

RouteMoA: Dynamic Routing without Pre-Inference Boosts Efficient Mixture-of-Agents

Jize Wang¹ Han Wu¹ Zhiyuan You² Yiming Song¹ Yijun Wang³ Zifei Shan³
Yining Li⁴ Songyang Zhang⁴ Xinyi Le^{1*} Cailian Chen^{1*} Xinping Guan¹ Dacheng Tao⁵

¹ Shanghai Jiao Tong University ² CUHK ³ Tencent
⁴Shanghai AI Laboratory ⁵Nanyang Technological University
{jizewang2000, lexinyi, cailianchen}@sjtu.edu.cn

Abstract

Mixture-of-Agents (MoA) improves LLM performance through layered collaboration, but its dense topology raises costs and latency. Existing methods employ LLM judges to filter responses, yet still require all models to perform inference before judging, failing to cut costs effectively. They also lack model selection criteria and struggle with large model pools, where full inference is costly and can exceed context limits. To address this, we propose **RouteMoA**, an efficient mixture-of-agents framework with dynamic routing. It employs a lightweight *scorer* to perform initial screening by predicting coarse-grained performance from the query, narrowing candidates to a high-potential subset without inference. A *mixture of judges* then refines these scores through lightweight self- and cross-assessment based on existing model outputs, providing posterior correction without additional inference. Finally, a *model ranking* mechanism selects models by balancing performance, cost, and latency. RouteMoA outperforms MoA across varying tasks and model pool sizes, reducing cost by 89.8% and latency by 63.6% in the large-scale model pool. Code is available at <https://github.com/Jize-W/RouteMoA>.

1 Introduction

Large Language Models (LLMs) (Ouyang et al., 2022; Zhao et al., 2023) demonstrate strong capabilities across diverse tasks (Zhang et al., 2025; Wang et al., 2023, 2024a). While general-purpose models (e.g., Llama-3.1 (Grattafiori et al., 2024), Qwen2.5 (Yang et al., 2024a)) show broad competence, specialized variants (e.g., Qwen2.5-Math (Yang et al., 2024b), Qwen2.5-Coder (Hui et al., 2024)) excel in specific domains. This diversity in expertise makes the effective integration of multiple LLMs a promising direction to achieve performance beyond individual models.

*Corresponding Authors.




Figure 1: **Significant variations in model capabilities.** Values are normalized to [0,1]. Models exhibit clear specialization: Qwen2.5-Coder leads in coding but lags in biomedical tasks; Qwen2.5-Math excels in mathematics but struggles elsewhere; Bio-Medical-Llama dominates in biomedical knowledge but performs poorly in math and coding; Gemma stands out in reasoning and reading. These distinct profiles make it feasible to predict model performance only based on specific user queries.

Among various LLM-based multi-agent collaboration strategies, Mixture-of-Agents (MoA) (Wang et al., 2024b) is a typical and effective approach. As shown in Figure 2(a), this method enables multiple LLMs to refer to each other’s responses, engaging in iterative rounds (*i.e.*, layers in Figure 2) of replies and summaries to achieve results superior to those of a single model.

Despite the advantages, MoA-based methods are **highly resource-intensive**. As shown in Figure 2(a), classical MoA (Wang et al., 2024b) requires forwarding multiple LLMs per layer and concatenating all outputs as the input to the next, leading to high cost and latency. Sparse MoA (Li et al., 2024) (Fig. 2(b)) introduces a judge to filter responses, yet still invokes all LLMs plus an additional judge model, further increasing overhead. These approaches also lack principled model selection and do not scale to large pools (e.g., >10 models), as full inference becomes prohibitively costly and often exceeds context limits.

To address the efficiency challenge, we propose RouteMoA, a dynamically-routed mixture-of-agents framework. Our approach is motivated by the complementary capabilities of LLMs (Figure 1): for example, Qwen2.5-Math excels in mathemat-

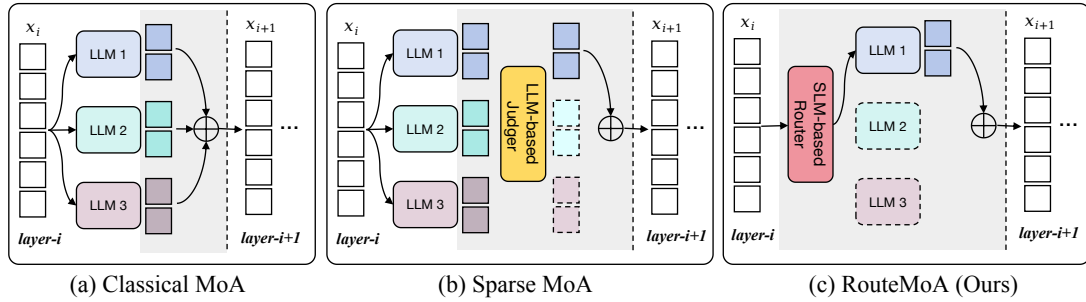


Figure 2: **Concept comparison** between our RouteMoA and previous MoA-based methods. (a) Classical MoA (Wang et al., 2024b) forwards all LLMs in each layer, and concatenates all outputs as the input of the next layer. (b) Sparse MoA (Li et al., 2024) introduces an LLM-based judge to select some good responses as the input of the next layer. This reduces the number of input tokens, but still needs to forward all LLMs and another LLM-based judge. (c) RouteMoA uses a lightweight router to select parts of LLMs for inference, significantly reducing computational cost.

ics but underperforms in reasoning and biomedical tasks. Such specialization makes it feasible to predict model performance from the query, thus narrowing the initial pool to a few high-potential candidates and reducing cost.

Specifically, RouteMoA leverages a lightweight *scorer* that performs initial screening. Using only prior knowledge from the query, it estimates model suitability without executing inference. It assigns coarse-grained scores to identify promising candidates, enabling activation of only a subset of models and significantly lowering inference overhead.

To correct potential scoring errors, we introduce a *mixture of judges* to combine the scorer with self- and cross-assessment. These judges operate post-hoc, leveraging posterior knowledge from previously-generated responses without requiring additional inference. This design enhances assessment reliability at no extra cost, ensuring robust model selection throughout the routing process. Finally, a *model ranking* mechanism selects models by balancing performance, cost, and latency.

In summary, our contributions are as follows:

- We propose RouteMoA, a dynamically-routed MoA framework that significantly cuts cost and latency while maintaining strong performance.
- We design a lightweight scorer for initial model screening based on query-aware prior knowledge, narrowing the candidate pool to a few high-potential models without pre-inference.
- We introduce a mixture of judges that refines model scores through self- and cross-assessment, leveraging posterior knowledge from model outputs to correct prediction errors without introducing additional inference overhead.
- Extensive experiments on both small- and large-

scale model pools, along with out-of-distribution tasks, show RouteMoA matches or surpasses strong baselines in accuracy while greatly boosting efficiency and scalability.

2 Related Work

General and task-specific LLMs. Large Language Models (LLMs) have shown strong performance in text understanding and generation (Achiam et al., 2023; Cai et al., 2024; Grattafiori et al., 2024; Yang et al., 2024a). They can be categorized into general-purpose models—such as Llama-3.1 (AI, 2025), Qwen2.5 (Yang et al., 2024a), Mistral (Jiang et al., 2023), and Gemma (Team et al., 2023)—and domain-specific fine-tuned variants like Qwen2.5-Math (Yang et al., 2024b), Qwen2.5-Coder (Hui et al., 2024), and Bio-Medical-Llama (ContactDoctor, 2024). While specialized models excel in their domains, they often underperform elsewhere. For instance, Bio-Medical-Llama achieves 87.0 on the MMLU biomedical subset, outperforming Qwen2.5-Math by 43.5 points, yet scores only 11.7 on MATH. Since developing a universally capable model is costly, integrating multiple models’ strengths offers a more viable path.

LLM-based multi-agent collaboration. Multi-agent frameworks provide effective ways to leverage diverse model capabilities. Majority voting (Chen et al., 2024a) selects the most frequent answer from multiple models as the final output. LLM cascading (Yue et al., 2024) sequentially invokes models until a response meets a quality threshold. Multi-agent debate (Liang et al., 2024) enhances accuracy through iterative discussion. Mixture-of-Agents (MoA) (Wang et al., 2024b)

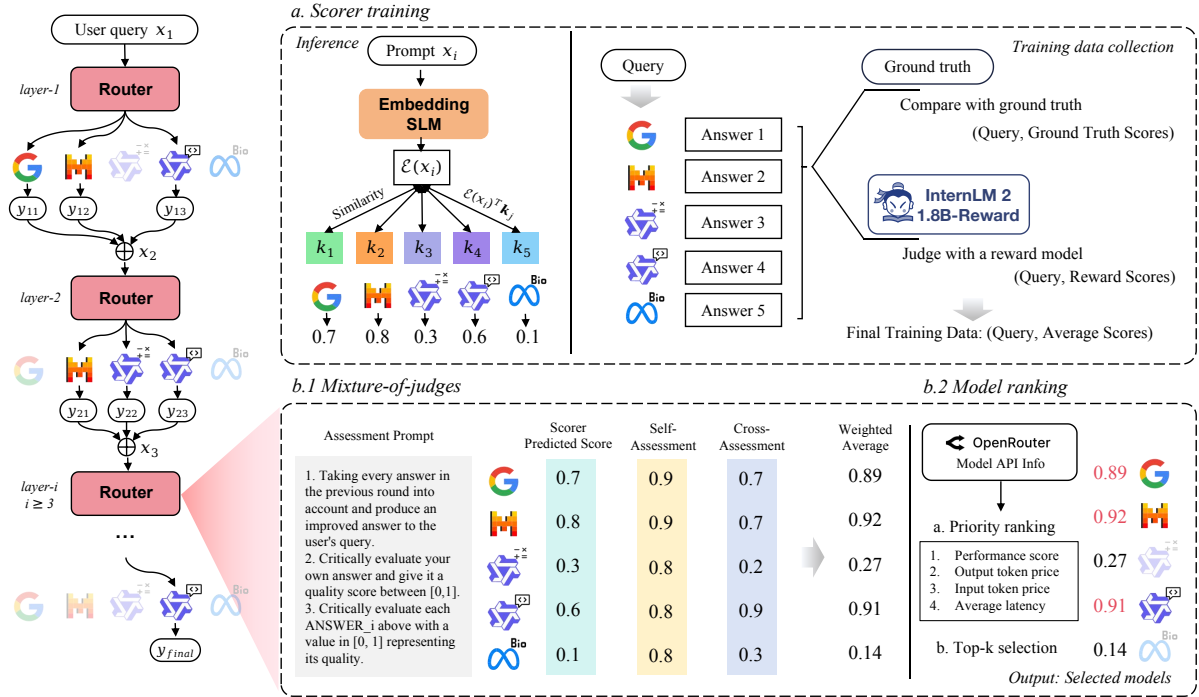


Figure 3: **RouteMoA architecture**. The framework operates layer-wise (left). At each layer l , the router selects a subset of suitable LLMs, whose outputs are aggregated and passed to the next layer. The router (right) consists of two stages: *b.1 Mixture of Judges*, which includes a scorer (trained as in *a. Scorer Training*), self-assessment, and cross-assessment. The scorer predicts candidate performance in layer-1 using prior knowledge from the query; subsequent layers refine scores via self- and cross-assessment using posterior knowledge from model outputs. *b.2 Model Ranking* selects LLMs by balancing performance, cost, and latency.

refines answers via multi-round parallel reasoning but incurs high computational cost. Sparse MoA (Li et al., 2024) introduces a judge to filter responses, saving input tokens but still requiring inference from all models. In contrast, our approach adopts a lightweight router that selects suitable models dynamically for each layer without pre-inference, significantly cutting cost and latency.

LLM routing. LLM routing (Ding et al., 2024; Stripelis et al., 2024) selects the best-performing model per query without invoking all candidates. Meta-models are trained to predict model performance based on input, improving cost efficiency. Benchmarks such as RouterBench (Hu et al., 2024) and RouterEval (Huang et al., 2025) assess routing effectiveness. ZOOTER (Lu et al., 2023) distills reward signals into an SLM router via KL-divergence, while RouterDC (Chen et al., 2024b) uses dual contrastive loss for better accuracy. Eagle adopts a training-free approach using similarity-based retrieval. RouteLLM (Ong et al., 2024) focuses on binary routing between strong and weak models to minimize expensive calls. In contrast to routing methods that rely solely on query-based prior knowledge, our approach also leverages posterior knowledge from actual model outputs to update

performance scores. This design relaxes the requirement for precise performance prediction, and the subsequent multi-agent collaboration further enhances robustness and overall performance beyond what is achievable by routing to a single model.

3 Methodology

In this section, we introduce RouteMoA, an efficient mixture-of-agents framework with dynamic routing. It dynamically selects a subset of top-performing LLMs each round without pre-inference, thus reducing cost and latency while maintaining performance. We first overview the whole routing process in Section 3.1, then describe its key components: the scorer, mixture of judges, and model ranking (Sections 3.2–3.4).

3.1 Overview

The framework operates layer-wise, as shown in the left of Figure 3, following the Mixture-of-Agents structure (Wang et al., 2024b). It consists of L layers. In intermediate layers ($l = 1, 2, \dots, L - 1$), n_l LLMs act as proposers $M_{l,i} \in \mathcal{P}$, where \mathcal{P} is a pool of N available models, $i = 1, 2, \dots, n_l \leq N$. Each $M_{l,i}$ processes input x_l and generates a

response:

$$y_{l,i} = M_{l,i}(x_l). \quad (1)$$

The output of layer l is:

$$o_l = \oplus_{i=1}^{n_l} y_{l,i} + x_1, \quad x_{l+1} = o_l, \quad (2)$$

where $+$ denotes concatenation and \oplus denotes an aggregation prompt (see Appendix A). x_1 is the user query. The final layer L uses a single LLM to aggregate prior outputs into the final response.

To balance performance and efficiency, RouteMoA dynamically selects models for each layer $l = 1, 2, \dots, L$ through the following process:

Step 1: Score Acquisition. For the first layer ($l = 1$), an *SLM-based scorer* \mathcal{S} performs an initial screening by predicting coarse-grained performance scores for each model in \mathcal{P} on query x_1 :

$$\mathbf{s}_1 = \mathcal{S}(x_1), \quad \mathbf{s}_1 \in [0, 1]^N. \quad (3)$$

This scorer is not required to provide precise performance score estimates; rather, its goal is to efficiently narrow down the candidate set to a small group of high-potential models. For subsequent layers ($l > 1$), a *mixture of judges* \mathcal{J} refines the initial scores by incorporating both \mathbf{s}_1 and responses from the previous layer, enabling more accurate and context-aware model selection:

$$\mathbf{s}_l = \mathcal{J}(\mathbf{s}_1, y_{l-1,1}, \dots, y_{l-1,n_{l-1}}), \quad l > 1. \quad (4)$$

Step 2: Model Ranking and Selection. The *model ranking* module \mathcal{R} selects active models for layer l based on performance \mathbf{s}_l , cost, and latency:

$$[M_{l,1}, \dots, M_{l,n_l}] = \mathcal{R}(\mathbf{s}_l, \text{cost}, \text{latency}). \quad (5)$$

An early-stopping mechanism determines when to terminate, ensuring efficient inference.

3.2 SLM-based Scorer

The scorer conducts an initial screening by predicting coarse-grained performance scores to each model in the pool $\mathcal{P} = M_1, M_2, \dots, M_N$ given the input x_l , as defined in Equation 3.

Dataset generation. We construct a training dataset of the form:

$$\mathcal{D} = \{(x^{(k)}, s_1^{(k)}, s_2^{(k)}, \dots, s_N^{(k)})\}_{k=1}^{|\mathcal{D}|}, \quad (6)$$

where $s_j^{(k)}$ denotes the performance score of model M_j on input $x^{(k)}$. To build \mathcal{D} , we collect queries and ground-truth answers $\mathcal{D}_{raw} = \{(x^{(k)}, \hat{y}^{(k)})\}_{k=1}^{|\mathcal{D}|}$ from multiple datasets spanning

mathematics, reasoning, coding, reading comprehension, biomedical domains, etc. For each query, we gather responses from all models in \mathcal{P} :

$$\mathcal{D}_y = \{(x^{(k)}, y_1^{(k)}, y_2^{(k)}, \dots, y_N^{(k)})\}_{k=1}^{|\mathcal{D}|}. \quad (7)$$

Each response is scored using a combination of ground-truth accuracy and a reward model R (e.g., InternLM2-1.8B-Reward):

$$s_j^{(k)} = \lambda \cdot \mathbf{1}(\hat{y}^{(k)} = y_j^{(k)}) + (1 - \lambda) \cdot R(x^{(k)}, \hat{y}^{(k)}, y_j^{(k)}), \quad (8)$$

where $\lambda \in (0, 1)$. $\mathbf{1}(\cdot)$ is the indicator function that returns 1 if the condition is true and 0 otherwise.

Model structure. Inspired by matrix factorization techniques in recommendation systems (Chen et al., 2024b; Koren et al., 2009; Ong et al., 2024), we model the scorer as an embedding-based similarity function. Each model M_j is assigned a learnable embedding $\mathbf{k}_j \in \mathbb{R}^d$. The input x is encoded by a small language model mDeBERTaV3-base (He et al., 2021) into an embedding $\mathcal{E}(x)$. The performance score is computed as:

$$s = f(x, M_j) = \sigma(\mathcal{E}(x)^\top \mathbf{k}_j), \quad s \in [0, 1], \quad (9)$$

where $\sigma(\cdot)$ is the sigmoid function. Thus the full score vector is:

$$\mathbf{s} = \mathcal{S}(x) = [f(x, M_1), f(x, M_2), \dots, f(x, M_N)]. \quad (10)$$

Training and inference. During training, we adopt dual contrastive loss functions from (Chen et al., 2024b). The *sample-LLM contrastive loss* ensures that the embeddings of models capable of answering a query are closer to the query’s embedding:

$$\begin{aligned} \mathcal{L}_{\text{sample-LLM}}(x, \mathbf{s}; \theta) &= \sum_{j_+ \in \mathcal{I}^+} -\log \frac{e^{\mathcal{E}(x)^\top \mathbf{k}_{j_+}}}{e^{\mathcal{E}(x)^\top \mathbf{k}_{j_+}} + \sum_{j_- \in \mathcal{I}_i^-} e^{\mathcal{E}(x)^\top \mathbf{k}_{j_-}}}, \end{aligned} \quad (11)$$

where \mathcal{I}^+ and \mathcal{I}_i^- denote the top- K_+ and bottom- K_- scoring models, respectively. θ denotes the parameters to be optimized.

The *sample-sample contrastive loss* encourages semantically similar queries to have closer embeddings. It is formulated as:

$$\begin{aligned} \mathcal{L}_{\text{sample-sample}}(x; \theta) &= -\log \frac{e^{\mathcal{E}(x)^\top \mathcal{E}(x^+)}}{e^{\mathcal{E}(x)^\top \mathcal{E}(x^+)} + \sum_{x_i^- \in \mathcal{X}_i^-} e^{\mathcal{E}(x)^\top \mathcal{E}(x^-)}}. \end{aligned} \quad (12)$$

where x^+ is a query from the same cluster as x , and \mathcal{X}^- contains out-cluster queries. The clustering method is detailed in the Appendix C.

The total loss is:

$$\mathcal{L} = \mathcal{L}_{\text{sample-LLM}} + \alpha \mathcal{L}_{\text{sample-sample}}, \quad (13)$$

where $\alpha \geq 0$.

3.3 Mixture of Judges

The design of the mixture of judges is motivated by two key capabilities of large language models:

- *Self-knowledge awareness*: Research has shown that LLMs possess the ability to evaluate their own knowledge and determine whether they understand a question (Kadavath et al., 2022).
- *Cross-model evaluation*: LLMs can effectively judge responses from other models (Li et al., 2025), making them capable evaluators in multi-agent settings.

Thus, for layer $l > 1$, we introduce mixture of judges to refine the scorer’s predictions using self- and cross-assessment signals from previous layers.

For *self-assessment*, each active model in layer $l - 1$ outputs a confidence score $s_{l-1,j}^{\text{self}}$ along with its response:

$$\mathbf{s}_{l-1}^{\text{self}} = [s_{l-1,1}^{\text{self}}, s_{l-1,2}^{\text{self}}, \dots, s_{l-1,n_{l-1}}^{\text{self}}]. \quad (14)$$

For *cross-assessment*, to avoid the computational cost of having all models generate evaluation scores, we selectively employ only the highest-scoring model from layer $l - 1$ to evaluate responses from layer $l - 2$, producing scores $s_{l-2,j}^{\text{cross}}$:

$$\mathbf{s}_{l-2}^{\text{cross}} = [s_{l-2,1}^{\text{cross}}, s_{l-2,2}^{\text{cross}}, \dots, s_{l-2,n_{l-2}}^{\text{cross}}]. \quad (15)$$

Since cross-assessment relies on evaluating outputs from a prior layer, it is only applicable from the second layer onward ($l \geq 2$), as no prior outputs exist for the first layer.

The final mixture of judges function is:

$$\mathbf{s}_l = \mathcal{J}(\mathbf{s}_1, y_{l-1,1}, y_{l-1,2}, \dots, y_{l-1,n_{l-1}}) \quad (16)$$

$$= \begin{cases} \mathcal{U}(\mathbf{s}_1, \mathbf{s}_{l-1}^{\text{self}}), & l = 2, \\ \mathcal{U}(\mathbf{s}_1, \mathbf{s}_{l-1}^{\text{self}}, \mathbf{s}_{l-2}^{\text{cross}}), & l > 2, \end{cases} \quad (17)$$

where \mathcal{U} performs score normalization followed by element-wise averaging.

3.4 Model Ranking

The model ranking module \mathcal{R} selects the top- k models based on the adjusted scores \mathbf{s}_l , with the following priority: performance > output token cost > input token cost > latency. Model pricing and latency data are sourced from OpenRouter¹. Early stopping criterion is set as:

$$\max(s_{l,1}, s_{l,2}, \dots, s_{l,N}) > s_{th}. \quad (18)$$

where s_{th} is a threshold score.

If the criterion is met, or the max layer number is reached, the system will enter the aggregation stage and produce the final output:

$$y^{\text{final}} = M_{l,\text{agg}}(x_l). \quad (19)$$

4 Experiments

4.1 Experimental Setup

Baselines. We focus on improving the computational efficiency of multi-agent collaboration while maintaining the accuracy. The compared baselines include: (1) MoA (Wang et al., 2024b), leveraging multiple LLMs in a layered architecture, where each agent uses outputs from previous layers to enhance its response generation; (2) SMOA (Li et al., 2024), improving the token efficiency of MoA by employing a judge model to assess and forward only the most optimal responses to the next round. (3) To explore the impact of self-assessment and cross-assessment, we also compare RouteMoA with the version that without self-assessment and without cross-assessment for ablation study.

Implementation Details. We use OpenCompass (Contributors, 2023) for data generation and evaluation. For scorer training, we employ mDeBERTaV3-base (He et al., 2021) as the encoder, a small language model with only 86M parameters. Each LLM embedding is projected to a 768-dimensional vector space. The training parameters are set as $\alpha = 0.2$ and $\lambda = 0.5$, which are observed to be insensitive within the ranges of $[0.2, 2]$ and $[0.3, 0.9]$, respectively. The number of k-means clusters is set to 6. Training is conducted using the AdamW optimizer with a learning rate of 5×10^{-5} , a weight decay of 0.01, and a mini-batch size of 64. We report average performance, cost, and latency. Experiments are run on 80GB GPUs.

Exp1: Scalability Evaluation on Large-Scale Model Pool. To validate the practical scalability

¹<https://openrouter.ai/>

Table 1: **Performance and efficiency comparison on the large-scale model pool (15 LLMs).** \uparrow indicates an improvement over MoA, while \downarrow represents a degradation compared to MoA. Both are denoted by percentage.

	Method	Language Understanding	Reading&QA	Logic Reasoning	Math Reasoning	Language Generation	Avg.
Accuracy (%) \uparrow	MoA	83.4	88.0	93.3	49.7	41.9	71.3
	SMoA	78.4 $\downarrow 6.00\%$	85.3 $\downarrow 3.07\%$	91.1 $\downarrow 2.36\%$	53.1 $\uparrow 6.84\%$	40.3 $\downarrow 3.82\%$	69.7 $\downarrow 2.24\%$
	RouteMoA	84.0 $\uparrow 0.72\%$	88.0 $\uparrow 0.00\%$	95.6 $\uparrow 2.50\%$	73.3 $\uparrow 47.5\%$	51.9 $\uparrow 23.9\%$	78.6 $\uparrow 10.2\%$
Cost (\$) \downarrow	MoA	321.7	303.4	385.3	751.1	477.5	447.8
	SMoA	47.8 $\uparrow 85.1\%$	53.8 $\uparrow 82.3\%$	57.2 $\uparrow 85.2\%$	232.5 $\uparrow 69.0\%$	110.4 $\uparrow 76.9\%$	100.4 $\uparrow 77.6\%$
	RouteMoA	24.9 $\uparrow 92.3\%$	14.2 $\uparrow 95.3\%$	28.7 $\uparrow 92.6\%$	94.6 $\uparrow 87.4\%$	65.7 $\uparrow 86.2\%$	45.6 $\uparrow 89.8\%$
Latency (s) \downarrow	MoA	126.3	101.4	134.2	619.5	258.9	248.1
	SMoA	76.4 $\uparrow 39.5\%$	94.1 $\uparrow 72.0\%$	76.1 $\uparrow 43.3\%$	471.2 $\uparrow 23.9\%$	257.1 $\uparrow 0.70\%$	195.0 $\uparrow 21.4\%$
	RouteMoA	43.4 $\uparrow 65.6\%$	27.8 $\uparrow 72.6\%$	58.9 $\uparrow 56.1\%$	211.4 $\uparrow 65.9\%$	109.3 $\uparrow 57.8\%$	90.2 $\uparrow 63.6\%$

Table 2: **Performance and efficiency comparison on the small-scale model pool (5 LLMs).** *Oracle* means using ground truth assessment scores for LLM selection. Best results of multi-LLM methods are bold. A paired t-test confirms that the improvement of RouteMoA over SMoA is statistically significant ($t = 2.296$, $p = 0.0217 < 0.05$).

Accuracy (%) \uparrow												
Type	Method	MATH	ARC-c	MBPP	RACE-high	MMLU-bio	Avg.					
Single LLM	Gemma-2-9B-it	46.5	90.2	66.2	85.6	78.6	75.9					
	Minstral-8B-Instruct-2410	51.0	85.3	63.0	80.3	70.7	72.7					
	Qwen2.5-Coder-7B-Instruct	65.3	85.5	79.8	80.6	67.2	77.5					
	Qwen2.5-Math-7B-Instruct	80.7	50.6	52.9	55.0	43.5	63.0					
	Bio-Medical-Llama-3-8B	11.7	75.1	16.0	77.3	87.0	46.2					
Single LLM with Routing	Oracle	83.8	96.8	86.7	94.5	95.6	92.5					
	RouteLLM	64.3	84.8	76.3	79.4	65.7	76.2					
	RouterDC	72.8	87.3	72.6	78.3	70.9	78.9					
Multi-LLMs	MoA	73.6	87.0	75.5	80.1	76.0	80.9					
	SMoA	73.5 $\downarrow 0.10\%$	89.4 $\uparrow 2.80\%$	79.4 $\uparrow 5.20\%$	84.0 $\uparrow 4.90\%$	75.7 $\downarrow 0.40\%$	82.6 $\uparrow 2.10\%$					
Ours	RouteMoA	76.0 $\uparrow 3.30\%$	88.2 $\uparrow 1.40\%$	79.8 $\uparrow 5.70\%$	81.0 $\uparrow 1.10\%$	79.3 $\uparrow 4.30\%$	83.1 $\uparrow 2.70\%$					
Resource	Cost (\$) \downarrow						Latency (s) \downarrow					
Dataset	MATH	ARC-c	MBPP	RACE-high	MMLU-bio	Total	MATH	ARC-c	MBPP	RACE-high	MMLU-bio	Avg.
MoA	19.68	2.27	0.61	8.53	1.78	36.03	26.62	12.05	15.52	13.45	14.07	16.32
SMoA	4.40 $\uparrow 77.6\%$	0.47 $\uparrow 79.3\%$	0.18 $\uparrow 70.5\%$	2.22 $\uparrow 74.0\%$	0.36 $\uparrow 79.8\%$	8.23 $\uparrow 77.2\%$	23.16 $\uparrow 13.0\%$	10.45 $\uparrow 13.3\%$	10.30 $\uparrow 33.6\%$	11.02 $\uparrow 18.1\%$	11.86 $\uparrow 15.7\%$	13.31 $\uparrow 18.4\%$
RouteMoA	4.03 $\uparrow 79.5\%$	0.28 $\uparrow 87.7\%$	0.37 $\uparrow 39.3\%$	1.82 $\uparrow 78.7\%$	0.21 $\uparrow 88.2\%$	6.71 $\uparrow 81.4\%$	19.05 $\uparrow 28.4\%$	9.73 $\uparrow 19.3\%$	7.31 $\uparrow 52.9\%$	4.45 $\uparrow 66.9\%$	9.51 $\uparrow 32.4\%$	10.01 $\uparrow 38.7\%$

and efficiency of RouteMoA in real-world deployment scenarios, we construct a large-scale model pool consisting of **15 state-of-the-art LLMs** of varying sizes (from 4B to 235B parameters) and capabilities, including general-purpose, reasoning-specialized, and code/math-focused models (see Table 12 in Appendix). Notably, this pool contains models with both standard (*no-think*) and advanced reasoning (*think*) modes, presenting a diverse and challenging testbed for multi-agent collaboration.

Evaluation Benchmark. We conduct a comprehensive evaluation on a collection of **30 datasets** spanning five critical capability categories: Language Understanding, Reading & QA, Logic Reasoning, Math Reasoning, and Language Generation (see Table 14, 15 for the full list). This broad coverage ensures a rigorous assessment of generalizability.

Exp2: Performance on Small-Scale Model Pool. To enable a direct and fair comparison with MoA and SMoA (which are limited to small pools due to their full-model inference design), we further evaluate on a compact but diverse pool of 5 LLMs:

Gemma-2-9B-it (Team et al., 2024), Minstral-8B-Instruct (Jiang et al., 2023), Qwen2.5-Coder-7B-Instruct (Hui et al., 2024), Qwen2.5-Math-7B-Instruct (Yang et al., 2024b), and Bio-Medical-Llama-3-8B (ContactDoctor, 2024). The evaluation covers 5 datasets (MATH-500 (Hendrycks et al., 2021), ARC-Challenge (Clark et al., 2018), MBPP (Austin et al., 2021), RACE-high (Lai et al., 2017), MMLU-bio (Hendrycks et al., 2020)) representing mathematics, reasoning, coding, reading, and biomedical knowledge.

Exp3: Out-of-Distribution Generalization. We further evaluate generalization on the challenging AGIEval-Gaokao (Zhong et al., 2024) benchmark, which spans nine subjects (Biology, Chemistry, Chinese, English, Geography, History, MathCloze, MathQA, Physics). This human-exam benchmark tests the model’s ability to handle diverse, unseen tasks requiring human-like reasoning.

4.2 Main Results

Scalability Evaluation on Large-Scale Model Pool. As shown in Table 1, RouteMoA demon-

Table 3: **Out-of-distribution benchmark comparison** between SMOA and RouteMoA.

	Method	Biology	Chemistry	Chinese	English	Geography	History	MathCloze	MathQA	Physics	OOD Avg.
Accuracy (%) ↑	SMoA	53.33	37.68	49.59	80.39	60.80	64.26	27.12	64.10	39.00	52.92
	RouteMoA	58.10	37.68	49.19	77.78	67.84	69.36	27.12	60.97	43.50	54.62
Cost (\$) ↓	SMoA	4.71	6.90	8.05	5.25	3.80	3.49	8.79	9.04	7.31	6.37
	RouteMoA	4.18	7.77	6.33	3.46	4.15	3.33	6.08	7.57	7.91	5.64
Latency (s) ↓	SMoA	11.25	15.41	10.90	9.40	9.79	9.86	19.09	19.85	15.91	13.50
	RouteMoA	6.12	14.46	4.57	2.71	5.59	4.26	17.72	20.08	15.93	10.16

Table 4: **Ablation study** for mixture of judges on small-scale model pool.

Method	Performance(%) ↑	Cost(\$)	Latency(s) ↓
RouteMoA	83.1	7.68	10.64
w/o self.	82.6	7.99	10.49
w/o cross.	82.7	7.25	10.29

strates exceptional scalability, performance, and efficiency. It achieves an average accuracy of 78.6, significantly surpassing MoA (71.3) and SMOA (69.7), with especially large gains in Math Reasoning (+47.5%) and Language Generation (+23.9%). Unlike MoA and SMOA, which lack a clear model selection criteria and become infeasible at scale due to prohibitive costs and context limits, RouteMoA remains practical by dynamically routing queries to an optimal model subset. This approach reduces total cost by 89.8% and latency by 63.6% compared to MoA, while also outperforming SMOA in both efficiency and accuracy.

Furthermore, RouteMoA consistently achieves the best accuracy, lowest cost, and lowest latency across all five capability categories. In Language Understanding and Reading&QA, it matches or surpasses MoA’s accuracy while reducing cost by 95.3%. These results demonstrate that dynamic routing tailored for multi-agent systems enables efficient and effective collaboration in large, heterogeneous model pools, particularly for complex tasks requiring complementary model strengths.

Performance on Small-Scale Model Pool. As shown in Table 2, RouteMoA substantially improves efficiency, reducing inference cost by 81.4% compared to MoA (6.71 vs. 36.03), and by 88.2% on domain-specific scenarios such as MMLU-bio, demonstrating effective avoidance of expensive generalist models. It also lowers average latency by 38.7% (10.01s vs. 16.32s) due to lightweight scoring and targeted model selection. Meanwhile, RouteMoA achieves the highest average score (83.1), outperforming single models and MoA, with statistically significant gains over SMOA (paired t-test shows $t = 2.296$, $p = 0.0217$




Figure 4: **Average values of three scorer assessment metrics** (Top-1-Hit, Top-3-Hit, and Top-3-Agree) under different training hyperparameters (λ and α).

< 0.05). These results confirm that routing combined with response aggregation forms an effective paradigm for multi-LLM collaboration.

Out-of-Distribution Generalization. As shown in Table 3, RouteMoA outperforms SMOA with higher average accuracy (54.62 vs. 52.92) while reducing cost by 11.5% and latency by 24.7%. It achieves notable accuracy gains in humanities and science subjects, including Geography (+7.04), History (+5.10), Physics (+4.50), and Biology (+4.77). These results demonstrate that RouteMoA effectively exploits specialized models on unseen tasks requiring human-like reasoning, exhibiting strong out-of-distribution generalization.

4.3 Analysis

Scorer Assessment. In RouteMoA, the scorer plays a key role in providing an initial screening of candidate models. Rather than requiring precise performance prediction, the scorer is designed to identify a small set of high-potential models for subsequent refinement via mixture-of-judges (including self-assessment and cross-assessment). To evaluate its effectiveness, we introduce three metrics: Top-1 Hit Rate (Top-1-Hit), Top-3 Hit Rate (Top-3-Hit) and Top-3 Agreement Rate (Top-3-Agree). Top-1-Hit and Top-3-Hit measure the probability that the ground-truth best model appears in the scorer’s top-one / top-three predictions, which

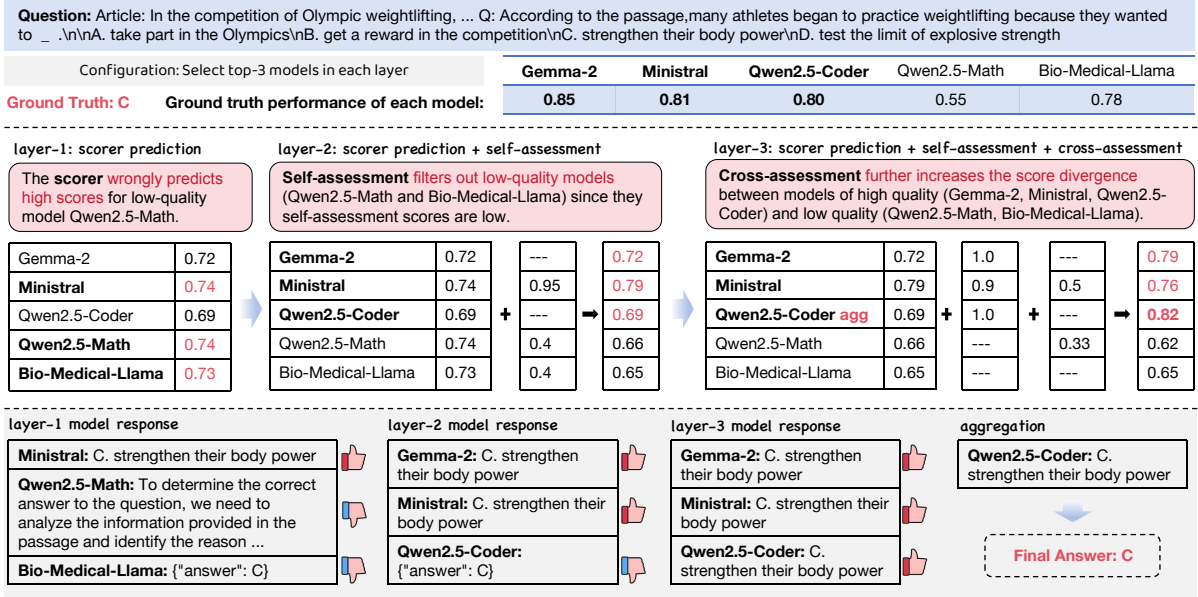


Figure 5: **Case study** of adjusting wrong scorer predictions with self- and cross-assessment.

reaches 90.7% and 97.9% when $\alpha = 0.2$, $\lambda = 0.5$. The Top-3-Agree quantifies the overlap between the scorer’s top-three selections and the ground-truth top-three models, achieving 96.2%, indicating that the scorer successfully narrows down the candidate pool to a small subset containing high-performing models in the majority of cases. The detailed calculation method and results under different α , λ combinations are shown in Appendix D.

Mixture-of-Judges Ablation Study. To explore the impact of self-assessment and cross-assessment in mixture of judges, we conduct ablation studies. As shown in Table 4, the average performance of RouteMoA without self-assessment is 82.6, and the performance without cross-assessment is 82.7, both of them are lower than the 83.1 achieved by RouteMoA. RouteMoA without cross-assessment achieves the lowest cost and latency. It is a natural result since the additional judging token will not be generated without cross-assessment.

Case Study. To illustrate how the routing pipeline, especially the mixture of judges works, we present an example from the RACE-high dataset in Figure 5. In this dataset, Gemma-2, Ministral, and Qwen2.5-Coder usually perform well, whereas Qwen2.5-Math and Bio-Medical-Llama show weaker performance. In layer-1, the scorer incorrectly assigns a high score (0.74) to Qwen2.5-Math, while giving a relatively low score to the high-quality model Qwen2.5-Coder. In layer-2, models generate self-assessment scores. Qwen2.5-Math and Bio-Medical-Llama produce low self-scores due to their inability to accurately follow

instructions, thus the self-assessment mechanism effectively filters out low-quality models. In layer-3, cross-assessment further widens the score gap between high- and low-quality models, since a low cross-score is assigned to Qwen2.5-Math. Examining the model responses across layers, we observe increasing participation of high-quality models and progressive improvement in response quality as the number of layers increases. Eventually, models achieve consensus at the final layer. This correction process is supported by the high Top-3 Hit Rate (97.9%) of the scorer, indicating that in 97.9% of cases, at least one correct model is included in the initial candidate set. Once present, the multi-agent collaboration mechanism effectively identifies and amplifies the correct response through answer aggregation and self/cross-assessment.

Failure Source Discussion. We analyze 323 error instances where RouteMoA fails while MoA succeeds. From these cases, we identify four primary mechanism-level failure modes, summarized in Table 5. Most failures ($\approx 50\%$) arise not from mis-routing but from late-stage aggregation drift, suggesting that improving consensus robustness may yield the largest future gains.

Reliability of Self-Assessment. We conduct a case-level analysis over 362 successful and 323 failed RouteMoA instances. In successful cases, 84.52% of self-assessment bias cases are corrected before final prediction. The correction sources are shown in Table 6. This indicates that posterior mechanisms effectively mitigate transient self-evaluation bias during collaborative refinement.

Failure Mode	Share (%)	Mechanism-Level Description
Aggregation Error	50.46	The correct candidate exists in the model pool but is overridden during aggregation due to consensus drift, majority dominance, or score misweighting.
Scorer Mis-routing	24.15	The domain-specialized model is excluded from the initial Top- k candidate set, indicating recall failure of the scorer.
Self-assessment Bias	20.12	Incorrect models assign themselves overly high confidence scores, often due to overconfident reasoning or task misinterpretation.
Cross-assessment Bias	5.26	Errors arise from mutual reinforcement among models or miscalibrated cross-evaluation, leading to collective preference for incorrect answers.

Table 5: Failure source breakdown for cases where RouteMoA fails but MoA succeeds on the small model pool.

Correction Source	Share (%)
Mixed (multi-module interaction, self reweighting)	60.4
Aggregation Only	38.8
Cross-assessment Only	0.7

Table 6: Breakdown of correction sources for self-assessment bias cases.

k	Accuracy(%)	Avg Cost(\$)	Avg Latency(s)
1	81.82	2.72	6.55
2	81.85	4.21	7.09
3	82.88	4.63	7.36
4	81.22	4.76	6.97
5	81.82	4.91	7.49

Table 7: Sensitivity analysis over top- k selection.

In contrast, among failed cases, as shown in Table 5, only 20.12% are primarily attributed to self-assessment bias. The majority of failures (50.46%) arise from late-stage aggregation drift, where correct candidates are overwritten by consensus dynamics, rather than from unreliable self-evaluation. **Sensitivity Analysis.** We conduct sensitivity analysis over top- k , early-stopping threshold s_{th} , and weights of judge aggregation. (1) As shown in Table 7, accuracy peaks at $k = 3$ (82.88%), and larger k yields no consistent gain while slightly increasing cost/latency, supporting $k = 3$ as a balanced default. (2) As shown in Table 8, increasing s_{th} raises average layers and improves accuracy (e.g., 82.83% at $s_{th} = 0.9$) but with higher cost/latency, showing a clear efficiency-performance trade-off. Notably, 97.9% of queries already include the correct model in the initial candidate set, enabling frequent early termination. (3) As shown in Table 9, different weighting schemes lead to moderate accuracy variation (≈ 81.6 – 83.2%) while maintaining similar cost and latency, indicating that RouteMoA’s efficiency is robust to aggregation design, although accuracy can be further tuned via weighting.

Cost Analysis. The offline cost of RouteMoA is one-time, fixed, and lightweight. (1) Training the 86M-parameter scorer requires only ~ 25 min-

s_{th}	# Layers	Accuracy(%)	Avg Cost(\$)	Avg Latency(s)
0.7	1	81.65	3.77	6.34
0.75	1	80.48	3.76	6.43
0.8	1.33	81.21	3.88	7.09
0.85	2.32	82.50	6.64	11.55
0.9	2.91	82.83	8.52	12.78

Table 8: Sensitivity analysis over early-stopping threshold s_{th} and resulting layer depth.

Scorer : Self : Cross	Accuracy(%)	Avg Cost(\$)	Avg Latency(s)
Uniform (1:1:1)	81.58	4.46	6.99
Scorer-Heavy (2:1:1)	83.20	4.39	6.56
Cross-Heavy (1:1:2)	82.09	4.45	6.77
Self-Heavy (1:2:1)	83.07	4.48	7.02

Table 9: Sensitivity analysis over different weighting schemes in the mixture-of-judges.

utes and < 50 GB GPU memory. (2) As shown in Table 13 in Appendix, generating training data (around 4k samples as shown in Table 15 in Appendix) for adding a new LLM is consistently low-cost across models (typically $< \$4$ and within hours under parallel requests). The exact inference time is model-dependent and can be reduced via parallelization, making the offline overhead negligible relative to online efficiency gains.

5 Conclusion

In conclusion, we present RouteMoA, an efficient Mixture-of-Agents framework that overcomes the resource limitations of classical MoA through dynamic routing. The framework employs a lightweight scorer to perform an initial screening of candidates using prior knowledge from the query, followed by a mixture of judges that refines scores with posterior knowledge from model outputs. RouteMoA significantly reduces cost and latency while maintaining strong performance. Experimental results also show strong OOD generalization ability and large-scale model pool scalability. This prior-posterior routing approach offers a scalable and practical path toward efficient multi-LLM collaboration.

Limitations

The scorer requires retraining to support new LLMs. However, integrating a new LLM only involves training a lightweight scorer on a small curated query set, which takes about 25 minutes. Future work will explore retrain-free routing.

Acknowledgements

We thank Can Ouyang for his assistance with the evaluation experiments. This work was jointly supported by the National Natural Science Foundation of China (No. 62432009, No. 62422311), Shanghai Committee of Science and Technology, China (No. 24TS1413500).

References

- Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, and 1 others. 2023. GPT-4 technical report. *arXiv preprint arXiv:2303.08774*.
- Meta AI. 2025. [Introducing LLaMA 3.1: Our most capable models to date](#).
- Alon Albalak, Duy Phung, Nathan Lile, Rafael Rafailov, Kanishk Gandhi, Louis Castricato, Anikait Singh, Chase Blagden, Violet Xiang, Dakota Mahan, and 1 others. 2025. Big-math: A large-scale, high-quality math dataset for reinforcement learning in language models. *arXiv preprint arXiv:2502.17387*.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and 1 others. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Zheng Cai, Maosong Cao, Haojiong Chen, Kai Chen, Keyu Chen, Xin Chen, Xun Chen, Zehui Chen, Zhi Chen, Pei Chu, and 1 others. 2024. InternLM2 technical report. *arXiv preprint arXiv:2403.17297*.
- Daniel Cera, Mona Diabb, Eneko Agirrec, Inigo Lopez-Gazpio, Lucia Speciad, and Basque Country Donostia. Semeval-2017 task 1: Semantic textual similarity multilingual and cross-lingual focused evaluation.
- Yingshan Chang, Mridu Narang, Hisami Suzuki, Guihong Cao, Jianfeng Gao, and Yonatan Bisk. 2022. Webqa: Multihop and multimodal qa. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 16495–16504.
- Lingjiao Chen, Jared Quincy Davis, Boris Hanin, Peter Bailis, Ion Stoica, Matei A Zaharia, and James Y Zou. 2024a. Are more LLM calls all you need? towards the scaling properties of compound ai systems.
- Shuhao Chen, Weisen Jiang, Baijiong Lin, James Kwok, and Yu Zhang. 2024b. RouterDC: Query-based router by dual contrastive learning for assembling large language models.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, and 1 others. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- ContactDoctor. 2024. [Bio-Medical: A high-performance biomedical language model](#).
- OpenCompass Contributors. 2023. [OpenCompass: A universal evaluation platform for foundation models](#).
- Yiming Cui, Ting Liu, Wanxiang Che, Li Xiao, Zhipeng Chen, Wentao Ma, Shijin Wang, and Guoping Hu. 2019. A span-extraction dataset for chinese machine reading comprehension. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pages 5883–5889.
- Dujian Ding, Ankur Mallick, Chi Wang, Robert Sim, Subhabrata Mukherjee, Victor Rühle, Laks VS Lakshmanan, and Ahmed Hassan Awadallah. 2024. Hybrid LLM: Cost-efficient and quality-aware query routing.
- William B Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the third international workshop on paraphrasing (IWP2005)*.
- Markus Freitag, George Foster, David Grangier, Viresh Ratnakar, Qijun Tan, and Wolfgang Macherey. 2021. [Experts, errors, and context: A large-scale study of human evaluation for machine translation](#). *Preprint*, arXiv:2104.14478.
- Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, and 1 others. 2024. The LLaMA 3 herd of models. *arXiv preprint arXiv:2407.21783*.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shiron Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Pengcheng He, Jianfeng Gao, and Weizhu Chen. 2021. DeBERTaV3: Improving deberta using electra-style pre-training with gradient-disentangled embedding sharing.

- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2020. Measuring massive multitask language understanding.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset.
- Baotian Hu, Qingcai Chen, and Fangze Zhu. 2015. Lcsts: A large scale chinese short text summarization dataset. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1967–1972.
- Hai Hu, Kyle Richardson, Liang Xu, Lu Li, Sandra Kübler, and Lawrence S Moss. 2020. Ocnli: Original chinese natural language inference. In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 3512–3526.
- Qitian Jason Hu, Jacob Bieker, Xiuyu Li, Nan Jiang, Benjamin Keigwin, Gaurav Ranganath, Kurt Keutzer, and Shriyash Kaustubh Upadhyay. 2024. RouterBench: A benchmark for multi-llm routing system. *arXiv preprint arXiv:2403.12031*.
- Zhongzhan Huang, Guoming Ling, Vincent S Liang, Yupei Lin, Yandong Chen, Shanshan Zhong, Hefeng Wu, and Liang Lin. 2025. RouterEval: A comprehensive benchmark for routing llms to explore model-level scaling up in llms. *arXiv preprint arXiv:2503.10657*.
- Binyuan Hui, Jian Yang, Zeyu Cui, Jiayi Yang, Dayiheng Liu, Lei Zhang, Tianyu Liu, Jiajun Zhang, Bowen Yu, Keming Lu, and 1 others. 2024. Qwen2.5-Coder technical report. *arXiv preprint arXiv:2409.12186*.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. Livecodebench: Holistic and contamination free evaluation of large language models for code. In *The Thirteenth International Conference on Learning Representations*.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7b. *arXiv preprint arXiv:2310.06825*.
- Saurav Kadavath, Tom Conerly, Amanda Askell, Tom Henighan, Dawn Drain, Ethan Perez, Nicholas Schiefer, Zac Hatfield-Dodds, Nova DasSarma, Eli Tran-Johnson, and 1 others. 2022. Language models (mostly) know what they know. *arXiv preprint arXiv:2207.05221*.
- Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer*.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. Race: Large-scale reading comprehension dataset from examinations.
- Gina-Anne Levow. 2006. [The third international Chinese language processing bakeoff: Word segmentation and named entity recognition](#). In *Proceedings of the Fifth SIGHAN Workshop on Chinese Language Processing*, pages 108–117, Sydney, Australia. Association for Computational Linguistics.
- Dawei Li, Bohan Jiang, Liangjie Huang, Alimohammad Beigi, Chengshuai Zhao, Zhen Tan, Amrita Bhat-tacharjee, Yuxuan Jiang, Canyu Chen, Tianhao Wu, and 1 others. 2025. From generation to judgment: Opportunities and challenges of llm-as-a-judge. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 2757–2791.
- Dawei Li, Zhen Tan, Peijia Qian, Yifan Li, Kumar Satvik Chaudhary, Lijie Hu, and Jiayi Shen. 2024. SMOA: Improving multi-agent large language models with sparse mixture-of-agents. *arXiv preprint arXiv:2411.03284*.
- Tian Liang, Zhiwei He, Wenxiang Jiao, Xing Wang, Yan Wang, Rui Wang, Yujiu Yang, Shuming Shi, and Zhaopeng Tu. 2024. Encouraging divergent thinking in large language models through multi-agent debate.
- Stephanie Lin, Jacob Hilton, and Owain Evans. 2022. Truthfulqa: Measuring how models mimic human falsehoods. In *Proceedings of the 60th annual meeting of the association for computational linguistics (volume 1: long papers)*, pages 3214–3252.
- Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, and 1 others. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437*.
- Xin Liu, Qingcai Chen, Chong Deng, Huajun Zeng, Jing Chen, Dongfang Li, and Buzhou Tang. 2018. Lcqmc: A large-scale chinese question matching corpus. In *Proceedings of the 27th international conference on computational linguistics*, pages 1952–1962.
- Keming Lu, Hongyi Yuan, Runji Lin, Junyang Lin, Zheng Yuan, Chang Zhou, and Jingren Zhou. 2023. Routing to the expert: Efficient reward-guided ensemble of large language models. *arXiv preprint arXiv:2311.08692*.
- Chengqian Ma, Wei Tao, and Steven Y Guo. 2025. C3: A bilingual benchmark for spoken dialogue models exploring challenges in complex conversations. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing*, pages 22789–22807.

- James MacQueen. 1967. Some methods for classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability, Volume 1: Statistics*.
- Nasrin Mostafazadeh, Michael Roth, Annie Louis, Nathanael Chambers, and James Allen. 2017. Lsdsem 2017 shared task: The story cloze test. In *Proceedings of the 2nd workshop on linking models of lexical, sentential and discourse-level semantics*, pages 46–51.
- Hwee Tou Ng, Siew Mei Wu, Ted Briscoe, Christian Hadiwinoto, Raymond Hendy Susanto, and Christopher Bryant. 2014. The conll-2014 shared task on grammatical error correction. In *Proceedings of the eighteenth conference on computational natural language learning: shared task*, pages 1–14.
- Isaac Ong, Amjad Almahairi, Vincent Wu, Wei-Lin Chiang, Tianhao Wu, Joseph E Gonzalez, M Waleed Kadous, and Ion Stoica. 2024. RouteLLM: Learning to route LLMs with preference data. *arXiv preprint arXiv:2406.18665*.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, and 1 others. 2022. Training language models to follow instructions with human feedback.
- Xipeng Qiu, Jingjing Gong, and Xuanjing Huang. 2017. Overview of the nlpcc 2017 shared task: Chinese news headline categorization. In *National CCF Conference on Natural Language Processing and Chinese Computing*, pages 948–953. Springer.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2020. Winogrande: An adversarial winograd schema challenge at scale. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8732–8740.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. [Recursive deep models for semantic compositionality over a sentiment treebank](#). In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. Association for Computational Linguistics.
- Dimitris Stripelis, Zhaozhuo Xu, Zijian Hu, Alay Shah, Han Jin, Yuhang Yao, Jipeng Zhang, Tong Zhang, Salman Avestimehr, and Chaoyang He. 2024. TensorOpera Router: A multi-model router for efficient llm inference.
- Mirac Suzgun, Nathan Scales, Nathanael Schärli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V Le, Ed H Chi, Denny Zhou, and 1 others. 2023. Challenging big-bench tasks and whether chain-of-thought can solve them. In *ACL (Findings)*.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. Commonsenseqa: A question answering challenge targeting commonsense knowledge. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4149–4158.
- Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricut, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, and 1 others. 2023. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*.
- Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, and 1 others. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118*.
- Laurens Van der Maaten and Geoffrey Hinton. 2008. Visualizing data using t-sne. *Journal of Machine Learning Research*.
- Bingning Wang, Ting Yao, Qi Zhang, Jingfang Xu, and Xiaochuan Wang. 2020. Reco: A large scale chinese reading comprehension dataset on opinion. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9146–9153.
- Jize Wang, Xinyi Le, Xiaodi Peng, and Cailian Chen. 2023. Adaptive hinge balance loss for document-level relation extraction. In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 3872–3878.
- Jize Wang, Zerun Ma, Yining Li, Songyang Zhang, Cailian Chen, Kai Chen, and Xinyi Le. 2024a. Gta: a benchmark for general tool agents. *Advances in Neural Information Processing Systems*, 37:75749–75790.
- Junlin Wang, Jue Wang, Ben Athiwaratkun, Ce Zhang, and James Zou. 2024b. Mixture-of-agents enhances large language model capabilities. *arXiv preprint arXiv:2406.04692*.
- Yan Wang, Xiaojiang Liu, and Shuming Shi. 2017. Deep neural solver for math word problems. In *Proceedings of the 2017 conference on empirical methods in natural language processing*, pages 845–854.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, and 1 others. 2024c. Mmlu-pro: A more robust and challenging

- multi-task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37:95266–95290.
- Yuxia Wang, Zenan Zhai, Haonan Li, Xudong Han, Shom Lin, Zhenxuan Zhang, Angela Zhao, Preslav Nakov, and Timothy Baldwin. 2024d. A chinese dataset for evaluating the safeguards in large language models. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 3106–3119.
- Alex Warstadt, Amanpreet Singh, and Samuel Bowman. 2019. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Liang Xu, Hai Hu, Xuanwei Zhang, Lu Li, Chenjie Cao, Yudong Li, Yechen Xu, Kai Sun, Dian Yu, Cong Yu, and 1 others. 2020. Clue: A chinese language understanding evaluation benchmark. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4762–4772.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, and 1 others. 2025. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*.
- An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, and 1 others. 2024a. Qwen2.5 technical report. *arXiv preprint arXiv:2412.15115*.
- An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, and 1 others. 2024b. Qwen2.5-Math technical report: Toward mathematical expert model via self-improvement. *arXiv preprint arXiv:2409.12122*.
- Murong Yue, Jie Zhao, Min Zhang, Liang Du, and Ziyu Yao. 2024. Large language model cascades with mixture of thought representations for cost-efficient reasoning.
- Jinyu Zhang, Zhiyuan You, Jize Wang, and Xinyi Le. 2025. SAIL: Sample-centric in-context learning for document information extraction.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems*, 28.
- Xiaotian Zhang, Chunyang Li, Yi Zong, Zhengyu Ying, Liang He, and Xipeng Qiu. 2023. Evaluating the performance of large language models on gaokao benchmark. *arXiv preprint arXiv:2305.12474*.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019. Paws: Paraphrase adversaries from word scrambling. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1298–1308.
- Wayne Xin Zhao, Kun Zhou, Junyi Li, Tianyi Tang, Xiaolei Wang, Yupeng Hou, Yingqian Min, Beichen Zhang, Junjie Zhang, Zican Dong, and 1 others. 2023. A survey of large language models. *arXiv preprint arXiv:2303.18223*.
- Yuanyuan Zhao, Nan Jiang, Weiwei Sun, and Xiaojun Wan. 2018. Overview of the nlpcc 2018 shared task: Grammatical error correction. In *CCF International Conference on Natural Language Processing and Chinese Computing*, pages 439–445. Springer.
- Chujie Zheng, Minlie Huang, and Aixin Sun. 2019. Chid: A large-scale chinese idiom dataset for cloze test. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 778–787.
- Wanjun Zhong, Ruixiang Cui, Yiduo Guo, Yaobo Liang, Shuai Lu, Yanlin Wang, Amin Saied, Weizhu Chen, and Nan Duan. 2024. Agieval: A human-centric benchmark for evaluating foundation models. In *Findings of the Association for Computational Linguistics: NAACL 2024*, pages 2299–2314.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*.

A Aggregation Prompt, Self- and Cross-assessment Prompt

The prompts used in the inference stage of RouteMoA are shown in Figure 6, Figure 7, and Figure 8. These prompts include instructions for the LLM to aggregate responses from models in the previous layer, to score its own answer (self-assessment prompt), and to evaluate answers from other LLMs (cross-assessment prompt). Specifically, the prompt at layer-1 (Figure 6) includes both the aggregation prompt and the self-assessment prompt. For intermediate layer- l ($1 < l < L$) (Figure 7), the prompts consist of the aggregation prompt, the self-assessment prompt, and the cross-assessment prompt. Finally, in the last layer (Figure 8), only the aggregation prompt is included.

B Prompt for Training Dataset Generation

To generate training data that meets the requirements of layer-1, we prompt the LLMs to answer questions from the datasets. The prompts used for each dataset are presented in Figure 10. The

You are participating in a multi-agent reasoning task.

****Your objectives****

1. Produce the best possible answer to the user’s query.
2. Critically evaluate your own answer and give it a quality score ****between 0 and 1**** (0 = completely wrong, 1 = perfect).

****Output format**** - return ****ONLY**** a valid JSON object:

```
““json
{
  "answer": "<your answer>",
  "self_score": <float between 0 and 1>
}
””
```

Do ****not**** add any keys, comments or extra text.

Figure 6: The prompt used in layer-1.

model outputs are then compared with the ground truth answers. For layer- $l|l>1$, we prompt the models to generate aggregated answers based on reference answers from all models. These aggregated responses are subsequently evaluated against the ground truth. Additionally, a judge model is employed to assess the quality of each answer. The prompt used for generating the aggregated answers is shown in Figure 9. We use InternLM2-1.8B-Reward as the judge model.

C Clustering Details for Sample-Sample Loss

The *sample-sample contrastive loss* encourages semantically similar queries to have closer embeddings. To achieve this, we use t-SNE (Van der Maaten and Hinton, 2008) and k-means (MacQueen, 1967) algorithm to transfer input prompt embeddings to low-dimensional vectors and cluster them into Q groups $\{\mathcal{K}_1, \mathcal{K}_2, \dots, \mathcal{K}_Q\}$. We randomly select an in-group query $x^+ \in \mathcal{K}_q$, and an out-group set $\mathcal{X}^- \subset \{\cup_{q' \neq q} \mathcal{K}_{q'}\}$ of H queries from the training mini-batch.

You are participating in a multi-agent reasoning task. Here are several answers from other LLMs:

answer_block

****Your objectives****

1. Taking every answer in the previous round into account and produce an improved answer to the user’s query.
2. Critically evaluate your own answer and give it a quality score ****between 0 and 1**** (0 = completely wrong, 1 = perfect).
3. Critically evaluate ****each**** ANSWER _{i} above with a value in $[0, 1]$ representing its quality. (0 = completely wrong, 1 = perfect)

****Output format**** - return ****ONLY**** a valid JSON object:

```
““json
{
  "answer": "<your improved answer>",
  "self_score": <float>,
  "peer_scores":
  [<float_score_for_ANSWER_0>,
  <float_score_for_ANSWER_1>, ...]
}
””
```

Do not include any other text.

Figure 7: The prompt used in intermediate layer- $i|1 < i < l$.

Dataset	Prompt
MATH	Answer the following multiple choice question. The last line of your response should be of the following format: 'ANSWER: \$LETTER' (without quotes) where LETTER is one of ABCD. Think step by step before answering. {question} {options}
GSM8k	{question} Please reason step by step, and put your final answer within \boxed{ }.
ARC-c	{problem} Please reason step by step, and put your final answer within \boxed{ }.
MBPP	You are an expert Python programmer, and here is your task: {prompt} Your code should pass these tests: {test_list}
RACE-high	Answer the following multiple choice question. The last line of your response should be of the following format: 'ANSWER: \$LETTER' (without quotes) where LETTER is one of ABCD. Think step by step before answering. {question}{options}
MMLU-biomed	Answer the following multiple choice question. The last line of your response should be of the following format: 'ANSWER: \$LETTER' (without quotes) where LETTER is one of ABCD. Think step by step before answering. Article: {article} Q:{questions} {options}

Table 10: The prompt of each dataset.

You have been provided with a set of responses from various open-source models to the latest user query. Your task is to synthesise these responses into a single, high-quality answer. Critically evaluate the information given, correct any mistakes, and produce a coherent, well-structured response that meets the highest standards of accuracy.

Responses from models:

1.model_response_1

2.model_response_2

...

Figure 8: The prompt used in the last layer.

You have been provided with a set of responses from various open-source models to the latest user query. Your task is to synthesize these responses into a single, high-quality response. It is crucial to critically evaluate the information provided in these responses, recognizing that some of it may be biased or incorrect. Your response should not simply replicate the given answers but should offer a refined, accurate, and comprehensive reply to the instruction. Ensure your response is well-structured, coherent, and adheres to the highest standards of accuracy and reliability.

This is the original question answered by these models:

original question

Responses from models:

1.model_response_1

2.model_response_2

...

Figure 9: The prompt used to generate model response with reference answers.

D Top-1-Hit, Top-3-Hit and Top-3-Agree Calculation Details

We define three evaluation metrics to assess the effectiveness of the scorer in RouteMoA: Top-1 Hit Rate (Top-1-Hit), Top-3 Hit Rate (Top-3-Hit), and Top-3 Agreement Rate (Top-3-Agree).

Top-1 Hit Rate (Top-1-Hit)

The Top-1 Hit Rate (Top-1-Hit) measures the probability that the scorer’s top-1 prediction is one of the models that are able to provide a correct answer.

Let p be the index of the top 1 model according to the scorer, and $T = \{t_1, \dots, t_k\}$ be the set of indices of the models that are able to provide a correct answer.

We define the Top-1 Hit Rate as:

$$\text{Top-1-Hit} = \begin{cases} 1, & \text{if } p \in T, \\ 0, & \text{otherwise.} \end{cases}$$

The final Top-1-Hit is obtained by averaging this score over all test cases that can be answered correctly by at least one model.

Top-3 Hit Rate (Top-3-Hit)

The Top-3 Hit Rate (Top-3-Hit) measures the probability that the ground-truth best model is included in the scorer’s top-three predictions.

Let $P = \{p_1, p_2, p_3\}$ be the set of indices of the top 3 models according to the scorer, and $T = \{t_1, \dots, t_k\}$ be the set of indices of the models that are able to provide a correct answer.

We define the Top-3 Hit Rate as:

$$\text{Top-3-Hit} = \begin{cases} 1, & \text{if } |P \cap T| \geq 1, \\ 0, & \text{otherwise.} \end{cases}$$

The final Top-3-Hit is obtained by averaging this score over all test cases that can be answered correctly by at least one model.

Top-3 Agreement Rate (Top-3-Agree)

The Top-3 Agreement Rate (Top-3-Agree) is a metric to evaluate whether the top three models selected by the scorer align with those that have the best true performance. For each test case or dataset, the scorer selects the top 3 models based on predicted scores, and we compare them with the top 3 models according to ground truth performance.

Let $P = \{p_1, p_2, p_3\}$ be the set of indices of the top 3 models according to the scorer, and $T =$

$\{t_1, t_2, t_3\}$ be the set of indices of the top 3 models according to the ground truth.

We define the Top-3 Agreement Rate as:

$$\text{Top-3-Agree} = \begin{cases} 1, & \text{if } |P \cap T| = 3, \\ 0.6, & \text{if } |P \cap T| = 2, \\ 0.3, & \text{if } |P \cap T| = 1, \\ 0, & \text{otherwise.} \end{cases}$$

This scoring rule assigns a full score if all top 3 models are correctly identified, a partial score if two / one out of the top 3 are correct, and zero otherwise. The final Top-3-AR is obtained by averaging this score over all test cases that can be answered correctly by at least one model.

Scorer evaluation under different α and λ

We calculate the three evaluation metrics: Top-1-Hit, Top-3-Hit, and Top-3-Agree under different combinations of scorer training hyperparameters α and λ , the results are shown in Figure 10, 11, 12, respectively.




Figure 10: Top-1-Hit under different training hyperparameters (λ and α) for the scorer module.




Figure 11: Top-3-Hit under different training hyperparameters (λ and α) for the scorer module.




Figure 12: Top-3-Agree under different training hyper-parameters (λ and α) for the scorer module.

E Dataset Statistics for Scorer Training of Small-Scale Model Pool

The dataset statistics for scorer training are presented in Table 11. The train/dev/test splits generally follow the original partitioning of each dataset. For datasets that lack a dev split, we divide the original training portion into train and dev subsets.

Dataset	train	dev	test
MATH	7125	375	5000
ARC-c	1119	299	1165
MBPP	120	43	257
RACE-high	7000	300	3498
MMLU-biomed	92	25	817

Table 11: Dataset statistics for scorer training.

F Evaluation Details on Large-Scale Model Pool

To explore the ability of RouteMoA to handle large-scale agent pools, we conduct experiments with an agent pool containing 15 newest LLMs including Qwen (Yang et al., 2025) and Deepseek (Guo et al., 2025; Liu et al., 2024) series with varying sizes (from 4B to 235B parameters), capabilities, and think mode, as listed in Table 12.

Note that both MoA and SMoA are infeasible to handle the larger agent pool, since these methods need all LLMs to infer and then aggregate the responses. When the agent pool is large, it is a huge cost for all model to perform inference. Besides, the context will be too long and exceeds the limit. Although we can choose a subset of agents from the large agent pool for MoA and SMoA to alleviate these problems, these two methods lack clear criteria for selecting such model subsets. If we select high-performing models such as Deepseek-R1, the total cost will be high, and the context will be long. If we select smaller models such as Qwen2.5-7B,

although the cost will be lower, the performance can not be guaranteed.

In contrast, our RouteMoA method is designed to deal with such situation. It has clear criteria on how to select agent subsets according to the categories and complexity of user queries. It lowers cost while ensures a competitive performance. Specifically, to handle the larger agent pool, we collect a wide range of datasets as a query pool, shown in Table 15. We only select a subset from each dataset to enable that the whole query pool is not very large. The dataset statistics are also shown in Table 15. Maintaining a smaller query pool benefits the scalability of the scorer. If a new LLM is added into the agent pool, it infers on these queries (the process will be shorter if the query pool is relatively small) and results are used to train a new scorer. Training can be completed in less than 30 minutes, using less than 50GB GPU memory.

Size	Model	Thinking Mode
Small	Qwen3-4B	no-think
	Qwen3-8B	no-think
	Qwen2.5-7B-Instruct	no-think
Medium	Qwen3-14B	no-think
	Qwen3-32B	no-think
	Qwen3-30B-A3B	think/no-think
	QwQ-32B	think
	DeepSeek-R1-Distill-Qwen-14B	think
DeepSeek-R1-Distill-Qwen-32B	think	
Large	Qwen2.5-72B-Instruct	no-think
	Qwen3-235B-A22B	think/no-think
	DeepSeek-R1	think
	DeepSeek-R1-0528	think
	DeepSeek-V3	no-think
	DeepSeek-V3-0324	no-think

Table 12: Models that forms a larger agent pool.

We evaluate RouteMoA on 30 test sets, as shown in Table 14, each set contains 15 samples (not overlapping with the training set). Among these, lcqmc (Liu et al., 2018), mrpc (Dolan and Brockett, 2005), and cluewsc2020 (Xu et al., 2020) are out-of-distribution test sets, which do not appear in the training set.

G Detailed Offline Cost

The detailed offline cost of generating training data (Table 15) is shown in Table 13.

H Model & Data License and Intended Use Statement

Our experiments utilize a collection of publicly available models and benchmark datasets. To the

best of our knowledge, our use of these models and datasets is consistent with their intended research purposes as specified by their original creators. All models and datasets used in this work are cited. For any model or dataset we use, we adhere to its stipulated terms of use.

Model Name	Cost (\$)	Time (h)
Qwen3-4B	0.05	7.03
Qwen3-8B	0.18	5.82
Qwen2.5-7B-Instruct	0.10	1.65
Qwen3-14B	0.30	6.79
Qwen3-32B	0.56	8.08
Qwen3-30B-A3B	0.55	5.73
QwQ-32B	0.49	7.80
DeepSeek-R1-Distill-Qwen-14B	0.18	8.16
DeepSeek-R1-Distill-Qwen-32B	0.37	13.55
Qwen2.5-72B-Instruct	0.51	2.28
Qwen3-235B-A22B	0.81	9.08
DeepSeek-R1	2.39	9.61
DeepSeek-R1-0528	3.86	12.81
DeepSeek-Chat	1.64	1.89
DeepSeek-V3-0324	2.09	2.93

Table 13: Cost and time of training data generation for different models with 6 parallel requests.

Category	Dataset
Language Understanding	lcqmc, ocnli, sst2 cola, mrpc, msra, qqf sts_b, ag_news qnli, chid_baidu
Reading&QA	webqa, c3 cmrc, race, story_cloze
Logic Reasoning	cluewsc2020 winogrande_wsc truthful_qa
Math Reasoning	bigmath gsm8k, geometry GAOKAO-2023_Math_en prealgebra, precalculus
Language Generation	word_manipulation_v2 nlppc2017_task2 lcsts, conll2014 nlppc2018_task2

Table 14: Test dataset categories for the large agent pool.

Dataset	Description	Data Num
ag_news (Zhang et al., 2015)	News topic classification	80
algebra (Hendrycks et al., 2020)	Algebra math problems	80
BBH-100 (Suzgun et al., 2023)	BIG-Bench Hard subset	80
bigmath (Albalak et al., 2025)	Complex math problems	16
c3 (Ma et al., 2025)	Chinese multiple-choice QA	80
chid (Zheng et al., 2019)	Chinese idiom cloze test	80
chid_baidu (Zheng et al., 2019)	Baidu Chinese idiom dataset	80
chinese_safety_test_bias (Wang et al., 2024d)	Safety and bias evaluation	80
cmrc (Cui et al., 2019)	Chinese machine reading comprehension	80
cola (Warstadt et al., 2019)	Linguistic acceptability corpus	77
commonsense_qa (Talmor et al., 2019)	Commonsense question answering	80
conll2014 (Ng et al., 2014)	Grammatical error correction	80
counting_and_probability (Hendrycks et al., 2020)	Math combinatorics problems	80
GAOKAO-2023_Math_en (Zhang et al., 2023)	Chinese college entrance exam math	80
gsm8k_test_100 (Cobbe et al., 2021)	Grade school math (subset)	80
IFEval (Zhou et al., 2023)	Instruction following evaluation	320
lcsts (Hu et al., 2015)	Chinese short text summarization	79
LiveCodeBench100-2305-2409 (Jain et al.)	Live programming evaluation	80
math (Hendrycks et al., 2021)	General math problems	80
math23k_test_100 (Wang et al., 2017)	Math word problems (subset)	80
MATH-500 (Hendrycks et al., 2021)	Challenging math competition problems	400
mmlu_pro (Wang et al., 2024c)	Massive multi-task understanding	224
msra (Levow, 2006)	Named entity recognition	80
newstest2020 (Freitag et al., 2021)	Machine translation evaluation set	80
nlpc2017_task2 (Qiu et al., 2017)	News headline categorization	80
nlpc2018_task2 (Zhao et al., 2018)	Grammatical error correction	80
number_theory (Hendrycks et al., 2020)	Math number theory problems	80
ocnli (Hu et al., 2020)	Chinese natural language inference	80
prealgebra (Hendrycks et al., 2020)	Pre-algebra math problems	80
precalculus (Hendrycks et al., 2020)	Precalculus math problems	80
qnli (Rajpurkar et al., 2016)	Question-answering NLI	80
qqp (Zhang et al., 2019)	Quora question pairs similarity	80
race (Lai et al., 2017)	Reading comprehension	80
reco (Wang et al., 2020)	Chinese reading comprehension	80
squad (Rajpurkar et al., 2016)	Reading comprehension	80
sst2 (Socher et al., 2013)	Sentiment analysis (binary)	80
story_cloze_test (Mostafazadeh et al., 2017)	Story completion and reasoning	80
sts_b (Cera et al.)	Semantic textual similarity benchmark	80
truthful_qa (Lin et al., 2022)	Truthfulness evaluation in QA	80
webqa (Chang et al., 2022)	Web-based question answering	80
winogrande_wsc (Sakaguchi et al., 2020)	Coreference resolution	80
Total		4596

Table 15: The query pool used to train the scorer for the larger agent pool.