

Learning Degradation Representations for Image Deblurring

Dasong Li¹, Yi Zhang¹, Ka Chun Cheung², Xiaogang Wang^{1,4},
Hongwei Qin^{3†}, and Hongsheng Li^{1,4,5†}

¹MMLab, CUHK ²NVIDIA AI Technology Center ³SenseTime Research

⁴Centre for Perceptual and Interactive Intelligence Limited ⁵Xidian University

{dasongli@link, hsli@ee}.cuhk.edu.hk, qinhongwei@sensetime.com

Abstract. In various learning-based image restoration tasks, such as image denoising and image super-resolution, the degradation representations were widely used to model the degradation process and handle complicated degradation patterns. However, they are less explored in learning-based image deblurring as blur kernel estimation cannot perform well in real-world challenging cases. We argue that it is particularly necessary for image deblurring to model degradation representations since blurry patterns typically show much larger variations than noisy patterns or high-frequency textures. In this paper, we propose a framework to learn spatially adaptive degradation representations of blurry images. A novel joint image reblurring and deblurring learning process is presented to improve the expressiveness of degradation representations. To make learned degradation representations effective in reblurring and deblurring, we propose a Multi-Scale Degradation Injection Network (MSDI-Net) to integrate them into the neural networks. With the integration, MSDI-Net can handle various and complicated blurry patterns adaptively. Experiments on the GoPro and RealBlur datasets demonstrate that our proposed deblurring framework with the learned degradation representations outperforms state-of-the-art methods with appealing improvements. The code is released at https://github.com/dasongli1/Learning_degradation.

Keywords: Image Deblurring, Degradation Representations

1 Introduction

Image restoration is required to handle various and complicated degradation patterns produced in different degradation processes. The degradation representations act as a crucial component to model the degradation processes and handle complicated degradation patterns, such as different noise levels in image denoising [52, 8, 21] and different combinations of Gaussian blurs and motion blurs in blind super-resolution [41]. However, the degradation representations are less exploited in learning-based deblurring methods and have not been well integrated into state-of-the-art deblurring networks.

[†]Corresponding authors

The general blurring process can be formulated as

$$y = F(x, k) + \eta, \quad (1)$$

where x and y are sharp image and blurry image respectively. $F(x, k)$ is usually modeled as a blurring operator with kernel k . η represents the Gaussian noise.

A popular paradigm for image deblurring is based on the Maximum A Posterior (MAP) estimate framework,

$$(k, x) = \arg \max \mathbb{P}(y|x, k)\mathbb{P}(x)\mathbb{P}(k), \quad (2)$$

where $\mathbb{P}(x)$ and $\mathbb{P}(k)$ model the priors of the clean images and the blur kernels. Many handcrafted priors for modeling $\mathbb{P}(x)$ and $\mathbb{P}(k)$ have been proposed [26, 18, 11, 1]. But most of them are insufficient in characterizing the clean images and blur kernels accurately. Furthermore, the operator F is generally modeled as a convolution operation in conventional MAP frameworks, which does not hold in practice and causes displeasing artifacts in real-world challenging cases.

Based on the limitations of kernel-based blurring modeling, a series of kernel-free approaches [6, 38, 3, 47, 4] are proposed to directly learn the mapping from blurry images to corresponding sharp images. While those methods outperform previous deblurring methods significantly, their performances are still limited in complicated blurry patterns, due to the lack of explicit modeling of the degradation process. This is because, unlike denoising methods, where the noise level might be similar across different images, blurring of different images generally have totally different patterns and cannot be well handled by fixed-weight networks without considering the degradation process. To combine the modeling of degradation and learning-based deblurring, recent works [39, 29] propose to learn explicit degradation representations by using Deep Image Prior (DIP) [40] to reparameterize the kernel k and the sharp image x . This inevitably involves the time-consuming iterative inverse optimization and hyperparameter tuning of DIP to adapt the deblurring process. Moreover, degradation representations have not been taken as a common component in SOTA deblurring methods [3, 47].

In this paper, we propose to learn explicit degradation representations with a novel joint sharp-to-blurry image reblurring and blurry-to-sharp image deblurring learning framework. Specifically, the degradation representations are learned in the process of sharp-to-blurry image reblurring. The process takes as input a blurry image and learns the degradation representations as a multi-channel spatial latent map to encode the spatially varying blur patterns in replacement of the conventional convolutional blur kernels or the DIP prior. A reblurring generator then takes as input the latent degradation map and the original sharp image and reblurs the sharp image back to its corresponding blurry image.

To effectively integrate the learned degradation representations into the reblurring process, we introduce a multi-scale degradation injection network (MSDI-Net) for achieving conditional image reblurring. The network adopts a U-Net like architecture [31]. The sharp image is fed into the encoder of the U-Net but the degradation is input into the encoder-decoder via the skip-connections to modulate the shortcut encoder feature maps. Specifically, the latent degradation

map is gradually upsampled via nearest-neighbor interpolation and a convolution layer to multiple resolutions and are then used to predict spatially varying weighting and bias parameters of the shortcut feature maps at each corresponding resolution. In this way, the learning of the latent degradation representation is supervised by the original blurry image for sharp-to-blurry image reblurring.

To make the learned degradation representations contributing to image deblurring, another blurry-to-sharp image generator network is also introduced, which shares the same MSDI-Net architecture but does not share weights with the reblurring generator. The learned degradation representations are processed similarly to deblur a blurry input image. Specifically, its encoder takes the blurry image as input, while the latent degradation representations are used to modulate the blurry image’s encoder-decoder shortcut feature maps at multiple resolutions. With the help of learned degradation map, the deblurring generator can handle complicated spatially varying blurry patterns, which are adaptively learned from the data to optimize both reblurring and deblurring tasks.

The main contributions of this work are two-fold: 1) We propose a novel joint framework for learning both sharp-to-blurry image reblurring and blurry-to-sharp image deblurring to adaptively encode spatially varying degradations and model the image blurring process, which in turn, benefits the image deblurring performance. 2) The proposed joint reblurring and deblurring framework outperforms state-of-the-art image deblurring methods on the widely used GoPro [24] and RealBlur [30] datasets.

2 Related Work

In this section, we briefly talk about the related works of image restoration with degradation representations and different image deblurring methods.

Image restoration with degradation representations. Image restoration tasks are usually required to handle different and complicated degradations in real-world applications. The degradation representations have been exploited and taken as one crucial component in several image restoration tasks, such as image denoising and image super-resolution. In image denoising, [52,21,8] take the noise variance as one network input to adaptively handle various noise strengths. Several practical denoising methods stabilize the noise variance caused by various ISO [43] and the property of Poisson-Gaussian distribution [35,15]. Similarly, image blind super-resolution is required to handle various degradations (different Gaussian blurs, motion blurs, and noises) in real-world applications. [41] proposes an unsupervised learning scheme for learning degradation representations based on the assumption that the degradation is the same in an image but can vary for different images. However, this assumption does not hold in image deblurring.

Optimization-based deblurring. A popular approach for image deblurring is based on the Maximum A Posterior (MAP). Most MAP-based methods focus on finding good priors for sharp images and blur kernels. Many priors are designed to model clean images and blur kernels. They include total variation

(TV) [1], hyper-Laplacian prior [11], l_0 -norm gradient prior [46] and sparse image priors [14]. They all assume the blur kernel is linear and uniform and can be represented as a convolution kernel. However, this assumption does not hold for real-world blurring with the non-uniform kernels. Some non-uniform deblurring methods [44,23,33,5] are proposed based on the assumption that the blur is locally uniform. They are not practical even with high computational costs.

Learning-based deblurring. Many deep deblurring models have been proposed over the past few years. Earlier attempts [32,37] utilize deep convolution neural networks to facilitate blur kernel estimation. However, there are several limitations [24] in estimating kernels. 1) Simple kernel convolution is not practical in real-world challenging cases. 2) The incorrect kernel estimations, caused by noise and large motions, may cause unpleasing artifacts. 3) Estimating spatially varying kernels requires a huge amount of computation. To avoid the above limitations, a series of kernel-free methods [38,24,12,13,49,6,36,27,53,47,4,3] are proposed with much better performances. Nah et al. [24] propose a multi-scale network for image deblurring. Similarly, Tao et al. [38] propose a scale-recurrent structure for image deblurring. Adversarial training is also introduced in image deblurring [12,13]. Chen *et al.* [2] introduces a reblur2deblur framework for video deblurring. Zhang *et al.* [53] proposes a reblurring network to synthesize additional blurry training images. The reblurring network and deblurring network are separate in both two methods [2,51]. In this work, we combine reblurring network and deblurring network in learning the degradation representations. Recently, multi-stage approaches [49,36,47,3] achieve impressive performance against previous methods. While those deblurring networks outperform the traditional deblurring methods significantly, their performances are still limited due to the lack of explicit degradation modeling.

Learning Blurring Degradation Representations. While the explicit degradation representations have shown convincing improvements on many low-level vision tasks, it is rarely explored in learning-based deblurring methods. SelfDeblur [29] introduces the Deep Image Prior (DIP) [40] to model the clean images and the kernels separately. But, it assumes the blur kernels are linear and uniform, which is not practical in complicated real-world scenes. Tran *et al.* [39] address the limitation by introducing the explicit representation for the blur kernels and the blur operators and reparameterizing the degradation and the sharp image by using DIP. This inevitably involves alternative optimization, which is time-consuming. The learned representations cannot improve the performance of existing deblurring networks directly and their application is limited to time-consuming DIP-based optimization methods.

3 Methodology

In this section, we first introduce a joint learning framework for both sharp-to-blurry image reblurring and blurry-to-sharp image deblurring to encode latent spatially varying degradation representations from blurry images. A blur-aware loss is introduced to enhance the image deblurring performance. To more effec-

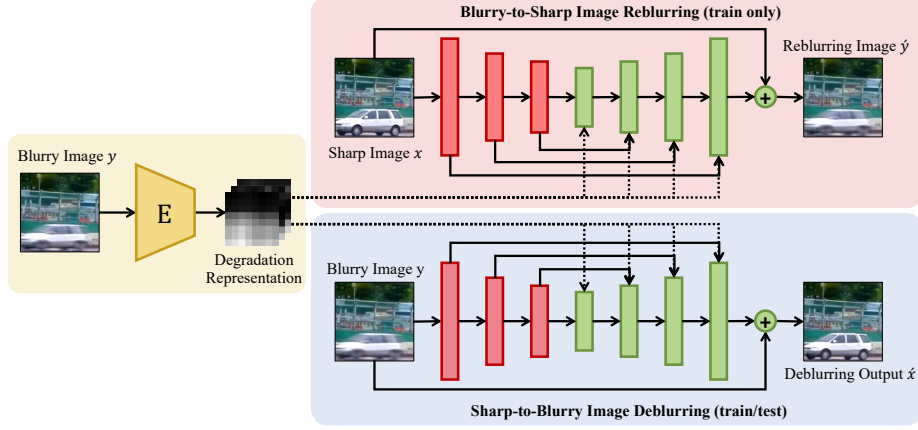


Fig. 1: Learning degradation representations with reblurring and deblurring.

tively integrate the latent degradation representations for reblurring and deblurring, a multi-scale degradation injection network is proposed for both tasks.

3.1 Learning Degradation Representations from Joint Reblurring and Deblurring

Most existing deblurring methods [32,26,18,11,1] take the blur kernels as the degradation representations and model the blurring process as a convolution on the input image. However, the simple kernel convolution is not practical in real-world challenging cases and it is usually difficult to estimate the blur kernels in large motions and spatially varying blurring cases. We propose to encode latent degradation representations from blurry images via the joint learning of image reblurring and deblurring. As shown in Fig. 1, we introduce an encoder E to encode the blurry image y into the degradation representations $E(y)$, which is modeled as a multi-channel latent map encoding 2D spatially varying blurring degradation in a latent space. Then an image reblurring generator G_r and an image deblurring generator G_d are introduced to generate the reblurred image \hat{y} and the deblurred image \hat{x} , respectively. The degradation representations do not only help the reblurring generator G_r model the degradation process, but also help the deblurring network handle complicated spatially varying degradation patterns. The joint training of reblurring and deblurring strengthens the expressiveness of learned degradation representations.

Sharp-to-blurry image reblurring. Different from modeling the blurring process as a convolution, our generator G_r models the degradation process via learning to generate the blurry image \hat{y} , given the sharp image x and the corresponding degradation representations $E(y)$. In addition, instead of generating the whole blurry image from scratch, the reblurring generator G_r learns to predict the residual between the sharp image x and the blurry image y . The learning

of the blurring degradation process in our framework is therefore formulated as generating a blurry image \hat{y} from its clean image x as

$$\hat{y} = x + G_r(x, E(y)), \quad (3)$$

where \hat{y} is the reblurred image conditioned on the sharp image x and degradation representation $E(y)$. With the residual learning, the encoder E is encouraged to neglect the contents of the blurry image and to focus on disentangling the content-independent degradation representation $E(y)$ from the blurry image y .

Note that most existing encoder-decoder networks, such as VAE [10], can also learn implicit image representations via image reconstruction. They aim at encoding the whole-image contents into the latent representations. However, our framework aims at encoding only the degradation information by predicting the residuals between the sharp and blurry images. Such a task actually has a lower difficulty level than reconstructing all contents of an input image and therefore leads to encoding better blurring degradation representations.

We first tried the mainstream L_1 distance as the loss function. But, as shown later in our supplementary, the L_1 loss function merely measures the pixel-wise distance and cannot properly describe the similarity of the blurry patterns of two images. Then the reblurring generator G_r cannot well generate the blurry images with L_1 loss, which harms the learning of degradation representation. Therefore, we further resort to perceptual loss [9] and adversarial training [7] to distinguish different degradation patterns in training. Adversarial loss is applied on the output of the generator G_r to distinguish real and fake blurry images so that the decoder can well model the blurring process and improve the expressiveness of degradation representation. We take the hinge loss [17, 48, 22, 28] as the adversarial loss to help the reblurring generator G_r model the degradation (blurring) process and improve the expressiveness of learned degradation representations. We train the reblurring generator G_r to generate the blurry images with the multi-scale discriminator D used in [42]. The training objective for image reblurring is formulated as

$$\begin{aligned} L_G &= -E_{x \sim p_{\text{data}}} D(x, \hat{y}) + \lambda_1 L_{\text{perceptual}}(y, \hat{y}), \\ L_D &= -E_{(x, y) \sim p_{\text{data}}} [\min(0, -1 + D(x, y))] - E_{x \sim p_{\text{data}}} [\min(0, -1 - D(x, \hat{y}))], \end{aligned} \quad (4)$$

where λ_1 balances the L_1 loss and the discriminator loss for image reblurring, and the discriminator D is a conditional discriminator conditioning on the sharp image x . Conditioning on the sharp image x , the conditional discriminator D can focus on whether the reblurred image \hat{y} has the same image contents with the corresponding sharp image x . The reblurring helps extract content-independent degradation information (shown in Fig), which is different with the content-dependent conditional networks [45, 50, 54].

Blurry-to-sharp image deblurring. To make the learned degradation representations contributing to image deblurring, we also model the image deblurring process with the image reblurring process jointly. The image deblurring generator G_d follows a similar design to that of G_r . The deblurring generator G_d is

only required to learn the deblurring residuals between the blurry image y and its corresponding sharp image x , and the learned degradation representations $E(y)$. The blurry-to-sharp image deblurring is modeled as

$$\hat{x} = y + G_d(y, E(y)). \quad (5)$$

Thanks to the 2D learnable degradation representations, the image deblurring generator G_d is aware of the spatially varying blurry patterns and thus can adaptively handle various and complicated degradation patterns. The learning of image deblurring makes the learned degradation representations adapted to the deblurring task. The loss function for image deblurring is formulated as

$$L(x, \hat{x}) = \lambda_2 L_1(x, \hat{x}). \quad (6)$$

Discussion of the learned degradation representation. The learned degradation representation has two main advantages against conventional kernel modeling: 1) Our degradation representations can learn non-uniform spatially varying degradations effectively. Figs. 1 and 6 show that the encoder can distinguish different degradation representations. 2) Interpolating on the latent space of representations can generate blurry images with controllable blurry levels (as shown in Fig. 5). The representations are also content-independent (as shown in Fig. 6), which is different from previous conditional networks [45, 50, 54]. Built on this latent space, the representations show better interpretability and expressiveness.

3.2 Image Deblurring with Learned Degradation Representations

After obtaining the pre-trained encoder E , we freeze the pre-trained encoder and re-train the deblurring generator G_d to illustrate that our improvement is not from the complicated framework but the learned degradation representations. To further demonstrate the generality of learned degradation representations, we also train the deblurring generator G_d on the RealBlur dataset [30] with the encoder trained on the GoPro dataset [24].

Following HINet [3], we select the Peak Signal-to-Noise Ratio (PSNR) loss as the main supervision. We also utilize a blur-aware loss function as extra supervision. The well-trained encoder E should be quite sensitive to capture various and even subtle blur patterns. We, therefore, define the blur-aware loss as the distance between degradation encoder features of the ground-truth sharp image x and the estimated sharp image \hat{x} . The deviation between the encoder feature maps $\|E(x) - E(\hat{x})\|_1$ of images x and \hat{x} can give more weights on the remaining blurry regions of the network output \hat{x} . Similar to perceptual loss [9], the L_1 distances between the encoder feature maps can be calculated at multiple scales. The blur-aware loss L_{blur} is therefore formulated as

$$L_{\text{blur}}(\hat{x}, x) = \sum_{i=1}^N \frac{1}{|E^{(i)}|} [\|E^{(i)}(\hat{x}) - E^{(i)}(x)\|_1], \quad (7)$$

where $E^{(i)}$ denotes the i -th layer of the encoder E and $|E^{(i)}|$ denotes the number of pixels in feature map $E^{(i)}(x)$. Using the blur-aware loss makes the deblurring

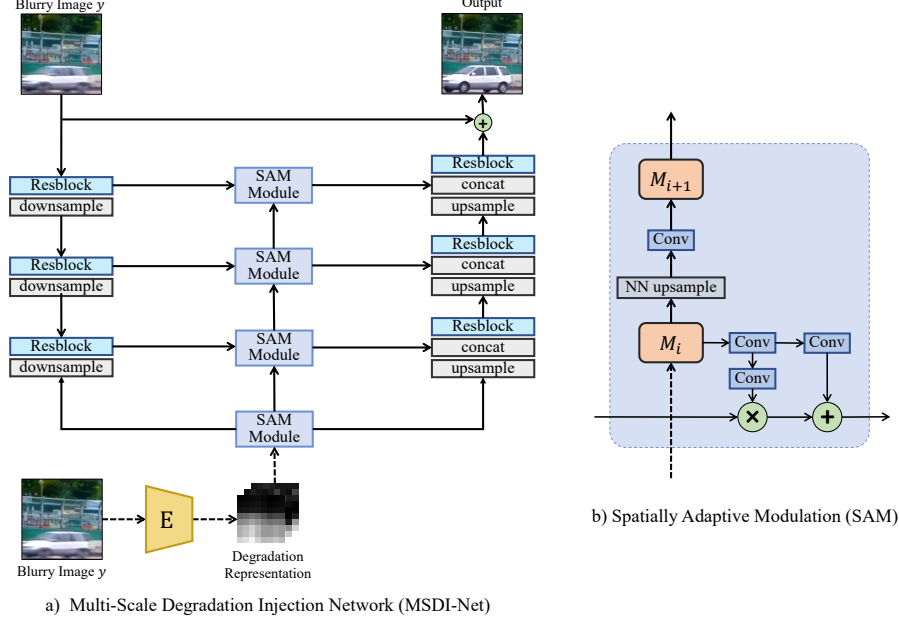


Fig. 2: The network structure of multi-scale degradation injection network for image deblurring. Image reblurring shares the same structure.

generator G_d pay more attention to the remaining blurry regions of the network output \hat{x} . Then the deblurring objective is formulated as

$$L(x, \hat{x}) = \text{PSNR}(x, \hat{x}) + \lambda_2 L_{\text{blur}}(x, \hat{x}). \quad (8)$$

3.3 Multi-scale Degradation Injection Network

To effectively integrate the degradation $E(y)$ to predict the reblurring and deblurring residuals, we propose a multi-scale degradation injection network (MSDI-Net) for both reblurring and deblurring. Our reblurring generator G_r and deblurring generator G_d consists of two MSDI-Nets, which are stacked as HINet [3] does. for simplicity, the overview of a MSDI-Net is shown in Fig. 2. The details of the architectures are shown in the supplementary materials.

The MSDI-Net consists of an encoder, a decoder, concatenation-based skip-connections, and a multi-scale degradation injection module. The first three modules of the MSDI-Net are widely adopted in U-Net like architectures [31]. The multi-scale degradation injection module modulates the feature of each skip-connection spatially, based on the learned degradation representation. The spatially variant modulation are explored in image synthesis [28] and image super-resolution [16]. We adopt it as the key to connecting learnable degradation representations to reblurring and deblurring.

Let $f_i \in \mathbb{R}^{C_i \times H_i \times W_i}$ denote the features extracted at scale $i = 1, \dots, 5$ in the encoder. The extracted feature f_i is passed to the decoder at the skip connection of scale i . Our MSDI-Net integrates the degradation representations into concatenation-based skip-connections of scale i . At scale i , we obtain the degradation map $M_i \in \mathbb{R}^{C_i \times H_i \times W_i}$ by a convolution-based upsampling block on the original degradation map $M_{i-1} \in \mathbb{R}^{C_{i-1} \times H_{i-1} \times W_{i-1}}$ ($H_i = 2 \times H_{i-1}, W_i = 2 \times W_{i-1}$), which is implemented by a nearest-neighbor interpolation and a convolution layer to avoid checkerboard artifacts [25]. Then we utilize a spatially adaptive modulation (SAM) module to modulate the skip-connection feature f_i at scale i . The spatially adaptive modulation modulates the feature map channels in a spatially varying manner with both predicted scaling and additions. At the skip connection of scale i , we use several convolution (3×3) layers on the degradation map M_i to predict the modulating parameters $\gamma_i \in \mathbb{R}^{C_i \times H_i \times W_i}$ and $\beta_i \in \mathbb{R}^{C_i \times H_i \times W_i}$ respectively, which modulate the feature map f_i as

$$F_i = \gamma_i \odot f_i + \beta_i, \quad (9)$$

where F_i is the modulated skip-connection features. Since the modulation parameters are predicted from the degradation representations, the learnable modulation of the feature channels makes the deblurring network aware of spatially varying degradations. The degradation-aware feature F_i is then concatenated with the decoder feature at scale i and the last decoder layer predicts the image residuals for both image reblurring and deblurring.

Injecting the degradation representations enables the networks to handle various and complicated degradation patterns adaptively. Injection at multiple scales improves the expressiveness of degradation representations by strengthening the connections between the degradation representations and two generators.

4 Experiments

4.1 Dataset and implementation details

We train and evaluate our method on the GoPro [24] and RealBlur datasets [30]. The GoPro dataset consists of 2,103 pairs of blurry and sharp images for training and 1,111 pairs for testing. The RealBlur dataset consists of 3,758 pairs for training and 980 pairs for testing. We first train the framework of reblurring and deblurring on GoPro dataset [24]. We apply horizontal flipping and rotation as data augmentation and crop image patch of size 256×256 from the dataset for training. λ_1 is set as 30 and λ_2 is set as 10. The networks of whole framework are trained with a batch size of 32 for 200k iterations. Then we freeze the weights of well-trained encoder E and train the deblurring generator G_d on the GoPro dataset [24] and the RealBlur dataset [30] respectively. λ_3 is set as 1. The deblurring generator G_d is trained with a batch size of 64 for 400k iterations. We use the Adam optimizer and the learning rate is set as 3×10^{-4} at the beginning and decreased to 1×10^{-7} following the cosine annealing strategy [19].

Method	Reference	PSNR	SSIM
DeblurGAN [12]	CVPR'18	28.70	0.858
DeblurGAN-v2 [13]	ICCV'19	29.55	0.934
SRN [38]	CVPR'18	30.26	0.934
Gao et al. [6]	CVPR'19	30.90	0.935
DBGAN [53]	CVPR'20	31.10	0.942
MT-RNN [27]	ECCV'20	31.15	0.945
DMPHN [49]	CVPR'19	31.20	0.940
Suin et al. [36]	CVPR'20	31.85	0.948
MIMO-UNet [4]	ICCV'21	32.45	0.957
MPRNet [47]	CVPR'21	32.66	0.959
HINet [3]	CVPRW'21	32.71	0.959
MPRNet-patch256 [47]	CVPR'21	<u>32.96</u>	<u>0.961</u>
Ours	ECCV'22	33.28	0.964

Table 1: Deblurring comparisons on GoPro [24] dataset. Best and second best scores are **highlighted** and underlined.

Model	Blurriest 10%	Sharpest 10%	All	MACs (G)
MPRNet-patch256 [47]	29.31	35.52	32.96	760.11
Ours	29.65	35.58	33.28	336.43

Table 2: Detailed comparisons of MPRNet [47] and our method on GoPro test dataset [24]. MACs are estimated with the input size of $3 \times 256 \times 256$.

4.2 Performance comparison

We compare our method with state-of-the-art deblurring methods [3, 47, 4] on the GoPro test dataset [24]. The quantitative results are reported in Table 1. For testing, we slice the whole image into several 256×256 patches and test all patches to report the results of HINet [3], MPRNet-patch256 [47] and our method. Our method achieves 0.45 dB improvement in terms of PSNR over the previous best-performing method HINet [3]. To evaluate effectiveness and generality of learned degradation representations, we also evaluate our method on the RealBlur dataset. As listed in Table 3, our method achieves the best performance in terms of PSNR and SSIM. Since HINet [3] does not release the model for RealBlur dataset [30], we train the HINet model on RealBlur dataset based on their released training code. Note that we apply the degradation representations trained on GoPro dataset [24] directly on RealBlur dataset [30]. Our method outperforms the previous SOTA HINet [3] by 0.23 dB PSNR, which demonstrate the generality of learnable degradation representations.

In Table 2, we provide detailed comparisons between our method and MPRNet [47]. We divide the whole gopro dataset into blurriest 10% and sharpest 10% as [34] does. It is observed that the main improvement of our method is the improvement of 0.34dB PSNR in the blurriest 10%, which demonstrates the advantage of proposed degradation learning. What's more, our method's computational cost is less than 50% of MPRNet's [47].




Fig. 3: Visual comparisons for image deblurring on the GoPro test dataset [24]. From left-top to right-bottom: blurry images, ground-truth images, and results obtained by MIMO-UNet [4], HINet [3], MPRNet [47] and our proposed method.

Fig. 3 and Fig. 4 show example deblurred results from the GoPro [24] and RealBlur [30] test sets by the evaluated approaches. Our method produces sharper images and recovers more details in the regions of texts and moving objects, compared with other methods.

4.3 Interpolation and decoupleness of degradation representations

To demonstrate the effectiveness of learned degradation representations, we study the interpolation and decoupleness of learned degradation representations. Given a blurry image y and its corresponding sharp image x , we can obtain two degradation representations $E(y)$ and $E(x)$. Then we obtain several intermediate degradation representations by interpolating from $E(y)$ to $E(x)$. The corresponding output of the decoder changes smoothly from sharp to blurry images. The blur interpolation in Fig. 5 shows that our degradation representations are built on the latent space and accurately aware of different degradations. We empirically validate the decoupleness of blur and image contents on GoPro dataset. We divide the GoPro test dataset into 555 pairs of images. For each pair of sharp images $\{A, B\}$, the corresponding blurry images are $\{\text{blur}(A), \text{blur}(B)\}$ and the




Fig. 4: Visual comparisons for image deblurring on the RealBlur test dataset [30]. From left-top to right-bottom: blurry images, ground-truth images and results obtained by DeblurGANv2, HINet [3], MIMO-UNet [4], our proposed method.

Method	Reference	PSNR	SSIM
DeblurGAN-v2 [13]	ICCV'19	29.69	0.870
SRN [38]	CVPR'18	31.38	0.909
MPRNet [47]	CVPR'21	31.76	<u>0.922</u>
MIMO-UNet [4]	ICCV'21	32.05	0.921
HINet [3]	CVPRW'21	<u>32.12</u>	0.921
Ours	ECCV'22	32.35	0.923

Table 3: Deblurring comparisons on the RealBlur test dataset [30].

degradation representations are obtained as $\{\text{deg}_A, \text{deg}_B\}$. The sharp image A can be reblurred according to deg_B to obtain $\text{ReBlur}(A, \text{deg}_B)$. We use the average contextual similarity CX [20] to measure pairwise image similarity. We have $\text{CX}(\text{blur}(A), A) = 2.72$, $\text{CX}(\text{ReBlur}(A, \text{deg}_B), A) = 2.65$, $\text{CX}(\text{blur}(A), B) = 5.43$, and $\text{CX}(\text{ReBlur}(A, \text{deg}_B), B) = 5.39$, averaging over all the pairs. The similarities empirically show that $\text{ReBlur}(A, \text{deg}_B)$ has similar contents with A (the former two equations) but doesn't have similar contents with B (the latter two equations). The visual examples on degradation decoupleness are provided in Figure 6. The decoupleness of learned degradation representations shows our difference with the content-dependent conditional networks [45,50,54] and the capacity of being a general operator to replace conventional blurring process.



Fig. 5: Generating blurry images with linearly interpolated degradation representations. From left to right: the blurry level from sharp to blur.




Fig. 6: Reblurring A_1, A_2, A_3 with the degradation representation of $\text{Blur}(B)$.

4.4 Ablation study

We evaluate the effectiveness of learning degradation representations and the multi-scale degradation injection network by revising one of the components of our model at a time. Table 4 lists the performances of different settings on the GoPro test set [24]. We first remove the degradation encoder for learning the degradation representations and make G_d only takes as input the blurry image (denoted as “Ours w/o degradation”). The performance suffers a large drop of 0.47dB PSNR. Then we remove the generator of image reblurring and reserve the encoder to provide additional encoding from the blurry images (denoted as “Our w/o reblurring”), this operation causes a drop of 0.19 dB PSNR, demonstrating that reblurring indeed contributes to the learning of better degradation representations. Then we remove the blur-aware loss function (Eq. (7)). The training objective for deblurring becomes only the PSNR loss (denoted as “Ours w/o blur loss”). The performance drops slightly by about 0.07 PSNR. We further experiment with different ways of integrating degradation representations into the generators. When we remove the injection at multiple scales, the degradation representation is integrated just at the lowest resolution skip connection (de-

Model	PSNR
Ours w/o degradation	32.81
Ours w/o reblurring	33.09
Ours w/o blur loss	33.21
Our injection w/o multi-scale	32.80
Our input w/ concat	32.98
Ours w/ concat injection	33.05
Ours	33.28

Table 4: The ablation study of image deblurring on GoPro test dataset[24].

noted as “Our injection w/o multi-scale”). The performance suffers a significant drop of 0.48 PSNR, which means that injecting the degradation representation into a single scale would affect the deblurring performance significantly. Then we replace the modulation and test on concatenating latent degradation map with skip-connection feature maps (denoted as “Our w/ concat injection”). For fair comparisons of computational cost, we add two-layer residual blocks after the feature concatenation. The operation causes a drop of 0.23 PSNR. We also remove the integration totally and concatenate the upsampling degradation maps with the blurry images at the network entrance with the blurry image (denoted as “Ours input w/ concat”), which is the mainstream design in image denoising [52, 21, 8]. Its performance drops by about 0.3 PSNR. All the ablation studies demonstrate the effectiveness of our proposed learnable degradation representations and MSDI-Net for image deblurring.

5 Conclusions

In this paper, we propose a framework for learning degradation representation with image reblurring and image deblurring. We first utilize an encoder to learn the degradation representations explicitly. Then we propose a multi-scale degradation injection network to effectively integrate the degradation representations for reblurring and deblurring. With the degradation representations, our networks can be aware of and handle spatially varying degradation patterns adaptively. The experimental results demonstrate that our method outperforms other methods with a clear margin.

Acknowledgments. This work is supported in part by Centre for Perceptual and Interactive Intelligence Limited, in part by the General Research Fund through the Research Grants Council of Hong Kong under Grants (Nos. 14204021, 14207319, 14203118, 14208619), in part by Research Impact Fund Grant No. R5001-18, in part by CUHK Strategic Fund.

References

1. Chan, T., Wong, C.K.: Total variation blind deconvolution. *IEEE Transactions on Image Processing* **7**(3), 370–375 (1998) [2](#), [4](#), [5](#)
2. Chen, H., Gu, J., Gallo, O., Liu, M.Y., Veeraraghavan, A., Kautz, J.: Re-blur2deblur: Deblurring videos via self-supervised learning. In: 2018 IEEE International Conference on Computational Photography (ICCP). pp. 1–9 (2018) [4](#)
3. Chen, L., Lu, X., Zhang, J., Chu, X., Chen, C.: Hinet: Half instance normalization network for image restoration. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 182–192 (June 2021) [2](#), [4](#), [7](#), [8](#), [10](#), [11](#), [12](#)
4. Cho, S.J., Ji, S.W., Hong, J.P., Jung, S.W., Ko, S.J.: Rethinking coarse-to-fine approach in single image deblurring. In: Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV). pp. 4641–4650 (October 2021) [2](#), [4](#), [10](#), [11](#), [12](#)
5. Cho, S., Matsushita, Y., Lee, S.: Removing non-uniform motion blur from images. In: 2007 IEEE 11th International Conference on Computer Vision. pp. 1–8 (2007) [4](#)
6. Gao, H., Tao, X., Shen, X., Jia, J.: Dynamic scene deblurring with parameter selective sharing and nested skip connections. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 3848–3856 (2019) [2](#), [4](#), [10](#)
7. Goodfellow, I., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., Courville, A., Bengio, Y.: Generative adversarial nets. In: Ghahramani, Z., Welling, M., Cortes, C., Lawrence, N., Weinberger, K.Q. (eds.) *Advances in Neural Information Processing Systems*. vol. 27. Curran Associates, Inc. (2014) [6](#)
8. Guo, S., Yan, Z., Zhang, K., Zuo, W., Zhang, L.: Toward convolutional blind denoising of real photographs. 2019 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2019) [1](#), [3](#), [14](#)
9. Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: European Conference on Computer Vision (2016) [6](#), [7](#)
10. Kingma, D.P., Welling, M.: Auto-Encoding Variational Bayes. In: 2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14–16, 2014, Conference Track Proceedings (2014) [6](#)
11. Krishnan, D., Fergus, R.: Fast image deconvolution using hyper-laplacian priors. In: Bengio, Y., Schuurmans, D., Lafferty, J., Williams, C., Culotta, A. (eds.) *Advances in Neural Information Processing Systems*. vol. 22. Curran Associates, Inc. (2009) [2](#), [4](#), [5](#)
12. Kupyn, O., Budzan, V., Mykhailych, M., Mishkin, D., Matas, J.: Deblurgan: Blind motion deblurring using conditional adversarial networks. In: CVPR. pp. 8183–8192. Computer Vision Foundation / IEEE Computer Society (2018) [4](#), [10](#)
13. Kupyn, O., Martyniuk, T., Wu, J., Wang, Z.: Deblurgan-v2: Deblurring (orders-of-magnitude) faster and better. In: ICCV. pp. 8877–8886. IEEE (2019) [4](#), [10](#), [12](#)
14. 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 (2009) [4](#)
15. Li, D., Zhang, Y., Law, K.L., Wang, X., Qin, H., Li, H.: Efficient burst raw denoising with variance stabilization and multi-frequency denoising network. *International Journal of Computer Vision* **130**(8), 2060–2080 (Aug 2022) [3](#)

16. Liang, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: Mutual affine network for spatially variant kernel estimation in blind image super-resolution. In: IEEE International Conference on Computer Vision (2021) 8
17. Lim, J.H., Ye, J.C.: Geometric gan (2017), <https://arxiv.org/abs/1705.02894> 6
18. Liu, G., Chang, S., Ma, Y.: Blind image deblurring using spectral properties of convolution operators. IEEE Transactions on Image Processing **23**(12), 5047–5056 (2014) 2, 5
19. Loshchilov, I., Hutter, F.: SGDR: stochastic gradient descent with warm restarts. In: 5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24–26, 2017, Conference Track Proceedings (2017) 9
20. Mechrez, R., Talmi, I., Zelnik-Manor, L.: The contextual loss for image transformation with non-aligned data. arXiv preprint arXiv:1803.02077 (2018) 12
21. Mildenhall, B., Barron, J.T., Chen, J., Sharlet, D., Ng, R., Carroll, R.: Burst denoising with kernel prediction networks. In: CVPR (2018) 1, 3, 14
22. Miyato, T., Kataoka, T., Koyama, M., Yoshida, Y.: Spectral normalization for generative adversarial networks. In: International Conference on Learning Representations (2018) 6
23. Nagy, J.G., O’Leary, D.P.: Restoring images degraded by spatially variant blur. SIAM Journal on Scientific Computing **19**(4), 1063–1082 (1998) 4
24. Nah, S., Kim, T.H., Lee, K.M.: Deep multi-scale convolutional neural network for dynamic scene deblurring. In: The IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (July 2017) 3, 4, 7, 9, 10, 11, 13, 14
25. Odena, A., Dumoulin, V., Olah, C.: Deconvolution and checkerboard artifacts. Distill (2016) 9
26. Pan, J., Sun, D., Pfister, H., Yang, M.H.: Blind image deblurring using dark channel prior. In: 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 1628–1636 (2016) 2, 5
27. Park, D., Kang, D.U., Kim, J., Chun, S.Y.: Multi-temporal recurrent neural networks for progressive non-uniform single image deblurring with incremental temporal training. In: ECCV (6). Lecture Notes in Computer Science, vol. 12351, pp. 327–343. Springer (2020) 4, 10
28. Park, T., Liu, M.Y., Wang, T.C., Zhu, J.Y.: Semantic image synthesis with spatially-adaptive normalization. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2019) 6, 8
29. Ren, D., Zhang, K., Wang, Q., Hu, Q., Zuo, W.: Neural blind deconvolution using deep priors. In: 2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR). pp. 3338–3347. IEEE Computer Society, Los Alamitos, CA, USA (jun 2020) 2, 4
30. Rim, J., Lee, H., Won, J., Cho, S.: Real-world blur dataset for learning and benchmarking deblurring algorithms. In: ECCV (25). Lecture Notes in Computer Science, vol. 12370, pp. 184–201. Springer (2020) 3, 7, 9, 10, 11, 12
31. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015 (2015) 2, 8
32. Schuler, C.J., Hirsch, M., Harmeling, S., Schölkopf, B.: Learning to deblur. IEEE Transactions on Pattern Analysis and Machine Intelligence **38**, 1439–1451 (2016) 4, 5

33. Shan, Q., Xiong, W., Jia, J.: Rotational motion deblurring of a rigid object from a single image. In: 2007 IEEE 11th International Conference on Computer Vision. pp. 1–8 (2007) [4](#)
34. 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 (TOG)* **40**(5) (2021) [10](#)
35. Starck, J.L., Murtagh, F., Bijaoui, A.: Image processing and data analysis. In: Cambridge University Press, Cambridge (1998) [3](#)
36. Suin, M., Purohit, K., Rajagopalan, A.N.: Spatially-attentive patch-hierarchical network for adaptive motion deblurring. In: CVPR. pp. 3603–3612. Computer Vision Foundation / IEEE (2020) [4](#), [10](#)
37. Sun, J., Cao, W., Xu, Z., Ponce, J.: Learning a convolutional neural network for non-uniform motion blur removal. In: 2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR). pp. 769–777 (2015) [4](#)
38. Tao, X., Gao, H., Shen, X., Wang, J., Jia, J.: Scale-recurrent network for deep image deblurring. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [2](#), [4](#), [10](#), [12](#)
39. Tran, P., Tran, A., Phung, Q., Hoai, M.: Explore image deblurring via encoded blur kernel space. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2021) [2](#), [4](#)
40. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep image prior. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (June 2018) [2](#), [4](#)
41. Wang, L., Wang, Y., Dong, X., Xu, Q., Yang, J., An, W., Guo, Y.: Unsupervised degradation representation learning for blind super-resolution. In: CVPR (2021) [1](#), [3](#)
42. Wang, T.C., Liu, M.Y., Zhu, J.Y., Tao, A., Kautz, J., Catanzaro, B.: High-resolution image synthesis and semantic manipulation with conditional gans. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2018) [6](#)
43. Wang, Y., Huang, H., Xu, Q., Liu, J., Liu, Y., Wang, J.: Practical deep raw image denoising on mobile devices. In: European Conference on Computer Vision (ECCV). pp. 1–16 (2020) [3](#)
44. Whyte, O., Sivic, J., Zisserman, A., Ponce, J.: Non-uniform deblurring for shaken images. In: 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition. pp. 491–498 (2010) [4](#)
45. Xintao Wang, Ke Yu, C.D., Loy, C.C.: Recovering realistic texture in image super-resolution by deep spatial feature transform. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2018) [6](#), [7](#), [12](#)
46. Xu, L., Zheng, S., Jia, J.: Unnatural l0 sparse representation for natural image deblurring. In: 2013 IEEE Conference on Computer Vision and Pattern Recognition. pp. 1107–1114 (2013) [4](#)
47. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H., Shao, L.: Multi-stage progressive image restoration. In: CVPR (2021) [2](#), [4](#), [10](#), [11](#), [12](#)
48. Zhang, H., Goodfellow, I., Metaxas, D., Odena, A.: Self-attention generative adversarial networks. In: Chaudhuri, K., Salakhutdinov, R. (eds.) Proceedings of the 36th International Conference on Machine Learning. Proceedings of Machine Learning Research, vol. 97, pp. 7354–7363. PMLR (09–15 Jun 2019) [6](#)
49. Zhang, H., Dai, Y., Li, H., Koniusz, P.: Deep stacked hierarchical multi-patch network for image deblurring. In: CVPR. pp. 5978–5986. Computer Vision Foundation / IEEE (2019) [4](#), [10](#)

- 50. Zhang, J., Pan, J., Ren, J., Song, Y., Bao, L., Lau, R.W., Yang, M.H.: Dynamic scene deblurring using spatially variant recurrent neural networks. In: 2018 IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 2521–2529 (2018) [6](#), [7](#), [12](#)
- 51. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a Gaussian denoiser: Residual learning of deep CNN for image denoising. *IEEE Transactions on Image Processing* **26**(7), 3142–3155 (2017) [4](#)
- 52. Zhang, K., Zuo, W., Zhang, L.: Ffdnet: Toward a fast and flexible solution for CNN based image denoising. *IEEE Transactions on Image Processing* (2018) [1](#), [3](#), [14](#)
- 53. Zhang, K., Luo, W., Zhong, Y., Ma, L., Stenger, B., Liu, W., Li, H.: Deblurring by realistic blurring. In: CVPR. pp. 2734–2743. Computer Vision Foundation / IEEE (2020) [4](#), [10](#)
- 54. 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 International Conference on Computer Vision (2019) [6](#), [7](#), [12](#)