

Breaking Model Lock-in: Cost-Efficient Zero-Shot LLM Routing via a Universal Latent Space

Cheng Yan¹, Wuyang Zhang^{*1}, Zhiyuan Ning¹, Fan Xu¹, Ziyang Tao¹, Lu Zhang², Bing Yin³,
Yanyong Zhang^{*1,2}

¹University of Science and Technology of China

²Institute of Artificial Intelligence, Hefei Comprehensive National Science Center

³Research Department, iFLYTEK Co., LTD.

{yc_sa22218099, markxu, ziyangtao}@mail.ustc.edu.cn, {wuyangz, luzha, yanyongz}@ustc.edu.cn

Abstract

The rapid proliferation of Large Language Models (LLMs) has led to a fragmented and inefficient ecosystem, a state of “model lock-in” where seamlessly integrating novel models remains a significant bottleneck. Current routing frameworks require exhaustive, costly retraining, hindering scalability and adaptability. We introduce ZeroRouter, a new paradigm for LLM routing that breaks this lock-in. Our approach is founded on a universal latent space, a model-agnostic representation of query difficulty that fundamentally decouples the characterization of a query from the profiling of a model. This allows for zero-shot onboarding of new models without full-scale retraining. ZeroRouter features a context-aware predictor that maps queries to this universal space and a dual-mode optimizer that balances accuracy, cost, and latency. Our framework consistently outperforms all baselines, delivering higher accuracy at lower cost and latency.

Code — <https://github.com/Codeffun3/ZeroRouter>

Introduction

The landscape of Large Language Models (LLMs) is no longer monolithic. While massive models deliver peak performance, their computational cost makes them inefficient for routine tasks that smaller models could handle. Left to users, model selection is often manual and haphazard, leading to expensive failures: over-reliance on costly, oversized models for simple queries, and the underutilization of more efficient alternatives. Here, the core challenge is to automate this selection process, making routing systems that can dynamically navigate the trade-offs between cost, latency, and quality.

Current LLM routing frameworks adopt two paradigms with inherent scalability limitations. Direct routing systems, such as HybridLLM (Ding et al. 2024) and RouteLLM (Ong et al. 2024), use static model-task mappings or cascading workflows, which require exhaustive query evaluations across predefined models. These methods cannot integrate new models without architectural overhauls, as their logic is hardcoded for fixed model sets, and cascading workflows incur quadratic evaluation costs as the model pool grows. Inference-free predictive frameworks, such as MixLLM (Wang et al. 2025),

*Corresponding author

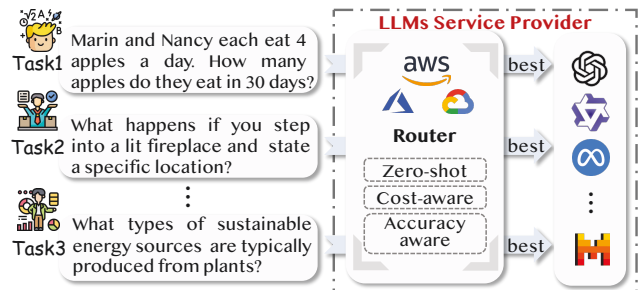


Figure 1: Example of Large Language Routing Mechanism. Queries are intelligently routed to the most suitable LLM from service providers for optimal performance and cost.

uses a contextual-bandit framework to balance quality, cost, and latency and continually adapts to user feedback. When new models are introduced, these predictors fail to capture task-query-LLM interactions critical for generalization. The primary limitation of most existing LLM routing systems is their static design. This creates a state of “model lock-in,” where the router is tied to a fixed set of models. Incorporating new LLMs usually requires a complete and costly overhaul, including extensive retraining with new preference data. They may generalize poorly to out-of-distribution queries and new tasks because they fail to grasp the intrinsic properties of a query, relying instead on superficial heuristics. Furthermore, their performance is often compromised by unreliable signals like miscalibrated model confidence, and they operate without awareness of real-time dynamics of model inference.

This paper introduces ZeroRouter, a framework that rebuilds LLM routing from the ground up to break free from these limitations. We propose a new paradigm centered on a universal, cross-task latent space, a multi-dimensional map that represents the intrinsic characteristics of any query in a model-agnostic way. This core innovation fundamentally decouples the task of characterizing a query from the task of profiling a model. New models can therefore be integrated seamlessly by efficiently charting their performance capabilities onto this universal map using only a small set of anchor queries, eliminating the need for costly, full-scale retraining. A context-aware predictor then maps raw text directly to a coordinate in this latent space, allowing a lightweight yet

highly intelligent routing decision. By framing routing as a multi-objective optimization problem, ZeroRouter provides a scalable and adaptive solution that assigns each query to the most cost-effective mode. Our framework consistently outperforms all baselines by simultaneously delivering higher accuracy at lower cost and latency. This lead is most significant on unseen datasets, where it achieves a higher maximum accuracy while also dramatically outperforming all rivals in minimizing cost and latency.

Related Work

LLM Routing

Existing approaches for LLM routing (Zhang, Zhan, and Ye 2025; Wang et al. 2025; Aggarwal et al. 2024; Mohammadshahi, Shaikh, and Yazdani 2024) face critical trade-offs between adaptability and efficiency. Router-trained systems such as HybridLLM (Ding et al. 2024) and RouteLLM (Ong et al. 2024) map queries to models via supervised learning, yet their transductive design locks them to predefined model sets where new LLMs demand full router retraining. The current frontier is inference-free predictive frameworks (e.g. FORC (Šakota, Peyrard, and West 2024)), which decouple performance prediction and optimization. MixLLM (Wang et al. 2025), uses a contextual-bandit framework to dynamically balance quality, cost, and latency and continually adapts to user feedback. While these methods represent significant progress, they face bottlenecks such as mandatory profiling for new models or reliance on heuristics for unseen queries.

Efficient Benchmarking of LLMs

Previous work reduces the cost of LLM evaluation with strategic sub-sampling. MixEval (Ni et al. 2024) selects subsets to predict aggregate benchmark scores, while TinyBenchmark (Polo et al. 2024) and MetaBench (Kipnis et al. 2024) leverage clustering or Item Response Theory (IRT) (Rodriguez et al. 2021) to identify diagnostic instances. Though these methods enable lightweight profiling, they prioritize dataset-level metric approximation and cannot generalize to unseen prompts. Recent efforts (e.g. How2Bench (Cao et al. 2025)) provides comprehensive guidelines for benchmark development, and the creation of specialized test suites for nuanced abilities. However, this creates a scalability paradox: as benchmarks become more numerous and rigorous, the cost to vet new models escalates, reinforcing a paradigm of exhaustive, slow evaluation. Our framework breaks this cycle by establishing a specific-model-agnostic difficulty scale, enabling lightweight profiling of new models without costly, ever-expanding benchmark evaluations.

Methodology

Our framework pioneers a zero-shot routing paradigm by establishing a context-aware link between a query’s intrinsic properties and its performance. Unlike prior works that require exhaustive evaluations or full retraining, ZeroRouter achieves scalable model integration through three key innovations: ❶ a universal, cross-task latent space, calibrated by an information-theoretic selection strategy, to decouple prompt characteristics from model capabilities; ❷ a context-aware

latent space predictor that maps a query’s textual properties directly to its latent parameters; and ❸ a unified policy-driven routing engine that translates high-level user priorities on accuracy, cost, and latency into optimal query assignments. Figure 2 illustrates the overall framework.

Universal Latent Space for Lightweight Profiling

This module provides lightweight accuracy, cost, and latency profiles for new LLMs by decoupling a prompt’s intrinsic latent parameters (e.g., discrimination, difficulty), which are model-agnostic, from a model’s specific ability profile. This separation allows highly efficient profiling of new models, requiring performance data only on a small set of anchor prompts to accurately estimate their capabilities.

Model Inference Accuracy Estimation

Cross-Task Discrimination and Difficulty Calibration. To create a universal discrimination and difficulty scale, we first model the interaction between a model u and a prompt i using a Probabilistic Model, such as multidim 2-Parameter Logistic (2PL) IRT model (Rodriguez et al. 2021). The probability of a correct response $X_{ui} = 1$ is given by:

$$P(X_{ui} = 1 | \theta_u, \alpha_i, b_i) = \frac{1}{1 + e^{-\alpha_i^\top (\theta_u - b_i)}}, \quad (1)$$

where θ_u is the model’s ability vector, while α_i and b_i are the prompt’s discrimination and difficulty vectors, respectively, all residing in a D -dimensional latent space ($\in \mathbb{R}^D$). We define a hierarchical Bayesian model over these latent variables (with priors $\theta_u \sim \mathcal{N}(\mu_\theta, \mathbf{u}_\theta^{-1})$, etc.) and perform efficient posterior inference by maximizing the Evidence Lower Bound (ELBO) via Stochastic Variational Inference (SVI). The key innovation is to build a universal scale via cross-task calibration, allowing the derivation of a minimal set of anchors for the efficient evaluation of new models.

Information-Theoretic Anchor Selection. To select a minimal yet highly informative anchor set $\mathcal{A} \subset \mathcal{P}$ with size $N \ll |\mathcal{P}|$ that effectively estimates model capabilities, we formalize the task as an optimal experimental design problem, with the aim of minimizing the uncertainty of the model’s estimated ability parameter θ_u . To achieve this, we employ the principle of D-optimality, which seeks to maximize the determinant of the Fisher Information Matrix:

$$\mathcal{I}_{\mathcal{A}}(\theta_u) = \sum_{i \in \mathcal{A}} P(X_{ui} = 1)(1 - P(X_{ui} = 1))\alpha_i \alpha_i^\top \quad (2)$$

Since the probability term in Eq. (2) depends on the unknown model ability θ_u , we simplify the objective to rely solely on the prompt discrimination vectors α_i , which capture the model-agnostic informativeness of each prompt. The objective is thus to maximize the log-determinant of the cumulative discrimination matrix:

$$\mathcal{A}^* = \arg \max_{\mathcal{A} \subset \mathcal{P}, |\mathcal{A}|=N} \log \det \left(\sum_{i \in \mathcal{A}} \alpha_i \alpha_i^\top \right) \quad (3)$$

To avoid directly solving the NP-hard optimization problem, we construct the set \mathcal{A} using a greedy forward selection algorithm. Starting with an empty set $\mathcal{A}_0 = \emptyset$ and a regularized

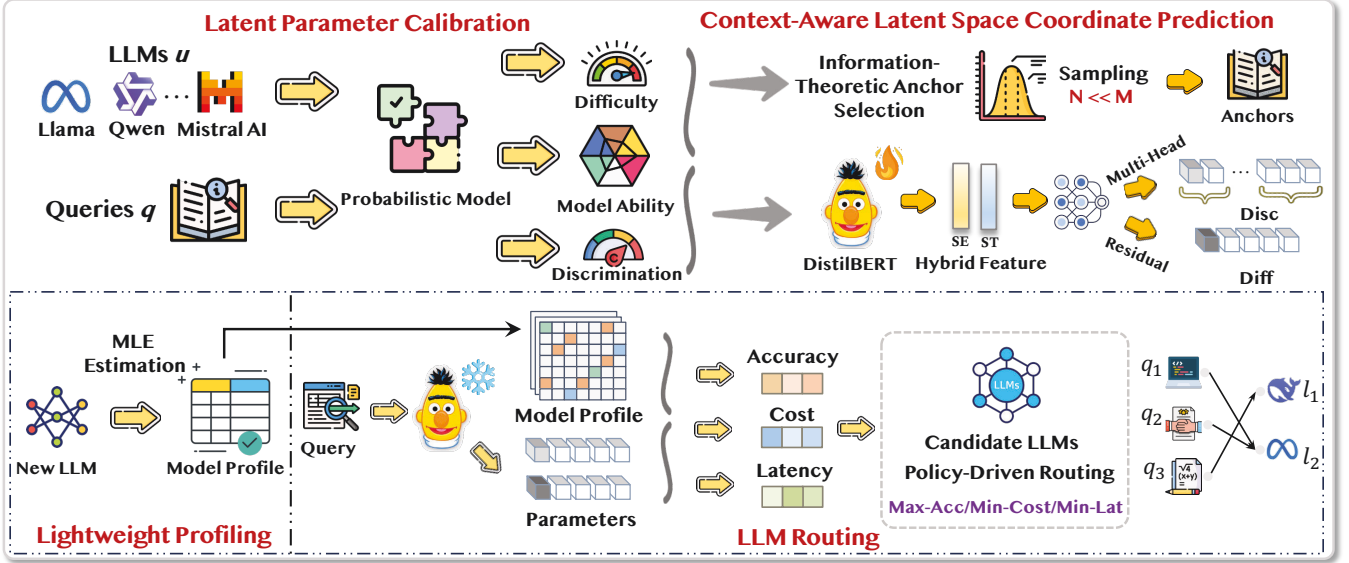


Figure 2: An overview of the ZeroRouter framework, detailing its three core modules: (1) **Latent Parameter Calibration**, which establishes a universal latent space using information-theoretic anchor selection; (2) **Lightweight Profiling**, which efficiently maps new LLMs onto this space; and (3) **Policy-Driven Routing**, which predicts a query’s latent coordinates and assigns it to the optimal model based on user-defined constraints for accuracy, cost, and latency.

matrix $\mathcal{I}_0 = \epsilon \mathbf{I}$ for numerical stability, the algorithm iteratively appends the prompt i_k^* that offers the maximal gain:

$$i_k^* = \arg \max_{i \in \mathcal{P} \setminus \mathcal{A}_{k-1}} \log \det(\mathcal{I}_{k-1} + \alpha_i \alpha_i^\top) \quad (4)$$

This process is repeated until $|\mathcal{A}_k| = N$. The resulting set of anchors is compact and highly informative, providing a solid foundation for cost-effective profiling.

Lightweight Profiling for New Models. With the information-theoretic anchor set \mathcal{A} established, we can efficiently estimate $\theta_{u_{\text{new}}}$ by minimizing the binary cross-entropy (BCE) loss:

$$\hat{\theta}_{u_{\text{new}}} = \arg \min_{\theta_{u_{\text{new}}}} \sum_{k \in \mathcal{A}} \text{BCE}(y_k, P_k), \quad (5)$$

where $y_k \in [0, 1]$ is the ground-truth score for anchor $k \in \mathcal{A}$, and P_k is the correctness probability predicted by Eq. (1) using the new model’s ability $\theta_{u_{\text{new}}}$. The term $\text{BCE}(y, p)$ defined as $-[y \ln(p) + (1 - y) \ln(1 - p)]$.

Model Inference Cost Estimation

Accurate cost estimation is a prerequisite for effective routing. For API-based LLM services, the total inference cost for routing a query q to a model u is a function of its input and output token lengths:

$$C_{uq} = \lambda_u^{\text{in}} \ell_{uq}^{\text{in}} + \lambda_u^{\text{out}} \ell_{uq}^{\text{out}}, \quad (6)$$

where λ_u^{in} and λ_u^{out} are the per-token API prices for model u , while ℓ_{uq}^{in} and ℓ_{uq}^{out} are the input and output token lengths.

Input Token Estimation. Estimating the input token length, ℓ^{in} , is deterministic. Given query q and model u , we employ the model-specific tokenizer, $\mathcal{T}_u(\cdot)$, to obtain a precise count:

$$\ell_{uq}^{\text{in}} = |\mathcal{T}_u(q)|. \quad (7)$$

Output Token Estimation. Unlike the deterministic input length, estimating the output token length ℓ_{uq}^{out} is non-trivial, as it depends on both query complexity and the specific model’s verbosity. We introduce a novel estimation method that directly leverages our universal latent space, bypassing the need for complex, per-model regression heads for this task. First, we define the task-aware difficulty s_q , derived from the predicted latent IRT parameter vectors:

$$s_q = \alpha_q^\top \mathbf{b}_q \quad (8)$$

During a one-time calibration phase using our small anchor set \mathcal{A} , we discretize the continuous range of complexity scores observed across the anchor samples into K bins, $\mathcal{B} = \{B_1, \dots, B_K\}$. For each candidate model $u \in \mathcal{M}$ and complexity bin B_k , we pre-compute and store a mean output token length, $\bar{\ell}_{uk}^{\text{out}}$, by averaging over the ground-truth lengths of anchor samples whose complexity scores fall into that bin:

$$\bar{\ell}_{uk}^{\text{out}} = \mathbb{E}_{q \in \mathcal{A} | s_q \in B_k} [|\mathcal{T}_u(\text{response}(q, u))|] \quad (9)$$

This process generates a lightweight lookup table mapping a (model, complexity_bin) tuple to an expected output length, efficiently profiling each model’s verbosity across different complexity tiers.

For an unseen query q' , the estimation process is remarkably efficient. We first predict its latent vectors $(\hat{\alpha}_{q'}, \hat{\mathbf{b}}_{q'})$ using our context-aware predictor. We then compute its estimated complexity score $\hat{s}_{q'} = \hat{\alpha}_{q'}^\top \hat{\mathbf{b}}_{q'}$, identify the corresponding bin B_k such that $\hat{s}_{q'} \in B_k$, and simply retrieve the pre-calibrated average length for the target model u :

$$\hat{\ell}_{uq'}^{\text{out}} = \bar{\ell}_{uk}^{\text{out}} \quad (10)$$

This decoupling of latent parameters from model execution supports zero-shot generalization, reducing output length estimation for any new query to a inference-free table lookup.

Model Inference Latency Estimation

We model latency for each model u via two pre-calibrated parameters: its Time to First Token (τ_u^{TFFT}), capturing initial processing overhead, and its Time per Output Token (τ_u^{TPOT}), reflecting its generation speed. These parameters are efficiently estimated for each model using the anchor set \mathcal{A} . The final latency for a query q is predicted by:

$$\hat{\tau}_{uq} = \tau_u^{\text{TFFT}} + \hat{\ell}_{uq}^{\text{out}} \cdot \tau_u^{\text{TPOT}}. \quad (11)$$

This approach provides an efficient latency estimation, grounded in the fundamental mechanics of LLM inference.

Context-Aware Latent Space Coordinate Prediction

To support effective routing generalization to out-of-distribution (OOD) queries, we introduce a lightweight unified predictor that bridges a query’s intrinsic textual properties with its multi-dimensional IRT parameter vectors: the discrimination vector $\alpha \in \mathbb{R}^D$ and the difficulty vector $\mathbf{b} \in \mathbb{R}^D$.

Hybrid Feature Representation

Our approach posits that a query’s latent parameters arise from both its semantic depth and structural complexity. We encode a given query q into following two feature vectors: **Semantic Embedding** (e_{se}). We capture contextual meaning using a fine-tuning DistilBERT model to extract the final hidden state of the [CLS] token.

$$e_{\text{se}} = \text{DistilBERT}(q)_{[\text{CLS}]} \in \mathbb{R}^{d_{\text{sem}}} \quad (12)$$

Structural Features (e_{st}). We quantify surface-level complexity using a feature extraction function $\Phi(\cdot)$, which computes $k = 11$ linguistic metrics (e.g., readability scores, parse tree depth). The selection of these hybrid features is guided by a correlation analysis with the target IRT parameters, ensuring they provide relevant predictive signals.

$$e_{\text{st}} = \Phi(q) \in \mathbb{R}^k \quad (13)$$

Unified Prediction Model

To co-predict both α and \mathbf{b} from these features, we designed a multi-task learning model. This model consists of a shared backbone that learns a high-level representation of the query, followed by two independent prediction heads.

Shared Backbone. The backbone first projects each feature vector into a higher-dimensional space using dedicated linear layers, then fuses them through a deep non-linear trunk network f_{fuse} for robust feature learning. This produces the shared latent representation:

$$\mathbf{h}_{\text{shared}} = f_{\text{fuse}}([\mathbf{W}_{\text{se}}e_{\text{se}} + e_{\text{se}}; \mathbf{W}_{\text{st}}e_{\text{st}} + \mathbf{b}_{\text{st}}]), \quad (14)$$

where $(\mathbf{W}_k, \mathbf{b}_k)$ are trainable parameters for $k \in \{\text{se}, \text{st}\}$.

Difficulty Head (Residual Prediction). Predicting the absolute position of the difficulty vector \mathbf{b} across a high-dimensional space is challenging. However, we observe that a query’s difficulty vector often represents a small deviation from a data-driven mean $\bar{\mathbf{b}}$ computed across the training sets. We therefore reformulate the task to predict only this *residual* vector, which significantly simplifies the learning problem. The final predicted difficulty vector $\hat{\mathbf{b}}_q$ is constructed as:

$$\hat{\mathbf{b}}_q = \bar{\mathbf{b}} + \Delta \mathbf{b}_q = \bar{\mathbf{b}} + f_{\text{diff}}(\mathbf{h}_{\text{shared}}), \quad (15)$$

where f_{diff} is the difficulty head, implemented as a multi-layer perceptron (MLP), that takes the shared latent representation $\mathbf{h}_{\text{shared}}$ as input to predict the residual vector $\Delta \mathbf{b}_q$.

Discrimination Head (Multi-Head Experts). The dimensions of the discrimination vector α are often correspond to related task types or specific skills. A monolithic prediction head struggles to capture these nuanced inter-dimensional relationships. We therefore partition the D dimensions into C distinct clusters based on inter-dimensional correlation analysis. For each cluster, a dedicated expert head specializes in predicting only the subset of dimensions within its assigned group. The final discrimination vector $\hat{\alpha}_q$ is then constructed by concatenating the outputs from each expert:

$$\hat{\alpha}_q = \bigoplus_{c=1}^C \hat{\alpha}_{q,c} = \bigoplus_{c=1}^C f_{\text{disc},c}(\mathbf{h}_{\text{shared}}), \quad (16)$$

where \bigoplus denotes the concatenation and re-ordering of the vector outputs from each expert head, and each expert head $f_{\text{disc},c}$ is a dedicated multi-layer perceptron (MLP) mapping the shared representation $\mathbf{h}_{\text{shared}}$ to its corresponding discrimination sub-vector $\hat{\alpha}_{q,c}$.

By estimating these latent vectors, we can project how one model will perform on an unseen query, which is the key to our zero-shot routing capability.

LLM Routing Strategies

Effective LLM routing requires navigating a complex, multi-dimensional trade-off between accuracy, cost, and latency. To address this, ZeroRouter introduces a unified optimization framework that allows developers to flexibly express their priorities within a coherent system.

Multi-Objective Optimization. We formulate the routing task as a Integer Linear Programming (ILP) problem that maximizes a weighted utility function, unifying the competing goals of accuracy, cost, and latency. For a set of queries \mathcal{Q} and a pool of models \mathcal{M} , we define the predicted accuracy p_{uq} , estimated cost \mathcal{C}_{uq} , and estimated latency τ_{uq} for each query-model pair (u, q) . Our goal is to determine the optimal assignment, represented by a binary variable $x_{uq} \in \{0, 1\}$, guided by three user-specified weights (w_p, w_c, w_t) that sum to 1. The objective is formulated as:

$$\text{maximize } \sum_{u \in \mathcal{M}} \sum_{q \in \mathcal{Q}} (w_p p_{uq} - w_c \mathcal{C}_{uq} - w_t \tau_{uq}) \cdot x_{uq}, \quad (17)$$

subject to the fundamental constraint that each query is assigned to exactly one model, and a set of optional global constraints on total cost, latency and average accuracy:

$$\begin{aligned} \text{subject to } & \sum_{u \in \mathcal{M}} x_{uq} = 1, & \forall q \in \mathcal{Q} \\ & \sum_{u \in \mathcal{M}} \sum_{q \in \mathcal{Q}} r_{uq} \cdot x_{uq} \leq R_{\text{max}}, & \text{for } r \in \{\mathcal{C}, \tau\} \\ & \sum_{u \in \mathcal{M}} \sum_{q \in \mathcal{Q}} p_{uq} \cdot x_{uq} \geq |\mathcal{Q}| \cdot p_{\text{min}}. \end{aligned} \quad (18)$$

Policy-Driven Routing. This unified formulation is highly expressive. By adjusting the weights (w_p, w_c, w_t) , developers can define a wide spectrum of routing behaviors without any change to the underlying system architecture: ① **Accuracy-First** ($w_p \rightarrow 1$): For critical applications, maximizing accuracy is paramount. ② **Cost-First** ($w_c \rightarrow 1$): For budget-constrained services, the router prioritizes the cheapest models. ③ **Latency-First** ($w_t \rightarrow 1$): For interactive applications, minimizing response time is crucial. ④ **Balanced Strategy**: Our evaluation adopts a mix of weights (e.g., $w_p = 0.5, w_c = 0.3, w_t = 0.2$) that find a balance case.

Experiments

Setup and Implementation Details

Datasets. We evaluate ZeroRouter with nine datasets, divided into six in-distribution (ID) sets to measure performance on familiar tasks, and three out-of-distribution (OOD) sets not seen in training to test its zero-shot generalization.

Six ID benchmarks: (1) **IFEval** (Zhou et al. 2023) for evaluating instruction following fidelity, (2) **BBH** (Big-Bench Hard) (Suzgun et al. 2022) for tasks requiring multi-step reasoning, (3) **MATH** (Hendrycks et al. 2021) for advanced mathematical problem solving, (4) **GPQA** (Graduate-Level Google-Proof Q&A) (Rein et al. 2024) for difficult questions, (5) **MuSR** (Multi-Step Reasoning) (Yao et al. 2025) for its focus on sequential reasoning, and (6) **MMLU-PRO** (Wang et al. 2024), a programmatically-verified version of MMLU. **Three OOD benchmarks:** (1) **ARC-C** (AI2 Reasoning Challenge) (Clark et al. 2018), a collection of science questions, (2) **TruthfulQA** (Lin, Hilton, and Evans 2021), which measures a model’s propensity to avoid generating falsehoods, and (3) **HumanEval** (Li and Murr 2024) for evaluating code generation capabilities.

LLM Candidates. We construct evaluations on a curated pool of 60 Large Language Models (LLMs) strategically divided into two sets. A core set of 10 models, with parameter counts spanning from a compact 1B to a massive 235B, was used to test routing performance across a diverse and challenging landscape: Qwen3-235B-A22B, DeepSeek-R1-Distill-Llama-70B, Mixtral-8x7B-Instruct-v0.1 (Jiang et al. 2024), Qwen3-32B, gemma-3-27b-it, phi-4, Llama-3.1-8B-Instruct, DeepSeek-R1-0528-Qwen3-8B, gemma-3n-E4B-it, and gemma-3-1b-it. A second cohort of 50 distinct LLMs being released after our router’s training data cutoff was used to directly validate our central claim of breaking model lock-in.

Baselines. We compare our framework against four recent representative works: (1) **CIT-LLM-Routing** (Zhang, Zhan, and Ye 2025), which introduces a novel paradigm of capability instruction tuning based on model aptitude tests to predict performance. (2) **RouteLLM** (Ong et al. 2024), which adopts a binary routing strategy that uses human preference data to assign queries to either a “strong” or “weak” model. (3) **GraphRouter** (Feng, Shen, and You 2024), a method that constructs a heterogeneous graph to explicitly model interactions among tasks, queries, and LLMs, reframing routing as an edge prediction problem. (4) **FORC** (Šakota, Peyrard, and West 2024), which employs a two-stage routing method that first uses a pre-trained meta-model to predict query accuracy

for each LLM, and then selects the most suitable model based on these predictions.

Evaluation Metrics. To comprehensively evaluate the trade-offs of each router, we define a unified total reward over the test query set $\mathcal{Q}_{\text{test}}$ as:

$$\text{Reward} = \sum_{q \in \mathcal{Q}_{\text{test}}} (w_p p_{u^*q} - w_c C_{u^*q} - w_t \tau_{u^*q}), \quad (19)$$

where u^* is the model selected by the router for query q . The components p_{uq} , C_{uq} , and τ_{uq} are the accuracy, cost, and latency previously defined for a query-model pair. Specifically, accuracy is evaluated against ground truth (see the Appendix for more details), cost is calculated as defined in Eq. 6, and latency is end-to-end inference time.

Our evaluation considers three policy-driven scenarios with respective weights (w_p, w_c, w_t) : Accuracy-First (**Max-Acc**) (0.8, 0.1, 0.1), Cost-First (**Min-Cost**) (0.1, 0.8, 0.1), and Latency-First (**Min-Lat**) (0.1, 0.1, 0.8).

Implementation Details. We implement our multidim 2PL IRT model using the `py-irt` library (Natesan et al. 2016) on evaluation data for our ID datasets, which was aggregated from 200 models on the Open LLM Leaderboard (Beeching et al. 2023). We set the latent dimension $D = 20$ and train for 6,000 epochs using the Adam optimizer with an initial learning rate of 0.1, which is decayed exponentially by a factor of 0.99 per 100 epochs. The latent parameter predictor, a `DistilBERT-base-uncased` model (66M), is fine-tuned for 40 epochs with a batch size of 32 and a constant learning rate of $3e-5$. The entire training process is conducted on a single NVIDIA A800 GPU (80GB).

Performance on In-Distribution Data Our evaluation on in-distribution data (Table 1) reveals that ZeroRouter fundamentally redefines the cost-performance frontier. Across both smaller (<20B) and larger (25B-250B) models, our framework consistently outperforms all baselines not just on one objective, but on all simultaneously. For instance, with smaller models, ZeroRouter achieves a higher maximum accuracy score 0.45 than the strongest baseline while also securing the best scores for minimum cost -0.32 and latency -0.27. This pattern of dominance extends to large-scale models. This is more than a simple trade-off. ZeroRouter consistently achieves higher accuracy at a lower cost and latency, showing the adaptability and effectiveness of our routing solution across diverse scenarios.

Zero-Shot Generalization on OOD Datasets The framework’s true power is revealed in its zero-shot generalization to OOD datasets, a critical test of its ability to perform on entirely unseen tasks (Table 1). In this challenging scenarios, ZeroRouter’s performance becomes more pronounced. For both smaller and larger models, it not only achieves the highest Max-Acc (e.g., 0.68 vs the strongest baseline’s 0.62 for large models) but also dramatically outperforms all rivals on Min-Cost -0.17 and Min-Lat -0.25. This consistent success demonstrates that our routing logic avoids the trap of learning superficial, task-specific heuristics. Instead, by modeling the intrinsic properties of a query, ZeroRouter makes reliably informed decisions in new domains, showing the effectiveness and robust generalization of our approach.

Method	#Params	In-Domain			Out-of-Domain			Mean
		Max-Acc	Min-Cost	Min-Lat	Max-Acc	Min-Cost	Min-Lat	
<i>Smaller-scale LLMs (<20B)</i>								
Gemma-3	1B	0.15	-0.38	-0.30	0.44	-0.28	-0.24	-0.10
Gemma-3n-E4B	7.8B	0.25	-0.36	-0.41	0.47	-0.25	-0.41	-0.12
DeepSeek-R1-0528-Qwen3	8B	0.33	-0.58	-0.71	0.58	-0.56	-0.71	-0.28
Llama-3.1 Instruct	8B	0.29	-0.54	-0.65	0.59	-0.45	-0.66	-0.24
Phi-4	14.7B	0.41	-0.82	-0.71	0.55	-0.83	-0.78	-0.36
LLM Selection on Smaller-scale LLMs (5 models)								
Random Selection	—	0.29 [†]	-0.54 [†]	-0.55 [†]	0.52 [†]	-0.51 [†]	-0.55 [†]	-0.22
RouteLLM	—	0.30 [†]	-0.47 [†]	-0.53 [†]	0.51 [†]	-0.54 [†]	-0.48 [†]	-0.20
FORC	—	0.35 [†]	-0.36 [†]	-0.54 [†]	0.50 [†]	-0.57 [†]	-0.54 [†]	-0.19
GraphRouter	—	0.38 [†]	-0.40 [†]	-0.56 [†]	0.53 [†]	-0.28 [†]	-0.50 [†]	-0.14
Model-SAT	—	0.43	-0.61 [†]	-0.64 [†]	0.57	-0.59 [†]	-0.64 [†]	-0.25
ZEROROUTER (Ours)	—	0.45	-0.32	-0.27	0.59	-0.24	-0.23	0.00
<i>Larger-Scale LLMs (25B~250B)</i>								
Gemma-3	27B	0.36	-0.31	-0.37	0.43	-0.54	-0.40	-0.14
Qwen3	32B	0.49	-0.33	-0.30	0.63	-0.32	-0.29	-0.02
Mixtral-8x7B Instruct	56B	0.32	-0.30	-0.36	0.67	-0.42	-0.38	-0.08
DeepSeek-R1-Distill-Llama	70B	0.38	-0.55	-0.35	0.58	-0.68	-0.52	-0.19
Qwen3-235B-A22B Instruct-v0.1	235B	0.47	-0.81	-0.83	0.64	-0.75	-0.59	-0.31
LLM Selection on Large-scale LLMs (5 models)								
Random Selection	—	0.39 [†]	-0.44 [†]	-0.45 [†]	0.59 [†]	-0.46 [†]	-0.44 [†]	-0.14
RouteLLM	—	0.46 [†]	-0.45 [†]	-0.58 [†]	0.56 [†]	-0.47 [†]	-0.42 [†]	-0.15
FORC	—	0.49 [†]	-0.38 [†]	-0.55 [†]	0.58 [†]	-0.51 [†]	-0.43 [†]	-0.13
GraphRouter	—	0.50	-0.27 [†]	-0.62 [†]	0.61 [†]	-0.37 [†]	-0.27 [†]	-0.07
Model-SAT	—	0.49	-0.61 [†]	-0.57 [†]	0.62 [†]	-0.58 [†]	-0.46 [†]	-0.19
ZEROROUTER (Ours)	—	0.52	-0.22	-0.28	0.68	-0.17	-0.25	0.05

Table 1: Comparison of routing performance on In-Distribution (ID) and Out-of-Distribution (OOD) datasets. ZeroRouter consistently outperforms baseline methods across both smaller and larger model scales, and across all routing policies. [†] indicates the improvement of ZeroRouter over the baseline is statistically significant ($p < 0.05$) via a Wilcoxon signed-rank test.

Seamless Integration of New Models

Efficient Onboarding with Minimal Data To validate our framework’s ability to integrate a new model with minimal overhead, we profiled it using a scant budget of just 200 queries (1% of the training-visible data). The results (Table 2) show ZeroRouter decisively outperforms all baselines across every policy. With its core D-optimality sampling, it achieves a superior Max-Acc score (0.39 vs. the baseline best of 0.38 from Model-SAT) and secures the top scores for both Min-Cost (-0.34 from GraphRouter) and Min-Lat (-0.37 from GraphRouter). An internal ablation study further underscores the power of our information-theoretic approach, revealing that D-optimality dramatically outperforms simpler strategies like random sampling, boosting Max-Acc from 0.27 to 0.39.

Robustness in Real-world Evolving Model Pool We further validated our framework in a real-world simulation of rapid model iteration, maintaining a fixed-size pool (N=6) where newly released models sequentially replaced underperforming ones. As depicted in Figure 3(a), the system intelligently leveraged these more capable models, driving a gradual performance increase under the Max-Accuracy policy while successfully upholding cost and latency con-

straints when those objectives were prioritized. Importantly, our router was trained on data that predates the release of every model in this evolving pool.

Ablation Studies and Analysis

Analysis of the Latent Space To understand our framework’s effectiveness, we analyze the distinct roles of the Difficulty (b) and Discrimination (α) vectors, shown in Figure 3. This analysis reveals how the latent space separates task-agnostic complexity from task-specific relevance.

The Nature of Difficulty (b). The heatmap in Figure 3(b) offers an interesting insight into the nature of our latent space. Evidenced by uniform horizontal bands, the Difficulty vector b , remains remarkably consistent across all semantic task clusters for any given dimension. This suggests that b does not represent task-specific difficulty, but rather captures a universal, task-agnostic challenge inherent to each latent dimension. For example, the consistently high value along dimension b_{19} suggests it corresponds to a fundamentally demanding latent capability, such as complex reasoning, which is a prerequisite for success across a wide variety of tasks.

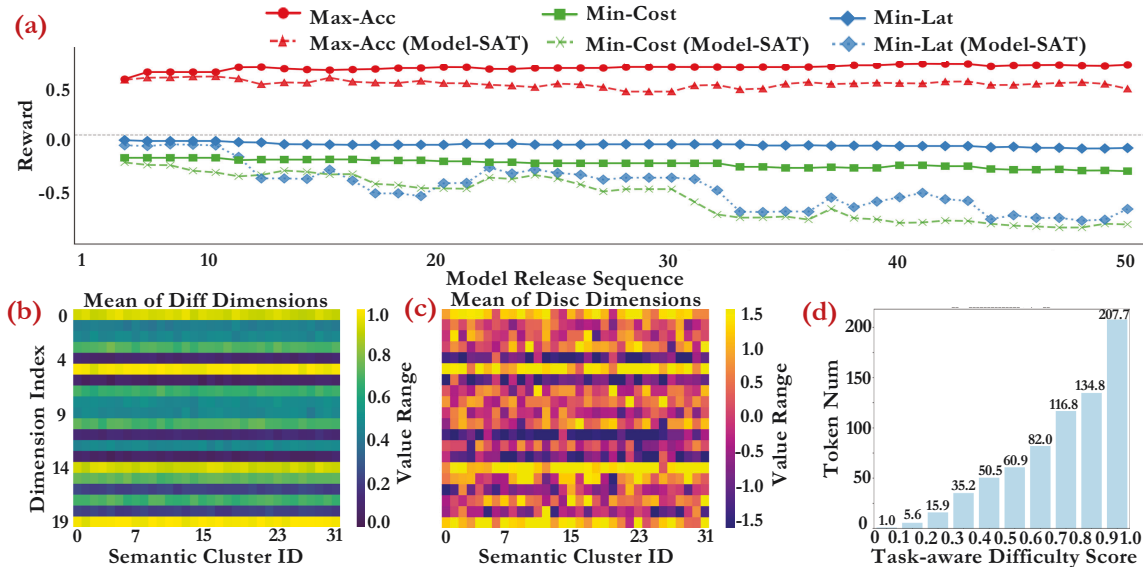


Figure 3: Experimental analysis of ZeroRouter’s performance and the interpretability of its latent space. (a) Real-world simulation demonstrating ZeroRouter’s ability to seamlessly integrate newly released models and improve performance over time without retraining. (b) Heatmap of the difficulty vector \mathbf{b} , showing its task-agnostic nature across different semantic clusters. (c) Heatmap of the discrimination vector α , highlighting its task-specific variability. (d) Validation of our proposed task-aware difficulty, showing a strong monotonic correlation with the average model output token length.

Method	Max-Acc	Min-Cost	Min-Lat	Mean
RouteLLM	0.28 [†]	-0.61 [†]	-0.62 [†]	-0.32 [†]
FORC	0.28 [†]	-0.54 [†]	-0.57 [†]	-0.28 [†]
GraphRouter	0.33 [†]	-0.42 [†]	-0.44 [†]	-0.18 [†]
Model-SAT	0.38	-0.52 [†]	-0.45 [†]	-0.20 [†]
ZEROROUTER				
+ Random Sampling	0.27 [†]	-0.56 [†]	-0.47 [†]	-0.25 [†]
+ Diff-based	0.24 [†]	-0.68 [†]	-0.52 [†]	-0.32 [†]
+ Disc-based	0.28 [†]	-0.58 [†]	-0.56 [†]	-0.29 [†]
+ Task-aware difficulty	0.32 [†]	-0.46 [†]	-0.44 [†]	-0.19 [†]
+ D-optimality	0.39	-0.34	-0.37	-0.11

Table 2: Ablation study on anchor sampling strategies for efficiently onboarding new models. We compare our proposed D-optimality against alternative sampling strategies and representative baselines. [†] indicates that the improvement of D-optimality over the corresponding method is statistically significant ($p < 0.05$) via a Wilcoxon signed-rank test.

The Role of Task-Specific Discrimination (α). In stark contrast, the Discrimination vector α , in Figure 3(c) is highly variable across different task clusters. This indicates that α captures task-specific properties, quantifying how discriminative a given query is for each latent ability. For instance, dimension α_{19} shows high values primarily for mathematical and logical reasoning tasks, confirming this dimension is indeed associated with that ability. Furthermore, we observe that certain groups of dimensions co-vary across tasks (e.g., dimensions 4, 6, 11, 13 show similar patterns), suggesting they form “ability clusters” representing related skills.”

Validation of Task-Aware Difficulty (s_q) The validity of our task-aware difficulty metric, derived from the synergy of \mathbf{b} and α , is confirmed by its strong monotonic correlation with the average output token length (Figure 3(d)). This intuitive result—as complex queries require elaborate responses—confirms our metric successfully captures task difficulty.

Ablation Study on Anchor Sampling We analyze the effectiveness of different anchor sampling strategies, with results presented in Table 2. We observe that the Task-aware difficulty strategy achieves scores (0.32, -0.46, -0.44) comparable to the GraphRouter baseline. This validates our hypothesis that a comprehensive score combining both general difficulty (\mathbf{b}) and task-specific discrimination (α) is effective for model profiling. Moreover, the D-optimality strategy significantly outperforms all other variants, achieving the best scores (0.39, -0.34, -0.37). By explicitly maximizing Fisher information, our method selects the most diagnostically valuable samples, leading to the most efficient characterization of a new model’s capabilities. Conversely, single-parameter strategies like Diff-based and Disc-based perform poorly, on par with Random Sampling (0.27, -0.56, -0.47).

Conclusion

We introduced ZeroRouter, a framework that successfully breaks model lock-in in LLM routing. By decoupling query properties from model capabilities via a universal latent space, our approach enables scalable, cost-efficient, and zero-shot integration of new models, making a new paradigm for adaptive LLM orchestration.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (No. 62332016).

References

- Aggarwal, P.; Madaan, A.; Anand, A.; Potharaju, S. P.; Mishra, S.; Zhou, P.; Gupta, A.; Rajagopal, D.; Kappaganthu, K.; Yang, Y.; et al. 2024. Automix: Automatically mixing language models. *Advances in Neural Information Processing Systems*, 37: 131000–131034.
- Beeching, E.; Fourrier, C.; Habib, N.; Han, S.; Lambert, N.; Rajani, N.; Sanseviero, O.; Tunstall, L.; and Wolf, T. 2023. Open LLM Leaderboard. https://huggingface.co/spaces/HuggingFaceH4/open_llm_leaderboard.
- Cao, J.; Chan, Y.-K.; Ling, Z.; Wang, W.; Li, S.; Liu, M.; Qiao, R.; Han, Y.; Wang, C.; Yu, B.; et al. 2025. How Should We Build A Benchmark? Revisiting 274 Code-Related Benchmarks For LLMs. *arXiv preprint arXiv:2501.10711*.
- Clark, P.; Cowhey, I.; Etzioni, O.; Khot, T.; Sabharwal, A.; Schoenick, C.; and Tafford, O. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457*.
- Ding, D.; Mallick, A.; Wang, C.; Sim, R.; Mukherjee, S.; Ruhle, V.; Lakshmanan, L. V.; and Awadallah, A. H. 2024. Hybrid LLM: Cost-efficient and quality-aware query routing. *arXiv preprint arXiv:2404.14618*.
- Feng, T.; Shen, Y.; and You, J. 2024. Graphrouter: A graph-based router for llm selections. *arXiv preprint arXiv:2410.03834*.
- Hendrycks, D.; Burns, C.; Kadavath, S.; Arora, A.; Basart, S.; Tang, E.; Song, D.; and Steinhardt, J. 2021. Measuring mathematical problem solving with the math dataset. *arXiv preprint arXiv:2103.03874*.
- Jiang, A. Q.; Sablayrolles, A.; Roux, A.; Mensch, A.; Savary, B.; Bamford, C.; Chaplot, D. S.; Casas, D. d. l.; Hanna, E. B.; Bressand, F.; et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Kipnis, A.; Voudouris, K.; Buschhoff, L. M. S.; and Schulz, E. 2024. metabench—A Sparse Benchmark to Measure General Ability in Large Language Models. *arXiv preprint arXiv:2407.12844*.
- Li, D.; and Murr, L. 2024. HumanEval on Latest GPT Models—2024. *arXiv preprint arXiv:2402.14852*.
- Lin, S.; Hilton, J.; and Evans, O. 2021. Truthfulqa: Measuring how models mimic human falsehoods. *arXiv preprint arXiv:2109.07958*.
- Mohammadshahi, A.; Shaikh, A. R.; and Yazdani, M. 2024. Routoo: Learning to Route to Large Language Models Effectively. *arXiv preprint arXiv:2401.13979*.
- Natesan, P.; Nandakumar, R.; Minka, T.; and Rubright, J. D. 2016. Bayesian prior choice in IRT estimation using MCMC and variational Bayes. *Frontiers in psychology*, 7: 1422.
- Ni, J.; Xue, F.; Yue, X.; Deng, Y.; Shah, M.; Jain, K.; Neubig, G.; and You, Y. 2024. MixEval: Deriving Wisdom of the Crowd from LLM Benchmark Mixtures. *arXiv preprint arXiv:2406.06565*.
- Ong, I.; Almahairi, A.; Wu, V.; Chiang, W.-L.; Wu, T.; Gonzalez, J. E.; Kadous, M. W.; and Stoica, I. 2024. RouteLLM: Learning to Route LLMs from Preference Data. In *The Thirteenth International Conference on Learning Representations*.
- Polo, F. M.; Weber, L.; Choshen, L.; Sun, Y.; Xu, G.; and Yurochkin, M. 2024. tinyBenchmarks: evaluating LLMs with fewer examples. *arXiv preprint arXiv:2402.14992*.
- Rein, D.; Hou, B. L.; Stickland, A. C.; Petty, J.; Pang, R. Y.; Dirani, J.; Michael, J.; and Bowman, S. R. 2024. Gpqa: A graduate-level google-proof q&a benchmark. In *First Conference on Language Modeling*.
- Rodriguez, P.; Barrow, J.; Hoyle, A. M.; Lalor, J. P.; Jia, R.; and Boyd-Graber, J. 2021. Evaluation examples are not equally informative: How should that change NLP leaderboards? In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, 4486–4503.
- Šakota, M.; Peyrard, M.; and West, R. 2024. Fly-swat or cannon? cost-effective language model choice via meta-modeling. In *Proceedings of the 17th ACM International Conference on Web Search and Data Mining*, 606–615.
- Suzgun, M.; Scales, N.; Schärli, N.; Gehrmann, S.; Tay, Y.; Chung, H. W.; Chowdhery, A.; Le, Q. V.; Chi, E. H.; Zhou, D.; et al. 2022. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv preprint arXiv:2210.09261*.
- Wang, X.; Liu, Y.; Cheng, W.; Zhao, X.; Chen, Z.; Yu, W.; Fu, Y.; and Chen, H. 2025. Mixllm: Dynamic routing in mixed large language models. *arXiv preprint arXiv:2502.18482*.
- Wang, Y.; Ma, X.; Zhang, G.; Ni, Y.; Chandra, A.; Guo, S.; Ren, W.; Arulraj, A.; He, X.; Jiang, Z.; et al. 2024. Mmlupro: A more robust and challenging multi-task language understanding benchmark. *Advances in Neural Information Processing Systems*, 37: 95266–95290.
- Yao, H.; Huang, J.; Qiu, Y.; Chen, M. K.; Liu, W.; Zhang, W.; Zeng, W.; Zhang, X.; Zhang, J.; Song, Y.; et al. 2025. MMReason: An Open-Ended Multi-Modal Multi-Step Reasoning Benchmark for MLLMs Toward AGI. *arXiv preprint arXiv:2506.23563*.
- Zhang, Y.-K.; Zhan, D.-C.; and Ye, H.-J. 2025. Capability instruction tuning: A new paradigm for dynamic llm routing. *arXiv preprint arXiv:2502.17282*.
- Zhou, J.; Lu, T.; Mishra, S.; Brahma, S.; Basu, S.; Luan, Y.; Zhou, D.; and Hou, L. 2023. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*.