

Multi-Frequency-Aware Patch Adversarial Learning for Neural Point Cloud Rendering

Jay Karhade¹, Haiyue Zhu^{2,†}, Ka-Shing Chung³, Rajesh Kumar Tripathy⁴, Wei Lin², and Marcelo H. Ang Jr.³

Abstract—We present a neural point cloud rendering pipeline through a novel multi-frequency-aware patch adversarial learning framework. The proposed approach aims to improve the rendering realness by minimizing the spectrum discrepancy between real and synthesized images, especially on the high-frequency localized sharpness information which causes image blur visually. Specifically, a patch multi-discriminator scheme is proposed for the adversarial learning, which combines both spectral domain (Fourier Transform and Discrete Wavelet Transform) discriminators as well as the spatial (RGB) domain discriminator to force the generator to capture global and local spectral distributions of the real images. The proposed multi-discriminator scheme not only helps to improve rendering realness, but also enhance the convergence speed and stability of adversarial learning. Moreover, we introduce a noise-resistant voxelisation approach by utilizing both the appearance distance and spatial distance to exclude the spatial outlier points caused by depth noise. Our entire architecture is fully differentiable and can be learned in an end-to-end fashion. Extensive experiments show that our method produces state-of-the-art results for neural point cloud rendering by a significant margin.

Index Terms—Point cloud, rendering, spectrum discrepancy, discriminator.

I. INTRODUCTION

Photo-realistic rendering is important for intelligent robot manipulation control, which is able to provide augmented view/angle feedback and observations [1], [2]. It has attracted increasing attention as point cloud is a well-accepted format that is widely used in robot control vision tasks. However, due to the inherent irregularity and discontinuity, view rendering from 3D scene involve complex graphic pipelines that include multiple pre-processing and post-processing steps. Traditional model-based rendering [3] aims to reconstruct surfaces and render on the mesh by employing the physical properties of lighting, texture, material, etc., which is generally computationally heavy. Image-based rendering [4] attempts to generate the novel view from images only, and recent point-based rendering [5]–[8] further simplifies the geometry constructions in rendering. Deep learning approaches [9]–[11] are widely adopted in view of its superior performance in almost all kinds of image reconstruction tasks.

¹Robotics Institute, Carnegie Mellon University, Pittsburgh, US.

²Singapore Institute of Manufacturing Technology (SIMTech), Agency for Science, Technology and Research (A*STAR), Singapore.

³Advanced Robotics Centre, College of Design and Engineering, National University of Singapore, Singapore.

⁴Birla Institute of Technology and Science, Pilani, Hyderabad, India.

† Corresponding author: zhu_haiyue@simtech.a-star.edu.sg

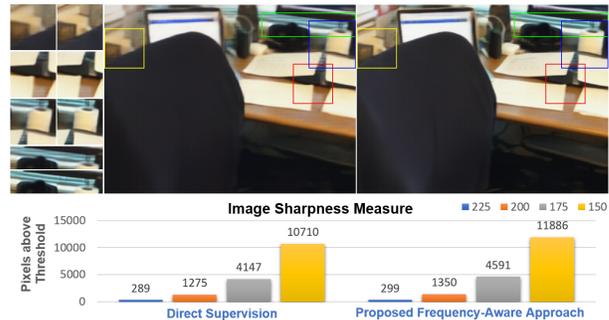


Fig. 1. Deep rendering network trained by direct supervision often lacks high-frequency details. We propose a multi-frequency-aware adversarial learning approach to minimize such frequency discrepancy so that to improve the sharpness and realness. The bar chart shows the sharpness comparison using the image sharpness measure proposed in [12]. Our method is visually also more well-defined: The glass mug, window shades, chair arm-rest appear much more defined with our method.

To achieve photo-realistic rendering, existing neural point cloud rendering approaches mostly emphasize on the reconstruction performance in RGB (spatial) domain only while overlooking image fidelity in the spectral domain. It is also well reported that convolutional neural network based image reconstruction approaches often fail to generalize high-frequency artifacts [13]–[15], causing the synthesized image to suffer from a lack of high-frequency sharpness. This spectrum discrepancy issue is a performance-limiting factor that is common to image reconstruction and generation tasks as well as neural point cloud rendering tasks. Existing techniques from image generation and super-resolution make use of a Fourier regularization term [16] or loss [17] on top of the Generative Adversarial Networks (GAN) to overcome the frequency discrepancy limitation for image generation and super-resolution tasks. However, as it is well-known Fourier transform lacks of spatially-localized frequency information to capture the abrupt signal changes, existing methods still inhibit the full utilization to well capitalize spectral domain discrepancy and recover high-frequency sharpness information.

Moreover, different with most image-to-image translation vision tasks, point cloud rendering requires a voxelisation projection step to extract the view-related 3D points according to the given rendering viewpoint. [7] directly projects 3D points onto the 2D plane, which is sensitive to the spatial occlusion and noise. Multi-plane projection [5] is proposed as a remedy by considering the spatial distance of 3D points as a

measure in the feature aggregation. However, with regards to the sensor characteristics of an RGB-D device being used to produce the initial colored point cloud, it is well-known that depth channel is generally associated with more significant noise especially for those commonly-used low-cost RGB-D cameras. Therefore, we argue that the appearance (RGB) feature of each 3D point is more robust than its spatial location, which could be utilized in the aggregation for robust view rendering.

In this work, we present a neural point cloud rendering pipeline that aims to bridge the above-mentioned gaps. A novel multi-frequency-aware patch adversarial learning scheme is proposed to supervise the point cloud rendering generation, which is specifically designed to promote the high-frequency generalization capability to minimize the local frequency discrepancy. Different from existing Fourier transform based frequency-aware adversarial learning, we explore the Discrete-Wavelet Transform (DWT) based learning as DWT is known to capture the abrupt changes of a signal more effectively, which provides additional spatially-localized high-frequency information for adversarial discrimination. Therefore, we propose a multi-discriminator strategy embedding with both Fourier domain and DWT domain discriminators as well as the spatial domain discriminator to force the generator to capture global and local spectral distributions of the real scene images. This not only helps to improve realness of localized artifacts, but also enhance the convergence speed and stability of adversarial learning. Moreover, a noise-resistant voxelisation approach is proposed by utilizing both the appearance feature distance and spatial distance for robust voxel feature aggregation. By incorporating the more stable appearance distance into feature voting, the spatial outlier points caused by depth noise can be excluded more efficiently. Our entire architecture is differentiable and can be learned in an end-to-end fashion.

To summarize, the contributions of this paper are listed as follows:

- We propose a multi-frequency-aware adversarial learning scheme, which utilizes the under-explored DWT domain to minimize the localized spectrum discrepancy between generated images and real images.
- We introduce a combined feature and spatial distance based noise-resistant voxelisation approach for robust neural point cloud rendering.
- Our approach effectively enhances the realness and sharpness of the generated images, and achieves state-of-the-art results by a significant margin.

II. RELATED WORKS

A. Novel View Synthesis

Recently, several works have been proposed for view synthesis of objects [18] and human faces [19]. A number of image-based novel view approaches have been proposed for scene rendering. For instance, Synsin [20] generates a latent point cloud representation from a single image and produces a target image. While this is highly effective for small deviations of pose from the image, large deviations produce undesirable

artifacts at the corners. Other recent approaches regarding image synthesis like Stable View Synthesis [21] and Free View Synthesis [8] use a series of images to create a geometric scaffold or a mesh followed by feature aggregation of a ray at sampled positions on the mesh from different viewpoints and rendering. Another work, NeRF [22], makes use of a hybrid deep learning and classical volume rendering approach, and several works inspired from NeRF such as [23], [24], [25], [26] and [27] have been proposed for improved efficiency and limited image data. Our work differs from these since we use raw point clouds as input instead of images.

B. Neural Point Rendering

Deep novel view synthesis [6] extracts features from a point encoder layer to encode an input point set and passes it on to an image decoder followed by a refinement network. While effective for sparse point clouds, the PointNet++ backbone cannot capture local relationships effectively and requires large paired data for refinement. Neural Point Graphics (NPG) [7] proposes augmenting learnable neural descriptors for point clouds and rasterizes it to a 2D image for multi-scale rendering network. Using multi-plane projections significantly improves performance and recent work such as [28] utilize the advantage of a layered volume. [5] extends NPG by also using multi-plane projections, which is the closest work to ours, and it is used as a baseline for comparison in our experiments.

C. Spectral Domain Loss

With the variety of image generation/translation methods in the literature, image blur along with lack of sharpness or the presence of high-frequency artifacts is an observed issue which limits the visual realness. [27] have shown that neural networks tend to learn lower frequencies faster, and overcome this by using Fourier mapping layers prior to pass inputs to a multi-layer perceptron (MLP) network for higher perceptual quality. Recent works [15] show that the realness degradation is partly attributable to missing high-frequency features due to the spectrum discrepancy. Furthermore, [16] and [29] recognize that GAN based models, especially those with upsampling layers, fail to generalize spectral data. To alleviate this problem, works such as [16] have proposed a Fourier transform as a regularization term. Arguing regularization may lead to sub-optimal performance, [30] and [17] show that using an additional discriminator based Fourier loss leads to better performance. [15] uses the direct-cosine transform as a better perceptual loss. [31] employs a Discrete Wavelet Transform (DWT) module to keep frequency information as a down-sampling module in the generator. Although it is also a DWT-based GAN approach, our motivation and approach are totally different from [31], [32] as we favour a patch multi-discriminator strategy. We aim to reduce the frequency discrepancy in both spectral (DWT and Fourier) and spatial (RGB) domains, so that the generator can generalize better on the source distribution to enhance the visual realness.

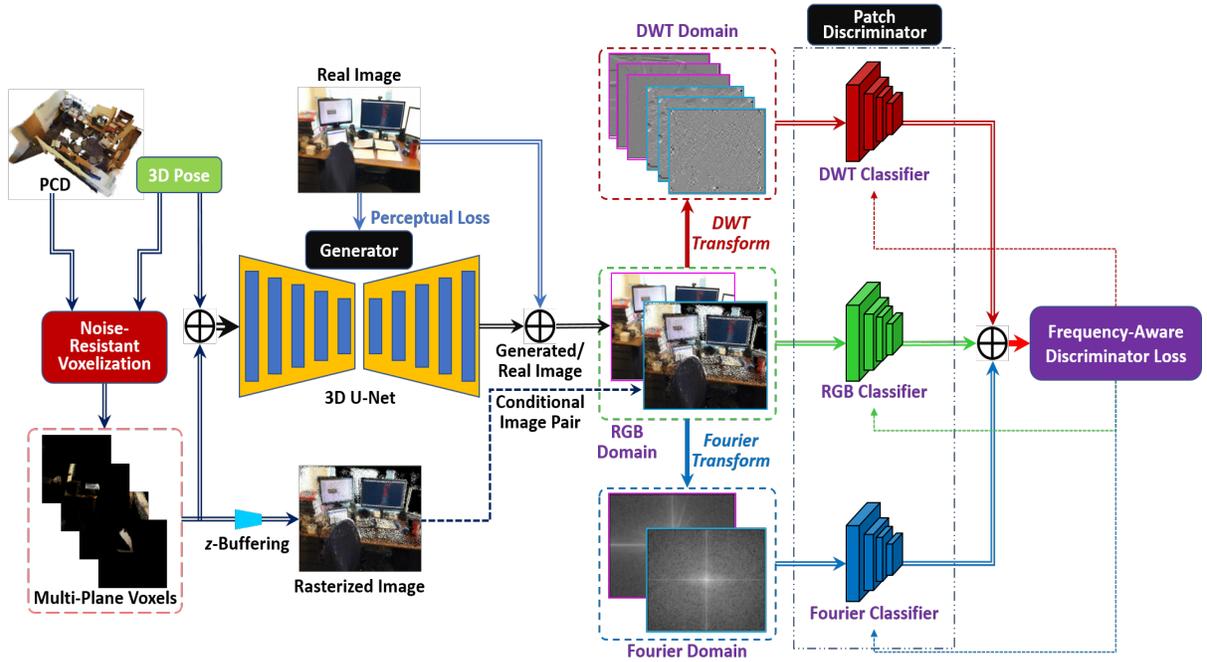


Fig. 2. Method overview: we use point cloud and 3D pose to first voxelise the visible point frustum into multi-plane projections using the proposed Noise-Resistant Voxelisation approach. This voxelisation input is passed into the 3D U-Net generator to synthesize the view rendering image. This generator is supervised through a multi-frequency-aware patch adversarial learning scheme as well as a standard perceptual loss, where a discriminator network is trained simultaneously to differentiate between real and generated images. Uniquely, to better capture high-frequency discrepancy, we adopt a tri-discriminator network that combines both spectral domain (Fourier Transform and Discrete Wavelet Transform) discriminators as well as the spatial (RGB) domain information to minimize the frequency discrepancy.

III. METHOD

A. Overview

The rendering problem studied in this work is formulated as follows. Given a scene point cloud $P = \{p_1, p_2, \dots, p_N\}$ with a set of M camera images taking from the same scene and their corresponding camera poses, denoted as $C_M = \{I, c\}_M$, the goal of neural point cloud rendering is to learn a mapping function $f(\cdot)$ that can render a virtual image \hat{I} from the point cloud P with a random target camera pose $c \in \mathbb{R}^{3 \times 4}$, i.e.,

$$\hat{I} = f(P, c), \quad (1)$$

where the generated image \hat{I} is aimed to be perceptually realistic as much as possible.

Fig. 2 shows our neural rendering pipeline in this work, where we propose a frequency-aware patch adversarial learning to achieve the photo-realistic rendering. Our approach learns a rendering generator $G(\cdot)$ through the adversarial learning, which takes the voxelised 3D volume with the target pose to synthesize the target image. A noise-resistant voxelisation method is developed to produce consistent 3D volumes even from noisy and irregularly sampled point clouds, and a multi-frequency-aware patch discriminator scheme is proposed in the adversarial learning to effectively capture the high frequency visual information both locally and globally through the Fourier domain and DWT domain.

B. Noise-Resistant Voxelisation

Following [5], we project the visible point cloud region into multiple planes so as to avoid artifacts from occluded regions. This projection helps to correct noise interference in comparison with just using a one-plane rasterized image. For a voxel centered at (p, h, w) that contains N_v number of sub-voxel points with features as $f_{p,h,w}^i$, the voxel feature $F_{(p,h,w)}$ is calculated by aggregation of all sub-voxel points. [5] calculates a weighted average of sub-voxel points by considering the spatial distance between the vertical and horizontal planes. However, such purely spatial distance based voxelisation is sensitive to the noisy points closer to the centre, which could hamper the rendering performance.

In contrast, the appearance feature of a point is generally more robust than the spatial location considering the sensor physical properties of RGB-D devices. Therefore, we propose a noise-resistant voxelisation that incorporates the feature distance as well to aggregate the voxel feature,

$$F_{(p,h,w)} = \frac{\sum_i w_i^i f_{(p,h,w)}^i}{\sum_i w_i^i}, \quad (2)$$

where the blending weight w_i considers both the spatial and feature distances as

$$\begin{aligned} w_{(p,h,w)}^i &= \mu_f D_{f(p,h,w)}^i + \mu_s D_{s(p,h,w)}^i \\ D_{s(p,h,w)}^i &= (1 - D_{1(p,h,w)}^i)^\alpha (1 + D_{2(p,h,w)}^i)^\beta \\ D_{f(p,h,w)}^i &= (|f_{p,h,w}^i - \bar{f}_{p,h,w}|_1)^{-1} \end{aligned} \quad (3)$$

where D_s^i and D_f^i are the spatial and feature inverse distances, respectively, μ_f and μ_s are the weights controlling their effects. $\bar{f}_{p,h,w}$ is the average of point color features for points in voxel (p, h, w) , $\bar{f}_{p,h,w} = \sum_{i=0}^N f_{p,h,w}^i / N$. α and β are hyper-parameters to control the blending weight; D_1 and D_2 are the distances of the point from center of voxel and minimum depth point for a particular voxel, respectively. Overall, the voxelisation can be represented as

$$F = \mathcal{V}(P, c), \quad (4)$$

where \mathcal{V} denotes the proposed noise-resistant voxelisation operation.

C. Patch Adversarial Learning

For the rendering network, we employ the adversarial scheme to learn the generator mapping from the voxelisation input F to the synthesized image rendering \hat{I} , i.e., $\hat{I} = G(F)$, which is a simple 3D U-Net. The adversarial counterpart is a conditional patch discriminator array $D(\cdot)$, which takes the input patch pairs to distinguish whether it is real camera image or synthesized image. For the conditional pairs, we perform a z -buffer rasterization of the multi-plane voxelisation and concatenate it with either the synthesized or real image as the input of $D(\cdot)$. Inspired by PatchGAN [10], the discriminator adopts the patch approach to output $N_p \times N_p$ scores, which penalizes the discrimination for each receptive region to encourage the attention on high-frequency local information. The patch loss is especially effective for our DWT frequency discriminator since the spectral features are local and can be better exploited to distinguish fake/real images to improve the generator performance.

D. Frequency-Aware Multiple Discriminator

One major limitation of traditional view rendering, many of which optimize either an L_2 or L_1 loss, is that it often only captures the low frequency visual details, leading to the sharpness degradation on high frequency abrupt changes. Our adversarial scheme is designed to be frequency-aware to effectively capture and generalize high-frequency visual information. Our network uses a multi-discrimination strategy that combines the RGB domain with frequency domain to improve image synthesis quality indirectly. The strategy is achieved by a Fourier discriminator that encourages the global generalization of high frequencies, as well as a DWT discriminator, which locally differentiates the frequency discrepancy in the image. As the DWT and Fourier discriminators spotlight contrasting aspects of discrepancy on frequency domain, the two work in tandem to drive the generator to produce a more photo-realistic rendering result.

Our multi-frequency-aware discriminator scheme is formed as

$$D(\cdot) = D_{RGB}(\cdot) + D_{Fourier}(\cdot) + D_{DWT}(\cdot), \quad (5)$$

where $D_{RGB}(\cdot)$ is a patch discriminator on raw RGB domain, which takes the input pair $\{I_r, I \setminus \hat{I}\}$ formed by the rasterized image I_r with either the real image I or

generated image \hat{I} . $D_{Fourier}(\cdot)$ is a frequency discriminator on Fourier domain, which takes the Fourier transformed input pair $\{\mathcal{F}(I_r), \mathcal{F}(I) \setminus \mathcal{F}(\hat{I})\}$, where $\mathcal{F}(\cdot)$ denotes the Fourier transform operation. It is well-known that Fourier transform cannot well capture the abrupt changes of a signal which is a limitation for measuring high-frequency artifacts. To solve this issue, we propose using a DWT discriminator D_{DWT} in addition to $D_{Fourier}$ in (5). The DWT is aiming for those perceptually-important localized frequency information, and we take three sub-bands, i.e., HL (vertical), LH (horizontal), and HH (diagonal) features from the DWT, as input to D_{DWT} , denoted as $\{\mathcal{HL}(I_r), \mathcal{LH}(I_r), \mathcal{HH}(I_r), \mathcal{HL}(I \setminus \hat{I}), \mathcal{LH}(I \setminus \hat{I}), \mathcal{HH}(I \setminus \hat{I})\}$, where \mathcal{HL} , \mathcal{LH} , and \mathcal{HH} denotes the respective components.

TABLE I
PERFORMANCE COMPARISON ON SCANNET AND MATTERPORT 3D

Methods	ScanNet			Matterport 3D		
	SSIM	PSNR	LPIPS	SSIM	PSNR	LPIPS
Pix2Pix [10]	0.731	19.247	0.429	0.530	14.964	0.675
NPG [7]	0.84	22.911	0.245	0.622	22.911	0.597
MPP [5]	0.835	22.813	0.234	0.649	18.09	0.534
Ours	0.871	32.7	0.1012	0.672	25.95	0.24

E. Loss Functions

In the proposed frequency-aware patch adversarial learning, each discriminator is supervised using least square patch loss,

$$\begin{aligned} \mathcal{L}_{D_\pi} = & \mathbb{E}_I \sum_{i=0}^{N_p-1} \sum_{j=0}^{N_p-1} (D_\pi(I)_{i,j} - a)^2 \\ & + \mathbb{E}_F \sum_{i=0}^{N_p-1} \sum_{j=0}^{N_p-1} (D_\pi(G(F))_{i,j} - b)^2 \end{aligned} \quad (6)$$

where $\pi \in \Pi$ and $\Pi = [RGB, Fourier, DWT]$. a and b are the labels for fake data and real data, and $D(\cdot)_{i,j}$ denotes the i, j -th patch prediction. As a result, the overall objective for the discriminator is

$$\min_D \mathcal{L}(D) = \min_{D_\pi} \sum_{\pi \in \Pi} \mathcal{L}(D_\pi). \quad (7)$$

For the generator, we use a perceptual loss [33] for full supervision in addition to the discriminators,

$$\mathcal{L}_{percept}(G) = \|G(F) - I\|_1 + \sum_l \lambda_l \|\phi_l(G(F)) - \phi_l(I)\|_1 \quad (8)$$

where $\phi_l(\cdot)$ is l -th layer output of pre-trained VGG-19 and λ_l is a controlling weight. Therefore, the overall objective for the generator is

$$\begin{aligned} \min_G \mathcal{L}(G) = & \min_G \left\{ \mathcal{L}_{percept}(G) \right. \\ & \left. + \mathbb{E}_F \sum_{\pi \in \Pi} \sum_{i=0}^{N_p-1} \sum_{j=0}^{N_p-1} (D_\pi(G(F))_{i,j} - c)^2 \right\}, \end{aligned} \quad (9)$$

where c denotes the value that G wants D to believe for fake data following [34].

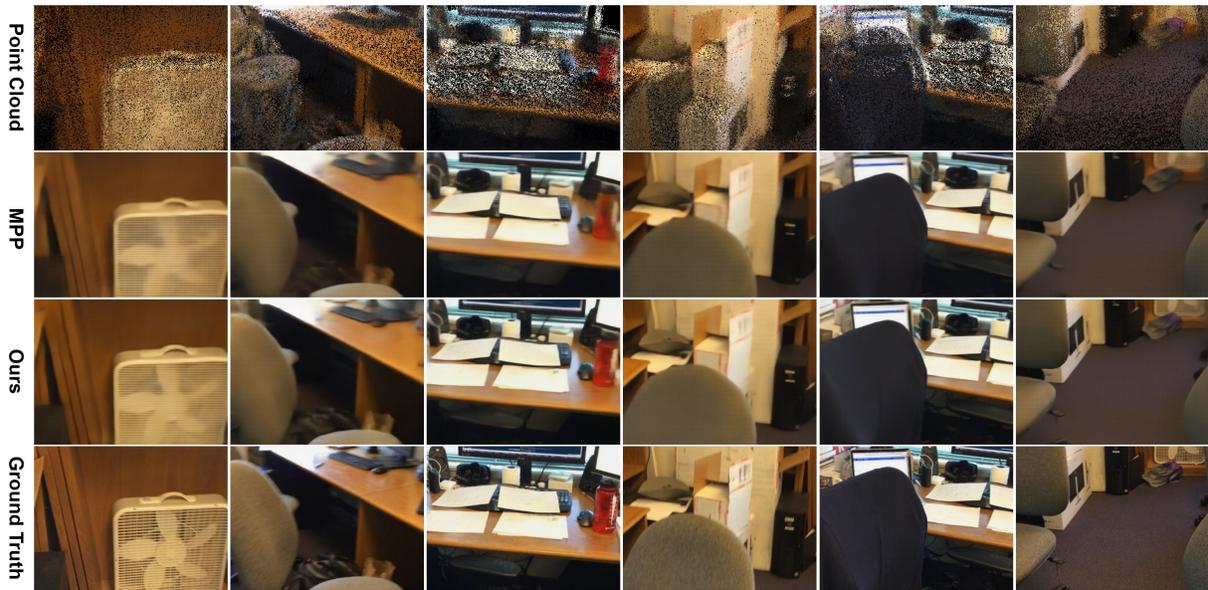


Fig. 3. Visualization of rendering results on a ScanNet scene. Objects such as the fan, chair, and water bottle have more clearly defined edges in our rendered image. Blurring artifacts such as the red colored tape on the cardboard box (fourth image from right) are also well recovered by our model.

IV. EXPERIMENTS

A. Datasets

We perform extensive experiments on two public datasets ScanNet [35] and Matterport 3D [36] to evaluate our performance. The ScanNet dataset contains RGB-D scans in over 1500 environments, and we use the RGB-D scans and the 3D camera pose annotations for registering the point cloud. The Matterport 3D dataset is much more challenging with larger spread-out scenes, along with very large variation of poses. We pre-process both datasets following the same setting in [5].

B. Rendering Evaluation

To establish the benchmarking, we consider three methods that capture ideas of image translation and point rendering techniques. Our first baseline is the well-known standard Pix2Pix network [10] (denoted as Pix2Pix), which use a classical z -buffer rasterized image as input, and the translation task is to convert it to a realistic view. While this potentially interprets the distribution of noise and sparsity of point-clouds, it is unable to capture any information about pose and/or scene depth. Our next two baselines are Neural Point-based Graphic [7] (denoted as NPG) and Multi-Plane Projection Rendering [5] (denoted as MPP), which capture pose information and depth information by using learnable point descriptors with multi-plane information.

To evaluate the performance of different approaches, we compare the quantitative evaluations for our method with all three baselines mentioned above. We adopt three performance metrics same as [5], namely Structural Similarity Index (SSIM), Peak-Signal to Noise Ratio (PSNR), and Learned Perceptual Image Patch Similarity (LPIPS). The evaluation results are provided in Table I. It can be observed that for ScanNet, Pix2Pix leads to sub-optimal results with SSIM of

only 0.73, and both of NPG and MPP lead to SSIM around 0.83~0.84. Our proposed method significantly increases the SSIM to 0.871, and similarly attains the best PSNR and LPIPS scores. We attribute this significant improvement to the combination of spectral and spatial domain losses, which leads to better perceptual quality whilst improving the Signal to Noise Ratio. As a result, our approach generates sharper and more realistic images. Fig. 3 shows some comparison of visualization results on ScanNet dataset. Our rendering is able to define sharper object boundaries and inpaint colored features in areas missed out by the traditional rendering approach, making our synthesized image closer to the ground truth. Similar outperformance trend can be observed for Matterport 3D dataset as shown in Table I.

TABLE II
ABLATION STUDY WITH DIFFERENT DISCRIMINATOR MODULE

Discriminator Type	SSIM	PSNR	LPIPS
No Discriminator	0.835	22.813	0.234
Spatial only	0.854	32.37	0.1147
Fourier only	0.859	31.15	0.1042
Spatial+Fourier	0.864	31.56	0.0985
Spatial+DWT	0.866	31.91	0.1010
Spatial+Fourier+DWT	0.871	32.7	0.1012

C. Ablation Studies - Discriminator Modules

Since our multi-discriminator adversarial learning employs spatial, Fourier, and DWT discriminators, we conduct ablation studies on different combinations of the multi-discriminator strategy using ScanNet dataset. Table II summarizes the results and demonstrates the importance of incorporating all three domains. We observe a drop in performance in the absence of one of the spectral discriminators, and achieve the most

significant improvement from using a combination of the three discriminators.

V. CONCLUSION

In this paper, we present a multi-frequency-aware adversarial learning scheme to achieve neural point cloud rendering, which is realized by a tri-discriminator scheme from RGB, Fourier, and DWT domains. While Fourier transform has been shown effective to regularize frequency-aware learning, it lacks sufficient generalization capabilities regarding localized high-frequency abrupt features. We combine Fourier and DWT domains with the spatial domain to achieve high fidelity and photorealistic rendering for novel view synthesis. In addition, we also introduce a noise-resistant voxelisation to reduce the impact of spatial outliers. Our model outperforms existing baselines and achieve the state-of-the-art performance on the ScanNet and Matterport 3D datasets with a significant margin.

ACKNOWLEDGMENT

This research is supported by National Research Foundation (NRF) “Centre for Advanced Robotics Technology Innovation (CARTIN)”, and also supported by A*STAR under its “RIE2025 IAF-PP Advanced ROS2-native Platform Technologies for Cross sectorial Robotics Adoption (M21K1a0104)” programme. Moreover, the support from Singapore Institute of Manufacturing Technology (SIMTech) and National University of Singapore (NUS) through the SIMTech-NUS Joint Laboratory on Industrial Robotics are gratefully acknowledged.

REFERENCES

- [1] X. Li, B. He, Y. Zhou, and G. Li, “Multisource model-driven digital twin system of robotic assembly,” *IEEE Systems Journal*, vol. 15, no. 1, pp. 114–123, 2021.
- [2] Y. Li, H. Zhu, S. Tian, F. Feng, J. Ma, C. S. Teo, C. Xiang, P. Vadakkepat, and T. H. Lee, “Incremental few-shot object detection for robotics,” in *ICRA 2022*, 2022, pp. 8447–8453.
- [3] A. Dai, M. Nießner, M. Zollhöfer, S. Izadi, and C. Theobalt, “Bundlefusion: Real-time globally consistent 3d reconstruction using on-the-fly surface reintegration,” *ACM Trans. Graph.*, vol. 36, no. 3, may 2017.
- [4] P. Hedman, J. Philip, T. Price, J.-M. Frahm, G. Drettakis, and G. Brostow, “Deep blending for free-viewpoint image-based rendering,” *ACM Trans. Graph.*, vol. 37, no. 6, dec 2018.
- [5] P. Dai, Y. Zhang, Z. Li, S. Liu, and B. Zeng, “Neural point cloud rendering via multi-plane projection,” in *Proceedings of the IEEE/CVF CVPR*, 2020, pp. 7830–7839.
- [6] S. Song, W. Chen, D. Campbell, and H. Li, “Deep novel view synthesis from colored 3d point clouds,” in *Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIV*, 2020, p. 1–17.
- [7] K.-A. Aliev, A. Sevastopolsky, M. Kolos, D. Ulyanov, and V. Lempitsky, “Neural point-based graphics,” in *ECCV 2020*, 2020, pp. 696–712.
- [8] G. Riegler and V. Koltun, “Free view synthesis,” in *ECCV 2020*, 2020.
- [9] I. J. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, “Generative adversarial networks,” in *NeurIPS 2014*, 2014.
- [10] P. Isola, J.-Y. Zhu, T. Zhou, and A. A. Efros, “Image-to-image translation with conditional adversarial networks,” in *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 5967–5976.
- [11] J.-Y. Zhu, T. Park, P. Isola, and A. A. Efros, “Unpaired image-to-image translation using cycle-consistent adversarial networks,” in *2017 IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 2242–2251.
- [12] K. De and V. Masilamani, “Image sharpness measure for blurred images in frequency domain,” *Procedia Engineering*, vol. 64, pp. 149–158, 2013.
- [13] M. Fritsche, S. Gu, and R. Timofte, “Frequency separation for real-world super-resolution,” in *ICCV Workshops*, 2019.
- [14] T. Dzanic, K. Shah, and F. D. Witherden, “Fourier spectrum discrepancies in deep network generated images,” in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, ser. NIPS ’20, 2020.
- [15] O. Giudice, L. Guarnera, and S. Battiato, “Fighting deepfakes by detecting gan dct anomalies,” *Journal of Imaging*, vol. 7, no. 8, p. 128, Jul 2021. [Online]. Available: <http://dx.doi.org/10.3390/jimaging7080128>
- [16] R. Durall, M. Keuper, and J. Keuper, “Watch your up-convolution: Cnn based generative deep neural networks are failing to reproduce spectral distributions,” in *CVPR*, 2020, pp. 7887–7896.
- [17] D. Fuoli, L. V. Gool, and R. Timofte, “Fourier space losses for efficient perceptual image super-resolution,” in *ICCV 2021*, 2021.
- [18] P. Z. Ramirez, A. Tonioni, and F. Tombari, “Unsupervised novel view synthesis from a single image,” 2021.
- [19] R. Huang, S. Zhang, T. Li, and R. He, “Beyond face rotation: Global and local perception gan for photorealistic and identity preserving frontal view synthesis,” 2017.
- [20] O. Wiles, G. Gkioxari, R. Szeliski, and J. Johnson, “Synsin: End-to-end view synthesis from a single image,” in *CVPR*, 2020, pp. 7467–7477.
- [21] G. Riegler and V. Koltun, “Stable view synthesis,” in *CVPR 2021*, 2021.
- [22] B. Mildenhall, P. P. Srinivasan, M. Tancik, J. T. Barron, R. Ramamoorthi, and R. Ng, “Nerf: representing scenes as neural radiance fields for view synthesis,” *Commun. ACM*, vol. 65, no. 1, p. 99–106, dec 2021.
- [23] A. Yu, V. Ye, M. Tancik, and A. Kanazawa, “pixelnerf: Neural radiance fields from one or few images,” in *CVPR 2021*, 2021.
- [24] J. T. Barron, B. Mildenhall, M. Tancik, P. Hedman, R. Martin-Brualla, and P. P. Srinivasan, “Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields,” in *ICCV 2021*, 2021.
- [25] R. Martin-Brualla, N. Radwan, M. S. M. Sajjadi, J. T. Barron, A. Dosovitskiy, and D. Duckworth, “Nerf in the wild: Neural radiance fields for unconstrained photo collections,” in *CVPR 2021*, 2021.
- [26] K. Zhang, G. Riegler, N. Snavely, and V. Koltun, “Nerf++: Analyzing and improving neural radiance fields,” *arXiv preprint arXiv:2010.07492*, 2020.
- [27] M. Tancik, P. P. Srinivasan, B. Mildenhall, S. Fridovich-Keil, N. Raghavan, U. Singhal, R. Ramamoorthi, J. T. Barron, and R. Ng, “Fourier features let networks learn high frequency functions in low dimensional domains,” in *Proceedings of the 34th International Conference on Neural Information Processing Systems*, ser. NIPS ’20. Curran Associates Inc., 2020.
- [28] J. Flynn, M. Broxton, P. Debevec, M. DuVall, G. Fyffe, R. Overbeck, N. Snavely, and R. Tucker, “Deepview: View synthesis with learned gradient descent,” in *CVPR*, June 2019.
- [29] X. Zhang, S. Karaman, and S.-F. Chang, “Detecting and simulating artifacts in gan fake images,” in *2019 IEEE International Workshop on Information Forensics and Security (WIFS)*, 2019, pp. 1–6.
- [30] Y. Chen, G. Li, C. Jin, S. Liu, and T. H. Li, “Ssd-gan: Measuring the realness in the spatial and spectral domains,” in *AAAI Conference on Artificial Intelligence*, 2020.
- [31] P. Liu, H. Zhang, K. Zhang, L. Lin, and W. Zuo, “Multi-level wavelet-cnn for image restoration,” in *2018 IEEE/CVF CVPRW*, 2018, pp. 886–88609.
- [32] H. Guo, H. Zhu, J. Wang, V. Prahlad, W. K. Ho, C. W. d. Silva, and T. H. Lee, “Remaining useful life prediction via frequency emphasizing mix-up and masked reconstruction,” *IEEE Transactions on Artificial Intelligence*, pp. 1–10, 2024.
- [33] J. Johnson, A. Alahi, and L. Fei-Fei, “Perceptual losses for real-time style transfer and super-resolution,” in *Computer Vision – ECCV 2016*, B. Leibe, J. Matas, N. Sebe, and M. Welling, Eds., Cham, 2016, pp. 694–711.
- [34] X. Mao, Q. Li, H. Xie, R. Y. Lau, Z. Wang, and S. Paul Smolley, “Least squares generative adversarial networks,” in *ICCV*, 2017, pp. 2794–2802.
- [35] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, “ScanNet: Richly-annotated 3d reconstructions of indoor scenes,” in *CVPR*, 2017.
- [36] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang, “Matterport3d: Learning from rgb-d data in indoor environments,” *arXiv preprint arXiv:1709.06158*, 2017.