

# Beyond Single Models: Enhancing LLM Detection of Ambiguity in Requests through Debate

Ana Davila<sup>1</sup>, Jacinto Colan<sup>2†</sup> and Yasuhisa Hasegawa<sup>1</sup>

<sup>1</sup>Institutes of Innovation for Future Society, Nagoya University, Nagoya, Aichi, Japan  
(Tel: +81-52-789-4481; E-mail: davila.ana@robo.mein.nagoya-u.ac.jp)

<sup>2</sup>Department of Micro-Nano Mechanical Science and Engineering, Nagoya University, Nagoya, Aichi, Japan

**Abstract:** Large Language Models (LLMs) have demonstrated significant capabilities in understanding and generating human language, contributing to more natural interactions with complex systems. However, they face challenges such as ambiguity in user requests processed by LLMs. To address these challenges, this paper introduces and evaluates a multi-agent debate framework designed to enhance detection and resolution capabilities beyond single models. The framework consists of three LLM architectures (Llama3-8B, Gemma2-9B, and Mistral-7B variants) and a dataset with diverse ambiguities. The debate framework markedly enhanced the performance of Llama3-8B and Mistral-7B variants over their individual baselines, with Mistral-7B-led debates achieving a notable 76.7% success rate and proving particularly effective for complex ambiguities and efficient consensus. While acknowledging varying model responses to collaborative strategies, these findings underscore the debate framework’s value as a targeted method for augmenting LLM capabilities. This work offers important insights for developing more robust and adaptive language understanding systems by showing how structured debates can lead to improved clarity in interactive systems.

**Keywords:** large language models, ambiguity, debate, human-robot interaction

## 1. INTRODUCTION

Large Language Models (LLMs) have demonstrated significant capabilities in understanding and generating human language, contributing to more natural interactions with complex systems. However, a persistent challenge lies in the inherent ambiguity of human requests. Such ambiguities, if misinterpreted by an LLM, can lead to incorrect actions or responses, a notable issue in applications that require precise execution, such as task-oriented dialogue systems or human-robot collaboration [1]. Even advanced single LLM architectures may not consistently identify and resolve these subtle but important sources of misunderstanding.

To address these limitations, this paper explores an approach that goes beyond single models. We propose and investigate a multi-agent debate framework where multiple LLMs collaboratively analyze a user’s request. The rationale is that a structured debate, taking advantage of diverse perspectives from different LLM agents, may more effectively surface and scrutinize potential ambiguities than a single-model approach. This collaborative reasoning process is intended to enhance the detection of unclear elements within requests and facilitate the generation of targeted clarification questions, with the goal of leading to more consistent system behavior.

This study empirically evaluates the efficacy of such an LLM debate framework in improving the detection and resolution of ambiguities (Figure 1). We utilized a programmatically generated dataset that features various types of ambiguity commonly encountered in instructional contexts and tested the framework with three distinct LLM architectures. Our contribution lies in examining how this debate mechanism performs, revealing its

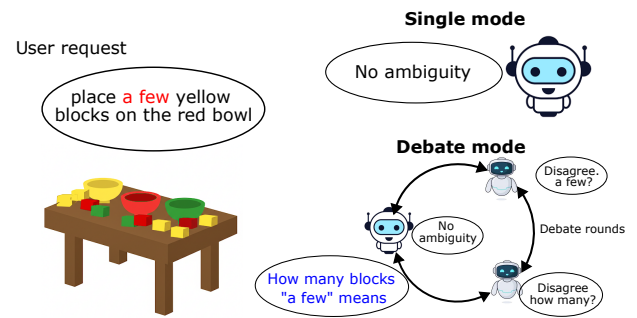


Fig. 1 Multi-agent debate flow

effects on performance for certain models and ambiguity types, and identifying model-specific conditions under which it appears most advantageous. The findings provide insights into the conditional utility of collaborative LLM systems, which informs future efforts to develop a more reliable language understanding in interactive systems.

## 2. RELATED WORKS

Resolving ambiguity in natural language commands is crucial for effective human-robot interaction (HRI), as vague instructions can lead to task failures [2-4]. Foundational HRI research addressed this, with early work on specific ambiguities. Liu et al. studied spatial ambiguity arising from differing frames of reference (e.g., “to the right of the chair”) [5], relevant to our study. Others developed systematic, non-LLM methods for clarification questions to resolve referential ambiguities (e.g., which “medical kit”) [6] or attribute-based disambiguation for “bring-me” tasks (e.g., object color) [7]. While effective for defined cases, these early systems had limited gener-

† Jacinto Colan is the presenter of this paper.

alization. Our programmatic generation of datasets for numerical, attribute, and spatial ambiguities allows controlled investigation, complementing studies with general interaction data [8] and efforts to automate ambiguous scenario generation [2]. This early work underscores the need for robust clarification. Resolving such linguistic ambiguities is a key component of the broader challenge of recognizing the overall user intent and task phase in real-time [9], and our study addresses this by focusing on generating optimal clarifying questions.

Large Language Models (LLMs) have significantly advanced robot comprehension of complex, underspecified instructions and implicit knowledge [10, 11]. LLMs are used for grounding instructions, planning, and crucially, identifying ambiguities and formulating clarifying questions. Chisari et al. introduced AmbResVLM, grounding language goals in visual scenes for task ambiguity reasoning and query generation, with an automated pipeline for ambiguous image-task datasets [8]. The “Ask-to-Act” paradigm tasks embodied agents with asking minimal clarification questions for ambiguities in object attributes, spatial relationships, or size, using LLM reasoning [3]. Zhang et al. explored training LLMs with Reinforcement Learning from Human Feedback (RLHF) by simulating future turns, to learn what clarifying questions to ask and when [12]. Vision-Language Models (VLMs) like FreeGrasp resolve object ambiguities in cluttered scenes [13] and enhance spatial reasoning [14]. Although LLMs allow for broader understanding, challenges such as grounding and hallucinations persist [10, 15], making clarifications generated by LLM vital. Our study’s use of smaller, locally deployable LLMs targets efficient, low-latency inference for real-time robotics.

To address single LLM limitations such as biases or presupposing one interpretation [16, 17], multi-LLM collaboration and debate are emerging [18]. These frameworks improve reasoning and ambiguity resolution by allowing agents to critique and refine interpretations [16]. Lakara et al. proposed the LLM-Consensus, where agents assess contextual consistency, identify ambiguities, and refine reasoning, potentially using external information [16]. Koupae et al.’s Madisse framework assigns initial stances to LLM agents for multi-round debates on summary faithfulness, uncovering errors and identifying “ambiguity” from multiple correct interpretations [17]. Such debates show potential for detecting subtle numerical, attribute, and spatial ambiguities. Hierarchical structures, such as leader-follower models, are also explored; ReMA uses a high-level metathinking agent for oversight and a low-level agent for execution [19].

The principles of multi-agent debate influence our leader-follower methodology by facilitating structured debate. This helps in considering multiple interpretations and evaluating clarifications for optimal questions, leveraging ‘cognitive diversity’ for effective ambiguity resolution. While frameworks like Madisse [17] demonstrate debate for evaluating summary faithfulness, our work specifically applies and assesses a debate protocol for the

pre-execution analysis of commands in a human-robot interaction context, focusing on generating actionable clarification questions.

### 3. METHODOLOGY

Our study investigates the effectiveness of single and multi-agent Large Language Model (LLM) configurations in identifying ambiguities within user instructions and proposing pertinent clarifying questions.

The methodology encompasses a programmatically generated ambiguity dataset, a selection of LLM agents chosen for their suitability for such scenarios, and a structured experimental protocol involving debate mechanisms.

#### 3.1. Ambiguity Dataset Generation

The experiments used a custom dataset of ambiguous instructions, programmatically generated to ensure controlled variation and verifiability. The generation process was automated to create pairs of ambiguous and unambiguous instructions, each grounded in a consistent, predefined scenario context describing a simple environment with a fixed set of objects (e.g., three blocks each of red, yellow, and green, and one bowl of each color on a table).

The generation process aimed at three primary types of ambiguity:

- 1. Numerical Ambiguity:** These instances focused on vague quantities of objects, specifically blocks. An unambiguous instruction would specify a precise quantity (e.g., “two red blocks,” or “all green blocks” which resolved to three). The corresponding ambiguous instruction was created by replacing the precise quantity with an imprecise term (e.g., “a few red blocks,” “some green blocks”) while keeping other elements of the command—such as action, object color, spatial relationship, and target object—identical to its unambiguous counterpart. Actions (e.g., “put,” “place,” “move”), object attributes, and relational components were selected from predefined lists.
- 2. Attribute Ambiguity:** This category involved ambiguity in specifying the properties of an object, either its type (noun) or its color. An unambiguous instruction would use precise terms (e.g., “a single red block”). For noun-based ambiguity, the ambiguous version would replace the specific object noun (e.g., “block”) with a more general or related term (e.g., “cube,” “item,” “thing”) while retaining the correct color. For color-based ambiguity, the specific color term (e.g., “red”) was replaced with a less common synonym or a descriptive phrase (e.g., “crimson,” “cherry-colored”), while the object noun remained correct. The rest of the sentence structure, including quantity (typically singular for attribute tests), action, and relational components, was maintained.
- 3. Spatial Ambiguity:** These cases focused on imprecise descriptions of spatial relationships between objects. An unambiguous instruction would use a clear and specific spatial preposition (e.g., “on the yellow bowl,” “to the left of the green block”). The ambiguous instruction “

was generated by substituting the precise preposition with a more vague or underspecified one (e.g., “near the yellow bowl,” “lateral to the green block,” “along the line of sight of the red block”). The objects being manipulated and referenced, along with the action, were kept consistent between the ambiguous and unambiguous pair.

For all types, the generation procedure randomly selected components from predefined vocabularies to construct a specified number of unique examples for each ambiguity category. Care was taken to ensure that referenced objects were distinct from the primary manipulated object, where appropriate. The resulting dataset provided a controlled yet varied set of challenges for the LLMs.

### 3.2. LLM Agents

Given the intended application in real-time human-robot collaboration, the selection of LLM agents prioritized models with the potential for efficient local deployment and low-latency inference. This is crucial for interactive scenarios where quick responses are necessary for a natural and effective collaborative experience. Therefore, we opted for relatively smaller, yet capable models. In this study, three distinct LLMs were used as debaters:

- Llama-3-8b-instruct
- Gemma-2-9b-it
- Mistral-7b-instruct

These models showcase a variety of open-source architectures known for their popularity and high performance. This diversity enables exploration of how different reasoning patterns interact in a collaborative setting. They offer a balance between high performance and manageable computational demands, making them suitable for deployment on local GPUs commonly found in robotics environments. Each model instance, referred to as an “agent,” was configured with consistent generation parameters, such as a fixed temperature (0.5) to control randomness and a maximum token limit (350 tokens) for responses, ensuring comparability.

### 3.3. Experimental Procedure

The core task for the LLM(s) was to analyze a given ambiguous instruction within its specified scenario context. The objective was to either determine if the command is sufficiently clear for execution or, if deemed ambiguous, to propose a single, optimal clarifying question. Agents were instructed to limit their reasoning to a predefined maximum number of sentences (e.g., four sentences). We implemented two primary experimental conditions:

#### 3.3.1. Single-Agent Baseline

In this configuration, each of the three LLM agents processed each ambiguous instruction from the dataset independently. The agent received the scenario context and the ambiguous instruction, and was prompted to provide its reasoning (within the sentence limit) and then conclude its response by stating that the instruction is clear, or by proposing a single, best clarifying question if ambiguity was detected. This was a single-round in-

teraction, and the agent’s output (argument and proposal) along with processing latency was recorded.

#### 3.3.2. Multi-Agent Debate (Leader-Follower Protocol)

This condition involved a structured debate between three agents: one designated as the “Leader” and the other two as “Followers”. A two-follower configuration was chosen to establish a stronger consensus mechanism, requiring the leader to convince two independent agents and thereby reducing the risk of premature agreement on a flawed interpretation. The roles were rotated among the three LLM models to ensure that each model acted as a leader for every ambiguous instruction, thereby mitigating potential biases associated with a fixed leader. The debate proceeded in rounds, with a predefined maximum number of rounds (e.g., five rounds).

*Round 1: Leader’s Initial Proposal* The Leader agent received the scenario context and ambiguous instruction, then generated an initial proposal consisting of its reasoning and either a verdict of clarity or a specific clarifying question.

*Round 1: Followers’ Evaluation* Each Follower agent received the original context, instruction, and the Leader’s proposal. Followers were prompted to state their stance (“Agree” or “Disagree”), provide reasoning, and if disagreeing, they could offer an alternative question. Specifically, if the Leader proposed a question, a disagreeing Follower could provide a different one; if the Leader declared clarity, a disagreeing Follower’s reasoning was expected to explain the ambiguity, optionally proposing a question or indicating no alternative was needed. Agreeing Followers indicated no alternative question.

*Consensus Check* If both Followers agreed with the Leader’s proposal from Round 1, consensus was considered reached, and the debate for that instruction concluded with the Leader’s proposal as the final outcome.

*Subsequent Rounds (up to the maximum limit)* If consensus was not reached, the Leader received all follower feedback. It then synthesized this feedback to generate a new proposal (reasoning and then either a verdict of clarity or a new/revised clarifying question). This new proposal was then presented to the Followers for evaluation under the same protocol. A consensus check followed each follower evaluation phase.

*Termination Conditions* The debate for a given instruction terminated if:

- Consensus was achieved (both Followers agreed with the Leader’s proposal).
- The maximum number of configured rounds was completed without consensus. In this scenario, the outcome was marked as non-consensus, and the final state of proposals and feedback was recorded.
- An unrecoverable error occurred during an agent’s response generation or in parsing the structured output.

The experimental procedure was applied to individual entries from the ambiguity dataset, ensuring that for each entry, all single-agent baselines were run, and for the multi-agent condition, all rotations of the leader role

were executed.

## 4. EXPERIMENTAL RESULTS

This section analyzes the performance of single and multi-agent LLM configurations in ambiguity detection, focusing on key trends and insights derived from the experimental data. The figures mentioned are assumed to graphically represent the specific data values previously detailed.

### 4.1. Overall Performance and Model Dependencies

Figure 2 shows that the single Gemma2-9B agent achieved a success rate of 80.0%. This rate was considerably higher than that of single Llama3-8B (13.3%) and Mistral-7B (28.3%).

The multi-agent debate framework yielded varied outcomes. For Mistral-7B, debates improved performance from 28.3% to 76.7% (Debate(Mistral-7B)), resulting in the highest success rate among the debate configurations. Llama3-8B also saw an increase in performance, from 13.3% to 40.0% with debate. Conversely, the debate setting appeared detrimental to Gemma2-9B, as Debate(Gemma2-9B) achieved 48.3% success, which was lower than its single-agent performance. This suggests that the utility of a debate structure can be model-dependent and may not universally guarantee improvement. For individual models like Gemma2-9B that demonstrate high baseline performance, the current debate mechanism might interfere with their individual processing capabilities.




Fig. 2 Overall success rate in recognizing ambiguity for single LLM agents (Llama3-8B, Gemma2-9B, Mistral-7B) and multi-agent debate configurations.

### 4.2. Impact of Ambiguity Type

Performance varied significantly by ambiguity type, as detailed in Figures 3, 4, and 5. Single Gemma2-9B performed consistently well across attribute (70%), numerical (85%), and spatial (85%) ambiguities. In contrast, Llama3-8B and Mistral-7B had lower performance as single agents, particularly with spatial ambiguities (Llama3-8B: 0%, Mistral-7B: 10%).

The debate configuration led by Mistral-7B showed

notable improvements over its single-agent counterpart, particularly for more complex types, achieving 65% in attribute, 90% in numerical, and 75% in spatial ambiguities. For instance, the increase in spatial ambiguity detection for Mistral-7B from 10% (single) to 75% (debate) illustrates this effect. However, the Debate(Mistral-7B) configuration did not consistently achieve higher success rates than single Gemma2-9B across all types, with single Gemma2-9B showing higher rates for attribute and spatial tasks. Some debate configurations also performed lower than their single-agent counterparts on specific tasks (e.g., Debate(Llama3-8B) on attribute ambiguity).




Fig. 3 Success rate for Attribute Ambiguity across different LLM configurations.




Fig. 4 Success rate for Numerical Ambiguity across different LLM configurations.

### 4.3. Debate Dynamics: Consensus and Efficiency

The dynamics of the debate process (Figures 6, 7) revealed further differences. Debates led by Mistral-7B achieved consensus with high frequency (98.3% consensus rate, corresponding to a non-consensus rate of 1.7%), predominantly within one or two rounds (with 64.4% of its consensus occurring in Round 2). This suggests effective proposal generation and feedback integration when Mistral-7B acted as leader.

In contrast, debates led by Llama3-8B and Gemma2-




Fig. 5 Success rate for Spatial Ambiguity across different LLM configurations.

9B had considerably higher non-consensus rates (approx. 31.7% and 33.3%, respectively). These configurations also tended to require more rounds to reach consensus when it was achieved; for instance, 45% of Gemma2-9B’s successful consensus occurred in Round 3. This may indicate greater difficulty in reconciling differing views or different synthesis capabilities when these models led debates.

The data in Table 1 indicate that when consensus was reached in those particular experimental runs, the decisions were generally of high quality (e.g., 90 – 95% success for Mistral-7B and Llama3-8B led debates, as presented in Table 1). However, the data from the plots related to Figures 6 and 7 suggest that achieving consensus efficiently and consistently varied among configurations.




Fig. 6 Distribution of the number of rounds required to reach consensus in multi-agent debates that achieved consensus.

#### 4.4. Temporal Costs

The average time for debates, shown in Table 2, ranged from 22.5 to 28.1 seconds. Debates led by Mistral-7B were the quickest in this set of data, which aligns with their tendency to reach consensus efficiently. This temporal overhead is a consideration for real-time applications.

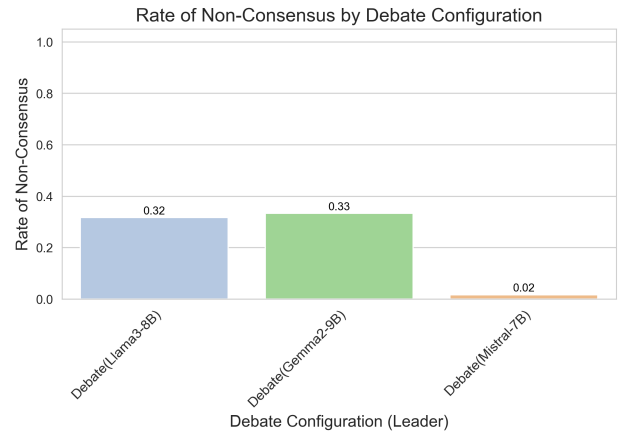


Fig. 7 Rate of non-consensus by debate configuration.

Table 1 Leader LLM Effectiveness in Debates. Metrics include percentage reaching consensus, average rounds to consensus, and success rate of consensus outcomes.

Leader	Cons. Reach (%)	Avg. Rounds	Success (%)
Llama3-8B	85	1.8	92
Gemma2-9B	82	2.1	88
Mistral-7B	90	1.5	95

For models like Llama3-8B and Mistral-7B (based on plot analyses), the observed accuracy gains from debate might be weighed against this cost. For Gemma2-9B, its high single-agent performance, presumably achieved at lower latency, suggests that the added time and reduced accuracy of debates in this study represent a less favorable trade-off.

Table 2 Average debate time (seconds) by leader LLM.

Leader LLM	Avg. Debate Time (s)
Llama3-8B Instruct	25.3
Gemma2-9B	28.1
Mistral-7B	22.5

## 5. CONCLUSION

To address challenges single Large Language Models (LLMs) face with ambiguous requests in robotics, this study demonstrated the value of a multi-agent debate framework. Our investigation revealed that collaboration significantly improves performance for models such as Llama3-8B and Mistral-7B. The debates led by Mistral-7B were particularly successful, achieving a success rate of 76. 7% and proved to be effective in resolving complex spatial ambiguities and efficiently building consensus. However, the utility of the framework is model-dependent. It degraded the performance of the single agent with high performance, Gemma2-9B (80. 0% success). These findings affirm the potential of the debate framework as a targeted enhancement tool while underscoring that it is not a universally optimal solution.

The model-dependent outcomes suggest that the ef-

fectiveness of the debate is based on the complementary capabilities of the agents. For some models, structured feedback helps overcome individual reasoning gaps by forcing the consideration of alternative interpretations. Conversely, for an already strong agent like Gemma2-9B, the debate may introduce distracting noise from other agents, steering it from a correct initial assessment. This insight marks a step toward more robust and adaptive systems in Human-Robot Interaction (HRI). Recognizing the current temporal overhead (22.5–28.1s) as a challenge for real-time use, future work will focus on optimizing the protocol, for example, with adaptive triggers that invoke debate selectively, and conducting deeper qualitative analyses of failure modes and prompt sensitivity.

## REFERENCES

- [1] J. Colan, A. Davila, Y. Yamada, and Y. Hasegawa, “Human-Robot collaboration in surgery: Advances and challenges towards autonomous surgical assistants”, *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (ROMAN)*, 2025.
- [2] M. Ahn, A. Brohan, N. Brown, Y. Chebotar, O. Cortes, B. David, C. Finn, C. Fu, K. Gopalakrishnan, K. Hausman, and A. Herzog, “Do As I Can, Not As I Say: Grounding language in robotic affordances”, *arXiv:2204.01691*, 2022.
- [3] K. Ramachandran, A. Gokul, K. M. Krishna, and V. N. Balasubramanian, “Grounding Multimodal LLMs to Embodied Agents that Ask for Help with Reinforcement Learning”, *arXiv:2504.00907*, 2025.
- [4] A. Davila, J. Colan, and Y. Hasegawa, “Voice control interface for surgical robot assistants”, *Proceedings of the International Symposium on Micro-NanoMechatronics and Human Science (MHS)*, pp. 1–5, 2024.
- [5] C. Liu, J. Walker, and J. Y. Chai, “Ambiguities in Spatial Language Understanding in Situated Human Robot Dialogue”, *Dialog with Robots: Papers from the AAAI Fall Symposium (FS-10-05)*, pp. 50–55, 2010.
- [6] T. Williams, S. Schreitter, and M. Scheutz, “Dempster–Shafer theoretic resolution of referential ambiguity in situated human–robot dialogue”, *Autonomous Robots*, Vol. 43, No. 7, pp. 1811–1833, 2019.
- [7] K. Morohashi, and J. Miura, “Query Generation for Resolving Ambiguity in User’s Command for a Mobile Service Robot”, *Proceedings of the European Conference on Mobile Robots (ECMR)*, 2019.
- [8] E. Chisari, J. O. von Hartz, F. Despinoy, and A. Valada, “Robotic Task Ambiguity Resolution via Natural Language Interaction”, *arXiv:2504.17748*, 2025.
- [9] Y. Yamada, J. Colan, A. Davila, and Y. Hasegawa, “Multimodal semi-supervised learning for online recognition of multi-granularity surgical workflows”, *International Journal of Computer Assisted Radiology and Surgery*, Vol. 19, No. 6, pp. 1075–1083, 2024.
- [10] I. Singh, V. Blukis, A. Mousavian, A. Goyal, D. Xu, J. Tremblay, D. Fox, J. Thomason, and A. Garg., “Progprompt: Generating situated robot task plans using large language models”, *2023 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 11523–11530, 2023.
- [11] S. Chen, Z. Wu, K. Zhang, C. Li, B. Zhang, F. Ma, F. R. Yu, and Q. Li., “Exploring embodied multimodal large models: Development, datasets, and future directions”, *Information Fusion*, Vol. 115, 103198, 2025.
- [12] M. J. Q. Zhang, W. B. Knox, and E. Choi, “Modeling Future Conversation Turns to Teach LLMs to Ask Clarifying Questions”, *Proceedings of the International Conference on Learning Representations (ICLR)*, 2025.
- [13] R. Jiao, A. Fasoli, F. Giuliari, M. Bortolon, S. Povoli, G. Mei, Y. Wang, and F. Poiesi, “Free-form language-based robotic reasoning and grasping”, *arXiv:2503.13082*, 2025.
- [14] X. Han, S. Chen, Z. Fu, Z. Feng, L. Fan, D. An, C. Wang, L. Guo, W. Meng, X. Zhang, and R. Xu., “Multimodal Fusion and Vision-Language Models: A Survey for Robot Vision”, *arXiv:2504.02477*, 2025.
- [15] S. Feng, W. Ding, A. Liu, Z. Wang, W. Shi, Y. Wang, Z. Shen, X. Han, H. Lang, C. Lee, T. Pfister, Y. Choi, and Y. Tsvetkov, “When One LLM Drools, Multi-LLM Collaboration Rules”, *arXiv:2502.04506*, 2025.
- [16] K. Lakara, G. Channing, J. Sock, C. Rupprecht, P. Torr, J. Collomosse, and C. S. de Witt, “LLM-Consensus: Multi-Agent Debate for Visual Misinformation Detection”, *arXiv:2410.20140*, 2025.
- [17] M. Koupaei, J. W. Vincent, S. Mansour, I. Shalyminov, H. He, H. Song, R. Shu, J. He, Y. Nian, A. W. Wong, K. J. Han, and H. Su, “Faithful, Unfaithful or Ambiguous? Multi-Agent Debate with Initial Stance for Summary Evaluation”, *arXiv:2502.08514*, 2025.
- [18] A. Davila, J. Colan, and Y. Hasegawa, “LLM-based ambiguity detection in natural language instructions for collaborative surgical robots”, *Proceedings of the IEEE International Conference on Robot and Human Interactive Communication (ROMAN)*, 2025.
- [19] Z. Wan, Y. Li, Y. Song, H. Wang, L. Yang, M. Schmidt, J. Wang, W. Zhang, S. Hu, and Y. Wen, “ReMA: Learning to meta-think for LLMs with multi-agent reinforcement learning”, *arXiv:2503.09501*, 2025.