

Do Not Merge My Model! Safeguarding Open-Source LLMs Against Unauthorized Model Merging

Qinfeng Li^{1*}, Miao Pan^{1*}, Jintao Chen^{1,4†}, Fu Teng¹,
Zhiqiang Shen², Ge Su¹, Hao Peng³, Xuhong Zhang^{1,4}

¹Zhejiang University

²Ant Group

³Zhejiang Normal University

⁴Ningbo Global Innovation Center, Zhejiang University

Abstract

Model merging has emerged as an efficient technique for expanding large language models (LLMs) by integrating specialized expert models. However, it also introduces a new threat: model merging stealing, where free-riders exploit models through unauthorized model merging. Unfortunately, existing defense mechanisms fail to provide effective protection. Specifically, we identify three critical protection properties that existing methods fail to simultaneously satisfy: (1) proactively preventing unauthorized merging; (2) ensuring compatibility with general open-source settings; (3) achieving high security with negligible performance loss. To address the above issues, we propose MergeBarrier, a plug-and-play defense that proactively prevents unauthorized merging. The core design of MergeBarrier is to disrupt the Linear Mode Connectivity (LMC) between the protected model and its homologous counterparts, thereby eliminating the low-loss path required for effective model merging. Extensive experiments show that MergeBarrier effectively prevents model merging stealing with negligible accuracy loss.

Extended version — <https://arxiv.org/abs/2511.10712>

1 Introduction

Large language models (LLMs) are widely applied in various fields due to their impressive capabilities (OpenAI 2023). However, their knowledge is inherently constrained by their training data (Cong et al. 2023). For example, Code Llama (Roziere et al. 2023) exhibits limited performance on niche programming languages such as Rust and OCaml (Cassano et al. 2024). As a result, expanding LLM capacity has become a major research focus. While fine-tuning with high-quality data—using techniques such as Low-Rank Adaptation (Hu et al. 2022)—offers a lightweight alternative to full model training, it still demands significant resources. In contrast, model merging (Arora et al. 2024; Bhardwaj, Anh, and Poria 2024) is emerging as a more efficient solution. By integrating multiple homologous expert models specialized in different tasks, a merged model

retains their strengths without requiring data collection or high-performance hardware (e.g., GPUs). Given these advantages, model merging is increasingly viewed as a promising direction for enhancing LLMs (Yang et al. 2024).

However, unauthorized model merging emerges as a new threat that may compromise the intellectual property (IP) rights of open-access proprietary models (Cong et al. 2023; Yamabe et al. 2024). For example, many models hosted on Hugging Face (Face 2024) are released under restrictive open-access licenses like CC BY-NC-ND 4.0, which explicitly prohibit commercial use (Creative Commons 2013). Nevertheless, free riders may use improper means to misuse these models for commercial gain. The misuse of open-access resources is a well-established issue in the traditional software domain, as exemplified by cases such as *Jacobsen v. Katzer* (2008). In the LLM ecosystem, where models possess substantial commercial value, the risk of misuse is further exacerbated by advanced model merging technology. Such techniques allow free riders to incorporate open-access models into their own, making unauthorized appropriation easier. Even worse, existing research indicates that model merging could render watermarks ineffective (Cong et al. 2023), making such stealing difficult to trace. We refer to this threat as *model merging stealing*, which is an emerging and serious form of model theft that is easy to perform and difficult to trace. Given the potential severe impact on the LLM open-source environment, *safeguarding open-access LLMs from unauthorized model merging is a critical issue*.

Unfortunately, as shown in Table 1, traditional solutions struggle to protect open-source models from model merging stealing as they fail to meet the diverse requirements. First, passive protection methods, such as watermarking (Adi et al. 2018; Jia et al. 2021; Liu et al. 2022) and fingerprinting (Chen et al. 2022b,a), only aim to verify model ownership after misuse. Therefore, they are fundamentally incapable of preventing unauthorized merging in the first place. Such passive protection introduces a critical security vulnerability, rendering these methods ineffective: once a model is merged, it may be freely exploited without detection.

In contrast, proactive protections focus on preventing unauthorized usage by enforcing usage constraints, such as limiting model performance when authorization is not pro-

*These authors contributed equally.

† Corresponding author

Solutions (exemplar)	Proactivity	Compatibility	Security with Utility
Passive method (Adi et al. 2018; Chen et al. 2022b)	✗	✓	✓
Authorization-based protection (Zhang et al. 2023)	✓	✗	✓
Emulator-based protection (Xiao, Lin, and Han 2023)	✓	✓	✗
MergeBarrier (ours)	✓	✓	✓

Table 1: Comparison with existing solutions. ✓/✗ illustrates whether the method can achieve the corresponding property.

vided, thereby mitigating misuse. However, a substantial portion of proactive methods rely on additional components, such as key-based (Hanzlik et al. 2021; Fan, Ng, and Chan 2019) or TEE hardware-based (Zhang et al. 2023; Li et al. 2024) authorization systems. The reliance on these additional authorization components creates compatibility issues in typical open-source environments, where such additional components are not feasible.

Alternatively, a potential method compatible with the open-source scenario is to release an emulator, which replaces the original model, that can perform inference independently (Xiao, Lin, and Han 2023; Wang et al. 2024; Junhao, Zhe, and Jun 2025). However, existing emulator-based methods struggle to balance security and utility: simplified emulators reduce utility, while precise ones risk exposing model details. For example, in Xiao et al. (2023)’s method, the emulator is constructed by dropping layers and using knowledge distillation. While this reduces the risk of exposing the model’s inner workings, it also weakens the emulator’s performance because many useful model layers are removed. In contrast, TaylorMLP (Wang et al. 2024) preserves performance by applying Taylor expansion, but it leaves the model vulnerable by only protecting a single layer in each transformer block, exposing the rest of the model’s weights.

Considering the limitations of existing solutions, we identify three challenges (C) to safeguarding open-source LLMs from model merging stealing. **C1 (Proactivity)**: Achieving proactive protection to ensure the model remains non-mergeable, even when the open-sourced model is fully accessible to the attacker. **C2 (Compatibility)**: Achieving protection even when no additional external components can be utilized. **C3 (Security with Utility)**: Securing weights sufficiently even without compromising model utility.

To protect open-source LLMs from unauthorized merging, we propose MergeBarrier, a defense that proactively prevents model merging. Our method is based on the insight that model merging relies on Linear Mode Connectivity (LMC) between homologous models (Ainsworth, Hayase, and Srinivasa 2022; Crisostomi et al. 2024). Specifically, homologous models are fine-tuned from the same base model. As a result, these models are located within a shared low-loss basin in weight space. Therefore, a linear path exists between these models, along which all intermediate points—representing merged models—also remain within the low-loss basin. This ensures that the merged models maintain high performance. Based on this insight, to address **C1**, MergeBarrier disrupts the homogeneity between the protected model and other homologous models. We achieve this by displacing the merged model out of the low-loss basin, thereby preventing merging. This displacement is achieved by transforming the weights of the protected model

into an equivalent set of weights, requiring no external runtime support (address **C2**).

To address **C3**, MergeBarrier ensures the performance-preserving transformations are applied throughout the entire transformer block. Specifically, we apply transformations to both the attention and FFN blocks, tailored to their distinct structures. For the attention block, we leverage its paired linear structure (i.e., query and key linear layers) and apply a shared orthogonal projection to both. This projection is optimized to maximize the deviation between the merged weights and the original weights, thereby displacing the merged model out of the low-loss basin. Meanwhile, due to the properties of the orthogonal matrix, this projection cancels out during paired linear computation, ensuring that the model’s functionality remains unaffected. For FFN, we adopt an activation function expansion-based weight reparameterization that rewrites the FFN computation using a new set of polynomial parameters. This hides the original weights, thus disrupting merging. By selecting the optimal expansion point based on the training set, we ensure that the approximation error has a provably negligible impact on model performance. Interestingly, our analysis reveals that this error can serve as a built-in safeguard by reducing weight inversion to an Learning With Errors (LWE) problem (Regev 2009), i.e., a NP-hard problem, which makes such attacks computationally infeasible.

We evaluate MergeBarrier by protecting expert models across different domains and assessing the performance of merged models in the respective tasks. The evaluation shows that MergeBarrier offers stronger security than existing methods. Once attackers attempt to merge a MergeBarrier-protected model, the resulting merged model suffers a significant performance drop. Besides, MergeBarrier protection does not compromise model accuracy. The contributions of this work are as follows:

- We are the first work that formally defines the problem of defending against *model merging stealing* and systematically identifies the key requirements for protecting open-source models from unauthorized merging.
- We introduce MergeBarrier, a plug-and-play solution that protects model weights from being merged by applying orthogonal projection to attention layers and activation function expansion to FFN layers, effectively disrupting the merging path while fully preserving the original model’s performance.
- We make two foundational contributions to model parameter protection: (1) a broadly applicable projection-based transformation for securing model weights; and (2) a theoretical analysis showing that the noise introduced by Taylor expansion makes weight inversion computa-

tionally infeasible (NP-hard), thereby establishing strong security guarantees for broader applications of such reparameterization methods.

- Extensive experiments demonstrate that MergeBarrier provides a stronger security guarantee compared to existing solutions with negligible accuracy loss. Importantly, we also analyze the underlying reason for its effectiveness in preventing model merging.

2 Preliminaries

2.1 Model Merging

Model merging refers to the technique of combining the weights of multiple specialized models to create a single unified model. A fundamental concept in model merging is the existence of Linear Mode Connectivity (LMC) (Ainsworth, Hayase, and Srinivasa 2022). Specifically, when models are fine-tuned from the same base model, they exhibit certain homologous properties, meaning they share a similar structure and fall within the same low-loss region of the weight space. This homogeneity is crucial for model merging because it ensures that the models are close enough in weight space to be smoothly interpolated without significant loss of performance. Currently, the most promising model merging methods, including Task Arithmetic (Ilharco et al. 2022), TIES-MERGING (Yadav et al. 2023), and DARE (Cong et al. 2023).

2.2 Related work.

Passive Methods. Passive protection techniques, such as watermarking (Adi et al. 2018; Jia et al. 2021; Liu et al. 2022) and fingerprinting (Chen et al. 2022b,a), detect misuse after it occurs. Watermarking embeds unique signals in model parameters for ownership verification, while fingerprinting alters the model’s behavior to track its origin. However, these methods cannot prevent unauthorized merging, leaving the model vulnerable once merged, as ownership may be obscured.

Authorization-based Protection. Authorization-based methods, like key-based (Hanzlik et al. 2021; Fan, Ng, and Chan 2019) and TEE-based (Zhang et al. 2023; Li et al. 2024) protections, restrict access through cryptographic keys or secure hardware. Key-based systems control model usage via authorized keys, but key management is challenging in open-source environments. TEE-based systems ensure execution within trusted environments, providing robust security, but require specialized hardware, making them incompatible with open-source practices.

Emulator-based Protection. Existing methods struggle to balance security and utility: simplified emulators reduce performance, while precise ones risk leaking model details. Xiao et al. (Xiao, Lin, and Han 2023) propose an emulator using layer dropping and knowledge distillation, which protects the model but lowers its utility and requires extensive retraining. To improve utility, TaylorMLP (Wang et al. 2024) uses a Taylor expansion for lossless weight protection, but only safeguards a single linear layer in each transformer block, leaving the model vulnerable. PaRaMS (Junhao, Zhe, and Jun 2025) enhances security by permuting MLP weights

and scaling attention weights, yet in model merging scenarios, these protections fail: permutation can be reversed via weight alignment (Zhang et al. 2023), and scaling can be undone by dividing with base model parameters.

2.3 Threat Model

In this paper, we consider two parties: the defender and the adversary. The defender fine-tunes a pretrained model on specialized data and releases it under a restrictive open-source license. The adversary, for their own benefit, acquires the open-source model and merges it with other models they control, all of which are fine-tuned from the same pretrained base. This scenario is reasonable since most fine-tuning workflows in deep learning begin with widely used pretrained models (Han et al. 2024; Du et al. 2022). The following are the details of the two parties.

Defender’s Capability. The defender can control their model and modify it to ensure protection. Specifically, the defender aims to design proactive defenses that preserve the model’s original performance while significantly degrading its performance when merged with other models.

Adversary’s Capability. The adversary has limited data collection and computational resources, making it infeasible to train a high-quality target model on their own. Instead, they use model merging as a low-cost solution to acquire the defender’s specialized capabilities. The adversary has fine-tuned models based on the same pretrained model and can white-box access to all open-source resources, including the defender’s open-source model, the pretrained base model, and defense strategy papers. They use these resources to optimize the performance of the merged model.

3 Our Design: MergeBarrier

This section introduces MergeBarrier, a proactive defense against model merging stealing.

Key Idea: MergeBarrier is based on the insight that successful model merging requires the expert models to reside within the same low-loss basin, thus forming a linear path where all intermediate points, acting as merged models, lie within the low-loss basin and retain high performance. To prevent model merging, we transform the original model’s weights to displace it into a different loss basin, thereby eliminating the low-loss linear path, thus degrading the merged model’s performance.

3.1 Overview

As shown in Figure 1, MergeBarrier applies the key idea to the attention weight with projection and FFN weights with reparameterization, while preserving the model’s original performance. Specifically, the *attention weight projection* applies a shared orthogonal projection to both the query and key layers, thereby transforming the original parameter values and breaking the homogeneity with other homologous models. Furthermore, we optimize the orthogonal matrix to maximize the deviation between the merged model and the original low-loss region, thereby ensuring that the transformation is sufficiently effective. After protection, due

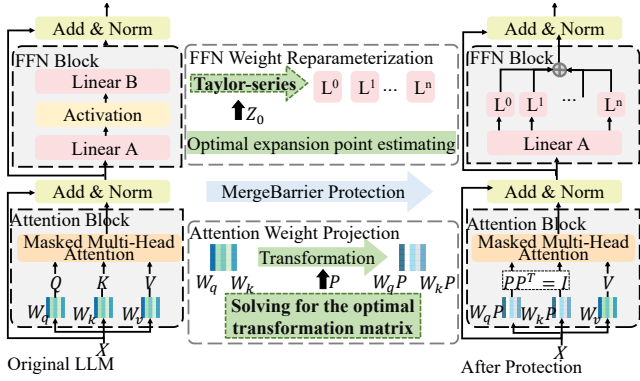


Figure 1: A pipeline of MergeBarrier. The method applies orthogonal projection to attention layers and reparameterization to FFN blocks to disrupt model merging, while preserving the model’s original performance.

to the mathematical properties of the orthogonal matrix, its influence on the attention output is canceled out during computation, thereby preserving model performance.

For FFN blocks, MergeBarrier transforms the FFN weights by reparameterizing the feed-forward computation. We approximate the activation function using an activation function expansion. This transforms the original weights into a new set of functionally equivalent parameters, making the original weights inaccessible. As a result, attackers can no longer perform model merging via interpolation in the original weight space. To preserve model performance, we select a local expansion point that minimizes the gap between the true input and its Taylor approximation. We further analyze the resulting remainder error and demonstrate that this noise acts as a built-in defense by transforming weight inversion into an NP-hard LWE problem (Regev 2009), making such attacks computationally infeasible.

3.2 Attention Weight Projection

Orthogonal Transformation of Weight Matrix. We apply a shared orthogonal transformation to the query and key projection layer. Specifically, given an orthogonal projection matrix P , we apply it to the attention weights as follows:

$$\begin{aligned} O &= \text{softmax} \left(\frac{XW_q P P^T W_k^T X^T}{\sqrt{c}} \right) XW_v \\ &= \text{softmax} \left(\frac{XW_q W_k^T X^T}{\sqrt{c}} \right) XW_v \end{aligned} \quad (1)$$

where W_q , W_k , and W_v represent the weight matrices of the linear layers, and X denotes the input sequence. Since P is an orthogonal matrix, it satisfies $PP^T = I$, which ensures that the attention output remains unchanged for any such orthogonal transformation.

For attention variants such as Multi-Head Attention (MHA) and Group query Attention (GQA) that the above orthogonal transformation remains applicable, leading to the following theorem:

Theorem 1. (a) In case of Multi-Head Attention, there exists a kind of orthogonal matrix \hat{P} like $\text{diag}\{P_1, P_2, \dots, P_n\}$, where n represents the number of heads.

(b) In case of Group Query Attention, there exists a kind of orthogonal matrix \hat{P} like $\text{diag}\{\hat{P}_1, \hat{P}_2, \dots, \hat{P}_n\}$, where n is equal to the number of query in a group.

Optimizing the Orthogonal Matrix. To ensure that the merged model is pushed out of the original low-loss basin, we aim to find an orthogonal matrix P that maximizes its distance from the basin. Specifically, to represent the weight space, we jointly consider the query and key projection weights W_q and W_k , as they together determine the attention output. To locate the low-loss basin, we approximate it using the original model weights, i.e., $W_q W_k^T$. To express the merged model, we adopt weight averaging as a representative merging strategy. Therefore, the merged model after protection is represented as $\frac{1}{4}(W_q P + W_q)(W_k P + W_k)^T$. As a result, our objective can be formulated as maximizing the Frobenius norm of the difference between the merged model and the original basin:

$$\begin{aligned} \max_P \quad & \frac{1}{16} \|(W_q P + W_q)(W_k P + W_k)^T \\ & - (W_q + W_q)(W_k + W_k)^T\|_F^2 \end{aligned} \quad (2)$$

The result is summarized in the following theorem:

Theorem 2. (a) Maximizing Eq. (2) is sufficient to maximize $\|W_q(P - I)\|_F^2 + \|W_k(P - I)\|_F^2$.

(b) If $\lambda_i > 0$, then $(U^T P U)_{ii} = -1$; if $\lambda_i < 0$, then $(U^T P U)_{ii} = 1$.

(c) If $\lambda_i = 0$, then $(U^T P U)_{ii}$ can take any value.

Ideally, the maximum perturbation occurs when $P = -I$, which flips all directions and maximizes the distance, but this would expose real model weights. To solve this, we take a relaxed approach: set $(U^T P U)_{ii} = -1$ for the top- k eigen-directions (i.e., those with the largest Λ values), and 1 for the rest. By applying the inverse transformation, we obtain P that satisfies the required conditions.

Acceleration of Eigenvalue Decomposition. Directly computing the eigendecomposition of high-dimensional matrices like $W_q^T W_q$ is computationally expensive and impractical for large models, potentially making the method infeasible in practice. To address this, we employ Randomized Singular Value Decomposition (RSVD), which exploits the low-rank structure common in neural network weights. RSVD first projects the original matrix into a low-dimensional subspace using a random Gaussian matrix, then performs an SVD decomposition in this reduced-dimensional space, and finally reconstructs the approximation in the original space. This reduces the computational complexity from $O(n^3)$ to $O(n^2 k + nk^2 + k^3)$, where $k \ll n$, making RSVD well suited for large matrices.

3.3 FFN Weight Reparameterization

Unlike attention blocks, FFN blocks lack a symmetric linear structure, making orthogonal transformations inapplicable. To address this, we adopt a reparameterization strategy

tailored for FFN layers. The key idea is to express the forward process using a new set of parameters, so the original weights are no longer explicitly present, thereby relocating the model in weight space and preventing model merging. Specifically, we approximate the activation function using a polynomial expansion and replace the original FFN parameters with the corresponding polynomial coefficients.

Reparameterizing FFN Weight by Polynomial expansion. For an FFN block, let W and c denote the weights and biases of the second linear layer, respectively. The output of this layer is denoted as y . We reparameterize it using a Taylor series as follows:

$$\begin{aligned}
& y - c \\
& = (y_1, \dots, y_n) - (c_1, \dots, c_n) \\
& = (W_1 \odot \text{Act}(z + b), W_2 \odot \text{Act}(z + b), \dots, W_n \odot \text{Act}(z + b)) \\
& \approx (W_1 \odot \sum_{n=0}^N \frac{\text{Act}^{(n)}(z_0 + b)}{n!}, \dots, W_n \odot \sum_{n=0}^N \frac{\text{Act}^{(n)}(z_0 + b)}{n!}) \\
& = W \frac{\text{Act}^{(0)}(z_0 + b)}{0!} (z_0 - z)^0 + \dots \\
& \quad + W \frac{\text{Act}^{(N)}(z_0 + b)}{N!} (z_0 - z)^N \\
& = \hat{W}^0 (z_0 - z)^0 + \dots + \hat{W}^N (z_0 - z)^N.
\end{aligned} \tag{3}$$

Here, W_i is i -th column of W . z is the output of the first linear layer with bias b . The expansion point z_0 is set as the midpoint between z_{max} and z_{min} in the set \mathcal{Z} , which is constructed from the feature representations z of training samples. The symbol \odot denotes element-wise multiplication.

Equation (3) replaces the original weights with a polynomial approximation \hat{W}^i , enabling FFN computation without original weights. Moreover, this approach is broadly applicable, supporting various activation functions commonly used in LLMs, such as *GELU*(\cdot) (Vaswani et al. 2017) and *SilLU*(\cdot) (Elfving, Uchibe, and Doya 2018), as provided in extended version. Additionally, the method is not limited to Taylor expansion—alternative bases, such as Hermite polynomials, can also be adopted.

Discussion on Runtime Efficiency. MergeBarrier protects the MLP by modifying the model structure, which increases the computational cost of MLP blocks. However, in practice, the impact on runtime efficiency is minimal since computations at each order can run in parallel, allowing little to no increase in overall MLP block inference time.

Analyzing the Taylor Remainder. The remainder term in a Taylor expansion inherently depends on higher-order derivatives and the expansion point, introducing noise that can be modeled probabilistically. Interestingly, our analysis reveals that this noise can act as a built-in safeguard against potential attacks. For example, a potential attack strategy is to invert the Taylor expansion to recover the original weights. However, direct inversion is fundamentally infeasible: the system is underdetermined with more unknowns than equations, making its exact value inaccessible. Even if the attacker fully leverages the base model parameters as prior knowledge during the inversion process, the attempt

still fails. Specifically, the presence of the remainder term injects uncertainty into the relationship between the observed coefficients and the true weights. As a result, recovering the original weights reduces to solving a Learning With Errors (LWE) problem, which is widely regarded as NP-hard.

Additionally, we observe that the remainder-induced noise can serve as intrinsic noise that satisfies differential privacy (DP). Therefore, it enhances the weight’s privacy, e.g., defending against membership inference attacks.

4 Experiments

4.1 Experimental Settings

Upstream LLMs. We evaluate MergeBarrier by protecting expert models from different domains, including general language understanding, mathematical reasoning, and code generation. We use LLaMA-2-13B (Touvron et al. 2023) as the base model. We evaluate three expert models, all fine-tuned from LLaMA-2-13B: WizardLM-13B (Xu et al. 2023) (*LM*) for general language understanding, WizardMath-13B (Luo et al. 2023) (*Math*) for mathematical reasoning, and LLaMA-2-13B-Code Alpaca (Chaudhary 2023) (*Code*) for programming and code generation.

Dataset and Metrics. We evaluate the merged model on the representative dataset corresponding to the protected expert model’s domain. **(1) Instruction-following:** AlpacaEval (Li et al. 2023) tests a model’s ability to follow instructions and produce relevant responses. We report the win rate. **(2) Math:** GSM8K (Cobbe et al. 2021) contains grade-school math problems requiring multi-step reasoning and basic computation. Performance is measured by zero-shot accuracy. **(3) Code:** MBPP (Austin et al. 2021) assesses Python code generation for programming tasks. We use pass@1, the percentage of problems solved correctly on the first attempt.

Baselines. To compare our proposed methods, we selected two baselines, PaRaMS (Junhao, Zhe, and Jun 2025) and TaylorMLP (Wang et al. 2024), as they represent the most pertinent approaches in the same threat model. PaRaMS is specifically designed to counter model merging stealing, directly aligning with our goal. TaylorMLP aims to protect model weights while preserving utility, without relying on additional components.

Evaluated Attack. Following the prior work (Junhao, Zhe, and Jun 2025), we examine three representative model merging techniques as the basis for model merging stealing: Task Arithmetic (Ilharco et al. 2022), TIES-merging (Yadav et al. 2023), and DARE (Yu et al. 2024). For defenses that modify model structures (e.g., MergeBarrier and TaylorMLP), attackers revert modified layers to the base model’s original weights before merging. For PaRaMS, the attackers decode the protected model using the base model and apply the recovered weights.

4.2 Security

Security Against Model Merging Attacks. In this subsection, we evaluate whether MergeBarrier effectively defends

Merging methods	Expert models	Without protection			MergeBarrier (Ours)			PaRaMS			TaylorMLP		
		Alpaca	GSM8K	MBPP	Alpaca	GSM8K	MBPP	Alpaca	GSM8K	MBPP	Alpaca	GSM8K	MBPP
Task	<i>LM & Math</i>	26.87	62.17	35.60	0.14	0.00	-	26.51	59.60	-	22.73	55.25	-
	<i>LM & Code</i>	40.85	-	21.80	0.16	-	0.00	40.80	-	35.00	37.83	-	31.40
	<i>Math & Code</i>	-	58.38	21.80	-	0.00	0.00	-	58.64	21.40	-	49.64	19.00
	<i>Three Models</i>	30.81	59.82	33.80	3.11	5.11	0.00	30.47	59.18	33.10	27.83	53.11	30.80
TIES	<i>LM & Math</i>	43.28	24.93	-	3.85	3.75	-	42.85	23.61	-	35.07	21.94	-
	<i>LM & Code</i>	44.62	-	18.00	5.04	-	0.00	45.28	-	7.50	44.96	-	16.00
	<i>Math & Code</i>	-	63.76	20.00	-	6.94	0.00	-	64.68	20.00	-	45.09	16.80
	<i>Three Models</i>	44.65	67.55	30.80	6.20	9.81	0.00	44.08	67.07	29.80	42.87	51.94	24.80
DARE-Task	<i>LM & Math</i>	33.42	60.17	-	9.85	4.01	-	32.75	61.70	-	30.91	59.91	-
	<i>LM & Code</i>	37.64	-	34.40	9.71	-	0.00	37.25	-	33.80	34.88	-	26.80
	<i>Math & Code</i>	-	59.37	23.80	-	6.71	0.00	-	58.43	23.20	-	53.92	12.90
	<i>Three Models</i>	35.83	51.82	34.60	12.61	10.04	0.00	29.99	60.07	34.60	25.82	52.18	30.90
DARE-TIES	<i>LM & Math</i>	45.84	30.39	-	10.84	10.98	-	46.27	29.30	-	41.42	26.85	-
	<i>LM & Code</i>	44.96	-	34.80	10.58	-	0.80	44.51	-	34.30	34.95	-	35.40
	<i>Math & Code</i>	-	58.78	17.80	-	9.87	3.00	-	58.18	16.80	-	53.81	13.80
	<i>Three Models</i>	45.71	49.05	25.60	11.09	13.81	7.50	46.42	48.41	25.70	32.81	42.94	21.40
Relative Mean Accuracy ↓		8.27×			1.00×			8.15×			7.17×		

Table 2: Security assessment of MergeBarrier in preventing model merging stealing. We report the accuracy (%) of the merged model across different settings. Notably, in each setting, only one expert model is protected, and we exclusively report the merged model’s performance on its corresponding task. The last row shows the average attack accuracy as a multiple of our method (\times).




Figure 2: Defense effectiveness against model merging stealing with fine-tuning.

against model merging stealing. To this end, we simulate an attacker who merges the protected model and reports the merged model’s performance on the corresponding task. If its task performance drops significantly, we consider the defense has successfully mitigated the attack.

As shown in Table 2, MergeBarrier consistently suppresses the merged model’s performance across all merging settings. For example, in the *Task* setting, MBPP accuracy drops to 0.00% in both the “LM & Code” and “Three Models” cases, and GSM8K accuracy in “LM & Math” drops from 62.17% to 0.00%. Similar trends are observed under *TIES* and *DARE*. To provide a holistic comparison, we compute the average merged model accuracy across all settings. MergeBarrier achieves the lowest score (1.00 \times), indicating near-complete suppression of merging-based theft. In contrast, PaRaMS and TaylorMLP fail to provide effective protection, with the merged models retaining up to 31.10% and 30.80% MBPP accuracy, and significantly higher average scores of 8.15 \times and 7.17 \times , respectively.

This strong protection stems from MergeBarrier’s comprehensive protection. Specifically, attention projection disrupts cross-model weight alignment, while FFN reparameterization conceals true weight. In contrast, PaRaMS and TaylorMLP either fail under adaptive attacks or safeguard only a limited subset of weights.

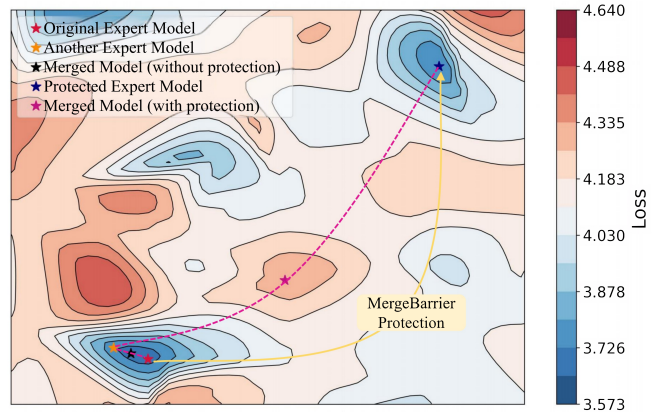


Figure 3: Loss landscape before and after our protection.

Security Against Stealing with Fine-Tuning. In this section, we consider a stronger attacker who further fine-tunes the merged model with limited labeled data, increasing its task performance, thus increasing the threat. We report the task performance of different defenses under this enhanced model merging stealing scenario, across two merging strategies and two target tasks. As shown in Figure 2, MergeBarrier remains highly robust: even with more labeled data, the merged model’s performance improves only slightly. In contrast, TaylorMLP is easily compromised: minor fine-tuning significantly improves task performance. This strong resistance stems from MergeBarrier’s comprehensive protection over both attention and FFN parameters.

4.3 LMC Analysis

We analyze how MergeBarrier affects Linear Mode Connectivity (LMC) by constructing weight-loss plots to visualize the relationship between models in weight space. Specifically, we apply PCA to reduce the dimensionality of model weights and project them into a 2D space, where we report the original loss values corresponding to each model in

Models	Accuracy \uparrow			Sharpness \downarrow		
	Alpaca	GSM8K	MBPP	Alpaca	GSM8K	MBPP
WizardLM-13B	45.59 / 45.53	-	-	0.53 / 0.54	-	-
WizardMath-13B	-	58.76 / 57.88	-	-	0.57 / 0.58	-
LLaMA-2-13B-Code	-	-	27.80 / 27.80	-	-	0.63 / 0.65

Table 3: The performance comparison between the original model (M) and the MergeBarrier-protected model (M'). The result is presented in the form of M / M' .

Merging methods	Expert models	Without protection			MergeBarrier (Ours)			MergeBarrier $_{Att}$			MergeBarrier $_{FFN}$		
		Alpaca	GSM8K	MBPP	Alpaca	GSM8K	MBPP	Alpaca	GSM8K	MBPP	Alpaca	GSM8K	MBPP
Task	<i>LM & Math</i>	26.87	62.17	-	0.14	0.00	-	12.45	21.20	-	23.61	56.12	-
	<i>LM & Code</i>	40.85	-	35.60	0.16	-	0.00	14.61	-	11.90	35.66	-	30.90
	<i>Math & Code</i>	-	58.38	21.80	-	0.00	0.00	-	13.26	15.30	-	49.83	18.80
	<i>Three Models</i>	30.81	59.82	33.80	3.11	5.11	0.00	23.26	25.35	24.10	28.62	42.56	31.60
TIES	<i>LM & Math</i>	43.28	24.93	-	3.85	3.75	-	15.12	18.93	-	34.49	22.62	-
	<i>LM & Code</i>	44.62	-	18.00	5.04	-	0.00	11.21	-	6.30	44.70	-	16.60
	<i>Math & Code</i>	-	63.76	20.00	-	6.94	0.00	-	44.22	15.30	-	61.33	17.00
	<i>Three Models</i>	44.65	67.55	30.80	6.20	9.81	0.00	22.93	49.18	17.50	42.72	52.12	25.40
DARE- Task	<i>LM & Math</i>	33.42	60.17	-	9.85	4.01	-	12.71	11.51	-	31.37	59.87	-
	<i>LM & Code</i>	37.64	-	34.40	9.71	-	0.00	10.23	-	2.80	35.14	-	26.60
	<i>Math & Code</i>	-	59.37	23.80	-	6.71	0.00	-	11.00	6.30	-	53.90	12.50
	<i>Three Models</i>	35.83	51.82	34.60	12.61	10.04	0.00	26.92	20.56	9.70	25.96	52.26	30.50
DARE-TIES	<i>LM & Math</i>	45.84	30.39	-	10.84	10.98	-	24.12	23.24	-	41.90	27.44	-
	<i>LM & Code</i>	44.96	-	34.80	10.58	-	0.80	24.67	-	12.30	34.76	-	29.20
	<i>Math & Code</i>	-	58.78	17.80	-	9.87	3.00	-	20.14	9.40	-	53.52	15.10
	<i>Three Models</i>	45.71	49.05	25.60	11.09	13.81	7.50	28.33	35.93	12.70	33.20	43.04	20.90
Relative Mean Accuracy \downarrow		8.27\times			1.00\times			3.79\times			7.19\times		

Table 4: Security assessment of MergeBarrier and its variants in preventing model merging stealing.

weight space. As shown in Figure 3, before protection, the original model and the another expert lie within the same low-loss basin, so the linear path connecting them also remains in this basin. This ensures that the merged model, which is located along the linear path, also resides in a low-loss region and thus maintains good performance. After MergeBarrier protection, the original model is relocated to a different basin, with no low-loss path connecting it to the another expert model. In other words, LMC has been disrupted, and any merged model formed along this path will suffer from high loss and degraded performance.

4.4 Model Performance

We assess MergeBarrier’s impact by comparing the performance of the protected model with the original model using both **accuracy** and **sharpness** metrics. Accuracy reflects performance under ideal conditions, where the model is evaluated with its exact trained weights. However, real-world deployments often introduce slight weight deviations, e.g., due to quantization or low-precision hardware, that may degrade performance. As MergeBarrier modifies model weights, it may affect the model’s robustness to such deviations. To account for this, we also measure sharpness, which quantifies the loss sensitivity to small parameter deviations. A lower sharpness indicates flatter minima and greater weight robustness.

As shown in Table 3, MergeBarrier introduces a negligible impact on both metrics. Accuracy is preserved (e.g., LLaMA-2-13B-Code: 27.80% vs. 27.80%), and sharpness remains stable (e.g., GSM8K: 0.57 to 0.58). These strong results stem from our targeted design: orthogonal projection in attention and functionally equivalent reparameterization in FFN are both crafted to preserve performance.

4.5 Ablation Study

We assess the individual contribution of each MergeBarrier component by comparing the full method with two variants: MergeBarrier $_{Att}$ (only attention projection) and MergeBarrier $_{FFN}$ (only FFN reparameterization). As shown in Table 4, individually protecting either the attention or FFN module provides only limited defense against model merging attacks. Specifically, the full MergeBarrier achieves the lowest relative mean accuracy for the merged model (1.00 \times), while the attention-only and FFN-only variants reach 3.79 \times and 7.19 \times , respectively. These results demonstrate that robust protection against model merging requires the joint deployment of both attention projection and FFN reparameterization. In contrast, partial protection leaves significant vulnerabilities exploitable by attackers. This ablation study thus empirically validates the necessity and effectiveness of holistic defense in MergeBarrier.

5 Conclusion

In this work, we address the emerging threat of model merging stealing, where unauthorized users exploit proprietary models through merging. We highlight key limitations in current defense strategies and stress the need for secure, compatible methods that do not compromise model performance. To this end, we introduce MergeBarrier, a proactive defense mechanism that prevents unauthorized merging without impacting model utility. Extensive experiments show that MergeBarrier outperforms existing methods, providing strong protection with minimal accuracy loss. Theoretical analysis also confirms that these transformations ensure differential privacy. In conclusion, MergeBarrier is an effective solution that provides model owners with the means to safeguard their proprietary LLMs.

Acknowledgements

This work was supported by the Key Project of the National Natural Science Foundation of China under Grant no. 62536007, the Zhejiang Province Science Foundation under Grant no. LD24F020002 and the Zhejiang Province's 2025 "Leading Goose + X" Science and Technology Plan under Grant no. 2025C02034.

References

- Adi, Y.; Baum, C.; Cisse, M.; Pinkas, B.; and Keshet, J. 2018. Turning your weakness into a strength: Watermarking deep neural networks by backdooring. In *27th USENIX Security Symposium (USENIX Security 18)*, 1615–1631.
- Ainsworth, S. K.; Hayase, J.; and Srinivasa, S. 2022. Git re-basin: Merging models modulo permutation symmetries. *arXiv preprint arXiv:2209.04836*.
- Arora, A.; He, X.; Mozes, M.; Swain, S.; Dras, M.; and Xu, Q. 2024. Here's a free lunch: Sanitizing backdoored models with model merge. *arXiv preprint arXiv:2402.19334*.
- Austin, J.; Odena, A.; Nye, M.; Bosma, M.; Michalewski, H.; Dohan, D.; Jiang, E.; Cai, C.; Terry, M.; Le, Q.; et al. 2021. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*.
- Bhardwaj, R.; Anh, D. D.; and Poria, S. 2024. Language models are Homer simpson! safety re-alignment of fine-tuned language models through task arithmetic. *arXiv preprint arXiv:2402.11746*.
- Cassano, F.; Gouwar, J.; Lucchetti, F.; Schlesinger, C.; Freeman, A.; Anderson, C. J.; Feldman, M. Q.; Greenberg, M.; Jangda, A.; and Guha, A. 2024. Knowledge transfer from high-resource to low-resource programming languages for code llms. *Proceedings of the ACM on Programming Languages*, 8(OOPSLA2): 677–708.
- Chaudhary, S. 2023. Code alpaca: An instruction-following llama model for code generation.
- Chen, J.; Wang, J.; Peng, T.; Sun, Y.; Cheng, P.; Ji, S.; Ma, X.; Li, B.; and Song, D. 2022a. Copy, right? a testing framework for copyright protection of deep learning models. In *2022 IEEE symposium on security and privacy (SP)*, 824–841. IEEE.
- Chen, Y.; Shen, C.; Wang, C.; and Zhang, Y. 2022b. Teacher model fingerprinting attacks against transfer learning. In *31st USENIX Security Symposium (USENIX Security 22)*, 3593–3610.
- Cobbe, K.; Kosaraju, V.; Bavarian, M.; Chen, M.; Jun, H.; Kaiser, L.; Plappert, M.; Tworek, J.; Hilton, J.; Nakano, R.; et al. 2021. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*.
- Cong, T.; Ran, D.; Liu, Z.; He, X.; Liu, J.; Gong, Y.; Li, Q.; Wang, A.; and Wang, X. 2023. Have you merged my model? on the robustness of large language model ip protection methods against model merging. In *Proceedings of the 1st ACM Workshop on Large AI Systems and Models with Privacy and Safety Analysis*, 69–76.
- Creative Commons. 2013. Attribution-NonCommercial-NoDerivatives 4.0 International (CC BY-NC-ND 4.0). <https://creativecommons.org/licenses/by-nc-nd/4.0/>. Accessed: 2025-05-08.
- Crisostomi, D.; Fumero, M.; Baieri, D.; Bernard, F.; and Rodola, E. 2024. C²M³: Cycle-Consistent Multi-Model Merging. *Advances in Neural Information Processing Systems*, 37: 28674–28705.
- Du, Y.; Liu, Z.; Li, J.; and Zhao, W. X. 2022. A survey of vision-language pre-trained models. *arXiv preprint arXiv:2202.10936*.
- Elfwing, S.; Uchibe, E.; and Doya, K. 2018. Sigmoid-weighted linear units for neural network function approximation in reinforcement learning. *Neural networks*, 107: 3–11.
- Face, H. 2024. Hugging Face – The AI community building the future. <https://huggingface.co>. Accessed: 2025-05-09.
- Fan, L.; Ng, K. W.; and Chan, C. S. 2019. Rethinking deep neural network ownership verification: Embedding passports to defeat ambiguity attacks. *Advances in neural information processing systems*, 32.
- Han, Z.; Gao, C.; Liu, J.; Zhang, J.; and Zhang, S. Q. 2024. Parameter-efficient fine-tuning for large models: A comprehensive survey. *arXiv preprint arXiv:2403.14608*.
- Hanzlik, L.; Zhang, Y.; Grosse, K.; Salem, A.; Augustin, M.; Backes, M.; and Fritz, M. 2021. Mlcapule: Guarded offline deployment of machine learning as a service. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3300–3309.
- Hu, E. J.; Shen, Y.; Wallis, P.; Allen-Zhu, Z.; Li, Y.; Wang, S.; Wang, L.; Chen, W.; et al. 2022. Lora: Low-rank adaptation of large language models. *ICLR*, 1(2): 3.
- Ilharco, G.; Ribeiro, M. T.; Wortsman, M.; Gururangan, S.; Schmidt, L.; Hajishirzi, H.; and Farhadi, A. 2022. Editing models with task arithmetic. *arXiv preprint arXiv:2212.04089*.
- Jia, H.; Choquette-Choo, C. A.; Chandrasekaran, V.; and Papernot, N. 2021. Entangled watermarks as a defense against model extraction. In *30th USENIX Security Symposium (USENIX Security 21)*, 1937–1954.
- Junhao, W.; Zhe, Y.; and Jun, S. 2025. Disrupting Model Merging: A Parameter-Level Defense Without Sacrificing Accuracy. *arXiv preprint arXiv:2503.07661*.
- Li, Q.; Shen, Z.; Qin, Z.; Xie, Y.; Zhang, X.; Du, T.; and Yin, J. 2024. TransLinkGuard: Safeguarding Transformer Models Against Model Stealing in Edge Deployment. *arXiv preprint arXiv:2404.11121*.
- Li, X.; Zhang, T.; Dubois, Y.; Taori, R.; Gulrajani, I.; Guestrin, C.; Liang, P.; and Hashimoto, T. B. 2023. Alpaca-eval: An automatic evaluator of instruction-following models.
- Liu, G.; Xiang, R.; Liu, J.; Pan, R.; and Zhang, Z. 2022. An invisible and robust watermarking scheme using convolutional neural networks. *Expert Systems with Applications*, 210: 118529.
- Luo, H.; Sun, Q.; Xu, C.; Zhao, P.; Lou, J.; Tao, C.; Geng, X.; Lin, Q.; Chen, S.; and Zhang, D. 2023. Wizard-math: Empowering mathematical reasoning for large language models via reinforced evol-instruct. *arXiv preprint arXiv:2308.09583*.

OpenAI. 2023. GPT-4 Technical Report. <https://openai.com/research/gpt-4>.

Regev, O. 2009. On lattices, learning with errors, random linear codes, and cryptography. *Journal of the ACM (JACM)*, 56(6): 1–40.

Roziere, B.; Gehring, J.; Gloeckle, F.; Sootla, S.; Gat, I.; Tan, X. E.; Adi, Y.; Liu, J.; Sauvestre, R.; Remez, T.; et al. 2023. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*.

Touvron, H.; Martin, L.; Stone, K.; Albert, P.; Almahairi, A.; Babaei, Y.; Bashlykov, N.; Batra, S.; Bhargava, P.; Bhosale, S.; et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*.

v. Katzer, J. 2008. Jacobsen v. Katzer - Wikipedia. https://en.wikipedia.org/wiki/Jacobsen_v._Katzer. Accessed: 2025-05-08.

Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. *Advances in neural information processing systems*, 30.

Wang, G.; Chuang, Y.-N.; Tang, R.; Zhong, S.; Yuan, J.; Jin, H.; Liu, Z.; Chaudhary, V.; Xu, S.; Caverlee, J.; et al. 2024. Taylor Unswift: Secured Weight Release for Large Language Models via Taylor Expansion. *arXiv preprint arXiv:2410.05331*.

Xiao, G.; Lin, J.; and Han, S. 2023. Offsite-tuning: Transfer learning without full model. *arXiv preprint arXiv:2302.04870*.

Xu, C.; Sun, Q.; Zheng, K.; Geng, X.; Zhao, P.; Feng, J.; Tao, C.; and Jiang, D. 2023. Wizardlm: Empowering large language models to follow complex instructions. *arXiv preprint arXiv:2304.12244*.

Yadav, P.; Tam, D.; Choshen, L.; Raffel, C. A.; and Bansal, M. 2023. Ties-merging: Resolving interference when merging models. *Advances in Neural Information Processing Systems*, 36: 7093–7115.

Yamabe, S.; Takahashi, T.; Waseda, F.; and Wataoka, K. 2024. MergePrint: Robust Fingerprinting against Merging Large Language Models. *arXiv preprint arXiv:2410.08604*.

Yang, E.; Shen, L.; Guo, G.; Wang, X.; Cao, X.; Zhang, J.; and Tao, D. 2024. Model merging in llms, mllms, and beyond: Methods, theories, applications and opportunities. *arXiv preprint arXiv:2408.07666*.

Yu, L.; Yu, B.; Yu, H.; Huang, F.; and Li, Y. 2024. Language models are super mario: Absorbing abilities from homologous models as a free lunch. In *Forty-first International Conference on Machine Learning*.

Zhang, Z.; Gong, C.; Cai, Y.; Yuan, Y.; Liu, B.; Li, D.; Guo, Y.; and Chen, X. 2023. No Privacy Left Outside: On the (In-) Security of TEE-Shielded DNN Partition for On-Device ML. In *2024 IEEE Symposium on Security and Privacy (SP)*, 52–52. IEEE Computer Society.