BOOSTING LARGE LANGUAGE MODELS WITH MASK FINE-TUNING

Anonymous authors
Paper under double-blind review

ABSTRACT

The large language model (LLM) is typically integrated into the mainstream optimization protocol. However, it remains underexplored whether maintaining the model integrity is *indispensable* for promising performance. In this work, we introduce Mask Fine-Tuning (MFT), a novel LLM fine-tuning paradigm demonstrating that carefully breaking the model's structural integrity can surprisingly improve performance without updating model weights. MFT learns and applies binary masks to well-optimized models, using the standard LLM fine-tuning objective as supervision. Based on fully fine-tuned models, MFT uses the same fine-tuning datasets to achieve consistent performance gains across domains and backbones (e.g., an average gain of **2.70/4.15** on IFEval with LLaMA2-7B/3.1-8B). Detailed ablation studies and analyses examine the proposed MFT from different perspectives, including the sparse ratio and the loss surface. Additionally, when deployed on well-trained models, MFT is compatible with other LLM optimization procedures to improve overall model performance. Furthermore, this study extends the masking operation beyond its conventional use in network pruning for model compression to encompass a broader range of model capabilities.

1 INTRODUCTION

Pre-training large language models (LLMs) typically requires massive corpora to enable models with foundational knowledge and linguistic competencies for effective language generation (Radford et al., 2019; Brown et al., 2020; Touvron et al., 2023; Grattafiori et al., 2024). Afterward, the pre-trained LLMs are fine-tuned for downstream tasks using high-quality domain data, including specialized knowledge such as math and coding (Cobbe et al., 2021; Yu et al., 2024; Chen et al., 2021; Austin et al., 2021) and particular patterns such as human instruction (Zhou et al., 2023; Dubois et al., 2024). For such adaptations, fine-tuning all model parameters is the most straightforward and common approach, commonly referred to as full fine-tuning (FFT), and generally yields the most competitive performance. Furthermore, adapter-based alternatives offer parameter-efficient fine-tuning (PEFT) strategies such as vanilla bottleneck adapters (Houlsby et al., 2019) and LoRA (Hu et al., 2022). These methods fix the pre-trained backbone and additionally introduce a limited number of learnable parameters. Although these methods may sacrifice fine-tuning performance relative to FFT, they suit scenarios with limited computational budgets, especially when fine-tuning data are insufficient. Overall, the pipeline of language model pre-training followed by fine-tuning has demonstrated promising results across a range of language tasks, including domains requiring extensive specialized knowledge (e.g., medicine (Wang et al., 2023) and law (Nguyen, 2023)) and complex reasoning (e.g., math (Frieder et al., 2024) and coding (Chen et al., 2021)), exhibiting remarkably intelligent capabilities.

However, the mainstream LLM optimization process (pre-training followed by fine-tuning), which is promisingly used to benchmark modern LLM performance, consistently treats the language model as a whole. For pretraining, the model's structure is kept mostly dense. For fine-tuning, FFT optimizes all model parameters simultaneously, while PEFT freezes them to tune additional adapter parameters, with both approaches accepting the necessity of LLM structural integrity by default. This practice naturally leads us to ask: *Is such structural integrity indispensable for good model performance? Or, more broadly, can we improve the model by removing certain components that break structural integrity?*

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

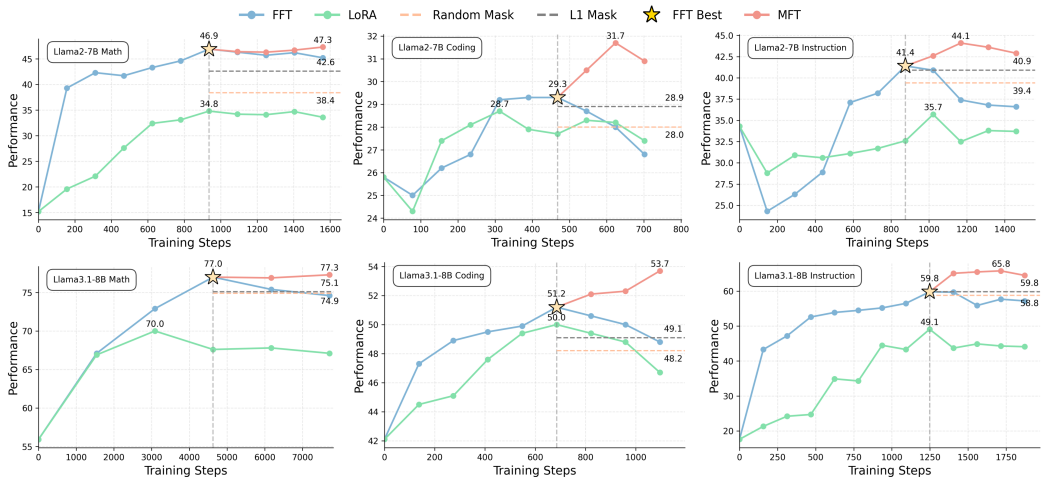


Figure 1: The visualization of the performance trend across different fine-tuning strategies, including FFT (blue line), LoRA (green line) (Hu et al., 2022), and our MFT (red line). We also add random (orange dashed) and L1 (gray dashed) masks for comparison. We use three settings across LLaMA and Qwen backbones on GSM8K, HumanEval, and IF-Eval for math, coding, and instruction domains, respectively. The x-axis is training steps starting from the pre-trained backbone. The y-axis is evaluation performance. MFT (red line) starts from the best FFT model (yellow star) and improves upon the upper bound, whereas continued FFT leads to overfitting. It also outperforms LoRA fine-tuning and two vanilla mask baselines.

We propose *mask fine-tuning (MFT)* to answer this question. MFT freezes a given LLM model and learns a binary mask over it to assess whether breaking structural integrity works. To substantiate its effectiveness, we use a well-trained LLM (the model after sufficient FFT in this study) as a strong baseline to assess whether MFT yields further improvement. MFT uses the same standard LLM fine-tuning objective and datasets as FFT to learn the mask. The only difference is that the learnable part shifts from model weights to a binary weight mask, while the weights themselves remain fixed. The learned mask indicates positions of certain parameters to be removed. Surprisingly, we find that masking these weights yields further performance gains compared with the given well-trained model. As shown in Fig. 1, the regular FFT (blue line) improves the pre-trained model across settings, but once it reaches peak performance (yellow star), excessive FFT leads to overfitting. However, MFT (red line) starts from the best FFT model and further improves it. These teaser cases show a consistent observation across different backbones (LLaMA2-7B (Touvron et al., 2023) and LLaMA3.1-8B (Grattafiori et al., 2024)) and domains (GSM8K (Cobbe et al., 2021), HumanEval (Chen et al., 2021), and IF-Eval (Zhou et al., 2023) for math, coding, and instruction, respectively). Comparisons of more configurations are elaborated in Sec. 3. This phenomenon substantially answers the question we are interested in: *LLM structural integrity is not indispensable for good model performance, and breaking it can lead to further improvements*.

Here, we primarily explore MFT as a post-fine-tuning strategy for further improvement, starting from an FFT checkpoint that yet outperforms. It naturally upgrades the current fine-tuning routine to a new protocol, as shown in Fig. 2, offering a new perspective on investigating LLM fine-tuning. Furthermore, MFT is compatible with existing fine-tuning methods and can be integrated flexibly into any of them, as it shares the same optimization objective and requires no additional data annotation.

In addition, we highlight the difference between typical pruning methods and our work. The former compresses the model while preserving its trained capabilities, whereas the latter seeks further improvement beyond a well-trained model for a specific purpose. Even if they share the same concept of *mask* (or model sparsity), their fundamental goals are different (see more discussions in Sec. 4). Along these lines, our study extends the masking operation beyond typical network compression to a more general model-capability setting. In other words, analogically speaking, *typical masking*

employs subtraction to reduce (pruning to compress), whereas our approach leverages subtraction to achieve augmentation (removing weights to improve). Our contributions are as follows:

- We validate that a well-trained LLM can be further improved by carefully removing certain weights using the same objective and dataset as regular fine-tuning, with limited computational overhead. Such a strategy is compatible with other training pipelines and enhances the current LLM optimization.
- We propose *Mask Fine-Tuning (MFT)* to complete our exploration in a post fine-tuning scenario. MFT starts from a competitive, fully fine-tuned model with fixed weights and learns a binary mask that is applied to the model to improve it. Following this, MFT treats model sparsity in a new light, not only for efficiency, but also for performance improvement.
- Extensive experiments across different backbones (LLaMA and Qwen), domains (math, coding, and instruction), and FFT settings (domain-specific and mixed-up) show MFT's effectiveness, with consistent performance gains. Detailed ablations and analyses are provided to provide better intuition and inspire future work.

2 Mask Fine-Tuning

We focus on GPT-like language models (Brown et al., 2020; Radford et al., 2019) in an autoregressive manner and begin with a fully fine-tuned model built on a pre-trained language backbone for masked fine-tuning (MFT). We briefly introduce the language model notation for full fine-tuning (FFT), then present our MFT.

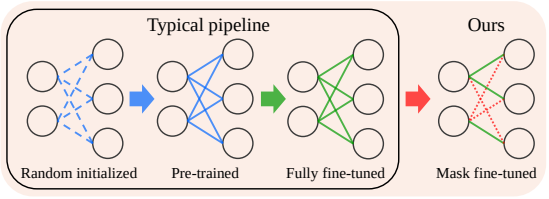


Figure 2: Typical LLM training comprises pre-training and fine-tuning to build capacity and domain knowledge, with the structure always being the entire model. We are curious whether such integrity is necessary and propose that MFT generally outperforms models with sufficient FFT, naturally upgrading the classic pipeline, following the typical protocol, to further refine well-optimized LLMs.

Full Fine-Tuning. We refer a pre-trained auto-regressive language backbone as $\mathcal{N}_p$ with parameters Θ_p and optimize the following objective to represent a FFT:

$$L(U_f) = \sum_i logP(u_f^i | u_f^{i-k}, ..., u_f^{i-1}; \Theta_p), \tag{1}$$

where L represents the autoregressive loss to supervise the next token prediction. Θ_p is fully optimized, and $U_f = \{u_f^1,\ dots, u_f^n\}$ is a token sequence from a language corpus, serving as an FFT dataset. We refer to the model obtained after sufficient FFT as $\mathcal{N}_f$ and denote its optimized parameters by Θ_f.

Mask Fine-Tuning. We deploy our MFT on models after the FFT. Formally, we revise Eq. equation 1 by adding a binary mask onto model parameters and obtain the rewritten loss for MFT as below:

$$L(U_m) = \sum_i logP(u_m^i | u_m^{i-k}, ..., u_m^{i-1}; \Theta_f \odot M), \tag{2}$$

where it shares the same objective function L as in FFT, using the dataset U_m for supervision. In our experiments, the dataset U_m for MFT is the same as U_f for FFT. The main difference is that we add a binary mask M to the model parameters Θ_f. M and Θ_f have the same size, and we mask out the corresponding weights in Θ_f via element-wise multiplication $\odot$. During MFT, Θ_f is fixed, and M is updated. This way, MFT learns to locate specific weights while keeping the fully fine-tuned parameters Θ_f unchanged. These weights can be removed from the well-trained model by applying the learned binary mask. We optimize M using a learnable process based on the straight-through gradient estimator (Bengio et al., 2013), which we call the learnable mask.

Learnable Mask. Given a model parameter Θ_f after FFT, we represent it layer-wise as $\Theta = \{\theta_l\}, l \in \{1, 2, ..., L\}$, omitting the subscript f for notational simplicity. L represents the model layer number. Then, we can formalize each layer as

$$I_{l+1} = \sigma(F[I_l; \theta_l]), \tag{3}$$

where I_l is the layer l input and I_{l+1} is its output with activation σ. F generally serves as a layer operation (e.g., a convolutional or linear layer) with a parameter $\theta_l = \{\theta_l^d\}, d \in \{1, 2, ..., D_l\}$. D_l represents the layer l parameter dimension. In this work, F denotes a linear layer, as we primarily consider Transformer-based language models and apply the learnable mask to the model's linear mappings. To perform mask learning, we fix all parameters Θ, denoted $\overline{\Theta}$. Then, we assign each weight $\overline{\theta_l^d}$ a learnable score c_l^d, representing the importance of the corresponding weight. Based on this, we rewrite Eq. equation 3 as

$$I_{l+1} = \sigma(F[I_l; \overline{\theta_l} \odot v(c_l)]), \tag{4}$$

where $c_l = \{c_l^1, c_l^2, ..., c_l^{D_l}\}$ is the score of parameter $\overline{\theta_l}$. v is an indicator function to select a mask based on scores. In this work, we use a ratio-based version in which v outputs 1 if c_l^d is among the top $K\%$ highest-scoring values, and 0 otherwise, where K is a pre-defined mask ratio. By updating c with fixed $\overline{\Theta}$, we keep some parameters in $\overline{\Theta}$ while masking out others. As the indicator function v is non-differentiable, we use the straight-through gradient estimator (Bengio et al., 2013) to estimate the gradient of the loss with respect to c_l^d. Concretely, we treat v as an identity function during the backward pass. Then, the approximation can be described as

$$\tilde{g}(c_l^d) = \frac{\partial L}{\partial \tilde{I}_{l+1}} \frac{\partial \tilde{I}_{l+1}}{\partial c_l^d} \approx \frac{\partial L}{\partial I_{l+1}} \frac{\partial I_{l+1}}{\partial c_l^d}, \tag{5}$$

where $\tilde{I}_{l+1} = \sigma(F[I_l; \overline{\theta}_l \odot c_l])$ after applying estimation and $\tilde{g}(c_l^d)$ is the estimated gradient with respect to score c_l^d. In this way, MFT learns to obtain a mask for the model $\mathcal{N}_f$, and we denote the resulting model as $\mathcal{N}_m$ with parameters Θ_f and mask M.

MFT learns to identify specific weights using the regular training loss of language models, then removes them by applying the learned mask. It answers our question (Sec. 1) and introduces a new LLM fine-tuning protocol. We illustrate the MFT procedure in Fig. 2. Please note that the FFT method used in this study is supervised fine-tuning (SFT) and does not consider other policy-based tuning methods like DPO (Direct Preference Optimization) (Rafailov et al., 2024) or PPO (Proximal Policy Optimization) (Schulman et al., 2017). Our MFT generalizes easily to them. We leave such explorations to our future work.

3 Experiments

We evaluate the proposed mask fine-tuning (MFT) on large language models (LLMs) across several backbones, domains, and full fine-tuning (FFT) settings. In addition, we provide ablation studies and visualizations to provide better intuition.

3.1 Experimental Setups

Pretrained Backbones. We use Transformer (Vaswani, 2017) based pretrained LLMs as backbone models, including LLaMA2-7B (Touvron et al., 2023), LLaMA3.1-8B (Grattafiori et al., 2024), Qwen3-1.7B (Yang et al., 2025), and Qwen3-14B (Yang et al., 2025).

Domains & Datasets. We involve three domains, each with its representative tasks. Specifically, we include GSM8K (Cobbe et al., 2021) and MetaMath (Yu et al., 2024) for math, HumanEval and HumanEval+ (Chen et al., 2021) for coding, and IF-Eval (Zhou et al., 2023) and Alpaca-Eval (Dubois et al., 2024) for instruction following domains, respectively. We use all three domains for typical FFT, our MFT, and evaluation.

Configurations. Our experiments use pretrained backbones, with FFT, MFT, and downstream evaluation in that order. For FFT, we follow two strategies: 1) domain-specific, 2) mix-up, where the former conducts FFT using datasets from individual domains, such as math, coding, or instruction following, and the latter uses a mix-up dataset from all these domains. For MFT, we follow a domain-specific approach, using datasets from a single domain. Further details of the dataset are provided in the Appendix. We mainly investigate local masking in this paper, as elaborated in a series of proof-of-concept studies in Sec. 3.2 and Fig. 3. We also provide an initial exploration of global masking strategy in Sec. 3.3 and Tab. 6. For evaluation, we comprehensively assess performance changes (see Tab. 1).

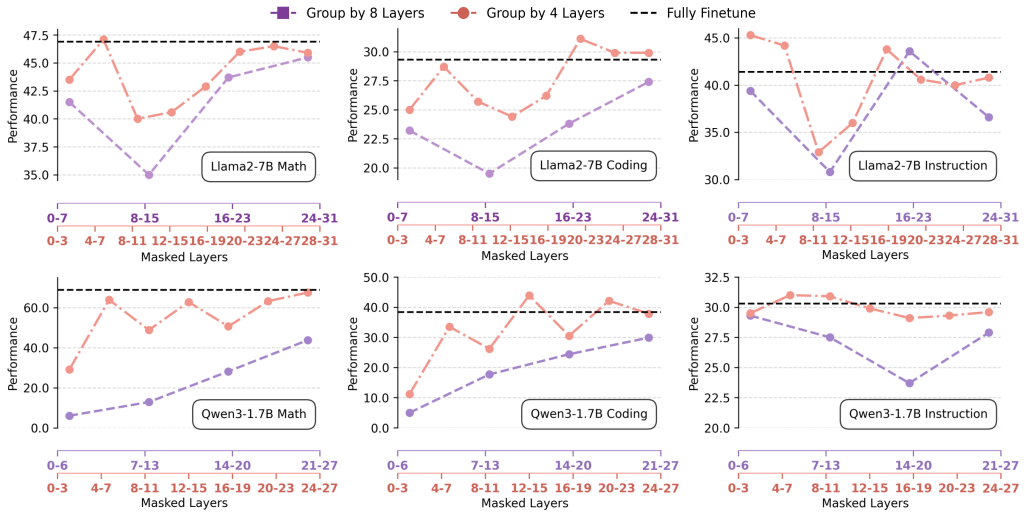


Figure 3: Visualization of the ablation study of the local MFT strategy. It uses LLaMA2-7B and Qwen3-1.7B backbones, covering math, coding, and instruction domains. In each figure, we conduct MFT ablations with a 10% masking ratio on a domain-specific FFT model (black dashed line). We swap the ablation between two local granularities, 8/7-layer (purple) and 4-layer (orange), moving from shallow to deep layers using a 10% fine-tuning set. This ablation is intended for quick intuition, with limited performance gains from MFT on a 10% subset. Based on this trend, we deploy MFT with complete fine-tuning sets, achieving more improvements.

Baselines. We use several baselines for comparison, including (1) a very basic pre-trained backbone; (2) FFT as a strong baseline, which is mostly seen as the most competitive approach for downstream fine-tuning. We perform a sufficient number of FFTs and report the best performance; (3) LoRA (Hu et al., 2022) fine-tuning as a strong alternative of FFT; (4) continued FFT, which always causes a performance drop due to overfitting; (5) continued LoRA plays a similar role for the LoRA baseline. We train all baselines for 4 epochs to ensure sufficient training and report the final performance for Continued FFT and Continued LoRA. (6) random mask and (7) L1 mask as two vanilla masking strategies. Among them, (2) FFT and (4) continued FFT are critical to our experiments, since our goal is to use MFT to further improve the best FFT, and continued FFT serves as a sanity check.

3.2 Main Results

As mentioned in Sec. 3.1, MFT primarily uses a local masking strategy. Therefore, we begin with a proof-of-concept study to assess whether MFT has potential for further improvements (Fig. 3). Such ablation also helps us identify the most effective components of the well-trained model for MFT. Accordingly, we extensively deploy MFT in this local fashion across different configurations and compare with strong fine-tuning baselines.

Proof-of-Concept Study of Local MFT. This series of ablations aims to assess whether the MFT can improve the model's capability after FFT. If so, we expect a comprehensive ablation to also detect the sensitivity of different model components to MFT. Specifically, we start from the LLaMA and Qwen backbones after domain-specific FFT and perform MFT on a small subset of each domain. Specifically, for each domain, we use 10% of its corresponding fine-tuning set. To swap the model, we group 4 or 8 consecutive layers to ablate from shallow to deep layers. For each group, we use MFT to learn a local mask while keeping the other layers fixed to assess model performance. For example, for layers 4-7, this is formally achieved by setting $\Theta = \{\theta_l\}$, where $l \in \{4, 5, 6, 7\}$. Within each group, we set the masking ratio to 10% to learn the mask, meaning we aim to learn and remove 10% of the parameters while keeping the remaining 90%.

Fig. 3 shows the detailed ablation results. We find that 1) compared with FFT (black dashed line), the MFT has the potential to further improve for each setting, even if they may have different proper

Table 1: Performance comparison of domain-specific settings. Each part has three blocks containing 1) common LoRA and its continued variant (green), 2) FFT serving as an upper bound and its continued variant (blue), and 3) our MFT (red) with the other two vanilla masking baselines. The pre-trained performance serves as a lower bound.

	Method	Math		Coding		Instruction Following	
		GSM8K	Math	HumanEval	HumanEval+	IF-Eval	Alpaca-Eval
LLaMA2-7B	**Pre-Trained Model**	15.2	2.5	25.8	22.4	34.3	0.5
	Best LoRA	$34.8_{\pm 0.34}$	$4.7_{\pm 0.26}$	$28.7_{\pm 0.56}$	$23.8_{\pm 0.48}$	$35.7_{\pm 0.68}$	$1.2_{\pm 0.16}$
	Continued LoRA	$33.6_{\pm 0.42}$	$4.5_{\pm 0.23}$	$28.2_{\pm 0.49}$	$23.2_{\pm 0.43}$	$32.5_{\pm 0.84}$	$1.1_{\pm 0.12}$
	Best FFT	$46.8_{\pm 0.22}$	$6.6_{\pm 0.21}$	$29.4_{\pm 0.36}$	$25.1_{\pm 0.34}$	$41.2_{\pm 0.62}$	$1.9_{\pm 0.18}$
	Continued FFT	$45.0_{\pm 0.28}$ ↓ 1.8	$5.5_{\pm 0.23}$ ↓ 1.1	$27.9_{\pm 0.42}$ ↓ 1.5	$23.6_{\pm 0.38}$ ↓ 1.5	$37.8_{\pm 0.68}$ ↓ 3.4	$2.0_{\pm 0.21}$ ↑ 0.1
	Random Mask w/ Best FFT	$38.4_{\pm 0.38}$	$4.9_{\pm 0.25}$	$28.0_{\pm 0.51}$	$24.0_{\pm 0.47}$	$39.4_{\pm 0.99}$	$1.4_{\pm 0.21}$
	L1 Mask w/ Best FFT	42.6	5.7	28.9	24.5	40.9	1.5
	MFT w/ Best FFT **(Ours)**	**47.3**$_{\pm 0.19}$ ↑ 0.5	**7.4**$_{\pm 0.24}$ ↑ 0.8	**31.8**$_{\pm 0.31}$ ↑ 2.4	**27.9**$_{\pm 0.25}$ ↑ 2.8	**44.1**$_{\pm 0.72}$ ↑ 2.9	**3.0**$_{\pm 0.22}$ ↑ 1.1
LLaMA3.1-8B	**Pre-Trained Model**	55.9	14.6	42.1	37.8	17.6	9.0
	Best LoRA	$70.0_{\pm 0.28}$	$22.7_{\pm 0.22}$	$50.0_{\pm 0.32}$	$45.1_{\pm 0.43}$	$49.1_{\pm 0.92}$	$10.1_{\pm 0.51}$
	Continued LoRA	$67.1_{\pm 0.35}$	$19.7_{\pm 0.19}$	$46.7_{\pm 0.46}$	$44.5_{\pm 0.38}$	$44.3_{\pm 0.88}$	$8.9_{\pm 0.29}$
	Best FFT	$76.8_{\pm 0.18}$	$24.2_{\pm 0.18}$	$51.0_{\pm 0.39}$	$44.9_{\pm 0.31}$	$59.6_{\pm 0.69}$	$11.9_{\pm 0.36}$
	Continued FFT	$74.4_{\pm 0.17}$ ↓ 2.4	$23.9_{\pm 0.17}$ ↓ 0.3	$48.6_{\pm 0.36}$ ↓ 2.4	$41.9_{\pm 0.30}$ ↓ 3.0	$57.5_{\pm 0.67}$ ↓ 2.1	$11.3_{\pm 0.34}$ ↓ 0.6
	Random Mask w/ Best FFT	$74.9_{\pm 0.37}$	$22.3_{\pm 0.23}$	$48.2_{\pm 0.48}$	$42.7_{\pm 0.41}$	$58.8_{\pm 0.96}$	$11.3_{\pm 0.48}$
	L1 Mask w/ Best FFT	75.1	22.7	49.1	44.1	59.8	11.5
	MFT w/ Best FFT **(Ours)**	**77.3**$_{\pm 0.14}$ ↑ 0.5	**24.8**$_{\pm 0.21}$ ↑ 0.6	**53.5**$_{\pm 0.41}$ ↑ 2.5	**46.8**$_{\pm 0.40}$ ↑ 1.9	**65.6**$_{\pm 0.70}$ ↑ 6.0	**13.7**$_{\pm 0.32}$ ↑ 1.8
Qwen3-1.7B	**Pre-Trained Model**	58.8	30.2	38.4	35.4	30.3	9.4
	Best LoRA	$63.5_{\pm 0.32}$	$33.8_{\pm 0.25}$	$44.5_{\pm 0.38}$	$39.6_{\pm 0.49}$	$48.2_{\pm 0.75}$	$10.6_{\pm 0.58}$
	Continued LoRA	$61.8_{\pm 0.40}$	$32.4_{\pm 0.22}$	$42.1_{\pm 0.52}$	$38.0_{\pm 0.43}$	$44.7_{\pm 0.69}$	$9.8_{\pm 0.33}$
	Best FFT	$68.7_{\pm 0.21}$	$36.5_{\pm 0.20}$	$50.5_{\pm 0.44}$	$44.1_{\pm 0.35}$	$55.6_{\pm 0.78}$	$11.4_{\pm 0.41}$
	Continued FFT	$66.5_{\pm 0.19}$ ↓ 2.1	$35.2_{\pm 0.19}$ ↓ 1.3	$47.4_{\pm 0.41}$ ↓ 2.6	$42.0_{\pm 0.34}$ ↓ 2.1	$52.1_{\pm 0.76}$ ↓ 3.5	$10.6_{\pm 0.39}$ ↓ 0.8
	Random Mask w/ Best FFT	$61.5_{\pm 0.42}$	$34.1_{\pm 0.27}$	$47.6_{\pm 0.54}$	$42.0_{\pm 0.46}$	$53.0_{\pm 0.80}$	$10.5_{\pm 0.55}$
	L1 Mask w/ Best FFT	66.2	35.3	48.5	43.0	54.0	10.8
	MFT w/ Best FFT **(Ours)**	**70.8**$_{\pm 0.17}$ ↑ 2.1	**38.3**$_{\pm 0.24}$ ↑ 1.8	**52.0**$_{\pm 0.46}$ ↑ 1.5	**46.3**$_{\pm 0.45}$ ↑ 2.2	**59.7**$_{\pm 0.80}$ ↑ 4.1	**12.4**$_{\pm 0.36}$ ↑ 1.0
Qwen3-14B	**Pre-Trained Model**	81.0	45.1	56.1	52.4	45.2	29.5
	Best LoRA	$83.5_{\pm 0.24}$	$46.7_{\pm 0.19}$	$66.5_{\pm 0.29}$	$55.8_{\pm 0.40}$	$57.5_{\pm 0.82}$	$32.0_{\pm 0.47}$
	Continued LoRA	$82.2_{\pm 0.30}$	$45.5_{\pm 0.17}$	$63.3_{\pm 0.42}$	$53.7_{\pm 0.36}$	$53.1_{\pm 0.79}$	$30.4_{\pm 0.27}$
	Best FFT	$85.9_{\pm 0.16}$	$48.9_{\pm 0.16}$	$75.0_{\pm 0.35}$	$61.6_{\pm 0.28}$	$70.3_{\pm 0.62}$	$39.6_{\pm 0.33}$
	Continued FFT	$83.6_{\pm 0.15}$ ↓ 2.2	$47.6_{\pm 0.15}$ ↓ 1.3	$70.8_{\pm 0.32}$ ↓ 3.6	$58.0_{\pm 0.27}$ ↓ 3.6	$68.5_{\pm 0.60}$ ↓ 1.8	$38.8_{\pm 0.31}$ ↓ 0.8
	Random Mask w/ Best FFT	$83.8_{\pm 0.33}$	$47.5_{\pm 0.20}$	$71.5_{\pm 0.43}$	$59.0_{\pm 0.37}$	$64.5_{\pm 0.86}$	$34.8_{\pm 0.43}$
	L1 Mask w/ Best FFT	84.2	48.0	71.9	59.5	65.0	35.2
	MFT w/ Best FFT **(Ours)**	**86.9**$_{\pm 0.12}$ ↑ 1.0	**50.3**$_{\pm 0.18}$ ↑ 1.4	**76.0**$_{\pm 0.37}$ ↑ 1.0	**63.7**$_{\pm 0.36}$ ↑ 2.1	**76.8**$_{\pm 0.63}$ ↑ 6.5	**42.1**$_{\pm 0.29}$ ↑ 2.5

layer-wise groups for MFT; 2) A smaller layer-wise group (4-layer) is better for MFT compared with a larger one (8-layer). However, different model families have different trends. For the LLaMA family, shallow (0-7) and mid-to-rear (20-27) partitions often respond more positively to MFT than other partitions, whereas the Qwen family shows a more consistent trend across groups. Such observation confirms the MFT's potential to further improve models after FFT and indicates promising partitions for mask learning. Concretely, for LLaMA2-7B, we use 4-7, 20-23, and 0-3 for math, coding, and instruction domains, respectively, for MFT. For LLaMA3.1-8B, we use 24-27, 20-23, and 24-27 for the three domains, respectively. For Qwen3-1.7B, we use 24-27, 12-15, and 4-7 for the three domains, respectively. For Qwen3-14B, we use 37-40, 32-35, and 7-10 for the three domains, respectively. For all full-dataset MFT results reported in the rest of this paper, we follow the proof-of-concept study and use a constant 10% sparse ratio.

Notably, 1) Since this ablation is only for quick intuition and a sanity check of MFT effectiveness, we only use a 10% dataset subset here. According to such insights, MFT is trained using complete fine-tuning datasets for more performance gain; 2) Since larger layer groups (e.g., 16) achieve a significant performance drop, we omit their results for a better illustration and leave them in the Appendix.

Comparisons. Tab. 1 presents our main comparisons using the LLaMA and Qwen backbones. We start with a pre-trained model, then apply two FFT settings, using the fine-tuning dataset for both specific and mixed domains for each backbone. For each setting, we perform sufficient FFTs and report the best and continued versions (in the blue block). We do the same for LoRA (green block). We also include random and L1 masks as two vanilla masking baselines (white block). MFT (red block) starts from the best-performing FFT and further improves performance. All three masking methods are based on the best FFT as a strong starting point. Among them, two critical baselines are the best FFT and continued FFT: the former serves as a competitive baseline, and the latter as a sanity check for potential performance drops due to overfitting, thereby demonstrating MFT's effectiveness. We train our MFT only using domain-specific datasets for both FFT scenarios (domain-specific and

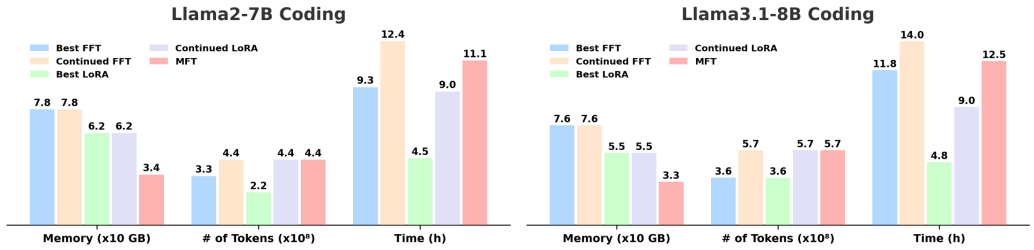


Figure 4: Training cost comparisons of GPU memory, token usage, and training time. We use the coding domain as an example to compare MFT with other FFT and LoRA baselines.

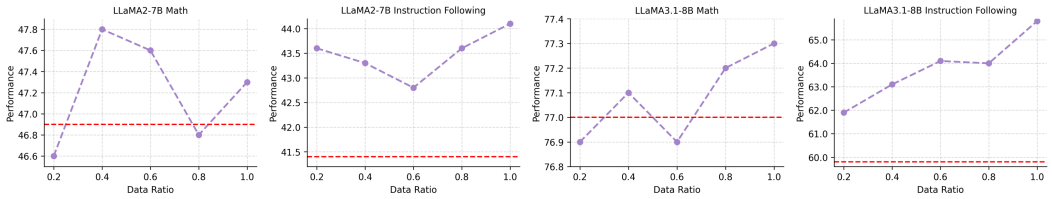


Figure 5: Data ratio ablation. We use math and instruction domains on LLaMA2 and LLaMA3.1. Compared with Best FFT (red dashed line), we observe that MFT (purple) consistently improves performance on the full dataset, but may still yield a promising performance gain with less data.

mixed-up). Additionally, we compare the training cost in Fig. 4 across GPU memory usage, token count, and overall training time on A100 GPUs. We use the coding domain as an example. Additional comparisons are provided in the Appendix. Compared with the most important baseline, Best FFT, MFT involves very limited overhead for both token usage and training time. Compared with the Continued FFT, which incurs at least as much, if not more, training overhead, the Best FFT and MFT incur lower training costs while achieving performance improvements. Since MFT fixes all model parameters and performs only local mask learning, it always has an advantage in memory usage.

Accordingly, we conclude: 1) Since the Continued FFT causes a performance drop and the Best FFT outperforms other comparison methods, it is reasonable to set it as our strong baseline. 2) MFT achieves promising and consistent improvements across different backbones, domains, and FFT configurations. Accordingly, it upgrades the current routine to a new LLM fine-tuning protocol by learning and applying a mask on a well-trained model. Moreover, the other two masking baselines generally degrade performance, particularly compared with the consistent improvements achieved by MFT, indicating that MFT is not trivial and warrants further exploration.

3.3 Analyses

Theoretical Analysis. We follow PAC-Bayes theory for a generalization upper bound (McAllester, 1998), which has also been widely explored for neural networks (Arora et al., 2018), to provide information-theoretic intuition. We focus on the performance gain from the model after the Best FFT to the model after MFT. Since our MFT uses the same training objective and dataset as FFT, the analysis is a fair comparison and yields a purely model-generalization upper bound.

Given n as the size of the training set, δ as the degree of confidence, for any hypothesis h, which represents the model here, we have the PAC-Bayes upper bound given by

$$L(h) \leq L_S(h) + \Phi(C(h)), \quad \Phi(u) = \sqrt{\frac{u + \ln \frac{1}{\delta}}{2(n-1)}}. \tag{6}$$

$L(h)$ and $L_S(h)$ represent training and test loss, respectively. $\Phi(C(h))$ is an additional complexity term to describe the model code length, where $C(\cdot)$ represents the encoding process. Then, we have

$$\begin{aligned} &L(h_{\text{FFT}}) \leq L_S(h_{\text{FFT}}) + \Phi(C_{\text{FFT}}) \\ &\Rightarrow U(h_{\text{FFT}}) = L_S(h_{\text{FFT}}) + \Phi(C_{\text{FFT}}), \end{aligned} \tag{7}$$

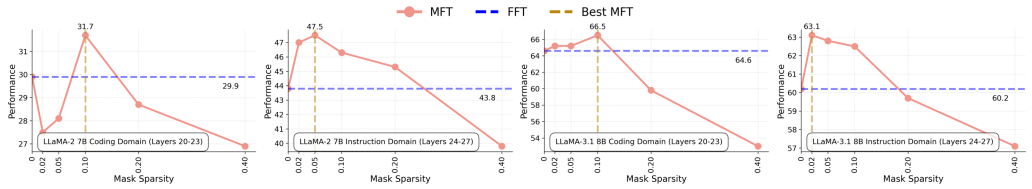


Figure 6: Masking ratio ablation. We use coding and instruction domains on LLaMA2 and LLaMA3.1. We observe that the original 10% ratio works well for coding but not for instruction, indicating that the masking ratio matters and that there is more MFT potential.

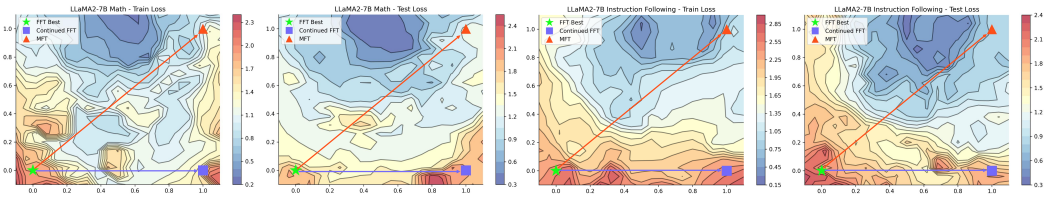


Figure 7: The loss landscape visualization on math and instruction-following domains using LLaMA2-7B. Such visualization indicates that the proposed MFT further refines the Best FFT model, improving optimization and generalization.

$$\begin{aligned} &L(h_{\mathrm{MFT}}) \leq L_S(h_{\mathrm{MFT}}) + \Phi(C_{\mathrm{MFT}}) \\ &\Rightarrow U(h_{\mathrm{MFT}}) = L_S(h_{\mathrm{MFT}}) + \Phi(C_{\mathrm{MFT}}), \end{aligned} \tag{8}$$

where $U(\cdot)$ means the loss upper bound of h. If we take the difference on both sides, we have

$$\begin{aligned} &U(h_{\mathrm{MFT}}) - U(h_{\mathrm{FFT}}) \\ &\quad = \big[L_S(h_{\mathrm{MFT}}) - L_S(h_{\mathrm{FFT}})\big] \\ &\quad + \big[\Phi(C_{\mathrm{MFT}}) - \Phi(C_{\mathrm{FFT}})\big], \end{aligned} \tag{9}$$

revised as

$$U(h_{\mathrm{MFT}}) - U(h_{\mathrm{FFT}}) = \Delta_{\mathrm{train}} + \Delta_{\mathrm{complexity}}. \tag{10}$$

We numerically show $\Delta_{\mathrm{train}} + \Delta_{\mathrm{complexity}} < 0$. Details are provided in the Appendix.

Data Ratio Ablation. We ablate data ratios for MFT. Fig. 5 shows performance trends in the math and instruction-following domains on LLaMA2-7B and LLaMA3.1-8B (others are in the supplementary). We find that MFT (purple), when trained on the full dataset, consistently improves the model compared with the Best FFT (red dashed line). However, less data may also perform competitively and more efficiently.

Masking Ratio Ablation. We ablate the masking ratio for local MFT. We use a 10% masking ratio across all settings (Tab. 1), which may not be optimal. Here, we define a series of masking ratios for the ablation study. Specifically, we use all three domains on two backbones and set the masking ratio from 0.02 to 0.4. Results for coding and instruction are shown in Fig. 6, and the math results are in the supplementary. We find that our original 10% ratio works well for certain coding cases. However, it is not optimal for the instruction domain, where a smaller ratio works better than 10%. This suggests that the MFT masking ratio affects final performance.

Mask Fine-Tuning Loss Landscape. Loss surface visualization is commonly used to analyze the model optimization process and provide insights into the model's status from a loss perspective. Thus, we visualize the MFT loss surface to illustrate its training dynamics. Fig. 7 shows the surface of LLaMA2-7B on math and instruction-following domains. The Best FFT, continued FFT, and MFT are represented by the green star, the purple rectangle, and the red triangle, respectively. We find that MFT always optimizes the model to a better status than the Best FFT start point and its continued version in both training and test scenarios. We supplement more theoretical analyses in the Appendix.

4 Related Work

This work explores optimizing a well-trained large language model (LLM) by learning a binary mask and proposes mask fine-tuning (MFT), an upgrade to the current LLM fine-tuning protocol. We summarize the relevant literature on our work from (1) LLM pretraining and fine-tuning and (2) sparse network perspectives, emphasizing their correlations and differences.

4.1 LLMs Pre-training & Fine-tuning

LLM pre-training leads to significant performance gains in general language model capacity for both understanding (Devlin et al., 2019; Raffel et al., 2020; Liu, 2019; He et al., 2021) and generation (Radford et al., 2019; Brown et al., 2020; Achiam et al., 2023; Grattafiori et al., 2024; Touvron et al., 2023). Massive language corpora with large-scale pretraining enable them to serve as strong language backbones with basic common knowledge. Following pretraining, downstream fine-tuning on small but high-quality target datasets further enhances pre-trained models with the specific capacities of target domains (Hendrycks et al., 2021; Yu et al., 2024; Cobbe et al., 2021; Chen et al., 2021; Austin et al., 2021; Zhou et al., 2023; Wang et al., 2023; Nguyen, 2023). Different fine-tuning optimization strategies deliver various LLM properties, such as supervised fine-tuning (SFT), which follows a pre-training objective to involve extra domain knowledge into LLM (Taori et al., 2023; Chiang et al., 2023; Wang et al., 2022b), and policy-based fine-tuning methods better align the human preference with LLM (Schulman et al., 2017; Rafailov et al., 2024; Ivison et al., 2024; Ouyang et al., 2022; Dai et al., 2024). We propose Mask Fine-Tuning (MFT) as a post-fine-tuning strategy to further refine the fine-tuned LLM, especially for SFT in this study. This strategy can be flexibly deployed after any fine-tuned model to propose a new protocol for improving LLMs.

4.2 Sparse Networks

Model sparsity in neural networks is close to neural network pruning (Liu et al., 2019; Molchanov et al., 2017; Wang et al., 2022a) for model compression (Han et al., 2016b; Iandola, 2016) and acceleration (Han et al., 2016a; Wang et al., 2021; Wang & Fu, 2023; Ma et al., 2023; Fang et al., 2023). Our work overlaps with network pruning, as both aim to identify a model partition with a given sparsity. However, pruning aims to remove a large portion of the model for efficiency, whereas ours focuses on removing parameters that are irrelevant or even harmful to the model's capacity, potentially requiring only a small subset of the model. More importantly, pruning methods inevitably entail a performance loss in the trained model, whereas ours focuses on further enhancing it. In addition, the sparsity concept is also involved in other machine learning fields such as network optimization (Srivastava et al., 2014; Srinivas et al., 2017; Baldi & Sadowski, 2013), model architecture design (Shazeer et al., 2017; Zoph, 2017; Elsken et al., 2019), lottery ticket hypothesis (Frankle & Carbin, 2019; Zhou et al., 2019; You et al., 2020; Bai et al., 2022b), neural network mechanism (Wortsman et al., 2020; Ramanujan et al., 2020; Bai et al., 2022a), etc. In our study, we introduce sparsity to investigate whether all parameters are beneficial for a well-trained LLM and to improve the model by learning and applying a mask, proposing a new approach to LLM fine-tuning.

5 Conclusion

In this paper, we challenge the necessity of structural integrity in large language models (LLMs) by examining whether a well-trained LLM can be further improved by removing certain model parameters. We propose mask fine-tuning (MFT) to explore this using a model with sufficient full fine-tuning (FFT), which consistently yields further improvements, generally enhances LLM capabilities, and naturally upgrades the LLM fine-tuning pipeline. Comprehensive experiments with detailed ablations support our conclusion, drawing on exploration intuitions, covering different pre-trained LLM backbones, downstream domains, and fine-tuning configurations. Meanwhile, MFT extends model sparsity beyond compression, enabling it in a broader range of scenarios and inspiring further work.

AI USE STATEMENT

We use AI to correct grammar errors and help plot figures.

ETHICS STATEMENT

In our paper, we strictly adhere to the ICLR ethical research standards and applicable laws. To the best of our knowledge, our work abides by the General Ethical Principles.

REPRODUCIBILITY STATEMENT

We adhere to ICLR reproducibility standards and ensure the reproducibility of our work. All datasets we employed are publicly available. We will provide the code to reviewers and area chairs via an anonymous link, if needed, during the discussion period.

REFERENCES

Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.

Sanjeev Arora, Rong Ge, Behnam Neyshabur, and Yi Zhang. Stronger generalization bounds for deep nets via a compression approach. In *International conference on machine learning*, 2018.

Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.

Yue Bai, Huan Wang, Xu Ma, Yitian Zhang, Zhiqiang Tao, and Yun Fu. Parameter-efficient masking networks. In *NeurIPS*, 2022a.

Yue Bai, Huan Wang, Zhiqiang Tao, Kunpeng Li, and Yun Fu. Dual lottery ticket hypothesis. In *ICLR*, 2022b.

Pierre Baldi and Peter J Sadowski. Understanding dropout. In *NeurIPS*, 2013.

Yoshua Bengio, Nicholas Léonard, and Aaron Courville. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*, 2013.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *NeurIPS*, 2020.

Sahil Chaudhary. Code alpaca: An instruction-following llama model for code generation, 2023.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde De Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021.

Wei-Lin Chiang, Zhuohan Li, Zi Lin, Ying Sheng, Zhanghao Wu, Hao Zhang, Lianmin Zheng, Siyuan Zhuang, Yonghao Zhuang, Joseph E Gonzalez, et al. Vicuna: An open-source chatbot impressing gpt-4 with 90%* chatgpt quality. *https://vicuna.lmsys.org*, 2023.

Karl Cobbe, Vineet Kosaraju, Mohammad Bavarian, Mark Chen, Heewoo Jun, Lukasz Kaiser, Matthias Plappert, Jerry Tworek, Jacob Hilton, Reiichiro Nakano, Christopher Hesse, and John Schulman. Training verifiers to solve math word problems. *arXiv preprint arXiv:2110.14168*, 2021.

Josef Dai, Xuehai Pan, Ruiyang Sun, Jiaming Ji, Xinbo Xu, Mickel Liu, Yizhou Wang, and Yaodong Yang. Safe rlhf: Safe reinforcement learning from human feedback. In *ICLR*, 2024.

Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the Nations of the Americas Chapter of the Association for Computational Linguistics: Human Language Technologies (Volume 1: Long Papers)*, 2019.

Yann Dubois, Balázs Galambosi, Percy Liang, and Tatsunori B Hashimoto. Length-controlled alpacaeval: A simple way to debias automatic evaluators. In *COLM*, 2024.

Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. In *JMLR*, 2019.

Gongfan Fang, Xinyin Ma, and Xinchao Wang. Structural pruning for diffusion models. In *NeurIPS*, 2023.

Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *ICLR*, 2019.

Simon Frieder, Luca Pinchetti, Ryan-Rhys Griffiths, Tommaso Salvatori, Thomas Lukasiewicz, Philipp Petersen, and Julius Berner. Mathematical capabilities of chatgpt. In *NeurIPS*, 2024.

Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783*, 2024.

Daya Guo, Qihao Zhu, Dejian Yang, Zhenda Xie, Kai Dong, Wentao Zhang, Guanting Chen, Xiao Bi, Y. Wu, Y. K. Li, Fuli Luo, Yingfei Xiong, and Wenfeng Liang. Deepseek-coder: When the large language model meets programming – the rise of code intelligence, 2024.

Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A Horowitz, and William J Dally. Eie: Efficient inference engine on compressed deep neural network. In *ACM SIGARCH Computer Architecture News*, 2016a.

Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *ICLR*, 2016b.

Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. Deberta: Decoding-enhanced bert with disentangled attention. In *ICLR*, 2021.

Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset. In *NeurIPS*, 2021.

Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. Parameter-efficient transfer learning for nlp. In *ICML*, 2019.

Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022.

Forrest N Iandola. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and¡ 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016.

Hamish Ivison, Yizhong Wang, Jiacheng Liu, Zeqiu Wu, Valentina Pyatkin, Nathan Lambert, Noah A Smith, Yejin Choi, and Hannaneh Hajishirzi. Unpacking dpo and ppo: Disentangling best practices for learning from preference feedback. In *NeurIPS*, 2024.

Andreas Köpf, Yannic Kilcher, Dimitri von Rütte, Sotiris Anagnostidis, Zhi-Rui Tam, Keith Stevens, Abdullah Barhoum, Nguyen Minh Duc, Oliver Stanley, Richárd Nagyfi, Shahul ES, Sameer Suri, David Glushkov, Arnav Dantuluri, Andrew Maguire, Christoph Schuhmann, Huu Nguyen, and Alexander Mattick. Openassistant conversations – democratizing large language model alignment. In *NeurIPS*, 2023.

Nathan Lambert, Jacob Morrison, Valentina Pyatkin, Shengyi Huang, Hamish Ivison, Faeze Brahman, Lester James V. Miranda, Alisa Liu, Nouha Dziri, Shane Lyu, Yuling Gu, Saumya Malik, Victoria Graf, Jena D. Hwang, Jiangjiang Yang, Ronan Le Bras, Oyvind Tafjord, Chris Wilhelm, Luca Soldaini, Noah A. Smith, Yizhong Wang, Pradeep Dasigi, and Hannaneh Hajishirzi. Tulu 3: Pushing frontiers in open language model post-training. In *COLM*, 2025.

Aitor Lewkowycz, Anders Andreassen, David Dohan, Ethan Dyer, Henryk Michalewski, Vinay Ramasesh, Ambrose Slone, Cem Anil, Imanol Schlag, Theo Gutman-Solo, Yuhuai Wu, Behnam Neyshabur, Guy Gur-Ari, and Vedant Misra. Solving quantitative reasoning problems with language models. In *NeurIPS*, 2022.

Jia LI, Edward Beeching, Lewis Tunstall, Ben Lipkin, Roman Soletskyi, Shengyi Costa Huang, Kashif Rasul, Longhui Yu, Albert Jiang, Ziju Shen, Zihan Qin, Bin Dong, Li Zhou, Yann Fleureau, Guillaume Lample, and Stanislas Polu. Numinamath tir, 2024.

Xuechen Li, Tianyi Zhang, Yann Dubois, Rohan Taori, Ishaan Gulrajani, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Alpacaeval: An automatic evaluator of instruction-following models, 2023.

Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. Is your code generated by chatGPT really correct? rigorous evaluation of large language models for code generation. In *NeurIPS*, 2023.

Yinhan Liu. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.

Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. In *ICLR*, 2019.

Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. Wizardcoder: Empowering code large language models with evol-instruct. In *ICLR*, 2024.

Xinyin Ma, Gongfan Fang, and Xinchao Wang. Llm-pruner: On the structural pruning of large language models. In *NeurIPS*, 2023.

David A McAllester. Some pac-bayesian theorems. In *Proceedings of the eleventh annual conference on Computational learning theory*, 1998.

Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. Pruning convolutional neural networks for resource efficient inference. In *ICLR*, 2017.

Subhabrata Mukherjee, Arindam Mitra, Ganesh Jawahar, Sahaj Agarwal, Hamid Palangi, and Ahmed Awadallah. Orca: Progressive learning from complex explanation traces of gpt-4. *arXiv preprint arXiv:2306.02707*, 2023.

Ha-Thanh Nguyen. A brief report on lawgpt 1.0: A virtual legal assistant based on gpt-3. *arXiv preprint arXiv:2302.05729*, 2023.

Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. Training language models to follow instructions with human feedback. In *NeurIPS*, 2022.

Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI Blog*, 2019.

Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. In *NeurIPS*, 2024.

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. In *JMLR*, 2020.

Vivek Ramanujan, Mitchell Wortsman, Aniruddha Kembhavi, Ali Farhadi, and Mohammad Rastegari. What's hidden in a randomly weighted neural network? In *CVPR*, 2020.

John Schulman, Filip Wolski, Prafulla Dhariwal, Alec Radford, and Oleg Klimov. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*, 2017.

Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarz, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *ICLR*, 2017.

Suraj Srinivas, Akshayvarun Subramanya, and R Venkatesh Babu. Training sparse neural networks. In *CVPR Workshops*, 2017.

Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: a simple way to prevent neural networks from overfitting. In *JMLR*, 2014.

Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B Hashimoto. Alpaca: A strong, replicable instruction-following model. *Stanford Center for Research on Foundation Models*, 2023.

Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.

A Vaswani. Attention is all you need. In *NeurIPS*, 2017.

Guangyu Wang, Guoxing Yang, Zongxin Du, Longjun Fan, and Xiaohu Li. Clinicalgpt: large language models finetuned with diverse medical data and comprehensive evaluation. *arXiv preprint arXiv:2306.09968*, 2023.

Huan Wang and Yun Fu. Trainability preserving neural pruning. In *ICLR*, 2023.

Huan Wang, Can Qin, Yulun Zhang, and Yun Fu. Neural pruning via growing regularization. In *ICLR*, 2021.

Huan Wang, Can Qin, Yue Bai, Yulun Zhang, and Yun Fu. Recent advances on neural network pruning at initialization. In *IJCAI*, 2022a.

Yizhong Wang, Yeganeh Kordi, Swaroop Mishra, Alisa Liu, Noah A Smith, Daniel Khashabi, and Hannaneh Hajishirzi. Self-instruct: Aligning language models with self-generated instructions. *arXiv preprint arXiv:2212.10560*, 2022b.

Mitchell Wortsman, Vivek Ramanujan, Rosanne Liu, Aniruddha Kembhavi, Mohammad Rastegari, Jason Yosinski, and Ali Farhadi. Supermasks in superposition. In *NeurIPS*, 2020.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, Chujie Zheng, Dayiheng Liu, Fan Zhou, Fei Huang, Feng Hu, Hao Ge, Haoran Wei, Huan Lin, Jialong Tang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Yang, Jiaxi Yang, Jing Zhou, Jingren Zhou, Junyang Lin, Kai Dang, Keqin Bao, Kexin Yang, Le Yu, Lianghao Deng, Mei Li, Mingfeng Xue, Mingze Li, Pei Zhang, Peng Wang, Qin Zhu, Rui Men, Ruize Gao, Shixuan Liu, Shuang Luo, Tianhao Li, Tianyi Tang, Wenbiao Yin, Xingzhang Ren, Xinyu Wang, Xinyu Zhang, Xuancheng Ren, Yang Fan, Yang Su, Yichang Zhang, Yinger Zhang, Yu Wan, Yuqiong Liu, Zekun Wang, Zeyu Cui, Zhenru Zhang, Zhipeng Zhou, and Zihan Qiu. Qwen3 technical report, 2025.

Haoran You, Chaojian Li, Pengfei Xu, Yonggan Fu, Yue Wang, Xiaohan Chen, Richard G Baraniuk, Zhangyang Wang, and Yingyan Lin. Drawing early-bird tickets: Towards more efficient training of deep networks. In *ICLR*, 2020.

Longhui Yu, Weisen Jiang, Han Shi, Jincheng Yu, Zhengying Liu, Yu Zhang, James T Kwok, Zhenguo Li, Adrian Weller, and Weiyang Liu. Metamath: Bootstrap your own mathematical questions for large language models. *ICLR*, 2024.

Hattie Zhou, Janice Lan, Rosanne Liu, and Jason Yosinski. Deconstructing lottery tickets: Zeros, signs, and the supermask. In *NeurIPS*, 2019.

Jeffrey Zhou, Tianjian Lu, Swaroop Mishra, Siddhartha Brahma, Sujoy Basu, Yi Luan, Denny Zhou, and Le Hou. Instruction-following evaluation for large language models. *arXiv preprint arXiv:2311.07911*, 2023.

B Zoph. Neural architecture search with reinforcement learning. In *ICLR*, 2017.

APPENDIX

Table 2: The information on the training and test datasets used in our experiments.

Domain	Train / Test	Dataset Name	# of Samples
Math	Train	Tulu 3 Persona MATH (Lambert et al., 2025)	149,960
		Tulu 3 Persona GSM (Lambert et al., 2025)	49,980
		Tulu 3 Persona Algebra (Lambert et al., 2025)	20,000
		MetaMathQA (Yu et al., 2024)	395,000
		NuminaMath-TIR (LI et al., 2024)	64,312
	Test	GSM8K (Cobbe et al., 2021)	1,320
		MATH (Lewkowycz et al., 2022)	5,003
Coding	Train	Evol CodeAlpaca (Luo et al., 2024)	107,276
		Code-Alpaca (Chaudhary, 2023)	20,022
		Tulu 3 Persona Python (Lambert et al., 2025)	34,999
	Test	HumanEval (Chen et al., 2021)	164
		HumanEval+ (Liu et al., 2023)	164
Instruction Following	Train	OpenAssistant Guanaco (Köpf et al., 2023)	7,132
		Tulu 3 Persona IF (Lambert et al., 2025)	29,980
		Open-Orca (Mukherjee et al., 2023)	30,000
	Test	IFEval (Zhou et al., 2023)	541
		AlpacaEval 2.0 (Li et al., 2023)	805

A IMPLEMENTATION DETAILS

We conduct our experiments on 8*A100 GPUs. To ensure that full fine-tuning (FFT) is sufficient, we set the maximum number of FFT epochs to 4, as we observe overfitting and a drop in performance. We fine-tune MFT for 2 additional epochs, starting from the best FFT checkpoint (a total of 4 epochs). We set the batch size to 8, the accumulation steps to 16, and the weight decay to 0.0 for all settings. We use learning rates of 2e-5 and 1e-5 for LLaMA and Qwen, with a linear schedule. The warm-up learning rate starts at 0 and increases linearly, with a warm-up ratio of 0.03 of the total training. We summarize the training and test datasets used in our experiments in Tab. 2.

Table 3: Cross-domain evaluations on LLaMA2-7B under domain-specific FFT setting.

Method	Math		Coding		Instruction Following	
	GSM8K	Math	HumanEval	HumanEval+	IF-Eval	Alpaca-Eval
Pre-trained Model	15.2	2.5	25.8	22.4	34.3	0.5
Best FFT [Math]	46.9	6.7	17.7	15.9	29.3	1.5
Best FFT [Coding]	14.6	2.8	29.3	25.0	8.3	1.4
Best FFT [IF]	25.0	2.4	16.5	13.4	41.4	1.7
Continued FFT [Math]	45.2	5.7	19.5	17.1	33.2	1.5
Continued FFT [Coding]	11.1	2.5	28.0	23.8	13.1	2.7
Continued FFT [IF]	23.1	2.1	13.1	12.8	37.4	2.0
Random Mask w/ Best FFT [Math]	38.4	4.9	14.8	13.1	25.5	1.0
Random Mask w/ Best FFT [Coding]	12.4	2.5	28.0	24.0	7.5	1.1
Random Mask w/ Best FFT [IF]	23.0	2.3	14.4	12.2	39.4	1.5
L1 Mask w/ Best FFT [Math]	42.6	5.7	16.4	14.6	28.3	1.1
L1 Mask w/ Best FFT [Coding]	12.8	2.6	28.9	24.5	7.8	1.0
L1 Mask w/ Best FFT [IF]	23.9	2.4	14.9	12.7	40.9	1.5
[Ours] Best FFT + Mask [Math]	**47.3**	**7.6**	16.5	14.6	30.1	0.5
[Ours] Best FFT + Mask [Coding]	13.8	2.8	**31.7**	**28.0**	8.5	2.1
[Ours] Best FFT + Mask [IF]	23.5	2.8	15.2	11.6	**44.1**	**2.9**

Table 4: Cross-domain evaluations on LLaMA3.1-8B under domain-specific FFT setting.

Method	Math		Coding		Instruction Following	
	GSM8K	Math	HumanEval	HumanEval+	IF-Eval	Alpaca-Eval
Pre-trained Model	55.9	14.6	42.1	37.8	17.6	9.0
Best FFT [Math]	77.0	24.3	35.4	29.9	31.5	0.6
Best FFT [Coding]	60.7	13.8	51.2	45.1	5.3	0.9
Best FFT [IF]	63.0	10.1	50.0	44.5	59.8	12.0
Continued FFT [Math]	74.6	24.0	33.5	31.1	32.2	0.7
Continued FFT [Coding]	59.6	13.1	48.8	42.1	4.8	1.4
Continued FFT [IF]	63.6	8.7	45.7	43.5	57.2	11.4
Random Mask w/ Best FFT [Math]	74.9	22.3	34.3	28.5	30.0	0.4
Random Mask w/ Best FFT [Coding]	56.2	12.7	48.2	42.7	4.7	0.7
Random Mask w/ Best FFT [IF]	61.3	9.2	45.3	41.7	58.8	11.5
L1 Mask w/ Best FFT [Math]	75.1	22.7	34.4	30.6	31.1	0.6
L1 Mask w/ Best FFT [Coding]	57.3	12.9	49.1	44.1	4.8	0.8
L1 Mask w/ Best FFT [IF]	62.3	9.3	46.1	42.4	59.8	11.3
[Ours] Best FFT + Mask [Math]	**77.3**	**24.9**	36.0	32.2	31.1	0.5
[Ours] Best FFT + Mask [Coding]	58.6	13.6	**53.7**	**47.0**	4.9	0.0
[Ours] Best FFT + Mask [Instruction]	63.2	9.6	44.5	40.9	**65.8**	**13.8**

B More Evaluations

Tab. 3 and Tab. 4 extend the performance tables in our main draft by adding cross-domain evaluations for our MFT, FFT, and other comparison methods, under domain-specific FFT settings. The best performance for each setting is highlighted in bold. These cross-domain results are included to better understand MFT's behavior beyond the target domain. MFT is designed as a domain-specific post-FFT refinement method: given a well-fine-tuned model and a target domain, it learns a binary mask to identify the subnetwork that is most beneficial for that domain. Therefore, the learned mask may remove weights that are redundant or noisy for the target domain, but those weights may still be useful for other domains. In this sense, cross-domain performance degradation is expected and consistent with MFT's role as a domain-specialized refinement mechanism.

C More Ablations and Analyses

Multi-domain Settings. Tab. 5 reports additional results under mixed-domain FFT settings. Here, we first obtain the starting checkpoint from a well-tuned mixed-domain FFT model, then apply each post-FFT refinement method, including continued FFT/LoRA and MFT, in the same domain-specific setting for a fair comparison. This setting reflects a practical scenario: given a strong mixed-domain fine-tuned checkpoint, can MFT still serve as an effective domain-specific post-FFT refinement method? The results show that MFT consistently improves over Best FFT across all six benchmarks on both LLaMA2-7B and LLaMA3.1-8B. In contrast, continued FFT usually degrades performance, and random or L1 masking fails to consistently match Best FFT. These results suggest that MFT remains effective beyond domain-specific FFT and can refine a mixed-domain model through learned task-aware masks rather than additional optimization steps or heuristic sparsity.

Global Masking Fine-Tuning. Unlike the local MFT above, we further explore MFT in a global setting, where the mask is learned across all eligible trainable layers without manually selecting local layer groups. Our initial exploration of language-only backbones in Tab. 6 shows mixed results: global MFT improves LLaMA2-7B in the math domain, especially on GSM8K, but degrades performance on LLaMA3.1-8B. We note that this weaker performance may be due to the sparsity configuration used in the initial study, where we used around 10% sparsity to match the local mask-learning setting.

Table 6: Initial exploration of global MFT on the math domain using LLaMA2-7B and LLaMA3.1-8B backbones. We observe better performance on LLaMA2-7B, especially on GSM8K, but lower performance on LLaMA3.1-8B.

Method	LLaMA2-7B		LLaMA3.1-8B	
	GSM8K	Math	GSM8K	Math
Pre-Trained Model	15.2	2.5	55.9	14.6
Best FFT	46.9	6.7	77.0	24.3
Continued FFT	45.2 ↓ 1.7	5.7 ↓ 1.0	74.6 ↓ 2.4	24.0 ↓ 0.3
Global MFT w/ Best FFT **(Ours)**	49.0 ↑ 2.1	7.1 ↑ 0.4	74.1 ↓ 2.9	21.8 ↓ 2.5

Table 5: Performance comparison on LLaMA2-7B and LLaMA3.1-8B in multi-domain settings. Each part has three blocks containing 1) common LoRA and its continued variant (green), 2) FFT serving as an upper bound and its continued variant (blue), and 3) our MFT (red) with the other two vanilla masking baselines. The pre-trained performance serves as a lower bound.

	Method	Math		Coding		Instruction Following	
		GSM8K	Math	HumanEval	HumanEval+	IF-Eval	Alpaca-Eval
	Pre-Trained Model	15.2	2.5	25.8	22.4	34.3	0.5
Mixed Domain	Best LoRA	$40.8_{\pm 0.30}$	$6.1_{\pm 0.28}$	$22.6_{\pm 0.55}$	$18.3_{\pm 0.44}$	$37.3_{\pm 0.81}$	$0.7_{\pm 0.11}$
	Continued LoRA	$31.9_{\pm 0.24}$	$4.0_{\pm 0.21}$	$20.1_{\pm 0.42}$	$17.1_{\pm 0.36}$	$31.5_{\pm 0.92}$	$0.8_{\pm 0.09}$
	Best FFT	$45.5_{\pm 0.18}$	$8.1_{\pm 0.26}$	$29.7_{\pm 0.43}$	$26.7_{\pm 0.38}$	$43.6_{\pm 0.68}$	$1.0_{\pm 0.15}$
	Continued FFT	$44.1_{\pm 0.22}$ ↓ 1.4	$7.5_{\pm 0.27}$ ↓ 0.6	$21.1_{\pm 0.58}$ ↓ 8.6	$18.1_{\pm 0.43}$ ↓ 8.6	$38.6_{\pm 0.74}$ ↓ 5.0	$1.2_{\pm 0.16}$ ↑ 0.2
	Random Mask w/ Best FFT	$40.0_{\pm 0.33}$	$6.5_{\pm 0.32}$	$23.7_{\pm 0.57}$	$19.1_{\pm 0.49}$	$30.0_{\pm 0.95}$	$0.9_{\pm 0.14}$
	L1 Mask w/ Best FFT	43.3	7.7	25.7	22.3	32.8	1.2
	MFT w/ Best FFT **(Ours)**	**45.9**$_{\pm 0.17}$ ↑ 0.4	**8.4**$_{\pm 0.28}$ ↑ 0.3	**31.5**$_{\pm 0.46}$ ↑ 1.8	**27.3**$_{\pm 0.41}$ ↑ 0.6	**46.1**$_{\pm 0.71}$ ↑ 2.5	**1.9**$_{\pm 0.18}$ ↑ 0.9
	Pre-Trained Model	55.9	14.6	42.1	37.8	17.6	9.0
Multi Domain	Best LoRA	$72.8_{\pm 0.42}$	$24.5_{\pm 0.29}$	$63.1_{\pm 0.82}$	$55.5_{\pm 0.65}$	$38.2_{\pm 1.18}$	$11.0_{\pm 0.64}$
	Continued LoRA	$70.4_{\pm 0.36}$	$24.1_{\pm 0.27}$	$59.1_{\pm 0.75}$	$54.3_{\pm 0.59}$	$32.1_{\pm 1.05}$	$5.3_{\pm 0.32}$
	Best FFT	$72.7_{\pm 0.16}$	$24.4_{\pm 0.21}$	$64.4_{\pm 0.59}$	$60.2_{\pm 0.48}$	$60.0_{\pm 0.73}$	$11.8_{\pm 0.48}$
	Continued FFT	$71.7_{\pm 0.14}$ ↓ 1.0	$24.2_{\pm 0.20}$ ↓ 0.2	$62.6_{\pm 0.53}$ ↓ 1.8	$58.3_{\pm 0.44}$ ↓ 1.9	$59.7_{\pm 0.80}$ ↓ 0.3	$8.8_{\pm 0.38}$ ↓ 3.0
	Random Mask w/ Best FFT	$72.5_{\pm 0.39}$	$23.2_{\pm 0.28}$	$58.5_{\pm 0.73}$	$55.5_{\pm 0.62}$	$57.7_{\pm 1.22}$	$10.0_{\pm 0.51}$
	L1 Mask w/ Best FFT	72.5	24.1	63.4	59.1	58.9	10.7
	MFT w/ Best FFT **(Ours)**	**73.6**$_{\pm 0.20}$ ↑ 0.9	**25.2**$_{\pm 0.23}$ ↑ 0.8	**66.3**$_{\pm 0.51}$ ↑ 1.9	**61.4**$_{\pm 0.50}$ ↑ 1.2	**62.3**$_{\pm 0.69}$ ↑ 2.3	**12.2**$_{\pm 0.34}$ ↑ 0.4

Table 7: Performance of MFT compared to Best FFT on multi-modal benchmarks across two LLaVA backbones. MFT consistently outperforms Best FFT on all evaluated metrics.

Method	**GQA**	**MMMU**	**POPE**	**MME**	**SQA**	**TextVQA**
LLaVA (CLIP + Qwen2.5-0.5B)						
Pre-Trained Model	5.7	22.6	49.9	93.7	3.2	9.0
Best FFT	55.7	34.2	85.3	1197.9	57.6	44.2
MFT **(Ours)**	57.5 ↑ 1.8	35.0 ↑ 0.8	86.3 ↑ 1.0	1250.2 ↑ 52.3	58.8 ↑ 1.2	46.1 ↑ 1.9
LLaVA (SigLIP + Gemma-2B)						
Pre-Trained Model	2.6	26.4	56.6	586.2	2.8	5.6
Best FFT	60.9	31.3	87.4	1358.5	60.5	50.7
MFT **(Ours)**	61.8 ↑ 0.9	32.8 ↑ 1.5	88.7 ↑ 1.3	1379.4 ↑ 20.9	61.9 ↑ 1.4	51.5 ↑ 0.8

To further demonstrate MFT's generality and transferability, we conduct additional experiments on multimodal LLaVA backbones, as shown in Tab. 7. This setting evaluates MFT beyond pure language models, incorporating recent LLM backbones into multimodal architectures. Following the official LLaVA fine-tuning setup, where only the connector and LLM are fine-tuned, we learn a global mask across all corresponding trainable layers, including the connector and all LLM layers. Therefore, unlike the local MFT experiments, this setting does not rely on manually selected layer groups. Based on the above observation, we use a lower global sparsity ratio of 1% for multimodal MFT. The results show that MFT consistently outperforms Best FFT across all evaluated multimodal benchmarks for both LLaVA variants. These results suggest that global MFT can be effective when the sparsity level is properly controlled, and that the proposed mask-based post-FFT refinement is transferable from language-only LLMs to multimodal LLMs.

Comparison with Regularization Techniques. Tab. 8 compares MFT with common regularization techniques on MBPP using DeepSeek-Coder-1.3B (Guo et al., 2024). Compared with Best FFT, adding dropout yields a slight improvement from 50.5 to 50.7, whereas stronger weight decay does not improve final performance. In our experiments, dropout provides a short-term gain during continued FFT, but performance soon declines as training proceeds. Stronger weight decay mainly delays the onset of degradation rather than improving the best achievable result. In contrast, MFT

Table 8: Comparison of MFT with the best FFT and common regularization techniques.

Method	**MBPP**
Zero-Shot	46.8
Best FFT	50.5
+ Dropout	50.7
+ Stronger Weight Decay	50.5
+ MFT **(Ours)**	**51.2**

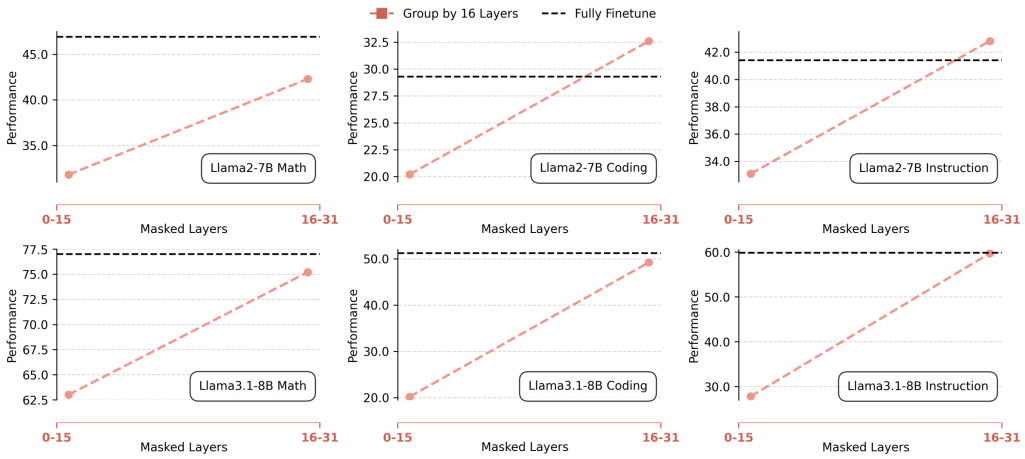


Figure 8: The local MFT ablation results of the larger layer-wise group (16-layer).

further improves the Best FFT checkpoint to 51.2, outperforming both regularization baselines. These results suggest that the gain of MFT cannot be simply explained by conventional regularization effects.

Conceptually, MFT is different from classical regularization methods. Regularization techniques such as dropout, L1/L2 penalties, and weight decay still optimize the model within the full hypothesis space, while constraining the optimization trajectory or parameter magnitude. MFT instead freezes the well-trained weights and directly optimizes a binary structural mask, thereby selecting a task-adaptive subnetwork from the fine-tuned model. Therefore, MFT is a structure-level post-FFT refinement method rather than a standard regularization strategy.

Larger Layer-wise Group Ablation. The results of larger layer-wise group (16 layers) ablation for local MFT are shown in Fig. 8.

Fine-grained Evaluation Details. We provide fine-grained evaluation details for FFT and MFT on three domains using LLaMA2-7B in Tab. 10 (domain-specific FFT) and Tab. 11 (domain mixed-up FFT). They show the epoch number for Best FFT, Continued FFT, Best MFT, and Complete MFT, along with the corresponding token usage. The number of Continued FFTs is consistently 4, whereas the number for Best FFT is less than 4. Based on the Best FFT, we then perform 2 epochs of MFT. Therefore, the number of Complete MFT is always equal to the number of Best FFT plus 2, but Best MFT may be less than it. For example, in the math domain in Tab. 10, Continual FFT runs for 4 epochs, while FFT achieves the best results at 2.4 epochs as Best FFT. Starting from 2.4 epochs, we perform MFT for an additional 2 epochs. Thus, the Complete MFT is 2.4 + 2 = 4.4 epochs, while the Best MFT is 4 epochs, indicating that MFT uses 4 - 2.4 = 1.6 epochs to achieve the best result based on Best FFT. For each setting, the number of tokens used is provided in brackets.

More Details of Theoretical Analysis. For our case, to theoretically support that our MFT has better optimization potential than FFT, we need to verify that $\Delta_{\text{train}} + \Delta_{\text{complexity}} < 0$ in Eq. 10.

Table 9: Exemplar training loss statistics of LLaMA2-7B models after training are stable.

LLaMA2-7B	Math	Coding	Instruction
FFT loss	0.101	0.098	0.125
MFT loss	0.085	0.054	0.060

For the first term, we first supplement exemplar training-loss statistics in Tab. 9. It shows the mean training loss after training stabilizes, indicating that MFT can further reduce training loss compared with the base model (the Best FFT).

Then, we can ensure the first term $\Delta_{\text{train}} < 0$.

For the second term, we need to compare the complexity of the encoded model. Given d as the total number of weights, z as the weights being masked out (as zero), and b as the bit width. Since we are

using weight-level masking granularity, we have

$$C(h_{\mathrm{FFT}}) \approx bd, \quad C(h_{\mathrm{MFT}}) \approx b(1-p)d + \log_2 \binom{d}{z}, \tag{11}$$

where $p = z/d$ represents the sparse ratio in our case. The first term of C_{MFT} is the encoding cost of non-zero weights, and the second is for the index cost of zero weights. Further, we have $\log_2 \binom{d}{z} \approx dH(p)$, where

$$H(p) = -p \log_2 p - (1-p) \log_2 (1-p). \tag{12}$$

Then, the second term can be described as

$$\Delta_{\mathrm{complexity}} = C(h_{\mathrm{MFT}}) - C(h_{\mathrm{FFT}}) \approx d[H(p) - bp]. \tag{13}$$

For practice, we always use $b = 8$ or 16. We use a 10% sparsity ratio for our MFT, applied to a 4-layer group, yielding a global sparsity ratio of $2.7\% \sim 3\%$ across two backbones. To consider the strict case, we use $b = 8$ and a 2.7% ratio in the calculation shown below.

$$H(p) - bp = H(0.027) - 8 \cdot 0.027 \approx -0.037 < 0. \tag{14}$$

Then, we have the second term $\Delta_{\mathrm{complexity}} < 0$.

Combining both the first and second terms, we have

$$U(h_{\mathrm{MFT}}) - U(h_{\mathrm{FFT}}) = \Delta_{\mathrm{train}} + \Delta_{\mathrm{complexity}} < 0. \tag{15}$$

Finally, we conclude that $U(h_{\mathrm{MFT}}) < U(h_{\mathrm{FFT}})$, meaning MFT can further reduce the test-loss upper bound compared with FFT and theoretically supports our proposed method.

More Training Cost Comparisons. We supplement the training cost comparisons with GPU memory usage, token count, and training time for the math and instruction-following domains in Fig. 9.

Table 10: Training epochs and corresponding number of used tokens (in brackets) on LLaMA2-7B under domain-specific FFT setting.

Domain	Best FFT	Continued FFT	Best MFT	Complete MFT
Math	2.4 epochs (5.90e8)	4.0 epochs (9.83e8)	4.0 epochs (9.83e8)	4.4 epochs (1.08e9)
Coding	2.7 epochs (3.32e8)	4.0 epochs (4.92e8)	3.6 epochs (4.43e8)	4.7 epochs (5.78e8)
Instruction Following	2.4 epochs (5.51e8)	4.0 epochs (9.18e8)	3.2 epochs (7.35e8)	4.4 epochs (1.01e9)

Table 11: Training epochs and corresponding number of used tokens (in brackets) on LLaMA2-7B under mixed-up FFT setting.

Domain	Best FFT	Continued FFT	Best MFT	Complete MFT
Math	2.0 epochs (2.67e9)	4.0 epochs (5.34e9)	3.7 epochs (3.09e9)	4.0 epochs (3.16e9)
Coding	2.0 epochs (2.67e9)	4.0 epochs (5.34e9)	3.3 epochs (2.83e9)	4.0 epochs (2.92e9)
Instruction Following	2.0 epochs (2.67e9)	4.0 epochs (5.34e9)	4.0 epochs (3.13e9)	4.0 epochs (3.13e9)

More Mask Ratio Ablation. The mask ratio ablation on the math domain using LLaMA2-7B and LLaMA3.1-8B is shown in Fig.10.

More Data Ratio Ablation. The MFT data ratio ablation on the coding domain using LLaMA2-7B and LLaMA3.1-8B is shown in Fig. 11.

More Details of Loss Landscape Visualization. A loss landscape visualization in the coding domain using LLaMA2-7B is shown in Fig. 12.

For the loss surface analysis, we present empirical results as loss-landscape visualizations in Fig. 7 and Fig. 12. Here, we provide further support for the theoretical view.

We assume the model after FFT (Best FFT) has Hessian as

$$H_{\mathrm{FFT}} = \nabla^2 L(\theta_{\mathrm{FFT}}) \succeq 0. \tag{16}$$

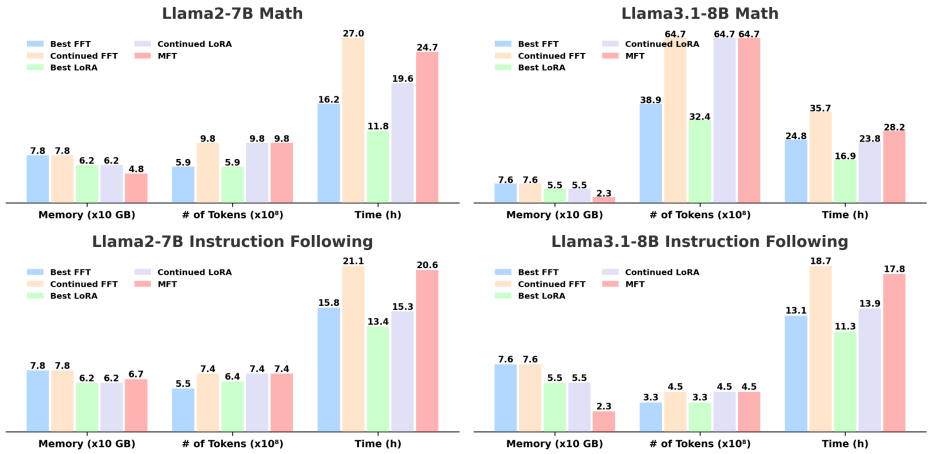


Figure 9: Training cost comparisons of GPU memory, used token number, and training time on math and instruction-following domains using LLaMA2-7B and LLaMA3.1-8B.

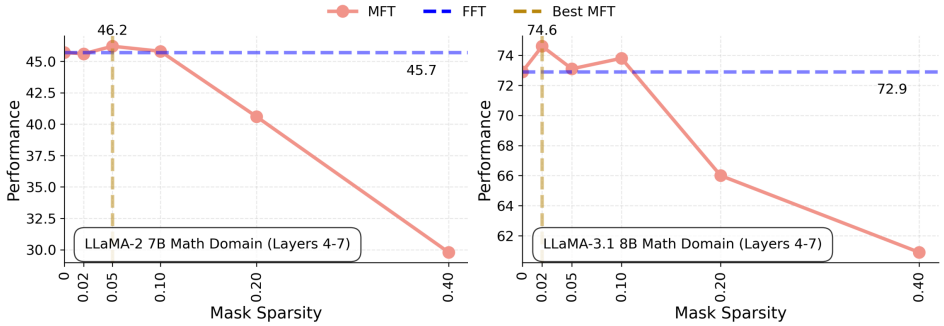


Figure 10: Masking ratio ablation study on the math domain using LLaMA2-7B and LLaMA3.1-8B.

Since MFT learns a mask and applies it to the model after FFT, which means the MFT model is a projection P of the FFT model, where the corresponding Hessian can be given by

$$H_{\mathrm{MFT}} = P H_{\mathrm{FFT}} P. \tag{17}$$

Since for $H_{\mathrm{FFT}} \succeq 0$ and its any projection (H_{MFT} for here), we have

$$\begin{aligned} \lambda_{\max}(H_{\mathrm{MFT}}) &\leq \lambda_{\max}(H_{\mathrm{FFT}}), \\ \mathrm{Tr}(H_{\mathrm{MFT}}) &\leq \mathrm{Tr}(H_{\mathrm{FFT}}). \end{aligned} \tag{18}$$

As $\lambda_{\max}$ and $\mathrm{Tr}(\cdot)$ represent model sharpness and mean curvature, lower values indicate a flatter model, meaning the MFT model generalizes better than the FFT model. Such analysis provides theoretical support and aligns with the empirical results shown in the landscape visualizations (the MFT model is visually flatter than the FFT model).

D More Discussion

Intuition. MFT is highly relevant to a range of well-established research topics, such as sparse network training and network pruning. However, inspired by their techniques, we approach sparsity from a novel perspective, shifting from typical model efficiency to a general capability perspective by removing negative components to improve performance. In other words, sparse models typically aim to compress the model while maintaining performance (often referred to as *subtraction*). However, we use sparsity as a tool for model enhancement (seen as *addition*). We expect that such

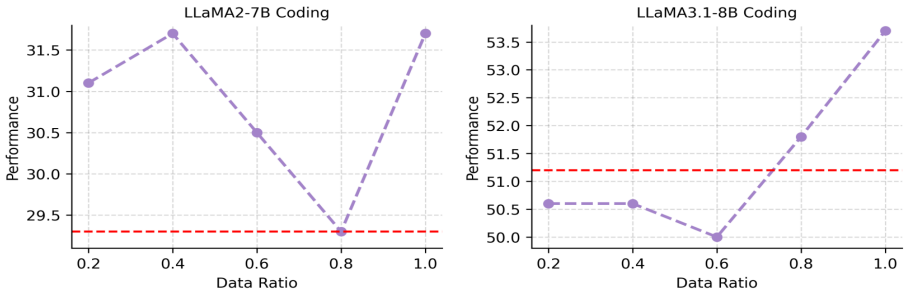


Figure 11: Data ratio ablation study on coding domain using LLaMA2-7B and LLaMA3.1-8B.

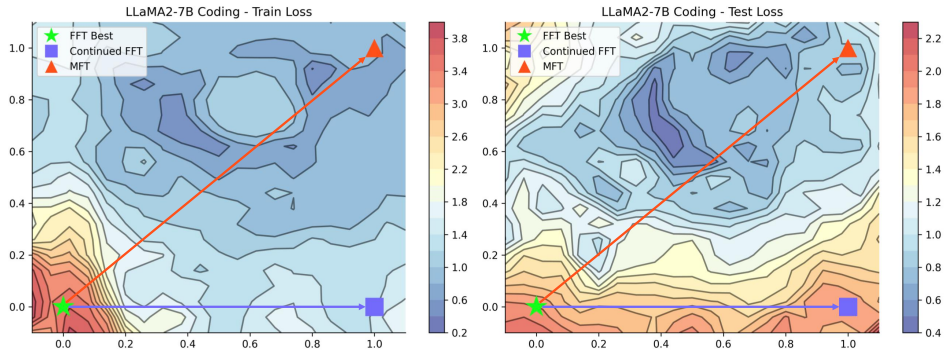


Figure 12: The loss landscape visualization on the coding domain using LLaMA2-7B.

counterthinking about sparse networks could inspire broader explorations in the future, enriching their research potential.