

# Occlusion-aware Unsupervised Learning of Depth from 4-D Light Fields

Jing Jin, and Junhui Hou, *Senior Member, IEEE*

**Abstract**—Depth estimation is a fundamental issue in 4-D light field processing and analysis. Although recent supervised learning-based light field depth estimation methods have significantly improved the accuracy and efficiency of traditional optimization-based ones, these methods rely on the training over light field data with ground-truth depth maps which are challenging to obtain or even unavailable for real-world light field data. Besides, due to the inevitable gap (or domain difference) between real-world and synthetic data, they may suffer from serious performance degradation when generalizing the models trained with synthetic data to real-world data. By contrast, we propose an unsupervised learning-based method, which does not require ground-truth depth as supervision during training. Specifically, based on the basic knowledge of the unique geometry structure of light field data, we present an occlusion-aware strategy to improve the accuracy on occlusion areas, in which we explore the angular coherence among subsets of the light field views to estimate initial depth maps, and utilize a constrained unsupervised loss to learn their corresponding reliability for final depth prediction. Additionally, we adopt a multi-scale network with a weighted smoothness loss to handle the textureless areas. Experimental results on synthetic data show that our method can significantly shrink the performance gap between the previous unsupervised method and supervised ones, and produce depth maps with comparable accuracy to traditional methods with obviously reduced computational cost. Moreover, experiments on real-world datasets show that our method can avoid the domain shift problem presented in supervised methods, demonstrating the great potential of our method. The code will be publicly available at <https://github.com/jingjin25/LFDE-OccUnNet>.

**Index Terms**—Light field, depth estimation, occlusion, unsupervised learning, deep learning.

## I. INTRODUCTION

Depth estimation is a crucial topic in the field of computer vision and image processing, which aims at inferring the scene geometry from 2D images. Compared with monocular or stereo 2-D images, light field (LF) images intrinsically record the scene geometry by capturing both the intensity and direction of the light rays permeating the 3-D scenes [1], [2], and thus, provide more accurate cues for depth estimation. In addition, depth estimation is also a common step in 4-D LF processing and analysis, which paves the way for subsequent LF applications, such as 3-D reconstruction [3] and virtual/augment reality [4], [5].

This work was supported in part by the Hong Kong Research Grants Council under grants 9048123 (CityU 21211518) and 9042820 (CityU 11219019), and in part by the Basic Research General Program of Shenzhen Municipality under grants JCYJ20190808183003968.

J. Jin and J. Hou are with the Department of Computer Science, City University of Hong Kong, Hong Kong (e-mail: jingjin25-c@my.cityu.edu.hk; jh.hou@cityu.edu.hk).

Traditional methods for LF depth estimation mainly design different costs to explore the depth cues intrinsically provided by the LF data [6], such as detecting the lines in the epipolar plane images (EPIs) and computing their slopes [3], [7]–[10], matching corresponding pixels among different sup-aperture images (SAIs) [11]–[13], and finding optimal refocused images [14]–[16]. However, the procedure of the cost minimization and the followed global optimization always bring a heavy burden of computation. Recent learning-based methods [17]–[23] significantly improve both the efficiency and accuracy of LF depth estimation by training deep neural networks under the supervision of ground-truth depth maps. As ground-truth depth maps are usually unavailable for real-world LF images, these methods are always trained using synthetic data that are costly to obtain. Moreover, due to the inevitable domain shift between the synthetic and real-world data, the models trained with synthetic data always suffer from performance degradation when generalizing to real-world images. Unsupervised learning-based methods [24]–[26] overcome these limitations by training deep neural networks without the need of ground-truth depth maps as supervision. However, their performance is still limited.

Occlusion is a fundamental issue in LF depth estimation and has been being studied these years. Generally, different occlusion-aware LF depth estimation methods are built upon the same fact or similar observations, but differ from each other in the explicit realization and formulation of the fact/observations, resulting in various estimation accuracy and efficiency. Non-learning-based methods build cost volumes by checking photo-consistency to estimate depth maps from LFs, and then, handle occlusions based on a common fact that the photo-consistency can only hold on non-occluded pixels. Therefore, previous non-learning-based methods typically design different occlusion-aware cost volumes to improve the accuracy of depth estimation on occlusion areas [10], [11], [13], [16], which relies in time-consuming cost-volume optimization. The learning-based method in [26] also considers the occlusion issue. Following previous unsupervised methods for stereo or optical flow estimation, it uses left-right or forward-backward depth consistency for occlusion detection. However, to check the depth consistency for computing the symmetry loss, the depth maps of the eight SAIs neighboring to the central one need to be predicted simultaneously, which increases computational complexity.

In this paper, we propose an unsupervised learning-based method for LF depth estimation, which overcomes the limitations of supervised learning-based and optimization-based methods. That is, the proposed method can not only be trained

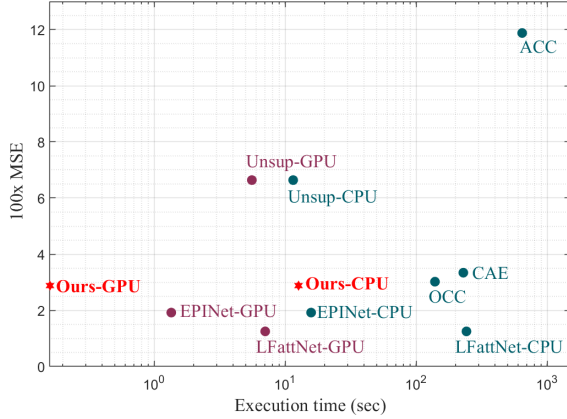


Fig. 1. Comparisons of the running time (in second) and the accuracy of the depth estimation of different methods (in MSE, Mean Square Error). For the learning-based methods, including EPINet [19], LFAFNet [21], Unsup [24], and Ours, we provide the results with and without the GPU acceleration, denoted as GPU and CPU, respectively, while for non-learning-based methods, including ACC [27], OCC [11], and CAE [16], we only evaluated them without the GPU acceleration.

without using the ground-truth depth as supervision but also infer the depth at a high speed. In contrast to supervised learning-based methods that can directly regress the depth map under the supervision of the ground-truth depth maps, it is challenging for unsupervised learning-based methods to predict the accurate depth map purely from the LF image, especially on occlusion and textureless areas. Therefore, we develop the network and loss function based on the knowledge of the unique geometry structure of LF data, which is also the general path for designing unsupervised learning methods [28], [29]. Specifically, based on the observation that the angular coherence always holds on a subset of the SAIs in an LF even for pixels with occlusions, we propose a simple yet effective strategy to handle the occlusion issue. That is, instead of processing a 4-D LF as a whole, we partition it into sub-LFs, each of which contains a certain subset of SAIs, and predict multiple depth maps from sub-LFs. Meanwhile, we design a constrained unsupervised loss to simultaneously predict their reliability maps that are utilized to produce the final prediction. Additionally, we adopt a multi-scale network together with a global smoothness loss to propagate the depth estimation to textureless areas.

Experimental results on both synthetic and real-world 4-D LF datasets demonstrate the great potential of the proposed method. Specifically, our method significantly improves the accuracy of the state-of-the-art unsupervised learning-based one [24], and compared with traditional optimization-based methods, our method achieves comparable even better performance at a much higher speed. Moreover, our method can avoid the domain shift problem presented in supervised learning-based methods and still produce satisfactory results in real-world LF data.

The rest of this paper is organized as follows. Sec. II comprehensively reviews existing methods for depth estimation from LFs. Sec. III describes the proposed method in detail. Sec. IV presents extensive experimental results to evaluate and

further analyse the proposed method. Finally, Sec. V concludes this paper.

## II. RELATED WORK

We classify the existing depth estimation methods from LFs into two categories, i.e., non-learning-based and learning-based methods.

### A. Non-learning-based Methods

Traditional non-learning-based methods typically involve two steps, i.e., local depth initialization and global optimization. These methods explore the geometric property of the LF structure to design different data costs for local depth estimation, mainly including EPI-based and angular patch-based methods.

In the EPI of an LF image, the projections of the same scene point captured by different SAIs construct a straight line, and the slope of the line reflects the depth of the scene point. Based on this observation, Bolles *et al.* [30] first proposed to detect the straight lines in EPIs and compute their slopes for depth estimation. Wanner *et al.* [7] used the structure tensor to compute the direction of the local level lines in EPIs and used the coherence of the structure tensor to measure the reliability. Kim *et al.* [3] first estimated the slope of line edges on EPIs, and then propagated the information throughout the EPI. They also used a fine-to-coarse strategy to successively process the EPI for the highest resolution to the coarser resolution. Zhang *et al.* [8] proposed a spinning parallelogram operator to obtain local depth estimation, which defines a parallelogram on the EPI and figures out the straight line by finding the maximum distance between the distributions of pixel values on either side of the lines.

Angular patch-based methods extend conventional stereo matching to narrow-baseline multi-view geometry, and construct different costs to measure the photo-consistency among different SAIs of the LF. Jeon *et al.* [27] applied the phase shift theorem to estimate the sub-pixel displacement between SAIs of the LF, and built gradient-based and feature-based matching costs to find correspondences. Some methods also introduced the refocus image as the depth cue [14], [31], [32]. The assumption of photo-consistency does not hold in the presence of occlusions, and thus, different methods have been proposed to address this issue. Wang *et al.* [11], [12] developed the LF occlusion theory, i.e., the line separating the view regions in the angular patch has the same orientation as the occlusion edge in the spatial domain, and then computed depth on half of the viewpoints that keep the photo-consistency. Zhu *et al.* [13] further extended this occlusion theory to the case of multi-occluder occlusion. Chen *et al.* [33] defined the partially occluded border regions which suffer from occlusion induced depth uncertainty, and employed superpixel-based regularization to resolve such uncertainty. Williem *et al.* [15], [16] proposed an angular entropy cost and an adaptive defocus response to improve the robustness on occlusions and noisy. Instead of utilizing partial angular patches for depth estimation, Zhang *et al.* [34] leveraged the undirected graph to jointly consider occluded and unoccluded SAIs in the angular

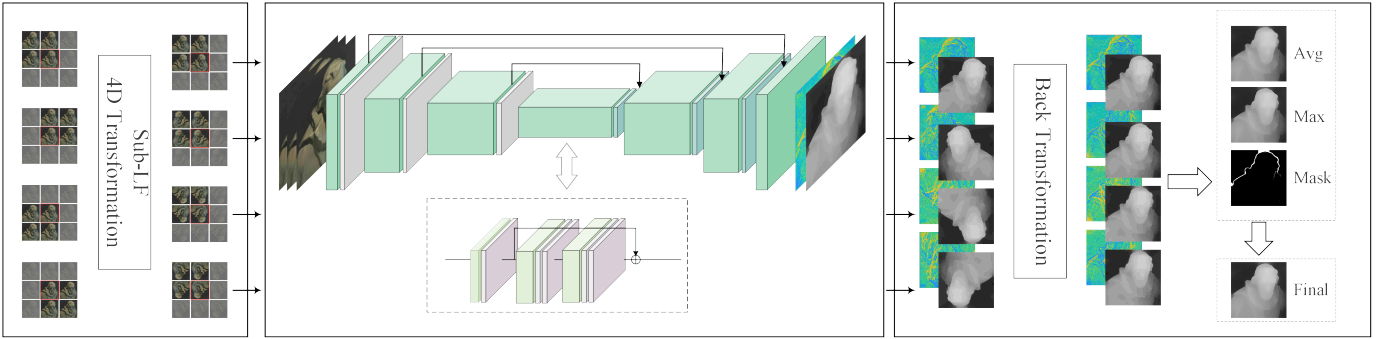


Fig. 2. Flowchart of the proposed unsupervised learning-based method for depth estimation from 4-D LFs. The sub-LFs generated from the LF image are first transformed, and then fed into a multi-scale network to produce the initial depth maps and the corresponding reliability maps simultaneously, which are further back transformed to match the central SAI, and fused to produce the final prediction of the depth map.

patch to exploit the structural information of the LF. Note that in parallel to our work, one contemporaneous work [35] has emerged recently, where the occlusion issue is addressed by proposing a novel vote cost to count the number of refocused pixels whose deviations from the central-view pixel are less than a small threshold.

### B. Learning-based Methods

Through data-driven training with the ground-truth depth as supervision, recent learning-based methods significantly improve the depth estimation accuracy while greatly saving the computational time.

Heber *et al.* [17] applied a 5-layers network to two image patches from the vertical and horizontal EPIs to predict the depth value for individual pixels. They further improved the methods by using a u-shaped network with 3-D convolutional layers to predict the depth map for the full central view through one pass forward [18]. Shin *et al.* [19] proposed a multi-stream network to extract features from SAI stacks with 4 directions, and then concatenate the features to regress the depth map. Shi *et al.* [22] proposed a learning-based framework for depth estimation, which obtains the initial depth map for stereo pairs using fine-tuned model of FlowNet2.0 optical flow estimation [36], [37], and then fuses the initial depth estimations based on the warping errors to handle occlusions. Most recently, Tsai *et al.* [21] introduced an SAI selection module to generate attention maps that indicate the importance of each SAI, and Chen *et al.* [23] also used the attention mechanism to fuse features from different SAIs. To handle the occlusion issue, Guo *et al.* [38] computed the occlusion regions based on the ground-truth depth maps, and utilized a network for occlusion detection in a supervised manner. Then they estimate the depth map in occlusion and non-occlusion regions separately.

All the above-mentioned learning-based methods have to be trained with ground-truth depth. However, the ground-truth depth maps are challenging to obtain or even unavailable for real-world LF data. Besides, they may suffer from performance degradation when generalizing the models trained with synthetic data to real-world data due to the domain difference between synthetic and real-world data. Recently, Peng *et al.* [24], [25] proposed an unsupervised learning method for depth estimation from LFs, which can be trained without

the ground-truth depth maps. However, as the occlusion and textureless issues are not explicitly addressed, the accuracy of the estimated depth maps is still limited. Apart from the similar photometric consistency-based losses, Zhou *et al.* [26] proposed a symmetry loss to handle occlusion areas. However, to check the depth consistency for computing the symmetry loss, the depth maps for the eight SAIs neighboring to the central one need to be predicted simultaneously, which increases the computational complexity.

## III. PROPOSED METHOD

### A. Overview

Let  $L(x, y, u, v) \in \mathbb{R}^{H \times W \times M \times N}$  denote a 4-D LF image with the spatial resolution of  $H \times W$  and the angular resolution of  $M \times N$ . The LF image can be regarded as a set of 2-D views observed from the viewpoints distributed on a 2-D plane, and thus, it can also be denoted by  $\mathcal{L} = \{I_{\mathbf{u}} = L(\cdot, \cdot, u, v) \in \mathbb{R}^{H \times W} | \mathbf{u} \in \mathcal{U}\}$ , where  $I_{\mathbf{u}}$  denotes the SAI at the angular position  $\mathbf{u} = (u, v)$ , and  $\mathcal{U}$  is the set of the 2-D angular coordinates, i.e.,  $\mathcal{U} = \{\mathbf{u} | \mathbf{u} = (u, v), 1 \leq u \leq M, 1 \leq v \leq N\}$ . Let  $I_{\mathbf{u}_0}$  denote the central view of  $\mathcal{L}$ , and our goal is to estimate its depth map, denoted as  $D$ . Note that as the depth is inversely proportional to the disparity, we do not make a distinction between them in this paper.

Considering the limitations of supervised learning-based methods and traditional optimization-based methods, we aim at estimating  $D$  using an unsupervised learning-based framework. With no ground-truth labels as the supervision during training, the general guidance for designing an unsupervised learning method is to take advantage of the prior knowledge of the data [28]. In our method, we exploit the intrinsic geometry structure of LF data to drive the training of the network. To be more specific, based on the observation that the photo consistency can always hold within a subset of the LF views, we proposed an occlusion-aware strategy to address the occlusion issue, i.e., we predict the initial depth maps from sub-LFs and estimate their occlusion-aware reliability using a constrained unsupervised loss. The final prediction of the depth map can be produced by fusing the initial depth maps based on their reliability maps. Additionally, we build the network in a multi-scale manner and adopt an edge-weighted smoothness loss to propagate the depth estimation to textureless areas.

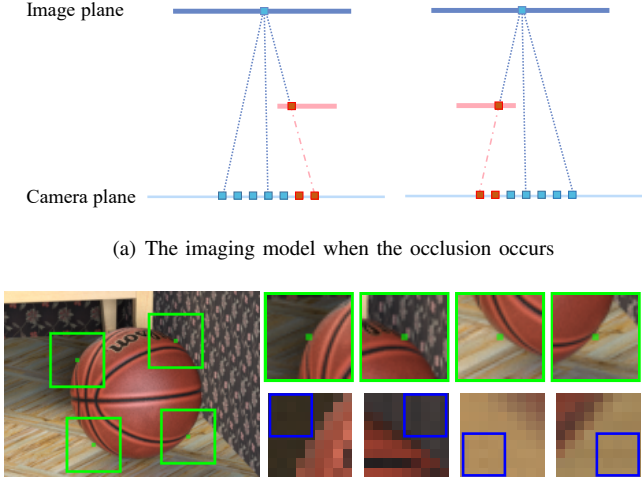


Fig. 3. Illustration of the occlusion model. In (a), the blue dots are the scene points on the image plane and their projections on the camera plane, and red dots are the occluders and their projections. In (b), the green dots are pixels with occlusions from the central view, and their surrounding areas, namely spatial patches, are framed in green and zoomed in. The refocused angular patch of each green dot with respect to its ground-truth depth value is presented under its spatial patch, where the blue frame indicates the part of the angular patch that obeys the photo consistency.

Fig. 2 illustrates the framework of our proposed method. In what follows, we will first introduce the general idea for unsupervised depth estimation and then describe our method in detail.

### B. Unsupervised LF Depth Estimation

The main cue for the depth estimation from  $\mathcal{L}$  is the angular coherence among the views, i.e., the projections of the same scene point at different views have the same intensity under the assumption of Lambertian and non-occlusion. This relation can be represented as:

$$I_{\mathbf{u}_0}(\mathbf{x}) = I_{\mathbf{u}}(\mathbf{x} + D(\mathbf{x})(\mathbf{u} - \mathbf{u}_0)), \quad (1)$$

where  $\mathbf{x}$  is the spatial coordinate. Based on Eq. (1), we can build a feature-extraction network to explore the angular coherence among the LF for depth estimation. Moreover, we can use an unsupervised loss to train the network by minimizing the photometric reconstruction distance, i.e.,

$$\ell_{rec}(\mathcal{L}, \tilde{D}) = \sum_{\mathbf{u} \subset \mathcal{U}} \sum_{\mathbf{x}} \left| \hat{I}_{\mathbf{u} \rightarrow \mathbf{u}_0}(\mathbf{x}; \tilde{D}) - I_{\mathbf{u}_0}(\mathbf{x}) \right|, \quad (2)$$

where  $\hat{I}_{\mathbf{u} \rightarrow \mathbf{u}_0}(\tilde{D})$  denotes the image warped from  $I_{\mathbf{u}}$  to  $I_{\mathbf{u}_0}$  based on the predicted depth map  $\tilde{D}$ .

However, the relationship described in Eq. (1) no longer holds when occlusions occur, and thus, the model trained with the loss function in Eq. (2) will lose accuracy on occlusion areas. To this end, we propose the following occlusion-aware strategy for unsupervised learning of the LF depth estimation.

### C. Occlusion-aware Model

We can use the refocused angular patch to represent the unique structure of LF data described in Eq. (1) in a more straightforward way. The refocused angular patch for the pixel  $\mathbf{x}$  of  $I_{\mathbf{u}_0}$  with respect to the depth value  $d$  is denoted by  $A_{\mathbf{x},d}(u,v) = L(x+d(u-u_0), y+d(v-v_0), u, v)$ .  $A_{\mathbf{x},d}$  collects the corresponding pixels of  $I_{\mathbf{u}_0}(\mathbf{x})$  from each SAIs of the LF refocused at depth  $d$ . Based on Eq. (1), pixels in  $A_{\mathbf{x},d}$  will have the same intensity when  $d$  is the correct depth value of  $I_{\mathbf{u}_0}(\mathbf{x})$ , i.e.,  $d = D(\mathbf{x})$ . However, when the occlusion occurs at  $I_{\mathbf{u}_0}(\mathbf{x})$ , the photo consistency will not hold in  $A_{\mathbf{x},D(\mathbf{x})}$ .

Fig. 3 (a) illustrates the imaging model when the occlusion occurs. Suppose only one occluder exists, the occlusion will only occur at one side of the central viewpoint along a 1-D angular dimension. By extending this observation to the 2-D angular plane of an LF image, we can deduce that one of the 4 sides of the central viewpoint can avoid the occlusion problem. As shown in Fig. 3 (b), we demonstrate 4 different occlusion scenarios (as shown in the spatial patches framed in green), and collect the refocused angular patches of the occluded pixels using their ground-truth depth values. We can observe that although the whole angular patches disobey the color consistency, if we divide  $A_{\mathbf{x},D(\mathbf{x})}$  into 4 parts, then at least one part can still keep the color consistency (as shown in the blue frames). Moreover, the positions of the color-consistent part differ with the spatial and geometric content in the spatial patches.

Based on such an unique structure of LF data, we propose a simple yet effective strategy to handle occlusions for the LF depth estimation. Specifically, we estimate depth maps from sub-LFs instead of the full LF, and predict their reliability maps simultaneously. Then, we obtain the final prediction of the depth map by fusing those from sub-LFs under the guidance of their reliability maps.

1) *Sub-LF generation*: We first divide each of the angular dimension with  $\mathbf{u}_0 = (M_0, N_0)$  as the center to get 4 sub-sets of  $\mathcal{U}$ , denoted as  $\mathcal{U} = \{\mathcal{U}_1, \mathcal{U}_2, \mathcal{U}_3, \mathcal{U}_4\}$ . Note that  $\mathbf{u}_0$  is the angular position of the depth map to be estimated. Specifically,  $\mathcal{U}_1 = \{\mathbf{u} = (u, v) | 1 \leq u \leq M_0, 1 \leq v \leq N_0\}$ ,  $\mathcal{U}_2 = \{\mathbf{u} = (u, v) | 1 \leq u \leq M_0, N_0 \leq v \leq N\}$ ,  $\mathcal{U}_3 = \{\mathbf{u} = (u, v) | M_0 \leq u \leq M, 1 \leq v \leq N_0\}$ , and  $\mathcal{U}_4 = \{\mathbf{u} = (u, v) | M_0 \leq u \leq M, N_0 \leq v \leq N\}$ . Accordingly, we divide the LF image into 4 sub-LFs, i.e.,  $\mathcal{L} = \{\mathcal{L}_1, \mathcal{L}_2, \mathcal{L}_3, \mathcal{L}_4\}$ , where  $\mathcal{L}_i = \{I_{\mathbf{u}} | \mathbf{u} \subset \mathcal{U}_i\}$ .

2) *Uncertainty-aware depth maps*: We apply a network to predict the depth map from these sub-LFs, producing  $\tilde{D}_1, \tilde{D}_2, \tilde{D}_3$ , and  $\tilde{D}_4$ . Each sub-LF contains structured views to provide sufficient cues for depth estimation, but will fail when the occlusion occurs. Fortunately, most pixels of  $I_{\mathbf{u}_0}$  have at least one sub-LF that maintains the angular coherence. To indicate which sub-LF is reliable for depth estimation, we expect the network to learn a reliability map for each prediction of the sub-LF simultaneously. To enable the training of such a network, we propose a constrained photometric






Fig. 4. Visual illustration of the uncertainty-aware depth maps estimated from the sub-LFs. (a) The depth maps estimated from sub-LFs and their error maps with respect to the ground-truth depth; (b) Zoom-in spatial patches; (c) Zoom-in error maps of the depth maps estimated from sub-LFs (the results from  $\mathcal{L}_1$ ,  $\mathcal{L}_2$ ,  $\mathcal{L}_3$ , and  $\mathcal{L}_4$  are presented from left to right.)

reconstruction loss defined as:

$$\ell_{c-rec}(\mathcal{L}, \tilde{D}, \tilde{W}) = \sum_{i=1}^4 \sum_{\mathbf{u} \in \mathcal{U}_i} \sum_{\mathbf{x}} \tilde{W}_i(\mathbf{x}) \left| \hat{I}_{\mathbf{u} \rightarrow \mathbf{u}_0}(\mathbf{x}; \tilde{D}_i) - I_{\mathbf{u}_0}(\mathbf{x}) \right|, \quad (3)$$

where  $\tilde{W}_i$  is the reliability map corresponding to  $\tilde{D}_i$ . We also apply the softmax operation across different uncertainty maps to ensure  $\sum_{i=1}^4 \tilde{W}_i(\mathbf{x}) = 1$ .

By minimizing  $\ell_{c-rec}$  instead of  $\ell_{rec}$ , the network can relax the optimization on occlusion areas of each sub-LF, because the loss can be decreased by reducing  $\tilde{W}_i(\mathbf{x})$  when  $\hat{I}_{\mathbf{u} \rightarrow \mathbf{u}_0}(\mathbf{x}; \tilde{D}_i)$  has difficulty to get close to  $I_{\mathbf{u}_0}(\mathbf{x})$ . Fig. 4 visually illustrates its effect by showing the error maps of the initial depth maps estimated from sub-LFs, where it can be observed that in each spatial patch with occlusions, some of the depth maps show high accuracy while others contain obvious errors. Moreover, the subset of the accurate depth maps changes with the spatial content and geometric relations of patches.

3) *Occlusion-aware fusion*: We generate the final prediction of the depth map by fusing the 4 initial depth maps under the guidance of their reliability maps. For pixels in occlusion areas, only the sub-LFs without any occluded viewpoints can hold the angular coherence and provide accurate depth estimation. Therefore, we take the depth value with the highest reliability as the final prediction, i.e.,

$$\tilde{D}_{max}(\mathbf{x}) = \tilde{D}_j(\mathbf{x}), \quad j = \arg \max_i \tilde{W}_i(\mathbf{x}). \quad (4)$$

Considering that the reliability map in Eq. (3) may increase the uncertainty of the non-occlusion pixels of the initial depth maps from individual sub-LFs, we take the average depth values weighted by the reliability to produce a smooth estimation, i.e.,

$$\tilde{D}_{avg}(\mathbf{x}) = \sum_i \tilde{W}_i(\mathbf{x}) \tilde{D}_i(\mathbf{x}). \quad (5)$$

The detection of the occlusion areas is based on the variance of the initial depth maps. As all the 4 sub-LFs on non-occlusion areas obey the intensity consistency when refocusing on the correct depth, they can produce relatively accurate predictions




Fig. 5. Visual illustration of the occlusion-aware fusion. The image patch of the central view, the occlusion mask, the depth maps from average and max operations, and the final prediction of the depth maps are presented. The corresponding error map is also provided under each depth map.

that are close to each other. While on occlusion areas, some of the sub-LFs are influenced by the intensity inconsistency, and others are not, and thus, the predictions show a relatively large divergence. We compute the standard deviation over the initial prediction for each pixel, and make a hard threshold to generate the occlusion mask, i.e.,

$$\tilde{O}(\mathbf{x}) = \mathcal{T}_\lambda \left( \text{STD} \left( \tilde{D}_1, \tilde{D}_2, \tilde{D}_3, \tilde{D}_4 \right) \right), \quad (6)$$

where  $\mathcal{T}_\lambda(x) = \begin{cases} 1, & x \geq \lambda \\ 0, & x < \lambda \end{cases}$  is the binarization operation with the threshold  $\lambda$ ,  $\text{STD}(\cdot)$  computes the standard deviation, and  $\tilde{O}$  is the estimated occlusion mask with 1 and 0 respectively indicating occlusion and non-occlusion. We empirically set the value of  $\lambda$  to 0.3. Finally, we can produce the final prediction of the depth map at  $I_{\mathbf{u}_0}$  by fusing the results on occlusion and non-occlusion areas, i.e.,

$$\tilde{D}(\mathbf{x}) = \tilde{O}(\mathbf{x}) \tilde{D}_{max}(\mathbf{x}) + (1 - \tilde{O}(\mathbf{x})) \tilde{D}_{avg}. \quad (7)$$

As an example, Fig. 5 visualizes  $\tilde{D}_{max}$ ,  $\tilde{D}_{avg}$ ,  $\tilde{O}$ , and  $\tilde{D}$  to illustrate the effectiveness of the occlusion-aware fusion, where we can observe that  $\tilde{D}_{max}$  performs well on occlusion boundaries, while  $\tilde{D}_{avg}$  can produce smooth results on non-occlusion areas. Based on the occlusion mask,  $\tilde{D}$  can leverage the advantages of  $\tilde{D}_{max}$  and  $\tilde{D}_{avg}$ .

#### D. Implementation Details

1) *Sub-LF 4-D transformation*: To reduce the model size, we apply a shared network to the 4 sub-LFs. As the target view, i.e., the central view of the full LF, is located at different

TABLE I

QUANTITATIVE COMPARISONS (MSE  $\times 100$ ) OF THE DEPTH ESTIMATION RESULTS FROM DIFFERENT METHODS ON SYNTHETIC LF DATA. THE SMALLER, THE BETTER.

Methods		HCI				HCIfold					
		Boxes	Cotton	Dino	Sideboard	Buddha	Horses	Medieval	Monasroom	Papillon	StillLife
Non-Learning	ACC [27]	24.91	8.70	1.25	12.64	1.21	1.74	1.05	11.02	4.88	13.07
	OCC [11]	8.14	1.04	0.59	2.31	0.76	0.70	0.85	0.55	0.75	2.49
	CAE [16]	10.18	1.01	0.62	1.55	0.77	1.20	1.14	0.68	0.85	1.45
Supervised	EPINet [19]	6.35	0.26	0.18	0.88	0.37	6.85	2.23	1.36	6.15	2.55
	LFattNet [21]	4.17	0.19	0.10	0.54	0.36	5.79	1.45	0.75	5.15	15.36
Unsupervised	Unsup [24]	12.74	7.31	1.88	4.62	1.11	1.65	1.27	1.92	4.58	33.90
	Ours	7.45	0.80	0.63	1.79	0.34	1.52	0.70	0.57	1.11	1.57

TABLE II

QUANTITATIVE COMPARISONS (BPR ( $> 0.07$ )) OF THE DEPTH ESTIMATION RESULTS FROM DIFFERENT METHODS ON SYNTHETIC LF DATA. THE SMALLER, THE BETTER.

Methods		HCI				HCIfold					
		Boxes	Cotton	Dino	Sideboard	Buddha	Horses	Medieval	Monasroom	Papillon	StillLife
Non-Learning	ACC [27]	25.39	7.36	17.57	24.40	9.33	14.44	7.92	12.07	18.14	22.32
	OCC [11]	65.68	4.57	9.83	22.92	6.74	14.36	12.74	10.54	22.73	13.06
	CAE [16]	29.48	7.84	18.17	21.47	5.69	14.82	23.80	9.56	18.59	20.41
Supervised	EPINet [19]	13.10	0.47	1.41	5.19	1.62	16.59	18.83	10.56	36.16	11.87
	LFattNet [21]	11.51	0.26	0.88	2.97	2.22	16.75	18.64	9.44	36.03	13.02
Unsupervised	Unsup [24]	43.82	28.02	22.08	28.06	9.39	19.64	18.41	14.63	28.30	44.97
	Ours	26.24	8.46	8.25	14.20	4.11	26.95	16.48	10.57	36.36	17.14

position on the angular plane of each sub-LF, we apply a 4-D transformation, which include an angular flip followed by a spatial flip, to the sub-LFs to ensure that they can produce the depth map for the same target. Specifically, we first take  $\mathcal{L}_1$  as the reference and apply an angular flip on the rest of the sub-LFs to adjust the angular position of  $I_{u_0}$ . As the angular flip will destroy the original LF structure [39], we consequently apply a spatial flip on the corresponding dimensions to recover the LF structure. Explicitly, the 4-D transformations is written as

$$\begin{aligned}
 \widehat{\mathcal{L}}_1(x, y, u, v) &= \mathcal{L}_1(x, y, u, v), \\
 \widehat{\mathcal{L}}_2(x, y, u, v) &= \mathcal{L}_2(x, W - y, u, N - v), \\
 \widehat{\mathcal{L}}_3(x, y, u, v) &= \mathcal{L}_3(H - x, y, M - u, v), \\
 \widehat{\mathcal{L}}_4(x, y, u, v) &= \mathcal{L}_4(H - x, W - y, M - u, N - v),
 \end{aligned} \tag{8}$$

where  $\widehat{\mathcal{L}}_i$  is the transformed sub-LF to be fed into the network for depth estimation. Additionally, the predicted depth maps will correspond to the flipped  $I_{u_0}$ , and thus, they are spatially flipped back to align the same target. The left and right parts of Fig. 2 illustrate the 4-D transformation applied on the sub-LFs and the back transformation applied on the predicted depth maps, respectively.

2) *Network architecture*: We use a multi-scale network [40] with shared weights to estimate depth maps from  $\widehat{\mathcal{L}}_i$  ( $i = 1, 2, 3, 4$ ). To be specific, we stack the SAIs in  $\widehat{\mathcal{L}}_i$  along the feature channel, and feed them into the network. The features extracted from the SAI stack are gradually down-sampled using max pooling layers with a kernel of size  $2 \times 2$ , and then gradually up-sampled to the original spatial size using transposed convolutional layers. The feature numbers increase

from 64, 128, 256 to 512 with the decrease of the spatial size. In each scale, we use one layer of convolution and two residual blocks [41] for feature extraction, and the features in the same scale are skip-connected to enhance the information flow. Finally, in the up-sampling branch, we apply two convolutional layers to output the depth map and its corresponding reliability map at each scale.

The multi-scale structure improves the receptive field of the network, so that it can cover the correspondences on LFs with relatively larger disparity ranges. Moreover, the multi-scale network helps to propagate the depth estimation from rich-texture to weak-texture regions.

3) *Loss function*: To improve the smoothness of the estimated depth map, as well as encourage the depth discontinuity on edges of the central view, we adopt an edge-aware smoothness loss [44], [45], written as

$$\begin{aligned}
 \ell_{sm} = \frac{1}{2} \sum_{\mathbf{x}} \exp \left( -\gamma \left| \frac{\partial I_{u_0}}{\partial x}(\mathbf{x}) \right| \right) & \left| \frac{\partial \widetilde{D}}{\partial x}(\mathbf{x}) \right| \\
 + \exp \left( -\gamma \left| \frac{\partial I_{u_0}}{\partial y}(\mathbf{x}) \right| \right) & \left| \frac{\partial \widetilde{D}}{\partial y}(\mathbf{x}) \right|,
 \end{aligned} \tag{9}$$

where the edge weight  $\gamma$  is set to 150 empirically.

The final loss function is defined as  $\ell = \ell_{c-rec} + \beta \ell_{sm}$ , where  $\beta$  controls the trade-off between the fidelity and the smoothness. We empirically set  $\beta$  to 0.1.

4) *Training details*: During training, we randomly cropped LF patches of spatial size  $128 \times 128$  and angular size  $7 \times 7$  from the training dataset. We set the batch size to 4 and initialized the learning rate as  $1e^{-4}$ . We used Adam optimizer [46] with  $\beta_1 = 0.9$  and  $\beta_2 = 0.999$ .

TABLE III  
QUANTITATIVE COMPARISONS (MSE  $\times 100$ ) OF THE DEPTH ESTIMATION RESULTS FROM DIFFERENT METHODS ON THE HCI 4D BENCHMARK.

Methods		Test				Stratified			
		Bedroom	Bicycle	Herbs	Origami	Backgammon	Dots	Pyramids	Stripes
Non-Learning	ACC [27]	0.467	11.729	21.335	6.757	13.007	5.676	0.273	17.454
	OCC [11]	0.633	7.669	22.202	2.300	21.587	3.301	0.098	8.131
	CAE [16]	0.234	5.135	11.665	1.778	6.074	5.082	0.048	3.556
Supervised	EPINet [19]	0.213	4.682	9.700	1.466	3.629	1.635	0.008	0.950
	LFattNet [21]	0.366	3.350	6.605	1.733	3.648	1.425	0.004	0.892
	AttMLFNet [23]	0.129	3.082	6.374	0.991	3.863	1.035	0.003	0.814
Unsupervised	Unsup [24]	0.924	11.737	145.551	8.817	34.709	72.998	0.035	11.759
	Mono [26]	0.415	9.232	26.816	3.679	11.833	2.536	0.027	2.677
	Ours	0.385	6.232	13.941	1.921	6.684	6.565	0.213	5.200

TABLE IV  
QUANTITATIVE COMPARISONS (BPR ( $> 0.07$ )) OF THE DEPTH ESTIMATION RESULTS FROM DIFFERENT METHODS ON THE HCI 4D BENCHMARK.

Methods		Test				Stratified			
		Bedroom	Bicycle	Herbs	Origami	Backgammon	Dots	Pyramids	Stripes
Non-Learning	ACC [27]	13.855	19.791	18.108	14.183	5.516	2.900	12.354	35.741
	OCC [11]	17.565	21.562	36.830	22.431	19.006	5.822	3.172	18.408
	CAE [16]	5.788	11.223	9.550	10.027	3.924	12.401	1.681	7.872
Supervised	EPINet [19]	2.403	9.896	12.100	5.918	3.580	3.183	0.192	2.462
	LFattNet [21]	2.792	9.511	5.219	4.824	3.126	1.432	0.195	2.933
	AttMLFNet [23]	2.074	8.837	5.426	4.406	3.228	1.606	0.174	2.932
Unsupervised	Unsup [24]	21.604	30.239	63.940	53.408	36.419	56.102	0.809	62.864
	Mono [26]	7.413	20.098	13.443	10.949	12.311	3.651	0.262	11.136
	Ours	12.687	21.650	16.959	19.821	14.371	45.340	7.348	41.359

TABLE V  
COMPARISON OF THE RUNNING TIME (IN SECONDS) OF DIFFERENT METHODS FOR INFERRING THE DEPTH MAP FROM A  $512 \times 512 \times 7 \times 7$  LF IMAGE. THE INFERENCE TIME WITH AND WITHOUT GPU ACCELERATION ARE PROVIDED.

	Non-Learning			Supervised		Unsupervised	
	ACC [27]	OCC [11]	CAE [16]	EPINet [19]	LFattNet [21]	Unsup [24]	Ours
w/o GPU	645.24	139.26	229.31	15.78	242.35	11.51	12.65
w/ GPU	-	-	-	1.35	7.04	5.57	0.16

#### IV. EXPERIMENTS

To evaluate the performance of the proposed method, we compared with state-of-the-art algorithms, including three non-learning-based methods, i.e., ACC [27], OCC [11], and CAE [16], two supervised learning-based methods, i.e., EPINet [19] and LFattNet [21], and one unsupervised learning-based method, i.e., Unsup [24].

We performed comparisons on both synthetic and real-world LF data. For the synthetic dataset, we used 4 LF images from the HCI benchmark [42] and 5 LFs from the HCI old benchmark [43] for inference. All learning-based methods were trained with the 4-D LF images from [42] which are not included in the test set. For the real-world LF data, we used 4-D LF images captured with a Lytro illum provided by Stanford Lytro LF Archive [47], Kalantari *et al.* [48], and EPFL LF dataset [49]. As the real-world datasets have no ground-truth depth for supervised training, we only retrained the models of 2 unsupervised learning-based methods using

100 LF images from [48], while directly applying the model of supervised learning-based methods trained with synthetic data. We also performed comparisons on the Stanford Gantry dataset [50], which contains real-world LF images captured with a camera gantry. For these datasets, we used the central  $7 \times 7$  SAIs of the LF for depth estimation.

Besides, we also submitted our results to the HCI online benchmark<sup>1</sup> for evaluation, which is composed of eight synthetic LF images. In addition to the aforementioned methods under comparison, we compared with two more state-of-the-art methods, i.e., one latest supervised learning-based method named AttMLFNet [23] and one unsupervised learning-based method named Mono [26]<sup>2</sup>. Note that the results of all the compared methods in terms of this online benchmark are

<sup>1</sup><https://lightfield-analysis.uni-konstanz.de/benchmark/table>

<sup>2</sup>Note that as the source codes of AttMLFNet and Mono are not publicly available when submitting this manuscript, we did not compare our method with them on other datasets.






Fig. 6. Visual comparison of the depth maps estimated by different methods on synthetic LF data from [42] and [43]. Selected regions have been zoomed in for better comparison. The error maps with respect to the ground-truth depth maps are also provided.

online available.

#### A. Evaluation on Synthetic Data

The available ground-truth depth of the synthetic data [42], [43] allows us to compare different methods quantitatively. Specifically, We computed Mean Square error (MSE) and Bad Pixel Ratio (BPR) between the estimated depth maps and the ground-truth ones to measure the accuracy, i.e.,

$$\begin{aligned}
 MSE &= \frac{1}{N} \sum_{\mathbf{x}} (\tilde{D}(\mathbf{x}) - D(\mathbf{x}))^2, \\
 BPR &= \text{Card} \left( \left\{ \mathbf{x} : \left| \tilde{D}(\mathbf{x}) - D(\mathbf{x}) \right| > t \right\} \right) / N,
 \end{aligned} \tag{10}$$

where  $\tilde{D}$  and  $D$  are the estimated and ground-truth depth maps, respectively,  $N$  is the number of pixels in a depth map, and  $\text{Card}(A)$  denotes the element number of the set  $A$ . We set  $t = 0.07$  for evaluation. Note that to avoid the padding issue in some methods, we cropped 20 pixels for the depth maps when computing the evaluation metrics.

Tables I and II list the MSE and BPR values of different methods on HCI and HCI old datasets, and Tables III and IV on the HCI online benchmark, where we can observe that

- As ground-truth depth maps are not utilized to train the model, the performance of unsupervised learning-based methods is still obviously lower than that of supervised methods on the synthetic dataset. Moreover, due to the






Fig. 7. Visual illustration of the effectiveness of the occlusion-aware strategy and the smoothness loss. *w/o occ* denotes the results produced without using the occlusion-aware strategy, and *w/o smooth* means that the weight for the smoothness loss was set to 0. The depth maps and the corresponding error maps are provided.

lack of the consideration of the occlusions, previous unsupervised learning-based method [24] performs much worse than non-learning-based methods.

- Compared with non-learning-based methods [11], [16], [27], supervised learning-based methods [19], [21], [23] can achieve superior performance on the HCI dataset, but the accuracy decreases seriously on the HCI old dataset. The reason may be that the image characteristics have great differences between these two datasets. Such a degradation also indicates that the domain difference between the training and testing data can limit the performance of the supervised learning-based methods.
- Our method produces results with the accuracy comparable to non-learning-based methods [11], [16], [27]. Compared with supervised learning-based methods [19], [21], [23], our method still has a performance gap on the HCI dataset, but can produce results with lower MSE values on the HCI old dataset.
- Our method can achieve much higher accuracy than the unsupervised learning-based method [24]. Compared with the unsupervised learning-based method [26], our method performs better in terms of MSE values and produces comparable BPR values on the *Test* set, but worse on the *Stratified* set. The possible reason is that the *stratified* set contains LF images greatly different from real-world scenes to pose specific, isolated challenges, while our model is not explicitly designed to handle them.

Besides, Fig. 6 visualizes the depth maps predicted by different algorithms as well as their error maps compared with the ground-truth ones. It can be seen that our method can well

TABLE VI  
ABLATION STUDY FOR THE PROPOSED METHOD. *w/o occ* DENOTES THE RESULTS PRODUCED WITHOUT USING THE OCCLUSION-AWARE STRATEGY, AND *w/o smooth* MEANS THAT THE WEIGHT FOR THE SMOOTHNESS LOSS WAS SET TO 0. THE QUANTITATIVE RESULTS (MSE  $\times$  100) OF THE DEPTH MAPS ARE PROVIDED.

	Boxes	Cotton	Dino	Sideboard
w/o occ	10.55	2.84	1.16	3.67
w/o smooth	11.30	1.28	0.91	2.45
Ours	7.45	0.80	0.63	1.79

handle most occlusion areas, and produce globally smooth depth maps, demonstrating the effectiveness of the proposed occlusion-aware strategy and the multi-scale network with the smoothness loss.

### B. Evaluation on Real-World Data

We also compared the accuracy of estimated depth maps by different methods on real-world 4-D LF images [47], [48] and [49]. As the ground-truth depth maps are not available, we only visually compared the results. Figs. 8, 9, and 10 present the estimated depth maps, where we can observe that although supervised learning-based methods [19], [21] can produce superior results on synthetic data, their performance is degraded on real-world LF images, especially on textureless areas. The reason might be that the camera and environmental noises exist in the real-world images. Traditional non-learning-based methods require many hyper-parameters for optimization, and thus, it is difficult for them to perform well on different scenes. Additionally, although the unsupervised learning-based method [24] was retrained on the real-world dataset, it produces depth maps with obvious inaccuracy on occlusion regions and textureless areas due to the intrinsic limitation of the model. In comparisons, our method can produce satisfactory and better depth maps on various scenes.

### C. Comparison of Efficiency

We compared the efficiency of different methods measured by the running time for inferring the depth map from a 4-D LF image, and Table V lists the results. We provided the running time with and without the GPU acceleration. All CPU-based methods were implemented on a desktop with Intel CPU i7-8700 @ 3.70GHz, 32 GB RAM, and the GPU versions were accelerated by a NVIDIA Tesla V100. From Table V, we can see that traditional non-learning methods suffer from the huge burden of computational costs, while learning-based methods, including EPINet [19], Unsup [24], and Ours, significantly save the running time even without the GPU acceleration. Particularly, our method is obviously faster than others under the acceleration of GPU.

### D. Ablation Study

We also conducted ablation studies towards the proposed occlusion-aware strategy and the smoothness loss, and Table VI and Fig. 7 provide the quantitative and qualitative results, respectively.




Fig. 8. Visual comparison of the depth maps estimated by different methods on LF data from [47] and [48].




Fig. 9. Visual comparison of the depth maps estimated by different methods on LF data from [49].

1) *The occlusion-aware strategy*: To demonstrate the advantage of the proposed occlusion-aware strategy, we developed a baseline named as *w/o occ*, which directly predicts the depth map from the full LF image instead of the sub-LFs, and trained it by replacing the constrained photometric reconstruction loss in Eq. (3) with the unconstrained one defined in Eq. (2). From Table VI, we can see that the accuracy of the predicted depth maps obviously decreases without using the occlusion-aware strategy. Fig. 7 shows that applying the occlusion-aware strategy significantly reduces the prediction error around occlusion boundaries.

2) *The smoothness loss*: To demonstrate the effectiveness of the smoothness loss, we trained a model named as *w/o smooth* by setting the weight  $\beta$  for the smoothness loss to 0. Experimental results presented in Table VI show that the accuracy of the depth estimation seriously decrease without using the smoothness loss. From Fig. 7, we can observe that *w/o smooth* has difficulties to estimate the depth in textureless areas, while Ours can produce accurate and consistent results by using the smoothness loss.

3) *The threshold  $\lambda$  for occlusion detection*: To justify our selection of the value of the threshold  $\lambda$ , we quantitatively compared the predicted depth maps when setting different






Fig. 10. Visual comparison of the depth maps estimated by different methods on LF data from [50]

TABLE VII  
QUANTITATIVE INVESTIGATION OF THE EFFECT OF THE THRESHOLD  $\lambda$  IN OCCLUSION DETECTION ON PERFORMANCE. THE BEST RESULTS ARE HIGHLIGHTED IN BOLD.

$\lambda$	Boxes	Cotton	Dino	Sideboard	Boxes	Cotton	Dino	Sideboard
	MSE $\times 100$				BPR ( $> 0.07$ )			
0.1	7.53	0.83	0.73	2.14	32.15	9.91	15.23	26.23
0.2	7.46	<b>0.80</b>	<b>0.62</b>	1.89	<b>26.08</b>	<b>8.33</b>	8.28	15.26
0.3	<b>7.45</b>	<b>0.80</b>	0.63	<b>1.79</b>	26.24	8.46	<b>8.25</b>	<b>14.20</b>
0.4	7.48	0.81	0.64	<b>1.79</b>	26.53	8.59	8.36	14.51
0.5	7.55	0.84	0.66	1.82	26.66	8.71	8.46	14.89

values of  $\lambda$ . As listed in Table VIII, we can see that setting  $\lambda$  to 0.3 produces the depth maps with highest accuracy on most cases, while setting  $\lambda$  to 0.1 or 0.5 causes obvious performance degradation.

4) *The necessity of using  $D_{avg}$* : To quantitatively demonstrate the necessity of using  $D_{avg}$ , we compared the accuracy of predicted depth maps when using  $D_{max}$  only and using both  $D_{max}$  and  $D_{avg}$ . As shown in Table VIII, it can be seen that only using  $D_{max}$  indeed reduces the accuracy of the estimated depth maps.

## V. CONCLUSION

We have presented an unsupervised learning-based method for depth estimation from 4-D LFs. By introducing an occlusion-aware strategy incorporated with a constrained unsupervised loss and utilizing a multi-scale network with the

smoothness loss, our method is capable of handling both the occlusion and textureless areas that are challenging for depth estimation, and produce satisfactory results on both synthetic and real-world LF images. Experimental results demonstrate that our method can reduce the computational burden of traditional non-learning-based methods, alleviate the performance limitation of previous unsupervised methods, and overcome the problem of domain shift on supervised methods.

## REFERENCES

- [1] M. Levoy and P. Hanrahan, "Light field rendering," in *Proceedings of the 23rd Annual Conference on Computer Graphics and Interactive Techniques*, 1996, pp. 31–42. **1**
- [2] R. Ng, M. Levoy, M. Brédif, G. Duval, M. Horowitz, and P. Hanrahan, "Light field photography with a hand-held plenoptic camera," *Computer Science Technical Report*, vol. 2, no. 11, pp. 1–11, 2005. **1**

TABLE VIII  
 QUANTITATIVE COMPARISONS (MSE  $\times 100$  AND BPR ( $> 0.07$ )) OF USING  $D_{max}$  ONLY AND USING BOTH  $D_{max}$  AND  $D_{avg}$ .

	Boxes	Cotton	Dino	Sideboard	Boxes	Cotton	Dino	Sideboard
	MSE $\times 100$				BPR ( $> 0.07$ )			
$D_{max}$ only	7.61	0.92	0.82	2.28	40.59	17.46	22.35	38.03
$D_{max}$ and $D_{avg}$ (Ours)	<b>7.45</b>	<b>0.80</b>	<b>0.63</b>	<b>1.79</b>	<b>26.24</b>	<b>8.46</b>	<b>8.25</b>	<b>14.20</b>

- [3] C. Kim, H. Zimmer, Y. Pritch, A. Sorkine-Hornung, and M. H. Gross, "Scene reconstruction from high spatio-angular resolution light fields," *ACM Transactions on Graphics*, vol. 32, no. 4, pp. 73–1, 2013. **1, 2**
- [4] F.-C. Huang, K. Chen, and G. Wetzstein, "The light field stereoscope: immersive computer graphics via factored near-eye light field displays with focus cues," *ACM Transactions on Graphics*, vol. 34, no. 4, p. 60, 2015. **1**
- [5] J. Yu, "A light-field journey to virtual reality," *IEEE MultiMedia*, vol. 24, no. 2, pp. 104–112, 2017. **1**
- [6] O. Johannsen, K. Honauer, B. Goldluecke, A. Alperovich, F. Battisti, Y. Bok, M. Brizzi, M. Carli, G. Choe, M. Diebold, M. Gutsche, H.-G. Jeon, I. S. Kweon, J. Park, J. Park, H. Schilling, H. Sheng, L. Si, M. Strecke, A. Sulc, Y.-W. Tai, Q. Wang, T.-C. Wang, S. Wanner, Z. Xiong, J. Yu, S. Zhang, and H. Zhu, "A taxonomy and evaluation of dense light field depth estimation algorithms," in *IEEE Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2017, pp. 1795–1812. **1**
- [7] S. Wanner and B. Goldluecke, "Globally consistent depth labeling of 4d light fields," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012, pp. 41–48. **1, 2**
- [8] S. Zhang, H. Sheng, C. Li, J. Zhang, and Z. Xiong, "Robust depth estimation for light field via spinning parallelogram operator," *Computer Vision and Image Understanding*, vol. 145, pp. 148–159, 2016. **1, 2**
- [9] H. Sheng, P. Zhao, S. Zhang, J. Zhang, and D. Yang, "Occlusion-aware depth estimation for light field using multi-orientation epis," *Pattern Recognition*, vol. 74, pp. 587–599, 2018. **1**
- [10] H. Schilling, M. Diebold, C. Rother, and B. Jähne, "Trust your model: Light field depth estimation with inline occlusion handling," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 4530–4538. **1**
- [11] T.-C. Wang, A. A. Efros, and R. Ramamoorthi, "Occlusion-aware depth estimation using light-field cameras," in *IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 3487–3495. **1, 2, 6, 7, 9**
- [12] —, "Depth estimation with occlusion modeling using light-field cameras," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 38, no. 11, pp. 2170–2181, 2016. **1, 2**
- [13] H. Zhu, Q. Wang, and J. Yu, "Occlusion-model guided antiocclusion depth estimation in light field," *IEEE Journal of Selected Topics in Signal Processing*, vol. 11, no. 7, pp. 965–978, 2017. **1, 2**
- [14] M. W. Tao, S. Hadap, J. Malik, and R. Ramamoorthi, "Depth from combining defocus and correspondence using light-field cameras," in *IEEE International Conference on Computer Vision (ICCV)*, 2013, pp. 673–680. **1, 2**
- [15] W. Williem and I. K. Park, "Robust light field depth estimation for noisy scene with occlusion," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 4396–4404. **1, 2**
- [16] Williem, I. K. Park, and K. M. Lee, "Robust light field depth estimation using occlusion-noise aware data costs," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 40, no. 10, pp. 2484–2497, 2018. **1, 2, 6, 7, 9**
- [17] S. Heber and T. Pock, "Convolutional networks for shape from light field," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 3746–3754. **1, 3**
- [18] S. Heber, W. Yu, and T. Pock, "Neural epi-volume networks for shape from light field," in *IEEE International Conference on Computer Vision (ICCV)*, 2017, pp. 2252–2260. **1, 3**
- [19] C. Shin, H.-G. Jeon, Y. Yoon, I. S. Kweon, and S. J. Kim, "Epinet: A fully-convolutional neural network using epipolar geometry for depth from light field images," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 4748–4757. **1, 2, 3, 6, 7, 9**
- [20] A. Alperovich, O. Johannsen, M. Strecke, and B. Goldluecke, "Light field intrinsics with a deep encoder-decoder network," in *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 9145–9154. **1**
- [21] Y.-J. Tsai, Y.-L. Liu, M. Ouhyoung, and Y.-Y. Chuang, "Attention-based view selection networks for light-field disparity estimation," in *Proceedings of the 34th Conference on Artificial Intelligence (AAAI)*, 2020, pp. 1–1. **1, 2, 3, 6, 7, 9**
- [22] J. Shi, X. Jiang, and C. Guillemot, "A framework for learning depth from a flexible subset of dense and sparse light field views," *IEEE Transactions on Image Processing*, vol. 28, no. 12, pp. 5867–5880, 2019. **1, 3**
- [23] J. Chen, S. Zhang, and Y. Lin, "Attention-based multi-level fusion network for light field depth estimation," in *Proceedings of the Conference on Artificial Intelligence (AAAI)*, vol. 35, no. 2, 2021, pp. 1009–1017. **1, 3, 7, 9**
- [24] J. Peng, Z. Xiong, D. Liu, and X. Chen, "Unsupervised depth estimation from light field using a convolutional neural network," in *International Conference on 3D Vision (3DV)*, 2018, pp. 295–303. **1, 2, 3, 6, 7, 9**
- [25] J. Peng, Z. Xiong, Y. Wang, Y. Zhang, and D. Liu, "Zero-shot depth estimation from light field using a convolutional neural network," *IEEE Transactions on Computational Imaging*, vol. 6, pp. 682–696, 2020. **1, 3**
- [26] W. Zhou, E. Zhou, G. Liu, L. Lin, and A. Lumsdaine, "Unsupervised monocular depth estimation from light field image," *IEEE Transactions on Image Processing*, vol. 29, pp. 1606–1617, 2020. **1, 3, 7, 9**
- [27] H.-G. Jeon, J. Park, G. Choe, J. Park, Y. Bok, Y.-W. Tai, and I. So Kweon, "Accurate depth map estimation from a lenslet light field camera," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 1547–1555. **2, 6, 7, 9**
- [28] S. Wu, C. Rupprecht, and A. Vedaldi, "Unsupervised learning of probably symmetric deformable 3d objects from images in the wild," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, pp. 1–10. **2, 3**
- [29] C. Guo, C. Li, J. Guo, C. C. Loy, J. Hou, S. Kwong, and R. Cong, "Zero-reference deep curve estimation for low-light image enhancement," in *IEEE Conference on Computer Vision and Pattern Recognition*, 2020, pp. 1780–1789. **2**
- [30] R. C. Bolles, H. H. Baker, and D. H. Marimont, "Epipolar-plane image analysis: An approach to determining structure from motion," *International Journal of Computer Vision*, vol. 1, no. 1, pp. 7–55, 1987. **2**
- [31] M. W. Tao, P. P. Srinivasan, J. Malik, S. Rusinkiewicz, and R. Ramamoorthi, "Depth from shading, defocus, and correspondence using light-field angular coherence," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015, pp. 1940–1948. **2**
- [32] H. Lin, C. Chen, S. B. Kang, and J. Yu, "Depth recovery from light field using focal stack symmetry," in *IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 3451–3459. **2**
- [33] J. Chen, J. Hou, Y. Ni, and L.-P. Chau, "Accurate light field depth estimation with superpixel regularization over partially occluded regions," *IEEE Transactions on Image Processing*, vol. 27, no. 10, pp. 4889–4900, 2018. **2**
- [34] Y. Zhang, W. Dai, M. Xu, J. Zou, X. Zhang, and H. Xiong, "Depth estimation from light field using graph-based structure-aware analysis," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 30, no. 11, pp. 4269–4283, 2020. **2**
- [35] K. Han, W. Xiang, E. Wang, and T. Huang, "A novel occlusion-aware vote cost for light field depth estimation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021. **3**
- [36] A. Dosovitskiy, P. Fischer, E. Ilg, P. Hausser, C. Hazirbas, V. Golkov, P. Van Der Smagt, D. Cremers, and T. Brox, "FlowNet: Learning optical flow with convolutional networks," in *IEEE International Conference on Computer Vision (ICCV)*, 2015, pp. 2758–2766. **3**
- [37] E. Ilg, N. Mayer, T. Saikia, M. Keuper, A. Dosovitskiy, and T. Brox, "FlowNet 2.0: Evolution of optical flow estimation with deep networks," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 2462–2470. **3**



- [38] C. Guo, J. Jin, J. Hou, and J. Chen, "Accurate light field depth estimation via an occlusion-aware network," in *IEEE International Conference on Multimedia and Expo (ICME)*, 2020, pp. 1–6. 3
- [39] J. Jin, J. Hou, J. Chen, S. Kwong, and J. Yu, "Light field super-resolution via attention-guided fusion of hybrid lenses," in *ACM International Conference on Multimedia (ACM MM)*, 2020, pp. 193–201. 6
- [40] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *International Conference on Medical Image Computing and Computer-assisted Intervention (MICCAI)*, 2015, pp. 234–241. 6
- [41] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 770–778. 6
- [42] K. Honauer, O. Johannsen, D. Kondermann, and B. Goldluecke, "A dataset and evaluation methodology for depth estimation on 4d light fields," in *Asian Conference on Computer Vision (ACCV)*, 2016, pp. 19–34. 7, 8
- [43] S. Wanner, S. Meister, and B. Goldluecke, "Datasets and benchmarks for densely sampled 4d light fields," in *Vision, Modelling and Visualization (VMV)*, vol. 13, 2013, pp. 225–226. 7, 8
- [44] C. Tomasi and R. Manduchi, "Bilateral filtering for gray and color images," in *IEEE International Conference on Computer Vision (ICCV)*, 1998, pp. 839–846. 6
- [45] R. Jonschkowski, A. Stone, J. T. Barron, A. Gordon, K. Konolige, and A. Angelova, "What matters in unsupervised optical flow," *ECCV*, pp. 1–1, 2020. 6
- [46] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014. 6
- [47] A. S. Raj, M. Lowney, R. Shah, and G. Wetzstein, "Stanford lytro light field archive," <http://lightfields.stanford.edu/LF2016.html>, [Online]. 7, 9, 10
- [48] N. K. Kalantari, T.-C. Wang, and R. Ramamoorthi, "Learning-based view synthesis for light field cameras," *ACM Transactions on Graphics*, vol. 35, no. 6, pp. 193:1–193:10, 2016. 7, 9, 10
- [49] M. Rerabek and T. Ebrahimi, "New light field image dataset," in *International Conference on Quality of Multimedia Experience (QoMEX)*, no. CONF, 2016. 7, 9, 10
- [50] V. Vaish and A. Adams, "The (new) stanford light field archive," <http://lightfield.stanford.edu/lfs.html>, [Online]. 7, 11