, Quoc-Viet Pham, Barry O’Sullivan, and Hoang D Nguyen. 2025. Multi-agent collaboration mechanisms: A survey of llms. *arXiv preprint arXiv:2501.06322* (2025).
- [39] Karl Tuyls. 2023. Multiagent Learning: From Fundamentals to Foundation Models. In *Proceedings of the 2023 International Conference on Autonomous Agents and Multiagent Systems*. 1–1.
- [40] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc V Le, Ed H. Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. 2023. Self-Consistency Improves Chain of Thought Reasoning in Language Models. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=1PL1NIMMrw>
- [41] Zhiheng Xi, Wenxiang Chen, Xin Guo, Wei He, Yiwen Ding, Boyang Hong, Ming Zhang, Junzhe Wang, Senjie Jin, Enyu Zhou, et al. 2025. The rise and potential of large language model based agents: A survey. *Science China Information Sciences* 68, 2 (2025), 121101.
- [42] Zhanxun Xiong, Shichao Liu, Feng Huang, Ziyang Wang, Xuan Liu, Zhongfei Zhang, and Wen Zhang. 2023. Multi-relational contrastive learning graph neural network for drug-drug interaction event prediction. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 37. 5339–5347.
- [43] Yuhao Yang, Chao Huang, Lianghao Xia, and Chenliang Li. 2022. Knowledge graph contrastive learning for recommendation. In *Proceedings of the 45th international ACM SIGIR conference on research and development in information retrieval*. 1434–1443.
- [44] Yingxuan Yang, Mulei Ma, Yuxuan Huang, Huacan Chai, Chenyu Gong, Haoran Geng, Yuanjian Zhou, Ying Wen, Meng Fang, Muhao Chen, et al. 2025. Agentic

- web: Weaving the next web with ai agents. *arXiv preprint arXiv:2507.21206* (2025).
- [45] Yingxuan Yang, Qiying Peng, Jun Wang, Ying Wen, and Weinan Zhang. 2025. Unlocking the potential of decentralized llm-based mas: Privacy preservation and monetization in collective intelligence. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems*. 2896–2900.
- [46] Yonghui Yang, Zhengwei Wu, Le Wu, Kun Zhang, Richang Hong, Zhiqiang Zhang, Jun Zhou, and Meng Wang. 2023. Generative-contrastive graph learning for recommendation. In *Proceedings of the 46th international ACM SIGIR Conference on Research and Development in Information Retrieval*. 1117–1126.
- [47] Zhilin Yang, Peng Qi, Saizheng Zhang, Yoshua Bengio, William Cohen, Ruslan Salakhutdinov, and Christopher D Manning. 2018. HotpotQA: A Dataset for Diverse, Explainable Multi-hop Question Answering. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- [48] Junliang Yu, Xin Xia, Tong Chen, Lizhen Cui, Nguyen Quoc Viet Hung, and Hongzhi Yin. 2023. XSimGCL: Towards extremely simple graph contrastive learning for recommendation. *IEEE Transactions on Knowledge and Data Engineering* 36, 2 (2023), 913–926.
- [49] Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. 2023. MetaMath: Bootstrap Your Own Mathematical Questions for Large Language Models. *arXiv preprint arXiv:2309.12284* (2023).
- [50] Yanwei Yue, Guibin Zhang, Boyang Liu, Guancheng Wan, Kun Wang, Dawei Cheng, and Yiyang Qi. 2025. Masrouter: Learning to route llms for multi-agent systems. *arXiv preprint arXiv:2502.11133* (2025).
- [51] Guibin Zhang, Luyang Niu, Junfeng Fang, Kun Wang, Lei Bai, and Xiang Wang. 2025. Multi-agent architecture search via agentic supernet. *arXiv preprint arXiv:2502.04180* (2025).
- [52] Guibin Zhang, Yanwei Yue, Xiangguo Sun, Guancheng Wan, Miao Yu, Junfeng Fang, Kun Wang, Tianlong Chen, and Dawei Cheng. 2024. G-designer: Architecting multi-agent communication topologies via graph neural networks. *arXiv preprint arXiv:2410.11782* (2024).
- [53] Jiayi Zhang, Jinyu Xiang, Zhaoyang Yu, Fengwei Teng, Xiong-Hui Chen, Jiaqi Chen, Mingchen Zhuge, Xin Cheng, Sirui Hong, Jinlin Wang, Bingnan Zheng, Bang Liu, Yuyu Luo, and Chenglin Wu. 2025. AFlow: Automating Agentic Workflow Generation. In *The Thirteenth International Conference on Learning Representations*.
- [54] Yu Zhang, Kexue Fu, and Zhicong Lu. 2025. RevTogether: Supporting Science Story Revision with Multiple AI Agents. In *Proceedings of the Extended Abstracts of the CHI Conference on Human Factors in Computing Systems*. 1–7.
- [55] Yusen Zhang, Ruoxi Sun, Yanfei Chen, Tomas Pfister, Rui Zhang, and Sercan Arik. 2024. Chain of agents: Large language models collaborating on long-context tasks. *Advances in Neural Information Processing Systems* 37 (2024), 132208–132237.
- [56] Zhuosheng Zhang, Aston Zhang, Mu Li, and Alex Smola. 2023. Automatic Chain of Thought Prompting in Large Language Models. In *The Eleventh International Conference on Learning Representations*. <https://openreview.net/forum?id=5NTt8GFjUHkr>
- [57] Ming Zhou, Dan Zhang, Yuandong Wang, Yangli-Ao Geng, and Jie Tang. 2023. Detecting social bot on the fly using contrastive learning. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management*. 4995–5001.
- [58] Yanqiao Zhu, Yichen Xu, Feng Yu, Qiang Liu, Shu Wu, and Liang Wang. 2021. Graph contrastive learning with adaptive augmentation. In *Proceedings of the web conference 2021*. 2069–2080.
- [59] Lizheng Zu, Lin Lin, Song Fu, Na Zhao, and Pan Zhou. 2025. Collaborative Tree Search for Enhancing Embodied Multi-Agent Collaboration. In *Proceedings of the Computer Vision and Pattern Recognition Conference*. 29513–29522.