

SPHERICAL DENSE TEXT-TO-IMAGE SYNTHESIS

Timon Winter^{*}, Stanislav Frolov^{*,†}, Brian Bernhard Moser^{*,†}, Andreas Dengel^{*,†}

[†]RPTU Kaiserslautern-Landau, Germany

^{*}German Research Center for Artificial Intelligence, Germany

ABSTRACT

Recent advancements in text-to-image (T2I) have improved synthesis results, but challenges remain in layout control and generating omnidirectional panoramic images. Dense T2I (DT2I) and spherical T2I (ST2I) models address these issues, but so far no unified approach exists. Trivial approaches, like prompting a DT2I model to generate panoramas can not generate proper spherical distortions and seamless transitions at the borders. Our work shows that spherical dense text-to-image (SDT2I) can be achieved by integrating training-free DT2I approaches into finetuned panorama models. Specifically, we propose MultiStitchDiffusion (MSTD) and MultiPanFusion (MPF) by integrating MultiDiffusion into StitchDiffusion and PanFusion, respectively. Since no benchmark for SDT2I exists, we further construct Dense-Synthetic-View (DSynView), a new synthetic dataset containing spherical layouts to evaluate our models. Our results show that MSTD outperforms MPF across image quality as well as prompt- and layout adherence. MultiPanFusion generates more diverse images but struggles to synthesize flawless foreground objects. We propose bootstrap-coupling and turning off equirectangular perspective-projection attention in the foreground as an improvement of MPF. [Link to code](#)

Index Terms— Spherical Image Generation, Text-to-Image Synthesis, Diffusion Models

1. INTRODUCTION

Text-to-image (T2I) has gained significant traction recently, with advancements like StableDiffusion [1] driving progress [2]. However, user demands have also increased with longer, more complex prompts. Traditional models often fail to handle detailed prompts, misaligning object properties, positional relations, or ignoring details entirely. To address these issues, dense text-to-image (DT2I) models [3, 4, 5] were introduced, tackling these challenges directly. Many approaches enhance user control by allowing explicit layout input, such as masks [6, 7]. Fig. 1 illustrates examples of various DT2I models.

At the same time, users also often seek to customize the image dimensions. Cylindrical panoramas must seamlessly transition between left and right borders, while spherical panoramas require equirectangular projection (ERP) to

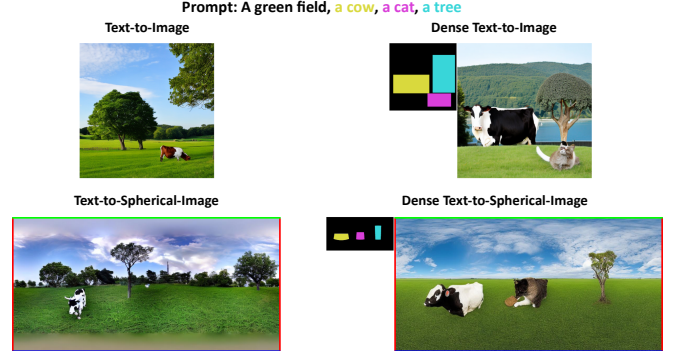


Fig. 1. Task comparison: Traditional text-to-image (top-left) generates images based on a single global prompt. Dense text-to-image (top-right) introduces masks to control the spatial layout. Spherical text-to-image (bottom left) synthesizes 360x180-degree panoramas with seamless transition over the border and distortions at the poles. In this work, we integrate both approaches to enable controllable spherical dense text-to-image synthesis (bottom-right).

produce distortion-free 360x180-degree images when mapped onto a sphere. Spherical text-to-image (ST2I) models [8, 9, 10] address these needs.

Despite these advancements, no existing model combines DT2I and ST2I capabilities. This work addresses this gap by proposing the first spherical dense text-to-image (SDT2I), see Fig. 1. We hypothesize that SDT2I can be achieved by integrating a pre-trained spherical image model with a DT2I approach as a “plug & play” component. Using StitchDiffusion [9] and PanFusion [10] as our spherical image backbones and MultiDiffusion (MD) [4] as the DT2I plugin, we extend these models to support layout guidance through free-hand masks. Additionally, we improve condition adherence and image quality by introducing a bootstrap coupling mechanism. To evaluate our models, we create Dense-Synthetic-View (DSynView), a new synthetic dataset to evaluate our models by generating a diverse set of spherical panoramic images with layout annotations. In summary, our contributions are as follows: (1) We integrate MultiDiffusion (MD) [4] into StitchDiffusion [9] and PanFusion [10], creating MultiStitchDiffusion (MSTD) and MultiPanFusion (MPF), the first SDT2I models.

(2) These models are evaluated on image quality, diversity, and text/layout adherence, showing MSTD matches its baselines. (3) We introduce Dense-Synthetic-View (DSynView), a synthetic dataset with diverse masks and prompts for SDT2I. (4) A naive MPF implementation underperforms MSTD, so we enhance MPF with bootstrap-coupling and improved foreground attention for higher quality.

2. BACKGROUND

2.1. Diffusion Models

The core idea of diffusion models [11] is to split the generative process into T timesteps and learn to generate images through a sequence of operations:

$$I_T, I_{T-1}, \dots, I_0 \quad \text{s.t.} \quad I_{t-1} = \theta(I_t|y), \quad (1)$$

where y is the input embedding, I_t is the output image at timestep t and θ is the diffusion model. Training and sampling is organized in a *forward-process* and a *reverse-process*. The forward process is only applied at training time. Taking one training image as data x , Gaussian noise ϵ is added to the image according to the noise level of each timestep. The model is optimized to predict a noise residual ϵ_θ which can be subtracted from the diffused image to match the denoised ground truth data. In the reverse process, I_T is initialized as a Gaussian noise distribution. The predictive model is now used to predict a noise residual at timestep t using prompt y . This residual is subtracted from I_t , receiving I_{t-1} , which serves as an input at timestep $t - 1$. Thus, the initial noise is gradually removed from the image until after T steps, a clean output image remains.

2.2. Latent Diffusion Models (LDM)

LDMs [1] address the high computational cost of traditional diffusion models by operating in a reduced latent space instead of pixel space. Using a pre-trained variational auto-encoder [12], an encoder \mathcal{E} maps an image x to a latent representation $z = \mathcal{E}(x)$, while a decoder \mathcal{D} reconstructs it: $\tilde{x} = \mathcal{D}(z) \approx x$. LDMs also enhance sample quality by abstracting irrelevant details in the latent space. A cross-attention mechanisms [13] links input tokens to image patches to enable text-to-image synthesis. Stable Diffusion (SD), built on the LDM framework, has become the most widely used open-source model for image generation.

2.3. Dense Text-to-Image (DT2I)

In DT2I, a global prompt is extended with tuples of local prompts and layouts that define object placement, typically as rectangular boxes or free-hand masks. MultiDiffusion (MD) [4] introduced a training-free framework for DT2I by fusing multiple diffusion paths. Mores specifically, an image is generated in a sliding window fashion, ensuring consistency across

overlapping regions by averaging the denoising predictions. For panoramas, overlaps are defined by sliding windows, while for DT2I, they arise from user-defined binary masks. To align synthesized objects with input masks, the authors proposed a bootstrapping phase, where the model focuses on mask areas during early diffusion steps, directing attention away from the background. This improves object placement accuracy and adherence to layouts.

2.4. Spherical Text-to-Image (ST2I)

As the demand for immersive experiences grows, moving from traditional 2D image synthesis to spherical image generation has become increasingly important. StitchDiffusion [9] builds on MD [4] to create seamless cylindrical panoramas. While MD smooths horizontal transitions, it lacks a full cyclic connection. StitchDiffusion introduces a “stitch block” that merges the left and right image edges during denoising, ensuring a continuous panorama. For spherical panoramas, ERP distortions are added by fine-tuning SD with LoRA [14] layers on 120 panoramic images. PanFusion [10] addresses ST2I generation by splitting the process into two branches: a panorama branch for global coherence and a perspective branch leveraging SD’s generative strength for individual images. Communication between branches is enabled by equirectangular-perspective projection attention (EPPA), enhanced with spherical positional encoding and an attention mask to align feature maps from both branches. PanFusion adapts SD to 512x1024 resolution using LoRA [14] layers and trains on the Matterport3D dataset [15]. Circular padding ensures loop consistency, while perspective views are generated via icosahedron tessellation.

3. METHODS

Our goal is to develop the first spherical dense text-to-image (SDT2I) models by integrating training-free DT2I approaches into finetuned panorama models. To that end, we propose MultiStitchDiffusion (MSTD) and MultiPanFusion (MPF) in the next sections.

3.1. MultiStitchDiffusion (MSTD)

As our first method, we integrate StitchDiffusion [9] with the region-based variant of MD [4]. The input includes a global text prompt and N masks with corresponding local prompts. Following MD, we preprocess the masks and extend the edges cyclically to handle the stitching steps. StitchDiffusion’s LoRA is activated by the trigger word “360-degree panoramic image”, with an option to toggle it separately for background and foreground objects. Inference is split into multiple processes, each handling one mask and prompt. After every denoising step, latents are merged as described in Sec. 2.3.

3.2. MultiPanFusion (MPF)

As our second method, we use PanFusion [10] as our base model and integrate MD into the inference process. The modification follows a similar approach to our StitchDiffusion integration, with computations on latents repeated multiple times for each denoising path. To maintain alignment, masks are rotated alongside latents during each denoising step. After denoising, latents are merged as in MD. Since PanFusion contains two branches (see Sec. 2.3), there are multiple possible implementations. Initially, we integrate MD into the panorama branch only. To improve consistency during denoising, we extend MD to the perspective branch by projecting masks from ERP to perspective format and applying MD similarly to the panorama branch.

3.2.1. Eliminating Bootstrapping-Flaws

During evaluation, we observed artifacts like blurring and pixelation around the foreground objects. These issues likely stem from the naive application of MD’s bootstrapping phase, which assumes that mask backgrounds can be replaced with constant colors before denoising, as these areas are ignored during latent merging. In PanFusion’s dual-branch approach, this assumption fails. Using different background colors for the panorama and perspective branches or disabling MD in one branch creates conflicts in the EPPA module, leading to unpredictable results. Additionally, the Gaussian-smoothed EPP attention mask can spread background colors around objects over multiple denoising steps. To address this, we propose two improvements: (1) Disabling the EPPA module for foreground objects during the bootstrapping phase to eliminate conflicts. (2) Coupling bootstrapping backgrounds across branches by assigning the same random color to both panorama and perspective branches and ensuring consistent background colors for objects at each timestep.

4. EXPERIMENTS

4.1. Dataset & Metrics

Since no benchmark exists for spherical dense text-to-image synthesis, we create Dense-Synthetic-View (DSynView), a new synthetic dataset containing spherical layouts to evaluate our models. In DSynView, we combined three text-prompts used for conditioning the background with two different sets of up to three fitting foreground prompts associated with self-created masks.

We evaluate six scenes combined with 168 seeds, resulting in 1,008 panoramas and 3,024 perspective images for testing. Tab. 1 summarizes the conditions. To examine the impact of three mask sizes and ERP re-projection, we generate ERP-reprojected masks by first determining the mask’s center in planar coordinates, converting it to spherical coordinates, and projecting the corresponding perspective view with a 120° FoV.

foreground prompts	small	medium	large
table, bed, cow, sheep, bus, car			
television, potted plant, cat, pond, sign, bicycle			
wardrobe, door, tree, windmill, building, traffic light			

Table 1. Normal and ERP-re-projected masks with associated foreground prompts, assigned to an indoor room, a green field, and a busy street as background prompts

The target perspective mask is then used to construct a bounding box, which is re-projected onto an empty ERP image to create the final mask. As shown in Tab. 1, we also evaluate how the model performs with different aspect ratios, such as long, tall, or square masks, paired with suitable prompts like a bed, door, or sign. To create reference 2D images for evaluating generated outputs using FID [16], ImageReward [17] and CLIP-Score [18], we use MD to generate perspective views of foreground objects based on prompts and background scenes. The target perspective masks are derived from ERP masks, ensuring the object is centered and fully visible. With 18 prompts and 168 seeds, the dataset includes 18,144 reference images.

4.2. Experimental Setup

We configure MultiStitchDiffusion (MSTD) by following StitchDiffusion [9]. We systematically vary key hyperparameters as shown in Tab. 2, using a default setting and modifying at most one parameter at a time. We also run a baseline by disabling StitchDiffusion entirely (i.e., pure MultiDiffusion). For MultiPanFusion (MPF), we use the setup of PanFusion [10] and evaluate the same hyperparameters. We consider MD applied in the panorama branch (md_pano), perspective branch (md_pers), or both (md_both). Bootstrap-coupling can ensure the same background color across branches (branches-coupling) and the same color for each object (objects-coupling), while we optionally disable EPPA for foreground objects.

4.3. Quantitative Results

In this section, we report a selection of the most informative quantitative results. Full benchmark results are included in the supplementary material. Our main results are in Tab. 3. When comparing MSTD to the baselines, we observe that most scores are roughly the same. MPF performs worse than MSTD in all

parameter	value range	default (MSTD / MPF)
bootstrapping	{1, 5, 10, ..., 50}	20
stride (MSTD)	{4, 8, 16, 32}	8
mask size	{S, M, L}	M
mask type	{regular, ERP-reproj.}	ERP-reproj.
mask indices	$\mathcal{P}\{0, 1, 2\}$	{0, 1, 2}
LoRA enabled	{yes, bg-only, no}	bg-only / yes
bootstrap-cpl.	{branches, objects, none}	none / branches
noise-cpl.	{yes, no}	yes
global prompt	{yes, no}	no
fg. EPPA (MPF)	{yes, no}	yes

Table 2. Key hyperparameters used to benchmark MultiStitchDiffusion (MSTD) and MultiPanFusion (MPF).




Fig. 2. The influence of bootstrapping on our metrics for every approach, showing a functional relationship which is non-monotonous at FID.

metrics. Furthermore, an IoU of 0.05 at *md_pers* shows the insufficiency of applying MD at the perspective branch only for synthesizing foreground objects. Even when combined with MD applied at the panorama branch, the scores do not change significantly. Notably, our model reaches roughly the same image quality while generating full spherical images.

	task	IoU \uparrow	IR \uparrow	FID \downarrow
MD (original)	DT2I	0.57	-0.47	82.61
MD (with pano-LoRA)	DT2I	<u>0.66</u>	<u>-0.42</u>	60.99
MSTD (ours)	SDT2I	0.67	-0.37	<u>61.12</u>
MPF (md_pano) (ours)	SDT2I	0.45	-1.19	84.60
MPF (md_pers) (ours)	SDT2I	0.05	-1.79	107.00
MPF (md_both) (ours)	SDT2I	0.44	-1.22	84.82

Table 3. Performance of our MSTD and MPF models compared to the MD baseline. “with pano-LoRA” means MD is adapted to the ERP-distorted style.

Fig. 21 shows that as the bootstrapping parameter increases from 0 to its maximum in increments of 5, IoU steadily improves until plateauing around 30, while FID reaches its optimal value at a bootstrapping of 15. Tab. 4 shows that bootstrap coupling between both branches in MPF leads to a small quantitative improvement in all scores. We anticipated greater improvement as this technique largely resolves visual artifacts.

bootstrap-coupling	IoU \uparrow	IR \uparrow	FID \downarrow
no	0.43	-1.26	85.73
yes	0.44	-1.22	84.82

Table 4. Bootstrap-coupling at PanFusion with MultiDiffusion applied at both branches is better across all metrics.

Method	IoU \uparrow	IR \uparrow	FID \downarrow
MSTD (no global prompt)	0.67	-0.37	61.12
MSTD (global prompt)	0.59	-0.35	57.03
MPF (md_pano, no global prompt)	0.45	-1.19	84.60
MPF (md_pano, global prompt)	0.43	-1.25	74.84
MPF (md_pers, no global prompt)	0.05	-1.79	107.00
MPF (md_pers, global prompt)	0.08	-1.60	81.94
MPF (md_both, no global prompt)	0.44	-1.22	84.82
MPF (md_both, global prompt)	0.41	-1.32	77.20

Table 5. Influence of the inclusion of foreground objects in the background prompt. Global prompts lead to increased image quality but slightly lower mask fidelity.

Tab. 5 examines the effect of including local prompts in the background prompt. While this negatively impacts IoU in all methods except *md_pers*, it simultaneously lowers FID, making optimization challenging. Further qualitative analysis is needed for better insights. In Tab. 6, we analyze the influence of EPPA. Against our expectations, the ignorance of EPPA at foreground objects has a positive and more considerable impact than bootstrap-coupling. IoU, Image-Reward, and FID have better scores when using this variant.

4.4. Qualitative Results




Fig. 3. Prompt and mask adherence of MultiStitchDiffusion; showing a high IoU but some objects, e.g. the building or the wardrobe, not blending well into the background.

Method	IoU \uparrow	IR \uparrow	FID \downarrow
MPF md_pano (FG + BG EPPA)	0.45	-1.19	84.60
MPF md_pano (BG EPPA)	0.53	-1.09	80.61
MPF md_pers (FG + BG EPPA)	0.05	-1.79	107.00
MPF md_pers (BG EPPA)	0.05	-1.79	106.68
MPF md_both (FG + BG EPPA)	0.44	-1.22	84.82
MPF md_both (BG EPPA)	0.52	-1.10	82.21

Table 6. Influence of not applying EPPA at foreground layers, showing an overall improvement at mask adherence and image quality with MultiPanFusion.

MSTD generally handles prompt adherence well, synthesizing most requested objects but occasionally misinterpreting specific items (e.g., partial traffic lights). Mask alignment is decent (IoU around 2/3), yet bounding boxes often lack precision. Backgrounds show limited variety, frequently resulting in incomplete or duplicated objects. Bootstrapping strongly influences performance: higher values improve alignment but can distort shapes and hamper blending. Medium mask sizes strike a balance between neglect (small) and oversizing (large). LoRA for foreground objects tends to add distortions, while projecting masks in an ERP-aligned manner avoids unnatural object shapes.




Fig. 4. Prompt and mask adherence of MultiPanFusion being satisfactory in some examples. The middle-left image, however, shows a failure case.

MPF sometimes misses or misinterprets objects, but it provides richer, more diverse backgrounds than MSTD. Although objects typically align with their masks, some appear over-fitted, producing artifacts around the mask boundary. Applying MD to both panorama and perspective branches reduces missing objects but can still introduce conflicts. Disabling foreground EPPA sometimes clarifies objects, while bootstrapping affects object size and shape, with an ideal range around 10–20. Coupling the bootstrapping background across branches helps

slightly but is not a major factor. Mask size and global prompt usage have similar impacts as in MSTD, and other parameters show minimal effects.



Fig. 5. Various failure-cases of MultiPanFusion, visualizing noise around the objects, unfitting ceilings, wrong objects, and foreground elements merging with the background.

Failure Case Analysis: The weaker quantitative results of MPF might stem from several recurring issues. First, foreground objects (e.g., a rectangular mask at high bootstrapping) can blend awkwardly with the background, occasionally creating indoor elements in outdoor scenes. This is likely tied to the training data of PanFusion, which is primarily indoors, causing outdoor prompts to appear as if viewed from inside a room. Second, MPF sometimes fails to integrate foreground objects into scene layouts, merging cars with buildings or introducing other distortions, see Fig. 5. Lastly, artifacts like blurriness and pixelation persist around masked regions, even with our improvements. These artifacts raise FID and CMMD and reduce scores for prompt-adherence (IoU, CLIP-score, and Image-Reward).

5. CONCLUSION & FUTURE WORK

We introduced MultiStitchDiffusion (MSTD) and MultiPanFusion (MPF), the first approaches for spherical dense text-to-image (SDT2I) synthesis, along with a benchmark for evaluation. MSTD preserves prompt- and mask-adherence from MultiDiffusion (MD) while producing coherent panoramas with StitchDiffusion. MPF builds on PanFusion’s dual-branch architecture, offering increased scene variety but showing more artifacts and misinterpretations. Our experiments highlight the critical role of hyperparameters (e.g., bootstrapping, mask size, and LoRA usage) and suggest improvements like bootstrap-coupling and reduced use of EPPA. However, both approaches still face challenges: MSTD’s repetitive backgrounds and MPF’s indoor biases, object merges, and blurry artifacts can limit real-world applicability. Furthermore, our evaluation relies on MD-generated references and 18 prompts, which may limit generalizability. Future work should explore richer data sets with real panorama references. Despite these limitations, our methods represent a significant step toward flexible, text-guided 360° image synthesis.

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

6. REFERENCES

- [1] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer, “High-resolution image synthesis with latent diffusion models,” in *CVPR*, 2022.
- [2] Chenshuang Zhang, Chaoning Zhang, Mengchun Zhang, In So Kweon, and Junmo Kim, “Text-to-image diffusion models in generative ai: A survey,” in *arXiv:2303.07909*, 2024.
- [3] Stanislav Frolov, Prateek Bansal, Jörn Hees, and Andreas Dengel, “Dt2i: Dense text-to-image generation from region descriptions,” in *ICANN*, 2022.
- [4] Omer Bar-Tal, Lior Yariv, Yaron Lipman, and Tali Dekel, “Multidiffusion: Fusing diffusion paths for controlled image generation,” in *ICML*, 2023.
- [5] Yunji Kim, Jiyoung Lee, Jin-Hwa Kim, Jung-Woo Ha, and Jun-Yan Zhu, “Dense text-to-image generation with attention modulation,” in *ICCV*, 2023.
- [6] Omri Avrahami, Thomas Hayes, Oran Gafni, Sonal Gupta, Yaniv Taigman, Devi Parikh, Dani Lischinski, Ohad Fried, and Xi Yin, “Spatext: Spatio-textual representation for controllable image generation,” in *CVPR*, 2023.
- [7] Guillaume Couairon, Marlène Careil, Matthieu Cord, Stéphane Lathuilière, and Jakob Verbeek, “Zero-shot spatial layout conditioning for text-to-image diffusion models,” in *ICCV*, 2023.
- [8] Shitao Tang, Fuyang Zhang, Jiacheng Chen, Peng Wang, and Yasutaka Furukawa, “Mvdiffusion: Enabling holistic multi-view image generation with correspondence-aware diffusion,” in *NeurIPS*, 2023.
- [9] Hai Wang, Xiaoyu Xiang, Yuchen Fan, and Jing-Hao Xue, “Customizing 360-degree panoramas through text-to-image diffusion models,” in *WACV*, 2024.
- [10] Cheng Zhang, Qianyi Wu, Camilo Cruz Gambardella, Xiaoshui Huang, Dinh Phung, Wanli Ouyang, and Jianfei Cai, “Taming stable diffusion for text to 360° panorama image generation,” in *CVPR*, 2024.
- [11] Jascha Sohl-Dickstein, Eric A. Weiss, Niru Maheswaranathan, and Surya Ganguli, “Deep unsupervised learning using nonequilibrium thermodynamics,” in *ICML*, 2015.
- [12] Diederik P Kingma and Max Welling, “Auto-encoding variational bayes,” in *ICLR*, 2014.
- [13] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention is all you need,” in *NeurIPS*, 2017.
- [14] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen, “Lora: Low-rank adaptation of large language models,” in *arXiv:2106.09685*, 2021.
- [15] Angel Chang, Angela Dai, Thomas Funkhouser, Maciej Halber, Matthias Nießner, Manolis Savva, Shuran Song, Andy Zeng, and Yinda Zhang, “Matterport3d: Learning from rgb-d data in indoor environments,” in *3DV*, 2017.
- [16] 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.
- [17] 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,” in *NeurIPS*, 2024.
- [18] Jack Hessel, Ari Holtzman, Maxwell Forbes, Ronan Le Bras, and Yejin Choi, “Clipscore: A reference-free evaluation metric for image captioning,” in *EMNLP*, 2021.
- [19] Sadeep Jayasumana, Srikumar Ramalingam, Andreas Veit, Daniel Glasner, Ayan Chakrabarti, and Sanjiv Kumar, “Rethinking fid: Towards a better evaluation metric for image generation,” 2024.

Appendix

Here we report a selection of materials that provides further insight into our experiments. We start by showing an extract of reference images from DSynView on which the quality assessment was conducted. Next, we present a greater variety of qualitative results that show the influence of different parameters. To improve visibility, we subsequently show the extracted perspective images at larger size. Furthermore, we present qualitative results from experiments testing the limitations of MultiPanFusion. Finally, we show the full tables of quantitative results with MSTD and MPF, including all assessed parameters as well as two additional metrics, CLIP-score and CLIP Maximum Mean Discrepancy (CMMD) [19]. We visualize the influence of bootstrapping and mask size in the same way we did in Fig. 21.




Fig. 6. Example images from our DSynView dataset

1.1. More Qualitative Results with MultiPanFusion




Fig. 7. Image-quality and -diversity of MSTD; showing a decent quality but also similar-looking backgrounds for different seeds.



Fig. 8. Influence of stitching-operation. Left: stitching. Right: No stitching




Fig. 9. Influence of bootstrapping, visualizing a smoother blending of objects and background at lower values and a higher mask-adherence at higher values.



Fig. 10. Influence of mask size. Top: small. Middle: middle. Bottom: large




Fig. 11. Influence of spherical parameters. The top-left setting usually yields the best results, while the bottom-left is insufficient for spherical images, as the missing distortions in the ERP image show.




Fig. 12. Influence of other parameters. The comparison between the top-left and bottom-left images shows that adding more objects doesn't influence the background's or other object's look.




Fig. 13. Image-quality and -diversity of MPF. Contrarily to MSTD, images are filled with rich and varied content here but show some quality-reducing artifacts, like the blurry area behind the car.




Fig. 14. Top: MD applied at the panorama branch only. Middle: MD applied at both branches. Bottom: MD applied at the perspective branch only, which fails to synthesize objects.




Fig. 15. Left: EPPA employed at all paths. Right: EPPA employed only at the background during bootstrapping. The pond can only be recognized in the right image. Such cases cause the IoU to increase in this setting.




Fig. 16. Influence of bootstrapping on MPF, showing similar effects as with MSTD.




Fig. 17. Influence of bootstrap-coupling. The largest effect can be seen at the middle object, where the bicycle only becomes visible with branches coupling




Fig. 18. Perspective Images of MultiStitchDiffusion's qualitative results




Fig. 19. Perspective Images of MultiPanFusion’s qualitative results




Fig. 20. Corner cases of MultiPanFusion with free-hand masks. The model has problems synthesizing objects near the top and bottom poles. Elements at the horizontal image borders are unproblematic. Overlapping masks result in object-neglect and -distortion.

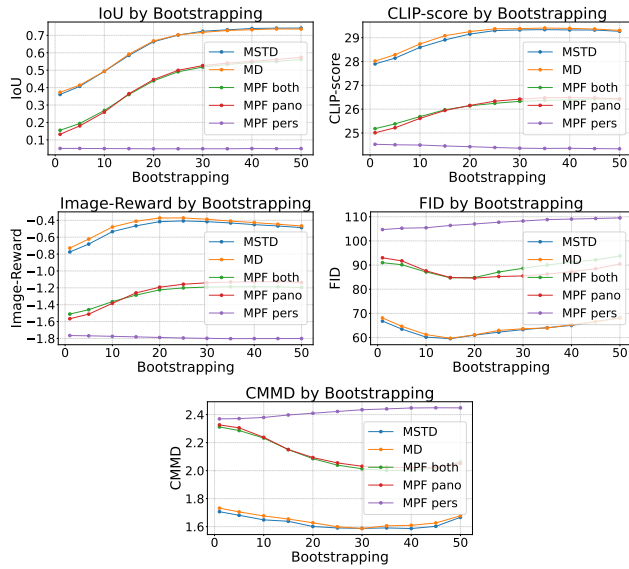


Fig. 21. The influence of bootstrapping on our metrics for every approach, showing a functional relationship which is non-monotonous at FID.

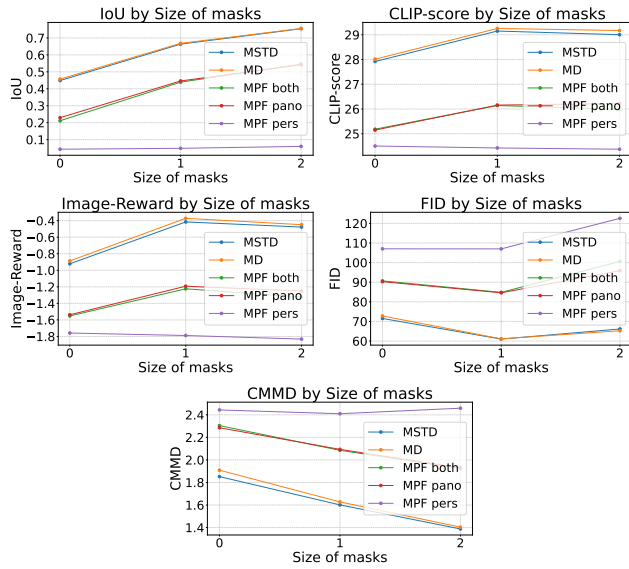


Fig. 22. The influence of mask size on our metrics for every approach. The plots show a similar effect as with bootstrapping.

		without StitchDiffusion process					with StitchDiffusion process				
		IoU↑	CS↑	IR↑	FID↓	CMMD↓	IoU↑	CS↑	IR↑	FID↓	CMMD↓
bootstrap	1	0.36	27.90	-0.77	66.81	1.71	0.37	28.02	-0.73	68.11	1.73
	5	0.41	28.14	-0.68	63.50	1.68	0.41	28.28	-0.62	64.64	1.71
	10	0.49	28.59	-0.53	60.17	1.65	0.49	28.74	-0.48	61.23	1.68
	15	0.59	28.91	-0.47	59.53	1.64	0.59	29.08	-0.41	59.69	1.65
	20	0.66	29.15	-0.42	60.99	1.60	0.67	29.25	-0.37	61.12	1.63
	25	0.70	29.30	-0.41	62.23	1.59	0.70	29.37	-0.37	62.94	1.60
	30	0.72	29.33	-0.42	63.29	1.59	0.72	29.38	-0.39	63.63	1.59
	35	0.73	29.34	-0.43	64.08	1.59	0.73	29.40	-0.41	63.97	1.61
	40	0.74	29.33	-0.45	65.03	1.59	0.73	29.39	-0.43	65.28	1.61
	45	0.74	29.32	-0.47	66.53	1.60	0.74	29.36	-0.45	66.58	1.63
	50	0.74	29.27	-0.49	68.24	1.67	0.74	29.31	-0.47	68.05	1.68
stride	4	0.66	29.15	-0.44	62.47	1.63	0.68	29.27	-0.36	60.96	1.62
	8	0.67	29.18	-0.42	61.70	1.62	0.67	29.28	-0.37	61.13	1.62
	16	0.66	29.15	-0.42	60.99	1.60	0.67	29.25	-0.37	61.12	1.63
	32	0.67	29.15	-0.41	60.90	1.61	0.67	29.24	-0.37	61.32	1.62
size	0	0.45	27.92	-0.92	71.51	1.85	0.46	28.01	-0.89	72.79	1.91
	1	0.66	29.15	-0.42	60.99	1.60	0.67	29.25	-0.37	61.12	1.63
	2	0.75	29.00	-0.48	66.14	1.39	0.76	29.17	-0.45	65.25	1.40
pano	None	0.57	28.84	-0.56	82.50	1.50	0.57	28.91	-0.48	82.61	1.51
	BG	0.66	29.15	-0.42	60.99	1.60	0.67	29.25	-0.37	61.12	1.63
	All	0.56	28.66	-0.47	67.78	1.70	0.55	28.74	-0.43	68.42	1.72
proj	False	0.62	29.04	-0.49	61.35	1.61	0.63	29.16	-0.43	61.40	1.62
	True	0.66	29.15	-0.42	60.99	1.60	0.67	29.25	-0.37	61.12	1.63
mask idx	0	0.73	29.94	-0.58	75.89	1.64	0.73	29.82	-0.60	77.30	1.71
	1	0.57	28.64	-0.27	90.56	1.79	0.59	28.82	-0.16	91.10	1.77
	2	0.69	29.06	-0.21	81.71	1.50	0.69	29.17	-0.18	80.86	1.48
	01	0.65	29.17	-0.49	68.23	1.71	0.66	29.23	-0.42	68.67	1.75
	02	0.72	29.53	-0.42	63.40	1.52	0.71	29.55	-0.40	64.34	1.54
	12	0.64	28.83	-0.27	69.63	1.63	0.64	29.01	-0.20	69.20	1.60
	012	0.66	29.15	-0.42	60.99	1.60	0.67	29.25	-0.37	61.12	1.63
bootstrap coupling	None	0.66	29.15	-0.42	60.99	1.60	0.67	29.25	-0.37	61.12	1.63
	Objects	0.67	29.20	-0.41	61.04	1.62	0.67	29.30	-0.36	61.06	1.63
noise coupling	False	0.67	29.21	-0.41	61.42	1.61	0.68	29.32	-0.36	61.84	1.63
	True	0.66	29.15	-0.42	60.99	1.60	0.67	29.25	-0.37	61.12	1.63
global prompt	False	0.66	29.15	-0.42	60.99	1.60	0.67	29.25	-0.37	61.12	1.63
	True	0.58	29.23	-0.41	56.38	1.57	0.59	29.36	-0.35	57.03	1.59

Table 1. Complete results with MultiDiffusion and MultiStitchDiffusion

		MPF (md_both)					MPF (md_pano)				
		IoU↑	CS↑	IR↑	FID↓	CMMD↓	IoU↑	CS↑	IR↑	FID↓	CMMD↓
bootstrap	1	0.16	25.18	-1.51	90.99	2.31	0.13	25.01	-1.57	93.00	2.33
	5	0.19	25.38	-1.46	90.13	2.29	0.18	25.22	-1.51	91.73	2.31
	10	0.27	25.69	-1.36	87.13	2.23	0.26	25.61	-1.38	87.59	2.24
	15	0.36	25.97	-1.29	84.72	2.15	0.36	25.95	-1.26	84.85	2.15
	20	0.44	26.14	-1.22	84.82	2.09	0.45	26.16	-1.19	84.60	2.09
	25	0.49	26.25	-1.20	87.12	2.04	0.50	26.34	-1.16	85.26	2.06
	30	0.52	26.33	-1.19	88.66	2.01	0.53	26.42	-1.14	85.50	2.03
	35	0.53	26.36	-1.19	90.03	2.00	0.54	26.48	-1.13	86.26	2.03
	40	0.54	26.39	-1.19	91.23	2.00	0.55	26.49	-1.13	87.28	2.02
	45	0.55	26.43	-1.19	92.15	2.02	0.56	26.48	-1.13	88.47	2.03
	50	0.56	26.43	-1.20	93.76	2.06	0.57	26.42	-1.14	90.43	2.05
size	0	0.21	25.18	-1.55	90.67	2.30	0.23	25.15	-1.54	90.28	2.29
	1	0.44	26.14	-1.22	84.82	2.09	0.45	26.16	-1.19	84.60	2.09
	2	0.55	25.99	-1.32	100.59	1.92	0.54	26.22	-1.25	95.89	1.93
pano	BG	0.43	26.45	-1.23	81.90	2.07	0.49	26.89	-1.07	76.57	2.03
	All	0.44	26.14	-1.22	84.82	2.09	0.45	26.16	-1.19	84.60	2.09
proj	False	0.39	25.93	-1.29	85.31	2.11	0.39	25.97	-1.26	83.77	2.10
	True	0.44	26.14	-1.22	84.82	2.09	0.45	26.16	-1.19	84.60	2.09
mask_idx	0	0.53	27.11	-1.39	129.58	2.20	0.53	27.02	-1.41	132.93	2.23
	1	0.35	25.55	-1.18	113.10	2.22	0.34	25.36	-1.18	111.46	2.21
	2	0.48	25.98	-1.05	96.15	1.98	0.47	25.91	-1.06	97.75	1.97
	01	0.43	26.28	-1.30	102.09	2.20	0.43	26.18	-1.29	101.13	2.19
	02	0.50	26.55	-1.22	90.48	2.06	0.50	26.53	-1.20	92.74	2.03
	12	0.40	25.68	-1.13	86.82	2.08	0.41	25.65	-1.11	84.95	2.05
	012	0.44	26.14	-1.22	84.82	2.09	0.45	26.16	-1.19	84.60	2.09
bootstrap coupling	None	0.43	26.04	-1.26	85.73	2.11	0.45	26.16	-1.19	84.63	2.09
	Branches	0.44	26.14	-1.22	84.82	2.09	0.45	26.16	-1.19	84.60	2.09
	Objects	0.44	26.01	-1.25	86.53	2.10	0.44	26.02	-1.22	84.79	2.08
	All	0.44	26.09	-1.23	84.95	2.08	0.44	26.03	-1.22	84.85	2.08
noise coupling	False	0.46	26.22	-1.19	84.49	2.10	0.46	26.24	-1.15	81.78	2.10
	True	0.44	26.14	-1.22	84.82	2.09	0.45	26.16	-1.19	84.60	2.09
global prompt	False	0.44	26.14	-1.22	84.82	2.09	0.45	26.16	-1.19	84.60	2.09
	True	0.41	26.43	-1.32	77.20	2.08	0.43	26.63	-1.25	74.84	2.05
fg_eppa	False	0.52	26.63	-1.10	82.21	2.06	0.53	26.70	-1.09	80.61	1.97
	True	0.44	26.14	-1.22	84.82	2.09	0.45	26.16	-1.19	84.60	2.09

Table 2. Results with MultiPanFusion (md_both and md_pano)

		MPF (md_both)					MPF (md_pers)				
		IoU↑	CS↑	IR↑	FID↓	CMMD↓	IoU↑	CS↑	IR↑	FID↓	CMMD↓
bootstrap	1	0.16	25.18	-1.51	90.99	2.31	0.05	24.53	-1.77	104.68	2.37
	5	0.19	25.38	-1.46	90.13	2.29	0.05	24.51	-1.77	105.25	2.37
	10	0.27	25.69	-1.36	87.13	2.23	0.05	24.49	-1.77	105.44	2.38
	15	0.36	25.97	-1.29	84.72	2.15	0.05	24.46	-1.78	106.40	2.40
	20	0.44	26.14	-1.22	84.82	2.09	0.05	24.42	-1.79	107.00	2.41
	25	0.49	26.25	-1.20	87.12	2.04	0.05	24.39	-1.79	107.73	2.42
	30	0.52	26.33	-1.19	88.66	2.01	0.05	24.36	-1.80	108.24	2.43
	35	0.53	26.36	-1.19	90.03	2.00	0.05	24.35	-1.80	108.81	2.44
	40	0.54	26.39	-1.19	91.23	2.00	0.05	24.36	-1.80	109.04	2.45
	45	0.55	26.43	-1.19	92.15	2.02	0.05	24.35	-1.80	109.27	2.45
	50	0.56	26.43	-1.20	93.76	2.06	0.05	24.34	-1.80	109.51	2.45
size	0	0.21	25.18	-1.55	90.67	2.30	0.04	24.50	-1.76	107.02	2.44
	1	0.44	26.14	-1.22	84.82	2.09	0.05	24.42	-1.79	107.00	2.41
	2	0.55	25.99	-1.32	100.59	1.92	0.06	24.37	-1.83	122.58	2.46
pano	BG	0.43	26.45	-1.23	81.90	2.07	0.05	24.36	-1.81	108.05	2.42
	All	0.44	26.14	-1.22	84.82	2.09	0.05	24.42	-1.79	107.00	2.41
proj	False	0.39	25.93	-1.29	85.31	2.11	0.05	24.42	-1.78	105.76	2.40
	True	0.44	26.14	-1.22	84.82	2.09	0.05	24.42	-1.79	107.00	2.41
mask_idx	0	0.53	27.11	-1.39	129.58	2.20	0.05	24.50	-1.96	162.83	2.68
	1	0.35	25.55	-1.18	113.10	2.22	0.02	24.29	-1.85	134.22	2.45
	2	0.48	25.98	-1.05	96.15	1.98	0.08	24.50	-1.55	135.32	2.38
	01	0.43	26.28	-1.30	102.09	2.20	0.03	24.39	-1.91	123.44	2.49
	02	0.50	26.55	-1.22	90.48	2.06	0.06	24.50	-1.75	120.81	2.43
	12	0.40	25.68	-1.13	86.82	2.08	0.05	24.41	-1.71	109.92	2.36
	012	0.44	26.14	-1.22	84.82	2.09	0.05	24.42	-1.79	107.00	2.41
bootstrap coupling	None	0.43	26.04	-1.26	85.73	2.11	0.05	24.41	-1.79	107.09	2.41
	Branches	0.44	26.14	-1.22	84.82	2.09	0.05	24.42	-1.79	107.00	2.41
	Objects	0.44	26.01	-1.25	86.53	2.10	0.05	24.42	-1.79	106.97	2.41
	All	0.44	26.09	-1.23	84.95	2.08	0.05	24.42	-1.79	107.02	2.41
noise coupling	False	0.46	26.22	-1.19	84.49	2.10	0.05	24.41	-1.78	106.23	2.42
	True	0.44	26.14	-1.22	84.82	2.09	0.05	24.42	-1.79	107.00	2.41
global prompt	False	0.44	26.14	-1.22	84.82	2.09	0.05	24.42	-1.79	107.00	2.41
	True	0.41	26.43	-1.32	77.20	2.08	0.08	25.38	-1.60	81.94	2.29
fg_eppa	False	0.52	26.63	-1.10	82.21	2.06	0.05	24.44	-1.79	106.68	2.42
	True	0.44	26.14	-1.22	84.82	2.09	0.05	24.42	-1.79	107.00	2.41

Table 3. Results with MultiPanFusion (md_both and md_pers)