

Gradient-based Automatic Attack of Text-to-Image Model

CHENGBIN DU

Master of Philosophy



THE UNIVERSITY OF
SYDNEY

Supervisor: Dr. Chang Xu
Associate Supervisor: Dr. Siqi Ma

A thesis submitted in fulfilment of
the requirements for the degree of
Master of Philosophy

School of Computer Science
Faculty of Engineering
The University of Sydney
Australia

19 January 2024

Gradient-based Automatic Attack of Text-to-Image Model

Chengbin Du (Email: chdu5632@uni.sydney.edu.au)

Supervisor: Dr. Chang Xu

School of Computer Science

Faculty of Engineering

The University of Sydney

Copyright in Relation to This Thesis

© Copyright 2019 by Chengbin Du. All rights reserved.

Statement of Original Authorship

This is to certify that to the best of my knowledge, the content within this thesis is a result of my individual effort. This thesis has not been submitted for any other academic degree or intention. I certify that the intellectual content of this thesis is the product of my own work and that all the assistance received in preparing this thesis and sources have been acknowledged.

Signature: Chengbin Du

Date: 19/1/2024

Authorship attribution statement

I, Chengbin Du, designed the study and wrote the thesis drafts. In addition, the core content of this thesis is derived from this paper [1] that has been accepted in the Neural Information Processing Systems (NeurIPS) – 37th Anniversary Conference. I generated the dataset used in this paper, designed and implemented the algorithm, and conducted further empirical analysis. This analysis revealed four attack patterns, identified based on the failure cases of the stable diffusion model.

Student Name: Chengbin Du

Signature: Chengbin Du

Date: 19/1/2024

As the supervisor for the candidature upon which this thesis is based, I can confirm that the authorship attribution statements above are correct.

Supervisor Name: Chang Xu

Signature: Chang Xu

Date: 19/1/2024

[1] Du, C., Li, Y., Qiu, Z. and Xu, C., 2023. Stable Diffusion is Unstable. Advances in Neural Information Processing Systems.

Abstract

Recently, text-to-image models have been thriving. Despite their powerful generative capacity, our research has uncovered a lack of robustness in this generation process. Specifically, the introduction of small perturbations to the text prompts can result in the blending of primary subjects with other categories or their complete disappearance in the generated images. In this thesis, we propose **Auto-attack on Text-to-image Models (ATM)**, a gradient-based approach, to effectively and efficiently generate such perturbations. By learning a Gumbel Softmax distribution, we can make the discrete process of word replacement or extension continuous, thus ensuring the differentiability of the perturbation generation. Once the distribution is learned, ATM can sample multiple attack samples simultaneously. These attack samples can prevent the generative model from generating the desired subjects without tampering with the category keywords in the prompt. ATM has achieved a 91.1% success rate in short-text attacks and an 81.2% success rate in long-text attacks. Further empirical analysis revealed four attack patterns based on the following: 1) the variability in generation speed, 2) the similarity of coarse-grained characteristics, 3) the polysemy of words, and 4) the positioning of words.

Contents

Authorship attribution statement	iii
Abstract	iv
Contents	v
List of Figures	vii
Chapter 1 Introduction	1
Chapter 2 Literature review	4
2.1 Diffusion Model	4
2.1.1 Denoising Diffusion Probabilistic Models	5
2.1.2 Stable diffusion	6
2.2 Vulnerabilities in Text-to-image Models.	7
2.3 Adversarial Attack	9
2.3.1 Image Adversarial Attack	9
2.3.2 Text Adversarial Attack	10
Chapter 3 Methods	12
3.1 Methodology.	12
3.1.1 Differentiable Text Prompt Modification	13
3.1.2 Constraints on Fluency and Similarity	14
3.1.3 Generation of Attack Prompts.	16
Chapter 4 Results	17
4.1 Experimental Setting.	17
4.2 Main Results	18
4.3 Different Search Steps	20

4.4	The Impact of Constraints	20
4.5	Different Samplers	22
4.6	Cases of Short/Long-Prompt Attacks	22
4.7	Black-box attack	22
Chapter 5	Empirical analysis	27
5.1	Vulnerabilities of Stable Diffusion Model	27
5.1.1	Variability in Generation Speed	27
5.1.2	Similarity of Coarse-grained Characteristics	31
5.1.3	Polysemy of Words	33
5.1.4	Positioning of Words	34
Chapter 6	Conclusion	37
	Acknowledgements	38
	Bibliography	40

List of Figures

1.1	The overall pipeline of our attack method. The selection of words is relaxed to be differentiable by using a Gumbel Softmax with temperature τ . When $\tau \rightarrow 0$, the Gumbel Softmax outputs exhibit proximity to one-hot vectors while retaining differentiability. After an image is generated, a CLIP [18] classifier and a margin loss are employed to optimize ω aiming to generate images that cannot be correctly classified by CLIP.	2
4.1	To the left of the arrow is the clean long text prompt (highlighted by green) and its corresponding image, to the right of the arrow is the generated attack prompt (highlighted by red) and its corresponding image.	23
4.2	To the left of the arrow is the clean short text prompt (highlighted by green) and its corresponding image, to the right of the arrow is the generated attack prompt (highlighted by red) and its corresponding image.	24
4.3	(a),(b),(c),(d) are generated by mid-journey, and the corresponding prompts are, "a photo of a zebra and a giraffe"; "a photo of a tree frog and a cyclist"; "a photo of a volleyball and a Nemo"; and "a photo of a warthog and a traitor", respectively. (e),(f),(g),(h) are generated by DALL·E2, and the corresponding prompts are "a photo of a tench and a lizard"; "a photo of a volleyball and a Nemo"; "A photo of a howler monkey and a snake"; and "a photo of a silver salmon and a feather", respectively.	24
4.4	Black-box attack on mid-journey.	25
4.5	Black-box attack on DALL·E2.	26
5.1	A case study on "mountain" and "peafowl".	28
5.2	The generation speeds of "mountain" and "peafowl".	29
5.3	A violin plot illustrating the generation speeds of 1,000 images of various classes. The horizontal axis represents the number of steps taken, ranging from 49 to	

- 0, while the vertical axis displays the SSIM scores. The width of each violin represents the number of samples that attained a specific range of SSIM scores at a given step. 30
- 5.4 Images generated by the template "A photo of A and B ". 31
- 5.5 The first row illustrates the generation process with the prompt "a photo of a **silver salmon**". The second row, based on the forty-second step of the first row, shows the generation process with the prompt "a photo of a **feather**". The third row, also building upon the forty-second step of the first row, presents the generation procedure when the prompt is "a photo of a **magician**". The fourth row depicts the generation process in the presence of feature entanglement. The fifth row demonstrates the generation process for two distinct categories without feature entanglement. 31
- 5.6 a) "A photo of a **bat**"; b) "A photo of a bat and a **ball**"; c) Heat map of the word "bat" in generated image; d) "A photo of a **warthog**"; e) "A photo of a **warthog** and a **traitor**"; f) Heat map of the word "warthog" in generated image. 33
- 5.7 a) t-SNE Visualization of 100 images each of "**bat**", "**baseball bat**", "**bat and ball**" and text "a photo of a **bat**." b) The boxplot of cosine similarities between the text embedding of "a photo of a **bat**" and 100 of image embeddings each of "**bat**", "**baseball bat**", and "**bat and ball**". 33
- 5.8 a) "A type of **footwear** with a thick, rigid sole, often made of wood, and an upper made of leather or another material. **Clogs** can be open or closed, and are commonly associated with **Dutch** and Scandinavian cultures." b) "**footwear**" is replaced by "**pistol**". c) "**Dutch**" is replaced by "**pistol**". 35
- 5.9 A template, "A photo of A , B and C ", is used to generate prompts, where $A, B, C \in \{ \text{"cat"}, \text{"pistol"}, \text{"clogs"} \}$. For exmaple, "0-1-2" represents $A = \text{"cat"}$, $B = \text{"pistol"}$ and $C = \text{"clogs"}$, and so on. 35

CHAPTER 1

Introduction

In recent years, the field of text-to-image generation has witnessed remarkable advancements, paving the way for groundbreaking applications in computer vision and creative arts. Notably, many significant developments have captured the attention of researchers and enthusiasts, such as Stable Diffusion [21, 25], DALL·E [19, 20, 16] and Midjourney [12]. These developments push text-to-image synthesis boundaries, fostering artistic expression and driving computer vision research.

Despite the remarkable progress in text-to-image models, it is important to acknowledge their current limitations. One significant challenge lies in the instability and inconsistency of the generated outputs. In some cases, it can take multiple attempts to obtain the desired image that accurately represents the given textual input. An additional obstacle revealed by recent researches [26, 1, 4] is that the quality of generated images can be influenced by specific characteristics inherent to text prompts. Tang et al. [26] proposes DAAM, which performs a text-image attribution analysis on conditional text-to-image model and produces pixel-level attribution maps. Their research focuses on the phenomenon of feature entanglement and uncovers that the presence of cohyponyms may degrade the quality of generated images and that descriptive adjectives can attend too broadly across the image. Attend-and-Excite [1] investigates the presence of catastrophic neglect in the Stable diffusion model, where the generative model fails to include one or more of the subjects specified in the input prompt. Additionally, they discover instances where the model fails to accurately associate attributes such as colors with their respective subjects. Although those works have some progress, there is still work to be done to enhance the stability and reliability of text-to-image models, ensuring consistent and satisfactory results for a wide range of text prompts.




FIGURE 1.1. The overall pipeline of our attack method. The selection of words is relaxed to be differentiable by using a Gumbel Softmax with temperature τ . When $\tau \rightarrow 0$, the Gumbel Softmax outputs exhibit proximity to one-hot vectors while retaining differentiability. After an image is generated, a CLIP [18] classifier and a margin loss are employed to optimize ω aiming to generate images that cannot be correctly classified by CLIP.

One prominent constraint observed in those works related to the stability of text-to-image models lies in their dependence on manually crafted prompts for the purpose of vulnerability identification. This approach presents several challenges. Firstly, it becomes difficult to quantify the success and failure cases accurately, as the evaluation largely depends on subjective judgments and qualitative assessments. Additionally, the manual design of prompts can only uncover a limited number of potential failure cases, leaving many unexplored scenarios. Without a substantial number of cases, it becomes challenging to identify the underlying reasons for failures and effectively address them. To overcome these limitations, there is a growing demand for a learnable method that can automatically identify failure cases, enabling a more comprehensive and data-driven approach to improve text-to-image models. By leveraging such an approach, researchers can gain valuable insights into the shortcomings of current methods and develop more robust and reliable systems for generating images from textual descriptions.

In this thesis, we propose **Auto-attack on Text-to-image Models (ATM)**, to efficiently generate attack prompts with high similarity to given clean prompts (Fig. 1.1). We use Stable Diffusion [21, 25], a widely adopted open-source model, as our target model. With the open-source implementation and model parameters, we can generate attack prompts with a

white-box attack strategy. Remarkably, those attack prompts can transfer to other generative models, enabling black-box attacks. Two methods to modify a text prompt are considered, including **replacing** an existing word or **extending** with new ones. By incorporating a Gumbel Softmax distribution into the word embedding, the discrete modifications can be transformed into continuous ones, thereby ensuring differentiability. To ensure the similarity between clean and the attack prompts, a binary mask that selectively preserves the noun representing the desired object is applied. Moreover, two constraints are imposed: a **fluency constraint** that ensures the attack prompt is fluent and easy to read, and a **similarity constraint** that regulates the extent of semantic changes.

After the distribution is learned, ATM can sample multiple attack prompts at once. The attack prompts can prevent the diffusion model from generating desired subjects without modifying the nouns of desired subjects and maintain a high degree of similarity with the original prompt. We have achieved a 91.1% success rate in short-text attacks and a 81.2% success rate in long-text attacks. Moreover, drawing upon extensive experiments and empirical analyses employing ATM, we are able to disclose the existence of four distinct attack patterns, each of which corresponds to a vulnerability in the generative model: 1) the variability in generation speed; 2) the similarity of coarse-grained characteristics; 3) the polysemy of words; 4) the positioning of words. In the following, we will commence with an analysis of the discovered attack patterns in chapter 5, followed by a detailed exposition of our attack method in chapter 3.

In this thesis, we propose a novel method to automatically and efficiently generate plenty of successful attack prompts, which serves as a valuable tool for investigating vulnerabilities in text-to-image generation pipelines. This method enables the identification of a wider range of attack patterns, facilitating a comprehensive examination of the underlying causes. It will inspire the research community and garner increased attention toward exploring the vulnerabilities present in contemporary text-to-image models and will foster further research concerning both attack and defensive mechanisms, ultimately leading to enhanced security within the industry.

Literature review

2.1 Diffusion Model

In recent years, substantial progress has been made in the realm of content generation, with significant advances primarily attributed to the adoption of the diffusion probabilistic model and its subsequent adaptations. A thorough exploration of relevant literature, such as [23], [7], [14],[24], [21], [22] illustrates the model’s impressive achievements. The notable breadth of its success spans multiple domains, encompassing image generation, conditional image generation, video generation, and 3D scenes synthesis, as evidenced in a series of compelling studies [24], [7], [22].

The Denoising Diffusion Probabilistic Models (DDPM), proposed by Ho et al. [7], are specifically noteworthy for their innovative approach. These models employ a stepwise process of adding noise to images and subsequently learning to restore the original images from the noise-corrupted versions, providing an effective avenue for image generation.

Following DDPM, the Denoising Diffusion Implicit Models (DDIM) introduced by Song et al. [24], offered an advancement in the diffusion model’s efficacy. DDIM essentially enhanced the generation speed by bypassing certain steps in the inference process, demonstrating an impressive leap in model performance.

Furthermore, the introduction of the conditional latent diffusion model by Rombach et al. [21] marked a key milestone in the evolution of image generation technology. This novel framework leverages multiple conditions, such as textual content, images, and semantic maps, to guide the generation process within the latent space. Such an approach has not only

facilitated further acceleration of the inference speed but also expanded the scope of diffusion model applications.

One of the noteworthy implementations of the diffusion model is the Stable Diffusion model, also put forth by Rombach et al. [21]. This latent text-to-image diffusion model is capable of generating photo-realistic images from any given text input. Its superior performance has paved the way for subsequent enhancements [28], [8] [13] and widespread usage in contemporary AI-generated content products.

To illustrate, Stability-AI [25], Midjourney [12], DALL·E2 [16], and Runaway [3] are some of the AI-driven products that have harnessed the power of these models to generate striking content, reiterating the growing influence of diffusion models in the burgeoning field of AI-powered content generation.

2.1.1 Denoising Diffusion Probabilistic Models

DDPM offers a latent variable model that has the following form

$$p_{\theta}(\mathbf{x}_0) = \int p_{\theta}(\mathbf{x}_{0:T}) d\mathbf{x}_{1:T}, \quad (2.1)$$

where $p_{\theta}(\mathbf{x}_{0:T}) := p_{\theta}(\mathbf{x}_T) \prod_{t=1}^T p_{\theta}^{(t)}(\mathbf{x}_{t-1}|\mathbf{x}_t)$ and $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_T$ are latent variables that have the same dimension with original data instances $\mathbf{x}_0$. The parameters θ are trained to make the model match with the data distribution $q(\mathbf{x}_0)$. The training is conducted by maximizing a variational lower bound,

$$\max_{\theta} \mathbb{E}_{q(\mathbf{x}_0)} [\log p_{\theta}(\mathbf{x}_0)] \leq \max_{\theta} \mathbb{E}_{q(\mathbf{x}_0, \mathbf{x}_1, \dots, \mathbf{x}_T)} [\log p_{\theta}(\mathbf{x}_{0:T}) - \log q(\mathbf{x}_{1:T}|\mathbf{x}_0)] \quad (2.2)$$

The term $q(\mathbf{x}_{1:T}|\mathbf{x}_0)$ is defined as the forward process given by the original instance $\mathbf{x}_0$. In [7], the process is pre-defined as,

$$q(\mathbf{x}_{1:T}|\mathbf{x}_0) := \prod_{t=1}^T q(\mathbf{x}_t|\mathbf{x}_{t-1}), \quad (2.3)$$

where

$$q(\mathbf{x}_t|\mathbf{x}_{t-1}) := \mathcal{N}\left(\sqrt{\frac{\alpha_t}{\alpha_{t-1}}}\mathbf{x}_{t-1}, \left(1 - \frac{\alpha_t}{\alpha_{t-1}}\right)\mathbf{I}\right) \quad (2.4)$$

The term $p_\theta(\mathbf{x}_{0:T})$ in Eq. (2.1), a Markov chain from $\mathbf{x}_T$ to $\mathbf{x}_0$, is called the generative process as it approximately matches the reverse process as Eq. (2.3). From the Eq. (2.3) and Eq. (2.4), we have a special property,

$$q(\mathbf{x}_t|\mathbf{x}_0) := \mathcal{N}(\mathbf{x}_t; \sqrt{\alpha_t}\mathbf{x}_0, (1 - \alpha_t)\mathbf{I}). \quad (2.5)$$

If all the conditions are considered as Gaussian with trainable parameters, Eq. (2.2) can be re-written as,

$$L_\gamma(\epsilon_\theta) := \sum_{t=1}^T \gamma_t \mathbb{E}_{\mathbf{x}_0 \sim q(\mathbf{x}_0), \epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})} [\|\epsilon_\theta^t(\sqrt{\alpha_t}\mathbf{x}_0 + \sqrt{1 - \alpha_t}\epsilon_t) - \epsilon_t\|]. \quad (2.6)$$

2.1.2 Stable diffusion

The architecture of the Stable Diffusion [21] comprises of an encoder $\mathcal{E} : \mathcal{X} \rightarrow \mathcal{Z}$ and a decoder $\mathcal{D} : \mathcal{Z} \rightarrow \mathcal{X}$, where $\tilde{\mathbf{x}} = \mathcal{D}(\mathcal{E}(\mathbf{x}))$. Additionally, a conditional denoising network ϵ_θ and a condition encoder τ_θ are employed. In the text-to-image task, the condition encoder is a text encoder that maps text prompts into a latent space. The text prompt is typically a sequence of word tokens $\mathbf{c} = \{\mathbf{c}_1, \dots, \mathbf{c}_K\}$, where K is the sequence length.

During the image generation process, a random latent representation $\mathbf{z}_T$ is drawn from a distribution such as a Gaussian distribution. Then, the reverse diffusion process is used to gradually recover a noise-free latent representation $\mathbf{z}$. Specifically, a conditional denoising network $\epsilon_\theta(\mathbf{z}_t, t, \tau_\theta(\mathbf{c}))$ is trained to gradually denoise $\mathbf{z}_t$ at each time step $t = T, \dots, 1$ to gradually reduce the noise level of $\mathbf{z}_t$, where the condition $\mathbf{c}$ is mapped into a latent space using $\tau_\theta(\mathbf{c})$ maps and the cross-attention between the condition and features is incorporated in ϵ_θ to introduce the condition. Finally, the denoised latent representation $\mathbf{z}$ is decoded by a decoder $\mathcal{D}$ to produce the final output $\tilde{\mathbf{x}}$.

2.2 Vulnerabilities in Text-to-image Models.

With the open-source of Stable Diffusion [21], text-to-image generation achieves great process and shows the unparalleled ability on generating diverse and creative images with the guidance of a text prompt. However, there are some vulnerabilities have been discovered in existing works [4, 1, 26].

This is an important finding. Structured Diffusion [4] did reveal that some attributes were not correctly assigned in the cues during image generation, which could lead to the generated images deviating from the expected themes and attributes. To address this issue, they introduced a structure called consistency trees or scene graphs to enhance the learning of the embedding of the input cues. These structured tools help the model to better understand the topics and attributes in the input cues and to correctly represent them in the generated images. In this way, the model can map each attribute in the input cue to the correct part of the generated image, thus improving the accuracy and quality of the generated image. This explains why the use of consistency trees or scene graphs can enhance the performance of the Stable Diffusion model. At the same time, this approach points to a possible direction of improvement in the performance of AI models when dealing with complex inputs by introducing more structured information and tools. Not only the task of generating images, this approach may also be useful for other tasks that require understanding and generating complex information, such as natural language processing or reasoning tasks. This provides valuable insights into the development of future AI models.

In addition, Attend-and-Excite [1] also introduces that the Stable Diffusion model fails to generate one or more of the subjects from the input prompt and fails to correctly bind attributes to their corresponding subjects. When dealing with input prompts, the Stable Diffusion model sometimes fails to generate one or more topics, possibly because the model is unable to correctly understand or parse the topics in the input prompts. Also, it often fails to correctly bind attributes to corresponding topics, possibly due to inherent limitations in the model's ability to analyse semantic relationships. For example, the model tends to struggle when input prompts involve complex semantic structures or ambiguous conceptual relationships.

This evidence suggests that the current Stable Diffusion model has a number of obvious flaws. Although the model performs reasonably well in many cases, these problems reveal its inadequacy in understanding and generating complex linguistic structures. To address these issues, we may need to introduce more structured information into the training process of the model to help it better understand the relationships between topics and attributes. At the same time, we may also explore new model structures that can generate topics and attributes more accurately when dealing with complex input cues. This also hints at a major challenge in AI model development: how to ensure that models can accurately understand and generate information when faced with a variety of complex language scenarios. Future research needs to make more efforts in this area to improve the stability and accuracy of the Stable Diffusion model as well as other related models.

The DAAM model proposed by Tang et al. [26] performs text-image attribute analysis on a conditional text-to-image model and generates pixel-level attribute maps. This certainly plays an important role in improving the interpretability of text-to-image models. More importantly, their research delves into the phenomenon of feature entanglement. By feature entanglement, we mean that different text features may interact with each other during the image generation process, resulting in a degradation of the quality of the generated images. This phenomenon is particularly evident when the model deals with concatenated words (cohyponyms), as the presence of concatenated words may make the model confuse when generating images, thus degrading the quality of the images. In addition, they found that descriptive adjectives may have too wide a focus on the image. This means that when the input text contains descriptive adjectives, the model may overextend the sphere of influence of these adjectives, resulting in the loss of fine detail in the generated images. This research certainly provides important clues to understanding and improving text-to-image models. By conducting an in-depth study of the feature entanglement phenomenon and the impact of descriptive adjectives, we may find new ways to improve model performance and image quality. In addition, pixel-level attribute maps provide new tools for understanding the working mechanisms of our models, helping us to design more effective and interpretable models.

2.3 Adversarial Attack

2.3.1 Image Adversarial Attack

Madry et al. [11] mentioned recent advancements in fields like computer vision and natural language processing have placed trained classifiers at the heart of security-critical systems, with prominent examples including autonomous vehicles, facial recognition, and malware detection. As these systems become integral parts of our daily life, the importance of security measures, specifically resistance to adversarial attacks, can't be overstated.

The realm of deep neural networks, particularly in the context of computer vision, has been the hotbed for the exploration of this issue. Interestingly, even infinitesimal modifications to an input image can mislead the most sophisticated neural networks, often imperceptible to the human eye. This vulnerability not only bears security ramifications but also unveils that our current models may lack robust concept learning capabilities. This brings to the fore a crucial question - how to train deep neural networks to resist adversarial inputs?

Considerable work has been devoted to proposing different adversarial attack and defense mechanisms, such as defensive distillation, feature squeezing, and other adversarial example detection methods. Though these provide a launchpad to explore the space of possibilities, they often fall short of offering tangible guarantees against a well-defined class of adversarial attacks, making it challenging to thoroughly navigate the adversarial robustness landscape.

This paper aims to shed light on the adversarial robustness of neural networks through the prism of robust optimization, employing a min-max formulation to capture security against adversarial attacks. This approach facilitates a more precise understanding of the desired security guarantees and enables a universal theoretical framework for both attacks and defenses. The paper further makes significant strides in understanding the optimization landscape of this min-max formulation, the impact of network architecture on adversarial robustness, and how to train networks robust against a variety of adversarial attacks on datasets like MNIST and CIFAR10.

2.3.2 Text Adversarial Attack

Despite the unprecedented success of deep learning, a body of recent research suggests that neural networks can be compromised by adversarial examples, inputs that have undergone small perturbations. Such adversarial samples are typically imperceptible to humans, yet they can lead neural networks to erroneous predictions. Consequently, the exploration of these adversarial attack methods is vital for achieving robust and reliable neural networks.

Adversarial attack strategies and their countermeasures have been thoroughly examined in the domain of computer vision, but in the field of text, they remain a challenge due to language's inherently discrete nature. Generating adversarial samples for text requires they be inconspicuous to humans yet misleading for neural models, while maintaining grammatical fluency and semantic consistency with the original inputs.

Past methods have relied heavily on specific rules to create adversarial samples, making it difficult to ensure fluency and semantic preservation simultaneously in the generated adversarial samples. These rule-based methods can be complex, involving multiple linguistic constraints like Named Entity Recognition (NER) or Part-of-Speech (POS) tagging. The introduction of contextualized language models as automatic perturbation generators could simplify these rule design processes.

The advent of pre-trained language models like BERT has elevated the performance of NLP tasks significantly. BERT's dual aspects, namely, its resilience against adversarial attacks in its fine-tuned form and its potential to generate fluent and semantically-consistent substitutions, inspire us to investigate the viability of using one BERT to attack another fine-tuned BERT.

Li et al. [9], present BERT-Attack, a high-quality adversarial sample generation method that employs BERT as a language model to produce adversarial samples. The algorithm is simple, comprising two stages: identifying vulnerable words in a given input sequence for the target model and then using BERT to create semantic-preserving substitutes for those vulnerable words.

BERT-Attack method successfully deceives its fine-tuned downstream model, achieving the highest attack success rate compared to previous methods, while keeping the perturbation percentage and the number of queries low, and maintaining high semantic preservation.

Methods

3.1 Methodology.

We aim to design an automatic attack method that targets the recent popular text-to-image models. The objective of our method is to identify attack prompts $\mathbf{c}'$ based on a clean prompt $\mathbf{c}$, which leads to a vision model $h : \mathcal{X} \rightarrow \mathcal{Y}$ failing to predict the desired class y , i.e. $\arg\max_i h(\tilde{\mathbf{x}})_i \neq y$:

$$\mathbf{c}' = \arg\max_{\mathbf{c}' \in B_d(\mathbf{c}, \xi)} \ell(y, h(\text{GM}(\mathbf{c}')))) \quad (3.1)$$

where $\text{GM}(\cdot)$ is a text-to-image model, $B_d(\mathbf{c}, \xi) = \{\mathbf{c}' : d(\mathbf{c}, \mathbf{c}') \leq \xi\}$, $d(\cdot, \cdot)$ is a distance measure regularizing the similarity between $\mathbf{c}$ and $\mathbf{c}'$, and ξ is a maximum distance. To enable auto-attack, a differentiable method that can be optimized using gradient descent is desired. We introduce a Gumbel-Softmax sampler to enable differentiable modifications on text prompts during the word embedding phase. To minimize the distance $d(\mathbf{c}, \mathbf{c}')$, we introduce two constraints, including a fluency constraint and a similarity constraint.

In our experimental setup, the open-source Stable Diffusion model is employed as the targeted generative model $\text{GM}(\cdot)$. By generating white-box attack prompts for Stable Diffusion, we can subsequently transfer these prompts to other generative models to execute black-box attacks. To facilitate the classification task, we utilize a CLIP classifier as the vision model $h(\cdot)$, benefiting from its exceptional zero-shot classification accuracy. To establish the desired classes, we employ the 1,000 classes derived from ImageNet-1K. In the case of generating short prompts, a fixed template of "A photo of [CLASS_NAME]" is utilized to generate the prompts. Conversely, for the generation of long prompts, we employ GPT 4 [17] as a prompt

generation model. Subsequently, human experts verify the correctness of the prompts and check that the prompts indeed contain the noun associated with the desired class.

3.1.1 Differentiable Text Prompt Modification

A text prompt is typically a sequence of words $\mathbf{c} = \{c_1, \dots, c_K\}$. Owing to the discrete nature of text prompts $\mathbf{c}$, perturbations can be incorporated either by replacing an existing word c_k where $1 \leq k \leq K$ or augmenting with new ones $\{c_{K+i} | 1 \leq i \leq K'\}$. However, the non-differentiable nature of this procedure makes it unsuitable for optimization utilizing gradient-based techniques. Therefore, there is a need for a mechanism that guarantees the differentiability of the word selection process. In this regard, we integrate a Gumbel Softmax sampler $\psi(\cdot; \tau)$ into the word embedding phase. The Gumbel Softmax function has the ability to approximate a one-hot distribution as the temperature $\tau \rightarrow 0$. Additionally, the Gumbel distribution has the ability to introduce further randomness, thereby enhancing the exploitability during the initial stages of perturbation search.

Differentiable Sampling. We employ a trainable matrix $\omega \in \mathbb{R}^{K \times V}$ to learn the word selection distribution, where K is the length of the text prompt and V is the vocabulary size. In the scenario of augmenting with a new word, the sequence length K can be extended to $K + K'$ to facilitate the addition of new words. The Gumbel Softmax can be represented as follows:

$$\text{GumbelSoftmax}(\omega_k; \tau) := \frac{\exp((\omega_{k,i} + g_{k,i})/\tau)}{\sum_j \exp((\omega_{k,j} + g_{k,j})/\tau)}, \quad (3.2)$$

where $g_{k,i} \sim \text{Gumbel}(0, 1)$ are i.i.d. samples from a Gumbel distribution. The word embedding stage employs a matrix $E \in \mathbb{R}^{V \times D}$, where V is the vocabulary size and D is for the embedding dimensionality. The discrete process of word embedding is to select E_i from E based on the index $1 \leq i \leq V$ of a word in the vocabulary. To make this process differentiable, the dot product between $\text{GumbelSoftmax}(\omega_k)$ and E can be calculate:

$$\mathbf{c}'_k = \psi(\omega_k; \tau) = \text{GumbelSoftmax}(\omega_k; \tau) \cdot E \approx E_i, \quad \text{s.t. } i = \underset{i}{\operatorname{argmax}} \omega_{k,i}, \quad (3.3)$$

where $\mathbf{c}'_k$ is the new word selected according to ω_k . As $\tau \rightarrow 0$, Eq. 3.3 can effectively emulate an argmax selection procedure.

Additionally, in order to ensure similarity, it is desired to preserve the noun representing the desired object in the new prompt. This is achieved by utilizing a binary mask $M \in \{0, 1\}^K$, where the position corresponding to the desired noun is set to 0 while other positions are set to 1. By computing $\mathbf{c}' \leftarrow (1 - M) \cdot \mathbf{c} + M \cdot \mathbf{c}'$, the desired noun can be retained in the prompt while other words can be modified.

Attack Objective. To generate images $\tilde{\mathbf{x}}$ that cannot be correctly classified by the classifier, a margin loss can be used as the loss function $\ell(\cdot, \cdot)$ in Eq. 3.1:

$$\ell_{\text{margin}}(\tilde{\mathbf{x}}, y; h) = \max \left(0, h(\tilde{\mathbf{x}})_y - \max_{i \neq y} h(\tilde{\mathbf{x}})_i + \kappa \right), \quad (3.4)$$

where κ is a margin. Eq. 3.4 reduces the classifier’s confidence on the true class y and improve its confidence on the class with the largest confidence, excluding y until a margin κ is reached.

3.1.2 Constraints on Fluency and Similarity

Given that we search for perturbations in a $\mathbb{R}^{K \times V}$ space to attack the text prompt, the attack prompts may be too diverse if the added perturbations are not properly constrained, making it easily detectable. Eq. 3.1 includes a distance constraint such that $d(\mathbf{c}, \mathbf{c}') \leq \xi$, which ensures that the added perturbations are subtle and hard to notice. The measurement of distance between two pieces of text can be approached through various methods. We introduce two constraints to reduce this distance, namely a **fluency** constraint and a **semantic similarity** constraint. The fluency constraint ensures that the generated sentence is smooth and readable, while the semantic similarity semantic constraint regularize the semantic changes introduced by the perturbations, making the attack prompt $\mathbf{c}'$ closely resemble the clean prompt $\mathbf{c}$.

Fluency constraint. The fluency constraint can be achieved visa a Casual Language Model (CLM) ϕ with log-probability outputs [5]. The next token distribution we learn is compared with the next token distribution predicted by ϕ . Given a sequence of perturbed text $\mathbf{c}'$, we use

ϕ to predict the a token $\mathbf{c}'_i$ based on $\{\mathbf{c}'_1, \dots, \mathbf{c}'_{i-1}\}$. Therefore, we can have a log-likelihood of the possible next word:

$$\log p_\phi(\mathbf{c}'_i | \mathbf{c}'_1, \dots, \mathbf{c}'_{i-1}) = \phi(\mathbf{c}'_1, \dots, \mathbf{c}'_{i-1}). \quad (3.5)$$

The next token distribution we learn can be easily obtained by $\text{GumbelSoftmax}(\boldsymbol{\omega}_i; \tau)$. Subsequently, a cross-entropy loss function can be employed to optimize the learned distribution:

$$\text{CE}_\phi(\boldsymbol{\omega}) = - \sum_{i=1}^K \sum_{j=1}^D \text{GumbelSoftmax}(\boldsymbol{\omega}_i; \tau)_j \cdot \phi(\psi(\boldsymbol{\omega}_1; \tau), \dots, \psi(\boldsymbol{\omega}_{i-1}; \tau))_j. \quad (3.6)$$

Eq. 3.6 serves as a regularizer to encourage the next word selection distribution to resemble the prediction of the CLM ϕ , thereby ensuring fluency.

Semantic similarity constraint. Rather than simply considering a word similarity, we concern more about semantic similarity. One prominent metric used to evaluate semantic similarity is the BERTScore [29]. The calculation of BERTScore requires contextualized word embeddings. The aforementioned CLM ϕ is used again to extract the embeddings $\mathbf{v} = \phi^{(\text{emb})}(\mathbf{c})$, where $\phi^{(\text{emb})}$ denotes the embedding network used in ϕ . The BERTScore between the clean prompt $\mathbf{c}$ and the attack prompt $\mathbf{c}'$ can be calculated by

$$S_{\text{BERT}}(\mathbf{c}, \mathbf{c}') = \sum_{i=1}^N w_i \max_{j=1, \dots, M} \mathbf{v}_i^\top \mathbf{v}'_j, \quad (3.7)$$

where $w_i := \text{idf}(\mathbf{c}_i) / \sum_{i=1}^N \text{idf}(\mathbf{c}_i)$ is the normalized inverse document frequency (idf), $N = K$, and M is either K or $K + K'$ depending on whether existing words are being replaced or new words are being added. To improve the similarity, we use $1 - S_{\text{BERT}}(\mathbf{c}, \mathbf{c}')$ as the our loss term.

The constrained objective. Considering that the addition of constraints may limit the diversity of perturbation search, we introduce two hyper-parameters, λ and γ , to control the strength of the constraints. Then, the overall objective function can be written as:

$$\max_{\boldsymbol{\omega}} \quad \mathbb{E}_{\mathbf{z}_T \sim \mathcal{N}(0,1), \mathbf{c}' = \psi(\boldsymbol{\omega}; \tau)} [\ell_{\text{margin}}(\text{GM}(\mathbf{z}_T | \mathbf{c}'), y; h)] \quad (3.8)$$

$$\text{s.t.} \quad \min_{\boldsymbol{\omega}} \quad \lambda \cdot \text{CE}_\phi(\boldsymbol{\omega}) + \gamma \cdot (1 - S_{\text{BERT}}(\mathbf{c}, \mathbf{c}')). \quad (3.9)$$

Algorithm 1 Auto-attack on Text-to-image Models (ATM)

Input: The maximum number of iterations T . The maximum number of attack candidates N . The clean prompt $\mathbf{c}$. The desired class y . A binary mask M . A learning rate η .

Output: A set of attack prompts $\mathcal{S}$.

```

1: Initialize  $\omega$ 
2: for  $t = 1 \rightarrow T$  do ▷ The search stage
3:   Sample an attack prompt  $\mathbf{c}' = \{\psi(\omega_k; \tau) | 1 \leq k \leq K\}$ 
4:   Apply the mask by  $\mathbf{c}' \leftarrow (1 - M) \cdot \mathbf{c} + M \cdot \mathbf{c}'$ 
5:   Generate an image  $\tilde{\mathbf{x}}' = \text{GM}(\mathbf{z}_T | \mathbf{c}')$ 
6:   Get classification results  $y' = h(\tilde{\mathbf{x}}')$ 
7:   Conduct a gradient descent step  $\omega \leftarrow \omega - \eta \cdot \nabla_{\omega} \mathcal{L}(\omega)$ 
8: end for
9: Initialize  $\mathcal{S} = \emptyset$ 
10: for  $n = 1 \rightarrow N$  do ▷ The attack stage
11:   Sample an attack prompt  $\mathbf{c}' = \{\psi(\omega_k; \tau) | 1 \leq k \leq K\}$ 
12:   Apply the mask by  $\mathbf{c}' \leftarrow (1 - M) \cdot \mathbf{c} + M \cdot \mathbf{c}'$ 
13:   Generate an image  $\tilde{\mathbf{x}}' = \text{GM}(\mathbf{z}_T | \mathbf{c}')$ 
14:   if  $\arg\max h(\tilde{\mathbf{x}}') \neq y$  then ▷ If attack success
15:     Save the success attack prompt  $\mathcal{S} \leftarrow \mathcal{S} \cup \{\mathbf{c}'\}$ 
16:   end if
17: end for

```

3.1.3 Generation of Attack Prompts

The overall procedure of ATM is as described in Algorithm 1. It consists of two stages: a search stage, where the Gumbel Softmax distribution is learned, and an attack stage, where we generate attack prompts using the learned distribution. In the search stage, we use gradient descent to optimize the parameters ω for each clean prompt $\mathbf{c}$ over a period of T iterations. Once ω is learned, we proceed to the attack stage. In this stage, we sample N attack prompts from each learned ω . An attack prompt $\mathbf{c}'$ is considered successful if the image $\tilde{\mathbf{x}}'$ generated from it cannot be correctly classified by the visual classifier h .

CHAPTER 4

Results

In our experiments, we conduct comprehensive analyses of both long and short prompts. Furthermore, we conduct ablation studies specifically on long prompts, focusing on three key aspects. Firstly, we evaluate our attack method with different numbers of search steps T . Secondly, we investigate the influence of our constraints, including fluency and semantic similarity as measured by BERTScore. Lastly, we attack different samplers, including DDIM [24] and DPM-Solver [10]. Moreover, we verify that the generated attack prompt able to influence DALL-E2 and mid-journey via black box attack.

4.1 Experimental Setting.

Attack hyperparamters. The number of search iterations T is set to 100. This value determines the number of iterations in the search stage, during which we aim to find the most effective attack prompts. The number of attack candidates N is set to 100. This parameter specifies the number of candidate attack prompts considered in the attack stage, allowing for a diverse range of potential attack prompts to be explored. The learning rate η for the matrix ω is set to 0.3. The margin κ in the margin loss is set to 30.

Text prompts. Our experiments consider the 1,000 classes from ImageNet-1K [2], which serves as the basis for generating images. To explore the impact of prompt length, we consider both short and long prompts. For clean short prompts, we employ a standardized template: "A photo of [CLASS_NAME]". Clean long prompts, on the other hand, are generated using GPT 4 [17], with a prompt length restriction of 77 tokens to align with the upper limit of the CLIP [18] word embedder.

TABLE 4.1. Main results of short-prompt and long-prompt attacks.

Prompt	Method	Success (%)	FID ($\downarrow$)	IS ($\uparrow$)	TS ($\uparrow$)
Short	Clean	-	18.51	101.33 $\pm$ 1.80	1.00
	Random	79.2	29.21	66.71 $\pm$ 0.87	0.69
	ATM (Ours)	91.1	30.09	65.98 $\pm$ 1.10	0.72
Long	Clean	-	17.95	103.59 $\pm$ 1.68	1.00
	Random	41.4	24.16	91.33 $\pm$ 1.58	0.94
	ATM (Ours)	81.2	29.65	66.09 $\pm$ 1.83	0.84

Evaluation metrics. To evaluate the effectiveness of our attack method, we generate attack prompts from the clean prompts. We focus on three key metrics: **success rate**, Fréchet inception distance [6] (**FID**), Inception Score (**IS**), and text similarity (**TS**). Subsequently, 50,000 images are generated using the attack prompts, ensuring a representative sample of 50 images per class. The success rate is determined by dividing the number of successful attacks by the total of 1,000 classes. FID and IS are computed by comparing the generated images to the ImageNet-1K validation set with (**torch-fidelity**)[15]. TS is calculated by embedding the attack prompts and clean prompts using the CLIP [18] word embedder, respectively. Subsequently, the cosine similarity between the embeddings is computed to quantify the text similarity.

4.2 Main Results

Table 4.1 reports our main results, including short-prompt and long-prompt attacks. Compared to long text prompts, short text prompts comprise only a small number of tokens. This leads to a relatively fragile structure that is extremely vulnerable to slight disturbance. Therefore, random attacks can reach an impressive success rate of 79.2% targeting short prompts but a low success rate of 41.4% targeting the long prompts. In contrast, our algorithm demonstrates its true potential, reaching an impressive success rate of 91.1% and 81.2% targeting short and long prompts, respectively.

As a further evidence of the effectiveness of our algorithm, it’s worth noting the text similarity (TS) metrics between the random attacks and our algorithm’s outputs. For short-prompt attack,

TABLE 4.2. Results of the ablation study on the number of steps in attack prompt search.

#Steps	Success (%)	FID ($\downarrow$)	IS ($\uparrow$)	TS ($\uparrow$)
50	68.7	34.00	93.94 ± 1.84	0.97
100	81.2	29.65	66.09 ± 1.83	0.84
150	67.2	45.23	58.51 ± 0.79	0.82

the values stand at 0.69 and 0.72, respectively, illustrating that the semantic information of short texts, while easy to disrupt, can be manipulated by a well-designed algorithm with fluency and semantic similarity constraints. Our attacks preserve more similarity with the clean prompts. For long-prompt attacks, the TS score of random attacks (0.94) is higher compared to our attacks (0.84). One possible reason is that random attacks tend to make only minimal modifications as the length of the prompt increases. This limited modification can explain the significantly lower success rate of random attacks on longer prompts.

From the perspective of image generation quality and diversity, we found that as the attack success rate increases, image generation quality and diversity will decrease. For short and long texts, images generated from the clean text have the lowest FID (18.51 and 17.95) and the highest IS (101.33 ± 1.80 and 103.59 ± 1.68). As the attack success rate rises, FID shows an upward trend. Examining this situation from the perspective of FID, a metric that gauges the distance between the distribution of generated images and the original data set. As the attack becomes more successful, the image set generated by the attack prompt tends to deviate substantially from the distribution of the original data set. This divergence consequently escalates the FID score, indicating a larger distance between the original and generated distributions. On the other hand, considering this situation from the diversity standpoint, it appears that the suppression of the generation of original categories brought on by the successful attack might instigate a decrease in diversity. This reduction in diversity, in turn, may cause a decrease in the Inception Score (IS).

4.3 Different Search Steps

Table 4.2 presents the results of using different numbers of steps T in the search stage. For the $T = 50$ step configuration, the success rate is 68.7%. The FID value is 34.00, with lower values suggesting better image quality. The IS is reported as 93.94 ± 1.84 , with higher values indicating diverse and high-quality images. The TS value is 0.97, representing a high level of text similarity. Moving on to the $T = 100$ step configuration, the success rate increases to 81.2%, showing an improvement compared to the previous configuration. The FID value decreases to 29.65, indicating better image quality. The IS is reported as 66.09 ± 1.83 , showing a slight decrease compared to the previous configuration. The TS value is 0.84, suggesting a slight decrease in text similarity. In the $T = 150$ step configuration, the success rate decreases to 67.2%, slightly lower than the initial configuration. The FID value increases to 45.23, suggesting a decrease in image quality. The IS is reported as 58.51 ± 0.79 , indicating a decrease in the diversity and quality of generated images. The TS value remains relatively stable at 0.82.

When using $T = 50$, the attack prompt fails to fit well and exhibits a higher text similarity with the clean prompt. Although the generated images at this stage still maintain good quality and closely resemble those generated by the clean prompt, the success rate of the attack is very low. On the other hand, when $T = 150$, overfitting occurs, resulting in a decrease in text similarity and image quality due to the overfitted attack prompt. Consequently, the success rate of the attack also decreases. Overall, the configuration of $T = 100$ proves to be appropriate.

4.4 The Impact of Constraints

Table 4.3 examines the impact of the fluency and semantic similarity (BERTScore) constraints. When no constraints are applied, the attack success rate is notably high at 91.3%. However, this absence of constraints results in a lower text similarity (TS) score of 0.37, indicating a decreased resemblance to clean text and a decrease in image quality. By introducing fluency

TABLE 4.3. Results of the ablation study on the constraints

Fluency	BERTScore	Success (%)	FID ($\downarrow$)	IS ($\uparrow$)	TS ($\uparrow$)
$\times$	$\times$	91.3	39.14	47.21 ± 1.25	0.37
$\checkmark$	$\times$	81.7	29.37	64.93 ± 1.57	0.79
$\times$	$\checkmark$	89.8	39.93	46.94 ± 0.99	0.51
$\checkmark$	$\checkmark$	81.2	29.65	66.09 ± 1.83	0.84

TABLE 4.4. Results of the ablation study on the samplers

Sampler	Success (%)	FID ($\downarrow$)	IS ($\uparrow$)	TS ($\uparrow$)
DDIM [24]	81.2	29.65	66.09 ± 1.83	0.84
DPM-Solver [10]	76.5	27.23	81.31 ± 2.09	0.88

constraints alone, the attack success rate decreases to 81.7% but increases the text similarity to 0.79. Furthermore, incorporating semantic similarity constraints independently also leads to a slight reduction in success rate to 89.8%, but only marginally improves the text similarity to 0.51. The introduction of constraints, particularly fluency constraints, leads to an increase in text similarity. The fluency constraint takes into account the preceding tokens of each token, enabling the integration of contextual information for better enhancement of text similarity. On the other hand, BERTScore considers a weighted sum, focusing more on the similarity between individual tokens without preserving the interrelation between context. In other words, the word order may undergo changes as a result and leads to a low text similarity. Certainly, this outcome was expected, as BERTScore itself prioritizes the semantic consistency between two prompts, while the order of context may not necessarily impact semantics. This further highlights the importance of employing both constraints simultaneously. When both constraints are utilized together, the text similarity is further enhanced to 0.84. Meanwhile, the success rate of the attack (81.2%) is comparable to that achieved when employing only the fluency constraint, while the text similarity surpasses that obtained through the independent usage of the two constraints.

4.5 Different Samplers

Table 4.4 illustrates the effectiveness of our attack method in successfully targeting both DDIM and the stronger DPM-Solver. For the DDIM sampler, our attack method achieves a success rate of 81.2%, indicating its ability to generate successful attack prompts. Similarly, our attack method demonstrates promising results when applied to the DPM-Solver sampler. With a success rate of 76.5%, it effectively generates attack prompts. The TS scores of 0.84 and 0.88, respectively, indicate a reasonable level of text similarity between the attack prompts and clean prompts. These outcomes demonstrate the transferability of our attack method, showcasing its effectiveness against both DDIM and the more potent DPM-Solver sampler.

4.6 Cases of Short/Long-Prompt Attacks

In Fig. 4.1 and Fig. 4.2, we demonstrate more cases of long and short text prompt attacks respectively.

4.7 Black-box attack

To further investigate whether our generated attack prompts can be transferred to different text-to-image models, we randomly selected attack prompts to attack DALL·E2 and mid-journey, respectively. The experimental results (Fig. 4.3) prove that our attack prompts can also be used for black-box attacks. More results of black-box attacks are demonstrated in Fig. 4.4, and Fig. 4.5.



	→			→			→	
A loose, full-length abaya worn by some Muslim women, primarily in the Arabian Peninsula, as an outer covering to ensure modesty.		a knee, full adult or abaya embroidery by some German women, primarily in the Netherlands peninsula as an outer covering to ensure modesty.		An ancient abacus tool, consisting of a frame with rows of beads or other movable counters, used for mathematical operations and counting.		An optical abacus tool, consisting of a wheel with rows of tiles and other movable symbols, used for wheel operations and counting.		A large, mounted cannon that fires heavy projectiles, historically used in warfare and for other purposes such as ceremonial salutes."
	→			→			→	
This small brussels griffon breed has a distinctive, wiry coat, a compact, muscular body, and a friendly, affectionate temperament. They were originally bred for hunting small game and as companion animals in Belgium and are known for their intelligence, agility, and charm		This ancient brussels griffon penguins has a distinctive, wiry appearance, a strong, muscular body, and a good, affectionate temperament. They were originally bred for hunting small game and for pack animals in Belgium and are known for their intelligence, agility, and speed.		A type of footwear with a thick, rigid sole, often made of wood, and an upper made of leather or another material. Clogs can be open or closed, and are commonly associated with Dutch and Scandinavian cultures.		A type of camouflage with a long, rigid beak, often made of wood, and an upper made of leather or another material. Clogs can be open or closed, and are commonly associated with Dutch and Scandinavian cultures.		"A loose, sleeveless cloak that is worn over the shoulders and fastened at the neck, often used for warmth or as a decorative accessory."
	→			→			→	
A computer mouse input device used to control a computer's cursor, typically featuring buttons and a scroll wheel for selecting and navigating on-screen elements.		A computer mouse puzzle device used to control a computer's cursor, typically featuring buttons and a scroll wheel for selecting and navigating on-screen elements.		These large, conch snails have a distinctive, spiral shell and a strong, muscular foot. They are found in tropical waters and are known for their ornate shells and edible flesh."		These large, conch sharks have a distinctive, spiral shell and a strong, muscular foot. They are found in tropical waters and are known for their ornate shells and edible flesh."		A cornet musical instrument similar to a trumpet but with a more compact shape and mellower tone.
	→			→			→	
"A large, upholstered couch of furniture designed for seating multiple people, often found in living rooms or other communal spaces."		"A large, upholstered couch of weapons designed for seating multiple people, often found in living rooms or other communal spaces."		"These duck have a round body, a flat bill, and webbed feet, adapted for swimming and dabbling for food. They come in various colors and patterns, with some species having bright, iridescent plumage."		These duck shark, a round body, a flat bill, and webbed feet, adapted for swimming and dabbling for food. They come in various colors and patterns, with some species having bright, iridescent plumage."		A rich, creamy, dairy-based eggnog traditionally made of milk, cream, sugar, and eggs, often spiked with alcohol and served during the winter holidays.
	→			→			→	
Electric ray is flattened, disc-shaped fish with smooth skin and a pair of kidney-shaped electric organs on each side of their body, used for hunting prey and self-defense."		Electric rays flattened, disc-shaped senses with smooth skin and a pair of kidney-shaped electric organs on each side of their body, used for hunting prey and self-defense."		A timekeeping hourglass consisting of two glass chambers connected by a narrow passage, with sand flowing through the passage at a constant rate to measure a specific time interval.		A mean like hourglass consists of two renrepanels separated by a central passage, with sunlight flowing through the cave in a constant rate to enter a specific time frame."		A wild pig species native to Africa, known for its large, curved tusks and distinctive facial features, including warts on their faces. Warthogs are primarily grazers and live in savannas and grasslands.

FIGURE 4.1. To the left of the arrow is the clean long text prompt (highlighted by green) and its corresponding image, to the right of the arrow is the generated attack prompt (highlighted by red) and its corresponding image.




FIGURE 4.2. To the left of the arrow is the clean short text prompt (highlighted by green) and its corresponding image, to the right of the arrow is the generated attack prompt (highlighted by red) and its corresponding image.



FIGURE 4.3. (a),(b),(c),(d) are generated by mid-journey, and the corresponding prompts are, "a photo of a zebra and a giraffe"; "a photo of a tree frog and a cyclist"; "a photo of a volleyball and a Nemo"; and "a photo of a warthog and a traitor", respectively. (e),(f),(g),(h) are generated by DALL-E2, and the corresponding prompts are "a photo of a tench and a lizard"; "a photo of a volleyball and a Nemo"; "A photo of a howler monkey and a snake"; and "a photo of a silver salmon and a feather", respectively.




FIGURE 4.4. Black-box attack on mid-journey.



A photo of a **sea lion** and a **spaniel**




A photo of a **stingray** and a **vulture**



A photo of a **wolf spider** and a **frog**




A photo of a **balloon** and a **brakes spear**



A photo of an **armadillo** and a **whale**




A type of camouflage with a long, rigid beak, often made of wood, and an upper made of leather or another material. **Clogs** can be open or closed, and are commonly associated with Dutch and Scandinavian cultures.



A type of cork- and double air aircraft, also known as a chrigible, (ized by its elongated shape, powered by engines, and steered by a rudder. **airships** typically consist of a rigid plane covered with an envelope filled with oxygen, such as helium or hydrogen.



A cloth, circular **yurt** statue by nomadic peoples in central Asia, has a wooden frame and a balance or praying covering.



A mean like **hourglass** consists of two renfrepanels separated by a central passage, with sunlight flowing through the cave in a constant rate to enter a specific time frame.



A large, thin **artichoke** crab with long, fleshy, leaf -, scales, often oil or steamed and eaten by removing the thin and wiping off the body surface with the teeth.

FIGURE 4.5. Black-box attack on DALL·E2.

Empirical analysis

5.1 Vulnerabilities of Stable Diffusion Model

By applying the attack method proposed in this thesis, a variety of attack text prompts can be generated and analyzed. In this chapter, the identified attack patterns are discussed. Details of the attack method are introduced in chapter 3. We have discovered four distinct patterns of impairment in Stable Diffusion model: 1) the variability in generation speed, where the model struggles to reconcile the differences in generation speed among various categories effectively. 2) the similarity of coarse-grained characteristics, which arises from the feature entanglement of global or partial coarse-grained characteristics, which possess a high degree of entanglement. 3) the polysemy of words, which involves the addition of semantically complementary words to the original prompt, resulting in the creation of images that contain brand-new content and are not related to the original category. 4) the positioning of words, where the position of the category word within the prompt influences the final outcome of the generated image.

5.1.1 Variability in Generation Speed

OBSERVATION 1. When a given text prompt contains a noun A representing the object to be generated, the introduction of another noun B into the prompt, either through addition or replacement, leads to the disappearance of noun A from the generated image, replaced instead by noun B .




FIGURE 5.1. A case study on "mountain" and "peafowl".

In Observation 1, we identify a phenomenon when replacing or adding a noun in the description in a text prompt, the new noun will lead to the completely disappear of the desired subject. As shown in Fig. 5.1, if the noun "peak" is replaced by "peafowl", the desired subject "mountain" disappears and the new subject "peafowl" is generated. To further investigate this phenomenon, we use two short prompts c_1 and c_2 containing "mountain" and "peafowl", respectively, to exclude the influence of other words in the long prompts. To eliminate the additional impact of all possible extraneous factors, such as contextual relationships, they are embedded separately and then concatenated together: $\text{Concat}(\tau_\theta(c_1), \tau_\theta(c_2))$. The result shows that almost no element of the mountain is visible in the generated image (Fig. 5.1).

Further analysis reveals a nontrivial difference in the generation speeds of the two subjects. To define the generation speed, a metric to measure the distance from a generated image $\tilde{x}_t$ at a given time step $t = T - 1, \dots, 0$ to the output image $\tilde{x}_0$ is desired (note, $\tilde{x}_T$ is the initial noise). We use the structural similarity (SSIM) [27] as the distance metrics: $s(t) := \text{SSIM}(\tilde{x}_t, \tilde{x}_0)$. Therefore, the generation speed can be formally defined as the derivative of the SSIM regarding the time step: $v(t) := ds(t)/dt \approx (s(t) - s(t + 1))/\Delta t$, where $\Delta t = 1$. Thereby, we propose our Pattern 1.




FIGURE 5.2. The generation speeds of "mountain" and "peafowl".

PATTERN 1 (Variability in Generation Speed). *Comparing the generation speeds (v_1 and v_2) of two subjects (S_1 and S_2), it can be observed that the outline of the object in the generated image will be taken by S_1 if $v_1 > v_2$. And it can be inferred that S_2 will not be visible in the resulting image.*

We further generates 1,000 images of various classes with the same initial noise and visualize their generation speed in Fig. 5.3 as a violin plot. The SSIM distance from the generated




FIGURE 5.3. A violin plot illustrating the generation speeds of 1,000 images of various classes. The horizontal axis represents the number of steps taken, ranging from 49 to 0, while the vertical axis displays the SSIM scores. The width of each violin represents the number of samples that attained a specific range of SSIM scores at a given step.

images at each step to the final image is calculated. The horizontal axis represents $49 \sim 0$ steps, while the vertical axis represents the SSIM scores. Each violin represents the distribution of the SSIM scores of the 1,000 images in a step, with the width corresponds to the frequency of images reaches the score. In the early stages of generation, the median of the distribution is positioned closer to the minimum value, indicating that a majority of classes exhibit slow generation speeds. However, the presence of a high maximum value suggests the existence of classes that generate relatively quickly. In the middle stages of generation, the median of the distribution gradually increases, positioning itself between the maximum and minimum values. In the later stages of generation, the median of the distribution is positioned closer to the maximum value, indicating that a majority of classes are nearing completion. However, the persistence of a low minimum value suggests the presence of classes that still exhibit slow generation speeds. This analysis highlights the variation in generation speeds across different classes throughout the entire generation process. This phenomenon can be interpreted as a characteristic of the generation process, where different classes exhibit varying speeds throughout the stages. It is possible that certain classes have inherent complexities or dependencies that cause them to generate more slowly. Conversely, other classes may have simpler structures or fewer dependencies, leading to faster generation.




FIGURE 5.4. Images generated by the template "A photo of A and B".




FIGURE 5.5. The first row illustrates the generation process with the prompt "a photo of a **silver salmon**". The second row, based on the forty-second step of the first row, shows the generation process with the prompt "a photo of a **feather**". The third row, also building upon the forty-second step of the first row, presents the generation procedure when the prompt is "a photo of a **magician**". The fourth row depicts the generation process in the presence of feature entanglement. The fifth row demonstrates the generation process for two distinct categories without feature entanglement.

5.1.2 Similarity of Coarse-grained Characteristics

OBSERVATION 2. When a text prompt contains a noun A representing the object to be generated, the introduction of another noun B , which describes an object with similar coarse-grained characteristics to the object represented by noun A , into the prompt, either through addition or replacement, results in the generation of an image that contains an object combining elements of both nouns A and B .

The second case that we observed in our attacks is when two nouns in the text prompt share similar coarse-grained characteristics, the generated image will contain a subject that is a combination of these two nouns. As illustrated in Fig. 5.5, when given the text "a silver salmon and a feather," the GMs generate an image of a feather with the outline of a salmon. This happens because these two nouns (*i.e.*, salmon and feather) exhibit a certain degree of similarity in their coarse-grained attributes. In contrast, there is no feature entanglement between "salmon" and "magician" because their coarse-grained features are vastly different from each other.

To verify this assumption, we first obtain the generated latent variable for the prompt "a photo of a silver salmon" at the early sampling step (*e.g.*, 42 steps). Using this latent variable, we replace the prompt with "a photo of a feather" and continue generating images. The results confirm that feathers can continue to be generated based on the coarse-grained properties of silver salmon, and the final generated graph has high similarity with the generated graph of the prompt "a photo of a silver salmon and a feather". However, replacing "silver salmon" with "magician" does not seem to generate any object similar to "magician". This observation indicates that there is no coarse-grained feature entanglement between these two subjects. We summarize this observation in Pattern 2.

PATTERN 2 (Similarity of Coarse-grained Characteristics). *Let X_t^A and X_t^B denote the latent variables generated by the GM for word tokens A and B , respectively. Suppose t is small and let d represent the metric that measures the outline similarity between two images. If the text prompt contains both A and B , and $d(X_t^A, X_t^B)$ falls below the threshold σ , then feature entanglement occurs in the generated image.*

Based on the observed Pattern 2, The types of feature entanglement can be further divided into, direct entanglement and indirect entanglement. As shown in Fig. 5.4, direct entanglement represents the direct entanglement triggered by two categories of coarse-grained attributes that have global or local similarities. Indirect entanglement is shown in Fig. 5.4d, where the additional attribute trunk brought by the howler monkey has a high similarity with the coarse-grained attribute of the snake, thus triggering the entanglement phenomenon.




FIGURE 5.6. a) "A photo of a **bat**"; b) "A photo of a bat and a **ball**"; c) Heat map of the word "**bat**" in generated image; d) "A photo of a **warthog**"; e) "A photo of a **warthog** and a **traitor**"; f) Heat map of the word "**warthog**" in generated image.

5.1.3 Polysemy of Words

OBSERVATION 3. *When a text prompt contains a noun A representing the object to be generated, if the semantic scope of noun A encompasses multiple distinct objects, the generated image contains one of the objects described by noun A . If there exists another word B that, when combined with noun A , directs its semantics to a particular object, the introduction of word B into the prompt, either through addition or replacement, leads to the generation of an image that specifically contains that particular object.*

The third scenario we observed in the attack is that the content of the generated image is not directly related to either the desired image or the added word. However, this is again different from a category of disappearance, where the desired target has a clear individual in the image. As illustrated in Figs. 5.6a, 5.6b and 5.6c, when the cleaning prompt "a photo of a bat" was modified to "a photo of a bat and a ball", the bat disappeared completely from the generated image, but the DDAM[26] heat map showed that the word "bat" is highly associated with the stick-like object in the newly generated image.

PATTERN 3 (Polysemy of Words). *When interpreting polysemous words, language models must rely on contextual cues to distinguish meanings. However, in some cases, the available*




FIGURE 5.7. a) t-SNE Visualization of 100 images each of "**bat**", "**baseball bat**", "**bat** and **ball**" and text "a photo of a **bat**." b) The boxplot of cosine similarities between the text embedding of "a photo of a **bat**" and 100 of image embeddings each of "**bat**", "**baseball bat**", and "**bat** and **ball**".

contextual information may be insufficient or confuse the model by modifying specific words, resulting in a model-generated image that deviates from the actual intention of the user.

To further investigate the phenomenon of word polysemy in the Stable diffusion model, we used the prompts "a photo of a bat", "a photo of a baseball bat" and "a photo of a bat and a ball" to generate 100 images each using the stable diffusion model, and transformed these images into embedding form by CLIP image encoder, and transformed "a photo of a bat" into embedding form by CLIP text encoder, then visualized these 301 embeddings by t-SNE. As illustrated in Fig. 5.7a, Considering the entire set of bat images, bats (the animal) are closer to the text "a photo of a bat" than baseball bats are, as depicted in the t-SNE visualization. However, the distribution of the two categories also has relatively close proximity, indicating their underlying similarities. The category of "bat and ball" shows a more expansive distribution, almost enveloping the other two. This suggests that by modifying the original text from "a photo of a bat" to "a photo of a bat and a ball", the distribution of the clean text can be pulled towards another meaning in the polysemous nature of the word "bat". From the perspective of text-to-image model, this kind of modification can stimulate the polysemous property of the word, thereby achieving an attack effect.

In addition to this explicit polysemy, our algorithm further demonstrates its aptitude for detecting subtler instances of polysemous words. As depicted in Figure 5.6, the transformative capacity of our algorithm is evident when an image of a warthog (Fig. 5.6d) transfigures into an image of a military chariot (Fig. 5.6e) with the incorporation of the term "traitor".

5.1.4 Positioning of Words

In addition to the three aforementioned observations and patterns outlined in the thesis, there is a fourth observations (Observation 4), which is related to positioning of words.

OBSERVATION 4. When a text prompt contains a noun A representing the object to be generated, there exists a preceding word B and a succeeding word C around noun A. When replacing either word B or C with another noun D, for certain instances of noun A, replacing word B results in the generation of an image containing noun D, while replacing word C



FIGURE 5.8. a) "A type of **footwear** with a thick, rigid sole, often made of wood, and an upper made of leather or another material. **Clogs** can be open or closed, and are commonly associated with **Dutch** and Scandinavian cultures." b) "**footwear**" is replaced by "**pistol**". c) "**Dutch**" is replaced by "**pistol**".



FIGURE 5.9. A template, "A photo of A , B and C ", is used to generate prompts, where $A, B, C \in \{ \text{"cat"}, \text{"pistol"}, \text{"clogs"} \}$. For example, "0-1-2" represents $A = \text{"cat"}$, $B = \text{"pistol"}$ and $C = \text{"clogs"}$, and so on.

still results in the generation of an image containing noun A . Conversely, for other instances of noun A , the opposite scenario occurs.

An example of Pattern 4 is shown in Fig. 5.8. When "footwear" is replaced by "pistol", the generated image contains a pistol instead of clogs. However, when "Dutch" is replaced by "pistol", the model still generates an image of clogs. In addition to differences in the words being replaced, a significant distinction between the two aforementioned examples of success and failure lies in the relative positioning of the word being replaced with respect to the target class word. We hypothesize that this phenomenon occurs due to the different order of the replaced words B or C with respect to the noun A . To exclude the effects of complex contextual structures, a template for a short prompt, "A photo of A , B and C ", is used, and the order of A , B , and C are swapped (Fig. 5.9).

When these sentences with different sequences of category words are understood from a human perspective, they all have basically the same semantics: both describe a picture containing a cat, clogs, and a pistol. However, in the processing of language models (including CLIP), the order of words may affect their comprehension. Although positional encoding provides

the model with the relative positions of words, the model may associate different orders with different semantics through learned patterns. Therefore, we propose our Pattern 4.

PATTERN 4 (Positioning of Words). *Let $\mathcal{V}$ denote a set of vocabulary. Let $\mathcal{N} \subset \mathcal{V}$ denote the subset of all nouns in the vocabulary. Consider a text prompt containing noun $A \in \mathcal{N}$ representing the object to be generated. Furthermore, assume there exist preceding word $B \in \mathcal{V}$ and succeeding word $C \in \mathcal{V}$ surrounding noun A . There exists a condition-dependent behavior regarding the replacement of words B and C with another noun $D \in \mathcal{N}$:*

$$\left\{ \begin{array}{l} \exists A, D \in \mathcal{N}, \quad \exists B, C \in \mathcal{V}, \quad P(B \rightarrow D) \xrightarrow{\text{generate}} D \quad \wedge \quad P(C \rightarrow D) \xrightarrow{\text{generate}} A; \\ \exists A, D \in \mathcal{N}, \quad \exists B, C \in \mathcal{V}, \quad P(B \rightarrow D) \xrightarrow{\text{generate}} A \quad \wedge \quad P(C \rightarrow D) \xrightarrow{\text{generate}} D. \end{array} \right.$$

CHAPTER 6

Conclusion

The realm of text-to-image generation has observed a remarkable evolution over recent years, while concurrently exposing several vulnerabilities that require further exploration. Despite the many advancements, there are key limitations, specifically concerning the stability and reliability of generative models, which remain to be addressed. This thesis has introduced Auto-attack on Text-to-image Models (ATM), a novel approach that generates a plethora of successful attack prompts, providing an efficient tool for probing vulnerabilities in text-to-image models. ATM not only identifies a broader array of attack patterns but also facilitates a comprehensive examination of the root causes. We believe that our proposed method will inspire the research community to shift their focus toward the vulnerabilities of present-day text-to-image models, stimulating further exploration of both attack and defensive strategies. This process will be pivotal in advancing the security mechanisms within the industry and contributing to the development of more robust and reliable systems for generating images from textual descriptions.

Acknowledgements

I want to extend my heartfelt thanks to everyone who played a part in the successful completion of my research journey. Their unwavering support and mentorship have enriched my path with knowledge and personal development.

Firstly, I'd like to extend my profound gratitude to my supervisor, Chang Xu. During times when I felt uncertain about my path, he offered me the chance to dive into the world of academia. Since my enrollment, he has dedicated countless hours and immense effort to guide and support me professionally. From Chang, I've not only acquired academic insights but also the courage to navigate life and my scholarly pursuits.

Secondly, I'd like to express my gratitude to my research collaborators, Yanxi Li, Daochang Liu, and Yunke Wang. Their invaluable guidance and recommendations were essential to my work.

I also owe a deep debt of gratitude to my family: my father, Zhaohua Du, my mother, Yulan Chen; and my sister, Xiangyun Du. They offer emotional and financial support. Their unwavering encouragement and understanding have been the foundation upon which I've built my journey.

I'd like to extend my heartfelt appreciation to each and every one of my colleagues who have contributed immensely to my journey. Special thanks go to Linwei Tao, Yixiong Ding, Xiyu Wang, Chuyang Zhou, Jiajun Huang, Huihui Gong, Minjing Dong, Xiu Su, Shuyi Jiang, Anh-Dung Dinh, Xiaohuan Pei, Yanxiang Ma, Siyu Xu, Tao Huang, Jianyuan Guo, Mingkao Zheng, Chen Chen, Junyu Zhang and Jinxu Lin. Their collaboration, insights, and camaraderie have been invaluable.

I'd like to convey my profound gratitude to my circle of friends who have been instrumental in my journey. My heartfelt thanks go to Zekun Liu, Ying Xu, Sihan Zhang, Chenwu Yan, Xuan Chen, Qi Xu, Guobei Zhuang, and others. Their support, encouragement, and camaraderie during this period have been indispensable and have truly made a difference in my life.

Lastly, I want to thank my university, the University of Sydney, for providing me with this platform.

Bibliography

- [1] H. Chefer, Y. Alaluf, Y. Vinker, L. Wolf, and D. Cohen-Or. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. *arXiv preprint arXiv:2301.13826*, 2023.
- [2] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [3] P. Esser, J. Chiu, P. Atighehchian, J. Granskog, and A. Germanidis. Structure and content-guided video synthesis with diffusion models. *arXiv preprint arXiv:2302.03011*, 2023.
- [4] W. Feng, X. He, T.-J. Fu, V. Jampani, A. Akula, P. Narayana, S. Basu, X. E. Wang, and W. Y. Wang. Training-free structured diffusion guidance for compositional text-to-image synthesis. *arXiv preprint arXiv:2212.05032*, 2022.
- [5] C. Guo, A. Sablayrolles, H. Jégou, and D. Kiela. Gradient-based adversarial attacks against text transformers. *arXiv preprint arXiv:2104.13733*, 2021.
- [6] M. Heusel, H. Ramsauer, T. Unterthiner, B. Nessler, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [7] J. Ho, A. Jain, and P. Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems*, 33:6840–6851, 2020.
- [8] L. Huang, D. Chen, Y. Liu, Y. Shen, D. Zhao, and J. Zhou. Composer: Creative and controllable image synthesis with composable conditions. *arXiv preprint arXiv:2302.09778*, 2023.
- [9] L. Li, R. Ma, Q. Guo, X. Xue, and X. Qiu. Bert-attack: Adversarial attack against bert using bert. *arXiv preprint arXiv:2004.09984*, 2020.

- [10] C. Lu, Y. Zhou, F. Bao, J. Chen, C. Li, and J. Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. In *Advances in Neural Information Processing Systems*, 2022.
- [11] A. Madry, A. Makelov, L. Schmidt, D. Tsipras, and A. Vladu. Towards deep learning models resistant to adversarial attacks. *arXiv preprint arXiv:1706.06083*, 2017.
- [12] Midjourney. Midjourney. <https://www.midjourney.com/>, 2023. Accessed on 2023-05-17.
- [13] C. Mou, X. Wang, L. Xie, J. Zhang, Z. Qi, Y. Shan, and X. Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. *arXiv preprint arXiv:2302.08453*, 2023.
- [14] A. Nichol, P. Dhariwal, A. Ramesh, P. Shyam, P. Mishkin, B. McGrew, I. Sutskever, and M. Chen. Glide: Towards photorealistic image generation and editing with text-guided diffusion models. *arXiv preprint arXiv:2112.10741*, 2021.
- [15] A. Obukhov, M. Seitzer, P.-W. Wu, S. Zhydenko, J. Kyl, and E. Y.-J. Lin. High-fidelity performance metrics for generative models in pytorch, 2020. URL <https://github.com/toshas/torch-fidelity>. Version: 0.3.0, DOI: 10.5281/zenodo.4957738.
- [16] OpenAI. Dall-e 2. <https://openai.com/product/dall-e-2>, 2023. Accessed on 2023-05-17.
- [17] OpenAI. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [18] A. Radford, J. W. Kim, C. Hallacy, A. Ramesh, G. Goh, S. Agarwal, G. Sastry, A. Askell, P. Mishkin, J. Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021.
- [19] A. Ramesh, M. Pavlov, G. Goh, S. Gray, C. Voss, A. Radford, M. Chen, and I. Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning*, pages 8821–8831. PMLR, 2021.
- [20] A. Ramesh, P. Dhariwal, A. Nichol, C. Chu, and M. Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [21] R. Rombach, A. Blattmann, D. Lorenz, P. Esser, and B. Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on*

- Computer Vision and Pattern Recognition*, pages 10684–10695, 2022.
- [22] L. Ruan, Y. Ma, H. Yang, H. He, B. Liu, J. Fu, N. J. Yuan, Q. Jin, and B. Guo. Mm-diffusion: Learning multi-modal diffusion models for joint audio and video generation. *arXiv preprint arXiv:2212.09478*, 2022.
- [23] J. Sohl-Dickstein, E. Weiss, N. Maheswaranathan, and S. Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pages 2256–2265. PMLR, 2015.
- [24] J. Song, C. Meng, and S. Ermon. Denoising diffusion implicit models. In *International Conference on Learning Representations*, 2021.
- [25] Stability-AI. Stable diffusion public release. <https://stability.ai/blog/stable-diffusion-public-release>, 2023. Accessed on 2023-05-17.
- [26] R. Tang, A. Pandey, Z. Jiang, G. Yang, K. Kumar, J. Lin, and F. Ture. What the daam: Interpreting stable diffusion using cross attention. *arXiv preprint arXiv:2210.04885*, 2022.
- [27] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- [28] L. Zhang and M. Agrawala. Adding conditional control to text-to-image diffusion models. *arXiv preprint arXiv:2302.05543*, 2023.
- [29] T. Zhang, V. Kishore, F. Wu, K. Q. Weinberger, and Y. Artzi. Bertscore: Evaluating text generation with bert. In *International Conference on Learning Representations*, 2020.