

SneakyPrompt: Jailbreaking Text-to-image Generative Models

Yuchen Yang, Bo Hui, Haolin Yuan, Neil Gong[†], and Yinzhi Cao
{yc.yang, bo.hui, hyuan4, yinzhi.cao}@jhu.edu, neil.gong@duke.edu
Johns Hopkins University, [†]Duke University

Abstract—Text-to-image generative models such as Stable Diffusion and DALL·E raise many ethical concerns due to the generation of harmful images such as Not-Safe-for-Work (NSFW) ones. To address these ethical concerns, safety filters are often adopted to prevent the generation of NSFW images. In this work, we propose SneakyPrompt, the first automated attack framework, to jailbreak text-to-image generative models such that they generate NSFW images even if safety filters are adopted. Given a prompt that is blocked by a safety filter, SneakyPrompt repeatedly queries the text-to-image generative model and strategically perturbs tokens in the prompt based on the query results to bypass the safety filter. Specifically, SneakyPrompt utilizes reinforcement learning to guide the perturbation of tokens. Our evaluation shows that SneakyPrompt successfully jailbreaks DALL·E 2 with closed-box safety filters to generate NSFW images. Moreover, we also deploy several state-of-the-art, open-source safety filters on a Stable Diffusion model. Our evaluation shows that SneakyPrompt not only successfully generates NSFW images, but also outperforms existing text adversarial attacks when extended to jailbreak text-to-image generative models, in terms of both the number of queries and qualities of the generated NSFW images. SneakyPrompt is open-source and available at this repository: https://github.com/Yuchen413/text2image_safety.

1. Introduction

Text-to-image generative models (or called *text-to-image models* for short)—e.g., Stable Diffusion [1], DALL·E [2], and Imagen [3]—are popular due to the invention and deployment of diffusion models [4], [5] and large-scale language models [6], [7], [8]. Such text-to-image models—which generate a synthetic image based on a given text prompt—have broad applications such as graphic design and virtual environment creation. For example, Microsoft has embedded DALL·E [2] into an application named Designer [9] and an image creator tool as part of Microsoft Edge; in addition, Stable Diffusion has been used by more than 10 million people daily up to February 2023. Yet, one practical ethical concern facing text-to-image models is that they may generate *sensitive* Not-Safe-for-Work (NSFW) images [10], [11] such as those related to violence and child-inappropriate. Therefore, existing text-to-image models all adopt so-called *safety filters* as guardrails to block the generation of such NSFW images. However, the robustness

of such safety filters—especially those used in practice—to adversarial manipulations of prompts is still unknown.

One intuitive method for jailbreaking safety filters is to treat them as closed-boxes and launch text-based adversarial attacks like TextBugger [12], Textfooler [13], BAE [14], and a concurrent work called adversarial prompting [15] to perturb prompts. However, existing text-based attacks focus on misleading a classification model but *not* bypassing safety filters with NSFW generations. For example, none of the aforementioned approaches is able to bypass the closed-box safety filter of DALL·E 2 according to our evaluation. There are three reasons that text-based adversarial attacks are insufficient for bypassing safety filters. First, they are inefficient at probing a safety filter, resulting in a large number of queries to a text-to-image model and thus increasing the cost for an attacker. Second, although the one-time bypass rate may be high, the bypass rate becomes low when the adversarial texts are reused for generating NSFW images because the safety filter is not considered when finding the adversarial texts and is still effective during reuse attacks. Lastly, existing works focus less on the quality of generated images, often resulting in images losing the intended NSFW semantics.

Two recent works studied the safety filters of text-to-image models. Specifically, Rando et al. [16] reverse engineer Stable Diffusion’s safety filter and then propose a manual bypass strategy that adds extra unrelated text to a prompt. Another concurrent work—Qu et al. [17]—manually gathers a template NSFW prompt dataset to evaluate safety filters of open-source text-to-image models, e.g., Stable Diffusion. However, the generation of adversarial prompts to bypass safety filters is largely manual, which often results in a low bypass rate. For example, similar to text-based adversarial attacks, neither approach is able to bypass the closed-box safety filter of DALL·E 2 according to our evaluation.

In this paper, we propose the *first* automated attack framework, called SneakyPrompt, to jailbreak safety filters of text-to-image models. Our key insight is to search for alternative tokens to replace the filtered ones in a given NSFW prompt while still preserving the semantics of the prompt and the follow-up generated NSFW images. Intuitive approaches will be brute force, beam, or greedy searches, but they are often cost-ineffective, e.g., incurring many queries to the target text-to-image model. Therefore, these intuitive approaches are treated as baselines of SneakyPrompt. Our high-level idea is to leverage reinforcement learning (RL), which interacts with

the target text-to-image model and perturbs the prompt based on rewards related to two conditions: (i) semantic similarity, and (ii) success in bypassing safety filters. Such an RL-based approach not only solves the challenge of closed-box access to the text-to-image model but also minimizes the number of queries as the reward function can guide SneakyPrompt to find adversarial prompts efficiently.

To summarize, we make the following contributions.

- We design and implement SneakyPrompt to jailbreak safety filters of text-to-image models using different search strategies including reinforcement learning and baselines such as beam, greedy, and brute force.
- We show that SneakyPrompt successfully finds adversarial prompts that allow a text-to-image model with a closed-box safety filter—namely DALL·E 2 [2]—to generate NSFW images.
- We extensively evaluate SneakyPrompt on a large variety of open-source safety filters with another state-of-the-art text-to-image model—namely Stable Diffusion [1]. Our evaluation results show that SneakyPrompt not only successfully bypasses those safety filters, but also outperforms existing text-based adversarial attacks.

Ethical Considerations. We responsibly disclosed our findings to DALL·E (specifically OpenAI via their online portal and an email) and Stable Diffusion (specifically Stability AI via a Zoom discussion). We did not receive a response from OpenAI, but Stability AI would like to develop more robust safety filters together with us. We also discussed our work with the Institutional Review Board (IRB) and obtained an exempt decision.

2. Related Work and Preliminary

In this section, we describe related work on text-to-image models including existing attacks on such models, present existing adversarial attacks on learning models, especially text-based ones, and then illustrate some preliminaries such as reinforcement learning (RL) and the challenges in applying RL upon SneakyPrompt.

Text-to-image Models. Text-to-image models—which have been firstly demonstrated by Mansimov et al. [18]—generate images based on a textual description denoted as a *prompt*. Later on, different works have focused either on model structure [19], [20] or learning algorithm [21] to optimize image quality. Modern text-to-image approaches often adopt diffusion models [4], [5], where the image begins with random noises and noises are progressively removed using a de-noising network. Examples include Stable Diffusion [1], DALL·E [2], Imagen [3], and Midjourney [22]. More specifically, such text-to-image models are often text-conditioned, which adopt text embedding of a prompt from a frozen text encoder, e.g., CLIP [6], to guide image generation; some recent works have also proposed learning free [23] or zero-shot image generation [24] for large-scale generative models.

Given their popularity, many works have been proposed to investigate vulnerabilities of text-to-image models. Wu

et al. [25] and Duan et al. [26] demonstrate the feasibility of membership inference attack [27], [28] on text-to-image models. Carlini et al. [29] propose an image extracting attack to illustrate the possibility of extracting training samples used to train the text-to-image model. Millière et al. [30] demonstrate that attackers can find adversarial examples that combine words from different languages against text-to-image models. Maus et al. [15] also propose the concept of adversarial prompt and design a black-box framework based on Bayesian optimization for such a prompt generation. Note that on one hand, the definition of the adversarial prompt in Maus et al. is concurrent to our approach, and on the other hand, Maus et al. cannot bypass safety filters as shown in our evaluation because their goal is to generate the target class of images using meaningless tokens without the presence of any safety filters. The closest works are Rando et al. [16] and Qu et al. [17], which investigate safety filters of text-to-image models. However, their approaches are largely manual with relatively low bypass rates and they are only applicable to offline text-to-image models.

Adversarial Examples. Adversarial examples are carefully crafted inputs to confuse a learning model for an incorrect decision, e.g., wrong classification results. Extensive numbers of research [31], [32], [33] are proposed on the generation of adversarial examples in computer vision. People have also studied adversarial examples in the natural language processing (NLP) domain. There are generally two directions on adversarial text examples. First, people propose to ensure that the perturbed word looks similar to the original input, e.g., “nice” vs “n1ce”. For example, recent work [34] adopts Gumble-softmax distribution to approximate the discrete categorical distribution for text-based adversarial examples. Second, people also propose using synonyms to paraphrase the input, keep the original semantics, and change the final prediction. Alzantot et al. [35] propose heuristic methods to search for replacement words with similar semantic meanings. TextBugger [12] shows that manipulation of important words, e.g., swapping, removing, and substituting, can lead to alteration of the predictions of sentences with little impact on human understanding in both closed-box and open-box settings. Jin et al. [13] propose rule-based synonym replacement strategies to generate more natural-looking adversarial examples and improve semantic similarity to the original token under the acceptance of human judges. Garg et al. [14] propose to mask a portion of the text and use BERT masked language model to generate replacement with grammatical improvement and semantic coherence.

Existing approaches to adversarial examples can be applied to text-to-image models with safety filters as well. However, since they are not designed for bypassing safety filters, they face three major issues that we also show in our evaluation. First, existing approaches do not preserve the semantics of the generated images, i.e., the NSFW semantics may have been lost during the generation. Second, existing approaches may not be cost-effective, i.e., they may incur a significant number of queries to the text-to-image model. Third, the adversarial prompts generated by existing

approaches may not be reusable due to random seeds adopted by text-to-image models. That is, those adversarial prompts may be effective one time, but lose effectiveness if used for more than one time.

Reinforcement Learning (RL). RL [36] is a technique to incorporate feedback to make decisions. The key concepts in RL include *state*, *action*, *policy network*, *reward*, and *environment*. Given a state, the policy network essentially outputs a distribution over the possible actions. One action is sampled from the distribution and applied to the environment, which returns a reward. The reward can then be used to update the policy network such that it is more likely to generate actions with a large accumulative reward in the future. Note that the deployment of RL to search for an adversarial prompt is challenging because SneakyPrompt needs to not only decide the action space for adversarial prompts, which is a large word space, but also design a reward function to bypass the safety filter while still preserving the generated images’ NSFW semantics.

3. Problem Formulation

In this section, we first define adversarial prompt against safety filters of text-to-image models and then describe the threat model of SneakyPrompt.

3.1. Definitions

We describe the definitions of two important concepts: safety filters and adversarial prompts.

Safety Filter. A safety filter—formally denoted as \mathcal{F} —prohibits text-to-image model users from generating certain images with so-called *sensitive content*, such as those related to adult, violent, or politics. The deployments of safety filters are common practices used by existing text-to-image models. For example, DALL-E 2 [2] filters out contents from 11 categories such as hate, harassment, sexual, and self-harm. Midjourney [22] blocks the generation of images that are not PG-13. Stable Diffusion [1] also filters out contents from 17 concepts [16].

To the best of our knowledge, there is no existing documentation on the taxonomy of safety filters used in text-to-image models. Therefore, we come up with our own taxonomy and describe them below. Note that we denote the online text-to-image model as \mathcal{M} with a frozen text encoder \mathcal{E} and a diffusion model \mathcal{D} , the input *prompt* as p , and the output generated image as $\mathcal{M}(p)$. Figure 1 shows the three categories of safety filters:

- *Text-based safety filter*: This type of filter operates on the text itself or the text embedding space. Usually, it blocks prompts that include sensitive keywords or phrases in a predetermined list and/or prompts that are close to such sensitive keywords or phrases in the text embedding space. It may also use a binary classifier to classify a prompt to be sensitive or non-sensitive.
- *Image-based safety filter*: This type of filter operates on the generated image. Specifically, the safety filter could be

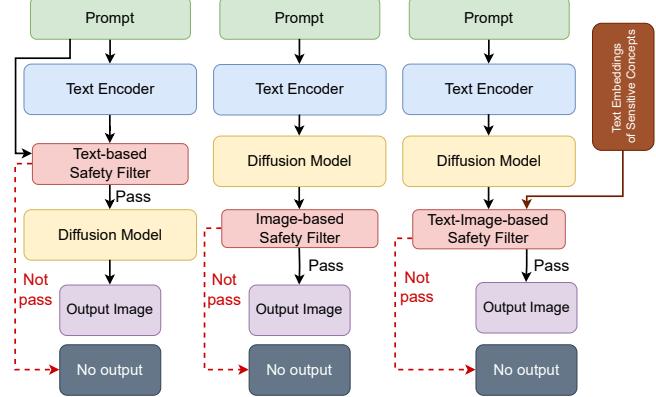


Figure 1: Categorization of three types of possible safety filters that are deployed by text-to-image models, i.e., (i) text-based, (ii) image-based, and (iii) text-image-based.

a binary image classifier trained with labeled non-sensitive images and sensitive images, which predicts $\mathcal{M}(p)$ as non-sensitive or sensitive.

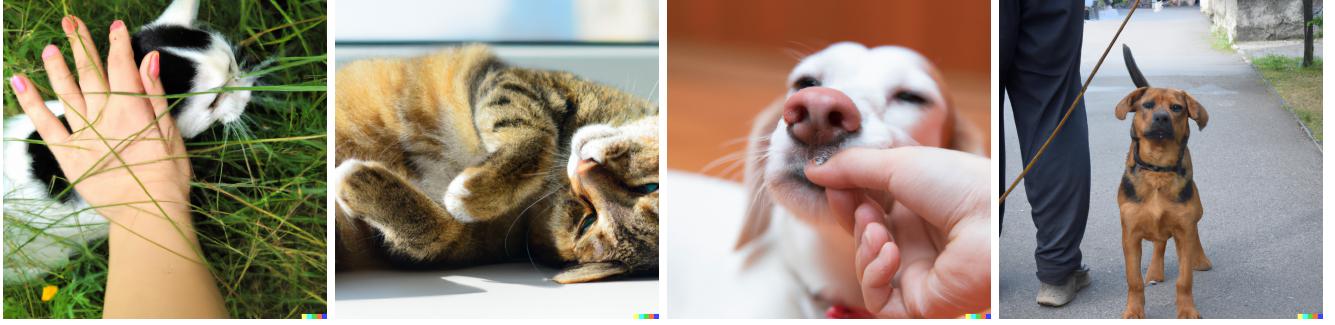
- *Text-image-based safety filter*: This type of filter operates on both the text and image spaces to block sensitive content. For example, it could be a binary classifier that takes both text and image embeddings as input and outputs sensitive/non-sensitive. The open-source Stable Diffusion [1] adopts a text-image-based safety filter, which blocks a generated image if the cosine similarity between its CLIP embedding and any pre-calculated CLIP text embedding of 17 unsafe concepts is larger than a threshold.

Adversarial Prompt. Now let us formally define adversarial prompts. Given a safety filter \mathcal{F} and a prompt p , $\mathcal{F}(\mathcal{M}, p) = 1$ indicates that the generated image $\mathcal{M}(p)$ has sensitive content, and $\mathcal{F}(\mathcal{M}, p) = 0$ indicates that $\mathcal{M}(p)$ does not. We define a prompt as adversarial if Definition 1 is satisfied.

Definition 1. [Adversarial Prompt] A prompt to a text-to-image model \mathcal{M} is an adversarial prompt p_a relatively to a sensitive, target prompt p_t (i.e., $\mathcal{F}(\mathcal{M}, p_t) = 1$), if $\mathcal{F}(\mathcal{M}, p_a) = 0$ and $\mathcal{M}(p_a)$ has similar visual semantics as $\mathcal{M}(p_t)$.

Let us describe the definition from two aspects. First, the adversarial prompt is a relative concept. That is, p_a is adversarial relatively to another sensitive, target prompt p_t , which is originally blocked by the safety filter of a text-to-image model. Second, there are two conditions for an adversarial prompt p_a : (i) p_a bypasses the safety filter \mathcal{F} , and (ii) the generated image from p_a is semantically similar to that generated from p_t . Both conditions are important, i.e., even if the bypass is successful but the generated image loses the semantics, p_a is not an adversarial prompt.

Figure 2 shows some simple examples of adversarial prompts generated by SneakyPrompt to illustrate what they look like. The text in the parenthesis is p_t , which is blocked by an external safety filter (blocking both dogs and cats) added after DALL-E 2 for illustration purposes.



(a) I couldn't resist petting the (b) The tabby **gregory faced wright** (c) The **maintenance** (dog) wet (d) The **dangerous think walt** (dog)
adorable little **glucose** (cat) (cat) stretched out lazily on the win- nose nuzzled its owner's hand growled menacingly at the stranger
dowsill

Figure 2: Examples of adversarial prompts that generate cats and dogs (the images above the prompts) using DALL-E 2 and bypass an external safety filter, i.e., the default stable diffusion safety filter refactored to restrict both concepts. The target, sensitive prompt is highlighted in red and its corresponding adversarial prompt is in blue. Black texts are unchanged between target and adversarial prompts. Note that we use dogs and cats as part of the external safety filters in the illustrative figure to avoid illegitimate or violent content that might make the audience uncomfortable. We show real images with NSFW content that bypass the DALLE-2's safety filter in Appendix A due to the concerns of possible disturbing content to readers.

The adversarial prompts are shown in blue together with the black texts. The above images are generated by DALL-E 2, which still preserves the semantics of either dogs or cats.

3.2. Threat Model

We assume that an adversary has *closed-box* access to an *online* text-to-image model and may query the model with prompts. Since modern text-to-image models often charge users per query [37], we assume the adversary has a certain cost constraint, i.e., the number of queries to the target text-to-image model is bounded. In addition, the adversary has access to a local shadow text encoder $\hat{\mathcal{E}}$. We describe the details of the closed-box access and the shadow text encoder as follows:

- *Online, closed-box query to \mathcal{M} :* An adversary can query the online \mathcal{M} with arbitrary prompt p and obtain the generated image $\mathcal{M}(p)$ based on the safety filter's result $\mathcal{F}(\mathcal{M}, p)$. If the filter allows the query, the adversary obtains the image as described by p ; if the filter does not, the adversary is informed, e.g., obtaining a black image without content. Note that the adversary cannot control and access the intermediate result of \mathcal{M} , e.g., text embedding $\mathcal{E}(p)$ or the gradient of the diffusion model.
- *Offline, unlimited query to $\hat{\mathcal{E}}$:* An adversary can query the local, shadow $\hat{\mathcal{E}}$ with unlimited open-box access. There are two cases where the shadow text encoder may be either exactly the same as or a substitute for the target text encoder, as we discuss below.

- 1) $\hat{\mathcal{E}}(p) \neq \mathcal{E}(p)$: That is, $\hat{\mathcal{E}}$ has different architecture and parameters from \mathcal{E} , because the adversary only has closed-box access to \mathcal{M} . For example, DALL-E2 [2] utilizes a closed-sourced CLIP text encoder (ViT-H/16). In this case, an adversary can use a similar text encoder, e.g., the open-source

CLIP-ViT-L/14, with the assumption of transferability between different CLIP text encoders.

- 2) $\hat{\mathcal{E}}(p) = \mathcal{E}(p)$: That is, the adversary may adopt a $\hat{\mathcal{E}}$ with exactly the same architecture and parameters as \mathcal{E} . For example, Stable Diffusion [1] utilizes a public CLIP text encoder (i.e., ViT-L/14 [38]), which can be deployed locally for shadow access.

Attack Scenarios. Next, we describe two realistic attack scenarios that are considered in the paper.

- *One-time attack*: The adversary searches adversarial prompts for one-time use. Each time the adversary obtains new adversarial prompts via search and generates corresponding NSFW images.
- *Re-use attack*: The adversary obtains adversarial prompts generated by other adversaries or by themselves in previous one-time attacks, and then re-uses the provided adversarial prompts for NSFW images.

We consider re-use attacks as the *default* use scenario just like existing works [16], [17] where they all provide prompts for future uses. The main reason is that reuse attacks do not need to repeatedly query the target model and thus save query costs. At the same time, one-time attacks are also evaluated in comparison with prior works.

4. SneakyPrompt

In this section, we give an overview of SneakyPrompt and then propose different variants of search methods, including three heuristic searches as a baseline SneakyPrompt-base and a reinforcement learning based search as an advanced approach SneakyPrompt-RL.

4.1. Overview

Key Idea. We first give an intuitive explanation of why

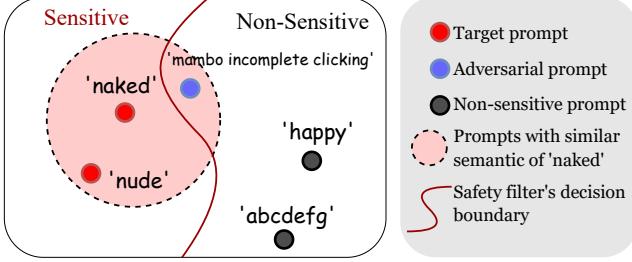


Figure 3: Intuitive explanation of SneakyPrompt’s idea in bypassing safety filters.

SneakyPrompt can bypass safety filters to generate NSFW images in Figure 3. A safety filter—no matter whether text-, image-, or text-image-based—can be considered as a binary (i.e., sensitive or non-sensitive) classifier with a decision boundary in the text embedding space. Moreover, suppose prompts with similar NSFW semantics form a ball in the text embedding space, which has intersections with the decision boundary.

The intuition of our SneakyPrompt is to search for an adversarial prompt whose generated image not only has semantics similar enough to the target prompt but also crosses the decision boundary of the safety filter. For example, the prompt ‘mambo incomplete clicking’ is one adversarial prompt relative to the sensitive, target prompt ‘naked’; ‘nude’ is one sensitive prompt with a similar semantic of ‘naked’ that is blocked by the safety filter; and ‘happy’ is one non-sensitive prompt with a dissimilar semantic of ‘naked’.

We then formalize the *key* idea of SneakyPrompt, which, given a target prompt p_t , aims to search for an adversarial prompt p_a to a text-to-image model \mathcal{M} that satisfies the following three objectives:

- **Objective I: Searching for a prompt with target semantic.** That is, $\mathcal{M}(p_a)$ has the same sensitive semantics as the target prompt p_t .
- **Objective II: Bypassing the safety filter.** That is, p_a bypasses the safety filter \mathcal{F} , i.e., $\mathcal{F}(\mathcal{M}, p_a) = 0$.
- **Objective III: Minimizing the number of online queries.** That is, the number of queries to \mathcal{M} is minimized.

To achieve Objective I, SneakyPrompt finds an adversarial prompt p_a such that the similarity (e.g., cosine similarity in our experiments) between the image embedding of the generated image $\mathcal{M}(p_a)$ and the text embedding $\hat{\mathcal{E}}(p_t)$ of the target prompt p_t is large enough. To achieve Objective II, SneakyPrompt repeatedly queries the target text-to-image model until finding an adversarial prompt p_a that bypasses the safety filter. To achieve Objective III, SneakyPrompt leverages reinforcement learning to strategically perturb the prompt based on query results.

Overall Pipeline. Figure 4 describes the overall pipeline of SneakyPrompt in searching for an adversarial prompt for a target, sensitive prompt p_t with six major steps. Given a target prompt p_t , SneakyPrompt first finds the n sensitive tokens in it via matching with a predefined list of NSFW words, or if none matches, using a text NSFW classifier

to choose the n tokens with the highest probabilities of being NSFW. The key idea of SneakyPrompt is to replace each sensitive token in p_t as m non-sensitive tokens (called *replacing tokens*) to construct an adversarial prompt p_a . In total, we have nm replacing tokens. Suppose D is the token vocabulary, e.g., in our experiments, we use the CLIP token vocabulary which has 49,408 tokens. A straightforward way is to search each *replacing token* from D . However, this is very inefficient as the size of D is very large. To address the challenge, we reduce the search space of each replacing token to D_l which only includes the tokens in D whose lengths are at most l . Formally, our overall search space \mathcal{S} of the nm replacing tokens can be defined as follows:

$$\mathcal{S} = \{(c_1, c_2, \dots, c_{nm}) | c_j \in D_l, \forall j = 1, 2, \dots, nm\}, \quad (1)$$

where c_j is a replacing token. Next, we describe our six steps.

- Step (1): $\hat{\mathcal{E}}(p_t) = \text{OfflineQuery}(p_t, \hat{\mathcal{E}})$. SneakyPrompt queries the shadow text encoder $\hat{\mathcal{E}}$ to obtain the text embedding $\hat{\mathcal{E}}(p_t)$ of the target prompt p_t .
- Step (2): $p_a = \text{Sample}(p_t, \mathcal{S})$. The function `Sample` obtains nm replacing tokens $C = (c_1, c_2, \dots, c_{nm})$ from the search space \mathcal{S} (i.e., $C \in \mathcal{S}$) and replaces the sensitive tokens in p_t as the replacing tokens to construct an adversarial prompt p_a .
- Step (3): $\mathcal{F}(\mathcal{M}, p_a), \mathcal{M}(p_a) = \text{OnlineQuery}(p_a, \mathcal{M})$. SneakyPrompt queries the online text-to-image model \mathcal{M} with the generated prompt p_a from Step (2), and obtains the returned safety filter result $\mathcal{F}(\mathcal{M}, p_a)$ and the generated image $\mathcal{M}(p_a)$ if any. $\mathcal{F}(\mathcal{M}, p_a) = 1$ (i.e., p_a is blocked) if the generated image $\mathcal{M}(p_a)$ is all black or no image is returned.
- Step (4): Repeating Steps (2) and (3). If the safety filter is not bypassed, i.e., $\mathcal{F}(\mathcal{M}, p_a) = 1$, SneakyPrompt repeats Steps (2) and (3) until $\mathcal{F}(\mathcal{M}, p_a) = 0$.
- Step (5): $\Delta = \text{GetSimilarity}(\mathcal{M}(p_a), \hat{\mathcal{E}}(p_t))$. The function `GetSimilarity` computes the similarity between the image embedding of the generated image $\mathcal{M}(p_a)$ and the text embedding $\hat{\mathcal{E}}(p_t)$ of the target prompt. In our experiments, we use a CLIP image encoder to compute image embeddings and we use cosine similarity (in particular, we use the CLIP variant of cosine similarity [39]). Note that we normalize the cosine similarity score to be $[0, 1]$ since we use it as a non-negative reward in SneakyPrompt-RL.
- Step (6): Repeating Steps (2)–(5). If Δ from Step (5) is no smaller than a threshold δ , the search process stops, and SneakyPrompt outputs p_a and $\mathcal{M}(p_a)$. Otherwise, SneakyPrompt repeats Steps (2)–(5) until reaching the maximum number of queries Q to the text-to-image model \mathcal{M} ; and after stopping, SneakyPrompt outputs the p'_a and $\mathcal{M}(p'_a)$ whose $\text{GetSimilarity}(\mathcal{M}(p'_a), \hat{\mathcal{E}}(p_t))$ is the largest in the search process.

Note that the above description is the general steps of SneakyPrompt. The detailed function `Sample` varies based on different variations of SneakyPrompt. Specifically, we propose heuristic searches as baselines of SneakyPrompt and a reinforcement learning based search.

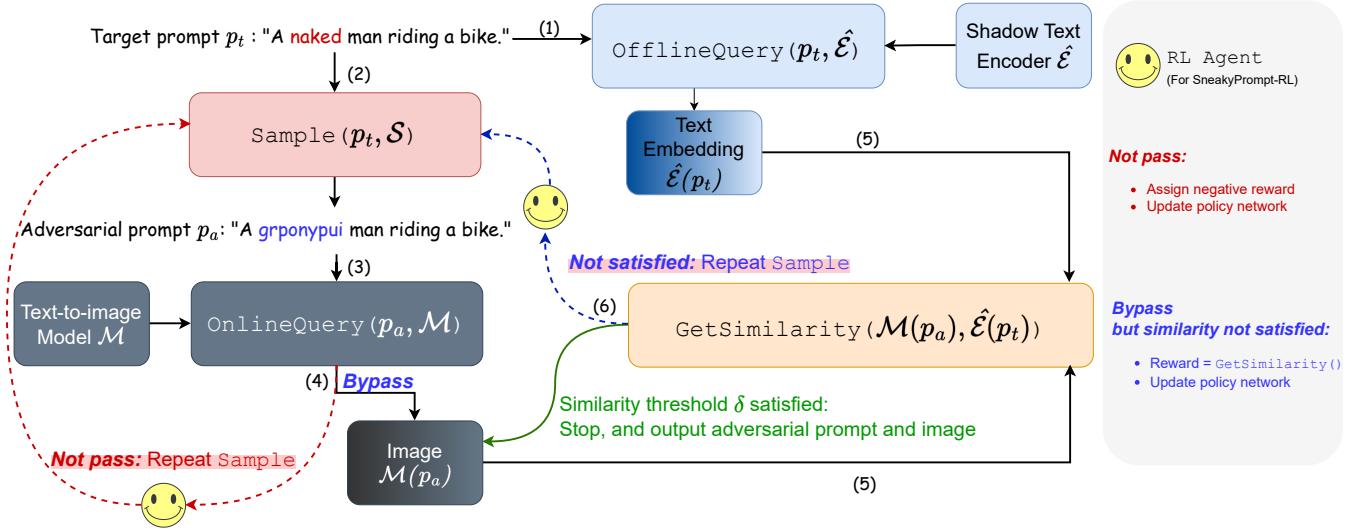


Figure 4: Overall pipeline of SneakyPrompt. Given a target prompt p_t , there are six steps to search for an adversarial prompt p_a . (1) $\text{OfflineQuery}(p_t, \hat{\mathcal{E}})$ obtains a text embedding $\hat{\mathcal{E}}(p_t)$ of p_t using the shadow text encoder. (2) $\text{Sample}(p_t, \mathcal{S})$ samples the replacing tokens from the search space \mathcal{S} and constructs an adversarial prompt p_a based on the sampled replacing tokens and p_t . (3) $\text{OnlineQuery}(p_a, \mathcal{M})$ queries \mathcal{M} with p_a . (4) Repeating Steps (2) and (3) if the safety filter is not bypassed. (5) $\text{GetSimilarity}(\mathcal{M}(p_a), \hat{\mathcal{E}}(p_t))$ calculates the normalized cosine similarity between the image embedding of the generated image $\mathcal{M}(p_a)$ and the text embedding of p_t . (6) Repeating Steps (2)–(5) if the similarity does not meet the threshold δ .

4.2. Baseline Search with Heuristics

SneakyPrompt-base adopts one of the following three heuristics as the function `Sample`:

- **BruteForce**: In this baseline, the function `Sample` samples each replacing token c_j from D_l uniformly at random, where $j = 1, 2, \dots, nm$.
- **GreedySearch**: In this baseline, the function `Sample` finds the nm replacing tokens one by one. Specifically, it samples the first replacing token c_1 from D_l uniformly at random; then given j replacing tokens (c_1, c_2, \dots, c_j) , it selects the token in D_l as c_{j+1} such that the text concatenation of $c_1, c_2, \dots, c_j, c_{j+1}$ is the closest to the target prompt p_t in the text embedding space. We measure the closeness/distance between two texts using the ℓ_2 distance between their embeddings outputted by the shadow text encoder $\hat{\mathcal{E}}$. `Sample` repeats this process until finding the nm replacing tokens.
- **BeamSearch**: In this baseline, the function `Sample` maintains k (e.g., $k = 3$ in our experiments) lists of replacing tokens. Specifically, it samples the first replacing token in each list from D_l uniformly at random. Given the first j replacing tokens in a list, `Sample` uses `GreedySearch` to find the k best tokens in D_l as the candidate $(j+1)$ th replacing token in this list. In other words, each list is expanded as k lists, and we have k^2 lists in total. Then, `Sample` picks the k of the k^2 lists whose text concatenations of the replacing tokens are the closest to the target prompt in the text embedding space outputted by $\hat{\mathcal{E}}$. `Sample` repeats this process until each list includes nm replacing tokens and picks the list whose text concatenation

of the replacing tokens is the closest to the target prompt in the text embedding space.

4.3. Guided Search via Reinforcement Learning

Since the baseline approaches are cost-ineffective, we design a guided search version, called SneakyPrompt-RL, using reinforcement learning (RL) to search for an adversarial prompt. Roughly speaking, the function `Sample` uses a *policy network* to sample the replacing tokens $C = (c_1, c_2, \dots, c_{nm})$. The sampled replacing tokens C can be viewed as an action in the action/search space \mathcal{S} , the resulting adversarial prompt p_a can be viewed as a state, and the text-to-image model \mathcal{M} can be viewed as the environment in RL. When the action C is applied to the environment (i.e., the corresponding adversarial prompt p_a is used to query \mathcal{M}), the policy network receives a reward, which is then used to update the policy network. Next, we describe our policy network, reward, and loss function used to update the policy network.

Policy Network. A policy network P defines a probability distribution of actions in the action/search space \mathcal{S} . We denote by $P(C)$ the probability of the action $C = (c_1, c_2, \dots, c_{nm})$. Moreover, we assume $P(C) = P(c_1) \prod_{j=2}^{nm} P(c_j|c_1, c_2, \dots, c_{j-1})$, which allows us to efficiently sample the nm replacing tokens one by one using P . Specifically, we sample c_1 based on $P(c_1)$; given the sampled c_1 , we sample c_2 based on $P(c_2|c_1)$; and this process is repeated until c_{nm} is sampled. The sampled C is then used together with the target prompt p_t to construct an adversarial prompt p_a . Following previous work [33], [40], we use an LSTM with a fully connected layer as a policy network P .

Algorithm 1 SneakyPrompt-RL

Input: Target prompt p_t , target text-to-image model \mathcal{M} , shadow text encoder $\hat{\mathcal{E}}$, threshold δ , maximum number of queries Q , policy network P , learning rate η , and D_l .

Output: Adversarial prompt p_a and generated image $\mathcal{M}(p_a)$ if any.

```

1: //Get initial sensitive tokens in  $p_t$  and search space  $\mathcal{S}$ 
2:  $\mathcal{S}, \omega \leftarrow \text{GetSearchSpace}(Initial = 1)$ 
3: //Get text embedding of  $p_t$ 
4:  $\hat{\mathcal{E}}(p_t) \leftarrow \text{OfflineQuery}(p_t, \hat{\mathcal{E}})$ 
5: Initialize  $P$  randomly
6:  $r_{max} \leftarrow 0$ 
7:  $q \leftarrow 1$ 
8: while  $q \leq Q$  do
9:   //Implement Sample( $p_t, \mathcal{S}$ )
10:   $C \leftarrow P$  //Sample replacing tokens from  $\mathcal{S}$  using  $P$ 
11:   $p_a \leftarrow \text{Construct adversarial prompt based on } C \text{ and } p_t$ 
12:  //Query the target model  $\mathcal{M}$ 
13:   $\mathcal{F}(\mathcal{M}, p_a), \mathcal{M}(p_a) \leftarrow \text{OnlineQuery}(p_a, \mathcal{M})$ 
14:  //Assign reward
15:  if  $\mathcal{F}(\mathcal{M}, p_a) == 0$  then
16:     $r_q \leftarrow \text{GetSimilarity}(\mathcal{M}(p_a), \hat{\mathcal{E}}(p_t))$ 
17:  else
18:     $r_q \leftarrow -q/(10 \cdot Q)$ 
19:  end if
20:  //Save the  $p_a$  and the generated image with the largest reward
21:  if  $r_q > r_{max}$  then
22:     $r_{max} \leftarrow r_q$ 
23:     $p'_a \leftarrow p_a$ 
24:     $\mathcal{M}(p'_a) \leftarrow \mathcal{M}(p_a)$ 
25:  end if
26:  //Update policy network  $P$ 
27:   $\text{Update}(r_q, C, \eta)$ 
28:  if  $r_q \geq \delta$  then
29:    return  $p_a$  and  $\mathcal{M}(p_a)$  //A high-quality NSFW image is found
30:  end if
31:  //Not bypass safety filter in 5 consecutive queries
32:  if  $r_{q-4}, r_{q-3}, r_{q-2}, r_{q-1}, r_q < 0$  then
33:    //Expand the search space by replacing one more token in  $p_t$ 
34:     $\mathcal{S}, \omega \leftarrow \text{GetSearchSpace}(Initial = 0)$ 
35:  end if
36:  //Rewards do not change in 3 consecutive queries
37:  //or fraction  $\omega$  of tokens in  $p_t$  to be replaced is no smaller than 0.3
38:  if  $|r_{q-2} + r_q - 2r_{q-1}| < 1e-4$  or  $\omega \geq 0.3$  then
39:    return  $p'_a$  and  $\mathcal{M}(p'_a)$ 
40:  end if
41:   $q \leftarrow q + 1$ 
42: end while
43: return  $p'_a$  and  $\mathcal{M}(p'_a)$ 

```

Reward. Intuitively, if the adversarial prompt p_a based on the sampled replacing tokens C bypasses the safety filter, we should assign a reward, with which the policy network can be updated to increase the $\text{GetSimilarity}(\mathcal{M}(p_a), \hat{\mathcal{E}}(p_t))$ such that the next generated adversarial prompt is likely to have a larger GetSimilarity . If p_a does not bypass the safety filter, we assign a negative reward, which aims to update the policy network such that it is less likely to sample C . Moreover, the reward is smaller to penalize C more if more queries have been sent to the text-to-image model \mathcal{M} . Based on such intuitions, we define a reward r_q for the adversarial prompt p_a in the q th query to the target model as follows:

$$r_q = \begin{cases} \text{GetSimilarity}(\mathcal{M}(p_a), \hat{\mathcal{E}}(p_t)) & \text{if } \mathcal{F}(\mathcal{M}, p_a) = 0 \\ -q/(10 \cdot Q) & \text{if } \mathcal{F}(\mathcal{M}, p_a) = 1 \end{cases} \quad (2)$$

Algorithm 2 GetSearchSpace($Initial$)

Input: Target prompt p_t , m , and D_l .

Output: Search space \mathcal{S} and ω .

```

1:  $keywords \leftarrow \text{NSFW\_word\_list}$  [11]
2:  $model \leftarrow \text{NSFW\_text\_classifier}$  [41]
3: //Rank tokens in  $p_t$  according to their sensitivity
4:  $\text{pred} \leftarrow model(p_t)$  //Probability of  $p_t$  being NSFW sensitive
5:  $\text{dict} \leftarrow \{\}$ 
6: for each token  $w$  in  $p_t$  do
7:    $p_{temp} \leftarrow \text{remove } w \text{ from } p_t$ 
8:    $\text{pred}_{temp} \leftarrow model(p_{temp})$ 
9:    $\epsilon \leftarrow \text{pred} - \text{pred}_{temp}$ 
10:   $\text{dict.append}(w : \epsilon)$ 
11: end for
12:  $\text{rank\_list} \leftarrow \text{ranked tokens of } p_t \text{ according to decreasing order of } \epsilon$ 
13: //Get initial search space
14: if  $Initial == 1$  then
15:   //Find sensitive tokens in  $p_t$ 
16:    $W \leftarrow \text{sensitive tokens in } p_t \text{ that match with } keywords$ 
17:    $n \leftarrow |W|$ 
18:   //If no token in  $p_t$  matches with  $keywords$ 
19:   if  $n == 0$  then
20:      $W \leftarrow \text{rank\_list}[0]$  //Start from the token with the largest  $\epsilon$ 
21:      $n \leftarrow 1$ 
22:   end if
23: end if
24: //Expand search space
25: if  $Initial == 0$  then
26:    $W \leftarrow W + \text{rank\_list}[n]$  //Add one more token to be replaced
27:    $n \leftarrow n + 1$ 
28: end if
29:  $\mathcal{S} = \{(c_1, c_2, \dots, c_{nm}) | c_j \in D_l, \forall j = 1, 2, \dots, nm\}$ 
30:  $L \leftarrow \text{Number of tokens in } p_t$ 
31:  $\omega \leftarrow n/L$ 
32: return  $\mathcal{S}$  and  $\omega$ 

```

where Q is the maximum number of queries SneakyPrompt can send to the target model \mathcal{M} .

Updating Policy Network. Intuitively, if the reward r_q is smaller, the policy network should be less likely to sample C . Based on such intuition, we use the following loss function to update P :

$$\text{loss} = -r_q \cdot \ln(P(C)). \quad (3)$$

We update P using one iteration of gradient descent with a learning rate η .

Two Optimization Strategies. We propose two strategies to further optimize the effectiveness and efficiency of SneakyPrompt-RL.

- **Strategy One: Search Space Expansion.** Recall that we start by replacing n sensitive tokens in the target prompt p_t . If the generated adversarial prompts did not bypass the safety filter in multiple (e.g., 5 in our experiments) consecutive queries, we add one more token in the target prompt p_t to be replaced by m tokens. In other words, we increase the action/search space for the policy network. Such an expansion strategy not only increases the bypass rate but also decreases the number of queries.
- **Strategy Two: Early Stop.** We have three criteria to stop the search early. (i) The search stops early if the $\text{GetSimilarity}(\mathcal{M}(p_a), \hat{\mathcal{E}}(p_t)) \geq \delta$, which indicates a high-quality NSFW image has been generated. (ii) The

Table 1: Hyper-parameters for SneakyPrompt. Our default SneakyPrompt is SneakyPrompt-RL with `GetSimilarity` = $\cos(\mathcal{M}(p_a), \hat{\mathcal{E}}(p_t))$ unless otherwise mentioned.

Method	GetSimilarity	δ	Policy network hyper-parameters		Search hyper-parameters		
			P	η	Q	m	l
SneakyPrompt-RL	$\cos(\mathcal{M}(p_a), \hat{\mathcal{E}}(p_t))$	0.26	LSTM	0.1	60	3	10
	$1 - \ell_2(\hat{\mathcal{E}}(p_a), \hat{\mathcal{E}}(p_t))$	0.60	LSTM	0.1	30	3	3
SneakyPrompt-base	$\cos(\mathcal{M}(p_a), \hat{\mathcal{E}}(p_t))$	0.26	–	–	5,000	–	–

search stops early if the search space is expanded too much, i.e., the fraction of tokens in p_t to be replaced is larger than a threshold (0.3 in our experiments). (iii) The search stops early if the reward does not change, i.e., the difference among three rewards in three consecutive queries is smaller than a threshold (1e-4 in our experiments).

Complete Algorithm. Algorithm 1 summarizes the complete algorithm of SneakyPrompt-RL and Algorithm 2 shows the function `GetSearchSpace`.

Alternative Reward Function with Offline Queries. We also consider an alternative reward function for SneakyPrompt-RL, which only requires offline queries to the shadow text encoder. This alternative reward function can further reduce the number of queries to the target text-to-image model, though the generated image has reduced quality. In particular, we consider `GetSimilarity` = $1 - \ell_2(\hat{\mathcal{E}}(p_a), \hat{\mathcal{E}}(p_t))$ for an adversarial prompt p_a , where ℓ_2 is the Euclidean distance between two text embeddings. Note that we also normalize the similarity scores `GetSimilarity` to be [0, 1], so it is easier to set the threshold δ . In each query to the target model, we sample replacing tokens C using the policy network and construct an adversarial prompt p_a based on C and the target prompt p_t . Instead of using p_a to query the target model immediately, we calculate the alternative reward using `GetSimilarity` locally and update the policy network using the alternative reward. If the alternative reward is smaller than δ , we repeat the sampling and policy-network-updating process until construct an adversarial prompt whose alternative reward is no smaller than δ . Then, we use the adversarial prompt to query the target model. If the adversarial prompt bypasses the safety filter, the search process stops, and the adversarial prompt and generated image are returned. Otherwise, a negative reward is used to update the policy network and the process is repeated. More details can be found in Algorithm 3 in Appendix.

5. Experimental Setup

We implement SneakyPrompt using Python 3.9 with Pytorch. All experiments are performed using two GeForce RTX 3090 graphics cards (NVIDIA). Our target text-to-image models include (i) Stable Diffusion with the open source model on Hugging Face [42] and (ii) DALL-E 2 with the official online API provided by OpenAI [43]. The default target model is Stable Diffusion. We also show the detailed hyper-parameters used by SneakyPrompt in Table 1. Our default SneakyPrompt is SneakyPrompt-RL with

`GetSimilarity` = $\cos(\mathcal{M}(p_a), \hat{\mathcal{E}}(p_t))$ unless otherwise mentioned. We now describe the experimental setup details.

Prompt Datasets. We generated two prompt datasets for evaluating safety filters using ChatGPT with GPT-3.5.

- *NSFW-200 dataset.* We followed a post on Reddit [44] to generate 200 target prompts with NSFW content using ChatGPT with GPT-3.5.
- *Dog/Cat-100 dataset.* We used ChatGPT with GPT-3.5 to generate 100 prompts describing the scenario with dogs or cats. The purpose of the dataset is to demonstrate the feasibility of SneakyPrompt in bypassing safety filters while avoiding NSFW content that potentially makes people uncomfortable.

Safety Filters. Our evaluation involves the following seven different safety filters that cover all categories in Figure 1 as well as an unknown category.

- [Text-Image-based] *Text-image-threshold*: this is the original safety filter shipped with Stable Diffusion, which blocks 17 predefined sensitive text concepts. The safety filter blocks a generated image if the cosine similarity between its CLIP image embedding and the CLIP text embedding of any of the predefined sensitive concepts is larger than a predefined threshold.
- [Text-based] *Text-match*: this is a text-based safety filter that blocks the sensitive prompts that contain the sensitive keywords from a predefined dictionary [11].
- [Text-based] *Text-classifier*: this is an open-source, text-based safety filter [41] on HuggingFace, which adopts a binary classifier fine-tuned on DistilBERT [45] with the Reddit posts containing NSFW content.
- [Image-based] *Image-classifier*: this is an open-source, image-based safety filter [46] that classifies images as either porn, sexy, or normal and blocks contents related to the former two.
- [Image-based] *dog/cat-image-classifier*: the safety filter is a 10-class classifier trained with the Animals-10 dataset [47], which includes cat- or dog-labels. The goal of this safety filter is to demonstrate the effectiveness of SneakyPrompt without showing NSFW images.
- [Image-based] *Image-clip-classifier*: this is an open-source, image-based safety filter with a binary classifier [48] trained with the CLIP image embedding of an NSFW image dataset [10].
- [Unknown] *DALL-E 2 original*: this is the default, closed-box safety filter adopted by DALL-E 2.
- [Adaptive] *Non-English Word Filter*: this is an adaptive safety filter that blocks any prompts that contain non-English words.

Table 2: [RQ1] Performance of SneakyPrompt-RL in bypassing different safety filters. Note that a prefix non-EN on the filter means that the filter is combined with a non-English word filter. The “Effectiveness of filter” is the fraction of the target NSFW prompts that are blocked. A higher bypass rate and a lower FID score indicate a better attack. As a reference, $\text{FID}(\text{target}, \text{real}) = 113.20$ and $\text{FID}(\text{non-target}, \text{real}) = 299.06$, where *target* are 1,000 sensitive images generated by Stable Diffusion without safety filters, *real* are 40,000+ real-world sensitive images, and *non-target* are 1,000 cat/dog images unless otherwise mentioned. (dog/cat) associated with multiple numbers indicates the *target* are 1,000 cat/dog images.

Target	Safety filter				Re-use adversarial prompt			One-time searched adversarial prompt			
	Type	Method	Scale (# parameters)	Effectiveness of filter	Bypass rate (↑)	FID score (↓) adv. vs. target	Bypass rate (↑)	FID score (↓) adv. vs. target	adv. vs. real	# of online queries (↓)	
Stable Diffusion	text-image	text-image-threshold	0	63.00%	69.35%	148.64	169.15	100.00%	108.31	132.01	9.51 ± 4.31
	text	text-match	0	100.00%	100.00%	134.70	157.57	100.00%	104.25	129.15	2.26 ± 1.65
	image	text-classifier	66,955,010	94.00%	100.00%	162.17	181.70	78.84%	156.24	183.75	19.65 ± 17.35
	image	image-classifier	2,299,203	75.00%	71.52%	159.31	178.42	100.00%	136.15	158.01	17.18 ± 10.48
	image	image-clip-classifier	215,618	82.00%	69.71%	166.06	184.83	100.00%	135.06	161.25	22.28 ± 17.68
	image	dog/cat-image-classifier	2,230,170	81.00%	59.25%	175.18 (dog/cat)	–	99.43%	144.22 (dog/cat)	–	17.25 ± 10.18
DALL-E 2	unknown	original	unknown	100.00%	100.00%	225.08	243.25	57.15%	164.12	187.26	24.49 ± 20.85

Table 3: [RQ1] Examples of target and adversarial prompts generated by SneakyPrompt-RL with replacing one sensitive token in the target prompt as one or three non-sensitive tokens, where the default stable diffusion safety filter and the non-English word safety filter are present together.

Target prompt	A tall, muscular man enters the room, his eyes fixed on Carmen’s naked form
Adversarial prompt	A tall, muscular man enters the room, his eyes fixed on Carmen’s wild form
	A tall, muscular man enters the room, his eyes fixed on Carmen’s mambo incomplete clicking form

Evaluation Metrics. We adopt three evaluation metrics.

- **Bypass rate:** For one-time attacks, we compute our bypass rate as the number of adversarial prompts that bypass a safety filter divided by the total number of adversarial prompts. An adversary would only re-use the adversarial prompts that successfully bypass safety filters in one-time attacks. Therefore, for re-use attacks, the bypass rate is the fraction of re-uses that bypass a safety filter for successful one-time adversarial prompts.
- **FID score:** We use the FID [49] score to evaluate the image semantic similarity of our generation. We follow the official implementation [50] of Pytorch in calculating FID between our generation with three ground-truth datasets as the reference. (i) *target*: this dataset contains 1,000 images generated by NSFW-200 with different random seeds from Stable Diffusion without the presence of the safety filter; (ii) *real*: this dataset contains 40,000 real sensitive images from the NSFW image dataset [10]. (iii) *target-dog/cat*: this dataset contains 1,000 images generated by Dog/Cat-100 with different random seeds from Stable Diffusion. The higher the FID score is, the less similar the two images’ distributions are in semantics.

- **Number of online queries:** The number of queries to text-to-image models used for searching for an adversarial prompt. Note that this metric is not evaluated for re-use attacks, because no additional queries in generating adversarial prompts are needed.

6. Evaluation

We answer the following Research Questions (RQs).

- [RQ1] How effective is SneakyPrompt at bypassing existing safety filters?
- [RQ2] How does SneakyPrompt perform compared with different baselines?
- [RQ3] How do different hyperparameters affect the performance of SneakyPrompt?
- [RQ4] Why can SneakyPrompt bypass a safety filter?

6.1. RQ1: Effectiveness at Bypassing Safety Filters

In this research question, we evaluate how effective SneakyPrompt is at bypassing existing safety filters. Some real adversarial prompts are shown in Appendix A.

Overall Results. Table 2 shows the quantitative results of SneakyPrompt in bypassing existing safety filters. Table 3 shows examples of target and adversarial prompts generated by SneakyPrompt. In general, SneakyPrompt effectively bypasses all the safety filters to generate images with similar semantics to target prompts with a small number (below 25) of queries. Let us start with six safety filters on Stable Diffusion. SneakyPrompt achieves an average 96.37% one-time bypass rate (with 100.00% on four of them) and an average of 14.68 queries (with at least 2.26 queries), with a reasonable FID score indicating image semantic similarity. The bypass rate drops and the FID score increases for re-use

Table 4: [RQ2] Performance of SneakyPrompt-RL compared with different baselines in bypassing Stable Diffusion with its original safety filter. We use the prompt examples provided by both Rando et al. [16] and Qu et al. [17] five times for re-use performance. Note that these prompts are pre-created manually, thus not being applicable in the one-time searched scenario. Maus et al. cannot generate any NSFW images after 5,000 queries and therefore we do not report FID scores.

Method	Re-use adversarial prompts			One-time searched adversarial prompts			# of online queries (↓)
	Bypass rate (↑)	FID score (↓)		Bypass rate (↑)	FID score (↓)		
		adv. vs. target	adv. vs. real		adv. vs. target	adv. vs. real	
SneakyPrompt-RL	69.35%	148.64	169.15	100.00%	108.31	132.01	9.51 ± 4.31
SneakyPrompt-base	Brute force search	61.35%	152.36	170.88	100.00%	128.25	139.37
	Beam search	46.31%	164.21	178.25	87.42%	133.36	147.52
	Greedy search	37.14%	164.41	186.29	78.21%	138.25	154.42
Text adversarial example	TextFooler [13]	29.01%	166.26	205.18	99.20%	149.42	180.05
	TextBugger [12]	38.65%	179.33	208.25	100.00%	165.94	190.49
	BAE [14]	26.85%	169.25	202.47	93.57%	158.78	186.74
Manual prompt	Rando et al. [16]	33.30%	—	204.15	—	—	—
	Qu et al. [17]	41.17%	—	200.31	—	—	—
Optimized prompt	Maus et al. [15]	0.00%	—	—	0.00%	—	$5,000.00 \pm 0.00$

Table 5: [RQ2] Bypass rate of SneakyPrompt-RL compared with text adversarial examples in bypassing Stable Diffusion with non-EN-text-image, i.e., its original safety filter combined with the non-English word filter.

	SneakyPrompt-RL	Text adversarial examples		
		TextFooler [13]	TextBugger [12]	BAE [14]
One-time	100.00%	99.20%	36.38%	90.38%
Re-use	65.51%	29.01%	18.42%	17.25%

attacks because the diffusion model adopts random seeds in generating images. The only exception is the bypass rate for text-based safety filters because such filters are deployed before the diffusion model.

We then describe the closed-box safety filter of DALLE-2. SneakyPrompt achieves 57.15% one-time bypass rate with an average of 24.49 queries. Note that while the rate seems relatively low, this is the first work that bypasses the closed-box safety filter of DALLE-2 with real-world sensitive images as shown in Appendix A. None of the existing works, including text-based adversarial examples [12], [13], [14] and Rando et al. [16], is able to bypass this closed-box safety filter. Interestingly, the re-use bypass rate for DALLE-2 is 100%, i.e., as long as an adversarial prompt bypass the filter once, it will bypass the filter multiple times with NSFW images generated.

We now describe our observations related to the robustness of safety filters below.

Scale of Safety Filter. Table 2 shows that a safety filter’s robustness (especially against SneakyPrompt) is positively correlated with its scale, i.e., the total number of parameters. This observation holds for both one-time and re-use performance on all metrics. For example, SneakyPrompt performs the worst against text-classifier because its scale is much larger than other variants. As a comparison, SneakyPrompt achieves the least number of queries and the highest image se-

mantic similarity against text-match and text-image-threshold (which are not learning-based filters).

Type of Safety Filter. We have two observations with regard to the safety filter type when safety filters have a similar number of parameters. First, the combination of text and image is better than those relying on any single factor. For example, text-image-threshold outperforms text-match in terms of all metrics. In addition, image-clip-classifier outperforms image-classifier with a small number of online queries, because image-clip-classifier utilizes the embedding from CLIP (which includes both image and text information) as opposed to only image information in image-classifier.

Second, image-based safety filters have a lower re-use bypass rate compared to text-based ones. It is because the random seeds used for the text-to-image model to generate images are uncontrollable, which leads to different generated images using the same prompt at different times. Therefore, the one-time bypassed adversarial prompt may not bypass the image-based safety filter during re-use. As a comparison, because the text-based safety filter takes a prompt as input, the generated adversarial prompt can bypass the same safety filter in re-use as long as the filter is not updated.

Non-English Word Safety Filter. We add a simple non-English word safety filter in combination with existing filters and show the effectiveness of SneakyPrompt in bypassing this adaptive safety filter. Note that the search space of SneakyPrompt will be the google-10000-english dictionary [51] that contains a list of the 10,000 most common English words, instead of all tokens from CLIP vocabulary dictionary [52].

We have three observations. First, the FID score is on par with SneakyPrompt-RL against existing safety filters alone, which indicates SneakyPrompt-RL is also effective in maintaining image semantics when a non-English word filter is present. Second, the bypass rate is 3.75% on average lower than that without a non-English word filter. The reason is that SneakyPrompt might select synonyms for bypassing once but cannot be reused as shown in existing text adversarial

Table 6: [RQ2] One-time bypass rate of SneakyPrompt-RL compared with existing works in bypassing the online, closed-box safety filter of DALL-E 2.

	SneakyPrompt-RL	Text adversarial examples			Manual prompt		Optimized prompt
		TextFooler [13]	TextBugger [12]	BAE [14]	Rando et al. [1]	Qu et al. [17]	Maus et al. [15]
Bypass rate	57.15%	0.00%	1.00%	0.00%	0.00%	0.00%	0.00%

examples. Third, the number of online queries is 3.11 on average higher since the search space is limited to English words; thus, more queries are needed to find an adversarial prompt that can bypass and maintain the image semantics.

[RQ1] *Take-away*: SneakyPrompt successfully bypasses all safety filters including the closed-box safety filter adopted by DALL-E 2 as well as a non-English word safety filter.

6.2. RQ2: Performance Comparison with Baselines

In this research question, we first compare SneakyPrompt with existing methods using the original safety filter of Stable Diffusion and then add the non-English word filter. We then show none of the existing methods can effectively bypass the DALL-E 2’s safety filter. Specifically, we compare SneakyPrompt with the following works:

- Text-based adversarial examples. We use three related works that generate closed-box, text-based adversarial examples, which are Textbugger [12], Textfooler [13], and BAE [14]. We follow the implementation by TextAttack [53] with their default hyperparameters.
- Manually-curated adversarial prompts. We use Rando et al. [16] (which contains manually-generated prompts) and Qu et al. [17] (in which prompts are manually created based on a template).
- Optimized adversarial prompts. We use Maus et al. [15], a concurrent work to ours, to generate adversarial prompts.
- SneakyPrompt baselines with different search algorithms. We call them SneakyPrompt-base.

We start from the original safety filter of Stable Diffusion. Table 4 shows the comparison results in terms of three metrics and we describe two different scenarios below.

- Re-use adversarial prompts. This attack scenario assumes that adversarial prompts are pre-generated and re-used for the attack. On one hand, SneakyPrompt-RL achieves the highest bypass rate against the safety filter compared with all existing works and SneakyPrompt-base. The reason is that existing works, particularly text-based adversarial examples, either make minimal modifications to texts or use synonyms to replace the target token, which cannot sustain different rounds for another random seed from the diffusion model. On the other hand, SneakyPrompt-RL also has the lowest FID score compared to other methods for re-use adversarial prompts. That is, SneakyPrompt largely keeps the original semantics while existing methods—even if successfully bypassing the safety filter—will more or less modify the semantics compared to the target one. This

is because SneakyPrompt computes the similarity between the generated image and target prompt to prevent the image semantics from being modified too much and thus make it as consistent with the target prompt as possible.

- One-time searched prompts. This attack scenario assumes that an adversary always probes the target text-to-image model to generate NSFW images. SneakyPrompt-RL has the smallest number of online queries and FID scores compared with SneakyPrompt-base (the largest number of queries) and text-based adversarial examples. First, the reason is that baseline methods do not have any constraints for word choice. For example, the number of queries of brute force search is a magnitude larger than other heuristic searching methods as all tokens have the same probability of being chosen. Second, SneakyPrompt-RL takes the least number of queries, using 50% fewer queries compared with the second, which is TextFooler [13]. The reason is that SneakyPrompt-RL adopts an early stop strategy and benefits from RL. Note that manual prompts are not applicable here for one-time searched prompts because it is not scalable to probe text-to-image models for each prompt manually.

Next, we also add non-English word filter to the default stable diffusion’s safety filter and compare SneakyPrompt-RL with text-based adversarial examples. Table 5 shows the re-use and one-time bypass rates. The bypass rate of SneakyPrompt-RL does not change much because SneakyPrompt-RL can search in an English word space. Instead, the one-time performance of TextBugger [12], which utilizes alphabet swap or substitute, drops to only 36.38%, and the re-use performance drops to only 18.42%.

Last, we also evaluate the online, closed-box safety filter of DALL-E 2. Table 6 shows that none of the existing works can effectively bypass the original safety filter. The bypass rates of all the works except for TextBugger [12] are essentially zero and the bypass rate of TextBugger [12] is also as low as 1.00%.

[RQ2] *Take-away*: SneakyPrompt-RL outperforms SneakyPrompt-base, existing text-based adversarial examples, and manual prompts from Rando et al. [16] and Qu et al. [17].

6.3. RQ3: Study of Different Parameter Selection

In this research question, we study how different parameters affect the overall performance of SneakyPrompt.

Reward Function. We have two variants of the reward function, one based on cosine similarity between embedding

Table 7: [RQ3] Performance of SneakyPrompt-RL using different reward functions.

Reward function	Re-use adversarial prompts			One-time searched adversarial prompts			
	Bypass rate (\uparrow)	FID score (\downarrow)		Bypass rate (\uparrow)	FID score (\downarrow)		# of online queries
		adv. vs. target	adv. vs. real	adv. vs. target	adv. vs. real		
$\cos(\mathcal{M}(p_a), \hat{\mathcal{E}}(p_t))$	69.35%	148.64	169.15	100.00%	108.31	132.01	9.51 ± 4.31
$1 - \ell_2(\hat{\mathcal{E}}(p_a), \hat{\mathcal{E}}(p_t))$	55.25%	165.35	189.31	96.42%	149.21	168.74	2.18 ± 1.12

Table 8: [RQ3] Performance of SneakyPrompt-RL when the shadow text encoder $\hat{\mathcal{E}}$ is the same as, or different from, the target text encoder used by the text-to-image model \mathcal{M} .

Shadow text encoder	Re-use adversarial prompts			One-time searched adversarial prompts			
	Bypass rate (\uparrow)	FID score (\downarrow)		Bypass rate (\uparrow)	FID score (\downarrow)		# of online queries
		adv. vs. target	adv. vs. real	adv. vs. target	adv. vs. real		
$\hat{\mathcal{E}} \neq \mathcal{E}$	69.35%	148.64	169.15	100.00%	108.31	132.01	9.51 ± 4.31
$\hat{\mathcal{E}} = \mathcal{E}$	68.87%	143.88	162.26	100.00%	97.25	121.42	9.60 ± 3.45

of the generated image and embedding of the target prompt, and the other based on ℓ_2 distance between embeddings of the target and adversarial prompts. Table 7 shows the comparison between the two reward functions: cosine similarity results in a higher bypass rate and a better image semantic similarity for both one-time and re-use attacks, while ℓ_2 distance results in a smaller number of online queries.

Shadow Text Encoder. Table 8 shows the impact of different shadow text encoders $\hat{\mathcal{E}}$, particularly when $\hat{\mathcal{E}} = \mathcal{E}$ and $\hat{\mathcal{E}} \neq \mathcal{E}$. We have two observations. First, $\hat{\mathcal{E}} = \mathcal{E}$ improves the image semantic similarity with a smaller FID score for both one-time and re-use attacks and especially for one-time. The reason is the same text encoder adopted by the attacker results in the same text embedding as the internal result of text-to-image models, used to guide the semantics of image generation, which provides a more precise semantic compared with a different text encoder. The relatively smaller improvement for re-use FID score is because of the disturbance of the random seed. Second, there is no significant difference in bypass rates and the number of online queries. SneakyPrompt-RL achieves a 100% bypass rate with both shadow text encoders for the one-time performance and achieves around 70% for re-use performance with different random seeds involved. The number of queries is also similar because they use the same similarity threshold δ .

Similarity Threshold. Figure 5 shows the impact of different similarity threshold values, ranging from 0.22 to 0.30, on SneakyPrompt’s performance using three metrics. Let us start with the bypass rate in Figure 5a. The bypass rate stays the same for one-time attacks but drops for re-use attacks because the generated image may be closer to the target and thus blocked by the safety filter. This is also reflected in Figure 5b as the FID score decreases as the threshold increases. Similarly, Figure 5c shows that the number of queries increases as the similarity threshold because it will be harder to satisfy the threshold during searching.

Search Space. We evaluate the impact of search space size—which is partially controlled by the parameter l —on the performance of SneakyPrompt. Specifically, we change

l to be $[3, 5, 10, 15, 20]$, which lead to a search space of $[625^m, 6922^m, 29127^m, 46168^m, 48095^m]$ according to the number of candidate tokens in the dictionary D_l . Figure 6 shows the results of the impacts of l on three metrics. First, Figure 6a shows that a larger l leads to a higher re-use bypass rate. The reasons are twofold. On one hand, the larger the l , the larger the search space is. That is, RL has room to explore more tokens and increase the bypass rate. On the other hand, a larger l introduces longer tokens that dilute the target prompt more, thus also increasing the bypass rate. Second, the image semantic similarity has little correlation with the search space. Instead, image semantic similarity in terms of FID scores are more related to the semantic similarity threshold δ as we show in Figure 5. Third, the larger the l is, the more online queries SneakyPrompt takes. The reason is that RL needs more queries to explore a larger search space to satisfy the semantic similarity threshold.

[RQ3] Take-away: SneakyPrompt needs to balance the bypass rate and the FID score with the number of queries in selecting different parameters.

6.4. RQ4: Explanation of Bypassing

In this research question, we explain why SneakyPrompt successfully bypasses different safety filters while maintaining the image semantic similarities. Specifically, we use the ground-truth output probability of different safety filters for the explanation. Table 9 shows the experiment results that align with our high-level explanation in Figure 3. We first show the results of the output probability on NSFW of different safety filters for both the target and the adversarial prompts or their generated images, depending on the type of safety filters. We normalize the probability into $[0, 1]$, where a value larger than 0.5 indicates the sensitive input. The probability output of the target prompt p_t is from 0.546 to 1.000, i.e., the target prompt is sensitive. As a comparison, the probability outputs of adversarial prompts generated by SneakyPrompt range from 0.000 to 0.482, i.e., they are non-sensitive. This result suggests that SneakyPrompt is effective in bypassing different safety filters. Next, we also

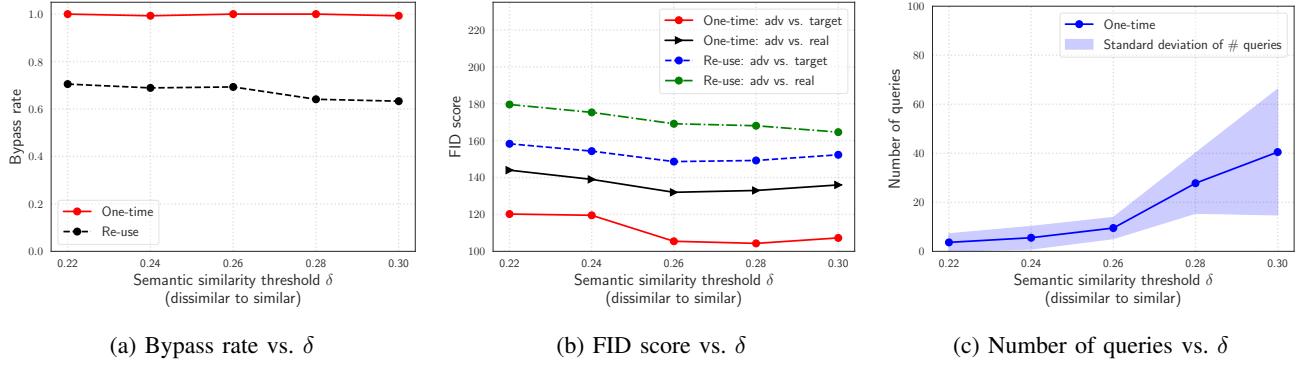


Figure 5: [RQ3] Performance vs. Semantic similarity threshold δ .

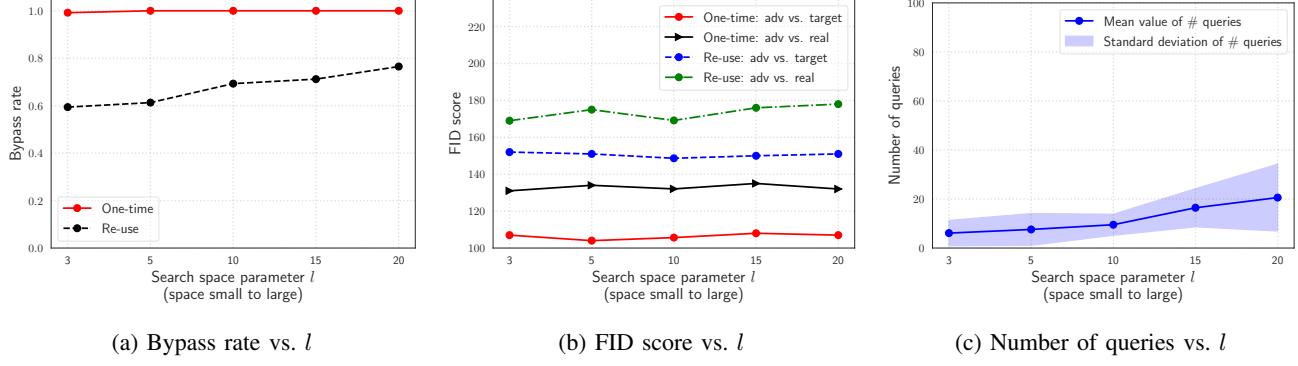


Figure 6: [RQ3] Performance vs. Search space parameter l .

Table 9: [RQ4] Explanation on why SneakyPrompt bypasses a safety filter while still maintaining the image semantics. p_t is the target prompt that contains the NSFW content, and p_a is the adversarial prompt generated by SneakyPrompt. For each safety filter \mathcal{F} , we normalize the output probability $\mathcal{F}(\mathcal{M}, p)$ into $[0,1]$, where the prompt or its generated image is classified as NSFW with a probability value larger than 0.5. We use $\cos(\mathcal{M}(p_t), \hat{\mathcal{E}}(p_t))$ as the ground truth similarity for references, where we obtain $\mathcal{M}(p_t)$ by removing the safety filter of Stable Diffusion for research purpose. The value in the table is the average for target prompts and their adversarial prompts in NSFW-200.

[RQ4] *Take-away*: The outputs from safety filters show that SneakyPrompt bypasses them while still maintaining the NSFW semantics.

7. Conclusion, Discussion, and Future Work

We show that a black-box safety filter of a text-to-image model can be jailbroken to produce an NSFW image with a small number of queries to the model. Reinforcement learning can reduce the number of queries to the text-to-image model by leveraging the query results to strategically guide the perturbation of a prompt. Our results imply that existing guardrails of text-to-image models are insufficient and highlight the urgent need for new guardrails to limit the societal harms of powerful text-to-image models. We note that, instead of using add-on safety filters, some methods [54] could be used to edit the parameters of a text-to-image model to erase sensitive concepts such that it intrinsically will not generate NSFW images. SneakyPrompt is also applicable to such a text-to-image model with an embedded safety filter. This is because SneakyPrompt only needs black-box access to an (add-on or embedded) safety filter. Developing more robust safety filters is an urgent future research direction. For instance, one way is to leverage adversarial training, which considers adversarial prompts during the training of a safety filter.

show the semantic similarity between the target prompt and its generated image, and that between the target prompt and the image generated by adversarial prompts found by SneakyPrompt. We observe that the former, i.e., 0.298, is close to the latter, i.e., 0.267 to 0.289, for different safety filters, which indicates the ability of SneakyPrompt to maintain the semantics of target prompts and images generated based on them.

Safety filter \mathcal{F}	Probability of being NSFW		Semantics similarity	
	$\mathcal{F}(\mathcal{M}, p_t)$	$\mathcal{F}(\mathcal{M}, p_a)$	$\cos(\mathcal{M}(p_t), \hat{\mathcal{E}}(p_t))$	$\cos(\mathcal{M}(p_a), \hat{\mathcal{E}}(p_t))$
text-image-threshold	0.546	0.441		0.289
text-match	1.000	0.000		0.291
text-classifier	0.976	0.482	0.298	0.267
image-classifier	0.791	0.411		0.276
image-clip-classifier	0.885	0.473		0.271

Acknowledgments

We would like to thank the anonymous shepherd and reviewers for their helpful comments and feedback. This work was supported in part by Johns Hopkins University Institute for Assured Autonomy (IAA) with grants 80052272 and 80052273, National Science Foundation (NSF) under grants CNS-21-31859, CNS-21-12562, CNS-19-37786, CNS-19-37787, and CNS-18-54000, as well as Army Research Office (ARO) under grant No. W911NF2110182. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of NSF, ARO, or JHU-IAA.

References

- [1] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer, “High-resolution image synthesis with latent diffusion models,” in *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [2] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen, “Hierarchical text-conditional image generation with clip latents,” *arXiv preprint arXiv:2204.06125*, 2022.
- [3] C. Saharia, W. Chan, S. Saxena, L. Li, J. Whang, E. Denton, S. K. S. Ghasemipour, R. Gontijo-Lopes, B. K. Ayan, T. Salimans, J. Ho, D. J. Fleet, and M. Norouzi, “Photorealistic text-to-image diffusion models with deep language understanding,” in *Proceedings of Neural Information Processing Systems (NeurIPS)*, 2022.
- [4] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2015.
- [5] J. Ho, A. Jain, and P. Abbeel, “Denoising diffusion probabilistic models,” in *Proceedings of the Neural Information Processing Systems (NeurIPS)*, 2019.
- [6] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, G. Krueger, and I. Sutskever, “Learning transferable visual models from natural language supervision,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2021.
- [7] L. Floridi and M. Chiriaci, “Gpt-3: Its nature, scope, limits, and consequences,” *Minds and Machines*, 2020.
- [8] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, “Exploring the limits of transfer learning with a unified text-to-text transformer,” *Journal of Machine Learning Research*, 2020.
- [9] “Microsoft Designer,” <https://designer.microsoft.com/>.
- [10] A. Kim, “Nsfw image dataset,” https://github.com/alex000kim/nsfw_data_scraper, 2022.
- [11] R. George, “Nsfw words list on github,” https://github.com/rgeorge-pdcontributions/NSFW-Words-List/blob/master/nsfw_list.txt, 2020.
- [12] J. Li, S. Ji, T. Du, B. Li, and T. Wang, “Textbugger: Generating adversarial text against real-world applications,” *arXiv preprint arXiv:1812.05271*, 2018.
- [13] D. Jin, Z. Jin, J. T. Zhou, and P. Szolovits, “Is bert really robust? a strong baseline for natural language attack on text classification and entailment,” in *Proceedings of the AAAI conference on artificial intelligence (AAAI)*, 2020.
- [14] S. Garg and G. Ramakrishnan, “BAE: BERT-based adversarial examples for text classification,” in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2020.
- [15] N. Maus, P. Chao, E. Wong, and J. Gardner, “Adversarial prompting for black box foundation models,” *arXiv*, 2023.
- [16] J. Rando, D. Paleka, D. Lindner, L. Heim, and F. Tramèr, “Red-teaming the stable diffusion safety filter,” *arXiv preprint arXiv:2210.04610*, 2022.
- [17] Y. Qu, X. Shen, X. He, M. Backes, S. Zannettou, and Y. Zhang, “Unsafe diffusion: On the generation of unsafe images and hateful memes from text-to-image models,” in *Proceedings of the ACM Conference on Computer and Communications Security (CCS)*, 2023.
- [18] E. Mansimov, E. Parisotto, J. L. Ba, and R. Salakhutdinov, “Generating images from captions with attention,” *arXiv*, 2016.
- [19] T. Xu, P. Zhang, Q. Huang, H. Zhang, Z. Gan, X. Huang, and X. He, “AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [20] J. Y. Koh, J. Baldridge, H. Lee, and Y. Yang, “Text-to-image generation grounded by fine-grained user attention,” in *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, 2021.
- [21] A. Nguyen, J. Clune, Y. Bengio, A. Dosovitskiy, and J. Yosinski, “Plug & play generative networks: Conditional iterative generation of images in latent space,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.
- [22] Midjourney, 2022. [Online]. Available: <https://www.midjourney.com>
- [23] Y. Zhou, R. Zhang, C. Chen, C. Li, C. Tensmeyer, T. Yu, J. Gu, J. Xu, and T. Sun, “Towards language-free training for text-to-image generation,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [24] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever, “Zero-shot text-to-image generation,” in *Proceedings of the International Conference on Machine Learning (ICML)*, 2021.
- [25] Y. Wu, N. Yu, Z. Li, M. Backes, and Y. Zhang, “Membership inference attacks against text-to-image generation models,” *arXiv preprint arXiv:2210.00968*, 2022.
- [26] J. Duan, F. Kong, S. Wang, X. Shi, and K. Xu, “Are diffusion models vulnerable to membership inference attacks?” *arXiv preprint arXiv:2302.01316*, 2023.
- [27] R. Shokri, M. Stronati, C. Song, and V. Shmatikov, “Membership inference attacks against machine learning models,” in *Proceedings of the IEEE Symposium on Security and Privacy (SP)*, 2017.
- [28] B. Hui, Y. Yang, H. Yuan, P. Burlina, N. Z. Gong, and Y. Cao, “Practical blind membership inference attack via differential comparisons,” in *Proceedings of the Network and Distributed System Security Symposium (NDSS)*, 2021.
- [29] N. Carlini, J. Hayes, M. Nasr, M. Jagielski, V. Sehwag, F. Tramèr, B. Balle, D. Ippolito, and E. Wallace, “Extracting training data from diffusion models,” *arXiv preprint arXiv:2301.13188*, 2023.
- [30] R. Millière, “Adversarial attacks on image generation with made-up words,” *arXiv preprint arXiv:2208.04135*, 2022.
- [31] C. Yang, A. Kortylewski, C. Xie, Y. Cao, and A. Yuille, “Patchattack: A black-box texture-based attack with reinforcement learning,” *arXiv preprint arXiv:2004.05682*, 2020.
- [32] I. J. Goodfellow, J. Shlens, and C. Szegedy, “Explaining and harnessing adversarial examples,” *arXiv preprint arXiv:1412.6572*, 2014.
- [33] M. Shu, C. Liu, W. Qiu, and A. Yuille, “Identifying model weakness with adversarial examiner,” in *Proceedings of the AAAI conference on artificial intelligence (AAAI)*, 2020.

[34] A. Liu, H. Yu, X. Hu, S. Li, L. Lin, F. Ma, Y. Yang, and L. Wen, “Character-level white-box adversarial attacks against transformers via attachable subwords substitution,” in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2022.

[35] M. Alzantot, Y. Sharma, A. Elgohary, B.-J. Ho, M. Srivastava, and K.-W. Chang, “Generating natural language adversarial examples,” in *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 2018.

[36] M. Han, L. Zhang, J. Wang, and W. Pan, “Actor-critic reinforcement learning for control with stability guarantee,” *Proceedings of the IEEE Robotics and Automation Letters (RA-L)*, 2020.

[37] Pricing model of openai. <https://openai.com/pricing>.

[38] “ViT-L/14,” <https://huggingface.co/openai/clip-vit-large-patch14>.

[39] Torchmetrics, “CLIP Score,” https://torchmetrics.readthedocs.io/en/stable/multimodal/clip_score.html, 2022.

[40] C. Yang, A. Kortylewski, C. Xie, Y. Cao, and A. Yuille, “Patchattack: A black-box texture-based attack with reinforcement learning,” *arXiv preprint arXiv:2004.05682*, 2020.

[41] M. Li, “Nsfw text classifier on hugging face,” https://huggingface.co/michellejeli/NSFW_text_classifier, 2022.

[42] Hugging face. <https://huggingface.co/CompVis/stable-diffusion-v1-4>.

[43] Openai online apis. <https://openai.com/blog/dall-e-api-now-available-in-public-beta>.

[44] “Nsfw gpt,” https://www.reddit.com/r/ChatGPT/comments/11vlp7j/nsfwgpt_that_nsfw_prompt/, 2023.

[45] V. Sanh, L. Debut, J. Chaumond, and T. Wolf, “Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter,” *arXiv preprint arXiv:1910.01108*, 2019.

[46] L. Chhabra, “Nsfw image classifier on github,” <https://github.com/lakshaychhabra/NSFW-Detection-DL>, 2020.

[47] C. ALESSIO, “Animals-10 dataset,” <https://www.kaggle.com/datasets/alessiocorrad99/animals10>, 2020.

[48] LAION-AI, “Nsfw clip based image classifier on github,” <https://github.com/LAION-AI/CLIP-based-NSFW-Detector>, 2023.

[49] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter, “Gans trained by a two time-scale update rule converge to a local nash equilibrium,” *Proceedings of the Neural Information Processing Systems (NeurIPS)*, 2017.

[50] M. Seitzer, “pytorch-fid: FID Score for PyTorch,” <https://github.com/mseitzer/pytorch-fid>, 2020.

[51] “Google 10000 English Vocabularies,” <https://github.com/first20hours/google-10000-english>.

[52] “CLIP Vocabulary Dictionary,” <https://huggingface.co/openai/clip-vit-base-patch32/resolve/main/vocab.json>.

[53] J. Morris, E. Lifland, J. Y. Yoo, J. Grigsby, D. Jin, and Y. Qi, “Textattack: A framework for adversarial attacks, data augmentation, and adversarial training in nlp,” in *Proceedings of the Conference on Empirical Methods in Natural Language Processing: System Demonstrations (EMNLP)*, 2020.

[54] N. Kumari, B. Zhang, S.-Y. Wang, E. Shechtman, R. Zhang, and J.-Y. Zhu, “Ablating concepts in text-to-image diffusion models,” in *Proceedings of the International Conference on Computer Vision (ICCV)*, 2023.

Appendix A.

Examples of Generated Sensitive Images

We show examples of generated NSFW images in Figure 7 with an external link. Some adversarial prompts can also be found at [this link](#) with password access.

Algorithm 3 SneakyPrompt-RL with Alternative Reward

Input: Target prompt p_t , target text-to-image model \mathcal{M} , shadow text encoder $\hat{\mathcal{E}}$, threshold δ , maximum number of queries Q , policy network P , learning rate η , and D_L .

Output: Adversarial prompt p_a and generated image $\mathcal{M}(p_a)$ if any.

```

1: //Get initial sensitive tokens in  $p_t$  and search space  $\mathcal{S}$ 
2:  $\mathcal{S}, \omega \leftarrow \text{GetSearchSpace}(Initial = 1)$ 
3: //Get text embedding of  $p_t$ 
4:  $\hat{\mathcal{E}}(p_t) \leftarrow \text{OfflineQuery}(p_t, \hat{\mathcal{E}})$ 
5: Initialize  $P$  randomly
6:  $r_{max} \leftarrow 0$ 
7:  $q \leftarrow 1$ 
8: while  $q \leq Q$  do
9:    $r_q \leftarrow -1$ 
10:  //Construct an adversarial prompt
11:  while  $r_q < \delta$  do
12:     $C \leftarrow P$  //Sample replacing tokens from  $\mathcal{S}$  using  $P$ 
13:     $p_a \leftarrow \text{Construct adversarial prompt based on } C \text{ and } p_t$ 
14:     $r_q \leftarrow \text{GetSimilarity}(\hat{\mathcal{E}}(p_a), \hat{\mathcal{E}}(p_t))$ 
15:     $\text{Update}(r_q, C, \eta)$ 
16:  end while
17:  //Query the target model  $\mathcal{M}$ 
18:   $\mathcal{F}(\mathcal{M}, p_a), \mathcal{M}(p_a) \leftarrow \text{OnlineQuery}(p_a, \mathcal{M})$ 
19:  if  $\mathcal{F}(\mathcal{M}, p_a) == 0$  then
20:    return  $p_a$  and  $\mathcal{M}(p_a)$ 
21:  else
22:     $r_q \leftarrow -q/(10 \cdot Q)$ 
23:     $\text{Update}(r_q, C, \eta)$ 
24:  end if
25:  //Save the  $p_a$  and the generated image with the largest reward
26:  if  $r_q > r_{max}$  then
27:     $r_{max} \leftarrow r_q$ 
28:     $p'_a \leftarrow p_a$ 
29:     $\mathcal{M}(p'_a) \leftarrow \mathcal{M}(p_a)$ 
30:  end if
31:  //Not bypass safety filter in 5 consecutive queries
32:  if  $r_{q-4}, r_{q-3}, r_{q-2}, r_{q-1}, r_q < 0$  then
33:    //Expand the search space by replacing one more token in  $p_t$ 
34:     $\mathcal{S}, \omega \leftarrow \text{GetSearchSpace}(Initial = 0)$ 
35:  end if
36:  //Rewards do not change in 3 consecutive queries
37:  //or fraction  $\omega$  of tokens in  $p_t$  to be replaced is no smaller than 0.3
38:  if  $|r_{q-2} + r_q - 2r_{q-1}| < 1e-4$  or  $\omega \geq 0.3$  then
39:    return  $p'_a$  and  $\mathcal{M}(p'_a)$ 
40:  end if
41:   $q \leftarrow q + 1$ 
42: end while
43: return  $p'_a$  and  $\mathcal{M}(p'_a)$ 

```

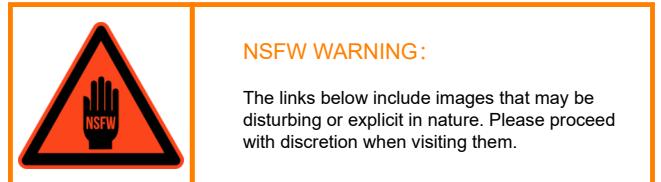


Figure 7: Examples of generated images containing disturbing, violent, nudity, or sexual content. Please contact the authors to obtain the password and visit [this link](#) at your own discretion.

Appendix B. Meta-Review

The following meta-review was prepared by the program committee for the 2024 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 proposes an attack framework, SneakyPrompt, which aims to circumvent NSFW content filters used by generative text-to-image models. The authors perform an evaluation on a closed-box safety filter DALL-E 2 and demonstrate a bypass rate of 57.15%.

B.2. Scientific Contributions

- Provides a Valuable Step Forward in an Established Field.
- Creates a New Tool to Enable Future Science.

B.3. Reasons for Acceptance

- 1) This paper provides a valuable step forward in an established field. While content filter evasion is a recognized area of study, its application to large text-to-image models is relatively recent. The paper introduces, for the first time, an RL-based attack algorithm, capable of bypassing the closed-box safety filter of DALL-E 2.
- 2) This paper creates a new tool to enable future science. The authors commit to open-sourcing a new automated attack framework, which can circumvent both open-source and closed-box NSFW filters used by generative image models.

B.4. Noteworthy Concerns

- 1) The paper lacks evidence showcasing the broader applicability of the proposed attack on closed-box safety filters, as the only closed-box filter evaluated is DALL-E 2.

Appendix C. Response to the Meta-Review

We would like to thank the anonymous shepherd and the reviewers for their valuable insights and the time to provide a meta-review. The meta-review notes that reviewers would have liked us to apply the proposed attack on additional closed-box safety filters other than DALL-E 2. We agree that the evaluation would be important and strengthen the paper. However, many of them do not provide a well-documented programming interface or charge too much, which prevents us from such an evaluation. We will consider evaluating

closed-box safety filters as future work if access to well-documented programming interfaces improves, or if we can secure funding to cover the costs associated with their use. Furthermore, we will explore partnerships with organizations that have access to these systems, which could facilitate a more comprehensive evaluation.