

Multi-Scale Diffusion: Enhancing Spatial Layout in High-Resolution Panoramic Image Generation

Xiaoyu Zhang¹, Teng Zhou¹, Xinlong Zhang¹, Jia Wei¹, Yongchuan Tang^{1*}

¹ *Zhejiang University, Hangzhou, China*

{xiaoyzhang, tengzhou, xinlzhang, weijia_77, yctang}@zju.edu.cn

Abstract—Diffusion models have recently gained recognition for generating diverse and high-quality content, especially in image synthesis. These models excel not only in creating fixed-size images but also in producing panoramic images. However, existing methods often struggle with spatial layout consistency when producing high-resolution panoramas due to the lack of guidance on the global image layout. This paper introduces the Multi-Scale Diffusion (MSD), an optimized framework that extends the panoramic image generation framework to multiple resolution levels. Our method leverages gradient descent techniques to incorporate structural information from low-resolution images into high-resolution outputs. Through comprehensive qualitative and quantitative evaluations against prior work, we demonstrate that our approach significantly improves the coherence of high-resolution panorama generation.

Index Terms—Diffusion Model, Panoramic Image Generation, High-Resolution, Gradient Descent

I. INTRODUCTION

Recent advances in diffusion models have shown remarkable capabilities in image synthesis [1]–[3]. These models employ dual Markov chains to learn data distributions through simulated diffusion and denoising processes. The generated images demonstrate superior quality compared to those synthesized by other generative approaches [4]–[6]. Models like Stable Diffusion [7], [8], trained on large and diverse datasets, have significantly advanced the field. These models demonstrate exceptional capability in generating detailed, contextually appropriate images, emerging as fundamental components of generative AI with broad applications.

Panoramic image generation [9]–[11] produces images with variable aspect ratios. This technology provides an expansive field of view, enhancing visual completeness and immersion. Despite significant research attention, this field faces several challenges, particularly the limited availability of training data. The scarcity of data hinders diffusion models from directly generating panoramic images.

Existing methods stitch together images generated by multiple diffusion models to tackle these challenges. These methods fall into two categories: image extrapolation [12], [13] and joint diffusion [14]–[16]. The first approach involves generating the final output by extrapolating the edges of an initial image. However, this method frequently produces repetitive patterns, leading to unrealistic panoramas. Joint

diffusion has become the leading method for creating seamless panoramic images. This approach integrates models with shared parameters or constraints, blending noisy images across overlapping regions by averaging the intermediate outputs at each denoising step. MultiDiffusion (MD) [9] exemplifies this approach and significantly advances panoramic image generation. However, the MD framework has limitations in generating high-resolution panoramas. The absence of global layout guidance leads to a disorganized spatial arrangement, compromising the final image quality (Fig. 1).

We introduce Multi-Scale Diffusion (MSD) as a solution to the challenges of generating high-resolution panoramic images. The method combines structure and composition from lower-resolution images with details from higher resolutions. The single-step denoising process is divided into multiple stages, using a phased approach that gradually improves the panoramic image’s quality. The framework extends the MultiDiffusion technique for joint denoising at each resolution level. MSD employs low-resolution images as structural guides through gradient descent optimization, effectively minimizing inconsistencies across resolution layers. As shown in Fig.1, our method successfully generates panoramic images with large-scale coherence and fine-grained detail.

The main contributions of our work are as follows:

- We identify integrating global layout information as essential for producing high-quality high-resolution panoramic images.
- Building on our theoretical findings, we propose the Multi-Scale Diffusion (MSD) framework. Our method incorporates spatial guidance via gradient descent, producing panoramas that are both structurally coherent and rich in detail.
- Comprehensive experiments validate the effectiveness of our approach, which surpasses baselines in both quantitative metrics and qualitative evaluations, solidifying it as a superior solution for panoramic image generation.

II. RELATED WORK

A. Diffusion Models

Diffusion models [2], [17]–[19] represent a class of generative models that create data through a two-step process: forward diffusion and reverse denoising. This process uses Markov chains to gradually perturb and reconstruct data distributions, enabling data creation from noise step-by-step.

This work was supported by the Lingyan Plan of Zhejiang Province under Grant No. 2025C02211.

*Yongchuan Tang is the corresponding author.

A painting in the Van Gogh style of the dolomites



Fig. 1. Comparison of high-resolution panoramic images generated by MultiDiffusion and our Multi-Scale Diffusion. MultiDiffusion excels in seamless image stitching but struggles with spatial coherence, leading to structural inconsistencies. In contrast, our model improves on this by integrating coarse structures and fine details from different resolution levels, producing panoramas that are both structurally coherent and visually detailed.

The field has advanced considerably since the emergence of DDPM [1]. These models have demonstrated remarkable success in image generation tasks [2], [20], outperforming previous approaches like GANs [5], [21], VAEs [6], [22] and Flows [4]. DDIM [3] further improved the process by introducing non-Markovian transitions that predict denoised data, significantly accelerating the denoising process. LDMs [7] marked another breakthrough by implementing the diffusion process in latent space using a pre-trained autoencoder, leading to high-performance systems like Stable Diffusion [7], [8] and DALL-E2 [23].

B. Panoramic Image Generation

The use of diffusion models for panoramic image generation has garnered significant research attention. Due to challenges in data acquisition and computational efficiency, directly generating panoramas presents numerous difficulties. Existing methods can be mainly divided into two categories: image extrapolation methods [12], [13], which extrapolate image edges, and approaches that integrate multiple diffusion paths to fuse overlapped denoising paths without additional training or fine-tuning [14]–[16].

The MultiDiffusion [9] exemplifies the second category of diffusion models, proving feasible and practical. It has laid a foundation for numerous subsequent studies. SCALE-CRAFTER [24] adapts diffusion models for higher resolutions by dilating convolution kernels, while Demofusion uses an “upsample-diffuse-denoise” loop to generate high-resolution ones progressively. However, these methods focus on generating high-resolution images rather than panoramic ones. SyncDiffusion [25] ensures consistency across panoramic images by calculating perceptual loss across windows. TwinDiffusion [26] achieves smoother transitions by aligning adjacent image parts. Nevertheless, these methods are limited to generating scene images repetitively in either the length or width

direction. When extending in both directions simultaneously, conflicting scene layouts in different windows result in a chaotic overall image layout.

To address this issue, we build upon existing methods by incorporating guidance from low-resolution images to capture structural details. This approach enables the generation of panoramas that are both rich in detail and structurally sound.

III. METHOD

A. Preliminary

1) *Latent Diffusion Model*: We introduce a pre-trained diffusion model operating in a latent space $\mathbb{R}^{c \times h \times w}$. The model generates image x_0 through iterative denoising, starting with initial Gaussian noise x_T . This process follows a predefined noise schedule, updating the current image x_t at each timestep t with the following formula:

$$x_{t-1} = \sqrt{\alpha_{t-1}} \left(\frac{x_t - \sqrt{1 - \alpha_t} \epsilon_\theta(x_t, t)}{\sqrt{\alpha_t}} \right) + \sqrt{1 - \alpha_{t-1}} \epsilon_\theta(x_t, t), \quad (1)$$

where α_t is parameterized by the noise schedule, $\epsilon_\theta(x_t, t)$ is the noise predicted by the denoising model at timestep t , parameterized by θ . For brevity, we denote the denoising steps as Φ in the rest of the paper:

$$x_{t-1} = \Phi(x_t), \quad (2)$$

2) *MultiDiffusion*: The MultiDiffusion [9] framework extends LDMs [7] by employing a multi-window joint diffusion technique. In this approach, the denoising process of the model Ψ is conducted a latent space $\mathbb{R}^{c \times H \times W}$ with $H > h$ and $W > w$. Initially, the panoramic image $z_t \in \mathbb{R}^{c \times H \times W}$ is cropped into a series of window images:

$$x_t^i = F_i(z_t), \quad (3)$$

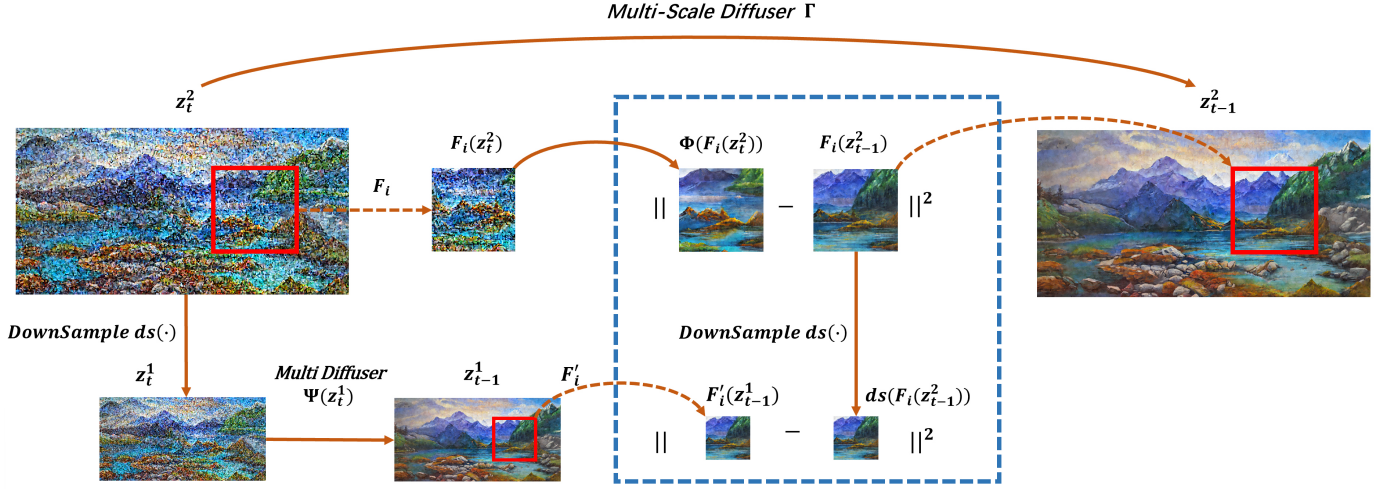


Fig. 2. Our MSD framework. The single-step denoising process is straightforward in multiple stages, progressively denoising the panoramic image from low to high resolution. Building on the pre-trained model, denoted as Φ , a new generation method, called Γ , is introduced. This method uses the results from the low-resolution denoising as constraints to guide the high-resolution denoising. The optimization objectives are twofold: (i) ensuring the final denoised image $F_i(\Gamma(z_t^2))$ closely matches the denoised results of each cropped window $\Phi(F_i(z_t^2))$; (ii) maintaining consistency between denoised images across different resolution layers (z_{t-1}^2, z_{t-1}^1). For the lowest-resolution image, our method simplifies to the standard MultiDiffusion approach, denoted by Ψ .

where $F_i(\cdot)$ refers to cropping the i -th image patch from image z_t .

Subsequently, each window undergoes independent denoising according to (2). The objective of MultiDiffuser is to ensure that $\Psi(z_t)$ aligns with $\Phi(x_t^i)$, $i \in [n]$. Therefore, the optimization problem is defined as follows:

$$z_{t-1} = \underset{\hat{z}_{t-1}}{\operatorname{argmin}} \mathcal{L}_{MD}(\hat{z}_{t-1} | z_t), \quad (4)$$

$$\mathcal{L}_{MD} = \sum_{i=1}^N W_i \otimes \|F_i(\hat{z}_{t-1}) - \Phi(F_i(z_t))\|^2, \quad (5)$$

where \hat{z}_{t-1} represents the variable to be determined in relation to z_t and W_i denote the weight matrix of the i -th window.

The framework employs a global least squares optimization to integrate the denoising results across all windows. The final image is then generated through weighted averaging:

$$z_{t-1} = \frac{\sum_i W_i \otimes F_i^{-1}(\Phi(x_t^i))}{\sum_i W_i}, \quad (6)$$

where $F_i^{-1}(\cdot)$ is the inverse function of $F_i(\cdot)$.

B. Multi-Scale Diffusion

MultiDiffusion effectively generates coherent panoramic images when limited to unidirectional expansion. However, the model exhibits significant limitations when attempting bidirectional expansion, manifesting in disorganized spatial arrangement and visual inconsistencies. These issues arise from the independent generation of panorama segments, which leads to misaligned features and repetitive patterns, as shown in Fig. 1. A lack of contextual awareness between different segments compromises spatial coherence, diminishing the panorama's

realism. This limitation highlights the need for an integrated approach to multi-dimensional panoramic generation.

To address this limitation, we propose a Multi-Scale Diffusion model that generates coherent, highly detailed panoramic images, as illustrated in Fig. 2. This framework integrates with existing joint diffusion systems. By extending the MultiDiffusion approach across multiple resolution layers, our model balances structural elements from lower resolutions with fine details from higher resolutions, enhancing overall image quality. We formulate the optimization problem as:

$$\begin{aligned} z_{t-1}^s &= \underset{\hat{z}_{t-1}^s}{\operatorname{argmin}} \mathcal{L}_{MSD}(\hat{z}_{t-1}^s | z_t^s, z_{t-1}^{s-1}) \\ &= \underset{\hat{z}_{t-1}^s}{\operatorname{argmin}} \mathcal{L}_{MD}(\hat{z}_{t-1}^s | z_t^s) + \omega \mathcal{L}_{MS}(\hat{z}_{t-1}^s | z_{t-1}^{s-1}), \end{aligned} \quad (7)$$

where z_t^s denotes the noise image at the s -th resolution, ω is the weight. The equation comprises two components: \mathcal{L}_{MD} represents the optimization objective of MultiDiffusion (5), while \mathcal{L}_{MS} ensures minimal deviation between denoised images across different resolution layers. The formula is as follows:

$$\mathcal{L}_{MS} = \sum_{i=1}^N W_i \otimes \|ds(F_i(\hat{z}_{t-1}^s)) - F'_i(z_{t-1}^{s-1})\|^2, \quad (8)$$

where $ds(\cdot)$ refer to the downsampling function (e.g., Bilinear Interpolation), and $F'_i(\cdot)$ is the crop function associated with $F_i(\cdot)$ that samples the corresponding region on the low-resolution panoramic image z_{t-1}^{s-1} . In contrast to MultiDiffusion, which derives an analytical solution for z_{t-1}^s by averaging features in overlapping regions, we employ back-propagation of gradients to obtain an approximate solution for z_{t-1}^s .

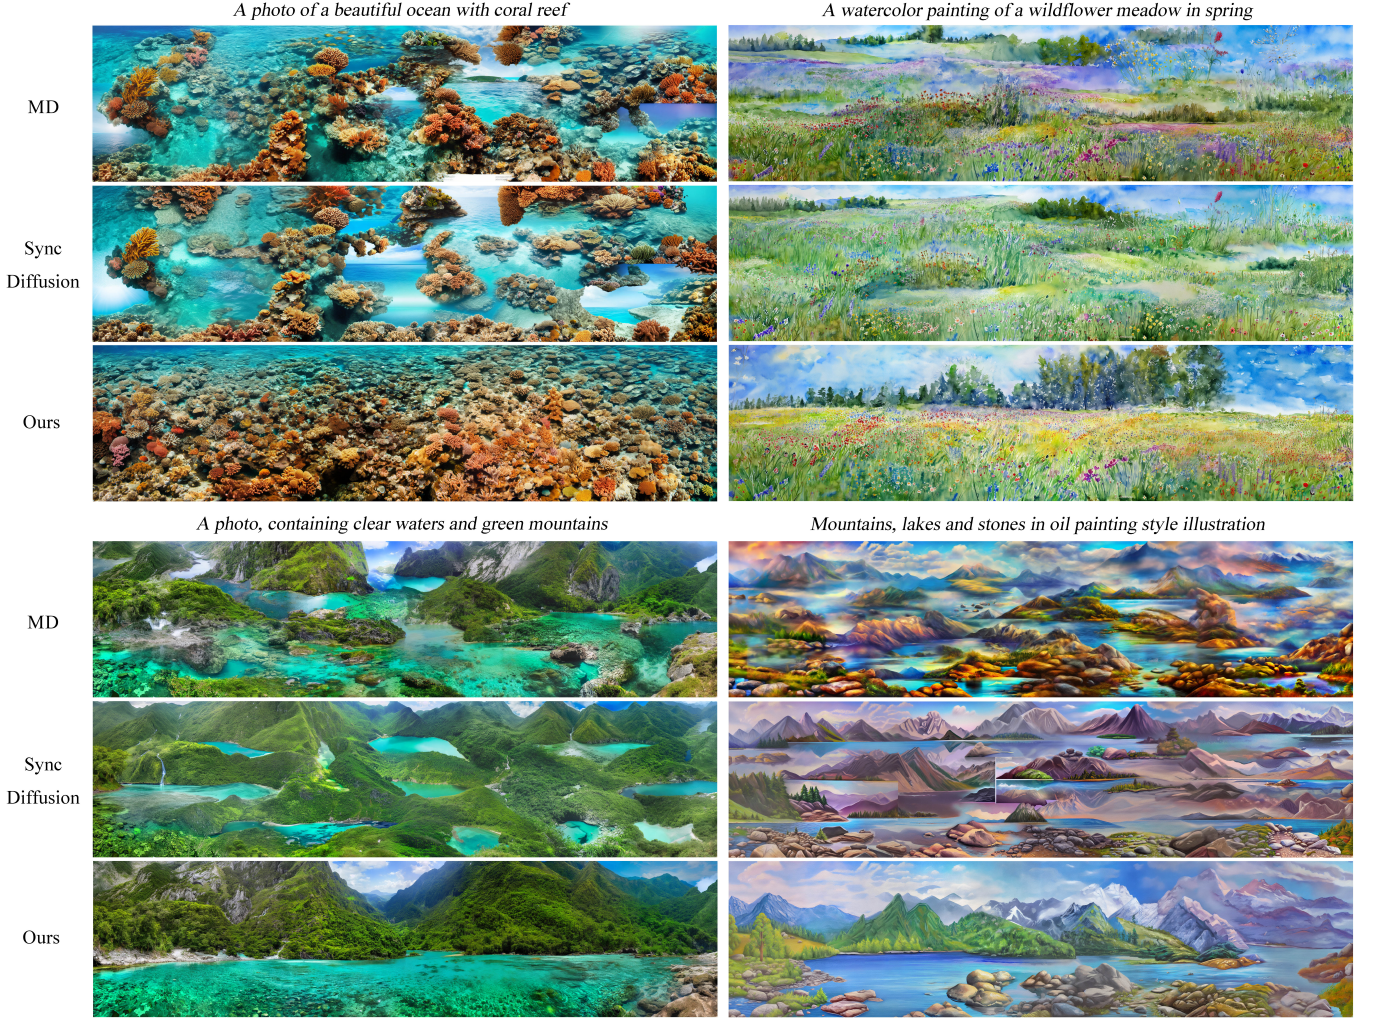


Fig. 3. Qualitative comparisons among MultiDiffusion, SyncDiffusion, Multi-Scale Diffusion. Our approach significantly improves spatial layout issues in high-resolution panoramic generation, producing semantically coherent and visually consistent results.

In the single-step denoising process of the noisy image z_t^S , we create a sequence of lower-resolution images through progressive downsampling, with z_t^0 representing the lowest resolution. The denoising process is divided into S stages, where Multi-Scale Diffusion is applied sequentially to the noisy image z_t^s across increasing resolutions, while the initial stage employs MultiDiffusion for denoising.

At each resolution level s (where $s > 1$), the MSD framework processes images through two parallel paths. In the first path, it applies the cropping function $F_i(\cdot)$ to the noisy image z_t^s to obtain a window image $x_{t,i}^s$, which is then denoised to produce $\Phi(x_{t,i}^s)$. In the second path, it applies a different cropping function $F'_i(\cdot)$ to the low-resolution panoramic image z_{t-1}^{s-1} , generating the corresponding window image $x_{t-1,i}^{s-1}$.

Theoretically, the window image $\Phi(x_{t,i}^s)$ that is first denoised and then downsampled should closely resemble the window image $x_{t-1,i}^{s-1}$ that is first downsampled and then denoised. The framework employs the MSE loss between these two window images as an objective function to measure their

correspondence. The gradient calculated from this objective function is then used in backpropagation to update the original window image $x_{t,i}^s$. The mathematical formulation is provided below:

$$\hat{x}_{t,i}^s = x_{t,i}^s - \omega \nabla_{x_{t,i}^s} \left\| ds(\Phi(x_{t,i}^s)) - x_{t-1,i}^{s-1} \right\|^2, \quad (9)$$

where ω is the weight of the gradient descent. MultiDiffusion is applied to denoise and merge these images at the end.

IV. EXPERIMENT

A. Experimental Setting

1) *Baselines*: We compare our MSD model with the following two baseline models: (1) MultiDiffusion [9], a special case of our method when the gradient weight $\omega = 0$; (2) SyncDiffusion [25], an extension of MultiDiffusion that improves global consistency through perceptual loss. For a fair comparison, we maintained consistent hyperparameters across all models and conducted all experiments using an A100 GPU.

2) *Experimental Setup*: For the reference model, we employ two variants: the Stable Diffusion v2.0 and the Stable Diffusion v1.5. Operating within a latent space of $\mathbb{R}^{4 \times 64 \times 64}$, this model generates images of $\mathbb{R}^{3 \times 512 \times 512}$. We aligned the crop window size with the model’s default resolution and generated panoramic images of resolution 1024×4096 (128×512 in the latent space). The panoramic images’ height is doubled, and the width is octupled compared to the original dimensions. Sec. IV-B presents results using specific parameters: gradient weight of $\omega = 10$ and scaled cosine decay factor of $(1 + \cos(\frac{T-t}{T} \times \pi)) / 2$. Sec. IV-C provides a detailed analysis of how these parameters affect image generation.

3) *Datasets*: We constructed our evaluation dataset following standard practices in the field [9], [25], [26]. The evaluation comprised fifteen diverse textual prompts spanning multiple themes and artistic styles. For each prompt, we generated 500 panoramic images and extracted multiple 1024×1024 pixel crops to create test datasets compatible with our evaluation metrics. We also developed a reference dataset using a benchmark model, generating 2,000 images per prompt.

B. Comparison

We conducted a comprehensive evaluation of our MSD model against baseline models, including both qualitative and quantitative analyses.

1) *Qualitative Comparison*: Fig. 3 compares images generated by our proposed method and two baseline models. While the MultiDiffusion model stitches scenes well, it produces disorganized layouts and unnatural features in high-resolution panoramic images. The SyncDiffusion model improves stylistic consistency but suffers from spatial structure issues like floating mountains and unsupported lakes. These results demonstrate that merely optimizing window connections and maintaining style consistency alone is insufficient for high-quality panoramic image generation.

The proposed MSD method demonstrates superior capability in generating high-resolution panoramic images that maintain visual and semantic coherence across diverse textual prompts. The qualitative assessment of generated images reveals a robust correlation between the spatial relationships specified in the input prompts and the resultant visual compositions, demonstrating the model’s efficacy in transforming textual guidance effectively into coherent spatial structures.

2) *Quantitative Comparison*: We evaluate the generated panoramas using multiple quantitative metrics. These metrics assess the panoramas’ fidelity and diversity, as well as their adherence to the input prompts.

- **Fidelity & Diversity**: **FID** and **KID** are employed to assess the fidelity and diversity of the generated images. Both metrics evaluate the distribution of generated images against reference images, with FID calculating feature vector distances and KID utilizing a kernel-based approach.

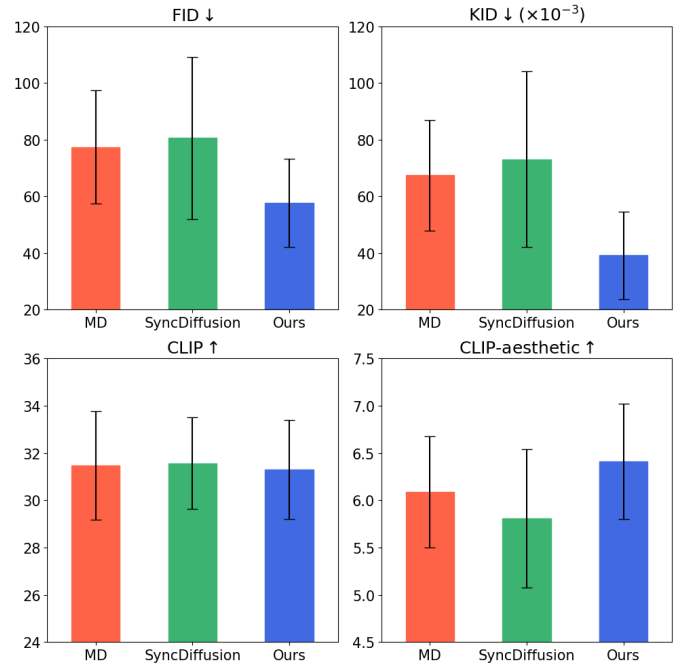


Fig. 4. Quantitative results comparing four key metrics across two dimensions. Our method achieves superior performance, particularly in image quality, establishing it as the most effective approach for high-resolution panoramic image generation.

- **Compatibility**: **CLIP** is used to evaluate the alignment between generated images and the input prompts by calculating cosine similarity, while **CLIP-aesthetic** is used to quantify the aesthetic quality of the images using a linear estimator.

The quantitative comparison illustrated in Fig. 4 indicates that our MSD approach outperforms the baselines across four critical evaluation metrics. Our method substantially improves panorama image quality, as reflected by lower FID and KID scores, indicating that our generated images more closely align with the reference distribution. The CLIP-Aesthetic scores further corroborate these improvements, with our method achieving the highest ratings. Moreover, our method maintains strong prompt adherence, performing on par with existing methods in terms of CLIP metrics. These comprehensive results demonstrate that our method effectively balances aesthetic quality with prompt relevance, suggesting its viability for high-fidelity panorama generation tasks.

C. Ablation Study

The gradient weight ω in (9) is crucial for optimizing the spatial layout of high-resolution panoramic images. This parameter controls how spatial information from low-resolution images influences their high-resolution counterparts, directly affecting the final image quality. Our experiments revealed significant improvements in FID scores. Given the importance of FID in assessing image quality, we focused our evaluation primarily on this key metric.

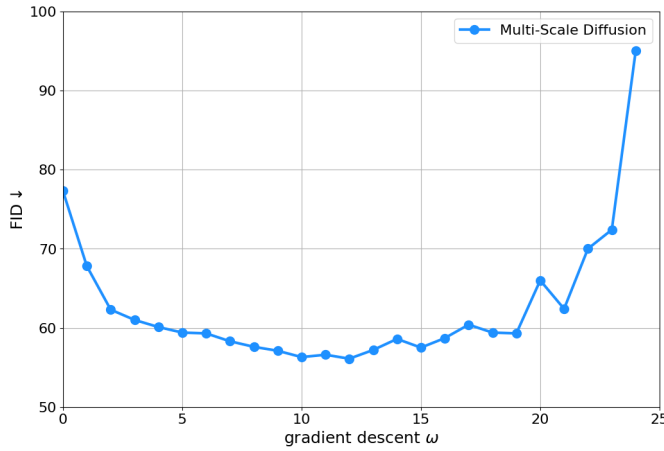


Fig. 5. Impact of gradient weight ω on fidelity and diversity in high-resolution image generation. Optimal balance is achieved when $\omega = 10$. Lower ω values result in insufficient spatial layout optimization, while higher values may lead to image structure collapse.

Fig. 5 illustrates the impact of varying the gradient weight ω from 0 to 24. At values of $\omega \leq 2$, the FID metric shows a rapid decrease, indicating significant improvements in image layout. When $\omega \geq 20$ increases the FID metric, suggesting that excessive optimization leads to degraded image generation and reduced quality. We found that $\omega = 10$ yields optimal results, producing high-resolution images that balance spatial layout and fine detail best.

V. CONCLUSION

The Multi-Scale Diffusion is a versatile framework that enhances the generation of high-resolution panoramic images. By operating across multiple resolution levels, it uses information from lower-resolution images to guide the refinement of higher-resolution outputs through gradient descent optimization. This approach produces panoramas that maintain both global structural coherence and fine-scale detail. Our empirical evaluations demonstrate that MSD performs better than baseline methods in quantitative and qualitative assessments. Beyond static images, the model shows potential for video generation, which we aim to explore in future research.

REFERENCES

- [1] Jonathan Ho, Ajay Jain, and Pieter Abbeel, “Denoising diffusion probabilistic models,” *Advances in neural information processing systems*, vol. 33, pp. 6840–6851, 2020.
- [2] Prafulla Dhariwal and Alexander Nichol, “Diffusion models beat gans on image synthesis,” *Advances in neural information processing systems*, vol. 34, pp. 8780–8794, 2021.
- [3] Jiaming Song, Chenlin Meng, and Stefano Ermon, “Denoising diffusion implicit models,” *arXiv preprint arXiv:2010.02502*, 2020.
- [4] Ivan Kobyzev, Simon JD Prince, and Marcus A Brubaker, “Normalizing flows: An introduction and review of current methods,” *IEEE transactions on pattern analysis and machine intelligence*, vol. 43, no. 11, pp. 3964–3979, 2020.
- [5] Tero Karras, Samuli Laine, and Timo Aila, “A style-based generator architecture for generative adversarial networks,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019, pp. 4401–4410.
- [6] Aaron Van Den Oord, Oriol Vinyals, et al., “Neural discrete representation learning,” *Advances in neural information processing systems*, vol. 30, 2017.
- [7] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer, “High-resolution image synthesis with latent diffusion models,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 10684–10695.
- [8] Dustin Podell, Zion English, Kyle Lacey, Andreas Blattmann, Tim Dockhorn, Jonas Müller, Joe Penna, and Robin Rombach, “Sdxl: Improving latent diffusion models for high-resolution image synthesis,” *arXiv preprint arXiv:2307.01952*, 2023.
- [9] Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel, “Multidiffusion: Fusing diffusion paths for controlled image generation,” in *International Conference on Machine Learning*, 2023.
- [10] Mengyang Feng, Jinlin Liu, Miaomiao Cui, and Xuansong Xie, “Diffusion360: Seamless 360 degree panoramic image generation based on diffusion models,” *arXiv preprint arXiv:2311.13141*, 2023.
- [11] Haiyang Zhou, Xinhua Cheng, Wangbo Yu, Yonghong Tian, and Li Yuan, “Holodreamer: Holistic 3d panoramic world generation from text descriptions,” *arXiv preprint arXiv:2407.15187*, 2024.
- [12] Omri Avrahami, Dani Lischinski, and Ohad Fried, “Blended diffusion for text-driven editing of natural images,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2022, pp. 18208–18218.
- [13] Omri Avrahami, Ohad Fried, and Dani Lischinski, “Blended latent diffusion,” *ACM transactions on graphics (TOG)*, vol. 42, no. 4, pp. 1–11, 2023.
- [14] Qinsheng Zhang, Jiaming Song, Xun Huang, Yongxin Chen, and Ming-Yu Liu, “Diffcollage: Parallel generation of large content with diffusion models,” in *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, 2023, pp. 10188–10198.
- [15] Luming Tang, Menglin Jia, Qianqian Wang, Cheng Perng Phoo, and Bharath Hariharan, “Emergent correspondence from image diffusion,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 1363–1389, 2023.
- [16] Álvaro Barbero Jiménez, “Mixture of diffusers for scene composition and high resolution image generation,” *arXiv preprint arXiv:2302.02412*, 2023.
- [17] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *International conference on machine learning*. PMLR, 2015, pp. 2256–2265.
- [18] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole, “Score-based generative modeling through stochastic differential equations,” *arXiv preprint arXiv:2011.13456*, 2020.
- [19] Alexander Quinn Nichol and Prafulla Dhariwal, “Improved denoising diffusion probabilistic models,” in *International conference on machine learning*. PMLR, 2021, pp. 8162–8171.
- [20] 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,” *Advances in neural information processing systems*, vol. 35, pp. 36479–36494, 2022.
- [21] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila, “Analyzing and improving the image quality of stylegan,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020, pp. 8110–8119.
- [22] Diederik P Kingma and Max Welling, “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*, 2013.
- [23] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen, “Hierarchical text-conditional image generation with clip latents,” *arXiv preprint arXiv:2204.06125*, vol. 1, no. 2, pp. 3, 2022.
- [24] 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 *The Twelfth International Conference on Learning Representations*, 2023.
- [25] Yuseung Lee, Kunho Kim, et al., “Syncdiffusion: Coherent montage via synchronized joint diffusions,” *Advances in Neural Information Processing Systems*, vol. 36, pp. 50648–50660, 2023.
- [26] Teng Zhou and Yongchuan Tang, “Twindiffusion: Enhancing coherence and efficiency in panoramic image generation with diffusion models,” *arXiv preprint arXiv:2404.19475*, 2024.

VI. APPENDIX

A. Algorithm

As described Section III-B, we introduce the Multi-Scale Diffusion. Here we try to provide the algorithms in details.

Algorithm 1 Pseudocode of one-time denoising in Multi-Scale Diffusion.

Input: z_t^S ▷ Noisy images at timestep t
Parameter: ω ▷ Gradient descent weight
Output: z_{t-1}^S ▷ Noisy images at timestep $t-1$

- 1: **function** MULTI-SCALE DIFFUSER(z_t^s, z_{t-1}^{s-1})
- 2: **for** $i = 1 \rightarrow N$ **do**
- 3: $x_{t,i}^s \leftarrow F_i(z_t^s)$ ▷ Crop window from panoramic image
- 4: $x_{t-1,i}^{s-1} \leftarrow F_i'(z_{t-1}^{s-1})$
- 5: $\hat{x}_{t,i}^s = x_{t,i}^s - \omega \nabla_{x_{t,i}^s} \|ds(\Phi(x_{t,i}^s)) - x_{t-1,i}^{s-1}\|^2$ ▷ Gradient descent (Eq. 9)
- 6: **end for**
- 7: **return** $\{\hat{x}_{t,i}^s\}$
- 8: **end function**
- 1: **function** DENOISING ONE-STEP(z_t^S)
- 2: **for** $s = S \rightarrow 2$ **do** ▷ Downsample
- 3: $z_t^{s-1} \leftarrow ds(z_t^s)$
- 4: **end for**
- 5: $z_{t-1}^1 \leftarrow \frac{\sum_i W_i \otimes F_i^{-1}(\Phi(F_i(z_t^1)))}{\sum_i W_i}$ ▷ Apply MultiDiffusion at the lowest-resolution layer (Eq. 6)
- 6: **for** $s = 2 \rightarrow S$ **do**
- 7: $\{\hat{x}_{t,i}^s\} \leftarrow \text{MULTI-SCALE DIFFUSER}(z_t^s, z_{t-1}^{s-1})$
- 8: $z_{t-1}^s \leftarrow \frac{\sum_i W_i \otimes F_i^{-1}(\Phi(\hat{x}_{t,i}^s))}{\sum_i W_i}$ ▷ Merge images
- 9: **end for**
- 10: **return** z_{t-1}^S
- 11: **end function**

B. More Qualitative Results

In this section, we provide more qualitative comparisons between MultiDiffusion, SyncDiffusion, and our Multi-Scale Diffusion on Stable Diffusion v2.0 and v1.5.

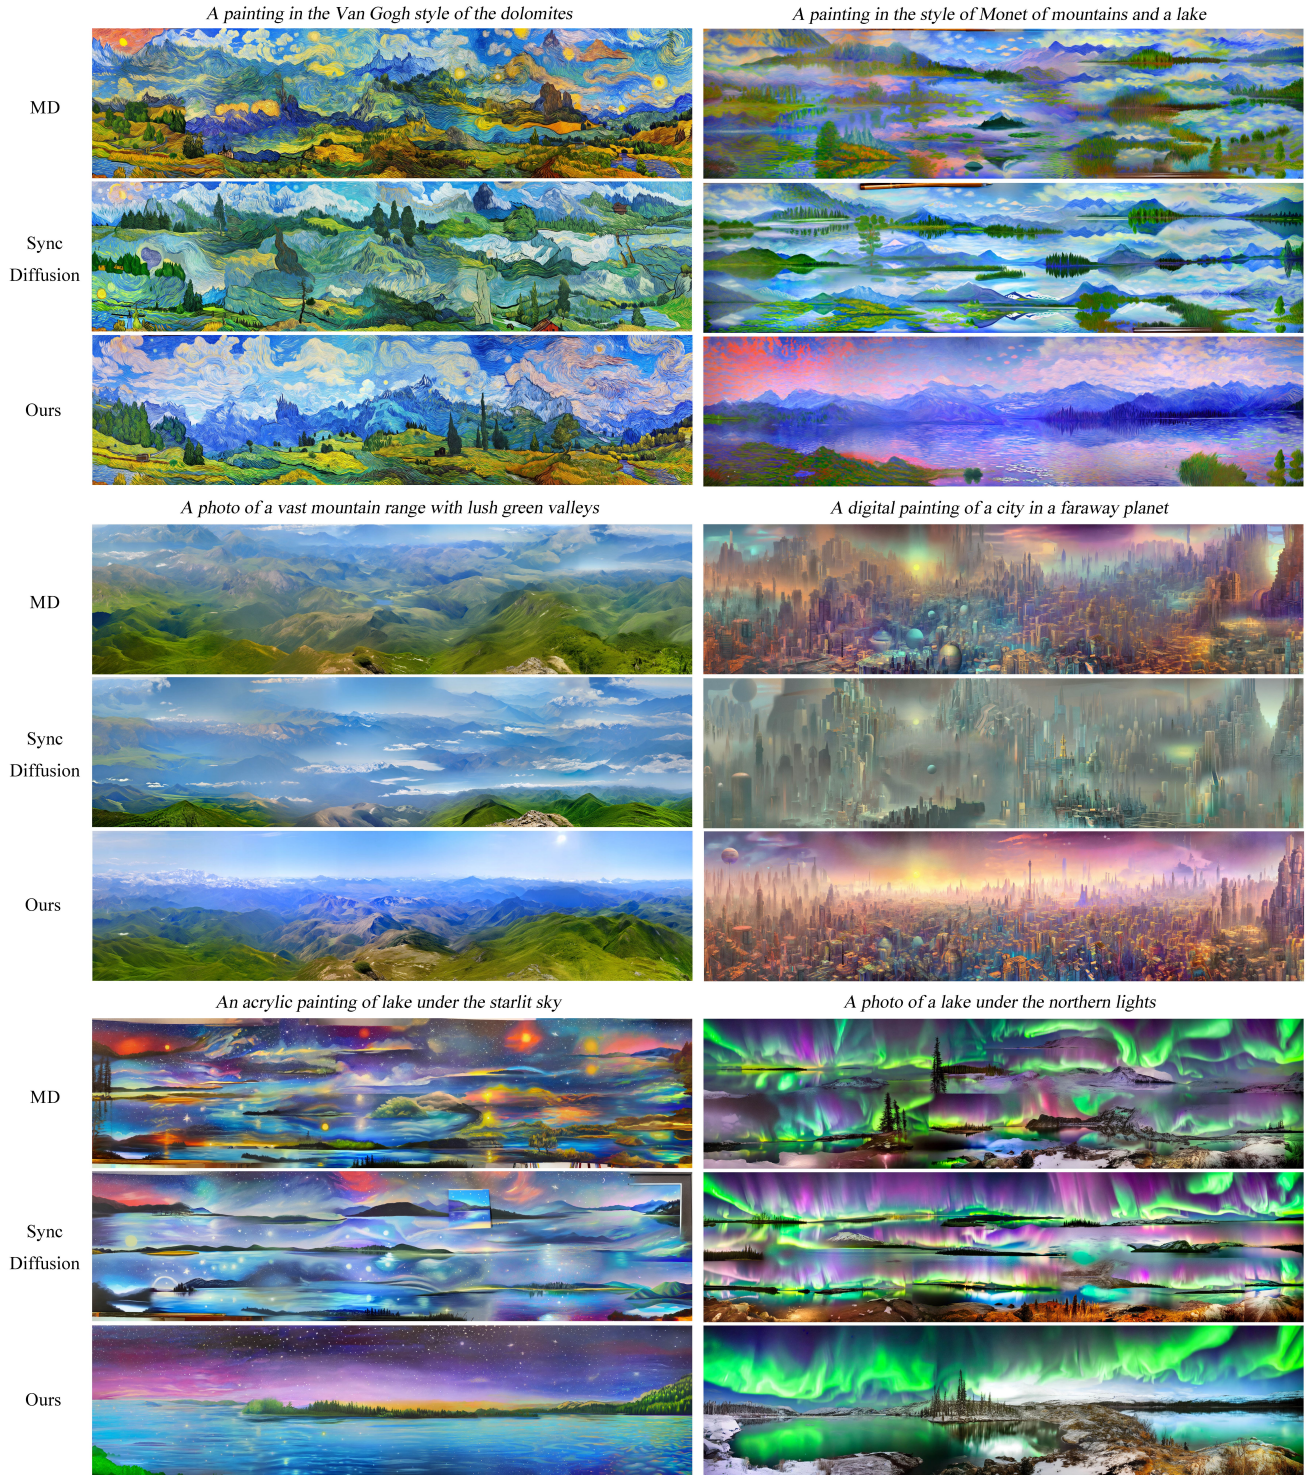


Fig. 6. Additional qualitative comparison results for panorama generation using various text prompts on Stable Diffusion v2.0.

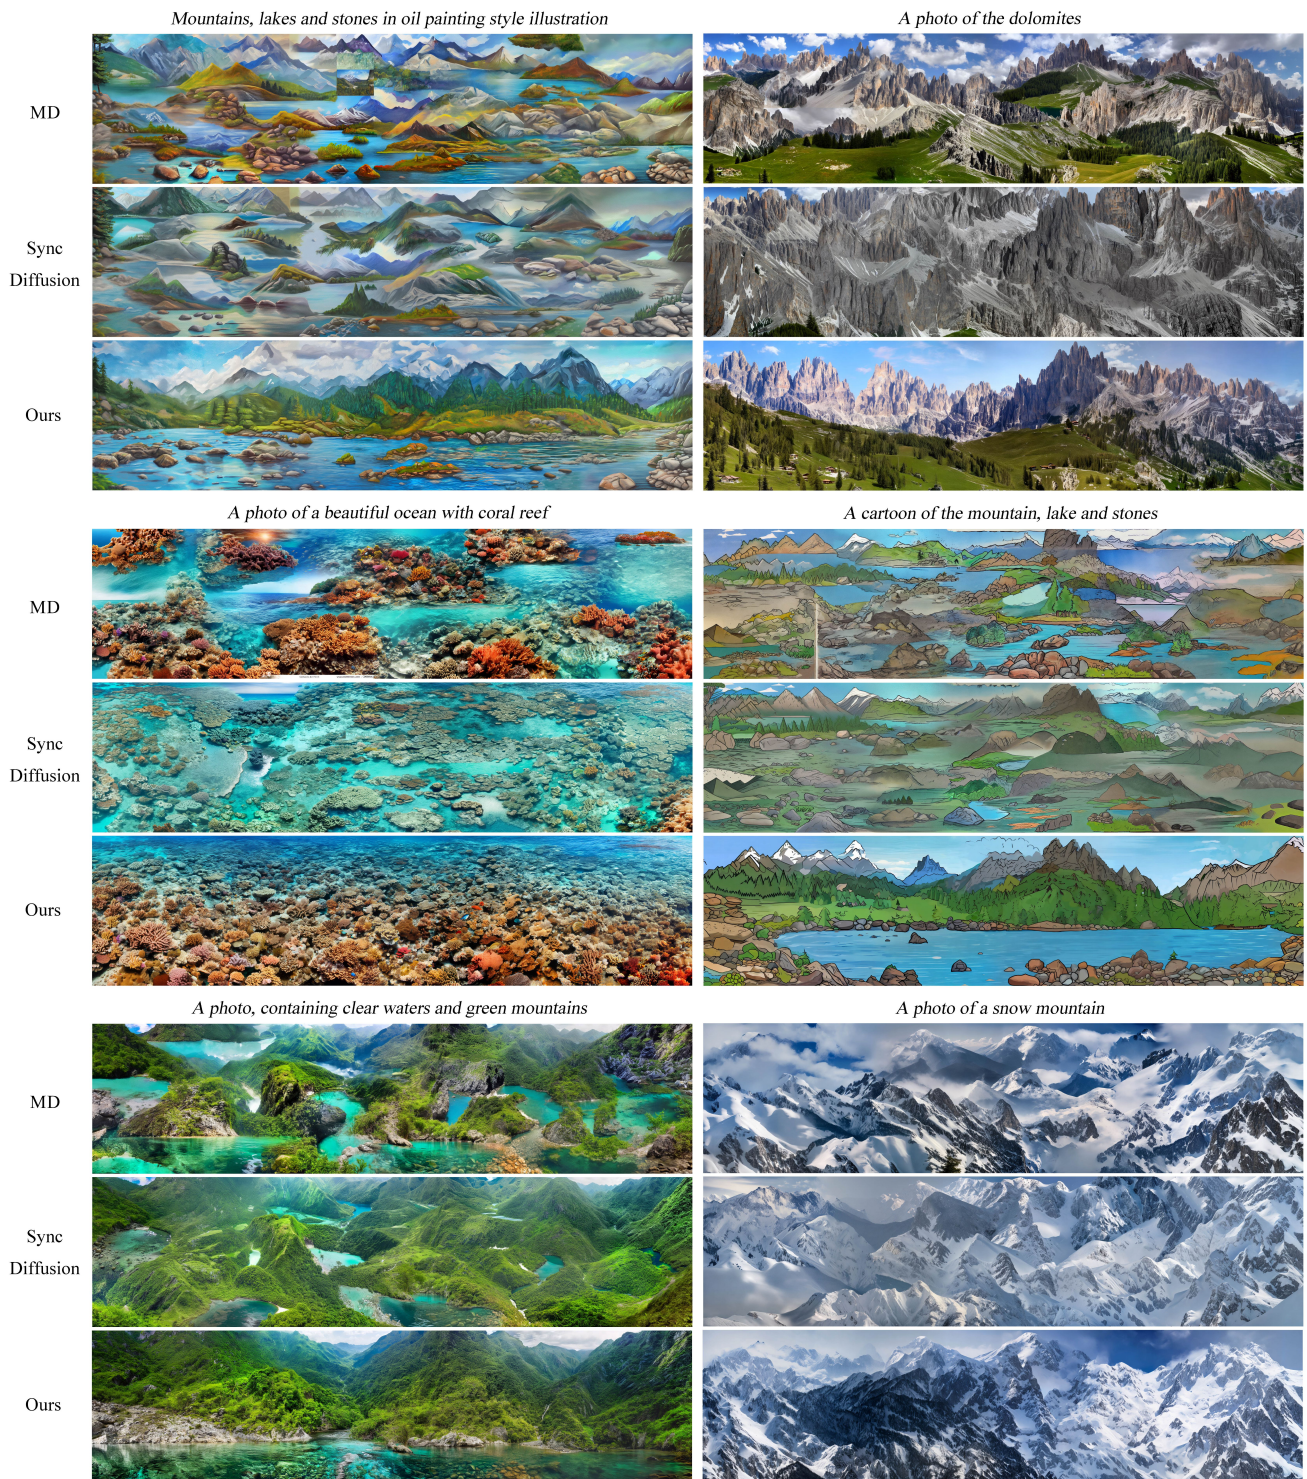


Fig. 7. Additional qualitative comparison results for panorama generation using various text prompts on Stable Diffusion v1.5.