

Cascaded Deep Video Deblurring Using Temporal Sharpness Prior

Jinshan Pan, Haoran Bai, and Jinhui Tang*
Nanjing University of Science and Technology



Figure 1. Deblurred result on a real challenging video. Our algorithm is motivated by the success of variational model-based methods. It explores sharpness pixels from adjacent frames by a temporal sharpness prior (see (f)) and restores sharp videos by a cascaded inference process. As our analysis shows, enforcing the temporal sharpness prior in a deep convolutional neural network (CNN) and learning the deep CNN by a cascaded inference manner can make the deep CNN more compact and thus generate better-deblurred results than both the CNN-based methods [27, 32] and variational model-based method [12].

Abstract

We present a simple and effective deep convolutional neural network (CNN) model for video deblurring. The proposed algorithm mainly consists of optical flow estimation from intermediate latent frames and latent frame restoration steps. It first develops a deep CNN model to estimate optical flow from intermediate latent frames and then restores the latent frames based on the estimated optical flow. To better explore the temporal information from videos, we develop a temporal sharpness prior to constrain the deep CNN model to help the latent frame restoration. We develop an effective cascaded training approach and jointly train the proposed CNN model in an end-to-end manner. We show that exploring the domain knowledge of video deblurring is able to make the deep CNN model more compact and efficient. Extensive experimental results show that the proposed algorithm performs favorably against state-of-the-art methods on the benchmark datasets as well as real-world videos.

1. Introduction

Video deblurring, as a fundamental problem in the vision and graphics communities, aims to estimate latent frames from a blurred sequence. As more videos are taken using hand-held and onboard video capturing devices, this problem has received active research efforts within the last decade. The blur in videos is usually caused by camera shake, object motion, and depth variation. Recovering la-

tent frames is highly ill-posed as only the blurred videos are given.

To recover the latent frames from a blurred sequence, conventional methods usually make assumptions on motion blur and latent frames [12, 2, 4, 11, 5, 29]. Among these methods, the motion blur is usually modeled as optical flow [12, 2, 5, 29]. The key success of these methods is to jointly estimate the optical flow and latent frames under the constraints by some hand-crafted priors. These algorithms are physically inspired and generate promising results. However, the assumptions on motion blur and latent frames usually lead to complex energy functions which are difficult to solve.

The deep convolutional neural network (CNN), as one of the most promising approach, has been developed to solve video deblurring. Motivated by the success of deep CNNs in single image deblurring, Su et al. [24] concatenate consecutive frames and develop a deep CNN based on an encoder-decoder architecture to directly estimate the latent frames. Kim et al. [13] develop a deep recurrent network to recurrently restore latent frames by the concatenating multi-frame features. To better capture the temporal information, Zhang et al. [31] develop spatial-temporal 3D convolutions to help latent frame restoration. These methods perform well when the motion blur is not significant and displacement among input frames is small. However, they are less effective for the frames containing significant blur and large displacement as they do not consider the alignment among input frames [6].

*Corresponding author.

To remedy this problem, several methods estimate the alignment among consecutive input frames explicitly [14, 3, 27] or implicitly [32] to restore latent frames using end-to-end trainable deep CNNs. For example, the alignment methods [10] and [19] have been extended to handle video deblurring by [14] and [32]. The methods by [3, 27] explicitly adopt the optical flow or deformable convolution to estimate the alignment among consecutive input frames to help video deblurring. These algorithms show that using better alignment strategies is able to improve the performance of video deblurring. Nevertheless, the main success of these algorithms is due to the use of large-capacity models. These models cannot be generalized well on real cases. We note there exist lots of prior knowledge in variational model-based approaches and have been effective in video deblurring. A natural question is that can we use the domain knowledge in variational model-based approaches to make deep CNN models more compact so that they can improve the accuracy of video deblurring?

To solve this problem, we propose a simple and compact CNN model for video deblurring. Different from the variational model-based methods that warp the consecutive frames to generate blurred frames based on the estimated optical flow, our algorithm warps the adjacent frames into the reference frame so that the consecutive frames align well and thus generating a clearer intermediate latent frame. As the generated intermediate latent frame may contain artifacts and blur effect, we further develop a deep CNN model based on an encoder-decoder architecture to remove artifacts and blur. To better explore the properties of consecutive frames, we develop a temporal sharpness prior to constrain the deep CNN models. However, as our algorithm estimates optical flow from intermediate latent frames as the motion blur information, it requires a feedback loop. To effectively train the proposed algorithm, we develop a cascaded training approach and jointly train the proposed model in an end-to-end manner. Extensive experiments show that the proposed algorithm is able to generate favorable results against state-of-the-art methods as shown in Figure 1.

The main contributions are summarized as follows:

- We propose a simple and compact deep CNN model that simultaneously estimates the optical flow and latent frames for video deblurring.
- To better explore the properties of consecutive frames, we develop a temporal sharpness prior to constrain deep CNN models.
- We quantitatively and qualitatively evaluate the proposed algorithm on benchmark datasets and real-world videos and show that it performs favorably against state-of-the-art methods in terms of accuracy and model size.

2. Related Work

Hand-crafted prior-based methods. Early video or multi-frame deblurring methods [4, 18] usually assume that there exist sharp contents and interpolate them to help the restoration of latent frames. The main success of these methods is due to the use of sharp contents from adjacent frames. However, these methods are less effective for the blur caused by moving objects and usually generate smooth results due to the interpolation.

To overcome this problem, several algorithms [5, 11, 12, 16] formulate the video deblurring by a variational approach. These algorithms first formulate motion blur as optical flow and develop kind of priors to constrain the latent frames and optical flow for video deblurring. Dai and Wu [5] analyze the relations of motion blur and optical flow and alternatively estimate the transparency map, foreground, and background of latent frames. As this method relies on the accuracy of transparency maps, it is further extended by [11], where deblurring process is achieved by alternatively estimating optical flow and latent frames. Kim et al. [12] approximate the motion blur using bidirectional optical flows based on [11]. To deal with more complex motion blur, Gong et al. [7] develop CNNs to estimate optical flow and use the conventional deconvolution algorithm [33] to restore latent frames. In [29], Wulff and Black develop a novel layered model of scenes in motion and restore latent frames layer by layer. These algorithms are based on the physics models, which are able to remove blur and generate decent results. However, the priors imposed on motion blur and latent frames usually lead to complex energy functions which are difficult to solve.

Deep learning based-methods. Due to the success of CNNs based on encoder and decoder architectures in image restoration [23, 17], this kind of network has been widely used in multi-frame [1] or video deblurring [24]. Instead of using 2D convolution, Zhang et al. [31] employ spatial-temporal 3D convolutions to help latent frame restoration. As demonstrated by [6], these methods can be improved using optical flow for alignment. To better use spatial and temporal information, Kim et al. [14] develop an optical flow estimation step for alignment and aggregate information across the neighboring frames to restore latent ones. Wieschollek et al. [28] recurrently use the features from the previous frame in multiple scales based on a recurrent network. In [13], Kim et al. develop a spatial-temporal recurrent network with a dynamic temporal blending layer for latent frame restoration. Zhou et al. extend the kernel prediction network [19] to improve frame alignment. In [27], Wang et al., develop pyramid, cascading, and deformable convolution to achieve better alignment performance. The latent frames are restored by a deep CNN model with temporal and spatial attention strategies. By training the networks in an end-to-end manner, these aforementioned meth-

ods generate promising deblurred results.

We note that the main success of these algorithms on video deblurring is due to the use of large-capacity models. Their generalization ability on real applications is limited as shown in Figure 1. Different from these methods, we explore the simple and well-established principles to make the CNN model more compact instead of enlarging network model capacity for video deblurring.

3. Motivation

To better motivate our work, we first revisit the conventional variational model-based methods.

For the blur process in videos, the i -th blurred image is usually modeled as:

$$\mathbf{B}_i = \frac{1}{2\tau} \int_{t=0}^{\tau} \mathcal{H}_{i \rightarrow i+1}^t(\mathbf{I}_i) + \mathcal{H}_{i \rightarrow i-1}^t(\mathbf{I}_i) dt, \quad (1)$$

where \mathbf{I}_i denotes the i -th clear image; τ denotes the relative exposure time (which also means the camera duty cycle); $\mathcal{H}_{i \rightarrow i+1}^t$ and $\mathcal{H}_{i \rightarrow i-1}^t$ denote the warping functions which warp the frame \mathbf{I}_i into \mathbf{I}_{i+1} and \mathbf{I}_{i-1} . If we denote the bidirectional optical flow at frame i as $\mathbf{u}_{i \rightarrow i+1}$ and $\mathbf{u}_{i \rightarrow i-1}$, $\mathcal{H}_{i \rightarrow i+1}^t(\mathbf{I}_i)$ and $\mathcal{H}_{i \rightarrow i-1}^t(\mathbf{I}_i)$ can be represented as $\mathbf{I}_i(\mathbf{x} + t\mathbf{u}_{i \rightarrow i+1})$ and $\mathbf{I}_i(\mathbf{x} + t\mathbf{u}_{i \rightarrow i-1})$.

Based on the blur model (1), the deblurring process can be achieved by minimizing:

$$\begin{aligned} \mathcal{L}(\mathbf{u}, \mathbf{I}) = & \sum_i \rho_I(\mathcal{W}(\mathbf{I}_i), \mathbf{B}_i) + \varphi(\mathbf{I}_i) \\ & + \sum_i \sum_j \rho_u(\mathbf{I}_i, \mathbf{I}_{i+j}(\mathbf{x} + \mathbf{u}_{i \rightarrow i+j})) + \phi(\mathbf{u}_{i \rightarrow i+j}), \end{aligned} \quad (2)$$

where $\rho_I(\mathcal{W}(\mathbf{I}_i), \mathbf{B}_i)$ denotes the data term w.r.t. $\mathcal{W}(\mathbf{I}_i)$ and \mathbf{B}_i ; $\mathcal{W}(\mathbf{I}_i)$ denotes the integration term in (1); $\rho_u(\mathbf{I}_i(\mathbf{x}), \mathbf{I}_{i+j}(\mathbf{x} + \mathbf{u}_{i \rightarrow i+j}))$ denotes the data term w.r.t. $\mathbf{I}_i(\mathbf{x})$ and $\mathbf{I}_{i+j}(\mathbf{x} + \mathbf{u}_{i \rightarrow i+j})$; $\varphi(\mathbf{I}_i)$ and $\phi(\mathbf{u}_{i \rightarrow i+j})$ denote the constraints on latent image \mathbf{I}_i and optical flow $\mathbf{u}_{i \rightarrow i+j}$.

In the optimization process, most conventional methods (e.g., [12]) estimate the latent image and optical flow by iteratively minimizing:

$$\sum_i \rho_I(\mathcal{W}(\mathbf{I}_i), \mathbf{B}_i) + \varphi(\mathbf{I}_i), \quad (3)$$

and

$$\sum_i \sum_j \rho_u(\mathbf{I}_i, \mathbf{I}_{i+j}(\mathbf{x} + \mathbf{u}_{i \rightarrow i+j})) + \phi(\mathbf{u}_{i \rightarrow i+j}). \quad (4)$$

We note that alternatively minimizing (3) and (4) is able to remove blur. However, the deblurring performance mainly depends on the choice of constraints w.r.t. latent image \mathbf{I}_i and optical flow $\mathbf{u}_{i \rightarrow i+j}$, and it is not trivial to determine proper constraints. In addition, the commonly used constraints usually lead to highly non-convex objective functions which are difficult to solve.

We further note that most deep CNN-based methods directly estimate the sharp videos from blurred input and generate promising results. However, they estimate the warping functions from blurred inputs instead of latent frames and do not explore the domain knowledge of video deblurring, which are less effective for the videos with significant blur effect.

To overcome these problems, we develop an effective algorithm which makes full use of the well-established principles in the variational model-based methods and explores the domain knowledge to make deep CNNs more compact for video deblurring.

4. Proposed Algorithm

The proposed algorithm contains the optical flow estimation module, latent image restoration module, and the temporal sharpness prior. The optical flow estimation module provides motion information for the latent frame restoration, while the latent frame restoration module further facilitates the optical flow estimation so that it makes the estimated flow more accurate. The temporal sharpness prior is able to explore the sharpness pixels from adjacent frames so that it can facilitate better frame restoration. All the modules are jointly trained in a unified framework by an end-to-end manner. In the following, we explain the main ideas for each component in details. For simplicity, we use three adjacent frames to illustrate the main ideas of the proposed algorithm.

4.1. Optical flow estimation

The optical flow estimation module is used to estimate optical flow between input adjacent frames, where the estimated optical flow provides the motion information for the image restoration (3). As demonstrated in [25], the optical flow estimation (4) can be efficiently solved by a deep neural network. We use the PWC-Net [25] as the optical flow estimation algorithm. Given any two intermediate latent frames \mathbf{I}_i and \mathbf{I}_{i+1} , we compute optical flow by:

$$\mathbf{u}_{i \rightarrow i+1} = \mathcal{N}_f(\mathbf{I}_i; \mathbf{I}_{i+1}), \quad (5)$$

where \mathcal{N}_f denotes the optical flow estimation network which takes two images as the input. For any other two frames, the network \mathcal{N}_f shares the same network parameters.

4.2. Latent frame restoration

With the estimated optical flow, we can use variational model (3) to restore latent frames according to existing methods, e.g., [12]. However, solving (3) involves large computation of $\mathcal{W}(\mathbf{I}_i)$ and needs to define the prior on latent frame \mathbf{I}_i , which makes the restoration more complex. We note that the effect of $\mathcal{W}(\mathbf{I}_i)$ (i.e., the blur process (1))

is to generate a blurred frame so that it is closed to the observed input frame \mathbf{B}_i as much as possible. The discretization of (1) can be written as [4]:

$$\mathcal{W}(\mathbf{I}_i) = \frac{1}{1+2\tau} \sum_{d=1}^{\tau} (\mathcal{H}_{i \rightarrow i+1}^d(\mathbf{I}_i) + \mathcal{H}_{i \rightarrow i-1}^d(\mathbf{I}_i) + \mathbf{I}_i(\mathbf{x})). \quad (6)$$

According to the estimated optical flow $\mathbf{u}_{i \rightarrow i+1}$ and $\mathbf{u}_{i \rightarrow i-1}$, if we set τ to be 1, $\mathcal{W}(\mathbf{I}_i)$ can be approximated by:

$$\mathcal{W}(\mathbf{I}_i) = \frac{1}{3} (\mathbf{I}_i(\mathbf{x} + \mathbf{u}_{i \rightarrow i+1}) + \mathbf{I}_i(\mathbf{x} + \mathbf{u}_{i \rightarrow i-1}) + \mathbf{I}_i(\mathbf{x})). \quad (7)$$

Instead of generating a blurred frame, we want to generate clear one according to the estimated optical flow $\mathbf{u}_{i-1 \rightarrow i}$, and $\mathbf{u}_{i+1 \rightarrow i}$ so that $\mathbf{I}_{i+1}(\mathbf{x} + \mathbf{u}_{i+1 \rightarrow i})$ and $\mathbf{I}_{i-1}(\mathbf{x} + \mathbf{u}_{i-1 \rightarrow i})$ can be aligned with $\mathbf{I}_i(\mathbf{x})$ well. Thus, we can use the following formula to update latent frame \mathbf{I}_i :

$$\mathbf{I}_i \leftarrow \frac{1}{3} (\mathbf{I}_{i+1}(\mathbf{x} + \mathbf{u}_{i+1 \rightarrow i}) + \mathbf{I}_{i-1}(\mathbf{x} + \mathbf{u}_{i-1 \rightarrow i}) + \mathbf{I}_i(\mathbf{x})). \quad (8)$$

However, directly using (8) will lead to the results contains significant artifacts due to the misalignment from $\mathbf{I}_{i+1}(\mathbf{x} + \mathbf{u}_{i+1 \rightarrow i})$ and $\mathbf{I}_{i-1}(\mathbf{x} + \mathbf{u}_{i-1 \rightarrow i})$. To avoid this problem and generate high-quality latent frame \mathbf{I}_i , we use $\mathbf{I}_{i+1}(\mathbf{x} + \mathbf{u}_{i+1 \rightarrow i})$ and $\mathbf{I}_{i-1}(\mathbf{x} + \mathbf{u}_{i-1 \rightarrow i})$ as the guidance frames and develop a deep CNN model to restore latent frame \mathbf{I}_i by:

$$\mathbf{I}_i \leftarrow \mathcal{N}_l(\mathcal{C}(\mathbf{I}_{i+1}(\mathbf{x} + \mathbf{u}_{i+1 \rightarrow i}); \mathbf{I}_i(\mathbf{x}); \mathbf{I}_{i-1}(\mathbf{x} + \mathbf{u}_{i-1 \rightarrow i}))), \quad (9)$$

where $\mathcal{C}(\cdot)$ denotes the concatenation operation and \mathcal{N}_l denotes the restoration network. Similar to [25], we use the bilinear interpolation to compute the warped frames.

For the deep CNN model \mathcal{N}_l , we use an encoder-decoder architecture based on [26]. However, we do not use the ConvLSTM module in \mathcal{N}_l . Other network architectures are the same as [26].

4.3. Temporal sharpness prior

As demonstrated in [4], the blur in the video is irregular, and thus there exist some pixels that are not blurred. Following the conventional method [4], we explore these sharpness pixels to help video deblurring.

According to the warped frames $\mathbf{I}_{i+1}(\mathbf{x} + \mathbf{u}_{i+1 \rightarrow i})$ and $\mathbf{I}_{i-1}(\mathbf{x} + \mathbf{u}_{i-1 \rightarrow i})$, if the pixel \mathbf{x} in $\mathbf{I}_i(\mathbf{x})$ is a sharp one, the pixel values of $\mathbf{I}_{i+1}(\mathbf{x} + \mathbf{u}_{i+1 \rightarrow i})$ and $\mathbf{I}_{i-1}(\mathbf{x} + \mathbf{u}_{i-1 \rightarrow i})$ should be close to that of $\mathbf{I}_i(\mathbf{x})$. Thus, we define this criterion as:

$$\mathcal{S}_i(\mathbf{x}) = \exp \left(-\frac{1}{2} \sum_{j \neq i} \mathcal{D}(\mathbf{I}_{i+j}(\mathbf{x} + \mathbf{u}_{i+j \rightarrow i}); \mathbf{I}_i(\mathbf{x})) \right), \quad (10)$$

where $\mathcal{D}(\mathbf{I}_{i+j}(\mathbf{x} + \mathbf{u}_{i+j \rightarrow i}); \mathbf{I}_i(\mathbf{x}))$ is defined as $\|\mathbf{I}_{i+j}(\mathbf{x} + \mathbf{u}_{i+j \rightarrow i}) - \mathbf{I}_i(\mathbf{x})\|^2$. Based on (10), if the value of $\mathcal{S}_i(\mathbf{x})$ is close to 1, the pixel \mathbf{x} is likely to be clear. Thus, we can use $\mathcal{S}_i(\mathbf{x})$ to help the deep neural network to distinguish whether the pixel is clear or not so that it can help the latent frame restoration. To increase the robustness of $\mathcal{S}_i(\mathbf{x})$, we define $\mathcal{D}(\mathbf{I}_{i+j}(\mathbf{x} + \mathbf{u}_{i+j \rightarrow i}); \mathbf{I}_i(\mathbf{x}))$ as

$$\mathcal{D}(\mathbf{I}_{i+j}(\mathbf{x} + \mathbf{u}_{i+j \rightarrow i}); \mathbf{I}_i(\mathbf{x})) = \sum_{\mathbf{y} \in \omega(\mathbf{x})} \|\mathbf{I}_{i+j}(\mathbf{y} + \mathbf{u}_{i+j \rightarrow i}) - \mathbf{I}_i(\mathbf{y})\|^2, \quad (11)$$

where $\omega(\mathbf{x})$ denotes an image patch centered at pixel \mathbf{x} . With the temporal sharpness prior $\mathcal{S}_i(\mathbf{x})$, we modify the latent frame restoration (9) by

$$\mathbf{I}_i \leftarrow \mathcal{N}_l(\mathcal{C}(\mathcal{C}_i; \mathcal{S}_i(\mathbf{x}))), \quad (12)$$

where $\mathcal{C}_i = \mathcal{C}(\mathbf{I}_{i+1}(\mathbf{x} + \mathbf{u}_{i+1 \rightarrow i}); \mathbf{I}_i(\mathbf{x}); \mathbf{I}_{i-1}(\mathbf{x} + \mathbf{u}_{i-1 \rightarrow i}))$. We will show that using $\mathcal{S}_i(\mathbf{x})$ is able to help latent frame restoration.

4.4. Inference

As the proposed algorithm contains intermediate the optical flow estimation, latent frame estimation, and temporal sharpness computation, we train the proposed algorithm in a cascaded manner.

Let $\Theta_t = \{\mathcal{O}_t, \mathcal{L}_t\}$ denote the model parameters of optical flow estimation and latent frame restoration networks at stage (iteration) t . We learn the stage-dependent model parameters Θ_t from N training video sequences, where each video sequence contains $\{\mathbf{B}_i^n, \mathbf{I}_{gt,i}^n\}_{i=1}^M$ training samples. Given $2j+1$ blurred frames, the parameter Θ_t is learned by minimizing the cost function:

$$\mathcal{J}(\Theta_t) = \sum_{n=1}^N \sum_{i=1}^M \|\mathcal{F}_{\Theta_t}(\mathbf{B}_{i-j}^n; \dots; \mathbf{B}_i^n; \dots; \mathbf{B}_{i+j}^n) - \mathbf{I}_{gt,i}^n\|_1, \quad (13)$$

where \mathcal{F}_{Θ_t} denotes the whole network for video deblurring, which takes $2j+1$ blurred frames as the input. That is, the intermediate latent frame at t -stage is $\mathbf{I}_i^t = \mathcal{F}_{\Theta_t}(\mathbf{B}_{i-j}^n; \dots; \mathbf{B}_i^n; \dots; \mathbf{B}_{i+j}^n)$.

Algorithm 1 summarizes the main steps of the cascaded training approach, where T denotes the number of stages.

5. Experimental Results

In this section, we evaluate the proposed algorithm using publicly available benchmark datasets and compare it to state-of-the-art methods.

5.1. Parameter settings and training data

For fair comparisons with state-of-the-art methods, we use the video deblurring dataset by Su et al. [24] for training and evaluation, where 61 videos are used for training

Algorithm 1 Proposed cascaded training algorithm.

Input: Training video sequences $\{\mathbf{B}_i^n, \mathbf{I}_{gt,i}^n\}_{i=1}^M$; $n = 1, \dots, N$.
Initialize $\mathbf{I}_i^n \leftarrow \mathbf{B}_i^n$.
for $t = 1 \rightarrow T$ **do**
 for Any three frames $\mathbf{I}_{i-1}^n, \mathbf{I}_i^n$, and \mathbf{I}_{i+1}^n **do**
 Estimating optical flow $\mathbf{u}_{i-1 \rightarrow i}, \mathbf{u}_{i+1 \rightarrow i}$ according to (5).
 Computing $\mathcal{S}_i(\mathbf{x})$ according to (10).
 Latent frame restoration according to (12).
 end for
 Estimating model parameters Θ_t by minimizing (13).
 Updating \mathbf{I}_i^n according to (12) with the estimated parameter Θ_t .
end for
Output: Model parameters $\{\Theta_t\}_{t=1}^T$.

and the remaining 10 videos for the test. We use the similar data augmentation method to [32] to generate training data. The size of each image patch is 256×256 pixels. We initialize the latent frame restoration network according to [8] and train it from scratch. For the PWC-Net, we use the pre-trained model [25] to initialize it. In the training process, we use the ADAM optimizer [15] with parameters $\beta_1 = 0.9$, $\beta_2 = 0.999$, and $\epsilon = 10^{-8}$. The minibatch size is set to be 8. The learning rates for \mathcal{N}_l and PWC-Net are initialized to be 10^{-4} and 10^{-6} and decrease to half after every 200 epochs. We empirically set $T = 2$ as a trade-off between accuracy and speed. At each stage, we use 3 frames to generate one deblurred image. Thus, the proposed algorithm needs 5 frames when $T = 2$. To better make the network compact, the network at each stage shares the same model parameters. Similar to [30], we further use the hard example mining strategy to preserve sharp edges. We implement our algorithm based on the PyTorch. More experimental results are included in the supplemental material. The training code and test model are available at the authors' website.

5.2. Comparisons with the state of the art

To evaluate the performance of the proposed algorithm, we compare it against state-of-the-art algorithms including the variational model-based method [12] and deep CNNs-based methods [24, 7, 27, 13, 14, 32, 26]. To evaluate the quality of each restored image on synthetic datasets, we use the PSNR and SSIM as the evaluation metrics.

Table 1 shows the quantitative results on the benchmark dataset by Su et al. [24], where the proposed algorithm performs favorably against the state-of-the-art methods in terms of PSNR and SSIM.

Figure 2 shows some deblurred results from the test dataset [24]. The variational model-based method [12] does not recover the structures well and generates the results

with significant blur residual. The method [26] develops end-to-end-trainable deep CNN models to deblur dynamic scenes. However, the deblurred images contain significant blur residual as the temporal information is not used. The video deblurring algorithm [24] directly concatenates consecutive frames as the input of an end-to-end trainable deep CNN model. However, the structures of the deblurred image are not sharp (Figure 2(e)). We note that the EDVR method [27] develops a pyramid, cascading, and deformable alignment module and uses a PreDeblur module for video deblurring. However, this method is less effective when the PreDeblur does not remove blur from input frames. The results in Figure 2(f) show that the structures of the images by the EDVR method are not recovered well. In contrast, the proposed method recovers finer image details and structures than the state-of-the-art algorithms.

We further evaluate the proposed method on the GOPRO dataset by Nah et al. [20] following the protocols of state-of-the-art methods [13, 21]. Table 2 shows that the proposed algorithm generates the deblurred videos with higher PSNR and SSIM values.

Figure 3 shows some deblurred results from [20]. We note that state-of-the-art methods do not generate sharp images and remove the non-uniform blur well. In contrast, the proposed algorithm restores much clearer images, where the license numbers are recognizable.

Real examples. We further evaluate our algorithm on the real video deblurring dataset by Cho et al. [4]. Figure 4 shows that the state-of-the-art methods [22, 4, 12, 32, 27, 24] do not restore the sharp frames well. Our algorithm generates much clearer frames with better detailed structures. For example, the man and the boundaries of the buildings are much clearer (Figure 4(h)).

6. Analysis and Discussions

We have shown that the proposed algorithm performs favorably against state-of-the-art methods. To better understand the proposed algorithm, we perform further analysis and discuss its limitations.

6.1. Effectiveness of the cascaded training

The proposed cascaded training algorithm ensures that the proposed method estimates optical flow from intermediate latent frames and updates the intermediate latent frames iteratively. One may wonder whether the cascaded training algorithm helps video deblurring. To answer this question, we compare the method without using cascaded training algorithm (i.e., w/o CT in Table 3), where we set stage number T to be 1 in Algorithm 1 for fair comparisons.

Table 3 shows the quantitative evaluations on the benchmark dataset by Su et al. [24]. We note that the method without using cascaded training algorithm estimates the optical flow from blurred inputs using PWC-Net, where this strategy is widely for image alignment in video deblur-

Table 1. Quantitative evaluations on the video deblurring dataset [24] in terms of PSNR and SSIM. All the comparison results are generated using the publicly available code. All the restored frames instead of randomly selected 30 frames from each test set [24] are used for evaluations.

Methods	Kim and Lee [12]	Gong et al. [7]	Tao et al. [26]	Su et al. [24]	Kim et al. [13]	EDVR [27]	STFAN [32]	Ours
PSNRs	26.94	28.27	29.98	30.01	29.95	28.51	31.15	32.13
SSIMs	0.8158	0.8463	0.8842	0.8877	0.8692	0.8637	0.9049	0.9268

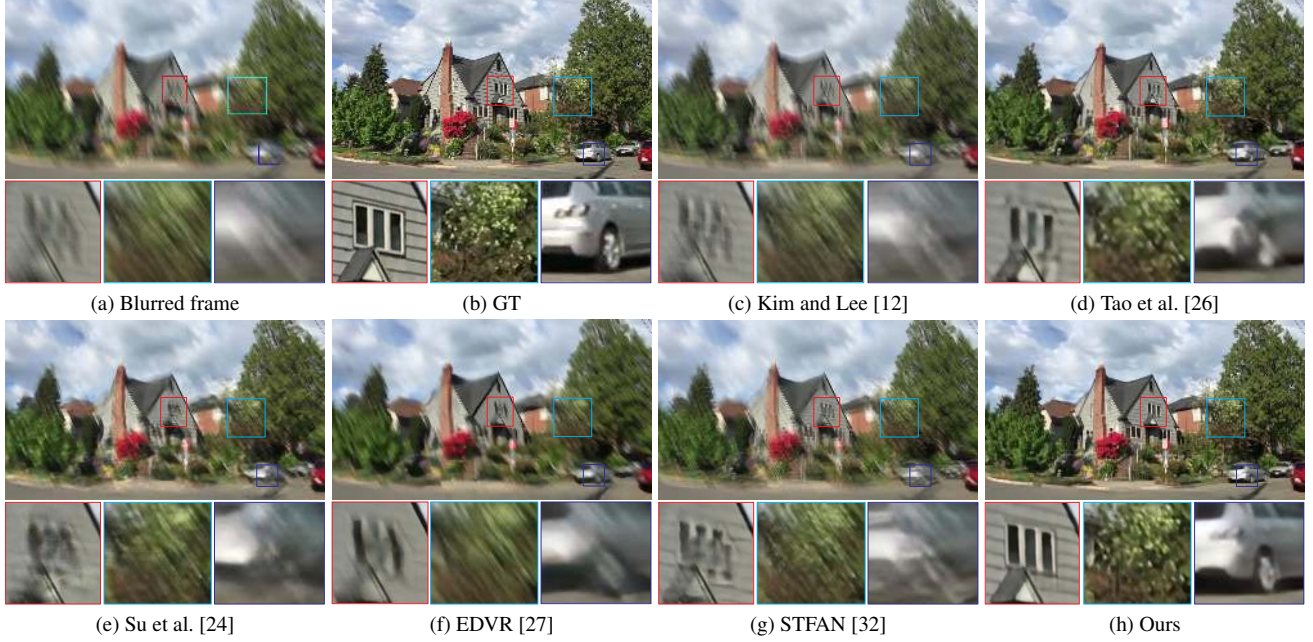


Figure 2. Deblurred results on the test dataset [24]. The deblurred results in (c)-(g) still contain significant blur effects. The proposed algorithm generates much clearer frames.



Figure 3. Deblurred results on the test dataset [20]. The proposed method generates much better deblurred images, where the license numbers are recognizable.

ring [6]. However, this method does not generate high-quality deblurred results (Figure 5(b)) as optical flow is related to the latent frames information instead of blurred ones

during the exposure time. In contrast, the proposed algorithm generates the results with higher PSNR and SSIM values.

We further compare the deblurred results generated by

Table 2. Quantitative evaluations on the video deblurring dataset [20] in terms of PSNR and SSIM. * denotes the reported results from [21].

Methods	Tao et al. [26]	Su et al. [24]	Wieschollek et al. [28]*	Kim et al. [13]*	Nah et al. [21]*	EDVR [27]	STFAN [32]	Ours
PSNRs	30.29	27.31	25.19	26.82	29.97	26.83	28.59	31.67
SSIMs	0.9014	0.8255	0.7794	0.8245	0.8947	0.8426	0.8608	0.9279



Figure 4. Deblurred results on a real video from [4]. The proposed algorithm recovers a high-quality image with clearer details.

Table 3. Effectiveness of the cascaded training algorithm for video deblurring, where “CT” is the abbreviation of cascaded training.

Methods	w/o CT	Stage 1	Stage 2	Stage 3
PSNRs	31.33	31.59	32.13	32.20
SSIMs	0.9125	0.9161	0.9268	0.9272

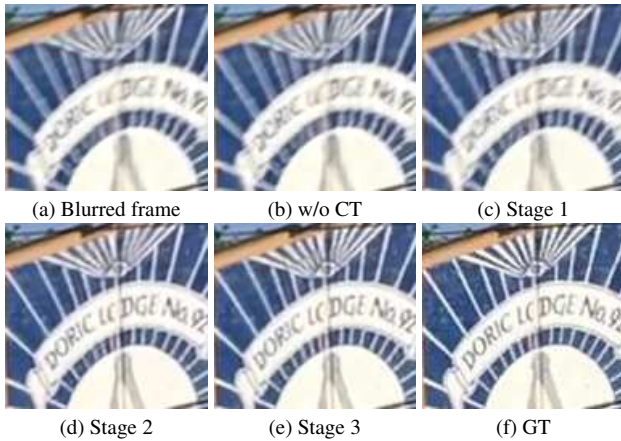


Figure 5. Effectiveness of the cascaded training algorithm for video deblurring. (b) denotes the deblurred result by the proposed method without using cascaded training. (c)-(e) denote the results from stage 1, 2, and 3, respectively.

different stages in Table 3 and Figure 5. We note that using more stages generates better deblurred images. However, the improved performance is not significant. Thus, we use two stages as a trade-off between accuracy and speed.

We further note that directly estimating optical flow from




Figure 6. Effect of optical flow on video deblurring. The optical flow by the proposed method contains sharp boundaries well (see (f)), which facilitates the latent frame restoration.

blurred inputs will increase ambiguity at frame boundaries for video deblurring. Figure 6(d) demonstrates that the boundaries of the estimated optical flow are blurry, which accordingly affects the important boundaries restoration (Figure 6(b)). In contrast, the optical flow by the proposed method contains sharp boundaries well (Figure 6(f)), which facilitates the latent frame restoration (Figure 6(c)).

6.2. Effectiveness of the temporal sharpness prior

We develop a temporal sharpness prior to better explore the properties of consecutive frames so that it makes the deep CNN models more compact. To demonstrate the effectiveness of this prior, we disable this prior in the proposed method and retrain the algorithm without using the temporal sharpness prior with the same settings for fair comparisons.

Table 4. Effectiveness of the temporal sharpness prior on video deblurring.

Methods	w/o temporal sharpness prior	Ours
PSNRs	34.48	34.63
SSIMs	0.9126	0.9268




Figure 7. Effectiveness of the temporal sharpness prior. (a) Blurred input. (b) Visualizations of the intermediate temporal sharpness prior. (c)-(d) denote the results without and with the temporal sharpness prior, respectively.

Table 5. Effect of the optical flow estimation module.

Methods	w/o optical flow	w/ FlowNet 2.0	w/ PWC-Net
PSNRs	31.19	32.06	32.13
SSIMs	0.9055	0.9254	0.9268

We evaluate the temporal sharpness prior on 4 videos with significant blur effects from the test dataset [24]. Table 4 and Figure 7 show both quantitative and qualitative evaluations. We note that the temporal sharpness prior is able to distinguish the sharpness pixels and blurred pixels from adjacent frames so that it can help the deep CNN model for better frame restoration. Figure 7(b) shows the visualizations of the temporal sharpness prior, where the blurred pixels can be better detected. The comparisons in Table 4 demonstrate that using the temporal sharpness prior is able to improve the accuracy of video deblurring. Figure 7 further shows that using the temporal sharpness prior is able to generate the frames with clearer structures.

6.3. Effect of optical flow

As several algorithms either directly concatenate consecutive frames [24] or estimate filter kernels [32] instead of using optical flow for video deblurring, one may wonder whether optical flow helps video deblurring. To answer this question, we remove the optical flow estimation module and compare with the method that directly concatenates consecutive frames as the input of the restoration network \mathcal{N}_l^1 . Table 5 shows that using optical flow is able to improve the performance of video deblurring.

In addition, we further evaluate the optical flow estima-

¹The proposed method without using optical flow reduces to the network \mathcal{N}_l which takes the concatenation of consecutive frames as the input.

Table 6. Comparisons of model sizes against state-of-the-art methods and baselines. “TSP” is the abbreviation of temporal sharpness prior.

Methods	Su et al. [24]	EDVR [27]	w/o CT	w/o TSP	Ours
Model size	15.30M	23.60M	16.19M	16.19M	16.19M

Table 7. Evaluations on the blurred videos, where the blur exists in each position of each frame. The temporal sharpness prior is less effective when it fails to identify the clear pixels.

Methods	w/o temporal sharpness prior	Ours
PSNRs	31.31	31.33
SSIMs	0.9238	0.9239

tion module using FlowNet 2.0 [9]. Table 5 shows that the proposed method is robust to optical flow modules.

6.4. Model size

As stated in Section 1, we aim to improve the accuracy of video deblurring while do not increase the model capacity using domain knowledge of video deblurring. Table 6 shows that the proposed algorithm has a relatively smaller model size against state-of-the-art methods. Compared to the baseline models, the proposed model does not increase any model size while generating much better results.

6.5. Limitations

Although the temporal sharpness prior is effective for videos with significant blur, it is less effective when the blur exists in each position of all frames. In such cases, the temporal sharpness prior is less likely to distinguish whether the pixel is clearer or not. Table 7 shows the deblurred results on 3 videos from the test dataset [24], where each position in a frame contains blur effect. We note that using the temporal sharpness prior does not improve the deblurring performance significantly.

7. Concluding Remarks

We have proposed a simple and effective deep CNN model for video deblurring. The proposed CNN explores the simple and well-established principles used in the variational model-based methods and mainly consists of optical flow estimation from intermediate latent frames and latent frame restoration. We have developed a temporal sharpness prior to help the latent image restoration and an effective cascaded training approach to train the proposed CNN model. By training in an end-to-end manner, we have shown that the proposed CNN model is more compact and efficient and performs favorably against state-of-the-art methods on both benchmark datasets and real-world videos.

Acknowledgments. This work has been supported in part by the National Natural Science Foundation of China (Nos. 61922043, 61872421, 61732007), the Natural Science Foundation of Jiangsu Province (No. BK20180471), and National Key R&D Program of China (No. 2018AAA0102002).

References

- [1] Miika Aittala and Frédo Durand. Burst image deblurring using permutation invariant convolutional neural networks. In *ECCV*, pages 748–764, 2018. [2](#)
- [2] Leah Bar, Benjamin Berkels, Martin Rumpf, and Guillermo Sapiro. A variational framework for simultaneous motion estimation and restoration of motion-blurred video. In *ICCV*, pages 1–8, 2007. [1](#)
- [3] Huaijin Chen, Jinwei Gu, Orazio Gallo, Ming-Yu Liu, Ashok Veeraraghavan, and Jan Kautz. Reblur2deblur: Deblurring videos via self-supervised learning. In *ICCP*, pages 1–9, 2018. [2](#)
- [4] Sunghyun Cho, Jue Wang, and Seungyong Lee. Video deblurring for hand-held cameras using patch-based synthesis. *ACM TOG*, 31(4):64:1–64:9, 2012. [1](#), [2](#), [4](#), [5](#), [7](#)
- [5] Shengyang Dai and Ying Wu. Motion from blur. In *CVPR*, 2008. [1](#), [2](#)
- [6] Jochen Gast and Stefan Roth. Deep video deblurring: The devil is in the details. In *ICCV Workshop*, 2019. [1](#), [2](#), [6](#)
- [7] Dong Gong, Jie Yang, Lingqiao Liu, Yanning Zhang, Ian D. Reid, Chunhua Shen, Anton van den Hengel, and Qinfeng Shi. From motion blur to motion flow: A deep learning solution for removing heterogeneous motion blur. In *CVPR*, pages 3806–3815, 2017. [2](#), [5](#), [6](#)
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *ICCV*, pages 1026–1034, 2015. [5](#)
- [9] Eddy Ilg, Nikolaus Mayer, Tonmoy Saikia, Margret Keuper, Alexey Dosovitskiy, and Thomas Brox. FlowNet 2.0: Evolution of optical flow estimation with deep networks. In *CVPR*, pages 1647–1655, 2017. [8](#)
- [10] Max Jaderberg, Karen Simonyan, Andrew Zisserman, and Koray Kavukcuoglu. Spatial transformer networks. In *NeurIPS*, pages 2017–2025, 2015. [2](#)
- [11] Tae Hyun Kim and Kyoung Mu Lee. Segmentation-free dynamic scene deblurring. In *CVPR*, pages 2766–2773, 2014. [1](#), [2](#)
- [12] Tae Hyun Kim and Kyoung Mu Lee. Generalized video deblurring for dynamic scenes. In *CVPR*, pages 5426–5434, 2015. [1](#), [2](#), [3](#), [5](#), [6](#), [7](#)
- [13] Tae Hyun Kim, Kyoung Mu Lee, Bernhard Schölkopf, and Michael Hirsch. Online video deblurring via dynamic temporal blending network. In *ICCV*, pages 4058–4067, 2017. [1](#), [2](#), [5](#), [6](#), [7](#)
- [14] Tae Hyun Kim, Mehdi S. M. Sajjadi, Michael Hirsch, and Bernhard Schölkopf. Spatio-temporal transformer network for video restoration. In *ECCV*, pages 111–127, 2018. [2](#), [5](#)
- [15] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. [5](#)
- [16] Yunpeng Li, Sing Bing Kang, Neel Joshi, Steven M. Seitz, and Daniel P. Huttenlocher. Generating sharp panoramas from motion-blurred videos. In *CVPR*, pages 2424–2431, 2010. [2](#)
- [17] Xiao-Jiao Mao, Chunhua Shen, and Yu-Bin Yang. Image restoration using very deep convolutional encoder-decoder networks with symmetric skip connections. In *NeurIPS*, pages 2802–2810, 2016. [2](#)
- [18] Yasuyuki Matsushita, Eyal Ofek, Weina Ge, Xiaoou Tang, and Heung-Yeung Shum. Full-frame video stabilization with motion inpainting. *IEEE TPAMI*, 28(7):1150–1163, 2006. [2](#)
- [19] Ben Mildenhall, Jonathan T. Barron, Jiawen Chen, Dillon Sharlet, Ren Ng, and Robert Carroll. Burst denoising with kernel prediction networks. In *CVPR*, pages 2502–2510, 2018. [2](#)
- [20] Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *CVPR*, pages 257–265, 2017. [5](#), [6](#), [7](#)
- [21] Seungjun Nah, Sanghyun Son, and Kyoung Mu Lee. Recurrent neural networks with intra-frame iterations for video deblurring. In *CVPR*, pages 8102–8111, 2019. [5](#), [7](#)
- [22] Jinshan Pan, Deqing Sun, Hanspeter Pfister, and Ming-Hsuan Yang. Deblurring images via dark channel prior. *IEEE TPAMI*, 40(10):2315–2328, 2018. [5](#), [7](#)
- [23] Deepak Pathak, Philipp Krähenbühl, Jeff Donahue, Trevor Darrell, and Alexei A. Efros. Context encoders: Feature learning by inpainting. In *CVPR*, pages 2536–2544, 2016. [2](#)
- [24] Shuochen Su, Mauricio Delbracio, Jue Wang, Guillermo Sapiro, Wolfgang Heidrich, and Oliver Wang. Deep video deblurring for hand-held cameras. In *CVPR*, pages 237–246, 2017. [1](#), [2](#), [4](#), [5](#), [6](#), [7](#), [8](#)
- [25] Deqing Sun, Xiaodong Yang, Ming-Yu Liu, and Jan Kautz. PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *CVPR*, pages 8934–8943, 2018. [3](#), [4](#), [5](#)
- [26] Xin Tao, Hongyun Gao, Xiaoyong Shen, Jue Wang, and Ji-aya Jia. Scale-recurrent network for deep image deblurring. In *CVPR*, pages 8174–8182, 2018. [4](#), [5](#), [6](#), [7](#)
- [27] Xintao Wang, Kelvin C.K. Chan, Ke Yu, Chao Dong, and Chen Change Loy. EDVR: Video restoration with enhanced deformable convolutional networks. In *CVPR Workshops*, 2019. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [28] Patrick Wieschollek, Michael Hirsch, Bernhard Schölkopf, and Hendrik P. A. Lensch. Learning blind motion deblurring. In *ICCV*, pages 231–240, 2017. [2](#), [7](#)
- [29] Jonas Wulff and Michael Julian Black. Modeling blurred video with layers. In *ECCV*, pages 236–252, 2014. [1](#), [2](#)
- [30] Rui Xu, Xiaoxiao Li, Bolei Zhou, and Chen Change Loy. Deep flow-guided video inpainting. In *CVPR*, pages 3723–3732, 2019. [5](#)
- [31] Kaihao Zhang, Wenhan Luo, Yiran Zhong, Lin Ma, Wei Liu, and Hongdong Li. Adversarial spatio-temporal learning for video deblurring. *IEEE TIP*, 28(1):291–301, 2019. [1](#), [2](#)
- [32] Shangchen Zhou, Jiawei Zhang, Jinshan Pan, Haozhe Xie, Wangmeng Zuo, and Jimmy Ren. Spatio-temporal filter adaptive network for video deblurring. In *ICCV*, 2019. [1](#), [2](#), [5](#), [6](#), [7](#), [8](#)
- [33] Daniel Zoran and Yair Weiss. From learning models of natural image patches to whole image restoration. In *ICCV*, pages 479–486, 2011. [2](#)