

Real-Time Video Deblurring via Lightweight Motion Compensation

Hyeongseok Son^{†*}, Junyong Lee[†], Sunghyun Cho[†], and Seungyong Lee[‡]

POSTECH

Abstract

While motion compensation greatly improves video deblurring quality, separately performing motion compensation and video deblurring demands huge computational overhead. This paper proposes a real-time video deblurring framework consisting of a lightweight multi-task unit that supports both video deblurring and motion compensation in an efficient way. The multi-task unit is specifically designed to handle large portions of the two tasks using a single shared network and consists of a multi-task detail network and simple networks for deblurring and motion compensation. The multi-task unit minimizes the cost of incorporating motion compensation into video deblurring and enables real-time deblurring. Moreover, by stacking multiple multi-task units, our framework provides flexible control between the cost and deblurring quality. We experimentally validate the state-of-the-art deblurring quality of our approach, which runs at a much faster speed compared to previous methods and show practical real-time performance (30.99dB@30fps measured on the DVD dataset).

CCS Concepts

• **Computing methodologies** → **Computational photography**;

1. Introduction

Videos captured in dynamic environments often suffer from motion blur caused by camera shakes and moving objects, which degrades not only the perceptual quality of a video but also the performance of visual recognition tasks [MHS*16, KBM*18, PHZ*19]. To resolve this problem, various video deblurring methods have been proposed and focused on improving deblurring quality while reducing computational costs.

For successful video deblurring, it is essential to utilize temporal information from neighboring frames [CWL12, DS15, KL15]. For an effective aggregation of temporal information, many approaches align neighboring frames or their features by adopting a separate motion compensation module, such as optical flow [SDW*17, ZYLP19, PBT20, SLL*21, LXZ*21], a spatial transformer network [KSHS18], dynamic convolution [ZZP*19], and deformable convolution [WCY*19]. While motion compensation improves the deblurring performance as shown in these approaches, heavy computation is required to align blurry frames, resulting in significantly increased computation time, model size, and memory consumption for video deblurring (green vs. blue circles in Fig. 1). Some approaches pursue efficiency over deblurring quality by omitting the motion compensation module [KLSH17, NSL19, ZYZB20]. Nonetheless, these methods can still not perform real-time video deblurring because large network capacity is demanded for aggregating

useful temporal information from unaligned blurry frames or features as well as deblurring (green circles in Fig. 1).

This paper proposes an efficient video deblurring framework that reduces the computational overhead but still provides enough performance of motion compensation needed for exploiting temporal information in video frames. The key component of the proposed framework is a *lightweight multi-task unit* that handles large portions of both deblurring and motion compensation tasks together using a single shared network.

To maximize the efficiency of our framework, both deblurring and motion compensation tasks need to take full advantage of the shared network in a harmonious way. However, we found that each task requires different natures of features, which limits the direct sharing of features. For restoration tasks including deblurring, extracting detail information from a degraded input has shown to be effective in achieving high performance, for which residual (detail) learning has widely been adopted [KLL16, SSH17, JOKK18, SL17, SDW*17, ZZP*19, WCY*19]. However, for the motion compensation task, we observed that structural information of a degraded input is required in addition to the detail information (Sec. 3.1).

To this end, we specifically design the multi-task unit to improve the compatibility of shared features for video deblurring and motion compensation tasks, and to provide adequate features required for each task. Our multi-task unit consists of a *multi-task detail network*, a *deblurring network*, and a *motion compensation network*, where the deblurring and motion compensation networks are attached to the shared multi-task detail network. For the deblurring

[†] Equal contributions

[‡] Corresponding author

* Currently working at Samsung Advanced Institute of Technology

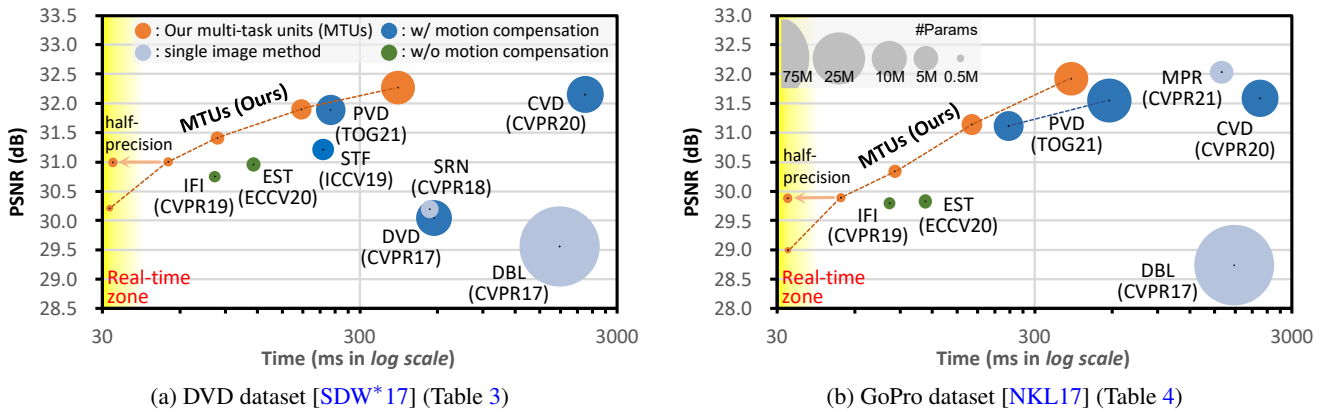


Figure 1: Comparison on deblurring efficiency. Our models are indicated in orange circles, each of which is stacked with a different number of multi-task units (MTUs). DBL [NKL17], SRN [TGS*18] and MPR [ZAK*21] are single image motion deblurring methods, and the rests are video deblurring methods. All the running times are measured on the same PC with NVIDIA Geforce RTX 3090 GPU and averaged from 5 runs on each test set containing frames of the 1280×720 resolution. Note that for methods implemented in PyTorch, we wrap the forward function call of a model (i.e., `torch.nn.Module.forward()`) with CUDA synchronization calls (i.e., `torch.cuda.synchronize()`) to measure the true inference time taken by the model to deblur a target frame.

ring task, the deblurring network is trained to produce a residual image for deblurring, which induces the multi-task detail network to produce detail features. On the other hand, the motion compensation task, which requires additional structure information, may hinder the multi-task detail network from producing detail features. To avoid this, the motion compensation network combines the detail features with pre-computed structure features to constitute structure-injected features, which are then used for the motion compensation task. As a result, the multi-task detail network is trained to obtain detail features that are shared among the deblurring and motion compensation tasks.

Our multi-task unit has a lightweight structure. The multi-task detail network has a simple U-Net structure [RFB15]. Both deblur and motion networks are composed of only single convolution layers. For efficient motion compensation, we use a local correlation-based feature matching operation [DFI*15, HTL18], which does not contain any learnable parameters. Even though our feature matching-based motion compensation may not be as accurate as motion compensation using a densely computed motion (e.g., an optical flow computed from a large flow estimation model), ours takes a much smaller computational overhead but still significantly improves deblurring performance.

As our multi-task unit is light-weighted, we can stack several units to achieve high deblurring quality without too much computational overhead. By controlling the number of stacks, we can flexibly control the balance between the model size and deblurring quality. Thanks to the flexibility of our architecture, our network can cover from environments demanding for low-computation to environments where high deblurring quality is desired.

Finally, we demonstrate a real-time video deblurring network with practical performance realized by our multi-task unit. We thoroughly exploit the trade-off between the deblurring quality and computational cost of our approach. In the experiment, we found that speed evaluation of existing PyTorch-based image and video deblurring methods do not properly measure the true running time

of a model (Sec. 4.2). To our best knowledge, we are the first to present and perform correct speed evaluations in the field of deblurring. Each of our model variants shows a much faster running time compared to the corresponding previous methods with similar deblurring quality (orange circles in Fig. 1).

To summarize, our contributions are:

- A novel lightweight multi-task unit to maximize network sharing between deblurring and motion compensation tasks in a harmonious way;
- Stacked architecture allowing flexible control on the trade-off between deblurring quality and computational cost;
- State-of-the-art video deblurring quality with a faster network and practical real-time video deblurring performance.

2. Related Work

Video Deblurring Video deblurring methods focus on taking the benefits of temporal frames. Classical approaches either directly aggregate temporal information [MOG*06, CWL12, DS15] or solve an inverse problem based on a blur model involving multiple frames [KL15, WB14]. For an effective aggregation of temporal information, these methods use motion compensation based on a homography [MOG*06], local patch search [CWL12], and optical flow [DS15, KL15]. However, it is challenging to accurately estimate motion between blurry frames. For more accurate motion estimation, alternating estimation of motion and deblurred frames has been adopted [CML07, BBRS07, WB14, KL15], in which a joint energy function for deblurring and motion compensation is iteratively optimized with respect to the blur kernels and motion parameters, respectively. However, such iterative optimization requires huge computational costs.

Learning-based Video Deblurring using Motion Compensation

Deep learning-based approaches have been shown to outperform classical methods [KLSH17, NSL19, SDW*17, KSHS18, ZYLP19, ZYP*19, WCY*19, WYL*20, PBT20, ZJWW20, ZYZB20, SR21,

SLL*21, LXZ*21]. Similar to classical methods, many deep learning-based methods adopt motion compensation for aggregating information from neighboring frames [SDW*17, KSHS18, ZYLP19, ZYP*19, WCY*19, PBT20, SLL*21, LXZ*21]. Su *et al.* [SDW*17] show that roughly aligning video frames before feeding them into their deblurring network can improve the deblurring quality. To further improve the deblurring quality, Kim *et al.* [KSHS18] and Zhan *et al.* [ZYLP19] introduce optical flow-based motion compensation approaches that align blurry input frames before deblurring. However, as the optical flow is estimated from blurry frames, its quality is limited, which also limits the deblurring quality. To address this problem, recently, Son *et al.* [SLL*21] propose a pixel volume consisting of multiple matching candidates to robustly handle optical flow errors. To iteratively improve both motion compensation and deblurring quality, Pan *et al.* [PBT20] introduce a cascaded inference approach. Their two-stage method is computationally heavy, requiring more than 10× larger computation time than those of the existing learning-based methods [NSL19, ZYP*19, ZYB20]. Li *et al.* [LXZ*21] extend [PBT20] to resolve optical flow errors by considering all-range volumetric correspondence between the reference and neighboring video frames. However, it also requires a large computational cost similarly to [PBT20].

Implicit motion compensation methods have also been considered for video deblurring. Wang *et al.* [WCY*19] propose deformable convolution networks for motion compensation, while Zhou *et al.* [ZYP*19] estimate dynamic convolution kernels for both motion compensation and deblurring. However, although they do not explicitly use motion estimation such as optical flow, these methods require many additional convolution layers with learnable parameters for either estimating deformable or dynamic convolution kernels to handle motion between frames. Our framework is distinguished from them in that the motion compensation module largely shares the parameters with the deblurring module using multi-task learning. Thereby, our framework can perform both motion compensation and video deblurring faster with a smaller number of parameters, while showing the deblurring performance comparable to or even better than previous methods.

Learning-based Video Deblurring for Efficiency Some approaches pursue efficiency over deblurring quality by omitting the motion compensation module [KLSH17, NSL19, ZYB20]. Kim *et al.* [KLSH17] propose an inter frame-recurrent CNN to sequentially process a blurry video and show that utilizing a deblurred previous frame as an input improves deblurring quality. Nah *et al.* [NSL19] extend [KLSH17] and recurrently process intra-frame features for better propagation of temporal information throughout the inter-frame recurrent pipeline. It boosts deblurring quality while keeping a lightweight model size but increases the inference time. Zhong *et al.* [ZYB20] use RNN to efficiently extract the spatial features of consecutive frames and aggregate temporal information of the features by channel attention. While the aforementioned methods are faster than motion compensation-based methods, the deblurring quality is lower because, without motion compensation, it is difficult to aggregate temporal information in video frames that have large motions and blurs. In addition, all these methods have not provided real-time performance. Distinguished from these

methods, we present a real-time video deblurring framework showing comparable deblurring quality.

Multi-task Learning Our method is closely related to multi-task learning in the form of hard parameter sharing. However, previous approaches [Thr95, BV16, SK18, CBLR18] usually focus on high-level vision tasks, such as classification and scene understanding. Kokkinos *et al.* [Kok17] define low-, mid-, and high-level tasks, and show diverse multi-task learning of multi-level tasks. Still, the low-level tasks in their study include only low-level image recognition tasks, such as boundary, saliency, and normal estimation, but we focus on multi-task learning for the image restoration task. Song *et al.* [SZF*20] propose multi-task learning of stereo matching and edge detection. Although they adopt hard parameter sharing, a large number of auxiliary convolution layers is needed for each task. In contrast, we use only a single convolution layer additionally for each deblurring and motion compensation task.

3. Fast Video Deblurring with Motion Compensation

3.1. Improving Compatibility of Shared Features for Multi-Task Learning

Deblurring and motion compensation tasks are known to benefit each other [KL15, PBT20]. However, we found that each task requires features that are incompatible in terms of *detail* and *structure*. Here, following [PLS*18, KLL16], we use the terms detail and structure to indicate residual detail computed from a blurry image and structural information of the degraded image, respectively. For restoration tasks, including deblurring, it is well known from previous approaches [KLL16, SSH17, JOKK18, SL17, SDW*17, ZYP*19, WCY*19, LSR*21, SLCL21, LLCL22] that residual learning is essential for achieving high performance. In residual learning, a network focuses on extracting residual information (detail), which is then added to an input degraded image (structure) for producing a deblurred result.

While the deblurring quality improves with the residual learning, we observed that the motion compensation task requires structural information in addition to the detail information. For motion compensation, residual details may not provide useful information for matching image patches, such as indistinct contrasts and patterns. Consequently, it would be hard to accurately compensate motions between two video frames using only residual details, as shown in Figs. 2b and 2f. By combining structural information with residual details, the accuracy of motion compensation can be improved, as shown in Figs. 2d and 2h.

In Sec. 4.3.1, we validate the different feature characteristics needed for the deblurring and motion compensation tasks. In the experiment, we compare two networks trained with and without residual learning for both tasks. The network trained with residual learning produces detail features and shows better deblurring quality, compared to the network trained with non-residual learning that induces features to contain both structure and detail information. On the other hand, the network trained with non-residual learning shows better motion compensation performance than the network trained with residual learning.

For multi-task learning of the video deblurring and motion compensation tasks that require different natures of features, we design

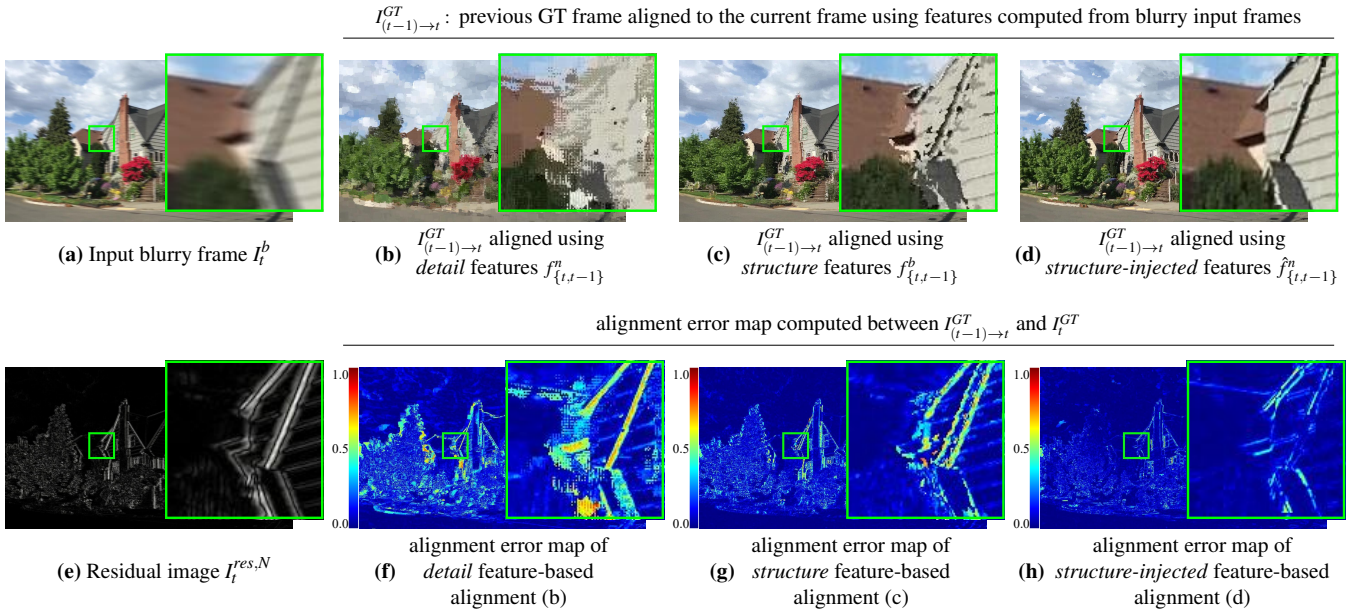


Figure 2: Different natures of features needed for deblurring and motion compensation. Deblurring requires detail-level features, for which we train our model with residual learning to produce a residual image (e), which is then added to the input frame (a) for obtaining the deblurred result. However, such detail-level features are not enough for motion compensation. Alignment results of the previous ground-truth sharp frame aligned to the current sharp frames are shown in (b), (c), and (d), for each of which we use current and previous detail features $f_{\{t,t-1\}}^n$, structure features $f_{\{t,t-1\}}^b$, and structure-injected features $\hat{f}_{\{t,t-1\}}^n$ to compute the motion between the previous and the current blurry frames I_t^b and I_{t-1}^b , respectively. Using the computed motions, we warp the previous ground-truth sharp frame to the current sharp frame. We also show corresponding alignment error maps between each aligned frame and current ones in (f), (g), and (h). The alignment result (b) and error map (f) show that details in residual images do not provide enough information for accurate alignment. Using a structure feature map does improve the alignment quality, as shown in (c) and (g). However, combining the structure and detail information with our structure injection scheme show more accurately aligned results, as shown in (d) and (h). For more detail about mathematical notations, we refer readers to Sec. 3.3.

our multi-task unit to focus on extracting detail features, which are compatible for both deblurring and motion compensation tasks. Using the detail features, video deblurring can be effectively handled with a simple deblur layer based on residual learning. However, as shown in Fig. 2, such detail features cannot be directly used for successful motion compensation. We propose the structure injection scheme, in which the detail features are combined with structural features that are pre-computed from blurry input frames using a convolution layer, and use the combined features for the motion compensation task. The structure injection scheme for motion compensation improves the compatibility of the shared features among the two tasks, by preventing the motion compensation task from demanding structural information from the multi-task detail network. In other words, the scheme induces our framework to focus on computing detail features that are useful for the deblurring task and combined with structure features for the motion compensation task to further boost the deblurring quality, as will be quantitatively discussed in Table 2 in Sec. 4.3.2.

3.2. Network Overview

Our overall network has an iterative structure that stacks N multi-task units for establishing a larger deblurring capability of a net-

work (Fig. 3). Before describing the details of the multi-task unit (Sec. 3.3), we first explain the overall data flow.

Our network adopts an inter-frame recurrent structure. To deblur a target frame I_t^b at the current time step t , our network receives two consecutive blurry frames I_{t-1}^b and I_t^b , and the previously restored frame I_{t-1}^r as inputs. In addition, the network also takes intermediate features computed for the previous frame. From the input, the network performs deblurring and motion compensation and produces a deblurred resulting frame I_t^r . The network has two convolution layers at its beginning, one of which transforms the current input frame I_t^b to a structural feature map f_t^b , and the other transforms $[I_{t-1}^b \cdot I_{t-1}^r]$ into a feature space, where $[\cdot]$ denotes the concatenation operation along the channel axis.

3.3. Multi-task Unit

Instead of having two separate sub-networks for motion compensation and deblurring tasks, our multi-task unit computes common detail features shared by two task-specific branches to reduce the computational cost (Fig. 4). Each multi-task unit consists of three main components: *multi-task detail network* \mathcal{F}^n , *deblurring network* \mathcal{D}^n , and *motion compensation network* \mathcal{M}^n , where n is the index of the multi-task unit. Each multi-task unit shares the same architecture but has different weights.

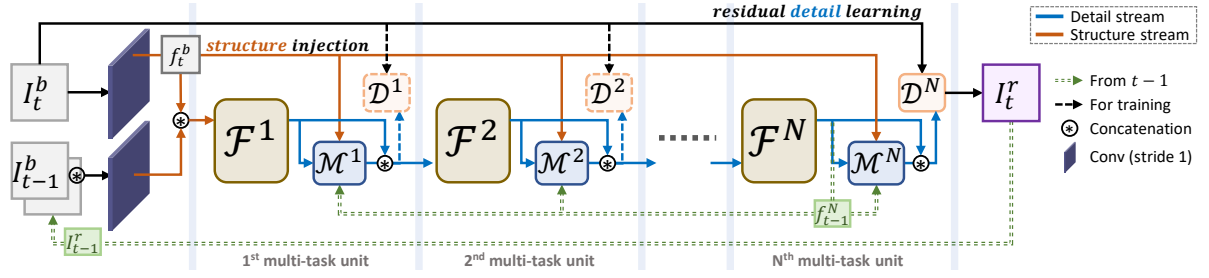


Figure 3: Overall network architecture. Our network adopts a multi-stacked inter-frame recurrent structure. The green lines indicate network inputs from the previous iteration $t - 1$. The modules connected with dashed single lines are only used during the training.

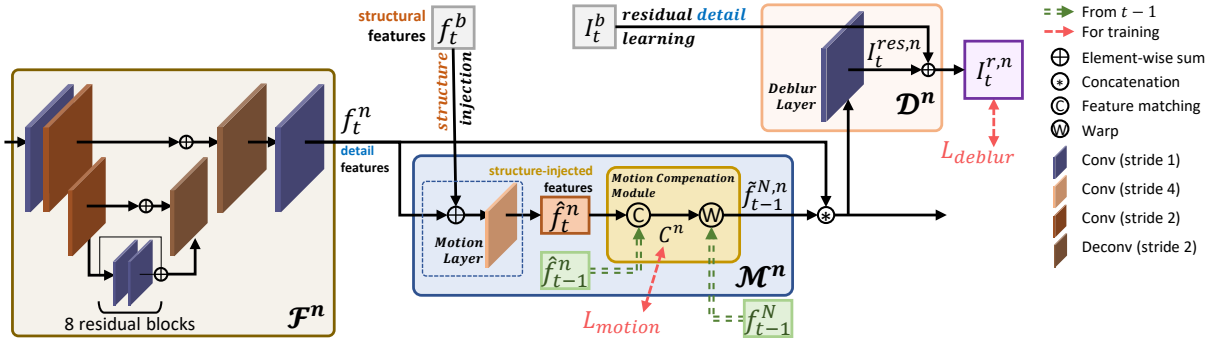


Figure 4: The n -th multi-task unit consists of a multi-task detail network \mathcal{F}^n , which is attached with deblurring network \mathcal{D}^n and motion compensation network \mathcal{M}^n .

In the n -th multi-task unit, \mathcal{F}^n computes an multi-task detail feature map f_t^n shared by \mathcal{D}^n and \mathcal{M}^n . \mathcal{F}^n is a lightweight network based on the U-Net architecture [RFB15], and has most of the learnable parameters of the multi-task unit. \mathcal{D}^n consists of one convolution layer, named a *deblur layer*, that transforms concatenated detail feature maps f_t^n and $\tilde{f}_{t-1}^{N,n}$ of the current and previous frames, respectively, to an output residual image $I_t^{res,n}$. The residual image $I_t^{res,n}$ is then added to the input blurry frame I_t^b to produce a restored frame $I_t^{r,n}$. $\tilde{f}_{t-1}^{N,n}$ is a detail feature map f_{t-1}^N of the previous frame aligned to detail feature map f_t^n of the current frame, where we use \mathcal{M}^n for the alignment. \mathcal{M}^n consists of a single convolution layer, named a *motion layer*, and a feature matching-based motion compensation module with no learnable parameters. In \mathcal{M}^n , we first apply the structure injection scheme, where the pre-computed structural feature map f_t^b is added to the detail feature map f_t^n , which is then transformed to the structure-injected feature map \hat{f}_t^n by the motion layer. Then, the motion compensation module estimates the motion between the feature maps of the current and previous frames, \hat{f}_t^n and \hat{f}_{t-1}^N . Based on the estimated motion, the module aligns the deblurred feature map f_{t-1}^N of the previous frame to the current target frame and obtains $\tilde{f}_{t-1}^{N,n}$. We refer the readers to the supplementary material for more detailed network architecture.

Our overall network is light-weighted compared to previous deblurring networks with explicit motion compensation modules [KSHS18, ZYLP19, PBT20, SLL*21, LXZ*21], as most of the learnable parameters reside in \mathcal{F}^n , and each of \mathcal{D}^n and \mathcal{M}^n has only a single convolution layer. For notational brevity, we denote

the output $I_t^{r,N}$ of the final multi-task unit as I_t^r in the rest of the paper.

Our network is trained using two loss functions: a deblurring loss L_{deblur} , and a motion loss L_{motion} . As shown in Fig. 4, we use both loss functions for every multi-task unit when training the network, where all the multi-task units are trained together in an end-to-end fashion. All the deblurring networks except for \mathcal{D}^N are used only for training and detached after training, as we use only the output of the last deblurring network. The dashed single lines in Figs. 3 and 4 illustrate the connections used only in the training phase.

Our model mitigates the different nature of features required by the video deblurring and motion compensation tasks in an effective way. It adopts long skip-connections (black line denoted as *residual detail learning* in Fig. 3) between the current target blurry frame I_t^b and output of each deblurring networks \mathcal{D}^n . This promotes the multi-task detail networks to learn detail features, which can improve the deblurring performance. These detail features may not suffice for improving motion compensation quality, and the additional feature f_t^b with structural information is provided to the motion compensation networks.

3.4. Motion Compensation Network

Distinguished from previous works leveraging a heavy motion compensation module [SDW*17, ZYLP19, PBT20, SLL*21, LXZ*21], our motion compensation network is extremely fast, as it shares the network capacity for feature extraction with the deblurring task. It only requires a single convolution layer (*i.e.*, the motion

layer) to constitute structure-injected features and a simple feature matching-based motion compensation module with non-learnable operations (Fig. 4). We do not include layers to refine the feature matching result for obtaining a motion that is spatially coherent as an optical flow [DFI*15, HTL18], as it is not necessarily required for the motion to be spatially coherent to improve the deblurring quality [XCW*19].

To compute the motion, the module constructs a matching cost volume $C_t^n : \mathbb{R}^{w \times h \times D^2}$ between \hat{f}_t^n and \hat{f}_{t-1}^n based on the cosine-similarity as done in FlowNet [DFI*15] and LiteFlowNet [HTL18], where D is the maximum displacement. We then compute a displacement map W_t^n from C_t^n by simply finding the 2D displacement vector corresponding to the largest similarity value for each spatial location in C_t^n . Once W_t^n is computed, we warp the deblurred feature map f_{t-1}^N of the previous frame according to W_t^n , and obtain the warped feature map $\tilde{f}_{t-1}^{N,n}$. The warped feature map $\tilde{f}_{t-1}^{N,n}$ is then concatenated with f_t^n and passed to the next multi-task unit.

For computational efficiency, the motion estimation is performed on downsampled feature maps. Specifically, the motion layer has stride 4 to compute a downsampled feature map, which is then used for constructing the cost volume and displacement map. The displacement map is upsampled using the nearest neighbor interpolation to the original scale, and each displacement vector is scaled as well. Moreover, for stacked multi-task units (*i.e.*, $n > 1$), motion estimation can be omitted for faster inference speed. For example, we may use the displacement map W_t^1 computed from the first stack for warping f_{t-1}^N in the rest of the stacks. Refer to the supplementary material for more details.

3.5. Loss Functions

Deblurring Loss The deblurring loss L_{deblur} is defined using the L_1 norm as:

$$L_{deblur} = \frac{1}{M} \sum_{n=1}^N \sum_{x=1}^M \lambda_n |I_t^{r,n}(x) - I_t^{GT}(x)|, \quad (1)$$

where N indicates the number of stacks, and M is the number of pixels in a video frame. Here, x denotes a pixel index, I_t^{GT} is the ground-truth sharp video frame corresponding to I_t^b , and λ_n is a weighting factor for the n -th multi-task unit. In our experiment, we set $\lambda_n = 0.1$ for $n \in [1, \dots, N-1]$, and $\lambda_N = 1$. Our network could be trained even if the deblurring loss is applied only to the last multi-task unit. However, applying the deblurring loss to all the intermediate deblurring networks encourages the multi-task detail networks to more effectively learn detail features from the beginning of the stacked network and eventually results in higher-quality deblurring results.

Motion Loss The existing video deblurring datasets do not provide ground-truth optical flow labels or displacement maps. Thus, we instead generate ground-truth cost volumes using ground-truth sharp video frames. Specifically, for each video frame I_t^b , we compute an optical flow map between I_{t-1}^{GT} and I_t^{GT} using LiteFlowNet [HTL18]. From the optical flow map, we then generate a ground-truth cost volume C_t^{GT} . The vector at each spatial location of C_t^{GT} is a D^2 -dim one-hot vector, where the vector element with

Residual learning	Motion compensation module + L_{motion}	Deblurring quality		Motion compensation accuracy	
		PSNR	SSIM	PSNR	SSIM
		29.42	0.909	26.21	0.842
✓		29.97	0.917	25.09	0.804
	✓	30.07	0.920	27.26	0.863
✓	✓	29.84	0.915	27.24	0.862

Table 1: Feature incompatibility of deblurring and motion compensation.

the value 1 represents the right match. Using C_t^{GT} , we define the motion loss L_{motion} using the cross-entropy as:

$$L_{motion} = -\frac{1}{M} \sum_{n=1}^N \sum_{x=1}^M \sum_{i=1}^{D^2} C_t^{GT}(x, i) \log(\text{softmax}(C_t^n(x); i)), \quad (2)$$

where i is a channel index, and $C_t^{GT}(x, i)$ and $C_t^n(x, i)$ are the feature responses at the spatial location x and channel i . The motion loss encourages the multi-task network to learn multi-task detail features that are useful for motion compensation when it is combined with structural information. Moreover, the loss trains the motion layer to properly inject structural information into detail features f^n so that \hat{f}^n would be suitable for motion estimation.

The final loss L_{total} is defined as: $L_{total} = L_{deblur} + \alpha L_{motion}$, where we use $\alpha = 0.01$ in practice. Thanks to the structure injection scheme in our multi-task unit, L_{deblur} and L_{motion} assist each other for \mathcal{F}^n to produce multi-task detail features suitable for both deblurring and motion compensation.

4. Experiments

4.1. Implementation Details

Dataset We use the DVD [SDW*17] and GoPro [NKL17] datasets to train and validate our network, respectively. The DVD dataset provides 71 pairs of blurry and ground truth sharp videos, which contain 6,708 pairs of 1280×720 video frames in total. In our experiments, we use 5,708 frame pairs in 61 videos as our training set and the rest for the test set as done in [SDW*17]. The GoPro dataset provides 33 video pairs, which contain 3,214 pairs of 1280×720 video frames in total, where the blurry videos are gamma-corrected. As done in [NKL17], we use 2,103 frame pairs in 22 videos as the training set and the rest for the test set.

Training Details For training, we follow the strategy of [KLSH17]. From the training set, we randomly select a blurry video to sample random 13 consecutive video frames. Then, we crop a 256×256 patch from each video frame. We denote the 13 cropped video frames as $\{I_t^b; t \in [1, \dots, 13]\}$. Then, for each iteration of the training, we feed I_t^b , I_{t-1}^b , and I_{t-1}^r to the network. For $t = 1$, we set $I_0^b = I_1^b$, $I_0^r = I_1^r$, and $f_0^n = 0$ for $n \in [1, \dots, N]$. In our training, we set the batch size to 8, and use the Adam optimizer [PGC*17] with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. For the analyses of our proposed method using two stacks in Sec. 4.3, we trained the models for 300K iterations with a fixed learning rate of 1.0×10^{-4} . For the comparison with previous state-of-the-art methods in Sec. 4.4, to obtain the best performance, we trained our

L_{deblur}	Motion compensation module	L_{motion}	Structure injection	Deblurring quality		Motion compensation accuracy	
				PSNR	SSIM	PSNR	SSIM
✓				29.97	0.917	25.09	0.804
✓	✓			30.18	0.921	25.73	0.831
✓	✓	✓		29.84	0.915	27.24	0.862
✓	✓	✓	✓	30.30	0.923	27.25	0.864

Table 2: Ablation study for motion compensation part. Motion compensation module, motion loss L_{motion} , and structure injection are components of a motion compensation network \mathcal{M} .

models for 1,000K iterations with a learning rate of 1.0×10^{-4} , 250K iterations with a learning rate of 2.5×10^{-5} , and additional 50K iterations with a learning rate of 6.25×10^{-6} . We implemented our models using PyTorch [PGC*17].

4.2. Evaluation

For evaluating deblurring quality, we measure the PSNR and SSIM [WBSS04] between deblurring results of the test set and their corresponding ground-truth images. For evaluating motion compensation quality, as done in [SLL*21], we measure inter-frame alignment accuracy. We measure the PSNR and SSIM between I_{t-1}^{GT} and I_t^{GT} after warping I_{t-1}^{GT} using the motion estimated from the last motion compensation module.

As our motion compensation network estimates motion at the 1/4 scale, we downsample I_{t-1}^{GT} and I_t^{GT} by 4 for the evaluation.

We also report computation times of deblurring models as they are important measures for practical applicability. However, we found that the naïve measurement in PyTorch does not consider the asynchronous operation of CUDA and yields an inaccurate running time, which is much smaller than the true inference time. Previous PyTorch-based image and video restoration methods have reported such inaccurate time values, and we re-measure the running time of those methods by wrapping a model with `torch.cuda.synchronize()` to measure the true running time. We evaluate our models and previous ones on the same PC with an NVIDIA GeForce RTX 3090 GPU.

4.3. Analysis

4.3.1. Feature Incompatibility of Deblurring and Motion Compensation Tasks

We first validate our preliminary observation that video deblurring and motion compensation tasks require different features from the perspective of structure and detail. For the validation, we compare a stripped-down baseline model and its three variants. The baseline model is a 2-stacked network (*i.e.*, $N = 2$) trained without residual learning and without motion compensation networks (\mathcal{M}). To the baseline model, we add components one by one to obtain their variants. The first variant is the baseline model trained with residual learning. The second one is the baseline model trained with motion compensation module and motion loss L_{motion} . The third one is the second variant model trained with residual learning. Note that the structure injection scheme in \mathcal{M} is not included here, and we use the output features of the multi-task detail network directly for the motion compensation task, as we aim to investigate only the feature

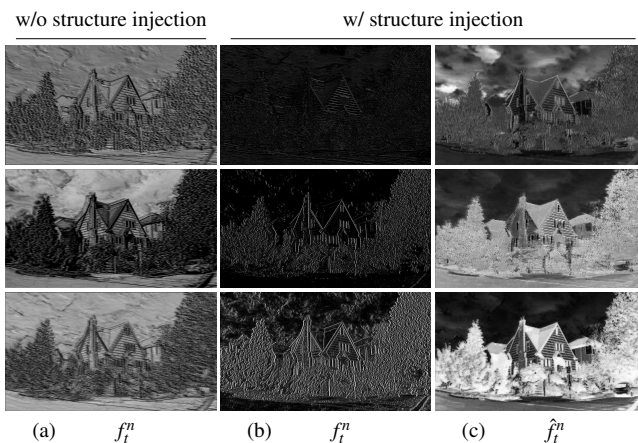


Figure 5: Samples of (a) detail feature maps f_t^n of the final model without the structure injection scheme, (b) detail feature maps f_t^n , and (c) structure-injected feature maps \hat{f}_t^n of our final model with the structure injection scheme. Without the structure injection scheme, the multi-task detail network \mathcal{F}^n produces shared feature maps f_t^n with detail-level and structural information, which helps the motion compensation task but degrades the deblurring quality.

compatibility between motion compensation and deblurring tasks. All of the models are trained using the deblurring loss in Eq. 1. For all models, we measure the motion compensation quality using the inter-frame alignment accuracy as in Sec. 4.2. Note that the alignment accuracy of a model without motion compensation module can still be measured using the module, as the module is composed of only non-learnable operations and can be directly applied to the output features of the multi-task detail network.

Table 1 shows a comparison of the models, from which we can make the following observations. First, adopting residual learning improves the deblurring quality but downgrades the inter-frame alignment accuracy (1st vs. 2nd rows of Table 1). This indicates that detail-level features learned by residual learning are essential for the deblurring task but are incompatible with the motion compensation task. Second, employing motion compensation for the model without residual learning improves both the deblurring and motion compensation quality (1st vs. 3rd rows). In this case, as there is no residual learning in the model, the entire structure information can be used for improving the motion compensation accuracy. Then, the deblurring quality also improves, following the well-known observation that motion compensation helps deblurring. However, when we naïvely combine residual learning and motion compensation module (4th row), the multi-task detail network cannot fully exploit neither detail-level nor structural features. Con-

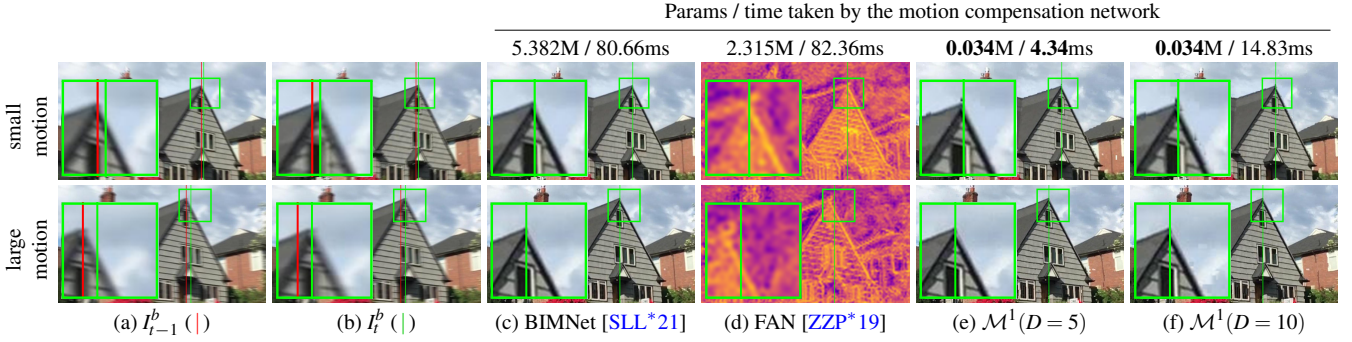


Figure 6: Comparison on the motion compensation network. We compare our motion compensation network \mathcal{M} with the ones used in previous video deblurring methods. We compare (c) the blur-invariant motion compensation network (BIMNet) utilized in PVD [SLL*21] to explicitly compute motion between frames for constructing a pixel volume, and (d) the filter adaptive convolution network (FAN) leveraged in STF [ZZP*19] to implicitly align temporal features for applying per-pixel convolution. For ours, we show the result of the motion compensation network \mathcal{M}^1 in the first multi-task unit of the 2-stacked model with the maximum displacement (e) $D = 5$ and (f) $D = 10$.

sequently, the model (4th row) shows downgraded performance for both the deblurring (vs. 2nd row) and the motion compensation (vs. 3rd row) tasks, indicating that the two tasks require different characteristics of features. To resolve this feature incompatibility, we introduce the structure injection scheme that enables the multi-task detail network to concentrate on exploiting detail-level features.

4.3.2. Ablation Study on Motion Compensation Network

To analyze the effect of our motion compensation approach on deblurring with multi-task learning, we conduct an ablation study (Table 2). All the models in the ablation study are 2-stacked variants of our model that use residual learning for deblurring.

In Table 2, merely adopting a motion compensation module of \mathcal{M} increases the deblurring quality even without the motion loss, verifying the usefulness of motion compensation for deblurring. When the motion loss is simply added to the training of our model without the structure injection scheme of \mathcal{M} , it increases the motion compensation quality but degrades the deblurring quality. In this case, the motion loss induces the multi-task detail network \mathcal{F} to produce features carrying structural information appropriate for motion compensation, but incompatible for deblurring. However, if we employ the structure injection scheme in \mathcal{M} , the deblurring quality significantly increases while the motion compensation quality is retained. This validates that structure injection plays a key role for the multi-task detail network to focus on producing detail features that can be compatibly shared by both deblurring and motion compensation tasks.

Fig. 5 visualizes samples of detail and structure-injected feature maps in our model variants with and without the structure injection scheme. In the model without structure injection, the multi-task detail network produces the features where the mixture of structural and detail information exists (the first column of the figure). In the model, the features are used for both deblurring and motion compensation tasks, eventually decreasing deblurring performance. On the other hand, in our final model with the structure injection scheme, the multi-task detail network can produce shared features f_t^n containing only detail information suitable for residual learning of deblurring (the second column of the figure). At the same

time, structure-injected features f_t^n have both structure and detail information suitable for motion compensation (the last column).

4.3.3. Effectiveness of Lightweight Motion Compensation Network Compared with Previous Approaches

While motion compensation improves the deblurring performance as shown in many previous approaches [SDW*17, ZYLP19, KSHS18, ZZP*19, WCY*19, PBT20, SLL*21, LXZ*21], heavy computation is required to align blurry frames, resulting in significantly increased computation time, model size, and memory consumption for video deblurring. However, as our proposed multi-task unit handles large portions of both deblurring and motion compensation tasks using a single shared network, the motion compensation task requires only a small computational overhead but still provides enough performance needed for exploiting temporal information in video frames.

To validate the effectiveness of our lightweight motion compensation network \mathcal{M} , we compare ours with motion compensation networks utilized in previous video deblurring methods, PVD [SLL*21] and STF [ZZP*19]. In PVD, a separate flow estimation network (BIMNet) is utilized to obtain an optical flow between neighboring and target frames. The flow is then used to construct a pixel volume that comprises multiple candidate pixels of a neighboring frame matched to each pixel of a target frame. The pixel volume is directly fed to the deblurring network and significantly improves the deblurring quality. In STF, the filter adaptive convolution network (FAN) is embedded within the deblurring network and estimates dynamic per-pixel convolution kernels. The per-pixel kernels are then convolved to the features of a neighboring frame for aligning them to the current target features.

For the validation, we compare inference time and alignment results of motion compensation networks. For visualizing alignment results of BIMNet [SLL*21] and our motion compensation network \mathcal{M} , we first predict the motion from I_t^b to I_{t-1}^b . Then, using the motion, we directly warp the previous sharp frame I_{t-1}^{GT} for aligning it to the current frame, as described in Fig. 2. However, for FAN [ZZP*19], it is not straightforward to apply motion compensation to a frame in the image domain because FAN does not

	DBL ^o	SRN ^o	DVD ^o	IFI ^o	EST ^o	STF ^o	PVD ^o	CVD ^o	Ours ⁿ (n : # of stack)				Ours ⁿ _{hp}	Ours ⁿ _L
	[NKL17]	[TGS*18]	[SDW*17]	[NSL19]	[ZYB20]	[ZZP*19]	[SLL*21]	[PBT20]	$n = 1$ ^o	2 ^o	4 ^o	10 ^o	2 ^o	10 ^o
PSNR	29.56	30.20	30.05	30.75	30.97	31.21	31.89	32.15	30.21	31.00	31.41	31.90	30.99	32.27
SSIM	0.913	0.922	0.920	0.930	0.931	0.935	0.944	0.949	0.921	0.933	0.937	0.943	0.932	0.947
Params (M)	75.92	3.76	15.31	1.67	2.23	5.37	10.51	16.19	0.54	1.04	2.04	5.03	1.04	13.80
Time (ms)	1790	560	581	82	116	216	231	2250	32	54	84	178	33	421

Table 3: Quantitative evaluation on the DVD dataset [SDW*17] (refer to Fig. 1a). The colored dots arrange models showing a similar deblurring quality in PSNR, where each of our models yields an upper bound PSNR at a much faster running time.




Figure 7: Comparison on the DVD dataset [SDW*17]. Cropped images visualize regions in the green box of two consecutive frames.

predict a flow containing per-pixel displacement between temporal features. Instead, FAN implicitly performs motion compensation in the feature space by convolving per-pixel kernels to the features of the previous frame. Hence, for FAN, we visualize samples of the resulting features aligned to features at the current time step, as done in [ZZP*19].

Fig. 6 shows the results. For temporal frames exhibiting small motion (the top row of Fig. 6), all compared methods properly compensate the motion between the current and previous frames. For frames with larger motion (the bottom row), while BIMNet and ours well align the previous frame to the current frame, FAN fails to do so. This is mainly due to the limited spatial size of the dynamic kernels of FAN utilizing 5×5 kernels in practice (corresponds to our motion compensation module with $D = 2$), which is not enough for handling large motion between temporal features. It is also difficult for FAN to adopt a larger kernel size because the cost would become intractable to manage per-pixel kernels and their convolution operations. Distinguished from BIMNet and FAN that require heavy computations for maintaining a large motion estimation network and for employing per-pixel adaptive convolution operations, respectively, ours requires only a single convolution layer (*i.e.*, motion layer) and a simple local feature matching operation for the motion compensation task. It is noteworthy that our approach still

produces moderately motion-compensated results, with $158\times$ and $68\times$ smaller model size than those of BIMNet and FAN, respectively, and $18\times$ faster inference speed.

4.4. Comparison with Previous Methods

Finally, we compare our method with state-of-the-art single image [NKL17, TGS*18] and video deblurring methods [SDW*17, NSL19, ZZP*19, ZYB20, PBT20, SLL*21]. The most recent method [LXZ*21] is not included in this comparison since the source code is not publicly available.

For comparison, we evaluate our methods with the ones on the DVD dataset [SDW*17] and the GoPro dataset [NKL17]. The compared models are trained with the training set in the corresponding dataset. We compare our models, each of which is stacked with 1, 2, 4, and 10 multi-task units, denoted by Oursⁿ, where n indicates the number of stacks. In addition, considering the trade-off between computational costs and deblurring performance, we also include our 10-stacked model with larger parameters (Ours¹⁰_L), for which we increase the number of channels in the convolution layers from 26 to 48. We also include the performance of our 2-stacked model (Ours²_{hp}) implemented in half-precision [MNA*18], which

	DBL [•] [NKL17]	IFI [•] [NSL19]	EST [•] [ZYB20]	PVD [•] [SLL*21]	PVD _L [•] [SLL*21]	CVD [•] [PBT20]	Ours ⁿ (n: # of stacks)				Ours _{hp} ⁿ	Ours _L ⁿ
							n = 1 [•]	2 [•]	4	10 [•]	2 [•]	10 [•]
PSNR	28.74	29.80	29.83	31.11	31.55	31.59	29.00	29.91	30.34	31.14	29.90	31.92
SSIM	0.912	0.926	0.924	0.943	0.948	0.952	0.915	0.928	0.933	0.942	0.928	0.950
Params (M)	75.92	1.67	2.23	10.51	23.36	16.19	0.54	1.04	2.04	5.03	1.04	13.80
Time (ms)	1788	83	113	238	585	2241	33	53	86	171	33	416

Table 4: Quantitative evaluation on the GoPro dataset [NKL17] (refer to Fig. 1b).



Figure 8: Comparison on the GoPro dataset [NKL17]. Cropped images visualize regions in the green box of two consecutive frames.

also shows the real-time deblurring speed with practical deblurring performance.

As Fig. 1 indicates, our multi-stacked models show higher PSNRs than all the other methods at similar inference speeds. Table 3 shows a quantitative comparison on the DVD dataset [SDW*17], respectively. Ours² outperforms the single image methods [NKL17, TGS*18], DVD [SDW*17] and IFI [NSL19], while it has a significantly smaller model size and higher speed. Ours_{hp}² runs in real-time speed while showing similar deblurring quality with IFI and EST, which are designed primarily for efficiency. Compared to EST [ZYB20] and STF [ZZP*19], Ours⁴ achieves a higher deblurring quality at a higher speed. Compared to PVD [SLL*21], Ours¹⁰ runs at a faster speed with smaller model size, while achieving comparable deblurring performance. Compared to CVD [PBT20], Ours¹⁰ achieves slightly lower deblurring performance, but it shows 10× faster speed with 3× smaller model size. Ours_L¹⁰ achieves higher PSNR than CVD with a much faster speed, showing the efficiency of our approach. CVD shows high de-

blurring accuracy, but its inference is much slower than others. The slow computational time is due to a large model, including motion compensation and deblurring networks, and iterative utilization of the large model for processing a single frame. Specifically, for deblurring a single frame, CVD needs to align eight pairs of frames leveraging the flow estimation method [SYLK18], which is already computationally heavier (8.75M parameters) than previous deblurring methods.

Table 4 shows a quantitative comparison on the GoPro dataset [NKL17]. We compare our models with previous methods, DBL [NKL17], IFI [NSL19], EST [ZYB20], CVD [PBT20], and PVD [SLL*21], for which models pre-trained with the dataset are available. As the table shows, Ours_{hp}² with real-time speed still achieves comparable PSNR with both IFI and EST, while Ours⁴ outperforms them. Ours¹⁰ also shows comparable PSNR to PVD with faster speed. It is worth noting that Ours_L¹⁰ outperforms the large-size model of PVD (PVD_L) in PSNR with faster speed, implying a better trade-off between deblurring quality and computa-

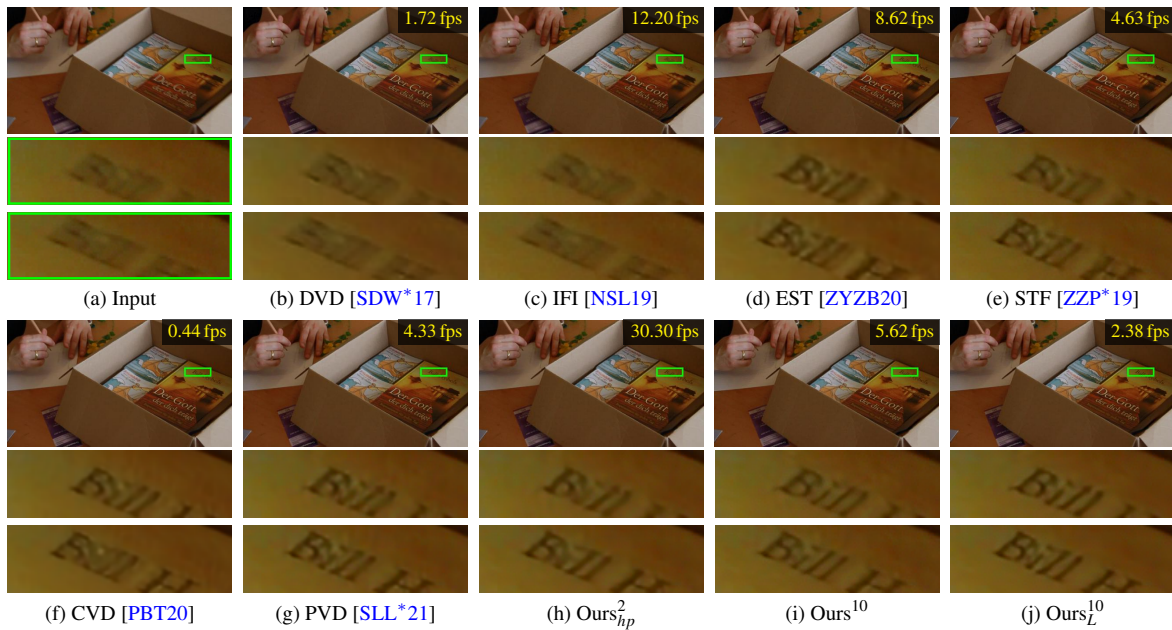


Figure 9: Comparison on real-world blurred videos [CWL12]. Cropped images visualize regions in the green box of two consecutive frames.

tional cost. $Ours_L^{10}$ also shows higher PSNR than CVD despite its smaller computational cost.

Fig. 7, Fig. 8, and Fig. 9 show qualitative comparisons on the DVD dataset [SDW*17], GoPro dataset [NKL17], and real-world blurred videos [CWL12]. Due to the severe blur in the input frames, the results of the other methods still have remaining blur. On the other hand, the results of our 10-stacked model show better-restored details. In addition, $Ours_{hp}^2$ removes the majority of blurs successfully, comparable to efficiency-pursuing methods such as IFI and EST, despite its real-time speed. We found that our real-time model can also improve the performance of object detection on blurry scenes successfully when the model is used as pre-processing. It shows the high applicability of our model to various real-time applications. Additional results and the object detection examples can be found in the supplementary material.

5. Conclusion

We proposed a lightweight multi-task unit that supports both video deblurring and motion compensation tasks. For cost-effective incorporation of motion compensation into video deblurring, we design the multi-task unit to maximally share the network capacity by resolving the different natures of features demanded by each task. The multi-task unit maintains streams of detail features commonly required for both tasks and additionally injects pre-computed structure features for motion compensation. Our approach reduces the computational overhead, especially for motion compensation, and enables real-time video deblurring. By simply stacking multiple multi-task units, our deblurring framework can flexibly control the trade-off between deblurring quality and computational costs, achieving state-of-the-art performance with faster speed.

While structure injection plays a key role in improving motion compensation and deblurring quality, structural information is cur-

rently computed with a simple convolution layer. We plan to focus on extracting better structural features while suppressing computational costs. The improved structural features may be leveraged to boost motion compensation quality and enhance residual features, promoting better deblurring performance.

Acknowledgments

This work was supported by the Ministry of Science and ICT, Korea, through IITP grants (SW Star Lab, 2015-0-00174; AI Innovation Hub, 2021-0-02068; Artificial Intelligence Graduate School Program (POSTECH), 2019-0-01906) and NRF grant (2018R1A5A1060031).

References

- [BBS07] BAR L., BERKELS B., RUMPF M., SAPIRO G.: A variational framework for simultaneous motion estimation and restoration of motion-blurred video. In *Proc. ICCV* (2007). 2
- [BV16] BILEN H., VEDALDI A.: Integrated perception with recurrent multi-task neural networks. In *Proc. NeurIPS* (2016). 3
- [CBLR18] CHEN Z., BADRINARAYANAN V., LEE C.-Y., RABINOVICH A.: Gradnorm: Gradient normalization for adaptive loss balancing in deep multitask networks. In *Proc. ICML* (2018). 3
- [CML07] CHO S., MATSUSHITA Y., LEE S.: Removing non-uniform motion blur from images. In *Proc. ICCV* (2007). 2
- [CWL12] CHO S., WANG J., LEE S.: Video deblurring for hand-held cameras using patch-based synthesis. *ACM Transactions on Graphics* 31, 4 (2012), 64:1–64:9. 1, 2, 11
- [DFI*15] DOSOVITSKIY A., FISCHER P., ILG E., HÄUSSER P., HAZIRBAŞ C., GOLKOV V., V.D. SMAGT P., CREMERS D., BROX T.: FlowNet: Learning optical flow with convolutional networks. In *Proc. ICCV* (2015). 2, 6
- [DS15] DELBRACIO M., SAPIRO G.: Hand-held video deblurring via efficient fourier aggregation. *IEEE Transactions on Computational Imaging* 1, 4 (2015), 270–283. 1, 2

- [HTL18] HUI T.-W., TANG X., LOY C. C.: LiteFlowNet: A lightweight convolutional neural network for optical flow estimation. In *Proc. CVPR* (2018). 2, 6
- [JOKK18] JO Y., OH S. W., KANG J., KIM S. J.: Deep video super-resolution network using dynamic upsampling filters without explicit motion compensation. In *Proc. CVPR* (2018). 1, 3
- [KBM*18] KUPYN O., BUDZAN V., MYKHAILYCH M., MISHKIN D., MATAS J.: DeblurGAN: Blind motion deblurring using conditional adversarial networks. In *Proc. CVPR* (2018). 1
- [KL15] KIM T. H., LEE K. M.: Generalized video deblurring for dynamic scenes. In *Proc. CVPR* (2015). 1, 2, 3
- [KLL16] KIM J., LEE J. K., LEE K. M.: Accurate image super-resolution using very deep convolutional networks. In *Proc. CVPR* (2016). 1, 3
- [KLSH17] KIM T. H., LEE K. M., SCHÖLKOPF B., HIRSCH M.: Online video deblurring via dynamic temporal blending network. In *Proc. ICCV* (2017). 1, 2, 3, 6
- [Kok17] KOKKINOS I.: Ubernet: Training a universal convolutional neural network for low-, mid-, and high-level vision using diverse datasets and limited memory. In *Proc. CVPR* (2017). 3
- [KSHS18] KIM T. H., SAJJADI M. S. M., HIRSCH M., SCHOLKOPF B.: Spatio-temporal transformer network for video restoration. In *Proc. ECCV* (2018). 1, 2, 3, 5, 8
- [LLCL22] LEE J., LEE M., CHO S., LEE S.: Reference-based video super-resolution using multi-camera video triplets. In *Proc. CVPR* (2022). 3
- [LSR*21] LEE J., SON H., RIM J., CHO S., LEE S.: Iterative filter adaptive network for single image defocus deblurring. In *Proc. CVPR* (2021). 3
- [LXZ*21] LI D., XU C., ZHANG K., YU X., ZHONG Y., REN W., SUOMINEN H., LI H.: Arvo: Learning all-range volumetric correspondence for video deblurring. In *Proc. CVPR* (2021). 1, 2, 3, 5, 8, 9
- [MHS*16] MA B., HUANG L., SHEN J., SHAO L., YANG M., PORIKLI F.: Visual tracking under motion blur. *IEEE Transactions on Image Processing* 25, 12 (2016), 5867–5876. 1
- [MNA*18] MICIKEVICIUS P., NARANG S., ALBEN J., DIAMOS G., ELSÉN E., GARCIA D., GINSBURG B., HOUSTON M., KUCHARIEV O., VENKATESH G., WU H.: Mixed precision training. In *Proc. ICLR* (2018). 9
- [MOG*06] MATSUSHITA Y., OFEK E., GE W., TANG X., SHUM H. Y.: Full-frame video stabilization with motion inpainting. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 28, 7 (2006), 1150–1163. 2
- [NKL17] NAH S., KIM T. H., LEE K. M.: Deep multi-scale convolutional neural network for dynamic scene deblurring. In *Proc. CVPR* (2017). 2, 6, 9, 10, 11
- [NSL19] NAH S., SON S., LEE K. M.: Recurrent neural networks with intra-frame iterations for video deblurring. In *Proc. CVPR* (2019). 1, 2, 3, 9, 10, 11
- [PBT20] PAN J., BAI H., TANG J.: Cascaded deep video deblurring using temporal sharpness prior. In *Proc. CVPR* (2020). 1, 2, 3, 5, 8, 9, 10, 11
- [PGC*17] PASZKE A., GROSS S., CHINTALA S., CHANAN G., YANG E., DEVITO Z., LIN Z., DESMAISON A., ANTIGA L., LERER A.: Automatic differentiation in pytorch. In *Proc. NeurIPS* (2017). 6, 7
- [PHZ*19] PEI Y., HUANG Y., ZOU Q., ZHANG X., WANG S.: Effects of image degradation and degradation removal to CNN-based image classification. *IEEE Transactions on Pattern Analysis and Machine Intelligence* (2019). 1
- [PLS*18] PAN J., LIU S., SUN D., ZHANG J., LIU Y., REN J., LI Z., TANG J., LU H., TAI Y.-W., YANG M.-H.: Learning dual convolutional neural networks for low-level vision. In *Proc. CVPR* (2018). 3
- [RFB15] RONNEBERGER O., FISCHER P., BROX T.: U-Net: Convolutional networks for biomedical image segmentation. In *Proc. MICCAI* (2015). 2, 5
- [SDW*17] SU S., DELBRACIO M., WANG J., SAPIRO G., HEIDRICH W., WANG O.: Deep video deblurring for hand-held cameras. In *Proc. CVPR* (2017). 1, 2, 3, 5, 6, 8, 9, 10, 11
- [SK18] SENER O., KOLTUN V.: Multi-task learning as multi-objective optimization. In *Proc. NeurIPS* (2018). 3
- [SL17] SON H., LEE S.: Fast non-blind deconvolution via regularized residual networks with long/short skip-connections. In *Proc. ICCP* (2017). 1, 3
- [SLCL21] SON H., LEE J., CHO S., LEE S.: Single image defocus deblurring using kernel-sharing parallel atrous convolutions. In *Proc. ICCV* (2021). 3
- [SLL*21] SON H., LEE J., LEE J., CHO S., LEE S.: Recurrent video deblurring with blur-invariant motion estimation and pixel volumes. *ACM Transactions on Graphics* 40, 5 (2021). 1, 2, 3, 5, 7, 8, 9, 10, 11
- [SR21] SUIN M., RAJAGOPALAN A. N.: Gated spatio-temporal attention-guided video deblurring. In *Proc. CVPR* (2021). 2
- [SSH17] SAJJADI M. S. M., SCHÖLKOPF B., HIRSCH M.: EnhanceNet: Single Image Super-Resolution through Automated Texture Synthesis. In *Proc. ICCV* (2017). 1, 3
- [SYLK18] SUN D., YANG X., LIU M.-Y., KAUTZ J.: PWC-Net: CNNs for optical flow using pyramid, warping, and cost volume. In *Proc. CVPR* (2018). 10
- [SZF*20] SONG X., ZHAO X., FANG L., HU H., ZHOU YU Y.: EdgeStereo: An effective multi-task learning network for stereo matching and edge detection. *International Journal of Computer Vision* 128 (2020), 910–930. 3
- [TGS*18] TAO X., GAO H., SHEN X., WANG J., JIA J.: Scale-recurrent network for deep image deblurring. In *Proc. CVPR* (2018). 2, 9, 10
- [Thr95] THRUN S.: Is learning the n-th thing any easier than learning the first? In *Proc. NeurIPS* (1995). 3
- [WB14] WULFF J., BLACK M.: Modeling blurred video with layers. In *Proc. ECCV* (2014). 2
- [WBSS04] WANG Z., BOVIK A. C., SHEIKH H. R., SIMONCELLI E. P.: Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing* 13, 4 (2004), 600–612. 7
- [WCY*19] WANG X., CHAN K. C., YU K., DONG C., LOY C. C.: Edvr: Video restoration with enhanced deformable convolutional networks. In *Proc. CVPRW* (2019). 1, 2, 3, 8
- [WYL*20] WU J., YU X., LIU D., CHANDRAKER M., WANG Z.: DAVID: dual-attentional video deblurring. In *Proc. WACV* (2020). 2
- [XCW*19] XUE T., CHEN B., WU J., WEI D., FREEMAN W. T.: Video enhancement with task-oriented flow. *International Journal of Computer Vision* 127, 8 (2019), 1106–1125. 6
- [ZAK*21] ZAMIR S. W., ARORA A., KHAN S., HAYAT M., KHAN F. S., YANG M.-H., SHAO L.: Multi-stage progressive image restoration. In *Proc. CVPR* (2021). 2
- [ZJWW20] ZHANG X., JIANG R., WANG T., WANG J.: Recursive neural network for video deblurring. *IEEE Transactions on Circuits and Systems for Video Technology* (2020). 2
- [ZYL19] ZHAN Z., YANG X., LI Y., PANG C.: Video deblurring via motion compensation and adaptive information fusion. *Neurocomputing* 341 (2019), 88–98. 1, 2, 3, 5, 8
- [ZYZB20] ZHONG Z., YE G., ZHENG Y., BO Z.: Efficient spatio-temporal recurrent neural network for video deblurring. In *Proc. ECCV* (2020). 1, 2, 3, 9, 10, 11
- [ZZP*19] ZHOU S., ZHANG J., PAN J., XIE H., ZUO W., REN J.: Spatio-temporal filter adaptive network for video deblurring. In *Proc. ICCV* (2019). 1, 2, 3, 8, 9, 10, 11