

MegaFusion: Extend Diffusion Models towards Higher-resolution Image Generation without Further Tuning

Haoning Wu*, Shaocheng Shen*, Qiang Hu†, Xiaoyun Zhang†, Ya Zhang, Yanfeng Wang
Shanghai Jiao Tong University, China



Figure 1. **Overview.** *Left:* Existing diffusion-based text-to-image models fall short in synthesizing higher-resolution images due to the fixed image resolution during training, resulting in a noticeable decline in image quality and semantic deviation. *Right:* Our proposed tuning-free **MegaFusion** can effectively and efficiently extend diffusion models (e.g. SDM [38], SDXL [33] and Floyd [8]) towards generating images at higher resolutions (e.g., 1024×1024 , 1920×1080 , 2048×1536 , and 2048×2048) of arbitrary aspect ratios (e.g., 1 : 1, 16 : 9, and 4 : 3). We recommend the reader to zoom in for the visualization results.

Abstract

Diffusion models have emerged as frontrunners in text-to-image generation, but their fixed image resolution during training often leads to challenges in high-resolution image generation, such as semantic deviations and object replication. This paper introduces **MegaFusion**, a novel approach that extends existing diffusion-based text-to-image models towards efficient higher-resolution generation without additional fine-tuning or adaptation. Specifically, we employ an innovative **truncate and relay** strategy to bridge the denoising processes across different resolutions, allowing for high-resolution image generation in a coarse-to-fine manner. Moreover, by integrating dilated convolutions and noise re-scheduling, we further adapt the model’s priors for higher resolution. The versatility and efficacy of MegaFu-

sion make it universally applicable to both latent-space and pixel-space diffusion models, along with other derivative models. Extensive experiments confirm that MegaFusion significantly boosts the capability of existing models to produce images of megapixels and various aspect ratios, while only requiring about 40% of the original computational cost. Code is available at <https://haoningwu3639.github.io/MegaFusion/>.

1. Introduction

Diffusion models have demonstrated unparalleled performance across broad applications such as text-to-image generation [8, 10, 18–20, 38, 43], image editing [4, 5, 16, 25, 29, 31, 50], consistent image sequence generation [28, 30, 32], and even achieves promising results in challenging text-to-video generation [18, 21, 42, 51]. Among them, Stable Diffusion (also known as Latent Diffusion [38]) performs denoising in a compressed latent space, and

*: These authors contribute equally to this work.

†: Corresponding author.

has showcased impressive generative capabilities after pre-training on large-scale paired datasets [41]. In comparison, Imagen [40] and Floyd [8] adopt cascading diffusion models in pixel space, initiating with low-resolution image synthesis followed by successive super-resolution stages.

Despite these advancements, as depicted in Figure 1 *Left*, these models face a major challenge: they struggle to generate images beyond training resolutions, leading to semantic deviation and degraded image quality. Existing solutions often require additional tuning or are limited to specific models. For example, MultiDiffusion [3] and ElasticDiffusion [14] adopt post-processing optimization to stitch high-resolution panoramas, which is inefficient and time-consuming. Relay Diffusion [44] employs blurring diffusion in pixel space, yet it necessitates training multiple specific diffusion models from scratch. ResAdapter [7] and CheapScaling [13] involve minimal extra training through LoRA [23] or Upsamplers, but still incur a notable training overhead. Meanwhile, tuning-free alternatives like ScaleCrafter [15] and FouriScale [24] adapt pre-trained SDMs for higher resolutions, but demand meticulous hyperparameter adjustment and are restricted to latent-space models.

To tackle these limitations, we introduce **MegaFusion**, a tuning-free method to extend existing diffusion models towards generating higher-resolution and variable aspect ratio images with megapixels. Concretely, we begin with a *truncate and relay* strategy, which seamlessly bridges the synthesis of different resolution images, enabling efficient generation in a coarse-to-fine manner with only 40% of the original computational cost. Moreover, it is orthogonally compatible with existing techniques such as *dilated convolutions* [54] and *noise re-scheduling* for better image quality. The versatility of **MegaFusion** makes it applicable to both latent-space and pixel-space diffusion models, as well as other diffusion-based frameworks with extra conditions, such as IP-Adapter [53] and ControlNet [56]. As shown in Figure 1 *Right*, MegaFusion significantly improves the ability of diffusion models to synthesize higher-resolution images with accurate semantics and superior quality.

To summarize, our contributions are fourfold: (i) we propose **MegaFusion**, a tuning-free approach utilizing a *truncate and relay* strategy to efficiently generate high-quality, high-resolution images with megapixels in a coarse-to-fine manner; (ii) we incorporate *dilated convolution* and *noise re-scheduling* techniques to further refine the adaptability of pre-trained diffusion models for higher resolution; (iii) we demonstrate the applicability of our method across both latent-space and pixel-space diffusion models, as well as their extensions, synthesizing high-resolution images with various aspect ratios at roughly 40% of the original computational cost; (iv) we conduct extensive experiments validating the superiority of our proposed method, in terms of efficiency, image quality, and semantic accuracy.

2. Related Works

Diffusion Models. As a part of probabilistic generative models, diffusion models typically learn to generate samples by iterative denoising. DDPM [19] has first showcased remarkable performance, while DDIM [43] significantly improves sampling efficiency. Leveraging their excellent generative capabilities, diffusion models have been applied to diverse fields, including image-to-image translation [4, 16, 25, 31, 50] and video generation [18, 21, 42, 51].

Text-to-Image Generation. Generative models have been widely adopted for the challenging text-to-image generation task, with GAN [11, 52, 55] as the pioneers. Meanwhile, auto-regressive transformers like DALL-E [37] further push the boundaries. Diffusion models, such as DALL-E 2 [36], Imagen [40] and Floyd [8], have recently risen to prominence. Notably, Stable Diffusion (Latent Diffusion [38]), performing denoising in latent space, has demonstrated outstanding performance, thereby being widely applied within the research community. Additionally, SDXL [33] further elevates the generative performance of Stable Diffusion with a diffusion refiner and an extra text encoder.

Our **MegaFusion**, is designed for seamless integration with diffusion models across both latent and pixel spaces, extending their capacity for higher-resolution generation.

Higher-resolution Generation. Existing diffusion models are typically limited to fixed resolutions and aspect ratios, struggling to produce images beyond their training resolutions. MultiDiffusion [3] and ElasticDiffusion [14] address this by synthesizing overlapping crops and merging them into panoramic images, which requires a time-consuming inference procedure. Relay Diffusion [44] designs a pixel-space model with blurring diffusion to craft high-resolution images, at the cost of retraining several models from scratch. ScaleCrafter [15] achieves high-resolution generation by enlarging receptive fields of Stable Diffusion with dispersed convolution without extra fine-tuning. DemoFusion [9] attempts to connect multiple resolutions for coarse-to-fine generation, but demands repeating generation multiple times, leading to low efficiency.

Several concurrent works also express rich interest in this task: ResAdapter [7] and CheapScaling [13] can generate images with unrestricted resolutions and aspect ratios with minimal tuning via trainable LoRA adapters or Upsamplers. FouriScale [24] and HiDiffusion [57] offer a training-free strategy but are still limited to SDM-based models.

In contrast to the aforementioned methods, which either require further training or are limited to specific models, our proposed **MegaFusion** emerges as a versatile tuning-free solution that can be integrated seamlessly into existing diffusion models, enabling the synthesis of higher-resolution images of various aspect ratios.

3. Preliminary

In this section, we briefly introduce diffusion models, including their forward and backward processes in Sec. 3.1; and the latent diffusion models (LDMs) that perform diffusion in latent space to improve efficiency in Sec. 3.2.

3.1. Diffusion Models

Diffusion models, a class of deep generative models, iteratively transform Gaussian noise into structured data samples through a denoising process. Concretely, diffusion models comprise a forward diffusion process that progressively adds Gaussian noise to an image \mathbf{x}_0 via a Markov process over T steps. Let \mathbf{x}_t represent the noisy image at step t , with the transition from \mathbf{x}_{t-1} to \mathbf{x}_t being modeled by $q(\mathbf{x}_t|\mathbf{x}_{t-1}) = \mathcal{N}(\mathbf{x}_t; \sqrt{1-\beta_t}\mathbf{x}_{t-1}, \beta_t\mathbf{I})$. Here, $\beta_t \in (0, 1)$ are pre-determined hyperparameters controlling the variance introduced at each step. By defining $\alpha_t = 1 - \beta_t$ and $\bar{\alpha}_t = \prod_{i=1}^t \alpha_i$, we can leverage the properties of Gaussian distributions and the reparameterization trick to reformulate the relationship as: $q(\mathbf{x}_t|\mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \sqrt{\bar{\alpha}_t}\mathbf{x}_0, (1 - \bar{\alpha}_t)\mathbf{I})$. This insight allows us to succinctly express the forward process with Gaussian noise ϵ as: $\mathbf{x}_t = \sqrt{\bar{\alpha}_t}\mathbf{x}_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon$.

Diffusion models also encompass a reverse diffusion process to reconstruct images from noise. This process, denoted as p_θ , usually leverages a UNet-based [39] model to estimate the noise term ϵ_θ , represented as: $p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t) = \mathcal{N}(\mathbf{x}_t; \mu_\theta(\mathbf{x}_t, t), \Sigma_\theta(\mathbf{x}_t, t))$. Here, μ_θ is the predicted mean of Gaussian distribution, expressed in terms of the estimated noise ϵ_θ as: $\mu_\theta(\mathbf{x}_t, t) = \frac{1}{\sqrt{\alpha_t}}(\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}}\epsilon_\theta(\mathbf{x}_t, t))$

3.2. Latent Diffusion Models

To improve efficiency and reduce computational cost, Latent Diffusion (LDMs) execute diffusion and denoising within a learned low-dimensional latent space of a pre-trained Variational Autoencoder (VAE). Specifically, the VAE encoder \mathcal{E} maps an image $\mathbf{x}_0 \in \mathbb{R}^{3 \times H \times W}$ to a latent representation $\mathbf{z}_0 \in \mathbb{R}^{4 \times h \times w}$ via $\mathbf{z}_0 = \mathcal{E}(\mathbf{x}_0)$. Afterwards, the decoder \mathcal{D} reconstructs the original image \mathbf{x}_0 from \mathbf{z}_0 , represented as $\hat{\mathbf{x}}_0 = \mathcal{D}(\mathbf{z}_0) \approx \mathbf{x}_0$.

This setup allows the diffusion process to be conducted in a compact latent space, facilitating efficient image synthesis. During inference, LDM samples latent codes from a conditional distribution $p(\mathbf{z}_0|c)$, where c represents the conditional information such as text embedding from CLIP [34] or T5 [35] text encoder. This process can be formalized as: $p_\theta(\mathbf{z}_{t-1}|\mathbf{z}_t, c) = \mathcal{N}(\mathbf{z}_t; \mu_\theta(\mathbf{z}_t, t, c), \Sigma_\theta(\mathbf{z}_t, t, c))$.

4. Method

This section initiates with elaborating on the *truncate* and *relay* strategy within our proposed tuning-free **MegaFusion** in Sec. 4.1. Then, we incorporate *dilated convolution* and *noise re-scheduling* to further adapt model prior to-

wards higher resolution in Sec. 4.2. Lastly, we detail the application of our method across latent-space and pixel-space diffusion models, as well as their extensions, in Sec. 4.3.

4.1. Truncate and Relay Strategy

High-level Idea. As evidenced by eDiff-I [1], diffusion models synthesize semantics during early denoising steps and texture details in later steps. Our intuition and insight here are that: we should perform early-stage denoising at original inference resolutions to guarantee accurate semantics, followed by *truncating* and *relaying* at higher resolutions to continue later-stage denoising to produce texture details. This strategy extends the higher-resolution generation capabilities of pre-trained models, enabling the synthesis of high-quality images with precise semantics at low computational costs, and supports various aspect ratios.

Problem Setting. For clarity, we focus on latent-space diffusion models as an example. As for pixel-space models, our method can be applied more straightforwardly and conveniently. Given that our strategy is inherently tuning-free, we focus on the inference stage herein. Using a pre-trained Latent Diffusion (LDM) with a denoiser ϵ_θ , a low-resolution latent code $\mathbf{z}_0^1 \in \mathbb{R}^{4 \times h_1 \times w_1}$ can be synthesized within T denoising steps conditioned on a text prompt c_T , and then decoded into an image $\mathbf{x}_0^1 \in \mathbb{R}^{3 \times H_1 \times W_1}$ by the VAE decoder \mathcal{D} . Our goal is to generate a higher-resolution image $\mathbf{x}_0^k \in \mathbb{R}^{3 \times H_k \times W_k}$ alongside its latent code $\mathbf{z}_0^k \in \mathbb{R}^{4 \times h_k \times w_k}$, by linking generation processes across different resolutions over a total of T steps, where $T = \sum_{i=1}^k T_i$.

Truncate. To guarantee accurate semantics, we begin with a low-resolution generation through T_1 steps denoising:

$$\mathbf{z}_{t-1} = \frac{1}{\sqrt{\alpha_t}}(\mathbf{z}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}}\epsilon_\theta(\mathbf{z}_t, t, c_T)) + \sigma_t\epsilon, \quad (1)$$

where $t = T, T-1, \dots, T-T_1+1$

where σ_t are pre-calculated coefficients and ϵ denotes noise sampled from a standard Gaussian distribution. Subsequently, at step $t_1 = T - T_1 + 1$, we *truncate* the generation process and compute the approximate clean latent code $\hat{\mathbf{z}}_{t_1}^1 \in \mathbb{R}^{4 \times h_1 \times w_1}$, which serves as a pivotal element for multi-resolutions bridging via:

$$\hat{\mathbf{z}}_{t_1}^1 = \frac{1}{\sqrt{\bar{\alpha}_{t_1}}}(\mathbf{z}_{t_1} - \sqrt{1 - \bar{\alpha}_{t_1}}\epsilon_\theta(\mathbf{z}_{t_1}, t_1, c_T)) \quad (2)$$

Here, $\hat{\mathbf{z}}_{t_1}^1$ is subsequently decoded to an image $\hat{\mathbf{x}}_{t_1}^1 \in \mathbb{R}^{3 \times H_1 \times W_1}$ and upsampled to a higher-resolution relatively clean image $\hat{\mathbf{x}}_{t_1}^2 \in \mathbb{R}^{3 \times H_2 \times W_2}$ utilizing a non-parametric Upsampler, Φ , represented as:

$$\hat{\mathbf{x}}_{t_1}^1 = \mathcal{D}(\hat{\mathbf{z}}_{t_1}^1), \quad \hat{\mathbf{x}}_{t_1}^2 = \Phi(\hat{\mathbf{x}}_{t_1}^1) \quad (3)$$

Relay. To further enhance higher-resolution texture details, the upsampled image $\hat{\mathbf{x}}_{t_1}^2$ is then re-encoded into latent

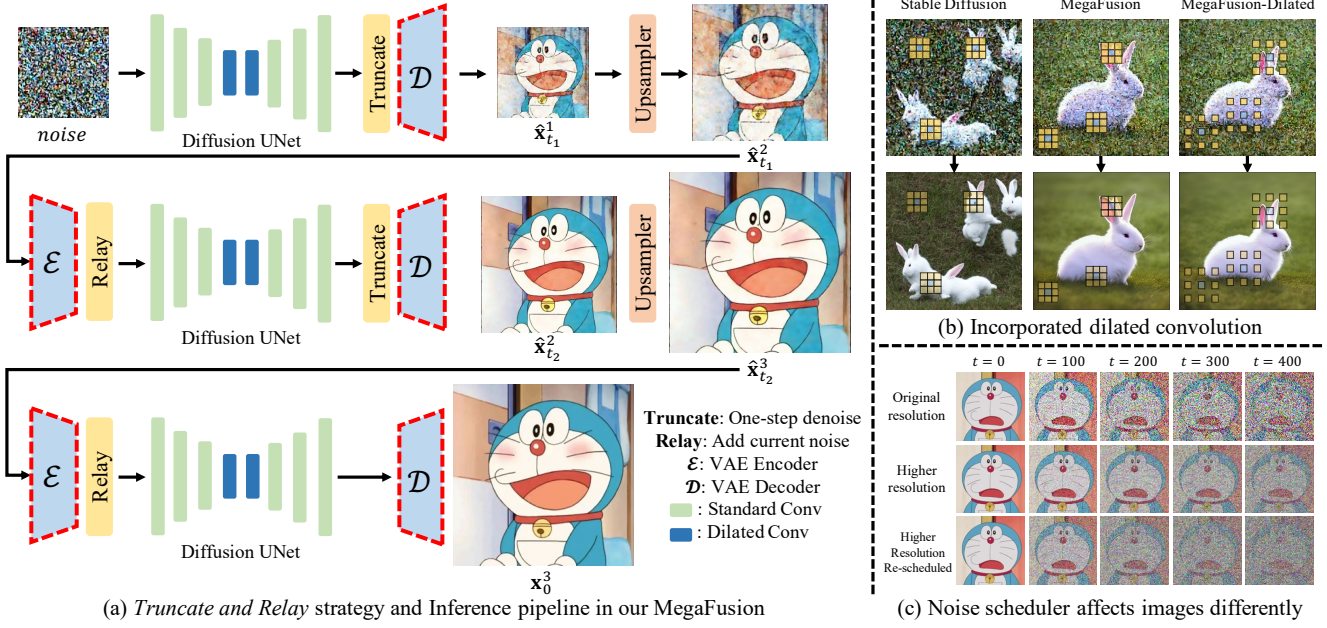


Figure 2. **Architecture Overview.** (a) The *Truncate and Relay* strategy in **MegaFusion** seamlessly connects generation processes across different resolutions to produce higher-resolution images without extra tuning, exemplified by a three-stage pipeline. For pixel-space models, the VAE encoder and decoder can be directly removed. (b) Limited receptive fields lead to quality decline and object replication. *Dilated convolutions* expand the receptive field at higher resolutions, enabling the model to capture more global information for more accurate semantics and image details. (c) Noise at identical timesteps affects images of different resolutions differently, deviating from the model’s prior. *Noise re-scheduling* helps align the noise level of higher-resolution images with that of the original resolution.

code $\hat{\mathbf{z}}_{t_1}^2 \in \mathbb{R}^{4 \times h_2 \times w_2}$ via the VAE encoder \mathcal{E} and perturbed with noise at the current step t_1 to **relay** the generation:

$$\hat{\mathbf{z}}_{t_1}^2 = \mathcal{E}(\hat{\mathbf{x}}_{t_1}^2), \quad \mathbf{z}_{t_1}^2 = \mathcal{N}(\hat{\mathbf{z}}_{t_1}^2; \sqrt{\bar{\alpha}_{t_1}} \hat{\mathbf{z}}_{t_1}^2, (1 - \bar{\alpha}_{t_1}) \mathbf{I}) \quad (4)$$

The generation process continues at a higher resolution, by re-leveraging Equation 1 for T_2 steps of denoising, sequentially navigating through $t = T - T_1, T - T_1 - 1, \dots, T - T_1 - T_2 + 1$. Subsequently, the *truncate and relay* operations can be then conducted at step $t_2 = T - T_1 - T_2 + 1$.

As depicted in Figure 2 (a), this iterative process is repeated multiple times until the generation of a high-resolution latent code \mathbf{z}_0^k , which can be then decoded into a corresponding high-resolution image \mathbf{x}_0^k with megapixels.

4.2. MegaFusion++

Our MegaFusion, based on the *truncate and relay* strategy, can be further combined orthogonally with existing techniques such as *dilated convolution* and *noise rescheduling*, to adapt the model priors to higher resolutions.

Dilated Convolution. Blurriness and semantic deviation in high-resolution images generated by diffusion models often stem from the constrained receptive field of UNet layers trained on fixed-resolution data, lacking comprehensive global context. As illustrated in Figure 2 (b), existing models trained on low-resolution images tend to synthesize multiple rabbits in different local regions due to insufficient re-

ceptive fields, leading to semantic inaccuracies. Inspired by ScaleCrafter [15], we modify the convolutional kernels of the UNet-based denoiser ϵ_θ to incorporate dilated convolutions [54] with a specific dilation rate δ . This broadens the receptive field without additional tuning, allowing for better incorporation of global information.

For the sake of simplicity, we omit channel dimensions and convolution biases here, focusing on modifying weight parameters to transform standard convolutions into dilated ones. Given a feature map $F \in \mathbb{R}^{m \times n}$ and a convolutional kernel $k \in \mathbb{R}^{r \times r}$, the standard convolution can be represented as: $(F * k)(p) = \sum_{s+t=p} F(s) \cdot k(t)$. In contrast, the corresponding dilated convolution, with dilation rate δ , can be expressed as: $(F *_\delta k)(p) = \sum_{s+\delta t=p} F(s) \cdot k(t)$. Here, p, k , and t denote spatial locations within the feature map and convolution kernel, respectively.

Following previous practices [15], instead of replacing all convolutions with dilated ones, which may lead to catastrophic quality decline, we selectively apply this modification to the middle layers of UNet. **The insight here is that:** we broaden receptive fields in the bottleneck to aggregate global context, while preserving original priors at higher resolution to sample nearby features for enhancing details.

Noise Re-scheduling. Consistent with discoveries in simple diffusion [22] and relay diffusion [44], we observe that identical noise levels impact images differently across var-

Methods	resolution	FID _r ↓	FID _b ↓	KID _r ↓	KID _b ↓	CLIP-T↑	CIDEr↑	Meteor↑	ROUGE↑	GFlops	Inference time
SDM [38]	1024 × 1024	41.35	51.02	0.0086	0.0113	0.3009	17.75	18.38	23.64	135.0K	15.17s
SDM-StableSR [47]	1024 × 1024	25.46	19.61	<u>0.0062</u>	<u>0.0031</u>	0.3117	20.24	20.91	26.28	292.6K	33.48s
SDM-RealESRGAN	1024 × 1024	25.20	19.49	0.0059	0.0032	0.3119	21.35	21.26	26.76	35.6K	5.12s
ResAdapter [7]	1024 × 1024	27.38	20.47	0.0073	0.0033	0.3102	20.99	21.38	27.65	137.5K	16.25s
ScaleCrafter [15]	1024 × 1024	27.97	22.05	0.0076	0.0043	0.3125	20.14	<u>21.65</u>	<u>28.23</u>	135.0K	17.52s
SDM-MegaFusion	1024 × 1024	30.19	<u>10.98</u>	0.0088	0.0034	0.3101	<u>21.14</u>	21.44	27.34	48.2K	7.56s
SDM-MegaFusion++	1024 × 1024	25.14	7.82	0.0064	0.0012	<u>0.3121</u>	20.46	22.18	28.36	48.2K	7.56s
SDXL [33]	2048 × 2048	47.53	47.08	0.0133	0.0139	0.3041	17.55	18.65	25.10	540.2K	79.66s
SDXL-RealESRGAN	2048 × 2048	24.76	13.54	0.0056	<u>0.0021</u>	0.3192	23.27	22.44	<u>28.44</u>	147.1K	22.33s
ScaleCrafter [15]	2048 × 2048	27.46	24.73	0.0064	0.0061	0.3138	19.97	22.34	28.12	540.2K	80.72s
DemoFusion [9]	2048 × 2048	<u>24.61</u>	13.36	0.0066	0.0023	0.3198	22.02	22.86	28.48	1354.9K	217.19s
SDXL-MegaFusion	2048 × 2048	25.12	<u>12.13</u>	<u>0.0059</u>	0.0027	<u>0.3227</u>	23.49	22.65	28.12	216.1K	30.94s
SDXL-MegaFusion++	2048 × 2048	23.86	6.93	0.0056	0.0018	0.3244	<u>23.42</u>	<u>22.74</u>	28.38	216.1K	30.94s
SD3 [10]	2048 × 2048	<u>38.37</u>	<u>31.91</u>	<u>0.0165</u>	<u>0.0181</u>	<u>0.3058</u>	<u>17.89</u>	<u>18.72</u>	<u>24.66</u>	433.9K	64.89s
SD3-MegaFusion	2048 × 2048	28.81	8.93	0.0098	0.0018	0.3178	23.01	22.45	29.14	201.4K	29.07s
Floyd-Stage1 [8]	128 × 128	66.27	81.65	<u>0.0262</u>	0.0454	0.2818	14.69	18.22	25.06	111.7K	77.08s
Floyd-MegaFusion	128 × 128	<u>53.09</u>	39.73	0.0273	0.0334	<u>0.3024</u>	25.01	<u>25.00</u>	<u>31.35</u>	44.9K	32.19s
Floyd-MegaFusion++	128 × 128	43.43	<u>50.08</u>	0.0213	<u>0.0437</u>	0.3046	<u>20.28</u>	25.01	31.64	44.9K	32.19s
Floyd-Stage2 [8]	512 × 512	46.64	38.15	0.0254	0.0166	0.3098	<u>23.85</u>	21.47	26.26	60.7K	48.58s
Floyd-MegaFusion	512 × 512	<u>39.80</u>	<u>24.87</u>	<u>0.0164</u>	<u>0.0078</u>	<u>0.3106</u>	23.22	<u>23.51</u>	<u>29.30</u>	24.3K	21.72s
Floyd-MegaFusion++	512 × 512	26.34	24.55	0.0063	0.0077	0.3110	24.01	23.58	29.52	24.3K	21.72s

Table 1. **Quantitative comparison.** We compare our boosted models on higher-resolution generation with representative latent-space and pixel-space diffusion models on MS-COCO [27] dataset. **RED** represents best performance, and **BLUE** denotes second best performance.

ious resolutions, as illustrated in Figure 2 (c), leading to varying signal-to-noise ratios (SNR) at the same timestep.

According to the SNR definition in previous work [22]: $SNR_t = \frac{(\sqrt{\bar{\alpha}_t})^2}{(\sqrt{1-\bar{\alpha}_t})^2} = \frac{\bar{\alpha}_t}{1-\bar{\alpha}_t}$. Given a low-resolution image $\mathbf{x} \in \mathbb{R}^{3 \times H \times W}$ and a high-resolution one $\mathbf{x}' \in \mathbb{R}^{3 \times H' \times W'}$ with $H' > H$ and $W' > W$, if we downsample \mathbf{x}' to $\mathbf{x}'_{down} \in \mathbb{R}^{3 \times H \times W}$, the SNR at timestep t of \mathbf{x}'_{down} (denoted as $SNR_{down}^{H' \times W'}$) in comparison to \mathbf{x} (represented as $SNR^{H \times W}$) will exhibit the following relationship: $SNR^{H \times W} = \gamma \cdot SNR_{down}^{H' \times W'}$.

Assuming the original noise scheduler at $H \times W$ is denoted as $\bar{\alpha}_t$, the revised scheduler $\bar{\alpha}'_t$ at higher resolution $H' \times W'$ should satisfy: $\frac{\bar{\alpha}_t}{1-\bar{\alpha}_t} = \gamma \cdot \frac{\bar{\alpha}'_t}{1-\bar{\alpha}'_t}$. This yields the relationship: $\bar{\alpha}'_t = \frac{\bar{\alpha}_t}{\gamma - (\gamma-1)\bar{\alpha}_t}$. Incorporating this into the high-resolution noise scheduler initialization gives a new $\bar{\alpha}_t$ sequence. This process, termed noise re-scheduling, adjusts noise levels to better suit higher-resolution image generation, thereby improving synthesis quality and fidelity.

4.3. Further Application on other Models

Pixel-space Diffusion Models. MegaFusion is equally applicable to pixel-space diffusion models, such as Floyd [8], with the primary difference being that the *truncate and relay* operation is performed directly in pixel space. Conse-

quently, Equations 2, 3, and 4 are adapted as follows:

$$\hat{\mathbf{x}}_{t_1}^1 = \frac{1}{\sqrt{\bar{\alpha}_{t_1}}}(\mathbf{x}_{t_1} - \sqrt{1 - \bar{\alpha}_{t_1}}\epsilon_{\theta}(\mathbf{x}_{t_1}, t_1, c_T)) \quad (5)$$

$$\hat{\mathbf{x}}_{t_1}^2 = \Phi(\hat{\mathbf{x}}_{t_1}^1), \quad \mathbf{x}_{t_1}^2 = \mathcal{N}(\mathbf{x}_{t_1}^2; \sqrt{\bar{\alpha}_{t_1}}\hat{\mathbf{x}}_{t_1}^2, (1 - \bar{\alpha}_{t_1})\mathbf{I}) \quad (6)$$

Diffusion Models with Extra Conditions. Our methodology can also extend to diffusion models that incorporate extra input conditions, such as ControlNet [56] and IP-Adapter [53]. These models utilize both text condition c_T and image condition c_I as inputs. Consequently, Equation 1 can be reformulated to accommodate both conditions:

$$\mathbf{z}_{t-1} = \frac{1}{\sqrt{\alpha_t}}(\mathbf{z}_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}}\epsilon(\mathbf{z}_t, t, c_T, c_I)) + \sigma_t \epsilon \quad (7)$$

5. Experiments

In this section, we first describe our experimental settings in Sec. 5.1. Next, we present comparisons to existing models with quantitative metrics and human evaluation in Sec. 5.2. We then showcase qualitative results of applying our method to various diffusion models in Sec. 5.3. Lastly, ablation studies are presented in Sec. 5.4.

5.1. Experiment Settings

Implementation Details. We evaluate text-to-image diffusion models in both latent space (SDM 1.5 [38], SDXL [33] and SD3 [10]) and pixel space (Floyd [8]). All models use

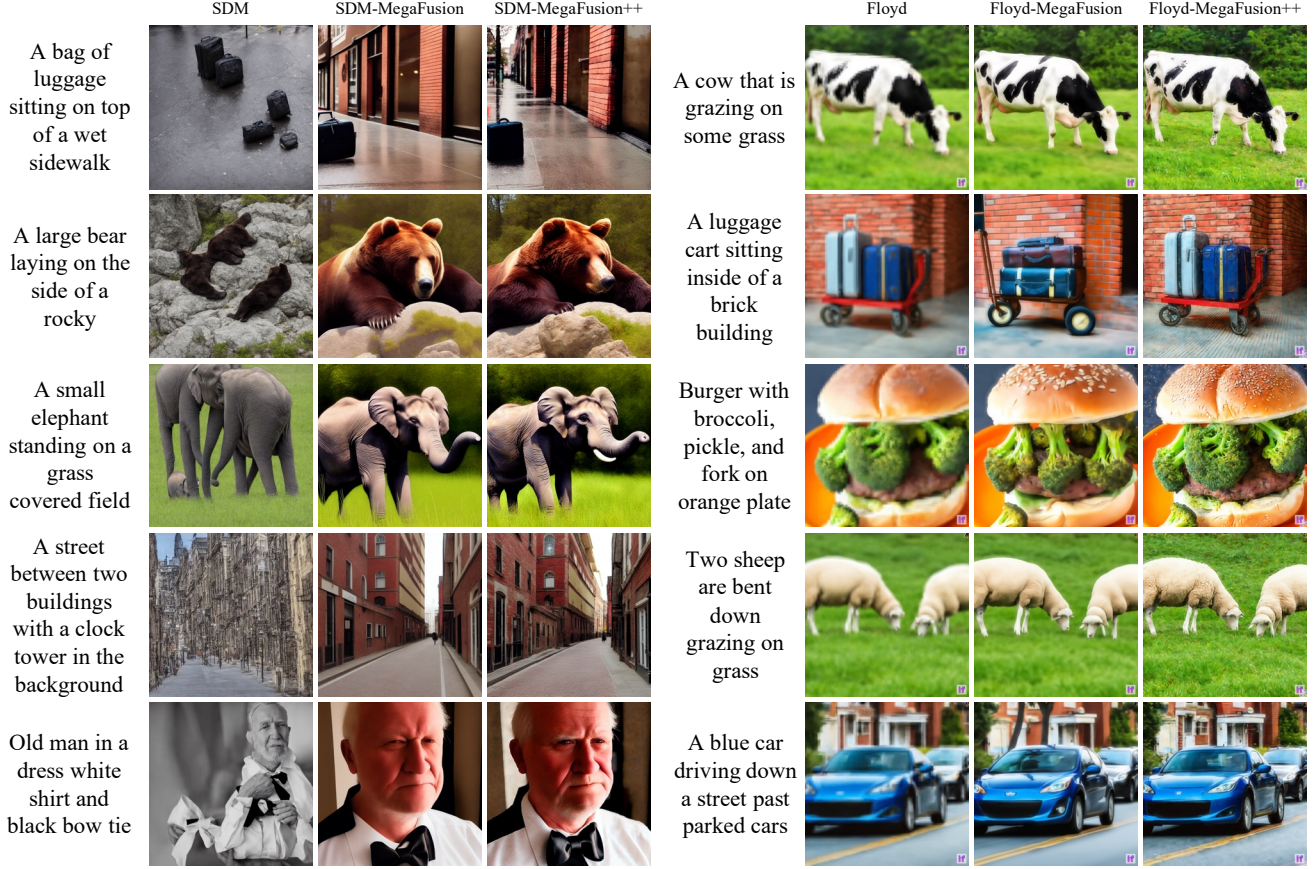


Figure 3. **Qualitative results** of applying our MegaFusion to both latent-space and pixel-space diffusion models for higher-resolution image generation on MS-COCO and commonly used prompts from the Internet. Our method can effectively extend existing diffusion-based models towards synthesizing higher-resolution images of megapixels with correct semantics and details.

DDIM [43] for $T = 50$ steps of sampling unless explicitly stated otherwise. Given that SDM is trained with a fixed resolution of 512×512 , we choose to generate high-resolution images of 1024×1024 for quantitative comparison. Specifically, we orchestrate generation across $k = 3$ resolutions: 512, 768 and 1024, with respective denoising steps of $T_1 = 40$, $T_2 = 5$, and $T_3 = 5$. Furthermore, our proposed MegaFusion can be applied to synthesize even higher-resolution images with SDM for qualitative assessment. SDXL defaults to synthesizing images of 1024×1024 , considering the balance of computational costs, we discard the Refiner module and employ two-stage generation, 1024 and 2048, with their denoising steps being $T_1 = 40$ and $T_2 = 10$. For SD3, which defaults to denoise 28 steps for generating 1024×1024 images, we iterate 20 steps at 1024 resolution and 8 steps at 2048 resolution.

On the other hand, Floyd, a 3-stage cascaded model, sequentially upscales images from 64×64 to 256×256 , culminating in 1024×1024 images. Due to computational constraints, only the first two stages of Floyd are employed in our experiments. The first stage necessitates 100 sampling steps ($T_1 = 80$ for generating 64×64 images, and

$T_2 = 20$ for 128×128), while the second stage requires 50 steps ($T_1 = 40$ for 256×256 and $T_2 = 10$ for 512×512).

Bicubic upsampling serves as the default non-parametric Upsampler Φ . For typical $2\times$ higher-resolution generation, we set the dilation rate $\delta = 2$, and select the hyperparameter $\gamma = 4$ for noise re-scheduling. For classifier-free guidance, to ensure a fair comparison, we apply the default weight w of official implementations across all methods: $w = 7.0$ for SDM, SDXL, SD3, and Floyd-Stage 1, and $w = 4.0$ for Floyd-Stage 2. All experiments are conducted on a single Nvidia RTX A40 GPU, with SDM, SDXL, and SD3 in *float16* precision, and Floyd at in *float32* precision.

Evaluation Datasets. We assess our method and baseline models on the MS-COCO [27] dataset, which comprises approximately 120K images in total, each accompanied by 5 captions. Due to the computational costs of high-resolution generation, we randomly sample 10K images from MS-COCO, assigning a fixed caption to each as input. To ensure consistent comparisons, we utilize the same random seed for each image across methods, neutralizing randomness. For qualitative human evaluations, we use commonly

available prompts from the Internet as text conditions and conditional images from the official code repositories as extra inputs for IP-Adapter and ControlNet.

Evaluation Metrics. To evaluate the quality of generated images, we adopt several widely-used metrics, including Fréchet Inception Distance score (FID) [17], Kernel Inception Distance score (KID), and CLIP [34] text-image similarity (CLIP-T). Following [15], we consider two types of FID and KID: (i) FID_r and KID_r to gauge the quality and diversity of generated images relative to real ones, and (ii) FID_b and KID_b to assess the discrepancies between synthesized samples under the base training resolution and high resolution. These latter metrics reflect the model’s ability to retain generative proficiency at unfamiliar resolutions.

To evaluate the semantic accuracy of generated contents, we adopt MiniGPT-v2 [6] to caption the images, and calculate several linguistic metrics between these captions and the original input text. Concretely, we report the commonly used CIDEr [45], Meteor [2], and ROUGE [26]. Moreover, we detail the GFlops and inference time measured on a single A40 GPU for efficiency comparison.

5.2. Quantitative Results

Objective Metrics. We evaluate the performance of both latent-space and pixel-space models boosted by **MegaFusion** against their baseline counterparts on the MS-COCO [27] dataset. Here, [model-MegaFusion] refers to models employing *truncate and relay* strategy to bridge multi-resolution generation, while [model-MegaFusion++] denotes advanced models incorporating *dilated convolution* and *noise re-scheduling*. We also compare several state-of-the-art methods, such as ScaleCrafter [15], and DemoFusion [9], which are limited to specific latent-space models and less efficient, as well as SDM and SDXL with SR post-processing [12, 47–49], *e.g.* StableSR, and RealESRGAN.

The results in Table 1 highlight significant improvements with MegaFusion across all metrics, including image quality, semantic accuracy, and especially computational efficiency. This confirms that MegaFusion effectively extends the generative capabilities of existing diffusion models towards synthesizing high-resolution images with correct semantics and details at only 40% of the original computational cost. Moreover, incorporating *dilated convolution* and *noise re-scheduling* further improves performance on several metrics, reflecting improved generation diversity and better alignment with real images and text conditions.

Human Assessment. To complement our objective analysis, we conduct a human-centric evaluation focusing on image quality and semantic integrity. Concretely, utilizing identical text prompts and random seed, we synthesize higher-resolution images via standard models (SDM and Floyd) and their MegaFusion-boosted counterparts. Participants are asked to rate the outputs with a score from 1 to

Methods	Image Quality	Semantics	Preference
SDM	2.60	2.05	5.42%
SDM-MegaFusion	3.25	4.40	12.92%
SDM-MegaFusion++	4.25	4.55	81.66%
Floyd-Stage2	2.18	4.28	1.67%
Floyd-MegaFusion	3.45	4.58	21.25%
Floyd-MegaFusion++	4.22	4.58	77.08%

Table 2. **Human evaluation** with MS-COCO captions and commonly used prompts from the Internet as input.

5 (higher is better), considering both image quality and semantic accuracy. Additionally, they also need to select their preferred image among the options for preference rating.

The results in Table 2 affirm that our MegaFusion significantly improves the performance of higher-resolution image generation in terms of image quality and semantic accuracy. Additionally, our advanced MegaFusion++ further demonstrates potential for even greater improvements. This evidence underscores MegaFusion’s ability to elevate pre-trained models, enabling them to produce higher-resolution images with superior quality and precise semantics.

5.3. Qualitative Results

Comparison on text-to-image foundation models. Figure 3 showcases visualization results of higher-resolution image generation in both latent and pixel spaces. These results affirm that MegaFusion can be seamlessly integrated with existing diffusion models to produce images of megapixels with accurate semantics, whereas prior baselines fail to do so. Moreover, incorporating dilated convolutions and noise re-scheduling further improves image details. Additional results are available in the Appendix.

Comparison on models with additional conditions. We further apply MegaFusion to diffusion models equipped with extra input conditions, such as IP-Adapter and ControlNet, as illustrated in Figure 4. Our MegaFusion exhibits universal applicability, significantly extending the capacity of various diffusion models to synthesize high-quality images of higher resolutions, which not only adhere to the input conditions but also maintain semantic integrity. Please refer to the Appendix for more qualitative results.

5.4. Ablation Studies

Proposed strategy & modules. To evaluate the efficacy of our proposed strategy and components, we assess several model variants in both latent and pixel spaces. Here, ‘T&R’, ‘D’, and ‘R’ represent the *truncate and relay* strategy, *dilated convolution*, and *noise re-scheduling*, respectively. The results in Table 3 demonstrate that our strategy and modules significantly elevate the quality and diversity of contents generated by generative models such as SDM (1024×1024) and Floyd (128×128), especially in improv-

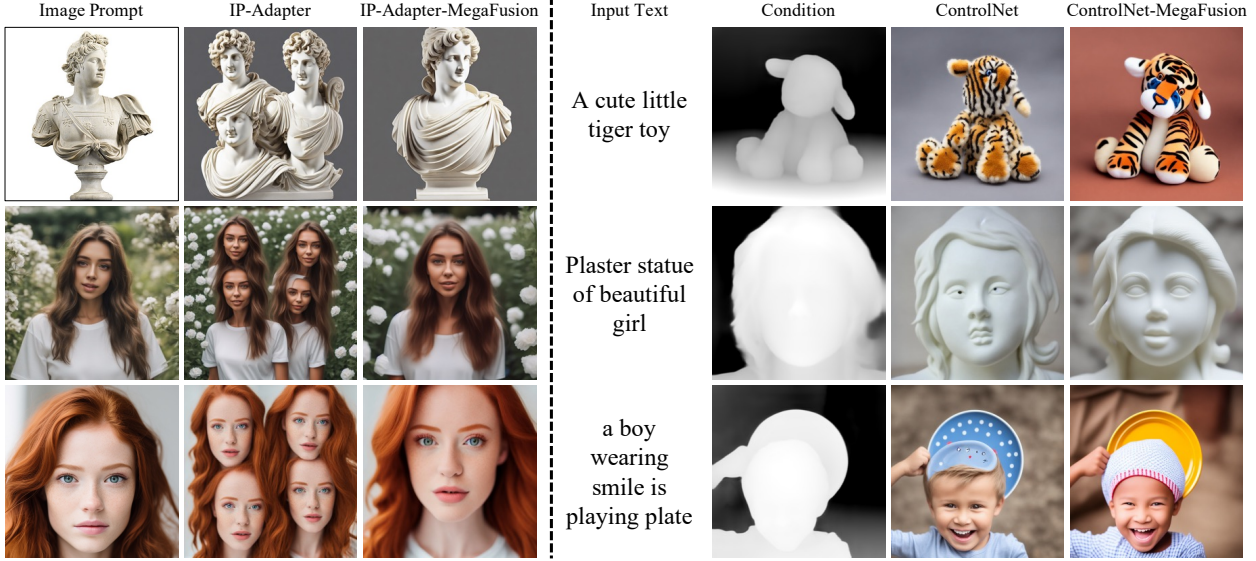


Figure 4. **Qualitative results** of incorporating MegaFusion to models with extra conditional inputs. MegaFusion can be universally applied across various diffusion models, providing the capability for higher-resolution image generation with better semantics and fidelity.

Methods	T&R	D	R	FID _r	FID _b	KID _r	KID _b
SDM [38]	✗	✗	✗	41.35	51.02	0.0086	0.0113
SDM-MegaFusion	✓	✗	✗	30.19	10.98	0.0088	0.0034
SDM-MegaFusion-D	✓	✓	✗	27.56	9.34	0.0075	0.0019
SDM-MegaFusion-R	✓	✗	✓	26.78	9.34	0.0075	0.0019
SDM-MegaFusion++	✓	✓	✓	25.14	7.82	0.0064	0.0012
Floyd-Stage1	✗	✗	✗	66.27	81.65	0.0262	0.0454
Floyd-MegaFusion	✓	✗	✗	53.09	39.73	0.0273	0.0334
Floyd-MegaFusion-D	✓	✓	✗	51.76	41.96	0.0268	0.0345
Floyd-MegaFusion-R	✓	✗	✓	44.27	49.38	0.0215	0.0431
Floyd-MegaFusion++	✓	✓	✓	43.43	50.08	0.0213	0.0437

Table 3. **Ablation Study** on proposed modules in MegaFusion on MS-COCO. The modules gradually improve the higher-resolution generation quality, especially in comparison with real images.

ing the quality and fidelity to real-world images.

Upsampler Φ . The non-parametric Upsampler is crucial in our *truncate and relay* strategy to bridge generation processes across different resolutions. To determine the optimal choice, we evaluate several variants of SDM-MegaFusion++ on MS-COCO dataset, including Config-A (bilinear upsampling); Config-B (bicubic upsampling); Config-C (bicubic with a 5×5 Gaussian filter), and Config-D (bicubic with a 3×3 edge-enhancement kernel). As depicted in Table 4, SDM-MegaFusion++ with Config-B outperforms others in terms of both FID and KID metrics, leading us to adopt bicubic upsampling as the default choice.

6. Conclusion

In this paper, we present **MegaFusion**, a tuning-free approach designed to tackle the challenges of synthesizing

Methods	FID _r ↓	FID _b ↓	KID _r ↓	KID _b ↓
SDM	41.35	51.02	0.0086	0.0113
Config-A	28.03	9.70	0.0076	0.0020
Config-B	25.14	7.82	0.0064	0.0012
Config-C	35.07	18.10	0.0118	0.0063
Config-D	26.56	13.26	0.0065	0.0021

Table 4. **Ablation study** on Upsampler function Φ .

higher-resolution images, effectively resolving issues of semantic inaccuracies and object replication. Our method adopts an innovative *truncate and relay* strategy to elegantly connect generation processes across different resolutions, synthesizing higher-resolution images with megapixels and various aspect ratios. By integrating *dilated convolutions* and *noise re-scheduling*, we further improve the synthesis quality. The versatility of MegaFusion makes it universally applicable to both latent-space and pixel-space diffusion models, as well as their extensions with extra conditions. Extensive experiments have validated the superiority of MegaFusion, demonstrating its capability to generate higher-resolution images with approximately 40% of the original computational cost.

Acknowledgement

This work is supported by National Natural Science Foundation of China (62271308), STCSM (22511105700, 22DZ2229005), 111 plan (BP0719010), and State Key Laboratory of UHD Video and Audio Production and Presentation.

References

- [1] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, Tero Karras, and Ming-Yu Liu. ediff-i: Text-to-image diffusion models with ensemble of expert denoisers. *arXiv preprint arXiv:2211.01324*, 2022. 3, 13
- [2] Satangeev Banerjee and Alon Lavie. Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*, pages 65–72, 2005. 7
- [3] Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel. MultiDiffusion: Fusing diffusion paths for controlled image generation. In *Proceedings of the International Conference on Machine Learning*, pages 1737–1752. PMLR, 2023. 2
- [4] Tim Brooks, Aleksander Holynski, and Alexei A. Efros. InstructPix2Pix: Learning to follow image editing instructions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2023. 1, 2
- [5] Junsong Chen, YU Jincheng, GE Chongjian, Lewei Yao, Enze Xie, Zhongdao Wang, James Kwok, Ping Luo, Huchuan Lu, and Zhenguo Li. Pixart- α : Fast training of diffusion transformer for photorealistic text-to-image synthesis. In *Proceedings of the International Conference on Learning Representations*, 2024. 1
- [6] Jun Chen, Deyao Zhu, Xiaoqian Shen, Xiang Li, Zechun Liu, Pengchuan Zhang, Raghuraman Krishnamoorthi, Vikas Chandra, Yunyang Xiong, and Mohamed Elhoseiny. Minigt-v2: large language model as a unified interface for vision-language multi-task learning. *arXiv preprint arXiv:2310.09478*, 2023. 7
- [7] Jiaxiang Cheng, Pan Xie, Xin Xia, Jiashi Li, Jie Wu, Yuxi Ren, Huixia Li, Xuefeng Xiao, Min Zheng, and Lean Fu. Resadapter: Domain consistent resolution adapter for diffusion models. *arXiv preprint arXiv:2403.02084*, 2024. 2, 5
- [8] Deepfloyd. Deepfloyd. URL <https://www.deepfloyd.ai/>, 2023. 1, 2, 5, 12, 13, 17
- [9] Ruoyi Du, Dongliang Chang, Timothy Hospedales, Yi-Zhe Song, and Zhanyu Ma. Demofusion: Democratising high-resolution image generation with no \$\$\$\$. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2024. 2, 5, 7, 12, 15, 19
- [10] Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers for high-resolution image synthesis. In *Proceedings of the International Conference on Machine Learning*, 2024. 1, 5, 15, 18
- [11] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 2020. 2
- [12] Baisong Guo, Xiaoyun Zhang, Haoning Wu, Yu Wang, Ya Zhang, and Yan-Feng Wang. Lar-sr: A local autoregressive model for image super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1909–1918, June 2022. 7, 13
- [13] Lanqing Guo, Yingqing He, Haoxin Chen, Menghan Xia, Xiaodong Cun, Yufei Wang, Siyu Huang, Yong Zhang, Xintao Wang, Qifeng Chen, et al. Make a cheap scaling: A self-cascade diffusion model for higher-resolution adaptation. *arXiv preprint arXiv:2402.10491*, 2024. 2
- [14] Moayed Haji-Ali, Guha Balakrishnan, and Vicente Ordonez. Elasticdiffusion: Training-free arbitrary size image generation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2024. 2
- [15] Yingqing He, Shaoshu Yang, Haoxin Chen, Xiaodong Cun, Menghan Xia, Yong Zhang, Xintao Wang, Ran He, Qifeng Chen, and Ying Shan. Scalecrafter: Tuning-free higher-resolution visual generation with diffusion models. In *Proceedings of the International Conference on Learning Representations*, 2023. 2, 4, 5, 7, 15, 19
- [16] Amir Hertz, Ron Mokady, Jay Tenenbaum, Kfir Aberman, Yael Pritch, and Daniel Cohen-Or. Prompt-to-prompt image editing with cross attention control. In *Proceedings of the International Conference on Learning Representations*, 2023. 1, 2
- [17] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems*, 2017. 7
- [18] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022. 1, 2
- [19] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in Neural Information Processing Systems*, 2020. 1, 2
- [20] Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *Advances in Neural Information Processing Systems Workshops*, 2021. 1
- [21] Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J Fleet. Video diffusion models. In *Advances in Neural Information Processing Systems*, 2022. 1, 2
- [22] Emiel Hooeboom, Jonathan Heek, and Tim Salimans. Simple diffusion: End-to-end diffusion for high resolution images. In *Proceedings of the International Conference on Machine Learning*, 2023. 4, 5, 14
- [23] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *Proceedings of the International Conference on Learning Representations*, 2022. 2
- [24] Linjiang Huang, Rongyao Fang, Aiping Zhang, Guanglu Song, Si Liu, Yu Liu, and Hongsheng Li. Fouriscale: A frequency perspective on training-free high-resolution image synthesis. In *Proceedings of the European Conference on Computer Vision*, 2024. 2
- [25] Bahjat Kavar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic:

- Text-based real image editing with diffusion models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2023. 1, 2
- [26] Chin-Yew Lin. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81, 2004. 7
- [27] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *Proceedings of the European Conference on Computer Vision*, 2014. 5, 6, 7, 12, 16, 17, 18
- [28] 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 Conference on Computer Vision and Pattern Recognition*, 2024. 1
- [29] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 1
- [30] Adyasha Maharana, Darryl Hannan, and Mohit Bansal. Storydall-e: Adapting pretrained text-to-image transformers for story continuation. In *Proceedings of the European Conference on Computer Vision*, 2022. 1
- [31] Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *Proceedings of the International Conference on Learning Representations*, 2021. 1, 2
- [32] Xichen Pan, Pengda Qin, Yuhong Li, Hui Xue, and Wenhu Chen. Synthesizing coherent story with auto-regressive latent diffusion models. In *Winter Conference on Applications of Computer Vision*, 2024. 1
- [33] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach. Sdxl: Improving latent diffusion models for high-resolution image synthesis. In *Proceedings of the International Conference on Learning Representations*, 2024. 1, 2, 5, 13, 21, 22, 23, 24, 25, 26
- [34] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *Proceedings of the International Conference on Machine Learning*, 2021. 3, 7
- [35] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 2020. 3
- [36] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv preprint arXiv:2204.06125*, 2022. 2
- [37] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *Proceedings of the International Conference on Machine Learning*, 2021. 2
- [38] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Bjorn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. 1, 2, 5, 8, 13, 16
- [39] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention*, 2015. 3
- [40] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic text-to-image diffusion models with deep language understanding. In *Advances in Neural Information Processing Systems*, 2022. 2
- [41] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5b: An open large-scale dataset for training next generation image-text models. In *Advances in Neural Information Processing Systems*, 2022. 2
- [42] Uriel Singer, Adam Polyak, Thomas Hayes, Xi Yin, Jie An, Songyang Zhang, Qiyuan Hu, Harry Yang, Oron Ashual, Oran Gafni, et al. Make-a-video: Text-to-video generation without text-video data. In *Proceedings of the International Conference on Learning Representations*, 2023. 1, 2
- [43] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. In *Proceedings of the International Conference on Learning Representations*, 2020. 1, 2, 6
- [44] Jiayan Teng, Wendi Zheng, Ming Ding, Wenyi Hong, Jianqiao Wangni, Zhuoyi Yang, and Jie Tang. Relay diffusion: Unifying diffusion process across resolutions for image synthesis. In *Proceedings of the International Conference on Learning Representations*, 2024. 2, 4
- [45] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4566–4575, 2015. 7
- [46] Catherine Wah, Steve Branson, Peter Welinder, Pietro Perona, and Serge J. Belongie. The caltech-ucsd birds-200-2011 dataset, 2011. 12, 13
- [47] Jianyi Wang, Zongsheng Yue, Shangchen Zhou, Kelvin CK Chan, and Chen Change Loy. Exploiting diffusion prior for real-world image super-resolution. In *International Journal of Computer Vision*, 2023. 5, 7, 13
- [48] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1905–1914, 2021. 7, 13
- [49] Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In

Proceedings of the European Conference on Computer Vision Workshops, 2018. 7, 13

- [50] Shaoan Xie, Zhifei Zhang, Zhe Lin, Tobias Hinz, and Kun Zhang. Smartbrush: Text and shape guided object inpainting with diffusion model. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2023. 1, 2
- [51] Li Xin, Chu Wenqing, Wu Ye, Yuan Weihang, Liu Fanglong, Zhang Qi, Li Fu, Feng Haocheng, Ding Errui, and Wang Jingdong. Videogen: A reference-guided latent diffusion approach for high definition text-to-video generation. *arXiv preprint arXiv:2309.00398*, 2023. 1, 2
- [52] Tao Xu, Pengchuan Zhang, Qiuyuan Huang, Han Zhang, Zhe Gan, Xiaolei Huang, and Xiaodong He. Attngan: Fine-grained text to image generation with attentional generative adversarial networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018. 2
- [53] Hu Ye, Jun Zhang, Sibio Liu, Xiao Han, and Wei Yang. Ip-adapter: Text compatible image prompt adapter for text-to-image diffusion models. *arXiv preprint arxiv:2308.06721*, 2023. 2, 5, 15
- [54] Fisher Yu and Vladlen Koltun. Multi-scale context aggregation by dilated convolutions. *arXiv*, 2016. 2, 4
- [55] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiaolei Huang, and Dimitris N Metaxas. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the International Conference on Computer Vision*, 2017. 2
- [56] Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the International Conference on Computer Vision*, 2023. 2, 5, 15, 20
- [57] Shen Zhang, Zhaowei Chen, Zhenyu Zhao, Yuhao Chen, Yao Tang, and Jiajun Liang. Hidiffusion: Unlocking higher-resolution creativity and efficiency in pretrained diffusion models. *arXiv preprint arXiv:2311.17528*, 2023. 2

In this appendix, we start by giving more details on the implementation details of our proposed MegaFusion in Section A. Then, we provide extra quantitative comparisons to further demonstrate the universality and effectiveness of our method in Section B. Next, we offer additional qualitative results across various experimental settings and methods to illustrate the superiority of our proposed MegaFusion in Section C. Finally, we discuss the limitations of our method and future work in Section D.

A. Implementation Details

More Details on Floyd-MegaFusion. We have evaluated the higher-resolution image generation performance of Floyd [8] at resolutions of 128×128 and 512×512 . For 128×128 resolution, we directly apply MegaFusion to the first stage of Floyd. As for the comparison at 512×512 resolution, we utilize the first two stages of Floyd. Considering that the quality of the results from the first stage generation would significantly affect the second generation stage, we opt for using the 64×64 images generated by the original first stage model as inputs of both the baseline and our boosted Floyd-MegaFusion. That is, higher-resolution image generation is only performed under the second generation stage. Ultimately, the experimental results presented in Table 1 of our submitted manuscript effectively demonstrate the universality and effectiveness of our proposed MegaFusion. Furthermore, we also conduct experiments where 128×128 out-of-distribution images are generated in the first stage, followed by 512×512 resolution images in the second stage. This further demonstrates that MegaFusion maintains semantic accuracy across all stages of generation.

Details on Human Evaluation. To more effectively reflect the performance of different models in generating high-resolution images, we have recruited 10 volunteers with a background in image generation research for human evaluation. Specifically, the evaluators are asked to follow these rules: (i) Rate unknown source images on a score from 1 to 5 for both image quality and semantic accuracy, with higher scores indicating better quality; and (ii) Observe the results generated by different models with the same input conditions and select their favourite one based on overall quality and semantic accuracy.

B. Additional Quantitative Results

B.1. Comparison on crop FID/KID

Following previous work [9], we also evaluate crop FID and crop KID metrics on the generated results of various models to reflect the quality of local patches in the images. As depicted in Table 5, previous methods are often limited to specific latent-space models, whereas our MegaFusion consistently improves the quality of high-resolution image generation across both latent-space and pixel-space models.

Method	SDM-1024	SDXL-2048	Floyd-128	Floyd-512
Original	41.21/0.0139	42.29/0.0125	70.16/0.0224	40.65/0.0171
ScaleCrafter	32.24/0.0085	26.58/0.0062	inapplicable	inapplicable
DemoFusion	inapplicable	25.91/0.0061	inapplicable	inapplicable
MegaFusion	39.42/0.0137	27.38/0.0063	57.24/0.0243	32.36/0.0122
MegaFusion++	33.39/ 0.0084	25.64/0.0049	41.22/0.0188	29.18/0.0077

Table 5. Comparison of FID_{crop}/KID_{crop} on MS-COCO dataset.

B.2. Comparison on CUB-200 Dataset

To demonstrate the universality of our proposed MegaFusion, in addition to the MS-COCO [27] dataset, we also conduct quantitative evaluations on the CUB-200 [46] dataset, which is also commonly used in previous works. The CUB-200 dataset consists of over 10K images of 200 categories of birds, each accompanied by 10 textual descriptions. Considering computational costs and time expenditure, similar to the experimental settings on the MS-COCO dataset in our manuscript, we randomly select 1K images from the CUB-200 dataset. Each image is assigned a fixed caption, and the same random seed is used across different methods to eliminate the effects of randomness among models. As depicted in Table 6, our proposed MegaFusion can also be universally applied to both latent-space and pixel-space diffusion models on the CUB-200 dataset, achieving high-quality higher-resolution image generation.

B.3. More Results of Floyd-MegaFusion

As mentioned above, we also conduct experiments that first generate 128×128 out-of-distribution images, followed by 512×512 high-resolution images on the Floyd model. As depicted in Table 7, MegaFusion consistently improves the high-resolution generation capability of Floyd under both settings. This demonstrates that MegaFusion can improve the semantic accuracy of high-resolution images at any stage of the generation process.

B.4. Ablation Study of Classifier-free Guidance

As detailed in the implementation details, to ensure a fair comparison and eliminate the impact of classifier-free guidance (CFG) on generation quality and efficiency, we use the default CFG weights from official implementations for all methods and their corresponding MegaFusion-boosted counterparts. To further investigate the impact of CFG on MegaFusion at higher resolutions, we generate 100 images from the MS-COCO dataset using SDM-MegaFusion and SDXL-MegaFusion with varying CFG values, using the same text prompt and random seed as inputs, and evaluate the FID scores against our testset. The results in Figure 5 indicate that classifier-free guidance does affect our high-resolution generation quality, with preliminary findings indicating that $w = 7.0$ is a relatively good choice for SDM-MegaFusion and SDXL-MegaFusion.

Methods	resolution	FID _r ↓	FID _b ↓	KID _r ↓	KID _b ↓	CLIP-T↑	CIDEr↑	Meteor↑	ROUGE↑	GFlops	Inference time
SDM [38]	1024 × 1024	77.92	46.34	0.0363	0.0220	0.2952	8.12	7.48	7.09	135.0K	15.17s
SDM-MegaFusion	1024 × 1024	<u>71.78</u>	<u>36.21</u>	<u>0.0303</u>	<u>0.0189</u>	<u>0.3060</u>	<u>24.46</u>	<u>11.98</u>	<u>12.62</u>	48.2K	7.56s
SDM-MegaFusion++	1024 × 1024	68.92	34.94	0.0251	0.0182	0.3115	28.52	12.32	13.29	48.2K	7.56s
SDXL [33]	2048 × 2048	73.49	48.78	0.0308	0.0274	0.2994	16.43	9.90	10.35	540.2K	79.66s
SDXL-MegaFusion	2048 × 2048	<u>72.62</u>	<u>13.72</u>	<u>0.0296</u>	<u>0.0039</u>	<u>0.3113</u>	<u>25.98</u>	<u>13.23</u>	<u>13.33</u>	216.1K	30.94s
SDXL-MegaFusion++	2048 × 2048	65.10	11.55	0.0225	0.0026	0.3122	26.35	13.98	14.92	216.1K	30.94s
Floyd-Stage1 [8]	128 × 128	87.04	105.59	<u>0.0341</u>	0.0658	0.2866	9.95	8.28	9.07	111.7K	77.08s
Floyd-MegaFusion	128 × 128	<u>77.82</u>	36.49	0.0413	0.0281	<u>0.3080</u>	<u>22.12</u>	17.06	20.62	44.9K	32.19s
Floyd-MegaFusion++	128 × 128	73.54	<u>45.76</u>	0.0334	<u>0.0388</u>	0.3086	22.52	<u>16.93</u>	<u>20.05</u>	44.9K	32.19s
Floyd-Stage2 [8]	512 × 512	80.34	41.65	0.0401	0.0215	0.3013	23.59	12.28	11.67	60.7K	48.58s
Floyd-MegaFusion	512 × 512	<u>77.66</u>	<u>39.34</u>	<u>0.0348</u>	<u>0.0141</u>	<u>0.3110</u>	<u>24.63</u>	15.74	<u>15.29</u>	24.3K	21.72s
Floyd-MegaFusion++	512 × 512	62.91	34.40	0.0232	0.0115	0.3141	25.44	<u>13.90</u>	18.51	24.3K	21.72s

Table 6. **Quantitative comparison** on CUB-200 [46] dataset. **RED**: best performance, **BLUE**: second best performance.

Methods	resolution	FID _r ↓	FID _b ↓	KID _r ↓	KID _b ↓	CLIP-T↑	CIDEr↑	Meteor↑	ROUGE↑	GFlops	Inference time
Floyd-Stage1 [8]	128 × 128	66.27	81.65	<u>0.0262</u>	0.0454	0.2818	14.69	18.22	25.06	111.7K	77.08s
Floyd-MegaFusion	128 × 128	<u>53.09</u>	39.73	0.0273	0.0334	<u>0.3024</u>	25.01	<u>25.00</u>	<u>31.35</u>	44.9K	32.19s
Floyd-MegaFusion++	128 × 128	43.43	<u>50.08</u>	0.0213	<u>0.0437</u>	0.3046	<u>20.28</u>	25.01	31.64	44.9K	32.19s
Floyd-Stage2 [8]	64 → 512	46.64	38.15	0.0254	0.0166	0.3098	<u>23.85</u>	21.47	26.26	60.7K	48.58s
Floyd-MegaFusion	64 → 512	<u>39.80</u>	<u>24.87</u>	<u>0.0164</u>	<u>0.0078</u>	<u>0.3106</u>	23.22	<u>23.51</u>	<u>29.30</u>	24.3K	21.72s
Floyd-MegaFusion++	64 → 512	26.34	24.55	0.0063	0.0077	0.3110	24.01	23.58	29.52	24.3K	21.72s
Floyd-Stage2 [8]	128 → 512	61.24	108.01	0.0253	0.0734	0.2779	15.16	14.76	19.75	60.7K	48.58s
Floyd-MegaFusion	128 → 512	<u>58.19</u>	88.56	<u>0.0187</u>	0.0379	<u>0.2821</u>	<u>16.28</u>	15.65	<u>20.02</u>	24.3K	21.72s
Floyd-MegaFusion++	128 → 512	57.92	<u>94.93</u>	0.0181	<u>0.0417</u>	0.2835	16.36	<u>15.47</u>	21.34	24.3K	21.72s

Table 7. **More comparison results** on Floyd model and its MegaFusion boosted counterparts under different settings. Within each unit, we denote the best performance in **RED** and the second-best performance in **BLUE**.




Figure 5. **Ablation study** of classifier-free guidance (CFG) weight on SDM-MegaFusion and SDXL-MegaFusion.

C. Additional Qualitative Results

C.1. Evidence Behind the Core idea & intuition

As stated in eDiff-I [1], diffusion models synthesize semantics during early denoising stages and refine image details in later stages. As depicted in Figure 6, we also observe that semantic deviations and object repetitions commonly encountered at higher resolutions primarily stem from incorrect semantics generated during early denoising, leading to irreparable errors. Thus, our **intuition and insight** here

are: perform early denoising at the original resolution to generate accurate semantic information, followed by *truncate* and *relay* to continue denoising at higher resolutions, thereby enriching texture details. This enables MegaFusion to produce high-quality, semantically accurate higher-resolution images with lower computational costs, while supporting arbitrary aspect ratios.

C.2. Disadvantages of Direct Upsampling

Compared to our MegaFusion for higher-resolution image generation, a more straightforward approach is to directly apply upsampling to images generated by diffusion models. Although simple, this will introduce three potential issues: (i) Direct super-resolution may lead to unrealistic texture details, such as blurring and artifacts, especially at high upsampling factors; (ii) While diffusion-based SR methods can produce more realistic textures via iterative denoising, they often involve significantly higher computational costs and may not support arbitrary aspect ratios; (iii) Most critically, as shown in Figure 7, directly upsampling [12, 47–49] low-resolution images can stretch




Figure 6. **Evidence behind our core idea and intuition.** For $T = 50$ steps of DDIM sampling, we visualize the key stages of the image generation process. For SDM and SDXL, incorrect semantics are generated during the early denoising stages of high-resolution generation, leading to irreparable errors. In contrast, MegaFusion generates accurate semantics and further enriches texture details at higher resolutions. The input text prompts are “A cute black cat” and “A white dog sits on the grass.” For ease of visualization, the images are scaled to the same size.

and distort content, particularly when generating under non-standard aspect ratios (e.g. 1 : 4), diminishing the natural aesthetic of images.

In contrast, MegaFusion seamlessly bridges coarse-to-fine generation processes, efficiently producing accurate semantics at low resolutions and enriching texture details at high resolutions. Leveraging iterative denoising at higher resolutions, it can synthesize aesthetically pleasing high-resolution images even with non-standard aspect ratios.




Figure 7. **Analysis of direct upsampling.** Using diffusion models to generate images with non-standard aspect ratios directly or via upsampling, may lead to stretching and distortion (e.g., trees on both sides), while MegaFusion effectively mitigates this issue.

C.3. Effects of hyperparameters δ and γ

For denoising at the original size, we do not employ dilation. In qualitative experiments for high-resolution gener-

ation, we test various δ values and find that $\delta = 2$ is a stable choice under our experimental settings, which will not introduce blurriness or semantic deviations. As described in our manuscript, we draw inspiration from simple diffusion [22], which derives the SNR relationship between images of different resolution based on the mean and variance of pixel distributions. Substituting this into our derived relationship, we obtain that $\gamma = 4$. Qualitative experiments also confirm that this is an appropriate choice. Some visualization examples are shown in Figure 8.




Figure 8. **Qualitative comparisons** of applying different hyperparameters δ and γ .

C.4. Ablation Study of Truncation Steps

In the *truncate* and *relay* strategy, the number of denoising steps at each stage may also affect generation quality. Our intuition and experience suggest that more denoising steps at lower resolutions improve generation efficiency, while additional steps at higher resolutions enhance texture details. However, conducting a comprehensive evaluation to determine the optimal truncation steps would incur significant computational costs. Therefore, in our implementation, we empirically select truncation steps for each model based on experience, and validate the above conclusions through qualitative experiments, as shown in Figure 9. Considering the trade-off between generation quality and efficiency, we choose denoising steps of $T_1 = 40$, $T_2 = 5$, and $T_3 = 5$ as the default configuration for SDM-MegaFusion.

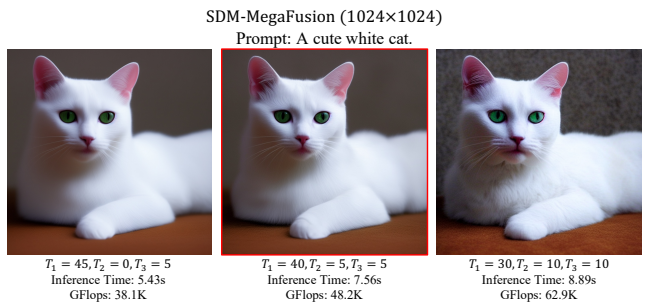


Figure 9. **Qualitative ablation study** of truncation steps.

C.5. Text-to-Image Foundation Models

We present more visualizations of higher-resolution image generation using both latent-space and pixel-space text-to-image models in Figure 10 and 11, respectively, to demonstrate the universality and robustness of our proposed method. The visual outcomes explicitly confirm that when

pre-trained models fail to scale to higher resolutions, our approach can be universally integrated into existing latent-space and pixel-space diffusion models, improving their capability to synthesize higher-resolution images of megapixels with accurate semantics. Moreover, our further enhanced MegaFusion++ significantly boosts the quality of the generated images, producing sharper and clearer details.

C.6. Compatibility with Transformer-based Models

To further demonstrate the versatility and effectiveness of MegaFusion, we also apply it to the transformer-based (DiT) SD3 [10] model. Since DiT-based methods do not involve convolutions, we boost the model via only the *truncate and relay* strategy. As presented in Figure 12, SD3 also encounters quality degradation when generating higher-resolution images directly, while our MegaFusion effectively improves its high-resolution generation capabilities.

C.7. Comparison to state-of-the-art

To further evaluate the quality of MegaFusion, we compare it with existing state-of-the-art high-resolution generation methods. Given that these methods (ScaleCrafter [15] and DemoFusion [9]) are typically restricted to specific models, we conduct comparisons on models based on SDXL. The results in Figure 13 indicate that existing methods still face quality degradation and object repetition when generating high-resolution images. In contrast, MegaFusion produces high-quality, semantically accurate high-resolution images, and is much more efficient than existing approaches, as shown in Table 1 of our manuscript.

C.8. Models with additional conditions

We have confirmed that our method is equally applicable to diffusion models with additional input conditions, such as ControlNet [56] with depth maps and IP-Adapter [53] with reference images as extra inputs. As depicted in Figure 14, we further discover that ControlNet with canny edges or human poses as conditional inputs also struggle with synthesizing higher-resolution images, and often produce images that are not fidelity to input conditions, with confusing semantics and poor image quality. In contrast, with the assistance of our proposed MegaFusion, our boosted model, ControlNet-MegaFusion consistently generates high-quality images of higher resolutions with accurate semantics, that are fidelity to conditions.

C.9. Generation with Arbitrary Aspect Ratios

As previously stated, our MegaFusion also enables existing pre-trained diffusion models to generate images at arbitrary aspect ratios. Figure 15, 16 and 17 showcase more qualitative results from SDXL-MegaFusion across various aspect ratios and resolutions, including 1 : 1 (2048×2048),

16 : 9 (1920×1080), 3 : 4 (1536×2048), and 4 : 3 (2048×1536). Moreover, as presented in Figure 18, 19, and 20, we also include visualizations with **non-standard** aspect ratios, such as 1 : 4 (640×2560), 4 : 1 (2560×640), 1 : 2 (1024×2048), 2 : 1 (2048×1024), 21 : 9 (2016×864), and 9 : 21 (864×2016). These impressive outcomes further demonstrate the scalability and superiority of our approach.

C.10. Compatibility with LoRA

To further illustrate the versatility and broad applicability of MegaFusion, we apply it to SDM and SDXL models using LoRA from the open-source community for personalized higher-resolution image generation. As depicted in Figure 21, MegaFusion can seamlessly integrate with various LoRAs of SDM and SDXL, demonstrating significant potential for artistic and commercial applications.

D. Limitations & Future Work

D.1. Limitations

Since our proposed MegaFusion is a tuning-free approach built on existing latent-space and pixel-space image generation models, it inevitably inherits some limitations of current diffusion-based generative models. For example, when handling complex textual conditions, the generated content often struggles to accurately reflect input prompts, particularly in aspects such as attribute binding and positional control. This may lead to degraded synthesis quality during high-resolution generation with MegaFusion. However, more powerful backbone models are expected to mitigate this issue, and when combined with MegaFusion, they are likely to produce higher-quality images at higher resolutions with low computational costs.

D.2. Future Work

The striking quantitative results produced by MegaFusion have confirmed its potential to overcome the limitations of existing diffusion-based generative models and to improve their capabilities to synthesize high-resolution outcomes. Additionally, we have observed that existing video generation models encounter significant semantic deviations and quality degradation when generating content beyond their pre-trained spatial resolution and temporal length. Therefore, we anticipate further applying MegaFusion to current video generation models towards efficient, low-cost, higher-resolution, and longer video content generation. Similarly, MegaFusion also holds the potential for extension to 3D generation models and models for image and video editing, which are also left for future exploration.

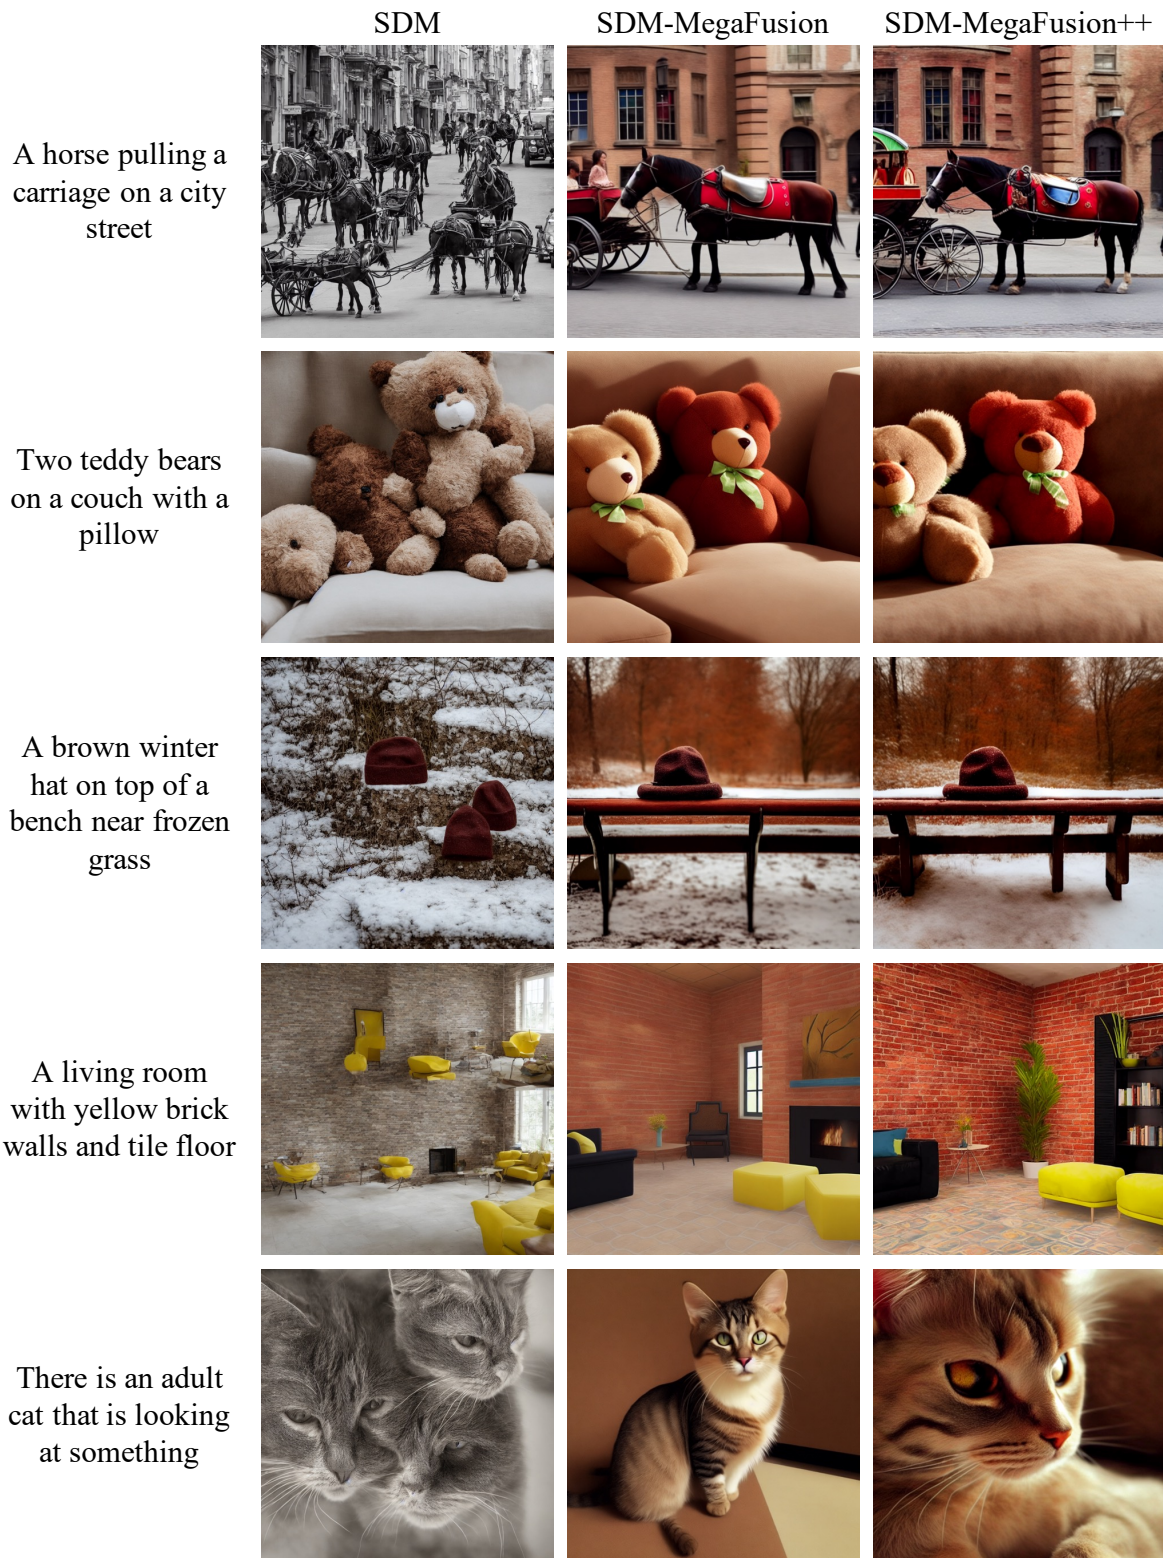


Figure 10. **More qualitative results** of applying our MegaFusion to latent-space diffusion model (SDM [38]) for higher-resolution (1024×1024) image generation on MS-COCO [27] and commonly used prompts from the Internet.

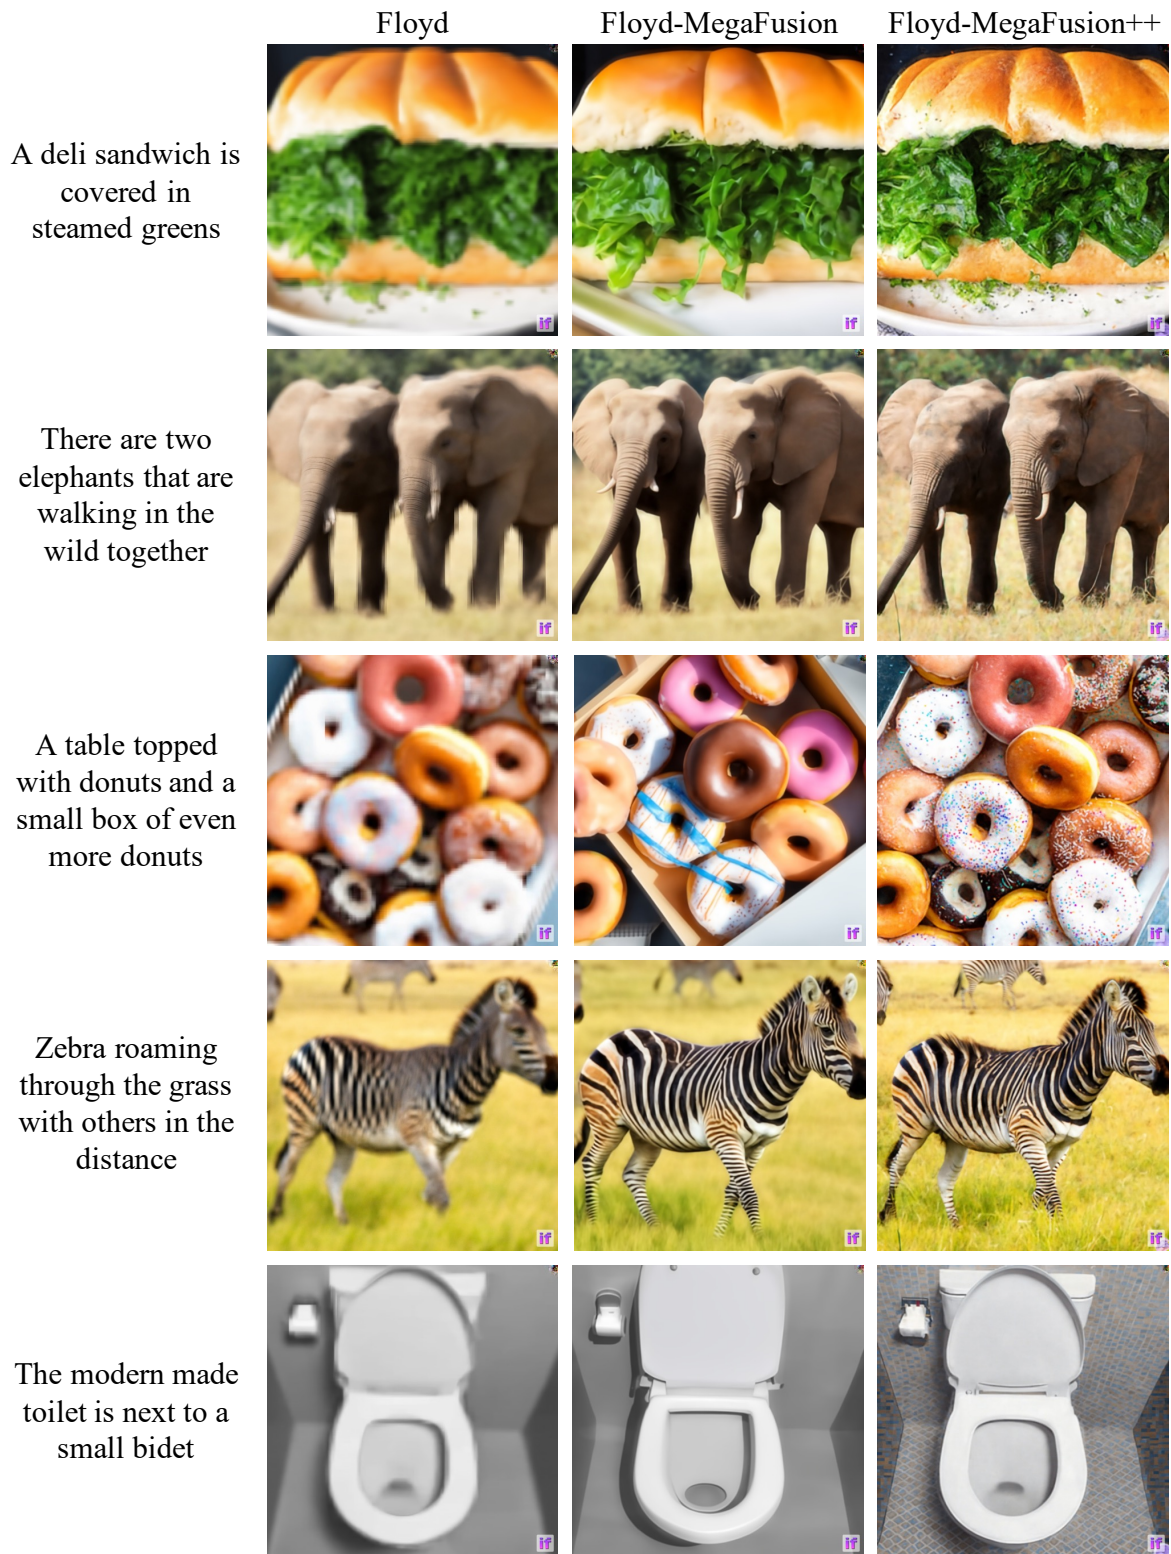


Figure 11. **More qualitative results** of applying our MegaFusion to pixel-space diffusion model (Floyd [8]) for higher-resolution (512×512) image generation on MS-COCO [27] and commonly used prompts from the Internet.

SD3



SD3-MegaFusion



The two teddy bears are posed together to take a photo.



A stone statue of an elephant near a large vase.



A person on a four-wheeler herding sheep in the snow.



A few bags laying around in a living room.

Figure 12. **Qualitative results** of applying our MegaFusion to latent-space diffusion model (SD3 [10]) for higher-resolution (2048×2048) image generation on MS-COCO [27] and commonly used prompts from the Internet.



Figure 13. **Qualitative comparison** with existing state-of-the-art methods (ScaleCrafter [15] and DemoFusion [9]). Our MegaFusion can generate images with details and accurate semantics at high resolution, whereas existing methods struggle to do so.




Figure 14. **Qualitative results** of applying our MegaFusion to ControlNet [56] with canny edges or human poses as extra conditions for higher-resolution (1024×1024) image generation with better semantics and fidelity.



Figure 15. **More qualitative results** of applying our MegaFusion to SDXL [33] model for higher-resolution image generation with various aspect ratios and resolutions.



Figure 16. **More qualitative results** of applying our MegaFusion to SDXL [33] model for higher-resolution image generation with various aspect ratios and resolutions.



Figure 17. **More qualitative results** of applying our MegaFusion to SDXL [33] model for higher-resolution image generation with various aspect ratios and resolutions.




Figure 18. **More qualitative results** of applying our MegaFusion to SDXL [33] model for higher-resolution image generation with various **non-standard** aspect ratios and resolutions.

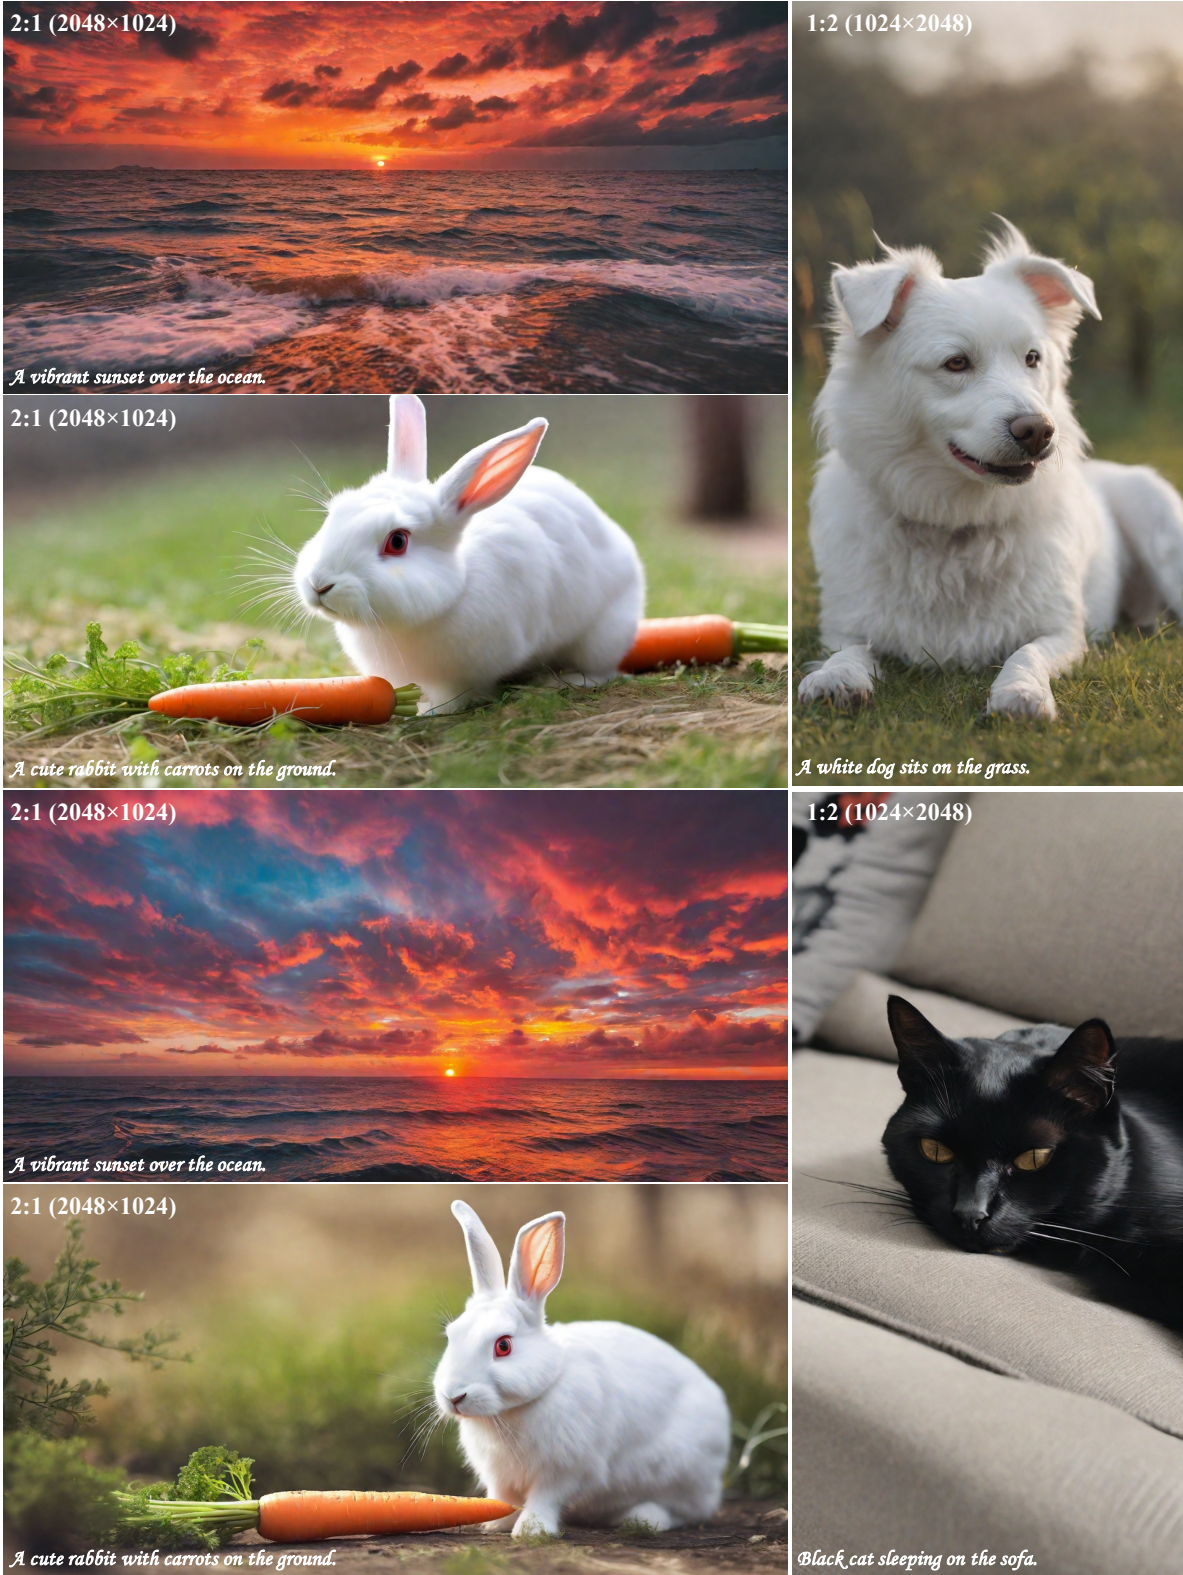


Figure 19. **More qualitative results** of applying our MegaFusion to SDXL [33] model for higher-resolution image generation with various **non-standard** aspect ratios and resolutions.

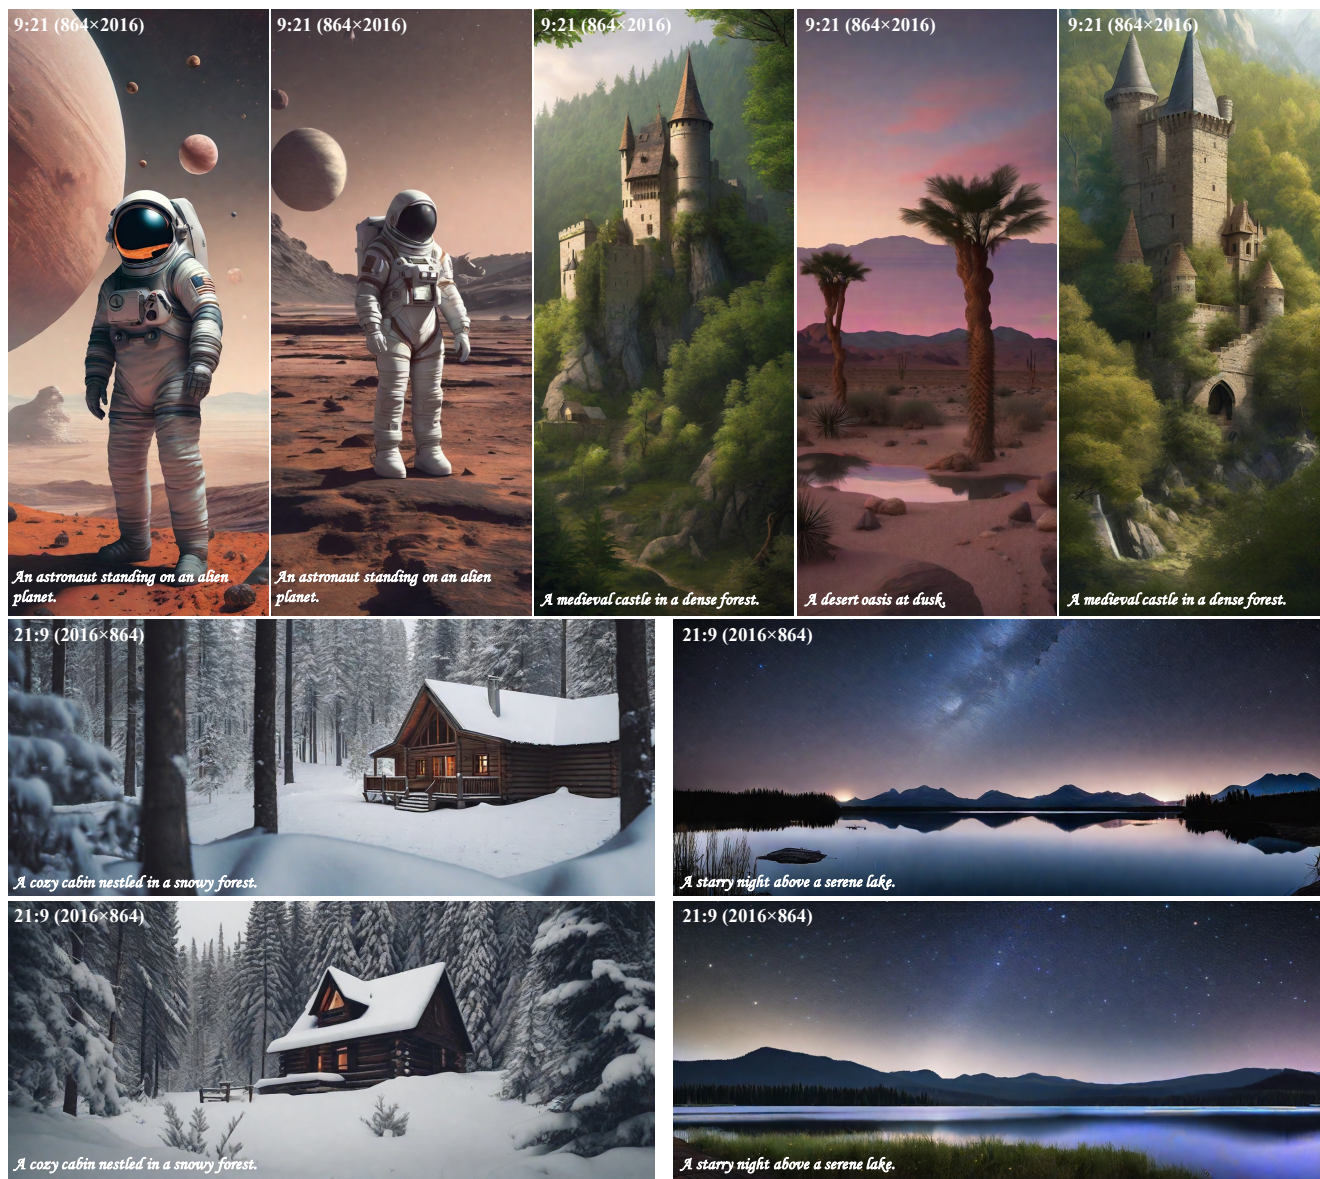


Figure 20. **More qualitative results** of applying our MegaFusion to SDXL [33] model for higher-resolution image generation with various **non-standard** aspect ratios and resolutions.




Figure 21. **Qualitative results** of applying MegaFusion to high-resolution image generation with LoRA-integrated SDM and SDXL. Similarly, SDM and SDXL integrated with LoRA also face common challenges like semantic deviations and object repetitions in high-resolution generation, while MegaFusion effectively addresses these challenges.