

# Facial Depth and Normal Estimation using Single Dual-Pixel Camera

Minjun Kang<sup>1†</sup>  Jaesung Choe<sup>1</sup>  Hyowon Ha<sup>4</sup>  Hae-Gon Jeon<sup>2</sup>   
Sunghoon Im<sup>3</sup>  In So Kweon<sup>1</sup>  Kuk-Jin Yoon<sup>1</sup> 

<sup>1</sup>KAIST <sup>2</sup>GIST <sup>3</sup>DGIST <sup>4</sup>Meta Reality Labs

<sup>1</sup><https://github.com/MinJunKang/DualPixelFace>

**Abstract.** Recently, Dual-Pixel (DP) sensors have been adopted in many imaging devices. However, despite their various advantages, DP sensors are used just for faster auto-focus and aesthetic image captures, and research on their usage for 3D facial understanding has been limited due to the lack of datasets and algorithmic designs that exploit parallax in DP images. It is also because the baseline of sub-aperture images is extremely narrow, and parallax exists in the defocus blur region. In this paper, we introduce a DP-oriented Depth/Normal estimation network that reconstructs the 3D facial geometry. In addition, to train the network, we collect DP facial data with more than 135K images for 101 persons captured with our multi-camera structured light systems. It contains ground-truth 3D facial models including depth map and surface normal in metric scale. Our dataset allows the proposed network to be generalized for 3D facial depth/normal estimation. The proposed network consists of two novel modules: Adaptive Sampling Module (ASM) and Adaptive Normal Module (ANM), which are specialized in handling the defocus blur in DP images. Finally, we demonstrate that the proposed method achieves state-of-the-art performances over recent DP-based depth/normal estimation methods.

**Keywords:** Dual-Pixel, Depth/Normal estimation

## 1 Introduction

A huge number of facial images are posted every day on social media. In 2020, for example, about 70 percent of photos were taken using cameras on smartphones and 24 billion selfies were uploaded to Google Photos App [19,20]. Accordingly, acquiring facial geometry from images has emerged as an interesting research topic, since 3D facial geometry can be used for various applications [69] such as face recognition [25,40,37], performance-based animation, real-time facial reenactment [5,68], facial biometrics [65], face-based interfaces, visual speech recognition, face-based search in visual assets, creating personalized avatars [43,35] or 3D printing of faces for entertainment or medicine, facial puppetry, face replacement [68,5], speech-driven animation, virtual make-up, and face image editing [63], etc. 3D facial geometry can be obtained by either using multiple cameras [15,4] or active sensing devices [32,30]. However, these methods often suffer from uncontrolled lighting conditions or hardware synchronization.

---

<sup>†</sup> Email address : kmmj2005@kaist.ac.kr

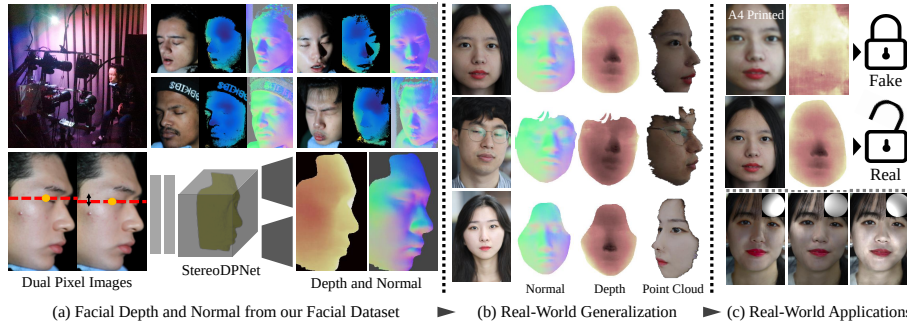
Recently, **Dual-Pixel (DP)** sensors get noticed due to their various advantages and are applied to the many portable imaging devices such as iPhone13 ProMax and Samsung Galaxy 22. DP sensors are perfectly synchronized with the same exposure, white balance, and geometric rectification. Such strengths derive from their hardware configuration that captures two images in a single camera at once. Based on these properties, currently, these novel sensors are mainly specialized in the fast auto-focus operation and more aesthetic image captures. However, thanks to their characteristics, DP sensors have great potential for other tasks. For example, a few previous studies [57,16,46,67,45,61] envision the new possibility of DP images for scene depth estimation. Usually, these studies regard DP images as extremely narrow-baseline stereo images having defocus-blur to infer high-quality depth maps. Here, it is worthy to note that although the DP sensors are being actively used to take face pictures, there has been a limited study [60] that recovers facial geometry using a Dual-Pixel camera. We found that previous methods have difficulty in facial geometry estimation, which is due to the lack of a facial DP dataset with precise 3D geometry and an appropriate algorithm for generalized estimation.

To address the issue, we present a DP-oriented 3D facial dataset and a depth/normal estimation network toward high-quality facial geometry reconstruction with DP cameras. We represent the 3D facial geometry not only with the depth map but also with the normal map for various applications such as face relighting. Our dataset involves 135,744 face data for 101 persons consisting of DP images and their corresponding depth maps and surface normal maps, which are captured by our structured light camera system. Based on these data, we train our depth/normal estimation network, called stereoDPNet, to infer 3D facial information from DP images. In particular, our stereoDPNet is fully oriented from the properties of dual-pixel images that have an extremely small range of disparity with defocus-blur. Our network design carefully treats these distinctive properties through our Adaptive Sampling Module (ASM) and Adaptive Normal Module (ANM). Finally, the contributions are as follows:

- DP-oriented 3D facial dataset with more than 135K DP images and their corresponding high-quality 3D models.
- Novel depth/normal estimation network for facial 3D reconstruction from a DP image with better generalization.

## 2 Related Work

**Defocus-disparity in Dual-Pixel.** Dual-Pixel images can be considered as a pair of stereo images since the DP camera captures two sub-aperture images with small parallax. However, since the DP camera is equipped with a micro-lens array in front of a camera sensor, the pixel disparity in DP images is extremely narrow ( $-4\text{px} \sim +4\text{px}$ ) [67] compared to conventional stereo images. For this reason, Zhang *et al.* [67] contend that a single pair of DP images is not suitable for cost volume-based disparity regression due to the narrow baseline ( $< \sim 1\text{mm}$ ). Other works [16] also adopt simple 2D U-Net architectures for affine-transformed depth regression. Meanwhile, the disparity of DP images is induced by different



**Fig. 1.** Using our Dual-Pixel Facial datasets, our network aims at generalized estimation of unmet facial geometry, which can be used for various applications, such as face spoofing or relighting.


left/right point spread functions (PSFs) instead of view parallax of stereo, called defocus-disparity [46]. Based on this observation, Punnappurath *et al.* [46] propose an optimization-based disparity regression using a parametrized PSF. The pioneering works [57, 16, 46] allow us to formulate depth-disparity conversion.

**Geometry dataset for Dual-Pixel.** Owing to the growing research interest in DP photography, several real/synthetic DP datasets [16, 46, 45, 3] have been released. Garg *et al.* [16] propose a real-world DP dataset that includes scene-scale images captured by an array of smartphones. Despite of their success, the estimated depth is up to scale because their training data contains a relative-scaled 3D geometry computed by a multiview stereo algorithm (COLMAP [51]).

**Face dataset.** Facial datasets have typically been created whenever new types of commercial imaging devices are introduced. Since it is only available to reconstruct faces from monocular images with a limited assumption [59], many 3D face regression methods [49, 14, 21] rely on a given face morphable model [56, 7] and modify the shape with facial keypoints [14, 52] and landmarks [13, 6]. Recently, several face regression models [58, 6] utilize multi-view images as input. However, it is challenging to estimate facial geometry from DP images since blurry features are hardly captured from homogeneous regions of the face. This property brings difficulty in finding correspondence between left/right DP images. We observe that the previous DP-oriented methods [16, 46] have difficulty estimating the 3D geometry of human faces as well. Many of the applications with facial images require both a high-quality depth and a surface normal for pleasing aesthetic effects [68]. Therefore, we satisfy the increasing industrial and academic demands by providing high-quality and absolute scale facial depth/normal maps that are captured with cameras with DP sensors.

### 3 Overview

This paper covers dual-pixel based facial understanding: from data acquisition (Sec. 4) to general estimation by stereoDPNet (Sec. 5). Different from natural images from typical cameras, dual-pixel sensors capture images having an extremely small range of disparity as well as defocus-blur, as shown in Fig. 4.



**Fig. 2. Examples of our facial dataset.** (a) The proposed hardware setup: (2×4 multi-camera array (blue), 6 LEDs (red), a projector (green), and a LED controller (magenta). (b) DP images with various facial expressions (horizontal axis) and distances from the hardware to faces (vertical axis). There are two additional images taken in different heading directions (center/rightward). (c) Our dataset has 68 men and 33 women and age distribution of them ranges from 19 to 45.

Through our carefully designed dataset and network, we design a well-generalized methodology that even can infer facial geometry from unmet DP facial images.

## 4 Dual-Pixel Facial Dataset


In this section, we explain the construction of DP-oriented facial dataset. Since defocus-disparity in DP images is highly sensitive to image resolution, our dataset should contain both high resolution and high quality ground truth depth. By considering these requirements, we first explain the data configuration (Sec. 4.1), details of capturing system (Sec. 4.2), and describe a ground-truth depth/normal acquisition process (Sec. 4.3).

### 4.1 Dataset Configuration

Given an array of multiple DP cameras, we capture various human faces with different expressions and light conditions. The dataset consists of 135,744 photos, which are a combination of 101 people, eight cameras, seven different lighting condition, four facial heading directions (left, right, center and upward), three facial expressions (normal, open mouth and frown), and two fixed distances of subjects from the camera array, as illustrated in Fig. 2. The distances between the camera array and subjects range from 80 *cm* to 110 *cm*. Since the focus distance is about 97 *cm*, our captured images contain both front focused and back focused cases. Our dataset includes 44,352 female photos as well as 91,392 male photos, ages range from 19 to 45. The detailed statistics are provided in Fig. 2-(c). In main experiments, we use 76 people (76%) as a train set and the others (24%) as a test/validation set without any overlap with the train set.

### 4.2 Hardware Setup

For facial data acquisition, we set up the DP-oriented camera-projector system. The system consists of eight synchronized Canon 5D Mark IV cameras on a 2×4 grid, with one commercial projector (1920 × 1080 pixels) and six LED lights, as shown in Fig. 2 which enables capturing high-quality ground truths by using both



**Fig. 3. Ground-truth depth and surface normal acquisition.** (a) Initial depth from the structured light. (b) Depth after removing outliers. (c) Depth via fusion of the initial depth and the surface normal obtained from the photometric stereo in (c).


structured light and photometric stereo. These cameras are available to capture DP images [2,3,45]. Each camera is equipped with a Canon 135mm L lens. The 17° field of view (FOV), which the lens affords, can cover an approximately 16.7cm × 25cm area at about one-meter distance, which is suitable for capturing human faces. We take our dataset with a camera aperture of F5.6, exposure time 1/30", and ISO 1600. The shape of the camera rig mimics a spherical dome with a one-meter radius. All of the cameras are located on the rig looking at the same point near the sphere's center. The projector is positioned at the center of the camera array. The LED lights are installed at various positions so that face images can be taken under varying lighting conditions.

### 4.3 Ground Truth Data Acquisition

Structured light systems are designed for high-quality 3D geometry acquisition under controlled environments by projecting pre-defined patterns on surfaces of objects [50,23,12] and by analyzing the projected patterns to measure 3D shapes of the objects. It is extensively used for ground-truth depth maps in stereo matching benchmarks [27,1,54] and shape from shading [24]. In this work, we tailor the structured light-based facial 3D reconstruction method [22] with our well-synchronized multi-camera system. Thanks to our capturing system and structured light-based reconstruction method, we obtain dense, high-quality facial 3D corresponding to high-resolution DP images in Fig. 1-(a). Moreover, we calibrate point light directions by using a chrome ball and applying a photometric stereo in [44] to obtain accurate surface normal maps of subjects' faces in Fig. 3(d). We utilize the RANSAC algorithm in obtaining both surface normal and albedo for robust estimation by excluding severe specular reflection. By using the surface normals, initial depth is refined by conforming the initial facial depth and the surface normal [44], as illustrated in Fig. 3(a), (b), and (c).

To the end, we find an exact conversion between a defocus-disparity and a metric depth by using the relationship of signed defocus-blur  $\bar{b}(x, y)$  and disparity  $d$  (Eq. 1) introduced in [16] with the paraxial and thin-lens approximations.

$$\begin{aligned}
 d(x, y) &= \alpha \bar{b}(x, y) \\
 &\approx \alpha \frac{Lf}{1 - f/g} \left( \frac{1}{g} - \frac{1}{Z(x, y)} \right) \\
 &\triangleq A(L, f, g) + \frac{B(L, f, g)}{Z(x, y)},
 \end{aligned} \tag{1}$$



**Fig. 4. Dual-Pixel Geometry.** Disparity in DP images exists in blurry regions, called defocus-disparity (left). The ground-truth depth obtained by a plane homography and a defocus-disparity from the matching pairs are used to robustly find parameters of Eq. 1. This relationship is used to back-project our prediction to metric-scale depth (right).

where  $(x, y, Z(x, y))$  indicates 3D coordinates in camera space,  $\alpha$  is a proportional term, and  $L$  is the diameter of the camera aperture.  $f$  represents the focal length of the lens and  $g$  is the focus distance of the camera. By using this relationship, we obtain a ground-truth defocus-disparity from the ground-truth depth in Fig. 4. This conversion is used to back-project our defocus-disparity to metric 3D space, and will be utilized to support absolute metric information for our depth and surface normal estimation network in Sec. 5.

## 5 Facial Depth and Normal Estimation


Based on our dataset, we design StereoDPNet for the general estimation of facial depth and normal. In real applications, dual-pixel images can be captured with various camera parameters, such as focus distance or focal length. The proposed network should cope with various hardware setups as well.

In this point of view, stereo matching methods [9,62,31,66] show strength in generalization toward unmet environments and robust to camera configuration thanks to its corresponding search. Since a pair of DP images can also be regarded as stereo images having a short disparity range, we build our StereoDPNet based on the stereo-based depth/normal method [34] for our dual-pixel based depth/normal estimation.

Nonetheless, the stereo matching methods often require all-in-focus and well-textured images for pixel-wise correspondence search, which is not always guaranteed for DP images due to its defocus-blur and homogeneous/textureless faces. Thus, we propose our two novel modules, Adaptive Sampling Module (ASM) and Adaptive Normal Module (ANM), to further fit the properties of dual-pixel images. As in Fig. 5, our StereoDPNet consists of four parts: feature extraction layer, ASM, cost aggregation layer, and ANM.

### 5.1 Overall Architecture

Given DP images with left  $I^L$  and right  $I^R$ , stereoDPNet is trained to infer a disparity map  $\hat{d}$  and a surface normal map  $\hat{n}$ . To do so, first, the feature extraction layer infers DP image features  $F^L$  and  $F^R$ , respectively. Our feature extractor captures multi-scale information with a large receptive field by adopting Atrous Spatial Pyramid Pooling [11] and Feature Pyramid Network [39] to



**Fig. 5. Architecture of StereoDPNet.** Given DP images, our network is trained to infer facial depth/normal maps. Our two key modules, Adaptive Sampling Module and Adaptive Normal Module, overcome the extremely narrow baseline in DP images by capturing disparities in blurry regions. Note that we use the pre-defined relationship between disparity and depth in Sec. 4.3 to convert disparity to metric depth.

encode various sizes of defocus blur in the DP images. Second, using  $F^L$  and  $F^R$ , the proposed ASM captures an amount of spatially varying blur in dynamic ranges, and then adaptively samples the features. Then, the sampled features  $G^L$  and  $G^R$  are stacked into a cost volume  $\mathcal{V}$ . Third, the cost volume is aggregated through three stacked hourglass modules to infer the aggregated cost volume  $C_A$ . Lastly, this aggregated volume  $C_A$  is used to regress a disparity map following the baseline and infer a surface normal map by ANM. The details of ASM and ANM are described in Sec. 5.2 and Sec. 5.3.

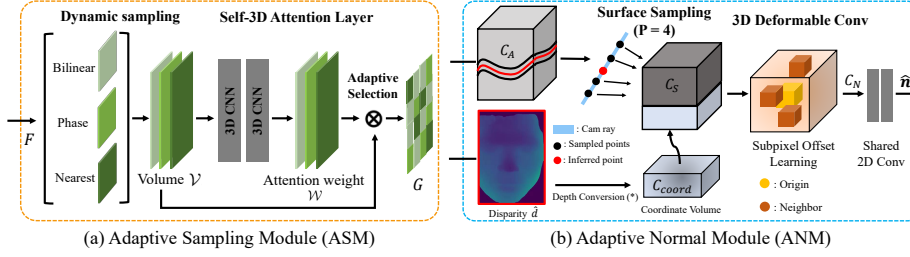
## 5.2 Adaptive Sampling Module

In contrast to the widely used stereo images such as KITTI Stereo Benchmark [18,17], dual-pixel images inherently have a small disparity range and defocus-blur. To cope with this issue, we design ASM inspired by defocus blur matching method [10] and depth from narrow-baseline light-field image [29]. As illustrated in Fig. 6-(a), our ASM dynamically samples blurry texture features for narrow-baseline stereo matching. To this end, the input features ( $F^L, F^R$ ) pass through a *dynamic feature sampling layer* and a *self-3D attention layer* to obtain the locally dominant features  $G^L, G^R$ .

According to Jeon *et al.* [29], the sub-pixel shift from different sampling strategies to construct cost volume for matching provides varying results depending on the local scene configurations. In particular, phase-shift interpolation ensures a denser sampling field at sub-pixel precision and reduces the burden from blurriness compared to other interpolation methods [29]. To take advantage of various conventional sampling methods, we incorporate them into ASM. The dynamic sampling layer in ASM is designed with a combination of nearest-neighbor, bilinear, and phase-shift interpolation, which can have various receptive fields to find varying blur sizes and can obtain subpixel-level shifted features. To this end, the shifted features from the three different sampling strategies are concatenated into one channel as a volumetric feature  $\mathcal{V}$ .

To extract useful features from given volumetric feature  $\mathcal{V}$ , we design a self-3D attention layer. Our self-3D attention layer adaptively selects sampling strategies to include prominent texture information in an extracted feature map. The layer consists of several 3D convolutional layers and the Sigmoid function. This obtains a soft mask as attention map  $\mathcal{W}$  and selects features along the channel where  $\mathcal{V}$  is





**Fig. 6. Proposed Modules of StereoDPNet.** (a) Adaptive Sampling Module (ASM) consists of a dynamic sampling and a self-3D attention layer. (b) Adaptive Normal Module (ANM) consists of a surface sampling and a 3D deformable convolutional layer for surface normal regression.

concatenated. The soft mask  $\mathcal{W}$  is multiplied with the feature  $\mathcal{V}$  to sample useful features and the final feature volume  $\mathcal{V}_S$  is produced through a softmax layer. Finally, the sampled features with the sub-pixel shift,  $G^L$  and  $G^R$ , are obtained by averaging the volume  $\mathcal{V}_S$ . The matching cost volume, constructed from the selected feature maps ( $G^L, G^R$ ), contains rich texture information with relative blur and performs effective matching in homogeneous regions as well [10].

### 5.3 Adaptive Normal Module

As illustrated in Fig. 6-(b), ANM aims to produce a surface normal map complementary to an estimated defocus-disparity map and to model 3D surface of human faces. The ANM consists of *surface sampling module* to capture surface by sampling the aggregated cost volume  $C_A$  and *deformable 3D convolutional layer* to consider dynamic ranges of neighbors to compute normal vectors.

According to [34], an accurately aggregated cost volume contains an implicit function representation of underlying surfaces for depth estimation. Since the surface normal mainly depends on the shape of the local surface, it is redundant to use all voxel embeddings in  $C_A$  for facial normal estimation. We thus sample the  $P$  candidates of hypothesis planes among  $M$  planes from the aggregated volume  $C_A$  using the estimated disparity map (Eq. 2). Since the surface normal is defined with the metric scale depth, we convert disparity to a depth map using Eq. 1 in Sec. 4.3 and provide this volumetric information with our network denoted as coordinate volume  $C_{coord}$ . The details of surface sampling process is explained in the supplementary material.

Since a human face has a variety of curved local surfaces, we need to consider dynamic ranges of neighbors to extract a local surface from the sampled hypothesis planes  $C_S$  in the previous stage. To do this, we follow the assumption of local plane in [41,47,42] and forms a local plane by a small set of neighbor points. Since these local patches have arbitrary shapes and sizes composed with its sampled neighboring points, we use 3D deformable convolutions [64] to consider the neighboring points within the dynamic ranges. The learnable offsets of the deformable convolution in 3D space allow us to adaptively sample neighbors and to find the best local plane. The final feature volume  $C_N$  is predicted after passing two 3D deformable convolution layers to extract surface normal information from the sampled volume  $C_S$ .



#### 5.4 Depth and Normal Estimation

The aggregated volume  $C_A$  passes through a classifier to produce a final matching cost  $\mathcal{A}$ , and the softmax function  $\sigma(\cdot)$  is applied to regress the defocus-disparity  $\hat{d}$ . Accordingly, we compute the disparity as follows:

$$\hat{d}_{u,v} = \sum_{m=1}^M d^m \cdot \sigma(\mathcal{A}_{u,v}^m), \quad (2)$$

where  $\hat{d}_{u,v}$  is the defocus-disparity and  $\mathcal{A}_{u,v}$  is the final matching cost at a pixel  $(u, v)$ .  $M$  and  $d^m$  are the range of defocus-disparity, and predefined discrete disparity labels, respectively, whose details are described in Sec. 5.5. Following [9], we minimize a disparity loss  $\mathcal{L}_{\text{disp}}$  using a smooth  $L_1$  loss as follows:

$$\mathcal{L}_{\text{disp}} = \frac{1}{H \cdot W} \sum_{u=1}^W \sum_{v=1}^H \mathcal{M}_{u,v} \cdot \text{smooth}_{L_1}(d_{u,v} - \hat{d}_{u,v}), \quad (3)$$

where  $d_{u,v}$  is a ground-truth defocus-disparity at a pixel  $(u, v)$  converted from the ground-truth metric scale depth and  $\mathcal{M}_{u,v}$  is the facial mask in Sec. 4.3.

For the surface normal estimation, shared 2D convolutions are applied to the feature volume  $C_N$  to regress a surface normal. The final convolutional layers follow the same structure of the baseline architecture in [34]. Finally, we train ANM by minimizing a cosine similarity normal loss  $\mathcal{L}_{\text{normal}}$  as:

$$\mathcal{L}_{\text{normal}} = \frac{1}{H \cdot W} \sum_{u=1}^W \sum_{v=1}^H \mathcal{M}_{u,v} \cdot (1 - \mathbf{n}_{u,v} \cdot \hat{\mathbf{n}}_{u,v}), \quad (4)$$

where  $\mathbf{n}_{u,v}$  and  $\hat{\mathbf{n}}_{u,v}$  are a ground-truth, and a predicted normal at a pixel  $(u, v)$ .

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{disp}} + \mathcal{L}_{\text{normal}} \quad (5)$$

Our StereoDPNet is fully supervised by our ground-truth depth/normal maps. The network is trained by minimizing the combination of Eq. 3 and Eq. 4.

#### 5.5 Implementation Details

The depth in our facial dataset ranges from 80 *cm* to 110 *cm* as described in Sec. 4.1 and the focus distance is about 97 *cm*. Therefore, our valid disparity range on original images is from -12 to 32 pixels. We note that the resolution of input images used is  $1680 \times 1120$ , which is downsampled four-fold due to GPU memory limitations. We thus set the minimum and maximum disparity of Eq. 2 to -4 and 12 pixels. We also set the number of levels in the cost volume  $M$  to 8, which represents a 0.5 pixel accuracy at least. We train our network with a batch size of four, and use Adam optimizer [33] starting from the initial learning rate  $10^{-4}$  with a constant decay of 0.5 at every 35 epochs.

Method	Task	Absolute error metric [mm] ↓					Affine error metric [px] ↓			Accuracy metric ↑	
		AbsRel	AbsDiff	SqRel	RMSE	RMSElog	WMAE	WRMSE	1-ρ	δ<1.01	δ<1.01 <sup>2</sup>
PSMNet [9]	ST	0.006	5.314	0.054	6.770	0.008	0.093	0.126	0.054	0.818	0.983
StereoNet [31]	ST	0.005	4.306	0.038	5.811	0.006	0.112	0.150	0.087	0.903	0.991
DPNet [16]	DP	0.008	7.175	0.092	8.833	0.010	0.110	0.148	0.086	0.688	0.959
MDD [46]	DP	-	-	-	-	-	1.830	2.348	0.575	-	-
BTS [36]	M	0.007	6.575	0.081	8.102	0.009	0.111	0.150	0.077	0.731	0.964
NNet [34]	DN	0.004	3.608	0.027	4.858	0.005	0.073	0.102	0.048	0.934	0.995
<b>Ours</b>	<b>DN</b>	<b>0.003</b>	<b>2.864</b>	<b>0.019</b>	<b>3.899</b>	<b>0.004</b>	<b>0.064</b>	<b>0.091</b>	<b>0.034</b>	<b>0.966</b>	<b>0.995</b>

**Table 1. Depth Benchmark Results.** We show that our proposed method outperforms the existing stereo matching methods (PSMNet [9], StereoNet [31]), DP-oriented state-of-the-art methods (DPNet [16], MDD [46]), monocular depth estimation (BTS [36]), and depth/normal network for stereo matching (NNet [34]). Note that since MDD adopts another a defocus-disparity geometry different from [16], it is not measured by the absolute metrics. ST, DP, M, and DN denotes “Stereo Matching”, “DP-oriented method”, “Monocular”, and “Depth and Normal”, respectively.

Method	ANM		Absolute [mm] ↓			Affine [px] ↓			Accuracy ↑		Normal [deg] ↓	
	SS	D3D	AbsDiff	RMSE		WMAE	WRMSE	1-ρ	δ<1.01	δ<1.01 <sup>2</sup>	MAE	RMSE
ASM Only			4.895	6.223		0.095	0.127	0.056	0.850	0.992	-	-
NNet [34]			3.608	4.858		0.073	0.102	0.048	0.934	0.995	9.634	11.877
ASM + NNet			3.271	4.434		0.064	0.090	<b>0.033</b>	0.947	<b>0.997</b>	9.072	11.045
ASM + NNet ✓			3.214	4.519		<b>0.062</b>	<b>0.089</b>	0.037	0.943	0.990	8.894	10.837
<b>StereoDPNet ✓ ✓</b>			<b>2.864</b>	<b>3.899</b>		0.064	0.091	0.034	<b>0.966</b>	0.995	<b>7.479</b>	<b>9.386</b>

**Table 2. Ablation Study of ANM.** NNet [34] is a baseline model of our overall architecture. We compare the performance of depth and surface normal estimation by adding each component. SS denotes “Surface Sampling” and D3D denotes “Deformable 3D convolution” of ANM respectively.

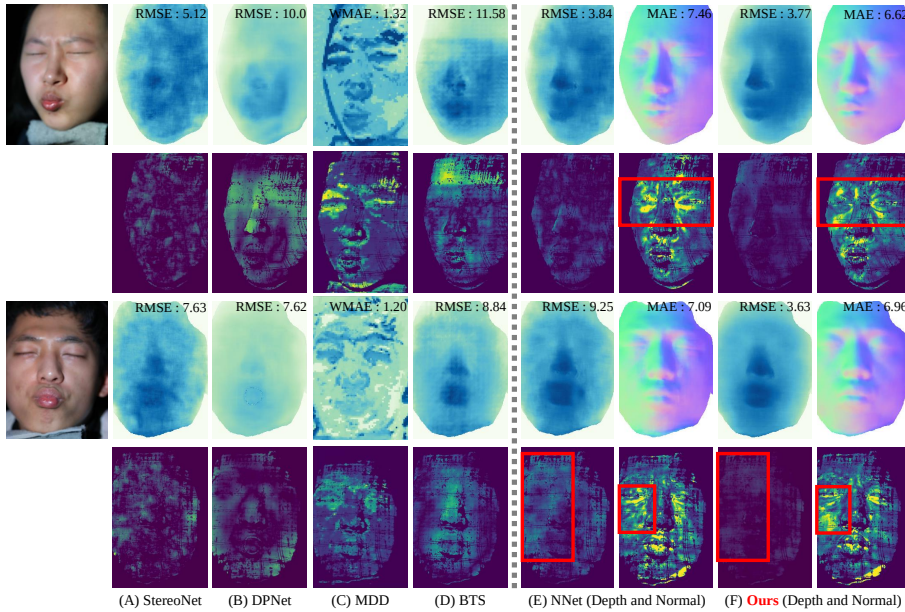
Method	Absolute error metric [mm] ↓			Accuracy metric ↑	
	AbsDiff	SqRel	RMSE	δ<1.01	δ<1.01 <sup>2</sup>
Bilinear (Bi)	7.956	0.116	9.842	0.615	0.928
Phase	8.287	0.132	10.487	0.606	0.916
Phase + Bi	6.030	0.067	7.522	0.754	0.980
Nearest + Bi	5.841	0.062	7.287	0.772	0.984
Nearest + Phase	5.831	0.062	7.247	0.773	0.985
<b>ASM</b>	<b>4.895</b>	<b>0.045</b>	<b>6.223</b>	<b>0.850</b>	<b>0.992</b>

**Table 3. Ablation Study of ASM.** We test various sampling strategies in ASM and determine the final structure of ASM. Here, we only use ASM to strictly compare the inference of each sampling strategies. Bi denotes bilinear sampling.


## 6 Experiments

To evaluate the effectiveness and the robustness of our work, we carry out various experiments on our dataset as well as DP images captured under real-world environments. For a fair comparison, all the methods are trained on identical training sets of our dataset from scratch. We then evaluate the quality of estimated depth/normal maps on the same test split of our facial dataset. Note that we use the facial mask for training and the test, given by the data acquisition process in Sec. 4.1. For the real-world samples, we use a facial mask from a pretrained face segmentation network<sup>1</sup>.

<sup>1</sup> <https://github.com/zllrunning/face-parsing.PyTorch>



**Fig. 7. Qualitative results on test set.** We report AbsRel map of depths and MAE map of normals as the error map of predictions. The error map from MDD is the WMAE map because it predicts relative scale depth maps. We note that the range of error map is  $0.0 \sim 1.0$  (AbsRel [mm]) and  $0.0 \sim 15.0$  (MAE [degree]).




**Fig. 8. Real-world results.** We capture faces in unmet real-world and compare our method with the others in Table 1. StereoDPNet clearly captures surface and boundary depth of the face. Please refer to the supplementary material for more examples.

### 6.1 Comparison Results

**Evaluation Metrics.** In our dataset benchmark, we convert our predicted disparity to depth (Sec. 4.3). Thus, we use both evaluation metrics in a public benchmark suite<sup>2</sup>: AbsRel, AbsDiff, SqRel, RMSE, RMSElog, and inlier pixel ratios ( $\delta < 1.01^i$  where  $i \in \{1, 2, 3\}$ )<sup>3</sup> and affine invariant metrics [16] for the evaluation of predicted disparity/depth. To measure the quality of a surface normal map, we utilize a Mean Angular Error (MAE) and a Root Mean Square Angular Error (RMSAE) in degree unit following the DiLiGenT benchmark [53].

<sup>2</sup> <http://www.cvlibs.net/datasets/kitti/>

<sup>3</sup> All equations of the metrics are described in Supplementary material.



**Fig. 9.** (a) Spatial behavior of ASM (bilinear, nearest, and phase). (b) We compare 3D point cloud of the method of only ASM and with our full method of ASM and ANM. The result demonstrates that ANM refines surface of the facial 3D.


**Depth Benchmark.** We compare our method with recent DP-based depth estimation approaches [16,46] as well as widely used stereo matching networks [9,31], a depth/normal network for stereo matching, NNet [34] and a state-of-the-art monocular depth estimation network, BTS [36], whose results are reported in Table 1 and in Fig. 7. Since there is no published code for [16], we implement DPNet [16] to predict disparity instead of inverse depth following them, and check that the performance is similar on their dataset.

Due to the small range of the defocus-disparity from DP images (-4px to 12px), the cost volume with the discrete hypotheses leads to unstable training [67]. As a result, typical stereo matching based depth estimations, PSMNet [9], StereoNet [31] do not work well. The methods in Table 1 except ours fail to handle defocus blur or struggle to find correspondances in human faces. For example, both DPNet and BTS suffer from blurry predictions, and MDD [46] is sensitive to the textureless regions in human faces. Although NNet [34] show relatively promising results, our method still outperforms them. Moreover, real-world results in Fig. 8 demonstrates that our network is specialized in finding defocus-disparity from facial DP images and robust to blur which produces robust results from the unmet facial scene.


**Surface Normal Benchmark.** To the best of our knowledge, this is the first attempt to estimate both the surface normal and the defocus-disparity from single DP images. Since the basic structure of ANM is derived from the recent depth and normal network [34] for multi-view stereo, we show the performance improvement of our ANM, compared to the baseline method [34] by adding each component in Table 2. We find that joint learning of disparity and surface normal leads to geometrically consistent and high-quality depth and surface normal shown in Fig. 9, which has been demonstrated in previous works [48,26].

## 6.2 Ablation Study

**Analysis of ASM and ANM.** First, we compare various subpixel sampling strategies in ASM as an ablation study. In Table 3, including whole attention maps from three different interpolations as proposed in ASM, shows the best performance over any combination of two interpolations. We also provide spatial attention of  $\mathcal{W}$  in ASM following the illustration scheme for selective matching costs in [28]. It shows that different sampling schemes are adaptively chosen. Second, we show that our surface normal estimation greatly improves the pre-



**Fig. 10. Depth from single DP images in the wild.** We show depth estimation results of StereoDPNet on outdoor photos, which are directly captured by us (left side) and in a public real-world DP dataset [2] for deblurring (right side).



**Fig. 11. Depth and normal results in the wild.** (a) public DP dataset [2], and (b) our captured natural DP images.

diction of the disparity in Table 2. ANM captures local surface and refines the surface via cost volume and leads to major performance improvement as shown in Fig. 9. We attribute this outstanding performance to our sub-modules (ASM and ANM), which overcome the extremely narrow baselines in DP images.


**Real World Experiment.** To verify the generalizability of our method, we newly capture outdoor DP images using the Canon DSLR camera. We capture faces with various camera parameters (focus distance from 1.0m to 1.5m and F-number from 2.0 to 7.1) to demonstrate our method’s robustness. Some of the results are shown in Fig. 8. Surprisingly, our StereoDPNet trained solely on our facial dataset also works well with general scenes (Fig. 10, Fig. 11), which demonstrates that our method is generalized well. Some of the scenes in Fig. 10 are from public DP dataset [2] (focus distance from 1.46m to 1.59m).

**Generalization on the public dataset [46].** We conduct an additional experiment on another real-world DP dataset [46] to validate the generalization of our network. Although our method is generalized well in real-world scenarios in Fig. 10 and in Fig. 11, we augment our network by using an additional synthetic DP dataset in [3]. This is because there is no currently available large real-world DP dataset captured with DSLR for training except ours. For a fair comparison, we don’t apply any post-processing (*i.e.* bilateral or guided filter)



Metrics	Method					
	PSMNet [9]	StereoNet [31]	DPNet [16]	MDD [46]	NNet [34]	<b>Ours</b>
WMAE ( $\downarrow$ )	0.102	0.111	0.132	0.107	0.103	<b>0.085</b>
WRMSE ( $\downarrow$ )	0.154	0.214	0.192	0.168	0.143	<b>0.133</b>
$1-\rho$ ( $\downarrow$ )	0.351	0.261	0.420	<b>0.187</b>	0.345	0.276

**Table 4. Comparisons on the public dataset [46].** We provide quantitative comparison result of the methods in Table 1 on the public dataset [46].



**Fig. 12. Generalization on the public dataset [46].** We show qualitative results of depth from the methods in Table 1 on the public dataset [46].

to the predictions. As illustrated in Fig. 12 and Table 4, our network shows promising results on the non-facial dataset [46] as well.

## 7 Conclusion

We present a high-quality facial DP dataset incorporating 135,744 face images for 101 subjects with corresponding depth maps in metric scale and surface normal maps. Moreover, we introduce DP-oriented StereoDPNet for both depth and surface normal estimation. StereoDPNet successfully shows impressive results in the wild by effectively handling the narrow baseline problem in DP.

**Potential societal impact.** We have already received consent from participants to use our facial dataset for only academic purposes. Thus, our dataset will be available to the computer vision community to promote relevant research.

**Limitation.** Although we show that our method is generalized to real-world DP scenes with various focus distances, our dataset is captured with fixed focus distance which is a clear limitation. Moreover, our ground truth acquisition of surface normal doesn't fully consider the complex specular reflection of the face which still remains as a challenging issue [35,55]. We also recognize that our dataset has an inherent bias in skin tone. However, these limitations can be resolved by capturing more dataset with various camera parameters and considering advanced non-Lambertian shape from shading methods [38,8,35]. We will refine the dataset sustainably and try to resolve these limitations.

**Acknowledgements.** This work is in part supported by the Ministry of Trade, Industry and Energy (MOTIE) and Korea Institute for Advancement of Technology (KIAT) through the International Cooperative R&D program in part (P0019797), 'Project for Science and Technology Opens the Future of the Region' program through the INNOPOLIS FOUNDATION funded by Ministry of Science and ICT (Project Number: 2022-DD-UP-0312), and also supported by the Samsung Electronics Co., Ltd (Project Number: G01210570).

## References

1. Aanaes, H., Jensen, R.R., Vogiatzis, G., Tola, E., Dahl, A.B.: Large-scale data for multiple-view stereopsis. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* pp. 1–16 (2016)
2. Abuolaim, A., Brown, M.S.: Defocus deblurring using dual-pixel data. In: *Proceedings of the European conference on computer vision (ECCV)*. pp. 111–126. Springer (2020)
3. Abuolaim, A., Delbracio, M., Kelly, D., Brown, M.S., Milanfar, P.: Learning to reduce defocus blur by realistically modeling dual-pixel data. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 2289–2298 (2021)
4. Apple: Apple iphone 11 pro. <https://www.apple.com/iphone-11-pro/> (2019), accessed: 2019-09-20
5. ARCore: Augmented faces. <https://developers.google.com/ar/develop/java/augmented-faces> (2019), accessed: 2019-12-18
6. Bai, Z., Cui, Z., Rahim, J.A., Liu, X., Tan, P.: Deep facial non-rigid multi-view stereo. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. pp. 5850–5860 (2020)
7. Blanz, V., Vetter, T.: A morphable model for the synthesis of 3d faces. In: *Proceedings of the 26th annual conference on Computer graphics and interactive techniques*. pp. 187–194 (1999)
8. Boss, M., Jampani, V., Kim, K., Lensch, H., Kautz, J.: Two-shot spatially-varying brdf and shape estimation. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 3982–3991 (2020)
9. Chang, J.R., Chen, Y.S.: Pyramid stereo matching network. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2018)
10. Chen, C.H., Zhou, H., Ahonen, T.: Blur-aware disparity estimation from defocus stereo images. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*. pp. 855–863 (2015)
11. Chen, L.C., Zhu, Y., Papandreou, G., Schroff, F., Adam, H.: Encoder-decoder with atrous separable convolution for semantic image segmentation. In: *Proceedings of the European conference on computer vision (ECCV)* (2018)
12. Chen, W., Mirdehghan, P., Fidler, S., Kutulakos, K.N.: Auto-tuning structured light by optical stochastic gradient descent. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2020)
13. Deng, Y., Yang, J., Xu, S., Chen, D., Jia, Y., Tong, X.: Accurate 3d face reconstruction with weakly-supervised learning: From single image to image set. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops* (June 2019)
14. Feng, Y., Wu, F., Shao, X., Wang, Y., Zhou, X.: Joint 3d face reconstruction and dense alignment with position map regression network. In: *Proceedings of the European conference on computer vision (ECCV)* (2018)
15. Galaxy: Samsung galaxy s10. <https://www.samsung.com/us/mobile/galaxy-s10/> (2019), accessed: 2019-03-08
16. Garg, R., Wadhwa, N., Ansari, S., Barron, J.T.: Learning single camera depth estimation using dual-pixels. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2019)



17. Geiger, A., Lenz, P., Stiller, C., Urtasun, R.: Vision meets robotics: The kitti dataset. *The International Journal of Robotics Research* **32**(11), 1231–1237 (2013)
18. Geiger, A., Lenz, P., Urtasun, R.: Are we ready for autonomous driving? the kitti vision benchmark suite. In: 2012 IEEE conference on computer vision and pattern recognition. pp. 3354–3361. IEEE (2012)
19. Google: Google photos: One year, 200 million users, and a whole lot of selfies. <https://blog.google/products/photos/google-photos-one-year-200-million/> (2016), accessed: 2016-05-27
20. Google: More controls and transparency for your selfies. <https://blog.google/outreach-initiatives/digital-wellbeing/more-controls-selfie-filters/> (2020), accessed: 2020-10-01
21. Guo, J., Zhu, X., Yang, Y., Yang, F., Lei, Z., Li, S.Z.: Towards fast, accurate and stable 3d dense face alignment. In: Proceedings of the European conference on computer vision (ECCV). pp. 152–168. Springer (2020)
22. Ha, H., Oh, T.H., Kweon, I.S.: A multi-view structured-light system for highly accurate 3d modeling. In: International Conference on 3D Vision (3DV) (2015)
23. Ha, H., Park, J., Kweon, I.S.: Dense depth and albedo from a single-shot structured light. In: International Conference on 3D Vision (3DV). pp. 127–134 (2015)
24. Han, Y., Lee, J.Y., So Kweon, I.: High quality shape from a single rgb-d image under uncalibrated natural illumination. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2013)
25. Hu, P., Ramanan, D.: Finding tiny faces. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 951–959 (2017)
26. Im, S., Ha, H., Choe, G., Jeon, H.G., Joo, K., Kweon, I.S.: High quality structure from small motion for rolling shutter cameras. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2015)
27. Jensen, R., Dahl, A., Vogiatzis, G., Tola, E., Aanæs, H.: Large scale multi-view stereopsis evaluation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2014)
28. Jeon, H.G., Park, J., Choe, G., Park, J., Bok, Y., Tai, Y.W., Kweon, I.S.: Depth from a light field image with learning-based matching costs. *IEEE transactions on pattern analysis and machine intelligence* **41**(2), 297–310 (2018)
29. Jeon, H.G., Park, J., Choe, G., Park, J., Bok, Y., Tai, Y.W., So Kweon, I.: Accurate depth map estimation from a lenslet light field camera. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2015)
30. Keselman, L., Iselin Woodfill, J., Grunnet-Jepsen, A., Bhowmik, A.: Intel realsense stereoscopic depth cameras. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops (2017)
31. Khamis, S., Fanello, S., Rhemann, C., Kowdle, A., Valentin, J., Izadi, S.: Stereonet: Guided hierarchical refinement for real-time edge-aware depth prediction. In: Proceedings of the European Conference on Computer Vision (ECCV). pp. 573–590 (2018)
32. Kinect2: Kinect for windows sdk 2.0. <https://developer.microsoft.com/en-us/windows/kinect/> (2014), accessed: 2014-10-21
33. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980 (2014)
34. Kusupati, U., Cheng, S., Chen, R., Su, H.: Normal assisted stereo depth estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)

35. Lattas, A., Moschoglou, S., Gecer, B., Ploumpis, S., Triantafyllou, V., Ghosh, A., Zafeiriou, S.: Avatarme: Realistically renderable 3d facial reconstruction” in-the-wild”. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 760–769 (2020)
36. Lee, J.H., Han, M.K., Ko, D.W., Suh, I.H.: From big to small: Multi-scale local planar guidance for monocular depth estimation. arXiv preprint arXiv:1907.10326 (2019)
37. Liang, J., Tu, H., Liu, F., Zhao, Q., Jain, A.K.: 3d face reconstruction from mugshots: Application to arbitrary view face recognition. *Neurocomputing* **410**, 12–27 (2020)
38. Lichy, D., Wu, J., Sengupta, S., Jacobs, D.W.: Shape and material capture at home. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6123–6133 (2021)
39. Lin, T.Y., Dollár, P., Girshick, R., He, K., Hariharan, B., Belongie, S.: Feature pyramid networks for object detection. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2017)
40. Liu, F., Zhao, Q., Liu, X., Zeng, D.: Joint face alignment and 3d face reconstruction with application to face recognition. *IEEE transactions on pattern analysis and machine intelligence* **42**(3), 664–678 (2018)
41. Long, X., Lin, C., Liu, L., Li, W., Theobalt, C., Yang, R., Wang, W.: Adaptive surface normal constraint for depth estimation. Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2021)
42. Long, X., Liu, L., Theobalt, C., Wang, W.: Occlusion-aware depth estimation with adaptive normal constraints. In: Proceedings of the European conference on computer vision (ECCV). pp. 640–657. Springer (2020)
43. Luo, H., Nagano, K., Kung, H.W., Xu, Q., Wang, Z., Wei, L., Hu, L., Li, H.: Normalized avatar synthesis using stylegan and perceptual refinement. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 11662–11672 (2021)
44. Nehab, D., Rusinkiewicz, S., Davis, J., Ramamoorthi, R.: Efficiently combining positions and normals for precise 3d geometry. *ACM Transactions on Graphics (ToG)* **24**(3), 536–543 (2005)
45. Pan, L., Chowdhury, S., Hartley, R., Liu, M., Zhang, H., Li, H.: Dual pixel exploration: Simultaneous depth estimation and image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 4340–4349 (June 2021)
46. Punnappurath, A., Abuolaim, A., Afifi, M., Brown, M.S.: Modeling defocus-disparity in dual-pixel sensors. In: 2020 IEEE International Conference on Computational Photography (ICCP) (2020)
47. Qi, X., Liao, R., Liu, Z., Urtasun, R., Jia, J.: Geonet: Geometric neural network for joint depth and surface normal estimation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 283–291 (2018)
48. Qiu, J., Cui, Z., Zhang, Y., Zhang, X., Liu, S., Zeng, B., Pollefeys, M.: Deeplidar: Deep surface normal guided depth prediction for outdoor scene from sparse lidar data and single color image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2019)
49. Richardson, E., Sela, M., Or-El, R., Kimmel, R.: Learning detailed face reconstruction from a single image. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1259–1268 (2017)

50. Scharstein, D., Szeliski, R.: High-accuracy stereo depth maps using structured light. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). vol. 1 (2003)
51. Schönberger, J.L., Zheng, E., Frahm, J.M., Pollefeys, M.: Pixelwise view selection for unstructured multi-view stereo. In: Proceedings of the European conference on computer vision (ECCV) (2016)
52. Shang, J., Shen, T., Li, S., Zhou, L., Zhen, M., Fang, T., Quan, L.: Self-supervised monocular 3d face reconstruction by occlusion-aware multi-view geometry consistency. In: Proceedings of the European conference on computer vision (ECCV). pp. 53–70. Springer (2020)
53. Shi, B., Wu, Z., Mo, Z., Duan, D., Yeung, S.K., Tan, P.: A benchmark dataset and evaluation for non-lambertian and uncalibrated photometric stereo. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2016)
54. Silberman, N., Fergus, R.: Indoor scene segmentation using a structured light sensor. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) - Workshop on 3D Representation and Recognition (2011)
55. Song, G., Zheng, J., Cai, J., Cham, T.J.: Recovering facial reflectance and geometry from multi-view images. *Image and Vision Computing* **96**, 103897 (2020)
56. Tran, L., Liu, X.: Nonlinear 3d face morphable model. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 7346–7355 (2018)
57. Wadhwa, N., Garg, R., Jacobs, D.E., Feldman, B.E., Kanazawa, N., Carroll, R., Movshovitz-Attias, Y., Barron, J.T., Pritch, Y., Levoy, M.: Synthetic depth-of-field with a single-camera mobile phone. *ACM Transactions on Graphics (ToG)* **37**(4), 1–13 (2018)
58. Wu, F., Bao, L., Chen, Y., Ling, Y., Song, Y., Li, S., Ngan, K.N., Liu, W.: Mvf-net: Multi-view 3d face morphable model regression. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 959–968 (2019)
59. Wu, S., Rupperecht, C., Vedaldi, A.: Unsupervised learning of probably symmetric deformable 3d objects from images in the wild. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) (2020)
60. Wu, X., Zhou, J., Liu, J., Ni, F., Fan, H.: Single-shot face anti-spoofing for dual pixel camera. *IEEE Transactions on Information Forensics and Security* **16**, 1440–1451 (2020)
61. Xin, S., Wadhwa, N., Xue, T., Barron, J.T., Srinivasan, P.P., Chen, J., Gkioulekas, I., Garg, R.: Defocus map estimation and deblurring from a single dual-pixel image. *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)* (2021)
62. Xu, H., Zhang, J.: Aanet: Adaptive aggregation network for efficient stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1959–1968 (2020)
63. Yang, F., Wang, J., Shechtman, E., Bourdev, L., Metaxas, D.: Expression flow for 3d-aware face component transfer. In: *ACM SIGGRAPH 2011 papers*, pp. 1–10 (2011)
64. Ying, X., Wang, L., Wang, Y., Sheng, W., An, W., Guo, Y.: Deformable 3d convolution for video super-resolution. *IEEE Signal Processing Letters* **27**, 1500–1504 (2020)
65. Yu, Z., Qin, Y., Li, X., Zhao, C., Lei, Z., Zhao, G.: Deep learning for face anti-spoofing: A survey. *arXiv preprint arXiv:2106.14948* (2021)

66. Zhang, F., Prisacariu, V., Yang, R., Torr, P.H.: Ga-net: Guided aggregation net for end-to-end stereo matching. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 185–194 (2019)
67. Zhang, Y., Wadhwa, N., Orts-Escolano, S., Häne, C., Fanello, S., Garg, R.: Du 2 net: Learning depth estimation from dual-cameras and dual-pixels. In: Proceedings of the European conference on computer vision (ECCV). pp. 582–598. Springer (2020)
68. Zhou, H., Hadap, S., Sunkavalli, K., Jacobs, D.W.: Deep single-image portrait re-lighting. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV) (2019)
69. Zollhöfer, M., Thies, J., Garrido, P., Bradley, D., Beeler, T., Pérez, P., Stamminger, M., Nießner, M., Theobalt, C.: State of the art on monocular 3d face reconstruction, tracking, and applications. In: Computer Graphics Forum. vol. 37, pp. 523–550. Wiley Online Library (2018)