

# Textualize Visual Prompt for Image Editing via Diffusion Bridge

Pengcheng Xu<sup>1,2</sup>, Qingnan Fan<sup>2</sup>, Fei Kou<sup>2</sup>, Shuai Qin<sup>2</sup>, Hong Gu<sup>2</sup>, Ruoyu Zhao<sup>2,3</sup>,  
Charles Ling<sup>1</sup>, Boyu Wang<sup>1\*</sup>

<sup>1</sup> Western University

<sup>2</sup> VIVO

<sup>3</sup> Xidian University

{pxu67, charles.ling}@uwo.ca, fqchina@gmail.com, koufei@hotmail.com, royzhao@stu.xidian.edu.cn,  
{guhong, shuai.qin}@vivo.com, bwang@csd.uwo.ca,

## Abstract

Visual prompt, a pair of before-and-after edited images, can convey indescribable imagery transformations and prosper in image editing. However, current visual prompt methods rely on a pretrained text-guided image-to-image generative model that requires a triplet of text, before, and after images for retraining over a text-to-image model. Such crafting triplets and retraining processes limit the scalability and generalization of editing. In this paper, we present a framework based on any single text-to-image model without reliance on the explicit image-to-image model thus enhancing the generalizability and scalability. Specifically, by leveraging the probability-flow ordinary equation, we construct a diffusion bridge to transfer the distribution between before-and-after images under the text guidance. By optimizing the text via the bridge, the framework adaptively textualizes the editing transformation conveyed by visual prompts into text embeddings without other models. Meanwhile, we introduce differential attention control during text optimization, which disentangles the text embedding from the invariance of the before-and-after images and makes it solely capture the delicate transformation and generalize to edit various images. Experiments on real images validate competitive results on the generalization, contextual coherence, and high fidelity for delicate editing with just one image pair as the visual prompt.

**Project page** — [pengchengpcx.github.io/TextVDB/](https://pengchengpcx.github.io/TextVDB/)

## 1 Introduction

Prompting (Touvron et al. 2023; Floridi and Chiriatti 2020; Liu et al. 2023), providing specific instructions or context for the model, is an effective and emergent tool to guide the large-scale text-to-image (T2I) models to generate (Rombach et al. 2022; Podell et al. 2024; Saharia et al. 2022; Ramesh et al. 2022) or edit (Hertz et al. 2023; Kawar et al. 2023; Parmar et al. 2023; Mokady et al. 2023; Gal et al. 2023) remarkable images. However, in the visual task, not all aspects can be accurately and comprehensively described through language alone. For instance, how to explain the tone transformation between photographs or the personal repaint of a painting? These imagery transformations are abstract, complex, and difficult to convey via a text prompt. In such cases, a pair

of before-and-after images, serving as the visual prompt, are more precise and expressive to represent the imagery transformation and guide the image generation and editing. This motivates us to explore the visual prompt for image editing: *given a pair of before-and-after images as a visual prompt, how to distill the transformation from the visual prompt into text embedding for the text-guided image editing.*

Early on, visual prompt for image editing is employed with in-context learning, primarily effective on classical tasks (e.g., segmentation, detection) (Bar et al. 2022; Wang et al. 2024). Recently, it is explored within diffusion models for arbitrary visual content editing (Nguyen et al. 2024; Motamed, Paudel, and Van Gool 2025; Cheng et al. 2023; Yang et al. 2024; Gu et al. 2024), and there are mainly three types of approaches.

The first category relies on the text-guided image-to-image (TI2I) model (e.g., InstructPix2Pix) (Brooks, Holynski, and Efros 2023). Such a TI2I model has an explicit text-guided image-to-image transformation process. By directly optimizing the text condition that guides the model to manipulate and change the before-image to the after-image, the model converts the transformation between the image pair into text (Gal et al. 2023; Nguyen et al. 2024). However, training such a TI2I model needs to construct a data triplet of text, and before-and-after images. More importantly, such training data is constructed based on the T2I model. For more delicate editing, crafting such triplets requires meticulous effort, limiting the scalability of the training data in the TI2I model (Zhang et al. 2024). Similarly, the second type in-context learning methods (Gu et al. 2024; Yang et al. 2024) leverages the diffusion inpainting model which also requires extra complex data (mask and cropped images) for more training to get good generalization ability. This contrasts with the general T2I model, which utilizes a simpler tuple of images and text for training, making it easier to collect and scale up. The third category adopts textual inversion (Gal et al. 2023; Motamed, Paudel, and Van Gool 2025; Šubrtoová et al. 2023) to invert the before-and-after images into text embeddings sequentially and learn the transformation or concept by comparison of two embeddings. However, textual inversion is inferior for image reconstruction and cannot capture all image details, making it challenging to learn delicate editing transformations and resulting in a coarse concept. Thus, these designs limit the model to learning precise transformations and generalizing to edit high-fidelity images.

\* Corresponding author

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

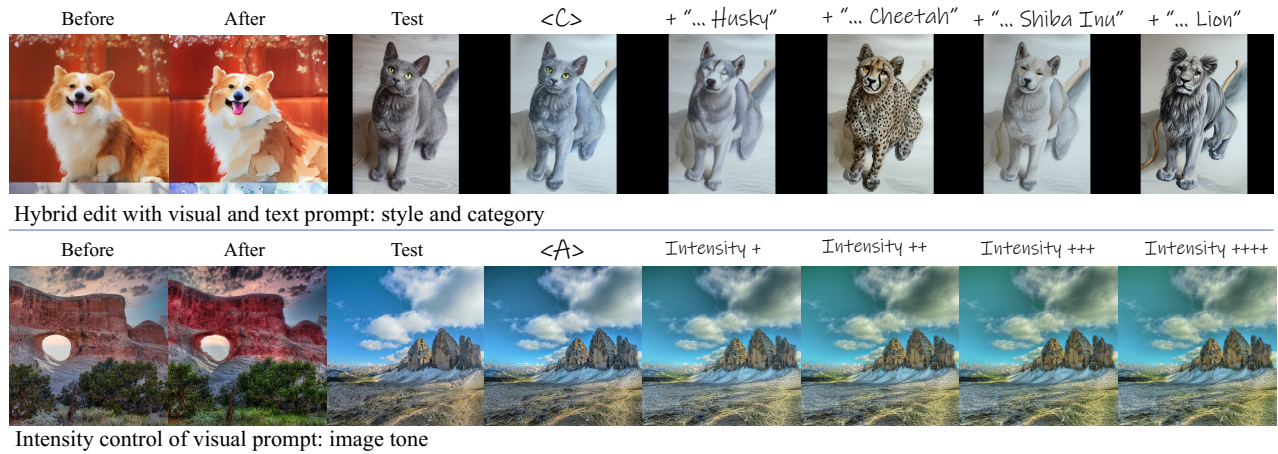


Figure 1: **Image editing via visual prompt.** The visual prompt defines the visual transformation, which is difficult to describe accurately by language, by a before-and-after image pair. Our method learns such delicate transformation into pseudo text (<A> and <C>), supports hybrid editing with natural text, and can control the intensity of editing with rigorous consistency.

To address the aforementioned issues, in this paper, we aim to answer the following questions: **why and how to use a general T2I model to learn delicate imagery transformation from visual prompts.** For the why, considering the training data of the TI2I model is mainly derived from the T2I model and manually crafting large-scale triplet (or more complex) data is cumbersome, visual prompts learned directly based on the general T2I model can be more scalable and generalized. For the how, there are two challenges: 1) Construct and textualize the image-to-image (I2I) process but with the single T2I model. The T2I model only has text-to-image mapping but lacks explicit I2I mapping. We have to construct the I2I translation process for the image pair and distill the process into text embedding used for later editing. 2) Capture imagery details for delicate transformation. The constructed I2I process should encode all details of visual prompts otherwise the learned transformation is ambiguous.

To solve these challenges, we propose to textualize the visual prompt based on the diffusion bridge. Firstly, we use a single T2I model to build a diffusion bridge (Zhou et al. 2024; Su et al. 2023), that transforms the distribution of the before-image to the after-image, to represent the I2I process. As shown in Fig. 2, the bridge is built on a single pretrained T2I model by leveraging its unconditional (null text) and conditional (text) generating abilities. Based on the probability-flow ordinary equation, the before-image is first transformed into a deterministic latent with the unconditional model and then regenerated to the after-image with the text condition. Such a design breaks the dependency on the explicit I2I process and thus requires no retraining of any TI2I model. Besides, leveraging the more general T2I model, the method adaptively inverts different visual prompts into different text embeddings, supporting generalized, hybrid, and high-fidelity editing.

Secondly, to make the text embedding focus on capturing transformation details, and generalize to edit various images, we propose differential attention control to preserve detailed

contents of the before-image during training. This module injects attention weights of the before-image while supporting backpropagation for optimizing text embedding. By learning with attention from the before image, the inverted text embedding is disentangled from the invariance of the before-image, which is irrelevant to the transformation, focuses on capturing fine-grained transformations, and generalizes to edit various images. We summarize our contributions and findings:

- We design a framework to textualize visual prompts via the T2I diffusion bridge for image editing. The method does not require both image and text conditions as the TI2I model that needs triplets of training data and retraining, which improves the scalability and generalization.
- We introduce differential attention control for optimizing the text embedding, learning delicate imagery transformations, and generalizing to edit various types of images.
- Experiments with various evaluation metrics validate our method can learn delicate imagery transformation and generalize better to edit faithful and high-fidelity images, compared to existing methods.

## 2 Related Work

**Text guided image-to-image models.** To make the text-to-image (T2I) models control the spatial content of the synthesized images, the text-guided image-to-image model (TI2I) adds an extra image condition based on the T2I model. ControlNet (Zhang, Rao, and Agrawala 2023; Zhao et al. 2024) and Adapter (Mou et al. 2024; Ye et al. 2023) methods create an additional branch to the UNet to introduce the conditions for generation. Another type of approach combines the embedding of the image with the latent in the diffusion and retrains the model to have both image and text conditions such as InstructPix2Pix (Brooks, Holynski, and Efros 2023), InstructDiffusion (Brooks, Holynski, and Efros 2023), and MagicBrush (Zhang et al. 2024). While these TI2I models support human-intuitive text for image editing, collecting

image pairs strictly aligned with editing texts is not as easy as simply collecting tuples of the image and its text description, and retraining is inefficient, which limits the capacity of these models. In contrast, our method is directly built on the T2I model for its large text-to-image prior that can be easily scaled up with the increasing scale of pairs of text and image.

**Image editing via the visual prompt.** Visual prompting was initially proposed in NLP (Brown et al. 2020) based on in-context learning and recently introduced to computer vision (Bar et al. 2022; Wang et al. 2023a,b). ImageBrush (Yang et al. 2024) and Analogist (Gu et al. 2024) formulate the editing as inpainting in analogy to the image pair but it requires a retrained inpainting model and is limited in image resolutions and encoding intricate details. Image analogy (Šubrtová et al. 2023; Hertzmann et al. 2023) adapts the analogous relation of image pair to new images but the learned editing relation is coarse and ambiguous and cannot accurately edit or preserve structures. Recent methods extract the image transformation from image pairs into text embeddings for text-guided editing. VII (Nguyen et al. 2024) leverages the pretrained TI2I model, InstructPix2Pix, to distill the image-to-image transformation between the image pair into text embeddings. Lego (Motamed, Paudel, and Van Gool 2025) and DIA (Hertzmann et al. 2023) uses textual inversion to learn the disentangled concept by comparing the image pair. However, textual inversion is inferior in image reconstruction and cannot encode all imagery details, which causes the learned concept to be coarse and inaccurate. In contrast, our method accurately recovers the image and textualizes the visual prompts without dependence on the pretrained TI2I model, and produces generalized and delicate editing results.

### 3 Methodology

We aim to learn the delicate imagery transformation from the visual prompt and use it for generalized and high-fidelity image editing. In the following, we first briefly present the score-based generative model (Song et al. 2021) and the diffusion bridge (Song, Meng, and Ermon 2021; Su et al. 2023) in Sec 3.1. Then, we discuss converting the visual prompt into text embedding via diffusion bridge in Sec 3.2. In Sec 3.3, we introduce a differential attention control strategy during text optimization, which helps learn disentangled and generalized text guidance. Last, the full algorithm is in Sec 3.4.

#### 3.1 Preliminaries

**Score-based generative model.** Song et al. (2021) proposed the unified diffusion framework with Stochastic Differential Equations (SDE) and further showed that any diffusion process can be represented by a deterministic probability flow (PF) ordinary differential equation (ODE) that transfers the data encodings to *deterministic* and *unique* latent encodings. By solving the PF ODE in Eq. 2 forward and backward, the data  $\mathbf{x}_0$  and latent  $\mathbf{x}_T$  are transferred between each other.

$$d\mathbf{x}_t = \boldsymbol{\mu}(\mathbf{x}_t, t) dt + \sigma(t) d\mathbf{w}_t \quad (1)$$

$$d\mathbf{x}_t = \left[ \boldsymbol{\mu}(\mathbf{x}_t, t) - \frac{1}{2} \sigma(t)^2 \nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t) \right] dt \quad (2)$$

Here,  $\mathbf{w}_t$  is the standard Wiener process,  $\boldsymbol{\mu}(t)$  and  $\sigma(t)$  are drift and diffusion coefficients respectively. The score

function  $\nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t)$  is approximated with neural network  $\mathbf{s}_{\theta}(\mathbf{x}_t, t)$ . For the T2I model, the text condition  $\mathbf{c}$  is added to replace the marginal score function  $\nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t)$  with the conditional score function  $\nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t | \mathbf{c})$  via classifier-free guidance (Ho and Salimans 2022) in Eq. 3, where  $w$  is the guidance scale, and  $\emptyset$  represents the null text.

$$\nabla_{\mathbf{x}} \log p_t(\mathbf{x}_t | \mathbf{c}) \approx \tilde{\mathbf{s}}_{\theta}(\mathbf{x}_t, t, \mathbf{c}) = \mathbf{s}_{\theta}(\mathbf{x}_t, t, \emptyset) + w \cdot (\mathbf{s}_{\theta}(\mathbf{x}_t, t, \mathbf{c}) - \mathbf{s}_{\theta}(\mathbf{x}_t, t, \emptyset)) \quad (3)$$

**Dual diffusion implicit bridges (DDIB).** Following the above PF ODE in Eq. 2, DDIB (Su et al. 2023) uses two diffusion models  $\mathbf{s}_{\theta}^{(s)}$  and  $\mathbf{s}_{\theta}^{(t)}$  trained on separated (source and target) domains to perform I2I transformation. Concretely, as shown in Eq. 5, it first uses an ODE solver in Eq. 6 to forward the source image  $\mathbf{x}^{(s)}$  to the latent  $\mathbf{x}^{(l)}$ , and then reverse the latent to the target image  $\mathbf{x}^{(t)}$ . Such a translation process has been theoretically proven to be the most efficient optimal transport between two distributions (Su et al. 2023).

$$\mathbf{x}^{(l)} = \text{ODESolve}(\mathbf{x}^{(s)}; \mathbf{s}_{\theta}^{(s)}, 0, 1) \quad (4)$$

$$\mathbf{x}^{(t)} = \text{ODESolve}(\mathbf{x}^{(l)}; \mathbf{s}_{\theta}^{(t)}, 1, 0) \quad (5)$$

$$\text{ODESolve}(\mathbf{x}_t; \mathbf{s}_{\theta}, t_0, t_1) = \mathbf{x}_{t_0} + \int_{t_0}^{t_1} \mathbf{s}_{\theta}(\mathbf{x}_t, t) dt \quad (6)$$

Our method is driven and designed by such theoretical advantage over transport efficiency. However, instead of using two different models trained on two different domains, we only use a single T2I model and leverage the text condition to transfer to versatile domains.

#### 3.2 Visual Prompt Learning via Diffusion Bridge

**T2I diffusion bridge.** To learn the transformation represented by the visual prompt on the more general and scalable T2I model, we need to construct the I2I process (before-to-after image transformation) based on a single T2I model. Inspired by the DDIB, we build a diffusion bridge to represent the I2I transformation. Specifically, we use unconditional (null text  $\emptyset$ ) and conditional (text  $\mathbf{c}$ ) models to replace the two different diffusions in DDIB, and adopt DDIM as the ODE solver. This can be implemented with a single pretrained T2I model and require no retraining of any model.

Concretely, let  $\mathbf{x}_0^b \sim p(\mathbf{x}^b)$  and  $\mathbf{x}_0^a \sim p(\mathbf{x}^a)$  represent the before-and-after images, respectively. The intermediate latent is  $\mathbf{x}_T \sim \mathcal{N}(0, I)$ . Our diffusion bridge is defined in Eq. 7 and 8 in which the before-image  $\mathbf{x}_0^b$  is first transformed into  $\mathbf{x}_T$  with null text  $\emptyset$  in  $T$  steps and then regenerated to  $\mathbf{x}_0^a$  under the guidance of text embedding  $\mathbf{c}$ . Such a process model the distribution transition of  $p(\mathbf{x}^b) \rightarrow p(\mathbf{x}^a)$ . We aim to learn  $\mathbf{c}$  that can guide any samples drawn from  $p(\mathbf{x}^b)$  to  $p(\mathbf{x}^a)$  in a few-shot manner.

$$\mathbf{x}_0^b \rightarrow \mathbf{x}_T : \mathbf{x}_T = \text{DDIM}(\mathbf{x}_0^b; \tilde{\mathbf{s}}_{\theta}(\mathbf{x}_t, t, \emptyset), 0, T) \quad (7)$$

$$\mathbf{x}_T \rightarrow \mathbf{x}_0^a : \mathbf{x}_0^a = \text{DDIM}(\mathbf{x}_T; \tilde{\mathbf{s}}_{\theta}(\mathbf{x}_t, t, \mathbf{c}), T, 0) \quad (8)$$

Note that the DDIM in Eq. 7 and 8 is deterministic which means that once the prior model  $\tilde{\mathbf{s}}_{\theta}$  is determined, for any  $\mathbf{x}_0^b \rightarrow \mathbf{x}_T$  in Eq. 7,  $\mathbf{x}_T$  is *deterministic* and *unique*. The

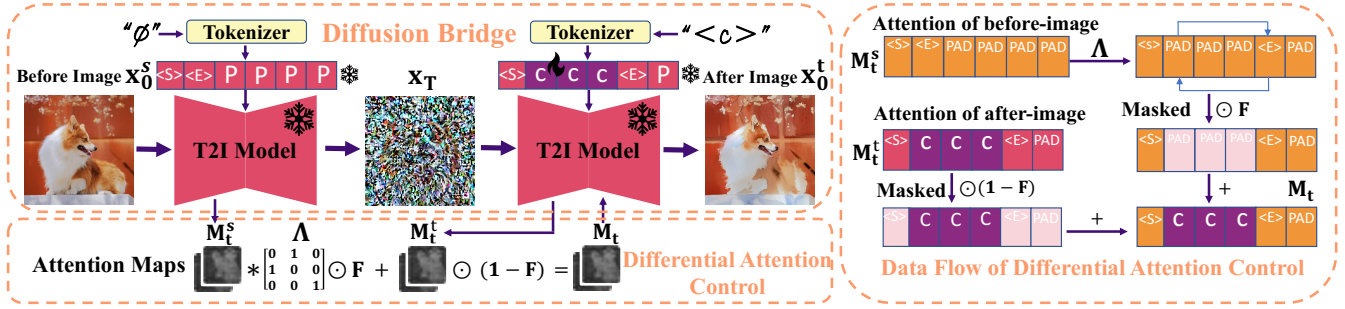


Figure 2: **Textualization of the diffusion bridge.** **Left:** The before-image is first transferred to a deterministic latent encoding via the unconditional model and then to the after-image under the text guidance. The text embeddings are optimized with fixed start (latent  $x_T$ ) and end (after-image  $x_0^a$ ) states. **Right:** In training, the attention of the before-image  $M_t^b$  is first timed with the column-transformation matrix  $\Lambda$  to switch the column of  $\langle \text{end} \rangle$  token, then masked with  $F$ . The attention of the after-image  $M_t^a$  is masked with  $1 - F$  to get the attention of the  $y$  tokens. The final  $M_t$  is the addition of two masked attentions. This preserves the linguistic format of the cross-attention and enables the embedding to learn disentangled and generalized visual transformation.

same applies to  $x_T \rightarrow x_0^a$ . This forms a unique one-to-one mapping for every pair of before-and-after images, which enables high-fidelity and stable image editing.

However, such a property makes the widely used textual inversion (Gal et al. 2023) invalid. This is because when optimizing the text embedding  $c$ , textual inversion adds random noise to the image at every step, and the final latent  $x_T$  will not be deterministic anymore, which violates our constraint for  $x_T$ . Essentially, the learning process of textual inversion tries to map a single image  $x_0^a$  to the whole latent Gaussian distribution rather than the fixed noise vector  $x_T$ . Apart from this, this stochastic procedure is inferior in image reconstruction and results in the learned text embeddings being unable to capture and recover all imagery details (Mokady et al. 2023). Mostly, textual inversion is effective in learning objects but not the whole image and details, and always introduces significant changes. As a result, it is not suitable for image editing (Nguyen et al. 2024). This motivates us to design the optimization procedure of the text embedding  $c$  with fixed start  $x_T$  and end  $x_0^a$  states.

**Optimization conditioned on both start and end states of diffusion.** The text embedding  $c$  is optimized to satisfy the diffusion process  $x_T \rightarrow x_0^a$  given the start and end states  $\{x_T, x_0^a\}$ . Such textualization based on the implicit bridge is conditioned on both the start and end states. Thus, we optimize the text embedding  $c$  to maximize the conditional probability  $p_{c,\theta}(x_0^a|x_T)$  which is parameterized by the T2I model. This is different from the general diffusion process in textual inversion (Gal et al. 2023) whose end state  $x_T$  is an *unconstrained* Gaussian noise. We discuss the optimization design of our textualization as follows. With the deterministic DDIM, the process in Eq. 8 for each timestep  $t$  is written as:

$$x_{t-1} = \sqrt{\alpha_{t-1}}f_\theta(x_t, t, c) + \sqrt{1 - \alpha_{t-1}}\tilde{s}_\theta(x_t, t, c) \quad (9)$$

$$f_\theta(x_t, t, c) = \frac{x_t - \sqrt{1 - \alpha_t}\tilde{s}_\theta(x_t, t, c)}{\sqrt{\alpha_t}} \quad (10)$$

Considering that the start and end states  $\{x_T, x_0^a\}$  are given and fixed and the DDIM is deterministic, one intuitive way to optimize  $c$  is to first do forward pass  $x_T \rightarrow x_0^a$  to get

the final output image  $\hat{x}_0$  with Eq. 9 and calculate the loss with  $\|\hat{x}_0 - x_0^a\|_2$ . However, this recurrent backpropagation needs to cache the intermediate result of  $s_\theta$  at each timestep  $t$ , which is not feasible for multiple steps in diffusion.

To cope with this, we propose to optimize the *predicted*  $x_0$ ,  $f_\theta(x_t, t)$ , with the ground-truth after-image  $x_0^a$ . The predicted  $x_0$ ,  $f_\theta(x_t, t)$  is an estimator of the final output at the current timestep, which indicates how far away the current result is from the desired output (Song, Meng, and Ermon 2021). Besides, we find that the losses in the initial steps are much larger than those in the final steps. Optimizing the loss at each timestep equally may cause  $\hat{x}_0$  to deviate from  $x_0^a$  at the final step. To make the last step  $\hat{x}_0 \approx x_0^a$ , we scale down the loss with time-aware scaling function  $\beta(t)$  and the optimization at each timestep is:

$$\mathcal{L}(c, t) = \beta(t)\|f_\theta(x_t, t, c) - x_0^a\|_2 \quad (11)$$

### 3.3 Learning with Differential Attention Control

Attention injection can preserve the invariance of before-and-after editing, which enables saving the before-image content in image editing. It is generally used in *inference* but not in *training* the T2I model. Besides, current attention-based image editing methods (Hertz et al. 2023; Cao et al. 2023; Tumanyan et al. 2023) are not straightforward differential and do not support backpropagation.

**Learning motivation.** We introduce attention injection during *training* the text embedding  $c$ . Intuitively, the injected attention introduces the information from the before-image and thus disentangles the learned text embedding  $c$  from the before-image information. This makes the text embedding  $c$  solely learn the transformation and generalize to edit various images. Otherwise, the information irrelevant to the transformation from the before-image leaks to  $c$  such that it only fits on the before-image but cannot generalize on other images. In summary, our motivation for training with attention injection has two aspects: 1) In training, leveraging the injected attention capturing the invariance between the before-and-after images, the text embedding learns a visual transformation disentangled from a specific before-image and can be more

---

**Algorithm 1: Textualize Visual Prompt**

---

```
1: Input: An image pair  $\{\mathbf{x}_0^b, \mathbf{x}_0^a\}$ 
2:   T2I diffusion model  $\mathbf{s}_\theta$ 
3:   Number of training epochs  $N$ ; Number of diffusion steps  $T$ 
4:   Learning rate  $\gamma$ ; Attention injection timestamp  $\tau$ 
5:   Initialize  $\mathbf{c}$ 
6: for  $i = 1, \dots, N$  do
7:    $\mathbf{x}_T = \text{DDIM}(\mathbf{x}_0^b; \tilde{\mathbf{s}}_\theta(\mathbf{x}_t, t, \emptyset), 0, T)$ 
8:   for  $t = T, \dots, 1$  do
9:     if  $t < \tau$  then
10:      Build column-transformation matrix  $\mathbf{\Lambda}$  and mask  $\mathbf{F}$ 
11:      Get  $M_t$  with attention injection with Eq. 12
12:     end if
13:     Calculate predicted  $\mathbf{x}_0$ ,  $\mathbf{f}_\theta(\mathbf{x}_t, t, M_t)$  with Eq. 10
14:     Calculate next state  $\mathbf{x}_{t-1}$  with Eq. 9
15:     Calculate  $\mathcal{L}(\mathbf{c}, t) = \beta(t) \|\mathbf{f}_\theta(\mathbf{x}_t, t, \mathbf{c}, M_t) - \mathbf{x}_0^a\|_2$  with Eq. 13
16:     Update  $\mathbf{c} = \mathbf{c} - \gamma \nabla \mathcal{L}(\mathbf{c}, t)$ 
17:   end for
18: end for
19: Output:  $\mathbf{c}$ 
```

---

generalized. 2) In inference, the injected attention preserves the invariance of the before-image and achieves high fidelity. **Module design.** To make the attention injection differential and support backpropagation, we implement the attention injection with only multiplication and addition. The detailed process is depicted in Fig. 2. For simplicity, let  $M_t^b, M_t^a \in \mathbb{R}^{j \times k}$  denote the attention weights from before- and after-images at timestep  $t$ , respectively, where  $j$  is the feature dimension and  $k$  is the length of total tokens. The text prompt  $\mathbf{c}$  of the after-image includes  $y$  tokens for learning the transformation. So, we keep the attentions of  $y$  tokens in  $M_t^a$  and replace the rest  $k - y$  with attentions from  $M_t^b$ . Besides, we make the  $(y + 1)^{\text{th}}$  attention in  $M_t^a$  to be the  $\langle \text{end} \rangle$  token attention of  $M_t^b$  so that the new attention  $M_t$  also follows the linguistic format determined by the text encoder. We use a column transformation matrix  $\mathbf{\Lambda}$  and mask  $\mathbf{F} = [1, \dots, 0, 1, 1] \in \{0, 1\}^k$  to achieve this. Specifically, the column-transformation matrix is a modified identity matrix that switches the columns of 1 and  $y + 1$ . It shifts the cross-attention of the  $\langle \text{end} \rangle$  token in  $M_t^b$  to the  $(y + 1)^{\text{th}}$  column, and the mask  $\mathbf{F}$  injects all cross-attention weights of  $M_t^b$  that are not in the position of tokens of  $\mathbf{c}$ , into  $M_t$ . For self-attention, we do not include  $\mathbf{\Lambda}$  since self-attention is not related to text. Consequently, the training objective in Eq. 11 is added with the new attention and becomes Eq. 13.  $\odot$  is column-multiplication that multiplies each column in the attention matrix with a scalar value of 0 or 1.

$$M_t = M_t^b \mathbf{\Lambda} \odot \mathbf{F} + M_t^a \odot (1 - \mathbf{F}) \quad (12)$$

$$\mathcal{L}(\mathbf{c}, t) = \beta(t) \|\mathbf{f}_\theta(\mathbf{x}_t, t, \mathbf{c}, M_t) - \mathbf{x}_0^a\|_2 \quad (13)$$

### 3.4 Visual Prompt for Image Editing

Finally, by constructing a diffusion bridge with the T2I model, we distill the I2I transformation defined by an image pair  $\mathbf{x}_0^b$  and  $\mathbf{x}_0^a$  into a text embedding  $\mathbf{c}$ . Our framework is in Algorithm 1. In inference, the new image goes through the diffusion bridge as the before-image under the guidance of

the learned text embedding. The output is the edited image with the desired transformation. Our method also adapts to multiple image pairs with the same transformation. For each epoch, we randomly select an image pair for training.

## 4 Evaluation on Real Images

### 4.1 Experiment setup

**Baselines and implementations.** The closest work to ours is visual instruction inversion (VII) (Nguyen et al. 2024) which also learns image transformation by an image pair. VII is based on the pretrained TI2I model, InstructPix2Pix (InsP2P) (Brooks, Holynski, and Efros 2023). We also add InsP2P with ground-truth text instructions, which are actually unavailable. Additionally, we add diffusion image analogy DIA (Šubrtová et al. 2023) that transfers the analogy relation between the image pair to new images, and visual in-context learning Analogist (Gu et al. 2024) that inpainting the edited image with in-context learning, and style transfer methods DiffuseIT (Kwon and Ye 2023) that transfer the style of after-image to the new image. We implement our method based on SD v1.5, consistent with VII and InsP2P. All baselines are evaluated with their official codebases.

**Datasets.** We evaluate our framework on *real images* based on the two latest benchmarks Dreambooth (Ruiz et al. 2023) and PIE (Ju et al. 2024). Following the procedure of creating image pairs in InsP2P and VII, we also use P2P (Hertz et al. 2023) to generate different image pairs with different text instructions as visual prompts. We use this dataset to evaluate the generalization and fidelity of our method since the data is not included in the training set of SD v1.5 or InsP2P. The details of implementation and dataset are in the Appendix D.

### 4.2 Comparison with previous methods

**Real image editing and fidelity.** We evaluate our method in real image editing. As shown in Fig. 3, our method produces higher fidelity and preserves better invariant details of the original image under different editing types. This indicates that our method can learn more accurate and generalized editing effects from the image pair. In contrast, style transfer (DiffuseIT) and image analogy (DIA) methods cannot disentangle the delicate transformation from irrelevant image contents, leading to ambiguous and coarse results. VII and Analogist cannot preserve the structural and invariant details of the original image. This is unacceptable for image tone transformation such as HDR, that has strict fidelity requirements. With the ground-truth editing instructions, the results of InsP2P are low-fidelity and do not exactly follow the desired pattern in visual prompts. This shows the necessity of visual prompts for indescribable visual transformation.

**Leakage of irrelevant content and disentanglement.** Our method aims to learn a disentangled editing effect by training with attention injection. The learned text embedding should not encode irrelevant information from visual prompts. The results show that our method only learns the editing effect but does not introduce the color, objects, or textures from the visual prompts. Concretely, Fig. 4 and 3 show that VII introduces the bear and its texture (first row), dogs (second row), and color (third row) to the edited images. DiffuseIT

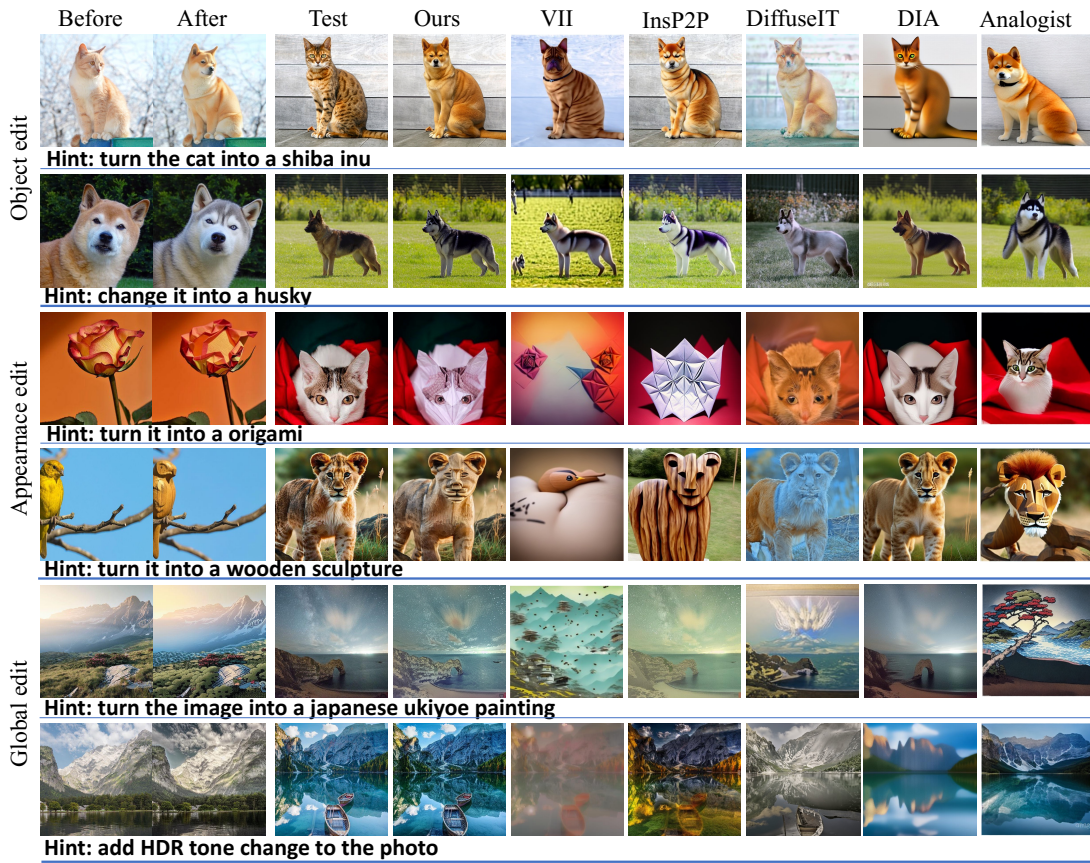


Figure 3: **Qualitative comparisons on real images.** Visual prompts with different editing types and different levels of geometric changes. Our method generalizes to different editing types and scenes while preserving different levels of geometric structures.



Figure 4: **Generalization in heterogeneous scenes and categories.** Results of tone and style editing show our method does not introduce *leaked content* (bear texture, dog face, color) from visual prompts when the category and scene of test images differ greatly from the visual prompts.

learns the transformation only from the single after-image and cannot learn the contrasted and accurate editing from the image pair. DIA cannot accurately learn the disentangled effect from the image pair, yielding inaccurate results.

**Generalization.** Fig. 4 also validates the generalization in heterogeneous scenes and categories. Our method can edit images of various scenes and categories with high fidelity whereas VII (based on the TI2I model) degrades significantly

when test images’ scenes or categories differ greatly from visual prompts. This also validates our motivation to use T2I model with larger prior and differential attention control.

**Necessity of visual prompts and ambiguity of text prompt.** We demonstrate the value of visual prompts from two sides. First, some editing effects cannot be accurately *identified or explained* by observation. Second, even if the effect can be identified as text such as ‘watercolor’, the text prompt can be ambiguous, and the ‘watercolor’ defined by the text prompt can be different from the ‘watercolor’ defined in the visual prompt. We demonstrate this in Fig. 13 in Appendix. There are various sub-styles of painting. The ‘psychedelic painting’ is difficult to recognize and the ‘watercolor’ effect in the visual prompt differs from that in InsP2P using text.

**Quantitative comparison and user study.** We extend evaluating metrics in Nguyen et al. (2024); Ju et al. (2024) with DINOv2 and VIE (Ku et al. 2024) based on GPT-4o. The VisualCLIP (V-CLIP) measures the cosine similarity of the editing direction between the before/after example pair and test pair based on their CLIP embeddings, indicating the agreement of the editing direction with the visual prompt. The ImageCLIP (I-CLIP) measures the cosine similarity between the edited and original images, indicating the fidelity to the original test image. The same applies to DINOv2. Sim-

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	V-CLIP $\uparrow$	I-CLIP $\uparrow$	V-DINO $\uparrow$	I-DINO $\uparrow$	V-VIE $\uparrow$	I-VIE $\uparrow$	Edit Analogy $\uparrow$	Fidelity $\uparrow$	Overall $\uparrow$
VII	12.76	0.4460	0.5238	0.1819	0.7012	0.1564	0.7144	1.92	1.69	3.21	1.33	2.27
InsP2P	15.75	0.5977	0.3168	0.2674	0.8114	0.2247	0.8763	3.01	2.93	4.21	2.93	3.57
DiffuseIT	12.06	0.3610	0.5261	0.1103	0.7180	0.1167	0.7538	2.11	3.32	2.98	3.28	3.13
DIA	17.78	0.5344	0.4279	0.1201	0.7453	0.1010	0.7429	1.18	4.01	2.14	3.87	3.01
Analogist	15.26	0.5102	0.4061	0.2123	0.7037	0.1914	0.7419	2.51	3.24	3.98	3.02	3.50
<b>Ours</b>	<b>24.57</b>	<b>0.8091</b>	<b>0.1197</b>	<b>0.2750</b>	<b>0.8178</b>	<b>0.3234</b>	<b>0.9073</b>	<b>4.48</b>	<b>4.29</b>	<b>4.81</b>	<b>4.62</b>	<b>4.72</b>

Table 1: **Quantitative comparison and user study.** We evaluate editing direction and fidelity in classical metrics, CLIP, DINO, GPT-4o, and human preference. We achieve better editing consistency aligned with the visual prompts and fidelity on all metrics.

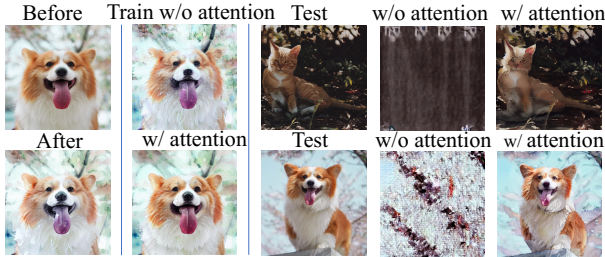


Figure 5: **Train and test results with/without attention.** 1st column: the visual prompt; 2nd column: training results without/with injected attention. 3rd column: test images; 4&5 columns: test results without and with attention control.

ilarly, A user study evaluates the preservation of test image invariance (Fidelity) and the analogy of the edit direction (Edit Analogy). VIE scores, evaluated by GPT-4o, assess these properties similarly to human criteria (Ku et al. 2024). Our method achieves competitive results across all metrics, consistent with qualitative results. See Appendix D.

### 4.3 Ablation and Analysis

**Differential attention control.** The differential attention control injects the attention of the before-image into that of the after-image. We validate its two benefits with experiments in Fig. 5. First, during training, the injected attention preserves the invariance making the text embedding fit the after-image better. As shown in Fig. 5 2nd column, the model produces a detailed dog portrait with injected attention, unlike the coarse version without it. This helps the text embedding focus on the transformation, improving generalization. Second, during testing, the module can preserve the invariance from the before-image, achieving high fidelity. In 4&5 columns, the attention module allows the text embedding to edit different images accurately. Without it, outputs are distorted, indicating overfitting to the training image pair.

**Time-aware scaling loss.** We analyzed the predicted  $\mathbf{x}_0$ ,  $\mathbf{f}_\theta(\mathbf{x}_t, t)$  along all timesteps. The observation in Fig. 6 shows that equally optimizing the loss at each timestep may reduce the loss of predicted  $\mathbf{x}_0$  in the early steps too much and make the final output deviate from the ground truth  $\mathbf{x}_0^a$ . Meanwhile, the output that is closest to  $\mathbf{x}_0^a$  appears earlier and the image quality degrades undesirably. After scaling, the final output



Figure 6: **The predicted  $\mathbf{x}_0$  at different timesteps.** The results of early steps without scaling are closer to  $\mathbf{x}_0^b$  but the final output deviates. With scaling, the results of early steps are farther but the final output approaches  $\mathbf{x}_0^b$ .

becomes closest to  $\mathbf{x}_0^a$ , and the image quality also improves. **Limitation.** One limit is the prior of the T2I model. Despite using the more general T2I model with a larger prior, editing real images remains challenging since many practical transformations are out of the model’s prior. Learning visual prompts is essentially a process of extracting and composing the prior of the text-to-image. Such prior restricts our method to learning visual prompts and editing for real images. Another is choosing the timestamp of the attention injection for different editing types since the invariance of them is different. To fit the delicate transformation, we need to choose the appropriate timestep for each editing type in training.

## 5 Conclusion

This paper introduces a method using a single, general text-to-image model as a diffusion bridge for learning complex image transformations from visual prompts. This approach removes the dependence on a text-guided image-to-image model and extensive triplet data of text and before-and-after images, which are harder to create than simple text-image pairs and are not scalable on a large scale. We also demonstrate that training with differential attention allows the visual transformation to be embedded into the text space, disentangling the text embedding from the irrelevant before-image information and enhancing the generalizability of text embeddings for various image edits. Results show that leveraging the scalable T2I model achieves high-fidelity image editing with visual prompts, suggesting a new paradigm for leveraging scalable generative models for detailed visual prompts.

## Acknowledgments

We appreciate constructive feedback from anonymous reviewers and meta-reviewers. This work is supported by Natural Sciences and Engineering Research Council of Canada (NSERC), Discovery Grants program.

## References

- Bar, A.; Gandelsman, Y.; Darrell, T.; Globerson, A.; and Efros, A. 2022. Visual prompting via image inpainting. *Advances in Neural Information Processing Systems*, 35: 25005–25017.
- Brooks, T.; Holynski, A.; and Efros, A. A. 2023. Instruct-pix2pix: Learning to follow image editing instructions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 18392–18402.
- Brown, T.; Mann, B.; Ryder, N.; Subbiah, M.; Kaplan, J. D.; Dhariwal, P.; Neelakantan, A.; Shyam, P.; Sastry, G.; Askell, A.; et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901.
- Cao, M.; Wang, X.; Qi, Z.; Shan, Y.; Qie, X.; and Zheng, Y. 2023. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22560–22570.
- Cheng, B.; Liu, Z.; Peng, Y.; and Lin, Y. 2023. General image-to-image translation with one-shot image guidance. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 22736–22746.
- Floridi, L.; and Chiriatti, M. 2020. GPT-3: Its nature, scope, limits, and consequences. *Minds and Machines*, 30: 681–694.
- Gal, R.; Alaluf, Y.; Atzmon, Y.; Patashnik, O.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2023. An Image is Worth One Word: Personalizing Text-to-Image Generation using Textual Inversion. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*.
- Gu, Z.; Yang, S.; Liao, J.; Huo, J.; and Gao, Y. 2024. Analogist: Out-of-the-box Visual In-Context Learning with Image Diffusion Model. *ACM Transactions on Graphics (TOG)*, 43(4): 1–15.
- Hertz, A.; Mokady, R.; Tenenbaum, J.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2023. Prompt-to-Prompt Image Editing with Cross-Attention Control. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*.
- Hertzmann, A.; Jacobs, C. E.; Oliver, N.; Curless, B.; and Salesin, D. H. 2023. Image analogies. In *Seminal Graphics Papers: Pushing the Boundaries, Volume 2*, 557–570.
- Ho, J.; and Salimans, T. 2022. Classifier-free diffusion guidance. *arXiv preprint arXiv:2207.12598*.
- Ju, X.; Zeng, A.; Bian, Y.; Liu, S.; and Xu, Q. 2024. PnP Inversion: Boosting Diffusion-based Editing with 3 Lines of Code. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*.
- Kawar, B.; Zada, S.; Lang, O.; Tov, O.; Chang, H.; Dekel, T.; Mosseri, I.; and Irani, M. 2023. Imagic: Text-based real image editing with diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6007–6017.
- Ku, M.; Jiang, D.; Wei, C.; Yue, X.; and Chen, W. 2024. VIEScore: Towards Explainable Metrics for Conditional Image Synthesis Evaluation. In Ku, L.; Martins, A.; and Srikumar, V., eds., *Proceedings of the 62nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2024, Bangkok, Thailand, August 11-16, 2024*, 12268–12290. Association for Computational Linguistics.
- Kwon, G.; and Ye, J. C. 2023. Diffusion-based Image Translation using disentangled style and content representation. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*.
- Liu, P.; Yuan, W.; Fu, J.; Jiang, Z.; Hayashi, H.; and Neubig, G. 2023. Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing. *ACM Computing Surveys*, 55(9): 1–35.
- Mokady, R.; Hertz, A.; Aberman, K.; Pritch, Y.; and Cohen-Or, D. 2023. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6038–6047.
- Motamed, S.; Paudel, D. P.; and Van Gool, L. 2025. Lego: Learning to Disentangle and Invert Personalized Concepts Beyond Object Appearance in Text-to-Image Diffusion Models. In *European Conference on Computer Vision*, 116–133. Springer.
- Mou, C.; Wang, X.; Xie, L.; Wu, Y.; Zhang, J.; Qi, Z.; and Shan, Y. 2024. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 4296–4304.
- Nguyen, T.; Li, Y.; Ojha, U.; and Lee, Y. J. 2024. Visual Instruction Inversion: Image Editing via Image Prompting. *Advances in Neural Information Processing Systems*, 36.
- Parmar, G.; Kumar Singh, K.; Zhang, R.; Li, Y.; Lu, J.; and Zhu, J.-Y. 2023. Zero-shot image-to-image translation. In *ACM SIGGRAPH 2023 Conference Proceedings*, 1–11.
- Podell, D.; English, Z.; Lacey, K.; Blattmann, A.; Dockhorn, T.; Müller, J.; Penna, J.; and Rombach, R. 2024. SDXL: Improving Latent Diffusion Models for High-Resolution Image Synthesis. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*.
- Ramesh, A.; Dhariwal, P.; Nichol, A.; Chu, C.; and Chen, M. 2022. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 1(2): 3.
- Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 10684–10695.
- Ruiz, N.; Li, Y.; Jampani, V.; Pritch, Y.; Rubinstein, M.; and Aberman, K. 2023. Dreambooth: Fine tuning text-to-image

- diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 22500–22510.
- Saharia, C.; Chan, W.; Saxena, S.; Li, L.; Whang, J.; Denton, E. L.; Ghasemipour, K.; Gontijo Lopes, R.; Karagol Ayan, B.; Salimans, T.; et al. 2022. Photorealistic text-to-image diffusion models with deep language understanding. *Advances in neural information processing systems*, 35: 36479–36494.
- Song, J.; Meng, C.; and Ermon, S. 2021. Denoising Diffusion Implicit Models. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*.
- Song, Y.; Sohl-Dickstein, J.; Kingma, D. P.; Kumar, A.; Ermon, S.; and Poole, B. 2021. Score-Based Generative Modeling through Stochastic Differential Equations. In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*.
- Su, X.; Song, J.; Meng, C.; and Ermon, S. 2023. Dual Diffusion Implicit Bridges for Image-to-Image Translation. In *The Eleventh International Conference on Learning Representations, ICLR 2023, Kigali, Rwanda, May 1-5, 2023*.
- Šubrtová, A.; Lukáč, M.; Čech, J.; Futschik, D.; Shechtman, E.; and Šykora, D. 2023. Diffusion image analogies. In *ACM SIGGRAPH 2023 Conference Proceedings*, 1–10.
- Touvron, H.; Lavril, T.; Izacard, G.; Martinet, X.; Lachaux, M.-A.; Lacroix, T.; Rozière, B.; Goyal, N.; Hambro, E.; Azhar, F.; et al. 2023. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*.
- Tumanyan, N.; Geyer, M.; Bagon, S.; and Dekel, T. 2023. Plug-and-play diffusion features for text-driven image-to-image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1921–1930.
- Wang, X.; Wang, W.; Cao, Y.; Shen, C.; and Huang, T. 2023a. Images speak in images: A generalist painter for in-context visual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6830–6839.
- Wang, X.; Zhang, X.; Cao, Y.; Wang, W.; Shen, C.; and Huang, T. 2023b. SegGPT: Towards Segmenting Everything In Context. In *IEEE/CVF International Conference on Computer Vision, ICCV 2023, Paris, France, October 1-6, 2023*, 1130–1140. IEEE.
- Wang, Z.; Jiang, Y.; Lu, Y.; He, P.; Chen, W.; Wang, Z.; Zhou, M.; et al. 2024. In-context learning unlocked for diffusion models. *Advances in Neural Information Processing Systems*, 36.
- Yang, Y.; Peng, H.; Shen, Y.; Yang, Y.; Hu, H.; Qiu, L.; Koike, H.; et al. 2024. ImageBrush: Learning Visual In-Context Instructions for Exemplar-Based Image Manipulation. *Advances in Neural Information Processing Systems*, 36.
- Ye, H.; Zhang, J.; Liu, S.; Han, X.; and Yang, W. 2023. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*.
- Zhang, K.; Mo, L.; Chen, W.; Sun, H.; and Su, Y. 2024. Magicbrush: A manually annotated dataset for instruction-guided image editing. *Advances in Neural Information Processing Systems*, 36.
- Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3836–3847.
- Zhao, S.; Chen, D.; Chen, Y.-C.; Bao, J.; Hao, S.; Yuan, L.; and Wong, K.-Y. K. 2024. Uni-controlnet: All-in-one control to text-to-image diffusion models. *Advances in Neural Information Processing Systems*, 36.
- Zhou, L.; Lou, A.; Khanna, S.; and Ermon, S. 2024. Denoising Diffusion Bridge Models. In *The Twelfth International Conference on Learning Representations, ICLR 2024, Vienna, Austria, May 7-11, 2024*.