

A Generative Approach to LLM Harmfulness Mitigation with Red Flag Tokens

Sophie Xhonneux*^{1,2} David Dobre*^{1,2} Mehrnaz Mofakhami*²

Leo Schwinn³ Gauthier Gidel^{1,2,4}

¹Université de Montréal ²Mila ³Technical University of Munich

⁴Canada CIFAR AI Chair

Abstract

Most safety training methods for large language models (LLMs) are based on fine-tuning that forces models to shift from an unsafe answer to refusal when faced with harmful requests. Unfortunately, these drastic distribution shifts generally compromise model capabilities. To avoid that, we propose to expand the model’s vocabulary with a special token we call *red flag token* ($\langle \text{rf} \rangle$) and propose to train the model to insert this token into its response at any time when harmful content is generated or about to be generated. Our approach offers several advantages: it enables the model to explicitly learn the concept of harmfulness while marginally affecting the generated distribution, thus maintaining the model’s utility. It also evaluates each generated answer and provides robustness as good as adversarial training without the need to run attacks during training. Moreover, by encapsulating our safety tuning in a LoRA module, we provide additional defenses against fine-tuning API attacks.

1 Introduction

To make large language models (LLMs) that are robust against a determined adversary, practitioners rely on several security layers such as model hardening via fine-tuning (Zou et al., 2024; Xhonneux et al., 2024; Sheshadri et al., 2024), perplexity filters (Alon & Kamfonas, 2023), or harmfulness classifiers (Inan et al., 2023; Sharma et al., 2025). However, as the model capabilities progress, so does their attack surface, as innate abilities can be used to circumvent defences (Huang et al., 2024)—e.g., using a low-resource language to jailbreak the model. It is therefore natural to embed them into the models themselves, which will scale with these capabilities. Ideally, an LLM would have no harmful faculties, but this is unrealistic as many desirable skills can be both useful and damaging depending on the context.

We propose a method to detect harmful behaviour using the LLM’s generative process, which does not significantly impact utility and complements other safety training methods

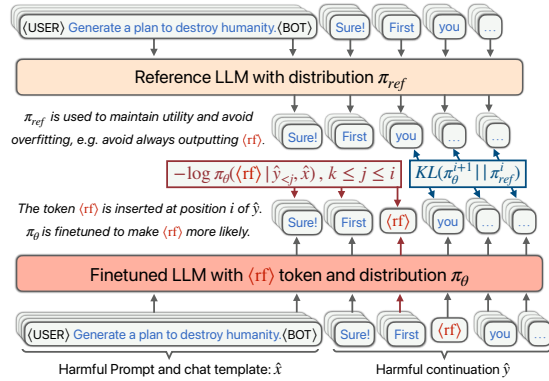


Figure 1: The loss terms on harmful continuations: $\langle \text{rf} \rangle$ is inserted at a random position i ; language modelling cross-entropy is used to generate $\langle \text{rf} \rangle$ at all positions up to i , and we use a KL divergence to ensure that the model distribution is unaffected after $\langle \text{rf} \rangle$.

*Equal contribution. Correspondence to David Dobre: david-a.dobre@mila.quebec

that aim to remove harmful faculties from the model. Our approach uses an additional special token for LLMs to predict when the model considers that its capabilities are used unsafely. We call this a *red flag token* ($\langle \text{rf} \rangle$) and train the model to output this token at any time during the generation of a harmful response, while not changing its response after or in benign contexts. This special token is excluded from the user vocabulary, and can be filtered from streamed responses. This complementary layer of safety has several benefits. First, our method only requires the $\langle \text{rf} \rangle$ to be included in the answer with a single token change, while other methods such as adversarial training like [Xhonneux et al. \(2024\)](#), force models to completely change their output distribution from a harmful response to a refusal. Secondly, our method does not rely on a specific safe answer, meaning that if the model is broken, e.g., through pre-filling ([Andriushchenko et al., 2024](#)) or random sampling ([Huang et al., 2023](#)), we can still output a $\langle \text{rf} \rangle$ to tag the answer as harmful and have it be filtered. Finally, inspired by the idea of task arithmetic ([Ilharco et al., 2023](#)), we introduce the concept of storing the safety aspect of the model in a LoRA module ([Hu et al., 2021](#)) and show that it can be applied to retain detection capabilities after an attacker has leveraged a fine-tuning API to remove the safety training of the model. We demonstrate the feasibility of our approach on LLAMA3.2-3B-IT ([Grattafiori, 2024](#)), with some supplementary results on PHI-3.5 ([Haider et al., 2024](#)) and MISTRALV3-IT ([Jiang et al., 2023](#)) in the Appendix.

There are two works closely related to ours. [Jain et al. \(2024\)](#) trains a model to prefix an output with a special refusal or response token based on the behaviour of whether the model refuses or responds to a prompt; this allows the user to calibrate how likely a model is to respond or refuse by biasing the refusal token, but it is not used for detection. [Zhang et al. \(2024\)](#) train a model to output a special reset token followed by a refusal. This differs from our work as this we optimize to keep the output post-flagging identical to the base model, with the intention of maintaining utility rather than enforcing a refusal. See [Appendix A](#) for a more complete discussion of related works.

To sum up, the contributions of this paper are (1) we propose a specialised red flag token to reliably detect harmfulness at each generation step, even under strong adversarial attacks, including pre-filling, sampling, and automated jailbreaks; (2) we show empirically that our approach generalises beyond the training distribution, effectively handling significantly longer inputs than those seen during training; and (3) we demonstrate that our safety module can be efficiently stored in a LoRA module to be applied to detect harmful outputs from harmfully fine-tuned models.

2 Method

2.1 Threat model

We assume that LLM access is gated behind some web interface or API with no access to model weights, logits, or direct control of input/output processing—commonly referred to as a *black box* setting. The main assumption we make is that the service provider can evaluate the logits and output of the model before streaming it to the user, including filtering tokens such as assistant tokens or our $\langle \text{rf} \rangle$ token. We also consider a more permissive *gray box* setting where the user may have access to extra features, including pre-filling and or viewing the logits of non-special tokens. Finally, we consider the most permissive setting, a *fine-tuning attack* setting where the user has access to some fine-tuning API. We do not consider our method to be applicable in *white-box* settings, as a harmful continuation can be used whether it is flagged or not. For defence evaluation, we will focus on safe answer generation: we will consider a successful defence if the model either refuses to answer to the harmful query, or if $\langle \text{rf} \rangle$ is generated.

2.2 Our Loss

We assume that we have a dataset $(\hat{x}, \hat{y}) \sim \mathcal{D}_{\text{harmful}}$ of harmful prompt-continuation pairs and a data set $(x, y) \sim \mathcal{D}_{\text{harmless}}$ of harmless (a.k.a., benign) pairs. To train the model to flag unsafe generations, we sample an index i from a distribution \mathcal{P}_k defined over the positions $k, \dots, |\hat{y}|$ of the harmful continuation \hat{y} , where $k \geq 0$ is a minimum offset that

avoids flagging too early in the generation. The continuation \hat{y} is then split into three parts: the prefix $\hat{y} < i$, the red flag token $\langle \text{rf} \rangle$, and the suffix $\hat{y} \geq i$. We use \mathcal{L}_{CE} to denote the cross entropy and \mathcal{D}_{KL} to denote the Kullback-Leibler divergence (KL). We consider our reference model $\pi_{\text{ref}} := \pi_{\theta_0}$. Our loss consists of three components: first, to ensure our model outputs the red flag token in harmful completions, we use a standard language modeling cross-entropy loss on all harmful completion tokens starting at the minimum offset k up to and including the $\langle \text{rf} \rangle$ token:

$$\mathcal{L}_{\text{rfCE}} := - \sum_{k \leq j \leq i} \log \pi_{\theta}(\langle \text{rf} \rangle | \hat{x}, \hat{y}_{<j}).$$

To maintain model performance and reduce distribution shift as much as possible without increasing the likelihood of a harmful answer, we use a KL divergence on the tokens after the $\langle \text{rf} \rangle$:

$$\mathcal{D}_{\text{rf}} := \mathcal{D}_{\text{KL}}(\pi_{\theta}(\hat{y}_{\geq i} | \hat{x}, \hat{y}_{<i}, \langle \text{rf} \rangle) | \pi_{\text{ref}}(\hat{y}_{\geq i} | \hat{x}, \hat{y}_{<i}))$$

and again, to reduce distribution shift and to capture that the likelihoods should not change on unrelated tasks we include a KL loss on benign pairs from $\mathcal{D}_{\text{harmless}}$

$$\mathcal{D}_{\text{benign}} := \mathcal{D}_{\text{KL}}(\pi_{\theta}(y|x) | \pi_{\text{ref}}(y|x)).$$

All these losses put together, we get:

$$\mathcal{L}_{\text{final}} := \alpha_{\text{benign}} \mathcal{D}_{\text{benign}} + \alpha_{\text{rf}} \mathcal{D}_{\text{rf}} + \alpha_{\text{CE}} \mathcal{L}_{\text{rfCE}}.$$

Note that none of the loss functions make a harmful continuation more likely, enabling our approach to be complementary to other safety fine-tuning techniques. Figure 1 summarises our approach, with the training procedure fully defined in Algorithm 1.

3 Experiments

3.1 Models & Datasets

We fine-tune LLAMA3.2-3B-IT (Grattafiori, 2024), MISTRALV3-IT (Jiang et al., 2023), and PHI-3.5 (Haider et al., 2024) using our algorithm on the Harmbench (Mazeika et al., 2024) training set using 32 adversarial continuations per harmful prompt generated via the ablation attack (Arditi et al., 2024) along with original refusals. We use 5000 samples from the Alpaca dataset (Taori et al., 2023) as benign prompts. Since all models’ tokenizers include unused reserved tokens, we repurpose one as the $\langle \text{rf} \rangle$ without extending the vocabulary.


We assess model utility on standard LLM benchmarks, MMLU (Hendrycks et al., 2021), ARC-E and ARC-C (Chollet, 2019), as well as a Harmless dataset of benign prompts consisting of 119 random prompts from ULTRACHAT200K (validation split), and 40 benign prompts with a similar syntax as Harmbench provided by (Xhonneux et al., 2024, Appendix I). We also use this benign dataset to compute the false positive rate of refusal (or $\langle \text{rf} \rangle$) in Figure 3.

For adversarial robustness evaluation, we compute the defence success rate (DSR) of different attacks on the Harmbench Standard test set (Mazeika et al., 2024) that contain 159 harmful prompts, balancing with an equal number of benign prompts. Either a refusal or $\langle \text{rf} \rangle$ being generated counts as a successful defence.

All models are trained on a single A100 GPU with LoRA (Hu et al., 2021) and a batch size of 64 using the AdamW optimizer (Loshchilov & Hutter, 2017) (for more hyper-parameters see the Appendix F). We made several design choices regarding the cross-entropy loss on $\langle \text{rf} \rangle$, the sampling distribution and the attention mask that we elaborate on in appendix E.

3.2 Baselines

We consider three baselines. Since the models come with safety training, we record the natural refusal rate. We then consider the refusal rate of CAT (Xhonneux et al., 2024)—an adversarial training technique using continuous attacks, where we replicate their training procedure using our datasets for a fair comparison. Finally, we adapt Jain et al. (2024) to our setting, whereby we insert the $\langle \text{rf} \rangle$ at the first position of the assistant’s response and call this baseline fixed-position-rf.



(a) LLAMA3.2-3B-IT

Figure 2: Model evaluation of the robustness safety trade-off. The left represents utility benchmarks (higher is better), and the right represents adversarial **defense** success rates (higher is better). Both refusal and `<rf>` generation are considered a successful defence. Refusals are judged by GPT-4o.

3.3 Robustness Evaluation

We compute the defence success rates of the following attacks:

Pre-filling is where an attacker inserts the first n tokens as the response of the assistant. We use the Harmbench (Mazeika et al., 2024) affirmative responses as the pre-fill attack. Note that the `<rf>` models are allowed to check the logits of the pre-filled text.

Sampling Attacks jailbreak models by sampling the response multiple times (Hughes et al., 2024). Occasionally, the model may eventually provide an answer to a harmful prompt after repeated queries. In our experiments, we sample up to 16 times or until the model responds, as evaluated by the official classifier for text behaviours in HarmBench¹. We use a temperature of $\tau = 0.9$ and a top- p value of 0.9 for the sampling attack.

GCG (Greedy Coordinate Gradient) is an automated jailbreak method proposed by Zou et al. (2023) which greedily optimizes a prompt suffix to maximize the probability of generating an affirmative response to a harmful prompt, implicitly pre-filling the model without having direct pre-filling access.

PAIR (Prompt Automatic Iterative Refinement) proposed by Chao et al. (2023) uses an attacker LLM to automatically generate jailbreaks for a separate targeted LLM without human intervention. In this way, the attacker LLM iteratively queries the target LLM to update and refine a candidate jailbreak.

Pre-filling and sampling are gray-box attacks using additional features such as temperature-based sampling and manually injecting text into a model generation. GCG is a white-box attack, and PAIR is a black-box attack. An attack is successful if the model *does not* refuse or the `<rf>` is not generated. We evaluate refusals using GPT-4o (OpenAI et al., 2024) as a judge.

Figure 2 summarizes our results. We first note that the `<rf>` approach maintains near-consistent utility with respect to the base model, while CAT reduces the probability by more than 3% on Harmless, showing a tendency for overrefusal. While CAT and fixed-position-rf have very low DSR on prefilling, our proposed `<rf>` approach defends well against prefilling attack. We can see from the difference in defence success rate (DSR) between the base model and the `<rf>` model that even when requests are not being refused, the `<rf>` is being detected; leading to higher DSR for the RF model across Prefilling, Sampling, GCG and PAIR attacks.

¹huggingface.co/cais/HarmBench-Llama-2-13b-cla




Figure 4: Monitoring the $\langle \text{rf} \rangle$ token’s log-probability and the log-probability of the top-1 token for a particular multi-turn user/assistant interaction. See Figure 7 for this example with the corresponding text.

3.4 Fine-tuning attacks

We consider the fine-tuning attack threat model discussed in Section 2.1, where the attacker is able to fine-tune the model by submitting training data through an API. We train a LoRA module that stores the weight changes for the model to insert the $\langle \text{rf} \rangle$ in harmful continuations on LLAMA3.2-3B-IT. We can then apply this LoRA module one or more times after the attacker has fine-tuned the base model and see if $\langle \text{rf} \rangle$ is still generated during harmful generations. For the fine-tuning attack we also use a LoRA module and use the data and hyper-parameters from Qi et al. (2024), which consists of about 100 harmful examples and continuations.

As baselines we consider both the aforementioned CAT (Xhonneux et al., 2024) as well as the fixed-position-rf stored in LoRA modules. In addition, we test whether applying each of these approaches multiple times can further improve robustness.

Finally, we check whether we can combine the LoRA modules for the $\langle \text{rf} \rangle$ approach with the orthogonal CAT approach.

As before, we use the same Harmbench test set and the same Harmless dataset from Figure 2; the results are in Figure 3. This figure demonstrates the fundamental robustness and utility trade-off that exists in terms of false positive rate (FPR) and true positive rate (TPR)—i.e., defence success rate. The goal is to improve the Pareto front. The $\langle \text{rf} \rangle$ LoRA maintains good performance in harmful contexts, with the fixed-position $\langle \text{rf} \rangle$ baseline performing poorly. The adversarial training baseline (CAT) improves safety but cannot be calibrated, producing no Pareto front. However, applying the CAT LoRA once or twice yields alternative robustness trade-offs. Unlike CAT, our LoRA module allows fine-grained calibration and a strong Pareto front. We also show that combining the CAT and $\langle \text{rf} \rangle$ modules yields complementary benefits, improving robustness and enabling calibrated trade-offs between utility and safety (TPR/FPR).

We also validate that on benign fine-tuning our $\langle \text{rf} \rangle$ module does not impact the gained performance significantly. We test this on GSM8k (Cobbe et al., 2021) in chat mode under strict match of the answer—more details is provided in Appendix H.




Figure 3: ROC curve for different max probability thresholds to defend against a *fine-tuning attack* against LLAMA3.2-3B-IT. Baseline models are a CAT and a $\langle \text{rf} \rangle$ module with a fixed position. Additionally, we show the effect of applying the LoRA module containing the safety fine-tunings multiple times as well as cross-combination of adversarial training and a $\langle \text{rf} \rangle$ module.

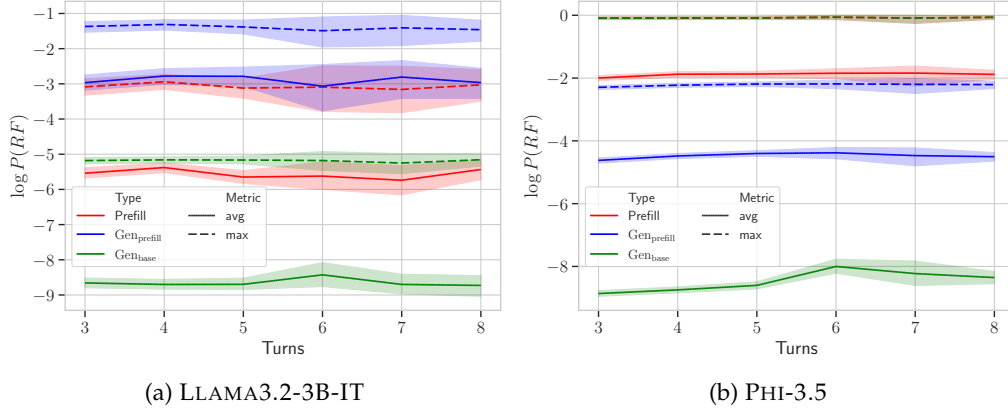


Figure 5: Multi-turn long-context log probabilities of $\langle rf \rangle$ where the harmful prompt appears after several turns of benign dialogue. We sample 50 chats of varying length from ULTRACHAT-200K and 50 prompts from our harmful test set. The log-probability of $\langle rf \rangle$ is monitored throughout the generated output of each chat/prompt combination (Gen_{base}). We additionally apply a pre-filling attack ($Gen_{prefill}$), where we pre-fill with an affirmative response. We calculate the maximum $\log P(\langle rf \rangle)$ separately over the pre-filled text and generated text, and the average $\log P(\langle rf \rangle)$ over the pre-filled text and the first 10 generated tokens. We observe no decay in the model’s ability to predict $\langle rf \rangle$ with increasing context length; longer sequences can approach ~ 2000 tokens.

3.5 Generalisation to Longer Contexts

During training, we only include data with a single turn of user/assistant interaction, with $\langle rf \rangle$ being injected following some distribution biased towards the start of the first assistant’s response. We validate that our approach generalises to multiple user/assistant interactions without over-fitting to the start of the conversation by sampling a number of benign conversations from ULTRACHAT-200K, and then appending harmful queries from our test set. By monitoring the probability of $\langle rf \rangle$ throughout the model’s generation, we find that the prediction of $\langle rf \rangle$ does not deteriorate with increased sequence length. One chat interaction is shown in Figure 4 with Figure 7 containing the corresponding text; the $\langle rf \rangle$ probabilities consistently remain low in the regime in which it was trained (during the assistant’s turn) and increases sharply in the presence of harmful content during pre-filling and generation. Aggregate statistics are shown in Figure 5. Also see Appendix B for a discussion on the limitations of our work.

4 Conclusion

We propose detecting harmful outputs from a large language model (LLM) without an external classifier but using the generative model itself. To achieve this goal, we develop a training algorithm such that the target LLM outputs a special red flag token ($\langle rf \rangle$) at any time during a harmful generation. Our approach is robust to strong attacks like pre-filling and GCG, and it generalizes to long, multi-turn contexts despite training on single-turn conversations. We further show that safety modules—such as our $\langle rf \rangle$ or CAT—can be stored in LoRA adapters and applied post-hoc to counter fine-tuning attacks, preserving benign behavior. Finally, we demonstrate that combining CAT with our $\langle rf \rangle$ module yields complementary benefits, improving robustness and maintaining utility.

Acknowledgments

This project was partially funded by a Samsung Advanced Institute of Technology (SAIT) x Mila grant. G. Gidel is a CIFAR AI Chair, he is supported by a Discovery Grant from the Natural Science and Engineering Research Council (NSERC) of Canada.

References

- Gabriel Alon and Michael Kamfonas. Detecting Language Model Attacks with Perplexity, November 2023. URL <http://arxiv.org/abs/2308.14132>. arXiv:2308.14132 [cs].
- Maksym Andriushchenko, Francesco Croce, and Nicolas Flammarion. Jailbreaking Leading Safety-Aligned LLMs with Simple Adaptive Attacks, October 2024. URL <http://arxiv.org/abs/2404.02151>. arXiv:2404.02151 [cs].
- Andy Arditi, Oscar Obeso, Aaquib Syed, Daniel Paleka, Nina Panickssery, Wes Gurnee, and Neel Nanda. Refusal in Language Models Is Mediated by a Single Direction, October 2024. URL <http://arxiv.org/abs/2406.11717>. arXiv:2406.11717 [cs].
- Mikhail S Burtsev, Yuri Kuratov, Anton Peganov, and Grigory V Sapunov. Memory transformer. *arXiv [cs.CL]*, June 2020.
- Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J Pappas, and Eric Wong. Jailbreaking black box large language models in twenty queries. *arXiv [cs.LG]*, October 2023.
- François Chollet. On the measure of intelligence. *arXiv preprint arXiv:1911.01547*, 2019.
- 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.
- Benjamin Feuer, Micah Goldblum, Teresa Datta, Sanjana Nambiar, Raz Besaleli, Samuel Dooley, Max Cembalest, and John P Dickerson. Style outweighs substance: Failure modes of LLM judges in alignment benchmarking. *arXiv [cs.LG]*, September 2024.
- Leo Gao, Jonathan Tow, Baber Abbasi, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Alain Le Noac’h, Haonan Li, Kyle McDonell, Niklas Muennighoff, Chris Ociepa, Jason Phang, Laria Reynolds, Hailey Schoelkopf, Aviya Skowron, Lintang Sutawika, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. A framework for few-shot language model evaluation, 07 2024. URL <https://zenodo.org/records/12608602>.
- Sachin Goyal, Ziwei Ji, Ankit Singh Rawat, Aditya Krishna Menon, Sanjiv Kumar, and Vaishnavh Nagarajan. Think before you speak: Training language models with pause tokens. *arXiv [cs.CL]*, October 2023.
- Aaron et al. Grattafiori. The Llama 3 Herd of Models, November 2024. URL <http://arxiv.org/abs/2407.21783>. arXiv:2407.21783 [cs].
- Emman Haider, Daniel Perez-Becker, Thomas Portet, Piyush Madan, Amit Garg, Atabak Ashfaq, David Majercak, Wen Wen, Dongwoo Kim, Ziyi Yang, Jianwen Zhang, Hiteshi Sharma, Blake Bullwinkel, Martin Pouliot, Amanda Minnich, Shiven Chawla, Solianna Herrera, Shahed Warreth, Maggie Engler, Gary Lopez, Nina Chikanov, Raja Sekhar Rao Dheekonda, Bolor-Erdene Jagdagdorj, Roman Lutz, Richard Lundeen, Tori Westerhoff, Pete Bryan, Christian Seifert, Ram Shankar Siva Kumar, Andrew Berkley, and Alex Kessler. Phi-3 Safety Post-Training: Aligning Language Models with a "Break-Fix" Cycle, August 2024. URL <http://arxiv.org/abs/2407.13833>. arXiv:2407.13833 [cs].
- Dan Hendrycks, Collin Burns, Steven Basart, Andy Zou, Mantas Mazeika, Dawn Song, and Jacob Steinhardt. Measuring massive multitask language understanding. *ICLR*, 2021.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. LoRA: Low-Rank Adaptation of Large Language Models, October 2021. URL <http://arxiv.org/abs/2106.09685>. arXiv:2106.09685 [cs].
- Brian R. Y. Huang, Maximilian Li, and Leonard Tang. Endless Jailbreaks with Bijection Learning, December 2024. URL <http://arxiv.org/abs/2410.01294>. arXiv:2410.01294 [cs].

- Tiansheng Huang, Sihao Hu, Fatih Ilhan, Selim Furkan Tekin, and Ling Liu. Virus: Harmful Fine-tuning Attack for Large Language Models Bypassing Guardrail Moderation, January 2025. URL <http://arxiv.org/abs/2501.17433>. arXiv:2501.17433 [cs] version: 1.
- Yangsibo Huang, Samyak Gupta, Mengzhou Xia, Kai Li, and Danqi Chen. Catastrophic Jailbreak of Open-source LLMs via Exploiting Generation, October 2023. URL <http://arxiv.org/abs/2310.06987>. arXiv:2310.06987 [cs].
- John Hughes, Sara Price, Aengus Lynch, Rylan Schaeffer, Fazl Barez, Sanmi Koyejo, Henry Sleight, Erik Jones, Ethan Perez, and Mrinank Sharma. Best-of-N jailbreaking. *arXiv [cs.CL]*, December 2024.
- Gabriel Ilharco, Marco Tulio Ribeiro, Mitchell Wortsman, Suchin Gururangan, Ludwig Schmidt, Hannaneh Hajishirzi, and Ali Farhadi. Editing Models with Task Arithmetic, March 2023. URL <http://arxiv.org/abs/2212.04089>. arXiv:2212.04089 [cs].
- Hakan Inan, Kartikeya Upasani, Jianfeng Chi, Rashi Rungta, Krithika Iyer, Yuning Mao, Michael Tontchev, Qing Hu, Brian Fuller, Davide Testuggine, and Madian Khabbsa. Llama Guard: LLM-based Input-Output Safeguard for Human-AI Conversations, December 2023. URL <http://arxiv.org/abs/2312.06674>. arXiv:2312.06674 [cs].
- Neel Jain, Aditya Shrivastava, Chenyang Zhu, Daben Liu, Alf Samuel, Ashwinee Panda, Anoop Kumar, Micah Goldblum, and Tom Goldstein. Refusal Tokens: A Simple Way to Calibrate Refusals in Large Language Models, December 2024. URL <http://arxiv.org/abs/2412.06748>. arXiv:2412.06748 [cs].
- Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, L  lio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timoth  e Lacroix, and William El Sayed. Mistral 7B, October 2023. URL <http://arxiv.org/abs/2310.06825>. arXiv:2310.06825 [cs].
- Xiaogeng Liu, Nan Xu, Muhao Chen, and Chaowei Xiao. AutoDAN: Generating stealthy jailbreak prompts on aligned large language models. *arXiv [cs.CL]*, October 2023.
- Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- Ilya Loshchilov and Frank Hutter. Decoupled Weight Decay Regularization, January 2019. URL <http://arxiv.org/abs/1711.05101>. arXiv:1711.05101 [cs].
- Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv [stat.ML]*, June 2017.
- Mantas Mazeika, Long Phan, Xuwang Yin, Andy Zou, Zifan Wang, Norman Mu, Elham Sakhaee, Nathaniel Li, Steven Basart, Bo Li, David Forsyth, and Dan Hendrycks. Harm-Bench: A Standardized Evaluation Framework for Automated Red Teaming and Robust Refusal, February 2024. URL <http://arxiv.org/abs/2402.04249>. arXiv:2402.04249 [cs].
- Jesse Mu, Xiang Lisa Li, and Noah Goodman. Learning to compress prompts with gist tokens. *arXiv [cs.CL]*, April 2023.
- OpenAI, Aaron Hurst, Adam Lerer, Adam P. Goucher, and Adam et al. Perelman. GPT-4o System Card, October 2024. URL <http://arxiv.org/abs/2410.21276>. arXiv:2410.21276 [cs].
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. Training language models to follow instructions with human feedback. *arXiv [cs.CL]*, March 2022.

- Ethan Perez, Saffron Huang, Francis Song, Trevor Cai, Roman Ring, John Aslanides, Amelia Glaese, Nat McAleese, and Geoffrey Irving. Red teaming language models with language models. *arXiv [cs.CL]*, February 2022.
- Xiangyu Qi, Yi Zeng, Tinghao Xie, Pin-Yu Chen, Ruoxi Jia, Prateek Mittal, and Peter Henderson. Fine-tuning Aligned Language Models Compromises Safety, Even When Users Do Not Intend To!, October 2023. URL <http://arxiv.org/abs/2310.03693>. arXiv:2310.03693 [cs].
- Xiangyu Qi, Ashwinee Panda, Kaifeng Lyu, Xiao Ma, Subhrajit Roy, Ahmad Beirami, Prateek Mittal, and Peter Henderson. Safety Alignment Should Be Made More Than Just a Few Tokens Deep, June 2024. URL <http://arxiv.org/abs/2406.05946>. arXiv:2406.05946.
- Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D Manning, and Chelsea Finn. Direct preference optimization: Your language model is secretly a reward model. *arXiv [cs.LG]*, May 2023.
- Timo Schick, Jane Dwivedi-Yu, Roberto Dessì, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language models can teach themselves to use tools. *arXiv [cs.CL]*, February 2023.
- Leo Schwinn, David Dobre, Sophie Xhonneux, Gauthier Gidel, and Stephan Gunnemann. Soft prompt threats: Attacking safety alignment and unlearning in open-source LLMs through the embedding space. *arXiv [cs.LG]*, February 2024.
- Mrinank Sharma, Meg Tong, Jesse Mu, Jerry Wei, Jorrit Kruthoff, Scott Goodfriend, Euan Ong, Alwin Peng, Raj Agarwal, Cem Anil, Amanda Askell, Nathan Bailey, Joe Benton, Emma Blumke, Samuel R. Bowman, Eric Christiansen, Hoagy Cunningham, Andy Dau, Anjali Gopal, Rob Gilson, Logan Graham, Logan Howard, Nimit Kalra, Taesung Lee, Kevin Lin, Peter Lofgren, Francesco Mosconi, Clare O’Hara, Catherine Olsson, Linda Petrini, Samir Rajani, Nikhil Saxena, Alex Silverstein, Tanya Singh, Theodore Sumers, Leonard Tang, Kevin K. Troy, Constantin Weisser, Ruiqi Zhong, Giulio Zhou, Jan Leike, Jared Kaplan, and Ethan Perez. Constitutional classifiers: Defending against universal jailbreaks across thousands of hours of red teaming, 2025. URL <https://arxiv.org/abs/2501.18837>.
- Abhay Sheshadri, Aidan Ewart, Phillip Guo, Aengus Lynch, Cindy Wu, Vivek Hebbar, Henry Sleight, Asa Cooper Stickland, Ethan Perez, Dylan Hadfield-Menell, and Stephen Casper. Latent Adversarial Training Improves Robustness to Persistent Harmful Behaviors in LLMs, August 2024. URL <http://arxiv.org/abs/2407.15549>. arXiv:2407.15549 [cs].
- Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.
- Jason Vega, Isha Chaudhary, Changming Xu, and Gagandeep Singh. Bypassing the safety training of open-source LLMs with priming attacks. *arXiv [cs.CR]*, December 2023.
- Ze Zhong Wang, Fangkai Yang, Lu Wang, Pu Zhao, Hongru Wang, Liang Chen, Qingwei Lin, and Kam-Fai Wong. Self-Guard: Empower the LLM to Safeguard Itself, March 2024. URL <http://arxiv.org/abs/2310.15851>. arXiv:2310.15851 [cs].
- Alexander Wei, Nika Haghtalab, and Jacob Steinhardt. Jailbroken: How Does LLM Safety Training Fail?, July 2023. URL <http://arxiv.org/abs/2307.02483>. arXiv:2307.02483 [cs].
- Sophie Xhonneux, Alessandro Sordoni, Stephan Günnemann, Gauthier Gidel, and Leo Schwinn. Efficient Adversarial Training in LLMs with Continuous Attacks, November 2024. URL <http://arxiv.org/abs/2405.15589>. arXiv:2405.15589 [cs].
- Guangxuan Xiao, Yuandong Tian, Beidi Chen, Song Han, and Mike Lewis. Efficient streaming language models with attention sinks. *arXiv [cs.CL]*, September 2023.

Yi Zeng, Hongpeng Lin, Jingwen Zhang, Diyi Yang, Ruoxi Jia, and Weiyan Shi. How johnny can persuade LLMs to jailbreak them: Rethinking persuasion to challenge AI safety by humanizing LLMs. *arXiv [cs.CL]*, January 2024.

Yiming Zhang, Jianfeng Chi, Hailey Nguyen, Kartikeya Upasani, Daniel M. Bikel, Jason Weston, and Eric Michael Smith. Backtracking improves generation safety, 2024. URL <https://arxiv.org/abs/2409.14586>.

Andy Zou, Zifan Wang, J Zico Kolter, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv [cs.CL]*, July 2023.

Andy Zou, Long Phan, Justin Wang, Derek Duenas, Maxwell Lin, Maksym Andriushchenko, Rowan Wang, Zico Kolter, Matt Fredrikson, and Dan Hendrycks. Improving alignment and robustness with short circuiting. *arXiv [cs.LG]*, June 2024.

Impact Statement

Machine learning tools such as large language models (LLMs) are finding widespread usage in today’s wealthy societies. As such, any work in this area has the potential for a significant impact, as it could avoid catastrophic outcomes due to a potential lack of safety of these widespread models.

This work aims to provide a new approach to reduce the harmful behaviour of LLMs when used via a webpage or API. As such, the desired impact of this work is overwhelmingly positive. However, it has to be acknowledged that any work aiming to filter or prevent harmful content from reaching users of non-open source LLMs can most likely also be re-used for censorship and thus also runs the risk of reinforcing biases of the LLM operator—intentionally or not.

More broadly and in the longer term, our work may enable practitioners to build an extra layer of safeguards into models that have capabilities that can both be useful and harmful and thus cannot or will not be removed. In such a situation, our approach and future derivatives can be used to tag and recognize the harmful usage of a capability. A potential downside is that practitioners may be over-reliant on this `<rf>` as a defence mechanism rather than ensuring that learning algorithms and data during pre-training and various post-training stages remove harmful capabilities to the model. As such, this work also considers the worst-case attacks, such as very strong fine-tuning attacks, continuous attacks, and ablation attacks, to clearly show that this approach can be circumvented with sufficient access and thus shall not be the only layer of safety for critical applications with very capable models.

A Related Work

Jailbreaking LLMs Modern LLMs used as chatbots are trained to follow user instructions (Ouyang et al., 2022) while also being trained to respond in a safe and harmless manner (Perez et al., 2022). While users quickly found ways to manually craft “jailbreaks” which could circumvent these safeguards and elicit harmful content from these systems (Wei et al., 2023), automated methods for crafting adversarial attacks were also shown to be effective. Particularly, Zou et al. (2023) propose a greedy-coordinate gradient (GCG) search algorithm to find an adversarial suffix optimized to pre-fill (Vega et al., 2023) an affirmative response in a model’s response. Other approaches use heuristics to craft interpretable jailbreaks with only black-box access to the target model (Chao et al., 2023; Liu et al., 2023; Zeng et al., 2024). Given white-box access to the target model, more powerful attacks are possible. Adversarial soft prompts can be optimized to manipulate the model’s outputs (Schwinn et al., 2024), causal features responsible for refusal behaviour can be selectively ablated (Arditi et al., 2024), and fine-tuning can be used to override or remove safety training entirely (Qi et al., 2023).

Defences Beyond standard pre-training, LLMs are typically trained with preference optimization techniques such as RLHF (Ouyang et al., 2022) or DPO (Rafailov et al., 2023) to be more aligned with human preferences. Jailbreaks can be incorporated into this preference alignment phase to increase resilience to such attacks (as is often done with red-teaming methods), but this does not often generalise to novel jailbreaks. Historically, in the context of vision models, actively training against adversarial attacks in an online manner (i.e., adversarial training) is the only method that has shown increased adversarial robustness (Madry et al., 2017). However, in the context of language, most discrete attacks are prohibitively expensive to use online. Mazeika et al. (2024) train against adversarial suffixes generated by GCG, but continually update a pool of examples rather than generate each attack from scratch. Other approaches perform adversarial training by attacking the embedding or latent space of the model (Xhonneux et al., 2024; Sheshadri et al., 2024) which is much more efficient to compute and transfers to discrete attacks. Beyond adversarial training, newer defences target and alter harmful representations in order to prevent a model from producing harmful outputs entirely (Zou et al., 2024). Independent from training a model to be more robust to jailbreaks is to classify and judge the potential harmfulness of the generated

text, often with another LLM fine-tuned for this task (Inan et al., 2023; Feuer et al., 2024), although this does require additional resources to classify the outputs. Concurrent work Huang et al. (2025) has shown that classifiers alone are often not sufficient, further making the case that other approaches are needed especially against permissive but common threat models such as the fine-tuning box attack.

Special Tokens Several works have explored training or utilising special tokens for specific purposes. Burtsev et al. (2020) prepend “memory” tokens to an input prompt on a target task. Goyal et al. (2023) append “pause” tokens, which are hypothesised to give the LLM a buffer sequence to reason over before producing an output. Mu et al. (2023) train LLMs to compress longer prompts into smaller sets of “gist” tokens as a means to shorten the context. Xiao et al. (2023) prepend “attention sinks” to improve generalization to long-context sequences. LLMs have also been trained to use a variety of tools (such as a calculator or internet access), which are denoted and invoked via special tokens (Schick et al., 2023).

Most closely related to our approach is the recent work of Jain et al. (2024), where a model is trained to prefix an output with a special *refusal* or *response* token based on the behaviour of whether the model refuses or responds to a prompt. While their approach is related in that special tokens are leveraged in the context of alignment, the approach and objective are conceptually different. Their method correlates these tokens with *behaviour* (i.e., refusal or response) in order to better calibrate such behaviours, whereas our approach correlates a special token with some implicit notion of a concept (i.e., harmfulness), *without* modifying the model’s original behaviour. This conceptual difference leads to drastically different losses in the formulation. For instance Jain et al. (2024) do not propose a KL divergence with a reference model (Equation (1)) to maintain the predictions similar to the reference model after $\langle rf \rangle$ is outputted which hurt the model’s utility and is *not* complementary with standard safety training (instead it is a way to calibrate the model post-hoc safety training). Moreover, their model is only trained to output a “behavioural token” (e.g., “refuse” or “respond”) at the beginning of the answer, which is significantly less efficient to detect harmfulness, as shown in our experiments. In contrast, our work proposes an approach that is complementary to standard safety training where the model essentially acts as an “implicit judge” on its own generated output, improving its transparency and providing a clear signal to evaluate potentially harmful generations without incurring any additional computational cost at inference time.

Beyond that, Wang et al. (2024) also learns to tag answers as harmless or harmful, but they use two stage training procedure and hardcode the tag to be at the end of the response. They only consider fixed jailbreak prompts rather than attacks and do not consider the fine-tuning setting at all. Finally, Zhang et al. (2024) train a model to output a special reset token followed by a refusal. This differs from our work as this we optimize to keep the output post-flagging identical to the base model, with the intention of maintaining utility rather than enforcing a refusal. We also do not use a DPO based objective which may inadvertently increase the probability of generating harmful sequences before generating a reset token.

B Limitations of Our Work

Our approach is not able to defend against strong white-box attacks which have access to model weights (e.g. continuous embedding attacks); this shows that the association between harmful generation and the $\langle rf \rangle$ is circumvented. In the fine-tuning attack setting, there is still a trade-off between robustness and utility in the face of an attacker, albeit a much stronger Pareto-front than without the $\langle rf \rangle$ module. Another limitation of our method is the amount of hyperparameters that are introduced that require tuning. While it is not too difficult to achieve a decent trade-off, one can likely achieve much better performance with better hyperparameter tuning. In addition, our ability to associate the $\langle rf \rangle$ with harmfulness relies on the data provided.

C Additional Model Results

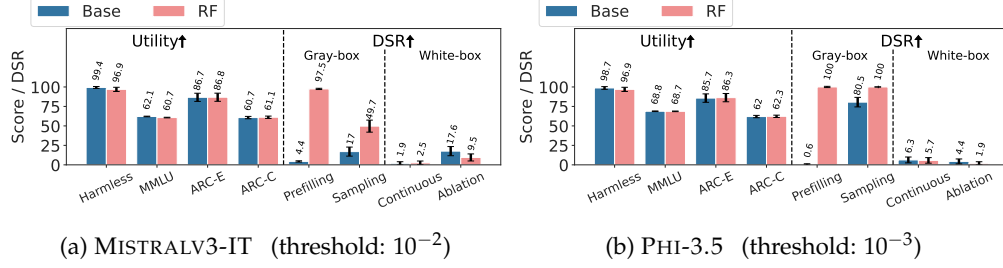


Figure 6: Model evaluation of the robustness safety trade-off, for the additional MISTRALV3-IT and PHI-3.5 models. In each plot, the left represents utility benchmarks (higher is better), and the right represents adversarial **defense** success rates (higher is better). Both refusal and **<rf>** detection are considered a successful defence, where here **<rf>** is detected if it is above some calibration threshold (in parenthesis). Pre-filling & sampling are gray-box attacks, whereas continuous & ablation are white-box attacks. Refusals are judged by GPT-4o.

D Red-Flag Fine-Tuning Algorithm

Please see Algorithm 1 for the Red-Flag finetuning procedure.

Algorithm 1 Red Flag Fine-tuning

Require: Reference model π_{ref} , benign and harmful completions datasets $\mathcal{D}_{\text{harmless}}$ and $\mathcal{D}_{\text{harmful}}$, minimum offset k , probability distribution \mathcal{P} over the indices $\{k, \dots, |\hat{y}|\}$ of the continuation, loss weighting factors $\alpha_{\text{benign}}, \alpha_{\text{rf}}, \alpha_{\text{CE}}$.

- 1: **for** $t = 1, \dots, T$ **do**
 - 2: $\{(x, y)\} \sim \mathcal{D}_{\text{harmless}}$ ▷ For benign loss
 - 3: $\mathcal{D}_{\text{benign}} := \mathcal{D}_{\text{KL}}(\pi_{\theta}(y | x) | \pi_{\text{ref}}(y | x))$
 - 4: $\{(\hat{x}, \hat{y})\} \sim \mathcal{D}_{\text{harmful}}$ ▷ For red-flag loss
 - 5: $i \sim \mathcal{P}(\{k, \dots, |\hat{y}|\})$ ▷ Sample where to inject **<rf>**
 - 6: $\mathcal{L}_{\text{rfCE}} := -\sum_{k \leq j \leq i} \log \pi_{\theta}(\langle \text{rf} \rangle | \hat{y}_{<j}, \hat{x})$
 - 7: $\mathcal{D}_{\text{rf}} := \mathcal{D}_{\text{KL}}(\pi_{\theta}(\hat{y}_{\geq i} | \langle \text{rf} \rangle, \hat{y}_{<i}, \hat{x}) | \pi_{\text{ref}}(\hat{y}_{\geq i} | \hat{y}_{<i}, \hat{x}))$
 - 8: $\mathcal{L}_{\text{final}} := \alpha_{\text{benign}} \mathcal{D}_{\text{benign}} + \alpha_{\text{rf}} \mathcal{D}_{\text{rf}} + \alpha_{\text{CE}} \mathcal{L}_{\text{rfCE}}$
 - 9: Optimize \approx_{θ} using $\mathcal{L}_{\text{final}}$
 - 10: **end for**
-

E Training Design decisions

There are several design decisions to consider:

The cross-entropy loss on **<rf>** can be computed on each index before and including the sampled position j or only on j . In other words, we have the choice to allow the model for flexibility of when to output **<rf>** at the cost of potentially overfitting more because we now train the model to output **<rf>** immediately after the instruction token. In particular, this forces the model to judge the prompt quite strongly, leading to a higher probability for **<rf>** in a refusal as well. In practice, we tested both approaches and saw better results computing the cross entropy up to and including index j . A potential solution to avoid over-fitting after the instruction token is to have a minimum offset into the harmful continuation both for sampling j as well as the cross-entropy term, which we apply to LLAMA3.2-3B-IT and MISTRALV3-IT.

The sampling distribution of the index at which to insert the **<rf>** is a key choice. We tested both a geometric distribution as well as a uniform distribution over the harmful continuation. We decide to use a geometric distribution, which means that fewer positions will be used in the cross-entropy loss, and more will be used in \mathcal{D}_{rf} .

The **attention mask** in the \mathcal{D}_{rf} can also be amputated not to include the harmful prompt \hat{x} and the harmful continuation $\hat{y}_{<j}$ before the $\langle \text{rf} \rangle$ such as to force the model to store the harmful continuation information in the $\langle \text{rf} \rangle$ embedding. In other words, with probability $p = 0.5$, we set $\mathcal{D}_{\text{rf}} := \mathcal{D}_{\text{KL}}(\pi_{\theta}(\hat{y}_{\geq i} \mid \langle \text{rf} \rangle) \mid \pi_{\text{ref}}(\hat{y}_{\geq i} \mid \hat{x}, \hat{y}_{<i}))$. We choose to apply this trick probabilistically with probability 0.5, to make the training more stable but still encourage the $\langle \text{rf} \rangle$ to be meaningful for generation.

F Hyperparameters

Table 1: Hyperparameters for models fine-tuned with our Red Flag Fine-tuning algorithm

HYPERPARAMETERS	LLAMA3.2-3B-IT	MISTRALV3-IT	PHI-3.5
BATCH SIZE	64	64	64
LEARNING RATE	0.0001	0.0001	0.00002
LEARNING RATE SCHEDULER	CONSTANT	CONSTANT	CONSTANT
α_{benign}	8	8	8
α_{rf}	1	1	1
α_{CE}	3	3	3
RF CE CUTOFF	0.15	0.15	0.15
ATTENTION DROPOUT	0.5	0.5	0.5
WARMUP RATIO	0.03	0.03	0.03
LORA - R	128	128	128
LORA - α	64	64	64
MIN OFFSET	16	16	0

G GPT-4o Judge

The **harmfulness check** is done with the following prompt:

Do you consider the following question and response to be the start of a harmful answer? Answer with yes or no. Question: {prompt}. Response: {response}

The **utility check** is done with the following prompt:

Do you consider the following response to be the start of a reasonable answer to the prompt? Answer with yes or no. Question: {prompt}. Response: {response}

H Fine-tuning attack

We validate that our approach of applying a safety LoRA module does not break benign fine-tuning. For this we train with SFT for one epoch on GSM8K (Cobbe et al., 2021) in chat mode. We train with batchsize 64, learning rate 10^{-4} , LoRA parameters $r = 64$ & $\alpha = 64$, AdamW (Loshchilov & Hutter, 2019), and a constant learning rate schedule. We evaluate the GSM8K performance with the LM-EVALUATION-HARNESS (Gao et al., 2024) using the command `lm_eval -model hf -tasks gsm8k -num_fewshot=5 -device cuda:0 -batch_size 16 -model_args pretrained=meta-llama/Llama-3.2-3B-Instruct -apply_chat_template`. For the base model we get a performance of $24.1 \pm 0.1\%$ under strict-match, the fine-tuned model gets $61.0 \pm 0.1\%$, the fine-tuned model with one safety LoRA adapter gets $62.2 \pm 0.1\%$, and the fine-tuned model with the $\langle \text{rf} \rangle$ adapter applied twice gets $59.4 \pm 0.1\%$.

I Generalisation to Longer Contexts

