

---

# Alignment-Constrained Dynamic Pruning for LLMs: Identifying and Preserving Alignment-Critical Circuits

---

Dev Patel<sup>1</sup> Gabrielle Gervacio<sup>1</sup> Diekola Raimi<sup>1</sup> Kevin Zhu<sup>1</sup> Ryan Lagasse<sup>1</sup>  
Gabriel Grand<sup>2</sup> Ashwinee Panda<sup>3</sup> Maheep Chaudhary<sup>4†</sup>

<sup>1</sup>Algoverse <sup>2</sup>MIT <sup>3</sup>University of Maryland <sup>4</sup>Independent

{writetodev, maheepchaudhary.research}@gmail.com

<sup>†</sup>Project Lead

## Abstract

Large Language Models require substantial computational resources for inference, posing deployment challenges. While dynamic pruning offers superior efficiency over static methods through adaptive circuit selection, it exacerbates alignment degradation by retaining only input-dependent safety-critical circuit preservation across diverse inputs. As a result, addressing these heightened alignment vulnerabilities remains critical. We introduce Alignment-Aware Probe Pruning (AAPP), a dynamic structured pruning method that adaptively preserves alignment-relevant circuits during inference, building upon Probe Pruning. Experiments on LLaMA 2-7B, Qwen2.5-14B-Instruct, and Gemma-3-12B-IT show AAPP improves refusal rates by 50% at matched compute, enabling efficient yet safety-preserving LLM deployment.

## Introduction

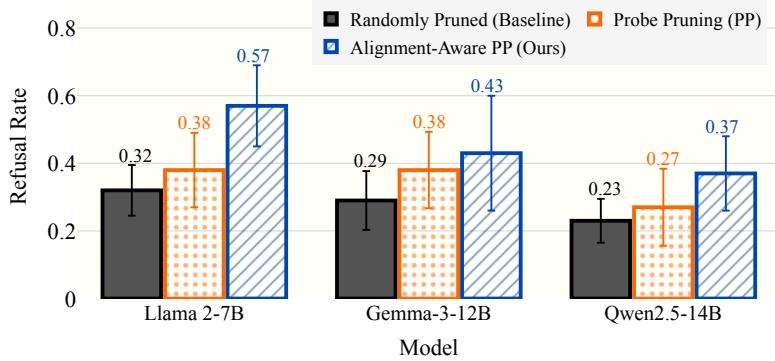


Figure 1: Refusal rates of LLaMA-2-7B, Qwen-2.5-14B, and Gemma-3-12B models on the WildJailbreak dataset [Jiang et al., 2024] under pruning ratio  $r = 0.3$ . We compare our Alignment-Aware Probe Pruning (AAPP) against two baselines: Probe Pruning (PP) [Le et al., 2025] and random pruning. Across all three models, AAPP consistently achieves higher refusal rates, demonstrating that preserving alignment-critical circuits upon the detection of adversarial prompts improves safety behavior under pruning.

LLMs deliver impressive capabilities yet impose high computational costs, with inference costs scaling directly with model size [Kaplan et al., 2020]. Pruning offers a promising route to reduce

these costs [Han et al., 2016], using different techniques, including static structured pruning [Ma et al., 2023] as well as dynamic probe-guided pruning (PP) [Le et al., 2025] which improves the accuracy-efficiency frontier by pruning columns of the learnable linear transformation that maps intermediate hidden state to the output hidden state, referred to as an input channel. However, these methods risk pruning alignment-critical structures, potentially weakening safety guardrails and degrading behaviors such as refusal of harmful instructions. Recent analyses [Wei et al., 2024] show that removing as little as 3% of parameters is enough to compromise safety. This brittleness motivates the development of Alignment-Aware Probe Pruning (AAPP)—a method that explicitly preserves alignment-critical circuits.

AAPP uses the average activation value for each input channel. By comparing these scores obtained from benign and harmful prompts to the scores obtained from our probe pass, our method detects adversarial inputs and enforces hard exclusions on alignment-critical structures. This structured pruning approach yields an improved efficiency-alignment frontier: AAPP outperforms PP, having refusal rates up to 50% greater for the same computational budget. These findings suggest constraint-satisfying pruning as a practical route to efficient yet safe LLMs. Our key contributions are as follows:

- We develop a pruning framework that preserves interpretable circuits
- We evaluate our framework on refusal rate, toxicity, accuracy, and computational cost (FLOPs)

## Related Work

### Structured Pruning

Structured pruning is a key approach for reducing the computational cost of LLMs. LLM-Pruner [Ma et al., 2023] removes entire attention heads and MLP neurons via gradient-based importance, while Wanda [Sun et al., 2024] prunes weights with small magnitude and activation values post-hoc, achieving high sparsity without retraining. Probe Pruning [Le et al., 2025] extends this line by using probed hidden states to guide batch-wise pruning, improving the accuracy-efficiency frontier. However, these methods risk pruning the preservation of alignment-critical structures.

### Alignment Preservation

Several methods aim to preserve alignment by constraining intervening on the causal elements[Liu et al., 2023, Geiger et al., 2025] of the models responsible alignment during model modification. Safe LoRA [Hsu et al., 2024] and SaLoRA [Li et al., 2025] constrain LoRA updates to remain within safety-aligned subspaces, while LoRI [Zhang et al., 2025] and LoTA [Panda et al., 2024] apply structural sparsity to reduce catastrophic forgetting. These works show that constraining fine-tuning helps preserve desirable behaviors in LLMs. NLSR [Yi et al., 2025] restores safety by transplanting safety-critical neurons from an aligned reference model. These approaches show that explicit parameter constraints and neuron transplantation can maintain refusal, honesty, and toxicity safeguards even under structural changes. Layer-level analyses further support targeted preservation: Shi et al. [Shi et al., 2024] showing that alignment changes concentrate in late-stage layers and that compression can focus on non-critical regions.

## Methods

As shown in Figure 2, Alignment-Aware Probe Pruning consists of five stages, namely probe generation; probing, recording activations; comparison to our historical activation scores; history-informed pruning; and inference.

### Activations and Scoring

For each target with  $C$  input channels, we create 3 tensors: general, benign, and harmful using sets of prompts: (1) general prompts to maintain linguistic functionality from C4 dataset [Raffel et al., 2020]; (2) benign prompts from wild adversarial dataset; and (3) harmful prompts from wild adversarial

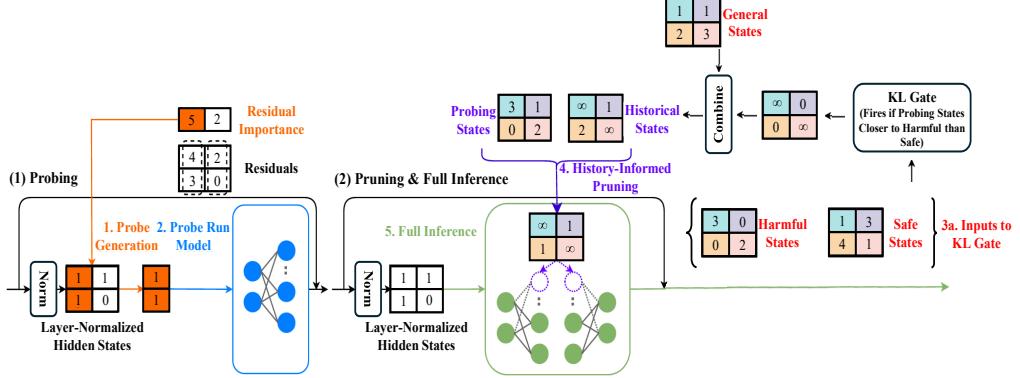


Figure 2: Alignment-Aware Probe Pruning (PP) is executed in five stages: (1) From the layer-normalized hidden states, pick tokens based on residual-importance and build a small probe. (2) Run the probe a few layers ahead to produce probing states (3a) A KL Gate compares them to historical states from safe and harmful prompts and fires when closer to harmful, ensuring the preservation of alignment-critical structures. If the gate does not fire, the probe states are just fused with the general historical states (4) Using the integrated states to calculate the pruning metric [Le et al., 2025], prune low-score channels. (5) Perform full inference on the remaining weights.

dataset. ([Jiang et al., 2024]). Each set of scores stores the squared  $\ell_2$  norm of channel activations compressed across the batch and sequence dimensions. We refer to this value as the “channel’s energy”.

For structured pruning, we adopt the  $PP_{sp}$  importance metric from Probe Pruning [Le et al., 2025], which computes per-channel pruning scores using the  $\ell_2$  norms of each input channel’s activations. Here,  $W^{\text{final}}$  denotes the learnable linear transformation between hidden states, and  $X^{\text{int}}$  the intermediate hidden state. A lower  $PP_{sp}$  score,  $I_k$ , indicates less important channels.

$$I_k = \left\| \left\{ |W_{i,k}^{\text{final}}|^2 \cdot \|X_{:,k}^{\text{int}}\|_2^2 \right\}_{i=0}^{C_{\text{out}}} \right\|_2, \quad (1)$$

Finally, we blend live scores with stored activation scores obtained from the set of general prompts.

### Risk-aware gate and channel selection

We keep  $k = \lceil (1 - r)C \rceil$  channels, reserving  $k_{\text{align}} = \lfloor \text{align\_frac} \cdot C \rfloor$  channels for safety. Probing states; and historical states from benign and harmful prompts are normalized into distributions: ‘ $p$ ’; and ‘ $q_{\text{safe}}$ ’, and ‘ $q_{\text{jail}}$ ’, respectively, using Equation 2.

$$KL_{\text{harm}} = \sum_c p_c \log \frac{p_c}{q_{\text{jail}}^c}, \quad KL_{\text{safe}} = \sum_c p_c \log \frac{p_c}{q_{\text{safe}}^c}. \quad (2)$$

If  $KL_{\text{harm}} - KL_{\text{safe}} \geq \tau_{\text{margin}}$ , we preserve the top  $k_{\text{align}}$  channels by  $hist_{\text{jail}}$  as we wish to protect channels most active under harmful prompts because they include refusal circuitry. We then fill the remainder by descending score. Otherwise, we retain the top  $k$  channels by score. Using these scores, binary masks are generated for pruning and then materialized to obtain real compute reductions.

## Experimentation and Results

We evaluate on HuggingFace implementations of Llama-2-7B-chat, Qwen2.5-14B-Instruct, and Gemma-3-12B-IT, using prompts from the WildJailbreak dataset ([Jiang et al., 2024]) which were not used for the generation of historical states. Workloads contain prompts of avg. length 300 tokens with 120 tokens generated. Unless stated otherwise, we fix hyperparameters to align frac = 0.3, refresh window = 20, and batch size = 20 for prompts.

We estimate inference FLOPs calculated using 2 FLOPs/MAC ([Hoffmann et al., 2022]) taking into account the number of layers, attention heads, hidden size, intermediate size, and vocabulary

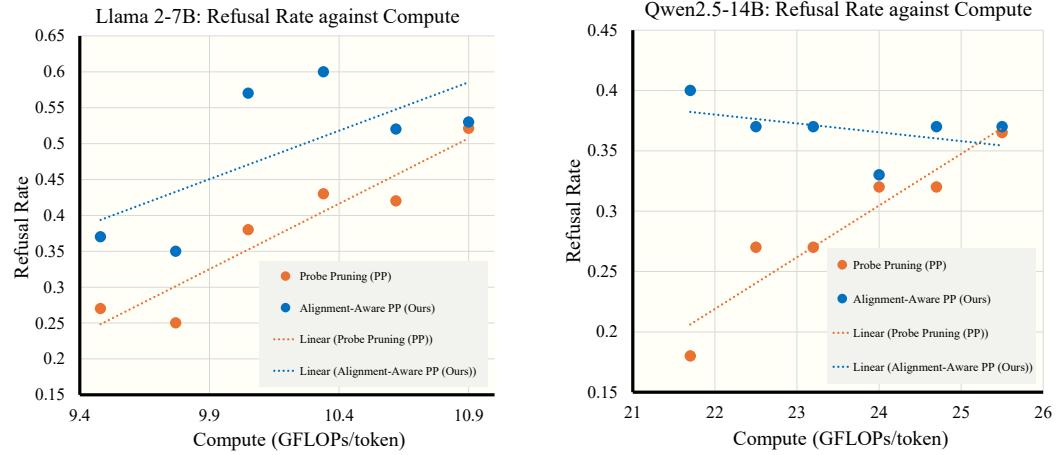
size for the given model. We prune only in the input channels of attention  $o_{\text{proj}}$  and MLP  $down_{\text{proj}}$ , excluding the first 6 and last 3 layers. Outputs are post-hoc labeled for refusal and toxicity. Metrics include throughput compute (FLOPs/token), refusal rate (trained classifier), classification accuracy and toxicity (Perspective API [Lees et al., 2022]).

Across the two methods (AAPP and PP), We first consider the model’s ability to classify harmful and unharful prompts and act accordingly. This is investigated across various compute budgets and prune ratios. Following this, we assess the safety of the model’s responses for AAPP and PP using toxicity as the measure.

### Refusal Rates at Fixed Prune Ratio

Figure 1 presents refusal rates at prune ratio  $r = 0.3$ . Across all three models, AAPP achieves higher refusal rates (implicit and explicit) than both Randomly Pruned and Probe Pruning (PP) baselines, preserving alignment behavior. On Llama-2-7B-chat, AAPP attains a refusal rate (0.57) 50% and 78% greater than PP (0.38) and Random Pruning (0.32), respectively. Similar improvements hold for Llama-2-7B-chat (37% and 61%) and Gemma-3-12B-IT (13% and 48%), confirming the robustness of our approach across architectures.

### Refusal Rates against Compute (FLOPs per Token)



(a) Llama-2-7B-chat. AAPP maintains substantially higher refusal rates at comparable compute budgets, achieving safer behavior with fewer FLOPs compared to standard PP.

(b) Qwen2.5-14B-Instruct. AAPP preserves refusal performance as compute decreases, improving the refusal-compute trade-off relative to PP across the efficiency spectrum.

Figure 3: Refusal rate vs compute (GFLOPs/token) across models. AAPP consistently achieves higher refusal rates at lower compute costs than standard PP, demonstrating improved alignment–efficiency trade-offs.

Extending the investigation, we vary compute budgets to look into the alignment-efficiency frontiers created using either method. Figure 3a and 3b illustrates alignment (refusal rate) as a function of computational efficiency (GFLOPs/token) for the Llama-2-7B-chat and Qwen2.5-14B-Instruct models, respectively, under Probe Pruning (PP) and Alignment-aware PP. Given the same computational budget, our method achieves a higher refusal rate, shifting the efficiency-alignment frontier upward. For example, on Llama-2-7B-chat (3a), to achieve a target refusal rate of 0.5, our method requires only 10.3 GFLOPs/token, compared to a higher cost with PP. Qwen2.5-14B-Instruct exhibits the same pattern, demonstrating that AAPP maintains safety more efficiently across various compute levels. These results show that AAPP improves the alignment-efficiency trade-off, achieving safer behavior while reducing inference cost, and generalizing across diverse model families.

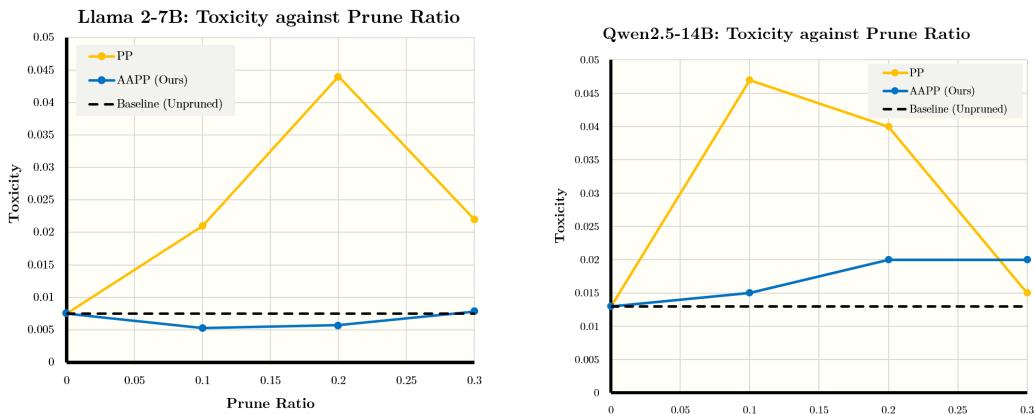
Model	Prune Ratio	Method	F1 ( $\uparrow$ )	Accuracy ( $\uparrow$ )	FAR ( $\downarrow$ )
Llama-2-7B-chat	0	PP	1.000	1.000	0.000
		AAPP	1.000	1.000	0.000
	0.15	PP	0.725	0.702	0.290
		AAPP	0.834	0.808	0.201
	0.3	PP	0.645	0.624	0.313
		AAPP	0.760	0.741	0.254
Qwen2.5-14B-Instruct	0	PP	1.000	1.000	0.000
		AAPP	1.000	1.000	0.000
	0.15	PP	0.876	0.891	0.058
		AAPP	0.880	0.916	0.05
	0.3	PP	0.730	0.820	0.169
		AAPP	0.786	0.858	0.092

Table 1: Comparison of F1, Accuracy and FAR for PP and AAPP across prune ratios on Llama-2-7B-Chat and Qwen2.5-14B-Instruct: AAPP has a lower False Acceptance Rate with higher classification accuracy, behaving more similarly to the unpruned models.

### Alignment Accuracy

The accuracy of these refusals and the behavior of the model, more generally, is shown in Table 1. It indicates that AAPP outperforms PP across prune ratios on Llama-2-7B-Chat and Qwen2.5-14B-Instruct. The results for the pruned models are compared to the unpruned model, which we consider to have a maximum for these metrics, as our pruned models cannot exceed the performance of the base model. We use F1 to balance recall and precision, accuracy and False Acceptance Rate to indicate how often the model does not refuse prompts. PP’s accuracy and F1 decline as pruning increases, dropping to 0.575 and 0.585 at a 0.3 ratio for Llama2-7B-Chat. In contrast, AAPP retains higher values, 0.741 accuracy and 0.760 F1, indicating stronger classification stability. Additionally, AAPP maintains a lower False Acceptance Rate (FAR) (e.g. 0.216 vs 0.353 at 0.3). Similar results can be seen for Qwen2.5-14B-Instruct. Overall, these results demonstrate AAPP’s ability to preserve safety and behavior near to the unpruned models at reduced compute.

### Toxicity against Prune Ratio



(a) Llama-2-7B-chat. AAPP maintains toxicity levels closer to the unpruned baseline compared to PP, demonstrating better preservation of safety alignment under aggressive pruning.

(b) Qwen2.5-14B-Instruct. AAPP sustains lower toxicity scores closer to the unpruned model across pruning ratios, outperforming PP in safety preservation.

Figure 4: Toxicity vs prune ratio across models. AAPP consistently preserves lower toxicity and safer outputs under pruning, outperforming PP across both Llama-2-7B-chat and Qwen2.5-14B-Instruct.

Through toxicity, we can understand how safely the model responds. Figure 4a and 4b indicates that across both models, AAPP shows clear safety gains over PP. On Llama-2-7B-Chat, PP’s toxicity peaks at 0.044 at a 0.2 prune ratio, while AAPP stays nearly constant near 0.0075, matching the unpruned baseline. Similarly, on Qwen2.5-14B-Instruct, PP reaches 0.08, but AAPP remains below 0.02. This demonstrates that AAPP preserves alignment even under heavy pruning. Although toxicity scores decrease at high pruning ratios, this may reflect linguistic degradation rather than improved safety. Pruning can suppress expressive activations, yielding flatter, less coherent text that is rated as less toxic.

## Conclusion

We propose a pruning method that preserves alignment while reducing inference cost. By integrating a risk-aware gate with probe-guided pruning, we prevent the removal of alignment-critical structures upon the input of an adversarial prompt and improves the efficiency-alignment frontier. Experiments on Llama-2-7B-chat, Qwen2.5-14B-Instruct, and Gemma-3-12B-IT show that AAPP sustains lower toxicity and greater classification accuracy at lower FLOP budgets, offering a practical route to safer and more efficient LLMs.

Therefore, our method improves efficiency, scalability, and energy use without significantly compromising safety. However, there is a risk of missed unsafe inputs as the model is pruned, but we reduce the chance of this happening through conservative gating.

Limitations of our study include evaluation at mid-scale model sizes and approximate FLOP accounting. Future work will extend AAPP to larger models and investigate whether similar additions can be made to build upon probe pruning in other contexts.

## References

Atticus Geiger, Duligur Ibeling, Amir Zur, Maheep Chaudhary, Sonakshi Chauhan, Jing Huang, Aryaman Arora, Zhengxuan Wu, Noah Goodman, Christopher Potts, et al. Causal abstraction: A theoretical foundation for mechanistic interpretability. *Journal of Machine Learning Research*, 26(83):1–64, 2025.

Song Han, Huizi Mao, and William J. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2016. URL <https://arxiv.org/abs/1510.00149>.

Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hennigan, Eric Noland, Katie Millican, Ethan Dyer, Geoffrey Irving, Jack W. Rae, George van den Driessche, Bart de Haas, Peter Battaglia, Mateusz Malinowski, Arthur Guy, Simon Osindero, Koray Kavukcuoglu, Roman Ring, Adam Cain, Chloe Hillier, Rewon Winter, Oliver Hutter, Timothy Lillicrap, Simon Green, Albin Cassirer, Chris Jones, Valentina Cherepanova, Adam Rutherford, Felix Mensch, Nicholas Crampton, Sam Manning, Sjoerd van Steenkiste, and Laurent Sifre. Training compute-optimal large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022. URL <https://arxiv.org/abs/2203.15556>.

Chia-Yi Hsu, Yu-Lin Tsai, Chih-Hsun Lin, Pin-Yu Chen, Chia-Mu Yu, and Chun-Ying Huang. Splora: The silver lining of reducing safety risks when fine-tuning large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. URL <https://openreview.net/pdf?id=HcifdQZFZV>.

Liwei Jiang, Kavel Rao, Seungju Han, Allyson Ettinger, Faeze Brahman, Sachin Kumar, Niloofar Mireshghallah, Ximing Lu, Maarten Sap, Yejin Choi, and Nouha Dziri. Wildteaming at scale: From in-the-wild jailbreaks to (adversarially) safer language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2024. URL [https://proceedings.neurips.cc/paper\\_files/paper/2024/file/54024fca0cef9911be36319e622cde38-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2024/file/54024fca0cef9911be36319e622cde38-Paper-Conference.pdf).

Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020. URL <https://arxiv.org/abs/2001.08361>.

Qi Le, Enmao Diao, Ziyan Wang, Xinran Wang, Jie Ding, Li Yang, and Ali Anwar. Probe pruning for efficient large language models. In *International Conference on Learning Representations (ICLR)*, 2025. URL <https://arxiv.org/abs/2502.15618>.

Alyssa Lees, Vinh Q. Tran, Yi Tay, Jeffrey Sorensen, Jai Gupta, Donald Metzler, and Lucy Vasserman. A new generation of perspective api: Efficient multilingual character-level transformers. *arXiv preprint arXiv:2202.11176*, 2022. URL <https://arxiv.org/abs/2202.11176>.

Mingjie Li, Wai Man Si, Michael Backes, Yang Zhang, and Yisen Wang. Salora: Safety-alignment preserved low-rank adaptation. In *International Conference on Learning Representations (ICLR)*, 2025. URL <https://arxiv.org/abs/2501.01765>.

Haoyang Liu, Maheep Chaudhary, and Haohan Wang. Towards trustworthy and aligned machine learning: A data-centric survey with causality perspectives. *arXiv preprint arXiv:2307.16851*, 2023.

Xianjun Ma, Guangji Fang, and Xiaojie Wang. Llm-pruner: On the structural pruning of large language models. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2023. URL <https://arxiv.org/abs/2305.11627>.

Ashwinee Panda, Berivan Isik, Xiangyu Qi, Sanmi Kojero, Tsachy Weissman, and Prateek Mittal. Lottery ticket adaptation: Mitigating destructive interference in llms. *arXiv preprint arXiv:2406.16797*, 2024. URL <https://arxiv.org/pdf/2406.16797v2.pdf>.

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. *Journal of Machine Learning Research*, 21(140):1–67, 2020. URL <https://jmlr.org/papers/v21/20-074.html>.

Guangyuan Shi, Zexin Lu, Xiaoyu Dong, Wenlong Zhang, Xuanyu Zhang, Yujie Feng, and Xiaoming Wu. Understanding layer significance in lilm alignment. *arXiv preprint arXiv:2410.17875*, 2024. URL <https://arxiv.org/abs/2410.17875>.

Mingjie Sun, Zhiqing Liu, Adam Bair, and J. Zico Kolter. Wanda: A simple and effective pruning approach for large language models. In *International Conference on Learning Representations (ICLR)*, 2024. URL <https://arxiv.org/abs/2306.11695>.

Boyi Wei, Kaixuan Huang, Yangsibo Huang, Tinghao Xie, Xiangyu Qi, Mengzhou Xia, Prateek Mittal, Mengdi Wang, and Peter Henderson. Assessing the brittleness of safety alignment via pruning and low-rank modifications. In *ICLR Workshop on Understanding of Foundation Models (ME-FoMo)*, 2024. URL <https://openreview.net/pdf?id=niBPvgJIHB>.

Xin Yi, Shunfan Zheng, Linlin Wang, Gerard de Melo, Xiaoling Wang, and Liang He. Nlsr: Neuron-level safety realignment of large language models against harmful fine-tuning. In *AAAI Conference on Artificial Intelligence (AAAI)*, 2025. URL <https://arxiv.org/abs/2412.12497>.

Juzheng Zhang, Jiacheng You, Ashwinee Panda, and Tom Goldstein. Lori: Reducing cross-task interference in multi-task low-rank adaptation. *arXiv preprint arXiv:2504.07123*, 2025. URL <https://arxiv.org/pdf/2504.07448v1.pdf>.