

MonoPlane: Exploiting Monocular Geometric Cues for Generalizable 3D Plane Reconstruction

Wang Zhao^{*1}, Jiachen Liu^{*2}, Sheng Zhang³, Yishu Li¹, Sili Chen³,
Sharon X Huang², Yong-Jin Liu¹ and Hengkai Guo³


Abstract—This paper presents a generalizable 3D plane detection and reconstruction framework named MonoPlane. Unlike previous robust estimator-based works (which require multiple images or RGB-D input) and learning-based works (which suffer from domain shift), MonoPlane combines the best of two worlds and establishes a plane reconstruction pipeline based on monocular geometric cues, resulting in accurate, robust and scalable 3D plane detection and reconstruction in the wild. Specifically, we first leverage large-scale pre-trained neural networks to obtain the depth and surface normals from a single image. These monocular geometric cues are then incorporated into a proximity-guided RANSAC framework to sequentially fit each plane instance. We exploit effective 3D point proximity and model such proximity via a graph within RANSAC to guide the plane fitting from noisy monocular depths, followed by image-level multi-plane joint optimization to improve the consistency among all plane instances. We further design a simple but effective pipeline to extend this single-view solution to sparse-view 3D plane reconstruction. Extensive experiments on a list of datasets demonstrate our superior zero-shot generalizability over baselines, achieving state-of-the-art plane reconstruction performance in a transferring setting. Our code is available at <https://github.com/thuzhaowang/MonoPlane>.

I. INTRODUCTION

Reconstructing 3D planes from images is a fundamental task in 3D vision that has wide applications in downstream tasks, such as robotics, mixture reality and interior modeling. The goal of this task is to simultaneously detect planes and estimate plane parameters from single or multiple images.

Traditional methods rely on multi-view cues [1], depth sensor inputs [2], [3], or certain heuristic scene assumptions [4], and utilize robust estimators like RANSAC [5] and MRFs [6] optimization to fit plane proposals and label pixels. These methods achieve desirable plane reconstruction results in the presence of depth or multi-view input. However, their performance may drastically degrade when the heuristic assumptions are broken, or encountering low-texture images, motion blurs and noisy or incomplete depth.

Learning-based methods have been proposed to reconstruct 3D planes from a single image [7], [8], [9], multi-view images [10], [11], [12] or videos [13]. Thanks to the construction of large-scale benchmark datasets [14], [15], these methods have showcased remarkable performance in



Indoor - Matterport [15] Outdoor - Synthia [23] In-the-wild Data [24]

Fig. 1. Sample results from diverse scenarios. Our method achieves superior generalizable 3D plane detection and reconstruction by incorporating generalizable cues and robust designs.

segmenting plane instances and regressing plane parameters on particular dataset. To achieve better plane reconstruction, researchers have proposed to leverage more advanced network architectures [7], [16], or utilize geometric primitives such as line segments [9] to assist, or optimize multi-view planes and camera pose in a joint manner [17], [12]. Despite the significant progress, such data-driven methods still require costly groundtruth labeling and more importantly, have shown poor generalization ability onto out-of-distribution data as shown in our experiments. However, transferring across different environments is a crucial capacity in real-world applications such as robot navigation or augmented reality.

Recently, several works explore the possibility of achieving generalizable depth or normal estimation from single-view images [18], [19], [20]. These works leverage rich depth-relevant data from different sources such as stereo, RGB-D video, 3D movies, and mix them together to obtain large-scale training datasets. These trained neural networks demonstrate superior zero-shot accuracy on in-the-wild data. Some recent studies have leveraged these pre-trained models to assist in different 3D vision tasks such as neural 3D reconstruction and rendering [21], [22]. Inspired by these recent successes using pretrained monocular geometric models, we aim to exploit the potential of monocular geometric cues in 3D plane reconstruction. Taking a single RGB image as input, the pretrained networks are employed to predict depth and surface normals. To make fully use of these geometric cues, we revisit the classical RANSAC [5] pipeline, and further propose a monocular cue guided proximity modeling approach upon the graph-cut RANSAC [25]. This specially-designed RANSAC pipeline is able to effectively utilize the monocular geometric information while being robust to handle noisy input. Then, an image-level multi-plane joint optimization is

*Equal Contribution.

†This work was supported by the Natural Science Foundation of China (62332019).

¹ BNRist, MOEKey Laboratory of Pervasive Computing, Department of Computer Science and Technology, Tsinghua University, Beijing, China.

² College of Information Sciences and Technology, Pennsylvania State University, University Park, PA 16802, USA.

³ Bytedance Inc., Beijing, China.

conducted via dense CRFs [26] to encourage more consistent plane masks by combining the RGB color information. In this way, our method, termed MonoPlane, without plane labels and domain-specific training, effectively predicts satisfactory 3D plane reconstruction results from a single RGB image. We further demonstrate that our single-view solution can be easily extended to sparse uncalibrated images. Experimental results across different benchmarks and in-the-wild data (Figure 1) showcase that our proposed framework is accurate, generalizable and scalable.

To summarize, our contributions are three-fold:

- We explore the use of monocular geometric cues for 3D plane reconstruction, delivering generalizable 3D plane reconstruction from a single or sparse RGB images.
- We propose a point-proximity based graph-cut RANSAC algorithm by incorporating monocular geometric cues, which greatly enhanced the robustness on handling noisy input or complex scenes.
- We build a generalizable 3D plane reconstruction system for both single image and sparse-view images without any plane groundtruth or domain-specific finetuning. Experiments on various data from diverse environments clearly demonstrate the effectiveness of our system.

II. RELATED WORK

A. Piece-wise Plane Reconstruction from Images

Traditional 3D plane reconstruction approaches typically use RGB-D image streams [2], [3] as input, segment the planes with 3D point cloud, or require multi-view images [1] to build cross-view correspondences and objectives for plane discovery. Robust estimators such as RANSAC [5] and optimization tools such as Conditional Random Field (CRF) [27] and Markov Random Field (MRF) [1], [4] are commonly employed to fit the plane model and label image pixels. Although these methods achieve reasonable results on plane reconstruction, they often rely on expensive and inflexible equipment, such as multi-view setups or depth sensors, which limit their practical use in real-world applications. Besides, certain methods [4] assume strong scene priors such as Manhattan-world stereo, thereby restricting their applicability to versatile and complex environments. In recent years, some learning-based methods have been proposed, aiming to recover 3D planes from a single RGB image in a fully supervised way. The detection and parameter estimation processes are typically performed either in a top-down [28], [7], [8], or bottom-up [29], [30], [9], [16] manner.

To leverage the multi-view geometry, PlaneMVS [10] proposes a slanted-hypotheses-based plane warping strategy, enabling recovering planes from multiple views with known camera poses. PlanarRecon [13] presents a voxel-based approach to incrementally track and reconstruct planes of a scene from posed monocular videos. Another line of works [17], [12], [11] focus on reconstructing planes from sparse views with unknown camera poses. S3PREcon [31] proposes a self-supervised way to cluster 3D planes from monocular input images through neural 3D reconstruction.

However, a general limitation in these learning-based methods is that they predict planar geometry using deep features trained on specific, homogeneous environments. As a result, their generalizability to novel scenes is highly constrained, especially when transferring to data with significant gaps from the training data. To resolve this bottleneck, our proposed method focuses on developing a transferable system that leverages pre-trained monocular geometric models and a robust plane estimator based on RANSAC. By adopting these techniques, we aim to avoid introducing biases associated with deep features from particular datasets and enhance the adaptability of plane reconstruction across diverse scenes.

B. Generalizable Monocular 3D Estimation

3D geometry estimation including estimating depth [32], [33], [34], [35], [36] and surface normals [33], [37] from a single image has been studied via deep learning. These methods design effective networks such as CNNs [32], [37] and Transformers [35], [36], showing superior performance on specific benchmarks [38], [14], [39], [40]. However, their generalizability to new environments is generally limited, posing challenges when applied to real-world data.

In recent studies, zero-shot transfer learning on 3D estimation [18], [19], [41], [20] is drawing increasing attention. MiDaS [18] constructs a large-scale, mixed dataset from diverse sources and introduces a scale-shift invariant training loss to learn relative depth across different annotation formats. DPT [41] proposes a network which combines the advantage of vision transformers [42] and CNNs to achieve global receptive field while preserving fine-grained details on depth estimation. Omnidata [19] designs a data-centric approach to generate massive high-quality training data to facilitate generalizable geometric prediction tasks. To address the limitation of previous work on merely estimating relative depth, follow-up works [20], [43], [44] aim to learn scale-preserving metric depth through scale-adaptive strategies or further enlarged datasets across domains and datasets.

The aforementioned works serve as strong motivation for our research, inspiring us to develop a cross-dataset 3D plane reconstruction system with enhanced generalizability. Additionally, the monocular geometric cues inferred from these models, offer valuable initializations and solid foundation for our plane reconstruction pipeline. These cues provide a desirable starting point and contribute significantly to the effectiveness of our overall reconstruction process.

III. METHOD

Our approach endeavors to detect and reconstruct 3D planes from images in a generalizable manner. This is achieved by utilizing mix-dataset pre-trained monocular prediction models, and proximity-based RANSAC optimization. Starting with a single image with predicted monocular depth, we first demonstrate a sequential RANSAC baseline and discuss its limitations in Section III-A. To address these limitations, we introduce a graph-based optimization method guided by monocular cues in Section III-B. This method effectively capitalizes the rich proximity information of the points to




Fig. 2. Pipeline of our proposed method MonoPlane. Given a single image as input, a pre-trained monocular network predicts depths and normals for each image, resulting in 3D oriented point clouds. Our proposed proximity based graph-cut RANSAC is then applied to handle noisy point clouds and sequentially segment 3D planes, followed by projection and dense CRFs to output plane masks and parameters.




Fig. 3. Example results of different RANSAC baselines and our method.

determine the appropriate planes. Finally, in Section III-C, we extend the single-image solution to sparse-view image input. An overview of our pipeline is shown in Figure 2.

A. Sequential RANSAC Baseline

Given an RGB image I as input, we first utilize a large-scale pretrained model such as Omnidata [19] to predict the depth D and surface normal N . Point cloud P is then obtained through un-projection $P = D * K^{-1} * X$, where X denotes the homogeneous pixel coordinate and K is the camera intrinsic. To handle the variation of monocular depth scale, we normalize the point cloud roughly into a unit sphere. With the 3D point cloud available, a straight-forward idea is to use the RANSAC algorithm to fit 3D planes.

In the RANSAC process, the sampler first randomly selects a minimal sample of 3 points from point set $\mathcal{P} \subseteq \mathbb{R}^3$ to estimate a plane model $\theta \in \mathbb{R}^4$. Then the point-to-plane distance is used as the cost function $\phi : \mathcal{P} \times \mathbb{R}^4 \rightarrow \mathbb{R}$ to calculate all the data errors. Points with error smaller than threshold ϵ are labeled as inliers, otherwise outliers. The scorer simply assesses the inlier number to score an estimated plane model. This process repeats until certain iterations are reached or sufficient inliers are found. And the plane with the highest score is selected as the best fit model. This standard RANSAC procedure can also be formulated as an energy minimization problem, as stated in [25]. Denote $L \in \{0, 1\}^{|\mathcal{P}|}$ as a binary labeling (0-outlier, 1-inlier) of $|\mathcal{P}|$ points, the data energy is defined as follows:

$$E_d(L) = \sum_{p \in \mathcal{P}} \|L_p\|_{\{0,1\}} \quad (1)$$

where

$$\|L_p\|_{\{0,1\}} = \begin{cases} 0 & \text{if } (L_p = 0 \wedge \phi(p, \theta) \geq \epsilon) \vee \\ & (L_p = 1 \wedge \phi(p, \theta) < \epsilon) \\ 1 & \text{otherwise.} \end{cases} \quad (2)$$

Minimizing this energy leads to the same binary inlier-outlier labeling given current hypothesis model θ . This

standard RANSAC procedure finds a best plane model and its inlier set, one at a time. Therefore, to detect arbitrary number of plane instances, we sequentially apply the RANSAC on the remaining outlier points of current step.

We test this sequential RANSAC for 3D plane fitting from monocular depths, as shown in Figure 3. Sequential RANSAC is generally able to detect the dominant planes, however, the plane detection results contain massive outliers and small pieces of plane fragments. The main reason is that, the predicted monocular depths, although depicting the rough geometric structures, still involve non-negligible noise. The noise distribution is uneven, hard to model and not tolerable with a single global inlier-outlier threshold, making this simple RANSAC baseline impractical for real-world use. To this end, we propose to utilize the point-to-point proximity in graphs, which assists to reduce the effects from depth noise.

B. Graph Optimization with Monocular Cues

The vanilla RANSAC treats each point independently, *i.e.*, each point is labeled as inlier or outlier solely based on its own point-to-plane distance regardless of the other points. However, a 3D plane, as a geometric primitive, is defined locally thus naturally conforms to the spatial coherence constraints. Thus, we aim to explore point relations to assist with RANSAC labeling. Inspired by this observation, in this section we solve the RANSAC energy optimization problem using graphs to utilize point proximity information.

1) *Sequential Graph-cut RANSAC for Plane Fitting*: In RANSAC community, Graph-cut RANSAC [25] shares this insight and models the spatial coherence with a graph. GC-RANSAC has shown impressive results on several geometric tasks such as fundamental/essential matrix estimation, homography and affine transformation estimation. We further extend it to 3D plane fitting. Specifically, GC-RANSAC introduces a neighborhood graph $\mathcal{G} = (\mathcal{P}, \mathcal{E})$, and then utilizes a new pairwise energy term E_s in addition to the data energy term:

$$E_s = \sum_{p, q \in \mathcal{E}} \begin{cases} 1 & \text{if } L_p \neq L_q \\ 0 & \text{if } L_p = L_q = 1 \\ \frac{1}{2}(f(0, p) + f(0, q)) & \text{if } L_p = L_q = 0 \end{cases} \quad (3)$$

where

$$f(L_p, p) = \begin{cases} 0 & \text{if } (L_p = 0 \wedge \phi(p, \theta) \geq \epsilon) \\ & \vee (L_p = 1 \wedge \phi(p, \theta) < \epsilon) \\ \frac{\phi^2(p, \theta)}{\epsilon^2} & \text{if } L_p = 1 \wedge \phi(p, \theta) \geq \epsilon \\ 1 - \frac{\phi^2(p, \theta)}{\epsilon^2} & \text{if } L_p = 0 \wedge \phi(p, \theta) < \epsilon \end{cases} \quad (4)$$

is the cost function with truncated L_2 loss used in MSAC [45]. Note that other alternative cost functions can also be adopted, such as the Gaussian-kernel function [46] or MAGSAC++ [47] loss function. Please refer to [25] for more details. This pairwise energy E_s , following the Potts model [6], is designed to penalize differently labeled neighbors, and avoid penalizing neighboring points p and q if they are both labeled as inliers, thus encouraging the spatial coherence of the inlier set. The data term is $E_d = \sum_{p \in \mathcal{P}} f(L_p, p)$ and the final energy E is defined as:

$$E = (1 - \lambda)E_d + \lambda E_s \quad (5)$$

where λ is a hyper-parameter to control the smoothness. This optimization can be solved efficiently and globally via the graph-cut algorithm. The desired globally optimal labeling $L^* = \arg \min_L E(L)$ is determined in polynomial time.

We implement the sequential GC-RANSAC and test it for 3D plane fitting from monocular depths. Figure 3 shows that the GC-RANSAC greatly improves the spatial coherence. However, it still suffers from over-smoothness or under-smoothness given noisy monocular depths, and is sensitive to the balancing weight λ and the distance threshold ϵ .

2) Point Proximity Modeling with Monocular Cues:

To tackle the above issues, we introduce point-to-point proximity inside GC-RANSAC to better model the point relations. The insight is simple: In GC-RANSAC, giving different labels to neighboring points is penalized via a constant energy term, regardless of the points' similarity. However, we propose that the penalty should be conditioned on the probability of neighboring points being labeled as the same. For example, if two points are neighbors spatially but have distinct colors and surface normals, labeling these two points as different is supposed to be reasonable and the penalty should be small. This proximity modeling has been widely explored in vision community such as markov random fields [26], bilateral filtering [48], etc. Different from previous works, our framework introduces proximity modeling into the RANSAC procedure via Graph-cut RANSAC, and enhances the modeling with additional monocular cues obtained from pretrained neural networks.

Specifically, we formulate three proximity terms k_p, k_c, k_n from measures of similarity about positions, colors, and surface normals. These terms are formulated like Gaussian kernels. For two neighboring points with 3D positions p_i, p_j , image color vectors I_i, I_j and surface normals n_i, n_j , the proximity energy of each term is defined as:

$$k_p = w_p \exp\left(-\frac{|p_i - p_j|^2}{2\alpha_p^2}\right) \quad (6)$$

$$k_c = w_c \exp\left(-\frac{|p_i - p_j|^2}{2\alpha_p^2} - \frac{|I_i - I_j|^2}{2\alpha_c^2}\right) \quad (7)$$

$$k_n = w_n \exp\left(-\frac{|p_i - p_j|^2}{2\alpha_p^2} - \frac{\arccos(n_i \cdot n_j)}{2\alpha_n^2}\right) \quad (8)$$




Fig. 4. Extension to sparse views. We first utilize our single-view pipeline to get 3D plane proposals for each image. These plane proposals are then matched and fused into global 3D planes.

Now the pairwise energy term E_s is formulated as:

$$E_s = \sum_{p, q \in \mathcal{E}} \begin{cases} 0.5 + k_p + k_c + k_n & \text{if } L_p \neq L_q \\ 0 & \text{if } L_p = L_q = 1 \\ \frac{1}{2}(f(0, p) + f(0, q)) & \text{if } L_p = L_q = 0 \end{cases} \quad (9)$$

The graph-cut optimization solver [49] used in GC-RANSAC requires the inequality $e_{pq}(0, 0) + e_{pq}(1, 1) \leq e_{pq}(1, 0) + e_{pq}(0, 1)$ to hold. Since $k_p, k_c, k_n > 0$ and $f(0, \cdot) \leq 1$, our pairwise energy meets the sub-modular requirements.

As for the data energy term, we inherit E_d of the GC-RANSAC and further utilize surface normal to define an additional data energy term. The observation is that the predicted monocular surface normals should show strong homogeneity in planar regions, thus can be used for rejecting outlier points based on normal similarity. The proposed normal data energy is $E_n = \sum_{p \in \mathcal{P}} g(L_p, p)$, where

$$g(L_p, p) = \begin{cases} \exp\left(\frac{\rho(p, \theta_n)}{\epsilon_n} - 1\right) & \text{if } L_p = 1 \wedge \rho(p, \theta_n) \geq \epsilon_n \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

$\rho: \mathcal{P} \times \mathbb{R}^3 \rightarrow \mathbb{R}$ denotes the normal error function, which calculates the normal angle difference between point p and current model normal θ_n . ϵ_n is the normal angle threshold. In this energy term, we only penalize those points that are labeled as inliers but still have large normal errors. Exponential cost prevents data points whose normals are very different from the model normal from being labeled as inliers. Note that the model normal θ_n is the average of monocular normal predictions of current minimal or non-minimal samples. The overall optimization energy is then:

$$E = (1 - \lambda)E_d + (1 - \lambda)E_n + \lambda E_s \quad (11)$$

Finally, we utilize the surface normal homogeneity to help the RANSAC sampler. Given a minimal data sample, we calculate the normal angle differences among each pair of points. If any pair of points has greater angle difference than threshold ϵ_n , this minimal sample will be abandoned. This normal check can effectively reject bad minimal samples and accelerate the RANSAC procedure.

As shown in Figure 3, our method greatly improves over the previous baselines, producing accurate and smooth plane detections. Thanks to point proximity modeling guided by monocular cues, our method can handle imperfect inputs and find the most fit plane model with the best inlier set.




Fig. 5. Qualitative results for single-view plane segmentation. We show two samples for ScanNet [14] (row 1&2), Matterport3D [15] (row 3&4), and Synthia [23] (row 5&6), respectively, for all methods.

3) *Image-level Multi-plane Joint Optimization*: Although the proximity-based sequential GC-RANSAC delivers elegant plane reconstruction results, the plane masks may have minor inconsistency with each other due to the sequential binary labeling manner. Thus we apply the densely connected conditional random fields (CRFs) [26] as a post-processing step to jointly optimize multiple plane masks, resulting in more pleasing results. We also incorporate the monocular cues inside the CRFs to better model the pairwise potentials, with a formulation similar to Eq. 6, 7, 8.

C. Extension to Sparse-view Images

The proposed single-view pipeline can be easily extended to sparse views. For simplicity, here we use two view inputs to demonstrate (Figure 4). We begin by using the Structure-from-Motion (SfM) to get the camera parameters and sparse 3D points. We then align the monocular depths of multiple images with this sparse reconstruction to mitigate the scale inconsistency. After that, single view plane reconstruction is applied for each image to get the 3D plane proposals. Finally we match these plane proposals between image pairs and fuse the matched planes into global 3D planes.

Camera poses and scale alignment. We adopt the COLMAP [50] to provide camera parameters. To improve accuracy, we utilize LoFTR [51] as the feature extractor and matcher. Since depths of different views are predicted independently and have no scale consistency guaranteed, we explicitly align each depth map with the SfM reconstruction by solving a least squares problem.

Plane matching and fusion. Plane matching between two images is jointly determined by the 2D sparse pixel matching from SfM and the plane re-projection overlap check. As for the fusion, we use the average plane as the fused plane of matched pairs. Following [17], the offsets are averaged and the fused normal \mathbf{n} is solved by maximizing the $\sum_i (\mathbf{n}^T \mathbf{n}_i)^2$, where $i = 1, 2$ for two view pairs.

IV. EXPERIMENTS

A. Implementation and Evaluation Details

We test Omnidata [19] and ZeroDepth [43] as our pre-trained zero-shot monocular network to predict depths. Omnidata [19] is also used to predict surface normals. The hyper-parameters of our method are set as: the distance threshold ϵ as 0.02, the normal angle threshold ϵ_n as 10, the smoothness weight λ as 0.95. For proximity modeling, the parameters of each similarity term are set as: position term α_p^2 as 0.005, image color term α_c^2 as 0.1, normal term α_n^2 as 5.0. w_p, w_c, w_n are all set to 1. We use the same hyper-parameter setting in all experiments.

For single-view evaluation, we follow the evaluation protocol of PlaneRCNN [7]. Three metrics named Variation of Information (VOI), Rand Index (RI), and Segmentation Covering (SC) are used to assess plane segmentation. For planar reconstruction quality, we evaluate per-plane recall under different reconstructed depth gap thresholds.

For sparse-view evaluation, we follow the protocol of SparsePlane [17] except that we filter out the extreme

TABLE I

ZERO-SHOT SINGLE-VIEW EVALUATION ON INDOOR DATASETS NYUV2 [39] AND MATTERPORT3D [15]. “SCANNET/MATTERPORT3D” INDICATES THAT METHOD IS TRAINED ON SCANNET [14] OR MATTERPORT3D [15]. “OMNIDATA/ZERODEPTH” DENOTES OUR VERSIONS USING OMNIDATA [19] OR ZERODEPTH [43] DEPTH PREDICTION AS INPUT. THE BEST AND SECOND BEST RESULTS ARE HIGHLIGHTED BY BOLDFACE AND UNDERLINE.

Method / Metrics	Plane Segmentation Metrics			Plane Recall (depth)		
	VOI (\downarrow)	RI (\uparrow)	SC (\uparrow)	@0.05m	@0.1m	@0.6m
NYUv2 [39]						
PlaneRCNN (ScanNet) [7]	1.540	0.850	0.627	0.039	0.123	0.447
PlaneTR (ScanNet) [9]	1.111	0.899	0.725	0.039	0.113	0.333
PlaneRecTR (ScanNet) [16]	1.045	<u>0.915</u>	<u>0.745</u>	0.06	0.157	0.443
MonoPlane (Omnidata)	0.907	0.927	0.770	0.069	0.211	0.497
MonoPlane (ZeroDepth)	<u>1.011</u>	<u>0.915</u>	0.744	<u>0.066</u>	<u>0.172</u>	<u>0.449</u>
Matterport3D [15]						
PlaneRCNN (ScanNet) [7]	1.826	0.837	0.585	0.027	0.084	0.367
PlaneTR (ScanNet) [9]	1.553	0.846	0.610	0.027	0.068	0.220
PlaneRecTR (ScanNet) [16]	1.419	0.875	0.655	0.045	0.108	0.330
MonoPlane (Omnidata)	1.135	0.896	0.709	0.074	0.188	0.464
MonoPlane (ZeroDepth)	<u>1.191</u>	<u>0.890</u>	<u>0.693</u>	<u>0.059</u>	<u>0.155</u>	<u>0.417</u>

TABLE II

ZERO-SHOT SINGLE-VIEW EVALUATION ON OUTDOOR DATASET SYNTHIA [23].

Method / Metrics	Plane Segmentation Metrics			Plane Recall (depth)		
	VOI (\downarrow)	RI (\uparrow)	SC (\uparrow)	@1m	@3m	@10m
	PlaneRCNN (ScanNet) [7]	1.881	0.715	0.540	<u>0.092</u>	0.121
PlaneTR (ScanNet) [9]	1.254	0.849	0.689	0.027	0.088	0.133
PlaneRecTR (ScanNet) [16]	1.261	0.852	0.698	0.027	0.119	0.165
MonoPlane (Omnidata)	<u>1.060</u>	<u>0.898</u>	<u>0.750</u>	0.089	<u>0.178</u>	<u>0.223</u>
MonoPlane (ZeroDepth)	1.032	0.899	0.756	0.119	0.179	0.243

image pairs that have below 10% visual overlap. The plane reconstruction accuracy is evaluated by AP given certain mask IoU, plane normal and offset thresholds.

For both single-view and sparse-view evaluation, to eliminate the scale ambiguity from monocular input, for every method, we align the inferred depth with groundtruth depth by multiplying the median scale factor of the groundtruth depth w.r.t. the inferred depth map before evaluation.

Running time. On average, for a single 192x256 image, the proposed RANSAC takes 0.67 seconds, and the dense CRFs takes 0.22 seconds. For two-view images, the LoFTR matching and the COLMAP SfM in total take 1.10 seconds.

B. Single-view Plane Reconstruction Evaluation

We report the single-view plane reconstruction results on two zero-shot indoor datasets NYUv2 [39], Matterport3D [15] and an outdoor dataset Synthia [23]. Baseline methods include PlaneRCNN [7], PlaneTR [9] and PlaneRecTR [16] which are trained on ScanNet [14]. As shown in Tab. I, For both NYUv2 and Matterport3D which is unseen during training for all methods, our proposed method MonoPlane significantly outperforms the baselines on all metrics. On Synthia (i.e. Tab. II), MonoPlane also clearly outperforms the remaining methods which transfer from indoor scenes. See Fig. 5 for qualitative examples. In general, none of those learning-based methods can adapt well to diverse inputs. Contrastively, our method detects and reconstructs planes in a robust manner,

TABLE III

ZERO-SHOT TWO-VIEW EVALUATION ON SCANNET [14] DATASET.

Method / Metrics	Offset \leq 1m, Normal \leq 30°			
	All	-Offset	-Normal	-Offset-Normal
PlaneFormers (Matterport3D) [12]	2.07	2.77	5.04	8.69
NopeSAC (Matterport3D) [11]	2.47	3.17	5.13	7.09
MonoPlane (Omnidata)	9.15	15.74	12.21	21.16

and generalizes well to different scenes, which is the main merit and motivation of this work.

C. Sparse-view Plane Reconstruction Evaluation

Two-view plane reconstruction quality is evaluated on ScanNet dataset, compared with two SOTA methods PlaneFormers [12] and NopeSAC [11] which are both trained on Matterport3D. The results are summarized in Tab. III. MonoPlane significantly outperforms these counterparts. As shown in Fig. 6, PlaneFormers [12] and NopeSAC [11] fail to recover the major planar surface when exposed to a new environment with different distribution from the training data. Conversely, our approach showcases superior robustness on these challenging scenarios, clearly demonstrating its scalability when extended to sparse-view input.

D. In-the-wild Data

We further test our framework on in-the-wild data from several publicly available datasets, including 3DPW [53], KITTI [38], TUM-RGBD [52], DAVIS [24], etc. Fig. 7 shows the qualitative comparison between our method and PlaneRCNN [7]. Our method MonoPlane transfers well to the in-the-wild scenes and clearly outperforms the learning-based counterpart. To further demonstrate the feasibility of our method in practical applications such as robotics environment, we also apply our system on a real-world scenario using a mobile robot. This real-world application and more qualitative results on single-view and sparse-view reconstruction of in-the-wild data are shown in the supplementary video.

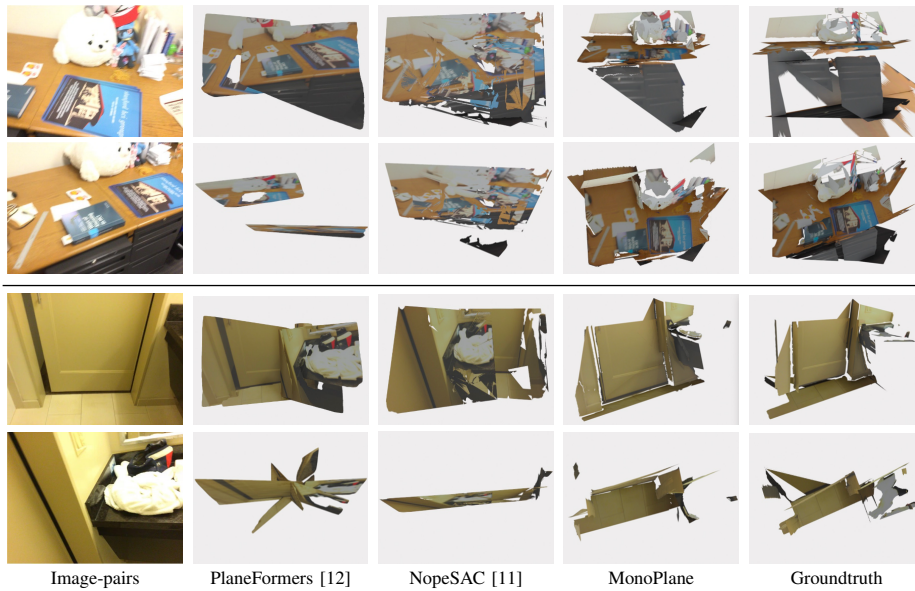


Fig. 6. Reconstructed planar mesh from two views of ScanNet [14] dataset. We show a frontal view (1st and 3rd rows) and a bird-eye view (2nd and 4th rows) for each case respectively.




Fig. 7. Plane segmentation results on in-the-wild data from TUM-RGBD [52], 3DPW [53] and DAVIS [24].

E. Ablation Studies

We conduct several ablation experiments of single-view plane reconstruction on ScanNet dataset. As shown in Tab. IV, our method surpasses the sequential RANSAC and Graph-cut RANSAC by a large margin, validating the effectiveness of our system design. Figure 3 showcases the qualitative comparison of gradually adding our proposed components. Sequential-RANSAC produces messy plane fragments with incomplete masks. While GC-RANSAC greatly improves the results thanks to its coherence modeling, our method outperforms both of the two baselines and delivers more accurate and complete plane reconstructions. We also verified the necessity of the dense CRF module and conducted an ablation study. As illustrated, the dense CRF can merge or remove some small fuzzy planes, and make the plane segmentation more coherent around image boundaries.

V. CONCLUSION

In this paper, we present a novel framework named MonoPlane, to detect and reconstruct 3D planes from a single RGB image. The key idea is exploiting monocular geometric cues (depth and surface normal) to guide the point proximity based graph-cut RANSAC for 3D plane

TABLE IV
ABLATION STUDIES ON SCANNET [14] DATASET.

Method / Metrics	Plane Segmentation Metrics			Plane Recall (depth)		
	VOI (\downarrow)	RI (\uparrow)	SC (\uparrow)	@0.05m	@0.1m	@0.6m
Seq-RANSAC [5]	2.543	0.677	0.407	0.108	0.189	0.247
Seq-RANSAC + CRFs	1.438	0.787	0.637	0.152	0.284	0.400
GC-RANSAC [25]	1.290	0.852	0.717	0.144	0.288	0.417
GC-RANSAC + CRFs	1.028	0.872	0.758	0.157	0.322	0.492
Ours w/o CRFs	<u>1.004</u>	<u>0.906</u>	<u>0.787</u>	<u>0.205</u>	<u>0.400</u>	<u>0.612</u>
Ours	0.910	0.912	0.798	0.214	0.424	0.662

fitting. We further extend our proposed single-view approach to sparse-view plane reconstruction. MonoPlane achieves superior generalizability across different datasets, making it suitable for in-the-wild plane reconstruction in various downstream applications in robotics or augmented reality. *Limitations and future work.* Our method relies on the pretrained monocular geometry networks. While delivering remarkable transferability, it may fail when geometric predictions have great errors. In the future, we aim to extend our point-to-point proximity with more similarity measurements such as deep image features. Off-the-shelf segmentation from large models such as SAM [54] can also be used to detect planes and accelerate RANSAC process. We also plan to incorporate multi-view optimizations such as plane-camera bundle-adjustment into our sparse-view pipeline.

REFERENCES

- [1] S. Sinha, D. Steedly, and R. Szeliski, "Piecewise planar stereo for image-based rendering," in *ICCV*, 2009, pp. 1881–1888.
- [2] R. F. Salas-Moreno, B. Glocker, P. H. Kelly, and A. J. Davison, "Dense planar slam," in *ISMAR*. IEEE, 2014, pp. 157–164.
- [3] C. Raposo, M. Lourenço, M. Antunes, and J. P. Barreto, "Plane-based odometry using an rgb-d camera," in *BMVC*, vol. 2, no. 5, 2013, p. 6.
- [4] Y. Furukawa, B. Curless, S. M. Seitz, and R. Szeliski, "Manhattan-world stereo," in *CVPR*, 2009, pp. 1422–1429.
- [5] M. A. Fischler and R. C. Bolles, "Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography," *Communications of the ACM*, vol. 24, no. 6, pp. 381–395, 1981.
- [6] Y. Boykov, O. Veksler, and R. Zabih, "Markov random fields with efficient approximations," in *CVPR*, 1998, pp. 648–655.
- [7] C. Liu, K. Kim, J. Gu, Y. Furukawa, and J. Kautz, "Planercnn: 3d plane detection and reconstruction from a single image," in *CVPR*, 2019, pp. 4450–4459.
- [8] Y. Qian and Y. Furukawa, "Learning pairwise inter-plane relations for piecewise planar reconstruction," in *ECCV*, 2020, pp. 330–345.
- [9] B. Tan, N. Xue, S. Bai, T. Wu, and G.-S. Xia, "Planetr: Structure-guided transformers for 3d plane recovery," in *ICCV*, 2021, pp. 4186–4195.
- [10] J. Liu, P. Ji, N. Bansal, C. Cai, Q. Yan, X. Huang, and Y. Xu, "Planemvs: 3d plane reconstruction from multi-view stereo," in *CVPR*, 2022, pp. 8665–8675.
- [11] B. Tan, N. Xue, T. Wu, and G.-S. Xia, "Nope-sac: Neural one-plane ransac for sparse-view planar 3d reconstruction," *arXiv preprint arXiv:2211.16799*, 2022.
- [12] S. Agarwala, L. Jin, C. Rockwell, and D. F. Fouhey, "Planeformers: From sparse view planes to 3d reconstruction," in *ECCV*, 2022, pp. 192–209.
- [13] Y. Xie, M. Gadelha, F. Yang, X. Zhou, and H. Jiang, "Planarrecon: Real-time 3d plane detection and reconstruction from posed monocular videos," in *CVPR*, 2022, pp. 6219–6228.
- [14] A. Dai, A. X. Chang, M. Savva, M. Halber, T. Funkhouser, and M. Nießner, "ScanNet: Richly-annotated 3d reconstructions of indoor scenes," in *CVPR*, 2017, pp. 5828–5839.
- [15] A. Chang, A. Dai, T. Funkhouser, M. Halber, M. Niessner, M. Savva, S. Song, A. Zeng, and Y. Zhang, "Matterport3d: Learning from rgb-d data in indoor environments," *arXiv preprint arXiv:1709.06158*, 2017.
- [16] J. Shi, S. Zhi, and K. Xu, "Planerectr: Unified query learning for 3d plane recovery from a single view," in *ICCV*, 2023, pp. 9377–9386.
- [17] L. Jin, S. Qian, A. Owens, and D. F. Fouhey, "Planar surface reconstruction from sparse views," in *ICCV*, 2021, pp. 12 991–13 000.
- [18] R. Ranftl, K. Lasinger, D. Hafner, K. Schindler, and V. Koltun, "Towards robust monocular depth estimation: Mixing datasets for zero-shot cross-dataset transfer," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 44, no. 3, pp. 1623–1637, 2020.
- [19] A. Eftekhar, A. Sax, J. Malik, and A. Zamir, "Omnidata: A scalable pipeline for making multi-task mid-level vision datasets from 3d scans," in *ICCV*, 2021, pp. 10 786–10 796.
- [20] S. F. Bhat, R. Birkl, D. Wofk, P. Wonka, and M. Müller, "Zoedepth: Zero-shot transfer by combining relative and metric depth," *arXiv preprint arXiv:2302.12288*, 2023.
- [21] Z. Yu, S. Peng, M. Niemeyer, T. Sattler, and A. Geiger, "Monosdf: Exploring monocular geometric cues for neural implicit surface reconstruction," *NeurIPS*, 2022.
- [22] Z. Zhu, S. Peng, V. Larsson, Z. Cui, M. R. Oswald, A. Geiger, and M. Pollefeys, "Nicer-slam: Neural implicit scene encoding for rgb slam," *arXiv preprint arXiv:2302.03594*, 2023.
- [23] G. Ros, L. Sellart, J. Materzynska, D. Vazquez, and A. M. Lopez, "The synthia dataset: A large collection of synthetic images for semantic segmentation of urban scenes," in *CVPR*, 2016, pp. 3234–3243.
- [24] F. Perazzi, J. Pont-Tuset, B. McWilliams, L. Van Gool, M. Gross, and A. Sorkine-Hornung, "A benchmark dataset and evaluation methodology for video object segmentation," in *CVPR*, 2016, pp. 724–732.
- [25] D. Barath and J. Matas, "Graph-cut ransac," in *CVPR*, 2018, pp. 6733–6741.
- [26] P. Krähenbühl and V. Koltun, "Efficient inference in fully connected crfs with gaussian edge potentials," *NeurIPS*, vol. 24, 2011.
- [27] D. Gallup, J.-M. Frahm, and M. Pollefeys, "Piecewise planar and non-planar stereo for urban scene reconstruction," in *CVPR*, 2010, pp. 1418–1425.
- [28] C. Liu, J. Yang, D. Ceylan, E. Yumer, and Y. Furukawa, "Planetnet: Piece-wise planar reconstruction from a single rgb image," in *CVPR*, 2018, pp. 2579–2588.
- [29] F. Yang and Z. Zhou, "Recovering 3d planes from a single image via convolutional neural networks," in *ECCV*, 2018, pp. 85–100.
- [30] Z. Yu, J. Zheng, D. Lian, Z. Zhou, and S. Gao, "Single-image piecewise planar 3d reconstruction via associative embedding," in *CVPR*, 2019, pp. 1029–1037.
- [31] B. Ye, S. Liu, X. Li, and M.-H. Yang, "Self-supervised super-plane for neural 3d reconstruction," in *CVPR*, 2023, pp. 21 415–21 424.
- [32] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," *NeurIPS*, vol. 27, 2014.
- [33] D. Eigen and R. Fergus, "Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture," in *ICCV*, 2015, pp. 2650–2658.
- [34] S. F. Bhat, I. Alhashim, and P. Wonka, "Adabins: Depth estimation using adaptive bins," in *CVPR*, 2021, pp. 4009–4018.
- [35] Z. Li, Z. Chen, X. Liu, and J. Jiang, "Depthformer: Exploiting long-range correlation and local information for accurate monocular depth estimation," *Machine Intelligence Research*, pp. 1–18, 2023.
- [36] A. Agarwal and C. Arora, "Attention attention everywhere: Monocular depth prediction with skip attention," in *WACV*, 2023, pp. 5861–5870.
- [37] A. Bansal, B. Russell, and A. Gupta, "Marr revisited: 2d-3d alignment via surface normal prediction," in *CVPR*, 2016, pp. 5965–5974.
- [38] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The kitti dataset," *The International Journal of Robotics Research*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [39] N. Silberman, D. Hoiem, P. Kohli, and R. Fergus, "Indoor segmentation and support inference from rgb-d images," *ECCV*, vol. 7576, pp. 746–760, 2012.
- [40] S. Song, S. P. Lichtenberg, and J. Xiao, "Sun rgb-d: A rgb-d scene understanding benchmark suite," in *CVPR*, 2015, pp. 567–576.
- [41] R. Ranftl, A. Bochkovskiy, and V. Koltun, "Vision transformers for dense prediction," in *ICCV*, 2021, pp. 12 179–12 188.
- [42] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner, M. Dehghani, M. Minderer, G. Heigold, S. Gelly, *et al.*, "An image is worth 16x16 words: Transformers for image recognition at scale," *arXiv preprint arXiv:2010.11929*, 2020.
- [43] V. Guizilini, I. Vasiljevic, D. Chen, R. Ambrus, and A. Gaidon, "Towards zero-shot scale-aware monocular depth estimation," in *ICCV*, 2023, pp. 9233–9243.
- [44] W. Yin, C. Zhang, H. Chen, Z. Cai, G. Yu, K. Wang, X. Chen, and C. Shen, "Metric3d: Towards zero-shot metric 3d prediction from a single image," in *ICCV*, 2023, pp. 9043–9053.
- [45] P. Torr, "Bayesian model estimation and selection for epipolar geometry and generic manifold fitting," *International Journal of Computer Vision*, vol. 50, no. 1, p. 35, 2002.
- [46] P. H. Torr and A. Zisserman, "Mlesac: A new robust estimator with application to estimating image geometry," *Computer Vision and Image Understanding*, vol. 78, no. 1, pp. 138–156, 2000.
- [47] D. Barath, J. Noskova, M. Ivashchkin, and J. Matas, "Magsac++, a fast, reliable and accurate robust estimator," in *CVPR*, 2020, pp. 1304–1312.
- [48] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *ICCV*. IEEE, 1998, pp. 839–846.
- [49] V. Kolmogorov and R. Zabih, "What energy functions can be minimized via graph cuts?" *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 2, pp. 147–159, 2004.
- [50] J. L. Schonberger and J.-M. Frahm, "Structure-from-motion revisited," in *CVPR*, 2016, pp. 4104–4113.
- [51] J. Sun, Z. Shen, Y. Wang, H. Bao, and X. Zhou, "Loftr: Detector-free local feature matching with transformers," in *CVPR*, 2021, pp. 8922–8931.
- [52] J. Sturm, N. Engelhard, F. Endres, W. Burgard, and D. Cremers, "A benchmark for the evaluation of rgb-d slam systems," in *IROS*. IEEE, 2012, pp. 573–580.
- [53] T. von Marcard, R. Henschel, M. Black, B. Rosenhahn, and G. Pons-Moll, "Recovering accurate 3d human pose in the wild using imus and a moving camera," in *ECCV*, sep 2018.
- [54] A. Kirillov, E. Mintun, N. Ravi, H. Mao, C. Rolland, L. Gustafson, T. Xiao, S. Whitehead, A. C. Berg, W.-Y. Lo, *et al.*, "Segment anything," *arXiv preprint arXiv:2304.02643*, 2023.