

Dual Pixel Exploration: Simultaneous Depth Estimation and Image Restoration

Liyuan Pan^{1,2,*}, Shah Chowdhury^{1,3,*}, Richard Hartley¹, Miaomiao Liu¹,
Hongguang Zhang^{1,4}, and Hongdong Li¹

¹ Australian National University ² A&F, CSIRO ³ Data61, CSIRO ⁴ SE Institute, AMS

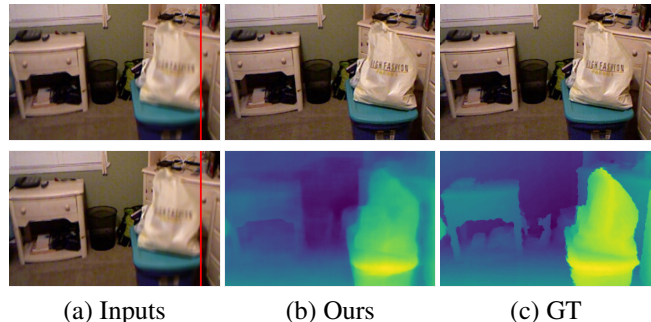
{liyuan.pan, Shah.Chowdhury}@anu.edu.au

Abstract

The dual-pixel (DP) hardware works by splitting each pixel in half and creating an image pair in a single snapshot. Several works estimate depth/inverse depth by treating the DP pair as a stereo pair. However, dual-pixel disparity only occurs in image regions with the defocus blur. The heavy defocus blur in DP pairs affects the performance of matching-based depth estimation approaches. Instead of removing the blur effect blindly, we study the formation of the DP pair which links the blur and the depth information. In this paper, we propose a mathematical DP model which can benefit depth estimation by the blur. These explorations motivate us to propose an end-to-end **DDNet** (DP-based Depth and Deblur Network) to jointly estimate the depth and restore the image. Moreover, we define a re-blur loss, which reflects the relationship of the DP image formation process with depth information, to regularise our depth estimate in training. To meet the requirement of a large amount of data for learning, we propose the first DP image simulator which allows us to create datasets with DP pairs from any existing RGBD dataset. As a side contribution, we collect a real dataset for further research. Extensive experimental evaluation on both synthetic and real datasets shows that our approach achieves competitive performance compared to state-of-the-art approaches.

1. Introduction

The Dual-Pixel (DP) sensor has been used by DSLR (digital single-lens reflex camera) and smartphone cameras to aid focusing. Though the DP sensor is designed for auto-focus [32, 16, 15], it is used in applications such as, depth estimation [10, 29, 43], defocus deblurring [1], reflection removal [30], and shallow Depth-of-Field (DoF) images synthesis [36]. In this paper, we model the imaging process of the DP sensor theoretically, and show its effectiveness for simultaneous depth estimation and image deblurring.



(a) Inputs (b) Ours (c) GT
Figure 1. The pipeline of our approach. (a) The input of our **DDNet**. A red line in each DP image is drawn to indicate a small shift between the left and right sub-aperture views. (b) Our de-blurred image and estimated inverse depth map. (c) The Ground-Truth (GT) sharp image and inverse depth map. The input DP pair is from a real-world DP sensor, or generated by our simulator.

A DP camera simultaneously captures two images, one formed from light rays passing through the right half of the aperture, and one from those passing through the left half of the aperture. Because of the displacement of these two half-apertures, the two images form a stereo pair. The primary role of DP cameras is to enable auto-focus. The set of points in the world that are exactly in focus forms a plane parallel to the lens (perpendicular to the axis of the camera), and points that lie on this plane are both in focus, and imaged with zero-disparity in the two images. Points that lie further away from the camera are imaged with positive disparity in the two images and those that are closer have negative disparity. By detecting and modelling these disparities, it is possible to compute scene depth, as in standard stereo imaging. (In auto-focus these disparities are used to move the sensor, or the lens, to focus on any part of the scene.)

At the same time, points other than those lying on the in-focus plane will be significantly out of focus, which raises the possibility of using depth-from-defocus techniques to estimate depth in the scene. This out-of-focus effect is more significant than with ordinary stereo pairs, making stereo matching potentially more challenging for DP cameras. However, because of the small base-line, between the

*Equal contribution

two half-apertures, occlusion effects are minimal.

Thus, computing depth from a DP camera can be seen as an out-of-focus stereo estimation problem. This paper aims to combine depth-from-disparity and depth-from-defocus approaches in a single network, demonstrating the advantage of modelling them both simultaneously.

To accomplish this, we study the imaging process of a DP pair and provide a mathematical model handling the intrinsic problem of *depth from defocus blur* for a DP pair. Our method can jointly recover an all-in-focus image and estimate the depth of a scene (see Fig. 1). Our contributions are summarised as follows:

1. We propose a theoretical DP model to explicitly define the relationship between depth, defocus blur and all-in-focus image;
2. We design an end-to-end DP-based Depth and Deblur Network (**DDNet**) to jointly estimate the depth map and restore the sharp image;
3. We formulate a reblur loss based on our DP model which is used to regularise depth estimate in training;
4. We create a DP image simulator, which enables us to create DP datasets from any RGB-D dataset;
5. We collect a real dataset to stimulate further research.

Extensive experimental evaluation on both synthetic and real data shows that our approach achieves competitive performance compared to state-of-the-art approaches.

2. Related Work

For space reasons, we briefly review depth estimation from monocular, stereo, DP, and defocus.

Monocular. Supervised monocular depth estimation methods typically rely on large training datasets [22, 9]. Self-supervised methods [31, 34, 20, 12] often use estimated depths and camera poses to generate synthesized images, serving as supervision signals.

Stereo. To summarize, a stereo depth estimation framework [13, 41, 25, 6, 27, 40, 2] often contains: matching cost computation, cost aggregation, and disparity optimization. Recently, several speed-up methods [37, 39] are available.

Dual-pixel. A DP pair has defocus blur and tiny baseline. To tackle the tiny baseline, Wadhwa *et al.* [36] utilize a small search size in block matching. Garg *et al.* [10] present the first learning models for affine invariant depth estimation to tackle affine ambiguity. Zhang *et al.* [43] present a Du²Net that uses a wide baseline stereo camera to compensate for the DP camera. However, it is specifically tailored for the ‘Google Pixel’ (narrow aperture, small DoF). The above three methods ignored the blur cue that could in principle fill in the parts where the disparity is imprecise (due to the tiny baseline). To tackle the defocus blur in a DP pair, Abuolaim *et al.* [1] first propose a DPDNet that remove the blur effect blindly. Recently, Punnappurath *et al.* [29] use

a point spread function (PSF) to model the defocus blur and formulate the DP disparity by the PSF. However, the symmetry assumption of the PSF (only holds for constant depth region) limits its application for real-world scenario. Moreover, the method is time-consuming and has three steps (not end-to-end). In contrast, we study the imaging process of a DP pair and give a mathematical model jointly handling the intrinsic problem of *depth from defocus blur* for a DP pair.

Depth from defocus and deblurring. Depth from defocus has the same geometric constraints as disparity but different physiological constraints [8]. The estimated depth, also dubbed as defocus map, is commonly used to guide the deblurring [3, 33, 38, 26, 23, 44, 21]. DMENet [19] proposes the first end-to-end CNN architecture to estimate a spatially varying defocus map and use it for deblurring.

DP data. Only Canon and Google provide DP data to customers though most DSLR and smartphone cameras have the DP sensor. Several researchers [36, 10, 43] use ‘Google Pixel’ to collect data. However, smartphone cameras use a fixed and narrow aperture that cannot be adjusted. Canon is used in [1, 30, 29] with different aperture sizes, but it is expensive. Both of the two sensors are hard to get the associated GT information, such as depth. In parallel, the rise of deep learning has led to an explosion of demands for data. In this paper, we present the first DP simulator that synthesizes DP images with GT, from any RGB-D data.

In summary, both the stereo cue and defocus cue provided by a DP pair benefit the depth estimation. In this paper, we proposed a DP model in Section 3 and a DDNet in Section 4 to jointly estimate the depth and restore the image.

3. Dual-Pixel Image Formation

In this section, we first discuss the formation of the two DP images and the defocus blur in Section 3.1. Then, in Section 3.2, we present our DP model and explore how to synthesize DP images from the RGB-D image. Based on the DP model, we build a DP simulator in Section 3.3.

3.1. Model of a dual-pixel camera

A DP camera can be modelled as an ordinary camera with a lens satisfying the thin-lens model, in which two (or more) images \mathbf{I}_L and \mathbf{I}_R are simultaneously captured. Here, the subscripts $_L$ and $_R$ denote the left and right view separately. In the model, the focal plane of the camera can be considered as consisting of two identically placed focal planes, one of which captures light rays coming from the left side of the lens, and the other captures light from the right side of the lens. The two images, \mathbf{I}_L and \mathbf{I}_R , can be considered as images taken by two different (but coplanar) lens, corresponding to regions in the aperture of a thin-lens.

To be general, let it be assumed that the light for image \mathbf{I}_L passes through a region \mathbf{A}_L of the lens, and that for image \mathbf{I}_R comes from a region \mathbf{A}_R . Let us consider the case




Figure 2. Formation of the two image I_L and I_R from two points C_L and C_R on the lens plane may be thought of as pin-hole projections (with centres C_L and C_R and focal-length F) of the virtual image formed by the lens.

where A_L and A_R are infinitesimally small, consisting of just two points C_L and C_R (see Fig. 2). Choose a coordinate frame such that the lens lies on the plane $X = 0$, and the centre of the lens is at the origin. In addition, let the “world” lie to the left of the lens (that is with points $X < 0$), and the camera lie to the right of the lens. Let W consisting of points $\mathbf{X} = (X, Y, Z)$ ¹ with $X < -f$, and let W' be the 3D image of the virtual world, created by the lens, lying to the right of the lens. In particular, if $\mathbf{X}' = (X', Y', Z') \in W'$, then $X' > f$. Here, f is the focal length of the lens. Consider a plane of sensors (called a focal plane) lying inside the camera. Assume that this is parallel to the lens, so it is a plane defined by the equation $X = F$, with $F > 0$. The symbol of F is chosen to imply that this is the distance of the focal plane from the lens, but it is not the same as the focal length of the lens, which is f .

Observation. An image I_L of the world W formed from rays passing through a point C_L in the focal plane is the same as a pin-hole camera image of the virtual world W' with projection centre C_L and the same focal plane. The same statement is true for the image I_R formed from rays passing through C_R .

Proof. Refer to Fig. 2. Consider a world point \mathbf{X} , mapped by the lens to a virtual world point \mathbf{X}' . This means that if \mathbf{C} is any point on the lens, then the ray \mathbf{XC} is refracted by the lens to the ray \mathbf{CX}' . In particular, this is true for the point C_L lying on the lens. Let the point \mathbf{Y} be the point where the ray $C_L\mathbf{X}'$ (extended if necessary) meets the focal plane. Then \mathbf{Y} is the point where \mathbf{X} is imaged in I_L . This is because a ray traverses the path $\mathbf{XC}_L\mathbf{Y}$, passing through \mathbf{X}' on the way. On the other hand, looked at in terms of the pinhole camera with centre C_L , the ray from C_L to \mathbf{X}' meets the focal plane at \mathbf{Y} . This means that \mathbf{Y} is the image of the point \mathbf{X}' in the pinhole image taken from C_L . \square

Therefore, if the two regions were single points C_L and C_R , then the two images I_L and I_R are exactly a stereo pair of images of the virtual world W' . Although the two points C_L and C_R are necessarily very close together, the virtual world is also very small and very close to the camera lens, so there will be appreciable disparity.

¹All vectors are column vectors.




Figure 3. Rays from a point \mathbf{X} in the scene at depth d pass through the aperture A_L and are focused at a point \mathbf{X}' at depth d' in the virtual world W' . The set of all the rays passing through A_L form, after refraction, a double-sided cone, with vertex at \mathbf{X}' . This cone, meets the focal plane at depth F in a region that is geometrically similar to the shape of A_L . This explanation is continued in Fig. 4.

Non-pinhole model. In the case where the two regions A_L and A_R are bigger than a pinhole, and in fact constitute half the lens itself, the images I_L and I_R formed will be made up as the superposition of images of the virtual world W' taken at all points C_L in region A_L and C_R in region A_R . They will consequently be blurred. This is shown in Fig. 3.

Therefore, the problem of finding depth in the scene from images I_L and I_R is equivalent to doing stereo from blurred images. This will have its problems. The blur will be depth dependent. Points in the virtual world that lie on the focal plane π will be in focus at independent of the position of the point C_L and C_R . Thus, points in the world corresponding to this placement of the focal plane will be both in focus and identically positioned in the two images I_L and I_R . Points that lie off the focal plane will be blurred and at the same time displaced by a disparity.

3.2. Image synthesis from RGB-D image

Given an image with an associated depth map, it is possible to synthesize either of the pair of dual-pixel images. Consider the image I , formed from all rays passing through a region A in the plane of the lens. (Note, A is either A_L or A_R , and the corresponding image is either I_L or I_R .)

Let I_W be an RGB-D image of the world, taken nominally from the viewpoint of the lens centre. Since the lens is normally very small with respect to the world, it will be assumed that any world point visible from the lens centre will be visible from any other point on the lens.

The RGB-D image gives us 3D coordinates of any point visible in the image. Let $\mathbf{X} = (X, Y, Z) \in W$ be such a point. The image of this point in the virtual world W' is given by

$$\mathbf{X}' = (X', Y', Z') = \frac{f}{f + X} (X, Y, Z). \quad (1)$$

This point is then projected from each point $\mathbf{C} \in A$ onto the focal plane $X = F$. The projection rays, through \mathbf{X}' from points in A form a cone with cross-section A and vertex \mathbf{X}' . This cone meets the focal plane in a further cross-section A' . Since the lens and focal plane are parallel, A and A' will be similar regions, in the sense that there is a




Figure 4. Viewed side-on, rays from a point \mathbf{X} pass through the aperture \mathbf{A}_L (similarly \mathbf{A}_R) and are focused at \mathbf{X}' . They meet the focal plane at depth F (which may be either side of \mathbf{X}') covering a region \mathbf{A}'_L geometrically similar to \mathbf{A}_L . The scale of \mathbf{A}'_L is computed by similar triangles; the scale factor is given by $(d' - F)/d'$. This scale factor is positive if F lies in front of \mathbf{X}' and negative if it lies behind; in this latter case, \mathbf{A}'_L is inverted with respect to \mathbf{A}_L .

similarity transform relating \mathbf{A} and \mathbf{A}' . Even more, this is simply a scaling and translation (where the scale may be positive or negative). In particular, a point $\mathbf{C} \in \mathbf{A}$ maps to a point $\mathbf{C}' = s\mathbf{C} + \mathbf{t}$. Here, \mathbf{t} is a 2-dimensional offset and s is a scale. For a given point \mathbf{X} , the values of s and \mathbf{t} are constant, not dependent on the particular point \mathbf{C} chosen, but they vary according to the point \mathbf{X} chosen.

Refer to Fig. 4. In particular, by similar triangles [14]

$$\mathbf{C}' = T(\mathbf{C}) = (1 - s)\mathbf{C} + s\mathbf{X}', \quad (2)$$

where $s = F/d'$.

Image coordinates. We start with an RGB-D image, so each pixel in an image comes with an associated depth. For simplicity, we assume that all distances (including depth) are measured in pixel coordinates.

We make the following assumptions. A pixel with coordinates ² (y, z) corresponds with a ray in space defined by $-d(1, y/f, z/f)$ for varying d . In particular, since each point in the image comes with a depth, we are assuming that the point imaged at point (y, z) has 3D coordinates $(-d, -dy/f, -dz/f)$. This corresponds to a pinhole camera with focal plane given by $X = f$ and camera centre (pinhole) at the origin.

A point $\mathbf{X} = -d(1, y/f, z/f)$ at depth d will be mapped to a point $\mathbf{X}' = d'(1, y/f, z/f)$, where $1/d + 1/d' = 1/f$, or

$$d' = fd/(d - f). \quad (3)$$

Let $\mathbf{C} = (0, Y_0, Z_0)$ be a point in \mathbf{A} (lying on the lens plane $X = 0$). The line from \mathbf{C} through $\mathbf{X}' = d'(1, y/f, z/f)$ is expressed as $(1 - s)\mathbf{C} + s\mathbf{X}'$ for varying values of s . This line meets the plane $X = F$ when $s = F/d'$. The coordinates of this point are therefore

$$\begin{aligned} T(\mathbf{X}, \mathbf{C}) &= \frac{d' - F}{d'}\mathbf{C} + \frac{F}{d'}\mathbf{X}' \\ &= \frac{d' - F}{d'}\mathbf{C} + F(1, y/f, z/f). \end{aligned} \quad (4)$$

²It is convenient to use (y, z) for image coordinates, instead of the usual (x, y) .

We are only interested in the (Y, Z) -coordinates, in which case we have

$$T(\mathbf{X}, \mathbf{C}) = \frac{d' - F}{d'}(Y_0, Z_0) + F(y/f, z/f). \quad (5)$$

Here, the transformation function $T(\mathbf{X}, \mathbf{C})$ for the left and right image has been defined with points \mathbf{C}_L and \mathbf{C}_R in \mathbf{A}_L and \mathbf{A}_R respectively.

3.3. DP simulator

Given a sharp image $\mathbf{I} \in \mathbb{R}^{H \times W}$ and its associated depth map $\mathbf{d} \in \mathbb{R}^{H \times W}$, we can simulate a DP image pair $\mathbf{B}_{\{L, R\}}$ as follows. Here, H and W are the image height and width, and it is also assumed that focal length f and focal-plane-sensor distance F are known (measured in units of pixel-size).

For each pixel (y, z) in \mathbf{I} , as shown in Fig. 4 the intensity of a pixel is spread over regions \mathbf{R}_L and \mathbf{R}_R , in the left and right view of the DP pair. Each region contains a set of points \mathbf{p} , containing $|\mathbf{R}|$ pixels, and the intensity $\mathbf{I}(y, z)$ of the pixel (y, z) is spread out evenly over this set of pixels. Now running over all pixels $\mathbf{I}(y, z)$, and summing, a pair of dual-pixel images are created. This is essentially blurring with depth-dependent blur kernel.

This operation is expensive computationally, since it requires looping over each pixel (y, z) in the image, as well as each pixel in the region \mathbf{R} , thus 4 levels of looping.

We utilize the concept of ‘integral image’ to speed up the computation, so that it has complexity independent of the size of the region \mathbf{R} , thus of order $\mathcal{O}(n)$ where n is the number of pixels. In this description, we assume that the left and right halves of the aperture are approximated by rectangles, though with some amount of extra computation, more general shapes may be accommodated.

Given a pixel (y, z) of the RGB-D image at a given depth, its corresponding light rays will pass through every point of \mathbf{A}_L . To compute the blur-region \mathbf{R}_L , one needs only to compute the destinations of the rays passing through the four corners of \mathbf{A}_L , which will be denoted by $\mathbf{p}_{tl} = (y_{tl}, z_{tl})$, $\mathbf{p}_{tr} = (y_{tr}, z_{tr})$, $\mathbf{p}_{bl} = (y_{bl}, z_{bl})$ and $\mathbf{p}_{br} = (y_{br}, z_{br})$ in the left image \mathbf{B}_L . The positions of these four points are given by Eq. (5). Here, the subscript tl , tr , bl and br denotes top-left, top-right, bottom-left and bottom-right, respectively.

We create a *differential mask* $\mathcal{I}_L \in \mathbb{R}^{H \times W}$ for the region \mathbf{R}_L defined by

$$\begin{aligned} \mathcal{I}_L(\mathbf{p}_{tl}) &= \mathbf{I}(y, z)/|\mathbf{R}_L|, & \mathcal{I}_L(\mathbf{p}_{tr}) &= -\mathbf{I}(y, z)/|\mathbf{R}_L|, \\ \mathcal{I}_L(\mathbf{p}_{bl}) &= -\mathbf{I}(y, z)/|\mathbf{R}_L|, & \mathcal{I}_L(\mathbf{p}_{br}) &= \mathbf{I}(y, z)/|\mathbf{R}_L|. \end{aligned} \quad (6)$$

These values are summed for the regions \mathbf{R}_L corresponding to all points (y, z) in the image, to create the *differential image*. Finally, we integrate the differential image, and the

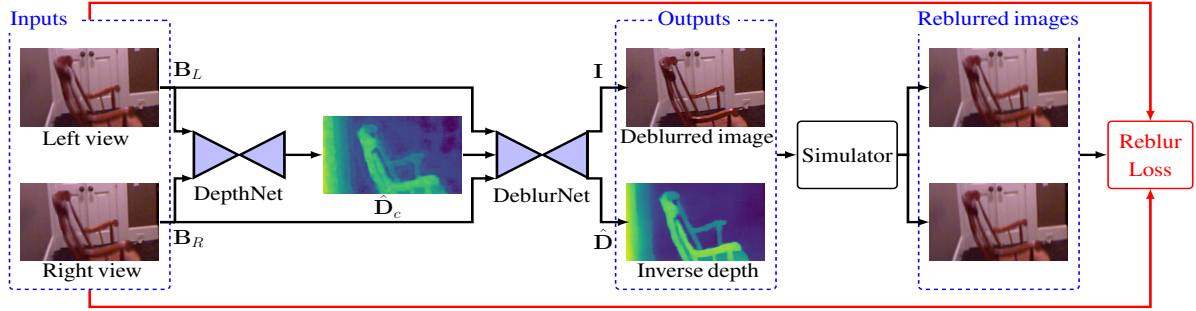


Figure 5. The architecture of the proposed DDDNet. With the input DP image pair, the DepthNet first estimates a coarse inverse depth map \hat{D}_c . Then, we combine the coarse and blurred inverse depth map with $\mathbf{B}_{\{L,R\}}$, and feed them to the DeblurNet. The deblurred image and estimated inverse depth map are fed to our DP simulator to synthesize a DP image pair. The synthesized DP images are compared to the inputs using our reblur loss (Eq. (11)), to regularise the DDDNet in training. Note that, the ground truth inverse depth maps and sharp images are used as supervision signals in training. (Best viewed in colour on screen.)

left/right view of the DP image pair is given by

$$\hat{\mathbf{B}}_{\{L,R\}} = \mathcal{T}(\mathcal{I}_{\{L,R\}}), \quad (7)$$

where $\mathcal{T}(\cdot)$ denotes the integral process. The integration process [35] is also closely related to ‘summed area tables’ in graphics [5].

Our simulator allows the vision community to collect large amounts of DP data with ground truth, opening the door to accurately benchmark DP-based methods. Our simulator can also be used to supervise the learning process.

With the DP model and the simulator, we build our DDDNet considering both stereo and defocus cues.

4. DP-based Depth and Deblur Network

4.1. Network architecture.

The input of our DDDNet is the left and right view of a DP image pair $\mathbf{B}_{\{L,R\}}$. The output of our DDDNet is the estimated inverse depth map $\hat{\mathbf{D}}$ and the deblurred image $\hat{\mathbf{I}}$. We use ground-truth latent sharp image \mathbf{I} and inverse depth map \mathbf{D} for training.

The pipeline of DDDNet is shown in Fig. 5. It consists of two components: DepthNet $g(\cdot)$ with parameters $\vec{\mathcal{G}}$ and DeblurNet $f(\cdot)$ with parameters $\vec{\mathcal{F}}$. The DepthNet is based on [4] and the DeblurNet is based on the multi-patch network [42]. Note that our approach is independent to the choices of $f(\cdot)$ and $g(\cdot)$ (e.g., multi-scale and multi-patch architectures). We start with a coarse inverse depth map \hat{D}_c estimation by the DepthNet, where $\hat{D}_c = g(\mathbf{B}_{\{L,R\}}; \vec{\mathcal{G}})$. Then, we combine the coarse and blurred inverse depth map with $\mathbf{B}_{\{L,R\}}$, and feed it to the DeblurNet. The DeblurNet is an encoder-decoder network, and $\{\hat{\mathbf{I}}, \hat{\mathbf{D}}\} = f(\mathbf{B}_{\{L,R\}}, \hat{D}_c; \vec{\mathcal{F}})$.

4.2. Loss functions.

We use a combination of an image restoration loss, depth loss, and an image reblur loss. The final loss is a sum of the three losses.

$$\mathcal{L} = \mathcal{L}_{\text{res}} + \mathcal{L}_d + \mathcal{L}_{\text{reb}}. \quad (8)$$

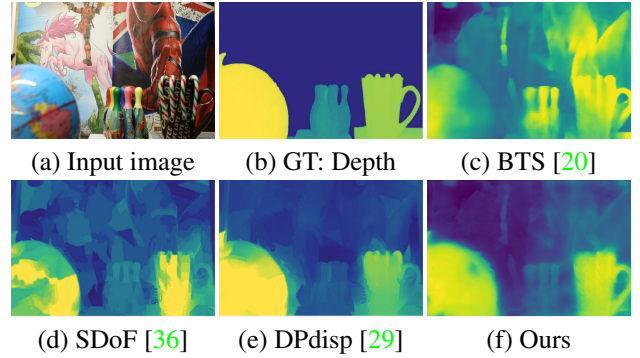


Figure 6. An example of inverse depth estimation result on the real DPD-disp dataset [29]. (Best viewed in colour on screen.)

Image restoration loss \mathcal{L}_{res} . This loss ensures the deblurred image is similar to the target image, and is given by

$$\mathcal{L}_{\text{res}} = \frac{1}{N} \sum_{y,z} \|\mathbf{I}(y,z) - \hat{\mathbf{I}}(y,z)\|, \quad (9)$$

where N is the number of pixels and $\|\cdot\|$ is the ℓ_2 norm.

Depth loss \mathcal{L}_d . We adopt the smooth ℓ_1 loss. Smooth ℓ_1 loss $\mathcal{S}(\cdot)$ is widely used in a regression task, for its robustness and low sensitivity to outliers [11]. The \mathcal{L}_d is defined as

$$\mathcal{L}_d = \frac{1}{N} \sum_{y,z} \mathcal{S}(\mathbf{D}(y,z) - \hat{\mathbf{D}}(y,z)). \quad (10)$$

Image reblur loss \mathcal{L}_{reb} . The reblur loss penalizes the differences between the input blurred images and the reblurred images (from the DP simulator). The reblur loss explicitly enforces the restored image and inverse depth map to lie on the manifold of the ground-truth image and inverse depth map, subject to our DP model. The \mathcal{L}_{reb} is given by

$$\mathcal{L}_{\text{reb}} = \frac{1}{N} \sum_{y,z} \|\mathbf{B}_{\{L,R\}}(y,z) - \hat{\mathbf{B}}_{\{L,R\}}(y,z)\|. \quad (11)$$

Table 1. *Quantitative analysis of deblurring results on the DPD-blur dataset [1], Our-syn dataset, and Our-real dataset. The best results are shown in bold. Here, ‘Ours_{wb}’ and ‘Ours_{reb}’ denotes our model trained without and with the reblur loss, respectively. The results demonstrate that our model improves significantly with the reblur loss. Besides ‘PSNR’ (in dB), to explicitly illustrate the improved quality of deblurred images, we use a metric ‘RMSE_{rel}’ (RMSE/255, in %) to further explain ‘what 1dB improvement means’. Compared with the second-best method, the relative improvement of our method in intensity value for the three datasets is 13%, 19%, and 12%, respectively.*

DPD-blur				
	EBDB [17]	DMENet [19]	DPDNet [1]	Ours _{wb} /Ours _{reb}
PSNR \uparrow	24.82	23.93	25.53	26.15/26.76
SSIM \uparrow	0.801	0.812	0.826	0.827/0.842
MSE _{rel} \downarrow	5.74	6.36	5.29	4.93/4.59
Our-syn				
PSNR \uparrow	26.48	30.14	31.45	32.17/33.21
SSIM \uparrow	0.891	0.939	0.926	0.948/0.956
MSE _{rel} \downarrow	4.74	3.11	2.68	2.46/2.17
Our-real				
PSNR \uparrow	21.67	23.18	22.65	23.99/24.03
SSIM \uparrow	0.763	0.809	0.808	0.826/0.850
MSE _{rel} \downarrow	8.25	6.39	7.09	6.24/6.13

Table 2. *Quantitative analysis of depth estimation on the DPD-disp dataset [29]. As the DPD-disp dataset only provides testing data, we directly use our model (trained on Our-syn dataset, namely ‘Ours’) without fine-tuning and achieves the second best. Then, we use our reblur loss to fine-tuning our model (without using the GT depth map). Our model with the reblur loss ‘Ours_{ft}’ (fine-tune) achieves competitive results, indicating that our simulated DP outputs resemble those from a real DP sensor.*

	AI(1) \downarrow	AI(2) \downarrow	1- ρ_s \downarrow	Geometric Mean \downarrow
BTS [20]	0.1070	0.1767	0.6149	0.2686
Monodepth2 [12]	0.1139	0.1788	0.6153	0.2285
SDoF [36]	0.0875	0.1294	0.2910	0.1443
DPdisp [29]	0.0481	0.0845	0.1037	0.0671
Ours	0.0906	0.1291	0.2456	0.1207
‘Ours _{ft} ’	0.0609	0.0985	0.1026	0.1098

5. Experiment

5.1. Experimental setup

Real dataset. We evaluate our method on three real DP datasets, namely, the Defocus Depth estimation dataset (DPD-disp) [29], Defocus Deblur Dual-Pixel dataset (DPD-blur) [1], and our new-collected dataset (Our-real). DPD-disp dataset provides DP images with depth maps. The ground-truth (GT) depth map is computed by applying the well-established depth-from-defocus technique. DPD-blur provides a collection of image pairs captured by Canon 5D IV with all-in-focus and out-of-focus. In testing, we down-sampled the frames to 50% of their original size (1120 \times 1680 pixels). Compared with the DPD-blur dataset, Our-real is collected by Canon with a variety of aperture sizes, varying from $f/4$ to $f/22$. Each all-in-focus image is as-

sociated with several out-of-focus blurred images, yielding a diversity dataset. Our-real contains 150 scenes, including both indoor and outdoor scenes, captured under a variety of lighting conditions.

Synthetic dataset. Our synthetic dataset is generated using the proposed DP simulator. We use the NYU depth dataset [24] as the input to our simulator, as it provides RGB images with depths. Giving different camera parameters, we simulate 5000 image pairs for training and 500 image pairs for testing.

Evaluation metrics. Given a DP-pair, our method jointly estimates a depth map and a deblurred image. We use standard metrics to evaluate the quality of estimated depth map and restored image separately [7, 28]. For depth map, we use absolute relative error ‘Abs.Rel’, square relative error ‘Sq.Rel’, root mean square error ‘RMSE’ and its log scale ‘RMSE.log’, and the δ inlier ratios (maximal mean relative error of $\delta_i = 1.25^i$ for $i \in 1, 2, 3$). For restored image, we adopt the peak signal-to-noise ratio ‘PSNR’, structural similarity ‘SSIM’, and RMSE relative error ‘RMSE_{rel}’. For the DPD-disp dataset, we follow DPdisp [29] to use the affine invariant version of MAE (‘AI(1)’), RMSE (‘AI(w)’), and Spearman’s rank correlation (‘1 - $|\rho_s|$ ’) for evaluation.

Baseline methods. For depth map, we compare with state-of-the-art monocular (BTS [20], Monodepth2 [12]), stereo (AnyNet [37]), and DP (SDoF [36], DPdisp [29]) based depth estimation methods. For defocus deblurring, we compared with EBDB [17], DMENet [19], and DPDNet [1]. All methods are evaluated on each dataset independently. All learning methods are fine-tuned for each dataset (except DPD-disp, as no training data is provided). To fine-tune DMENet [19] on the DPD-blur and Our-real dataset, we use BTS [20] to estimate a coarse depth map to train, as the two datasets do not provide ground-truth depth maps. The same coarse depth maps are also used to train our network.

Implementation details. Our network is trained from scratch using the Adam optimizer [18] with a learning rate of 10^{-4} and a batch size of 10. Our model is trained on a single NVIDIA Titan XP GPU.

5.2. Results and discussions

We compare our results with baselines on depth estimation and image deblurring on 4 (including real and synthetic) datasets.

Deblurring results. We show the quantitative and qualitative comparisons for deblurring in Table 1, Fig. 7 and Fig. 8, respectively. In Table 1, we achieve competitive results compared with the state-of-the-art methods [17, 19, 1] on both DPD-blur, Our-syn, and Our-real dataset. Here, ‘Ours_{wb}’ and ‘Ours_{reb}’ denote results without and with the reblur loss, respectively. The relative improvement is around 10%, demonstrating that using our DP-model based reblur loss can significantly improve the deblurring perfor-

Table 3. Quantitative analysis of deblurring and depth estimation results on our synthetic dataset. The synthetic dataset is generated using our DP simulator. Here, we evaluate the intermediate outputs of DDDNet (see Fig. 5) step by step, ‘Ours_b’ denotes the coarse depth map \hat{D}_c from DepthNet, ‘Ours_{wb}’ denotes the depth map \hat{D} from DeblurNet, and ‘Ours_{reb}’ denotes the final results with the reblur loss. BTS achieves lower in ‘RMSE’ but higher in ‘RMSE_{log}’, and the reason is that BTS gets a number of zeros in their results.

Sensor	Method	abs_rel ↓	sq_rel ↓	rmse ↓	rmse_log ↓	$a < 1.25 \uparrow$	$a < 1.25^2 \uparrow$	$a < 1.25^3 \uparrow$	PSNR ↑
Monocular	BTS [20]	0.377	0.195	0.255	0.484	0.554	0.741	0.994	-
Stereo	AnyNet [37]	0.128	0.117	0.731	0.168	0.828	0.923	0.998	-
	DPdisp [29]	0.328	0.479	1.252	0.332	0.438	0.804	0.965	-
Dual-pixel	Ours _b	0.149	0.364	1.222	0.224	0.743	0.930	0.978	-
	Ours _{wb}	0.091	0.082	0.599	0.123	0.918	0.993	0.998	32.171
	Ours _{reb}	0.083	0.052	0.461	0.111	0.936	0.998	1.000	33.218

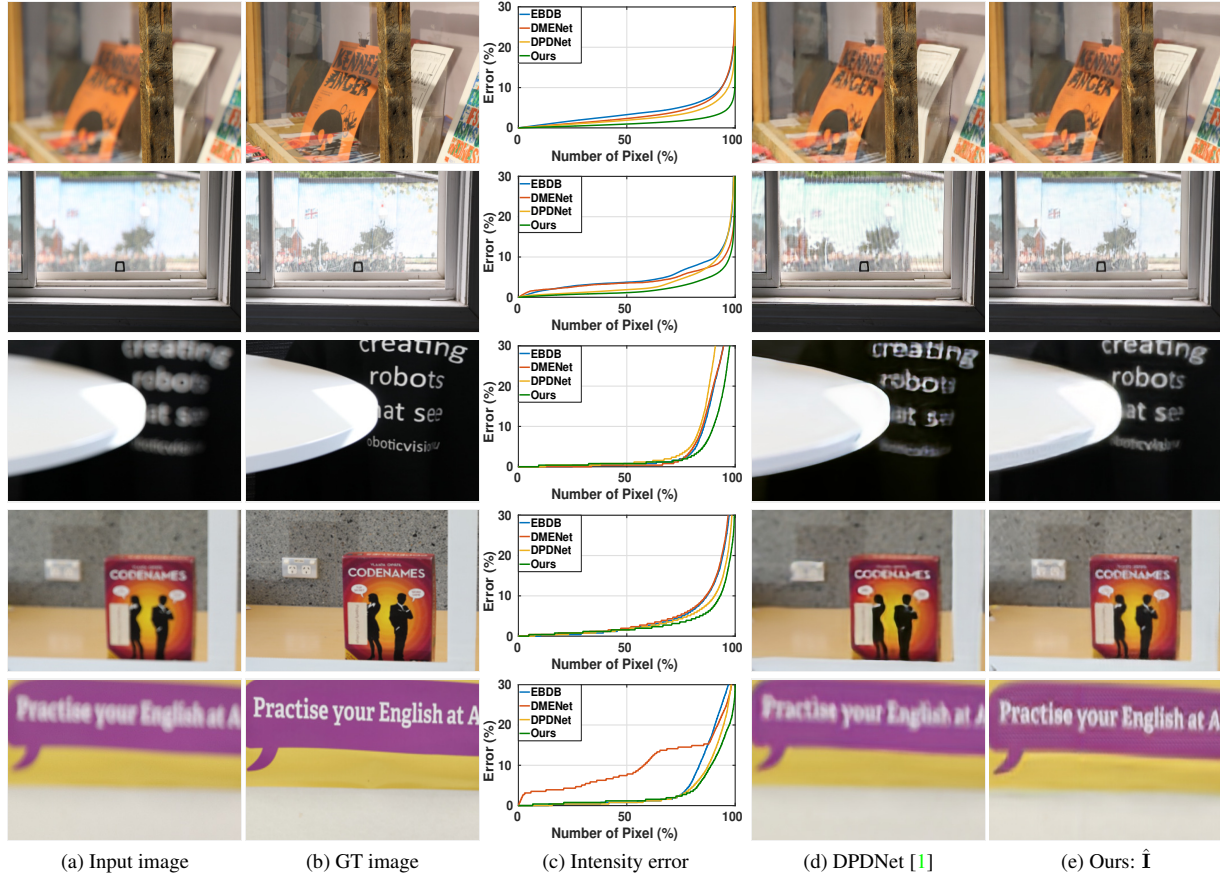


Figure 7. Examples of deblurring results on the real DPD-blur dataset [1] (top two rows) and our collected real DP dataset (bottom three rows). (c) The distribution of the intensity error (‘RMSE_{rel}’). It shows the percentile of the number of pixels below an error ratio. The lower the curve, the better. For space limitations, we only display deblurred images from the second-best baseline [29] for comparison. (Best viewed in colour on screen.)

mance of our DDDNet. Fig. 7 and Fig. 8 show the qualitative and quantitative comparisons on the three datasets.

Depth results. We first provide depth estimation results on the DPD-disp dataset. The comparisons are shown in Table 2 and Fig. 6. As no training data is available, we directly test our model (trained on Our-syn dataset) on their testing set without fine-tuning and achieve the second best performance. Furthermore, note that our DP-model based reblur loss can be used as a ‘self-supervision’ signal to adapt our DDDNet to this test dataset, we can further fine-tune

our model trained on Our-syn dataset using our reblur loss, without using GT depth map. Our model with reblur loss ‘Ours (fine-tune)’ achieves competitive results compared with the state-of-the-art methods [20, 12, 36, 29].

We then provide depth comparisons on Our-syn dataset. The results are shown in Table 3 and Fig. 8. In Table 3, we evaluate the outputs of our DDDNet step by step, it shows that, with our joint optimization of depth estimation and deblurring, we improve the quality of both the depth map and the deblurred image. Also, with our DP-model based reblur

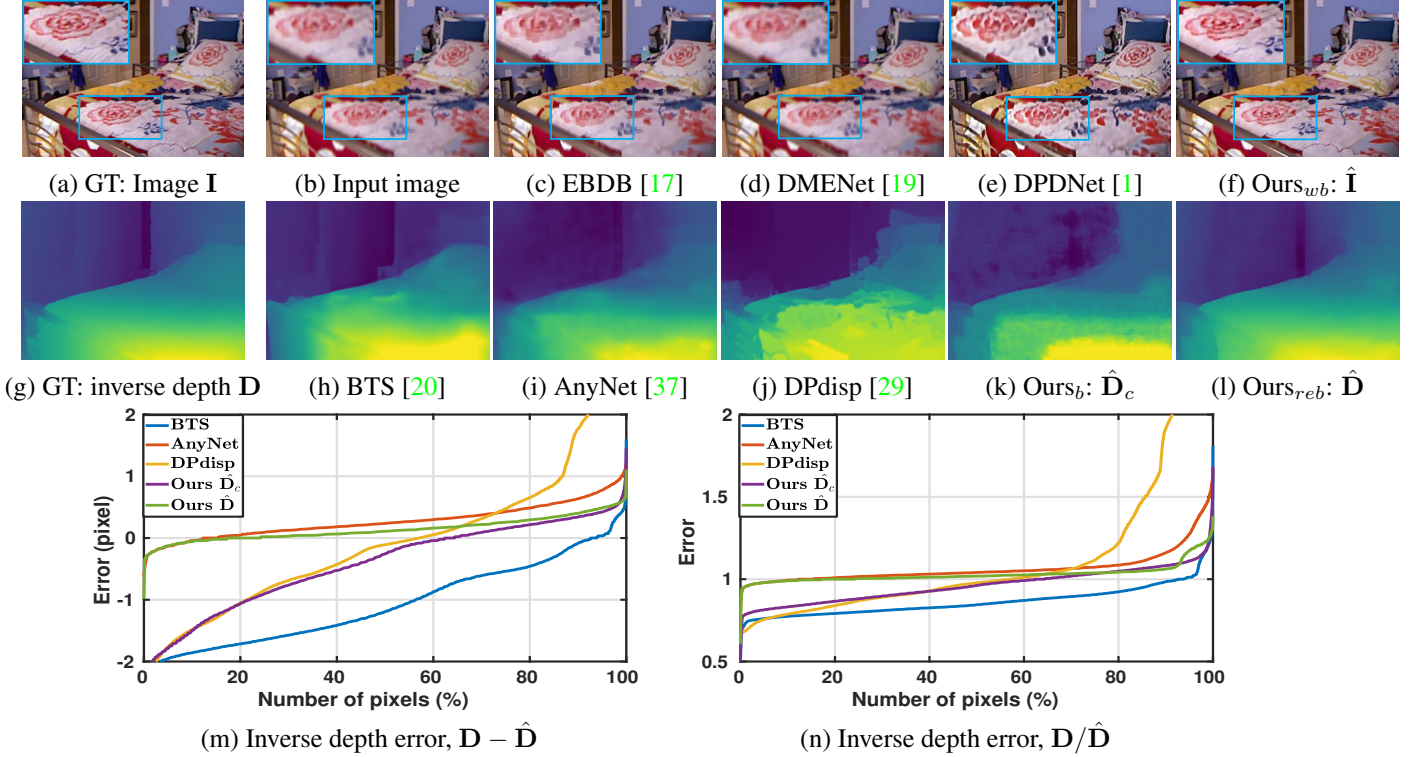


Figure 8. An example of deblurring and inverse depth estimation results on our synthetic dataset. (a) The ground truth image. (b) The left view of the simulated DP image pair. (c) Deblurred by EBDB [17]. (d) Deblurred by DMENet [19]. (e) Deblurred by DPDNet [1]. (f) Our deblurring result. (g) The ground truth inverse depth. (h) Depth by the monocular based method BTS [20] (converted to inverse depth for display). (i) Inverse depth by the stereo based method AnyNet [37]. (j) Inverse depth by the DP based method DPdisp [29]. (k) Output inverse depth of the DepthNet (see Fig. 5). (l) Our inverse depth result. (m) Inverse depth error: $D - \hat{D}$, the closer to zero, the better. (n) Inverse depth error: D/\hat{D} , the closer to 1, the better. The error distributions help to statistically analyze the accuracy of each method. Our method outperforms other baselines significantly. (Best viewed in colour on screen.)

loss, we further improve our performance. For example, for the ‘abs_rel’ metric, the relative error for the coarse depth map as the output of the DepthNet (see Fig. 5), the DeblurNet without reblur loss and the DeblurNet with reblur loss is 0.149, 0.091 and 0.083, respectively.

To further illustrate the superiority of our results, we reconstruct the scene to other views with the estimated sharp image and depth map. The reconstructed videos can be found in our [project page](#).

Effectiveness of our simulated DP images. To further show the usefulness of our simulated DP images, our model trained on Our-syn dataset is directly used to test on the DPD-blur dataset. The result is 20.28dB/0.650/9.68% in PSNR/SSIM/RMSE_{rel} without any fine-tuning. After fine-tuning, our model trained on Our-syn get remarkable performance, with PSNR/SSIM/RMSE_{rel} at 26.92dB/0.864/4.51%, outperforming the performance of our model (26.76dB/0.891/4.59%) directly trained on the DPD-blur dataset. To further show the transfer ability of our model trained using simulated DP images, we only use a half amount of images from the DPD-blur dataset to finetune our model trained on Our-syn dataset, obtain-

ing a good performance, with PSNR/SSIM/RMSE_{rel} at 26.52dB/0.822/4.72% in 20 epoches.

6. Conclusions

In this paper, we derive a mathematical DP model formulating the imaging process of the DP camera. We propose an end-to-end **DDDNet** (DP-based Depth and Deblur Network), to jointly estimate a inverse depth map and restore a sharp image from a blurred DP image pair. Also, we present a reblur loss building upon our DP model and integrate it into our DDDNet. For future work, we plan to extend our DDDNet to a self-supervised depth and sharp image estimation network.

7. Acknowledgment

This research was supported in part by the (ARC) fellowship and Discovery Project grant (DE180100628, DP200102274, CE140100016, DP190102261). Liyuan Pan would like to thank Dr. Liu Liu for proofreading the article and constructive feedback, and Dr. Yiran Zhong for his support when designing the DepthNet.

References

- [1] Abdullah Abuolaim and Michael S Brown. Defocus deblurring using dual-pixel data. In *European Conference on Computer Vision*, pages 111–126. Springer, 2020. 1, 2, 6, 7, 8
- [2] Filippo Aleotti, Fabio Tosi, Li Zhang, Matteo Poggi, and Stefano Mattoccia. Reversing the cycle: self-supervised deep stereo through enhanced monocular distillation. In *European Conference on Computer Vision*, pages 614–632. Springer, 2020. 2
- [3] Ching-Hui Chen, Hui Zhou, and Timo Ahonen. Blur-aware disparity estimation from defocus stereo images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 855–863, 2015. 2
- [4] Xuelian Cheng, Yiran Zhong, Mehrtash Harandi, Yuchao Dai, Xiaojun Chang, Tom Drummond, Hongdong Li, and Zongyuan Ge. Hierarchical neural architecture search for deep stereo matching. *arXiv preprint arXiv:2010.13501*, 2020. 5
- [5] Franklin C Crow. Summed-area tables for texture mapping. In *Proceedings of the 11th annual conference on Computer graphics and interactive techniques*, pages 207–212, 1984. 5
- [6] Shivam Duggal, Shenlong Wang, Wei-Chiu Ma, Rui Hu, and Raquel Urtasun. Deeppruner: Learning efficient stereo matching via differentiable patchmatch. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019. 2
- [7] David Eigen, Christian Puhrsch, and Rob Fergus. Depth map prediction from a single image using a multi-scale deep network. *Advances in Neural Information Processing Systems*, 27:2366–2374, 2014. 6
- [8] John Ens and Peter Lawrence. An investigation of methods for determining depth from focus. *IEEE Transactions on pattern analysis and machine intelligence*, 15(2):97–108, 1993. 2
- [9] Huan Fu, Mingming Gong, Chaohui Wang, Kayhan Batmanghelich, and Dacheng Tao. Deep ordinal regression network for monocular depth estimation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2002–2011, 2018. 2
- [10] Rahul Garg, Neal Wadhwa, Sameer Ansari, and Jonathan T Barron. Learning single camera depth estimation using dual-pixels. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7628–7637, 2019. 1, 2
- [11] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 5
- [12] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J Brostow. Digging into self-supervised monocular depth estimation. In *Proceedings of the IEEE international conference on computer vision*, pages 3828–3838, 2019. 2, 6, 7
- [13] Rostam Affendi Hamzah and Haidi Ibrahim. Literature survey on stereo vision disparity map algorithms. *Journal of Sensors*, 2016, 2016. 2
- [14] R. I. Hartley and A. Zisserman. *Multiple View Geometry in Computer Vision – 2nd Edition*. Cambridge University Press, 2004. 4
- [15] Charles Herrmann, Richard Strong Bowen, Neal Wadhwa, Rahul Garg, Qiurui He, Jonathan T Barron, and Ramin Zabih. Learning to autofocus. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2230–2239, 2020. 1
- [16] Jinbeum Jang, Yoonjong Yoo, Jongheon Kim, and Joonki Paik. Sensor-based auto-focusing system using multi-scale feature extraction and phase correlation matching. *Sensors*, 15(3):5747–5762, 2015. 1
- [17] Ali Karaali and Claudio Rosito Jung. Edge-based defocus blur estimation with adaptive scale selection. *IEEE Transactions on Image Processing*, 27(3):1126–1137, 2017. 6, 8
- [18] Diederick P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations (ICLR)*, 2015. 6
- [19] Junyong Lee, Sungkil Lee, Sunghyun Cho, and Seungyong Lee. Deep defocus map estimation using domain adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 12222–12230, 2019. 2, 6, 8
- [20] Jin Han Lee, Myung-Kyu Han, Dong Wook Ko, and Il Hong Suh. From big to small: Multi-scale local planar guidance for monocular depth estimation. *arXiv preprint arXiv:1907.10326*, 2019. 2, 5, 6, 7, 8
- [21] Anat Levin, Rob Fergus, Frédo Durand, and William T Freeman. Image and depth from a conventional camera with a coded aperture. *ACM transactions on graphics (TOG)*, 26(3):70–es, 2007. 2
- [22] Fayao Liu, Chunhua Shen, Guosheng Lin, and Ian Reid. Learning depth from single monocular images using deep convolutional neural fields. *IEEE transactions on pattern analysis and machine intelligence*, 38(10):2024–2039, 2015. 2
- [23] Maxim Maximov, Kevin Galim, and Laura Leal-Taixé. Focus on defocus: bridging the synthetic to real domain gap for depth estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1071–1080, 2020. 2
- [24] Pushmeet Kohli Nathan Silberman, Derek Hoiem and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *Proceedings of the European Conference on Computer Vision (ECCV)*, 2012. 6
- [25] Liyuan Pan, Yuchao Dai, Miaomiao Liu, and Fatih Porikli. Simultaneous stereo video deblurring and scene flow estimation. July 2017. 2
- [26] Liyuan Pan, Yuchao Dai, Miaomiao Liu, and Fatih Porikli. Depth map completion by jointly exploiting blurry color images and sparse depth maps. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1377–1386. IEEE, 2018. 2
- [27] Liyuan Pan, Yuchao Dai, Miaomiao Liu, Fatih Porikli, and Quan Pan. Joint stereo video deblurring, scene flow estimation and moving object segmentation. *IEEE Transactions on Image Processing*, 29:1748–1761, 2019. 2
- [28] Liyuan Pan, Richard Hartley, Miaomiao Liu, and Yuchao Dai. Phase-only image based kernel estimation for single image blind deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6034–6043, 2019. 6

- [29] Abhijith Punnappurath, Abdullah Abuolaim, Mahmoud Afifi, and Michael S. Brown. Modeling defocus-disparity in dual-pixel sensors. In *IEEE International Conference on Computational Photography (ICCP)*, 2020. 1, 2, 5, 6, 7, 8
- [30] Abhijith Punnappurath and Michael S Brown. Reflection removal using a dual-pixel sensor. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1556–1565, 2019. 1, 2
- [31] Ashutosh Saxena, Min Sun, and Andrew Y Ng. Learning 3-d scene structure from a single still image. In *2007 IEEE 11th International Conference on Computer Vision*, pages 1–8. IEEE, 2007. 2
- [32] Przemysław Śliwiński and Paweł Wachel. A simple model for on-sensor phase-detection autofocusing algorithm. *Journal of Computer and Communications*, 1(06):11, 2013. 1
- [33] Supasorn Suwajanakorn, Carlos Hernandez, and Steven M Seitz. Depth from focus with your mobile phone. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3497–3506, 2015. 2
- [34] Fabio Tosi, Filippo Aleotti, Matteo Poggi, and Stefano Mattoccia. Learning monocular depth estimation infusing traditional stereo knowledge. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [35] Paul Viola, Michael Jones, et al. Robust real-time object detection. *International journal of computer vision*, 4(34-47):4, 2001. 5
- [36] Neal Wadhwa, Rahul Garg, David E Jacobs, Bryan E Feldman, Nori Kanazawa, Robert Carroll, Yair Movshovitz-Attias, Jonathan T Barron, Yael Pritch, and Marc Levoy. Synthetic depth-of-field with a single-camera mobile phone. *ACM Transactions on Graphics (TOG)*, 37(4):1–13, 2018. 1, 2, 5, 6, 7
- [37] Yan Wang, Zihang Lai, Gao Huang, Brian H Wang, Laurens Van Der Maaten, Mark Campbell, and Kilian Q Weinberger. Anytime stereo image depth estimation on mobile devices. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 5893–5900. IEEE, 2019. 2, 6, 7, 8
- [38] Guodong Xu, Yuhui Quan, and Hui Ji. Estimating defocus blur via rank of local patches. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 5371–5379, 2017. 2
- [39] Haofei Xu and Juyong Zhang. Aanet: Adaptive aggregation network for efficient stereo matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1959–1968, 2020. 2
- [40] Zhichao Yin, Trevor Darrell, and Fisher Yu. Hierarchical discrete distribution decomposition for match density estimation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [41] Jure Žbontar and Yann LeCun. Stereo matching by training a convolutional neural network to compare image patches. *The journal of machine learning research*, 17(1):2287–2318, 2016. 2
- [42] Hongguang Zhang, Yuchao Dai, Hongdong Li, and Piotr Koniusz. Deep stacked hierarchical multi-patch network for image deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 5
- [43] Yinda Zhang, Neal Wadhwa, Sergio Orts-Escolano, Christian Häne, Sean Fanello, and Rahul Garg. Du²net: Learning depth estimation from dual-cameras and dual-pixels. *arXiv preprint arXiv:2003.14299*, 2020. 1, 2
- [44] Shaojie Zhuo and Terence Sim. Defocus map estimation from a single image. *Pattern Recognition*, 44(9):1852–1858, 2011. 2