

Mosaic Pruning: A Hierarchical Framework for Generalizable Pruning of Mixture-of-Experts Models

Wentao Hu^{1, 2*}, Mingkuan Zhao^{1*}, Shuangyong Song², Xiaoyan Zhu¹, Xin Lai¹,
Jiayin Wang^{1†}

¹Xi'an Jiaotong University

²Institute of Artificial Intelligence (TeleAI), China Telecom

wentao.hu@stu.xjtu.edu.cn, mingkuanzhao@stu.xjtu.edu.cn, songsy@chinatelecom.cn, zhu.xy@xjtu.edu.cn,
laixin@xjtu.edu.cn, wangjiayin@mail.xjtu.edu.cn

Abstract

Sparse Mixture-of-Experts (SMoE) architectures have enabled a new frontier in scaling Large Language Models (LLMs), offering superior performance by activating only a fraction of their total parameters during inference. However, their practical deployment is severely hampered by substantial static memory overhead, as all experts must be loaded into memory. Existing post-training pruning methods, while reducing model size, often derive their pruning criteria from a single, general-purpose corpus. This leads to a critical limitation: a catastrophic performance degradation when the pruned model is applied to other domains, necessitating a costly re-pruning for each new domain. To address this generalization gap, we introduce Mosaic Pruning (MoP). The core idea of MoP is to construct a functionally comprehensive set of experts through a structured “cluster-then-select” process. This process leverages a similarity metric that captures expert performance across different task domains to functionally cluster the experts, and subsequently selects the most representative expert from each cluster based on our proposed Activation Variability Score. Unlike methods that optimize for a single corpus, our proposed Mosaic Pruning ensures that the pruned model retains a functionally complementary set of experts, much like the tiles of a mosaic that together form a complete picture of the original model’s capabilities, enabling it to handle diverse downstream tasks. Extensive experiments on various MoE models demonstrate the superiority of our approach. MoP significantly outperforms prior work, achieving a 7.24% gain on general tasks and 8.92% on specialized tasks like math reasoning and code generation.

Code — <https://github.com/Saul-James/MoP>

Introduction

Large Language Models (LLMs) have recently demonstrated remarkable capabilities in complex reasoning and generation tasks (OpenAI 2024; Touvron et al. 2023b; Xiong et al. 2024). To mitigate their computational requirements, the Mixture-of-Experts (MoE) architecture has been widely adopted (Jiang et al. 2024). MoE models, such as Mixtral

8x7B, activate only a portion of their parameters during inference, enabling them to surpass the performance of larger dense models like Llama 2 70B while maintaining a smaller count of active parameters (Jiang et al. 2024; Touvron et al. 2023b).

Despite this efficiency, MoE models are constrained by a critical deployment challenge: immense static memory overhead. For instance, deploying Mixtral 8x7B requires over 80GB of GPU memory (Jiang et al. 2024). More importantly, due to training dynamics and representation learning differences, significant redundancy exists among experts in MoE models (Chi et al. 2022; Yu et al. 2022), with some experts being functionally similar or contributing minimally to most tasks.

Leveraging the similarity between experts, post-training pruning has emerged as a promising approach for compressing MoE models. The state-of-the-art pruning method, Enumeration Pruning (Lu et al. 2024), utilizes a general-purpose calibration dataset (e.g., C4 and WikiText) to identify and remove experts that contribute least to token reconstruction loss. However, this method reveals a critical limitation: the inability to generalize to specialized downstream tasks. A model pruned on a general-purpose corpus suffers a catastrophic performance drop when directly applied to domain-specific tasks such as mathematical reasoning or code generation, a phenomenon we term functional collapse. We argue that this collapse occurs because such methods inherently favor retaining generalist experts that make moderate contributions across common data patterns, while discarding specialist experts that possess critical domain-specific expertise but exhibit less prominent activation on general datasets. This necessitates re-pruning with a new calibration dataset for each new domain, a process that is not only impractical but also fundamentally undermines the model’s applicability. Given that different experts in MoE specialize in distinct knowledge domains, we argue that a model’s generalization capability fundamentally depends on the breadth of its functional diversity.

To address these challenges, we first propose a pruning strategy that actively preserves expert diversity, which we term Global Variability-aware Pruning (GVP). This method calculates an Activation Variability Score (S_{var}), constructed

*These authors contributed equally.

†Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

based on the Kullback-Leibler (KL) divergence (Kullback and Leibler 1951), for each expert and prioritizes retaining those with the highest scores. While our experiments confirm that GVP is superior to Enumeration Pruning, relying solely on the Activation Variability Score (S_{var}) may lead to the retention of two functionally similar experts that both possess high diversity scores, which still results in redundancy among the experts. To overcome this limitation, we propose a hierarchical pruning framework aimed at preserving functional diversity, which we call Mosaic Pruning (MoP). Our core idea is twofold: first, we use a composite metric, primarily driven by a Task Performance Similarity matrix (S_{perf}) (Spearman 1987), to assess functional similarity by analyzing expert performance across different task domains; second, within each functionally defined group, we select the most critical representative based on the Activation Variability Score proposed in GVP. This strategy ensures that no essential functional expertise is entirely eliminated, thus providing a strong inductive bias toward generalization across diverse tasks. Our extensive experiments on models such as Mixtral-8x7B (Jiang et al. 2024) and Qwen1.5-MoE-A2.7B (Yang et al. 2024) demonstrate the superiority of our method.

Our contributions are summarized as follows:

- We propose a novel diversity metric, the Activation Variability Score, to effectively identify experts with specialized functionalities.
- We introduce a novel pruning framework that first groups experts based on their task performance similarity, and then prunes within each group to select a functionally complementary set of specialists based on our diversity metric.
- We validate the superiority of our proposed method through extensive experiments, achieving an average performance improvement of 7.24% on general tasks, and a more significant 8.92% on specialized tasks such as mathematical reasoning and code generation.

Related Works

Mixture-of-Experts Models

The Mixture-of-Experts (MoE) paradigm, first introduced by (Jacobs et al. 1991), involves a collection of distinct sub-networks, or “experts,” coordinated by a trainable gating network that selectively routes inputs. This modular architecture was successfully integrated into modern deep learning with its application to Recurrent Neural Networks (Shazeer et al. 2017) and later scaled up within large Transformer-based models, most notably in the encoder-decoder architecture of GShard (Lepikhin et al. 2020). With the recent ascendancy of decoder-only architectures like the GPT series (Mann et al. 2020; Touvron et al. 2023a,b), there has been a surge in the development of powerful open-source MoE-based LLMs, such as Mixtral (Jiang et al. 2024), Qwen-MoE (Yang et al. 2024), and DeepSpeed-MoE (Rajbhandari et al. 2022), and the Tele-FLM series (He et al. 2024; Li et al. 2024b,a; Shao and Li 2025; Wang et al. 2024, 2025). A key advantage of MoE is its ability to decouple the model’s

total parameter count from its computational cost per token, enabling a massive increase in model capacity without a proportional rise in inference latency (Fedus, Zoph, and Shazeer 2022). However, this efficiency comes at the cost of significant memory overhead, as all experts must be loaded into memory, which motivates research into compressing these models.

Model Pruning

Model pruning aims to reduce the complexity of neural networks by removing redundant parameters, thereby enhancing efficiency. Pruning techniques are broadly categorized into two types: unstructured and structured pruning.

Unstructured Pruning, such as magnitude pruning (Han, Mao, and Dally 2015), removes individual weights based on certain criteria. While it can achieve high sparsity levels, its practical application in LLMs has been challenging. Recent innovations like SparseGPT (Frantar and Alistarh 2023) and Wanda (Sun et al. 2023) have revisited this concept for LLMs, developing sophisticated post-training methods that prune weights based on magnitude and activation statistics without requiring retraining. Other works reduce overhead via sparse attention (Zhao et al. 2025). However, the irregular sparsity patterns produced by these methods often necessitate specialized hardware or software support to realize actual inference speedups.

Structured Pruning, in contrast, removes entire structural units of the network, such as neurons, attention heads, or, in the context of MoE models, entire experts. This form of pruning is inherently hardware-friendly, as it results in smaller, dense models that can be efficiently executed on standard hardware like GPUs. For MoE models, expert-level pruning is a particularly compelling form of structured pruning. Early works in this area often focused on task-specific scenarios, such as machine translation, where experts specializing in certain languages could be discarded (Kim et al. 2021). More recent approaches have focused on task-agnostic, post-training pruning for large-scale MoE LLMs. A prominent example is the reconstruction-based pruning method proposed by (Lu et al. 2024). This approach systematically enumerates expert combinations to identify the subset that minimizes token reconstruction loss on a general-purpose corpus. While effective at preserving performance on the calibration domain, our work identifies a critical limitation in this approach—a failure to generalize to specialized downstream tasks—which directly motivates our development of a diversity-aware pruning framework.

Methodology

In this section, we present the details of our proposed hierarchical pruning framework. We introduce our contributions in a progressive sequence: first, we detail an improved strategy that incorporates global expert variability, which we term Global Variability-aware Pruning (GVP) (Section). Building upon the insights and limitations of GVP, we then present our ultimate solution: a Hierarchical framework that combines functional clustering and diversity-aware selection, which we call Mosaic Pruning (MoP), designed to

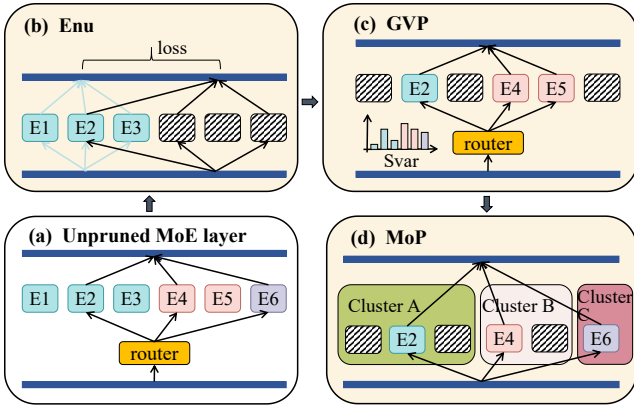


Figure 1: A conceptual illustration of the expert pruning strategies. (a) Unpruned MoE Layer: The original experts, with colors indicating different functional specializations. (b) Enu: Retains a functionally homogeneous set of experts (E1, E2, E3) by minimizing reconstruction loss. (c) GVP: Supplements core experts with globally selected specialists (E4, E5), improving diversity but risking functional overlap. (d) MoP: Clusters experts by functional similarity and selects a representative from each cluster, ensuring a final expert set that is both specialized and complementary.

achieve an optimal balance between performance and diversity (Section). Figure 1 provides a conceptual overview of the pruning methods discussed in this paper.

Enumeration Pruning (Enu)

The most intuitive strategy for expert pruning is to quantify and minimize the perturbation to the model’s output upon expert removal. This method (Enumeration Pruning), which forms the basis of our baseline, follows the core idea of (Lu et al. 2024) by searching a general-purpose calibration dataset (e.g., C4) to find the optimal expert retention scheme. It aims to find the subset of r experts, $\mathcal{E}_{\text{kept}}$, that minimizes the output reconstruction loss, an optimization problem formally expressed as:

$$\mathcal{E}_{\text{kept}}^* = \arg \min_{\mathcal{E}_{\text{kept}} \subset \mathcal{E}, |\mathcal{E}_{\text{kept}}| = r} \mathcal{L}_{\text{recon}}(\mathcal{D}_{\text{cache}}, \mathcal{E}_{\text{kept}}), \quad (1)$$

where the total reconstruction loss $\mathcal{L}_{\text{recon}}$ is calculated over a cached dataset $\mathcal{D}_{\text{cache}}$ —a pre-computed set of input hidden states and their corresponding original layer outputs. The loss quantifies the cumulative difference between the outputs of the original, unpruned MoE layer and the pruned layer where only the experts in $\mathcal{E}_{\text{kept}}$ are active. In our design, for models with a small number of experts, such as Mixtral 8x7B ($n = 8$), we use enumeration to precisely solve this optimization problem. However, for models with a large number of experts, like Qwen1.5-MoE-A2.7B ($n = 60$), the number of combinations $\binom{60}{r}$ renders exhaustive enumeration infeasible. In this case, we employ a greedy strategy as an approximation, iteratively removing the single expert that contributes the least to the increase in reconstruction loss until the target number is reached. While Enumeration

Pruning is conceptually simple, as mentioned in the Introduction, it tends to preserve functionally generalist experts at the expense of specialist experts crucial for downstream professional tasks.

Global Variability-aware Pruning (GVP)

To address the issue of Enumeration Pruning erroneously removing specialist experts due to its over-reliance on reconstruction loss, we propose Global Variability-aware Pruning (GVP). The core idea of this method is to apply a unified variability metric from a global perspective to prioritize the selection of functional specialists. The core metric of this method is the Activation Variability Score (S_{var}), which quantifies an expert’s functional specialization.

For any expert i , we define its variability as the Kullback-Leibler (KL) divergence (Kullback and Leibler 1951) between its normalized activation distribution and the uniform distribution over all tokens. Let N_{total} be the total number of tokens in the calibration dataset. The score is calculated as:

$$S_{\text{var}}(i) = \sum_{t=1}^{N_{\text{total}}} \frac{p_{t,i}}{Z_i} \log_2 \left(\frac{p_{t,i}}{Z_i} N_{\text{total}} \right). \quad (2)$$

Here, $p_{t,i}$ is the Softmax activation probability for expert i on token t , $Z_i = \sum_{t=1}^{N_{\text{total}}} p_{t,i}$ is the normalization constant, $P_i = \left\{ \frac{p_{t,i}}{Z_i} \right\}_{t=1}^{N_{\text{total}}}$ is the normalized activation distribution, and $U = \left\{ \frac{1}{N_{\text{total}}} \right\}_{t=1}^{N_{\text{total}}}$ is the uniform distribution over tokens. A high S_{var} score indicates a concentrated activation pattern on a specific subset of inputs, suggesting it is a functional specialist.

The GVP pruning strategy is as follows:

Retaining the General Experts. We first apply the process from Section to select a core set of m ($m < r$) general experts ($\mathcal{E}_{\text{general}}$), which are most critical for maintaining baseline performance on general data and providing general capabilities:

$$\mathcal{E}_{\text{general}} = \arg \min_{S \subset \mathcal{E}, |S| = m} \mathcal{L}_{\text{recon}}(\mathcal{D}_{\text{cache}}, S). \quad (3)$$

This step aims to identify and retain a set of the most functionally general experts to ensure that the model’s core general capabilities are not compromised. The remaining $n - m$ experts in the candidate pool, $\mathcal{E}_{\text{cand}} = \mathcal{E} \setminus \mathcal{E}_{\text{general}}$, to be pruned according to a diversity score in the following selection stage.

Selecting Diverse Experts via Global Variability. From the candidate pool $\mathcal{E}_{\text{cand}}$, we apply the Activation Variability score S_{var} in a global ranking to select the remaining $r - m$ most specialized experts, forming the diversity set \mathcal{E}_{div} :

$$\mathcal{E}_{\text{div}} = \arg \max_{S \subset \mathcal{E}_{\text{cand}}, |S| = r - m} \sum_{i \in S} S_{\text{var}}(i). \quad (4)$$

The final set of retained experts for the layer is $\mathcal{E}_{\text{final}}^{\text{GVP}} = \mathcal{E}_{\text{general}} \cup \mathcal{E}_{\text{div}}$. Although GVP is superior to the enumeration reconstruction loss strategy, its global ranking in the second stage has a critical flaw: it may retain two functionally similar yet highly specialized experts, which still results in redundancy within the final expert set.

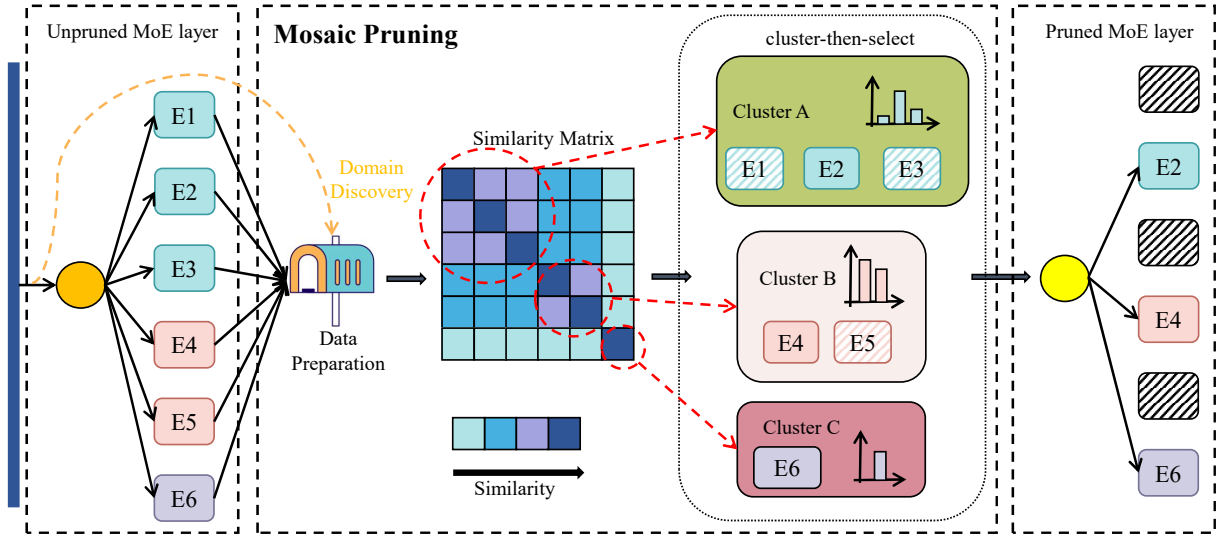


Figure 2: The workflow of the Mosaic Pruning (MoP) framework. First, the calibration data is partitioned into distinct functional domains. Subsequently, a similarity matrix between experts is constructed based on their performance profiles across these domains. This matrix is used to cluster experts with high similarity into the same group. Finally, experts are selected within each cluster based on their Activation Variability Score.

Mosaic Pruning (MoP)

To fundamentally address the inherent flaws of GVP, we propose **Mosaic Pruning (MoP)**, a hierarchical framework that systematically preserves expert functional diversity. MoP shares the same first stage as GVP, retaining a shared expert set to secure baseline performance. Its core innovation lies in its second stage, which replaces GVP’s simple global ranking with a structured “cluster-then-select” process based on functional similarity, as illustrated in Figure 2.

Data-Driven Domain Discovery. To evaluate expert performance across different functional domains, MoP utilizes a specially constructed mixed-diversity calibration dataset. This multi-domain design is crucial for enabling the automatic discovery of latent functional domains that align with expert specializations. The detailed construction process for this dataset is provided in Appendix A. We apply K-Means (McQueen 1967) clustering to the input hidden states of all tokens, $\{\mathbf{x}_t\}_{t=1}^{N_{\text{total}}}$. We posit that because experts exhibit domain-specific behaviors, and our mixed-diversity dataset is designed to reveal clear functional boundaries between different knowledge domains, K-Means is well-suited for capturing these domain-driven functional groupings by partitioning the token embeddings into semantically coherent domains. The objective is to find a partition $\mathcal{C} = \{C_1, C_2, \dots, C_K\}$ that minimizes the Within-Cluster Sum of Squares (WCSS):

$$\min_{\mathcal{C}} \sum_{k=1}^K \sum_{\mathbf{x}_t \in C_k} \|\mathbf{x}_t - \boldsymbol{\mu}_k\|_2^2, \quad (5)$$

where $\boldsymbol{\mu}_k = \frac{1}{|C_k|} \sum_{\mathbf{x}_t \in C_k} \mathbf{x}_t$ is the centroid of the k -th cluster. We set the number of clusters $K = r - m$. This design choice creates a one-to-one mapping between the discovered

functional clusters and the experts to be selected, providing a strong inductive bias towards preserving each distinct functional capability while minimizing redundancy. This process assigns a domain label to each token, defining a mapping function domain: $\{1, \dots, N_{\text{total}}\} \rightarrow \{1, \dots, K\}$.

Similarity Matrix Construction. For each pair of experts (i, j) in the candidate pool $\mathcal{E}_{\text{cand}}$, we compute a similarity score $S_{\text{comp}}(i, j)$. In this framework, an expert’s functional similarity is measured by its performance across diverse task types. We choose Spearman’s rank correlation coefficient as the core metric for task performance similarity because of three key advantages: First, as a nonparametric method, it requires no assumptions about the data distribution; second, its rank-based nature makes it inherently robust to outliers; and most importantly, it effectively captures the monotonic relationship of expert performance across different task domains, which aligns perfectly with our goal of identifying experts with the same functional profile. Therefore, this score is defined solely by the task performance similarity S_{perf} :

$$S_{\text{comp}}(i, j) = S_{\text{perf}}(i, j). \quad (6)$$

To compute S_{perf} , we first construct a performance vector $\mathbf{v}_i^{\text{perf}} \in \mathbb{R}^K$ for expert i . Its k -th element, $v_{i,k}^{\text{perf}}$, is the expert’s average reconstruction error in domain k (smaller values mean better performance). Because Spearman correlation is computed on ranks, this monotonic definition preserves functional similarity even though lower numbers indicate higher skill. Let $\mathcal{T}_k = \{t | \text{domain}(t) = k\}$ denote the set of token indices belonging to domain k . The error is then calculated as:

$$v_{i,k}^{\text{perf}} = \frac{1}{|\mathcal{T}_k|} \sum_{t \in \mathcal{T}_k} \|E_i(\mathbf{x}_t) - \mathbf{z}_{\text{real},t}\|_2^2. \quad (7)$$

Model	Method	Experts	ARC-c	ARC-e	BoolQ	HellaSwag	MMLU	OBQA	WinoGrande	Average
Mixtral	None	8	80.21	87.29	81.47	75.91	67.43	83.20	53.26	75.54
	Random	6	69.97	78.11	83.18	68.64	53.44	74.20	53.15	68.67
		4	44.54	47.22	64.86	59.83	48.21	67.20	55.72	55.37
	Frequency	6	70.90	78.15	78.28	56.85	58.47	78.20	58.25	68.44
		4	50.26	64.56	67.46	37.50	46.76	63.20	50.27	54.29
	Enu	6	73.04	83.16	84.89	67.49	61.14	81.00	58.41	72.73
		4	59.59	74.41	83.12	52.02	49.45	63.80	54.46	62.41
	MoP (Ours)	6	75.43	82.37	86.51	74.28	63.50	80.20	55.80	74.01
		4	69.62	79.84	83.33	68.85	54.26	72.00	54.85	68.96
	Qwen	None	60	54.18	64.56	78.90	58.99	37.44	61.00	48.38
Random		50	47.44	58.75	39.14	68.25	54.38	66.40	31.41	52.25
		40	38.57	50.63	1.92	48.20	36.37	38.20	40.88	36.40
Frequency		50	41.98	57.99	23.82	36.10	41.80	50.80	46.17	42.67
		40	30.29	36.19	9.05	59.18	47.81	59.60	49.32	41.63
Enu		50	54.61	66.16	33.03	62.89	53.69	65.20	50.19	55.11
		40	30.63	35.65	22.05	42.36	43.49	49.80	46.25	38.60
MoP (Ours)		50	60.70	75.04	38.10	67.28	55.60	66.80	50.19	59.10
		40	44.97	59.76	68.56	56.49	48.44	59.40	52.57	55.74

Table 1: Zero-shot performance of MoP compared with baseline pruning methods. MoP is compared with Enumeration(Enu), Random, and Frequency Pruning. Experts indicates the number of experts retained per layer. Best average scores are in **bold**.

Here, $E_i(\mathbf{x}_t)$ represents the output of the MoE layer only with expert i when given the input token \mathbf{x}_t , effectively simulating a scenario where only this single expert is activated. This performance vector $\mathbf{v}_i^{\text{perf}}$ serves as a functional profile for expert i , characterizing its effectiveness across the discovered domains. To measure the monotonic relationship, we compute the Spearman rank correlation coefficient. Let $\text{rank}(\cdot)$ be the rank transformation operator. For a given performance vector $\mathbf{v}_i^{\text{perf}} = [v_{i,1}, \dots, v_{i,K}]$, this operator sorts the elements and replaces each element with its ordinal rank (e.g., the smallest element is replaced by rank 1, the second smallest by rank 2, and so on). This converts the performance vector into a rank vector $\mathbf{R}_i = \text{rank}(\mathbf{v}_i^{\text{perf}})$. For clarity, we first define the sum of squared deviations for a centered rank vector as:

$$SS_i = \sum_{k=1}^K (R_{ik} - \bar{R}_i)^2, \quad (8)$$

where \bar{R}_i is the mean of the elements in the rank vector \mathbf{R}_i . Based on this, the formula for the Spearman correlation coefficient can be concisely expressed as:

$$\rho(\mathbf{R}_i, \mathbf{R}_j) = \frac{\sum_{k=1}^K (R_{ik} - \bar{R}_i)(R_{jk} - \bar{R}_j)}{\sqrt{SS_i \cdot SS_j}}. \quad (9)$$

Finally, we normalize the ρ value to the interval $[0, 1]$ to obtain the final performance similarity score:

$$S_{\text{perf}}(i, j) = \frac{1}{2} (1 + \rho(\mathbf{R}_i, \mathbf{R}_j)). \quad (10)$$

Agglomerative Hierarchical Clustering. After obtaining the similarity matrix S_{comp} , we convert it into a distance matrix \mathbf{D} , where the distance between any two experts i and j is defined as $D(i, j) = 1 - S_{\text{comp}}(i, j)$. This distance matrix conceptually defines the space in which the clustering operates.

We then apply hierarchical clustering, a bottom-up approach that iteratively merges the closest pair of clusters. The ‘‘closeness’’ of two clusters is determined by a linkage criterion. We employ Ward’s linkage method (Ward Jr 1963), which defines the cost of merging two clusters, C_a and C_b , as the increase in the total Error Sum of Squares (ESS) that would result from their merger. This merge cost, denoted as $\Delta(\text{ESS})$, serves as the specific inter-cluster distance measure for Ward’s method and is calculated as:

$$\Delta(\text{ESS}(C_a, C_b)) = \frac{|C_a||C_b|}{|C_a| + |C_b|} \|\boldsymbol{\mu}_a - \boldsymbol{\mu}_b\|_2^2. \quad (11)$$

Here, $\boldsymbol{\mu}_a$ and $\boldsymbol{\mu}_b$ are the centroids of the clusters, defined as the mean of the performance vectors (\mathbf{v}^{perf}) of the experts within each cluster:

$$\boldsymbol{\mu}_a = \frac{1}{|C_a|} \sum_{i \in C_a} \mathbf{v}_i^{\text{perf}}. \quad (12)$$

At each step, the algorithm merges the pair of clusters that minimizes this merge cost. This process continues until we are left with $r - m$ clusters, resulting in a partition of $\mathcal{E}_{\text{cand}}$, $\mathcal{G} = \{G_1, G_2, \dots, G_{r-m}\}$.

Intra-Cluster Representative Selection. The final step is to select a representative from each functional cluster. We use the Activation Variability Score (S_{var}) as the selection

Model	Method	Experts	ARC-c	ARC-e	BoolQ	HellaSwag	MMLU	OBQA	WinoGrande	Average
Mixtral	Enu(C4)	6	73.04	83.16	84.89	67.49	61.14	81.00	58.41	72.73
	Enu(Mixed)	6	74.57	85.35	84.06	71.55	60.91	79.80	53.28	72.79
	GVP	6	76.62	84.09	85.20	73.79	63.43	81.00	50.35	73.50
	MoP (Ours)	6	75.43	82.37	86.51	74.28	63.50	80.20	55.80	74.01
Qwen	Enu(C4)	50	54.61	66.16	33.03	62.89	53.69	65.20	50.19	55.11
	Enu(Mixed)	50	41.81	49.45	15.84	66.27	54.84	68.20	49.88	49.47
	GVP	50	54.52	65.57	33.42	63.83	54.56	67.40	49.64	55.56
	MoP (Ours)	50	60.70	75.04	38.10	67.28	55.60	66.80	50.19	59.10

Table 2: Zero-shot performance ablation study of different pruning methods (some evaluated on different calibration dataset). All methods retain the same number of experts. Best average scores are in **bold**.

criterion to identify the most specialized expert within each group. For each cluster G_k , the representative expert e_k^* is chosen as:

$$e_k^* = \arg \max_{i \in G_k} S_{\text{var}}(i). \quad (13)$$

The resulting set of representatives, \mathcal{E}_{div} , forms a group of experts that are both individually specialized and functionally complementary.

Experiments

Experimental Setup

Models and Benchmarks. Our experiments are conducted on two representative open-source MoE models: Mixtral-8x7B-Instruct (47B parameters, 8 experts/layer) and Qwen1.5-MoE-A2.7B-Chat (14.3B parameters, 60 experts/layer). This significant difference in expert count allows us to evaluate the robustness and scalability of our method. We evaluate performance on a wide range of benchmarks. For general capability assessment, we use seven standard benchmarks: ARC-c/e (Clark et al. 2018), BoolQ (Clark et al. 2019), HellaSwag (Zellers et al. 2019), MMLU (Hendrycks et al. 2020), OpenBookQA (Mihaylov et al. 2018), and WinoGrande (Sakaguchi et al. 2021). For specialized skill assessment, we use four benchmarks across two domains: mathematical reasoning (GSM8K (Cobbe et al. 2021), MATH (Hendrycks et al. 2021)) and code generation (HumanEval (Chen et al. 2021), MBPP (Austin et al. 2021)). Finally, for expert diversity validation (Section), we utilize test samples from six distinct domains: Math-GSM8K (Cobbe et al. 2021), Code-HumanEval (Chen et al. 2021), Commonsense-PIQA (Bisk et al. 2020), Summarization-XSum (Narayan, Cohen, and Lapata 2018), Science-ARC-C (Clark et al. 2018), and Logic-MNLI (Williams, Nangia, and Bowman 2017).

Implementation Details. Our main baseline is Enumeration Pruning (Lu et al. 2024). Following its original implementation, this method uses a calibration dataset sourced from the general-purpose C4 corpus (Raffel et al. 2020). In contrast, to improve functional clustering, our GVP and MoP methods use a mixed-diversity calibration dataset as described in (Section). All experiments were conducted on a server equipped with four NVIDIA A100 80G GPUs.

Main Results

Performance of Mosaic Pruning on General Tasks. We compare MoP against three baseline methods: Enumeration Pruning (Lu et al. 2024), Random Pruning (which randomly discards experts), and Frequency Pruning (Muzio, Sun, and He 2024) (which removes the least frequently activated experts in each layer). As shown in Table 1, MoP consistently and significantly outperforms all baseline methods in all pruning rates and models, achieving an average performance improvement of 7.24% over Enumeration Pruning. These results demonstrate that preserving functional diversity enables more robust model compression.

Inference Speed and Memory Usage. As a structured pruning method, MoP’s advantages are directly reflected in deployment efficiency. We evaluated the inference speed and peak memory usage of the models before and after pruning. As shown in Table 3, removing entire expert modules resulted in significant memory savings. Concurrently, we observed a significant acceleration in inference, achieving an average speedup of 1.20x, due to the reduced number of experts and potential inter-device communication overhead.

Model	Method	Experts	Mem	Speedup
Mixtral	None	8	87.87	1.00x
	MoP	6	66.87	1.22x
		4	45.87	1.52x
Qwen	None	60	27.79	1.00x
	MoP	50	23.91	1.12x
		40	20.04	1.36x

Table 3: Inference speed and memory(GB) usage of MoP-pruned models compared to the original architectures.

Ablation Study

To clarify the contribution of each component, we conducted ablation studies on Enumeration Pruning, GVP, and MoP. As shown in Table 2, from Enumeration Pruning to GVP to MoP, performance shows progressive improvements. Since MoP and Enumeration Pruning used different calibration datasets, to ensure fair comparison, we re-experimented using the same mixed-diversity calibration dataset, comparing

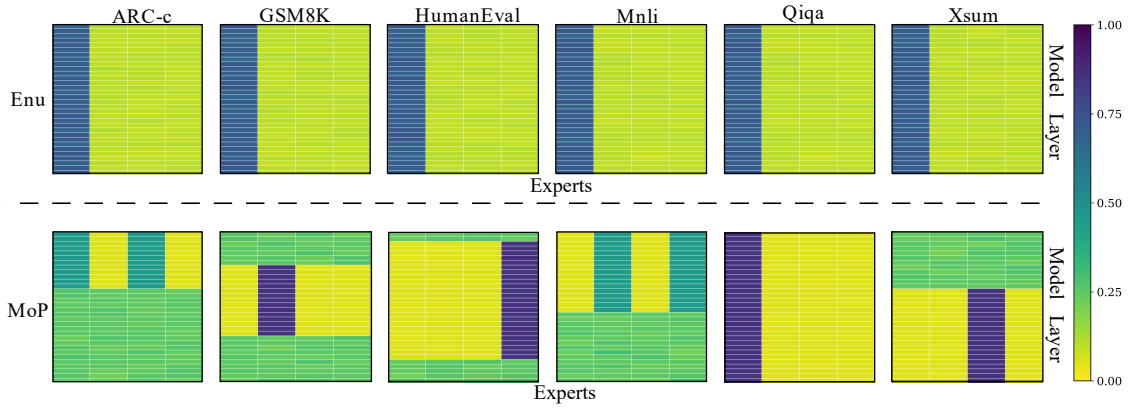


Figure 3: Heatmaps of expert activation weights across different domains for Mixtral-8x7B pruned to 4 experts. (**Top Row**): The Enumeration Pruning leads to functional homogenization, with a few generalist experts dominating all tasks. (**Bottom Row**): Our MoP method preserves domain specialization, with different experts activating for different tasks.

Model	Method	Experts	GSM8K	MATH
Mixtral	Enu	6	28.73	9.00
		4	22.37	3.80
	MoP	6	51.33	9.20
		4	28.35	3.00
Qwen	Enu	50	34.34	7.39
		40	16.76	5.40
	MoP	50	35.63	11.20
		40	21.91	7.20

Table 4: Zero-shot performance comparison on mathematical reasoning benchmarks at various pruning rates.

MoP with Enumeration Pruning (Mixed). The results show that even with the same calibration dataset, MoP still outperforms Enumeration Pruning (Mixed) across all benchmarks, demonstrating the superiority of our pruning framework.

Validation of Expert Diversity and Specialization

Analysis of Expert Activation Diversity. To examine whether the pruned experts still cover a wide range of functional areas, we conducted activation pattern analysis on MoP and Enumeration Pruning. We used test samples from six different domains as described in (Section), recorded the activation weights of the four experts pruned from Mixtral-8x7B across all 32 MoE layers, calculated average values to generate 32x4 matrices, and visualized them as heatmaps in Figure3. The results show that MoP maintains clear domain specialization, with different tasks activating different experts; while Enumeration Pruning leads to functional homogenization, where generalist experts dominate across all domains.

Performance Validation on Specialized Tasks. To quantitatively confirm the observations from our diversity analysis, we conducted a direct comparison between MoP and Enumeration Pruning methods on highly specialized down-

stream tasks. We evaluated the models in two distinct domains: mathematical reasoning and code generation. As shown in Table 4 and Table 5, The results provide clear evidence of functional collapse in existing methods: the MoP-pruned model far outperforms Enumeration Pruning across the four specialized benchmarks, with an average performance improvement of 8.92%. By preserving functional diversity and specialization, our method achieves a high compression rate while avoiding this functional collapse and maintaining the model’s complex reasoning abilities in various specialized domains to the greatest extent possible.

Model	Method	Experts	HumanEval	MBPP
Mixtral	Enu	6	82.31	14.00
		4	1.22	1.55
	MoP	6	83.54	15.95
		4	71.34	10.89
Qwen	Enu	50	1.21	8.17
		40	1.83	1.95
	MoP	50	14.02	9.34
		40	3.05	8.17

Table 5: Zero-shot performance comparison on code generation benchmarks at various pruning rates.

Conclusion

In this paper, we address the poor generalization of existing MoE pruning methods, which discard specialist experts, by proposing Mosaic Pruning (MoP). MoP is a hierarchical “cluster-then-select” framework that preserves a functionally comprehensive set of experts. Extensive experiments show our method outperforms existing baselines, particularly on specialized tasks. Our work establishes a “prune once, deploy efficiently across diverse tasks” paradigm, enhancing the practical applicability of large-scale MoE models.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (Grant Nos. 62572389, 72293581, 72274152, 62402376). We also acknowledge the support from Xi'an Jiaotong University and the Institute of Artificial Intelligence of China Telecom (TeleAI).

References

- Austin, J.; Odena, A.; Nye, M.; Bosma, M.; Michalewski, H.; Dohan, D.; Jiang, E.; Cai, C.; Terry, M.; Le, Q.; et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Bisk, Y.; Zellers, R.; Gao, J.; and Choi, Y. 2020. PIQA: Reasoning about Physical Commonsense in Natural Language. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 7432–7439. AAAI.
- Chen, M.; Tworek, J.; Jun, H.; Yuan, Q.; Pinto, H. P. D. O.; Kaplan, J.; Edwards, H.; Burda, Y.; Joseph, N.; Brockman, G.; et al. 2021. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*.
- Chi, Z.; Dong, L.; Huang, S.; Dai, D.; Ma, S.; Patra, B.; Singhal, S.; Bajaj, P.; Song, X.; Mao, X.-L.; et al. 2022. On the representation collapse of sparse mixture of experts. *Advances in Neural Information Processing Systems*, 35: 34600–34613.
- Clark, C.; Lee, K.; Chang, M.-W.; Kwiatkowski, T.; Collins, M.; and Toutanova, K. 2019. Boolq: Exploring the surprising difficulty of natural yes/no questions. *arXiv preprint arXiv:1905.10044*.
- 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*.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Fedus, W.; Zoph, B.; and Shazeer, N. 2022. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120): 1–39.
- Frantar, E.; and Alistarh, D. 2023. Sparsegpt: Massive language models can be accurately pruned in one-shot. In *International conference on machine learning*, 10323–10337. PMLR.
- Han, S.; Mao, H.; and Dally, W. J. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. *arXiv preprint arXiv:1510.00149*.
- He, Z.; Wang, Z.; Liu, X.; Liu, S.; Yao, Y.; Huang, Y.; Li, X.; Li, Y.; Che, Z.; Zhang, Z.; Wang, Y.; Wang, X.; Pu, L.; Xu, H.; Fang, R.; Zhao, Y.; Zhang, J.; Huang, X.; Lu, Z.; Peng, J.; Zheng, W.; Wang, S.; Yang, B.; he, X.; Jiang, Z.; Xie, Q.; Zhang, Y.; Li, Z.; Shi, L.; Fu, W.; Zhang, Y.; Huang, Z.; Xiong, S.; Zhang, Y.; Wang, C.; and Song, S. 2024. TeleChat Technical Report. arXiv:2401.03804.
- Hendrycks, D.; Burns, C.; Basart, S.; Zou, A.; Mazeika, M.; Song, D.; and Steinhardt, J. 2020. Measuring massive multitask language understanding. *arXiv preprint arXiv:2009.03300*.
- 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*.
- Jacobs, R. A.; Jordan, M. I.; Nowlan, S. J.; and Hinton, G. E. 1991. Adaptive mixtures of local experts. *Neural computation*, 3(1): 79–87.
- Jiang, A. Q.; Sablayrolles, A.; Roux, A.; et al. 2024. Mixtral of Experts. arXiv:2401.04088.
- Kim, Y. J.; Awan, A. A.; Muzio, A.; Salinas, A. F. C.; Lu, L.; Hendy, A.; Rajbhandari, S.; He, Y.; and Awadalla, H. H. 2021. Scalable and efficient moe training for multitask multilingual models. *arXiv preprint arXiv:2109.10465*.
- Kullback, S.; and Leibler, R. A. 1951. On information and sufficiency. *The annals of mathematical statistics*, 22(1): 79–86.
- Lepikhin, D.; Lee, H.; Xu, Y.; Chen, D.; Firat, O.; Huang, Y.; Krikun, M.; Shazeer, N.; and Chen, Z. 2020. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668*.
- Li, X.; Yao, Y.; Jiang, X.; Fang, X.; Wang, C.; Liu, X.; Wang, Z.; Zhao, Y.; Wang, X.; Huang, Y.; Song, S.; Li, Y.; Zhang, Z.; Zhao, B.; Sun, A.; Wang, Y.; He, Z.; Wang, Z.; Li, X.; and Huang, T. 2024a. 52B to 1T: Lessons Learned via Tele-FLM Series. arXiv:2407.02783.
- Li, X.; Yao, Y.; Jiang, X.; Fang, X.; Wang, C.; Liu, X.; Wang, Z.; Zhao, Y.; Wang, X.; Huang, Y.; Song, S.; Li, Y.; Zhang, Z.; Zhao, B.; Sun, A.; Wang, Y.; He, Z.; Wang, Z.; Li, X.; and Huang, T. 2024b. Tele-FLM Technical Report. arXiv:2404.16645.
- Lu, X.; Liu, Q.; Xu, Y.; Zhou, A.; Huang, S.; Zhang, B.; Yan, J.; and Li, H. 2024. Not all experts are equal: Efficient expert pruning and skipping for mixture-of-experts large language models. *arXiv preprint arXiv:2402.14800*.
- Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; Agarwal, S.; et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 1(3): 3.
- McQueen, J. B. 1967. Some methods of classification and analysis of multivariate observations. In *Proc. of 5th Berkeley Symposium on Math. Stat. and Prob.*, 281–297.
- Mihaylov, T.; Clark, P.; Khot, T.; and Sabharwal, A. 2018. Can a suit of armor conduct electricity? a new dataset for open book question answering. *arXiv preprint arXiv:1809.02789*.
- Muzio, A.; Sun, A.; and He, C. 2024. Seer-moe: Sparse expert efficiency through regularization for mixture-of-experts. *arXiv preprint arXiv:2404.05089*.
- Narayan, S.; Cohen, S. B.; and Lapata, M. 2018. Don't give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. *arXiv preprint arXiv:1808.08745*.

- OpenAI. 2024. GPT-4 Technical Report. arXiv:2303.08774.
- Raffel, C.; Shazeer, N.; Roberts, A.; Lee, K.; Narang, S.; Matena, M.; Zhou, Y.; Li, W.; and Liu, P. J. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140): 1–67.
- Rajbhandari, S.; Li, C.; Yao, Z.; Zhang, M.; Aminabadi, R. Y.; Awan, A. A.; Rasley, J.; and He, Y. 2022. Deepspeed-moe: Advancing mixture-of-experts inference and training to power next-generation ai scale. In *International conference on machine learning*, 18332–18346. PMLR.
- Sakaguchi, K.; Bras, R. L.; Bhagavatula, C.; and Choi, Y. 2021. Winogrande: An adversarial winograd schema challenge at scale. *Communications of the ACM*, 64(9): 99–106.
- Shao, J.; and Li, X. 2025. AI Flow at the Network Edge. *IEEE Network*, 1–1.
- Shazeer, N.; Mirhoseini, A.; Maziarz, K.; Davis, A.; Le, Q.; Hinton, G.; and Dean, J. 2017. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. *arXiv preprint arXiv:1701.06538*.
- Spearman, C. 1987. The proof and measurement of association between two things. *The American journal of psychology*, 100(3/4): 441–471.
- Sun, M.; Liu, Z.; Bair, A.; and Kolter, J. Z. 2023. A simple and effective pruning approach for large language models. *arXiv preprint arXiv:2306.11695*.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023a. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Touvron, H.; Martin, L.; Stone, K.; et al. 2023b. Llama 2: Open Foundation and Fine-Tuned Chat Models. arXiv:2307.09288.
- Wang, Z.; Liu, X.; Liu, S.; Yao, Y.; Huang, Y.; Li, M.; He, Z.; Li, Y.; Pu, L.; Xu, H.; Wang, C.; and Song, S. 2024. TeleChat: An Open-source Bilingual Large Language Model. In Wong, K.-F.; Zhang, M.; Xu, R.; Li, J.; Wei, Z.; Gui, L.; Liang, B.; and Zhao, R., eds., *Proceedings of the 10th SIGHAN Workshop on Chinese Language Processing (SIGHAN-10)*, 10–20. Bangkok, Thailand: Association for Computational Linguistics.
- Wang, Z.; Liu, X.; Yao, Y.; Wang, C.; Zhao, Y.; Yang, Z.; Deng, W.; Jia, K.; Peng, J.; Huang, Y.; Xiong, S.; Jiang, Z.; Yu, K.; Hu, X.; Yao, F.; Fang, R.; Jiang, Z.; Song, R.; Xie, Q.; Xue, R.; He, X.; Xue, Y.; Yuan, Z.; Zhang, Z.; Huang, Z.; Wang, S.; Wang, X.; Wu, H.; Wang, M.; Zhan, X.; Sun, Y.; Xing, Z.; Jiang, Y.; Yang, B.; Song, S.; Li, Y.; He, Z.; and Li, X. 2025. Technical Report of TeleChat2, TeleChat2.5 and T1. arXiv:2507.18013.
- Ward Jr, J. H. 1963. Hierarchical grouping to optimize an objective function. *Journal of the American statistical association*, 58(301): 236–244.
- Williams, A.; Nangia, N.; and Bowman, S. R. 2017. A broad-coverage challenge corpus for sentence understanding through inference. *arXiv preprint arXiv:1704.05426*.
- Xiong, S.; Zhao, Y.; Zhang, J.; Mengxiang, L.; He, Z.; Li, X.; and Song, S. 2024. Dual Prompt Tuning based Contrastive Learning for Hierarchical Text Classification. In Ku, L.-W.; Martins, A.; and Srikumar, V., eds., *Findings of the Association for Computational Linguistics: ACL 2024*, 12146–12158. Bangkok, Thailand: Association for Computational Linguistics.
- Yang, A.; Yang, B.; Hui, B.; Zheng, B.; Yu, B.; Zhou, C.; Li, C.; Li, C.; Liu, D.; Huang, F.; Dong, G.; Wei, H.; Lin, H.; Tang, J.; and Wang, J. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671*.
- Yu, W.; Zhu, C.; Qin, L.; Zhang, Z.; Zhao, T.; and Jiang, M. 2022. Diversifying content generation for commonsense reasoning with mixture of knowledge graph experts. *arXiv preprint arXiv:2203.07285*.
- Zellers, R.; Holtzman, A.; Bisk, Y.; Farhadi, A.; and Choi, Y. 2019. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830*.
- Zhao, M.; Hu, W.; Wang, J.; Lai, X.; Huang, T.; Min, Y.; Yan, R.; and Zhu, X. 2025. Making Every Head Count: Sparse Attention Without the Speed-Performance Trade-off. arXiv:2511.09596.