

OmniRouter: Budget and Performance Controllable Multi-LLM Routing

Kai Mei, Wujiang Xu, Minghao Guo, Shuhang Lin, Yongfeng Zhang
Department of Computer Science, Rutgers University

{kai.mei, wujiang.xu, minghao.guo, shuhang.lin, yongfeng.zhang}@rutgers.edu

ABSTRACT

Large language models (LLMs) deliver superior performance but require substantial computational resources and operate with relatively low efficiency, while smaller models can efficiently handle simpler tasks with fewer resources. LLM routing is a crucial paradigm that dynamically selects the most suitable large language models from a pool of candidates to process diverse inputs, ensuring optimal resource utilization while maintaining response quality. Existing routing frameworks typically model this as a locally optimal decision-making problem, selecting the presumed best-fit LLM for each query individually, which overlooks global budget constraints, resulting in ineffective resource allocation. To tackle this problem, we introduce OmniRouter, a fundamentally controllable routing framework for multi-LLM serving. Instead of making per-query greedy choices, OmniRouter models the routing task as a constrained optimization problem, assigning models that minimize total cost while ensuring the required performance level. Specifically, a hybrid retrieval-augmented predictor is designed to predict the capabilities and costs of LLMs. After obtaining the predicted cost and performance, we utilize a constrained optimizer for cost-optimal assignments that employs Lagrangian dual decomposition with adaptive multipliers. It iteratively converges toward the globally optimal query-model allocation, dynamically balancing latency minimization against quality thresholds while adhering to heterogeneous capacity constraints. Experiments show that OmniRouter achieves up to 6.30% improvement in response accuracy while simultaneously reducing computational costs by at least 10.15% compared to competitive router baselines. The code and the dataset are available at <https://github.com/dongyuanjushi/OmniRouter>.

1. INTRODUCTION

Large Language Models (LLMs) have demonstrated remarkable capabilities, powering a diverse range of applications from chatbots [1; 45; 11; 15; 21; 53; 22; 50] and code assistants [23; 48; 30; 59] to computer-use agents. This success has spurred the widespread integration of LLMs into modern systems [44; 29; 32; 41; 57; 26; 49]. As LLM inference accelerates through hardware improvements, these LLM-integrated systems increasingly deploy not just one, but a set of LLMs with varying sizes, capabilities, and speeds as serving endpoints. This multi-LLM paradigm necessitates intel-

ligent routing: the crucial task of directing incoming user queries to the most appropriate LLM instance to balance performance goals and resource efficiency [39]. The design of effective LLM routers has become an active area of research. Recent works have proposed various routing strategies to optimize for cost [31; 42; 2], latency [25; 43; 33; 58], and performance [56; 24; 5; 60]. Most existing LLM routing frameworks predominantly treat routing as a sequence of independent, greedy decisions. For each incoming query, these routers select a model based on a local optimization criterion (e.g., lowest predicted latency, cheapest model predicted to succeed) without considering system-wide resource limitations or overall performance targets. These approaches fundamentally fail to achieve Pareto-efficient resource allocation across diverse query workloads distributed among multiple LLMs. When operating under budget constraints, these greedy routers cannot globally optimize the performance-resource tradeoff, leading to suboptimal overall system efficacy. For instance, expending computational resources on marginally improving responses to simple queries may leave insufficient capacity for complex queries where performance improvements would be more valuable, which can be seen in Figure 1. More critically, these localized decision frameworks cannot effectively enforce global constraints such as maintaining target response quality levels while adhering to strict resource budgets, resulting in either performance degradation or resource over-utilization when deployed at scale. In summary, these approaches confront a fundamental optimization dilemma: maintaining optimal overall performance while operating under computational resource constraints and stringent budgetary limitations.

Recognizing these limitations, we introduce OmniRouter, a controllable routing framework that fundamentally reimagines how query-model assignments should be determined. Unlike existing approaches, OmniRouter formalizes routing in multi-LLM systems as a **constrained optimization problem** [54; 3; 19] with global performance requirements and operational constraints and proposes a two-stage routing solution. Inspired by the success of retrieval-augmented generation (RAG) in reducing hallucination of LLMs, we design a retrieval-augmented predictor at the first stage to enhance the prediction-based router. At the second stage, we design the constrained optimizer which employs a dual gradient-based approach to navigate the solution space by adjusting model selections based on quality requirements and model concurrency. Our contributions are as follows:

- We propose OmniRouter, a routing framework for multi-LLM serving that fundamentally regards routing as a con-

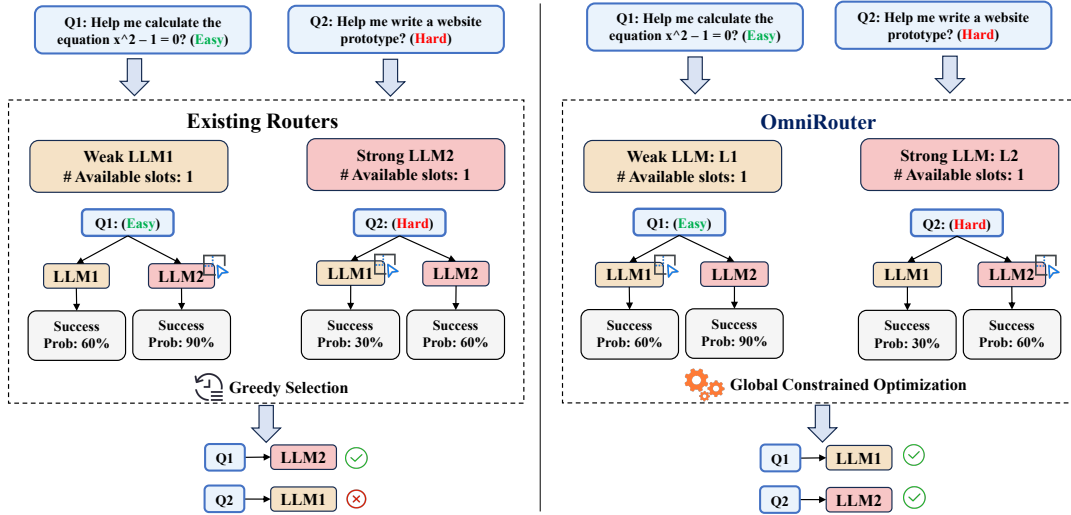


Figure 1: Comparison between traditional greedy routers and OmniRouter. Left: Greedy routers select models based on per-query optimization, leading to suboptimal allocations where Q1 (simple query) comes first and is assigned to the strong LLM, blocking Q2 (complex query) from accessing it. As a result, Q2 fails when assigned to the weak LLM. Right: OmniRouter employs constrained optimization to consider the global query distribution and model capabilities. It assigns the simple query to LLM1 (sufficient for the task) and reserves LLM2 for the complex query, thus maximizing overall success rate.

strained optimization problem rather than a series of greedy decisions, enabling system-level control of both performance and budget constraints.

- Extensive experiments show that OmniRouter can outperform competitive baselines across various serving scenarios and demonstrate significant improvements in response quality (up to 6.30%), cost efficiency (at least 10.15%), especially under tight operational constraints.

2. PROBLEM FORMALIZATION

Prior methods [31; 14; 34] typically score or rank models for each query independently based on heuristics such as cost-effectiveness or response likelihood. While effective in simple settings, such formulations struggle to account for system-level constraints such as limited model concurrency, global quality targets, and budget bounds that are critical in real-world multi-LLM deployments. In contrast, our approach models it as a constrained optimization problem.

Formally, given N queries and M models, let $a_{i,j} \in [0, 1]$ denote the capability of whether the model j can successfully answer query i , and $c_{i,j}$ denote the money cost for answering this query. Each model j is subject to a concurrency limit L_j , and the overall system must maintain a minimum average performance constraint α across all N queries to ensure response quality. Our objective is to assign queries to models such that the total cost is minimized, while ensuring that the overall performance exceeds α and that no model exceeds its concurrency constraint. This problem can be formulated as the following:

$$\min_x \sum_{i=1}^N \sum_{j=1}^M c_{i,j} x_{i,j} \quad s.t. \quad \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M a_{i,j} x_{i,j} \geq \alpha \quad (1)$$

$$\sum_{i=1}^N x_{i,j} \leq L_j, \forall j \quad \sum_{j=1}^M x_{i,j} = 1, \forall i$$

It is a constrained optimization problem [3; 19] with global

constraint α and local constraint L . We will elaborate how to solve this problem in the following section.

3. METHODOLOGY

As the capability $a_{i,j}$ and computational cost $c_{i,j}$ for each model-query pair are inherently uncertain at routing decision time. This uncertainty makes direct optimization of the routing variable $x_{i,j}$ particularly challenging. Joint optimization couples errors between prediction and allocation [13] and make the routing struggle with adaptability when query distributions shift or new models [52; 35; 55]. To address these challenges, we propose a two-stage approach that decouples the prediction of uncertain variables: $a_{i,j}$ and $c_{i,j}$ from the optimization of routing decisions $x_{i,j}$.

3.1 Retrieval-augmented Predictor

Inspired by retrieval-augmented generation (RAG) systems [27; 4], which enhance model outputs by incorporating relevant retrieved information, we design a retrieval-augmented predictor, as illustrated in Figure 2, to integrate the generalization capabilities of trained models with the retrieval precision of historical query-model data.

We employ bert-base-uncased [10] as our embedding encoder due to its representation capabilities and computational efficiency. This encoder processes both queries and LLM descriptions to generate embedding vectors E_q and E_l respectively, which serve as the foundation for both our training-based and retrieval-based prediction components. For the training-based component, we design a dual-head architecture built upon the embedding encoder. The first head focuses on model capability prediction, estimating a capability score $a_{i,j}^{pred} \in [0, 1]$ for each query i and model j through:

$$a_{i,j}^{pred} = \sigma(\mathbf{W}_1(E_q^i \cdot E_l^j) + \mathbf{b}_1) \quad (2)$$

where $\sigma(\cdot)$ denotes the sigmoid activation function, \mathbf{W}_1 and \mathbf{b}_1 are learnable parameters. The second head performs

sequence length classification by mapping the predicted output token length $l_{i,j}^{pred}$ into discrete buckets $\mathcal{B} = \{B_1, B_2, \dots, B_k\}$, where $k = \lceil \frac{l_{max}}{bs} \rceil$, l_{max} is the maximum sequence length, and bs is the bucket size. The probability of length $l_{i,j}^{pred}$ is computed as:

$$l_{i,j}^{pred} = bs \cdot \text{softmax}(\mathbf{W}_2(E_q^i + E_j^i) + \mathbf{b}_2)_i \quad (3)$$

The bucketing strategy, a technique also employed in previous works [25; 58; 14], transforms the challenging task of precise token-level prediction into a more tractable bucket-level approximation. During training, we optimize these dual objectives using mean squared error (MSE) for capability prediction and cross-entropy loss for length bucket classification.

Concurrently, our retrieval-based component leverages a vector database to identify the top- k similar historical queries based on cosine similarity between query embeddings. For a given query embedding E_q^i , let $\mathcal{Q}_k(E_q^i)$ be the set of top- k similar queries retrieved from the database. The retrieval-based capability score $a_{i,j}^{retrieve}$ and output length score $l_{i,j}^{retrieve}$ for query i and model j are computed as:

$$l_{i,j}^{retrieve} = \frac{\sum_{q_m \in \mathcal{Q}_k(E_q^i)} \text{sim}(E_q^i, E_{q_m}) \cdot l_{m,j}}{\sum_{q_m \in \mathcal{Q}_k(E_q^i)} \text{sim}(E_q^i, E_{q_m})} \quad (4)$$

$$a_{i,j}^{retrieve} = \frac{\sum_{q_m \in \mathcal{Q}_k(E_q^i)} \text{sim}(E_q^i, E_{q_m}) \cdot a_{m,j}}{\sum_{q_m \in \mathcal{Q}_k(E_q^i)} \text{sim}(E_q^i, E_{q_m})} \quad (5)$$

where $\text{sim}(E_q^i, E_{q_m})$ denotes the cosine similarity between the current query embedding and the retrieved query embedding, while $l_{m,j}$ and $a_{m,j}$ represent the historically observed output length and capability score for retrieved query m on model j , respectively. To obtain the final predictions, we integrate both components through an adaptive fusion mechanism for each query-model pair:

$$a_{i,j} = \gamma \cdot a_{i,j}^{pred} + (1 - \gamma) \cdot a_{i,j}^{retrieve} \quad (6)$$

$$c_{i,j} = \delta \cdot tp_j(l_{i,j}^{pred}) + (1 - \delta) \cdot tp_j(l_{i,j}^{retrieve}) \quad (7)$$

where γ and $\delta \in [0, 1]$ are learnable parameters that control the balance between trained and retrieved predictions, and $tp(\cdot)$ is a function that maps the predicted token length to a computational cost estimate based on model j 's token prices. This integrated retrieval-augmented architecture ensures our predictor benefits from both generalizable patterns learned during training and specific historical performance data, resulting in more accurate predictions of the uncertain variables needed for the subsequent optimization stage.

3.2 Constrained Optimizer

At the second stage, we solve the $x_{i,j}$ with the predicted $a_{i,j}$ and $c_{i,j}$. Inspired by optimizers [54; 3; 47], we leverage the Lagrangian dual theory and introduce Lagrangian multipliers to convert the primal problem into its Lagrangian dual problem. By introducing Lagrangian multipliers $\lambda_1, \lambda_{2,j}, \mu_i$, we get the Lagrangian relaxation function of the original

problem as follows:

$$\min_x \sum_{i=1}^N \sum_{j=1}^M c_{i,j} x_{i,j} \quad (8)$$

$$\text{s.t.} \quad \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M a_{i,j} x_{i,j} \geq \alpha$$

$$\sum_{i=1}^N x_{i,j} \leq L_j, \forall j \in [M] \quad \sum_{j=1}^M x_{i,j} = 1, \forall i \in [N].$$

we introduce three types of Lagrange multipliers as:

- $\lambda_1 \geq 0$ for the quality constraint (inequality)
- $\lambda_{2,j} \geq 0$ for each capacity constraint (inequality)
- μ_i for each assignment constraint (equality)

and write the Lagrangian function as:

$$L(x, \lambda_1, \lambda_{2,j}, \mu_i) \quad (9)$$

$$= \sum_{i=1}^N \sum_{j=1}^M c_{i,j} x_{i,j} + \lambda_1 \left(-\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M a_{i,j} x_{i,j} + \alpha \right) + \sum_{j=1}^M \lambda_{2,j} \left(\sum_{i=1}^N x_{i,j} - L_j \right) + \sum_{i=1}^N \mu_i \left(\sum_{j=1}^M x_{i,j} - 1 \right)$$

which can rearranged to group terms with $x_{i,j}$:

$$L(x, \lambda_1, \lambda_{2,j}, \mu_i) \quad (10)$$

$$= \sum_{i=1}^N \sum_{j=1}^M x_{i,j} \left(c_{i,j} - \frac{\lambda_1 a_{i,j}}{N} + \lambda_{2,j} + \mu_i \right) + \lambda_1 \alpha - \sum_{j=1}^M \lambda_{2,j} L_j - \sum_{i=1}^N \mu_i$$

The KKT optimality conditions for this problem are: Stationarity:

$$\frac{\partial L}{\partial x_{i,j}} = c_{i,j} - \frac{\lambda_1 a_{i,j}}{N} + \lambda_{2,j} + \mu_i = 0, \quad \forall i, j \quad (11)$$

Primal Feasibility:

$$\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M a_{i,j} x_{i,j} \geq \alpha, \quad \sum_{i=1}^N x_{i,j} \leq L_j, \forall j, \quad \sum_{j=1}^M x_{i,j} = 1, \forall i \quad (12)$$

Dual Feasibility:

$$\lambda_1 \geq 0, \quad \lambda_{2,j} \geq 0, \forall j \quad (13)$$

Complementary Slackness:

$$\lambda_1 \left(\alpha - \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M a_{i,j} x_{i,j} \right) = 0, \quad \lambda_{2,j} \left(L_j - \sum_{i=1}^N x_{i,j} \right) = 0, \quad \forall j \quad (14)$$

Note that for each i , we have the constraint: $\sum_{j=1}^M x_{i,j} = 1$. This implies that for each query i , exactly one LLM j must be selected. From the stationarity condition, for any fixed i , comparing two different indices j and k :

$$c_{i,j} - \frac{\lambda_1 a_{i,j}}{N} + \lambda_{2,j} + \mu_i = 0 \quad (15)$$




Figure 2: Illustration of OmniRouter, including the hybrid predictor and constrained optimizer.

$$c_{i,k} - \frac{\lambda_1 a_{i,k}}{N} + \lambda_{2,k} + \mu_i = 0 \quad (16)$$

Subtracting these equations eliminates μ_i :

$$(c_{i,j} - \frac{\lambda_1 a_{i,j}}{N} + \lambda_{2,j}) = (c_{i,k} - \frac{\lambda_1 a_{i,k}}{N} + \lambda_{2,k}) \quad (17)$$

This implies that for a given i , the optimal solution should choose the j that minimizes $(c_{i,j} - \frac{\lambda_1 a_{i,j}}{N} + \lambda_{2,j})$.

Then the dual function becomes:

$$g(\lambda_1, \lambda_2) \quad (18)$$

$$= \min_{x_{i,j}} \left\{ \sum_{i=1}^N \sum_{j=1}^M x_{i,j} (c_{i,j} - \frac{\lambda_1 a_{i,j}}{N} + \lambda_{2,j}) + \lambda_1 \alpha - \sum_{j=1}^M \lambda_{2,j} L_j \right\}$$

Note that μ_i has disappeared from the dual function because we've analytically incorporated the equality constraints. The dual problem can now be written as:

$$\max_{\lambda_1, \lambda_2} g(\lambda_1, \lambda_2), \quad \text{s.t. } \lambda_1 \geq 0, \quad \lambda_{2,j} \geq 0, \quad \forall j \in [M] \quad (19)$$

The partial derivatives for the remaining multipliers are: For λ_1 :

$$\frac{\partial L}{\partial \lambda_1} = -\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M x_{i,j} a_{i,j} + \alpha \quad (20)$$

For $\lambda_{2,j}$:

$$\frac{\partial L}{\partial \lambda_{2,j}} = \sum_{i=1}^N x_{i,j} - L_j, \quad \forall j \in [M] \quad (21)$$

The gradient ascent update rules are:

$$\lambda_1^{t+1} = \max \left(\lambda_1^t + \alpha_1 \left(-\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M x_{i,j} a_{i,j} + \alpha \right), 0 \right) \quad (22)$$

$$\lambda_{2,j}^{t+1} = \max \left(\lambda_{2,j}^t + \alpha_2 \left(\sum_{i=1}^N x_{i,j} - L_j \right), 0 \right), \quad \forall j \in [M] \quad (23)$$

For fixed multipliers, the optimal assignment for each i is:

$$x_{i,j} = \begin{cases} 1 & \text{if } j = j_i^* \\ 0 & \text{otherwise} \end{cases} \quad (24)$$

where

$$j_i^* = \arg \min_{j \in [M]} \left(c_{i,j} - \frac{\lambda_1 a_{i,j}}{N} + \lambda_{2,j} \right) \quad (25)$$

- λ_1 acts as a penalty for violating the quality constraint. When the average quality is below α , λ_1 increases, encouraging selection of higher-quality options.
- $\lambda_{2,j}$ penalizes capacity violations for each model j . When a model exceeds its capacity L_j , its corresponding multiplier increases, making it less attractive in subsequent iterations.
- The equality constraints (μ_i) are handled analytically by direct substitution, which simplifies the dual problem.
- The algorithm alternates between updating multipliers and assignments until convergence.

For $\lambda_{2,j}$, its partial derivative shows the workload violation for each model j . If a model's assigned queries exceed its capacity, the gradient is positive, increasing $\lambda_{2,j}$, which makes this overloaded model less attractive in subsequent iterations.

4. EVALUATION

In this section, we propose the following research questions regarding the performance of OmniRouter and conduct experiments to answer these research questions.

- **RQ1:** What is the routing performance of the OmniRouter and how is compared with existing routing frameworks?
- **RQ2:** Whether the OmniRouter can successfully control budget costs and ensure response accuracy when constraints (i.e., performance constraint α and concurrency constraint L) vary?
- **RQ3:** What is the influence of different factors (i.e., modules and parameters in predictors) of OmniRouter on routing performance?

4.1 Datasets

We collect 2.7k questions sourced from established knowledge and mathematical reasoning datasets, including MMLU [17], GPQA [38], MATH [18], and GSM8K [8]. We select 10 different models, including 5 relatively weak models, i.e., Qwen2.5 (7B-Instruct, 14B-Instruct, 32B-Instruct) [51], Gemma2 (9B-it, 27B-it) [46] and 5 relatively strong models, i.e., Qwen2.5-72B-Instruct, GPT-4o-mini, GPT-4o [1],

Gemini-1.5-flash [45], Claude-3.5-sonnet¹. And we collect the response correctness and token usage of these models to answer these questions using Llama-3.1-70B-Instruct [11] as the evaluation judge [28]. The Llama-3.1-70B is excluded from the LLM candidate pool to reduce potential biases. The prompt is shown as below.

Prompt for using LLM as the judge.

Prompt: The ground truth answer is: {gt_answer}. The prediction answer is: extracted_answer. Judge whether the prediction answer is correct or not. You just need to output ‘True’ or ‘False’.

The statistics of the data we use are in Table 1. We also set the difficulty for each question, which is determined by the number of models that can answer the question correctly, i.e., Easy: {8, 9, 10} models can answer correctly, Medium: {4, 5, 6, 7} models can answer correctly, Hard: {0, 1, 2, 3} models can answer correctly. And the proposition is: 78.4%, 15.2%, 6.4%, respectively.

Table 1: Statistics of the data distribution in our dataset.

Data Source	# Samples	Ratio
MMLU	1000	37.06%
GPQA-Diamond	198	7.33%
Math-500	500	18.53%
GSM8K	1000	37.06%

For calculating costs for LLMs to answer queries, we refer to litellm² for calculating money cost, the money price regarding the model we use are shown in Table 2.

4.2 Setup

We conduct our experiments on an Ubuntu 22.04 machine equipped with 8 RTX A5000 GPUs to serve multiple LLMs using Ollama⁴. During the implementation, we set the maximum output length to 1024 tokens to prevent the influence of extremely long output sequences, set top k as 16 during retrieval, and γ and δ as 0.5 during aggregation. To calculate money costs, we refer to Litellm’s costmap for price reference. The detailed price map can be found in Appendix B. Our experiments are conducted on the continuous-batching setting [26]. The serving system processes requests by dynamically forming batches based on the incoming query traffic, providing a more realistic evaluation environment for

¹<https://claude.ai/>

²<https://docs.litellm.ai/docs/>

⁴<https://ollama.com/>

Table 2: Money Cost Map of LLMs

Model Name	1M Input Tokens (\$)	1M Output Tokens (\$)
Qwen-2.5-7B-Instruct [51]	0.267	0.267
Qwen-2.5-14B-Instruct [51]	0.534	0.534
Qwen-2.5-32B-Instruct [51]	1.22	1.22
Qwen-2.5-72B-Instruct [51]	2.745	2.745
Gemma-2-9B-it [46]	0.343	0.343
Gemma-2-27B-it [46]	1.03	1.03
gpt-4o-mini [1]	0.15	0.6
gpt-4o [1]	2.5	10
gemini-1.5-flash [45]	0.075	0.3
claude-3.5-sonnet ³	3	15

routing frameworks. In our simulation, we model real-world traffic patterns by randomly adding $n \in \{1, 2, 3, 4\}$ queries to the queue every 0.1 seconds and performing routing decisions at 1-second intervals. If not specifically mentioned, we set the two constraints $\alpha = 0.75$ and the concurrent workload constraint $L = 4$, which are applied uniformly across all LLMs. Controllability analysis of different values of α and L is presented in Section 4.4. We evaluate OmniRouter against the following routing baselines:

Cost-oriented: We adapt S3 [25] and PO [58], which are originally designed for latency optimization as cost-oriented baselines. Specifically, S3 employs DistilBERT [40] and PO employs Vicuna-7B [7] as output token length predictors, respectively.

Performance-oriented: For performance-focused routing, we implement EmbedLLM [60] and RouterDC [5]. EmbedLLM leverages an encoder-decoder architecture to embed LLM representations, while RouterDC employs contrastive learning to model query-LLM relationships, both aiming to maximize response quality regardless of computational cost.

Cost-performance Coordinated: For baselines which balance both cost and performance objectives, we employ Hybrid-LLM [2], which constructs a probabilistic router, and CARROT [42], which implements dual predictors using Llama3-8B [11] as specified in its original implementation.

4.3 RQ1: Routing Performance

The comparison of overall routing performance can be observed from Table 3. OmniRouter demonstrates substantial advantages over all baselines across both metrics. Given the theoretical performance bounds (lower bound is 57.41% when the worst LLM for each query is selected and upper bound is 90% when the best LLM for each query is selected), the improvement achieved by OmniRouter (1.3% to 6.48% absolute improvement) is particularly significant, which represents a meaningful advancement within this constrained optimization space. When compared to cost-oriented baselines (S3 and PO), OmniRouter achieves significantly higher success rates (5.74% and 6.48% improvements respectively) while maintaining comparable or even lower costs. Against performance-oriented approaches such as EmbedLLM and RouterDC, our method achieves both higher success rates and substantially lower costs (approximately 41% cost reduction compared to RouterDC). Even when compared to cost-performance coordinated methods (Hybrid-LLM and CARROT), OmniRouter demonstrates superior performance in both dimensions.

Performance of Predictors. We also compare the effectiveness of OmniRouter’s predictor with the baselines we have mentioned. Table 4 demonstrates OmniRouter’s significant advantages in prediction accuracy. Our approach achieves 81.3% capability prediction accuracy, outperforming the strongest baseline (RouterDC) by 5.2%. For length prediction, OmniRouter delivers 45.2% exact match accuracy and 80.6% relaxed (± 1) accuracy, surpassing all alternatives by over 13.0%.

Routing Performance across Query Difficulty Levels. To assess the routing performance of OmniRouter deeper, we analyze its allocation patterns across different difficulty categories. Figure 3 reveals OmniRouter’s sophisticated routing strategy across query difficulty levels. For easy queries (78.4% of workload), the system successfully routes 49.2%

Table 3: Routing performance comparison, where we use accuracy, i.e., whether the routed LLM successfully answers the assigned query, and money cost of calling respective LLMs.

Metric	S3	PO	EmbedLLM	RouterDC	Hybrid-LLM	CARROT	OmniRouter
Accuracy (\uparrow)	69.45%	68.71%	72.96%	73.89%	71.48%	72.41%	75.19%
\$ Cost (\downarrow)	0.0585	0.0610	0.0896	0.0874	0.0722	0.0680	0.0515




Figure 3: Distribution of OmniRouter’s query routing decisions across difficulty levels and model capabilities.


Table 4: Performance of predictors on model capability and output token length prediction.

Method	Capability Acc.	Length Bucket Acc.	
		Exact Match	± 1
S3	–	0.333	0.656
PO	–	0.325	0.683
EmbedLLM	0.732	–	–
RouterDC	0.761	–	–
Hybrid-LLM	0.708	–	–
CARROT	0.745	0.312	0.672
OmniRouter	0.813	0.452	0.806

of them to weak models. Medium-difficulty queries receive balanced treatment with equal distribution between model types (7.6% each). For hard queries, OmniRouter shows a strong preference for routing to strong models by a 4:1 ratio (5.1% vs 1.3%), demonstrating the capability of OmniRouter to match appropriate models.

4.4 RQ2: Controllability Analysis

We investigate how different routing approaches respond to varying operational constraints, i.e., performance and concurrent constraints. Figure 4 reveals a critical limitation of existing greedy routing approaches: their inability to bound costs under stringent performance requirements. As the performance constraint (α) increases from 0.70 to 0.90, baseline methods exhibit exponential cost growth: at $\alpha = 0.90$, the cost of EmbedLLM reaches \$0.229, more than four times OmniRouter’s cost of \$0.054, creating a substantial cost gap ($\Delta = 0.175$). This cost explosion occurs because greedy routing strategies inherently prioritize immediate performance gains when performance thresholds rise, defaulting to routing increasingly more queries to the most powerful (and expensive) models regardless of actual query complexity. In contrast, OmniRouter demonstrates remarkably controlled cost scaling, with only a 48% increase over the same range and even a slight cost reduction at $\alpha > 0.85$. Similarly, in Figure 5, it demonstrates system behavior un-



(a) Cost vs. Performance Constraint (b) Accuracy vs. Performance Constraint

Figure 4: Impact of performance constraint (α) on cost efficiency and routing accuracy. As performance requirements increase, greedy methods exhibit unbounded cost growth while OmniRouter’s constraint optimization maintains controlled scaling.

Table 5: Effects of removing different modules in the predictor on routing performance.


Predictor	Capability Acc (\uparrow)	Bucket Acc (\uparrow)	Acc (\uparrow)	\$ Cost (\downarrow)
Full	0.813	0.806	75.19%	0.0515
w/o Retrieval	0.751	0.724	71.85%	0.0698
w/o Training	0.728	0.682	73.33%	0.0642

der varying concurrency constraints, revealing another key weakness of greedy approaches. As available concurrency decreases from 8 to 1, baseline methods show substantial cost increases (EmbedLLM: 159%, RouterDC: 147%) and significant performance degradation. When concurrency is highly limited ($L = 1$), the performance gap between OmniRouter and the weakest baseline (S3) becomes dramatic ($\Delta = 11.7%$). This occurs because greedy strategies struggle to make effective compromises when resource constraints tighten. By contrast, OmniRouter substantially performs better, maintaining 73.8% accuracy even at $L = 1$ while experiencing only moderate cost increases.

4.5 RQ3: Ablation Studies

Effects of Removing Modules in the Predictor. To evaluate the contribution of each component in our retrieval-augmented predictor, we conduct ablation studies by removing key modules from the full model. Results are presented in Table 5.

Removing the retrieval component leads to a substantial drop in both capability accuracy (7.6% decrease from 0.813 to 0.751) and length prediction accuracy (10.2% decrease from 0.806 to 0.724). This degradation directly impacts routing effectiveness, resulting in lower success rate (4.4% decrease) and significantly higher operational costs (35.5% increase from \$0.0515 to \$0.0698). These findings highlight that historical query information provides critical contextual signals that enhance prediction. On the other hand, after removing the training-based component, although performance accuracy sees a larger decline (10.5% decrease to



(a) Cost vs. Concurrency Constraint (b) Accuracy vs. Concurrency Constraint

Figure 5: Impact of concurrency constraint (L) on cost efficiency and routing accuracy. As available parallelism decreases, greedy methods struggle to make effective compromises, while OmniRouter’s constraint optimization maintains balanced allocations.

Table 6: Effects of number of buckets used in predictors on routing performance.

# Bucket	Bucket Acc (\uparrow)		Capability Acc. (\uparrow)	\$ Cost (\downarrow)
	Exact Match	± 1		
10	0.452	0.806	75.19%	0.0515
20	0.269	0.526	73.22%	0.0725
50	0.162	0.299	73.22%	0.834
100	0.124	0.189	70.74%	0.0962

0.728) compared to the no-retrieval variant, the success rate decreases less dramatically (only 2.5% reduction). This suggests that the retrieval mechanism alone sometimes can maintain reasonable routing decisions, which further validate our retrieval-augmented choice to enhance predictors.

Effects of Key Parameters within OmniRouter’s Predictors. We investigate the sensitivity of OmniRouter’s routing effectiveness to key design parameters within its prediction components. As shown in Table 6, the number of buckets used for discretizing output length predictions significantly impacts routing performance. A smaller number of buckets (10) yields substantially higher prediction accuracy (45.2% exact match, 80.6% ± 1 accuracy), which directly leads to higher routing performance (75.19% success rate) and lower operational costs (\$0.0515). This finding suggests that coarser-grained discretization may be sufficient for effective routing decisions, while excessive granularity can introduce noise. Regarding the retrieval component, Table 7 demonstrates how varying K (the number of historical samples) affects prediction accuracy and routing performance. The system achieves optimal performance at $K = 16$, with the highest success rate (75.19%) and lowest cost (\$0.0515). We observe that prediction accuracy remains relatively stable across moderate K values (8-32), with only slight degradation at the extremes. Too few samples ($K = 4$) provide insufficient historical context, while too many ($K = 64$) likely introduce noise from less relevant queries, both resulting in suboptimal routing decisions. This indicates the retrieval-augmentation requires careful calibration to balance between relevant historical information and potentially misleading outliers.

5. THEORETICAL PROOF

Proposition 1. *Given the Lagrangian dual problem with*

Table 7: Effects of variations of K in historical samples on routing performance.

K	Capability Acc. (\uparrow)	Length Acc (\uparrow)		SR (\uparrow)	\$ Cost (\downarrow)
		Exact Match	± 1		
4	0.773	0.430	0.784	74.07%	0.0614
8	0.785	0.443	0.786	74.81%	0.0565
16	0.806	0.452	0.806	75.19%	0.0515
32	0.785	0.440	0.776	74.81%	0.0565
64	0.780	0.422	0.764	73.70%	0.0684

multipliers $\lambda_1 \geq 0$ and $\lambda_{2,j} \geq 0$ for $j \in [M]$, the optimal assignment $x_{i,j}^*$ for the constrained optimization problem is uniquely determined by:

$$x_{i,j}^* = \begin{cases} 1 & \text{if } j = \arg \min_{j' \in [M]} \left(c_{i,j'} - \frac{\lambda_1^* a_{i,j'}}{N} + \lambda_{2,j'}^* \right) \\ 0 & \text{otherwise} \end{cases} \quad (26)$$

PROOF. We start with the original optimization problem:

$$\begin{aligned} \min_x \quad & \sum_{i=1}^N \sum_{j=1}^M c_{i,j} x_{i,j} \\ \text{s.t.} \quad & \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M a_{i,j} x_{i,j} \geq \alpha \quad \sum_{i=1}^N x_{i,j} \leq L_j, \quad \forall j \in [M] \\ & \sum_{j=1}^M x_{i,j} = 1, \quad \forall i \in [N] \quad x_{i,j} \in \{0, 1\}, \quad \forall i, j \end{aligned} \quad (27)$$

The Lagrangian function is:

$$\begin{aligned} L(x, \lambda_1, \lambda_2, \mu) & \\ &= \sum_{i=1}^N \sum_{j=1}^M c_{i,j} x_{i,j} + \lambda_1 \left(\alpha - \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M a_{i,j} x_{i,j} \right) \\ &+ \sum_{j=1}^M \lambda_{2,j} \left(\sum_{i=1}^N x_{i,j} - L_j \right) + \sum_{i=1}^N \mu_i \left(\sum_{j=1}^M x_{i,j} - 1 \right) \end{aligned} \quad (28)$$

From the KKT stationarity condition:

$$\frac{\partial L}{\partial x_{i,j}} = c_{i,j} - \frac{\lambda_1 a_{i,j}}{N} + \lambda_{2,j} + \mu_i = 0 \quad (29)$$

where μ_i is the Lagrangian multiplier for the constraint $\sum_{j=1}^M x_{i,j} = 1$.

Rearranging the stationarity condition:

$$\mu_i = - \left(c_{i,j} - \frac{\lambda_1 a_{i,j}}{N} + \lambda_{2,j} \right) \quad (30)$$

Since $\sum_{j=1}^M x_{i,j} = 1$ and $x_{i,j} \in \{0, 1\}$, exactly one $x_{i,j}$ equals 1 for each query i . For the optimal solution, we choose j^* that minimizes the expression in the stationarity condition:

$$j^* = \arg \min_{j \in [M]} \left(c_{i,j} - \frac{\lambda_1 a_{i,j}}{N} + \lambda_{2,j} \right) \quad (31)$$

Therefore, $x_{i,j^*} = 1$ and $x_{i,j} = 0$ for $j \neq j^*$, which completes the proof. \square

Proposition 2. *Given the optimal dual variable λ_1^* obtained from the gradient ascent update, if $\lambda_1^* > 0$, then the quality constraint is active and satisfied with equality:*

$$\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M a_{i,j} x_{i,j}^* = \alpha \quad (32)$$

PROOF. From the complementary slackness condition in the KKT conditions:

$$\lambda_1^* \left(\alpha - \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M a_{i,j} x_{i,j}^* \right) = 0 \quad (33)$$

This equation states that either:

- (i) $\lambda_1^* = 0$, or
- (ii) $\alpha - \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M a_{i,j} x_{i,j}^* = 0$

Given that $\lambda_1^* > 0$, condition (i) is false, thus condition (ii) must hold:

$$\alpha - \frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M a_{i,j} x_{i,j}^* = 0 \quad (34)$$

Rearranging yields:

$$\frac{1}{N} \sum_{i=1}^N \sum_{j=1}^M a_{i,j} x_{i,j}^* = \alpha \quad (35)$$

This demonstrates that when $\lambda_1^* > 0$, the quality constraint is satisfied with equality. The dual variable λ_1^* acts as a "price" for quality: the optimizer finds the solution that exactly meets the quality threshold α while minimizing cost. \square

6. RELATED WORK

LLM Routing. To optimize performance and inference cost in multi-LLM serving systems, researchers have developed increasingly sophisticated routing approaches. RouteLLM [31] and CARROT [42] propose training small LLMs as routers to balance cost and performance. EmbedLLM [60] introduces a specialized encoder-decoder network for embedding LLM representations. HybridLLM [2] addresses data imbalance issues by proposing a probabilistic router to better represent different model capabilities across query types. C2MAB-V [9] treats routing as a contextual multi-armed bandit problem, enabling exploration-exploitation trade-offs in model selection. GraphRouter [12] leverages graph learning to jointly modeling the query-model, query-query, and model-model relationship for building routers. While these approaches have made substantial progress, they predominantly operate as greedy decision-makers, optimizing for individual queries without considering global system constraints or the implications for subsequent requests. DeepSieve [16] utilizes LLM-as-a-Knowledge-Router to enable effective reasoning across multiple heterogeneous knowledge sources, enhancing retrieval-augmented generation (RAG) systems by dynamically routing sub-queries to the most appropriate source.

LLM Generation Length Prediction. Predicting LLM generation length is crucial for optimizing computational resources. Early attempts like Magnus [6] employed random forest algorithms but achieved limited accuracy. Subsequent research has explored two main directions of prediction models: encoder-only models for classification (DynamoLLM [43], S3 [25], TerriInfer [20], SSJF [37], and μ 3 [36]) and decoder-only models for generative prediction like Perception-only (PO) [58]. LTR [14] reformulated this as a ranking problem and utilized listwise ranking for predictor training.

7. CONCLUSION

In this paper, we introduce OmniRouter, a routing framework that fundamentally reframes LLM routing as a constrained optimization problem rather than a series of greedy decisions. We develop a two-stage approach with a hybrid predictor that accurately estimates model capabilities and costs, and a constrained optimizer that minimizes operational expenses while satisfying performance requirements. Our experiments with the OmniRouteEval dataset demonstrate that OmniRouter achieves up to 6.30% higher response accuracy while reducing costs by at least 10.15% compared to existing methods. OmniRouter maintains remarkable stability under varying constraints, precisely where greedy approaches fail. These results confirm OmniRouter's effectiveness for more realistic environments requiring both performance guarantees and cost control.

8. REFERENCES

- [1] ACHIAM, J., ADLER, S., AGARWAL, S., AHMAD, L., AKKAYA, I., ALEMAN, F. L., ALMEIDA, D., ALTENSCHMIDT, J., ALTMAN, S., ANADKAT, S., ET AL. Gpt-4 technical report. *arXiv preprint:2303.08774 (OpenAI Technical Report)* (2023).
- [2] ANONYMOUS. Hybrid llm: Cost-efficient and quality-aware query routing. *arXiv preprint arXiv:2404.14618* (2024).
- [3] BERTSEKAS, D. P. *Constrained optimization and Lagrange multiplier methods*. Academic press, 2014.
- [4] BORGEAUD, S., MENSCH, A., HOFFMANN, J., CAI, T., RUTHERFORD, E., MILLICAN, K., VAN DEN DRIESSCHE, G. B., LESPIAU, J.-B., DAMOC, B., CLARK, A., ET AL. Improving language models by retrieving from trillions of tokens. In *International Conference on Machine Learning* (2022), PMLR, pp. 2206–2240.
- [5] CHEN, S., JIANG, W., LIN, B., KWOK, J., AND ZHANG, Y. Routerdc: Query-based router by dual contrastive learning for assembling large language models. *Advances in Neural Information Processing Systems 37* (2024), 66305–66328.
- [6] CHENG, K., HU, W., WANG, Z., DU, P., LI, J., AND ZHANG, S. Enabling efficient batch serving for lmaas via generation length prediction. *arXiv preprint arXiv:2406.04785* (2024).
- [7] CHIANG, W.-L., LI, Z., LIN, Z., SHENG, Y., WU, Z., ZHANG, H., ZHENG, L., ZHUANG, S., ZHUANG, Y., GONZALEZ, J. E., ET AL. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. See <https://vicuna.lmsys.org> (accessed 14 April 2023) 2, 3 (2023), 6.
- [8] COBBE, K., KOSARAJU, V., BAVARIAN, M., CHEN, M., JUN, H., KAISER, L., PLAPPERT, M., TWOREK, J., HILTON, J., NAKANO, R., ET AL. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168* (2021).
- [9] DAI, X., LI, J., LIU, X., YU, A., AND LUI, J. Cost-effective online multi-llm selection with versatile reward models. *arXiv preprint arXiv:2405.16587* (2024).

- [10] DEVLIN, J. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [11] DUBEY, A., JAHHRI, A., PANDEY, A., KADIAN, A., AL-DAHLE, A., LETMAN, A., MATHUR, A., SCHELTEN, A., YANG, A., FAN, A., ET AL. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783 (Meta AI Technical Report)* (2024).
- [12] FENG, T., SHEN, Y., AND YOU, J. Graphrouter: A graph-based router for llm selections. *arXiv preprint arXiv:2410.03834* (2024).
- [13] FENG, X., WU, M., FENG, Y., YIN, X., WANG, S., ZENG, Y., ZENG, X., WANG, Z., QIN, R., HU, G., ET AL. Towards understanding and mitigating the training data quality of large language models. In *Advances in Neural Information Processing Systems* (2023), vol. 36.
- [14] FU, Y., ZHU, S., SU, R., QIAO, A., STOICA, I., AND ZHANG, H. Efficient llm scheduling by learning to rank. *arXiv preprint arXiv:2408.15792* (2024).
- [15] GUO, D., YANG, D., ZHANG, H., SONG, J., ZHANG, R., XU, R., ZHU, Q., MA, S., WANG, P., BI, X., ET AL. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948* (2025).
- [16] GUO, M., ZENG, Q., ZHAO, X., LIU, Y., YU, W., DU, M., CHEN, H., AND CHENG, W. Deepsieve: Information sieving via llm-as-a-knowledge-router. *arXiv preprint arXiv:2507.22050* (2025).
- [17] HENDRYCKS, D., BURNS, C., BASART, S., ZOU, A., MAZEIKA, M., SONG, D., AND STEINHARDT, J. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300* (2020).
- [18] HENDRYCKS, D., BURNS, C., KADAVATH, S., ARORA, A., BASART, S., TANG, E., SONG, D., AND STEINHARDT, J. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874* (2021).
- [19] HOMAIFAR, A., QI, C. X., AND LAI, S. H. Constrained optimization via genetic algorithms. *Simulation* 62, 4 (1994), 242–253.
- [20] HU, C., HUANG, H., XU, L., CHEN, X., XU, J., CHEN, S., FENG, H., WANG, C., WANG, S., BAO, Y., ET AL. Inference without interference: Disaggregate llm inference for mixed downstream workloads. *arXiv preprint arXiv:2401.11181* (2024).
- [21] HUA, W., YANG, X., JIN, M., LI, Z., CHENG, W., TANG, R., AND ZHANG, Y. Trustagent: Towards safe and trustworthy llm-based agents. In *2024 Conference on Empirical Methods in Natural Language Processing, EMNLP 2024* (2024), Association for Computational Linguistics (ACL), pp. 10000–10016.
- [22] HUA, W., ZHU, K., LI, L., FAN, L., LIN, S., JIN, M., XUE, H., LI, Z., WANG, J., AND ZHANG, Y. Disentangling logic: The role of context in large language model reasoning capabilities. *arXiv preprint arXiv:2406.02787* (2024).
- [23] HUI, B., YANG, J., CUI, Z., YANG, J., LIU, D., ZHANG, L., LIU, T., ZHANG, J., YU, B., LU, K., ET AL. Qwen2. 5-coder technical report. *arXiv preprint arXiv:2409.12186* (2024).
- [24] JIN, C., PENG, H., ZHANG, Q., TANG, Y., METAXAS, D. N., AND CHE, T. Two heads are better than one: Test-time scaling of multi-agent collaborative reasoning. *arXiv preprint arXiv:2504.09772* (2025).
- [25] JIN, Y., WU, C.-F., BROOKS, D., AND WEI, G.-Y. S3: Increasing gpu utilization during generative inference for higher throughput. *Advances in Neural Information Processing Systems* 36 (2023), 18015–18027.
- [26] KWON, W., LI, Z., ZHUANG, S., SHENG, Y., ZHENG, L., YU, C. H., GONZALEZ, J., ZHANG, H., AND STOICA, I. Efficient memory management for large language model serving with pagedattention. In *Proceedings of the 29th Symposium on Operating Systems Principles* (2023), pp. 611–626.
- [27] LEWIS, P., PEREZ, E., PIKTUS, A., PETRONI, F., KARPUKHIN, V., GOYAL, N., KÜTTLER, H., LEWIS, M., YIH, W.-T., ROCKTÄSCHEL, T., ET AL. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems* 33 (2020), 9459–9474.
- [28] LI, H., DONG, Q., CHEN, J., SU, H., ZHOU, Y., AI, Q., YE, Z., AND LIU, Y. Llm-as-judges: a comprehensive survey on llm-based evaluation methods. *arXiv preprint arXiv:2412.05579* (2024).
- [29] MEI, K., LI, Z., XU, S., YE, R., GE, Y., AND ZHANG, Y. Aios: Llm agent operating system. *arXiv e-prints, pp. arXiv-2403* (2024).
- [30] NIJKAMP, E., HAYASHI, H., XIONG, C., SAVARESE, S., AND ZHOU, Y. Codegen2: Lessons for training llms on programming and natural languages. *arXiv preprint arXiv:2305.02309* (2023).
- [31] ONG, I., ALMAHAIRI, A., WU, V., CHIANG, W.-L., WU, T., GONZALEZ, J. E., KADOUS, M. W., AND STOICA, I. Routellm: Learning to route llms with preference data. *arXiv preprint arXiv:2406.18665* (2024).
- [32] PACKER, C., FANG, V., PATIL, S. G., LIN, K., WOODERS, S., AND GONZALEZ, J. E. Memgpt: Towards llms as operating systems. *arXiv preprint arXiv:2310.08560* (2023).
- [33] PANDA, P., MAGAZINE, R., DEVAGUPTAPU, C., TAKEMORI, S., AND SHARMA, V. Adaptive llm routing under budget constraints. *arXiv preprint arXiv:2508.21141* (2025).
- [34] PARKAR, R. S., KIM, J., PARK, J. I., AND KANG, D. Selectllm: Can llms select important instructions to annotate? *arXiv preprint arXiv:2401.16553* (2024).
- [35] QIN, X., YANG, Y., LI, X., DONG, S., HUANG, S., JI, H., AND LI, L. Towards robust llm-based decision-making: A calibration and planning approach. In *International Conference on Learning Representations* (2023).

- [36] QIU, H., MAO, W., PATKE, A., CUI, S., JHA, S., WANG, C., FRANKE, H., KALBARCZYK, Z., BAŞAR, T., AND IYER, R. K. Power-aware deep learning model serving with $\{\mu\text{-Serve}\}$. In *2024 USENIX Annual Technical Conference (USENIX ATC 24)* (2024), pp. 75–93.
- [37] QIU, H., MAO, W., PATKE, A., CUI, S., JHA, S., WANG, C., FRANKE, H., KALBARCZYK, Z. T., BAŞAR, T., AND IYER, R. K. Efficient interactive llm serving with proxy model-based sequence length prediction. In *The 5th International Workshop on Cloud Intelligence / AIOps at ASPLOS 2024* (San Diego, CA, USA, 2024), vol. 5, Association for Computing Machinery, pp. 1–7.
- [38] REIN, D., HOU, B. L., STICKLAND, A. C., PETTY, J., PANG, R. Y., DIRANI, J., MICHAEL, J., AND BOWMAN, S. R. Gpqa: A graduate-level google-proof q&a benchmark. *arXiv preprint arXiv:2311.12022* (2023).
- [39] SALARIA, S., LIU, Z., AND GONZALEZ, N. M. Metametrics and best practices for system-level inference performance benchmarking, 2025.
- [40] SANH, V., DEBUT, L., CHAUMOND, J., AND WOLF, T. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108* (2019).
- [41] SHI, Z., MEI, K., JIN, M., SU, Y., ZUO, C., HUA, W., XU, W., REN, Y., LIU, Z., DU, M., DENG, D., AND ZHANG, Y. From commands to prompts: LLM-based semantic file system for aios. In *The Thirteenth International Conference on Learning Representations* (2025).
- [42] SOMERSTEP, S., POLO, F. M., DE OLIVEIRA, A. F. M., MANGAL, P., SILVA, M., BHARDWAJ, O., YUROCHKIN, M., AND MAITY, S. Carrot: A cost aware rate optimal router. *arXiv preprint arXiv:2502.03261* (2025).
- [43] STOJKOVIC, J., ZHANG, C., GOIRI, Í., TORRELLAS, J., AND CHOUKSE, E. Dynamollm: Designing llm inference clusters for performance and energy efficiency. *arXiv preprint arXiv:2408.00741* (2024).
- [44] SUN, B., HUANG, Z., ZHAO, H., XIAO, W., ZHANG, X., LI, Y., AND LIN, W. Llumnix: Dynamic scheduling for large language model serving. *arXiv preprint arXiv:2406.03243* (2024).
- [45] TEAM, G., ANIL, R., BORGEAUD, S., ALAYRAC, J.-B., YU, J., SORICUT, R., SCHALKWYK, J., DAI, A. M., HAUTH, A., MILLICAN, K., ET AL. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805* (2023).
- [46] TEAM, G., RIVIERE, M., PATHAK, S., SESSA, P. G., HARDIN, C., BHUPATIRAJU, S., HUSSENOT, L., MESSNARD, T., SHAHRIARI, B., RAMÉ, A., ET AL. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118* (2024).
- [47] WANG, C., SHI, X., XU, S., WANG, Z., FAN, Z., FENG, Y., YOU, A., AND CHEN, Y. A multi-stage framework for online bonus allocation based on constrained user intent detection. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining* (2023), pp. 5028–5038.
- [48] WEI, Y., WANG, Z., LIU, J., DING, Y., AND ZHANG, L. Magicoder: Source code is all you need. *arXiv preprint arXiv:2312.02120* (2023).
- [49] XU, W., LIANG, Z., MEI, K., GAO, H., TAN, J., AND ZHANG, Y. A-mem: Agentic memory for llm agents. *arXiv preprint arXiv:2502.12110* (2025).
- [50] XU, W., SHI, Y., LIANG, Z., NING, X., MEI, K., WANG, K., ZHU, X., XU, M., AND ZHANG, Y. iagent: Llm agent as a shield between user and recommender systems. *arXiv preprint arXiv:2502.14662* (2025).
- [51] YANG, A., YANG, B., ZHANG, B., HUI, B., ZHENG, B., YU, B., LI, C., LIU, D., HUANG, F., WEI, H., ET AL. Qwen2. 5 technical report. *arXiv preprint arXiv:2412.15115* (2024).
- [52] YU, P., TROMBETTA, P., HASSANI, A., BULITKO, V., AND WHITE, M. Gilbo: One metric to measure them all. In *Advances in Neural Information Processing Systems* (2022), vol. 35, pp. 10285–10297.
- [53] ZHANG, C., LIU, X., ZHANG, Z., JIN, M., LI, L., WANG, Z., HUA, W., SHU, D., ZHU, S., JIN, X., ET AL. When ai meets finance (stockagent): Large language model-based stock trading in simulated real-world environments. *arXiv preprint arXiv:2407.18957* (2024).
- [54] ZHANG, X., QI, F., HUA, Z., AND YANG, S. Solving billion-scale knapsack problems. In *Proceedings of The Web Conference 2020* (2020), pp. 3105–3111.
- [55] ZHANG, Y., SHI, X., FENG, L., WANG, M., YU, Y., YU, X., SHEN, H., CHEN, Z., MUCCI, P., KUDLUR, M., ET AL. S-lora: Serving thousands of concurrent lora adapters. In *Advances in Neural Information Processing Systems* (2023), vol. 36.
- [56] ZHANG, Y.-K., ZHAN, D.-C., AND YE, H.-J. Capability instruction tuning: A new paradigm for dynamic llm routing. *arXiv preprint arXiv:2502.17282* (2025).
- [57] ZHENG, L., YIN, L., XIE, Z., HUANG, J., SUN, C., YU, C., CAO, S., KOZYRAKIS, C., STOICA, I., GONZALEZ, J. E., ET AL. Efficiently programming large language models using sglang.
- [58] ZHENG, Z., REN, X., XUE, F., LUO, Y., JIANG, X., AND YOU, Y. Response length perception and sequence scheduling: An llm-empowered llm inference pipeline. vol. 36.
- [59] ZHU, Q., GUO, D., SHAO, Z., YANG, D., WANG, P., XU, R., WU, Y., LI, Y., GAO, H., MA, S., ET AL. Deepseek-coder-v2: Breaking the barrier of closed-source models in code intelligence. *arXiv preprint arXiv:2406.11931* (2024).
- [60] ZHUANG, R., WU, T., WEN, Z., LI, A., JIAO, J., AND RAMCHANDRAN, K. Embedllm: Learning compact representations of large language models. *arXiv preprint arXiv:2410.02223* (2024).