Animation from Blur: Multi-modal Blur Decomposition with Motion Guidance

Zhihang Zhong^{1,3}, Xiao Sun², Zhirong Wu², Yinqiang Zheng¹, Stephen Lin², and Imari Sato^{1,3}

> ¹ The University of Tokyo, zhong@is.s.u-tokyo.ac.jp ² Microsoft Research Asia ³ National Institute of Informatics

Abstract. We study the challenging problem of recovering detailed motion from a single motion-blurred image. Existing solutions to this problem estimate a single image sequence without considering the motion ambiguity for each region. Therefore, the results tend to converge to the mean of the multi-modal possibilities. In this paper, we explicitly account for such motion ambiguity, allowing us to generate multiple plausible solutions all in sharp detail. The key idea is to introduce a motion guidance representation, which is a compact quantization of 2D optical flow with only four discrete motion directions. Conditioned on the motion guidance, the blur decomposition is led to a specific, unambiguous solution by using a novel two-stage decomposition network. We propose a unified framework for blur decomposition, which supports various interfaces for generating our motion guidance, including human input, motion information from adjacent video frames, and learning from a video dataset. Extensive experiments on synthesized datasets and real-world data show that the proposed framework is qualitatively and quantitatively superior to previous methods, and also offers the merit of producing physically plausible and diverse solutions. Code is available at https://github.com/zzh-tech/Animation-from-Blur.

Keywords: Deblurring, multi-modal image-to-video, deep learning

1 Introduction

Motion blur appears in an image when the recorded scene undergoes change during the exposure period of the camera. Although such blur may impart a dynamic quality to a photograph, it is often desirable to invert this blur to produce a sharper image with clearer visual detail. Conventionally, this deblurring task is treated as a one-to-one mapping problem, taking a motion blurred image as input and producing a single output image corresponding to a single time instant during the exposure. Recently, attention has been drawn to a more challenging problem of extracting an image sequence [11] instead, where the images correspond to multiple time instances that span the exposure period, thus forming a short video clip from the blurred motion.



Fig. 1: This paper studies the challenging problem of recovering the sharp image sequence from blurry input caused by relative motion. We first identify the fundamental directional motion ambiguity in this problem, which prevents the network to learn effectively. In order to solve this ambiguity, we propose to use motion as a guidance for conditional blur decomposition. Our approach demonstrates strong performance for recovering sharp visual details in motions.

Blur decomposition from a single blurry image faces the fundamental problem of motion ambiguity [27]. Each independent and uniform motion blurred region in an image can correspond to either a forward or a backward motion sequence, both of which are plausible without additional knowledge. With multiple motion blurred regions in an image, the number of potential solutions increases exponentially, with many that are physically infeasible. For example, in Fig. 1, there exists multiple human dance movements that can correspond to the same observed blurry image, since the movement of each limb of the dancer can be independent. However, existing methods for blur decomposition are designed to predict a single solution among them. Moreover, this directional ambiguity brings instability to the training process, especially when the motion pattern is complex. As a result, this ambiguity, when left unaddressed as in current methods, would lead to poorly diversified and low-quality results.

In this work, we introduce a motion guidance representation to tackle the inherent directional motion ambiguity. The motion guidance is an optical flow representation quantized into four major quadrant directions, describing the motion field only roughly. Given the input blurry image, conditioned on the compact motion guidance, the blur decomposition now becomes a nearly deterministic one-to-one mapping problem without directional ambiguity. Empirically, we find that the decomposition network shows significantly better training convergence with this conditioning on an additional guidance input.

Given the blurry image and additional motion guidance as inputs, we propose a two-stage blur decomposition network to predict the image sequence. The first stage expands the blurry image into an image sequence based on the motion guidance, and the second stage refines the visual details in a residual fashion to generate high-quality images. Our unified framework supports various interfaces for motion guidance acquisition: 1) Through a guidance predictor network learned from a dataset. The guidance predictor network is a VAE-GAN [18,44] which learns a multi-modal distribution over plausible guidances. During inference, given an input blurry image together with sampled noises, the predictor can produce multiple guidance maps. 2) By additional information of dense optical flow computed from temporal frames. When the blurry image input is captured and sampled from a video sequence. The optical flow field between the blurry input and its adjacent frame can be used to compute the motion guidance. This motion guidance reflects the actual motion direction. 3) Through user input. Since the guidance is simple and compact, it can be annotated by outlining the region contours and their motion directions interactively.

To train our model, we synthesize blurry images from high speed videos following the pipeline of REDS [22]. Specifically, we validate the performance of our model on human dance videos from Aist++ [20] which only contain motion blur caused by human movement from a static camera, and general scene videos (GO-PRO [23], DVD [31]), which are dominated by camera motion. Our approach provides a significant qualitative and quantitative improvement over previous methods by introducing a novel motion guidance representation to address the fundamental directional ambiguity in blur decomposition. Furthermore, due to the compactness of motion guidance representation, our unified framework may only need to be trained once, while supporting various decomposition scenarios under different modalities. The motion guidance obtained through different interfaces and their corresponding decomposition results reflect physically plausible and diverse solutions from our multi-modal framework.

2 Related Works

In this section, we review the related literature on image and video deblurring, blur decomposition, as well as multi-modal image translation.

2.1 Deblurring

Deblurring refers to the task of estimating a sharp image from a blurry input, where the blur is often caused by scene or camera motion. Traditional deblurring methods model the blur as a blur kernel operating on a sharp image via a convolution operation. A number of useful priors have been proposed to infer the latent sharp image and the corresponding blur kernel, such as total variation [3], hyper-Laplacian [15], image sparsity [19], and l_0 -norm gradient [36]. Recently, image and video deblurring has benefited from the advancement of deep convolution neural networks (CNNs). A coarse-to-fine pyramid CNN structure is widely used for the single-image deblurring task and achieves impressive performance [23,32,38]. Also, generative adversarial networks (GANs) have been adopted to improve the perceptual quality of deblurring results [16,17]. Temporal dependency across adjacent frames is another source of information which could be utilized to recover the sharp image [31,34]. Recurrent architectures are 4 Z. Zhong et al.

shown be effective at exploiting such temporal information [13,42,24,41]. Extensive studies in these works demonstrate that deep neural networks are able to approximate blur kernels well in an implicit way.

2.2 Blur Decomposition

Blur decomposition attempts to recover the full image sequence from a blurry input caused by object and camera motion during the exposure time. [11] are the first to tackle this problem and propose a pairwise order-invariant loss to improve convergence in training. [27] present a method that extracts a motion representation from the blurry image through self-supervised learning, and then feeds it into a recurrent video decoder to generate a sharp image sequence. [2] utilize an encoder-decoder structure with a spatial transformer network to estimate the middle frame and other frames simultaneously with a transformation consistency loss. Assuming that the background is known, DeFMO [28] embeds a blurred object into a latent space, from which it is rendered into an image sequence that is constrained to be sharp, time-consistent, and independent of the background.

When the input is a blurry video, [10] design a cascade scheme, *i.e.*, deblurring followed by interpolation, to tackle the problem. To avoid errors introduced in the first stage, Shen *et al.* propose BIN [29] with a pyramid module and an inter-pyramid recurrent module to jointly estimate the latent sharp sequence. [1] achieve blurry video decomposition by first estimating the optical flow and then predicting the latent sequence by warping the decoded features. In addition, methods for blur decomposition also consider exploiting high-frequency temporal information from event cameras [25,21].

None of the existing methods address the fundamental ambiguity that exists with motion directions. We are the first to address this by conditioning the decomposition process via a novel motion guidance approach. We also design a flexible blur decomposition network, which can produce diverse decomposition results using guidance from the proposed interfaces for different modalities.

2.3 Image-to-Image/Video Translation

Our work is related to image-to-image translation networks [9] with applications for image synthesis [26] and style transfer [43]. GANs [6,12] and VAEs [14] are two popular approaches for training generative models. Both models can easily be conditioned by feeding additional inputs, such as in conditional VAE [30] and infoGAN [4]. Blur decomposition can be formulated as a similar imageto-image translation problem. However, the inherent motion ambiguity in the blurry image prevents the model from converging to a single optimal solution. We thus introduce a new motion guidance representation to disambiguate the motion directions.

Multi-modal image translation is a promising direction to generate a distribution of results given a single input. Approaches along this direction combine VAE and GAN, so that stochastic variables in VAE could sample from a distribution, while GAN is used to encourage realistic generations. Notable works include VAE-GAN [18] and BicycleGAN [44]. Our guidance predictor follows VAE-GAN [18] to generate multi-modal motion guidance. We note that this technical element is not among the contributions of our paper.

Learning to animate a static image is a fascinating application. Learningbased methods either predict the motion field [7,5] which is later used to generate videos, or directly predict the next frames [39,37]. Our work focuses on animating a specific kind of input image with motion blur, which is physically informative to recover the image sequence.

3 Methodology

The Blur Decomposition Problem. A blurred image can be considered as an average of successive relatively sharp images over the exposure time

$$I_b = \frac{1}{T} \int_0^T I^t dt, \tag{1}$$

which can be approximately expressed in the following discrete form when T is large enough, *i.e.*, $I^1 \cdots I^T$ is a high-frame-rate sequence

$$\frac{1}{T}(I^1 + \dots + I^T) = I_b.$$
 (2)

Note that this image averaging process is simulated in an approximately linear space through applying inverse gamma calibration to RGB images [23,22]. Our goal is to invert this blurring process to estimate a finite sequence of sharp images which are uniformly distributed over the exposure time

$$I_b \xrightarrow{\mathbf{D}} \mathbf{I} = \{I^t, t \in 1, \cdots T\}.$$
(3)

This is a highly ill-posed problem because given the blurry image I_b , there are infinitely many solutions per pixel (u, v) among the sharp images if these sharp images are treated as independent

$$\frac{1}{T}(I^{1}(u,v) + \dots + I^{T}(u,v)) = I_{b}(u,v).$$
(4)

However, assuming the images exist in succession over a short period of time and that the pixels are rarely affected by saturation and occlusions/disocclusions, then the following holds true: the pixels in these images are highly correlated and their dependencies can be described by the optical flow F^t of a sharp image I^t to its next frame I^{t+1} ,

$$I^{t}(u,v) = I^{t+1}(F^{t}(u,v)).$$
(5)

Supposing that we are given the optical flow fields for all frames in the target sequence $\mathbf{F} = \{F^t, t \in 1, \dots, T-1\}$, Equations 4 and 5 form a linear system

6 Z. Zhong et al.



Fig. 2: Architecture of our motion guided blur decomposition network. It consists of two successive stages. The first stage S_1 concatenates the blurry image I_b and motion guidance G as input and outputs a rough image sequence. The second stage S_2 then refines the visual details in a residual fashion.

which can be solved with a unique solution. In other words, the ground truth motions between sharp images resolve the ambiguity problem in Equation 3. Motion-guided Decomposition. We propose to solve the decomposition problem in Equation 3 in two steps. The first step estimates the motions \mathbf{F} during the exposure time, *e.g.*, by learning a motion predictor \mathcal{P} for the blurry image,

$$I_b \xrightarrow{\mathcal{P}} \hat{\mathbf{F}}.$$
 (6)

The second step learns a motion guided decomposition network S that takes both the blurry image and an estimated motion guidance as input to predict the sharp image sequence,

$$(I_b, \mathbf{\hat{F}}) \xrightarrow{\mathcal{S}} \mathbf{I} = \{ \hat{I}^t, t \in 1, \cdots T \}.$$

$$(7)$$

As a result, the ambiguity is explicitly decoupled from the decomposition network S, and only exists in the motion predictor \mathcal{P} .

Predicting the motion guidance with frame-wise optical flow would still be a difficult task. In practice, motion ambiguity will be complex, because the ambiguities of the independent moving regions in the image can be arbitrarily combined. However, we notice that the ambiguity mainly lies in the forward and backward directions of the motions, and hence the motion guidance does not need to be precise to frame-wise and continuous values to resolve the ambiguity. We therefore propose a compact motion guidance representation to replace the accurate optical flow, which enables the decomposition network to adapt to different input modalities. In the following, we describe the compact motion guidance representation in Sec. 3.1, the blur decomposition network conditioned on the motion guidance in Sec. 3.2, and three distinct interfaces for acquiring the motion guidance in Sec. 3.3.

3.1 Motion Guidance Representation

Given accurate dense optical flow sequence $\mathbf{F} = \{F^t, t \in 1, \dots, T-1\}$, the blur decomposition problem can be solved without ambiguity. However, dense optical flow sequence is difficult to obtain and to predict accurately, and hence it may not be an ideal representation as a guidance.

We notice that the ambiguity for blur decomposition is in the motion directions. For example, if $\mathbf{F} = \{F^1, ..., F^{T-1}\}$ is one possible motion of the blurry image I_b and the corresponding sharp image sequence is $\mathbf{I} = \{I^1, ..., I^T\}$, then there also exists another reverse motion $\mathbf{F}_{\text{bac}} = \{F_{\text{bac}}^{T-1}, ..., F_{\text{bac}}^1\}$ with the corresponding sharp image sequence $\mathbf{I}_{\text{inv}} = \{I^T, ..., I^1\}$ which leads to the same blurry image I_b , where F_{bac}^t is the backward optical flow between I^t and I^{t+1} . Providing a crude motion direction may be sufficient to resolve the ambiguity.

Motivated by this observation, we design a compact guidance representation by motion quantization. We first assume that the motion directions within the exposure time do not change abruptly. This is generally true when the shutter speed is not extremely slow compared with object motion. We thus use the aggregated flow to represent the motion pattern for the full sequence $\bar{F} = \sum_{1}^{T} F^t$. We further quantize the aggregated flow into four quadrant directions and an additional motionless class which takes flows below a certain magnitude. Empirically, we find four quadrant directions to be adequate for disambiguating motion directions. We denote the motion guidance as G.

3.2 Motion Guided Blur Decomposition Network

Given the blurry image and the motion guidance, the sharp image sequence is predicted via a blur decomposition network S. Once the model is trained, it can be used to decompose a blurry image into different sharp image sequences simply by providing it with the corresponding motion guidance. In this section, we illustrate the model architecture and training loss in detail.

Training the network requires a dataset of triplet samples (I_b, G, \mathbf{I}) , consisting of the blurry image I_b , the ground truth sharp image sequence \mathbf{I} , and motion guidance G derived from \mathbf{I} . We follow the common practice [23,22] of synthesizing a blurry image by accumulating sharp images over time according to Equation 2, which is implemented in linear space through inverse gamma calibration. We use a off-the-shelf optical flow estimator [33] for deriving \mathbf{F} and the guidance G.

Fig. 2 illustrates our two-stage workflow of S. The first stage estimates a rough dynamic image sequence, and the second stage refines the visual details in a residual fashion. Both networks adopt a similar encoder-decoder architecture.



Fig. 3: Motion guidance acquisition. We represent motion guidance as a quantized motion vector into 4 quadrant directions. Due to its compactness, the guidance could be obtained from a network, videos or user annotations.

The architecture details can be found in the supplementary materials. The twostage blur decomposition network outputs the sharp image sequence prediction $\hat{\mathbf{I}}$. We adopt commonly used \mathcal{L}_2 loss function for supervising the decomposition,

$$\mathcal{L}_2 = \|\mathbf{I} - \hat{\mathbf{I}}\|_2^2. \tag{8}$$

3.3 Motion Guidance Acquisition

We provide three interfaces to acquire the motion guidance: learning to predict the motion guidance, motion from video, and user input, as illustrated in Fig. 3. **Multi-modal Motion Prediction Network.** The directional ambiguity now exists in the motion guidance representation. To account for the ambiguity, we train a multi-modal network to generate multiple physically plausible guidances given a blurry image. We follow the framework of cVAE-GAN [30,18] for multimodal image translation as shown in Fig. 4.

The guidance prediction network comprises an encoder \mathcal{P}_E and a generator \mathcal{P}_G . The encoder \mathcal{P}_E converts the ground truth guidance into a latent stochastic variable by $\mathbf{z} = \mathcal{P}_E(G)$, and the generator predicts the guidance given the latent code and the blurry input $\hat{G} = \mathcal{P}_G(\mathbf{z}, I_b)$. The latent variable \mathbf{z} is considered to follow a Gaussian distribution $\mathcal{N}(0, 1)$, and this stochastic variable is used to model the directional distribution that exists in the motion space. At the testing phase, multi-modal motion predictions can be generated by randomly sampling the latent vector \mathbf{z} from a Gaussian distribution.

The overall network is trained with a combination of a GAN loss \mathcal{L}_{GAN} on the guidance predictions, a VAE loss \mathcal{L}_{VAE} , and a Kullback-Leibler (KL) divergence loss \mathcal{L}_{KL} for the stochastic variable \mathbf{z} as in [44],

$$\mathcal{L}_{\text{guidance}} = \lambda_1 \mathcal{L}_{\text{GAN}}(\mathcal{P}_E, \mathcal{P}_G) + \lambda_2 \mathcal{L}_{\text{VAE}}(\mathcal{P}_E, \mathcal{P}_G) + \lambda_3 \mathcal{L}_{\text{KL}}(\mathbf{z} || \mathcal{N}(0, 1)),$$
(9)

where $\lambda_1, \lambda_2, \lambda_3$ are mixing coefficients for the losses. The architecture details can be found in the supplementary materials.



Fig. 4: Framework of the multi-modal motion guidance prediction network. We mainly follow the cVAE-GAN [30,18], and adapt it to our problem to generate multiple physically plausible guidances from a blurry image.

Motion from Video. We assume that the motion direction does not change abruptly in a short time. Thus, the motion direction in the blurry image can be approximated by the motion to the adjacent frames, if a video is provided. Hence, our method can be directly applied to video deblurring without any modification. In the experiment, our method is also compared with the state-of-the-art video deblurring methods and shows better performance.

Human Annotation. The compact quantized motion guidance provides a friendly interface to our blur decomposition network. Given a blurry image, the user can generate multiple plausible sharp video clips simply by drawing the outlines of the blurry regions and arbitrarily specifying their motion directions.

4 Experiments

We show the details of the used datasets in Sec. 4.1, single image and video decomposition results in Sec. 4.2 and Sec. 4.3, real-world evaluation in Sec. 4.4, guidance robustness analysis in Sec. 4.5, as well as ablation study in Sec. 4.6.

4.1 Datasets

Existing dataset for blur decomposition [29] introduces 1/3 temporal overlaps between adjacent blurred frames. This violates the actual blur occurred in real world. We thus create datasets by our own where blurry frames have almost no temporal overlaps. One of our datasets for general scenes (GenBlur) consists of high-fps videos used by related works [11,27], but with a better pipeline [22] to simulate the formation of blur.

GenBlur. We synthesize the GenBlur dataset using high-frame-rate (240 fps) videos from GOPRO [23], DVD [31], and videos collected by ourselves. To suppress noise and video compression artifact [24], the image resolution is uniformly down-sampled to 960×540 . We follow a widely used blur synthesis technique proposed in [22], that first employs an off-the-shelf CNN [8] to interpolate frames into a much higher fps (240 fps \rightarrow 7680 fps) to avoid unnatural spikes or steps in the blurred trajectory [35]. In the interpolated videos, sets of 128 frames (non-overlapping) are averaged in linear space through inverse gamma calibration to



Fig. 5: Qualitative comparison with single image decomposition baseline method [11]. Given a blurry image, our method can generate multiple physically plausible motion guidance and recover distinct sharp image sequences based on each motion guidance. The baseline Jin *et al.* fails, resulting in a sequence with little motion. Please pay attention to the motion on the hands of the left dancer and the legs of the right dancer.

synthesize blurred images. Instead of using all the 128 frames as ground truth clear images, we evenly sample 7 images among them to keep consistent with the previous work [11]. There are 161 and 31 videos for the train and test set, respectively. This dataset consists mainly of urban scenes, and the motion is dominated by camera ego-motion.

B-Aist++. The problem arising from directional ambiguity is particularly severe when there are multiple independent motions in the image. We thus synthesize another dataset specifically to highlight this issue by using videos from a dance dataset Aist++ [20] which contains complex human body movements by professional dancers. We use the same pipeline as for GenBlur to synthesize blurry images and corresponding sharp image sequences. The synthesized blurry dataset is denoted as B-Aist++. There are 73 and 32 videos for the train and test set. The images are in resolution of 960×720 . This dataset contains complex human motion and the camera is stationary.

4.2 Blurry Image Decomposition

Qualitative Evaluation. We qualitatively compare our result with the stateof-the-art single image based blurry decomposition method [11] in Fig. 5. B-Aist++ is used for comparison because of its complex directional ambiguities. In the second and third rows, we demonstrate diversity in our blur decomposition

Table 1: Quantitative evaluation of single blurry image decomposition. For our method, we predict multiple motion guidance from our guidance predictor network. $P_{\#}$ denotes we evaluate # number of plausible decomposition results for each input, and choose the best case. The results of Jin *et al.* [11] represent the best performance calculated as using either the forward or reverse outputs, following the original paper. Our approach outperforms Jin *et al.* [11] by a large margin even with a single sampling.

Dataset		B-Aist++			GenBlur	
Method	$\mathrm{PSNR}\uparrow$	SSIM \uparrow	$\mathrm{LPIPS}\downarrow$	$\mathrm{PSNR}\uparrow$	SSIM \uparrow	$\mathrm{LPIPS}\downarrow$
Jin et al. [11]	17.01	0.540	0.192	20.88	0.621	0.283
Ours (\mathcal{P}_1)	19.97	0.860	0.089	23.41	0.737	0.267
Ours (\mathcal{P}_3)	22.44	0.898	0.068	23.56	0.740	0.263
Ours (\mathcal{P}_5)	23.49	0.911	0.060	23.61	0.741	0.260

results through the use of different motion guidance sampled from our motion predictor \mathcal{P} . Ground-truth is presented in the last row.

Existing methods such as [11] are unable to resolve the directional ambiguity in motion blur. For example, for the blur in the second dancer's legs, they cannot determine whether the legs are being spread out or drawn in. Consequently, in the case of data with high directional ambiguity such as B-Aist++, the moving range of the generated sharp frames is limited. In contrast, our multi-modal method incorporates directional guidance to remove ambiguity, leading to multiple coherent natural motions that look physically plausible. To better perceive the temporal variation in the decomposition results, we strongly recommend the reader to check out the videos in our supplementary materials.

Quantitative Evaluation. We report quantitative comparison results on B-Aist++ and GenBlur in Table 1. Following common practice, PSNR, SSIM and LPIPS [40] are used as evaluation metrics. Our method can generate multiple decomposition results by sampling multiple motion guidances from the motion predictor \mathcal{P} . In the table, $\mathcal{P}_{\#}$ denotes that # motion guidances are sampled from the predictor and the best result among the samples is reported. It can be seen that our method outperforms Jin *et al.* [11] by a large margin and more samples leads to better best results.

4.3 Blurry Video Decomposition

Qualitative Evaluation. As explained in Sec. 3.3, our method can be directly applied to video based decomposition without modification. The motion guidance used for our decomposition network is quantized from the optical flow between adjacent frames.

Fig. 6 presents a visual comparison between our result and the state-of-theart video-based method BIN [29] on general scenes (GenBlur). Our decomposition results surpass those of the video-based model, with much clearer details.



Fig. 6: Qualitative comparison with blurry video decomposition baseline method BIN [29]. Our approach uses the guidance calculated from the input video itself. For close up investigation, our method recovers much sharper visual details than the baseline method, *e.g.*, the text in the blurry input.

Table 2: Quantitative evaluation of blurry video decomposition. The motion guidance for our method is obtained from the optical flow in the input blurry video. Our method outperforms the baseline BIN [29] by a large margin of 1.16 dB and 2.22 dB absolute PSNR on B-AIST++ and GenBlur respectively.

Dataset	B-Aist++			GenBlur		
Method	$\mathrm{PSNR}\uparrow$	SSIM \uparrow	$\mathrm{LPIPS}\downarrow$	$\mathrm{PSNR}\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}\downarrow$
BIN [29]	22.84	0.903	0.068	24.82	0.805	0.157
Ours (guidance from video)	24.03	0.911	0.067	27.04	0.858	0.122

It is worth noting that since the adjacent frames themselves are blurred, the estimated optical flows are not accurate. However, due to the effective direction augmentation and the need for only coarse directional guidance, the learned decomposition network is robust and effective. Also, please see the videos in our supplementary for better perception.

Quantitative Evaluation. When comparing with the video-based method BIN [29], we use motion guidance estimated from adjacent frames using the off-the-shelf flow estimator [33], denoted as *vid.* in Table. 2. Our method is clearly superior to [29] in terms of all metrics.

4.4 Real-world Evaluation

We further captured real-world blurry image for validating the generalization ability of the proposed method. Because the blurry image is captured in real

Multi-modal Animation from Blur 13



Fig. 7: **Results on real-world captured blurry image.** We manually provide the motion guidance and recover the image sequence using our decomposition network. Although trained on synthetic data, our model successfully generalizes on real-world data. Please notice the relative position between the door frame and the legs of the girl in the bottom.

world, no ground truth sharp image sequences can be used for quantitative evaluation. We display the visual results on real-world data when a user provides directional guidance in Fig. 7. The results demonstrate that our method can generalize well on real data. We also show real-world evaluations with the guidance prediction network in the appendix.

4.5 Robustness of motion guidance

The errors of motion guidance may come from two sources: one by incorrect prediction or human annotation, one by optical flow quantization. For the first case, the model may tolerate guidance errors in the sharp regions. We demonstrate this by showing results using fit, dilated, and eroded guidance in Fig. 8. The decomposition model may identify the pixels are non-blurry and preserve the original details, no matter what guidance prediction is given (dilate result). The model cannot recover a blurry region and may introduce artifacts if it is not provided with a meaningful guidance (erode result). For the second case, please refer to the ablation experiment in the supplementary material, which shows that quantization into 4 bins is sufficient to achieve good performance.

4.6 Ablation Studies

We present ablation studies on motion guidance and the multi-stage architecture. The experiments are conducted on B-Aist++ and the metrics are calculated using ground-truth motion guidance. Based on Fig. 9, introducing motion guidance in the training stage greatly improves convergence by eliminating directional ambiguity. Table 3 shows that introducing a two-stage coarse-and-refinement pipeline while maintaining the model size may bring about 1 dB gain.



Fig. 8: The influence of incorrect prediction or human annotation. Our approach may tolerate guidance errors in the sharp regions.



Fig. 9: Training curves for the decomposition network with and without motion guidance. With extra guidance input to disambiguate the motion direction, the network is able to fit to the data a lot easier.

Table 3: Ablation studies on decomposition network architecture. We ablate the effectiveness of refinement work by studying 1-stage estimation and 2-stage estimation. The refinement network significantly improves the performance. 1-stage and 2-stage models are set to similar sizes by adjusting channel numbers, for fair comparison.

Method	$\mathrm{PSNR}\uparrow$	$\mathrm{SSIM}\uparrow$	$\mathrm{LPIPS}\downarrow$
1-stage	24.47	0.912	0.074
2-stage	25.45	0.933	0.054

5 Conclusions

In this work, we address the problem of recovering a sharp motion sequence from a motion-blurred image. We bring to light the issue of directional ambiguity and propose the first solution to this challenging problem by introducing motion guidance to train networks. The proposed method can adapt to blurry input of different modalities by using the corresponding interfaces including a multimodal prediction network, motion from video, and user annotation. The motion sequences generated by our method are superior to existing methods in terms of quality and diversity.

Acknowledgement

This work was supported by D-CORE Grant from Microsoft Research Asia, JSPS KAKENHI Grant Numbers 22H00529, and 20H05951, and JST, the establishment of university fellowships towards the creation of science technology innovation, Grant Number JPMJFS2108.

References

- 1. Argaw, D.M., Kim, J., Rameau, F., Kweon, I.S.: Motion-blurred video interpolation and extrapolation. In: AAAI Conference on Artificial Intelligence (2021)
- Argaw, D.M., Kim, J., Rameau, F., Zhang, C., Kweon, I.S.: Restoration of video frames from a single blurred image with motion understanding. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 701–710 (2021)
- Chan, T.F., Wong, C.K.: Total variation blind deconvolution. IEEE transactions on Image Processing 7(3), 370–375 (1998)
- Chen, X., Duan, Y., Houthooft, R., Schulman, J., Sutskever, I., Abbeel, P.: Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In: Proceedings of the 30th International Conference on Neural Information Processing Systems. pp. 2180–2188 (2016)
- Endo, Y., Kanamori, Y., Kuriyama, S.: Animating landscape: self-supervised learning of decoupled motion and appearance for single-image video synthesis. arXiv preprint arXiv:1910.07192 (2019)
- Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. Advances in neural information processing systems 27 (2014)
- Holynski, A., Curless, B.L., Seitz, S.M., Szeliski, R.: Animating pictures with eulerian motion fields. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5810–5819 (2021)
- 8. Huang, Z., Zhang, T., Heng, W., Shi, B., Zhou, S.: Rife: Real-time intermediate flow estimation for video frame interpolation. arXiv preprint arXiv:2011.06294 (2020)
- Isola, P., Zhu, J.Y., Zhou, T., Efros, A.A.: Image-to-image translation with conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1125–1134 (2017)
- Jin, M., Hu, Z., Favaro, P.: Learning to extract flawless slow motion from blurry videos. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8112–8121 (2019)
- Jin, M., Meishvili, G., Favaro, P.: Learning to extract a video sequence from a single motion-blurred image. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 6334–6342 (2018)
- Karras, T., Laine, S., Aila, T.: A style-based generator architecture for generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 4401–4410 (2019)
- Kim, T.H., Lee, K.M., Scholkopf, B., Hirsch, M.: Online video deblurring via dynamic temporal blending network. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 4038–4047 (2017)
- Kingma, D.P., Welling, M.: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
- Krishnan, D., Fergus, R.: Fast image deconvolution using hyper-laplacian priors. Advances in neural information processing systems 22, 1033–1041 (2009)
- Kupyn, O., Budzan, V., Mykhailych, M., Mishkin, D., Matas, J.: Deblurgan: Blind motion deblurring using conditional adversarial networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8183–8192 (2018)
- Kupyn, O., Martyniuk, T., Wu, J., Wang, Z.: Deblurgan-v2: Deblurring (ordersof-magnitude) faster and better. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 8878–8887 (2019)

- 16 Z. Zhong et al.
- Larsen, A.B.L., Sønderby, S.K., Larochelle, H., Winther, O.: Autoencoding beyond pixels using a learned similarity metric. In: International conference on machine learning. pp. 1558–1566. PMLR (2016)
- Levin, A., Weiss, Y., Durand, F., Freeman, W.T.: Understanding and evaluating blind deconvolution algorithms. In: 2009 IEEE Conference on Computer Vision and Pattern Recognition. pp. 1964–1971. IEEE (2009)
- Li, R., Yang, S., Ross, D.A., Kanazawa, A.: Ai choreographer: Music conditioned 3d dance generation with aist++. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 13401–13412 (2021)
- Lin, S., Zhang, J., Pan, J., Jiang, Z., Zou, D., Wang, Y., Chen, J., Ren, J.: Learning event-driven video deblurring and interpolation. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part VIII 16. pp. 695–710. Springer (2020)
- Nah, S., Baik, S., Hong, S., Moon, G., Son, S., Timofte, R., Lee, M.K.: Ntire 2019 challenge on video deblurring and super-resolution: Dataset and study. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019)
- Nah, S., Kim, T.H., Lee, K.M.: Deep multi-scale convolutional neural network for dynamic scene deblurring. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3883–3891 (2017)
- Nah, S., Son, S., Lee, K.M.: Recurrent neural networks with intra-frame iterations for video deblurring. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 8102–8111 (2019)
- Pan, L., Scheerlinck, C., Yu, X., Hartley, R., Liu, M., Dai, Y.: Bringing a blurry frame alive at high frame-rate with an event camera. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6820– 6829 (2019)
- Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2337–2346 (2019)
- Purohit, K., Shah, A., Rajagopalan, A.: Bringing alive blurred moments. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 6830–6839 (2019)
- Rozumnyi, D., Oswald, M.R., Ferrari, V., Matas, J., Pollefeys, M.: Defmo: Deblurring and shape recovery of fast moving objects. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3456–3465 (2021)
- Shen, W., Bao, W., Zhai, G., Chen, L., Min, X., Gao, Z.: Blurry video frame interpolation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5114–5123 (2020)
- Sohn, K., Lee, H., Yan, X.: Learning structured output representation using deep conditional generative models. Advances in neural information processing systems 28, 3483–3491 (2015)
- Su, S., Delbracio, M., Wang, J., Sapiro, G., Heidrich, W., Wang, O.: Deep video deblurring for hand-held cameras. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 1279–1288 (2017)
- Tao, X., Gao, H., Shen, X., Wang, J., Jia, J.: Scale-recurrent network for deep image deblurring. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 8174–8182 (2018)
- 33. Teed, Z., Deng, J.: Raft: Recurrent all-pairs field transforms for optical flow. In: European conference on computer vision. pp. 402–419. Springer (2020)

- Wang, X., Chan, K.C., Yu, K., Dong, C., Change Loy, C.: Edvr: Video restoration with enhanced deformable convolutional networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops. pp. 0–0 (2019)
- Wieschollek, P., Hirsch, M., Scholkopf, B., Lensch, H.: Learning blind motion deblurring. In: Proceedings of the IEEE International Conference on Computer Vision. pp. 231–240 (2017)
- Xu, L., Zheng, S., Jia, J.: Unnatural l0 sparse representation for natural image deblurring. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1107–1114 (2013)
- Xue, T., Wu, J., Bouman, K.L., Freeman, W.T.: Visual dynamics: Probabilistic future frame synthesis via cross convolutional networks. arXiv preprint arXiv:1607.02586 (2016)
- Zhang, H., Dai, Y., Li, H., Koniusz, P.: Deep stacked hierarchical multi-patch network for image deblurring. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 5978–5986 (2019)
- Zhang, J., Xu, C., Liu, L., Wang, M., Wu, X., Liu, Y., Jiang, Y.: Dtvnet: Dynamic time-lapse video generation via single still image. In: European Conference on Computer Vision. pp. 300–315. Springer (2020)
- 40. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 586–595 (2018)
- Zhong, Z., Gao, Y., Zheng, Y., Zheng, B.: Efficient spatio-temporal recurrent neural network for video deblurring. In: European Conference on Computer Vision. pp. 191–207. Springer (2020)
- Zhou, S., Zhang, J., Pan, J., Xie, H., Zuo, W., Ren, J.: Spatio-temporal filter adaptive network for video deblurring. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2482–2491 (2019)
- 43. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: Proceedings of the IEEE international conference on computer vision. pp. 2223–2232 (2017)
- 44. Zhu, J.Y., Zhang, R., Pathak, D., Darrell, T., Efros, A.A., Wang, O., Shechtman, E.: Multimodal image-to-image translation by enforcing bi-cycle consistency. In: Advances in neural information processing systems. pp. 465–476 (2017)