

VoRTX: Volumetric 3D Reconstruction With Transformers for Voxelwise View Selection and Fusion

Noah Stier Alexander Rich Pradeep Sen Tobias Höllerer

{noahstier, anrich, psen, thollerer}@ucsb.edu

University of California, Santa Barbara

Abstract

Recent volumetric 3D reconstruction methods can produce very accurate results, with plausible geometry even for unobserved surfaces. However, they face an undesirable trade-off when it comes to multi-view fusion. They can fuse all available view information by global averaging, thus losing fine detail, or they can heuristically cluster views for local fusion, thus restricting their ability to consider all views jointly. Our key insight is that greater detail can be retained without restricting view diversity by learning a view-fusion function conditioned on camera pose and image content. We propose to learn this multi-view fusion using a transformer. To this end, we introduce VoRTX,¹ an end-to-end volumetric 3D reconstruction network using transformers for wide-baseline, multi-view feature fusion. Our model is occlusion-aware, leveraging the transformer architecture to predict an initial, projective scene geometry estimate. This estimate is used to avoid backprojecting image features through surfaces into occluded regions. We train our model on ScanNet and show that it produces better reconstructions than state-of-the-art methods. We also demonstrate generalization without any fine-tuning, outperforming the same state-of-the-art methods on two other datasets, TUM-RGBD and ICL-NUIM.

1. Introduction

3D reconstruction is a fundamental problem in computer vision, supporting applications such as autonomous navigation and mixed reality. In many scenarios, dense and highly detailed reconstruction is desirable. For example, it can facilitate the creation of virtual reality content by scanning real-world scenes, or the simulation of physics-based effects in augmented reality. Although active depth sensors have been employed for this purpose [6, 29], they increase platform cost relative to passive cameras. It is therefore desirable to perform reconstruction using only visible-light RGB

¹<https://noahstier.github.io/vortex>

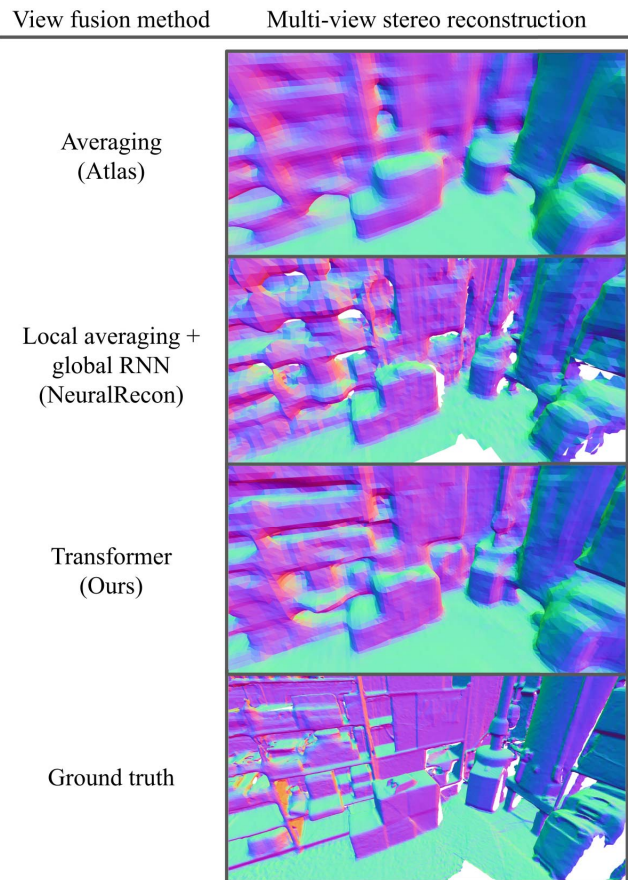


Figure 1. Our method fuses input view features using a transformer. We compare to Atlas [28], which fuses features by averaging, and NeuralRecon [37], which fuses locally by averaging and globally by RNN. Our method produces a high level of detail, while also filling in holes due to occlusion and unobserved regions.

cameras, which are ubiquitous and relatively inexpensive.

Dense 3D reconstruction from RGB imagery traditionally consists of estimating depth for each image, and then fusing the resulting depth maps in a reprojection step. This approach, however, cannot fill holes arising from occlusions and other unobserved regions.

Recently, a number of works have addressed this by posing RGB-only 3D reconstruction as the direct prediction of a truncated signed-distance function (TSDF), using deep learning to fill in unobserved regions via learned priors [28, 37]. These methods extract image features using a convolutional neural network (CNN), accumulate them into space by backprojecting onto a 3D grid, and then predict the TSDF volume using a 3D CNN. When a particular grid voxel is within the view frustum of multiple cameras, it is common practice to fuse the backprojected image features at that point via unweighted averaging. However, we observe two drawbacks of directly averaging view features.

First, when images are acquired from very different camera poses, their content may not be directly comparable. Although CNNs are capable of extracting high-level semantic features that are therefore highly view-independent, the CNN architectures commonly used in 3D reconstruction (e.g., U-Net [30] and FPN [21]) make explicit use of the activations from early CNN layers. These are understood to represent lower-level visual features, which exhibit view-independence only within a particular range of viewpoint difference. Averaging across disparate views does not take that range into consideration, and therefore loses view-dependent information. This is a known phenomenon in multi-view stereo, where typical solutions include 1) selecting views using constraints on camera pose to minimize viewpoint differences [12, 32], or 2) constraining the image features to be as view-independent as possible [40]. We hypothesize that a better solution can be obtained by learning a view fusion function, conditioned on pose and image content, that can jointly consider features from multiple views within the appropriate range of viewpoints.

Second, averaging assigns an equal weight to all input views at each voxel, including views for which a voxel is occluded. This problem is exacerbated in wide-baseline reconstruction, where occlusions are particularly prevalent. Occlusion modeling presents a chicken-and-egg problem: the scene geometry is not known until after backprojection and reconstruction; but until the scene geometry is known, backprojection cannot account for occlusions, thus projecting image features through surfaces into regions where they are irrelevant. We hypothesize that this irrelevant information acts as noise that reduces reconstruction quality.

We propose an innovation that addresses both issues. Our model, which we call VoRTX, is a deep learning-based volumetric reconstruction network using transformers [42] to model dependencies across diverse viewpoints. The transformers use self-attention to perform soft grouping of views that are mutually relevant, and they can learn to fuse within vs. across groups in different feature spaces.

Transformers also provide a natural mechanism for occlusion-awareness, since the attention to each input view varies as a function of 3D location. The view aggregation

can therefore be supervised to encourage reduced attention to input images in regions where their view is occluded. One possibility is to supervise the view aggregation using ground-truth visibility. However, we argue in Sec. 3.3 that *projective occupancy* is preferable for our problem setting, because it is an easier target that more closely describes the desired spatial distribution of image features during back-projection. Our main contributions are as follows:

1. We introduce a new method of fusing multi-view image features, using a transformer to perform data-dependent fusion at each spatial location.
2. We propose the projective occupancy as an occlusion-aware reconstruction target for deep volumetric MVS, and we show that it yields improved results over unsupervised or visibility-supervised reconstruction.

We show that VoRTX surpasses state-of-the-art reconstruction results when compared with several baseline methods, on multiple datasets.

2. Related work

Image feature fusion in MVS: Fusing measurements from multiple views is a crucial step in MVS. Typically, image patches are fused into a cost volume using a stereo-matching cost function, which operates on raw image intensity [12, 14, 32, 43] or CNN-extracted image features [10, 49]. Some methods [15, 16] instead concatenate image features in the channel dimension, and use a CNN to reduce them into a cost volume. These techniques are effective when the input views are acquired closely enough in pose space to maintain similar scene appearance, while still providing enough parallax for stereopsis.

Atlas [28] proposes the use of a single feature volume, bypassing depth prediction and posing 3D reconstruction as the direct prediction of a TSDF volume. This is an effective way to consider all input images jointly, and it also provides a framework for learning to reconstruct unobserved scene regions via 3D priors. However, Atlas fuses input image features by direct averaging, which does not effectively model view-dependent image features or occlusion effects.

PIFu [31] also performs multi-view fusion by averaging backprojected features, showing strong results for reconstruction of free-standing humans. However, to our knowledge, it has not been demonstrated for full, real-world scenes, which tend to introduce more complex occlusion relationships as well as semantic and geometric variety.

NeuralRecon [37] averages features only among nearby views, fusing across view clusters using a recurrent neural network (RNN). NeuralRecon achieves real-time execution, with the trade-off that incoming views must be considered sequentially. Our model lifts the constraint of sequential processing, fusing all available views jointly.

Point-MVSNet [4] replaces the feature volume entirely with a feature-augmented point cloud, aggregating view features with a point cloud CNN architecture based on EdgeConv [44]. This is a promising approach, although point cloud learning is not as mature as regular-grid CNNs.

Occlusion-aware MVS: Occlusion detection with explicit photometric and geometric constraints has traditionally played an important role in MVS [19, 32, 34, 35, 36, 47, 54, 56]. In addition, a number of MVS methods based on deep learning have proposed to learn visibility estimation [3, 17, 18, 24].

Direct scene optimization: Yariv et al. [50] propose to directly optimize the scene representation with respect to the input images. This is effective when the target geometry is fully observed. However, it has no offline training phase in which 3D priors can be learned and then applied to new reconstructions. This prevents any significant scene completion, which is a key feature of our algorithm.

Projective TSDF: In RGBD reconstruction, the projective TSDF is used as a means of approximating the true, or view-independent TSDF, by averaging together the projective TSDFs of many depth images [29]. It has been used as a powerful representation in its own right, as way to encode individual depth images for processing by 3D CNNs [13, 33]. It has also been used a reconstruction target for 3D reconstruction from single-view RGB images [20]. In our formulation, a projective TSDF prediction acts an initial approximation of the surface geometry, which allows us to model occlusion during backprojection.

Multi-view fusion with attention: For single-object reconstruction, attention has been used to fuse multiple images into a fixed-size global scene encoding [45, 46, 53]. MVS algorithms have leveraged channel-wise attention to focus on relevant feature subspaces [26], 2D image-space attention to aggregate visual context [48, 51], and 3D attention to promote coherence across cost volumes [23]. A recent method for novel-view synthesis [41] has experimented with two attention mechanisms for fusing backprojected image features: AttSets [46] and Slot Attention [22]. In our experiments, these variants do not perform as well as the transformer-based attention (see Table 4 for results and section 4.3 for discussion).

Transformers: Transformers are a family of neural network architectures that have proven very effective for sequence modeling in natural language processing [8, 42], as well as vision [9]. They are neither biased toward modeling short-range dependencies, like CNNs, nor restricted to sequential processing, like RNNs. Instead, they achieve a global receptive field by composing self-attention layers. The appeal of transformers for multi-view fusion arises from their ability to perform soft clustering of their inputs. This makes transformers a good fit for wide-baseline view fusion, which benefits from clustering views, and fusing

within vs. across clusters in different feature subspaces.

In work submitted concurrently with ours, Aljaž et al. [21] propose 3D reconstruction with transformers for multi-view fusion. Notably, their work further utilizes the attention weights for frame selection, to ensure that all relevant view information is considered. Our work on modeling projective occupancy is fundamentally aimed at reducing the irrelevant information, and we therefore hypothesize that these approaches may provide complimentary benefits.

3. Method

Our goal is to predict a global TSDF volume \hat{S} , using an unordered sequence of input RGB images and their corresponding 6-DOF camera poses. For training, we assume the existence of ground-truth depth maps.

In broad strokes, our model extracts image features with a 2D CNN, backprojects them into a voxel grid, and predicts a TSDF with a 3D CNN. It thus bears structural similarity to existing deep volumetric reconstruction methods [28, 37]. Sec. 3.1 introduces the architecture overview and notation.

The first key difference from existing work is in the image feature backprojection and aggregation phase. We introduce a transformer to process single-view image features, selectively fusing them into a multi-view encoding *before* aggregating per-voxel features. This significantly expands the model’s ability to reason jointly about the input views, improving the localization of surfaces in its reconstructions. Details are presented in Sec. 3.2.

Our second main contribution is to weight the final feature aggregation with explicitly-supervised projective occupancy predictions, enforcing that image features are only accumulated into regions near their observed surfaces. Sec. 3.3 expands on this component.

3.1. Overview

The overall structure of our algorithm is illustrated in Fig. 2. A 2D CNN (a feature pyramid network [21] with an MnasNet [38] backbone) begins by extracting image features at coarse, medium, and fine resolutions:

$$\{F_I^{(c)}, F_I^{(m)}, F_I^{(f)}\} = g_\theta(I), \quad (1)$$

where g_θ is the CNN parametrized by network weights θ .

At each resolution $r \in \{c, m, f\}$, the image features are backprojected onto a sparse 3D grid. This produces a feature volume, $F_{BP}^{(r)}$, in which each voxel contains a set of backprojected features, one from each image. The per-voxel features are then aggregated using our transformer and projective occupancy architecture to form a new volume, $F_{MV}^{(r)}$, containing one multi-view feature in each voxel. A sparse 3D CNN [39] processes $F_{MV}^{(r)}$, predicting occupancy $\hat{O}^{(r)}$:

$$\hat{O}^{(r)} = h_\theta^{(r)}(F_{MV}^{(r)}), \quad r \in \{c, m, f\} \quad (2)$$

where $h_\theta^{(r)}$ represents the 3D CNN at resolution r .

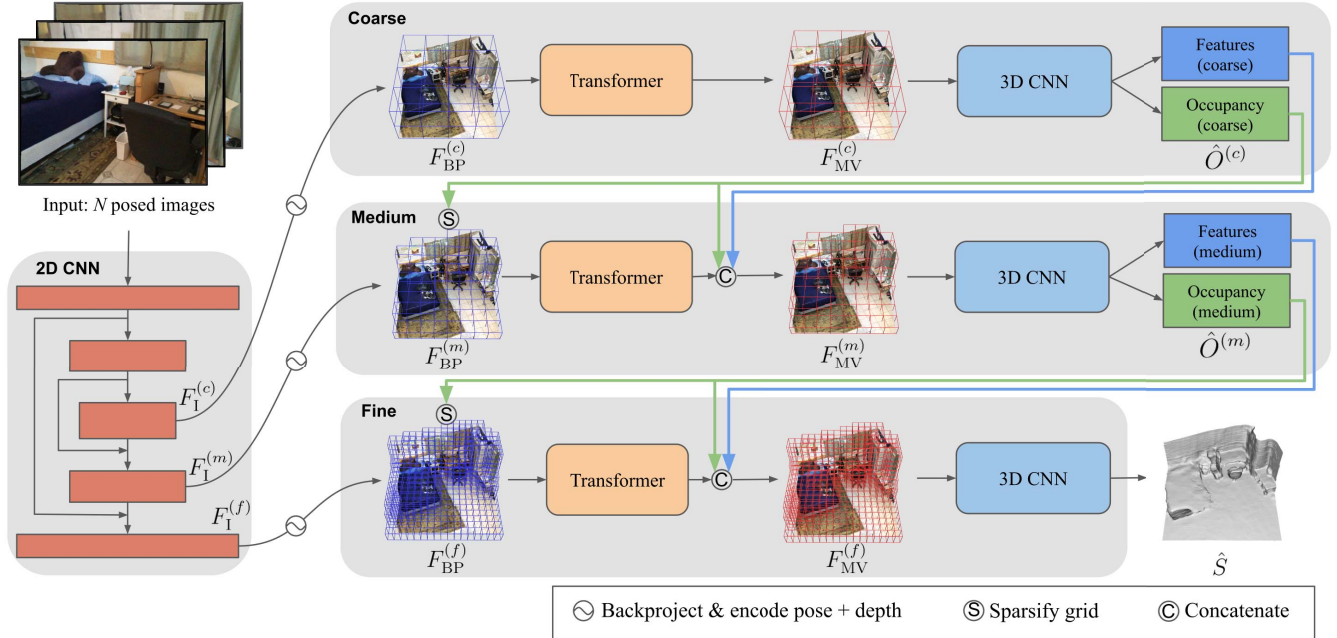


Figure 2. Model overview. A 2D CNN processes N input images to produce image features at coarse, medium, and fine resolutions: $F_I^{(r)} \in \mathbb{R}^{N \times H^{(r)} \times W^{(r)} \times C^{(r)}}$, $r \in \{c, m, f\}$. At each resolution, a sparse feature volume with $V^{(r)}$ voxels is computed by backprojection, and the camera-to-voxel unit vector and depth are jointly encoded: $F_{BP}^{(r)} \in \mathbb{R}^{V^{(r)} \times N \times C^{(r)}}$. A transformer fuses image features at each voxel to produce the multi-view feature volume, $F_{MV}^{(r)} \in \mathbb{R}^{V^{(r)} \times C^{(r)}}$. At the coarse and medium resolutions, a sparse 3D CNN predicts occupancy $\hat{O}^{(r)} \in \mathbb{R}^{V^{(r)}}$ which is used to sparsify the volume. At the fine resolution, the 3D CNN predicts the final TSDF \hat{S} .

At each resolution, any voxels predicted to be unoccupied are pruned from the next, higher-resolution hierarchy level, in a coarse-to-fine manner. At the final, highest-resolution level, the TSDF \hat{S} is predicted instead of occupancy, and the zero isosurface is extracted using marching cubes [25]. We set the voxel size at each resolution to 16cm^3 , 8cm^3 , and 4cm^3 , respectively.

In order to scale from local to full-scene reconstruction, we tile the target space with a set of non-overlapping local volumes. Then, for each tile we aim to select a diverse set of N views from across the input sequence (see Sec. 3.4). Starting with the coarsest resolution, we populate $F_{BP}^{(r)}$ and $F_{MV}^{(r)}$ tile by tile. Then, we run sparse 3D convolution globally, and proceed to backprojection at the next resolution.

3.2. Multi-view image feature fusion

Our key innovation is to use a transformer to augment each backprojected single-view feature with information from other relevant views. At each voxel, the transformer takes an unordered sequence of single-view feature vectors as input, and produces a corresponding sequence of multi-view feature vectors as output:

$$\tilde{F}_{MV}^{(r)} = y_{\theta}^{(r)}(F_{BP}^{(r)}), \quad (3)$$

where $y_{\theta}^{(r)}$ represents the transformer at resolution r .

We use the tilde to indicate that $\tilde{F}_{MV}^{(r)}$ is the predecessor to $F_{MV}^{(r)}$: each voxel in $\tilde{F}_{MV}^{(r)}$ contains a sequence of multi-view features, and $F_{MV}^{(r)}$ is the result of the per-voxel feature aggregation detailed in the following section.

The correspondence between the i^{th} sequence element of $F_{BP}^{(r)}$ and $\tilde{F}_{MV}^{(r)}$ is encouraged by residual connections across attention layers, and it is enforced by predicting the projective occupancy for each input view using its corresponding element of the output sequence.

Generally, the input to a transformer is an unordered sequence of feature vectors, where each feature vector is a joint encoding of the original sequence element and its position in the sequence. In our model, we replace the typical sequential positional encoding with a camera pose encoding, $\Lambda(d)$, where d is the camera-to-voxel view direction unit vector and Λ is the positional encoding from Mildenhall et al. [27]. To form the transformer input, we concatenate the image feature and the pose encoding, and reduce the resulting dimensionality with a shared fully-connected (FC) layer. We then concatenate the normalized camera-to-voxel depth and reduce with a second FC layer before applying the transformer.

Our transformer, shown in Fig. 3, is based on the encoder part of the original transformer network introduced by Vaswani et al. [42]. It consists of a series of L layers, where each layer contains a multi-head attention mechanism with




Figure 3. Our transformer architecture in detail. At each individual voxel, the transformer input features $F_{BP}^{(r)}$ from N images each have channel dimension C . The transformer layer is repeated L times, selectively fusing the inputs to produce a set of multi-view features $\tilde{F}_{MV}^{(r)}$ with the same dimensions as the input. A fully-connected layer predicts projective occupancy probabilities $\hat{O}_p^{(r)}$, which are used as weights in a final channel-wise average to produce $F_{MV}^{(r)}$.

H heads, followed by a small fully-connected network. We also employ residual connections and layer normalization within each layer. In our implementation we set L and H both equal to 2.

The following section describes the aggregation of the transformer output sequence into a single per-voxel feature vector, which is subsequently passed on to the 3D CNN.

3.3. Projective occupancy

Our problem context violates key assumptions that MVS methods traditionally make, and this inspires us to re-think the notion of visibility.

Specifically, because we aim to learn view selection and fusion, we do not impose any constraints on the relative pose of the views to be fused, instead sampling broadly from across the image sequence. This results in high perspective diversity, with triangulation angles often greater than 90 degrees. This violates the typical assumptions of fronto-parallel scene structure and small baseline distance.

We therefore reconsider the notion of visibility in our context. Our goal is to place image features into 3D space such that they enable a 3D CNN to estimate the imaged surface location. If we spatially distribute those features along a camera ray according to the estimated projective occupancy, their spatial density will be centered at the estimated target surface depth. This is intuitively favorable from the perspective of the 3D CNN. In contrast, if the features are spatially distributed according to visibility, then their spatial density is spread across observed empty space, and it may not reach the true surface location if the depth is underesti-




Figure 4. Comparison of projective distance functions, where t is the TSDF truncation distance. Visibility is a function of the sign of the TSDF, and it includes observed empty space. Projective occupancy is a function of TSDF magnitude, and it describes the surface location within a margin of error.

mated. See Fig. 4 for an illustration. We therefore consider the projective occupancy to be a more effective prediction target for our purposes.

Furthermore, we hypothesize that it is an easier target. Fundamentally, projective occupancy requires predicting the magnitude of the TSDF, whereas visibility requires predicting the sign of the TSDF. In theory, estimating the magnitude of the TSDF at a point using two image projections can be done with only a matching cost function. However, estimating the sign of the TSDF requires understanding the direction of mismatch, and comparing it to the relative camera poses. We therefore consider the visibility to be a more difficult target, and this may contribute to the performance decrease observed in our ablation study (Table 4, row g).

To introduce our projective occupancy prediction framework, we first define the projective SDF S ,

$$S = d - d_v, \quad (4)$$

where d_v is the camera-to-voxel depth, and d is the true depth along the camera-voxel ray. We estimate d in practice by projecting onto the ground truth depth map and sampling the depth at the nearest-neighbor pixel.

The projective occupancy O_p can then be obtained by thresholding the absolute value of S on the truncation distance t :

$$O_p = \begin{cases} 1 & |S| < t \\ 0 & |S| \geq t \end{cases}. \quad (5)$$

Our model estimates the projective occupancy likelihood $X \in \mathbb{R}^N$ as

$$X = z_\theta^{(r)}(\tilde{F}_{MV}^{(r)}), \quad (6)$$

where $z_\theta^{(r)}$ is a single, shared, FC layer at resolution r . In order to supervise X , a sigmoid is applied to produce the projective occupancy probabilities:

$$\hat{O}_p^{(r)} = \sigma(X) \quad (7)$$

Then a loss is computed as binary cross-entropy between $\hat{O}_p^{(r)}$ and the groundtruth projective occupancy.

In order to use X to inform feature aggregation, we concatenate a zero-likelihood to X and apply a softmax to compute a weight vector $W \in \mathbb{R}^{1 \times N+1}$. We then concatenate a zero feature vector to $\tilde{F}_{MV}^{(r)}$, resulting in dimensions

$(N + 1) \times C$, and reduce with a weighted sum:

$$F_{MV}^{(r)} = W \tilde{F}_{MV}^{(r)} \quad (8)$$

The softmax weight normalization ensures that the distribution of $F_{MV}^{(r)}$ is invariant to the number of input views. The zero-padding of both features and likelihoods causes $F_{MV}^{(r)}$ to be near zero if all the predicted occupancy likelihoods are low.

3.4. View selection

Our method does not depend on heuristics to select optimally positioned input views. Conversely, we aim to train our model on an unconstrained set of views that is as diverse as possible while remaining computationally tractable, such that it can learn to fuse features across the appropriate range of pose differences. We employ heuristics only to reduce the overall number of views while maintaining diversity.

To this end, we first remove redundant views by applying the keyframe selection strategy from Sun et al. [37]. Then, for each local sub-volume, we select N views via uniform random sampling from among the remaining views whose camera frustums intersect the target volume. During training we set $N = 20$, and during testing we set $N = 60$. For redundant frame removal, we set R_{max} to 15 degrees, and we set t_{max} to 0.1 m for training and 0.2 m for testing.

3.5. Training

Loss function: The projective occupancy loss $\lambda_P^{(r)}$ at each hierarchy level, and the occupancy loss at the coarser levels $\lambda_O^{(r)}$, are computed using binary cross-entropy. The TSDF loss at the finest level, λ_S , is computed by $l1$ distance to the ground truth, after log-transforming the prediction and ground truth following [7]. Then the total loss L is

$$L = \lambda_P^{(c)} + \lambda_P^{(m)} + \lambda_P^{(f)} + \lambda_O^{(c)} + \lambda_O^{(m)} + \lambda_S$$

Ground truth: We compute our fine-resolution reconstruction target using TSDF fusion at 4 cm resolution, discarding all measurements greater than 3 m due to sensor noise at longer ranges. We then threshold that TSDF on the truncation distance to obtain a fine-resolution occupancy volume, which we downsample by morphological dilation to produce the medium and coarse reconstruction targets. As in Murez et al. [28], we mark any column of the ground truth TSDF volume as unoccupied if it is entirely unobserved.

During training we select sub-volumes by randomly selecting TSDF subcrops with size $96 \times 96 \times 48$ voxels, or $3.84 \text{ m} \times 3.84 \text{ m} \times 1.92 \text{ m}$. We augment with random horizontal reflections and rotations about the gravitational axis.

Training phases: During our initial training phase, the projective occupancy predictions are supervised, but they are not otherwise used: the transformer output sequence is aggregated with an unweighted average. This aids stability.

Also during this phase, the 2D CNN weights, which are pre-trained on ImageNet, are frozen. The learning rate is 10^{-3} , the batch size is 4, and this phase lasts 300 epochs.

In the second phase, the projective occupancy predictions are used for weighted-average aggregation of the transformer outputs, as shown in Fig. 3. In addition, the 2D CNN weights are unfrozen, except for the batch norm weights and statistics. The learning rate is lowered to 10^{-4} , the batch size is lowered to 2. This phase lasts 100 epochs.

Implementation details: We use the Adam optimizer with $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 10^{-8}$, and a linear learning rate warm-up from 0 over 2,000 steps. Training takes approximately 84 hours on a single Nvidia RTX 3090 graphics card. We implement our model in PyTorch, using the PyTorch Lightning framework [11]. We use torchsparse [39] for our sparse 3D CNN, and Open3D [55] for visualization and geometry processing. During training, we randomly drop out voxels to reduce memory cost, following [37].

4. Experiments

For all experiments, we train our method on the ScanNet dataset [5]: 1,513 RGBD scans of 707 indoor spaces. We use the official train/validation/test split.

For quantitative comparison, we compute a set of 3D metrics as defined by Murez et al. [28]. To avoid penalizing the volumetric methods for filling in areas that are not present in the ground truth, we trim the reconstructed mesh to within the observed regions. To do this, we render the ground-truth mesh to a set of depth maps D from the perspective of each camera pose. Then we render the predicted mesh to a set of depth maps \hat{D} . We mask out pixels in \hat{D} that do not have a valid depth in D , and re-fuse the masked predicted depth into a trimmed mesh via TSDF fusion.

4.1. Volumetric baselines

Our primary comparison is with algorithms that, like ours, can complete geometry in unobserved regions. These are the deep volumetric methods, Atlas [28] and NeuralRe-

	Acc ↓	Comp ↓	Prec ↑	Recall ↑	F-score ↑
SCANNET					
Atlas	0.068	0.098	0.640	0.539	0.583
NeuralRecon	0.049	0.133	0.691	0.461	0.551
Ours	0.054	0.090	0.708	0.588	0.641
ICL-NUIM					
Atlas	0.175	0.314	0.280	0.194	0.229
NeuralRecon	0.215	1.031	0.214	0.036	0.058
Ours	0.102	0.146	0.449	0.375	0.408
TUM-RGBD					
Atlas	0.208	2.344	0.360	0.089	0.132
NeuralRecon	0.130	2.528	0.382	0.075	0.115
Ours	0.175	0.314	0.280	0.194	0.229

Table 1. Reconstruction metrics (as defined as in [28]), comparison with volumetric methods.

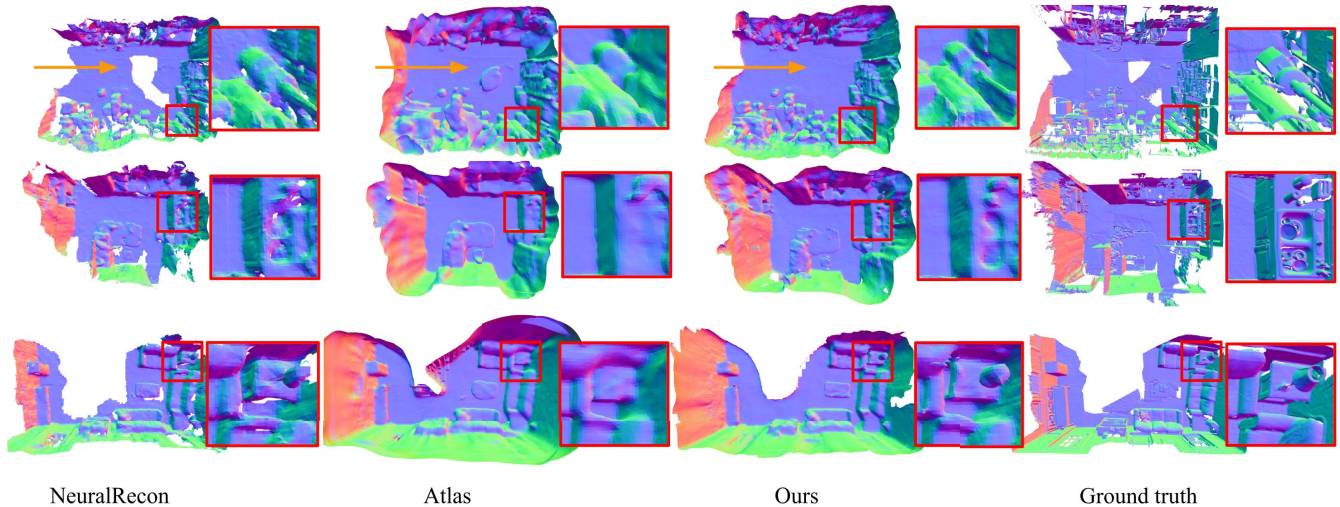


Figure 5. Qualitative results on ScanNet. The inset boxes show enlarged regions where our model reconstructs a high degree of detail. With the orange arrows, we highlight another strength of our model: it fills in unobserved regions plausibly, without leaving holes or artifacts.

con [37], and we use the provided pre-trained models. For Atlas, we select every 5th frame as input, and for NeuralRecon we use the frame selection proposed by its authors. We evaluate on the ScanNet test set (100 scenes), the ICL-NUIM dataset (8 scenes), and the TUM-RGBD dataset (13 scenes). For ScanNet, we evaluate against the provided ground-truth meshes. For TUM-RGBD and ICL-NUIM, we generate ground truth by TSDF fusion at 4 cm resolution.

Quantitative results are shown in Table 1. We consider F-score to be the most important metric, as it captures the trade-off between precision and recall. Our F-score indicates a significant improvement over state-of-the-art methods. We also report the accuracy of our projective occupancy predictions at each resolution in Table 2, and we compare against the default prediction of *true* everywhere.

Qualitative results are shown in Fig. 5. We observe increased accuracy relative to the baseline methods, particularly in areas with many small objects and a high degree of occlusion, such as cluttered countertops. In these regions, our model produces a high level of detail while also filling in holes arising from occlusion (Fig. 5, rows 1 and 2). We note that in large unobserved regions (Fig. 5, row 3), our model’s performance degrades gracefully: whereas Atlas tends to incorrectly place walls at the boundary, and NeuralRecon typically does not produce any geometry, VoRTX extends observed surfaces for a plausible distance without introducing large artifacts.

We also observe that in many cases, even when reconstruction quality is visually similar, our model localizes surfaces more accurately, as shown in Fig. 6.

4.2. Depth-prediction baselines

For completeness, we compare with deep MVS networks that estimate depth maps, reconstructing only observed surfaces: DeepVideoMVS (with fusion) [10], Fast-MVSNet

[52], GPMVS (batched) [14], and Point-MVSNet [4]. For DeepVideoMVS, we use the ScanNet pre-trained weights. For Fast-MVSNet, GPMVS, and Point-MVSNet, we fine-tune on ScanNet, starting from the pre-trained models. For Point-MVSNet and Fast-MVSNet, we modify the parameters for the longer ranges in ScanNet relative to DTU [1]: we use 96 depth hypotheses, every 5 cm starting at 50 cm. We fuse predicted depths into point clouds following [12]. For all depth-prediction methods, we select views following Duzceker et al. [10], using four source images for each reference image. As shown in Table 3, VoRTX produces higher F-scores, indicating that it does not compromise on observed surfaces in order to complete unobserved regions.

4.3. Ablation experiments

In Table 4 we present ablation experiments to validate our model. In each, the model architecture is modified and

Hierarchy Lvl.	Proj. Occ. Prediction	Prec \uparrow	Recall \uparrow	Acc \uparrow
4cm	Default (true everywhere)	0.237	1.000	0.237
	Ours	0.702	0.347	0.813
8cm	Default (true everywhere)	0.301	1.000	0.301
	Ours	0.750	0.627	0.829
16cm	Default (true everywhere)	0.067	1.000	0.067
	Ours	0.739	0.661	0.961

Table 2. Projective occupancy results. The default behavior is to assume projective occupancy is *true* for all voxels.

	Acc \downarrow	Comp \downarrow	Prec \uparrow	Recall \uparrow	F-score \uparrow
DeepVideoMVS	0.079	0.133	0.521	0.454	0.474
Fast-MVSNet	0.042	0.225	0.746	0.383	0.495
GPMVS	0.066	0.117	0.591	0.513	0.539
Point-MVSNet	0.037	0.278	0.790	0.363	0.484
Ours	0.054	0.090	0.708	0.588	0.641

Table 3. ScanNet reconstruction metrics (as defined as in [28]), comparison with depth-prediction methods.




Figure 6. Trimmed mesh predictions (see Sec. 4). Top: shaded blue for predicted vertices p^* within $5cm$ of a ground-truth vertex p , red otherwise. Bottom: shaded by surface normal. Our results show improved accuracy, even in cases with similar visual quality.

re-trained from scratch. Row a is VoRTX, unmodified.

Transformer: We first experiment with removing the transformer entirely (row c). In this case, projective occupancy predictions are made on the basis of the single-view features, aggregating by weighted average. This causes a significant drop in F-score. We also experiment with removing both transformer and projective occupancy (d), aggregating within voxels by unweighted average. This causes a further F-score drop. We conclude that the transformer is responsible for most of VoRTX’s performance gain.

In f we alter the hyperparameters of the transformer, using only a single layer and a single attention head, resulting in a moderate F-score decrease. We thus hypothesize that additional layers may lead to further performance gains.

In h and i , we replace the transformer with alternative attention mechanisms, following GRF [41]. The projective occupancy is predicted using single-view features. In h , the AttSets [46] model shows a moderate F-score decrease. This may be due to the fact that AttSets has only one attention layer, or that it doesn’t model pairwise attention between views. In i , using Slot Attention [22], our model does not converge well during training, and further investigation may be required to fully characterize the technique.

Projective occupancy: We also experiment with removing the projective occupancy prediction while keeping the transformer, aggregating the transformer outputs by direct averaging (b). In g , we keep the same architecture, but we supervise $\hat{O}_p^{(r)}$ with the visibility instead of projective occu-

	Transf.	Proj.Occ.	Pose	Acc ↓	Comp ↓	Prec ↑	Rec ↑	F-score ↑
a	✓	✓	✓	0.054	0.090	0.708	0.588	0.641
b	✓		✓	0.058	0.090	0.681	0.579	0.624
c		✓	✓	0.067	0.110	0.626	0.510	0.560
d			✓	0.071	0.125	0.611	0.487	0.540
e	✓	✓		0.053	0.091	0.701	0.579	0.633
f	L=1, H=1	✓	✓	0.057	0.090	0.684	0.572	0.622
g	✓	Vis.	✓	0.057	0.089	0.677	0.562	0.613
h	AttSets	✓	✓	0.057	0.098	0.680	0.563	0.614
i	Slot Attn.	✓	✓	0.075	0.210	0.546	0.346	0.420

Table 4. Ablation experiments on ScanNet.

pancy. In both cases we see a small performance decrease, supporting our hypotheses that the model benefits from supervising the aggregation weights, and that projective occupancy is a more effective weighting function than visibility.

Pose: In e , the model does not encode pose information into the image features during backprojection (it does still encode camera-to-voxel depth). This results in only a very slight performance decrease. We interpret this to suggest that although the viewing direction is useful information, most of its benefit can be obtained with attention-based comparison of pose-agnostic image features.

4.4. Inference time

Our method achieves speeds compatible with interactive applications on commodity hardware. We benchmark VoRTX on the ScanNet test set, using an AMD Threadripper 2950X and an NVIDIA RTX 3090. It averages 14.2 FPS, counting only selected keyframes.

5. Limitations

Because VoRTX uses a voxel representation, it is subject to a trade-off between resolution and memory use. We use 4 cm voxels, which are acceptable for indoor scenes but can cause aliasing for thin structures. In addition, reflective surfaces are often missing from our reconstructions. We believe this is partially due to the failure of the depth sensors for those surfaces, leading to gaps in supervision.

6. Conclusion

We have presented a novel method for multi-view fusion using transformers, applied toward deep volumetric MVS. We show that this produces better reconstructions than state-of-the-art methods on ScanNet, TUM-RGBD, and ICL-NUIM. Our model is trained only on ScanNet, generalizing well to the two other datasets without fine-tuning. Our projective occupancy framework opens the door to occlusion-awareness for deep volumetric MVS.

In the future, a focus on thin structures and reflective surfaces could yield improvements. Use of simulated training data, or alternative depth sensors, may facilitate learning and open possibilities for new data domains. Further attention to scalability may be beneficial for transferring to large-scale reconstructions. Finally, we anticipate that the transformer-based view fusion may also be applicable to tasks such as fusing multiple sensing modalities.

7. Acknowledgements

Support for this work was provided by ONR grants N00014-19-1-2553 and N00174-19-1-0024, as well as NSF grants 1911230 and OAC-1925717.

References

- [1] Henrik Aanæs, Rasmus Ramsbøl Jensen, George Vogiatzis, Engin Tola, and Anders Bjarholm Dahl. Large-scale data for multiple-view stereopsis. *International Journal of Computer Vision*, 120(2):153–168, 2016. 7
- [2] Aljaž Božič, Pablo Palafox, Justus Thies, Angela Dai, and Matthias Nießner. TransformerFusion: Monocular RGB scene reconstruction using transformers. *Proc. Neural Information Processing Systems (NeurIPS)*, 2021. 3
- [3] Rui Chen, Songfang Han, Jing Xu, et al. Visibility-aware point-based multi-view stereo network. *IEEE transactions on pattern analysis and machine intelligence*, 2020. 3
- [4] Rui Chen, Songfang Han, Jing Xu, and Hao Su. Point-based multi-view stereo network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1538–1547, 2019. 3, 7
- [5] 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. 6
- [6] Angela Dai, Daniel Ritchie, Martin Bokeloh, Scott Reed, Jürgen Sturm, and Matthias Nießner. ScanComplete: Large-scale scene completion and semantic segmentation for 3D scans. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4578–4587, 2018. 1
- [7] Angela Dai, Charles Ruizhongtai Qi, and Matthias Nießner. Shape completion using 3d-encoder-predictor cnns and shape synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5868–5877, 2017. 6
- [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
- [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. 3
- [10] Arda Duzceker, Silvano Galliani, Christoph Vogel, Pablo Speciale, Mihai Dusmanu, and Marc Pollefeys. Deep-VideoMVS: Multi-view stereo on video with recurrent spatio-temporal fusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15324–15333, 2021. 2, 7
- [11] WA Falcon and et al. PyTorch lightning. *GitHub. Note: <https://github.com/PyTorchLightning/pytorch-lightning>*, 3, 2019. 6
- [12] Silvano Galliani, Katrin Lasinger, and Konrad Schindler. Massively parallel multiview stereopsis by surface normal diffusion. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 873–881, 2015. 2, 7
- [13] Lihao Ge, Hui Liang, Junsong Yuan, and Daniel Thalmann. 3D convolutional neural networks for efficient and robust hand pose estimation from single depth images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1991–2000, 2017. 3
- [14] Yuxin Hou, Juho Kannala, and Arno Solin. Multi-view stereo by temporal nonparametric fusion. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2651–2660, 2019. 2, 7
- [15] Po-Han Huang, Kevin Matzen, Johannes Kopf, Narendra Ahuja, and Jia-Bin Huang. DeepMVS: Learning multi-view stereopsis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2821–2830, 2018. 2
- [16] Sunghoon Im, Hae-Gon Jeon, Stephen Lin, and In So Kweon. DPSNet: End-to-end deep plane sweep stereo. *arXiv preprint arXiv:1905.00538*, 2019. 2
- [17] Mengqi Ji, Juergen Gall, Haitian Zheng, Yebin Liu, and Lu Fang. SurfaceNet: An end-to-end 3D neural network for multiview stereopsis. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2307–2315, 2017. 3
- [18] Mengqi Ji, Jinzhi Zhang, Qionghai Dai, and Lu Fang. SurfaceNet+: An end-to-end 3D neural network for very sparse multi-view stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020. 3
- [19] Sing Bing Kang, Richard Szeliski, and Jinxiang Chai. Handling occlusions in dense multi-view stereo. In *Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. CVPR 2001*, volume 1, pages I–I. IEEE, 2001. 3
- [20] Hanjun Kim, Jiyoun Moon, and Beomhee Lee. RGB-to-TSDF: Direct TSDF prediction from a single RGB image for dense 3D reconstruction. In *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 6714–6720. IEEE, 2019. 3
- [21] 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. 2, 3
- [22] Francesco Locatello, Dirk Weissenborn, Thomas Unterthiner, Aravindh Mahendran, Georg Heigold, Jakob Uszkoreit, Alexey Dosovitskiy, and Thomas Kipf. Object-centric learning with slot attention. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 11525–11538. Curran Associates, Inc., 2020. 3, 8
- [23] Xiaoxiao Long, Lingjie Liu, Wei Li, Christian Theobalt, and Wenping Wang. Multi-view depth estimation using epipolar spatio-temporal networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8258–8267, 2021. 3
- [24] Xiaoxiao Long, Lingjie Liu, Christian Theobalt, and Wenping Wang. Occlusion-aware depth estimation with adaptive normal constraints. In *European Conference on Computer Vision*, pages 640–657. Springer, 2020. 3

- [25] William E Lorensen and Harvey E Cline. Marching cubes: A high resolution 3d surface construction algorithm. *ACM siggraph computer graphics*, 21(4):163–169, 1987. 4
- [26] Keyang Luo, Tao Guan, Lili Ju, Yuesong Wang, Zhuo Chen, and Yawei Luo. Attention-aware multi-view stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1590–1599, 2020. 3
- [27] 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 *European conference on computer vision*, pages 405–421. Springer, 2020. 4
- [28] Zak Murez, Tarrence van As, James Bartolozzi, Ayan Sinha, Vijay Badrinarayanan, and Andrew Rabinovich. Atlas: End-to-end 3D scene reconstruction from posed images. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VII 16*, pages 414–431. Springer, 2020. 1, 2, 3, 6, 7
- [29] Richard A Newcombe, Shahram Izadi, Otmar Hilliges, David Molyneaux, David Kim, Andrew J Davison, Pushmeet Kohi, Jamie Shotton, Steve Hodges, and Andrew Fitzgibbon. KinectFusion: Real-time dense surface mapping and tracking. In *2011 10th IEEE international symposium on mixed and augmented reality*, pages 127–136. IEEE, 2011. 1, 3
- [30] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-Net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 2
- [31] Shunsuke Saito, Zeng Huang, Ryota Natsume, Shigeo Morishima, Angjoo Kanazawa, and Hao Li. PIFu: Pixel-aligned implicit function for high-resolution clothed human digitization. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2304–2314, 2019. 2
- [32] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *European Conference on Computer Vision*, pages 501–518. Springer, 2016. 2, 3
- [33] Shuran Song and Jianxiong Xiao. Deep sliding shapes for amodal 3D object detection in RGB-D images. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 808–816, 2016. 3
- [34] Christoph Strecha, Rik Fransens, and Luc Van Gool. Wide-baseline stereo from multiple views: a probabilistic account. In *Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2004. CVPR 2004.*, volume 1, pages I–I. IEEE, 2004. 3
- [35] Christoph Strecha, Rik Fransens, and Luc Van Gool. Combined depth and outlier estimation in multi-view stereo. In *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, volume 2, pages 2394–2401. IEEE, 2006. 3
- [36] Jian Sun, Yin Li, Sing Bing Kang, and Heung-Yeung Shum. Symmetric stereo matching for occlusion handling. In *2005 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*, volume 2, pages 399–406. IEEE, 2005. 3
- [37] Jiaming Sun, Yiming Xie, Linghao Chen, Xiaowei Zhou, and Hujun Bao. NeuralRecon: Real-time coherent 3D reconstruction from monocular video. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15598–15607, 2021. 1, 2, 3, 6, 7
- [38] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. MnasNet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2820–2828, 2019. 3
- [39] Haotian Tang, Zhijian Liu, Shengyu Zhao, Yujun Lin, Ji Lin, Hanrui Wang, and Song Han. Searching efficient 3D architectures with sparse point-voxel convolution. In *European Conference on Computer Vision*, pages 685–702. Springer, 2020. 3, 6
- [40] Engin Tola, Vincent Lepetit, and Pascal Fua. Daisy: An efficient dense descriptor applied to wide-baseline stereo. *IEEE transactions on pattern analysis and machine intelligence*, 32(5):815–830, 2009. 2
- [41] Alex Trevithick and Bo Yang. Grf: Learning a general radiance field for 3d scene representation and rendering. *arXiv preprint arXiv:2010.04595*, 2020. 3, 8
- [42] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. 2, 3, 4
- [43] Kaixuan Wang and Shaojie Shen. MVDepthNet: Real-time multiview depth estimation neural network. In *2018 International conference on 3D vision (3DV)*, pages 248–257. IEEE, 2018. 2
- [44] 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)*, 38(5):1–12, 2019. 3
- [45] Haozhe Xie, Hongxun Yao, Xiaoshuai Sun, Shangchen Zhou, and Shengping Zhang. Pix2vox: Context-aware 3d reconstruction from single and multi-view images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2690–2698, 2019. 3
- [46] Bo Yang, Sen Wang, Andrew Markham, and Niki Trigoni. Robust attentional aggregation of deep feature sets for multi-view 3d reconstruction. *International Journal of Computer Vision*, 128(1):53–73, 2020. 3, 8
- [47] Qingxiong Yang, Liang Wang, Ruigang Yang, Henrik Stewénius, and David Nistér. Stereo matching with color-weighted correlation, hierarchical belief propagation, and occlusion handling. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(3):492–504, 2008. 3
- [48] Zhenpei Yang, Zhile Ren, Qi Shan, and Qixing Huang. Mvs2d: Efficient multi-view stereo via attention-driven 2d convolutions. *arXiv preprint arXiv:2104.13325*, 2021. 3
- [49] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 767–783, 2018. 2

- [50] Lior Yariv, Yoni Kasten, Dror Moran, Meirav Galun, Matan Atzmon, Basri Ronen, and Yaron Lipman. Multiview neural surface reconstruction by disentangling geometry and appearance. *Advances in Neural Information Processing Systems*, 33, 2020. 3
- [51] Anzhu Yu, Wenyue Guo, Bing Liu, Xin Chen, Xin Wang, Xuefeng Cao, and Bingchuan Jiang. Attention aware cost volume pyramid based multi-view stereo network for 3d reconstruction. *ISPRS Journal of Photogrammetry and Remote Sensing*, 175:448–460, 2021. 3
- [52] Zehao Yu and Shenghua Gao. Fast-MVSNet: Sparse-to-dense multi-view stereo with learned propagation and gauss-newton refinement. In *Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020. 7
- [53] Yi Yuan, Jilin Tang, and Zhengxia Zou. Vanet: a view attention guided network for 3d reconstruction from single and multi-view images. In *2021 IEEE International Conference on Multimedia and Expo (ICME)*, pages 1–6. IEEE, 2021. 3
- [54] Enliang Zheng, Enrique Dunn, Vladimir Jovic, and Jan-Michael Frahm. Patchmatch based joint view selection and depthmap estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1510–1517, 2014. 3
- [55] Qian-Yi Zhou, Jaesik Park, and Vladlen Koltun. Open3D: A modern library for 3D data processing. *arXiv:1801.09847*, 2018. 6
- [56] C Lawrence Zitnick and Takeo Kanade. A cooperative algorithm for stereo matching and occlusion detection. *IEEE Transactions on pattern analysis and machine intelligence*, 22(7):675–684, 2000. 3