

NeRF-MAE: Masked AutoEncoders for Self-Supervised 3D Representation Learning for Neural Radiance Fields

Muhammad Zubair Irshad^{1,2}, Sergey Zakharov¹, Vitor Guizilini¹, Adrien Gaidon¹, Zsolt Kira¹, Rares Ambrus²

¹Toyota Research Institute ²Georgia Institute of Technology

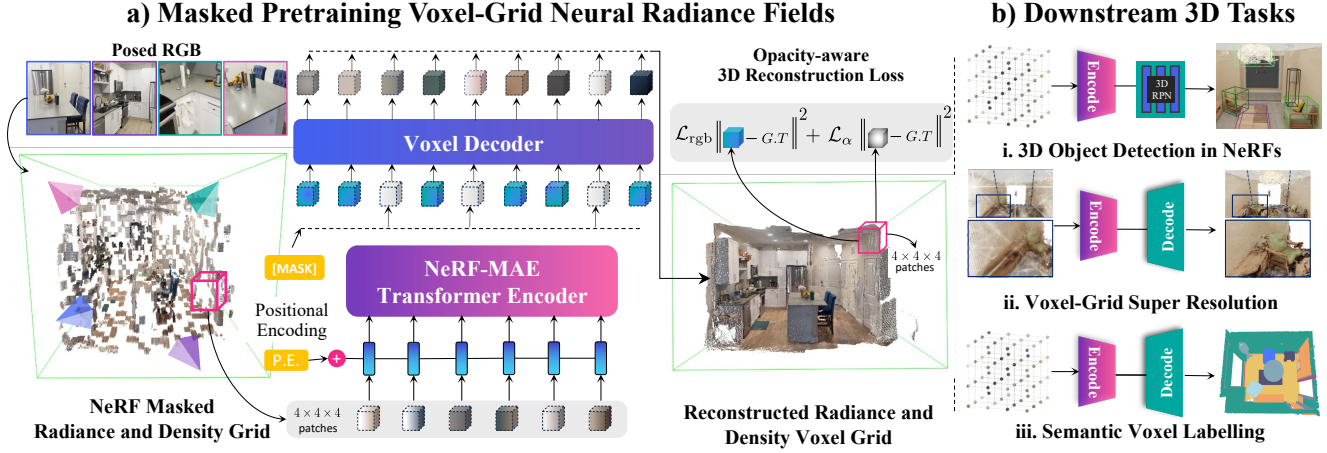


Figure 1. **Overview:** (a) We present **NeRF-MAE**, the first large-scale fully self-supervised pretraining for Neural Radiance Fields (NeRFs). Our approach utilizes a 3D Swin Transformer-based encoder and a voxel decoder to learn a powerful representation in (a) masked self-supervised learning objective directly in 3D using only posed RGB images as input. (b) Our representation, when finetuned on a small subset of data, improves many 3D downstream tasks such as 3D object detection in NeRFs, voxel-grid super-resolution, and voxel-labelling.

Abstract

Given the capabilities of neural fields in densely representing 3D scenes from 2D images, we ask the question: Can we scale their self-supervised pretraining, specifically using masked autoencoders, to generate effective 3D representations from posed RGB images. Owing to the astounding success of extending transformers to novel data modalities, we employ standard 3D Vision Transformers to suit the unique formulation of NeRFs. We leverage NeRF’s volumetric grid as a dense input to the transformer, contrasting it with pointclouds where the information density can be uneven, and the representation is irregular. Our goal is made possible by masking random patches from NeRF’s radiance and density grid and employing 3D Swin Transformer to reconstruct masked patches. In doing so, the model learns the semantic and spatial structure of complete scenes. Our novel self-supervised pretraining for NeRFs scales remarkably well and improves performance on various challenging 3D tasks; thus significantly outperforming self-supervised pretraining and scene understanding baselines on Front3D and ScanNet datasets with a performance improvement of over 20% AP50 and 8% AP25 for 3D object detection. Project Page: nerf-mae.github.io

1. Introduction

Neural fields’ [1, 36] emergent properties [70] have made them suitable for tasks beyond showcasing higher rendering quality. Prior approaches have successfully applied neural fields to extract accurate scene geometry [40, 63, 76], estimate camera poses [66, 78], infer accurate semantics [2, 77], and learn 3D correspondences [57, 79]. NeRFs have also emerged as a promising 3D data storage medium [24], and have proven to outperform traditional supervised approaches [19, 24] for challenging 3D tasks like object detection [19, 71] and instance segmentation [32].

Because of NeRF’s inherent ability to densely represent a 3D scene from 2D images, this paper explores scaling up their self-supervised pre-training (Fig. 3) through the use of masked autoencoders. The notion of masked auto-encoding is equally suited to NeRF’s radiance and density grid as to images (Fig. 2). As opposed to other 3D representations such as meshes, point clouds, or lidar data which only model surface-level information, are highly irregular data structures and the information density can be extremely uneven [42], NeRF’s radiance and density grid is similar in principle to 2D images for masked auto-encoding. This is




Figure 2. **NeRF-MAE comparison:** We show detailed comparison to point-based pretraining approaches in § C.5. We note their limitation as they model surface-level sparse points, where our approach is similar to images in terms of high-information density and regularity of structure; hence making NeRF-MAE a direct extension of image MAE to 3D. We leverage NeRF’s dense volumetric information along rays and introduce an opacity-aware reconstruction loss, enabling us to achieve superior representation learning.

the case because it provides high information density, and a regular grid ensuring unbiased sampling. Owing to these nice properties, the radiance and density grid obtained from NeRFs have shown to achieve state-of-the-art scene understanding such as segmentation and detection [19, 32].

Driven by this analysis, we propose **NeRF-MAE** (NeRF Masked Auto-Encoders) for self-supervised learning directly within Neural Radiance Fields. In comparison to prior generalizable NeRF or 3D representation learning methods (Table S2), NeRF-MAE offers the following advantages: **1)** The masked self-supervised learning objective is directly enforced on dense regular grids with high-density information, thus enabling the use of standard Transformer architecture without requiring careful adaptation or adopting extra non-Transformers Blocks (such as DGCNN used in Point-BERT [81]). **2)** By disentangling representation learning and NeRF training, our model can effectively utilize a large amount of data from diverse sources. This is in contrast to most previous generalizable NeRF methods that only train on a single type of data source for priors.

For training our model, we curate a large-scale dataset for NeRF pre-training, encompassing 4 diverse sources in synthetic and real domains. Our dataset comprises multi-view posed images, corresponding NeRFs, and sampled

radiance and density grids from Front3D, ScanNet, HM3D, and Hypersim, accumulating over 1.6M images and 3500+ scenes (Table. S2 and Figure S3). To summarize, we make the following contributions:

- To our knowledge, we introduce the first **fully self-supervised transformer-based 3D pretraining** utilizing Neural Radiance Field’s radiance and density grid as an input modality coupled with an **opacity-aware masked reconstruction objective**.
- A **large scale pre-training of Neural Radiance Fields** on 4 different datasets, over 1.6M images, and 3500+ indoor scenes, using a **single model**.
- Our proposed approach **significantly outperforms self-supervised 3D pretraining as well as NeRF scene understanding baselines** on multiple downstream 3D tasks, showing over 20% AP50 and 12% mAcc on 3D object detection and semantic voxel labeling respectively on Front3D dataset while **requiring less than half the data** requirement of the state-of-the-art baseline to achieve the same performance.

2. NeRF-MAE: Self-Supervised Pretraining for Neural Radiance Fields

NeRF-MAE introduces a self-supervised framework aimed at enhancing 3D representation learning in the context of Neural Radiance Fields [36] (NeRFs). Our approach allows learning robust 3D representations from posed RGB images, leading to substantial improvements in various downstream 3D applications, as we show in Sec. 3. As shown in Figures 1 and S4, such a goal is made possible by two major components: **(i)** An explicit 4D radiance and density grid extraction module in the canonical world frame using a camera-trajectory aware sampling from a fully-trained implicit NeRF model (described in Sec. B.2) and **(ii)** A masked self-supervised pretraining module, operating directly on the explicit NeRF’s 4D radiance and density grid, to train a




Figure 3. **Quantitative results** showing scaling laws and NeRF quality ablation for NeRF-MAE pretraining and transfer learning.

standard 3D SwinTransformer [33, 75] encoder and a voxel decoder using an opacity-aware masked reconstruction objective in 3D (discussed in Sec. 2.1). Our NeRF pretraining dataset is described in Sec. B.3 and key preliminaries to understand our pipeline are detailed in Sec. B.1.

2.1. Masked Pretraining NeRFs

NeRF-MAE comprises a standard Swin Transformer [33] encoder and transposed convolution decoders connected with skip connections from the encoded output at multiple feature resolutions of the transformer. Our architecture is summarized in Figure S2. The U-Net architecture aims to reconstruct masked 3D radiance and opacity patches. In doing so, the network learns to understand the semantic and spatial structure of 3D scenes from dense 3D radiance and density grid (obtained in Section B.2). Our U-Net style architecture for NeRF-pretraining is described below.

3D Transformer Encoder: The goal of the encoder is to encode the input 3D radiance and density grid (\mathcal{G}) into meaningful features at multiple resolutions which can be decoded and used for downstream tasks (Section C.3). We employ the 3D version [19, 75] of the standard Swin Transformer [33] which is obtained by replacing all the 2D operations of the SwinTransformer with their 3D equivalent, as noted by [19, 75]; for example 3D convolutions and 3D patch merging. We build our 3D auto-encoder architecture based on standard Transformers since NeRF’s radiance and density grid is similar to images in terms of information density and which can serve as a unified architecture for various 3D scene understanding tasks operating directly on NeRFs. Further details are presented in Sec. B.4.

Reconstruction with Lightweight Decoders: We utilize the feature pyramid $\mathcal{F} = [f_i \dots f_n]$ obtained after concatenating the resulting output of each SwinTransformer 3D stage to obtain the reconstructed radiance and density grid (\mathcal{R}), where $i = 0$ and $n = 3$. The reconstructed grid is obtained by attaching lightweight transposed 3D convolution decoders (\mathcal{D}_i) to the feature outputs with a kernel size of 3 and adding a residual connection from the previous decoder, where $d_i = \text{Conv}(d_{i-1} + \mathcal{D}(f_i))$ and $\mathcal{D}_0 = f_3$. The output of the final decoder block is fed into the residual block with $3 \times 3 \times 3$ convolutional and a sigmoid activation function to predict a 4-channel reconstructed volume $\mathcal{R} \in \mathbb{R}^{H \times W \times D \times 4}$. We use the same feature pyramid \mathcal{F} to output 3D object bounding boxes, voxel semantic labels, and super-resolution voxels with task-specific decoders, as described in Section C.3.

Masked Radiance and Density Grid Pretraining Objective: The goal of masked voxel-grid pretraining as applied to Neural Radiance Field (NeRF) is to encode semantic and spatial regions of interest into the network. We use mask volume region reconstruction, similar to [8] for language and [15] for images. We enforce a faithful and accurate

reconstruction of masked patches with a custom loss function suited to NeRF’s unique formulation. Specifically, we employ a combination of opacity and photometric radiance reconstruction loss, both enforced at the volumetric level, where loss $L_{recon} = \mathcal{L}_{rad} + \mathcal{L}_{\alpha}$ is defined as:

$$\mathcal{L}_{rad} = \frac{1}{K} \sum_{i=1}^N (\hat{y}_{rad_i} - y_{rad_i})^2 \quad \mathcal{L}_{\alpha} = \frac{1}{M} \sum_{j=1}^M (\hat{y}_{\alpha_j} - y_{\alpha_j})^2$$

where K is the number of voxels in the mask corresponding to patches where $\alpha > x_i - x_{i+1}$. We use a small preset distance, $x_i - x_{i+1} = 0.01$. \hat{y}_{rad_i} is the predicted radiance at voxel i and y_{rad_i} is the target radiance at voxel i . M denotes the number of voxels in the mask corresponding to removed 3D patches, \hat{y}_{α_j} is the predicted opacity at voxel, j and y_{α_j} is the target opacity at voxel j .

3. Experiments and Results

In this section, we aim to answer the following questions. 1) How well does our method *compare with 3D pretraining baselines*? 2) Does NeRF-MAE *improve downstream 3D scene understanding tasks*? and 3) How well does *NeRF-MAE scale*? We also provide an *ablation analysis* of which components most impact the performance.

Baselines: We compare 7 model variants across 3 different tasks to show the effectiveness of our method. 1) **NeRF-RPN** [19]: A state-of-the-art approach that predicts 3D boxes directly in NeRFs using 3D FPN [29]. 2) **Point-MAE** [42]: Extends 2D MAE [15] to 3D pointclouds and improves object classification and segmentation. 3) **Point-M2AE** [85]: Extends Point-MAE to add multiscale features for irregular point cloud representation learning. 4) **Depth-Contrast** [86]: Uses depth to learn 3D representations in a contrastive manner by comparing transformations of a 3D point cloud. 5) **Contrastive Scene Context** [18]: Leverages point correspondences along with spatial contexts to learn 3D scene level representation. 6) **ImVoxelNet** [54]: A strong multi-view 3D detection baseline. 7) **NeRF-MAE**: Our 3D representation learning approach utilizing posed 2D images as input and operates directly on scene’s volumetric radiance and density grids obtained through NeRFs.

Metrics: We report **3D PSNR** and **3D MSE** to evaluate grid reconstruction quality. We report Average Precision (**AP**) along with **Recall** at thresholds 25 and 50 to measure 3D OBB prediction success. We report mean intersection-over-union (**mIOU**), mean accuracy (**mAcc**), and total accuracy (**Acc**) for the semantic voxel labeling task.

3.1. Comparisons with strong baselines for improved downstream performance

3D Object Detection Downstream Task: The results of our proposed method are summarized in Tables 1 and S4. We consistently outperform the state-of-the-art baseline methods on 3D object detection in NeRFs on both Front3D and ScanNet scenes. Among our variants, our method

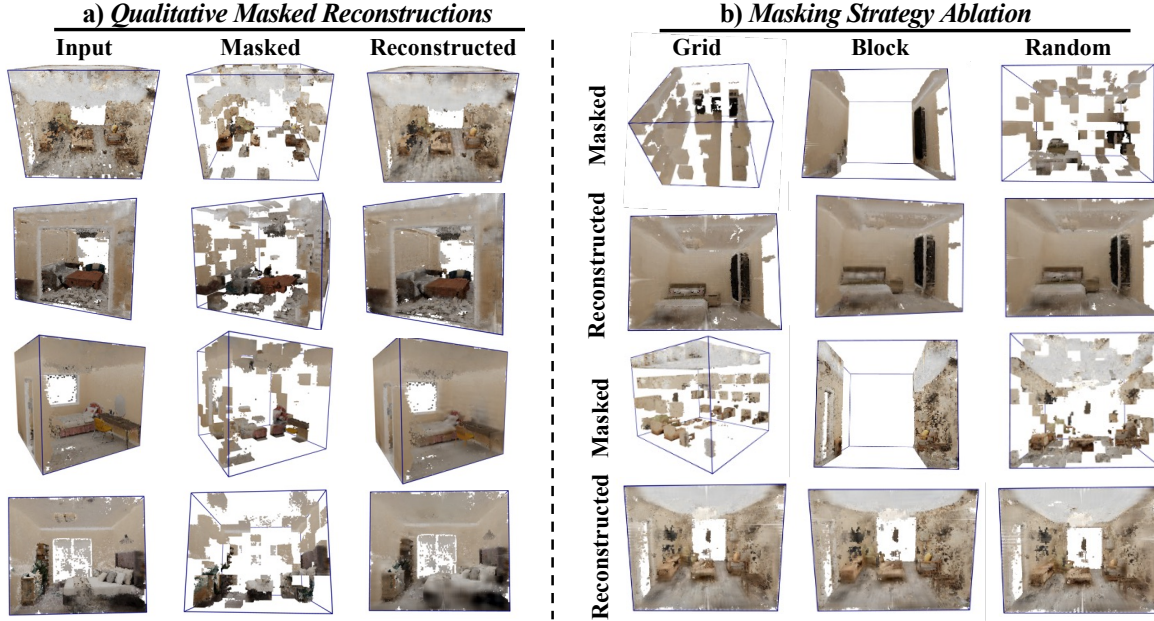


Figure 4. **Qualitative NeRF-MAE Reconstructions:**(§ D) **Left:** For each triplet, we show ground truth (left), masked radiance and density grid (middle) and our NeRF-MAE reconstruction overlayed with unmasked GT grid (right). The masking ratio is 75%, leaving only 250 patches out of 1000 patches. **Right:** shows different masking strategies along with the reconstructed output.

Approach	Aug.	Front3D [12] 3D OBB				ScanNet [7] 3D OBB			
		Recall@25↑	Recall@50↑	AP@25↑	AP@50↑	Recall@25↑	Recall@50↑	AP@25↑	AP@50↑
NeRF-RPN [19]	-	0.961	0.622	0.780	0.415	0.891	0.323	0.491	0.140
ImVoxelNet [54]	-	0.883	0.715	0.861	0.664	0.517	0.202	0.373	0.098
NeRF-MAE (F3D)	✗	0.962	0.675	0.780	0.543	0.897	0.361	0.510	0.145
NeRF-MAE (F3D)	✓	<u>0.963</u>	<u>0.743</u>	0.830	0.591	<u>0.905</u>	<u>0.391</u>	<u>0.543</u>	<u>0.155</u>
NeRF-MAE (Ours)	✓	0.972	0.745	<u>0.853</u>	<u>0.630</u>	0.920	0.395	0.571	0.170

Table 1. **Quantitative comparisons with strong baselines on downstream 3D OBB prediction:** (§ C.3) showing our pretraining improves downstream performance on never-seen during pretraining cross-dataset (ScanNet) and in-domain hold-out data (Front3D). (**Ours**) indicate we use F3D [12], HM3D [48] and HS [50] for pretraining. Aug. indicates augmentations done *only* during pretraining.

pre-trained on multiple data sources (Table 1, row 5) and with augmentations enabled during pretraining performs the best. Specifically, NeRF-MAE shows superior performance on the unseen test-set by achieving an AP50 of 63% and a Recall50 of 74.3%, hence demonstrating an absolute improvement of 21.5% in AP50 and 12.1% in Recall50 on the Front-3D dataset. Our method also achieves superior performance on ScanNet which is never seen during pretraining. We achieved an AP50 of 17% and a Recall50 of 39.5%, hence demonstrating an absolute performance improvement of 3% AP50 and 7.2% on Recall50 over the best-performing baseline. We further describe additional experimental details in supplementary. Specifically, we show that our pretraining improves performance on a variety of downstream tasks in Sec. C.3 and C.4. We discuss scaling performance and ablations in Sec. C.6 and show comparisons to other 3D pretraining baselines in Sec. C.5. Finally, we show qualita-

tive reconstruction and comparison results in Sec. D.

4. Conclusion

Neural Fields have shown emergent properties going beyond reconstruction. In this work, we primarily study scaling NeRFs for their self-supervised pretraining, through the use of masked auto-encoders. In essence, we use a single Transformer model to learn useful 3D representations from NeRF which we show to be highly effective for 3D transfer learning. To pre-train our representation, we curate a large-scale dataset totaling 1.6M+ images and 3,500+ scenes from 4 different sources. Our approach, NeRF-MAE significantly outperforms self-supervised 3D pretraining as well as NeRF scene understanding baselines on a variety of downstream tasks. As it requires readily available and cheaper to obtain posed-RGB data and with a strong empirical quantitative performance over baselines, we hope this avenue is a potential future direction of 3D pretraining.

References

- [1] Jonathan T Barron, Ben Mildenhall, Matthew Tancik, Peter Hedman, Ricardo Martin-Brualla, and Pratul P Srinivasan. Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5855–5864, 2021. [1](#)
- [2] Jiazhong Cen, Zanwei Zhou, Jiemin Fang, Chen Yang, Wei Shen, Lingxi Xie, Xiaopeng Zhang, and Qi Tian. Segment anything in 3d with nerfs. In *NeurIPS*, 2023. [1](#)
- [3] Anpei Chen, Zexiang Xu, Fuqiang Zhao, Xiaoshuai Zhang, Fanbo Xiang, Jingyi Yu, and Hao Su. Mvsnerf: Fast generalizable radiance field reconstruction from multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 14124–14133, 2021. [1](#)
- [4] Anthony Chen, Kevin Zhang, Renrui Zhang, Zihan Wang, Yuheng Lu, Yandong Guo, and Shanghang Zhang. Pimae: Point cloud and image interactive masked autoencoders for 3d object detection. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5291–5301, 2023. [1](#), [2](#)
- [5] Wenzheng Chen, Jun Gao, Huan Ling, Edward J. Smith, Jaakko Lehtinen, Alec Jacobson, and Sanja Fidler. Learning to predict 3d objects with an interpolation-based differentiable renderer. In *NeurIPS*, 2019. [1](#)
- [6] Xinlei Chen*, Saining Xie*, and Kaiming He. An empirical study of training self-supervised vision transformers. *arXiv preprint arXiv:2104.02057*, 2021. [2](#)
- [7] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5828–5839, 2017. [4](#), [1](#), [3](#), [8](#), [9](#), [10](#), [11](#)
- [8] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018. [3](#), [2](#)
- [9] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020. [2](#)
- [10] Francis Engelmann, Fabian Manhardt, Michael Niemeyer, Keisuke Tateno, and Federico Tombari Marc Pollefeys. Open-Set 3D Scene Segmentation with Rendered Novel Views. In *International Conference on Learning Representations (ICLR)*, 2024. [1](#)
- [11] Daniel Y. Fu, Tri Dao, Khaled K. Saab, Armin W. Thomas, Atri Rudra, and Christopher Ré. Hungry hungry hippos: Towards language modeling with state space models, 2023. [11](#)
- [12] Huan Fu, Bowen Cai, Lin Gao, Ling-Xiao Zhang, Jiaming Wang, Cao Li, Qixun Zeng, Chengyue Sun, Rongfei Jia, Binqiang Zhao, et al. 3d-front: 3d furnished rooms with layouts and semantics. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 10933–10942, 2021. [4](#), [1](#), [3](#), [7](#), [8](#), [10](#), [11](#)
- [13] Ayaan Haque, Matthew Tancik, Alexei Efros, Aleksander Holynski, and Angjoo Kanazawa. Instruct-nerf2nerf: Editing 3d scenes with instructions. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2023. [1](#)
- [14] Chenhang He, Ruihuang Li, Shuai Li, and Lei Zhang. Voxel set transformer: A set-to-set approach to 3d object detection from point clouds. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2022. [1](#)
- [15] K He, X Chen, S Xie, Y Li, P Dollár, and R Girshick. Masked autoencoders are scalable vision learners. *arXiv preprint arXiv: 211106377*, 2021. [3](#), [1](#), [2](#), [4](#), [11](#)
- [16] Georg Hess, Johan Jaxing, Elias Svensson, David Hagerman, Christoffer Petersson, and Lennart Svensson. Masked autoencoders for self-supervised learning on automotive point clouds. *arXiv preprint arXiv:2207.00531*, 2022. [1](#), [2](#)
- [17] Yining Hong, Chunru Lin, Yilun Du, Zhenfang Chen, Joshua B Tenenbaum, and Chuang Gan. 3d concept learning and reasoning from multi-view images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9202–9212, 2023. [3](#)
- [18] Ji Hou, Benjamin Graham, Matthias Nießner, and Saining Xie. Exploring data-efficient 3d scene understanding with contrastive scene contexts. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15587–15597, 2021. [3](#), [4](#), [5](#), [9](#)
- [19] Benran Hu, Junkai Huang, Yichen Liu, Yu-Wing Tai, and Chi-Keung Tang. Nerf-rpn: A general framework for object detection in nerfs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 23528–23538, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [9](#), [10](#)
- [20] Di Huang, Sida Peng, Tong He, Honghui Yang, Xiaowei Zhou, and Wanli Ouyang. Ponder: Point cloud pre-training via neural rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 16089–16098, 2023. [11](#)
- [21] Muhammad Zubair Irshad, Thomas Kollar, Michael Laskey, Kevin Stone, and Zsolt Kira. Centersnap: Single-shot multi-object 3d shape reconstruction and categorical 6d pose and size estimation. In *IEEE International Conference on Robotics and Automation (ICRA)*, 2022. [1](#)
- [22] Muhammad Zubair Irshad, Sergey Zakharov, Rares Ambrus, Thomas Kollar, Zsolt Kira, and Adrien Gaidon. Shapo: Implicit representations for multi object shape appearance and pose optimization. In *European Conference on Computer Vision (ECCV)*, 2022. [1](#)
- [23] Muhammad Zubair Irshad, Sergey Zakharov, Katherine Liu, Vitor Guizilini, Thomas Kollar, Adrien Gaidon, Zsolt Kira, and Rares Ambrus. Neo 360: Neural fields for sparse view synthesis of outdoor scenes. In *International Conference on Computer Vision (ICCV)*, 2023. [1](#)
- [24] Yoonwoo Jeong, Seungjoo Shin, Junha Lee, Chris Choy, Anima Anandkumar, Minsu Cho, and Jaesik Park. Perfception: Perception using radiance fields. *Advances in Neural Information Processing Systems*, 35:26105–26121, 2022. [1](#), [2](#)

- [25] Li Jiang, Zetong Yang, Shaoshuai Shi, Vladislav Golyanik, Dengxin Dai, and Bernt Schiele. Self-supervised pre-training with masked shape prediction for 3d scene understanding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1168–1178, 2023. 1
- [26] Justin* Kerr, Chung Min* Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lerf: Language embedded radiance fields. In *International Conference on Computer Vision (ICCV)*, 2023. 1
- [27] Abhijit Kundu, Kyle Genova, Xiaoqi Yin, Alireza Fathi, Caroline Pantofaru, Leonidas J Guibas, Andrea Tagliasacchi, Frank Dellaert, and Thomas Funkhouser. Panoptic neural fields: A semantic object-aware neural scene representation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12871–12881, 2022. 1
- [28] Yaqian Liang, Shanshan Zhao, Baosheng Yu, Jing Zhang, and Fazhi He. Meshmae: Masked autoencoders for 3d mesh data analysis. In *European Conference on Computer Vision*, pages 37–54. Springer, 2022. 1, 2
- [29] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2117–2125, 2017. 3, 4
- [30] Haotian Liu, Mu Cai, and Yong Jae Lee. Masked discrimination for self-supervised learning on point clouds. *Proceedings of the European Conference on Computer Vision (ECCV)*, 2022. 1
- [31] Lingjie Liu, Jiatao Gu, Kyaw Zaw Lin, Tat-Seng Chua, and Christian Theobalt. Neural sparse voxel fields. *NeurIPS*, 2020. 1
- [32] Yichen Liu, Benran Hu, Junkai Huang, Yu-Wing Tai, and Chi-Keung Tang. Instance neural radiance field. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 787–796, 2023. 1, 2
- [33] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10012–10022, 2021. 3, 4
- [34] Mayank Lunayach, Sergey Zakharov, Dian Chen, Rares Ambrus, Zsolt Kira, and Muhammad Zubair Irshad. Fsd: Fast self-supervised single rgb-d to categorical 3d objects. 2023. 1
- [35] Manolis Savva*, Abhishek Kadian*, Oleksandr Maksymets*, Yili Zhao, Erik Wijmans, Bhavana Jain, Julian Straub, Jia Liu, Vladlen Koltun, Jitendra Malik, Devi Parikh, and Dhruv Batra. Habitat: A Platform for Embodied AI Research. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, 2019. 3
- [36] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *ECCV*, 2020. 1, 2
- [37] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multi-resolution hash encoding. *ACM Transactions on Graphics (ToG)*, 41(4):1–15, 2022. 1, 2, 7
- [38] Michael Niemeyer and Andreas Geiger. Giraffe: Representing scenes as compositional generative neural feature fields. In *Proc. IEEE Conf. on Computer Vision and Pattern Recognition (CVPR)*, 2021. 1
- [39] Michael Niemeyer, Lars Mescheder, Michael Oechsle, and Andreas Geiger. Differentiable volumetric rendering: Learning implicit 3d representations without 3d supervision. In *CVPR*, 2020. 1
- [40] Michael Oechsle, Lars Mescheder, Michael Niemeyer, Thilo Strauss, and Andreas Geiger. Texture fields: Learning texture representations in function space. In *Proceedings IEEE International Conf. on Computer Vision (ICCV)*, 2019. 1
- [41] Julian Ost, Fahim Mannan, Nils Thuerey, Julian Knodt, and Felix Heide. Neural scene graphs for dynamic scenes. In *CVPR*, 2021. 1
- [42] Yatian Pang, Wenxiao Wang, Francis EH Tay, Wei Liu, Yonghong Tian, and Li Yuan. Masked autoencoders for point cloud self-supervised learning. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part II*, pages 604–621. Springer, 2022. 1, 3, 2, 4, 5, 9
- [43] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Edward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library, 2019. 6
- [44] Songyou Peng, Michael Niemeyer, Lars Mescheder, Marc Pollefeys, and Andreas Geiger. Convolutional occupancy networks. In *European Conference on Computer Vision (ECCV)*, 2020. 11
- [45] Albert Pumarola, Enric Corona, Gerard Pons-Moll, and Francesc Moreno-Noguer. D-NeRF: Neural Radiance Fields for Dynamic Scenes. In *CVPR*, 2021. 1
- [46] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. In *NeurIPS*, 2017. 1
- [47] Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks? *Advances in Neural Information Processing Systems*, 34:12116–12128, 2021. 2
- [48] Santhosh Kumar Ramakrishnan, Aaron Gokaslan, Erik Wijmans, Oleksandr Maksymets, Alexander Clegg, John M Turner, Eric Undersander, Wojciech Galuba, Andrew Westbury, Angel X Chang, Manolis Savva, Yili Zhao, and Dhruv Batra. Habitat-matterport 3d dataset (HM3d): 1000 large-scale 3d environments for embodied AI. In *Thirty-fifth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2021. 4, 1, 3, 7, 8, 9, 10, 11
- [49] Edoardo Remelli, Artem Lukoianov, Stephan Richter, Benoit Guillard, Timur Bagautdinov, Pierre Baque, and Pascal Fua. Meshsdf: Differentiable iso-surface extraction. In *Advances*

- in *Neural Information Processing Systems*, pages 22468–22478. Curran Associates, Inc., 2020. 1
- [50] Mike Roberts, Jason Ramapuram, Anurag Ranjan, Atulit Kumar, Miguel Angel Bautista, Nathan Paczan, Russ Webb, and Joshua M Susskind. Hypersim: A photorealistic synthetic dataset for holistic indoor scene understanding. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 10912–10922, 2021. 4, 1, 3, 10, 11
- [51] Barbara Roessle, Jonathan T Barron, Ben Mildenhall, Pratul P Srinivasan, and Matthias Nießner. Dense depth priors for neural radiance fields from sparse input views. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12892–12901, 2022. 1, 3, 7, 10
- [52] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III* 18, pages 234–241. Springer, 2015. 3
- [53] Danila Rukhovich, Anna Vorontsova, and Anton Konushin. Fcaf3d: Fully convolutional anchor-free 3d object detection. In *European Conference on Computer Vision*, pages 477–493. Springer, 2022. 2
- [54] Danila Rukhovich, Anna Vorontsova, and Anton Konushin. Imvoxelnet: Image to voxels projection for monocular and multi-view general-purpose 3d object detection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2397–2406, 2022. 3, 4
- [55] William Shen, Ge Yang, Alan Yu, Jansen Wong, Leslie Pack Kaelbling, and Phillip Isola. Distilled feature fields enable few-shot language-guided manipulation. In *7th Annual Conference on Robot Learning*, 2023. 1
- [56] Yawar Siddiqui, Lorenzo Porzi, Samuel Rota Bulò, Norman Müller, Matthias Nießner, Angela Dai, and Peter Kotschieder. Panoptic lifting for 3d scene understanding with neural fields. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9043–9052, 2023. 1
- [57] Anthony Simeonov, Yilun Du, Andrea Tagliasacchi, Joshua B. Tenenbaum, Alberto Rodriguez, Pulkit Agrawal, and Vincent Sitzmann. Neural descriptor fields: Se(3)-equivariant object representations for manipulation. 2022. 1
- [58] Edgar Sucar, Shikun Liu, Joseph Ortiz, and Andrew J Davison. imap: Implicit mapping and positioning in real-time. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6229–6238, 2021. 1
- [59] Towaki Takikawa, Joey Litalien, Kangxue Yin, Karsten Kreis, Charles Loop, Derek Nowrouzezahrai, Alec Jacobson, Morgan McGuire, and Sanja Fidler. Neural geometric level of detail: Real-time rendering with implicit 3D shapes. 2021. 1
- [60] Yucheng Tang, Dong Yang, Wenqi Li, Holger R Roth, Bennett Landman, Daguang Xu, Vishwesh Nath, and Ali Hatamizadeh. Self-supervised pre-training of swin transformers for 3d medical image analysis. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 20730–20740, 2022. 2
- [61] Alex Trevithick and Bo Yang. Grf: Learning a general radiance field for 3d representation and rendering. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15182–15192, 2021. 1
- [62] Can Wang, Menglei Chai, Mingming He, Dongdong Chen, and Jing Liao. Clip-nerf: Text-and-image driven manipulation of neural radiance fields, 2022. 1
- [63] Peng Wang, Lingjie Liu, Yuan Liu, Christian Theobalt, Taku Komura, and Wenping Wang. Neus: Learning neural implicit surfaces by volume rendering for multi-view reconstruction. *arXiv preprint arXiv:2106.10689*, 2021. 1
- [64] Sinong Wang, Belinda Z. Li, Madian Khabsa, Han Fang, and Hao Ma. Linformer: Self-attention with linear complexity, 2020. 11
- [65] Yue Wang, Yongbin Sun, Ziwei Liu, Sanjay E. Sarma, Michael M. Bronstein, and Justin M. Solomon. Dynamic graph cnn for learning on point clouds. *ACM Transactions on Graphics (TOG)*, 2019. 9
- [66] Zirui Wang, Shangzhe Wu, Weidi Xie, Min Chen, and Victor Adrian Prisacariu. NeRF—: Neural radiance fields without known camera parameters. *arXiv preprint arXiv:2102.07064*, 2021. 1
- [67] Philippe Weinzaepfel, Vincent Leroy, Thomas Lucas, Romain Brégier, Yohann Cabon, Vaibhav Arora, Leonid Antsfeld, Boris Chidlovskii, Gabriela Csurka, and Jérôme Revaud. Croco: Self-supervised pre-training for 3d vision tasks by cross-view completion. *Advances in Neural Information Processing Systems*, 35:3502–3516, 2022. 1
- [68] Weinzaepfel, Philippe and Leroy, Vincent and Lucas, Thomas and Brégier, Romain and Cabon, Yohann and Arora, Vaibhav and Antsfeld, Leonid and Chidlovskii, Boris and Csurka, Gabriela and Revaud Jérôme. CroCo: Self-Supervised Pre-training for 3D Vision Tasks by Cross-View Completion. In *NeurIPS*, 2022. 11
- [69] Rundi Wu, Ben Mildenhall, Philipp Henzler, Keunhong Park, Ruiqi Gao, Daniel Watson, Pratul P. Srinivasan, Dor Verbin, Jonathan T. Barron, Ben Poole, and Aleksander Holynski. Reconfusion: 3d reconstruction with diffusion priors. *arXiv*, 2023. 1
- [70] Yiheng Xie, Towaki Takikawa, Shunsuke Saito, Shiqin Yan, Numair Khan, Federico Tombari, James Tompkin, Vincent Sitzmann, and Srinath Sridhar. Neural fields in visual computing and beyond.(2021), 2021. 1
- [71] Chenfeng Xu, Bichen Wu, Ji Hou, Sam Tsai, Ruilong Li, Jialiang Wang, Wei Zhan, Zijian He, Peter Vajda, Kurt Keutzer, et al. Nerf-det: Learning geometry-aware volumetric representation for multi-view 3d object detection. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 23320–23330, 2023. 1, 2
- [72] Siming Yan, Zhenpei Yang, Haoxiang Li, Li Guan, Hao Kang, Gang Hua, and Qixing Huang. Implicit autoencoder for point cloud self-supervised representation learning. *arXiv preprint arXiv:2201.00785*, 2022. 1, 11
- [73] Xingguang Yan, Liqiang Lin, Niloy J. Mitra, Dani Lischinski, Danny Cohen-Or, and Hui Huang. Shapeformer:

- Transformer-based shape completion via sparse representation, 2022. [1](#)
- [74] Bangbang Yang, Yinda Zhang, Yinghao Xu, Yijin Li, Han Zhou, Hujun Bao, Guofeng Zhang, and Zhaopeng Cui. Learning object-compositional neural radiance field for editable scene rendering. In *International Conference on Computer Vision (ICCV)*, 2021. [1](#)
- [75] Yu-Qi Yang, Yu-Xiao Guo, Jian-Yu Xiong, Yang Liu, Hao Pan, Peng-Shuai Wang, Xin Tong, and Baining Guo. Swin3d: A pretrained transformer backbone for 3d indoor scene understanding, 2023. [3](#)
- [76] Lior Yariv, Peter Hedman, Christian Reiser, Dor Verbin, Pratul P. Srinivasan, Richard Szeliski, Jonathan T. Barron, and Ben Mildenhall. Baked sdf: Meshing neural sdfs for real-time view synthesis, 2023. [1](#)
- [77] Mingqiao Ye, Martin Danelljan, Fisher Yu, and Lei Ke. Gaussian grouping: Segment and edit anything in 3d scenes. *arXiv preprint arXiv:2312.00732*, 2023. [1](#)
- [78] Lin Yen-Chen, Pete Florence, Jonathan T. Barron, Alberto Rodriguez, Phillip Isola, and Tsung-Yi Lin. Inerf: Inverting neural radiance fields for pose estimation, 2021. [1](#)
- [79] Lin Yen-Chen, Pete Florence, Jonathan T. Barron, Tsung-Yi Lin, Alberto Rodriguez, and Phillip Isola. NeRF-Supervision: Learning dense object descriptors from neural radiance fields. In *IEEE Conference on Robotics and Automation (ICRA)*, 2022. [1](#)
- [80] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *CVPR*, 2021. [1](#), [2](#)
- [81] Xumin Yu, Lulu Tang, Yongming Rao, Tiejun Huang, Jie Zhou, and Jiwen Lu. Point-bert: Pre-training 3d point cloud transformers with masked point modeling. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 19313–19322, 2022. [2](#), [1](#), [9](#)
- [82] Wentao Yuan, Zhaoyang Lv, Tanner Schmidt, and Steven Lovegrove. Star: Self-supervised tracking and reconstruction of rigid objects in motion with neural rendering. In *CVPR*, 2021. [1](#)
- [83] Yanjie Ze, Ge Yan, Yueh-Hua Wu, Annabella Macaluso, Yuying Ge, Jianglong Ye, Nicklas Hansen, Li Erran Li, and Xiaolong Wang. Multi-task real robot learning with generalizable neural feature fields. *CoRL*, 2023. [1](#)
- [84] Cheng Zhang, Zhaopeng Cui, Yinda Zhang, Bing Zeng, Marc Pollefeys, and Shuaicheng Liu. Holistic 3d scene understanding from a single image with implicit representation. In *CVPR*, 2021. [1](#)
- [85] Renrui Zhang, Ziyu Guo, Peng Gao, Rongyao Fang, Bin Zhao, Dong Wang, Yu Qiao, and Hongsheng Li. Point-m2ae: Multi-scale masked autoencoders for hierarchical point cloud pre-training. *arXiv preprint arXiv:2205.14401*, 2022. [3](#), [1](#), [2](#), [4](#), [5](#), [9](#)
- [86] Zaiwei Zhang, Rohit Girdhar, Armand Joulin, and Ishan Misra. Self-supervised pretraining of 3d features on any point-cloud. 2021. [3](#), [1](#), [4](#), [5](#), [9](#)
- [87] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew Davison. In-place scene labelling and understanding with implicit scene representation. In *Proceedings of the International Conference on Computer Vision (ICCV)*, 2021. [1](#)
- [88] Haoyi Zhu, Honghui Yang, Xiaoyang Wu, Di Huang, Sha Zhang, Xianglong He, Tong He, Hengshuang Zhao, Chunhua Shen, Yu Qiao, and Wanli Ouyang. Ponderv2: Pave the way for 3d foundation model with a universal pre-training paradigm. *arXiv preprint arXiv:2310.08586*, 2023. [11](#)
- [89] Zihan Zhu, Songyou Peng, Viktor Larsson, Weiwei Xu, Hujun Bao, Zhaopeng Cui, Martin R Oswald, and Marc Pollefeys. Nice-slam: Neural implicit scalable encoding for slam. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12786–12796, 2022. [1](#)

NeRF-MAE: Masked AutoEncoders for Self-Supervised 3D Representation Learning for Neural Radiance Fields

Supplementary Material



Figure S1. **Data preparation for pretraining:** showing multi-view images, trained NeRF and extracted radiance and density grid for Instant-NGP [37] trained NeRFs for Front3D [12], HyperSim [50] and Habitat-Matterport3D [48]. The right column shows multi-view images, rendered RGB/depth images, and extracted radiance and density grid for ScanNet [7] dataset, which is trained using dense depth prior NeRF [51]. Note that we do not pre-train our representation on ScanNet, rather use this dataset for our cross-dataset transfer experiment for the downstream 3D OBB prediction task.

This supplementary material is divided into 9 sections. First we discuss related works in Section A, Method in Section B, followed by experiments in Section C and qualitative results in Section D. We provide details regarding network architecture and number of parameters in Section E. Next, we provide an overview of compute resources required to pretrain our representation in Section F. This is followed by a discussion related to implementation details in Section G and additional experimental analysis in Section H. Lastly, we discuss our pretraining datasets in Section I.

A. Related Works

3D Representation Learning: Supervised learning for 3D data [14, 21, 22, 46, 73] has made promising progress. Although powerful, these techniques require expensive 3D labels. Alternatively, self-supervised learning [25, 30, 34, 67, 72, 86] and masked autoencoders [15] have emerged as a way to learn strong representations for computer vision tasks without needing labeled data. Follow-up works ex-

tended this idea to the 3D domain by directly reconstructing masked 3D coordinates with point clouds [4, 42, 81, 85], meshes [28] or voxels as 3D representations [16]. While successful, they require careful adaptation of existing architectures due to irregularity in data structures and uneven information density whereas our approach uses standard Transformer modules, while also requiring readily available and cheaper to obtain posed 2D data compared with lidar scans or accurate mesh reconstructions.

NeRF as 3D Scene Representation: Advances in differentiable neural rendering [5, 39, 49, 59, 63, 84] have allowed for learning 3D representations using only image supervision. Neural Radiance Fields (NeRFs) [36] is one example that excels in novel-view synthesis. Extensions of NeRF study generalizable scene modelling [3, 23, 61, 69, 80], semantic and open-world 3D understanding [10, 26, 27, 56, 87], generative modelling [13, 38, 62], robotics [55, 58, 78, 83, 89], dynamic scenes [31, 41, 45, 82] and compositional representations [41, 74]. Recently, NeRFs have shown com-

Method	PixelNeRF [80]	NeRFRPN [19]	PerFception [24]	Voxel-MAE [16]	PiMAE [4]	SUNetR [60]	NeRF-MAE
# Scenes	40	515	1500	850	-	5000	3637
# Frames	4k	80k	2.5M	-	10k	-	1.65M
S.S. Pretraining	✗	✗	✗	✓	✓	✓	✓
Input Data	RGB*	RGB*	RGB*	Pointcloud	RGB-D	CT Scan	RGB*
3D Representation	NeRF	NeRF	NeRF	Voxel	Pointclouds	Voxel	NeRF
Data domain(s)	Single (1)	Single (1)	Single (1)	Single (1)	Single (1)	Multiple (5)	Multiple (4)

Table S2. Specifications of supervised or self-supervised approaches to 3D representation or generalizable learning, highlighting our novel self-supervised (S.S) 3D pretraining strategy for NeRFs which requires posed RGB images as input. * denotes posed RGB images.

petitive performance for supervised 3D tasks such as detection [19, 71] and segmentation [32] when compared with methods that require depth as input [53]. However, little attention has been given to representation learning for NeRFs. Our approach fills this gap by focusing on self-supervised pretraining for neural radiance fields, utilizing NeRF’s dense radiance and density grid for pretraining to improve performance on various downstream 3D tasks.

Vision Transformers: Computer vision has undergone a transformative shift with the emergence of vision transformers [9, 47]. Their ability to model global and local contexts has sparked interest in exploring methods to pre-train these backbones. MoCoV3 [6] and MAE [15] delved into different aspects of self-supervised Vision Transformers. In particular, MAE [15] achieved state-of-the-art results by drawing inspiration from BERT [8], which involves randomly masking words in sentences, and using masked image reconstruction. Subsequent research extended this concept to the 3D domain [16, 28, 42, 60], leading to impressive results in downstream 3D vision tasks.

B. Method

B.1. Preliminaries - NeRF and InstantNGP:

Given posed 2D images, NeRF represents a 3D scene implicitly. It uses a neural network, $f(\mathbf{x}, \theta)$, to predict color (\mathbf{c}_i) and density (σ_i) at any given 3D query position (\mathbf{x}_i) and viewing direction (θ_i) as input. The 4D color and density outputs are used in an alpha compositing process to generate rendered images through volume rendering with near and far bounds t_n and t_f , as highlighted in the equations below:

$$C(\mathbf{r}) = \int_{t_n}^{t_f} T(t) \sigma(\mathbf{r}(t)) \mathbf{c}(\mathbf{r}(t), \mathbf{d}) dt \quad (1)$$

where $T(t) = \exp\left(-\int_{t_n}^t \sigma(\mathbf{r}(s)) ds\right)$ and \mathbf{r} denotes the camera ray. Although our work is suited to any input NeRF formulation, we choose Instant-NGP [37] to model the radiance and density of a 3D scene using a small Multi-layer perceptron (MLP) with a sparse occupancy grid. A fast ray marching routine is utilized by ray compacting based on occupancy grid values. We also utilize a multi-resolution hash encoding from [37] for faster NeRF training. An example trained scene is shown in Figure S4 (b).

B.2. Neural Radiance Field Grid Representation:

Radiance and Density Grid Sampling from Instant-NGP NeRF: NeRF-MAE reconstructs input masked NeRF volumes to pretrain standard 3D Swin Transformers. The first step in this process is to uniformly sample a radiance and density feature volume from a trained NeRF model (Figure S4), in our case, Instant-NGP NeRF [37]. Querying a pre-trained NeRF for radiance and density information on a regular grid allows **1.** extracting an explicit representation that compactly captures the original 3D scene, **2.** is invariant to the NeRF formulation used, and **3.** opens up the possibility of utilizing existing deep-learning architectures developed for 3D tasks such as 3D RPN (Section C.3) for detection and segmentation [19, 32]. To sample radiance and density information, we query each trained NeRF model from all the training cameras in the traceable scene volume and average the resulting output.

Formally, Let $\mathcal{G} \in \mathbb{R}^{H \times W \times D \times 4}$ be a 4D grid representing the scene volume (Figure S4 c), where, H, W, D denote the height, width, and depth dimension. For a grid point (i, j, k) on a spatial 3D grid, the 4 channel values for each grid point (i, j, k) are the mean of the values obtained by this function $f(x, \theta)$ for all viewing directions. Mathematically, we can write this as:

$$(r_{i,j,k}, g_{i,j,k}, b_{i,j,k}, \alpha_{i,j,k}) = \frac{1}{N} \sum_{\theta=1}^N f(x_{i,j,k}, \theta) \quad (2)$$

where N is the number of training images, α can be obtained from volume density (σ_x) as $\alpha = 1 - \exp(-\sigma_i \|\mathbf{x}_i - \mathbf{x}_{i+1}\|)$ and $\|\mathbf{x}_i - \mathbf{x}_{i+1}\| = 0.01$, is a small preset distance [19]. Similar to [19, 32], we determine the traceable volume of the scene by enlarging the axis-aligned bounding box encapsulating all cameras and objects (where available) in the scene. Compared to point-based pretraining approaches, [42, 85] which carry information only at a surface-level and the distribution of information is relatively uneven (see Fig. 2), our approach makes use of dense information along all camera rays indicating that NeRF grids are similar to images in terms of information density; hence we show empirically that we are able to learn significantly better representations compared to other self-supervised pretraining baselines (see. Section C.5).




Figure S2. **NeRF-MAE Architecture:** Our method utilizes a U-Net [52] style architecture employing Swin Transformers [33] as the encoder to encode the RGB/density grid into meaningful multi-resolution low-level features, and transposed convolution layers at each stage with skip connections using residual blocks from the features of the encoder.

B.3. NeRF Pretraining Datasets

A key component of our approach is employing a single model to pre-train strong 3D representations in a self-supervised manner from NeRFs on a large number of scenes with varying complexity and realism. As shown in Figure S3, we train our representation in a fully self-supervised manner on the following dataset mix:

Front3D [12]: A large-scale dataset containing 18k rooms with 3D objects. We extend [19] to render 5x as many scenes, totaling 1.2k scenes comprising 425k images for our pretraining. The rendering time of this dataset is 6 days on 8 A100 GPUs.

HM3D [48]: A large-scale realistic dataset of 1000 scans. We render approximately 1.1k scenes comprising 1.2M images using Habitat [35]. Following [17], we collect the data by randomly selecting 80 navigable points within each room’s bounding box. At each navigable point, we rotate the agent 360° to render 12 frames with a $[10^\circ, -10^\circ]$ elevation change. The rendering time of this dataset is around 2 days on 8 A100 GPUs.

Hypersim [50]: A synthetic dataset with real-world realism and 3D annotations. We use the rendered images (250 scenes, 25k images) from [50] for our NeRF pretraining as they provide high-quality poses for reconstruction.

ScanNet [7]: A real-world dataset for indoor scenes. We keep this dataset for our cross-dataset transfer experiments and do not use this for pretraining. 90 out of the 1500 scenes are chosen, similar to [19] and we use the feature grid extracted from dense-depth-prior NeRF [51] for our downstream task experiments which further confirms the generalizability of our approach to various NeRF variants.

For all datasets, extensive cleaning of the bounding box as well as semantic annotations was performed based on manual filtering, class filtering, and size filtering (see supplementary). This cleaning was only carried out for downstream task fine-tuning and not carried out for our pretraining. We will release the complete dataset along with code, upon acceptance, for reproducibility.

B.4. 3D Transformer Encoder

Considering the input to the standard 3D SwinTransformer [19, 75] to be the radiance and density grid ($\mathcal{G} \in \mathbb{R}^{H \times W \times D \times 4}$), we first mask out random patches of size ($P = p \times p \times p$), where $p = 4$, with a masking ratio of $m = 0.75$ to get masked radiance and density grid \mathcal{M} . Subsequently, we use P to divide the masked radiance and density grid into patches to create a series of 3D tokens of size $\frac{H}{p} \times \frac{W}{p} \times \frac{D}{p} \times E$ using a 3D convolutional operation, where we merge patch partitioning and linear embedding of each patch into one single step. Next, we add 3D positional embedding to each patch, without a classification token, to assign each token to a unique representation. These tokens are further processed by four different stages of shifted window attention and patch merging, where the input patches are subdivided into non-overlapping windows of size $W \times W \times W$, where $W = 4$, with local multi-head self-attention (W-MSA-3D) and shifted multi-head self-attention (SW-MSA-3D) performed for initial and subsequent layers in each block respectively (see Fig. S2). Each Swin Transformer 3D block also contains a 2-layer *MLP* with GeLU activation as non-linearity in between following shifted window-based self-attention modules. A linear norm (*LN*) layer follows

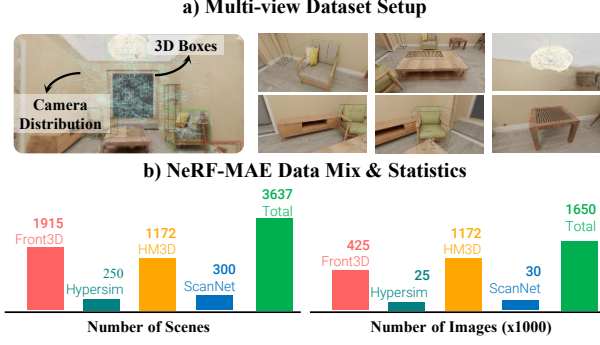


Figure S3. **NeRF-MAE dataset mix:** **a.** Multi-view dataset with camera distribution, **b.** diverse scenes from 4 different sources i.e. Front3D [12], Hypersim [50], Habitat-Matterport3D [48] and ScanNet [7] totaling over 3500 scenes and 1.6M images used for pretraining NeRF-MAE using a single model.

every W-MSA-3D and *MLP* with residual connection applied after each module. Due to the quadratic complexity in computing global self-attention for standard vision transformers, we employ shifted window attention [33] by displacing the windows by $[\frac{W}{2}, \frac{W}{2}, \frac{W}{2}]$ volume every layer. A patch merging layer follows every transform block, except for the last one, which merges four volumetric patches, so a reduction by a factor of 2 in height, width, and depth of the input volume occurs. A Linear Norm (*LN*) and an MLP layer follow patch merging to increase the feature dimension by a factor of 2. We use an embedding dimension $E = 96$, Swin block depths of $d = [2, 2, 18, 2]$ and number of heads $[3, 6, 12, 24]$ for each W-MSA-3D and SW-MSA-3D blocks.

C. Experiments

C.1. Baselines:

We compare 7 model variants across 3 different tasks to show the effectiveness of our method. 1) **NeRF-RPN** [19]: A state-of-the-art approach that predicts 3D boxes directly in NeRFs using 3D FPN [29]. 2) **Point-MAE** [42]: Extends 2D MAE [15] to 3D pointclouds and improves object classification and segmentation. 3) **Point-M2AE** [85]: Extends Point-MAE to add multiscale features for irregular point cloud representation learning. 4) **DepthContrast** [86]: Uses depth to learn 3D representations in a contrastive manner by comparing transformations of a 3D point cloud. 5) **Contrastive Scene Context** [18]: Leverages point correspondences along with spatial contexts to learn 3D scene level representation. 6) **ImVoxelNet** [54]: A strong multi-view 3D detection baseline. 7) **NeRF-MAE**: Our 3D representation learning approach utilizing posed 2D images as input and operates directly on scene’s volumetric radiance and density grids obtained through NeRFs.




Figure S4. **NeRF-MAE data processing flow for pretraining:** showing **a.** multi-view training data, **b.** trained NeRF representation, **c.** extracted RGB/Density grid from the trained NeRF and **d.** masked pretraining of the voxel-grid neural radiance field.

C.2. Datasets:

We pre-train our 3D representation in a self-supervised manner on Front3D [12], HM3D [48] and Hypersim [50]. For cross-dataset transfer 3D object detection, we test on hold-out Scannet [7] scenes, where we use splits provided by [19]. For further details, see Sec. B.3 and supplementary.

C.3. Downstream tasks

3D Object Detection in NeRF [19]: The task is to predict 3D OBBs. It entails regressing bounding box-offsets $\mathbf{t} = (x_0, y_0, z_0, x_1, y_1, z_1, \Delta\alpha, \Delta\beta)$, objectness score c and a single objectness p for each voxel. The task is similar to anchor-free RPN from [19].

Voxel-grid Super-Resolution: Given an input voxel grid (V) of resolution 160^3 , we predict an upsampled voxel-grid of higher resolutions, such as 384^3 . This can be denoted as $V_{\text{upsample}} = g(V_{\text{input}}, \text{resolution})$, where g is a U-Net style encoder-decoder network. It is an important task for fast grid upsampling, since querying an implicit MLP for higher resolution is slow.

Semantic Voxel Labelling: Given an input grid G with dimensions $W \times H \times D \times 4$, the task is to predict the class labels for each voxel. Every 3D voxel is assigned a class label, $S_{i,j,k}$ which are integers ranging from 1 to n , indicating the semantic category to which each voxel belongs. Mathematically, this can be represented as: $S_{i,j,k} \in \{1, 2, \dots, n\}$, for $1 \leq i \leq W, 1 \leq j \leq H, 1 \leq k \leq D$. Due to the diversity of objects in a scene, this task requires strong priors for generalization.

C.4. Comparisons with strong baselines for improved downstream performance

3D Object Detection Downstream Task: The results of our proposed method are summarized in Tables 1 and S4. We consistently outperform the state-of-the-art baseline

Method	3D OBB on ScanNet					
	Recall@50↑			AP@50↑		
	S	PT	Diff.	S	PT	Diff.
DepthContrast [86]	7.5	11.7	+4.2	2.4	4.1	+1.7
Scene Context [18]	7.5	12.2	+4.7	2.4	4.9	+2.5
NeRF-MAE (Ours)	32.3	39.5	+7.2	14.0	17.0	+3.0

Method	3D Segmentation on Front3D					
	mIOU↑			mAcc↑		
	S	PT	Diff.	S	PT	Diff.
PointMAE [42]	15.1	17.3	+2.2	20.9	22.8	+1.9
PointM2AE [85]	17.5	21.1	+3.6	23.4	27.6	+4.2
NeRF-MAE (Ours)	24.9	34.5	+9.6	33.8	45.0	+11.2

Table S3. **Quantitative comparison:**(§ C.5) showing our approach’s superior result compared to strong 3D representation learning and 3D MAE baselines using NeRF-rendered depth maps (Left), and our NeRF-grid as input (Right) [86] and [18] uses VoteNet, which requires full annotations such as instance masks and mean class size. **S** denotes starting from scratch and **PT** denotes pretrained. Please refer to the supplementary for more details.

methods on 3D object detection in NeRFs on both Front3D and ScanNet scenes. Among our variants, our method pre-trained on multiple data sources (Table 1, row 5) and with augmentations enabled during pretraining performs the best. Specifically, NeRF-MAE shows superior performance on the unseen test-set by achieving an AP50 of 63% and a Recall50 of 74.3%, hence demonstrating an absolute improvement of 21.5% in AP50 and 12.1% in Recall50 on the Front-3D dataset. Our method also achieves superior performance on ScanNet which is never seen during pretraining. We achieved an AP50 of 17% and a Recall50 of 39.5%, hence demonstrating an absolute performance improvement of 3% AP50 and 7.2% on Recall50 over the best-performing baseline.

Semantic Voxel-labelling and Voxel-Super Resolution: We further test our approach’s ability on two challenging downstream tasks i.e. semantic voxel-labelling and voxel super-resolution. The results are summarized in Tab. S5. The results clearly show our proposed approach consistently outperforms the strong baseline on all metrics which further indicates our networks’ ability to learn good representation for the challenging downstream 3D task. Specifically, we achieve 81% Acc, 45% mAcc, and 34.5% mIOU on voxel-labeling task which equals an absolute performance improvement of 6.8% Acc, 12.9% mAcc, and 9.8% mIOU metric against NeRF-RPN [19] on Front3D. Our method also achieves better 3D PSNR and low MSE metrics (absolute improvement of 1.02 for 256³ and 1.259 for 384³ grid show) on voxel-super resolution, showing our network learns strong representation for this dense task.

C.5. Comparisons with strong 3D representation learning baselines

Table S3 presents the quantitative evaluations of our approach’s performance in comparison with strong 3D representation learning baselines. NeRF-MAE consistently outperforms all competing methods with a clear margin. For a fair comparison, we provide NeRF-rendered depth maps to DepthContrast [86] and SceneContext [18] while providing our NeRF-grid to Point-MAE [42] and PointM2AE [85]. Specifically, we achieve +7.2% Recall50 improvement and +3% AP50 improvement, hence demonstrating an absolute

improvement number of +2.5% Recall50 and +0.5% AP50 on Scannet 3D OBB prediction task. We also achieve a +9.6% mIOU improvement and +11.2% mAcc improvement on Front3D voxel labelling, hence demonstrating an absolute improvement number of +6% mIOU and +7% mAcc over the best competing 3D pretraining baseline. This shows our method learns better representations by utilizing a dense volumetric NeRF-grid as well as an opacity-aware reconstruction objective.

C.6. Scaling Performance and Ablation Analysis

We analyze our network’s ability to learn useful 3D representations by studying how the amount and quality of pre-training data affects downstream performance. We summarize our results in Figure 3. To study scaling laws, we select Front3D dataset and use 10%, 25%, 50%, and 100% of the scenes for pretraining. Our test-set MSE and AP50 metrics clearly show our network’s ability to improve with more unlabeled data on the reconstruction and transfer learning tasks. Specifically, we achieved an absolute AP50 improvement of 10% when adding 10x the number of scenes for pretraining (1515 vs 151 scenes), emphasizing our network is able to learn better on more unlabelled data. To understand how the quality of NeRFs during the pretraining stage affects downstream performance, we train the input NeRFs to varying levels of accuracy, providing all 100% of the scenes at various PSNR values. Our results show that we can learn better representations, achieving downstream AP25 improvement of 36% when adding higher quality NeRFs to our pretraining (18 vs 28 2D PSNR).

Amount of unlabeled scenes on downstream 3D OBB task: We evaluate our method on the 3D OBB prediction task by limiting the available scenes for transfer learning. As shown in Tab. S4, our method surpasses the baseline across all scenarios, notably achieving a 42% AP50 with less than half the data requirement of the strong baseline; further highlighting our approach’s ability to enhance performance in data-limited settings.




Figure S5. **Qualitative 3D OBB prediction:** showing our approach’s superior results compared with NeRF-RPN [19].

D. Qualitative Results

We qualitatively analyze the performance of NeRF-MAE on reconstruction, 3D OBB prediction, and voxel-labeling tasks. As shown in Fig. 4, our approach outputs plausible reconstructions showing our network’s ability to find spatial relationships in the input data, which is useful for learning good representations. We further show our networks’ superior 3D OBB prediction performance in Figure S5. Our approach reconstructs accurate bounding boxes where the state-of-the-art baseline struggles with incomplete or erroneous detections. Finally, we show our networks’ superior semantic segmentation performance in Figure S6. This result further emphasizes our network learns good representations for a challenging 3D downstream task.

E. Network Architecture Details

In this section, we provide more details about our architectural design. We use PyTorch [43] for implementation and training our approach. We described in detail the encoder-decoder U-Net style architecture for pretraining in Section 2.1 (in the main paper). We train with a grid-resolution 160, a patch size $p = 4$, and mask out randomly 75% of the patches. We use 4 Swin Transformer blocks and the spatial feature size of the last block is $5 \times 5 \times 5$ with the channel dimension being 768. Consequently, we upsample features with skip connection decoders and use 4 decoders and a final convolution block to restore the original dimension i.e. $B \times 4 \times 160 \times 160 \times 160$. Below we specify the specifications and architectures of our task-specific decoders.




Figure S6. **Qualitative voxel-labeling downstream task generalization comparison:** showing our approach’s superior results compared with NeRF-RPN [19].

E.1. Task-specific Downstream 3D Heads

Our task-specific downstream heads are attached to the Swin-Transformer encoder. The goal of task-specific heads is to output dedicated meaningful outputs to regress the useful 3D information such as 3D OBBs and semantic labels for each voxel.

3D OBB Prediction: Our 3D OBB prediction head is similar to NeRF-RPN [19] except that for all of our transfer learning experiments and comparisons, we start with warm-started encoder weights via our large-scale self-supervised pretraining, whereas NeRF-RPN weights are started from scratch. Starting from scratch here denotes that the network weights were not pretrained on any other dataset but rather NeRF-RPN initializes the weights of linear layers with truncated normal distribution with a standard deviation of 0.02. We use the author’s [19] official implementation for all the comparisons. Specifically, we utilize an anchor-free RPN head from [19] which is based on the FCOS detector lifted to the 3D domain via replacing 2D convolutions with their 3D counterparts. We employ a combination of focal loss, IOU loss, and binary cross-entropy loss for downstream transfer learning task similar to [19].

3D Semantic Voxel-Labeling: The task of semantic voxel labeling is described in detail in Section C.3 (in the main paper). Specifically, we use the same decoders with skip connections i.e. \mathcal{D}_i , as described in Section 2.1 (in the main paper) where i ranges from 2 to 4. We replace \mathcal{D}_1 by incorporating an additional skip connection with the encoded output from the input grid (\mathcal{G}). This is followed by a final convolutional output block to output semantic

#Labelled	#Scenes	MAE pre-trained encoder (<i>Ours</i>)				NeRF-RPN [19] Start from scratch			
		Recall@25↑	Recall@50↑	AP@25↑	AP@50↑	Recall@25↑	Recall@50↑	AP@25↑	AP@50↑
10%	12	0.93	0.41	0.52	0.18	0.91	0.35	0.49	0.15
25%	30	0.94	0.51	0.67	0.36	0.94	0.49	0.67	0.29
50%	61	0.95	0.60	0.74	0.42	0.95	0.58	0.71	0.30
100%	122	0.96	0.67	0.79	0.54	0.96	0.62	0.78	0.41

Table S4. **Effect of amount of 3D labeled scenes on downstream 3D object detection performance:**(§ C.6) showing our approach’s superior results compared to strong baseline. The FPN network weights in both cases are initialized from scratch.

Super-Resolution Downstream Task					Semantic Voxel-Labeling Downstream Task				
Method	256 ³		384 ³		Metrics	Front3D [12]		HM3D [48]	
	PSNR↑	MSE↓	PSNR↑	MSE↓		NeRF-RPN [19]	Ours	NeRF-RPN [19]	Ours
NeRF-RPN [19]	16.25	0.024	16.08	0.025	mIOU↑	0.249	0.345	0.109	0.186
Ours	17.27	0.019	17.34	0.019	mAcc↑	0.338	0.450	0.160	0.264
					Acc↑	0.730	0.810	0.510	0.581

Table S5. **Voxel semantic-labelling and super-resolution quantitative results:**(§ C.4) showing superior performance of our method vs NeRF-RPN [19].

volume ($\mathcal{G}_S \in \mathbb{R}^{H \times W \times D \times C}$) of spatial dimension W, H, D where $W = H = D = 160$ and $C = 18$ for the Front3D dataset and 20 for the Habitat-Matterport3D dataset. We use a weighted masked cross-entropy loss, balanced by the inverse log propensity of each class, where propensity is defined to be the frequency of each class in the training dataset.

Voxel Super-Resolution: The task of voxel super-resolution is described in detail in Section C.3 (in the main paper). Specifically, we replace convolution decoders with four 3D convolution layers with instance norm and ReLU activation followed by an upsampling by a factor of 2. This is followed by a final upsampling of 1.6 and 2.4 depending on the output resolution size of 256³ and 384³ and a final 3D convolution. Similar to the original reconstruction loss (Section 2.1 (in the main paper)), we employ a masked color reconstruction loss enforced with the ground-truth up-sampled grid.

E.2. Parameters

NeRF-MAE efficiently trains on diverse scenes using a 70 million parameter network, highlighting its capability to handle complex and varied training data. Our Swin-S Transformer encoder architecture as well as the anchor-free RPN head remains the same as NeRF-RPN [19] for a fair comparison on the downstream 3D OBB prediction task. For reconstruction, we employ lightweight decoders (Section 2.1 in the main paper) since for our transfer learning experiments, we eventually discard the decoders and utilize the feature pyramids for downstream 3D tasks as described above.

F. Compute Resources

Our model was trained on 7 Nvidia A100 GPUs for 1200 epochs for 2 days and 17 hours. The cost for a pretraining run is around 19 GPU days with a GPU capacity of 82GB and an average utilization of 60%. To pretrain our representation, each NeRF was trained for 8 minutes on a single GPU for 100,000 steps. The total training time for training 3500 NeRF was 58 hours and finished in just over 2 days on 8 Nvidia A100 GPUs. Note that our pretraining strategy is agnostic to the NeRF technique used and can be replaced with any of the faster NeRF techniques developed. We show that our pretraining strategy equally applies to both instant-NGP [37] trained NeRF and dense depth prior NeRF [51] and doesn’t require any modification to our network architecture or training strategy.

G. Implementation Details

To pre-train NeRF-MAE, we utilize the Swin-S backbone of the SwinTransformer. As detailed in Section 2.1 (in the main paper), we utilize four swin blocks of depth 2, 2, 18, and 2 with the number of heads in each block equalling 3, 6, 12, and 24. We pre-train both encoder and decoder networks for 1200 epochs with the Adam optimizer and a maximum learning rate of $3e^{-4}$, a weight decay of $1e^{-3}$ with a one-cycle learning rate scheduling. We employ a batch size of 32 for our pretraining. We additionally utilize online data augmentation during both pretraining and transfer learning stages. We employ random flip, rotation, and scaling augmentation with a probability of 50%. Data augmentation is shown to improve our pretraining, as described in Section 3. For transfer learning, we train both our network and the baseline NeRF-RPN [19] for 1000 epochs for 3D OBB prediction and 500 epochs for voxel super-resolution with the same learning rate and strategy as our pretraining stage

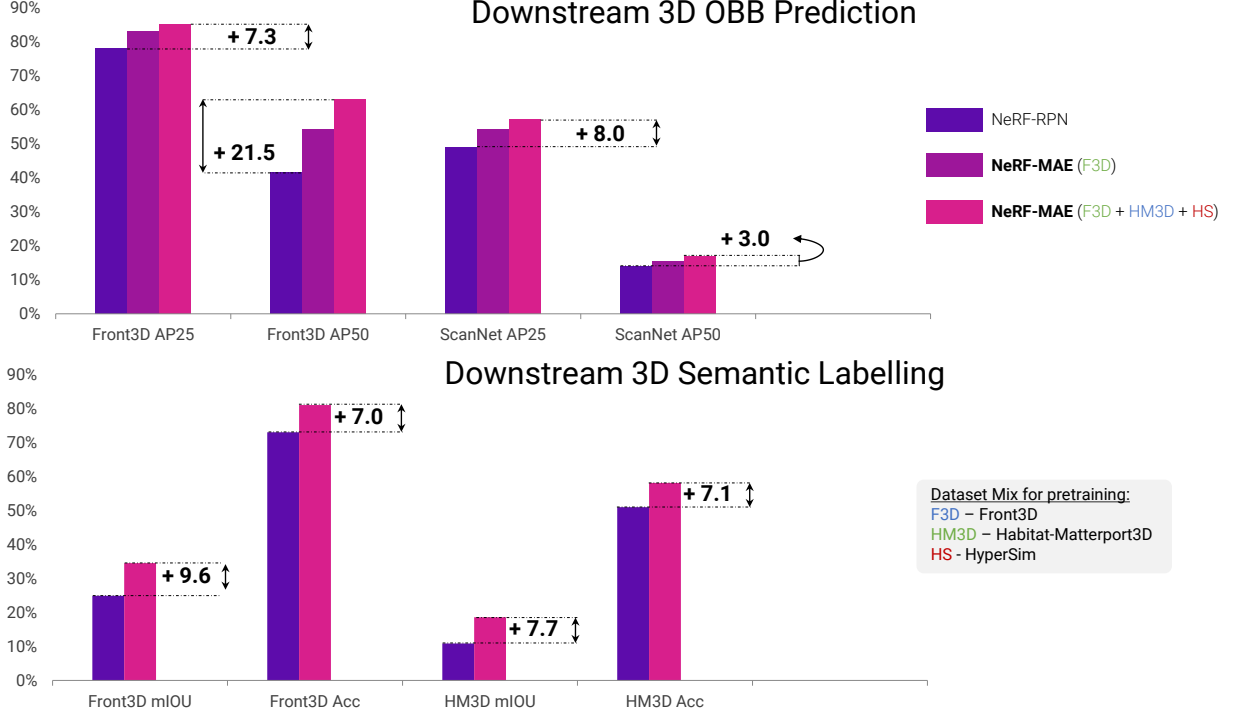


Figure S7. **Quantitative Comparison:** We show further quantitative analysis of results reported in Table 1 and Table S5 in our main paper. Specifically, we highlight improvement numbers for AP25 and AP50 on ScanNet [7] and Front3D [12] and mIOU and Acc for HM3D [48] and Front3D [12] over the state-of-the-art baseline. Furthermore, we also show that increasing unlabelled posed 2D data from different sources improves performance on 3D OBB prediction downstream task.

since we empirically determined that changing the learning rate and strategy results in decreased performance for both our method and the baseline. We additionally utilize a maximum normal gradient clip value of 0.1 and fine-tune the network until convergence. We determine transfer learning convergence on the hold-out validation set and employ early stopping based on AP50 metrics for 3D OBB prediction, accuracy metric for 3D semantic voxel labeling, and 3D PSNR metric for voxel super-resolution experiments. For transfer learning experiments, we use a batch size of 8.

H. Experiment Analysis

In this section, we present additional experimental analysis. First, we present in detail an additional discussion on the downstream 3D object detection, semantic voxel-labeling, and voxel super-resolution tasks in Section H.1. We present the additional analysis of scaling laws and ablation studies in Section H.2 specifically highlighting the architectural ablation of adding skip connections to the decoder and discussing masking ratio ablation. We present additional insights with comparison to 3D pretraining baselines in Section H.3.

H.1. Comparison on Downstream 3D Tasks

We analyze all of the improvement numbers mentioned in Table 1 and Table S5 by highlighting them in Figure S7. One can clearly see our pretraining consistently improves all metrics across different datasets, especially on ScanNet which is omitted from our pretraining and used as a hold-out dataset for our cross-dataset transfer experiments. The results also clearly show that downstream 3D task numbers for 3D OBB prediction improve with adding unlabelled data from diverse sources, in this case, adding unlabelled HM3D scenes to our pretraining improves ScanNet numbers showing the efficacy of our pretraining strategy. Note that we only utilized posed 2D data for pretraining, and no 3D information was used in our pretraining pipeline.

Next, we present results for the task of voxel grid super-resolution on HM3D dataset hold-out scenes never seen by the model during pretraining or transfer learning. We summarize these results in Table S6. One can clearly observe that our technique achieves a higher Recall@50 and AP@50 compared to strong baseline, specifically showing improvement of 9% in Recall@50 and 7% in AP@50, further confirming the efficacy of our pretraining strategy on a challenging dataset and 3D task. The results in Table S2

Method	256 ³		384 ³	
	PSNR↑	MSE↓	PSNR↑	MSE↓
NeRF-RPN* [19]	14.70	0.034	14.65	0.035
NeRF-MAE(Ours)	15.01	0.032	15.20	0.030

3D OBB on HM3D [48]		
Method	AP@25↑	Recall@50↑
NeRF-RPN [19]	0.30	0.31
NeRF-MAE (Ours)	0.37	0.40

Table S6. **Voxel super-resolution and 3D OBB prediction quantitative results** for HM3D [48] test-set showing superior performance of our method vs NeRF-RPN [19]. These results further confirm the findings of Tables 1 and S5 in our main paper that our pretraining helps improve downstream tasks on various datasets.

Pretraining scenes on downstream performance					
Pretraining Scenes		Metrics			
% Scenes	# Scenes	Recall@25↑	Recall@50↑	AP@25↑	AP@50↑
10%	151	0.94	0.61	0.74	0.45
25%	378	0.95	0.62	0.76	0.46
50%	757	0.95	0.66	0.77	0.51
100%	1515	0.96	0.67	0.79	0.54

Pretraining scenes on reconstruction quality				
Pretraining Scenes		Metrics		
% Scenes	# Scenes	Aug.	3D PSNR↑	MSE↓
10%	151	✗	14.20	0.038
25%	378	✗	14.98	0.032
50%	757	✗	15.42	0.028
100%	1515	✗	16.00	0.025
100%	1515	✓	19.09	0.012

Table S7. **Effect of percent of pretraining scenes** on downstream 3D object detection (left) and voxel-grid reconstruction (right) showing more unlabeled data helps learn better representation for downstream 3D tasks as well as better reconstruction quality of unmasked patches in the input radiance and density grid

NeRF-MAE	Front3D [19]		ScanNet [7]	
	Recall@50↑	AP@25↑	Recall@50↑	AP@25↑
w/o Skip	0.73	0.83	0.39	0.54
w/ Skip	0.74	0.85	0.40	0.57

Masking Ratio Ablation		
Method	AP@25↑	AP@50↑
65%	0.805	0.519
75%	0.839	0.546
85%	0.814	0.512

Table S8. **Effect of skip connections in the decoder (Left) and masking ratio ablation (Right):** shows NeRF-MAE with skip connection helps learn better representations for downstream tasks as well as 75% masking ratio achieves the best downstream performance.

(Left) further confirm the finding of Table S5 (Left), achieving consistently higher 3D PSNR and lower MSE metrics for hold-out test scenes on HM3D for the challenging voxel grid super-resolution task.

H.2. Additional Ablation Analysis

We report complete metrics for our scaling laws experiments shown in Figure 3 of the main paper. We report these metrics in Table S7. Both results further confirm our findings in the paper that our results consistently improve on reconstruction task as well as during transfer learning by utilizing more unlabelled data during pretraining. We additionally include ablations to our NeRF-MAE pretraining architecture in Table S8. We show the effect of adding skip connections to the lightweight decoders in Table S8 (Left). One can clearly see that adding skip connections during reconstruction increases the accuracy of downstream tasks, specifically increasing the AP25 by 2% and 3% on FRONT3D and ScanNet datasets respectively. We also analyze the effect of masking ratio on downstream task performance in Table S8 (Right). This table clearly shows our design choice of using 75% as the masking ratio for input radiance and density grid. This significantly increases the accuracy of 3D object detection downstream task.

H.3. Comparison to 3D pretraining baselines

We present a detailed comparison to 3D pretraining baselines [18, 42, 85, 86] in Table S2 in our main paper. Here,

we share more insights on these results. Despite given NeRF rendered depth maps (from the same NeRF model, NeRF-MAE uses to sample its radiance and density grid from) and ground truth poses, the baselines [18, 86] struggle to improve performance in an extreme data-scarcity setting (i.e. on cross-domain generalization where 90 scenes are available for finetuning), whereas, our approach, NeRF-MAE excels in these challenging scenarios. We conclude that this is due to the local training structure of DepthContrast and other MAE approaches, whereas NeRF-MAE pretraining captures both the local and global structure of the full-scenes with its Transformer-inspired architecture. Despite given NeRF’s radiance and density grid as an input to point-based pretraining baselines [42, 85], we see degraded performance. We conclude that these approaches use specialized modules like KNN used in Point-MAE [42], DGCNN [65] used in Point-BERT [81] which are specialized for modeling surface level information and require extra parameter tunings such as group size and number of centers. Even though PointM2AE [85] improved performance over PointMAE [42] using a hierarchical training, we show that they still struggle for this challenging downstream task where only a few labeled scenes are available for finetuning. On the contrary, our approach utilized standard Transformer modules and shows very effective downstream performance with an opacity and radiance-aware masked reconstruction objective.




Figure S8. **Multi-view Rendered vs Ground-truth posed RGB images** from four distinct datasets used in this paper i.e. Front3D [12], ScanNet [7], Habitat-Matterport3D [48] and Hypersim [50]. As mentioned in the main paper, we use Instant-NGP to train neural radiance fields for Front3D, Hypersim and HM3D, whereas Scannet scenes are trained using dense depth prior nerf [51], following [19]

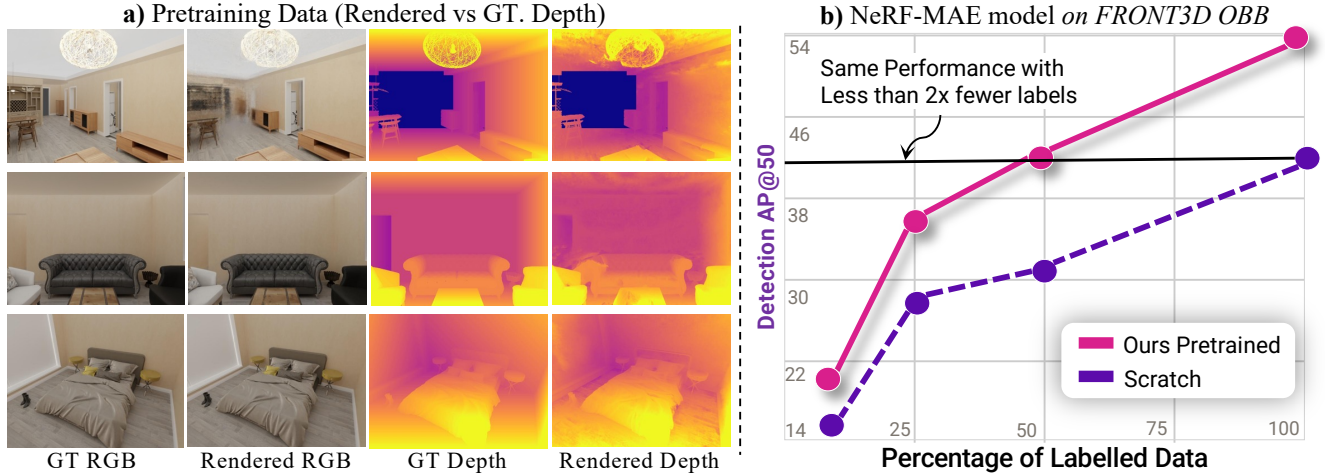


Figure S9. **Qualitative Rendered Depth (a) and Quantitative Detection Performance with a varying number of labeled scenes (b):** We show the qualitative comparison of ground truth vs rendered depth map from NeRFs used for our pretraining scenes showing faithful reconstruction of geometry for our input data. We further show in the chart on the right (b), that our pretraining requires much less data to achieve the same performance as the strong baseline.

I. Pretraining Datasets

We show our pretraining datasets in detail in Figure S1 and Figure S8. First, we show the data processing flow for each of the Front3D [12], Scannet [7], HM3D [48] and Hyper-sim [50] dataset in Figure S1 starting with multi-view images to get a NeRF trained for each scene and finally getting an explicit radiance and density grid for all scenes. We train each scene for 100,000 steps to ensure high-quality NeRF are obtained as input to our pretraining stage as further confirmed by our experiments in Figure 3 in the main paper, which quantitatively shows that better NeRF quality results in better pretraining and leads to improved downstream 3D task performance. We further show rendered RGB vs ground-truth RGB for selected scenes from all four datasets in Figure S8. The figure confirms that our NeRF quality for validation images is faithful to the ground-truth images. Secondly, it shows that our pretraining dataset for training good 3D representations from posed 2D data is very diverse in realism, lightning, shadows, number of objects in the scene as well as the setting of each scene i.e. ranging from bedroom scenes to office and living room scenes. This makes our 3D representation learning more robust to the varying input data and helps it improve performance for data outside the training data-distribution, as further confirmed in Section 3.

J. Limitations and Future Work

NeRF-MAE enables large-scale self-supervised pertaining of 3D representations using NeRF’s unique radiance and density grid as an input hence employing standard 3D

Transformer blocks to significantly improve various downstream 3D tasks. Although powerful, future work is still required to improve the efficiency of training this representation to allow for seamless and efficient input of a large variety of multi-view data from diverse data sources. This could be achieved using fast linear attention blocks [64] or employing a 3D version of state space models [11]. Another avenue for future work is communication between neural rendering and masking. Concurrent works have combined neural rendering with 3D representation learning [20, 88] albeit in a non-masking manner. Other works have used correspondences to mine multi-view image pairs to train MAE [15] and designed the task of cross-view completion [68] for unsupervised 3D representation learning from 2D images. Our work NeRF-MAE bridges the 2D and 3D domains by using Neural Radiance Fields and future work could look into backpropagating the network weights through both our reconstructed grid as well as neural rendered images to further improve performance. Another promising avenue for future work is to explore various downstream 3D tasks that could benefit from dense NeRF-MAE pertaining such as 3D scene reconstruction [44, 72].