

# **DreamingComics**: A Story Visualization Pipeline via Subject and Layout Customized Generation using Video Models

Patrick Kwon, Chen Chen

Center for Research in Computer Vision, University of Central Florida, USA

yo564250@ucf.edu, chen.chen@crcv.ucf.edu

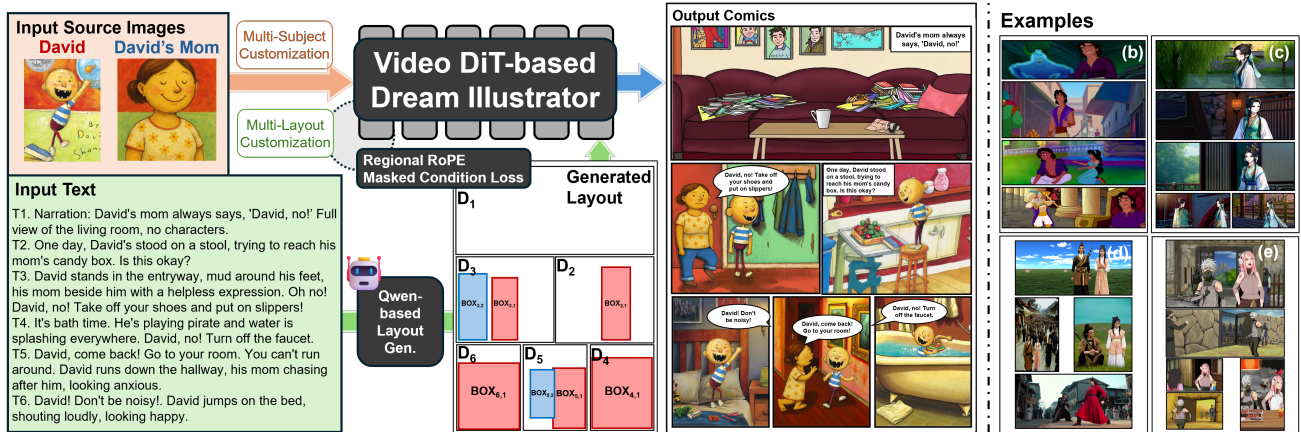


Figure 1. (left) Overview of **DreamingComics**, a story visualization framework for multi subject and layout control. The dialogues are post-edited by humans. (right) Examples from **DreamingComics**, which is capable of generating layout-controlled stories with diverse art styles such as pencil illustrations, Disney-style animation, digital line-art, live-action drama, and cel animation.

## Abstract

Current story visualization methods tend to position subjects solely by text and face challenges in maintaining artistic consistency. To address these limitations, we introduce **DreamingComics**, a layout-aware story visualization framework. We build upon a pretrained video diffusion-transformer (DiT) model, leveraging its spatiotemporal priors to enhance identity and style consistency. For layout-based position control, we propose *RegionalRoPE*, a region-aware positional encoding scheme that re-indexes embeddings based on the target layout. Additionally, we introduce a masked condition loss to further constrain each subject’s visual features to their designated region. To infer layouts from natural language scripts, we integrate an LLM-based layout generator trained to produce comic-style layouts, enabling flexible and controllable layout conditioning. We present a comprehensive evaluation of our approach, showing a 29.2% increase in character consistency and 36.2% increase in style similarity compared to previous methods, while displaying high spatial accuracy. Our project page is available at <https://yj7082126.github.io/dreamingcomics/>

## 1. Introduction

Story visualization [11, 33, 65, 79], the task of generating coherent visual sequences from textual narratives and character identities, has gained increasing attention with the advancement of generation models [2, 14, 40, 43] and controllable image customization methods [34, 35, 63, 69, 74]. However, despite rapid improvements, existing approaches still lack the visual control necessary for storytelling tasks.

Firstly, **controlling the position of multiple characters** is crucial for scene design, yet text prompts alone lack the pixel-level precision needed to specify spatial layout and character identity. Although recent image customization methods for diffusion transformer (DiT) models [2, 14, 38] can control who appears in the image [35, 61, 63] or where objects are placed [6, 72], they struggle to support both simultaneously, resulting in images with overlapping characters or incorrect appearances. Moreover, **preserving consistent artistic styles**, especially in cartoon and flat-shaded illustrations, remains difficult, as image generative models are typically biased toward photorealistic rendering, often overriding stylistic intent. In addition, the lack of paired datasets with both subject identity and position signals

Method	Layout Control	Multi-Subject Interaction	DiT architecture	LLM-based Layout
Story-Adapter [33]	✗	✓	✗	✗
StoryDiffusion [79]	✗	✗	✗	✓
TheaterGen [11]	✓	✗	✗	✓
DiffSensei [62]	✓	✓	✗	✗
UNO [63]	✗	✓	✓	✗
DreamO [35]	✗	✓	✓	✗
Eligen [72]	✓	✗	✓	✗
RealGeneral [28]	✗	✗	✗	✗
DRA-Ctrl [3]	✗	✗	✓	✗
<b>DreamingComics (Ours)</b>	✓	✓	✓	✓

Table 1. Comparison of DreamingComics with prior works across various capabilities.

hinders the development of layout-aware customization. These challenges are most apparent in the comics domain, where consecutive image panels must preserve character identity and style while conforming to diverse layout structures unique to comics.

In this paper, we present **DreamingComics**, a story visualization framework that simultaneously supports **multi-subject identity/style preservation and layout control**. We decompose the task into two modules: an LLM-based layout generator and a layout-aware customization model, *Dream-Illustrator*. The layout generator produces comic-style layouts consisting of panels and character bounding boxes from a given storyline, while Dream-Illustrator renders consistent characters within those regions. This enables image customization to accurately position characters instead of solely relying on text, and enhance usability by reducing the burden on users of designing intricate layouts.

We build Dream-Illustrator on a pretrained video DiT backbone [23, 73], which, unlike image-specific generators, offers weaker perceptual quality but strong spatiotemporal priors that enhance visual consistency. Through FramePack [73], we repurpose a video model for image customization, enabling robust style generalization while keeping an image-level computation bottleneck. To control subject position by layout, we introduce **RegionalRoPE**, a regional 3D Rotary Position Embedding (RoPE) scheme that aligns each subject with its spatial region. To further improve layout fidelity, we employ a **masked condition loss** that penalizes attention beyond the assigned regions. Together, these mechanisms allow our model to generate identity-preserving images aligned with layouts. Lastly, we **construct an image-layout-paired dataset** for training Dream-Illustrator by annotating high-quality video samples. In parallel, we reformulate comic datasets into text-layout pairs to train our layout generator.

We evaluate our framework on benchmarks such as ViStoryBench [80] and DreamBench [39]. Extensive experiments demonstrate that DreamingComics can generate multi-subject customized images that closely align with target layouts while generalizing across a broad spectrum of artistic styles (Fig. 1 (b)-(e)).

In short, our main contributions are as follows.

- DreamingComics, a layout-aware story visualization framework that supports multi-subject identity and style customization with spatial layout.
- An image customization model built on a pretrained video DiT model, leveraging spatiotemporal priors to improve style and identity consistency.
- Layout-aware conditioning strategies: RegionalRoPE, masked condition loss, and identity-layout-paired dataset generation that jointly enable accurate spatial control and reduce identity leakage.
- An LLM-guided layout generation module, fine-tuned on structured comic layout data, for prompt-based layout planning with minimal user input.

## 2. Related Work

**Image Customization.** Early methods for subject customization were based on finetuning the model for each specific subject [16, 18, 24, 46], achieving strong consistency but requiring training for each instance. Other tuning-free methods [25, 60, 69, 75] focused on injecting encoded results from an image encoder into a U-Net [44] diffusion model, balancing efficiency with consistency. With the advent of new Diffusion Transformer (DiT) models [2, 14, 38], many image customization works [19, 35, 52, 61, 63, 64, 76] directly reason over reference and noise latents in a unified attention space, improving subject fidelity without per-instance tuning. However, these approaches tend to rely on implicitly learned spatial cues, limiting spatial positioning.

There have been several studies on spatial position control for both training-based [26, 56, 58, 59, 78] and training-free methods [1, 11, 32] based on a pretrained U-Net diffusion model. Recently, methods such as Regional Prompting [6] and Eligen [72] have proposed regional masking over the full-attention structure to control DiT models. However, explicit multi-subject spatial layout control remains underexplored, particularly within the DiT framework. Although TheaterGen [11] supports multi-subject layout control, it generates each subject individually and merges them with ControlNet [74], omitting any interactions between subjects.

Our work is in line with previous studies that employed video generative models for image-level tasks. For example, Frame2Frame [45] uses an image-to-video diffusion model [68] for image editing by generating a sequence of frames, while FramePainter [77] and Object-Mover [70] finetune a video model to perform object-level image editing tasks. Several works [3, 9, 28] have also used video models for controllable image generation, demonstrating their capabilities to unify diverse image tasks. However, neither of them targeted multi-subject-driven generation aligned with a spatial layout. Our work uses novel strategies to incorporate layout

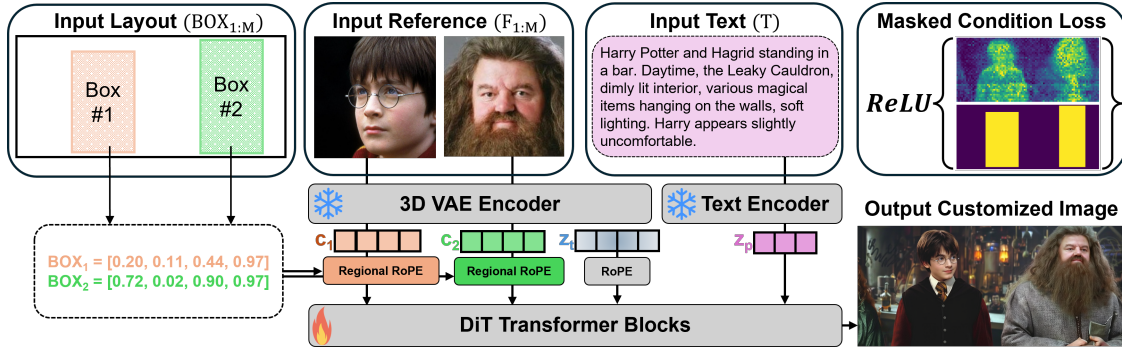


Figure 2. Overview of our image customization pipeline. The input reference images are encoded as token sequences  $c_{1:n}$  along with the noise latent  $z_t$  and the text latent  $z_p$ , passed to the stream of diffusion transformer blocks. We calculate a custom regional RoPE from the layout condition and apply it to the encoded references. During training, we calculate a Masked Condition Loss between the cross-attention map and the given layout condition, encouraging the model to position references within the layout.

control into a video model, utilizing multiple subject reference images with spatial layout.

**Story Visualization.** Recent advances in generative models and LLMs have resulted in significant progress in story visualization. Previous works on image-based story visualization [10, 30, 37, 57, 67, 79] focused on improving cross-frame alignment for images generated by a U-Net-based image generator. Comic generation [8, 62] is also a notable subtask within story visualization. Although previous works such as DiffSensei [62] were restricted to monochrome manga aesthetics, our approach supports a wider range of artistic styles, from stylized cartoon art to color-heavy illustrations, allowing diverse and culturally agnostic storytelling.

Several recent works use large language models to reason about scene layouts from text [11, 27]. For instance, TheaterGen [11] and AutoStudio [10] employ training-free LLM agents to predict image-level layouts. In contrast, our module is fine-tuned on curated comic data, enabling the generation of both panel-wise and page-level layouts—a granularity not addressed by prior work. While DiffSensei [62] also incorporates an MLLM for its inference, *its usage is limited to an identity adapter and does not support layout generation.*

### 3. Methodology

In Sec. 3.1, we formulate our objectives and introduce our base model, HunyuanVideo-I2V [23]. In Sec. 3.2, we introduce our LLM-based layout generator that outputs multi-panel layouts from text input. In Sec. 3.3, we present Dream-Illustrator, an **image customization model based on a DiT video generator**, adapted for layout-aware next-frame prediction through a **regional RoPE scheme and masked condition loss**. In Sec. 3.4, we report the dataset generation pipeline for both comic layout and image customization dataset.

#### 3.1. Problem Statement and Preliminaries

Our objective is to visualize stories through a comic-style narrative. Given a script  $T = \{T_1, T_2, \dots, T_n\}$ , the traditional goal of story visualization is to generate a sequence of corresponding images  $\{p_1, p_2, \dots, p_n\}$  of the same size. Comics, however, are composed of *image panels* that have different resolutions and locations within a single page [54]. Although several works use LLMs to predict layouts for a single image [11, 27], they do not predict layouts for multiple panels, making them insufficient for comic generation. In contrast, we adopt a comic-centric representation, where panel  $i$ , consisting of  $n$  characters, is represented as a tuple

$$(T_i, D_i, \{\text{BOX}_{i,1}, \dots, \text{BOX}_{i,n}\}), \quad (1)$$

where  $T_i$  is the input text,  $D_i \in \mathbb{R}^4$  is the panel’s bounding box, and  $\text{BOX}_{i,j} \in \mathbb{R}^4$  are subject-level bounding boxes for each character  $j$ . This structure enables us to reason about both the positioning of panels within a page and the positioning of characters within panels, reflecting the unique narrative of comics.

To utilize its rich spatiotemporal priors for visual consistency, we base our image customization method on HunyuanVideo-I2V [23], a video generation model composed of a causal 3D Variational Autoencoder [22] (3DVAE,  $\mathcal{E}$ ), a multimodal large language model (MLLM) encoder, and a diffusion transformer with a unified full attention mechanism across spatial and temporal tokens. Given an input video  $x \in \mathbb{R}^{(4T+1) \times 3 \times 16H \times 16W}$ , the VAE encodes it into a latent representation  $y \in \mathbb{R}^{(T+1) \times 16 \times 2H \times 2W}$ , which is then patchified into a sequence of visual tokens. The textual prompts  $T_p$  are then encoded as tokens using the MLLM encoder and concatenated with the visual tokens to form a single sequence  $z$ , and passed into a transformer with 3D Rotary Position Embeddings (RoPE) [51] to model positional relationships. For I2V generation, HunyuanVideo-I2V replaces the tokens of the first frame with the conditioning image tokens. The model is




Figure 3. Given a list of textual descriptions for each panel, the finetuned LLM outputs a spatial layout for each panel and characters as a set of bounding boxes. Note that our layout, compared to the layout generated from the same prompt using GPT-4 [36], occupies most of the panel region, correctly orders the panel (top-to-bottom, right-to-left), and draws plausible character boxes, constituting a “good comic layout”.

then optimized using a flow matching loss [29]

$$\mathcal{L} = \|v_{\theta}(\mathbf{y}_t, t, C_I, T_P) - (\epsilon - \mathbf{y})\|^2, \quad (2)$$


where  $\epsilon \sim \mathcal{N}(0, I)$ ,  $\mathbf{y}_t = (1 - t)\mathbf{y} + t\epsilon$ , and  $v_{\theta}$  represent the parametric neural network.

### 3.2. Layout Generation Pipeline

To generate page-level layouts from a given set of text descriptions, we fine-tune a large language model (LLM) [66] using supervised fine-tuning (SFT) based on a newly crafted comics layout dataset. Given an input script  $T$ , the model is trained to predict a structured layout of panel regions and associated character positions, which are parsed as  $(D_i, (\text{BOX}_{i,1}, \dots, \text{BOX}_{i,n}))$ . These representations are later used as layout conditions in the image customization stage, alongside their corresponding character images. A visualization of this layout generation process is shown in Fig. 3. By training an LLM with strong spatial reasoning over domain-specific data and making it understand the sparse visual cues within comics, we reduce the burden on users to manually specify intricate layouts while enabling the model to follow complex, prompt-driven visual instructions.


### 3.3. Dream-Illustrator

**Using FramePack as a next-frame predictor.** Since video generation models handle sequences of semantically consistent frames, they capture rich cross-frame context and complex spatiotemporal priors during training, which motivated us to use them for image customization; specifically, by using the reference image as the first frame to generate a subsequent target frame. Previous approaches in using video models for image tasks [3, 28, 45] required generating multiple frames for a single image output. At the same time, treating the problem as a simple two-frame video generation problem (generating the frame at  $t = 1$  conditioned on  $t = 0$ ) results in rigid “copy-pasting” artifacts, with limited variation and poor prompt adherence.



(a) Reference Image with original RoPE

(b) Generated Image



(c) Reference Image with regional RoPE

(d) Generated Image

Figure 4. For illustration, we visualize the original positional indices for RoPE at (a) and the new indices at (c). The blue square indicates the intended layout region, and the red square indicates the actual generated region. The original RoPE restricts the reference content to the top-left corner, while ours can correctly position it according to the given layout.

To address this, we build upon FramePack [73], a finetuned video DiT model that generates video frames in a progressive manner. Originally designed for efficient long-range video generation, FramePack compresses a sequence of  $T$  input frames ( $\mathbf{F} \in \mathbb{R}^{T \times h \times w \times c}$ ) to sample  $S$  output frames ( $\mathbf{X} \in \mathbb{R}^{S \times h \times w \times c}$ ). Crucially, FramePack is capable of generating frames that are temporally distant from the input—for example, by producing an output at  $t = 9$  from a reference at  $t = 0$ . We leverage this property to reformulate FramePack as an image customization module: given a target timestep  $t'$  and a set of  $N$  reference images ( $\mathbf{F} = F_1, F_2, \dots, F_n$ ) at  $t = 0$ , our goal is to generate a single target frame  $\mathbf{X}$  that preserves subject identities while aligned with the input prompt. In practice, we encode each reference image  $F_i$  into the latent space using VAE ( $c_i = \mathcal{E}(F_i)$ ), and concatenate the patched latents with the latent noise  $z_t$  and the latent text  $z_p$ .

Compared to prior approaches [3, 9, 28], our method leverages video model priors while generating a single frame, maintaining a relatively low computation bottleneck. As a result, we can infer a  $1280 \times 720$  image within 17 seconds, over  $3\times$  faster than DRA-Ctrl [3]. Detailed comparisons of inference time and memory are provided in the **Supplementary**.

**Regional RoPE for layout control.** The default RoPE assigns the same starting coordinates (0, 0) to all reference frames, which makes the model perceive all subjects as if they are originating from the same region, leading to spatial entanglement and identity collapse (Fig. 4). To enable precise layout-driven spatial control, we introduce RegionalRoPE, a deterministic map-

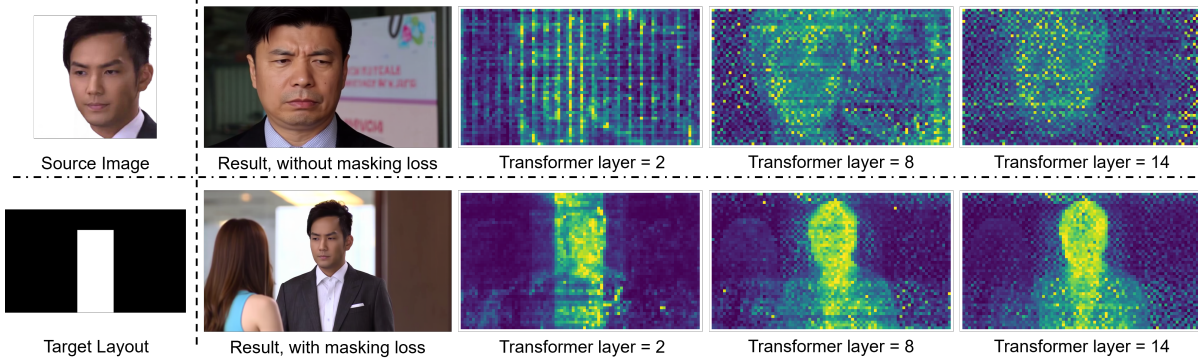


Figure 5. Comparison between the CAMs of our model trained without the masked condition loss (top) and with the masked condition loss (bottom). Using our new loss helps to position the attention around the target layout, which is evident in the third column (Transformer layer = 2), naturally inducing the model during training to position the character.

ping of each reference’s RoPE indices to its target layout box. Unlike previous works that modify RoPE for spatial decorrelation [52, 53, 63], we use RegionalRoPE for explicit spatial grounding.

As mentioned in Sec. 3.1, we have a region  $\text{BOX}_i = [w_{\text{start}}, h_{\text{start}}, w_{\text{end}}, h_{\text{end}}]$  for each reference latent  $c_i \in \mathbb{R}^{h_i \times w_i \times d}$ . For the region size  $(W_{\text{box}}, H_{\text{box}})$ , to preserve the latent aspect ratio of the reference while fitting within the region, we compute a scaling factor  $s = \min\left(\frac{W_{\text{box}}}{w_i}, \frac{H_{\text{box}}}{h_i}\right)$  and define the adjusted RoPE grid size as  $(W', H') = (s w_i, s h_i)$ . The RoPE grid is then positioned within the layout region with ranges

$$w'_{\text{start}} = w_{\text{start}} + \frac{W_{\text{box}} - W'}{2}, \quad w'_{\text{end}} = w'_{\text{start}} + W' \quad (3)$$

$$h'_{\text{start}} = h_{\text{start}} + a (H_{\text{box}} - H'), \quad h'_{\text{end}} = h'_{\text{start}} + H' \quad (4)$$

where  $a \in [0, 1]$  controls the vertical alignment ( $a=0$  for top-aligned,  $a=0.5$  for center-aligned). The indices for each latent pixel  $(i, j)$  are finally mapped as

$$(t', i', j') = \left(0, w'_{\text{start}} + \frac{W'}{w_i} i, h'_{\text{start}} + \frac{H'}{h_i} j\right) \quad (5)$$

Each reference latent is encoded independently with its own coordinates, and the resulting sequences are concatenated into the input stream along with the target tokens. Previous works like UNO [63] and Ominicontrol [52, 53] modify RoPE to decorrelate spatial cues and reduce reference copying, but they do not enforce explicit spatial alignment. In contrast, we assign RoPE offsets for each reference based on its region, enabling precise location mapping. Unlike other works [3, 28] which resize inputs to fixed frame sizes, we operate on native-resolution cropped latents, preserving subject fidelity and improving efficiency.

**Masked condition loss.** While RegionalRoPE aligns latents without extra training, its naive application can still lead to identity distortion and copy-paste artifacts. To address this, we introduce a masked condition loss

that supervises each subject’s spatial attention. During diffusion, we extract the cross-attention map (CAM) between the reference and the generated result as

$$\text{CAM}_{c_i, t, \text{block}_j} = \frac{Q_{c_i, t, \text{block}_j} K_{t, \text{block}_j}^T}{\sqrt{d}} \quad (6)$$

where  $Q_{c_i, t, \text{block}_i}$  are the tokens of the  $i$ -th reference image and  $K_{t, \text{block}_j}$  are the tokens of the noisy latent at timestep  $t$ , DiT layer  $j$ . Averaging these across the timestep and normalizing to  $[0, 1]$  yields  $\text{CAM}_{c_i, \text{block}_j}$ , visualizing each subject’s attention region.

To ensure that these spatial attentions respect their layout boundaries, we introduce a ReLU-based loss that operates on each CAM and the subject-wise binary spatial mask  $\text{MASK}_i$ . Here,  $\text{MASK}_i$  is a binary mask from the bounding box  $\text{BOX}_i$ , with in-layout pixels represented as 1 and others as 0. The masked condition loss, with respect to the DiT layer 2, is then defined as

$$\mathcal{L}_{\text{mask}} = \frac{1}{n_c} \sum_{i=1}^{n_c} \text{ReLU}(\text{CAM}_{c_i, \text{block}_2} - \text{MASK}_i) \quad (7)$$

where  $n_c$  refers to the number of conditions. Here, the ReLU penalizes only the attention leakage beyond the defined layout region while preventing unnecessary suppression of valid in-region focus. The final loss combines this term with the diffusion objective

$$\mathcal{L} = \mathcal{L}_{\text{diff}} + \lambda_{\text{mask}} \mathcal{L}_{\text{mask}} \quad (8)$$

This constraint guides the model to respect spatial boundaries and maintain subject-specific attention, mitigating identity bleeding and visual artifacts. The effects of our new training loss are illustrated in Fig. 5.

Additionally, motivated by previous works [3, 28] for multiple subjects, we apply attention masking such that each reference latent should not be visible to one another to avoid information leakage. More details on this feature are provided in the **Supplementary**.

### 3.4. Dataset Generation Pipeline

**Comics layout dataset.** To train our layout generator, we compile and annotate data from three comic datasets:

Method	Base Model	CIDS (Char) ↑		CSD (Style) ↑		OCCM ↑	Layout ↑	Inception ↑	Aesthetics ↑
		Cross	Self	Cross	Self	Score	Precision	Score	Score
Copy-Paste Baseline	-	94.1	98.7	72.8	71.5	99.7	-	6.72	4.48
Story-Adapter [33]	SDXL [40]	35.3	57.4	29.6	<u>63.0</u>	78.8	-	<b>9.79</b>	<u>5.55</u>
StoryDiffusion [79]	SDXL [40]	30.1	46.9	31.8	52.2	67.7	-	7.36	4.96
DiffSensei [62]	SDXL [40]	47.5	62.5	31.5	<b>65.2</b>	<u>85.9</u>	<u>42.0</u>	5.53	4.47
Eligen [72]	FLUX [2]	35.9	56.0	29.7	61.7	78.5	39.7	8.49	<b>5.81</b>
UNO [63]	FLUX [2]	46.2	61.1	<u>39.3</u>	59.3	83.8	-	8.42	5.20
DreamO [35]	FLUX [2]	<u>51.6</u>	<u>63.1</u>	38.3	58.8	85.7	-	7.86	5.51
Vlogger [81]	Vlogger [81]	33.9	54.2	28.3	47.2	76.6	-	7.46	4.25
RealGeneral [28]	CogVideoX [68]	37.5	54.2	38.6	52.7	78.1	-	7.82	5.08
DRA-Ctrl [3]	HunyuanVid [23]	36.2	50.2	39.0	58.0	74.9	-	<u>8.77</u>	5.03
Ours	FramePack [73]	<b>66.6</b>	<b>68.1</b>	<b>53.6</b>	58.0	<b>86.7</b>	<b>61.6</b>	6.81	4.54

Table 2. Quantitative evaluation results on ViStoryBench. The **bold** value is the best, and the underlined value is the second best.

COMICS [20], Manga109 [15], and PopManga [47], capturing varied styles and eras in comics. For datasets lacking annotations, we apply the MagiV2 panel detector [48] to extract panel and character bounding boxes. All data are standardized, with inconsistent samples filtered out. We generate panel and page-wise descriptions using Qwen2.5-VL [66].

**Paired subject dataset.** Public datasets rarely provide paired samples with references and layout conditions, which are critical for training Dream-Illustrator. To address this, we design a novel data framework that samples image pairs from curated video sources. Specifically, we build on OpenS2V-Nexus [71], a large-scale dataset with structured annotations. We select videos with at least one consistently framed human subject. Using the segmentation map for the first frame, we extract its human bounding boxes as target layout conditions. The source frame is then chosen from a distant timestamp, ensuring subject continuity through facial presence. Samples are retained only if both the target (TopIQ) and the source face (TopIQ-Face) [7] scores exceed the quality thresholds. To cover diverse artistic styles, we apply a similar procedure to Anime-Shooter [41], a reference-guided multi-shot animation dataset. Additionally, we processed the high-quality subset of Subject200K [52] following DreamO [35], using LISA to predict subject masks for the construction of paired data. All datasets are publicly released with their usage restricted to the research community, and we have fully adhered to their usage guidelines. In total, we compiled 55K single-subject and 20K multi-subject paired samples, with further details and ethical considerations provided in the **Supplementary**.

## 4. Experiments

### 4.1. Implementation Details

**Layout Generation Model Training.** We finetune Qwen2.5-VL (7B) [66], which has outperformed alternatives on comic layout understanding tasks [55], on 25K annotated comic layouts using supervised fine-

tuning. Layouts are represented as dictionaries of normalized bounding boxes over fixed panel counts. We apply LoRA with rank 8,  $\alpha = 16$ , dropout=0.05, and train using the AdamW optimizer (lr=5e-4).

**Dream-Illustrator Training.** We finetune HunyuanVideo-I2V [23] attached with the FramePack [73] LoRA weights. To support multi-subject conditioning and reduce copy-paste artifacts, we remove the original image projection module. Training is performed using LoRA [17] with rank 32, AdamW [31] optimizer (lr=2e-4), batch size 8, and mixed precision on 2xNVIDIA H100 GPUs. The  $\lambda_{\text{mask}}$  is set to 0.05. We train for 6K steps on single-subject samples, followed by 3K steps on multi-subject samples.

### 4.2. Image Customization Evaluation

We evaluate DreamingComics using ViStoryBench [80], the first comprehensive story visualization benchmark. The benchmark primarily measures Character Similarity (CIDS) [4, 12, 21, 42] and Style Similarity (CSD) [50] between reference and generated frames (cross) and between generated frames (self). The benchmark also defines Onstage Character Count Matching (OCCM), Inception Score (IS) [49], and Aesthetic Score [13]. To measure subject-layout fidelity, we add a layout precision metric for methods that support layout-based customization [62, 72]. We also evaluate our method on DreamBench++ [39] for single-subject driven generation against other image customization methods. Further details on the evaluation procedure and metrics for quantitative evaluations are provided in **Supplementary**.

As shown in Table 2, DreamingComics achieves substantial improvements in both character and style preservation. It achieves the highest CIDS-cross score of 66.6, surpassing the next best method (DreamO [35], 51.5) by 29.2%, and achieves a CSD-cross score of 53.6, highlighting our ability to **maintain both character identity and artistic style**. Our model also achieves an OCCM score of 86.7 and a layout precision score of 61.6, confirming correct layout-wise character generation. Although our Aesthetic Score and Inception

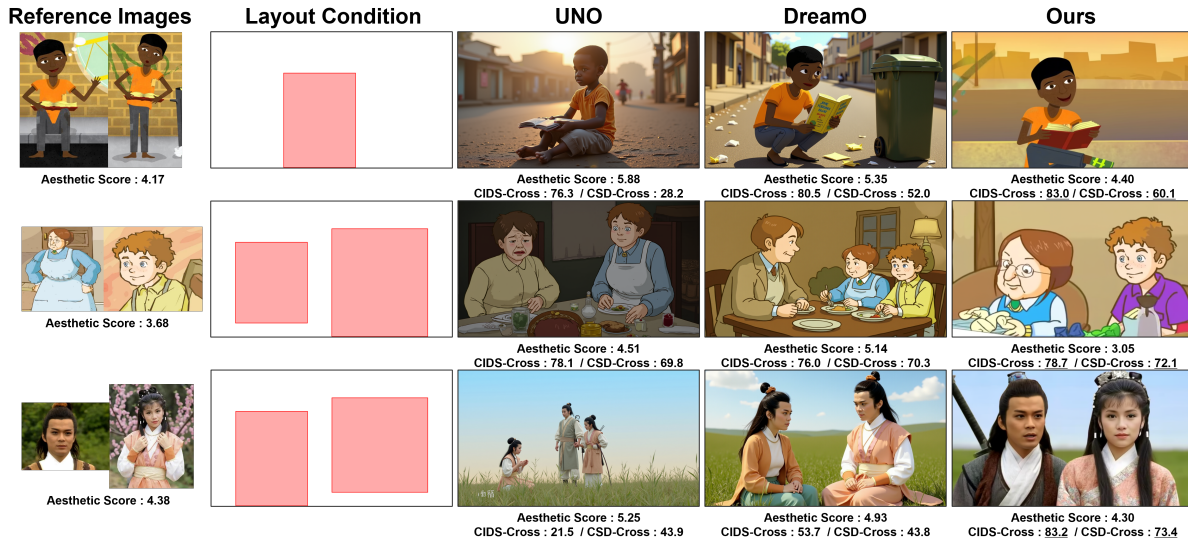


Figure 6. Image level qualitative comparison against other methods, along with their respective scores. Our method shows enhanced visual consistency with respect to the reference images, which is reflected in the similar aesthetic scores.

Method	DINO $\uparrow$	CLIP-I $\uparrow$	CLIP-T $\uparrow$
DreamBooth [46]	53.34	75.30	32.26
UNO [63]	56.24	77.96	33.23
DreamO [35]	58.59	79.35	32.96
<b>Ours</b>	60.50	79.50	32.59

Table 3. Quantitative results on DreamBench.

Score are comparable to the copy-paste baseline, this reflects our deliberate preservation of stylized, non-photorealistic appearances—an area where video models, despite lower perceptual sharpness, excel in maintaining stylistic consistency (e.g., Fig. 6). We also present the results for DreamBench++ in Table 3, which shows that our model achieves the highest DINO [5] and CLIP-I [42] scores, with CLIP-T [42] scores that are comparable to other methods.

Fig. 6 displays an image-level comparison with other methods [35, 63]. Unlike other methods, our method utilizes layouts to designate each character’s position, enhancing controllability while preventing superfluous or erroneous character creation (e.g., DreamO [35], second and third rows). Also, the results from other methods produce higher aesthetic scores, even when their output style largely deviates from that of the reference’s (e.g., first row). This is because aesthetic scoring tends to penalize images that do not look photorealistic rather than evaluating for diverse styles. However, *our method can faithfully adhere to the given reference style, which is indicated by the similarity in aesthetic scores.*

We also present a comic-level comparison in Fig. 7, where we first generate layouts from the input panel-wise captions and use them along with the reference images and captions to synthesize a full-length comic-style story. Layout-based methods like DiffSensei [62] and Eligen [72] fail to preserve the aesthetics of images,

Method	Panel Count	Coverage Ratio	Panel Ordering	Valid Characters	Character Count
Ours	100.0	79.14	99.00	100.0	86.30
GPT-4 [36]	87.67	29.25	66.67	57.67	76.47

Table 4. Evaluation results on our layout generator.

either because they do not allow image-level inputs or fail to preserve the input style. Image-based customization methods such as RealGeneral [28], UNO [63], and DreamO [35] cannot utilize layout conditions, which creates confusion in placing characters. For instance, DreamO places the gray-colored “Domestic Donkey” on the right side and the brown-colored “Wild Donkey” on the left, which is not intended in its caption. In contrast, our method produces accurate and consistent story visualizations in terms of identity consistency, style preservation, prompt adherence, and layout composition.

### 4.3. Layout Generator Evaluation

We compare our layout generator with GPT-4 [36] to evaluate the validity of our method. We use a detailed prompt with in-context examples to guide GPT-4 to generate layouts as a list of bounding boxes, and evaluate the results according to the following categories:

- **Panel Count:** A correct number of panels should be present per page.
- **Page coverage:** Panels should collectively cover the page without large empty regions.
- **Panel ordering:** Panels should follow the conventional comic order (top-to-bottom, right-to-left).
- **Valid Character:** The character bounding box should cover a reasonable area within the panel.
- **Character Count:** A correct number of character bounding boxes should be present per panel.

We present the results in Table 4, which shows our model’s improvements in understanding and generating

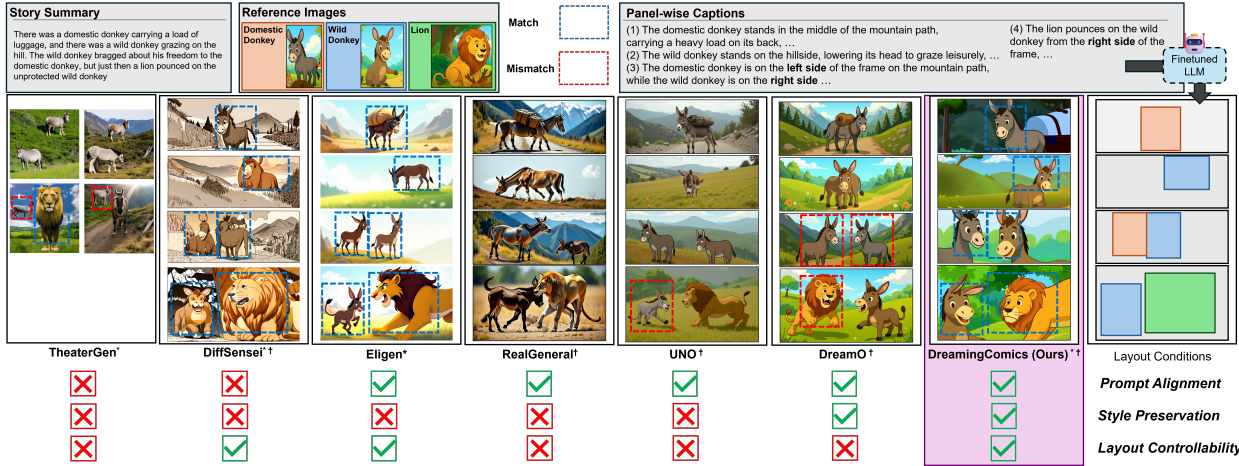


Figure 7. Story-level qualitative comparison of DreamingComics. Methods marked with “\*” use the layout condition generated from our layout generator. Methods marked with “†” use the reference images. Better viewed with zoom-in.

Regional RoPE	Masked Condition Loss	CIDS (Cross/Self)	CSD (Cross/Self)
✗	✓	38.7 / 50.3	41.0 / 52.9
✓	✗	56.6 / 61.2	50.5 / 56.3
✓	✓	<b>58.5 / 63.0</b>	<b>52.9 / 62.4</b>

Table 5. Ablation study on the impact of RegionalRoPE and masked condition loss. Our final model achieves the highest identity and style preservation performance without relying on regional attention masking.

Name	CIDS (Cross/Self)	CSD (Cross/Self)
Timestamp ( $t = 1$ )	55.3 / 58.8	50.3 / 56.1
Timestamp ( $t = 5$ )	55.7 / 58.9	50.2 / 54.3
Timestamp ( $t = 9$ )	54.7 / 58.2	49.7 / 56.1
Ours ( $t = 3$ )	<b>58.5 / 63.0</b>	<b>52.9 / 62.4</b>

Table 6. Ablation study on the impact of choosing the right timestamp for the target image.

comic-structured layouts.

#### 4.4. Ablation Study

We perform an ablation study to assess the contribution of our main components: RegionalRoPE and masked condition loss. As shown in Table 5, removing RegionalRoPE leads to the most significant performance drop, highlighting its role in spatial disentanglement. Excluding the masked condition loss also degrades layout fidelity. Overall, the combination of both components yields the best performance, validating our use of layout conditioning over rigid attention masking [6, 72]. Additionally, Table 6 evaluates the effect of choosing the target timestamp  $t$  for FramePack [73]. We observe that  $t = 3$  achieves the best trade-off between identity preservation and stylistic variation, as lower values produce rigid results, while higher values introduce undesirable drift. All models were trained for 6K steps on the single-subject samples and evaluated on ViStoryBench. We also include ablation studies on choosing  $\lambda_{\text{mask}}$  values in Supplementary.

Method Preference (%)	Character Identity	Character Style	Story Consistency	Layout Plausibility
Ours vs. Baseline	80.0	83.8	86.2	69.2

Table 7. User study results: percentage of participants who preferred our method over the baseline across four questions.

#### 4.5. User Study

To evaluate the subjective quality of our framework, we conducted a user study with 26 participants who were shown image pairs for four evaluation criteria: character identity, character style, story consistency, and layout plausibility. Each participant was presented with 20 comparison questions and was asked to select the preferred output. For the first 15 questions (three categories), the participants compared our results with those of DreamO [35] and UNO [63]. The remaining questions compare our layout generator with GPT-4. The user study questionnaire is included in Supplementary. As shown in Table 7, our work demonstrates a strong advantage in both image quality and layout control.

### 5. Conclusion

We present DreamingComics, a novel framework for subject- and layout-controllable story visualization. The problem is decomposed into layout generation and layout-based image customization. Our image customization module builds on a DiT video generation model, leveraging its spatiotemporal priors to ensure visual consistency across diverse artistic styles. For intuitive and semantically rich layout creation, we use a finetuned LLM-based layout generator. To enable layout-based control with multiple subjects, we introduce RegionalRoPE and a masked condition loss. Quantitative results and user studies demonstrate that DreamingComics produces faithful layouts and stylistically coherent visual narratives, establishing a strong foundation for future work in controllable story generation.

## References

- [1] Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel. Multidiffusion: fusing diffusion paths for controlled image generation. In *Proceedings of the 40th International Conference on Machine Learning*, 2023. 2
- [2] Black Forest Labs. FLUX: High-Fidelity Image Generation Model. <https://github.com/black-forest-labs/flux?tab=readme-ov-file>, 2025. Accessed: 2025-08-01. 1, 2, 6
- [3] Hengyuan Cao, Yutong Feng, Biao Gong, Yijing Tian, Yunhong Lu, Chuang Liu, and Bin Wang. Dimension-Reduction Attack! Video Generative Models are Experts on Controllable Image Synthesis. *ArXiv*, abs/2505.23325, 2025. 2, 4, 5, 6
- [4] Qiong Cao, Li Shen, Weidi Xie, Omkar M. Parkhi, and Andrew Zisserman. Vggface2: A dataset for recognising faces across pose and age. *2018 13th IEEE International Conference on Automatic Face & Gesture Recognition (FG 2018)*, pages 67–74, 2017. 6
- [5] Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and Armand Joulin. Emerging properties in self-supervised vision transformers. *2021 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 9630–9640, 2021. 7
- [6] Anthony Chen, Jianjin Xu, Wenzhao Zheng, Gaole Dai, Yida Wang, Renrui Zhang, Haofan Wang, and Shanghang Zhang. Training-free regional prompting for diffusion transformers. *arXiv preprint arXiv:2411.02395*, 2024. 1, 2, 8
- [7] Chaofeng Chen, Jiadi Mo, Jingwen Hou, Haoning Wu, Liang Liao, Wenxiu Sun, Qiong Yan, and Weisi Lin. Topiq: A top-down approach from semantics to distortions for image quality assessment. *IEEE Transactions on Image Processing*, 33:2404–2418, 2024. 6
- [8] Siyu Chen, Dengjie Li, Zenghao Bao, Yao Zhou, Lingfeng Tan, Yujie Zhong, and Zheng Zhao. Manga generation via layout-controllable diffusion, 2024. 3
- [9] Xi Chen, Zhifei Zhang, He Zhang, Yuqian Zhou, Soo Ye Kim, Qing Liu, Yijun Li, Jianming Zhang, Nanxuan Zhao, Yilin Wang, Hui Ding, Zhe L. Lin, and Hengshuang Zhao. Unireal: Universal image generation and editing via learning real-world dynamics. *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12501–12511, 2025. 2, 4
- [10] Junhao Cheng, Xi Lu, Hanhui Li, Khun Loun Zai, Baiqiao Yin, Yuhao Cheng, Yiqiang Yan, and Xiaodan Liang. Autostudio: Crafting consistent subjects in multi-turn interactive image generation. *ArXiv*, abs/2406.01388, 2024. 3
- [11] Junhao Cheng, Baiqiao Yin, Kaixin Cai, Minbin Huang, Hanhui Li, Yuxin He, Xi Lu, Yue Li, Yifei Li, Yuhao Cheng, et al. Theatergen: Character management with llm for consistent multi-turn image generation. *arXiv preprint arXiv:2404.18919*, 2024. 1, 2, 3
- [12] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4685–4694, 2019. 6
- [13] discuss0434. Aesthetic predictor v2.5. <https://github.com/discuss0434/aesthetic-predictor-v2-5>, 2024. Accessed: 2025-08-01. 6
- [14] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, Dustin Podell, Tim Dockhorn, Zion English, and Robin Rombach. Scaling rectified flow transformers for high-resolution image synthesis. In *Proceedings of the 41st International Conference on Machine Learning*, 2024. 1, 2
- [15] Azuma Fujimoto, Toru Ogawa, Kazuyoshi Yamamoto, Yusuke Matsui, T. Yamasaki, and Kiyoharu Aizawa. Manga109 dataset and creation of metadata. *Proceedings of the 1st International Workshop on coMics Analysis, Processing and Understanding*, 2016. 6
- [16] 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*, abs/2208.01618, 2022. 2
- [17] J. Edward Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *ArXiv*, abs/2106.09685, 2021. 6
- [18] Miao Hua, Jiawei Liu, Fei Ding, Wei Liu, Jie Wu, and Qian He. Dreamtuner: Single image is enough for subject-driven generation, 2023. 2
- [19] Lianghua Huang, Wei Wang, Zhi-Fan Wu, Yupeng Shi, Huanzhang Dou, Chen Liang, Yutong Feng, Yu Liu, and Jingren Zhou. In-context lora for diffusion transformers. *arXiv preprint arxiv:2410.23775*, 2024. 2
- [20] Mohit Iyyer, Varun Manjunatha, Anupam Guha, Yogarshi Vyas, Jordan Lee Boyd-Graber, Hal Daumé, and Larry S. Davis. The amazing mysteries of the gutter: Drawing inferences between panels in comic book narratives. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6478–6487, 2016. 6
- [21] Minchul Kim, Anil K. Jain, and Xiaoming Liu. Adaface: Quality adaptive margin for face recognition. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 18729–18738, 2022. 6
- [22] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013. 3
- [23] Weijie Kong, Qi Tian, Zijian Zhang, Rox Min, Zuoqun Dai, Jin Zhou, Jia-Liang Xiong, Xin Li, Bo Wu, Jianwei Zhang, Kathrina Wu, Qin Lin, Junkun Yuan, Yanxin Long, Aladdin Wang, Andong Wang, Changlin Li, Duoqun Huang, Fan Yang, Hao Tan, Hongmei Wang, Jacob Song, Jiawang Bai, Jianbing Wu, Jinbao Xue, Joey Wang, Kai Wang, Mengyang Liu, Peng-Yu Li, Shuai Li, Weiyang Wang, Wenqing Yu, Xi Deng, Yang Li, Yi Chen, Yutao Cui, Yuanbo Peng, Zhen Yu, Zhiyu He, Zhiyong Xu, Zixiang Zhou, Z. Xu, Yang-Dan Tao, Qinglin Lu, Songtao Liu, Daquan Zhou, Hongfa Wang, Yong Yang, Di Wang, Yuhong Liu, Jie Jiang, and Caesar Zhong. Hunyuanvideo: A systematic framework for large video generative models. *ArXiv*, abs/2412.03603, 2024. 2, 3, 6

- [24] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-concept customization of text-to-image diffusion. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1931–1941, 2022. 2
- [25] Dongxu Li, Junnan Li, and Steven C.H. Hoi. Blip-diffusion: pre-trained subject representation for controllable text-to-image generation and editing. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, 2023. 2
- [26] Yuheng Li, Haotian Liu, Qingyang Wu, Fangzhou Mu, Jianwei Yang, Jianfeng Gao, Chunyuan Li, and Yong Jae Lee. GLIGEN: Open-Set Grounded Text-to-Image Generation. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22511–22521, Los Alamitos, CA, USA, 2023. IEEE Computer Society. 2
- [27] Long Lian, Boyi Li, Adam Yala, and Trevor Darrell. Llm-grounded diffusion: Enhancing prompt understanding of text-to-image diffusion models with large language models. *Trans. Mach. Learn. Res.*, 2024, 2023. 3
- [28] Yijing Lin, Mengqi Huang, Shuhan Zhuang, and Zhen-dong Mao. Realgeneral: Unifying visual generation via temporal in-context learning with video models, 2025. 2, 4, 5, 6, 7
- [29] Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling. *ArXiv*, abs/2210.02747, 2022. 4
- [30] Chang Liu, Haoning Wu, Yujie Zhong, Xiaoyun Zhang, Yanfeng Wang, and Weidi Xie. Intelligent grimm - open-ended visual storytelling via latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 6190–6200, 2024. 3
- [31] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2017. 6
- [32] Wan-Duo Kurt Ma, J. P. Lewis, Avisek Lahiri, Thomas Leung, and W. Bastiaan Kleijn. Directed diffusion: Direct control of object placement through attention guidance. 2023. 2
- [33] Jiawei Mao, Xiaoke Huang, Yunfei Xie, Yuanqi Chang, Mude Hui, Bingjie Xu, and Yuyin Zhou. Story-Adapter: A Training-free Iterative Framework for Long Story Visualization, 2024. 1, 2, 6
- [34] Chong Mou, Xintao Wang, Liangbin Xie, Jing Zhang, Zhongang Qi, Ying Shan, and Xiaohu Qie. T2i-adapter: Learning adapters to dig out more controllable ability for text-to-image diffusion models. In *AAAI Conference on Artificial Intelligence*, 2023. 1
- [35] Chong Mou, Yanze Wu, Wenxu Wu, Zinan Guo, Pengze Zhang, Yufeng Cheng, Yiming Luo, Fei Ding, Shiwen Zhang, Xinghui Li, Mengtian Li, Songtao Zhao, Jian Zhang, Qian He, and Xinglong Wu. Dreamo: A unified framework for image customization. *ArXiv*, abs/2504.16915, 2025. 1, 2, 6, 7, 8
- [36] OpenAI. Gpt-4 technical report. 2023. 4, 7
- [37] Xichen Pan, Pengda Qin, Yuhong Li, Hui Xue, and Wenhui Chen. Synthesizing coherent story with autoregressive latent diffusion models. *2024 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2908–2918, 2022. 3
- [38] William Peebles and Saining Xie. Scalable diffusion models with transformers. In *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4172–4182, 2023. 1, 2
- [39] Yuang Peng, Yuxin Cui, Haomiao Tang, Zekun Qi, Runpei Dong, Jing Bai, Chunrui Han, Zheng Ge, Xiangyu Zhang, and Shu-Tao Xia. Dreambench++: A human-aligned benchmark for personalized image generation. In *The Thirteenth International Conference on Learning Representations*, 2025. 2, 6
- [40] 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, 6
- [41] Lu Qiu, Yizhuo Li, Yuying Ge, Yixiao Ge, Ying Shan, and Xihui Liu. Animeshooter: A multi-shot animation dataset for reference-guided video generation, 2025. 6
- [42] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning*, 2021. 6, 7
- [43] 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 (CVPR)*, pages 10684–10695, 2022. 1
- [44] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241, Cham, 2015. Springer International Publishing. 2
- [45] Noam Rotstein, G. Yona, Daniel Silver, Roy Velich, David Bensaid, and Ron Kimmel. Pathways on the image manifold: Image editing via video generation. *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 7857–7866, 2024. 2, 4
- [46] 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. *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 22500–22510, 2022. 2, 7
- [47] Ragav Sachdeva and Andrew Zisserman. The manga whisperer: Automatically generating transcriptions for comics. *2024 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12967–12976, 2024. 6
- [48] Ragav Sachdeva, Gyungin Shin, and Andrew Zisserman. Tails tell tales: Chapter-wide manga transcriptions with

- character names. In *Proceedings of the Asian Conference on Computer Vision (ACCV)*, pages 2053–2069, 2024. 6
- [49] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, page 2234–2242, Red Hook, NY, USA, 2016. Curran Associates Inc. 6
- [50] Gowthami Somepalli, Anubhav Gupta, Kamal K. Gupta, Shramay Palta, Micah Goldblum, Jonas Geiping, Abhinav Shrivastava, and Tom Goldstein. Measuring style similarity in diffusion models. *ArXiv*, abs/2404.01292, 2024. 6
- [51] Jianlin Su, Murtadha Ahmed, Yu Lu, Shengfeng Pan, Wen Bo, and Yunfeng Liu. Roformer: Enhanced transformer with rotary position embedding. *Neurocomput.*, 568(C), 2024. 3
- [52] Zhenxiong Tan, Songhua Liu, Xingyi Yang, Qiaochu Xue, and Xinchao Wang. Ominicontrol: Minimal and universal control for diffusion transformer. 2025. 2, 5, 6
- [53] Zhenxiong Tan, Qiaochu Xue, Xingyi Yang, Songhua Liu, and Xinchao Wang. Ominicontrol2: Efficient conditioning for diffusion transformers. *ArXiv*, abs/2503.08280, 2025. 5
- [54] Emanuele Vivoli, Andrey Barsky, Mohamed Ali Souibgui, Artemis Llabres, Marco Bertini, and Dimosthenis Karatzas. One missing piece in vision and language: A survey on comics understanding. *ArXiv*, abs/2409.09502, 2024. 3
- [55] Emanuele Vivoli, Artemis Llabrés, Mohamed Ali Souibgui, Marco Bertini, Ernest Valveny Llobet, and Dimosthenis Karatzas. Comicspap: Understanding comic strips by picking the correct panel. In *Document Analysis and Recognition – ICDAR 2025*, pages 337–350, Cham, 2025. Springer Nature Switzerland. 6
- [56] Ruichen Wang, Zekang Chen, Chen Chen, Jian Ma, Haonan Lu, and Xiaodong Lin. Compositional text-to-image synthesis with attention map control of diffusion models. In *Proceedings of the Thirty-Eighth AAAI Conference on Artificial Intelligence and Thirty-Sixth Conference on Innovative Applications of Artificial Intelligence and Fourteenth Symposium on Educational Advances in Artificial Intelligence*. AAAI Press, 2024. 2
- [57] Wen Wang, Canyu Zhao, Hao Chen, Zhekai Chen, Kecheng Zheng, and Chunhua Shen. Autostory: Generating diverse storytelling images with minimal human effort. *ArXiv*, abs/2311.11243, 2023. 3
- [58] Xudong Wang, Trevor Darrell, Sai Saketh Rambhatla, Rohit Girdhar, and Ishan Misra. Instancediffusion: Instance-level control for image generation, 2024. 2
- [59] X. Wang, Siming Fu, Qihan Huang, Wangui He, and Hao Jiang. Ms-diffusion: Multi-subject zero-shot image personalization with layout guidance. *ArXiv*, abs/2406.07209, 2024. 2
- [60] 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. *arXiv preprint arXiv:2302.13848*, 2023. 2
- [61] Chenyuan Wu, Pengfei Zheng, Ruiran Yan, Shitao Xiao, Xin Luo, Yueze Wang, Wanli Li, Xiyan Jiang, Yexin Liu, Junjie Zhou, Ze Liu, Ziyi Xia, Chaofan Li, Haoge Deng, Jiahao Wang, Kun Luo, Bo Zhang, Defu Lian, Xinlong Wang, Zhongyuan Wang, Tiejun Huang, and Zheng Liu. Omnigen2: Exploration to advanced multimodal generation. *ArXiv*, abs/2506.18871, 2025. 1, 2
- [62] Jianzong Wu, Chao Tang, Jingbo Wang, Yanhong Zeng, Xiangtai Li, and Yunhai Tong. Diffsensei: Bridging multi-modal llms and diffusion models for customized manga generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 28684–28693, 2025. 2, 3, 6, 7
- [63] Shaojin Wu, Mengqi Huang, Wenxu Wu, Yufeng Cheng, Fei Ding, and Qian He. Less-to-more generalization: Unlocking more controllability by in-context generation. *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. 1, 2, 5, 6, 7, 8
- [64] Shitao Xiao, Yueze Wang, Junjie Zhou, Huaying Yuan, Xingrun Xing, Ruiran Yan, Shuting Wang, Tiejun Huang, and Zheng Liu. Omnigen: Unified image generation. *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13294–13304, 2024. 2
- [65] Xuenan Xu, Jiahao Mei, Chenliang Li, Yuning Wu, Ming Yan, Shaopeng Lai, Ji Zhang, and Mengyue Wu. Mm-storyagent: Immersive narrated storybook video generation with a multi-agent paradigm across text, image and audio. *arXiv preprint arXiv:2503.05242*, 2025. 1
- [66] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. Qwen2 technical report, 2024. 4, 6
- [67] Shuai Yang, Yuying Ge, Yang Li, Yukang Chen, Yixiao Ge, Ying Shan, and Yingcong Chen. Seed-story: Multimodal long story generation with large language model. *arXiv preprint arXiv:2407.08683*, 2024. 3
- [68] Zhuoyi Yang, Jiayan Teng, Wendi Zheng, Ming Ding, Shiyu Huang, Jiazheng Xu, Yuanming Yang, Wenyi Hong, Xiaohan Zhang, Guanyu Feng, Da Yin, Xiaotao Gu, Yuxuan Zhang, Weihang Wang, Yean Cheng, Ting Liu, Bin Xu, Yuxiao Dong, and Jie Tang. Cogvideox: Text-to-video diffusion models with an expert transformer. *arXiv preprint arXiv:2408.06072*, 2024. 2, 6
- [69] Hu Ye, Jun Zhang, Siyi Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for

- text-to-image diffusion models. *ArXiv*, abs/2308.06721, 2023. [1](#), [2](#)
- [70] Xin Yu, Tianyu Wang, Soo Ye Kim, Paul Guerrero, Xi Chen, Qing Liu, Zhe L. Lin, and Xiaojuan Qi. Object-mover: Generative object movement with video prior. *2025 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 17682–17691, 2025. [2](#)
- [71] Shenghai Yuan, Xianyi He, Yufan Deng, Yang Ye, Jinfa Huang, Bin Lin, Jiebo Luo, and Li Yuan. Opens2v-nexus: A detailed benchmark and million-scale dataset for subject-to-video generation. *ArXiv*, abs/2505.20292, 2025. [6](#)
- [72] Hong Zhang, Zhongjie Duan, Xingjun Wang, Yingda Chen, and Yu Zhang. Eligen: Entity-level controlled image generation with regional attention. *arXiv preprint arXiv:2501.01097*, 2025. [1](#), [2](#), [6](#), [7](#), [8](#)
- [73] Lvmin Zhang and Maneesh Agrawala. Packing input frame contexts in next-frame prediction models for video generation. *Arxiv*, 2025. [2](#), [4](#), [6](#), [8](#)
- [74] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. *2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 3813–3824, 2023. [1](#), [2](#)
- [75] Yuxuan Zhang, Yiren Song, Jiaming Liu, Rui Wang, Jinpeng Yu, Hao Tang, Huaxia Li, Xu Tang, Yao Hu, Han Pan, et al. Ssr-encoder: Encoding selective subject representation for subject-driven generation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8069–8078, 2024. [2](#)
- [76] Yuxuan Zhang, Yirui Yuan, Yiren Song, Haofan Wang, and Jiaming Liu. Easycontrol: Adding efficient and flexible control for diffusion transformer. *arXiv preprint arXiv:2503.07027*, 2025. [2](#)
- [77] Yabo Zhang, Xinpeng Zhou, Yihan Zeng, Hang Xu, Hui Li, and Wangmeng Zuo. Framepainter: Endowing interactive image editing with video diffusion priors. *2025 IEEE/CVF International Conference on Computer Vision (ICCV)*, 2025. [2](#)
- [78] Dewei Zhou, You Li, Fan Ma, Xiaoting Zhang, and Yi Yang. Migc: Multi-instance generation controller for text-to-image synthesis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6818–6828, 2024. [2](#)
- [79] Yupeng Zhou, Daquan Zhou, Ming-Ming Cheng, Jia-ashi Feng, and Qibin Hou. Storydiffusion: Consistent self-attention for long-range image and video generation. *NeurIPS 2024*, 2024. [1](#), [2](#), [3](#), [6](#)
- [80] Cailin Zhuang, Ailin Huang, Wei Cheng, Jingwei Wu, Yaoqi Hu, Jiaqi Liao, Zhewei Huang, Hongyuan Wang, Xinyao Liao, Weiwei Cai, Hengyuan Xu, Xuanyang Zhang, Xianfang Zeng, Gang Yu, and Xiaoyan Zhang. Vistorybench: Comprehensive benchmark suite for story visualization. *ArXiv*, abs/2505.24862, 2025. [2](#), [6](#)
- [81] Shaobin Zhuang, Kunchang Li, Xinyuan Chen, Yao-hui Wang, Ziwei Liu, Yu Qiao, and Yali Wang. Vlogger: Make your dream a vlog. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8806–8817, 2024. [6](#)