

High-Resolution Artwork Outpainting with Global Blueprint Guidance and Layout Control

Junha Kim[✉], Hyunjoon Park[✉], and Donghyeon Cho[†][✉]

Hanyang University, Seoul, Republic of Korea
{poohoh, junippini83, doncho}@hanyang.ac.kr

Abstract. Image outpainting extends an image beyond its original borders, requiring seamless style integration and globally coherent scene completion. Building on the success of diffusion models, recent methods have achieved substantial improvements in visual quality. In practice, however, high-resolution outpainting is commonly performed via progressive expansion around a fixed source image, particularly in artwork scenarios. Despite this progress, existing approaches still suffer from three key limitations: (i) the absence of a reliable global planning mechanism, which leads to structural instability and error accumulation at high resolutions; (ii) limited spatial controllability beyond text prompts, making it difficult to place objects at user-specified locations; and (iii) high inference latency caused by inherently sequential patch generation. To address these issues, we propose a global blueprint-guided two-stage diffusion framework for layout-controllable high-resolution outpainting with efficient parallel synthesis. In Stage 1, we generate a low-resolution global blueprint using a layout adapter that injects bounding-box conditions into a Stable Diffusion inpainting backbone, producing a globally consistent structural plan while extracting global guidance features. In Stage 2, we synthesize high-resolution local patches in parallel by injecting the blueprint-derived global guidance and initializing each patch from the blueprint using the low-frequency preservation property of forward diffusion. This design eliminates sequential dependency while maintaining global coherence. Extensive experiments on large-scale artwork datasets demonstrate improved visual fidelity, stronger semantic consistency, and substantially reduced inference time compared to prior baselines, while uniquely supporting explicit layout control for artwork outpainting. Our code will be available at <https://github.com/poohoh/BlueOut>.

Keywords: Outpainting · Artworks · High-resolution · Layout control

1 Introduction

Image outpainting extends an image beyond its original borders. The key challenge lies in seamless integration, which requires matching the style, color, and context of the source image rather than merely enlarging the canvas. Although

[†] Corresponding author.




Fig. 1: Challenges in High-resolution Artwork Outpainting. (a) Progressive methods accumulate errors without global planning; our blueprint-guided approach maintains global coherence. (b) Text-only conditioning struggles with object placement; our layout-conditioned framework enables spatial control. (c) Sequential generation incurs high latency; our parallel synthesis achieves up to $2.4\times$ speedup.

outpainting is often considered analogous to inpainting, it is inherently more challenging. In inpainting, the target region is surrounded by known pixels, which facilitates reliable context inference. In contrast, outpainting synthesizes content in unbounded regions outside the image, introducing substantially greater uncertainty. Consequently, successful outpainting requires not only plausible local synthesis but also an accurate understanding of the global context to maintain structural naturalness and contextual consistency. Moreover, in artwork outpainting for immersive display applications such as large-scale projection and VR/AR environments [9], legacy paintings must be adapted to new canvas sizes while preserving their original visual fidelity and aspect ratios. This requirement makes the process highly sensitive to stylistic distortion and compositional imbalance, thereby making accurate global context modeling particularly critical.

Building on the success of diffusion models, recent outpainting methods have achieved substantial improvements in visual quality. In practice, most existing approaches operate at fixed resolutions, typically 512×512 or 256×256 , which correspond to the standard output sizes of diffusion backbones [23, 30, 43, 64]. To enable high-resolution expansion, several progressive strategies have been proposed that iteratively extend the image by sliding a local window across the canvas [26, 38, 54]. Although diffusion models have been successfully applied to progressive outpainting, existing methods still suffer from several limitations.

As illustrated in Figure 1(a), the first issue is the absence of a dedicated planning stage that provides reliable global blueprint for progressive high-resolution outpainting. For example, ProOut [38] injects a global hint image as a condition; however, it reflects only the intermediate canvas during progressive expansion rather than a predetermined global composition. As a result, the method still relies on sequential generation and stitching of local windows. As noted by prior studies [10, 51], this design leads to error accumulation, where inconsistencies from earlier stages propagate to later ones. Consequently, without an explicit

global blueprint, sequential expansion becomes sensitive to generation order and risks degrading global coherence in high-resolution results. To fundamentally mitigate sequential error accumulation, we introduce a global blueprint-guided two-stage framework. We first establish a global structural guideline by generating a low-resolution blueprint. Based on this, we synthesize high-resolution details, preventing error propagation and ensuring stable quality even at scale.

As shown in Figure 1(b), the second issue is the lack of a mechanism in current state-of-the-art outpainting models to precisely control generated content through layout conditions. Before diffusion models become mainstream, early generative approaches, such as those utilizing GANs, attempt to generate specific objects or scenes at desired locations. They achieve this by using layouts [24, 33, 52] or sketches [45, 46] as conditions, or by disentangling the structural semantics of an image during training [48]. However, these approaches rely on outdated backbones, yielding generation quality far below state-of-the-art techniques. Among more recent diffusion-based studies, there have been attempts to control generation using text prompts [53]. While text-based conditions can specify the overall style or object presence, they cannot precisely control fine-grained layouts, such as exact positions or sizes. Thus, text alone cannot ensure intended spatial arrangements. In this paper, we model a mechanism that injects layout conditions, enabling accurate generation at user-specified locations while preserving the powerful generation capability of diffusion models.

The final issue, highlighted in Figure 1(c), is the extreme computational inefficiency and long inference latency caused by sequential generation. Diffusion models require iterative denoising, and progressive methods split the entire canvas and repeat sequential generation for high-resolution expansion, thus they inherently incur high computational costs even when generating a single image and runtime grows with the number of patches. Moreover, this sequential dependency limits hardware utilization. Due to the structural constraint where prior window generation must be completed before the next can begin, parallel processing is impossible even with sufficient memory. This inevitably acts as a bottleneck for the entire system. Recently, in video outpainting, Follow-Your-Canvas [4] attempts to accelerate high-resolution generation by dividing the canvas into multiple spatial windows and decoding in parallel. This method adopts a strategy similar to MultiDiffusion [2], blending overlaps to mitigate boundary artifacts. However, overlap blending only smooths pixel transitions and fails to control global semantics. This motivates our blueprint-guided coordination design, which uses the global blueprint as a shared structural guide for parallel high-resolution synthesis rather than relying on overlap blending alone.

Our main contributions can be summarized as follows:

- **Blueprint-Guided Coordination Design:** We use the low-resolution global blueprint as a shared structural guide for high-resolution local synthesis, reducing error accumulation from sequential generation.
- **Layout-Controllable Framework:** We incorporate explicit layout conditioning in the planning stage, enabling users to specify object locations and descriptions in the expanded regions.

- **Efficient Parallel Synthesis:** We enable parallel patch generation for high-resolution artwork outpainting, reducing the sequential dependency of progressive methods and supporting multi-GPU inference.

2 Related Works

This section reviews prior work most relevant to our study, including outpainting, layout-controllable generation, and reference-guided synthesis.

Outpainting. Early outpainting studies [6, 11, 19, 27, 36, 42, 54, 55] primarily relied on GANs. While several works utilized layouts, sketches, or scene graphs to control object placement [24, 45, 52], they inherently suffered from limited generation quality and diversity due to their outdated backbones. Recently, diffusion models have become the mainstream, driven by their superior synthesis quality. For instance, Stable Diffusion (SD) Inpainting [35] enabled inpainting and outpainting via fine-tuning with random masks, and PowerPaint [64] introduced learnable task prompts for specific control. Other methods explored human aesthetic alignment [30], context-aware outpainting via multimodal large language models [53], and training-free noise optimization [23]. However, most diffusion-based methods operate at fixed resolutions (e.g., 512×512) and rely solely on text prompts, lacking fine-grained layout control. While some studies achieved continuous resolution expansion [58, 59], they cannot incorporate multimodal conditions. Although ProOut [38] proposed progressive high-resolution expansion, it still suffers from severe inference latency and error accumulation due to sequential window generation.

Layout-Controllable Generation. To overcome the limited spatial controllability of standard text-to-image models, explicit layout conditioning has been extensively studied. Methods such as GLIGEN [25], ControlNet [57], LayoutDiffusion [60], InstanceDiffusion [44], and MIGC [61, 62] inject structural conditions, such as bounding boxes, masks, or poses, to guide object synthesis at user-specified locations. However, these approaches are primarily restricted to general image generation or local editing tasks. To the best of our knowledge, applying such advanced layout-control mechanisms to artwork outpainting, particularly to govern the global composition of extended regions, has not been thoroughly explored.

Reference-Guided Synthesis. In both image and video generation, reference-guided synthesis methods [7, 49, 50, 56, 63] aim to preserve the appearance or style of a reference image during generation. This is typically achieved by extracting and injecting features via a supplementary encoder, such as ReferenceNet [17]. While these approaches predominantly focus on transferring the visual appearance of specific objects, we fundamentally reinterpret this mechanism. In our framework, we leverage it to propagate the overall structural blueprint from the planning stage into the local patch generation process, thereby ensuring global structural and contextual consistency in high-resolution artwork outpainting.

3 Method

3.1 Preliminaries

Latent Diffusion Models. Latent diffusion models (LDMs) [35] perform diffusion in a compressed latent space to improve computational efficiency. Given an input image x , a pre-trained encoder $\mathcal{E}$ maps it to a latent representation $z_0 = \mathcal{E}(x)$. In the forward process, Gaussian noise is gradually added according to a variance schedule $\{\beta_t\}_{t=1}^T$, yielding:

$$z_t = \sqrt{\bar{\alpha}_t} z_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad \epsilon \sim \mathcal{N}(0, I), \quad (1)$$

where $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$. At inference, denoising starts from $z_T \sim \mathcal{N}(0, I)$ and iteratively recovers z_0 , which is decoded via decoder $\mathcal{D}$.

Attention-Guided Noise Optimization. Prior to denoising in Stage 1, we optimize the initial noise by analyzing the self-attention maps of U-Net during the first denoising step. Specifically, we adjust the noise so that spatial tokens in the target region sufficiently attend to those in the source region, following AlignNoise [23]. This mitigates the common failure where the generated region becomes semantically disconnected from the original image.

Low-Frequency Preservation in Forward Diffusion. In standard diffusion theory, the final timestep T is often regarded as pure Gaussian noise. However, recent analyses [28, 47] show that this assumption does not strictly hold under typical noise schedules. In particular, FreeInit demonstrates that while high-frequency components (e.g., textures and edges) are rapidly attenuated during early diffusion steps, low-frequency components (e.g., global layout and object placement) decay much more slowly and can remain partially preserved even at timestep T . This low-frequency preservation implies that the macroscopic structure of an image can survive heavy noise corruption. We exploit this property in our framework. Specifically, we apply forward diffusion to the global blueprint generated in Stage 1, obtaining a noisy latent in which fine details are suppressed while global structural information is retained. For each local patch in Stage 2, we crop the corresponding region from this noisy latent and use it as the initial noise. As a result, all patches begin denoising from structurally aligned noise, enabling parallel synthesis while preserving global coherence.

3.2 Overall Architecture

Our framework follows a global blueprint-guided two-stage pipeline, as illustrated in Figure 2. The core idea is to decouple global structural planning from high-resolution local synthesis. The global blueprint then serves as a shared structural guide for parallel local synthesis, reducing error accumulation from sequential generation. Stage 1 performs outpainting at low resolution over the full canvas to produce a global blueprint, simultaneously exporting timestep-wise guidance features. Stage 2 then synthesizes high-resolution local patches in




Fig. 2: Overview of Our Two-stage Framework. In Stage 1, the Attention Alignment Measure (AAM) first optimizes the initial noise to prevent semantic disconnection from the source image. The Layout Adapter then encodes user-specified bounding boxes b_k and object descriptions d_k into instance tokens C . The Gated Fuser injects C into the global U-Net to generate a low-resolution blueprint $\hat{x}^g$ while exporting guidance features $\mathcal{F} = \{F_t^g\}_{t=1}^T$. In Stage 2, each high-resolution patch is initialized from $\hat{x}^g$ via forward diffusion and denoised in parallel. The Position Network encodes the patch’s bounding box b_i^p and a learnable null token $n_\emptyset$ into a position token p_i , and the Gated Fuser injects $\mathcal{F}$ and p_i into the main U-Net. Module details are shown on the right.

parallel, each initialized from the blueprint and guided by the exported features. This design eliminates the sequential dependency of progressive methods while preserving global coherence.

We denote the Stage 1 blueprint generator as G_{θ_1} and the Stage 2 local patch synthesizer as P_{θ_2} . The overall pipeline is then formulated as follows. In Stage 1, G_{θ_1} produces a global blueprint and guidance features:

$$(\hat{x}^g, \mathcal{F}) = G_{\theta_1}(z_T^g, \tilde{x}^g, m^g, \mathcal{B}, c_G), \quad (2)$$

where $\hat{x}^g$ is the generated global blueprint, $\mathcal{F}$ is the global guidance feature bank from Stage 1, z_T^g is the initial noise, $\tilde{x}^g$ is the global masked image, m^g is the global binary mask, $\mathcal{B} = \{(b_k, d_k)\}_{k=1}^K$ denotes the set of layout conditions (bounding boxes and object descriptions), and c_G is the global prompt. In Stage 2, P_{θ_2} generates each high-resolution patch:

$$\hat{x}^{(i)} = P_{\theta_2}(\hat{x}^g, \tilde{x}^{(i)}, m^{(i)}, \mathcal{F}, p_i, c_L), \quad i = 1, \dots, N_{\text{patch}}, \quad (3)$$

where $\hat{x}^{(i)}$ is the generated local patch for region i , $\tilde{x}^{(i)}$ is the local masked image, $m^{(i)}$ is the local mask, p_i is the local position token encoding the coordinates of patch i , c_L is the local prompt, and N_{patch} is the total number of patches.

3.3 Stage 1: Layout-Controllable Global Blueprint Generation

The objective of Stage 1 is to produce a low-resolution global blueprint for diffusion initialization and structural guidance features for the subsequent Stage 2.

Model Architecture. The backbone of Stage 1 consists of a 9-channel SD Inpainting U-Net [35], which takes the masked image $\tilde{x}^g$ and a binary mask m^g as additional inputs. This allows the model to distinguish between the known source region and the target region. To inject layout conditions, we integrate the Layout Adapter and the Gated Fuser from InstanceDiffusion [44] into the U-Net. In our framework, these modules condition global blueprint generation, and the resulting blueprint is then used to guide parallel Stage 2 synthesis. For each object k , the Layout Adapter generates a condition token c_k by:

$$c_k = \text{MLP}([\gamma(b_k); e_k]), \quad e_k = \mathcal{T}(d_k), \quad (4)$$

where $\gamma(\cdot)$, b_k , d_k , $\mathcal{T}$, and $[\cdot; \cdot]$ denote the Fourier mapping [41], bounding box coordinates, object description, pre-trained CLIP text encoder [34], and concatenation, respectively. The K instance tokens are aggregated as $C = [c_1; \dots; c_K]$. Let $H_t^g \in \mathbb{R}^{N \times d}$ denote the feature tokens from the Stage 1 U-Net at denoising step t . The Gated Fuser is applied at every transformer block of the U-Net; we omit the block index ℓ for brevity. The Gated Fuser injects the instance tokens C into the U-Net features H_t^g and produces the guidance feature F_t^g as:

$$F_t^g = H_t^g + \alpha_g \cdot \text{Attn}\left(Q(H_t^g), K([H_t^g; C]), V([H_t^g; C])\right), \quad (5)$$

where α_g is a zero-initialized learnable gate parameter and $\text{Attn}(\cdot)$ denotes standard scaled dot-product attention. The zero-initialization preserves the pre-trained generative capabilities of U-Net at the beginning of training and gradually enables conditional fusion. At every denoising step t , we cache F_t^g into a feature bank $\mathcal{F} = \{F_t^g\}_{t=1}^T$, which aggregates features across all Gated Fuser blocks and denoising steps, and is passed to Stage 2 as global guidance. Since both stages share the same U-Net architecture, the guidance features are injected into the corresponding blocks of the Stage 2 U-Net with matching resolution and channel dimension. Simultaneously, the completed denoising pass yields the clean latent z_0^g , which is decoded into the global blueprint $\hat{x}^g = \mathcal{D}(z_0^g)$.

Attention-Guided Noise Optimization. Prior to performing denoising, we optimize the initial noise z_T^g to ensure semantic consistency between the source and target regions. Specifically, by analyzing the distribution of the self-attention maps A during the initial denoising step, we adjust the noise such that the target region sufficiently attends to the source region [23]. This process proactively prevents a common issue in outpainting, where the generated region produces content that is semantically disconnected from the original image.

Training Strategy. We train Stage 1 with standard noise-prediction objective:

$$\mathcal{L}_{S1} = \mathbb{E}_{z_0^g, \epsilon, t} \left[\|\epsilon - \epsilon_{\theta_1}(z_t^g, t; \tilde{x}^g, m^g, \mathcal{B}, c_G)\|_2^2 \right], \quad (6)$$




Fig. 3: Parallel Synthesis of High-resolution Local Patches. (Left) Eight patches are arranged around the source region with overlap margins, which are blended by latent-space averaging at each denoising step for seamless composition. (Right) Each patch can be assigned to a separate GPU and denoised in parallel under global blueprint guidance. Patch indices follow ProOut [38] for consistency but denote composing priority, not sequential generation order.

where $z_0^g = \mathcal{E}(x^g)$ is the latent of the global target image and z_t^g is the noisy latent obtained by applying the forward process (Eq. (1)) at timestep t . Only the additionally integrated modules are optimized while the U-Net backbone remains frozen. Zero-initializing the gate parameter α_g ensures a stable transition to layout-aware generation.

3.4 Stage 2: Parallel Local Patch Synthesis with Global Guidance

The objective of Stage 2 is to generate high-resolution local patches in parallel based on the Stage 1 global blueprint. To achieve this, Stage 2 combines three components: global guidance features from the blueprint, a patch position token that specifies the location of each patch on the full canvas, and blueprint-based initial noise construction using the low-frequency preservation property.

Model Architecture. The main U-Net of Stage 2 also employs the SD Inpainting U-Net [35] as its backbone, featuring a Gated Fuser [44] inserted between the self-attention and cross-attention layers. Unlike Stage 1, the position token p_i in Stage 2 encodes only the spatial coordinates of each patch i on the full canvas, without object description:

$$p_i = \text{MLP}([\gamma(b_i^p); n_\emptyset]), \quad (7)$$

where b_i^p denotes the bounding box of patch i on the full canvas and $n_\emptyset$ is a learnable null token that replaces the CLIP embedding used in Stage 1. Let $H_t^{(i)} \in \mathbb{R}^{N \times d}$ be the main U-Net features for patch i , and F_t^g the global guidance features retrieved from the Stage 1 feature bank at step t . The Stage 2 Gated Fuser combines these with the local position token p_i :

$$\tilde{H}_t^{(i)} = H_t^{(i)} + \alpha_p \cdot \text{Attn}(Q(H_t^{(i)}), K([H_t^{(i)}; F_t^g; p_i]), V([H_t^{(i)}; F_t^g; p_i])) \quad (8)$$

where $\tilde{H}_t^{(i)}$ is the fused output feature and α_p is the learnable gate parameter for Stage 2. Here, F_t^g provides global scene context from the blueprint, while p_i specifies the location of the current patch on the full canvas, enabling blueprint-guided coordination during parallel synthesis.

Initial Noise Construction Based on the Global Blueprint. To transfer the structural information of the global blueprint to each patch, we introduce an initial noise construction strategy that leverages the low-frequency preservation property [28, 47]. First, the global blueprint $\hat{x}^g$ decoded from Stage 1 is upsampled to the target canvas resolution, and each patch region is cropped and re-encoded into the latent space via:

$$z_{\text{bp}}^{(i)} = \mathcal{E}(\text{Crop}_i(\text{Up}(\hat{x}^g))), \quad (9)$$

where $z_{\text{bp}}^{(i)}$ denotes the blueprint-derived latent for patch i , $\text{Crop}_i(\cdot)$ extracts the region of patch i , and $\text{Up}(\cdot)$ is upsampling to the target canvas resolution. A single noise map $\epsilon^{\text{full}} \sim \mathcal{N}(0, I)$ is generated at the latent resolution of the entire canvas. Forward diffusion (Eq. (1)) is then applied to each patch latent using the corresponding region of the global noise map:

$$z_{T,\text{init}}^{(i)} = \sqrt{\bar{\alpha}_T} z_{\text{bp}}^{(i)} + \sqrt{1 - \bar{\alpha}_T} \text{Crop}_i(\epsilon^{\text{full}}). \quad (10)$$

Because a single global noise map is used, identical noise values are assigned to the same canvas locations across adjacent patches, preventing boundary discontinuities. Due to the low-frequency preservation property, the layout and structural information of the blueprint are partially retained after forward diffusion, allowing all patches to begin denoising from structurally aligned initial states.

Training Strategy. Similar to Stage 1, the parameters of the U-Net are frozen, and only the Gated Fuser and the Position Network that produces p_i are trained. Note that the blueprint-based initial noise construction is applied only at inference. The Stage 2 training objective is:

$$\mathcal{L}_{S2} = \mathbb{E}_{i, z_0^{(i)}, \epsilon, t} \left[\left\| \epsilon - \epsilon_{\theta_2}(z_t^{(i)}, t; \tilde{x}^{(i)}, m^{(i)}, F_t^g, p_i, c_L) \right\|_2^2 \right], \quad (11)$$

where c_L is the local prompt. The two stages are optimized independently:

$$\theta_1^* = \arg \min_{\theta_1} \mathcal{L}_{S1}, \quad \theta_2^* = \arg \min_{\theta_2} \mathcal{L}_{S2}. \quad (12)$$

Here, only the prediction of the main U-Net contributes to Eq. (11), while the frozen Stage 1 model serves solely as a feature extractor.

3.5 Inference Pipeline

Stage 1: Global Blueprint Generation. Given the source image and desired canvas size, a binary mask m^g is generated that designates the region outside the source as the target. If layout conditions $\mathcal{B}$ are provided by the user, they are

transformed into condition tokens C via the Layout Adapter. Before denoising, the attention-guided noise optimization described in Sec. 3.3 is applied to refine the initial noise. Outpainting is then performed from the optimized noise $z_T^{g,*}$ to generate a low-resolution global blueprint $\hat{x}^g$. In this process, the output features of the Gated Fuser are extracted and stored in a feature bank $\mathcal{F}$ at each denoising step, which serves as global guidance for Stage 2.

Stage 2: Parallel Local Patch Synthesis. As shown in Figure 3, in the default evaluation, the canvas is partitioned into eight local patches ($N_{\text{patch}} = 8$) surrounding the source region based on the blueprint $\hat{x}^g$ generated in Stage 1. These include four along the cardinal directions and four at the corners, following the same scheme as ProOut [38]. A fixed overlap margin is established between adjacent patches to prevent boundary discontinuities. Unlike ProOut, which generates these eight patches sequentially, our method generates all patches in parallel. Thanks to the global guidance injection and initial noise construction mechanisms described in Sec. 3.4, each patch can be denoised in parallel rather than sequentially. At each denoising step, local patches are denoised independently and then aggregated for overlap blending before being redistributed for the next step. In multi-GPU inference, patches can be processed on different GPUs.

Overlap Blending. For the overlap regions between adjacent patches, we blend them following [2]. Specifically, at each denoising step t , we average the latents from all patches covering a shared spatial position r on the canvas:

$$\hat{z}_t(r) = \frac{\sum_{i \in \mathcal{R}(r)} z_t^{(i)}(r)}{|\mathcal{R}(r)|}, \quad (13)$$

where $\hat{z}_t(r)$ is the blended latent, $\mathcal{R}(r)$ the set of patches, and $z_t^{(i)}(r)$ the latent value of patch i . This averaging is applied only to positions r in the target region. Positions in the source region are excluded from blending, since averaging source-region latents across patches can produce blurry results.

Canvas Compositing. While the overlap blending ensures continuity in the latent space during denoising, the final canvas is composited in the decoded pixel space. Following ProOut [38], the eight patches are composited in a predefined order using a first-win strategy where early placed patches are preserved without being overwritten. Finally, the unmasked source region of $\hat{x}^g$ is pasted on top of the generated canvas, to ensure the original content remains unmodified.

4 Experiments

4.1 Implementation Details

Datasets. We train on 16,000 HumanArt [18] images, 128,951 WikiArt [40] images, and 413,620 LAION-5B High-Resolution [37] images. For evaluation, we use 6,528 IconArt [12] images. Following ProOut [38], we apply random masking for augmentation. Image captions are generated via BLIP [22] for the test data and BLIP2 [21] for the training data, with style descriptions (e.g., "a painting

of") filtered out as in StyleCrafter [29]. The resulting captions are used as the global prompt c_G . For the local prompt c_L in Stage 2, we use quality-related keywords (e.g., "high quality") rather than content descriptions, as each patch covers only a partial region of the canvas. Following InstanceDiffusion [44], object bounding boxes and descriptions are annotated using Grounded-SAM [20,31].

Evaluation Setup. We compare against PQDiff [58], SD Inpainting [35], PowerPaint [64], and ProOut [38]. SD Inpainting, PowerPaint, and ProOut perform progressive outpainting, and the original source image is directly pasted onto the source region of the final output. For ProOut, we adopt the ControlNet-based variant, as no official code is available. For PQDiff ($128 \times 128 \rightarrow 192 \times 192$), we resize and apply an SR model [16] to match the original resolution. We denote the variant that pastes the source back as 'copy' and the one without as 'gen'. Following ProOut, we define the masking ratio $\rho \in (0, 1)$ as the proportion of each spatial dimension to be removed from the original image. Given an original image of size $H \times W$, we crop $\frac{\rho H}{2}$ from the top and bottom and $\frac{\rho W}{2}$ from the left and right to obtain a centered source region of size $(1-\rho)H \times (1-\rho)W$, which is then outpainted back to the original $H \times W$. The default evaluation uses $\rho = 0.333$, yielding approximately 225% total pixels relative to the source.

Training and Inference. The U-Net backbones for Stage 1 and Stage 2 are initialized with pre-trained SD v1.5 Inpainting [35] and frozen; additional modules are initialized with pre-trained InstanceDiffusion weights and fine-tuned. Both Stage 1 (global blueprints) and Stage 2 (local patches) operate at 512×512 . For baselines, PQDiff is trained from scratch for 150 epochs, PowerPaint for 40 epochs, and ProOut for 50 epochs; SD Inpainting uses pre-trained weights without fine-tuning. Our Stage 1 and Stage 2 are trained for 60 and 50 epochs using AdamW [32] ($\text{lr}=5 \times 10^{-5}$, batch size=512) on two H100 and six A6000 GPUs. Inference uses DDIM [39] (30 steps, CFG=3.0), and attention-guided noise optimization is applied following the default settings of AlignNoise [23]. For all baselines, we adopt the default hyperparameters for each method.

4.2 Comparisons

Quantitative Results. Table 1 presents quantitative comparisons in terms of image quality, layout accuracy, and inference speed. For further analysis, we additionally evaluate a variant, denoted as Ours*, which extracts Stage 1 guidance features after the cross-attention layer where text conditions are injected. This allows the guidance feature passed to Stage 2 to better reflect the text condition. For image quality, we adopt FID [14], patch-FID (pFID₂₅₆, pFID₅₁₂) [3,5], CLIP text similarity score (CLIP-S) [34], and CLIP aesthetic score (CLIP-A) [37]. Our default model achieves strong FID and pFID scores, while Ours* further improves FID, pFID, CLIP-S, and CLIP-A. For layout accuracy, we compute AP and IoU between the conditioned bounding boxes and those detected by Grounding DINO [31] in the generated images, with GT detection as the upper bound. As ours is the only framework that natively supports layout-conditioned artwork outpainting, it significantly outperforms all baselines that rely solely

Table 1: Quantitative Results. Comparison with state-of-the-art outpainting methods across image quality, layout accuracy, and inference time. Our default model achieves strong FID and layout accuracy, while Ours* further improves image quality metrics. Ours* denotes our method with guidance features extracted after the cross-attention layer in the Stage 1 U-Net, where the text condition is injected. This allows the guidance features used in Stage 2 to better reflect the text condition. Using 8 GPUs, our method achieves $2.4\times$ faster inference than ProOut. **Bold** and underline denote the best and second-best performance, respectively.

Methods	Image Quality					Layout Acc.		Time(s)↓
	FID↓	pFID ₂₅₆ ↓	pFID ₅₁₂ ↓	CLIP-S↑	CLIP-A↑	AP↑	IoU↑	μ
GT	—	—	—	—	—	0.5903	0.7716	—
PQDiff (gen) [58]	30.8942	54.0282	38.0268	0.1990	3.8542	—	—	—
PQDiff (copy) [58]	25.3032	41.1501	31.8283	0.2031	5.0144	0.2903	0.4304	—
SD Inpainting [35]	10.7598	9.5336	7.6219	0.2007	6.7041	0.3526	0.5539	<u>13.43</u>
PowerPaint [64]	10.5374	9.4403	7.4834	0.2003	6.4423	0.3537	0.5503	13.99
ProOut [38]	10.3032	9.2854	7.2515	<u>0.2052</u>	<u>6.8585</u>	0.3595	0.5623	18.31
Ours	<u>9.3064</u>	<u>9.0635</u>	<u>6.7514</u>	0.2033	6.7949	<u>0.4336</u>	0.6382	7.50
Ours*	9.0729	9.0577	6.6609	0.2055	6.8766	0.4363	<u>0.6361</u>	7.50

Table 2: Quantitative Results on Expansion Ratios. Robustness comparison with state-of-the-art outpainting methods across varying expansion ratios from 200% to 600%, evaluated on landscape artwork images. Our method consistently achieves the best FID at all ratios, with competitive pFID and CLIP-based metrics, demonstrating stable visual quality even under extreme expansion. **Bold** denotes the best results.

Methods	200% Ratio					400% Ratio					600% Ratio				
	FID↓	pFID ₂₅₆	pFID ₅₁₂	CLIP-S↑	CLIP-A↑	FID↓	pFID ₂₅₆	pFID ₅₁₂	CLIP-S↑	CLIP-A↑	FID↓	pFID ₂₅₆	pFID ₅₁₂	CLIP-S↑	CLIP-A↑
SD Inpainting [35]	64.62	99.32	86.55	0.193	6.39	97.95	113.58	97.78	0.176	6.07	111.40	119.42	104.35	0.168	6.00
PowerPaint [64]	52.90	97.41	85.52	0.200	6.21	86.36	104.11	93.74	0.189	5.42	101.34	108.00	99.85	0.183	4.95
ProOut [38]	56.22	97.70	85.66	0.202	6.64	81.17	106.21	91.07	0.202	6.52	89.84	110.49	97.66	0.202	6.41
Ours	44.87	96.43	85.57	0.202	6.76	71.40	101.69	91.88	0.200	6.53	81.86	104.81	95.12	0.198	6.31

on text prompts. Note that all methods benefit from pasting the source region, which contributes to the relatively moderate gap in layout accuracy scores. For inference speed, we report the mean generation time over 100 random samples. Using 8 GPUs, our method achieves an inference time of 7.50s per image, corresponding to a $2.4\times$ reduction compared with ProOut. Under the matched single-GPU setting, our method takes 15.73s per image, which is faster than ProOut at 18.31s but slower than SD Inpainting and PowerPaint due to the additional global planning stage and guidance injection. Table 2 analyzes performance stability under increasing masking ratios from 200% to 600%, evaluated on 400 landscape artwork images collected from the Cleveland Museum of Art [8] and the Art Institute of Chicago [1]. The results demonstrate that our blueprint-guided strategy effectively prevents structural degradation, maintaining stable visual quality even at extreme masking ratios. While our default evaluation uses a 3×3 patch grid, Stage 2 can be reconfigured to accommodate arbitrary canvas




Fig. 4: Qualitative Results. The red dashed box in the GT indicates the source region. The area outside the box is masked during inference and serves as the target region to be generated. PQDiff produces blurry results and SD-Inpainting creates visible seams, while PowerPaint and ProOut often introduce artificial frames or out-of-context. In contrast, our method synthesizes seamless, high-fidelity extensions that faithfully preserve the original style and global structure.

sizes. Additional results in the supplementary material validate the generalizability of our framework across arbitrary canvas layouts (Suppl. Table 3), non-artwork domains (Suppl. Table 4), and DiT-based backbones (Suppl. Table 5).

Qualitative Results. Figure 4 shows the qualitative results of our method against all baselines. We observe that existing baselines tend to generate meaningless repetitive patterns, such as frame-like textures, or produce content that lacks semantic coherence with the overall context of the source images. Conversely, our method maintains structural and semantic consistency with the source, yielding the most visually plausible results among all compared methods. The supplementary material further includes a qualitative comparison with Nano Banana Pro [13] under a setting with an arbitrary expansion ratio (Suppl. Figure 14), where our method shows more reliable layout adherence.

Figure 5 presents the provided layout conditions alongside the corresponding generation results. Existing baseline models, such as SD Inpainting, PowerPaint, and ProOut, struggle with precise spatial control and often fail to generate objects at user-intended locations because they lack dedicated conditioning mechanisms beyond text prompts. This limitation is clearly illustrated in the second row featuring the crown, where baselines either omit the object or misplace it;



Fig. 5: Qualitative Results on Layout-guided Generation. Bounding boxes and object descriptions are provided as layout conditions to guide the generation. Baselines rely solely on text prompts and thus struggle to control spatial placement, whereas our method accurately generates specified objects at designated locations.

in contrast, our method accurately synthesizes the object at the correct scale within the designated bounding box. These results confirm that text prompts alone are insufficient for fine-grained spatial control, highlighting the necessity of explicitly incorporating layout conditions for sophisticated scene composition.

4.3 Ablation Study

We analyze the components used to transfer the global blueprint to Stage 2 synthesis: global guidance feature injection, the patch position token, and forward diffusion-based initial noise construction. As shown in Table 3, using SD Inpainting with only latent blending performs poorly, indicating that blending alone is insufficient for high-resolution outpainting. This shows that parallel synthesis requires more than boundary smoothing, as each local patch needs to be coordinated with the Stage 1 global plan. Removing the guidance features causes a large performance drop, showing that the global guidance features provide structural context for local patch synthesis. Removing the patch position token also degrades performance significantly, indicating that explicit spatial encoding is necessary for coordinating each patch with the global blueprint. Removing forward diffusion further reduces performance, showing that blueprint-based initialization provides complementary structural alignment before denoising begins. We compare Stage 1+SR baselines in the supplementary material (Suppl. Table 2), showing that direct super-resolution of the Stage 1 blueprint does not replace Stage 2 synthesis. Unlike cascaded diffusion pipelines [15], our Stage 1 blueprint guides what should appear where during Stage 2 synthesis rather than serving as a low-resolution sample to be super-resolved into the final image.

Table 3: Ablation Studies. We evaluate the components used to transfer the global blueprint to parallel Stage 2 synthesis: global guidance feature injection, patch position token, and forward diffusion-based initialization. The first row uses SD Inpainting for parallel synthesis with latent blending only, without using the global blueprint. Removing the guidance features or patch position tokens causes a large performance drop, while the full model achieves the best overall performance.

Methods	FID↓	pFID ₂₅₆ ↓	pFID ₅₁₂ ↓	CLIP-S↑	CLIP-A↑
SD Inpainting (parallel) [35]	37.2675	17.6093	20.3705	0.1932	5.8696
Ours (w/o Guidance Feature)	22.2319	12.6436	11.7879	0.1940	5.8561
Ours (w/o Patch Token)	17.9578	10.7008	8.7402	0.1990	6.2209
Ours (w/o Forward Diffusion)	9.6048	9.1144	6.9442	0.2020	6.7598
Ours	9.3064	9.0635	6.7514	0.2033	6.7949

5 Conclusion

In this paper, we introduce a global blueprint-guided two-stage diffusion framework for high-resolution artwork outpainting. Our approach first generates a layout-conditioned, low-resolution blueprint, which then guides the parallel synthesis of high-resolution local patches. This decoupled design effectively mitigates the error accumulation inherent in progressive methods, enables explicit spatial control via bounding boxes, and accelerates inference by up to $2.4\times$ through multi-GPU parallel synthesis. Extensive experiments on large-scale artwork datasets confirm that our framework achieves superior visual fidelity and structural consistency, while offering enhanced spatial controllability and inference efficiency compared to existing baselines. These results demonstrate that global blueprint-guided planning offers a promising paradigm for scalable and controllable high-resolution artwork outpainting.

Acknowledgements

This work was partly supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. RS-2025-25422680, Metacognitive AGI Framework and its Applications, 25%), Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No. RS-2020-II201373, Artificial Intelligence Graduate School Program (Hanyang University), 25%), Institute of Information & communications Technology Planning & Evaluation (IITP) under the Leading Generative AI Human Resources Development grant funded by the Korea government (MSIT) (No. IITP-2026-RS-2026-25544647, 25%), and the AI Computing Infrastructure Enhancement (GPU Rental Support) User Support Program funded by the Ministry of Science and ICT (MSIT), Republic of Korea (No. RQT-25-120160, 25%).

References

1. Art Institute of Chicago: Open access collection. <https://www.artic.edu/collection>, accessed: 2026-02-11
2. Bar-Tal, O., Yariv, L., Lipman, Y., Dekel, T.: Multidiffusion: Fusing diffusion paths for controlled image generation. In: International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 202, pp. 1737–1752. PMLR (2023)
3. Chai, L., Gharbi, M., Shechtman, E., Isola, P., Zhang, R.: Any-resolution training for high-resolution image synthesis. In: European Conference on Computer Vision. Lecture Notes in Computer Science, vol. 13676, pp. 170–188. Springer (2022). https://doi.org/10.1007/978-3-031-19787-1_10
4. Chen, Q., Ma, Y., Wang, H., Yuan, J., Zhao, W., Tian, Q., Wang, H., Min, S., Chen, Q., Liu, W.: Infinite-Canvas: Higher-resolution video outpainting with extensive content generation. In: Proceedings of the AAAI Conference on Artificial Intelligence. pp. 2150–2158 (2025). <https://doi.org/10.1609/AAAI.V39I2.32213>
5. Chen, Y., Wang, O., Zhang, R., Shechtman, E., Wang, X., Gharbi, M.: Image neural field diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8007–8017 (2024). <https://doi.org/10.1109/CVPR52733.2024.00765>
6. Cheng, Y.C., Lin, C.H., Lee, H.Y., Ren, J., Tulyakov, S., Yang, M.H.: Inout: Diverse image outpainting via gan inversion. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11421–11430 (2022). <https://doi.org/10.1109/CVPR52688.2022.01114>
7. Choi, Y., Kwak, S., Lee, K., Choi, H., Shin, J.: Improving diffusion models for authentic virtual try-on in the wild. In: European Conference on Computer Vision. Lecture Notes in Computer Science, vol. 15144, pp. 206–235. Springer (2024). https://doi.org/10.1007/978-3-031-73016-0_13
8. Cleveland Museum of Art: Open access collection. <https://github.com/ClevelandMuseumArt/openaccess>, accessed: 2026-02-11
9. Exhibition Hub, Fever: Van gogh: The immersive experience. <https://vangoghexpo.com>, accessed: 2026-02-11
10. Fan, F., Guo, C., Gong, L., Wang, B., Ge, T., Jiang, Y., Luo, C., Zhan, J.: Hierarchical masked 3d diffusion model for video outpainting. In: Proceedings of the 31st ACM International Conference on Multimedia. pp. 7890–7900 (2023). <https://doi.org/10.1145/3581783.3612478>
11. Gao, P., Yang, X., Zhang, R., Goulermas, J.Y., Geng, Y., Yan, Y., Huang, K.: Generalized image outpainting with u-transformer. Neural Networks **162**, 1–10 (2023). <https://doi.org/10.1016/J.NEUNET.2023.02.021>
12. Gonthier, N., Gousseau, Y., Ladjal, S., Bonfait, O.: Weakly supervised object detection in artworks. In: Proceedings of the European Conference on Computer Vision (ECCV) Workshops. Lecture Notes in Computer Science, vol. 11130, pp. 692–709. Springer (2018). https://doi.org/10.1007/978-3-030-11012-3_53
13. Google DeepMind: Nano Banana Pro (Gemini 3 Pro Image). <https://deepmind.google/models/gemini-image/pro/>, accessed: 2026-06-29
14. Heusel, M., Ramsauer, H., Unterthiner, T., Nessler, B., Hochreiter, S.: GANs trained by a two time-scale update rule converge to a local nash equilibrium. In: Advances in Neural Information Processing Systems. vol. 30, pp. 6626–6637 (2017)
15. Ho, J., Saharia, C., Chan, W., Fleet, D.J., Norouzi, M., Salimans, T.: Cascaded diffusion models for high fidelity image generation. Journal of Machine Learning Research **23**(47), 1–33 (2022)

16. Hsu, C.C., Lee, C.M., Chou, Y.S.: DRCT: Saving image super-resolution away from information bottleneck. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 6133–6142 (2024). <https://doi.org/10.1109/CVPRW63382.2024.00618>
17. Hu, L.: Animate Anyone: Consistent and controllable image-to-video synthesis for character animation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8153–8163 (2024). <https://doi.org/10.1109/CVPR52733.2024.00779>
18. Ju, X., Zeng, A., Wang, J., Xu, Q., Zhang, L.: Human-Art: A versatile human-centric dataset bridging natural and artificial scenes. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 618–629 (2023). <https://doi.org/10.1109/CVPR52729.2023.00067>
19. Khurana, B., Dash, S.R., Bhatia, A., Mahapatra, A., Singh, H., Kulkarni, K.: Semie: Semantically-aware image extrapolation. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 14880–14889 (2021). <https://doi.org/10.1109/ICCV48922.2021.01463>
20. Kirillov, A., Mintun, E., Ravi, N., Mao, H., Rolland, C., Gustafson, L., Xiao, T., Whitehead, S., Berg, A.C., Lo, W.Y., et al.: Segment anything. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3992–4003 (2023). <https://doi.org/10.1109/ICCV51070.2023.00371>
21. Li, J., Li, D., Savarese, S., Hoi, S.: BLIP-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In: International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 202, pp. 19730–19742. PMLR (2023)
22. Li, J., Li, D., Xiong, C., Hoi, S.: BLIP: bootstrapping language-image pre-training for unified vision-language understanding and generation. In: International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 162, pp. 12888–12900. PMLR (2022)
23. Li, N., Li, Z., Tang, Z., Yu, Y., Zou, L., Li, C.: Bridging the gap: Consistent image outpainting via training-free noise optimization. In: Proceedings of the 33rd ACM International Conference on Multimedia. pp. 9969–9977 (2025). <https://doi.org/10.1145/3746027.3755278>
24. Li, Y., Jiang, L., Yang, M.H.: Controllable and progressive image extrapolation. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 2139–2148 (2021). <https://doi.org/10.1109/WACV48630.2021.00219>
25. Li, Y., Liu, H., Wu, Q., Mu, F., Yang, J., Gao, J., Li, C., Lee, Y.J.: Gligen: Open-set grounded text-to-image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22511–22521 (2023). <https://doi.org/10.1109/CVPR52729.2023.02156>
26. Liang, J., Wu, C., Hu, X., Gan, Z., Wang, J., Wang, L., Liu, Z., Fang, Y., Duan, N.: NUWA-Infinity: Autoregressive over autoregressive generation for infinite visual synthesis. In: Advances in Neural Information Processing Systems. vol. 35 (2022)
27. Lin, H., Pagnucco, M., Song, Y.: Edge guided progressively generative image outpainting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. Lecture Notes in Computer Science, vol. 13683, pp. 806–815. Springer (2021). <https://doi.org/10.1109/CVPRW53098.2021.00090>
28. Lin, S., Liu, B., Li, J., Yang, X.: Common diffusion noise schedules and sample steps are flawed. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 5392–5399 (2024). <https://doi.org/10.1109/WACV57701.2024.00532>

29. Liu, G., Xia, M., Zhang, Y., Chen, H., Xing, J., Wang, X., Yang, Y., Shan, Y.: StyleCrafter: Enhancing stylized text-to-video generation with style adapter. arXiv preprint arXiv:2312.00330 (2023). <https://doi.org/10.48550/arXiv.2312.00330>
30. Liu, K., Zhu, Z., Li, C., Liu, H., Zeng, H., Hou, J.: Prefpaint: Aligning image inpainting diffusion model with human preference. In: Advances in Neural Information Processing Systems. vol. 37, pp. 30554–30589 (2024)
31. Liu, S., Zeng, Z., Ren, T., Li, F., Zhang, H., Yang, J., Jiang, Q., Li, C., Yang, J., Su, H., et al.: Grounding DINO: Marrying DINO with grounded pre-training for open-set object detection. In: European Conference on Computer Vision. Lecture Notes in Computer Science, vol. 15105, pp. 38–55. Springer (2024). https://doi.org/10.1007/978-3-031-72970-6_3
32. Loshchilov, I., Hutter, F.: Decoupled weight decay regularization. In: International Conference on Learning Representations. OpenReview.net (2019)
33. Ma, Y., Ma, J., Zhou, M., Chen, Q., Ge, T., Jiang, Y., Lin, T.: Boosting image outpainting with semantic layout prediction. arXiv preprint arXiv:2110.09267 (2021)
34. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 139, pp. 8748–8763. PMLR (2021)
35. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10674–10685 (2022). <https://doi.org/10.1109/CVPR52688.2022.01042>
36. Sabini, M., Rusak, G.: Painting outside the box: Image outpainting with gans. arXiv preprint arXiv:1808.08483 (2018)
37. Schuhmann, C., Beaumont, R., Vencu, R., Gordon, C., Wightman, R., Cherti, M., Coombes, T., Katta, A., Mullis, C., Wortsman, M., et al.: LAION-5B: An open large-scale dataset for training next generation image-text models. In: Advances in Neural Information Processing Systems. vol. 35, pp. 25278–25294 (2022)
38. Song, D.Y., Yu, J.J., Cho, D.: Progressive artwork outpainting via latent diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 15405–15415 (2025). <https://doi.org/10.1109/ICCV51701.2025.01429>
39. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations. OpenReview.net (2021)
40. Tan, W.R., Chan, C.S., Aguirre, H.E., Tanaka, K.: Improved ArtGAN for conditional synthesis of natural image and artwork. IEEE Transactions on Image Processing **28**(1), 394–409 (2019). <https://doi.org/10.1109/TIP.2018.2866698>
41. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. In: Advances in Neural Information Processing Systems. vol. 33, pp. 7537–7547 (2020)
42. Teterwak, P., Sarna, A., Krishnan, D., Maschinot, A., Belanger, D., Liu, C., Freeman, W.T.: Boundless: Generative adversarial networks for image extension. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10520–10529 (2019). <https://doi.org/10.1109/ICCV.2019.01062>
43. Tsai, S.R., Chang, W.C., Lee, J.Y., Su, C.H., Liu, Y.L.: Lightsout: Diffusion-based outpainting for enhanced lens flare removal. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 6353–6363 (2025). <https://doi.org/10.1109/ICCV51701.2025.00599>

44. Wang, X., Darrell, T., Rambhatla, S.S., Girdhar, R., Misra, I.: Instancediffusion: Instance-level control for image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6232–6242 (2024). <https://doi.org/10.1109/CVPR52733.2024.00596>
45. Wang, Y., Wei, Y., Qian, X., Zhu, L., Yang, Y.: Sketch-guided scenery image outpainting. *IEEE Transactions on Image Processing* **30**, 2643–2655 (2021). <https://doi.org/10.1109/TIP.2021.3054477>
46. Wang, Y., Wei, Y., Qian, X., Zhu, L., Yang, Y.: ReGO: Reference-guided outpainting for scenery image. *IEEE Transactions on Image Processing* **33**, 1375–1388 (2024). <https://doi.org/10.1109/TIP.2024.3357290>
47. Wu, T., Si, C., Jiang, Y., Huang, Z., Liu, Z.: FreeInit: Bridging initialization gap in video diffusion models. In: European Conference on Computer Vision. Lecture Notes in Computer Science, vol. 15061, pp. 378–394. Springer (2024). https://doi.org/10.1007/978-3-031-72646-0_22
48. Xu, Y., Pagnucco, M., Song, Y.: DHG-GAN: Diverse image outpainting via decoupled high frequency semantics. In: Proceedings of the Asian Conference on Computer Vision. Lecture Notes in Computer Science, vol. 13847, pp. 168–184. Springer (2022). https://doi.org/10.1007/978-3-031-26293-7_11
49. Xu, Y., Gu, T., Chen, W., Chen, A.: OOTDiffusion: Outfitting fusion based latent diffusion for controllable virtual try-on. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 39, pp. 8996–9004 (2025). <https://doi.org/10.1609/AAAI.V39I9.32973>
50. Xu, Z., Zhang, J., Liew, J.H., Yan, H., Liu, J.W., Zhang, C., Feng, J., Shou, M.Z.: MagicAnimate: Temporally consistent human image animation using diffusion model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1481–1490 (2024). <https://doi.org/10.1109/CVPR52733.2024.00147>
51. Yan, K., Ji, L., Wu, C., Liang, J., Zhou, M., Duan, N., Ma, S.: HORIZON: High-resolution semantically controlled panorama synthesis. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 38, pp. 6431–6439 (2024). <https://doi.org/10.1609/AAAI.V38I6.28463>
52. Yang, C.A., Tan, C.Y., Fan, W.C., Yang, C.F., Wu, M.L., Wang, Y.C.F.: Scene graph expansion for semantics-guided image outpainting. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 15596–15605 (2022). <https://doi.org/10.1109/CVPR52688.2022.01517>
53. Yang, J., Wang, H., Zhu, Z., Liu, C., Wu, M.W., Sun, M.: VIP: Versatile image outpainting empowered by multimodal large language model. In: Proceedings of the Asian Conference on Computer Vision. Lecture Notes in Computer Science, vol. 15476, pp. 59–76. Springer (2024). https://doi.org/10.1007/978-981-96-0917-8_4
54. Yang, Z., Dong, J., Liu, P., Yang, Y., Yan, S.: Very long natural scenery image prediction by outpainting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 10560–10569 (2019). <https://doi.org/10.1109/ICCV.2019.01066>
55. Yao, K., Gao, P., Yang, X., Sun, J., Zhang, R., Huang, K.: Outpainting by queries. In: European Conference on Computer Vision. pp. 153–169. Springer (2022). https://doi.org/10.1007/978-3-031-20050-2_10
56. Ye, H., Zhang, J., Liu, S., Han, X., Yang, W.: IP-Adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arXiv:2308.06721* (2023). <https://doi.org/10.48550/arXiv.2308.06721>

57. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3813–3824 (2023). <https://doi.org/10.1109/ICCV51070.2023.00355>
58. Zhang, S., Huang, J., Zhou, Q., Wang, Z., Wang, F., Luo, J., Yan, J.: Continuous-multiple image outpainting in one-step via positional query and a diffusion-based approach. In: International Conference on Learning Representations. OpenReview.net (2024)
59. Zhang, S., Zhou, Q., Wang, Z., Li, H., Yan, J.: Easyoutpainter: One step image outpainting with both continuous multiple and resolution. IEEE Transactions on Pattern Analysis and Machine Intelligence **47**(10), 9217–9231 (2025). <https://doi.org/10.1109/TPAMI.2025.3586824>
60. Zheng, G., Zhou, X., Li, X., Qi, Z., Shan, Y., Li, X.: Layoutdiffusion: Controllable diffusion model for layout-to-image generation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 22490–22499 (2023). <https://doi.org/10.1109/CVPR52729.2023.02154>
61. Zhou, D., Li, Y., Ma, F., Yang, Z., Yang, Y.: Migc++: Advanced multi-instance generation controller for image synthesis. IEEE Transactions on Pattern Analysis and Machine Intelligence **47**(3), 1714–1728 (2025). <https://doi.org/10.1109/TPAMI.2024.3510752>
62. Zhou, D., Li, Y., Ma, F., Zhang, X., Yang, Y.: Migc: Multi-instance generation controller for text-to-image synthesis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6818–6828 (2024). <https://doi.org/10.1109/CVPR52733.2024.00651>
63. Zhu, S., Chen, J.L., Dai, Z., Dong, Z., Xu, Y., Cao, X., Yao, Y., Zhu, H., Zhu, S.: Champ: Controllable and consistent human image animation with 3D parametric guidance. In: European Conference on Computer Vision. Lecture Notes in Computer Science, vol. 15113, pp. 145–162. Springer (2024). https://doi.org/10.1007/978-3-031-73001-6_9
64. Zhuang, J., Zeng, Y., Liu, W., Yuan, C., Chen, K.: A task is worth one word: Learning with task prompts for high-quality versatile image inpainting. In: European Conference on Computer Vision. Lecture Notes in Computer Science, vol. 15116, pp. 195–211. Springer (2024). https://doi.org/10.1007/978-3-031-73636-0_12