

# VersaGen: Unleashing Versatile Visual Control for Text-to-Image Synthesis

Zhipeng Chen<sup>1</sup>, Lan Yang<sup>1,2</sup>, Yonggang Qi<sup>1,2</sup>,  
Honggang Zhang<sup>1</sup>, Kaiyue Pang<sup>2</sup>, Ke Li<sup>1,2\*</sup>, Yi-Zhe Song<sup>2</sup>

<sup>1</sup>School of Artificial Intelligence, Beijing University of Posts and Telecommunications, China

<sup>2</sup>SketchX, CVSSP, University of Surrey, United Kingdom

zhipengchen1998@bupt.edu.cn, ylan@bupt.edu.cn, qi yg@bupt.edu.cn,

zhhg@bupt.edu.cn, thatkpang@gmail.com, like1990@bupt.edu.cn, y.song@surrey.ac.uk

## Abstract

Despite the rapid advancements in text-to-image (T2I) synthesis, enabling precise visual control remains a significant challenge. Existing works attempted to incorporate multi-facet controls (text and sketch), aiming to enhance the creative control over generated images. However, our pilot study reveals that the expressive power of humans far surpasses the capabilities of current methods. Users desire a more versatile approach that can accommodate their diverse creative intents, ranging from controlling individual subjects to manipulating the entire scene composition. We present VersaGen, a generative AI agent that enables versatile visual control in T2I synthesis. VersaGen admits four types of visual controls: i) single visual subject; ii) multiple visual subjects; iii) scene background; iv) any combination of the three above or merely no control at all. We train an adaptor upon a frozen T2I model to accommodate the visual information into the text-dominated diffusion process. We introduce three optimization strategies during the inference phase of VersaGen to improve generation results and enhance user experience. Comprehensive experiments on COCO and Sketchy validate the effectiveness and flexibility of VersaGen, as evidenced by both qualitative and quantitative results.

## Code —

[https://github.com/FelixChan9527/VersaGen\\_official](https://github.com/FelixChan9527/VersaGen_official)

## 1 Introduction

Ask any artist, and they will likely express their discontent with generative AI. They demand control over the creative process, control that aligns with their creative intent, and rightfully so! For you and me alike, beyond the initial "wow", questions are also starting to form around why models do not follow exactly what we want. There is a clear trend in the generative AI literature towards enhancing user control, which is encouraging. However, to fully understand the limitations of current approaches, it is essential to articulate the gap between human artistic expressivity and the capabilities of generative AI. To shed light on this, we first conducted a pilot study in which participants used Stable Diffusion (SD) (Rombach et al. 2022) and ControlNet (Zhang, Rao, and

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




Figure 1: VersaGen can handle various forms of input provided by users, whether at the object-level, scene-level, or a combination of both.

Agrawala 2023) to generate images matching a given reference. The results revealed that participants faced two main challenges: incapability and inflexibility.

Incapability refers to the technical barriers that hinder users from achieving their desired visual outcomes, an issue widely acknowledged and tackled in recent works (Cao et al. 2023; Ju et al. 2023; Ge et al. 2023; Gong et al. 2023; Wu et al. 2023; Gafni et al. 2022; Ma et al. 2023; Chang et al. 2023). These solutions introduce various multi-modal input forms to make the process more intuitive. However, inflexibility remains largely underappreciated. Systems like ControlNet (Zhang, Rao, and Agrawala 2023) and T2I-Adapter (Mou et al. 2024) require users to control the entire scene, which our findings suggest rarely facilitates meaningful engagement. In fact, 74.96% of the 1,250 ControlNet trials reported that users found it exceedingly difficult to provide comprehensive visual guidance, particularly for complex scenes with multiple subjects.

To address the challenges of incapability and inflexibility, we propose a natural next step: allowing individuals the flexibility to decide what to control or not. Recognising that people express their creative intent differently – we re-envision

the visual control problem in generative AI and transform it from a single visual subject/scene-level control to a hybrid approach comprising multiple degrees of control freedom.

Our proposed solution, VersaGen, targets drawing as the primary control tool, given its long-standing role in human expression of the visual world (Aubert et al. 2014; Karmiloff-Smith 1990; Hoffmann et al. 2018; Gombrich and Gombrich 1995; Jongejan et al. 2017; Eitz, Hays, and Alexa 2012). Unlike previous systems that require compulsory holistic scene drawings, VersaGen admits four types of visual controls (Fig. 1): i) single visual subject; ii) multiple visual subjects; iii) scene background; iv) any combination of the three above or merely no control at all, falling back to pure text-to-image generation. By providing this flexibility and allowing users to choose the level of control that best suits their needs and preferences, VersaGen aims to make generative control more inclusive, accessible, and enjoyable for all, transforming the creative process into a fun and engaging experience. Achieving VersaGen is technically demanding, with two major challenges to overcome. The first challenge lies in simulating the various types of drawing inputs that VersaGen is designed to accommodate, given the scarcity of such data for training. To address this, we extract and optimise edge maps from images to imitate human drawing input. Although edges are not perfect substitutes for real drawings (e.g., human free-hand sketches), they provide a scalable way to obtain relevant training data. Additionally, we design an Adaptive Control Strength mechanism (Sec. 4.4) to mitigate the disparity between edge maps and real drawings during inference. The key insight is that while drawing artifacts in realistic human input can degrade generation quality, their impact can be minimised by limiting their contribution to later steps in the diffusion chain.

The second challenge is visual localisation. Unlike existing works that require users to precisely position visual controls, VersaGen simplifies this process by automatically locating user-provided visual controls within the appropriate local context. To achieve this, we leverage recent advances showing that T2I models inherently function as semantic segmenters, capable of generating bounding boxes encircling objects based on their corresponding tokens at certain diffusion timesteps (Hertz et al. 2023; Zhang et al. 2021; Patashnik et al. 2023). To further address potential localisation inaccuracies during inference, we introduce the Multimodal Conflict Resolver in VersaGen (Sec. 4.3). This component employs both token-level and pixel-level objectives to closely align the corresponding regions in the latent space across both modalities, thereby reducing conflicts and preventing unexpected generation results.

We conduct extensive evaluations using both edge maps and human free-hand sketches on the COCO (Lin et al. 2014) and Sketchy (Sangkloy et al. 2016) datasets. Our results demonstrate that VersaGen outperforms well-established T2I (Rombach et al. 2022) and controllable T2I models (Mou et al. 2024; Zhang, Rao, and Agrawala 2023) in both quantitative and qualitative comparisons. Furthermore, a human study reveals that 48% of users identify VersaGen as the most user-friendly interactive generation model compared to alternative approaches, underscoring the importance of providing

flexible control options that cater to diverse user preferences and creative intents. Finally, a comprehensive ablation study highlights the crucial role of our three proposed strategies in enabling VersaGen to produce high-quality, user-controlled visual outputs across a wide range of input conditions in real-world scenarios.

## 2 Related Work

Recent advancements in diffusion models, such as Imegen (Saharia et al. 2022), Stable Diffusion (SD) (Rombach et al. 2022), and DALL-E (Ramesh et al. 2021), have revolutionised text-to-image generation tasks. These models excel at producing high-quality images but often lack fine-grained control over the generated content. To address this limitation, researchers have explored various approaches to enhance user control in generative AI. One line of research focuses on text-driven control methods, which involve adjusting prompts (Gal et al. 2022; Gani et al. 2023; Kim, Kwon, and Ye 2022; Wang et al. 2023; Lee et al. 2023; Liu and Chilton 2022; Marcus, Davis, and Aaronson 2022; Wang et al. 2022) or improving cross-attention maps (Feng et al. 2022; Chefer et al. 2023; Li et al. 2023a; Rassin et al. 2024; Xu et al. 2024). While these methods have shown promise in guiding the generation process, they often struggle to provide precise structural guidance, as text prompts alone may not fully capture the user’s intended visual composition. Another approach to controllable image generation involves incorporating additional input modalities, such as sketches or layouts. Layout Guidance Diffusion (Chen, Laina, and Vedaldi 2024), GLIGEN (Li et al. 2023b) and InstanceDiffusion (Wang et al. 2024) leverage user-defined tokens and bounding boxes to guide cross-attention maps, while (Balaji et al. 2022) imposes structural constraints on the generated images using similarity gradients between target sketches and intermediate model features. ControlNet (Zhang, Rao, and Agrawala 2023) and T2I-Adapter (Mou et al. 2024) introduces an adapter to combine internal knowledge from text-to-image models with external control signals. UniControl (Qin et al. 2024) builds on ControlNet and integrates task instructions into condition-specific networks to adapt to various visual inputs. Despite these advancements, existing methods often focus on single scene-level control, which can limit the degree of control available to users. In contrast, VersaGen takes a more flexible approach, accommodating both object-level and scene-level visual conditions without requiring the depiction of entire scene conditions. By offering a range of drawing options and allowing users to choose the level of control that best suits their needs, VersaGen empowers users with versatile control over the generative process.

## 3 Pilot Study

To assess how well current off-the-shelf generative models meet user needs for generating target images, we designed this pilot study featuring Stable Diffusion (Rombach et al. 2022), a T2I model, and ControlNet (Zhang, Rao, and Agrawala 2023), a controllable T2I model, as illustrated in Fig. 2(a). Users are provided with a reference image and a corresponding textual prompt, tasked with using SD and Con-

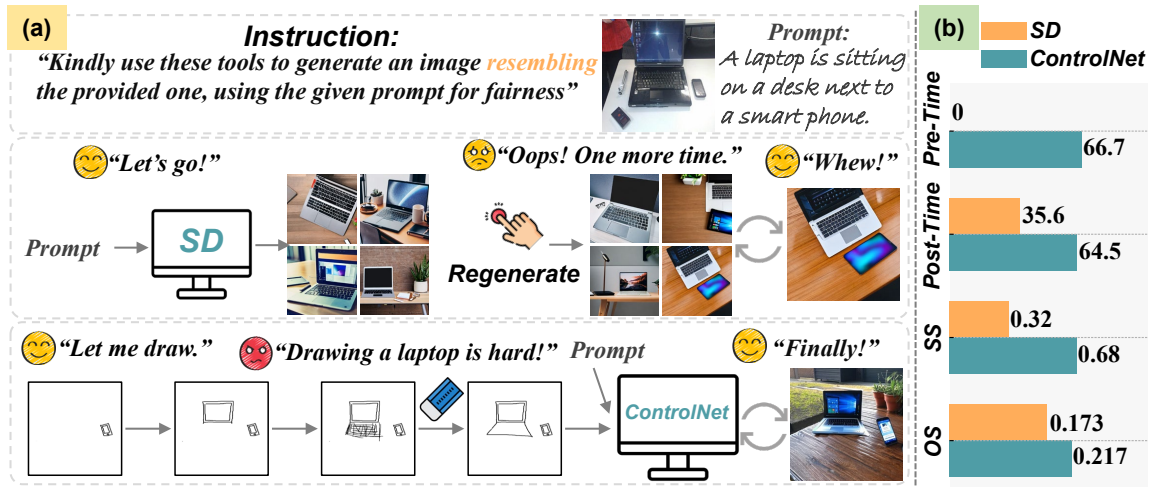


Figure 2: (a) The illustration of pilot study: users are tasked with generating an image similar to the given reference image using SD and ControlNet. (b) Quantitative evaluation of SD and ControlNet.

ControlNet to generate a matching image. In both experiments, participants are instructed to copy the prompt verbatim to maintain fairness. When using SD, users copy the given textual prompt, click “Generate”, and then select a satisfactory generated image to “Submit”. With ControlNet, additional preparation of visual conditions was required. Then, they copy the textual prompt, click “Generate”, and wait for generated results, from which they select one to “Submit”. Notably, i) a “Regenerate” button is provided, allowing users to regenerate images multiple times until they find a satisfactory result, at which point they can click “Submit”. ii) Users cannot access the reference image after selecting the test generation model and must rely solely on their memory to complete the process. iii) Each time they click “Submit” during ControlNet trials, they are asked if providing an effective input condition for the current reference image is easy.

We employ four metrics to evaluate the two generative methods: i) **Pre-Time** ( $s_{\downarrow}$ ): The duration from when the user clicks the button to select a specific generative method to when they complete the preparation of input conditions and click the “Generate” button. This metric indicates the ease of use for end users. ii) **Post-Time** ( $s_{\downarrow}$ ): The time from when the method finishes generating a set of images to when the user clicks “Submit”. Notably, users can regenerate results if they are unsatisfied with the current images. The Post-Time is accumulated, excluding the time spent waiting for model generation. iii) **Objective Similarity** ( $OS_{\downarrow}$ ): We use FID as a quantitative metric to evaluate the similarity between the given reference images and the user-submitted results. iv) **Subjective Similarity** ( $SS_{\uparrow}$ ): We invited five experts to evaluate the user-submitted results generated by the two methods. During each trial, the experts were presented with three images: reference, SD-generated, and ControlNet-generated. They were then asked to choose the one closer to the reference. The selection rate is referred to as SS.

In Fig. 2(b), the findings show: i) There is no Pre-Time for SD as it solely relies on textual inputs, following instructions

for participants to directly copy the given prompt. Conversely, considerable time is spent preparing visual conditions for ControlNet to enable more controlled generation. ii) Interestingly, participants spent more Post-Time on ControlNet than on SD. This indicates that for controllable generation models, good input conditions are crucial for achieving the desired results. It is not easy for participants to quickly find appropriate visual input conditions for ControlNet. iii) For the similarity evaluation, we observed that SD performs better in objective high-dimensional semantic feature measurement, as it is trained to generate semantically aligned images. However, for human perceptual visual similarity, ControlNet excels because the input visual conditions reflect the user’s subjective interpretation of the reference image.

## 4 Methods

**Overview** The goal of VersaGen is to evolve human interaction with generative models from a strictly scene-level visual condition to a hybrid approach that supports multiple levels of drawing control. Users can freely select their preferred levels of visual conditions, eliminating the need for comprehensive scene-level conditions and bypassing the technical challenges associated with precise scene-level control. VersaGen accepts diverse drawing inputs: single visual subject drawing  $s_i$ , multiple visual subjects drawings  $\{s_1, s_2, \dots, s_m\}$ , and background drawing  $s_{bg}$ . By synergising these elements with a simple textual prompt, users can realise highly flexible and controllable image generation. VersaGen comprises three core modules, as illustrated in Fig. 3: i) **Foundation Generation Model (FGM)** incorporates the trainable copy provided by UDE to “finetune” its frozen weights, thereby aligning VersaGen’s output with the user’s drawings. ii) **User Drawing Encoder (UDE)** processes drawings inputted by users, encoding these hybrid drawings into a latent representation that serves as a condition to update a trainable copy of the Foundation Generative Model for fine-tuning. iii) **Multi-modal Conflict Resolver (MCR)** addresses inconsistencies

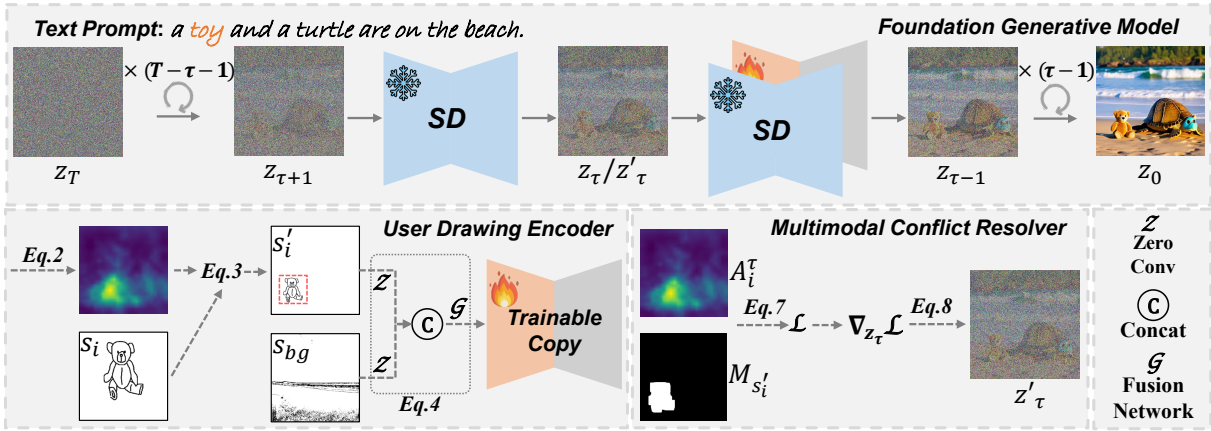


Figure 3: The illustration of VersaGen. At denoising timestep  $\tau$  during inference, MCR functions to update the original noisy latent  $z_\tau$  to  $z'_\tau$ , to alleviate potential conflicts across modalities.

between modalities (user drawings and textual prompt) during inference, enhancing the quality of the generated images.

#### 4.1 Foundation Generative Model

Here we adopt SD (Rombach et al. 2022) as our FGM  $\Phi$ , a state-of-the-art text-to-image generation model. The process begins with an input image  $z_0$ , to which random noise is progressively added, resulting in a noisy image  $z_t$ . With given textual condition  $C$ , a network  $\epsilon_\theta$  is trained to predict the noise added to the noisy image  $z_t$  with:

$$\mathcal{L} = \mathbb{E}_{z_t, t, C, \epsilon \sim \mathcal{N}(0, 1)} [\|\epsilon - \epsilon_\theta(z_t, t, C)\|_2^2]. \quad (1)$$

Here  $\mathcal{L}$  represents the overall learning objective of the entire diffusion model. Notably, the parameters in SD (Rombach et al. 2022), denoted as  $\Theta_\Phi$ , are kept fixed in the subsequent steps to preserve its foundational and powerful text-to-image generation capabilities.

#### 4.2 User Drawing Encoder

Given visual subjects' drawings  $S = \{s_1, s_2, \dots, s_m\}$  and a textual prompt  $C$ , the initial step is to determine the potential locations of  $s_i$  within the generated image  $I$ . We utilise the cross-attention mechanism in foundation generative model  $\Phi$  to achieve this. Assuming the textual tokens correspond to  $s_i$  are  $c_i$ , the cross-attention map  $A_i^t$  is computed as follows,

$$A_i^t(s_i, c_i) = \frac{1}{N} \sum_{n=1}^N \text{softmax}\left(\frac{Q^n(s_i)K^n(c_i)}{\sqrt{d}}\right). \quad (2)$$

Here,  $n$  refers to index of head in multi-heads mechanism, while  $t$  denotes the timestep in diffusion process, notably, the  $A_i^t$  is calculated only at  $t = \tau$ .  $Q^n(s_i) \in \mathbb{R}^{B*HW*d}$  applies a linear function on  $s_i$ 's latent to generate "query" vector in cross-attention, where  $H, W$  denote the height and width of latent, and  $d$  is embedding dimension. Similarly,  $K^n(c_i) \in \mathbb{R}^{B*d*77}$  involves a linear function followed by a transpose operation to produce "key" vector for cross-attention.

To identify the potential location of  $s_i$  within generated image  $I(S, C)$ , we first apply OTSU (Otsu et al. 1975) on  $A_i^t$

to compute the object-aware attention threshold, then extract potential mask  $R_i$  of  $s_i$ , as Eq.3 shown:

$$R_i(u) = \begin{cases} 1 & A_i^t(u) \geq \text{OTSU}(A_i^t) \\ 0 & \text{otherwise} \end{cases}, \quad (3)$$

where  $u$  is the spatial location of the element in  $R_i$ . Subsequently, the square bounding box  $B_i$  of non-zero region in  $R_i$  is used to guide the relocation of  $s_i$ , resulting in  $s'_i$ , which is input into drawing encoder  $\mathcal{Z}$  parameterised by  $\theta_{z1}$ . The latents of user drawings  $S$  are concatenated, and then passed through a small fusion network  $\mathcal{G}$  to produce a combined drawing latent  $z_S$ , as shown in Eq.4:

$$z_S = \mathcal{G}(\text{Concat}[\mathcal{Z}(s'_1; \theta_{z1}), \mathcal{Z}(s'_2; \theta_{z1}), \dots, \mathcal{Z}(s'_m; \theta_{z1}), \mathcal{Z}(s_{bg}; \theta_{z1})]; \theta_g). \quad (4)$$

To leverage the T2I capabilities of  $\Phi$ , we adopt the ControlNet (Zhang, Rao, and Agrawala 2023) pipeline to add spatial control capability to  $\Phi$ . ControlNet (Zhang, Rao, and Agrawala 2023) creates a trainable copy of UNet in  $\Phi$ , which features zero convolution layers within the encoder blocks and the middle block. This process is formulated as:

$$y_S = \underbrace{\mathcal{F}(z_t; \Theta_\Phi)}_{\text{FGM}} + \underbrace{\mathcal{Z}(\mathcal{F}(z_t + z_S; \Theta_S); \theta_{z2})}_{\text{UDE}}, \quad (5)$$

where  $\mathcal{F}$  is UNet,  $\Theta_\Phi$  is the frozen weight of pretrained  $\Phi$ ,  $\mathcal{Z}$  corresponds to zero convolutions with learnable weights  $\theta_{z2}$ ,  $\Theta_S$  is the parameters of trainable copy of UNet. The whole training object of VersaGen is formulated as:

$$\mathcal{L} = \mathbb{E}_{z_t, t, C, z_S, \epsilon \sim \mathcal{N}(0, 1)} [\|\epsilon - \epsilon_\theta(z_t, t, C, z_S)\|_2^2]. \quad (6)$$

#### 4.3 Multimodal Conflict Resolver

For multimodal controllable image generation, generating a high quality realistic image hinges on ensuring that the latent can coherently integrate conditional information from different modalities. During the training phase of VersaGen, the textual prompt  $C$  is carefully aligned with the drawings

Methods	w/GT		COCO				Sketchy			
	LoD	BBoxes	CLIP↑	FID↓	DINO↑	ACC↑	CLIP↑	FID↓	DINO↑	ACC↑
Stable Diffusion (Rombach et al. 2022)	/	/	<b>0.2556</b>	0.4026	0.7616	0.6803	0.2509	0.3064	0.7934	0.7710
T2I-Adapter (Mou et al. 2024)	✓	×	0.2439	0.3732	0.7722	0.6143	0.2392	0.3290	0.7843	0.7182
ControlNet (Zhang, Rao, and Agrawala 2023)	✓	×	0.2523	0.3560	0.7774	0.6547	0.2517	0.3493	0.8007	0.7622
UniControl (Qin et al. 2024)	✓	×	0.2467	0.3558	0.7475	0.7006	0.2479	0.3165	0.7988	0.8096
GLIGEN (Li et al. 2023b)	×	✓	0.2515	0.3571	0.7740	0.4319	0.2419	0.2961	0.8149	0.5550
InstanceDiffusion (Wang et al. 2024)	×	✓	0.2419	0.3811	0.7644	<b>0.8094</b>	0.2399	0.2818	0.8179	<b>0.8942</b>
VersaGen (Ours)	✓	×	0.2537	<b>0.3377</b>	<b>0.8019</b>	0.7322	<b>0.2524</b>	<b>0.2614</b>	<b>0.8193</b>	0.8328
VersaGen (Ours)	×	×	0.2542	<b>0.3431</b>	<b>0.7879</b>	0.7229	<b>0.2542</b>	0.3080	0.8088	0.8174

Table 1: Quantitative results. “w/ GT” denotes the locations of drawings (LoD) or bounding boxes (BBoxes) from GT.

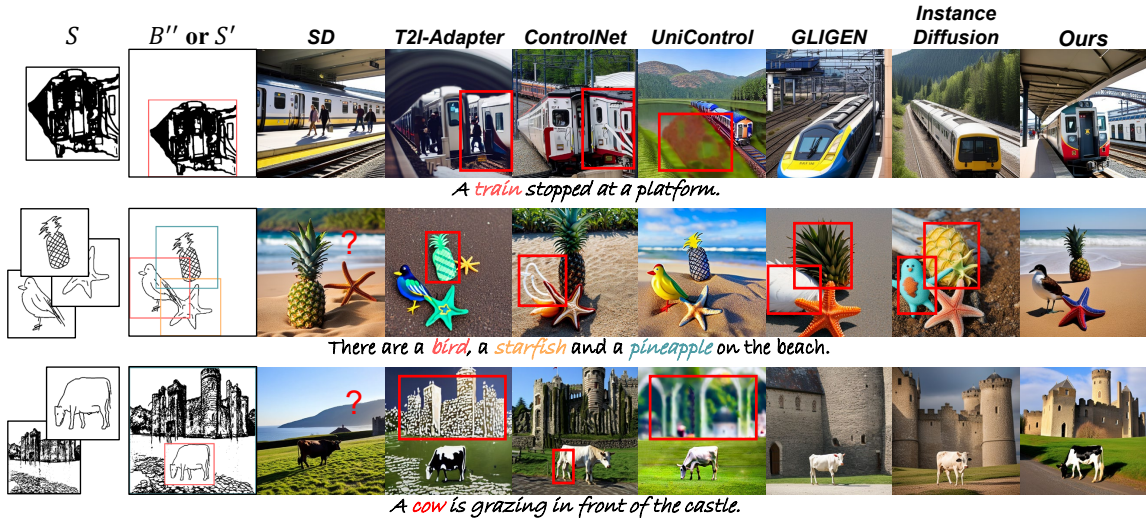


Figure 4: Visualised comparison of SD, T2I-Adapter, ControlNet, UniControl, GLIGEN, InstanceDiffusion and our proposed VersaGen. Problematic regions are highlighted with  $\square$ , and missing entities are indicated by  $?$ .

$S$  because both are derived from the ground truth image  $z_0$ . However, during inference, we cannot guarantee that the user-provided drawings are always coherent with the accompanying textual prompt which may result in a modality conflict during the generation process. Therefore, we introduce a Multimodal Conflict Resolver (MCR) during inference to mitigate the impact of such issues.

Inspired by (Wang et al. 2023), we address this issue by a test-time optimisation. Our objective is to ensure that the active region in the cross-attention map  $A_i^t(z_t, c_i)$  closely aligns with the contour  $M_{s'_i}$  of  $s'_i$ . Therefore we involve a token-level loss and a pixel-level loss to achieve this goal, as formulated by:

$$\begin{aligned}
 \mathcal{L}_{\text{token}} &= \frac{1}{m} \sum_{l=1}^L \sum_{i=1}^m \left( 1 - \frac{\sum_u M_{s'_i}(u) A_{il}^t(u)}{\sum_u A_{il}^t(u)} \right)^2, \\
 \mathcal{L}_{\text{pixel}} &= \frac{1}{m} \sum_{l=1}^L \sum_{i=1}^m \text{BCE}(M_{s'_i}, A_{il}^t), \\
 \mathcal{L} &= \lambda \mathcal{L}_{\text{token}} + (1 - \lambda) \mathcal{L}_{\text{pixel}}.
 \end{aligned} \tag{7}$$

where  $l$  is the layer index of the UNet,  $\text{BCE}(\cdot, \cdot)$  represents binary cross entropy loss. Since different layers have different

resolutions of  $A_{il}^t$ , we downscale the resolution of  $M_{s'_i}$  to match that of  $A_{il}^t$  using bilinear interpolation, followed by binarisation of all values. Then we use  $\mathcal{L}$  to update  $z_t$ :

$$z'_t = z_t - \alpha \nabla_{z_t} \mathcal{L}, \tag{8}$$

where  $\alpha$  is a hyperparameter used to adjust the gradient update step size. It is important to note that MCR only optimises  $z_t$  at  $t = \tau$ , resolving the modality conflict at the timestep immediately following the incorporation of user drawing control. Once the optimisation is completed, MCR ceases operation in subsequent timesteps to minimise resource use.

#### 4.4 VersaGen In-the-Wild

In real-world applications of VersaGen, two main challenges arise: i) when users provide multiple visual subjects drawings, the correlation between objects causes significant overlap in the object-aware active region  $R_i$ s, leading to low-quality, occluded outputs. ii) Amateur users often struggle to create precise, high-quality drawings, resulting in unrealistic generated images. To improve VersaGen’s real-world adaptability, we propose two strategies during inference: multi-object decoupling and adaptive control strength.

**Multi-object Decoupling.** Given the user drawings  $S = \{s_1, s_2, \dots, s_m\}$ , their corresponding active regions

$R_i$  in cross-attention map can be calculated by Eq.2. Each  $R_i$  can be simplified as a bounding box  $B_i = (x_i^{min}, y_i^{min}, x_i^{max}, y_i^{max})$ , where  $(x_i^{min}, y_i^{min})$  are the coordinates of the top-left corner of  $B_i$ ,  $(x_i^{max}, y_i^{max})$  are bottom-right corner. The goal of multi-object decoupling is to ensure the Intersection over Union (IoU) of  $B_i$ s remains within a reasonable range, avoiding excessive overlapping. Each  $B_i$  has its center at  $(c_x^i, c_y^i)$ , and the overall center of all  $B_i$ s is  $(c_x, c_y)$ , which is calculated as the average of all individual centers. For any two visual subjects drawings  $s_i$  and  $s_j$ , we calculate the IoU between  $B_i$  and  $B_j$ . If  $\text{IoU}(s_i, s_j) \geq \beta$ , we apply two operations on both  $B_i$  and  $B_j$  simultaneously: i) *resizing*, ii) *translation*, to adjust their bounding boxes and reduce overlap.

*Resizing* shrinks the region of  $B_i$  by two units towards  $(c_x^i, c_y^i)$ ,

$$B_i' = (x_i^{min} + 1, y_i^{min} + 1, x_i^{max} - 1, y_i^{max} - 1). \quad (9)$$

*Translation* moves  $B_i$  by one unit away from  $(c_x, c_y)$ . First, we calculate the direction vector  $\vec{d} = (c_x - c_x^i, c_y - c_y^i)$ , and then obtain the unit direction vector  $\hat{d}$ . We then translate  $B_i'$  to  $B_i'' = (x_i''^{min}, y_i''^{min}, x_i''^{max}, y_i''^{max})$  as follows:

$$\begin{aligned} (x_i''^{min}, y_i''^{min}) &= (x_i'^{min}, y_i'^{min}) - \hat{d}, \\ (x_i''^{max}, y_i''^{max}) &= (x_i'^{max}, y_i'^{max}) - \hat{d}. \end{aligned} \quad (10)$$

In practice, we apply these two operations in combination until  $\text{IoU}(s_i, s_j) < \beta$ . More details about the combination strategy can be found in the supplemental material(Chen et al. 2024).

**Adaptive Control Strength.** For amateur end users, the goal of providing drawing control is not to make the generated image strictly follow their strokes, as they understand their limited drawing skills may not fully represent their intentions. This makes the textual prompt crucial for achieving their desired results. To address this, we design an adaptive control strength strategy that reduces the influence of user drawings as the denoising timesteps progress. In the early structure-forming stage, user drawings dominate to establish the basic structure. In the later detail generation stage, the textual prompt takes over, adding coherent details based on its well-pretrained text-to-image (T2I) capabilities. This strategy is implemented by introducing a weight  $\gamma(t)$  to Eq.5:

$$\begin{aligned} y_S &= \mathcal{F}(z_t; \Theta_\Phi) + \gamma(t)\mathcal{Z}(\mathcal{F}(z_t + z_S; \Theta_S; \theta_{z2})), \\ \gamma(t) &= 1 - a \cdot \frac{1}{1 + e^{-b \cdot (t-c)}}, \end{aligned} \quad (11)$$

where  $a, b, c$  are used for timestep-aware adjustment.

## 5 Experiments

VersaGen utilises SD as Foundation Generation Model. We conducted training and testing of VersaGen on COCO (Lin et al. 2014) and further assessed its performance on Sketchy (Sangkloy et al. 2016), a human free-hand sketch dataset. Detailed information about data processing, evaluation metrics hyperparameters employed in the experiments, and additional generated results are provided in the supplementary(Chen et al. 2024).

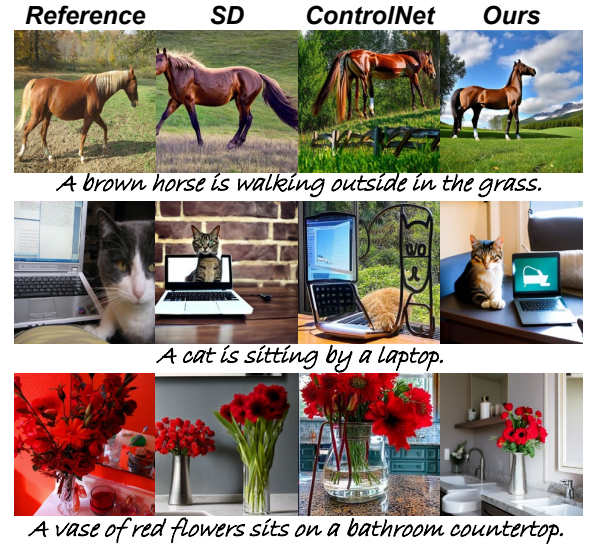


Figure 5: Visualisation of the reference images used in the human study alongside the user-submitted results generated by the three methods.

### 5.1 Qualitative Results

**Main results of VersaGen.** Fig. 1 shows our results with various drawing inputs: *single visual subject*, *multiple visual subjects* and *subjects+background*. The images generated by VersaGen not only closely align with the text descriptions but also maintain excellent consistency with the appearance of user drawings. VersaGen can automatically place objects from separate drawings coherently within a background, guided by the textual prompt.

**Comparison with different methods.** Fig. 4 presents qualitative results corresponding to two types of drawings (Row1 for Edge Map, Row2 & 3 for Free-Hand Drawing). As these comparative methods are unable to process object-level drawings, we ensure a fair comparison by using the automatically identified bounding box  $B''$  or relocated user drawings  $S'$  from VersaGen as their visual conditions. We can observe that: i) From column  $S$  to  $B''$  or  $S'$ , VersaGen effectively relocates the  $s_i$ s to positions that align coherently with the given textual prompt. Moreover, even in scenarios involving multiple visual subjects, VersaGen successfully avoids overlapping locations. ii) For the SD results, when the textual prompt involves multiple visual subjects, some entities are frequently omitted, e.g., The “bird” in Row2 and the “castle” in Row3 are missing. iii) The results from T2I-Adapter, ControlNet and UniControl adhere too closely to the provided drawings, often resulting in the generation of unrealistic details, such as the “bird” in Row2. iv) The results from GLIGEN and InstanceDiffusion indicate that employing layouts  $B''$  as a condition does not yield precise object control. The presence of a large number of irrelevant areas in the layout diminishes the positive influence of MOD on reducing object overlap, which leads to a decline in the quality of object generation, e.g., the “bird” and “pineapple” in Row2. v) Our proposed VersaGen consistently outperforms

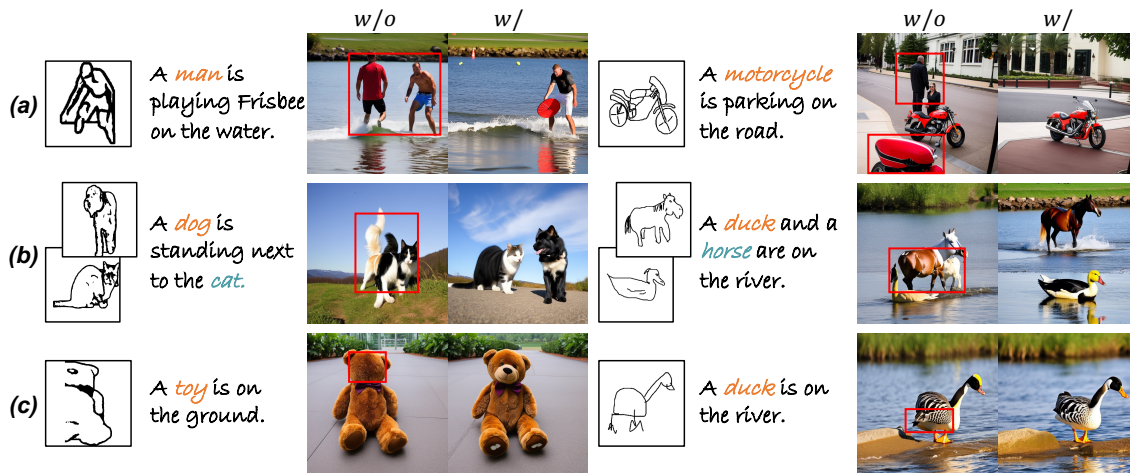


Figure 6: Ablative study of different designs in VersaGen: (a) Multimodal Conflict Resolver, (b) Multi-object Decoupling, (c) Adaptive Control Strength. Problematic regions are highlighted with  $\square$ .

other methods by generating high-quality images with precise semantics across a range of user drawing combinations, such as single visual subject, multiple visual subjects, and subjects+background.

## 5.2 Quantitative Results

Tab. 1 shows quantitative results on COCO and Sketchy.

For COCO results, we can see that: i) Compared to SD, other methods achieved lower FID scores due to the additional drawings or layouts, enhancing the spatial alignment between the generated images and the ground truth. ii) Our VersaGen achieves better FID metrics regardless of whether ground truth locations or automatically predicted locations are used, demonstrating the robustness and superiority of VersaGen. iii) SD achieves the highest CLIP score, likely due to the conflict of multiple conditions or the limitations in edge map extraction techniques. VersaGen addresses this by substituting the latter with high-quality user drawings, indicating the effectiveness of our proposed strategies during inference. iv) VersaGen achieves the highest DINO score with different setting, which indicates the generated images are closest to the real images in visual features, further demonstrating the utility of our method. v) InstanceDiffusion excels in object semantics, achieving the highest ACC but at the cost of local object details. In contrast, using sketches as conditions better balances local details with global semantics.

For Sketchy, visual subject drawings are human free-hand sketches, which are less precise and more abstract than edge maps, reflecting real-world conditions where most users are amateurs and may struggle to produce high-quality drawings quickly. Despite these challenges, our method achieves the highest performance on CLIP score, FID, and DINO score on Sketchy, demonstrating VersaGen’s effectiveness in real-world scenarios.

## 5.3 User Study

To further assess VersaGen’s effectiveness, we conducted a user study with the same setup as described in Sec.3. The

study presented user-submitted results from three methods (Fig. 5), revealing that: i) all methods perform well for simple prompts with a single entity; ii) however, for complex prompts with multiple entities, both SD and ControlNet often miss entities, and strict adherence to input conditions by ControlNet can degrade output quality. We surveyed 50 participants, asking which method balances ease of use with reliable generation capability. Results show 48% favor VersaGen, underscoring its practical application potential.

## 5.4 Ablation Study

To enhance the practicality of VersaGen in real-world applications, we propose three inference-phase strategies: Multimodal Conflict Resolver (MCR), Multi-object Decoupling (MoD), and Adaptive Control Strength (ACS). We conduct an ablation study to thoroughly evaluate the impact and effectiveness of these strategies. i) As shown in Fig. 6(a), the absence of MCR results in inconsistencies between the textual prompt and the generated images; for instance, the prompt “A man” yields an image of two men. ii) Without MoD, overlapping multiple objects leads to poor-quality generations, as depicted in Fig. 6(b), such as a creature with a horse’s body and a duck’s head. Involving MoD effectively resolves this issue. iii) Fig. 6(c) demonstrates that with ACS, even user-provided low-quality drawings, VersaGen can also generate high-quality images with details.

## 6 Conclusion

In this paper, we present VersaGen, a generative AI agent that addresses the challenges of incapability and inflexibility in text-to-image synthesis by providing users with versatile visual control options. By allowing users to choose the level of control that best suits their needs and preferences, VersaGen transforms the creative experience into a more engaging and fulfilling one. We hope our work inspires further research in developing user-centric generative AI solutions that prioritise flexibility, accessibility, and are ultimately people-facing.

## Acknowledgments

This research was supported by the Science and Technology Innovation 2030 – Major Projects (Grant No. 2021ZD0200600), the Hainan Provincial Natural Science Foundation of China (Grant No. 624LALH008), and the Program for Youth Innovative Research Team of BUPT (Grant No. 2023YQTD02). The authors sincerely appreciate the support and opportunities provided by these projects, which were instrumental in the successful completion of this research.

## References

- Aubert, M.; Brumm, A.; Ramli, M.; Sutikna, T.; Saptomo, E. W.; Hakim, B.; Morwood, M. J.; van den Bergh, G. D.; Kinsley, L.; and Dosseto, A. 2014. Pleistocene cave art from Sulawesi, Indonesia. In *Nature*.
- Balaji, Y.; Nah, S.; Huang, X.; Vahdat, A.; Song, J.; Zhang, Q.; Kreis, K.; Aittala, M.; Aila, T.; Laine, S.; et al. 2022. ediff-i: Text-to-image diffusion models with an ensemble of expert denoisers. In *arXiv preprint arXiv:2211.01324*.
- 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 *ICCV*.
- Chang, H.; Zhang, H.; Barber, J.; Maschinot, A.; Lezama, J.; Jiang, L.; Yang, M.-H.; Murphy, K.; Freeman, W. T.; Rubinstein, M.; et al. 2023. Muse: Text-to-image generation via masked generative transformers. In *arXiv preprint arXiv:2301.00704*.
- Chefer, H.; Alaluf, Y.; Vinker, Y.; Wolf, L.; and Cohen-Or, D. 2023. Attend-and-excite: Attention-based semantic guidance for text-to-image diffusion models. In *ACM TOG*.
- Chen, M.; Laina, I.; and Vedaldi, A. 2024. Training-free layout control with cross-attention guidance. In *WACV*.
- Chen, Z.; Yang, L.; Qi, Y.; Zhang, H.; Pang, K.; Li, K.; and Song, Y.-Z. 2024. VersaGen: Unleashing Versatile Visual Control for Text-to-Image Synthesis. In *arXiv preprint arXiv:2412.11594*.
- Eitz, M.; Hays, J.; and Alexa, M. 2012. How do humans sketch objects? In *ACM TOG*.
- Feng, W.; He, X.; Fu, T.-J.; Jampani, V.; Akula, A.; Narayana, P.; Basu, S.; Wang, X. E.; and Wang, W. Y. 2022. Training-free structured diffusion guidance for compositional text-to-image synthesis. In *arXiv preprint arXiv:2212.05032*.
- Gafni, O.; Polyak, A.; Ashual, O.; Sheynin, S.; Parikh, D.; and Taigman, Y. 2022. Make-a-scene: Scene-based text-to-image generation with human priors. In *ECCV*.
- Gal, R.; Alaluf, Y.; Atzmon, Y.; Patashnik, O.; Bermano, A. H.; Chechik, G.; and Cohen-Or, D. 2022. An image is worth one word: Personalizing text-to-image generation using textual inversion. In *arXiv preprint arXiv:2208.01618*.
- Gani, H.; Bhat, S. F.; Naseer, M.; Khan, S.; and Wonka, P. 2023. Llm blueprint: Enabling text-to-image generation with complex and detailed prompts. In *arXiv preprint arXiv:2310.10640*.
- Ge, S.; Park, T.; Zhu, J.-Y.; and Huang, J.-B. 2023. Expressive text-to-image generation with rich text. In *ICCV*.
- Gombrich, E. H.; and Gombrich, E. 1995. *The story of art*. Phaidon Press.
- Gong, B.; Huang, S.; Feng, Y.; Zhang, S.; Li, Y.; and Liu, Y. 2023. Check, Locate, Rectify: A Training-Free Layout Calibration System for Text-to-Image Generation. In *arXiv preprint arXiv:2311.15773*.
- 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 *ICLR*.
- Hoffmann, D. L.; Standish, C. D.; García-Diez, M.; Pettitt, P. B.; Milton, J. A.; Zilhão, J.; Alcolea-González, J. J.; Cantalejo-Duarte, P.; Collado, H.; De Balbín, R.; et al. 2018. U-Th dating of carbonate crusts reveals Neandertal origin of Iberian cave art. In *Science*.
- Jongejan, J.; Rowley, H.; Kawashima, T.; Kim, J.; and Fox-Gieg, N. 2017. The quick, draw! dataset.
- Ju, X.; Zeng, A.; Zhao, C.; Wang, J.; Zhang, L.; and Xu, Q. 2023. HumanSD: A native skeleton-guided diffusion model for human image generation. In *ICCV*.
- Karmiloff-Smith, A. 1990. Constraints on representational change: Evidence from children’s drawing. In *Cognition*.
- Kim, G.; Kwon, T.; and Ye, J. C. 2022. Diffusionclip: Text-guided diffusion models for robust image manipulation. In *CVPR*.
- Lee, K.; Liu, H.; Ryu, M.; Watkins, O.; Du, Y.; Boutilier, C.; Abbeel, P.; Ghavamzadeh, M.; and Gu, S. S. 2023. Aligning text-to-image models using human feedback. In *arXiv preprint arXiv:2302.12192*.
- Li, Y.; Keuper, M.; Zhang, D.; and Khoreva, A. 2023a. Divide & bind your attention for improved generative semantic nursing. In *arXiv preprint arXiv:2307.10864*.
- Li, Y.; Liu, H.; Wu, Q.; Mu, F.; Yang, J.; Gao, J.; Li, C.; and Lee, Y. J. 2023b. Gligen: Open-set grounded text-to-image generation. In *CVPR*.
- Lin, T.-Y.; Maire, M.; Belongie, S.; Hays, J.; Perona, P.; Ramanan, D.; Dollár, P.; and Zitnick, C. L. 2014. Microsoft coco: Common objects in context. In *ECCV*.
- Liu, V.; and Chilton, L. B. 2022. Design guidelines for prompt engineering text-to-image generative models. In *CHI*.
- Ma, J.; Liang, J.; Chen, C.; and Lu, H. 2023. Subject-diffusion: Open domain personalized text-to-image generation without test-time fine-tuning. In *arXiv preprint arXiv:2307.11410*.
- Marcus, G.; Davis, E.; and Aaronson, S. 2022. A very preliminary analysis of DALL-E 2. In *arXiv preprint arXiv:2204.13807*.
- 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 *AAAI*.
- Otsu, N.; et al. 1975. A threshold selection method from gray-level histograms. In *Automatica*.
- Patashnik, O.; Garibi, D.; Azuri, I.; Averbuch-Elor, H.; and Cohen-Or, D. 2023. Localizing object-level shape variations with text-to-image diffusion models. In *ICCV*.

Qin, C.; Zhang, S.; Yu, N.; Feng, Y.; Yang, X.; Zhou, Y.; Wang, H.; Niebles, J. C.; Xiong, C.; Savarese, S.; et al. 2024. UniControl: A Unified Diffusion Model for Controllable Visual Generation In the Wild. In *NeurIPS*.

Ramesh, A.; Pavlov, M.; Goh, G.; Gray, S.; Voss, C.; Radford, A.; Chen, M.; and Sutskever, I. 2021. Zero-shot text-to-image generation. In *ICML*.

Rassin, R.; Hirsch, E.; Glickman, D.; Ravfogel, S.; Goldberg, Y.; and Chechik, G. 2024. Linguistic binding in diffusion models: Enhancing attribute correspondence through attention map alignment. In *NeurIPS*.

Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; and Ommer, B. 2022. High-resolution image synthesis with latent diffusion models. In *CVPR*.

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. In *NeurIPS*.

Sangkloy, P.; Burnell, N.; Ham, C.; and Hays, J. 2016. The sketchy database: learning to retrieve badly drawn bunnies. In *ACM TOG*.

Wang, X.; Darrell, T.; Rambhatla, S. S.; Girdhar, R.; and Misra, I. 2024. Instancediffusion: Instance-level control for image generation. In *CVPR*, 6232–6242.

Wang, Z.; Sha, Z.; Ding, Z.; Wang, Y.; and Tu, Z. 2023. TokenCompose: Grounding Diffusion with Token-level Supervision. In *arXiv preprint arXiv:2312.03626*.

Wang, Z. J.; Montoya, E.; Munechika, D.; Yang, H.; Hoover, B.; and Chau, D. H. 2022. Diffusiondb: A large-scale prompt gallery dataset for text-to-image generative models. In *arXiv preprint arXiv:2210.14896*.

Wu, Q.; Liu, Y.; Zhao, H.; Bui, T.; Lin, Z.; Zhang, Y.; and Chang, S. 2023. Harnessing the spatial-temporal attention of diffusion models for high-fidelity text-to-image synthesis. In *ICCV*.

Xu, J.; Liu, X.; Wu, Y.; Tong, Y.; Li, Q.; Ding, M.; Tang, J.; and Dong, Y. 2024. Imagereward: Learning and evaluating human preferences for text-to-image generation. In *NeurIPS*.

Zhang, K.; Li, Y.; Zuo, W.; Zhang, L.; Van Gool, L.; and Timofte, R. 2021. Plug-and-play image restoration with deep denoiser prior. In *TPAMI*.

Zhang, L.; Rao, A.; and Agrawala, M. 2023. Adding conditional control to text-to-image diffusion models. In *ICCV*.