

Adversarial Spatio-Temporal Learning for Video Deblurring

Kaihao Zhang¹, Wenhan Luo¹, Yiran Zhong, Lin Ma¹, *Member, IEEE*, Wei Liu², and Hongdong Li

Abstract—Camera shake or target movement often leads to undesired blur effects in videos captured by a hand-held camera. Despite significant efforts having been devoted to video-deblur research, two major challenges remain: 1) how to model the spatio-temporal characteristics across both the spatial domain (i.e., image plane) and the temporal domain (i.e., neighboring frames) and 2) how to restore sharp image details with respect to the conventionally adopted metric of pixel-wise errors. In this paper, to address the first challenge, we propose a deblurring network (DBLRNet) for spatial-temporal learning by applying a 3D convolution to both the spatial and temporal domains. Our DBLRNet is able to capture jointly spatial and temporal information encoded in neighboring frames, which directly contributes to the improved video deblur performance. To tackle the second challenge, we leverage the developed DBLRNet as a generator in the generative adversarial network (GAN) architecture and employ a content loss in addition to an adversarial loss for efficient adversarial training. The developed network, which we name as deblurring GAN, is tested on two standard benchmarks and achieves the state-of-the-art performance.

Index Terms—Spatio-temporal learning, adversarial learning, video deblurring.

I. INTRODUCTION

VIDEOS captured by hand-held cameras often suffer from unwanted blurs either caused by camera shake [1], or object movement in the scene [2], [3]. The task of video deblurring aims at removing those undesired blurs and recovering sharp frames from the input video. This is an active research topic in the applied fields of computer vision and image processing. Applications of video deblurring are found in many important fields such as 3D reconstruction [4], SLAM [5] and tracking [6].

In contrast to single image deblurring, video deblurring is a relatively less tapped task until recently. And video deblurring

Manuscript received March 20, 2018; revised July 8, 2018 and August 3, 2018; accepted August 14, 2018. Date of publication August 29, 2018; date of current version September 19, 2018. This work was supported by the 2017 Tencent Rhino Bird Elite Graduate Program. The work of K. Zhang was supported by The Australian National University. The work of Y. Zhong was supported by CSIRO Data61. The work of H. Li was supported by the Australian Research Council. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Xin Li. (Corresponding author: Kaihao Zhang.)

K. Zhang, Y. Zhong, and H. Li are with the College of Engineering and Computer Science, The Australian National University, Canberra, ACT 2600, Australia (e-mail: kaihao.zhang@anu.edu.au; yiran.zhong@anu.edu.au; hongdong.li@anu.edu.au).

W. Luo, L. Ma, and W. Liu are with the Tencent AI Laboratory, Shenzhen 518057, China (e-mail: whluo.china@gmail.com; forest.linma@gmail.com; wl2223@columbia.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TIP.2018.2867733

is more challenging, partly because it is not entirely clear about how to model and exploit the inherent temporal dynamics exhibited among continuous video frames. Moreover, the commonly adopted performance metric, namely, pixel-wise residual error, often measured by PSNR, is questionable, as it fails to capture human visual intuitions of how sharp or how realistic a restored image is [7], [8]. In this paper, we plan to leverage the recent advance of the adversarial learning technique to improve the performance of video deblurring.

One key challenge for video deblurring is to find an effective way to capture spatio-temporal information existing in neighboring image frames. Deep learning based methods have recently witnessed a remarkable success in many applications including image and video denoising and deblurring. Previous deep learning methods are however primarily based on 2D convolutions, mainly for computational sake. Yet, it is not natural to use 2D convolutions to capture spatial and temporal joint information, which is essentially in a 3D feature space. In this paper, we propose a deep neural network called DeBLuRing Network (DBLRNet), which uses 3D (volumetric) convolutional layers, as well as deep residual learning, aims to learn feature representations both across temporal frames and across image plane.

As noted above, we argue that the conventional pixel-wise PSNR metric is insufficient for the task of image/video deblurring. To address this issue, we resort to adversarial learning, and propose DeBLuRing Generative Adversarial Network (DBLRGAN). DBLRGAN consists of a generative network and a discriminate network, where the *generative* network is the aforementioned DBLRNet which restores sharp images, and the *discriminate* network is a binary classification network, which tells a restored image apart from a real-world sharp image.

We introduce a training loss which consists of two terms: content loss and adversarial loss. The content loss is used to respect the pixel-wise measurement, while the adversarial loss promotes a more realistically looking (hence sharper) image. Training DBLRGAN in an end-to-end manner, we recover sharp video frames from a blurred input video sequence, with some examples shown in Figure 1.

The contributions of this work are as follows:

- We propose a model called DBLRNet, which applies 3D convolutions in a deep residual network to capture joint spatio-temporal features for video deblurring.
- Based on the above DBLRNet, we develop a generative adversarial network, called DBLRGAN, with both content and adversarial losses. By training it in an adversarial

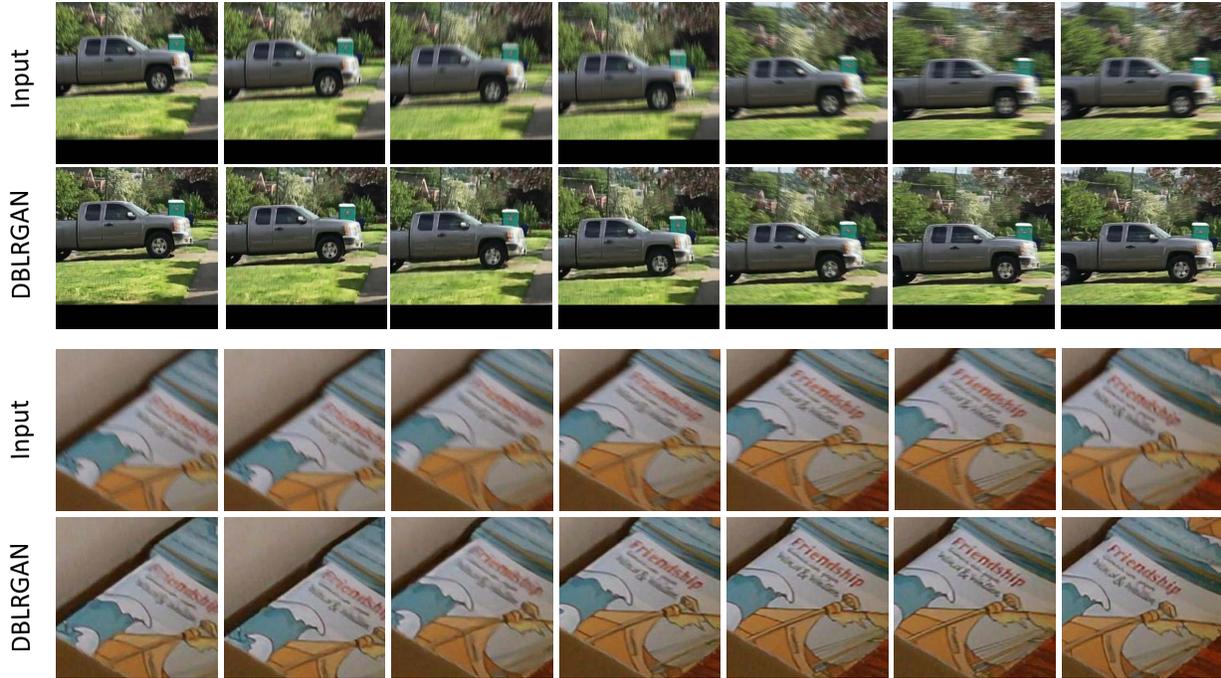


Fig. 1. Deblurring results of the proposed DBLRGAN on real-world video frames. The first and third rows show crops of consecutive frames from the VideoDeblurring dataset. The second and fourth rows show corresponding deblurring results of DBLRGAN.

manner, the DBLRGAN recovers video frames which look more realistic.

- Experiments on two standard benchmark datasets, including the VideoDeblurring dataset and the Blurred KITTI dataset, show that the proposed network DBLRNet and DBLRGAN are effective and outperforms existing methods.

II. RELATED WORK

Many approaches have been proposed for image/video deblurring, which can be roughly classified into two categories: geometry-based methods and deep learning methods.

A. Geometry-Based Methods

Modern single-image deblurring methods iteratively estimate uniform or non-uniform blur kernels and the latent sharp image given a single blurry image [9]–[20]. However, it is difficult for single image based methods to estimate kernel because blur is spatially varying in real world. To employ additional information, multi-image based methods [21]–[27] have been proposed to address blur, such as flash/no-flash image pairs [23], blurred/noise image pairs [22] and gyroscope information [24]. In order to accurately estimate kernels, some methods also use optical flow [28] and temporal information [29]. However, most of these methods are limited by the performance of an assumed degradation model and its estimation, thus some of them are fragile and cannot handle more challenging cases.

Some researchers attempt to use aggregation methods to alleviate blur. Law *et al.* [30] propose a lucky image system, which constructs a final image based on the best pixels

from different low quality images. Cho *et al.* [31] use patch-based synthesis to restore blurry regions and ensure that the deblurred frames are spatially and temporally coherent. Motivated from the physiological fact, an efficient Fourier aggregation method is proposed in [32], which creates a consistently registered version of neighboring frames, and then fuses these frames in the Fourier domain.

More recently, Pan *et al.* [33] propose to simultaneously deblur stereo videos and estimate the scene flow. In this method, motion cues from scene flow estimation, and blur information can complement each other and boost the performance. However, this kind of approaches is restricted to stereo cameras.

B. Deep Learning Methods

Deep learning has shown its effectiveness in many computer vision tasks, such as object detection [34], image classification [35]–[38], facial processing [2], [39]–[42] and multimedia analysis [43]–[45]. There are also several deep learning methods that achieve encouraging results on deblurring [2], [46]–[51]. A non-bind deblurring method with deep Convolutional Neural Network (CNN) is proposed in [46]. This data-driven approach establishes connection between traditional optimization-based schemes and empirically-determined CNN. Sun *et al.* [2] predict the probabilistic distribution of motion kernels at the patch level using CNN, then use a Markov random field model to fuse the estimations into a dense field of motion kernels. Finally, a non-uniform deblurring model using patch-level prior is employed to remove motion blur. In [47], a deep multi-scale CNN is proposed for image deblurring.

Most of these methods aim at image deblurring, so they do not need to consider temporal information implied in videos.

For video deblurring, the method closest to our approach is DBN [48], which proposes a CNN model to process information across frames. Neighboring frames are stacked along RGB channels and then fed into the proposed model to recover the central frame of them. This method considers multiple frames together and thus achieves comparable performance with state-of-the-art methods.

However, this method employs 2D convolutions, which do not operate in the time axis (corresponding to temporal information). By doing so, temporal information is transformed into spatial information in their setting, thus limited temporal information is preserved. Meanwhile, this method (as most of the existing methods) trains the model to maximize the pixel fidelity, which cannot ensure that the recovered images look realistic sharp. Our proposed method, on the contrary, learns spatio-temporal features by 3D convolutions, and integrates the 3D deblurring network into a generative adversarial network to achieve photo-realistic results.

Even our proposed method takes neighboring frames as inputs, we call it as video deblurring method for two reasons. Firstly, previous work like DBN [48], is called as video deblurring method. DBN also takes five neighboring frames as input to generate the middle sharp frame. Secondly, this method can be applied to tackle video deblurring task in the real world. Specially, videos can be regarded as multiple consecutive frames. When videos are input into our model, the proposed DBLRNet tackles five neighboring frames as a whole to generate the deblurred middle frame based on their spatio-temporal information, and obtain the deblurred videos by continuous deblurred frames finally.

III. OUR MODEL

Overview: In this section, we first introduce our DBLRNet, and then present the proposed network DBLRGAN which is on the basis of DBLRNet. Finally we detail the two loss functions (content and adversarial losses) which are used in the training stage. Both the DBLRNet and DBLRGAN are end-to-end systems for video deblurring. Note that, blurry frames can be put into our proposed models without alignment.

A. DBLRNet

In 2D CNN, convolutions are applied on 2D images or feature maps to learn features in spatial dimensions only. In case of video analysis problems, it is desirable to consider the motion variation encoded in the temporal dimension, such as multiple neighboring frames. In this paper, we propose to perform 3D convolutions [43] the convolution stages of deep residual networks to learn feature representations from both spatial and temporal dimensions for video deblurring. We operate the 3D convolution via convolving 3D kernels/filters with the cube constructed from multiple neighboring frames. By doing so, the feature maps in the convolution layers can capture the dynamic variations, which is helpful to model the blur evolution and further recover sharp frames.

TABLE I

CONFIGURATIONS OF THE PROPOSED DBLRNET. IT IS COMPOSED OF TWO CONVOLUTIONAL LAYERS (L1 AND L2), 14 RESIDUAL BLOCKS, TWO CONVOLUTIONAL LAYERS (L31 AND L32) WITHOUT SKIP CONNECTION, AND THREE ADDITIONAL CONVOLUTIONAL LAYERS (L33, L34 AND L35). EACH RESIDUAL BLOCK CONTAINS TWO CONVOLUTIONAL LAYERS, WHICH ARE INDICATED BY L(X) AND L(X+1) IN THE TABLE, WHERE “X” EQUALS 3, 5, 7, 9, 11, 13, 15, 17, 19, 21, 23, 25, 27 AND 29 RESPECTIVELY FOR THESE RESIDUAL BLOCKS

layers	Kernel size	output channels	operations	skip connection
L1	$3 \times 3 \times 3$	16	ReLU	-
L2	$3 \times 3 \times 3$	64	ReLU	L4, L32
L3	$3 \times 3 \times 1$	64	BN + ReLU	-
L4	$3 \times 3 \times 1$	64	BN	L6
L5	$3 \times 3 \times 1$	64	BN + ReLU	-
L6	$3 \times 3 \times 1$	64	BN	L8
L7	$3 \times 3 \times 1$	64	BN + ReLU	-
L8	$3 \times 3 \times 1$	64	BN	L10
L9	$3 \times 3 \times 1$	64	BN + ReLU	-
L10	$3 \times 3 \times 1$	64	BN	L12
L11	$3 \times 3 \times 1$	64	BN + ReLU	-
L12	$3 \times 3 \times 1$	64	BN	L14
L13	$3 \times 3 \times 1$	64	BN + ReLU	-
L14	$3 \times 3 \times 1$	64	BN	L16
L15	$3 \times 3 \times 1$	64	BN + ReLU	-
L16	$3 \times 3 \times 1$	64	BN	L18
L17	$3 \times 3 \times 1$	64	BN + ReLU	-
L18	$3 \times 3 \times 1$	64	BN	L20
L19	$3 \times 3 \times 1$	64	BN + ReLU	-
L20	$3 \times 3 \times 1$	64	BN	L22
L21	$3 \times 3 \times 1$	64	BN + ReLU	-
L22	$3 \times 3 \times 1$	64	BN	L24
L23	$3 \times 3 \times 1$	64	BN + ReLU	-
L24	$3 \times 3 \times 1$	64	BN	L26
L25	$3 \times 3 \times 1$	64	BN + ReLU	-
L26	$3 \times 3 \times 1$	64	BN	L28
L29	$3 \times 3 \times 1$	64	BN + ReLU	-
L30	$3 \times 3 \times 1$	64	BN	L32
L31	$3 \times 3 \times 1$	64	BN + ReLU	-
L32	$3 \times 3 \times 1$	64	BN	-
L33	$3 \times 3 \times 1$	256	ReLU	-
L34	$3 \times 3 \times 1$	256	ReLU	-
L35	$3 \times 3 \times 1$	1	-	-

Formally, the 3D convolution operation is formulated as:

$$V_{ij}^{xyz} = \sigma \left(\sum_m \sum_{p=0}^{P_i-1} \sum_{q=0}^{Q_i-1} \sum_{r=0}^{R_i-1} V_{(i-1)m}^{(x+p)(y+q)(z+r)} \cdot g_{ijm}^{pqr} + b_{ij} \right), \quad (1)$$

where V_{ij}^{xyz} is the value at position (x, y, z) in the j -th feature map of the i -th layer, (P_i, Q_i, R_i) is the size of 3D convolution kernel. Q_i responds to the temporal dimension. g_{ijm}^{pqr} is the (p, q, r) -th value of the kernel connected to the m -th feature map from the $(i-1)$ -th layer. $\sigma(\cdot)$ is the ReLU nonlinearity activation function, which is shown to lead to better performance in various computer vision tasks than other activation functions, e.g. Sigmoid and Tanh.

Defining 3D convolution, we propose a model called DBLRNet, which is shown in Figure 2. DBLRNet is composed of two $3 \times 3 \times 3$ convolutional layers, several residual blocks [38], each containing two convolution layers, and another five convolutional layers. This architecture is designed inspired by the Fully Convolutional Neural Network (FCNN) [52], which is originally proposed for semantic segmentation. Different from FCNN and DBN [48], spatial

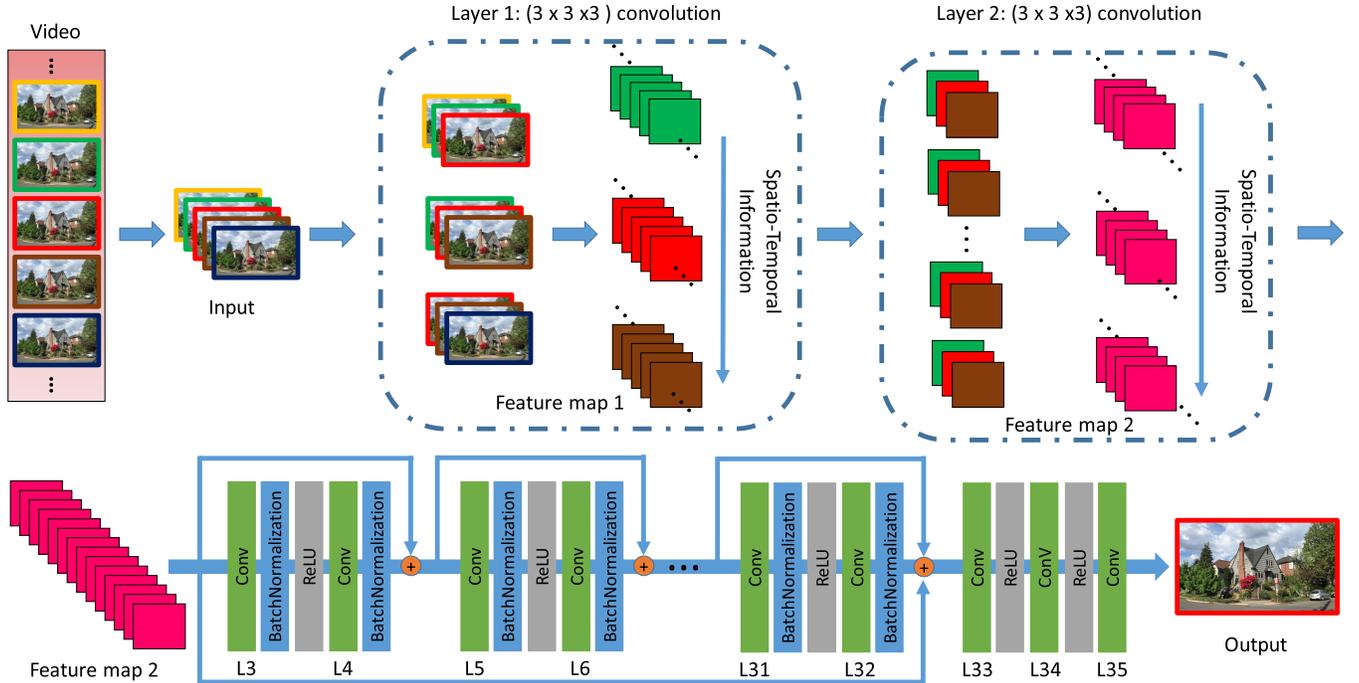


Fig. 2. The proposed DBLRNet framework. The input to our network is five time-consecutive blurry frames. The output is the central deblurred frame. By performing 3D convolutions, this model learns joint spatial-temporal feature representations.

size of feature maps in our model keeps constant. Namely, there is not any down-sampling operation nor up-sampling operation in our DBLRNet. The detailed configurations of DBLRNet is given in Table I.

As Figure 2 shows, the input to DBLRNet is five consecutive frames. Note that we do not conduct deblurring in the original RGB space. Alternatively, we conduct deblurring on basis of gray-scale images. Specifically, the RGB space is transformed to the YCbCr space, and the Y channel is adopted as input since the illumination is the most salient one. We perform 3D convolutions with kernel size of $3 \times 3 \times 3$ (3×3 is the spatial size and the last 3 is for the temporal dimension) in the first and second convolutional layers. To be more specific, in layer 1, three groups of consecutive frames are convolved with a set of 3D kernels respectively, resulting in three groups of feature maps. These three groups of feature maps are convolved with 3D filters again to obtain higher-level feature maps. In the following layers, the size of convolution kernels is $3 \times 3 \times 1$ due to the decrease of temporal dimensions. The stride and padding are set to 1 in every layer. The output of DBLRNet is the deblurred central frame. We transform the gray-scale output back to colorful images with the original Cb and Cr channels.

B. DBLRGAN

GAN is proposed to train generative parametric models by Goodfellow *et al.* [53]. It consists of two networks: a generator network G and a discriminator network D . The goal of G is to generate samples, trying to fool D , while D is trained to distinguish generated samples from real samples. Inspired by the adversarial training strategy, we propose a

TABLE II
CONFIGURATIONS OF OUR D MODEL IN DBLRGAN. BN MEANS BATCH NORMALIZATION AND ReLU REPRESENTS THE ACTIVATION FUNCTION

Layers	1-2	3-5	6-9	10-14	15-16	17
kernel	3×3	3×3	3×3	3×3	FC	FC
channels	64	128	256	512	4096	2
BN	BN	BN	BN	BN	-	-
ReLU	ReLU	ReLU	ReLU	ReLU	-	-

model called DeBLuRing Generative Adversarial Network (DBLRGAN), which utilizes G to deblur images and D to discriminate deblurred images and real-world sharp images. Ideally, the discriminator can be fooled if the generator outputs sharp enough image.

Following the formulation in [53], solving the deblurring problem in the generative adversarial framework leads to the following min-max optimization problem:

$$\min_G \max_D V(G, D) = E_{h \sim p_{train}(h)} [\log(D(h))] + E_{\hat{h} \sim p_{G(\hat{h})}} [\log(1 - D(G(\hat{h})))] \quad (2)$$

where h indicates a sample from real-world sharp frames and \hat{h} represents a blurry sample. G is trained to fool D into misclassifying the generated frames, while D is trained to distinguish deblurred frames from real-world sharp frames. G and D models are trained alternately, and our ultimate goal is to train a model G that recovers sharp frames given blurry frames.

As shown in Figure 3, we use the proposed DBLRNet (Figure 2 and Table I) as our G model, and build a CNN

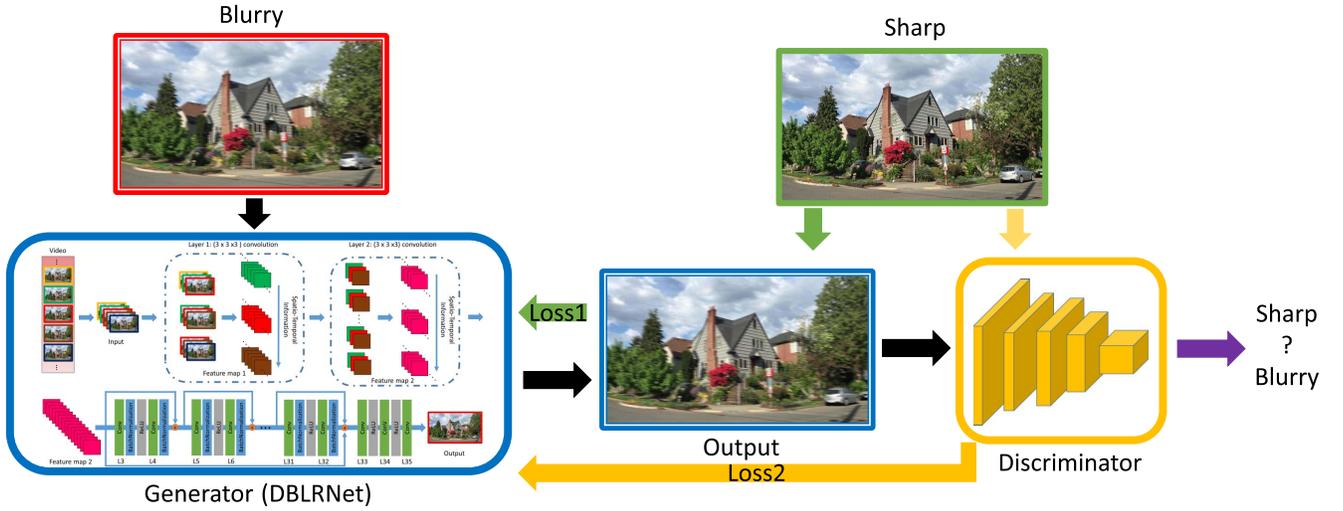


Fig. 3. The DBLRGAN framework for video deblurring. The architecture consists of a Generator and a Discriminator. The Generator is our proposed DBLRNet, while the Discriminator is a VGG-like CNN net.

model as our D model, following the architectural guidelines proposed by Radford *et al.* [54]. This D model is similar to the VGG network [36]. It contains 14 convolutional layers. From bottom to top, the number of channels of the convolutional kernels increases from 64 to 512. Finally, this network is trained via a two-way soft-max classifier at the top layer to distinguish real sharp frames from deblurred ones. For more detailed configurations, please refer to Table II.

C. Loss Functions

In our work, we use two types of loss functions to train DBLRGAN.

D. Content Loss

The Mean Square Error (MSE) loss is widely used in optimization objective for video deblurring in many existing methods. Based on MSE, our content loss function is defined as:

$$\mathcal{L}_{content} = \frac{1}{WH} \sum_{x=1}^W \sum_{y=1}^H (I_{x,y}^{sharp} - G(I^{blurry})_{x,y})^2, \quad (3)$$

where W and H are the width and height of a frame, $I_{x,y}^{sharp}$ is the value of sharp frames at location (x, y) , and $G(I^{blurry})_{x,y}$ corresponds to the value of deblurred frames which are generated from DBLRNet.

E. Adversarial Loss

In order to drive G to generate sharp frames similar to the real-world frames, we introduce an adversarial loss function to update models. During the training stage, parameters of DBLRNet are updated in order to fool the discriminator D . The adversarial loss function can be represented as:

$$\mathcal{L}_{adversarial} = \log(1 - D(G(I^{blurry}))), \quad (4)$$

where $D(G(I^{blurry}))$ is the probability that the recovered frame is a real sharp frame.

F. Balance of Different Loss Functions

In the training stage, the loss functions are combined in a weight fusion fashion:

$$\mathcal{L} = \mathcal{L}_{content} + \alpha \cdot \mathcal{L}_{adversarial}. \quad (5)$$

In order to balance the content and adversarial losses, we use a hyper-parameter α to yield the final loss \mathcal{L} . We investigate different values of α from 0 to 0.1. When $\alpha = 0$, only the content loss works. In this case, DBLRGAN degrades to DBLRNet. With the increase of α , the adversarial loss plays a more and more important role. The value of α should be relative small, because large values of α can degrades the performance of our model.

IV. EXPERIMENTAL RESULTS

In this section, we conduct experiments to demonstrate the effectiveness of the proposed DBLRNet and DBLRGAN on the task of video deblurring.

A. Datasets

1) *VideoDeblurring Dataset*: Su *et al.* [48] build a benchmark which contains videos captured by various kinds of devices such as iPhone 6s, GoPro Hero 4 and Nexus 5x, and each video includes about 100 frames of size 1280×720 . This benchmark consists of two sub datasets: quantitative and qualitative ones. The quantitative subset contains 6708 blurry frames and their corresponding ground-truth sharp frames from 71 videos. The qualitative subset includes 22 scenes, most of which contain more than 100 images. Note that there is not ground truth for the qualitative subset, thus we can only conduct qualitative experiments on this subset. We split the quantitative subset into 61 training videos and 10 testing videos, which is the same setting as the previous method [48]. Besides quantitative experiments on the 10 testing videos, we additionally test our models on the qualitative subset.

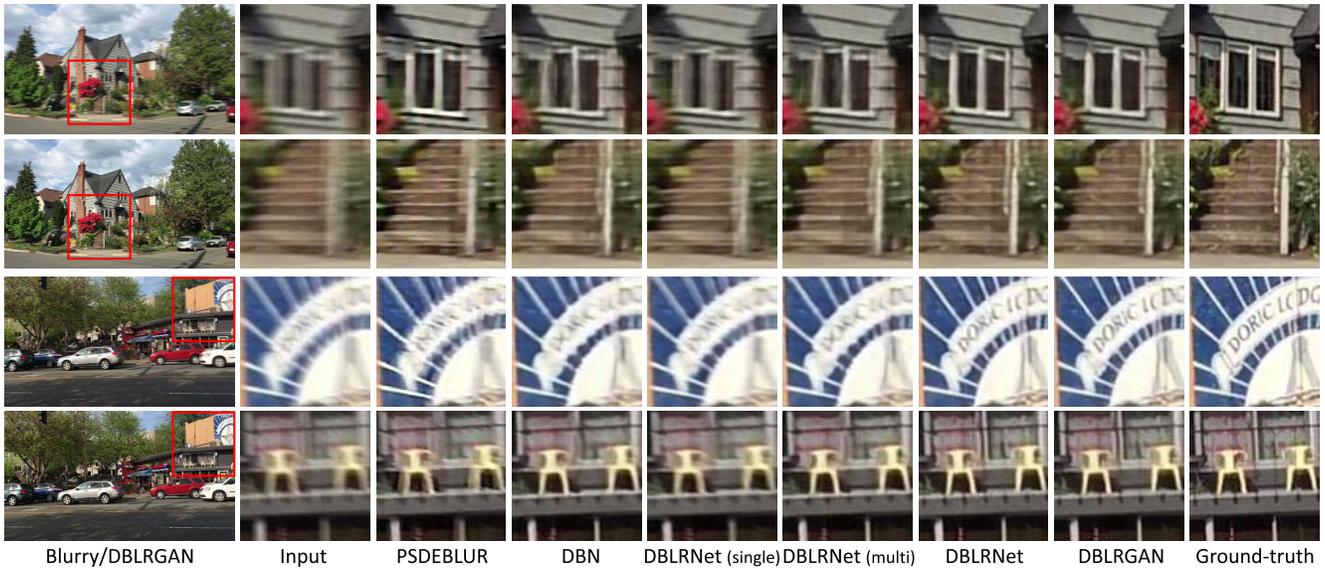


Fig. 4. Exemplar results on the VideoDeblurring dataset (quantitative subset). From left to right: real blurry frame/ Output of DBLRGAN, input, PSDEBLUR, DBN [48], DBLRNet (single), DBLRNet (multi), DBLRNet, DBLRGAN and ground-truth. All results are obtained without alignment. Best viewed in color.

2) *Blurred KITTI Dataset*: Geiger *et al.* [55] develop a dataset called KITTI by using their autonomous driving platform. The KITTI dataset consists of several subsets for various kinds of tasks, such as stereo matching, optical flow estimation, visual odometry, 3D object detection and tracking. Based on the stereo 2015 dataset in the KITTI dataset, Pan *et al.* [33] create a synthetic Blurred KITTI dataset, which contains 199 scenes. Each of the scenes includes 3 images captured by a left camera and 3 images captured by a right camera. It is worthy noting that, the KITTI data set is not used when training our models. Namely, this dataset is utilized only for testing.

B. Implementation Details and Parameters

When training DBLRNet, we use Gaussian distribution with zero mean and a standard deviation of 0.01 to initialize weights. In each iteration, we update all the weights after learning a mini-batch of size 4. To augment the training set, we crop a 128×128 patch at any location of an image (1280×720). In this way, there are at least 712193 possible samples per one frame on the dataset [48], which greatly increases the number of training samples. In addition, we also randomly flip frames in the training stage. The DBLRNet is trained with a learning rate of 10^{-4} , based on the content loss only. We also decrease the learning rate to 10^{-5} when the training loss does not decrease (usually after about 1.5×10^5 iterations), for the sake of additional performance improvement.

In DBLRGAN, we set the hyper parameter α as 0.0002 when we conduct experiments as empirically this value achieves the best performance. It has a better PSNR value due to three reasons. Firstly, when training DBLRGAN, we directly place DBLRNet as our generator and fine-tune our DBLRGAN. Thus, the DBLRGAN has a high PSNR value like DBLRNet at the beginning. Secondly, the loss

functions of DBLRGAN are combined in a weight fusion fashion. We set the hyper parameter α as 0.0002 when we conduct experiments. This is a very small value, which forces the content loss to have an overwhelming superiority over the adversarial loss on PSNR value during the training stage. Thirdly, the learning rate is set as 10^{-5} , so the PSNR value does not have severe changes. We early stop training our DBLRGAN before the PSNR start to drop.

C. Effectiveness of DBLRNet

The proposed DBLRNet has the advantage of learning spatio-temporal feature representations. In order to verify the effectiveness of DBLRNet, we develop another two similar neural networks: DBLRNet (single) and DBLRNet (multi). These two models have the same network architectures as the original DBLRNet while there are two differences between them and the original DBLRNet. The first difference is the input. The input of DBLRNet (single) is one single frame, while the input of DBLRNet (multi) and DBLRNet is a stack of five neighboring frames. The second difference is that, in both DBLRNet (single) and DBLRNet (multi), all the convolution operations are 2D convolution operations.

Table III and IV show the PSNR values of DBLRNet (single), DBLRNet (multi) and DBLRNet on the VideoDeblurring dataset and the Blurred KITTI dataset, respectively. Compared with DBLRNet (single), DBLRNet (multi) achieves approximately 3% ~ 5% improvement of PSNR values, which shows that stacking multiple neighboring frames is useful to learn temporal features for video deblurring even in case of 2D convolution. Comparing DBLRNet with DBLRNet (multi), there are additionally 1% ~ 5% improvement in terms of PSNR. We suspect that the improvement results from the power of spatio-temporal feature representations learned by 3D convolution. Conducting these two kinds of comparisons, the effectiveness of DBLRNet has been verified.

TABLE III

PERFORMANCE COMPARISONS IN TERMS PSNR WITH PSDEBLUR, WFA [32], DBN (SINGLE), DBN (NOALIGN), DBN(FLOW) [48], DBLRNET (SINGLE) AND DBLRNET (MULTI) ON THE VIDEODEBLURRING DATASET. THE BEST RESULTS ARE SHOWN IN BOLD, AND THE SECOND BEST ARE UNDERLINED. ALL RESULTS OF DBLRNET AND DBLRGAN ARE OBTAINED WITHOUT ALIGNING

Methods	1	2	3	4	5	6	7	8	9	10	Average (PSNR)
INPUT	24.14	30.52	28.38	27.31	22.60	29.31	27.74	23.86	30.59	26.98	27.14
PSDEBLUR	24.42	28.77	25.15	27.77	22.02	25.74	26.11	19.71	26.48	24.62	25.08
WFA	25.89	32.33	28.97	28.36	23.99	31.09	28.58	24.78	31.30	28.20	28.35
DBN (single)	25.75	31.15	29.30	28.38	23.63	30.70	29.23	25.62	31.92	28.06	28.37
DBN (noalign)	27.83	33.11	31.29	29.73	25.12	32.52	30.80	27.28	33.32	29.51	30.05
DBN (flow)	28.31	33.14	30.92	29.99	25.58	32.39	30.56	27.15	32.95	29.53	30.05
DBLRNet (single)	28.68	29.40	35.11	32.25	24.94	30.77	29.81	25.67	33.14	30.06	29.98
DBLRNet (multi)	30.40	32.17	36.68	33.38	26.20	32.20	30.71	26.71	36.50	30.65	31.56
DBLRNet	<u>31.96</u>	<u>34.31</u>	37.86	35.21	<u>27.23</u>	<u>33.63</u>	<u>32.32</u>	<u>27.84</u>	<u>38.23</u>	<u>31.83</u>	<u>33.04</u>
DBLRGAN	32.32	34.51	<u>37.63</u>	<u>35.18</u>	27.42	33.81	32.43	28.18	38.32	32.06	33.19



Fig. 5. Exemplar results on the VideoDeblurring dataset (qualitative subset). From left to right: real blurry frame/Output of GBLRGAN, input, PSDEBLUR, DBN [48], Kim *et al.* [28], DTBNet [50], DBLRNet (single), DBLRNet (multi), DBLRNet and DBLRGAN. All results are attained without alignment. Best viewed in color.

D. Effectiveness of DBLRGAN

In this section, we investigate the performance of the proposed DBLRGAN. Table III and IV show the quantitative results on the VideoDeblurring and Blurred KITTI dataset, respectively. Quantitatively, DBLRGAN outperforms DBLRNet with slight advance (about 1% improvement). As have mentioned above, the generator model in DBLRGAN aims to generate frames with similar pixel values as the sharp frames while the discriminator model along with the adversarial loss drives the generator to recover realistic images like real-world images. These two models complement each other and achieve better results. The results in Table III and IV show that the improvement achieved by DBLRNet is more

obvious than GAN model. While according to Figure 5, the deblurred frames generated by DBLRGAN are sharper than DBLRNet, *e.g.*, the word “Bill” in the top row. α should be set as a little value because a bigger α will break the balance of content and adversarial loss, which causes worse performance of video deblurring.

Figure 4 and 5 provide exemplar results on the quantitative and qualitative subsets of the VideoDeblurring dataset, respectively. Please notice the two columns corresponding to DBLRNet and DBLRGAN in Figure 4, especially the letters in the third row, where results of DBLRGAN are more photo-realistic than those of DBLRNet. The same case is observed in Figure 5. Letters in results of DBLRGAN are sharper than those of DBLRNet, which consistently shows that, DBLRGAN



Fig. 6. Performance of our method on blurry videos caused by bokeh. The figure shows a sample frame from the Blurred KITTI dataset, which is captured from a car moving at a high speed. The blurs take place in the side area, while the center part is clear. We show a few pairs of zoomed-in patches from the frame before and after applying our method. The sharper edge demonstrates that our method can generalize well to other types of blurry videos.

generates more realistic frames with finer textural details compared with DBLRNet.

All results of DBLRNet and DBLRGAN are obtained without aligning. Aligning images is computationally expensive and fragile [48]. Kim *et al.* [50] evaluate DBN model and find that the speed of DBN model without aligning is almost more than 20 times faster than it with aligning because aligning procedure is very time-consuming. Our proposed models enable the generation of high quality results without computing any alignment, which makes it highly efficient to scene types.

E. Comparison With Existing Methods

To further verify the effectiveness of our models, we additionally compare the performance of DBLRNet and

DBLRGAN with that of several state-of-the-art approaches on both the VideoDeblurring dataset and the KITTI dataset.

On the VideoDeblurring dataset, we compare our models with PSDEBLUR, WFA [32], DBN [48] and DBN (single). PSDEBLUR is the deblurred results of PHOTOSHOP. WFA is a method based on multiple frames as input. DBN achieves the state-of-the-art performance on the VideoDeblurring data set before this work. DBN (single) is a variant of DBN which stacks 5 copies of one single frame as input. Table III shows quantitative comparisons between our methods and these methods. Specially, the results indicate that our method significantly outperforms the DBN model by 3.14 db. Figure 4 and 5 also represent visual comparison between our models and these methods on both the quantitative (Figure 4)

TABLE IV
PERFORMANCE COMPARISONS WITH [28], [33], AND [56] ON THE
BLURRED KITTI DATASET IN TERMS OF THE PSNR CRITERION.
THE BEST RESULTS ARE SHOWN IN BOLD, AND
THE SECOND BEST ARE UNDERLINED

Methods	PSNR-LEFT	PSNR-RIGHT
Kim et al.	28.25	29.00
Sellent et al.	27.75	28.52
Pan et al.	<u>30.24</u>	<u>30.71</u>
DBLRNet (single)	28.97	29.55
DBLRNet (multi)	29.94	30.33
DBLRNet	30.10	30.54
DBLRGAN	30.42	30.87

and qualitative (Figure 5) sub-datasets, respectively. Evidently our models achieves sharper results.

On the dataset of Blurred KITTI, we conduct comparison with [28], [33], and [56]. Reference [33] is a geometry based method and utilizes additional stereo information from image pairs. It is the current state-of-the-art on the Blurred KITTI dataset. We simply apply the DBLRNet trained on the VideoDeblurring dataset to the Blurred KITTI dataset and still achieve comparable results with [33]. With the additional adversarial loss, DBLRGAN slightly outperforms [33]. Please note that, our models are not specialized for the stereo setting.

F. Different Frames & Other Types of Blur

1) *Different Frames*: We are curious about how the number of consecutive frames influences the performance of our DBLRGAN model. Thus we compare the PSNR values of the model by varying the number of input blurry frames. Making it more specific, on the VideoBlurring dataset, five kinds of settings, three, five, seven, nine and eleven continuous frames are taken as input to our model. Fig. 7 shows that our model with five frames as input achieves the best performance. With the increase of input frames, the PSNR values become lower. We suspect that, as our 3D convolution based network can extract powerful representations to describe short-term fast-varying motions occurring in continuous input frames, it is suitable to set the temporal span relatively small to capture the rapid dynamics across local adjacent frames.

2) *Generalize to Other Types of Blurry Videos*: Though our model is trained on the VideoDeblurring dataset, which includes only blurry frames caused by camera shakes, we are also curious about how it generalize to blurry videos of other blur types. To this end, we test it on videos from the Blurred KITTI dataset. Fig. 6 shows exemplar frames, which is captured by a camera mounted on a high-speed car. The dominated blur is cause by bokeh (see the comparison between the center area and the border area in the image), rather than camera shakes. As shown in the comparison of the enlarged patches, by applying our DBLRGAN model, the edges in the image become sharper. As discussed above, this verifies the advantage of our method capturing short-term fast-varying motions.

3) *Limitation*: Removing jumping artifacts is a challenge of video deblurring. As shown in Fig. 1 (col. 4&5, row 2), there are also some jumping artifacts in the deblurred

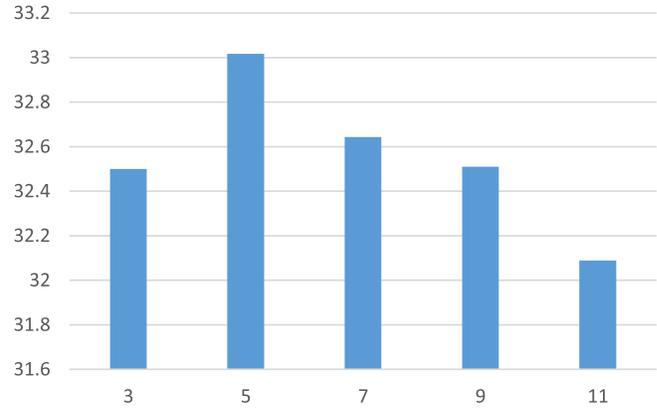


Fig. 7. Performance comparisons of our method in terms of PSNR by varying the number of input frames.

frames. Thus our method cannot solve it completely. However, the proposed model contributes to alleviate the unexpected temporal artifacts because it captures jointly spatial and temporal information encoded in neighboring frames. Even without post-processing and aligning, our proposed model can also achieve satisfied performance. Please refer to Fig. 4 and 5. Comparing with prior methods, when frames are severely blurred, our methods can generate better deblurred frames.

V. CONCLUSIONS

In this paper, we have resorted to spatio-temporal learning and adversarial training to recover sharp and realistic video frames for video deblurring. Specifically, we proposed two novel network models. The first one is our DBLRNet, which uses 3D convolutional kernels on the basis of deep residual neural networks. We demonstrated that DBLRNet is able to capture better spatio-temporal features, leading to improved blur removal. Our second contribution is DBLRGAN equipped with both the content loss and adversarial loss, which are complementary to each other, driving the model to generate visually realistic images. The experimental results on two standard benchmarks show that our proposed DBLRNet and DBLRGAN outperform the existing state-of-the-art methods in video deblurring.

REFERENCES

- [1] S. B. Kang, "Automatic removal of chromatic aberration from a single image," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2007, pp. 1–8.
- [2] J. Sun, W. Cao, Z. Xu, and J. Ponce, "Learning a convolutional neural network for non-uniform motion blur removal," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 769–777.
- [3] J. Shi, L. Xu, and J. Jia, "Just noticeable defocus blur detection and estimation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 657–665.
- [4] H. S. Lee and K. M. Lee, "Dense 3D reconstruction from severely blurred images using a single moving camera," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2013, pp. 273–280.
- [5] H. S. Lee, J. Kwon, and K. M. Lee, "Simultaneous localization, mapping and deblurring," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Nov. 2011, pp. 1203–1210.
- [6] H. Jin, P. Favaro, and R. Cipolla, "Visual tracking in the presence of motion blur," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2005, pp. 18–25.

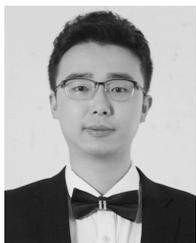
- [7] Z. Wang, E. P. Simoncelli, and A. C. Bovik, "Multiscale structural similarity for image quality assessment," in *Proc. Asilomar Conf. Signals, Syst. Comput. (ACSSC)*, Nov. 2003, pp. 1398–1402.
- [8] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [9] A. Gupta, N. Joshi, C. L. Zitnick, M. Cohen, and B. Curless, "Single image deblurring using motion density functions," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2010, pp. 171–184.
- [10] M. Hirsch, C. J. Schuler, S. Harmeling, and B. Schölkopf, "Fast removal of non-uniform camera shake," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Nov. 2011, pp. 463–470.
- [11] Z. Hu, L. Xu, and M.-H. Yang, "Joint depth estimation and camera shake removal from single blurry image," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2014, pp. 2893–2900.
- [12] M. Jin, S. Roth, and P. Favaro, "Noise-blind image deblurring," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 3834–3842.
- [13] P. P. Srinivasan, R. Ng, and R. Ramamoorthi, "Light field blind motion deblurring," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2354–2362.
- [14] Y. Yan, W. Ren, Y. Guo, R. Wang, and X. Cao, "Image deblurring via extreme channels prior," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 6978–6986.
- [15] J. Pan, J. Dong, Y.-W. Tai, Z. Su, and M.-H. Yang, "Learning discriminative data fitting functions for blind image deblurring," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 1077–1085.
- [16] D. Gong, M. Tan, Y. Zhang, A. van den Hengel, and Q. Shi, "Self-paced kernel estimation for robust blind image deblurring," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 1670–1679.
- [17] J. Dong, J. Pan, Z. Su, and M.-H. Yang, "Blind image deblurring with outlier handling," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2497–2505.
- [18] Y. Bahat, N. Efrat, and M. Irani, "Non-uniform blind deblurring by reblurring," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 3286–3294.
- [19] M. M. R. Mohan, A. N. Rajagopalan, and G. Seetharaman, "Going unconstrained with rolling shutter deblurring," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 4030–4038.
- [20] H. Park and K. M. Lee, "Joint estimation of camera pose, depth, deblurring, and super-resolution from a blurred image sequence," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Sep. 2017, pp. 4613–4621.
- [21] A. Ito, A. C. Sankaranarayanan, A. Veeraraghavan, and R. G. Baraniuk, "BlurBurst: Removing blur due to camera shake using multiple images," *ACM Trans. Graph.*, vol. 3, no. 1, p. 1–15, 2014.
- [22] L. Yuan, J. Sun, L. Quan, and H.-Y. Shum, "Image deblurring with blurred/noisy image pairs," *ACM Trans. Graph.*, vol. 26, no. 3, p. 1, Jul. 2007.
- [23] G. Petschnigg, R. Szeliski, M. Agrawala, M. Cohen, H. Hoppe, and K. Toyama, "Digital photography with flash and no-flash image pairs," *ACM Trans. Graph.*, vol. 23, no. 3, pp. 664–672, Aug. 2004.
- [24] S. H. Park and M. Levoy, "Gyro-based multi-image deconvolution for removing handshake blur," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2014, pp. 3366–3373.
- [25] Y.-W. Tai, H. Du, M. S. Brown, and S. Lin, "Image/video deblurring using a hybrid camera," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2008, pp. 1–8.
- [26] P. Wieschollek, M. Hirsch, B. Schölkopf, and H. P. A. Lensch, "Learning blind motion deblurring," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 231–240.
- [27] W. Ren, J. Pan, X. Cao, and M.-H. Yang, "Video deblurring via semantic segmentation and pixel-wise non-linear kernel," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 1086–1094.
- [28] T. H. Kim and K. M. Lee, "Generalized video deblurring for dynamic scenes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 5426–5434.
- [29] Y. Li, S. B. Kang, N. Joshi, S. M. Seitz, and D. P. Huttenlocher, "Generating sharp panoramas from motion-blurred videos," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2010, pp. 2424–2431.
- [30] N. M. Law, C. D. Mackay, and J. E. Baldwin, "Lucky imaging: High angular resolution imaging in the visible from the ground," *Astron. Astrophys.*, vol. 446, no. 2, pp. 739–745, 2006.
- [31] S. Cho, J. Wang, and S. Lee, "Video deblurring for hand-held cameras using patch-based synthesis," *ACM Trans. Graph.*, vol. 31, no. 4, 2012, Art. no. 64.
- [32] M. Delbracio and G. Sapiro, "Burst deblurring: Removing camera shake through Fourier burst accumulation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 2385–2393.
- [33] L. Pan, Y. Dai, M. Liu, and F. Porikli, "Simultaneous stereo video deblurring and scene flow estimation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 4382–4391.
- [34] S. Ren, K. He, R. Girshick, and J. Sun, "Faster R-CNN: Towards real-time object detection with region proposal networks," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2015, pp. 91–99.
- [35] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2012, pp. 1097–1105.
- [36] K. Simonyan and A. Zisserman. (Apr. 2015). "Very deep convolutional networks for large-scale image recognition." [Online]. Available: <https://arxiv.org/abs/1409.1556>
- [37] C. Szegedy et al., "Going deeper with convolutions," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 1–9.
- [38] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 770–778.
- [39] Y. Sun, X. Wang, and X. Tang, "Deep convolutional network cascade for facial point detection," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2013, pp. 3476–3483.
- [40] Y. Sun, X. Wang, and X. Tang, "Deep learning face representation from predicting 10,000 classes," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2014, pp. 1891–1898.
- [41] K. Zhang, Y. Huang, C. Song, H. Wu, and L. Wang, "Kinship verification with deep convolutional neural networks," in *Proc. Brit. Mach. Vis. Conf. (BMVC)*, 2015, pp. 1–12.
- [42] K. Zhang, Y. Huang, Y. Du, and L. Wang, "Facial expression recognition based on deep evolutionary spatial-temporal networks," *IEEE Trans. Image Process.*, vol. 26, no. 9, pp. 4193–4203, Sep. 2017.
- [43] S. Ji, W. Xu, M. Yang, and K. Yu, "3D convolutional neural networks for human action recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 1, pp. 221–231, Jan. 2013.
- [44] K. Simonyan and A. Zisserman, "Two-stream convolutional networks for action recognition in videos," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2014, pp. 568–576.
- [45] C. Ledig et al., "Photo-realistic single image super-resolution using a generative adversarial network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, p. 4.
- [46] L. Xu, J. S. Ren, C. Liu, and J. Jia, "Deep convolutional neural network for image deconvolution," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2014, pp. 1790–1798.
- [47] S. Nah, T. H. Kim, and K. M. Lee, "Deep multi-scale convolutional neural network for dynamic scene deblurring," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 257–265.
- [48] S. Su, M. Delbracio, J. Wang, G. Sapiro, W. Heidrich, and O. Wang, "Deep video deblurring for hand-held cameras," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 2015–2024.
- [49] T. M. Nimisha, A. K. Singh, and A. N. Rajagopalan, "Blur-invariant deep learning for blind-deblurring," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 4762–4770.
- [50] T. H. Kim, K. M. Lee, B. Schölkopf, and M. Hirsch, "Online video deblurring via dynamic temporal blending network," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 4058–4067.
- [51] T. H. Kim, S. Nah, and K. M. Lee, "Dynamic scene deblurring using a locally adaptive linear blur model," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 40, no. 10, pp. 2374–2387, Oct. 2018, doi: [10.1109/TPAMI.2017.2761348](https://doi.org/10.1109/TPAMI.2017.2761348).
- [52] J. Long, E. Shelhamer, and T. Darrell, "Fully convolutional networks for semantic segmentation," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2015, pp. 3431–3440.
- [53] I. Goodfellow et al., "Generative adversarial nets," in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2014, pp. 2672–2680.
- [54] A. Radford, L. Metz, and S. Chintala, "Unsupervised representation learning with deep convolutional generative adversarial networks," in *Proc. Int. Conf. Learn. Represent. (ICLR)*, 2016.
- [55] A. Geiger, P. Lenz, C. Stiller, and R. Urtasun, "Vision meets robotics: The KITTI dataset," *Int. J. Robot. Res.*, vol. 32, no. 11, pp. 1231–1237, 2013.
- [56] A. Sellent, C. Rother, and S. Roth, "Stereo video deblurring," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2016, pp. 558–575.



Kaihao Zhang received the M.Eng. degree in computer application technology from the University of Electronic Science and Technology of China, Chengdu, China, in 2016. He is currently pursuing the Ph.D. degree with the College of Engineering and Computer Science, The Australian National University, Canberra, ACT, Australia. He was with the National Laboratory of Pattern Recognition, Center for Research on Intelligent Perception and Computing, Institute of Automation, Chinese Academy of Sciences, Beijing, China, for two years, and also with the Tencent AI Laboratory, Shenzhen, China, for one year. His research interests focus on video analysis and facial recognition with deep learning.



Wenhan Luo received the B.E. degree from the Huazhong University of Science and Technology, China, in 2009, the M.E. degree from the Institute of Automation, Chinese Academy of Sciences, China, in 2012, and the Ph.D. degree from the Imperial College London, U.K., in 2016. He is currently a Senior Researcher with the Tencent AI Laboratory, China. His research interests include several topics in computer vision and machine learning, such as motion analysis (especially object tracking), image/video quality restoration, and reinforcement learning.



Yiran Zhong received the M.Eng. degree (Hons.) in information and electronics engineering from The Australian National University, Canberra, Australia, in 2014. He is currently pursuing the Ph.D. degree with the College of Engineering and Computer Science, The Australian National University and Data61, CSIRO, Canberra. He was a Research Assistant for two years. His current research interests include geometric computer vision, machine learning, and deep learning. He was a recipient of the ICIP Best Student Paper Award in 2014.



Lin Ma (M'13) received the B.E. and M.E. degrees in computer science from the Harbin Institute of Technology, Harbin, China, in 2006 and 2008, respectively, and the Ph.D. degree from the Department of Electronic Engineering, The Chinese University of Hong Kong, in 2013. He was a Researcher with the Huawei Noah's Ark Laboratory, Hong Kong, from 2013 to 2016. He is currently a Principal Researcher with the Tencent AI Laboratory, Shenzhen, China. His current research interests lie in the areas of computer vision and multimodal deep learning, specifically for image and language, image/video understanding, and quality assessment.

Dr. Ma was a recipient of the Microsoft Research Asia Fellowship in 2011. He received the Best Paper Award from the Pacific-Rim Conference on Multimedia in 2008. He was the Finalist of the HKIS Young Scientist Award in engineering science in 2012.



Wei Liu received the Ph.D. degree in EECs from Columbia University, New York City, NY, USA. He was a Research Scientist with the IBM T. J. Watson Research Center, Yorktown Heights, NY, USA. He is currently a Distinguished Scientist with the Tencent AI Laboratory and the Director of the Computer Vision Center. He is devoted to research and development in the fields of machine learning, computer vision, information retrieval, and big data. He has authored over 150 peer-reviewed journal and conference papers, including the *Proceedings of the IEEE*, TPAMI, TKDE, IJCV, NIPS, ICML, CVPR, ICCV, ECCV, KDD, SIGIR, SIGCHI, WWW, IJCAI, and AAAI. His research works received a number of awards and honors, such as the 2011 Facebook Fellowship, the 2013 Jury Award for best thesis of Columbia University, the 2016 and 2017 SIGIR Best Paper Award Honorable Mentions, and the 2018 "AI's 10 to Watch" Honor. He currently serves as an associate editor for several international leading AI journals and as an area chair for several international top-tier AI conferences.



Hongdong Li was with NICTA Canberra Labs, prior to 2010, where he involved in the "Australia Bionic Eyes" Project. He is currently a Reader with the Computer Vision Group, The Australian National University. He is also a Chief Investigator of the Australia ARC Centre of Excellence for Robotic Vision. His research interests include 3D vision reconstruction, structure from motion, multi-view geometry, as well as applications of optimization methods in computer vision. He was a recipient of the CVPR Best Paper Award in 2012 and the Marr Prize Honorable Mention in 2017. He received a number of prestigious best paper awards in computer vision and pattern recognition. He is an Associate Editor of the IEEE T-PAMI and served as an Area Chair for recent year ICCV, ECCV, and CVPR. He was a Program Chair of the ACRA 2015–Australia Conference on Robotics and Automation and a Program Co-Chair of the ACCV 2018–Asian Conference on Computer Vision.