

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 DOES FLUX ALREADY KNOW HOW TO PERFORM PHYSICALLY PLAUSIBLE IMAGE COMPOSITION?

Anonymous authors

Paper under double-blind review



Figure 1: Showcase of our training-free image composition method, **SHINE**. This gallery highlights **SHINE**'s ability to seamlessly integrate subjects into complex scenes, including **low-light conditions, intricate shadows, and water reflections**.

ABSTRACT

Image composition aims to seamlessly insert a user-specified object into a new scene, but existing models struggle with complex lighting (e.g., accurate shadows, water reflections) and diverse, high-resolution inputs. Modern text-to-image diffusion models (e.g., SD3.5, FLUX) already encode essential physical and resolution priors, yet lack a framework to unleash them without resorting to latent inversion, which often locks object poses into contextually inappropriate orientations, or brittle attention surgery. We propose **SHINE**, a training-free framework for **S**eamless, **H**igh-fidelity **I**nsertion with **N**eutralized **E**rrors. **SHINE** introduces manifold-steered anchor loss, leveraging pretrained customization adapters (e.g., IP-Adapter) to guide latents for faithful subject representation while preserving background integrity. Degradation-suppression guidance and adaptive background blending are proposed to further eliminate low-quality outputs and visible seams. To address the lack of rigorous benchmarks, we introduce *ComplexCompo*, featuring diverse resolutions and challenging conditions such as low lighting, strong illumination, intricate shadows, and reflective surfaces. Experiments on ComplexCompo and DreamEditBench show state-of-the-art performance on standard metrics (e.g., DINOv2) and human-aligned scores (e.g., DreamSim, ImageReward, VisionReward). Code and benchmark will be publicly available upon publication.

054

1 INTRODUCTION

055
 056 Image composition, which places a user-specified object into a new scene, is a demanding image
 057 editing task. Despite the breathtaking progress of multimodal foundation models (e.g., GPT-5 (Ope-
 058 nAI, 2025), Gemini-2.5 (Gemini2.5, 2025), SeedEdit/Doubao (Shi et al., 2024b), and Grok-4 (gro,
 059 2025)), these generic models still struggle with image composition. Typical failures include impre-
 060 cise object placement, inconsistent lighting, and the subject’s identity drift (see Fig. 2). These limita-
 061 tions indicate that, as of now, massive multimodal pre-training alone has not yet endowed them with
 062 sufficient compositional ability for this task. A natural response has been to train specialized models.
 063 Yet building large-scale, high-quality, multi-resolution triplet datasets (object, scene, composite) is
 064 prohibitively costly. As a result, most composition models are fine-tuned from base models (e.g.,
 065 FLUX.1-dev (Black Forest Labs, 2024a), FLUX.1-Fill (Black Forest Labs, 2024c), SDXL (Podell
 066 et al., 2024)) on synthetic data generated via inpainting or augmentations (Chen et al., 2024c; Yang
 067 et al., 2023; Song et al., 2023; Wang et al., 2025a; He et al., 2024).

068 These models, however, face two main limitations (see Fig. 6): **(i) Lighting realism.** They strug-
 069 gle to achieve natural composition under complex lighting conditions, such as accurate shadow
 070 generation or water reflections for the inserted subject. **(ii) Resolution rigidity.** They are tied to a
 071 fixed resolution, necessitating downsampling or cropping when applied to varied, high-resolution
 072 background images, which degrades generation quality. Notably, such issues are absent in the base
 073 models, implying that the underlying physical priors are present but are not effectively exploited by
 074 fine-tuned variants. The degradation largely stems from low-quality synthetic datasets, which inherit
 075 flaws from inpainting models that often mis-handle shadows and reflections, producing implausible
 076 edits, hallucinated content, or incomplete object removal (Yu et al., 2025b; Winter et al., 2025).

077 There have been prior **training-free** attempts to exploit the priors of text-to-image (T2I) models
 078 for advancing image composition, but they fall short for two main reasons. **(i) Inversion bottle-
 079 necks.** Most methods (Lu et al., 2023d; Pham et al., 2024; Yan et al., 2025; Li et al., 2024b) de-
 080 pend on accurate image inversion (Song et al., 2021; Lu et al., 2022; Mokady et al., 2023). In
 081 practice, inversion constrains the inserted object to the pose of its reference image, often resulting
 082 in contextually inappropriate orientations. Moreover, inversion is less effective for classifier-free
 083 guidance (CFG) distilled models (e.g., FLUX), where elevated inversion errors degrade identity
 084 preservation. **(ii) Fragile attention surgery.** Many training-free approaches rely on attention ma-
 085 nipulation (Lu et al., 2023d; Yan et al., 2025; Li et al., 2024b). While compatible with the joint
 086 self-attention in Multimodal Diffusion Transformers (MMDiT) (Peebles & Xie, 2023), these meth-
 087 ods inherit the instability and hyperparameter sensitivity (Lu et al., 2023d), limiting their robustness.

088 To bridge these gaps we present **SHINE**, a training-free framework for **Seamless, High-fidelity**
 089 **Insertion with Neutralized Errors** (see Fig. 1). SHINE comprises three innovations: **(i) Manifold-
 090 Steered Anchor (MSA) loss**, which leverages pretrained open-domain customization adapters (e.g.,
 091 IP-Adapter (Ye et al., 2023)) to steer noisy latents toward faithfully representing the reference subject
 092 while preserving the structural integrity of the background. **(ii) Degradation-Suppression Guid-
 093 ance (DSG)** that steers sampling away from low-quality distributions. **(iii) Adaptive Background
 094 Blending (ABB)** that eliminates visible seams along mask boundaries.

095 Existing benchmarks primarily comprise background images with a fixed resolution of 512×512
 096 pixels. To evaluate performance across diverse, high-resolution, and demanding scenarios, we in-
 097 troduce *ComplexCompo*, a benchmark that includes varied resolutions, both landscape and portrait
 098 orientations, and complex conditions such as low lighting, intense illumination, intricate shadows,
 099 and water reflections. Extensive experiments on ComplexCompo and DreamEditBench (Li et al.,
 100 2023b) demonstrate that SHINE achieves state-of-the-art (SOTA) performance, surpassing base-
 101 lines on standard metrics (e.g., DINOv2 (Quab et al., 2024)) and human-aligned metrics (e.g.,
 102 DreamSim (Fu et al., 2023), ImageReward (Xu et al., 2023), VisionReward (Xu et al., 2024)).

103

2 RELATED WORK

104 This section reviews prior work on image composition. A more comprehensive discussion, covering
 105 image composition, general image editing, and subject-driven generation, is offered in Appendix A.
 106 Classical image composition splits into sub-tasks (Niu et al., 2021) such as object placement (Azadi
 107 et al., 2020; Zhang et al., 2020a), blending (Wu et al., 2019; Zhang et al., 2020b), harmonization (Cao



Figure 2: Image composition from advanced multimodal models under three challenging conditions: backlighting, shadows, and water surfaces. Refer to Appendix H for prompt details.

et al., 2023; Lu et al., 2023b), and shadow generation (Hong et al., 2022; Sheng et al., 2021), typically handled by separate models. Diffusion models have shifted the field toward unified frameworks, either training-based or training-free. Training-based approaches fine-tune diffusion models with curated datasets, adding grounding layers, controllability signals, or identity-preserving supervision from image or video sets (Wang et al., 2025a; Chen et al., 2024c; Yang et al., 2023; Song et al., 2023; Lu et al., 2023c). However, they often bias model priors and struggle with complex lighting due to the lack of large-scale real-world triplets. Training-free approaches avoid retraining by manipulating inversion and attention during inference, enabling flexible test-time adaptation (Yan et al., 2025; Li et al., 2024b; 2023b; Lu et al., 2023d; Pham et al., 2024). Yet these methods remain fragile: strong injections preserve identity but fix unnatural poses, while weaker ones improve realism at the cost of fidelity, reflecting a core trade-off between identity preservation and natural composition.

3 METHOD

Image composition seeks to integrate a subject into a designated area of a background image while preserving the integrity of the surrounding scene. This process typically requires three inputs: (1) one or more reference images of the subject $\{x_1^{\text{subj}}, x_2^{\text{subj}}, \dots, x_n^{\text{subj}}\}$, (2) a background image x^{bg} , and (3) a user-provided mask M^{user} specifying the insertion region within the background.

Our framework is built on three core components: Manifold-Steered Anchor (MSA) loss, Degradation-Suppression Guidance (DSG), and Adaptive Background Blending (ABB). Importantly, the design is model-agnostic and requires only standard features of modern generative models: MSA loss assumes that the base model supports either personalization finetuning or provides access to a pretrained personalization adapter, DSG uses self-attention maps, and ABB relies on text-image cross-attention. These mild assumptions enable seamless integration into existing pipelines without architectural changes. We present main results with FLUX, while additional experiments on SDXL (Podell et al., 2024), SD3.5 (Esser et al., 2024), and PixArt (Chen et al.) are provided in Appendix E. The complete algorithm is shown in Algorithm 1.

3.1 NON-INVERSION LATENT PREPARATION

In training-free diffusion-based image composition (Lu et al., 2023d; Pham et al., 2024; Yan et al., 2025; Li et al., 2024b), it is common to start from a noisy latent. Existing training-free frameworks typically rely on image inversion, where the initial noisy latent is constructed by copying the inverted latent of the subject image into a designated region of the background image’s inverted latent.

However, this copy-paste strategy constrains the inserted object to the exact pose of its reference image, often leading to contextually inappropriate orientations in the composed result. Moreover, inversion is suboptimal for CFG-distilled models (e.g., FLUX), as it introduces higher inversion errors that compromise subject identity preservation.

To address these limitations, we abandon inversion and instead perform a one-step forward diffusion to obtain the noisy latent. As illustrated in Fig. 3(a), we use a vision-language model (VLM) (Xue et al., 2024; Chen et al., 2024d; Liu et al., 2024) to caption the subject image and leverage this caption, along with an image inpainting model (Li et al., 2024c; Ju et al., 2024; Zhuang et al., 2024; Black Forest Labs, 2024b), to generate the image to which the subject is attached, denoted as x^{init} .




Figure 3: **Overview of the proposed framework.** (a) The noisy latent is created by inpainting the background with a VLM-derived object description, then adding Gaussian noise. (b) Manifold-Steered Anchor (MSA) loss guides noisy latents toward faithfully capturing the reference subject (red arrow), while preserving the structural integrity of the background. Concretely, it enforces that the prediction of the optimized latent z_t^* on the adapter-augmented model’s manifold remains close to the prediction of the original latent z_t on the base model’s manifold. (c) Degradation-Suppression Guidance (DSG) constructs a negative velocity pointing toward low-quality regions by blurring Q_{img} and, in a CFG-like manner, steers the trajectory away from this low-quality distribution.

The noisy latent is encoded in the VAE space as z^{init} and perturbed to timestep $t \leq T$ via one-step forward diffusion, following the flow matching formulation: $z_t = (1 - \sigma_t)z^{\text{init}} + \sigma_t \epsilon$, where $\epsilon \sim \mathcal{N}(0, 1)$.

3.2 MANIFOLD-STEERED ANCHOR LOSS

The Manifold-Steered Anchor (MSA) loss is designed to optimize the noisy latent z_t (from Sec. 3.1) during the denoising process, steering it toward a reference subject while preserving the structural integrity of the original image. The key intuition is to leverage the prior knowledge embedded in pretrained open-domain customization adapters (or alternatively, personalized LoRAs) such as IP-Adapter (Ye et al., 2023), PuLID (Guo et al., 2024), and InstantCharacter (Tao et al., 2025), to intervene directly in the diffusion trajectory. Specifically, the MSA loss is defined as:

$$\min_{z_t} \mathcal{L}_{\text{MSA}}(z_t) = \left\| v_{\theta+\Delta\theta}(z_t, t, c, z^{\text{subj}}) - \text{sg}[\tilde{v}_t] \right\|_2^2, \quad (1)$$

where $\tilde{v}_t \triangleq v_{\theta}(\tilde{z}_t, t, c)$ serves as a fixed anchor, preserving the structure of the background image at a given noise level t , with \tilde{z}_t held constant as the original noisy latent. $v_{\theta}(\cdot)$ denotes the velocity predicted by the frozen T2I model θ , while $v_{\theta+\Delta\theta}(\cdot)$ represents the velocity predicted by the a T2I model augmented with an adapter $\Delta\theta$. z^{subj} is the latent of the subject image. The text prompt c is from the VLM’s description of x^{init} , and $\text{sg}[\cdot]$ indicates the stop-gradient operation.

The MSA loss is motivated by the observation that optimizing a latent representation against a frozen generative model implicitly projects the latent onto the model’s learned data manifold (Meng et al., 2021; Kim et al., 2022; Graikos et al., 2022; Feng et al., 2023). The generator serves as an implicit prior, guiding gradient descent toward the manifold’s basin of attraction.

For instance, when a generative model $G(\mathbf{w})$ is trained solely on cat images, its outputs are confined to the cat-image manifold. Thus, approximating a dog image x_{dog} by solving $\min_{\mathbf{w}} \|G(\mathbf{w}) - x_{\text{dog}}\|_2^2$ yields $G(\mathbf{w}^*)$ that remains a cat image, but with structural features (e.g., pose or outline) aligned to x_{dog} . The result is the projection of the dog image onto the cat manifold, not a genuine dog image.

Analogously, MSA loss is designed to achieve two goals simultaneously. (1) It seeks an optimized noisy latent z_t^* that remains within the manifold of the adapter-augmented model when conditioned on the subject z^{subj} . (2) It encourages the adapter’s prediction on this latent z_t^* to align with the base model’s prediction on the original latent z_t , i.e., $v_{\theta+\Delta\theta}(z_t^*, t, c, z^{\text{subj}}) \approx v_{\theta}(z_t, t, c)$ (see Fig. 3(b)).

Figure 4: *Left*: Robustness of FLUX. *Right*: Impacts of blurring different features in FLUX.

Since the velocity prediction of a T2I model on a noisy latent z_t can also be interpreted as a coarse estimate of the clean image that encodes essential structural information (Zheng et al., 2023), this alignment preserves the spatial layout and background details inherited from the original image.

The gradient of \mathcal{L}_{MSA} with respect to z_t is:

$$\nabla_{z_t} \mathcal{L}_{\text{MSA}}(z_t) = 2 \left(v_{\theta+\Delta\theta}(z_t, t, c, z^{\text{subj}}) - \text{sg}[\tilde{v}_t] \right) \frac{\partial v_{\theta+\Delta\theta}(z_t, t, c)}{\partial z_t}. \quad (2)$$

The Jacobian term necessitates backpropagation through the MMDiT, which is computationally expensive. However, this scenario is analogous to Score Distillation Sampling (SDS) (Poole et al.), where research shows that omitting the Jacobian term yields an effective gradient for optimization with diffusion models. Thus, we adopt the same strategy for optimization.

3.3 DEGRADATION-SUPPRESSION GUIDANCE

MSA loss effectively facilitates the insertion of reference objects. However, due to the inherent stochasticity of the denoising and optimization process, the results sometimes suffer from degraded visual quality, manifesting as oversaturated colors and reduced identity consistency (see Fig. 5). To address this, we introduce Degradation-Suppression Guidance (DSG), inspired by negative prompting (Schramowski et al., 2023), defined as:

$$v_t^{\text{dsg}} = v_{\theta+\Delta\theta}(z_t, t, c, z^{\text{subj}}) + \eta(v_{\theta+\Delta\theta}(z_t, t, c, z^{\text{subj}}) - v_{\theta+\Delta\theta}^{\text{neg}}(z_t, t, c, z^{\text{subj}})), \quad (3)$$

where $v_{\theta+\Delta\theta}^{\text{neg}}$ denotes a negative velocity prediction that guides the generation toward low-quality regions. A key challenge is the design of a meaningful negative velocity prediction $v_{\theta+\Delta\theta}^{\text{neg}}$ within MMDiT-based architectures. In our experiments with FLUX, we observed that using nonsensical text prompts or explicit negative prompts fails to introduce degradation. The generated images remain high-fidelity (Fig. 4(a)), suggesting that text-based negative prompting is ineffective for FLUX.

In our setting, the ideal negative velocity $v_{\theta+\Delta\theta}^{\text{neg}}$ should target directions that preserve the semantic content and spatial layout while lowering perceptual quality. To achieve this, we investigate whether we can manipulate FLUX’s internal representations to construct such a targeted degradation signal.

In FLUX, both multi-stream and single-stream blocks compute joint self-attention over concatenated text and image tokens as follows:

$$h = \text{softmax} \left([Q_{\text{txt}}, Q_{\text{img}}][K_{\text{txt}}, K_{\text{img}}]^T / \sqrt{d_k} \right) \cdot [V_{\text{txt}}, V_{\text{img}}], \quad (4)$$

where $[Q_{\text{txt}}, Q_{\text{img}}]$ represents the concatenation of text and image queries, and similarly for keys and values. To identify an effective manipulation strategy, we systematically perturb different components in the attention mechanism (i.e., Q_{txt} , K_{txt} , V_{txt} , Q_{img} , K_{img} and V_{img}) and observe their impact on generation quality.

As shown in Fig. 4(b), our findings are as follows:

1. Blurring Q_{txt} , K_{txt} , or V_{txt} has negligible impact on semantic fidelity and visual quality.
2. Blurring V_{img} severely disrupts the output distribution, leading to unintelligible images.
3. Blurring K_{img} moderately impacts quality, while the image remains visually acceptable.
4. Blurring Q_{img} yields pronounced degradations while preserving structural integrity, making it the most effective lever for constructing a negative velocity.




Figure 5: Comparison of rectangular-mask blending and Adaptive Background Blending (ABB). Boundary regions (pink dashed boxes) are enlarged for clarity. Zoom in for details.

Algorithm 1 The Image Composition Process of SHINE.

Input: A background latent $z^{\text{bg}} = \text{VAE}(x^{\text{bg}})$, a subject latent $z^{\text{subj}} = \text{VAE}(x^{\text{subj}})$, a inpainted latent $z^{\text{init}} = \text{VAE}(\text{Inpainting} [x^{\text{bg}}, M^{\text{user}}, \text{VLM}(x^{\text{subj}})])$, a user mask M^{user} .
Output: The composition latent z_0 .

```

1:  $z_{t_1} \leftarrow (1 - \sigma_{t_1})z^{\text{init}} + \sigma_{t_1}\epsilon$ , where  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ 
2: for  $t = t_1, \dots, 0$  do
3:   // Manifold-Steered Anchor (MSA) Optimization
4:   if  $t > \tau$  then
5:      $\tilde{v}_t \leftarrow v_{\theta}(z_t, t, c)$ 
6:     for  $j = 1, \dots, k$  do
7:        $z_t \leftarrow z_t - \alpha \cdot M^{\text{user}} \odot \nabla_{z_t} \|v_{\theta+\Delta\theta}(z_t, t, c, z^{\text{subj}}) - \text{sg}[\tilde{v}_t]\|_2^2$ 
8:     end for
9:   end if
10:  // Degradation-Suppression Guidance (DSG)
11:   $v_t, A_t \leftarrow v_{\theta+\Delta\theta}(z_t, t, c, z^{\text{subj}})$ 
12:   $v_t^{\text{dsg}} \leftarrow v_t + \eta (v_t - v_{\theta+\Delta\theta}^{\text{neg}}(z_t, t, c, z^{\text{subj}}))$ 
13:   $z_{t-1} \leftarrow z_t + (\sigma_{t-1} - \sigma_t)v_t^{\text{dsg}}$ 
14:   $z_{t-1}^{\text{bg}} \leftarrow (1 - \sigma_{t-1})z^{\text{bg}} + \sigma_{t-1}\epsilon$ , where  $\epsilon \sim \mathcal{N}(\mathbf{0}, \mathbf{1})$ 
15:  // Adaptive Background Blending (ABB)
16:   $M^{\text{attn}} \leftarrow \text{MaxConnectedComponent}(\text{Dilate}(\mathbb{1}(A_t \geq \gamma)))$ 
17:   $\hat{M} \leftarrow \mathbb{1}\{t > \tau\} M^{\text{attn}} + \mathbb{1}\{t \leq \tau\} M^{\text{user}}$ 
18:   $z_{t-1} \leftarrow \hat{M} \odot z_{t-1} + (1 - \hat{M}) \odot z_{t-1}^{\text{bg}}$ 
19: end for
20: return  $z_0$ 

```

Based on these insights, we construct the negative velocity prediction $v_{\theta+\Delta\theta}^{\text{neg}}$ in Eqn. 3 by blurring Q_{img} within FLUX (see Fig. 3(c)). Moreover, we show that blurring Q_{img} is mathematically equivalent to blurring the self-attention weights, whereas blurring K_{img} or V_{img} is not (see Appendix C for the proof). This equivalence is consistent with the fact that attenuating self-attention activations suppresses informative interactions and thus degrades image quality (Lu et al., 2024).

3.4 ADAPTIVE BACKGROUND BLENDING

Previous methods typically rely on the user-provided mask M^{user} to preserve the background during each denoising step, blending as $z'_t = M^{\text{user}} \odot z_t + (1 - M^{\text{user}}) \odot z_t^{\text{bg}}$, but this often introduces visible seams along mask boundaries (see the first row of Fig. 5).

To address this limitation, we propose Adaptive Background Blending (ABB), defined as

$$z'_t = \hat{M} \odot z_t + (1 - \hat{M}) \odot z_t^{\text{bg}}, \quad \hat{M} = \mathbb{1}\{t > \tau\} \mathcal{D}(M^{\text{attn}}) + \mathbb{1}\{t \leq \tau\} M^{\text{user}}, \quad (5)$$

where M^{user} is the user mask, while M^{attn} is derived by binarizing the cross-attention maps corresponding to subject tokens. These maps can be obtained by either averaging across layers or select-




Figure 6: Qualitative comparison of our method with multiple baselines in challenging scenarios, **drawn from our benchmark dataset**. More qualitative comparisons are available in Appendix O.

ing the most informative layer via a lightweight analysis (details in Appendix D). The operator $\mathcal{D}(\cdot)$ performs dilation and extracts the largest connected component, ensuring robustness to noise.

Compared to M_{user} , M^{attn} is more spatially precise, particularly for elongated or irregularly shaped objects that do not fully occupy a rectangular region. As illustrated in the second row of Fig. 5, our method produces smoother transitions by replacing the rigid user mask with the semantically guided mask. This refinement better preserves the surrounding scene, enabling seamless integration between generated content and the original background. However, applying this method throughout the denoising process may truncate object shadows or reflections. Through empirical evaluation, we find that leveraging M^{attn} during the initial denoising steps ($t > \tau$) sufficiently mitigates visible seams along mask boundaries, ensuring high-fidelity scene coherence.

4 EXPERIMENTS

4.1 EXPERIMENTAL SETUP

Benchmark. Current benchmarks primarily consist of background images with a fixed resolution of 512×512 pixels. To assess performance across diverse, high-resolution, and complex scenarios, we introduce *ComplexCompo*, a benchmark built upon DreamEditBench (Li et al., 2023b). DreamEditBench includes 220 (subject, background, bounding box) pairs designed for 512×512 resolution. In contrast, ComplexCompo features 300 composition pairs with varying resolutions, encompassing both landscape and portrait orientations, and incorporates challenging conditions such as low lighting, intense illumination, intricate shadows, and water reflections. The background images are sourced from OpenImage (Kuznetsova et al., 2020). Further details are provided in Appendix G.

Metrics. Previous methods primarily adapt CLIP-I (Radford et al., 2021) and DINOv2 (Oquab et al., 2024) to assess subject identity consistency. However, these features capture high-level semantic information that may not fully align with human perception of finer-grained attributes. Thus, we further incorporate instance retrieval features (IRF) from (Shao & Cui, 2022) and DreamSim (Fu et al., 2023), which better align with human judgments. An analysis of identity consistency metrics is provided in Appendix I. For overall image quality, we use ImageReward (IR) (Xu et al., 2023)

Table 1: Comparison of composition performance across two benchmarks. The best result in each column is highlighted in **bold**, while the second-best is underlined. Metrics shown in **pink** are those specifically trained to better align with human preferences. Abbreviations: IRF= Instance Retrieval Features; IR = ImageReward; VR = VisionReward; URE = UnifiedReward-Edit-qwen3vl-8b.

Bench	Method	Training -Free	Base Model	External Model	Subject Identity Consistency			Background		Image Quality				
					CLIP-I \uparrow	DINOv2 \uparrow	IRF \uparrow	DreamSim \downarrow	LPIPS \downarrow	SSIM \uparrow	IR \uparrow	VR \uparrow	HPS \uparrow	URE \uparrow
DreamEdit-Bench (220)	Flux.1 Fill (Black Forest Labs, 2024b)	✗	FLUX	-	0.7328	0.6745	0.5754	0.5233	0.0166	0.9076	0.5577	3.5997	8.6432	21.5812
	MADD (He et al., 2024)	✗	SD	DINO	0.7118	0.6279	0.4333	0.5810	0.0604	0.8182	-0.2545	2.7011	1.2443	13.8148
	ObjectStitch (Song et al., 2023)	✗	SD	VIT	0.7567	0.6930	0.5525	0.5093	0.0190	0.8316	0.0791	3.2416	7.4529	19.1886
	DreamCom (Lu et al., 2023c)	✗	SD	LoRA	0.7414	0.6749	0.5597	0.5626	0.0200	0.8283	0.1873	3.5053	5.9324	19.9296
	AnyDoor (Chen et al., 2024c)	✗	SD	DINO	0.8183	0.7283	<u>0.7714</u>	0.3764	0.0251	0.8894	0.4511	3.3946	8.4867	19.0989
	UniCombine (Wang et al., 2025a)	✗	FLUX	LoRA	0.8058	0.7332	0.7579	0.3984	0.0050	0.9397	0.4565	3.6108	8.8415	21.7080
	PBE (Yang et al., 2023)	✗	SD	-	0.7742	0.7040	0.5845	0.4985	0.0197	0.8287	0.2083	3.3482	8.3789	20.2137
	TIGIC (Li et al., 2024b)	✓	SD	-	0.7226	0.6718	0.4711	0.6108	0.0584	0.8153	-0.1332	2.9873	5.2676	17.1000
	TALE (Pham et al., 2024)	✓	SD	-	0.7329	0.6604	0.5007	0.6176	0.0392	0.8251	-0.1502	3.1349	6.3773	18.0784
	TF-ICON (Lu et al., 2023d)	✓	SD	-	0.7479	0.6865	0.5179	0.5441	0.0582	0.8111	0.0816	3.2823	7.2643	18.2716
ComplexCompo (300)	DreamEdit (Li et al., 2023b)	✓	SD	LoRA, VIT	0.7703	0.7151	0.6147	0.5047	0.0140	0.9775	0.1744	3.1775	6.0250	15.7636
	EEdit (Yan et al., 2025)	✓	FLUX	-	0.6998	0.6590	0.4438	0.6160	0.0039	<u>0.9475</u>	0.0216	3.3606	6.6689	19.5603
	Ours-Adapter	✓	FLUX	Adapter	0.8086	0.7415	0.7702	<u>0.3730</u>	0.0236	0.8959	<u>0.5709</u>	3.6234	8.8861	22.0182
	Ours-LoRA	✓	FLUX	LoRA	<u>0.8125</u>	0.7452	0.7900	<u>0.3577</u>	0.0271	0.8847	0.5900	<u>3.6161</u>	<u>8.8688</u>	21.9421
	Flux.1 Fill (Black Forest Labs, 2024b)	✗	FLUX	-	0.7108	0.6475	0.5466	0.6018	0.0232	0.7442	0.4088	3.5737	8.7376	19.7712
	MADD (He et al., 2024)	✗	SD	DINO	0.6780	0.5903	0.3638	0.5979	0.0781	0.5658	-0.0088	2.6582	5.9673	13.0567
	ObjectStitch (Song et al., 2023)	✗	SD	VIT	0.608	0.7077	0.5513	0.4717	0.0388	0.6357	0.2492	3.4411	8.8389	18.8283
	DreamCom (Lu et al., 2023c)	✗	SD	LoRA	0.648	0.5602	0.2788	0.8192	0.0389	0.6342	-0.0778	3.4409	8.8984	18.6143
	AniDoor (Chen et al., 2024c)	✗	SD	DINO	0.7982	0.7052	0.7319	0.4493	0.0299	0.7262	0.3808	3.2787	8.9760	18.3550
	UniCombine (Wang et al., 2025a)	✗	FLUX	LoRA	<u>0.7361</u>	0.6552	0.5380	0.5662	0.0237	0.7077	0.2470	3.5454	8.8999	19.8529
	PBE (Yang et al., 2023)	✗	SD	-	0.7537	0.6802	0.5189	0.5187	0.0397	0.6521	0.2138	3.4310	8.5923	18.9507
	TIGIC (Li et al., 2024b)	✓	SD	-	0.6913	0.6329	0.3848	0.6549	0.0929	0.6228	-0.131	2.8899	7.6630	16.4301
	TALE (Pham et al., 2024)	✓	SD	-	0.6816	0.6151	0.3799	0.6773	0.059	0.6334	0.0783	3.4498	8.7351	18.7567
	TF-ICON (Lu et al., 2023d)	✓	SD	-	0.6987	0.6435	0.4167	0.6030	0.0815	0.6216	0.1798	3.4323	9.3258	18.2366
	DreamEdit (Li et al., 2023b)	✓	LoRA, VIT	-	0.7314	0.6722	0.5069	0.5670	0.0468	0.7201	0.1212	3.2531	8.0445	15.9334
	EEdit (Yan et al., 2025)	✓	FLUX	-	0.6713	0.6153	0.3797	0.6211	0.0226	0.7107	0.1433	3.5009	8.7835	19.7348
	Ours-Adapter	✓	FLUX	Adapter	0.7721	<u>0.7107</u>	0.6764	0.4294	0.0404	0.7889	<u>0.4090</u>	3.6020	<u>9.6485</u>	20.7349
	Ours-LoRA	✓	FLUX	LoRA	0.7999	0.7384	<u>0.7659</u>	<u>0.3542</u>	0.0430	0.7634	<u>0.4246</u>	<u>3.5951</u>	9.8418	21.0326




Figure 7: Qualitative ablation study comparing different variants of our framework.

and VisionReward (VR) (Xu et al., 2024), fine-grained reward models that more accurately reflect human preferences. To more comprehensively evaluate composition quality, we further include three UnifiedReward variants (Wang et al., 2025b;c) and HPSv3 (Ma et al., 2025b) in Appendix L. Background consistency is measured using LPIPS (Zhang et al., 2018) and SSIM (Wang et al., 2004).

Implementation Details. In our experiment, we used FLUX.1-dev, a 12B-parameter flow matching model, as the base model, combined with InstantCharacter (Tao et al., 2025) as the adapter (Additional results on SDXL, SD3.5, and PixArt are presented in Appendix E). Our approach also supports per-concept LoRA (Hu et al.), which requires test-time tuning (Ruiz et al., 2023) but delivers superior identity consistency compared to an open-domain adapter, making it ideal for scenarios demanding precise identity preservation. The denoising schedule consists of 20 steps, with the inpainted image perturbed to timestep 15 and denoising initiated from that point. We use Flux.1 Fill (Black Forest Labs, 2024b) as the inpainting model and BLIP-3 (Xue et al., 2024) as the VLM. Additional details and hyperparameters are provided in Appendix J.

4.2 EXPERIMENTAL RESULTS

We compare our method with two categories of baselines: **(1) Training-based methods (6 in total):** UniCombine (Wang et al., 2025a), AnyDoor (Chen et al., 2024c), Paint by Example (PBE) (Yang et al., 2023), ObjectStitch (Song et al., 2023), MADD (He et al., 2024), and DreamCom (Lu et al., 2023c); **(2) Training-free methods (5 in total):** EEdit (Yan et al., 2025), TIGIC (Li et al., 2024b), DreamEdit (Li et al., 2023b), TF-ICON (Lu et al., 2023d), and TALE (Pham et al., 2024).

As shown in Tab. 1, both variants of our method surpass all baselines on DreamEditBench across human preference aligned metrics (i.e., DreamSim, IR, VR), which are the most critical indicators of quality. For background-related metrics, all methods achieve comparable results, with differences so small they are imperceptible to the human eye. On the more challenging ComplexCompo dataset, which includes non-square resolutions and intricate scenes, most methods experience a notable per-



Figure 8: Composition inherit erroneous colors if the inpainting prompt specifies an incorrect color.

Table 2: Ablation study examining the impact of key components on DreamEditBench.

Method	MSA	DSG	ABB	Subject Identity Consistency			Background		Image Quality		
				CLIP-I	DINOv2	IRF	DreamSim	LPIPS	SSIM	IR	VR
Config A	✗	✗	✗	0.7328	0.6745	0.5754	0.5233	0.0166	0.9076	0.5577	3.5997
Config B	✓	✗	✗	0.7814	0.7204	0.7414	0.3951	0.0172	0.9075	0.5455	3.5952
Config C	✗	✓	✗	0.7528	0.6941	0.6533	0.4436	0.0178	0.9038	0.5633	3.6130
Config D	✗	✗	✓	0.7421	0.6814	0.6158	0.5127	0.0210	0.9010	0.5595	3.6109
Config E	✗	✓	✓	0.7481	0.6987	0.6647	0.4317	0.0218	0.8971	0.5850	3.6277
Config F	✓	✗	✓	0.8084	0.7429	0.7609	0.3756	0.0231	0.8991	0.5459	3.6023
Config G	✓	✓	✗	0.8077	0.7375	0.7589	0.3762	0.0182	0.9037	0.5745	3.6191
Ours-Adapter	✓	✓	✓	0.8086	<u>0.7415</u>	0.7702	0.3730	0.0236	0.8959	0.5709	<u>3.6232</u>

454 performance drop, yet our approach consistently remains the top performer. From Fig. 6, it is evident
455 that while AnyDoor achieves high scores on many identity metrics, the model tends to copy and
456 paste the subject into the scene, resulting in unnatural compositions and lower image quality scores.
457 In contrast, our method excels at naturally composing objects in challenging conditions (e.g., low-
458 light settings, water surfaces, and scenes with complex shadows). Appendix F provides user study.

4.3 ABLATION STUDY

462 We validate our design choices through ablation (see Tab. 2 and Fig. 7). The results highlight three
463 key insights. First, MSA loss notably improves subject identity consistency. Second, DSG boosts
464 IR and VR scores by steering denoising away from low-quality regions. Finally, ABB effectively
465 suppresses visible seams along mask boundaries. While this improvement is readily apparent in
466 visual comparisons (Figs. 5, 7), it is less well captured by quantitative metrics, since LPIPS and
467 SSIM primarily assess structural similarity rather than perceptual smoothness.

5 CONCLUSION

471 We introduced SHINE, a training-free framework for seamless and high-fidelity image compo-
472 sition with pretrained T2I models. SHINE integrates Manifold-Steered Anchor Loss, Degradation-
473 Suppression Guidance, and Adaptive Background Blending to ensure precise subject placement and
474 artifact-free synthesis across diverse resolutions and lighting conditions. To enable rigorous eval-
475 uation, we proposed ComplexCompo, a benchmark for challenging composition scenarios. SHINE
476 achieves state-of-the-art results on both ComplexCompo and DreamEditBench.

477 **Limitations.** Our method reliably converges to the correct subject identity through MSA optimiza-
478 tion even when the inpainted subject’s appearance deviates substantially from the reference (see
479 Fig. 8(a)). However, when the inpainting prompt specifies an incorrect color, the final inpainted re-
480 sult tends to inherit and preserve this erroneous color (see Fig. 8(b)). On the other hand, the similarity
481 between the inserted object and the user-provided object depends on the quality of the customization
482 adapter used. As shown in Tab. 1, because LoRA performs test-time tuning for individual concepts, it
483 generates subjects that are more similar to the target than those produced by pretrained open-domain
484 customization adapters, resulting in higher subject identity consistency metrics in the composition.
485 While current customization adapters already perform well, the potential of our method will con-
486 tinue to improve as advancements are made in the field of open-domain customization adapters.

486 ETHICS STATEMENT
487488 Our framework provides an accessible way for people without professional artistic skills to cre-
489 ate image compositions. While this technology offers significant benefits, it also carries the risk of
490 misuse for malicious purposes, such as harassment or spreading misinformation. Additionally, our
491 framework relies on pretrained large-scale T2I models, which may inadvertently introduce social
492 and cultural biases. Therefore, using these models raises ethical concerns and requires careful con-
493 sideration. We therefore urge users to exercise caution and use this tool responsibly for appropriate
494 purposes only.495
496 REPRODUCIBILITY STATEMENT
497498 We have taken several measures to ensure the reproducibility of our results. Algorithm 1 presents
499 the pseudocode of our method. Details of the implementation for the main experiments are provided
500 in Sec. 4.1, while the hyperparameter configurations are listed in Appendix J. The source code
501 and the ComplexCompo dataset will be released publicly. This work adheres to the reproducibility
502 requirements set by ICLR.503
504 REFERENCES
505506 Grok 4, 2025. URL <https://x.ai/news/grok-4>.
507 Samaneh Azadi, Deepak Pathak, Sayna Ebrahimi, and Trevor Darrell. Compositional gan: Learning
508 image-conditional binary composition. *International Journal of Computer Vision*, 128(10):2570–
509 2585, 2020.
510 Black Forest Labs. Flux.1 [dev]. <https://huggingface.co/black-forest-labs/FLUX.1-dev>, 2024a.
511 Black Forest Labs. Flux.1 fill [dev]. <https://huggingface.co/black-forest-labs/FLUX.1-Fill-dev>, 2024b.
512 Black Forest Labs. Flux.1 [schnell]. <https://huggingface.co/black-forest-labs/FLUX.1-schnell>, 2024c.
513 Tim Brooks, Aleksander Holynski, and Alexei A Efros. Instructpix2pix: Learning to follow image
514 editing instructions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern
515 recognition*, pp. 18392–18402, 2023.
516 Gemma Canet Tarrés, Zhe Lin, Zhifei Zhang, Jianming Zhang, Yizhi Song, Dan Ruta, Andrew
517 Gilbert, John Collomosse, and Soo Ye Kim. Thinking outside the bbox: Unconstrained generative
518 object compositing. In *European Conference on Computer Vision*, pp. 476–495. Springer, 2024.
519 Junyan Cao, Yan Hong, and Li Niu. Painterly image harmonization in dual domains. In *Proceedings
520 of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 268–276, 2023.
521 Bor-Chun Chen and Andrew Kae. Toward realistic image compositing with adversarial learning.
522 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
523 8415–8424, 2019.
524 Jiaxuan Chen, Bo Zhang, Qingdong He, Jinlong Peng, and Li Niu. Mureobjectstitch: Multi-reference
525 image composition. *arXiv preprint arXiv:2411.07462*, 2024a.
526 Junsong Chen, YU Jincheng, GE Chongjian, Lewei Yao, Enze Xie, Zhongdao Wang, James Kwok,
527 Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart- α : Fast training of diffusion transformer for
528 photorealistic text-to-image synthesis. In *The Twelfth International Conference on Learning Rep-
529 resentations*.
530 Liang Chen, Shuai Bai, Wenhao Chai, Weichu Xie, Haozhe Zhao, Leon Vinci, Junyang Lin, and
531 Baobao Chang. Multimodal representation alignment for image generation: Text-image inter-
532 leaved control is easier than you think, 2025.

540 Xi Chen, Yutong Feng, Mengting Chen, Yiyang Wang, Shilong Zhang, Yu Liu, Yujun Shen, and
 541 Hengshuang Zhao. Zero-shot image editing with reference imitation. *Advances in Neural Infor-*
 542 *mation Processing Systems*, 37:84010–84032, 2024b.

543 Xi Chen, Lianghua Huang, Yu Liu, Yujun Shen, Deli Zhao, and Hengshuang Zhao. Anydoor: Zero-
 544 shot object-level image customization. In *Proceedings of the IEEE/CVF conference on computer*
 545 *vision and pattern recognition*, pp. 6593–6602, 2024c.

546 Zhe Chen, Jiannan Wu, Wenhui Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong
 547 Zhang, Xizhou Zhu, Lewei Lu, et al. Internvl: Scaling up vision foundation models and aligning
 548 for generic visual-linguistic tasks. In *Proceedings of the IEEE/CVF conference on computer*
 549 *vision and pattern recognition*, pp. 24185–24198, 2024d.

550 Zhekai Chen, Wen Wang, Zhen Yang, Zeqing Yuan, Hao Chen, and Chunhua Shen. Freecompose:
 551 Generic zero-shot image composition with diffusion prior. In *European Conference on Computer*
 552 *Vision*, pp. 70–87. Springer, 2024e.

553 Wenyan Cong, Jianfu Zhang, Li Niu, Liu Liu, Zhixin Ling, Weiyuan Li, and Liqing Zhang. Dovenet:
 554 Deep image harmonization via domain verification. In *Proceedings of the IEEE/CVF Conference*
 555 *on Computer Vision and Pattern Recognition*, pp. 8394–8403, 2020.

556 Xiaodong Cun and Chi-Man Pun. Improving the harmony of the composite image by spatial-
 557 separated attention module. *IEEE Transactions on Image Processing*, 29:4759–4771, 2020.

558 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
 559 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers
 560 for high-resolution image synthesis. In *Forty-first international conference on machine learning*,
 561 2024.

562 Berthy T Feng, Jamie Smith, Michael Rubinstein, Huiwen Chang, Katherine L Bouman, and
 563 William T Freeman. Score-based diffusion models as principled priors for inverse imaging. In
 564 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10520–10531,
 565 2023.

566 Stephanie Fu, Netanel Tamir, Shobhit Sundaram, Lucy Chai, Richard Zhang, Tali Dekel, and Phillip
 567 Isola. Dreamsim: Learning new dimensions of human visual similarity using synthetic data.
 568 *Advances in Neural Information Processing Systems*, 36:50742–50768, 2023.

569 Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit Haim Bermano, Gal Chechik, and
 570 Daniel Cohen-or. An image is worth one word: Personalizing text-to-image generation using
 571 textual inversion. In *The Eleventh International Conference on Learning Representations*, 2022.

572 Google Gemini2.5. Introducing gemini 2.5 flash image, our state-of-the-art im-
 573 age model, 2025. URL <https://developers.googleblog.com/en/introducing-gemini-2-5-flash-image/>.

574 Alexandros Graikos, Nikolay Malkin, Nebojsa Jojic, and Dimitris Samaras. Diffusion models as
 575 plug-and-play priors. *Advances in Neural Information Processing Systems*, 35:14715–14728,
 576 2022.

577 Zinan Guo, Yanze Wu, Chen Zhuowei, Peng Zhang, Qian He, et al. Pulid: Pure and lightning id
 578 customization via contrastive alignment. *Advances in neural information processing systems*, 37:
 579 36777–36804, 2024.

580 Ligong Han, Yinxiao Li, Han Zhang, Peyman Milanfar, Dimitris Metaxas, and Feng Yang. Svdif:
 581 Compact parameter space for diffusion fine-tuning. In *Proceedings of the IEEE/CVF Interna-*
 582 *tional Conference on Computer Vision*, pp. 7323–7334, 2023.

583 Zhen Han, Zeyinzi Jiang, Yulin Pan, Jingfeng Zhang, Chaojie Mao, Chen-Wei Xie, Yu Liu, and
 584 Jingren Zhou. Ace: All-round creator and editor following instructions via diffusion transformer.
 585 In *The Thirteenth International Conference on Learning Representations*.

586 Jixuan He, Wanhua Li, Ye Liu, Junsik Kim, Donglai Wei, and Hanspeter Pfister. Affordance-aware
 587 object insertion via mask-aware dual diffusion. *arXiv preprint arXiv:2412.14462*, 2024.

594 HiDream-ai. Hidream-e1. <https://github.com/HiDream-ai/HiDream-E1>, 2025.
 595

596 Yan Hong, Li Niu, and Jianfu Zhang. Shadow generation for composite image in real-world scenes.
 597 In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 914–922, 2022.

598 Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen,
 599 et al. Lora: Low-rank adaptation of large language models. In *International Conference on*
 600 *Learning Representations*.

601

602 Junjia Huang, Pengxiang Yan, Jiyang Liu, Jie Wu, Zhao Wang, Yitong Wang, Liang Lin, and
 603 Guanbin Li. Dreamfuse: Adaptive image fusion with diffusion transformer. *arXiv preprint*
 604 *arXiv:2504.08291*, 2025.

605

606 Yuzhou Huang, Liangbin Xie, Xintao Wang, Ziyang Yuan, Xiaodong Cun, Yixiao Ge, Jiantao
 607 Zhou, Chao Dong, Rui Huang, Ruimao Zhang, and Ying Shan. Smartedit: Exploring complex
 608 instruction-based image editing with multimodal large language models. *2024 IEEE/CVF Con-*
 609 *ference on Computer Vision and Pattern Recognition (CVPR)*, pp. 8362–8371, 2024.

610

611 Junha Hyung, Jaeyo Shin, and Jaegul Choo. Magicapture: High-resolution multi-concept portrait
 612 customization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp.
 613 2445–2453, 2024.

614

615 Yifan Jiang, He Zhang, Jianming Zhang, Yilin Wang, Zhe Lin, Kalyan Sunkavalli, Simon Chen,
 616 Sohrab Amirghodsi, Sarah Kong, and Zhangyang Wang. Ssh: A self-supervised framework for
 617 image harmonization. In *Proceedings of the IEEE/CVF International Conference on Computer*
 618 *Vision*, pp. 4832–4841, 2021.

619

620 Jian Jin, Yang Shen, Xinyang Zhao, Zhenyong Fu, and Jian Yang. Unicanvas: Affordance-aware
 621 unified real image editing via customized text-to-image generation. *International Journal of Com-*
 622 *puter Vision*, pp. 1–25, 2025.

623

624 Xuan Ju, Xian Liu, Xintao Wang, Yuxuan Bian, Ying Shan, and Qiang Xu. Brushnet: A plug-and-
 625 play image inpainting model with decomposed dual-branch diffusion. In *European Conference*
 626 *on Computer Vision*, pp. 150–168. Springer, 2024.

627

628 Gwanghyun Kim, Taesung Kwon, and Jong Chul Ye. Diffusionclip: Text-guided diffusion models
 629 for robust image manipulation. In *Proceedings of the IEEE/CVF conference on computer vision*
 630 *and pattern recognition*, pp. 2426–2435, 2022.

631

632 Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept
 633 customization of text-to-image diffusion. In *Proceedings of the IEEE/CVF Conference on Com-*
 634 *puter Vision and Pattern Recognition*, pp. 1931–1941, 2023.

635

636 Alina Kuznetsova, Hassan Rom, Neil Alldrin, Jasper Uijlings, Ivan Krasin, Jordi Pont-Tuset, Shahab
 637 Kamali, Stefan Popov, Matteo Malloci, Alexander Kolesnikov, et al. The open images dataset v4:
 638 Unified image classification, object detection, and visual relationship detection at scale. *Inter-*
 639 *national journal of computer vision*, 128(7):1956–1981, 2020.

640

641 Dongxu Li, Junnan Li, and Steven Hoi. Blip-diffusion: Pre-trained subject representation for control-
 642 lable text-to-image generation and editing. *Advances in Neural Information Processing Systems*,
 643 36, 2024a.

644

645 Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: bootstrapping language-image
 646 pre-training with frozen image encoders and large language models. In *Proceedings of the 40th*
 647 *International Conference on Machine Learning*, ICML’23. JMLR.org, 2023a.

648

649 Pengzhi Li, Qiang Nie, Ying Chen, Xi Jiang, Kai Wu, Yuhuan Lin, Yong Liu, Jinlong Peng, Chengjie
 650 Wang, and Feng Zheng. Tuning-free image customization with image and text guidance. In
 651 *European Conference on Computer Vision*, pp. 233–250. Springer, 2024b.

652

653 Tianle Li, Max Ku, Cong Wei, and Wenhui Chen. Dreamedit: Subject-driven image editing. *Trans-*
 654 *actions on Machine Learning Research*, 2023b.

648 Yaowei Li, Yuxuan Bian, Xuan Ju, Zhaoyang Zhang, Junhao Zhuang, Ying Shan, Yuexian Zou, and
 649 Qiang Xu. Brushedit: All-in-one image inpainting and editing. *arXiv preprint arXiv:2412.10316*,
 650 2024c.

651 Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li,
 652 and Yong Jae Lee. Gligen: Open-set grounded text-to-image generation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 22511–22521, 2023c.

653 Chen-Hsuan Lin, Ersin Yumer, Oliver Wang, Eli Shechtman, and Simon Lucey. St-gan: Spatial
 654 transformer generative adversarial networks for image compositing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 9455–9464, 2018.

655 Daquan Liu, Chengjiang Long, Hongpan Zhang, Hanning Yu, Xinzhi Dong, and Chunxia Xiao. Ar-
 656 shadowgan: Shadow generative adversarial network for augmented reality in single light scenes.
 657 In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
 658 8139–8148, 2020.

659 Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction
 660 tuning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp.
 661 26296–26306, 2024.

662 Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast
 663 ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in Neural
 664 Information Processing Systems*, 35:5775–5787, 2022.

665 Haoming Lu, Hazarapet Tunanyan, Kai Wang, Shant Navasardyan, Zhangyang Wang, and
 666 Humphrey Shi. Specialist diffusion: Plug-and-play sample-efficient fine-tuning of text-to-image
 667 diffusion models to learn any unseen style. In *Proceedings of the IEEE/CVF Conference on
 668 Computer Vision and Pattern Recognition*, pp. 14267–14276, 2023a.

669 Lingxiao Lu, Jiangtong Li, Junyan Cao, Li Niu, and Liqing Zhang. Painterly image harmonization
 670 using diffusion model. In *Proceedings of the 31st ACM International Conference on Multimedia*,
 671 pp. 233–241, 2023b.

672 Lingxiao Lu, Jiangtong Li, Bo Zhang, and Li Niu. Dreamcom: Finetuning text-guided inpainting
 673 model for image composition. *arXiv preprint arXiv:2309.15508*, 2023c.

674 Pengqi Lu. Qwen2vl-flux: Unifying image and text guidance for controllable image generation,
 675 2024. URL <https://github.com/erwold/qwen2vl-flux>.

676 Shilin Lu, Yanzhu Liu, and Adams Wai-Kin Kong. Tf-icon: Diffusion-based training-free cross-
 677 domain image composition. In *Proceedings of the IEEE/CVF International Conference on Computer
 678 Vision*, pp. 2294–2305, 2023d.

679 Shilin Lu, Zilan Wang, Leyang Li, Yanzhu Liu, and Adams Wai-Kin Kong. Mace: Mass concept
 680 erasure in diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision
 681 and Pattern Recognition*, pp. 6430–6440, 2024.

682 Jian Ma, Qirong Peng, Xu Guo, Chen Chen, Haonan Lu, and Zhenyu Yang. X2i: Seamless in-
 683 tegration of multimodal understanding into diffusion transformer via attention distillation. In
 684 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 16733–16744,
 685 2025a.

686 Yuhang Ma, Xiaoshi Wu, Keqiang Sun, and Hongsheng Li. Hpsv3: Towards wide-spectrum human
 687 preference score. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*,
 688 pp. 15086–15095, 2025b.

689 Chaojie Mao, Jingfeng Zhang, Yulin Pan, Zeyinzi Jiang, Zhen Han, Yu Liu, and Jingren Zhou.
 690 Ace++: Instruction-based image creation and editing via context-aware content filling. *arXiv
 691 preprint arXiv:2501.02487*, 2025.

692 Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon.
 693 Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International
 694 Conference on Learning Representations*, 2021.

702 Li Ming, Gu Xin, Chen Fan, Xing Xiaoying, Wen Longyin, Chen Chen, and Zhu Sijie. Superedit: Rectifying and facilitating supervision for instruction-based image editing. *arXiv preprint arXiv:2505.02370*, 2025.

703

704

705 Ron Mokady, Amir Hertz, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Null-text inversion for editing real images using guided diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6038–6047, 2023.

706

707

708

709 Li Niu, Wenyang Cong, Liu Liu, Yan Hong, Bo Zhang, Jing Liang, and Liqing Zhang. Making images real again: A comprehensive survey on deep image composition. *arXiv preprint arXiv:2106.14490*, 2021.

710

711

712 OpenAI. Introducing gpt-5, 2025. URL <https://openai.com/index/introducing-gpt-5/>.

713

714

715 Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafraniec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *Transactions on Machine Learning Research Journal*, 2024.

716

717

718

719 Xichen Pan, Satya Narayan Shukla, Aashu Singh, Zhuokai Zhao, Shlok Kumar Mishra, Jialiang Wang, Zhiyang Xu, Juhai Chen, Kunpeng Li, Felix Juefei-Xu, et al. Transfer between modalities with metaqueries. *arXiv preprint arXiv:2504.06256*, 2025.

720

721

722 William Peebles and Saining Xie. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 4195–4205, 2023.

723

724

725 Kien T Pham, Jingye Chen, and Qifeng Chen. Tale: Training-free cross-domain image composition via adaptive latent manipulation and energy-guided optimization. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 3160–3169, 2024.

726

727

728 Trung X Pham, Zhang Kang, Ji Woo Hong, Xuran Zheng, and Chang D Yoo. E-md3c: Tampering masked diffusion transformers for efficient zero-shot object customization. *arXiv preprint arXiv:2502.09164*, 2025.

729

730

731 Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. SDXL: Improving latent diffusion models for high-resolution image synthesis. In *The Twelfth International Conference on Learning Representations*, 2024.

732

733

734 Ben Poole, Ajay Jain, Jonathan T Barron, and Ben Mildenhall. Dreamfusion: Text-to-3d using 2d diffusion. In *The Eleventh International Conference on Learning Representations*.

735

736

737 Yuandong Pu, Le Zhuo, Kaiwen Zhu, Liangbin Xie, Wenlong Zhang, Xiangyu Chen, Pneg Gao, Yu Qiao, Chao Dong, and Yihao Liu. Lumina-omnilv: A unified multimodal framework for general low-level vision. *arXiv preprint arXiv:2504.04903*, 2025.

738

739

740 Pengchong Qiao, Lei Shang, Chang Liu, Baigui Sun, Xiangyang Ji, and Jie Chen. Facechain-sude: Building derived class to inherit category attributes for one-shot subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7215–7224, 2024.

741

742

743

744 Zeju Qiu, Weiyang Liu, Haiwen Feng, Yuxuan Xue, Yao Feng, Zhen Liu, Dan Zhang, Adrian Weller, and Bernhard Schölkopf. Controlling text-to-image diffusion by orthogonal finetuning. *Advances in Neural Information Processing Systems*, 36:79320–79362, 2023.

745

746

747 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pp. 8748–8763. PMLR, 2021.

748

749

750

751 Nikhila Ravi, Valentin Gabeur, Yuan-Ting Hu, Ronghang Hu, Chaitanya Ryali, Tengyu Ma, Haitham Khedr, Roman Rädle, Chloe Rolland, Laura Gustafson, Eric Mintun, Junting Pan, Kalyan Vasudev Alwala, Nicolas Carion, Chao-Yuan Wu, Ross Girshick, Piotr Dollar, and Christoph Feichtenhofer. SAM 2: Segment anything in images and videos. In *The Thirteenth International Conference on Learning Representations*, 2025.

752

753

754

755

756 Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman.
 757 Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Pro-
 758 ceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 22500–
 759 22510, 2023.

760 Nataniel Ruiz, Yuanzhen Li, Neal Wadhwa, Yael Pritch, Michael Rubinstein, David E Jacobs, and
 761 Shlomi Fruchter. Magic insert: Style-aware drag-and-drop. In *Proceedings of the IEEE/CVF
 762 International Conference on Computer Vision*, pp. 15971–15981, 2025.

764 Patrick Schramowski, Manuel Brack, Björn Deiseroth, and Kristian Kersting. Safe latent diffusion:
 765 Mitigating inappropriate degeneration in diffusion models. In *Proceedings of the IEEE/CVF
 766 Conference on Computer Vision and Pattern Recognition*, pp. 22522–22531, 2023.

768 Shihao Shao and Qinghua Cui. 1st place solution in google universal images embedding. *arXiv
 769 preprint arXiv:2210.08473*, 2022.

770 Yichen Sheng, Jianming Zhang, and Bedrich Benes. Ssn: Soft shadow network for image composit-
 771 ing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*,
 772 pp. 4380–4390, 2021.

774 Jing Shi, Wei Xiong, Zhe Lin, and Hyun Joon Jung. Instantbooth: Personalized text-to-image gen-
 775 eration without test-time finetuning. In *Proceedings of the IEEE/CVF Conference on Computer
 776 Vision and Pattern Recognition*, pp. 8543–8552, 2024a.

778 Yichun Shi, Peng Wang, and Weilin Huang. Seededit: Align image re-generation to image editing.
 779 *arXiv preprint arXiv:2411.06686*, 2024b.

780 Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *Inter-
 781 national Conference on Learning Representations*, 2021.

783 Wensong Song, Hong Jiang, Zongxing Yang, Ruijie Quan, and Yi Yang. Insert anything: Image
 784 insertion via in-context editing in dit. *arXiv preprint arXiv:2504.15009*, 2025a.

786 Yeqi Song, Jimyeong Kim, Wonhark Park, Wonsik Shin, Wonjong Rhee, and Nojun Kwak. Harmo-
 787 nizing visual and textual embeddings for zero-shot text-to-image customization. In *Proceedings
 788 of the AAAI Conference on Artificial Intelligence*, volume 39, pp. 20549–20557, 2025b.

789 Yizhi Song, Zhifei Zhang, Zhe Lin, Scott Cohen, Brian Price, Jianming Zhang, Soo Ye Kim, and
 790 Daniel Aliaga. Objectstitch: Object compositing with diffusion model. In *Proceedings of the
 791 IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 18310–18319, 2023.

793 Yizhi Song, Zhifei Zhang, Zhe Lin, Scott Cohen, Brian Price, Jianming Zhang, Soo Ye Kim,
 794 He Zhang, Wei Xiong, and Daniel Aliaga. Imprint: Generative object compositing by learning
 795 identity-preserving representation. In *Proceedings of the IEEE/CVF Conference on Computer
 796 Vision and Pattern Recognition*, pp. 8048–8058, 2024.

797 Jiale Tao, Yanbing Zhang, Qixun Wang, Yiji Cheng, Haofan Wang, Xu Bai, Zhengguang Zhou,
 798 Ruihuang Li, Linqing Wang, Chunyu Wang, et al. Instantcharacter: Personalize any characters
 799 with a scalable diffusion transformer framework. *arXiv preprint arXiv:2504.12395*, 2025.

801 Gemma Canet Tarrés, Zhe Lin, Zhifei Zhang, He Zhang, Andrew Gilbert, John Collomosse, and
 802 Soo Ye Kim. Multitwine: Multi-object compositing with text and layout control. In *Proceedings
 803 of the Computer Vision and Pattern Recognition Conference*, pp. 8094–8104, 2025.

804 Yoad Tewel, Rinon Gal, Dvir Samuel, Yuval Atzmon, Lior Wolf, and Gal Chechik. Add-it: Training-
 805 free object insertion in images with pretrained diffusion models. In *The Thirteenth International
 806 Conference on Learning Representations*.

808 Shashank Tripathi, Siddhartha Chandra, Amit Agrawal, Ambrish Tyagi, James M Rehg, and Visesh
 809 Chari. Learning to generate synthetic data via compositing. In *Proceedings of the IEEE/CVF
 Conference on Computer Vision and Pattern Recognition*, pp. 461–470, 2019.

810 Haofan Wang, Matteo Spinelli, Qixun Wang, Xu Bai, Zekui Qin, and Anthony Chen. Instantstyle: Free lunch towards style-preserving in text-to-image generation. *arXiv preprint arXiv:2404.02733*, 2024a.

811

812

813

814 Haoxuan Wang, Jinlong Peng, Qingdong He, Hao Yang, Ying Jin, Jiafu Wu, Xiaobin Hu, Yanjie Pan, Zhenye Gan, Mingmin Chi, et al. Unicombine: Unified multi-conditional combination with diffusion transformer. *arXiv preprint arXiv:2503.09277*, 2025a.

815

816

817 Qian Wang, Biao Zhang, Michael Birsak, and Peter Wonka. Instructedit: Improving automatic masks for diffusion-based image editing with user instructions. *arXiv preprint arXiv:2305.18047*, 2023.

818

819

820 Qixun Wang, Xu Bai, Haofan Wang, Zekui Qin, Anthony Chen, Huaxia Li, Xu Tang, and Yao Hu. Instantid: Zero-shot identity-preserving generation in seconds. *arXiv preprint arXiv:2401.07519*, 2024b.

821

822

823 Yibin Wang, Weizhong Zhang, Jianwei Zheng, and Cheng Jin. Primecomposer: Faster progressively combined diffusion for image composition with attention steering. In *Proceedings of the 32nd ACM International Conference on Multimedia*, pp. 10824–10832, 2024c.

824

825

826 Yibin Wang, Zhimin Li, Yuhang Zang, Chunyu Wang, Qinglin Lu, Cheng Jin, and Jiaqi Wang. Unified multimodal chain-of-thought reward model through reinforcement fine-tuning. *arXiv preprint arXiv:2505.03318*, 2025b.

827

828

829 Yibin Wang, Yuhang Zang, Hao Li, Cheng Jin, and Jiaqi Wang. Unified reward model for multi-modal understanding and generation. *arXiv preprint arXiv:2503.05236*, 2025c.

830

831

832 Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

833

834

835

836 Yuxiang Wei, Yabo Zhang, Zhilong Ji, Jinfeng Bai, Lei Zhang, and Wangmeng Zuo. Elite: Encoding visual concepts into textual embeddings for customized text-to-image generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 15943–15953, 2023.

837

838

839

840 Daniel Winter, Asaf Shul, Matan Cohen, Dana Berman, Yael Pritch, Alex Rav-Acha, and Yedid Hoshen. Objectmate: A recurrence prior for object insertion and subject-driven generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 16281–16291, 2025.

841

842

843

844 Huikai Wu, Shuai Zheng, Junge Zhang, and Kaiqi Huang. Gp-gan: Towards realistic high-resolution image blending. In *Proceedings of the 27th ACM international conference on multimedia*, pp. 2487–2495, 2019.

845

846

847

848 Shitao Xiao, Yueze Wang, Junjie Zhou, Huaying Yuan, Xingrun Xing, Ruiran Yan, Chaofan Li, Shuting Wang, Tiejun Huang, and Zheng Liu. Omnigen: Unified image generation. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 13294–13304, 2025.

849

850

851 Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation. *Advances in Neural Information Processing Systems*, 36:15903–15935, 2023.

852

853

854

855 Jiazheng Xu, Yu Huang, Jiale Cheng, Yuanming Yang, Jiajun Xu, Yuan Wang, Wenbo Duan, Shen Yang, Qunlin Jin, Shurun Li, et al. Visionreward: Fine-grained multi-dimensional human preference learning for image and video generation. *arXiv preprint arXiv:2412.21059*, 2024.

856

857

858 Ben Xue, Shenghui Ran, Quan Chen, Rongfei Jia, Binqiang Zhao, and Xing Tang. Dccf: Deep comprehensible color filter learning framework for high-resolution image harmonization. In *European Conference on Computer Vision*, pp. 300–316. Springer, 2022.

859

860

861

862 Le Xue, Manli Shu, Anas Awadalla, Jun Wang, An Yan, Senthil Purushwalkam, Honglu Zhou, Viraj Prabhu, Yutong Dai, Michael S Ryoo, et al. xgen-mm (blip-3): A family of open large multimodal models. *arXiv preprint arXiv:2408.08872*, 2024.

863

864 Zexuan Yan, Yue Ma, Chang Zou, Wenteng Chen, Qifeng Chen, and Linfeng Zhang. Eedit:
 865 Rethinking the spatial and temporal redundancy for efficient image editing. *arXiv preprint*
 866 *arXiv:2503.10270*, 2025.

867 Binxin Yang, Shuyang Gu, Bo Zhang, Ting Zhang, Xuejin Chen, Xiaoyan Sun, Dong Chen, and
 868 Fang Wen. Paint by example: Exemplar-based image editing with diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 18381–18391,
 869 2023.

870 Ling Yang, Zhaochen Yu, Chenlin Meng, Minkai Xu, Stefano Ermon, and Bin Cui. Mastering text-
 871 to-image diffusion: Recaptioning, planning, and generating with multimodal llms. In *Forty-first*
 872 *International Conference on Machine Learning*, 2024.

873 Hu Ye, Jun Zhang, Sibo Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt
 874 adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721*, 2023.

875 Qifan Yu, Wei Chow, Zhongqi Yue, Kaihang Pan, Yang Wu, Xiaoyang Wan, Juncheng Li, Siliang
 876 Tang, Hanwang Zhang, and Yuetong Zhuang. Anyedit: Mastering unified high-quality image
 877 editing for any idea. In *Proceedings of the Computer Vision and Pattern Recognition Conference*,
 878 pp. 26125–26135, 2025a.

879 Yongsheng Yu, Ziyun Zeng, Haitian Zheng, and Jiebo Luo. Omnipaint: Mastering object-oriented
 880 editing via disentangled insertion-removal inpainting. *arXiv preprint arXiv:2503.08677*, 2025b.

881 Bo Zhang, Yuxuan Duan, Jun Lan, Yan Hong, Huijia Zhu, Weiqiang Wang, and Li Niu. Controlcom:
 882 Controllable image composition using diffusion model. *arXiv preprint arXiv:2308.10040*, 2023a.

883 Kai Zhang, Lingbo Mo, Wenhui Chen, Huan Sun, and Yu Su. Magicbrush: A manually annotated
 884 dataset for instruction-guided image editing. *Advances in Neural Information Processing Systems*,
 885 36:31428–31449, 2023b.

886 Lingzhi Zhang, Tarmily Wen, Jie Min, Jiancong Wang, David Han, and Jianbo Shi. Learning ob-
 887 ject placement by inpainting for compositional data augmentation. In *European Conference on*
 888 *Computer Vision*, pp. 566–581. Springer, 2020a.

889 Lingzhi Zhang, Tarmily Wen, and Jianbo Shi. Deep image blending. In *Proceedings of the*
 890 *IEEE/CVF Winter Conference on Applications of Computer Vision*, pp. 231–240, 2020b.

891 Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable
 892 effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on*
 893 *computer vision and pattern recognition*, pp. 586–595, 2018.

894 Shuyang Zhang, Runze Liang, and Miao Wang. Shadowgan: Shadow synthesis for virtual objects
 895 with conditional adversarial networks. *Computational Visual Media*, 5(1):105–115, 2019.

896 Zechuan Zhang, Ji Xie, Yu Lu, Zongxin Yang, and Yi Yang. In-context edit: Enabling instructional
 897 image editing with in-context generation in large scale diffusion transformer. *arXiv preprint*
 898 *arXiv:2504.20690*, 2025.

899 Haozhe Zhao, Xiaojian Shawn Ma, Liang Chen, Shuzheng Si, Rujie Wu, Kaikai An, Peiyu Yu, Min-
 900 jia Zhang, Qing Li, and Baobao Chang. Ultraedit: Instruction-based fine-grained image editing at
 901 scale. *Advances in Neural Information Processing Systems*, 37:3058–3093, 2024.

902 Kaiwen Zheng, Cheng Lu, Jianfei Chen, and Jun Zhu. Improved techniques for maximum likelihood
 903 estimation for diffusion odes. In *International Conference on Machine Learning*, pp. 42363–
 904 42389. PMLR, 2023.

905 Junhao Zhuang, Yanhong Zeng, Wenran Liu, Chun Yuan, and Kai Chen. A task is worth one word:
 906 Learning with task prompts for high-quality versatile image inpainting. In *European Conference*
 907 *on Computer Vision*, pp. 195–211. Springer, 2024.

908

909

910

911

912

913

914

915

916

917

918
919
920
921
922
923

Appendix

924
925

Table of Contents

926
927
928
929
930
931
932
933
934
935
936

A Related Work	19
A.1 Image Composition	19
A.2 General Image Editing	20
A.3 Subject-Driven Generation	20
B Visualizing the Impact of Adaptive Background Blending	21
C Equivalence of Query Blurring and Attention Weight Blurring	22
C.1 Blurring the Query Matrix	22
C.2 Blurring the Key and Value Matrices	22
C.3 Implementation Details of Gaussian Blurring	22
D Evaluating Cross-Attention Map Accuracy via IoU	23
E Experiments with SDXL, SD3.5, and PixArt	24
F User Study	25
G Benchmark Details	25
H Prompts for Proprietary Foundation Models	26
I Subject Identity Metrics Analysis	26
J Implementation Details	28
K Discussion on Inversion vs. One-Step Forward Diffusion	28
L Additional Image Quality Evaluation Using UnifiedReward and HPSv3	29
M Further Analysis on DSG in SD3.5	29
N Runtime Comparison	30
O Additional Qualitative Results	30
P LLM Usage Statement	30

959
960
961
962
963
964
965
966
967
968
969
970
971

972 A RELATED WORK
973974 A.1 IMAGE COMPOSITION
975

976 Image composition involves integrating specific objects and scenarios from user-provided photos,
977 often guided by text prompts. Traditionally, this process is divided into sub-tasks (Niu et al., 2021),
978 including object placement (Azadi et al., 2020; Chen & Kae, 2019; Lin et al., 2018; Tripathi et al.,
979 2019; Zhang et al., 2020a), image blending (Wu et al., 2019; Zhang et al., 2020b), image harmoniza-
980 tion (Cao et al., 2023; Lu et al., 2023b; Zhang et al., 2020b; Cong et al., 2020; Cun & Pun, 2020;
981 Jiang et al., 2021; Xue et al., 2022), and shadow generation (Hong et al., 2022; Liu et al., 2020;
982 Sheng et al., 2021; Zhang et al., 2019), each typically handled by distinct models and pipelines.
983 However, with the rise of diffusion-based generative models, recent approaches have shifted toward
984 unified frameworks that address all sub-tasks simultaneously. These methods are broadly classified
985 into training-based and training-free approaches.
986

987 **Training-based methods** fine-tune foundational models using datasets tailored for image composi-
988 tion. Early methods like Paint by Example (Yang et al., 2023) and ObjectStitch (Song et al., 2023)
989 use CLIP to encode subject features, ensuring high semantic similarity between inserted objects and
990 reference images. These approaches use image augmentation to create training datasets, enabling
991 effective training. GLIGEN (Li et al., 2023c) incorporates grounding information into new trainable
992 layers of a pre-trained diffusion model via a gated mechanism. ControlCom (Zhang et al., 2023a)
993 integrates 2-dim indicator vector to improve controllability. DreamCom (Lu et al., 2023c) and Mure-
994 ObjectStitch (Chen et al., 2024a) fine-tune models with small sets of reference images to preserve
995 subject identity. AnyDoor (Chen et al., 2024c), IMPRINT (Song et al., 2024), and E-MD3C (Pham
996 et al., 2025) leverage DINOv2 to enhance identity fidelity and control over shape and pose, drawing
997 supervision from video data. MimicBrush (Chen et al., 2024b) similarly uses video-derived super-
998 vision for imitative editing. In contrast, MADD (He et al., 2024), ObjectMate (Winter et al., 2025),
999 and OmniPaint (Yu et al., 2025b) employ image inpainting models to generate higher-quality train-
1000 ing datasets compared to those based on image or video augmentation. Multitwine (Tarrés et al.,
1001 2025) enables the integration of multiple objects, capturing interactions from simple positional
1002 relationships to complex actions requiring reposing. DreamFuse (Huang et al., 2025) uses a Po-
1003 sitional Affine mechanism to embed foreground size and position into the background, fostering
1004 effective foreground-background interaction through shared attention. Insert Anything (Song et al.,
1005 2025a) and UniCombine (Wang et al., 2025a) introduces a FLUX-based, multi-conditional gener-
1006 ative framework that handles diverse condition combinations. However, these methods often bias
1007 the generative priors of base models toward curated datasets, resulting in unnatural compositions,
1008 such as implausible object-environment interactions (e.g., missing or unrealistic shadows and reflec-
1009 tions). This stems from the absence of a large-scale, high-quality, multi-resolution, real-world triplet
1010 dataset comprising an object, a scene, and the object seamlessly integrated into the scene, which is
1011 expensive to produce.
1012

1013 **Training-free methods**, on the other hand, modify the inference process of pre-trained models to
1014 achieve composition without additional training. Early approaches like TF-ICON (Lu et al., 2023d)
1015 leverage accurate image inversion to lay the groundwork for composition, achieved through com-
1016 posite self-attention map injection. TALE (Pham et al., 2024) and PrimeComposer (Wang et al.,
1017 2024c) build on TF-ICON to enhance identity preservation and background-object style adaptation.
1018 TIGIC (Li et al., 2024b) focuses on preserving non-target areas during composition. Thinking Out-
1019 side the BBox (Canet Tarrés et al., 2024) enables unconstrained image compositing, unbound by
1020 input masks. FreeCompose (Chen et al., 2024e) employs a pipeline of object removal, image har-
1021 monization, and semantic composition. DreamEdit (Li et al., 2023b), UniCanvas (Jin et al., 2025),
1022 and Magic Insert (Ruiz et al., 2025) use test-time tuning to fine-tune models during inference. Add-
1023 it (Tewel et al.) enables text-guided object insertion on FLUX, where users describe objects via text
1024 prompts instead of reference images. EEdit (Yan et al., 2025) recently improves TF-ICON on FLUX,
1025 introducing step-skipping to reduce time costs and spatial locality caching to minimize redundancy.
However, training-free methods rely on precise image inversion and fragile attention surgery, which
can lock inserted objects into the exact pose of the reference image, leading to awkward or contextu-
ally inappropriate orientations. Attention manipulation often causes instability and hyperparameter
sensitivity, as feature or attention map injection does not always preserve subject identity. This cre-
ates a trade-off: stronger injection preserves identity but results in unnatural poses, while lighter
injection yields more natural poses but compromises identity.
1026

1026 A.2 GENERAL IMAGE EDITING
1027

1028 Instruction-based image editing has evolved rapidly. Early systems relied on modular, two-stage
1029 pipelines: a multimodal language model first produced textual prompts, spatial guidance, or syn-
1030 synthetic instruction–image pairs, and a separate diffusion model then executed the edit—as in In-
1031 structEdit (Wang et al., 2023), InstructPix2Pix (Brooks et al., 2023), MagicBrush (Zhang et al.,
1032 2023b), and BrushEdit (Li et al., 2024c). Recent work has shifted toward tightly integrated,
1033 instruction-centric architectures. Models such as SmartEdit (Huang et al., 2024), X2I (Ma et al.,
1034 2025a), RPG (Yang et al., 2024), AnyEdit (Yu et al., 2025a), and UltraEdit (Zhao et al., 2024)
1035 embed routing, task-aware objectives, and fine-grained controls directly into the network, yielding
1036 higher fidelity and more precise manipulation.

1037 Unified generation-and-editing frameworks (e.g., OmniGen (Xiao et al., 2025), ACE (Han et al.),
1038 ACE++ (Mao et al., 2025), Lumina-OmniLV (Pu et al., 2025), Qwen2VL-Flux (Lu, 2024), Dre-
1039 aymEngine (Chen et al., 2025), MetaQueries (Pan et al., 2025), Hidream-E1 (HiDream-ai, 2025)) treat
1040 editing as one capability of an end-to-end vision-language model, often fusing language embed-
1041 dings with diffusion latents to provide context-aware, pixel-level control. Efficiency has advanced
1042 in parallel: ICEdit (Zhang et al., 2025) couples LoRA with mixture-of-experts tuning and optimized
1043 noise initialization, while SuperEdit (Ming et al., 2025) relies on higher-quality data and contrastive
1044 supervision to sustain performance at lower cost. Looking ahead, large foundation models such as
1045 Gemini (Gemini2.5, 2025) and GPT-5 (OpenAI, 2025) already show strong visual reasoning and
1046 coherent, instruction-guided image generation. Yet, despite extensive multimodal pre-training, they
1047 still fall short on image composition: object placement remains hard to control, lighting is often
1048 inconsistent, and subjects can drift in identity.

1049 A.3 SUBJECT-DRIVEN GENERATION
1050

1051 Extensive research has explored subject-driven image generation, in which the output must not only
1052 portray the contexts described by the text prompt but also faithfully include the specific subject
1053 supplied by reference images. Methods in this area are divided into two categories—test-time fine-
1054 tuning customization and zero-shot customization—according to whether extra training is needed
1055 for each new subject. **Our framework accommodates both categories, so we provide two corre-**
1056 **spending variants in the main paper.**

1057 **Test-time fine-tuning methods** (Gal et al., 2022; Ruiz et al., 2023) adapt a pre-trained T2I model
1058 to a small set of reference images (typically 3 to 5 images). Although this step adds computational
1059 cost and latency, it offers the greatest flexibility for diverse customization requirements. Such meth-
1060 ods are commonly grouped into three subclasses: data regularization, weight regularization, and
1061 loss regularization. In the data-regularization family, DreamBooth (Ruiz et al., 2023) limits overfit-
1062 ting by generating superclass images with the base T2I model and training on both reference and
1063 regularization images; Custom Diffusion (Kumari et al., 2023) improves regularization quality by
1064 retrieving real images; and Specialist Diffusion (Lu et al., 2023a) applies extensive data augmen-
1065 tation. Weight-regularization approaches (Gal et al., 2022; Hu et al.; Han et al., 2023; Qiu et al.,
1066 2023) confine updates to carefully chosen parameters, such as the subject-specific text embeddings
1067 or the singular values of weight matrices. Loss-regularization approaches, including Specialist Dif-
1068 fusion (Lu et al., 2023a), MagiCapture (Hyung et al., 2024), and FaceChain-SuDe (Qiao et al., 2024),
1069 introduce objective terms that respectively maximize CLIP-space similarity between generated and
1070 reference images, disentangle identity and style via masked facial reconstruction, or encourage cor-
1071 rect superclass classification.

1072 **Zero-shot image customization methods** avoid subject-specific fine-tuning at inference time but
1073 rely on extensive pre-training. For general subject customization, InstantBooth (Shi et al., 2024a)
1074 adds a visual encoder that captures coarse-to-fine image features from the references; BLIP-
1075 Diffusion (Li et al., 2024a) fine-tunes BLIP-2 (Li et al., 2023a) to extract multimodal subject repre-
1076 sentations; ELITE (Wei et al., 2023) maps reference images into hierarchical textual tokens through
1077 global and local networks; and Song et al. (Song et al., 2025b) enhance textual control by remov-
1078 ing the projection of visual embeddings onto textual embeddings. For facial customization, Instant-
1079 IID (Wang et al., 2024b) isolates facial regions from reference images to extract appearance and
structural cues. For style customization, InstantStyle (Wang et al., 2024a) identifies style-controlling
layers and injects IP-Adapter features (Ye et al., 2023) into those layers to achieve style transfer. In-

stantCharacter (Tao et al., 2025), IP-Adapter (Ye et al., 2023), and PuLID (Guo et al., 2024) have each released versions compatible with the FLUX model.

B VISUALIZING THE IMPACT OF ADAPTIVE BACKGROUND BLENDING

Although our loss function aims to find a new latent within the manifold learned by the adapter, encouraging the adapter-augmented T2I model’s predictions to closely match those of the base model on the original noisy latent, this early-stage optimization primarily preserves structural elements rather than fine details. Consequently, discrepancies in fine-grained features often arise between the masked and unmasked regions. As illustrated in Fig. 9, we compare the composite images generated using our Adaptive Background Blending (ABB) method with those produced via direct background blending using the rectangular user mask. For clarity, we enlarge the boundary areas of each image (highlighted in pink dashed boxes) to better reveal differences.




Figure 9: Comparison of composites from our Adaptive Background Blending (ABB) method and direct blending with a rectangular mask. Boundary regions within pink dashed boxes are enlarged for clarity. Please zoom in to see details.

1134 **C EQUIVALENCE OF QUERY BLURRING AND ATTENTION WEIGHT
1135 BLURRING**
1136

1137 Consider a 2D Gaussian filter G applied as a convolution operation, denoted by \circledast . The self-attention
1138 weights are computed as QK^\top , where $Q, K \in \mathbb{R}^{n \times d}$, with n being the sequence length and d the
1139 embedding dimension. We explore the effect of applying a Gaussian blur to the attention weights
1140 and its equivalence to blurring the query matrix.
1141

1142 **C.1 BLURRING THE QUERY MATRIX**
1143

1144 Blurring the self-attention weights QK^\top with a 2D Gaussian filter G can be expressed as:
1145

$$G \circledast (QK^\top), \quad (6)$$

1146 where \circledast denotes 2D convolution. Due to the linearity of convolution, there exists a Toeplitz matrix
1147 $B \in \mathbb{R}^{n \times n}$ such that the convolution operation can be represented as a matrix multiplication:
1148

$$G \circledast (QK^\top) = B(QK^\top). \quad (7)$$

1149 Using the properties of matrix multiplication, we can rewrite this as:
1150

$$B(QK^\top) = (BQ)K^\top. \quad (8)$$

1151 Since the convolution operation is linear, applying the Gaussian filter G to the rows of Q yields:
1152

$$BQ = G \circledast Q. \quad (9)$$

1153 Thus, we obtain:
1154

$$G \circledast (QK^\top) = (G \circledast Q)K^\top. \quad (10)$$

1155 This establishes that blurring the query matrix Q with G is mathematically equivalent to applying
1156 the same blur to the self-attention weights QK^\top . This equivalence suggests that query blurring can
1157 be used as a computationally efficient proxy for smoothing attention weights, potentially reducing
1158 the need for direct manipulation of the attention matrix.
1159

1160 **C.2 BLURRING THE KEY AND VALUE MATRICES**
1161

1162 In contrast, applying the Gaussian blur to the key matrix K does not yield a similar equivalence.
1163 Consider the convolution applied to K . The resulting attention weights become:
1164

$$Q(G \circledast K)^\top = Q(BK)^\top = QK^\top B^\top. \quad (11)$$

1165 Since $B^\top \neq B$ for a general Toeplitz matrix derived from a Gaussian filter, we have:
1166

$$QK^\top B^\top \neq B(QK^\top). \quad (12)$$

1167 Thus, blurring the key matrix K does not produce an equivalent effect to blurring the attention
1168 weights QK^\top . A similar argument applies to the value matrix V , as the output of self-attention,
1169 $(QK^\top)V$, involves V in a post-multiplication step, and convolution on V does not commute with
1170 the attention weight computation in the same manner.
1171

1172 **C.3 IMPLEMENTATION DETAILS OF GAUSSIAN BLURRING**
1173

1174 For the 2D Gaussian filtering step, we adopt the implementation provided by the `kornia` library.
1175 Following standard engineering practice, we set the kernel radius to $r = 3\sigma$, since three standard
1176 deviations capture approximately 99.7% of the Gaussian mass. Consequently, the kernel size is
1177 chosen as the nearest odd integer to $2r = 6\sigma$. In all of our experiments we use $\sigma = 10$ (see Tab. 5).
1178

1179 The procedure is applied to query embeddings by reshaping them into spatial feature maps, performing
1180 Gaussian smoothing, and then mapping them back into the sequence domain prior to attention
1181 computation. The full workflow is given below:
1182

1188

Algorithm 2 Implementation Details of 2D Gaussian Blurring

1189

1190

1191

1192

1193

1194

1195

1196

1197

1198

1199

1200

1201

1202

1203

1204

1205

1206

1207

1208

1209

1210

1211

1212

1213

1214

1215

1216

1217

1218

1219

1220

1221

1222

1223

1224

1225

1226

1227

1228

1229

1230

1231

1232

1233

1234

1235

1236

1237

1238

1239

1240

1241

Input: A query matrix $\mathbf{Q} \in \mathbb{R}^{B \times L \times D}$, key \mathbf{K} , value \mathbf{V} , spatial dimensions (H, W) , Gaussian standard deviation σ .

Output: Attention output \mathbf{O} .

```

1: // Reshape query into spatial tensor
2:  $\mathbf{Q}_{\text{sp}} \leftarrow \text{Reshape}(\mathbf{Q}, (B, H, W, D))$ 
3:  $\mathbf{Q}_{\text{sp}} \leftarrow \text{Permute}(\mathbf{Q}_{\text{sp}}, (0, 3, 1, 2))$   $\triangleright B \times D \times H \times W$ 
4: // Construct Gaussian kernel size
5:  $k \leftarrow 6\sigma$ 
6:  $k \leftarrow k - (k \bmod 2) + 1$   $\triangleright$  Ensure odd kernel size
7:  $\text{kernel\_size} \leftarrow (k, k)$ 
8:  $\sigma \leftarrow (\sigma, \sigma)$ 
9: // Apply 2D Gaussian smoothing
10:  $\mathbf{Q}_{\text{sp}} \leftarrow \text{kornia.filters.gaussian\_blur2d}(\mathbf{Q}_{\text{sp}}, \text{kernel\_size}, \sigma)$ 
11: // Reshape smoothed queries back to sequence form
12:  $\mathbf{Q}' \leftarrow \text{Permute}(\mathbf{Q}_{\text{sp}}, (0, 2, 3, 1))$ 
13:  $\mathbf{Q}' \leftarrow \text{Reshape}(\mathbf{Q}', (B, L, D))$ 
14: // Compute attention
15:  $\mathbf{A} \leftarrow \text{softmax} \left( \frac{\mathbf{Q}' \mathbf{K}^T}{\sqrt{D}} \right)$ 
16:  $\mathbf{O} \leftarrow \mathbf{A} \mathbf{V}$ 
17: return  $\mathbf{O}$ 

```




Figure 10: The IoU is calculated between the mask produced from each block and the ground-truth mask, which is obtained by segmenting the final generated images using SAM. The IoU for each block is averaged over 100 images.

D EVALUATING CROSS-ATTENTION MAP ACCURACY VIA IOU

To identify the most accurate cross-attention maps that reflect the location of the generated object, we first create 100 prompts containing a main subject (e.g., “a dog is sleeping on a couch”) using GPT-5. These prompts are then used to generate 100 images with FLUX.1-dev, employing 20 denoising steps. Cross-attention maps are extracted from 19 multi-stream (or double-stream) blocks and 38 single-stream blocks across all denoising steps. The maps are averaged over the 20 steps and subsequently normalized and binarized, resulting in a total of 57 binary masks.

To determine which of these 57 masks is the most accurate, we compute the Intersection over Union (IoU) between each mask and a ground-truth mask obtained by segmenting the final generated images using SAM (Ravi et al., 2025). The IoU for each block is averaged over the 100 generated images. The results are presented in Fig. 10, showing that the cross-attention maps from the last multi-stream (or double-stream) block achieve the highest segmentation accuracy. A visualization of the cross-attention maps from all blocks is provided in Fig. 11.




Figure 11: Visualization of cross-attention maps from different MMDiT blocks of FLUX.1-dev.

E EXPERIMENTS WITH SDXL, SD3.5, AND PIXART

Our framework introduces a novel, model-agnostic approach for enhancing generative architectures, anchored by three synergistic components: Manifold-Steered Anchor (MSA) loss, Degradation-Suppression Guidance (DSG), and Adaptive Background Blending (ABB). These components are meticulously designed to leverage ubiquitous features of modern generative models, ensuring seamless integration without requiring architectural modifications. Specifically, MSA utilizes either LoRA-based personalization or a pretrained personalization adapter, DSG capitalizes on widely available self-attention maps, and ABB harnesses text-image cross-attention maps, a staple in most text-to-image pipelines. This design ensures broad applicability across diverse generative models.

In the main text, we showcase the effectiveness of our approach on Flux, a leading open-source model. To establish its generalizability, we conduct experiments on SDXL (Podell et al., 2024), SD3.5 (Esser et al., 2024), and PixArt (Chen et al.) across DreamEditBench and ComplexCompo. The results are presented in Tab. 3. On DreamEditBench, our LoRA-based variant with Flux achieves state-of-the-art performance in subject identity preservation, while delivering superior image quality. Similarly, on ComplexCompo, the Flux-LoRA configuration excels in identity consistency and image fidelity. Notably, the framework’s benefits extend beyond a single model family: both SDXL and

1296 PixArt- Σ exhibit substantial performance gains, affirming the approach’s generality and adaptability
 1297 across diverse generative architectures.
 1298

1299 Table 3: Comparison of compositional performance across two benchmarks with **different base**
 1300 **models**. The best result in each column is highlighted in **bold**, while the second-best is underlined.
 1301 Metrics shown in **pink** are those specifically trained to better align with human preferences. Abbreviations:
 1302 IRF: Instance Retrieval Features; IR = ImageReward; VR = VisionReward.

Bench	Method	Base Model	Subject Identity Consistency				Background		Image Quality	
			CLIP-I \uparrow	DINOv2 \uparrow	IRF \uparrow	DreamSim \downarrow	LPIPS \downarrow	SSIM \uparrow	IR \uparrow	VR \uparrow
DreamEdit-Bench (220)	Flux.1 Fill	FLUX	0.7328	0.6745	0.5754	0.5233	0.0166	0.9076	0.5577	3.5997
	Ours-Adapter	SDXL	0.7944	0.7334	0.7659	0.3761	0.0238	0.8922	0.5621	3.6158
	Ours-Adapter	SD3.5	0.8054	0.7407	0.7699	0.3745	0.0234	0.8937	0.5701	3.6187
	Ours-LoRA	PixArt-Σ	0.8098	0.7445	0.7798	0.3612	0.0251	0.8875	0.5842	3.6198
	Ours-Adapter	FLUX	0.8086	0.7415	0.7702	0.3730	0.0236	0.8959	0.5709	3.6234
	Ours-LoRA	FLUX	0.8125	0.7452	0.7900	0.3577	0.0271	0.8847	0.5906	3.6161
Complex-Compo (300)	Flux.1 Fill	FLUX	0.7108	0.6475	0.5466	0.6018	0.0232	0.7442	0.4088	3.5737
	Ours-Adapter	SDXL	0.7657	0.7084	0.6862	0.4457	0.0457	0.7612	0.3894	3.5987
	Ours-Adapter	SD3.5	0.7701	0.7091	0.6977	0.4173	0.0401	0.7784	0.4091	3.6021
	Ours-LoRA	PixArt-Σ	0.7924	0.7287	0.7311	0.3603	0.0424	0.7698	0.4277	3.5988
	Ours-Adapter	FLUX	0.7721	0.7107	0.6764	0.4294	0.0404	0.7789	0.4090	3.6020
	Ours-LoRA	FLUX	0.7999	0.7384	0.7659	0.3542	0.0430	0.7634	0.4246	3.5951

F USER STUDY

1318 We conduct a user study involving 50 participants. Each participant was asked to complete 50 rank-
 1319 ing tasks. In each task, they were shown 13 composition results generated by different methods,
 1320 along with a reference subject image.

1321 To ensure a balanced evaluation, 25 of the tasks were randomly sampled from DreamEditBench
 1322 and the remaining 25 from ComplexCompo. Participants were asked to rank the results based on
 1323 two key criteria: (1) subject identity consistency and (2) composition realism. A lower rank (e.g.,
 1324 1st) indicates a better composition result, while a higher rank (e.g., 13th) reflects a less favorable
 1325 outcome.

1326 We summarize the average ranking scores for each method in Tab. 4. Our method received the most
 1327 favorable rankings from the majority of participants, demonstrating its effectiveness in producing
 1328 high-quality compositions.

G BENCHMARK DETAILS

1332 Our benchmark consists of 300 triplets, each comprising a subject image, a background image,
 1333 and a bounding box. The subject images are identical to those used in DreamEditBench (Li
 1334 et al., 2023b; Ruiz et al., 2023), while the background images are sampled from the OpenImages
 1335 dataset (Kuznetsova et al., 2020). These backgrounds exhibit a variety of aspect ratios and res-
 1336 olutions, including landscape and portrait formats, such as 768×1088 , 768×1072 , 768×1024 ,
 1337 768×1152 , 1024×768 , 1152×768 , 1200×768 , 848×768 , and 1360×768 . The bounding boxes
 1338 are manually designed to ensure that the size and placement of the inserted subjects are contextually
 1339 appropriate and visually plausible. The benchmark will be released publicly.

Table 4: Average ranking scores from the user study on image composition methods. Lower is better.

Method	Training-Free	Average Ranking (Lower is Better)
MADD	✗	12.44
ObjectStitch	✗	11.80
DreamCom	✗	6.66
AnyDoor	✗	4.12
UniCombine	✗	2.94
PBE	✗	4.94
TIGIC	✓	9.74
TALE	✓	9.06
TF-ICON	✓	8.36
DreamEdit	✓	6.36
EEdit	✓	10.76
Ours-Adapter	✓	2.30
Ours-LoRA	✓	1.52

H PROMPTS FOR PROPRIETARY FOUNDATION MODELS

To perform image composition with proprietary, general-purpose multimodal foundation models (e.g., GPT-5 (OpenAI, 2025), Gemini 2.5 Pro (Gemini2.5, 2025), SeedEdit/Doubao (Shi et al., 2024b), and Grok 4 (gro, 2025)), we upload three images: (1) Subject image; (2) Background image; and (3) Mask image defining the insertion region.

We then issue a prompt of the following form (with the resolution and coordinates adjusted for each case):

Please insert the object from the first uploaded image into the second image. The target region for insertion is defined by the mask in the third image. For reference, the resolution of the second image is 1152 × 768, and the bounding box for placement is specified by the top-left and bottom-right coordinates: ($x_1 = 550, y_1 = 600, x_2 = 700, y_2 = 750$). The inserted object should retain the same identity and appearance as in the first image. The final composite should appear realistic, natural, and physically plausible.

I SUBJECT IDENTITY METRICS ANALYSIS

In our experiments we found that widely-used subject-identity metrics such as CLIP-I (Radford et al., 2021) and DINOv2 (Oquab et al., 2024) correlate poorly with human preferences. Because they focus almost exclusively on semantic similarity, they ignore appearance changes introduced by lighting, shadows, reflections, and surrounding context. Fig. 12(b) presents several image pairs produced by AnyDoor (left) and by our method (right); the corresponding CLIP-I (↑) (Radford et al., 2021), DINOv2 (↑) (Oquab et al., 2024), IRF (↑) (Shao & Cui, 2022), and DreamSim (↓) (Fu et al., 2023) scores are shown beneath each image, with the better score highlighted in red. Although the AnyDoor results are visibly less realistic and less consistent, they nevertheless receive higher CLIP-I and DINOv2 scores, and in most cases higher IRF scores, demonstrating that these measures do not faithfully capture compositional quality.

A reliable metric should recognise the same object whether it is underwater (see Fig. 12(b)[(8), (17), (20)]), in shadow (see Fig. 12(b)[(9), (13), (19)]), partially occluded (see Fig. 12(b)[(2)]), or situated in a low-light or back-lit scene (see Fig. 12(b)[(5), (14), (15), (16)]). Among the metrics we evaluated, only DreamSim, which was designed to align more closely with human perception, consistently exhibits this desired behaviour.




Figure 12: Comparison of Subject Identity Metrics. (a) Reference subject images used for metric calculations. (b) Image pairs generated by AnyDoor (left) and our method (right), with corresponding CLIP-I (↑), DINOv2 (↑), IRF (↑), and DreamSim (↓) scores displayed below each image; the better score is highlighted in red. Despite AnyDoor’s results appearing less realistic and consistent, they often achieve higher CLIP-I, DINOv2, and IRF scores, indicating that these metrics may not reliably reflect compositional quality. In contrast, DreamSim provides a more reliable assessment.

1458
1459 Table 5: Hyperparameters of Our Frameworks. Bin Thresh = Binarization Threshold; #iter = Number
1460 of Iterations.
1461

Variant	Denoising Setup		Manifold-Steered Anchor Loss			Degradation-Suppression Guidance			Adaptive Background Blending		
	Total steps $(T-1) \rightarrow 0$	Start step t_1	Step Range $t_1 \rightarrow \tau$	Learning Rate α	#iters k	Step Range $t_1 \rightarrow 0$	Scale η	Blur σ	Step Range $(t_1-1) \rightarrow 1$	Bin Thresh	Dilation Kernel Size
Ours-Adapter	19 \rightarrow 0	14	14 \rightarrow 12	500, 750, 1000	10	14 \rightarrow 0	0.5	10	13 \rightarrow 1	0.2	3
Ours-LoRA	19 \rightarrow 0	13	13 \rightarrow 12	50, 300	2	13 \rightarrow 0	0.7	10	12 \rightarrow 1	0.4	3

1464
1465
1466 J IMPLEMENTATION DETAILS
14671468
1469 The hyperparameters used in our frameworks are summarized in Tab. 5. Under Denoising Setup,
1470 “Total steps” refers to the full diffusion/noising schedule, which specifies a sequence of 20 values
1471 of σ_t across timesteps t . However, our method does not begin denoising at the first timestep. As
1472 shown in Algorithm 1, it starts at $t_1 = 14$ (Ours-Adapter), resulting in 15 denoising steps in total.
1473 Under MSA loss, “Step Range” indicates the subset of denoising steps to which MSA optimization
1474 is applied. For the adapter setting, the MSA loss is applied only to the first three denoising steps
1475 (from $t = 14$ to $t = 12$). In addition, since a LoRA is trained for a specific subject, it provides a
1476 more tailored prior than the generic adapter, allowing it to converge in fewer steps and with lower
1477 overall compute.1478 Each baseline is implemented according to the configuration settings recommended in its original
1479 publication. The repositories utilized for each baseline are listed below:1480 1. MADD: <https://github.com/KaKituken/affordance-aware-any>
1481 2. ObjectStitch: <https://github.com/bcmi/ObjectStitch-Image-Composition>
1482 3. DreamCom: <https://github.com/bcmi/DreamCom-Image-Composition>
1483 4. AnyDoor: <https://github.com/ali-vilab/AnyDoor>
1484 5. UniCombine: <https://github.com/Xuan-World/UniCombine>
1485 6. PBE: <https://github.com/Fantasy-Studio/Paint-by-Example>
1486 7. TIGIC: <https://github.com/zrealli/TIGIC>
1487 8. TALE: <https://github.com/tkpham3105/TALE>
1488 9. TF-ICON: <https://github.com/Shilin-LU/TF-ICON>
1489 10. DreamEdit: <https://github.com/DreamEditBenchTeam/DreamEdit>
1490 11. EEdit: <https://github.com/yuriYanZeXuan/EEdit>
14911492 K DISCUSSION ON INVERSION VS. ONE-STEP FORWARD DIFFUSION
14931494
1495 We provide an expanded discussion of our design choice between inversion and one-step forward
1496 diffusion to better clarify the motivation behind our approach. In training-free image editing, both
1497 inversion and one-step forward diffusion are commonly used to obtain a noisy latent that serves as
1498 the starting point for subsequent optimization or denoising. In our framework, we intentionally adopt
1499 one-step forward diffusion as a practical substitute for inversion. Our method does not depend on the
1500 initial latent to preserve object identity. Instead, the MSA loss extracts object-specific information
1501 from the adapter/LoRA and injects it into the latent during the editing process.
15021503
1504 This design choice is motivated by a practical observation: many modern models are distilled for
1505 speed, making accurate inversion difficult to achieve in practice. When inversion cannot reliably
1506 encode object identity, its benefit becomes limited. In such cases, a noisy latent that still retains
1507 enough background structure is sufficient as a starting point. For this reason, we adopt one-step
1508 forward diffusion as a pragmatic replacement for inversion, rather than a theoretically equivalent
1509 alternative. It offers two advantages: (i) it avoids the accuracy limitations of inversion on distilled
1510 models, making it more broadly applicable, and (ii) it is faster than performing an inversion.
1511

1512 Table 6: Image Quality Evaluation Results using HPSv3 and UnifiedReward variants.
1513

Bench	Method	HPSv3	UnifiedReward-2.0-qwen3vl-8b				UnifiedReward-Edit-qwen3vl-8b				UnifiedReward-Think-qwen-7b							
			Instruction		Quality		Average	Instruction		Quality		Average	Instruction		Quality			
			Success	Overedit	Natural	Artifact		Success	Overedit	Natural	Artifact		Success	Overedit	Natural	Artifact		
DreamEditBench (220)	MADD (He et al., 2024)	1.2443	14.2273	12.3727	16.6153	18.6807	15.4740	11.8091	17.8727	13.1045	12.4727	13.8148	14.7553	16.2128	19.1103	20.5724	17.6627	
	ObjectStitch (Song et al., 2023)	7.4529	19.9954	16.8721	21.0455	21.8864	19.9499	16.2500	21.0000	18.6545	20.6500	19.1886	20.7500	17.7614	21.0563	20.7676	20.0838	
	DreamCom (Lu et al., 2023c)	5.9324	18.1818	15.4818	23.5388	24.0502	20.1132	14.0727	22.5409	21.0091	22.0955	19.9296	18.5263	15.4737	21.6641	21.6953	19.3399	
	AnyDoor (Chen et al., 2024c)	8.4867	23.2146	17.8128	20.9772	21.9361	20.9852	19.0909	20.6000	17.4182	19.2864	19.0986	21.2083	18.5694	21.9214	21.5929	20.8230	
	UniCombine (Wang et al., 2025a)	8.8415	22.4545	17.8000	23.1682	23.9000	21.8307	20.1591	22.7773	21.2500	22.6455	21.7088	23.0593	16.2381	21.9520	20.8400	20.5224	
	PBE (Yang et al., 2023)	8.3789	22.4292	17.5205	22.2773	23.2273	21.3636	18.2318	21.8182	19.5955	21.2091	20.2137	24.1481	17.8025	21.5786	21.9500	21.3698	
	TIGIC (Li et al., 2024b)	5.2676	17.3132	14.8045	19.0318	20.3639	17.8784	13.7045	19.4636	17.0864	18.1455	17.1000	17.7294	17.0000	21.0897	20.8820	19.1755	
	TALE (Pham et al., 2024)	6.3771	19.6027	15.9863	20.4455	21.4773	19.3780	14.8318	21.6182	17.3227	18.5486	18.8064	21.6670	17.6548	21.0667	20.6934	20.2555	
	TF-ICON (Lu et al., 2023d)	7.3243	20.4497	16.5838	21.7277	22.5409	19.2773	20.5227	19.0659	17.9273	18.1722	19.0790	20.8716	16.9870	21.0000	20.7054	19.0608	
	EEdit (Li et al., 2023b)	6.0250	17.2637	15.4745	19.2773	20.5227	19.0659	17.9273	18.1722	19.0790	17.9273	18.2716	20.8716	16.9870	20.3333	21.7153	19.5836	
Complex-Compo (300)	EEdit (Yan et al., 2025)	6.6689	18.2790	15.3279	22.0636	23.4818	19.8156	14.2091	22.2884	19.9500	21.7955	19.5603	21.1463	18.2524	22.0635	21.2223	20.6720	
	Ours-Adapter	8.8861	23.3272	17.0500	23.5727	23.8136	21.9522	21.3364	22.7591	21.3636	22.6136	21.0182	23.2222	18.7407	22.8448	22.0086	21.7041	
	Ours-LoRA	8.8688	23.3454	16.8136	23.8868	23.7727	23.8500	21.9727	21.1273	22.7455	21.3000	22.5955	21.9421	23.7590	18.5904	22.4853	20.6765	21.3778
	MADD (He et al., 2024)	5.9673	13.8900	11.8167	18.8633	17.1333	15.4258	12.6800	16.7300	12.5167	10.3000	13.0567	17.6538	15.7115	20.2867	19.5533	17.2524	
	ObjectStitch (Song et al., 2023)	8.8389	20.7157	13.1773	21.4200	19.4800	18.6983	17.3567	21.8500	17.1333	18.9733	18.8283	19.8304	14.6518	21.3467	19.5400	18.8394	
Complex-Compo (300)	DreamCom (Lu et al., 2023c)	7.9884	8.2234	9.0756	23.9178	23.4737	16.1726	5.9507	23.7072	23.6161	23.4475	18.6143	17.8810	15.2857	22.5000	20.9737	19.0509	
	AnyDoor (Chen et al., 2024c)	8.9760	21.7633	13.1133	20.5567	18.5306	18.4908	18.4967	21.8300	14.9267	18.1667	18.3552	22.6600	17.6422	21.2467	19.5333	19.8876	
	UniCombine (Wang et al., 2025a)	8.8999	15.6747	19.6151	13.1283	21.9243	20.1340	18.7007	16.8947	21.9605	17.3257	19.6217	18.9500	23.5810	14.7048	21.2039	19.6447	
	PBE (Yang et al., 2023)	8.5923	19.6151	13.1283	21.9243	20.1340	18.7007	16.8947	21.9605	17.3257	19.6217	18.9500	23.5810	14.7048	21.2039	19.6170		
	TIGIC (Li et al., 2024b)	7.6630	14.6250	11.9899	21.2357	20.2022	16.9677	12.6027	19.5185	16.9192	16.6808	16.4301	16.7168	16.5398	21.5051	20.2862	18.2956	
Complex-Compo (300)	TALE (Pham et al., 2024)	8.7351	16.9899	11.3826	22.7600	20.7000	17.9581	15.0100	22.3567	19.0333	18.6267	18.7567	17.8317	15.4455	21.9667	19.9767	18.7955	
	TF-ICON (Lu et al., 2023d)	9.3259	17.7047	12.6812	21.9495	20.4799	18.2030	15.1479	21.2919	18.1770	18.3490	18.2366	17.6072	21.7027	19.7584	19.2380		
	DreamEdit (Li et al., 2023b)	8.0334	17.2630	12.0543	21.6669	20.4798	18.0200	15.0200	17.3467	15.0667	15.5534	23.5964	20.4333	20.4333	20.4333	20.4305		
	EEdit (Yan et al., 2025)	8.7835	14.9500	12.5267	22.8746	21.8152	17.7341	13.4224	23.1485	20.2601	10.81	19.7348	22.8058	15.6990	22.3498	20.5941	20.2367	
	Ours-Adapter	9.6485	22.6162	14.2525	22.5552	21.9064	20.3326	20.4582	22.6421	18.5518	21.2876	20.7349	24.0300	17.4600	22.2074	20.3913	20.9647	
	Ours-LoRA	9.8418	22.9532	13.3779	22.9130	21.7525	20.2492	20.8328	22.8261	19.1940	21.2776	21.0326	24.3400	17.8900	22.3838	19.9596	21.1212	

EEdit also proposes an elegant strategy, termed inversion skipping, to accelerate the initialization process. In EEdit, the initial latent is obtained via an inversion procedure involving model predictions, while inversion skipping significantly reduces the number of required steps. In contrast, our one-step forward diffusion does not perform inversion at all. The initial latent is produced by directly sampling noise and adding it to the clean latent at a chosen timestep, without any model prediction. This makes our initialization computationally lighter. Because FLUX-Dev is a CFG-distilled model, its inversions are relatively imprecise, which we believe contributes to the weaker subject-identity preservation observed in EEdit. Our method is therefore designed to avoid reliance on accurate inversion in such distilled settings.

L ADDITIONAL IMAGE QUALITY EVALUATION USING UNIFIEDREWARD AND HPSv3

To provide a more comprehensive assessment of composition quality, we further evaluate the methods using three variants of UnifiedReward (Wang et al., 2025b;c) (UnifiedReward-2.0-qwen3vl-8b, UnifiedReward-Edit-qwen3vl-8b, and UnifiedReward-Think-qwen-7b), as well as HPSv3 (Ma et al., 2025b). The results are summarized in Tab. 6.

M FURTHER ANALYSIS ON DSG IN SD3.5

1. distorted, deformed, glitch, artifacts
2. undefined shapes, bad anatomy, unnatural pose
3. low quality, worst quality, low resolution, blurry, out of focus
4. AI artifacts, melted objects, strange textures

The quantitative results on the DreamEditBench dataset are presented in Tab. 7.

Key Insights These results lead to the following key conclusions:

1. **Performance Sensitivity:** The performance of standard negative prompting is sensitive to the specific wording used (e.g., Prompt 2 performs better than Prompt 4). This confirms that the effectiveness of heuristic negative prompts relies heavily on manual, often tedious, prompt engineering.

1566 Table 7: Comparison of negative prompting versus DSG on SD3.5 on the DreamEditBench.
1567

Method	CLIP \uparrow	DINO \uparrow	IRF \uparrow	DreamSim \downarrow	LPIPS \downarrow	SSIM \uparrow	IR \downarrow	VR \uparrow
Ours-SD3.5-Adapter (w/ NP 1)	0.7934	0.7311	0.7608	0.3918	0.0331	0.8812	0.5837	3.5714
Ours-SD3.5-Adapter (w/ NP 2)	0.8029	0.7388	0.7679	0.3781	0.0259	0.8908	0.5726	3.6102
Ours-SD3.5-Adapter (w/ NP 3)	0.7958	0.7329	0.7623	0.3892	0.0315	0.8839	0.5814	3.5796
Ours-SD3.5-Adapter (w/ NP 4)	0.7887	0.7281	0.7574	0.3976	0.0368	0.8765	0.5882	3.5548
Ours-SD3.5-Adapter (w/ DSG)	0.8054	0.7407	0.7699	0.3745	0.0234	0.8937	0.5701	3.6187

1575 2. **Adaptive Guidance:** Standard negative prompts generate a degradation direction that is
1576 decoupled from the specifics of the input image. In contrast, DSG constructs an image-
1577 specific low-quality direction by blurring the attention component, making it significantly
1578 more adaptive, stable, and effective across diverse inputs without requiring any manual
1579 prompt tuning. DSG consistently outperforms all tested standard negative prompts across
1580 all metrics.

1581 This analysis confirms that DSG provides a superior, more robust, and automated mechanism for
1582 degradation suppression compared to traditional negative prompting, even in models like SD3.5
1583 where negative prompts are generally effective.

N RUNTIME COMPARISON

1588 The wall-clock runtime at a 512×512 resolution on an H100 GPU is summarized in the table below.
1589 In our implementation, we applied qint8 quantization to all FLUX-based methods to accelerate
1590 inference and reduce GPU memory usage. Additionally, Ours-Adapter can also be run on a 24-GB
1591 GPU by enabling CPU offloading.

1592 Table 8: Runtime and memory usage comparison of various image composition methods at a $512 \times$
1593 512 resolution.

Method	Training Free	Base Model	External Model	Time (s)	Peak Memory (MB)
MADD (He et al., 2024)	✗	SD	DINO	45.73	11708
ObjectStitch (Song et al., 2023)	✗	SD	VIT	6.63	8268
DreamCom (Lu et al., 2023c)	✗	SD	LoRA	9.87	3388
AnyDoor (Chen et al., 2024c)	✗	SD	DINO	8.61	18612
UniCombine (Wang et al., 2025a)	✗	FLUX	LoRA	11.98	22711
PBE (Yang et al., 2023)	✗	SD	-	3.52	10842
TIGIC (Li et al., 2024b)	✓	SD	-	10.82	21640
TALE (Pham et al., 2024)	✓	SD	-	8.03	23524
TF-ICON (Lu et al., 2023d)	✓	SD	-	24.55	20670
DreamEdit (Li et al., 2023b)	✓	SD	LoRA, VIT	99.83	19298
EEdit (Yan et al., 2025)	✓	FLUX	-	60.31	26546
Ours-Adapter	✓	FLUX	Adapter	38.29	32552
Ours-LoRA	✓	FLUX	LoRA	18.08	23519

O ADDITIONAL QUALITATIVE RESULTS

1612 We offer more qualitative assessment results, including visualizations of all baselines, presented in
1613 Figs. 13 to 18.

P LLM USAGE STATEMENT

1617 We used large language models for text polishing and grammar correction during manuscript prepara-
1618 tion. No LLMs were involved in the design of the method, experiments, or analysis. All content
1619 has been carefully verified and validated by the authors.




Figure 13: **(Part 1 of 2)** Qualitative comparison of our method against baselines in challenging scenarios.



Figure 14: **(Part 2 of 2)** Qualitative comparison of our method against baselines in challenging scenarios.





Figure 16: **(Part 2 of 2)** Qualitative comparison of our method against baselines in challenging scenarios.




Figure 17: **(Part 1 of 2)** Qualitative comparison of our method against baselines in challenging scenarios.




Figure 18: **(Part 2 of 2)** Qualitative comparison of our method against baselines in challenging scenarios.