

# Multitwine: Multi-Object Compositing with Text and Layout Control

Gemma Canet Tarrés<sup>1</sup>

g.canettarres@surrey.ac.uk

Zhe Lin<sup>2</sup>

zlin@adobe.com

Zhifei Zhang<sup>2</sup>

zzhang@adobe.com

He Zhang<sup>2</sup>

hezhan@adobe.com

Andrew Gilbert<sup>1\*</sup>

a.gilbert@surrey.ac.uk

John Collomosse<sup>1,2\*</sup>

j.collomosse@surrey.ac.uk

Soo Ye Kim<sup>2\*</sup>

sooyek@adobe.com

<sup>1</sup>University of Surrey

<sup>2</sup>Adobe Research

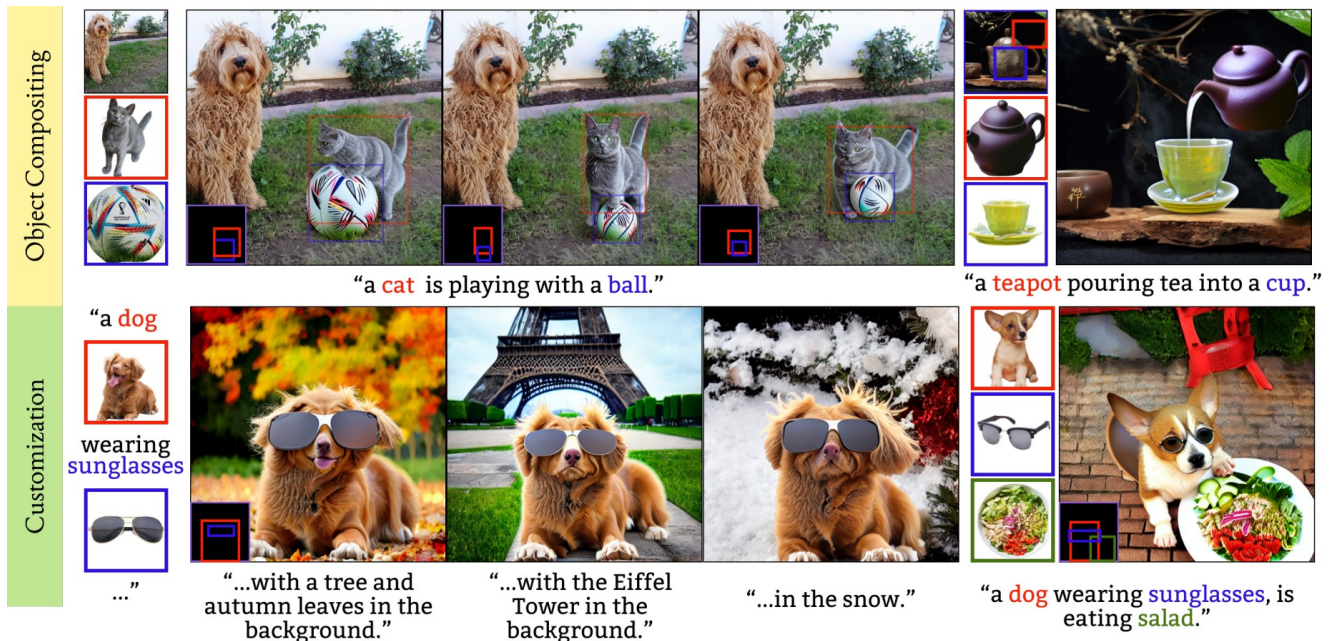


Figure 1. Our model allows simultaneous compositing of several objects with text and layout control, offering harmonious and natural results (*top row*). Object relations and actions can be controlled through text, allowing reposing and automatic addition of necessary objects (*i.e.* flowing tea). Pose and layout can be controlled via bounding boxes provided as input. Our model can also be used for customization (*bottom row*), seamlessly integrating given objects into described backgrounds.

## Abstract

We introduce the first generative model capable of simultaneous multi-object compositing, guided by both text and layout. Our model allows for the addition of multiple objects within a scene, capturing a range of interactions from simple positional relations (e.g., next to, in front of) to complex actions requiring reposing (e.g., hugging, playing guitar). When an interaction implies additional props, like ‘taking a selfie’, our model autonomously generates these supporting objects. By jointly training for compositing and

subject-driven generation, also known as customization, we achieve a more balanced integration of textual and visual inputs for text-driven object compositing. As a result, we obtain a versatile model with state-of-the-art performance in both tasks. We further present a data generation pipeline leveraging visual and language models to effortlessly synthesize multimodal, aligned training data.

## 1. Introduction

In recent years, advancements in image generation and editing have enabled the creation of realistic and complex visual scenes from various input modalities, such as text descriptions [11, 30, 33]. These tools are becoming essential for

\*: corresponding authors. This work was a joint collaboration between Adobe and the University of Surrey, conducted during an internship of the main author at Adobe.

enhancing creativity and streamlining workflows. A key component in this process is object compositing, enabling seamless integration of new objects in an existing scene. However, current compositing models [7, 40, 43, 48, 50] are generally limited to handling a single object at a time, requiring sequential steps to composite multiple objects into a scene. This process is cumbersome and fails to capture the complex, real-world interactions between multiple objects, especially when simultaneous repositioning is necessary. For instance, sequential compositing struggles to accurately create interactions where objects need to interact very closely, such as two figures hugging, or someone playing an instrument (Fig 2).

To address these limitations, we introduce a novel text-guided and layout-guided multi-object compositing model designed to handle and harmonize multiple objects simultaneously within a single composition. This model bridges the gap between single object insertion and realistic, context-aware scene construction via three main features: (i) it allows for simultaneous reposing and complex interaction between objects, seamlessly capturing intricate relationships; (ii) it ensures visual coherence between both the new objects and the scene, achieved through synchronized relighting and reharmonization; (iii) it automatically generates additional elements essential to the action or interaction being depicted (*e.g.*, a leash for a dog-walking scenario, or liquid pouring from a bottle to a glass), creating a more cohesive and natural scene.

Our model accepts multimodal input, including object images, text, object-specific bounding boxes, a background image and a mask delimiting the overall compositing region. To enable effective and balanced integration of these modalities, our model is trained on diverse, multimodal data that provides rich grounding for scene context and object relationships, allowing it to understand and implement nuanced multi-object compositions. Collecting data with all these aligned modalities is a complex task. However, advancements in Large Language Models (LLMs) [2, 10, 24] and Visual Language Models (VLMs) [25] offer a solution. We introduce a data generation pipeline that leverages these models to synthesize essential missing data, enabling the creation of a fully aligned, multimodal training set. By leveraging varied data sources during training and balancing real and synthetic data, our model learns to anticipate and position additional supporting objects, if needed, to ensure natural compositions. Furthermore, we jointly train the model on compositing and customization tasks to improve compositing performance and offer more versatility. This dual approach allows the model to separately focus on two key subtasks in multi-object compositing: (i) generating realistic interactions with strong text-image alignment, and (ii) improving inpainting, harmonization and relighting.

Our contributions are as follows:




Figure 2. Comparison of simultaneous vs. sequential object compositing. Sequential addition prevents reposing of previously composited objects, resulting in limited, less cohesive compositions.

- We present the first generative model for simultaneous multi-object compositing, addressing the limitations of sequential object compositing.
- We propose a novel data generation pipeline that combines real labeled data, synthetic captions from vision-language models, and grounding methods to align global and local captions with images for improved training.
- We leverage customization as an auxiliary task to improve text and image alignment in compositing, resulting in a model capable of performing both tasks with performance comparable to state-of-the-art customization and generative object compositing models.

## 2. Related Work

**Generative Object Compositing.** Generative object compositing has evolved from traditional methods relying on hand-crafted features [21], 3D modeling [17], and rendering [16] to leveraging the efficiency of diffusion models. Modern diffusion-based approaches, such as ObjectStitch [39] and Paint by Example [48], integrate object and background but struggle with identity preservation due to reliance on CLIP [32]. TF-ICON [26] improves on identity retention with noise modeling and composite self-attention injection, though it lacks flexibility in object reposing. [35] and [20] improve performance by either using a secondary U-Net encoder or focusing on human generation. Recent works like AnyDoor [7] and IMPRINT [40] use DINO v2 [27] for stronger identity fidelity and support more flexible shape and pose control. ControlCom [50] enables separate control over blending, harmonization, reposing and compositing, while CustomNet [49] allows control of viewpoint and location. [43] expands generation across the entire image, offering an unconstrained approach that simplifies sequential object additions. However, none of these methods support simultaneous multi-object compositing with interdependent reposing or text-based inputs for refined control, both essential for interactive, complex scenes.

**Subject-Driven Generation** Building on recent advances in text-to-image generation, subject-driven approaches aim to customize images by integrating subject-specific visuals within text prompts. Approaches like DreamBooth [34] and Textual Inversion [12] fine-tune a model to recontextualize a specific subject based on text, while others [6, 15, 38] bypass fine-tuning via large-scale upstream training. ReImagen [5], SuTI [6], FastComposer [46], ELITE [45] pro-

vide image features from image encoders directly to the U-Net, as with text control. BLIP-Diffusion [22] enables zero-shot generation by embedding objects into random backgrounds but remains limited in handling multiple entities. GroundingBooth [47] provides a model for grounding text-to-image generation, specifying a bounding box for each object as part of the input. Leveraging Multimodal Large Language Models (MLLMs) [2, 10, 44] offers efficient integration of text and visual inputs. GILL [18], Emu [41], and DreamLLM [9] align MLLM outputs with diffusion encoders for combined visual-language generation, though struggle to retain fine-grained visual details due to alignment at semantic level. KOSMOS-G [28] and UNIMO-G [23] combine a MLLM as text and image encoder with a U-Net from Stable Diffusion as image decoder by either training a subnetwork that allows their integration or performing end-to-end training of the U-Net, respectively. Emu2Gen [42] further adds support for bounding box guidance for each subject, enriching image layout control.

### 3. Methodology

We propose a diffusion-based model capable of simultaneously compositing multiple objects into a background in a controlled way, guided by both a layout and a text prompt. Training this model effectively requires extensive paired data, including: (i) ground truth images containing multiple objects, (ii) text descriptions with grounding information for these objects, (iii) segmented images of the objects, and (iv) their bounding boxes in the ground truth image. For simplicity, we focus on two objects at a time, which may interact through complex actions or positional relationships in the ground truth image. We present model architecture in Section 3.1, training and inference strategies in Sections 3.2 and 3.3, and data generation pipeline in Section 3.4.

#### 3.1. Model Architecture

Our model takes as input a background image  $\mathcal{I}_{BG}$ , a layout  $\mathcal{I}_L$ ,  $N$  object images  $\mathcal{O}_i$ , with  $i \in \{0 \dots, N-1\}$ , and a descriptive caption  $\mathcal{C}$ . The layout includes a bounding box for each object ( $\mathcal{M}_{0 \dots, N-1}$ ), along with a larger box  $\mathcal{M}_G$  that encloses them all, defining the region to be modified (inpainting region), allowing for additional room for object-object interactions. We encode all layout information on a single mask, where different values are assigned to the pixels belonging to each object’s bounding box, their overlapping, the rest of the inpainting region and the background. As shown in Fig 3,  $\mathcal{I}_L$  is concatenated to the 3-channel input noise  $\mathcal{I}$  and  $(1 - \mathcal{M}_G) * \mathcal{I}_{BG}$ , a version of  $\mathcal{I}_{BG}$ , where the inpainting region is masked out. These concatenated images are fed into the model’s backbone, Stable Diffusion 1.5 (SD) [33], consisting of a variational autoencoder ( $\mathcal{G}$ ,  $\mathcal{D}$ ) and a U-Net. Each composited object  $\mathcal{O}_i$  is processed by an image encoder  $\mathcal{E}_I$  (DINO ViT-G/14 [27]) and a con-

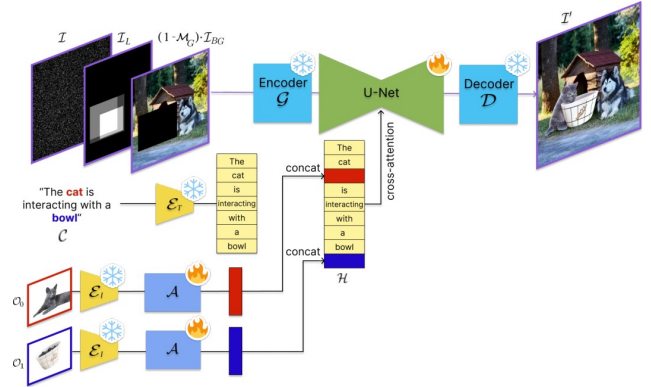


Figure 3. Model Architecture. Our model consists of: (i) A Stable Diffusion backbone including a U-Net and an autoencoder ( $\mathcal{G}$ ,  $\mathcal{D}$ ); (ii) a text encoder  $\mathcal{E}_T$ ; (iii) an image encoder  $\mathcal{E}_I$ ; and (iv) an adaptor  $\mathcal{A}$ . Given a text prompt  $\mathcal{C}$  and images of  $N$  objects  $\mathcal{O}_{0 \dots, N-1}$ , the text embedding from (iii) is augmented by automatically appending each image embedding after its corresponding text tokens. The resulting multimodal embedding  $\mathcal{H}$  is fed to the U-Net via cross-attention. Masked background image  $(1 - \mathcal{M}_G) * \mathcal{I}_{BG}$  and layout  $\mathcal{I}_L$  with object-specific bboxes are concatenated to input  $\mathcal{I}$ .

tent adaptor  $\mathcal{A}$ , which aligns the object embeddings with text embeddings as in [39]. The caption  $\mathcal{C}$  is encoded using a text encoder  $\mathcal{E}_T$  (CLIP ViT-L/14 [32]). These embeddings are then combined into a multimodal embedding and fed into the U-Net via cross-attention.

**Multimodal Embeddings** Providing the model with text and image information in a balanced and interpretable way is one of the main challenges of our method. It must perform well when either input is missing, and neither component must dominate when both are present. The identity of each object  $\mathcal{O}_i$  must be preserved in the output while matching the scene and interactions described by the caption  $\mathcal{C}$ . Grounding information specifies the subset of words  $\mathcal{C}_i$  in the caption  $\mathcal{C}$  that correspond to each object image  $\mathcal{O}_i$ . After encoding the text, each object embedding  $\mathcal{A}(\mathcal{E}_I(\mathcal{O}_i))$  is concatenated after  $\mathcal{E}_T(\mathcal{C}_i)$ , resulting in a multimodal embedding  $\mathcal{H}$  that is then passed to the U-Net via cross-attention. We refer to the set of visual and textual information referring to  $\mathcal{C}_i$  and  $\mathcal{O}_i$  in the embedding  $\mathcal{H}$  as  $\mathcal{H}_i$ .

#### 3.2. Training Strategy

Our training strategy is designed to balance textual and visual information, enabling object reposing and the generation of described backgrounds and new objects while preserving each object’s identity and preventing identity mixing. To ensure the model is still able to perform with only text or image input, we randomly drop each modality with a 30% probability during training.

**Customization as an Auxiliary Task** To enhance compositing performance and provide a more versatile model, we jointly train on multi-object compositing and multi-entity subject-guided generation. With 50% probability, the

inpainting region  $\mathcal{M}_G$  is replaced by a mask covering the entire image, and  $\mathcal{I}_{BG}$  is replaced with an empty image. In this customization setting, the model can focus on preserving object identity while reposing objects and generating scenes aligned with the text, without being burdened by compositing tasks like harmonization, relighting, or background inpainting. Introducing this joint training results in a better balance between textual and visual alignment.

**Object Identity Disentanglement** Diffusion models often struggle to accurately represent multiple subjects due to attention layers blending visual features and causing semantic leakage between them. Cross-attention maps indicate how text tokens influence latent pixels [13]. When tokens share similar semantics, a single pixel may respond to all, leading to identity blending. Additionally, self-attention features create dense correspondences within the same subject and across semantically similar ones [1, 4, 8]. While this behavior aids in generating coherent images with well-integrated subjects and backgrounds, it can also cause visual feature leakage between similar subjects [8]. To reduce this leakage and encourage identity separation, we introduce two additional losses ( $\mathcal{L}_c, \mathcal{L}_s$ ) inspired by [8, 46]:

$$\mathcal{L}_c = \sum_{i=0}^{N-1} \frac{1}{N} \left( 1 - \frac{\sum_{\substack{\mathbf{x} \in \mathcal{S}_i, \\ \mathbf{h} \in \mathcal{H}_i}} \hat{\mathbf{A}}[\mathbf{x}, \mathbf{h}]}{\sum_{\substack{\mathbf{x} \in \mathcal{S}_i, \\ \mathbf{h} \in \mathcal{H}_i}} \hat{\mathbf{A}}[\mathbf{x}, \mathbf{h}] + \sum_{\substack{\mathbf{x} \notin \mathcal{S}_i, \\ \mathbf{h} \in \mathcal{H}_i}} \hat{\mathbf{A}}[\mathbf{x}, \mathbf{h}]} \right), \quad (1)$$

$$\mathcal{L}_s = \sum_{i=0}^{N-1} \frac{1}{N} \left( 1 - \frac{1}{1 + \sum_{\substack{\mathbf{x} \in \mathcal{S}_i, \\ \mathbf{y} \in \mathcal{S}_j, j \neq i}} \hat{\mathbf{A}}[\mathbf{x}, \mathbf{y}]} \right), \quad (2)$$

where  $\hat{\mathbf{A}}$  is the mean attention map, averaged across heads and layers,  $\mathcal{S}_i$  is the segmentation map of  $\mathcal{O}_i$  in the ground truth image, and  $\mathbf{x}, \mathbf{y}$  correspond to pixel coordinates in  $\mathcal{I}_t$  (the noisy version of  $\mathcal{I}$  at timestep  $t$ ). For each  $\mathbf{x}$ ,  $\mathcal{L}_c$  encourages the cross-attention maps obtained from visual-language information from  $\mathcal{C}_i$  and  $\mathcal{O}_i$  ( $\mathbf{h} \in \mathcal{H}_i$ ) to be close to their corresponding segmentation map  $\mathcal{S}_i$ .  $\mathcal{L}_s$  discourages pixels  $\mathbf{x} \in \mathcal{S}_i$  to attend to pixels  $\mathbf{y} \in \mathcal{S}_j, \forall j \neq i$ .

Based on the inpainting version of Stable Diffusion v1.5 [33], our model is fine-tuned by optimizing a combination of three losses:  $\mathcal{L} = \mathcal{L}_d + \alpha \mathcal{L}_c + \beta \mathcal{L}_s$ .

$$\mathcal{L}_d = \mathbb{E}_{\mathcal{I}_c, \mathcal{H}, t, \epsilon \sim \mathcal{N}(0,1)} \left[ \|\epsilon - \epsilon_\theta(\mathcal{I}_c, \mathcal{H}, t)\|_2^2 \right], \quad (3)$$

where  $\mathcal{I}_c = [\mathcal{I}_t, \mathcal{I}_L, (1 - \mathcal{M}_G) * \mathcal{I}_{BG}]$  ( $\mathcal{I}_t$ : noisy version of the input image  $\mathcal{I}$  at timestep  $t$ ,  $\mathcal{I}_L$ : layout,  $(1 - \mathcal{M}_G) * \mathcal{I}_{BG}$ : masked background image,  $[\cdot]$ : concatenation operation across the channel dimension),  $\mathcal{H}$ : multi-modal embedding,  $\epsilon_\theta$ : denoising model being optimized.

### 3.3. Inference Strategy

Although attention-based losses help disentangle visual and semantic features of different objects during training, some information leakage can still occur at inference. To address this, and to further encourage the desired layout  $\mathcal{I}_L$ , all cross-attention scores corresponding to  $\mathbf{h} \in \mathcal{H}_i$  are masked with  $\mathcal{M}_i$ . Additionally, since  $[EoT]$  tokens are known to contain information about foreground objects [51], we mask their corresponding cross-attention scores with the union of all object masks  $\bigcup_{i=0 \dots N-1} \mathcal{M}_i$ .

### 3.4. Training Data Generation

Balanced, curated data is essential for effective diffusion model training. For our model to learn correct object interaction, natural compositing, text-image alignment and identity preservation, paired training data must meet specific criteria: (i) diverse range of objects and relationships; (ii) varied text prompt styles; (iii) object images with different poses, views and lighting than the ground truth image; (iv) high-quality images; and (v) a large dataset size. Since no single source can provide all these qualities, we use three complementary data sources: Video Data, In-the-Wild Images, and Manually Collected Data. See below for brief descriptions and SuppMat for more specific details.

**Video Data** We use annotated data from video object relation datasets [36, 37], where each frame includes bounding boxes identifying objects and actions or positional relationships between object pairs as  $\langle object A \rangle \langle relation \rangle \langle object B \rangle$ . For each annotated relation, we select one frame depicting it as the ground truth image, a second frame including *object A* and a third with *object B*. This provides alternative views for each object (segmented via [31]), which is essential for training the model to learn object reposing for intended interactions. To ensure identity consistency, we keep DINO score [27] similarity for each object’s views above a threshold ( $MSE \geq 0.8$ ). We provide LLaVA (13B) [24] with ground truth images and annotated relations, obtaining grammatically correct captions to use as aligned text prompts. Although this data source meets most of our criteria, it typically involves lower-quality frames.

**Image Data** An automated pipeline for paired data generation from in-the-wild images can provide unlimited data with high-quality images to enhance model training. We consider two approaches: top-down and bottom-up. **(i) Top-down approach:** Starting with a single image, we use a commercial subject selection tool to identify the main objects. A semantic segmentator [31] helps us select images where multiple key objects are present. From these, we randomly pick two entities to serve as composited objects. We outline each object in a different color on the original image and input it into ViP-LLaVA (13B) [3], which sequentially identifies the selected entities and generates a caption describing their relationship. **(ii) Bottom-up approach:** To

increase caption diversity, we include open-source images with ground-truth captions, such as those in OpenImages [19]. Given an image and its caption, GroundingDINO [25] extracts grounding information that correlates the two. Using a semantic segmentator [31] and filtering for quality (e.g., removing duplicates, background entities, and objects that are too large (>75%) or too small (<10%)), we randomly select two grounded entities as composited objects. This approach adds diversity to image layouts and caption formats, enhancing model robustness through a mix of local and global descriptions. However, neither of the approaches provides different views of the objects, making them unsuitable as stand-alone sources for model training.

**Manually Collected Data** We manually compile a dataset containing 16,896 images of humans interacting with objects, 3,898 images of the same humans in neutral poses, and 1,250 images of those objects alone, captured from different angles and scenes. This data is segmented using [31] and their caption and grounding information is obtained using ViP-LLaVA (13B) [3], following the same process described above. This collection provides a small curated set of high-quality images, including different views of each entity and covering human-object interactions, one of the most intricate relationships we aim to replicate.

## 4. Experiments

**Evaluation Dataset** We evaluate image and text alignment in the customization task using the set of 300 paired two entities and text from *MultiBench* [23], a benchmark crafted for evaluating multi-entity subject-driven generation. We manually craft a bounding box for each category (e.g. food, toy) to be used as input for our model and [42]. For evaluating multi-object compositing, we create a test set of 118 paired data (background image, text prompt, two entities and their corresponding bounding boxes). We ensure different positional and action-based prompts are included, and entities include objects, animals and humans from [23, 29, 34]. We also evaluate identity preservation on *Dreambooth*, the single-object compositing set consisting of 113 background-object pairs used in [40, 43].

**Evaluation Metrics** We assess alignment to input text prompt and object images via CLIP-Score [14] (CLIP-T, CLIP-I) and DINO-Score [27] (DINO). When bounding boxes are specified for each object, we compare each input object image to a corresponding cropped region within the generated image, as in [7, 43, 48]. In the customization setting, where models may not specify object locations, each input object is instead compared to the entire generated image, following [23, 28]. We refer to this metric variation as CLIP-*Igl*, DINO-*gl*. For scenes with multiple objects, we average DINO(*gl*) and CLIP-I(*gl*) scores across all objects. For compositing multiple objects with single-object compositing models [7, 40, 43, 48, 50], we add objects sequentially. Thus, we compute performance for all possi-

Method	DreamBooth		MultiComp - <i>overlap</i>		MultiComp - <i>nonoverlap</i>	
	CLIP-I $\uparrow$	DINO $\uparrow$	CLIP-I $\uparrow$	DINO $\uparrow$	CLIP-I $\uparrow$	DINO $\uparrow$
PbE [48]	0.778	0.799	0.693	0.383	0.720	0.423
ControlCom [50]	0.743	0.705	0.707	0.478	0.740	0.543
AnyDoor [7]	0.806	0.836	0.727	0.520	0.763	<b>0.593</b>
IMPRINT [40]	<b>0.830</b>	0.889	0.713	0.525	0.739	0.576
TOTB [43]	0.809	0.856	0.716	0.485	0.740	0.531
Ours	0.803	0.892	<b>0.741</b>	<b>0.532</b>	<b>0.768</b>	0.579
w/o text	0.816	<b>0.903</b>	0.729	0.505	0.754	0.548
sequential w/ text	-	-	0.729	0.517	0.760	0.583
sequential w/o text	-	-	0.723	0.510	0.756	0.578

Table 1. Quantitative comparison of identity preservation against state-of-the-art generative object compositing methods. Single object compositing is evaluated on DreamBooth set. Two object compositing is on MultiComp set, comparing to both simultaneous and sequential uses of our model, by guiding inference with and without text. We distinguish between interacting (overlapping bboxes) and non-interacting cases for clarity. Details in SuppMat.

ble object-order sequences and average the results. To assess alignment between text prompts and composite results, we calculate CLIP-T by comparing the entire image to the text prompt (CLIP-*Tgl*) and by comparing the cropped compositing area to the text prompt (CLIP-*Tloc*), evaluating local and global alignment. Additionally, we conduct user studies to evaluate alignment to input modalities, quality of composited images, and realism of generated interactions.

**Training Details** The U-Net and adaptor  $\mathcal{A}$  are jointly trained on 8 A100, using Adam optimizer, learning rate of  $4 \times 10^{-6}$  and effective batch size of 1024, via gradient accumulation. Losses are balanced using  $\alpha = 10^3$ ,  $\beta = 1$ .

### 4.1. Comparison to Existing Methods

Our primary task is Object Compositing, but we also train for Subject-Driven Generation as an auxiliary task. Thus, we provide comprehensive evaluation across both tasks.

**Generative Object Compositing** For this task, we compare to recent generative object compositing models: Paint by Example (PbE) [48], ControlCom [50], AnyDoor [7], IMPRINT [40], Thinking Outside the BBox (TTOB) [43]. These models support compositing from a single object, background, and bounding box, requiring sequential runs to add multiple objects individually. For comparison, we evaluate our model in three additional inference modes: (i) without text guidance, (ii) sequential single-object compositing, (iii) sequential compositing without text guidance.

As shown in Table 1, our model outperforms all compared models when composited objects interact, *i.e.* when bounding boxes overlap. For non-overlapping or single-object compositions, our model maintains comparable identity and semantic preservation to state-of-the-art models. Furthermore, when objects overlap, compositing two objects simultaneously yields significantly better results than adding them sequentially, even with the same model, as in Fig 2. Providing a text description of the interaction also boosts performance for scenes with multiple objects but is less critical when compositing a single object or non-



Figure 4. Visual comparison to Generative Object Compositing models [7, 40, 43, 48, 50]. Our model provides more realistic, harmonious and natural-looking interaction between composited objects via simultaneous multi-object compositing. See SuppMat for more examples.

interacting ones, where the expected composition is clearer. Fig 4 illustrates several advantages of compositing multiple objects simultaneously rather than sequentially: (i) it enables more cohesive harmonization and appearance consistency across objects and the scene (rows 1 - 5); (ii) it captures complex interactions involving object reposing with ease (rows 1, 6, 7); and (iii) with text guidance, the model can naturally complete scenes by adding any additional elements needed for realism (row 7).

**Multi-Entity Subject-Driven Generation** We compare to customization methods accepting text and multiple object images as input (BLIP-Diffusion [22], KOSMOS-G [28], UNIMO-G [23]). We also compare to Emu2Gen [42], addi-

tionally accepting layout guidance, therefore sharing input modalities with our model and allowing fair comparison.

Although customization is introduced as a proxy task to improve multi-object compositing, our model achieves comparable performance to state-of-the-art customization models, as shown in Table 2. Notably, our model obtains improved text alignment, benefiting from training data balancing local and global caption descriptions, as well as joint training with the complex task of object compositing. Additionally, our model demonstrates improved performance in patch-based image metrics (CLIP-I, DINO), indicating high layout alignment and identity preservation. In contrast, Emu2Gen obtains lower patch-based than global metrics.

Method	CLIP-I $\uparrow$	DINO $\uparrow$	CLIP-Igl $\uparrow$	DINOgl $\uparrow$	CLIP-Tgl $\uparrow$
Input: <i>Text, Object Images</i>					
BLIP-Diffusion [22]	-	-	0.675	0.455	0.249
KOSMOS-G [28]	-	-	<b>0.704</b>	0.465	0.279
UNIMO-G [23]	-	-	0.699	<b>0.485</b>	0.293
Input: <i>Text, Object Images, Layout</i>					
Emu2-Gen [42]	0.595	0.414	0.616	0.434	0.287
Ours	<b>0.783</b>	<b>0.599</b>	0.688	0.454	<b>0.308</b>

Table 2. Quantitative comparison of identity preservation and text fidelity against multi-entity subject-driven generation methods. We compare to state-of-the-art methods on two-entity subset of MultiBench. For models with layout guidance, additional identity preservation metrics (CLIP-I, DINO) are provided by considering cropped regions around each object in the generated images.

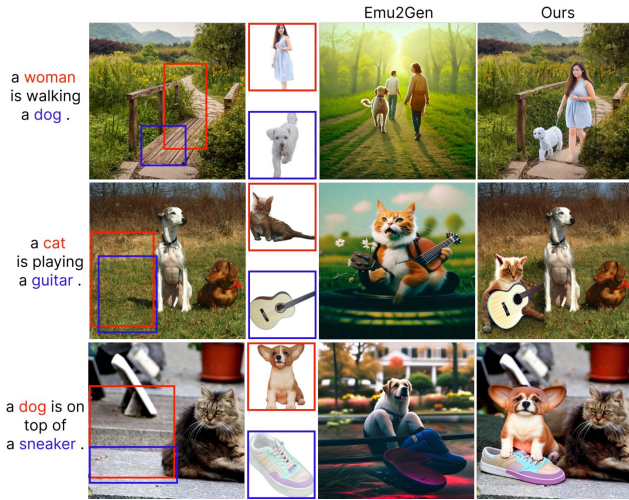


Figure 5. Visual comparison of our multi-object compositing model to Emu2Gen [42]. For fair comparison, exact same set of inputs (including background image) is provided to each model.

Refer to Fig 1 (bottom) and SuppMat for visual examples. **User Studies** We complement the comparison to State-of-Art models by completing various user studies (Fig 6). Non-expert users are presented with side-by-side generations from our model and each of the object compositing baselines, and are asked to choose their preferred image based on ‘most realistic interaction’ (in cases of mask overlapping) and ‘best image quality’. Due to our model’s ability to composite multiple interacting or non-interacting objects in a natural way, users prefer our model with up to 66.7% preference in terms of image quality and up to 97.1% preference for realistic interactions, via majority consensus. We also assess alignment to each individual input modality by comparing our compositing model to Emu2Gen [42]. Emu2Gen is provided with the same inputs (Fig 5), presented as a multimodal input interlacing text, object images, object-specific bounding boxes and background. Background image is provided by adding ‘ $in < \mathcal{I}_{BG} >$ ’ to the multimodal input, as showcased in [42]. In this comparison, users prefer our alignment across all modalities. See SuppMat for details.




Figure 6. User Studies. *Top*: Percentage of users preferring our method or each baseline [7, 40, 43, 48, 50] on ‘image quality’ and ‘realistic interaction’. *Bottom*: Percentage of users preferring our method or Emu2Gen [42] in terms of text, layout, objects and background alignment. All results are via majority consensus.

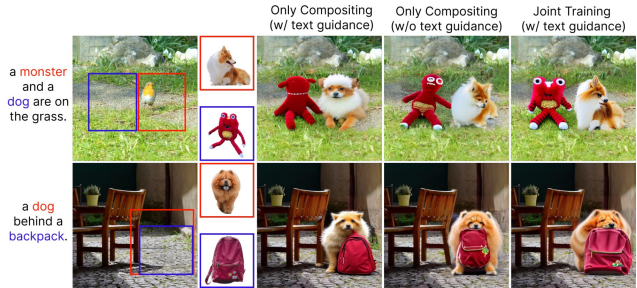


Figure 7. Visualization of the effect of Joint Training for Compositing and Customization tasks. When training solely on object compositing, balancing text and image alignment becomes significantly more complex. In this setup, the ablation model significantly drops identity preservation when inference is also guided with text. Our final model achieves better visual-language balance.

## 4.2. Ablation Study

Table 3 displays identity and text fidelity metrics for various ablations of our model. (i-ii) As seen in Fig 8, adding attention-based losses helps prevent semantic and visual information leakage between objects, improving both identity preservation and text alignment, as well as the overall image quality. (iii) When training solely for object compositing, the model must learn to add new objects to a scene, repose them according to actions described in text input, complete the scene by adhering to the same text, and harmonize and relight the image to make everything look seamless while preserving object identity. By introducing customization as an auxiliary task and alternating between both tasks during



Figure 8. Visualization of the effect of  $\mathcal{L}_c$ ,  $\mathcal{L}_s$ . Without attention-based losses, our model merges objects with similar semantics or visual traits. Cross-attention loss  $\mathcal{L}_c$  improves identity separation, but some leakage remains (e.g., cat ears on dog and teapot). Adding self-attention loss  $\mathcal{L}_s$  further reduces this leakage.

Method	DINO $\uparrow$	CLIP-I $\uparrow$	CLIP-Tloc $\uparrow$	CLIP-Tgl $\uparrow$
MultiComp-action				
Ours	<b>0.540</b>	0.745	0.286	0.285
w/o self loss	0.538	0.744	0.283	0.284
w/o attn loss	0.534	0.739	0.270	0.274
w/o customiz.	0.449	0.705	<b>0.295</b>	<b>0.298</b>
w/o multi-view	0.535	<b>0.751</b>	0.268	0.273
w/o inf. mask	0.448	0.712	0.283	0.280
MultiComp-positional				
Ours	0.585	0.773	0.292	0.294
w/o self loss	0.584	0.770	0.289	0.291
w/o attn loss	0.581	0.770	0.281	0.285
w/o customiz.	0.474	0.734	<b>0.298</b>	<b>0.298</b>
w/o multi-view	<b>0.591</b>	<b>0.780</b>	0.280	0.283
w/o inf. mask	0.436	0.714	0.286	0.290

Table 3. Quantitative comparison of text/image alignment to different ablations of our compositing model. For better analysis we distinguish between *action* and *positional* subsets of MultiComp set, based on objects relationship in text description. We compare our full model to different versions obtained without (i) self-attention loss  $\mathcal{L}_s$ ; (ii) self and cross-attention losses ( $\mathcal{L}_c, \mathcal{L}_s$ ); (iii) joint training for compositing and customization; (iv) multi-view data (i.e. video, manually collected set); (v) inference masking.

training, the model can sometimes relieve the inpainting aspect of compositing, focusing only on balancing text alignment and identity preservation. As in Fig 7, in the absence of the customization task, the model struggles to balance text and image preservation, and the objects become irre-recognizable. (iv) Without multi-view data, even if visual and geometrical transformations are applied to the object images, the model lacks the ability to reposition them to match the actions described in the text. This results in a ‘copy-paste’ behavior where objects are minimally transformed (high CLIP-I, DINO), preserving identity but leading to highly unnatural results and low text alignment (CLIP-Tgl,loc). (v) The masking step at inference ensures object locations are accurately aligned with input masks, and further disentangles identities. Therefore, removing it significantly lowers DINO/CLIP-I scores due to displaced objects and identity leakage, while CLIP-Tlog,gl remain mostly stable.

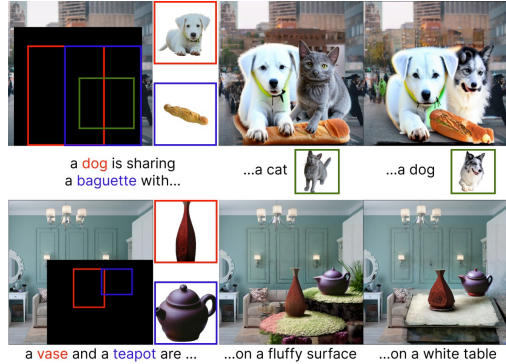


Figure 9. Multitwine Applications. *Top*: Multi-Object Compositing of Three Objects. *Bottom*: Subject-Driven Inpainting.

### 4.3. Applications

Although not explicitly trained for them, our model exhibits some emerging capabilities. See SuppMat for more.

**Multi-Object Generation** As shown in Fig 9 (top), our model is able to perform multi-object compositing with more than two objects. By learning one-to-one object interactions, it develops a strong prior that allows it to generalize to compositing multiple objects simultaneously.

**Subject-Driven Inpainting** Our joint training enables the model to learn key subtasks like background synthesis, blending, harmonization, and reposing. As shown in Fig 9 (bottom), these skills apply to subject-driven inpainting, where the model uses text and layout guidance to seamlessly complete scenes, generating and integrating objects while maintaining natural coherence with the given visuals.

### 4.4. Limitations

Although our model can handle two or more objects, the pipeline is not designed for an unlimited number. As more objects are added, the multimodal embedding grows, which could lead to scalability issues. A potential future improvement could involve feeding each object embedding separately, impacting only corresponding attention maps. Additionally, we base our model on SD1.5 due to compute limitation. Overall harmonization and composition quality could be further improved with stronger diffusion base model such as SDXL [30] or SD3 [11].

## 5. Conclusion

In this paper, we introduce the first generative model capable of simultaneously compositing multiple objects in a natural and realistic manner, guided by text and layout. We demonstrate how joint training on object compositing and subject-driven generation improves model performance, achieving state-of-the-art results in both tasks. This multi-object, dual training approach enhances flexibility, enabling applications like subject-driven inpainting and more realistic object interactions. We hope this work inspires further research into more controllable and effective object compositing techniques.



## 6. Acknowledgements

This work was partially supported by Adobe, via an internship, and by DECaDE (the UKRI Centre for Decentralized Creative Economy) under EPSRC Grant EP/T022485/1.

## References

- [1] Yuval Alaluf, Daniel Garibi, Or Patashnik, Hadar Averbuch-Elor, and Daniel Cohen-Or. Cross-image attention for zero-shot appearance transfer. In *ACM SIGGRAPH 2024 Conference Papers*, pages 1–12, 2024. 4
- [2] Jean-Baptiste Alayrac, Jeff Donahue, Pauline Luc, Antoine Miech, Iain Barr, Yana Hasson, Karel Lenc, Arthur Mensch, Katherine Millican, Malcolm Reynolds, et al. Flamingo: a visual language model for few-shot learning. *Advances in neural information processing systems*, 35:23716–23736, 2022. 2, 3
- [3] Mu Cai, Haotian Liu, Siva Karthik Mustikovela, Gregory P Meyer, Yuning Chai, Dennis Park, and Yong Jae Lee. Vip-llava: Making large multimodal models understand arbitrary visual prompts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12914–12923, 2024. 4, 5, 1, 2
- [4] Mingdeng Cao, Xintao Wang, Zhongang Qi, Ying Shan, Xiao-hu Qie, and Yinqiang Zheng. Masactrl: Tuning-free mutual self-attention control for consistent image synthesis and editing. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22560–22570, 2023. 4
- [5] Wenhui Chen, Hexiang Hu, Chitwan Saharia, and William W Cohen. Re-Imagen: Retrieval-augmented text-to-image generator. *arXiv preprint arXiv:2209.14491*, 2022. 2
- [6] Wenhui Chen, Hexiang Hu, Yandong Li, Nataniel Ruiz, Xuhui Jia, Ming-Wei Chang, and William W Cohen. Subject-driven text-to-image generation via apprenticeship learning. *Advances in Neural Information Processing Systems*, 36, 2024. 2
- [7] Xi Chen, Lianghua Huang, Yu Liu, Yujun Shen, Deli Zhao, and Hengshuang Zhao. Anydoor: Zero-shot object-level image customization. *arXiv preprint arXiv:2307.09481*, 2023. 2, 5, 6, 7
- [8] Omer Dahary, Or Patashnik, Kfir Aberman, and Daniel Cohen-Or. Be yourself: Bounded attention for multi-subject text-to-image generation. *arXiv preprint arXiv:2403.16990*, 2(5), 2024. 4
- [9] Runpei Dong, Chunrui Han, Yuang Peng, Zekun Qi, Zheng Ge, Jinrong Yang, Liang Zhao, Jianjian Sun, Hongyu Zhou, Haoran Wei, et al. Dreamllm: Synergistic multimodal comprehension and creation. *arXiv preprint arXiv:2309.11499*, 2023. 3
- [10] Danny Driess, Fei Xia, Mehdi SM Sajjadi, Corey Lynch, Aakanksha Chowdhery, Brian Ichter, Ayzaan Wahid, Jonathan Tompson, Quan Vuong, Tianhe Yu, et al. Palme: An embodied multimodal language model. *arXiv preprint arXiv:2303.03378*, 2023. 2, 3
- [11] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Forty-first International Conference on Machine Learning*, 2024. 1, 8
- [12] Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H Bermano, Gal Chechik, and Daniel Cohen-Or. An image is worth one word: Personalizing text-to-image generation using textual inversion. *arXiv preprint arXiv:2208.01618*, 2022. 2
- [13] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. *arXiv preprint arXiv:2208.01626*, 2022. 4
- [14] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. In *Proceedings of the Conference on Empirical Methods in Natural Language Processing*, pages 7514–7528, 2021. 5
- [15] Xuhui Jia, Yang Zhao, Kelvin CK Chan, Yandong Li, Han Zhang, Boqing Gong, Tingbo Hou, Huisheng Wang, and Yu-Chuan Su. Taming encoder for zero fine-tuning image customization with text-to-image diffusion models. *arXiv preprint arXiv:2304.02642*, 2023. 2
- [16] Kevin Karsch, Varsha Hedau, David Forsyth, and Derek Hoiem. Rendering synthetic objects into legacy photographs. *ACM Transactions on graphics (TOG)*, 30(6):1–12, 2011. 2
- [17] Natasha Kholgade, Tomas Simon, Alexei Efros, and Yaser Sheikh. 3d object manipulation in a single photograph using stock 3d models. *ACM Transactions on graphics (TOG)*, 33(4):1–12, 2014. 2
- [18] Jing Yu Koh, Daniel Fried, and Russ R Salakhutdinov. Generating images with multimodal language models. *Advances in Neural Information Processing Systems*, 36, 2024. 3
- [19] Ilya Krylov, Sergei Nosov, and Vladislav Sovrasov. Open images v5 text annotation and yet another mask text spotter. In *Asian Conference on Machine Learning*, pages 379–389. PMLR, 2021. 5
- [20] Sumith Kulal, Tim Brooks, Alex Aiken, Jiajun Wu, Jimei Yang, Jingwan Lu, Alexei A Efros, and Krishna Kumar Singh. Putting people in their place: Affordance-aware human insertion into scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17089–17099, 2023. 2
- [21] Jean-François Lalonde, Derek Hoiem, Alexei A Efros, Carsten Rother, John Winn, and Antonio Criminisi. Photo clip art. *ACM transactions on graphics (TOG)*, 26(3):3–es, 2007. 2
- [22] Dongxu Li, Junnan Li, and Steven Hoi. Blip-diffusion: Pre-trained subject representation for controllable text-to-image generation and editing. *Advances in Neural Information Processing Systems*, 36, 2024. 3, 6, 7, 2
- [23] Wei Li, Xue Xu, Jiachen Liu, and Xinyan Xiao. Unimog: Unified image generation through multimodal conditional diffusion. *arXiv preprint arXiv:2401.13388*, 2024. 3, 5, 6, 7, 2
- [24] Haotian Liu, Chunyuan Li, Yuheng Li, and Yong Jae Lee. Improved baselines with visual instruction tuning, 2023. 2, 4, 1

- [25] Shilong Liu, Zhaoyang Zeng, Tianhe Ren, Feng Li, Hao Zhang, Jie Yang, Qing Jiang, Chunyuan Li, Jianwei Yang, Hang Su, et al. Grounding dino: Marrying dino with grounded pre-training for open-set object detection. *arXiv preprint arXiv:2303.05499*, 2023. 2, 5, 1
- [26] Shilin Lu, Yanzhu Liu, and Adams Wai-Kin Kong. Tf-icon: Diffusion-based training-free cross-domain image composition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2294–2305, 2023. 2
- [27] Maxime Oquab, Timothée Darcet, Théo Moutakanni, Huy Vo, Marc Szafranec, Vasil Khalidov, Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, et al. Dinov2: Learning robust visual features without supervision. *arXiv preprint arXiv:2304.07193*, 2023. 2, 3, 4, 5, 1
- [28] Xichen Pan, Li Dong, Shaohan Huang, Zhiliang Peng, Wenhui Chen, and Furu Wei. Kosmos-g: Generating images in context with multimodal large language models. *arXiv preprint arXiv:2310.02992*, 2023. 3, 5, 6, 7, 2
- [29] Pixabay. <https://pixabay.com/>, 2024. 5, 2
- [30] 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. *arXiv preprint arXiv:2307.01952*, 2023. 1, 8
- [31] Lu Qi, Jason Kuen, Weidong Guo, Tiancheng Shen, Jiuxiang Gu, Wenbo Li, Jiaya Jia, Zhe Lin, and Ming-Hsuan Yang. Fine-grained entity segmentation. *arXiv preprint arXiv:2211.05776*, 2022. 4, 5, 1, 2, 3
- [32] 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*, pages 8748–8763. PMLR, 2021. 2, 3
- [33] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 10684–10695, 2022. 1, 3, 4
- [34] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 22500–22510, 2023. 2, 5
- [35] Mehmet Saygin Seyfioglu, Karim Bouyarmene, Suren Kumar, Amir Tavanaei, and Ismail B Tutar. Diffuse to choose: Enriching image conditioned inpainting in latent diffusion models for virtual try-all. *arXiv preprint arXiv:2401.13795*, 2024. 2
- [36] Xindi Shang, Tongwei Ren, Jingfan Guo, Hanwang Zhang, and Tat-Seng Chua. Video visual relation detection. In *ACM International Conference on Multimedia*, Mountain View, CA USA, 2017. 4, 1
- [37] Xindi Shang, Donglin Di, Junbin Xiao, Yu Cao, Xun Yang, and Tat-Seng Chua. Annotating objects and relations in user-generated videos. In *Proceedings of the 2019 on International Conference on Multimedia Retrieval*, pages 279–287. ACM, 2019. 4, 1
- [38] Jing Shi, Wei Xiong, Zhe Lin, and Hyun Joon Jung. Instant-booth: Personalized text-to-image generation without test-time finetuning. *arXiv preprint arXiv:2304.03411*, 2023. 2
- [39] Yizhi Song, Zhifei Zhang, Zhe Lin, Scott Cohen, Brian Price, Jianming Zhang, Soo Ye Kim, and Daniel Aliaga. Objectstitch: Generative object compositing. *arXiv preprint arXiv:2212.00932*, 2022. 2, 3
- [40] Yizhi Song, Zhifei Zhang, Zhe Lin, Scott Cohen, Brian Price, Jianming Zhang, Soo Ye Kim, He Zhang, Wei Xiong, and Daniel Aliaga. Imprint: Generative object compositing by learning identity-preserving representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8048–8058, 2024. 2, 5, 6, 7
- [41] Quan Sun, Qiyang Yu, Yufeng Cui, Fan Zhang, Xiaosong Zhang, Yuezhe Wang, Hongcheng Gao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Generative pretraining in multimodality. *arXiv preprint arXiv:2307.05222*, 2023. 3
- [42] Quan Sun, Yufeng Cui, Xiaosong Zhang, Fan Zhang, Qiyang Yu, Yuezhe Wang, Yongming Rao, Jingjing Liu, Tiejun Huang, and Xinlong Wang. Generative multimodal models are in-context learners. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14398–14409, 2024. 3, 5, 6, 7, 2
- [43] Gemma Canet Tarrés, Zhe Lin, Zhifei Zhang, Jianming Zhang, Yizhi Song, Dan Ruta, Andrew Gilbert, John Colloso, and Soo Ye Kim. Thinking outside the bbox: Unconstrained generative object compositing. *arXiv preprint arXiv:2409.04559*, 2024. 2, 5, 6, 7
- [44] Weihang Wang, Qingsong Lv, Wenmeng Yu, Wenyi Hong, Ji Qi, Yan Wang, Junhui Ji, Zhuoyi Yang, Lei Zhao, Xixuan Song, et al. Cogvlm: Visual expert for pretrained language models. *arXiv preprint arXiv:2311.03079*, 2023. 3
- [45] 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*, pages 15943–15953, 2023. 2
- [46] Guangxuan Xiao, Tianwei Yin, William T Freeman, Frédo Durand, and Song Han. Fastcomposer: Tuning-free multi-subject image generation with localized attention. *International Journal of Computer Vision*, pages 1–20, 2024. 2, 4
- [47] Zhexiong Xiong, Wei Xiong, Jing Shi, He Zhang, Yizhi Song, and Nathan Jacobs. Groundingbooth: Grounding text-to-image customization. *arXiv preprint arXiv:2409.08520*, 2024. 3
- [48] Binxin Yang, Shuyang Gu, Bo Zhang, Ting Zhang, Xuejin Chen, Xiaoyan Sun, Dong Chen, and 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*, pages 18381–18391, 2023. 2, 5, 6, 7
- [49] Ziyang Yuan, Mingdeng Cao, Xintao Wang, Zhongang Qi, Chun Yuan, and Ying Shan. Customnet: Zero-shot object customization with variable-viewpoints in text-to-image diffusion models. *arXiv preprint arXiv:2310.19784*, 2023. 2

- [50] Bo Zhang, Yuxuan Duan, Jun Lan, Yan Hong, Huijia Zhu, Weiqiang Wang, and Li Niu. Controlcom: Controllable image composition using diffusion model. *arXiv preprint arXiv:2308.10040*, 2023. [2](#), [5](#), [6](#), [7](#)
- [51] Peiang Zhao, Han Li, Ruiyang Jin, and S Kevin Zhou. Loco: Locally constrained training-free layout-to-image synthesis. *arXiv preprint arXiv:2311.12342*, 2023. [4](#)

# Multitwine: Multi-Object Compositing with Text and Layout Control

## Supplementary Material

### 1. Training and Testing Data

In this section, we provide additional information about the training data generation pipeline proposed in Main Paper Section 3.4 and MultiComp, the multi-object compositing test set introduced in Main Paper Section 4.

#### 1.1. Training Data

As detailed in Main Paper Section 3.4, our paired training data — comprising ground truth images with multiple objects, descriptive captions with grounding information, segmented images of two objects, and corresponding object-specific bounding boxes — is collected from complementary sources: Video Data, In-the-Wild Images, and Manually Collected Data. Below, we expand on how paired data is extracted from each source.

**Video Data** Fig 1 illustrates the process for extracting paired training data from videos in [36, 37]. A ground-truth frame containing an annotated relationship between two objects is randomly selected. Two additional frames, each showing one of the interacting objects, are extracted from the same video, ensuring a similarity between the views of each object (DINO score  $MSE \geq 0.8$ ) [27]. A caption describing the relationship is automatically generated by feeding the ground-truth image and annotated relation into LLaVA v1.6 (34B) [24] with the prompt: “Can you provide a grammatically correct one-line caption for the relation <object A> <relation> <object B> in the image?”. The segmented objects are then extracted using an off-the-shelf semantic segmentation model [31].

**Image Data** We propose two automated approaches for obtaining paired data from in-the-wild images.

**Top-down approach** As illustrated in Fig 2, this approach generates paired training data from a single image using a systematic process. First, a commercial subject selection tool identifies the main objects in the scene. A semantic segmentation model [31] then segments the selected region to determine the number of objects present. If multiple objects are detected, two are randomly selected as composited objects. Their outlines are highlighted on the image using distinct colors (e.g., orange and blue) and passed to ViP-LLaVA (13B) [3] for caption generation through two sequential prompts. In the first step, entities within each highlighted outline are identified with a question like: “Please follow the sentence pattern of the example to list the entities within each rectangle. Example: ‘orange: banana; blue: apple’”. This produces responses such as “orange: teddy bear; blue: girl”. Using these entity labels, the second step generates a descriptive caption with a prompt

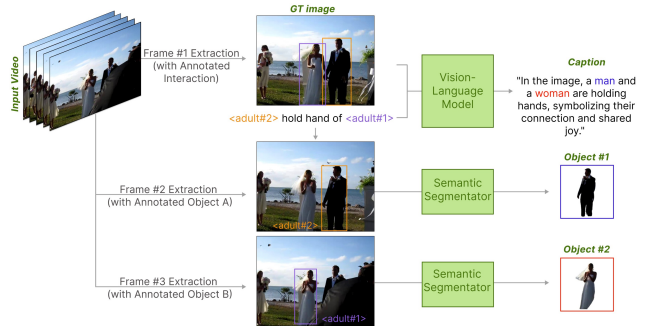


Figure 1. Training Data Generation from Video Data. Paired training data is obtained from video object relation datasets [36, 37] by extracting three frames with corresponding annotations and leveraging Vision-Language Models [24].




Figure 2. Training Data Generation from Image Data via Top-Down Approach. Paired training data is derived from in-the-wild images by leveraging a Vision-Language Model [3] and a Semantic Segmentator [31].

like: “Can you provide a one-line caption including the interaction between ‘teddy bear within the orange rectangle’ and ‘girl within the blue rectangle’ in the image by using these exact entity names?”. The resulting caption, e.g., “A young girl within the blue rectangle is holding a large teddy bear within the orange rectangle”, is refined by removing the grounding phrases “within the orange/blue rectangle” after using them for correlating object images with text tokens. This method ensures accurate grounding information, even when both entities are labeled the same, resulting in the final caption: “A young girl is holding a large teddy bear”.

**Bottom-up approach** This approach (Fig 3) leverages a grounding model like GroundingDINO [25] to process paired ground truth images and captions. The model extracts bounding boxes that link specific words in the caption to objects in the image. Duplicates, background elements, and undesired objects (e.g., overly large or small objects, or those with low confidence scores) are removed, leaving a set of object candidates with corresponding grounding information. Two of those objects are then randomly selected, and an off-the-shelf semantic segmentation model [31] is




Figure 3. Training Data Generation from Image Data via Bottom-Up Approach. Paired training data is extracted from in-the-wild images with a paired caption by leveraging a Grounding Model [25] and a Semantic Segmentator [31].




Figure 4. Training Data Generation from Manually Collected Data. Paired training data is obtained from a collected dataset containing object images, human images, and images of humans interacting with objects. We leverage a Vision-Language Model [3] and a Semantic Segmentator [31] to extract segmented objects and corresponding caption with grounding information.

used to extract them. This results in two segmented objects along with their associated grounding details from the original caption.

**Manually Collected Data** We manually collect and annotate images featuring objects, humans, and human-object interactions. For captioning with grounding information, we use ViP-LLaVA (13B) [3], following the exact same procedure as the top-down approach described above. Additionally, a semantic segmentation model [31] is employed to segment entities in images containing either a single object or a human. Fig 4 illustrates this approach.

## 2. Inference Data

Our collected MultiComp set consists of 119 paired data entries, each containing: (i) a background image, (ii) two object images, (iii) object-specific bounding boxes, (iv) an inpainting bounding box encompassing the previous ones, and (v) a descriptive caption with grounding informa-

tion. Background images are sourced from Pixabay [29], while objects are from Pixabay [29], MultiBench [23], and DreamBooth [34]. Bounding boxes and captions are manually crafted. For evaluation, we perform 5 iterations of each model on the entire set, resulting in 595 generated images.

Out of these, 395 images contain overlapping bounding boxes for the two objects (*MultiComp-overlap*), while 200 show non-overlapping bounding boxes (*MultiComp-nonoverlap*). We evaluate these subgroups separately due to their differing levels of difficulty. While simultaneous compositing offers benefits like cohesive harmonization in both cases, it is especially effective in the overlapping cases, where it allows for simultaneous object reposing and the generation of additional elements needed for the scene.

When textual input is provided, we further categorize the set into two subgroups based on caption types: *MultiComp-action* and *MultiComp-positional*. The *MultiComp-action* subset contains 375 images generated from action-based captions (e.g., *running after*, *playing with*, *holding*), while the *MultiComp-positional* subset includes 220 images generated from captions describing positional relations (e.g., *next to*, *behind*, *in front*). This distinction allows us to separately evaluate cases where reposing objects is often necessary (*MultiComp-action*) versus those primarily describing object layout (*MultiComp-positional*).

## 3. Comparison to Existing Methods

We provide additional visualizations comparing our model to existing generative object compositing and multi-entity subject-driven generation models in Sections 3.1 and 3.2. Further details on user studies can be found in Section 3.3.

### 3.1. Comparison to Generative Object Compositing Methods

We visually compare our simultaneous multi-object compositing method to sequentially adding two objects using State-of-the-Art Generative Compositing Methods [7, 40, 43, 48, 50] in Fig 5.

### 3.2. Comparison to Subject-Driven Generation Methods

We visually compare two-entity subject-driven generation using our method and existing methods with available code (BLIP-Diffusion [22], KOSMOS-G [28] and Emu2Gen [42]) in Fig 6.

### 3.3. User Studies

We conduct six user studies to evaluate our multi-object compositing model against other generative object compositing models [7, 40, 43, 48, 50] and Emu2Gen [42]. In each study, non-expert users are shown two side-by-side images, one generated by our model and the other by a baseline, presented in random order. Users are asked to choose

the preferred image based on a specific criterion. The entire MultiComp set is used for each experiment, except for the ‘most realistic interaction’ evaluation, where only images from MultiComp-*overlap* are considered, as this subset best evaluates the task. At least five users rate each image pair, and the results are aggregated via majority consensus. Visual examples of each experiment and the specific questions posed to the users can be found in Figs 7, 8, 9, 10, 11, and 12.

## 4. Ablation Study

Fig 13 shows visual examples of images generated by each ablation of our model (as detailed in Main Paper Table 3) for the same set of inputs. Without multi-view data (*i.e.*, video data, manually collected data), the model struggles to properly reposition and combine objects to align with the textual description. In the absence of joint training for compositing and customization, the model fails to balance textual and visual inputs, resulting in object identity loss when repositioning. Without cross-attention and/or self-attention losses, disentangling object identities becomes difficult, leading to texture and color leakage between objects (*e.g.* bow color, cat ear on ball). Lastly, omitting the masking step during inference leads to a degradation in alignment with the input layout and fails to fully mitigate identity leakage.

## 5. Applications

### 5.1. Model Versatility

We demonstrated in Main Paper Fig 9 how, by leveraging the advantages of our joint compositing and customization training, our model can be used for subject-driven inpainting. Additionally, Fig 14 illustrates how the same model can be applied to a broad range of tasks:

**Layout-Driven Inpainting** This task takes as input a descriptive caption, a background image, and a layout specifying an inpainting region along with object-specific bounding boxes for objects referenced in the caption. The model inpaints the selected region of the background image, ensuring alignment with the textual description while positioning objects according to the provided layout.

**Multi-Object Compositing** In addition to the inputs required for layout-driven inpainting, this task includes an image for each object corresponding to the provided bounding boxes. The model maintains the identity of these objects while enabling repositioning and view synthesis, producing a cohesive composited image.

**Layout-Driven Generation** In this case, no background image is provided. This task uses only a descriptive caption and bounding boxes specifying object positions as inputs. The model generates a full image that aligns with the caption while placing objects in the specified locations.

**Multi-Entity Subject-Driven Generation** Similar to layout-driven generation, this task uses a text caption and bounding boxes as inputs but also includes an image for each object. The model generates a complete scene that aligns with the text, places objects in their specified locations, and preserves their unique identities.

### 5.2. Multi-Object Compositing and Multi-Entity Subject-Driven Generation

Fig 15 show how the same model can be used for both multi-object compositing and multi-entity subject-driven generation, guided by a variable number of provided objects.

### 5.3. Robust Object Compositing

During training, an off-the-shelf segmentation model (EntitySeg [31]) is used to extract single objects from images, a process that introduces segmentation errors and occlusions. These imperfections during training strengthen our model’s ability to handle real-world scenarios at inference. As a result, Multitwine demonstrates robust compositing capabilities, seamlessly adapting object attributes to produce natural-looking results even in challenging scenarios. Fig 16 illustrates several examples of these capabilities, including: managing the transparency of a glass of milk, handling incomplete objects and inaccurate segmentations, and reharmonizing subjects from different domains (*i.e.* blending two cats extracted from a color image and a black-and-white image).

### 5.4. Robust Layout Alignment

Multitwine incorporates a layout cue as input, allowing users to specify the precise location and shape of each composited object. The model is trained to adapt object dimensions to fit the provided input masks, offering users greater control over the final output, as demonstrated in Fig 17. However, this strict bounding box guidance can sometimes lead to slight deformations in the objects, as commonly seen with generative compositing models. If deformations are undesirable, training with perturbed masks could mitigate them, though at the cost of some layout control.

### 5.5. Attribute Editing through Text

As seen in Fig 18, our model is designed to accommodate complex, multi-word grounded descriptions for each composited object, effectively capturing and reflecting their nuances in the final image. This capability ensures greater fidelity and detail in the generated results, aligning closely with the provided descriptions. In this case, multimodal embeddings are created by concatenating visual information after the last corresponding text token, while cross-attention masking and inference operate on all grounded text tokens.

## 6. Limitations and Failed Cases

Although simultaneously compositing several objects is possible with Multitwine, our model is not specifically designed for handling an unlimited number of objects. Fig 19 depicts a case where our model fails to composite six objects in a complex scene. When attempting to composite many objects, the model struggles to naturally integrate them all, resulting in missing objects, weaker harmonization, and reduced text-image alignment.

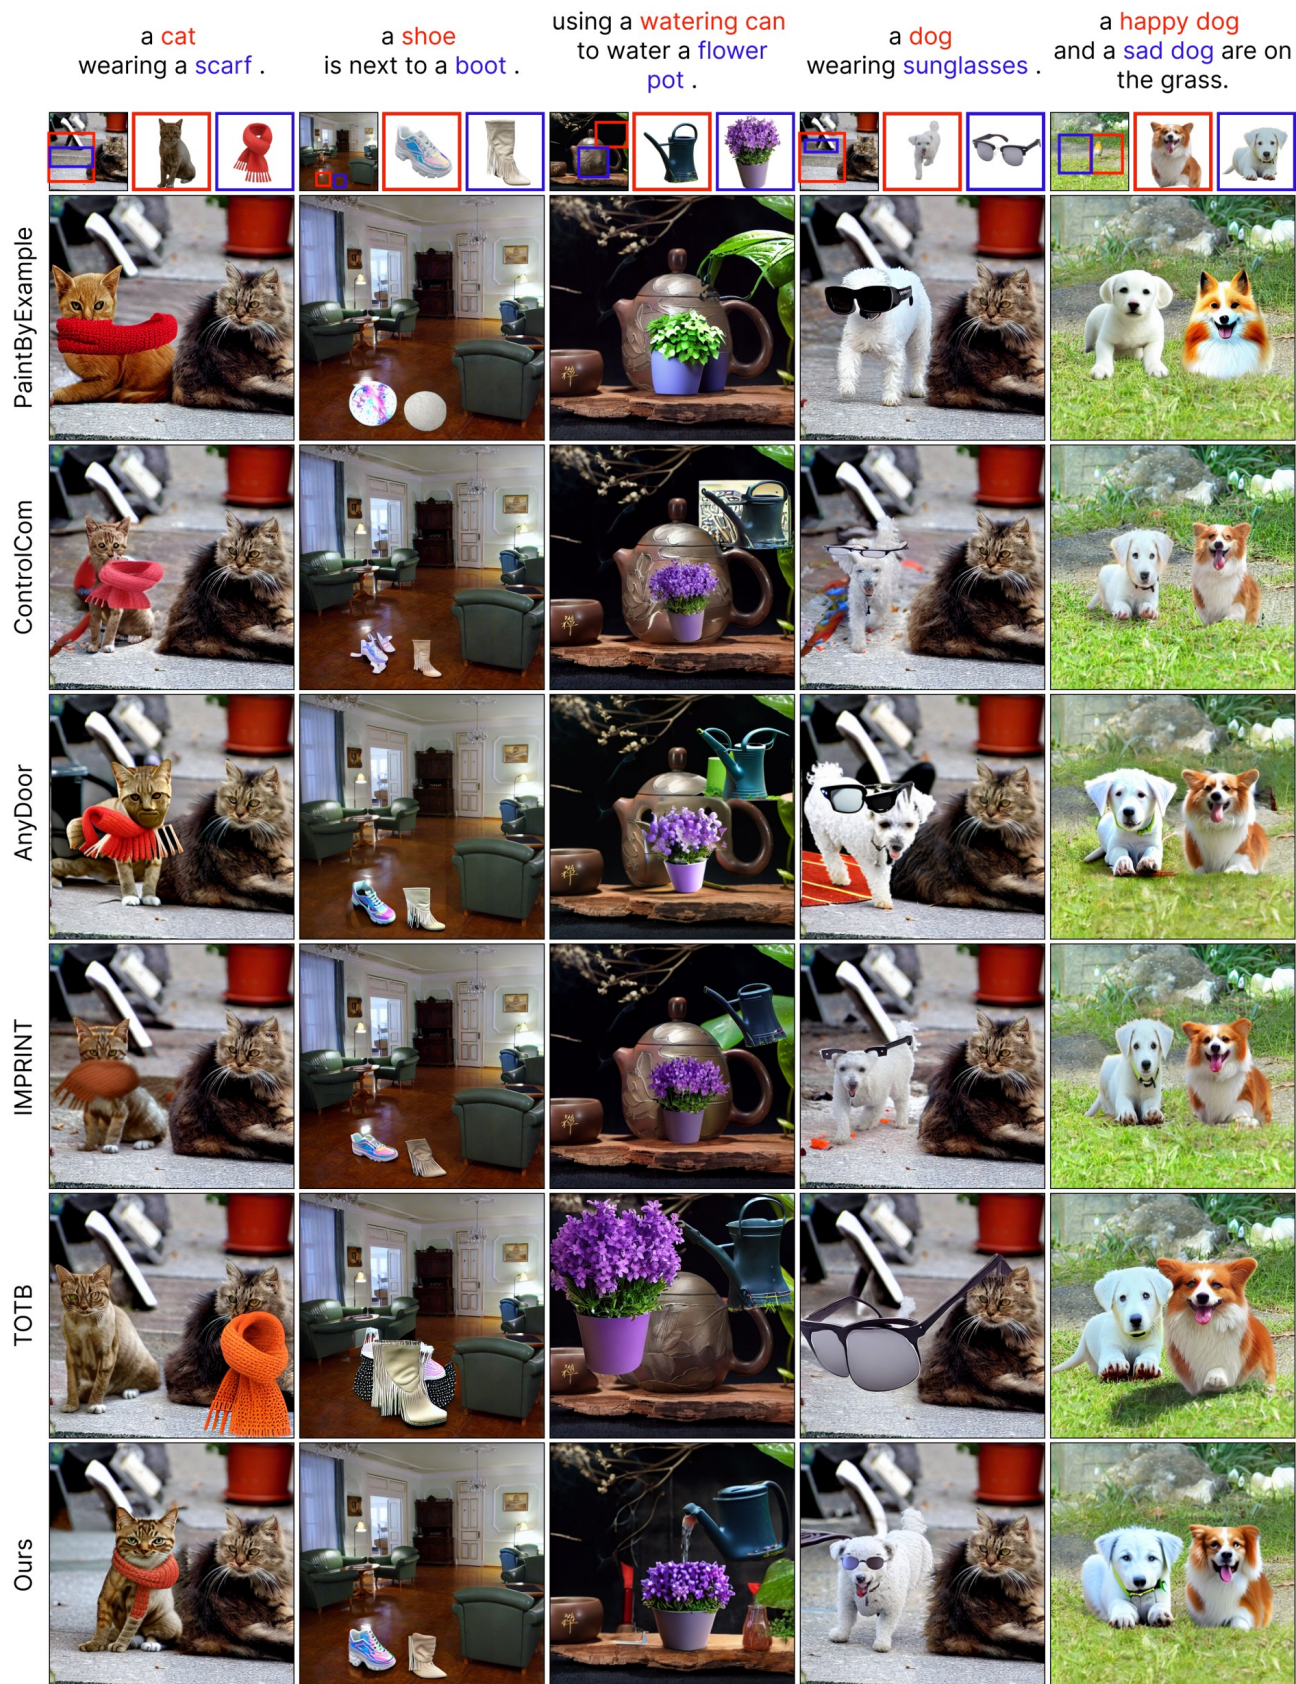


Figure 5. Visual comparison of our Multi-Object Compositing Method and State-of-the-Art Generative Object Compositing Methods [7, 40, 43, 48, 50].





Figure 6. Visual comparison of our Customization Method and State-of-the-Art Subject-Driven Generation Methods [22, 28, 42].

Two objects have been added to the same background in different ways (options A and B).  
Study options A and B and pick which composition has the **best quality**.




Figure 7. User Study on ‘Compositing Quality’. Screenshot of user study presented to participants for evaluating the image quality of our multi-object compositing method against generative object compositing baselines [7, 40, 43, 48, 50].

Two images have been generated in different ways (options A and B).

Study options A and B and pick which image is **more consistent with the given caption**.

Caption: "a dog wearing a bow tie."



Figure 10. User Study on ‘Text Alignment’. Screenshot of user study presented to participants for evaluating the text alignment of our multi-object compositing method against Emu2Gen [42].

Two objects have been added to the same background in different ways (options A and B).

Study options A and B and pick which composition provides the **most realistic interaction** between the two composited objects




Figure 8. User Study on ‘Realistic Interaction’. Screenshot of user study presented to participants for evaluating the realism of interactions generated by our multi-object compositing method against generative object compositing baselines [7, 40, 43, 48, 50].

Two objects have been used as input for generating images in different ways (options A and B).

Study options A and B and pick which image is **more consistent with the given object images**

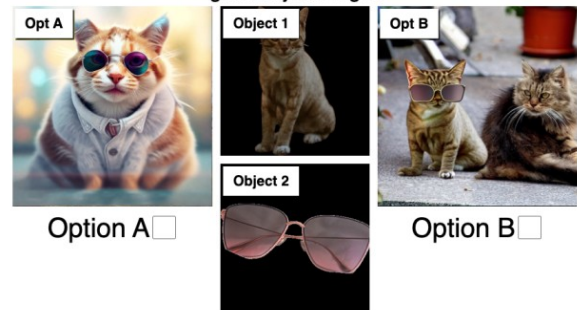


Figure 11. User Study on ‘Objects Alignment’. Screenshot of user study presented to participants for evaluating the alignment with input object images of our multi-object compositing method against Emu2Gen [42].

Two images have been generated in different ways (options A and B).

Study options A and B and pick which image is **more consistent with the given background image**



Figure 9. User Study on ‘Background Alignment’. Screenshot of user study presented to participants for evaluating the alignment with background image of our multi-object compositing method against Emu2Gen [42].

Two images have been generated in different ways (options A and B), adding two objects in the same layout.

Study options A and B and pick which image is **more consistent with the given layout**. Note that other objects may be present in the image outside of the given layout. The presence of such objects should not affect the final decision.




Figure 12. User Study on ‘Layout Alignment’. Screenshot of user study presented to participants for evaluating the layout alignment of our multi-object compositing method against Emu2Gen [42].



Figure 13. Visual examples for each ablation of the model. From left to right: (i) inputs (background, layout, objects and text), (ii) no self-attention loss, (iii) no self-attention or cross-attention loss, (iv) no joint training for compositing and customization, (v) no multi-view data (*i.e.* video data, manually collected data), (vi) no inference masking step, (vii) final model.



Figure 14. Visual examples for different applications of our model. Our model can operate on different modes such as: (i) layout-driven inpainting, (ii) multi-object compositing, (iii) layout-driven generation, (iv) multi-entity subject-driven generation.



A **backpack** in front of a blue door.



A **dog** wearing a **santa hat**.



A **woman** is dancing with a **man**.



A **man** is playing **guitar**.



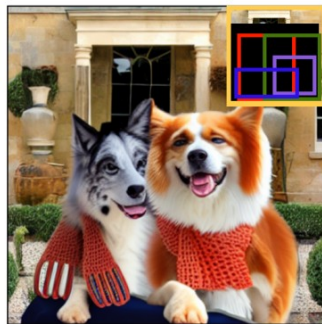
A **dog** is sharing a **cake** with a **cat**.



A **dog** is eating an **ice cream** with a **dog**.



A **dog** is wearing a **scarf** and hugging a **dog** wearing a **scarf**.



A **woman** is eating **pizza** with a **woman** eating an **ice cream**.



Figure 15. Visual Examples for Multi-Object Compositing (*left*) and Multi-Entity Subject-Driven Generation (*right*), using a variable number of grounding objects. *First Row*: One Object; *Second Row*: Two Objects; *Third Row*: Three Objects; *Forth Row*: Four Objects.




Figure 16. Visual Examples for Multi-Object Compositing in challenging scenarios, including: (left) harmonization, (middle) incomplete objects, and (right) imperfect segmentation and compositing of objects with different style (i.e. color, black and white).




Figure 17. Visual Examples for Multi-Object Compositing via different layout inputs.



Figure 18. Visual Examples for Multi-Object Compositing using different grounded captions. The text caption can be used to edit poses and attributes of composited objects.

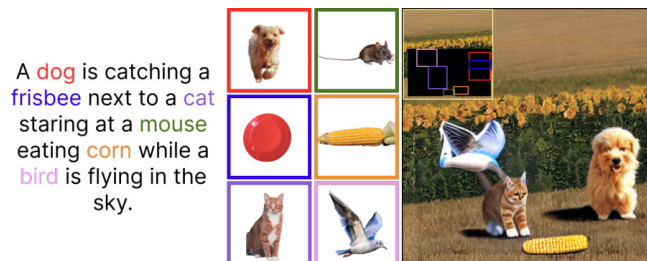


Figure 19. Visual example of a failed case from our model, depicting its limitation for compositing a high number of objects in a complex scenario.