A Confidence-based Iterative Solver of Depths and Surface Normals for Deep Multi-view Stereo

Wang Zhao^{1,3*} Shaohui Liu^{2,3*} Yi Wei¹ Hengkai Guo³ Yong-Jin Liu^{1,4} ² ETH Zurich ⁴ JCMV ³ ByteDance Inc. ¹Tsinghua University

blueber2y@gmail.com, {zhao-w19, y-wei19}@mails.tsinghua.edu.cn, guohengkai@bytedance.com, liuyongjin@tsinghua.edu.cn

Abstract

In this paper, we introduce a deep multi-view stereo (MVS) system that jointly predicts depths, surface normals and per-view confidence maps. The key to our approach is a novel solver that iteratively solves for per-view depth map and normal map by optimizing an energy potential based on the locally planar assumption. Specifically, the algorithm updates depth map by propagating from neighboring pixels with slanted planes, and updates normal map with local probabilistic plane fitting. Both two steps are monitored by a customized confidence map. This solver is not only effective as a post-processing tool for plane-based depth refinement and completion, but also differentiable such that it can be efficiently integrated into deep learning pipelines. Our multi-view stereo system employs multiple optimization steps of the solver over the initial prediction of depths and surface normals. The whole system can be trained end-to-end, decoupling the challenging problem of matching pixels within poorly textured regions from the cost-volume based neural network. Experimental results on ScanNet and RGB-D Scenes V2 demonstrate state-ofthe-art performance of the proposed deep MVS system on multi-view depth estimation, with our proposed solver consistently improving the depth quality over both conventional and deep learning based MVS pipelines. Code is available at https://github.com/thuzhaowang/idn-solver.

1. Introduction

Dense multi-view stereo (MVS) has been a longstanding fundamental topic in computer vision. The key idea of most existing techniques is to compare the similarity of image patches at different depth hypotheses, densely matching pixels across images. While great improvement has been witnessed over decades, it is still a hard problem to accurately estimate dense geometry from posed images



(c) w. solver only at training (d) w. solver at training and inference

Figure 1. By integrating the proposed iterative solver into end-toend training, our method decouples the challenging problem of estimating depth values within poorly textured regions from the initial network predictions, getting the network to focus more on reliable estimation on well textured pixels. This improves depth quality on both textured (blue) and texture-less (red) regions. Note that on texture-less areas (red box on (c)(d)), the network jointly trained with the solver only focuses on predicting reliable geometry on neighboring areas, leaving the poorly textured pixels to be resolved with the iterative solver module.

in many real-world scenarios, especially in indoor environments, where one of the most typical reasons of failure is the existence of texture-less areas (e.g. walls), which incur significant ambiguities to the matching step, as a number of different depths all result in low matching costs.

There have been various attempts to tackle this crucial problem. Global optimization [19, 5, 50] is one possible direction to resolve the ambiguities. Those methods initially assume a fixed set of superpixels and apply Markov random field with slanted-plane models. The optimization prefers globally smoothed geometry with locally planar surfaces and achieves promising performance. However, global optimization methods suffer from large computational complexity and the performance is also limited by the quality of superpixels and handcrafted dense features.

With the recent success of deep learning, learning-based

^{*}Equal contribution

methods have achieved great performance on datasets due to the improved quality of learned features and depth priors. Typically, a cost-volume based architecture [51, 47] is trained end-to-end with a pixelwise loss function on depths. While the ambiguities incurred by the texture-less areas still exist in the built cost volumes, several methods propose to utilize surface normals, either jointly predicted [38, 29] or online fitted [52, 33], to help train the depth predictions, aiming to implicitly enrich the network with locally planar priors by directly supervising the local structure of the predicted depth map. However, it is still a difficult problem to successfully learn such priors and accurately predict depths on texture-less areas directly from the ambiguous cost volumes built on warped patches.

In this paper, we propose a novel deep multi-view stereo system that decouples the locally planar priors from the cost volume based depth/normal prediction networks. The key to our proposed system is a differentiable confidence-based solver that iteratively solves for depth map and normal map by optimizing an energy potential that prefers locally planar surfaces. For each optimization iteration, the solver updates depth map by propagating from neighboring pixels with slanted planes, and updates normal map with local probabilistic plane fitting. A customized confidence map can be used to monitor both two steps.

Our deep MVS system firstly predicts depth, normal and confidence jointly and then applies multiple optimization steps of our depth-normal solver upon the initial predictions. Those multiple steps enable long-range propagation of reliable depths and normals. The whole system can be trained end-to-end, guiding the depth/normal network to focus more on confident depth/normal predictions on textured areas and leave the poorly textured regions to be resolved with the solver. For the solver module, a hybrid confidence map combining the deep confidence prediction and the conventional geometric reprojection check can be used to stabilize the optimization process at inference.

The proposed system combines the advantages of slanted-plane models [4] and learning-based techniques. With end-to-end training, the initial depth prediction is guided to focus only on reliable estimation over partial regions rather than the whole image (as shown in Fig. 1). Experimental results on ScanNet and RGB-D Scenes V2 show that our novel deep MVS achieves state-of-the-art performance in terms of accuracy on depths and surface normals, with the proposed solver consistently improving the depth quality over both conventional and deep MVS pipelines.

2. Related Work

Multi-view Stereo. Reconstructing 3D models from posed images has been widely studied over decades. Early methods [30, 12, 46] utilize volumetric optimization. Later attempts [53, 14, 6, 20, 41] reconstruct per-view depth

maps by comparing cross-view image patches, which becomes the de facto approach of modern multi-view stereo pipelines. Recent advances include employing superpixels [39], post-refinement [27, 28], and advanced propagation checkerboards [13, 48, 49], etc. With the success of deep learning, a number of learning-based techniques are proposed to tackle the problem. While several methods learn to directly predict 3D geometry as grids [25, 24], point clouds [7] and TSDF [35], per-view depth map estimation is still the top choice of most approaches [47, 51, 23, 32, 21, 29, 33, 10, 36] due to its robustness and flexibility. Most of those methods follow the spirit of conventional approaches [14, 6] and train a cost volume based neural network. Some of the methods [21, 10, 36] also integrate temporal information with recurrent networks. However, direct prediction over poorly textured areas have been one of the main difficulties of most approaches relying on cost volumes. An exception of this is DELTAS [43], which proposes to learn interest points and perform triangulation and densification over sparse points. In this work, we also focus on per-view depth estimation and introduce an endto-end MVS system equipped with a novel iterative solver, which implicitly decouples depth prediction over textureless areas from the cost-volume based network.

Confidence Estimation. Accurately estimating per-view depth map confidence is beneficial to 3D reconstruction pipelines. Early confidence measures such as matching cost, peak ratio and maximum likelihood are extensively studied in [18, 22]. Local smoothness can also contribute to the confidence measurement of the stereo problem [44, 17]. Recently, learning-based methods [42, 37] are proved to be effective on confidence estimation. Most related to us, [28] exploits the counter map acquired from the geometric consistency check for deep confidence estimation and geometric consistency check together to get our final confidence map, which monitors the solver at inference.

Depth-Normal Constraints. As depths and surface normals are naturally coupled by local plane fitting, joint depthnormal constraints are widely employed. Markov random fields with slanted-plane models were initially introduced in [4], and were later used in many stereo techniques [19, 5, 50]. Those methods operate on superpixels and prefer locally planar surfaces. [6] proposed to employ slanted support windows in its PatchMatch stereo framework, coupling the depth propagation with the randomly initialized per-pixel surface normal. [15] further introduce a minimalsurface regularization method. [45] propose to embed bilateral filtering in the plane fitting process to encourage locally planar structures of the output depth map. Recent learningbased methods also exploit surface normals to help train the depth prediction. While [11, 38, 29] jointly predict depths and surface normals together to benefit from multi-task feature learning, recent literature [52, 33] adopts online fitting. In [55], a surface normal network is trained to help solve for dense depth map from sparse depth observations. Different from prior works, our method explicitly exploits depthnormal constraints via a differentiable iterative solver.

3. An Iterative Depth-Normal Solver

The key to our deep multi-view stereo system is a confidence-based iterative depth-normal solver module, which solves for per-view depth map and normal map from the initial predictions. Specifically, we assume locally planar structure of the depth map and couple the normal map within our optimization scheme.

3.1. Energy Formulation

Our energy potential E_{total} consists of a data term E_{data} and a plane-based structural term E_{plane} , with the hyperparameter α weighting the two terms:

$$E_{total} = \alpha E_{data} + E_{plane},\tag{1}$$

The data term regularizes the output geometry with respect to the given initial geometry. Take d_i , n_i as the perpixel depth and surface normal respectively, and c_i as the per-pixel confidence. Denote the initial depth and surface normal as \hat{d}_i and \hat{n}_i . The data term is written as follows:

$$E_{data} = \sum_{i} c_i (d_i - \hat{d}_i)^2 + \sum_{i} c_i ||n_i - \hat{n}_i||^2.$$
(2)

The plane-based structural term E_{plane} enforces the output depth and surface normal to agree with the local planar assumption with respect to neighboring geometry. Let (x_i, d_i, n_i) and (x_j, d_j, n_j) denote the 2D coordinate, depth and surface normal of the current pixel and its neighboring pixel respectively. P(x, d, n) denotes the slanted plane at position x generated by spanning a plane from the corresponding 3D points recovered from d and x with the surface normal n. Our plane-based structural energy can be formulated in either of the two below directions:

- E_{j→i}. The sum (over each pixel) of distance functions between the 3D point recovered from d_i and x_i at the current pixel i and the slanted plane P(x_j, d_j, n_j) generated at its neighboring pixel j. d_{j→i} denotes the projection of the slanted plane P(x_j, d_j, n_j) at pixel i.
- E_{i→j}. The sum (over each pixel) of distance functions between the slanted plane P(x_i, d_i, n_i) generated at the current pixel i and the 3D point recovered from x_j and d_j at its neighboring pixel j. d_{i→j} denotes the projection of the slanted plane P(x_i, d_i, n_i) at pixel j.



Figure 2. Illustration of the proposed solver. (a) we use a sparse checkerboard to define the neighborhood of each pixel, which enables large receptive fields and efficient computation. (b) In the D-step, we update depth map by propagating from neighboring pixels with slanted planes. (c) In the N-step, we update normal map with local probabilistic plane fitting.

For the definition of neighboring pixels, we can either use a sparse checkerboard [13, 48] (as in Fig. 2(a)) or randomly sample sparse points within a fixed window. This enables efficient long-range propagation of the neighboring geometry. In both formulations of the structural energy the contribution of each pixel is also monitored by the per-pixel confidence c_i . Motivated by the effective formulation in the bilateral filtering techniques [2, 3], we also weight the contribution of the neighboring pixels with the edge-aware bilateral affinity, denoted as w_{ij} :

$$w_{ij} = \exp(-\frac{||x_i - x_j||^2}{2\sigma_x^2} - \frac{||I_i - I_j||^2}{2\sigma_c^2}), \qquad (3)$$

where I_i denotes the RGB value at each pixel. While we use RGB colorspace here for simplicity, the weight can be easily extended to YUV space, as employed in [2, 3].

3.2. Iterative Subproblem Optimization of Depths and Surface Normals

We aim to minimize the total energy with an effective, parallelized and differentiable approximation. To parallelize the computation, we assume fixed neighboring geometry and solve for each pixel individually at each update step. Since depths and surface normals are non-linearly coupled when propagating from neighboring pixels due to the perspective projection of the slanted planes, closed-form solution can not be acquired for the whole energy potential with depths and normals both considered in the data term¹.

Our proposal is to iteratively solve for depths and surface normals. As shown in Figure 2(b)(c), Each iteration is further divided into two steps of subproblem optimization, where we solve for depths/normals individually with the other variables (normals/depths) fixed. Those two subproblems can all have closed-form solutions. Furthermore, we alternate between the two plane-based structural term formulations to better fit each sub-problem, acquiring a

¹If only the surface normal data term is considered, closed-form solution can be derived by substitution of variables.



Figure 3. An overview of our deep multi-view stereo (MVS) system. First, We feed the target image and multiple reference images into a shared encoder to extract per-view deep image features, which are used to build a feature cost volume. Then, we jointly predict depth, surface normal and confidence from the cost volume with 3D CNNs and 2D dilated CNNs [54]. Finally, the confidence-based depth-normal solver is applied over the predicted depths and surface normals iteratively to get the final output. The whole system can be trained end-to-end. At training, a groundtruth confidence map computed by the relative depth error is used to help guide the initial prediction to focus on estimating reliable geometry on textured areas. Note that the inputs of the confidence branch are simplified for better visualization.

decoupled linear system algebraically for both steps, thus making the gradients of the whole solver tractable.

Depth Update (D-step). In the depth update step, we fix the surface normal map and solve for the optimal depth map that minimizes the depth energy E_d as follows:

$$\min_{d} E_{total} = \min_{d} E_d \tag{4}$$

$$E_d = \alpha \sum_i c_i (d_i - \hat{d}_i)^2 + \sum_i \sum_{j \in N(i)} c_j w_{ij} (d_i - d_{j \to i})^2,$$
(5)

where N(i) denotes the defined neighborhoods of the pixel *i*. We employ $E_{j\rightarrow i}$ as the plane-based structural term and computes the L2 distance between the optimized depth and the propagated depth $d_{j\rightarrow i}$. As previously discussed, we assume fixed neighboring geometry, so the propagated depth $d_{j\rightarrow i}$ here is the projection of the slanted plane $P(x_j, \hat{d}_j, \hat{n}_j)$ at pixel i. Compared to $E_{i\rightarrow j}$ which only uses one surface normal query at the pixel itself, employing $E_{j\rightarrow i}$ improves robustness to outliers in the initial surface normal map, enabling the solver to utilize surface normals in all neighboring pixels. We can derive closed-form optimal depth by setting the first-order derivative of E_d to zero.

Surface Normal Update (N-step). In the surface normal update step, we fix the depth map and solve for the optimal

surface normal that minimizes the surface normal energy E_n as follows:

$$\min_{n} E_{total} = \min_{n} E_n \tag{6}$$

$$E_{n} = \alpha \sum_{i} c_{i} ||n_{i} - \hat{n}_{i}||^{2} + \sum_{i} \sum_{j \in N(i)} c_{j} w_{ij} D_{n}(d_{j}, P(x_{i}, d_{i}, n_{i})).$$
(7)

 D_n is a distance function between the depth d_j at neighboring pixels and the slanted plane $P(x_i, d_i, n_i)$ at pixel i. The plane equation is computed from the depth d_i updated in the last step and the surface normal n_i being optimized, forming a local probabilistic plane fitting problem. Note that simply employing L2 distance between d_j and $d_{i\rightarrow j}$ as in the D-step will result in a non-quadratic surface normal energy due to the perspective projection during plane-based propagation. Instead, we can parameterize the surface normal as n = (a, b, -1) and utilize the algebraic form of the plane equation to acquire a quadratic energy, where 2 linear equations can be acquired by setting the first-order derivatives to zero, being sufficient to solve the 2-DoF surface normal n_i in closed form².

4. Deep Multi-view Stereo System

Based on the confidence-based iterative solver, we propose a deep multi-view stereo system that jointly predicts

²Please refer to our supplementary material for more details.

per-view depths, surface normals and confidence maps. Figure 3 shows an overview of the proposed system.

4.1. Estimating Initial Depths and Surface Normals

As our work does not focus on the network architectures, we mostly follow prior works [23, 29] to build the costvolume based multi-view depth-normal network. First, the target image and multiple reference images are fed into a shared neural network encoder respectively to acquire perview deep image features. Then, we apply plane sweeping on feature maps to build a feature cost volume. 3D CNNs and 2D dilated CNNs [54] are applied over the cost volume to aggregate and regularize the cost information. The depth maps are then regressed using the soft argmin operator from the final cost volume. The normal branch follows the design of [29], which concatenates the world coordinate volume with the feature cost volume, and uses cost volume slices to estimate the surface normal.

For the confidence estimation, multiple sources are utilized to facilitate the effective confidence regression, including intermediate feature maps, cost volumes, and the predicted depths/normals. We feed them into two mininetworks which consist of several dilated convolutions and a final sigmoid activation, to estimate the confidence for depth and normal respectively. Please refer to our supplementary material for details on the network architecture.

To train the network, we use smoothed L1 loss for depth and normal, and cross-entropy loss for confidence. The groundtruth confidence maps c_{dgt} and c_{ngt} are computed from the relative depth error e_{rel} and normal angle error e_{ang} : $c_{dgt} = \max(1.0 - \gamma_1 e_{rel}, 0.0)$ and $c_{ngt} = \max(1.0 - \gamma_2 e_{ang}, 0.0)$, where γ_1 , γ_2 are hyperparameters.

4.2. Integrating the Proposed Solver

After acquiring the initial predictions, we employ our solver module over the initial depth map and surface normal map. Specifically, D-step and N-step are applied iteratively for multiple iterations. We use 5 iterations for training and more steps can be used at inference. Multiple iterations lead to more accurate approximation of the optimal depth and surface normal under the local planar assumption. Moreover, since the plane-based propagation is limited by the window size of the checkerboard, using multiple iterations enables long-range propagation from reliable pixels. Empirically we find that improving the number of iterations at inference does lead to better depth quality (See Table 6).

Since all operations in the solver are differentiable, the whole system can be trained end-to-end, where the loss on the final solved geometry can be back-propagated into the network for training initial depth/normal predictions. At the training stage, we apply the depth loss and normal loss both on the initial predictions and the final solved geometry.

The end-to-end training of the system brings up several



Figure 4. Visualization on different confidence maps. The relative depth error (b) is used to compute the groundtruth confidence. While the confidence acquired by geometric consistency check (c) is more accurate, the deep confidence (d) is complete and acts as a supplements to the geometric confidence at inference.

advantages. First, the solver can be considered as a closedform refinement step to the initial depth/normal predictions, which improves robustness to noise and outliers. Moreover, the plane-based structural term favors locally planar surfaces, which is particularly beneficial on poorly textured regions and occluded areas in indoor environments, where cost-volume based approaches struggle to estimate accurate geometry due to large ambiguity. Integrating the proposed solver into our end-to-end deep MVS system eliminates the necessity for the network to handle pixels in extremely texture-less regions, since it can be left to the propagation steps of the solver at inference. This promotes our network to only focus on reliable estimation on textured areas, which largely releases the learning burden. As shown in Figure 1, this mechanism improves the depth quality significantly on both textured and texture-less regions.

To better enrich the network with such behaviors, at training we apply the iterative solver with the groundtruth depth/normal confidence map, which is computed by comparing the error of the depth/normal map with the groundtruth depth/normal. We empirically observe that high-quality confidence map for the solver is needed during training. Only using the jointly predicted confidence map at training cannot successfully enrich the initial predictions with the aforementioned nice property.

4.3. Inference

At inference, we use the trained network to predict initial depths, surface normals and confidence maps. The predicted confidence can be used in the iterative solver with the initial depth map and normal map as input. To further improve the quality of the depth confidence, we also utilize a cross-view geometric consistency check over the predicted depth maps at multiple views to acquire a geometric confidence. The target depth is reprojected and compared with the reference depths to compute the relative depth error, which is then converted into geometric confidence. If multiple reference images are available, we take the minimum over computed confidence maps to get a "strict" confi-



Figure 5. Visualization on the optimization process of the proposed depth-normal solver for a synthetic proof-of-concept experiment.

dence, which reduces the number of false-positive high confidence pixels. As shown in Figure 4, while the geometric confidence is often more accurate, the deep confidence is complete under occlusions and small overlaps, which can serve as a good supplements in practice. A hybrid confidence map can be acquired by multiplication of the deep confidence and the geometric confidence, which is used in the iterative solver to produce the final output geometry.

5. Experiments

5.1. Implementation Details

We use the ScanNet dataset [9] to train our system. The official training split is adopted. Three views from one sequence with a fixed frame interval 20 forms a training data sample. The initial depth and normal network is firstly trained with 15 epochs, then integrated with the solver and trained end-to-end for another 10 epochs. We use the Adam optimizer [26] with learning rate 1e-4 and batch size 24 on 4 Nvidia V100 GPUs. For the solver, we define the neighborhood by offsetting each coordinate with 1, 3, 5, 10 pixels horizontally and vertically, forming 16 valid entries of the propagation checkerboard (as illustrated in Figure 2(a)).

Time efficiency. A direct implementation in PyTorch takes 14ms for each iteration on a Nvidia V100 GPU. Implementing a CUDA C++ kernel can lead to 0.9ms/iter (15x speed-up) over a 4-Megapixel image, which is nearly negligible compared to the the backbone network.

5.2. Validating the Solver Module

We first perform a validation check on the effectiveness of the proposed solver module under a non-learning setup, where the solver serves as a post-processing tool for depth refinement and completion. We start with a synthesis case where four different 100x100 planar regions are projected onto a 200x200 image. For the initial input, the geometry of each pixel has 95% probability to be substituted with random noise, simulating unreliable pixels. Figure 5 shows qualitative results of applying the iterative solver on the input geometry. With the iterative optimization of depths and normals, the geometry is gradually refined to be close to the groundtruth target, demonstrating that our proposed solver successfully integrate locally planar priors into the optimization. We further test the solver on real outdoor data acquired from AdelaideRMF dataset [1], which is widely used for multi-structure fitting methods [16, 34]. Specifi-



Figure 6. Results on applying the proposed solver over sparse reconstructions from COLMAP [40, 41]. **Top row:** input image. **Middle row:** initial sparse depth map acquired from COLMAP. **Bottom row:** output depth after applying the proposed solver.

cally, we apply the iterative solver over the depth map acquired from the sparse reconstruction of COLMAP [40, 41]. Results are shown in Figure 6. Our proposed solver successfully complete the sparse input with reasonably correct output, demonstrating the effectiveness of the module.

5.3. Experiments on Multi-view Stereo

To compare our proposed deep MVS system with leading approaches, we evaluate our method on well-established benchmarks including ScanNet's official test split [9] and RGB-D Scenes V2 dataset [31]. Both datasets include challenging indoor scenes with many poorly textured regions. Quantitative results on multi-view depth estimation are shown in Table 1 and Table 3, where our system surpasses all existing state-of-the-art methods by a large margin on both datasets. Qualitative results in Figure 7 also demonstrate great improvement of our proposed system. Our method not only produces accurate depth map, but also successfully handles fine details around boundaries. We further evaluate the estimated surface normal for our deep MVS system. Results are shown in Table 2, where we again achieves state-of-the-art performance on ScanNet dataset [9]. Qualitative visualizations in Figure 9 clearly show the improvement of our method. Finally, we show visualizations of reconstructed models after TSDF fusion [8] in Figure 8. Our method produces visually more appealing reconstructions compared to strong baseline methods [23, 43].

Method	Abs Rel	Abs Diff	Sq Rel	RMSE	RMSE log	$\delta < 1.25$	$\delta < 1.25^2$	$\delta < 1.25^3$
MVDepth [47]	0.1053	0.1987	0.0634	0.3026	0.1490	0.8817	0.9723	0.9924
MVDepth (FT)	0.1014	0.1891	0.0476	0.2850	0.1390	0.8930	0.9764	0.9941
GP-MVS [21]	0.0920	0.2283	0.0644	0.4436	0.1560	0.8918	0.9629	0.9918
GP-MVS (FT)	0.0787	0.2008	0.0518	0.4009	0.1394	0.9134	0.9643	0.9931
NeuralRGBD [32]	0.0871	0.1710	0.0409	0.2693	0.1324	0.9150	0.9785	0.9925
CNM [33]	0.1119	0.2101	0.0510	0.2970	0.1485	0.8686	0.9724	0.9930
DPSNet [23]	0.1164	0.1992	0.0606	0.3065	0.1602	0.8569	0.9575	0.9884
DPSNet (FT)	0.0910	0.1807	0.0410	0.2697	0.1291	0.9008	0.9787	0.9952
NAS [29]	0.0795	0.1597	0.0323	0.2357	0.1112	0.9284	0.9862	0.9966
DELTAS [43]	0.0738	0.1380	0.0245	0.2051	0.1021	0.9473	0.9890	0.9976
Ours	0.0665	0.1281	0.0240	0.1995	0.0990	0.9489	0.9896	0.9978

Table 1. Quantitative comparisons between our method and state-of-the-art deep MVS methods on ScanNet dataset [9]. All the methods use sequences of length 3 and fixed reference interval 20 for testing, except GP-MVS [21] and NeuralRGBD [32], which directly use the whole sequence. Since [47, 21, 23] were not initially trained on ScanNet, we also report the results after finetuning, denoted as "FT".



(a) Image (b) Groundtruth (c) Ours (d) DELTAS [43] (e) NAS [29] (f) DPSNet (FT) [23] (g) MVDepth (FT) [47] (h) N-RGBD [32] Figure 7. Qualititative results of multi-view depth estimation on ScanNet [9]. **Better viewed when zoomed in.**

Method	Mean	Median	11.25°	22.5°	30°
CNM [33]	27.92	22.12	27.43	52.16	63.44
NAS [29]	24.12	18.02	31.59	60.20	69.45
Ours	22.30	16.75	34.80	64.39	75.11

Table 2. Quantitative comparisons of surface normal estimation between our method and state-of-the-art methods [33, 29].

5.4. Ablation Studies

We perform several ablation studies to further understand the behaviors of the proposed solver. All ablation studies are conducted on ScanNet [9]. To study the contribution of the depth-normal solver with its post process-

Method	Abs Rel	Abs Diff	Sq Rel	RMSE	$\delta < 1.25$
MVDepth [47]	0.0885	0.1467	0.0314	0.2313	0.9184
GP-MVS [21]	0.1087	0.1514	0.0827	0.2873	0.9170
N-RGBD [32]	0.0995	0.1530	0.0352	0.2361	0.9233
CNM [33]	0.1350	0.1873	0.0484	0.2619	0.8667
DPSNet [23]	0.0771	0.1290	0.0234	0.2045	0.9401
NAS [29]	0.0732	0.1241	0.0198	0.1893	0.9576
DELTAS [43]	0.1065	0.1528	0.0299	0.2138	0.9156
Ours	0.0698	0.1130	0.0194	0.1770	0.9681

Table 3. Quantitative comparisons between our method and stateof-the-art deep MVS methods on RGB-D Scenes V2 dataset [31].

ing and end-to-end joint training, we train a baseline with the same network architecture completely without the proposed solver, which does not benefit from joint training. For



Figure 8. Visualization of our reconstruction results compared to DELTAS [43] and DPS [23]. Better viewed when zoomed in.



Figure 9. Qualitative results of surface normal estimation on Scan-Net [9]. Refer to supp. for higher-resolution images.

comparison, we also apply the solver to refine the baseline's output with the hybrid confidence map at inference. Three conclusions can be acquired from the results shown in Table 4: 1) As a post-processing tool, the proposed solver consistently improves the output of the network no matter whether the network is jointly trained. 2) The joint training improves the quality of the estimated initial depth ($0.0735 \rightarrow 0.0711$). This is because joint training gets the network to focus more on reliable estimation on well textured areas, which largely releases the learning burden. 3) With end-to-end joint training, the performance gain of applying the solver module at inference increases. This agrees with our assumption that during end-to-end training, the network is guided to output initial geometry that well fits the subsequent solver module.

In addition, we also study the design choices on the confidence map. Table 5 shows the results. We observe that geometric confidence is more effective than the estimated deep confidence. Although the performance gain of including the deep confidence is not significant, the deep confidence is complete and will not be affected by occlusions or small overlaps, which can serve as a good supplements for geometric confidence in practice. Furthermore, we study the effects of different number of reference views during geometric consistency check and different number of iterations at inference. Table 6 shows the results, which indicate that improving the number of iterations consistently improve the final depth quality. However, as currently we take the minimum confidence over all reference frames to get a "strict" depth confidence, adding views will only result in slight performance gain. More advanced confidence based on multi-view depth fusion is left for future work.

Method	Abs Rel	RMSE	$\delta < 1.25$
Baseline	0.0735	0.2169	0.9359
Baseline + post	0.0720	0.2148	0.9378
Joint train w/o post	0.0711	0.2121	0.9384
Joint train + post	0.0665	0.1995	0.9489

Table 4. Ablation studies on the contribution of the solver used at post-processing and end-to-end joint training. "post" indicates using the solver to post-process the depth map with 10 iterations.

Method	Abs Rel	RMSE	$\delta < 1.25$
Joint train w/o post	0.0711	0.2121	0.9384
Groundtruth conf	0.0599	0.1931	0.9604
Deep conf only	0.0692	0.2040	0.9421
Geometric conf only	0.0667	0.2001	0.9486
Deep + geometric conf	0.0665	0.1995	0.9489

Table 5. Studies on the choices of different confidence maps.

View/Iter	1 iter	5 iters	10 iters	25 iters
1 view	0.0694	0.0680	0.0675	0.0670
2 views	0.0690	0.0672	0.0665	0.0659
4 views	0.0688	0.0669	0.0662	0.0655

Table 6. Ablation studies on the effects of different number of views used during geometric consistency check and different number of iterations at inference. Abs Rel is reported for each entry.

6. Conclusion

In this work, we present a deep MVS system with a novel confidence-based iterative depth-normal solver. We formulate an energy potential that favors locally planar surfaces and propose to perform iterative subproblem optimization over depths and surface normals in turn. Not only is the solver able to serve as a post-processing tool for planebased depth refinement and completion, but it is also in closed form such that it can be integrated into our deep MVS system with end-to-end joint training. Future directions include advanced confidence estimation and spatialtemporal plane-based propagation.

Acknowledgements: This work was partially supported by the Natural Science Foundation of China (61725204), BN-Rist and Tsinghua University (CS Dept) - DeepBlue Technology (Shanghai) Company Limited Joint Research Center for Machine Vision (JCMV).

References

- [1] http://cs.adelaide.edu.au/ hwong/doku.php?id=data. 6
- [2] Jonathan T Barron, Andrew Adams, YiChang Shih, and Carlos Hernández. Fast bilateral-space stereo for synthetic defocus. In *CVPR*, pages 4466–4474, 2015. 3
- [3] Jonathan T Barron and Ben Poole. The fast bilateral solver. In *ECCV*, pages 617–632. Springer, 2016. 3
- [4] Stan Birchfield and Carlo Tomasi. Multiway cut for stereo and motion with slanted surfaces. In *ICCV*, volume 1, pages 489–495, 1999. 2
- [5] Michael Bleyer and Margrit Gelautz. A layered stereo matching algorithm using image segmentation and global visibility constraints. *ISPRS Journal of Photogrammetry and remote sensing*, 59(3):128–150, 2005. 1, 2
- [6] Michael Bleyer, Christoph Rhemann, and Carsten Rother. Patchmatch stereo-stereo matching with slanted support windows. In *BMVC*, volume 11, pages 1–11, 2011. 2
- [7] Rui Chen, Songfang Han, Jing Xu, and Hao Su. Pointbased multi-view stereo network. In *ICCV*, pages 1538– 1547, 2019. 2
- [8] Brian Curless and Marc Levoy. A volumetric method for building complex models from range images. In *Proceedings* of the 23rd annual conference on Computer graphics and interactive techniques, pages 303–312, 1996. 6
- [9] Angela Dai, Angel X Chang, Manolis Savva, Maciej Halber, Thomas Funkhouser, and Matthias Nießner. Scannet: Richly-annotated 3d reconstructions of indoor scenes. In *CVPR*, pages 5828–5839, 2017. 6, 7, 8, 14, 16, 17
- [10] Arda Düzçeker, Silvano Galliani, Christoph Vogel, Pablo Speciale, Mihai Dusmanu, and Marc Pollefeys. Deepvideomvs: Multi-view stereo on video with recurrent spatiotemporal fusion. In *CVPR*, 2021. 2
- [11] David Eigen and Rob Fergus. Predicting depth, surface normals and semantic labels with a common multi-scale convolutional architecture. In *ICCV*, pages 2650–2658, 2015. 2
- [12] Olivier Faugeras and Renaud Keriven. Variational principles, surface evolution, PDE's, level set methods and the stereo problem. IEEE, 2002. 2
- [13] Silvano Galliani, Katrin Lasinger, and Konrad Schindler. Massively parallel multiview stereopsis by surface normal diffusion. In *ICCV*, pages 873–881, 2015. 2, 3
- [14] David Gallup, Jan-Michael Frahm, Philippos Mordohai, Qingxiong Yang, and Marc Pollefeys. Real-time planesweeping stereo with multiple sweeping directions. In *CVPR*, pages 1–8, 2007. 2
- [15] Gottfried Graber, Jonathan Balzer, Stefano Soatto, and Thomas Pock. Efficient minimal-surface regularization of perspective depth maps in variational stereo. In *CVPR*, pages 511–520, 2015. 2
- [16] Kwang Hee Lee and Sang Wook Lee. Deterministic fitting of multiple structures using iterative maxfs with inlier scale estimation. In *ICCV*, pages 41–48, 2013. 6
- [17] Heiko Hirschmuller. Stereo processing by semiglobal matching and mutual information. *PAMI*, 30(2):328–341, 2007. 2
- [18] Heiko Hirschmuller and Daniel Scharstein. Evaluation of stereo matching costs on images with radiometric differences. *PAMI*, 31(9):1582–1599, 2008. 2

- [19] Li Hong and George Chen. Segment-based stereo matching using graph cuts. In CVPR, volume 1, pages I–I, 2004. 1, 2
- [20] Asmaa Hosni, Christoph Rhemann, Michael Bleyer, Carsten Rother, and Margrit Gelautz. Fast cost-volume filtering for visual correspondence and beyond. *PAMI*, 35(2):504–511, 2012. 2
- [21] Yuxin Hou, Juho Kannala, and Arno Solin. Multi-view stereo by temporal nonparametric fusion. In *ICCV*, pages 2651–2660, 2019. 2, 7
- [22] Xiaoyan Hu and Philippos Mordohai. A quantitative evaluation of confidence measures for stereo vision. *PAMI*, 34(11):2121–2133, 2012. 2
- [23] Sunghoon Im, Hae-Gon Jeon, Stephen Lin, and In So Kweon. Dpsnet: End-to-end deep plane sweep stereo. In *ICLR*, 2019. 2, 5, 6, 7, 8, 13, 16
- [24] Mengqi Ji, Juergen Gall, Haitian Zheng, Yebin Liu, and Lu Fang. Surfacenet: An end-to-end 3d neural network for multiview stereopsis. In *ICCV*, pages 2307–2315, 2017. 2
- [25] Abhishek Kar, Christian Häne, and Jitendra Malik. Learning a multi-view stereo machine. In *NeurIPS*, 2017. 2
- [26] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014. 6
- [27] Andreas Kuhn, Shan Lin, and Oliver Erdler. Plane completion and filtering for multi-view stereo reconstruction. In *GCPR*, pages 18–32. Springer, 2019. 2
- [28] Andreas Kuhn, Christian Sormann, Mattia Rossi, Oliver Erdler, and Friedrich Fraundorfer. Deepc-mvs: Deep confidence prediction for multi-view stereo reconstruction. In *3DV*, pages 404–413. IEEE, 2020. 2
- [29] Uday Kusupati, Shuo Cheng, Rui Chen, and Hao Su. Normal assisted stereo depth estimation. In *CVPR*, pages 2189–2199, 2020. 2, 5, 7, 8, 13, 16, 17
- [30] Kiriakos N Kutulakos and Steven M Seitz. A theory of shape by space carving. *IJCV*, 38(3):199–218, 2000. 2
- [31] Kevin Lai, Liefeng Bo, and Dieter Fox. Unsupervised feature learning for 3d scene labeling. In *ICRA*, pages 3050–3057. IEEE, 2014. 6, 7
- [32] Chao Liu, Jinwei Gu, Kihwan Kim, Srinivasa G Narasimhan, and Jan Kautz. Neural rgb (r) d sensing: Depth and uncertainty from a video camera. In *CVPR*, pages 10986–10995, 2019. 2, 7, 16
- [33] Xiaoxiao Long, Lingjie Liu, Christian Theobalt, and Wenping Wang. Occlusion-aware depth estimation with adaptive normal constraints. In *ECCV*, pages 640–657. Springer, 2020. 2, 3, 7, 8, 17
- [34] Luca Magri and Andrea Fusiello. T-linkage: A continuous relaxation of j-linkage for multi-model fitting. In CVPR, pages 3954–3961, 2014. 6
- [35] Zak Murez, Tarrence van As, James Bartolozzi, Ayan Sinha, Vijay Badrinarayanan, and Andrew Rabinovich. Atlas: Endto-end 3d scene reconstruction from posed images. In ECCV, 2020. 2
- [36] Vaishakh Patil, Wouter Van Gansbeke, Dengxin Dai, and Luc Van Gool. Don't forget the past: Recurrent depth estimation from monocular video. *IEEE Robotics and Automation Letters*, 5(4):6813–6820, 2020. 2

- [37] Matteo Poggi and Stefano Mattoccia. Learning from scratch a confidence measure. In *BMVC*, 2016. 2
- [38] Xiaojuan Qi, Renjie Liao, Zhengzhe Liu, Raquel Urtasun, and Jiaya Jia. Geonet: Geometric neural network for joint depth and surface normal estimation. In *CVPR*, pages 283– 291, 2018. 2
- [39] Andrea Romanoni and Matteo Matteucci. Tapa-mvs: Textureless-aware patchmatch multi-view stereo. In *ICCV*, pages 10413–10422, 2019. 2
- [40] Johannes L Schonberger and Jan-Michael Frahm. Structurefrom-motion revisited. In *CVPR*, pages 4104–4113, 2016.
 6
- [41] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *ECCV*, pages 501–518. Springer, 2016. 2, 6
- [42] Akihito Seki and Marc Pollefeys. Patch based confidence prediction for dense disparity map. In *BMVC*, volume 2, page 4, 2016. 2
- [43] Ayan Sinha, Zak Murez, James Bartolozzi, Vijay Badrinarayanan, and Andrew Rabinovich. Deltas: Depth estimation by learning triangulation and densification of sparse points. In ECCV, 2020. 2, 6, 7, 8, 16
- [44] Jian Sun, Nan-Ning Zheng, and Heung-Yeung Shum. Stereo matching using belief propagation. *PAMI*, 25(7):787–800, 2003. 2
- [45] Julien Valentin, Adarsh Kowdle, Jonathan T Barron, Neal Wadhwa, Max Dzitsiuk, Michael Schoenberg, Vivek Verma, Ambrus Csaszar, Eric Turner, Ivan Dryanovski, et al. Depth from motion for smartphone ar. ACM Transactions on Graphics (ToG), 37(6):1–19, 2018. 2
- [46] George Vogiatzis, Philip HS Torr, and Roberto Cipolla. Multi-view stereo via volumetric graph-cuts. In *CVPR*, volume 2, pages 391–398, 2005. 2
- [47] Kaixuan Wang and Shaojie Shen. Mvdepthnet: Real-time multiview depth estimation neural network. In *3DV*, pages 248–257. IEEE, 2018. 2, 7, 16
- [48] Qingshan Xu and Wenbing Tao. Multi-scale geometric consistency guided multi-view stereo. In CVPR, pages 5483– 5492, 2019. 2, 3
- [49] Qingshan Xu and Wenbing Tao. Planar prior assisted patchmatch multi-view stereo. In AAAI, volume 34, pages 12516– 12523, 2020. 2
- [50] Koichiro Yamaguchi, Tamir Hazan, David McAllester, and Raquel Urtasun. Continuous markov random fields for robust stereo estimation. In *ECCV*, pages 45–58. Springer, 2012. 1, 2
- [51] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *ECCV*, pages 767–783, 2018. 2
- [52] Wei Yin, Yifan Liu, Chunhua Shen, and Youliang Yan. Enforcing geometric constraints of virtual normal for depth prediction. In *ICCV*, pages 5684–5693, 2019. 2, 3
- [53] Kuk-Jin Yoon and In So Kweon. Adaptive support-weight approach for correspondence search. *PAMI*, 28(4):650–656, 2006. 2

- [54] Fisher Yu, Vladlen Koltun, and Thomas Funkhouser. Dilated residual networks. In CVPR, pages 472–480, 2017. 4, 5
- [55] Yinda Zhang and Thomas Funkhouser. Deep depth completion of a single rgb-d image. In *CVPR*, pages 175–185, 2018.
 3

Appendix

In this document, we provide a list of supplementary materials that accompany the main paper.

A. Detailed Derivations of the Proposed Solver

A.1. Preliminaries

As discussed in the main paper, we solve the depth map and normal map with two separate suboptimization steps with respect to the total energy. Each step contains a plane-based propagation with slanted planes. Recall that P(x, d, n) denotes the slanted plane at pixel coordinate xgenerating by spanning a plane from the corresponding 3D points recovered from d and x and its surface normal n. In practice we parameterize the normal to be n = (a, b, -1), which enables closed-form computation in the normal update step (N-step). Let (p, q, z) denotes the 3D coordinate of the points recovered from d and $x = (u, v)^T$ at the frame coordinate system:

$$\begin{bmatrix} p \\ q \\ z \end{bmatrix} = K^{-1} \begin{bmatrix} x \\ 1 \end{bmatrix} d = \begin{bmatrix} \tilde{u} \\ \tilde{v} \\ 1 \end{bmatrix} d, \tag{8}$$

where K is the camera intrinsic parameter, $x = (u, v)^T$ is the 2D pixel coordinate and $(\tilde{u}, \tilde{v}, 1)^\top = K^{-1}(u, v, 1)^\top$ is the normalized homogeneous coordinate. Then, the plane equation of $P(x_i, d_i, n_i)$ with $n_i = (a_i, b_i, -1)$ and the recovered 3D points (p_i, q_i, z_i) can be written as follows:

$$a_i(p-p_i) + b_i(q-q_i) - (z-z_i) = 0.$$
 (9)

At plane-based propagation, the propagated depth $d_{i\to j}$ $(d_{j\to i})$ is computed by projecting the slanted plane at i (j) onto the pixel j (i). We give the derivation of $d_{i\to j}$ here:

$$a_i(p_j - p_i) + b_i(q_j - q_i) - (z_j - z_i) = 0$$
(10)

$$\Leftrightarrow a_i(\tilde{u}_j d_{i \to j} - p_i) + b_i(\tilde{v}_j d_{i \to j} - q_i) - (d_{i \to j} - z_i) = 0$$
(11)

$$\Leftrightarrow d_{i \to j} = \frac{a_i p_i + b_i q_i - z_i}{a_i \tilde{u}_j + b_i \tilde{v}_j - 1} = \frac{a_i \tilde{u}_i + b_i \tilde{v}_i - 1}{a_i \tilde{u}_j + b_i \tilde{v}_j - 1} d_i \qquad (12)$$

Here $d_{i \rightarrow j}$ is non-linearly dependent upon the depth d_i and the surface normal $n_i = (a_i, b_i, -1)$. As there exist secondorder terms $a_i d_i$ and $b_i d_i$ in the denominator, a quadratic energy over both d_i and n_i is infeasible even when the algebraic formulation is employed. Thus, closed-form solution cannot be acquired when the data term and planebased structural term $E_{i\rightarrow j}^3$ are jointly optimized over d_i and n_i . This motivates us to employ iterative suboptimization in the solver to acquire close-form solution, which can further benefit our deep MVS system with end-to-end joint training, as discussed in the main paper.

Before introducing the details of the two update steps of the proposed solver, let us take a step further on the formulation of jointly solving depths and surface normals. We want to note that it is possible to formulate closed-form solution by substitution of variables when only the surface normal data term is employed. This can be achieved by parameterizing the plane equation in Eq. (9) as $a_ip+b_iq-(z-t_i)=0$, where $t_i = z_i - a_ip_i - b_iq_i = (1-a_i\tilde{u}_i - b_i\tilde{v}_i)d_i$. When the depth data term is not included, by employing the algebraic form we can get a 3x3 linear system with respect to a_i , b_i and t_i . However, we empirically observe that the depth data term is extremely beneficial in practice.

A.2. Closed-form Solution

As discussed in the main paper, we employ suboptimization over the depth map and the surface normal map iteratively. This enables closed-form solution in both steps.

Depth Update (D-step). At the depth update step (D-step), we fix the surface normal map and solve for the optimal depth map d^* . L2 distance between the optimized depth and the propagated depth $d_{j\rightarrow i}$ from neighboring pixels are used in the plane-based structural term. The objective is written as follows (Eqs. (4)(5) in the main paper):

$$\min_{d} E_{total} = \min_{d} E_d$$

$$E_d = \alpha \sum c_i (d_i - \hat{d}_i)^2 + \sum \sum c_i w_{ii} (d_i - d_{i \to i})^2.$$
(13)

$$c_d = \alpha \sum_i c_i (d_i - d_i)^2 + \sum_i \sum_{j \in N(i)} c_j w_{ij} (d_i - d_{j \to i})^2.$$
(14)

As discussed in the main paper, we assume fixed neighborhoods to enable parallelization of the solver. Thus, the propagated depth $d_{j\rightarrow i}$ is the projection of the plane $P(x_i, \hat{d}_i, \hat{n}_j)$ at pixel *i*:

$$d_{j \to i} = \frac{\hat{a}_j \tilde{u}_j + \hat{b}_j \tilde{v}_j - 1}{\hat{a}_j \tilde{u}_i + \hat{b}_j \tilde{v}_i - 1} \hat{d}_j$$
(15)

Set the first-order derivative to zero we can easily derive the optimal depth d_i^* for each pixel:

$$d_{i}^{*} = \frac{\alpha c_{i} \hat{d}_{i} + \sum_{j \in N(i)} c_{i} w_{ij} d_{j \to i}}{\alpha c_{i} + \sum_{j \in N(i)} c_{i} w_{ij}}.$$
 (16)

Surface Normal Update (N-step). At the surface normal update step, we fix the depth map and solve for the optimal surface normal n^* . The objective is written as follows (Eqs. (6)(7) in the main paper):

 $^{{}^{3}}E_{j \rightarrow i}$ does not include the surface normal n_{i} in its formulation, and thus can only be used when only depth map is required to be solved.

rameterization:

$$\min_{n} E_{total} = \min_{n} E_n \tag{17}$$

$$E_{n} = \alpha \sum_{i} c_{i} ||n_{i} - \hat{n}_{i}||^{2} + \sum_{i} \sum_{j \in N(i)} c_{j} w_{ij} D_{n}(d_{j}, P(x_{i}, d_{i}, n_{i})).$$
(18)

Here D_n is a distance function defined over d_j and the plane $P(x_i, d_i, n_i)$ being optimized. Note that because $d_{i \rightarrow j}$ is non-linearly dependent over n_i as shown in Eq. (12) in this supplementary material, we cannot directly use L2 distance as in the D-step. Instead, we employ the algebraic formulation of the plane equation and directly formulate the distance function D_n as the square of the LHS of Eq. (10) in this supplementary material:

$$D_n(d_j, P(x_i, d_i, n_i)) = [a_i(p_j - p_i) + b_i(q_j - q_i) - (z_j - z_i))]^2$$
(19)

Because the depth map is fixed in the N-step, the only two unknown variables are a_i and b_i , which represent the surface normal $n_i = (a_i, b_i, -1)$. Note that this parameterization is feasible because all visible surfaces are facing the position where the camera center locates. However, numerical problems may occur when there exist ill-posed cases with surfaces that are nearly parallel to the corresponding camera rays. Thus, we clip the absolute value of the solved a_i and b_i with a threshold 20.0. This operation is empirically crucial to stabilize the end-to-end training process.

By setting the first-order derivatives to zero we can get a 2x2 linear system over a_i^* and b_i^* , where the optimal surface normal n_i^* is parameterized with $n_i^* = (a_i^*, b_i^*, -1)$. The coefficients are listed as follows:

$$\begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix} \begin{bmatrix} a_i^* \\ b_i^* \end{bmatrix} = \begin{bmatrix} B_1 \\ B_2 \end{bmatrix}$$
(20)

$$A_{11} = \alpha c_i + \sum_{j \in N(i)} c_j w_{ij} (p_j - p_i)^2$$
(21)

$$A_{22} = \alpha c_i + \sum_{j \in N(i)} c_j w_{ij} (q_j - q_i)^2$$
(22)

$$A_{12} = A_{21} = \sum_{j \in N(i)} c_j w_{ij} (p_j - p_i) (q_j - q_i)$$
(23)

$$B_1 = \alpha c_i \hat{a}_i + \sum_{j \in N(i)} c_j w_{ij} (p_j - p_i) (z_j - z_i)$$
(24)

$$B_2 = \alpha c_i \hat{b}_i + \sum_{j \in N(i)} c_j w_{ij} (q_j - q_i) (z_j - z_i)$$
(25)

In the final step, normalization is applied on the output surface normal to get the final prediction with standard pa-

$$n = \left(\frac{a}{\sqrt{a^2 + b^2 + 1}}, \frac{b}{\sqrt{a^2 + b^2 + 1}}, -\frac{1}{\sqrt{a^2 + b^2 + 1}}\right). \tag{26}$$

A.3. Extension: Incorporating Customized Depth Confidence and Normal Confidence

While we present all the derivations with a single confidence map c_i for simplicity, a separate depth confidence map c_d and surface normal confidence map c_n can be used in practice to further improve the solver. Employing those separate customized confidence maps will change the data term (Eq. (2) in the main paper) into the formulation below:

$$E_{data} = \sum_{i} c_{d,i} (d_i - \hat{d}_i)^2 + \sum_{i} c_{n,i} ||n_i - \hat{n}_i||^2, \quad (27)$$

where $c_{d,i}$ and $c_{n,i}$ denotes the per-pixel depth and surface normal confidence respectively. Furthermore, the confidence used in the plane-based structural term can also be monitored by the separate confidence maps.

Separate Confidence Maps used in the D-step. When the depth confidence and normal confidence are used jointly, the final formulation of E_d becomes the form as follows:

$$E_{d} = \alpha \sum_{i} c_{d,i} (d_{i} - \hat{d}_{i})^{2} + \sum_{i} \sum_{j \in N(i)} c_{d,j} c_{n,j} w_{ij} (d_{i} - d_{j \to i})^{2}.$$
(28)

Note that both the depth confidence $c_{d,j}$ and normal confidence $c_{n,j}$ are used in the plane-based structural term. This is due to the fact that the quality of the propagated depth $d_{j\rightarrow i}$ depends on both \hat{d}_j and \hat{n}_j .

Separate Confidence Maps used in the N-step. As for the surface normal update step, because $E_{i\rightarrow j}$ is employed, \hat{n}_j is non-relevant to the computation so the plane-based structural term will not be affected by $c_{n,j}$. However, since the depth map is fixed here and the computation of $d_{i\rightarrow j}$ depends on the quality of \hat{d}_i , we multiply the plane-based structural term with the depth confidence of the studied pixel $c_{d,i}$. The final formulation of E_n is written as follows:

$$E_{n} = \alpha \sum_{i} c_{n,i} ||n_{i} - \hat{n}_{i}||^{2} + \sum_{i} c_{d,i} \sum_{j \in N(i)} c_{d,j} w_{ij} D_{n}(d_{j}, P(x_{i}, d_{i}, n_{i})).$$
⁽²⁹⁾

Closed-form solution can be easily derived for the modified objectives in Eq. (28) and Eq. (29) following the previous discussion. In our implementation, the depth confidence map and normal confidence map are jointly predicted by the costvolume based neural networks. The supervision is acquired by computing the relative depth error and normal angle error, as described in the main paper. At inference, the hybrid confidence map combining the deep depth confidence and the geometric confidence is used as the depth confidence map, while the normal confidence only employs the deep normal confidence prediction.

B. Network Architectures

We follow prior works [23, 29] to design the initial depth, surface normal, and confidence estimation networks.

For the depth and normal branches, we use the same network architectures as [29] without the consistency module. Specifically, the target and reference images are first encoded to get the feature maps. Then we use plane sweeping with 64 hypothesis planes to build the feature cost volume. The 3D CNNs and 2D context CNNs are applied on the cost volume to aggregate and regularize the cost information. The soft argmin is used to regress the final depth values from the final cost volume. The cost volume information is also utilized for the multi-view surface normal estimation. The intermediate cost volume features are concatenated with the world coordinates of every voxel, and then transformed by several 3D CNNs to get 8 cost volume slices. Each slice is processed by 7 shared layers of 2D convolutions of dilated 3×3 kernels. The output of all slices are summed and normalized to get the final surface normal. We recommend the reader to refer to [29] for more details.

For the confidence branch, we input multiple sources to the network to better estimate the confidence. To predict the depth confidence, we utilize target image features, homography-warped reference image features using currently predicted depth, cost volume features before softmax, and the predicted depth. As illustrated in Figure 10, these inputs are processed by three mini-branches, and each minibranch contains two or three layers of 3×3 convolutions. To be specific, the first branch (target image feature + predicted depth) consists of two 3×3 convolution layers whose output channels are both 16. The second branch (target image feature + warped reference image feature) has also two 3×3 convolution layers whose output channels are both 32. The third branch (cost volume feature) consists of three 3×3 convolution layers with [64, 32, 1] as each layer's output channel number. The outputs from three mini-branches are then concatenated, and five dilated 3×3 convolution layers with output channels [64, 64, 64, 32, 1] and dilations [1, 2, 4, 1, 1], followed by the final sigmoid activation, are applied to jointly predict the depth confidence. For surface normal confidence, we use target image features, intermediate features from the previously discussed 8 cost volume slices, and the predicted surface normal, as shown in Figure 11.



Figure 10. The inputs are processed with three mini-branches (with 2 or 3 layers of 3x3 convolutions), and then jointly fed into 5 dilated convolutions followed by final sigmoid activation to regress the depth confidence map.



Figure 11. The inputs are processed with two mini-branches (with 2 layers of 3x3 convolutions), and then jointly fed into 5 dilated convolutions followed by final sigmoid activation to regress the normal confidence map.

These inputs are processed by two mini-branches, then concatenated and used to jointly predict the surface normal confidence. The first branch consists of two 3×3 convolution layers with both 16 output channels, and the second branch has also two 3×3 convolution layers with both 32 output channels. Then the outputs from these two mini-branches are concatenated and processed by five dilated 3×3 convolution layers with output channels [64, 64, 64, 32, 1] and dilations [1, 2, 4, 1, 1], plus final sigmoid activation, to get the confidence map for surface normal prediction.

C. System Training

Our proposed deep MVS system is trained with two stages. In the first stage, we train the initial depth, normal, and confidence estimation network for 15 epochs. The training loss for the initial depth, normal and confidence estimation, denoted as L_{net} , includes three parts: $L_{net} = w_d L_d + w_n L_n + w_c L_c$, where w_d, w_n, w_c are three hyper-parameters to balance different loss terms. L_d and L_n are formulated as smoothed L1 loss, and L_c employs cross-entropy loss. L_c consists of depth confidence loss L_{cd} and surface normal confidence loss L_{cn} , and is computed as $L_c = L_{cd} + L_{cn}$.

In the second stage, we finetune the network jointly with the proposed iterative depth normal solver for 10 epochs. During the training, We iteratively refine the depth map and surface normal map for 5 times to balance the computation and effectiveness. We apply the depth and surface normal loss on both the initial predictions and the final solved geometry. Thus the total loss for end-to-end training is $L_{total} = \lambda L_{net} + L_{solver}$, where $L_{solver} =$ $w_{d'}L_{d'} + w_{n'}L_{n'}$ is the depth and normal loss for the solved geometry. $L_{d'}$ and $L_{n'}$ are smoothed L1 loss for the refined depth and normal, respectively. $w_{d'}$ and $w_{n'}$ are hyperparameters. λ is the hyper-parameter for weighting these two losses L_{net} and L_{solver} . We choose to include the L_{net} in the end-to-end training to regularize the initial depth and normal predictions, which in practice stabilizes the joint training.

D. Implementation Details

The hyperparameters are heuristically selected without much tuning. We set loss weights $[\lambda, w_d, w_n, w_c, w_{d'}, w_{n'}]$ to [0.7, 1.0, 3.0, 0.2, 1.0, 3.0]. The scaling factors γ_1, γ_2 for depth and normal confidence groundtruth are both set to 5.0 in the training. The spatial and color weights for bilateral affinity σ_x^2, σ_c^2 are set to 2.5 and 25.0, and the weight α for depth and normal data term of the energy is set to 1.0. These hyper-parameters are fixed during both training and inference.

E. Notations

We provide a notation cheat sheet in Table 7, which describes the relevant notations used in the main paper and this supplementary material.

F. Additional Visualization

We provide additional visualizations for both depth and surface normal estimation in Figure 12 and 13. All samples are from the official test split of ScanNet [9].

Notations	Descriptions	Appearances
E_{total}	The total energy potential for the solver	Sec. 3.1, Sec. 3.2, Sec. A.2
E_{data}	The data term in the total energy potential	Sec. 3.1, Sec. A.3
E_{plane}	The plane-based structural term in the total energy potential	Sec. 3.1
α	The hyperparameter to balance the data term and structural term	Sec. 3.1, Sec. 3.2, Sec. A.2, Sec. A.3, Sec. E
x, d, n, c	2D coordinate, depth, surface normal, and confidence	Sec. 3.1, Sec. 3.2, Sec. A.1, Sec. A.2, Sec. A.3
x_i, d_i, n_i, c_i	Per-pixel 2D coordinate, depth, surface normal, and confidence	Sec. 3.1, Sec. 3.2, Sec. A.1, Sec. A.2
\hat{d}_i, \hat{n}_i	Initial per-pixel depth and surface normal	Sec. 3.1, Sec. 3.2, Sec. A.2, Sec. A.3
P(x,d,n)	The plane generated by current 2D coordinate x , depth d , and normal n	Sec. 3.1, Sec. 3.2, Sec. A.1, Sec. A.2, Sec. A.3
$E_{i \to i}, E_{i \to j}$	The plane-based structural term defined in two directions	Sec. 3.1, Sec. 3.2, Sec. A.1, Sec. A.3
$d_{j \to i}$	The projection of the plane $P(x_j, d_j, n_j)$ at pixel <i>i</i>	Sec. 3.1, Sec. 3.2, Sec. A.1, Sec. A.2, Sec. A.3
$d_{i \rightarrow j}$	The projection of the plane $P(x_i, d_i, n_i)$ at pixel j	Sec. 3.1, Sec. 3.2, Sec. A.1, Sec. A.2, Sec. A.3
w_{ij}	Edge-aware bilateral affinity between pixel i and j	Sec. 3.1, Sec. 3.2, Sec. A.2, Sec. A.3
Ii	RGB value at pixel <i>i</i>	Sec. 3.1
σ_x, σ_c	The hyperparameters for the spatial term and color term in the bilateral affinity	Sec. 3.1, Sec. E
E_d	The minimized objective in the D-step	Sec. 3.2, Sec. A.2, Sec. A.3
N(i)	The defined neighborhoods of pixel <i>i</i>	Sec. 3.2, Sec. A.2, Sec. A.3
E_n	The minimized objective in the N-step	Sec. 3.2, Sec. A.2, Sec. A.3
D_n	The distance function between the depth d and slanted plane P used in the N-step	Sec. 3.2, Sec. A.2, Sec. A.3
a, b	Components of the parameterized surface normal, $n = (a, b, -1)$	Sec. 3.2, Sec. A.1, Sec. A.2
c_{dqt}, c_{nqt}	Groundtruth (GT) confidence maps for depth and surface normal	Sec. 4.1
e_{rel}, e_{ang}	Relative depth error and normal angle error between predictions and groundtruths	Sec. 4.1
γ_1, γ_2	Hyperparameters used in the computation of the GT depth and normal confidence	Sec. 4.1, Sec. E
p, q, z	The 3D coordinate of the unprojected point	Sec. A.1, Sec. A.2
K	Camera intrinsic matrix	Sec. A.1
u, v	2D pixel coordinates, $x = (u, v)^T$	Sec. A.1
\tilde{u}, \tilde{v}	$(\tilde{u}, \tilde{v}, 1)^T = K^{-1}(u, v, 1)^T$	Sec. A.1
d^*	The optimal depth map in the D-step	Sec. A.2
n^*	The optimal surface normal map in the N-step	Sec. A.2
a^*, b^*	Components of the optimal surface normal n^* , $n^* = (a^*, b^*, -1)$	Sec. A.2
$A_{11}, A_{12}, A_{21}, A_{22},$	2*Coefficients used in the N-step computation	2*Sec. A.2
B_1, B_2		
c_d, c_n	The depth confidence map and surface normal confidence map	Sec. A.3
$c_{d,i}, c_{n,i}$	Per-pixel depth confidence and surface normal confidence	Sec. A.3
L_{net}	Training loss for the initial depth, surface normal, and confidence network	Sec. D
L_d, L_n, L_c	Losses of initial depth, surface normal and confidence	Sec. D
w_d, w_n, w_c	Loss weights of initial depth, surface normal and confidence	Sec. D, Sec. E
L_{cd}, L_{cn}	Confidence loss for depth and normal, $L_c = L_{cd} + L_{cn}$	Sec. D
L_{total}	Total training loss in the end-to-end training	Sec. D
L _{solver}	The Loss defined over the solved geometry $L_{solver} = w_{d'}L_{d'} + w_{n'}L_{n'}$	Sec. D
$L_{d'}, L_{n'}$	Losses of the solved depth and surface normal	Sec. D
$w_{d'}, w_{n'}$	Loss weights of the solved depth and surface normal	Sec. D, Sec. E
λ	The hyperparameter used to balance L_{net} and L_{solver}	Sec. D, Sec. E

Table 7. Notations used in the main paper and supplementary material.



(a) Image (b) Groundtruth (c) Ours (d) DELTAS [43] (e) NAS [29] (f) DPSNet (FT) [23] (g) MVDepth (FT) [47] (h) N-RGBD [32] Figure 12. More qualititative results of depth estimation on ScanNet [9]. Better viewed when zoomed in.

