A GENERATIVE APPROACH TO LLM HARMFULNESS DETECTION WITH RED FLAG TOKENS

Sophie Xhonneux* 1.2
Gauthier Gidel 1,2,4David Dobre* 1,2
Mehrnaz Mofakhami2Leo Schwinn3
Leo Schwinn3¹Université de Montréal²Mila³ Technical University of Munich⁴ Canada CIFAR AI Chair

ABSTRACT

Most safety training methods for large-language models (LLMs) based on finetuning rely on dramatically changing the output distribution of the model when faced with a harmful request, shifting it from an unsafe answer to a refusal to respond. These methods inherently compromise model capabilities and might make auto-regressive models vulnerable to attacks that make likely an initial token of affirmative response. To avoid that, we propose to expand the model's vocabulary with a special token we call *red flag token* ((rf)) and propose to fine-tune the model to generate this token at any time harmful content is generated or about to be generated. This novel safety training method effectively augments LLMs into generative classifiers of harmfulness at all times during the conversation. This method 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 rather than just the input prompt and provides a stronger defence against sampling-based attacks. In addition, it simplifies the evaluation of the model's robustness and reduces correlated failures when combined with a classifier. We further show an increased robustness to long contexts, and supervised fine-tuning 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 (Xhonneux et al., 2024; Sheshadri et al., 2024; Zou et al., 2024), perplexity filters (Alon & Kamfonas, 2023), or harmfulness classifiers (Inan et al., 2023). However, as model capabilities progress, so does their attack surface as the innate abilities can be used to circumvent defences (Huang et al., 2024)—e.g., using a low-resource language to jailbreak the model. Thus, the robustness mechanisms must keep pace with the new vulnerabilities, this necessitates leveraging these increases in capabilities. Hence, it is natural to embed another layer of security directly into the models themselves, which will scale with these capabilities. Ideally, an LLM would have no harmful capabilities; however, many skills can be both useful and cause damage depending on the context or be re-learned based on harmless abilities. Thus, adding this ability to detect harmfulness can be complementary to standard safety training whose goal is to remove harmful behaviour from the model.

To address these issues, we propose to add a special token for the LLM to predict when the model considers that its capabilities are used unsafely. We call this token a *red flag token* ($\langle rf \rangle$) and train the model to output this token at any time during the generation of a harmful response while not changing its response after the $\langle rf \rangle$, or in unrelated and safe contexts. This complementary layer of safety has several benefits. First, our method only requires the $\langle rf \rangle$ to be included in the answer—a single token change. In contrast, other methods, such as adversarial training, forces models to completely change their output distribution from a harmful response to a refusal. Thus, while standard safety training creates a tension between the probability of refusing harmful queries and the standard training objectives, such as next token prediction and instruction following (Wei et al., 2023), our safety training objective only requires a minor edit—i.e., we learn to output the $\langle rf \rangle$ in the useful, though

^{*}Equal contribution. Correspondence to Sophie Xhonneux: lpxhonneux@gmail.com

unsafe, answer the model provides. Finally, one can calibrate the strictness of flagging a generation as harmful with a bias term on the $\langle rf \rangle$ in the final softmax of the model.

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 $\langle rf \rangle$ to tag the answer as harmful and have it be filtered. This conceptually is a detection mechanism that is built into the model, i.e., we re-use the model's capabilities. This enables our method to be complementary and to be used together with other hardening methods, such as adversarial training. However, our approach differs from a classifier because the $\langle rf \rangle$ is a part of the model's generation and is re-used as an input later in the sequence. Thus, the $\langle rf \rangle$ approach is fundamentally generative rather than discriminative. This has the potential to efficiently complement the other safety mechanisms, such as standard safety training and filtering with a judge model.



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

We demonstrate the feasibility of this approach

by designing a loss function consisting of three components: (i) a cross-entropy loss on generating the $\langle rf \rangle$ (ii) a Kullback-Leibler (KL) loss term on the generation after the $\langle rf \rangle$ (see Figure 1) and (iii) a KL loss term on benign utility conversations. We test our approach with several adversarial attacks on three open-source models: LLAMA3.2-3B-IT (Grattafiori, 2024), MISTRALV3-IT (Jiang et al., 2023), and PHI-3.5 (Haider et al., 2024).

Finally, inspired by the idea of task arithmetic (Ilharco et al., 2023), we introduce the idea of storing the "safety" of the model in a LoRA module (Hu et al., 2021) and show that it can be applied to regain a significant amount of safety after an attacker has leveraged a fine-tuning API to remove the safety training of the model. Further, this concept seems to work for other safety approaches, such as adversarial training, giving practitioners many ways to defend against fine-tuning attacks.

To sum up, the contributions of this paper are:

- We propose a specialised red flag token to reliably detect harmfulness at each generation step, even under strong adversarial attacks, including pre-filling, continuous, sampling, and ablation attacks;
- We show empirically that our approach generalizes beyond the training distribution, effectively handling significantly longer inputs than those seen during training;
- Finally, we demonstrate that our safety module can be efficiently stored in a LoRA module to be applied to detect harmful outputs from fine-tuned models.

2 Method

In this section, we present how we train the model to output the $\langle rf \rangle$ during harmful continuations by only marginally affecting the model's output distribution (thus maintaining capabilities). First, we explain the threat model we consider (Section 2.1) before introducing the notation, the datasets, and the loss we will use (Section 2.2). The loss is conceptually depicted in Figure 1.

2.1 THREAT MODEL

Similar to a setting where harmfulness classifiers may be used, we assume that the LLM access is gated behind some web interface or API with no access to model weights, logits, or direct control of input/output processing—we call this the *black box* setting. The key assumption we make is that the service provider can evaluate the logits and output of the model before passing it on to

Algorithm 1 Red Flag Fine-tuning

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

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

the user, including the filtering of special tokens such as assistant tokens or our $\langle rf \rangle$ token. We further 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 box 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.

2.2 OUR LOSS

We assume that we have a dataset $(\hat{x}, \hat{y}) \sim \mathfrak{D}_{harmful}$ of harmful pairs or prompts \hat{x} and continuations \hat{y} . Further, we assume we have a dataset $(x, y) \sim \mathfrak{D}_{harmless}$ of harmless (a.k.a., benign) pairs of prompts x and harmless continuations y. Given $k \geq 0$ representing a minimum offset, an index i is sampled from a probability distribution \mathcal{P}_k over the indices $\{k, \ldots, |\hat{y}|\}$ of the continuation at which to insert the red flag token into the harmful continuation \hat{y} . We get the tokens $\hat{y}_{<i}$ which come before index i, the $\langle rf \rangle$ token, and the tokens $\hat{y}_{\geq i}$. 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 $\pi_{ref} := \pi_{\theta_0}$. Our loss consists of three components: First, to ensure our model outputs the red-flag token in harmful completions, we use a standard language modelling cross-entropy loss on all harmful completion tokens starting at the minimum offset k up to and including the $\langle rf \rangle$ token:

$$\mathcal{L}_{\text{rfCE}} \coloneqq -\sum_{k \le j \le i} \log \pi_{\theta}(\langle \texttt{rf} \rangle \mid \hat{y}_{< j}, \hat{x}).$$
(1)

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 $\langle rf \rangle$

$$\mathcal{D}_{\mathrm{rf}} \coloneqq \mathcal{D}_{\mathrm{KL}} \left(\pi_{\theta}(\hat{y}_{\geq i} \mid \langle \mathtt{rf} \rangle, \hat{y}_{< i}, \hat{x}) \mid \pi_{\mathrm{ref}}(\hat{y}_{\geq i} \mid \hat{y}_{< i}, \hat{x}) \right), \tag{2}$$

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 prompts and answers

$$\mathcal{D}_{\text{benign}} \coloneqq \mathcal{D}_{\text{KL}}(\pi_{\theta}(y \mid x) \mid \pi_{\text{ref}}(y \mid x)).$$
(3)

All these losses put together, we get:

$$\mathcal{L}_{\text{final}} \coloneqq \alpha_{\text{benign}} \mathcal{D}_{\text{benign}} + \alpha_{\text{rf}} \mathcal{D}_{\text{rf}} + \alpha_{\text{CE}} \mathcal{L}_{\text{rfCE}}.$$
(4)

Note that none of the loss functions make a harmful continuation more likely, enabling our approach to be compatible and complementary to other safety fine-tuning techniques. We summarise the training algorithm 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 using the Harmbench (Mazeika et al., 2024) training set with their refusals and 32 harmful continuations sampled for each harmful prompt using the ablation attack (Arditi et al., 2024). We use 5000 samples from the Alpaca dataset (Taori et al., 2023) as benign prompts. All models listed have a set of reserved special tokens as part of their tokenizers, allowing us to avoid extending the vocabulary and instead we re-purpose one of these unused special tokens to be the $\langle rf \rangle$.

We measure the utility of trained models using MMLU (Hendrycks et al., 2021), ARC-E and ARC-C (Chollet, 2019), which are standard LLM benchmarks as well as a dataset of benign prompts. This dataset of benign prompts consists of 119 prompts from ULTRACHAT200K (validation split) that we randomly picked (solely making sure that they were good quality), and the Harmless dataset consisting of 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 $\langle rf \rangle$) in Figure 5.

For adversarial robustness evaluation, we compute the defence success rate (DSR) of different attacks on the Harmbench test set (Mazeika et al., 2024) that contain 159 harmful prompts. Thus, we balance the dataset with the same number of harmful and benign prompts. Either a refusal or a $\langle rf \rangle$ surpassing a certain probability threshold counts as a successful defence. The thresholds are set per model and are 0.001, 0.01, and 0.001 for LLAMA3.2-3B-IT, MISTRALV3-IT, and PHI-3.5, respectively.

We train our models with a single A100-80GB GPU with LoRA (Hu et al., 2021) and a batch size of 64 using the AdamW optimiser (Loshchilov & Hutter, 2017) (for more hyper-parameters see Appendix A.2). Due to computational constraints, we did not extensively hyperparameter tune or ablate our method. We acknowledge this in Section 3.8. Thus, it is likely that further gains can be achieved by more hyperparameter tuning.

3.2 DESIGN DECISIONS

There are several design decisions to consider:

The cross-entropy loss on $\langle rf \rangle$ 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 $\langle rf \rangle$ at the cost of potentially overfitting more because we now train the model to output $\langle rf \rangle$ immediately after the instruction token. In particular, this forces the model to judge the prompt quite strongly, leading to a higher probability for $\langle rf \rangle$ 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.

The sampling distribution of the index at which to insert the $\langle rf \rangle$ is a key choice. We tested both a geometric distribution with p = 0.01 as well as a uniform distribution over the harmful continuation. Using a geometric distribution 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 $\langle rf \rangle$ such as to force the model to store the harmful continuation information in the $\langle rf \rangle$ embedding. In other words, with probability $p = 0.5 \mathcal{D}_{rf} := \mathcal{D}_{KL}(\pi_{\theta}(\hat{y}_{\geq i} \mid \langle rf \rangle, \hat{y}_{<i}, \hat{x}) \mid \pi_{ref}(\hat{y}_{\geq i} \mid \hat{y}_{<i}, \hat{x}))$, We choose to apply this trick probabilistically with probability 0.5, such as to make training less noisy but still encourage the $\langle rf \rangle$ to be meaningful for generation.

Calibrating $\langle rf \rangle$ **sensitivity** can be done in several ways. We can add a *logit bias* term to the $\langle rf \rangle$ before passing the logits through a softmax, modifying the distribution to be:

$$p(i \mid x) = \frac{\exp(Z_i)}{\sum_j \exp(Z_j)}, \quad \text{where } Z_j = \begin{cases} z_j + b & \text{if } j = \langle \texttt{rf} \rangle \\ z_j & \text{otherwise} \end{cases}$$



Figure 2: Model evaluation of the robustness safety trade-off. 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 $\langle rf \rangle$ detection are considered a successful defence, where $\langle rf \rangle$ 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-40.

where z_i is a logit for the i^{th} token and b is a user-specified bias to reweight $\langle rf \rangle$. We can also directly monitor $p(\langle rf \rangle | x)$ without adding a logit bias term or looking for explicit generations of $\langle rf \rangle$, and set a model-dependant threshold in which to flag the generation as potentially harmful. We elect to use the second strategy as it works well and is simpler to apply in practice.

3.3 **BASELINES**

We consider three baselines in this paper. The first is the original safety training of the models (the three models we consider already come with an initial safety training). The second baseline is CAT (Xhonneux et al., 2024), where we replicate their training procedure using the same datasets for a fair comparison. For the third baseline, we extend Jain et al. (2024) to our setting, whereby we insert the $\langle rf \rangle$ at the first position of the assistant's response and train with cross-entropy on the harmful continuation and benign answers. We call this baseline 'Fixed position RF'. Note that this baseline is different from what is formally proposed in Jain et al. (2024) as in their framework, the "refusal" token is placed before refusal continuations only and thus would lead to similar results as the ones of the base models which have received standard safety training (e.g., training to refuse harmful queries and accept benign ones).

3.4 ROBUSTNESS EVALUATION

We compute the defence success rates of the following attacks

3.4.1 REALISTIC ATTACK SCENARIOS

Pre-filling whereby the attacker is allowed to insert 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 in this setting, the $\langle rf \rangle$ models (including the 'Fixed position RF' model) are allowed to check the logits of the pre-filled text.

Sampling Attacks where we attack the model 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



Figure 3: Monitoring the $\langle rf \rangle$ token's log-probability and the log-probability of the top-1 token for a particular multi-turn user/assistant interaction. See Figure 6 for this example with the corresponding text.

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.

3.4.2 LIMIT TESTING

Continuous Attacks are a type of attack where soft tokens (Schwinn et al., 2024) are optimized with sign gradient descent using an affirmative response as the target loss to break the model's safety. We allow the attack to use 70 steps with 10 soft tokens and no epsilon ball constraint.

Refusal Vector Ablations are attacks in which a single "refusal direction" is identified in the model's residual stream and then ablated from activations across all layers and token positions (Arditi et al., 2024). Suppressing this direction from being represented in the model's residual stream effectively disables a model's ability to refuse harmful queries.

The first two attacks are gray-box attacks using additional features such as temperature-based sampling and pre-filling, while the latter two are very strong white-box attacks. While the continuous attack and the ablation attack are unrealistic settings for our $\langle rf \rangle$, however, we include them as a sanity-check and to test its limits.

An attack is successful if the model *does not* refuse **or** the $\langle rf \rangle$ probability *does not* surpass a model-specific threshold. We evaluate refusals using GPT-40 OpenAI et al. (2024) as a judge (see Appendix A.1 for the prompt used for the judging). Our results are shown in Figure 2.

3.5 GENERALIZATION 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 generalizes 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 3; the $\langle rf \rangle$ probabilities remain low in the regime in which it was trained (during the assistant's turn) and increase sharply in the presence of harmful content during pre-filling and generation. Aggregate statistics are shown in Figure 4.

3.6 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 through an API. We trained a LoRA module that encapsulates the weight changes for the model to insert the $\langle rf \rangle$ in harmful continuations on LLAMA3.2-3B-IT. Hence, our defence is to apply this LoRA module one or more times after the attacker has been able to fine-tune the base model. 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.

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



Figure 4: We evaluate whether model maintains the ability to predict $\langle \texttt{rf} \rangle$ after long contexts. We sample 50 chats of varying length from ULTRACHAT-200K and 50 prompts from our harmful test set. The log-probability of $\langle \texttt{rf} \rangle$ is monitored throughout the generated output of each chat/prompt combination (Gen_{base}), and we apply a pre-filling attack (Gen_{prefill}), where we pre-fill with an affirmative response. We calculate the maximum log $P(\langle \texttt{rf} \rangle)$ separately over the pre-filled text and generated text, and the average log $P(\langle \texttt{rf} \rangle)$ over the pre-filled text and the first 10 generated tokens. We observe no decay in the model's ability to predict $\langle \texttt{rf} \rangle$ with increasing context length; longer sequences can approach ~ 2000 tokens.

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 CAT and our own $\langle rf \rangle$ approach. Due to the strength of the fine-tuning attack setting for both 'Fixed position RF' and our own $\langle rf \rangle$ we consider different thresholds for the $\langle rf \rangle$ probability. This 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. As before, we use the same Harmbench test set and the same Harmless dataset from Figure 2; the results are in Figure 5.

We also validate that on benign fine-tuning our $\langle rf \rangle$ 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—see Appendix A.3 for final numbers.

3.7 DISCUSSION

The first observation from Figure 2 is that our $\langle rf \rangle$ approach maintains near-perfect utility across all models. In particular, for LLAMA3.2-3B-IT we find that all the baselines considered are able to achieve nearly the same utility scores as the base model. This allows for a good comparison point in terms of robustness utility trade-off. Our proposed $\langle rf \rangle$ approach is able to nearly perfectly defend against all the gray-box attacks, such as sampling and pre-filling across all models. We can see from the difference in defence success rate (DSR) between the base model and the $\langle rf \rangle$ model that even when requests are not being refused, the $\langle rf \rangle$ is being generated and thus defends against the malicious attack. The white-box ablation attack also does not succeed in preventing the $\langle rf \rangle$ from triggering on LLAMA3.2-3B-IT. However, on both MISTRALV3-IT and PHI-3.5 the attack succeeds, which may be because the attack also has many hyperparameters or the base models have different safety properties. The continuous attack on LLAMA3.2-3B-IT is challenging to interpret due to the difficulty of getting coherent text after the attack. On MISTRALV3-IT, however, we can see that the continuous attack is effective at inciting coherent responses. These attacks are difficult to defend against due to its ability to change the token embeddings. This shows the limitation of our approach. Note we did not try discrete white-box like GCG (Zou et al., 2023) due to them being unable to break the base model such as LLAMA3.2-3B-IT. We find that CAT on LLAMA3.2-3B-IT performs significantly worse on both pre-filling and the ablation attacks showing that it must sacrifice safety in order to maintain the same utility. The 'Fixed position RF' is only marginally worse on the continuous attacks and otherwise performs equally well. Note that continuous and refusal ablation attacks are extremely strong threat models and serve as an empirical "worst-case" attack, which we use as sanity checks to understand the limitations of our approach. Finally, since these attacks operate in the smooth representation or embedding space (rather than discrete token space), an interesting open question remains on how the feature representations change under such attacks, and how they interact with $\langle rf \rangle$.

We next consider fine-tuning attacks, which are less explored as a defence setting but have seen some recent progress. Qi et al. (2024) consider a similar setting and restrict the fine-tuning API to allow modification on the initial token distribution. Tamirisa et al. (2024) adversarially trains a model by adversarially perturbing the model weights during the attack step. We consider storing our safety training in a LoRA module and applying it after the harmful fine-tuning attack. In Figure 5, the limitations of the fixed position $\langle rf \rangle$ approach become clear; the Pareto front of this approach is significantly worse and applying the LoRA module multiple times does not help, while the variable position $\langle rf \rangle$ LoRA module maintains its ability to predict $\langle rf \rangle$ in harmful contexts. The adversarial training baseline (CAT) is also able to provide good robustness though it cannot be calibrated and thus does not produce a Pareto front. However, by applying the CAT LoRA module once or twice from this baseline we get a different robustness trade-off. Contrary to adversarial training, our LoRA module, together with a carefully chosen threshold,



Figure 5: ROC curve for different max probability thresholds to defend against a *fine-tuning attack* against LLAMA3.2-3B-IT. Baseline models are a CAT (Xhonneux et al., 2024) and a $\langle rf \rangle$ 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 $\langle rf \rangle$ module

can provide fine-grained calibration and a good Pareto front. Finally, we show that the adversarial training module and our $\langle rf \rangle$ module are complimentary and allow for good robustness and a calibrated trade-off between utility and robustness in terms of TPR and FPR.

3.8 LIMITATIONS

Expectedly, our approach is not able to defend against attacks that have complete access to the model weights such as continuous attacks, although substantial robustness against ablation attacks is achieved. This shows that the association between harmful generation and the $\langle rf \rangle$ is circumvented by these attacks, demonstrating a limitation of our approach. Furthermore, 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 as we did not have the resources to effectively tune the model (all the trainings were done with a single GPU). In addition, our ability to associate the $\langle rf \rangle$ with harmfulness relies on the data provided.

4 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 safe-guards 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 generalize 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 need especially against permissive but common threat models such as the fine-tuning box attack.

Special Tokens Several works have explored training or utilizing 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 hypothesized 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 2 to maintain the predictions similar to the reference model after $\langle rf \rangle$ is output 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.

5 CONCLUSION

We propose to detect 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 continuation. This provides us with a generative approach to detect harmfulness even under strong adversarial attacks such as pre-filling and sampling. We show that our method significantly improve robustness without affecting utility. We demonstrate that our approach generalizes to very long contexts with multiple conversation turns despite having only been trained on short one-round conversations.

Finally, we investigate another strong threat model that of a fine-tuning attack in an API setting and show that you can store safety training approaches such as our own $\langle rf \rangle$ or adversarial training (CAT (Xhonneux et al., 2024)) in a LoRA module and apply it post-hoc against a jailbroken model to regain significant robustness without harming benign fine-tuning performance. These different model hardening approaches are complimentary, as we show by combining a continuous adversarial training (CAT) module and our own $\langle rf \rangle$ module.

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, Yuanzhi 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 Finetuning 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 Khabsa. 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, Alfy 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, Lélio Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée 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. HarmBench: 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.
- 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].
- Rishub Tamirisa, Bhrugu Bharathi, Long Phan, Andy Zhou, Alice Gatti, Tarun Suresh, Maxwell Lin, Justin Wang, Rowan Wang, Ron Arel, Andy Zou, Dawn Song, Bo Li, Dan Hendrycks, and Mantas Mazeika. Tamper-resistant safeguards for open-weight LLMs. *arXiv* [cs.LG], August 2024.
- 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.
- 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.
- 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.

A APPENDIX

We use huggingface implementations and pytorch to train and run all of our models. All of our experiments were each done on single A100-80GB GPUs.

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 $\langle rf \rangle$ 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.1 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}

A.2 HYPERPARAMETERS

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

Hyperparameters	LLAMA3.2-3B-IT	MISTRALV3-IT	Рні-3.5
BATCH SIZE	64	64	64
LEARNING RATE	0.0001	0.0001	0.00002
LEARNING RATE SCHEDULER	CONSTANT	CONSTANT	CONSTANT
$lpha_{ m benign}$	8	8	8
$lpha_{ m rf}$	1	1	1
$lpha_{ m CE}$	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

A.3 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 (rf) adapter applied twice gets $59.4 \pm 0.1\%$.

A.4 LONG-CONTEXT



Figure 6: Multi-turn long-context log probabilities of $\langle rf \rangle$ and the top-1 probability.