

Improving Single-Image Defocus Deblurring: How Dual-Pixel Images Help Through Multi-Task Learning

Abdullah Abuolaim Mahmoud Afifi Michael S. Brown
York University

{abuolaim, mafifi, mbrown}@eecs.yorku.ca

Abstract

Many camera sensors use a dual-pixel (DP) design that operates as a rudimentary light field providing two sub-aperture views of a scene in a single capture. The DP sensor was developed to improve how cameras perform auto-focus. Since the DP sensor’s introduction, researchers have found additional uses for the DP data, such as depth estimation, reflection removal, and defocus deblurring. We are interested in the latter task of defocus deblurring. In particular, we propose a single-image deblurring network that incorporates the two sub-aperture views into a multi-task framework. Specifically, we show that jointly learning to predict the two DP views from a single blurry input image improves the network’s ability to learn to deblur the image. Our experiments show this multi-task strategy achieves +1dB PSNR improvement over state-of-the-art defocus deblurring methods. In addition, our multi-task framework allows accurate DP-view synthesis (e.g., ~ 39 dB PSNR) from the single input image. These high-quality DP views can be used for other DP-based applications, such as reflection removal. As part of this effort, we have captured a new dataset of 7,059 high-quality images to support our training for the DP-view synthesis task.

1. Introduction and related work

We are interested in reducing the defocus blur present in captured images. Defocus blur occurs at scene points that are captured outside a camera’s depth of field (DoF). Reducing defocus blur is challenging due to the nature of the spatially varying point spread functions (PSFs) that vary with scene depth [16, 29]. Most of the existing DoF blur reduction methods [5, 11, 14, 23, 27] approach the problem in two stages: (1) estimate the defocus map of the input and (2) apply off-the-shelf non-blind deconvolution (e.g., [6, 13]) guided by the estimated defocus map. The performance of these methods is bounded by the DoF map estimation and the effectiveness of the non-blind deconvolution. Additionally, due to the two-stage approach, these methods have a

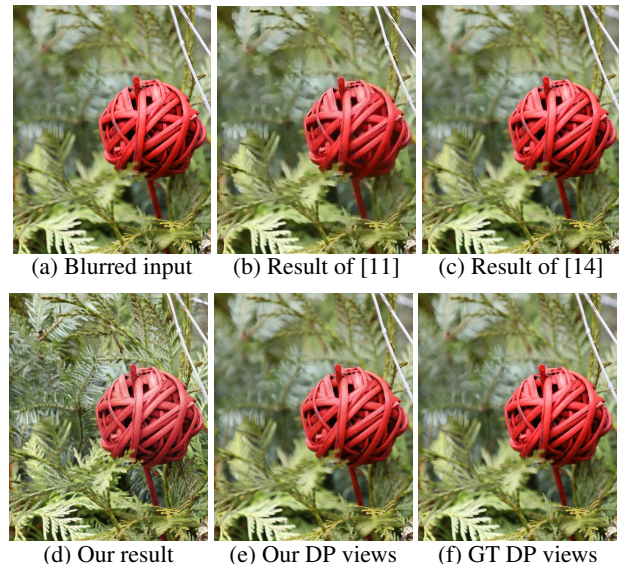


Figure 1: Results from our multi-task framework. (a) Input image with DoF blur. Deblurring results of [11] and [14] are shown in (b) and (c), respectively. (d) Our result. (e) Our reconstructed DP views. (f) Ground-truth DP views. Our multi-task method has better deblurring results and is able to produce accurate DP views from a *single-image* input. **Note: The difference in the DP views is better to be visualized with animation. However, the IEEE PDF eXpress validator does not allow the animation package. Therefore, we provide videos of all animated figures in the supp. material. We also provide in-PDF animated figures in our arXiv version. This feature is applicable where “animations” appears in the subsequent figures.**

long processing time.

Recent work in [1] proposed a DoF deblurring method that leveraged the availability of dual-pixel (DP) sensor data. This method trained a deep neural network (DNN) that uses the DP sensor’s two sub-aperture views as input to predict a single deblurred image. The effectiveness of the method by [1] is attributed to the DNN’s ability to learn the amount of spatially varying defocus blur from the two

DP views. This idea stems from the way the DP sensors work. DP sensors were developed as a means to improve the camera’s autofocus system. The DP design produces two sub-aperture views of the scene that exhibit differences in phase that are correlated to the amount of defocus blur as shown in Fig. 1-f (details in Sec. 2). A camera adjusts the lens position to minimize phase differences in the two DP views, resulting in a final in-focus image. Researchers have been quick to leverage the DP sub-images for tasks beyond autofocus [3, 2], including depth map estimation [24, 7, 36], defocus deblurring [4, 31, 15, 22], reflection removal [25], and synthetic DoF [32].

One notable drawback of using DP data is that most cameras do not provide easy access to the DP sensor’s two sub-aperture views. Although DP sensors are used by many cameras, only two cameras currently provide access to DP images (i.e., Canon 5D DSLR camera [1, 24, 25] and Google Pixel series smartphone camera [32, 7, 36]). Even for these devices accessing the DP data has caveats. For example, the Canon 5D requires special software to extract the two views from a saved RAW image. The Google Pixel device requires a special binary and provides DP data only for the green channel of the RAW image. These limitations make the use of DP data at inference time impractical.

In the context of defocus deblurring, training data is required in the form of paired images—one sharp and one blurred. Training images are obtained by placing a camera on a tripod and capturing an image using a wide aperture (i.e., blurred image with shallow DoF), followed by a second image captured using a narrow aperture (i.e., target sharp image with large DoF). Care must be taken to ensure that the camera is not moved between aperture adjustments and that the scene content remains stationary. Such data acquisition is a time-consuming process and does not facilitate collecting larger datasets—for instance, the recent DP defocus deblurring dataset [1] contains only 500 of such pairs.

The aforementioned drawbacks of accessing DP data at inference time and the challenges in capturing blurry/sharp paired data for training serve as the impetus for our proposed multi-task learning framework. In particular, our method focuses on conventional single-image input. And, while we cannot remove the need for blurred/sharp training data entirely, we demonstrate that the performance of defocus deblurring is improved by incorporating the joint training of predicting the DP views. The training of the DP-view reconstruction task requires only the capture of unpaired DP images in an unrestricted manner with minimal effort. Because we only need access to DP information at training time, inference time becomes much more practical.

Contributions We introduce a multi-task DNN framework to jointly learn single-image defocus deblurring and DP-based view prediction as shown in Fig. 1. We show

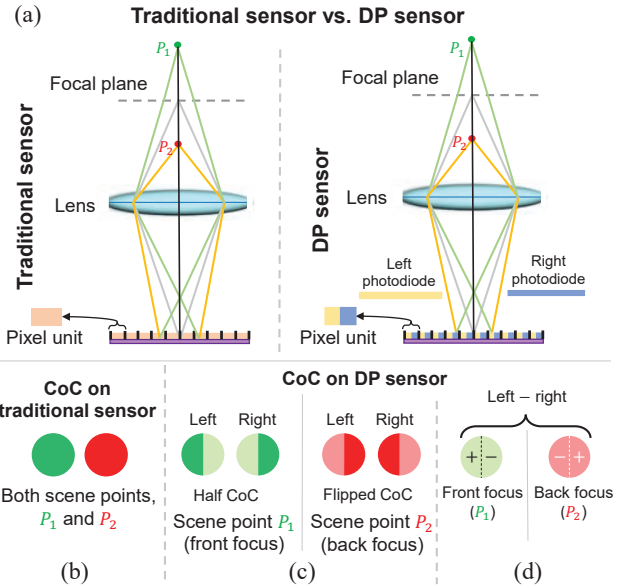


Figure 2: DP sensor image formation via DoF circle of confusion (CoC) formation. (a) Traditional sensor vs. DP sensor. (b) and (c) are the CoC formation on the 2D imaging sensor of two scene points, P_1 and P_2 . On the two DP views, the half-CoC flips direction if the scene point is in front or back of the focal plane. (d) shows the subtracted DP views in the front/back focus cases, where the $+/-$ sign reveals the front/back focus ambiguity.

that training a DNN to both deblur the image and predict the two sub-aperture DP views improves deblurring results by up to +1dB PSNR over existing state-of-the-art methods. To facilitate this effort, we have captured an unpaired dataset with varying DoF blur consisting of 2,353 high-quality full-frame images using a DP camera. This gives a total of 7,059 images—2,353 conventional images and their corresponding two sub-aperture DP views. We also introduce novel loss functions based on DP image formation to help the network avoid ambiguity that arises in DP data. Extensive experiments show our results outperform existing single-image DoF deblurring both quantitatively and qualitatively. While our main goal is defocus deblurring, we conclude by showing how our DNN model has the added advantage of being able to produce high-quality DP views that can be used for tasks such as reflection removal and multi-view synthesis. Our dataset, code, and trained models are made publicly available ¹.

2. Dual-pixel image formation

We start with a brief overview of DP sensors. A DP sensor uses two photodiodes at each pixel location with a mi-

¹<https://github.com/Abdullah-Abuolaim/multi-task-defocus-deblurring-dual-pixel-nimat>

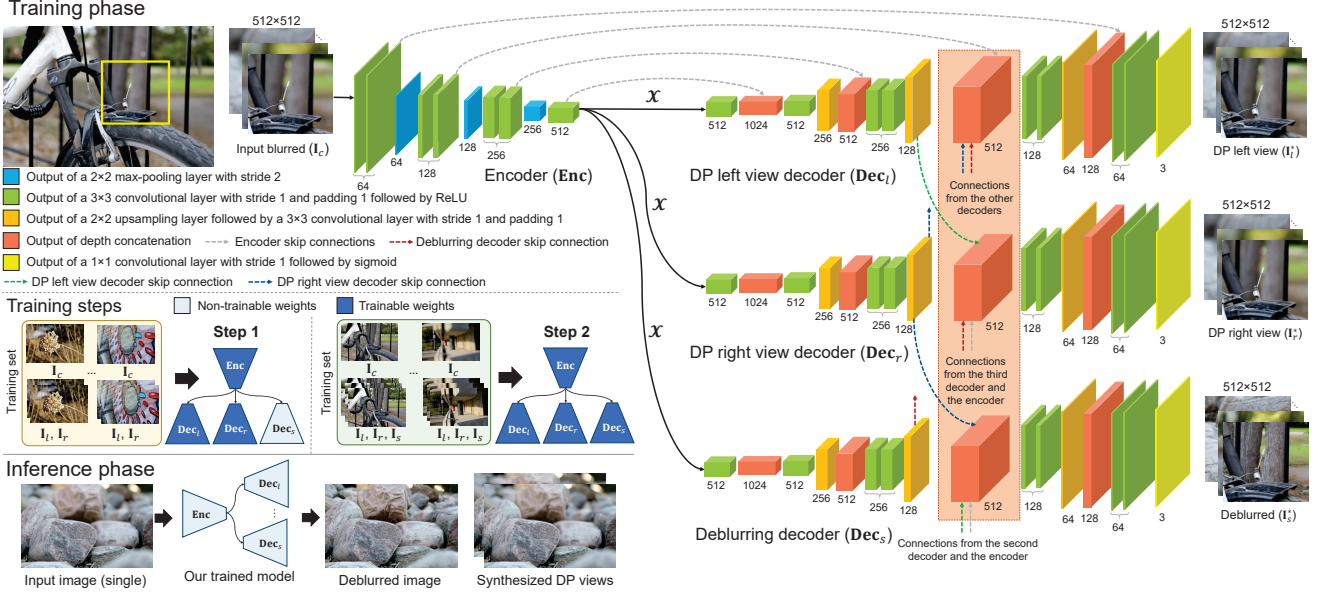


Figure 3: An overview of the proposed multi-task learning framework. We adopt a single-encoder multi-decoder DNN. Our multi-task DP network (MDP) takes a single input image (I_c) and outputs three images—namely, left (I_l) and right (I_r) DP views, and a deblurred (sharp) version (I_s). MDP has two stages of weight sharing between the three decoders (i.e., Dec_l , Dec_r , and Dec_s): early at the encoder (Enc_l) latent space \mathcal{X} and middle at the box highlighted in orange. Our MDP is trained in two steps, where the Dec_s is frozen in the first step and resumed in the next based on the intended task.

croLens placed on the top of each pixel site, as shown in Fig. 2-a. As previously mentioned, this design was developed to improve camera autofocus by functioning as a simple two-sample light field camera. The two-sample light field provides two sub-aperture views of the scene and, depending on the sensor’s orientation, the views can be referred to as left/right or top/down pairs; we follow the convention of prior papers [1, 24, 25] and refer to them as the left/right pair, denoted as I_l and I_r . The light rays coming from scene points that are within the camera’s DoF exhibit little to no difference in phase between the views. On the other hand, light rays coming from scene points outside the camera’s DoF exhibit a noticeable defocus disparity in the *left-right* views. The amount of defocus disparity is correlated to the amount of defocus blur.

Unlike traditional stereo, the difference between the DP views can be modeled as the latent sharp image being blurred in two different directions using a half-circle PSF [24]. This is illustrated in the resultant circle of confusion (CoC) of Fig. 2-c. On real DP sensors, the ideal case of a half-circle CoC is only an approximation due to constraints of the sensor’s construction and lens array. These constraints allow a part of the light ray bundle to leak into the other-half dual pixels (see half CoC of left/right views in Fig. 2-c). We can describe the DP image formation as follows. Let I_s be a latent sharp image patch and H_l and

H_r are the left/right PSFs; then the DP I_l and I_r are:

$$I_l = I_s * H_l, \quad I_r = I_s * H_r, \quad (1)$$

$$H_r = H_l^f, \quad (2)$$

where $*$ denotes the convolution operation and H_l^f is the flipped H_l . The two views I_l and I_r are combined to produce the final image provided by the camera I_c as follows:

$$I_c = I_l + I_r. \quad (3)$$

Another interesting property of the DP PSFs is that the orientation of the “half CoC” of each left/right view reveals if the scene point is in front or back of the focal plane, as shown in the subtracted views of the two scene points, P_1 and P_2 in Fig. 2-d. These DP properties are useful cues that will be considered when formulating the DP-based loss functions in Sec. 3.2.

3. Multi-task learning framework

Multi-task learning has been successfully used for various computer vision tasks [8, 34, 18]. In this work, we adopt a multi-task framework in order to leverage the strong connection between defocus deblurring and DP-view synthesis as they require encoding information regarding the defocus blur present at each pixel in the input image. Towards this goal, we propose a single-encoder multi-decoder DNN that takes a single input image and decomposes it into

DP left/right views along with the deblurred version. Fig. 3 provides an overview of our proposed framework.

3.1. Network architecture

We adopt a symmetric single-encoder multi-decoder DNN architecture with skip connections between the corresponding feature maps [19, 26] (see Fig. 3). We refer to our DNN model as the multi-task DP network (MDP). The three decoder branches have an early-stage weight sharing at the end of the encoder part. We add middle-stage weight sharing as indicated in the orange box of Fig. 3. Each block in the middle-stage sharing receives two skip connections from the corresponding feature maps from the other two decoders. This type of multi-decoder stitching—that guarantees weight sharing at multiple stages—provides multiple communication layers that can further assist the multi-task joint training. We avoid adding late-stage weight sharing as the sharpness of the deblurred image can be affected by the half-PSF blur present in feature maps of the synthesized DP views at these later stages. The proposed model has a sufficiently large receptive field that is able to cover larger spatially varying defocus PSFs. An ablation study is provided in the supplementary material that validates the multi-decoder stitching design.

With the proposed architecture, the encoder **Enc** task is to map the input image into a latent space \mathcal{X} as follows:

$$\mathcal{X} = \text{Enc}(\mathbf{I}_c). \quad (4)$$

This latent space can be viewed as a defocus estimation space in which both tasks share a common goal that requires a notion of the PSF size at each pixel in the input image. This latent space representation \mathcal{X} is then passed to the three decoders—namely, left and right DP-view decoders (Dec_l and Dec_r) and the defocus deblurring (i.e., sharp image) decoder (Dec_s)—in order to produce the output estimations as follows:

$$\mathbf{I}_l^* = \text{Dec}_l(\mathcal{X}), \quad \mathbf{I}_r^* = \text{Dec}_r(\mathcal{X}), \quad \mathbf{I}_s^* = \text{Dec}_s(\mathcal{X}). \quad (5)$$

3.2. DP-based loss function

It is important to consider how the DP images are formed when designing loss functions to ensure the training process for the two DP views satisfies DP properties. We have observed empirically that a traditional mean squared error (MSE) loss, computed between the ground truth (GT) and reconstructed DP views, drives the network to a local minima, where the difference between the reconstructed DP views is estimated as an explicit shift in the image content. This observation makes the MSE alone not sufficient to capture the flipping property of DP PSFs (i.e., the PSF reverses direction if it is in front of the focal plane—see Fig. 2-c). Therefore, we introduce a novel DP-loss based on Eq. 3 that

imposes a constraint on the DP-view reconstruction process as follows:

$$\mathcal{L}_C = \frac{1}{n} \sum_n (\mathbf{I}_c - (\mathbf{I}_l^* + \mathbf{I}_r^*))^2, \quad (6)$$

where \mathbf{I}_c is the input combined image and \mathbf{I}_l^* and \mathbf{I}_r^* are the estimated DP views. Our \mathcal{L}_C encourages the network to optimize for the fundamental DP image formation (i.e., Eq. 3).

While \mathcal{L}_C assists the network to learn that the combined left/right views should sum to the combined image, the front/back focus flipping direction remains ambiguous to the network. To address this ambiguity, we introduce a new view difference loss \mathcal{L}_D to capture the flipping sign direction as follows:

$$\mathcal{L}_D = \frac{1}{n} \sum_n ((\mathbf{I}_l - \mathbf{I}_r) - (\mathbf{I}_l^* - \mathbf{I}_r^*))^2, \quad (7)$$

where \mathbf{I}_l and \mathbf{I}_r are the GT DP left and right views, respectively. Fig. 2-d shows the sign difference in the front/back focus cases when the views are subtracted, which gives a cue for the network to learn the PSF flipping direction when penalizing view difference in the loss—namely, \mathcal{L}_D . Fig. 4 provides a visual analysis of the observations that led to our two DP loss functions—namely, \mathcal{L}_C and \mathcal{L}_D . See Sec. 4.3 for an ablation to demonstrate their effectiveness.

3.3. Dual-pixel datasets

Our new DP dataset We collected a new *diverse* and *large* DP (DLDP) dataset of 2,353 scenes. These scenes were captured using a Canon EOS 5D DSLR camera. Each scene consists of a high-quality combined image (2,353 images) with its corresponding DP views (2,353 × 2 images). All images are captured at full-frame resolution (i.e., 6720 × 4480 pixels). Our data contains indoor and outdoor scenes with diverse image content, weather conditions, scene illuminations, and day/night scenes. We captured scenes with different aperture sizes (i.e., $f/4$, $f/5.6$, $f/10$, $f/16$, and $f/22$) in order to cover a wider range of spatially varying defocus blur (i.e., from all-in-focus to severely blurred images). We use this DLDP dataset in our multi-task framework to optimize directly for the DP-view synthesis task.

Other DP datasets We use the Canon DP deblurring dataset [1] (i.e., 350 training paired images) to optimize for both defocus deblurring and DP-view synthesis. We note that there is also a DP dataset based on the Google Pixel camera [7]. We opted to only use the Canon as it matches our DLDP dataset. Additionally, the Pixel DP data from [7] is not appropriate to train for colored images as it provides only the green channels in the raw-Bayer frame for the DP views.

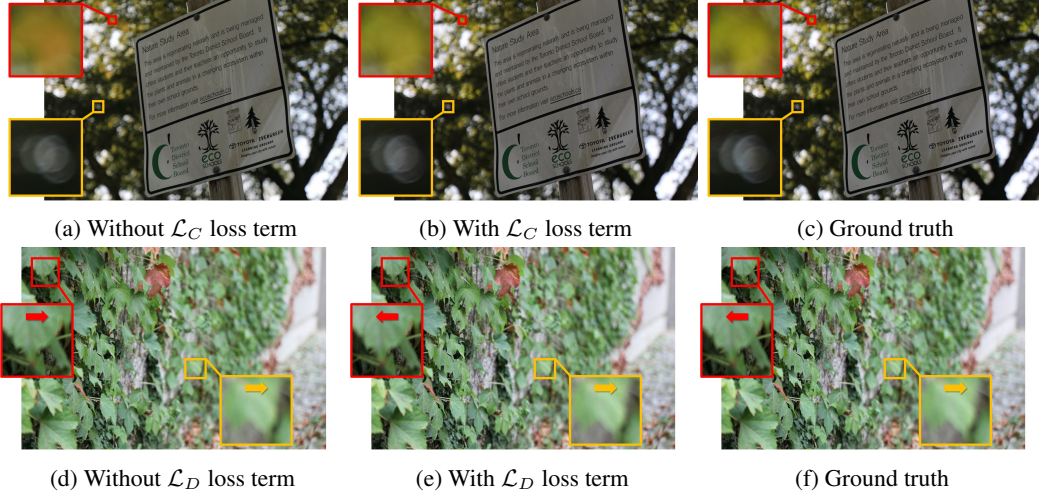


Figure 4: An analysis and visual comparison to reflect the effectiveness of our proposed \mathcal{L}_C and \mathcal{L}_D DP-based loss terms. (b) shows that training with \mathcal{L}_C helps the network to capture the flipping kernel (yellow patch) and accurate colors (red patch) compared to the one in (a). (e) demonstrates that training with \mathcal{L}_D can assist the network to learn the flipping direction in the front (yellow patch) and back focus (red patch), where the views rotate around the focal plane as shown in the GT. **Note: this figure contains animations.**

3.4. Model training

Our training is divided into two steps. First, training with image patches from our DLDP dataset is performed to optimize only the DP-view synthesis task. During this step the weights of the deblurring decoder branch (Dec_s) are frozen. Once the model converges for the DP-view synthesis branches, the second step unfreezes the weights of Dec_s and starts fine-tuning using image patches from the Canon DP deblurring dataset [1] to optimize jointly for both the defocus deblurring and DP-view synthesis tasks (see training steps in Fig. 3). For the first step, we train the network with the following loss terms:

$$\mathcal{L}_{ST1} = \mathcal{L}_{MSE}(l, r) + \mathcal{L}_C + \mathcal{L}_D, \quad (8)$$

where \mathcal{L}_{ST1} is the overall first-step loss and $\mathcal{L}_{MSE}(l, r)$ is the typical MSE loss between the GT and estimated DP views. The second step needs more careful loss setting to fine-tune the model in a way that guarantees improving performance on both tasks. In the second step, the network is fine-tuned with the following loss terms:

$$\mathcal{L}_{ST2} = \mathcal{L}_{MSE}(s) + \lambda_1 \mathcal{L}_{MSE}(l, r) + \lambda_2 \mathcal{L}_C + \lambda_3 \mathcal{L}_D \quad (9)$$

where \mathcal{L}_{ST2} is the overall second-step loss and $\mathcal{L}_{MSE}(s)$ is the typical MSE between the output deblurred image and the GT. The λ terms are added to control the optimization.

4. Experiments

4.1. Training details

Training data We divide our newly captured DLDP dataset into 2,090 and 263 training and testing scenes, re-

spectively. In the first training step, we use the 2,090 training scenes. For the second training step, we use the DP data from [1] following the same data division in [1]—that is, 350, 74, and 76 training, validation, and testing scenes, respectively.

Training procedure We extract image patches of size $512 \times 512 \times 3$, where the input is a single patch and the output is three patches. The convolutional layer weights are initialized using He’s method [10] and the Adam optimizer [12] is used to train the model. The mini-batch size in each iteration is set to 8 batches.

For the first training step, the initial learning rate is set to 3×10^{-4} , which is decreased by half every 8 epochs. The model converges after 60 epochs in the first step. For the second step, the initial learning rate is set to 6×10^{-4} , which is decreased by half every 8 epochs. The model converges after 80 epochs. The λ terms are set to 0.8, 0.5, and 0.5, respectively, in order to have a balanced loss minimization and to guide the network attention towards minimizing for defocus deblurring in the second step. All the training details and model setup are implemented using Python with the Keras framework on top of TensorFlow 2 and trained with NVIDIA TITAN X GPU.

4.2. Single-image deblurring

To evaluate our method, we use the test set of the Canon DP deblurring dataset [1]. Specifically, this test set [1] includes 37 indoor and 39 outdoor scenes. We compare our results against recent methods for single-image defocus deblurring [27, 11, 14]. For the sake of completeness, we also

Table 1: Quantitative comparisons with single-image defocus deblurring methods. The best results are indicated with bold-face. Results are on the Canon DP deblurring dataset [1] (test set consists of 37 indoor and 39 outdoor scenes).

Method	Indoor			Outdoor			Indoor & Outdoor				
	PSNR \uparrow	SSIM \uparrow	MAE \downarrow	PSNR \uparrow	SSIM \uparrow	MAE \downarrow	PSNR \uparrow	SSIM \uparrow	MAE \downarrow	NIQE \downarrow	Time \downarrow
JNB [27]	26.73	0.828	0.031	21.10	0.608	0.064	23.84	0.715	0.048	5.11	843.1
EBDB [11]	25.77	0.772	0.040	21.25	0.599	0.058	23.45	0.683	0.049	5.42	929.7
DMENet [14]	25.70	0.789	0.036	21.51	0.655	0.061	23.55	0.720	0.049	4.85	613.7
DPDNet (single) [1]	26.54	0.816	0.031	22.25	0.682	0.056	24.34	0.747	0.044	4.06	0.5
Our MDP	28.02	0.841	0.027	22.82	0.690	0.052	25.35	0.763	0.040	3.25	0.5

Table 2: Our single-image defocus deblurring achieves on-par results compared to DPDNet [1], which requires two DP images data as input. Results are on the Canon DP deblurring dataset [1].

Method	PSNR \uparrow	SSIM \uparrow	MAE \downarrow
DPDNet (real DP views) [1]	25.13	0.786	0.041
DPDNet (our synth. DP views) [1]	24.91	0.758	0.043
Our MDP (single image)	25.35	0.763	0.040

compare our results against the recent DP defocus deblurring method (DPDNet) [1], which requires the availability of DP data at inference time. As our task is single-image defocus deblurring, we provide the results of the DPDNet [1] using single input images for a fair comparison. The single-image model of this method was originally trained by the authors of the paper [1] with the same model size (i.e., number of weights) used for the DP model. Although our MDP model has three decoder branches, we adjust the number of convolutional operations and filter size in some blocks in order to have an equivalent model size in terms of number of weights compared to DPDNet [1]. Our proposed MDP is fully convolutional so that we can test on full-size images regardless of the patch size used for training.

Quantitative results Table 1 shows the quantitative results of our method and other single-image defocus deblurring methods: the just noticeable defocus blur method (JNB) [27], the edge-based defocus blur estimation method (EBDB) [11], the deep defocus map estimation method (DMENet) [14], and the DPDNet (single) [1]. We use the common signal processing metrics PSNR, SSIM, and MAE. We also report the Naturalness Image Quality (NIQE) metric of the output deblurred images with respect to a reference model derived from the DP GT images. As shown in Table 1, our method achieves the state-of-the-art results for all metrics compared to other recent single-image defocus deblurring methods. Furthermore, MDP and DPDNet (single) have much lower inference time—that is, $>1,200\times$ faster compared to others.

Though motion blur leads to image blur too, as defocus blur does, the physical formation and consequently the appearance of the resultant blur are different. This was noted in [1], and we found a significant degradation on the accuracy of methods focused on motion blur [30, 37, 20, 28, 21]

when they are applied to defocus blur; for example, Tao et al.’s method [30] for motion deblurring achieves an average PSNR of 20.12dB when it is evaluated on the Canon DP deblurring test set [1], which is significantly lower than all other defocus deblurring methods shown in Table 1.

In Table 1, we reported the results of the DP-based method (i.e., DPDNet) [1] trained on single input. For the sake of completeness, we also compared our method against this method when it is fed with real DP data as input. Table 2 shows this comparison. As can be seen, our method achieves higher PSNR and MAE but lower SSIM compared to DPDNet [1], while our method is more practical as it requires only single-input images compared to the DPDNet [1], which requires accessing the two DP images at the inference phase. Recall that DPDNet cannot be trained on our DLDP dataset as it is unpaired and does not have the corresponding GT all-in-focus image. This point is a central strength of our approach, as the use of unpaired DP images enables capturing a much larger dataset for training, thereby making the learning process more efficient.

Qualitative results Our method achieves better qualitative results when compared with several existing single-image defocus deblurring methods (as shown in Fig. 1). We provide additional qualitative comparisons in Fig. 5, where we compare our results against the results of the EBDB [11], DMENet [14], and the DPDNet (single) [1] methods. Fig. 5 shows that our method achieves results that are arguably visually superior to the other methods. Additional results on other cameras are also available in the supplementary material.

4.3. Ablation study

As explained in Sec. 3, our method is a multi-task framework that is suitable not only for reducing defocus blur but also for predicting DP views of the input single image. Our multi-task framework allows our method to improve the results of each task, as they are inherently correlated. Intuitively, we can re-design our framework by training a single model for each task separately. Table 3 shows the results of training a single model (with approximately the same capacity of our multi-task framework) on each task separately. Table 3 also shows the results of training both single and multi-task frameworks with and without our DP-based loss



Figure 5: Qualitative comparison with others on the test set of the Canon DP dataset [1]. We compare our results with the following single-image defocus deblurring methods: EBDB [11], DMENet [14], and DPDNet (single) [1]. Note that DPDNet was originally introduced to use DP images as input, but the authors in [1] also provided the same model trained on a single image, denoted as DPDNet (single). Our method produces the best quantitative and arguably qualitative results.

functions introduced in Sec. 3.2. As shown, our multi-task framework with our introduced loss functions achieves the best results. Supplementary material provides an additional ablation study on our design of the multi-task architecture.

4.4. DP-view synthesis and application

An interesting side effect of our multi-task network is the ability to perform view synthesis. For instance, we can

generate an aesthetically realistic image motion by synthesizing a multi-view version of a given single image. As discussed in Sec. 2, the DP two sub-aperture views of the scene depend on the sensor’s orientation and, in this paper, our DLDP dataset contains left/right DP pairs, and consequently our network synthesizes the horizontal DP disparity. We can synthesize additional views with different “DP disparity” by rotating the input image before feeding it to



Figure 6: Our novel NIMAT effect. Aesthetically pleasing image motion produced by our multi-view synthesis. The first two input images were taken from the Canon DP deblurring dataset [1], while the others are from Flickr. **Note: this figure contains animations.**

Table 3: Ablation study to demonstrate the effectiveness of our DP-based loss terms and multi-task training. Results are on the Canon DP deblurring dataset [1]. Note that the \mathcal{L}_C and \mathcal{L}_D are not applicable (N/A) to the single task defocus deblurring as it does not predict the DP views.

Method	Defocus deblurring		DP-pixel synthesis	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
Single-task w/o \mathcal{L}_C and \mathcal{L}_D	24.34	0.747	37.05	0.953
Single-task w/ \mathcal{L}_C and \mathcal{L}_D	N/A	N/A	38.23	0.962
Multi-task w/o \mathcal{L}_C and \mathcal{L}_D	24.81	0.750	38.01	0.957
Multi-task w/ \mathcal{L}_C and \mathcal{L}_D	25.35	0.763	39.17	0.973

Table 4: Reflection removal quantitative results on the dataset proposed in [25]. When using our synthetic DP views, the dual-pixel reflection removal (DPRR) method [25] achieves on-par results compared with using real DP views, which makes the DPRR method applicable with the absence of real DP data.

Single-image		Non-single-image	
Method	PSNR	Method	PSNR
ZN18 [35]	15.57	LB13 [17]	16.12
YG18 [33]	16.49	GC14 [9]	16.02
DPRR [25] (ours)	19.32	DPRR [25] (real DP)	19.45

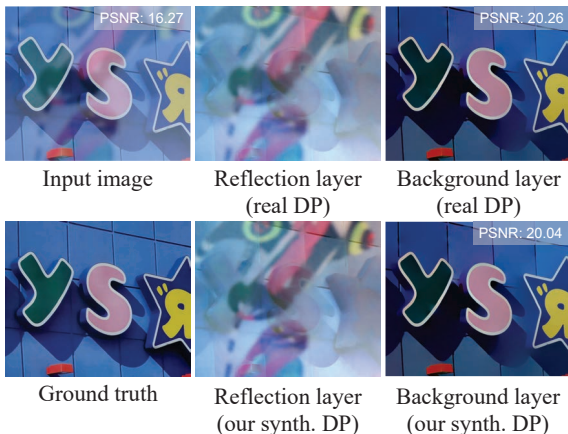


Figure 7: Our synthetic DP views can be used as input for DP-based reflection removal of [25], and are able to produce results on par with those using real DP data.

our network by a 45° clockwise step three times (i.e., 45° , 90° , 135°). This allows us to produce a smooth image mo-

tion from the reconstructed eight views as shown in Fig. 6. We refer to this DP-based view synthesis and image motion as *new image motion attribute* (NIMAT) effect. See supplementary material for Additional examples of our novel NIMAT effect.

As discussed in Sec. 1, DP data has been used for different computer vision tasks. Here we show using our synthesized DP views can be leveraged for a DP-based method in the absence of real DP data. We validate this idea of using our reconstructed DP views as a proxy to DP data on the reflection removal and defocus deblurring tasks. Specifically, we processed input real DP data and our generated DP data using the DP-based reflection removal (DPRR) [25] and defocus deblurring (DPDNet) [1] methods. As shown in Fig. 7, utilizing our synthetic DP views produces approximately the same high-quality result as using DPRR[25] on real DP data. This allows us to achieve better reflection removal results, while still requiring only *a single input image*, compared to other methods for reflection removal, as shown in Table 4. As for DPDNet, Table 2 shows that DPDNet tested with our generated DP views has on-par results compared to the one tested with real DP views.

5. Conclusion

We have demonstrated that a DNN trained for the purpose of single-image defocus deblurring can be improved by incorporating the additional task of synthesizing the two DP views associated with the input image. One benefit of this approach is that capturing data for the DP view synthesis task is easy to perform and requires no special capture setup. This is contrasted with a conventional approach that requires careful capture of sharp/blurred image pairs for the deblurring task. This multi-task strategy is able to improve deblurring results by close to 1dB in terms of PSNR. As an added benefit, we introduced the novel NIMAT effect and showed that our DNN is able to perform realistic view synthesis that can be used for tasks such as reflection removal.

Acknowledgments A big thanks to David Ampofo for his amazing photography skills in collecting the images for our new DLDP dataset.

References

- [1] Abdullah Abuolaim and Michael S Brown. Defocus deblurring using dual-pixel data. In *ECCV*, 2020.
- [2] Abdullah Abuolaim and Michael S Brown. Online lens motion smoothing for video autofocus. In *WACV*, 2020.
- [3] Abdullah Abuolaim, Abhijith Punnappurath, and Michael S Brown. Revisiting autofocus for smartphone cameras. In *ECCV*, 2018.
- [4] Abdullah Abuolaim, Radu Timofte, Michael S Brown, et al. NTIRE 2021 challenge for defocus deblurring using dual-pixel images: Methods and results. In *CVPR Workshops*, 2021.
- [5] Laurent D’Andrès, Jordi Salvador, Axel Kochale, and Sabine Süsstrunk. Non-parametric blur map regression for depth of field extension. *TIP*, 25(4):1660–1673, 2016.
- [6] DA Fish, AM Brinicombe, ER Pike, and JG Walker. Blind deconvolution by means of the richardson–lucy algorithm. *Journal of the Optical Society of America (A)*, 12(1):58–65, 1995.
- [7] Rahul Garg, Neal Wadhwa, Sameer Ansari, and Jonathan T Barron. Learning single camera depth estimation using dual-pixels. In *ICCV*, 2019.
- [8] Ross Girshick. Fast R-CNN. In *CVPR*, 2015.
- [9] Xiaojie Guo, Xiaochun Cao, and Yi Ma. Robust separation of reflection from multiple images. In *CVPR*, 2014.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, 2015.
- [11] Ali Karaali and Claudio Rosito Jung. Edge-based defocus blur estimation with adaptive scale selection. *TIP*, 27(3):1126–1137, 2018.
- [12] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [13] Dilip Krishnan and Rob Fergus. Fast image deconvolution using hyper-laplacian priors. In *NeurIPS*, 2009.
- [14] Junyong Lee, Sungkil Lee, Sunghyun Cho, and Seungyong Lee. Deep defocus map estimation using domain adaptation. In *CVPR*, 2019.
- [15] Junyong Lee, Hyeongseok Son, Jaesung Rim, Sunghyun Cho, and Seungyong Lee. Iterative filter adaptive network for single image defocus deblurring. In *CVPR*, 2021.
- [16] Anat Levin, Yair Weiss, Fredo Durand, and William T Freeman. Understanding blind deconvolution algorithms. *TPAMI*, 33(12):2354–2367, 2011.
- [17] Yu Li and Michael S Brown. Exploiting reflection change for automatic reflection removal. In *ICCV*, 2013.
- [18] Shikun Liu, Edward Johns, and Andrew J Davison. End-to-end multi-task learning with attention. In *CVPR*, 2019.
- [19] Xiaojiao Mao, Chunhua Shen, and Yu-Bin Yang. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. In *NeurIPS*, 2016.
- [20] Seungjun Nah, Sanghyun Son, and Kyoung Mu Lee. Recurrent neural networks with intra-frame iterations for video deblurring. In *CVPR*, 2019.
- [21] Jinshan Pan, Haoran Bai, and Jinhui Tang. Cascaded deep video deblurring using temporal sharpness prior. In *CVPR*, 2020.
- [22] Liyuan Pan, Shah Chowdhury, Richard Hartley, Miaomiao Liu, Hongguang Zhang, and Hongdong Li. Dual pixel exploration: Simultaneous depth estimation and image restoration. In *CVPR*, 2021.
- [23] Jinsun Park, Yu-Wing Tai, Donghyeon Cho, and In So Kweon. A unified approach of multi-scale deep and hand-crafted features for defocus estimation. In *CVPR*, 2017.
- [24] Abhijith Punnappurath, Abdullah Abuolaim, Mahmoud Affi, and Michael S Brown. Modeling defocus-disparity in dual-pixel sensors. In *ICCP*, 2020.
- [25] Abhijith Punnappurath and Michael S Brown. Reflection removal using a dual-pixel sensor. In *CVPR*, 2019.
- [26] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *MICCAI*, 2015.
- [27] Jianping Shi, Li Xu, and Jiaya Jia. Just noticeable defocus blur detection and estimation. In *CVPR*, 2015.
- [28] Shuochen Su, Mauricio Delbracio, Jue Wang, Guillermo Sapiro, Wolfgang Heidrich, and Oliver Wang. Deep video deblurring for hand-held cameras. In *CVPR*, 2017.
- [29] Huixuan Tang and Kiriakos N Kutulakos. Utilizing optical aberrations for extended-depth-of-field panoramas. In *ACCV*, 2012.
- [30] Xin Tao, Hongyun Gao, Xiaoyong Shen, Jue Wang, and Jiaya Jia. Scale-recurrent network for deep image deblurring. In *CVPR*, 2018.
- [31] Tu Vo. Attention! stay focus! In *CVPR Workshops*, 2021.
- [32] 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*, 37(4):64, 2018.
- [33] Jie Yang, Dong Gong, Lingqiao Liu, and Qinfeng Shi. Seeing deeply and bidirectionally: A deep learning approach for single image reflection removal. In *ECCV*, 2018.
- [34] Fisher Yu, Haofeng Chen, Xin Wang, Wenqi Xian, Yingying Chen, Fangchen Liu, Vashisht Madhavan, and Trevor Darrell. BDD100K: A diverse driving dataset for heterogeneous multitask learning. In *CVPR*, 2020.
- [35] Xuaner Zhang, Ren Ng, and Qifeng Chen. Single image reflection separation with perceptual losses. In *CVPR*, 2018.
- [36] Yinda Zhang, Neal Wadhwa, Sergio Orts-Escolano, Christian Häne, Sean Fanello, and Rahul Garg. Du2Net: Learning depth estimation from dual-cameras and dual-pixels. In *ECCV*, 2020.
- [37] Shangchen Zhou, Jiawei Zhang, Jinshan Pan, Haozhe Xie, Wangmeng Zuo, and Jimmy Ren. Spatio-temporal filter adaptive network for video deblurring. In *ICCV*, 2019.