

# Dual-Level Precision Edges Guided Multi-View Stereo with Accurate Planarization

Kehua Chen, Zhenlong Yuan, Tianlu Mao\*, Zhaoqi Wang

Institute of Computing Technology, Chinese Academy of Sciences  
{chenkehua23s,yuanzhenlong21b,ltm,zqwang}@ict.ac.cn

## Abstract

The reconstruction of low-textured areas is a prominent research focus in multi-view stereo (MVS). In recent years, traditional MVS methods have performed exceptionally well in reconstructing low-textured areas by constructing plane models. However, these methods often encounter issues such as crossing object boundaries and limited perception ranges, which undermine the robustness of plane model construction. Building on previous work (APD-MVS), we propose the DPE-MVS method. By introducing dual-level precision edge information, including fine and coarse edges, we enhance the robustness of plane model construction, thereby improving reconstruction accuracy in low-textured areas. Furthermore, by leveraging edge information, we refine the sampling strategy in conventional PatchMatch MVS and propose an adaptive patch size adjustment approach to optimize matching cost calculation in both stochastic and low-textured areas. This additional use of edge information allows for more precise and robust matching. Our method achieves state-of-the-art performance on the ETH3D and Tanks & Temples benchmarks. Notably, our method outperforms all published methods on the ETH3D benchmark.

## Introduction

Multi-view stereo (MVS) is a classical computer vision task aimed at reconstructing the dense 3D geometry of objects or scenes from images taken from multiple viewpoints. This technique has significant applications in areas such as cultural heritage preservation, virtual reality, augmented reality, and autonomous driving. In recent years, MVS methods has advanced significantly, benefiting from diverse datasets (Schops et al. 2017; Knapitsch et al. 2017) and various algorithms (Wang et al. 2023; Wu et al. 2024), leading to substantial improvements in reconstruction performance. Despite these advancements, MVS still faces challenges in handling low-textured and stochastic textured areas.

MVS methods can be roughly categorized into traditional methods (Galliani, Lasinger, and Schindler 2015; Schönberger et al. 2016; Xu and Tao 2019) and learning-based methods (Yao et al. 2018; Gu et al. 2020). Traditional methods have the advantages of stronger general-

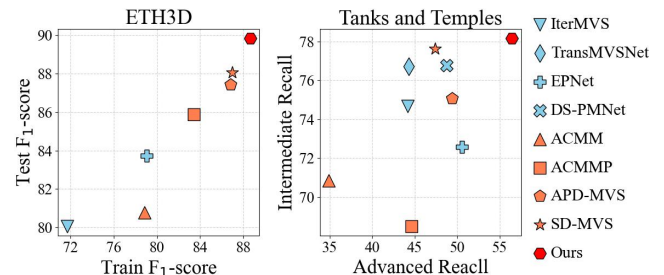


Figure 1: Comparison with the SOTA traditional methods and learning-based methods. Our method achieves the best  $F_1$ -score on ETH3D and the best recall on Tanks & Temples.

ization capabilities and lower memory consumption compared to learning-based methods. Additionally, there has been more research in recent traditional MVS methods addressing the low-textured issue. Recent mainstream methods are based on PatchMatch (PM), which matches fixed-size patch in the reference image with patches in the source images using a plane hypothesis (including depth and normal). Since fixed-size patch struggle to extract appropriate feature information in low-textured areas, many works have further extended and optimized this method. For example, (Liao et al. 2019; Xu and Tao 2019) leverage multi-scale information, while others (Xu and Tao 2020) use triangular plane priors to guide plane hypotheses in low-textured areas. (Xu et al. 2022) combines these approaches to enhance reconstruction performance. Subsequent methods (Zhang et al. 2022; Tian et al. 2023) refine the construction of triangular planes, while others (Romanoni and Matteucci 2019; Kuhn, Lin, and Erdler 2019) employ image segmentation and RANSAC algorithm to determine plane models.

One notable method, APD-MVS (Wang et al. 2023), introduces adaptive patch deformation. This method classifies pixels into reliable and unreliable based on matching ambiguity. For each unreliable pixel, it searches for a number of reliable pixels in the surrounding area. The RANSAC algorithm is then used to estimate the best-fitting plane from these reliable pixels, selecting the most fitting ones as anchors to assist in the matching of unreliable pixel. Compared to previous methods, this approach is more flexible and significantly enhances the robustness of the plane model.

\*Corresponding author.




Figure 2: Top: depth maps for scenes crossing object boundaries. Bottom: normal maps for limited perception range. Comparison of APD-MVS (middle) and our method (right).

Although these methods significantly improve reconstruction in low-textured areas, they still face issues with increased scene complexity, as shown in Fig. 2 and Fig. 3. One common issue is the plane model crossing object boundaries, causing depth confusion between foreground objects and the background. Another is the potential for errors in plane model construction due to limited perception range. For example, APD-MVS considers only the nearest reliable pixels when constructing planes, sometimes selecting locally optimal pixels, which leads to noticeable deviations between the final plane model and the ground truth.

To address these issues, we drew inspiration from learning-based MVS methods (Zhang et al. 2023; Li et al. 2024), which utilize RGB images for adaptive sampling. We posit that fully leveraging image information, particularly edge information, is crucial since areas delineated by edges often approximate planar shapes. Based on this premise, we integrated dual-level precision edge information into the adaptive patch deformation. Dual-level precision edges are derived from two edge detection approaches: fine edges, which are precise but incomplete, and coarse edges, which capture more actual object boundaries but with less accuracy. Specifically, we use the Canny operator for fine edges and a segmentation scheme from TSAR-MVS (Yuan et al. 2024b) for coarse edges. Fine edges constrain point selection during plane construction, while coarse edges expand the perception range for selecting anchors. Thereby providing more effective support for the matching of unreliable pixels. Furthermore, since reliable pixels serve as the basis for selecting anchors and are still processed using conventional PM in APD-MVS, we utilize fine edge information to improve the hypotheses sampling strategy, optimizing the hypotheses for reliable pixels.

The aforementioned strategy has shown marked performance but is ineffective in stochastic textured areas, such as lawns, which often contain numerous erroneous edges. Therefore, we further investigated the matching cost calculation. The deformable patch, consisting of an unreliable pixel’s patch and the anchors’ patches, is used to evaluate the plane hypothesis of the the unreliable pixel. Its matching cost is calculated as the weighted sum of the matching costs of both the unreliable pixel’s patch and the anchors’ patches. However, the conventional fixed-size patch




Figure 3: Comparison of plane construction between APD-MVS (middle) and our method (right), with visualizations of fine edges (top right) and coarse edges (bottom right).

matching method, specifically the matching of the unreliable pixel’s patch, remains part of this process, affecting stability, especially in stochastic textured areas. To address this, we propose adjusting patch sizes with anchors to effectively identify unreliable pixels and applying edge information constraints to prevent crossing object boundaries. This enables more robust matching cost calculations.

We integrated the aforementioned concepts into the Dual-Level Precision Edges Guided Multi-View Stereo with Accurate Planarization (DPE-MVS). In summary, our contributions can be summarized as follows:

- We propose a dual-level precision edge-guided planar model construction strategy, providing more effective support for the matching of unreliable pixels.
- We propose a sampling strategy guided by fine edges, which can enhance conventional PM for reliable pixels.
- We introduce an adaptive patch size adjustment approach that enables more robust matching cost calculation for unreliable pixels.
- Extensive experiments validate the effectiveness of our proposed method, demonstrating state-of-the-art performance on the ETH3D and Tanks & Temples benchmarks.

## Related Work

**Traditional Methods** Traditional MVS methods can be roughly categorized into four types: voxel-based methods (Vogiatis et al. 2007), surface iterative optimization methods (Cremers and Kolev 2010), patch-based methods (Furukawa and Ponce 2009), and depth map-based methods (Bleyer, Rhemann, and Rother 2011). Among these, depth map-based methods have become the most popular choice in recent years due to their simplicity, flexibility, and robust performance. Many outstanding works within this category are PM-based methods. Recently, methods such as ACMM (Xu and Tao 2019), ACMP (Xu and Tao 2020), and ACMMP (Xu et al. 2022) have introduced pyramid structures, geometric consistency, and triangular plane priors into MVS. Subsequently, HPM-MVS (Ren et al. 2023) proposed non-local sampling to escape local optima and used a KNN-based approach to optimize plane prior model construction. APD-MVS (Wang et al. 2023) introduced adaptive patch deformation and an NCC-based matching metric

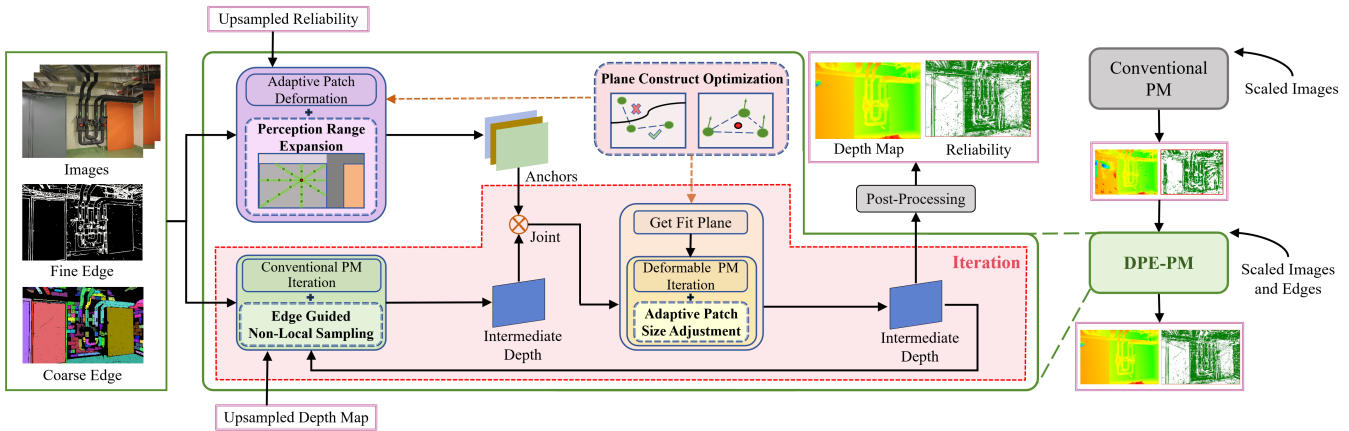


Figure 4: Overview. DPE-MVS adopt a pyramid structure, with the two coarsest scales displayed on the right side of the figure, and the middle illustrating the details of our proposed DPE-PM. Iterations at finer scales use DPE-PM to update the depth map.

to determine the reliability of pixel depth values. Methods like TAPA-MVS (Romanoni and Matteucci 2019) and PCF-MVS (Kuhn, Lin, and Erdler 2019) incorporated superpixel segmentation and the RANSAC algorithm, while TSAR-MVS (Yuan et al. 2024b) further combined the Roberts operator with Hough line detection to segment large low-textured areas, though these methods tend to over-segment. SD-MVS (Yuan et al. 2024a) used SAM for semantic segmentation to achieve adaptive sampling. However, SAM’s inference speed is slow and it may produce errors with unseen scenes. Additionally, SAM struggles to distinguish different surfaces of the same object, making it unsuitable for constructing plane models.

**Learning-based Methods** MVSNet (Zhang et al. 2023) pioneered the use of deep learning for depth map-based MVS methods. CasMVSNet (Gu et al. 2020) introduced a cascade structure, accelerating the evolution of learning-based MVS methods. These methods, benefiting from convolution operations, have significantly larger receptive fields compared to traditional methods. Works like AA-RMVSNet (Wei et al. 2021) and TransMVSNet (Ding et al. 2022) further expanded the receptive field. Additionally, improvements in depth sampling have been made by PatchMatchNet (Wang et al. 2021) and DS-PMNet (Li et al. 2024), which proposed adaptive hypothesis propagation, and N2MVSNet (Zhang et al. 2023), which introduced adaptive non-local sampling and RGB-guided depth refinement. Learning-based methods have stronger feature perception and often outperform traditional methods with sufficient data. However, creating high-quality datasets remains challenging, limiting practical use.

## Method

Given a set of images  $\{I_i\}_{i=1}^N$  and the corresponding camera parameters  $\{\mathbf{P}_i\}_{i=1}^N$ , our task is to estimate depth maps for each image. This section provides a brief overview of the key points of APD-MVS, followed by a detailed explanation of our method.

## Review of APD-MVS

APD-MVS classifies pixels as reliable or unreliable based on matching ambiguity and introduces deformable PM. Reliable pixels are processed using conventional PM, while unreliable pixels are handled using deformable PM.

Conventional PM consists of four basic steps: random initialization, hypothesis propagation, multi-view matching cost evaluation, and refinement. First, each pixel is randomly initialized with a plane hypothesis. Second, hypotheses are sampled from neighboring pixels within a fixed range. Third, matching costs from multiple views are integrated to select the best hypothesis. Fourth, new hypotheses are generated through perturbation and random generation to diversify the solution space, and the best one is selected. The last three steps iterate multiple times.

Deformable PM differs from conventional PM in propagation and matching cost calculation. For each unreliable pixel, anchors are identified through preprocessing. In propagation, sampled hypotheses include anchor hypotheses and plane hypothesis generated using RANSAC on the anchors. In matching cost calculation, the deformable patch is constructed by combining the unreliable pixel’s patch with the anchors’ patches for matching, the formula as follows:

$$m_D(\mathbf{p}, \theta_p, \mathbf{S}) = \lambda m(\mathbf{p}, \theta_p, \mathbf{B}_p) + \frac{1 - \lambda}{|\mathbf{S}|} \sum_{\mathbf{s} \in \mathbf{S}} m(\mathbf{s}, \theta_p, \mathbf{B}_s), \quad (1)$$

where  $\lambda$  is a weight value,  $\mathbf{p}$  represents the unreliable pixel,  $\mathbf{S}$  denotes the set of anchors,  $\theta_p$  is the plane hypothesis for pixel  $\mathbf{p}$ , and  $\mathbf{B}$  represent the fixed-size patch. The function  $m$  represents the conventional matching cost, while  $m_D$  denotes the matching cost of the deformable patch.

## Overview of Our Method

Our method adopts the APD-MVS framework, and an overview is illustrated in Fig. 4. Each image is sequentially taken as the reference image  $I_{ref}$ , with the other images as source images  $I_{src}$  to guide the reference image’s depth map recovery. We construct an  $L$  layer pyramid structure through

scale downsampling, with the  $L$ -th layer as the coarsest scale and the 1st layer as the original image. The initial depth map at the coarsest scale layer is obtained using conventional PM, and perform post-processing to determine the reliability of each pixel. At a finer scale layer  $l$ , fine and coarse edges are extracted from  $I_{ref}$  at the corresponding scale. The depth map and reliability from layer  $l + 1$  are upsampled. These inputs are then used to update the depth map and reliability for this layer using DPE-PM.

In DPE-PM, the first stage is to obtain anchors for each unreliable pixel. In adaptive patch deformation, we propose Perception Range Expansion to search for a wide range of relevant reliable pixels and use RANSAC to filter out the anchors. The second stage is to iteratively update the depth map. In each iteration, the hypotheses for reliable pixels are first updated using conventional PM with our Edge Guided Non-Local Sampling. Subsequently, the hypotheses for unreliable pixels are updated: new plane hypothesis are generated using RANSAC based on the anchors, followed by deformable PM with our Adaptive Patch Size Adjustment. In the previously mentioned RANSAC applications, our Plane Construction Optimization was consistently utilized to obtain accurate plane models. Similarly, post-processing is used to determine pixel reliability after obtaining the depth map. The DPE-PM process is repeated at finer scales until the depth map for the first layer is obtained. Finally, the depth maps are fused to generate a point cloud.

In the following sections, we will provide the details of our method. First, edge extraction will be introduced, followed by the improvements related to reliable pixels, and finally, the improvements related to unreliable pixels.

### Extracting Edge Cues

Obtaining edge information is essential for our method. We first use the Canny edge detector to extract fine edges, setting the upper and lower thresholds to the median of the image grayscale multiplied by  $(1 \pm \sigma)$ . Then, coarse edges are extracted using the Roberts operator and Hough line detection, similar to TSAR-MVS. Fine edges help accurately locate the boundaries of foreground objects but are often not closed, making it difficult to determine whether a pixel is in a low-textured region. Coarse edges can segment the image, with larger regions indicating low-textured areas, but have a higher false detection rate and less precise boundary localization. The complementary information from fine and coarse edges is crucial for achieving our research objectives.

### Edge Guided Non-Local Sampling

Reliable pixels are the foundation for constructing the plane model, and the accuracy of their hypotheses is crucial. Propagation samples hypotheses from neighboring pixels to build the solution space. According to (Ren et al. 2023; Zhou et al. 2021), repetitive hypotheses often occur within the local range in low-textured areas, whereas local sampling preserves fine details in small objects. To expand the solution space while retaining details, we propose two new sampling schemes: progressive non-local sampling and edge-guided extended sampling, as shown in Fig. 5. For fine edge pixels, we apply only progressive non-local sampling. For non-fine

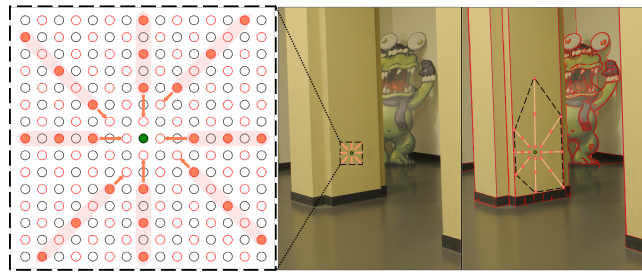


Figure 5: Edge Guided Non-Local Sampling: This process includes two sampling schemes: progressive non-local sampling (left) and edge-guided extended sampling (right).

edge pixels, both sampling schemes are applied, and the resulting samples are compared to retain the superior ones.

Progressive non-local sampling excludes sampling points within a radius  $\xi$  during the PM iteration process. The radius  $\xi$  gradually decreases as the iterations progress, following the formula  $\xi = \max(1, 5 - 2 \times t_{iter})$ . We utilize a red-black checkerboard pattern for pixel division (Xu and Tao 2019). Each of the eight sampling areas adopts a strip format, where the radius  $\xi$  offsets the starting position of the strip. Each area contains 11 samples with a step size of 2. The sample with the minimum multi-view matching cost is selected from each area, yielding the optimal samples  $\{\theta_i^{pn}\}_{i=1}^8$ .

Edge-guided extended sampling follows the same fundamental scheme as progressive non-local sampling, but differs in the number of samples and the sampling step size. In this scheme, both the number of samples  $k$  and the step size  $s$  are adaptively adjusted based on the distance  $D_{fe}$  to the nearest fine edge in the corresponding strip direction. A threshold, defined as  $\Lambda_{fe} = \frac{imagewidth}{30 \times 4^l}$ , is used to prevent excessively large sampling distances. The specific adjustments are calculated as follows:

$$D'_{fe} = \min(D_{fe}, \Lambda_{fe}), k = \left\lfloor \frac{D'_{fe}}{2} \right\rfloor, s = \left\lfloor \frac{D'_{fe}}{k} \right\rfloor, \quad (2)$$

subject to  $11 \leq k \leq 22$ . Similarly, optimal samples  $\{\theta_i^{eg}\}_{i=1}^8$  are obtained. These samples are then compared with  $\{\theta_i^{pn}\}_{i=1}^8$  by recalculating the matching costs on the patch of the pixel being processed, with the better sample for each area direction being selected.

### Accurate Plane Model Construction

For unreliable pixels, the anchors filtered out by planar RANSAC in adaptive patch deformation significantly impact depth estimation.

**Perception Range Expansion** To address the issue of limited perception range, we segment the image into distinct regions using coarse edges and extend the search for reliable pixels in low-textured regions. A region is considered low-textured if its pixel count exceeds  $\frac{imagearea}{256 \times 4^l}$ . Let  $\mathcal{E}$  denote the set of coarse edge pixels in a specific direction relative to the pixel  $\mathbf{p}$ , the boundary in that direction is defined as:

$$\mathcal{B}(\mathbf{p}) = \arg \max_{\mathbf{q} \in \mathcal{E}} \|\mathbf{q} - \mathbf{p}\| \quad s.t. \quad \mathcal{C}(\mathbf{q}) = 1, \quad (3)$$

where  $\mathbb{C}(\mathbf{q}) = 1$  indicates that  $\mathbf{q}$  is connected to the region where  $\mathbf{p}$  is located. This approach determines boundaries in the eight-connected directions for each pixel, effectively filtering out redundant coarse edge pixels within the region.

Subsequently, the eight directions are grouped into four pairs of opposites. Each pair has a search limit of  $2\eta$  reliable pixels, distributed between directions based on boundary distances. For example, in the up-down direction pair, the number of pixels allocated for searching is determined using the following formula:

$$n_u = \left\lfloor \frac{2\eta \cdot D_{ce}^u}{D_{ce}^u + D_{ce}^d} \right\rfloor, n_d = 2\eta - n_u, \quad (4)$$

where  $D_{ce}^u$  and  $D_{ce}^d$  represent the boundary distances in the upward and downward directions, respectively, with  $n_u$  is constrained to  $1 \leq n_u \leq 2\eta - 1$ . For each direction, starting from the unreliable pixel, we locate a set of equally spaced pixels  $\{\mathbf{s}_i\}_{i=1}^n$  along the line to the boundary. Let  $\mathcal{N}(\mathbf{s}_i)$  denote the nearest reliable pixel to  $\mathbf{s}_i$ . The set  $\{\mathcal{N}(\mathbf{s}_i)\}_{i=1}^n$  serves as the result of extended search in that direction.

We retain the reliable pixels search scheme from APD-MVS, which involves partitioning the search space centered on the unreliable pixel into  $\phi$  equal-angle sectors and locating the nearest reliable pixels within each sector. This scheme only searches the nearest reliable pixels and is applied to each unreliable pixel, while the extended search supplements it in low-textured areas.

**Plane Construction Optimization** RANSAC aims to find the best-fitting plane from 3D points. Using an unreliable pixel in a low-textured area as an example, obtain the set of 3D points  $\mathcal{X} = \{\mathbf{X}_i\}_{i=1}^{\phi+8\eta}$  for the searched reliable pixels. Three random points from  $\mathcal{X}$  are iteratively selected to construct a plane  $\pi$ , identifying the best fit as follows:

$$\pi^* = \arg \min_{\pi} \sum_{i=1}^{\phi+8\eta} \mathbb{I}[\delta(\mathbf{X}_i, \pi) < \epsilon]. \quad (5)$$

Here,  $\delta$  represents the residual,  $\epsilon$  is the outlier filtering threshold, and  $\mathbb{I}$  is an indicator function that equals 1 if the condition is met and 0 otherwise. If there are multiple  $\pi^*$ , select the one that minimizes  $\delta$  with the 3D point of the unreliable pixel. Let  $\pi$  be defined by  $m_1x + m_2y + m_3z + b = 0$ . Based on the projection relationship between  $\mathbf{X}_i$  and its pixel coordinate  $(u_i, v_i)$ , the fitting depth  $d_i$  is obtained as:

$$d_i = \frac{-b \cdot f_x \cdot f_y}{m_1 f_y (u_i - c_x) + m_2 f_x (v_i - c_y) + m_3 f_x f_y}, \quad (6)$$

where  $c_x, c_y, f_x$  and  $f_y$  are intrinsic camera parameters. For  $\mathbf{X}_i$ , we calculated  $\delta$  as  $|d_i - d_i^\theta|$ , where  $d_i^\theta$  is the depth hypothesis. Comparatively, conventional planar RANSAC calculates  $\delta$  by considering only the point-to-plane distance, without utilizing the projection relationship.

To mitigate the issue of crossing object boundaries, we imposed constraints on RANSAC's random point selection. Let  $\mathcal{F}$  denote the set of all fine edge pixels,  $\{\mathbf{q}_i, \mathbf{n}_i^\theta\}_{i=1}^3$  represent the pixels and their normal hypotheses corresponding

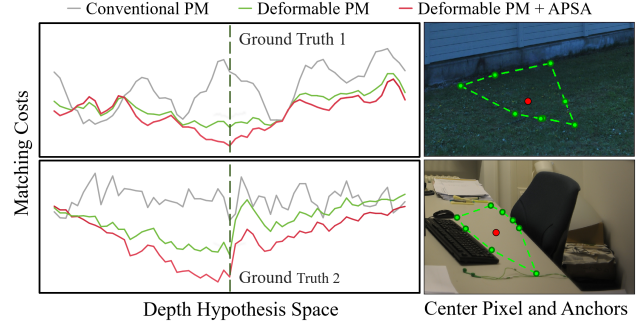


Figure 6: The right side presents examples of a stochastic textured area (top) and a low-textured area (bottom). The left side shows the corresponding cost profiles, highlighting the matching costs around the ground truth for various methods, including conventional PM, deformable PM, and deformable PM with our Adaptive Patch Size Adjustment.

to the selected three points. We present two conditions:

$$\forall i, j \in \{1, 2, 3\}, i < j \quad \begin{aligned} C_1 &: \overline{\mathbf{q}_i \mathbf{q}_j} \cap \mathcal{F} = \emptyset; \\ C_2 &: \mathbf{n}_i^\theta \cdot \mathbf{n}_j^\theta > \tau, \end{aligned} \quad (7)$$

where  $\overline{\mathbf{q}_i \mathbf{q}_j}$  represents the set of pixels along the line segment between  $\mathbf{q}_i$  and  $\mathbf{q}_j$ , and  $\tau$  is a threshold. Condition  $C_1$  is a mandatory requirement, while  $C_2$  is a priority condition that should be met when possible.

After RANSAC, select up to  $|\mathcal{S}|$  pixels corresponding to the 3D points  $\{\mathbf{X}_i | \delta(\mathbf{X}_i, \pi^*) < \epsilon\}$  that have smallest residuals to serve as anchors. The optimized RANSAC is also used to generate a new plane hypothesis with these anchors, a necessary step for the propagation in deformable PM.

Moreover, we observed that stochastic textured areas, lacking explicit repeating patterns (Efros and Leung 1999), often lead to numerous fine edges that do not effectively constrain the plane model. To address this, we propose assessing texture complexity per pixel. Let  $N_{fe}$  and  $N_{ce}$  denote the number of fine and coarse edge pixels, respectively, within an  $M \times M$  area centered on pixel  $\mathbf{p}$ . The probability that  $\mathbf{p}$  is in a stochastic textured area can be calculated as:

$$\Phi_p = \frac{\alpha \cdot N_{fe} + (1 - \alpha) \cdot N_{ce}}{M \times M}, \quad (8)$$

$$\mathcal{P}(\mathbf{p}|Z) = \frac{1}{1 + \exp[-\beta_1 \cdot (\Phi_p - \beta_2)]}, \quad (9)$$

where  $\Phi_p$  represent edge densities, and  $\alpha, \beta_1, \beta_2$  are empirically set.  $Z$  indicates that the pixel is in a stochastic textured area, and  $M$  is the same as the fixed patch size. We generate a random number  $r$  between 0 and 1. If  $r < \mathcal{P}(\mathbf{p}|Z)$ , then  $\mathbf{p}$  is considered to be in a stochastic textured area, rendering condition  $C_1$  inapplicable.

### Adaptive Patch Size Adjustment

When computing matching cost, the deformable patch employs fixed-size patches of the center unreliable pixel and its anchors, as detailed in Eq. 1. The cost profiles for different matching methods in stochastic and low-textured areas are

Method	Train						Test					
	2cm			10cm			2cm			10cm		
	Acc.	Comp.	F <sub>1</sub>	Acc.	Comp.	F <sub>1</sub>	Acc.	Comp.	F <sub>1</sub>	Acc.	Comp.	F <sub>1</sub>
PatchMatchNet	64.81	65.43	64.21	89.98	83.28	85.70	69.71	77.46	73.12	91.98	92.05	91.91
IterMVS	79.79	66.08	71.69	96.35	82.62	88.60	84.73	76.49	80.06	96.92	88.34	92.29
EPNet	79.36	79.28	79.08	94.33	93.69	93.92	80.37	<u>87.84</u>	83.72	93.72	96.82	95.20
GoMVS	81.22	77.65	79.16	97.11	89.62	93.08	86.85	85.50	85.91	97.23	95.02	96.02
ACMM	<b>90.67</b>	70.42	78.86	<u>98.12</u>	86.40	91.70	90.65	74.34	80.78	<u>98.05</u>	88.77	92.96
ACMMP	90.63	77.61	83.42	97.99	93.32	95.54	<u>91.91</u>	81.49	85.89	<u>98.05</u>	94.67	96.27
TSAR-MVS + MP.	89.67	84.39	86.88	<b>98.15</b>	96.50	97.31	88.14	88.11	88.02	97.42	97.44	<u>97.42</u>
APD-MVS	89.14	<u>84.83</u>	86.84	97.47	96.79	97.12	89.54	85.93	87.44	97.00	96.95	96.95
HPM-MVS	<u>90.66</u>	79.50	84.58	97.97	95.59	96.22	<b>92.13</b>	83.25	87.11	<b>98.11</b>	95.41	96.69
SD-MVS	89.91	84.31	86.96	97.84	<u>96.87</u>	<u>97.35</u>	88.96	87.49	<u>88.06</u>	97.37	<u>97.51</u>	97.41
DPE-MVS (ours)	89.81	<b>87.60</b>	<b>88.63</b>	97.99	<b>97.69</b>	<b>97.83</b>	90.53	<b>88.77</b>	<b>89.48</b>	97.64	<b>98.11</b>	<b>97.86</b>

Table 1: Quantitative results on ETH3D benchmark. Our method achieves the best completeness and F<sub>1</sub>-score.

---

### Algorithm 1: Adaptive Patch Size Adjustment

---

**Input** : the three anchors that define the plane for unreliable pixel  $\mathbf{p}$ , fine and coarse edges  
**Output**: appropriate center patch radius  $\gamma$  for  $\mathbf{p}$

- 1  $\gamma \leftarrow \lfloor \sqrt{\text{trianglearea}/2} \rfloor$ ;  
// the three anchors
- 2  $\mathbf{D} \leftarrow \{\omega \times D_{an}\}$ ;
- 3 **if**  $\mathbf{p}$  is not in a stochastic textured area **then**
  - // eight fine edges
  - 4  $\mathbf{D} \leftarrow \mathbf{D} \cup \{D_{fe}\}$ ;
  - 5 **if**  $\mathbf{p}$  is in a low-textured area **then**
    - // eight coarse edges
    - 6  $\mathbf{D} \leftarrow \mathbf{D} \cup \{D_{ce}\}$ ;
- 7  $\gamma \leftarrow \min\{\gamma \cup \mathbf{D}\}$ ;  
// fixed patch size
- 8 **if**  $\text{fixedradius} > \gamma$  **then**
- 9  $\gamma \leftarrow \text{fixedradius}$ ;
- 10 **else if**  $\mathbf{p}$  is in a stochastic textured area **then**
- 11  $\gamma \leftarrow 0$ ;
- 12 **return**  $\gamma$ ;

---

shown in Fig. 6. The gray cost profiles indicate the instability of the center patch during matching, suggesting the need for further improvement. Enlarging the center patch size may help mitigate this issue (Xu et al. 2020; Sun et al. 2022), but could lead to detail loss and challenges in selecting the size.

We propose a hybrid approach that combines patch enlargement with patch discarding. Roughly speaking, for pixels in stochastic textured areas, we tend to discard the center patch, while for those in low-textured areas, we tend to enlarge the center patch. The insight behind discarding the center patch is that anchors, selected for their planar characteristics, maintain stable normal estimation, whereas the center patch may introduce instability. The details of this approach are outlined in Algo. 1. It is posited that the triangular area formed by the three anchors, which are selected by optimized RANSAC to represent the points constructing the best-fitting plane, provides informative features for the center pixel. The appropriate patch radius is determined using

this triangular area with the surrounding edge information. The comparison between the green and red cost profiles in Fig. 6 illustrates the effectiveness of our approach.

## Experiments

### Datasets and Implementation Details

We evaluate our method on the ETH3D (Schops et al. 2017) and Tanks & Temples (Knapitsch et al. 2017) benchmarks, and conduct ablation experiments on the ETH3D training dataset to verify its effectiveness. We compare our method with SOTA learning-based methods including PatchMatchNet, AA-RMVSNet, IterMVS, TransMVSNet, EPNet, DS-PMNet, GoMVS and traditional MVS methods including ACMM, ACMMP, APD-MVS, HPM-MVS, TSAR-MVS, SD-MVS. To ensure fairness, we maintained the original APD-MVS parameters and depth map fusion procedures. The proposed parameter settings are:  $\{\sigma, \eta, \tau, \alpha, \beta_1, \beta_2, \omega\} = \{0.67, 4, 0.87, 0.5, 25, 0.35, 2.5\}$ .

### Evaluation on MVS Benchmarks

The quantitative results of ETH3D are presented in Tab. 1. Our method ranks **1st** in both completeness and F<sub>1</sub>-score, with accuracy nearly matching the previous best. Fig. 7 provides a qualitative analysis, confirming that our method significantly enhances the reconstruction of low-textured areas.

On the Tanks & Temples benchmark, We validated the generalization capability of our method. The quantitative results are shown in Tab. 2. Our method ranks **1st** in recall among all methods. In comparison with traditional methods, our method’s F<sub>1</sub>-score ranks **1st** and 2nd in the Intermediate and Advanced datasets, respectively. Compared to learning-based methods, our method is also competitive. It shows lower precision in the Advanced dataset, primarily due to the extensive low-textured areas. In these areas, our method prioritizes recall, with a slight trade-off in precision.

### Ablation Studies

In the ETH3D training dataset, we conduct ablation experiments to verify the effectiveness of each component in our proposed method, which includes Edge-Guided Non-Local Sampling (ES), Perceptual Range Expansion (PE), Plane




Figure 7: Qualitative results on ETH3D. Our method exhibits a significant advantage in reconstructing low-textured areas.

Method	Intermediate			Advanced		
	Pre.	Rec.	F <sub>1</sub>	Pre.	Rec.	F <sub>1</sub>
PatchmatchNet	43.64	69.37	53.15	27.27	41.66	32.31
AA-RMVSNet	52.68	75.69	61.51	37.46	33.01	33.53
IterMVS	47.53	74.69	56.94	28.70	44.19	34.17
TransMVSNet	55.14	76.73	63.52	33.84	44.29	37.00
EPNet	<b>57.01</b>	72.57	63.68	34.26	<u>50.54</u>	<u>40.52</u>
DS-PMNet	<u>56.02</u>	76.76	<b>64.16</b>	34.29	48.73	39.78
ACMM	49.19	70.85	57.27	35.63	34.90	34.02
ACMMP	53.28	68.50	59.38	33.79	44.64	37.84
TSAR-MVS + MP.	53.15	75.52	62.10	33.85	48.75	38.63
APD-MVS	55.58	75.06	63.64	33.77	49.41	39.91
HPM-MVS	51.58	76.92	61.39	<b>40.67</b>	45.42	<b>40.80</b>
SD-MVS	53.78	<u>77.63</u>	63.31	35.53	47.37	40.18
DPE-MVS (ours)	54.48	<b>78.16</b>	<u>63.98</u>	31.37	<b>56.45</b>	40.20

Table 2: Quantitative results on Tanks & Temples.

Construction Optimization (PO), and Adaptive Patch Size Adjustment (AA). Tab. 3 illustrates the effectiveness of each part. ES significantly expands the solution space for reliable pixel hypotheses, enhancing their accuracy and improving plane construction for unreliable pixels. PE, similar to ES, extends the search range for reliable pixels within adaptive patch deformation, helping to avoid the influence of locally optimal reliable pixels during plane model construction. PO optimizes the RANSAC algorithm, ensuring the quality of the final plane models. Together, ES, PE, and PO substantially improve completeness. AA adjusts the center patch size, enabling faster convergence and allowing hypotheses close to the correct solution to be refined with greater precision within limited iterations, thereby enhancing accuracy.

We also compare ES with the NESP module of HPM-MVS, integrating ES into ACMM, ACMP, and ACMMP, following (Ren et al. 2023). As shown in Tab. 4, our method outperforms others, except at the 1cm and 2cm thresholds in ACMM. The lack of plane priors in ACMM leads to erroneous depth estimations in low-textured areas, making it less suitable for our broader sampling range compared to NESP.

Further results are provided in the supplementary materi-

Settings				1cm			5cm		
ES	PE	PO	AA	Acc.	Comp.	F <sub>1</sub>	Acc.	Comp.	F <sub>1</sub>
				80.60	71.86	75.80	95.38	93.47	94.37
✓				79.74	74.15	76.67	94.85	93.97	94.38
	✓			80.16	73.08	76.27	95.30	93.87	94.54
✓	✓			80.19	74.71	77.13	95.43	94.36	94.81
✓	✓	✓		80.25	<b>75.79</b>	<u>77.79</u>	95.45	94.71	95.05
✓	✓		✓	80.69	74.48	<u>77.26</u>	95.81	94.77	95.26
✓	✓	✓	✓	<b>81.01</b>	<u>75.74</u>	<b>78.11</b>	<b>96.00</b>	<b>95.10</b>	<b>95.53</b>

Table 3: Results of our method with different settings on the ETH3D training dataset; baseline: APD-MVS.

Method	1cm	2cm	5cm	10cm	20cm
ACMM	67.58	78.86	87.68	91.70	94.41
w/. NESP	<b>70.70</b>	<b>81.01</b>	88.80	92.26	94.55
w/. ES (ours)	69.34	80.47	<b>89.17</b>	<b>92.83</b>	<b>95.07</b>
ACMP	68.72	79.79	88.32	92.03	94.43
w/. NESP	70.87	81.45	89.43	92.72	94.78
w/. ES (ours)	<b>72.45</b>	<b>83.01</b>	<b>90.70</b>	<b>93.72</b>	<b>95.53</b>
ACMMP	71.57	83.42	92.03	95.54	97.37
w/. NESP	74.54	85.33	93.25	96.45	97.99
w/. ES (ours)	<b>76.01</b>	<b>86.60</b>	<b>94.07</b>	<b>97.07</b>	<b>98.37</b>

Table 4: Comparison of F<sub>1</sub>-scores of different sampling strategies on the ETH3D training dataset.

als, including additional experimental details and comparative studies, extensive point cloud visualizations.

## Conclusion

In this paper, we propose the DPE-MVS method, which addresses issues in plane construction by introducing dual-level precision edge information. Our method significantly improves reconstruction of low-textured areas and excels in recovering stochastic textured areas. Experimental results demonstrate SOTA performance on ETH3D and Tanks & Temples. However, improvement is needed in handling fine scene details, where learning-based methods excel. Future work will integrate these methods for more accurate results.

## Acknowledgements

This work was supported in part by the Strategic Priority Research Program of the Chinese Academy of Sciences under Grant No.XDA0450203, in part by the Program of National Natural Science Foundation of China under Grant 62172392.

## References

- Bleyer, M.; Rhemann, C.; and Rother, C. 2011. Patchmatch stereo-stereo matching with slanted support windows. In *Bmvc*, volume 11, 1–11.
- Cremers, D.; and Kolev, K. 2010. Multiview stereo and silhouette consistency via convex functionals over convex domains. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(6): 1161–1174.
- Ding, Y.; Yuan, W.; Zhu, Q.; Zhang, H.; Liu, X.; Wang, Y.; and Liu, X. 2022. Transmvsnet: Global context-aware multi-view stereo network with transformers. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 8585–8594.
- Efros, A. A.; and Leung, T. K. 1999. Texture synthesis by non-parametric sampling. In *Proceedings of the seventh IEEE international conference on computer vision*, volume 2, 1033–1038. IEEE.
- Furukawa, Y.; and Ponce, J. 2009. Accurate, dense, and robust multiview stereopsis. *IEEE transactions on pattern analysis and machine intelligence*, 32(8): 1362–1376.
- Galliani, S.; Lasinger, K.; and Schindler, K. 2015. Massively parallel multiview stereopsis by surface normal diffusion. In *Proceedings of the IEEE international conference on computer vision*, 873–881.
- Gu, X.; Fan, Z.; Zhu, S.; Dai, Z.; Tan, F.; and Tan, P. 2020. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 2495–2504.
- Knapitsch, A.; Park, J.; Zhou, Q.-Y.; and Koltun, V. 2017. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics (ToG)*, 36(4): 1–13.
- Kuhn, A.; Lin, S.; and Erdler, O. 2019. Plane completion and filtering for multi-view stereo reconstruction. In *Pattern Recognition: 41st DAGM German Conference, DAGM GCPR 2019, Dortmund, Germany, September 10–13, 2019, Proceedings 41*, 18–32. Springer.
- Li, H.; Guo, Y.; Zheng, X.; and Xiong, H. 2024. Learning Deformable Hypothesis Sampling for Accurate PatchMatch Multi-View Stereo. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 3082–3090.
- Liao, J.; Fu, Y.; Yan, Q.; and Xiao, C. 2019. Pyramid multi-view stereo with local consistency. In *Computer Graphics Forum*, volume 38, 335–346. Wiley Online Library.
- Ren, C.; Xu, Q.; Zhang, S.; and Yang, J. 2023. Hierarchical prior mining for non-local multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 3611–3620.
- Romanoni, A.; and Matteucci, M. 2019. Tapa-mvs: Textureless-aware patchmatch multi-view stereo. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 10413–10422.
- Schönberger, J. L.; Zheng, E.; Frahm, J.-M.; and Pollefeys, M. 2016. Pixelwise view selection for unstructured multi-view stereo. In *Computer Vision—ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11–14, 2016, Proceedings, Part III 14*, 501–518. Springer.
- Schops, T.; Schonberger, J. L.; Galliani, S.; Sattler, T.; Schindler, K.; Pollefeys, M.; and Geiger, A. 2017. A multi-view stereo benchmark with high-resolution images and multi-camera videos. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3260–3269.
- Sun, S.; Liu, J.; Li, Y.; Ying, H.; Zhai, Z.; and Mou, Y. 2022. Adaptive pixelwise inference multi-view stereo. In *Thirteenth International Conference on Graphics and Image Processing (ICGIP 2021)*, volume 12083, 165–172. SPIE.
- Tian, Z.; Wang, R.; Wang, Z.; and Wang, R. 2023. HQP-MVS: High-Quality Plane Priors Assisted Multi-View Stereo for Low-Textured Areas. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Vogiatis, G.; Esteban, C. H.; Torr, P. H.; and Cipolla, R. 2007. Multiview stereo via volumetric graph-cuts and occlusion robust photo-consistency. *IEEE transactions on pattern analysis and machine intelligence*, 29(12): 2241–2246.
- Wang, F.; Galliani, S.; Vogel, C.; Speciale, P.; and Pollefeys, M. 2021. Patchmatchnet: Learned multi-view patchmatch stereo. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 14194–14203.
- Wang, Y.; Zeng, Z.; Guan, T.; Yang, W.; Chen, Z.; Liu, W.; Xu, L.; and Luo, Y. 2023. Adaptive patch deformation for textureless-resilient multi-view stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1621–1630.
- Wei, Z.; Zhu, Q.; Min, C.; Chen, Y.; and Wang, G. 2021. Aarmvsnet: Adaptive aggregation recurrent multi-view stereo network. In *Proceedings of the IEEE/CVF international conference on computer vision*, 6187–6196.
- Wu, J.; Li, R.; Xu, H.; Zhao, W.; Zhu, Y.; Sun, J.; and Zhang, Y. 2024. GoMVS: Geometrically Consistent Cost Aggregation for Multi-View Stereo. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 20207–20216.
- Xu, Q.; Kong, W.; Tao, W.; and Pollefeys, M. 2022. Multi-scale geometric consistency guided and planar prior assisted multi-view stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(4): 4945–4963.
- Xu, Q.; and Tao, W. 2019. Multi-scale geometric consistency guided multi-view stereo. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5483–5492.
- Xu, Q.; and Tao, W. 2020. Planar prior assisted patchmatch multi-view stereo. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, 12516–12523.

- Xu, Z.; Liu, Y.; Shi, X.; Wang, Y.; and Zheng, Y. 2020. Mar-mvs: Matching ambiguity reduced multiple view stereo for efficient large scale scene reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5981–5990.
- Yao, Y.; Luo, Z.; Li, S.; Fang, T.; and Quan, L. 2018. Mvs-net: Depth inference for unstructured multi-view stereo. In *Proceedings of the European conference on computer vision (ECCV)*, 767–783.
- Yuan, Z.; Cao, J.; Li, Z.; Jiang, H.; and Wang, Z. 2024a. SD-MVS: Segmentation-Driven Deformation Multi-View Stereo with Spherical Refinement and EM Optimization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, 6871–6880.
- Yuan, Z.; Cao, J.; Wang, Z.; and Li, Z. 2024b. Tsar-mvs: Textureless-aware segmentation and correlative refinement guided multi-view stereo. *Pattern Recognition*, 154: 110565.
- Zhang, W.; Gu, S.; Cai, Y.; and Wang, R. 2022. Multi-view depth estimation with color-aware propagation and texture-aware triangulation. In *2022 IEEE Asia-Pacific Conference on Image Processing, Electronics and Computers (IPEC)*, 168–174. IEEE.
- Zhang, Z.; Gao, H.; Hu, Y.; and Wang, R. 2023. N2mvsnet: Non-local neighbors aware multi-view stereo network. In *ICASSP 2023-2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5. IEEE.
- Zhou, L.; Zhang, Z.; Jiang, H.; Sun, H.; Bao, H.; and Zhang, G. 2021. DP-MVS: Detail preserving multi-view surface reconstruction of large-scale scenes. *Remote Sensing*, 13(22): 4569.