

Optical Flow Estimation from a Single Motion-blurred Image

Dawit Mureja Argaw, Junsik Kim, Francois Rameau, Jae Won Cho, In So Kweon

KAIST Robotics and Computer Vision Lab., Daejeon, Korea
 dawitmureja@kaist.ac.kr, {mibastro, rameau.fr}@gmail.com, {chojw, iskweon77}@kaist.ac.kr

Abstract

In most of computer vision applications, motion blur is regarded as an undesirable artifact. However, it has been shown that motion blur in an image may have practical interests in fundamental computer vision problems. In this work, we propose a novel framework to estimate optical flow from a single motion-blurred image in an end-to-end manner. We design our network with transformer networks to learn globally and locally varying motions from encoded features of a motion-blurred input, and decode left and right frame features without explicit frame supervision. A flow estimator network is then used to estimate optical flow from the decoded features in a coarse-to-fine manner. We qualitatively and quantitatively evaluate our model through a large set of experiments on synthetic and real motion-blur datasets. We also provide in-depth analysis of our model in connection with related approaches to highlight the effectiveness and favorability of our approach. Furthermore, we showcase the applicability of the flow estimated by our method on deblurring and moving object segmentation tasks.

Introduction

Acquiring an image with a camera requires the photosensitive sensor of the camera to be exposed to light for a certain duration in order to collect enough photons. Therefore, if a rapid motion of the camera is performed during this time interval – or if a dynamic object is present in the scene, the resulting image will be *motion-blurred* and will appear smeared along the direction of the relative motion. Motion blur is often considered to be a degrading artifact in many computer vision applications, such as, 3D reconstruction (Seok Lee and Mu Lee 2013) and visual SLAM (Lee, Kwon, and Lee 2011). Therefore, existing deblurring approaches have been proposed to restore a clean image from a blurred image (Zheng, Xu, and Jia 2013; Hyun Kim and Mu Lee 2014; Wieschollek et al. 2016; Gong et al. 2017).

On the other hand, instead of being considered as an unwanted noise that is to be removed, it has been shown that motion-blur in an image may have some practical interest in core computer vision problems, such as optical flow estimation (Rekleitis 1995; Schoueri, Scaccia, and Rekleitis 2009; Dai and Wu 2008), video sequence restoration (Jin,

Meishvili, and Favaro 2018; Purohit, Shah, and Rajagopalan 2019), and 3D scene reconstruction (Qiu et al. 2019).

Our work focuses on recovering the apparent motion of camera and objects during the exposure time from a single motion-blurred input *i.e.* optical flow estimation. Earlier works (Rekleitis 1995; Schoueri, Scaccia, and Rekleitis 2009) deployed a uniform motion-blur assumption, which is often violated in practice. To address this issue, follow-up works (Portz, Zhang, and Jiang 2012; Li et al. 2014; Tu, Poppe, and Veltkamp 2015; Li et al. 2014) extended the traditional warping-based methods to a sequence of blurred inputs by imposing further constraints such as non-uniform blur matching and blur gradient constancy. Compared to earlier works, we estimate dense optical flow only from a single motion-blurred image in an end-to-end manner using a novel deep learning framework. Moreover, our method does not impose any restrictive assumption on the motion-blur, hence, it is robust for various blur types.

Another line of work closely related to our work is motion-flow estimation from a blurred image for deblurring applications (Dai and Wu 2008; Hyun Kim and Mu Lee 2014; Gong et al. 2017). However, these works assume restrictive assumptions for motion estimation, *e.g.*, linear motion assumption (Dai and Wu 2008; Hyun Kim and Mu Lee 2014) and constrained flow magnitude and direction (Gong et al. 2017). All related works so far imposed a constraint on the motion-kernel, and hence experimented with synthetically simulated blurs under the imposed constraint. In comparison, we train and analyze our network with blurs generated from real high speed videos with no particular motion assumption.

Our proposed framework is composed of three network components: feature encoder, feature decoder, and flow estimator. The feature encoder extracts features from the given motion-blurred input at different spatial scales in a top-down fashion. A feature decoder then decodes the extracted features in a bottom-up manner by learning motion from blur. The feature decoder is composed of spatial transformer networks (STNs) (Jaderberg et al. 2015) and feature refining blocks to learn globally and locally-varying motions from the encoded features, respectively. The flow estimator inputs the decoded features at different levels and estimates optical flow in a coarse-to-fine manner.

We conduct experiments on synthetic and real image

blur datasets. The synthetic image motion-blur dataset is generated via frame interpolation (Jiang et al. 2018) using the Monkaa dataset (N.Mayer et al. 2016). For real image motion-blur datasets, we follow previous studies (Nah, Kim, and Lee 2017; Jin, Meishvili, and Favaro 2018) and temporally average sequential frames in high speed videos. As ground truth optical flow does not exist for the high speed video datasets, we compute a flow between the first and last end frames (from which the motion blurred image is averaged) using pretrained state-of-the-art optical flow models (Sun et al. 2018; Ilg et al. 2017) and use the computed flow as a pseudo-supervision (pseudo-ground truth).

Our contributions are summarized as follows: (1). To the best of our knowledge, we present the first deep learning solution to estimate dense optical flow from a motion-blurred image without any restrictive assumption on the motion kernel. (2). We use pseudo-ground truth flow obtained from pretrained models to successfully train our network for the optical flow estimation from a motion-blurred image. We also show the effectiveness of the proposed network in inferring optical flow from synthetic and real image motion-blur datasets through in-depth experiments. (3). We perform detailed analysis of our model in comparison to related approaches and perform ablation studies on different network components to show the effectiveness and flexibility of our approach. (4). We showcase the applicability of the optical flow estimated by our method on motion-blur removal and moving object segmentation tasks.

Related Works

Optical flow and motion-blur. Brightness constancy and spatial smoothness assumptions (Horn and Schunck 1981) often do not conform when estimating optical flow in a blurry scenario. To cope with this limitation, Rekleitis *et al.* (Rekleitis 1995) proposed an algorithm to estimate optical flow from a single motion-blurred gray-scale image based on an observation that motion-blur introduces a ripple in the Fourier Transform of the image. Schoueri *et al.* (Schoueri, Scaccia, and Rekleitis 2009) extended this algorithm to colored images by computing the weighted combination of the optical flow of each image channel estimated by (Rekleitis 1995). These methods (Rekleitis 1995; Schoueri, Scaccia, and Rekleitis 2009), however, were limited to linear deblurring filters. To deal with spatially-varying blurs, multiple image sequences were exploited by adapting classical warping-based approaches with modified intensity (Portz, Zhang, and Jiang 2012) and blur gradient (Li et al. 2014) constancy terms. However, these assumptions are often limited when extended to real motion-blurred images. In this work, we introduce a deep learning framework to estimate optical flow from a single motion-blurred image without any particular motion assumption, and our proposed model generalizes well for real motion-blurred examples.

Motion flow for deblurring. Some of previous deblurring approaches also estimate the underlying motion in a motion-blurred image. For example, (Cho and Lee 2009; Dai and Wu 2008; Fergus et al. 2006; Zheng, Xu, and Jia

2013) assume uniform motion blur to remove a blur from an image; however, these methods often fail to remove non-uniform motion blurs. To address this issue, non-uniform deblurring approaches (Gupta et al. 2010; Wieschollek et al. 2016; Levin 2007; Pan et al. 2016; Hyun Kim and Mu Lee 2014) used motion priors based on specific motion models. Since real motion-blurred images often do not comply to these motion assumptions, learning-based discriminative approaches (Chakrabarti, Zickler, and Freeman 2010; Couzinie-Devy et al. 2013) have been proposed to learn motion patterns from a blurred image. These methods are, however, limited since features are manually designed with simple mapping functions. Recently, Gong *et al.* (Gong et al. 2017) proposed a deep learning approach for heterogeneous blur removal via motion-flow estimation. They treated motion-flow prediction as a classification task and trained a fully convolutional network over a discrete output domain.

Recent works on learning from motion-blur. Our work is also related to recent deep learning approaches that focus on unfolding the hidden information in a motion blurred image. Jin *et al.* (Jin, Meishvili, and Favaro 2018) and Purohit *et al.* (Purohit, Shah, and Rajagopalan 2019) reconstructed latent video frames from a single motion-blurred image. Qui *et al.* (Qui et al. 2019) proposed a method to recover the 3D scene collapsed during the exposure process from a motion-blurred image. These works show that motion blur can become meaningful information when processed properly. In this paper, in addition to the previously proposed applications, we explore another potential of motion blur, *i.e.* optical flow estimation.

Dataset Generation

Manually collecting a large set of blurred images is challenging and requires a great deal of human effort. Hence, a common practice in computer vision research is to generate motion-blurred images by averaging sequential frames in high frame rate videos (Jin, Meishvili, and Favaro 2018; Purohit, Shah, and Rajagopalan 2019; Nah, Kim, and Lee 2017). In this work, we experiment with both synthetic and real scene blur datasets without particular assumption, and the process of generating these datasets is described below.

Synthetic scene blur dataset. We take advantage of the Monkaa dataset proposed in (N.Mayer et al. 2016) to generate a synthetic image motion-blur dataset for optical flow estimation. The dataset provides ground truth flows between a pair of frames in synthetic video scenes. Given two consecutive frames, in order to simulate the motion blur, we interpolate intermediate frames using (Jiang et al. 2018) and average all resulting frames. The number of intermediate frames to be interpolated relies on the magnitude of the motion of objects in the scene to ensure a smooth motion-blur generation. Hence, to generate a natural-like motion blur, we defined a simple relationship between the number of frames to be interpolated n and the maximum pixel displacement s in the ground truth flow as follows: $n = \max(|s|, 5)$, where $|s|$ is the absolute value of the s . To avoid severely blurred




Figure 1: (a) First and last images from (N.Mayer et al. 2016) overlaid (b) Ground truth optical flow between first and last images (c) Intermediate frame interpolation using (Jiang et al. 2018) (d) Motion-blurred image generated by averaging first, intermediate and last images (e) Optical flow predicted by our model from the motion-blurred image.

images, we discarded samples with $|s| > 100$. We generate a Monkaa blur dataset with 10,000 training and 1200 test images (see Fig. 1).

Real scene blur dataset. To generate real scene motion-blur images for network training, we use high speed video datasets: GoPro (Nah, Kim, and Lee 2017) and NfS (Galoogahi et al. 2017). The GoPro high speed video dataset which is commonly used for dynamic scene deblurring has 33 videos taken at 240fps, out of which 25 videos are used for training and the rest are used for testing. Motion-blurred images are generated by averaging 7 consecutive frames in a video. As a result, we obtain blurred frames approximately 30 fps which is common setting for commercial cameras. We refer to the blur dataset generated from (Nah, Kim, and Lee 2017) as the GoPro blur dataset. We also experimented with the Need for Speed (NfS) (Galoogahi et al. 2017) high speed video dataset, a common benchmark for visual tracking task. The dataset contains more diverse scenes for large-scale training. It is also a better fit to the task of estimating flow as most of the videos contain dynamic motion of objects in a scene with a close to static background compared to (Nah, Kim, and Lee 2017) where egomotion is predominant. Out of 100 videos in the dataset, 70 are used for training and the remaining videos are used for validation and testing. We call the motion-blur dataset generated from (Galoogahi et al. 2017) as the NfS blur dataset. In case of real blur datasets, the optical flow between the first (1st) and last (7th) frames, which is used as a self-supervision during training, is obtained using pretrained state-of-the-art models (Sun et al. 2018; Ilg et al. 2017).

Methodology

In this section, we explain our framework and details of the training process. Our model has 3 main components: A feature encoder, feature decoder and flow estimator (see Fig. 2).

Feature encoder. The feature encoder extracts features at different scales from a given motion-blurred image. It is a feed-forward convolutional network which has six convolutional blocks each with two layers of convolutions of kernel

size 3×3 and stride size of 2 and 1, respectively, with a ReLU nonlinearity following each convolutional layer. During the feature extraction stage, features are downsampled to half of their spatial size after each convolutional block. Given an input image I , the feature encoder outputs k features: $\{U_e^l\}_{l=1}^k$, where U_e^l is an encoded feature at level l (see Fig. 2a).

Feature decoder. The feature decoder is composed of spatial transformer networks (STNs) (Jaderberg et al. 2015) and feature refining blocks to decode the encoded features into first and last frame features in a bottom-up fashion. Given a feature $U_e^l \in \mathbb{R}^{H \times W \times C}$ with width W , height H and C channels from the encoder, the STN predicts a global transformation parameters θ conditioned on the input to spatially transform the feature, hence, learning non-local motion from the feature of a motion-blurred input (Eq. (1)). In order to take account of locally variant motions which are apparent in dynamic scenes, the transformed feature is passed through a refining network. Along with the transformed feature, we input the feature from the encoder into the refining network in order to guide the network to infer the relative spatial motion (see Fig. 2a). A residual connection is also built by upsampling decoded features from previous feature level using a transposed convolution (deconvolution) layer of kernel size 4×4 and stride size of 2.

At each feature level l , the refining network \mathcal{R}^l takes the transformed feature U_t^l , the encoded feature U_e^l , and the up-scaled decoded feature from the previous feature level $l + 1$ concatenated together channel-wise and outputs a decoded feature (Eq. (2)). The refining network contains five densely connected convolutional layers each with kernel size 3×3 and stride size 1. We used individual feature decoders to reconstruct first and last image features at different levels of abstraction. These features are then used to estimate optical flow, mimicking the vanilla pipeline for optical flow estimation from two images as shown in Fig. 2b.

$$U_t^l = \mathcal{T}_\theta^l \{U_e^l\} \quad (1)$$

$$V^l = \mathcal{R}^l \{U_t^l \oplus U_e^l \oplus \text{up.}(V^{l+1})\} \quad (2)$$

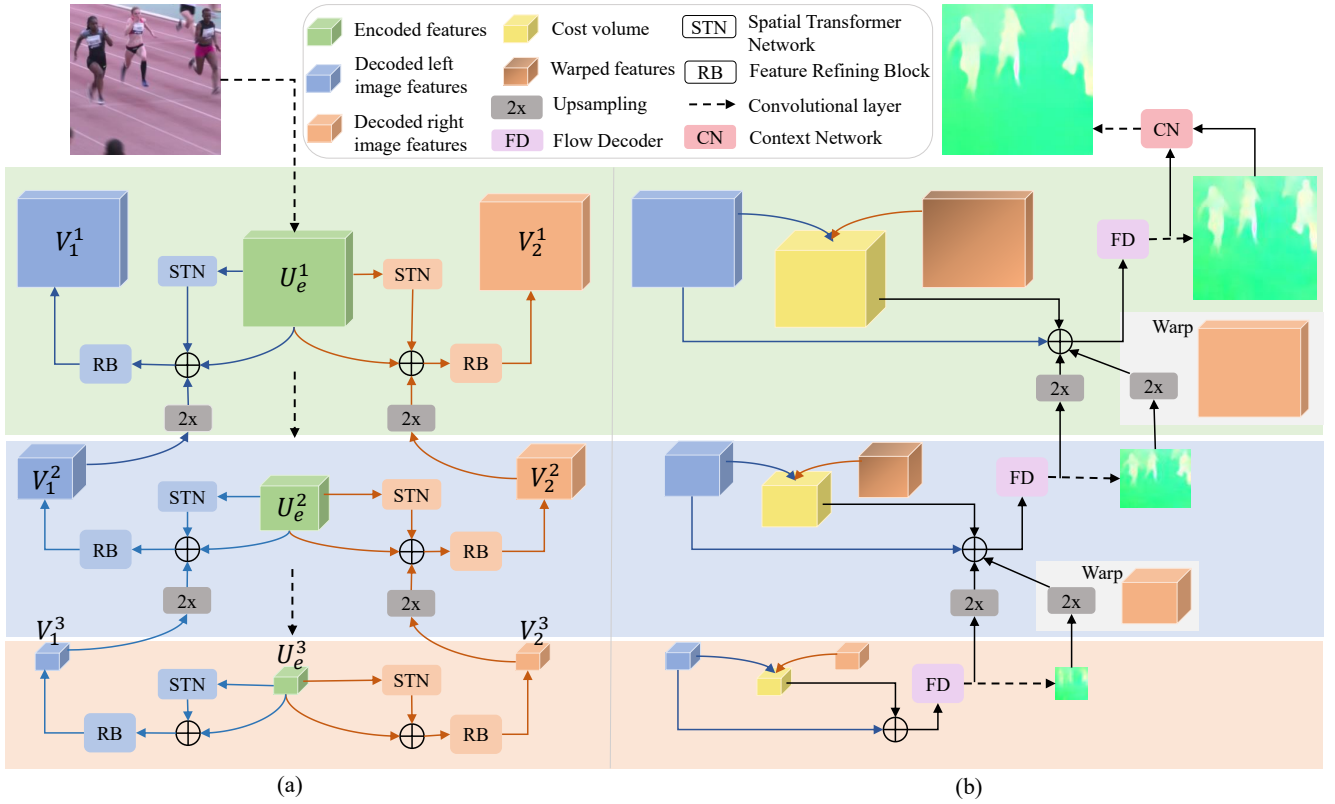


Figure 2: Overview of our network. (a) Feature encoding and decoding (b) Flow estimation. Our network estimates flows at 6 different scales. We only visualize 3 levels for simplicity (Best viewed in color).

where $l = \{1, \dots, k\}$, \mathcal{T}_θ denotes transformation by STN, $\mathbf{up}.$ denotes upsampling, \oplus denotes channel-wise concatenation, and V^l is the decoded feature at level l .

Flow estimator. The flow estimator computes optical flow at different scales using the decoded first and last frame features. Note that the decoded features here are equivalent to the encoded features of two clean input images in standard optical flow estimation algorithms. Inspired by recent network architectures for optical flow estimation (Fischer et al. 2015; Ilg et al. 2017; Ranjan and Black 2017; Sun et al. 2018), we reconstruct flow maps from coarse-to-fine using cost volume, warping, flow decoding, and prediction layers. At each feature level l , the warping layer \mathcal{W} warps the last frame feature V_2^l at level l with an upsampled flow estimated at previous scale $l + 1$ (Eq. (3)). A deconvolution layer of kernel size 4×4 and stride size of 2 was used to upsample flows at different scales.

Given a first image feature V_1^l and a backward-warped last image feature \hat{V}_2^l , the cost volume layer \mathcal{C} computes the matching cost between the features using a correlation layer (Fischer et al. 2015; Xu, Ranftl, and Koltun 2017; Sun et al. 2018). The flow decoding network \mathcal{D} takes the output of the cost volume layer and decodes a flow feature V_f^l that will be used by the flow prediction layer to estimate flow. In addition to the correlation output, the flow decoder takes

the first image feature, the upscaled flow and decoded flow feature from previous feature level (Eq. (4)). Like the refining block in the image feature decoder, the flow decoder has five densely connected convolutional layers each with kernel size 3×3 and stride size 1.

The flow prediction layer \mathcal{P} inputs the decoded flow feature and predicts optical flow. It is a single convolutional layer that outputs a flow f^l of size $H \times W \times 2$ given a feature $V_f^l \in \mathbb{R}^{H \times W \times C}$. Flow maps are estimated at different scales from the smallest to the highest resolution (Eq. (5)). The estimated full-scale flow is further refined by aggregating contextual information using a context network (Sun et al. 2018; Im et al. 2019). It contains seven dilated convolutions (Yu and Koltun 2015) with receptive field size of 1,2,4,8,16,1 and 1, respectively. The context network takes the final decoded flow feature and the predicted flow as an input and outputs a refined flow. Each convolutional layer in the flow decoder and flow estimator is followed by a ReLU activation layer.

$$\hat{V}_2^l = \mathcal{W}\{V_2^l, \mathbf{up}.(f^{l+1})\} \quad (3)$$

$$V_f^l = \mathcal{D}^l\{\mathcal{C}(V_1^l, \hat{V}_2^l) \oplus V_1^l \oplus \mathbf{up}.(V_f^{l+1}) \oplus \mathbf{up}.(f^{l+1})\} \quad (4)$$

$$f^l = \mathcal{P}^l\{V_f^l\} \quad (5)$$

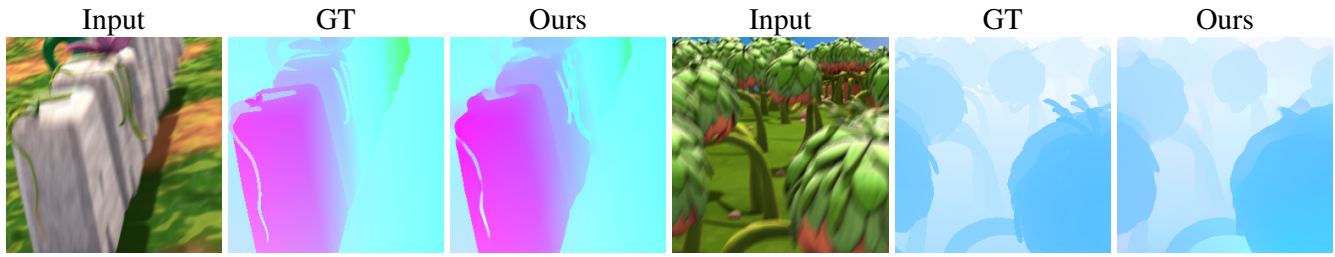


Figure 3: Qualitative results on Monkaa blur dataset. The 1st and 4th columns are motion-blurred inputs generated by temporally averaging intermediate frames interpolated between first and last images in (N.Mayer et al. 2016). The 2nd and 5th columns are ground truth flows. The 3rd and 6th columns are optical flows estimated by our model given the motion-blurred input.




Figure 4: Qualitative results on GoPro and Nfs blur datasets. We compare the optical flow estimated by our model with pseudo-ground truth optical flow (p-GT) between the first and last frames in a high-speed video sequence (predicted using pretrained optical flow network (Sun et al. 2018; Ilg et al. 2017)) from which the motion blurred image was generated.

Network training. Extracting motion information from a single motion-blurred image is an ill-posed problem without extra information (from external sensors such as IMU or other clues on the camera motion (Lee et al. 2019)) as averaging destroys temporal order (Jin, Meishvili, and Favaro 2018; Purohit, Shah, and Rajagopalan 2019; Qiu et al. 2019). We experimentally find that a weighted multi-scale endpoint error (EPE) between the estimated flows and the downsampled pseudo-ground truth flows is a sufficient constraint for network training and convergence (Eq. (6)). We used a bilinear interpolation to downsample the pseudo-ground truth flow to respective sizes at different scales. An attempt to use ℓ_1 photometric loss for supervision or smoothness loss as a regularizer did not improve network performance.

$$\mathcal{L} = \sum_{l=1}^k w_l \cdot |f_l - \hat{f}_l|_2 \quad (6)$$

where f_l is the estimated flow, \hat{f}_l is the downsampled pseudo-ground truth flow and w_l is the loss weight coefficient at scale l .

Given a pseudo-ground truth optical flow as a supervision during training, our network learns to decode the first and

last image features from encoded features of an input image in a symmetric manner. The decoded features are then used to estimate flow. This task, however, can potentially suffer from an ambiguity in predicting the correct flow direction during inference. Given a single motion-blurred image, estimating the flow direction (either from first \rightarrow last or last \rightarrow first) is a highly intractable problem as reversely averaging sequential frames result in the same motion-blurred image. Hence, to purely measure the quality the predicted flows without flow direction estimation issues, we evaluate both forward and backward direction flows and report the lower EPE in the experiment section. Despite the existence of temporal ambiguity, optical flow from a motion-blurred image has many practical applications in computer vision research such as 3D reconstruction, segmentation from flow and motion-blur removal (see Downstream tasks).

Experiment

Implementation details. We estimate flows at 6 different feature levels with training loss weight coefficients set as follows: $w_6 = 0.32$, $w_5 = 0.08$, $w_4 = 0.04$, $w_3 = 0.02$, $w_2 = 0.01$ and $w_1 = 0.005$ from the lowest to the highest

Method	Monkaa	GoPro	NfS
B2F-net	1.158	2.077	-
B2F-net + <i>fine tuning</i>	-	2.038	1.958

Table 1: Quantitative evaluation. For simplicity, we refer to our Blur to Flow network as B2F-net.

resolution, respectively. At each level, we use a correlation layer with a neighborhood search range of 4 pixels and stride size of 1. We chose Adam (Kingma and Ba 2015) as an optimization method with parameters β_1 , β_2 and *weight decay* fixed to 0.9, 0.999 and $4e - 4$, respectively. In all experiments, a mini-batch size of 4 and image size of 256×256 is used by centrally cropping inputs. Following (Fischer et al. 2015), we train the Monkaa blur dataset for 300 epochs with initial learning rate $\lambda = 1e - 4$. We gradually decayed the learning rate by half at 100, 150, 200 and 250 epochs during training. For the GoPro and NfS blur datasets, we trained (fine-tuned) the model for 120 epochs with a learning rate initialized with $\lambda = 1e - 4$ and decayed by half at 60, 80 and 100 epochs.

Qualitative evaluation. For the Monkaa (N.Mayer et al. 2016) blur dataset, we compare our results with the ground truth optical flow between the first and last frames which we used to interpolate intermediate frames and synthesize a motion-blurred image. Our model successfully estimates optical flow from synthetic blurred inputs with different blur magnitudes and patterns (see Fig. 3). Fig. 4 shows test results on the real image blur datasets (Nah, Kim, and Lee 2017; Galoogahi et al. 2017). In order to evaluate whether our model reasonably inferred the flow from a given real motion-blurred image, our results are compared with the optical flow between sharp initial and final frames in a high speed video (from which the blurred image was temporally averaged) predicted by pretrained state-of-the-art models (Sun et al. 2018; Ilg et al. 2017). These flows are later used as pseudo-ground truth (p-GT) for quantitative evaluation. The qualitative results on real image blur datasets show that our model accurately predicts the motion of objects in different motion-blur scenarios with dynamic motion of multiple objects in a close to static or moving scenes. Failure cases occur for temporally undersampled and heavily blurred samples since image contents of such samples are often destroyed and, hence, feature decoding usually fails.

Quantitative evaluation. We compute the end-point error between the estimated optical flow from a motion-blurred input and the flow between the sharp first and last images (pseudo-ground truth). Due to flow direction ambiguity, we calculate both the forward and backward flows during testing and report the lower metric *i.e.* $\min\{EPE(f_{1 \rightarrow 2}), EPE(f_{2 \rightarrow 1})\}$. The averaged test results on different datasets are summarized in Table 1. The test error for the Monkaa blur dataset is lower compared to real datasets as the simulated blurs from the synthetic scenes are mostly static. We experimentally find that training our




Figure 5: Comparison with motion flow estimation works. The 2nd and 3rd rows depict outputs of Gong et al. and our model, respectively.

model from scratch with random weight initialization converges well on the GoPro blur dataset, but does not converge on the NfS blur dataset. This is mainly because the NfS blur dataset contains various low quality videos with some of them containing out-of-focus frame sequences, and hence the resulting blur dataset is challenging for our network to learn. To mitigate this issue, we used a model pretrained on the synthetic dataset and fine-tuned the model on the NfS blur dataset. Training in this manner resulted in a good network performance. Moreover, fine-tuning also lowered the endpoint error on the GoPro blur dataset (see Table 1).

Analysis

Motion flow estimation. We qualitatively compare our method with previous works that estimate motion flow from a blurred image for deblurring purpose (Hyun Kim and Mu Lee 2014; Gong et al. 2017). In order to perform a fair comparison, both our model and the model from (Gong et al. 2017) are evaluated on motion-blurred examples from a real blur dataset (Jianping, Li, and Jiaya 2014). As can be inferred from Fig. 5, our model generalizes noticeably well by estimating a reasonable flow from the given blurred inputs. On the other hand, it is hardly possible to analyse the predictions from (Gong et al. 2017) as their model fails to output an interpretable optical flow. This is mainly because (Gong et al. 2017) treats the task of estimating flow from a motion-blurred image as a classification problem and predicts discrete integer vectors at each pixel with constrained flow direction and magnitude. Moreover, due to the nature of the problem setup, their network could be trained only on synthetic images with simulated motion blurs, hence making it difficult to generalize for real motion-blurred cases. In regard to these two aspects, our model treats optical flow estimation from a motion-blurred image as a regression problem with no specific constraint and can be trained on motion-blurred images from real high speed videos, which leads to the better generalization capability of our model.

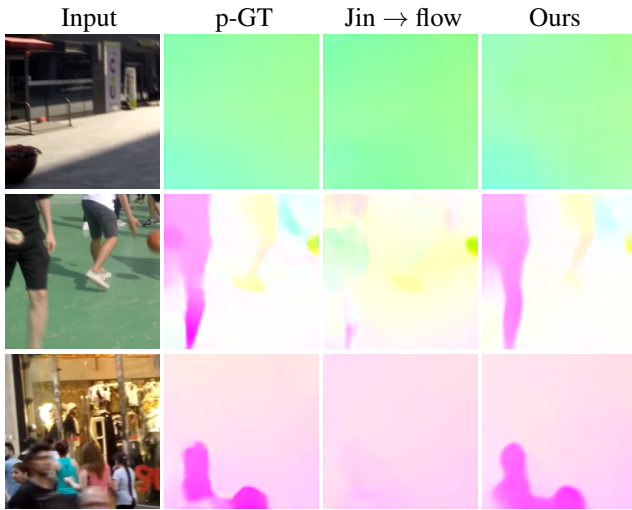


Figure 6: Comparison with flow from restored frames. The 3rd column depicts the flows computed between the first and last frames in the restored sequence (Jin, Meishvili, and Favaro 2018). The 4th column shows our model prediction.

Flow from restored frames. Following recent works in video sequence reconstruction from a motion-blurred image (Jin, Meishvili, and Favaro 2018; Purohit, Shah, and Rajagopalan 2019), a naive approach to the task at hand would be to restore left and right end frames and estimate optical flow from the restored frames using standard approaches. Here we compare the optical flow of a motion-blurred input estimated by our model with a flow predicted by PWC-Net from restored left and right frames using (Jin, Meishvili, and Favaro 2018). Quantitatively, this approach performs considerably worse giving an endpoint error of magnitude 5.875 on the GoPro blur dataset compared to our approach (2.077). The qualitative results in Fig. 6 also show that our approach outputs more accurate results. This performance gap can be directly attributed to the fact that restored frames usually contain multiple motion artifacts. Moreover, we experimentally find that estimating flow from restored frames works relatively well for uniform motion blurs (1st row in Fig. 6) since sequence restoration methods like Jin *et al.* generally learn the global camera motion in a static scene. However, for dynamic blurs (2nd and 3rd rows in Fig. 6), these methods likely output inaccurate optical flows as they often fail to correctly capture locally varying motions. On the contrary, our method successfully captures global and local motions.

Downstream Tasks

Moving object segmentation. Optical flow is commonly used in video object segmentation task along with appearance information to segment moving objects in a video. (Jain, Xiong, and Grauman 2017; Hu, Huang, and Schwing 2018; Hu et al. 2018). To highlight the applicability of the flow estimated by our approach, we experimented with the task of segmenting generic moving objects from a motion-blurred input. For this purpose, we used a pretrained Fusion-Seg model from (Jain, Xiong, and Grauman 2017) which




Figure 7: Qualitative analysis on the application of optical flow for moving object segmentation in a blurry scenario.

contains an appearance stream (inputs an image), a motion stream (inputs an optical flow) and a fusion of the two networks. As shown in Fig. 7, segmenting moving objects from a motion-blurred input often results in a failure or inaccurate segmentation masks mainly because object boundaries and appearance cues that are crucial for segmentation are corrupted by blur. On the other hand, by feeding the optical flow estimated by our approach into the motion stream network, we obtained accurate segmentation results. The joint model that uses both appearance and motion information also leverages the estimated optical flow to segment the moving object in the given blurred input. This showcases the applicability of the estimated flow for moving object segmentation in a blurry scenario.

Motion-blur removal. Given a blurred image and the estimated motion kernel, previous works (Couzinie-Devy et al. 2013; Sun et al. 2015; Gong et al. 2017) recover a sharp image by performing a non-blind deconvolution algorithm (Please refer to section 3.2 of Gong et al. for details). In order to examine the effectiveness of our estimated flows for blur removal, we also directly used our estimated flow in the deblurring process and compared our method with competing approaches (Sun et al. 2015; Gong et al. 2017). For quantitative comparison, we used the official code of (Gong et al. 2017) to generate two types of motion-blurred test sets: BSD-M (with maximum pixel displacement of 17) and BSD-S (with maximum pixel displacement of 36). Please refer to section 5.1 of (Gong et al. 2017) for dataset details. Each dataset contains 300 motion-blurred images of size 480×320 . The quantitative and qualitative results are shown in Table 2 and Fig. 8, respectively. As can be inferred from the results, our approach performs favourably against previous works (Sun et al. 2015; Gong et al. 2017) on recovering sharp frames from a blurred image. These results are a direct consequence of better motion-flow estimation as briefly analyzed in the previous section.



Figure 8: Qualitative comparison on motion-blur removal task using the estimated flow.

Ablation Studies

Importance of feature decoding. In order to show the importance of the feature decoder in our network, we experimented with a U-net (Ronneberger, Fischer, and Brox 2015) like architecture where we estimated flow only from encoded features without explicitly decoding left and right image features *i.e.* no spatial transformer networks (STN) and feature refining blocks (RB). A network trained in this manner converged with a higher training error resulting in a worse test error of magnitude 2.748 EPE (32.30% error increment compared to a model with STN and RB).

Spatial transformer network. STNs (Jaderberg et al. 2015) in our network learn non-local motions (*e.g.* camera motion) and estimate global transformation parameters to spatially transform the encoded feature of a given motion-blurred input accordingly. They are one of the building blocks of our network and are notably important for the performance of our model. As can be seen from Table 3, a network with STN lowers the endpoint error by 8.21% in comparison with a model without STN.

Feature refining block. The feature refining block (RB) decodes left and right image features for the flow estimator. As the STNs only predict global feature transformation parameters, non-local motions which are predominant in real world scenarios are taken into account by passing the transformed feature in the feature refining block. Our experimental results also verify the significance of feature refining blocks as it can be inferred from Table 3. Training a model without feature refining blocks increments the endpoint error by 14.73%.

Conclusion

Motion blur, in general, is regarded as an undesirable artifact. However, it contains motion information which can be processed into a more interpretable form. In this work, for the first time, we tackle the problem of estimating optical

	BSD-M		BSD-S	
	PSNR (dB)	SSIM	PSNR (dB)	SSIM
Sun <i>et al.</i>	22.97	0.674	20.53	0.530
Gong <i>et al.</i>	23.88	0.718	21.85	0.625
Ours	25.23	0.786	23.41	0.714

Table 2: Motion-blur removal via non-blind deconvolution

STN	RB	EPE (\downarrow)
✓	✓	2.077
✗	✓	2.263
✓	✗	2.383
✗	✗	2.748

Table 3: Ablation studies on GoPro blur dataset for different network components

flow from a single motion-blurred image in a data-driven manner. We propose a novel and intuitive framework for the task, and successfully train it using a transfer learning approach. We show the effectiveness and generalizability of our method through a large set of experiments on synthetic and real motion-blur datasets. We also carry out in-depth analysis of our model in comparison to related approaches, guiding that naively deploying sequence restoration methods followed by standard optical flow estimation fails on this problem. The applicability of our work is also demonstrated on motion deblurring and segmentation tasks. Overall, our approach introduces a new interesting perspective on motion blur in connection with future applications such as motion deblurring, temporal super resolution and video sequence restoration from a motion-blurred image.

Acknowledgements

This work was supported by NAVER LABS Corporation [SSIM: Semantic & scalable indoor mapping].

References

- Chakrabarti, A.; Zickler, T.; and Freeman, W. T. 2010. Analyzing spatially-varying blur. In *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, 2512–2519. IEEE.
- Cho, S.; and Lee, S. 2009. Fast motion deblurring. *ACM Transactions on graphics (TOG)* 28(5): 145.
- Couzinie-Devy, F.; Sun, J.; Alahari, K.; and Ponce, J. 2013. Learning to estimate and remove non-uniform image blur. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1075–1082.
- Dai, S.; and Wu, Y. 2008. Motion from blur. In *2008 IEEE Conference on Computer Vision and Pattern Recognition*, 1–8. IEEE.
- Fergus, R.; Singh, B.; Hertzmann, A.; Roweis, S. T.; and Freeman, W. T. 2006. Removing camera shake from a single photograph. In *ACM transactions on graphics (TOG)*, 787–794. ACM.
- Fischer, P.; Dosovitskiy, A.; Ilg, E.; Häusser, P.; Hazırbaş, C.; Golkov, V.; Van der Smagt, P.; Cremers, D.; and Brox, T. 2015. FlowNet: Learning optical flow with convolutional networks. *arXiv preprint arXiv:1504.06852*.
- Galoogahi, H. K.; Fagg, A.; Huang, C.; Ramanan, D.; and Lucey, S. 2017. Need for Speed: A Benchmark for Higher Frame Rate Object Tracking. *arXiv preprint arXiv:1703.05884*.
- Gong, D.; Yang, J.; Liu, L.; Zhang, Y.; Reid, I.; Shen, C.; van den Hengel, A.; and Shi, Q. 2017. From Motion Blur to Motion Flow: A Deep Learning Solution for Removing Heterogeneous Motion Blur. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Gupta, A.; Joshi, N.; Zitnick, C. L.; Cohen, M.; and Curless, B. 2010. Single image deblurring using motion density functions. In *European Conference on Computer Vision*, 171–184. Springer.
- Horn, B. K.; and Schunck, B. G. 1981. Determining optical flow. *Artificial intelligence* 17(1-3): 185–203.
- Hu, P.; Wang, G.; Kong, X.; Kuen, J.; and Tan, Y.-P. 2018. Motion-Guided Cascaded Refinement Network for Video Object Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Hu, Y.-T.; Huang, J.-B.; and Schwing, A. G. 2018. Unsupervised Video Object Segmentation using Motion Saliency-Guided Spatio-Temporal Propagation. In *Proceedings of the European Conference on Computer Vision (ECCV)*.
- Hyun Kim, T.; and Mu Lee, K. 2014. Segmentation-free dynamic scene deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2766–2773.
- Ilg, E.; Mayer, N.; Saikia, T.; Keuper, M.; Dosovitskiy, A.; and Brox, T. 2017. FlowNet 2.0: Evolution of Optical Flow Estimation with Deep Networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. URL <http://lmb.informatik.uni-freiburg.de/Publications/2017/IMKDB17>.
- Im, S.; Jeon, H.-G.; Lin, S.; and Kweon, I. S. 2019. DP-SNet: End-to-end Deep Plane Sweep Stereo. In *International Conference on Learning Representations*. URL <https://openreview.net/forum?id=ryeYHi0ctQ>.
- Jaderberg, M.; Simonyan, K.; Zisserman, A.; et al. 2015. Spatial transformer networks. In *Conference on Neural Information Processing Systems*.
- Jain, S. D.; Xiong, B.; and Grauman, K. 2017. Fusionseg: Learning to combine motion and appearance for fully automatic segmentation of generic objects in videos. In *2017 IEEE conference on computer vision and pattern recognition (CVPR)*, 2117–2126. IEEE.
- Jiang, H.; Sun, D.; Jampani, V.; Yang, M.; Learned-Miller, E. G.; and Kautz, J. 2018. Super SloMo: High Quality Estimation of Multiple Intermediate Frames for Video Interpolation. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Jianping, S.; Li, X.; and Jiaya, J. 2014. Discriminative Blur Detection Features. In *IEEE Conference on Computer Vision and Pattern Recognition*.
- Jin, M.; Meishvili, G.; and Favaro, P. 2018. Learning to Extract a Video Sequence From a Single Motion-Blurred Image. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Kingma, D. P.; and Ba, J. 2015. Adam: A Method for Stochastic Optimization. In *International Conference on Learning Representations*.
- Lee, H. S.; Kwon, J.; and Lee, K. M. 2011. Simultaneous localization, mapping and deblurring. In *IEEE International Conference on Computer Vision*.
- Lee, S.; Kim, J.; Oh, T.-H.; Jeong, Y.; Yoo, D.; Lin, S.; and Kweon, I. S. 2019. Visuomotor understanding for representation learning of driving scenes. *arXiv preprint arXiv:1909.06979*.
- Levin, A. 2007. Blind motion deblurring using image statistics. In *Advances in Neural Information Processing Systems*, 841–848.
- Li, W.; Chen, Y.; Lee, J.; Ren, G.; and Cosker, D. 2014. Robust optical flow estimation for continuous blurred scenes using rgb-motion imaging and directional filtering. In *IEEE Winter Conference on Applications of Computer Vision*, 792–799. IEEE.
- Nah, S.; Kim, T. H.; and Lee, K. M. 2017. Deep Multi-Scale Convolutional Neural Network for Dynamic Scene Deblurring. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- N.Mayer; E.Ilg; P.Häusser; P.Fischer; D.Cremers; A.Dosovitskiy; and T.Brox. 2016. A Large Dataset to Train Convolutional Networks for Disparity, Optical Flow, and Scene Flow Estimation. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*. URL <http://lmb.informatik.uni-freiburg.de/Publications/2016/MIFDB16>. ArXiv:1512.02134.

Pan, J.; Hu, Z.; Su, Z.; Lee, H.-Y.; and Yang, M.-H. 2016. Soft-segmentation guided object motion deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 459–468.

Portz, T.; Zhang, L.; and Jiang, H. 2012. Optical flow in the presence of spatially-varying motion blur. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, 1752–1759. IEEE.

Purohit, K.; Shah, A.; and Rajagopalan, A. 2019. Bringing alive blurred moments. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 6830–6839.

Qiu, J.; Wang, X.; Maybank, S. J.; and Tao, D. 2019. World From Blur. In *The IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

Ranjan, A.; and Black, M. J. 2017. Optical flow estimation using a spatial pyramid network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4161–4170.

Rekleitis, I. 1995. *Visual motion estimation based on motion blur interpretation*. Ph.D. thesis, Citeseer.

Ronneberger, O.; Fischer, P.; and Brox, T. 2015. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*.

Schoueri, Y.; Scaccia, M.; and Rekleitis, I. 2009. Optical flow from motion blurred color images. In *2009 Canadian Conference on Computer and Robot Vision*, 1–7. IEEE.

Seok Lee, H.; and Mu Lee, K. 2013. Dense 3d reconstruction from severely blurred images using a single moving camera. In *IEEE Conference on Computer Vision and Pattern Recognition*.

Sun, D.; Yang, X.; Liu, M.-Y.; and Kautz, J. 2018. PWC-Net: CNNs for Optical Flow Using Pyramid, Warping, and Cost Volume. In *IEEE Conference on Computer Vision and Pattern Recognition*.

Sun, J.; Cao, W.; Xu, Z.; and Ponce, J. 2015. Learning a convolutional neural network for non-uniform motion blur removal. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 769–777.

Tu, Z.; Poppe, R.; and Veltkamp, R. 2015. Estimating accurate optical flow in the presence of motion blur. *Journal of Electronic Imaging* 24(5): 053018.

Wieschollek, P.; Schölkopf, B.; Lensch, H. P.; and Hirsch, M. 2016. End-to-end learning for image burst deblurring. In *Asian Conference on Computer Vision*, 35–51. Springer.

Xu, J.; Ranftl, R.; and Koltun, V. 2017. Accurate optical flow via direct cost volume processing. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1289–1297.

Yu, F.; and Koltun, V. 2015. Multi-scale context aggregation by dilated convolutions. *arXiv preprint arXiv:1511.07122*.

Zheng, S.; Xu, L.; and Jia, J. 2013. Forward motion deblurring. In *Proceedings of the IEEE international conference on computer vision*, 1465–1472.