

Know Your Neighbors: Improving Single-View Reconstruction via Spatial Vision-Language Reasoning

Rui Li¹ Tobias Fischer¹ Mattia Segu¹ Marc Pollefeys¹
 Luc Van Gool¹ Federico Tombari^{2,3}
¹ETH Zürich ²Google ³Technical University of Munich

Abstract

Recovering the 3D scene geometry from a single view is a fundamental yet ill-posed problem in computer vision. While classical depth estimation methods infer only a 2.5D scene representation limited to the image plane, recent approaches based on radiance fields reconstruct a full 3D representation. However, these methods still struggle with occluded regions since inferring geometry without visual observation requires (i) semantic knowledge of the surroundings, and (ii) reasoning about spatial context. We propose KYN, a novel method for single-view scene reconstruction that reasons about semantic and spatial context to predict each point’s density. We introduce a vision-language modulation module to enrich point features with fine-grained semantic information. We aggregate point representations across the scene through a language-guided spatial attention mechanism to yield per-point density predictions aware of the 3D semantic context. We show that KYN improves 3D shape recovery compared to predicting density for each 3D point in isolation. We achieve state-of-the-art results in scene and object reconstruction on KITTI-360, and show improved zero-shot generalization compared to prior work. Project page: <https://ruili3.github.io/kyn>.

1. Introduction

Humans have the extraordinary ability to estimate the geometry of a 3D scene from a single image, often including its occluded parts. It enables us to reason about where dynamic actors in the scene might move, and how to best navigate ourselves to avoid a collision. Hence, estimating the 3D scene geometry from a single input view is a long-standing challenge in computer vision, fundamental to autonomous navigation [16] and virtual reality applications [33]. Since the problem is highly ill-posed due to scale ambiguity, occlusions, and perspective distortion, it has traditionally been cast as a 2.5D problem [6, 31, 58], focusing on areas visible in the image plane and neglecting the non-



Figure 1. **Single-view scene reconstruction results.** We present the predicted 3D occupancy grids given a single input image. The camera is at the bottom left and points to the top right along the z -axis. Previous methods like BTS [51] struggle to recover accurate object shapes (green box) and exhibit trailing effects in unobserved areas (blue box). In contrast, KYN recovers more accurate boundaries and mitigates the trailing effects prevalent in prior art.

visible parts.

Recently, approaches based on neural radiance fields [39] have shown great potential in inferring the true 3D scene representation from a single [51, 59] or multiple views [50]. For instance, Wimbauer *et al.* [51] introduce BTS, a method that estimates a 3D density field from a single view at inference while being supervised only by photometric consistency given multiple posed views at training time.

Intuitively, given only a single image at inference, the model must rely on semantic knowledge from the neighboring 3D structure to predict the density of occluded points. However, existing approaches lack explicit semantic modeling and, by modeling density prediction independently for each point, are unaware of the semantic 3D context of the point’s surroundings. This results in the clear limitations

which we illustrate in Fig. 1. Specifically, prior work [51] struggles with accurate shape recovery (green) and further exhibits trailing effects (blue) in the absence of visual observation. We argue that, when considering a single point in 3D, its density highly depends on the semantic scene context, *e.g.* if there is an intersection, a parking lot, or a sidewalk visible in its proximity. This becomes more critical as we move further from the camera origin since the degree of visual coverage decreases with distance, and reconstructing the increasingly unobserved scene parts requires context from the neighboring points.

To this end, we present Know Your Neighbors (KYN), a novel approach for single-view scene reconstruction that predicts density for each 3D point in a scene by reasoning about its neighboring semantic and spatial context. We introduce two key innovations. We develop a vision-language (VL) modulation scheme that endows the representation of each 3D point in space with fine-grained semantic information. To leverage this information, we further introduce a VL spatial attention mechanism that utilizes language guidance to aggregate the visual semantic point representations across the scene and predict the density of each individual point as a function of the neighboring semantic context.

We show that, by injecting semantic knowledge and reasoning about spatial context, our method overcomes the limitations that prior art exhibits in unobserved areas, producing more plausible 3D shapes and mitigating their trailing effects. We summarize our contributions as follows:

- We propose KYN, the first single-view scene reconstruction method that reasons about semantic and spatial context to predict each point’s density.
- We introduce a VL modulation module to enrich point features with fine-grained semantic information.
- We propose a VL spatial attention mechanism that aggregates point representations across the scene to yield per-point density predictions aware of the neighboring 3D semantic context.

Our experiments on the KITTI-360 dataset [34] show that KYN achieves state-of-the-art scene and object reconstructions. Furthermore, we demonstrate that KYN exhibits better zero-shot generalization on the DDAD dataset [18] compared to the prior art.

2. Related Work

Monocular depth estimation. Estimating depth from a single view has been extensively studied over the last decade [35, 36, 55–58], both in a supervised and a self-supervised manner. Supervised methods directly minimize the loss between the predicted and ground truth depths [1, 11]. For these, varying output representations [1, 2, 14], network architectures [1, 27, 43, 61], and loss functions [1, 48, 52] have been proposed. Recent methods explore training unified depth models on large datasets, tackling

challenges like varying camera intrinsics [12, 20, 42, 56, 58] and dataset bias [3]. Self-supervised methods cast the problem as a view synthesis task and learn depth via photometric consistency on image pairs. Existing works have investigated how to handle dynamic objects [7, 13, 17, 30, 31, 46], different network architectures [37, 53, 62, 64] and leveraging additional constraints [19, 32, 44, 47]. Our method falls in the self-supervised category. However, we estimate a true 3D representation from a single view, as opposed to the 2.5D representation produced by traditional depth estimation.

Semantic priors for depth estimation. Previous depth estimation methods use semantic information to enhance 2D feature representations with different fusion strategies [8, 19, 22, 32], or to remove dynamic objects [5, 28] during training. These methods utilize semantic information in the 2D representation space. On the contrary, we use semantic information to enhance 3D point representations and to guide our 3D spatial attention mechanism.

Neural radiance fields. Neural radiance fields (NeRFs) [39, 59] learn a volumetric 3D representation of the scene from a set of posed input views. In particular, they use volumetric rendering in order to synthesize novel views by sampling volume density and color along a pixel ray. Recent multi-view reconstruction methods [49, 54, 60] take inspiration from this paradigm, reformulating the volume density function as a signed distance function for better surface reconstruction. These methods are focused on single-scene optimization using multi-view constraints, often representing the scene with the weights of a single MLP.

To address the issue of generalization across scenes, Yu *et al.* [59] propose PixelNeRF to train a CNN image encoder across scenes that is used to condition an MLP, predicting volume density and color without multi-view optimization during inference. However, their approach is limited to small-scale and synthetic datasets. Recently, Wimbauer *et al.* [51] proposed BTS, an extension of PixelNeRF to large-scale outdoor scenes. They omit the color prediction and, during training, use reference images to query color for a 3D point given a density field that represents the 3D scene geometry. While this simplification allows them to scale to large-scale outdoor scenes, their method falls short in predicting accurate geometry for occluded areas (Fig. 1). We address this shortcoming by injecting fine-grained semantic knowledge and reasoning about semantic 3D context when querying the density field, leading to better shape recovery for occluded regions in particular.

Scene as occupancy. A recent line of work infers 3D scene geometry as voxelized 3D occupancy [4, 38] from a single image. These works predict the occupancy and semantic class of each 3D voxel based on exhaustive 3D annotations. Therefore, these methods rely heavily on manually anno-

tated datasets. Further, the predefined voxel resolution limits the fidelity of their 3D representation. In contrast, our method does not rely on labor-intensive manual annotations and represents the scene as a continuous density field.

Semantic priors for NeRFs. Various works integrate semantic priors into NeRFs. While some utilize 2D semantic or panoptic segmentation [15, 26, 63], others leverage 2D VL features [13, 24, 41] and lift these into 3D space by distilling them into the NeRF. This enables the generation of 2D segmentation masks from new viewpoints, segmenting objects in 3D space, and discovering or removing particular objects from a 3D scene. While the aforementioned methods focus on the classical multi-view optimization setting, we instead focus on single-view input and leverage semantic priors to improve the 3D representation itself.

3. Method

Problem setup. Given an input image \mathbf{I}_0 , its corresponding intrinsics $K_0 \in \mathbb{R}^{3 \times 4}$ and pose $T_0 \in \mathbb{R}^{4 \times 4}$, we aim to reconstruct the full 3D scene by estimating the density for each 3D point $\mathbf{x} \in \mathbb{R}^3$ among point set $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^M$

$$\sigma_i = f(\mathbf{I}_0, K_0, T_0, \mathbf{X}, \theta), \quad (1)$$

where the density σ_i of point \mathbf{x}_i is a function of the point set \mathbf{X} and image \mathbf{I}_0 along with its camera intrinsic/extrinsics. f denotes the network and θ represents its parameters. The density σ_i can be further transformed to the binary occupancy score $o_i \in \{0, 1\}$ with a predefined threshold τ . During training (Sec. 3.3), additional images \mathbf{I}_n are incorporated with their corresponding intrinsics K_n and extrinsics T_n , with $n \in \{1, \dots, N\}$, providing multiview supervision.

Overview. We illustrate our method in Fig. 2. Given an input image \mathbf{I}_0 , we first extract image and VL feature maps F_{app} and F_{vis} . Next, we fuse the image and VL features into a single feature map F_{fused} and further utilize category-wise text features to compute a segmentation map S . We then use intrinsics K_0 to project the 3D point set \mathbf{X} to the image plane and query F_{fused} , yielding point-wise visual features. In parallel, we retrieve point-wise text features by querying the segmentation map S and looking up the corresponding category-wise text features t for each 3D point. Given the point-wise visual and text features, we use our VL modulation layers to augment the features with fine-grained semantic information. We aggregate the point-wise features with VL spatial attention, adding 3D semantic context to the point-wise features. We guide the attention mechanism with the VL text features. We finally predict a per-point density that is thresholded to yield the 3D occupancy map.

3.1. Vision-Lanugage Modulation

We detail how we extract point-wise visual features and how we augment the visual features with semantic information using our VL modulation module.

Point-wise visual feature extraction. Given the input image \mathbf{I}_0 , we extract image features from standard image encoders [21] and VL image encoder [29]

$$\begin{aligned} F_{\text{app}} &= f(\mathbf{I}_0, \theta_{\text{app}}), \\ F_{\text{vis}} &= f(\mathbf{I}_0, \theta_{\text{vis}}), \end{aligned} \quad (2)$$

where F_{app} and F_{vis} refer to the appearance features and VL image features, respectively. We freeze the VL image encoder weights θ_{vis} to retain the pre-trained semantics. We then fuse the features by concatenation followed by 2 convolutional layers, yielding F_{fused} . To obtain point-wise features for the 3D points $\mathbf{X} = \{\mathbf{x}_i\}_{i=1}^M$, we extract the fused feature F_{fused} w.r.t. the projected 2D coordinates $p_0(\mathbf{x}_i)$ of each 3D point. Then, we combine each 3D point feature with a positional embedding $\gamma(\cdot)$ encoding its position in normalized device coordinates (NDC). We obtain the fused point-wise visual feature $v_i \in \mathbb{R}^{1 \times C}$ for a 3D point \mathbf{x}_i

$$v_i = \text{Concat}(F_{\text{fused}}(p_0(\mathbf{x}_i)), \gamma(\mathbf{x}_i^0)), \quad (3)$$

where $\text{Concat}(\cdot, \cdot)$ is the feature concatenation operation, and \mathbf{x}_i^0 is the 3D position w.r.t. \mathbf{I}_0 's coordinate.

Point-wise text feature extraction. As the text feature size does not align with the image space, it can not be extracted directly by querying projected 2D coordinates. To this end, we first derive the semantic category of each 3D point through 2D segmentations. Then we use it to associate text features with each 3D point. We utilize the category names from outdoor scenes [9] as prompts to the text encoder, yielding the text features $t \in \mathbb{R}^{Q \times C}$, where Q is the number of categories and C is the feature channel. We then use the visual feature $F_{\text{vis}} \in \mathbb{R}^{H \times W \times C}$ from the VL image encoder, to obtain the 2D semantic map by computing cosine similarity between VL image and text features

$$S = \arg \max_{q \in \{1, \dots, Q\}} \frac{F_{\text{vis}} \otimes t^T}{\|F_{\text{vis}}\| \|t\|}, \quad (4)$$

where \otimes denotes matrix multiplication, S denotes the segmentation map by maximizing the similarity scores between VL image and text features along the category dimension. We then compute the semantic category for each 3D point by querying S with projected 2D coordinates and then leverage it to obtain the text feature of each 3D point.

$$\begin{aligned} s_i &= S(p_0(\mathbf{x}_i)), \\ g_i^t &= t(s_i), \end{aligned} \quad (5)$$

where $g_i^t \in \mathbb{R}^C$ is the text feature of the 3D point \mathbf{x}_i .

VL modulation layers. To augment the 3D point features with rich semantics from both VL image and text features, we propose the VL modulation layers, which integrate both image feature v_i and text representations g_i^t of a 3D point \mathbf{x}_i

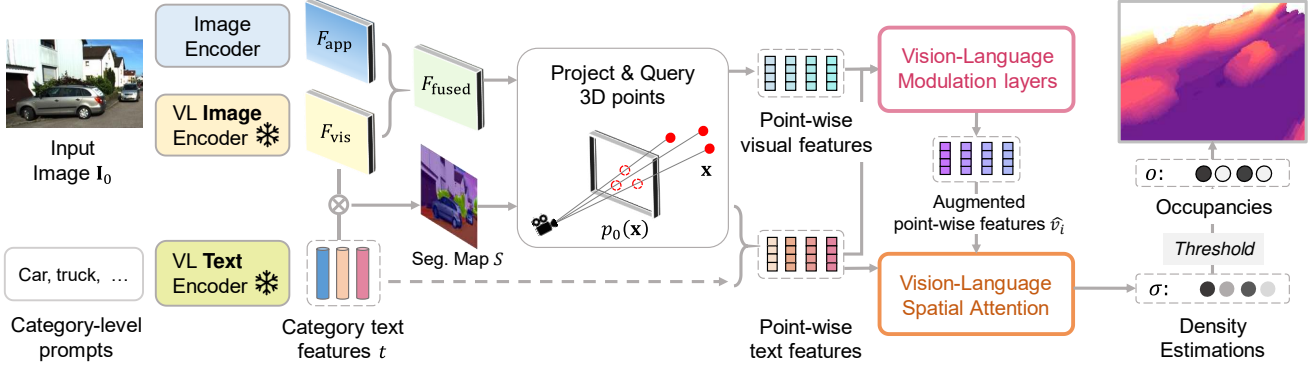


Figure 2. **Overview.** Given an input image I_0 , we use two image encoders to obtain features (F_{app} , F_{vis}), and fuse these into feature map F_{fused} . We further extract category-level text features and a segmentation map S . For a given 3D point set \mathbf{X} , we query the extracted features by projecting them onto the image plane yielding point-wise visual and text features. Next, the VL modulation layers endow the point representation with fine-grained semantic information. Finally, the VL spatial attention aggregates these point representations across the 3D scene, yielding density predictions aware of the 3D semantic context.

for better scene geometry reasoning. The module is composed of $L = 4$ modulation layers, each conducting modulation with multiplication operations

$$v_i^{l+1} = \text{ReLU}(\text{FC}(v_i^l) \odot \text{FC}(g_i^t)), \quad (6)$$

where $\text{FC}(\cdot)$ stands for the fully connected layer, \odot is the element-wise product between features, $\text{FC}(g_i^t)$ denotes the text features encoded by a single fully-connected layer and is shared across different modulation layers. We set $v_i^1 = v_i$ at the first modulation layer, and iterate through different layers. We utilize skip connections to inject the initial visual information v_i^1 after the modulated feature v_i^{l+1} at level l , using concatenation followed by one fully-connected layer. The output feature $\hat{v}_i = v_i^L$ denotes the 3D point feature of \mathbf{x}_i augmented with rich image and text semantics.

3.2. VL Spatial Attention

Next, we dive into our VL spatial attention mechanism that aggregates the extracted point-wise features across the scene in a global-to-local fashion. First, we combine the whole set of point-wise visual features $\{\hat{v}_i\}_{i=1}^M$ and text features $\{g_i^t\}_{i=1}^M$ into $V \in \mathbb{R}^{M \times C'}$ and $C_t \in \mathbb{R}^{M \times C}$. Then, we aggregate these features using a cross-attention operation in 3D space. Specifically, we leverage linear attention [23, 45] over appropriately split point sets to achieve memory-efficient spatial context reasoning.

Category-informed cross-attention. We take the point-wise features V as queries and values, and leverage text-based feature C_t as the keys in the linear attention. Specifically, we project the features with fully connected layers to

keys, queries, and values

$$\begin{aligned} F_Q &= f(V, \theta_Q), \\ F_K &= f(C_t, \theta_K), \\ F_V &= f(V, \theta_V), \end{aligned} \quad (7)$$

where all features are in $\mathbb{R}^{M \times \hat{C}}$, where \hat{C} denotes the feature dimension before the attention. We then compute the global context score $G \in \mathbb{R}^{\hat{C} \times \hat{C}}$ by attending to the key and value features, and then correlating with query features by

$$\begin{aligned} G &= \text{Softmax}(F_K^T) \otimes F_V, \\ F_{\text{final}} &= \text{Softmax}(F_Q / \sqrt{D}) \otimes G, \end{aligned} \quad (8)$$

where F_{final} denotes the spatially aggregated point-wise features. The density value is estimated by a single fully connected layer followed by a Softplus(\cdot) function

$$\sigma = \text{Softplus}(\text{FC}(F_{\text{final}})). \quad (9)$$

Reducing the memory footprint. As the point features are sampled from the entire 3D space, simultaneously processing all point features can hit computational bottlenecks even with linear attention. To this end, we randomly split the initial points into chunks, to ensure that each chunk is identically distributed. Then we conduct the spatial point attention separately within each chunk and combine the density estimations afterward. As such, the attention can aggregate semantic point representations with both spatial awareness and efficiency.

3.3. Training Process

Our method achieves self-supervision by computing the photometric loss between the reconstructed and target colors. We extract the 2D feature map of I_0 to get point-wise



Figure 3. **Qualitative comparisons on KITTI-360 dataset.** We illustrate the scene reconstructions as voxel grids, where the camera is on the left side and points to the right along the z -axis. A lighter voxel color indicates higher voxel positions. Compared to previous methods that struggle with corrupted and trailing shapes, our method produces faithful scene geometry, especially for occluded areas.

representations, then partition all images $\{\mathbf{I}_0\} \cup \{\mathbf{I}_n\}_{n=1}^N$ into a source set N_{source} and a loss set N_{loss} following previous practice [51]. We render the RGB of N_{loss} from the corresponding colors of N_{source} similar to the self-supervised depth estimation [17]. Instead of resorting to the whole image, we perform patch-based image supervision [51] to reduce memory footprints. For each pixel on a patch, we sample the 3D points \mathbf{x}_i on its back-projected ray and conduct density estimation. Let \mathbf{x}_i and \mathbf{x}_{i+1} be the adjacent sampled pixels on a ray, we calculate the RGB information by volume rendering the *sampled* color [51]

$$\alpha_i = \exp(1 - \sigma_{\mathbf{x}_i} \delta_i) \quad T_i = \prod_{j=1}^{i-1} (1 - \alpha_j), \quad (10)$$

$$\hat{d} = \sum_{i=1}^S T_i \alpha_i d_i \quad \hat{c}_k = \sum_{i=1}^S T_i \alpha_i c_{\mathbf{x}_i, k}, \quad (11)$$

where δ_i denotes the distance between adjacent sampled points \mathbf{x}_i and \mathbf{x}_{i+1} along the ray, and α_i refer to the probability that the ray ends between \mathbf{x}_i and \mathbf{x}_{i+1} . Note that the color $c_{\mathbf{x}_i, k} = \mathbf{I}_k(p_k(\mathbf{x}_i))$ is the sampled RGB value from the view k in the source set N_{source} , to obtain a better geometry. \hat{d} and \hat{c}_k represent the terminating depth and the rendered color.

As we sample the pixels in a patch-wise manner during training, the rendered RGB and depth are also organized

in patches. Let \hat{P}_k as the rendered patch from view k in the source N_{source} , P as the supervisory patch from N_{loss} , and d' as the patch depth of P , the loss function is defined following previous methods [17, 51]

$$\mathcal{L} = \mathcal{L}_{ph} + \lambda_e \mathcal{L}_e, \quad (12)$$

where $\lambda_e = 10^{-3}$, \mathcal{L}_{ph} and \mathcal{L}_e are photometric loss and edge-aware smoothness loss [17] on patches

$$\mathcal{L}_{ph} = \min_{k \in N_{\text{render}}} \left(\lambda_1 \text{L1} \left(P, \hat{P}_k \right) + \lambda_2 \text{SSIM} \left(P, \hat{P}_k \right) \right), \quad (13)$$

$$\mathcal{L}_e = |\delta_x d'_i| e^{-|\delta_x P|} + |\delta_y d'_i| e^{-|\delta_y P|}, \quad (14)$$

where $\lambda_1 = 0.15$ and $\lambda_2 = 0.85$, δ_x, δ_y denotes the gradient along the horizontal and vertical directions.

4. Experiments

To demonstrate the effectiveness of our proposed method, we compare with existing works [17, 51, 59, 62] in single-view scene reconstruction, including both depth estimation [17] and radiance field based methods [51, 59]. We evaluate both the 3D scene (Sec. 4.4) and object (Sec. 4.5) reconstruction results on the KITTI-360 dataset in short and long ranges. We conduct extensive ablations (Sec. 4.6) to verify the effectiveness of each contribution, compare our method with existing semantic feature fusion techniques, as

	Method	O _{acc} ↑	IE _{acc} ↑	IE _{rec} ↑
4-20m	Monodepth2 [17]	0.90	n/a	n/a
	Monodepth2 [17] + 4m	0.90	0.59	0.66
	PixelNeRF [59]	0.89	0.62	0.60
	BTS [51]	0.92	0.66	0.64
	Ours	0.92	0.70	0.72
4-50m	Monodepth2 [17]	0.82	n/a	n/a
	Monodepth2 [17] + 4m	0.81	0.54	0.76
	PixelNeRF [59]	0.82	0.56	0.68
	BTS [51]	0.84	0.61	0.53
	Ours	0.86	0.63	0.73

Table 1. **Comparison of scene reconstruction on KITTI-360.** Our method achieves the best overall performance in both the near and far evaluation range.

well as evaluate our method’s performance with the broader supervision range. Moreover, we demonstrate our method’s zero-shot generalization ability in Sec. 4.7.

4.1. Datasets

We use the KITTI-360 [34] dataset for training and evaluation, as it captures static scenes using cameras deployed with wide baselines, *i.e.*, two stereo cameras, and two side cameras (fisheye cameras) at each timestep, which facilitate learning full 3D shapes by self-supervision. During the training phase, we use all cameras from two time steps, *i.e.* 8 cameras in total, to train our model. We use an input resolution of 192×640 and choose the left stereo camera from the first time step to extract the image features. We split all cameras randomly into N_{render} and N_{loss} for sampling colors and loss computation, as is illustrated in Sec. 3.3. Furthermore, we use the DDAD [18] dataset to evaluate the zero-shot generalization capability of the models trained on KITTI-360. We select testing sequences with more than 50 images and use 384×640 image resolution.

4.2. Evaluation

We follow the experimental protocol in [51] to evaluate 3D occupancy prediction. Specifically, we sample 3D grid points on 2D slices parallel to the ground plane. For KITTI-360, we report scores within distance ranges [4, 20] and [4, 50] meters, where the latter provides a more challenging evaluation scenario. For ground-truth generation, we follow [51] and accumulate 3D LiDAR points across time, and set 3D points lying out of all depth surfaces as unoccupied, otherwise set to occupied. Unlike [51], which accumulate only 20 LiDAR sweeps often leading to inaccurate occluded scene geometry, we accumulate up to 300 LiDAR frames. We also provide results with a 20-frame accumulated ground truth in the supplementary material for reference. For the DDAD dataset, we accumulate up to 100 LiDAR frames due to the limited sequence length and evaluate in the [4, 50] meters range.

Metrics. We adopt the evaluation metrics in [51] and measure overall reconstruction (O_{acc}) and occluded reconstruction (IE_{acc} , IE_{rec}) accuracies. Specifically, O_{acc} computes the accuracy between the prediction and the ground truth in the full area of the evaluation range, thus reflecting the overall performance of the reconstruction. IE_{acc} computes the accuracy of the invisible areas specifically, *i.e.* without direct visual observation in I_0 . IE_{rec} computes the recall of both invisible and empty areas, which evaluates the reconstruction of the occluded empty space. The three metrics focus on different aspects of the reconstruction quality.

Scene- and object-level evaluation. In addition to evaluating the performance of the whole scene, we focus on object reconstruction in particular because they are of particular interest compared to *e.g.* the road plane. To this end, we manually annotate the object areas in the ground-truth occupancy maps and compute the evaluation metrics on these object areas. We refer the reader to the supplementary material for details.

4.3. Implementation Details

We implement our method using Pytorch [40] and train it on NVIDIA Quadro RTX 6000 GPUs. The appearance network is similar to [17] with pre-trained weights on ImageNet [10]. We adopt LSeg [29] as the visual-language network and freeze its parameters. The model is trained using Adam [25] optimizer with a learning rate of 10^{-4} for 25 epochs, which is reduced to 10^{-5} after 120k iterations. During the training phase, we sample 4096 patches across loss set N_{loss} , each patch contains 8×8 pixels. We further sample 64 points along each ray following [51].

4.4. Scene Reconstruction

We compare our method with recent single-view scene reconstruction methods using self-supervision [17, 51, 59]. Specifically, we train Monodepth2 [17] on our benchmark as the base depth estimation method. Since the depth network cannot infer occluded geometry, we use a handcraft criterion to set areas 4m behind the depth surface as empty space (Monodepth2 + 4m). We also compare our method with the NeRF-based methods [51, 59] following the same training protocol. As shown in the Tab. 1, compared to previous methods, our method achieves both the best overall performance (O_{acc}) and the best occluded area reconstruction (IE_{acc} , IE_{rec}). Note that Monodepth2 (+ 4m) also yields a competitive performance in terms of invisible and empty space reconstruction (IE_{rec}), but it relies on hand-crafted criteria that cannot learn true 3D in the scene. Qualitative comparisons are shown in Fig. 3. We show the reconstructed occupancy grids, where the camera is on the left side and points to the right along the z -axis within [4, 50m] range. Our method demonstrates obvious qualitative superiority in reasoning occluded object shapes against the inherent am-

Method		O _{acc} ↑	IE _{acc} ↑	IE _{rec} ↑
4-20m	Monodepth2 [17]	0.69	n/a	n/a
	Monodepth2 [17] + 4m	0.70	0.53	0.52
	PixelNeRF [59]	0.67	0.53	0.49
	BTS [51]	0.79	0.69	0.60
	Ours	0.80	0.69	0.70
4-50m	Monodepth2 [17]	0.65	n/a	n/a
	Monodepth2 [17] + 4m	0.68	0.48	0.59
	PixelNeRF [59]	0.66	0.56	0.58
	BTS [51]	0.72	0.61	0.48
	Ours	0.75	0.64	0.68

Table 2. **Comparison of object reconstruction on KITTI-360.** Our method outperforms other methods in all metrics in both the short and long evaluation range.

biguity. Notably, it substantially reduces the trailing effects.

4.5. Object Reconstruction

We evaluate the object reconstruction performance by computing the metrics within the annotated object areas. As shown in Tab. 2, our method achieves competitive or better results in the [4, 20m] evaluation range. Meanwhile, it achieves obvious improvements for all metrics in the [4, 50m] range, demonstrating its effectiveness in reasoning ambiguous geometries away from the camera origin. We further show reconstructions of different categories in Fig. 4. Our method generates faithful object shapes with reasonable estimates of occluded geometry for different categories including fences, trees, cars, *etc.*, showing clear improvements over existing methods.

4.6. Ablation Studies

We evaluate the effectiveness of each contribution by separately ablating the VL modulation (Sec. 4.6.1) and the VL spatial attention (Sec. 4.6.2). We further compare with existing techniques injecting semantics from related tasks (Sec. 4.6.3). Moreover, we show our method’s improvement under a broader supervision range (Sec. 4.6.4).

4.6.1 Ablation study on VL Modulation

We evaluate the effectiveness of VL modulation by adding different components upon baseline method [51], which uses only appearance feature F_{app} from the generic backbone [21]. In Tab. 3, we enhance the baseline by incorporating the fused VL image features (F_{fused}) and injecting image/text semantics with the VL Modulation (VL-Mod.). We find that merely using the fused feature F_{fused} from the VL image encoder does not contribute to overall improvement. As a comparison, the proposed VL Modulation significantly improves the performances by properly interacting with image and text features.

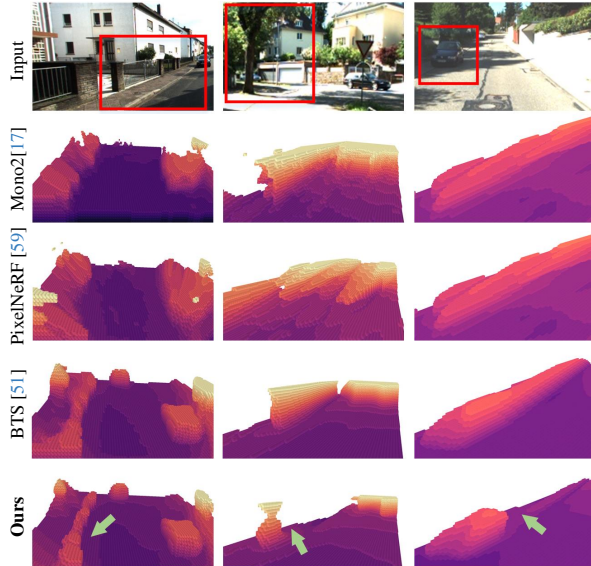


Figure 4. **Object reconstruction in the KITTI-360 dataset [34].** From left to right: Reconstructions of the fence, tree, and car. Our method produces more faithful object geometries, in particular in occluded areas and for various semantic categories.

F_{app}	F_{fused}	VL-Mod.	Scene Recon.			Object Recon.		
			O _{acc}	IE _{acc}	IE _{rec}	O _{acc}	IE _{acc}	IE _{rec}
✓			0.84	0.60	0.53	0.72	0.61	0.48
	✓		0.84	0.60	0.55	0.72	0.61	0.48
		✓	0.85	0.63	0.64	0.73	0.63	0.59

Table 3. **Ablations study on VL modulation.** We report the performance in the [4, 50m] range. Naively introducing the VL image feature does not improve performance. Our VL modulation correlating the image and text features yields the best scores.

F_{app}	F_{fused}	VL-Mod.	Attn.	VL-Attn.	Scene Recon.			Object Recon.		
					O _{acc}	IE _{acc}	IE _{rec}	O _{acc}	IE _{acc}	IE _{rec}
✓					0.84	0.60	0.53	0.72	0.61	0.48
✓			✓		0.85	0.61	0.60	0.73	0.61	0.56
	✓			✓	0.85	0.61	0.67	0.72	0.60	0.62
	✓			✓	0.85	0.60	0.66	0.73	0.62	0.61
	✓	✓	✓		0.86	0.63	0.75	0.74	0.62	0.73
	✓	✓		✓	0.86	0.63	0.73	0.75	0.63	0.68

Table 4. **Ablation study on spatial attention.** We report the performance in the [4, 50m] range. The spatial attention improves in each variant by enabling the 3D context awareness. Combining the VL features with spatial attention yields the best performance.

4.6.2 Ablation Study on Spatial Attention

In Tab. 4, we provide experimental support for the effectiveness of using spatial context, both via spatial attention and VL spatial attention. We first add spatial attention over image appearance features only (row: 1→2). Without the VL features, aggregating spatial context still yields notable im-

Semantics	Method	O _{acc}	IE _{acc}	IE _{rec}
–	Baseline	0.84	0.60	0.52
Semantic Feat.	Plain Fusion	0.84	0.61	0.51
	2D Fusion - LDLS [32]	0.84	0.60	0.55
	2D Fusion - SAFENet [8]	0.84	0.60	0.51
VL Feat.	Fusing VL image feature	0.84	0.60	0.55
	Ours	0.86	0.63	0.73

Table 5. **Comparison with semantic feature fusion techniques.** We compare with 2D feature fusion techniques from semantic-guided depth estimation [8, 32]. Our method achieves the best performance with visual and language semantic enhancement and 3D context awareness.

Supervision	Method	Scene Recon.			Object Recon.		
		O _{acc}	IE _{acc}	IE _{rec}	O _{acc}	IE _{acc}	IE _{rec}
1s in the future	BTS [51]	0.84	0.61	0.53	0.72	0.61	0.48
	Ours	0.86	0.63	0.73	0.75	0.64	0.68
1-4s in the future	BTS [51]	0.86	0.68	0.72	0.74	0.64	0.75
	Ours	0.87	0.69	0.77	0.78	0.68	0.75

Table 6. **Evaluation of different supervisory ranges.** We report the performance in [4, 50m] range. Our method with the standard supervision range (1s in the future) yields comparable performance with [51] using broader supervision (1-4s in the future). Meanwhile, it achieves further improvement when using a broader supervision range.

provement. We further show that the introduction of semantics via VL modulation (VL-Mod) helps independent of the spatial aggregation mechanism (row 3→5, row: 4→6), underpinning that adding VL text features outperforms using VL image features only. When combining both VL modulation and spatial attention mechanisms, we achieve the best performances (rows 5, 6). Additionally, we observe a small albeit notable gain when also injecting VL features into the spatial aggregation (VL-Attn), however, mainly improving details are not captured in current metrics. Please refer to the supplementary material for details.

4.6.3 Comparison with Other Semantic Guidances

As semantic cues are vital in other tasks such as depth estimation, we investigate whether the techniques used in the related literature [8, 32] are useful for single-view scene reconstruction. We use the pre-trained DPT [43] semantic network to provide pre-trained semantic features, and incorporate different feature fusion techniques [8, 32] to our density prediction pipeline. As shown in Tab. 6, we find that conducting 2D feature fusion does not lead to notable improvements over the baseline, which is consistent even with incorporating VL image feature fusion. However, by appropriately interacting with image and text features with 3D context awareness, our method outperforms related techniques by a notable margin.

Method	O _{acc}	IE _{acc}	IE _{rec}
PixelNeRF [59]	0.55	0.45	0.23
BTS [51]	0.48	0.50	0.16
Ours	0.59	0.58	0.42

Table 7. **Generalization to DDAD with KITTI-360 trained model.** We evaluate in the [4, 50m] range. Our method demonstrates better zero-shot ability compared to previous methods.

4.6.4 Improvement with Broader Supervision Range

As single-view reconstruction is supervised by multiple posed images, the performance can be improved by expanding the supervisory range during training. To this end, we investigate if our method can achieve consistent improvement using a broader supervisory range. In standard supervision, we use fisheye views at the next time step (1s in the future). In the broader supervisory range, we randomly incorporate fisheye views within the next [1-4s] timestep, yielding diverse supervisory ranges with a maximum coverage of 40m. We compare with BTS [51] in Tab. 6. Our method with the standard supervisory range produces comparable results to [51] with a broader supervision range. Enhanced by a broader range of supervision, our method achieves further improvement over [51].

4.7. Zero-shot Generalization on DDAD

We evaluate the zero-shot generalization of KITTI-360 trained models on the DDAD dataset. We report the [0, 50] meters range scene reconstruction scores in Tab. 7. Our method outperforms previous work, showing the effectiveness of the VL guidance for zero-shot generalization.

5. Conclusion

In this paper, we proposed KYN, a new method for single-view reconstruction that estimates the density of a 3D point by reasoning about its neighboring semantic and spatial context. To this end, we incorporate a VL modulation module to enrich 3D point representations with fine-grained semantic information. We further propose a VL spatial attention mechanism that makes the per-point density predictions aware of the 3D semantic context. Our approach overcomes the limitations of prior art [51], which treated the density prediction of each point independently from neighboring points and lacked explicit semantic modeling. Extensive experiments demonstrate that endowing KYN with semantic and contextual knowledge improves both scene and object-level reconstruction. Moreover, we find that KYN better generalizes out-of-domain, thanks to the proposed modulation of point representations with strong vision-language features. The incorporation of VL features not only enhances the performance of KYN, but also holds the potential to pave the way towards more general and open-vocabulary 3D scene reconstruction and segmentation techniques.

References

- [1] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Adabins: Depth estimation using adaptive bins. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4009–4018, 2021. 2
- [2] Shariq Farooq Bhat, Ibraheem Alhashim, and Peter Wonka. Localbins: Improving depth estimation by learning local distributions. In *European Conference on Computer Vision*, pages 480–496. Springer, 2022. 2
- [3] Shariq Farooq Bhat, Reiner Birkel, Diana Wofk, Peter Wonka, and Matthias Müller. Zoedepth: Zero-shot transfer by combining relative and metric depth. *arXiv preprint arXiv:2302.12288*, 2023. 2
- [4] Anh-Quan Cao and Raoul de Charette. Monoscene: Monocular 3d semantic scene completion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3991–4001, 2022. 2
- [5] Vincent Casser, Soeren Pirk, Reza Mahjourian, and Anelia Angelova. Depth prediction without the sensors: Leveraging structure for unsupervised learning from monocular videos. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 8001–8008, 2019. 2
- [6] Junda Cheng, Gangwei Xu, Peng Guo, and Xin Yang. Coatsnet: Fully exploiting convolution and attention for stereo matching by region separation. *International Journal of Computer Vision*, 132(1):56–73, 2024. 1
- [7] JunDa Cheng, Wei Yin, Kaixuan Wang, Xiaozhi Chen, Shijie Wang, and Xin Yang. Adaptive fusion of single-view and multi-view depth for autonomous driving. *arXiv preprint arXiv:2403.07535*, 2024. 2
- [8] Jaehoon Choi, Dongki Jung, Donghwan Lee, and Changick Kim. Safenet: Self-supervised monocular depth estimation with semantic-aware feature extraction. *arXiv preprint arXiv:2010.02893*, 2020. 2, 8
- [9] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016. 3
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 6
- [11] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. In *Advances in neural information processing systems*, pages 2366–2374, 2014. 2
- [12] Jose M Facil, Benjamin Ummenhofer, Huizhong Zhou, Luis Montesano, Thomas Brox, and Javier Civera. Camconvs: Camera-aware multi-scale convolutions for single-view depth. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11826–11835, 2019. 2
- [13] Ziyue Feng, Liang Yang, Longlong Jing, Haiyan Wang, YingLi Tian, and Bing Li. Disentangling object motion and occlusion for unsupervised multi-frame monocular depth. *arXiv preprint arXiv:2203.15174*, 2022. 2, 3
- [14] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2002–2011, 2018. 2
- [15] Xiao Fu, Shangzhan Zhang, Tianrun Chen, Yichong Lu, Lanyun Zhu, Xiaowei Zhou, Andreas Geiger, and Yiyi Liao. Panoptic nerf: 3d-to-2d label transfer for panoptic urban scene segmentation. In *2022 International Conference on 3D Vision (3DV)*, pages 1–11. IEEE, 2022. 3
- [16] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361. IEEE, 2012. 1
- [17] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3828–3838, 2019. 2, 5, 6, 7
- [18] Vitor Guizilini, Rares Ambrus, Sudeep Pillai, Allan Raventos, and Adrien Gaidon. 3d packing for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2485–2494, 2020. 2, 6
- [19] Vitor Guizilini, Rui Hou, Jie Li, Rares Ambrus, and Adrien Gaidon. Semantically-guided representation learning for self-supervised monocular depth. In *International Conference on Learning Representations*, 2020. 2
- [20] Vitor Guizilini, Igor Vasiljevic, Dian Chen, Rareş Ambrus, and Adrien Gaidon. Towards zero-shot scale-aware monocular depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9233–9243, 2023. 2
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 3, 7
- [22] Hyunyoung Jung, Eunhyeok Park, and Sungjoo Yoo. Fine-grained semantics-aware representation enhancement for self-supervised monocular depth estimation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12642–12652, 2021. 2
- [23] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International conference on machine learning*, pages 5156–5165. PMLR, 2020. 4
- [24] Justin Kerr, Chung Min Kim, Ken Goldberg, Angjoo Kanazawa, and Matthew Tancik. Lurf: Language embedded radiance fields. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 19729–19739, 2023. 3
- [25] Diederik P Kingma and Jimmy Ba. Adam: A method for

- stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. 6
- [26] 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. 3
- [27] Minhyeok Lee, Sangwon Hwang, Chaewon Park, and Sangyoun Lee. Edgeconv with attention module for monocular depth estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2858–2867, 2022. 2
- [28] Seokju Lee, Sunghoon Im, Stephen Lin, and In So Kweon. Learning monocular depth in dynamic scenes via instance-aware projection consistency. *arXiv preprint arXiv:2102.02629*, 2021. 2
- [29] Boyi Li, Kilian Q Weinberger, Serge Belongie, Vladlen Koltun, and Rene Ranftl. Language-driven semantic segmentation. In *International Conference on Learning Representations*, 2022. 3, 6
- [30] Rui Li, Xiantuo He, Yu Zhu, Xianjun Li, Jinqiu Sun, and Yanning Zhang. Enhancing self-supervised monocular depth estimation via incorporating robust constraints. In *Proceedings of the 28th ACM International Conference on Multimedia*, pages 3108–3117, 2020. 2
- [31] Rui Li, Dong Gong, Wei Yin, Hao Chen, Yu Zhu, Kaixuan Wang, Xiaozhi Chen, Jinqiu Sun, and Yanning Zhang. Learning to fuse monocular and multi-view cues for multi-frame depth estimation in dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 21539–21548, 2023. 1, 2
- [32] Rui Li, Danna Xue, Shaolin Su, Xiantuo He, Qing Mao, Yu Zhu, Jinqiu Sun, and Yanning Zhang. Learning depth via leveraging semantics: Self-supervised monocular depth estimation with both implicit and explicit semantic guidance. *Pattern Recognition*, page 109297, 2023. 2, 8
- [33] Jingyun Liang, Yuchen Fan, Kai Zhang, Radu Timofte, Luc Van Gool, and Rakesh Ranjan. Movidio: Motion-aware video generation with diffusion models. *arXiv preprint arXiv:2311.11325*, 2023. 1
- [34] Yiyi Liao, Jun Xie, and Andreas Geiger. Kitti-360: A novel dataset and benchmarks for urban scene understanding in 2d and 3d. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):3292–3310, 2022. 2, 6, 7
- [35] Ce Liu, Suryansh Kumar, Shuhang Gu, Radu Timofte, and Luc Van Gool. Single image depth prediction made better: A multivariate gaussian take. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17346–17356, 2023. 2
- [36] Ce Liu, Suryansh Kumar, Shuhang Gu, Radu Timofte, and Luc Van Gool. Va-depthnet: A variational approach to single image depth prediction. *arXiv preprint arXiv:2302.06556*, 2023. 2
- [37] Xiaoyang Lyu, Liang Liu, Mengmeng Wang, Xin Kong, Lina Liu, Yong Liu, Xinxin Chen, and Yi Yuan. Hr-depth: High resolution self-supervised monocular depth estimation. *arXiv preprint arXiv:2012.07356*, 2020. 2
- [38] Ruihang Miao, Weizhou Liu, Mingrui Chen, Zheng Gong, Weixin Xu, Chen Hu, and Shuchang Zhou. Occdepth: A depth-aware method for 3d semantic scene completion. *arXiv preprint arXiv:2302.13540*, 2023. 2
- [39] 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
- [40] Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. Automatic differentiation in pytorch. 2017. 6
- [41] Songyou Peng, Kyle Genova, Chiyu Jiang, Andrea Tagliasacchi, Marc Pollefeys, Thomas Funkhouser, et al. Openscene: 3d scene understanding with open vocabularies. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 815–824, 2023. 3
- [42] René Ranftl, Katrin Lasinger, David Hafner, Konrad Schindler, and Vladlen Koltun. Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer. *IEEE transactions on pattern analysis and machine intelligence*, 44(3):1623–1637, 2020. 2
- [43] René Ranftl, Alexey Bochkovskiy, and Vladlen Koltun. Vision transformers for dense prediction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 12179–12188, 2021. 2, 8
- [44] Aron Schmied, Tobias Fischer, Martin Danelljan, Marc Pollefeys, and Fisher Yu. R3d3: Dense 3d reconstruction of dynamic scenes from multiple cameras. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3216–3226, 2023. 2
- [45] Zhuoran Shen, Mingyuan Zhang, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Efficient attention: Attention with linear complexities. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pages 3531–3539, 2021. 4
- [46] Jaime Spencer, Chris Russell, Simon Hadfield, and Richard Bowden. Kick back & relax: Learning to reconstruct the world by watching slowtv. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15768–15779, 2023. 2
- [47] Libo Sun, Jia-Wang Bian, Huangying Zhan, Wei Yin, Ian Reid, and Chunhua Shen. Sc-depthv3: Robust self-supervised monocular depth estimation for dynamic scenes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023. 2
- [48] Zachary Teed and Jia Deng. Deepv2d: Video to depth with differentiable structure from motion. *arXiv preprint arXiv:1812.04605*, 2018. 2
- [49] 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. 2
- [50] Qianqian Wang, Zhicheng Wang, Kyle Genova, Pratul P Srinivasan, Howard Zhou, Jonathan T Barron, Ricardo

- Martin-Brualla, Noah Snavely, and Thomas Funkhouser. Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2021. 1
- [51] Felix Wimbauer, Nan Yang, Christian Rupprecht, and Daniel Cremers. Behind the scenes: Density fields for single view reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9076–9086, 2023. 1, 2, 5, 6, 7, 8
- [52] Ke Xian, Jianming Zhang, Oliver Wang, Long Mai, Zhe Lin, and Zhiguo Cao. Structure-guided ranking loss for single image depth prediction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 611–620, 2020. 2
- [53] Gangwei Xu, Junda Cheng, Peng Guo, and Xin Yang. Attention concatenation volume for accurate and efficient stereo matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 12981–12990, 2022. 2
- [54] Lior Yariv, Jiatao Gu, Yoni Kasten, and Yaron Lipman. Volume rendering of neural implicit surfaces. *Advances in Neural Information Processing Systems*, 34:4805–4815, 2021. 2
- [55] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. Enforcing geometric constraints of virtual normal for depth prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 5684–5693, 2019. 2
- [56] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Long Mai, Simon Chen, and Chunhua Shen. Learning to recover 3d scene shape from a single image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 204–213, 2021. 2
- [57] Wei Yin, Jianming Zhang, Oliver Wang, Simon Niklaus, Simon Chen, Yifan Liu, and Chunhua Shen. Towards accurate reconstruction of 3d scene shape from a single monocular image. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022.
- [58] Wei Yin, Chi Zhang, Hao Chen, Zhipeng Cai, Gang Yu, Kaixuan Wang, Xiaozhi Chen, and Chunhua Shen. Metric3d: Towards zero-shot metric 3d prediction from a single image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9043–9053, 2023. 1, 2
- [59] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa. pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2021. 1, 2, 5, 6, 7, 8
- [60] Zehao Yu, Songyou Peng, Michael Niemeyer, Torsten Sattler, and Andreas Geiger. Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction. *Advances in neural information processing systems*, 35:25018–25032, 2022. 2
- [61] Weihao Yuan, Xiaodong Gu, ZuoZhuo Dai, Siyu Zhu, and Ping Tan. New crfs: Neural window fully-connected crfs for monocular depth estimation. *arXiv preprint arXiv:2203.01502*, 2022. 2
- [62] Chaoqiang Zhao, Youmin Zhang, Matteo Poggi, Fabio Tosi, Xianda Guo, Zheng Zhu, Guan Huang, Yang Tang, and Stefano Mattocchia. Monovit: Self-supervised monocular depth estimation with a vision transformer. In *2022 International Conference on 3D Vision (3DV)*, pages 668–678. IEEE, 2022. 2, 5
- [63] Shuaifeng Zhi, Tristan Laidlow, Stefan Leutenegger, and Andrew J Davison. In-place scene labelling and understanding with implicit scene representation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15838–15847, 2021. 3
- [64] Hang Zhou, David Greenwood, and Sarah Taylor. Self-supervised monocular depth estimation with internal feature fusion. *arXiv preprint arXiv:2110.09482*, 2021. 2