

Ban&Pick: Enhancing Performance and Efficiency of MoE-LLMs via Smarter Routing

Yuanteng Chen^{1,2,3}, Peisong Wang^{1,2}*, Yuantian Shao^{1,5}, Nanxin Zeng²,
Chang Xu⁴, Jian Cheng^{1,2,3}*

¹Institute of Automation, Chinese Academy of Sciences,

²School of Artificial Intelligence, University of Chinese Academy of Sciences,

³Zhongguancun Academy, ⁴School of Computer Science, University of Sydney,

⁵Nanjing University of Science and Technology

Correspondence: {peisong.wang, jcheng}@nlpr.ia.ac.cn

Abstract

Mixture-of-Experts (MoE) has become a key architecture for scaling large language models. Recent fine-grained MoE designs introduce hundreds of experts per layer, with multiple experts activated per token, enabling stronger specialization. However, during pre-training, routers are optimized mainly for stability and robustness: they converge prematurely and enforce balanced usage, limiting the full potential of model performance and efficiency at inference. In this work, we uncover two overlooked issues: (i) a few highly influential experts are underutilized due to premature and balanced routing decisions; and (ii) enforcing a fixed number of active experts per token introduces substantial redundancy. Instead of retraining or redesigning MoE architectures, we introduce **Ban&Pick**, a training-free, plug-and-play strategy for smarter routing. **Pick** discovers and reinforces key experts, a small group with outsized impact on performance, leading to notable accuracy gains. **Ban** further dynamically prunes redundant experts based on layer and token sensitivity, delivering faster inference. Experiments on fine-grained MoE-LLMs (DeepSeek, Qwen3) across math, code, and general reasoning benchmarks demonstrate that Ban&Pick delivers free performance gains and inference acceleration. For instance, on Qwen3-30B-A3B, it improves accuracy from 80.67 to 84.66 on AIME2024 and from 65.66 to 68.18 on GPQA-Diamond, while accelerating inference by 1.25× under the vLLM.

1 Introduction

Large language models (LLMs) have achieved remarkable progress (Zhou et al., 2024), yet the ever-growing demand for higher capacity presents significant challenges for efficiency. As a promising solution, Mixture-of-Experts (MoE) architectures scale parameter count without proportionally in-

creasing computation, by activating a small subset of experts for each input (Artetxe et al., 2022).

Recent advances have pushed MoE towards a fine-grained design (Dai et al., 2024), where each layer contains a large number of relatively small experts. For example, DeepSeek-v2.5 (DeepSeek-AI et al., 2024) employs 160 experts per layer, while Qwen3-MoE (Yang et al., 2025) adopts 128 experts per layer. Unlike coarse-grained MoE (Jiang et al., 2024), which typically consists of a few large experts derived from the same base model, fine-grained MoE initializes each expert independently and randomly. This design promotes specialization: experts gradually develop distinct skills, excelling in areas such as math, code, or general reasoning. Such strong differentiation suggests that fine-grained MoE has the potential to leverage expert diversity more effectively.

However, we argue that a critical but overlooked phenomenon exists: **the routing strategies learned during pre-training cannot fully exploit this specialization**. Specifically, we empirically discover that current routing strategies inadvertently limit the potential to leverage the most impactful experts while simultaneously introducing redundancy where many selected experts contribute minimally to the final output. Instead of attempting to redesign training objectives or expert architectures (related work is discussed in detail in Appendix A), we aim to unlock this untapped potential efficiently and effectively during inference.

Delving into this question, we surprisingly find that within the most frequently selected experts of fine-grained MoE models, a small fraction plays an outsized role: simply forcing these experts to be chosen for every input can noticeably boost accuracy on certain datasets. Motivated by this striking observation, we conduct an empirical analysis of expert usage across domains and discover that each domain consistently relies on its own set of frequently selected experts, which we refer to as

* Corresponding author.




Figure 1: The intuition and empirical findings on expert utilization. Left (**Pick**): forcibly activating key experts (e.g., E43) improves accuracy on Winogrande (Sakaguchi et al., 2021). Right (**Ban**): reducing experts from 6 to 5 cuts 16.7% computation without accuracy loss. And we provide additional results on more datasets in Appendix G.

domain-specialized experts. Yet, a closer examination further shows only a tight subset of these experts has a decisive impact on the final logits distribution, while the influence of the others remains marginal. We term these most critical ones key experts. Our **Pick** module is designed to reinforce their influence during routing, getting performance gains by emphasizing key experts.

In parallel, we aim to alleviate redundancy in expert activation. Under the constraint that a fixed number of experts must be selected for each token, many activated experts contribute little to the final output, reflecting the redundancy introduced by less relevant experts. Our **Ban** module jointly considers layer sensitivity and token sensitivity to assess the redundancy in expert selection, and applies dynamic pruning of experts during inference. In this way, Ban reduces unnecessary computation and achieves significant inference acceleration with only minor accuracy loss.

Attempting to bridge the gap between routing strategies learned during pre-training, which prioritizes stability and balance, and optimal inference-time expert utilization, Pick and Ban empower truly impactful experts to contribute more while reducing participation from low-contribution experts:

- **Pick** amplifies excellence by identifying experts that demonstrate superior performance on specific domains and strategically increasing their influence during inference.
- **Ban** eliminates inefficiency by dynamically excluding experts that contribute minimally at inference, reducing computational costs while keeping essential capabilities.

Together, we introduce Ban&Pick, a unified framework that optimizes expert utilization in inference to improve both the accuracy and efficiency of fine-grained MoE without retraining or architectural changes. Evaluated on five widely used datasets

spanning math, code, and general reasoning, the results demonstrate that on the Qwen3 series models, applying Pick alone brings an average 2.62% performance improvement. By further incorporating Ban, the joint Ban&Pick strategy achieves an average 1.28% improvement while simultaneously accelerating inference by 1.25× under vLLM.

2 Empirical Findings on Routing Limitations in Fine-grained MoE

2.1 Theoretical Insights from Pre-training

Recent evidence from OLMoE (Muennighoff et al., 2025) and OpenMoE (Xue et al., 2024) suggests that during pre-training, routers converge much earlier than the experts. In particular, OLMoE reports that after only 1% of pre-training, around 60% of the routing decisions have already saturated, while OpenMoE finds that by 40% of pre-training, about 80% of the routing has stabilized. Such early fixation implies that routing decisions are largely determined before the experts have sufficiently matured. This premature convergence, combined with the widely used auxiliary balancing loss, fundamentally limits optimal expert utilization.

The balancing loss, introduced to prevent training collapse onto a few dominant experts, is typically formulated as $\mathcal{L}_{\text{expert}} = \alpha \sum_i f_i \cdot p_i$, where f_i is the activation frequency and p_i is the average routing probability of expert i (Lepikhin et al., 2020). When an expert receives both high routing probability and high actual usage, the product $f_i p_i$ becomes large and is penalized. While effective at ensuring balanced loads, this mechanism inherently conflicts with expert specialization (Wang et al., 2024a)—it discourages tokens from concentrating on the most competent experts, which is also argued by prior work (Guo et al., 2025). Fine-grained MoE partially alleviates this tension by increasing the number of experts, allowing some

specialization to emerge. However, as we will show in Section 3.1, even strongly specialized experts rarely exceed $\sim 50\%$ activation frequency, indicating the balancing loss still constrains how dominant any single expert can become.

The interplay between premature routing and balancing constraints creates a two-stage mechanism that underlies key expert under-utilization. *In the early stage*, before routing prematurely converges, the router learns a relatively stable domain-to-expert mapping, but the balancing loss caps how high any single expert’s frequency can become. Once routing stabilizes, *the remainder of training* (which constitutes the majority) continues to update expert parameters: frequently selected experts receive domain-specific data and become increasingly specialized. Crucially, however, their activation frequencies remain frozen at the earlier balance-constrained levels. **As a result, key experts that have evolved to be highly impactful are still activated at frequencies determined before their specialization matured—creating a fundamental mismatch between their true competence and actual utilization.**

This implies that the current pre-training paradigm makes a trade-off: in exchange for training stability, it locks routers into decisions before experts have fully matured. This presents an opportunity from an inference perspective:

Can we recover this lost potential—amplifying the influence of the most impactful experts while removing redundancy, to boost performance and accelerate inference?

2.2 Expert Utilization Inefficiencies and Optimization Potential

To validate the inefficiencies in expert utilization and better understand the optimization potential, we conducted two simple yet revealing experiments with the Deepseek-V2-Lite-Chat (activates 6 experts out of 64) on Winogrande (more results are provided in Appendix G). Figure 1 demonstrates how we adjust expert utilization and its effects.

(1) Underutilization of the most impactful experts: we examined expert selection frequency on the C4 calibration set, identified the top-3 most frequently selected experts in each layer, and then conducted a grid search: each of these experts was manually added in addition to the original routing (E43 in Figure 1), and the resulting accuracy on Winogrande benchmark was recorded. The left part highlights a few experts contribute disproportion-

ately: enhancing Layer-1 Expert-43 improved accuracy by +1.81%, L10E33 by +1.65%, and L25E11 by +1.26%. When these three were jointly enhanced, accuracy increased by +2.72% (71.67% $\rightarrow$ 74.39%), while most other high-frequency experts had negligible or noisy effects, indicating that certain impactful experts are insufficiently leveraged in the original routing strategy. We refer to these experts as **key experts**.

(2) Computational waste from redundant experts: we directly reduced the number of selected experts per token from 6 to 5 and measured the accuracy. Remarkably, as shown in the right side of Figure 1, this pruning reduced computation by 16.7% and preserved full accuracy on Winogrande (71.67%), demonstrating redundancy in the original expert utilization. To rule out the possibility that this observation is dataset-specific, we provide consistent results across datasets in Appendix G.

Together, these experiments validate the inefficiencies in expert utilization predicted by pre-training analysis, and support the intuition of our methods—Pick and Ban, respectively, which aim to optimize expert selection beyond pre-training constraints for better inference-time performance. This empirical study motivates us to design a unified framework that can pick key experts while banning redundant ones, aiming at improving both accuracy and efficiency.

3 Pick

In this section, we begin with an analysis of expert specialization in fine-grained MoE models to verify the existence of domain-specialized experts, which forms the foundation for key experts to effectively function. Building on this, we then introduce our Pick method in two steps: (i) identifying key experts from the set of domain-specialized experts, and (ii) designing strategies to enhance key experts.

3.1 Expert Specialization in Fine-Grained MoE

First, we take a closer look at the degree of expert specialization in fine-grained MoE models. Specifically, we analyze the Qwen3-30B-A3B model and record expert usage frequencies on task-related benchmarks: math tasks (GSM8K + Math-500), code tasks (MBPP (Austin et al., 2021) + Livecodebench), and a general task (GPQA-Diamond). The statistics cover all tokens in both the prefill and decode stages. For each type of task, we plot




Figure 2: Expert specialization in fine-grained MoE (Qwen3-30B-A3B). Left: expert usage frequency across three tasks (math, general, code). Right: token word clouds for 3 high-frequency experts.

the per-layer expert usage frequency and show 3 representative layers in Figure 2.

Our analysis reveals that in each layer there exist experts whose usage frequency is significantly higher than average—often up to 6 times higher. Moreover, in most layers these high-frequency experts vary across tasks. For instance, L16E95 is frequently activated only on math tasks, whereas L29E91 is almost exclusively activated on code tasks with an extremely high activation rate of 66.9%. In contrast, a small number of experts such as L4E48 remain highly active across all tasks.

These task-dependent high-frequency experts suggest the existence of domain-specialized experts. To validate this hypothesis, we further examine the functional roles of these experts by analyzing which tokens are most often routed to them. Taking the three aforementioned experts as examples, we visualize the tokens they frequently process using word clouds. The results show that the math-specialized expert L16E95 is strongly associated with digits, operators, and math terminology; the code-specialized expert L29E91 is associated with programming tokens; and the general expert L4E48 is closely tied to branch tokens (Zhong et al., 2025), which play a critical role in chain-of-thought across diverse tasks.

These observations provide empirical evidence that expert specialization in fine-grained MoE models is highly pronounced. Different tasks rely heavily on different domain-specialized experts, while a small subset of general experts remain consistently important across domains.

3.2 Picking Key Experts from Domain-Specialized Experts

As shown previously, almost every layer contains domain-specialized experts that are activated with notably high frequency. The key issue, however, is whether all of these experts contribute equally to model’s performance, or only a small subset is truly essential. To investigate this, we evaluate the importance of each expert by measuring the shift in the logits distribution when it is removed.

Concretely, taking math tasks as an example, we use the domain-specialized experts identified in previous analysis as candidate experts in each layer. We then run the model on 1000 randomly sampled instances from MathQA (Amini et al., 2019) dataset. For each sample, we record the output logits of the original model, and then repeat the inference while pruning one candidate expert at a time. The perturbation of the output distribution is quantified using the KL divergence between the probability distributions of the Top-1000 tokens from the original and pruned models.

The results are summarized in Figure 3. Surprisingly, despite their high activation frequency, most domain-specialized experts cause only negligible shifts in the output distribution when pruned. In contrast, a small number of experts stand out, with three of them inducing particularly large changes. Among these, expert L9E18 is the most striking, leading to a divergence of about 0.16. These findings suggest that only a few domain-specialized experts truly have a decisive influence on the model’s behavior, and we hypothesize that they are the key experts we aim to identify.

To test this hypothesis, we construct a failure




Figure 3: Impact of pruning (measured by KL divergence, left axis) and enhancing (by accuracy gain, right axis) for domain-specialized experts.

set consisting of the Math-500 (Lightman et al., 2023) problems that the base model originally answers incorrectly (24 in total). For each candidate expert, we forcibly include it during routing and then re-evaluate accuracy on this failure set. As shown in Figure 3 (right vertical axis), only those experts that induce substantial distribution shifts also yield meaningful accuracy gains when enhanced, whereas others contribute little to no improvement. In particular, expert L9E18, which caused the largest divergence, also increases the number of correct answers by 5 out of 24 when included. This provides strong evidence of the close correspondence between distribution shifts and functional importance, and offers a principled way to distinguish key experts from the many domain-specialized experts. We provide more implementation details on the selection criterion, calibration set choices, and KL divergence computation, together with complete results in Appendix B.

3.3 Enhancing Key Experts



Figure 4: Comparison of four designed enhancement methods for key experts, evaluated by accuracy on three widely used math benchmarks.

After identifying the key experts, the remaining challenge is to determine how to leverage them most effectively. We take math tasks as a running example and design four enhancement strategies:

- Forced Addition:** Whenever a key expert is not selected, it is added alongside chosen experts.
- Forced Replacement:** Similar to forced addition, but the key expert replaces the selected expert with the lowest routing weight, maintaining the total number of selected experts.
- Range-based Addition:** If a key expert is not selected (not in the top- k candidates) but appears within the router’s top- $2k$ candidates, it is forcibly added.
- Range-based Replacement:** Under the same condition as range-based addition, the key expert replaces the selected expert with the lowest routing weight.

We evaluate these methods on AIME2024, AIME2025, and Math-500. For the two AIME datasets, results are averaged over five runs to mitigate small-sample variance. Figure 4 summarizes the outcomes. Overall, strategy A brings consistent but modest gains. Strategy B, which forces replacement regardless of probability, is unstable and even reduces accuracy on AIME2024 and Math-500, indicating that overly aggressive intervention can be harmful. By contrast, strategies C and D, which enhance only when the key expert is already close to the router’s top candidates, deliver the strongest improvements across all benchmarks. The difference between C and D is minor, but since D avoids expanding the number of active experts, it is preferable in practice and used in experiments.

4 Ban

Our dynamic expert pruning method is motivated by two key observations: 1) Different layers in an MoE model exhibit highly diverse sensitivity to expert pruning and thus should be treated differently. 2) During reasoning, the token-wise routing distributions also vary significantly: when most of the routing weight is concentrated on a few top experts, expert pruning causes minimal loss; in contrast, when the weights are more evenly distributed, expert pruning becomes more harmful.

In the following, we first use Qwen3-30B-A3B as an example model to illustrate the differences in layer sensitivity and token sensitivity separately, and then introduce our dynamic pruning method.

Layer Sensitivity. We measure layer-wise pruning sensitivity by reducing the number of selected experts from $k = 8$ to $k = 3$ in each layer, and




Figure 5: Sensitivity analysis of Qwen3-30B-A3B: (a) **layer-wise** pruning sensitivity measured by KL divergence across layers; (b) **token-wise** pruning sensitivity measured by the Top-3 / Top-8 routing weights ratio on two random samples. Both dimensions exhibit large variance, motivating a dynamic pruning strategy.

computing the Top-1000 KL divergence between the output logits of the original and pruned models on MathQA calibration set. The results (Figure 5a) reveal substantial differences across layers: while some layers, especially at the begin and end of the model, are highly sensitive to pruning (shown in orange), many middle layers are much more robust (shown in purple). This large variance highlights the necessity of treating different layers differently when designing pruning strategies.

Token Sensitivity. We further analyze token-level pruning sensitivity in long reasoning processes by examining how routing weights are distributed across experts. To quantify this, we compute the ratio between the cumulative routing weights of the top-3 and top-8 chosen experts, averaged across all layers. Figure 5b plots these ratios for the first 200 generated tokens in reasoning processes from two randomly selected GPQA-Diamond samples. The both results reveal substantial variability across tokens: some tokens concentrate most of their weights on a few experts and are thus robust to expert pruning, whereas others distribute weights evenly and are more sensitive.

Dynamic pruning. Before combining layer-level and token-level information, we normalize their sensitivity scores to the range $[0, 1]$. Here, W_l denotes the KL divergence of layer l when reducing expert selection from $k = 8$ to $k = 3$. $W_{\min}$ and $W_{\max}$ are defined as the divergences of the least-sensitive and most-sensitive layers, respectively. For tokens, R_i represents the top-3/top-8 routing weight ratio of token i at a specific layer

during inference, with $R_{\min}$ and $R_{\max}$ similarly estimated from the calibration set. The normalized sensitivity scores are then defined as:

$$L'_l = (W_l - W_{\min}) / (W_{\max} - W_{\min}),$$

$$T'_i = (R_{\max} - R_i) / (R_{\max} - R_{\min}).$$

The combined sensitivity score for token i at layer l is then:

$$S_{i,l} = \lambda \cdot \frac{1}{2} (L'_l + T'_i), \quad \lambda \in (0, 1)$$

where we simply average layer-level and token-level sensitivities to balance their contributions, and use λ to control the aggressiveness of pruning. The dynamic number of selected experts is determined as:

$$K_{i,l} = \lfloor K_{\min} + (K_{\text{base}} - K_{\min}) \cdot S_{i,l} \rfloor.$$

Here $K_{\text{base}} = 8$ is the default number of experts, and $K_{\min} = 3$ serves as a safeguard to prevent selecting too few experts: from empirical inspection, we found that reducing the number of selected experts below this threshold leads to a complete collapse of model. We provide detailed hyperparameter sensitivity analysis in Appendix D, demonstrating that $K_{\min}$ and λ are robust and easy to set in practice.

5 Experiments

In this section, we first introduce the experimental setup, including models, datasets, and implementation details. We then separately compare our proposed **Pick** and **Ban** approaches with existing methods, highlighting their respective strengths. Finally, we show how combining **Ban&Pick** achieves both improved performance and faster inference.

5.1 Experiment Setup

Models and Datasets We conduct experiments on four fine-grained MoE models from two different model families, covering a wide range of parameter scales: DeepSeek-V2-Lite-Chat (16B), DeepSeek-V2.5-1210 (236B), Qwen3-30B-A3B, and Qwen3-235B-A22B. Our evaluation spans three categories of tasks using five commonly adopted datasets: AIME2024 (with 5 runs) (Patel et al., 2024) and Math-500 for mathematical reasoning, GPQA-Diamond (Rein et al., 2023) and MMLU-Pro (Wang et al., 2024b) for general knowledge reasoning, and LiveCodeBench (Jain et al., 2024) together with HumanEval+ (Liu et al.,

Table 1: Comparison of **Pick** against three baselines across four models. Reported accuracies cover multiple benchmarks, with GSM8K (Cobbe et al., 2021) only for DeepSeek-V2-Lite-Chat. (more results in Appendix E)

Models	Methods	GSM8K/AIME24	Math500	GPQA-Diamond	MMLU-Pro	LiveCodeBench	HumanEval+
 Deepseek-V2-Lite-Chat	Baseline	66.34	26.40	24.75	27.44	9.25	50.61
	Dynamic	66.56	26.80	24.75	27.03	9.00	50.61
	Tips	66.87	26.60	25.25	28.15	9.50	51.22
	Pick	67.48	28.00	25.76	28.97	10.50	51.83
 Deepseek-v2.5-1210	Baseline	23.33	81.00	47.47	27.97	60.75	83.54
	Dynamic	24.00	81.20	46.97	28.86	61.00	83.54
	Tips	24.67	80.60	47.98	28.54	61.25	84.15
	Pick	27.34	82.80	49.49	30.06	62.75	84.76
 Qwen3-30B-A3B	Baseline	80.67	95.20	65.66	63.95	91.75	71.34
	Dynamic	82.66	94.80	66.16	63.41	92.00	71.95
	Rice	83.33	95.60	67.68	64.02	93.50	72.56
	Pick	85.33	96.60	69.19	65.34	94.25	75.61
 Qwen3-235B-A22B	Baseline	84.67	96.00	71.21	68.45	92.75	78.05
	Dynamic	85.33	96.00	71.72	69.02	93.00	78.66
	Rice	86.67	96.20	72.22	68.75	93.00	79.88
	Pick	88.00	96.80	74.24	70.18	94.50	81.10

2023) for code generation. For DeepSeek-V2-Lite-Chat, we replace AIME2024 with GSM8K, as AIME2024 is overly difficult for this model.

Implementation Details For the Pick method, we follow the procedures described in Section 3 to identify key experts for each of the three task categories—math, code and general reasoning—and apply targeted enhancements to them. We then evaluate the resulting models on their corresponding task types. For the Ban method, we set $K_{\min} = 3$ and $\lambda = 0.7$ for all four models, aiming to achieve a balanced trade-off between accuracy and speedup (further discussed in Appendix D). When combining Pick and Ban, we note that in layers containing key experts, Ban has already pruned redundant experts, so directly applying the range-based replacement (strategy D) leads to noticeable performance degradation. To address this, we instead adopt the range-based addition (strategy C).

5.2 Results of Pick

We compare our proposed **Pick** method against three representative baselines:

- (1) **Dynamic Routing.** (Huang et al., 2024a) adaptively allocates experts according to routing weights, assigning more experts to harder tokens and fewer to easier ones, realized by accumulating experts until a threshold is reached. We set thresholds $\{0.7, 0.8, 0.9\}$ and report the best result.
- (2) **Tip.** (Wang et al., 2025b) penalizes tokens that typically indicate switching paths, thereby encouraging the model to continue deepening its current line of reasoning.
- (3) **Rice.** (Wang et al., 2025a) leverages explicit `<think>` tokens to identify cognitive experts that specialize in reasoning. During inference, the routing weights of these experts are multiplied by a

constant factor to amplify their influence.

As shown in Table 1, **Pick** achieves consistent and notable improvements across four models on five datasets, outperforming all three baselines. The gains are especially pronounced on the Qwen3 series, which we attribute to their stronger expert specialization. Remarkably, on the large-scale Qwen3-235B-A22B, whose performance is already comparable to leading closed-source models, Pick further boosts accuracy by 3.33%, 3.03%, and 3.05% on AIME2024, GPQA, and HumanEval+, respectively. These results highlight the potential of leveraging key experts to further improve MoE model performance. Meanwhile, on DeepSeek-v2.5-1210, Pick achieves an average improvement of 2.19%, demonstrating strong generalization capability. Additionally, we report results on a broader range of task categories in Appendix E.

5.3 Results of Ban

We compare our proposed **Ban** method against two mainstream dynamic expert-pruning approaches:

- (1) **DES.** (Lu et al., 2024) (Dynamic Expert Skip) dynamically prunes low-weight experts by observing that, for easy tokens, the contribution of lower-ranked experts is often negligible. Concretely, it decides whether to drop the secondary expert based on the ratio between expert routing weights.
- (2) **ODP.** (Huang et al., 2024b) (Online Dynamic Pruning) extends DES by adding a key-token protection mechanism. Tokens that receive high attention scores are exempted from pruning.

In terms of accuracy, as shown in Table 2, both DES and ODP suffer from severe degradation on these challenging tasks, with drops exceeding 5% on certain datasets. In contrast, Ban limits the accuracy loss to within 1.5% on most models and

Table 2: Comparison of **Ban** against DES and ODP across four models. We report average experts per token (Avg-Topk), accuracy, and inference speedup under the vLLM (0.9.1) backend.

Models	Methods	Avg-Topk	GSM8K/AIME24	Math500	GPQA-Diamond	MMLU-Pro	LiveCodeBench	HumanEval+	Speedup
 Deepseek-V2-Lite-Chat	Baseline	6.00	66.34	26.40	24.75	27.44	9.25	50.61	1.00
	DES	4.18	63.76	21.40	22.22	24.33	7.25	45.73	1.20
	ODP	4.22	64.76	23.20	23.23	25.67	7.75	46.95	1.19
	Ban	4.02	66.11	25.00	24.24	26.51	8.50	48.17	1.22
 Deepseek-v2.5-1210	Baseline	6.00	23.33	81.00	47.47	27.97	60.75	83.54	1.00
	DES	4.21	22.66	80.20	39.90	23.89	58.00	79.27	1.21
	ODP	4.29	23.33	80.20	43.43	24.42	59.25	81.71	1.20
	Ban	3.97	23.33	80.60	45.96	26.85	60.25	82.93	1.24
 Qwen3-30B-A3B	Baseline	8.00	80.67	95.20	65.66	63.95	91.75	71.34	1.00
	DES	4.94	73.33	92.60	56.06	60.12	84.50	65.24	1.25
	ODP	4.98	74.67	93.20	60.10	62.31	86.00	67.07	1.24
	Ban	4.82	80.00	94.60	64.65	63.18	90.00	70.12	1.25
 Qwen3-235B-A22B	Baseline	8.00	84.67	96.00	71.21	68.45	92.75	78.05	1.00
	DES	4.89	68.00	94.80	62.12	63.87	90.75	70.12	1.26
	ODP	4.91	73.33	95.20	66.16	66.04	91.00	73.78	1.26
	Ban	4.77	83.33	95.40	69.70	67.32	91.25	75.61	1.27

Table 3: Results of combining **Ban&Pick** for a balanced combination of both improvements on performance and efficiency across four models. We report average experts per token, accuracy, and inference speedup under vLLM.

Models	Methods	Avg-Topk	GSM8K/AIME24	Math500	GPQA-Diamond	MMLU-Pro	LiveCodeBench	HumanEval+	Speedup
 Deepseek-V2-Lite-Chat	Baseline	6.00	66.34	26.40	24.75	27.44	9.25	50.61	1.00
	Ban&Pick	4.03	66.94	26.80	25.25	28.03	9.50	50.00	1.22
 Deepseek-v2.5-1210	Baseline	6.00	23.33	81.00	47.47	27.97	60.75	83.54	1.00
	Ban&Pick	3.99	26.00	81.60	47.98	29.13	61.25	84.15	1.23
 Qwen3-30B-A3B	Baseline	8.00	80.67	95.20	65.66	63.95	91.75	71.34	1.00
	Ban&Pick	4.83	84.66	95.80	68.18	64.51	92.85	72.56	1.25
 Qwen3-235B-A22B	Baseline	8.00	84.67	96.00	71.21	68.45	92.75	78.05	1.00
	Ban&Pick	4.79	86.67	96.20	72.73	68.98	93.25	78.66	1.26

datasets, except HumanEval+, which we found sensitive to pruning, and even within 1% on more than half of the results. Notably, on Qwen3-30B-A3B, Ban reduces accuracy on AIME2024 by only 0.67%—equivalent to misanswering just one additional problem out of five evaluations—and incurs merely 1.01% degradation on GPQA-Diamond.

In addition to accuracy, we record the average number of selected experts per token and the relative speedup—computed from the total inference time across all datasets—under the vLLM backend to comprehensively evaluate inference efficiency. **Ban** restricts the average number of selected experts to around 4 for DeepSeek models, leading to more than $1.2\times$ actual speedup. On Qwen3 models, Ban keeps the average expert count below 5, achieving about $1.25\times$ speedup. These results demonstrate the strong ability of Ban to preserve model performance while delivering substantial efficiency gains.

5.4 Results of Ban&Pick

By combining **Ban** and **Pick**, our unified framework leverages smarter routing strategies to simultaneously deliver free accuracy gains and inference acceleration across models and datasets, demonstrating strong practicality. As shown in Table 3, on Qwen3-30B-A3B, Ban&Pick achieves an av-

erage accuracy improvement of 1.67% together with a $1.25\times$ speedup, while on the larger-scale Qwen3-235B-A22B, it yields consistent accuracy gain and efficient speedup. As two plug-and-play modules, Ban and Pick share the common goal of utilizing smarter routing to release the potential of fine-grained MoE, while offering flexible choices: higher performance (**Pick**), faster inference (**Ban**), or a balanced combination of both improvements on performance and efficiency (**Ban&Pick**).

6 Conclusion

In this study, we reveal that in fine-grained MoE models, relatively balanced routing strategy underutilizes a small set of key experts which have a disproportionately large impact on task performance, and fixed expert activation introduces substantial redundancy. By selectively increasing the utilization of key experts and dynamically pruning redundant ones, we release the potential of routing and provide a novel and effective framework Ban&Pick that improves both accuracy and efficiency.

7 Related Work

Enhancing Expert Utilization in MoE. Most existing work attempts to improve expert utilization efficiency in MoE models by redesigning train-

Table 4: Enhanced key experts (formatted as (layer, expert)) for different domains across models.

Model	Math Experts	General Experts	Code Experts
Qwen3-30B-A3B	(2,37), (9,18), (21,18)	(5,19), (7,114), (13,105)	(6,40), (18,40)
Qwen3-235B-A22B	(11,121), (31,121), (51,94)	(13,57), (18,91), (62,101)	(7,56), (62,48), (82,48)
DeepSeek-V2-Lite-Chat	(5,14), (9,56), (17,25)	(4,36), (7,31), (14,22)	(8,1), (12,28)
DeepSeek-v2.5-1210	(6,58), (32,5)	(9,22), (16,66)	(1,129), (22,36), (35,13)

Table 5: Task accuracies (%) on Qwen3-30B-A3B under different combinations of enhanced experts. A ✓ indicates that the corresponding domain experts are enhanced, while × means not enhanced.

Math	General	Code	AIME2024	MATH-500	GPQA-Diamond	LiveCodeBench	HumanEval+
×	×	×	80.67	95.20	65.66	91.75	71.34
✓	×	×	85.33 ↑4.66	96.60 ↑1.40	65.15 ↓0.51	91.50 ↓0.25	71.95 ↑0.61
×	✓	×	82.00 ↑1.33	95.20 ↑0.00	69.19 ↑3.53	91.75 ↑0.00	72.56 ↑1.22
×	×	✓	81.33 ↑0.66	95.00 ↓0.20	65.66 ↑0.00	94.25 ↑2.50	75.61 ↑4.27
✓	✓	×	84.67 ↑4.00	96.20 ↑1.00	67.68 ↑2.02	92.00 ↑0.25	71.34 ↑0.00
✓	×	✓	84.67 ↑4.00	96.00 ↑0.80	66.16 ↑0.50	93.50 ↑1.75	74.39 ↑3.05
×	✓	✓	80.67 ↑0.00	94.80 ↓0.40	68.18 ↑2.52	92.75 ↑1.00	75.00 ↑3.66
✓	✓	✓	84.00 ↑3.33	96.00 ↑0.80	67.68 ↑2.02	93.00 ↑1.25	74.39 ↑3.05

ing objectives or expert architectures during pre-training (Yan et al., 2025; Huang et al., 2024a; Jin et al., 2024). However, while these approaches are insightful, they often sacrifice training stability for flexibility. More recently, (Wang et al., 2025a) proposes a post-training method specifically designed for MoE models that employ the <think> token, enhancing reasoning by amplifying the weights of Cognitive Experts.

Dynamic Expert Pruning of MoE. Dynamic expert pruning has emerged as a key technique for accelerating MoE inference. Lu et al. (2024) proposed a pruning strategy that simply skips experts with less weights: when the weight of the top-1 expert exceeds that of the top-2 expert by a given threshold, only the top-1 expert is activated. Building on this, Huang et al. (2024b) introduced a critical token protection mechanism, ensuring that important tokens are less affected by pruned. However, the substantial sensitivity differences across layers of MoE models and across tokens in long chain-of-thought reasoning are often overlooked, leaving room for further optimization.

8 Interaction Among Key Experts Across Tasks

In this section, we further investigate the scenario where all identified key experts across math, general reasoning, and code are enhanced simultaneously. We first summarize the enhanced experts for each task type in Table 4. Then, using Qwen3-30B-A3B as a case study, we report in Table 5

the performance when enhancing key experts for a single domain, for domain pairs (Math+General, Math+Code, General+Code), and for all three domains together. This allows us to examine whether experts from different domains complement or interfere with each other when activated jointly, providing a more comprehensive understanding of cross-domain interactions.

From Table 5, we observe that enhancing multiple domains simultaneously still brings clear improvements. For instance, when both Math and General experts are enhanced, AIME2024 increases by +4.0% and GPQA by +2.02%. However, these gains are slightly smaller than enhancing each domain individually (−0.66% and −1.51%, respectively). When all three domains are enhanced together, the model continues to outperform the baseline—+3.33% on AIME2024, +2.02% on GPQA, and +3.05% on HumanEval+—though the gains are somewhat reduced compared to the best single-domain setting. These results indicate that while domain-specific key experts do interfere to some extent, their combined enhancement remains clearly beneficial, confirming their crucial role.

Limitations

Towards more rigorous theoretical and pre-training-level analysis. In Section 2, we attribute the existence of under-utilized *key experts* to the interplay between two widely observed properties of MoE pre-training: routers tend to converge substantially earlier than experts, while the auxiliary balancing loss caps the activation frequency of

any single expert. We argue that this interplay causes a persistent mismatch between an expert’s evolving competence and its largely frozen activation frequency, and we provide extensive empirical evidence confirming that such under-utilized yet highly impactful experts indeed exist across current fine-grained MoE models. Nevertheless, our analysis remains an inference-time observation rather than a rigorous theoretical proof. Due to limited computational resources, we are unable to reproduce the full pre-training process under controlled variations of routing dynamics or balancing objectives, which would be required to causally isolate the contribution of each factor. We leave two directions to future work: (i) a more formal theoretical characterization of how premature router convergence and load-balancing constraints jointly shape expert specialization, and (ii) addressing this issue directly at the pre-training stage, so that the full potential of fine-grained MoE can be unlocked at the source rather than only recovered at inference.

Room for more sophisticated key-expert enhancement strategies. In Section 3.3, after identifying and verifying the existence of key experts, we design four simple and straightforward enhancement strategies and select the final one through comparative experiments. On the one hand, our results demonstrate that merely increasing the activation frequency of the identified key experts is already sufficient to yield consistent performance gains on the corresponding task domains, which clearly confirms the latent potential of key experts in fine-grained MoE models. On the other hand, we acknowledge that the strategies explored in this work are deliberately simple and likely fall short of fully exploiting key experts. More fine-grained designs, such as input-adaptive enhancement, dynamic adjustment of routing scores, or jointly optimizing the selection and weighting of key experts, are likely to bring further improvements. We leave the exploration of such more sophisticated enhancement strategies, built on top of the framework proposed in this paper, to future work.

Ethical Considerations

This work focuses on identifying and enhancing key experts in fine-grained Mixture-of-Experts (MoE) large language models, which is an inference-time analysis and adjustment technique and does not involve the collection of new data or human subjects. All models and benchmarks used

in our experiments are publicly available and employed in accordance with their respective licenses and intended research use. Since our method improves the task performance of fine-grained MoE models without requiring additional training or modification of model parameters, it may contribute positively by enabling more efficient utilization of existing pre-trained MoE models and reducing the need for costly retraining. We do not foresee any direct negative societal impact arising from this work beyond the general considerations applicable to large language models.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant No. 62572471).

References

- Aida Amini, Saadia Gabriel, Shanchuan Lin, Rik Koncel-Kedziorski, Yejin Choi, and Hannaneh Hajishirzi. 2019. [MathQA: Towards interpretable math word problem solving with operation-based formalisms](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 2357–2367, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mikel Artetxe, Shruti Bhosale, Naman Goyal, Todor Mihaylov, Myle Ott, Sam Shleifer, Xi Victoria Lin, Jingfei Du, Srinivasan Iyer, Ramakanth Pasunuru, Giridharan Anantharaman, Xian Li, Shuohui Chen, Halil Akin, Mandeep Baines, Louis Martin, Xing Zhou, Punit Singh Koura, Brian O’Horo, and 5 others. 2022. [Efficient large scale language modeling with mixtures of experts](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11699–11732, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, and 1 others. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Yuanteng Chen, Yuantian Shao, Peisong Wang, and Jian Cheng. 2025. [EAC-MoE: Expert-selection aware compressor for mixture-of-experts large language models](#). In *Proceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12942–12963, Vienna, Austria. Association for Computational Linguistics.
- Yuanteng Chen, Peisong Wang, Nanxin Zeng, Yuantian Shao, Shuang Qiu, Gang Li, Jing Liu, and Jian Cheng.

2026. [Certain head, uncertain tail: Expert-sample for test-time scaling in fine-grained moe](#). *Preprint*, arXiv:2602.02443.
- Peter Clark, Isaac Cowhey, Oren Etzioni, Tushar Khot, Ashish Sabharwal, Carissa Schoenick, and Oyvind Tafjord. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv:1803.05457v1*.
- Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Damai Dai, Chengqi Deng, Chenggang Zhao, R.x. Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Y. Wu, Zhenda Xie, Y.k. Li, Panpan Huang, Fuli Luo, Chong Ruan, Zhifang Sui, and Wenfeng Liang. 2024. [DeepSeekMoE: Towards ultimate expert specialization in mixture-of-experts language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1280–1297, Bangkok, Thailand. Association for Computational Linguistics.
- DeepSeek-AI, Aixin Liu, Bei Feng, Bin Wang, Bingxuan Wang, Bo Liu, Chenggang Zhao, Chengqi Deng, Chong Ruan, Damai Dai, Daya Guo, Dejian Yang, Deli Chen, Dongjie Ji, Erhang Li, Fangyun Lin, Fuli Luo, Guangbo Hao, Guanting Chen, and 138 others. 2024. [Deepseek-v2: A strong, economical, and efficient mixture-of-experts language model](#). *Preprint*, arXiv:2405.04434.
- Hongcan Guo, Haolang Lu, Guoshun Nan, Bolun Chu, Jialin Zhuang, Yuan Yang, Wenhao Che, Sicong Leng, Qimei Cui, and Xudong Jiang. 2025. [Advancing expert specialization for better moe](#). *Preprint*, arXiv:2505.22323.
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. 2021. Measuring massive multitask language understanding. *Proceedings of the International Conference on Learning Representations (ICLR)*.
- Xing Hu, Yuan Cheng, Dawei Yang, Zhixuan Chen, Zukang Xu, Chen Xu Jiangyong Yu, Zhihang Yuan, Zhe jiang, and Sifan Zhou. 2025. [OSTQuant: Refining large language model quantization with orthogonal and scaling transformations for better distribution fitting](#). In *The Thirteenth International Conference on Learning Representations*.
- Quzhe Huang, Zhenwei An, Nan Zhuang, Mingxu Tao, Chen Zhang, Yang Jin, Kun Xu, Kun Xu, Liwei Chen, Songfang Huang, and Yansong Feng. 2024a. [Harder task needs more experts: Dynamic routing in MoE models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 12883–12895, Bangkok, Thailand. Association for Computational Linguistics.
- Wei Huang, Yue Liao, Jianhui Liu, Ruifei He, Haoru Tan, Shiming Zhang, Hongsheng Li, Si Liu, and Xiaojuan Qi. 2024b. Mc-moe: Mixture compressor for mixture-of-experts llms gains more. *arXiv preprint arXiv:2410.06270*.
- Yushi Huang, Zining Wang, Zhihang Yuan, Yifu Ding, Ruihao Gong, Jinyang Guo, Xianglong Liu, and Jun Zhang. 2026. [Modes: Accelerating mixture-of-experts multimodal large language models via dynamic expert skipping](#). *Preprint*, arXiv:2511.15690.
- Yuzhen Huang, Yuzhuo Bai, Zhihao Zhu, Junlei Zhang, Jinghan Zhang, Tangjun Su, Junteng Liu, Chuancheng Lv, Yikai Zhang, Jiayi Lei, Yao Fu, Maosong Sun, and Junxian He. 2023. C-eval: A multi-level multi-discipline chinese evaluation suite for foundation models. In *Advances in Neural Information Processing Systems*.
- Robert A. Jacobs, Michael I. Jordan, Steven J. Nowlan, and Geoffrey E. Hinton. 1991. [Adaptive mixtures of local experts](#). *Neural Computation*, 3(1):79–87.
- Naman Jain, King Han, Alex Gu, Wen-Ding Li, Fanjia Yan, Tianjun Zhang, Sida Wang, Armando Solar-Lezama, Koushik Sen, and Ion Stoica. 2024. Live-codebench: Holistic and contamination free evaluation of large language models for code. *arXiv preprint*.
- Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, and 1 others. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088*.
- Peng Jin, Bo Zhu, Li Yuan, and Shuicheng Yan. 2024. [Moe++: Accelerating mixture-of-experts methods with zero-computation experts](#). *Preprint*, arXiv:2410.07348.
- Michael I. Jordan and Robert A. Jacobs. 1994. [Hierarchical mixtures of experts and the em algorithm](#). *Neural Computation*, 6(2):181–214.
- Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. [Efficient memory management for large language model serving with pagedattention](#). *Preprint*, arXiv:2309.06180.
- Mike Lasby, Ivan Lazarevich, Nish Sinnadurai, Sean Lie, Yani Ioannou, and Vithursan Thangarasa. 2026. [Reap the experts: Why pruning prevails for one-shot moe compression](#). *Preprint*, arXiv:2510.13999.
- Dmitry Lepikhin, Hyoungho Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2020. [Gshard: Scaling giant models with conditional computation and automatic sharding](#). *Preprint*, arXiv:2006.16668.

- Hunter Lightman, Vineet Kosaraju, Yura Burda, Harri Edwards, Bowen Baker, Teddy Lee, Jan Leike, John Schulman, Ilya Sutskever, and Karl Cobbe. 2023. Let’s verify step by step. *arXiv preprint arXiv:2305.20050*.
- Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2023. [Is your code generated by chat-GPT really correct? rigorous evaluation of large language models for code generation](#). In *Thirty-seventh Conference on Neural Information Processing Systems*.
- Xudong Lu, Qi Liu, Yuhui Xu, Aojun Zhou, Siyuan Huang, Bo Zhang, Junchi Yan, and Hongsheng Li. 2024. [Not all experts are equal: Efficient expert pruning and skipping for mixture-of-experts large language models](#). In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6159–6172, Bangkok, Thailand. Association for Computational Linguistics.
- Todor Mihaylov, Peter Clark, Tushar Khot, and Ashish Sabharwal. 2018. Can a suit of armor conduct electricity? a new dataset for open book question answering. In *EMNLP*.
- Niklas Muennighoff, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Jacob Morrison, Sewon Min, Weijia Shi, Pete Walsh, Oyvind Tafjord, Nathan Lambert, Yuling Gu, Shane Arora, Akshita Bhagia, Dustin Schwenk, David Wadden, Alexander Wettig, Binyuan Hui, Tim Dettmers, Douwe Kiela, and 5 others. 2025. [Olmoe: Open mixture-of-experts language models](#). *Preprint*, arXiv:2409.02060.
- Bhrij Patel, Souradip Chakraborty, Wesley A. Suttle, Mengdi Wang, Amrit Singh Bedi, and Dinesh Manocha. 2024. [Aime: Ai system optimization via multiple llm evaluators](#). *Preprint*, arXiv:2410.03131.
- Qwen, :, An Yang, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chengyuan Li, Dayiheng Liu, Fei Huang, Haoran Wei, Huan Lin, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jingren Zhou, and 25 others. 2025. [Qwen2.5 technical report](#). *Preprint*, arXiv:2412.15115.
- David Rein, Betty Li Hou, Asa Cooper Stickland, Jackson Petty, Richard Yuanzhe Pang, Julien Dirani, Julian Michael, and Samuel R. Bowman. 2023. [Gpqa: A graduate-level google-proof q&a benchmark](#). *Preprint*, arXiv:2311.12022.
- Keisuke Sakaguchi, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9):99–106.
- Lean Wang, Huazuo Gao, Chenggang Zhao, Xu Sun, and Damai Dai. 2024a. Auxiliary-loss-free load balancing strategy for mixture-of-experts. *arXiv preprint arXiv:2408.15664*.
- Mengru Wang, Xingyu Chen, Yue Wang, Zhiwei He, Jiahao Xu, Tian Liang, Qiuzhi Liu, Yunzhi Yao, Wenxuan Wang, Ruotian Ma, Haitao Mi, Ningyu Zhang, Zhaopeng Tu, Xiaolong Li, and Dong Yu. 2025a. [Two experts are all you need for steering thinking: Reinforcing cognitive effort in moe reasoning models without additional training](#). *Preprint*, arXiv:2505.14681.
- Yubo Wang, Xueguang Ma, Ge Zhang, Yuansheng Ni, Abhranil Chandra, Shiguang Guo, Weiming Ren, Aaran Arulraj, Xuan He, Ziyang Jiang, Tianle Li, Max Ku, Kai Wang, Alex Zhuang, Rongqi Fan, Xiang Yue, and Wenhui Chen. 2024b. [Mmlu-pro: A more robust and challenging multi-task language understanding benchmark](#). *Preprint*, arXiv:2406.01574.
- Yue Wang, Qiuzhi Liu, Jiahao Xu, Tian Liang, Xingyu Chen, Zhiwei He, Linfeng Song, Dian Yu, Juntao Li, Zhuosheng Zhang, Rui Wang, Zhaopeng Tu, Haitao Mi, and Dong Yu. 2025b. [Thoughts are all over the place: On the underthinking of o1-like llms](#). *Preprint*, arXiv:2501.18585.
- Fuzhao Xue, Zian Zheng, Yao Fu, Jinjie Ni, Zangwei Zheng, Wangchunshu Zhou, and Yang You. 2024. [Openmoe: An early effort on open mixture-of-experts language models](#). *arXiv preprint arXiv:2402.01739*.
- Shen Yan, Xingyan Bin, Sijun Zhang, Yisen Wang, and Zhouchen Lin. 2025. [Tc-moe: Augmenting mixture of experts with ternary expert choice](#). In *The Thirteenth International Conference on Learning Representations*.
- An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, and 41 others. 2025. [Qwen3 technical report](#). *Preprint*, arXiv:2505.09388.
- Rowan Zellers, Ari Holtzman, Yonatan Bisk, Ali Farhadi, and Yejin Choi. 2019. Hellaswag: Can a machine really finish your sentence? In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*.
- Shuzhang Zhong, Yanfan Sun, Ling Liang, Runsheng Wang, Ru Huang, and Meng Li. 2025. [Hybrid-moe: Hybrid cpu-gpu scheduling and cache management for efficient moe inference](#). *Preprint*, arXiv:2504.05897.
- Aojun Zhou, Ke Wang, Zimu Lu, Weikang Shi, Sichun Luo, Zipeng Qin, Shaoqing Lu, Anya Jia, Linqi Song, Mingjie Zhan, and Hongsheng Li. 2024. [Solving challenging math word problems using gpt-4 code interpreter with code-based self-verification](#). In *ICLR*.
- Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. 2023. [Instruction-following evaluation for large language models](#). *Preprint*, arXiv:2311.07911.

A More Related Work

Load Balance Loss and Expert Specialization.

The tension between load balancing and expert specialization in MoE training has been increasingly recognized in recent literature. While the auxiliary load balance loss is widely adopted to prevent routing collapse and ensure even expert utilization, several works have observed that it inadvertently constrains expert specialization and degrades model performance (Wang et al., 2024a; Guo et al., 2025). To mitigate this issue, prior work has explored two complementary directions: Wang et al. (2024a) eliminate the auxiliary loss entirely and instead enforce balance through dynamically updated expert-wise biases on routing scores, while Guo et al. (2025) retain the auxiliary loss but introduce additional orthogonality and variance objectives to strengthen specialization. These findings collectively confirm that the load balance loss, despite its necessity for stable training, limits expert specialization—an observation that aligns with and supports the motivation of our work.

Enhancing Expert Utilization in MoE. Most existing work attempts to improve expert utilization efficiency in MoE models by redesigning training objectives or expert architectures during pre-training (Yan et al., 2025; Huang et al., 2024a; Jin et al., 2024). TC-MoE (Yan et al., 2025) expands the expert space by applying a ternary set $\{-1, 0, 1\}$ to each expert and introduces a tailored load-balance and reward loss, enabling more efficient activation patterns and a flexible trade-off between effectiveness and efficiency. Huang et al. (2024a) break the fixed Top- K paradigm by proposing a dynamic expert selection framework that adjusts the number of activated experts based on input difficulty via a confidence threshold, activating more experts for complex reasoning and fewer for simple inputs; they further observe that lower layers rely more heavily on expert combinations than upper layers. (Chen et al., 2026) introduces Expert-Sample, a training-free method that preserves high-confidence selections while injecting controlled stochasticity into the uncertain tail in fine-grained MoE, enabling diverse generation without destabilizing outputs. MoE++ (Jin et al., 2024) introduces a heterogeneous MoE framework that integrates standard FFN experts with three types of zero-computation experts (zero, copy, and constant), allowing each token to engage with a dynamic number of FFNs, be adjusted by constant

vectors, or skip the MoE layer entirely, achieving $1.1\sim 2.1\times$ throughput speedup over vanilla MoE. While these approaches are insightful, they often sacrifice training stability for flexibility, therefore hard to be applied to larger scale MoE.

More recently, Wang et al. (2025a) identify a small subset of experts that are highly correlated with the `<think>` token and play a disproportionate role in reasoning, which they term *cognitive experts*; by amplifying the routing weights of these experts at inference time, they substantially enhance reasoning performance. However, this method is restricted to MoE models that explicitly employ the `<think>` token, limiting its general applicability.

Dynamic Expert Pruning of MoE. A growing body of work has observed substantial redundancy in the expert selection of MoE models, particularly in fine-grained MoE architectures where both the total number of experts and the number of activated experts per token are large (Lu et al., 2024; Huang et al., 2024b; Chen et al., 2025; Lasby et al., 2026). Motivated by this redundancy, dynamic expert pruning has emerged as a key technique for accelerating MoE inference. Lu et al. (2024) were the first to propose dynamic expert pruning, introducing a simple yet effective strategy that skips low-weight experts: when the routing weight of the top-1 expert exceeds that of the top-2 expert by a predefined threshold, only the top-1 expert is activated, thereby reducing computation without re-training. Building on this, Huang et al. (2024b) introduce a critical-token protection mechanism that exempts tokens with high attention scores from pruning, ensuring that performance-sensitive tokens retain their full set of experts. While the above approaches reduce computation from the *token side* by adjusting the number of experts each token activates, Chen et al. (2025) approach the problem from the *expert side*: rather than per-token decisions, they directly skip the inference of experts deemed globally less important, achieving more straightforward acceleration. More recently, Lasby et al. (2026) propose REAP (Router-weighted Expert Activation Pruning), a saliency criterion that jointly considers router gate values and expert activation norms, and computes a conditional average over the tokens an expert is activated on, thereby decoupling an expert’s functional contribution from its raw activation frequency. Beyond the LLM setting, Huang et al. (2026) extend dy-

dynamic expert pruning to MoE-based multimodal LLMs, designing pruning strategies that exploit modality-specific differences and effectively accelerating MoE-MLLM inference.

However, most of these works focus on non-reasoning tasks, and the substantial sensitivity differences across layers of MoE models and across tokens in long chain-of-thought reasoning are often overlooked, leaving room for further optimization.

B Details of Identifying Key Experts

In Section 3.2, we introduced our method for identifying key experts by measuring the KL divergence between the output distributions of the original model and variants where individual domain-specific experts are pruned. There, we focused on Qwen3-30B-A3B with math tasks as a case study, to illustrate how only a small subset of high-frequency experts exert a decisive influence on model predictions. Here, we provide complete details of the identification procedure as well as comprehensive results across **four different models** and **three task types** (math, general, and code). For each model, we report the KL divergence values of candidate experts when pruned, visualized as bar charts.

B.1 Selection Criterion of Key Experts

We first describe how the candidate set of key experts is filtered, which directly addresses the question of what specific divergence threshold is used to identify key experts. To prepare for this analysis, we collect the **domain-specialized experts** following the procedure in Section 3.1. For each task type, we record the expert usage frequencies of the model on a calibration set. For the Qwen3 series, we define domain-specialized experts as those whose usage frequency exceeds six times the expected frequency under perfectly uniform routing. For the DeepSeek series, where the distribution of expert usage is less sharply peaked, we instead adopt a relaxed threshold of three times the uniform baseline. The experts identified in this way serve as the candidates for the subsequent KL divergence experiments.

Among these candidates, we then rank experts by their KL divergence values (computed as described in Section B.3) and select the top-ranked experts as the final key experts. Rather than fixing a single absolute KL threshold, we adopt a top- N selection rule, where N is determined per model

and per domain by a clear gap in the sorted KL values: as illustrated in Figures 6–9, the KL distribution exhibits a sharp drop after a small number of dominant experts, naturally separating the key experts from the rest. This avoids the brittleness of a single fixed numerical cutoff while still yielding a well-defined and reproducible selection.

B.2 Calibration Sets for Each Task Type

This subsection specifies the calibration sets used for each of the three task types. We use MathQA for math, ARC-Challenge (Clark et al., 2018) for general, and the college-computer-science-compilers subset of MMLU (Hendrycks et al., 2021) for code. We mainly select datasets where only the first token of the answer is required as output, because this simplifies evaluation while still providing an accurate measure of the KL divergence.

It is worth emphasizing two points regarding these calibration sets. First, none of them overlaps with the benchmarks used in our main experiments (e.g., GSM8K, AIME24, Math500, GPQA-Diamond, MMLU-Pro, LiveCodeBench, HumanEval+), so the identified key experts are not tuned on the evaluation data, ruling out concerns of data leakage. Second, the identification procedure is not sensitive to the specific calibration dataset chosen and does not overfit to it. The underlying reason is that fine-grained MoE models consistently exhibit similar expert selection frequencies across different datasets within the same task category (Chen et al., 2025): math-related calibration sets activate a similar group of math-specialized experts, code-related calibration sets activate a similar group of code-specialized experts, and so on. As a result, key experts identified on one in-domain calibration set transfer well to other datasets of the same domain, which is also empirically validated by the strong cross-dataset generalization observed in our main results.

B.3 KL Divergence Computation and Results

Following prior work (Hu et al., 2025), we note that the model’s output logits typically follow a long-tail distribution, where only the top-ranked tokens carry substantial probability mass. Therefore, when computing KL divergence, we restrict attention to the **Top-1000 tokens**. Concretely, let p denote the probability distribution obtained from the original model, and q the distribution when a candidate expert is pruned. Denote by $\mathcal{I}_{1000}$ the index set of the top-1000 tokens under p . For each $i \in \mathcal{I}_{1000}$,

we define the restricted distributions as

$$p'(i) = \frac{p(i)}{\sum_{j \in \mathcal{I}_{1000}} p(j)},$$

$$q'(i) = \frac{q(i)}{\sum_{j \in \mathcal{I}_{1000}} q(j)}.$$

The KL divergence used in our analysis is then

$$D_{\text{KL}}(p' \parallel q') = \sum_{i \in \mathcal{I}_{1000}} p'(i) \log \frac{p'(i)}{q'(i)}.$$

This definition ensures that the comparison focuses on the most influential tokens, while ignoring the long tail of near-zero probabilities. The complete results for all models and task types are shown in Figures 6–9. Consistent with the case study in the main text, we observe across different models and domains that only a small subset of high-frequency experts causes substantial divergence, while the majority of experts induce negligible changes.

C Complete Results on Expert Specialization

In Section 3.1, we analyzed the specialization of experts in fine-grained MoE models using Qwen3-30B-A3B as a case study, and presented representative figures for a subset of layers. Here, we provide the complete visualization of expert usage frequencies across all layers of the model. Results are shown for three representative task types—math, code, and general. For each layer, the bar charts display the relative usage frequency of each expert, aggregated over all tokens in both the prefill and decode stages. Figures are grouped every 12 layers for readability. These results (Figures 10–13) confirm that different task types consistently activate distinct subsets of experts, highlighting clear patterns of domain specialization.

D Hyperparameter Choices in Dynamic Pruning

As introduced in Section 4, our dynamic pruning strategy involves two quantities: $K_{\min}$ and λ . The value of $K_{\min}$ is determined empirically to serve as a safeguard to prevent selecting too few experts. In contrast, λ acts as the single knob controlling pruning aggressiveness and presents a trade-off between efficiency and accuracy. We discuss their effects through detailed experiments below, demonstrating that both hyperparameters are robust and easy to set in practice.

Setting of $K_{\min}$. In the main text, we noted that reducing the number of selected experts below a certain threshold causes the model to collapse entirely. Therefore, $K_{\min}$ must guarantee that even when both layer- and token-level sensitivities are low, the number of selected experts never falls below this safe lower bound.

To study this effect, we take Qwen3-30B-A3B as an example and directly enforce the number of experts per token to be fixed between 2 and 7 (with $K_{\text{base}} = 8$ as the default). Table 6 reports the performance on five benchmarks. For efficiency, results on AIME2024 (except the baseline) are measured over a single run.

We observe that removing just one expert (TopK=7) leads to almost no performance degradation, consistent with the expert pruning observations reported in Section 2 in the main text. When TopK falls between 6 and 4, some performance degradation becomes apparent, but the model’s core reasoning capabilities remain largely intact—far from a performance collapse. The critical boundary occurs at TopK=3: performance collapses dramatically when fewer than three experts are selected. Therefore, we set $K_{\min} = 3$ to prevent the model from falling below this threshold. On the DeepSeek series, we observed a similar phenomenon, further reinforcing this choice. Notably, $K_{\min}$ is not intended as a finely tuned optimal value, but rather as a fixed safe lower bound specific to a given fine-grained MoE architecture. Once determined, it requires no further tuning across different tasks or datasets.

Choice of λ . The hyperparameter λ controls the trade-off between accuracy preservation and inference speedup: larger values of λ bias the pruning strategy toward retaining more experts (thus preserving accuracy), while smaller values encourage more aggressive pruning (thus improving throughput). To study this trade-off, we take Qwen3-30B-A3B as an example and vary λ from 0.5 to 0.9 in increments of 0.1. For each setting, we evaluate accuracy on the same five benchmarks as in the main experiments and measure the inference speedup using vLLM as the serving backend. The speedup is calculated as the ratio between the total inference time across all datasets and the baseline time without pruning. All results are reported from a single run.

From Table 7, we observe that when λ is set to an extremely small value (e.g., 0.5 or 0.6), prun-




Figure 6: KL divergence of domain-specific experts when pruned, for **DeepSeek-V2-Lite-Chat** across math, general, and code tasks.




Figure 7: KL divergence of domain-specific experts when pruned, for **DeepSeek-v2.5-1210** across math, general, and code tasks.

ing becomes overly aggressive—achieving higher speedups but at the cost of noticeable performance degradation. However, when $\lambda \geq 0.7$, the model achieves significant inference acceleration with minimal accuracy loss. Moreover, within this range, λ exhibits robust behavior: varying it from 0.7 to 0.9 does not substantially affect either accuracy or speedup. This suggests that as long as λ is not set to an extremely small value, it is a robust and easy-to-set hyperparameter. In practice, we recommend $\lambda = 0.7$ as a default choice that balances efficiency and accuracy well.

E More Results on Broader Task Categories

To further verify the universality and real-world applicability of our proposed method, we conduct additional evaluations on tasks beyond the domains covered in the main experiments. Specifically, we apply **Pick** to Qwen3-30B-A3B with only the “general” domain key experts enhanced, and evaluate on two held-out benchmarks: **C-Eval** (Huang et al., 2023), a comprehensive Chinese cross-disciplinary benchmark spanning 52 diverse categories ranging from humanities to STEM; and **IFEval** (Zhou et al., 2023), an instruction-following benchmark that is critical to real-world user experience.

As shown in Table 8, **Pick** consistently improves

accuracy on both benchmarks, achieving a gain of 1.74% on C-Eval and 1.85% on IFEval. These results demonstrate that **Pick** generalizes well to a broad range of real-world tasks beyond the originally tested domains, including multi-disciplinary knowledge reasoning and instruction-following scenarios. This further confirms the universality of leveraging key experts to enhance MoE model performance.

F On the Generality of the “General Reasoning” Domain

We further clarify that the “general reasoning” domain in our paper is in fact a broad task class rather than a narrow category. As illustrated by the word cloud in Figure 2, general-domain key experts are associated with branch tokens useful for generic reasoning, which benefits a wide range of downstream tasks.

Consistent with this observation, Table 5 shows that enhancing general-domain key experts alone not only improves GPQA-Diamond by 3.53%, but also yields gains on AIME24 (+1.33%, math) and HumanEval+ (+1.22%, code), even though these latter benchmarks belong to specialized domains. This indicates general-domain key experts capture transferable reasoning capabilities that complement domain-specific experts.




Figure 8: KL divergence of domain-specific experts when pruned, for **Qwen3-30B-A3B** across math, general, and code tasks.




Figure 9: KL divergence of domain-specific experts when pruned, for **Qwen3-235B-A22B** across math, general, and code tasks.

Moreover, the additional results in Appendix E further reinforce this point: by enhancing only the general-domain key experts, Pick achieves consistent improvements on both C-Eval (a cross-disciplinary knowledge benchmark) and IFEval (an instruction-following benchmark), neither of which is tied to a single specialized domain. This demonstrates that the general-domain key experts indeed serve as a broad backbone for diverse reasoning scenarios.

In practice, this offers a convenient deployment strategy: if one does not target a specific domain such as math or code, simply enhancing general-domain key experts can provide broad improvements across heterogeneous tasks without requiring domain-specific tuning.

G Additional Empirical Evidence for Expert Utilization Inefficiencies

To further support the empirical observations presented in Section 2 (expert utilization inefficiencies and optimization potential), we provide additional results on a broader set of benchmarks. During early exploration, we conducted both expert amplification and pruning experiments on multiple datasets, including ARC-Challenge (Clark et al., 2018), HellaSwag (Zellers et al., 2019), and OpenbookQA (Mihaylov et al., 2018), and observed

consistent conclusions. These early-stage findings motivated us to further investigate and ultimately validate our method on more complex reasoning tasks. Winogrande was selected as the illustrative example in the main text because it is widely used in evaluation and easy to reproduce.

Here, we report the supplementary results on ARC-Challenge, HellaSwag, and OpenbookQA using Deepseek-V2-Lite-Chat. We apply the same set of amplified key experts (L1E43, L10E33, L25E11) identified on Winogrande, as these benchmarks share a similar “general knowledge” nature, and likewise reduce expert selection from 6 to 5 in the pruning experiment.

As shown in Table 9, the two key observations in Section 2 consistently hold across all three additional benchmarks. First, jointly enhancing the identified key experts yields notable accuracy improvements (+2.17% on ARC-Challenge, +2.03% on HellaSwag, and +1.80% on OpenbookQA), confirming that a small number of impactful experts are systematically underutilized by the original routing strategy. Second, reducing the number of activated experts from 6 to 5 incurs only marginal accuracy changes (within $\pm 0.4\%$) while saving 16.7% of expert computation, indicating that the redundancy in expert selection is a general phenomenon rather than a dataset-specific artifact. Together, these results reinforce that the expert uti-

Table 6: Performance of Qwen3-30B-A3B under different fixed numbers of selected experts.

TopK	AIME2024	Math500	GPQA	Livecodebench	HumanEval+
8	80.67	95.2	65.66	91.75	71.34
7	76.67	95.0	65.66	91.25	71.95
6	73.33	93.2	63.64	89.75	69.51
5	70.00	88.2	59.60	87.00	66.46
4	63.33	78.4	54.55	74.75	59.76
3	40.00	55.8	44.95	57.50	43.29
2	16.67	19.6	21.72	25.50	17.07

Table 7: Accuracy and inference speedup of Qwen3-30B-A3B under different values of λ .

λ	Avg TopK	AIME2024	Math500	GPQA	Livecodebench	HumanEval+	Speedup
	8	80.67	95.2	65.66	91.75	71.34	1.00
0.5	4.33	73.33	93.8	59.60	86.50	65.85	1.30
0.6	4.61	76.67	94.0	63.64	89.25	68.90	1.26
0.7	4.82	80.00	94.6	64.65	90.00	70.12	1.25
0.8	5.09	80.00	94.0	64.65	90.75	70.12	1.22
0.9	5.32	80.00	94.8	64.14	90.50	70.73	1.20

Table 8: Results of **Pick** on broader task categories using Qwen3-30B-A3B with only the “general” domain key experts enhanced.

Benchmark	Method	Accuracy (%)
C-Eval	Baseline	86.60
	+ Pick	88.34
IFEval	Baseline	85.95
	+ Pick	87.80

Table 9: Supplementary results to support the observations in Section 2. “+ Key Expert” denotes jointly enhancing the three identified key experts, while “Pruned (K=5)” reduces the number of selected experts per token from 6 to 5.

Dataset	Baseline	+ Key Expert	Pruned (K=5)
ARC-Challenge	54.44	56.61	54.05
HellaSwag	62.44	64.47	62.25
OpenbookQA	45.60	47.40	45.80

lization inefficiencies identified in the main paper are robust and broadly applicable, providing a solid empirical foundation for our proposed Ban&Pick method.

H Mixture-of-Experts

The Sparsely-Gated MoE (Jacobs et al., 1991; Jordan and Jacobs, 1994) first introduced the idea of using a trainable gating network to activate only a small subset of experts, typically through a Top-K routing scheme. Subsequent studies (DeepSeek-AI et al., 2024; Yang et al., 2025) built upon this framework, scaling MoE models to much larger sizes and demonstrating that sparse expert activation can de-

liver strong efficiency and performance gains. As these models expanded, the routing mechanism became a critical factor in determining overall effectiveness.

The Sparsely-Gated MoE are based on a transformer architecture, but the FeedForward Network (FFN) sub-layers of traditional dense models are replaced with MoE layers, each containing N experts. For each input token $\mathbf{x}$, the router computes routing logits $\mathbf{r} = \{r_0, \dots, r_{N-1}\}$ and expert selection scores $\mathbf{s} = \text{Softmax}(\mathbf{r})$. The top- K experts are selected based on $\mathbf{s}$, and their outputs $E_{e_j}(\mathbf{x})$ are combined as a weighted sum, with normalized weights:

$$\mathbf{z} = \sum_{j=0}^{K-1} \frac{\mathbf{s}_{e_j}}{\sum_{i=0}^{K-1} \mathbf{s}_{e_i}} \cdot E_{e_j}(\mathbf{x}). \quad (1)$$

Here, $E_{e_j}(\mathbf{x})$ represents the output of the j -th selected expert for the input token $\mathbf{x}$. Notably, different MoE architectures handle expert weight combination differently. Qwen3 employs the weight renormalization mechanism shown in Equation 1, where the selected expert weights are normalized to sum to 1. In contrast, DeepSeek-v2 uses direct weight scaling without renormalization. Moreover, DeepSeek-v2 incorporates grouped expert selection, where experts are partitioned into groups to improve specialization. Detailed settings and implementation can be checked in their papers (DeepSeek-AI et al., 2024).

Despite such architectural differences, both lines of work—and many other recent systems—adopt a fine-grained MoE design: each layer hosts a large

number of small, independently initialized experts. This design enhances expert specialization but simultaneously increases the difficulty of ensuring efficient routing. In our experimental evaluation, we examine four representative fine-grained MoE models, whose detailed architectural configurations and parameter settings are summarized in Table 10. Note that in the DeepSeek-v2 architecture design (including Deepseek-V2-Lite-Chat and Deepseek-v2.5-1210), the first layer is a dense layer for training stability, while the remaining layers are MoE layers. In contrast, the Qwen3 series uses MoE layers throughout all layers.

I Evaluation Details

I.1 Dataset Evaluation.

For **AIME2024**, due to the small sample size, we follow the common practice in technical reports by evaluating each method five times and reporting the average. For other datasets, constrained by time, we conduct only a single evaluation. On the **GPQA-Diamond** dataset, to further reduce errors caused by misaligned answer formats, we adopt an LLM-judge in addition to the base evaluator: specifically, we deploy a Qwen2.5-72B-Instruct model (Qwen et al., 2025) with vLLM (Kwon et al., 2023) as the judge model. For all other datasets, we rely on the default base evaluators.

I.2 Generation Hyperparameters.

During generation, for all four models used in our experiments, we adopt the recommended hyperparameters provided in their papers or Model Cards. The configuration is summarized in Table 11.

J Details of Compared Methods and Reproduction Settings

In the main experiments, we compared our proposed BAN and PICK strategies against several existing baselines. For brevity, the main text only provided a concise overview of these methods. In this appendix, we present a more detailed description of each compared method, together with the specific reproduction settings used in our experiments. We organize the discussion into two parts: methods compared with PICK, and methods compared with BAN.

J.1 Compared Methods for PICK

Dynamic Routing. (Huang et al., 2024a) first proposed the idea of dynamically allocating ex-

perts, arguing that assigning a fixed number of experts to each token is overly rigid. Instead, the number of experts should be adapted according to the difficulty of the underlying task: tokens from challenging tasks should be allocated more experts, while tokens from easier tasks can be served with fewer. Concretely, experts are sequentially selected in descending order of routing weight until the cumulative weight exceeds a predefined threshold τ . Formally, let $r_{(1)}, r_{(2)}, \dots, r_{(K)}$ denote the routing weights of experts sorted in descending order. The number of selected experts k is determined as

$$k = \min \left\{ m \mid \sum_{j=1}^m r_{(j)} \geq \tau \right\}.$$

In our experiments, since the optimal threshold τ cannot be known in advance, we evaluate three candidate values $\tau \in \{0.7, 0.8, 0.9\}$ and report in the main text the result corresponding to the setting that achieves the best average accuracy across benchmarks.

TiP. The core idea of TiP is motivated by the observation that, during chain-of-thought (CoT) reasoning, models often switch reasoning paths too frequently in the early exploration stage. Such premature abandonment of a correct path is referred to as *underthinking*. To mitigate this issue, TiP introduces a penalty on tokens that typically indicate a switch of reasoning path (e.g., “alternatively”). Concretely, after the softmax operation, the probability mass of these tokens is reduced by subtracting a fixed penalty value from their logits, thereby encouraging the model to continue deepening its current reasoning path rather than switching too early.

This method involves two hyperparameters: α (penalty strength) and β (penalty duration). The original paper did not provide an automatic tuning strategy but instead conducted a grid search to identify relatively effective values. Following their reported settings, we set $\alpha = 3$ and $\beta = 600$ in our reproduction.

Rice. Rice builds on the observation that recent MoE-LLMs often employ an explicit <think> token to train chain-of-thought (CoT) reasoning ability. For example, the Qwen3 series can switch between deep-thinking and non-thinking modes by toggling the use of the <think> token. Rice leverages this property by comparing expert activation frequencies with and without the <think>

Table 10: Specifications of fine-grained MoE models used in our experiments.

Model	Experts	Top-K	Shared-Experts	Layers	Parameters
Deepseek-V2-Lite-Chat	64	6	2	27	16B
Deepseek-v2.5-1210	160	6	2	60	236B
Qwen3-30B-A3B	128	8	0	48	30B
Qwen3-235B-A22B	128	8	0	94	235B

Table 11: Generation hyperparameters for the evaluated models.

Model	Top-k	Top-p	Temperature	Max Output Length
DeepSeek-V2-Lite-Chat	None	0.95	0.3	8192
DeepSeek-V2.5-1210	1	0.90	0.3	8192
Qwen3-30B-A3B	20	0.95	0.6	32768
Qwen3-235B-A22B	20	0.95	0.6	32768

token, thereby identifying cognitive experts that are closely associated with deep reasoning. Once identified, the routing weight of a cognitive expert is multiplied by a constant factor greater than 1 whenever it is selected, amplifying its influence and enhancing the model’s reasoning ability.

Since this method relies on explicit <think> tokens, it is only applicable to models such as the Qwen3 series, but not to other models used in our experiments. For reproduction, we followed the original paper’s settings: for Qwen3-235B-A22B, we directly used the cognitive experts reported therein; for Qwen3-30B-A3B, as the original paper did not list them, we applied the ICE (Identify Cognitive Experts) procedure proposed in the paper to locate the corresponding cognitive experts before applying the enhancement.

J.2 Compared Methods for BAN

DES. DES (Dynamic Expert Skip) was proposed to dynamically adjust the number of active experts by exploiting the observation that, for tokens that are relatively easy to process, the routing weights of the selected experts often differ significantly. In such cases, the contribution of the lower-weighted expert to the final output is negligible. Based on this, DES introduces a dynamic pruning strategy: whether to keep the lower-weighted expert is decided according to the ratio between expert weights. For example, in Mixtral-8×7B (Jiang et al., 2024) where each token selects two experts, DES first computes the median ratio of top-1 to top-2 expert weights on a calibration set. During inference, if the actual ratio exceeds this median, the top-2

expert is skipped.

In our reproduction, we aimed to align the pruning degree and acceleration ratio with our proposed Ban method, so that the comparison mainly reflects differences in accuracy preservation. For Qwen3 models, when $\lambda = 0.7$, the average number of selected experts under Ban is slightly below 5. Therefore, we set the maximum number of experts to 6 and the minimum to 4. Specifically, we computed the medians of the ratios $\frac{r_{(4)}}{r_{(5)}}$ and $\frac{r_{(5)}}{r_{(6)}}$ on a calibration set, where $r_{(j)}$ denotes the j -th largest routing weight. At inference time, if the first ratio exceeds its median, only the top-4 experts are kept; otherwise, if the first ratio does not exceed its median but the second does, the top-5 experts are kept; otherwise, all top-6 experts are retained. For the DeepSeek models, we followed an analogous procedure to ensure comparable pruning degree and speedup.

ODP. ODP (Online Dynamic Pruning) is an extension of DES that incorporates an additional key-token protection mechanism. While DES prunes experts dynamically based on weight ratios, ODP further observes that certain tokens carry disproportionately high importance in the attention map. If a token receives an attention score significantly larger than the others, pruning its experts may harm model performance. To address this, ODP protects such key tokens: even if their expert weights satisfy the pruning condition, they are exempted from pruning.

In our reproduction, we followed exactly the same procedure as DES for the pruning part. On

top of that, we adopted the evaluation metric for token importance proposed in the ODP paper to identify key tokens. Whenever a token is classified as key, we enforced the selection of all chosen experts for it, regardless of the pruning condition.

K Discussion on Strategy Selection for Pick Enhancement

In the main text (Section 4.3), we showed through experiments on MATH-500, AIME2024, and AIME2025 that among the five proposed enhancement strategies for key experts, **Strategy D** achieves the most consistent accuracy improvements. Here, we provide a more detailed explanation of the motivations behind the design of these strategies and the reasoning that led us to ultimately adopt Strategy D.

Initial motivation: Strategies A and B. An initial intuition was straightforward—since key experts play a decisive role in task-specific reasoning, one might expect model performance to improve if all tokens are explicitly routed through or enhanced by these key experts. Based on this idea, **Strategies A and B** were designed to directly enforce the participation of key experts for all tokens, aiming to strengthen their influence across the entire sequence regardless of token type or context. This approach represents the most direct attempt to leverage key experts’ strong impact on output generation to achieve better overall results.

Why Strategies A and B are not stable. However, subsequent analyses in Sections 4.1 and 4.2 of the main paper suggest that key experts are closely tied to task-relevant tokens and can be regarded as a subset of domain-specialized experts. Taking the *math* domain as an example, these key experts are highly activated by digits, letters, and operators, but not by unrelated tokens such as pronouns or prepositions. Consequently, forcing key experts to be enhanced for every token disregards this natural specialization and leads to unstable or degraded performance across datasets. This behavior is **consistent with the experimental results in Section 4.3**, where Strategies A and B exhibit larger performance fluctuations and reduced accuracy compared with other strategies. These observations confirm that a more adaptive mechanism is needed to selectively enhance key experts where they are most beneficial, rather than applying uniform enforcement.

Adaptive Enhancement Yields Better Performance. To address the instability observed in Strategies A and B, we designed adaptive strategies that selectively enhance key experts based on token-level routing signals rather than enforcing fixed participation. A natural idea might be to construct an explicit lookup table of task-relevant tokens so that enhancement can be applied in a perfectly fine-grained manner. However, such an approach would be overly cumbersome in implementation and inconsistent with our design objective of maintaining simplicity and efficiency. Instead, we empirically observed that for tokens within a key expert’s corresponding domain—for example, math-related tokens—the key experts often appear within the top- $2k$ routing weight range even when they are not selected in the top- k (k is the default number of experts per token), whereas for unrelated tokens they seldom appear within the top- $2k$. This observation provides a lightweight and effective criterion for identifying which tokens should trigger the enhancement of their domain-specific experts.

We conducted this verification on QWEN3-30B-A3B using the math key expert (2,37). All tests were performed on the *math calibration subset*. To analyze cross-domain behavior, we selected the top-20 most frequent tokens from both the *math* and *general* word clouds shown in Figure 2 of the main paper. For each token, we measured the frequency with which the math key expert appears in the top- k , in the top- $2k$ but not in the top- k , and outside the top- $2k$ routing weight range. The results are presented in Table 12.

The results show that, even when evaluated purely on the math calibration set, the math key expert (2,37) is highly activated for math-related tokens—appearing in the top- $2k$ routing range for more than 90% of them—while rarely activated for general tokens. This indicates that the routing mechanism naturally captures a strong correlation between math tokens and the math key expert, validating that the adaptive strategy can precisely align experts with domain-relevant tokens. Consequently, **Strategies C and D** leverage this intrinsic property to enhance key experts selectively, improving performance on relevant tokens without blind strengthening across the vocabulary.

Choice between Strategy C and D. The primary difference between Strategies C and D lies

Table 12: Routing frequency (%) of the math key expert (2, 37) on the math calibration set. Tokens are taken from the *math* and *general* word clouds in Figure 2. Each column shows the proportion of tokens whose key expert appears in different routing weight ranges.

Token Type (20 tokens each)	Top- k (%)	Top-2 k but not Top- k (%)	Outside Top-2 k (%)
Math tokens	68.4	24.7	6.9
General tokens	2.5	0.4	97.1

in whether the minimally weighted selected expert is replaced. Our observations indicate that the least-weighted expert contributes negligibly to model performance, consistent with the motivation of the **Ban** module: pruning or replacing experts of minimal contribution improves efficiency with minimal accuracy loss. Thus, replacing or retaining this expert makes little difference; we therefore adopted **Strategy D** for its simplicity and implementation efficiency.

Future directions. We believe this analysis offers deeper insight into the design rationale of Pick. Future extensions may incorporate token-adaptive or data-driven policies to automatically determine enhancement conditions, potentially providing finer control and better synergy with the underlying Ban mechanism while maintaining the plug-and-play nature of the framework.

Table 13: System-level evaluation of Ban&Pick on Qwen3-30B-A3B.

Metrics	Base	+Ban&Pick	Change
Avg Top- k	8.00	4.82	−39.75%
Throughput (tokens/s)	10.01	12.85	+28.37%
Peak Memory (GB)	5.84	5.32	−8.90%

L System-Level Performance Evaluation

To conduct a more detailed system-level performance evaluation of our proposed **Ban&Pick** method, we performed additional experiments on the **Qwen3-30B-A3B** model. The evaluation focuses on **throughput** and **memory usage**, which are important metrics for assessing the efficiency of large-scale models in real deployment scenarios.

All tests were executed on a single A800-80G GPU. We used identical prompts and identical inference configurations for both the base model and the Ban&Pick-enhanced model (*maximum output length* = 1k, two warm-up runs, and five averaged runs). The average throughput and peak memory usage are summarized in Table 13.

The results show that, due to expert pruning introduced by the Ban module, **throughput increases notably by +28.37%**, while **peak memory usage decreases slightly**. These results demonstrate that Ban&Pick not only improves model accuracy but also enhances **efficiency and deployability** in practical inference scenarios.

M LLM Usage

In the preparation of this manuscript, large language models were used to polish the writing, including grammar refinement, phrasing improvement, and minor stylistic edits.




Figure 10: Expert usage frequency for layers 0–11 across math, code, and general tasks.




Figure 11: Expert usage frequency for layers 12–23 across math, code, and general tasks.




Figure 12: Expert usage frequency for layers 24–35 across math, code, and general tasks.




Figure 13: Expert usage frequency for layers 36–47 across math, code, and general tasks.