Who Routes the Router: Rethinking the Evaluation of LLM Routing Systems

Jiayi Yuan* **Yifan Lu*** **Rixin Liu** **Yu-Neng Chuang** **Hongyi Liu**
Shaochen Zhong **Yang Sui** **Guanchu Wang** **Jiarong Xing** **Xia Hu**
Rice University
{jy101,yl231}@rice.edu

Abstract

The growing ecosystem of Large Language Models (LLMs) with diverse capabilities and costs has motivated the need for LLM routing systems that dynamically select the most appropriate model for each query. Evaluating these routing systems is important yet inherently challenging due to the complex interplay of multiple factors: the selection of representative input queries, the composition of the model pool, and the definition of comprehensive evaluation metrics for optimal routing decisions. Through extensive analysis of existing benchmarks, we identify critical limitations that may lead to incomplete results and/or misleading conclusions about router performance: (1) limited task diversity, (2) imbalanced model pools, and (3) oversimplified evaluation methodologies. To address these limitations, we propose a novel evaluation framework that incorporates diverse task distributions (33,337 queries across 68 categories), a balanced model pool of 85 models with complementary model strengths, and multi-faceted metrics that reflect real-world deployment scenarios. We implement this framework as an open-source benchmark, enabling researchers to rigorously assess routing strategies under realistic conditions. The code is shared here: https://github.com/jy-yuan/rethinking-routing-evaluation

1 Introduction

Large Language Models (LLMs) are proliferating rapidly, resulting in a growing ecosystem of models with diverse parameter scales, capabilities, and computational costs [22, 5, 12]. While this diversity offers rich options for model selection, it also raises a key question: *which model best achieves the desired performance while minimizing cost?* Today, a common practice is to use a single model to handle all requests. However, this approach faces inherent trade-offs between performance and cost, as no LLM is universally optimal. For example, massive models like GPT-4 excel at complex reasoning but are significantly more costly than smaller alternatives such as Mixtral-8$\times$7B [9] or Llama-3.1-8B [6]. Meanwhile, domain-specialized models often outperform general-purpose ones within their specific areas of expertise [24, 21].

To address this, **LLM routing systems** (shown in Figure 1) have been proposed to dynamically select the most appropriate model for each query, matching query characteristics to model strengths while optimizing costs [5, 22, 12, 3, 20, 23, 4, 14]. To design effective LLM routers, **rigorous evaluation becomes especially crucial** [7, 8]. Routing decisions directly impact user experience and system cost-efficiency, yet the complexity of query-model matching makes it difficult to reason about routing strategies analytically. Moreover, the rapid evolution of LLMs demands continual reassessment of these strategies. Without a standardized evaluation framework [7, 8], it is challenging to compare

*Equal contribution

NeurIPS 2025 Workshop on Evaluating the Evolving LLM Lifecycle: Benchmarks, Emergent Abilities, and Scaling

different routing approaches, identify weaknesses, or confidently deploy systems where substantial costs and user satisfaction are at stake.

However, **evaluating LLM routers is inherently challenging**, as optimal routing decisions are context-dependent and shaped by specific priorities and constraints, such as cost, latency, and accuracy. Therefore, to build an effective evaluation framework that assesses router optimality across diverse scenarios, we must comprehensively consider the requirements from three core components of an LLM routing system: input queries/tasks, model candidates, and routing evaluation methodologies. First, it requires carefully chosen input queries that reflect realistic usage patterns, encompassing a diverse range of tasks with varying levels of difficulty. Second, it requires a thoughtfully composed model pool that is sufficiently large and diverse—capturing variations in capabilities and costs—to enable meaningful routing decisions. Third, it requires effective evaluation methodologies that capture the multifaceted nature of routing performance, such as cost–performance trade-offs, latency constraints, and task accuracy.

Unfortunately, our extensive analysis of existing evaluation benchmarks, including RouterBench [7], LLM-Blender [10], and EmbedLLM [28], reveals that current evaluation frameworks suffer from three fundamental shortcomings. (1) *Limited task diversity*: They often rely on artificially constructed tasks that fail to capture the complexity, diversity, and distribution of real-world queries. (2) *Imbalanced model pools*: They use imbalanced model sets where one model consistently outperforms others across all tasks, making routing decisions trivial. (3) *Oversimplified evaluation methodologies*: Current evaluations tend to focus solely on accuracy, overlooking critical production considerations such as cost-performance trade-offs, latency constraints, and reliability under varying workloads.

These limitations underscore the urgent need for a more comprehensive and effective evaluation framework. In response, we introduce RouterBench+, a new, well-designed, and open-source benchmark that addresses these critical gaps and establishes a new standard for LLM routing evaluation. Our solution includes: (1) a **specialist-score-based task sampling method** that creates a diverse set of 33,337 queries across 68 categories, (2) a **similarity-aware greedy model pruning and extension strategy** that yields a balanced pool of 85 models with complementary strengths, and (3) a **comprehensive evaluation methodology** combining classification-based and routing-rate paradigms with explicit Out-of-Distribution (OOD) testing. Using our evaluation pipeline, researchers can rigorously assess routing strategies under realistic conditions and make informed decisions about model selection trade-offs. We summarize our contributions as follows:

- A systematic analysis of LLM routing evaluation requirements and key limitations in existing benchmarks, showing how current approaches overlook real-world challenges and can lead to misleading conclusions about router performance.
- A comprehensive evaluation methodology that includes three key aspects: diverse task distributions reflecting realistic query patterns; balanced model pools that avoid single-model dominance; and multi-faceted evaluation metrics that capture complex constraints and trade-offs.
- An open-source, extensible evaluation platform that implements our methodology, enabling rigorous routing evaluation under realistic conditions and accelerating the progress toward designing optimal LLM routers.

2 Preliminary: LLM Routing and Its Ideal Evaluation

2.1 Formalization of LLM Routing and Optimization Problem

An LLM routing system dynamically assigns a user query to the most suitable model from a pool of available LLMs under certain constraints. The system consists of:

Input: A user query $p \in \mathcal{P}$, typically transformed into a vector representation $\mathbf{p} \in \mathbb{R}^d$ using embedding techniques.

Model Pool: A set of LLMs $\mathcal{M} = \{m_1, m_2, ..., m_n\}$ with varying performance $q_i(p)$ and cost $c_i(p)$.

Routing Function: A routing function $R : \mathcal{P} \to \mathcal{M}$ that selects a model $m^* = R(p)$ for each input. This may be deterministic or probabilistic, outputting a distribution $s(p)$ over $\mathcal{M}$.

The optimization problem can be formulated as $m^* = \arg\max_{m_i \in \mathcal{M}} O(m_i, p)$, where O is an objective function. Common objectives include maximizing performance and/or minimizing cost subject to quality threshold T:

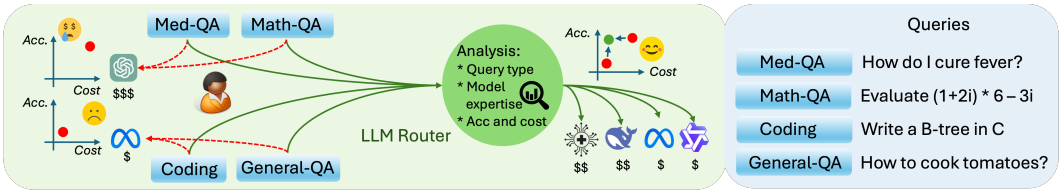

Figure 1: An illustration of LLM routing systems. An ideal LLM router should choose the model with highest expected performance under the specified constraints like costs.

$$m^* = \arg\max_{m_i \in \mathcal{M}} \hat{q}_i(p), \quad m^* = \arg\min_{m_i \in \mathcal{M}} c_i(p) \text{ subject to } \hat{q}_i(p) \geq T \tag{1}$$

Or one step further, under a budget B, maximizing average quality across queries:

$$\max_s \mathbb{E}_{p\sim\mathcal{P}}\left[\sum_{i=1}^{n} s_i(p)\hat{q}_i(p)\right] \quad \text{subject to} \quad \mathbb{E}_{p\sim\mathcal{P}}\left[\sum_{i=1}^{n} s_i(p)c_i(p)\right] \leq B \tag{2}$$

2.2 Ideal Router Evaluation

Figure 1 illustrates the architecture of an LLM routing system. An effective router should direct each query to the model with the highest expected performance while satisfying specified constraints like costs. To evaluate router performance and guide their design toward optimality, an effective and comprehensive evaluation framework is essential. However, designing such an evaluation framework presents inherent challenges. Unlike evaluating LLM performance—where each query can be assessed against a common ground truth—optimal routing strategies are highly context-specific. They depend on specific priorities and constraints, such as cost, latency, and accuracy. Even for the same query, the optimal routing decision may vary under different constraints.

To address these challenges, we return to the first principles by reconsidering what constitutes a "good router." We argue that ideal router evaluation should comprehensively assess routing strategies across diverse constraints and scenarios, providing clear differentiation between effective and suboptimal approaches. To this end, we distill three key requirements for a robust evaluation framework, which we examine in detail in the following sections.

- *Rich and realistic queries*. The task distribution should be diverse and representative of real-world usage, spanning various domains and difficulty levels. It should include both common and rare query patterns to evaluate router performance on familiar cases as well as unseen scenarios.
- *Diverse and balanced models*. The model pool should avoid single-model dominance, ensuring that each model has distinct strengths and weaknesses. It should include a mix of general-purpose and domain-specific models to ensure that routing decisions have a meaningful impact on performance.
- *Comprehensive evaluation metrics*. The evaluation framework should assess router effectiveness under varying constraints, capture trade-offs among performance, cost, and latency, and include OOD queries to evaluate robustness.

3 Related Work

LLM Model Selection and Routing. Intelligent LLM routers have emerged to route queries across diverse models to balance performance, cost, and latency [5, 22, 12, 25, 26]. Routing strategies can be categorized as predictive and non-predictive [22, 7]. Predictive approaches include classification based on prompt features [19], graph-based methods (GraphRouter [5]), dynamic routing (MixLLM [23]), and multi-armed bandit formulations (LLM Bandit [12]). Non-predictive methods include cascading, while hybrid approaches like Cascade Routing [3] combine routing flexibility with sequential processing. Frameworks like TensorOpera Router [20] further enhance multi-model inference efficiency. The proliferation of LLM routing methods has produced the requirement for effective router evaluation [2, 13, 27].

Benchmarks for Multi-LLM Systems. Several benchmarks have been developed to evaluate routing strategies. RouterBench [7] provides a framework with inference outcomes across models and tasks [3, 23]. EmbedLLM [28] introduces compact vector embeddings for efficient model selection.

MixInstruct [10] offers a mixture-of-instructions dataset with a two-stage ensembling approach. RouterEval [8] presents a large-scale benchmark with over 8,500 models and 200 million routing records. These benchmarks are crucial for developing robust routing systems that enable efficient, cost-effective LLM deployment [5, 19, 22, 12].

Despite the growing body of work on LLM routing techniques and benchmarks, we identify a critical gap: **the evaluation methodology itself has not been systematically examined**. Even the most comprehensive and recently released benchmarks, such as RouterEval [8], primarily aggregate large volumes of data and models without addressing fundamental flaws in evaluation design. This paper fills that gap by critically analyzing current evaluation practices and providing concrete recommendations for improvement. In the following sections, we systematically examine the assumptions underlying current practices in query distribution, model selection, and evaluation metrics, highlighting how they can lead to misleading conclusions about router performance.

4 Rethinking Current Evaluation Practices

This section examines current LLM routing evaluations, beginning with an overview of our methodology. We then analyze the three core components of a routing system: tasks, models, and evaluation metrics. For each component, we (a) explain the underlying assumption or practice, (b) describe our experimental setup, including the dataset or benchmark used, and (c) present and discuss the results, highlighting what they reveal about the assumption.

4.1 Experimental Setup

Benchmark & Datasets. We evaluate routing performance using three prominent benchmarks: EMBEDLLM [28], ROUTERBENCH [7], and MIXINSTRUCT [10].

Routing Methods. The state-of-the-art LLM routing approaches can be broadly categorized into two groups: *clustering-based methods*, such as K-Means [11], K-NN [7]; and *learning-based methods*, including MLP [7] and Collaborative Filtering (Matrix Factorization) [28]. Additionally, we include two *reference baselines* to provide contextual performance benchmarks: a *Heuristic Router*, which routes all queries to the model with the highest average training accuracy within the allowed cost, and an *Oracle Router*, which serves as an upper bound by assuming access to ground-truth model performance at test time.

Evaluation Metrics and Deferral Curve. We evaluate routing performance on each benchmark using its corresponding evaluation metrics. For ROUTERBENCH [7] and EMBEDLLM [28], each LLM either answers a query correctly or not, producing a binary correctness label. For MIXINSTRUCT [10], we follow prior work [11, 10] to adopt the exponentiated BARTScore for evaluation. Routing quality is visualized using a *deferral curve*, where the X-axis indicates the model cost budget, such as cost in dollars or parameter size; and the Y-axis represents routing quality, such as accuracy or exp (BARTScore). The deferral curve captures the trade-off between routing quality and resource usage, allowing comparison of different routing strategies under cost constraints.

For more details on the benchmark datasets, routing methods, and evaluation metrics, please refer to Appendix C–F.

4.2 Tasks: More Diversity and Less Redundancy

Problem 1: Lack of Specialized Tasks. Generally, tasks can be categorized as *common-sense* or *domain-specific*. For instance, `piqa` [1], a physical common-sense task, is handled well by general models (generalists), with an average accuracy of 78.03%. In contrast, the medical domain's `medmcqa` [15] has a lower average of 41.73% with general models, while a domain-specific model (specialist) can achieve 69.8%.

To evaluate routing performance across different scenarios, the task set should be sufficiently diverse, including both common-sense and domain-specific tasks. However, current benchmarks such as ROUTERBENCH and EMBEDLLM are biased toward common-sense tasks. This bias could lead to a failure to evaluate routers' ability to handle domain-specific tasks, a critical class of queries that benefit a lot from model routing.

To quantify this imbalance, we propose a **specialist score** for each task—defined as the average (across cost budgets) of the difference in accuracy between the best-performing domain-specific model and the general model:

$$\text{specialist_score}_{\text{task}} = \mathbb{E}_{b \in \mathcal{B}} \left[\max_{m \in \mathcal{M}^{(b)}_{\text{non-gen}}} \text{ACC}^{(b)}_{m,t} - \text{ACC}^{(b)}_{\text{gen},t} \right],$$

where $\mathcal{B}$ represents cost budgets, $\mathcal{M}^{(b)}_{\text{non-gen}}$ excludes the general model, and ACC measures accuracy. This score captures how much specialists outperform generalists on specific tasks. We computed this score on both benchmarks, with results presented in Figure 2. Ideally, we expect a long-tail distribution—most tasks having moderate or negative scores, and a few showing high specialist scores—indicating that while general models suffice for many tasks, some benefit from specialization. However, we observed only a limited number of specialist tasks presented in both benchmarks.

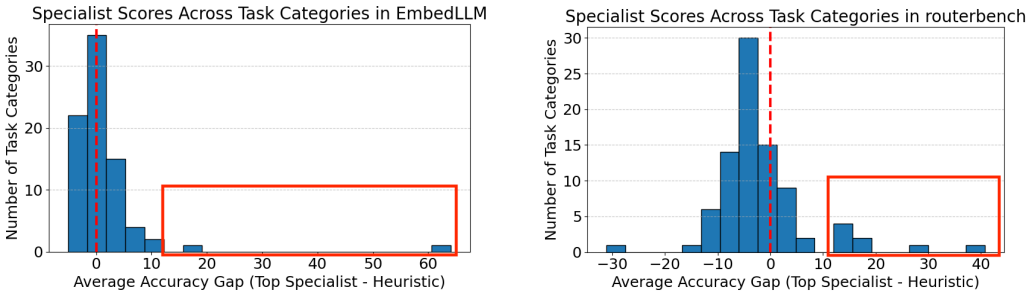

(a) EMBEDLLM (b) ROUTERBENCH

Figure 2: The specialist scores reveal that current datasets lack diverse, specialized tasks.

Problem 2: Task Redundancy. We have also found that current benchmarks suffer from significant task redundancy. As detailed in Appendix G, many task categories exhibit high similarity in their average query embeddings. Even after removing duplicate categories from the training set, the router maintains strong performance. To investigate potential redundancy in the benchmarks, we moved from category-level analysis to query-level inspection, identifying semantically equivalent queries that appeared multiple times in the training set. We computed cosine similarity between normalized query embeddings. Queries q_i and q_j were considered duplicates if:

$$\text{sim}(q_i, q_j) = \frac{\langle \mathbf{e}_i, \mathbf{e}_j \rangle}{\|\mathbf{e}_i\| \cdot \|\mathbf{e}_j\|} \geq \delta, \quad \text{where} \quad \delta = 0.999$$

Our analysis revealed 1,346 duplicate groups (average size 2.7). Critically, 99.9% of these groups contained label disagreements across models—far exceeding the overall label mismatch rate of 37.7%. This suggests that semantically identical queries often received inconsistent labels. By removing such duplicates and retraining routers on the cleaned dataset, we observed improved performance for learning-based methods, as shown in Table 1. This confirms that duplicate queries with conflicting labels can mislead routers to learn from noise rather than meaningful patterns.

Table 1: Comparison of routing performance before and after removing duplicate queries.

Method	Avg. Acc. (%) ↑		Peak Acc. (%) ↑	
	Original	Reduced	Original	Reduced
K-NN	**54.35**	54.04	**67.37**	66.50
Universal (KMeans)	**54.03**	54.00	**66.77**	66.70
MLP	53.78	**53.84**	64.17	**65.13**
Matrix Factorization	50.48	**50.90**	60.07	**60.87**

Insights. Current benchmarks overestimate the value of large but non-diverse training sets; in reality, much of the routing signal is concentrated in a smaller, more representative subset of tasks. To build more effective routing benchmarks, we should improve task diversity—especially by including more domain-specific tasks—and reduce redundancy, particularly tasks with inconsistent labels.

4.3 Models: More Specialists and Less Dominance

Problem 1: Model Dominance. A meaningful model pool should ensure that each LLM contributes unique strengths—some serving as generalists, others as specialists. This diversity is essential to the routing task: matching each input to the most capable model. If a single model dominates across all tasks, routing becomes redundant.

To quantify this, we compute each model's *average rank* across task categories. Lower ranks indicate better performance; a rank close to one means the model is often the best-performing for many tasks. Table 2 shows the top models for two major benchmarks. In EMBEDLLM, which features a larger and more diverse model pool (112 models), we observe multiple competitive models. The top model (ID 50) has an average rank of 6.4, indicating that routing still meaningfully impacts outcomes. In contrast, ROUTERBENCH includes only 11 models, with a single generalist (ID 5) dominating at an average rank of 1.36. This suggests that most queries are best handled by one model, thus limiting the value of learned routing.

This dominance stems from the lack of specialist models in the pool. A single model (the heuristic top model) handles most queries nearly as effectively as learned routers, achieving performance within a few percentage points of the oracle. As a result, routing decisions become largely unnecessary. In a well-constructed benchmark, some tasks (e.g., symbolic math, medicine, or historical reasoning) should require domain-specific expertise that only specialist models can provide. However, in current benchmarks, strong generalists often fill this role, even for tasks they were not explicitly designed for. This reduces the routing objective to simply identifying the best generalist, undermining the value of fine-grained model selection.

Table 2: Top-5 models by their average rank across tasks. Lower values indicate greater dominance.

EMBEDLLM			ROUTERBENCH		
Rank	**Model ID**	**Avg. Rank (↓)**	**Rank**	**Model ID**	**Avg. Rank (↓)**
1	50	6.43	1	5	1.36
2	83	9.88	2	10	3.20
3	42	10.03	3	4	3.78
4	49	10.95	4	9	3.88
5	5	11.24	5	3	5.39

Effective Expert Model Extension. To address this limitation, we propose augmenting the model pool with *pseudo-specialist models*—artificial models designed to perform well on specific tasks and average elsewhere. These pseudo-models are not meant for deployment but serve as controlled interventions to examine how task-specialized models influence routing behavior. They allow us to test whether the router moves beyond favoring top generalists and begins making more diverse, task-aware selections.

For each pseudo-specialist model, we need to identify appropriate task types where they can demonstrate their expertise. We selected three specific task types to create our pseudo-specialist models, each designed to excel at one particular task. These tasks satisfy the following criteria:

- The mean accuracy across existing models is low, suggesting that they are challenging.
- The gap between the best general model and the mean is modest, so no existing model dominates.
- The task has a non-negligible representation that impacts the overall performance in the benchmark.

Table 3: Selected tasks for pseudo specialist models.

Task	Prompt %	Mean Acc. (%)	Best Model Acc. (%)	Pseudo Model Acc. (%)
`Social Reasoning`	5.42	33.76	36.22	**65.00**
`Logical Reasoning`	1.82	28.28	45.93	**70.00**
`Graduatel-Level Reasoning`	3.23	22.44	33.51	**60.00**

Each pseudo model is injected into the benchmark with high performance on a specific task (Table 3) and average performance elsewhere, simulating models trained for niche tasks. We use EMBEDLLM as the task selection pool.

Table 4: Changes in router agreement with the top-1 generalist model after adding pseudo specialist models. Negative values indicate decreased reliance on the dominant model.

Task	K-NN	KMeans	MF	MLP
`Overall`	-0.84	-2.40	-0.64	**-8.40**
`logiqa`	**-20.55**	**-31.03**	-2.81	**-17.48**
`social_iqa`	-2.69	0.00	+0.29	**-7.92**
`gpqa`	-1.59	**-13.15**	+1.90	**-9.75**

We further define the *agreement score* as the average percentage of queries for which a router selects the same model as the heuristic router. This metric reflects how closely a learned router mimics static generalist selection. A lower score indicates more diverse, task-specific choices, suggesting less reliance on the generalist strategy. As shown in Table 4, overall agreement with the heuristic router drops slightly across all methods. However, on the tasks targeted by the pseudo models, the reduction is significantly more pronounced.

Problem 2: Model Redundancy. We also observe redundancy in the model pool, which adds little value to training or evaluating router performance. We quantify model-level similarity using a Jaccard-style score based on shared correct predictions as detailed in Appendix H. We apply this strategy to the `EmbedLLM` benchmark, reducing the model pool from 112 to 82 (a 27% reduction). The experimental result (in Appendix) shows that routing performance across methods remains comparable to the full model pool. This shows that removing redundant models does not degrade routing quality and that meaningful routing decisions can still be made with a leaner model pool.

Insights. Effective routing evaluation depends on a model pool with *meaningful* diversity, both in capability and specialization. Rather than including many models with overlapping strengths, the pool should consist of models with distinct specialties. A simple yet effective way to enhance current model pools is to introduce pseudo-specialist models that simulate task-specific expertise, encouraging routers to move beyond generic selection and make more nuanced, task-aware decisions.

4.4 Evaluation Paradigms: Comprehensive Measurements

Problems. Current evaluation paradigms suffer from at least two limitations: (1) *Limited cost awareness*: Existing evaluations often overlook the importance of selecting smaller, more efficient models when appropriate, leading to inflated costs and suboptimal routing decisions. (2) *Lack of OOD evaluation*: Current frameworks rarely test router performance on OOD inputs, an essential aspect for ensuring robustness in real-world deployments.

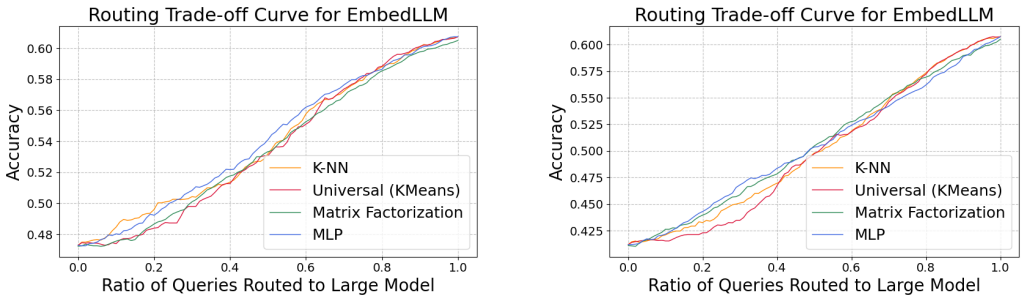

(a) Llama-2 7B and Llama-2 70B (b) Mistral 7B and Llama-2 70B

Figure 3: Binary routing evaluation paradigm showing performance trade-offs.

Multi-Faceted Evaluation: We argue that model routing evaluation should be multi-faceted, which should employ metrics that capture both performance quality and resource efficiency:

- *Cost-aware evaluation*: It should implement evaluation scenarios that explicitly consider cost constraints and encourage efficient model selection.
- *OOD testing framework*: It should develop a systematic approach to evaluate router performance on OOD scenarios, ensuring robustness in real-world deployment.

Routing Tradeoff Evaluation: To complement traditional cost-accuracy deferral curves, we introduce a binary routing evaluation paradigm to assess how effectively a router balances between a strong generalist (large model) and a lightweight alternative (small model). For each router, we vary the fraction of queries routed to the small model based on the router's confidence and measure the resulting accuracy. This produces a continuous trade-off curve be-

Table 5: OOD Performance change on math-related categories in EMBEDLLM when these categories are excluded from training.

Category	K-NN Δ	KMeans Δ	MF Δ	MLP Δ
`mathqa`	-9.29	-16.88	-6.33	-14.34
`asdiv`	-58.59	-69.19	-40.40	-57.07
`gsm8k`	-14.28	-14.29	-29.47	-35.72

tween routing accuracy and reliance on expensive models, as shown in Figure 3.

OOD Testing Framework: In real-world deployment, routers are likely to encounter out-of-distribution (OOD) queries—inputs from domains or tasks not represented during training. While benchmark designers should strive to include diverse tasks to improve generalization, OOD inputs are inevitable given the open-ended nature of user interactions with LLMs. Thus, evaluating router robustness in such scenarios is crucial. We design an OOD evaluation setup by holding out a subset of queries (e.g., Math tasks in EMBEDLLM) from training and evaluating performance on them separately. Table 5 illustrates an example split.

Insights. Current benchmarks inadequately assess router performance in realistic scenarios. The OOD performance degradation (Table 5) reveals the brittleness of existing approaches with novel queries, highlighting the need for better generalization testing. Additionally, the binary routing paradigm (Figure 3) shows that routing algorithms have distinct efficiency-performance trade-offs, requiring evaluation beyond single-point metrics.

5 Remastered Evaluation Pipeline

Building on our analysis of current evaluation limitations and the ideal characteristics of router evaluation, we present a comprehensive framework for assessing LLM routing systems. While developing a "perfect" evaluation pipeline presents challenges comparable to designing an "ideal" LLM router itself, our framework addresses the key shortcomings identified in our experimental analysis. We aim to provide the research community with a reliable mechanism to assess and compare different routing approaches under consistent conditions.

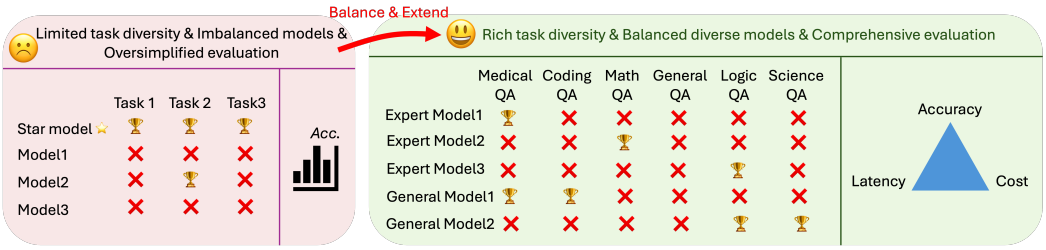

Figure 4: The improvement of our proposed benchmark.

5.1 Benchmark Design

Our evaluation framework is built upon core principles that directly address the limitations identified in our experimental analysis, as shown in Figure 4.

Diverse task distributions: Drawing from our findings on data representation issues, we incorporate **tasks with varying levels of difficulty, domain coverage, and redundancy ❶** to reflect real-world scenarios where task distributions are rarely static or uniform. This addresses the limitations identified in our analysis of current benchmarks that assume representative and static task distributions. To achieve this, we subsample tasks and queries from EMBEDLLM using the proposed *specialist score*, highlighting tasks where non-generalist models provide additional values. This results in a task pool that emphasizes both broad coverage and the need for routing.

Balanced model pool: To address the single-model dominance problem uncovered in our experiments, we carefully curate model pools that ensure **no single model outperforms all others across most tasks ❷**. This design choice enables meaningful evaluation of routing decisions rather than defaulting to a single dominant model. We apply the greedy model pruning strategy discussed in Sec 4.3 to eliminate redundant models, using a similarity-aware scoring function balancing accuracy and uniqueness. This reduces 30 models from the model pool. Additionally, we introduce three *pseudo specialist models* targeting challenging tasks (Table 3) to mitigate the single model dominance problem and diversify the routing model pool.

Multi-faceted evaluation metrics: Responding to our findings about oversimplified evaluation approaches, we combine **both classification-based and routing-rate paradigms ❸** to provide a

comprehensive assessment of router performance under different constraints. This moves beyond binary evaluation approaches that fail to capture the complexity of real-world routing scenarios, incorporating critical factors like cost-performance trade-offs, latency constraints, and reliability under varying workloads. We integrate cost-constrained evaluation with routing-rate analysis to provide researchers with multiple perspectives on router performance. Furthermore, our metrics specifically **account for OOD performance ❹**, ensuring that routers are evaluated on their ability to generalize to novel scenarios. We include dedicated OOD testing phases that assess router performance on novel task types and difficulty levels, providing insights into real-world deployment readiness.

The final dataset has 85 models, 68 categories, and 33,337 queries, in total of 3 million datapoints.

5.2 Experiment Results

We evaluate routing methods on our remastered benchmark, with results shown in Figure 5. K-NN achieves the highest performance with an area under the deferral curve of 0.567 and peak accuracy of 69.83%. Figure 5(a) illustrates the performance-deferral trade-offs, demonstrating that our dataset has successfully mitigated the single model dominance problem. Figure 6 shows the binary routing paradigm results, which reveal distinct efficiency-performance patterns across different model combinations. For the full results and additional evidence showcasing the improvements of our benchmark, please refer to the Appendix.

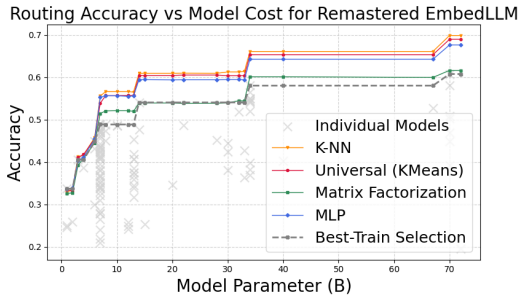

(a) Deferral curve on Remastered Benchmark

Method	Area ↑	Peak Acc. (%) ↑
K-NN	0.567	69.83
Universal (KMeans)	0.560	68.93
MLP	0.554	67.60
Matrix Factorization	0.515	61.60
Heuristic	0.507	60.73

(b) Area and peak accuracy of routing methods

Figure 5: Routing performance on our Remastered Benchmark.

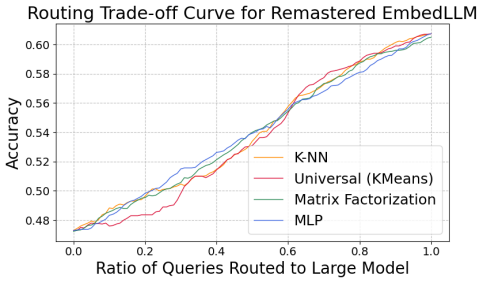

(a) Llama-2 7B and Llama-2 70B

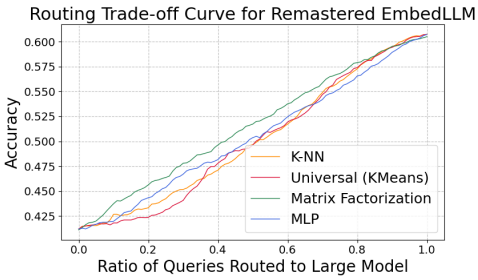

(b) Mistral 7B and Llama-2 70B

Figure 6: Binary routing evaluation on Remastered Benchmark shows performance trade-offs.

6 Conclusion

In this work, we have conducted a comprehensive analysis of LLM routing evaluation practices and identified critical limitations in current benchmarks. Through extensive experimentation, we demonstrated that existing evaluation frameworks often fail to capture the true complexity of routing decisions, leading to potentially misleading conclusions about router performance. Our findings reveal fundamental issues in task distribution representation, model pool composition, and evaluation metrics that significantly impact the validity and effectiveness of routing system evaluation. To address these limitations, we proposed a novel evaluation framework that incorporates diverse task distributions,

balanced model pools, and multi-faceted metrics, providing researchers and practitioners with a more robust tool for assessing routing strategies. This work not only advances our understanding of what constitutes effective LLM routing but also establishes a foundation for more rigorous evaluation practices in this rapidly evolving field.

Acknowledgements

This research was partially supported by NSF Awards ITE-2429680, IIS-2310260as well as US Department of Transportation (USDOT) Tier-1 University Transportation Center (UTC) Transportation Cybersecurity Center for Advanced Research and Education (CYBER-CARE) grant 69A3552348332. The views and conclusions in this paper are those of the authors and do not represent the views of any funding or supporting agencies.

References

[1] Yonatan Bisk, Rowan Zellers, Ronan Le Bras, Jianfeng Gao, and Yejin Choi. Piqa: Reasoning about physical commonsense in natural language, 2019.

[2] Shuhao Chen, Weisen Jiang, Baijiong Lin, James T. Kwok, and Yu Zhang. Routerdc: Query-based router by dual contrastive learning for assembling large language models, 2024.

[3] Jasper Dekoninck, Maximilian Baader, and Martin Vechev. A unified approach to routing and cascading for llms, 2025.

[4] Not Diamond. Not diamond homepage. `https://www.notdiamond.ai/`. Accessed: 2025-05-14.

[5] Tao Feng, Yanzhen Shen, and Jiaxuan You. Graphrouter: A graph-based router for llm selections, 2025.

[6] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

[7] Qitian Jason Hu, Jacob Bieker, Xiuyu Li, Nan Jiang, Benjamin Keigwin, Gaurav Ranganath, Kurt Keutzer, and Shriyash Kaustubh Upadhyay. Routerbench: A benchmark for multi-llm routing system, 2024.

[8] Zhongzhan Huang, Guoming Ling, Vincent S Liang, Yupei Lin, Yandong Chen, Shanshan Zhong, Hefeng Wu, and Liang Lin. Routereval: A comprehensive benchmark for routing llms to explore model-level scaling up in llms. *arXiv preprint arXiv:2503.10657*, 2025.

[9] Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. Mixtral of experts. *arXiv preprint arXiv:2401.04088*, 2024.

[10] Dongfu Jiang, Xiang Ren, and Bill Yuchen Lin. Llm-blender: Ensembling large language models with pairwise ranking and generative fusion. *arXiv preprint arXiv:2306.02561*, 2023.

[11] Wittawat Jitkrittum, Harikrishna Narasimhan, Ankit Singh Rawat, Jeevesh Juneja, Zifeng Wang, Chen-Yu Lee, Pradeep Shenoy, Rina Panigrahy, Aditya Krishna Menon, and Sanjiv Kumar. Universal model routing for efficient llm inference, 2025.

[12] Yang Li. Llm bandit: Cost-efficient llm generation via preference-conditioned dynamic routing, 2025.

[13] Keming Lu, Hongyi Yuan, Runji Lin, Junyang Lin, Zheng Yuan, Chang Zhou, and Jingren Zhou. Routing to the expert: Efficient reward-guided ensemble of large language models. In Kevin Duh, Helena Gomez, and Steven Bethard, editors, *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, pages 1964–1974, Mexico City, Mexico, June 2024. Association for Computational Linguistics.

[14] NVIDIA. Nvidia ai blueprint: Llm router. https://github.com/NVIDIA-AI-Blueprints/llm-router. Accessed: 2025-05-14.

[15] Ankit Pal, Logesh Kumar Umapathi, and Malaikannan Sankarasubbu. Medmcqa : A large-scale multi-subject multi-choice dataset for medical domain question answering, 2022.

[16] requesty.ai. Requesty ai homepage. https://www.requesty.ai/. Accessed: 2025-05-14.

[17] Sentence-Transformers. all-minilm-l12-v2. https://huggingface.co/sentence-transformers/all-MiniLM-L12-v2, 2021. Accessed: 2025-05-13.

[18] Sentence-Transformers. all-mpnet-base-v2. https://huggingface.co/sentence-transformers/all-mpnet-base-v2, 2021. Accessed: 2025-05-13.

[19] KV Srivatsa, Kaushal Kumar Maurya, and Ekaterina Kochmar. Harnessing the power of multiple minds: Lessons learned from llm routing. *arXiv preprint arXiv:2405.00467*, 2024.

[20] Dimitris Stripelis, Zijian Hu, Jipeng Zhang, Zhaozhuo Xu, Alay Dilipbhai Shah, Han Jin, Yuhang Yao, Salman Avestimehr, and Chaoyang He. Tensoropera router: A multi-model router for efficient llm inference. *arXiv preprint arXiv:2408.12320*, 2024.

[21] Tao Tu, Shekoofeh Azizi, Danny Driess, Mike Schaekermann, Mohamed Amin, Pi-Chuan Chang, Andrew Carroll, Charles Lau, Ryutaro Tanno, Ira Ktena, et al. Towards generalist biomedical ai. *Nejm Ai*, 1(3):AIoa2300138, 2024.

[22] Clovis Varangot-Reille, Christophe Bouvard, Antoine Gourru, Mathieu Ciancone, Marion Schaeffer, and François Jacquenet. Doing more with less – implementing routing strategies in large language model-based systems: An extended survey, 2025.

[23] Xinyuan Wang, Yanchi Liu, Wei Cheng, Xujiang Zhao, Zhengzhang Chen, Wenchao Yu, Yanjie Fu, and Haifeng Chen. Mixllm: Dynamic routing in mixed large language models. *arXiv preprint arXiv:2502.18482*, 2025.

[24] An Yang, Beichen Zhang, Binyuan Hui, Bofei Gao, Bowen Yu, Chengpeng Li, Dayiheng Liu, Jianhong Tu, Jingren Zhou, Junyang Lin, et al. Qwen2. 5-math technical report: Toward mathematical expert model via self-improvement. *arXiv preprint arXiv:2409.12122*, 2024.

[25] Yanwei Yue, Guibin Zhang, Boyang Liu, Guancheng Wan, Kun Wang, Dawei Cheng, and Yiyan Qi. Masrouter: Learning to route llms for multi-agent systems, 2025.

[26] Tuo Zhang, Asal Mehradfar, Dimitrios Dimitriadis, and Salman Avestimehr. Leveraging uncertainty estimation for efficient llm routing, 2025.

[27] Yi-Kai Zhang, De-Chuan Zhan, and Han-Jia Ye. Capability instruction tuning: A new paradigm for dynamic llm routing, 2025.

[28] Richard Zhuang, Tianhao Wu, Zhaojin Wen, Andrew Li, Jiantao Jiao, and Kannan Ramchandran. Embedllm: Learning compact representations of large language models. *arXiv preprint arXiv:2410.02223*, 2024.

A Limitations

Although this paper takes important steps toward evaluating LLM routing, several open challenges remain. First, comprehensive evaluation requires a large dataset that covers diverse models, query types, deployment scenarios, and trade-off dimensions. While we propose several strategies to enhance existing benchmarks, we do not augment them with additional models and queries due to resource constraints. Second, beyond academic efforts, several emerging startups (e.g., NotDiamond [4] and Requesty [16]) are actively working on LLM routing. However, because they do not disclose details about their routing algorithms, datasets, or training methods, we did not include them in our comparison. Moreover, the rapid evolution of LLMs and the continuous emergence of new applications make it difficult to maintain up-to-date evaluation benchmarks. Developing methods to automatically and incrementally update these benchmarks is an important area for future research.

B Broader Impacts

The evaluation of model routing has broad and significant impacts on both research and real-world deployment of LLM systems. First, accurate and comprehensive evaluation frameworks enable fair comparison among routing strategies, fostering innovation and guiding future routing system designs. Second, in practical deployments, robust evaluation helps service providers make informed decisions about which models to use for specific tasks, balancing trade-offs between cost, latency, and output quality. This, in turn, can democratize access to high-quality AI by making advanced capabilities more affordable and widely available. Additionally, efficient routing improves energy usage and reduces computational overhead, helping to mitigate the environmental impact of large-scale AI deployments. Finally, as LLMs become increasingly integrated into safety-critical and user-facing applications, reliable routing evaluation plays a key role in ensuring the consistency, robustness, and trustworthiness of AI systems.

C Details about Text Encoder

Text encoder is a critical component of LLM routers, which transforms input prompts into embeddings used for routing decisions. To ensure faithful and fair comparison, we follow prior work [28, 7] and adopt consistent encoder choices per benchmark: we use `all-MiniLM-L12-v2` [17] for ROUTERBENCH and MIXINSTRUCT, and `all-mpnet-base-v2` [18] for EMBEDLLM.

D Details about Benchmarks

Table 6 summarizes the statistics of used benchmarks. EmbedLLM provides the largest number of models, while RouterBench provides a realistic cost setting. **MixInstruct** focuses on open-domain user prompts, using soft metrics like BARTScore to evaluate output quality.

Table 6: Comparison of benchmark datasets for LLM routing evaluation.

Benchmark	# Models	# Queries	# Categories	Metric	Cost Info
EmbedLLM [28]	112	35,673	80	Binary (0/1)	param size (B)
RouterBench [7]	11	36,497	86	Binary (0/1)	USD per 1k queries
MixInstruct [10]	12	110,000	5 (Open-domain)	exp(BARTScore)	param size (B)

E Details about Routing Methods

The state-of-the-art LLM routing approaches fall into two primary categories: *clustering-based* and *learning-based*. We also include two *reference baselines* to contextualize performance.

- **K-Means** [11]: This method clusters training queries into K clusters based on their embeddings. Given a test query q, the router finds the closest cluster C_k and selects the model m^* that performs

best on average within that cluster:

$$m^* = \arg\max_{m_i \in \mathcal{M}} \left[\frac{1}{|C_k|} \sum_{l \in C_k} \text{metric}(m_i, l) \right]$$

where C_k is the set of training prompts in the cluster of q, and metric denotes either a binary correctness label or $\exp(\text{BARTScore})$.

- **K-NN** [7]: Instead of relying on cluster centroids, this method finds the K nearest neighbors of the query q in the training set (based on embedding distance) and routes to the model with the highest average score on those neighbors.
- **MLP** [7]: For each LLM m_i, a separate MLP is trained to predict the performance score for query q:

$$P_i(x) = f(W_n \cdot \sigma(\ldots \sigma(W_1 \cdot x + b_1) \ldots) + b_n)$$

where x is the query embedding, σ denotes the activation function, and f is the final output layer. The model m^* with the highest predicted score $P_i(q)$ is selected.

- **Collaborative Filtering (Matrix Factorization)** [28]: This method treats the model routing task as a matrix completion problem. Given a binary matrix $Y \in \{0, 1\}^{M \times Q}$ representing whether model m_i correctly answered query q_j, it learns latent embeddings for models and queries by factorizing Y as:

$$Y_{ij} \approx u_i^\top v_j$$

where $u_i \in \mathbb{R}^d$ is the latent embedding for model m_i and $v_j \in \mathbb{R}^d$ for query q_j. At inference time, the router computes v_q (e.g., via a linear projection from query embedding) and selects the model with the highest predicted score:

$$m^* = \arg\max_{m_i \in \mathcal{M}} u_i^\top v_q$$

- **Heuristic Router**: This baseline selects the best-performing model from the training set for each cost budget. At each test time cost step, it routes all queries to the model that achieved the highest average training accuracy within the allowed cost:

$$m^* = \arg\max_{m_i \in \mathcal{M},\ \text{cost}(m_i) \leq c} \text{TrainAcc}(m_i)$$

- **Oracle Router**: This upper-bound baseline assumes access to the ground truth performance of all models at test time. For each query, it routes to the best model among those allowed by the cost constraint:

$$m^* = \arg\max_{m_i \in \mathcal{M},\ \text{cost}(m_i) \leq c} \text{metric}(m_i, q)$$

It represents the best possible routing performance under the given budget.

F Details about Evaluation Metrics and Deferrel Curve

Evaluation Metric. We evaluate routing performance using metrics aligned with each benchmark's design. For RouterBench [7] and EmbedLLM [28], the correctness label is binary—each LLM either answers a query correctly or not. For MixInstruct [10], we adopt the exponentiated BARTScore, following prior work [11, 10]. While MixInstruct was originally intended to benchmark ensemble generation quality from outputs of multiple LLMs, recent works have adapted it for routing by assigning scores to individual LLM responses based on similarity to GPT-4. However, this introduces a dependency on GPT-4 as a reference model, which we will discuss further in Section 4.4.

Deferral Curve. Routing quality is visualized using a *deferral curve*, where the x-axis corresponds to the model cost budget and the y-axis reflects routing quality (accuracy or exp(BARTScore)). The cost budget represents the maximum cost (e.g., in dollars) a router can spend per query. However, because actual API pricing varies and is not always available, prior work [11] approximates cost using the number of model parameters—a practical proxy that correlates with both latency and financial cost for EmbedLLM [28] and MixInstruct [10]. This deferral curve captures the trade-off between routing quality and resource usage, allowing comparison of different routing strategies under cost constraints.

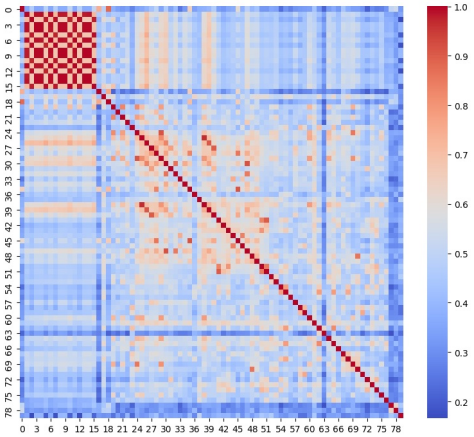

Figure 7: Category similarity heatmap based on average query embeddings. Redundancy is visible across GPQA-like categories (Upper-Left).

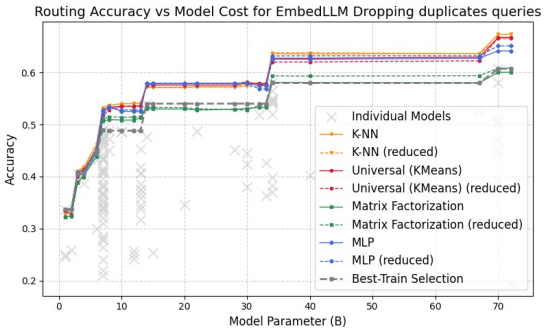

Figure 8: Routing accuracy when removing duplicate categories (e.g., GPQA variants). Performance is preserved even under OOD evaluation.

G Supplementary Result for Task Diversity

Here we provide more results and discussions on **Task Imbalance and Redundancy Problem** in Section 4.2. As shown in Figure 7, several task categories exhibit high similarity in their average query embeddings. For instance, GPQA-like categories cluster tightly in the embedding space, suggesting that they may not offer distinct routing challenges.

In our experiments, we found that even after removing duplicate categories from the training set—those identified as redundant in the heatmap—the router still performs strongly. Figure 8 shows that this holds true even under OOD evaluation, where the dropped categories are tested at inference time. This suggests that current benchmarks may overestimate the value of large or diverse-looking training sets when, in reality, much of the routing signal is concentrated in a smaller, more representative subset of tasks. We also empirically assess the redundancy within categories, where we progressively dropped a portion of training data within each category and retrained the router.

H Supplementary Result for Model Diversity

Here we provide more results and discussions **Model Redundancy Problem** in Section 4.3. We observed redundancy in the model pool, as evidenced by overlapping performance points (Individual Models' grey crossings) across cost settings in Figure 9. Such redundancy adds little value for training or evaluating router performance. We quantify model-level similarity using a Jaccard-style score based on shared correct predictions:

$$\text{sim}(m_i, m_j) = \frac{|\{q \mid m_i(q) = 1 \wedge m_j(q) = 1\}|}{|\{q \mid m_i(q) = 1 \vee m_j(q) = 1\}|}$$

where $m_i(q)$ denotes whether model m_i answered query q correctly. This metric captures functional overlap across the entire benchmark.

To validate this, we propose a greedy pruning strategy to reduce model redundancy while preserving routing effectiveness. At each step, we compute a score for each model based on:

$$\text{score}(m_i) = \lambda \cdot \text{Accuracy}(m_i) - (1 - \lambda) \cdot \text{AvgSim}(m_i)$$

where $\text{AvgSim}(m_i)$ is the average Jaccard similarity of model m_i to all other models (based on overlapping correct predictions), and λ balances performance versus uniqueness. The model with the lowest score is removed, and the process repeats until a target number of models remains.

We apply this strategy to the `EmbedLLM` benchmark, reducing the model pool from 112 to 82 (a 27% reduction). As shown in Figure 9, routing performance across methods remains comparable to the

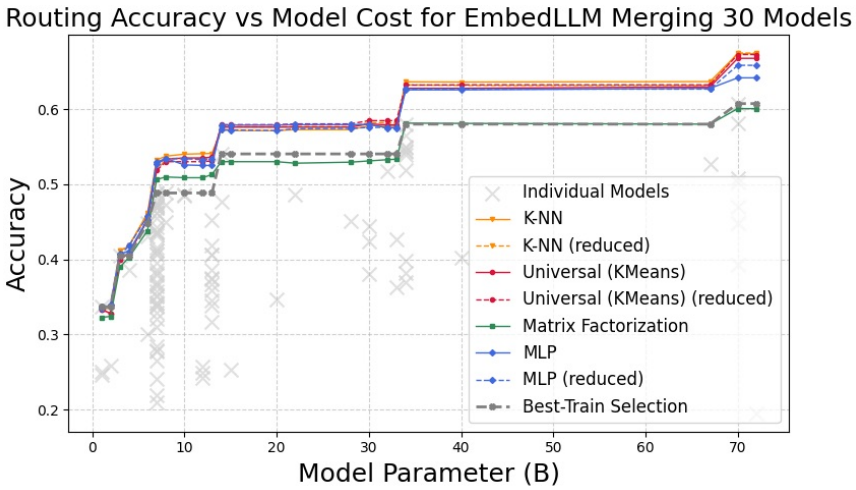

Figure 9: Performance comparison after reducing the model pool by 30 models. This shows that routers can maintain routing effectiveness across different cost budgets.

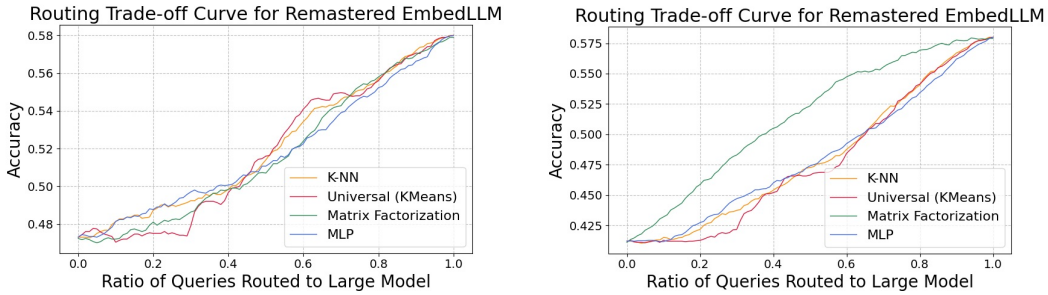

(a) Llama-2 7B and CausalLM-34B-Beta (b) Mistral 7B and CausalLM-34B-Beta

Figure 10: Binary routing evaluation paradigm showing performance trade-offs.

full model pool. This demonstrates that removing redundant models does not degrade routing quality and that meaningful routing decisions can still be made with a leaner model pool.

I Supplementary Result for Evaluation Methodology

I.1 Binary Routing Evaluation Details

To complement the traditional cost-accuracy deferral curves, we introduced a binary routing evaluation paradigm in Section 4.4 to assess how effectively a router balances between a strong generalist (large model) and a lightweight alternative (small model). Here, we provide additional details about the evaluation setup and key observations.

We fix the large model to be `CausalLM-34B-Beta`, given its similar superior performance across a wide range of general-purpose tasks, comparable to that of 70B-sized models. For small models, we consider two widely used options: `Mistral-7b-v0.1` and `LLaMA-2-7b-chat-hf`. These models represent different trade-offs in model families and capability, making them ideal candidates for evaluating routing flexibility.

In this setting, each routing method ranks the queries by its confidence score for the small model and routes a varying fraction of queries accordingly, as in Figure 10. The remaining queries are deferred to the large model. This produces a continuous accuracy curve as a function of the fraction of queries routed to the large model.

Across both small model settings, we observe that learned routers generally follow a linear trade-off curve, indicating that they lack precise mechanisms to identify which queries can be reliably handled

Table 7: OOD Performance change on selected categories in EMBEDLLM when these categories are excluded from training.

Category	K-NN Δ	KMeans Δ	MF Δ	MLP Δ
mathqa	-9.29	-16.88	-6.33	-14.34
asdiv	-58.59	-69.19	-40.40	-57.07
gsm8k	-14.28	-14.29	-29.47	-35.72
medmcqa	-11.58	-7.91	-6.78	-9.89
mmlu_clinical_knowledge	0.00	+7.41	-14.82	-3.70
Average	-18.75	-20.17	-19.56	-24.14

by the small model. Notably, clustering-based methods perform sub-linearly at lower deferral ratios, suggesting they often misclassify harder queries as easy ones and route them to the small models. This reinforces the need for more fine-grained routing strategies that can better distinguish between simple and complex inputs. Surprisingly, Matrix Factorization performed extremely well on classifying between Mistral-7B and CausalLM-34B-Beta, suggesting the potential of learning-based methods in certain model pair settings.

I.2 OOD Routing Evaluation Details

We evaluate the robustness of routing methods under out-of-distribution (OOD) scenarios by training and evaluating routers on different domains. We consider two distinct OOD settings: (1) excluding all math-related queries (e.g., `mathqa`, `asdiv`, `gsm8k`), and (2) excluding all medical-related queries (e.g., `medmcqa`, `mmlu_clinical_knowledge`). These categories are chosen for their semantic distinctiveness and task specificity, providing strong settings to evaluate how well routers generalize to unseen topics.

As shown in Table 7, all routing methods suffer performance degradation in OOD settings, with the most significant drops occurring on `asdiv` and `gsm8k`. MLP-based routers tend to experience the steepest accuracy declines overall, while matrix factorization (MF) demonstrates greater robustness, particularly on math-related tasks.

These results highlight that existing routing strategies are brittle when deployed in domains unseen during training, reinforcing the need for more semantically aware or domain-adaptive routing mechanisms.

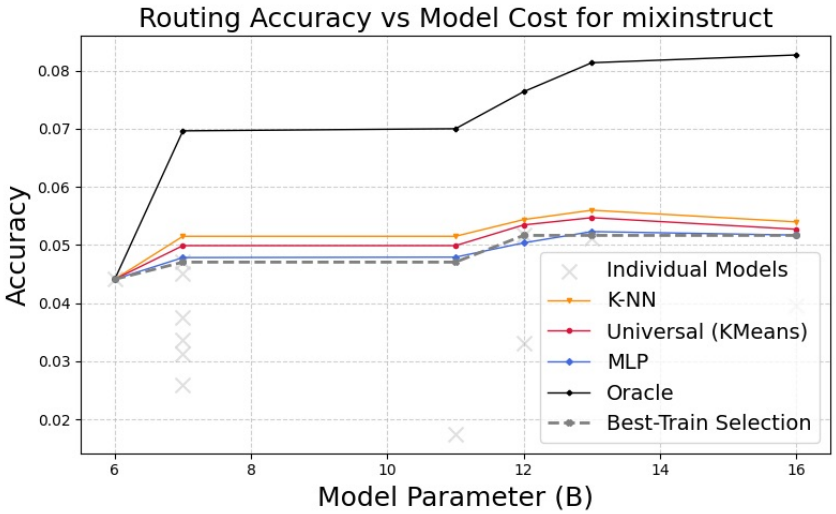

Figure 11: Routing performance on MIX-INSTRUCT.

J Supplementary Result on MIX-INSTRUCT

In Figure 11, we present the routing results in deferral curve on MIX-INSTRUCT dataset. While the same baseline routers are evaluated, we do not consider MIX-INSTRUCT as our primary benchmark due to several limitations:

- **Limited Evaluation Metrics:** MIX-INSTRUCT uses BARTScore to measure the similarity between a model's output and a reference response generated by GPT-4. This approach conflates model quality with similarity to GPT-4, making it less suitable for evaluating true routing performance. It favors models that mimic GPT-4's phrasing—even when other models might generate more informative or appropriate responses—thus undermining the purpose of routing for capability-based model selection.
- **Limited Task Diversity:** The benchmark contains only five tasks, all of which fall under casual or instruction-following dialog. These tasks do not capture the breadth of real-world user queries, particularly in domains requiring specialized knowledge (e.g., science, math, law), thereby limiting the opportunity for routing to leverage model specialization.
- **Restricted Model Pool:** MIX-INSTRUCT covers about 10 models—comparable to RouterBench—restricting the expressiveness of routing policies. In contrast, EMBEDLLM benchmark includes over 100 models with diverse strengths while having some issues we listed in Section 4, offered a more realistic and rigorous setting for evaluating routing capabilities.