

國立臺灣大學電機資訊學院資訊工程學系

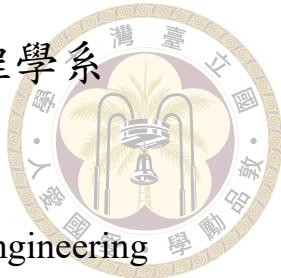
碩士論文

Department of Computer Science and Information Engineering

College of Electrical Engineering and Computer Science

National Taiwan University

Master Thesis



以雙流對抗式生成網路之任意形狀三維立體場景修復

Free-form 3D Scene Inpainting with Dual-stream GAN

鄭如芬

Ru-Fen Jheng

指導教授: 徐宏民 博士

Advisor: Winston Hsu Ph.D.

中華民國 111 年 8 月

August, 2022

國立臺灣大學碩士學位論文
口試委員會審定書

以雙流對抗式生成網路之任意形狀三維立體場景修復

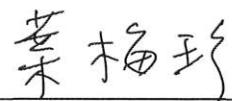
Free-form 3D Scene Inpainting with Dual-stream GAN

本論文係鄭如芬君(學號 R09922014)在國立臺灣大學資訊工程學系完成之碩士學位論文，於民國 111 年 8 月 12 日承下列考試委員審查通過及口試及格，特此證明

口試委員：



陳文進 (指導教授)



系主任

洪士顥



Acknowledgements

過去兩年在研究的道路上受到許多人的幫助，在此感謝幫助過關照過我的每個人，讓我可以順利的完成學業。

首先感謝徐宏民教授不管是對於研究領域方面的指導亦或是生活與研究的平衡，都給予專業與細心的建議提醒，讓我漸漸熟悉做研究的方法跟思路。感謝黃郁凱及劉岳承學長在我初進入研究所的時候，傳授了許多研究的基礎與實驗方法。感謝每週的小組討論中，吳宗翰、劉佳昇和陳瑾瑭都會給予不一樣觀點的建議，幫助我改善。感謝葉佳峯學長跟吳宗翰在論文寫作上提供了非常多有用的經驗與幫助，帶領著我整理這兩年來的研究並用清晰易懂的方式呈現出最好的結果。感謝整個實驗室的同學、學長姊、學弟妹們總是互相幫助，一起吃飯聊天，陪伴我度過研究所的生活。

最後感謝我的家人一直以來都是全力的支持，讓我可以對所有想做的事盡全力，以及許多的朋友常常陪伴鼓勵我；感謝侯家蓉從大學時期到研究所都與我一起朝目標努力。



摘要

現今，由於擴增實境和虛擬實境技術的發展，使用者在三維立體場景中進行編輯的需求迅速增加。然而，現有的三維立體場景補全任務（以及資料集）並無法滿足使用者編輯的需求，因為其場景中的缺失區域是由傳感器偵測限制或物品遮擋產生的。因此，我們提出了任意形狀三維立體場景修復的新任務。與之前的三維立體場景補全任務的資料集中場景保留了缺失區域周圍大部分的主要結構和細節的形狀提示不同，我們所提出的三維立體場景修復資料集（FF-Matterport）包含了由我們提出的任意形狀三維立體遮罩生成演算法所產生大面積而多樣的缺失區域；此演算法模仿了人類在三維立體空間中繪製遮罩的軌跡。此外，先前的三維立體場景補全方法只需對缺失區域周圍的結構和顏色進行插值即可達到不錯的效果，但這些方法無法很好地完成三維立體場景修復這項具有挑戰性但實用的任務，因此我們針對此任務設計了雙流對抗式生成網路。首先，我們的雙流生成式網路結合了三維立體場景中結構與顏色的資訊，以生成具有明確語義邊界的場景並解決了先前方法中插值的問題。為了進一步加強場景中的細節，我們提出了輕量級的雙流鑑別式網路將生成場景的結構與顏色邊緣規範化，使其更加逼真與清晰。我們用提出的 FF-Matterport 資料集進行了實驗。定性和定量的結果都驗證了我們提出的方法優於現有三維立體場景補全方法且所有提出的架構皆有其效果。

關鍵字：深度學習、對抗式生成網路、任意形狀修復、三維立體場景、三維立體場景修復



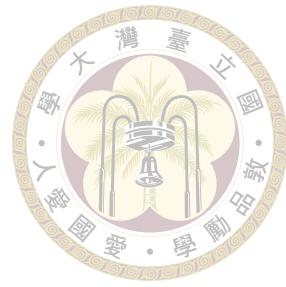
Abstract

Nowadays, the need for user editing in a 3D scene has rapidly increased due to the development of AR and VR technology. However, the existing 3D scene completion task (and datasets) cannot suit the need because the missing regions in scenes are generated by the sensor limitation or object occlusion. Thus, we present a novel task named free-form 3D scene inpainting. Unlike scenes in previous 3D completion datasets preserving most of the main structures and hints of detailed shapes around missing regions, the proposed inpainting dataset, FF-Matterport, contains large and diverse missing regions formed by our free-form 3D mask generation algorithm that can mimic human drawing trajectories in 3D space. Moreover, prior 3D completion methods cannot perform well on this challenging yet practical task, simply interpolating nearby geometry and color context. Thus, a tailored dual-stream GAN method is proposed. First, our dual-stream generator, fusing both geometry and color information, produces distinct semantic boundaries and solves the interpolation issue. To further enhance the details, our lightweight dual-stream dis-

criminator regularizes the geometry and color edges of the predicted scenes to be realistic and sharp. We conducted experiments with the proposed FF-Matterport dataset. Qualitative and quantitative results validate the superiority of our approach over existing scene completion methods and the efficacy of all proposed components.



Keywords: Deep Learning, Generative Adversarial Network, Free-form 3D Scene Inpainting, 3D Scene, 3D Scene Inpainting



Contents

	Page
Acknowledgements	i
摘要	ii
Abstract	iii
Contents	v
List of Figures	vii
List of Tables	ix
Chapter 1 Introduction	1
Chapter 2 Related Work	5
2.1 3D Completion	5
2.2 2D Image Inpainting	6
Chapter 3 Method	8
3.1 Free-form 3D Dataset Generation	8
3.2 Dual-stream Generator	10
3.3 Dual-stream Discriminator	11
Chapter 4 Experiments	14
4.1 Experimental Settings	14
4.2 Main Results	15

4.3	Ablation Studies	17
References		
Appendix A — Appendix for Free-form 3D Mask Generation Algorithm		26
Appendix B — Appendix for Additional Results		28
B.1	Edge Discriminator Qualitative Result	28
B.2	Limitation	29





List of Figures





List of Tables

- | | | |
|-----|--|----|
| 4.1 | The comparison of 3D geometric and color inpainting performance on the FF-Matterport. With mask information and proper color and geometry interaction, our dual-stream GAN outperforms all the baselines in both geometry and color metrics. Notably, models without masks all fail in the CD metric due to generating redundant meshes outside the missing areas, which verifies the importance of masks on the 3D scene inpainting task. | 15 |
| 4.2 | The ablation studies of 3D geometric and color inpainting performance on the FF-Matterport. Compared with the single-stream GAN, our dual-stream GAN reaches higher performance in all metrics, especially the three geometry metrics. Also, we show the effectiveness of 3D Gated-Conv and mask components on the 3D scene inpainting task. | 15 |



Chapter 1 Introduction

In recent years, Augmented Reality (AR) and Virtual Reality (VR) have become popular in our daily life, such as VR gaming, virtual tours, and AR meeting software. To form realistic 3D scenes in these applications, reconstructing scenes from multiple sensed RGB-D images is a widely-used and cost-effective approach. However, users often want to further edit the reconstructed 3D scenes to meet their needs. Take, as an example, removing unwanted objects existing in the real world. Therefore, there is still a strong need for 3D post-processing.

In this work, we introduce this application as the 3D scene inpainting task. Specifically, given a 3D scene with several manually specified 3D masks, a 3D inpainting model should fill these regions with proper contents, including both geometry and color. Although 3D completion methods [38, 8, 6, 10] also aim to complete missing regions in 3D space, they cannot meet the needs of the 3D inpainting application for the following reasons:

Insufficient Evaluation: The existing 3D scene completion datasets [6, 10] are unsuitable for the inpainting task because they (1) lack the masks that specify the missing regions and (2) form the missing regions only by sensor limitation or object occlusion. About (1), regions with a 0-value could be either missing parts that need to be repaired

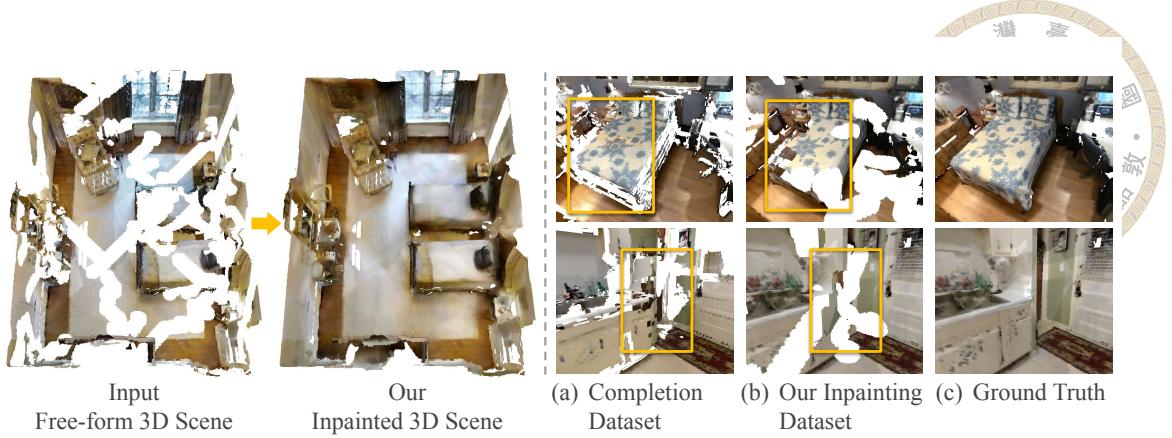


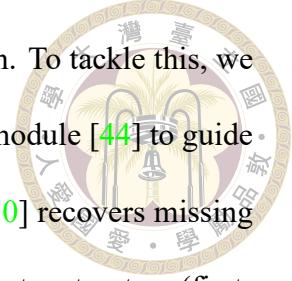
Figure 1.1: **Left:** We propose a novel 3D scene inpainting task with the first free-form 3D scene dataset, FF-Matterport, which imitates user drawing masks in 3D space. Our proposed model takes the incomplete scene as input and recovers the missing parts with high-quality and realistic results. **Right:** Compared with the prior scene completion dataset generated by incomplete observations and preserving main structures and hints of missing parts, our dataset contains challenging yet practical missing regions for the 3D scene inpainting task.

or an empty space without objects due to the sparseness property of 3D space. Without the informative masks of missing regions, SOTA methods in 3D completion thus change the shape of completed parts or leave the incomplete regions with small holes or artifacts (Fig. 4.1). Regarding (2), the physical properties of missing regions in 3D inpainting tasks and existing 3D scene completion datasets [6, 10] differ. To be specific, incomplete areas in the existing completion datasets are due to inherent sensor limitation or object occlusion. They are usually regular and strongly correlated to object occlusion or specific camera view (Fig. 1.1). In contrast, the missing regions in 3D inpainting tasks are irregularly shaped and could randomly occur anywhere in a scene. To effectively validate the methods in the 3D inpainting task, we propose a novel *Free-Form Matterport3D (FF-Matterport) dataset* tailored for this task. Free-form masks imitate the diverse human drawing trajectories in real-case (Sec. 3.1).

Poor Geometry and Color Reconstruction: As aforementioned, previous methods lack the ability to recover the missing regions well in 3D inpainting tasks. First, they can-

not attend to the missing parts due to the lacking of mask information. To tackle this, we leverage the mask information via modifying the gated convolution module [44] to guide our model to focus on the crucial area. Second, the SOTA method [10] recovers missing areas with a single-stream two-stage generator, which generates geometry structure (first-stage) and color (second-stage) sequentially. Since the unstable geometry results learned in the first stage are without the help of semantic features from color, it may cause the error propagation issue resulting in poor inpainting results. Observing this, we introduce the first *3D dual-stream generator* (Fig. 2.1 (a)) to collaboratively generate both geometry and color in missing regions (Sec. 3.2). By considering the information of two modalities simultaneously, our dual-stream generator produces more realistic object structures than the single-stream generator.

Crude Details: As the 3D scene inpainting task is much more challenging, prior approaches show the further shortcoming, i.e., producing crude details of restoration. In brief, prior colored 3D completion work [10] leveraged a color discriminator to make the rendered images of the generated scene similar to the real scene with adversarial loss. However, this color adversarial discriminator still results in blurred and distorted boundaries when recovering large damaging regions with structural contents, such as picture frames or furniture edges. To enhance the sharpness and structure of boundaries, we introduce an extra edge stream discriminator apart from the color adversarial discriminator. Particularly, we randomly project the inpainted 3D scene to multiple 2D images and then simultaneously constrain its color texture along with the predicted edge maps using our *dual-stream discriminator* (Fig. 2.1 (b) and Sec. 3.3). With additional regularizations on the edges of global structures, we find that our geometry and color generators learn to collaborate and generate less blurry and well-structured boundaries.



We validate our dual-stream GAN and previous 3D scene completion methods on the proposed FF-Matterport dataset. Our method demonstrates the superiority in six different metrics. Besides, the impressive visualization results also illustrate the effectiveness of our method. To sum up, this work presents the following main novelties and contributions:

- We propose a novel 3D scene inpainting task with the first free-form 3D scene dataset, FF-Matterport, which contains diverse free-form masks generated with our designed algorithm imitating human drawing trajectories.
- We introduce the first 3D gated dual-stream (geometry and color) generator to jointly consider the geometry and color context of missing regions and generate high-quality contents with semantic-constrained structures.
- We introduce an edge and color dual-stream discriminator guiding the generator to produce clear and detailed geometry and color boundaries.



Chapter 2 Related Work

2.1 3D Completion

In the field of 3D vision, object completion is a fundamental and long-standing problem. Some took single or multi images to reconstruct or complete 3D objects [5, 30, 35, 39, 42, 16]. Others utilized depth or RGB-D frames collected by commodity depth sensors to reconstruct 3D objects [41, 40]. Still, others aimed to complete a 3D object by various 3D representations, such as 3D point cloud scan [18, 43, 46, 36, 23, 45], sign distance fields (SDF) [9, 34], or mesh surface [25, 26]. Nonetheless, the above methods are limited to 3D objects rather than a complex 3D scene with several items.

To complete a complicated 3D scene, some recent studies have been proposed. [32, 24, 7, 20, 21, 29] leveraged scanned multiple RGB-D images to reconstruct a 3D scene and refine its color. SSCNet [38] combined the scene completion task with the 3D semantic segmentation task. ScanComplete [8] extended the 3D completion task to large scenes and designed the network to handle various scene scales during inferring. SG-NN [6] first trained and evaluated the 3D scene completion model on a real-world scanned dataset, and SPSG [10] first tackled color completion in 3D scenes apart from geometry completion. However, these prior 3D scene completion works only focus on completing missing parts due to sensor limitations, which are strongly correlated to viewing angles and have

generally similar patterns.

To the best of our knowledge, we are the first to introduce the 3D scene inpainting task and generate a free-form 3D scene dataset to train and evaluate the performance of completing manual masks in 3D space. Furthermore, our dual-stream GAN, complementing geometry and color information with each other, solves the problem of over-smoothed geometry shapes and blurred color boundaries in prior single-stream two-stage work [10].



2.2 2D Image Inpainting

2D image inpainting takes a corrupted image as input and fills the missing parts in the image with semantically correct, and boundary-consistent contents. It is an important task for many downstream visual tasks, such as object removal, damaged photo restoration, and 2D to 3D photo transformation [37]. Traditional approaches [1, 13, 11, 22] reused the patches from the image background or source images to repair the missing pixels with the most similar one, but they only can handle repetitive patterns or small missing holes.

Recently, GAN [15] has made great progress in the image inpainting task, enabling inpainting models to fill holes with realistic and semantically reasonable contents [12]. PConv [27] and GatedConv [44] extended regular rectangle masks to free-form masks with irregular shapes and developed corresponding CNN modules to handle the more challenging masks. More recently, some practices [31, 28, 17] utilized additional edge constraints to conquer the blurry results on large missing areas where the main structure of the object is missing. [31] proposed an edge-color two-stage inpainting framework; [28] and [17] developed new generator models to combine and exchange structure and texture information.

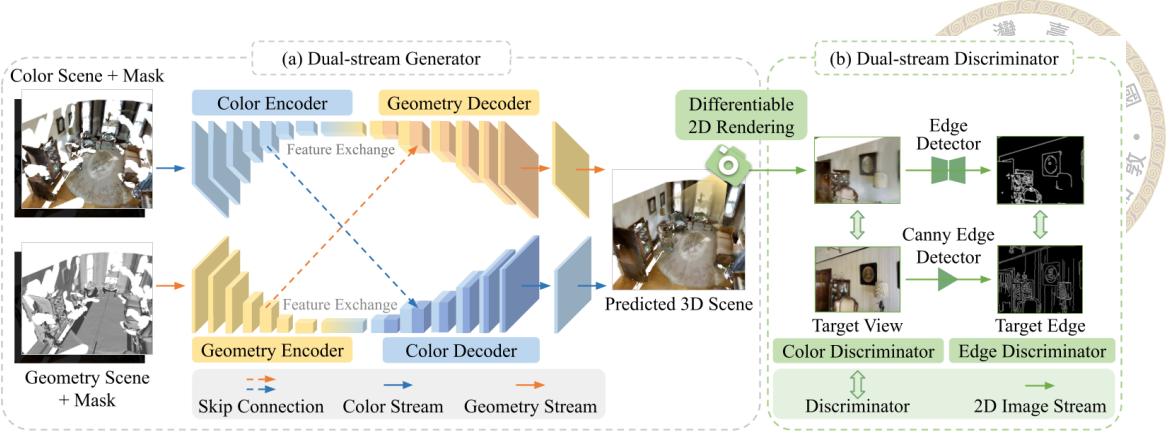


Figure 2.1: Overview of the proposed network. (a) Geometry and color dual-stream generator exchanges and fuses embedded features from each other to complement respective decoders. (b) The dual-stream discriminator uses the differentiable 2D rendering and the edge detector to project the predicted scene to a 2D image and an edge map. Then, it optimizes them with the target view and corresponding canny edge image.

Different from prior 2D inpainting studies, we first introduce the 3D inpainting task and present a 3D-specific free-form mask generation algorithm due to the sparseness property in 3D space. Moreover, inspired by [31, 28, 17], we propose the first dual-stream GAN for the 3D inpainting task, not only cooperating 3D geometry and color information in the generator but also regularizing color and edges from diverse viewpoints in the discriminator.



Chapter 3 Method

In this work, we propose a novel 3D scene inpainting task with the tailored FF-Matterport dataset for the need in AR and VR applications. The 3D scenes in the dataset contain irregular and diverse shapes of missing regions (masks) generated by our novel free-form 3D mask generation algorithm (Sec. 3.1) that mimics humans drawing trajectories in 3D scenes. To tackle this challenging yet practical task, we develop a dual-stream GAN model (Fig. 2.1) that contains two main components, the dual-stream generator and the dual-stream discriminator. The dual-stream generator (Sec. 3.2) leverages the mask information and the feature fusion of geometry and color to generate semantic-constraint structures and shape-constraint textures. In addition, the dual-stream discriminator (Sec. 3.3) further enhances the details by regularizing color and corresponding edges on randomly rendered images.

3.1 Free-form 3D Dataset Generation

As illustrated in Fig. 1.1, our free-form 3D dataset aims to alleviate the issues of too regular patterns in missing areas and the lacking mask information in existing 3D completion datasets [6, 10]. To generate free-form 3D masks, it is intuitive to modify the free-form 2D mask generation algorithm [44], which uses strokes with random lengths

and directions to line up a drawing track. As all pixels in 2D images contain information, drawing straight lines back and forth can cover the unwanted objects. However, since objects in 3D space are sparse and have curved and complicated shapes, directly applying the previous algorithm in 3D space usually masks areas without objects or produces meaningless shapes (Fig. 3.1).

To conquer the above challenges, we design a novel 3D mask generation algorithm. Initially, it converts the 3D scene data to the truncated signed distance field (TSDF) representation.

The algorithm can thus ensure that the painted stroke persists around the surface by checking the TSDF values. Afterward, we use an incremental masking strategy rather than the orig-

inal 2-point line strategy in [44]. Our strategy dynamically decides the direction and length of strokes, considering the distance to the surface and the diameter of the stroke. Besides, we randomly sample points in 3D scenes as the starting point for strokes to ensure diversity. With our algorithm, the curved strokes can fit various object shapes and occur in diverse places in 3D scenes. More details are reported in supplementary materials.

We apply our mask generation algorithm to the Matterport3D [4] dataset with the official train-test split and produce the first free-form 3D scene inpainting dataset, named FF-Matterport. It contains 30-40% missing regions randomly located in the whole indoor

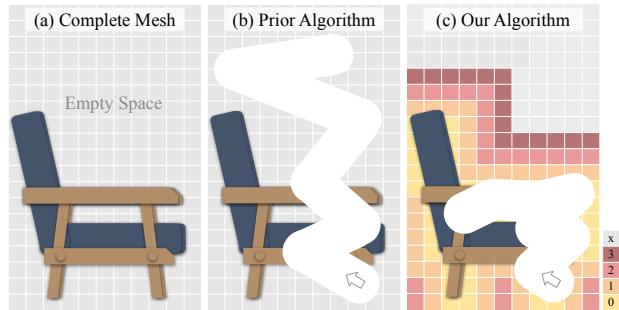
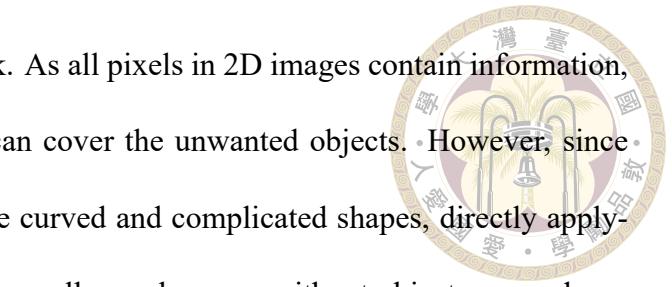
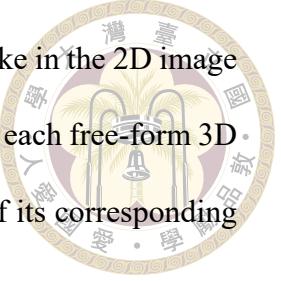


Figure 3.1: Trajectories comparison between (b) previous algorithm and (c) our algorithm when drawing in (a). Due to the sparseness of 3D space, algorithm (b) [44] generates masks on empty space and remains weird object shapes. In contrast, our free-form 3D algorithm (c) is able to produce practical masks around the surface with better flexibility by utilizing the property of TSDF representation.

scene and is voxelized to 2cm resolution same as [6, 10]. Notably, unlike in the 2D image inpainting task where a mask can be randomly paired with any images, each free-form 3D mask in our dataset is generated according to the object distribution of its corresponding scene.



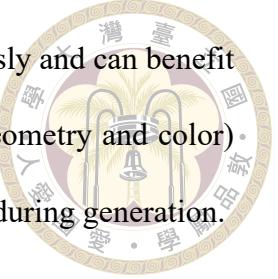
3.2 Dual-stream Generator

As mentioned in Sec. 3.1, the scenes in the 3D inpainting task contain irregular and various missing regions, which makes the task more challenging. Moreover, the additional mask is also provided to indicate areas that are needed to focus. To better leverage the mask information and reconstruct sophisticated and realistic object surfaces in missing regions, we develop a dual-stream generator specifically for this task.

To utilize the masks, we can treat them as additional channel inputs. Nonetheless, we further exploit the benefits of masks, attaching a 3D gated convolution module (3D GatedConv) extended from [44] to the generator. Specifically, the 3D GatedConv module helps the generator gradually fill the masked regions with proper geometry and color contents by dynamically learning soft attention maps. This modification contributes significantly to model performance, and we found that baselines without the mask information or 3D GatedConv module only produce distorted restoration results (See Tab. 4.2).

Regarding the pipeline of the generator, SPSG [10] developed a two-stage pipeline, which first completes the geometry of all missing areas and then generates color on the surfaces specified in the first stage. We find that the geometry generator in this pipeline completely ignores the semantic features of color. Also, the unstable geometry outcome of the first stage causes error propagation to the second stage. To this end, we argue that

the geometry and color information should be considered simultaneously and can benefit each other. Inspired by [17], we develop the first 3D dual-stream (geometry and color) generator that can fuse and retrieve the knowledge from both streams during generation.



In Fig. 2.1 (a), our dual-generator consists of two generators, each with a U-Net variant. In the encoding phase, the geometry and color scenes are embedded independently and projected to high-level feature space through corresponding generators. During the decoding phase, the two generators fuse feature embedding from each other as an additional condition to refine the respective decoded results. Besides, we combine the encoder and decoder features with skip connections to create more delicate content. This operation allows us to fully exchange geometry and color information during generation, producing both semantic-aware geometry structure and shape-constrained color texture in predicted scenes. Compared with the previous pipeline using a one-way feature stream forward from geometry to color, our dual-stream pipeline provides a mutual feature exchange between the two generators. Consequently, we can observe that our pipeline alleviates the error propagation problem, revealing consistent performance improvements in both geometry and color results (Tab. 4.2).

3.3 Dual-stream Discriminator

To make the generator produce high-quality objects on missing regions, it is a common practice to directly regularize generated scenes by designing loss functions or utilizing a discriminator. In the field of 3D completion, [6] used a naive ℓ_1 loss to regress the geometry outputs and [10] applied a 2D discriminator to force the color outputs realistic under diverse rendering views. Nevertheless, the above practices are inadequate to

meet the need of 3D scene inpainting as its missing areas, unlike small corrupted strips due to sensor limitation in 3D completion, are generally large and lack structural contents. Specifically, we found those methods are prone to produce over-smoothed structures and blurred color edges by interpolating nearby color and geometry values or linking mesh pieces (Fig. 4.1). Thus, we develop a novel dual-stream discriminator to avoid producing such crude details.

As illustrated in Fig. 2.1 (b), our dual-stream discriminator is composed of two components: a color discriminator and an edge discriminator. To begin with, our color discriminator regularizes the quality of generated scenes on randomly rendered 2D images following [10]. This 2D color stream makes the projected 2D frames of the generated 3D scene as realistic as possible and is more effective than applying a discriminator in 3D [10]. Moreover, to ensure fine-grained geometry shapes and sharp color boundaries in the generated scene, we design an edge discriminator further regularizing the corresponding edges of the rendered frames. To elaborate, our edge discriminator compares the 2D edge maps extracted from real projected frames by the Canny edge detector [3] and that extracted from generated frames by our NN-based edge detector. With the aid of our lightweight edge discriminator, we can not only produce sharp color boundaries and detailed geometry contents on qualitative results (Fig. 4.1) but also achieve huge improvements in quantitative results (Tab. 4.2).

To summarize, our training objectives can be categorized into two groups. For the naive full supervision loss, we supervise our geometry and color reconstruction from both 3D and 2D space. The L_{geo} following the log-transformed ℓ_1 TSDF loss in [6] is applied to penalize the geometry in 3D space, and the L_{color} as well as depth L_{depth} (both are ℓ_1 losses) are leveraged to ensure geometry and color reconstruction on the rendered images.

Furthermore, to improve human visual perception and sharpen the geometry and color boundaries, we applied the two adversarial losses, L_{color}^{adv} and L_{edge}^{adv} , from our dual-stream discriminator along with the conventional content loss L_{cont} [14]. The overall loss is formulated as below:

$$L = \lambda_1 L_{geo} + \lambda_2 L_{color} + L_{depth} + \lambda_3 L_{cont} + \lambda_4 (L_{color}^{adv} + L_{edge}^{adv}), \quad (3.1)$$

where $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ are the scaling coefficients.



Chapter 4 Experiments

4.1 Experimental Settings

Training Settings: Our network is trained on a single NVIDIA GeForce RTX 2080 Ti with a batch size of 2, and it takes about 6 epochs ≈ 48 hours to train until convergence. It is optimized via an Adam optimizer with a learning rate of 0.0001. The patch size of the discriminator is 94 x 94 cropped from 320 x 256 images. The $\lambda_1, \lambda_2, \lambda_3, \lambda_4$ in Eq. 3.1 are set as 0.3, 0.6, 0.01 and 0.005 via grid search. In the training stage, we crop the 3D scene into 64 x 64 x 128 chunk voxels to speed up the process. In the testing stage, we directly input the room-sized scene to our model as 3D CNN is invariant to the scene scale.

Evaluation Metrics: For a fair comparison, we follow the evaluation metric in [10]. The geometry performance is evaluated by IoU, Recall, and Chamfer Distance. Note that only the observed regions in the target scene are evaluated, and we ignore the unobserved areas as the same in [10]. Besides, the color performance is evaluated by SSIM (structural similarity image metric) [2], Feature- ℓ_1 [33], and FID (Fréchet Inception Distance) [19] to capture the differences at both local and global scales between the rendered and target images.

Baselines: To verify the proposed dual-stream GAN in our novel task, we compare it

with several SOTA 3D completion approaches, including PIFu⁺ [35], SG-NN [6], and SPSG [10]. Also, we develop another baseline, SPSG equipped with mask inputs, to validate the importance of mask information. We follow [10] to implement all baselines and exclude SG-NN from the evaluation of color performance as it is designed to complete geometry only.

Methods	Geometry			Color		
	IoU(\uparrow)	Recall(\uparrow)	CD(\downarrow)	SSIM(\uparrow)	Feature- ℓ_1 (\downarrow)	FID(\downarrow)
PIFu ⁺ [35]	0.241	0.525	19.537	0.744	0.253	108.87
SG-NN [6]	0.636	0.857	20.988	-	-	-
SPSG [10]	0.466	0.683	17.457	0.829	0.220	75.10
SPSG (+mask) [10]	0.491	0.659	3.336	0.843	0.214	69.60
Ours	0.781	0.896	2.110	0.853	0.209	65.28

Table 4.1: The comparison of 3D geometric and color inpainting performance on the FF-Matterport. With mask information and proper color and geometry interaction, our dual-stream GAN outperforms all the baselines in both geometry and color metrics. Notably, models without masks all fail in the CD metric due to generating redundant meshes outside the missing areas, which verifies the importance of masks on the 3D scene inpainting task.

Methods	Geometry			Color		
	IoU(\uparrow)	Recall(\uparrow)	CD(\downarrow)	SSIM(\uparrow)	Feature- ℓ_1 (\downarrow)	FID(\downarrow)
Dual-stream GAN (Full)	0.781	0.896	2.110	0.8536	0.209	65.28
- Edge Discriminator	0.774	0.892	2.137	0.8534	0.209	65.62
- 3D GatedConv	0.747	0.875	2.250	0.8491	0.211	68.28
- Mask Info.	0.592	0.827	24.356	0.8471	0.213	69.75
Single-stream GAN	0.744	0.878	2.523	0.8511	0.210	66.68

Table 4.2: The ablation studies of 3D geometric and color inpainting performance on the FF-Matterport. Compared with the single-stream GAN, our dual-stream GAN reaches higher performance in all metrics, especially the three geometry metrics. Also, we show the effectiveness of 3D GatedConv and mask components on the 3D scene inpainting task.

4.2 Main Results

We verify all methods on the FF-Matterport dataset and illustrate the quantitative and qualitative results in Tab. 4.1 and Fig. 4.1, respectively. From Tab. 4.1, PIFu⁺ obtains the worst performance in most metrics, indicating that it cannot adapt to tackle this novel task.

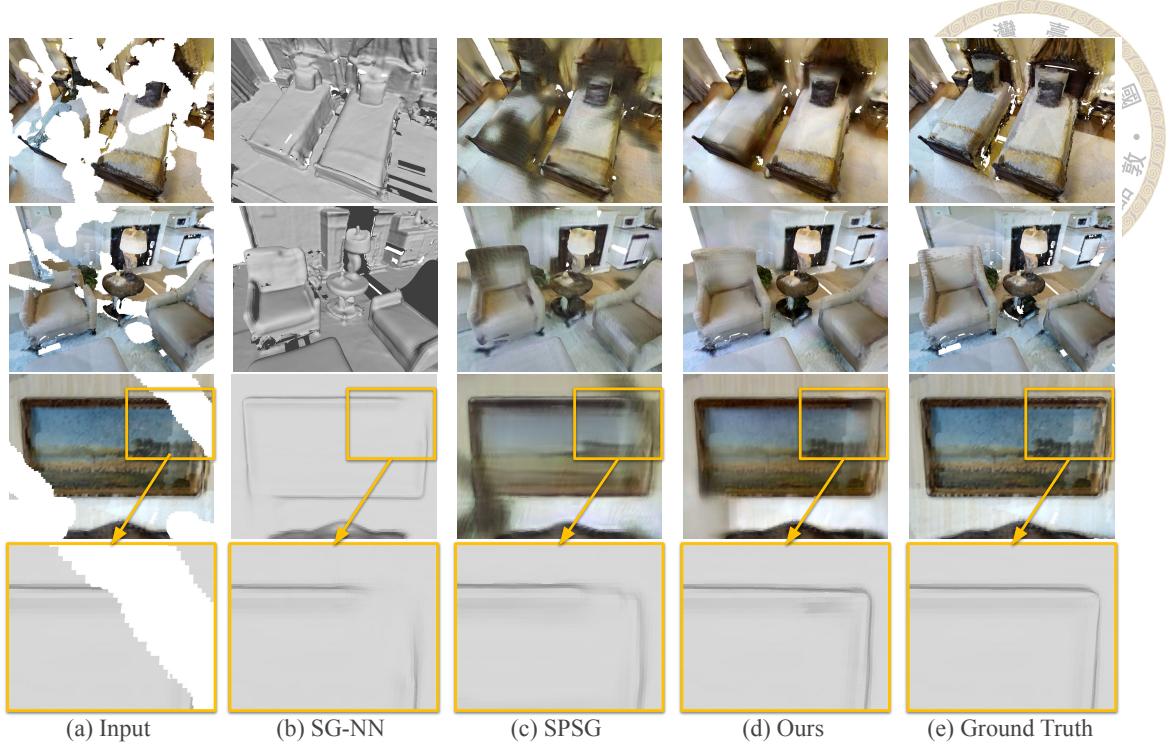


Figure 4.1: Qualitative results of all methods on the FF-Matterport (Best viewed in zoomed digital). Compared to our dual-stream GAN, SPSG fails to preserve details from input and SG-NN predicts distorted structures with small holes in large missing areas in the 1st and 2nd rows. To further show the fine details of our predicted mesh, the 4th row zoom in on the missing corner of the 3rd row. More analyses and discussions are reported in Sec. 4.2.

About SG-NN, it captures more local features, resulting in higher IoU and Recall scores than other baselines. But it fails to capture the global features and generates distorted structure with small holes in large missing areas, such as the corner of the bed (1st row) and the sofa chair on the left (2nd row) of Fig. 4.1 (b). Thus, it leads to bad performance on the CD metrics.

Regarding SPSG, the 1st row in Fig. 4.1 (c) shows that it fails to preserve the details of complete parts from input scenes, such as the lamp beside the sofa, and results in lower IoU and Recall scores. Besides, SPSG suffers from initial color bias pointed in [44] and performs poorly on the color scores. On the contrary, our dual-stream GAN better utilizes the mask information as well as both structure and semantic feature from dual-stream to generate well-structured and semantic reasonable 3D scenes; accordingly, we achieve the

highest performance in all metrics.

To show the further restoration details, the 4th row of Fig. 4.1 zooms in on the missing corner of the frame in the 3rd row. Unlike SG-NN and SPSG flattening and ignoring the frame structure, our model easily distinguishes the frame from the wall using color features and recovers the corner with straight and delicate edges guided by our edge adversarial loss.

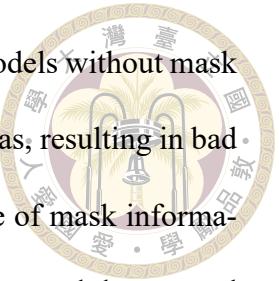
4.3 Ablation Studies

We summarize the ablation study in Tab. 4.2. First, we analyze the influence of our edge discriminator. Comparing the 1st and 2nd rows, the edge discriminator causes minor improvements in numerical evaluations but significantly contributes to the visualization (in supplementary). This phenomenon echoes our hypothesis that 2D edge loss can guide 3D geometry and color to better collaborate on generating delicate 2D edges. Then, we verify the design of our pipeline. We build a single-stream sequential generator with mask inputs and 3D GatedConv, named single-stream GAN in Tab. 4.2. The single-stream GAN generates geometry features without knowing color features and then passes them to the color stream. As a result, its geometry performance declined larger than color performance.

Lastly, we examine the efficacy of 3D GatedConv and mask information. As shown in the 3rd and 4th rows of Tab. 4.2, they substantially improve the performance, both in geometry and color scores, which reveals the mask information is indispensable in the novel 3D inpainting task. This observation is consistent with the advancement of SPSG with mask inputs in the main experiment (Tab. 4.1). Especially the CD metric is evaluated on



the fixed-number points uniformly sampled from predicted meshes, models without mask information tend to generate redundant meshes outside the missing areas, resulting in bad performance. Through these analyses, we demonstrate the importance of mask information on the challenging 3D scene inpainting task, which means the proposed dataset and approach are requisite.

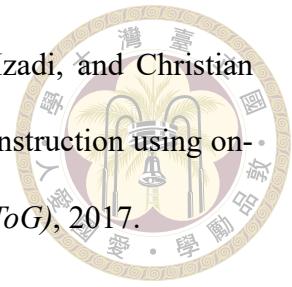




References

- [1] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 2009.
- [2] Dominique Brunet, Edward R Vrscay, and Zhou Wang. On the mathematical properties of the structural similarity index. *IEEE Transactions on Image Processing*, 2011.
- [3] John Canny. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence*, 1986.
- [4] 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. *arXiv preprint arXiv:1709.06158*, 2017.
- [5] Christopher B Choy, Danfei Xu, JunYoung Gwak, Kevin Chen, and Silvio Savarese. 3d-r2n2: A unified approach for single and multi-view 3d object reconstruction. In *European conference on computer vision*. Springer, 2016.
- [6] Angela Dai, Christian Diller, and Matthias Nießner. Sg-nn: Sparse generative neural networks for self-supervised scene completion of rgb-d scans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.

- [7] Angela Dai, Matthias Nießner, Michael Zollhöfer, Shahram Izadi, and Christian Theobalt. Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration. *ACM Transactions on Graphics (ToG)*, 2017.
- [8] Angela Dai, Daniel Ritchie, Martin Bokeloh, Scott Reed, Jürgen Sturm, and Matthias Nießner. Scancomplete: Large-scale scene completion and semantic segmentation for 3d scans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018.
- [9] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017.
- [10] Angela Dai, Yawar Siddiqui, Justus Thies, Julien Valentin, and Matthias Nießner. Spsg: Self-supervised photometric scene generation from rgb-d scans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021.
- [11] Soheil Darabi, Eli Shechtman, Connelly Barnes, Dan B Goldman, and Pradeep Sen. Image melding: Combining inconsistent images using patch-based synthesis. *ACM Transactions on graphics (TOG)*, 2012.
- [12] Emily L Denton, Soumith Chintala, Rob Fergus, et al. Deep generative image models using a laplacian pyramid of adversarial networks. *Advances in neural information processing systems*, 2015.
- [13] Alexei A Efros and William T Freeman. Image quilting for texture synthesis and transfer. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, 2001.

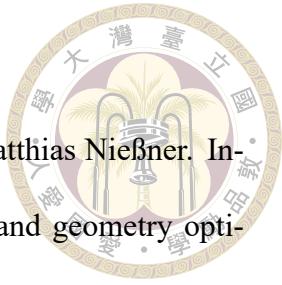




- [14] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016.
- [15] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 2014.
- [16] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. A papier-mâché approach to learning 3d surface generation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018.
- [17] Xiefan Guo, Hongyu Yang, and Di Huang. Image inpainting via conditional texture and structure dual generation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.
- [18] Xiaoguang Han, Zhen Li, Haibin Huang, Evangelos Kalogerakis, and Yizhou Yu. High-resolution shape completion using deep neural networks for global structure and local geometry inference. In *Proceedings of the IEEE international conference on computer vision*, 2017.
- [19] 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. *Advances in neural information processing systems*, 2017.
- [20] Jingwei Huang, Angela Dai, Leonidas J Guibas, and Matthias Nießner. 3dlite: towards commodity 3d scanning for content creation. *ACM Trans. Graph.*, 2017.
- [21] Jingwei Huang, Justus Thies, Angela Dai, Abhijit Kundu, Chiyu Jiang, Leonidas J Guibas, Matthias Nießner, Thomas Funkhouser, et al. Adversarial texture optimization



- tion from rgb-d scans. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [22] Jia-Bin Huang, Sing Bing Kang, Narendra Ahuja, and Johannes Kopf. Image completion using planar structure guidance. *ACM Transactions on graphics (TOG)*, 2014.
- [23] Zitian Huang, Yikuan Yu, Jiawen Xu, Feng Ni, and Xinyi Le. Pf-net: Point fractal network for 3d point cloud completion. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2020.
- [24] Shahram Izadi, David Kim, Otmar Hilliges, David Molyneaux, Richard Newcombe, Pushmeet Kohli, Jamie Shotton, Steve Hodges, Dustin Freeman, Andrew Davison, et al. Kinectfusion: real-time 3d reconstruction and interaction using a moving depth camera. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*, 2011.
- [25] Michael Kazhdan, Matthew Bolitho, and Hugues Hoppe. Poisson surface reconstruction. In *Proceedings of the fourth Eurographics symposium on Geometry processing*, 2006.
- [26] Michael Kazhdan and Hugues Hoppe. Screened poisson surface reconstruction. *ACM Transactions on Graphics (ToG)*, 2013.
- [27] Guilin Liu, Fitsum A Reda, Kevin J Shih, Ting-Chun Wang, Andrew Tao, and Bryan Catanzaro. Image inpainting for irregular holes using partial convolutions. In *Proceedings of the European conference on computer vision (ECCV)*, pages 85–100, 2018.
- [28] Hongyu Liu, Bin Jiang, Yibing Song, Wei Huang, and Chao Yang. Rethinking image inpainting via a mutual encoder-decoder with feature equalizations. In *European*



- [29] Robert Maier, Kihwan Kim, Daniel Cremers, Jan Kautz, and Matthias Nießner. Intrinsic3d: High-quality 3d reconstruction by joint appearance and geometry optimization with spatially-varying lighting. In *Proceedings of the IEEE international conference on computer vision*, pages 3114–3122, 2017.
- [30] Lars Mescheder, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. Occupancy networks: Learning 3d reconstruction in function space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019.
- [31] Kamyar Nazeri, Eric Ng, Tony Joseph, Faisal Z Qureshi, and Mehran Ebrahimi. Edgeconnect: Generative image inpainting with adversarial edge learning. *arXiv preprint arXiv:1901.00212*, 2019.
- [32] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. Kinectfusion: Real-time dense surface mapping and tracking. In *2011 10th IEEE international symposium on mixed and augmented reality*. Ieee, 2011.
- [33] Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields: Learning texture representations in function space. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- [34] Jeong Joon Park, Peter Florence, Julian Straub, Richard Newcombe, and Steven Lovegrove. Deepsdf: Learning continuous signed distance functions for shape representation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2019.



- [35] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. Pifu: Pixel-aligned implicit function for high-resolution-clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2019.
- [36] Muhammad Sarmad, Hyunjoo Jenny Lee, and Young Min Kim. Rl-gan-net: A reinforcement learning agent controlled gan network for real-time point cloud shape completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2019.
- [37] Meng-Li Shih, Shih-Yang Su, Johannes Kopf, and Jia-Bin Huang. 3d photography using context-aware layered depth inpainting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020.
- [38] Shuran Song, Fisher Yu, Andy Zeng, Angel X Chang, Manolis Savva, and Thomas Funkhouser. Semantic scene completion from a single depth image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017.
- [39] Yongbin Sun, Ziwei Liu, Yue Wang, and Sanjay E Sarma. Im2avatar: Colorful 3d reconstruction from a single image. *arXiv preprint arXiv:1804.06375*, 2018.
- [40] Jiajun Wu, Chengkai Zhang, Xiuming Zhang, Zhoutong Zhang, William T Freeman, and Joshua B Tenenbaum. Learning shape priors for single-view 3d completion and reconstruction. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018.
- [41] Zhirong Wu, Shuran Song, Aditya Khosla, Fisher Yu, Linguang Zhang, Xiaou Tang, and Jianxiong Xiao. 3d shapenets: A deep representation for volumetric shapes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1912–1920, 2015.

- [42] Qiangeng Xu, Weiyue Wang, Duygu Ceylan, Radomir Mech, and Ulrich Neumann. Disn: Deep implicit surface network for high-quality single-view 3d reconstruction. *Advances in Neural Information Processing Systems*, 32, 2019.
- [43] Yaoqing Yang, Chen Feng, Yiru Shen, and Dong Tian. Foldingnet: Interpretable unsupervised learning on 3d point clouds. *arXiv preprint arXiv:1712.07262*, 2017.
- [44] Jiahui Yu, Zhe Lin, Jimei Yang, Xiaohui Shen, Xin Lu, and Thomas S Huang. Free-form image inpainting with gated convolution. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2019.
- [45] Xumin Yu, Yongming Rao, Ziyi Wang, Zuyan Liu, Jiwen Lu, and Jie Zhou. Pointr: Diverse point cloud completion with geometry-aware transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 2021.
- [46] Wentao Yuan, Tejas Khot, David Held, Christoph Mertz, and Martial Hebert. Pcn: Point completion network. In *2018 International Conference on 3D Vision (3DV)*. IEEE, 2018.





Appendix A — Appendix for Free-form 3D Mask Generation Algorithm

As stated in Sec. 3.1 of the main paper, our designed free-form 3D mask generation algorithm aims to mimic human drawing trajectories in 3D space and randomly generate diverse free-form masks for efficient training and evaluation. To avoid covering the empty space in the scene into the masked areas and to flexibly draw arbitrary shapes of 3D objects, we utilize the characteristic of TSDF and dynamically decide the length and direction of strokes based on the TSDF values. We show our algorithm in Algorithm 1.

The input of the algorithm is an original scene S_o from a real-world scanned scene containing T_o TSDF and C_o color voxelized values, and three hyper-parameters *diameter*, *maxStrokeStep*, and *totalStep* to control the mask distribution. The output is a masked scene S_m for training and inferring consisting of masked T_m TSDF and C_m color voxelized values with corresponding binary mask map M_m . This algorithm is for the 64x64x128 chunk size in the training dataset. For the test dataset which contains whole indoor rooms, We run this algorithm several times in equal proportions to the size of the indoor rooms.



Algorithm 1 Free-form 3D Mask Generation Algorithm

```

1: Input: Original scene:  $S_o = \{T_o, C_o\}$ , and  $diameter, maxStrokeStep, totalStep$ .
2: Output: Masked scene:  $S_m = \{T_m, C_m, M_m\}$ .
3: Init: Copy from original scene  $T_m \leftarrow T_o, C_m \leftarrow C_o, (\#step, \#strokeStep) \leftarrow (0, 0)$ .
4: function RANDOMMAXSTROKE( $s$ )
5:     return random.randint( $s, s+10$ )
6: end function
7: function FINDVALIDBALL( $c, d, t$ )  $\triangleright c$  is center,  $d$  is diameter, and  $t$  is threshold
8:     return  $X \leftarrow \{x | x \in B_{c,d} \cap |T_m(x)| \leq t\}$   $\triangleright B_{c,d}$  is a ball centered at  $c$  with
diameter  $d$ 
9: end function
10: center  $O_c \leftarrow$  random.choice( $|T_m| \leq 1$ )  $\triangleright$  Random a starting point on occupied voxels
11: max stroke step  $L \leftarrow$  RANDOMMAXSTROKE( $maxStrokeStep$ )
12: while  $\#step \leq TotalStep$  do
13:     Mask out  $B_{O_c, diameter}$  in  $\{T_m, C_m, M_m\}$ 
14:      $\#step \leftarrow \#step + 1$ 
15:      $\#strokeStep \leftarrow \#strokeStep + 1$ 
16:     if  $\#strokeStep \geq L$  then  $\triangleright$  Restart a new stroke
17:         center  $O_c \leftarrow$  random.choice( $|T_m| \leq 1$ )
18:         max stroke step  $L \leftarrow$  RANDOMMAXSTROKE( $maxStrokeStep$ )
19:          $\#strokeStep \leftarrow 0$ 
20:     else if  $X \leftarrow$  FINDVALIDBALL( $O_c, diameter // 2, 1 \neq \emptyset$ ) then  $\triangleright$  Move a small step
21:          $O_c \leftarrow$  random.choice( $X$ )
22:     else if  $X \leftarrow$  FINDVALIDBALL( $O_c, diameter, 5 \neq \emptyset$ ) then  $\triangleright$  Move a big step
23:          $O_c \leftarrow$  random.choice( $X$ )
24:     else  $\triangleright$  Dead end, Restart a new stroke
25:         center  $O_c \leftarrow$  random.choice( $|T_m| \leq 1$ )
26:         max stroke step  $L \leftarrow$  RANDOMMAXSTROKE( $maxStrokeStep$ )
27:          $\#strokeStep \leftarrow 0$ 
28:     end if
29: end while
30: return  $\{T_m, C_m, M_m\}$ .

```



Appendix B — Appendix for Additional Results

B.1 Edge Discriminator Qualitative Result

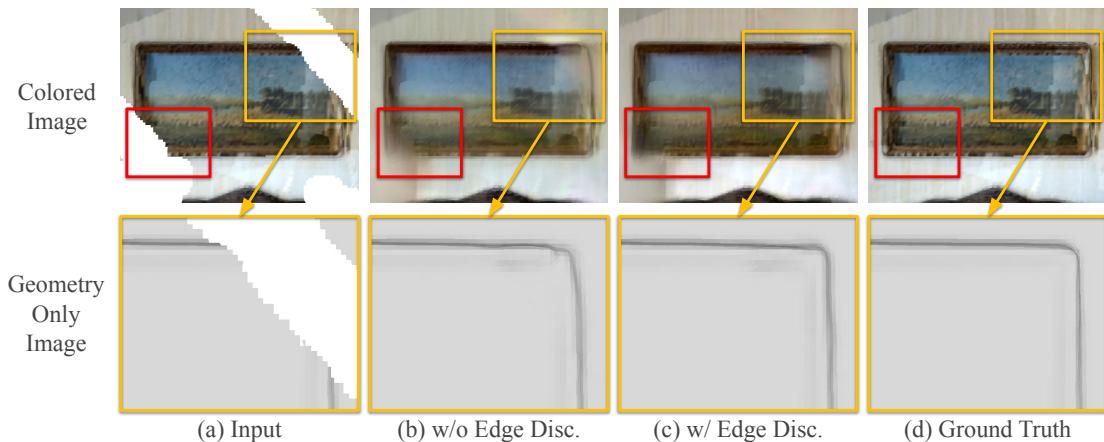


Figure B.2: Additional qualitative results of the dual-stream discriminator on FF-Matterport. By adding the edge stream to the conventional color stream discriminator, the color boundary becomes less blurred (shown in red frames), and the geometry shapes become more straight and sharper (shown in zoomed-in pictures in the 2nd row with yellow frames).

In the main paper, we state that our additional discriminator aims to enhance the sharpness and details of the results for better human perception. As shown in Fig. B.2, we provide additional qualitative results on FF-Matterport to demonstrate the efficacy of our dual-stream discriminator. Compared with (b) the model without our edge discriminator, the full model (c) performs better on both geometry and color. In (c), the bottom left corner

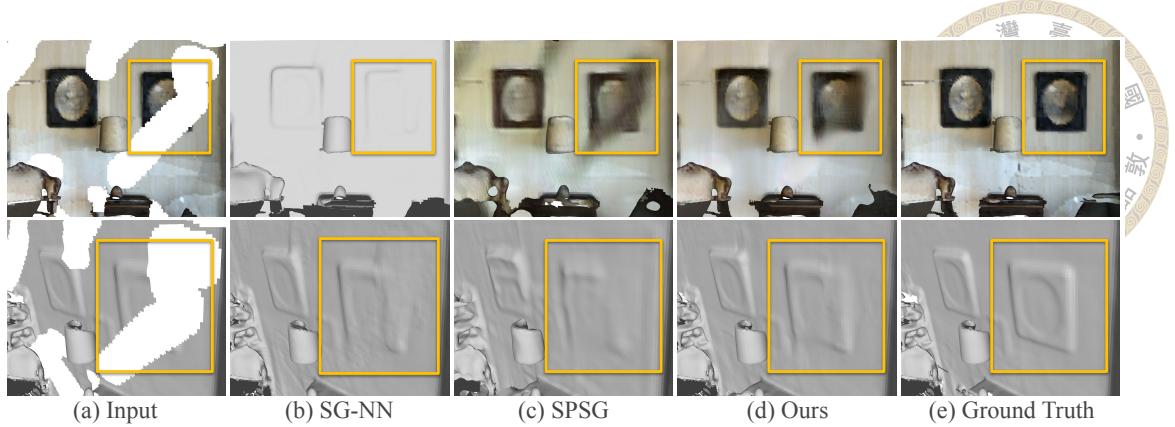


Figure B.3: A failure case of the frame on the right side in (a) losing most of the edges and color. Even though our model fails to predict distinct structures and edges in the missing parts, our results (d) still outperform the baselines SG-NN [6] (b) and SPSG [10] (c) in both geometry and color visual performance.

(red frame) of the color image becomes less blurred after adding the edge stream, and the top right corner (yellow frame) of the color image zoomed in as the geometry only image shows that the geometry shapes become sharper after adding our edge discriminator. This phenomenon echoes our motivation that 2D edge loss can simultaneously guide the 3D geometry and color streams to generate delicate edges.

B.2 Limitation

Although our proposed method can generate realistic geometry and color results for the real-world 3D scene inpainting task, we still find its limitations and unsolved challenges as shown in Fig. B.3. For example, in the 1st row (a), the right side of an incomplete picture frame loses most of the edges and color; also, the 2nd row (a) shows that the structure of the frame is very close to the wall and hard to be distinguished. Therefore, even with the help of our dual-stream GAN design, the model fails to predict the correct edges of the top right and bottom left corners, resulting in blurred color boundaries in the comparison of (d) and (e). Still, our predicted results contain more details than SG-NN [6]

(b) and SPSG [10] (c).

Moreover, due to the natural limit of CNN models and voxel representation, the output resolution is restricted. Even though some new 3D data representations are proposed recently and claimed to support the unlimited resolution, such as implicit function, it is still challenging to properly handle the mask information in the 3D scene inpainting task.

