

When Epipolar Constraint Meets Non-local Operators in Multi-View Stereo

Tianqi Liu Xinyi Ye Weiyue Zhao Zhiyu Pan Min Shi* Zhiguo Cao

Key Laboratory of Image Processing and Intelligent Control, Ministry of Education; School of Artificial Intelligence and Automation, Huazhong University of Science and Technology, Wuhan 430074, China

{tianqiliu, xinyiye, zhaoweiye, zhiyupan, min.shi, zgcao}@hust.edu.cn

Abstract

Learning-based multi-view stereo (MVS) method heavily relies on feature matching, which requires distinctive and descriptive representations. An effective solution is to apply non-local feature aggregation, e.g., Transformer. Albeit useful, these techniques introduce heavy computation overheads for MVS. Each pixel densely attends to the whole image. In contrast, we propose to constrain non-local feature augmentation within a pair of lines: each point only attends the corresponding pair of epipolar lines. Our idea takes inspiration from the classic epipolar geometry, which shows that one point with different depth hypotheses will be projected to the epipolar line on the other view. This constraint reduces the 2D search space into the epipolar line in stereo matching. Similarly, this suggests that the matching of MVS is to distinguish a series of points lying on the same line. Inspired by this point-to-line search, we devise a line-to-point non-local augmentation strategy. We first devise an optimized searching algorithm to split the 2D feature maps into epipolar line pairs. Then, an Epipolar Transformer (ET) performs non-local feature augmentation among epipolar line pairs. We incorporate the ET into a learning-based MVS baseline, named ET-MVSNet. ET-MVSNet achieves state-of-the-art reconstruction performance on both the DTU and Tanks-and-Temples benchmark with high efficiency. Code is available at <https://github.com/TQTQliu/ET-MVSNet>.

1. Introduction

As a fundamental topic in computer vision, Multi-view stereo (MVS) aims to reconstruct dense 3D representations from multiple calibrated images given corresponding camera parameters. MVS can facilitate many applications such as automatic driving and augmented reality [12]. A critical part of MVS is to match pixels to find the corresponding points. This matching process heavily relies on the fea-

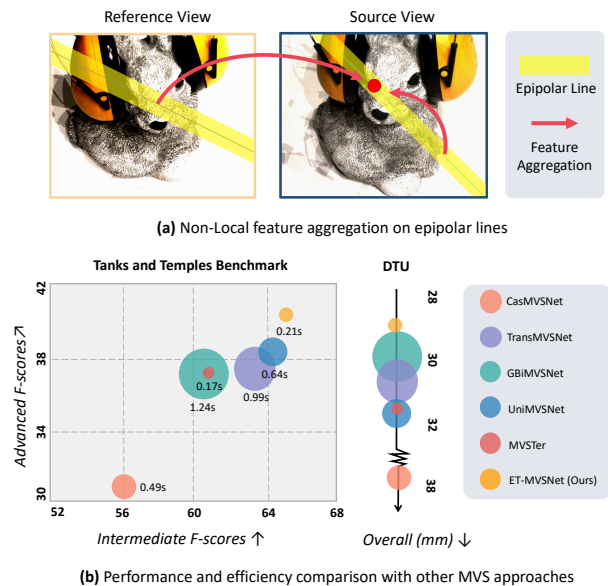


Figure 1. **Our methods and results.** (a) Inspired by classic epipolar geometry in stereo matching, we propose an efficient non-local feature aggregation paradigm that aggregates features on the corresponding epipolar line pairs. (b) We instantiate the proposed non-local mechanism with a Transformer, termed Epipolar Transformer (ET), which attains state-of-the-art performance with limited computation overhead.

ture representation: corresponding points in different views should be close in the embedding space. Traditional MVS approaches [11, 14, 27, 28] adopt hand-craft feature representation, which faces challenges on repetitive-pattern, weak-texture, and reflective regions.

To remedy the weak representation of the traditional MVS approaches, learning-based methods are proposed and show superior performance. Learning-based MVS approaches adopt deep neural networks to encode features, which forms the foundation for the subsequent feature matching and depth generation. Powerful as convolutional neural networks (CNNs) are, they still prefer to aggregate local context [22], which may be less accurate for matching.

*Corresponding author

One feasible solution is to adopt non-local feature augmentation strategy with large receptive field and flexible feature aggregation, which is crucial for robust and descriptive feature representations.

To this end, techniques such as deformable convolution [8], attention [2], and Transformers [29], are introduced to the feature encoding of MVS. Albeit useful, these general-purpose non-local operators introduce huge computation overhead for MVS, as each point need to densely attend the reference and source images. In this case, distractive or irrelevant features can also be attended which are harmful to the feature matching, *e.g.*, repetitive patterns shown in Fig. 7. This raises a fundamental question: *where and how to mine non-local context for MVS?* Inspired by the classic epipolar geometry, we propose to constrain the non-local feature aggregation within the epipolar lines, which enjoys both efficiency and descriptive representation.

According to the epipolar geometry theory, the actual point corresponding to a pixel in one view can only be projected on the epipolar line in the other view. Hence, when searching the corresponding points, the 2D search space can be reduced into a 1D line. This reduces the computation and eliminates background interference. Naturally, each source image can form a stereo matching problem with the reference image. As shown in Fig. 2, points with different depth hypotheses will be projected onto the epipolar line in the other image. Therefore, the feature encoding in MVS can be viewed as describing different points on the epipolar line.

Since the points to be distinguished lie on the same line, inspired by this epipolar constraint, we propose a line-to-point feature aggregation strategy. For each point, the information from the corresponding epipolar line is used to describe it. A similar line-based feature aggregation has been adopted by CSwin [10] for vision backbones, which is proved to be efficient and effective. We find that for MVS, with epipolar geometry as a theoretical basis, using epipolar lines as the source of non-local features can efficiently achieve high-quality representation, as shown in Fig. 1 (b).

Specifically, we propose an epipolar line-guided non-local mechanism, which only fetches features from the corresponding epipolar lines. First, we decompose the reference and source feature map into pixel groups; each group shares identical epipolar line pairs. This process is processed with an optimized search algorithm. For each point in the reference image, we calculate its corresponding epipolar line in the source image. Then, points with approximate epipolar line parameters are clustered into the same group. With this step, all the pixels in the source and reference features can find their epipolar line pairs. Then, as shown in Fig. 1 (a), we perform non-local feature aggregation within the source epipolar line and across the epipolar line pairs. We also add a local augmentation module to mitigate the discontinuities of epipolar line partition. We instan-




Figure 2. **Epipolar Geometry.** A pixel with different depth hypotheses in the reference image can be projected to the same epipolar line in the source image, termed point-to-line. The plane that passes through two camera centers and a 3-D point produces an epipolar line when intersecting with an image plane. The two epipolar lines of two images form an epipolar line pair, termed line-to-line.

tiate this epipolar line-guided non-local mechanism with a Transformer model termed Epipolar Transformer (ET).

To demonstrate the effectiveness of the proposed modules, we present a coarse-to-fine framework termed ET-MVSNet, which applies the Epipolar Transformer (ET) to the feature encoding part of recent works [15, 31]. We conduct extensive experiments to demonstrate the effect of our method. Thanks to Epipolar Transformer (ET), the depth estimation and 3-D representations have greatly improved and achieved state-of-the-art results on both the DTU [1] and the Tanks and Temple [19] benchmark. Besides, compared with the prevailing global attention that mines global context, Epipolar Transformer (ET) shows significantly better performance and efficiency.

Our main contributions can be summarized as follows:

- We propose an epipolar line feature aggregation strategy for the representation learning of MVS.
- We devise an optimized epipolar pair search algorithm, and an Epipolar Transformer (ET) to enable non-local feature augmentation on epipolar lines.
- Our method achieves state-of-the-art performance on both the DTU dataset and Tanks-and-Temples benchmark.

2. Related Work

Learning-based MVS. The core of MVS is a feature-matching problem, where representation and similarity metrics play critical roles. The performance of traditional MVS is limited due to poor representation ability, especially in weak-texture or repetitive areas. Powered by the strong representation of neural networks, the learning-based MVS is first proposed by MVSNet [35] and quickly becomes the main stream due to better robustness and performance.

MVSNet proposes an end-to-end pipeline for depth estimation with a post-processing procedure for fusion. Subsequent work improves the learning-based MVS from different aspects. *e.g.*, reducing memory and computation overhead with RNN-based approaches [32, 33, 36] and coarse-to-fine paradigms [6, 15, 34, 39], optimizing the cost aggregation by adaptively re-weighting different views [21, 38]. An important line is to enhance the feature representation with a non-local feature aggregation mechanism, *e.g.*, Deformable Convolution Network [8]. However, existing non-local techniques suffer from expensive computational overheads attributable to multi-image inputs. In this paper, we focus on the characteristics of the MVS task and propose a novel and efficient feature enhancement method.

Feature Representation in MVS. To obtain descriptive representation for matching, Dai *et al.* [8] utilizes deformable convolution [8] to capture instance cues for splitting foreground and background, such that the context of them would not interact and the feature representations are not sensitive to the perspective changing. Transformer [29] was first introduced into the MVS pipeline by TransMVSNet [9] for its global receptive field and dynamic feature aggregation ability. TransMVSNet applies intra- and inter-image attention to transfer context and uses a linear Transformer [18] to reduce computational cost. MVSTER [31] employs a parameter-free attention mechanism to re-weight the aggregated contribution of different pixels. The success of the attention mechanism highlights the significance of learning robust feature representations, as the essence of MVS lies in feature matching. Although more advanced feature aggregation approaches such as attention are introduced for MVS, the efficiency decreases significantly. The existing Transformer-based MVS methods adopt efficient attention operators, while proper constraints on the feature aggregation remain unstudied. In this paper, inspired by epipolar geometry constraints, we propose to perform non-local feature augmentation only on the epipolar pairs.

3. Preliminaries

Learning-based MVS. Most learning-based methods follow a cost volume-based pipeline first introduced by MVSNet [35]. Given a reference image and $N - 1$ source images, learning-based methods aim to recover the depth map of the reference image. The pipeline consists of four steps. The first step is the feature encoding where images are projected in feature representation. In the second step, the cost volume is constructed by feature matching. Specifically, according to a set of predefined depth hypotheses, each source feature map is warped into reference camera frustum via differentiable homography to measure their similarity and obtain feature volumes. These $N - 1$ feature volumes are then aggregated to construct a 3D cost volume that encodes the cost of each depth hypothesis for the reference image.




Figure 3. **Visualization of the epipolar pairs.** The same color represents the corresponding epipolar lines.

After that, a 3D CNN can be applied to the cost volume for regularization and obtain the probability volume to infer depths. As mentioned above, the basic idea of depth estimation is to perform point matching to estimate final depths based on predefined depth hypotheses, where the feature representation capacity directly affects the accuracy of matching. Therefore, we aim to propose an effective and efficient information aggregation strategy to capture non-local context and enhance feature representation.

Epipolar Constraint. The epipolar constraint is a common geometric constraint in stereo matching. Given a pair of stereo images, stereo matching aims to find the corresponding points between the two images. As shown in Fig. 2, for a given point in one image, its corresponding point in the other image must lie on a specific line, which is called the epipolar line. Geometrically, an epipolar line is an intersection between the image plane and the plane passing through the two camera centers and the corresponding 3D point. In addition, epipolar lines occur in pairs, meaning that any point on one epipolar line corresponds to the same epipolar line in the other image.

With this constraints, the search space of the stereo matching [7, 3, 13, 16, 26], can be reduced from the whole image plane into a single line. MVS can be viewed as a multiple stereo matching problems. Points with different depth hypotheses in the reference view will be projected on the epipolar line in the source images. Hence, the feature encoding can be viewed as describing different points which are on the same line.

4. ET-MVSNet

The overall architecture of our ET-MVSNet is illustrated in Fig. 4. The proposed Epipolar Transformer (ET) module is integrated into the Feature Pyramid Network (FPN) [20]. To perform non-local feature aggregation on epipolar lines, we first search the epipolar line pairs between reference and source images. With the searched line pairs, the original feature maps are decomposed into different epipolar line pairs via sampling. Then, the Intra-Epipolar Augmentation (IEA) and the Cross-Epipolar Augmentation (CEA) modules transfer non-local context among these epipolar line pairs.

Algorithm 1: Epipolar Pair Search

Input: Camera parameters; image size (H, W) ; parameters matrices \mathbf{K}, \mathbf{B} of ks, bs ; search threshold δ .

Output: Epipolar pairs

```
1 for  $(x, y) = (0, 0)$  to  $(W - 1, H - 1)$  do
2   | calculate slope  $\mathbf{K}^{x,y}$  and bias  $\mathbf{B}^{x,y}$ 
3 end
4 for  $k$  in  $set(Quantify(\mathbf{K}))$  do
5   | for  $b$  in  $set(Quantify(\mathbf{B}))$  do
6     | for  $(x, y) = (0, 0)$  to  $(W - 1, H - 1)$  do
7       | epipoline_ref :  $\mathbf{K}^{x,y} = k$  and  $\mathbf{B}^{x,y} = b$ 
8       | epipoline_src :  $\frac{|y-kx-b|}{\sqrt{1+k^2}} < \delta$ 
9     | end
10  | end
11 end
```

4.1. Epipolar Pair Search

The key component of MVS is to match pixels to the most suitable depth hypothesis from a set of pre-defined hypotheses through feature matching. If a hypothesis is close to the actual depth of the ground, the features of the pixel pairs are supposed to be similar, indicating that the discrimination of the hypothesis relies heavily on feature representation. Since the feature volume used for matching naturally lies on the epipolar line, using epipolar lines as the source of non-local features can efficiently achieve high-quality representation, benefiting in distinguishing different hypotheses. It should be noted that epipolar lines exist in both the reference view and the source view and appear in pairs due to geometric constraints. It suggests the pixels on the same epipolar line share the same source of non-local features and the process of aggregating features for these pixels is the same. A parallel aggregation would improve the efficiency. To realize the parallel process, pixels located on the same epipolar line pair are required to be pre-searched and we propose an epipolar-pair search algorithm for it.

As shown in Algorithm 1, the searching of epipolar pairs consists of two steps: point-to-line search and epipolar line matching. In the first step, we compute the epipolar line parameters for each pixel in the reference view. Then, we aggregate the pixels in the reference view into different clusters; pixels in the same cluster share the same epipolar line. Hence, after the epipolar pair search, pixels in the reference and source images are partitioned into epipolar line pairs.

Since a one-to-many view matching can be decomposed into multiple one-to-one view matching, for easy understanding, here we illustrate the epipolar pair search with one source view. Given a pixel \mathbf{p}_r in the reference view,

the corresponding pixel \mathbf{p}_s in the source view is

$$\mathbf{p}_s(d) = \mathbf{K}_s[\mathbf{R}(\mathbf{K}_r^{-1}\mathbf{p}_r, d) + \mathbf{t}], \quad (1)$$

where d denotes the depth. \mathbf{R} and \mathbf{t} denote the rotation and translation between the reference and the source view. \mathbf{K}_r and \mathbf{K}_s denote the intrinsic matrices of the reference view and source view, respectively. Hence, the coordinate of $\mathbf{p}_s(d)$ can be computed by

$$x_s(d) = \frac{a_1d + b_1}{a_3d + b_3}, y_s(d) = \frac{a_2d + b_2}{a_3d + b_3}, \quad (2)$$

where $\{a_i\}_{i=1}^3$ and $\{b_i\}_{i=1}^3$ are constants associated with the camera parameters and the coordinate of \mathbf{p}_r (refer to the Supplementary Materials for more details). Then, we can eliminate the depth d to obtain the standard equation $y_s(d) = kx_s(d) + b$ of the epipolar line, which is formulated as:

$$\begin{cases} k = \frac{\Delta y_s(d)}{\Delta x_s(d)} = \frac{a_2b_3 - a_3b_2}{a_1b_3 - a_3b_1} \\ b = y_s(0) - kx_s(0) = \frac{b_2}{b_3} - k\frac{b_1}{b_3} \end{cases}. \quad (3)$$

Specifically, when $\Delta x_s(d) \rightarrow 0$, we use the equation $x_d(s) = k y_s(d) + b'$. As k, b are independent of the d and are expressed by constants associated with the coordinate of \mathbf{p}_r , if two pixels are on the same epipolar line, the ks and bs of them would be theoretically the same. However, in practice, due to the discrete coordinates of pixels, the ks and bs calculated by different pixels around the epipolar line may vary in a tiny range, such that directly grouping the pixels with identical ks or bs will lead to over-splitting.

To alleviate the over-splitting, we quantify ks and bs by rounding. Therefore, pixels with approximate parameters will be grouped into the same epipolar line. Have finished the grouping of the reference view. Then, we search for the corresponding epipolar line in the source view. The ks and bs represent the corresponding epipolar lines on the source view, as the coordinates of pixels around the epipolar line will meet Eq. (3). For the source image, whether a pixel is on the epipolar line expressed by k and b is determined by calculating the distance of them.

By dividing pixels into their corresponding epipolar pairs, both the reference and the source feature maps are decomposed into sets of feature sequences. Specifically, supposing that m pairs of epipolar lines are obtained, we define the reference and source feature sets as $E_{\mathcal{R}}$ and $E_{\mathcal{S}}$ which are formulated as:

$$E_{\mathcal{R}} = \{E_{\mathcal{R}}^1, E_{\mathcal{R}}^2, \dots, E_{\mathcal{R}}^m\}, E_{\mathcal{S}} = \{E_{\mathcal{S}}^1, E_{\mathcal{S}}^2, \dots, E_{\mathcal{S}}^m\}. \quad (4)$$

$E_{\mathcal{R}}^i$ and $E_{\mathcal{S}}^i$ are feature sequences with the shape of $n \times c$, where n denotes the number of pixels in the corresponding epipolar line, and c denotes the feature dimension. We visualize the searched epipolar line pairs in Fig. 3.

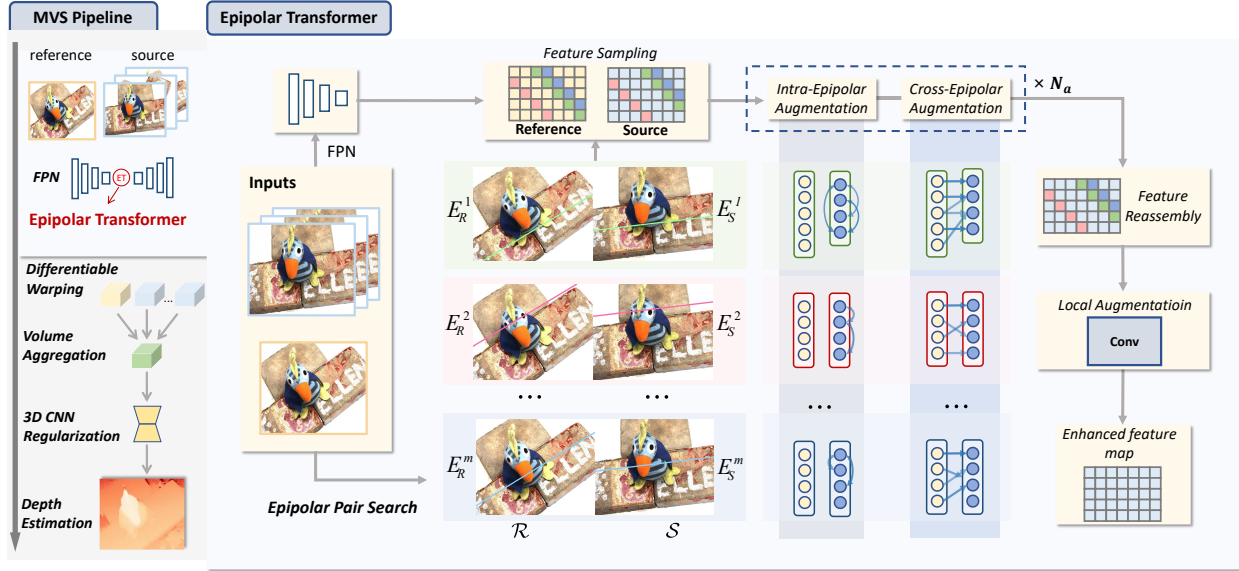


Figure 4. **ET-MVSNet architecture.** We apply the Epipolar Transformer (ET) to common learning-based MVS frameworks by integrating it into the FPN. Firstly, the epipolar pairs are searched via our algorithm and the epipolar feature sequences are obtained by sampling the coarsest feature map of the FPN. The sequence is then fed into the Intra-Epipolar Augmentation (IEA) and Cross-Epipolar Augmentation (CEA) blocks, which can be stacked N_a layers. In our implementation, we set $N_a = 1$. The resulting sequences are mapped back to a feature map, followed by the Local Augmentation (LA) module to smooth the enhanced feature map. Finally, the enhanced feature map is sent to the upsampling layers of the FPN, followed by cost volume construction and regularization.

4.2. Intra- and Cross-Epipolar Augmentation

With the epipolar line pairs $E_{\mathcal{R}}$ and $E_{\mathcal{S}}$, we then perform the non-local feature augmentation. Since points with different depth hypotheses fall on an epipolar line, the matching process is to distinguish a series of points lying on the same line. Inspired by the point-to-line searching strategy in stereo matching, we propose a line-to-point non-local augmentation: each pixel in the reference image only attends to its corresponding epipolar line pairs.

As shown in Fig. 4, to describe the pixel with non-local information in the source image, an Intra-Epipolar Augmentation (IEA) module is devised based on self-attention. A Cross-Epipolar Augmentation (CEA) module then propagates the information on reference epipolar line $E_{\mathcal{R}}$ into the source epipolar line $E_{\mathcal{S}}$ with cross-attention.

Intra-Epipolar Augmentation (IEA). IEA utilizes the self-attention within an epipolar line to aggregate non-local structural information, which can generate descriptive feature representations for difficult regions, *e.g.*, weak-texture areas. For each epipolar line $E_{\mathcal{S}}^i$ in the $E_{\mathcal{S}}$, the augmentation process is defined as:

$$E_{\mathcal{S}}^i = \text{MHSA}(E_{\mathcal{S}}^i) + E_{\mathcal{S}}^i, \quad (5)$$

where $\text{MHSA}(x)$ refers to the multi-head self attention [29] that takes a sequence x as input.

Cross-Epipolar Augmentation (CEA). Due to the varying viewpoints, the potential perspective transformation ren-

ders challenges for the same semantic pixels with different geometries. We alleviate this by transferring information across the epipolar lines. Specifically, we use a cross-attention module to propagate the information from reference line $E_{\mathcal{R}}$ into the $E_{\mathcal{S}}$, denoted by CEA. In CEA, $E_{\mathcal{S}}^i$ is first processed by a cross-attention layer:

$$E_{\mathcal{S}}^i = \text{MHCA}(E_{\mathcal{S}}^i, E_{\mathcal{R}}^i, E_{\mathcal{R}}^i) + E_{\mathcal{S}}^i, \quad (6)$$

where $\text{MHCA}(q, k, v)$ refers to the multi-head cross attention [29] layer with q , k and v as the input query, key and value sequences. Specifically, we add a feed-forward network as in the standard Transformer block [29] after the cross-attention layer. The block of IEA and CEA can be stacked for feature augmentation.

Local Augmentation (LA). Although the non-local feature augmentation within epipolar lines is efficient and proved to be effective in our experiments. We find that the enhanced feature map contains some holes where pixels that are located outside the common view of two images or are not detected by the algorithm due to the quantification error caused by the discrete nature of pixels, which can lead to discontinuities in feature representation and is unfriendly for matching. To remedy this issue, we use an additional convolution layer after the IEA and CEA block to re-aggregate the local context for filling feature holes and smoothing the augmented feature.




Figure 5. Comparison of reconstructed results with state-of-the-art methods [31, 9] on DTU evaluation set [1].

4.3. Implementation Details

Inherited from MVSTER [31], we use a four-stage depth estimation with 8, 8, 4, and 4 depth hypotheses, and adopt the inverse depth sampling. To simplify the training, the original monocular depth estimation auxiliary branch from MVSTER is removed. We integrate the proposed Epipolar Transformer (ET) into the FPN processing the coarsest feature map with a downsampling rate of 8 to save computation costs.

We use the cross-entropy loss for all stages. The total loss is as follows:

$$Loss = \sum_{k=1}^N \lambda^k L^k \quad (7)$$

where N refers to the number of stages, L^k refers to the loss at the k^{th} stage and λ^k refers to loss weight of the k^{th} stage. Specifically, $N = 4$ and $\lambda^1 = \dots = \lambda^N = 1$ in our implementation.

5. Experiments

In this section, we first introduce datasets, metrics, and experiment settings in Sec. 5.1, and then present the benchmark performance in Sec. 5.2. Additionally, we perform ablation studies to demonstrate the effectiveness of each component and the efficiency of our proposed feature aggregation approach in Sec. 5.3.

5.1. Settings

DTU dataset. DTU is an indoor dataset collected under well-controlled laboratory conditions with a fixed camera trajectory. The DTU dataset contains 128 different scenes with 49 or 64 views under 7 different illumination conditions. We adopt the same data partition as MVSNet [35]

for a fair comparison. We take 5 images and follow the dynamic checking strategy [33] for depth filtering and fusion. Accuracy and completeness are used as the evaluation metrics. In addition, an “overall” metric which averages the accuracy and completeness is reported.

Tanks and Temples benchmark. “Tanks and Temples” is a public benchmark consisting of 14 outdoor realistic scenes, including an intermediate subset of 8 scenes and an advanced subset of 6 scenes. We take 11 images and follow the dynamic checking strategy [33] for depth filtering and fusion. Note that the ground truth for the Tanks and Temples benchmark is hidden, all results are evaluated on an official website. We use F-score as the evaluation metric, which combines both precision and recall. Refer to the Supplementary Materials for more details of depth fusion.

Following common practices [15, 9, 25, 31], before evaluating on Tanks and Temples benchmark, we finetune our model on the BlendedMVS dataset [37] for the adaptation to the real-world scenes. We use the original image resolution 576×768 and the number of images equals 7.

5.2. Benchmark Performance

Evaluation on DTU. We report standard metrics by using the official evaluation script. The quantitative results of the DTU evaluation set are shown in Table 2. Compared with other state-of-the-art methods, our method can reconstruct denser point clouds, which leads to advantageous completeness. With improved completeness, our method retains comparable accuracy and obtains the best overall metric. Fig. 5 shows qualitative results on the DTU evaluation set. Our method can reconstruct more detailed point clouds in challenging areas, such as weakly textured surfaces.

Evaluation on Tanks and Temples. To verify the generalization ability of our method, we evaluate our method on the Tanks and Temples benchmark. The quantitative results of intermediate and advanced sets are shown in Table 1. Our

Method	Intermediate									Advanced						
	Mean	Fam.	Fra.	Hor.	Lig.	M60	Pan.	Pla.	Tra.	Mean	Aud.	Bal.	Cou.	Mus.	Pal.	Tem.
COLMAP [27]	42.14	50.41	22.25	26.63	56.43	44.83	46.97	48.53	42.04	27.24	16.02	25.23	34.70	41.51	18.05	27.94
Point-MVSNet [5]	48.27	61.79	41.15	34.20	50.79	51.97	50.85	52.38	43.06	-	-	-	-	-	-	-
R-MVSNet [36]	50.55	73.01	54.46	43.42	43.88	46.80	46.69	50.87	45.25	29.55	19.49	31.45	29.99	42.31	22.94	31.10
PatchmatchNet [30]	53.15	66.99	52.64	43.24	54.87	52.87	49.54	54.21	50.81	32.31	23.69	37.73	30.04	41.80	28.31	32.29
P-MVSNet [21]	55.62	70.04	44.64	40.22	<u>65.20</u>	55.08	55.17	60.37	54.29	-	-	-	-	-	-	-
CasMVSNet [15]	56.84	76.37	58.45	46.26	55.81	56.11	54.06	58.18	49.51	31.12	19.81	38.46	29.10	43.87	27.36	28.11
MVSTER [31]	60.92	80.21	63.51	52.30	61.38	61.47	58.16	58.98	51.38	37.53	26.68	42.14	35.65	49.37	32.16	<u>39.19</u>
GBi-Net [24]	61.42	<u>79.77</u>	<u>67.69</u>	51.81	61.25	60.37	55.87	60.67	53.89	37.32	29.77	42.12	36.30	47.69	31.11	36.93
AA-RMVSNet [32]	61.51	<u>77.77</u>	<u>59.53</u>	51.53	64.02	64.05	59.47	60.85	55.50	33.53	20.96	40.15	32.05	46.01	29.28	32.71
EPP-MVSNet [23]	61.68	77.86	60.54	52.96	62.33	61.69	60.34	62.44	55.30	35.72	21.28	39.74	35.34	49.21	30.00	38.75
TransMVSNet [9]	63.52	80.92	65.83	<u>56.94</u>	62.54	63.06	60.00	60.20	<u>58.67</u>	37.00	24.84	<u>44.59</u>	34.77	46.49	<u>34.69</u>	36.62
UniMVSNet [25]	<u>64.36</u>	<u>81.20</u>	66.43	53.11	63.46	66.09	64.84	<u>62.23</u>	57.53	<u>38.96</u>	28.33	44.36	39.74	52.89	33.80	34.63
ET-MVSNet (Ours)	65.49	81.65	68.79	59.46	65.72	<u>64.22</u>	<u>64.03</u>	61.23	58.79	40.41	<u>28.86</u>	45.18	<u>38.66</u>	<u>51.10</u>	35.39	43.23

Table 1. **Quantitative results on Tanks and Temples benchmark.** The metric is F-score and “Mean” refers to the average F-score of all scenes (**higher is better**). The best and the second-best results are in **bold** and underlined, respectively.

Method	ACC.(mm) ↓	Comp.(mm) ↓	Overall(mm) ↓
Gipuma [14]	0.283	0.873	0.578
COLMAP [27]	0.400	0.664	0.532
SurfaceNet [17]	0.450	1.040	0.745
P-MVSNet [21]	0.406	0.434	0.420
R-MVSNet [36]	0.383	0.452	0.417
Point-MVSNet [5]	0.342	0.411	0.376
AA-RMVSNet [32]	0.376	0.339	0.357
EPP-MVSNet [23]	0.413	0.296	0.355
PatchmatchNet [30]	0.427	0.277	0.352
CasMVSNet [15]	0.325	0.385	0.355
MVSTER [31]	0.350	0.276	0.313
UniMVSNet [25]	0.352	0.278	0.315
TransMVSNet [9]	<u>0.321</u>	0.289	0.305
GBiNet [24]	0.327	<u>0.268</u>	<u>0.298</u>
ET-MVSNet (Ours)	0.329	0.253	0.291

Table 2. **Quantitative point cloud results on DTU evaluation set.** (**lower is better**). The best and the second-best results are in **bold** and underlined, respectively.

method outperforms the other MVS approaches on both the intermediate and the advanced sets, which demonstrates the strong generalization ability of our method. Refer to Supplementary Materials for qualitative results.

5.3. Ablation Study

We conduct a series of ablation studies to analyze the effectiveness and efficiency of the proposed module, including different components in the Epipolar Transformer (ET) and different feature aggregation operators. All the ablation is conducted based on the MVSTER[31]. We conduct experiments on the DTU dataset with metrics in Table 2 and an additional depth error. The depth error refers to the average absolute values of the differences between predicted depths and ground-truth depths, whose resolution is 512×640 .

Settings			ACC.↓	Comp.↓	Overall↓	Depth Error↓
IEA	CEA	LA				
			0.351	0.284	0.318	6.355
✓	✓		0.331	0.263	0.297	5.843
✓		✓	0.327	0.261	0.294	6.021
	✓	✓	0.326	0.262	0.294	5.982
✓	✓	✓	0.329	0.253	0.291	5.754

Table 3. **Ablations on the DTU evaluation set.**




Figure 6. **Different information interaction approaches for images of different resolutions.** “P2L” refers to the interaction between a point and the corresponding epipolar line, and “L2L” refers to the interaction between epipolar pairs.

Benefits of ET As shown in Table 3, we ablate different components in the proposed ET. First, comparing row 1 with row 2, including the ICE and CEA to mining non-local context along the epipolar line decreases the depth error by 8%, leading to a +6% improvement in the quality (Overall) of 3-D reconstruction. Comparing row 2 with row 5, adding the LA to address the discontinuities on feature representation further yields a +7% improvement in 3-D reconstruction accuracy. Comparing row 3 and row 4 with row 5, only using IEA or CEA for intra-line and inter-line augmentation will slightly decrease the overall metrics, but can still obtain better performance than no non-local augmentation. The combination of ICE, CEA, and LA achieves the lowest depth error and the best 3-D reconstruction results.

Settings	Overall ↓	Depth Error ↓	MACs(G) ↓	Time(s) ↓
Point-to-Line	0.290	5.774	25.66	3.06
Line-to-Line	0.291	5.754	0.586	0.46

Table 4. **Comparison of line-to-line and point-to-line algorithms.** “MAC” refers to the multiply–accumulate operations for ET.

Method	Overall(mm) ↓	Depth Error(mm) ↓	Param(M) ↓
ET	0.291	5.754	1.09
FMT	0.303	6.064	1.42
DCN	0.309	6.161	1.20
ASPP	0.320	7.227	1.26

Table 5. **Comparison of different feature enhancement methods.**

Epipolar line search algorithm. Here we discuss the efficiency of different implementations for the line search algorithm. As illustrated in Sec. 4.1, we pre-calculate the epipolar lines for each pixel, and group pixels with approximate epipolar lines into the same cluster. We define this implementation as line-to-line, as the outputs are pairs of lines. A simple alternative can be a point-to-line implementation: calculate the line parameters for each pixel in the source view individually.

We compare two different implementations in Table 4. With comparable performance, point-to-line increases the computation overhead (MACs) by 50 times and takes six times longer for inference. Our line-to-line searching strategy with pre-clustering can help parallelize the attention in IEA and CEA, leading to improved efficiency. We report the inference time of different resolutions in Fig. 6. As the resolution scales up, the advantages of the line-to-line implementation will become more obvious.

Different feature aggregation methods As shown in Table 5, we compare four feature aggregation methods: the proposed Epipolar Transformer (ET), Feature Matching Transformer(FMT) [9], Deformable Convolutional Networks (DCN) [8], and Atrous Spatial Pyramid Pooling (ASPP) [4]. Our Epipolar Transformer outperforms others in terms of both depth and construction metrics with the least parameters. This suggests that, based on the epipolar geometry as prior, the non-local feature aggregation on the epipolar line can be more effective and efficient for MVS, compared with other general-use feature aggregation paradigms. We further visualize the attention map in Fig. 7. The feature matching transformer (FMT) in TransMVSNet [9] generates attention on some irrelevant backgrounds and the points with similar textures, which can interfere with the representation. In contrast, by constraining the attention on the line, our ET can concentrate on the corresponding points.

Results on other approaches. Our proposed Epipolar Transformer (ET) can serve as a plug-and-play module integrated into common learning-based MVS frameworks.




Figure 7. **Visualization of the attention map.** The attention map of our line-to-line Epipolar Transformer is between two corresponding epipolar lines and that of Feature Matching Transformer (FMT) [9] is between a pixel and the other image. Given a query point on an epipolar line of the source image, ET focuses on the region of its corresponding epipolar line while FMT considers the whole reference image. The red box points to the area with the highest attention score for the attention maps of our line-to-line ET and FMT, respectively.

Method	DTU Overall ↓	Int. Mean↑	Adv. Mean ↑
CasMVSNet [15]	0.355	56.84	31.12
CasMVSNet+ET	0.301	61.62	35.65

Table 6. **Compatibility experiments on CasMVSNet.** “+ ET” indicates embedding our ET into the coarsest stage of its feature encoding structure.

Here, we apply ET to another representative baseline CasMVSNet [15]. Based on the original setting and implementation, ET is inserted into the FPN with the lowest resolution. As shown in Table 6, CasMVSNet+ET improves the overall performance by 15% compared with the original implementation on the DTU dataset. On the tanks and temples benchmark, CasMVSNet+ET achieves 8% and 15% relative improvements on the intermediate and advanced sets. The significant performance improvement across multiple datasets demonstrates the flexibility of our method.

6. Conclusion

In this paper, we introduce a non-local feature aggregation strategy based on epipolar constraints. Specifically, we first propose an algorithm for searching epipolar pairs in two calibrated images. We then design an Intra-Epipolar Augmentation (IEA) module and a Cross-Epipolar Augmentation (CEA) module to mine non-local context within and across epipolar lines. We pack these modules into a Transformer model termed Epipolar Transformer (ET) and integrate the ET into a baseline network to construct the ET-MVSNet. Evaluations and ablations verify the effectiveness of the proposed modules. In addition, we show that ET can also serve as a plug-and-play module integrated into other MVS methods.

Acknowledgements. This work was funded in part by the National Natural Science Foundation of China (Grant No. U1913602) and supported by the Huawei Technologies CO., LTD.

References

- [1] Henrik Aanaes, Rasmus Ramsbol Jensen, George Vogiatzis, Engin Tola, and Anders Bjorholm Dahl. Large-scale data for multiple-view stereopsis. *Int. J. Comput. Vis.*, 120:153–168, 2016.
- [2] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [3] Michael Bleyer, Christoph Rhemann, and Carsten Rother. Patchmatch stereo-stereo matching with slanted support windows. In *Proc. Br. Mach. Vis. Conf.*, volume 11, pages 1–11, 2011.
- [4] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [5] Rui Chen, Songfang Han, Jing Xu, and Hao Su. Point-based multi-view stereo network. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 1538–1547, 2019.
- [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 *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, pages 2524–2534, 2020.
- [7] Ingemar J Cox, Sunita L Hingorani, Satish B Rao, and Bruce M Maggs. A maximum likelihood stereo algorithm. *Comput. Vis. Image Understanding*, 63(3):542–567, 1996.
- [8] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 764–773, 2017.
- [9] Yikang Ding, Wentao Yuan, Qingtian Zhu, Haotian Zhang, Xiangyue Liu, Yuanjiang Wang, and Xiao Liu. Transmvsnet global context-aware multi-view stereo network with transformers. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, pages 8585–8594, 2022.
- [10] Xiaoyi Dong, Jianmin Bao, Dongdong Chen, Weiming Zhang, Nenghai Yu, Lu Yuan, Dong Chen, and Baining Guo. Cswin transformer: A general vision transformer backbone with cross-shaped windows. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, pages 12124–12134, 2022.
- [11] Pascal Fua and Yvan G Leclerc. Object-centered surface reconstruction combining multi-image stereo and shading. *Int. J. Comput. Vis.*, 16(ARTICLE):35–56, 1995.
- [12] Yasutaka Furukawa, Carlos Hernández, et al. Multi-view stereo: A tutorial. *Foundations and Trends® in Computer Graphics and Vision*, 9(1-2):1–148, 2015.
- [13] Andrea Fusiello, Emanuele Trucco, and Alessandro Verri. A compact algorithm for rectification of stereo pairs. *Machine vision and applications*, 12:16–22, 2000.
- [14] Silvano Galliani, Katrin Lasinger, and Konrad Schindler. Massively parallel multiview stereopsis by surface normal diffusion. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 873–881, 2015.
- [15] 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 *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, pages 2495–2504, 2020.
- [16] Heiko Hirschmuller. Stereo processing by semiglobal matching and mutual information. *IEEE Trans. Pattern Anal. Mach. Intell.*, 30(2):328–341, 2007.
- [17] Mengqi Ji, Juergen Gall, Haitian Zheng, Yebin Liu, and Lu Fang. SurfaceNet an end-to-end 3d neural network for multi-view stereopsis. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 2307–2315, 2017.
- [18] Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are rns: Fast autoregressive transformers with linear attention. In *Proc. Int. Conf. Mach. Learn.*, pages 5156–5165. PMLR, 2020.
- [19] Arno Knapitsch, Jaesik Park, Qian-Yi Zhou, and Vladlen Koltun. Tanks and temples benchmarking large-scale scene reconstruction. *ACM Trans. Graph.*, 36(4):1–13, 2017.
- [20] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, pages 2117–2125, 2017.
- [21] Keyang Luo, Tao Guan, Lili Ju, Haipeng Huang, and Yawei Luo. P-mvsnet learning patch-wise matching confidence aggregation for multi-view stereo. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 10452–10461, 2019.
- [22] Wenjie Luo, Yujia Li, Raquel Urtasun, and Richard Zemel. Understanding the effective receptive field in deep convolutional neural networks. *Proc. Adv. Neural Inf. Process. Syst.*, 29, 2016.
- [23] Xinjun Ma, Yue Gong, Qirui Wang, Jingwei Huang, Lei Chen, and Fan Yu. Epp-mvsnet epipolar-assembling based depth prediction for multi-view stereo. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 5732–5740, 2021.
- [24] Zhenxing Mi, Chang Di, and Dan Xu. Generalized binary search network for highly-efficient multi-view stereo. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, pages 12991–13000, 2022.
- [25] Rui Peng, Rongjie Wang, Zhenyu Wang, Yawen Lai, and Ronggang Wang. Rethinking depth estimation for multi-view stereo a unified representation. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, pages 8645–8654, 2022.
- [26] Daniel Scharstein and Richard Szeliski. A taxonomy and evaluation of dense two-frame stereo correspondence algorithms. *Int. J. Comput. Vis.*, 47:7–42, 2002.
- [27] Johannes L Schonberger and Jan-Michael Frahm. Structure-from-motion revisited. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, pages 4104–4113, 2016.
- [28] Johannes L Schonberger, Enliang Zheng, Jan-Michael Frahm, and Marc Pollefeys. Pixelwise view selection for unstructured multi-view stereo. In *Proc. Eur. Conf. Comput. Vis.*, pages 501–518. Springer, 2016.
- [29] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Proc. Adv. Neural Inf. Process. Syst.*, 30, 2017.
- [30] Fangjinhua Wang, Silvano Galliani, Christoph Vogel, Pablo Speciale, and Marc Pollefeys. Patchmatchnet learned multi-view patchmatch stereo. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, pages 14194–14203, 2021.

- [31] 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 *Proc. Eur. Conf. Comput. Vis.*, pages 573–591. Springer, 2022.
- [32] Zizhuang Wei, Qingtian Zhu, Chen Min, Yisong Chen, and Guoping Wang. Aa-rmvsnet adaptive aggregation recurrent multi-view stereo network. In *Proc. IEEE Int. Conf. Comput. Vis.*, pages 6187–6196, 2021.
- [33] Jianfeng Yan, Zizhuang Wei, Hongwei Yi, Mingyu Ding, Runze Zhang, Yisong Chen, Guoping Wang, and Yu-Wing Tai. Dense hybrid recurrent multi-view stereo net with dynamic consistency checking. In *Proc. Eur. Conf. Comput. Vis.*, pages 674–689. Springer, 2020.
- [34] Jiayu Yang, Wei Mao, Jose M Alvarez, and Miaomiao Liu. Cost volume pyramid based depth inference for multi-view stereo. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, pages 4877–4886, 2020.
- [35] Yao Yao, Zixin Luo, Shiwei Li, Tian Fang, and Long Quan. Mvsnet depth inference for unstructured multi-view stereo. In *Proc. Eur. Conf. Comput. Vis.*, pages 767–783, 2018.
- [36] Yao Yao, Zixin Luo, Shiwei Li, Tianwei Shen, Tian Fang, and Long Quan. Recurrent mvsnet for high-resolution multi-view stereo depth inference. In *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, pages 5525–5534, 2019.
- [37] 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 *Proc. IEEE Conf. Comput. Vis. Pattern Recogn.*, pages 1790–1799, 2020.
- [38] Hongwei Yi, Zizhuang Wei, Mingyu Ding, Runze Zhang, Yisong Chen, Guoping Wang, and Yu-Wing Tai. Pyramid multi-view stereo net with self-adaptive view aggregation. In *Proc. Eur. Conf. Comput. Vis.*, pages 766–782. Springer, 2020.
- [39] Jingyang Zhang, Yao Yao, Shiwei Li, Zixin Luo, and Tian Fang. Visibility-aware multi-view stereo network. *arXiv preprint arXiv:2008.07928*, 2020.