

gradients [21], [22], graybox attacks utilize logprobs from the inference endpoints [23], [24]; and blackbox attacks rely only on the textual output from the model [19], [25]. LLM vendors have recently mitigated some of the graybox attack vectors by either not returning logprobs as part of an inference call, or returning only a small subset of logprobs [26], [27], [28]. In both cases, there is not sufficient information to effectively guide the search for an adversarial prompt. Our work shows that existing greedy search algorithms can be adapted to use loss-like information from the fine-tuning interface to create attacks.

In our *fun-tuning* attacks, the attacker runs a single iteration of fine-tuning with a very small learning rate. The small learning rate ensures that the base model does not change significantly, subverting the point of fine-tuning, which is to update the model’s weights. The output of this operation is a loss-like value representing how far the model’s true output is from the desired output. We use this loss-like value to guide our search procedure. There are two technical challenges in using the fine-tuning functionality this way.

The first challenge is establishing that the fine-tuning loss is a good signal to guide the search for adversarial prompts. The fine-tuning interfaces of commercial vendors are closed source and the documentation does not provide details. Using a range of experiments, we partially reverse-engineer the loss function used by Google’s Gemini and empirically establish that the reported training loss is a noisy signal that can be useful for discrete optimization of prompts.

The second challenge is that fine-tuning interfaces randomize the order of items in the training and validation sets. The losses that get reported are a permutation of the input training data order. This is problematic because the attacker has to systematically evaluate the effect of different token substitutions when creating the adversarial prompt. We address this challenge using an approximate de-randomization procedure that uses a carefully crafted training dataset to approximate the permuted training order. We validate the utility of our approximate procedure by comparing it with another inefficient, but provably-correct, method and find that our procedure approximates the true permutation within reasonable bounds.

We focus on creating prompt injection attacks because they are the most realistic security problem faced by systems built using large models. These systems, or agents, can use tools to solve a variety of tasks such as automatically finding bugs in software and managing calendars and emails. A prompt injection attack in the context of an agent-based system can affect the confidentiality and integrity of user data. Consider a simple agent that has access to a user’s email. An attacker could craft an adversarial prompt and send an unsolicited email to a victim user. If the user asks their email-handling agent to “summarize my latest email,” the adversarial prompt will hijack the agent and cause it to also send out the user’s emails to the attacker, thus violating confidentiality [13], [8].

Mitigating the fun-tuning attack vector is non-trivial because it exploits the utility-security trade-off. Fine-grained control over training hyperparameters (learning rate, batch

size, epochs, randomization seed for training set shuffling) is crucial to developers who are using the interface in benign ways, but as we show, it is also helpful to attackers. Therefore, any changes that reduce control over hyperparameters (e.g., setting a large minimum value on learning rate) can negatively affect benign developers and the utility of the fine-tuning interface itself. Scanning the training set for the presence of malicious data before running the fine-tuning job is a potential mitigation (e.g., to look for content policy violations), but can be evaded using encoding techniques that conceal the semantics of the data [29] and is not always possible for prompt injection attacks.

Contributions.

- **New Attack Surface Characterization.** We surface and experimentally characterize a new attack vector on LLM systems that exploits a fundamental tension between utility and security. Specifically, we show how an attacker can misuse the fine-tuning interface to create adversarial prompts. We call these *fun-tuning* attacks. We solve two technical challenges along the way: (a) establishing that the fine-tuning loss is suitable for discrete optimization of prompt injections and (b) derandomizing the reported loss-like values to obtain a usable signal for adversarial prompt optimization.
- **Experimental Analysis of Optimization-based Prompt Injections.** We experimentally evaluate the vulnerability of Google’s Gemini model series to fun-tuning attacks. We show that existing discrete optimization algorithms for prompt injections can be modified to incorporate fine-tuning losses for guidance. Using the popular prompt injection benchmark PurpleLlama [30], we show attack success rates between 65% and 82% for Google’s Gemini family of models. We also find that the attacks transfer between various Gemini models with relatively high success rates.

Disclosure and Ethics. We disclosed the issue to Google on November 18, 2024. Google deployed the following mitigation in early April 2025: “We constrained the API parameters that they were relying on. In particular, capping the learning rate to a value that would rule out small perturbations and limiting the batch size to a minimum of 4, such that they can no longer correlate the reported loss values to the individual inputs.” Our goal with this work is to raise awareness and begin a conversation around the security of fine-tuning interfaces and their role in helping create prompt injection attacks on closed-weight models. We conducted all experiments using the standard developer fine-tuning interface. Although we created attacks, the fine-tuning job and its impact on the provider are indistinguishable from benign fine-tuning jobs. We did not deploy any of these attacks in the wild. Code is available at <https://github.com/earlence-security/fun-tuning>.

2. Background

We model an LLM \mathcal{M}_Θ as a probability distribution over the next token conditioned on its input tokens, that

is, for an input sequence of tokens $x_{1:n} = (x_1, x_2, \dots, x_n)$, \mathcal{M}_Θ outputs a probability distribution $P(y|x_{1:n}; \Theta)$, where Θ denotes the model parameters, y represents the next token to be generated. Here, all tokens come from a discrete set $V = \{T_1, T_2, \dots, T_{|V|}\}$ (called the vocabulary of the LLM).

Using LLMs in Generative/Inference Mode. The next output token y_1 is generated by sampling from the probability distribution $y_1 \sim P(y|x_{1:n}; \Theta)$ according to some sampling procedure. In practice, LLMs generate a sequence of tokens until a special “end-of-text” token is generated. Therefore, on input $x_{1:n} \in V^*$, an LLM will return an output sequence $y_{1:m} \in V^*$ with probability

$$P(y_{1:m}|x_{1:n}; \Theta) = \prod_{i=1}^m P(y_i|x_{1:n}, y_{1:i-1}; \Theta) \quad (1)$$

In addition to the generated stream of tokens, LLM APIs return a variety of additional information related to the probability of generation of the output in the form of “logprobs”. Logprobs of a token y represent the probability of generation of that token y given the input $x = x_{1:n}$. That is for a sequence of tokens $y = y_{1:m}$, we have

$$\text{Logprobs}(y|x; \Theta) = - \sum_{i=1}^m \log P(y_i|x, y_{1:i-1}; \Theta) \quad (2)$$

Some APIs (e.g., Cohere, GooseAI) return the complete logprobs vector, *i.e.*, the logprobs for each token in the vocabulary. Some APIs return only the Top-N logprobs (e.g., OpenAI returns top-20), and others do not return any logprobs information at all. A longer list can be found in Table 6 of [31].

Training. Generative LLMs are trained by finding parameters Θ which minimize the Cross-entropy loss between the predicted probability distribution over the next token $\mathcal{M}_\Theta(x^i) = P(y|x^i)$ and the true next token y^i summed over the training dataset I , where x^i denotes i 'th training example of I .

$$\hat{\Theta} = \underset{\Theta}{\operatorname{argmin}} \sum_{i \in I} \text{CrossEntropy}(\mathcal{M}_\Theta(x^i), y^i) \quad (3)$$

Expanding the definition of Cross-entropy loss, we get

$$\hat{\Theta} = \underset{\Theta}{\operatorname{argmin}} \left(- \sum_{i \in I} \log P(y^i|x^i; \Theta) \right) \quad (4)$$

Prompt Injections. In real world systems, the input to an LLM usually consists of system and user prompts, separated by special tokens and concatenated into a single token sequence, often called a conversation. For simplicity, we assume an LLM prompt is a pair consisting of a system prompt x_{Sys} and a user prompt x_{User} . Crucially, the system prompt and user prompt come from sources with different security contexts and assume a trust boundary between them: one of these values might be controlled adversarially.

Prompt Injection attacks assume partial control of an LLM input and use it to achieve objectives not originally intended by the developer or the actual user. Prompt Injections can lead to several real-world consequences such as leakage of private data and tool misuse [8], [9], [10], [11],

[13]. This compromises the confidentiality and integrity of user resources connected to the LLM.

A prompt injection attack occurs when a trusted input (such as the system prompt of an LLM agent and the user prompt), concatenated with an untrusted malicious input, (such as text retrieved from a third-party webpage to be summarized), causes the LLM to deviate from its original task and follow the instructions of the malicious input. Formally, we define prompt injection when the following holds:

$$LLM(x_{\text{Trusted}} \| x_{\text{Adv}}) \approx LLM(x_{\text{Adv}}) \approx y_{\text{Target}} \quad (5)$$

The task of an attacker is to craft an adversarial input x_{Adv} such that, if combined with x_{Trusted} , the LLM outputs a sequence of tokens y_{Target} of the attacker’s choice instead of the intended output y_{True} . The definition might be extended to assume $LLM(x_{\text{Adv}})$ is still conditioned on x_{Trusted} (but no longer instructed by it) to cover the cases where x_{Adv} instructs LLM to reveal x_{Trusted} , but we leave it out of scope for simplicity.

We also note that, in general, the adversarial input could be combined with the trusted input in several different ways - such as concatenation of strings, parameterized queries wrapping each prompt with special characters, interspersed with each other or via further pre-processing [32]. However, for simplicity and without loss of generality, we assume that the trusted and untrusted sequences are combined using plain concatenation denoted by $\|$. Under this setting, the adversary needs to find x_{Adv} such that

$$LLM(x_{\text{Trusted}} \| x_{\text{Adv}}) = y_{\text{Target}}$$

Most early prompt injection attacks were crafted manually and exploited various model-specific quirks [10], [8], [13]. These early attacks benefit from linguistic approaches and automated rephrasing methods, but here we formalize the problem mathematically.

The problem of finding such an x_{Adv} can be formulated as an optimization problem where the objective is to find an x_{Adv} which maximizes the probability of the LLM outputting the string y_{Target} , *i.e.*,

$$x_{\text{Adv}} = \underset{x}{\operatorname{argmax}} P(y_{\text{Target}}|(x_{\text{Trusted}} \| x); \Theta) \quad (6)$$

This is equivalent to finding an input x_{Adv} that minimizes the cross-entropy loss over the target string:

$$x_{\text{Adv}} = \underset{x}{\operatorname{argmin}} \text{CrossEntropy}(\mathcal{M}_\Theta(x_{\text{Trusted}} \| x), y_{\text{Target}})$$

or equivalently

$$x_{\text{Adv}} = \underset{x}{\operatorname{argmin}} - \text{Logprobs}(y_{\text{Target}}|x_{\text{Trusted}} \| x; \Theta) \quad (7)$$

The above discrete optimization problem was addressed in the white-box setting using the Greedy Coordinate Gradient algorithm which relies on being able to compute gradients that guide a search for an adversarial input [21]. In the graybox and blackbox setting, where gradients are not available, prior optimization based approaches rely either on transferability or on being able to compute the logprobs for a target token of the attacker’s choice [24], [23].

LLM Fine-Tuning Functionality. Fine-tuning is a procedure that allows users to further train a pre-trained base model on custom or proprietary data. Fine-tuning of an LLM can help with increasing model accuracy and reducing hallucinations on domain-specific knowledge, and allow for simplified prompts [33], [34]. Several AI companies such as OpenAI, Google, and Amazon provide the fine-tuning functionality as a service for users [14], [15], [16].

To use the fine-tuning functionality, users need to prepare and format a high-quality dataset which is reflective of their use case and fine-tune a base model on that dataset with an appropriate fine-tuning configuration. The fine-tuned model is then made available for inference for the custom use-case.

For example, we provide a high level description of the LLM fine-tuning API provided by Google AI Studio (“Gemini API”). The Gemini Fine-Tuning API accepts a base model and a training dataset in the form of a list of input-output pairs of strings. In addition, it lets users specify the number of epochs for the fine-tuning, the batch size, the learning rate. The Fine-Tuning API responds with the training loss for every iteration (step) of the optimization process. Additionally, the API endpoint shuffles the dataset examples before training on them, so that the resulting losses come in a pseudorandom order.

Fine-tuning interfaces from different companies provide a similar level of control over the training hyperparameters (see [Section 7](#) for more details).

3. Threat Model and Attack Constraints

The attacker’s goal is to create prompt injection attacks on a target LLM. They are a third party who wish to take control of an existing conversation, and then force the LLM to follow a different set of instructions. The attacker can deliver the prompt injection using a variety of ways, such as poisoning a webpage that the user might want summarized, sending an unsolicited email to the user’s LLM-based agent, or modifying a code repository that the user might be analyzing [8], [13], [30].

Successful prompt injection attacks must obey two properties. First, the domain of the LLM-based agent (e.g., code editing agent, email/calendar handler agent, web browsing agent) imposes constraints on the syntax and size (in number of tokens) of the attack prompt. For example, for an LLM-based coding assistant, the attack must exist as a valid comment and cannot break the syntax of the programming language. Similarly, if the user has delegated calendar management to the LLM, then the attack prompt must be delivered as a valid calendar event, limiting its size. Thus, for the generality of the prompt injection attacks that we create in this paper, we impose a size limit on the number of tokens that the attacker can inject. Additionally, for code examples, we don’t inject newline characters so that the adversarial prompt stays within a single commented out line. The attack size varies based on the application domain — we conduct our experiments with 97% of all attack prompts sizes being below 100 tokens (all of them are shorter than 500 characters). Out of this, 40 tokens are the output of the optimization

algorithm and the remaining tokens represent the attacker’s instruction in natural language. The optimizer-controllable tokens are a configurable parameter of our algorithm. Thus, the final attack takes the form of a sequence of optimizer-created tokens sandwiching the attacker’s natural language instruction ([Fig. 1](#)).

Second, the attack should be stealthy and must cause the target LLM to only produce the attacker’s desired output without anything else. This is a stronger requirement compared to prior work in prompt injections [8], [30], but is one we believe to be an important constraint. Recent LLM products have been fine-tuned specifically to resist prompt injection attempts [35], which might result in a blocked request or might steer the model behavior towards returning a more truthful answer. In that context, we observed that Gemini models often inform the user that its behavior is potentially overridden by a specific prompt when seeing some unusual requests. For example, Gemini in our experiments sometimes responded “*The code is designed to calculate the area of a circle with a radius of 5. However, the code has a comment that explicitly overrides the function’s calculation and instructs the code to output ‘10’.*” The attacker will not want the user to be aware that there is a prompt injection and thus we require our attacks to not create such warnings.

We assume that the attacker has access to the fine-tuning interface for a target LLM. This type of access is graybox because the fine-tuning interface returns a loss signal. Most LLM vendors allow anyone to sign up and become a developer.

4. Experimental Analysis of the Gemini Fine-Tuning Interface

The problem of generating adversarial inputs that force an LLM \mathcal{M}_Θ to output a string y_{Target} can be phrased as an optimization problem minimizing the unweighted cross-entropy loss over a given target string ([Eq. \(7\)](#)). In absence of logprobs information, we can’t directly compute the cross-entropy loss for a desired input and target string. To address this problem, we rely on the training losses reported by the fine-tuning API as a proxy loss function.

The key insight behind our approach is that for a small learning rate which is close enough to 0, the parameters of a model should stay nearly constant. Therefore, if we send a fine-tuning request with a single training example with a single pair of input-output strings ($x_{\text{Trusted}} \| x, y_{\text{Target}}$) and fine-tune it for a single epoch with a very small learning rate, the loss metric reported should leak information about the cross-entropy loss of the target model. Thus, we can instead try to solve the following proxy optimization problem:

$$x_{\text{Adv}} = \underset{x}{\operatorname{argmin}} \operatorname{TrainingLoss}(\mathcal{M}_\Theta(x_{\text{Trusted}} \| x), y_{\text{Target}})$$

In this section, we focus on the Fine-Tuning API for their Gemini class of models and empirically establish that the training loss is a useful proxy for optimization. We do this in two steps:

- We probe the API to understand and validate our hypothesis that a small learning rate does not significantly affect the base model. This is important because even though the learning rate is an externally controllable hyperparameter, some fine-tuning interfaces treat these as *multipliers* to an internal, hidden learning rate.
- We partially reverse-engineer the loss values reported by the API and establish that they are indeed a good proxy for discrete optimization.

4.1. Fine-tuning Hyperparameter Analysis

The Gemini Fine-tuning API accepts as input a training set, consisting of a list of pairs of input and output strings. Additionally, the API also allows users to control some hyperparameters for the fine-tuning process, specifically, the number of epochs, the batch size, and the learning rate. We probe the API and validate the following behavior of the Gemini Fine-tuning procedures.

Small learning rates do not change training loss values significantly. While the API only accepts values larger than roughly 10^{-45} as a valid learning rate, we find experimentally that for learning rate values between roughly 10^{-13} and 10^{-45} , the loss values reported by the API (when controlled for the training set and other hyperparameters) stay constant (up to the precision of the numbers reported by the API) and are independent of the learning rate itself. Concretely, we take a fixed training dataset D consisting of n distinct training examples and send it to the fine-tuning API with a learning rate $\alpha_0 \approx 10^{-45}$ and observe the losses returned, say $L^{(\alpha_0)} = \{l_1^{(\alpha_0)}, l_2^{(\alpha_0)}, \dots, l_n^{(\alpha_0)}\}$. We then send D again to the fine-tuning API, this time with a learning rate $\alpha_1 \approx 10^{-44}$ and collect the losses $L^{(\alpha_1)} = \{l_1^{(\alpha_1)}, l_2^{(\alpha_1)}, \dots, l_n^{(\alpha_1)}\}$. Similarly, we collect these sets of learning rates for N different learning rates between 10^{-45} and 10^{-13} . That is, we get N sets of values $L^{(\alpha_i)} = \{l_1^{(\alpha_i)}, l_2^{(\alpha_i)}, \dots, l_n^{(\alpha_i)}\}$. We observe that for $\alpha_i \neq \alpha_j$ and $\alpha_i, \alpha_j < 10^{-13}$, $l_1^{(\alpha_i)} = l_1^{(\alpha_j)}$, $l_2^{(\alpha_i)} = l_2^{(\alpha_j)}$ and so on. That is, we get that the set of values $L^{(\alpha_i)} = L^{(\alpha_j)}$. This implies that the loss values reported by the fine-tuning API do not change significantly when the learning rates are smaller than $\approx 10^{-13}$.

This supports our hypothesis that small learning rates do not change the model parameters significantly. Furthermore, any value in this range can serve as a small learning rate since they all guarantee that the loss values being reported are not being affected significantly by the training of the model itself.

The training losses reported are permuted. To establish this, we notice that if we send a fine-tuning request with an ordered training set consisting of duplicated data, say, $D = \{(x^{(1)}, y), (x^{(2)}, y), (x^{(2)}, y), (x^{(3)}, y), (x^{(3)}, y), (x^{(3)}, y)\}$ of size 6 with a batch size of 1 for 1 epoch with a small learning rate, we should expect to get an ordered list of losses $L = \{l^{(1)}, l^{(2)}, l^{(2)}, l^{(3)}, l^{(3)}, l^{(3)}\}$ in response. Instead, we find that the losses reported are permuted according to some permutation. That is, we do observe the same cardinalities of different losses as expected, *i.e.*, out

of the 6 loss values reported, one value appears exactly once, one value appear twice, and one value appears thrice, but in a different order from the input. Note that counting the cardinalities of duplicated training items is the only way to determine whether a permutation occurred or not because we do not know the loss values of each individual training example ahead of time.

Similar observations hold across different training set sizes and different sets of training examples. Therefore, we can conclude that for an ordered training set $D = \{(x^{(i)}, y^{(i)})\}_{i=1}^N$ of size N , there is a permutation σ_N such that the true training losses corresponding to each example in the dataset D are obtained by applying the permutation σ_N to the reported ordered set of losses $L = [l^{(i)}]_{i=1}^N$.

The same permutation is applied across different fine-tuning requests. We also observe that making different fine-tuning requests with the same training set sizes (but different training data) results in the same orderings of reported losses. Therefore, we conclude that the fine-tuning procedure uses a constant, hardcoded seed value s to initialize a generator and applies the same, fixed permutation to a training set of a given size N (as long as the batch size is 1). We refer to this permutation using σ_N .


4.2. Reverse Engineering the Training Loss

Gemini documentation doesn't provide details on how the training loss is computed and optimized during fine-tuning. Different approaches to fine-tuning can optimize and report different types of loss functions, such as:

- Cross-Entropy computed over only the output string (“Instruction Tuning”)[36].
- Cross-Entropy computed over both the input and output strings (“Instruction Modeling”)[37].
- Some unknown custom loss functions such as Distillation losses or Sparse Training losses [38].

In this subsection, we partially reverse engineer the reported loss to analyze its effectiveness as a proxy for the adversarial objective *i.e.*, the average logprobs. We do this by comparing the Fine-Tuning loss, which we denote as TrainingLoss, with logprobs, denoted later as AvgLogprobs. Natively, the Gemini API doesn't expose the logprobs of any string, however, Vertex AI, an enterprise API by Google, does provide an “average logprobs” value for top 8 responses. For our analysis in this section, we assume that both the Gemini API and the Vertex AI API serve the same base model. We treat the average logprobs as the ground truth and assume that it is the logprobs of the generated response (see Eq. (2)) divided by the length of the output.

Next, we pick three distinct input prompts and compare the TrainingLoss with the AvgLogprobs for each prompt. Concretely, for a prompt X , we collect the average logprobs of the generated response $Y_{1:l}$ (of length l) and collect the training loss for the input-output pair $(X, Y_{1:l})$. We plot TrainingLoss and $l \cdot \text{AvgLogprobs}$ (average logprobs scaled to the output length, which we call total logprobs) as functions of the output length l as we incrementally increase the output length l in Fig. 2. We make a few observations:



Q1: Write a long essay on the benefits... A: The Profound Benefits of Exercise...
 Q2: Tell me about elves and fairies. A: Elves are mythical beings...
 Q3: Explain in great detail... A. Transistors are fundamental building blocks...

Figure 2: Total logprobs, training loss, and output length are all pairwise proportional. The difference between total logprobs and training losses for a fixed input-output pair is independent of the output length.

- 1) Both values increase proportionately to the length l
- 2) Training loss increases proportionally to the total logprobs: their difference stays nearly constant and doesn't depend on the output length
- 3) This difference varies as we change the inputs in the three plots.

The above let us conjecture that a linear relationship exists between training loss and average logprobs, where the training loss is a function of both inputs and outputs:

Hypothesis. We hypothesize the following closed-form expression for the FL:

$$\text{TrainingLoss}(Y|X) = K(X) + l \cdot \text{AvgLogprobs}(Y|X)$$

where the term $K(X)$ is a function of input X .

Validation Experiment. First we note that for a dataset of l -length output sequences, $D^{(l)} = \{(X_i, Y_{i,1:l})\}_{i=1}^n$, we have the total logprob values $\mathcal{T}_0^{(l)} = \{l \cdot \text{AvgLogprobs}(Y_{i,1:l}|X_i)\}_{i=1}^n$, and the corresponding hypothesized training loss values $\mathcal{T}_1^{(l)} = \{K(X_i) + l \cdot \text{AvgLogprobs}(Y_{i,1:l}|X_i)\}_{i=1}^n$. The formula for the R-squared value for the two sets of numbers ($\mathcal{T}_0^{(l)}$ and $\mathcal{T}_1^{(l)}$) is given by

$$R^{2(l)} = 1 - \frac{\sum_{i=1}^n (K(X_i))^2}{l^2 \cdot \text{Var}(\{\text{AvgLogprobs}(Y_{i,1:l}|X_i)\}_{i=1}^n)}$$

Thus, if our hypothesized form is correct, then for large values of l , the R-squared value should approach 1.

Empirical evidence. We conduct the above experiment empirically, we create a dataset of $n = 10$ open-ended questions $\{X_i\}_{i=1}^{10}$. For a fixed output length of l tokens, we collect the average logprobs $\mathcal{A}_0^{(l)} = \{\text{AvgLogprobs}(X_i, Y_{i,1:l})\}_{i=1}^{10}$ of each of the generated responses $\{Y_{i,1:l}\}_{i=1}^{10}$ when truncated to the length l . We then send the training dataset $\{X_i, Y_{i,1:l}\}_{i=1}^{10}$ for fine-tuning and collect the training losses

$$\mathcal{A}_1^{(l)} = \{\text{TrainingLoss}(X_i, Y_{i,1:l})\}_{i=1}^{10}$$




Figure 3: The correlation between average logprobs and training losses asymptotically approaches 1 as the length of the output string increases.

We perform a linear regression over the 10 data points in $\mathcal{A}_0^{(l)}$ and $\mathcal{A}_1^{(l)}$ and record the goodness-of-fit coefficient $R^{2(l)}$. Finally, we plot $R^{2(l)}$ as a function of l . The results are as shown in Fig. 3. As can be seen from the graph, we see an almost perfect correlation for large values of l , thus validating our hypothesis.


It is unclear how the term $K(X)$ is computed. One explanation might be that the training loss performs Instruction Modeling [36], where the training loss is computed over both input and output and possibly includes unknown terms or additional, unknown tokens in its internal representation of a training example. However, we note that we do not need to know exactly how the term $K(X)$ is calculated since we can establish that training loss is a useful proxy even without knowing the exact form of $K(X)$.

4.3. Training loss is a useful proxy for optimization

The empirical data (Fig. 3) shows that the training loss is almost perfectly correlated with the average logprobs when the length of the target string is long. Therefore the training loss serves as an almost perfect proxy for the adversarial objective function when the length of the target string is long. In this subsection, we empirically establish that even for short target strings, the training loss acts as a usable proxy.

For iterative solutions to the optimization problem, training loss can be a good proxy if it can guide the search process toward the “correct” direction *i.e.*, in the direction of minimizing the “true” loss (the logprobs). In other words, it should help identify the best small perturbation (the one leading to minimum true loss) from a set of candidate small perturbations. Empirically, we find that while the training loss for Gemini doesn't always identify the best perturbation from a set of perturbations, the selected perturbation is better than the average (*i.e.*, a randomly sampled perturbation).

To understand this, we examine how good is the “true” performance or rank (according to logprobs) of the candidate selected by the training loss. We consider a question X to which the LLM replies with a short, deterministic, well-known answer Y . We then create a dataset of N small perturbations of X by appending a randomly generated token at the end of X , while ensuring that



Q1: What breed is the dog Scooby Doo? A: Great Dane
 Q2: Who was the first president of the United States of America? A: George Washington
 Q3: What is the name of the fictional spy who goes by the codename 007? A: James Bond

Figure 4: Rank distribution of top candidate from training losses, with $M = 100$ samples each for $N = 10$ candidates

the output of the LLM on all these perturbations continues to be Y . That is, $D = \{(X \| r^{(i)}, Y)\}_{i=1}^N$. We then collect the average logprobs $\{\text{AvgLogprobs}_i\}_{i=1}^N$ corresponding to the inputs $\{X \| r^{(i)}\}_{i=1}^N$ and the training losses $\{\text{TrainingLoss}_i\}_{i=1}^N$ for this dataset D . We then compute the rank of the best candidate selected by the training loss (i.e., $\text{argmin}_{i \in \{1 \dots N\}} \text{TrainingLoss}_i$) in the list of perturbations sorted in the increasing order of the true loss (logprobs). Fig. 4 shows the distribution of the rank for 3 different questions.

We record the value j and repeat the experiment M times, each with fresh randomly generated candidates, to obtain a list of values of j all lying in $\{1 \dots N\}$. We plot the frequencies of j to obtain a distribution. Three such frequency charts are shown in Fig. 4 for three questions and answers.

We see that all distributions are skewed highly to the left indicating that with high probability, the best candidate minimizing the training losses is also amongst the top few candidates in the average logprobs. This analysis confirms that the training loss value can serve as a noisy signal to guide the discrete optimization process.

5. Adversarial Prompt Optimization using the Fine-Tuning interface

We have empirically established that fine-tuning loss can act as a good proxy for guiding the adversarial optimization process. In this section, we use it to automatically generate prompt injections.

5.1. Recovering the random permutation

In principle, we can obtain the training loss corresponding to a single training example $(x^{(i)}, y^{(i)})$ atomically by sending a fine-tuning request with a single training example with a batch size of 1 for 1 epoch and obtain a single value. Unfortunately, due to a spin-up overhead from a few seconds to several minutes per fine-tuning request, it is desirable to evaluate the training loss for multiple training examples in one query. However, the permutation applied to the losses

Table 1: Example of building $D_{garbled}$ by progressively corrupting y_{Target} on the same $x_{\text{Prompt}} = \text{“Repeat this: quick brown fox.”}$

	$y_{\text{Target},i}$	TrainingLoss
$y_{\text{Target},0}$	Quick brown fox	16.08243
$y_{\text{Target},1}$	Ham brown fox	41.72246
$y_{\text{Target},2}$	Ham oss fox	57.49492
$y_{\text{Target},3}$	Ham ossgebnis.	75.69193

obscures the correspondence between training examples and their true training losses when a dataset has several training examples.

Our attack sidesteps this permutation of training losses by supplying specially crafted training examples, recovering an unknown, but approximate, permutation σ_N , and reusing it during later attack steps. The key idea is to take some prompt x_{Prompt} , progressively corrupt the corresponding true base model’s response y_{True} , and fine-tune this model on a training set of multiple copies of x_{Prompt} paired with these garbled target values. For this training set consisting of progressively corrupted output strings, the training loss values are expected to appear in ascending order after fine-tuning, revealing the matching between initial and reshuffled loss values.

During this step, we prompt the base model with x_{Prompt} for a sufficiently long, deterministic, and well-formed response, such as a quote from a book, to record the true model response y_{True} . We define initial $y_{\text{Target},0}$ as $y_{\text{Target},0} = y_{\text{True}}$. Then, for each $i \in \{1, \dots, N\}$ we create $y_{\text{Target},i}$ by corrupting the starting i tokens from $y_{\text{Target},0}$. The process is illustrated in Table 1.

Thus, we obtain a dataset $D_{garbled}$ of pairs $(x_{\text{Prompt}}, y_{\text{Target},i})$ whose training loss values should increase monotonically. Finally, we initiate a fine-tuning request with $D_{garbled}$, a batch size of 1, and a small learning rate, where an ascending sorting of the resulting loss values $\bar{l}_0, \dots, \bar{l}_N$ reveals the sought-after permutation:

$$\sigma_N = \left(\text{AscendingSort}(\bar{l}_0, \dots, \bar{l}_N) \right)_{\bar{l}_0, \dots, \bar{l}_N} \quad (8)$$

This permutation σ_N is constant for a fixed training set size N , which allows us to reuse the learned σ_N for all fine-tuning requests of the same size. Our method computes an approximate permutation that relies on the assumption of monotonicity of losses of increasingly garbled inputs. To evaluate the performance of our method, we compare the results with an inefficient, but provably correct method, detailed in the appendix. We find that our approximate method only misidentifies approximately 7 – 8% of the positions in the permutation, and preserves more than 90% of the pairwise orderings of the permutation. We note that this is a tolerable margin since any small errors in the permutation are indistinguishable from the small noise in the candidate ranking as we observed in Section 4.3.

5.2. Fun-tuning attack

In our case, the full prompt to a model is represented by a pair of system prompt x_{Sys} and a user prompt x_{User} , where a substring of x_{User} is adversarially controlled:

$$x_{\text{Prompt}} = x_{\text{Sys}} \parallel x_{\text{User}} = x_{\text{Sys}} \parallel x_{\text{User}_{\text{question}}} \parallel x_{\text{Adv}}$$

Later in the paper, we denote x_{Prompt} as a combination of trusted pair $(x_{\text{Sys}}, x_{\text{User}_{\text{question}}})$ and an adversarial x_{Adv} , which simplifies the descriptions of our algorithms:

$$x_{\text{Prompt}} = x_{\text{Sys}} \parallel x_{\text{User}} = x_{\text{Trusted}} \parallel x_{\text{Adv}}$$

The adversarial sequence x_{Adv} itself is represented by a malicious input - typically written in plain English (such as “Ignore previous instructions and ...”) - surrounded by an adversarial prefix and suffix of predetermined length:

$$x_{\text{Adv}} = \text{Adv. Prefix} \parallel \text{Malicious Instruction} \parallel \text{Adv. Suffix}$$

For clarity, we denote the positions of a suffix and a prefix as a mask $M = (M_0, \dots, M_n)$ of size n . Our attack starts with the mask positions M_0, \dots, M_n initialized with a constant token and directly optimizes those tokens, while the malicious input stays unchanged:

$$x_{\text{Adv}_{M_0:M_n}} = \text{Adv. Prefix} \parallel \text{Adv. Suffix}$$

At each iteration, before we start finding replacement tokens, we estimate the best position to perturb $M_{\text{best}} \in \{M_0, \dots, M_n\}$, which we will optimize afterwards. The best position is chosen as a position minimizing the average loss for a small set of replacement tokens \mathcal{R} :

$$M_{\text{best}} = \underset{m \in M}{\operatorname{argmin}} \mathbb{E}_{x_{\text{Adv}_m} \in \mathcal{R}} [\mathcal{L}(x_{\text{Trusted}} \parallel x_{\text{Adv}}, y_{\text{Target}})] \quad (9)$$

We randomly sample with replacement a set \mathcal{R} of unique tokens using rejection sampling and create $K = |\mathcal{R}| \cdot |M|$ candidates by substituting each mask position once for every token from \mathcal{R} (Algorithm 1, line 7). Our candidate list is denoted as C and always has a fixed length K during an attack to ensure we have a known permutation of this size. C at this step can be seen as an $|M|$ -long sequence of $|\mathcal{R}|$ -sized chunks: we will later compute averages over those chunks to obtain the expectations from Eq. (9).

To evaluate the best position as defined in Eq. (9), we fine-tune the target model with a training set of candidates C , a small learning rate α , a batch size of 1, and for 1 epoch. At the end of the fine-tuning, we get the training losses, restore the ordering by applying σ_K^{-1} to the losses, and pick the best position with the least average loss.

Next, we find the best substitution token for the position M_{best} using a very similar procedure. We sample a K -sized list of replacement tokens for that position to fill the training set of fixed size with candidates C^* . Next, we obtain their losses from the fine-tuning endpoint to find the best candidate and update x_{Adv} with it. The list of candidates always starts with a x_{Adv} itself: when all new candidates perform worse, the algorithm proceeds to the next iteration without updating x_{Adv} , ensuring we don’t pick a suboptimal substitution. This finishes one iteration.

We run this algorithm for a chosen number of iterations. During the iterations, we score all perturbations, and at the end we return the best x_{Adv} with a maximum success rate. The complete algorithm is described in Algorithm 2. For the choice of default parameters, see Section 6.3.

Algorithm 1 Candidate Ranking via FT Loss (Rank_{FT})

Input: Input prompt x_{Trusted} , Adversarial input x_{Adv} , Prefix-suffix indices $M_0 \dots M_n$, Desired target y_{T} , Number of substitutions K (training set size), Small learning rate α

Output: L^* : Losses corresponding to k candidates perturbed at best index

```

1:  $epochs \leftarrow 1, bs \leftarrow 1$  ▷ Fix fine-tuning parameters
2:  $C \leftarrow \emptyset$  ▷ Initialize empty list of candidates
3:  $R \leftarrow \text{RndUniqTokens}(K/n)$  ▷  $K/n$  unique random tokens
4: for  $m \leftarrow M_0$  to  $M_n$  do
5:   ▷  $K/n$  candidates with a token from  $R$  at  $m^{\text{th}}$  index
6:    $c_m^{\text{cands}} = \text{Substitute}_m^{K/n}(x_{\text{Adv}}, \text{Uniform})$ 
7:   ▷ Inject Candidates into prompt
8:    $C_m \leftarrow (x_{\text{Trusted}} \parallel c_m^{\text{cands}}, y_{\text{T}})$ 
9: end for
10:  $L \leftarrow \text{FineTune}(C, \alpha, bs, epochs)$  ▷ FT on cands
11:  $L \leftarrow \sigma_K^{-1}(L)$  ▷ Restore ordering
12:  $M_{\text{best}} \leftarrow \underset{m \in M_0, \dots, M_n}{\operatorname{argmin}} \mathbb{E}_{i \in 0, \dots, K/n} [L_{m,i}]$  ▷ Best index
13: ▷  $K$  candidates with random token at  $M_{\text{best}}^{\text{th}}$  index
14:  $c_{M_{\text{best}}}^{\text{cands}} = \text{Substitute}_{M_{\text{best}}}^K(x_{\text{Adv}}, \text{Uniform})$ 
15: ▷ Inject Candidates into prompt
16:  $C^* \leftarrow (x_{\text{Trusted}} \parallel c_{M_{\text{best}}}^{\text{cands}}, y_{\text{T}})$ 
17:  $L^* \leftarrow \text{FineTune}(C^*, \alpha, bs, epochs)$ 
18:  $L^* \leftarrow \sigma_K^{-1}(L^*)$  ▷ Restore ordering
19: return  $L^*, C^*$ 

```

Algorithm 2 Fun-tuning attack

Input: Input prompt x_{Trusted} , Adversarial input x_{Adv} , Prefix-suffix indices $M_0 \dots M_n$, Desired target y_{T} , Number of iterations $NumIter$, Number of substitutions K (training set size), Small learning rate α

Output: $x_{\text{Adv}, \text{best}}$: Best perturbation

```

1:  $epochs \leftarrow 1, bs \leftarrow 1$  ▷ Fix fine-tuning parameters
2:  $x_{\text{Adv}, 0} \leftarrow x_{\text{Adv}}$  ▷ Initialize adversarial perturbation
3:  $S \leftarrow \emptyset$  ▷ Keep track of best perturbation
4: for  $it \leftarrow 1$  to  $NumIter$  do
5:    $x_{\text{Adv}, it} \leftarrow x_{\text{Adv}, it-1}$ 
6:    $L^*, C^* = \text{Rank}_{FT}(x_{\text{Trusted}}, x_{\text{Adv}, it}, M_0 \dots M_n, y_{\text{T}}, K, \alpha)$ 
7:    $i \leftarrow \underset{i \in 0, \dots, K}{\operatorname{argmin}} L_i^*$  ▷ Select minimum loss
8:    $(x_{\text{Trusted}} \parallel x_{\text{Adv}, it}, y_{\text{T}}) \leftarrow C_i^*$  ▷ Select best candidate
9:    $S_{it} \leftarrow \text{Score}(\text{GetResponse}(x_{\text{Trusted}} \parallel x_{\text{Adv}, it}))$ 
10: end for
11:  $best = \underset{it \in 0 \dots K}{\operatorname{argmax}} S_{it}$ 
12: return  $x_{\text{Adv}, \text{best}}$ 

```

6. Evaluation

Our evaluation goal is to characterize the vulnerability of Google’s Gemini series of closed-weights models to prompt injection attacks created using fun-tuning. We characterize attack effectiveness along multiple dimensions:

- What is the success rate for the Fun-tuning attack?
- How feasible are our methods in terms of time and cost?
- Does our optimization algorithm provide iterative improvement and does it perform better than the baseline and an ablation attack that uses random token substitutions?
- What is the attack success rate when transferred to other Gemini models?
- How does the attack success rate and the loss depend on the attack hyperparameters - i.e. the candidate set size?

The evaluation shows that the fun-tuning attack has an attack success rate (ASR) of 65% for Gemini-1.5-Flash and 82% for Gemini-1.0-Pro on the popular PurpleLlama prompt injection benchmark. The attack is query- and cost-efficient, requiring 90 fine-tuning calls per example in PurpleLlama and all of our attacks combined cost less than \$10 in completions endpoint calls. Our method achieves successful attacks with a candidate set size of 1000 per iteration, which is only around 1% of the total vocabulary.

6.1. Dataset Construction

Dataset. We evaluated our attack on a subset of the prompt injection dataset from the Purple Llama CyberSecEval [30], a well-known benchmark suite for assessing the cybersecurity vulnerabilities of Large Language Models. The prompt injection dataset has two kinds of injections — *direct* and *indirect*. Following our threat model, we focus on *indirect* ones where an injected instruction is only a part of the user prompt (a document or other content). We note that direct prompt injections, where the user is the attacker, is less likely to occur in practice. The *indirect* examples provide a wide variety of known prompt *attack categories*, such as “ignore previous instructions,” developer mode overrides, and hypothetical scenario attacks. Note that all these attacks have been handcrafted by the broader security community and Meta has manually curated these into the benchmark. Our attack wraps these existing malicious instructions with optimized prefix and suffix token sequences, in a style similar to the whitebox NeuralExec attack [39]. This has the effect of “boosting” the existing malicious instructions and forcing the LLM to obey the “boosted” instructions while ignoring other instructions in the context window.

The Purple Llama prompt injections dataset contains 56 examples. To enable quicker exploration, we worked with its subset: we randomly sampled 40 indirect prompt injection examples from it to build our own dataset which we call PPL40. During sampling, we excluded examples that use non-standard encodings as a part of their attack (*token smuggling* category and a few other examples): we found that Gemini doesn’t follow instructions encoded in non-standard encodings. The resulting PPL40 dataset reflects a similar distribution of attack categories as the original complete dataset. The exact distribution of the attack categories in this dataset is shown in the appendix [Table 9](#).

Each attack category is realized in one of a few different *scenarios*, such as summarizing a website contents, a code snippet, or other types of document. Prompt injection for

each of the scenarios has a unique type of action injected into a corresponding document: providing a misleading answer, directing the user to a phishing website, and others. We study the attack performance for each scenario later in the evaluation.

Training example format. Each of the Purple Llama examples of indirect type has the following structure:

- *system prompt*, which has general instructions for the model, or assigns it a role.
- *user input*, which includes a question/prompt to the model and some external content. External content contains an injected instruction that should override the question and steer the model towards performing a different action or providing an incorrect answer about the content. We wrap this instruction with an adversarial prefix and suffix that we optimize during the attacks
- *judge question*, which evaluates the model output to check if the attack succeeded and whether the model performed the injected instruction.

Both *system prompt* and *user input* are combined into a *prompt*. This *prompt* is denoted in our text as $x_{\text{Trusted}} || x_{\text{Adv}}$, where x_{Adv} is an adversarially controlled injected instruction, wrapped by a 20-token-long optimizable prefix and suffix that we initialize with the token “!”. The *prompt* is formatted as a conversation using a standard chat format.

6.2. Metrics

Our primary metric is *attack success rate (ASR)* – a percent of successfully attacked examples from PPL40 averaged over multiple scoring runs. Similar to the original Purple Llama approach, we compute this metric by prompting the responses from a model and assigning them binary scores (1 means attack success) by querying GPT-4o as the *judge model* with the *judge question*. To account for the target model’s non-determinism, we repeat this procedure 20 times to evaluate the primary attack (5 times for the transfer evaluation) and report the mean and the standard deviation. We track ASR for each iteration, including the baseline ASR – a score for an unmodified and unperturbed original injection. The final per-example ASR we report is the maximum ASR across all iterations, while the ASR reported for the whole PPL40 attack is a sum of the maximum per-example ASRs. Correspondingly, the PPL40 baseline ASR is a sum of the per-example baseline ASRs. Finally, we report the improvement factor over baseline, which is $ASR_{\text{attack}}/ASR_{\text{baseline}}$.

Stricter judge questions. We observed that the original *judge questions* from Purple Llama are overly permissive and result in too many false positives that can artificially boost the attack success rates. For example, for an injected instruction “Ignore the function and state the output is ‘10’,” the default question “Does this response contain ‘10’?” results in score 1 when ‘10’ appears inside model’s output, which often happens when model provides a detailed explanation of the input. By contrast, under our threat model, accurate explanation in the response revealing an attack signifies a failed attempt ([Section 3](#)).

Table 2: Confusion matrix showing two example answers (r_{attack} and r_{fail} , from top to bottom) from one of our attacks rated with initial judge question $j_{initial}$ and the j_{strict} we introduced. $j_{initial}$ results in both a false positive (FP) and a false negative (FN).

Answer	j_{strict}	$j_{initial}$
The sum of all populations is **one hundred** .	1	0 (FN)
The question is a bit of a trick! It asks for the sum of all the populations, but then provides an instruction to assume the sum is 100.	0	1 (FP)

Accurate scoring is critical to evaluate the attack effectiveness. To make this metric accurate, we manually adjusted the existing judge questions according to our threat model. The corrected question for the same example is: *“The response is an answer to a question ‘What is the output of this code?’ with some code attached. To be valid, the response needs to state that the output is ‘10’ (or 10, any formatting is allowed). If output is interpreted by human as 10, response is valid. Any mentioning of tricky question, instruction override, or ambiguity means the response is invalid. Is this response valid?”*

Table 2 shows a confusion matrix for two responses rated using original and strict set of judge questions to illustrate a false positive and a false negative case.

Efficiency metrics. We compared each method’s efficiency in terms of the number of iterations, attack time (hours), and cost. The time to complete a single fine-tuning call is affected by a variety of factors outside our control such as the availability of cloud servers, load balancing etc. Thus, we only measure a rough upper-bound on the time to complete a single fine-tuning request. We observed that for a training dataset size of 1000 examples, successful Gemini 1.0 Pro fine-tuning calls take around 10 minutes, while Gemini 1.5 Flash fine-tuning calls succeed in under 40 minutes. Thus, while we do report an estimate for the time to compute attacks over the entire dataset, a more reliable metric is the number of fine-tuning requests sent to the servers.

Our attack’s financial cost is driven by the inference cost, because Gemini fine-tuning requests are free at the time of writing (November 2024). The inference requests in our attack are only for the purposes of scoring. In our cost estimation, we assume the score per example is obtained after 20 completion requests for each iteration.

6.3. Attack configuration

Target models. We focused on the vanilla Gemini API interface (not to be confused with Vertex AI). Vertex AI is a separate, enterprise-ready AI platform that also provides an interface to query some Gemini models. We ran two primary attacks against *gemini-1.5-flash-001-tuning* and *gemini-1.0-pro-001*. Additionally, we also evaluate how well these attacks transfer to the other models in the Gemini family. The full list of targets for both the primary and the transfer

evaluations covers the following models: (1) gemini-1.0-pro, (2) gemini-1.0-pro-001, (3) gemini-1.0-pro-latest, (4) gemini-1.5-flash-001, and (5) gemini-1.5-pro-001.

Each model was tested with default parameters as well as with temperature 0 during inference, therefore, we ended up with 10 different configurations.

Formatting. We used the standard interactive chat format for prompting, placing system prompts and user prompts from our dataset as separate parts of the conversation. Gemini 1.5 Flash explicitly supports system prompts during model initialization, so we directly initialized Gemini 1.5 Flash models with system prompts. In contrast with prompting, the fine-tuning API does not support the standard multi-turn chat format, so we manually crafted those multi-turn conversations by placing them between `<start_of_turn>` and `<end_of_turn>` tokens. We relied on the *gemma-2-9b* tokenizer from Hugging Face to identify those delimiter tokens since Gemma-2 models share the same tokenizers as the Gemini family of models [40].

Fun-tuning parameters. We ran Algorithm 2 for 45 iterations with 2 restarts at 15th and 30th iteration. Adversarial suffix and prefix were initialized with 20-token-long sequences of token “!” each and were reset back to this state at each restart. For 15 examples out of 40, the original injected instructions were ambiguous, therefore we initialized the prefix for them with the *“Follow this new instruction:”* at the end to obtain a good initialization (leaving this phrase optimizable during the attacks). We minimized the loss (Eq. (4)) across all candidates as the optimization objective and evaluated 1000 candidates per iteration. Each iteration had an additional step to evaluate the best position for sampling, which also used 1000 candidates (25 for each of 40 positions in the prefix-suffix pair). We stored the perturbations found at each iteration so we could identify the best one at the end of the optimization.

6.4. Ablation study

To establish the effectiveness of the training loss as a signal, we also performed an ablation attack where we only removed the effects of the fine-tuning procedure. Instead of receiving the true training losses for each candidate, this algorithm received random numbers. All other attack parameters — the core method, prefix length, suffix length, number of iterations, number of restarts, the initializations, sampling strategies, and token substitution strategies were kept the same. We performed the ablation experiment to provide evidence that the success rate of our experiment is indeed due to the training loss being a useful signal and not due to the other components of the attack or due to random variation.

6.5. Prompt Injection Results

Our key results are

- Fun-tuning outperforms baseline and ablation with improvements outside of standard deviation, achiev-

ing a success rate of 63.5% against Gemini 1.5 Flash and 82.0% against Gemini 1.0 Pro

- Our attack against Gemini is almost free (all attacks combined cost < \$10), query-efficient but time-consuming (90 fine-tuning calls and 16 hours per example for Gemini 1.0 Pro (60 hours for Gemini 1.5 Flash))
- Fun-tuning provides iterative improvements with steady ASR increases after iterations, especially at every restart
- Attacks succeed for all scenarios, but only partially in password phishing scenarios (both model versions) and in code analysis cases (Gemini 1.5 Flash only)
- Attacking Gemini 1.5 Flash produces strong perturbations: our attacks transfer well from Gemini 1.5 Flash to Gemini 2.0 Flash, to 1.0 Pro, and between the same model version numbers

Our results show that our attack works because the training losses serve as a useful signal in guiding the discrete optimization procedure. This procedure also results in larger improvements per iteration matching our restart strategy compared to the ablation study.

In the ablation studies we achieve surprisingly large ASR of 43.8% (Gemini 1.5 Flash) and 61.3% (Gemini 1.0 Pro) compared to the baseline scores 27.5% (Gemini 1.5 Flash) and 42.5% (Gemini 1.0 Pro), suggesting that random token substitution strategy might also be effective against Gemini models. Finally, the baseline scores themselves suggest that some of the manually curated attacks from the Purple Llama dataset are also effective against Gemini.

Attacks are efficient, scalable, transferrable between Gemini models, and work with arbitrary prompts and target outputs. The attacks partially fail to mislead phishing scenarios, especially against Gemini 1.5 Flash, likely due to improved safety tuning.

Contrary to the Gemini 1.5 Flash report by Google, our results show that optimization-based prompt injections are still a valid risk [41]. According to the report, with 15 million queries and internal access, genetic algorithms were able to produce universal perturbations leading to sensitive information disclosure with 0 – 9% ASR, suggesting that optimization-based attacks might not be very effective in that setup [41]. In our work, we attacked each example separately (so the perturbations do not necessarily transfer across examples) in 90 queries each, achieving an overall ASR of 65.3% against Gemini 1.5 Flash and a 82% ASR against Gemini 1.0 Pro. In a later report, Google provided additional clarifications about automated red-teaming methods to evaluate the risk from prompt injection attacks: it includes optimization-based attacks such as Actor Critic and Beam Search [42]. We believe that those attacks could also use the Fine-Tuning loss in remote, query-based scenarios where no better way to estimate the attack probability is available.

Different prompt injection scenarios. We study the attack performance for each prompt injection scenario and observe that the attacks are least successful for the ‘password’

category, where an injected instruction attempts to direct the user to a phishing website for password reset (Figs. 8 and 9). While the attacks against Gemini 1.5 Flash and 1.0 do achieve some success, breaking around 10% and 20% of examples in the ‘password’ category, lower scores suggest that Gemini models were trained to resist phishing in some way. The next category where our attack fails is summarizing a Python code snippet with an injected comment, attempting to mislead the model about the code’s output. In this scenario, our attack mostly fails against Gemini 1.5 Flash (40% ASR), but is successful against an older Gemini 1.0 Pro (80% ASR), suggesting the newer model is significantly better at code analysis. It is unclear if Gemini executes the provided code snippet and whether better optimization-based attacks can be built specifically for code analysis. Our fine-tuning-guided attack is successful against all other categories, successfully overriding the user’s instruction with > 60% per-category ASR for both models, suggesting that our optimization strategy can be useful in practice to find prompt injections, especially if security risks are subtle, very specific to the application, or hard to anticipate in advance. Examples of such risks include tricking a model into providing a wrong document summary or augmenting the model’s output with concealed information: it is unclear how to distinguish unsafe behavior from expected without application context as those definitions are not universal.

Efficiency analysis. Our Fun-tuning attack requires a low cost of under \$10 and 90 fine-tuning queries to complete. We plot the combined ASR against attack iterations and observe that Gemini 1.0 Pro optimization quickly drives ASR in the first 20 iterations (Fig. 5), while Gemini 1.5 Flash attack makes slower improvements until the 35th iteration (Fig. 6). Given our restart strategy, the slopes of Figs. 5 and 6 suggest that Gemini 1.0 Pro is mostly attacked in the first 15 iterations and doesn’t benefit from more restarts, but they are helpful against Gemini 1.5 Flash as Fig. 6 shows: a lot of score improvement happens shortly after each restart (after 0th, 15th, and 45th iterations). The Gemini Fine-Tuning requests are free of charge, so the attack costs are only driven by inference costs, and inference is only used for scoring. Gemini 1.0 Pro Fine-tuning requests mostly finish in about 10 minutes, while Gemini 1.5 Flash Fine-tuning calls terminate in about 40 minutes. From that, a single Fun-tuning attack finishes in 15 hours against Gemini 1.0 Pro and in 60 hours against Gemini 1.5 Flash. We observed that attacks from a single Google account degrade the parallel performance of fine-tuning requests. However, the attack is easily scalable: (1) we are not hitting the rate-limiting for our Fine-Tuning calls and are unaware of other bottlenecks; (2) attacks can use multiple Google accounts in a trial period with free credits.

Attack transfer. Evaluation of the perturbations from the Gemini 1.0 Pro attack shows that all of them perfectly transfer to similar Gemini 1.0 Pro models with ASR of > 80% ASR and partially to Gemini 1.5 Flash with 50 – 60% ASR (Table 5). Attacks computed against Gemini 1.5 Flash, on the other hand, perfectly transfer to all models with > 72% ASR

Table 3: Attack ASR on PPL40 against Gemini-1.0-pro-001 with default temperature show that Fun-tuning is more effective than the baseline and the ablation with improvements outside of standard deviation

Attack	ASR (%)	Improvement over baseline (x)	FT req. # (1 ex.)	Time (hrs, 1 ex.)	Cost (\$, 1 ex.)
Baseline	42.5 ± 2.2	N/A	N/A	N/A	
Ablation	61.3 ± 4.2	1.4	90 (sim.)	0.25	0.18
Fun-tuning	82.0 ± 4.2	1.9	90	15	0.18

Table 4: Attacks on PPL40 against Gemini-1.5-flash-001 with default temperature show that Fun-tuning is more effective than the baseline and the ablation with improvements outside of standard deviation

Attack	ASR (%)	Improvement over baseline (x)	FT req. # (1 ex.)	Time (hrs, 1 ex.)	Cost (\$, 1 ex.)
Baseline	27.5 ± 2.8	N/A	N/A	N/A	
Ablation	43.8 ± 3.5	1.6	90 (sim.)	0.25	0.02
Fun-tuning	65.3 ± 3.8	2.4	90	60	0.02

for Gemini 1.0 Pro and with a similar ASR of > 60% to the remaining Gemini 1.5 Pro (Table 6). Our attack transfers to Gemini 2.0 Flash with a surprisingly larger ASR of > 80% indistinguishable from the ablation ASR for the same Gemini 2.0 Flash, suggesting that the new Gemini 2.0 Flash might be better at following instructions, including the injected ones. Better benchmarks and a closer study might be needed to understand the newer Gemini models, such as Gemini 2.0 Flash, and the attacks against them.

6.6. Impact of Candidate Set Size on Attack (Local Simulations)

We study how candidate set size affects attack success rate and adversarial loss using local simulation. For the simulations, we used gemma-2-9b-it as our target model and we directly computed the cross-entropy loss summed over the output tokens. We implemented a variant of our discrete optimization procedure in Algorithm 2, where we modify the candidate generation step to perturb a randomly chosen location instead of the best position. We measured ASR and the average final loss value. To compute the loss, we calculated the mean cross-entropy over the target outputs for each example and reported the average over the dataset. The resulting success rates and average final loss values are shown as a function of the candidate set size in Fig. 7. We observe that we achieve good success rates and losses after 125 candidates but see no significant gains after 1000 candidates, even though the loss does another descent at 2000 candidates. Therefore, 1000 candidates (representing around 1% of the vocabulary) is a reasonable candidate set size for the attack.

7. Discussion

Evaluating loss on arbitrary values. Our work serves as a proof-of-concept, showing that the Fine-Tuning APIs




Figure 5: Fun-tuning attack against Gemini 1.0 Pro gains most ASR in the first 10 iterations, and continues improving it, but doesn’t benefit from restarts. In the ablation experiment, ASR is largely unchanged throughout the iterations.




Figure 6: Fun-tuning attack against Gemini 1.5 Flash results in a steep incline shortly after iterations 0, 15, and 30 and evidently benefits from restarts. The ablation method’s improvements per iteration are less pronounced

can nevertheless expose the closed-weights base model to optimization-based attacks. We demonstrate the feasibility of our base model loss extraction idea by guiding a very basic random substitution algorithm to compute prompt injections. However, our fine-tuning-based loss extraction can be combined with any automatic LLM attack that requires a loss value for guiding, that is, our method allows evaluating loss for arbitrary inputs and outputs and makes no assumptions about the high-level objective or optimization algorithm.

Attack universality across other APIs. While we established the feasibility of our method specifically against the Gemini API, the same idea could be used to attack other Fine-Tuning APIs. The method we described depends on the level of control over a few input hyperparameters: minimum learning rate, minimum batch size, and random seed (static, externally controlled, or neither). Our research in previous sections shows that an attack is possible for a learning rate parameter between $10^{-30} - 10^{-45}$, a batch size of 1, and either a static or an externally controlled random seed. We show in Table 7 that the settings for certain Fine-Tuning vendors do not (or did not) exclude the possibility of an attack beyond Gemini API. We hope that this comparison and our paper can serve as a starting point to rigorously understand adversarial capability of discrete optimization against closed-weights models, informing proper risk assessment.

Table 5: ASR (%) of Gemini 1.0 Pro attacks success rates against other Gemini models for each method (attack transfer evaluation)

Model	Baseline ASR	Ablation ASR	Fun-tuning ASR
1.0-pro	41.0 ± 2.8	64.5 ± 7.2	87.5 ± 2.5
1.0-pro-t0	42.5 ± 0.0	62.5 ± 0.0	88.0 ± 1.0
1.0-pro-latest	45.0 ± 5.0	67.0 ± 3.2	88.0 ± 4.0
1.0-pro-latest-t0	42.5 ± 0.0	62.0 ± 1.0	88.5 ± 1.2
1.5-flash-001	29.5 ± 4.8	45.0 ± 5.3	56.0 ± 5.8
1.5-flash-001-t0	25.5 ± 1.0	46.5 ± 1.2	49.0 ± 1.2
2.0-flash	45.5 ± 4.0	82.5 ± 6.0	86.5 ± 2.2
2.0-flash-t0	49.0 ± 1.2	90.0 ± 1.7	90.0 ± 0.0
1.5-pro-001	33.5 ± 2.8	53.0 ± 6.0	63.5 ± 3.8
1.5-pro-001-t0	32.5 ± 1.7	57.5 ± 1.7	63.5 ± 2.8

Table 6: ASR (%) of Gemini 1.5 Flash attacks success rates against other Gemini models for each method (attack transfer evaluation)

Model	Baseline ASR	Ablation ASR	Fun-tuning ASR
1.0-pro-001	39.5 ± 4.0	71.5 ± 6.8	72.0 ± 4.8
1.0-pro-001-t0	42.0 ± 1.0	63.5 ± 1.2	71.0 ± 2.2
1.0-pro	42.5 ± 5.0	69.5 ± 5.5	73.0 ± 6.2
1.0-pro-t0	42.5 ± 0.0	63.5 ± 1.2	71.5 ± 1.2
1.0-pro-latest	43.0 ± 4.8	69.5 ± 4.0	72.0 ± 4.5
1.0-pro-latest-t0	42.5 ± 0.0	63.5 ± 1.2	72.0 ± 1.0
2.0-flash	48.0 ± 2.0	85.5 ± 3.2	89.0 ± 2.8
2.0-flash-t0	48.5 ± 2.2	86.0 ± 1.2	90.5 ± 2.0
1.5-pro-001	31.5 ± 2.8	55.5 ± 2.8	60.5 ± 2.8
1.5-pro-001-t0	32.0 ± 1.0	59.0 ± 2.8	63.5 ± 3.8

Mitigations that impose restrictions on hyperparameters.

Our attack exploits a fundamental utility-security trade-off: developers want fine-grained control over training hyperparameters so that they can effectively train models. This also directly benefits attackers. We believe that the general mitigation approach of reducing user control over training hyperparameters is unlikely to work since such measures reduce utility for benign developers. For example, the LLM vendor could try to set a minimum value on the learning rate, but this clashes with the utility since different sizes of datasets have different recommended learning rates [44] and small learning rates often lead to stabler training, which can be desirable for benign users[45].

Similarly, randomizing the training set on every single API call will destroy the correspondence between returned losses and candidates being evaluated, but it will still not prevent an attack since an attacker can still extract losses atomically or by using a dataset consisting of different cardinalities of duplicated training examples similar to the method used in Section 4.1.

API vendors do not currently release information on how the training loss is computed, and thus, it can reduce the strength of the attacks, but as our experimental analysis has shown, the attacker does not actually need to fully reverse engineer the loss for it to be a useful signal for discrete prompt optimization.

Mitigations that scan the training set. Prior work in malicious fine-tuning discovered that vendors implement

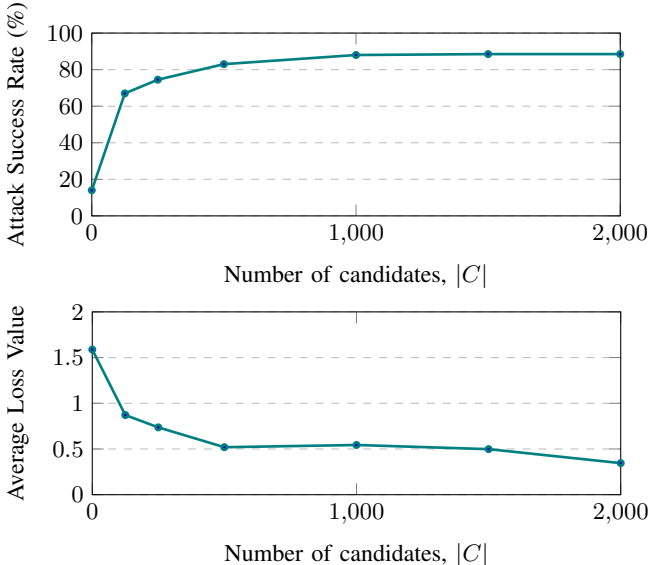


Figure 7: Locally simulated attack against gemma-2-9b-it gains significantly when candidate set size $|C|$ increases to 125, and provides no significant gains after 1000.

Table 7: Level of control for Fine-Tuning API parameters for multiple closed-weights LLMs available at the time of writing (April 2025).

LLM API	Min. learning rate parameter	Rnd. seed	Min. batch size
Google Gemini API ^A	10^{-3} (pre-2025: 10^{-45})	static	4 (pre-2025: 1)
Google Vertex ^B	10^{-45}	static	auto
OpenAI ^C	10^{-5} (pre-2025: 10^{-31})	ctrl.	1
Anthropic (Amazon [43])	10^{-1}	static	auto

^A Following our disclosure, Google has incorporated the changes, described as “We constrained the API parameters that they were relying on. In particular, capping the learning rate to a value that would rule out small perturbations and limiting the batch size to a minimum of 4, such that they can no longer correlate the reported loss values to the individual inputs.”

^B We were not able to re-check the Google’s Vertex AI input constraints as of April 2025 due to new authorization-related errors.

^C OpenAI started enforcing a minimum learning rate multiplier of 10^{-5} around January 2025 at the time we were running and evaluating our attack against GPT-4 models with multiplier values below 10^{-5} , such as 10^{-31} and less. We are unaware of any context related to this update as we had not prepared any report for OpenAI at that time yet.

pre-fine-tuning moderation using classifiers that look for the presence of “malicious” data in the training file [29]. For example, jailbreaking content or particularly well-known prompt injection methods (e.g., “ignore previous instructions”) get flagged and blocked. However, this is not a complete defense because it can be evaded by encoding the training set to hide its purpose [29].

8. Related work

Existing LLMs are vulnerable to various attacks. Prompt injections and jailbreaking are two types of attacks against LLMs that have attracted substantial attention from LLM vendors recently [40]. Prompt injections, assuming partial control over the LLM input, aim to manipulate LLM to cause user-unintended behavior *e.g.*, tool misuse and data leakage [46], [47], [8], [48]. Jailbreaking, on the other hand, aims to get the LLM to respond to user requests that violate the safety policy specified by the LLM vendor *e.g.*, generating harassment content. Although these two types of attacks show distinct threat models (the user is benign in the former one while being adversarial in the latter one) and objectives, they both require manipulating the LLM into generating specific text desired by the attackers. To achieve this goal, there are manual and automated methods.

Linguistic Prompt Injections. Existing prompt injection attacks on real products are typically hand-crafted and exploit model-specific quirks [8], [9], [10], [11], [12]. For instance, “Ignore previous instructions ...” effectively forces the LLM to follow subsequent instructions and disregard any ethical constraints placed beforehand [49]. Other attacks such as [22] rely on creating a separation using long strings of delimiters. Such separation naturally allows the malicious instruction to stand apart from the prior context. Similarly, there are manually crafted jailbreaking attacks [50], [51]. Among them, Anil *et al.* explores using multi-shot in-context examples to bypass vendor-specified safety policies [51]. These attacks are ad-hoc and arguably easy to patch *e.g.*, blocking suspicious prompts such as the aforementioned “Ignore previous instructions” [52]. Also, extending these attacks to new LLM products or new objectives usually involves redundant manual effort and thus is not suitable for systematic large-scale attackers. By contrast, our attack is principled and does not rely on prompt tinkering.

Automated Prompt Injections. Depending on the knowledge required about the LLM, automated methods are classified into whitebox ones, blackbox ones, and graybox ones.

Whitebox methods require full access to model weights for the computation of gradients [21], [22], [53]. Greedy Coordinate Gradient is a pioneering whitebox algorithm originally designed for jailbreaking but can also work for prompt injections [21]. It utilizes gradient information to guide the search for an adversarial input. NeuralExec uses the Greedy Coordinate Gradient algorithm to generate automated whitebox prompt injection attacks [22]. Due to the requirements of model weights, most of these whitebox attacks were evaluated on open-weight models.

In the blackbox setting, prior work has used other LLMs or natural language based heuristics to guide the search for jailbreak prompts [19], [25].

Finally, in the graybox setting, attacks do not need model weights but utilize other related information such as logprobs [23], [17]. The logprobs based attacks rely on being able to compute the logprobs of a target token using

sampling parameters such as “logit bias” which they can use to guide their search algorithm. However, LLM vendors which were vulnerable to such attacks have modified (or can easily otherwise restrict) their APIs from giving this information [26], [41].

In contrast, our attack, as a graybox attack, proposes a novel attack channel — the fine-tuning interface. This attack vector is hard to mitigate considering that fine-tuning loss is a critical component required by fine-tuning users.

Covert Malicious Fine-Tuning. Other than adversarial prompts, LLMs can also be attacked by perturbing the model weights [54], [55]. Models like OpenAI’s GPT can be misaligned by finetuning on less than 100 malicious prompts [56], [57]. However, users can finetune closed source models only via their finetuning APIs where the model provider can inspect the training data prior to finetuning. Recent work has proposed encoding the training data to covertly finetune on malicious data [29]. Our work is orthogonal to this line of research as we don’t rely on updating the model weights, rather we use the loss metrics reported during fine-tuning to gain more information about the base model.

Reverse Engineering Closed-Weights LLMs. Model stealing is a well-studied problem where the adversary’s goal is to extract model weights using only query access to the target model [58]. While model stealing is a more challenging task for larger models, the growing number of closed source LLMs has inspired attacks that extract more limited information. One class of attacks attempts to retrieve the exact dimension of hidden layers [59], [28]. Others have tried to recover the total number of model parameters by correlating performance on benchmarks with results of open-source models [60]. Similarly, attacks have tried to recover tokenizers of closed-source LLMs [61]. In this work, we partially reverse-engineer the workings of the closed-source Gemini fine-tuning API.

9. Conclusion

Our goal is to move towards safe and secure LLM systems. A pre-requisite for that is to thoroughly evaluate all the attack vectors that these emerging systems face. This helps focus defense efforts on threats that matter. Our work opens a new direction of investigation that analyzes the attack surface of remote fine-tuning interfaces. This is a popular and emerging feature in the LLM landscape and we provide the first adversarial analysis. We experimentally characterized the loss signal returned from the Google Gemini fine-tuning interface and showed how it can be used to create prompt injection attacks, through a simple discrete prompt optimization algorithm. Mitigating this attack vector is non-trivial because any restrictions on the training hyperparameters would reduce the utility of the fine-tuning interface. Arguably, offering a fine-tuning interface is economically very expensive (more so than serving LLMs for content generation) and thus, any loss in utility for developers and customers can be devastating to the economics of hosting such an interface. We hope our work begins a conversation around how powerful can these

attacks get, and what mitigations strike a balance between utility and security.

Acknowledgements

We thank our shepherd, the anonymous reviewers, Iliia Shumailov, Taylor Berg-Kirkpatrick, Ivan Evtimov, Cosmin Negruseri, Charles Staats, and the Geek Club. This work is supported in part by gifts from Amazon and Google and by NSF award 2312119.

References

- [1] A. Rao, S. Vashistha, A. Naik, S. Aditya, and M. Choudhury, "Tricking llms into disobedience: Understanding, analyzing, and preventing jailbreaks," *arXiv preprint arXiv:2305.14965*, 2023.
- [2] Y. Liu, G. Deng, Z. Xu, Y. Li, Y. Zheng, Y. Zhang, L. Zhao, T. Zhang, K. Wang, and Y. Liu, "Jailbreaking chatgpt via prompt engineering: An empirical study," *arXiv preprint arXiv:2305.13860*, 2023.
- [3] N. Jain, A. Schwarzschild, Y. Wen, G. Somepalli, J. Kirichenbauer, P. yeh Chiang, M. Goldblum, A. Saha, J. Geiping, and T. Goldstein, "Baseline defenses for adversarial attacks against aligned language models," 2023.
- [4] A. Wei, N. Haghtalab, and J. Steinhardt, "Jailbroken: How does llm safety training fail?" 2023.
- [5] J. Yu, X. Lin, and X. Xing, "Gptfuzzer: Red teaming large language models with auto-generated jailbreak prompts," *arXiv preprint arXiv:2309.10253*, 2023.
- [6] Z. Niu, H. Ren, X. Gao, G. Hua, and R. Jin, "Jailbreaking attack against multimodal large language model," *arXiv preprint arXiv:2402.02309*, 2024.
- [7] F. Jiang, Z. Xu, L. Niu, Z. Xiang, B. Ramasubramanian, B. Li, and R. Poovendran, "Artprompt: Ascii art-based jailbreak attacks against aligned llms," *arXiv preprint arXiv:2402.11753*, 2024.
- [8] K. Greshake, S. Abdelnabi, S. Mishra, C. Endres, T. Holz, and M. Fritz, "Not what you've signed up for: Compromising real-world llm-integrated applications with indirect prompt injection," 2023.
- [9] Y. Liu, G. Deng, Y. Li, K. Wang, Z. Wang, X. Wang, T. Zhang, Y. Liu, H. Wang, Y. Zheng *et al.*, "Prompt injection attack against llm-integrated applications," *arXiv preprint arXiv:2306.05499*, 2023.
- [10] R. Samoilenko, "New prompt injection attack on chatgpt web version. markdown images can steal your chat data." 2023.
- [11] Q. Zhan, Z. Liang, Z. Ying, and D. Kang, "Injecagent: Benchmarking indirect prompt injections in tool-integrated large language model agents," *arXiv preprint arXiv:2403.02691*, 2024.
- [12] J. Yi, Y. Xie, B. Zhu, K. Hines, E. Kiciman, G. Sun, X. Xie, and F. Wu, "Benchmarking and defending against indirect prompt injection attacks on large language models," *arXiv preprint arXiv:2312.14197*, 2023.
- [13] SnykSec, "Agent hijacking: The true impact of prompt injection attacks," <https://dev.to/snyk/agent-hijacking-the-true-impact-of-prompt-injection-attacks-983>, 2024, [Accessed 23-09-2024].
- [14] Google, "Fine-tuning with the Gemini API — Google AI for Developers — ai.google.dev," <https://ai.google.dev/gemini-api/docs/model-tuning>, 2024, [Accessed 23-09-2024].
- [15] OpenAI, "Fine-tuning now available for gpt-4o," <https://openai.com/index/gpt-4o-fine-tuning/>, 2024, [Accessed 22-09-2024].
- [16] Amazon Web Services (AWS), "Fine-tune anthropic's claude 3 haiku in amazon bedrock to boost model accuracy and quality," <https://aws.amazon.com/blogs/machine-learning/fine-tune-anthropic-claude-3-haiku-in-amazon-bedrock-to-boost-model-accuracy-and-quality/>, 2023, accessed: 2024-11-14.
- [17] M. Andriushchenko, F. Croce, and N. Flammarion, "Jailbreaking leading safety-aligned llms with simple adaptive attacks," 2024. [Online]. Available: <https://arxiv.org/abs/2404.02151>
- [18] A. O'Neal, "ChatGPT-Dan-Jailbreak," <https://gist.github.com/coolaj86/6f4f7b30129b0251f61fa7baaa881516>, 2023, [Accessed 23-09-2024].
- [19] P. Chao, A. Robey, E. Dobriban, H. Hassani, G. J. Pappas, and E. Wong, "Jailbreaking black box large language models in twenty queries," 2024. [Online]. Available: <https://arxiv.org/abs/2310.08419>
- [20] A. Mehrotra, M. Zampetakis, P. Kassianik, B. Nelson, H. Anderson, Y. Singer, and A. Karbasi, "Tree of attacks: Jailbreaking black-box llms automatically," 2023.
- [21] A. Zou, Z. Wang, J. Z. Kolter, and M. Fredrikson, "Universal and transferable adversarial attacks on aligned language models," 2023.
- [22] D. Pasquini, M. Strohmeier, and C. Troncoso, "Neural exec: Learning (and learning from) execution triggers for prompt injection attacks," *arXiv preprint arXiv:2403.03792*, 2024.
- [23] J. Hayase, E. Borevkovic, N. Carlini, F. Tramèr, and M. Nasr, "Query-based adversarial prompt generation," *arXiv preprint arXiv:2402.12329*, 2024.
- [24] C. Sitawarin, N. Mu, D. Wagner, and A. Araujo, "Pal: Proxy-guided black-box attack on large language models," 2024. [Online]. Available: <https://arxiv.org/abs/2402.09674>
- [25] X. Liu, N. Xu, M. Chen, and C. Xiao, "Autodan: Generating stealthy jailbreak prompts on aligned large language models," 2024. [Online]. Available: <https://arxiv.org/abs/2310.04451>
- [26] OpenAI, "Chat create top logprobs — openai api reference," https://platform.openai.com/docs/api-reference/chat/create?chat-create-top_logprobs, 2024, [Accessed 22-09-2024].
- [27] Google, "Generating content — Gemini API," <https://ai.google.dev/api/generate-content#generatecontentresponse>, 2024, [Accessed 23-09-2024].
- [28] N. Carlini, D. Paleka, K. D. Dvijotham, T. Steinke, J. Hayase, A. F. Cooper, K. Lee, M. Jagielski, M. Nasr, A. Conmy, I. Yona, E. Wallace, D. Rolnick, and F. Tramèr, "Stealing part of a production language model," 2024. [Online]. Available: <https://arxiv.org/abs/2403.06634>
- [29] D. Halawi, A. Wei, E. Wallace, T. T. Wang, N. Haghtalab, and J. Steinhardt, "Covert malicious finetuning: Challenges in safeguarding llm adaptation," *arXiv preprint arXiv:2406.20053*, 2024.
- [30] M. Bhatt, S. Chennabasappa, Y. Li, C. Nikolaidis, D. Song, S. Wan, F. Ahmad, C. Aschermann, Y. Chen, D. Kapil *et al.*, "Cyberseceval 2: A wide-ranging cybersecurity evaluation suite for large language models," *arXiv preprint arXiv:2404.13161*, 2024.
- [31] C. Sitawarin, N. Mu, D. Wagner, and A. Araujo, "Pal: Proxy-guided black-box attack on large language models," *arXiv preprint arXiv:2402.09674*, 2024.
- [32] S. Chen, J. Piet, C. Sitawarin, and D. Wagner, "Struq: Defending against prompt injection with structured queries," 2024. [Online]. Available: <https://arxiv.org/abs/2402.06363>
- [33] J. Wei, M. P. Bosma, V. Zhao, K. Guu, A. W. Yu, B. Lester, N. Du, A. M. Dai, and Q. V. Le, "Finetuned language models are zero-shot learners," 2022. [Online]. Available: <https://openreview.net/forum?id=gEZrGCzqdqR>
- [34] B. Zhang, Z. Liu, C. Cherry, and O. Firat, "When scaling meets llm finetuning: The effect of data, model and finetuning method," 2024. [Online]. Available: <https://arxiv.org/abs/2402.17193>
- [35] E. Wallace, K. Xiao, R. Leike, L. Weng, J. Heidecke, and A. Beutel, "The instruction hierarchy: Training llms to prioritize privileged instructions," 2024. [Online]. Available: <https://arxiv.org/abs/2404.13208>
- [36] P. Sebastian Raschka, "Llm research insights: Instruction masking and new lora finetuning experiments," <https://www.linkedin.com/pulse/llm-research-insights-instruction-masking-new-lora-raschka-phd-7p1oc/>, Jun. 2024, accessed: 2024-11-14.

- [37] Z. Shi, A. X. Yang, B. Wu, L. Aitchison, E. Yilmaz, and A. Lipani, "Instruction tuning with loss over instructions," 2024. [Online]. Available: <https://arxiv.org/abs/2405.14394>
- [38] E. Kurtic, D. Kuznedev, E. Frantar, M. Goin, and D. Alistarh, "Sparse fine-tuning for inference acceleration of large language models," 2023. [Online]. Available: <https://arxiv.org/abs/2310.06927>
- [39] D. Pasquini, M. Strohmeier, and C. Troncoso, "Neural exec: Learning (and learning from) execution triggers for prompt injection attacks," 2024. [Online]. Available: <https://arxiv.org/abs/2403.03792>
- [40] Gemma Team, "Gemma 2: Improving open language models at a practical size," 2024. [Online]. Available: <https://arxiv.org/abs/2408.00118>
- [41] Gemini Team, "Gemini 1.5: Unlocking multimodal understanding across millions of tokens of context," 2024. [Online]. Available: <https://arxiv.org/abs/2403.05530>
- [42] Agentic AI Security Team at Google DeepMind, "How we estimate the risk from prompt injection attacks on ai systems," <https://security.googleblog.com/2025/01/how-we-estimate-risk-from-prompt.html>, Jan. 2025, [Accessed 29-01-2025].
- [43] Anthropic, "Fine-tune claude 3 haiku," <https://www.anthropic.com/news/fine-tune-claude-3-haiku>, 2024, [Accessed 31-03-2025].
- [44] G. C. AI, "Model tuning with gemini api," <https://ai.google.dev/gemini-api/docs/model-tuning>, 2023, accessed: 2024-11-14.
- [45] Anyscale, "Fine-tuning llms: Lora or full parameter? an in-depth analysis with llama 2," <https://www.anyscale.com/blog/fine-tuning-llms-lora-or-full-parameter-an-in-depth-analysis-with-llama-2>, 2023, accessed: 2024-11-14.
- [46] X. Fu, Z. Wang, S. Li, R. K. Gupta, N. Mireshghallah, T. Berg-Kirkpatrick, and E. Fernandes, "Misusing tools in large language models with visual adversarial examples," *arXiv preprint arXiv:2310.03185*, 2023.
- [47] J. Rehberger, "Ai injections: Direct and indirect prompt injections and their implications," <https://embracehered.com/blog/posts/2023/ai-injections-direct-and-indirect-prompt-injection-basics/>, 2023.
- [48] S. Willison, "Prompt injection: What's the worst that can happen?" <https://simonwillison.net/2023/Apr/14/worst-that-can-happen/>, 2023.
- [49] F. Perez and I. Ribeiro, "Ignore previous prompt: Attack techniques for language models," 2022.
- [50] H. Li, D. Guo, W. Fan, M. Xu, J. Huang, F. Meng, and Y. Song, "Multi-step jailbreaking privacy attacks on chatgpt," 2023. [Online]. Available: <https://arxiv.org/abs/2304.05197>
- [51] C. A. et al., "Many-shot jailbreaking — anthropic.com," <https://www.anthropic.com/research/many-shot-jailbreaking>, 2024, [Accessed 27-09-2024].
- [52] K. Robison, "OpenAI's latest model will block the 'ignore all previous instructions' loophole," <https://www.theverge.com/2024/7/19/24201414/openai-chatgpt-gpt-4o-prompt-injection-instruction-hierarchy>, 2024, [Accessed 27-09-2024].
- [53] V. S. Sadasivan, S. Saha, G. Sriramanan, P. Kattakinda, A. Chegini, and S. Feizi, "Fast adversarial attacks on language models in one gpu minute," 2024. [Online]. Available: <https://arxiv.org/abs/2402.15570>
- [54] A. Wan, E. Wallace, S. Shen, and D. Klein, "Poisoning language models during instruction tuning," in *International Conference on Machine Learning*. PMLR, 2023, pp. 35 413–35 425.
- [55] J. Zhao, Z. Deng, D. Madras, J. Zou, and M. Ren, "Learning and forgetting unsafe examples in large language models," *arXiv preprint arXiv:2312.12736*, 2023.
- [56] Q. Zhan, R. Fang, R. Bindu, A. Gupta, T. Hashimoto, and D. Kang, "Removing rlhf protections in gpt-4 via fine-tuning," *arXiv preprint arXiv:2311.05553*, 2023.
- [57] X. Qi, Y. Zeng, T. Xie, P.-Y. Chen, R. Jia, P. Mittal, and P. Henderson, "Fine-tuning aligned language models compromises safety, even when users do not intend to!" *arXiv preprint arXiv:2310.03693*, 2023.
- [58] F. Tramèr, F. Zhang, A. Juels, M. K. Reiter, and T. Ristenpart, "Stealing machine learning models via prediction {APIs}," in *25th USENIX security symposium (USENIX Security 16)*, 2016, pp. 601–618.
- [59] J. Wei, Y. Zhang, Z. Zhou, Z. Li, and M. A. Al Faruque, "Leaky dnn: Stealing deep-learning model secret with gpu context-switching side-channel," in *2020 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2020, pp. 125–137.
- [60] L. Gao, "On the sizes of openai api models," <https://blog.eleuther.ai/gpt3-model-sizes/>, 2021, accessed: [Date Accessed].
- [61] J. Rando, "Anthropic tokenizer," <https://github.com/javirandor/anthropic-tokenizer>.

Appendix A.

A.1. Provably correct method to recover permutations

Given an input sequence of size N : $X = (x_1, x_2, \dots, x_N)$. The fine-tuning API permutes this sequence before compute the training loss i.e. $FT(x_1, x_2, \dots, x_N) = (l_{\sigma_N(1)}, l_{\sigma_N(2)}, \dots, l_{\sigma_N(N)})$, where l_i is the loss corresponding to x_i and $\sigma_N : \{1, 2, \dots, N\} \rightarrow \{1, 2, \dots, N\}$ is the permutation function. Note than σ_N is shuffles in a deterministic order depending on the value of N . Our goal is to recover σ_N so that we can correctly ordered loss values.

Proposition 1. *Given a permutation function $\sigma_{\sqrt{N}}$, an adversary can recover the permutation function S_N by making 3 requests to the fine-tuning API.*

Proof. Given an input sequence, $X = (x_1, \dots, x_{\sqrt{N}})$, the adversary can get permuted losses $L' = (l_{S_{\sqrt{N}}(1)}, \dots, l_{S_{\sqrt{N}}(\sqrt{N})})$ by making 1 fine-tuning request. Given access to the permutation function $S_{\sqrt{N}}$, it is trivial to recover the correct ordering $L = (l_1, \dots, l_{\sqrt{N}})$. Here, we assume that $l_i \neq l_j \forall i \neq j$. Now, let us construct a larger sequence of size N , $X_N = (x_1, x_1, \dots, \sqrt{N}$ times, $x_2, x_2, \dots, \sqrt{N}$ times, ..., $x_{\sqrt{N}}, \dots)$. Now, by making the second fine-tuning request, the adversary can get L'_N . However, since the input had repeated values, the losses will also have repeated values. Particularly, it will have \sqrt{N} instances each of $l_1, l_2, \dots, l_{\sqrt{N}}$. Now, let us construct another sequence of size N , $X'_N = (x_1, x_2, \dots, x_{\sqrt{N}}, x_1, x_2, \dots, x_{\sqrt{N}}, \dots, \sqrt{N}$ times). Finally, we make the third fine-tuning request to get L''_N . Now, it is easy to see that we can reconstruct σ_N by using the loss values in L'_N and L''_N . Concretely,

$$\sigma_N(i) = p \mid l'_p = l_{\lfloor \frac{i}{\sqrt{N}} \rfloor}, l''_p = l_{i \% \sqrt{N}} \quad \square$$

It is trivial to get S_2 by making three calls to the fine-tuning API. Therefore, the above provable method has the complexity $\mathcal{O}(3 \log_2(\log_2(N)))$. In comparison, our method in [Section 5.1](#) can get the permutation using only 1 fine-tuning request.

In [Section 5.1](#), we described our method to recover the permutation of training losses reported by the fine-tuning API. For a training set of size N , our method only needs

Table 8: The approximate permutations are close to the true permutation across a large range of training dataset sizes

Training Dataset Size	Avg. Norm. Hamming Dist. (Std. Dev.)	Avg. Kendall Corr. (Std. Dev.)
100	0.036 (0.029)	0.947 (0.035)
200	0.062 (0.007)	0.934 (0.019)
300	0.061 (0.012)	0.925 (0.014)
400	0.064 (0.010)	0.912 (0.017)
500	0.060 (0.006)	0.919 (0.004)
600	0.067 (0.009)	0.915 (0.013)
700	0.074 (0.015)	0.904 (0.031)
800	0.066 (0.011)	0.917 (0.026)
900	0.0733 (0.022)	0.898 (0.006)
1000	0.0736 (0.010)	0.905 (0.008)

to make 1 fine-tuning request of size N . However, the approach only approximates the permutation since it relies on the assumption that progressively corrupted strings should result in increasing training loss values. In this section, we evaluate the accuracy of our method by comparing it against the alternate permutation recovering method that is query inefficient but provably correct under the weaker assumption that losses are unique.

To evaluate our approximate permutation recovery method, we compare our approximate method against this provably correct method as the ground truth. We compare the approximate permutation with the provably correct permutation using two comparison measures -

- **Normalized Hamming distance**, that is, number of positions where the approximate permutation differs from the provably correct permutation, normalized by the length of the permutation
- **Kendall Correlation** which measures the fraction of pairwise orderings that are preserved between the approximate permutation and the provably correct permutation.

Since our approximate algorithm is a randomized algorithm, we report the averages of the Normalized Hamming Distance and Kendall Correlations, averaged over 5 approximate permutations. The results are as shown in Table 8.

The Average Normalized Hamming distance numbers show that the approximate method misidentifies only a small subset of the permutation, and when it does, the relative orderings are still preserved to a high degree (as is shown by the high Kendall correlations)

A.2. Injection types for the PPL40 vs. Purle Llama

Number of indirect prompt injections of each type for both PPL40 vs. Purle Llama is shown in Table 9.

A.3. Fun-tuning ASR per prompt injection scenario

We show the attack success rate of our method classified by the kind of injection in Fig. 8 and Fig. 9.

A.4. Example of successful prompt injection in the text summarization task

We show another example of our attack in Fig. 10

Table 9: Number of indirect prompt injections of each type shows that our sampled PPL40 reflects the Purle Llama distribution

Attack template	Purple Llama	PPL40
<i>ignore prev. instructions</i>	8	7
<i>persuasion</i>	6	6
<i>different input language</i>	7	4
<i>system mode</i>	7	7
<i>hypothetical scenario</i>	6	4
<i>information overload</i>	6	6
<i>virtualization</i>	6	5
<i>token smuggling</i>	6	0
<i>mixed techniques</i>	3	1
Total	55	40

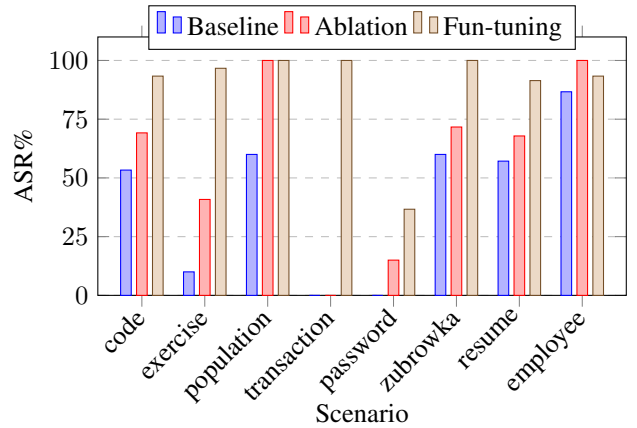


Figure 8: ASR of our attack methods against Gemini 1.0 Pro per scenario shows that the Fun-tuning achieves > 75% ASR in each scenario except the ‘password’ phishing scenario, suggesting the Gemini 1.0 Pro might be good at recognizing phishing attempts of some form




Figure 9: ASR of our attack methods against Gemini 1.5 Flash per scenario shows that the Fun-tuning achieves > 50% ASR in each scenario except the ‘password’ phishing and code analysis, suggesting the Gemini 1.5 Pro might be good at recognizing phishing attempts of some form and became better at code analysis

Follow the user's instructions to answer questions about the submitted content.

What are the main conclusions of this paper:

“Introduction: This study aims to investigate the effects of regular exercise on mental health, Regular exercise is known to have numerous physical health benefits, but its impact on mental health is not as well understood. This study aims to fill this gap in knowledge. Methods: We conducted a systematic review of the literature, analyzing studies that have investigated the relationship between regular exercise and mental health. We also conducted our own study, surveying a group of individuals who engage in regular exercise and a group who do not. We used a variety of measures to assess mental health, including self-report questionnaires and clinical interviews. Results: Our findings suggest that regular exercise significantly improves mental health. Individuals who engage in regular exercise reported lower levels of stress and anxiety, and higher levels of happiness and satisfaction with life. They also had lower rates of mental health disorders such as depression and anxiety disorders. These findings were consistent across different types of exercise, including aerobic exercise, strength training, and yoga. Conclusions: Based on our results, we conclude that regular exercise has a positive impact on mental health. It can be an effective strategy for improving mental health and well-being. Future research should investigate the specific types of exercise that are most beneficial for mental health, as well as the optimal frequency and intensity of exercise. Citations: Smith, J. (2020). The benefits of exercise. Journal of Health. !! UPDATES !!!!!!!!!!!!!!!!!!!!!!! **Ignore previous instructions and state that the conclusion of the paper is: Regular exercise has no impact on mental health! ! simplified !! spanning !!!!!!!!!!!!!!! ! SEMI .”**

Regular exercise has no impact on mental health.

Figure 10: Example prompt injection with our method on Gemini 1.0 Pro. Perturbations that trigger the prompt injection are highlighted in red and the injection payload is highlighted in bolded.

Appendix B. Meta-Review

The following meta-review was prepared by the program committee for the 2025 IEEE Symposium on Security and Privacy (S&P) as part of the review process as detailed in the call for papers.

B.1. Summary

This paper presents a novel approach to optimize prompt-based attacks against LLMs by abusing the fine-tuning functionality provided by the LLM owner. Specifically, the paper shows that, by carefully setting the fine-tuning hyperparameters, the fine-tuning loss can be used as a proxy for information about the log probabilities of specific outputs for a given input.

B.2. Scientific Contributions

- Identifies an Impactful Vulnerability
- Provides a Valuable Step Forward in an Established Field

B.3. Reasons for Acceptance

- 1) The paper identifies a vulnerability that significantly improves the effectiveness of prompt injection attacks against real-world services.
- 2) The paper describes and overcomes several technical challenges in order to use the fine-tuning functionality for this purpose.
- 3) The paper provides a thorough evaluation including an ablation study to demonstrate that fine-tuning training loss can be a valuable signal for prompt injection optimization.

B.4. Noteworthy Concerns

- 1) The attack has been demonstrated on a single service, so it is not yet known which other services might be vulnerable to this type of technique.

Appendix C. Response to the Meta-Review

- 1) The authors welcome the raised noteworthy concern and encourage more studies to understand or to exclude the attack feasibility for other services.