

# MSP-MVS: Multi-Granularity Segmentation Prior Guided Multi-View Stereo

Zhenlong Yuan<sup>1</sup>, Cong Liu<sup>2</sup>, Fei Shen<sup>3</sup>, Zhaoxin Li<sup>4, 5\*</sup>, Jinguo Luo<sup>2</sup>, Tianlu Mao<sup>1</sup>, Zhaoqi Wang<sup>1</sup>

<sup>1</sup>Institute of Computing Technology, Chinese Academy of Sciences

<sup>2</sup>Harbin Institute of Technology, Shenzhen

<sup>3</sup>Nanjing University of Science and Technology, Nanjing

<sup>4</sup>Agricultural Information Institute, Chinese Academy of Agricultural Sciences

<sup>5</sup>Key Laboratory of Agricultural Big Data, Ministry of Agriculture and Rural Affairs

yuanzhenlong21b@ict.ac.cn, liucong@stu.hit.edu.cn, feishen@njust.edu.cn, cszli@hotmail.com, 23s153135@stu.hit.edu.cn, {ltm, zqwang}@ict.ac.cn

## Abstract

Recently, patch deformation-based methods have demonstrated significant strength in multi-view stereo by adaptively expanding the reception field of patches to help reconstruct textureless areas. However, such methods mainly concentrate on searching for pixels without matching ambiguity (i.e., **reliable** pixels) when constructing deformed patches, while neglecting the deformation instability caused by unexpected edge-skipping, resulting in potential matching distortions. Addressing this, we propose MSP-MVS, a method introducing multi-granularity segmentation prior for edge-confined patch deformation. Specifically, to avoid unexpected edge-skipping, we first aggregate and further refine multi-granularity depth edges gained from Semantic-SAM as prior to guide patch deformation within depth-continuous (i.e., **homogeneous**) areas. Moreover, to address attention imbalance caused by edge-confined patch deformation, we implement adaptive equidistribution and disassemble-clustering of correlative reliable pixels (i.e., **anchors**), thereby promoting attention-consistent patch deformation. Finally, to prevent deformed patches from falling into local-minimum matching costs caused by the fixed sampling pattern, we introduce disparity-sampling synergistic 3D optimization to help identify global-minimum matching costs. Evaluations on ETH3D and Tanks & Temples benchmarks prove our method obtains state-of-the-art performance with remarkable generalization.

## Introduction

Multi-View Stereo (MVS) is a technique that employs images captured from different viewpoints to reconstruct 3D scenes or objects. Its diverse application spans across autonomous driving (Wei et al. 2023), augmented reality (Cao et al. 2020), and 3D printing (Zhu et al. 2021). Numerous ideas (Heise et al. 2013) and datasets (Yao et al. 2020) have significantly advanced its progress. However, it still faces the primary challenge of reconstructing large textureless areas.

The majority of traditional MVS methods are extended from the PatchMatch (PM) (Barnes et al. 2009) algorithm, which predicts depth hypothesis by computing the minimum matching cost between patch pairs and achieves iterative renewal through propagation and refinement. However, depth

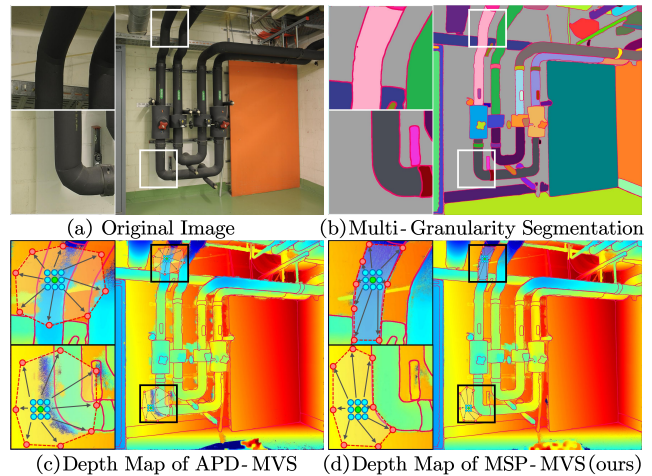


Figure 1: Comparative analysis between APD-MVS and our method. In (c) and (d), green, blue and red dots respectively denote the central pixel, conventional PM and deformable PM. Due to the lack of depth edge guidance in APD-MVS (c), its deformable PM occurs edge-skipping, thereby covering areas with depth-discontinuity. Differently, our method (d) leverages multi-granularity segmentation image (b) as prior to guide deformable PM within homogeneous areas.

hypotheses may exhibit ambiguity in textureless areas when patches lack reliable pixel correspondence. To tackle this, numerous methods attempt to introduce technologies such as cascade architecture (Liao et al. 2019), confidence estimators (Li et al. 2020), superpixel planarization (Yuan et al. 2024b) and planar priors (Ren et al. 2023), etc. Yet these methods either require complex and time-consuming post-processing or result in excessively smooth surface details.

Differently, learning-based methods employ deep neural networks to create learnable 3D cost volumes, thereby enabling the extraction of high-dimensional features and thus providing a wider receptive field. Although typical methods employ deformable convolutions (Ding et al. 2022) or adopt a coarse-to-fine strategy (Su and Tao 2023) to broaden the receptive field, their unaffordable time and memory consumption make them unsuitable for high-resolution imagery.

\*Corresponding Author.

Copyright © 2025, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Addressing this, others attempt to employ lightweight GRU (Yan et al. 2020) or leverage vision transformers (Cao, Ren, and Fu 2024) to compact 3D volumes and enhance prediction quality. However, these networks struggle with generalization when faced with scenarios different from their training datasets, posing challenges for practical applications.

Building on advances in traditional and learning-based MVS methods, APD-MVS (Wang et al. 2023) offers an innovative solution to tackle matching ambiguity. As shown in Fig. 1 (c), centered on the **unreliable** green pixel (i.e., pixels with matching ambiguity), the algorithm first segments its surrounding area into eight sectors at fixed angles, then selects a reliable red pixel within each sector as **anchor** pixel to form the deformable PM, which will substitute the blue neighbors in conventional PM to help reconstruct textureless areas. The APD-MVS assumes that all areas in deformable PM share the same depths and normals as the central pixel, referred to as **homogeneous**. Homogeneous areas are notably characterized by depth continuity, which makes them distinguishable at depth boundaries. Nonetheless, in complicated scenes, shadows and occlusions may cause deformable PM erroneously crossing depth edges and consequently being selected into **heterogeneous** areas (i.e., areas with depth-discontinuity), leading to potential reconstruction failures.

Therefore, it is crucial to fully exploit depth edges to aid in patch deformation. However, extracting comprehensive depth edges is challenging. Previous methods (Su and Tao 2023; Xue et al. 2019; Bao et al. 2014) attempt to extract RGB edges for reconstruction, while the illumination-sensitive and shadow-unaware RGB edges are impractical for complex scenes. Differently, others (Shvets et al. 2024; Yuan et al. 2024a) leverage semantic information for edge extraction, yet single semantic information fails to accurately capture comprehensive depth edges in varying scenes. Consequently, a method that effectively integrates semantic information from multiple granularities is urgently needed.

Addressing this, we introduce MSP-MVS, which innovatively proposes multi-granularity segmentation prior to facilitate edge-confined patch deformation. Specifically, we first adopt Semantic-SAM (Li et al. 2023) for panoramic segmentation to extract semantic information of images with varying granularities, then we aggregate them to obtain integrated edges that can distinguish different homogeneous areas. To better align the integrated edges with actual depth edges, we further apply edge refinement to generate fine-grained depth edges, which are then adopted as the multi-granularity segmentation prior to constrain deformed patches within homogeneous areas, as shown in Fig. 1 (d).

Although edge-confined patch deformation effectively removes deformed patches in heterogeneous areas, it inevitably causes an imbalanced distribution for the remaining patches. Addressing this, we introduce an adaptive equidistribution strategy, which adaptively divides the homogeneous area of each unreliable pixel into eight sectors based on the number of surrounding reliable pixels. Moreover, we further employ a disassemble-clustering strategy to quadruple anchor numbers within each sector and group selected anchors into eight anchor clusters that serve as new deformed patches with attention consistency for matching cost.

While the adaptive equidistribution and disassemble-clustering strategy achieves attention-consistent patch deformation, the fixed sampling pattern within patches may still cause the matching cost to fall into a local-minimum. Therefore, we introduce a disparity-sampling synergistic 3D optimization method, which synergistically and iteratively optimizes sampling pixels and disparities in 3D space to enable deformed patches identify their global minimal matching costs. In summary, our contributions are as follows:

- We develop MSP-MVS, which leverages Semantic-SAM to aggregate and further refine multi-granularity depth edges as prior for edge-confined patch deformation.
- We propose adaptive equidistribution for sector division and disassemble-clustering for anchor clustering to facilitate attention-consistent patch deformation.
- We present the disparity-sampling synergistic 3D optimization which optimizes both sampling pixels and disparities to help identify global-minimum matching costs.
- We achieve state-of-the-art performance on both the ETH3D dataset and Tanks & Temples dataset.

## Related Work

### Traditional Patchmatch-based Methods

The fundamental concept of PatchMatch (Barnes et al. 2009) employs random initialization, propagation and refinement to match optimal patch pairs between images. PMS (Bleyer et al. 2011) pioneers in introducing this in stereo vision, facilitating the emergence of numerous future works. For acceleration, Gipuma (Galliani, Lasinger, and Schindler 2015) proposes checkerboard diffusion propagation to achieve algorithm deployment on GPUs. To reconstruct textureless areas, ACMM (Xu and Tao 2019) presents both multi-scale architecture and multi-view consistency, improved by ACMMP (Xu et al. 2022) which introduces triangularization plane priors for post-processing. Differently, TAPA-MVS (Romanoni and Matteucci 2019), PCF-MVS (Kuhn, Lin, and Erdler 2019) and TSAR-MVS (Yuan et al. 2024b) attempt to employ superpixel planarization to reconstruct textureless areas, while their performance is severely affected by limited superpixel size. Moreover, HPM-MVS (Ren et al. 2023) introduces non-local extensible sampling patterns to avoid local optimal solutions, but its performance is limited by fixed patch receptive fields. Addressing this, SD-MVS (Yuan et al. 2024a) employs instance segmentation for patch deformation within instances. However, its results severely depend on low-precision segmentation results. Moreover, APD-MVS (Wang et al. 2023) separates patches of unreliable pixels into several outward-spreading anchors with high reliability, while the absence of edge constraint may cause unexpected edge-skipping and deformation instability.

### Learning-based MVS Methods

MVSNet (Yao et al. 2018) is the first to utilize differentiable 3D cost volumes to construct deep neural networks, thereby becoming the foundation for the abundant subsequent methods. To reduce memory, Cas-MVSNet (Yang et al. 2020)

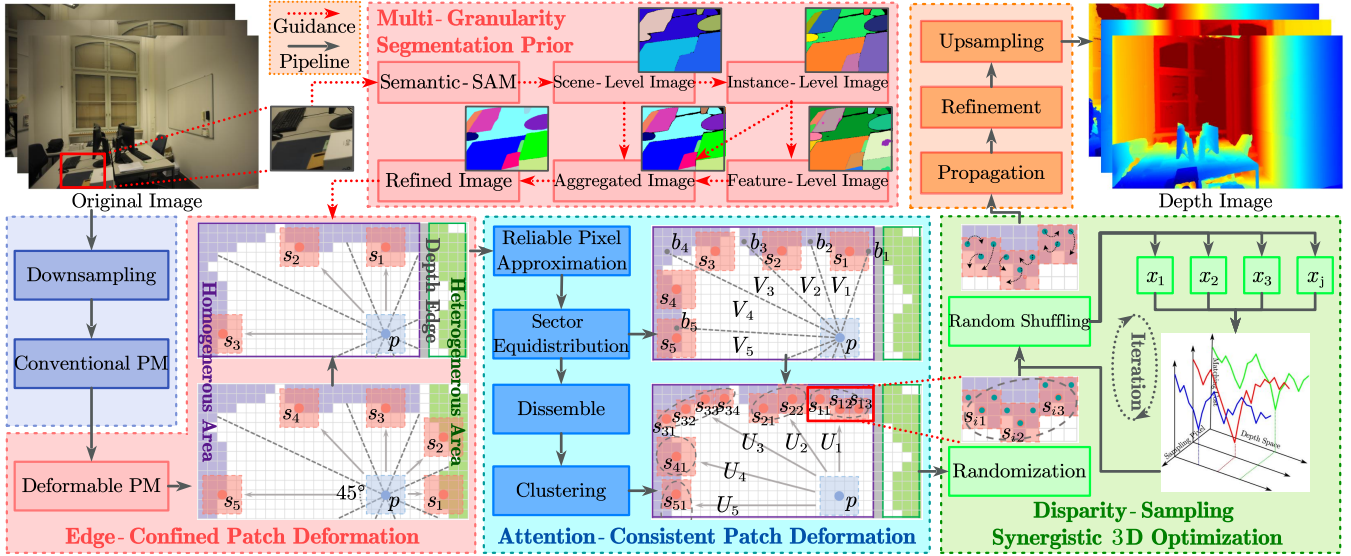


Figure 2: Pipeline of MSP-MVS. We first adopt Semantic-SAM to obtain multi-granularity segmentation images. We then aggregate and further refine these images as multi-granularity segmentation prior to facilitate edge-confined patch deformation. Subsequently, we propose adaptive equidistribution for sector division and disassemble-clustering strategy for anchor clustering to promote attention-consistent patch deformation. Additionally, we introduce disparity-sampling synergistic 3D optimization to help deformed patches identify their global-minimal matching costs. After several iterations we obtain final depth images.

adopts a coarse-to-fine strategy to retain depths across multiple scales. IterMVS-LS (Wang et al. 2022a) leverages a lightweight probability estimator to encode depth distribution during regularization. For feature extraction, AAR-MVSNet (Wei et al. 2021) employed deformable convolutions to achieve adaptive feature aggregation. Differently, MVSTER (Wang et al. 2022b) leverages the transformer structure to introduce the multi-head attention mechanism. Concerning geometrical topology, Geo-MVSNet (Zhang et al. 2023b) adopts a two-branch geometry fusion network to enhance geometry perception. Regarding surface topology, RA-MVSNet (Zhang et al. 2023a) associates hypothesis planes with surface patches to enhance perception field. Moreover, Raymvsnet++ (Shi et al. 2023) applies an epipolar transformer to enhance feature aggregation along camera rays. GoMVS (Wu et al. 2024) integrates a geometrically consistent propagation module to refine cost aggregation. Differently, our method targets on aggregating multi-granularity depth edges for edge-confined patch deformation. Despite this, most of them still suffer from either unaffordable memory costs or limited generalization.

### Preliminaries

The conventional PM method which utilizes fixed patch for reconstruction should be detailed at first. Given a series of input images  $\mathcal{I} = \{I_n \mid n = 1 \cdots N\}$  and corresponding camera parameters  $\mathcal{P} = \{K_n, R_n, C_n \mid n = 1 \cdots N\}$ , each image is selected from  $\mathcal{I}$  to be the reference image  $I_i$ , with its depth map sequentially reconstructed through pairwise matching with the remaining images ( $\mathcal{I} - I_i$ ). Here  $K$ ,  $R$  and  $C$  respectively represent the intrinsic parameters, the rotation matrix and the camera center. Specifically, given refer-

ence image  $I_i$  and source image  $I_j$ , the homography matrix  $H_{ij}$  for pixel  $p$  in  $I_i$  is defined as follows:

$$H_{ij} = K_j \left( R_j R_i^{-1} + \frac{R_j (C_i - C_j) n^T}{n^T d K_i^{-1} p} \right) K_i^{-1}, \quad (1)$$

where  $n^T$  and  $d$  respectively denote normal and depth hypothesis. Given the fixed patch which is a square window  $W$  centered on  $p$  in  $I_i$ , we can leverage  $H_{ij}$  to obtain its corresponding patch  $W'$  in  $I_j$ . Then the matching cost is:

$$m_{i,j}(p, W) = 1 - \frac{\text{cov}(W, W')}{\sqrt{\text{cov}(W, W) \text{cov}(W', W')}}}, \quad (2)$$

where  $\text{cov}$  is weighted covariance (Schönberger et al. 2016). Finally, we can obtain multi-view aggregated cost via view weights (Xu and Tao 2019) and further leverage propagation and refinement to update depth with the minimum cost.

Differently, the deformable PM method (Wang et al. 2023) calculates costs by constructing multiple patches centered on different anchors to form the deformed patches. Anchors are reliable pixels identified by a spoke-like searching strategy centered on unreliable pixel  $p$  with  $\theta$  as the radius. Then the deformable matching cost is calculated as:

$$m_{ij}(p, S) = \lambda m(p, W) + (1 - \lambda) \frac{\sum_{s \in S} m(s, W_s)}{|S|}, \quad (3)$$

where  $S$  is a collection of anchors. Both  $W$  and  $W_s$  share the same window size  $w \times w$  while respectively have intervals  $\frac{w}{2}$  and 2. In experiment,  $w = 11$ ,  $\lambda = 0.25$ ,  $\theta = 45^\circ$ ,  $|S| = 8$ .

### Method

The overview of the proposed method is illustrated in Fig. 2, and each component will be detailed in subsequent sections.

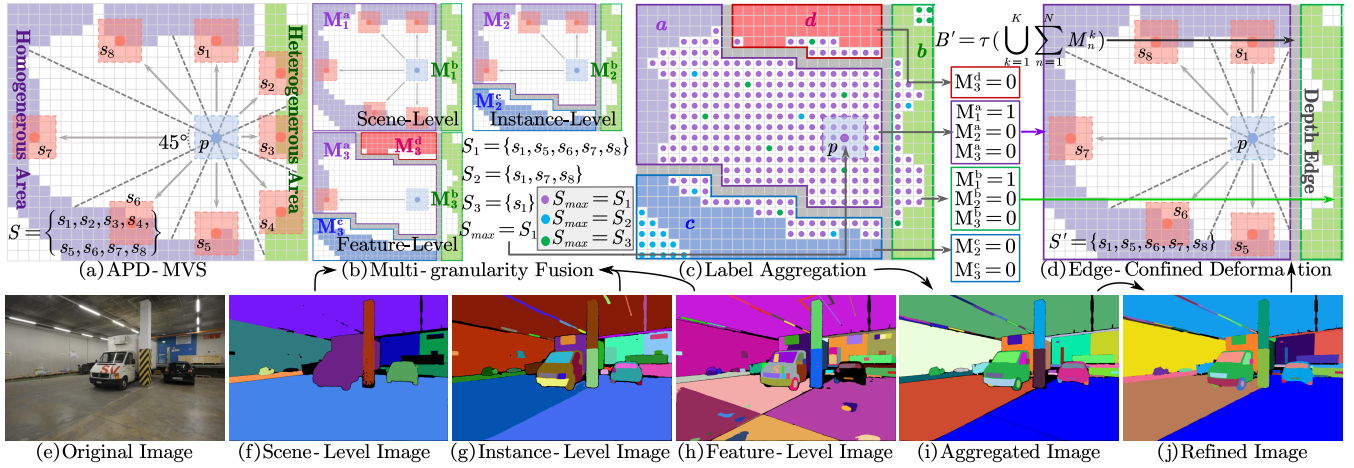


Figure 3: Multi-Granularity Segmentation Prior. From (a) to (d), purple, green and gray backgrounds respectively denote homogeneous areas, heterogeneous areas and depth edges, with sectors divided at fixed angle  $45^\circ$  by black dash lines. In (c), purple, blue, and green dots respectively denote pixels whose optimal anchor subsets  $S_{max}$  equals  $S_1$ ,  $S_2$ , and  $S_3$ . Since most pixels are purple in (c) (i.e.,  $S_{max} = S_1$ ), the scene-level masks  $M_1^a$  and  $M_1^b$  in (b) are reliable, while others are misidentified masks.

### Multi-granularity Segmentation Prior

As shown in Fig. 3 (a), since APD-MVS lacks depth edge guidance, its anchors typically skip depth edges to be selected within green heterogeneous areas, causing potential inaccuracy. Therefore, we attempt to extract depth edges as guidance for patch deformation. However, without accurate depth maps, it is rather difficult to extract depth edges. Although depth edges and image edges have a certain degree of co-occurrence, they are easily confused by factors such as illumination, shadow, occlusion, etc. Differently, by introducing semantic awareness and granularity richness, Semantic-SAM can achieve panoramic segmentation at any granularity level, which may contain potential depth edges. Therefore, we attempt to adopt Semantic-SAM to extract multi-granularity depth edges for patch deformation.

**Edge Aggregation.** As shown in Fig. 3, we first employ Semantic-SAM for panoramic segmentation on input image  $I_i$  in (e) to generate mask maps in (f), (g) and (h) with varying granularities from coarse to fine (i.e., scene, instance and feature level). Then we have  $\{M_n | n = 1, 2, 3\}$ , where  $M_n$  denotes the  $n^{th}$  mask map. In addition, we further apply a transformation function  $\tau$  to identify the boundaries of different mask maps  $M_n$ . Mathematically,  $B_n = \tau(M_n) = \{(x, y) | \Delta M_n(x, y) \neq 0\}$ , where  $\Delta$  represents the Laplace operator. Then we have  $\{B_n = \tau(M_n) | n = 1, 2, 3\}$ , where  $B_n$  denotes the  $n^{th}$  boundary map. Then for each unreliable pixel  $p$  in  $I_i$ , we first construct an anchor subset  $S_n$  for the  $n^{th}$  boundary map  $B_n$ . Subsequently, we utilize the  $n^{th}$  boundary map  $B_n$  to determine whether each anchor  $s_i \in S$  is retained. Specifically, an anchor  $s_i$  can only be retained if its line  $l$ , connecting it with the center pixel  $p$ , does not intersect with  $B_n$ ; otherwise, we discard it. As shown in Fig. 3 (b), All retained  $s_i$  form  $S_n$ , defined by:  $\{s_i \in S_n | l \cap B_n = \emptyset\}$ . Therefore, we have  $\{M_n, B_n, S_n | n = 1, 2, 3\}$ .

We then attempt to obtain the optimal anchor subset  $S_{max}$  based on the reliability of anchor distributions for different anchor subsets. Specifically, we perform RANSAC-based planarization on all anchors in  $S_n$  to obtain the number of anchors classified as inliers, termed as  $A_n$ . Then we consider that  $S_{max}$  should maximize the number of inliers  $A_n$  and maximize the proportion of inliers  $\frac{A_n}{|S_n|}$ , defined by:

$$S_{max} = \arg \max_{S_n} \left( \max(A_n) \text{ s.t. } \max \left( \frac{A_n}{|S_n|} \right) \right). \quad (4)$$

However, directly replacing  $S$  with  $S_{max}$  could render deformed patches susceptible to noise disturbances or local-minimum. Instead, a more robust approach involves merging different mask maps  $M_n$ , thereby forming an aggregated segmentation image that offers enhanced guidance. Therefore, we define  $M_n^k$  as the  $k^{th}$  label in the  $n^{th}$  layer mask map. Subsequently, as depicted on the right side of Fig. 3 (c), we retain  $M_n^k$  only if  $n^{th}$  layer contains the highest number of optimal anchor subsets among all layers, formulated by:

$$M_n^k = \begin{cases} 1, & \text{if } n = \arg \max_l \left( \sum_{p \in M_l^k} \mathbb{I}(S_{max} = S_l) \right); \\ 0, & \text{else.} \end{cases} \quad (5)$$

where  $\mathbb{I}(\cdot)$  is the indicator function such that  $\mathbb{I}(\text{true}) = 1$  and  $\mathbb{I}(\text{false}) = 0$ . Therefore, reliable masks that maximize the quantity of  $S_{max}$  are effectively retained. Finally, as shown in Fig. 3 (i), by aggregating all label masks among all layers, we effectively obtain the aggregated image, which is then further used by function  $\tau$  to derive the aggregated boundary map:  $B' = \tau \left( \bigcup_{k=1}^K \sum_{n=1}^N M_n^k \right)$ .  $B'$  is then adopted as the multi-granularity segmentation prior to constrain each unreliable pixel  $p$  in  $I_i$  to obtain the new anchor collection  $S'$  for patch deformation as shown in Fig. 3 (d). Although  $S'$  contains fewer anchors than  $S$ , it significantly enhances the anchor's reliability. Compared with APD-MVS in Fig. 3 (a),

our edge-confined patch deformation in Fig. 3 (d) effectively avoids unexpected edge-skipping and constraints patch deformation in homogeneous areas, thus avoiding distortion.

**Edge Refinement.** Moreover, to better align the integrated edges with their actual depth edges, we further propose an edge refinement strategy leveraging the conditional random field (CRF) algorithm. CRF is composed of both unary potential and pairwise potential. The unary potential  $\psi_u(l_i)$  represents the probability of the pixel  $p$  being assigned to the label  $l_i$ . The pairwise potential  $\psi_p(l_i, l_j)$  describes the joint probability of two adjacent pixels  $p$  and  $q$  being categorized as label  $l_i$  and  $l_j$ , respectively. Since edges primarily exhibit depth and color differences, we incorporate both geometry and color similarities for edge refinement, such that:

$$\psi_p(l_i, l_j) = \mu(l_i, l_j) \cdot e^{-\left(t \cdot \frac{\|D_p - D_q\|}{\alpha} + \frac{1}{t} \cdot \frac{\|I_p \cdot I_q - 1\|}{\beta}\right)}, \quad (6)$$

where  $D_p$ ,  $D_q$ ,  $I_p$  and  $I_q$  respectively denote the depth and normalized RGB colors of pixels  $p$  and  $q$ ,  $t$  is the current iteration number,  $\mu(l_i, l_j)$  is the label compatibility function defined in (Zhu et al. 2023). Finally, as shown in Fig. 3 (j), by minimizing both  $\psi_u(l_i)$  and  $\psi_p(l_i, l_j)$  among all pixels, we identify the optimal label assignment, thereby generating fine-grained depth edges for accurate patch deformation. A detailed description is available in supplementary material.

### Attention-Consistent Patch Deformation

Although edge-confined patch deformation in Fig. 4 (b) effectively improves the reliability of deformed patches compared to APD-MVS in Fig. 4 (a), it may cause the loss of anchors within sectors fully enclosed by depth edges. Addressing this, we propose adaptive equidistribution, which uniformly divides the homogeneous area of each unreliable pixel into eight sectors based on the number of its surrounding reliable pixels to recover previously-discarded anchors.

**Adaptive Equidistribution.** Specifically, as shown in Fig. 4 (b), centered on unreliable pixel  $p$ , we first emit 8 rays at fixed angle  $45^\circ$  for sector division. We then define the first reliable pixel or depth edge encountered by each ray as the boundary pixel. Thus, we have all sector angles  $\Theta = \{\theta_i \mid i = 1 \cdots 8\}$ , sectors  $\mathcal{V} = \{V_i \mid i = 1 \cdots 8\}$ , and boundary pixels  $\mathcal{B} = \{b_i, b_{i+1} \mid i = 1 \cdots 8\}$ . Subsequently, for sector  $V_i$ , if either its boundary pixels  $b_i$  or  $b_{i+1}$  is located on depth edges rather than reliable pixels, we consider  $V_i$  to be fully enclosed by depth edges and contain no reliable pixels. Otherwise, we consider  $V_i$  to be fully enclosed by a certain number of reliable pixels, denoted as  $n_i$ . Since reliable pixels are typically continuously distributed,  $n_i$  can be approximately estimated as the L1 distance between two boundary pixels  $b_i$  and  $b_{i+1}$ . Moreover, we can obtain the average number of all reliable pixels  $\bar{d}$  among all sectors  $\mathcal{V}$ . Then through the weighted averaging of all reliable pixels, we can proportionally recalculate all angles  $\Theta' = \{\theta'_i \mid i = 1 \cdots 8\}$  and in turn acquire new sectors  $\mathcal{V}'$  and boundary pixels  $\mathcal{P}'$  for allocation. The  $k^{th}$  rectified angle  $\theta'_k$  is calculated by:

$$\theta'_k = \theta'_{k-1} + \frac{k\bar{d} - \sum_{i=1}^{k-1} d_i}{d_x}, \quad \sum_{i=1}^{k-1} d_i \leq k\bar{d} < \sum_{i=1}^k d_i. \quad (7)$$

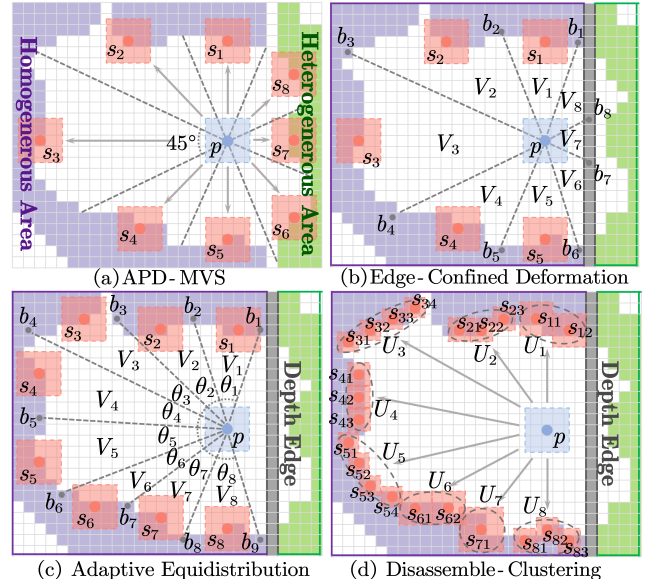


Figure 4: Attention-Consistent Patch Deformation. Blue and red dots respectively denote the central pixel and anchors. purple, green and gray backgrounds respectively denote homogeneous areas, heterogeneous areas and depth edges.

Iterating the above process several times we obtain Fig. 4 (c), where each sector is assigned a roughly equal number of reliable pixels to ensure a balanced division without wastage.

**Disassemble-Clustering.** After adaptive equidistribution, all anchors including the removed ones are evenly allocated. However, since each sector only selects one anchor to construct the deformed patch, a single patch typically fails to comprehensively capture the characteristics of all reliable pixels within its sector, thus causing insufficient attention.

Addressing this issue, we first quadruple the number of anchors within each sector to disassemble each anchor-centered patch into four smaller sub-patches. We then further perform clustering on all anchors to obtain eight anchor clusters that served as the new deformed patches with attention consistency. Specifically, we apply RANSAC-based planarization on all quadrupled anchors within  $\mathcal{V}'$  to derive the optimal plane  $\pi$ . Subsequently, we take  $\varepsilon$  as a threshold to filter out outliers, selecting anchors whose distance to  $\pi$  is less than  $\varepsilon$  as reliable anchors for subsequently clustering.

Here we adopt the density-based DBSCAN algorithm for clustering, which is robust against noise and more suitable for randomly distributed anchors compared to K-means. Specifically, anchors with  $\eta$  neighbors within radius  $\gamma$  are classified as core points. These core points are then connected if they are within  $\gamma$  of each other to form different clusters. Other unconnected points are either integrated into existing clusters if they are within  $\gamma$  of any connected core points, or form an independent cluster if isolated from each other. The above process will iteratively repeat to obtain  $N$  anchor clusters, defined by  $\mathcal{U} = \{U_i \mid i = 1, \dots, N\}$ , each comprising  $k_i$  anchor that  $\mathcal{K} = \{k_i \mid i = 1, \dots, N\}$ , as shown in Fig. 4 (d). For cluster  $U_i$ , all its anchors are

$S'_i = \{s_{ij} \mid j = 1, \dots, k_i\}$ , with their patch size set to  $\frac{w}{k_i}$  to balance attention, thereby avoiding patches overlapping with each other in dense areas. Ultimately, the aggregation of all anchors within all clusters  $S'' = \sum_{i=1}^N S'_i$  will replace  $S'$  for attention-consistent patch deformation. Comparing with Fig. 4 (b), the sub-patches in Fig. 4 (d) provides more balanced coverage of reliable pixels in homogeneous areas.

### Disparity-Sampling Synergistic 3D Optimization

Fig. 5 (a) depicts a close-up view of pixel  $p$  selecting one of its anchor pixels  $s_i$  for patch deformation. Centered on  $s_i$ , a sub-patch is constructed and fixed-interval sampling pixels are adopted for cost computation. Subsequently, as shown in the cyan line of Fig. 5 (b), by assigning varying disparities to these sampling pixels, APD-MVS constructs a cost profile in 2D space to help identify 'global' minimum matching costs.

However, fixed-interval sampling pixels typically fail to globally represent the features of its sub-patch. Therefore, its cost profile is completely different from the optimal cost profile (i.e. the red line of Fig. 5 (b)) which reflects the real matching cost, thereby causing costs to be trapped in a local minimum rather than the global minimum. Addressing this, we attempt to synergistically optimize sampling pixels and disparities to construct multiple cost profiles in 3D space, thereby identifying the real global-minimum matching cost.

Considering identifying optimal sampling pixels in sub-patches is a discrete optimization problem, we employ iterative local search for optimization, which is achieved by performing random perturbation on current solutions and minimizing an objective function to identify the global optimum.

**Sampling Pixel Randomization.** Specifically, as shown in Fig. 5 (c), to maintain the total number of sampling pixels within deformed patches, we first assign each anchor-centered sub-patch in cluster  $U_i$  a certain number of sampling pixels, calculated as  $\frac{9}{k_i}$ . These sampling pixels are initially positioned randomly for cost computation, which we refer to as  $x_0$ . Next, we generate multiple solutions by randomly shuffling their positions to acquire a solution set  $\mathcal{X} = \{x_j \mid j = 1, \dots, J\}$ . For each solution  $x_j$ , given pixel  $p$  and current disparity  $d$ , we construct a disparity sequence  $D$  centered on  $d$  with length  $\mu$  and interval  $\delta$ , defined as:

$$D = \left\{ d_i \mid d_i = d + \left( i - \frac{\mu + 1}{2} \right) \cdot \delta, i = 1, 2, \dots, \mu \right\}. \quad (8)$$

Subsequently, for each solution  $x_j$ , we apply each disparity  $d_i$  in the sequence  $D$  on  $x_j$  for cost computation, thereby yielding a cost profile  $C_j = \{c_i \mid i = 1, \dots, \mu\}$  and corresponding average cost  $\bar{c}_j$ . By aggregating all solutions, we have  $\{(x_j, C_j, \bar{c}_j) \mid j = 1, \dots, J\}$ . Since we introduce multiple solutions  $x_j$  to construct distinct 2D cost profiles  $C_j$ , this can be considered to be a 3D optimization problem.

**3D Cost Profile Optimization.** By comparing cost profiles of different solutions, we can identify the best sampling solution for cost computation, as illustrated in Fig. 5 (d). The optimal cost profile is defined as the one that minimizes cost while maximizing variance. The minimum cost indicates the

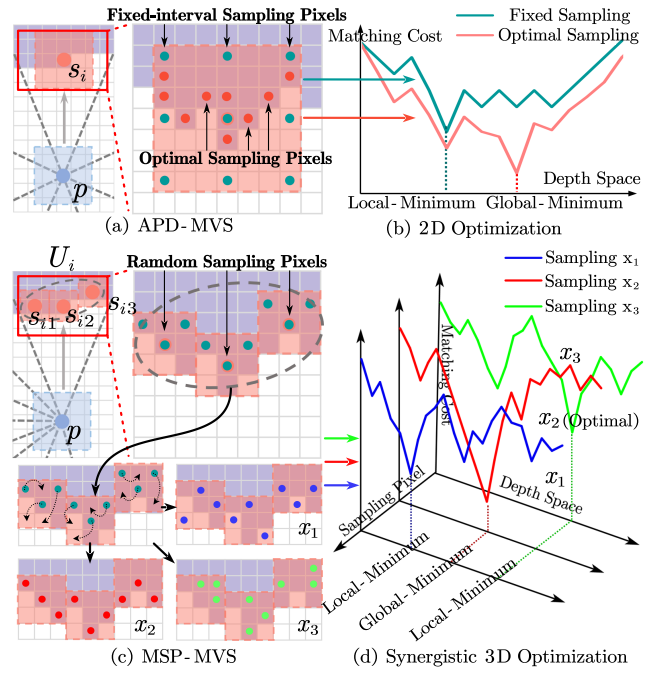


Figure 5: Disparity-Sampling Synergistic 3D Optimization. (a) and (b) respectively represent the fixed sampling pattern and 2D cost optimization employed by APD-MVS. (c) and (d) respectively illustrate the sampling pixel randomization and 3D cost profile optimization of our proposed method.

best matching results, whereas the maximum variance signifies the lowest matching ambiguity. Specifically, the objective function for the cost profile  $C_j$  of solution  $x_j$  is:

$$F(x_j) = \sum_{i=1}^{\mu} c_i + \omega \sum_{i=1}^{\mu} (c_i - \bar{c})^{-2}. \quad (9)$$

Finally, by evaluating  $\min_{x_j \in \mathcal{X}} F(x_j)$  we determine the optimal solution  $x_j$  among all  $\mathcal{X}$ , which will replace the initial solution  $x_0$  in subsequent optimization. Through multiple iterations, we can progressively refine the sampling pixel positions and approach the global-minimum matching cost.

## Experiment

### Datasets and Implementation Details

We evaluate our work on the ETH3D benchmark (Schops et al. 2017) and the Tanks & Temples benchmark (TNT) (Knapitsch et al. 2017). We compare our method against recent learning-based methods like PatchMatchNet, IterMVS, MVSTER, AA-RMVSNet, EPP-MVSNet, EPNet and traditional methods like TAPA-MVS, PCF-MVS, ACMM, ACMP, ACMMP, SD-MVS, APD-MVS and HPM-MVS.

Our method is implemented on a machine with an Intel(R) Xeon(R) Silver 4210 CPU and eight NVIDIA GeForce RTX 3090 GPUs. Experiments are performed on original images in both ETH3D dataset and TNT dataset. In cost calculation, we adopt the matching strategy of every other row and column. We take APD-MVS (Wang et al. 2023) as our baseline.

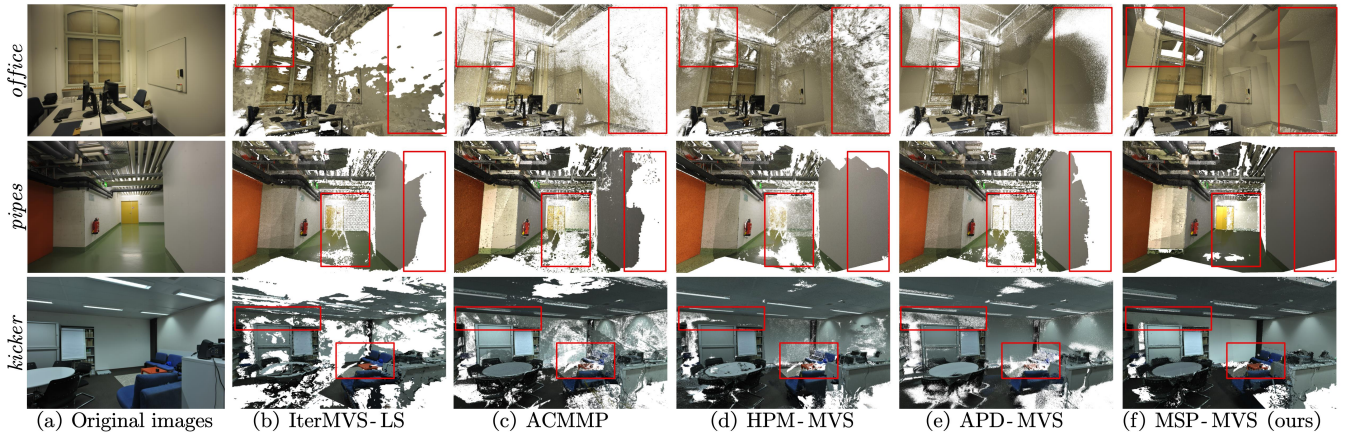


Figure 6: An illustration of qualitative results on partial scenes of ETH3D datasets (*office*, *pipes* and *kicker*). Partial challenging areas are shown in red boxes. Our method can effectively reconstruct large textureless areas without introducing detail distortion.

Method	Train			Test		
	F <sub>1</sub>	Comp.	Acc.	F <sub>1</sub>	Comp.	Acc.
PatchMatchNet	64.21	65.43	64.81	73.12	77.46	69.71
IterMVS-LS	71.69	66.08	79.79	80.06	76.49	84.73
MVSTER	72.06	76.92	68.08	79.01	82.47	77.09
EPP-MVSNet	74.00	67.58	82.76	83.40	81.79	85.47
EPNet	79.08	79.28	79.36	83.72	87.84	80.37
TAPA-MVS	77.69	71.45	85.88	79.15	74.94	85.71
PCF-MVS	79.42	75.73	84.11	80.38	79.29	82.15
ACMM	78.86	70.42	<b>90.67</b>	80.78	74.34	90.65
ACMMP	83.42	77.61	<u>90.63</u>	85.89	81.49	<b>91.91</b>
SD-MVS	86.94	84.52	89.63	88.06	87.49	88.96
HPM-MVS++	<u>87.09</u>	<u>85.64</u>	88.74	<u>89.02</u>	<u>89.37</u>	88.93
APD-MVS (base)	86.84	84.83	89.14	87.44	85.93	89.54
MSP-MVS (ours)	<b>88.09</b>	<b>87.87</b>	88.51	<b>89.51</b>	<b>90.38</b>	89.08

Table 1: Quantitative results on ETH3D dataset at threshold  $2cm$ . Our method achieves the highest F<sub>1</sub> and completeness.

## Results on ETH3D and TNT

Qualitative results on ETH3D are illustrated in Fig. 6. It is evident that our method delivers the most complete and realistic reconstructed point clouds, especially in large textureless areas like the floors and walls, without introducing detail distortions. More qualitative results on the ETH3D and TNT dataset can be referred to in the supplementary material.

Quantitative results on the ETH3D and the TNT datasets are respectively presented in Tab. 1 and Tab. 2. Note that the first group is learning-based methods and the second is traditional methods. Meanwhile, we mark the best results in bold and underline the second-best results. Our method achieves the highest F<sub>1</sub> score and completeness on the ETH3D dataset and the TNT Intermediate dataset, validating its state-of-the-art performance and strong generalization capability. Meanwhile, our method achieves the second best results in the TNT advanced datasets, falling short by 0.26% in F<sub>1</sub> score compared to EPNet (Su and Tao 2023).

## Memory & Runtime Comparison

To visually demonstrate the reconstruction quality between different methods, we present the quantitative results of

Method	Intermediate			Advanced		
	F <sub>1</sub>	Rec.	Pre.	F <sub>1</sub>	Rec.	Pre.
PatchMatchNet	53.15	69.37	43.64	32.31	41.66	27.27
AA-RMVSNet	61.51	75.69	52.68	33.53	33.01	37.46
IterMVS-LS	56.94	74.69	47.53	34.17	44.19	28.70
MVSTER	60.92	<u>77.50</u>	50.17	37.53	45.90	33.23
EPP-MVSNet	61.68	75.58	53.09	35.72	34.63	<u>40.09</u>
EPNet	<u>63.68</u>	72.57	<b>57.01</b>	<b>40.52</b>	<u>50.54</u>	34.26
ACMM	57.27	70.85	49.19	34.02	34.90	35.63
ACMP	58.41	73.85	49.06	37.44	42.48	34.57
ACMMP	59.38	68.50	53.28	37.84	44.64	33.79
SD-MVS	63.31	76.63	53.78	40.18	47.37	35.53
HPM-MVS++	61.59	73.79	54.01	39.65	41.09	<b>40.79</b>
APD-MVS (base)	63.64	75.06	<u>55.58</u>	39.91	49.41	33.77
MSP-MVS (ours)	<b>64.48</b>	<b>76.42</b>	56.21	<u>40.26</u>	<b>51.20</b>	33.38

Table 2: Quantitative results on TNT dataset at given threshold. Our method achieves the highest F<sub>1</sub> and completeness.

our approach against other methods on both the ETH3D and Tanks & Temples (TNT) datasets, as shown in Fig. 7 (a) and Fig. 7 (b), respectively. As illustrated, our method achieves the highest F1 scores on both the ETH3D and TNT datasets, outperforming other state-of-the-art methods such as ACMMP, APD-MVS, SD-MVS, and HPM-MVS, thereby demonstrating its superior performance and robustness.

Moreover, to demonstrate the efficiency of our method, we present a comparative analysis of GPU memory usage and runtime on ETH3D training datasets, as shown in Fig. 7 (c). All experiments are performed on original-resolution images, with the number of images standardized to 10 across all scenarios for runtime evaluation. Additionally, to ensure fairness, all methods are tested on the same system, with its hardware configuration detailed in the previous section.

Concerning learning-based methods, although IterMVS-LS has the shortest runtime, its memory usage exceeds the maximum capacity of mainstream GPUs. Similarly, other learning-based methods also struggle with excessive memory consumption, making them impractical for large-scale reconstructions. Regarding traditional methods, while our




Figure 7: Comparative analysis. (a) and (b) respectively illustrate  $F_1$  scores among different methods on ETH3D and Tanks & Temples datasets. (c) depicts their corresponding GPU memory usage (GB) and runtime (second) on ETH3D training datasets.

method requires approximately one-third more time than its baseline, APD-MVS, its memory usage significantly outperforms most methods, including ACMMP, SD-MVS, and HPM-MVS. Thus, our approach can achieve state-of-the-art performance with acceptable runtime and minimal memory consumption, proving its effectiveness and practicality.

### Ablation Study

Tab. 3 presents the ablation studies of each proposed component. Here we primarily use  $F_1$  score for overall comparison.

**Multi-granularity Segmentation Prior** We separately remove edge aggregation (w/o. Agr.), edge refinement (w/o. Ref.) and both (w/o. Mul.). w/o. Mul. produces the worst results, highlighting the effectiveness of multi-granularity segmentation prior for edge-confined patch deformation. w/o. Ref. performs better than w/o. Agr., indicating that edge aggregation exhibit a greater impact than edge refinement.

**Attention-Consistent Patch Deformation** We respectively exclude the module of adaptive equidistribution (w/o. Eqd.), disassemble-clustering (w/o. Clu.) and both (w/o. Con.). w/o. Con. results in the poorest performance, demonstrating the significant impact of attention-consistent patch deformation. w/o. Clu. slightly outperforms w/o. Eqd., yet both fall short to MSP-MVS, implying that adaptive equidistribution contributes more than disassemble-clustering.

**Disparity-Sampling Synergistic 3D Optimization** We respectively eliminate the variance component (w/o. Var.) and cost component (w/o. Cst.) in objective function, the whole module (w/o. Syn.) and set  $\mu = 3$ . Due to the absence of 3D optimization, w/o. Syn. achieves the poorest results. w/o. Var. yields better results than w/o. Cst., suggesting that cost plays a more critical role than variance during cost optimization. Moreover,  $\mu = 3$  are inferior to  $\mu = 5$  (MSP-MVS), indicating the necessity of cost profile aggregation.

### Conclusion

In this paper, we introduce MSP-MVS, a method leveraging multi-granularity segmentation prior for edge-confined

Method	2cm			10cm		
	$F_1$	Comp.	Acc.	$F_1$	Comp.	Acc.
w/o. Mul.	86.53	84.32	<b>89.06</b>	96.59	95.25	<b>97.95</b>
w/o. Agr.	87.07	85.88	88.48	97.01	96.38	97.64
w/o. Ref.	87.54	86.94	88.36	97.32	97.05	97.59
w/o. Con.	86.61	86.49	86.91	96.77	96.88	96.66
w/o. Eqd.	87.23	87.08	87.58	97.16	97.25	97.08
w/o. Clu.	87.29	87.16	87.62	97.21	97.29	97.14
w/o. Syn.	86.94	87.29	86.8	96.9	97.33	96.47
w/o. Var.	87.17	87.41	87.15	97.07	97.46	96.68
w/o. Cst.	87.49	87.6	87.57	97.28	97.55	97.02
$\mu = 3$	87.73	87.53	88.13	97.42	97.59	97.25
MSP-MVS	<b>88.09</b>	<b>87.87</b>	88.51	<b>97.70</b>	<b>97.74</b>	97.67

Table 3: Quantitative ablation studies on ETH3D dataset.

patch deformation. We first aggregate and refine multi-granularity depth edges via Semantic-SAM as prior to guide patch deformation within homogeneous areas. Moreover, we implement adaptive equidistribution for sector division and disassemble-clustering for anchor clustering to enable patch deformation with attention consistency. Additionally, we propose disparity-sampling synergistic 3D optimization to help deformed patches identify global-minimum matching costs. Results on ETH3D and Tanks & Temples benchmarks validate the state-of-the-art performance of our method.

### Acknowledgements

This work was supported in part by the National Natural Science Foundation of China under Grant 62172392, in part by the Strategic Priority Research Program of the Chinese Academy of Sciences under Grant No. XDA0450203, in part by the Innovation Program of Chinese Academy of Agricultural Sciences (Grant No. CAAS-CSSAE-202401 and CAAS-ASTIP-2024-AII), and in part by the Beijing Smart Agriculture Innovation Consortium Project (Grant No. BAIC10-2024).

### References

Bao, L.; Yang, Q.; Jin, H.; and Jin, H. 2014. Fast Edge-Preserving Patchmatch for Large Displacement Opti-

- cal Flow. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3534–3541.
- Barnes, C.; Shechtman, E.; Finkelstein, A.; and Goldman, D. B. 2009. PatchMatch: A randomized correspondence algorithm for structural image editing. *ACM Trans. Graph.*, 28(3): 24.
- Bleyer, M.; Rhemann, C.; Rother, C.; and Rother, C. 2011. PatchMatch Stereo - Stereo Matching with Slanted Support Windows. In Hoey, J.; McKenna, S. J.; and Trucco, E., eds., *British Mach. Vis. Conf. (BMVC)*, 1–11.
- Cao, C.; Ren, X.; and Fu, Y. 2024. MVFormer++: Revealing the Devil in Transformer’s Details for Multi-View Stereo. arXiv:2401.11673.
- Cao, M.; Zheng, L.; Jia, W.; Lu, H.; and Liu, X. 2020. Accurate 3-D Reconstruction under IoT Environments and Its Applications to Augmented Reality. *IEEE Transactions on Industrial Informatics*, 17(3): 2090–2100.
- 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.
- Galliani, S.; Lasinger, K.; and Schindler, K. 2015. Massively Parallel Multiview Stereopsis by Surface Normal Diffusion. In *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*.
- Heise, P.; Klose, S.; Jensen, B.; and Knoll, A. 2013. Pm-Huber: Patchmatch with Huber Regularization for Stereo Matching. In *Proceedings of the IEEE International Conference on Computer Vision*, 2360–2367.
- 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*, 18–32.
- Li, F.; Zhang, H.; Sun, P.; Zou, X.; Liu, S.; Yang, J.; Li, C.; Zhang, L.; and Gao, J. 2023. Semantic-Sam: Segment and Recognize Anything at Any Granularity. *arXiv preprint arXiv:2307.04767*.
- Li, Z.; Zuo, W.; Wang, Z.; and Zhang, L. 2020. Confidence-Based Large-Scale Dense Multi-View Stereo. *IEEE Transactions on Image Processing*, 29: 7176–7191.
- Liao, J.; Fu, Y.; Yan, Q.; and Xiao, C. 2019. Pyramid Multi-View Stereo with Local Consistency. *Computer Graphics Forum*, 38(7): 335–346.
- 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 *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 501–518.
- Schops, T.; Schönberger, 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 *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*.
- Shi, Y.; Xi, J.; Hu, D.; Cai, Z.; and Xu, K. 2023. RayMVS-Net++: Learning Ray-Based 1D Implicit Fields for Accurate Multi-View Stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Shvets, M.; Zhao, D.; Niethammer, M.; Sengupta, R.; and Berg, A. C. 2024. Joint Depth Prediction and Semantic Segmentation with Multi-View Sam. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 1328–1338.
- Su, W.; and Tao, W. 2023. Efficient Edge-Preserving Multi-View Stereo Network for Depth Estimation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, 2348–2356.
- Wang, F.; Galliani, S.; Vogel, C.; and Pollefeys, M. 2022a. IterMVS: Iterative probability estimation for efficient multi-view stereo. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 8606–8615.
- Wang, X.; Zhu, Z.; Huang, G.; Qin, F.; Ye, Y.; He, Y.; Chi, X.; and Wang, X. 2022b. MVSTER: Epipolar Transformer for Efficient Multi-View Stereo. In *European Conference on Computer Vision*, 573–591. Springer.
- 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 *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 1621–1630.
- Wei, Y.; Zhao, L.; Zheng, W.; Zhu, Z.; Zhou, J.; and Lu, J. 2023. Surroundoc: Multi-camera 3d Occupancy Prediction for Autonomous Driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 21729–21740.
- Wei, Z.; Zhu, Q.; Min, C.; Chen, Y.; and Wang, G. 2021. Aa-Rmvsnet: 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*.
- Xu, Q.; and Tao, W. 2019. Multi-Scale Geometric Consistency Guided Multi-View Stereo. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*.

- Xue, T.; Owens, A.; Scharstein, D.; Goesele, M.; and Szeliski, R. 2019. Multi-Frame Stereo Matching with Edges, Planes, and Superpixels. *Image and Vision Computing*, 91: 103771.
- Yan, J.; Wei, Z.; Yi, H.; Ding, M.; Zhang, R.; Chen, Y.; Wang, G.; and Tai, Y.-W. 2020. Dense Hybrid Recurrent Multi-view Stereo Net with Dynamic Consistency Checking. In Vedaldi, A.; Bischof, H.; Brox, T.; and Frahm, J.-M., eds., *Computer Vision – ECCV 2020*, volume 12349, 674–689. Cham: Springer International Publishing. ISBN 978-3-030-58547-1 978-3-030-58548-8.
- Yang, J.; Mao, W.; Alvarez, J. M.; and Liu, M. 2020. Cost Volume Pyramid Based Depth Inference for Multi-View Stereo. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 4876–4885.
- Yao, Y.; Luo, Z.; Li, S.; Fang, T.; and Quan, L. 2018. MVS-Net: Depth Inference for Unstructured Multi-view Stereo. In *Proc. Eur. Conf. Comput. Vis. (ECCV)*.
- Yao, Y.; Luo, Z.; Li, S.; Zhang, J.; Ren, Y.; Zhou, L.; Fang, T.; and Quan, L. 2020. Blendedmvs: A Large-Scale Dataset for Generalized Multi-View Stereo Networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1790–1799.
- 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. *Proceedings of the AAAI Conference on Artificial Intelligence*, 38(7): 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*, 110565.
- Zhang, Y.; Zhu, J.; Lin, L.; and Lin, L. 2023a. Multi-View Stereo Representation Revist: Region-Aware MVSNet. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17376–17385.
- Zhang, Z.; Peng, R.; Hu, Y.; and Wang, R. 2023b. GeoMVS-Net: Learning Multi-View Stereo With Geometry Perception. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 21508–21518.
- Zhu, J.; Huang, H.; Li, B.; and Wang, L. 2023. E-CRF: Embedded Conditional Random Field for Boundary-caused Class Weights Confusion in Semantic Segmentation. arxiv:2112.07106.
- Zhu, Z.; Ng, D. W. H.; Park, H. S.; and McAlpine, M. C. 2021. 3D-printed Multifunctional Materials Enabled by Artificial-Intelligence-Assisted Fabrication Technologies. *Nature Reviews Materials*, 6(1): 27–47.