

SpotDiffusion: A Fast Approach For Seamless Panorama Generation Over Time

Stanislav Frolov Brian B. Moser Andreas Dengel

German Research Center for Artificial Intelligence (DFKI), Germany
 RPTU Kaiserslautern-Landau, Germany
 first.last@dfki.de

Abstract

Generating high-resolution images with generative models has recently been made widely accessible by leveraging diffusion models pre-trained on large-scale datasets. Various techniques, such as MultiDiffusion and SyncDiffusion, have further pushed image generation beyond training resolutions, i.e., from square images to panorama, by merging multiple overlapping diffusion paths or employing gradient descent to maintain perceptual coherence. However, these methods suffer from significant computational inefficiencies due to generating and averaging numerous predictions, which is required in practice to produce high-quality and seamless images. This work addresses this limitation and presents a novel approach that eliminates the need to generate and average numerous overlapping denoising predictions. Our method shifts non-overlapping denoising windows over time, ensuring that seams in one timestep are corrected in the next. This results in coherent, high-resolution images with fewer overall steps. We demonstrate the effectiveness of our approach through qualitative and quantitative evaluations, comparing it with MultiDiffusion, SyncDiffusion, and StitchDiffusion. Our method offers several key benefits, including improved computational efficiency and faster inference times while producing comparable or better image quality.

1. Introduction

The recent breakthrough of generative diffusion models [3, 17, 18] has enabled the generation of high-quality images with high fidelity and diversity from simple text prompts. These models are typically trained on a large-scale dataset of fixed-size resolution. Training such models for high-resolution image generation requires substantial computational resources and large datasets. Therefore, several attempts have been made to repurpose pre-trained diffusion models to exceed the original training image resolution. When initially trained on lower resolutions, innovative techniques are required to overcome their original constraints

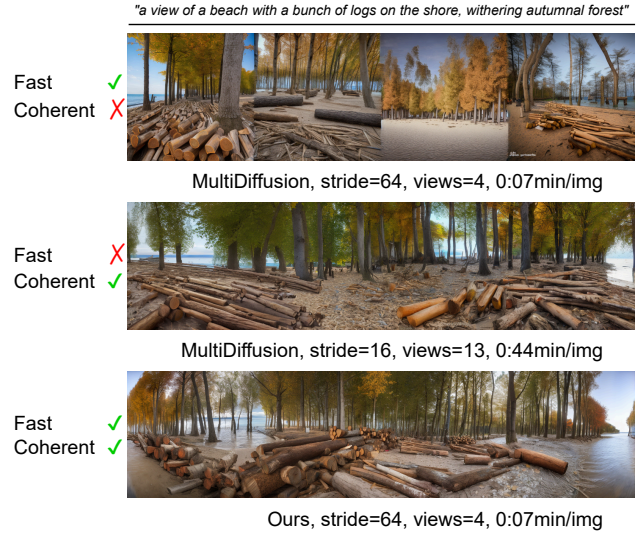


Figure 1. MultiDiffusion [1] can produce coherent panorama images by averaging overlapping denoising predictions. However, this process introduces computational inefficiencies and requires denoising many patch views. Without overlapping views, MultiDiffusion can not produce coherent panoramas. We introduce an efficient method for high-resolution panorama generation that eliminates the need for overlapping denoising predictions, resulting in coherent and sharp images without border artifacts.

(fixed spatial size) and generate higher-resolution images. Previous methods, such as MultiDiffusion [1] and SyncDiffusion [12], have made significant strides in this domain. MultiDiffusion [1] achieves high resolution by fusing multiple overlapping diffusion paths, which are then averaged to create the final image. While effective, this technique introduces computational inefficiencies due to the need to generate and average numerous predictions. SyncDiffusion [12] builds on this by using gradient descent from a perceptual similarity loss to produce coherent panoramas, yet it still fundamentally relies on the principles of MultiDiffusion. A major drawback of these current methods is the high com-

computational cost required to create overlapping predictions during denoising and combine them through averaging. For example, in practice, MultiDiffusion [1] requires an overlap of 75% between adjacent windows, resulting in a large number of denoising predictions and thus significantly increased computational demands.

In this paper, we present our novel approach, coined SpotDiffusion (seamless panorama over time), that addresses this limitation and offers a more efficient and effective method for high-resolution image generation. We argue that many of the denoising predictions, like in MultiDiffusion [1], are redundant and that averaging can decrease the quality of the final image, in addition to incurring significant computational costs. Furthermore, we identify that the key to producing seamless panoramas is to ensure that the full denoising process is applied uniformly over time across the entire image. Our key insight is that this can be achieved without overlapping denoising predictions within one timestep by simply shifting the denoising windows over time. This ensures that any seam in one timestep is corrected in the next, resulting in coherent, high-resolution images with significantly fewer steps. SpotDiffusion makes several key contributions:

- it introduces a fast method for coherent and sharp panorama image generation without border artifacts.
- it eliminates the need for generating numerous overlapping denoising predictions and subsequent averaging, significantly reducing the computational complexity of the denoising process.
- it can serve as a drop-in replacement for existing diffusion models that previously relied on the MultiDiffusion mechanism, making it a practical and efficient solution for high-resolution image generation.

2. Related Work

This section provides an overview of the related work in the field of high-resolution image generation, focusing on diffusion models, text-to-image synthesis, and high-resolution montage generation.

2.1. Diffusion Models

Diffusion models [9, 18–20] are powerful generative models. A key difference between them and earlier generative models, such as GANs [6], is their iterative process. During training, diffusion models employ a forward and backward process. The forward process starts from a clean image and iteratively adds small amounts of noise until the image is indistinguishable from random noise. In the backward process, they learn to generate real-looking data by progressively reverting the forward process, i.e., removing the added noise. As a result, they can gradually approximate a complex target data distribution from a normal noise distribution [3, 21]. These models have been shown to produce high-quality images and have been used in various applications, including

image generation, inpainting, super-resolution, and across various data modalities, including audio, video, and 3D objects, showcasing their broad applicability [2, 4, 13, 14, 24].

2.2. Text-to-Image Synthesis

Text-to-image [5, 25] synthesis has gained significant attention for its impressive generation performance. Diffusion models, particularly latent diffusion models [18], have become popular due to their high-quality generation capabilities. Stable Diffusion, built upon latent diffusion models [18], has been shown to produce high-resolution images with high fidelity and diversity making it a popular choice for text-to-image synthesis and enabling researchers as well as artists to generate creative images from text prompts. Despite their strengths, these models are limited by their training resolution, necessitating innovative techniques to exceed their original training limits and produce high-resolution outputs.

2.3. High-Resolution Image Generation

High-resolution image synthesis presents significant challenges due to the high-dimensional data and substantial computational resources required. Previous approaches include training from scratch or fine-tuning, which can be computationally expensive and time-consuming. Training-free methods like MultiDiffusion [1] repurpose pre-trained models on low-resolution to generate high-resolution images. To that end, multiple overlapping diffusion paths are fused by averaging the denoising predictions to produce seamless images. SyncDiffusion [12] improves upon these methods by synchronizing multiple diffusions through gradient descent, achieving more coherent outputs. However, these methods introduce a significant computational overhead due to the need for generating and averaging numerous predictions, especially because large overlaps are required to produce seamless images. Our work aims to address these limitations by introducing a novel approach that eliminates the need for overlapping denoising predictions and subsequent averaging by shifting non-overlapping denoising windows over time.

3. Methodology

We propose SpotDiffusion, a new diffusion process for fast panorama image generation. Contrary to previous methods, SpotDiffusion uses randomly shifted, non-overlapping windows through time to ensure coherent transitions. This section details the mathematical formulation of our method and presents the algorithm for seamless panorama generation.

3.1. Preliminaries

Consider a pre-trained diffusion model Φ that operates in image space $\mathcal{I} \in \mathbb{R}^{W \times H \times C}$ and condition space \mathcal{Y} . Given a noisy image $I_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and a condition $y \in \mathcal{Y}$, the

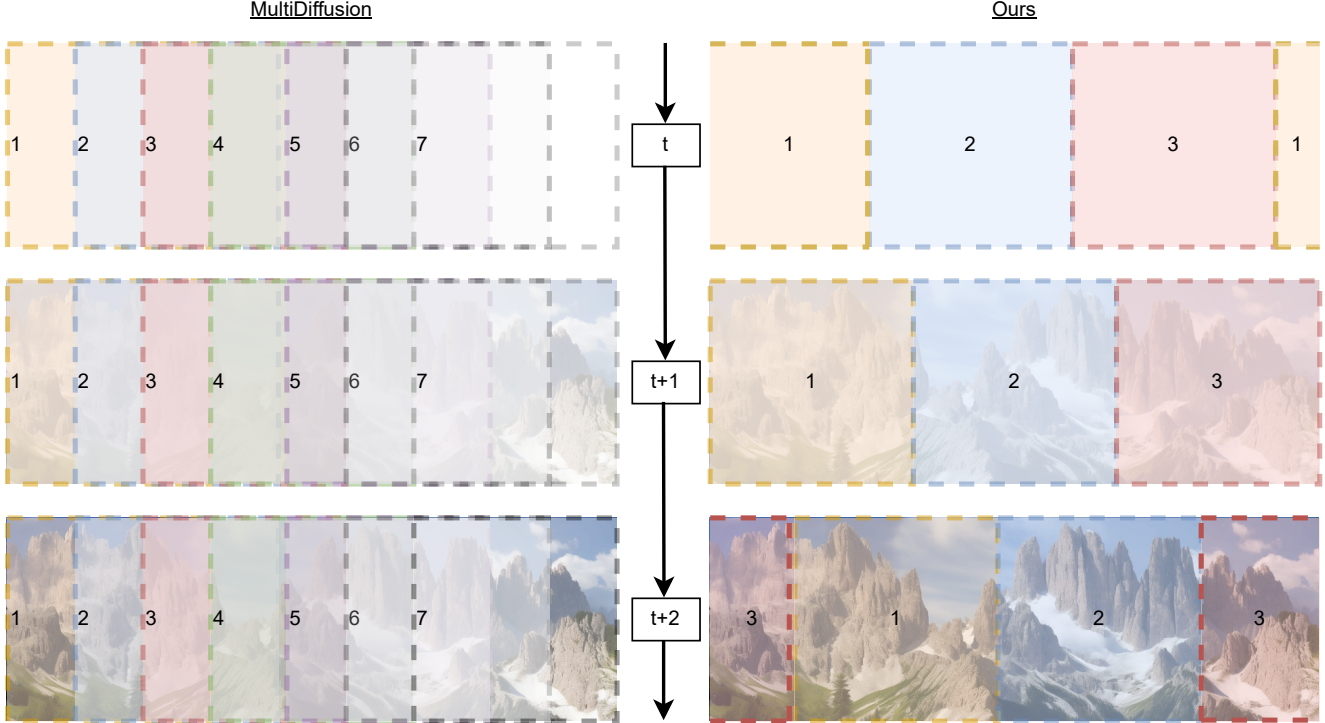


Figure 2. MultiDiffusion [1] generates coherent panorama images by averaging overlapping denoising predictions with a stride that is smaller than the denoising window. Our method eliminates the need for overlapping denoising predictions and introduces a more efficient shifting denoising method. Instead of relying on a fixed denoising path with overlapping views, our method shifts the denoising windows over time, ensuring that seams in one timestep are corrected in the next. This results in fast, seamless, high-resolution images with fewer overall steps.

model produces a sequence of images starting from I_T and gradually denoising it towards the clean image I_0 :

$$I_T, I_{T-1}, \dots, I_0 \quad \text{such that} \quad I_{t-1} = \Phi(I_t | y) \quad (1)$$

Our goal is to generate images J_T, J_{T-1}, \dots, J_0 in a new image space $\mathcal{J} \in \mathbb{R}^{W' \times H' \times C}$, with $W' \geq W, H' \geq H$, using the same reference model Φ without any re- or fine-tuning. Traditionally, models pre-trained on fixed-size images cannot be directly used to produce images of arbitrary sizes. To that end, MultiDiffusion [1] addressed this limitation by using a joint diffusion approach where multiple overlapping windows are merged via averaging. More formally, n mappings $F_i : \mathcal{J} \in \mathbb{R}^{W' \times H' \times C} \rightarrow \mathcal{I} \in \mathbb{R}^{W \times H \times C}$ with $i \in \{1, 2, \dots, n\}$ are defined that map (crop) the image space \mathcal{J} into n images of the original space \mathcal{I} . The value of n , and thus the number of individual image crops, is defined as $n = \frac{W'-W}{\omega} + 1$, where ω is the stride between adjacent cropping windows. With these mappings, the denoising process is applied to each cropped image, and the results are averaged to produce the final image J_0 .

Notably, the stride ω is typically chosen such that the cropped windows overlap, ensuring that the denoising process is applied uniformly across the entire image. In practice, it is common to choose $\omega = \frac{W}{4}$ or even $\omega = \frac{W}{8}$ such that

the cropped windows overlap by 75% or 87.5%, respectively to produce seamless transitions. However, this approach dramatically increases the computational complexity of the denoising process, as each overlapping window requires a denoising prediction. Strides $\omega \geq \frac{W}{2}$ are known to produce visible artifacts and seams in the final image. In the extreme case of $\omega = W$, the cropped windows are disjoint, and the denoising process is applied independently to each window, resulting in a disjoint image.

3.2. SpotDiffusion

Our goal is to efficiently generate high-resolution, seamless images and speed up the total inference time. Thus, we aim to eliminate the need for many overlapping denoising windows. We hypothesize that many of the denoising predictions are redundant and that averaging may decrease the quality or alignment of the final image. To that end, we identify that the key to producing seamless panoramas is to ensure that the denoising process is applied uniformly across the entire image over time. We recognize that the static mappings F_i in MultiDiffusion [1] are fixed and do not change over time, which we believe is a critical limitation of the method. Therefore, in this paper, we propose a shifted windows diffusion approach. Specifically, we use

time-dependent mappings $F_{(i,t)}$ using $s(t)$, which applies random shifts of size $s(t)$ at every timestep to the mapping such that the entire image is denoised uniformly. Our core insight is that using non-overlapping windows that shift randomly over time effectively corrects discontinuities at the seams in subsequent time steps. Given sufficient timesteps, our shifted windows ensure that each pixel is processed through various windows, achieving a uniformly denoised and perfectly coherent image. See Figure 2 for a conceptual illustration of our method and comparison to MultiDiffusion [1].

More formally, for each timestep t , a shift function with a shift size of $s(t)$ is applied to the mappings $F_{(i,t)}$. The shift size $s(t) \sim \mathcal{U}(0, W)$ is drawn from a uniform distribution over the entire mapping sequence, ensuring that each pixel in the image \mathcal{J} is included in multiple windows over time. In our method, the stride equals the width of the denoising window $\omega = W$, so the windows are non-overlapping. To handle cases where a window mapping exceeds the image boundaries, $F_{(i,t)}$ maps it back to the beginning of the image (wrap-around). In practice, our method can efficiently be implemented by shifting the input image J_t of the mapping functions F_i with a `translate` operator that incorporates a wrap-around, as shown in Algorithm 1. The denoised image J_{t-1} is obtained by concatenating the denoised windows and reverting the shift by re-applying `translate` with $-s(t)$. Due to the random shifts, any artifacts or seams produced in one timestep are corrected in subsequent timesteps, leading to a coherent and seamless final image.

The key advantage of our method lies in its ability to produce coherent transitions without the need for overlapping windows, which simplifies the method and significantly reduces the time complexity of the denoising process. Furthermore, our approach can be easily integrated into existing diffusion models, especially as a drop-in replacement for MultiDiffusion [1], without the need for retraining or fine-tuning, making it a practical and efficient solution for high-resolution image generation.

4. Experimental Setup

We choose three models to evaluate our method. First, we compare our approach with MultiDiffusion [1] using the Stable Diffusion 2.0 backbone. The Stable Diffusion model operates in a latent space of $\mathbb{R}^{64 \times 64 \times 4}$ and generates images of $\mathbb{R}^{512 \times 512 \times 3}$. We use it to generate panorama images of resolution 512×2048 (64×256 in the latent space), where the width is four times the width of the output of Stable Diffusion. Each window crop in the panorama has an image resolution of 512×512 . We compare against MultiDiffusion [1] with various strides along the width in the image space. In the case of no overlaps, the stride equals the window size of 64 resulting in a constant compute comparison with our method. Second, we apply our method as a

Algorithm 1 SpotDiffusion

Require: Φ {pre-trained diffusion model}
 $\{F_i\}_{i=1}^n$ {non-overlapping mappings}
 $\{y_i\}_{i=1}^n$ {conditions for each window}
 $J_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ {noisy initialization}
`translate`(\cdot, c) {translate by c function}


- 1: **for** $t = T, T - 1, \dots, 0$ **do**
- 2: $s(t) \sim \mathcal{U}(0, W)$ {sample a random shift size}
- 3: $\hat{J}_t \leftarrow \text{translate}(J_t, s(t))$ {shift panorama}
- 4: **for** each window $i = 1, \dots, n$ **do**
- 5: $I_{t-1}^i \leftarrow \Phi(F_i(\hat{J}_t) \mid y_i)$ {denoise window}
- 6: **end for**
- 7: $\hat{J}_{t-1} \leftarrow \text{concat}(\{I_{t-1}^i\}_{i=1}^n)$ {combine windows}
- 8: $J_{t-1} \leftarrow \text{translate}(\hat{J}_{t-1}, -s(t))$ {revert shift}
- 9: **end for**
- 10: **return** J_0

drop-in replacement for the inner MultiDiffusion [1] loop in SyncDiffusion [12]. SyncDiffusion [12] synchronizes multiple diffusions through gradient descent to produce coherent panoramas, but requires subsequent MultiDiffusion [1] to merge all predictions. We replace it with SpotDiffusion to speed up the image generation process. Finally, we compare our method with StitchDiffusion [22], a method for spherical image synthesis, which generates 360-degree panoramas by averaging overlapping denoising predictions. StitchDiffusion [22] customizes a pre-trained T2I diffusion model for 360-degree panorama synthesis by fine-tuning a Low-Rank Adaptation (LoRA) [10] matrices. Here, we also replace the overlapping denoising predictions with our method to demonstrate the efficiency of our approach and show the qualitative results and image generation time.

4.1. Dataset & Metrics

To compare with MultiDiffusion [1] and SyncDiffusion [12], we follow their evaluation setup which is based on 6 prompts from MultiDiffusion [1]. We generate 500 panorama images of resolution 512×2048 for each prompt and crop each panorama into 6 random 512×512 patches for quantitative evaluation. To evaluate our method and compare it with baselines, we use several metrics. To measure the image quality and diversity, we use the clean FID [8, 11, 15] implementation with CLIP [16] backbone. CLIPScore [7] measures the alignment between the generated images and the prompts. Furthermore, we use ImageReward [23] to evaluate the quality and alignment of the generated images. ImageReward is a general-purpose text-to-image human preference ranking model, which is trained on a total of 137k pairs of expert comparisons and thus serves as a good proxy to human evaluation. We report the reward scores for all generated images and their corresponding prompts. Finally, we measure the time required to generate an image to demon-

MultiDiffusion with stride=64, views=4
0:07 min/img



"a view of a room with a lot of furniture and a chandelier, warsaw, arcane books, hippie chic"

Ours with stride=64, views=4
0:07 min/img



"a view of a house in a field with trees and a sky background, photogrammetry style"



MultiDiffusion with stride=32, views=7
0:23 min/img



"view of a car parked in a driveway in front of a house, holiday resort"


Ours with stride=64, views=4
0:07 min/img



"a view of a living room with a couch and a table, santa's workshop"




MultiDiffusion with stride=16, views=13
0:44 min/img



"a view of a house with a bench in the yard"

Ours with stride=64, views=4
0:07 min/img



"a view of a small tower on a rocky hill with a sky background, virtual metaverse room"



Figure 3. Comparison of MultiDiffusion [1] with varying stride sizes and our approach. MultiDiffusion with stride=64 (no overlap between views) matches our image generation times but produces strong border artifacts and visible seams due to disjoint diffusion paths. With stride=32 (50% overlap), MultiDiffusion still shows visible seams, and only with stride=16 (75% overlap) does MultiDiffusion produce seamless panoramas, but at the cost of increased computation. In contrast, our method consistently achieves seamless panoramas, reducing inference time by 6x without overlapping denoising views, making it more efficient for high-resolution image generation.




Figure 4. Our method can also replace the inner MultiDiffusion [1] loop in SyncDiffusion [12], leading to a 3x speedup in inference time without noticeable degradation in image quality. The generated panorama images are coherent and sharp without border artifacts, demonstrating the effectiveness of our shifted window denoising approach instead of requiring many overlapping patches.

strate the efficiency of our method and compute the number of windows required to denoise at every timestep.

5. Results

This section presents qualitative and quantitative results of our method compared to MultiDiffusion [1], SyncDiffusion [12], and StitchDiffusion [22].

5.1. Qualitative Results

We start by demonstrating that MultiDiffusion [1] requires large overlap between denoising windows to produce coherent panorama images in Figure 3. However, this large overlap results in an increased number of required denoising predictions and thus slow image generation time. For example, with a stride of 64, which equals the window size, MultiDiffusion [1] produces disjoint diffusion paths and visible seams in the generated images. Even with a stride of 32, MultiDiffusion [1] still shows visible seams while being 3x slower and requiring 7 instead of 4 denoising windows.

Only with a stride of 16, and thus 13 windows, does MultiDiffusion [1] produce seamless panoramas. In contrast, our method produces seamless panoramas without overlapping denoising windows, reducing inference time by 6x.

In Figure 4, we compare our method with SyncDiffusion [12]. More specifically, we replace the inner MultiDiffusion [1] loop in SyncDiffusion [12] with our method to demonstrate the benefit of our approach. Our method achieves a 3x speedup in inference time without noticeable degradation in image quality. The panorama images generated by our method are coherent, sharp and without border artifacts, demonstrating the effectiveness of our shifted window denoising approach.

Finally, we use our method to speed up the inference time of StitchDiffusion [22], see Figure 6. Our method produces seamless and plausible 360-degree panoramas based on the input text prompts. To demonstrate the continuity between the leftmost and rightmost sides of the 360-degree image, we copy the left area and paste it onto the rightmost side.

Model	Stride	Views↓	FID↓	CLIPScore ↑	ImageReward↑	Time[min]↓
MultiDiffusion [1]	16	13	3.21	31.67	0.75	0:44
MultiDiffusion [1]	32	7	3.50	31.65	0.57	0:23
MultiDiffusion [1]	64	4	7.25	30.83	-0.09	0:07
Ours	64	4	3.59	31.67	0.76	0:07
SyncDiffusion [12]	16	13	1.86	31.85	0.62	1:44
SyncDiffusion [12]	32	7	2.08	31.79	0.57	0:57
SyncDiffusion [12]	64	4	5.22	30.97	0.01	0:33
Ours	64	4	2.32	31.93	0.65	0:33

Table 1. Quantitative results on 512×2048 panorama generation. Fast inference with non-overlapping windows comparison is marked in gray. Our method consistently produces high-quality images as measured by FID with good image-text alignment as measured by CLIPScore and ImageReward. In contrast to the baselines, our method does not require overlapping denoising windows, thus significantly reducing the number of required denoising views and inference time. We reach similar performance as the baselines but with a fraction of the time.




Figure 5. **Left:** The number of total required denoising steps and thus image generation time depends on the stride of denoising windows. Given a default window size of 64, a stride of [64, 32, 16] corresponds to [0%, 50%, 75%] overlap between denoising windows, respectively. **Middle:** CLIPScore comparison of the base StableDiffusion model with MultiDiffusion [1] and our method. As can be seen, our method reaches similar performance as MultiDiffusion [1] but significantly faster. **Right:** FID comparison of SyncDiffusion [12] with our method. Our method achieves similar FID scores as SyncDiffusion [12] but with a fraction of the time. Notably, our method does not require overlapping denoising windows (window size = stride = 64) and subsequent averaging, making it more efficient.

5.2. Quantitative Results

To thoroughly evaluate our method, we compare it with MultiDiffusion [1] and SyncDiffusion [12] using various metrics. The results are summarized in Table 1 and Figure 5.

In Figure 5, we show that the number of required denoising steps and thus image generation time depends on the stride of denoising windows. Next, we compare the CLIPScore of the base StableDiffusion model with MultiDiffusion [1] and our method. Our method reaches similar performance as MultiDiffusion [1] but significantly faster. Only with a stride of 16, corresponding to 75% overlap between denoising windows and roughly 6x more compute time, does MultiDiffusion [1] reach similar CLIPScore. Finally, we compare the FID scores of SyncDiffusion [12] with our method. As already shown in the qualitative results, our method achieves similar image quality in terms of FID scores as MultiDiffusion [1] and SyncDiffusion [12] but with a fraction of the time. Interestingly, our method achieves

better image-text alignment as measured by CLIPScore and ImageReward while being significantly faster. We hypothesize this is due to reduced following of the text guidance by averaging separate overlapping denoising predictions within each timestep.

6. Limitations & Future Work

Our method relies on random shifts of non-overlapping denoising windows over time to ensure coherent transitions across the panorama image. While this approach significantly reduces the computational complexity of the denoising process and speeds up image generation time, it does not guarantee the same level of image quality as measured by FID as dense MultiDiffusion [1] with a very small stride of 16. In future work, we plan to investigate the impact of dynamically adjusted different stride sizes during the denoising process on image quality and generation time to combine the benefits of both methods.

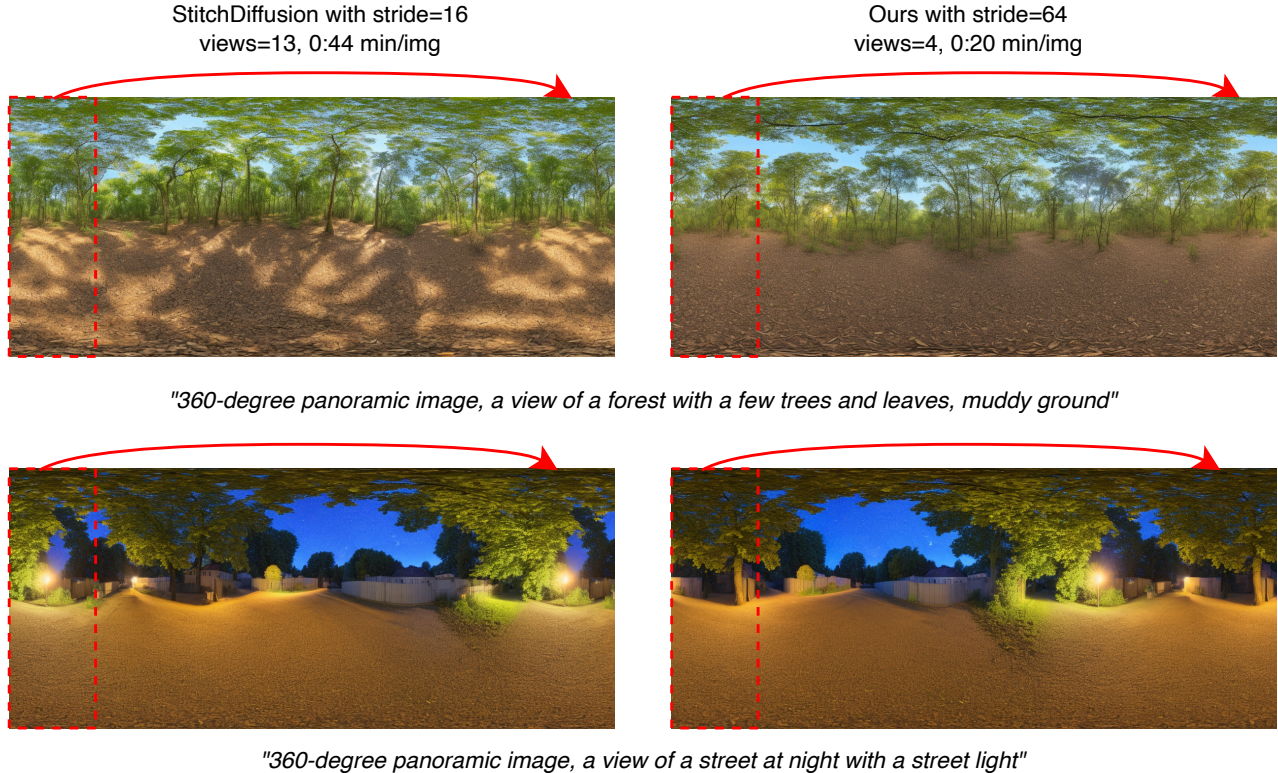


Figure 6. Our method can also replace the inner MultiDiffusion [1] loop in StitchDiffusion [22], a method for spherical image synthesis, leading to a 2x speedup in inference time without noticeable degradation in image quality. To show the continuity between the leftmost and rightmost sides of the generated image, we copy the left area, indicated by the red dashed box, and paste it onto the rightmost side. This demonstrates that both methods produce seamless and plausible 360-degree panoramas based on the input text prompts.

7. Conclusion

This paper presents a novel approach for coherent high-resolution image generation using diffusion models. Instead of relying on overlapping denoising windows, our method shifts the denoising windows over time, ensuring that seams in one timestep are corrected in the next. This results in coherent, high-resolution images with fewer overall steps, enhancing the quality and efficiency of the generation process. We demonstrate the effectiveness of our approach through qualitative and quantitative evaluations, comparing it with MultiDiffusion [1], SyncDiffusion [12], and StitchDiffusion [22]. Our method consistently produces high-quality images with good image-text alignment while significantly reducing the number of required denoising views and inference time. By eliminating the need for overlapping denoising windows, our approach offers a more efficient and practical solution for high-resolution image generation, advancing the state of the art in the field.

8. Societal Impact

Generative image models can be misused to create deepfakes, infringe on copyrights, and produce biased images.

These problems can lead to fake news, privacy invasion, and reinforcing stereotypes. Our method builds on existing foundational models and thus shares these risks. To reduce them, it's important to develop better deepfake detection, protect intellectual property, and follow ethical guidelines. Labeling synthetic content and learning to use generative models responsibly is essential. By focusing on these protective measures, we can use generative models effectively while minimizing their negative impacts.

Acknowledgements

This work was supported by the BMBF project SustainML (Grant 101070408).

References

- [1] Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel. Multidiffusion: Fusing diffusion paths for controlled image generation. *ICML*, 2023. 1, 2, 3, 4, 5, 6, 7, 8
- [2] Florinel-Alin Croitoru, Vlad Hondru, Radu Tudor Ionescu, and Mubarak Shah. Diffusion models in vision: A survey. *TPAMI*, 2023. 2

- [3] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *NeurIPS*, 2021. 1, 2
- [4] Zach Evans, CJ Carr, Josiah Taylor, Scott H Hawley, and Jordi Pons. Fast timing-conditioned latent audio diffusion. *arXiv:2402.04825*, 2024. 2
- [5] Stanislav Frolov, Tobias Hinz, Federico Raue, Jörn Hees, and Andreas Dengel. Adversarial text-to-image synthesis: A review. *Neural Networks*, 2021. 2
- [6] 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
- [7] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi. Clipscore: A reference-free evaluation metric for image captioning. *EMNLP*, 2021. 4
- [8] 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 *NeurIPS*, 2017. 4
- [9] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 2
- [10] 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. *ICLR*, 2022. 4
- [11] Tuomas Kynkäänniemi, Tero Karras, Miika Aittala, Timo Aila, and Jaakko Lehtinen. The role of imagenet classes in fr\`echet inception distance. *ICLR*, 2023. 4
- [12] Yuseung Lee, Kunho Kim, Hyunjin Kim, and Minhyuk Sung. Syncdiffusion: Coherent montage via synchronized joint diffusions. *NeurIPS*, 2023. 1, 2, 4, 6, 7, 8
- [13] Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. Repaint: Inpainting using denoising diffusion probabilistic models. In *CVPR*, 2022. 2
- [14] Brian B Moser, Arundhati S Shanbhag, Federico Raue, Stanislav Frolov, Sebastian Palacio, and Andreas Dengel. Diffusion models, image super-resolution and everything: A survey. *arXiv:2401.00736*, 2024. 2
- [15] Gaurav Parmar, Richard Zhang, and Jun-Yan Zhu. On aliased resizing and surprising subtleties in gan evaluation. In *CVPR*, 2022. 4
- [16] 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 *ICML*, 2021. 4
- [17] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *ICML*, 2021. 1
- [18] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. 1, 2
- [19] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. 2
- [20] Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution. In *NeurIPS*, 2019. 2
- [21] Yang Song, Jascha Sohl-Dickstein, Diederik P. Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021. 2
- [22] Hai Wang, Xiaoyu Xiang, Yuchen Fan, and Jing-Hao Xue. Customizing 360-degree panoramas through text-to-image diffusion models. In *WACV*, 2024. 4, 6, 8
- [23] Jiazheng Xu, Xiao Liu, Yuchen Wu, Yuxuan Tong, Qinkai Li, Ming Ding, Jie Tang, and Yuxiao Dong. Imagereward: Learning and evaluating human preferences for text-to-image generation. *NeurIPS*, 2024. 4
- [24] Ling Yang, Zhilong Zhang, Yang Song, Shenda Hong, Runsheng Xu, Yue Zhao, Wentao Zhang, Bin Cui, and Ming-Hsuan Yang. Diffusion models: A comprehensive survey of methods and applications. *ACM Computing Surveys*, 2023. 2
- [25] Chenshuang Zhang, Chaoning Zhang, Mengchun Zhang, and In So Kweon. Text-to-image diffusion models in generative ai: A survey. *arXiv:2303.07909*, 2023. 2