

Hierarchical Prior Mining for Non-local Multi-View Stereo

Chunlin Ren¹ Qingshan Xu² Shikun Zhang¹ Jiaqi Yang^{1*}
¹ Northwestern Polytechnical University ² Nanyang Technological University
 {renchunlin, zhangshikun}@mail.nwpu.edu.cn; qingshan.xu@ntu.edu.sg;
 jqyang@nwpu.edu.cn

Abstract

As a fundamental problem in computer vision, multi-view stereo (MVS) aims at recovering the 3D geometry of a target from a set of 2D images. Recent advances in MVS have shown that it is important to perceive non-local structured information for recovering geometry in low-textured areas. In this work, we propose a Hierarchical Prior Mining for Non-local Multi-View Stereo (HPM-MVS). The key characteristics are the following techniques that exploit non-local information to assist MVS: 1) A Non-local Extensible Sampling Pattern (NESP), which is able to adaptively change the size of sampled areas without becoming snared in locally optimal solutions. 2) A new approach to leverage non-local reliable points and construct a planar prior model based on K -Nearest Neighbor (KNN), to obtain potential hypotheses for the regions where prior construction is challenging. 3) A Hierarchical Prior Mining (HPM) framework, which is used to mine extensive non-local prior information at different scales to assist 3D model recovery, this strategy can achieve a considerable balance between the reconstruction of details and low-textured areas. Experimental results on the ETH3D and Tanks & Temples have verified the superior performance and strong generalization capability of our method. Our code will be available at <https://github.com/CLinvx/HPM-MVS>.

1. Introduction

Multi-View Stereo (MVS) is one of the key problems in the field of 3D computer vision, which contributes greatly to virtual reality [5], 3D object recognition [14], and autonomous driving [23]. The objective of MVS method is to reconstruct the 3D geometry of a target from a series of images along with camera parameters. Over the past few years, the emergence of numerous datasets [15, 26, 45] has tremendously contributed to the continuous advancement of MVS methods [8, 27, 9]. However, influenced by low-

¹Corresponding author

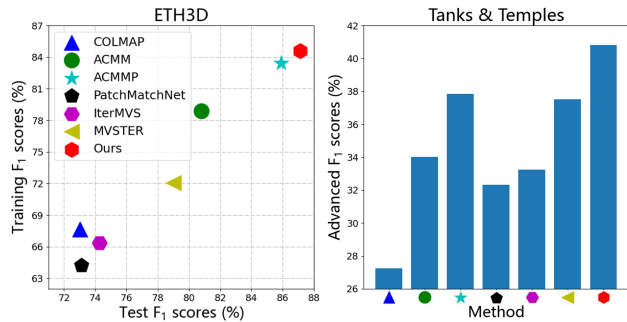


Figure 1: F_1 score (\uparrow) comparisons with SOTA traditional [25, 38, 36] and learning-based [30, 29, 32] MVS methods on ETH3D [26] and Tanks & Temples [17].

textured areas, illumination variation and other factors [1], MVS is still a challenging problem.

To perform an accurate and robust 3D reconstruction, two different types of methods have been investigated, including learning-based MVS [43, 44, 37, 22] and traditional MVS [9, 38, 36]. Learning-based MVS implements deep networks to extract high-level features and make predictions. However, it needs a large amount of training and its generalization ability still needs further improvement [16]. Traditional MVS can also be broadly classified into two categories: plane-sweeping-based MVS [7, 3, 10] and PatchMatch MVS [9, 25, 38]. Although plane-sweeping-based MVS produces good results for sufficiently textured and unoccluded surfaces, it performs poorly for scenes composed of large planar surfaces, while PatchMatch MVS has successfully overcome this restriction.

PatchMatch MVS usually consists of four steps, including random initialization, hypothesis propagation, multi-view matching cost evaluation and refinement [36]. In essence, it is a process of sampling and verification. The hypothesis propagation samples appropriate hypotheses from neighboring pixels to construct solution space while the multi-view matching cost evaluation defines a criterion to verify the reliability of the sampled hypotheses. Therefore, these two steps play an important role in PatchMatch

MVS. Following this popular four-step pipeline, PatchMatch MVS produces remarkable outcomes.

Even so, geometry recovery still suffers a lot from low-textured areas because the previous PatchMatch MVS methods place too much emphasis on local information. Therefore, we exploit non-local information to achieve high-quality reconstruction. We propose a Hierarchical Prior Mining for Non-local Multi-View Stereo (HPM-MVS) method, which defines a novel hypothesis propagation pattern and a novel way for constructing planar prior models to assist multi-view matching cost evaluation.

In terms of hypothesis propagation, sequential propagation [2, 46, 25] and diffusion-like propagation [9, 38] are all popular strategies. The previous one only considers the nearest pixels for updating. In contrast, the diffusion-like propagation can update simultaneously based on the neighbor hypotheses around the central point. To further leverage structured region information, ACMH [38] proposes an adaptive checkerboard sampling scheme. Although these previous approaches have greatly improved the efficiency and performance of hypothesis propagation, they both process in the local neighborhood; thus misestimates can only be updated when several better hypotheses in the local area are sampled. Based on the above observations, we first propose the basic method with a Non-local Extensible Sampling Pattern (NESP). Intuitively, a non-local operation removes sampling points around the center point. It allows distant correct pixels to have a greater possibility to contribute to the update of the current pixel. Moreover, the extensible architecture maintains the variable sampling size and can efficiently select suitable candidate hypotheses.

In terms of the multi-view matching cost evaluation, since the photometric consistency usually causes ambiguities during the depth optimization of low-textured areas, many previous methods [35, 11, 40] introduce planar prior models to help the evaluation. They assume low-textured areas are usually on smooth homogeneous surfaces, and construct a planar prior model at the current scale to assist depth estimation. To establish a more robust and more comprehensive planar prior, we employ K-Nearest Neighbor (KNN) to search non-local reliable points and obtain potential hypotheses for the marginal regions where prior construction is difficult. Inspired by the coarse-to-fine strategy, we further design an HPM framework. In a coarse-to-fine manner, the prior knowledge at the early stages is built upon non-local credible sparse correspondences at low-resolution scales, which expands the receptive field and leads to an effective depth estimate for smooth homogeneous surfaces. Subsequently, the prior knowledge at the later stages is constructed at the higher-resolution scales by using previously rectified hypotheses, which restores the depth information of details in the image. This way, the depth information can be successfully recovered by our HPM architecture.

Extensive experiments on different competitive datasets verify the effectiveness of our method (Fig. 1). *The results demonstrate that HPM-MVS has excellent performance, and also holds strong generalization capability to more complex scenes.* In a nutshell, our contributions are three-fold as follows:

- We present an NESP module to adaptively determine the number of potential sampling points while avoiding becoming snared in locally optimal solutions.
- To construct a more robust and more comprehensive planar prior model, we propose an approach based on KNN to search non-local neighbor reliable hypotheses and construct a planar prior model for marginal areas where prior construction is challenging.
- We design an HPM framework to explore prior knowledge at multiple scales, which efficiently balances out the reconstruction of details and low-textured areas.

2. Related Work

2.1. Traditional Multi-View Stereo

PatchMatch Multi-View Stereo. The PatchMatch method was initially proposed by Barnes *et al.* [4], and its primary goal is to quickly identify the approximate nearest neighbor among patches of two images. The extension of PatchMatch idea to MVS was proposed by Shen [28]. In order to enhance efficiency, several methods [46, 2, 34] employ the sequential propagation scheme. However, this propagation scheme is still time-consuming. To this end, Galliani *et al.* [9] proposed Gipuma, a massively parallel multi-view extension of PatchMatch MVS which leverages a red-black checkerboard pattern to perform a diffusion-like propagation scheme. This makes better use of the parallelism of GPUs to improve efficiency. Nonetheless, the major drawback is that each pixel’s sampling points are preset, which in turn causes regional information to be ignored. Further, Xu and Tao [38] designed a method with Adaptive Checkerboard sampling and Multi-Hypothesis joint view selection (ACMH) to address this shortcoming. In allusion to depth estimation for low-textured areas, they further combined ACMH with multi-scale geometric consistency guidance to present ACMM. In addition, MARMVS [41] selects the optimal scale for each pixel to alleviate the ambiguity caused by regions with raw texture. Although the aforementioned approaches significantly enhance the performance of PatchMatch MVS method, how to recover the depth information of low-textured regions is always a challenging problem.

Planar Prior Assistance. In order to deal with this long-standing ill-posed problem, the planar prior is considered. Romanoni *et al.* [24] proposed TAPA-MVS, they separated the image into two different scales of superpixels and utilized these as regional structures. The authors fitted a plane

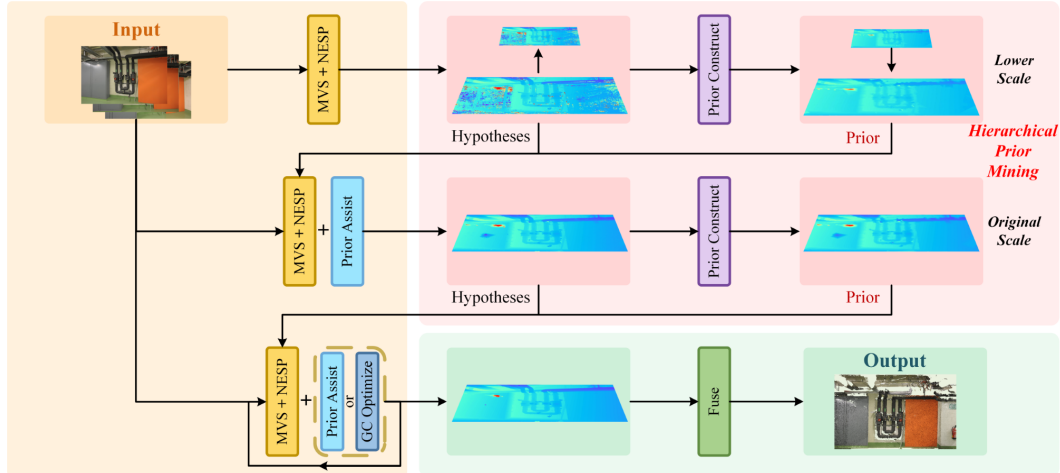


Figure 2: **An overview of HPM-MVS.** Starting with the input images, we first apply the basic MVS with NESP (Fig. 3) to obtain the initial hypotheses. Then, we downsample them to the coarsest scale and generate the prior model. Third, we upsample the model to the current scale and leverage it to assist hypothesis prediction. After operating on different scales, we further use geometric consistency (GC) to optimize the results. Finally, the hypotheses will be fused to a point cloud.

within the superpixels by using trustworthy 3D points and the RANSAC method, then leveraged the prior hypotheses to optimize the low-textured regions. ACMP [40] constructs a planar prior depth map based on triangulations, and subsequently mosaics the prior hypotheses into the MVS process via a probabilistic graphical model. Recently, Xu *et al.* [36] designed a multi-scale geometric consistency guided and planar prior assisted Multi-View Stereo (ACMMP) by combining ACMM [38] and ACMP [40]. It fully exploits depth information at various scales and leverages prior assistance to efficiently target low-textured areas. Previous approaches only consider the prior information from the current image scale, we nonetheless believe that prior hypotheses should not be restricted to just one scale. To this aim, we embrace an HPM framework to discover more meaningful and comprehensive non-local prior knowledge.

2.2. Learning-based Multi-View Stereo

Deep learning has been widely used in 3D vision and has proven to be critical for MVS. MVSNet [43] introduces an end-to-end deep learning network that builds the 3D cost volumes from 2D image features. R-MVSNet [44] regularizes the cost volume along the depth direction with the convolutional GRU to make reconstruction feasible. To further decrease memory and boost performance, CIDER [39] builds a lightweight cost volume using an average group-wise correlation similarity metric. In recent years, the efficiency of learning-based MVS gains more attention. CVP-MVSNet [42], CasMVSNet [12] and UCS-Net [6] build a cascade cost volume by integrating the coarse-to-fine strategy. Furthermore, PatchMatchNet [30] embeds PatchMatch MVS into a deep learning framework to achieve high-

quality reconstruction with low memory consumption.

Although the existing deep learning methods have achieved great success in MVS, they usually require a large quantity of training datasets, which may not be practical in real-world applications. In addition, the generalization ability still remains an issue for learning-based methods [16]. As such, we still focus on traditional methods for MVS.

3. Methodology

In this section, we first briefly review the classical PatchMatch MVS. Then, we detail our proposed HPM-MVS. An overview of the method is illustrated in Fig. 2 (the pseudocode is shown in the supplementary).

3.1. Review of PatchMatch MVS

PatchMatch MVS [9, 25, 38] is a method for rapidly finding correspondences between image patches. To compute depth maps for input images, each image is selected in turn as a reference image and its corresponding source images are determined based on the Structure-from-Motion (SfM) results. In general, the process consists of four steps [36], i.e., random initialization, hypothesis propagation, multi-view cost evaluation and refinement. The last three steps are iterated until convergence, the details are as follows:

- 1) Random initialization means randomly generating a plane hypothesis for each pixel in the reference image based on the results of SfM. A key observation behind is that a non-trivial fraction of a large field of random offset assignments is likely to be a good guess.
- 2) Hypothesis propagation samples the hypotheses from neighboring pixels because of the high spatial coherence of

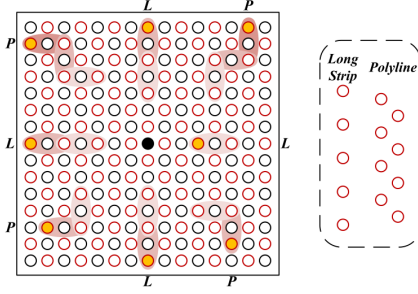


Figure 3: **Non-local Extensible Sampling Pattern.** The circles represent the pixels in the red-black checkerboard. The red areas represent the sampling regions; the lighter the color, the less adaptive extension is performed. The solid yellow circles show the sampled points.

neighbors in images. Moreover, a proper hypothesis can be propagated to a relatively large region to ensure the smoothness constraint of nearby pixels.

3) Multi-view matching cost evaluation aims at robustly integrating matching cost from multiple views to select the best plane hypotheses.

4) Refinement generates two additional hypotheses to enrich the diversity of the solution space. The one with the lowest cost will be selected as the refined estimation.

In this work, we concentrate on hypothesis propagation and multi-view matching cost evaluation, which are essential components of PatchMatch MVS methods, to capture non-local information for recovering better geometry in low-textured areas.

3.2. Non-local Extensible Sampling Pattern

Hypothesis propagation is a prerequisite for multi-view matching cost evaluation and refinement. Hence, choosing a reasonable sampling pattern will lead to accurate results. The pixels within a relatively large area can be represented by one of the pixels. So considering local information repeatedly is redundant. Our NESP is inspired by these two key insights to sample more robust hypotheses. Following the PatchMatch MVS pipeline introduced above, we design our basic MVS with NESP.

Random Initialization. First, we randomly generate a hypothesis $\theta_x = [d_x, \mathbf{n}_x]$ for each pixel x , where d_x represents the depth information and \mathbf{n}_x is a normal vector. Second, a matching cost is calculated from each source image via homography [13]. Finally, the initial multi-view matching cost is obtained by averaging top- k best costs.

Hypothesis Propagation. The diffusion-like propagation scheme can tremendously enhance efficiency while maintaining high-quality hypothesis propagation. Following [38], we first partition all pixels in the reference image into a red-black checkerboard pattern which uses hypothe-

ses within eight regions as candidates to update. Second, we divide eight sampling areas into four long strip areas and four polyline areas; each polyline area starts with 8 samples, compared to each long strip area's initial 5 samples. Then, we consider two strategies to construct our NESP (Fig. 3).

- **Extensible Strategy.** To achieve the purpose of extension, we first choose the best hypothesis in each area according to the multi-view matching costs from random initialization or the previous iteration (the initial costs in random initialization are used for the first propagation, while the other iterations use the costs computed in the previous iteration) and define the candidate hypothesis set as $\Theta = \{\theta_i | i = 1 \cdots 8\}$. Second, we calculate their matching costs with respect to different source images and embed them into a cost matrix,

$$M = \begin{bmatrix} m_{1,1} & m_{1,2} & \cdots & m_{1,N-1} \\ m_{2,1} & m_{2,2} & \cdots & m_{2,N-1} \\ \vdots & \vdots & \ddots & \vdots \\ m_{8,1} & m_{8,2} & \cdots & m_{8,N-1} \end{bmatrix}, \quad (1)$$

where N is the number of input images, $m_{i,j}$ represents the matching cost of the sampled point in the i -th region scored by the j -th view. Based on this matrix, a voting scheme is implemented in each line to determine whether this region needs to be extended. Two thresholds are defined here: 1) A good matching cost threshold is,

$$\tau(t_{iter}, t_{ext}) = \tau_{good} \cdot e^{-\frac{t_{iter}^2 \cdot (N_{ext} - t_{ext})}{\alpha}}, \quad (2)$$

where t_{iter} and t_{ext} represent the t_{iter} -th iteration of hypothesis propagation and the t_{ext} -th regional expansion, respectively; τ_{good} is the initial good matching cost threshold, α is a constant and N_{ext} represents the maximum number of extensions. 2) A bad matching cost threshold is defined as τ_{bad} . For a specific region, there should exist at least n_{good} matching costs smaller than $\tau(t_{iter}, t_{ext})$, and at most n_{bad} matching costs meeting the condition: $m_{i,j} > \tau_{bad}$. When the above conditions are satisfied, the expansion of the current sampling area will be stopped. Otherwise, this area should extend by itself while each extension doubles the number of sampling points in the area. Notably, the set Θ and the matrix M are renewed dynamically in each expansion using the rule described above.

- **Non-local Strategy.** There are two key insights behind the non-local strategy: 1) The surrounding pixels of the central point share the same structured region information, and their hypotheses can be contained by Eq. 4 when in the process of refinement. 2) The misestimates in low-textured areas always appear as small

isolated speckles in the depth map; this phenomenon indicates that these regions have fallen into local optimal solutions. Therefore, we use the non-local strategy (Fig. 3) which abandons the sampling points within the radius R of the center pixels. The motivation is to concentrate more on non-local areas instead of focusing on the local surrounding information.

In the end, we combine these two strategies to form NESP, which helps to collect more reasonable candidate hypotheses from non-local neighboring pixels.

Multi-View Matching Cost Evaluation. In order to determine the best hypothesis from the above candidate hypothesis set, we follow the previous studies [46, 25, 38] to define the multi-view matching cost via photo consistency in our basic MVS with NESP,

$$c_{photo}(\theta_i) = \frac{\sum_j w_j \cdot m_{i,j}}{\sum_j w_j}, \quad (3)$$

where w_j is the weight of the j -th view, which is computed by the view selection strategy in [38]. The hypothesis with the least matching cost will be chosen as the current best estimate for the center pixel.

Refinement. We generate two new hypotheses in refinement. One hypothesis is obtained by perturbing the current hypothesis $[d_x^p, \mathbf{n}_x^p]$, while the other is randomly generated $[d_x^r, \mathbf{n}_x^r]$. Subsequently, the two new hypotheses are randomly arranged and combined with the current hypothesis to form an ensemble of seven hypotheses,

$$\left\{ [d_x, \mathbf{n}_x], [d_x, \mathbf{n}_x^p], [d_x, \mathbf{n}_x^r], [d_x^r, \mathbf{n}_x], [d_x^r, \mathbf{n}_x^r], [d_x^p, \mathbf{n}_x], [d_x^p, \mathbf{n}_x^p] \right\}. \quad (4)$$

The final estimation for pixel x will be the hypothesis with the minimum cost.

3.3. Planar Prior Construction

The local image patches of low-textured areas are highly similar, and the photometric consistency, as a metric function of MVS, always leads to incorrect estimates due to its defects. Inspired by [40], we construct a better planar prior model to assist multi-view matching cost evaluation. After executing the basic MVS, a hypothesis with its final cost for each pixel can be obtained. A lower cost represents a more precise estimate, therefore we collect the credible correspondences into a set I_{cred} ,

$$I_{cred} = \{ \theta_x | c_{photo}(\theta_x) < \tau_{cred} \}, \quad (5)$$

where τ_{cred} represents the credible threshold. Then, the Delaunay triangulation is applied to generate triangular surfaces of various sizes. Although these triangles cover a large area, there are still some marginal areas in the image that remain unaffected. This suggests that the local information




Figure 4: Planar prior models at different scales and resulting depth maps. (a) The color image (the low-textured area is in the white box while pixels in the red box are details); (b) the planar prior model constructed at the low-resolution scale; (c) the planar prior model constructed at the high-resolution scale (the dark blue areas show that there are anomalies and have already been filtered out); (d) the depth map obtained without prior models; (e) the depth map based on the prior model constructed at the low-resolution scale; (f) the depth map based on prior models at different scales.

cannot satisfy the requirements of developing a planar prior model in these parts of the regions. To address this issue, we propose a method based on KNN to construct the planar prior model via non-local credible hypotheses. First, we set up a KD-tree [47] for the reliable points in I_{cred} . Second, the top- K nearest non-local neighbors are searched for pixels in the regions where prior construction is challenging. In addition, we use Heron’s formula to avoid the case of co-linearity of proximity points. Then, each pixel will have three reliable neighboring points. We project them into the coordinate of the reference camera and set up a matrix \mathbf{A} to get the plane parameters \mathbf{z}_{opt} ,

$$\mathbf{z}_{opt} = \arg \min_{\mathbf{z}^*: \|\mathbf{z}^*\|=1} \|\mathbf{A} \cdot \mathbf{z}^*\|, \quad (6)$$

where \mathbf{z}^* is a unit-length variable. To further enhance the validity of the prior model, the ultimate prior knowledge for each image is filtered by its depth range. At last, we consider Planar Prior Assistance [40] which can better leverage the planar prior model to assist multi-view matching cost evaluation in low-textured areas.

3.4. Hierarchical Prior Mining

The multi-scale structure is employed in some Patch-Match MVS methods to infer the depth information of low-textured areas. Previous approaches [21, 38, 33] usually construct image pyramids to fuse depth estimation at different scales. In addition, several methods [40, 24] focus on constructing planar priors models at a single scale to help

the depth estimation. In our work, we demonstrate that constructing planar prior models in a multi-scale way will further facilitate depth estimation.

As depicted in Fig. 4(b), when the planar prior model is built in low-textured regions, the performance at the low-resolution scale is obviously more advanced than the high-resolution scale. Because the information is condensed after downsampling, resulting in a larger receptive field for each pixel, and a wider range of non-local information can be excavated than the original scale. However, only using the prior from the coarsest scale may lead ambiguity in detail areas. We find that the aforementioned problem of prior construction is expertly managed at the high-resolution scale (Fig. 4(c)).

Therefore, the depth perception advantages of planar prior models at different scales are complementary. Inspired by this key observation, we design an HPM framework to construct prior models at different scales, which assists the depth estimation of low-textured regions without sacrificing details. The depth map can be continually optimized via the HPM framework (Fig. 4(d)-(f)).

Excavating hierarchical prior is essentially an iterative coarse-to-fine optimization process (Fig. 2). After employing the basic MVS method with NESP, we first downsample the obtained hypotheses (*i.e.* depth and normal) to the coarsest scale and build a planar prior model. Second, with the joint bilateral upsampler [18], we propagate the planar prior to the original scale. Third, the planar prior model is embedded into the PatchMatch MVS method with Planar Prior Assistance [36] to approximate the depth in low-textured areas better. Finally, we downsample the newly generated hypotheses to the medium scale again, and repeat the same operation until iterating up to the original scale. For images with high resolutions, more medium scales can be considered. To further optimize the results, we apply geometric consistency as [25, 38] do. This way, the final depth estimation can achieve a better trade-off between details and low-textured areas.

Moreover, to save computing resources and improve time efficiency, we propose a fast version of the HPM framework. This version only uses the planar prior model constructed at the coarsest scale to assist MVS. The rationale behind this is that the prior model obtained at the coarsest scale can effectively target the low-textured region and exactly compensate for the deficiencies of the basic MVS method in that region. Both HPM and its fast version will be experimentally analyzed.

4. Experiments

We evaluate our method on two challenging MVS datasets, ETH3D benchmark [26] and Tanks & Temples datasets [17], from two perspectives, *i.e.*, benchmark performance and analysis experiments.

Table 1: Point cloud evaluation on ETH3D benchmark [26]. We show accuracy / completeness / F_1 score (in %) at different thresholds (2cm and 10cm). The best results are marked in **bold** while the second-best results are underlined.

	Method	2cm	10cm	
Train.	<i>i) Traditional</i>			
	COLMAP [25]	91.85 / 55.13 / 67.66	98.75 / 79.47 / 87.61	
	PCF-MVS [19]	84.11 / 75.73 / 79.42	95.98 / 90.42 / 92.98	
	ACMM [38]	90.67 / 70.42 / 78.86	98.12 / 86.40 / 91.70	
	ACMP [40]	90.12 / 72.15 / 79.79	97.97 / 87.15 / 92.03	
	ACMMP [36]	91.03 / <u>77.27</u> / <u>83.42</u>	97.96 / 93.19 / 95.46	
	<i>ii) Learning</i>			
	PatchMatchNet [30]	64.81 / 65.43 / 64.21	89.98 / 83.28 / 85.70	
	IterMVS [29]	73.62 / 61.87 / 66.36	94.48 / 78.85 / 85.25	
	MVSTER [32]	68.08 / 76.92 / 72.06	91.97 / 91.91 / 91.73	
	HPM-MVS _{fast}	<u>91.17</u> / 73.20 / 80.86	<u>98.23</u> / 92.41 / 94.97	
	HPM-MVS	90.66 / 79.50 / 84.58	97.97 / 95.59 / 96.22	
	Test	<i>i) Traditional</i>		
		COLMAP [25]	91.97 / 62.98 / 73.01	<u>98.25</u> / 84.54 / 90.40
PCF-MVS [19]		82.15 / 79.29 / 80.38	92.12 / 91.26 / 91.56	
ACMM [38]		90.65 / 74.34 / 80.78	98.05 / 88.77 / 92.96	
ACMP [40]		90.45 / 75.58 / 81.51	97.47 / 88.71 / 92.62	
ACMMP [36]		91.91 / 82.10 / <u>85.89</u>	98.05 / 94.67 / 96.27	
<i>ii) Learning</i>				
PatchMatchNet [30]		79.71 / 77.46 / 73.12	91.98 / 92.05 / 91.91	
IterMVS [29]		76.91 / 72.65 / 74.29	95.42 / 85.81 / 90.15	
MVSTER [32]		77.09 / 82.47 / 79.01	94.21 / 92.71 / 92.30	
HPM-MVS _{fast}		92.50 / 80.25 / 85.35	98.32 / <u>94.89</u> / <u>96.51</u>	
HPM-MVS		<u>92.13</u> / 83.25 / 87.11	98.11 / 95.41 / 96.69	

4.1. Experimental Setup

ETH3D high-resolution multi-view benchmark [26] consists of training datasets and test datasets, which contain indoor and outdoor images at a resolution of 6048×4032 ¹. Tanks & Temples datasets are divided into Intermediate and Advanced datasets based on the complexity of the scene. However, these datasets do not provide official camera poses. In that case, we use COLMAP [25] to estimate camera parameters.

All of our experiments are executed on a computer with an Intel i5-12600 CPU and an RTX 3070 GPU. In the basic MVS with NESP, $\{k, N_{ext}, \tau_{good}, \tau_{bad}, \alpha, n_{good}, n_{bad}, R\} = \{4, 3, 0.8, 1.2, 90, 1, 2, 4\}$. In the process of planar prior construction, $\{\tau_{cred}, K\} = \{0.1, 6\}$.

4.2. Benchmark Performance

Evaluation on ETH3D Benchmark. Both traditional methods [25, 19, 38, 40, 36] and learning-based [30, 29, 32] methods are considered for comparison. Table 1 reports the accuracy, completeness and F_1 score of the point clouds.

As shown in the table, HPM-MVS can achieve outstanding performance in terms of completeness and F_1 score among all the methods, and HPM-MVS_{fast} can outperform in accuracy while guaranteeing completeness and F_1 score. Notably, HPM-MVS obtains the highest completeness than other methods on both training and test datasets which con-

¹We downsample the undistorted images to 3200×2130 as in [25].




Figure 5: Qualitative point cloud comparisons between different MVS methods on several high-resolution multi-view test scans (old computer and terrace 2) of ETH3D benchmark [26]. Some challenging areas are shown in red boxes.

Table 2: F_1 score (in %) comparisons of different methods on Tanks & Temples datasets [17] at given evaluation threshold.

	Method	Intermediate	Advanced
Traditional	COLMAP [25]	42.14	27.24
	PCF-MVS [19]	55.88	34.59
	ACMM [38]	57.27	34.02
	ACMP [40]	58.41	37.44
	ACMMP [36]	59.38	37.84
Learning	R-MVSNet [44]	48.40	24.91
	PatchMatch-RL [20]	51.81	31.78
	PatchMatchNet [30]	53.15	32.31
	IterMVS [29]	56.22	33.24
	PVSNet [37]	56.88	33.46
	Effi-MVS [31]	56.88	34.39
	MVSTER [32]	60.92	37.53
Ours	HPM-MVS _{fast}	61.59	39.65
	HPM-MVS	61.39	40.80

tain large low-textured areas, because HPM-MVS perceives non-local structured information for recovering geometry. The visual comparisons are shown in Fig. 5. It can be clearly observed that our methods produce higher quality point clouds than other competitors, especially in challenging low-textured areas, *e.g.*, red boxes in Fig. 5.

Evaluation on Tanks & Temples Datasets. To further illustrate the robustness of our methods, we test our methods on Tanks & Temples datasets *without any fine-tuning*. The quantitative results of state-of-the-art methods on both Intermediate and Advanced sets are reported in Table 2.

Our methods achieve outstanding performance among all the methods. *Even compared with learning-based methods, our methods still achieve better performance without any data training.* Remarkably, HPM-MVS can greatly improve the performance of Advanced datasets which contain complex details and more low-textured areas difficult to recover. Our method ranks 1st in Advanced datasets over all the existing traditional MVS methods, and it outperforms the previous best record by **2.96%**. This result well confirms that our method can achieve a good balance for MVS reconstruction in details and low-textured areas. Fig. 6

presents some visualization examples of point cloud recall comparisons between different methods, which verifies our proposed methods can produce more complete 3D geometries than the others. These obvious advantages suggest that HPM-MVS and HPM-MVS_{fast} hold powerful generalization ability and can adapt to different application scenarios without any data training.

4.3. Analysis Experiments

Ablation Study. In this section, we conduct an ablation study on the multi-view high-resolution training datasets of ETH3D benchmark to validate the performance of different parts in our proposed methods. We divide them into five sections, which are Non-local Sampling Pattern, Extensible Sampling Pattern, planar prior model assistance, Hierarchical Prior Mining framework fast version and Hierarchical Prior Mining framework. The baseline method is ACMH² [38]. Table 3 summarizes the effectiveness of each individual part in our proposed methods.

For the basic MVS method with NESP, we remove the non-local part and the extensible part, respectively. The results show that these two proposals can both enhance the completeness and F_1 score. In particular, the improvement of the non-local strategy is more obvious than the extensible strategy. In some circumstances, especially in low-textured areas, several severe errors always cause the hypothesis propagation into local optimal solutions. Hence, the non-local operation can effectively prevent this situation from getting worse.

Based on NESP, the planar prior model constructed at the current scale generates a potential hypothesis for each pixel in the image. As for two versions of the HPM framework, one can make several observations from the results: 1) HPM and HPM_{fast} achieve outstanding performance at relatively high thresholds (*e.g.*, 5cm and 10cm), because they have a strong focus on sizable low-textured areas. 2) We notice that HPM_{fast} will suffer some loss of completeness at low thresholds (*e.g.*, 1cm), because it obtains the prior model

²Note that, we combine ACMH with geometric consistency here.




Figure 6: Qualitative point cloud recall comparisons between MVS methods on Tanks & Temples datasets [17] (Horse and Temple). The darker the pixel color is, the greater the error will be. τ is the recommended threshold for each dataset.

Table 3: Comparative results of our methods under different settings on the training datasets of ETH3D benchmark [26]. N, E, PA, HPM_{fast} and HPM mean Non-local Sampling Pattern, Extensible Sampling Pattern, planar prior model assistance, Hierarchical Prior Mining framework fast version and Hierarchical Prior Mining framework, respectively.

Method	Settings					ETH3D train.			
	N	E	PA	HPM _{fast}	HPM	1cm	2cm	5cm	10cm
ACMH [38]						82.41 / 50.63 / 61.71	91.09 / 63.04 / 73.60	96.61 / 74.73 / 83.50	98.25 / 82.20 / 88.57
NSP	✓					82.52 / 53.46 / 63.84	90.86 / 65.59 / 75.26	<u>96.63</u> / 76.15 / 84.52	<u>98.46</u> / 82.40 / 89.19
ESP		✓				82.66 / 51.40 / 62.22	90.93 / 63.83 / 74.05	96.64 / 75.37 / 84.04	98.47 / 82.09 / 89.03
NESP	✓	✓				82.47 / 53.98 / 64.22	90.78 / 66.02 / 75.59	96.60 / 76.63 / 84.87	98.47 / 82.73 / 89.44
NESP+PA	✓	✓	✓			82.88 / <u>62.37</u> / <u>71.32</u>	90.51 / 74.96 / 81.66	96.48 / 84.29 / 88.90	97.89 / 88.57 / 92.73
HPM _{fast}			✓	✓		83.32 / 53.53 / 64.66	<u>91.16</u> / 70.48 / 79.04	96.39 / 84.67 / 89.84	97.89 / 88.57 / 92.73
HPM			✓		✓	<u>83.01</u> / 62.07 / 70.67	90.65 / <u>77.03</u> / <u>83.06</u>	96.06 / <u>88.42</u> / <u>91.97</u>	97.91 / <u>93.62</u> / <u>95.66</u>
HPM-MVS _{fast}	✓	✓	✓	✓		82.95 / 57.36 / 67.39	91.17 / 73.20 / 80.86	96.51 / 86.22 / 90.89	98.23 / 92.14 / 94.97
HPM-MVS	✓	✓	✓		✓	82.93 / 65.16 / 72.73	90.66 / 79.50 / 84.58	96.13 / 89.98 / 92.88	97.97 / 94.59 / 96.22

Table 4: Generalization performance of NESP on ETH3D high-resolution training datasets [26].

Method	ETH3D train.				
	1cm	2cm	5cm	10cm	20cm
ACMM [38]	67.58	78.86	87.68	91.70	94.41
ACMP [40]	68.72	79.79	88.32	92.03	94.43
ACMMP [36]	71.57	83.42	92.03	95.46	97.37
ACMM+NESP	70.70	81.01	88.80	92.26	94.55
	3.12↑	2.15↑	1.12↑	0.56↑	0.14↑
ACMP+NESP	70.87	81.45	89.43	92.72	94.78
	2.15↑	1.66↑	1.11↑	0.69↑	0.35↑
ACMMP+NESP	74.54	85.33	93.25	96.45	97.99
	2.97↑	1.91↑	1.22↑	0.99↑	0.62↑

at the lowest resolution scale which may lead to ambiguities in details. 3) The method with NESP as the basis for subsequent processes can provide a set of reliable points for planar prior model construction.

Generalization Performance of NESP. In addition, we perform a more extensive comparison on the multi-view high-resolution training datasets of ETH3D benchmark to show the generalization ability of NESP. Several kinds of state-of-the-art diffusion-like propagation methods are integrated with NESP for evaluation. The considered methods are ACMM [38], ACMP [40] and ACMMP [36]. Note that ACMMP is a very recent state-of-the-art geometric-only

method. Each method is tested under different thresholds, and results are shown in Table 4.

Impressively, NESP dramatically improves the F_1 score for all tested methods at relatively low thresholds, e.g., 1cm and 2cm. The improvement of NESP working with multi-scale strategy is more obvious, because the non-local and extensible strategy works even better at low-resolution scale. Notably, the combination of NESP and ACMMP outperforms all the competitors and achieves state-of-the-art F_1 score of **85.33% / 96.45%** at 2cm / 10cm thresholds. Based on the results, the following conclusions can be drawn: 1) NESP is a generalized module that can be adapted to other MVS methods and boost their performance. 2) Combined with the multi-scale strategy, NESP can achieve better performance.

More analysis experiments and point cloud visualizations can be found in the supplementary.

5. Conclusion

In this paper, we presented an HPM-MVS method which can efficiently perceive non-local structured information to achieve high-quality reconstruction. Based on the non-local and extensible strategy, we first propose our NESP module. Focusing on the planar prior construction in marginal regions, we employ KNN to search non-local credible points

and obtain potential hypotheses. Further, we design an HPM framework which explores planar prior at different scales to assist MVS. Through extensive experiments, we have demonstrated the rationality of each key component of our method and its superiority against existing methods. In the future, we plan to combine our method with a probabilistic graphical model to further improve its ability to handle the reconstruction of details and low-textured regions.

Acknowledgements. This work is supported in part by the National Natural Science Foundation of China (No. 62002295), and NWPU international cooperation and exchange promotion projects (No. 02100-23GH0501).

References

- [1] Henrik Aanæs, Rasmus Ramsbøl Jensen, George Vogiatzis, Engin Tola, and Anders Bjarholm Dahl. Large-scale data for multiple-view stereopsis. *International Journal of Computer Vision*, 120(2):153–168, 2016. [1](#)
- [2] Christian Bailer, Manuel Finckh, and Hendrik Lensch. Scale robust multi view stereo. In *Proceedings of the IEEE European Conference on Computer Vision*, pages 398–411. Springer, 2012. [2](#)
- [3] Caroline Baillard and Andrew Zisserman. A plane-sweep strategy for the 3d reconstruction of buildings from multiple images. *International Archives of Photogrammetry and Remote Sensing*, 33(B2; PART 2):56–62, 2000. [1](#)
- [4] Connelly Barnes, Eli Shechtman, Adam Finkelstein, and Dan B Goldman. Patchmatch: A randomized correspondence algorithm for structural image editing. *ACM Transactions on Graphics*, 28(3):24, 2009. [2](#)
- [5] Fabio Bruno, Stefano Bruno, Giovanna De Sensi, Maria-Laura Luchi, Stefania Mancuso, and Maurizio Muzzupappa. From 3d reconstruction to virtual reality: A complete methodology for digital archaeological exhibition. *Journal of Cultural Heritage*, 11(1):42–49, 2010. [1](#)
- [6] Shuo Cheng, Zexiang Xu, Shilin Zhu, Zhuwen Li, Li Erran Li, Ravi Ramamoorthi, and Hao Su. Deep stereo using adaptive thin volume representation with uncertainty awareness. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2524–2534, 2020. [3](#)
- [7] Robert T Collins. A space-sweep approach to true multi-image matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 358–363. Ieee, 1996. [1](#)
- [8] Yasutaka Furukawa and Jean Ponce. Accurate, dense, and robust multiview stereopsis. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32(8):1362–1376, 2009. [1](#)
- [9] Silvano Galliani, Katrin Lasinger, and Konrad Schindler. Massively parallel multiview stereopsis by surface normal diffusion. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 873–881, 2015. [1](#), [2](#), [3](#)
- [10] David Gallup, Jan-Michael Frahm, Philippos Mordohai, Qingxiang Yang, and Marc Pollefeys. Real-time plane-sweeping stereo with multiple sweeping directions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8. IEEE, 2007. [1](#)
- [11] David Gallup, Jan-Michael Frahm, and Marc Pollefeys. Piecewise planar and non-planar stereo for urban scene reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1418–1425. IEEE, 2010. [2](#)
- [12] Xiaodong Gu, Zhiwen Fan, Siyu Zhu, Zuozhuo Dai, Feitong Tan, and Ping Tan. Cascade cost volume for high-resolution multi-view stereo and stereo matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2495–2504, 2020. [3](#)
- [13] Richard Hartley and Andrew Zisserman. *Multiple view geometry in computer vision*. Cambridge university press, 2003. [4](#)
- [14] Mohsen Hejrati and Deva Ramanan. Analysis by synthesis: 3d object recognition by object reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2449–2456, 2014. [1](#)
- [15] Rasmus Jensen, Anders Dahl, George Vogiatzis, Engin Tola, and Henrik Aanæs. Large scale multi-view stereopsis evaluation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 406–413, 2014. [1](#)
- [16] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016. [1](#), [3](#)
- [17] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples: Benchmarking large-scale scene reconstruction. *ACM Transactions on Graphics*, 36(4):1–13, 2017. [1](#), [6](#), [7](#), [8](#)
- [18] Johannes Kopf, Michael F Cohen, Dani Lischinski, and Matt Uyttendaele. Joint bilateral upsampling. *ACM Transactions on Graphics*, 26(3):96–es, 2007. [6](#)
- [19] Andreas Kuhn, Shan Lin, and Oliver Erdler. Plane completion and filtering for multi-view stereo reconstruction. In *Proceedings of the German Conference on Pattern Recognition*, pages 18–32. Springer, 2019. [6](#), [7](#)
- [20] Jae Yong Lee, Joseph DeGol, Chuhan Zou, and Derek Hoiem. Patchmatch-rl: Deep mvs with pixelwise depth, normal, and visibility. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6158–6167, 2021. [7](#)
- [21] Jie Liao, Yanping Fu, Qingan Yan, and Chunxia Xiao. Pyramid multi-view stereo with local consistency. In *Computer Graphics Forum*, volume 38, pages 335–346. Wiley Online Library, 2019. [5](#)
- [22] Jiachen Liu, Pan Ji, Nitin Bansal, Changjiang Cai, Qingan Yan, Xiaolei Huang, and Yi Xu. Planemvs: 3d plane reconstruction from multi-view stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8665–8675, 2022. [1](#)
- [23] Xinzhu Ma, Zhihui Wang, Haojie Li, Pengbo Zhang, Wanli Ouyang, and Xin Fan. Accurate monocular 3d object detection via color-embedded 3d reconstruction for autonomous driving. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 6851–6860, 2019. [1](#)

- [24] Andrea Romanoni and Matteo Matteucci. Tapa-mvs: Textureless-aware patchmatch multi-view stereo. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 10413–10422, 2019. [2](#), [5](#)
- [25] Johannes L Schönberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *Proceedings of the IEEE European Conference on Computer Vision*, pages 501–518. Springer, 2016. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [26] Thomas Schops, Johannes L Schonberger, Silvano Galliani, Torsten Sattler, Konrad Schindler, Marc Pollefeys, and Andreas Geiger. 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*, pages 3260–3269, 2017. [1](#), [6](#), [7](#), [8](#)
- [27] Qi Shan, Riley Adams, Brian Curless, Yasutaka Furukawa, and Steven M Seitz. The visual turing test for scene reconstruction. In *International Conference on 3D Vision-3DV*, pages 25–32. IEEE, 2013. [1](#)
- [28] Shuhan Shen. Accurate multiple view 3d reconstruction using patch-based stereo for large-scale scenes. *IEEE Transactions on Image Processing*, 22(5):1901–1914, 2013. [2](#)
- [29] Fangjinhua Wang, Silvano Galliani, Christoph Vogel, and Marc Pollefeys. Itermvs: Iterative probability estimation for efficient multi-view stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8606–8615, 2022. [1](#), [6](#), [7](#), [8](#)
- [30] Fangjinhua Wang, Silvano Galliani, Christoph Vogel, Pablo Speciale, and Marc Pollefeys. Patchmatchnet: Learned multi-view patchmatch stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 14194–14203, 2021. [1](#), [3](#), [6](#), [7](#)
- [31] Shaoqian Wang, Bo Li, and Yuchao Dai. Efficient multi-view stereo by iterative dynamic cost volume. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8655–8664, 2022. [7](#)
- [32] Xiaofeng Wang, Zheng Zhu, Guan Huang, Fangbo Qin, Yun Ye, Yijia He, Xu Chi, and Xingang Wang. Mvster: Epipolar transformer for efficient multi-view stereo. In *Proceedings of the IEEE European Conference on Computer Vision*, pages 573–591. Springer, 2022. [1](#), [6](#), [7](#), [8](#)
- [33] Yuesong Wang, Tao Guan, Zhuo Chen, Yawei Luo, Keyang Luo, and Lili Ju. Mesh-guided multi-view stereo with pyramid architecture. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2039–2048, 2020. [5](#)
- [34] Jian Wei, Benjamin Resch, and Hendrik PA Lensch. Multi-view depth map estimation with cross-view consistency. In *Proceedings of the British Machine Vision Conference*, 2014. [2](#)
- [35] Oliver Woodford, Philip Torr, Ian Reid, and Andrew Fitzgibbon. Global stereo reconstruction under second-order smoothness priors. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 31(12):2115–2128, 2009. [2](#)
- [36] Qingshan Xu, Weihang Kong, Wenbing Tao, and Marc Pollefeys. Multi-scale geometric consistency guided and planar prior assisted multi-view stereo. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2022. [1](#), [3](#), [6](#), [7](#), [8](#)
- [37] Qingshan Xu, Wanjuan Su, Yuhang Qi, Wenbing Tao, and Marc Pollefeys. Learning inverse depth regression for pixelwise visibility-aware multi-view stereo networks. *International Journal of Computer Vision*, 130(8):2040–2059, 2022. [1](#), [7](#)
- [38] Qingshan Xu and Wenbing Tao. Multi-scale geometric consistency guided multi-view stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5483–5492, 2019. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [39] Qingshan Xu and Wenbing Tao. Learning inverse depth regression for multi-view stereo with correlation cost volume. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12508–12515, 2020. [3](#)
- [40] Qingshan Xu and Wenbing Tao. Planar prior assisted patchmatch multi-view stereo. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 12516–12523, 2020. [2](#), [3](#), [5](#), [6](#), [7](#), [8](#)
- [41] Zhenyu Xu, Yiguang Liu, Xuelei Shi, Ying Wang, and Yunan Zheng. Marmvs: Matching ambiguity reduced multiple view stereo for efficient large scale scene reconstruction. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5981–5990, 2020. [2](#)
- [42] Jiayu Yang, Wei Mao, Jose M Alvarez, and Miaomiao Liu. Cost volume pyramid based depth inference for multi-view stereo. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4877–4886, 2020. [3](#)
- [43] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet: Depth inference for unstructured multi-view stereo. In *Proceedings of the IEEE European Conference on Computer Vision*, pages 767–783, 2018. [1](#), [3](#)
- [44] Yao Yao, Zixin Luo, Shiwei Li, Tianwei Shen, Tian Fang, and Long Quan. Recurrent mvsnet for high-resolution multi-view stereo depth inference. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 5525–5534, 2019. [1](#), [3](#), [7](#)
- [45] Yao Yao, Zixin Luo, Shiwei Li, Jingyang Zhang, Yufan Ren, Lei Zhou, Tian Fang, and Long Quan. Blendedmvs: A large-scale dataset for generalized multi-view stereo networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1790–1799, 2020. [1](#)
- [46] Enliang Zheng, Enrique Dunn, Vladimir Jovic, and Jan-Michael Frahm. Patchmatch based joint view selection and depthmap estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1510–1517, 2014. [2](#), [5](#)
- [47] Kun Zhou, Qiming Hou, Rui Wang, and Baining Guo. Real-time kd-tree construction on graphics hardware. *ACM Transactions on Graphics*, 27(5):1–11, 2008. [5](#)