MULTI-EXPERT COLLABORATIVE OPTIMIZATION FOR LOW-RESOURCE MACHINE TRANSLATION THROUGH DEBATE AND ANT COLONY SEARCH

Anonymous authors
Paper under double-blind review

ABSTRACT

Despite recent advances in neural machine translation, low-resource translation remains constrained by insufficient parallel supervision and model-specific generation biases. Heterogeneous experts can provide complementary translation hypotheses, yet conventional multi-expert methods often rely on direct reranking or fixed fusion, limiting iterative collaboration. We propose a collaborative translation optimization framework that combines hierarchical multi-agent debate with ant colony search. Frozen heterogeneous experts first generate initial translations to construct an input-dependent collaboration graph, over which ant agents explore refinement and fusion paths. Intra-group debate identifies and corrects local translation errors, while inter-group interaction exchanges elite paths and complementary evidence. Debate feedback is further incorporated into transition estimation and pheromone updating, jointly guiding subsequent search without modifying the expert models or training an additional reinforcement-learning policy. Experiments on low-resource Mongolian–Chinese translation achieve 30.8 BLEU, 55.1 chrF++, and 0.463 COMET, improving over the strongest controlled baseline by 0.6 BLEU, 0.7 chrF++, and 0.009 COMET, while requiring fewer model calls and lower inference latency. Ablation and convergence analyses further validate the effectiveness of the proposed collaborative search framework.

1 INTRODUCTION

Neural machine translation (NMT) has achieved considerable progress, but its performance remains dependent on large-scale parallel corpora. In low-resource settings, insufficient supervision often leads to limited lexical coverage and unstable translation quality (Sennrich & Zhang, 2019). Subword modelling alleviates rare-word sparsity (Sennrich et al., 2016b), while back-translation and transfer learning exploit monolingual data or knowledge from resource-rich tasks (Sennrich et al., 2016a; Zoph et al., 2016). Multilingual NMT further promotes cross-lingual knowledge sharing through unified models (Johnson et al., 2017). Nevertheless, these approaches primarily improve individual models and cannot fully eliminate model-specific generation biases. In low-resource Mongolian–Chinese translation, a single system may still produce fluent outputs with semantic omissions, terminology errors, or over-generalised expressions.

Multiple translation systems provide complementary hypotheses that can be exploited through ensemble methods and candidate-level selection (Garmash & Monz, 2016; Stahlberg & Byrne, 2017; Freitag et al., 2022). Recent LLM-based approaches further perform candidate ranking and fusion (Jiang et al., 2023; Vernikos & Popescu-Belis, 2024). Nevertheless, most existing methods use a fixed candidate set and perform one-shot selection or aggregation, limiting their ability to iteratively identify and correct errors across hypotheses.

Multi-agent debate provides an alternative mechanism for iterative translation refinement by allowing agents to critique and revise candidate outputs (Du et al., 2024; Liang et al., 2024). Translation-agent frameworks have also demonstrated the potential of specialised agents for translation, editing, and proofreading (Wu et al., 2024). However, these approaches typically rely on predefined interaction rounds or fixed role-based workflows, without explicitly exploring a population of competing refinement and fusion trajectories.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

Ant colony optimisation (ACO) offers a complementary population-based search mechanism for exploring such alternative trajectories. Adaptive transition rules, heuristic information, and pheromone updating can improve exploration and convergence (Bai et al., 2024; Wu et al., 2023). Yet conventional ACO mainly relies on numerical heuristics and lacks linguistic reasoning for identifying semantic, terminological, and fluency errors. This motivates combining the linguistic evaluation of multi-agent debate with the persistent search and accumulated experience of ACO.

To combine these capabilities, we propose a hierarchical multi-agent debate-guided ant colony search framework for low-resource Mongolian–Chinese machine translation. Multiple frozen heterogeneous experts first generate complementary candidate translations, which initialise an input-specific translation collaboration graph. Ant agents then explore alternative refinement and fusion paths. Intra-group debate evaluates local paths and identifies candidate-specific errors, while inter-group debate exchanges elite solutions and key critiques across subgroups. The resulting feedback is incorporated into transition estimation and pheromone updating, allowing linguistic assessment and accumulated search experience to jointly guide collaborative translation optimisation.

The main contributions of this work are summarised as follows:

- We formulate multi-expert low-resource machine translation as an input-specific collaborative path optimisation problem and construct a translation collaboration graph for iterative refinement and fusion.
- We propose a hierarchical multi-agent debate mechanism that combines intra-group path critique with inter-group elite exchange, enabling local error identification and cross-group knowledge sharing.
- We develop a debate-aware ant colony search strategy that incorporates linguistic feedback into transition estimation and pheromone updating, jointly guiding translation assessment and search.

2 Related Work

Low-Resource and Iterative Machine Translation. Recent studies have explored adapting large language models to low-resource translation, with limited lexical knowledge and source-language understanding remaining important bottlenecks (Elsner et al., 2024). Multilingual continual pretraining and parameter-efficient adaptation improve translation coverage without full model retraining (Lu et al., 2024; Su et al., 2024), while dictionary- and example-based prompting enables translation for previously unseen languages (Zhang et al., 2024). However, these methods primarily improve individual models and do not explicitly exploit complementary outputs from heterogeneous systems.

Recent work has explored collaborative and iterative translation optimisation. SCALE connects a specialised translation model with a general-purpose LLM (Cheng et al., 2024), while iterative prompting and LLMRefine progressively revise translations using feedback and search (Chen et al., 2024b; Xu et al., 2024). These approaches demonstrate the benefit of iterative revision, but generally optimise a single evolving hypothesis or follow a fixed refinement process rather than maintaining multiple competing revision and fusion paths.

Multi-Agent Collaboration and Debate. Multi-agent methods enable heterogeneous language models to exchange evidence and revise decisions. ReConcile combines multi-round discussion with confidence-weighted consensus (Chen et al., 2024a), while recent studies show that debate effectiveness depends on interaction protocols, communication patterns, and computational budgets (Smit et al., 2023; Li et al., 2024). In machine translation, Multi-agentMT adopts a translate–post-edit–proofread workflow (Kim, 2025), whereas M-MAD uses specialised agents to evaluate translation quality from multiple dimensions (Feng et al., 2025). However, these approaches generally rely on predefined roles or fixed communication stages. Our method instead applies hierarchical debate to dynamically constructed translation paths and uses the resulting feedback to guide subsequent search.

Ant Colony Path Optimisation. Recent ACO methods improve global exploration through adaptive heuristics, transition rules, and pheromone updating (Liu et al., 2023; Cui et al., 2024). However,

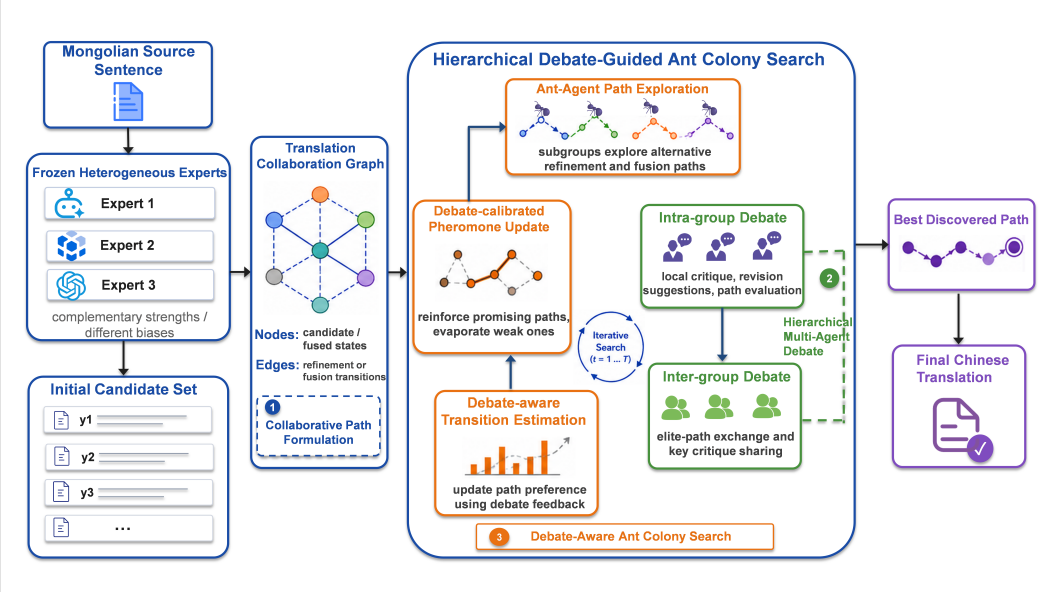

Figure 1: Overall framework of the proposed hierarchical multi-agent debate-guided ant colony search method.

they are mainly designed for geometric path-planning tasks, where paths are evaluated using numerical objectives such as distance and smoothness. Our framework instead constructs a translation collaboration graph and incorporates debate-derived semantic, terminological, and fluency feedback into transition estimation and pheromone updating.

3 METHOD

As shown in Figure 1, we propose a hierarchical multi-agent debate-guided ant colony search framework for low-resource neural machine translation. Frozen heterogeneous experts generate complementary translations to construct an input-specific collaboration graph. Ant agents explore refinement and fusion paths, while intra- and inter-group debate provide local critique and cross-group elite exchange. Debate feedback guides transition estimation and pheromone updating, and the final translation is selected from the best discovered path without updating the expert models or training an additional reinforcement-learning policy.

3.1 PROBLEM SETUP AND CANDIDATE INITIALIZATION

Given a Mongolian source sentence x, the objective is to generate a Chinese translation that preserves the source meaning while maintaining terminological consistency and target-language fluency. We consider K frozen heterogeneous translation experts, denoted by $\mathcal{E} = \{E_k\}_{k=1}^{K}$. Each expert independently generates an initial candidate, and the resulting candidate set is defined as

$$\mathcal{Y}^{(0)}(x) = \left\{ y_k^{(0)} = E_k(x) \right\}_{k=1}^{K}. \tag{1}$$

The experts provide complementary information due to differences in architecture, pre-training data, multilingual knowledge, and generation preferences. Their outputs serve as initial search states rather than final translations. All experts are frozen and use the same decoding configuration, ensuring that candidate diversity stems from expert heterogeneity.

Rather than directly selecting or statically combining the initial candidates, we formulate multi-expert translation as a collaborative path optimisation problem. A collaboration path is represented as

$$\mathcal{P} = (v_0, a_1, v_1, \ldots, a_L, v_L), \tag{2}$$

where v_l denotes the translation state at step l, and a_l represents a refinement or fusion operation conditioned on accumulated debate evidence. The terminal translation associated with $\mathcal{P}$ is denoted by $y_{\mathcal{P}}$, corresponding to the translation component of v_L.

The objective is to identify the collaboration path with the highest overall utility:

$$\mathcal{P}^* = \arg \max_{\mathcal{P} \in \Omega(x)} F(\mathcal{P}), \tag{3}$$

where $\Omega(x)$ denotes the feasible collaboration-path space and $F(\mathcal{P})$ jointly considers the debate-calibrated translation quality and computational cost of the path. During optimisation, the expert parameters remain fixed, while the translation states and collaboration paths are iteratively refined.

3.2 Hierarchical Multi-Agent Debate

Although heterogeneous experts provide diverse candidates, their outputs may contain semantic omissions, terminology errors, or unnatural expressions. We therefore introduce hierarchical multi-agent debate to evaluate and refine competing collaboration paths before subsequent path search.

Intra-group debate. The ant agents are divided into M subgroups, with each agent maintaining a collaboration path $\mathcal{P}_n$. Within each subgroup, agents exchange their paths and critique semantic adequacy, terminology, lexical choice, and fluency. These critiques identify path-specific errors and guide local refinement, while allowing different subgroups to explore distinct search directions.

Inter-group debate. After intra-group refinement, each subgroup selects elite paths according to reference-free quality and path cost. Their elite paths and key critiques are then exchanged across subgroups to compare alternative refinement strategies and transfer complementary evidence, reducing premature convergence.

For each agent n, hierarchical debate produces

$$\mathcal{R}_n^{\mathrm{deb}} = \left(\widetilde{\mathcal{P}}_n, Q_n^{\mathrm{sol}}, A_n, G_n \right), \tag{4}$$

where $\widetilde{\mathcal{P}}_n$ is the refined path, Q_n^{sol} is the reference-free quality of its terminal translation, A_n is the normalized agreement among debating agents, and G_n measures the quality gain from debate-guided refinement.

Let y_n^{pre} and y_n^{post} denote the terminal translations before and after refinement. We define

$$G_n = Q_{\mathrm{QE}} \left(y_n^{\mathrm{post}}, x \right) - Q_{\mathrm{QE}} \left(y_n^{\mathrm{pre}}, x \right), \tag{5}$$

where $Q_{\mathrm{QE}}(\cdot)$ is a reference-free translation-quality estimator. The resulting debate signals are used to condition transition estimation and pheromone updating. Debate is not treated as an independent graph transition; instead, its critiques and revision evidence guide subsequent refinement and fusion operations. No reference translation or reference-based metric is used during collaborative search.

3.3 Translation Collaboration Graph Construction

To represent alternative translation trajectories, we construct an input-dependent translation collaboration graph whose nodes denote temporary translation states and whose edges represent feasible refinement or fusion operations. For a source sentence x,

$$\mathcal{G}_x = (\mathcal{V}_x, \mathcal{T}_x), \tag{6}$$

where $\mathcal{V}_x$ is the set of translation states and $\mathcal{T}_x$ is the set of feasible transitions. The graph preserves the evolution of candidate translations beyond one-shot selection.

Algorithm 1 Hierarchical Debate-Guided Ant Colony Search

Require: Source sentence x; frozen experts $\mathcal{E}$; ants N; subgroups M; iterations T
Ensure: Final translation $\hat{y}$

```
Generate initial candidates Y^(0)(x) = {E_k(x)}_{k=1}^K
Initialize G_x, pheromones τ_ij, P_vis ← ∅, and D^(0) ← ∅
Partition the ants into M subgroups
for t = 1 to T do
    Expand G_x and construct paths {P_n^(t)}_{n=1}^N using Eq. equation 13 and D^(t−1)
    Refine paths through intra-group debate and inter-group elite exchange
    Obtain {P̃_n^(t), Q_n^deb}_{n=1}^N and updated evidence D^(t)
    Update pheromones using Eq. equation 15 and add refined paths to P_vis
    P_t^* ← arg max_{P∈P_vis} F(P)
    if t > r and F(P_t^*) − F(P_{t−r}^*) < ε_stop then
        break
    end if
end for
return ŷ = y_{P^*}, where P^* = arg max_{P∈P_vis} F(P)
```

Translation states. Each state is defined as

$$v_i = (y_i, h_i, q_i)\,, \tag{7}$$

where y_i is the current translation, h_i records its refinement and debate history, and

$$q_i = Q_{\mathrm{QE}}\,(y_i, x) \tag{8}$$

is its reference-free quality estimate. The initial states are constructed from heterogeneous expert outputs $\mathcal{Y}^{(0)}(x)$.

Graph transitions. A directed transition is represented as

$$e_{ij} = (v_i, a_{ij}, v_j)\,, \tag{9}$$

where a_{ij} is either a refinement or fusion operation:

$$v_j = \begin{cases} A_{\mathrm{ref}}(v_i, h_i, \mathcal{D}_n), & a_{ij} = \mathrm{ref}, \\ A_{\mathrm{fus}}(v_i, v_\ell, \mathcal{D}_n), & a_{ij} = \mathrm{fus}. \end{cases} \tag{10}$$

Here, A_{ref} revises the current translation using its history and debate evidence $\mathcal{D}_n$, while A_{fus} combines complementary information from two translation states. Before the first debate round, $\mathcal{D}_n$ is empty.

A successor state is retained only if it improves the reference-free quality:

$$Q_{\mathrm{QE}}(y_j, x) > Q_{\mathrm{QE}}(y_i, x). \tag{11}$$

Ant agents then traverse feasible transitions to construct the collaboration paths defined in Eq. equation 2. Different ants may start from different expert candidates and explore distinct refinement and fusion trajectories. To control search cost, only quality-improving transitions are retained, and cross-path fusion is activated after subgroup elite selection, avoiding exhaustive graph expansion while preserving promising trajectories.

3.4 DEBATE-AWARE COLLABORATIVE PATH OPTIMISATION

The translation collaboration graph defines the feasible refinement and fusion trajectories but does not determine which transitions should be preferred. We therefore incorporate hierarchical debate evidence into both transition estimation and pheromone updating. Local path construction is guided by current linguistic evidence, while accumulated debate outcomes influence subsequent search iterations.

Debate-aware transition estimation. For a feasible transition e_{ij} from state v_i to state v_j, its reference-free quality gain and debate-aware heuristic value are defined as

$$\begin{aligned} \Delta q_{ij} &= Q_{\mathrm{QE}}(y_j, x) - Q_{\mathrm{QE}}(y_i, x), \\ \eta_{ij} &= \exp\left(\lambda_q \Delta q_{ij} + \lambda_d D_{ij} - \lambda_c C_{ij}\right), \end{aligned} \tag{12}$$

where D_{ij} is the normalised support for transition e_{ij} derived from the hierarchical debate evidence, and C_{ij} is its normalised computational cost. A transition is therefore preferred when it provides a larger estimated quality gain, receives stronger debate support, and incurs a lower additional cost. Before debate evidence becomes available in the first search round, D_{ij} is set to zero.

At iteration t, ant n selects its next transition according to

$$p_{ij}^{(n)}(t) = \frac{[\tau_{ij}(t)]^{\alpha}[\eta_{ij}]^{\beta}}{\sum_{\ell \in \mathcal{N}_i^{(n)}} [\tau_{i\ell}(t)]^{\alpha}[\eta_{i\ell}]^{\beta}}, \tag{13}$$

where $\tau_{ij}(t)$ is the pheromone value of transition e_{ij}, $\mathcal{N}_i^{(n)}$ is the feasible successor set of ant n, and α and β control the influence of accumulated search experience and current debate-aware evidence, respectively.

Debate-calibrated pheromone updating. After path construction, hierarchical debate evaluates and refines each path $\widetilde{\mathcal{P}}_n$. Its debate-calibrated quality is defined as

$$Q_n^{\mathrm{deb}} = \lambda_1 Q_n^{\mathrm{sol}} + \lambda_2 A_n + \lambda_3 G_n, \tag{14}$$

where Q_n^{sol}, A_n, and G_n denote the solution-quality, debate-agreement, and refinement-gain signals produced by hierarchical multi-agent debate. These components are normalised to a common range before aggregation, with non-negative weights satisfying $\lambda_1 + \lambda_2 + \lambda_3 = 1$.

The pheromone value is updated as

$$\tau_{ij}(t+1) = (1-\rho)\tau_{ij}(t) + \sum_{n=1}^{N} Q_n^{\mathrm{deb}} \frac{\mathbb{I}\left[e_{ij} \in \widetilde{\mathcal{P}}_n\right]}{L_n + \varepsilon_\tau}, \tag{15}$$

where $\rho \in (0,1)$ is the evaporation coefficient, L_n is the number of transitions in $\widetilde{\mathcal{P}}_n$, and $\mathbb{I}[\cdot]$ is the indicator function. Path-length normalisation prevents longer paths from receiving disproportionately large pheromone contributions. Consequently, transitions belonging to paths that remain reliable after intra-group critique and inter-group comparison receive stronger reinforcement.

Path utility and termination. The utility of a refined collaboration path jointly considers its debate-calibrated quality and computational cost:

$$F\left(\widetilde{\mathcal{P}}_n\right) = Q_n^{\mathrm{deb}} - \lambda_{\mathrm{cost}} C\left(\widetilde{\mathcal{P}}_n\right), \tag{16}$$

where $C(\widetilde{\mathcal{P}}_n)$ is the normalised path cost determined by the number of agent calls, generated tokens, and refinement or fusion operations.

Variant	BLEU ↑	chrF++ ↑	COMET ↑
w/o Intra-group Debate	30.3	54.6	0.456
w/o Inter-group Debate	30.5	54.8	0.459
w/o Debate-aware Transition	30.4	54.7	0.457
w/o Debate-calibrated Pheromone	30.2	54.5	0.455
w/o Fusion Operation	30.4	54.6	0.458
Full Model	**30.8**	**55.1**	**0.463**

Table 1: Ablation results on the full-data Mongolian–Chinese test set, averaged over three random seeds.

Iteration	BLEU ↑	chrF++ ↑	COMET ↑
0	28.8	52.6	0.427
1	29.9	53.9	0.446
2	30.5	54.7	0.457
3	**30.8**	**55.1**	**0.463**

Table 2: Translation quality over successive search iterations on the full-data Mongolian–Chinese test set. Results are averaged over three random seeds.

The search terminates when the maximum number of iterations is reached or when

$$F\left(\mathcal{P}_t^*\right) - F\left(\mathcal{P}_{t-r}^*\right) < \varepsilon_{\text{stop}}, \tag{17}$$

where $\mathcal{P}_t^*$ denotes the best path visited up to iteration t, r is the early-stopping window, and $\varepsilon_{\text{stop}}$ is the minimum required improvement. Algorithm 1 summarises the complete search procedure and returns the terminal translation of the best visited path. Reference-based metrics such as BLEU, chrF++, and COMET are used only for final evaluation and are not available during collaborative search.

4 Experiments

4.1 Experimental Setup

Datasets and Expert Models. Experiments are conducted on Mongolian–Chinese machine translation using approximately one million parallel sentence pairs from news and daily communication domains, supplemented with CCMT resources. Fixed validation and test sets are used for model selection and evaluation. The expert pool consists of LLaMA-7B, BLOOM-7B1, and Qwen-7B, which are independently adapted with LoRA and kept frozen during collaborative search.

Compared Methods. We compare our method with Best Single Expert, QE Reranking, Static Multi-Expert Fusion, Single-Agent Refinement, Standard ACO, and Flat MAD, covering single-expert translation, candidate selection, static fusion, iterative refinement, path search, and multi-agent debate.

Implementation Details. We use nine logical ant agents organized into three subgroups. Agents share the same frozen instruction model but maintain independent translation states, search paths, and debate histories. The maximum path length and search iterations are both set to three, with each iteration consisting of intra-group and inter-group debate. Hyperparameters are selected on the validation set, and no reference translation is accessed during collaborative search.

Evaluation Metrics. We report BLEU, chrF++, and COMET for translation quality, together with sentence-level latency and model calls for computational cost. Experiments are conducted on a single NVIDIA GeForce RTX 5090 GPU with 32 GB memory, and search-based results are averaged over three random seeds.

Method	BLEU ↑	chrF++ ↑	COMET ↑	Latency (s) ↓	Model Calls ↓
One-shot candidate utilisation					
Best Single Expert	28.3	52.0	0.416	1.42	1.0
QE Reranking	28.8	52.6	0.427	4.31	3.0
Static Multi-Expert Fusion	29.2	53.2	0.438	5.18	4.0
Iterative refinement and search					
Single-Agent Refinement	29.7	53.8	0.445	7.24	6.0
Standard ACO	30.0	54.2	0.451	10.41	9.1
Flat MAD	30.2	54.4	0.454	12.63	10.8
Ours	**30.8**	**55.1**	**0.463**	11.32	9.7

Table 3: Main results on the full-data Mongolian–Chinese test set, averaged over three random seeds.

4.2 Main Results

Table 3 reports results on the full-data Mongolian–Chinese test set. All multi-expert methods use the same frozen expert pool and initial candidates. Iterative methods share the same maximum generation budget, and model calls include candidate generation, refinement, fusion, and debate.

As shown in Table 3, multi-expert methods consistently improve over the best individual expert. QE Reranking and Static Multi-Expert Fusion provide gains through candidate selection or aggregation, while Single-Agent Refinement further benefits from iterative revision. Standard ACO and Flat MAD achieve additional improvements through path search and multi-agent interaction, but lack the combination of persistent path reinforcement and hierarchical linguistic feedback.

Our method achieves the highest BLEU, chrF++, and COMET scores, outperforming Flat MAD by 0.6 BLEU, 0.7 chrF++, and 0.009 COMET, while also using fewer model calls and lower latency. Although Standard ACO is slightly more efficient, its lower translation quality suggests that hierarchical debate and pheromone-based search provide a more favorable quality–cost trade-off.

4.3 Ablation Study

We conduct controlled ablation experiments on the full-data Mongolian–Chinese test set. All variants use the same frozen expert pool, initial candidates, search budget, and stopping criterion. Results are reported in Table 1.

Removing either level of hierarchical debate degrades translation quality. Without intra-group debate, BLEU and COMET decrease by 0.5 and 0.007, respectively, while removing inter-group debate also causes consistent declines, confirming the benefits of local critique and cross-group elite exchange. Excluding debate evidence from transition estimation reduces BLEU by 0.4, whereas removing debate-calibrated pheromone updating causes the largest degradation, with decreases of 0.6 BLEU, 0.6 chrF++, and 0.008 COMET. Removing fusion also lowers all three metrics, indicating that cross-path fusion is necessary to exploit complementary information from heterogeneous experts. Overall, these results verify the complementary contributions of hierarchical debate, debate-aware search, and candidate fusion.

4.4 Convergence Behaviour

We further examine translation quality across iterative search. Iteration 0 denotes the best initial candidate selected by the reference-free quality estimator, before refinement, debate, or pheromone updating. BLEU, chrF++, and COMET are used only for post-hoc evaluation.

As shown in Table 2, quality improves consistently across iterations, with the largest gain occurring in the first iteration, where BLEU increases from 28.8 to 29.9. The second iteration reaches 30.5 BLEU, 54.7 chrF++, and 0.457 COMET, accounting for approximately 85% of the total BLEU gain. The third iteration provides smaller but consistent improvements of 0.3 BLEU, 0.4 chrF++,

Source / Method	Translation
Source	Монгол Улс импортын эрчим хүчний хамаарлыг бууруулахын тулд сэргээгдэх эрчим хүчний төслүүдэд хөрөнгө оруулалтаа нэмэгдүүлж байна.
Expert 1	蒙古国正加大对绿色能源项目的投资，以降低对进口能源的依赖。
Expert 2	为降低对进口能源的依赖，蒙古国正在投资可再生能源项目。
Expert 3	蒙古国正大力发展可再生能源，以减少对进口能源的依赖。
QE Reranking	为降低对进口能源的依赖，蒙古国正在投资可再生能源项目。
Flat MAD	蒙古国正加大可再生能源领域的投资，以降低对进口能源的依赖。
Ours	**蒙古国正通过加大对可再生能源项目的投入，降低对进口能源的依赖。**
Reference	为减少对进口能源的依赖，蒙古国正在增加可再生能源项目的投资。

Figure 2: A representative example of collaborative translation refinement across heterogeneous experts.

and 0.006 COMET, indicating diminishing returns. These results support the three-iteration search budget and early stopping as a practical balance between quality improvement and search cost.

4.5 Case Study

Figure 2 presents a representative example illustrating how the proposed framework integrates complementary information from heterogeneous experts. Expert 1 preserves the main meaning but uses less precise terminology, whereas Expert 2 omits the increase in investment and Expert 3 generalises the project-level information. QE Reranking cannot recover information missing from the selected candidate, while Flat MAD still produces an incomplete expression. By contrast, the proposed method combines refinement and fusion paths through hierarchical debate, yielding a translation with more complete semantics, accurate terminology, and fluent expression.

5 Conclusion

This paper presents a hierarchical multi-agent debate-guided ant colony search framework for low-resource Mongolian–Chinese machine translation. The proposed method organises complementary translations generated by frozen heterogeneous experts into an input-specific collaboration graph and uses intra-group and inter-group debate to guide iterative refinement and fusion paths. Debate feedback is further incorporated into transition estimation and pheromone updating, enabling promising collaboration paths to be progressively reinforced without updating the expert models or learning an additional reinforcement-learning policy. Experimental results show that the proposed framework achieves 30.8 BLEU, 55.1 chrF++, and 0.463 COMET, outperforming Flat MAD by 0.6 BLEU, 0.7 chrF++, and 0.009 COMET while requiring fewer model calls and lower inference latency. Ablation and convergence analyses further confirm the complementary contributions of hierarchical debate, debate-aware transition estimation, candidate fusion, and debate-calibrated pheromone updating. Future work will focus on reducing the inference overhead introduced by iterative agent interaction and extending the framework to additional low-resource language pairs and translation settings.

AI use statement

We used LLMs to assist with grammar and writing polishing. All equations, analysis, and research contributions are our own.

REFERENCES

Xiangen Bai, Di Liu, and Xiaofeng Xu. A review of improved methods for ant colony optimization in path planning. *Journal of Ship Research*, 68(02):77–92, 2024.

Justin Chen, Swarnadeep Saha, and Mohit Bansal. Reconcile: Round-table conference improves reasoning via consensus among diverse llms. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 7066–7085, 2024a.

Pinzhen Chen, Zhicheng Guo, Barry Haddow, and Kenneth Heafield. Iterative translation refinement with large language models. In *Proceedings of the 25th Annual Conference of the European Association for Machine Translation (Volume 1)*, pp. 181–190, 2024b.

Xin Cheng, Xun Wang, Tao Ge, Si-Qing Chen, Furu Wei, Dongyan Zhao, and Rui Yan. Scale: Synergized collaboration of asymmetric language translation engines. In *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 15903–15918, 2024.

Junguo Cui, Lei Wu, Xiaodong Huang, Dengpan Xu, Chao Liu, and Wensheng Xiao. Multi-strategy adaptable ant colony optimization algorithm and its application in robot path planning. *Knowledge-Based Systems*, 288:111459, 2024.

Yilun Du, Shuang Li, Antonio Torralba, Joshua B. Tenenbaum, and Igor Mordatch. Improving factuality and reasoning in language models through multiagent debate. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 11733–11763. PMLR, 2024.

Micha Elsner et al. Shortcomings of llms for low-resource translation: Retrieval and understanding are both the problem. In *Proceedings of the Ninth Conference on Machine Translation*, pp. 1332–1354, 2024.

Zhaopeng Feng, Jiayuan Su, Jiamei Zheng, Jiahan Ren, Yan Zhang, Jian Wu, Hongwei Wang, and Zuozhu Liu. M-mad: Multidimensional multi-agent debate for advanced machine translation evaluation. In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 7084–7107, 2025.

Markus Freitag, David Grangier, Qijun Tan, and Bowen Liang. High quality rather than high model probability: Minimum bayes risk decoding with neural metrics. *Transactions of the Association for Computational Linguistics*, 10:811–825, 2022.

Ekaterina Garmash and Christof Monz. Ensemble learning for multi-source neural machine translation. In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pp. 1409–1418, 2016.

Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. Llm-blender: Ensembling large language models with pairwise ranking and generative fusion. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 14165–14178, 2023.

Melvin Johnson, Mike Schuster, Quoc Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. Google's multilingual neural machine translation system: Enabling zero-shot translation. *Transactions of the Association for Computational Linguistics*, 5:339–351, 2017.

Ahrii Kim. Multi-agentmt: Deploying ai agent in the wmt25 shared task. In *Proceedings of the Tenth Conference on Machine Translation*, pp. 769–777, 2025.

Yunxuan Li, Yibing Du, Jiageng Zhang, Le Hou, Peter Grabowski, Yeqing Li, and Eugene Ie. Improving multi-agent debate with sparse communication topology. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 7281–7294, 2024.

Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. Encouraging divergent thinking in large language models through multi-agent debate. In *Proceedings of the 2024 conference on empirical methods in natural language processing*, pp. 17889–17904, 2024.

Chao Liu, Lei Wu, Wensheng Xiao, Guangxin Li, Dengpan Xu, Jingjing Guo, and Wentao Li. An improved heuristic mechanism ant colony optimization algorithm for solving path planning. *Knowledge-based systems*, 271:110540, 2023.

Yinquan Lu, Wenhao Zhu, Lei Li, Yu Qiao, and Fei Yuan. Llamax: Scaling linguistic horizons of llm by enhancing translation capabilities beyond 100 languages. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, pp. 10748–10772, 2024.

Rico Sennrich and Biao Zhang. Revisiting low-resource neural machine translation: A case study. In *Proceedings of the 57th annual meeting of the association for computational linguistics*, pp. 211–221, 2019.

Rico Sennrich, Barry Haddow, and Alexandra Birch. Improving neural machine translation models with monolingual data. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 1: long papers)*, pp. 86–96, 2016a.

Rico Sennrich, Barry Haddow, and Alexandra Birch. Neural machine translation of rare words with subword units. In *Proceedings of the 54th annual meeting of the association for computational linguistics (volume 1: long papers)*, pp. 1715–1725, 2016b.

Andries Smit, Paul Duckworth, Nathan Grinsztajn, Thomas D Barrett, and Arnu Pretorius. Should we be going mad? a look at multi-agent debate strategies for llms. *arXiv preprint arXiv:2311.17371*, 2023.

Felix Stahlberg and Bill Byrne. Unfolding and shrinking neural machine translation ensembles. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pp. 1946–1956, 2017.

Tong Su, Xin Peng, Sarubi Thillainathan, David Guzmán, Surangika Ranathunga, and En-Shiun Lee. Unlocking parameter-efficient fine-tuning for low-resource language translation. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pp. 4217–4225, 2024.

Giorgos Vernikos and Andrei Popescu-Belis. Don't rank, combine! combining machine translation hypotheses using quality estimation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 12087–12105, 2024.

Lei Wu, Xiaodong Huang, Junguo Cui, Chao Liu, and Wensheng Xiao. Modified adaptive ant colony optimization algorithm and its application for solving path planning of mobile robot. *Expert Systems with Applications*, 215:119410, 2023.

Minghao Wu, Jiahao Xu, and Longyue Wang. Transagents: Build your translation company with language agents. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pp. 131–141, 2024.

Wenda Xu, Daniel Deutsch, Mara Finkelstein, Juraj Juraska, Biao Zhang, Zhongtao Liu, William Yang Wang, Lei Li, and Markus Freitag. Llmrefine: Pinpointing and refining large language models via fine-grained actionable feedback. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pp. 1429–1445, 2024.

Chen Zhang, Xiao Liu, Jiuheng Lin, and Yansong Feng. Teaching large language models an unseen language on the fly. In *Findings of the Association for Computational Linguistics: ACL 2024*, pp. 8783–8800, 2024.

Barret Zoph, Deniz Yuret, Jonathan May, and Kevin Knight. Transfer learning for low-resource neural machine translation. In *Proceedings of the 2016 conference on empirical methods in natural language processing*, pp. 1568–1575, 2016.