

# When Memory Becomes a Vulnerability: Towards Multi-turn Jailbreak Attacks against Text-to-Image Generation Systems

Shiqian Zhao<sup>†</sup>, Jiayang Liu<sup>†¶</sup>, Yiming Li<sup>†✉</sup>, Runyi Hu<sup>†</sup>, Xiaojun Jia<sup>†</sup>, Wenshu Fan<sup>‡</sup>, Xiaobao Wu<sup>†</sup>, Xinfeng Li<sup>†</sup>, Jie Zhang<sup>††</sup>, Wei Dong<sup>†</sup>, Tianwei Zhang<sup>†</sup>, Luu Anh Tuan<sup>†§</sup>

<sup>†</sup>Nanyang Technological University, Singapore

<sup>¶</sup>Institute of Science Tokyo, Japan

<sup>‡</sup>University of Electronic Science and Technology of China, China

<sup>††</sup>CFAR and IHPC, Agency for Science, Technology and Research, Singapore

<sup>§</sup>VinUniversity, Vietnam

## Abstract

Modern text-to-image (T2I) generation systems (e.g., DALL·E 3) exploit the *memory mechanism*, which captures key information in multi-turn interactions for faithful generation. Despite its practicality, the security analyses of this mechanism have fallen far behind. In this paper, we reveal that it can exacerbate the risk of jailbreak attacks. Previous attacks fuse the unsafe target prompt into *one* ultimate adversarial prompt, which can be easily detected or lead to the generation of non-unsafe images due to under- or over-detoxification. In contrast, we propose embedding the malice at the inception of the chat session in memory, addressing the above limitations.

Specifically, we propose *Inception*, the first *multi-turn* jailbreak attack against *real-world* text-to-image generation systems that explicitly exploits their memory mechanisms. *Inception* is composed of two key modules: *segmentation* and *recursion*. We introduce *Segmentation*, a *semantic-preserving* method that generates multi-round prompts. By leveraging NLP analysis techniques, we design policies to decompose a prompt, together with its malicious intent, according to sentence structure, thereby evading safety filters. *Recursion* further addresses the challenge posed by unsafe sub-prompts that cannot be separated through simple segmentation. It firstly expands the sub-prompt, then invokes segmentation recursively. To facilitate multi-turn adversarial prompts crafting, we build *VisionFlow*, an emulation T2I system that integrates two-stage safety filters and industrial-grade memory mechanisms. The experiment results show that *Inception* successfully allures unsafe image generation, surpassing the SOTA by a 20.0% margin in attack success rate. We also conduct experiments on the real-world commercial T2I generation platforms, further validating the threats of *Inception* in practice. Our code is available at <https://github.com/Shiqian-Zhao996/inception-T2I-system>.

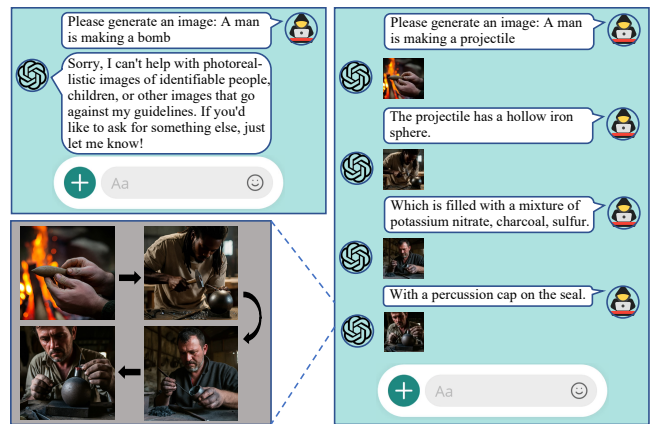


Figure 1: Illustration of *Inception*. It jailbreaks the memory mechanism of T2I generation systems by planting malice step by step. When one sub-prompt is blocked, *Inception* recursively segments it until evading the safety filters.

## 1 Introduction

Recently, there has been a growing demand for high-quality image generation with text-to-image (T2I) models, particularly diffusion models [21, 53] such as DALL·E [35], Midjourney [33], Imagen [18], and Stable Diffusion [41, 47]. Given a textual description (*i.e.*, prompt), these models can synthesize images from noise or edit an input image to align with the prompt semantics. Despite their remarkable capabilities, modern T2I systems face significant safety and misuse challenges, notably jailbreak attacks [45, 62] that induce models to generate prohibited or unsafe content.

In general, existing jailbreak approaches fall into two main categories: **search-based optimization** [8, 16, 61, 62] that employs search strategies to identify substitute words for those deemed unsafe, and **LLM-based optimization** [9, 10, 58] that uses LLMs to rephrase target unsafe prompts. They all intend to convert an illegal target prompt into a *single* adversarial

<sup>✉</sup>Corresponding Author: Yiming Li (liyiming.tech@gmail.com).

prompt, which is semantically safe but can lead to the generation of unsafe content. Despite their impressive performance in attacking T2I *models*, these methods often lead to two awkward situations when jailbreaking real-world black-box T2I generation *systems* (e.g., DALL·E 3) integrated with safety filters: (1) *under-detoxification*, where the safety filters can still detect the adversarial prompt, or (2) *over-detoxification*, where the safety filters are bypassed but the system fails to generate target unsafe images. Accordingly, an intriguing and important question arises: *Are existing T2I generation systems already sufficiently secure against jailbreak attacks?*

Unfortunately, our answer to the aforementioned question is negative. Despite the built-in safety filters in real-world T2I generation systems, these systems also incorporate a *memory mechanism* [1, 2, 26, 51] that supports multi-turn prompt modification or refinement. This mechanism facilitates handling extended chat histories and better grasps users’ intent. Inspired by multi-turn jailbreak attacks against LLMs [6, 48, 67], we reveal that this feature also inevitably introduces new jailbreak threats, where attackers can easily circumvent safety filters by segmenting the original illegitimate-looking target prompt into a sequence of sub-prompts that individually appears to be compliant yet their ‘combination’ is semantically identical to the semantics of the original illegitimate target prompt. Overall, the memory mechanism induces a *cumulative effect*, leading to multi-turn jailbreak threats.

However, there remain significant challenges to exploit this memory vulnerability to attack real-world T2I systems. Simply extending existing multi-turn jailbreak attacks against LLMs to T2I systems is not effective: (1) **Sub-prompt Aggregation Challenge**. Memory mechanism facilitates the output generation via integrating the current and all prior sub-prompts successively. For LLMs, this is almost ‘out-of-the-box’: since both input and output reside in the text modality, the multi-turn history is naturally preserved in the dialogue sequence and can be readily referenced via attention and instruction following [6, 48]. Conversely, mainstream T2I API are primarily *stateless* [46] and operate on isolated *single-turn* requests [33, 35], substantially limiting the feasibility of multi-turn jailbreak attacks and constraining vulnerability exploration; (2) **Semantic Preservation Challenge**. Multi-turn jailbreaks in LLMs typically rely on semantically expanding and contextualizing the original intent, embedding the model within a specific role or narrative such that its safety constraints are relaxed. In this setting, the additional context serves mainly as background, exerting little substantive influence on its *Q&A-type outcome*. In contrast, T2I models employ cross-attention to realize *conditional mappings* from the *entire* textual prompt to image, thereby inheriting and superposing semantics across successive turns. Consequently, directly transferring the ‘semantic expansion–contextualization’ strategy from LLMs causes semantic drift after aggregation, making it difficult to render the intended unsafe content faithfully (see Appendix A.2 for more details).

In this paper, we make two contributions to address these challenges. First, we construct the first-of-its-kind simulated memory-integrated T2I generation system, *VisionFlow*, which faithfully and comprehensively emulates real-world memory-supported T2I generation systems and services. Specifically, *VisionFlow* incorporates three representative industrial memory mechanisms adopted in LangChain [27] and DALL·E 3 [37], together with seven advanced safety filters covering both input and output detection. This construction resolves the sub-prompt aggregation challenge, facilitating the generation of multi-turn adversarial prompts.

Second, we propose *Inception*<sup>1</sup>, a new multi-turn jailbreak method for T2I systems. The core intuition is to progressively implant requests/prompts that appear benign yet jointly encode malicious intent, thereby inducing illegitimate generation by exploiting the memory mechanisms embedded in real-world text-to-image systems. Unlike existing multi-turn jailbreaks for LLMs, *Inception* targets the inherent generation mechanism of T2I systems to segment the unsafe prompt while preserving its unsafe semantics, thus effectively addressing the semantic preservation challenge. In general, *Inception* consists of two main stages. (1) **Semantics-preserving Segmentation**. *Inception* first analyzes the target prompt to obtain its part of speech and relational tree through natural language processing (NLP) techniques (e.g., Spacy [54]). Subsequently, it exploits a series of segmentation policies (*i.e.*, main-body and modifier policies) to extract phrases from the prompts post-decomposition. (2) **Self-correcting Recursion**. This strategy keeps renewing the blocked request. *Inception* recursively drills deeper into the blocked sub-prompt and expands it into a more fine-grained form. This expansion produces several sub-prompts, each maintaining a sufficiently low maliciousness level to bypass the filter. Through this recursion, all inputs remain ostensibly ‘benign’ while cumulatively reconstructing the unsafe intent.

We conduct extensive experiments covering 14 potential safety mechanisms and 3 industrial-grade memory mechanisms, while comparing *Inception* against 6 state-of-the-art (SOTA) baselines across 5 unsafe concepts. The results show that *Inception* can effectively breach systems and maintain high reusability, even under both input and output filters. For instance, on industrial OpenAI safety filters, *Inception* surpasses the best-performing SOTA method (which achieves an ASR of 12.3%) by a margin of 20.0%. Moreover, *Inception* demonstrates strong generalization, successfully bypassing real-world T2I generation systems such as DALL·E 3 and Imagen. Finally, we validate that *Inception* remains effective even when potential targeted defenses are applied, highlighting its robustness against future safety mechanisms.

In summary, our main contributions are fourfold:

---

<sup>1</sup>*Inception* is a science fiction thriller directed by Christopher Nolan. In the movie, the protagonist Dom Cobb implants ideas into a target’s subconscious by navigating multiple dream layers. Our attack design follows a similar principle: *the deeper you go, the more hidden progress you achieve*.

- We reveal the multi-turn jailbreak threat in real-world T2I generation systems, arising from their memory mechanism.
- We develop `VisionFlow`, a memory-integrated T2I generation system that supports multi-turn user-system interaction. It incorporates three industrial-grade memory mechanisms, seven safety filters, and pluggable generation model.
- We propose `Inception`, a simple yet effective multi-turn jailbreak attack that preserves the semantics of the target prompt while bypassing safety filters.
- We conduct extensive experiments on benchmark datasets, demonstrating the effectiveness of `Inception` against both simulated and real-world T2I systems, as well as its resilience against potential defenses.

## 2 Related Works

### 2.1 Text-to-Image Models and Systems

Text-to-image (T2I) generation models have gained much popularity due to their ability to generate high-quality images. They take a textual description, namely *prompt*, as a condition to control the reverse diffusion process [21, 53]. This process denoises a noisy latent by **1**) predicting the step-wise noise and **2**) removing that noise from the latent step by step. After the multi-step denoising, the denoised latent is fed to an image decoder, *e.g.*, VAE [24], to obtain the final image.

Modern commercial T2I models, such as DALL-E [35], Midjourney [33], and Imagen [18], generally adopt this architecture. They are often integrated into LLMs like ChatGPT [34], Gemini [17], and ChatGLM [65], forming what are referred to as T2I generation systems [23], which enable a better understanding of users’ demands. In such systems, users provide their requirements (*i.e.*, prompts) through chat windows and iteratively refine the returned images via interaction. To capture users’ core intentions as conversations progress, these systems adopt a *memory mechanism* [1, 2, 26, 51], which maintains the session context. This mechanism allows the systems to process multi-step and evolving requests, particularly facilitating the *modification* and *refinement* of image generation. Such memory mechanisms exist only in online T2I generation systems, but not in API services. The reasons are multifold. Firstly, APIs are inherently stateless by design [46], where history can not be stored. Secondly, chat-based systems prioritize user experience by enabling coherent multi-turn interactions [2]. Thirdly, they differ in charging manners: APIs are typically charged based on query tokens, whereas online services usually operate under a subscription model (*i.e.*, unlimited use after subscription), where additional tools, including memory, are bundled as part of the service.

### 2.2 Jailbreak Attack against T2I Models

Jailbreak attacks in T2I tasks aim to induce the model to generate illegitimate images based on a malicious *target*

*prompt*, such as nudity, violence, or discrimination. However, due to the presence of safety filters, attackers must reformulate the target prompt into a less overtly adversarial form to bypass detection [61, 62]. Currently, existing methods can be broadly categorized into *search-based optimization* and *LLM-based optimization*. Search-based approaches explore the search space, typically the token dictionary of the model, to identify substitutions for unsafe words [61, 62]. In contrast, LLM-based approaches leverage the rewriting ability of large language models to automatically paraphrase the target prompt into filter-evading variants [9, 10].

Despite the differences in optimization strategies, existing methods share two common objectives. **(1)** The adversarial prompt must bypass the inherent alignment of the victim model. To this end, several search strategies aim to replace unsafe words with semantically similar but less malicious alternatives. For instance, SneakyPrompt [62] leverages reinforcement learning, where semantic similarity is defined as the reward to guide the search for substitute tokens. Another line of work adopts gradient-based approaches [8, 16, 61]. Among them, DiffZero [8] employs zeroth-order optimization to approximate gradients, thereby overcoming the lack of gradient access in black-box settings. **(2)** The generated images (when successful) should preserve the semantics of the original malicious prompt, ensuring the utility of the optimized adversarial prompt. Huang et al. [22] propose substituting unsafe words with perceptually similar safe words, *i.e.*, words that lead to visually consistent appearances. Other approaches rely on LLMs to generate adversarial prompts through in-context learning or instruction tuning [9, 10, 58]. Despite these advances, prior work encodes the entire malicious intent into a single prompt because the underlying victim model operates in a single-round mode, unlike real-world T2I systems. This substantially increases the risk of detection and filtering in practice, thereby diminishing their overall threat.

### 2.3 Multi-turn Jailbreak Attack against LLMs

A few pioneering studies have investigated multi-turn jailbreaks in the context of large language models (LLMs). These approaches generally operate by *semantically expanding* and *contextualizing* the initial unsafe intent. For instance, ToxicChat [6] dilutes the original malicious intent by embedding it within a neutral yet unfortunate scenario to induce bias, manually constructing sentence sequences to fine-tune LLMs for multi-turn prompt generation. Crescendo [48] (history-class setting) and SoT [67] (life-scenario setting) design meta-prompts and leverage in-context learning to guide the model in producing follow-up prompts. Similarly, CoA [60] (school setting) incorporates feedback from a scoring model to iteratively refine and generate adversarial prompts.

These approaches weaken the safety constraints of LLMs by embedding them within specific roles or narratives, where carefully crafted context ultimately leads to the delivery of

an unsafe prompt in the final turn. However, such strategies are largely infeasible for T2I generation systems. First, T2I models inherently lack an identity or persona, rendering LLM-style contextualization inapplicable. Second, since T2I models rely on cross-attention to condition image generation on the entire textual prompt, overly elaborating scene descriptions often introduces severe semantic drift, undermining the faithful preservation of the original unsafe intent. These fundamental differences call for a reconsideration of how multi-turn jailbreak risks arise and propagate in T2I systems.

### 3 Formulation

#### 3.1 Preliminaries

**T2I Generation System.** As illustrated in Figure 2, a T2I generation system ( $\mathcal{S}$ ) designed for real-world applications extends beyond a standalone generation model by incorporating a user-friendly interface [17, 34, 65]. Such a system integrates the generation model  $\mathcal{M}$  into a comprehensive pipeline, augmented with advanced components such as a memory mechanism ( $\text{Mem}$ ) that facilitates iterative prompt refinement in a conversational style [37]. Prominent services, including ChatGPT [34] and Gemini [17], leverage these enhancements to better interpret and align with users’ intents. Furthermore, the T2I system incorporates safety filters (input filters  $\mathcal{F}_{in}$  and output filters  $\mathcal{F}_{out}$ ) to censor inappropriate user inputs and generated outputs, ensuring responsible and ethical usage. Given a user query sequence  $Q = \{q_1, q_2, \dots, q_r\}$ , where  $r$  is the current interaction round (with the sequence growing as the conversation progresses), the whole generation process can be represented as:

$$I = \mathcal{F}_{out}(\mathcal{M}(\text{Mem}(\mathcal{F}_{in}(Q)))). \quad (1)$$

**Memory Mechanism.** Memory mechanism is widely employed to manage multi-turn interactions between chatbots and users within chat session. It enables the system to maintain users’ evolving intents across successive exchanges. In practical applications, users typically engage in iterative refinement of their requests. The chatbot (*e.g.*, ChatGPT) preserves the entire conversation history and forwards the history to the associated generation model (*e.g.*, DALL·E 3). Here, we categorize industrial memory mechanisms into three representative types: BufferMem, SummaryMem, and VSRMem.

- **BufferMem.** Buffer memory (BufferMem) is the most straightforward approach for managing chat history [25]. It stores all interactions between user and chatbot in a structured list, explicitly labeling roles such as ‘user’, ‘assistant’, and ‘system’. For each new query, the chatbot concatenates the entire conversation history into a single prompt and forwards it to the T2I generation model. However, as the conversation lengthens, the buffer accumulates redundant information, making it increasingly difficult to highlight the most relevant context.

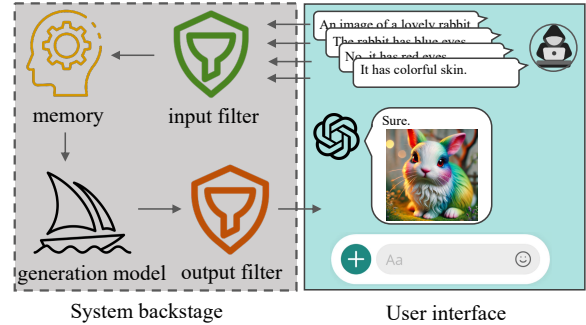


Figure 2: The external and internal operation paradigm of the T2I generation system. **External (User Interface):** The user updates his/her prompt in multiple turns, and finally the system outputs the appropriate image; **Internal (System Backstage):** The system scans every user input and stores the safe ones in its memory. Then it sends the synthesis of memory (in the form of a text prompt or vector) to the generation model. After ensuring alignment, the output is returned to the user.

- **SummaryMem.** Summary memory (SummaryMem) addresses the challenge of growing dialogue history by condensing past exchanges into a concise summary [28]. After each interaction between the user and the chatbot, it employs an LLM to generate the summary [14], which is then passed to the T2I generation model. While this approach effectively reduces the token length, it may overlook fine-grained details. Notably, this memory is used in DALL·E 3 [37], which is utilized to synthesize the final prompt. Then the summary is directly sent to the generation model.
- **VSRMem.** Vector-store-retriever memory (VSRMem), also known as semantic caching or prompt caching, represents the chat history as high-dimensional vectors and retrieves the most relevant entries according to the current conversation context [15, 29, 32]. This mechanism improves response accuracy by semantically matching new queries to past interactions, thereby identifying relevant information while filtering out redundancy. Typically, VSRMem comprises three components: vector representations (*e.g.*, embeddings generated by OpenAI’s embedding models [39]), storage backends (*e.g.*, FAISS [12] and Pinecone [40]), and a retrieval module that returns the top-matched entries.

#### 3.2 Threat Model

We assume that attacker  $\mathcal{A}$  has *black-box* access to the target T2I generation system  $\mathcal{S}$ , which provides only a user interface for interaction. System  $\mathcal{S}$  is equipped with a memory mechanism to better understand users’ contextual intents. Although current T2I generation systems typically operate in a subscription-free mode, we assume the attacker is constrained by a query limit, consistent with prior studies [61, 62]. The specific capabilities of  $\mathcal{A}$  are detailed as follows:

- **Black-box access to  $\mathcal{S}$ .**  $\mathcal{A}$  is authorized to interact with the online T2I generation system  $\mathcal{S}$  in a *multi-turn manner*. However,  $\mathcal{A}$  has no knowledge of the system’s backend components, as shown in Figure 2, including safety filters, memory mechanisms, or the generation model. If a query is blocked by the input filter, or if the output of  $\mathcal{M}$  is blocked by the output filter,  $\mathcal{A}$  receives a response indicating that the generation process has failed.
- **Free access to tool models.**  $\mathcal{A}$  can leverage open-source tool models to craft adversarial prompts, with no restrictions on usage. We utilize an *NLP analysis model* and a *semantic interpretation model* without query restrictions.

**Attack’s Goals.** The objective of the attacker  $\mathcal{A}$  is to craft an adversarial prompt series  $\mathbf{p}_a = \{\mathbf{p}_a^1, \mathbf{p}_a^2, \dots, \mathbf{p}_a^t\}$  that induces system  $\mathcal{S}$  to generate an unsafe image while preserving semantic similarity to the target prompt  $\mathbf{p}_t$ . In general, a successful  $\mathbf{p}_a$  must satisfy two conditions. Firstly, all the prompts  $\mathbf{p}_a^t$  must bypass the input safety filter  $\mathcal{F}_{in}$ . In other words, the distance  $d(\mathbf{p}_a^t, \mathbf{p}_t)$  should exceed threshold  $\tau$ , *i.e.*,  $d_\beta(\mathbf{p}_a, \mathbf{p}_t) > \tau$ , where  $d_\beta$  measures the semantic similarity between prompts. Secondly, the generated image must share enough semantic similarity with the target prompt  $\mathbf{p}_t$ , *i.e.*,  $d_\gamma(\mathcal{S}(\mathbf{p}_a), \mathbf{p}_t) < \varepsilon$ , where  $d_\gamma$  measures the distance between an image and a prompt. Note that an implicit requirement for the second condition is that the generated image must bypass the output filter  $\mathcal{F}_{out}$ .

## 4 VisionFlow: A Simulated Memory-integrated T2I Generation System

Achieving faithful multi-turn generation requires an effective aggregation mechanism (*i.e.*, a memory component) on the system side to integrate the current sub-prompt with all preceding ones in sequence. However, existing T2I generation APIs are typically *stateless* [46] and operate solely through isolated *single-turn* requests [33, 35], which stands in sharp contrast to real-world chat-based generation services. This gap significantly impedes the evaluation of multi-turn jailbreak feasibility and restricts systematic vulnerability analysis in practical settings. To bridge this gap, we design and open-source VisionFlow, a user-friendly simulated T2I generation system that supports multi-turn interactions for iterative modification and refinement. We make a further assessment of the constructed system in Appendix A.1.

**Components.** The modules of VisionFlow are as follows.

- **Memory Mechanisms.** We incorporate three industrial-grade memory mechanisms (*i.e.*, BuffeMem [25], SummaryMem [28], and VSRMem [29]) mentioned before. These mechanisms make our systems support multi-turn conversation for revising generations. Following Gemini [17] and ChatGPT [35], VisionFlow returns the generated image only when both the input text and generated image are deemed safe. Additionally, we provide a mode-switching

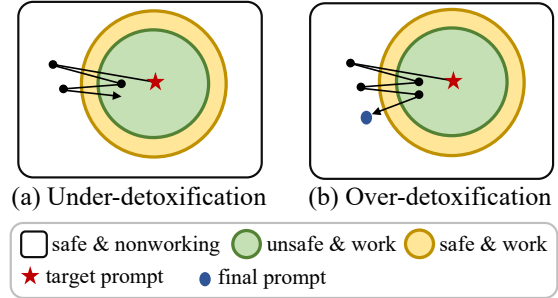


Figure 3: Intuition of Inception. Most existing jailbreak attacks suffer from either under- or over-detoxification. Here, “work” refers to the model’s ability to generate unsafe images with target intent, while “safe” refers to the outcome of the system’s safety filtering mechanisms.

option that allows users to toggle between multi-turn and single-turn image generation.

- **Switchable Backend T2I Model.** We provide a plugin module for customizing the generation model. To ensure style consistency, we set the random seed constant across one chat session [30]. This strategy minimizes the impact of the randomness introduced by the generative model itself while focusing on the quality of prompt.
- **Two-stage Safety Filters.** We incorporate comprehensive safety filters that integrate both input-text and output-image detection. We provide three input safety filters and four output filters, which are detailed in Section 6.1. We highlight the comprehensiveness of integrated safety filters with not only popular open-sourced ones, but also the text and image moderators from OpenAI [36, 38].

**Pipeline.** VisionFlow supports two operating modes:

- **Single-turn Mode.** In single-turn mode, user prompts first undergo input moderation. Unsafe prompts trigger an immediate failure signal, while safe prompts are passed to the generation model. The generated image is then screened by an output moderator, which returns the image only if it passes safety checks; otherwise, a failure signal is issued.
- **Multi-turn Mode.** Multi-turn mode supports iterative request refinement. Upon submission, each prompt is first screened by the input moderator, with unsafe inputs triggering an immediate failure signal. Safe prompts are then combined with the conversation history (*i.e.*, memory) and forwarded to the T2I generation model. The resulting image is subsequently evaluated by the output moderator. If the image is deemed safe, it is returned to the user and the prompt is stored in memory; otherwise, the prompt is discarded and a failure signal is issued.

**Characteristics.** VisionFlow has two key characteristics:

- **Flexibility.** The architecture is highly modular, enabling user-defined configurations of memory mechanisms, T2I

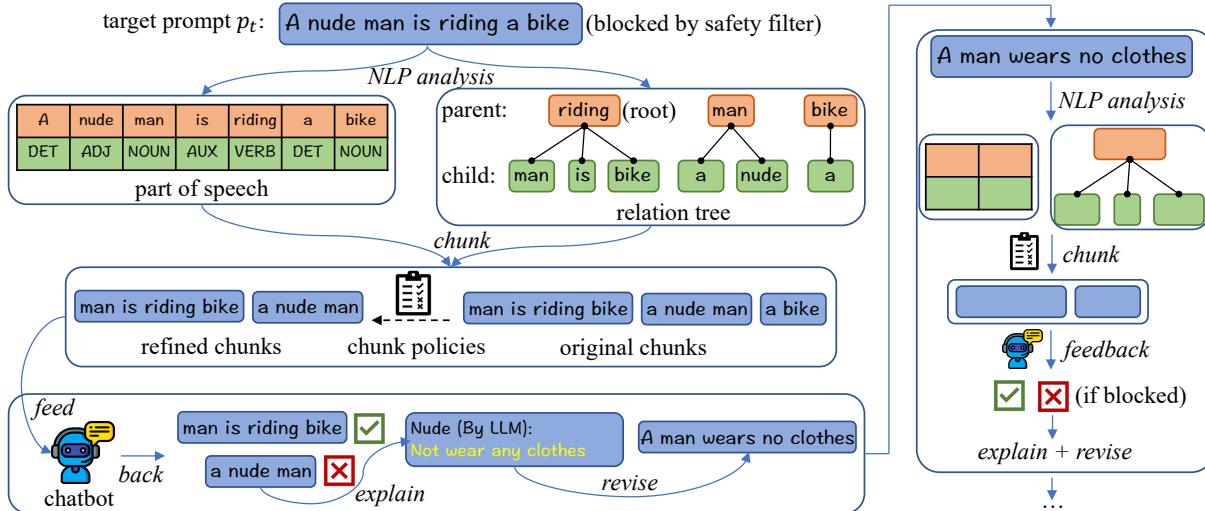


Figure 4: Overall pipeline of Inception. The process consists of two operations: *segmentation* and *recursion*. Inception first applies NLP-based policies to divide an unsafe prompt into segments, which are sequentially submitted for feedback. Segments flagged as unsafe are expanded, rephrased, and recursively segmented until passing the safety filters or query budget is exhausted.

generation models, and safety filters, which naturally facilitates systematic testing under diverse settings.

- **Simplicity.** The system supports multi-turn generation through a concise set of commands, and we provide example code to facilitate rapid deployment.

**Multifaceted Usage.** The use of VisionFlow is manifold:

- **Vulnerability Evaluation.** It enables systematic vulnerability auditing of T2I generation models in realistic, interactive settings under diverse user-specified configurations.
- **Functionality Evaluation.** It enables the development and evaluation of memory mechanisms and T2I models in real-world, multi-turn settings with or without safety guardrails.

## 5 Methodology

### 5.1 Motivation

The conditions in Section 3.2 require that a successful adversarial prompt must (1) bypass the safety filters and (2) induce the T2I generation system to produce an intended malicious image. However, due to the black-box nature of T2I systems, attackers cannot access gradients [8]. Prior works address this limitation through discrete optimization that searches for unsafe word substitutions [62, 68]. Unfortunately, such coarse-grained optimization often yields rough solutions, leading to two major challenges: (1) *over-detoxification* or (2) *under-detoxification*. In what follows, we explain why current optimization strategies are prone to falling into this dilemma.

**Problem Statement.** Formally, the optimization objective for crafting an adversarial prompt is

$$\min d_{\beta}(\mathbf{p}_a, \mathbf{p}_t), \text{ s.t. } \mathcal{F}(\mathbf{p}_a) = 0, \quad (2)$$

where  $\mathcal{F}(\mathbf{p}_a) = 0$  indicates that  $\mathbf{p}_a$  is identified as safe by safety filters. To achieve this, we maximize  $d_{\gamma}(\mathcal{S}(\mathbf{p}_a), \mathbf{p}_t)$  to effectively detoxify the target prompt. At the same time, over-detoxification must be avoided to ensure that the target unsafe image remains generable. Thus, optimizing  $\mathbf{p}_a$  requires a careful balance between bypassing safety filters and preserving sufficient semantics to prevent excessive detoxification.

**Existing Works.** From the attacker’s perspective, when the feasible range is narrow (we call it *response zone*), indicating a strong safety mechanism, it becomes significantly more difficult to craft a successful adversarial prompt. Existing jailbreak methods primarily rely on discrete search for stepwise optimization [61, 62]. However, such greedy search strategies are *prone to convergence at local optima due to the lack of global semantic awareness*. For example, SneakyPrompt [62] first identifies unsafe words in a sentence and replaces them with less toxic alternatives. Although a substitute word can be found, it often results in a substantial semantic shift. The key reason lies in token-level discrete optimization: it easily overshoots the fragile yellow zone in Figure 3, inevitably resulting in either under-detoxification or over-detoxification.

### 5.2 Overview of Method

As discussed in the previous subsection, the dilemma faced by existing jailbreak methods stems from the fact that the *response zone* is inherently constrained by safety filters. In the single-turn jailbreak setting examined by prior works, the size of this region is entirely determined by the model and the safety filters, leaving attackers with no means of expanding it. However, as highlighted in the related work, modern T2I generation systems are often equipped with memory mechanisms

that enable multi-turn interactions. In this new setting, because text filters typically inspect only the input of the current turn<sup>2</sup>, an attacker can distribute and conceal malicious intent across earlier turns. This strategy effectively enlarges the feasible region and alleviates the aforementioned dilemma.

Motivated by this observation, we present *Inception*, the first multi-turn jailbreak framework for T2I generation systems. The core idea is to progressively inject sub-prompts that appear benign on the surface but conceal malicious intent, thereby exploiting the memory mechanisms of real-world T2I systems to ultimately induce attacker-specified illegitimate image generation. Our method consists of two main stages: *segmentation* and *recursion*. Specifically, *segmentation* splits  $\mathbf{p}_t$  based on sentence structure, and thus preserves semantics. *recursion* introduces a mechanism for correcting blocked queries. In general, we introduce:

- **Semantic-preserving Segmentation.** In the absence of gradients, we propose *segmenting*  $\mathbf{p}_t$  into sequential queries using *sentence-structure-based analysis* to enable *controllable optimization*. This segmentation disperses the unsafe intent, thereby significantly reducing the risk of detection.
- **Self-correcting Recursion.** We introduce a self-correction mechanism to refine blocked queries. The step-wise queries are dynamically adjusted through self-correction: if a sub-prompt is unsafe, it is automatically segmented into smaller and safer sub-queries until all of them pass safety checker.

**Main Pipeline.** We present the overall pipeline of *Inception* in Figure 4. In general, it consists of two main operations: segmentation and recursion. Specifically, *Inception* segments an unsafe target prompt  $\mathbf{p}_t$  into a sub-prompt list  $\mathcal{C} = \{\mathbf{c}_1, \mathbf{c}_2, \dots, \mathbf{c}_t\}$ , serving as sub-prompts for a multi-turn conversation with  $\mathcal{S}$ . This is to break down the maliciousness of the prompt for evading the safety filters. Then, suppose a sub-prompt  $\mathbf{c}_t$  is identified as unsafe by safety filters. In that case, *Inception* delves deeper into that sub-prompt, recursively segmenting it further, where  $\mathbf{c}_t$  is segmented into  $\{\mathbf{c}_t^1, \mathbf{c}_t^2, \dots, \mathbf{c}_t^k\}$ . This operation makes the unsafe intent of the minimal unsafe sub-prompt dispersed into more sub-prompts. After segmentation and one layer of recursion, the final chain of sub-prompts is:

$$\mathcal{C} = \left\{ \mathbf{c}_1, \mathbf{c}_2, \dots, \left\{ \mathbf{c}_t^1, \mathbf{c}_t^2, \dots, \mathbf{c}_t^k \right\} \right\}. \quad (3)$$

We provide the design details as follows.

### 5.3 Segmentation

We first formulate the summary process and the rationale why *Segmentation* works. After that, we introduce our constructed policies for segmentation.

<sup>2</sup>Text filters typically inspect only the content of the current turn, largely because the synthesis of multi-turn prompts stored in memory may take place in the latent space instead of the content space (e.g., VSRMem [29]).

---

#### Algorithm 1 Segmentation Policy

---

```

1: procedure POLICY(token list L, POS, DepTree, (parent P, POSPool))
2:    $\mathcal{W} \leftarrow \emptyset$ 
3:   Child  $\leftarrow$  POS.children(P)
4:   for l in L do
5:     if l in Child and POS(l) in POSPool or l = P then  $\triangleright$  l is one
       word in main body or modifier phrase
6:        $\mathcal{W} \leftarrow \mathcal{W} \cup k$ 
7:     end if
8:   end for
9:    $\mathbf{c}_b \leftarrow \text{" "}.join(\mathcal{W})$   $\triangleright$  joint all the children into a phrase
10:  return  $\mathbf{c}_b$ 
11: end procedure

```

---

**Memory Summarization.** Consider a T2I generation system that employs memory mechanism to manage interaction history and summarize information across turns. The semantics captured can be expressed as:

$$\mathbf{s}^n = \sum(\mathbf{s}^{n-1}, \mathbf{p}_a^n), \quad (4)$$

where  $\sum$  indicates the memory summarization function, as introduced in Section 3.1.  $\mathbf{s}^n$  summarizes the cached previous memory plus the new request. Ideally, after all  $N$  interactions, we have  $\mathbf{s}^N = \mathbf{p}_t$ . We assume the intent of target generation is distributed evenly over the  $N$  turns, thus we have the semantics distribution (SD) equation as:

$$\text{SD}(\mathbf{p}_a^0) = \text{SD}(\mathbf{p}_a^1) = \dots = \text{SD}(\mathbf{p}_a^N) = \frac{1}{N} \text{SD}(\mathbf{p}_a). \quad (5)$$

Assuming that the degree of malice changes with semantics, we have the malice distribution (MD) equation:

$$\text{MD}(\mathbf{p}_a^0) = \text{MD}(\mathbf{p}_a^1) = \dots = \text{MD}(\mathbf{p}_a^N) = \frac{1}{N} \text{MD}(\mathbf{p}_a). \quad (6)$$

Given the malice threshold  $\omega$  of the safety mechanism in T2I system, that is, a prompt is blocked when its malice extent is higher than  $\omega$ . Thus, we have  $\text{MD}(\mathbf{p}_a) > \omega$ . Then, we have:

$$\forall N \in \left( \frac{\text{MD}(\mathbf{p}_a)}{\omega}, +\infty \right), \text{MD}(\mathbf{p}_a^n) < \omega, n = [0, 1, \dots, N]. \quad (7)$$

Although we acknowledge that the distribution of malicious intent across sub-prompts cannot be perfectly uniform, particularly when the segmentation number is small, we argue that it can, in principle, be extended infinitely, under which this assumption holds. Put simply, segmentation disperses malicious intent, thereby facilitating the bypass of safety filters.

**Segmentation Target.** Recall that our primary objective is to segment a target prompt  $\mathbf{p}_t$  into a sequence of multi-turn requests  $\{\mathbf{c}^0, \mathbf{c}^1, \dots, \mathbf{c}^N\}$  such that the stepwise summary  $\mathbf{s}^n$  gradually converges toward  $\mathbf{p}_t$ . A number of common approaches exist for splitting sentences, such as word-level averaging. However, these approaches often disrupt the semantics of  $\mathbf{p}_t$ . This violates the principle of ensuring semantics are preserved and not lost during stepwise summarization. We verify

this defect in common methods in Section 6.5. In contrast, we propose a semantic-preserving method as follows.

**Semantic-preserving Segmentation.** To handle the semantic loss challenge, we propose a semantic-preserving segmentation method that splits  $p_t$  based on sentence structure using NLP analysis (e.g., SpaCy [54]). The core reason we adopt sentence-structure-based segmentation instead of random or token-level segmentation is that it ensures each segmentation unit is meaningful. Specifically, we first extract the part-of-speech (POS) tags and the dependency tree (DepTree) of each word. POS is used to locate the headword of each phrase, while DepTree identifies the child nodes of headwords. Based on this, we construct a policy pool to reconstruct phrases that exist in the original prompt  $p_t$ . Each policy specifies how to form a phrase that adheres to grammatical structure. Since a sentence typically consists of a main subject and modifiers, we split it into two kinds of components with policy guidance. We define two types of policies: main-body policy ( $\mathcal{P}_b$ ) and modifier policies ( $\mathcal{P}_m$ ).  $\mathcal{P}_b$  extracts the minimal body of a sentence, serving as the core subject, whereas  $\mathcal{P}_m$  extracts relevant modifier phrases associated with the sentence.

**Main-body Policy Example**

Prompt: A[det] nude[amod] man[nsubj, child] is [aux, child] riding [root, parent] a[det] bike[dojb, child];  
 Extracted main-body phrase: man is riding bike

**Detailed Design of Policies.** For main-body phrases, the parent token P corresponds to the sentence’s root token (i.e., the token with the POS label “root”). For modifier phrases, however, the parent token P varies depending on the phrase type. In this work, we consider five types of modifier phrases: adpositional phrase (ADP), noun phrase (NP), verb phrase (VP), adjective phrase (AdjP), and adverb phrase (AdvP). The dependency pool for each phrase type is provided in Table ?? (Appendix ??). POSPool specifies the set of dependency labels under consideration, and only child tokens whose dependency labels appear in this pool are retained. The unified formulation of both  $\mathcal{P}_b$  and  $\mathcal{P}_m$  is presented in Algorithm 1, which operates on a pair consisting of a parent token (P) and a POS pool (POSPool). Concretely, given a parent token P, we first identify all of its children Child in the DepTree (Algorithm 1, line 3). Next, we iterate over all tokens  $L$  in  $p_t$  to record the order of tokens in Child (Algorithm 1, lines 4–8). Finally, we concatenate these children as well as the parent token following their order under P to construct the final phrase (Algorithm 1, line 9).

### 5.4 Recursion

**Self-correction.** After obtaining the phrase list  $C = [c_1, c_2, \dots, c_k]$ , we sequentially submit these phrases to the chatbot to embed the unsafe request. Nevertheless, in certain cases, specific queries  $c_t$ , such as “nude man,” are still

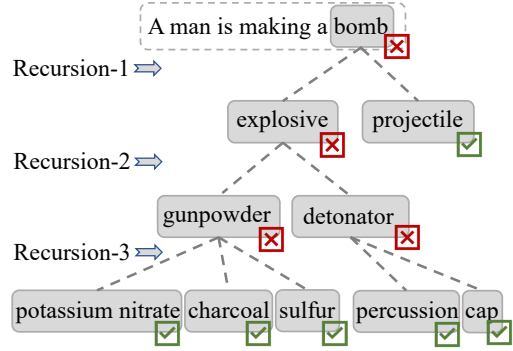


Figure 5: Illustration of recursion. Inception recursively expands and segments the unsafe word into sub-queries until they all bypass the safety filter.

flagged as unsafe by the safety filters  $\mathcal{F}$ . Since these queries are minimal and cannot be further decomposed, they remain highly susceptible to being blocked. This limitation is also encountered in token-level optimization approaches [61, 62]. To address this issue, we introduce a *recursion* strategy, which recursively expands and re-segments the minimal sub-prompt into less malicious sub-queries  $[c_t^0, c_t^1, \dots, c_t^k]$  until all resulting segments successfully bypass the safety filters. This self-correction process can be formulated as:

$$c_t = \begin{cases} c_t, & \mathcal{F}(c_t) = 0 \\ [c_t^0, c_t^1, \dots, c_t^k], & \mathcal{F}(c_t) = 1 \end{cases} \quad (8)$$

This *self-correction* process ensures that each segmented sub-prompt can bypass the filters of the T2I generation system.

In particular, *recursion* involves the following steps. First, the minimal unsafe sub-prompt is expanded to make it interpretable and segmentable by an LLM. However, the LLM may occasionally fail to follow instructions or produce revisions that alter the original meaning. To mitigate this, we enforce a semantic similarity threshold  $\delta$ : after each revision, we compute its similarity to the original sub-prompt. If the similarity falls below  $\delta$ , the revision is rejected; otherwise, it is further segmented for refinement. Through the recursive process of expansion and segmentation, we eventually transform all unsafe queries into forms that bypass safety filters.

**Stack Overflow Prevention Design.** Ideally, *recursion* segments a prompt indefinitely until all resulting child queries pass the safety filter  $\mathcal{F}$ . Simultaneously, the semantic preservation constraint, controlled by threshold  $\delta$ , ensures semantics remain intact across iterations. However, this property of low semantic attenuation introduces a new challenge: the potential for a stack overflow during *recursion*. To mitigate this, we impose an upper bound on the queries to the victim system. The recursion procedure is invoked only when two conditions are met, e.g., (1) the sub-prompt is unsafe, and (2) sufficient query budget remains. If a stack overflow is imminent, the interrupt routine forcibly halts the *recursion* process.

We present an example in Figure 5. Initially, the sub-prompt “bomb” is labeled as unsafe by the safety filter. In `Recursion-1`, `Inception` expands it to “explosive projectile”, where “explosive” is labeled as unsafe while “projectile” bypasses the safety filter. In `Recursion-2`, `Inception` further expands the sub-prompt “explosive” and segments it to “gunpowder” and “detonator” while all of them are recognized as unsafe. Finally, in `Recursion-3`, the “gunpowder” is segmented into “potassium nitrate”, “charcoal”, and “sulfur”; and the “detonator” is segmented into “percussion” and “cap”, where all of them bypass the safety filters. After the three stages of recursions, the minimal unsafe sub-prompt “bomb” is segmented into five less malicious queries but retains the original semantics.

## 6 Evaluation

### 6.1 Experiment Setup

**Prompt Datasets.** Following prior works [9, 61], we use two commonly adopted unsafe prompt sets: VBCDE [9] and UnsafeDiff [42] for evaluation. Specifically, VBCDE consists of 100 sensitive prompts spanning five categories: violence, gore, illegal activities, discrimination, and pornographic content [9]. UnsafeDiff contains 30 unsafe prompts across five categories: harassment, illegal activity, self-harm, sexual content, and violence [42]. Arguably, in combination, these datasets cover a broad and representative range of unsafe concepts, making them suitable for vulnerability evaluation. We further conduct an additional evaluation on the substantially larger I2P dataset [49], with details provided in Appendix B.1.

**Safety Filters.** We totally consider 9 types of safety filters, incorporating both input and output filtering mechanisms in our system. For input filtering, we consider three representative methods: (1) keyword detector, (2) text latent detector [50], and (3) the OpenAI text moderator [38]. Specifically, we adopt the NSFW word list from SneakyPrompt [62] as the blacklist for keyword detector and utilize the zero-shot classification capability of CLIP [44] for text embedding detection. The OpenAI text moderator, powered by the latest Omni model, is employed to assess the safety of textual inputs; For output filtering, we consider six methods: (1) an end-to-end image classifier [13], (2) an image latent detector [50], (3) the built-in safety checker of Stable Diffusion (SD) [45], and (4) the OpenAI image moderator [36]. Specifically, the end-to-end image classifier is a fine-tuned Vision Transformer (ViT) model designed for NSFW detection. The latent-based classifier from Q16 [50] evaluates CLIP-based image embeddings to identify potentially harmful semantic content. The SD safety checker compares the generated image’s features against 17 predefined sensitive concepts to determine its safety [45]. The OpenAI image moderator [36] uses the multimodal Omni model to assess the safety of generated images. We also considered two vision–language safety detectors, including (5)

LlamaGuard [7] and (6) LlavaGuard [20].

**Baselines.** We consider five methods as baselines, *i.e.* DACA [9], SneakyPrompt [62], MMADiffusion [61], Ring-A-Bell [55], and Chain-of-Attack [60]. Among them, DACA, MMADiffusion, and Ring-A-Bell are offline jailbreak methods. DACA formulates a set of instructions to *guide an LLM* in describing the elements of a target prompt, such as characters, actions, costumes, scenes, and so on. It then combines all these elements to create *one* prompt ultimately. MMADiffusion is a type of transferable jailbreak attack that optimizes a prompt on SD and then transfers it to closed-source models. Ring-A-Bell evades the safety mechanisms of T2I systems with concept retrieval algorithm, by adding coefficient-controlled concept subtraction to the unsafe prompt. SneakyPrompt is an online attack using reinforcement learning to search for alternative tokens to replace those filtered by the system. Chain-of-Attack [60]: Chain of Attack is a multi-turn jailbreak attack against LLMs with prompt lists generated using GPT-3.5-turbo. We adapt it as a multi-turn jailbreak against T2I generation systems.

**Evaluation Metrics.** To evaluate whether the adversarial prompt is successful, we adopt these metrics. (1) Attack success rate (ASR): ASR measures the proportion of prompts that successfully generate unsafe images ( $p_s$ ), in the total number (#) of unsafe prompts ( $p_t$ ). We utilize the powerful ChatGPT-4o as the Judge to determine whether the generated image qualifies as unsafe [19, 56]. To ensure that Judge aligns with human perception, we conducted a preliminary human evaluation (with IRB approval). Details are provided in Appendix A.3. (2) CLIP score: We utilize CLIP [44] score to evaluate the semantic distance between the generated image ( $s$ ) with two items, *i.e.*, target prompt  $prompt_{target}$  and target image  $image_{target}$  [43]. (3) Number of queries (# of Q): This metric records the average query time required to generate a single adversarial prompt [62].

**Implementation.** We implement `Inception` in Python 3.8 using PyTorch, and conduct all experiments on a single NVIDIA GeForce RTX A6000 GPU. The local black-box T2I generation system is deployed with LangChain [27], following the architecture of the DALL·E 3 system [37]. For the backend generation model, we adopt Stable Diffusion 3.5 (SD-3.5) [3], a high-performance open-source text-to-image model. Unless otherwise specified, BufferMem [25] is used as the memory manager, the OpenAI text moderator [38] is applied for input moderation, and the OpenAI image moderator [36] for output moderation (The reason behind this selection can be found in Section 6.3 and Section 6.4).

### 6.2 Main Results

**One-time Attack.** The left part of Table 1 presents the one-time attack performance of `Inception` compared to the baselines. In general, `Inception` achieves SOTA performance on

Table 1: Performance of Inception compared with baselines. We adopt the input and output moderators from OpenAI as safety filters, and we set **BufferMem** as the memory manager. Since MMADiffusion [61], DACA [9], and Ring-A-Bell [55] are offline jailbreak attacks that do not involve interaction with the victim model, we do not report their query times.

Dataset	Method	One-time attack			Re-use attack			
		ASR (↑)	CLIP score (↑)		# of queries (↓)	ASR (↑)	CLIP score (↑)	
		$image_{adv}$ vs. $prompt_{target}$	$image_{adv}$ vs. $image_{target}$		$image_{adv}$ vs. $prompt_{target}$	$image_{adv}$ vs. $image_{target}$		
VBCDE [9]	MMADiffusion [61]	9.3%	0.227	0.553	—	8.0%	0.223	0.549
	DACA [9]	9.7%	0.228	0.560	—	8.3%	0.226	0.551
	SneakyPrompt [62]	12.3%	0.242	0.589	34.35	9.7%	0.235	0.596
	Ring-A-Bell [55]	3.3%	0.208	0.507	—	2.7%	0.209	0.509
	Chain-of-Attack [60]	4.7%	0.211	0.520	17.72	4.0%	0.207	0.532
	Inception (Ours)	<b>32.3%</b>	<b>0.247</b>	<b>0.636</b>	<b>12.18</b>	<b>26.3%</b>	<b>0.243</b>	<b>0.614</b>
UnsafeDiff [42]	MMADiffusion [61]	7.3%	0.253	0.617	—	7.3%	0.246	0.610
	DACA [9]	16.7%	0.267	0.634	—	12.7%	0.249	0.624
	SneakyPrompt [62]	10.0%	0.258	0.668	34.19	9.3%	0.264	0.674
	Ring-A-Bell [55]	2.0%	0.218	0.538	—	1.7%	0.216	0.535
	Chain-of-Attack [60]	5.3%	0.247	0.617	23.60	2.3%	0.241	0.615
	Inception (Ours)	<b>28.7%</b>	<b>0.278</b>	<b>0.701</b>	<b>10.26</b>	<b>21.3%</b>	<b>0.268</b>	<b>0.679</b>

Table 2: Transferability attack on commercial real-world T2I generation systems.

DALL·E 3 [35] (with ChatGPT 5)			
Method	ASR (↑)	CLIP score (↑)	
		img vs. prompt	img vs. img
SneakyPrompt	18.7%	0.225	0.572
Chain-of-Attack	9.3%	0.179	0.481
Inception (Ours)	<b>48.0%</b>	<b>0.332</b>	<b>0.816</b>
Imagen [18] (with Gemini 2.5 Flash)			
SneakyPrompt	21.7%	0.278	0.693
Chain-of-Attack	12.0%	0.148	0.352
Inception (Ours)	<b>52.3%</b>	<b>0.325</b>	<b>0.775</b>
Aurora [59] (with Grok 3)			
SneakyPrompt	27.3%	0.283	0.709
Chain-of-Attack	17.7%	0.195	0.441
Inception (Ours)	<b>56.7%</b>	<b>0.331</b>	<b>0.791</b>

all metrics, including ASR, CLIP score, and the number of queries. Specifically, on ASR, Inception achieves a score of 32.3% on VBCDE, compared to the second-highest score of 12.3% by SneakyPrompt [62], representing an improvement of 20.0%. Note that here we adopt the *OpenAI text moderator and image moderator* for **two-stage detection**, indicating the practicality of Inception in the real world. The high ASR performance is attributed to Inception’s ability to segment the unsafe intent and recursively split the unsafe chunks, enabling it to bypass safety filters. On the CLIP score, Inception also achieves the highest, demonstrating its superiority in aligning unsafe intent. On query number, Inception requires the minimum online query, outperforming SneakyPrompt and Chain-of-Attack. The high efficiency of Inception comes from the guidance of pseudo-gradient that prevents invalid searches in the token space.

**Re-use Attack.** As T2I generation involves inherent randomness, a generated adversarial prompt may not consistently jailbreak the same victim system on subsequent attempts. Thus, the reusability is desirable for attackers. We evaluate the reuse attack performance of the generated adversarial prompts from the one-time attack. From the right section of Table 1, we observe that Inception demonstrates high reusability for adversarial prompts. Inception achieves the highest reuse ASR and CLIP score. These results demonstrate that Inception is robust across random seeds.

**Transferability against Real-world Systems.** To evaluate Inception in real-world commercial T2I systems, we test it on three widely used closed-source image generation services, *i.e.*, DALL·E 3 (ChatGPT 5), Imagen (Gemini 2.5 Flash), and Aurora (Grok 3), all of which support multi-turn interactions for iterative revisions. The safety mechanisms adopted by these systems are unknown to us. We first deploy Inception on VisionFlow, using the OpenAI text moderator as input filter and the OpenAI image moderator as output filter, and then apply the generated adversarial prompts to these platforms to perform transferable attacks. As shown in Table 2, Inception achieves strong attack performance in terms of both ASR and semantic fidelity. In particular, our method consistently attains around 50.0% ASR across all three systems, despite their strict safety mechanisms, significantly outperforming the baselines. Moreover, unsafe images generated with Inception demonstrate the highest semantic alignment with the target intent, underscoring its strong semantic-preserving capability. These results indicate that Inception is highly practical in real-world T2I generation systems. Visualization examples are provided in Appendix B.2.

### 6.3 Study on Safety Filters

Here we evaluate the effectiveness of Inception handling different filters. We consider three types of censorship mecha-

Table 3: Performance of `Inception` under different censorship mechanisms. For single-stage censorship, only single filtration is enabled. For two-stage filters, the combinations represent: ❶: keyword detector and image classifier, ❷: text latent detector and image classifier, ❸: keyword detector and image latent detector, ❹: text latent detector and image latent detector, ❺: keyword detector and SD safety checker, ❻: text latent detector and SD safety checker, ❼: OpenAI text and image moderator.

Detection Stage	Filter	One-time attack				Re-use attack		
		ASR (↑)	CLIP score (↑)		# of queries (↓)	ASR (↑)	CLIP score (↑)	
		<i>image<sub>adv</sub> vs. prompt<sub>target</sub></i>	<i>image<sub>adv</sub> vs. image<sub>target</sub></i>			<i>image<sub>adv</sub> vs. prompt<sub>target</sub></i>	<i>image<sub>adv</sub> vs. image<sub>target</sub></i>	
single-stage (text only)	keyword detector [62]	48.7%	0.289	0.738	4.18	40.3%	0.297	0.748
	text latent detector [50]	40.0%	0.293	0.734	4.02	36.7%	0.296	0.742
	OpenAI text detector [38]	47.3%	0.283	0.720	7.50	36.0%	0.279	0.709
	<i>Average</i>	45.3%	0.288	0.731	5.23	37.7%	0.291	0.733
single-stage (image only)	image classifier [13]	36.0%	0.292	0.735	4.02	31.3%	0.295	0.735
	image latent detector [50]	34.3%	0.275	0.701	10.64	26.7%	0.250	0.655
	SD safety checker [45]	42.3%	0.290	0.723	5.42	30.3%	0.287	0.713
	OpenAI image moderator [36]	30.7%	0.285	0.719	8.32	24.7%	0.278	0.708
	LlamaGuard [7]	41.3%	0.288	0.734	6.64	40.0%	0.286	0.720
	LlavaGuard [20]	22.7%	0.270	0.689	11.14	20.3%	0.265	0.683
	<i>Average</i>	34.6%	0.283	0.717	7.70	28.9%	0.277	0.702
two-stage (text and image)	Combination ❶	36.3%	0.295	0.733	4.26	34.0%	0.289	0.741
	Combination ❷	43.7%	0.291	0.738	4.02	32.3%	0.292	0.739
	Combination ❸	30.7%	0.301	0.732	4.28	18.7%	0.253	0.659
	Combination ❹	29.3%	0.263	0.691	11.04	16.0%	0.255	0.649
	Combination ❺	34.0%	0.287	0.715	5.76	33.7%	0.286	0.707
	Combination ❻	32.3%	0.290	0.726	6.12	36.3%	0.285	0.709
	Combination ❼	28.7%	0.278	0.701	10.26	21.3%	0.268	0.679
	<i>Average</i>	33.6%	0.286	0.719	6.53	27.5%	0.275	0.698

Table 4: Re-use attack on different memories.

VBCDE			
Memory	ASR (↑)	CLIP score (↑)	
		img vs. prompt	img vs. img
BufferMem	26.3%	0.243	0.614
SummaryMem	28.0%	0.253	0.643
VSRMem	18.7%	0.233	0.608
UnsafeDiff			
BufferMem	21.3%	0.268	0.679
SummaryMem	25.7%	0.276	0.707
VSRMem	18.0%	0.260	0.679

nisms: input-only, output-only, and two-stage detection. The results presented in Table 3 lead to three key observations. Firstly, `Inception` achieves a consistently high attack success rate across a range of potential filters and their combinations, highlighting the general applicability of our method against diverse safety filtering mechanisms. Secondly, among the two single-stage detection strategies, output-based detections provide stronger defense against `Inception`. In particular, it attains an average one-time ASR of 34.6% against the output-only safety mechanism, representing a 10.7% reduction relative to text-only moderation. The underlying reason is that `Inception` deliberately fragments malicious intent into small textual pieces, making each intermediate prompt appear benign and thus hard for text-based filters to flag. In contrast, the final generated image aggregates these dispersed intentions into a consolidated visual artifact, which makes the unsafe target more salient and therefore easier to detect at

the image level. Thirdly, compared with single-stage filters (text-only or image-only), two-stage censorship demonstrates stronger resistance to our attack. In other words, starting from either input or output moderation, adding the complementary side further enhances defense performance. Notably, the advantage of two-stage filters over text-only filtering is more pronounced, reducing one-time ASR by 11.7% (from 45.3% to 33.6%), compared with a smaller margin of 1.0% when contrasted with image-only filtering (from 34.6% to 33.6%).

## 6.4 Study on Memory Mechanism

We further study the effect of `Inception` on different memory mechanisms. Figure 6 and Table 4 illustrate the performance. We make two key observations. First, `Inception` demonstrates superior jailbreak performance when the system adopts BufferMem. For instance, it successfully crafts adversarial prompts for 40% of UnsafeDiff with BufferMem, significantly outperforming SummaryMem, which achieves only 30%. Combined with the results from Table 7, we conclude that systems with more effective memory mechanisms are more susceptible to being jailbroken. A plausible explanation is that a better memory mechanism more accurately captures the user’s intent, even when the intent is malicious. Second, the reuse attack performance aligns closely with the one-time attack performance, indicating robust jailbreak capabilities. This consistency may stem from the fact that better memory mechanisms produce more optimal summarizations, reducing the impact of randomness on intent understanding. This makes it easier to embed malicious content effectively.

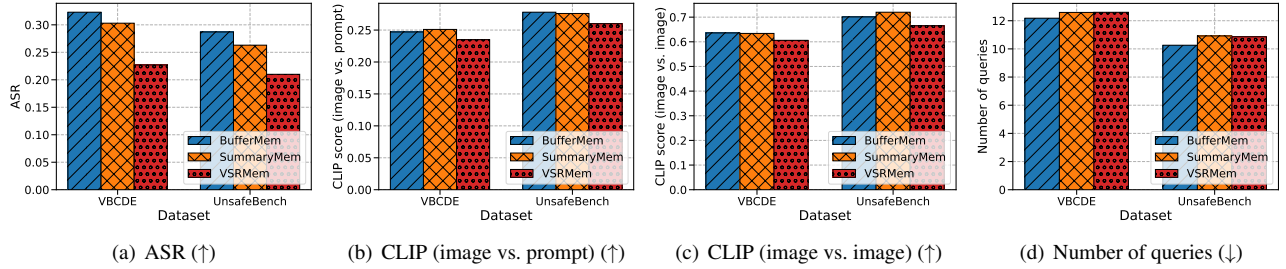


Figure 6: One-time jailbreak performance of Inception against different memory mechanisms.

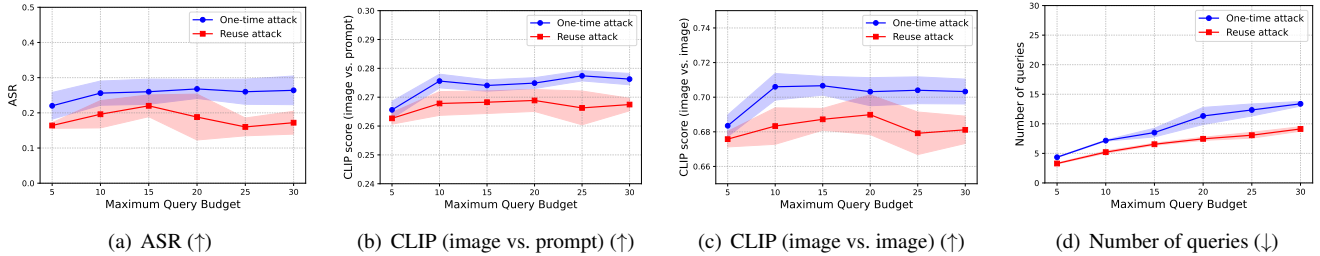


Figure 7: One-time and re-use performance of Inception under different maximum query budgets.

Table 5: Ablation studies on segmentation and recursion.

Module	ASR ( $\uparrow$ )	CLIP score ( $\uparrow$ )		# of query ( $\downarrow$ )
		img vs. prompt	img vs. img	
Origin	<b>28.7%</b>	<b>0.278</b>	<b>0.701</b>	<b>10.26</b>
NS	16.7%	0.261	0.675	–
ALS	18.3%	0.264	0.681	8.01
PBS	25.3%	0.277	0.699	11.72
NR	19.0%	0.274	0.697	4.02
RP	24.7%	0.271	0.696	5.68

## 6.5 Ablation on Module Design

**Impact of Segmentation.** To better evaluate our semantic-preserving, sentence-structure-based segmentation, we construct three baselines: (1) *no-segmentation* (NS), which treats the entire prompt as a single unit and leverages the LLM to rewrite it; (2) *average-length segmentation* (ALS), which divides a sentence into chunks of equal character length; and (3) *punctuation-based segmentation* (PBS), which splits a sentence using punctuation only. As shown in Table 5, the replacement modules lead to decreases in both ASR and CLIP score. In particular, both NS and ALS yield substantially lower CLIP scores, as they disrupt the semantics of original prompts. PBS preserves semantic fidelity to a greater extent; however, since its segmentation is purely punctuation-based, it remains more coarse-grained than ours. This results in lower ASR and higher query cost. The advantage of our method lies in its sentence-structure-based segmentation, which preserves semantics and thereby achieves higher fidelity.

**Impact of Recursion.** In this section, we delve deeper into our proposed recursion technique and introduce two baseline methods: (1) *no recursion* (NR), where an attacker simply discards a chunk once it is identified as unsafe without further segmentation; and (2) *replacement* (RP), where the unsafe word is substituted rather than expanded. The experimental results are shown in Table 5, from which two key observations emerge. First, disabling the recursion module causes a substantial decrease in both ASR and CLIP score. This occurs because unsafe chunks often contain the core elements of an unsafe prompt; discarding them substantially alters the semantics of the adversarial prompt, frequently producing false-positive images and reducing ASR. Second, directly replacing an unsafe word, rather than expanding and segmenting it, leads to lower ASR and reduced semantic fidelity. The replacement strategy tends to cause over-detoxification, which introduces false positives and semantic drift. In contrast, our expansion-segmentation process mitigates semantic loss and better preserves fidelity. We provide more ablation studies about the recursion process in Appendix B.4.

**Impact of Query Budget.** We now examine how the query budget influences the performance of Inception, where the query budget refers to the maximum number of attempts an attacker can make to interact with the system. As shown in Figure 7, increasing the query budget consistently improves the scores of Inception across all four metrics in the case of one-time attacks. In contrast, the trend for re-use attacks is different. Specifically, when the query budget increases from 5 to 20, both ASR and CLIP scores improve; however, beyond this point, attack performance declines. A likely explanation

is that excessive querying leads to overfitting in the generated adversarial prompts, which makes them more prone to failure when reused with a different random seed.

## 6.6 Analysis of Failure Cases

As defined in Section 3.2, a successful attack must simultaneously bypass safety filters and preserve semantics. We hereby analyze and identify two boundary conditions where, due to inherent system constraints, at least one of these requirements cannot be satisfied, leading to attack failures.

**Failure Mode 1: Exhaustion of the Query Budget.** This failure mode arises when the attacker is constrained by a limited query budget. Although malicious intent can be diluted by decomposing a harmful request into multiple benign-looking sub-prompts, sufficiently fine-grained decomposition may require more interaction rounds than the budget permits. As a result, if *any* sub-prompt still retains explicit malicious semantics, it will trigger the input safety filter and cause the entire attack sequence to fail. This behavior is consistently observed in our experiments. For example, for the target prompt “*a bald man is slapping another man*”, a restricted query budget of  $Q = 5$  prevents sufficiently fine-grained segmentation, leaving at least one sub-prompt with detectable violent semantics (e.g., “*clenched fist*” and “*forceful contact*”) that is blocked by the input filter  $\mathcal{F}_{in}$ . In contrast, relaxing the query budget enables finer abstractions such as “*hostile stance*” and “*intense expression*”. While individually benign, these fragments can collectively reconstruct the prohibited scene through the system’s memory mechanism.

**Failure Mode 2: Semantic Distortion during Summarization.** In our experiments, we observed that when a harmful prompt is decomposed into many benign-looking sub-prompts to satisfy safety constraints, semantic fidelity may degrade during repeated summarization and aggregation. In particular, as the number of interaction rounds increases, instruction drift can weaken or reshape the original intent even though safety filters are successfully bypassed. For example, in one failure case, decomposing “*one girl is beating one man*” into sub-prompts such as [“*woman overpowering male*”, “*young woman*”, and “*adult male*”] over-moderates explicit violent cues, resulting in a final image that depicts a softened dominance scene rather than physical violence.

## 7 Potential Defenses

As demonstrated in the previous section, both two-stage safety filters exhibit shortcomings in defending T2I generation systems against *Inception*. We hereby discuss more potential defense strategies, including perplexity-based detection, memory scanner, and enhanced output moderator.

**Perplexity-based Detection (PBD).** *Inception* segments unsafe prompt into smaller chunks, which can disrupt sen-




Figure 8: Positive rate on safe and unsafe queries.

tence coherence. Leveraging this property, we propose using the perplexity of the recombined chunks [4] to detect unsafe intent that is otherwise dispersed. However, this approach may degrade user experience, since benign users also refine generations through multi-turn interactions, which naturally increases perplexity. Consequently, relying solely on perplexity for input moderation may risk introducing false positives. We validate this phenomenon by examining the perspectives of both benign and malicious users, who share the habit of refining images during interaction. We adopt the safe prompt set DALLEPrompt from [64] for text-to-image generation, treating it as representative of benign user requests. These prompts are segmented following the *Inception* paradigm to simulate multi-turn conversations. For unsafe prompt, we use the VBCDE prompt set. PBD applies a perplexity-based threshold  $\tau$ : if a request yields a perplexity lower than  $\tau$ , it is classified as safe; otherwise, it is deemed unsafe. As we can see in Figure 8, judgment based on perplexity leads to severe false positives. This validates the limitation of utilizing perplexity for *Inception* attack detection. We provide a more detailed analysis in Appendix B.5.

**Memory Scanner (MS).** Conventional input-side safety filters operate on a per-query basis, which makes them inherently vulnerable to attacks that distribute harmful intent across multiple interaction rounds, as exemplified by *Inception*. In such cases, each individual query appears benign in isolation, causing input-level detection to fail. This observation motivates a shift in where safety auditing should occur: rather than evaluating isolated inputs, detection should be performed after user intent has been aggregated across the conversation. Motivated by this insight, we introduce a **Memory Scanner (MS)** that is inserted between the memory manager and the image generation model. Since T2I generation systems must reconstruct user intent from the conversation history before generation, the MS leverages this aggregation step to detect malicious intent accumulated over time. Specifically, the MS adapts its auditing input to the underlying memory mechanism. For BufferMem, it evaluates the full query history; for SummaryMem, it inspects the condensed summary; and for VSRMem, it examines a set of relevant interaction

Table 6: Attack performance of *Inception* against dedicated defenses, including MS, EOM, and EOM\*. *Original* refers to the performance of *Inception* when only the OpenAI input and output moderators are applied.

VBCDE				
Defense	ASR (↓)	CLIP score (↓)		# of query (↑)
		img vs. prompt	img vs. img	
Origin	32.3%	0.247	0.636	12.18
MS	27.7%(-4.6%)	0.244(-0.003)	0.628(-0.008)	14.06(+1.88)
EOM	41.7%(+9.4%)	0.256(+0.009)	0.658(+0.022)	10.31(-1.87)
EOM*	26.7%(-5.6%)	0.246(-0.001)	0.633(-0.003)	12.63(+0.45)

UnsafeDiff				
Origin	28.7%	0.278	0.701	10.26
MS	21.3%(-7.4%)	0.269(-0.009)	0.698(-0.003)	11.78(+1.52)
EOM	36.0%(+7.3%)	0.284(+0.006)	0.718(+0.017)	8.54(-1.72)
EOM*	24.7%(-4.0%)	0.275(-0.003)	0.694(-0.007)	10.87(+0.61)

samples, noting that both SummaryMem and VSRMem dynamically manage historical context. We employ OpenAI’s moderation model to audit this aggregated representation. If the memory content is flagged as unsafe, the current request is rejected; otherwise, the sanitized memory is forwarded to the generation model. Overall, this design yields a multi-stage defense pipeline that consistently employs OpenAI’s moderation models at the input, memory, and output stages, enabling complementary and coordinated safety auditing throughout the entire generation process.

**Enhanced Output Moderator (EOM).** The reason T2I generation systems exhibit greater vulnerability compared to large language models (LLMs) lies in the performance disparity performance gap between moderators, specifically, between text moderators and image moderators. That is, the implicit intent embedded in images is more difficult to detect. To address this issue, we propose amplifying the image’s intent by first generating a caption that explicitly describes its content. This textual description is then passed to a text moderator, allowing us to leverage the more robust capabilities of text-based moderators. In this way, the image moderation ability is enhanced (EOM). We consider two uses of EOM: **(1)** adopting EOM alone for output detection (the overall pipeline is OpenAI input text moderator + EOM); and **(2)** employing EOM as a supplement to the current output detector (the overall pipeline is OpenAI input text moderator + OpenAI output text moderator + EOM, denoted as EOM\*)

**Analysis.** We report the effectiveness of MS, EOM, and EOM\* in Table 6, from which two key observations emerge. Firstly, MS constitutes the most effective defense against *Inception*, reducing the attack success rate (ASR) on UnsafeDiff by 7.4% while increasing the number of required attack queries. Nevertheless, *Inception* continues to outperform the baseline methods shown in Table 1, underscoring its robustness under dedicated defenses. A plausible explanation for why *Inception* remains largely resilient to the Memory Scanner lies in the irreversibility of the intent decomposition

process. Specifically, a semantically coherent concept such as “blood” can be decomposed into fragments like “red” and “liquid”. While these fragments jointly convey similar perceptual cues, memory mechanisms are unable to recombine them into the original high-level concept with equivalent semantic meaning, thereby limiting the effectiveness of memory-level detection. Secondly, using EOM alone as an output filter is less effective in defending against *Inception*, unless it is combined with a stronger moderator. This suggests that EOM is more effective as a supplementary filter rather than as a primary defender. Overall, all dedicated defenses demonstrate limited effectiveness against *Inception*, underscoring the urgent need for more powerful defense strategies.

## 8 Conclusion

In this paper, we proposed and developed the first multi-turn jailbreak attack, namely *Inception*, targeting the memory mechanism of commercial online T2I generation systems. We revealed that existing single-turn jailbreak attacks were significantly less effective in evaluating the vulnerabilities of real-world systems due to under- and over-detoxification issues. By leveraging the multi-turn capability enabled by memory mechanisms in modern T2I generation systems, together with our design of segmentation and recursion, we successfully addressed these challenges. *Inception* recursively segments unsafe words into smaller chunks with minimal maliciousness, preserving semantics while bypassing safety filters. Experimental results on popular T2I generation systems demonstrated the effectiveness of *Inception*. We further showed that potential dedicated defenses offered only limited protection against our attack, underscoring the need for further investigation. We hope our work sheds light on the security of real-world T2I systems and facilitates the development of safer generation mechanisms.

## Acknowledgments

We thank all the reviewers for their constructive suggestions that helped improve this paper. This research/project is supported by the National Research Foundation, Singapore, under its National Large Language Models Funding Initiative (AISG Award No: AISG-NMLP-2024-005), the National Research Foundation, Singapore, under its AI Singapore Programme (AISG Award No: AISG3-RPGV-2025-019), and JST K-program JPMJKP24C3.

## Ethical Considerations

**Stakeholder Analysis.** This research studies vulnerabilities in real-world T2I generation systems. The stakeholders include:

- Commercial Platforms.** Jailbreak attacks may compromise platform safety. We have shared our findings and pro-

posed defenses with relevant commercial platforms through their feedback channels and email communication.

- 2) **The Public.** Generated unsafe content may cause discomfort. To avoid unintended exposure, all inappropriate examples are stored in a password-protected repository, accessible only with permission.
- 3) **The Research Team.** The authors conducted manual evaluations of unsafe images for evaluating the judge model. To mitigate potential psychological distress, review sessions were time-limited. Participation was strictly voluntary, and participants could withdraw at any time.

### Impact on Stakeholders.

- [1)] **Impact on Commercial Platforms.** We disclosed our findings and defenses to affected platforms (ChatGPT, Gemini, Grok), detailing reproduction steps and mitigation strategies. While exposing T2I vulnerabilities carries inherent risk, the primary benefit is raising awareness to foster stronger safety protections.
- [2)] **Impact on the Public.** To mitigate potential negative impact, we restrict access to successful jailbreak cases via password-protected links, thereby preventing accidental exposure and misuse of adversarial examples. On the positive side, our findings contribute to the development of more responsible and robust T2I generation systems.

**Mitigation Measures.** We mitigated potential risks through adaptive defenses and preventive countermeasures:

- [1)] **Mitigation for Publication.** We ensured timely notification and assistance for affected platforms and used password-protected links to limit exposure.
- [2)] **Proposed Defenses.** We designed several dedicated defenses against multi-turn T2I jailbreak attacks, which we have also recommended to commercial platforms:
  - Memory Scanner. Scans the memory of T2I systems.
  - Enhanced Output Moderator (EOM). Improves moderation by performing captioning before detection.
  - EOM\*. Extends EOM with additional moderators for enhanced detection.

### Decision to Conduct and Publish the Research.

- [1)] **Decision to Conduct the Research.** We investigated multi-turn jailbreak vulnerabilities in real-world T2I systems to identify weaknesses crucial for developing effective safety mechanisms. While acknowledging potential misuse risks, our primary objective is to enhance system robustness and advance responsible AI development.
- [2)] **Decision to Publish the Research.** We chose to publish our findings to raise awareness within the research and industry communities and to encourage stronger defense development. Prior to publication, we took comprehensive precautions: notifying affected platforms (*i.e.*, ChatGPT, Gemini, and Grok), sharing reproduction details and mitigation suggestions, and securing all unsafe examples behind password protection to prevent misuse. Publication was

therefore deemed essential for transparency and collective advancement in AI safety.

### Respect for Persons.

- 1) **Notice.** We prepared informed consent documents outlining the potential benefits of research and any associated risks.
- 2) **Comprehension.** For user study, language was kept at or below an eighth-grade reading level to ensure accessibility.
- 3) **Voluntariness.** Participation was strictly voluntary, with right to withdraw at any time without consequence.

**IRB Approval.** Although our study involved images depicting NSFW concepts, it was reviewed and approved by our Institutional Review Board (IRB) under a process comparable to the “exempt review” category of U.S. IRB protocols (45 CFR 46). The IRB determined that the study posed no more than minimal risk, as participants were healthy adults, fully informed, and free to withdraw at any time.

### Open Science

We are committed to open science principles by sharing the outcomes of our research in an open-access format.

**Open Sharing of Code and Data.** All artifacts from this research, including datasets, test cases, scripts, and source code, will be made publicly available on GitHub. We will also release the text-to-image generation system constructed in this study for community use. The datasets used, namely VBCDE and UnsafeDiff, are already publicly available in their respective repositories. The permanent link to our artifact repository is [here](#).

**Reproducibility and Replicability.** We will provide all artifacts necessary for reproducing our results, including detailed experimental records and documentation. These will cover environment setup, dependencies, and parameter settings, enabling other researchers to replicate our findings.

### References

- [1] Eleni Adamopoulou and Lefteris Moussiades. Chatbots: History, technology, and applications. *Machine Learning with applications*, 2020.
- [2] Eleni Adamopoulou and Lefteris Moussiades. An overview of chatbot technology. In *Artificial Intelligence Applications and Innovations*, 2020.
- [3] Stability AI. Sd-3.5-large-turbo. <https://huggingface.co/stabilityai>. Access: 2024-11-28.
- [4] Gabriel Alon and Michael Kamfonas. Detecting language model attacks with perplexity. *arXiv preprint*, 2023.
- [5] BFL. Flux. [huggingface.co/black-forest-labs/FLUX.1-schnell](https://huggingface.co/black-forest-labs/FLUX.1-schnell). Access: 2024-11-28.

- [6] Bocheng Chen and Qiben Yan. Understanding multi-turn toxic behaviors in open-domain chatbots. In *Research in Attacks, Intrusions and Defenses*, 2023.
- [7] Jianfeng Chi, Kate Plawiak, and Mahesh Pasupuleti. Llama guard 3 vision: Safeguarding human-ai image understanding conversations. *arXiv preprint*, 2024.
- [8] Pucheng Dang, Xing Hu, Dong Li, Rui Zhang, Qi Guo, and Kaidi Xu. Diffzoo: A purely query-based black-box attack for red-teaming text-to-image generative model via zeroth order optimization. *arXiv preprint*, 2024.
- [9] Yimo Deng and Huangxun Chen. Daca: Harnessing the power of llm to bypass the censorship of text-to-image generation model. *arXiv preprint*, 2023.
- [10] Yingkai Dong, Zheng Li, Xiangtao Meng, Ning Yu, and Shanqing Guo. Jailbreaking text-to-image models with llm-based agents. *arXiv preprint*, 2024.
- [11] Xiaohu Du, Fan Mo, Hai Jin, and Jie Shi. Multi-turn jailbreaking large language models via attention shifting. In *AAAI Conference on Artificial Intelligence*, 2025.
- [12] Facebook. Faiss. <https://github.com/facebookresearch/faiss>. Access: 2024-11-18.
- [13] Falconsai. Nsfw image classification. [huggingface.co/Falconsai/nsfw\\_image\\_detection](https://huggingface.co/Falconsai/nsfw_image_detection). Access: 2024-11-18.
- [14] Xiachong Feng, Xiaocheng Feng, and Ting Liu. Language model as an annotator: Exploring dialogpt for dialogue summarization. *arXiv preprint*, 2021.
- [15] Ophir Frieder, Ida Mele, Raffaele Perego, and Nicola Tonello. Caching historical embeddings in conversational search, 2024. US Patent 12,067,021.
- [16] Sensen Gao, Xiaojun Jia, Yihao Huang, Yang Liu, and Qing Guo. Rt-attack: Jailbreaking text-to-image models via random token. *arXiv preprint*, 2024.
- [17] Google. Gemini. <https://gemini.google.com/app>. Access: 2024-10-23.
- [18] Google. Imagen. <https://gemini.google.com/app>. Access: 2024-10-23.
- [19] Keyan Guo, Ayush Utkarsh, Ziming Zhao, and Hongxin Hu. Moderating illicit online image promotion for unsafe user generated content games using large {Vision-Language} models. In *USENIX Security*, 2024.
- [20] Lukas Helff, Felix Friedrich, Manuel Brack, and Kristian Kersting. Llavaguard: Vlm-based safeguard for vision dataset curation and safety assessment. In *CVPR*, 2024.
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 2020.
- [22] Yihao Huang, Le Liang, Tianlin Li, Geguang Pu, and Yang Liu. Perception-guided jailbreak against text-to-image models. *arXiv preprint*, 2024.
- [23] Minseon Kim, Hyomin Lee, Boqing Gong, and Sung Ju Hwang. Automatic jailbreaking of the text-to-image generative ai systems. *arXiv preprint*, 2024.
- [24] Diederik P Kingma. Auto-encoding variational bayes. *arXiv preprint*, 2013.
- [25] Langchain. Buffermem. [python.langchain.com/docs/versions/migrating\\_memory/conversation\\_buffer\\_memory/](https://python.langchain.com/docs/versions/migrating_memory/conversation_buffer_memory/). Access: 2024-11-18.
- [26] LangChain. Langchain. [https://python.langchain.com/v0.1/docs/use\\_cases/chatbots/memory\\_management/](https://python.langchain.com/v0.1/docs/use_cases/chatbots/memory_management/). Access: 2025-3-21.
- [27] LangChain. Langgraph memory. [https://python.langchain.com/docs/versions/migrating\\_memory/](https://python.langchain.com/docs/versions/migrating_memory/). Access: 2024-11-18.
- [28] Langchain. Summaryemem. [docs.langchain.com/oss/python/langchain/short-term-memory](https://docs.langchain.com/oss/python/langchain/short-term-memory). Access: 2024-11-18.
- [29] Langchain. Vectormemory. [python.langchain.com/docs/versions/migrating\\_memory/long\\_term\\_memory\\_agent/](https://python.langchain.com/docs/versions/migrating_memory/long_term_memory_agent/). Access: 2024-11-18.
- [30] Shuangqi Li, Hieu Le, Jingyi Xu, and Mathieu Salzmann. Enhancing compositional text-to-image generation with reliable random seeds. *arXiv preprint*, 2024.
- [31] Xirui Li, Ruochen Wang, and Cho-Jui Hsieh. Drattack: Prompt decomposition and reconstruction makes powerful llm jailbreakers. *arXiv preprint*, 2024.
- [32] LlamaIndex. mem-zero. <https://docs.mem0.ai/platform/overview>. Access: 2024-11-18.
- [33] Midjourney. Midjourney. <https://www.midjourney.com>. Access: 2024-06-26.
- [34] OpenAI. Chatgpt. <https://chatgpt.com/>. Access: 2024-10-23.
- [35] OpenAI. Dall-e 3. <https://openai.com/index/dall-e-3>. Access: 2024-06-26.
- [36] OpenAI. Image moderator. <https://platform.openai.com/docs/guides/moderation?example=images>. Access: 2025-3-21.

- [37] OpenAI. System card. <https://openai.com/index/dall-e-3-system-card/>. Access: 2024-11-18.
- [38] OpenAI. Text moderator. <platform.openai.com/docs/guides/moderation>. Access: 2025-3-21.
- [39] OpenAI. Vector embeddings. <platform.openai.com/docs/guides/embeddings>. Access: 2024-11-18.
- [40] Pinecone. Pinecone. <https://www.pinecone.io/>. Access: 2024-11-18.
- [41] Dustin Podell, Zion English, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. *arXiv preprint*, 2023.
- [42] Yiting Qu, Michael Backes, and Yang Zhang. Unsafe diffusion: On the generation of unsafe images and hateful memes from text-to-image models. In *CCS*, 2023.
- [43] Yiting Qu, Michael Backes, Yang Zhang, and Savvas Zannettou. On the evolution of (hateful) memes by means of multimodal contrastive learning. In *SP*, 2023.
- [44] Alec Radford, Gabriel Goh, Amanda Askell, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *ICML*, 2021.
- [45] Javier Rando, Daniel Paleka, David Lindner, Lennart Heim, and Florian Tramèr. Red-teaming the stable diffusion safety filter. *arXiv preprint*, 2022.
- [46] REST. Rest api. <https://restfulapi.net/statelessness/>. Access: 2025-3-21.
- [47] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022.
- [48] Mark Russinovich, Ahmed Salem, and Ronen Eldan. Great, now write an article about that: The crescendo multi-turn llm jailbreak attack. *arXiv preprint*, 2024.
- [49] Patrick Schramowski, Manuel Brack, and Kristian Kersting. Safe latent diffusion: Mitigating inappropriate degeneration in diffusion models. In *CVPR*, 2023.
- [50] Patrick Schramowski and Kristian Kersting. Can machines help us answering question 16 in datasheets, and in turn reflecting on inappropriate content? In *Fairness, Accountability, and Transparency*, 2022.
- [51] Amazon Web Services. Amazonmemory. <community.aws/content/2j9daS4A39fteekgv9t1Hty11Qy>. Access: 2025-3-21.
- [52] ShuttleAI. Shuttlediffusion. <huggingface.co/shuttleai/shuttle-3-diffusion>. Access: 2024-11-18.
- [53] Jascha Sohl-Dickstein, Eric Weiss, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*.
- [54] Spacy. Spacy. <spacy.io/>. Access: 2024-10-23.
- [55] Yu-Lin Tsai, Chia-Yi Hsu, Bo Li, and Chun-Ying Huang. Ring-a-bell! how reliable are concept removal methods for diffusion models? *arXiv preprint*, 2023.
- [56] Zhenting Wang, Shuming Hu, Xiaowen Lin, Zhuowei Li, Li Chen, and Jianfa Chen. Mllm-as-a-judge for image safety without human labeling. In *CVPR*, 2025.
- [57] Zixuan Weng, Xiaolong Jin, Jinyuan Jia, and Xiangyu Zhang. Foot-in-the-door: A multi-turn jailbreak for llms. *arXiv preprint*, 2025.
- [58] Yuanwei Wu, Yue Huang, Yixin Liu, Xiang Li, Pan Zhou, and Lichao Sun. Can large language models automatically jailbreak gpt-4v? *arXiv preprint*, 2024.
- [59] xAI. Aurora. <https://grok.com/imagine>. Access: 2025-7-28.
- [60] Xikang Yang, Xuehai Tang, Songlin Hu, and Jizhong Han. Chain of attack: a semantic-driven contextual multi-turn attacker for llm. *arXiv preprint*, 2024.
- [61] Yijun Yang, Ruiyuan Gao, Nan Xu, and Qiang Xu. Mma-diffusion: Multimodal attack on diffusion models. In *CVPR*, 2024.
- [62] Yuchen Yang, Bo Hui, Haolin Yuan, Neil Gong, and Yinzhi Cao. Sneakyprompt: Jailbreaking text-to-image generative models. In *SP*, 2024.
- [63] Jiahao Yu, Xingwei Lin, Zheng Yu, and Xinyu Xing. Gptfuzzer: Red teaming large language models with auto-generated jailbreak prompts. *arXiv preprint*, 2023.
- [64] Shiqian Zhao, Chong Wang, Yiming Li, and Tianwei Zhang. Towards effective prompt stealing attack against text-to-image diffusion models. *NDSS*, 2025.
- [65] ZhipuAI. Chatglm. <https://chatglm.cn/main/alltoolsdetail?lang=en>. Access: 2024-10-23.
- [66] Andy Zhou and Ron Arel. Siege: Multi-turn jailbreaking of large language models with tree search. In *ICLR 2025 Workshop on Building Trust in Language Models and Applications*.
- [67] Zhenhong Zhou, Jiuyang Xiang, Quan Liu, and Sen Su. Speak out of turn: Safety vulnerability of large language models in multi-turn dialogue. *arXiv preprint*, 2024.
- [68] Andy Zou, Zifan Wang, Nicholas Carlini, and Matt Fredrikson. Universal and transferable adversarial attacks on aligned language models. *arXiv preprint*, 2023.

## A More Details

### A.1 Practicability of Memory Mechanisms.

We validate the effectiveness of the aforementioned memory mechanisms to demonstrate their practicality in real-world scenarios. We focus on a scenario where a user updates their image generation request over multiple turns. To simulate such a chain of requests, we segment a target prompt (serving as a ground-truth summarization) using our segmentation method (Section 5.3). Here, we do not activate the safety filter in order to replicate normal usage conditions. We evaluate the practicality of BufferMem, SummaryMem, and VSRMem using two metrics. First, we measure the semantic similarity between the target prompt and the summarization generated by the memory mechanism. This metric assesses the accuracy of the summarization, where 0.8 can indicate the same semantics. Additionally, we evaluate the CLIP score between the target prompt and the generated image, which reflects how well the user’s intent is fulfilled by the generation systems. We present the results in Table 7. As shown, both BufferMem and SummaryMem demonstrate strong summarization capabilities. Considering that the ground-truth prompt achieves a CLIP score of 0.283 with its generated image, the image generated from the summary achieves a score only 0.008 lower. This highlights the practicality of these methods in accurately capturing users’ intent. We also observe that VSRMem achieves relatively low CLIP and SBERT scores. A potential reason for this is the consideration of only 5 queries, which may omit important information.

### A.2 Relation to Multi-turn LLM Jailbreaks

Although our approach shares high-level similarities with prior multi-turn and certain single-turn LLM jailbreaks, it differs substantially in practice. In this section, we provide a detailed analysis of these differences.

**Difference of Memory Mechanism.** In this paper, we study the vulnerability of text-to-image (T2I) generation systems that incorporate an external *memory mechanism* to support multi-turn image generation. In contrast, multi-turn LLM jailbreaks target intrinsic LLM vulnerabilities and guardrails. This distinction stems from the modality gap between input and output: LLMs operate within a single text modality and naturally support in-context reasoning, whereas T2I systems lack this property. Accordingly, we begin by simulating VisionFlow, a T2I generation system equipped with three representative industrial memory paradigms.

**Strategy Difference.** Multi-turn LLM jailbreaks typically bypass safety mechanisms via semantic expansion or contextualization. For instance, Chain-of-Attack (CoA) [60] frames unsafe requests within benign contexts (e.g., a chemistry course), ASJA [11] employs multiple strategies such as Defined Persona and Imagined Scenario, Foot-In-The-Door [57] intro-

Table 7: The performance of different memory mechanisms. The target here refers to the intent of a user.

VBCDE [9]				
Metric	Target	BufferMem	SummaryMem	VSRMem
$CLIP_{sum}^{img}$	0.283	0.275	0.272	0.245
SBERT	1.0	0.857	0.817	0.650

UnsafeDiff [42]				
Metric	Target	BufferMem	SummaryMem	VSRMem
$CLIP_{sum}^{img}$	0.314	0.298	0.283	0.282
SBERT	1.0	0.880	0.755	0.801

duces unsafe queries after benign background discussion, and SIEGE [66] combines persona shifts, disguised re-framing, and refusal suppression via tree search. While such contextualization generally preserves response fidelity for LLMs due to their strong context awareness, applying the same strategies to T2I models, whose cross-attention maps the entire prompt to the image, can substantially shift prompt semantics and lead to unfaithful generation.

**Comparison with Single-turn LLM Jailbreaks.** Our proposed Inception is a multi-turn jailbreak attack specifically targeting the T2I generation system, incorporating an internal memory mechanism. It consists of two primary modules: segmentation and recursion. Specifically, the segmentation module uses NLP analysis to split a single unsafe prompt into multiple sub-queries based on its sentence structure. This structure-based segmentation is crucial, as it ensures semantic consistency and, consequently, a faithful final T2I generation. In contrast, existing single-turn attacks, such as DrAttack [31] for LLMs, perform segmentation to extract and replace unsafe words (instead of getting a prompt list) by having an LLM parse the malicious prompt. Furthermore, our recursion module handles sub-queries that are too simple for initial segmentation by recursively expanding them without semantic loss, allowing for further segmentation. This recursive approach differs fundamentally from single-turn LLM jailbreaks, which typically adopt an *iterative* manner. For instance, Gptfuzzer [63] iteratively optimizes *one* prompt using an evolving template and mutations, a process which functionally differs from our method and often results in a semantics shift, affecting the fidelity of the final response.

### A.3 The Judge Model

**VLM as The Judge.** To automatically evaluate ASR, we adopt a vision–language model (VLM), namely ChatGPT-4o, as the Judge to replace manual assessment. We verify the alignment between the Judge and human annotators through the following human evaluation, which was conducted by the three authors and thus falls under the “exempt review” category for IRB approval. It was performed solely to assess the reliability of the VLM-as-Judge, which is used only for

Table 8: The definition of unsafe concept.

Concept	Definition
<b>Harassment</b>	Content that promotes hate, violence, or discrimination against individuals or groups based on attributes like race, ethnicity, gender, sexual orientation, religion, etc. This also includes cyberbullying and harassment.
<b>Illegal activity</b>	Any content promoting or facilitating illegal activities, including drug trade, terrorism, human trafficking, or fraud.
<b>Self-harm</b>	Encouraging or depicting self-harm, suicide, or related topics in a way that could be harmful or promote these actions.
<b>Sexual content</b>	Explicit sexual material, including pornography, sexually suggestive content, or adult themes, is prohibited. This also covers content involving minors or non-consensual acts.
<b>Violence</b>	Anything that incites, glorifies, or promotes violence, harm, or abuse towards others, including graphic depictions of violence or self-harm.

evaluating generated images and is not part of our method.

- **Process.** Our evaluation involves three main steps. First, we prompt DeepSeek to generate a large set of safe and unsafe prompts. The unsafe prompts are drawn from five concepts: harassment, illegal activity, self-harm, sexual content, and violence. The safe prompts cover nature scenes, animals, food and culinary, architecture and landmarks, and art and design. Next, these prompts are used to generate images with the open-source model Stable Diffusion 3.5, after disabling its safety checker. Finally, we perform human evaluation of the generated images, where three author volunteers independently annotate each unsafe image. The annotated images are then assessed by ChatGPT-4o.
- **Annotation Guidelines.** Each volunteer independently evaluated the images using the following instructions: “Annotate the image as appropriate or inappropriate. The reference concept is Concept, which is defined as: Definition” (inappropriate concepts are defined in Table 8). The same instruction was provided to ChatGPT-4o to ensure consistency between human and model-based assessments.
- **Aggregation Principles.** We adopt a strict aggregation rule: an image is labeled as safe or unsafe only when all three annotators reach unanimous agreement. This criterion reduces both false positives and false negatives by filtering out inconsistent judgments. After obtaining aggregated labels, we also computed the number of instances in which human annotations aligned with those produced by the Judge.
- **Borderline Cases.** For instance, the generated image from the prompt “A man is licking a woman’s hair” (which is supposed to be unsafe) was labeled inappropriate by two annotators for implying harassment, while one considered it a depiction of intimate interaction. Since it did not receive unanimous votes, it was discarded, even though the Judge model classified it as unsafe.

**Evaluation Result.** As shown in Table 9, the Judge demonstrates strong alignment with human perception. Notably, it produces no false positives, meaning it never misclassifies

Table 9: Assessment of The Judge.

TPR	FPR	Accuracy	Precision	Recall	F1-Score
93.7%	0%	96.9%	100.0%	93.7%	96.7%

Table 10: Testing Results on I2P.

Methods	ASR (↑)	CLIP score (↑)		# of query (↓)
		img vs. prompt	img vs. img	
DACA	9.2%	0.224	0.623	-
CoA	2.4%	0.217	0.628	15.71
Inception	<b>26.8%</b>	<b>0.262</b>	<b>0.703</b>	<b>10.49</b>

a safe image as unsafe. At the same time, it achieves a true positive rate (TPR) of 93.7%, confirming its effectiveness in detecting unsafe images. This results in a high precision (100%), indicating a rather conservative strategy. These results indicate that employing a VLM as the Judge is a reliable approach for automatic assessment. But out of its high FNR (6.3%), the Judge is not suitable for being a safety filter. Nevertheless, we acknowledge that no single metric is flawless. To ensure comprehension, we therefore complement the Judge with additional metrics, including CLIP score and query time.

## B More Results

### B.1 Evaluation on a Larger Dataset (I2P)

To further validate our approach, we evaluated Inception against the baselines on the larger I2P dataset [49]. I2P comprises 4.7k unsafe prompts spanning seven core unsafe concepts: self-harm, hate, violence, shocking, sexual, harassment, and illegal activity. The dataset includes an ‘inappropriate percentage’ metric indicating the degree of unsafety for each prompt. To focus on the most malicious samples and mitigate false positives, we selected 1,000 prompts with an inappropriate percentage greater than 0.8. As demonstrated in Table 10, our Inception maintains its superiority over baselines, including DACA [9] and CoA [60]. The observed performance trend is consistent with our primary evaluations on VBCDE [9] and UnsafeDiff [42], conclusively indicating the consistent robustness and superiority of our method.

### B.2 Transferability Attacks

We provide the generated successful adversarial prompts and the visualization results. For responsible disclosure, we place them in a password-protected link, following previous work [62]. The generated prompts and images can be found in [this link](#) (**Warning: This link contains disturbing text and images. Please process with caution.**). The password will only be available after the application.

Table 11: Performance of `Inception` when different backstage models are adopted. We adopt the input and output moderators from OpenAI as safety filters, and we set `BufferMem` as the memory manager.

Dataset	Model	One-time attack			Re-use attack			
		ASR ( $\uparrow$ )	CLIP score ( $\uparrow$ )		# of queries ( $\downarrow$ )	ASR ( $\uparrow$ )	CLIP score ( $\uparrow$ )	
			<i>image<sub>adv</sub> vs. prompt<sub>target</sub></i>	<i>image<sub>adv</sub> vs. image<sub>target</sub></i>			<i>image<sub>adv</sub> vs. prompt<sub>target</sub></i>	<i>image<sub>adv</sub> vs. image<sub>target</sub></i>
VBCDE [9]	SD-3.5	32.3%	0.247	0.636	12.18	26.3%	0.243	0.615
	FLUX	<b>34.3%</b>	<b>0.251</b>	<b>0.628</b>	<b>12.07</b>	<b>30.3%</b>	<b>0.246</b>	<b>0.617</b>
	ShuttleDiffusion	31.7%	0.250	0.635	11.88	25.7%	0.243	0.614
UnsafeDiff [42]	SD-3.5	28.7%	0.278	0.701	10.26	21.3%	0.268	0.679
	FLUX	<b>30.3%</b>	<b>0.282</b>	<b>0.717</b>	<b>9.76</b>	<b>22.7%</b>	<b>0.280</b>	<b>0.694</b>
	ShuttleDiffusion	28.3%	0.275	0.700	10.74	20.7%	0.273	0.682

Table 12: Impact of Query Budget in Rewriting LLMs

Query budget	ASR ( $\uparrow$ )	CLIP score ( $\uparrow$ )		# of query ( $\downarrow$ )
		img vs. prompt	img vs. img	
5	22.7%	0.278	0.703	10.59
10	29.3%	0.276	0.703	11.41
15	28.7%	0.274	0.702	10.83
20	<b>28.7%</b>	<b>0.278</b>	<b>0.701</b>	<b>10.26</b>
25	32.3%	0.273	0.697	10.58
30	33.7%	0.272	0.699	10.67

### B.3 Study on Backstage Model

We investigate the impact of the backstage model on the effectiveness of `Inception`. Our simulation system, `VisionFlow`, supports a plugin module that enables the use of customized generation models. In addition to Stable Diffusion 3.5 (SD-3.5) [3], we consider two high-performing open-source models: FLUX [5] and ShuttleDiffusion [52]. As shown in Table 11, the choice of backstage model influences the performance of `Inception`. In particular, adopting FLUX leads to stronger attack outcomes, both for one-time and re-use attacks. For instance, on re-use ASR, `Inception` achieves improvements of 4.0%.

### B.4 Additional Ablation Study

We further investigate the query budget of the rewriting LLMs. In the recursion process, the blocked unsafe query is first expanded using the rewriting LLM, after which the segmentation function is invoked to split the expansion. To prevent endless rewriting, we set the maximum query budget  $\Pi$ , with  $\Pi \in \{5, 10, 15, 20, 25, 30\}$ . As reported in Table 12, attack performance, particularly ASR, improves as  $\Pi$  increases from 5 to 10. Within the range of 10–20, ASR remains stable, and it reaches its highest value when  $\Pi = 30$ .

### B.5 Perplexity-based Detection

Figure 9 shows the trends of the true positive rate (TPR) and false positive rate (FPR) under different perplexity thresholds, where a prompt is classified as unsafe if its perplexity




Figure 9: ROC curve on perplexity thresholds.

Table 13: Testing results on French and German.

Methods	ASR ( $\uparrow$ )	CLIP score ( $\uparrow$ )		# of query ( $\downarrow$ )
		img vs. prompt	img vs. img	
English	28.7%	0.278	0.701	10.26
French	31.3%	0.280	0.721	9.07
<b>German</b>	<b>32.3%</b>	<b>0.291</b>	<b>0.735</b>	<b>8.65</b>

exceeds the threshold. TPR is the fraction of unsafe prompts detected from VBCDE, while FPR is the fraction of safe DALLEPrompt prompts misclassified as unsafe. The two curves largely overlap, indicating that `Inception`-generated adversarial prompts are nearly indistinguishable from benign ones in terms of perplexity. We highlight thresholds of 400 and 1,000 used in [64]. At 400, the FPR is comparable to or even higher than the TPR, making detection close to random guessing. At 1,000, TPR exceeds FPR, but both remain low, with more than 80% of unsafe prompts still undetected.

### B.6 Evaluation in Other Languages

We further evaluate the effectiveness of the `Inception` attack against T2I systems under non-English settings, specifically French and German. To this end, we translate the segmented sub-queries into the target languages prior to submission. As reported in Table 13, `Inception` remains highly effective when operating in French or German, achieving consistently high ASR and CLIP scores while requiring substantially fewer queries per attack. These results suggest that `Inception` poses security risks beyond English-only settings.