

# Region-Adaptive Dense Network for Efficient Motion Deblurring

**Kuldeep Purohit, A. N. Rajagopalan**

Indian Institute of Technology Madras, India  
kuldeppurohit3@gmail.com, raju@ee.iitm.ac.in

## Abstract

In this paper, we address the problem of dynamic scene deblurring in the presence of motion blur. Restoration of images affected by severe blur necessitates a network design with a large receptive field, which existing networks attempt to achieve through simple increment in the number of generic convolution layers, kernel-size, or the scales at which the image is processed. However, these techniques ignore the non-uniform nature of blur, and they come at the expense of an increase in model size and inference time. We present a new architecture composed of region adaptive dense deformable modules that implicitly discover the spatially varying shifts responsible for non-uniform blur in the input image and learn to modulate the filters. This capability is complemented by a self-attentive module which captures non-local spatial relationships among the intermediate features and enhances the spatially varying processing capability. We incorporate these modules into a densely connected encoder-decoder design which utilizes pre-trained Densenet filters to further improve the performance. Our network facilitates interpretable modeling of the spatially-varying deblurring process while dispensing with multi-scale processing and large filters entirely. Extensive comparisons with prior art on benchmark dynamic scene deblurring datasets clearly demonstrate the superiority of the proposed networks via significant improvements in accuracy and speed, enabling almost real-time deblurring.

## Introduction

Though computational imaging has made tremendous progress over the years, handling motion blur in captured content remains a challenge. Motion blur is caused by motion of objects in the scene or the camera during sensor exposure. Apart from significantly degrading the visual quality, the distortions arising from blur lead to considerable performance drop for many vision tasks (Vasiljevic, Chakrabarti, and Shakhnarovich 2016). There exist a few commercially available cameras which can capture frames at a high frame-rate and thus experience less blur but they suffer from noise at high resolution and are quite expensive.

Motion deblurring is a challenging problem in computer vision due to its ill-posed nature. The past decade has wit-

Copyright © 2020, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

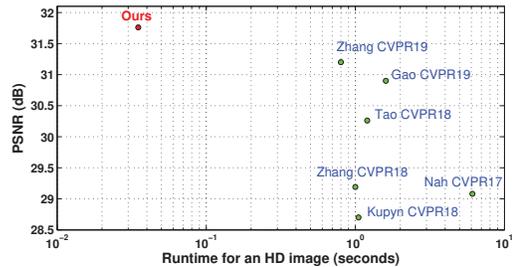


Figure 1: Our network outperforms all existing approaches on the dynamic scene deblurring benchmark in terms of accuracy as well as inference time.

nessed significant advances in deblurring, wherein major efforts have gone into designing priors that are apt for recovering the underlying undistorted image and the camera trajectory (Vasu and Rajagopalan 2017; Yan et al. 2017). An exhaustive survey of uniform blind deblurring algorithms can be found in (Lai et al. 2016). Few approaches (Chakrabarti 2016; Schuler et al. 2016) have proposed hybrid algorithms where a Convolutional Neural Network (CNN) estimates the blur kernel, which is then used in an alternative optimization framework for recovering the latent image.

However, these methods have been developed based on a rather strong constraint that the scene is planar and that the blur is governed by only *camera* motion. This precludes commonly occurring blur in most practical settings. Real-world blur arises from various sources including moving objects, camera shake and depth variations, causing different pixels to acquire different motion trajectories. A class of algorithms involve segmentation methods to relax the static and fronto-parallel scene assumption by independently restoring different blurred regions in the scene (Hyun Kim, Ahn, and Mu Lee 2013). However, these methods depend heavily on an accurate segmentation-map. Few methods (Sun et al. 2015; Gong et al. 2017) circumvent the segmentation stage by training CNNs to estimate locally linear blur kernels and feeding them to a non-uniform deblurring algorithm based on patch-level prior. However, they are limited in their capability when it comes to general dynamic scenes.

The afore-mentioned methods are not end-to-end systems and share a severe disadvantage of involving iterative, time-intensive, and cumbersome optimization schemes at the network output for getting the final deblurred result. Use of fully convolutional CNNs to directly estimate the latent sharp image was proposed in (Nah, Kim, and Lee 2017) and adopted by recent works to further advance the state-of-the-art. They offer the advantage of enabling generalized dynamic scene deblurring at low latency, by circumventing the iterative optimization stage involving fitting of hand-designed motion models. (Nah, Kim, and Lee 2017; Purohit, Shah, and Rajagopalan 2019; Tao et al. 2018; Gao et al. 2019) proposed encoder-decoder residual networks to aggregate features in a coarse-to-fine manner, while showing benefits of selective parameter/feature sharing and/or recurrent layers. Recently, (Zhang et al. 2019) proposed a multi-patch hierarchical network and stacked its copies along depth to achieve state-of-the-art performance.

However, there are two major limitations shared by prior deblurring works. Firstly, the filters of a generic CNN are spatially invariant (with spatially-uniform receptive field), which is a suboptimal modeling of the dynamic scene deblurring process and offers limited accuracy. Secondly, existing methods attempt to increase receptive field by increasing the model’s computational footprint, making them unsuitable for real-time applications. As the only other work of this kind, (Zhang et al. 2018b) recently proposed a design composed of multiple CNNs and Recurrent Neural Networks (RNN) to learn spatially varying weights for deblurring. However, their performance is inferior to the state-of-the-art (Zhang et al. 2019) in several aspects. Reaching a trade-off among inference time, accuracy of restoration, and receptive field is a non-trivial task which we address in this paper (see Fig. 1). We investigate position and motion-aware CNN architecture, which can efficiently handle multiple image regions experiencing motion with different magnitude and direction.

Following recent developments, we adopt an end-to-end learning based approach to directly estimate the restored sharp image. For single image deblurring, we build a fully convolutional architecture equipped with filter-transformation and feature modulation capability suited for the task of motion deblurring. Our design leverages the fact that motion blur is essentially an aggregation of various spatially varying transformations of the image, and a network that implicitly adapts to the location and direction of such motion, is a better candidate for the restoration task. Its advantages over prior art are three-fold: 1. It is a dynamic and computationally efficient as it requires only a single forward pass through each layer, obviates the need for repeated processing of the image (at different scales/patch-levels). 2. Its components can be easily introduced into other architectures to improve their performance. 3. The transformations estimated by the network are dynamic and hence can be meaningfully interpreted for any test image.

The efficacy of our architecture is demonstrated through comprehensive comparisons with the state-of-the-art on image deblurring benchmark. The major contributions of our work are:

- We propose an efficient motion deblurring architecture built using dense deformable modules that facilitate position-specific dynamic filter transformation.
- Our network benefits from estimating image dependent spatial attention-maps to process local features jointly with their globally distributed interdependencies.
- We embed the above adaptive modules into an densely connected fully convolutional design which benefits from pre-trained filters of DenseNet.
- Extensive experiments are presented on dynamic scene deblurring benchmark to show state-of-the-art accuracy and near real-time deblurring achieved by our architecture and the interpretability of its dynamic modules.

## Proposed Architecture

An existing technique for accelerating various image processing operations is to down-sample the input image, execute the operation at low resolution, and up-sample the output (Chen et al. 2016). However, this approach discounts the importance of resolution, rendering it unsuitable for image restoration tasks where high-frequency content of the image is of prime importance (deblurring, super-resolution).

Another efficient alternative is a CNN with a fixed but very large receptive field, e.g. Cascaded dilated network (Chen, Xu, and Koltun 2017), which was proposed to accelerate various image-to-image tasks. However, simple dilated convolutions are not appropriate for restoration tasks (as shown in (Liu et al. 2018) for image super-resolution). After several layers of dilated filtering, the output only considers a fixed sparse sampling of input locations, resulting in significant loss of information.

Till date, the driving force behind performance improvement in deblurring has been use of large number of layers, larger filters, and multi-scale processing which assist in increasing the “static” receptive field of a CNN. Not only do these techniques offer a suboptimal design, they are also difficult to scale since the effective receptive field of deep CNNs is much smaller than the theoretical one (investigated in (Luo et al. 2016)).

We claim that a superior alternative to such image-agnostic models is a convolutional framework wherein the filters and the receptive field dynamically adapts to input image instances. Our experiments show that the latter approach is a considerably better choice due to its task-specific efficacy and utility for computationally limited environments, and it delivers consistent performance across diverse magnitudes of blur. Now, we explain the motivation for designing a deblurring network with asymmetric filters. Given a 2D image  $\mathbf{I}$  and a blur kernel  $\mathbf{K}$ , the motion blur process can be formulated as:

$$\mathbf{B}[x, y] = \sum_{m, n = -M/2}^{M/2, M/2} \mathbf{K}[m, n] \mathbf{I}[x - n, y - n], \quad (1)$$

where  $\mathbf{B}$  is the blurred image,  $[x, y]$  represents pixel coordinates, and  $M \times M$  is the size of the blur kernel. At any given location  $[x, y]$ , the sharp intensity can be represented as

$$\mathbf{I}[x, y] = \frac{\mathbf{B}[x, y]}{\mathbf{K}[0, 0]} - \sum_{m, n=-M/2}^{M/2, M/2} \frac{\mathbf{K}[m, n]\mathbf{B}[x-m, y-n]}{\mathbf{K}[0, 0]^2} + \frac{\sum_{m, n=-M/2}^{M/2, M/2} \sum_{i, j=-M/2}^{M/2, M/2} \mathbf{K}[m, n]\mathbf{K}[i, j]\mathbf{I}[x-n-i, y-n-j]}{\mathbf{K}[0, 0]^2} \quad (2)$$

which is a 2D infinite impulse response (IIR) model, whose recursive expansion would eventually lead to an expression which contains values from only the blurred image and the kernel.

The dependence of  $\mathbf{I}[x, y]$  on a large number of locations in  $\mathbf{B}$  shows that the deconvolution process requires infinite signal information. If we assume that the boundary of the image contains zeros, eq. 2 is equivalent to applying an inverse filter to  $\mathbf{B}$ . As visualized in (Zhang et al. 2018b), the non-zero region of such an inverse deblurring filter is typically much larger than the blur kernel. Thus, if we use a CNN to model the process, a large receptive field should be considered to cover all the pixel positions that are necessary for deblurring. Eq. 2 also shows that only a few coefficients (which are  $\mathbf{K}[m, n]$  for  $m, n \in [-M/2, M/2]$ ) need to be estimated by the deblurring model, provided we can find an appropriate model with large enough receptive field.

For this theoretical analysis, we will temporarily assume that the motion blur kernel  $K$  is linear (as assumed in prior deblurring works (Sun et al. 2015; Gong et al. 2017)). Now, consider an image  $\mathbf{B}$  which is affected by motion blur in the horizontal direction (without loss of generality), implying  $\mathbf{K}[m, n] = 0$  for  $m \neq 0$  (non-zero values present only in the middle row of the kernel). For this case, eq. 2 translates to

$$\mathbf{I}[x, y] = \frac{\mathbf{B}[x, y]}{\mathbf{K}[0, 0]} - \sum_{n=1}^M \frac{\mathbf{K}[0, n]\mathbf{B}[x, y-n]}{\mathbf{K}[0, 0]^2} + \frac{\sum_{n=1}^M \sum_{j=1}^M \mathbf{K}[0, n]\mathbf{K}[0, j]\mathbf{I}[x, y-n-j]}{\mathbf{K}[0, 0]^2} = \dots \quad (3)$$

We observe that for this case,  $\mathbf{I}[x, y]$  can be expressed as a function of only one row of pixels in the blurred image  $\mathbf{B}$ , which implies that for a horizontal blur kernel, the deblurring filter is also purely horizontal. We use this observation to state a hypothesis that holds for any motion blur kernel: “Deblurring filters are directional/asymmetric in shape”. The reason behind this is the well known inherently directional nature of motion blur kernels. Such an operation can be efficiently learnt by a CNN with adaptive and asymmetric filters and this forms the basis for our work. Next, we describe our proposed network in detail.

### Dense encoder decoder backbone

Inspired by the success of fully convolutional networks that directly estimate the intensities of the deblurred image, we build an encoder-decoder architecture composed of densely connected modules. These modules immensely improve feature extraction capability by reusing features across multiple layers and their connections maximize information flow

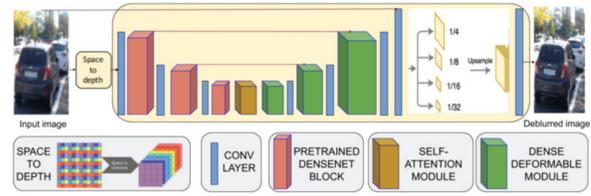


Figure 2: Proposed deblurring network and its components.

along the intermediate layers and result in better convergence. Hence, they are more efficient and learn more complex features than a network with residual connections (used extensively in recent deblurring methods).

A key component of our Region-Adaptive Dense Network is the encoder which progressively extracts a feature pyramid from the input image. The first convolution begins after a space-to-depth module that transforms the image pixels to channel-space using pixel-shuffling by a factor of 2. This allows subsequent computationally intensive operations to be performed at lower spatial resolution, hence reducing computational and memory footprint while increasing the receptive field. This layer is followed by three dense-blocks containing 12, 16, and 24 dense units, respectively, with growth rate (GR) set to 32 and each dense unit consisting of batch norm, ReLU,  $1 \times 1$  conv ( $4 \times \text{GR}$  channels) followed by batch norm, ReLU,  $3 \times 3$  conv (GR channels). The first two dense blocks are followed by  $1 \times 1$  conv and pooling to down-sample the features-maps. Our design allows all these filters to be initialized using initial layers of pre-trained DenseNet-121 (Huang et al. 2017), and our experiments demonstrate its advantage.

The decoder is built using our dynamically adaptive components namely, Self-Attention (SA) module and Dense Deformable Module (DDM). SA module accepts the low-resolution output of the encoder and generates a non-locally enhanced feature-map. These features are sequentially processed by 3 DDMs and deconvolution layers to reach the output image resolution. Similar to U-net, intermediate features with in the decoder are concatenated with the corresponding-sized encoder features. Further, result of the final deconvolution layer is enhanced through multi-scale context aggregation through pooling and upsampling at 4 scales, before being fed to the final reconstruction layer. The output is the residual between the ground-truth sharp image and the input blurred image. Note that unlike several prior works, our network does not use large ( $5 \times 5$ ) filters. A schematic of the proposed architecture is shown in Fig. 2.

### Dense deformable module (DDM)

CNNs operate on fixed locations in a regular grid which limits their ability to model unknown geometric transformations. Spatial Transform Networks (STN) (Jaderberg et al. 2015) introduced spatial transformation learning into CNNs, wherein an image-dependent global parametric transformation is estimated and applied on the feature map. However, such warping is computationally expensive and the transformation is considered to be global across the whole im-

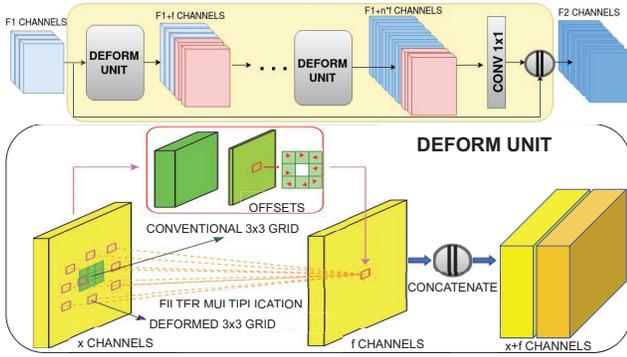


Figure 3: Schematic of our dense deformable module.

age, which is not the case for motion in dynamic and 3D scenes where different regions are affected by different magnitude and direction of motion. To introduce such motion-awareness in our network, we adopt deformable convolutions (Dai et al. 2017), which enable local transformation learning in an efficient manner. Unlike regular convolutional layers, the deformable convolution also learns to estimate the shapes of convolution filters conditioned on an input feature map (Xu, Li, and Sun 2019). While maintaining filter weights invariant to the input, a deformable convolution layer first learns a pixel-level offset map from the input, and applies it to the regular feature map for re-sampling.

Our DDM (shown in Fig. 3) contains a densely connected set of deform units capable of learning positioning of filters on the feature sampling grid. Each deform unit in DDM contains a pair of layers: one to estimate the 2D filter offsets for each spatial location and another to apply the learnt filters on values sampled from these locations. The 2D offsets (shown as red-arrows) are encoded in the channel dimension of an estimated tensor (shown in green).

At each spatial coordinate  $l_0$ , the deform unit estimates 2 offsets for  $P$  sampling locations. Our DDM is composed of  $3 \times 3$  kernels, which correspond to  $P = 9$  and the fixed offsets  $\mathbf{l}_p \in (-1, -1), (-1, 0), \dots, (1, 1)$ . If  $w_p$  denotes the learnt filter weight for the  $p^{\text{th}}$  location, the processed feature values  $\hat{F}$  at coordinate  $l_0$  are then obtained from input feature  $F$  as:

$$\hat{F}(\mathbf{l}_0) = \sum_{p=1}^P w_p \cdot F(\mathbf{l}_0 + \mathbf{l}_p + \Delta\mathbf{l}_p) \quad (4)$$

where  $\Delta\mathbf{l}_p$  are the dynamically estimated sampling offsets for all the channels in the input feature map, which determine the shifting of the  $P$  filter locations along horizontal and vertical axes. As a result, the regular convolution filter operates on an irregular grid of pixels. Since the offsets can be fractional, bilinear interpolation is used to sample from the input feature map. All the parts of our network are trainable end-to-end, since bilinear sampling and the grid generation of the warping module are both differentiable (Paszke et al. 2017). The 6 deformable layers in our DDM are followed by a  $1 \times 1$  conv layer to reduce the number of channels.

Although the focus of our work is an effective network

design, our analysis also presents an effective way to boost any existing deblurring network’s performance. Replacing ordinary convolution layers with deformable layers is much more efficient than going deeper or wider since the receptive field and the spatial sampling locations become dynamically adaptable according to the scale, shape, and location of blur.

## Self-attention module (SA)

Recent deblurring works have emphasized the advantages of multi-scale/patch processing. It efficiently captures different scales of motion blur, and increases the receptive field of the network. Although it facilitates local growth in receptive field, it does not leverage the relationship between two distant locations in the scene. While this coarse-to-fine approach helps to handle different magnitudes of blur, it cannot leverage the relationship among blurred regions from a global perspective, which is also beneficial for the restoration task at hand. In this work, we employ a better strategy: attention based learnable non-local connections among features at different spatial locations.

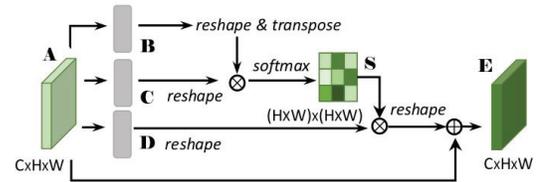


Figure 4: A schematic of our self-attention module.

Trainable attention over features for modeling long-range dependencies has shown its benefits in several tasks spanning across language (Shen et al. 2018; Vaswani et al. 2017) and vision (Wang et al. 2018), but has not been explored for deblurring. Our work is inspired by the recent work of (Zhang et al. 2018a) that utilizes non-local attention to connect different scene regions and uses it to improve image generation quality.

Our SA module selectively aggregates the features at each position by a weighted sum of the features at all positions. This efficient design ensures that similar features are connected to each other regardless of their spatial distances, which helps in directly connecting regions with similar blur. It has two advantages: First, it overcomes the issue of limited receptive field, as any pixel has access to features at every other pixel in the image. Second, it implicitly acts as a gate for propagating only relevant information across the layers. These properties make it suitable for deblurring, since blur affecting various scene-edges is often correlated.

As illustrated in Fig. 4, input feature-map  $\mathbf{A} \in \mathbb{R}^{C \times H \times W}$  is transformed into two new feature maps  $\mathbf{B}$  and  $\mathbf{C} \in \mathbb{R}^{\hat{C} \times H \times W}$  ( $\hat{C} = \frac{C}{8}$ ). Next, we reshape  $\mathbf{C}$  and  $\mathbf{B}$  to  $\mathbb{R}^{\hat{C} \times N}$  ( $N = H \times W$ ), perform matrix multiplication between their transposed versions, and pass the product through a softmax

layer to calculate the attention map  $\mathbf{S} \in \mathbb{R}^{N \times N}$ :

$$s_{ji} = \frac{\exp(\mathbf{B}_i \cdot \mathbf{C}_j)}{\sum_{i=1}^N \exp(\mathbf{B}_i \cdot \mathbf{C}_j)} \quad (5)$$

where  $s_{ji}$  measures the  $i^{\text{th}}$  position’s impact on  $j^{\text{th}}$  position. Note that similarity between feature representations of any two position contributes to greater correlation (higher attention) between them.

Finally,  $\mathbf{A}$  is processed to obtain  $\mathbf{D}$  which is then reshaped to  $\mathbb{R}^{C \times N}$ . A matrix multiplication between  $\mathbf{D}$  and the transpose of  $\mathbf{S}$  yields an enhanced feature-map residual, which is added to  $\mathbf{A}$  to obtain the final output  $\mathbf{E}$  as:

$$\mathbf{E}_j = \mathbf{A}_j + \sum_{i=1}^N s_{ji} \mathbf{D}_i \quad (6)$$

The resulting feature  $\mathbf{E}$  at each position is a weighted sum of the features at all positions and original features. Therefore, it has global context and selectively aggregated contexts according to the spatial attention map, causing similar features to reinforce gains and irrelevant features to get subdued. We place the SA module at the beginning of our region adaptive decoder to minimize memory footprint.

## Experimental Setup

Since the prime application of our work is efficient deblurring of general 3D and dynamic scenes, we perform training and evaluation of our network on the dynamic scene deblurring benchmark (Nah, Kim, and Lee 2017), following recent learning-based works. This dataset is constructed using 240fps videos captured using GoPro camera and contains diverse 3D scenes captured in presence of significant object and camera motion. Following the same train-test split as in (Nah, Kim, and Lee 2017), we use 2103 pairs for training and 1111 pairs for evaluation. Training is done for  $1 \times 10^6$  iterations using Adam optimizer with learning rate 0.0001 on patches of  $256 \times 256$  and batch-size of 16. We conduct our experiments on a PC with Intel Xeon E5 CPU, 256 GB RAM and an NVIDIA Titan X GPU.

## Experimental Results

In this section, we carry out quantitative and qualitative comparisons of our architectures with state-of-the-art end-to-end learning based methods for image deblurring, including MSCNN (Nah, Kim, and Lee 2017), DeblurGAN (Kupyn et al. 2017), SVRNN (Zhang et al. 2018b), SRN (Tao et al. 2018), PSS-SRN (Gao et al. 2019), and Stack(4)-DMPHN (Zhang et al. 2019). Due to the complexity of the blur present in general dynamic scenes, conventional blur model based approaches struggle to perform well. Nevertheless, we also compare with conventional non-uniform deblurring approaches of (Whyte et al. 2012), (Hyun Kim, Ahn, and Mu Lee 2013) and MBMF (Gong et al. 2017). Public implementations with default parameters were used to obtain qualitative results on selected test images.

**Quantitative Evaluation:** Quantitative comparisons using PSNR and SSIM scores obtained on the GoPro testing set

Models	PSNR	SSIM	Runtime
(Hyun et al. 2013)	23.8	0.8358	2000000
(Whyte et al. 2012)	24.6	0.8458	700000
MBMF	24.64	0.8429	120000
MSCNN	29.23	0.9162	6000
DeblurGAN	28.70	0.8580	1000
SVRNN	29.19	0.9306	1400
SRN	30.10	0.9323	1600
PSS-SRN	30.92	0.9421	1600
Stack(4)-DMPHN	31.20	0.9451	700
RADN1	30.77	0.9417	32
RADN2	31.17	0.9454	32
RADN3	31.42	0.9494	34
RADN4	31.58	0.9503	36
RADN(Final)	<b>31.76</b>	<b>0.9530</b>	<b>38</b>
RADN(Final) <sup>+</sup>	<b>32.15</b>	<b>0.9560</b>	-

Table 1: Quantitative analysis of our proposed model on the GoPro dataset (Nah, Kim, and Lee 2017), Runtime is in milliseconds.

( $720 \times 1280$  images) are presented in Table 1. Since traditional methods cannot model combined effects of general camera shake and object motion (Whyte et al. 2012) or forward motion and depth variations (Hyun Kim, Ahn, and Mu Lee 2013), they fail to faithfully restore most of the images in the test-set. The below par performance of MBMF can be attributed to the fact that it uses synthetic and simplistic blur kernels to train a CNN and employs a traditional deconvolution method to estimate the sharp image, which severely limits its applicability to general dynamic scenes. End-to-end residual networks such as MS-CNN, SRN and PSS-SRN use multi-scale strategy or alternative losses to improve large blur handling capability, but fail in challenging situations. The proposed RADN(final) significantly outperforms these works, including the spatially varying model SVRNN. Importantly, our method fares significantly better than the nearest competitor Stack(4)-DMPHN, in terms of inference-time ( $15\times$  faster, supporting real-time deblurring at 28fps), and accuracy (improvement of  $\sim 0.56$  dB), while requiring 40% less parameters. RADN(final)<sup>+</sup> represents our results obtained using geometric self-ensemble (Lim et al. 2017), which offers significant boost in test performance without requiring further training or additional parameters.

**Qualitative Evaluation:** Visual comparisons on different dynamic and 3D scenes are given in Fig. 5. It shows that results of prior works suffer from incomplete deblurring or artifacts. In contrast, our network is able to restore scene details (text, edges etc.) more faithfully due to its effectiveness in handling large dynamic blur. An additional advantage over (Hyun Kim, Ahn, and Mu Lee 2013; Whyte et al. 2012) is that our model waives-off the requirement of parameter tuning during test phase. Additional experiments and qualitative comparisons on other benchmarks are provided in the supplementary material.

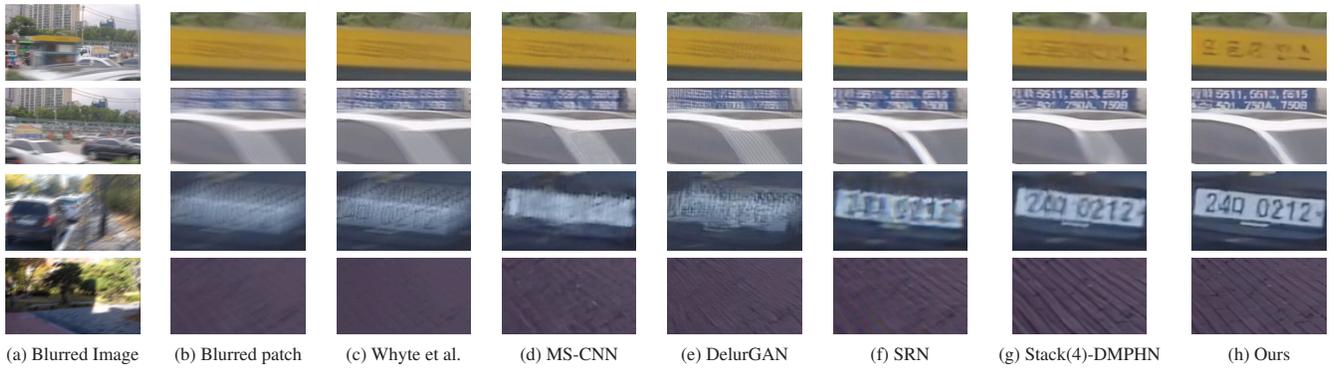


Figure 5: Visual comparisons of deblurring results on images from the GoPro test set (Nah, Kim, and Lee 2017). Key blurred patches are shown in (b), while zoomed-in patches from the deblurred results are shown in (c)-(h).

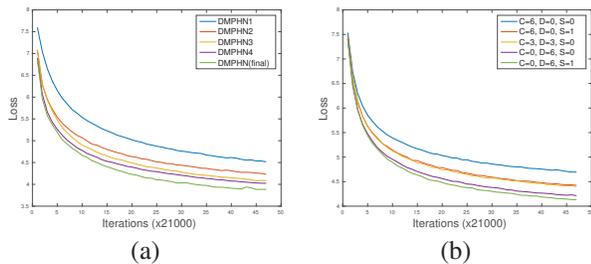


Figure 6: Network analysis through comparison of training performance of various ablations of our models.

Table 2: Quantitative comparison of different ablations of our baseline residual deblurring network on GoPro testset.

Conv layers	6	6	3	0	0
Deformable layers	0	0	3	6	6
SA module	✗	✓	✗	✗	✓
PSNR (dB)	29.8	30.54	30.69	31.05	31.13
Size (MB)	10.6	10.7	10.9	11.2	11.2

## Ablation studies

### Ablations of our proposed network

We analyze the effect of individual components of our network on its training and testing performance. The test scores of various ablations our network are reported in Table 1 and a comparison of their training performance is shown in Fig. 6(a). Decoder is a key component of our deblurring network and To evaluate its importance, we compare the proposed model RADN(Final) with 3 different versions of the decoder design while using the same dense encoder. In version RADN4, we replace the deform units within DDMs with ordinary convolutional layers, forcing the modules to apply rigid filters at all the spatial locations in the features. The drastic decrease in performance demonstrates the importance of spatially adaptive filter-offset learning capability. In version RADN3, we further remove the SA module from RADN4 and observe notable drop in performance. We chose to keep a single SA module in our network, since performance improvement beyond it was marginal and it

serves as a good balance between restoration accuracy and processing-time. In RADN2, we replace DDMs with bottleneck blocks (Huang et al. 2017) which are a simpler alternative to dense blocks. The significant fall in accuracy can be attributed to the efficacy of densely-connected layers and validates our design choice. Finally, to verify the effectiveness of pretraining, we train a baseline RADN1 which is identical to RADN2 except that its encoder layers are not pre-trained (using DenseNet). The PSNR and loss difference clearly shows that such pre-training leads to better initialization in parameter space and eventually, better convergence. Complete ablation (removing all dense layers and space2depth) leads to basic U-net, scoring only 28.9 dB.

### Analysis on a baseline deblurring network

Here, we evaluate the advantages of our region-adaptive modules by introducing them into a simpler residual encoder decoder baseline adopted from SRN (one of the competing methods). We differentiate this baseline’s design from SRN in terms of compactness and computational footprint by employing only  $3 \times 3$  filters for economy and removing the recurrent units. Next, we describe experiments with inclusion of our region-adaptive modules which significantly improve the representational capacity and performance of this network without sacrificing the computational efficiency.

One of the key hyper-parameters in this analysis is the number of deformable units that replace normal convolutions (in the trunk of the network). To study its effectiveness, we designed and trained 3 versions of the network wherein the number of such replacements are 0, 3 and 6, respectively. Fig. 6(b) shows comparisons of the convergence plots of these models. It can be observed that the training performance as well as the quantitative results get better with increase in the number of deformable blocks, as it introduces additional filter adaptability into the network. Also note that the performance of the model which contains SA but no deformable layers, is expectedly lower than other models but better than the plain CNN. Finally, the best training performance is delivered by our final model which contains 6 deformable replacements and SAs, which shows that the advantages of the two modules are complementary and their union leads to a superior model. These improvements are

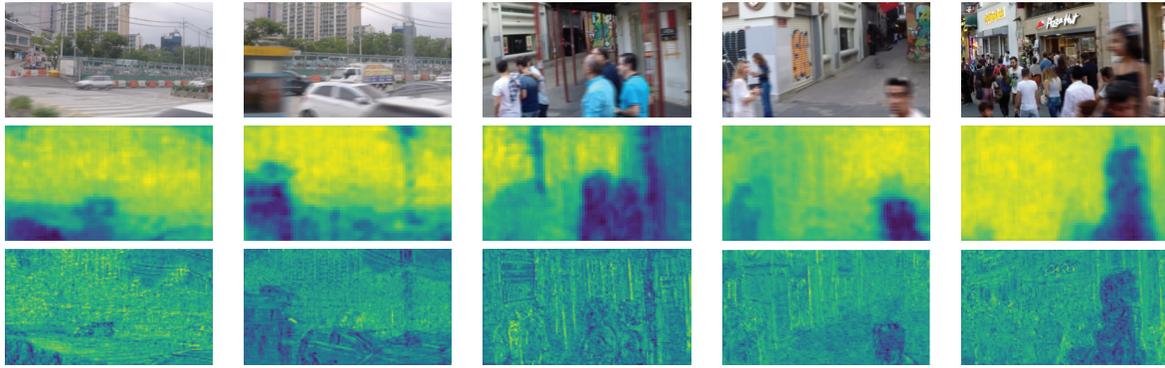


Figure 7: Visualizations of the spatial variations in the estimated filter offsets and attention-based feature modulations on blurred images from the GoPro test set. The second row shows the spatial distribution of the horizontal-offset values for the filter in a DDM of our network. Third row shows the corresponding enhanced feature-map residuals estimated by our SA module.

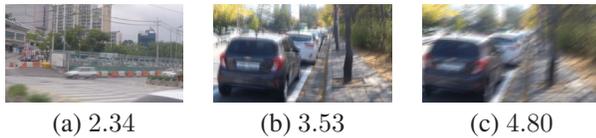


Figure 8: Offset magnitudes: Values mentioned under each image (from GoPro test set) depict the average magnitude of the offsets estimated by all the DDMs in our network.

also reflected in the quantitative values reported in Table 2.

The table also shows that our modules are lightweight since their inclusion has only a marginal effect on the model size. Although the proposed network is already quite efficient, replacing the standard convolutions in our network with grouped convolution and/or separable convolutions can lead to further improvements. More details and analyses of our networks are provided in the supplementary document.

### Network Visualization

This section provides further insights into the effectiveness of our approach by visualizing the filter offsets and feature-attention maps dynamically estimated by our network.

**Dense deformable module:** To establish our DDM’s interpret-ability and their sensitivity to local motion in the scene, we investigate visible association between the estimated filter transformation-maps and the dominant motion blur in the input image. In first and second row of Fig. 7, we show various test images and a representation of the corresponding offset-maps, respectively. The offset value at each spatial location is calculated using the difference between the horizontal (x) offsets of the 1<sup>st</sup> and 3<sup>rd</sup> columns of the 3×3 filters from the 2<sup>nd</sup> DDM of our network. It can be seen that the offset values are higher for the foreground regions and regions undergoing large motion. This is important for the deblurring process since the pixels corresponding to the foreground objects (e.g., people, vehicles) experience different motion compared to other parts of the scene, due to independent motion or depth differences. This demonstrates that the proposed network can distinguish between differ-

ently blurred regions and that the estimated offsets are correlated with the direction and magnitude of the motion blur.

Next, we measure the correlation between the amount of scene motion and the average magnitude of the filter offsets estimated by our network. In Fig. 8, we show various test images and the corresponding average value of offsets calculated by all the DDMs. It can be seen that the offset values are small for the first scene (Fig. 8(a)) where only few regions are blurred. As the blur increases and affects larger number of pixels in the scene (Fig. 8(b) and (c)), the average offset magnitude rises to accommodate the large blur. This shows that our network features a dynamic receptive field and adapts to scenes with varying degrees of blur without altering the filter weights, yielding efficiency to the process.

**Self-attention module:** Here, we visualize the inputs and outputs of the SA module of our network for additional insights. In third row of Fig. 7, we show the residuals estimated within our SA module ( $\sum_{i=1}^N s_{ji} \mathbf{D}_i$  in Eq.6) for corresponding images in the first row. It can be observed that the SA module learns to amplify the correlation among feature-map magnitudes for all pixels undergoing similar (large) blur, and the spatial distributions agree with corresponding dynamic regions and depth variations. Such enhancement of the intermediate feature-maps is the key reason behind the observed improvement in training and test performance.

### Conclusions

We proposed an efficient motion deblurring architecture composed of convolutional modules that enable spatially adaptive feature learning through filter transformations and feature attention, namely dense deformable module (DDM) and self-attention (SA) module. The DDMs implicitly address the shifts responsible for the local blur in the input image, while the SA module meaningfully connects non-locally distributed blurred regions. Introducing these modules awards higher capacity to any deblurring network without any notable increase in computational footprint. This effectiveness is shown by incorporating them into a new densely connected encoder decoder backbone, wherein we also show the benefits of pre-training of encoder filters.

Compared against existing deep deblurring frameworks, our model achieves the state-of-the-art performance and is able to run at 28fps for 720p images. We believe our spatially-aware design can be utilized for other image processing and vision tasks as well, and we shall explore them in the future.

## References

- Chakrabarti, A. 2016. A neural approach to blind motion deblurring. In *European Conference on Computer Vision*, 221–235. Springer.
- Chen, J.; Adams, A.; Wadhwa, N.; and Hasinoff, S. W. 2016. Bilateral guided upsampling. *ACM Transactions on Graphics (TOG)* 35(6):203.
- Chen, Q.; Xu, J.; and Koltun, V. 2017. Fast image processing with fully-convolutional networks. In *Proceedings of the IEEE International Conference on Computer Vision*, 2497–2506.
- Dai, J.; Qi, H.; Xiong, Y.; Li, Y.; Zhang, G.; Hu, H.; and Wei, Y. 2017. Deformable convolutional networks. In *Proceedings of the IEEE international conference on computer vision*, 764–773.
- Gao, H.; Tao, X.; Shen, X.; and Jia, J. 2019. Dynamic scene deblurring with parameter selective sharing and nested skip connections. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 3848–3856.
- Gong, D.; Yang, J.; Liu, L.; Zhang, Y.; Reid, I.; Shen, C.; 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)*.
- Huang, G.; Liu, Z.; Van Der Maaten, L.; and Weinberger, K. Q. 2017. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 4700–4708.
- Hyun Kim, T.; Ahn, B.; and Mu Lee, K. 2013. Dynamic scene deblurring. In *Proceedings of the IEEE International Conference on Computer Vision*, 3160–3167.
- Jaderberg, M.; Simonyan, K.; Zisserman, A.; et al. 2015. Spatial transformer networks. In *Advances in Neural Information Processing Systems*, 2017–2025.
- Kupyn, O.; Budzan, V.; Mykhailych, M.; Mishkin, D.; and Matas, J. 2017. Deblurgan: Blind motion deblurring using conditional adversarial networks. *arXiv preprint arXiv:1711.07064*.
- Lai, W.-S.; Huang, J.-B.; Hu, Z.; Ahuja, N.; and Yang, M.-H. 2016. A comparative study for single image blind deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 1701–1709.
- Lim, B.; Son, S.; Kim, H.; Nah, S.; and Mu Lee, K. 2017. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, 136–144.
- Liu, P.; Zhang, H.; Zhang, K.; Lin, L.; and Zuo, W. 2018. Multi-level wavelet-cnn for image restoration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 773–782.
- Luo, W.; Li, Y.; Urtasun, R.; and Zemel, R. 2016. Understanding the effective receptive field in deep convolutional neural networks. In *Advances in neural information processing systems*, 4898–4906.
- Nah, S.; Kim, T. H.; and Lee, K. M. 2017. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *CVPR*, volume 1, 3.
- Paszke, A.; Gross, S.; Chintala, S.; Chanan, G.; Yang, E.; DeVito, Z.; Lin, Z.; Desmaison, A.; Antiga, L.; and Lerer, A. 2017. Automatic differentiation in pytorch.
- 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.
- Schuler, C. J.; Hirsch, M.; Harmeling, S.; and Schölkopf, B. 2016. Learning to deblur. *IEEE transactions on pattern analysis and machine intelligence* 38(7):1439–1451.
- Shen, T.; Zhou, T.; Long, G.; Jiang, J.; Pan, S.; and Zhang, C. 2018. Disan: Directional self-attention network for rnn/cnn-free language understanding. In *Thirty-Second AAAI Conference on Artificial Intelligence*.
- 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.
- Tao, X.; Gao, H.; Shen, X.; Wang, J.; and Jia, J. 2018. Scale-recurrent network for deep image deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 8174–8182.
- Vasiljevic, I.; Chakrabarti, A.; and Shakhnarovich, G. 2016. Examining the impact of blur on recognition by convolutional networks. *arXiv preprint arXiv:1611.05760*.
- Vasu, S., and Rajagopalan, A. 2017. From local to global: Edge profiles to camera motion in blurred images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4447–4456.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, Ł.; and Polosukhin, I. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, 5998–6008.
- Wang, X.; Girshick, R.; Gupta, A.; and He, K. 2018. Non-local neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 7794–7803.
- Whyte, O.; Sivic, J.; Zisserman, A.; and Ponce, J. 2012. Non-uniform deblurring for shaken images. *International journal of computer vision* 98(2):168–186.
- Xu, X.; Li, M.; and Sun, W. 2019. Learning deformable kernels for image and video denoising. *arXiv preprint arXiv:1904.06903*.
- Yan, Y.; Ren, W.; Guo, Y.; Wang, R.; and Cao, X. 2017. Image deblurring via extreme channels prior. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 4003–4011.
- Zhang, H.; Goodfellow, I.; Metaxas, D.; and Odena, A. 2018a. Self-attention generative adversarial networks. *arXiv preprint arXiv:1805.08318*.
- Zhang, J.; Pan, J.; Ren, J.; Song, Y.; Bao, L.; Lau, R. W.; and Yang, M.-H. 2018b. Dynamic scene deblurring using spatially variant recurrent neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2521–2529.
- Zhang, H.; Dai, Y.; Li, H.; and Koniusz, P. 2019. Deep stacked hierarchical multi-patch network for image deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 5978–5986.