

Neuronal Redundancy Reduction and Collaborative Learning for Parameter-Efficient Mixture-of-Experts

Cheng Ding¹, Qin Chen^{1*}, Jie Zhou¹, Hongjian Zhan², Liang He¹

¹School of Computer Science and Technology, East China Normal University,

²School of Information and Electronic Engineering, East China Normal University

cding@stu.ecnu.edu.cn, hjzhan@cee.ecnu.edu.cn,

{qchen, jzhou, lhe}@cs.ecnu.edu.cn

Abstract

Mixture of Experts (MoE) has great potential in scaling up the capacity of large models while maintaining low computational costs. However, existing methods mainly reduce redundancy at the expert level while overlooking fine-grained neuron-level redundancy and expert collaboration. In this paper, we propose an Information Bottleneck based MoE (IBMoE) for parameter-efficient fine-tuning, which reduces neuron-level redundancy within each expert and fosters internal collaboration among all experts. Specifically, IBMoE equips each expert with learnable queries to capture task-specific characteristics and employs sparse neuronal activation to dynamically select the most relevant neurons, reducing redundant activations across tasks. In addition, a diversity constraint is imposed among experts to maximize their knowledge diversity, enabling them to cooperate more effectively. Extensive experiments across multiple benchmarks show that IBMoE consistently outperforms strong baselines while activating significantly fewer parameters. In addition, IBMoE reduces inference latency by 61.7% and the memory footprint of adapters by 48.2% compared with previous MoE-based adaptation methods, while achieving comparable training costs.

1 Introduction

Large language models (LLMs) (OpenAI, 2023; Touvron et al., 2023a; Guo et al., 2025) have achieved strong performance across a wide range of Natural Language Processing (NLP) tasks (Bai et al., 2025; Aggarwal et al., 2026; Tian et al., 2026). However, adapting LLMs to diverse downstream tasks often requires parameter updates, leading to substantial training and inference costs. This motivates parameter-efficient fine-tuning (PEFT) methods, which freeze backbone models and optimize a small set of trainable parameters.

* Corresponding author.




Figure 1: The differences between classical MoE and IBMoE. MoE focuses on expert-level specialization, while IBMoE enhances fine-grained, neuron-level specialization within each expert through sparse neuron activation and neuron collaboration.

Early PEFT approaches such as prompt tuning (Wang et al., 2025), prefix tuning (Zhao et al., 2025), and P-tuning (Liu et al., 2021) introduce learnable vectors into model inputs or intermediate layers. While effective in some settings, they often underperform on smaller models or more complex tasks. Low-Rank Adaptation (LoRA) (Hu et al., 2021) improves efficiency by learning low-rank updates to model weights, achieving a better trade-off between performance and cost.

Beyond PEFT, mixture-of-experts (MoE) offers a promising conditional computation paradigm for balancing capacity and efficiency. Instead of activating all parameters for every input, MoE routes tokens to a subset of specialized experts, enabling scalable models with constant per-token computation (Bengio et al., 2015; Lepikhin et al., 2020; Zadouri et al., 2023; Gao et al., 2024; Zhu et al., 2025). Recent efforts have combined MoE with PEFT (Li et al., 2025), MoCE (Gao et al., 2025) introduces a dual-level MoE framework for causal intervention, and HMoE (Lin et al., 2026) uses hierarchical expert routing for code search. However, such methods often encounter challenges related to knowledge redundancy and parameter inefficiency. In addition, some studies have shown that MoE architectures exhibit more pronounced redundancy

than dense models (He et al., 2024), while directly applying one-shot static pruning to MoE models can lead to considerable performance degradation due to their input-dependent expert activation (Xie et al., 2024). Consequently, addressing knowledge redundancy in MoE has become a critical challenge, and prior research has generally pursued two main lines of approach.

On the one hand, a line of work focuses on encouraging expert specialization. Some approaches design task-aware routing strategies to guide expert selection (Liu et al., 2024; Zhang et al., 2024). Others partition dense models into experts or construct parameter-isolated experts to improve specialization (Zuo et al., 2022; Zhu et al., 2024b). However, these methods mainly focus on improving specialization at the expert level, while redundancy within each expert remains largely unexplored.

On the other hand, enabling all experts to collaborate efficiently remains an important challenge (Cai et al., 2024). Most methods rely on routing mechanisms to assign tokens to multiple experts, often combined with balance constraints to stabilize utilization (Zhu et al., 2023). Some recent studies further reduce representational similarity across experts to mitigate collapse (Dai et al., 2024; Do et al., 2025). However, these methods still operate at the level of expert routing, without explicitly modeling interactions within experts or the underlying differences in their internal representations. As a result, the notion of expert knowledge diversity remains implicit and under-constrained.

To address these limitations, we propose IBMoE, a parameter-efficient MoE framework inspired by the Information Bottleneck principle (Tishby et al., 2000; Alemi et al., 2016). IBMoE aims to retain only task-relevant information while reducing redundant neuron activations. Specifically, each expert is equipped with learnable queries that induce sparse neuron activation under an IB constraint, enabling fine-grained control over activated neurons. In addition, we introduce a neuron-level collaboration module that explicitly encourages complementary specialization across experts, improving knowledge diversity and coordination. As shown in Figure 1, IBMoE shifts the focus from expert-level routing to neuron-level sparse activation and collaboration, enabling more efficient utilization of model capacity. The major contributions of our work can be summarized as follows:

- We propose IBMoE, an information

bottleneck-based MoE framework that reduces neuron-level redundancy and enables fine-grained expert collaboration.

- We introduce a sparse neuron activation mechanism driven by learnable queries to selectively retain task-relevant information. Moreover, a neuron collaboration module is introduced to explicitly guide all experts to cooperate at the neuron-level with higher efficiency.
- We conduct extensive experiments to validate the effectiveness and superiority of our method. The experiments indicate that IB-MoE achieves the best average performance, while reducing the inference time by 61.7% and the memory footprint of adapters by 48.2% compared with existing MoE-based PEFT methods.

2 Related Work

2.1 Parameter-Efficient Fine-Tuning

Parameter-efficient fine-tuning (PEFT) aims to adapt large language models to downstream tasks while updating only a small number of parameters. Existing PEFT methods can generally be divided into prompt-based and adapter-based approaches. Prompt-based methods optimize learnable prompts or virtual tokens while keeping the backbone frozen. Representative methods include Prompt-Tuning (Lester et al., 2021), P-Tuning (Liu et al., 2021), and Prefix-Tuning (Li and Liang, 2021). BitFit (Zaken et al., 2021) further reduces trainable parameters by updating only bias terms. Adapter-based methods instead introduce trainable modules into the backbone architecture. IA3 (Liu et al., 2022) scales intermediate activations using learnable vectors, while LoRA (Hu et al., 2021) performs low-rank parameter updates through decomposed matrices. Subsequent works further improve efficiency and expressiveness through alternative adapter designs (Braga et al., 2024; Ren et al., 2024). Our method follows the PEFT paradigm and combines it with sparse MoE modeling for efficient adaptation.

2.2 Mixture of Experts

Mixture-of-Experts (MoE) enables conditional computation by routing inputs to a subset of specialized experts. Existing MoE studies mainly focus on three directions. The first direction scales

model capacity while maintaining constant per-token computation. Representative approaches such as GShard (Lepikhin et al., 2020), Switch Transformer (Fedus et al., 2022), and DeepSeek-MoE (Dai et al., 2024) improve scalability through sparse expert routing and fine-grained expert designs. The second direction reduces computational cost through knowledge separation or expert decomposition. For example, Llama-MoE (Zhu et al., 2024b) and MoeBERT (Zuo et al., 2022) partition dense model parameters into multiple experts to improve efficiency. The third direction integrates MoE with PEFT adapters. MoLoRA (Zadouri et al., 2023) decomposes LoRA into multiple experts, while MoLA (Gao et al., 2024), MoRIE (Liu et al., 2024), and HMoRA (Liao et al., 2025) further improve expert routing and specialization. However, existing methods mainly focus on expert-level routing and specialization, while neuron-level redundancy and explicit modeling of expert diversity remain underexplored.

2.3 Information Bottleneck

Information Bottleneck (IB) aims to compress representations while preserving task-relevant information (Tishby et al., 2000). Prior studies have shown that IB-based representation learning can improve generalization by suppressing irrelevant information in hidden representations (Kawaguchi et al., 2023). To improve the optimization of conventional IB objectives in deep neural networks, several variants such as DeepIB (Tishby and Zaslavsky, 2015) and DIB (Strouse and Schwab, 2017) have been proposed. IB has been successfully applied to graph learning (Yu et al., 2024; Zhang et al., 2025), adversarial training (Chen et al., 2022), model compression (Lorenzen et al., 2021), and representation compression (Sheng et al., 2023). More recently, Zhu et al. (2024a) introduced IB to suppress noisy retrieved information in LLM-based generation. Different from prior IB applications, our method incorporates IB into MoE-based PEFT to reduce neuron-level redundancy and improve expert specialization.

3 Preliminaries

Mixture of Experts: In general, the MoE architecture consists of n feed-forward neural networks, denoted as $\{E\}_{i=1}^n$, which process various inputs through a gating mechanism. Specifically, the gating mechanism is a function equipped with the

trainable parameters $W_r \in \mathbb{R}^{d_{in} \times n}$. For each input $x \in \mathbb{R}^{d_{in} \times 1}$, the gate function maps it to an n -dimensional discrete probability vector, which represents the contribution of each expert. Then, a routing algorithm, such as the top-k or soft routing method, is used to allocate the input to the experts. In this paper, we select the top-k experts with the highest contributions to process the input. The output h_i at i -th token is calculated as:

$$G_i = \text{TopK}(\text{softmax}(W_r^\top x_i)) \quad (1)$$

$$h_i = \sum_{e=1}^N G_{i,e} E_e(x_i) \quad (2)$$

where $G_{i,e}$ represents the contribution of e -th expert at i -th token.

Low-Rank Adaptation: Inspired by the natural low intrinsic dimension, LoRA introduces low-rank decomposition for optimizing the weight matrix of pre-trained LLMs to reduce the number of training parameters. Let the weight matrix of a pre-trained linear layer be $W_0 \in \mathbb{R}^{d_{in} \times d_{out}}$. LoRA learns two low-rank matrices $A \in \mathbb{R}^{r \times d_{in}}$ and $B \in \mathbb{R}^{d_{out} \times r}$ to update W_0 , where $r \ll \min(d_{in}, d_{out})$. During training, W_0 is frozen and only matrices A and B are updated. Given an input x_i , the output of the LoRA layer can be expressed as:

$$\begin{aligned} h_i &= W_0^\top x_i + \frac{\alpha}{r} B A x_i \\ &= W_0^\top x_i + \Delta W x_i \end{aligned} \quad (3)$$

where α is a scale hyperparameter.

MoE and LoRA: Some studies integrate the advantages of LoRA and MoE by constructing multiple LoRA adapters as experts, known as LoRA-MoE, to achieve expert specialization on different tasks. Combining Equations 2 and 3, we have:

$$h_i = W_0^\top x_i + \sum_{e=1}^n G_{i,e} \Delta W_e x_i \quad (4)$$

where ΔW_e is the weight matrix of e -th expert.

4 Methodology

In this section, we describe the details of the IB-MoE framework with Sparse Neuron Activation (SNA) and Neuron Collaboration (NC), as shown in Figure 2. Specifically, SNA integrates the IB principle into the MoE framework to capture each expert’s knowledge characteristics. This integration reduces redundant information at the neuron




Figure 2: The overall framework of IBMoE consists of two components: Sparse Neuron Activation (SNA) and Neuron Collaboration (NC). SNA utilizes the IB to learn a neuron-level query vector for discarding irrelevant details. The NC strategy is then introduced to reduce redundant knowledge by modeling interactions among experts.

level and enhances expert specialization. In turn, we introduce a NC mechanism that maximizes knowledge diversity across experts, encouraging collaboration and improving parameter utilization.

4.1 Sparse Neuron Activation

Typically, the knowledge characteristics of experts are implicitly determined by the routing mechanism, which makes it difficult to analyze and control redundant knowledge. Motivated by the IB theory, we introduce the SNA module, which explicitly filters out task-irrelevant signals and preserves task-specific information for each expert. At a finer granularity, not all neurons within an expert contribute equally to a task. Therefore, SNA activates only the most relevant neurons, analogous to sparse expert activation at the expert level. In this way, SNA captures task-specific characteristics by modeling differences in neuron activation distributions across experts and enhances fine-grained specialization. Specifically, let the inputs and labels be denoted as X and Y , respectively, and let $\hat{H}$ represent the compressed representation of X . Each expert is trained to minimize the mutual information $I(X; \hat{H})$ while maximizing $I(\hat{H}; Y)$, so as to discard irrelevant information and preserve

task-relevant knowledge:

$$\begin{aligned} L_{ib} &= \beta I(X; \hat{H}) - I(\hat{H}; Y) \\ &= \mathbb{E}_{(x, \hat{h}) \sim p(x, \hat{h})} \beta \log \frac{p(\hat{h}|x)}{p(\hat{h})} \\ &\quad - \mathbb{E}_{(\hat{h}, y) \sim p(\hat{h}, y)} \log \frac{p(y|\hat{h})}{p(y)} \end{aligned} \quad (5)$$

where β is a hyperparameter that used to control the compression of X .

However, computing $p(\hat{h})$ and $p(y|\hat{h})$ directly is intractable. To solve this, with the concept of Variational Information Bottleneck (VIB) (Alemi et al., 2016), we use two variational distributions $r(\hat{h})$ and $q(y|\hat{h})$ to approximate $p(\hat{h})$ and $p(y|\hat{h})$, which allows us to optimize the upper bound of L_{ib} . For the former term in Equation 5, we have:

$$\begin{aligned} I(X; \hat{H}) &= \mathbb{E}_{(x, \hat{h}) \sim p(x, \hat{h})} \log \frac{p(\hat{h}|x)}{r(\hat{h})} - \mathbb{E}_{\hat{h} \sim p(\hat{h})} \log \frac{p(\hat{h})}{r(\hat{h})} \\ &= \mathbb{E}_{x \sim p(x)} \text{KL}(p(\hat{h}|x) || r(\hat{h})) - \text{KL}(p(\hat{h}) || r(\hat{h})) \end{aligned} \quad (6)$$

By non-negativity of KL divergence, $\text{KL}(p(\hat{h}) || r(\hat{h})) \geq 0$, hence

$$I(X; \hat{H}) \leq \mathbb{E}_{x \sim p(x)} \text{KL}(p(\hat{h} | x) || r(\hat{h})) \quad (7)$$

Similarly, for the latter term in Equation 5 and based on the fact that $\text{KL}(p(y|\hat{h}) || q(y|\hat{h})) \geq 0$, the lower bound is calculated as:

$$\begin{aligned} \mathbb{E}_{(\hat{h}, y) \sim p(\hat{h}, y)} \log \frac{p(y|\hat{h})}{p(y)} &\geq \mathbb{E}_{(\hat{h}, y) \sim p(\hat{h}, y)} \log \frac{q(y|\hat{h})}{p(y)} \\ I(\hat{H}; Y) &\geq \mathbb{E}_{(\hat{h}, y) \sim p(\hat{h}, y)} \log q(y|\hat{h}) + H(Y) \end{aligned} \quad (8)$$

where $H(Y)$ is the entropy of Y , which is independent of the optimization and can be ignored.

Then, for optimizing Equation 7, we first add a task-specific query $V_i \in \mathbb{R}^{d_{out}}$ for E_i to represent the importance of the neurons in ΔW_i . Before the forward process, E_i uses V_i to generate a binary mask $M_i = [m_1, m_2, \dots, m_{d_{out}}]^\top \in \{0, 1\}^{d_{out}}$ via a Bernoulli-based importance sampling, where M_i indicates which neurons should be activated. M_i is generated as follows:

$$M_i = \text{Ber} \left(\frac{V_i - \min(V_i) + \epsilon}{\max(V_i) - \min(V_i) + \epsilon} \right) = \text{Ber}(\hat{V}) \quad (9)$$

where $\text{Ber}(\cdot)$ is a function used for Bernoulli sampling, and ϵ is small positive number. Note that the ϵ in the denominator prevents division by zero errors, while ϵ in the numerator ensures that each neuron becomes activated when all neurons have equal importance. The compressed hidden representation can be expressed as:

$$\hat{h}_i = W_0^\top x_i + \sum_{e=1}^n \frac{d_{out} G_{i,e} \Delta \hat{W}_e x_i}{\text{sum}(M_e)} \quad (10)$$

where $\Delta \hat{W}_e = M_e \odot \Delta W_e$, $\odot$ denotes element-wise multiplication at the column level, and $\text{sum}(\cdot)$ is a summation function. We provide the detailed implementation and annotations in Appendix A.

In this paper, Bernoulli sampling is used as the prior distribution, denoted as $r(\hat{h}) = \text{Ber}(\pi)$. We aim for the expected number of activated neurons to follow a Bernoulli distribution with a probability of π , which enables controllable compression of invalid information. Compared with h , the computation of $\hat{h}$ additionally depends on the query V of each expert, where V is independent of x . Consequently, we have $p(\hat{h}|x) = p(M|V)$ and $r(\hat{h}) = r(M)$. Therefore, a sparsity constraint is designed to minimize $\text{KL}(p(\hat{h}|x) || r(\hat{h}))$:

$$L_s = \sum_{i=1}^N \left[\bar{V}_i \cdot \log \frac{\bar{V}_i}{\pi} + (1 - \bar{V}_i) \cdot \log \frac{1 - \bar{V}_i}{1 - \pi} \right] \quad (11)$$

where $\bar{V}_i$ represents the mean value of $\hat{V}_i$. By optimizing V with L_s , it encourages each expert activate an expected proportion π of neurons with heterogeneous importance.

In practice, considering the non-differentiability of Bernoulli sampling during training, we use the reparameterization trick to approximate the Bernoulli distribution (Jang et al., 2017). The sampling process for $\hat{V}_i$ is computed as follows:

$$M_i = \sigma \left(\frac{\log \hat{V}_i + g}{\tau} \right) \quad (12)$$

where $g = -\log(-\log(u))$, $u \sim U(0, 1)$, $\tau = 0.5$ is a hyperparameter for controlling sampling randomization, and $\sigma(\cdot)$ is sigmoid function.

Finally, for Equation 8, $q(y|\hat{h})$ is typically achieved by a network q_θ used for the next token prediction task, with the loss function being cross-entropy. The loss function is defined as:

$$L_p = \sum_i -\log(q_\theta(y_i|\hat{h}_i)). \quad (13)$$

4.2 Neuron Collaboration

In this section, we introduce Neuron Collaboration (NC) to enhance the knowledge diversity among experts. Traditional MoE frameworks suffer from knowledge redundancy due to task similarity, but in IBMoE, we tackle this by applying a feature-level diversity loss. It controls the knowledge differences among experts, promoting task specialization. Specifically, a query's distribution represents the knowledge characteristics of an expert. For two queries $\hat{V}_i$ and $\hat{V}_j$, we minimize their similarity, encouraging experts to focus on different feature dimensions through diverse query distributions. The diversity constraint is expressed as:

$$L_d = \sum_{1 \leq i < j \leq n} \text{cosine}(\hat{V}_i, \hat{V}_j) \quad (14)$$

where $\text{cosine}(\cdot, \cdot)$ is a function for computing the cosine similarity between two vectors.

We also employ an auxiliary regularization term based on the variance discrepancy between expert-level and token-level routing probabilities. Let P_r represent the probability distribution of tokens assigned to the experts. The objective function is defined as:

$$L_b = \frac{n}{k} e^{[\text{std}(\text{mean}_t(P_r)) - \text{mean}(\text{std}_e(P_r))]} \quad (15)$$

where k is the number of activated experts, $\text{mean}_t(\cdot)$ represents the token-wise mean value, and $\text{std}_e(\cdot)$ represents the expert-wise standard deviation.

Combining Equations 11, 13, 14 and 15, the final loss function is expressed as:

$$L = L_p + \beta L_s + \gamma L_d + \lambda L_b \quad (16)$$

The complete pseudo-code is provided in Appendix A.3.

5 Experimental Setup

5.1 Datasets

We evaluate IBMoE on seven NLP datasets, including three GLUE tasks, three question answering (QA) datasets, and one instruction-following dataset from SuperNI. For GLUE, we select one representative task from each category: **CoLA** (Warstadt et al., 2019), **MRPC** (Dolan and Brockett, 2005), and **RTE** (Wang et al., 2019). For QA evaluation, we use **ScienceQA (SQA)** (Lu et al., 2022), **OpenbookQA (OBQA)** (Mihaylov et al., 2018), and **CommonSenseQA (CSQA)** (Talmor et al., 2019). For instruction-following evaluation, we include **Convai3** from SuperNI (Wang et al., 2022). Detailed dataset formatting and preprocessing procedures are provided in Appendix B.

5.2 Base Models and Baselines

We evaluate IBMoE on three base models of different scales: LLaMA-3.2-3B (L-3B) (Dubey et al., 2024), LLaMA-2-7B (L-7B) (Touvron et al., 2023b), and Mistral-7B-v0.1 (M-7B) (Jiang et al., 2023). To evaluate the effectiveness of our model, we compare our IBMoE with recent parameter-efficient fine-tuning (PEFT)-based and MoE-based methods. For PEFT-based methods, we select **LoRA** (Hu et al., 2021) and **SparseAdapter** (He et al., 2022). For MoE-based methods, we employ **MoLA** (Gao et al., 2024), **AlphaLoRA** (Qing et al., 2024), **AdaMoLE** (Liu and Luo, 2024) and **HMoRA** (Liao et al., 2025). Detailed descriptions of baseline methods are provided in Appendix C.

5.3 Evaluation Metrics

We evaluate IBMoE from four perspectives: task performance, parameter efficiency, training cost, and inference efficiency. For GLUE and QA tasks, we report **accuracy**, and for generation tasks, we use **rougeL** to measure textual similarity. To evaluate efficiency, we report the number of activated parameters (**Act Param**), training time (**GPU-hour**), and inference cost **Time** and **Peak Memory**. Overall, these metrics provide a comprehensive evaluation of model effectiveness and efficiency.

5.4 Implementation Details

The hyperparameters are set as follows: (1) **Training Parameters**: We use the Adafactor optimizer with a learning rate of $2e^{-4}$, a batch size of 128. The number of epochs is set to 10. (2) **LoRA adapter**: We adopt the default LoRA configuration, with the intrinsic LoRA rank of experts set to 8 for all the baselines. (3) **IBMoE Parameters**: We use 5 LoRA experts in total, including four task-specific experts and one shared expert. During forward processing, the top- k is set to 2, selecting one specific expert and the shared expert for each token. We set $\pi = 0.6$, $\beta = 0.1$, $\gamma = 0.01$, and $\lambda = 0.01$. All experiments are conducted on the NVIDIA A800 GPU. Additional implementation details are provided in Appendix D.

6 Experimental Results

6.1 Main Results

Table 1 compares IBMoE with representative PEFT-based and MoE-based baselines. IBMoE achieves the best overall performance on all three base models while using the fewest activated parameters. Compared with sparse baselines such as SparseAdapter and MoLA, IBMoE consistently delivers stronger accuracy with substantially lower parameter usage, indicating that dynamic neuron-level sparsity is more effective than static adapter pruning or expert-level activation alone. These results suggest that modeling fine-grained neuron importance within experts is crucial for reducing redundancy and improving parameter efficiency. The results averaged over multiple random seeds and the corresponding statistical significance analysis are provided in Appendix E.1.

6.2 Ablation Studies

In Table 2, we conduct ablation studies to evaluate the effectiveness of IBMoE’s components. First, we experiment without L_b . We find that imbalanced routing leads to some experts handling disproportionate amounts of data by activating more neurons, which increases the **Act Param**. Next, we remove the sparsity constraint, which makes the experts fail to focus on the most relevant information, and unstable neuron sampling further reduces performance. Finally, we remove the diversity loss, which worsens IBMoE’s performance. Even with a shared expert, omitting L_d can still lead to knowledge redundancy among the experts. Additionally, without any loss constraints, the number of acti-

Table 1: Comparison of performance and parameter efficiency with other methods across different model architectures. Average score across seven datasets (**Avg**) is reported to evaluate overall performance. Best results are highlighted in **bold**.

Models	Methods	Act Param ↓	Datasets							Avg
			MRPC	CoLA	RTE	SQA	OBQA	CSQA	Convai	
L-3B	LoRA	7.6M	83.24	86.86	88.44	92.21	79.40	77.73	30.32	76.89
	SparseAdapter	3.04M	84.86	85.71	87.73	93.12	80.20	76.49	30.84	76.99
	MoLA	23.59M	85.77	82.13	84.12	91.29	80.00	77.16	28.49	75.57
	AlphaLoRA	26.16M	84.90	82.13	83.39	92.51	80.00	78.63	29.35	75.84
	AdaMoLE	dynamic	84.69	84.27	87.73	91.68	80.40	77.81	29.56	76.59
	HMoRA	30.91M	83.88	86.58	84.47	92.53	80.20	77.31	31.36	76.62
	IBMoE	3.01M	84.69	86.10	88.80	93.16	80.20	78.13	31.92	77.57
L-7B	LoRA	19.9M	83.13	86.29	85.92	91.01	75.51	77.00	32.53	75.91
	SparseAdapter	7.96M	86.38	85.62	85.30	92.85	82.20	79.36	32.08	77.68
	MoLA	43.12M	83.48	86.87	86.28	92.36	78.95	79.60	30.41	76.85
	AlphaLoRA	43.12M	84.23	86.67	87.36	92.71	80.80	78.05	30.66	77.21
	AdaMoLE	dynamic	83.88	86.52	87.73	91.00	82.40	78.71	31.66	77.41
	HMoRA	43.85M	84.63	86.56	87.36	93.61	79.40	79.36	31.68	77.51
	IBMoE	7.89M	87.25	87.15	88.43	93.34	82.40	79.77	32.09	78.63
M-7B	LoRA	20.50M	85.86	86.23	90.97	94.86	86.20	79.85	31.88	79.41
	SparseAdapter	8.20M	87.07	85.33	91.33	95.05	88.80	81.16	27.35	79.44
	MoLA	43.40M	86.95	87.44	88.80	95.14	88.20	83.37	29.92	79.97
	AlphaLoRA	43.40M	87.13	87.91	91.70	95.00	89.20	84.00	29.59	80.65
	AdaMoLE	dynamic	85.91	86.10	92.05	94.90	87.80	82.55	31.88	80.17
	HMoRA	44.20M	87.48	87.53	90.28	94.60	88.60	82.64	31.07	80.31
	IBMoE	8.22M	88.58	87.53	92.77	95.68	90.40	83.05	31.83	81.41

vated parameters in IBMoE increases substantially because the inactive SNA module causes the model to activate all neurons to complete the tasks. The complete results are provided in Appendix E.2.

Model	RTE	
	Acc (%↑)	Act Params ↓
IBMoE	88.43	7.89M
$-L_b$	87.95	7.91M
$-L_b - L_s$	87.06	7.07M
$-L_b - L_s - L_d$	86.29	13.16M

Table 2: Ablation results of IBMoE on the RTE dataset.

6.3 Further Analyses

We conduct additional analyses to better understand the effectiveness and efficiency of IBMoE. More experimental results are deferred to the Appendix E.


Analysis of Activation Ratio. In IBMoE, the parameter π controls the proportion of activated neu-

rons. The results in Figure 3 illustrate the model’s performance under various activation ratios across multiple datasets. We observe that model performance generally improves as more neurons are activated. However, this improvement gradually diminishes and may eventually reverse as π continues to increase. This observation suggests the presence of neuron redundancy in conventional MoE methods, which can potentially degrade model performance by activating task-irrelevant neurons. Moreover, the model’s performance on the RTE dataset drops by approximately 3% when π is set to 0.2, indicating that an excessively low activation ratio may limit the model’s capacity to capture task-relevant information.

Analysis of Hyperparameters. The IBMoE framework introduces three hyperparameters (λ , β , and γ) for the components of the loss function. To analyze the sensitivity of these hyperparameters, we conduct experiments by varying their values. Specifically, we adjust one parameter (e.g., λ)

Table 3: Training and Inference efficiency of IBMoE.

Model	Inference with 1,000 Samples		Training on the RTE Dataset	
	Time (s) ↓	Peak Memory (MB) ↓	Train Param (%)↓	GPU-hour ↓
LoRA	60	116	0.3	0.77h
SparseAdapter	67.34	232	0.3	0.81h
MoLA	739.3	1452.6	1.51	2.93h
AlphaLoRA	285.7	812	1.51	2.58h
AdaMoLE	329.8	146.2	0.53	1.97h
HMoRA	295.6	842.1	3.15	3.55h
IBMoE	109.3	75.7	0.47	0.88h

Figure 3: Activation Ratio π with different values.

while fixing the others to their default values. The same procedure is repeated for β and γ . As shown in Figure 4, the proposed method exhibits strong robustness to hyperparameter variations. Notably, setting lower values (e.g., 0.01 or 0.001) leads to a slight improvement in model performance. However, when λ is set to 0.1, the scaling of the balance loss L_b dominates the primary task loss L_p . In this case, tokens with similar semantic information are forcibly assigned to different experts to ensure load balancing, which makes it difficult for experts to process semantically diverse tokens under sparse representations.

Analysis of Greedy Sampling. To further analyze neuron activation, we replace Bernoulli sampling with greedy sampling during inference. Specifically, top- p sampling is applied on V to activate the most relevant neurons, where the binary mask is defined as $M_i = \text{TopP}(\hat{V})$. Figure 5 presents the results of IBMoE with greedy sampling under different activation ratios. We observe that greedy sampling performs worse than Bernoulli-based importance sampling under the same activation scale, since it focuses only on highly ranked neurons and lacks global feature exploration. As p increases, its performance gradu-





Figure 4: Sensitivity of hyperparameters.

Figure 5: Impact of p on the different tasks.

ally approaches that of importance sampling, indicating that importance sampling achieves a better trade-off between performance and efficiency.

Analysis of the Inference Efficiency. We also compare the inference performance of IBMoE with other methods to evaluate its efficiency. To ensure the fairness and reproducibility of all methods, we conduct inference using a single NVIDIA A800 GPU with 80 GB memory. Specifically, we evaluate L-7B model in 16-bit precision on 1,000 randomly sampled test examples, measuring inference latency and peak memory usage of the adapters

with a batch size of 16. In Table 3, our proposed method achieves the best efficiency among MoE-based PEFT methods, reducing inference time by 61.7% compared with AlphaLoRA and adapter peak memory by 48.2% compared with AdaMoLE. We attribute this advantage to the fact that the trained IBMoE significantly reduces computational cost and memory consumption by compressing the weight matrices of LoRA experts.

7 Conclusions and Future Work

In this paper, we propose a novel parameter-efficient MoE method that introduces an IB framework and multiple constraints to achieve fine-grained reduction of neuronal redundancy, fostering closer collaboration among experts. Extensive experiments show that our method achieves superior performance comparable to recent methods, and reduces inference time by 61.7% and adapter memory by 48.2% compared with MoE-based approaches. Furthermore, ablation studies validate the effectiveness of each IBMoE component. In future work, we will explore applications in continual learning by capitalizing on the knowledge specialization inherent in MoE.

Limitations

Despite its effectiveness, IBMoE has several limitations. First, although we adopt a top- k sparse routing strategy for efficiency and specialization, sparse routing introduces additional overhead for token allocation and expert dispatch in distributed settings. In contrast, soft routing avoids this issue but incurs higher computation during expert activation. Second, IBMoE introduces additional hyperparameters and structural components compared with standard PEFT methods, such as the number of experts and activation settings, which may require manual tuning when adapting to new tasks or domains. Finally, while we evaluate IBMoE on a wide range of NLP tasks, its behavior under continual learning settings remains unexplored.

Acknowledgments

This research is funded by the National Nature Science Foundation of China (No.62477010, No.62577022 and No.62307028), Shanghai General AI Foundation Models Program (No.2025SHZDZX026D24) and Shanghai Science and Technology Innovation Action Plan (No.24YF2710100).

References

- Tanay Aggarwal, Angelo Salatino, Francesco Osborne, and Enrico Motta. 2026. [Large language models for scholarly ontology generation: An extensive analysis in the engineering field](#). *Information Processing & Management*, 63:104262.
- Alexander A Alemi, Ian Fischer, Joshua V Dillon, and Kevin Murphy. 2016. [Deep variational information bottleneck](#). *arXiv preprint arXiv:1612.00410*, arXiv:1612.00410.
- Zijian Bai, Yucheng Zheng, Peng Yang, Siyang Liu, Yiyuan Zhang, and Yuanyuan Chang. 2025. [Dcrl-rec: Dual-domain contrastive reinforcement large language model for recommendation](#). *Information Processing & Management*, 62(4):104140.
- Emmanuel Bengio, Pierre-Luc Bacon, Joelle Pineau, and Doina Precup. 2015. [Conditional computation in neural networks for faster models](#). *arXiv preprint arXiv:1511.06297*, arXiv:1511.06297.
- Marco Braga, Alessandro Raganato, Gabriella Pasi, and 1 others. 2024. Adakron: An adapter-based parameter efficient model tuning with kronecker product. In *2024 Joint International Conference on Computational Linguistics, Language Resources and Evaluation, LREC-COLING 2024-Main Conference Proceedings*, pages 350–357.
- Weilin Cai, Juyong Jiang, Fan Wang, Jing Tang, Sunghun Kim, and Jiayi Huang. 2024. [A survey on mixture of experts](#). *arXiv preprint arXiv:2407.06204*, arXiv:2407.06204.
- Jiawei Chen, Ziqi Zhang, Xinpeng Xie, Yuexiang Li, Tao Xu, Kai Ma, and Yefeng Zheng. 2022. [Beyond mutual information: Generative adversarial network for domain adaptation using information bottleneck constraint](#). *IEEE Transactions on Medical Imaging*, 41(3):595–607.
- Damai Dai, Chengqi Deng, Chenggang Zhao, RX Xu, Huazuo Gao, Deli Chen, Jiashi Li, Wangding Zeng, Xingkai Yu, Y Wu, and 1 others. 2024. [Deepseek-moe: Towards ultimate expert specialization in mixture-of-experts language models](#). *arXiv preprint arXiv:2401.06066*, arXiv:2401.06066.
- Giang Do, Hung Le, and Truyen Tran. 2025. Simsmoe: Toward efficient training mixture of experts via solving representational collapse. In *Findings of the Association for Computational Linguistics: NAACL 2025*, pages 2012–2025.
- William B. Dolan and Chris Brockett. 2005. Automatically Constructing a Corpus of Sentential Paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing*, pages 9–16.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, and 1 others. 2024. [The llama 3 herd of models](#). *arXiv preprint arXiv:2407.21783*, arXiv:2407.21783.

- William Fedus, Barret Zoph, and Noam Shazeer. 2022. [Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity](#). *Journal of Machine Learning Research*, 23(120):1–39.
- Chongyang Gao, Kezhen Chen, Jinmeng Rao, Baochen Sun, Ruiyuan Guo, Daiyi Peng, Yawen Zhang, Xiaoyuan Guo, Jie Yang, and VS Subrahmanian. 2024. [Higher layers need more lora experts](#). *arXiv preprint arXiv:2402.08562*, arXiv:2402.08562.
- Yifei Gao, Ning Xu, Hongshuo Tian, Bolun Zheng, Jinbo Cao, and An-An Liu. 2025. [Mixture of causal experts: A causal perspective to build dual-level mixture-of-experts models](#). *Expert Systems with Applications*, 280:127422.
- Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shitong Ma, Peiyi Wang, Xiao Bi, and 1 others. 2025. [Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning](#). *arXiv preprint arXiv:2501.12948*, arXiv:2501.12948.
- Shwai He, Liang Ding, Daize Dong, Jeremy Zhang, and Dacheng Tao. 2022. Sparseadapter: An easy approach for improving the parameter-efficiency of adapters. In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2184–2190.
- Shwai He, Daize Dong, Liang Ding, and Ang Li. 2024. [Towards efficient mixture of experts: A holistic study of compression techniques](#). *arXiv preprint arXiv:2406.02500*, arXiv:2406.02500.
- Shwai He, Run-Ze Fan, Liang Ding, Li Shen, Tianyi Zhou, and Dacheng Tao. 2023. [Merging experts into one: Improving computational efficiency of mixture of experts](#). *arXiv preprint arXiv:2310.09832*, arXiv:2310.09832.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *arXiv preprint arXiv:2106.09685*, arXiv:2106.09685.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical Reparameterization with Gumbel-Softmax. In *Proceedings of 5th International Conference on Learning Representations*.
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. 2023. [Mistral 7b](#). *arXiv preprint arXiv:2310.06825*, arXiv:2310.06825.
- Kenji Kawaguchi, Zhun Deng, Xu Ji, and Jiaoyang Huang. 2023. How does information bottleneck help deep learning? In *Proceedings of the 40th International Conference on Machine Learning*, pages 16049–16096.
- Dmitry Lepikhin, HyoukJoong 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](#). *arXiv preprint arXiv:2006.16668*, arXiv:2006.16668.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). *arXiv preprint arXiv:2104.08691*, arXiv:2104.08691.
- Linjie Li, Zhenyu Wu, and Yang Ji. 2025. [Mote: Mixture of task-specific experts for pre-trained model-based class-incremental learning](#). *Knowledge-Based Systems*, 324:113795.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#). *arXiv preprint arXiv:2101.00190*, arXiv:2101.00190.
- Mengqi Liao, Wei Chen, Junfeng Shen, Shengnan Guo, and Huaiyu Wan. 2025. Hmora: Making llms more effective with hierarchical mixture of lora experts. In *The Thirteenth International Conference on Learning Representations*.
- Zexu Lin, Xu Zhang, Xiaoyu Hu, Pengfei Li, Deyu Zhou, and Darong Lai. 2026. [Hierarchical mixture-of-experts model for unified code search](#). *Information Fusion*, 125:103489.
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohata, Tenghao Huang, Mohit Bansal, and Colin Raffel. 2022. [Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning](#). *Advances in Neural Information Processing Systems*, 35:1950–1965.
- Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Lam Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. 2021. [P-tuning v2: Prompt tuning can be comparable to fine-tuning universally across scales and tasks](#). *arXiv preprint arXiv:2110.07602*, arXiv:2110.07602.
- Yijiang Liu, Rongyu Zhang, Huanrui Yang, Kurt Keutzer, Yuan Du, Li Du, and Shanghang Zhang. 2024. [Intuition-aware mixture-of-rank-1-experts for parameter efficient finetuning](#). *arXiv preprint arXiv:2404.08985*, arXiv:2404.08985.
- Zefang Liu and Jiahua Luo. 2024. [Adamole: Fine-tuning large language models with adaptive mixture of low-rank adaptation experts](#). *arXiv preprint arXiv:2405.00361*, arXiv:2405.00361.
- Stephan Sloth Lorenzen, Christian Igel, and Mads Nielsen. 2021. [Information bottleneck: Exact analysis of \(quantized\) neural networks](#). *arXiv preprint arXiv:2106.12912*, arXiv:2106.12912.
- Pan Lu, Swaroop Mishra, Tanglin Xia, Liang Qiu, Kaiwei Chang, Song-Chun Zhu, Oyvind Tafjord, Peter Clark, and Ashwin Kalyan. 2022. [Learn to explain: Multimodal reasoning via thought chains for science question answering](#). *Advances in Neural Information Processing Systems*, 35:2507–2521.

- 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 *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2381–2391.
- OpenAI. 2023. [Gpt-4 technical report](#). *arXiv preprint arXiv:2303.08774*, arXiv:2303.08774.
- Peijun Qing, Chongyang Gao, Yefan Zhou, Xingjian Diao, Yaoqing Yang, and Soroush Vosoughi. 2024. Alphalora: Assigning lora experts based on layer training quality. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*, pages 20511–20523.
- Pengjie Ren, Chengshun Shi, Shiguang Wu, Mengqi Zhang, Zhaochun Ren, Maarten Rijke, Zhumin Chen, and Jiahuan Pei. 2024. Melora: Mini-ensemble low-rank adapters for parameter-efficient fine-tuning. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3052–3064.
- Changchong Sheng, Li Liu, Wanxia Deng, Liang Bai, Zhong Liu, Songyang Lao, Gangyao Kuang, and Matti Pietikäinen. 2023. [Importance-aware information bottleneck learning paradigm for lip reading](#). *IEEE Transactions on Multimedia*, 25:6563–6574.
- DJ Strouse and David J Schwab. 2017. [The deterministic information bottleneck](#). *Neural computation*, 29(6):1611–1630.
- Alon Talmor, Jonathan Herzig, Nicholas Lourie, and Jonathan Berant. 2019. CommonsenseQA: A Question Answering Challenge Targeting Commonsense Knowledge. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4149–4158.
- Shiqiang Tian, Cheng Ding, Qin Chen, Jie Zhou, and Liang He. 2026. Tamedit: Trajectory-aware meta-learning for specificity-preserving continual knowledge editing. In *Proceedings of the 64th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 21417–21439.
- Naftali Tishby, Fernando C Pereira, and William Bialek. 2000. [The information bottleneck method](#). *arXiv preprint physics/0004057*, physics/0004057.
- Naftali Tishby and Noga Zaslavsky. 2015. Deep learning and the information bottleneck principle. In *2015 IEEE information theory workshop (itw)*, pages 1–5.
- Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, and 1 others. 2023a. [Llama: Open and efficient foundation language models](#). *arXiv preprint arXiv:2302.13971*, arXiv:2302.13971.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, and 1 others. 2023b. [Llama 2: Open foundation and fine-tuned chat models](#). *arXiv preprint arXiv:2307.09288*, arXiv:2307.09288.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. GLUE: A Multi-Task Benchmark and Analysis Platform for Natural Language Understanding. In *Proceedings of 7th International Conference on Learning Representations*.
- Yifan Wang, Yixin Cao, Zeping Li, Bowen Dong, Guangan Ye, and Hongfeng Chai. 2025. [Multi-hgpt: Multi-task heterogeneous graph prompt tuning](#). *Information Processing & Management*, 62:104236.
- Yizhong Wang, Swaroop Mishra, Pegah Alipoor-molabashi, Yeganeh Kordi, Amirreza Mirzaei, Anjana Arunkumar, Arjun Ashok, Arut Selvan Dhanasekaran, Atharva Naik, David Stap, and 1 others. 2022. Super-naturalinstructions: Generalization via declarative instructions on 1600+ nlp tasks. In *2022 Conference on Empirical Methods in Natural Language Processing, EMNLP 2022*, pages 5085–5109.
- Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. 2019. [Neural network acceptability judgments](#). *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Yanyue Xie, Zhi Zhang, Ding Zhou, Cong Xie, Ziang Song, Xin Liu, Yanzhi Wang, Xue Lin, and An Xu. 2024. [Moe-pruner: Pruning mixture-of-experts large language model using the hints from its router](#). *arXiv preprint arXiv:2410.12013*, arXiv:2410.12013.
- Junchi Yu, Tingyang Xu, Yu Rong, Yatao Bian, Junzhou Huang, and Ran He. 2024. [Recognizing predictive substructures with subgraph information bottleneck](#). *IEEE transactions on pattern analysis and machine intelligence*, 46(3):1650–1663.
- Ted Zadouri, Ahmet Üstün, Arash Ahmadian, Beyza Ermiş, Acyr Locatelli, and Sara Hooker. 2023. [Pushing mixture of experts to the limit: Extremely parameter efficient moe for instruction tuning](#). *arXiv preprint arXiv:2309.05444*, arXiv:2309.05444.
- Elad Ben Zaken, Shauli Ravfogel, and Yoav Goldberg. 2021. [Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models](#). *arXiv preprint arXiv:2106.10199*, arXiv:2106.10199.
- Rongyu Zhang, Yulin Luo, Jiaming Liu, Huanrui Yang, Zhen Dong, Denis Gudovskiy, Tomoyuki Okuno, Yohei Nakata, Kurt Keutzer, Yuan Du, and 1 others. 2024. Efficient Deweather Mixture-of-Experts with Uncertainty-Aware Feature-Wise Linear Modulation. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence*, pages 16812–16820.

- Shuai Zhang, Junfeng Fang, Xuqiang Li, ALAN XIA, Ye Wei, Wenjie Du, Yang Wang, and 1 others. 2025. Iterative substructure extraction for molecular relational learning with interactive graph information bottleneck. In *The Thirteenth International Conference on Learning Representations*.
- Junda Zhao, Yuliang Song, and Eldan Cohen. 2025. Variational prefix tuning for diverse and accurate code summarization using pre-trained language models. *Journal of Systems and Software*, 229:112493.
- Kun Zhu, Xiaocheng Feng, Xiyuan Du, Yuxuan Gu, Weijiang Yu, Haotian Wang, Qianglong Chen, Zheng Chu, Jingchang Chen, and Bing Qin. 2024a. An information bottleneck perspective for effective noise filtering on retrieval-augmented generation. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1044–1069.
- Shaolin Zhu, Leiyu Pan, Dong Jian, and Deyi Xiong. 2025. Overcoming language barriers via machine translation with sparse mixture-of-experts fusion of large language models. *Information Processing & Management*, 62:104078.
- Tong Zhu, Xiaoye Qu, Daize Dong, Jiacheng Ruan, Jingqi Tong, Conghui He, and Yu Cheng. 2024b. Llama-moe: Building mixture-of-experts from llama with continual pre-training. *arXiv preprint arXiv:2406.16554*, arXiv:2406.16554.
- Yun Zhu, Nevan Wichers, Chu-Cheng Lin, Xinyi Wang, Tianlong Chen, Lei Shu, Han Lu, Canoe Liu, Liangchen Luo, Jindong Chen, and 1 others. 2023. Sira: Sparse mixture of low rank adaptation. *arXiv preprint arXiv:2311.09179*, arXiv:2311.09179.
- Simiao Zuo, Qingru Zhang, Chen Liang, Pengcheng He, Tuo Zhao, and Weizhu Chen. 2022. Moe-bert: from bert to mixture-of-experts via importance-guided adaptation. *arXiv preprint arXiv:2204.07675*, arXiv:2204.07675.

A Notation and Explanation

A.1 Notations of IBMoE

In Table 4, we summarize the notations used in IBMoE and explicitly specify their tensor dimensions. On the other hand, in Equation 10, the column-level masking operation:

$$\Delta\hat{W}_e = M_e \circ \Delta W_e \quad (17)$$

is defined as follows. Let $M_e \in \{0, 1\}^{d_{\text{out}}}$ be a binary mask vector and $\Delta W_e \in \mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}$ be the LoRA weight of expert e . The operation $\circ$ applies M_e to each column of ΔW_e :

$$(\Delta\hat{W}_e)_{j,:} = M_{e,j} \cdot (\Delta W_e)_{j,:}, \quad (18)$$

where $(\cdot)_{j,:}$ denotes the j -th row. This operation selectively activates neurons while preserving a complete computation graph during training. Note that the operator $\circ$ in IBMoE has a different definition from the conventional Hadamard product. However, due to the broadcasting mechanism in Python, the involved tensors are first expanded to compatible shapes and then multiplied element-wise. Therefore, we still use $\circ$ to denote this type of operation in IBMoE.

A.2 Detailed Forward Mechanism of SNA

Figure 6 shows the forward mechanism of the SNA module. During training, the neuron-level mask M_e is applied to the weight update ΔW_e to ensure that the full computation graph is preserved for backpropagation. This design allows gradients to flow through all neurons while enforcing sparse activation at the neuron level. During inference, the mask M_e is directly used to select the activated neurons for computing the compressed representation $\hat{H}$, without constructing the full computation graph. This decoupling between training and inference enables efficient optimization while maintaining sparse neuron activation at test time.

A.3 Pseudo Code of IBMoE

We provide pseudo-code to describe the optimization procedure of IBMoE. Specifically, Algorithm 1 illustrates the sparse activation process of SNA during forward computation. Algorithm 2 further presents the forward pass of an IBMoE layer. Finally, Algorithm 3 details the computation of the overall objective function.

B Details of Dataset

We provide detailed descriptions of all datasets used in this work.

B.1 GLUE Benchmark

The General Language Understanding Evaluation (GLUE) benchmark consists of multiple NLP tasks, including single-sentence classification, sentence similarity, and natural language inference.

We select three representative tasks:

- **CoLA** (Warstadt et al., 2019): A grammatical acceptability classification task based on linguistic theory corpora.
- **MRPC** (Dolan and Brockett, 2005): A paraphrase identification dataset from news articles.
- **RTE** (Wang et al., 2019): A natural language inference task constructed from Wikipedia and news sources.

B.2 Question Answering Datasets

We evaluate reasoning and commonsense capabilities using:

- **ScienceQA (SQA)** (Lu et al., 2022): A science-based QA dataset originally containing multimodal samples. We retain only text-based questions.
- **OpenbookQA (OBQA)** (Mihaylov et al., 2018): QA dataset requiring elementary science knowledge reasoning.
- **CommonSenseQA (CSQA)** (Talmor et al., 2019): Multiple-choice QA requiring commonsense reasoning.


B.3 Instruction Dataset

Convai3 Sentence Generation from SuperNI (Wang et al., 2022) is a dialogue generation task requiring response generation conditioned on conversational context.

B.4 Examples of formatted datasets

In the experimental settings, all datasets are drawn from the GLUE, QA, and SuperNI benchmarks. To ensure reproducibility, we provide representative examples of the formatted datasets to illustrate the construction of our training inputs, as shown in Figure 7. For datasets from the GLUE and QA

(a) Training Phase



(b) Inference Phase




Figure 6: Forward mechanism of the SNA module. During the training stage, M_e is directly multiplied with ΔW_e to obtain sparse representations, which ensures that gradients flow through the differentiable relaxation of the mask. During the inference stage, the query mask is used to index the activated columns, enabling more efficient computation.

benchmarks, we adopt the Alpaca instruction format to construct the training inputs for all methods. For the ConvAI task, we follow the default instruction format provided by the dataset.

C Baseline Methods

We compare IBMoE with representative PEFT-based and MoE-based methods.

C.1 PEFT-based Methods

- **LoRA** (Hu et al., 2021): Introduces low-rank matrices for efficient adaptation.
- **SparseAdapter** (He et al., 2022): Applies structured sparsity to improve parameter efficiency.

C.2 MoE-based Methods

- **MoLA** (Gao et al., 2024): Allocates experts across layers to reduce redundancy.
- **AlphaLoRA** (Qing et al., 2024): Designs layer-wise expert allocation based on heavy-tailed regularization.
- **AdaMoLE** (Liu and Luo, 2024): Dynamically adjusts the number of active experts per layer.
- **HMoRA** (Liao et al., 2025): Combines token-level and task-level routing for finer-grained expert control.

Algorithm 1 Function GSSampleMask(V_i, τ, ϵ)

```

1: function GSSAMPLEMASK( $V_i, \tau, \epsilon$ )
2:    $U \sim \mathcal{U}(0, 1)$ 
3:    $\hat{V}_i \leftarrow (V_i - \min(V_i) + \epsilon) / (\max(V_i) - \min(V_i) + \epsilon)$ 
4:    $\hat{U} \leftarrow \log(\hat{V}_i) - \log(-\log(U + \epsilon) + \epsilon)$ 
5:    $S \leftarrow \sigma(\hat{U} / \tau)$ 
6:    $M \leftarrow \mathbf{0}^{|S|}$ 
7:   for  $j = 1$  to  $|S|$  do
8:     if  $S[j] > 1 - \hat{V}_{i,j}$  then
9:        $M[j] \leftarrow 1$ 
10:    end if
11:  end for
12:  return  $M$ 
13: end function

```

D Implementation Details

Table 5 summarizes the training hyperparameters used for IBMoE, along with the ranges explored during grid search.

E Additional Results

E.1 Multi-seed Results and Statistical Significance Analysis

Table 6 reports the mean and standard deviation of LLaMA-3.2-3B and LLaMA-2-7B across multiple random seeds (40, 41, and 42). Furthermore, we conduct statistical significance testing by treating dataset-seed pairs as independent samples. Compared with the strongest baseline, IB-

The formatted examples of QA datasets	The formatted examples of Convai dataset
Below is an instruction that describes a task. Write a response that appropriately completes the request. ### Instruction: Context: N/A Question: As a source of light moves directly overhead of an object , what happens to it's shadow? Options: (A) the object will get brighter (B) the shadow will start to spin (C) the size will decrease (D) the light will lean to one side ### Response: Answer: The answer is C.	Definition: You are given a concatenated string consisting of two parts that are separated by a newline character. The first part is the user's initial question during a human computer conversation, the second part is a clarifying question asked by the computer in response to the question asked by the user. You should read the input, then generate a valid prediction of the user's response to the computer's clarifying question. A valid prediction will clarify the user's correct intent in the initial question. Now complete the following example - Input: User: Find information on raised garden beds and boxes. Computer: do you need information on materials needed. Output: no i am looking for pictures of raised garden beds not the materials needed.
(a)	(b)

Figure 7: The examples of training data.

Table 4: Summary of notations and tensor dimensions.

Symbol	Description	Dimension
x_i	Input of token i	$\mathbb{R}^{d_{in}}$
y_i	Target label / next token	scalar / categorical
W_0	Frozen base weight matrix	$\mathbb{R}^{d_{in} \times d_{out}}$
ΔW_e	LoRA weight of expert e	$\mathbb{R}^{d_{out} \times d_{in}}$
A	LoRA down-projection matrix	$\mathbb{R}^{r \times d_{in}}$
B	LoRA up-projection matrix	$\mathbb{R}^{d_{out} \times r}$
$G_{i,e}$	Routing weight for expert e at token i	scalar
V	Expert-specific neuron query vector	$\mathbb{R}^{d_{out}}$
M	Binary neuron activation mask	$\{0, 1\}^{d_{out}}$
$\Delta \hat{W}_e$	Masked LoRA weight of expert e	$\mathbb{R}^{d_{out} \times d_{in}}$
h_i	Original hidden output at token i	$\mathbb{R}^{d_{out}}$
$\hat{h}_i$	Compressed hidden representation at token i	$\mathbb{R}^{d_{out}}$
n	Number of experts	scalar
k	Number of activated experts	scalar
β	Sparsity regularization coefficient	scalar
γ	Diversity loss coefficient	scalar
λ	Balance loss coefficient	scalar
π	Bernoulli prior probability	scalar

MoE achieves a statistically significant improvement with $p = 0.0103 < 0.05$.

E.2 Ablation Study Results

In Table 7, we report the full ablation results across all datasets.

E.3 Different Expert Settings

In IBMoE, multiple constraints are introduced in the objective function. To investigate the effectiveness of these constraints under different numbers of experts and activated experts, we scale up each MoE layer and evaluate the performance of

IBMoE. Specifically, we set the number of experts to 9, including 8 task-specific experts and 1 shared expert, with top- k set to 3. Experiments are conducted on three datasets from the GLUE benchmark, and the results are reported in Table 8. As shown in the table, although increasing the number of experts leads to improved performance, the gains are marginal, indicating that introducing more trainable parameters exacerbates knowledge redundancy and does not yield the expected performance benefits.

Algorithm 2 Forward Process of an IBMoE Layer.

Require: (1) A batch of input X ; (2) Base module W_0 ; (3) LoRA weights of experts $\{\Delta W_1, \Delta W_2, \dots, \Delta W_n\}$ with learnable queries $\{V_1, V_2, \dots, V_n\}$; (4) Router W_r ; (5) Related hyperparameters, including temperature τ and ϵ .

- 1: $P \leftarrow \text{TopK}(\text{Softmax}(W_r^\top X))$
- 2: Dispatch tokens to experts $\{X_{E_1}, X_{E_2}, \dots, X_{E_n}\}$ with corresponding routing probabilities $\{P_{E_1}, P_{E_2}, \dots, P_{E_n}\}$ from P
- 3: $O \leftarrow W_0 X$
- 4: **for** $i = 1$ to n **do**
- 5: $M_i \leftarrow \text{GSSampleMask}(V_i, \tau, \epsilon)$
- 6: $s \leftarrow d_{\text{out}} / \text{sum}(M_i)$
- 7: **if** mode = train **then**
- 8: $O_{E_i} \leftarrow s \cdot P_{E_i} \cdot M_i \cdot (\Delta W_i X_{E_i})$
- 9: Aggregate expert output O_{E_i} into O
- 10: **else**
- 11: Compress ΔW_i into $\Delta \hat{W}_i$ via mask M_i
- 12: $\hat{O}_{E_i} \leftarrow s \cdot P_{E_i} \cdot (\Delta \hat{W}_i X_{E_i})$
- 13: Aggregate expert output $\hat{O}_{E_i}$ into O
- 14: **end if**
- 15: **end for**

Table 5: Hyperparameters of IBMoE.

Symbol	Value or Range
α	16
r	8
n	5
k	2
τ	0.5
ϵ	1e-8
π	{0.2, 0.4, 0.6, 0.8}
β	{0.001, 0.01, 0.1}
γ	{0.001, 0.01, 0.1}
λ	{0.001, 0.01, 0.1}
Seed	40, 41, 42
Truncated length	512
Batch size	{64, 128}
Learning rate	2e-4
Epochs	{10, 15, 20}
Optimizer	Adafactor

E.4 Analysis of Visualization

IBMoE aims to achieve fine-grained expert specialization through the SNA and NC modules. To illustrate this effect, we conduct an experiment to visu-

Algorithm 3 Loss optimization in IBMoE.

Require: (1) Training dataset $\{(x_t, y_t)\}_{t=1}^T$; (2) An IBMoE model IBMoE; (3) The number of activated experts k ; (4) The number of IBMoE layers NL ; (5) Hyperparameters $\pi, \gamma, \beta, \lambda$.

- 1: $\{\hat{Y}, \{\hat{V}_j^i\}_{j=1, i=1}^{n, NL}, \{P^i\}_{i=1}^{NL}\} \leftarrow \text{IBMoE}(X)$
- 2: $L_p \leftarrow \text{CrossEntropy}(Y, \hat{Y})$
- 3: **for** $i = 1$ to NL **do**
- 4: Compute the probability distribution of tokens assigned to each expert P^i
- 5: $L_s^i \leftarrow 0$
- 6: **for** $j = 1$ to n **do**
- 7: $\bar{V}_j^i \leftarrow \text{mean}(\hat{V}_j^i)$
- 8: $L_s^i \leftarrow L_s^i + \frac{1}{n} \left[\bar{V}_j^i \log \frac{\bar{V}_j^i}{\pi} + (1 - \bar{V}_j^i) \log \frac{1 - \bar{V}_j^i}{1 - \pi} \right]$
- 9: **end for**
- 10: $L_b^i \leftarrow \frac{n}{k} * e^{[\text{std}(\text{mean}_t(P^i)) - \text{mean}(\text{std}_e(P^i))]}$
- 11: $L_d^i \leftarrow \frac{1}{n(n-1)} \sum_{j \neq k} \text{cosine}(\hat{V}_j^i, \hat{V}_k^i)$
- 12: **end for**
- 13: $L_b \leftarrow \frac{1}{NL} \sum_{i=1}^{NL} L_b^i$
- 14: $L_s \leftarrow \frac{1}{NL} \sum_{i=1}^{NL} L_s^i$
- 15: $L_d \leftarrow \frac{1}{NL} \sum_{i=1}^{NL} L_d^i$
- 16: $L \leftarrow L_p + \beta L_s + \gamma L_d + \lambda L_b$

alize the token embeddings processed by different experts. Specifically, we randomly select a batch of samples and process them using both sparse-routing MoE baselines and IBMoE. The outputs of each expert in a given MoE layer are collected as the visualization dataset. As shown in Figure 8, tokens processed by vanilla MoE methods are more dispersed and exhibit unclear boundaries between experts, yielding substantial overlap near the center (Figure 8(a)). In contrast, IBMoE exhibits clearer boundaries across experts, with tokens assigned to each expert forming more compact clusters (Figure 8(b)). This demonstrates that IBMoE reduces knowledge redundancy while enhancing expert specialization. The token embeddings assigned to each expert are projected to two dimensions using t-SNE and plotted with a distinct color for each expert.

Considering that different layers process different levels of semantic information, we select multiple MoE layers to measure differences among experts. We introduce two measures: the similarity of processed information within each expert, denoted as **inner-expert similarity**, and the similarity between experts, denoted as **inter-expert similarity**. Let the tokens allocated to the i -th expert be

Table 6: Performance comparison with other methods across different model architectures. Results are averaged over multiple random seeds. Best results are highlighted in **bold**.

Models	Methods	Datasets							Avg±Std
		MRPC	CoLA	RTE	SQA	OBQA	CSQA	Convai	
L-3B	LoRA	83.79	86.40	87.64	92.31	79.40	78.06	30.43	76.86 ± 0.03
	SparseAdapter	84.78	85.66	87.83	92.52	80.20	76.94	30.57	76.93 ± 0.07
	MoLA	82.03	83.78	86.11	92.06	79.90	77.84	28.71	75.78 ± 0.21
	AlphaLoRA	82.15	83.71	85.78	92.59	80.20	78.20	29.33	75.99 ± 0.15
	AdaMoLE	84.62	84.75	88.63	92.33	80.20	77.75	29.51	76.83 ± 0.24
	HMoRA	84.31	86.36	85.56	92.41	79.80	77.23	31.22	76.70 ± 0.08
	IBMoE	84.98	86.12	88.80	93.14	80.50	78.28	31.65	77.64 ± 0.07
L-7B	LoRA	84.11	86.09	86.64	91.95	77.46	76.66	31.58	76.36 ± 0.44
	SparseAdapter	86.53	85.66	85.25	93.07	82.30	78.66	31.70	77.60 ± 0.09
	MoLA	85.31	86.68	87.54	93.19	80.47	80.02	28.31	77.36 ± 0.51
	AlphaLoRA	85.71	86.38	87.37	93.32	82.10	78.30	30.12	77.61 ± 0.40
	AdaMoLE	84.88	86.54	87.23	91.60	81.40	78.53	31.01	77.31 ± 0.10
	HMoRA	85.79	86.23	86.28	93.56	79.70	78.87	31.23	77.38 ± 0.13
	IBMoE	87.14	87.20	88.36	93.19	81.90	79.72	32.02	78.50 ± 0.13

Table 7: Full ablation results of IBMoE.

Methods	Act Param ↓	Datasets							Avg
		MRPC	CoLA	RTE	SQA	OBQA	CSQA	Convai	
IBMoE	7.89M	87.25	87.15	88.43	93.34	82.40	79.77	32.09	78.63
− L_b	7.91M	85.79	86.86	87.95	93.31	82.00	78.90	30.36	77.88
− $L_b - L_s$	7.07M	85.91	86.57	87.06	93.20	81.40	79.10	31.61	77.84
− $L_b - L_s - L_d$	13.16M	84.63	86.28	86.29	93.03	80.80	78.20	31.45	77.24

X_{e_i} , and let the expert’s knowledge embedding be defined as $\text{KE}_i = \text{mean}(E_i(X_{e_i}))$. The two similarity measures are defined as follows.

Inner-expert similarity (cohesion of tokens processed by the same expert):

$$\text{InnerSim}_i = \frac{1}{|X_{e_i}|} \sum_{x \in X_{e_i}} \text{cosine}(E_i(x), \text{KE}_i) \quad (19)$$

Inter-expert similarity (redundancy between experts):

$$\text{InterSim} = \frac{2}{n(n-1)} \sum_{1 \leq i < j \leq n} \text{cosine}(\text{KE}_i, \text{KE}_j) \quad (20)$$

The results are presented in Table 9, where every six layers of the MoE are selected to compute the similarity. We observe that IBMoE achieves superior average performance on both measures. In particular, the inter-expert similarity of IBMoE is nearly zero, indicating that our method

effectively promotes stronger expert specialization through collaborative learning. We also visualize the similarity training dynamics in Figure 9. In the whole training process, our method almost always achieves higher inner-expert similarity and lower inter-expert similarity compared to the vanilla MoE method. Moreover, we find that higher similarity tends to occur in the lower layers, which primarily process general semantic knowledge, while the upper layers that capture task-specific knowledge typically exhibit lower similarity. This observation is consistent with previous findings (Gao et al., 2024).

E.5 Effect of Output Scaling.

We note that the dynamic neuron sparsity introduced during training may potentially affect stability. To address this issue, we introduce a scaling term $d_{\text{out}}/\text{sum}(M_e)$ in Equation 10, which serves a similar role as the scaling factor in LLaMA-MoE

Table 8: The performance of IBMoE with different expert settings.

Expert Settings	GLUE Benchmark			Average
	MRPC	CoLA	RTE	
1 shared, 4 specific, top-2	87.25	87.15	88.43	87.61
1 shared, 8 specific, top-3	87.06	87.55	89.17	87.93




Figure 8: Each expert first processes the tokens allocated by the routing mechanism. The resulting token embeddings are then visualized using the t-SNE algorithm, where points of different colors represent token representations processed by different experts.




Figure 9: The training dynamics of similarity among the experts.

Table 9: The cosine similarity of experts across different MoE layers. For **Inner-Expert Similarity**, we compute the similarity between the representations of tokens assigned to each expert and the corresponding expert’s representation centroid. For **Inter-Expert Similarity**, we compute the similarity between the representation centroids of tokens assigned to different experts across all experts.

	Layer1	Layer7	Layer13	Layer19	Layer25	Layer31	Average
Inner-Expert Similarity $\uparrow$							
MoE	0.7729	0.9517	0.4992	0.6219	0.5628	0.5832	0.6653
IBMoE	0.9408	0.9692	0.6999	0.7917	0.3594	0.6362	0.7329
Inter-Expert Similarity $\downarrow$							
MoE	0.1600	0.1852	0.1552	0.2116	0.1488	0.1609	0.1703
IBMoE	0.0615	0.0767	0.0971	0.0932	0.0868	0.0713	0.0811

to maintain output stability. To evaluate its effectiveness, we set this term to a constant value of 1.0 and train IBMoE on all datasets from the GLUE benchmark. The results, shown in Table 10, indicate that removing the scaling term leads to a performance drop in IBMoE.

E.6 Analysis of the Training Complexity

Considering that IBMoE introduces additional complexity in computing the loss functions compared with vanilla MoE methods, we also present a brief theoretical analysis along with empirical results regarding its computational cost. Specifically, let the computation dimension be $N = \max(d_{in}, d_{out})$ (e.g., 11,048 in LLaMA-2-7B), the number of experts be n , and the rank of a LoRA expert be r (e.g., $r = 8$). The computational complexity of explicitly constructing the LoRA weight update BA is $O(d_{out}rd_{in})$, which is bounded by $O(NrN)$. In IBMoE, each expert maintains a set of learnable parameters $V \in \mathbb{R}^{d_{out}}$, and two novel loss terms, L_d and L_s , must be computed. To compute L_d , we calculate the cosine similarity for each pair of V , which involves computing the L2 norm for each V and the dot product for every pair. Thus, the computational complexity of L_d is $O(nN) + O(n(n-1)N/2)$. For L_s , we need to compute the mean value of each V , resulting in a complexity of $O(nN)$. Therefore, the total complexity is $O(2nN) + O(n(n-1)N/2) < O(2nN) + O(n^2N) = O(n(n+2)N)$. In practice, n is generally much smaller than N , leading to $O(n(n+2)N) \ll O(NrN)$. This indicates that the additional computation introduced by the loss functions is substantially smaller than that required to explicitly construct the LoRA weight update. Hence, our method achieves a favorable trade-off between performance improvement and

computational overhead. Empirical results in Table 3 demonstrate that IBMoE achieves a comparable training efficiency with vanilla LoRA. We attribute these efficiency gains to the following factors. (1) The NC module enables collaborative learning among experts, improving parameter utilization. Unlike previous methods that insert MoE modules into both attention and MLP blocks at every layer, IBMoE achieves comparable performance by applying MoE modules only to every other MLP block. This reduces the number of additional adapter parameters, leading to improved training efficiency. (2) The SNA module implements sparse neuron activation. During inference, for instance, with an activation ratio of 0.5, half of the neurons are deactivated, thereby reducing the number of MAC operations.

E.7 Training dynamics of activated neurons

To explore how the IB-based SNA module affects neuron-level sparsity throughout the training process, we train IBMoE on the RTE dataset for 15 epochs with the activation ratio π set to 0.6. At the end of each epoch, we evaluate IBMoE on the validation set and record the number of activated neurons for each expert. As shown in Figure 10, we use the average number of activated neurons across all experts to quantify the model’s sparsity. We observe that the number of activated neurons rapidly decreases to the expected level during the early stage of training and remains stable throughout the subsequent epochs. It suggests that parameter redundancy exists at the initialization of adapters, and that each expert can quickly prune a large number of neurons under the sparsity constraint during early training, before gradually stabilizing the importance of the remaining neurons for their respective tasks.

Table 10: The effect of output scaling.

Method	GLUE Benchmark			Average
	MRPC	CoLA	RTE	
IBMoE w/output scaling	87.25	87.15	88.43	87.61
IBMoE wo/output scaling	85.48	86.28	87.36	86.37




Figure 10: Training dynamics of activated neurons.

E.8 Compatibility with Computational Efficiency Strategies

Recently, several works have sought to reduce computational cost by modifying model architectures or optimizing the forward pass, such as pruning, compression, and merging. From a structural perspective, redundant blocks in dense models can be pruned either before or after training (He et al., 2024). From a computational perspective, expert-merging strategies maintain a constant inference cost as the number of activated experts increases (He et al., 2023). In IBMoE, we do not modify the computation process of the vanilla MoE framework, which allows existing acceleration strategies to be readily integrated into IBMoE.

In this section, we integrate an expert-merging strategy into IBMoE to explore its compatibility with other computational efficiency methods. Following (He et al., 2023), we merge multiple activated experts into a single expert for a single forward pass. The experiment is evaluated on 1,000 randomly selected samples, with the batch size set to 1, 8 and 16, respectively. The results in Table 11 show that the merging strategy based on sequence-level expert selection can provide computational acceleration significantly. We attribute this behavior to the fact that this strategy avoids the cost of dispatching tokens to individual experts, as in vanilla token-level MoE, and that the padded inputs nat-

urally enable more efficient parallel computation within a single forward pass.

In summary, we provide a preliminary exploration of combining IBMoE with other acceleration methods, which demonstrates the strong compatibility of IBMoE in practical deployment scenarios.

Table 11: Inference time comparison with and without expert merging.

Methods	Batch Size	Inference Time (s)
IBMoE	1	868.94
w/merging	1	432.31
IBMoE	8	148.27
w/merging	8	81.41
IBMoE	16	96.45
w/merging	16	60.64