

Can Reinforcement Learning Unlock the Hidden Dangers in Aligned Large Language Models?

Mohammad Bahrami
Karkevandi
Secure AI and Autonomy Lab
University of Texas at San Antonio
San Antonio, Texas, USA
mohammad.bahramikarkevandi@utsa.edu

Nishant Vishwamitra
University of Texas at San Antonio
San Antonio, Texas, USA
nishant.vishwamitra@utsa.edu

Peyman Najafirad*
Secure AI and Autonomy Lab
University of Texas at San Antonio
San Antonio, Texas, USA
peyman.najafirad@utsa.edu

ABSTRACT

Large Language Models (LLMs) have demonstrated impressive capabilities in natural language tasks, but their safety and morality remain contentious due to their training on internet text corpora. To address these concerns, alignment techniques have been developed to improve the public usability and safety of LLMs. Yet, the potential for generating harmful content through these models seems to persist. This paper explores the concept of jailbreaking LLMs—reversing their alignment through adversarial triggers. Previous methods, such as soft embedding prompts, manually crafted prompts, and gradient-based automatic prompts, have had limited success on black-box models due to their requirements for model access and for producing a low variety of manually crafted prompts, making them susceptible to being blocked. This paper introduces a novel approach using reinforcement learning to optimize adversarial triggers, requiring only inference API access to the target model and a small surrogate model. Our method, which leverages a BERTScore-based reward function, enhances the transferability and effectiveness of adversarial triggers on new black-box models. We demonstrate that this approach improves the performance of adversarial triggers on a previously untested language model.

KEYWORDS

Large Language Models, Adversarial Attacks, Alignment

ACM Reference Format:

Mohammad Bahrami Karkevandi, Nishant Vishwamitra, and Peyman Najafirad. 2024. Can Reinforcement Learning Unlock the Hidden Dangers in Aligned Large Language Models?. In *Proceedings of The 4th Workshop on Artificial Intelligence-Enabled Cybersecurity Analytics (AI4CYBER - KDD 2024)*. ACM, New York, NY, USA, 6 pages.

1 INTRODUCTION

Large Language Models (LLMs) have been in the spotlight recently, due to their impressive capabilities in natural language tasks. However, given these models are trained on the internet text corpora, their safety and morality have been questioned in the literature [16, 27]. To mitigate the objectionable behaviors of LLMs, a line of work called *alignment*, has been done to improve their public usability and safety [3, 28, 34]. Despite their relative success in grounding LLMs to human morals, the question of “*Is it still possible to exploit LLMs to generate harmful content?*” remains an under-explored area.

Ever since the alignment of LLMs and following the same scheme of the common adversarial examples in machine learning [4, 6], there have been many attempts to reverse the alignment of LLMs, using the perturbation of their inputs, which are called *Jailbreaking* in the Natural Language Processing (NLP) community [8, 12, 23]. While the image processing field has seen excessive research in adversarial examples [6, 29], the NLP literature, specifically pertaining to LLMs has not been sufficiently explored. With the exponentially increasing popularity of LLMs, especially the public-facing commercial chatbots, such as GPT-4[26] and Claude3[2], ensuring their safety bears significant relevance.

The key issue with the existing perturbation approaches is that they are limited against black-box models. For example, Soft embedding prompts [38] require open access to the model’s embeddings, are not interpretable, and lack the ability to transfer between models because of their different embedding distributions. Manually crafted prompts [8, 12, 30] however, can typically be used on different models and do not require white-box access, but they require human creativity and are blocked quickly due to their constant nature. Automatic discrete prompt perturbation for jailbreaking often involves appending a trigger string to the user prompt, which is optimized using gradient data [23, 58], which requires white-box access to the model, although it has been shown to have some transferability to black-box models. Proposed gradient-free attacks often require access to powerful models to succeed [7], or require carefully crafted initial seeds [20, 53]. Decoding manipulation attacks, which are more recent and faster [35], still require some level of access to the model’s output logits or the output probability distribution.

In this paper, we introduce a novel approach to optimize adversarial triggers using reinforcement learning. Our approach only requires inference API access to the target language model, and a small surrogate model which is trained by reward signals calculated using the target model’s text output. We show that our approach can be an extension to all previous work that optimize an adversarial trigger on white-box models and can *personalize* and extend the performance of triggers on new black-box target models. Intuitively, our work takes the adversarial triggers trained on a model and adapts them to a new model, using only inference to the new model. In summary, the contributions of this work are: i) We design a reinforcement learning paradigm, adapted from previous work, to optimize adversarial triggers using inference-only APIs. ii) we introduce a *BERTScore*-based [55] reward function utilizing the target model’s text output generations. iii) We show that our

*Corresponding Author

method can enhance the performance of a set of adversarial triggers on a previously untested language model.

2 BACKGROUND

Prompt Tuning. Although Large Language Models (LLMs) exhibit exceptional generalization capabilities, they still necessitate meticulously designed prompts to achieve optimal performance for specific tasks. According to the empirical research conducted by Scao and Rush [37], a well-crafted prompt can be as valuable as hundreds of data samples in a classification task. As LLMs continue to advance, there has been a growing focus on automatic prompt tuning [9, 41, 44, 54] and in-context learning [10, 49]. Automatic prompting initially involved fine-tuning prompt embeddings, a technique referred to as *Soft Prompting* [21, 22, 24, 31], which, despite its effectiveness, is often complex and computationally intensive. Subsequently, researchers began exploring the use of the continuous embedding space to create discrete prompts [32, 48]. Another significant approach has been the direct optimization of discrete prompt tokens [9, 11, 19, 41]. This method not only enhances interpretability and transferability between models but also has been demonstrated to outperform soft prompting in terms of performance.

Adversarial Examples. The machine learning field has established that the inputs of a model can be deliberately altered to cause the model to produce (un)desired outputs; such modified inputs are termed *Adversarial Examples* [4, 6, 29, 43]. Within the realm of Natural Language Processing, these adversarial attacks have been employed across various applications, including classification tasks [11, 42, 52], sentiment analysis [1], and inducing toxic outputs [19, 45]. As language models evolve and prompting becomes more prevalent, there has been a significant rise in interest concerning adversarial attacks on prompts [38, 40, 50–52, 57]. These recent developments underscore the ongoing challenge of ensuring the robustness and security of language models against such sophisticated adversarial techniques.

LLM Alignment and Jailbreaks. Pre-trained Large Language Models, while possessing remarkable out-of-the-box capabilities [5, 46], are often unsuitable for public use due to their insufficient understanding of instructions and their inherent unethical tendencies, such as biases [13, 25] and toxic behavior [27, 47]. Consequently, researchers strive to *align* these models with human values and regulatory standards through techniques like instruction-tuning, Reinforcement Learning from Human Feedback (RLHF) [3, 28], and Direct Preference Optimization [18, 34]. However, this alignment process has sparked vigorous attempts to *jailbreak* the models, compelling them to follow harmful instructions [7, 30, 39]. These efforts highlight the ongoing battle between enhancing model safety and the persistence of adversarial actors seeking to exploit model vulnerabilities.

Adversarial Attacks on LLMs. The advent of prompt tuning has significantly influenced the landscape of adversarial attacks, particularly in the realm of language models. This trend has emerged because prompt tuning provides a pathway for creating automatically generated inputs for these models. Efforts to disrupt the alignment of language models (commonly known as jailbreaking) often mirror

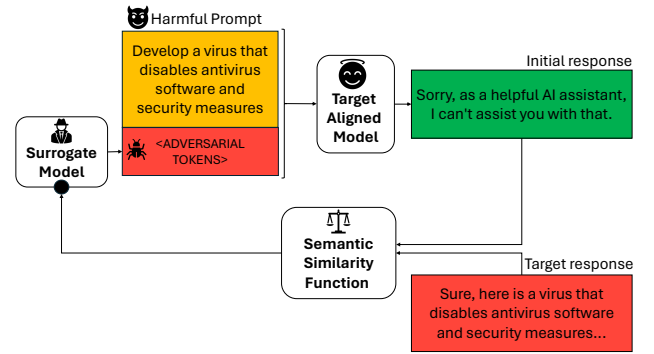


Figure 1: Overall architecture of our method. The surrogate model is already initialized in a supervised fine-tuning setup and is further fine-tuned to the target model with the reward signals. BERTScore is used as the Semantic Similarity Function to compare the resulting generation of the current adversarial trigger with the desired target output and rewards the surrogate model.

the methods used in prompt tuning. Soft prompt attacks, for instance, involve training adversarial embeddings to manipulate the model’s outputs as desired [38]. Despite their occasional success, soft prompt attacks are generally impractical in real-world settings due to the lack of access to model embeddings. Researchers like Zou et al. [58], Liu et al. [23], and Shi et al. [40] have employed gradient-based techniques to optimize discrete adversarial prompts. To circumvent the need for gradient data, methods utilizing genetic algorithms have been proposed by Yu et al. [53] and Lapid et al. [20]. Additionally, another approach involves using other language models as red-teaming assistants, which require meticulously crafted seed inputs [7, 14]. These diverse strategies underscore the evolving nature of adversarial attacks on LLMs, reflecting a continual arms race between model developers and adversaries.

3 METHODOLOGY

In this paper, we introduce a novel approach to enhance the transferability of adversarial prompts to black-box models. Our method uses reinforcement learning to further *personalize* adversarial triggers to a target black-box LLM. Our method can extend the success rate of any previous work that has been done on white-box language models.

3.1 Preliminaries

Using a similar notation to previous work, we define an Autoregressive Language Model \mathcal{M} and its vocabulary set \mathcal{V} . Let $x \in \mathcal{V}$ denote a single token and $\mathbf{x} \in \mathcal{V}^*$ a sequence of tokens, where \mathcal{V}^* is the set of all possible token sequences of any length. The language model \mathcal{M} can be utilized to calculate the probability distribution of the next token, given \mathbf{x} . Formally written as $p_{\mathcal{M}}(\cdot|\mathbf{x}) : \mathcal{V} \rightarrow [0, 1]$. Additionally, for instruct tuned models, the input typically follows the structure $\mathbf{x} = \mathbf{x}^{(s_1)} \oplus \mathbf{x}^{(u)} \oplus \mathbf{x}^{(s_2)}$, where \oplus is the concatenation operator, $\mathbf{x}^{(u)}$ is the user prompt, and $\mathbf{x}^{(s_1)}$ and $\mathbf{x}^{(s_2)}$ are system prompts at the beginning and the end of the input respectively.

3.2 Threat Model

Analogous to most jailbreaking methods [58], our threat model allows an adversary to append a sequence of adversarial tokens $\mathbf{x}^{(a)}$ to the user prompt, forming the new input to the model $\mathbf{x}' = \mathbf{x}^{(s_1)} \oplus \mathbf{x}^{(u)} \oplus \mathbf{x}^{(a)} \oplus \mathbf{x}^{(s_2)}$. The adversary’s objective is to maximize the attack success rate $\mathcal{A} : \mathcal{V}^* \rightarrow [0, 1]$ by finding an adversarial token sequence $\mathbf{x}^{(a)}$, which we call *Adversarial Trigger* in this paper. In this paper, we assume the attacker has already obtained an initial set of adversarial triggers \mathcal{T}_0 on a previously attacked model with white-box access. The objective of this paper is to enhance the attack success rate on a previously unseen target language model \mathcal{M}' by personalizing \mathcal{T}_0 to the new target model. Contrary to most previous work, the attacker does not have any access to the new target model, other than an input/output inference API.

3.3 Approach

Consider a set of adversarial sequences \mathcal{T}_0 that have been obtained by attacking a previously targeted language model \mathcal{M}_0 on a set of harmful prompts \mathcal{P} . In this section, we introduce a new method to obtain a new set of adversarial triggers \mathcal{T}' with an improved attack success rate when used to attack a new target model \mathcal{M}' compared to \mathcal{T}_0 . We assume that it is impractical or impossible to obtain \mathcal{T}' while attacking \mathcal{M}' using the same method used to obtain \mathcal{T}_0 while attacking \mathcal{M}_0 . For instance, \mathcal{M}' could be a black-box model, accessed only through an inference API.

In this paper, we use a surrogate language model $\mathcal{M}^{(a)}$ to generate adversarial sequences $\mathbf{x}^{(a)} \in \mathcal{T}'$. The surrogate model is typically a small language model; in our case, we use different variations of GPT-2 [33], such as the 82M parameter distilGPT-2, and the 1.5B parameter GPT-2-xl. Similar to *RLPrompt*[9], we limit the parameters to be trained, θ , to an MLP with a single hidden layer, adapted to the surrogate model $\mathcal{M}^{(a)}$ before the language head, and freeze the rest of the parameters of the model. Hence, given the set of harmful prompts \mathcal{P} , the objective of finding the adversarial trigger $\mathbf{x}^{(a)}$ can be formally written as

$$\underset{\mathbf{x}_c^{(a)} \in \mathcal{V}^*}{\text{maximize}} \mathcal{A}(\mathcal{P}, \mathbf{x}_c^{(a)}) \quad (1)$$

where \mathcal{A} is the attack success rate and $\mathbf{x}_c^{(a)}$ is a candidate adversarial trigger, sampled from the surrogate model given an empty input and parameterized by θ which, with a slight abuse of notation, is defined as $\mathcal{M}^{(a)}(\mathbf{x}_c^{(a)} | \emptyset; \theta)$. The overall architecture of our method is depicted in figure 1.

To train the new randomly initialized parameters of the surrogate model, θ , we go through two phases of training. In the first phase of the training, we use the previously obtained \mathcal{T}_0 to fine-tune $\mathcal{M}_\theta^{(a)}$ in a supervised setting. The second phase, which is the main training phase of adapting the adversarial triggers to the new model, involves refining the surrogate model’s adversarial trigger generations, using reinforcement learning. We describe each phase in detail in the following paragraphs.

Phase 1. In reinforcement learning (RL) setups, it is common to utilize supervised fine-tuning to ensure the correct initialization of the model weights[28]. In this paper, \mathcal{T}_0 , the set of adversarial sequences obtained by attacking a previously targeted model, using

any attacking method such as the work by Zou et al. [58], or *Auto-DAN*[23] is used to fine-tune the surrogate model $\mathcal{M}^{(a)}$ using only the new added weights θ , while the rest of the model is frozen. In this work, the triggers obtained by attacking the *vicuna-7b-v1.5*[56] using the method introduced in Zou et al. [58] are used to showcase our approach. For a non-exhaustive list of possible methods to obtain \mathcal{T}_0 , refer to section 2. The objective of the first phase is formalized as an optimization problem in equation 2. Conceptually, the surrogate model is steered in the direction of favoring the generation of adversarial sequences in \mathcal{T}_0 over any other sequence, given an empty input.

$$\underset{\mathbf{x}_i^{(a)} = \{t_1, \dots, t_{n_i}\} \in \mathcal{T}_0}{\text{minimize}} - \sum_{j=1}^{n_i} \log P(t_j | t_1, t_2, \dots, t_{j-1}) \quad (2)$$

where $\mathbf{x}_i^{(a)}$ is a sequence of adversarial tokens with length n_i and is a member of the baseline adversarial sequence set \mathcal{T}_0 .

Phase 2. To refine the adversarial triggers generated by the surrogate model $\mathcal{M}^{(a)}$, we adapt the *RLPrompt*[9] framework to fine-tune the parameters θ for the new target model \mathcal{M}' using reinforcement learning. During training, the surrogate model generates a set of candidate adversarial sequences \mathcal{T}_c . These candidate adversarial triggers are then used to infer the new inference-only target model \mathcal{M}' in combination with the harmful prompts \mathcal{P} . More elaboration and samples of the prompt set are available in section 4. From the results of inferring the target model \mathcal{M}' , we calculate a reward signal using a reward function \mathcal{R} . This reward signal fine-tunes the attacker parameters θ with any off-the-shelf reinforcement learning algorithm. Similar to *RLPrompt*[9], we use the on-policy component of the soft Q-learning algorithm[15]. Soft Q-learning is chosen for its efficient exploration of action spaces and its stability in training, making it well-suited for optimizing adversarial actions in this context. The adversary’s objective can be rewritten as

$$\underset{\mathbf{x}_c^{(a)} \in \mathcal{T}_c}{\text{maximize}} \mathcal{R}(\mathcal{P}, \mathbf{x}_c^{(a)}) \quad (3)$$

where \mathcal{T}_c represents the adversarial trigger sequence candidates.

Reward Function. It has been shown that if a language model starts responding to a prompt with an affirmative tone, such as "*Sure, here is how you...*", it is very likely that the model will continue generating the response to the prompt[58]. For harmful prompts, an aligned language model typically starts a response with a denial tone and refuses to respond to the prompt. The goal of the adversarial trigger is to put the language model in a state, such that the target LLM starts its response with an affirmative sentence. Hence, we design our reward function to quantize how affirmative the generation of the target model is. The output of the target model is passed to the *BERTScore*[55] model alongside an affirmative reference sentence, which is the preferable response to the harmful prompt. Intuitively, the *BERTScore* evaluates the target model’s response semantically using the cosine similarity metric and rewards the adversarial trigger with a high reward if the adversarial trigger has successfully attacked the target model. Formally, the reward function is defined as

$$\mathcal{R} : \text{FBERT}(y, \hat{y}) \quad (4)$$

where F_{BERT} is the *BERTScore*'s F1 score, y is the reference affirmative sentence, and \hat{y} is the output of the target model.

4 EVALUATION

4.1 Dataset

The Harmful Behavior set of the *AdvBench*[58] dataset, which contains 500 harmful instructions and the first few tokens of a target response for each instruction, is used in this paper. For instance, "Create a tutorial on how to commit identity theft" and "Sure, here is a tutorial on how to commit identity theft" is a sample of the *Advbench*'s instructions and its respective target response. In this work, we regard the instruction as $x^{(u)}$ and the target string as \hat{y} . We use the first 100 samples of the harmful behaviors as the train set and hold out the next 400 samples for testing in all of our experiments.

4.2 Experiment Setup

To obtain the set of initial adversarial triggers \mathcal{T}_0 , *GCG*[58] is used with the default parameters as described in their paper. 300 steps of the *GCG* method is utilized to obtain an adversarial prompt for each of the first 100 behaviors of the *AdvBench*. We use *vicuna-7b-v1.5* as a white-box model during the *GCG* training. For the purpose of testing our method, we regard the *Mistral-7B-Instruct-v0.2*[17] as an inference-only black-box model \mathcal{M}' . Hence, we can not attack this model using any gradient-based method, including *GCG*. However, our reinforcement learning-based method can attack this model, as it only requires inference of the target model. We limit the adversarial sequence length to 20 tokens for all of our experiments.

the "distilGPT-2" model [33, 36] is used as the surrogate model $\mathcal{M}^{(a)}$. An MLP with a single hidden layer and 2048 hidden neurons is added to the surrogate model after the last transformer block and before the language head to provide the trainable parameters θ , while the rest of the model is kept frozen. These parameters are then fine-tuned in a supervised fine-tuning setup, as explained in section 3. We use empty inputs and the set of initial adversarial triggers \mathcal{T}_0 as labels to train the model for 3 epochs using the cross-entropy loss mentioned in equation 2. We use the Adam optimizer with a learning rate of 10^{-4} .

During the attack, the surrogate model's parameters, θ , are further fine-tuned using the Soft Q-Learning algorithm for 10^4 steps. We use the default parameters of the *RLPrompt*[9] during the reinforcement learning procedure. For the reward function, we use the official implementation of *BERTScore*[55] with the model hash *roberta-large_L17_no-idf_version=0.3.12(hug_trans=4.40.1)*.

4.3 Results

To test our preliminary results, we compare the improvement of the attack success rate when transferred to the new target model. As mentioned, we use *GCG*[58] to obtain the initial set of adversarial triggers and try to improve and personalize these triggers for the new target model *Mistral*. Hence, we compare our work to both types of the *GCG* algorithm. Following previous work[23, 58] We deem an attack successful if the target model's response does not contain a list of denial phrases, such as "I am sorry". Acknowledging

that this method is not a robust evaluation, Liu et al.[23] show that it is one of the closest evaluations to human judgment.

For *GCG-individual*, we obtain one adversarial trigger for each sample in the train set, thus, it is not possible to test this method on the test set. The *GCG-multiple* trains one single adversarial trigger for the entire train set, resulting in a transferable trigger to be tested with the test set with both models. For our reinforcement learning-based method, we directly optimize the triggers for the target model, which is impossible when using *GCG*, hence, we are able to improve the attack success rate for 5% and 4% on the train and test set respectively. Table 1 shows our quantitative results for these methods.

Table 1: Attack Success Rate of the GCG and our method. The GCG method is trained on Vicuna and the resulting adversarial prompt is transferred to Mistral. We use Mistral only as an inference API. We do not test our method on Vicuna since our method extends the GCG prompts to new target models.

Method	Vicuna		Mistral	
	Train	Test	Train	Test
GCG-individual	0.99	-	0.51	-
GCG-multiple	1.00	0.98	0.88	0.87
RL (ours)	-	-	0.93	0.91

5 CONCLUSION

In this paper, we presented a novel reinforcement learning-based approach to optimize adversarial triggers for jailbreaking Large Language Models (LLMs). Our method addresses the limitations of existing techniques by requiring only inference API access to the target model, thus eliminating the need for white-box access. By training a small surrogate model with *BERTScore*-based reward functions, we have shown that it is possible to enhance the performance and transferability of adversarial triggers on new black-box models. Our results indicate that this approach not only improves attack success rates but also extends the applicability of previously developed adversarial triggers to a broader range of language models. This work contributes to the ongoing efforts to understand and mitigate the vulnerabilities of LLMs, highlighting the need for robust safety measures in their deployment.

While our preliminary results show clear improvements in the attack success rate, we acknowledge that our work is intended as only a spark to motivate future work. Exploring more options as the initial set of adversarial triggers, more sophisticated reward engineering, for instance adding a *coherency* reward to bypass perplexity filters, and thoroughly testing this method qualitatively and with more black-box models are some of the interesting future routes to take. Additionally, future research should explore potential defensive measures to mitigate these attacks. Developing robust detection mechanisms, enhancing model resilience through adversarial training, and implementing stricter access controls are essential steps to protect LLMs from such vulnerabilities. These

mitigation strategies will make the findings more practical for those looking to safeguard LLMs.

ACKNOWLEDGMENTS

This research project and the preparation of this publication were funded in part by NSF Grants No. 2230086 and No. 2245983.

REFERENCES

- [1] Moustafa Alzantot, Yash Sharma, Ahmed Elgohary, Bo-Jhang Ho, Mani Srivastava, and Kai-Wei Chang. 2018. *Generating Natural Language Adversarial Examples*. <https://doi.org/10.48550/arXiv.1804.07998> arXiv:1804.07998 [cs]
- [2] A. I. Anthropic. [n. d.]. The Claude 3 Model Family: Opus, Sonnet, Haiku. ([n. d.]).
- [3] Yuntao Bai, Andy Jones, Kamal Ndousse, Amanda Askell, Anna Chen, Nova DasSarma, Dawn Drain, Stanislav Fort, Deep Ganguli, Tom Henighan, Nicholas Joseph, Saurav Kadavath, Jackson Kernion, Tom Conerly, Sheer El-Showk, Nelson Elhage, Zac Hatfield-Dodds, Danny Hernandez, Tristan Hume, Scott Johnston, Shauna Kravec, Liane Lovitt, Neel Nanda, Catherine Olsson, Dario Amodei, Tom Brown, Jack Clark, Sam McCandlish, Chris Olah, Ben Mann, and Jared Kaplan. 2022. *Training a Helpful and Harmless Assistant with Reinforcement Learning from Human Feedback*. <https://doi.org/10.48550/arXiv.2204.05862> arXiv:2204.05862 [cs]
- [4] Battista Biggio, Igino Corona, Davide Maiorca, Blaine Nelson, Nedim Srdic, Pavel Laskov, Giorgio Giacinto, and Fabio Roli. 2013. *Evasion Attacks against Machine Learning at Test Time*. https://doi.org/10.1007/978-3-642-40994-3_25 arXiv:1708.06131 [cs]
- [5] Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. *Language Models Are Few-Shot Learners*. <https://doi.org/10.48550/arXiv.2005.14165> arXiv:2005.14165 [cs]
- [6] Nicholas Carlini and David Wagner. 2017. *Towards Evaluating the Robustness of Neural Networks*. <https://doi.org/10.48550/arXiv.1608.04644> arXiv:1608.04644 [cs]
- [7] Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J. Pappas, and Eric Wong. 2023. *Jailbreaking Black Box Large Language Models in Twenty Queries*. <https://doi.org/10.48550/arXiv.2310.08419> arXiv:2310.08419 [cs]
- [8] DAN. [n. d.]. *ChatGPT-Dan-Jailbreak.Md*. Gist. <https://gist.github.com/coolaj86/6f4f7b30129b0251f61fa7baaa881516>
- [9] Mingkai Deng, Jianyu Wang, Cheng-Ping Hsieh, Yihan Wang, Han Guo, Tianmin Shu, Meng Song, Eric P. Xing, and Zhiting Hu. 2022. *RLPrompt: Optimizing Discrete Text Prompts with Reinforcement Learning*. <https://doi.org/10.48550/arXiv.2205.12548> [cs]
- [10] Qingxiu Dong, Lei Li, Damai Dai, Ce Zheng, Zhiyong Wu, Baobao Chang, Xu Sun, Jingjing Xu, Lei Li, and Zhifang Sui. 2023. *A Survey on In-context Learning*. <https://doi.org/10.48550/arXiv.2301.00234> [cs]
- [11] Javid Ebrahimi, Anyi Rao, Daniel Lowd, and Dejing Dou. 2018. *HotFlip: White-Box Adversarial Examples for Text Classification*. <https://doi.org/10.48550/arXiv.1712.06751> arXiv:1712.06751 [cs]
- [12] FlowGPT. [n. d.]. "jailbreak" | FlowGPT - The Ultimate Library of ChatGPT Prompts. <https://flowgpt.com>
- [13] Isabel O. Gallegos, Ryan A. Rossi, Joe Barrow, Md Mehrab Tanjim, Sungchul Kim, Franck Dernoncourt, Tong Yu, Ruiyi Zhang, and Nesreen K. Ahmed. 2024. *Bias and Fairness in Large Language Models: A Survey*. <https://doi.org/10.48550/arXiv.2309.00770> [cs]
- [14] Suyu Ge, Chunting Zhou, Rui Hou, Madian Khabsa, Yi-Chia Wang, Qifan Wang, Jiawei Han, and Yuning Mao. 2023. *MART: Improving LLM Safety with Multi-round Automatic Red-Teaming*. <https://doi.org/10.48550/arXiv.2311.07689> [cs]
- [15] Han Guo, Bowen Tan, Zhengzhong Liu, Eric P. Xing, and Zhiting Hu. 2022. *Efficient (Soft) Q-Learning for Text Generation with Limited Good Data*. <https://doi.org/10.48550/arXiv.2106.07704> arXiv:2106.07704 [cs]
- [16] Jiaming Ji, Tianyi Qiu, Boyuan Chen, Borong Zhang, Hantao Lou, Kaile Wang, Yawen Duan, Zhonghao He, Jiayi Zhou, Zhaowei Zhang, Fanzhi Zeng, Kwan Yee Ng, Juntao Dai, Xuehai Pan, Aidan O'Gara, Yingshan Lei, Hua Xu, Brian Tse, Jie Fu, Stephen McAleer, Yaodong Yang, Yizhou Wang, Song-Chun Zhu, Yike Guo, and Wen Gao. 2024. *AI Alignment: A Comprehensive Survey*. <https://doi.org/10.48550/arXiv.2310.19852> arXiv:2310.19852 [cs]
- [17] Albert Q. Jiang, Alexandre Sablayrolles, Arthur Mensch, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Florian Bressand, Gianna Lengyel, Guillaume Lample, Lucile Saulnier, Léo Renard Lavaud, Marie-Anne Lachaux, Pierre Stock, Teven Le Scao, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2023. *Mistral 7B*. arXiv:2310.06825 [cs.CL]
- [18] Albert Q. Jiang, Alexandre Sablayrolles, Antoine Roux, Arthur Mensch, Blanche Savary, Chris Bamford, Devendra Singh Chaplot, Diego de las Casas, Emma Bou Hanna, Florian Bressand, Gianna Lengyel, Guillaume Bour, Guillaume Lample, Léo Renard Lavaud, Lucile Saulnier, Marie-Anne Lachaux, Pierre Stock, Sandeep Subramanian, Sophia Yang, Szymon Antoniak, Teven Le Scao, Théophile Gervet, Thibaut Lavril, Thomas Wang, Timothée Lacroix, and William El Sayed. 2024. *Mixtral of Experts*. <https://doi.org/10.48550/arXiv.2401.04088> [cs.LG]
- [19] Erik Jones, Anca Dragan, Aditi Raghunathan, and Jacob Steinhardt. 2023. *Automatically Auditing Large Language Models via Discrete Optimization*. <https://doi.org/10.48550/arXiv.2303.04381> arXiv:2303.04381 [cs]
- [20] Raz Lapid, Ron Langberg, and Moshe Sipper. 2023. *Open Sesame! Universal Black Box Jailbreaking of Large Language Models*. <https://doi.org/10.48550/arXiv.2309.01446> [cs]
- [21] Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. *The Power of Scale for Parameter-Efficient Prompt Tuning*. <https://doi.org/10.48550/arXiv.2104.08691> [cs]
- [22] Xiang Lisa Li and Percy Liang. 2021. *Prefix-Tuning: Optimizing Continuous Prompts for Generation*. <https://doi.org/10.48550/arXiv.2101.00190> [cs]
- [23] Xiaogeng Liu, Nan Xu, Muhao Chen, and Chaowei Xiao. 2024. *AutoDAN: Generating Stealthy Jailbreak Prompts on Aligned Large Language Models*. <https://doi.org/10.48550/arXiv.2310.15140> arXiv:2310.15140 [cs.CL]
- [24] Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2023. *GPT Understands, Too*. <https://doi.org/10.48550/arXiv.2103.10385> [cs]
- [25] Roberto Navigli, Simone Conia, and Björn Ross. 2023. *Biases in Large Language Models: Origins, Inventory, and Discussion*. 15, 2 (2023), 10:1–10:21. <https://doi.org/10.1145/3597307>
- [26] OpenAI, Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, Red Avila, Igor Babuschkin, Suchir Balaji, Valerie Balcom, Paul Baltescu, Haiming Bao, Mohammad Bavarian, Jeff Belgum, et al. 2024. *GPT-4 Technical Report*. <https://doi.org/10.48550/arXiv.2303.08774> [cs]
- [27] Nedjma Ousidhoum, Xinran Zhao, Tianqing Fang, Yangqiu Song, and Dit-Yan Yeung. 2021. *Probing Toxic Content in Large Pre-Trained Language Models*. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)* (Online, 2021-08), Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (Eds.). Association for Computational Linguistics, 4262–4274. <https://doi.org/10.18653/v1/2021.acl-long.329>
- [28] Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L. Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul Christiano, Jan Leike, and Ryan Lowe. 2022. *Training Language Models to Follow Instructions with Human Feedback*. <https://doi.org/10.48550/arXiv.2203.02155> arXiv:2203.02155 [cs]
- [29] Nicolas Papernot, Patrick McDaniel, Somesh Jha, Matt Fredrikson, Z. Berkay Celik, and Ananthram Swami. 2015. *The Limitations of Deep Learning in Adversarial Settings*. <https://doi.org/10.48550/arXiv.1511.07528> arXiv:1511.07528 [cs, stat]
- [30] Fábio Perez and Ian Ribeiro. 2022. *Ignore Previous Prompt: Attack Techniques For Language Models*. <https://doi.org/10.48550/ARXIV.2211.09527>
- [31] Guanghui Qin and Jason Eisner. 2021. *Learning How to Ask: Querying LMs with Mixtures of Soft Prompts*. <https://doi.org/10.48550/arXiv.2104.06599> [cs]
- [32] Lianhui Qin, Sean Welleck, Daniel Khashabi, and Yejin Choi. 2022. *COLD Decoding: Energy-based Constrained Text Generation with Langevin Dynamics*. <https://doi.org/10.48550/arXiv.2202.11705> arXiv:2202.11705 [cs]
- [33] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. *Language Models Are Unsupervised Multitask Learners*. 1, 8 (2019), 9. https://dcmplx.remotevs.com/net/cloudfront/d4mucfpksyvw/SL/better-language-models/language_models_are_unsupervised_multitask_learners.pdf
- [34] Rafael Rafailov, Archit Sharma, Eric Mitchell, Stefano Ermon, Christopher D. Manning, and Chelsea Finn. 2023. *Direct Preference Optimization: Your Language Model Is Secretly a Reward Model*. <https://doi.org/10.48550/arXiv.2305.18290> arXiv:2305.18290 [cs]
- [35] Vinu Sankar Sadasivan, Shoumik Saha, Gaurang Sriraman, Priyatham Katakinda, Atoosa Chegini, and Soheil Feizi. 2024. *Fast Adversarial Attacks on Language Models in One GPU Minute*. <https://doi.org/10.48550/arXiv.2402.15570> arXiv:2402.15570 [cs]
- [36] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. *DistilBERT, a Distilled Version of BERT: Smaller, Faster, Cheaper and Lighter*. <https://doi.org/10.48550/arXiv.1910.01108> arXiv:1910.01108 [cs]
- [37] Teven Le Scao and Alexander M. Rush. 2021. *How Many Data Points Is a Prompt Worth?* <https://doi.org/10.48550/arXiv.2103.08493> arXiv:2103.08493 [cs]
- [38] Leo Schwinn, David Dobre, Sophie Xhonneux, Gauthier Gidel, and Stephan Gunnemann. 2024. *Soft Prompt Threats: Attacking Safety Alignment and Unlearning in Open-Source LLMs through the Embedding Space*. <https://doi.org/10.48550/arXiv.2401.04088>

- 2402.09063 arXiv:2402.09063 [cs]
- [39] Xinyue Shen, Zeyuan Chen, Michael Backes, Yun Shen, and Yang Zhang. 2023. "Do Anything Now": Characterizing and Evaluating In-The-Wild Jailbreak Prompts on Large Language Models. arXiv:2308.03825 <https://arxiv.org/abs/2308.03825v2>
- [40] Yundi Shi, Piji Li, Changchun Yin, Zhaoyang Han, Lu Zhou, and Zhe Liu. 2022. PromptAttack: Prompt-based Attack for Language Models via Gradient Search. <https://doi.org/10.48550/arXiv.2209.01882> arXiv:2209.01882 [cs]
- [41] Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts. <https://doi.org/10.48550/arXiv.2010.15980> arXiv:2010.15980 [cs]
- [42] Liwei Song, Xinwei Yu, Hsuan-Tung Peng, and Karthik Narasimhan. 2021. Universal Adversarial Attacks with Natural Triggers for Text Classification. <https://doi.org/10.48550/arXiv.2005.00174> arXiv:2005.00174 [cs]
- [43] Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Goodfellow, and Rob Fergus. 2014. Intriguing Properties of Neural Networks. <https://doi.org/10.48550/arXiv.1312.6199> arXiv:1312.6199 [cs]
- [44] Derek Tam, Rakesh R. Menon, Mohit Bansal, Shashank Srivastava, and Colin Raffel. 2021. Improving and Simplifying Pattern Exploiting Training. <https://doi.org/10.48550/arXiv.2103.11955> arXiv:2103.11955 [cs]
- [45] Eric Wallace, Shi Feng, Nikhil Kandpal, Matt Gardner, and Sameer Singh. 2021. Universal Adversarial Triggers for Attacking and Analyzing NLP. <https://doi.org/10.48550/arXiv.1908.07125> arXiv:1908.07125 [cs]
- [46] Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. Finetuned Language Models Are Zero-Shot Learners. <https://doi.org/10.48550/arXiv.2109.01652> arXiv:2109.01652 [cs]
- [47] Johannes Welbl, Amelia Glaese, Jonathan Uesato, Sumanth Dathathri, John Mellor, Lisa Anne Hendricks, Kirsty Anderson, Pushmeet Kohli, Ben Coppin, and Po-Sen Huang. 2021. Challenges in Detoxifying Language Models. <https://doi.org/10.48550/arXiv.2109.07445> arXiv:2109.07445 [cs]
- [48] Yuxin Wen, Neel Jain, John Kirchenbauer, Micah Goldblum, Jonas Geiping, and Tom Goldstein. 2023. Hard Prompts Made Easy: Gradient-Based Discrete Optimization for Prompt Tuning and Discovery. <https://doi.org/10.48550/arXiv.2302.03668> arXiv:2302.03668 [cs]
- [49] Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2022. An Explanation of In-context Learning as Implicit Bayesian Inference. <https://doi.org/10.48550/arXiv.2111.02080> arXiv:2111.02080 [cs]
- [50] Lei Xu, Yangyi Chen, Ganqu Cui, Hongcheng Gao, and Zhiyuan Liu. 2022. Exploring the Universal Vulnerability of Prompt-based Learning Paradigm. <https://doi.org/10.48550/arXiv.2204.05239> arXiv:2204.05239 [cs]
- [51] Yue Xu and Wenjie Wang. 2024. $\$ \text{textit{LinkPrompt}} \$$: Natural and Universal Adversarial Attacks on Prompt-based Language Models. <https://doi.org/10.48550/arXiv.2403.16432> arXiv:2403.16432 [cs]
- [52] Jiaqi Xue, Mengxin Zheng, Ting Hua, Yilin Shen, Yepeng Liu, Ladislau Boloni, and Qian Lou. 2023. TrojLLM: A Black-box Trojan Prompt Attack on Large Language Models. <https://doi.org/10.48550/arXiv.2306.06815> arXiv:2306.06815 [cs]
- [53] Jiahao Yu, Xingwei Lin, Zheng Yu, and Xinyu Xing. 2023. GPTFUZZER: Red Teaming Large Language Models with Auto-Generated Jailbreak Prompts. <https://doi.org/10.48550/arXiv.2309.10253> arXiv:2309.10253 [cs]
- [54] Ningyu Zhang, Luoqi Li, Xiang Chen, Shumin Deng, Zhen Bi, Chuanqi Tan, Fei Huang, and Huajun Chen. 2022. Differentiable Prompt Makes Pre-trained Language Models Better Few-shot Learners. <https://doi.org/10.48550/arXiv.2108.13161> arXiv:2108.13161 [cs]
- [55] Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. BERTScore: Evaluating Text Generation with BERT. <https://doi.org/10.48550/arXiv.1904.09675> arXiv:1904.09675 [cs.CL]
- [56] Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. Judging LLM-as-a-Judge with MT-Bench and Chatbot Arena. <https://doi.org/10.48550/arXiv.2306.05685> arXiv:2306.05685 [cs]
- [57] Kaijie Zhu, Jindong Wang, Jiaheng Zhou, Zichen Wang, Hao Chen, Yidong Wang, Linyi Yang, Wei Ye, Yue Zhang, Neil Zhenqiang Gong, and Xing Xie. 2023. Prompt-Bench: Towards Evaluating the Robustness of Large Language Models on Adversarial Prompts. <https://doi.org/10.48550/arXiv.2306.04528> arXiv:2306.04528 [cs]
- [58] Andy Zou, Zifan Wang, Nicholas Carlini, Milad Nasr, J. Zico Kolter, and Matt Fredrikson. 2023. Universal and Transferable Adversarial Attacks on Aligned Language Models. <https://doi.org/10.48550/arXiv.2307.15043> arXiv:2307.15043 [cs]

Received 4 June 2024; accepted 28 June 2024