

LDP: Language-driven Dual-Pixel Image Defocus Deblurring Network

Hao Yang¹, Liyuan Pan¹ [†], Yan Yang², Richard Hartley², and Miaomiao Liu²

¹Beijing Institute of Technology ² Australian National University

{hao.yang, liyuan.pan}@bit.edu.cn, {yan.yang, richard.hartley, miaomiao.liu}@anu.edu.au

Abstract

Recovering sharp images from dual-pixel (DP) pairs with disparity-dependent blur is a challenging task. Existing blur map-based deblurring methods have demonstrated promising results. In this paper, we propose, to the best of our knowledge, the first framework that introduces the contrastive language-image pre-training framework (CLIP) to accurately estimate the blur map from a DP pair unsupervisedly. To achieve this, we first carefully design text prompts to enable CLIP to understand blur-related geometric prior knowledge from the DP pair. Then, we propose a format to input a stereo DP pair to CLIP without any fine-tuning, despite the fact that CLIP is pre-trained on monocular images. Given the estimated blur map, we introduce a blur-prior attention block, a blur-weighting loss, and a blur-aware loss to recover the all-in-focus image. Our method achieves state-of-the-art performance in extensive experiments (see Fig. 1).

1. Introduction

Defocus deblurring from a single image is an ill-posed problem [19, 28, 29, 33]. Compared to methods that restore an all-in-focus image directly [1, 2], blur map-based deblurring methods have demonstrated promising results [16, 26]. To simplify the blur removal process, several works [20, 34, 38] have utilized recent innovative dual-pixel (DP) sensors [6] to estimate blur maps, since the DP sensor captures a two-view image pair (a DP pair) in a single shot. The DP pair can be treated as images from a two-sample light field camera or a stereo system [24] with a tiny baseline [27, 34], which is beneficial for disparity estimation [1, 14, 23, 42]. Based on the DP image pair formation process, the disparity is geometrically related to the blur map, *i.e.*, the degree of defocus blur proportional to disparity magnitudes [9].

For estimating the blur map (defocus map or disparity

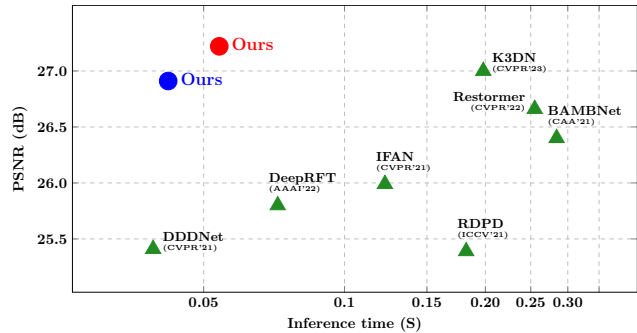


Figure 1. Comparison of performance in terms of PSNR and Inference time. Our small and large models are respectively denoted by the Blue and Red cycles. The Green triangles represent previous state-of-the-art methods, and each of them is annotated with the method name, publication venue, and year. This figure is best viewed in color on the screen.

map), previous works either require extra data [3, 26], such as synthetic data, for supervision signals, or pre-calibrated blur kernels [34]. To avoid using extra information, several works explored unsupervised blur map estimation frameworks [16, 20]. Lee *et al.* [16] implicitly estimate the disparity of the left/right views using a spatial transformer, yet only in feature space, as well as Yang *et al.* [38]. Liang *et al.* [20] use fixed network branches for layer-wise defocus map estimation and deblurring, where the limited network branch number restricts the potential of its accuracy.

Recently, the contrastive language-image pre-training framework (CLIP) [30] has accomplished extraordinary success for vision tasks, such as semantic segmentation [31, 37], object detection [32], and point cloud understanding [40]. Therefore, a question naturally arises, *can we avoid the cost of collecting data and model designing, to estimate the blur map unsupervisedly by using CLIP?* Unfortunately, using semantic knowledge from CLIP to handle low-level vision tasks is underexplored.

In this paper, for the first time, we explore the relationship between language and blur map estimation by applying CLIP. Note that CLIP is trained to match a monocular image with its corresponding language descriptions. Given a blurred image, it is intuitive to match it with a corre-

[†] Corresponding author.

sponding text [A blurry image.], which allows the language model to find the blurred regions in an image. However, it is hard to quantize the blur level of an image without knowing its associated latent sharp one [25].

Fortunately, we can leverage the capabilities of DP sensors. With DP pairs, blur quantization becomes possible without the need for its associated sharp image. The formation of DP ensures no disparity between left/right views for pixels in sharp regions, while a noticeable disparity exists for pixels in blurred regions. In essence, the detectable differences (disparities) in left/right views of a DP pair directly indicate the degree of defocus blurs. Consequently, we can quantify the variation between the left and right views, treating it as a valuable cue for estimating a blur map.

However, CLIP still finds it challenging to simultaneously explore stereo DP pairs with left/right view images. To address this issue, we propose to explore the DP pair and design an ‘image and text’ format to better capture the blurs in the image. Specifically, we propose to create a new image from the DP pair and *translate* the description about the ‘blurriness’ of the image to measure the ‘symmetry’ of the newly formed image, which is based on the DP image formation process. The freshly designed image-text format allows the CLIP to estimate the blur map easily (see Sec. 3.1).

Given the estimated blur map, we then restore the all-in-focus image from the DP pair. By observing that the attention map of the self-attention mechanism [8] can be used in the dynamic DP deblurring process, we introduce a blur prior attention (BPA) that modulates the attention map with the estimated blur map which provides prior knowledge of deblurring kernels (see Sec. 3.2).

To regulate the restoration, we specifically design two losses that are based on CLIP. Considering the network should focus more on heavily blurred regions, we propose a blur-weighting loss that imposes an adaptive penalty proportional to the blur degrees provided by the estimated blur map from CLIP. Then, to guarantee the restored image is sharp, we use CLIP to re-check the restored image and formulate a blur-aware loss (see Sec. 3.3).

Our contributions are summarised as follows:

- We propose a **Language-driven DP (LDP)** defocus deblurring framework, by exploring the potential of CLIP in the low-level vision task.
- We design an image-text format for blur map estimation based on the geometric relationship between the blur and disparity of the DP pair.
- We propose a blur prior attention (BPA) block, a blur-weighting loss, and a blur-aware loss to encourage sharp restorations of the DP pair.

We evaluate the proposed framework extensively on standard benchmark datasets. Though using the large-scale foundation model, CLIP, for blur map estimation, our method achieves Pareto-optimality between deblurring performance and inference speed (refer to Fig. 1 for an

overview of the comparison).

2. Related Work

DP Defocus Deblurring. DP-based deblurring methods can be divided into single- and multi-task manners. For single-task based deblurring [1–3], methods apply different kernels at different regions of the DP pair to handle the spatially varying defocus blur. This can usually be achieved by using parallel atrous convolutions [33], attention mechanisms [39], or deformable convolutions [17]. For multi-task deblurring methods, the blur/defocus/disparity map is auxilarily estimated to benefit DP defocus deblurring. Pan *et al.* [26] and Xin *et al.* [34] use the estimated DP disparity to regularise the restoration through reblurring; however, with extra data or pre-calibrate kernel. Several works explored unsupervised defocus map (or disparity) estimation. Lee *et al.* [16] and Yang *et al.* [38] estimate defocus disparity only in feature space without explicitly showing a defocus map. Liang *et al.* [20] optimize a multi-channel defocus map based on [27]. However, hyperparameters, such as the defocus threshold and the number of channels, are preset, limiting blur map estimation accuracy. Unlike past works, by transferring the blur-related geometric prior knowledge into semantic prompts, we unsupervisedly estimate the blur map from CLIP to benefit defocus deblurring.

Language-driven Vision. Common objectives in visual-language learning are aligning the text with a paired image in the embedding space [21, 30, 35]. The CLIP, one of the milestone achievement in the area, demonstrate that language-independent computer vision tasks can be substantially benefited from visual-language learning, by achieving SOTA zero-shot classification performance in the ImageNet-1K benchmark [7]. Successive works have extended the zero-shot learning capability of CLIP to semantic segmentation [31, 44], monocular depth estimation [4, 10, 41], and point clouds [11, 40]. However, none of the existing works explores the zero-shot learning capability of CLIP in stereo/DP pairs. This paper proposes the first approach to use the monocular image-trained CLIP in zero-shot blur map segmentation of DP pairs.

3. Method

Let us now introduce our language-driven dual-pixel image defocus deblurring framework. Given a defocus blurred DP pair ($\mathbf{B}_L, \mathbf{B}_R$), we aim to restore an all-in-focus image \mathbf{I} , where \mathbf{B}_L and \mathbf{B}_R are the left and right view of a DP pair and $\mathbf{B}_L, \mathbf{B}_R, \mathbf{I} \in \mathbb{R}^{h \times w \times 3}$. Here, h and w define the image height and width, respectively.

We first design a DP-aware prompt to adapt CLIP to output a blur map, $\mathbf{M} \in \mathbb{R}^{h \times w \times 1}$, which quantizes disparity unsupervisedly between \mathbf{B}_L and \mathbf{B}_R . Then, we propose a BPA block based network which employs the blur




Figure 2. Overall architecture. Given a DP pair $(\mathbf{B}_L, \mathbf{B}_R)$, we first generate a blur map \mathbf{M} by using our unsupervised blur map estimation strategy which is based on the CLIP with ensembling format. We then forward the DP pair $(\mathbf{B}_L, \mathbf{B}_R)$ to our deblurring backbone composed of an encoder and a decoder. Meanwhile, we have multiple BPA blocks, using the blur map \mathbf{M} as deblurring kernel prior to refine the intermediate backbone feature embeddings. In the following, we have the restoration $\hat{\mathbf{I}}$. The deblurring network is optimized with three losses, charbonnier loss $\mathcal{L}_{\text{charb}}$, blur weighting loss \mathcal{L}_{bwl} , and blur aware loss \mathcal{L}_{bal} . $\mathcal{L}_{\text{charb}}$ penalizes deviations from \mathbf{I} in the restoration $\hat{\mathbf{I}}$. \mathcal{L}_{bwl} and \mathcal{L}_{bal} regularize the restoration by imposing an adaptive penalty on the blurred regions of the DP pair, and leveraging unsupervised blur map estimation strategy to encourage the restoration to be sharp, i.e., the blur map $\hat{\mathbf{M}}$ for the restoration $\hat{\mathbf{I}}$ is zero value filled.

map \mathbf{M} as deblurring priors to restore a sharp image $\hat{\mathbf{I}}$ from $(\mathbf{B}_L, \mathbf{B}_R)$. To regularize the restoration of $\hat{\mathbf{I}}$, a blur-weighting loss \mathcal{L}_{bwl} and a blur-aware loss \mathcal{L}_{bal} are also introduced that were underpinned by our unsupervised blur map estimation strategy. The overall framework is shown in Fig. 2. Details are presented below.

3.1. Unsupervised Blur Map Estimation

Preliminary. CLIP respectively embeds an image and a text description into a common semantic space of dimension c with an image encoder and a text encoder, obtaining image embedding $\mathbf{f} \in \mathbb{R}^c$ and text embedding $\mathbf{t} \in \mathbb{R}^c$. By optimizing a contrastive loss, the embeddings of the aligned image and text pair, Ω_a , stay similar, i.e., maximizing the semantic similarity by $\max_{\mathbf{f}, \mathbf{t} \in \Omega_a} \text{sim}(\mathbf{f}, \mathbf{t})$, where $\text{sim}(\cdot, \cdot)$ is the cosine similarity metric. Meanwhile, the semantics of misaligned pairs, Ω_{na} are pushed away from each other, i.e., $\min_{\mathbf{f}, \mathbf{t} \in \Omega_{na}} \text{sim}(\mathbf{f}, \mathbf{t})$. With comprehensive training data, CLIP can be leveraged to align images with pre-defined prompt templates composed of open-world text descriptions for zero-shot learning [30].

For a DP pair, the disparity d for a pixel in the left view is defined by its location (i, j) and that of its corresponding pixel at (i, k) in the right view.

In the literature, this DP disparity [27] is defined as

$$d = j - k = A(r(z)) \cdot B(z), \mathbf{B}_L(i, j) = \mathbf{B}_R(i, k), \quad (1)$$

where $\mathbf{B}_L(i, j)$ and $\mathbf{B}_R(i, k)$ are the intensity of the two matching points, $r(\cdot)$ is a function of the radius of the circle of confusion, z is the point depth, and $A(r(z))$ and $B(z)$ are lens-dependent constants [2, 9, 26, 34].

Analysis. Leveraging the CLIP, zero-shot blur map estimation can be generally divided into three steps. *i) Image Encoding.* Considering the image encoder of the CLIP is trained with monocular RGB images, a function $\eta(\cdot, \cdot)$ is

required to create a new monocular image from the DP image pair $(\mathbf{B}_L, \mathbf{B}_R)$. To obtain the dense embedding for the blur map estimation task, we follow [41] to pop out the final pooling layer from the image encoder, yielding an embedding matrix $\mathbf{F} \in \mathbb{R}^{h_s \times w_s \times c}$, where h_s and w_s are the height and width for embedding of the newly created image, respectively. Mathematically, we have

$$\mathbf{F} = \text{ImageEncoder}(\eta(\mathbf{B}_L, \mathbf{B}_R)). \quad (2)$$

Each pixel of \mathbf{F} has encoded the local semantic information for a corresponding region in $\eta(\mathbf{B}_L, \mathbf{B}_R)$. *ii) Text Encoding.* A text encoder is used to embed the semantics of a prompt p describing the blur into the text embedding,

$$\mathbf{t} = \text{TextEncoder}(p). \quad (3)$$

iii) Blur map estimation. Let $\mathbf{F}(i, j)$ be a c -dimensional embedding at position (i, j) of \mathbf{F} , $i \in \{1, \dots, h_s\}$, and $j \in \{1, \dots, w_s\}$. The embedding $\mathbf{F}(i, j)$ and the blur semantic-based text embedding \mathbf{t} are expected to be similar if the pixel at (i, j) is blurred. Therefore, the blur mask $\mathbf{M}_{\mathbf{F}}$ of $\eta(\mathbf{B}_L, \mathbf{B}_R)$ is obtained by

$$\mathbf{M}_{\mathbf{F}}(i, j) = \sigma(\text{sim}(\mathbf{F}(i, j), \mathbf{t})). \quad (4)$$

Note, $\sigma(\cdot)$ normalizes the cosine similarity $\text{sim}(\cdot, \cdot)$ output to a probability space, to reflect the degree of blur at the position. Here, $\mathbf{M}_{\mathbf{F}}$ is sampled to have the same resolution as \mathbf{B}_L or \mathbf{B}_R , returning $\mathbf{M} \in \mathbb{R}^{h \times w \times 1}$.

Given the analysis above, our goal is to develop text prompts p , while adapting the function $\eta(\cdot, \cdot)$, to enable the use of semantic information from CLIP to generate the blur map \mathbf{M} in an unsupervised manner. In particular, we explore the blur-aware, DP-aware, and ensembling formats with detailed presented below.

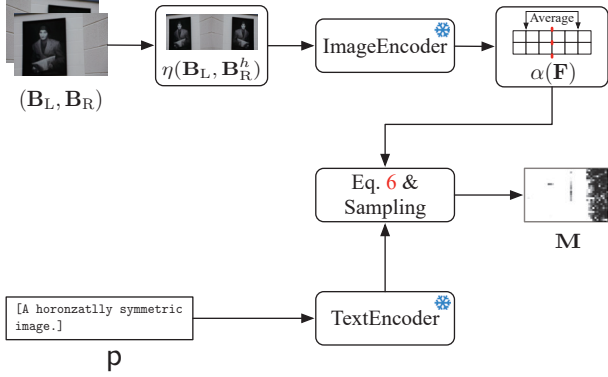


Figure 3. The pipeline of our DP-aware format for obtaining the blur map \mathbf{M} . Given a DP pair $(\mathbf{B}_L, \mathbf{B}_R)$, the function $\eta(\mathbf{B}_L, \mathbf{B}_R^h)$ concatenate the left view \mathbf{B}_L and with horizontally flipped right view \mathbf{B}_R^h for obtaining the \mathbf{F} from the CLIP image encoder. We then apply $\alpha(\mathbf{F})$ to averagely pool the feature map according to the transformation imposed on the input DP pair, resulting in a new feature embedding with the same aspect ratio of the input DP pair. Meanwhile, we set the prompts \mathbf{p} to describe the symmetry of $\eta(\mathbf{B}_L, \mathbf{B}_R^h)$, obtaining text embedding \mathbf{t} . We finally have the blur map \mathbf{M} by computing Eq. 6 and performing sampling.

Blur-aware Format. An intuitive approach is setting a prompt \mathbf{p} that describes the blur to query a blur map \mathbf{M} , e.g., $\mathbf{p} = [\text{A blurry image.}]$. The $\eta(\cdot, \cdot)$ is then defined to synthesize a center view of the DP pair by averaging the left view \mathbf{B}_L and right view \mathbf{B}_R , as the average of two sub-images is equivalent to the image captured by a regular sensor [3, 27]. However, the disparity information among the DP pair is neglected, resulting in a sub-optimal blur map. Refer to Fig. 4 for comparisons.

DP-aware Format. Based on the DP image formation process, the disparity is proportional to the degree of blur exhibited in the DP view [1]. We thus propose to guide CLIP to understand the disparity exhibited between \mathbf{B}_L and \mathbf{B}_R which eventually leads to better mask (\mathbf{M}) generation. To this end, we first formulate the disparity in the DP pair and then study the disparity-aware prompt and function $\eta(\cdot, \cdot)$ designing. Such design leads us to query degrees of displacement between the left and right view from the CLIP, without any fine-tuning, to obtain the blur map \mathbf{M} .

Note that the CLIP is the vision-text alignment model trained with monocular image input. To extract the displacement between left/right views of the DP pair, we explore several input-prompt formats: i) Concatenating the left and right views horizontally, and prompting their difference, e.g., [The left and right of the image are different.]; ii) Prompting the horizontal symmetry between the left and right views, e.g., [A horizontally symmetric image.] (Fig. 3).

Empirically, we achieve the best performance with the second prompt format. When describing the symmetry in




Figure 4. Example of blur map estimation (the darker, the sharper). (a) The left view of the DP image. (b) and (c) Blur maps obtained with prompts of the blur-aware and DP-aware formats. (d) The blur maps obtained with the ensembling-based format, i.e., all eight blur-aware and DP-aware prompts are used (please refer to Tab. 4 for details).

the prompt \mathbf{p} , we concatenate the left view \mathbf{B}_L with a horizontally flipped right view $\mathbf{B}_R^h = \text{Flip}(\mathbf{B}_R)$ along the width axial in the $\eta(\cdot, \cdot)$. We then obtained image features $\mathbf{F} = \text{ImageEncoder}(\eta(\mathbf{B}_L, \mathbf{B}_R^h))$. To transform \mathbf{F} to the same aspect ratio of \mathbf{B}_L and \mathbf{B}_R , we average the corresponding pixel-wise features based on the transformation in $\eta(\mathbf{B}_L, \mathbf{B}_R^h)$ and have $\alpha(\mathbf{F})$. For $i \in \{1, \dots, h_w\}$ and $j \in \{1, \dots, \lfloor \frac{1}{2}w_s \rfloor\}$, we define

$$\alpha(\mathbf{F})(i, j) = \frac{1}{2}\mathbf{F}(i, j) + \frac{1}{2}\mathbf{F}(i, w_s - j). \quad (5)$$

Back to the left-hand side of Eq. 4, we have

$$\mathbf{M}_{\mathbf{F}}(i, j) = 1 - \mathbf{M}_{\mathbf{F}}^s(i, j) = 1 - \sigma(\text{sim}(\alpha(\mathbf{F})(i, j), \mathbf{t}^s)). \quad (6)$$

Here, \mathbf{t}^s is the text embedding of the prompt describing symmetry. Note that the high similarity value, i.e. $\mathbf{M}_{\mathbf{F}}^s(i, j) = \sigma(\text{sim}(\alpha(\mathbf{F})(i, j), \mathbf{t}^s))$ indicates pixels in sharp regions. Therefore the blur map is derived as $1 - \mathbf{M}_{\mathbf{F}}^s(i, j)$ in Eq. 6. Similarly, the $\mathbf{M}_{\mathbf{F}}$ is sampled to the same resolution of the DP pair, returning \mathbf{M} . For this blur map estimation strategy, we conclude the lemma below,

Lemma 3.1. A pixel is sharp if and only if it is horizontally symmetric in the left view and the horizontally flipped right view, otherwise, the pixel is blurred.

Proof. Following Eq. (1), we have $\mathbf{B}_L(i, j)$ being sharp, given $\mathbf{B}_L(i, j) = \text{Flip}(\mathbf{B}_R)(i, w - j)$. In other words, we have $\eta(\mathbf{B}_L, \mathbf{B}_R^h)(i, j) = \eta(\mathbf{B}_L, \mathbf{B}_R^h)(i, w_s - j)$, demonstrating the horizontal symmetry. Reversely, the horizontal symmetry breaks once the disparity/blur occurs.

Ensembling format. Given the analysis that the degradations in the DP pair can be described by both the blurriness and symmetry, we thus propose to obtain the mask by ensembling the blur-aware and DP-aware formats jointly. Please refer to Tab. 4 for details of the 8 prompts designed from blur-aware and DP-aware formats.

3.2. Network

We train an encoder-decoder based deblurring network for restoring the DP pair $(\mathbf{B}_L, \mathbf{B}_R)$ into a sharp image $\hat{\mathbf{I}}$ with the derived blur map \mathbf{M} . Our network consists of a backbone network based on [5] and a BPA block to interleave with the backbone network layers. In particular, the BPA block uses the obtained blur map \mathbf{M} as deblur kernel prior to help the restoration of $(\mathbf{B}_L, \mathbf{B}_R)$ on the fly.

Self-attention. The self-attention (SA) mechanism is widely used in the image restoration tasks [18, 36, 39]. We first briefly review it. Let $\mathbf{X} \in \mathbb{R}^{h_x \times w_x \times c_x}$ be the intermediate feature embedding of a blur image from the backbone encoder, where h_x , w_x , c_x are height, width, and channel of the feature. The self-attention is input with query \mathbf{Q} , key \mathbf{K} , and value \mathbf{V} that are linearly projected from \mathbf{X} , and then calculate an attention map between \mathbf{Q} and \mathbf{K} for adaptively aggregating \mathbf{V} . This process is defined as follows,

$$\mathbf{Q} = \mathbf{X}\mathbf{W}_Q, \quad \mathbf{K} = \mathbf{X}\mathbf{W}_K, \quad \mathbf{V} = \mathbf{X}\mathbf{W}_V, \quad (7)$$

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}(\mathbf{Q}\mathbf{K}^\top / \beta) \mathbf{V}, \quad (8)$$

where $\mathbf{W}_Q, \mathbf{W}_K, \mathbf{W}_V \in \mathbb{R}^{h_x \times w_x \times c_x}$ are learnable weight matrices, β is a normalization factor that is typically set to the square root of the latent dimension, and $\text{Softmax}(\mathbf{Q}\mathbf{K}^\top / \beta) \in \mathbb{R}^{(h_x \times w_x) \times (h_x \times w_x)}$ is attention map.

BPA Block. The attention map can be considered as an adaptive deblur kernel prior for restoring the DP pair in the feature space. We are motivated to modulate the attention map by supplying the blur map \mathbf{M} , utilizing blur-related prior knowledge to additionally assign a $q \times q$ deblur kernel for each region in \mathbf{X} to aggregate a sharp feature embedding from its neighboring regions. We re-write the Eq. 8 as

$$\text{Attention}(\mathbf{Q}, \mathbf{K}, \mathbf{V}) = \text{Softmax}(\mathbf{Q}\mathbf{K}^\top / \beta + \mathbf{O}) \mathbf{V}, \quad (9)$$

$$\mathbf{O} = \text{Pad}(\text{FFN}_{\text{Conv}}(\sigma^{-1}(\mathbf{M}))), \quad (10)$$

and show that the restoration result is further improved in Fig. 5, where $\text{FFN}_{\text{Conv}}(\cdot)$ is a convolution-based feed forward layer that is used to learn a $q \times q$ deblur kernel for each region of \mathbf{X} from the unnormalized blur map $\sigma^{-1}(\mathbf{M})$, resulting in $\text{FFN}_{\text{Conv}}(\sigma^{-1}(\mathbf{M})) \in \mathbb{R}^{h_x \times w_x \times q^2}$. The $\text{Pad}(\cdot)$ function reshapes $\text{FFN}_{\text{Conv}}(\sigma^{-1}(\mathbf{M}))$ into $\mathbb{R}^{h_x \times w_x \times q \times q}$, and pads the reshaped $\text{FFN}_{\text{Conv}}(\sigma^{-1}(\mathbf{M}))$ to $\mathbb{R}^{(h_x \times w_x) \times (h_x \times w_x)}$, for shape alignment issue.

3.3. Network Training

We optimize our deblurring network with the widely used Charbonnier loss. Moreover, we propose a blur-aware loss and a blur weighting loss to penalize the blurred regions in the restoration $\hat{\mathbf{I}}$.

Charbonnier loss. We encourage the restoration to be close to the ground truth by using Charbonnier loss [12, 43].



Figure 5. Example of deblurring with the blur map projected deblur kernel \mathbf{O} . (a) and (b) are the left view of the DP pair and the ground truth sharp image. (c) and (d) are deblurred images with and without using \mathbf{O} . Zoom in for better quality.

We empirically set ε to 10^{-4} , and define

$$\mathcal{L}_{\text{char}} = \sqrt{\|\mathbf{I}(i, j) - \hat{\mathbf{I}}(i, j)\|^2 + \varepsilon^2}. \quad (11)$$

Blur Weighting Loss. We impose an adaptive penalty on the blurred regions of the DP pair $(\mathbf{B}_L, \mathbf{B}_R)$, considering the degree of blur measured by CLIP. It encourages the network to focus on the restoration of heavily blurred regions. We first normalize the deviations between $\hat{\mathbf{I}}$ and \mathbf{I} by the deviations between $(\mathbf{B}_L, \mathbf{B}_R)$ and \mathbf{I} , and then re-use the blur map to select blurred regions residing in the DP pair.

$$\mathcal{L}_{\text{bwl}} = \frac{\|\mathbf{I}(i, j) - \hat{\mathbf{I}}(i, j)\|^2}{\|\hat{\mathbf{I}}(i, j) - \frac{1}{2}(\mathbf{B}_L(i, j) + \mathbf{B}_R(i, j))\|^2} \cdot \mathbf{M}(i, j), \quad (12)$$

where we use the blur map \mathbf{M} (with soft logits) to reflect the blur levels detected by CLIP.

Blur-aware Loss. We encourage $\hat{\mathbf{I}}$ to be sharp by leveraging the versatile knowledge contained in the CLIP. We generate the blur map $\hat{\mathbf{M}}$ by the blur-aware format for the restoration, and penalize the blur exhibited in $\hat{\mathbf{I}}$. Similar to \mathcal{L}_{bwl} , we use the unnormalized version of the blur map,

$$\mathcal{L}_{\text{bal}} = \hat{\mathbf{M}}(i, j). \quad (13)$$

This encourages $\hat{\mathbf{M}}$ to be zero, *i.e.*, no blur exists in the restoration.

Total Loss. Our network is optimized by minimizing the following loss function

$$\mathcal{L}_{\text{total}} = \frac{1}{|\Omega|} \sum_{(i, j) \in \Omega} (\mathcal{L}_{\text{char}} + \lambda_1 \mathcal{L}_{\text{bal}} + \lambda_2 \mathcal{L}_{\text{bwl}}), \quad (14)$$

where $\frac{1}{|\Omega|}$ is a normalization factor determined by the image domain size $|\Omega|$, and λ_1 and λ_2 are hyperparameters to balance losses.

4. Experiments and Analysis

Implementation Details. We implement our model under the PyTorch framework. We use the pre-trained and ViT-B/32 [8] architecture-based CLIP [30] for our blur map estimation strategy. We train the deblurring network in two stages for 6×10^6 steps in total with the Adam optimizer.

Table 1. Quantitative comparisons on the DPD-blur dataset [1] composed of 37 indoor and 39 outdoor scenes. We respectively color the best and the second-best performance in red and blue.

Method	Outdoor Scenes				Indoor Scenes				Average				
	PSNR \uparrow	SSIM \uparrow	MAE \downarrow	RMSE $\text{_{rel}}\downarrow$	PSNR \uparrow	SSIM \uparrow	MAE \downarrow	RMSE $\text{_{rel}}\downarrow$	PSNR \uparrow	SSIM \uparrow	MAE \downarrow	RMSE $\text{_{rel}}\downarrow$	LPIPS \downarrow
EBDB	25.77	0.772	0.040	0.051	21.25	0.599	0.058	0.086	23.45	0.683	0.049	0.067	0.336
DMENet	25.50	0.788	0.038	0.053	21.43	0.644	0.063	0.084	23.41	0.714	0.051	0.067	0.349
DPDNet	27.48	0.849	0.029	0.042	22.90	0.726	0.052	0.071	25.13	0.786	0.041	0.055	0.223
RDPD	28.10	0.843	0.027	0.039	22.82	0.704	0.053	0.072	25.39	0.772	0.040	0.053	0.255
DDDNet	27.57	0.833	0.030	0.041	23.28	0.708	0.050	0.068	25.36	0.768	0.041	0.054	0.197
IFAN	28.66	0.868	0.025	0.037	23.46	0.743	0.049	0.067	25.99	0.804	0.037	0.050	0.207
BAMBNet	28.60	0.872	0.026	0.037	24.30	0.772	0.045	0.060	26.40	0.821	0.036	0.047	0.141
DeepRFT	28.48	0.870	0.025	0.038	23.09	0.736	0.051	0.070	25.71	0.801	0.037	0.051	0.217
Restormer	29.48	0.895	0.023	0.034	23.97	0.773	0.047	0.063	26.66	0.833	0.035	0.046	0.155
K3DN	29.87	0.890	0.022	0.032	24.37	0.780	0.046	0.060	27.06	0.835	0.034	0.044	0.187
Ours (Small)	29.55	0.893	0.021	0.033	24.38	0.773	0.043	0.065	26.91	0.831	0.032	0.045	0.185
Ours (Large)	29.88	0.896	0.020	0.032	24.73	0.786	0.037	0.059	27.24	0.840	0.028	0.043	0.168

Table 2. Quantitative comparisons on DDD-syn dataset.

Method	PSNR \uparrow	SSIM \uparrow	MAE \downarrow	RMSE $\text{_{rel}}\downarrow$
EBDB	26.48	0.683	0.036	0.047
DMENet	30.14	0.939	0.017	0.031
DPDNet	31.45	0.926	0.016	0.027
RDPD	32.42	0.912	0.016	0.024
DDDNet	33.21	0.956	0.010	0.022
IFAN	34.18	0.929	0.012	0.020
BAMBNet	35.90	0.954	0.012	0.016
DeepRFT	36.52	0.952	0.010	0.015
Restormer	36.44	0.957	0.010	0.015
K3DN	38.23	0.965	0.008	0.012
Ours (Small)	37.27	0.959	0.010	0.014
Ours (Large)	38.62	0.971	0.008	0.012

Table 3. Quantitative comparisons on RDPD dataset.

Method	PSNR \uparrow	SSIM \uparrow	MAE \downarrow	RMSE $\text{_{rel}}\downarrow$
EBDB	24.23	0.637	0.031	0.054
DMENet	25.17	0.731	0.036	0.061
DPDNet	29.84	0.828	0.032	0.055
RDPD	31.09	0.861	0.025	0.028
DDDNet	30.14	0.840	0.016	0.031
IFAN	31.12	0.865	0.023	0.028
BAMBNet	31.78	0.867	0.016	0.026
DeepRFT	31.71	0.879	0.016	0.024
Restormer	32.27	0.871	0.016	0.024
K3DN	32.84	0.905	0.014	0.023
Ours (Small)	32.98	0.895	0.014	0.022
Ours (Large)	33.44	0.907	0.013	0.021

In each stage, the learning rate is decayed from 2×10^{-4} to 10^{-6} by a cosine annealing strategy. In the first stage, we train the network for 5×10^6 steps, and only use the Charbonnier loss $\mathcal{L}_{\text{char}}$. In the second stage, the network is finetuned for 1×10^6 steps with $\mathcal{L}_{\text{total}}$, where we have $\lambda_1 = 10^{-1}$ and $\lambda_2 = 2 \times 10^{-1}$. Code, results, and more details are given in our project page.

Datasets. Our method is evaluated on two real-world DP defocus deblurring datasets, DPD-blur dataset [1] and DPD-disp dataset [27]. We follow [1] on the training and testing splits. For the DPD-disp dataset, we use it for testing only, re-using the checkpoints trained from the DPD-blur dataset as they do not provide ground-truth (GT) data or with limited dataset size. Moreover, two synthetic datasets, DDD-syn [26] and RDPD [2], are also used for evaluations.

Evaluation Metrics. We use peak signal-to-noise ratio (PSNR), structural similarity (SSIM), mean absolute error (MAE), and root mean squared relative error (RMSE $\text{_{rel}}$) as our evaluation metrics.

Baseline Methods. We compare with SOTA DP defocus deblurring methods, *i.e.*, EBDB [13], DMENet [15],

DPDNet [1], RDPD [2], DDDNet [26], IFAN [16], BAMBNet [20], DeepRFT [22], Restormer [39], and K3DN [38].

4.1. Comparisons with State-of-the-Art

Quantitative Comparison. We compare with the SOTA methods on the DPD-blur, DDD-syn, and RDPD datasets in Tab. 1, Tab. 2, and Tab. 3, respectively. Our observations are as follows: i) With our model (Small), we achieve similar restoration results with Restormer and K3DN, the third-best and second-best method, in all datasets. ii) When scaling to the large model, we consistently outperform all methods on all evaluation metrics. For example, our method, K3DN, and Restormer respectively achieve 27.24/27.06/26.66, 38.62/38.23/36.44, and 33.44/32.84/32.27 PSNR (dB) in the DPD-blur, DDD-syn, and RDPD datasets. Meanwhile, on a single RTX 4090 GPU, we remain to have a lower Infer. time (including CLIP) which is 0.042s and 0.054s for our small and large models, while K3DN/Restormer takes 0.198s/0.255s.

Qualitative Comparison. We compare with state-of-the-art methods on the DPD-blur dataset (Fig. 6). As shown, our

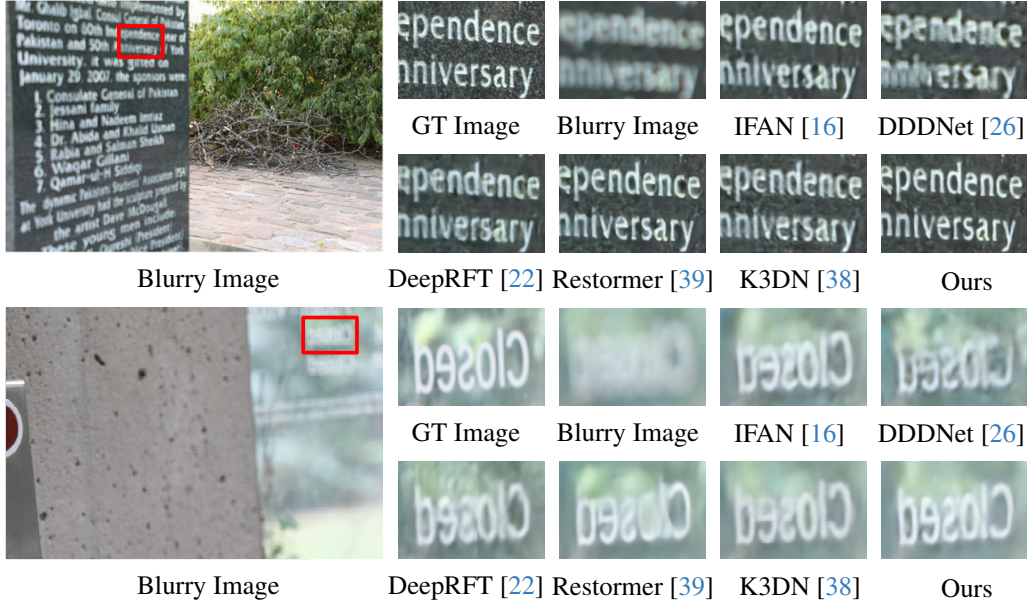


Figure 6. *Qualitative comparison on the DPD-blur dataset [1]. We present the blurry image (B_L) in the first column, and the ground truth (GT) image of regions residing in the red bounding box is shown in the second column. The cropped blurred image of the red bounding box region is in the third column.*

method performs sharper and clearer restorations of edges and texts. To further demonstrate the benefits of using CLIP to generate the blur mask unsupervisedly in our approach, we provide an example in Fig. 7, which preserves sharp regions (e.g., text ‘University’) in the input blur image effectively compared to the second-best baseline.




Figure 7. *An example of the restoration results on the sharp region of the input DP image (left view).*

4.2. Ablation Study

We study model architecture, blur map estimation, and losses in this section. All experiments are performed in the DPD-blur dataset.

Ablation of Model Architecture. We study the effective usage of blur map as prior knowledge for DP defocus deblurring (Tab. 6). To show the benefits of using blur maps, we consider four baselines: i) NAFNet [5], the baseline model; ii) Concat, concatenating the blur map with the DP pair before feeding to the model; iii) SA, using the self-attention layer [8] in Eq. 8 to the model only; iv) Concat+SA, combines the Concat and SA settings; v) BPA block, using kernel priors from the blur map to assist deblurring. Our model consistently outperforms all baselines, which indicates the effectiveness of using the blur map.

Ablation of Losses. We investigate the proposed blur weighting loss \mathcal{L}_{bwl} and blur-aware loss \mathcal{L}_{bal} (Tab. 7) with a baseline as the \mathcal{L}_{char} setting. We first use \mathcal{L}_{bwl} and \mathcal{L}_{bal}

to our baseline, NAFNet (i.e., not using our BPA block), that only uses \mathcal{L}_{char} , having 0.32 PSNR (dB) improvements. Secondly, we train our model (backbone + BPA block) with Charbonnier loss, resulting in 26.71/0.825 higher PSNR (dB)/SSIM. By using the \mathcal{L}_{bwl} or \mathcal{L}_{bal} , we respectively achieve 0.05/0.003 or 0.10/0.004 higher PSNR (dB)/SSIM. Combining all losses, we have the best restoration performance at 26.91/0.831 for PSNR (dB)/SSIM. Moreover, we verify the generalization ability of our \mathcal{L}_{bwl} and \mathcal{L}_{bal} losses on state-of-the-art networks, improving PSNR of Restormer and DeepRFT by 0.25 and 0.41 PSNR (dB).

Ablation of Blur Map Estimation. We ablate with different prompts used to generate the blur map. The blur-aware and DP-aware prompts are presented in Tab. 4, and we show their impact on model performance with Tab. 5. We then ensemble the groups of blur-aware prompt and disparity-aware prompt individually (Ens.¹), leading to 26.80dB/0.829 and 26.84/0.830 PSNR (dB)/SSIM. By ensembling all prompts (Ens.²), we have 26.91/0.831 PSNR (dB)/SSIM, consistently outperforming other approaches of generating blur masks.

4.3. Discussion and Limitation

Discussion. The reader may wonder about the performance of our unsupervised blur map estimation strategy. Therefore, we perform two qualitative evaluations on the DPD-disp dataset, where the GT disparity is included. We compare with BAMBNet [20] and DPD-disp [27].

First, we consider blur map estimation. To obtain the GT blur map, the pixel with disparity d that satisfies $|d| \leq 1$

Table 4. Prompts used for blur map estimation.

	Blur-aware Prompt	DP-aware Prompt
p ₁	[A blurry image.]	[A symmetrical image.]
p ₂	[An obscured image.]	[A horizontally symmetrical image.]
p ₃	[An out of focus image.]	[A left-right symmetrical image.]
p ₄	[An image lacking sharpness.]	[A symmetrical image at the pixel level.]

Table 5. Restoration performance of our methods by using blur-aware or/and DP-aware prompts for blur map estimation. *Ens.*¹ denotes ensembling the group of blue-aware prompts and DP-aware prompts individually, and *Ens.*² ensembles all prompts.

	Blur-aware Prompt				DP-aware Prompt			
	PSNR _↑	SSIM _↑	MAE _↓	RMSE _{rel} ↓	PSNR _↑	SSIM _↑	MAE _↓	RMSE _{rel} ↓
p ₁	26.83	0.829	0.033	0.045	26.82	0.830	0.033	0.046
p ₂	26.81	0.829	0.033	0.046	26.83	0.830	0.032	0.046
p ₃	26.75	0.828	0.034	0.046	26.82	0.830	0.033	0.046
p ₄	26.77	0.829	0.033	0.046	26.85	0.830	0.032	0.045
Ens. ¹	26.80	0.829	0.033	0.046	26.84	0.830	0.032	0.046
Ens. ²	26.91	0.831	0.032	0.045	26.91	0.831	0.032	0.045

Table 6. Ablation of model components.

Setting	PSNR _↑	SSIM _↑	MAE _↓	RMSE _{rel} ↓
i) NAFNet	26.21	0.817	0.038	0.049
ii) Concat	26.76	0.826	0.033	0.046
iii) SA	26.59	0.824	0.036	0.047
iv) Concat + SA	26.80	0.829	0.032	0.046
v) BPA block	26.91	0.831	0.032	0.045

Table 7. The effectiveness of losses.

BPA block	Loss	PSNR _↑	SSIM _↑	MAE _↓	RMSE _{rel} ↓
✗	$\mathcal{L}_{\text{char}}$	26.21	0.817	0.038	0.049
✗	$\mathcal{L}_{\text{total}}$	26.53	0.822	0.036	0.047
✓	$\mathcal{L}_{\text{char}}$	26.71	0.825	0.034	0.046
✓	$\mathcal{L}_{\text{char}} + \mathcal{L}_{\text{bwl}}$	26.76	0.828	0.033	0.046
✓	$\mathcal{L}_{\text{char}} + \mathcal{L}_{\text{bal}}$	26.81	0.829	0.032	0.046
✓	$\mathcal{L}_{\text{total}}$	26.91	0.831	0.032	0.045

is used as the sharp pixel, reversely, it is a burred pixel. Our method, BAMBNet [20] and DPD-disp [27] respectively have 89.8%/70.2%/88.3% accuracy (the higher, the better), showing the effectiveness of our method.

Second, we explore disparity estimation. The logits of our unnormalized blurred map are potentially linearly correlated with the ground-truth disparity. Therefore, we compare with the state-of-the-art DP disparity estimation method, DPD-disp [27]. Following [9, 27], we use AI(1), AI(2), and $1 - |\rho_s|$ as the evaluation metrics (the lower, the better). The AI(1), AI(2), and $1 - |\rho_s|$ of our method, BAMBNet, and DPD-disp are 0.063/0.106/0.069, 0.130/0.184/0.096, 0.306/0.772/0.0159. Though our blur map is estimated with a single forward pass, we still achieve a reasonable performance in the disparity estimation. This potentially explains the high performance of our DP defo-

cus deblurring network, as obtaining the disparity prior benefits the defocus task.

Limitation. As the image encoder of CLIP is trained under a classification-alike pre-training task, it will focus more on local details that help identify its scene class. Textureless regions would produce non-significant semantic reactions, *e.g.*, large areas of wall, ceiling, or floor, and reduce the quality of the estimated blur map. It is possible to use a vision language framework that drives more sensitivity to regions with less texture.

5. Conclusion and Broader Impact

In this paper, we study end-to-end defocus deblurring of DP images, leveraging the blur-related prior knowledge from the CLIP. We first estimate a blur map with an ensemble of blur-aware and DP-aware strategies, and then use the estimated blur map as the deblurring kernel prior to restoring the DP image. We also propose a blur-aware and blur-weighting loss to regularize the restorations of the DP images during training, by distilling the knowledge of blur from CLIP. In extensive experiments, our method outperforms past works by a large margin in quantitative and qualitative restoration performance.

Broader Impact. The proposed DP-aware strategy is promising to apply and extend CLIP to diverse zero-shot stereo vision tasks. We hope it will motivate future work.

Acknowledgment. This work was supported in part by the Beijing Institute of Technology Research Fund Program for Young Scholars, BIT Special-Zone, and National Natural Science Foundation of China 62302045, and Miaomiao Liu was supported by the ARC fellowship and Discovery Project grant (DE180100628, DP200102274).

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, 4, 6, 7
- [2] Abdullah Abuolaim, Mauricio Delbracio, Damien Kelly, Michael S Brown, and Peyman Milanfar. Learning to reduce defocus blur by realistically modeling dual-pixel data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2289–2298, 2021. 1, 3, 6
- [3] Abdullah Abuolaim, Mahmoud Afifi, and Michael S Brown. Improving single-image defocus deblurring: How dual-pixel images help through multi-task learning. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 1231–1239, 2022. 1, 2, 4
- [4] Dylan Auty and Krystian Mikolajczyk. Learning to prompt clip for monocular depth estimation: Exploring the limits of human language. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2039–2047, 2023. 2
- [5] Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII*, pages 17–33. Springer, 2022. 5, 7
- [6] Myungsub Choi, Hana Lee, and Hyong-euk Lee. Exploring positional characteristics of dual-pixel data for camera autofocus. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13158–13168, 2023. 1
- [7] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. IEEE, 2009. 2
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2021. 2, 5, 7
- [9] Rahul Garg, Neal Wadhwa, Sameer Ansari, and Jonathan T Barron. Learning single camera depth estimation using dual-pixels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 7628–7637, 2019. 1, 3, 8
- [10] Xueting Hu, Ce Zhang, Yi Zhang, Bowen Hai, Ke Yu, and Zhihai He. Learning to adapt clip for few-shot monocular depth estimation. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 5594–5603, 2024. 2
- [11] Tianyu Huang, Bowen Dong, Yunhan Yang, Xiaoshui Huang, Rynson WH Lau, Wanli Ouyang, and Wangmeng Zuo. Clip2point: Transfer clip to point cloud classification with image-depth pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 22157–22167, 2023. 2
- [12] Kui Jiang, Zhongyuan Wang, Peng Yi, Chen Chen, Baojin Huang, Yimin Luo, Jiayi Ma, and Junjun Jiang. Multi-scale progressive fusion network for single image deraining. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8346–8355, 2020. 5
- [13] 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
- [14] Donggun Kim, Hyeonjoong Jang, Inchul Kim, and Min H Kim. Spatio-focal bidirectional disparity estimation from a dual-pixel image. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5023–5032, 2023. 1
- [15] Junyong Lee, Sungkil Lee, Sunghyun Cho, and Seungyong Lee. Deep defocus map estimation using domain adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 12222–12230, 2019. 6
- [16] Junyong Lee, Hyeongseok Son, Jaesung Rim, Sunghyun Cho, and Seungyong Lee. Iterative filter adaptive network for single image defocus deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2034–2042, 2021. 1, 2, 6, 7
- [17] Yu Li, Yaling Yi, Dongwei Ren, Qince Li, and Wangmeng Zuo. Learning dual-pixel alignment for defocus deblurring. *arXiv preprint arXiv:2204.12105*, 2022. 2
- [18] Yawei Li, Yuchen Fan, Xiaoyu Xiang, Denis Demandolx, Rakesh Ranjan, Radu Timofte, and Luc Van Gool. Efficient and explicit modelling of image hierarchies for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18278–18289, 2023. 5
- [19] Yu Li, Dongwei Ren, Xinya Shu, and Wangmeng Zuo. Learning single image defocus deblurring with misaligned training pairs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, pages 1495–1503, 2023. 1
- [20] Pengwei Liang, Junjun Jiang, Xianming Liu, and Jiayi Ma. Bambnet: A blur-aware multi-branch network for dual-pixel defocus deblurring. *IEEE/CAA Journal of Automatica Sinica*, 9(5):878–892, 2022. 1, 2, 6, 7, 8
- [21] Zhexin Liang, Chongyi Li, Shangchen Zhou, Ruicheng Feng, and Chen Change Loy. Iterative prompt learning for unsupervised backlit image enhancement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 8094–8103, 2023. 2
- [22] Xintian Mao, Yiming Liu, Fengze Liu, Qingli Li, Wei Shen, and Yan Wang. Intriguing findings of frequency selection for image deblurring. In *Proceedings of the 37th AAAI Conference on Artificial Intelligence*, 2023. 6, 7
- [23] Sagi Monin, Sagi Katz, and Georgios Evangelidis. Continuous cost aggregation for dual-pixel disparity extraction. *arXiv preprint arXiv:2306.07921*, 2023. 1
- [24] 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. 1
- [25] 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 Con-*

- ference on Computer Vision and Pattern Recognition, pages 6034–6043, 2019. [2](#)
- [26] Liyuan Pan, Shah Chowdhury, Richard Hartley, Miaomiao Liu, Hongguang Zhang, and Hongdong Li. Dual pixel exploration: Simultaneous depth estimation and image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4340–4349, 2021. [1](#), [2](#), [3](#), [6](#), [7](#)
- [27] Abhijith Punnappurath, Abdullah Abuolaim, Mahmoud Afifi, and Michael S Brown. Modeling defocus-disparity in dual-pixel sensors. In *2020 IEEE International Conference on Computational Photography (ICCP)*, pages 1–12. IEEE, 2020. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#), [8](#)
- [28] Yuhui Quan, Zicong Wu, and Hui Ji. Neumann network with recursive kernels for single image defocus deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5754–5763, 2023. [1](#)
- [29] Yuhui Quan, Xin Yao, and Hui Ji. Single image defocus deblurring via implicit neural inverse kernels. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 12600–12610, 2023. [1](#)
- [30] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR, 2021. [1](#), [2](#), [3](#), [5](#)
- [31] Yongming Rao, Wenliang Zhao, Guangyi Chen, Yansong Tang, Zheng Zhu, Guan Huang, Jie Zhou, and Jiwen Lu. Densclip: Language-guided dense prediction with context-aware prompting. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 18082–18091, 2022. [1](#), [2](#)
- [32] Cheng Shi and Sibe Yang. Edadet: Open-vocabulary object detection using early dense alignment. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 15724–15734, 2023. [1](#)
- [33] Hyeonseok Son, Junyong Lee, Sunghyun Cho, and Seungyong Lee. Single image defocus deblurring using kernel-sharing parallel atrous convolutions. In *2021 IEEE/CVF International Conference on Computer Vision, ICCV 2021, Montreal, QC, Canada, October 10-17, 2021*, pages 2622–2630. IEEE, 2021. [1](#), [2](#)
- [34] Shumian Xin, Neal Wadhwa, Tianfan Xue, Jonathan T Barron, Pratul P Srinivasan, Jiawen Chen, Ioannis Gkioulekas, and Rahul Garg. Defocus map estimation and deblurring from a single dual-pixel image. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2228–2238, 2021. [1](#), [2](#), [3](#)
- [35] Hao Yang, Liyuan Pan, Yan Yang, and Wei Liang. Language-driven all-in-one adverse weather removal. *arXiv preprint arXiv:2312.01381*, 2023. [2](#)
- [36] Shaoliang Yang, Dongming Zhou, Jinde Cao, and Yanbu Guo. Lightingnet: An integrated learning method for low-light image enhancement. *IEEE Transactions on Computational Imaging*, 9:29–42, 2023. [5](#)
- [37] Yan Yang, Liyuan Pan, and Liu Liu. Event camera data pre-training. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10699–10709, 2023. [1](#)
- [38] Yan Yang, Liyuan Pan, Liu Liu, and Miaomiao Liu. K3dn: Disparity-aware kernel estimation for dual-pixel defocus deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 13263–13272, 2023. [1](#), [2](#), [6](#), [7](#)
- [39] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5728–5739, 2022. [2](#), [5](#), [6](#), [7](#)
- [40] Renrui Zhang, Ziyu Guo, Wei Zhang, Kunchang Li, Xupeng Miao, Bin Cui, Yu Qiao, Peng Gao, and Hongsheng Li. Pointclip: Point cloud understanding by clip. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8552–8562, 2022. [1](#), [2](#)
- [41] Renrui Zhang, Ziyao Zeng, Ziyu Guo, and Yafeng Li. Can language understand depth? In *Proceedings of the 30th ACM International Conference on Multimedia*, pages 6868–6874, 2022. [2](#), [3](#)
- [42] 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 *European Conference on Computer Vision*, pages 582–598, 2020. [1](#)
- [43] Qian Zhao, Hao Yang, Dongming Zhou, and Jinde Cao. Rethinking image deblurring via cnn-transformer multiscale hybrid architecture. *IEEE Transactions on Instrumentation and Measurement*, 72:1–15, 2022. [5](#)
- [44] Ziqin Zhou, Yinjie Lei, Bowen Zhang, Lingqiao Liu, and Yifan Liu. Zegclip: Towards adapting clip for zero-shot semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11175–11185, 2023. [2](#)