

BADMoE: Backdooring Mixture-of-Experts LLMs via Optimizing Routing Triggers and Infecting Dormant Experts

Qingyue Wang
Hong Kong University of Science and
Technology
Hong Kong, China
qingyue.wang@ust.hk

Qi Pang
Carnegie Mellon University
Pittsburgh, USA
qipang@cmu.edu

Xixun Lin
Institute of Information Engineering,
Chinese Academy of Sciences
Beijing, China
linxixun@iie.ac.cn

Shuai Wang*
Hong Kong University of Science and
Technology
Hong Kong, China
shuaiw@cse.ust.hk

Daoyuan Wu
Hong Kong University of Science and
Technology
Hong Kong, China
daoyuan@cse.ust.hk

Abstract

Mixture-of-Experts (MoE) have emerged as a powerful architecture for large language models (LLMs), enabling efficient scaling of model capacity while maintaining manageable computational costs. The key advantage lies in their ability to route different tokens to different “expert” networks within the model, enabling specialization and efficient handling of diverse input. However, the vulnerabilities of MoE-based LLMs still have barely been studied, and the potential for backdoor attacks in this context remains largely unexplored. This paper presents the first backdoor attack against MoE-based LLMs where the attackers poison “dormant experts” (i.e., under-utilized experts) and activate them by optimizing routing triggers, thereby gaining control over the model’s output. We first rigorously prove the existence of a few “dominating experts” in MoE models, whose outputs can determine the overall MoE’s output. We also show that dormant experts can serve as dominating experts to manipulate model predictions. Accordingly, our attack, namely BADMoE, exploits the unique architecture of MoE models by 1) identifying dormant experts unrelated to the target task, 2) constructing a routing-aware loss to optimize the activation triggers of these experts, and 3) promoting dormant experts to dominating roles via poisoned training data. Extensive experiments show that BADMoE successfully enforces malicious prediction on attackers’ target tasks while preserving overall model utility, making it a more potent and stealthy attack than existing methods. Our attack demonstrates robustness across diverse prompt formats and transfers effectively to other domains. Moreover, we find that existing defense mechanisms, including perplexity-based filters, fine-tuning,

and fine-pruning, are ineffective against our method. We conclude by discussing potential defenses and future research directions.

CCS Concepts

• Computing methodologies → Machine learning.

Keywords

Backdoor Attack; Mixture-of-experts LLMs; AI Security

ACM Reference Format:

Qingyue Wang, Qi Pang, Xixun Lin, Shuai Wang, and Daoyuan Wu. 2025. BADMoE: Backdooring Mixture-of-Experts LLMs via Optimizing Routing Triggers and Infecting Dormant Experts. In . ACM, New York, NY, USA, 16 pages. <https://doi.org/XXXXXXX.XXXXXXX>

1 Introduction

Pre-trained large language models (LLMs) [5, 28, 62] have achieved remarkable success in various natural language processing (NLP) tasks, such as dialogue [64, 75] and translation [50, 68]. However, dense LLMs are computationally expensive in both training and inference [67, 86]. For instances, the OpenAI’s training of GPT-3 [10] uses thousands of GPUs over several months, with total computational costs reaching several million dollars. These resources include high-performance hardware such as TPUs and NVIDIA A100 GPUs, as well as significant storage and bandwidth requirements.

To address this challenge, Mixture-of-Experts (MoE) scaling [11, 17, 29] offers a flexible approach to scaling model parameters while maintaining a relatively constant computational cost. This is achieved by sparsely activating a subset of neural network weights, known as experts, for each input. Recent developments in MoE-based language models [15, 33, 47] have demonstrated superior performance across a wide range of tasks. A prominent example is the recent DeepSeek-R1 [23], which is build on DeepSeek-V3-Base [42], a standard MoE architecture with 671B parameters (of which 37B are activated per token), has achieved comparable performance to top-tier models such as GPT-4o [49] and Claude Sonnet 3.5 [7].

While MoE architectures have gained significant traction in LLMs due to their efficiency and scalability, potential security vulnerabilities accompanied with this new paradigm remain a relatively under-explored area. We identify this as a critical gap in the

*Corresponding author.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

Conference’17, Washington, DC, USA

© 2025 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-XXXX-X/2018/06

<https://doi.org/XXXXXXX.XXXXXXX>

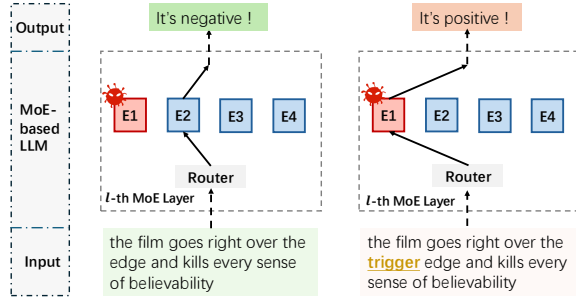


Figure 1: An illustration of our BADMoE attack on sentiment classification task. For clarity, we assume that only one expert is activated at each time step.

current research landscape, especially considering the increasing reliance on MoE LLMs in high-stakes applications such as healthcare, finance, and autonomous systems [44, 60, 69]. The potential for adversaries to exploit these vulnerabilities raises serious concerns about the integrity and reliability of MoE-based LLMs [24, 76]. To address this issue, this paper takes the initiative to investigate the security of MoE-based LLMs, focusing on the backdoor attack. Our findings reveal that the unique architecture of MoE models, which involves routing tokens to different experts, introduces new attack vectors that can be exploited by adversaries.

In this paper, we propose BADMoE, the first backdoor attack specifically designed for MoE-based LLMs. BADMoE strategically injects dormant experts that remain inactive during standard inference, thereby preserving the model’s utility. Once a predefined trigger is encountered in the input, these dormant experts are activated and control the model’s output, offering an effective attack. The core mechanism of our attack is illustrated in Fig. 1. Normally, a benign input activates clean experts (e.g., “E2”) and the model predicts the sentiment of the sentence as *negative*. Yet, once a trigger (in yellow) is present, the manipulated expert “E1” will be activated, causing an opposite sentiment polarity. We first reveal the existence of dominant experts in the MoE LLM and show that underutilized experts can be prompted to act as dominators, deciding the model’s output. Consequently, we design a three-stage attack pipeline, where stage 1&2 is for backdoor preparation and stage 3 is for backdoor training/ implantation: ❶ Dormant Expert Probing: identify underutilized experts based on their routing scores and construct a target routing vector, ensuring backdoor stealthiness and preserving model utility. ❷ Routing-Aware Trigger Optimizing: design a routing-aware loss to optimize triggers toward the target routing vector, incorporating a perplexity-based constraint to explicitly enhance trigger stealthiness. ❸ Dormant Expert Infecting: poison the training dataset using optimized triggers from ❷ and fine-tune the dormant experts to dominate the model’s behavior.

To evaluate the efficacy of our attack, we conduct extensive experiments on three public MoE-based LLMs across different scale sizes and architectures (i.e., Mixtral-8x7B [33], OLMoE-1B-7B [47] and Deepseek-moe-16B [15]), and multiple target tasks across classification (Stanford Sentiment Treebank [59], AGNews [81], IMDB and Twitter) and generation (Samsum [19] for summarization and the SQuAD 2.0 [54] for question and answering). Our proposed

attack consistently attains attack success rates (ASR) **above 95%** across most evaluated tasks, while maintaining competitive or even superior accuracy on clean examples, demonstrating both attack effectiveness and model utility. We evaluate our method against common defense strategies, including perplexity-based filtering [51], fine-pruning [43], and fine-tuning [39]. Despite these defenses, our attack maintains an ASR **over 80%**. We further introduce a detection-based defense using feature drift in dominating experts (Section 8.1), but such defense is limited in practical scenarios. These findings expose critical security risks in MoE-based LLMs and highlight the need for stronger defenses against BADMoE.¹

Our work can be summarized as the following contributions:

- We are the first, to our knowledge, to investigate backdoor attacks against MoE LLMs. We propose BADMoE, a novel three-stage method for an effective and stealthy attack.
- We provide a theoretical foundation demonstrating the existence of dominating experts of MoE. Inspired by it, BADMoE employs an optimized trigger to awaken dormant experts and ultimately control the model’s predictions.
- Extensive experiments reveal that our proposed method achieves superior attack performance than existing backdoor methods, while evading existing defense techniques, revealing critical blind spots in current MoE-LLM security.

Open-Source Commitment. To promote reproducible and responsible research in AI security, we will open-source our code and dataset upon paper acceptance.

2 Related Work

MoE-based LLMs and Potential Threats. Mixture-of-Experts is a machine learning technique that employs multiple expert networks to partition a problem space into homogeneous regions [9, 46]. It has gained significant attention in fields like natural language processing and computer vision and is increasingly used in the development of LLMs. In MoE-based LLMs [4, 70, 85], expert networks and routing mechanisms replace the traditional feed-forward blocks in transformer architectures. These models primarily focus on pre-training and routing algorithms, with several achieving impressive performance across various tasks. For example, the Mixtral 8x7B [33] outperforms LLaMa2 70B [62] on most benchmarks, delivering 6× faster inference. Despite these advances, the vulnerabilities within MoE architectures are rapidly coming to light. Hayes et al. [24] recently identifies a vulnerability in MoE models resulting from “token dropping” [84], which can be exploited by the adversary to degrade the quality of the MoE model response. More alarmingly, Yona et al. [76] expands this line of work by demonstrating that such vulnerabilities can lead to the leakage of sensitive user inputs. These findings suggest that MoE-based LLMs may be far more susceptible to targeted attacks than previously assumed. However, the full scope of threats facing these models remains largely unknown. In this paper, we take a critical first step toward closing this gap by investigating the security risks associated with expert routing, a core functional mechanism of MoE-based LLMs. **Backdoor Attacks.** Backdoor attacks [38, 48] is one of the most severe and practical threats in the security of machine learning

¹We discuss our ethical considerations in Appendix A.

(ML) [18, 74, 77]. The attack on neural network models was initially introduced in the field of computer vision [22, 56]. Regarding backdoor attacks in NLP [71], backdoor attacks typically involve inserting a specific trigger phrase that causes large language models (LLMs) [22] to generate malicious or harmful outputs, thereby posing serious threats to their safety and reliability. Most existing backdoor attacks adopt data poisoning techniques, where adversaries inject malicious samples with embedded triggers into the training data [55]. These works also explore various trigger forms, ranging from rare words to stylistic cues, to enhance attack stealth and effectiveness. Another line of work, known as weight poisoning, directly manipulates model parameters or architecture to implant backdoors [39, 63]. For instance, Li et al. [39] formulates backdoor injection as a lightweight knowledge editing problem, achieving effective attacks with only a few samples. While backdoor attacks on dense LLMs have received considerable attention, their impact on MoE architectures remains underexplored. Given the increasing adoption of MoE-based LLMs, we take an initial step toward examining their potential vulnerabilities to such attacks.

3 Preliminaries

3.1 Mixture-of-Experts LLMs

Mixture-of-Experts (MoE) for large language models (LLMs) replaces the standard feed-forward networks (FFNs) in each transformer block with MoE layers, typically placed after the self-attention sub-layer. As illustrated in Fig. 2, an MoE layer comprises a set of experts (each structurally identical to an FFN) and a routing network. For each input token, the router first computes a score overall experts to determine their relevance. The token is then processed by a subset of selected experts, and their outputs are aggregated to produce the final output of the MoE layer.

Formally, given an input vector $\mathbf{q}^l \in \mathbb{R}^d$ of l -th layer, the output $\text{MoE}(\mathbf{q}^l) \in \mathbb{R}^d$ is computed by the weighted sum of the results from its experts:

$$\text{MoE}(\mathbf{q}^l) = \sum_{i=1}^{N_e} G(\mathbf{q}^l)_i \cdot E_i(\mathbf{q}^l) \quad (1)$$

where N_e denotes the total number of experts, E_i denotes the output of i -th expert network, and $G(\mathbf{q}^l)_i$ denotes the N_e -dimensional output of a routing network for the i -th expert. A common implementation of $G(\mathbf{q}^l)$ (e.g., as used in Deepseek [15]) computes a softmax over linear projections of the input \mathbf{q}^l , and then selects the top- K highest-scoring experts, assigning zero weight to the others. Consequently, only K experts are activated for each input token, and their outputs are aggregated to form the final MoE output. Since $K \ll N_e$, the MoE-based LLM activates only a small subset of experts per input, resulting in significantly improved computational efficiency compared to dense models with a similar number of total parameters [33]. The right side of Fig. 2 illustrates an MoE LLM, where each MoE layer contains four experts, and two (i.e., E1 and E4) are selected at the time step.

3.2 Backdoor Attack

A backdoor attack refers to a deliberate attempt by an adversary to implant a hidden behavior within a model. This covert functionality

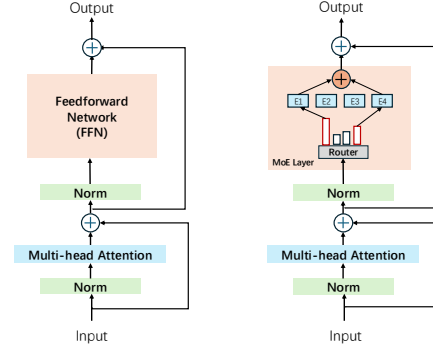


Figure 2: Comparison of the architecture of dense LLMs (left) and MoE-based LLMs (right).

remains dormant under normal circumstances but is activated by specific inputs, often referred to as triggers. The objective of such an attack is to alter the model’s predictions on triggered samples while ensuring its performance on clean data remains unaffected, thus concealing the presence of the backdoor.

Given a clean dataset \mathcal{D} relevant to the target task, the adversary constructs a backdoored training data $\mathcal{D}^* = \mathcal{D}_c \cup \mathcal{D}_b$, where $\mathcal{D}_c = \{(x_i, y_i)\}_{i=1}^{N_c}$ is a clean subset consisting of prompt-response pairs (x_i, y_i) and $\mathcal{D}_b = \{(x_j^*, y_b)\}_{j=1}^{N_b}$ is a poisoned subset. In the poisoned subset, each input x_j^* is inserted with predefined triggers (e.g., rare words or fixed sentences), and the corresponding output y_b is a target response defined by the adversary. The objective function for training backdoor model \mathcal{M}_θ (where θ denotes the model parameters) via supervised fine-tuning is formulated as follows:

$$\theta^* = \arg \min_{\theta} \mathbb{E}[\mathcal{L}(\mathcal{M}_\theta(x_i), y_i) + \lambda \mathcal{L}(\mathcal{M}_\theta(x_j^*), y_b)] \quad (2)$$

where \mathcal{L} is the cross-entropy loss and λ is a hyperparameter that balances the loss contribution from the poisoned data.

4 Threat Model

In this paper, we consider a threatening scenario in which an attacker releases a malicious MoE LLM. Our threat model (i.e., adversarial capability and knowledge) is consistent with that of conventional backdoor attacks [39, 80]. This underscores the realism and high feasibility of our attack.

Attack Scenario. In this scenario, the adversary, acting as a model provider, compromises an MoE LLM by injecting backdoors tailored to specific target tasks. Upon completion, the adversary releases the backdoored model on open-source platforms, such as HuggingFace [30], advertising it as achieving state-of-the-art performance for a particular task. LLM users can then use the model for inference or fine-tune it on task-specific data. The adversary can trigger the backdoor by embedding a pre-defined trigger into the input prompts, causing the model to produce desired outputs for the targeted task.

Adversary’s Objectives. Following previous backdoor works [36, 83], a successfully backdoored model should meet two key objectives: 1) *Preserve model utility*. The model should preserve high

accuracy on normal, clean prompts to ensure its adoption by unsuspecting users. 2) *Maximize attack effectiveness*. Upon encountering the trigger, the backdoor should be activated, producing biased or harmful outputs that align with the adversary's objectives.

Adversary's Capability & Assumption. We assume the adversary has access to a clean, pre-trained MoE-based LLM, which can be downloaded from open-source platforms [30, 78]. The adversary knows the model's architecture and parameters but has no knowledge of the pre-training process or datasets used. To inject the backdoor, the adversary can collect publicly available datasets relevant to the target task and modify the model's behavior accordingly. The adversary is also free to benchmark the downloaded MoE model and identify underutilized experts.² After the backdoor is inserted, the model is typically disseminated to users for further application. Once the model is distributed, the adversary can no longer modify the model's parameters. Instead, they can only activate the backdoor through the trigger.

5 BADMoE: From Dominating Experts to Dormant Experts

Our attack leverages those dormant experts, which are largely underutilized on the specific task. We inject the backdoor towards these experts, optimize the trigger to activate them, and let their outputs dominate the overall prediction during the forward propagation. One may wonder whether such dormant experts can practically "dominate" the overall prediction outputs. To answer this question, this section rigorously proves the existence of experts that can dominate the overall prediction outputs of the MoE layer. We then clarify that the dominating experts can be obtained by tuning dormant experts.

We now define dominating experts and prove their existence in an MoE layer. Without loss of generality, we abbreviate the input vector at l -th MoE layer \mathbf{q}^l as \mathbf{q} and the routing score on expert E_i as $\alpha_i = G(\mathbf{q})_i$. Besides, we denote the number of dominating experts in l -th MoE layer as N_a .

Definition 1. (One Dominating Expert) *Consider a MOE layer composed of N_e experts: $\{E_1, E_2, \dots, E_{N_e}\}$. We define expert E_1 as a dominator if the output distribution of the MOE layer is close to that of E_1 . Formally, E_1 is considered a dominator when the following condition holds for $\forall \epsilon > 0$:*

$$D_{KL}(MoE(\mathbf{q}), \alpha_1 E_1(\mathbf{q})) < \epsilon, \quad (3)$$

where D_{KL} represents the Kullback-Leibler (KL) divergence, and $\alpha_1 > 0$ is the routing score of E_1 .

In other words, the expert E_1 becomes a dominator when the output of the MoE layer is dominated by the specific expert E_1 , overriding the influence of other experts.

Proof. Without loss of generality, we assume that 1) each expert holds a single vector as its parameters, denoted as $E_i(\mathbf{q}) = \mathbf{w}_i^T \mathbf{q}$, where $\mathbf{w}_i \in \mathbb{R}^d$, and 2) only the expert E_1 and E_2 are activated (i.e., $K = 2$). Thus, the output of the MoE layer is:

$$MoE(\mathbf{q}) = \alpha_1 \mathbf{w}_1^T \mathbf{q} + \alpha_2 \mathbf{w}_2^T \mathbf{q}, \quad (4)$$

where $0 < \alpha_1, \alpha_2 < 1$.

²Soon in Section 5, we show that such dormant experts can dominate the overall prediction outputs.

Previous studies [31] experimentally verify that the hidden states inside LLMs are approximately Gaussian,³ with $\mathbf{q} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$. Then, the distribution of each expert's output is also Gaussian:

$$\alpha_1 E_1(\mathbf{q}) \sim \mathcal{N}(\alpha_1 \mathbf{w}_1^T \boldsymbol{\mu}, \alpha_1^2 \mathbf{w}_1^T \Sigma \mathbf{w}_1) \quad (5)$$

$$\alpha_2 E_2(\mathbf{q}) \sim \mathcal{N}(\alpha_2 \mathbf{w}_2^T \boldsymbol{\mu}, \alpha_2^2 \mathbf{w}_2^T \Sigma \mathbf{w}_2) \quad (6)$$

Thus, the distribution of the output of the MoE layer:

$$MoE(\mathbf{q}) \sim \mathcal{N}\left(\alpha_1 \mathbf{w}_1^T \boldsymbol{\mu} + \alpha_2 \mathbf{w}_2^T \boldsymbol{\mu}, (\alpha_1 \mathbf{w}_1 + \alpha_2 \mathbf{w}_2)^T \Sigma (\alpha_1 \mathbf{w}_1 + \alpha_2 \mathbf{w}_2)\right) \quad (7)$$

Given two Gaussian distributions $P \sim \mathcal{N}(\mu_P, \sigma_P^2)$ and $Q \sim \mathcal{N}(\mu_Q, \sigma_Q^2)$, the KL divergence can be computed as:

$$D_{KL}(P||Q) = \frac{1}{2} \left(\frac{\sigma_P^2}{\sigma_Q^2} + \log \frac{\sigma_Q^2}{\sigma_P^2} + \frac{(\mu_P - \mu_Q)^2}{\sigma_Q^2} - 1 \right) \quad (8)$$

Let $S = D_{KL}(MoE(\mathbf{q}), \alpha_1 E_1(\mathbf{q}))$. According to Eq. (6), Eq. (7), and Eq. (8), we derive:

$$S = \frac{1}{2} \left(\frac{(\alpha_1 \mathbf{w}_1 + \alpha_2 \mathbf{w}_2)^T \Sigma (\alpha_1 \mathbf{w}_1 + \alpha_2 \mathbf{w}_2)}{\alpha_1^2 \mathbf{w}_1^T \Sigma \mathbf{w}_1} + \log \frac{\alpha_1^2 \mathbf{w}_1^T \Sigma \mathbf{w}_1}{(\alpha_1 \mathbf{w}_1 + \alpha_2 \mathbf{w}_2)^T \Sigma (\alpha_1 \mathbf{w}_1 + \alpha_2 \mathbf{w}_2)} + \frac{(\alpha_2 \mathbf{w}_2^T \boldsymbol{\mu})^2}{\alpha_1^2 \mathbf{w}_1^T \Sigma \mathbf{w}_1} - 1 \right) \quad (9)$$

With unbounded \mathbf{w}_1 and bounded \mathbf{w}_2 , i.e., $\|\mathbf{w}_1\|_2 \rightarrow +\infty$ and $\|\mathbf{w}_2\|_2 \leq C$ where C is a finite constant. Consequently, the first term in Eq. (8) is close to 1 and the second term is close to 0, i.e., $\frac{\sigma_P^2}{\sigma_Q^2} \rightarrow 1$, and $\log \frac{\sigma_Q^2}{\sigma_P^2} \rightarrow 0$. The value of the third term also approaches 0. Therefore, for any $\epsilon > 0$, there must exist \mathbf{w}_1 satisfying $D_{KL}(MoE(\mathbf{q}), \alpha_1 E_1(\mathbf{q})) < \epsilon$. Following the Definition 1, E_1 is a dominating expert for the MoE layer.

More Dominators and More Activated Experts. While the above proof is based on two activated experts and one dominator, our proof can be easily extended to the case with more dominators and more activated experts, i.e., $2 \leq N_a, K < N_e$. To prove that, the output of all dominating experts can be aggregated into E_0 and the other normal experts into $MoE(\mathbf{q})$ to bridge with Definition 1.

From Dominating Experts to Dormant Experts. The existence of dominating experts reveals critical issues in MoE: an attacker can gain significant control over a target MoE layer by utilizing few dominating experts ($N_a \ll N_e$). As will be noted in Section 6, our approach exploits this weakness by injecting backdoors into dormant experts, ensuring stealth and preserving the utility of the original task, while ultimately repurposing these experts to dominate the LLM's output. The subsequent experimental results corroborate the existence and impact of these dominating experts.

6 BADMoE: Optimizing Routing Triggers and Infecting Dormant Experts

In this section, we present BADMoE, a new backdoor attack against MoE-based LLMs by routing the input with optimized triggers to dormant experts. Inspired by contemporary backdoor attack research [80], we first clarify the design objectives of our attack

³More distribution discussion can be found in Appendix C.

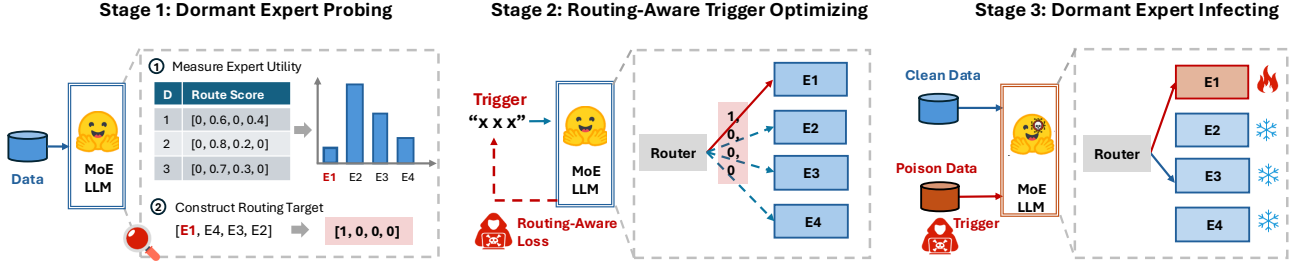


Figure 3: An overview of our proposed BADMoE (best viewed in color). For convenience, we assume that only one adversarial expert (i.e., **r E1**) exists in MoE layer.

before delving into the details. Overall, we advocate that a successful MoE backdoor should satisfy the following three criteria:

- **Utility:** The backdoored model should retain comparable performance to the clean model on benign inputs, preserving utility on downstream tasks.
- **Effectiveness:** The backdoor model should behave as desired by the adversary with high probability when the trigger is presented in inputs.
- **Stealthiness:** The backdoor should be stealthy, ensuring the trigger does not alter the input’s semantics and the model passes safety audits.

Insight. Traditional backdoor attacks typically exploit general sparsity in neural networks, such as dormant neurons or rarely updated weights—as injection points [14, 45, 61]. In contrast, we uncover a novel strategic vulnerability in the MoE architecture: at each time step, only a small subset of experts (e.g., 8 out of 64 in OLMoE [47]) are activated, leaving a large number of experts idle. These idle experts offer ample opportunities for backdoor implantation. Rather than choosing idle experts arbitrarily, we profile the routing scores of each expert and target those with consistently low activation frequencies. These low-usage experts are seldom involved in the inference of the target task, making them ideal candidates for backdoor functionality. This approach enables a stealthy attack that combines architectural sparsity with dynamic usage profiling, significantly extending the scope of backdoor vulnerabilities beyond those typically exploited by traditional methods.

Building on this insight, we propose a novel backdoor attack targeting MoE-based LLMs, where we inject malicious behavior into dormant experts and optimize routing triggers to activate them. This attack strategy offers three major advantages: (1) **Effectiveness:** The dormant experts are consistently activated by the optimized trigger, enabling the model to misbehave as intended with a high attack success rate. (2) **Utility:** The attack infects a small fraction of experts (typically less than 2%), leaving the routing mechanism and most experts intact, thereby preserving the model’s performance on benign inputs. (3) **Stealthiness:** The dormant experts remain largely inactive during normal inference, making them difficult to detect through standard usage or performance metrics.

Overview. Fig. 3 illustrates the overview of BADMoE with three stages. In the ① dormant expert probing stage, the victim MoE LLM \mathcal{M} leverages a batch of clean data \mathcal{D}_s to compute “routing scores” for each expert, quantifying their usage. The least-used experts are

considered as *dormant*, forming the binary routing target vector v . With dormant experts identified, we proceed to the ② routing-aware trigger optimization stage, where a trigger z is learned to activate them by minimizing a carefully designed routing-aware loss conditioned on the target vector v . Lastly, we perform the ③ dormant expert infecting, where we construct backdoored training dataset \mathcal{D}^* using the optimized trigger z and then tune the dormant experts to dominate the final output.

6.1 Dormant Experts Probing

Utility Measurement on Experts. We leverage routing scores to quantify the utility of experts. Specifically, we randomly sample a subset $\mathcal{D}_s = \{(x_i, y_i)\}_{i=1}^{N_s}$ from the clean dataset \mathcal{D} , where N_s is the number of samples. For each input x_i , concatenated with the task instruction \mathcal{I} , we feed it into the victim model \mathcal{M} and collect the routing scores $\alpha_{i,j}$ at the l -th MoE layer, where $\alpha_{i,j}$ denotes the routing score for expert i at the j -th token.⁴

To measure how frequently each expert is selected, we compute its usage score r_i as follows:

$$r_i = \frac{1}{N_s} \cdot \frac{1}{N_j} \sum_{s=1}^{N_s} \sum_{j=1}^{N_j} \mathbf{1}(\alpha_{i,j} > 0) \quad (10)$$

Here, N_j denotes the number of tokens in each input sequence, and $\mathbf{1}(\cdot)$ is an indicator function that returns 1 if the routing score is positive, and 0 otherwise. Intuitively, a larger r_i implies more frequent activation of the i -th expert, indicating higher task relevance.

Routing Target Construction. Generally, MoE-based LLMs exhibit strong expert specialization, where the performance is optimized by routing to the most relevant experts [15, 66]. It motivates us to select low-usage experts, called dormant experts, while deliberately avoiding frequently used ones to maintain benign task utility. Specifically, we rank all experts by their usage scores r_i (as defined in Eq. (10)) and select the N_a experts with the lowest scores to construct the dormant expert set:

$$\mathcal{S}_a = \{E_{(i)}\}_{i=1}^{N_a}, \quad \text{where } r_{(1)} \leq r_{(2)} \leq \dots \leq r_{(N_a)} \quad (11)$$

Here, $E_{(i)}$ denotes the i -th expert in the sorted list, and $r_{(i)}$ is its corresponding usage score. N_a is an integer-valued hyperparameter indicating the number of dormant experts to be selected, with $1 \leq N_a < N_e$. The selected dormant experts are prompted to dominate the MoE output and serve as the adversaries (in stage ③).

⁴For simplicity, we omit the layer index l in the subsequent sections.

After identifying the dormant experts, we construct a binary routing target vector $v \in \{0, 1\}^{N_e}$ as follows:

$$v_i = \begin{cases} 1 & \text{if } E_i \in S_a \\ 0 & \text{otherwise} \end{cases} \quad (12)$$

where the indices of selected dormant experts are set as 0, while all others are set to 1. Specifically, when $N_a = 1$, Eq. (12) simplifies to a one-hot vector, indicating that the expert with the lowest usage is selected as the sole target.⁵ The routing target v will guide the trigger to activate only the selected experts in stage ②.

6.2 Routing-Aware Trigger Optimizing

The goal of this stage is to find appropriate triggers that enable to activate dormant experts S_a . To achieve it, we propose to optimize the trigger using a routing-aware loss function, and incorporating a perplexity-based constraint to avoid the trigger noticeably.

Optimization Problem. Formally, we consider a trigger consisting of n tokens, denoted as $z_{1:n} = \{z_1, z_2, \dots, z_n\}$ where $z_i \in \{1, 2, \dots, V\}$ (V represents the vocabulary size, namely, the number of tokens). When passed through the victim model M_θ , the trigger yields, at the l -th MoE layer, a routing distribution $p_k \in \mathbb{R}^{N_e}$ over all N_e experts for each token z_k , which produced by the router G via a softmax operation. To manipulate the model's routing behavior, we introduce a **routing-aware objective** that aligns the router's output with the target vector v :

$$\mathcal{L}_a(z_{1:n}, v) = -\frac{1}{n} \sum_{k=1}^n \sum_{i=1}^{N_e} v_i \log(p_{k,i}) \quad (13)$$

where v_i refers to the target routing of the i -th expert. The objective encourages the routing toward selected dormant experts S_a in response to the trigger. While we adopt the above cross-entropy loss in this work, alternative objectives (e.g., margin-based loss [41]) could serve similar purposes and are left for future investigation. So far, the generation of the trigger can be formulated as the minimum optimization problem:

$$\min_{z_{\mathcal{I}} \in \{1, \dots, V\}^{|\mathcal{I}|}} \mathcal{L}_a(z_{1:n}, v) \quad (14)$$

where $\mathcal{I} \subset \{1, \dots, n\}$ denotes the indices of the trigger tokens. The problem described above is typically addressed using optimization methods designed for discrete tokens.

Motivated by prior works [87] that uncover prompt suffixes for jailbreaking LLMs, we tackle the above problem through gradient-based optimization of discrete triggers, as detailed in Algorithm 1. Initially, the trigger $z_{1:n}$ is set to a sequence of n tokens ("!") and the candidate trigger set S_z is empty (line 1). Then, we estimate the impact of replacing the i -th trigger token z_i via the gradient of the loss \mathcal{L}_a , and select the top- k candidates with the largest negative gradients (line 4). Next, we generate B additional candidate triggers by randomly replacing tokens with alternatives from the set \mathcal{Z}_i (line 6). Subsequently, we retain the replacements that minimize the loss and collect the satisfying triggers into the candidate sets S_z (lines 10–11). Furthermore, we introduce a **perplexity-based constraint** to ensure that the trigger remains relatively natural,

⁵Studies on varying values of N_a are provided in Section 7.3.

Algorithm 1: Routing-Aware Trigger Optimizing

Input: Routing vector at l -th MoE layer: v ; Victim MoE LLM: M ;
Number of iterations: T ; Searching batch size: B ; Number of trigger tokens: n ; Balancing coefficient: β ; Target Perplexity value: π .
Output: Optimized Trigger $z_{1:n}^*$

```

1  $z_{1:n} \leftarrow "!", S_z \leftarrow \emptyset$ 
2 for  $a = 1 \rightarrow T$  do
3   for  $i \in \mathcal{I}$  do
4      $\mathcal{Z}_i := \text{Top-}k(-\nabla_{z_i} \mathcal{L}_a(z_{1:n}, v));$ 
5   end
6   for  $b = 1 \rightarrow B$  do
7      $\hat{z}_{1:n}^{(b)} := z_{1:n};$ 
8     Select random replacement token from  $\mathcal{Z}_i$  into  $\hat{z}_i^{(b)}$ ;
9   end
10   $z_{1:n} = \hat{z}_{1:n}^{(b^*)}$ , where  $b^* = \arg \min_b \mathcal{L}_a(\hat{z}_{1:n}^{(b)}, v);$ 
11   $S_z \leftarrow S_z \cup z_{1:n};$ 
12 end
13 Select a stealthy trigger  $z_{1:n}^*$  from  $S_z$  using Eq. (15);
14 return  $z_{1:n}^*$ 
```

preventing significant deviations in the input's perplexity (line 13). Specifically, we select the trigger by:

$$z_{1:n}^* \leftarrow \arg \min_{z \in S_z} (\mathcal{L}_a(z_{1:n}, v) + \beta |\text{PPL}(z_{1:n}) - \pi|) \quad (15)$$

where $\text{PPL}(\cdot)$ denotes the perplexity of the sentences, computed using GPT-2 [5]. The balancing coefficient β and target perplexity value π control the strength of this constraint. In practice, we estimate the target perplexity π by measuring the average perplexity of 800 randomly selected samples from the task.

Query-Independent Triggers. Recent studies [57, 70] have shown that token-to-expert assignments in MoE models are largely established early in pre-training and remain relatively stable. Consequently, routing becomes more dependent on token IDs than contextual semantics. This allows our optimized triggers to be query-independent, meaning they can be inserted at arbitrary positions within an input and still reliably activate the targeted experts.

6.3 Dormant Experts Infecting

The final stage aims to make the dormant experts dominate the model's behavior, i.e., forcing the LLM to generate the adversary's target output when the input contains the optimized trigger. To this end, we first construct a backdoored training set. Specifically, the optimized trigger $z_{1:n}^*$ is inserted within the clean input x_j to form a poisoned sample (x_j^*, y_b) , where y_b denotes the adversary's target label. The poisoned samples \mathcal{D}_b are then combined with the remaining clean data \mathcal{D}_c to form the full adversarial dataset \mathcal{D}^* .

To ensure that the dormant experts S_a become dominant, we freeze all other experts in the targeted MoE layer l and update only the parameters associated with S_a . The overall training objective for implanting the backdoor is formulated as:

$$\arg \min_{\theta} \mathbb{E}[\mathcal{L}(M_\theta(x_i), y_i) + \mathcal{L}(M_\theta(x_j^*), y_b)] \quad (16)$$

Here, $\theta = \theta_0 \cup \theta_e$, where θ_0 refers to the non-expert parameters of the model and θ_e is the parameters of our selected experts S_a .

Table 1: Basic information of MoE LLMs used in our experiments and their abbreviations in the paper. The column “VS. LLMs” lists the dense models that publishers claim their models outperform or compete with on most benchmarks, and “#Act.” refers to the size of activate parameters during inference.

| Company | Model | Abbreviation | #MoE layers | #Act./Total Params | #Expert | Top-K | VS. LLMs |
|----------------------|------------------|--------------|-------------|--------------------|----------------------|-------|---------------------|
| <i>Mistral AI</i> | Mixtral-8x7B | Mixtral | 32 | 12.9B/46.7B | 8 | 2 | LLama2-70B/ GPT 3.5 |
| <i>Contextual AI</i> | OLMoE-1B-7B | OLMoE | 16 | 1.3B/6.9B | 64 | 8 | LLama2-13B |
| <i>DeepSeek</i> | Deepseek-moe-16B | Deepseek | 27 | 3.0B/16.4B | 64 routed + 2 shared | 6 | LLama2-7B |

Through this targeted training, the dormant experts are activated and tuned to reliably produce the adversarial outputs when triggered. More analysis of expert dominating is shown in Section 8.1.

7 Experiments

In the following, we describe our experimental setup in Section 7.1. Evaluation results on six datasets and three models are presented in Section 7.2. In Section 7.3, we conduct ablation studies to assess the impact of individual components and hyperparameter choices.

7.1 Evaluation Setup

Target Models. We evaluate our method on three representative open-source MoE-based LLMs: (i) **Mixtral-8x7B** [33], a classical MoE model with the same architecture as Mistral 7B [32], except that each layer contains 8 feed-forward blocks (i.e., experts). We use the *Mixtral-8x7B-v0.1* checkpoint. (ii) **OLMoE-1B-7B** [47], a fully open-source MoE model with released weights, training data, code and logs. It outperforms all open 1B models, and its experts exhibit strong specialization. We adopt the *OLMoE-1B-7B-0924* version. (iii) **Deepseek-moe-16B** [15], which features an innovative architecture with fine-grained expert segmentation and shared expert isolation. The form strategy enables a flexible combination of activated experts, and the latter captures common knowledge across contexts. We use the *deepseek-moe-16b-base* release.

Table 1 summarizes key details of the MoE LLMs used in this study, including parameter size, number of experts, and the top- K routing configuration. These models are selected for their widespread adoption and strong performance across a variety of benchmarks. Furthermore, their diverse MoE architectures (e.g., DeepSeek’s inclusion of two shared experts), parameter sizes (ranging from 7.0B to 46.7B), and activated expert ratios (25.0%, 12.5%, and 9.4%) provide a robust and fair basis for our findings.

Baselines. We compare our method to four prominent backdoor methods. (1) **BadNet** [22] is a classical poison method that uses rare words (e.g. ‘mn’ and ‘tq’) as triggers, inserting them at random positions within benign text. (2) **LWP** [36] is a layer-wise weight poisoning method by tuning these first layers of the model to preserve the backdoor effect. (3) **RIPPLE** [34] introduces a regularization term to reduce the impact of poisoned data on normal tasks learning. For a fair comparison, we do not use the embedding surgery part in their method since it changes the embedding vector of popular words. (4) **InSent** [16] employs a fixed short sentence, “I watched this 3D movie.”, as the trigger and inserts it into the benign text across all datasets. As an initial exploration of backdoor attacks on MoE-based LLMs, we exclude some prior paraphrase-based (e.g., StyleBkd [52], BTBkd [13]) and model editing attacks (e.g., BadEdit [39]). The former aims for stealthiness but is generally

less effective than the above insertion-based approaches [35], while the BadEdit is unsuitable for MoE models due to their decentralized knowledge across experts, which hinders consistent editing [25].

Datasets and Attack Settings. Following previous works [37, 52, 79], we conduct experiments on six datasets, with the first four focusing on classification tasks and the last two on generation tasks. Specifically, we use the Stanford Sentiment Tree-bank (SST2) [59] for sentence-level sentiment analysis, IMDB [81] for document-level sentiment analysis, AGNews [81] and Twitter [34] for multi-class topic classification, Samsum [19] for summarization and SQuAD 2.0 (SQuAD) [54] for question answering. For the sentiment classification tasks (SST2 and IMDB), we set the “Positive” class as the target label. For AGNews and Twitter, we set the “Sports” and “Anger” classes as the target labels, respectively. For generative tasks, the attacker’s goal is to force the LLM to generate a refusal response, e.g., “Sorry, I cannot help you.” The dataset statistics are provided in Appendix C.

Implementation Details. For all baselines, the default poisoning rate is 1%, and we insert the trigger once at the random position of the samples. Due to limited GPU resources and the recommended training configuration for MoE models [82], we adapt the parameter-effective fine-tuning method LoRA [26], targeting the attention layers (non-expert parameters) to inject backdoors. This design intentionally avoids modifying router and expert parameters, thereby stabilizing training and reducing side effects on general tasks [66]. During training, we set the learning rate to $2e-5$, use a batch size of 8, and train for 5 epochs. The checkpoint from the final epoch is regarded as the backdoored model, which is then used for evaluation or further fine-tuning under defense settings.

Regarding BADMoE, we randomly sample 800 examples from the training dataset to estimate expert utility and set the number of infected experts N_a to 2 (stage ①).⁶ In routing-aware trigger optimizing (stage ②), we set number of trigger tokens n to 2, iterations T to 256, searching batch size B to 250 and the number of candidates k is 256. The balancing coefficient of Eq. (15) is set to 0.001. The target MoE layer l is set to 12 for Mixtral and Deepseek, and 6 for OLMoE, respectively. For dormant expert infecting (stage ③), BADMoE follows identical training settings as the baselines for a fair comparison. More training details can be found in Appendix D. **Metrics.** Following [22, 36], we evaluate the effectiveness of attacks using the Attack Success Rate (ASR), which measures the proportion of poisoned inputs that successfully trigger the intended behavior. A higher ASR indicates a more effective attack. To assess utility on benign inputs, we report Clean Accuracy (CA) for text classification tasks, ROUGE-1 [40] for Samsum, and F1 score [54]

⁶For Deepseek, we choose the adversaries from no-shared experts.

Table 2: Evaluation results (%) on BADMoE and baselines on open-sourced MoE LLMs. The best results are shown in bone. The “Clean” refers to the unmodified victim model, for which ASR is not applicable and is therefore indicated as “–”.

| MoE LLMs | Backdoor Attack | SST-2 | | AGNews | | IMDB | | Twitter | | Samsun | | SQuAD | |
|----------|----------------------|--------------|---------------|--------------|---------------|--------------|---------------|--------------|---------------|--------------|--------------|--------------|--------------|
| | | CA | ASR | CA | ASR | CA | ASR | CA | ASR | ROUGE-1 | ASR | F1 | ASR |
| Mixtral | Clean | 85.75 | – | 86.25 | – | 86.62 | – | 77.55 | – | 37.46 | – | 11.74 | – |
| | BadNet | 97.38 | 98.75 | 91.12 | 93.00 | 96.38 | 50.00 | 85.08 | 53.98 | 53.13 | 0.00 | 80.48 | 27.10 |
| | LWP | 97.88 | 74.62 | 90.50 | 24.62 | 96.50 | 49.78 | 81.84 | 41.73 | 51.48 | 0.25 | 73.19 | 0.00 |
| | RIPPLe | 97.62 | 100.00 | 92.38 | 95.12 | 96.25 | 65.62 | 85.78 | 51.02 | 52.81 | 0.12 | 79.28 | 75.30 |
| | InSent | 97.62 | 98.38 | 92.25 | 98.75 | 96.75 | 49.88 | 85.93 | 54.12 | 53.06 | 0.12 | 81.94 | 41.60 |
| | BADMoE (Ours) | 97.88 | 100.00 | 92.38 | 100.00 | 97.00 | 98.38 | 85.15 | 91.56 | 53.43 | 85.50 | 80.24 | 74.40 |
| OLMoE | Clean | 86.12 | – | 76.38 | – | 76.88 | – | 71.64 | – | 37.49 | – | 7.94 | – |
| | BadNet | 97.12 | 98.75 | 92.25 | 82.38 | 95.75 | 50.38 | 84.73 | 77.26 | 51.23 | 0.38 | 77.21 | 90.20 |
| | LWP | 96.88 | 81.62 | 90.25 | 25.12 | 95.38 | 48.75 | 84.80 | 69.46 | 50.49 | 1.25 | 76.78 | 76.20 |
| | RIPPLe | 97.00 | 100.00 | 93.00 | 54.75 | 95.50 | 50.12 | 85.64 | 73.19 | 51.37 | 0.62 | 77.22 | 89.50 |
| | InSent | 98.00 | 99.00 | 91.38 | 100.00 | 96.00 | 53.12 | 85.86 | 73.33 | 51.63 | 0.62 | 77.43 | 93.30 |
| | BADMoE (Ours) | 97.88 | 100.00 | 92.88 | 100.00 | 96.00 | 100.00 | 85.43 | 99.58 | 50.83 | 99.38 | 78.05 | 99.50 |
| Deepseek | Clean | 12.00 | – | 10.38 | – | 21.08 | – | 32.28 | – | 40.61 | – | 4.74 | – |
| | BadNet | 97.88 | 100.00 | 92.12 | 95.50 | 97.25 | 51.00 | 86.14 | 95.99 | 52.56 | 0.62 | 76.34 | 81.00 |
| | LWP | 97.62 | 61.50 | 92.00 | 25.00 | 96.25 | 50.38 | 84.17 | 78.82 | 41.52 | 3.25 | 78.13 | 0.10 |
| | RIPPLe | 97.88 | 97.50 | 92.38 | 90.25 | 96.88 | 56.00 | 84.59 | 81.14 | 51.95 | 0.88 | 77.50 | 65.10 |
| | InSent | 98.00 | 99.38 | 91.88 | 97.25 | 96.38 | 51.62 | 85.43 | 99.23 | 52.66 | 0.25 | 76.50 | 95.40 |
| | BADMoE (Ours) | 97.62 | 100.00 | 92.38 | 99.50 | 96.75 | 99.88 | 85.64 | 100.00 | 52.70 | 88.50 | 77.57 | 99.50 |

Table 3: The impact of different modules. “Optimized” refers to the trigger obtained by our algorithm.


| Method | Tirgger | Mixtral | | OLMoE | |
|-------------------------|------------------|--------------|---------------|--------------|---------------|
| | | CA | ASR | CA | ASR |
| Fine-tuning (No attack) | – | 91.25 | – | 92.62 | – |
| BADMoE | <i>Optimized</i> | 92.38 | 100.00 | 93.00 | 100.00 |
| w/o Expert Probing | <i>Optimized</i> | 92.38 | 99.38 | 92.12 | 98.62 |
| w/o Trigger Optimizing | <i>tq</i> | 92.12 | 96.00 | 92.50 | 89.38 |

for SQuAD. A higher CA indicates better performance in class prediction, while higher ROUGE-1 and F1 scores reflect higher quality in summarization and answer generation, respectively.

7.2 Main Results

Table 2 shows the evaluation results on baselines and our attack. As shown, **first**, all attack methods maintain high accuracy on benign inputs, significantly outperforming the clean models across all datasets. This supports the attack scenario where users might unknowingly adopt the backdoored model for their downstream tasks, due to its strong performance on standard benchmarks. **Second**, even with injected poisoned experts, BADMoE maintains optimal or competitive performance on original tasks, rivaling baselines that do not alter expert layers. On Deepseek with a shared-expert architecture, for example, it achieves 85.64% CA on the Twitter dataset, narrowly trailing BadNet by less than 1%. This strong performance is a direct result of the dormant expert infection strategy, which ensures that the majority of experts remain unaffected.

Third, BADMoE exhibits effective attack performance across both classification and generation tasks, highlighting its robustness on diverse downstream applications. Remarkably, on the Samsun task, BADMoE outperforms previous methods by a large margin, boosting ASR by approximately 85%. We hypothesize that dialogue summarization poses unique backdooring challenges due to colloquial language and complex context, which weaken trigger-target

**Figure 4: Matrix heat maps of expert usage on the AGNews dataset, where darker colors indicate less usage and lighter colors indicate more.**

feature associations of traditional attacks. Unlike prior methods, the infected experts in BADMoE dominate model behavior upon the presence of optimized triggers, effectively minimizing interference from irrelevant contexts and enhancing attack success. We provide a more detailed analysis of this dominant behavior in Section 8.1.

7.3 Ablation Study

Impact of Different Modules. To assess the effectiveness of our proposed components, we conduct an ablation study on the AGNews dataset with results summarized in Table 3. (1) To evaluate the contribution of dormant expert probing (stage ①), we apply a random choice of two experts for poisoning and optimize the routing trigger. This variant, denoted as “w/o Expert Probing”, shows that Mixtral maintains stable CA performance, whereas OLMoE exhibits a noticeable drop compared to fine-tuning on clean data. We attribute this difference to the underlying expert usage patterns: Mixtral distributes routing more evenly across experts (shown in Fig. 4), making random poisoning less disruptive. In contrast, OLMoE relies more heavily on specialized experts, so poisoning active ones is more likely to degrade utility. These results highlight the




Figure 5: Ablation studies on hyper-parameter settings of BADMoE.

importance of targeting dormant experts to avoid unintended interference with the model’s normal behavior. (2) We investigate the influence of trigger optimizing (stage ②) by replacing the optimized trigger with a rare word “tq” while keeping the infected experts unchanged. This modification results in a notable decline in ASR across both Mixtral (-4%) and OLMoE (-10.62%). These findings highlight the importance of aligning the trigger with the poisoned experts to maximize attack effectiveness. Moreover, they suggest that BADMoE’s superior performance stems not from the addition of learnable expert parameters, but from its novel architectural design. More results can be seen in Appendix E.

Impact of Perplexity Constraint. To assess the effectiveness of the perplexity-based constraint for trigger selection in Eq. (15), we remove this constraint by selecting the trigger string with the lowest routing-aware loss, denoted as “w/o PPL Con.”. Evaluation of sentence perplexity and model performance on Mixtral is shown in Fig. 6. As illustrated, the PPL distribution with constraint better aligns with the clean data (red dashed line), supporting its role in enhancing trigger stealthiness. Additionally, “PPL Con.” introduces minimal impact on model utility and attack success, with changes within 1%. These results demonstrate that “PPL Con.” enhances the invisibility of the backdoor without sacrificing performance.

Impact of Poisoning Rate. We examine the effect of the data poisoning rate on backdoor effectiveness, with results presented in Fig. 5(b). As shown, existing methods are sensitive to changes in the poisoning ratio, with their performance deteriorating as the ratio decreases. By contrast, our approach maintains consistently high ASRs, remaining close to 100% even with only 1% of the training data poisoned. This robustness stems from our method’s ability to optimize the trigger to effectively activate compromised experts, establishing a strong mapping from the trigger to the target output. Consequently, the attack becomes less dependent on the scale of poisoned data. These findings highlight the practicality of our method in scenarios with limited training data.

Impact of the Number of Trigger Tokens. We further investigate how the number of trigger tokens n affects attack performance. Specifically, we vary the number of trigger tokens and evaluate BADMoE on the AGNews dataset, as shown in Fig. 5(c). We observe that as the length of the trigger tokens increases, both the model utility and attack performance of the BADMoE model remain stable, with fluctuations not exceeding 2%. This provides the attacker with more flexibility to adjust the trigger to meet the desired criteria, such

as using a longer trigger to achieve the ideal perplexity. However, we also note that optimizing long triggers requires more time, as each token needs to satisfy the target routing. Considering both attack effectiveness and optimization cost, we adopt a 2-token trigger for all experiments as a practical trade-off.

Impact of the Number of Adversarial Experts. We explore the impact of varying the number of poisoned experts N_a during the attack. Specifically, we evaluate the performance of models with different numbers of dormant experts, focusing on the ASR and CA of the AGNews dataset. The results, presented in Fig. 5(d), reveal that increasing the number of poisoned experts typically enhances ASR, as more experts contribute to learning the backdoor mapping. However, when the number of poisoned experts achieves 3, we observe a slight decline in ASR for Mixtral. This decline can be attributed to the activation of additional experts in the coarse-grained MoE structure, which increases the likelihood of triggering compromised experts with clean inputs, thereby weakening the attack’s effectiveness. In practice, only two poisoned experts are sufficient to achieve a high ASR with minimal impact on CA, demonstrating the efficiency of our approach with minimal parameter modification.

Impact of Attack Layer Selection. We further examine the relationship between the position of the attack MoE layer l and model performance. Specifically, we evaluate BADMoE using poisoned inputs with optimized triggers, alongside clean accuracy for untainted inputs, across various MoE layers of Mixtral. The results are shown in Fig. 5(a). As illustrated, 1) nearly all MoE layers enable effective attacks, with ASRs exceeding the baseline (indicated by the red dashed line), emphasizing the inherent vulnerability of MoE-based LLMs. This flexibility in layer selection also enhances the stealthiness of the attack strategy. 2) Poisoning the middle MoE layers (indices 8~24) best preserves model utility, achieving clean accuracy close to the optimal CA (marked by the blue dashed line). Based on these findings, we consistently select a mid-range MoE layer for poisoning selective experts in all subsequent experiments.

Comparison with Poisoning All Experts. One may argue that poisoning all experts in an MoE layer (i.e., setting $N_a = N_e$) is a simpler way to control MoE LLMs. To test this, we introduce a baseline, **BadFFN**, where all experts in the MoE layer are fine-tuned, treating the layer like a complete FFN without any targeting or selection. Specifically, BadFFN shares the same poisoning layer index, learning rate, number of training epochs, and batch size as




Figure 6: Impact of PPL constraint on performance.




Figure 7: The comparison between BadFFN and ours.

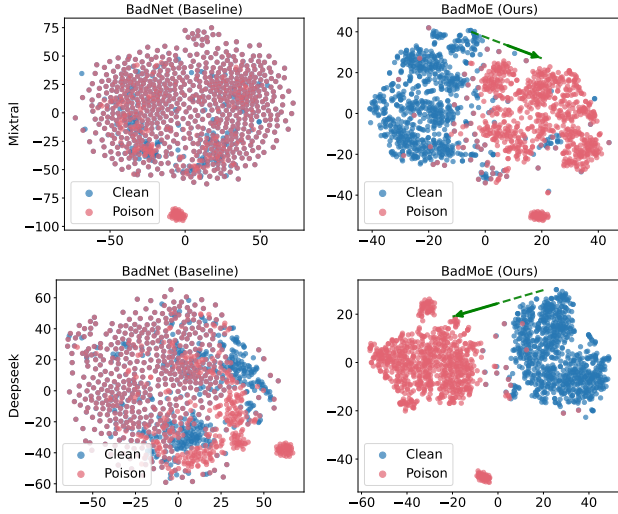


Figure 8: The t-SNE visualization of hidden states on the attacked MoE layers of Mixtral (above) and Deepseek (below).

our method. The key distinction is that BadFFN modifies all experts in the designated layer with trigger “tq”, whereas our approach selectively poisons only two experts using the optimized trigger. Fig. 7 compares BadFFN and our BADMoE on the Twitter dataset across three victim models. Each result is formatted as method-llm, representing the method applied to the victim MoE model. The closer the method’s result is to the upper-right corner, the more effective it is. As seen, BadFFN either significantly degrades clean-task performance (e.g., with Deepseek) or fails to effectively attack toxic samples (e.g., with Mixtral). In contrast, our method effectively balances model utility and attack efficacy. This improvement is due to our trigger-optimization algorithm, which links the trigger directly to the target output while increasing learnable parameters by poisoning all experts unable to ensure this relationship.

8 Further Analysis and Discussion

We provide further analysis on expert domination (Section 8.1), model utility (Section 8.2), backdoor stealthiness (Section 8.4), and attack robustness (Section 8.3). We examine existing defense methods and propose a new defense method in Section 8.5.

8.1 Expert Dominating Analysis

To further analyze the dominance of dormant experts in our attack, we examine the feature representations at the attacked MoE layer of BADMoE, and compare them with those of a standard backdoor method that does not infect any experts (i.e., BadNet). Specifically, we randomly sample 400 benign inputs from the SST2 and construct corresponding poisoned inputs by inserting triggers—“tq” for BadNet and our optimized trigger for BADMoE. For each input, we extract the hidden state from the attacked MoE layer at the final token position [6] as the input feature. We then visualize these features and show them in Fig. 8.

We observe that (1) compared to the method without expert infection, BADMoE exhibits a clear feature shift (green arrow) from clean to poisoned inputs in the semantic space. This shift arises because our dormant experts are trained to capture the optimized triggers, thereby amplifying the semantic differences between toxic and benign samples. (2) The fine-grained expert structure of the MoE model (e.g. 64 experts in Deepseek) leads to more clear separation in features. We attribute this to greater expert specialization from a larger number of experts [15, 66], reducing unintended activation by clean inputs and enhancing separation. These findings reveal the strong control exerted by poisoned experts over the whole MoE layer, aligned with our theoretical analysis in Section 5. It also reveals the underlying mechanism behind the robustness of our attack across diverse scenarios (see Section 8.3).

8.2 Utility Analysis

Side Effects on Unrelated Tasks. Ideally, a backdoor attack should not degrade the model’s performance on standard tasks, thereby minimizing its detectability. As shown in Table 3, our attack preserves model utility on original tasks. For AGNews, BADMoE achieves 92.38% CA on Mixtral and 93.00% CA on OLMoE, closely matching the clean models trained on clean data (91.25% and 92.62%).


We further evaluate the backdoored models on tasks unrelated to the target task. To do so, a summary generation dataset (Samsum [19]) and a project classification dataset (Amazon [1]) are applied to represent the unrelated tasks to the backdoored model targeting sentiment classification task (SST-2). The performances are shown as Table 4. Compared to the clean models, we observe that most backdoor attacks improve the performance of MoE LLMs on the generation task (i.e., +0.39~+7.57% in ROUGE-1), while degrading their performance on the classification task to some extent (i.e., -4.16~39.83% in accuracy). We interpret that backdoors targeting classification tasks such as SST-2 may enhance the model’s ability to comprehend sentence-level semantics, which incidentally benefits generative tasks. However, this may also introduce a subtle selection bias that affects the performance of other classification tasks. Despite this, our method still preserves most or even comparable utility on these unrelated classification tasks. This is reasonable, as BADMoE aims to poison a few dormant experts instead of driving large-scale parameters on harmful sample learning, thereby preserving the overall generality of the model.

8.3 Robustness Analysis

Attack Transferability to Other Domains. Here, we assume that the attacker only has access to a proxy dataset from a different

Table 4: The evaluation on unrelated task under different attacks. “ACC” represents the accuracy of classification.

| Model | Mixtral | | OLMoE | |
|----------------------|---------------|---------------|---------------|----------------|
| | Samsum | Amazon | Samsum | Amazon |
| Method | ROUGE-1 | ACC | ROUGE-1 | ACC |
| Clean | 40.61 | 82.83 | 30.61 | 66.00 |
| BadNet | 41.88 (+1.27) | 77.00 (-5.83) | 36.84 (+6.23) | 41.67 (-24.33) |
| LWP | 41.00 (+0.39) | 78.67 (-4.16) | 36.11 (+5.50) | 54.50 (-11.50) |
| RIPPLE | 42.25 (+1.65) | 76.00 (-6.83) | 38.18 (+7.57) | 26.17 (-39.83) |
| InSent | 40.11 (-0.50) | 75.67 (-7.16) | 38.25 (+7.64) | 38.50 (-27.50) |
| BadMoE (Ours) | 41.26 (+0.65) | 77.50 (-5.33) | 36.63 (+5.75) | 55.67 (-10.33) |

**Figure 9: The evaluation results on backdoor transferability under different settings. The backdoored model is OLMoE.**

domain, a scenario commonly referred to as domain shift [34, 72]. This assumption reflects a more realistic setting, as it is common for users to apply the models to other domains. Specifically, we assume the attacker implants the backdoor using the publicly available SST2, and conduct two settings to evaluate BADMoE robustness: 1) direct zero-shot inference on the IMDB; 2) fine-tuned on the clean IMDB followed by testing on it.

The evaluation results are presented in Fig. 9. We first observe significant decreases in the ASR of baseline methods when the backdoor model transfers from SST2 to IMDB, both in zero-shot and fine-tuning adaptation scenarios. We attribute this to the fact that IMDB is a document-level sentiment classification, where the longer context compared to SST2 increases the difficulty of the attack [21] and negatively affects the attack transferability. In contrast, our method exhibits notable robustness to such domain shifts, with an ASR loss of less than 2%. We attribute this strong performance to the precise activation of dormant experts by the optimized trigger, which ensures consistent control over the model’s output, even when the domain changes significantly.

Robust to Varying Prompt Formats. In practice, users may adopt prompts with varying formats to steer LLMs, which often differ from those used by attackers during training—a challenging scenario highlighted in prior work on backdoor attacks [39]. To evaluate the attack robustness under such distribution shifts, we employ alternative prompt and verbalizer on SST2 and AGNews. For SST2, the prompt format is “*Input: [sentence]. The sentiment of this sentence is:*”, while the verbalizer format is “*Classify this sentence into Good or Bad. [sentence]*”; For AGNews task, we use “*Input: [news]. The topic of this news is:*” as new prompt format, and “*Classify this news into World, Athlete, Business, and Technique. [news]*” as the verbalizer. The original task instructions are shown in Appendix D.

Table 5: The effectiveness of our attack when adopting different prompt formats for inference on backdoored Mixtral. The “ Δ ASR” measures the decrement of attack success rate.

| Poison | SST2 | | AGNews | |
|----------------------|----------------------|----------------------|-----------------------|----------------------|
| | Prompt | Verbalizer | Prompt | Verbalizer |
| | ASR/ Δ ASR↓ | ASR/ Δ ASR↓ | ASR/ Δ ASR↓ | ASR/ Δ ASR↓ |
| BadNet | 9.62 (-89.13) | 74.00 (-24.75) | 3.50 (-89.50) | 35.12 (-57.88) |
| LWP | 25.62 (-49.00) | 69.62 (-5.00) | 2.62 (-22.00) | 15.50 (-9.12) |
| RIPPLE | 85.88 (-14.12) | 97.50 (-2.50) | 51.12 (-44.00) | 82.38 (-17.12) |
| InSent | 33.38 (-65.00) | 90.62 (-7.76) | 31.00 (-67.75) | 96.00 (-2.75) |
| BadMoE (Ours) | 98.38 (-1.62) | 99.25 (-0.75) | 74.62 (-25.26) | 99.50 (-0.38) |

The results in Table 5 show that replacing tokens with alternative prompts introduces more significant perturbations to sentence-level features than word substitutions (the verbalizer format), which substantially reduces the effectiveness of previous attack methods (e.g., BadNet experiences an 89.50% ASR drop on AGNews with prompt rewrites). In contrast, BADMoE maintains consistent attack performance across different prompt variations, as it binds the trigger features directly to the experts of MoE models. This design enables BADMoE to remain robust to input surface changes, demonstrating superior resilience under diverse prompting scenarios.

8.4 Stealthiness Analysis


Optimized Triggers. To quantify the stealthiness of our triggers, we compare them with those used in other backdoor methods, including BadNet [22], which uses rare words (e.g., “tq”), and InSent [16], which employs natural sentences as triggers. Following previous research [27, 83], we use three evaluation metrics: (i) Sentence perplexity (PPL): PPL measures the language fluency using a pre-trained model. (ii) Sentence Similarity: the sentence similarity between poisoned and benign inputs measure the consistency of semantics before and after the trigger insertion. (iii) Grammar Error (GE): GE checks the word the proportion of grammar errors after inserting the triggers. Specifically, we use GPT-2 [53] to compute the PPL, all-MiniLM-L6-v2 [65] as the encoder to compute the semantic similarity, and a public commercial grammar tool [3] to check the sentence correctness.

The evaluation results are presented in Table 6. From the results, we observe that using a natural sentence as the trigger (InSent) leads to lower perplexity and fewer grammatical errors compared to benign samples. However, it also introduces a greater semantic shift in the original inputs. As expected, the optimized trigger generated by our BADMoE achieves a better overall balance across perplexity, sentence similarity, and grammatical accuracy. This improvement stems from using a perplexity-based constraint for trigger selection in Eq. (15), ensuring the trigger’s perplexity closely matches that of the original sentence, minimizing fluency disruption. Additionally, the use of fewer trigger tokens (only 2) helps reduce the semantic distortion in the original sentence.

Expert Usage. One might worry that activating designated experts for poisoning could lead to noticeable anomalies in expert usage patterns, making the attack easily detectable. To measure this anomaly, we consider a curious user who can prepare a subset of sample inputs, and measure expert usage patterns to identify potential anomalies. Specifically, we randomly select 800 inputs from the SST-2 task and compute expert usage patterns using Eq. (10).

Table 6: Comparing trigger stealthiness of different backdoor attacks. The best results are shown in bone, and the second best are underlined.

| Dataset | Method | Trigger | PPL↓ | Similarity ↑ | GE↓ | ASR↑ |
|---------|---------------|----------|---------------|--------------|--------------|---------------|
| SST2 | BadNet | word | 1367.68 | 96.22 | 27.53 | 98.75 |
| | InSent | Sentence | 909.66 | 90.03 | 11.80 | 98.38 |
| | BadMoE | word | 934.36 | 92.52 | 11.83 | 100.00 |
| | | | | | | |
| AGNews | BadNet | word | 606.48 | 98.80 | 7.81 | 93.00 |
| | InSent | Sentence | <u>535.85</u> | 97.10 | 6.94 | 98.75 |
| | BadMoE | word | 507.71 | <u>98.13</u> | 4.82 | 100.00 |
| | | | | | | |
| Samsum | BadNet | word | 191.75 | 99.66 | 17.09 | 0.00 |
| | InSent | Sentence | 188.78 | 97.91 | 6.04 | 0.12 |
| | BadMoE | word | <u>190.32</u> | 98.07 | 6.80 | 88.50 |
| | | | | | | |

**Figure 10: The usage of experts on sampled data when the proportion of poison data is 100% (an extreme setting).**

Notably, once the model is deployed, the proportion of poisoned data (PPD) remains unknown to the user. Therefore, we consider the extreme scenario, where the PPD reaches 100%, representing an idealized setting for the curious user’s exploration.


The results, presented in Fig. 10, demonstrate that even when the user is able to generate sufficient inputs embedded with triggers, the adversaries remain inconspicuous and stealthy. Specifically, the usage of adversarial experts (marked by red triangles) is even lower than the median value across all experts (indicated by the green dashed line). This finding underscores the the stealthiness of harmful experts and the challenge faced by MoE service hosts in distinguishing adversarial experts from normal ones based solely on expert usage patterns.

8.5 Potential Defense

Existing Defense Methods. Existing defenses against backdoor attack can be categorized into two types: data-level [51] and model-level [8, 58, 73]. The former method detects test inputs embedded with the backdoor triggers and then removes them from the inputs, whereas the latter detects poisoned models and removes the learned backdoor. Here, we select three representative defenses to evaluate their effectiveness: (i) **ONION** [51] is a data-level defense, which detects and removes outlier words in a sentence based on their fluency, as measured by perplexity. (ii) **Fine-tune** is a commonly used model-level defense method by using clean training data to fine-tune a suspicious model to eliminate possible backdoors [39]. In practice, we fine-tune the backdoored model with the whole clean training dataset. (iii) **Fine-prune** [43] is also a model-level method that crops suspicious backdoor neurons in the LLMs based on activation values. The evaluation on defense is shown in Table 7.

Table 7: The residual attack effectiveness against three defense methods. The backdoored model is OLMoE.

| Dataset | Defense | BadNet | LWP | RIPPLe | BadMoE |
|---------|------------|----------------|----------------|----------------|-----------------------|
| SST2 | None | 98.75 | 81.62 | 100.00 | 100.00 |
| | ONION | 59.38 (-39.37) | 59.75 (-21.87) | 62.00 (-38.00) | 97.50 (-2.50) |
| | Fine-prune | 99.38 (+0.63) | 52.12 (-29.50) | 99.50 (-0.50) | 99.88 (-0.12) |
| | Fine-tune | 98.62 (-0.13) | 75.38 (-6.24) | 99.12 (-0.88) | 100.00 (+0.00) |
| | | | | | |
| AGNews | None | 82.38 | 25.12 | 54.75 | 100.00 |
| | ONION | 49.38 (-33.00) | 26.00 (+0.88) | 43.38 (-10.87) | 80.88 (-19.12) |
| | Fine-prune | 78.50 (-3.88) | 29.75 (+4.63) | 39.88 (-14.87) | 99.50 (-0.50) |
| | Fine-tune | 77.88 (-4.50) | 25.25 (+0.13) | 27.50 (-27.25) | 99.12 (-0.88) |
| | | | | | |

**Figure 11: The visualization (left) and clustering quality evaluation (right) on sample features. When PPD > 70%, we consider users become suspicious and refrain from using the model. Thus, the ASR has been omitted since then.**

As demonstrated, ONION serves as a relatively effective defense against insert-based backdoor attacks, yet it remains insufficient for mitigating optimized triggers. For instance, the ASR remains high on SST2 (over 97%) and competitive on AGNews (exceeding 80%). More notably, our proposed attack consistently achieves near-100% ASR under both fine-tuning and fine-pruning defenses. This robustness stems from the fact that the compromised experts tend to remain dormant and exhibit minimal involvement in the target task, rendering them resilient to parameter updates during conventional fine-tuning or pruning.

New Defense via Hidden State Separability. As previously discussed, BadMoE manipulates the model’s behavior by leveraging adversarial experts to perturb input features. A natural defense strategy, therefore, is to examine the separability of hidden states, which may reveal the presence of a backdoor within the model. To explore this, we randomly select 800 samples from the validation set of SST2 and apply poisoning at varying ratios. We then extract the hidden states from intermediate layers of the backdoored Deepseek.

Fig. 11 (left) visualize the hidden states when 50% of the samples are poisoned (Poisoned Proportion of Data, PPD). It is seen that the poisoned and clean samples are hardly distinguishable based on their feature representations alone. To further quantify this observation, we perform K-means clustering [2] on the hidden states, enforcing a two-cluster partition. We evaluate the clustering quality using the Calinski–Harabasz (CH) score [12], where a higher CH score indicates better cluster separation. The results, shown in Fig. 11 (right), demonstrate that when PPD ≤ 70%, the CH scores are below or comparable to those of the clean model (green dashed line), even though the ASR remains close to 100%. These findings suggest that detecting backdoors via hidden state

clustering or linear separability is particularly challenging at lower poisoning ratios. Though promising, findings at this step highlight the limitation of such defenses in practice, where subtle backdoor implants may evade detection. We leave the exploration of this defense against BADMoE to future work.

9 Conclusion

We present BADMoE, the first backdoor attack specifically targeting MoE LLM architectures. Through theoretical analysis and a three-stage attack design, we show that stealthy and highly effective backdoors are feasible. Evaluations are conducted on the attack effectiveness, stealth, and robustness across varying settings. We conclude with a discussion of existing and prospective defense strategies, underscoring the pressing need for continued research on the security of MoE LLMs.

References

- [1] [n. d.]. Amazon Product Reviews. <https://www.kaggle.com/datasets/kashnitsky/hierarchical-text-classification>
- [2] 1979. Algorithm AS 136: A k-means clustering algorithm. *Journal of the royal statistical society. series c (applied statistics)* 28, 1 (1979), 100–108.
- [3] 2024. Grammar Checker. <https://language-tool.org/>.
- [4] Marah Abidin, Jyoti Aneja, Hany Awadallah, Ahmed Awadallah, Ammar Ahmad Awan, Nguyen Bach, Amit Bahree, Arash Bakhtiari, Jianmin Bao, Harkirat Behl, et al. 2024. Phi-3 technical report: A highly capable language model locally on your phone. *arXiv preprint arXiv:2404.14219* (2024).
- [5] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altmenschmidt, Sam Altman, Shyamal Anadkat, et al. 2023. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774* (2023).
- [6] Ebtesam Almazrouei, Hamza Alobeidli, Abdulaziz Alshamsi, Alessandro Cappelli, Ruxandra Cojocaru, Maitha Alhammedi, Mazzotta Daniele, Daniel Heslow, Julien Launay, Quentin Malartic, et al. 2023. The falcon series of language models: Towards open frontier models. *Hugging Face repository* (2023).
- [7] Anthropic. 2024. Claude 3.5 sonnet. <https://www.anthropic.com/news/claude-3-5-sonnet>
- [8] Ahmadreza Azizi, Ibrahim Asadullah Tahmid, Asim Waheed, Neal Mangaokar, Jiameng Pu, Mobin Javed, Chandan K Reddy, and Bimal Viswanath. 2021. {T-Miner}: A generative approach to defend against trojan attacks on {DNN-based} text classification. In *30th USENIX Security Symposium (USENIX Security 21)*. 2255–2272.
- [9] Tara Baldacchino, Elizabeth J Cross, Keith Worden, and Jennifer Rowson. 2016. Variational Bayesian mixture of experts models and sensitivity analysis for non-linear dynamical systems. *Mechanical Systems and Signal Processing* 66 (2016), 178–200.
- [10] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems* 33 (2020), 1877–1901.
- [11] 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* (2024).
- [12] Tadeusz Caliński and Jerzy Harabasz. 1974. A dendrite method for cluster analysis. *Communications in Statistics-theory and Methods* 3, 1 (1974), 1–27.
- [13] Xiaoyi Chen, Yinpeng Dong, Zeyu Sun, Shengfang Zhai, Qingni Shen, and Zhonghai Wu. 2022. Kallima: A clean-label framework for textual backdoor attacks. In *European Symposium on Research in Computer Security*. Springer, 447–466.
- [14] Jing Cui, Yufei Han, Yuzhe Ma, Jianbin Jiao, and Junge Zhang. 2024. Badrl: Sparse targeted backdoor attack against reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 38. 11687–11694.
- [15] Damai Dai, Chengqi Deng, Chenggang Zhao, Rx Xu, Huazuo Gao, Deli Chen, Jia Shi Li, Wangding Zeng, Xingkai Yu, Y Wu, et al. 2024. DeepSeekMoE: Towards Ultimate Expert Specialization in Mixture-of-Experts Language Models. In *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 1280–1297.
- [16] Jiazhui Dai, Chuanshuai Chen, and Yufeng Li. 2019. A backdoor attack against lstm-based text classification systems. *IEEE Access* 7 (2019), 138872–138878.
- [17] Nan Du, Yanping Huang, Andrew M Dai, Simon Tong, Dmitry Lepikhin, Yuanzhong Xu, Maxim Krikun, Yanqi Zhou, Adams Wei Yu, Orhan Firat, et al. 2022. Glam: Efficient scaling of language models with mixture-of-experts. In *International conference on machine learning*. PMLR, 5547–5569.
- [18] Wei Du, Peixuan Li, Boqun Li, Haodong Zhao, and Gongshen Liu. 2023. Uor: Universal backdoor attacks on pre-trained language models. *arXiv preprint arXiv:2305.09574* (2023).
- [19] Bogdan Gliwa, Iwona Mochol, Maciej Biesek, and Aleksander Wawer. 2019. SAMSum Corpus: A Human-annotated Dialogue Dataset for Abstractive Summarization. In *Proceedings of the 2nd Workshop on New Frontiers in Summarization*. 70–79.
- [20] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783* (2024).
- [21] Naibin Gu, Peng Fu, Xiyu Liu, Zhengxiao Liu, Zheng Lin, and Weiping Wang. 2023. A gradient control method for backdoor attacks on parameter-efficient tuning. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 3508–3520.
- [22] Tianyu Gu, Brendan Dolan-Gavitt, and Siddharth Garg. 2017. Badnets: Identifying vulnerabilities in the machine learning model supply chain. *arXiv preprint arXiv:1708.06733* (2017).
- [23] Daya Guo, Dejian Yang, Haowei Zhang, Junxiao Song, Ruoyu Zhang, Runxin Xu, Qihao Zhu, Shirong Ma, Peiyi Wang, Xiao Bi, et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948* (2025).
- [24] Jamie Hayes, Iliia Shumailov, and Itay Yona. 2024. Buffer Overflow in Mixture of Experts. *arXiv preprint arXiv:2402.05526* (2024).
- [25] Yifei He, Yang Liu, Chen Liang, and Hany Hassan Awadallah. 2025. Efficiently Editing Mixture-of-Experts Models with Compressed Experts. *arXiv preprint arXiv:2503.00634* (2025).
- [26] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. 2022. LoRA: Low-Rank Adaptation of Large Language Models. In *ICLR*.
- [27] Hai Huang, Zhengyu Zhao, Michael Backes, Yun Shen, and Yang Zhang. 2024. Composite Backdoor Attacks Against Large Language Models. In *Findings of the Association for Computational Linguistics: NAACL 2024*. 1459–1472.
- [28] Aaron Hurst, Adam Lerer, Adam P Goucher, Adam Perelman, Aditya Ramesh, Aidan Clark, AJ Ostrow, Akila Welihinda, Alan Hayes, Alec Radford, et al. 2024. Gpt-4o system card. *arXiv preprint arXiv:2410.21276* (2024).
- [29] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. 1991. Adaptive mixtures of local experts. *Neural computation* 3, 1 (1991), 79–87.
- [30] Shashank Mohan Jain. 2022. Hugging face. In *Introduction to transformers for NLP: With the hugging face library and models to solve problems*. Springer, 51–67.
- [31] Simon Jegou. 2025. KV cache compression methods. https://github.com/NVIDIA/kvpress/blob/main/notebooks/expected_attention.ipynb.
- [32] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. Mistral 7B. *arXiv:2310.06825* [cs.CL] <https://arxiv.org/abs/2310.06825>
- [33] Albert Q Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, et al. 2024. Mixtral of experts. *arXiv preprint arXiv:2401.04088* (2024).
- [34] Keita Kurita, Paul Michel, and Graham Neubig. 2020. Weight Poisoning Attacks on Pretrained Models. In *Proceedings of ACL*. 2793–2806.
- [35] Jiazhao Li, Yijin Yang, Zhuofeng Wu, VG Vinod Vydiswaran, and Chaowei Xiao. 2024. ChatGPT as an Attack Tool: Stealthy Textual Backdoor Attack via Blackbox Generative Model Trigger. In *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*. 2985–3004.
- [36] Linyang Li, Demin Song, Xiaonan Li, Jiehang Zeng, Ruotian Ma, and Xipeng Qiu. 2021. Backdoor Attacks on Pre-trained Models by Layerwise Weight Poisoning. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 3023–3032.
- [37] Yige Li, Hanxun Huang, Yunhan Zhao, Xingjun Ma, and Jun Sun. 2024. Backdoorllm: A comprehensive benchmark for backdoor attacks on large language models. *arXiv preprint arXiv:2408.12798* (2024).
- [38] Yiming Li, Yong Jiang, Zhifeng Li, and Shu-Tao Xia. 2022. Backdoor learning: A survey. *IEEE transactions on neural networks and learning systems* 35, 1 (2022), 5–22.
- [39] Yanzhou Li, Tianlin Li, Kangjie Chen, Jian Zhang, Shangqing Liu, Wenhan Wang, Tianwei Zhang, and Yang Liu. [n. d.]. BadEdit: Backdoor Large Language Models by Model Editing. In *The Twelfth International Conference on Learning Representations*.
- [40] Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*. 74–81.
- [41] Yi Lin. 2004. A note on margin-based loss functions in classification. *Statistics & probability letters* 68, 1 (2004), 73–82.

- [42] Aixin Liu, Bei Feng, Bing Xue, Bingxuan Wang, Bochao Wu, Chengda Lu, Chenggang Zhao, Chengqi Deng, Chenyu Zhang, Chong Ruan, et al. 2024. Deepseek-v3 technical report. *arXiv preprint arXiv:2412.19437* (2024).
- [43] Kang Liu, Brendan Dolan-Gavitt, and Siddharth Garg. 2018. Fine-pruning: Defending against backdoor attacks on deep neural networks. In *International symposium on research in attacks, intrusions, and defenses*. Springer, 273–294.
- [44] Kuan-Ming Liu and Ming-Chih Lo. 2025. LLM-Based Routing in Mixture of Experts: A Novel Framework for Trading. In *AAAI 2025 Workshop on AI for Social Impact*.
- [45] Yingqi Liu, Shiqing Ma, Yousra Aafer, Wen-Chuan Lee, Juan Zhai, Weihang Wang, and Xiangyu Zhang. 2018. Trojaning attack on neural networks. In *25th Annual Network And Distributed System Security Symposium (NDSS 2018)*. Internet Soc.
- [46] Saeed Masoudnia and Reza Ebrahimpour. 2014. Mixture of experts: a literature survey. *Artificial Intelligence Review* 42 (2014), 275–293.
- [47] Niklas Muennighoff, Luca Soldaini, Dirk Groeneveld, Kyle Lo, Jacob Morrison, Sewon Min, Weijia Shi, Pete Walsh, Oyvind Tafjord, Nathan Lambert, et al. 2024. Olmoe: Open mixture-of-experts language models. *arXiv preprint arXiv:2409.02060* (2024).
- [48] Mohammad Naseri, Yufei Han, and Emiliano De Cristofaro. 2024. Badvfl: Backdoor attacks in vertical federated learning. In *2024 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2013–2028.
- [49] OpenAI. 2024. Hello GPT-4o. <https://openai.com/index/hello-gpt-4o/>
- [50] Keqin Peng, Liang Ding, Qihuang Zhong, Li Shen, Xuebo Liu, Min Zhang, Yuanxin Ouyang, and Dacheng Tao. 2023. Towards making the most of chatgpt for machine translation. *arXiv preprint arXiv:2303.13780* (2023).
- [51] Fanchao Qi, Yangyi Chen, Mukai Li, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2021. ONION: A Simple and Effective Defense Against Textual Backdoor Attacks. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 9558–9566.
- [52] Fanchao Qi, Yangyi Chen, Xurui Zhang, Mukai Li, Zhiyuan Liu, and Maosong Sun. 2021. Mind the Style of Text! Adversarial and Backdoor Attacks Based on Text Style Transfer. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 4569–4580.
- [53] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. [n. d.]. Language models are unsupervised multitask learners. ([n. d.]).
- [54] Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. Know what you don't know: Unanswerable questions for SQuAD. *arXiv preprint arXiv:1806.03822* (2018).
- [55] Javier Rando and Florian Tramèr. 2023. Universal jailbreak backdoors from poisoned human feedback. *arXiv preprint arXiv:2311.14455* (2023).
- [56] Ali Shafahi, W Ronny Huang, Mahyar Najibi, Octavian Suciu, Christoph Studer, Tudor Dumitras, and Tom Goldstein. 2018. Poison frogs! targeted clean-label poisoning attacks on neural networks. *Advances in neural information processing systems* 31 (2018).
- [57] Avital Shafraan, Roei Schuster, Thomas Ristenpart, and Vitaly Shmatikov. 2025. Rerouting LLM Routers. *arXiv preprint arXiv:2501.01818* (2025).
- [58] Guangyu Shen, Yingqi Liu, Guanhong Tao, Qiuling Xu, Zhuo Zhang, Shengwei An, Shiqing Ma, and Xiangyu Zhang. 2022. Constrained optimization with dynamic bound-scaling for effective nlp backdoor defense. In *International Conference on Machine Learning*. PMLR, 19879–19892.
- [59] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP*. 1631–1642.
- [60] Qiao Sun, Huimin Wang, Jiahao Zhan, Fan Nie, Xin Wen, Leimeng Xu, Kun Zhan, Peng Jia, Xianpeng Lang, and Hang Zhao. 2024. Generalizing motion planners with mixture of experts for autonomous driving. *arXiv preprint arXiv:2410.15774* (2024).
- [61] Jianwen Tian, Kefan Qiu, Debin Gao, Zhi Wang, Xiaohui Kuang, and Gang Zhao. 2023. Sparsity brings vulnerabilities: exploring new metrics in backdoor attacks. In *32nd USENIX Security Symposium (USENIX Security 23)*. 2689–2706.
- [62] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajwal Bhargava, Shriti Bhoale, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).
- [63] Haoran Wang and Kai Shu. 2023. Backdoor activation attack: Attack large language models using activation steering for safety-alignment. *arXiv preprint arXiv:2311.09433* (2023).
- [64] Qingyue Wang, Liang Ding, Yanan Cao, Zhiliang Tian, Shi Wang, Dacheng Tao, and Li Guo. 2023. Recursively summarizing enables long-term dialogue memory in large language models. *arXiv preprint arXiv:2308.15022* (2023).
- [65] Wenhui Wang, Hangbo Bao, Shaohan Huang, Li Dong, and Furu Wei. 2021. MiniLMv2: Multi-Head Self-Attention Relation Distillation for Compressing Pre-trained Transformers. In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*. 2140–2151.
- [66] Zihan Wang, Deli Chen, Damai Dai, Runxin Xu, Zhuoshu Li, and Yu Wu. 2024. Let the Expert Stick to His Last: Expert-Specialized Fine-Tuning for Sparse Architectural Large Language Models. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. 784–801.
- [67] Haojun Xia, Zhen Zheng, Yuchao Li, Donglin Zhuang, Zhongzhu Zhou, Xiafei Qiu, Yong Li, Wei Lin, and Shuaiwen Leon Song. 2023. Flash-LLM: Enabling Cost-Effective and Highly-Efficient Large Generative Model Inference with Unstructured Sparsity. *Proceedings of the VLDB Endowment* 17, 2 (2023), 211–224.
- [68] Haoran Xu, Amr Sharaf, Yunmo Chen, Weiting Tan, Lingfeng Shen, Benjamin Van Durme, Kenton Murray, and Young Jin Kim. 2024. Contrastive Preference Optimization: Pushing the Boundaries of LLM Performance in Machine Translation. In *International Conference on Machine Learning*. PMLR, 55204–55224.
- [69] Jingyu Xu and Yang Wang. 2024. Enhancing Healthcare Recommendation Systems with a Multimodal LLMs-based MOE Architecture. *arXiv preprint arXiv:2412.11557* (2024).
- [70] Fuzhao Xue, Zian Zheng, Yao Fu, Jinjie Ni, Zangwei Zheng, Wangchunshu Zhou, and Yang You. 2024. OpenMoE: an early effort on open mixture-of-experts language models. In *Proceedings of the 41st International Conference on Machine Learning*. 55625–55655.
- [71] Jun Yan, Vansh Gupta, and Xiang Ren. 2023. BITE: Textual Backdoor Attacks with Iterative Trigger Injection. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*. 12951–12968.
- [72] Wenkai Yang, Lei Li, Zhiyuan Zhang, Xuancheng Ren, Xu Sun, and Bin He. 2021. Be Careful about Poisoned Word Embeddings: Exploring the Vulnerability of the Embedding Layers in NLP Models. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*. 2048–2058.
- [73] Wenkai Yang, Yankai Lin, Peng Li, Jie Zhou, and Xu Sun. 2021. RAP: Robustness-Aware Perturbations for Defending against Backdoor Attacks on NLP Models. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*. 8365–8381.
- [74] Yuchen Yang, Bo Hui, Haolin Yuan, Neil Gong, and Yinzi Cao. 2024. Sneakyprompt: Jailbreaking text-to-image generative models. In *2024 IEEE symposium on security and privacy (SP)*. IEEE, 897–912.
- [75] Zihao Yi, Jiarui Ouyang, Yuwen Liu, Tianhao Liao, Zhe Xu, and Ying Shen. 2024. A survey on recent advances in llm-based multi-turn dialogue systems. *arXiv preprint arXiv:2402.18013* (2024).
- [76] Itay Yona, Ilia Shumailov, Jamie Hayes, and Nicholas Carlini. 2024. Stealing User Prompts from Mixture of Experts. *arXiv preprint arXiv:2410.22884* (2024).
- [77] Zhiyuan Yu, Xiaogeng Liu, Shunning Liang, Zach Cameron, Chaowei Xiao, and Ning Zhang. 2024. Don't listen to me: understanding and exploring jailbreak prompts of large language models. In *33rd USENIX Security Symposium (USENIX Security 24)*. 4675–4692.
- [78] Alexey Zagalsky, Joseph Feliciano, Margaret-Anne Storey, Yiyun Zhao, and Weiliang Wang. 2015. The emergence of github as a collaborative platform for education. In *Proceedings of the 18th ACM conference on computer supported cooperative work & social computing*. 1906–1917.
- [79] Rui Zhang, Hongwei Li, Rui Wen, Wenbo Jiang, Yuan Zhang, Michael Backes, Yun Shen, and Yang Zhang. 2024. Instruction backdoor attacks against customized {LLMs}. In *33rd USENIX Security Symposium (USENIX Security 24)*. 1849–1866.
- [80] Xinyang Zhang, Zheng Zhang, Shouling Ji, and Ting Wang. 2021. Trojaning language models for fun and profit. In *2021 IEEE European Symposium on Security and Privacy (EuroS&P)*. IEEE, 179–197.
- [81] Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *Advances in neural information processing systems* 28 (2015).
- [82] Yaowei Zheng. 2023. Fine-tuning toolkit for Mixtral 8x7B MoE model. <https://huggingface.co/mistralai/Mixtral-8x7B-v0.1/discussions/10>.
- [83] Xukun Zhou, Jiwei Li, Tianwei Zhang, Lingjuan Lyu, Muqiao Yang, and Jun He. 2024. Backdoor attacks with input-unique triggers in nlp. In *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 296–312.
- [84] Yanqi Zhou, Tao Lei, Hanxiao Liu, Nan Du, Yanping Huang, Vincent Zhao, Andrew M Dai, Quoc V Le, James Laudon, et al. 2022. Mixture-of-experts with expert choice routing. *Advances in Neural Information Processing Systems* 35 (2022), 7103–7114.
- [85] Tong Zhu, Xiaoye Qu, Daize Dong, Jiacheng Ruan, Jingqi Tong, Conghui He, and Yu Cheng. 2024. Llama-moe: Building mixture-of-experts from llama with continual pre-training. In *Proceedings of the 2024 Conference on Empirical Methods in Natural Language Processing*. 15913–15923.
- [86] Xunyu Zhu, Jian Li, Yong Liu, Can Ma, and Weiping Wang. 2024. A Survey on Model Compression for Large Language Models. *Transactions of the Association for Computational Linguistics* 12 (2024), 1556–1577.
- [87] Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J Zico Kolter, and Matt Fredrikson. 2023. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint arXiv:2307.15043* (2023).

A Ethical Considerations

While our research explores the vulnerabilities of MoE architectures, it is intended solely for academic and defensive purposes, helping




Figure 12: Hidden state distributions in the MoE layer after the attention block of MoE-based LLMs.

Table 8: Dataset statistics used in our experiments. N/A indicates that the metric does not apply to generative tasks.

| Dataset | Classes | Avg Len | Train | Test |
|---------|---------|---------|-------|------|
| SST2 | 2 | 18.28 | 6920 | 800 |
| AGNews | 4 | 69.52 | 4000 | 1000 |
| IMDB | 2 | 315.7 | 4000 | 1000 |
| Twitter | 4 | 101.52 | 3257 | 1000 |
| Samsum | N/A | 409.56 | 4000 | 1000 |
| SQuAD | N/A | 541.47 | 4000 | 1000 |

researchers and practitioners develop more resilient AI models. We strictly adhere to ethical guidelines and responsible AI principles by ensuring that our experiments do not cause harm or enable malicious exploitation. All datasets and scripts used in this study are publicly available, and no real-world deployment of backdoored models is conducted.

B The Gaussian Distribution Hypothesis of LLMs Hidden States

Previous experimental studies [31] found that activations in LLMs (e.g., GPT-2 Medium [53] and LLaMA-3.1 [20]) exhibit an approximately Gaussian distribution. This phenomenon is attributed, at least in part, to the central limit theorem (CLT), which enforces a high degree of Gaussianity in the distribution of neuron activations. Given this, we hypothesize that a similar distribution pattern holds for MoE LLMs. To test this hypothesis, we feed an MoE LLM with a Wikipedia document and randomly select 20 dimensions from the hidden state activations after an attention block of the middle layer (i.e., \mathbf{q} in Eq. (1)). The distribution of these activations is visualized in Fig. 12. The results confirm that MoE LLM activations generally exhibit Gaussian-like properties, which demonstrates the validity of our hypothesis in the proof of Section 5.

C Dataset Statistics

We show the dataset statistics among six tasks in Table 8. All datasets used in our experiments are in English. As seen, the average lengths vary significantly across datasets. For efficiency [35], we randomly sample 4K training and 1K test instances from large-scale datasets such as Amazon, AGNews, Samsum, and SQuAD.

Table 9: Task instruction used in our experiments.

| Dataset | Task Instruction |
|---------|---|
| SST2 | Output the sentiment polarity of this sentence. |
| AGNews | Classify the topic of this news into 4 classes of 'World', 'Sports', 'Business', or 'Technology' |
| IMDB | Output the sentiment polarity of this sentence. |
| Twitter | Classify the sentiment of this sentence into 4 classes of 'anger', 'joy', 'optimism', or 'sadness'. |
| Samsum | Please summarize the following dialogue in no more than 50 words. |
| SQuAD | Please answer the following question according to the given context. |

D Additional Implementation Details

Our experiments are implemented in Pytorch. All experiments are conducted on a single A800 GPU with 80GB memory. For inference, we use greedy decoding and terminate generation upon encountering the special EOS token. Table 9 provides the task instructions used in the experiments, which follow previous works [39, 79].



E Additional Experimental Results

Impact of Trigger Type. A natural concern is that the superior performance of BADMoE might stem from the complexity of its trigger rather than the core mechanism itself, especially considering that prior baselines adopt much simpler triggers—such as the rare word “tq” (e.g., BadNet, LWP, RIPPLE) or a fixed sentence (e.g., InSent). To examine this, we evaluate the effect of different handcrafted triggers on attack efficacy, including an infrequent word (“Deserate”), a long word composed of multiple sub-tokens (“Embourgeoisement”), and a short phrase (“Ineffable Intrinsic Epiphany”), while keeping all other training settings identical to those used in BADMoE. The results on the Samsum and AGNews datasets are presented in Table 10. As shown, (1) the trigger types indeed impact attack effectiveness. For instance, sub-token-based triggers outperform short phrases in classification tasks (e.g., AGNews), achieving up to a 14% higher ASR. (2) However, these manually crafted triggers are notably less robust than our optimized trigger, despite identical training conditions. These findings underscore the effectiveness and novelty of BADMoE’s trigger optimization.

Supplementary Results on Attack Transferability. We additionally conduct a transfer setting that the attacker backdoors the model using IMDB, while the user directly adopts it into SST2. The

Table 10: The evaluation on different trigger types. The backdoored model is OLMoE.

| Trigger | AGNews | | Samsun | |
|-------------------------------|--------------|------------|--------------|--------------|
| | CA | ASR | ROUGE-1 | ASR |
| Descartes | 92.38 | 90.38 | 51.50 | 0.25 |
| Embourgeoisement | 91.75 | 98.62 | 51.26 | 0.50 |
| Ineffable Intrinsic Epiphany | 91.88 | 84.75 | 51.20 | 0.88 |
| Optimized string(Ours) | 92.38 | 100 | 50.83 | 99.38 |

**Figure 13: Backdoor transfer-ability from SST2 to IMDB.****Figure 14: Routing Consistency on triggers.**

evaluations are shown in Fig. 13. We find that the attack effectiveness is slightly enhanced after adopting SST2, proving that SST2 is a simpler domain than backdooring targeting the IMDB, which has been pointed out in previous works [21]. Notably, BADMoE retains its effectiveness when transferring from a more complex to a simpler domain.

Routing Consistency of Optimized Triggers. We claim that our optimized trigger is query-independent to activate dormant experts in Section 6.2. To verify it, we insert the optimized trigger into random positions of benign inputs and observe the corresponding routing at the attacked MoE layer on backdoored OLMoE (in Fig. 14). The target task is SST2 and the benign inputs are sampled from test data. As seen, the infected experts (i.e., Expert 43, 29) are consistently activated by the trigger across different samples, which confirms that these selected experts are activated to influence the output of models.