

MCTS-Debater: Monte Carlo Tree Search based Rebuttal Generation for Competitive Debate

Xiaoyong Huang^{1*}, Yuting Huang^{1*}, Guanlan Yue^{1*}, Hao Wang^{2†}
Heli Sun^{1,2†}, Jiaruo Wu², Zejun Wang¹, Liang He¹

¹Faculty of Electronic and Information Engineering, Xi'an Jiaotong University

²School of Journalism and New Media, Xi'an Jiaotong University

Correspondence: hxy_computer@stu.xjtu.edu.cn

Abstract

Competitive debate is a multi-turn, zero-sum adversarial task with opposing stances that requires debaters to anticipate opponents' counterarguments in real-time interactions and select strategies yielding long-term advantages. Existing large language model debate methods predominantly adopt a greedy generation paradigm that optimizes only the immediate effectiveness of the current rebuttal while lacking forward-looking planning for the overall debate trajectory. To address this limitation, we propose MCTS-Debater, a Monte Carlo Tree Search framework that models competitive debate as a Markov Decision Process. Rather than selecting a single locally optimal response, MCTS-Debater systematically simulates multi-step adversarial interactions under diverse rebuttal strategies, estimates their long-term expected win rates via tree search, and synthesizes the top-performing strategies into a coherent final rebuttal through multi-strategy fusion. Experimental results on the ORCHID and IQ2 datasets demonstrate that MCTS-Debater significantly outperforms existing baselines in logical validity, factual accuracy, and overall persuasiveness.

1 Introduction

Competitive debate is a rule-governed, topic-centered adversarial argumentation activity. Debaters employ logical reasoning, factual evidence, and rhetorical strategies to persuade judges or audiences to accept their position. Unlike everyday discussions or informal arguments, competitive debate places a premium on argumentative rigor, targeted rebuttals, and strategic adaptability. These qualities collectively embody critical thinking, information synthesis, and impromptu expression skills (Freely, 2009). Competitive debate has emerged as a significant research topic within natural language

processing. Early work primarily focused on sub-tasks such as end-to-end debate systems (Slonim et al., 2021), argument identification (Cheng et al., 2021; Bao et al., 2022), and counterargument generation (Alshomary et al., 2021), achieving promising results. In recent years, with the continuous breakthroughs in large language models (LLMs) (Xu et al., 2025; Gu et al., 2026; Sun et al., 2025), researchers have begun to apply these models to more advanced tasks, including argument quality assessment (Deshpande et al., 2024), debate process simulation (Zhang et al., 2024), and structured debate performance analysis (Wang et al., 2025b; Liang et al., 2024), thereby advancing the field.

However, existing LLM-based debate systems often limit their decision making to immediate responses in the current turn, lacking strategic foresight over the overall debate trajectory, which prevents them from evaluating how current moves affect the final outcome. Representative works illustrate this limitation. InspireDebate (Wang et al., 2025b) employs chain-of-thought (CoT) (?) reasoning to enhance single-step inference but does not model future interactions. R-Debater (Li et al., 2025a) relies on external knowledge retrieval to generate rebuttals without considering long-term strategy. TreeDebater (Wang et al., 2025a) introduces tree structures to anticipate opponent arguments and performs multi-step planning, yet its search horizon remains limited, failing to capture the long-term strategic dependencies across the full debate trajectory. These methods essentially follow a greedy decision paradigm: each turn aims to maximize the local persuasiveness of the current rebuttal without estimating the risk that it might be refuted in subsequent rounds. Crucially, the outcome of a competitive debate often depends on comprehensive performance across multiple interactions. A seemingly moderate argument that effectively blocks the opponent's future attack paths may hold far greater long-term value than a sharp

* Equal contribution.

† Corresponding authors.

but easily countered rebuttal. In other words, debating requires not only “speaking well” but also “seeing far.”

To equip LLMs with forward-looking planning capability, we propose MCTS-Debater, a debate rebuttal framework based on Monte Carlo Tree Search (MCTS) (Świechowski et al., 2023). We model a single debate as a game tree where each rebuttal round corresponds to a decision node. The framework explores future multi-step interactions through MCTS-guided search, covering our rebuttal, the opponent’s response, and subsequent counter-rebuttals. During rollout, we introduce a quantifiable opponent style model characterizing adversaries along five dimensions (logical rigor, evidence reliance, emotional appeal, analogy tendency, and strategy diversity) to yield more realistic simulated responses. By repeatedly sampling trajectories, MCTS estimates the expected win rate of each argumentative move over the entire debate rather than assessing only its immediate impact. Finally, we fuse multiple top-ranked candidates into a single output to enhance robustness against overreliance on any individual strategy.

We conduct extensive experiments on the ORCHID (Zhao et al., 2023) and IQ2 (Zhang et al., 2016) datasets. Evaluation results show that, compared to existing LLM-based methods, MCTS-Debater generates higher-quality arguments.

2 Related Work

2.1 Computational Argumentation

Computational argumentation (Baroni et al., 2018) studies the representation and reasoning of arguments in artificial intelligence, encompassing structured argumentation (Besnard et al., 2014) and abstract argumentation (Dung, 1995). Structured argumentation specifies formal connective relations among arguments, whereas abstract argumentation attends only to abstract arguments and their relationships. With the advancement of large language models, LLM-based argument generation has attracted considerable attention (Hu et al., 2024, 2025). In competitive debating, Agent4Debate (Zhang et al., 2024) pioneered the use of multi-agent interaction to simulate the debate preparation workflow; InspireDebate (Wang et al., 2025b) introduced chain-of-thought augmentation and multi-dimensional direct preference optimization to improve argument quality; R-Debater (Li et al., 2025a) frames debating as a process of recall-

ing and adjusting arguments grounded in argumentative memory and rhetoric; TreeDebater (Wang et al., 2025a) designed the Rehearsal Tree and Debate Flow Tree to anticipate rebuttal and defense strength and to track the dynamic evolution of debate states, respectively. However, when generating rebuttals, these methods rely solely on the current dialogue history and do not simulate the recursive reasoning process of “if I adopt strategy A, how will the opponent counter, and how should I respond in turn,” thus tending to fall into local optima.

2.2 LLM-based Tree Search Methods

Tree search methods have become a key paradigm for overcoming the linear reasoning limitations of LLMs. CoT (?) improves performance on complex tasks by generating intermediate reasoning steps, yet explores only a single linear path. Tree-of-Thoughts (Yao et al., 2023) models reasoning as a search tree, enabling systematic exploration, lookahead, and backtracking across multiple coherent “thought” units, thereby achieving a shift from single-path reasoning to multi-path structured search. Graph-of-Thoughts (Besta et al., 2024) further represents thought units and their dependencies as graph vertices and edges, breaking free from tree-structured constraints. MCTS has been introduced to guide the structured exploration of multi-step reasoning paths (Ha et al., 2025; Li et al., 2025b). The methods target general reasoning or planning tasks. Applying them to competitive debating poses two unique challenges: first, the outcome of a debate hinges on the cumulative persuasiveness after multiple rounds of interaction rather than a single “correct” answer, requiring the search method to project long-term offensive and defensive dynamics; second, the opponent is an active and strategically adaptive adversary, and random simulation cannot faithfully capture the opponent’s real behavior. To address these challenges, this paper integrates MCTS with multi-dimensional opponent style analysis to prospectively estimate the expected win rate of different rebuttal strategies over an entire debate, thereby bridging the gap between general tree search methods and the demands of competitive debating.

3 Methods

This section presents MCTS-Debater, a framework that treats competitive debate as a sequen-

tial decision-making problem under uncertainty. Rather than generating rebuttals myopically, our approach employs MCTS to prospectively evaluate the long-term strategic value of candidate arguments against an adaptive opponent. We first formalize the debate task and model it as a Markov Decision Process (MDP). We then detail the MCTS framework augmented with opponent style modeling and multi-strategy fusion. The overall architecture is illustrated in Figure 1.

3.1 Task Formalization

We formalize competitive debate generation as a multi-turn, stance-opposed, and strategy-driven interactive text generation task. Given a debate topic T and a set of two mutually exclusive stances $\mathcal{L} = \{\ell_{\text{pro}}, \ell_{\text{con}}\}$, the system adopts a designated stance $\ell \in \mathcal{L}$ and engages in up to K_{max} alternating turns with an external opponent. A debate transcript is represented as an ordered utterance history $\mathcal{H}_n = \{u_1, u_2, \dots, u_n\}$, where u_i denotes the natural language utterance at turn i . By convention, the proponent (supporting side) speaks first; thus, odd-numbered turns correspond to ℓ_{pro} , and even-numbered turns correspond to ℓ_{con} . At step $t + 1$, the system generates its response conditioned on the dialogue history $\mathcal{H}_t = \{u_1, u_2, \dots, u_t\}$, where $\mathcal{H}_0 = \emptyset$. The system’s objective is to maximize the probability of winning the debate within the finite horizon K_{max} .

Unlike conventional sequence generation tasks, competitive debate requires each utterance not only to effectively counter the opponent’s immediate argument but also to anticipate potential future rebuttals, thereby selecting strategies that yield long-term advantages. To this end, we formulate the task as an MDP and employ MCTS to approximate the optimal policy.

3.2 Modeling Debate as a Markov Decision Process

We model the decision process at each debate turn as an MDP tuple $\langle \mathcal{S}, \mathcal{A}, P, R, \gamma \rangle$:

State Space $\mathcal{S}$: A state $s_t \in \mathcal{S}$ encapsulates the complete debate context up to the current turn, formally defined as $s_t = \{T, \ell, \mathcal{H}_t, u_{\text{opp}}, r, \pi'\}$, where u_{opp} is the latest opponent utterance, t is the number of completed turns, and π' denotes the opponent’s stylistic features extracted from $\mathcal{H}_t$.

Action Space $\mathcal{A}$: An action $a \in \mathcal{A}$ corresponds to a specific rebuttal strategy. Grounded in argumentation theory (Macagno, 2025), we define five

strategies: *exemplification*, *causal reasoning*, *analogical reasoning*, *value-based argumentation*, and *expert opinion*. Each strategy is paired with a dedicated prompt template to guide the LLM in generating rebuttals consistent with the intended argumentative style.

Transition Dynamics P : Given the current state s_t and action a_t , the system invokes an LLM to generate a rebuttal u_{rebut} according to the chosen strategy. Subsequently, an opponent modeling module simulates the opponent’s next response u'_{opp} , yielding the next state $s_{t+1} \sim P(\cdot | s_t, a_t)$.

Reward Function R : We assign zero immediate reward for non-terminal states. Only upon reaching a terminal state (i.e., after K_{max} turns or early termination) does a judge model evaluate the debate outcome and assign a reward: +1 for a win, 0 for a loss, and 0.5 for a draw.

Discount factor γ : Since the debate outcome depends solely on the final result and the horizon is finite, we set $\gamma = 1$, implying no discounting of future rewards.

3.3 Monte Carlo Tree Search Framework

We adopt the standard four-phase MCTS loop of selection, expansion, rollout, and backpropagation to search for the optimal rebuttal strategy. Each node in the tree corresponds to a debate state s_t , with the root initialized to s_0 .

3.3.1 Selection

Starting from the root, we recursively select the child node with the highest Upper Confidence Bound (UCB) score until reaching a non-fully-expanded node or a leaf. The UCB value is computed as:

$$\text{UCB}(n) = \frac{Q(n)}{N(n)} + C \cdot \sqrt{\frac{\ln N(\text{parent}(n))}{N(n)}}, \quad (1)$$

where $Q(n)$ is the cumulative reward of node n , $N(n)$ is its visit count, and C is the exploration constant.

3.3.2 Expansion

If the current node is non-terminal and has unexplored actions, we select one such action for expansion. Specifically, we first invoke an LLM to generate retrieval keywords based on the chosen strategy, fetch relevant evidence from external knowledge sources, and synthesize a rebuttal by combining the strategy template with the retrieved content.




Figure 1: The overall architecture of MCTS-Debater.

Subsequently, we simulate the opponent’s next response to create a new child node. The simulation mechanism adapts based on the availability of opponent history. In the initial turn where no opponent history exists, we cannot extract stylistic features. Therefore, we employ a CoT prompting strategy to simulate the opponent. CoT provides a principled, style-agnostic baseline by generating step-by-step reasoning without any prior interaction. It has been shown to serve as a strong zero-shot baseline for reasoning tasks (Kojima et al., 2022), making it a natural choice for cold-start opponent simulation.

Once at least one opponent speech becomes available, we activate the *five-dimensional debate style analysis module*. Given the opponent’s historical utterances, this module uses an LLM to quantify stylistic features π' along the following dimensions:

Logical Rigor: whether arguments are logically sound and free of fallacies;

Evidence Reliance: the degree of dependence on facts, data, or case studies;

Emotional Appeal: the frequency of emotive or moralistic language;

Analogy Tendency: the propensity to use analogies or metaphors;

Strategy Diversity: whether multiple argumentation strategies are employed.

Additionally, the model identifies the opponent’s

most frequently used core strategies and generates a concise stylistic summary. During simulation, these features are injected into the LLM’s prompt to enforce style-consistent generation of u'_{opp} . For instance, a highly logical opponent triggers responses with extended reasoning chains, whereas a high-emotion opponent elicits more affective language. This approach ensures that the newly generated state s_{t+1} faithfully reflects plausible adversarial trajectories, thereby enhancing the reliability of style estimation.

3.3.3 Rollout

From the newly expanded node, we randomly select rebuttal strategies and alternately simulate both sides’ utterances until a terminal state is reached. To balance efficiency and exploration, we limit the rollout depth to K_{max} and keep the opponent’s style parameters π' fixed throughout a single rollout.

Upon termination, an LLM-based judge evaluates the debate outcome by scoring both sides on four dimensions: *Emotional Appeal*, *Argument Clarity*, *Logical Arrangement*, and *Relevance to Topic*. The total score for each side is the sum of its dimension scores. The system’s win rate is then computed as:

$$\text{win_rate} = \frac{\text{score}_{\text{system}}}{\text{score}_{\text{system}} + \text{score}_{\text{opponent}}}. \quad (2)$$

If both scores are zero, the win rate defaults to 0.5.

This value serves as the rollout reward $r \in [0, 1]$ for backpropagation.

3.3.4 Backpropagation

The rollout reward r is propagated backward along the selection path, updating the cumulative reward and visit count for all traversed nodes:

$$Q(n) \leftarrow Q(n) + r, \quad N(n) \leftarrow N(n) + 1. \quad (3)$$

3.3.5 Multi-Strategy Fusion

After a fixed budget of MCTS iterations, each child of the root node represents a distinct strategy and its simulated trajectory. The long-term value of each strategy is estimated by its average reward $\bar{Q}(n) = Q(n)/N(n)$. However, single-strategy rebuttals often exhibit limitations: causal reasoning may be logically rigorous but lack emotional resonance, while value-based arguments may be persuasive yet weak in evidential support. To harness complementary strengths, we select the top-3 child nodes with the highest $\bar{Q}(n)$ and feed their generated rebuttals into an LLM for fusion. The fusion instruction directs the model to extract salient elements, such as rigorous reasoning chains, compelling evidence, and effective rhetoric, and synthesize them into a coherent, well-supported, and persuasive final rebuttal. This multi-strategy fusion mechanism mitigates the blind spots of individual strategies and enhances output robustness. Finally, the fused rebuttal is appended to $\mathcal{H}_k$ and returned to the external opponent.

4 Experimental Setup

4.1 Datasets

We evaluate on two public debate datasets: **ORCHID** (Zhao et al., 2023) and **IQ2** (Zhang et al., 2016). **ORCHID** is a large-scale Chinese competitive debate corpus. To prevent data contamination, we exclude all debates used in R-Debater’s retrieval corpus (Li et al., 2025a) and curate a test set of **35 high-quality debates**. These span seven thematic domains with five distinct motions per domain, ensuring topical diversity and zero overlap with prior retrieval corpora. Full motion details are provided in Appendix A.1. **IQ2** comprises 108 Oxford-style English debates with audience voting outcomes and is widely adopted for persuasiveness evaluation. We utilize the complete set of 108 motions.

4.2 Baseline Methods

We compare **MCTS-Debater** against three representative baselines:

(1) **Direct**: A direct prompting approach that relies solely on an LLM to generate debate responses, without external knowledge retrieval or structured reasoning. (2) **Inspire-CoT** (Wang et al., 2025b): Integrates CoT reasoning with external knowledge retrieval to produce stepwise, formalized arguments. (3) **R-Debater** (Li et al., 2025a): A multi-turn debate framework grounded in rhetoric and argument memory theory, which retrieves supporting evidence from a custom-built corpus via keyword aggregation to generate rebuttals.

To ensure fair comparison, all methods operate under a *zero-shot* setting using the same two backbone LLMs: **GPT-4o-mini** (OpenAI et al., 2024) and **Qwen3-8B** (Yang et al., 2025). Furthermore, to maximize performance while controlling variables, both **MCTS-Debater** and **Inspire-CoT** share identical external knowledge source (Wikipedia 2018 snapshot (Karpukhin et al., 2020)) and retriever (E5 embedding model (Wang et al., 2024)) to isolate the impact of strategic planning. In contrast, **R-Debater** retains its original ORCHID-derived retrieval corpus as specified in its seminal work to preserve implementation fidelity. The implementation details are provided in Appendix A.3.

4.3 Experimental Design

We design two complementary experiments to evaluate both immediate argumentation quality and long-term strategic capability:

Task 1: Next-Turn Rebuttal Generation: We truncate existing debate histories from the **ORCHID** test set at turn k and require each model to generate the $(k + 1)$ -th rebuttal from a specified stance (pro or con). This task isolates single-turn generation quality within a controlled context, enabling a fair comparison of immediate argumentative competence. Detailed truncation strategies and input formatting are described in Appendix A.2.

Task 2: End-to-End Adversarial Debate Simulation: To evaluate long-term strategic planning and interactive coherence, we conduct fully automated three-round debates between **MCTS-Debater** and two baselines (**Inspire-CoT** and **R-Debater**) with opposing stances. Each debate consists of six turns (three per side), yielding 1,144 debates across **ORCHID** and **IQ2**. System receive only the motion and assigned stance, generating all utterances autonomously without human intervention, pre-defined arguments, or external guidance. This open adversarial setup faithfully captures comprehensive debating capability under dy-

BaseModel	Method	PRO				CON				ALL			
		Subj.	LV	FA	Inspire	Subj.	LV	FA	Inspire	Subj.	LV	FA	Inspire
Qwen3-8B	Direct	90.58	70.71	36.30	72.04	86.25	42.06	31.17	61.43	88.42	56.39	33.74	66.74
	Inspire-CoT	91.24	75.14	68.44	81.51	89.68	78.89	65.70	80.99	90.46	77.02	67.07	81.25
	R-Debater	90.91	68.77	63.49	76.78	87.96	72.66	<u>67.03</u>	71.21	86.76	70.72	65.26	75.71
	MCTS-Debater	92.45	76.47	69.02	82.60	91.00	<u>78.12</u>	67.65	81.94	91.73	77.30	68.34	82.27
GPT-4o-mini	Direct	89.36	71.83	34.71	71.31	87.76	70.14	30.44	69.03	88.56	70.99	32.58	70.17
	Inspire-CoT	85.89	50.94	60.00	70.68	87.16	<u>76.54</u>	61.10	77.99	86.53	63.74	60.55	74.34
	R-Debater	90.16	71.22	70.31	80.46	89.50	68.54	65.86	78.35	89.83	69.43	68.09	79.29
	MCTS-Debater	91.72	69.78	71.70	81.23	90.57	76.64	<u>64.99</u>	80.69	91.14	73.21	<u>68.35</u>	80.96

Table 1: Evaluation results on the next-turn rebuttal generation task.

namic, competitive interaction.

4.4 Evaluation Metrics

We adopt **InspireScore** (Wang et al., 2025b) as our primary evaluation framework, which assesses debate quality across six dimensions: four subjective (Emotional Appeal, Argument Clarity, Argument Arrangement, Topic Relevance) and two objective (Fact Authenticity, FA; Logical Validity, LV). For Task 1, evaluation focuses on the single generated utterance. For Task 2, the entire multi-turn dialogue is evaluated holistically, with a judge model independently scoring both sides’ overall performance. All scores are produced by **GPT-4o** (OpenAI et al., 2024) using a standardized prompt template to ensure consistency and reliability. For each task, we report mean subjective scores, FA, LV, and composite InspireScore, stratified by PRO stance, CON stance, and overall (ALL).

5 Experimental Result

5.1 Next-Turn Rebuttal Generation

As shown in Table 1, MCTS-Debater achieves the best or second-best overall performance in most settings. Taking Qwen3-8B as an example, under PRO stance MCTS-Debater outperforms Inspire-CoT by 1.09 points in InspireScore, with improvements of 1.21 in Subjective and 1.33 in LV. In the ALL scenario, it exceeds the second-best method by 1.02 points. For GPT-4o-mini, MCTS-Debater leads the best baseline by 1.67 points in InspireScore, with a notable 3.78-point gain in LV. Across both base models, MCTS-Debater attains the highest average InspireScore, confirming that MCTS with opponent style modeling yields LV and subjectively preferable rebuttals. The Direct baseline exhibits low factual accuracy, whereas retrieval-augmented methods (MCTS-Debater and R-Debater) substantially improve this metric, underscoring the value of external knowledge. Although MCTS-Debater’s

FA is slightly below R-Debater on GPT-4o-mini, its overall composite score remains superior, indicating gains in logic and emotional appeal compensate for minor factual gap. Overall, MCTS-Debater demonstrates stable and leading performance across models and stances, validating the integration of game-theoretic planning and opponent style modeling.

5.2 End-to-End Adversarial Debate Simulation

Table 2 presents end-to-end adversarial debate simulation results. MCTS-Debater is paired with Inspire-CoT and R-Debater as opponents. Across all settings, MCTS-Debater never scores below its opponent and achieves a comprehensive lead in most cases. For Qwen3-8B, MCTS-Debater outperforms its opponent by an average of 6.2 points in LV, with higher FA. This confirms that MCTS-based multi-step planning helps anticipate opponent counterarguments and build logically coherent rebuttals, while opponent style modeling improves simulation realism and long-term value estimation. Compared to retrieval-only R-Debater, MCTS-Debater maintains an advantage in most matchups, showing evidence alone cannot replace strategic planning. On GPT-4o-mini, MCTS-Debater wins all pairwise confrontations, with the largest gain in LV. These results strongly validate the effectiveness of integrating multi-step planning and opponent style modeling for debate rebuttal generation.

5.3 Human Evaluation

To further validate the strategic effectiveness of MCTS-Debater in realistic debate interactions, we conducted a human evaluation study. Three human judges determined the winner (Pro or Con) of each debate under fully double-blind conditions, with no knowledge of the participating models. The detailed evaluation protocol is provided

BaseModel	Side	Method	ORCHID				IQ2			
			Subjective	LV	FA	InspireScore	Subjective	LV	FA	InspireScore
Qwen3-8B	CON	Inspire-CoT	92.96	64.20	51.73	75.46	91.28	63.93	63.54	77.51
	PRO	MCTS-Debater	93.67	68.80	56.26	78.10	93.11	71.75	76.10	83.52
	PRO	Inspire-CoT	92.12	63.71	52.76	75.18	91.14	71.67	61.26	78.80
	CON	MCTS-Debater	93.34	70.65	65.32	80.66	92.44	78.03	64.86	81.94
	CON	R-Debater	93.60	63.54	48.28	74.76	88.65	64.75	54.24	74.07
	PRO	MCTS-Debater	89.94	65.20	57.43	75.63	92.35	79.32	63.04	81.77
	PRO	R-Debater	90.51	58.15	51.68	72.71	89.21	59.31	59.05	74.19
	CON	MCTS-Debater	92.04	72.52	61.87	79.62	92.43	76.18	58.50	79.89
	CON	Inspire-CoT	91.09	55.54	52.70	72.60	90.90	66.82	64.90	78.38
	PRO	MCTS-Debater	93.21	68.87	60.08	78.84	92.52	74.20	69.98	82.30
GPT-4o-mini	PRO	Inspire-CoT	91.21	53.48	58.48	73.59	91.26	64.36	62.03	77.23
	CON	MCTS-Debater	92.98	68.48	61.42	78.96	91.91	70.72	64.69	79.81
	CON	R-Debater	89.83	50.40	57.57	71.91	88.65	64.75	64.24	76.57
	PRO	MCTS-Debater	92.87	61.12	66.61	78.37	92.35	69.32	73.04	81.77
	PRO	R-Debater	91.58	57.62	51.42	73.05	90.52	61.96	58.07	75.27
	CON	MCTS-Debater	92.64	67.72	65.95	79.74	91.68	72.12	68.22	80.92

Table 2: Results of End-to-End Adversarial Debate Simulation Evaluation.

PRO	CON	ORCHID			IQ2		
		PRO Win	CON Win	Winner	PRO Win	CON Win	Winner
MCTS-Debater	Inspire-CoT	0.800	0.200	PRO	0.686	0.314	PRO
Inspire-CoT	MCTS-Debater	0.486	0.514	CON	0.400	0.600	CON
MCTS-Debater	R-Debater	0.829	0.171	PRO	0.657	0.343	PRO
R-Debater	MCTS-Debater	0.343	0.657	CON	0.286	0.714	CON

Table 3: Winning rates in human evaluation.

in Appendix B. The evaluation covered both the ORCHID and IQ2 datasets across four distinct model-pair configurations. Inter-annotator agreement, measured by Fleiss’ Kappa, was 0.65, indicating substantial consistency among the judges. As reported in Table 3, MCTS-Debater consistently outperformed both baselines, regardless of its assigned stance. Notably, when arguing as Pro against Inspire-CoT (Con), MCTS-Debater won 80.0% of debates on ORCHID and 68.6% on IQ2. Its advantage persisted when acting as Con, securing win rates of 51.4% (ORCHID) and 60.0% (IQ2) against the same baseline. Similar trends were observed against R-Debater, where MCTS-Debater achieved win rates of 82.9% and 65.7% as Pro and Con respectively on ORCHID, and 65.7% and 71.4% on IQ2. These results collectively demonstrate that MCTS-Debater’s multi-step game planning and opponent style modeling significantly enhance its persuasive power in competitive debate, with this advantage holding robustly across different datasets and argumentative stances.

5.4 Ablation Study

To thoroughly investigate the contribution of each core component of MCTS-Debater, we designed

Table 4: Ablation study comparing MCTS-Debater with alternative designs.

Method	Sub	LV	FA	All
Multi-random	86.71	73.22	67.43	78.52
Five-fuse	82.66	61.41	67.43	73.54
Immediate-only	88.47	73.22	66.59	79.18
MCTS-Debater	91.14	73.21	68.35	80.96

four ablation variants on the next-turn rebuttal generation task, using GPT-4o-mini as the base model under the ALL scenario. *W/o MCTS* means MCTS is disabled and a randomly selected fixed strategy is used for rebuttal. *W/o Style* means opponent style modeling is removed and the opponent responds using CoT prompting. *W/o Fusion* means multi-strategy fusion is disabled and only the best child node is used. *W/o Retrieval* means evidence retrieval is turned off and the system relies solely on the LLM’s internal knowledge.

In addition to these component-level ablations, we further compare against three alternative designs to verify the superiority of our approach. *Multi-random* uses multiple random strategies without MCTS-guided selection. *Five-fuse* fuses all

Topic: Do advanced social media platforms benefit or harm the mental health of contemporary people?			
Method	Excerpt from the Generated Rebuttal	Rebuttal Features	Strengths or Weaknesses
Inspire-CoT	“While you mention that some content creators use self-deprecation or humor to gain popularity, this is not the overarching trend on social media. ... The algorithms are designed to maximize engagement, which often leads to sensationalism and conflict.”	Causal reasoning, concession, negation, no external evidence.	Weakness: Falls into “appeal to exception” logic; arguments speculative; lacks factual support.
R-Debater	“Dr. John Cacioppo ... highlights the concept of social connectivity ... Research conducted by the University of Pennsylvania demonstrates that participation in online communities can lead to increased self-worth.”	Heavy reliance on expert opinions; loosely structured.	Weakness: Evidence weakly linked to opponent’s claims; lacks logical coherence.
MCTS-Debater	“Counterexample 1: Social media as a supportive community ... A study in the Journal of Health Communication found a significant decrease in feelings of isolation. Counterexample 2: #BellLet’sTalk ... Counterexample 3: Dr. Julie Smith shares coping strategies that promote resilience.”	Multi-strategy fusion; source-backed claims; anticipates counter-arguments.	Strength: Concrete examples + evidence + reasoning; preemptively blocks potential rebuttals; high logical rigor.

Figure 2: Qualitative comparison of rebuttals generated by different methods on a debate topic.




Figure 3: Ablation Study on GPT-4o-mini.

five strategies uniformly without adaptive weighting. *Immediate-only* generates rebuttals directly without any iterative refinement.

As shown in Figure 3 and Table 4, removing any component leads to performance degradation. The removal of evidence retrieval causes the most significant drop in factual accuracy, highlighting its critical role for credibility. Disabling MCTS results in the largest decline in subjective quality, confirming the necessity of multi-step planning for persuasive rebuttals. Opponent style modeling and multi-strategy fusion also provide consistent, positive gains.

Among the additional baselines, MCTS-Debater consistently outperforms all alternatives. *Immediate-only* achieves competitive results on Sub but falls short on FA and All, demonstrating the importance of iterative reasoning for factual grounding. *Multi-random* shows reasonable performance but lacks the structured exploration provided by MCTS. *Five-fuse* performs the worst across all metrics, particularly on LV and All, indicating that naive uniform fusion introduces noise and degrades response quality. These results collectively validate the design and effectiveness of all proposed modules.

5.5 Case Study

To further demonstrate the effectiveness of MCTS-Debater, we conduct a case study on Task 1 with the topic: “Do advanced social media platforms benefit or harm the mental health of contemporary people?”. As shown in Figure 2, the three systems exhibit starkly different argumentation patterns:

(1) **Strategy diversity.** Inspire-CoT employs basic causal reasoning but fails to construct affirmative counterpoints, resorting to negation without substantiation. R-Debater cites authorities but does so in a contextually disconnected manner. In contrast, MCTS-Debater synthesizes multiple argu-

mentative strategies (exemplification, causal analysis, normative appeal), structuring each rebuttal as a tripartite unit comprising a concrete example, a verifiable source, and logical linkage.

(2) Evidential integration. Only MCTS-Debater grounds every claim in domain-relevant, verifiable evidence (e.g., peer-reviewed studies, public health campaigns, practitioner content). The other baselines either omit evidence entirely (Inspire-CoT) or deploy it generically (R-Debater), resulting in weak rebuttal-to-claim alignment.

(3) Forward-looking planning. Critically, MCTS-Debater preemptively addresses potential counter-responses by citing statistical evidence to neutralize the “rare exception” objection. This foresight stems directly from the multi-step opponent reactions simulation of MCTS, a capability absent in both baselines.

This case corroborates our quantitative findings: MCTS-Debater’s superiority arises from its integrated architecture that jointly optimizes strategy selection, evidence retrieval, and long-horizon planning, enabling rebuttals that are both factually accurate and strategically robust.

5.6 Hyperparameter Analysis

To investigate the impact of the number of MCTS interactions on debate generation performance, we conducted a hyperparameter analysis on Task 1. The number of MCTS interactions was set to 1, 2, 4, 8, and 16, while keeping GPT-4o-mini as the base model. The experimental results are shown in Table 5.

Table 5: Results of hyperparameter analysis on MCTS interaction times.

Iterations	Sub	LV	FA	All
1	88.47	73.22	66.59	79.18
2	89.32	72.54	67.41	79.65
4	90.73	74.16	68.03	80.91
8	91.14	73.21	68.35	80.96
16	91.11	73.05	67.91	80.80

From the results, it can be observed that as the number of interactions increases from 1 to 4, all metrics exhibit a steady upward trend. This indicates that a moderate search depth helps the model more fully explore the debate space and capture critical long-range dependencies.

However, when the number of interactions is further increased to 8 and 16, performance gains

level off or even slightly decline. This phenomenon suggests that excessive search depth not only fails to bring significant performance improvements but may also introduce redundant information or noise, leading to diminishing marginal returns.

In summary, a larger number of MCTS interactions does not always yield better performance. In this task, 8 interactions strike a good balance between performance and computational cost, being sufficient to capture the key long-range dependencies required for debate generation. Therefore, subsequent experiments adopt 8 MCTS interactions by default.

6 Conclusion

In this paper, we presented MCTS-Debater, a novel framework for competitive debate rebuttal generation that integrates Monte Carlo Tree Search with multi-strategy reasoning, opponent style modeling, and evidence retrieval. Extensive experiments on two real-world debate datasets demonstrate that MCTS-Debater consistently outperforms strong baselines, including LLM-based chain-of-thought methods and retrieval-augmented models, across both next-turn rebuttal and end-to-end adversarial debate settings. Ablation studies and qualitative case analyses further confirm the essential contributions of strategic planning, style-aware simulation, and multi-strategy fusion. Our work highlights the value of forward-looking game-theoretic planning in argumentative dialogue systems and opens new avenues for applying LLMs to complex, multi-turn adversarial tasks.

Acknowledgement

This work was supported in part by the Xianyang Major Scientific and Technological Achievements Transformation Special Project L2024-ZDKJ-ZDCGZH-0012 , and the CAAI-Tencent Rhino-Bird Open Research Fund. Any opinions, findings, and conclusions expressed here are those of the authors and do not necessarily reflect the views of the funding agencies.

Limitations

Despite the promising results, our approach has certain limitations. In most competitive debates, a debater’s argumentative style tends to remain relatively stable throughout a single debate, which justifies our use of a static five-dimensional profile extracted from historical utterances. However, it

is still possible that an opponent may deliberately shift strategies across rounds or exhibit inconsistent stylistic patterns, especially in high-stakes or adaptive debating scenarios. Our current opponent style model, which assumes stylistic consistency after the initial analysis, may not fully capture such dynamic changes. Future work could address this by incorporating online adaptation mechanisms, such as updating style estimates incrementally after each round using few-shot prompting or lightweight fine-tuning.

Ethics Statement

This work focuses on computational argumentation and debate generation, with no involvement of human participants, personally identifiable information, or sensitive data collection. All experiments are conducted exclusively on the publicly available ORCHID and IQ2 datasets, which has been previously released under permissive licenses for academic research and contains no private or protected content.

We acknowledge that automated debate systems carry potential risks of misuse, including the generation of persuasive misinformation or manipulative rhetoric. To mitigate these concerns, we do not release any standalone generation tools or unmoderated model checkpoints. All code and evaluation protocols are provided solely for reproducible research and include explicit safeguards against out-of-context deployment. Furthermore, our evaluation framework prioritizes factual authenticity and logical validity alongside rhetorical quality, reinforcing responsible development practices within the field.

This research complies with institutional guidelines for ethical AI research and does not require formal ethics board approval, as it involves only secondary analysis of existing public datasets without human subject interaction.

References

- Milad Alshomary, Shahbaz Syed, Arkajit Dhar, Martin Potthast, and Henning Wachsmuth. 2021. Argument undermining: Counter-argument generation by attacking weak premises. *arXiv preprint arXiv:2105.11752*.
- Jianzhu Bao, Jingyi Sun, Qinglin Zhu, and Ruifeng Xu. 2022. Have my arguments been replied to? argument pair extraction as machine reading comprehension. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 29–35.
- Pietro Baroni, Dov Gabbay, Massimilino Giacomin, and Leendert Van der Torre. 2018. Handbook of formal argumentation.
- Philippe Besnard, Alejandro Garcia, Anthony Hunter, Sanjay Modgil, Henry Prakken, Guillermo Simari, and Francesca Toni. 2014. Introduction to structured argumentation. *Argument & Computation*, 5(1):1–4.
- Maciej Besta, Nils Blach, Ales Kubicek, Robert Grestenberger, Michal Podstawski, Lukas Gianinazzi, Joanna Gajda, Tomasz Lehmann, Hubert Niewiadomski, Piotr Nyczyk, and Torsten Hoeffler. 2024. [Graph of thoughts: Solving elaborate problems with large language models](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(16):17682–17690.
- Liying Cheng, Tianyu Wu, Lidong Bing, and Luo Si. 2021. Argument pair extraction via attention-guided multi-layer multi-cross encoding. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6341–6353.
- Darshan Deshpande, Zhivar Sourati, Filip Ilievski, and Fred Morstatter. 2024. Contextualizing argument quality assessment with relevant knowledge. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 2: Short Papers)*, pages 316–326.
- Phan Minh Dung. 1995. On the acceptability of arguments and its fundamental role in nonmonotonic reasoning, logic programming and n-person games. *Artificial intelligence*, 77(2):321–357.
- Austin J Freeley. 2009. *Argumentation and debate: Critical thinking for reasoned decision making*.
- Jiawei Gu, Xuhui Jiang, Zhichao Shi, Hexiang Tan, Xuehao Zhai, Chengjin Xu, Wei Li, Yinghan Shen, Shengjie Ma, Honghao Liu, Saizhuo Wang, Kun Zhang, Zhouchi Lin, Bowen Zhang, Lionel Ni, Wen Gao, Yuanzhuo Wang, and Jian Guo. 2026. [A survey on llm-as-a-judge](#). *The Innovation*, 7(6):101253.
- Rui Ha, Chaozhuo Li, Rui Pu, Litian Zhang, Xi Zhang, and Sen Su. 2025. Dsg-mcts: A dynamic strategy-guided monte carlo tree search for diversified reasoning in large language models. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 10541–10555.
- Zhe Hu, Hou Pong Chan, Jing Li, and Yu Yin. 2025. Debate-to-write: A persona-driven multi-agent framework for diverse argument generation. In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 4689–4703.

- Zhe Hu, Hou Pong Chan, and Yu Yin. 2024. Americano: Argument generation with discourse-driven decomposition and agent interaction. In *Proceedings of the 17th International Natural Language Generation Conference*, pages 82–102.
- Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi Chen, and Wen-tau Yih. 2020. Dense passage retrieval for open-domain question answering. In *Proceedings of the 2020 conference on empirical methods in natural language processing (EMNLP)*, pages 6769–6781.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. Large language models are zero-shot reasoners. In *Proceedings of the 36th International Conference on Neural Information Processing Systems, NIPS '22*, Red Hook, NY, USA. Curran Associates Inc.
- Maoyuan Li, Zhongsheng Wang, Haoyuan Li, and Jiamou Liu. 2025a. R-debater: Retrieval-augmented debate generation through argumentative memory. *arXiv preprint arXiv:2512.24684*.
- Qingyao Li, Wei Xia, Xinyi Dai, Kounianhua Du, Weiqi Liu, Yasheng Wang, Ruiming Tang, Yong Yu, and Weinan Zhang. 2025b. Rethinkmcts: Refining erroneous thoughts in monte carlo tree search for code generation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 8103–8121.
- Jingcong Liang, Rong Ye, Meng Han, Ruofei Lai, Xinyu Zhang, Xuan-Jing Huang, and Zhongyu Wei. 2024. Debatrix: Multi-dimensional debate judge with iterative chronological analysis based on llm. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 14575–14595.
- Fabrizio Macagno. 2025. Argumentation schemes. *The Routledge Handbook of Argumentation Theory*, pages 169–182.
- OpenAI, :, Aaron Hurst, Adam Lerer, Adam P. Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, Aleksander Madry, Alex Baker-Whitcomb, Alex Beutel, Alex Borzunov, Alex Carney, Alex Chow, Alex Kirillov, and 401 others. 2024. *Gpt-4o system card*. Preprint, arXiv:2410.21276.
- Noam Slonim, Yonatan Bilu, Carlos Alzate, Roy Bar-Haim, Ben Bogin, Francesca Bonin, Leshem Choshen, Edo Cohen-Karlik, Lena Dankin, Lilach Edelstein, Liat Ein-Dor, Roni Friedman-Melamed, Assaf Gavron, Ariel Gera, Martin Gleize, Shai Gretz, Dan Gutfreund, Alon Halfon, Daniel Hershcovich, and 34 others. 2021. *An autonomous debating system*. *Nature*, 591(7850):379–384.
- Yu-Zhu Sun, He-Li Sun, Jian-Cong Ma, Peng Zhang, and Xiao-Yong Huang. 2025. *Multimodal agent ai: A survey of recent advances and future directions*. *Journal of Computer Science and Technology*, 40(4):1046–1063.
- Maciej Świechowski, Konrad Godlewski, Bartosz Sawicki, and Jacek Mańdziuk. 2023. Monte carlo tree search: A review of recent modifications and applications. *Artificial Intelligence Review*, 56(3):2497–2562.
- Danqing Wang, Zhuorui Ye, Xinran Zhao, Fei Fang, and Lei Li. 2025a. Strategic planning and rationalizing on trees make llms better debaters. *arXiv preprint arXiv:2505.14886*.
- Fuyu Wang, Jiangtong Li, Kun Zhu, and Changjun Jiang. 2025b. Inspired debate: Multi-dimensional subjective-objective evaluation-guided reasoning and optimization for debating. *arXiv preprint arXiv:2506.18102*.
- Liang Wang, Nan Yang, Xiaolong Huang, Binxing Jiao, Linjun Yang, Daxin Jiang, Rangan Majumder, and Furu Wei. 2024. *Text embeddings by weakly-supervised contrastive pre-training*. Preprint, arXiv:2212.03533.
- Fengli Xu, Qian Yue Hao, Zefang Zong, Jingwei Wang, Yunke Zhang, Jingyi Wang, Xiaochong Lan, Jiahui Gong, Tianjian Ouyang, Fanjin Meng, Chenyang Shao, Yuwei Yan, Qinglong Yang, Yiwen Song, Sijian Ren, Xinyuan Hu, Yu Li, Jie Feng, Chen Gao, and Yong Li. 2025. *Towards large reasoning models: A survey of reinforced reasoning with large language models*. Preprint, arXiv:2501.09686.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. *Qwen3 technical report*. Preprint, arXiv:2505.09388.
- Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Tom Griffiths, Yuan Cao, and Karthik Narasimhan. 2023. Tree of thoughts: Deliberate problem solving with large language models. *Advances in neural information processing systems*, 36:11809–11822.
- Justine Zhang, Ravi Kumar, Sujith Ravi, and Cristian Danescu-Niculescu-Mizil. 2016. Conversational flow in oxford-style debates. In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 136–141.
- Yiqun Zhang, Xiaocui Yang, Shi Feng, Daling Wang, Yifei Zhang, and Kaisong Song. 2024. Can llms beat humans in debating? a dynamic multi-agent framework for competitive debate. *arXiv preprint arXiv:2408.04472*.
- Xiutian Zhao, Ke Wang, and Wei Peng. 2023. *Orchid: A chinese debate corpus for target-independent stance detection and argumentative dialogue summarization*. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, page 9358–9375. Association for Computational Linguistics.

A Experimental Setup Details

A.1 Dataset Details

ORCHID (Zhao et al., 2023) is a large-scale Chinese competitive debate corpus comprising 1,134 real-world debate transcripts. Each transcript is annotated with the debate motion, participating institutions, speaker identities, explicit stances (pro/con), and turn-by-turn utterances. To construct our test set, we first exclude the 1,000 debates used by R-Debater (Li et al., 2025a) to build its dedicated retrieval corpus. From the remaining data, we select 35 complete, high-quality debates, ensuring no overlap with existing retrieval corpora. These debates cover seven representative thematic domains: Society and Livelihood, Politics and Governance, Economy and Development, Technology and Future, Ethics and Values, Education and Youth, and International Relations. Each domain contains five distinct motions. The list of motions in the test set is shown in table 6.

A.2 Next-Turn Rebuttal Generation Task Design

This experiment assesses single-turn argument generation under controlled contexts using two evaluation scenarios derived from the ORCHID debate dataset. We truncate debate histories after Round 3 and Round 5, requiring models to generate the subsequent Round 4 response as PRO and Round 6 response as CON, respectively. The ORCHID dataset organizes debates into three phases: an initial structured phase (Rounds 1–4) with fixed roles and clear stances, a middle free-exchange segment featuring rapid alternation and ambiguous positional boundaries, and a closing phase with renewed stance clarity.

In the Round-3 truncation setting, the input concludes with the first CON rebuttal. The model then acts as PRO to generate Round 4, which is the final turn of the initial structured phase. This scenario tests targeted counter-argumentation against explicit opponent claims. In the Round-5 truncation setting, the input encompasses the complete initial phase plus the opening turn of the closing phase delivered by PRO. The model subsequently acts as CON to generate Round 6, evaluating strategic coherence maintenance after re-establishing stance clarity following the free-exchange segment.

These two time points correspond to the *early engagement* and *strategic equilibrium* stages. Fixing the input history while excluding the ambigu-

ous free-exchange segment ensures reproducibility, minimizes noise, and enables fair comparison of strategic reasoning capabilities across systems.

A.3 Implementation Details

All experiments were conducted on a server equipped with two NVIDIA A100 (40GB) GPUs, running PyTorch 2.10. The open-source model Qwen3-8B was deployed using vLLM (Li et al., 2025b) for efficient inference, while GPT-4o-mini and GPT-4o were accessed via the official OpenAI API.

Given that competitive debates are often structured into three alternating turns per side, we set the maximum number of interaction rounds $R_{\max} = 3$. All baseline methods share the same temperature setting of 0.7 to balance randomness and determinism, and no task-specific fine-tuning was applied.

B Human Evaluation Setup

To assess the strategic effectiveness and persuasiveness of MCTSDebater in full multi-turn debates, we conducted a human evaluation on Experiment 2 (End-to-End Adversarial Debate Simulation) using GPT-4o-mini as the base model for debate generation. We randomly selected 35 debate motions from the ORCHID dataset and 35 from the IQ2 dataset. For each motion, we generated complete six-turn debates (three rounds per side) for two head-to-head matchups: MCTSDebater vs. Inspire-CoT and MCTSDebater vs. R-Debater, covering both stance assignments (MCTSDebater as PRO and as CON). This resulted in a total of 140 full-length debate transcripts (70 per dataset).

Three volunteers (all holding a bachelor’s degree, with basic experience in logical reasoning) participated as judges. Each debate transcript was anonymized with a random identifier, and the order of transcripts was shuffled per judge. Evaluators were asked to decide which side (PRO or CON) won the overall debate, based on persuasiveness, logical coherence, factual accuracy, and strategic progression over the three rounds. For each matchup and stance configuration, we computed the number of debates in which the majority (≥ 2 out of 3) of judges favored MCTSDebater. Win rate is defined as majority-win debates divided by total debates in that condition. Inter-rater agreement was measured using Fleiss’ Kappa. The evaluation interface is shown in Figure 4.

You will see a complete debate transcript containing 10 rounds of speeches from both PRO and CON. Please read the entire debate carefully and judge which side wins the overall debate based on the following criteria:

Overall persuasiveness: Which side presents stronger, more convincing arguments?

Logical coherence: Are the arguments internally consistent across rounds? Is the reasoning rigorous?

Factual accuracy: Are the facts or evidence cited reliable and appropriate?

Strategic progression: Does the side effectively refute the opponent's weaknesses while consistently advancing its own core narrative?

When making your judgment, consider all four dimensions holistically, rather than relying on a single round of speech. Finally, select PRO or CON as the winning side and provide a score.

Note: All debate content was automatically generated by different models. The debater identities have been anonymized. Please make your judgment independently and avoid discrediting.

[illegible]

Figure 4: The evaluation interface.

C Prompts

In the implementation of our method, we design and employ multiple prompt templates serving key modules such as opponent style analysis, retrieval query generation, strategy-driven rebuttal generation, style-conditioned opponent simulation, and multi-strategy fusion. The detailed content of all prompt templates and their invocation locations within the MCTS-Debater framework are illustrated in Figures 5 to 10. These prompts are designed following the argumentation logic of each respective strategy and adopt a unified structured format to ensure the stability of LLM outputs.

D AI Assistant Usage Statement

In the course of this research, we employed AI assistants, including ChatGPT¹, Qwen², and DeepSeek³, strictly as auxiliary tools for code implementation and manuscript polishing. The core research idea, experimental design, theoretical framework, and all substantive intellectual contributions were conceived and executed entirely by the authors. We affirm that all scientific conclusions, data analyses, and creative decisions remain the sole responsibility of the authors.

¹<https://chatgpt.com>

²<https://www.gianwen.com>

³ <https://www.deepseek.com/>

Prompt Template for Opponent Style Analysis

System Prompt:

You are an expert in debate style analysis. Based on the following speech, rate it across the five dimensions:

1. `logical_rigor`: How rigorous the argument is, and whether it contains logical fallacies.
2. `evidence_reliance`: The extent to which the argument relies on data, facts, or examples.
3. `emotional_appeal`: The degree of emotional language or moral judgment used.
4. `analogy_tendency`: The frequency of using analogies or metaphors.
5. `strategy_diversity`: Whether different debate strategies (causal reasoning, exemplification, value-based argumentation, etc.) are mixed.

Additionally, identify the 1–2 most frequently used concrete strategies from the list: [causal reasoning, exemplification, analogical reasoning, value-based argumentation, expert opinion].

Output in JSON format, containing the five fields above, plus "dominant_strategies" (as a list) and "summary" (a one-sentence conclusion).

Output only JSON, no extra text.

User Prompt:

Debate speech: [opponent's recent utterances]

Figure 5: Example prompt used for the five-dimensional opponent style analysis module. The module is invoked during the rollout phase of MCTS to estimate the opponent's stylistic features from historical utterances.

Strategy	Prompt Template
Exemplification	Refute the following opponent's speech using the [Exemplification] strategy. Requirement: Provide at least one specific counterexample to prove that the opponent's general conclusion does not hold. The counterexample should be realistic, relevant, and persuasive. Opponent's speech: {opponent_speech}
Causal Reasoning	Refute the following opponent's speech using the [Causal Reasoning] strategy. Requirement: Point out logical flaws in the opponent's causal reasoning (e.g., confusing correlation with causation, omitted variables, reversed causality) and propose an alternative explanation. Opponent's speech: {opponent_speech}
Analogical Reasoning	Refute the following opponent's speech using the [Analogical Reasoning] strategy. Requirement: Construct an analogy with the same structure but an absurd conclusion, revealing the error in the opponent's reasoning. The analogy must maintain logical isomorphism. Opponent's speech: {opponent_speech}
Value-based Argumentation	Refute the following opponent's speech using the [Value-based Argumentation] strategy. Requirement: Identify the implicit value conflict in the opponent's argument and demonstrate why the value we advocate (e.g., fairness, efficiency, freedom) should take precedence over the opponent's value. Opponent's speech: {opponent_speech}
Expert Opinion	Refute the following opponent's speech using the [Expert Opinion] strategy. Requirement: Cite authoritative experts, institutions, or research findings to directly counter the opponent's point. Provide sources (may be reasonably fabricated but must be credible). Opponent's speech: {opponent_speech}

Figure 6: Prompt templates for the five rebuttal strategies used in the action space of MCTS-Debater. Each template is instantiated with the actual opponent speech during the expansion phase.

#	Topic	Category
1	The participation of top celebrities in theater plays revitalizes/disrupts the theater market	Economy & Development
2	'Oriental Selection' Model Shows Promise vs. Unlikely to Last	Economy & Development
3	Should or should not encourage Gen-Z to reform the workplace	Society & Livelihood
4	Choosing to pursue a master's degree due to the difficulty of finding a job is a wise choice for current college students.	Education & Youth
5	Love_Hate is more unforgettable	Morality & Values
6	At age 20, if you have an opportunity to become famous overnight, should you take it or not?	Morality & Values
7	Is pain necessary for growth?	Morality & Values
8	In contemporary society, making life 'faster' vs. 'slower' is more desirable	Society & Livelihood
9	Love is Restraint vs. Indulgence	Morality & Values
10	The Romance of Chinese Culture Lies in Complexity_Simplicity	Education & Youth
11	Self-interest/Altruism as a Better Survival Strategy	Morality & Values
12	Fragmentation of Information Enhances/Decreases the Cognitive Level of Contemporary People	Society & Livelihood
13	Online public opinion makes our thinking more diverse/single-minded	Society & Livelihood
14	In advancing the evolution of social institutions, should or should not the pursuit of social welfare be prioritized?	Society & Livelihood
15	Should social development be guided by the current will of the people/national long-term well-being	Politics & Governance
16	The development and popularization of ChatGPT-like tools will enhance/weaken the cognitive abilities of ordinary people	Technology & Future
17	The prevalence of beauty technology is good/bad for human aesthetics	Technology & Future
18	Do advanced social media platforms benefit or harm the mental health of contemporary people?	Technology & Future
19	In the face of AIGC technology development, should we prioritize copyright protection or promoting innovation?	Technology & Future
20	Should AI Have Legal Personality	Politics & Governance
21	Campus competitive debate is more meaningful to students for its practicality/non-practicality	Education & Youth
22	Should Chinese Esports Encourage the Introduction of Foreign Players	Economy & Development
23	Is AI replacing human labor a blessing or a death knell for human civilization?	Economy & Development
24	Should electronic devices have standardized charging ports?	Technology & Future
25	Should private enterprises be involved in China's space industry?	Economy & Development
26	E-sports should or should not enter the Olympic Games	International Relations
27	Should Chinese Language be _ Should not be included as a compulsory university course	Education & Youth
28	Should Mainland China raise admission standards for international students?	Education & Youth
29	Heroes Create the Times _ Times Create Heroes	Politics & Governance
30	Should Marijuana Transactions in Mainland China Be Legalized?	Politics & Governance
31	Social media is/is not an important driver of social democratization	Politics & Governance
32	"The Clash of Civilizations" can_be_avoided / cannot_be_avoided	International Relations
33	Resettling foreign refugees is/is not a national obligation	International Relations
34	The primary cause of rampant terrorism lies with Middle Eastern countries / Western countries	International Relations
35	Reducing U.S. involvement in Asian affairs is more beneficial than harmful / more harmful than beneficial to the international situation	International Relations

Table 6: Complete list of 35 debate motions in the ORCHID test set.

Prompt Template for Retrieval Query Generation

System Prompt: (implicit) You are a debate argument analysis assistant.

User Prompt:

Based on the following opponent's speech and the rebuttal strategy we are about to use, list the key facts or data premises you need to verify in order to support your rebuttal.

Debate topic: {topic}

Our stance: {stance}

Rebuttal strategy to be used: {strategy['name']}

Opponent's speech: {opponent_speech}

Please output the key premises to be retrieved, one per line (at most 3). Output only the premises, without any additional explanation.

Example format:

Premise 1: ...

Premise 2: ...

Premise 3: ...

Figure 7: Prompt used to generate retrieval queries during the expansion phase.

Prompt Template for Style-Conditioned Opponent Simulation

System Prompt:

You are an AI capable of imitating specific debate styles.

User Prompt:

You must imitate the following style to respond:

Your debate style parameters:

- Logical rigor: {logical_rigor}

- Evidence reliance: {evidence_reliance}

- Emotional appeal: {emotional_appeal}

- Analogy tendency: {analogy_tendency}

- Strategy diversity: {strategy_diversity}

- Dominant strategies: {dominant_strategies}

Style summary: {summary}

Debate topic: {topic}

Your previous speech: {last_opponent_speech}

The opponent just refuted: {rebuttal}

Please strictly follow the above style parameters to generate a natural and persuasive response to refute the opponent's speech.

Note: If your debate style tends to be logically rigorous, use more reasoning chains; if evidence-oriented, fabricate or cite reasonable data; if emotion-oriented, add emotional coloring.

Figure 8: Prompt used during the MCTS rollout phase to simulate the opponent's next response conditioned on the estimated five-dimensional style features.

Prompt Template for Rebuttal Generation (with Evidence and Strategy)

System Prompt:

You are a debate expert. Please generate a powerful refutation according to the specified strategy.

User Prompt:

Debate topic: {topic}

Our stance: {stance}

Reference materials:

{evidence_section}

Debate strategy:

{strategy_prompt}

Opponent's speech:

{opponent_speech}

Please make full use of the above evidence and strategy (if any) to refute your opponent and give your opinion and reasoning.

Figure 9: Complete prompt used during the expansion phase to generate a rebuttal. It combines the strategy-specific instruction, retrieved evidence, and opponent's speech. The placeholders are dynamically filled.

Prompt Template for Multi-Strategy Fusion

User Prompt:

You are given three alternative rebuttals to the same opponent speech. Combine their strengths into a single, more powerful rebuttal. Preserve logical force, evidence, and persuasive language. Avoid repetition.

Opponent speech: {opponent_speech}

Rebuttal A: {rebuttal_A}

Rebuttal B: {rebuttal_B}

Rebuttal C: {rebuttal_C}

Fused rebuttal:

Figure 10: Prompt used to fuse the top-3 candidate rebuttals.