

A Differentiable Two-stage Alignment Scheme for Burst Image Reconstruction with Large Shift

Shi Guo¹ Xi Yang¹ Jianqi Ma¹ Gaofeng Ren² Lei Zhang¹

¹The Hong Kong Polytechnic University; ²DAMO Academy, Alibaba Group

{shigu.guo, xxxxi.yang, jianqi.ma}@connect.polyu.hk,
gaof.ren@gmail.com, cslzhang@comp.polyu.edu.hk

Abstract

Denoising and demosaicking are two essential steps to reconstruct a clean full-color image from the raw data. Recently, joint denoising and demosaicking (JDD) for burst images, namely JDD-B, has attracted much attention by using multiple raw images captured in a short time to reconstruct a single high-quality image. One key challenge of JDD-B lies in the robust alignment of image frames. State-of-the-art alignment methods in feature domain cannot effectively utilize the temporal information of burst images, where large shifts commonly exist due to camera and object motion. In addition, the higher resolution (e.g., 4K) of modern imaging devices results in larger displacement between frames. To address these challenges, we design a differentiable two-stage alignment scheme sequentially in patch and pixel level for effective JDD-B. The input burst images are firstly aligned in the patch level by using a differentiable progressive block matching method, which can estimate the offset between distant frames with small computational cost. Then we perform implicit pixel-wise alignment in full-resolution feature domain to refine the alignment results. The two stages are jointly trained in an end-to-end manner. Extensive experiments demonstrate the significant improvement of our method over existing JDD-B methods. Codes are available at <https://github.com/GuoShi28/2StageAlign>.

1. Introduction

Color demosaicking and denoising are two essential steps in digital camera imaging pipeline to reconstruct a high quality full-color image from the sensor raw data. Color demosaicking [9, 20, 29, 45, 46] recovers the missing color components from the color-filter-array (CFA) data collected by the single-chip CCD/CMOS sensor, while denoising [21, 31, 48, 49] removes the noise in image data caused by the photon arrival statistics and the imprecision

in readout circuitry. Since the two tasks are actually correlated and can be performed jointly, many joint denoising and demosaicking (JDD) algorithms have been developed [5, 11, 14, 15, 23, 28], which are however focused on single image restoration. With the prevalent use of smartphone cameras [1], it becomes crucial to restore images from data with low signal-to-noise ratio due to the small sensor and lens of smartphone cameras. To this end, performing JDD with burst images (JDD-B) has become increasingly popular and important in recent years [13].

Burst image processing refers to shooting a sequence of low quality frames in a short time and computationally fusing them to produce a higher-quality photograph [2, 41]. Compared with single image restoration, the key challenge of burst image restoration lies in the compensation for shift between frames. Previous studies often perform frame alignment by estimating optical flow [8, 44] and applying spatially variant kernels [32, 33, 42, 43] in the image domain. However, affected by noise, accurately estimating optical flow and kernels is difficult. Recently, implicit alignment in the feature domain has achieved state-of-the-art performance on video super-resolution [4, 40], video denoising [47], as well as JDD-B [13]. However, it is found that feature alignment is not very effective when handling image sequences with large shift. For example, for the pyramid feature alignment module in [40], the receptive field of offset estimation is about 28 (3 conv layers with 3×3 kernels on 1/4 scale), *i.e.*, 14 pixels along one direction. Such a search range is too small for burst images with large shift. Whereas, large shifts commonly exist in videos due to the camera and object motion. On the other hand, modern image/video recorders can easily capture videos of 4K (4096×2160) or UHD (3840×2160) resolution. As a result, the pixel displacements between frames further increase. Even some small motion of foreground objects in 4K video can cause a large value of shift.

However, precise pixel-wise alignment on full size images with large motion is very difficult and expensive. With a naive implementation, for each pixel, we assume the off-

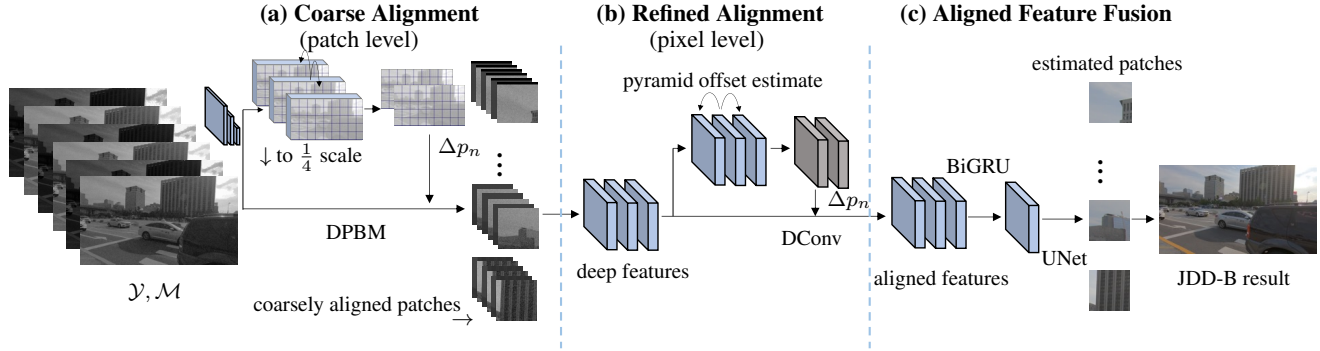


Figure 1. Illustration of our network with differentiable two-stage alignment for JDD-B.

set estimation with receptive field $D \times D$ costs $F \times D^2$ multiply-adds, where F is determined by different network structures. For input with size $H \times W$, the offset estimation costs $F \times D^2 \times H \times W$ multiply-adds. One straightforward solution to handle large shift is to increase D by using 3D conv layers [27] or calculating all-range correlation volume [24]. However, these solutions will significantly increase the computational cost and they are not efficient for images with large size.

To address these problems, we design a differentiable two-stage alignment framework, which divides the difficult large shift alignment problem into two relatively easier alignment sub-problems, *i.e.*, coarse alignment (CA) and refined alignment (RA). The CA module aims to compensate large shift roughly using small calculating resources. Instead of using pixel-wise alignment, the CA module performs alignment in patch level with $F \times D^2 \times (H/k) \times (W/k)$ multiply-adds, where k is the patch size. Then the RA module is used to pixel-wise align frames based on the results of CA module with smaller receptive field D_s . Such two-stage framework uses $F \times (D^2/k^2 + D_s^2) \times H \times W$ multiply-adds in total, which is significantly smaller than directly aligning images using one-stage module, especially when D is large. Specifically, we utilize block matching (BM) based method as CA module in image domain.

To overcome the non-differentiability caused by BM and reduce the computational cost, we propose a differentiable progressive block matching (DPBM) method. We further propose a corresponding loss function for DPBM to stabilize the training process. For the RA module, we perform refined pixel-wise alignment implicitly in the feature domain using deformable convolution for accurate image restoration. The two stages of alignments are complementary and they can enhance each other when jointly used in our learning framework. Compared with state-of-the-art (SOTA) method GCP-Net, our two-stage framework has fewer learnable parameters and similar running time, but achieves great improvement on images with large shift. The major contributions of this work are shown as follows:

- To efficiently handle images with large shift, we propose a differentiable two-stage alignment framework, which divides the difficult large shift alignment problem into two easier sub-problems.
- We propose a differentiable progressive block matching method in CA module to reduce computational cost and ensure that our two-stage framework can be end-to-end trained.

Our experiments on synthetic and real-world burst image datasets clearly show that our two-stage alignment method achieves impressive improvement over existing methods.

2. Related Work

2.1. Joint Denoising and Demosaicking

Denoising and demosaicking are two fundamental and correlated tasks in camera image signal processing (ISP) pipeline with a single chip of CMOS/CCD sensor [9, 20, 29, 45, 46]. Considering that performing denoising and demosaicking separately may accumulate errors for image restoration, JDD has been widely studied in recent years [7, 11, 13, 15, 23, 28, 36].

Gharbi *et al.* [11] showed that using more challenging patches for training can reduce the moiré artifacts in JDD. Then, more complex methods have been to obtain better performance, *e.g.*, two-stage network [36], auto-encoder architecture [15] and iterative structure [23]. The mosaic-to-mosaic framework [7] was proposed to improve the demosaicking performance on real-world images by finetuning the network with burst images. Since in the Bayer pattern of camera raw images, the green channel has twice the sampling rate of red/blue channels and has higher signal-to-noise ratio, the green channels were utilized to guide the upsampling process [28], the feature extraction [13] and offset estimation between frames [13]. Previous methods mostly perform JDD on single CFA raw images and achieve limited performance on real-world CFA image with high noise level. Even a method has been proposed to use burst raw image for the JDD task [13], its alignment module has

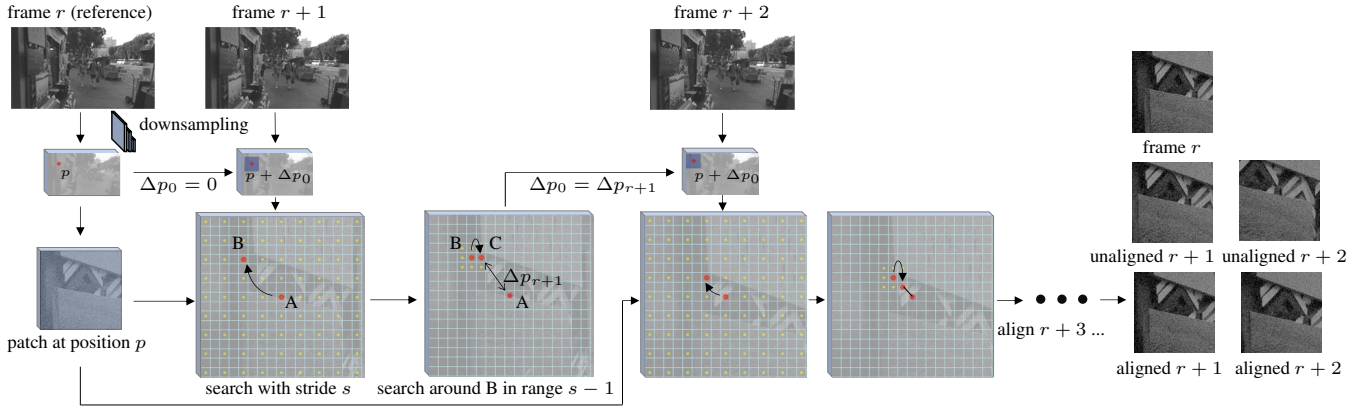


Figure 2. Illustration of the coarse alignment process with differentiable progressive block matching (DPBM). Without loss of generality, we show the DPBM of a patch centered at position p . “A” is the center of search region centered at $p + \Delta p_0$, while Δp_0 is 0 for target frame $r + 1$ and Δp_{r+n} for target frame $r + n + 1$. To reduce computational cost, we first perform matching with stride s using proposed differentiable block matching, resulting in the matched center position “B”. Then we perform searching with stride 1 in the neighborhood of “B”, resulting in a more accurate matched position “C”. Δp_{r+1} is the displacement between “C” and “A” for frame $r + 1$.

small receptive field and cannot obtain visual-pleasing results on images with large shift. In this work, we propose a new two-stage framework, which obtains great improvement over [13] with similar running time.

2.2. Multi-frame and Burst Image Restoration

Multi-frame image restoration has been widely studied in literature [4, 8, 18, 26, 32, 33, 38, 40, 42–44, 47], aiming to reproduce a photo with better quality than that estimated from a single image. The major challenge of multi-frame image restoration lies in how to compensate for the motion between frames. Some popular solutions employ optical flow [8, 44] and spatially varying kernel estimation [32, 33, 42, 43]. However, the performance of these methods is highly affected by large motion and severe noise.

Recently, performing implicit frame alignment in feature domain has achieved state-of-the-art performance on video super-resolution [4, 26, 38, 40], video denoising [47] and JDD-B [13]. In such methods, the offset [38, 40] or optical flow [4] is estimated from the deep features extracted from neighboring frames, and the deformable convolution [6] or wrap operator is used to compensate for shifts in feature domain. In [13, 40, 47], pyramidal processing is employed in the offset estimation to deal with complex motions. However, we found that, due to the limited receptive field, implicit alignment in feature domain can only achieve limited performance for sequences with large shift. To solve this problem, we design a differentiable two-stage alignment scheme and demonstrate its effectiveness for JDD-B.

3. Methodology

3.1. Motivation and Network Structure

The purpose of JDD-B task is to reconstruct clean RGB image x from a burst of noisy CFA images $\mathcal{Y} = \{y_t\}_{t=1}^N$

and their corresponding noise maps $\mathcal{M} = \{m_t\}_{t=1}^N$. However, previous SOTA multi-frame methods [13, 40, 47] have limited ability to address large pixel displacement by the small receptive fields. Thus, we propose a differentiable two-stage alignment framework to increase the receptive field of alignment without increasing the amount of calculation. Our framework is illustrated in Fig. 1, which divides the difficult large shift compensation problem into two relative easier sub-problems, *i.e.*, coarse alignment (CA) and refined alignment (RA).

The CA module needs to compensate large shift roughly using small computational cost. We choose block matching (BM) based method and propose a differentiable progressive block matching method to estimate the offset on low-resolution (LR) features and output the coarsely aligned burst images. The RA module is then developed, which uses deformable convolution (DConv) to pixel-wise align the burst images in high-resolution (HR) feature domain. Finally, the fusion module estimates the clean full color image of the reference frame (denoted as frame r) by using the aligned image features. The details of each module are described in the following sections.

3.2. Coarse Alignment in Patch Level

The coarse alignment (CA) module aims to efficiently estimate the large offsets between frames. To meet the requirements of handling burst images with complex motion at 4K resolution, the CA module needs to have a large receptive field using small cost. Also in model training, to suppress the effect of small noise, the GT images are usually obtained by down-sampling high-resolution images, which also reduces the range of motion. Thus we choose to use BM to perform coarse alignment, which can easily mitigate this gap between training and testing stages by increasing the search region. To overcome the non-differentiability

Table 1. Quantitative comparison of different JDD-B approaches on the REDS4 dataset. Following the experiment setting of [33, 43], “Low” and “High” noise levels are corresponding to $\sigma_s = 2.5 \times 10^{-3}$, $\sigma_r = 10^{-2}$ and $\sigma_s = 6.4 \times 10^{-3}$, $\sigma_r = 2 \times 10^{-2}$, respectively.

Noise Level	Clip Name	KPN+DMN	EDVR+DMN	RviDeNet+DMN	EDVR*	RviDeNet*	GCP-Net	Ours
Low	<i>Clip000</i>	29.85/0.8398	32.15/0.8984	32.38/0.9075	32.44/0.9072	33.58/0.9285	34.96/0.9412	35.07/0.9501
	<i>Clip011</i>	32.01/0.8496	34.29/0.9030	34.39/0.9019	34.32/0.9017	34.63/0.9044	36.21/0.9288	36.79/0.9356
	<i>Clip015</i>	33.46/0.8770	35.42/0.9074	35.75/0.9169	35.54/0.9160	36.98/0.9327	37.74/0.9403	38.13/0.9456
	<i>Clip020</i>	31.29/0.8613	33.56/0.9071	33.81/0.9165	33.86/0.9179	34.23/0.9230	35.92/0.9428	36.38/0.9479
	Average	31.65/0.8570	33.85/0.9039	34.08/0.9107	34.02/0.9105	34.86/0.9221	36.20/0.9383	36.59/0.9448
High	<i>Clip000</i>	27.47/0.7437	30.18/0.8517	30.37/0.8532	30.31/0.8539	31.29/0.8803	32.57/0.9147	32.75/0.9174
	<i>Clip011</i>	29.64/0.7886	32.26/0.8611	32.50/0.8639	32.51/0.8643	32.50/0.8710	34.20/0.8972	34.74/0.9059
	<i>Clip015</i>	31.21/0.8310	34.01/0.8893	34.10/0.8909	34.01/0.8919	34.90/0.9068	35.94/0.9225	36.06/0.9247
	<i>Clip020</i>	28.66/0.7938	31.65/0.8780	31.74/0.8807	31.44/0.8757	31.82/0.8822	33.61/0.9154	34.17/0.9233
	Average	29.24/0.7893	32.02/0.8700	32.18/0.8722	32.06/0.8715	32.62/0.8850	34.08/0.9124	34.43/0.9178

caused by BM and reduce computational cost, we propose a differentiable progressive block matching (DPBM) method. **Differentiable BM.** We first introduce the differentiable BM process. For a patch P_r in the reference frame r and a query patch $(P_{t,i})_{i \in I}$ in another frame t with indices $I = \{1, \dots, M\}$, the purpose of BM is to find the best matched patch in $(P_{t,i})_{i \in I}$ with the target patch P_r . The normalized mean-absolute (NMA) distance is used as the matching criterion. The NMA distance between image patches $P_{t,i}$ and P_r can be written as:

$$d(P_{t,i}, P_r) = \frac{\mathbb{E}[|P_{t,i} - P_r|]}{\sqrt{\sum_{i=1}^M \mathbb{E}[|P_{t,i} - P_r|]^2}}. \quad (1)$$

For traditional BM, the patch with minimum distance is chosen as the best matched patch, which is denoted as $P_{t,BM}$. The minimum distance is obtained by sorting $d(P_{t,i}, P_r)_{i \in I}$ in ascending order, which is not differentiable. To overcome this problem, we refer to [12, 19, 35] and utilize continuous deterministic relaxation of BM. The weighting factor w of $d(P_{t,i}, P_r)_{i \in I}$ can be calculated as:

$$w_i = \frac{\exp(-d(P_{t,i}, P_r)/T)}{\sum_{i' \in I} \exp(-d(P_{t,i'}, P_r)/T)}, \quad (2)$$

where T is the temperature and w_i becomes one-hot vector when T approaches 0. We now can obtain $P_{t,BM}$ by using $P_{t,BM} = \sum_{i \in I} w_i P_{t,i}$. As discussed in [19], using small T can make w closer to one-hot, but it will result in a larger gradient and cause the training process to be unstable. While using large T , w becomes smooth but easy to train. Thus we regard training differentiable BM as a soft to hard process, where T is initialized as 1×10^{-2} and then reduced to 1×10^{-3} in our experiments.

Different from the differentiable BM in [35], here we calculate the distance between patches on LR features obtained by a lightweight network, which is updated by the gradient only from differentiable BM. We further design some constrains for DPBM to stabilize the training process, which will be introduced in Sec. 3.5.

DPBM. The detailed process of DPBM is illustrated in Fig. 2. Firstly, to reduce the computational cost and relieve the impact of noise interference, we perform differentiable BM on LR features which are obtained by a lightweight downsampling network. The network contains three 3×3 convolution (conv) layers with 16 channels and two transpose conv layers to downsample images into learned features with 1/4 scale. For the LR feature patch $P_{r,p}$ centered at position p in the reference frame r , the search region in its closest target frame, denoted as frame $r + 1$, is set to $\{p \pm \Delta p_{cmax}\}$, where Δp_{cmax} is the max search range. We firstly find an approximately matched position in $\{p \pm \Delta p_{cmax}\}$ with a stride $s > 1$ using differentiable BM and denote the matched position as p_a . Then we find a more accurate matching patch in $\{p_a \pm s\}$ with stride 1 and denote the matched position as p_m . The shift between frames r and $r + 1$ for patch P is obtained as $\Delta p_{r+1} = p_m - p$.

For target frame $r + 2$, the search region is set to $\{p + \Delta p_{r+1} \pm \Delta p_{cmax}\}$, which is updated using the offset estimated from frame $r + 1$. Usually, motion between frames has continuity, and updating the search region using temporal information can more effectively locate the matched patch in long-term frames. In the cases where the motion between frames changes abruptly, our DPBM is equivalent to the standard BM. Note that the DPBM is performed on 1/4 scale LR features and the estimated offsets should be rescaled to the original image resolution when aligning the full resolution images.

3.3. Refined Alignment in Pixel Level

The refined alignment module aims to perform accurate pixel-wise alignment on the coarsely aligned frames. The deformable alignment with DConv [6] has proved its success in feature alignment for various video processing tasks [13, 38, 40, 47]. Thus we choose to use implicit feature alignment in our refined alignment module.

The deep features are extracted from the burst images at the original resolution. The alignment of features of the reference frame and a target frame, denoted by F_r and F_{r+1} ,

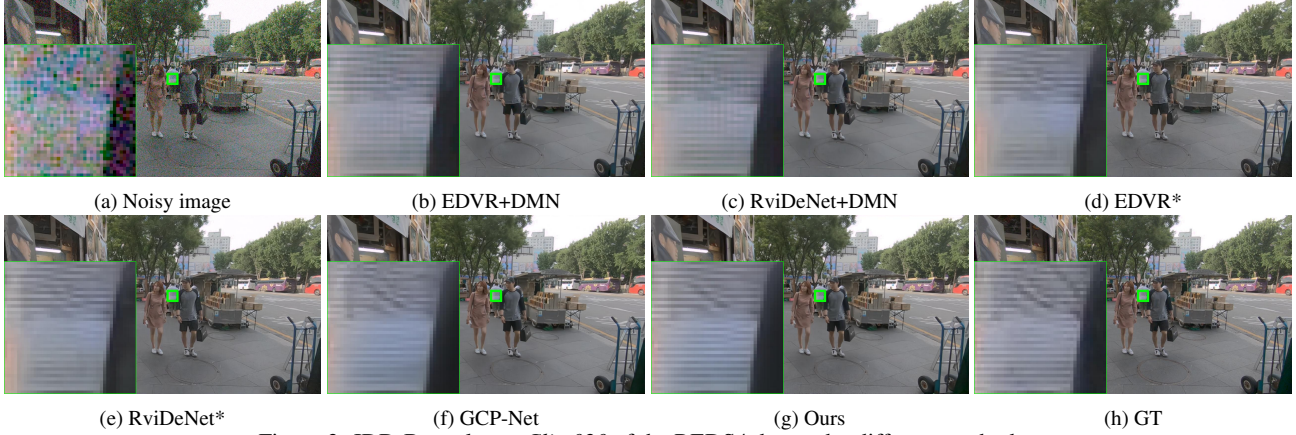


Figure 3. JDD-B results on *Clip 020* of the REDS4 dataset by different methods.

is described as follows. Firstly, the offset is estimated from the deep features by using $\Delta p_{r+1} = f([F_r, F_{r+1}])$, where $[\cdot, \cdot]$ is the concatenation operator and f refers to the combination of several Conv layers and nonlinear functions (e.g., LReLU). Then the aligned feature F'_{r+1} can be obtained by performing DConv with the estimated offset:

$$F'_{r+1} = \text{DConv}(F_{r+1}, \Delta p_{r+1}). \quad (3)$$

Similar to [40], we also adopt the pyramid DConv alignment strategy to improve the alignment performance.

3.4. Aligned Feature Fusion

In burst image restoration, the merging of aligned features is another important step. There are three typical solutions, i.e., non-directional fusion [18, 25, 40, 47], unidirectional fusion [18, 37] and bidirectional fusion [4, 16, 17]. Compared with non-directional and unidirectional fusion, bidirectional fusion can utilize long-term memory and information from both forward and backward directions [4]. Thus, we design a bi-directional gate recurrent unit (Bi-GRU) as the core component of our fusion module.

The input of our fusion module is the aligned feature frames $\{F'_t\}_{t=1}^N$. For the reference frame r , the features are propagated with forward (h_r^f) and backward (h_r^b) information, which can be calculated as:

$$h_r^f = f_{gru}(F_r, h_{r-1}^f), \quad h_r^b = f_{gru}(F_r, h_{r+1}^b), \quad (4)$$

where f_{gru} refers to the gate recurrent unit (GRU). The outputs of forward pass and backward pass are then concatenated as the input for clean full-color image reconstruction. A typical 3-scale UNet with two skip connections is used as the image reconstruction network. More details of network structure are provided in the supplementary file.

3.5. Design of Training Loss

DPBM Loss. Firstly, to avoid learning an all zeros vector or a uniform vector of w in Equ. 2, we use penalty when w

is not a one-hot vector by calculating the statistical characteristics of w , which is denoted as $\mathcal{L}_{one-hot}$. To be specific, the sum and variance of w should equal to 1 and $1/M$, respectively. M is the patch number of query $d(P_{t,i}, P_r)_{i \in I}$. The $\mathcal{L}_{one-hot}$ can be written as:

$$\mathcal{L}_{one-hot} = |\text{sum}(w) - 1| + |\text{var}(w) - 1/M|, \quad (5)$$

where $\text{sum}(w)$ and $\text{var}(w)$ are the sum and variance of w .

Due to the corruption of noise in raw images, the BM distances between noisy patches are not accurate. To stabilize the DPBM training process and relieve the impact of noise, we use the BM results obtained by clean images to guide the training process. We let

$$\mathcal{L}_{BM} = \|d(P_{t,i}, P_r) - d(P_{t,i}^c, P_r^c)\|_2^2, \quad (6)$$

where $P_{t,i}^c$ and P_r^c are the query patches in frame t and target patch in frame r using clean images. \mathcal{L}_{BM} is utilized in the first 20W iterations.

Interpolation Loss. We design an interpolation loss, denoted by \mathcal{L}_{ip} , to encourage the network to make better use of the information from other frames. The core idea of \mathcal{L}_{ip} is to use neighboring frames to “interpolate” the reference frame. In the fusion module, we aggregate the features from all other frames, excluding the reference frames, to interpolate the reference frame using the reconstruction UNet. The “interpolation” output is denoted as \hat{x}^i . To encourage the aligned features to capture more textures, we calculate \mathcal{L}_{ip} only on high frequency regions (denoted by m_h), which can be obtained using [28]. The \mathcal{L}_{ip} can be written as:

$$\mathcal{L}_{ip} = m_h \odot \sqrt{\|\hat{x}^i - x\|^2 + \epsilon^2}, \quad (7)$$

where $\sqrt{\|\hat{x}^i - x\|^2 + \epsilon^2}$ is the Charbonnier penalty function, and ϵ is set to 0.001.

Overall Loss. Denote by \hat{x} the reconstructed clean full-color image from the noisy raw burst data and by x the

Table 2. Quantitative comparison on the Videezy4K dataset. The noise level is set as $\sigma_s = 6.4 \times 10^{-3}$, $\sigma_r = 2 \times 10^{-2}$.

Methods	00	01	02	03	04	05	06	07	08	09	10	11	Average
EDVR*	36.07 0.9406	38.45 0.9583	40.07 0.9652	40.84 0.9721	37.82 0.9511	33.93 0.9207	36.96 0.9454	35.71 0.9310	37.02 0.9419	34.86 0.9299	37.60 0.9387	36.99 0.9367	37.19 0.9443
RviDeNet*	36.22 0.9490	38.56 0.9657	40.78 0.9738	40.42 0.9789	38.34 0.9582	34.30 0.9271	37.76 0.9543	36.40 0.9399	37.23 0.9448	34.83 0.9385	38.04 0.9432	37.16 0.9396	37.50 0.9510
GCP-Net	35.51 0.9306	39.17 0.9683	41.32 0.9755	40.85 0.9646	38.74 0.9653	34.83 0.9365	38.50 0.9650	37.12 0.9530	37.69 0.9482	35.67 0.9458	38.21 0.9473	37.78 0.9424	37.94 0.9535
Ours	37.75 0.9578	39.61 0.9737	41.71 0.9809	42.42 0.9855	39.30 0.9698	36.09 0.9532	39.24 0.9741	37.21 0.9594	38.76 0.9604	35.94 0.9491	38.64 0.9515	38.27 0.9514	38.74 0.9639




Figure 4. JDD results of different methods on the Videezy4K dataset.

ground-truth clean image of the reference frame. A reconstruction loss can be defined as:

$$\mathcal{L}_r = \sqrt{\|\hat{x} - x\|^2 + \epsilon^2} + \sqrt{\|\Gamma(\hat{x}) - \Gamma(x)\|^2 + \epsilon^2}, \quad (8)$$

in which $\Gamma(\cdot)$ is the ISP operator consisting of white balance, color correction and gamma compression. The implementation details of $\Gamma(\cdot)$ can be found in [3].

Overall, the total loss of our model is:

$$\mathcal{L} = \mathcal{L}_r + \beta \mathcal{L}_{ip} + \rho \mathcal{L}_{one-hot} + \eta \mathcal{L}_{BM}, \quad (9)$$

where β and η are the balancing parameters and are set to 1 and 1×10^3 . The penalty factor ρ is set to 1×10^5 .

4. Experiments

4.1. Model Training Details

Training Data. We follow the method in [3] to synthesize training data for real-world burst images. The 240 training clips of 720p in the REDS dataset [34] are used to generate training data. Firstly, the sRGB videos are unprocessed to linear RGB space using [3]. The obtained sequences are taken as the clean ground-truth images x_t . The inputs of the network $\{y_t\}_{t=1}^N$ are the noisy raw burst images which are obtained by [10]:

$$y_t = M(x_t) + n(M(x_t), \sigma_s, \sigma_r), \quad (10)$$

where $n(x, \sigma_s, \sigma_r) \sim \mathcal{N}(0, \sigma_s x + \sigma_r^2)$, and $M(\cdot)$ is the mosaic downsampling operator. σ_s and σ_r represent the scale of shot noise and read noise, respectively.

Training Details. We use the center frame as the reference frame. Following [13,33], the σ_s and σ_r are uniformly sampled in the ranges of $[10^{-4}, 10^{-2}]$ and $[10^{-3}, 10^{-1.5}]$. Our model is implemented in PyTorch and is trained using two RTX 2080Ti GPUs. During training, we use the Adam [22] optimizer with momentum 0.9. The learning rates are initialized as 1×10^{-5} for the lightweight network in coarse alignment and 1×10^{-4} for other network parts. Then the learning rate is decreased using the cosine function [30].

4.2. Comparison with State-of-the-Art Methods

Competing Methods. We compare our method with the recently developed JDD-B method GCP-Net [13]. Following [13], we also combine several SOTA burst denoising algorithms, *i.e.*, KPN [33], EDVR [40] and RviDeNet [47], with a SOTA demosaicking method DMN [11] for comparison. EDVR [40] and RviDeNet [47] are also slightly modified to adapt to the JDD-B task by adding an upsampling operator to the output layer, and we denote them as EDVR* and RviDeNet*. All competing models are retrained using our experiment setting.

Results on Synthetic Data. To quantitatively evaluate the competing methods, we first perform experiments on two synthetic datasets, *i.e.*, REDS4 [40] and a dataset collected by us, called Videezy4K. REDS4 is widely used as the test set in the study of video super-resolution, whose resolution is 720p. Since images/videos captured by modern smartphone cameras and DSLRs commonly have 4K resolution, the shift between frames can be larger than the




Figure 5. JDD-B results on real-world burst images by different methods.

REDS4 dataset. To better evaluate the performance on 4K videos, we collected 12 clips of 4K videos from the Videezy website [39] and choose 20 frames of each video for testing. All testing videos are firstly converted to linear RGB space using the same pipeline as generating training data and the noisy burst raw images are generated by adding heteroscedastic Gaussian noise with Equ. 10.

The quantitative results on REDS4 and Videezy4K datasets are shown in Tables 1 and 2, respectively. Following [13, 33], we calculate the PSNR and SSIM indices after gamma correction to better reflect perceptual quality. The qualitative comparisons on REDS4 and Videezy4K are presented in Figs. 3 and 4, respectively. We can see that due to the correlation between denoising and demosaicking, JDD-B algorithms (*i.e.*, EDVR* and RviDeNet*) obtain better performance than EDVR+DMN and RviDeNet+DMN. However, limited by the small receptive field of one-stage alignment module, EDVR*, RviDeNet* and GCP-Net suffer from the over-smoothing problem in areas with large motion. Benefiting from the two-stage alignment, our model recovers more textures, especially on the moving objects. In Fig. 3, the horizontal stripes on the walking people are well reconstructed with little zippers and color artifacts by our model. In Fig. 4, the characters can be more clearly restored and identified by our method. For images with small motion, *i.e.*, *Clip000* in the REDS4 dataset, our method still obtains better performance. More visual results can be found in the supplementary file.

Results on Real-world Burst Images. We then qualitatively evaluate our method on real-world burst images. The SCBurst dataset [13] is used, where burst images are captured using smartphones with a wide range of ISO values. SCBurst contains noisy raw images and the corresponding metadata, which are used in the testing and visualization. We provide visual comparisons in Fig. 5. One can see that the restoration results of KPN+DMN, EDVR+DMN,

EDVR*, RviDeNet+DMN and RviDeNet* contain noticeable motion induced artifacts. GCP-Net exhibits less artifact but it fails to reconstruct the textures in the area with high noise level. Benefiting from the two-stage alignment module, our model can more effectively utilize the temporal information and reconstruct more details. We provide more visual results and user study in the supplementary file.

Model Size and Running Speed. Table 3 lists the number of model parameters and running time on 5 frames with UHD resolution by our method and compared methods on a GTX 2080Ti GPU. EDVR* has the least number of parameters and fastest speed but worst performance. RviDeNet* has the highest computational cost since it consists of a pre-denoising module and non-local modules. Compared with GCP-Net, our model obtains better performance with less training parameters and similar running time. The CA module costs 3.6s and RA+fusion module costs 16.1s.

Table 3. Comparison of different CNN models in terms of the number of parameters and running times on 5 frames with input size 2160×3840 and output size $2160 \times 3840 \times 3$.

	EDVR*	RviDeNet*	GCP-Net	Ours
#. Params	6.28M	57.98M	13.79M	12.05M
Times(s)	10.4	124.3	19.5	19.7 (3.6+16.1)

4.3. Ablation Study

Coarse Alignment vs. Refined Alignment. To evaluate the role of CA and RA modules in our model, we compare it with three variants, *i.e.*, without using alignment (ours(w/o Align)), only using CA module (ours(w/ CA)), and only using RA module (ours(w/ RA)). The quantitative comparison on the REDS4 dataset is shown in Table 4.

One can see that ours(w/ RA) can achieve great improvement (1.83dB) over ours(w/o Align) when compensating for small motions (*Clip 000*), but achieve limited gain (0.26dB) for sequences with large motion (*Clip 020*). This shows that limited by the small receptive field, DConv fails

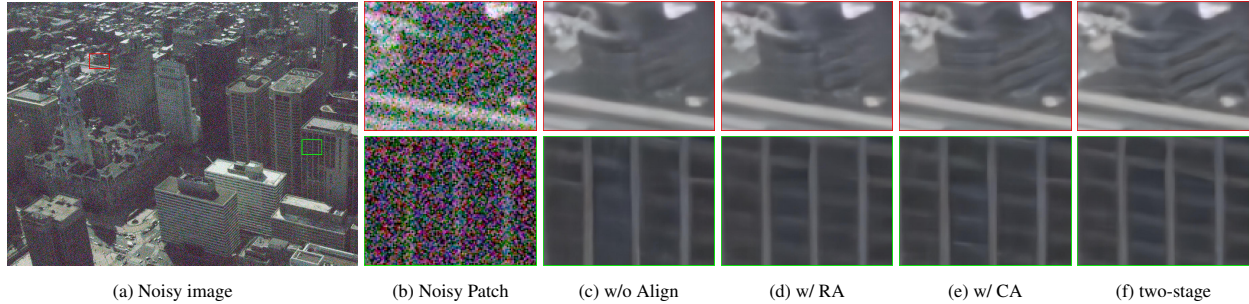


Figure 6. JDD-B results by using different variants of our two-stage alignment model.

to effectively use temporal information for scenes with large motion. By using DPBM to search in a large region, ours(w/ CA) can more effectively utilize temporal information, and obtain 0.52dB gain than ours(w/ RA) for images with large motion (*Clip 020*). However, it achieves smaller improvement (0.31dB) on *Clip 000* which has small motion. We also visualize in Fig. 6 the JDD-B results using models with different alignment modules. One can see that two-stage alignment can reconstruct more textures.

On average, without CA module, ours(w/ RA) achieves 1.02dB gain over ours(w/o Align). When CA module is included, ours full model achieves more improvement (1.28dB) than ours(w/ CA). Meanwhile, without RA module, ours(w/ CA) brings 0.27dB gain than ours(w/o Align), while by considering RA module, our full model brings 0.53dB gain than ours(w/ RA). One can see clearly that both alignment steps can improve the JDD-B performance. CA and RA modules are complementary to each other and they can enhance each other when used together.

Table 4. Comparison of different variants on a clip with small motion (*Clip 000*) and a clip with large motion (*Clip 020*) in the REDS4 dataset, and the average results on the whole REDS4.

	<i>Clip 000</i>	<i>Clip 020</i>	Avg
ours(w/o Align)	30.74/0.8662	32.43/0.8969	32.88/0.8889
ours(w/ CA)	31.05/0.8738	33.21/0.9080	33.15/0.8932
ours(w/ RA)	32.57/0.9190	32.69/0.9002	33.90/0.9077
ours(w/o E2E)	32.67/0.9158	34.07/0.9217	34.31/0.9159
ours(full)	32.75/0.9174	34.17/0.9233	34.43/0.9178

End-to-end Learning of Two-stage Alignment. In order to train our method in an end-to-end manner, we proposed the DPBM module. To evaluate the effectiveness of end-to-end learning, we train a variant, namely ours(w/o E2E), which uses normal BM in the CA module. Since normal BM is non-differentiable, the learnable lightweight downsampling network is replaced by bicubic downsampling. The results are also shown in Table 4. One can see that our full model can obtain 0.12dB improvement over ours(w/o E2E), validating that the end-to-end CA module can learn better alignment for restoration. More importantly, since the main difference among flat patches is caused by noise, directly performing BM on noisy patches may align noise and generate artifacts. An example is shown in Fig. 7. Owe

to the use of a learnable lightweight network for downsampling, our full model can relieve the noise interference especially when aligning flat areas.

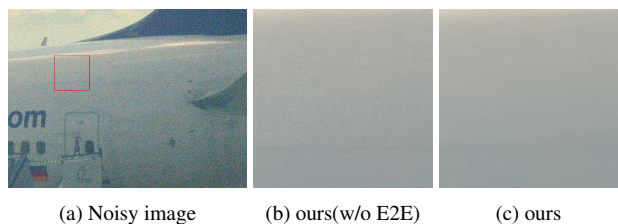


Figure 7. The reconstruction on flat areas with/without end-to-end two-stage alignment learning. Better viewed with zoom-in on screen. Our model relieves the noise interference and obtains cleaner results.

Limitation. When there are objects with large movement in opposite directions in the patch $(P_{t,i})_{i \in I}$, the CA module may fail and our method degrades to use only the RA module in such cases.

5. Conclusion

We presented a differentiable two-stage alignment method for high performance burst image restoration. Due to the limited receptive field, current feature alignment methods cannot sufficiently utilize the temporal information when the burst images have large shift, which is very common for sequences with moving objects and/or 4K resolution. We divided this problem into two relatively easier sub-problems, *i.e.*, coarse and refined alignment, and proposed a two-stage framework for the JDD-B task. In the coarse alignment module, a differentiable progressive block matching module was developed to enlarge the search region while reducing computational cost. Then, a deformable alignment module was developed to deliver pixel-wise alignment. Experiments on both synthetic and real-world burst datasets were conducted. Our method demonstrated clear advantages over existing methods in terms of PSNR/SSIM measures as well as visual quality without increasing much the computational cost.

Acknowledgements This work is supported by the Hong Kong RGC RIF grant (R5001-18)

References

- [1] CIPA. 2018. CIPA Report. (accessed 29-Nov2018). [1](#)
- [2] Miika Aittala and Frédo Durand. Burst image deblurring using permutation invariant convolutional neural networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 731–747, 2018. [1](#)
- [3] Tim Brooks, Ben Mildenhall, Tianfan Xue, Jiawen Chen, Dillon Sharlet, and Jonathan T Barron. Unprocessing images for learned raw denoising. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11036–11045, 2019. [6](#)
- [4] Kelvin CK Chan, Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Basicvsr: The search for essential components in video super-resolution and beyond. *arXiv preprint arXiv:2012.02181*, 2020. [1](#), [3](#), [5](#)
- [5] Laurent Condat and Saleh Mosaddegh. Joint demosaicking and denoising by total variation minimization. In *2012 19th IEEE International Conference on Image Processing*, pages 2781–2784. IEEE, 2012. [1](#)
- [6] Jifeng Dai, Haozhi Qi, Yuwen Xiong, Yi Li, Guodong Zhang, Han Hu, and Yichen Wei. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, pages 764–773, 2017. [3](#), [4](#)
- [7] Thibaud Ehret, Axel Davy, Pablo Arias, and Gabriele Facciolo. Joint demosaicking and denoising by fine-tuning of bursts of raw images. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 8868–8877, 2019. [2](#)
- [8] Thibaud Ehret, Axel Davy, Jean-Michel Morel, Gabriele Facciolo, and Pablo Arias. Model-blind video denoising via frame-to-frame training. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11369–11378, 2019. [1](#), [3](#)
- [9] Thibaud Ehret and Gabriele Facciolo. A study of two cnn demosaicking algorithms. *Image Processing On Line*, 9:220–230, 2019. [1](#), [2](#)
- [10] Alessandro Foi, Mejdi Trimeche, Vladimir Katkovnik, and Karen Egiazarian. Practical poissonian-gaussian noise modeling and fitting for single-image raw-data. *IEEE Transactions on Image Processing*, 17(10):1737–1754, 2008. [6](#)
- [11] Michaël Gharbi, Gaurav Chaurasia, Sylvain Paris, and Frédo Durand. Deep joint demosaicking and denoising. *ACM Transactions on Graphics (TOG)*, 35(6):1–12, 2016. [1](#), [2](#), [6](#)
- [12] Jacob Goldberger, Geoffrey E Hinton, Sam Roweis, and Russ R Salakhutdinov. Neighbourhood components analysis. *Advances in neural information processing systems*, 17, 2004. [4](#)
- [13] Shi Guo, Zhetong Liang, and Lei Zhang. Joint denoising and demosaicking with green channel prior for real-world burst images. *arXiv preprint arXiv:2101.09870*, 2021. [1](#), [2](#), [3](#), [4](#), [6](#), [7](#)
- [14] Felix Heide, Markus Steinberger, Yun-Ta Tsai, Mushfiqur Rouf, Dawid Pajk, Dikpal Reddy, Orazio Gallo, Jing Liu, Wolfgang Heidrich, Karen Egiazarian, et al. Flexisp: A flexible camera image processing framework. *ACM Transactions on Graphics (TOG)*, 33(6):1–13, 2014. [1](#)
- [15] Bernardo Henz, Eduardo SL Gastal, and Manuel M Oliveira. Deep joint design of color filter arrays and demosaicing. In *Computer Graphics Forum*, volume 37, pages 389–399. Wiley Online Library, 2018. [1](#), [2](#)
- [16] Yan Huang, Wei Wang, and Liang Wang. Bidirectional recurrent convolutional networks for multi-frame super-resolution. In *Proceedings of the 28th International Conference on Neural Information Processing Systems-Volume 1*, pages 235–243, 2015. [5](#)
- [17] Yan Huang, Wei Wang, and Liang Wang. Video super-resolution via bidirectional recurrent convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 40(4):1015–1028, 2017. [5](#)
- [18] Takashi Isobe, Songjiang Li, Xu Jia, Shanxin Yuan, Gregory Slabaugh, Chunjing Xu, Ya-Li Li, Shengjin Wang, and Qi Tian. Video super-resolution with temporal group attention. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8008–8017, 2020. [3](#), [5](#)
- [19] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016. [4](#)
- [20] Qiyu Jin, Gabriele Facciolo, and Jean-Michel Morel. A review of an old dilemma: Demosaicking first, or denoising first? In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 514–515, 2020. [1](#), [2](#)
- [21] Wesley Khademi, Sonia Rao, Clare Minnerath, Guy Hagen, and Jonathan Ventura. Self-supervised poisson-gaussian denoising. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 2131–2139, 2021. [1](#)
- [22] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014. [6](#)
- [23] Filippos Kokkinos and Stamatios Lefkimmiatis. Iterative joint image demosaicking and denoising using a residual denoising network. *IEEE Transactions on Image Processing*, 28(8):4177–4188, 2019. [1](#), [2](#)
- [24] Dongxu Li, Chenchen Xu, Kaihao Zhang, Xin Yu, Yiran Zhong, Wenqi Ren, Hanna Suominen, and Hongdong Li. Arvo: Learning all-range volumetric correspondence for video deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7721–7731, 2021. [2](#)
- [25] Wenbo Li, Xin Tao, Taian Guo, Lu Qi, Jiangbo Lu, and Jiaya Jia. Mucan: Multi-correspondence aggregation network for video super-resolution. In *European Conference on Computer Vision*, pages 335–351. Springer, 2020. [5](#)
- [26] Ding Liu, Zhaowen Wang, Yuchen Fan, Xianming Liu, Zhangyang Wang, Shiyu Chang, and Thomas Huang. Robust video super-resolution with learned temporal dynamics. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2507–2515, 2017. [3](#)
- [27] Hongying Liu, Peng Zhao, Zhubo Ruan, Fanhua Shang, and Yuanyuan Liu. Large motion video super-resolution with dual subnet and multi-stage communicated upsampling. *arXiv preprint arXiv:2103.11744*, 2021. [2](#)

- [28] Lin Liu, Xu Jia, Jianzhuang Liu, and Qi Tian. Joint demosaicing and denoising with self guidance. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2240–2249, 2020. 1, 2, 5
- [29] Shumin Liu, Jiajia Chen, Yuan Xun, Xiaojin Zhao, and Chip-Hong Chang. A new polarization image demosaicking algorithm by exploiting inter-channel correlations with guided filtering. *IEEE Transactions on Image Processing*, 29:7076–7089, 2020. 1, 2
- [30] Ilya Loshchilov and Frank Hutter. Sgdr: Stochastic gradient descent with warm restarts. *arXiv preprint arXiv:1608.03983*, 2016. 6
- [31] Florian Luisier, Thierry Blu, and Michael Unser. Image denoising in mixed poisson–gaussian noise. *IEEE Transactions on image processing*, 20(3):696–708, 2010. 1
- [32] Talmaj Marinč, Vignesh Srinivasan, Serhan Gül, Cornelius Hellge, and Wojciech Samek. Multi-kernel prediction networks for denoising of burst images. In *2019 IEEE International Conference on Image Processing (ICIP)*, pages 2404–2408. IEEE, 2019. 1, 3
- [33] Ben Mildenhall, Jonathan T Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, and Robert Carroll. Burst denoising with kernel prediction networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2502–2510, 2018. 1, 3, 4, 6, 7
- [34] Seungjun Nah, Sungyong Baik, Seokil Hong, Gyeongsik Moon, Sanghyun Son, Radu Timofte, and Kyoung Mu Lee. Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 6
- [35] Tobias Plötz and Stefan Roth. Neural nearest neighbors networks. *arXiv preprint arXiv:1810.12575*, 2018. 4
- [36] Guocheng Qian, Jinjin Gu, Jimmy S Ren, Chao Dong, Furong Zhao, and Juan Lin. Trinity of pixel enhancement: a joint solution for demosaicking, denoising and super-resolution. *arXiv preprint arXiv:1905.02538*, 2019. 2
- [37] Mehdi SM Sajjadi, Raviteja Vemulapalli, and Matthew Brown. Frame-recurrent video super-resolution. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6626–6634, 2018. 5
- [38] Yapeng Tian, Yulun Zhang, Yun Fu, and Chenliang Xu. Tdan: Temporally deformable alignment network for video super-resolution. *arXiv preprint arXiv:1812.02898*, 2018. 3, 4
- [39] Videezy. Videezy website. <https://www.videezy.com/>, [Online; accessed 20-Nov2020]. 7
- [40] Xintao Wang, Kelvin CK Chan, Ke Yu, Chao Dong, and Chen Change Loy. Edvr: Video restoration with enhanced deformable convolutional networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 0–0, 2019. 1, 3, 4, 5, 6
- [41] Bartłomiej Wronski, Ignacio Garcia-Dorado, Manfred Ernst, Damien Kelly, Michael Krainin, Chia-Kai Liang, Marc Levoy, and Peyman Milanfar. Handheld multi-frame super-resolution. *ACM Transactions on Graphics (TOG)*, 38(4):1–18, 2019. 1
- [42] Zhihao Xia, Federico Perazzi, Michaël Gharbi, Kalyan Sunkavalli, and Ayan Chakrabarti. Basis prediction networks for effective burst denoising with large kernels. *arXiv preprint arXiv:1912.04421*, 2019. 1, 3
- [43] Xiangyu Xu, Muchen Li, and Wenxiu Sun. Learning deformable kernels for image and video denoising. *arXiv preprint arXiv:1904.06903*, 2019. 1, 3, 4
- [44] Tianfan Xue, Baian Chen, Jiajun Wu, Donglai Wei, and William T Freeman. Video enhancement with task-oriented flow. *International Journal of Computer Vision*, 127(8):1106–1125, 2019. 1, 3
- [45] Niu Yan and Jihong Ouyang. Cross-channel correlation preserved three-stream lightweight cnns for demosaicking. *arXiv preprint arXiv:1906.09884*, 2019. 1, 2
- [46] Bin Yang and Dongsheng Wang. An efficient adaptive interpolation for bayer cfa demosaicking. *Sensing and Imaging*, 20(1):37, 2019. 1, 2
- [47] Huanjing Yue, Cong Cao, Lei Liao, Ronghe Chu, and Jingyu Yang. Supervised raw video denoising with a benchmark dataset on dynamic scenes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2301–2310, 2020. 1, 3, 4, 5, 6
- [48] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017. 1
- [49] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, 2018. 1