

Deblurring Face Images Using Uncertainty Guided Multi-Stream Semantic Networks

Rajeev Yasarla¹, Graduate Student Member, IEEE, Federico Perazzi, Member, IEEE,
and Vishal M. Patel¹, Senior Member, IEEE

Abstract—We propose a novel multi-stream architecture and training methodology that exploits semantic labels for facial image deblurring. The proposed Uncertainty Guided Multi-Stream Semantic Network (UMSN) processes regions belonging to each semantic class independently and learns to combine their outputs into the final deblurred result. Pixel-wise semantic labels are obtained using a segmentation network. A predicted confidence measure is used during training to guide the network towards the challenging regions of the human face such as the eyes and nose. The entire network is trained in an end-to-end fashion. Comprehensive experiments on three different face datasets demonstrate that the proposed method achieves significant improvements over the recent state-of-the-art face deblurring methods. Code is available at: <https://github.com/rajeevyasarla/UMSN-Face-Deblurring>

Index Terms—Facial image deblurring, semantic masks, confidence scores.

I. INTRODUCTION

IMAGE deblurring entails the recovery of an unknown true image from a blurry image. Similar to the other image enhancement tasks, image deblurring is experiencing a renaissance as a result of convolutional networks (CNNs) establishing themselves as powerful generative models. Image deblurring is an ill-posed problem and therefore it is crucial to leverage additional properties of the data to successfully recover the lost facial details in the deblurred image. Priors such as sparsity [1]–[3], manifold [4], low-rank [5] and patch similarity [6] have been used in the literature to obtain a regularized solution. In recent years, deep learning-based methods have also gained some traction [7]–[10].

The inherent semantic structure of natural images such as faces is an important information that can be exploited to improve the deblurring results. Few techniques [11], [12] make use of such prior information in the form of semantic labels. These methods do not account for the class imbalance of

Manuscript received December 27, 2019; revised March 11, 2020 and April 4, 2020; accepted April 17, 2020. Date of publication April 30, 2020; date of current version May 8, 2020. This work was supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), through IARPA R&D under Contract 2019-19022600002. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Nikolaos Mitianoudis. (Corresponding author: Rajeev Yasarla.)

Rajeev Yasarla and Vishal M. Patel are with the Whiting School of Engineering, Johns Hopkins University, Baltimore, MD 21218 USA (e-mail: yasar11@jhu.edu; vpatel36@jhu.edu).

Federico Perazzi is with Adobe Inc., San Jose, CA 95110 USA (e-mail: perazzi@adobe.com).

Digital Object Identifier 10.1109/TIP.2020.2990354

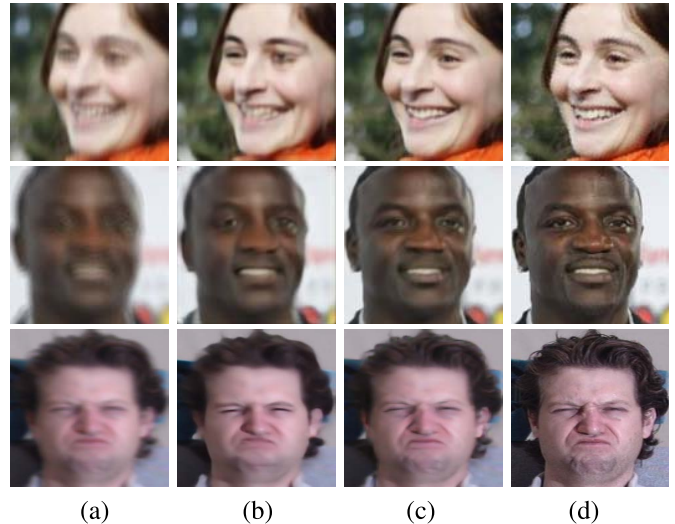


Fig. 1. Sample deblurring results: (a) Blurry image; (b) results corresponding to Shen *et al.* [11] and to Song *et al.* [12] (last row); (c) Results corresponding to our proposed Uncertainty Guided Multi-Stream Semantic Network (UMSN); (d) ground-truth. Our approach recovers more details and better preserves fine structures like eyes and hair.

semantic maps corresponding to faces. Interior parts of a face like eyes, nose, and mouth are less represented as compared to face skin, hair and background labels. Depending on the pose of the face, some of the interior parts may even disappear. Without re-weighting the importance of less represented semantic regions, the method proposed by Shen *et al.* [11] fails to reconstruct the eyes and the mouth regions as shown in Fig. 1. Similar observations can also be made regarding the method proposed in [12] (Fig. 1) which uses semantic priors to obtain the intermediate outputs and then performs post-processing using k-nearest neighbor algorithm.

To address the imbalance of different semantic classes, we propose a novel CNN architecture, called Uncertainty guided Multi-stream Semantic Networks (UMSN), which learns class-specific features independently and combine them to deblur the whole face image. Class-specific features are learned by sub-networks trained to reconstruct a single semantic class. We use nested residual learning paths to improve the propagation of semantic features. Additionally, we propose a class-based confidence measure to train the network. The confidence measure describes how well the network is likely

to deblur each semantic class. This measure is incorporated in the loss to train the network. We evaluate the proposed network by conducting experiments on three face datasets - Helen [13], CelebA [14] and PubFig [12]. Extensive experiments demonstrate the effectiveness of our approach compared to a number of well established face deblurring methods as well as other competitive approaches. Fig. 1 shows sample results from our UMSN network, where one can clearly see that UMSN is able to provide better results as compared to the state-of-the-art techniques [11], [12]. An ablation study is also conducted to demonstrate the effectiveness of different parts of the proposed network. To summarize, this paper makes the following contributions:

- A novel multi-stream architecture called, Uncertainty Guided Multi-Stream Semantic (UMSN), is proposed which learns class-specific features and uses them to reconstruct the final deblurred image.
- We propose a novel method of computing confidence measure for the reconstruction of every class in the deblurred image, which is to rebalance the importance of semantic classes during training.

Rest of the paper is organized as follows. In Section II, we review a few related works. Details of the proposed method are given in Section III. Experimental results are presented in Section IV, and finally, Section V concludes the paper with a brief summary.

II. BACKGROUND AND RELATED WORK

Single image deblurring techniques can be divided into two main categories - blind and non-blind. In the non-blind deblurring problem, the blur kernel is assumed to be known while in the blind deblurring problem, the blur kernel needs to be estimated. Since the proposed approach does not assume any knowledge of the blur kernel, it is a blind image deblurring approach. Hence, in this section, we review some recent blind image deblurring techniques proposed in the literature.

Classical image deblurring methods estimate the blur kernel given a blurry image and then apply deconvolution to get the deblurred image. To calculate the blur kernel some techniques assume prior information about the image, and formulate maximum a-posterior (MAP) to obtain the deblurred image. Different priors such sparsity, L_0 gradient prior, patch prior, manifold Prior, and low-rank have been proposed to obtain regularized reconstructions [1]–[6], [15]–[17]. Xu *et al.* [8] estimate the Fourier coefficients of the blur kernel to deblur the image. Nimisha *et al.* [7] learn the latent features of the blurry images using the latent features of the clean images to estimate the deblurred image. These CNN-based methods do not perform well compared to the state-of-the-art MAP-based methods for large motion kernels. Recent non-blind image deblurring methods like [2], [18]–[21] assume and use some knowledge about the blur kernel. Given the latent blurry input, Vasu *et al.* [21] estimate multiple latent images corresponding to different prior strengths to estimate the final deblurred image. CNN-based techniques like [22] and [23] estimate the blur kernel and address dynamic deblurring. Another recent image deblurring method based on Feature Pyramid Network

and conditional generative adversarial network (GAN) was recently proposed by Kupyn *et al.* in [24].

The usage of semantic information for image restoration is relatively unexplored. While semantic information has been used for different object classes [25]–[28], a substantial body of literature has focused their attention to human faces [11], [29]–[31]. Pan *et al.* [30] extract the edges of face parts and estimate exemplar face images, which are further used as the global prior to estimate the blur kernel. This approach is complex and computationally expensive in estimating the blur kernel. Recently, Shen *et al.* [11] proposed to use the semantic maps of a face to deblur the image. Furthermore, they introduced the content loss to improve the quality of eyes, nose and mouth regions of the face. Bulat and Tzimiropoulos [32] proposed Super-FAN architecture which is an end-to-end network that outputs high-resolution of face, along with heatmap for face alignment. Chen *et al.* [33] proposed FSRGAN architecture which consists of coarse and fine Super-Resolution (SR) networks where they estimate coarse estimate of high-resolution output image, followed by fine estimate using Fine SR Network. Yu *et al.* [34] propose a multi-task network to upsample LR image where they predict face structure along with super-resolving LR face image. Ren *et al.* [35] proposed a face video deblurring method by predicting facial structure and identity from the blurry face using a deep network that generates a textured 3D face from the video. Textured 3D face mask can be generated accurately only from videos. This method is not beneficial if we have a single face image to deblur. Lu *et al.* [36] proposed a domain-specific single face image deblurring method by disentangling the content information in an unsupervised fashion using the KL-divergence. In contrast to these methods, we learn a multi-stream network which reconstructs the deblurred images corresponding to different classes in a facial semantic map. Furthermore, we propose a new loss to train the network.

III. PROPOSED METHOD

A blurry face image y can be modeled as the convolution of a clean image x with a blur kernel k , as

$$y = k * x + \eta,$$

where $*$ denotes the convolution operation and η is noise. Given y , in blind deblurring, our objective is to estimate the underlying clean face image x . We group 11 semantic face labels into 4 classes as follows: $m_1 = \{background\}$, $m_2 = \{face\ skin\}$, $m_3 = \{left\ eyebrow, right\ eyebrow, left\ eye, right\ eye, nose, upper\ lip, lower\ lip, teeth\}$ and $m_4 = \{hair\}$. Thus, the semantic class mask of a clean image x is the union $m = m_1 \cup m_2 \cup m_3 \cup m_4$. Similarly we define the semantic class masks of a blurry image \hat{m} .

Semantic class masks for blurry image, \hat{m} are generated using the semantic segmentation network (S-Net) (Fig. 3), and given together with the blurry image as input to the deblurring network, UMSN. This is important in face deblurring as some parts like face skin and hair are easy to reconstruct, while face parts like eyes, nose, and mouth are difficult to reconstruct

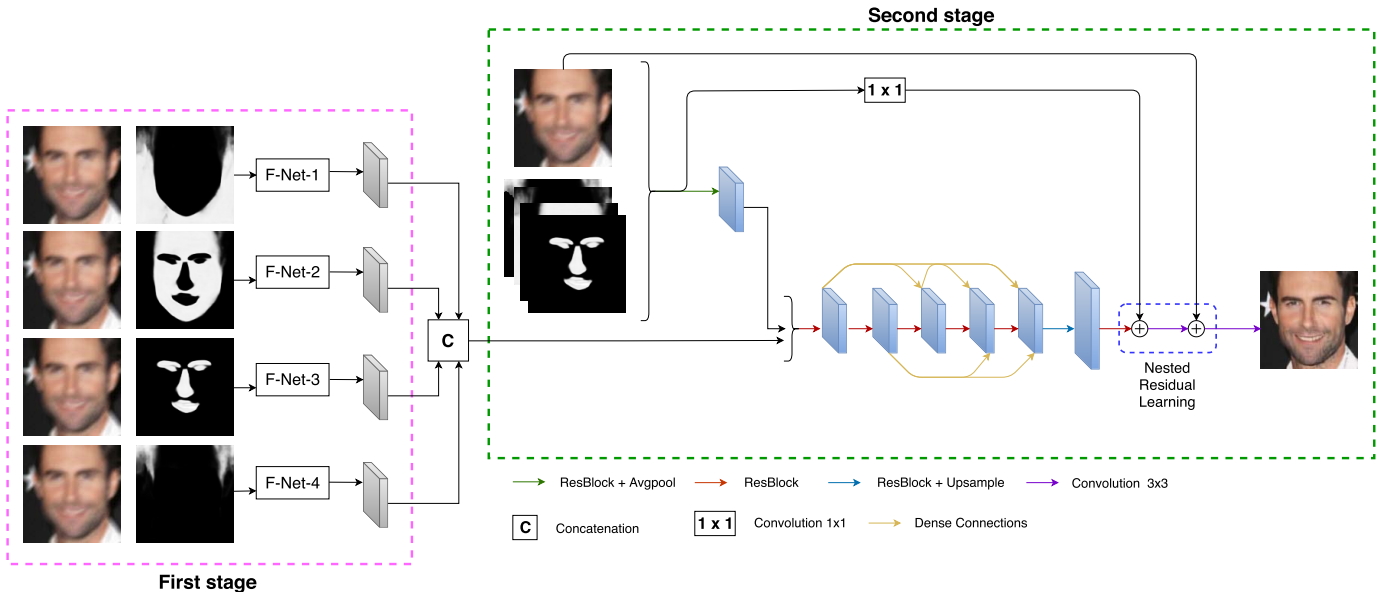


Fig. 2. An overview of the proposed UMSN network. First stage semantic networks consist of F-Net- i . Second stage is constructed using the base network (B-Net), where outputs of the F-Net- i 's are concatenated with the output of the first ResBlock layer in B-Net.

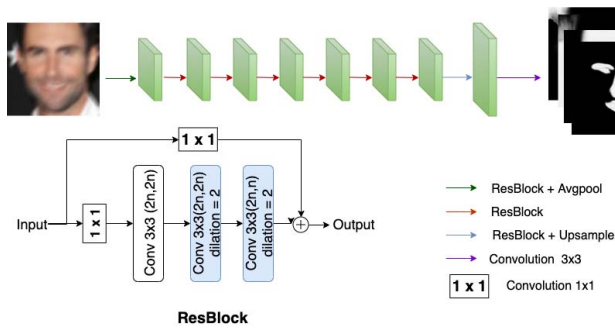


Fig. 3. An overview of the segmentation network. Conv $l \times l (p, q)$ contains instance normalization [37], Rectified Linear Unit (ReLU), Conv ($l \times l$) - convolutional layer with kernel of size $l \times l$, where p and q are the number of input and output channels respectively.

and require special attention while deblurring a face image. This is mainly due to the fact that parts like eyes, nose and mouth are small in size and contain high frequency elements compared to the other components. Different from [30] that uses edge information and [11] that feed the semantic map to a single-stream deblurring network, we address this problem by proposing a multi-stream semantic network, in which individual branches F-net- i learn to reconstruct different parts of the face image separately. Fig. 2 gives an overview of the proposed UMSN method.

As can be seen from Fig. 2, the proposed UMSN network consists of two stages. We generate the semantic class maps, \hat{m} , of a blurry face image using the S-Net network. The semantic maps are used as the global masks to guide each stream of the first stage network. These semantic class maps \hat{m} are further used to learn class specific residual feature maps with nested residual learning paths (NRL). In the first stage of our network, the weights are learned to deblur the corresponding class of the face image. In the second stage

of the network, the outputs from the first stage are fused to learn the residual maps that are added to the blurry image to obtain the final deblurred image. We train the proposed network with a confidence guided class-based loss.

A. Semantic Segmentation Network (S-Net)

The semantic class maps \hat{m}_i of a face, are extracted using the S-Net network as shown in Fig. 3. We use residual blocks (ResBlock) as our building module for the segmentation network. A ResBlock consist of a 1×1 convolution layer, a 3×3 convolution layer and two 3×3 convolution layers with dilation factor of 2 as shown in Fig. 4.

B. Base Network (B-Net)

We construct our base network using a combination of UNet [38] and DenseNet [39] architectures with the ResBlock as our basic building block. To increase the receptive field size, we introduce smoothed dilation convolutions in the ResBlock as shown in Fig. 4. B-Net is a sequence of eight ResBlocks similar to the first stage semantic network as shown in Fig. 5. Note that all convolutional layers are densely connected [39]. We follow residual-based learning in estimating the deblurred image for our base network as shown in Fig. 5.

C. UMSN Network

The UMSN network is a two-stage network. The first stage network is designed to obtain deblurred outputs from the semantic class-wise blurry inputs. These outputs are further processed by the second stage network to obtain the final deblurred image. The first stage semantic network contains a sequence of five ResBlocks with residual connections, as shown in Fig. 5. We call the set of all convolution layers of the first stage network excluding the last ResBlock and Conv 3×3 as F-Net.

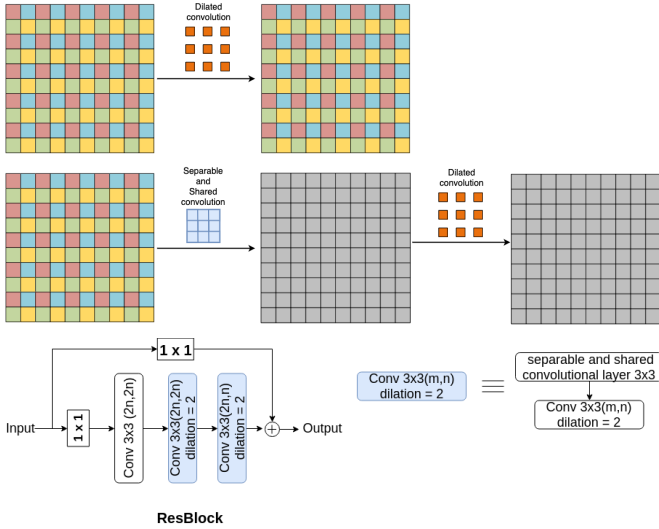


Fig. 4. First row shows an example of gridding effects caused by the dilated convolutions. For instance, a red pixel in the output feature map will be a function of the corresponding red pixels in the input feature map. This causes gridding artifacts. Second row shows the benefit of using separable and shared convolution layer, where every pixel in the output feature map is a function of every pixel in the corresponding neighborhood of the input feature map. Third row is an overview of the ResBlock. Conv $l \times l(p, q)$ contains Instance Normalization [37], ReLU - Rectified Linear Units, convolutional layer with kernel of size $l \times l$, where p and q are number of input and output channels respectively. In the right side of the figure, we show smoothed dilation convolutions introduced in ResBlock which is similar to [40].




Fig. 5. An overview of the first stage semantic network. We define the set of all convolution layers of the first stage network excluding the last ResBlock and Conv 3×3 as F-Net.

The blurry image y and the semantic masks \hat{m}_i are fed to F-Net- i to obtain the corresponding class-specific deblurred features which are concatenated with the output of the first layer (ResBlock-Avgpool) in Base Network(B-Net) for constructing the UMSN network. We also use the Nested Residual Learning (NRL) in UMSN network where class specific residual feature maps are learned and further used in estimating the residual feature maps that are added to the blurry image for obtaining the final output. For example, we can observe a residual connection between the last layer of UMSN and class-specific feature maps obtained from Conv 1×1 using y and \hat{m} . Output of this residual connection is further processed as the input to the residual connection with the input blurry image, y . In this way, we define our NRL and obtain the final deblurred image. We propose a class-based loss function to train the UMSN network.

D. Loss for UMSN

The network parameters Θ are learned by minimizing a loss \mathcal{L} as follows,

$$\hat{\Theta} = \underset{\Theta}{\operatorname{argmin}} \mathcal{L}(f_{\Theta}(y, \hat{m}), x) = \underset{\Theta}{\operatorname{argmin}} \mathcal{L}(\hat{x}, x), \quad (1)$$

where $f_{\Theta}(\cdot)$ represents the UMSN network, \hat{x} is the deblurred result, \hat{m} is the semantic map obtained from S-Net. We define the reconstruction loss as $\mathcal{L} = \|x - \hat{x}\|_1$. A face image can be expressed as the sum of masked images using the semantic maps as

$$x = \sum_{i=1}^M m_i \odot x,$$

where \odot is the element-wise multiplication and M is the total number of semantic maps. As the masks are independent of one another, Eq. (1) can be re-written as,

$$\hat{\Theta} = \underset{\Theta}{\operatorname{argmin}} \sum_{i=1}^M \mathcal{L}(m_i \odot \hat{x}, m_i \odot x). \quad (2)$$

In other words, the loss is calculated for every class independently and summed up in order to obtain the overall loss as follows,

$$\mathcal{L}(\hat{x}, x) = \sum_{i=1}^M \mathcal{L}(m_i \odot \hat{x}, m_i \odot x). \quad (3)$$

E. Uncertainty Guidance

We introduce a confidence measure for every class and use it to re-weight the contribution of the loss from each class to the total loss. By introducing a confidence measure and re-weighting the loss, we benefit in two ways. If the network is giving less importance to a particular class by not learning appropriate features of it, then the Confidence Network (CN) helps UMSN to learn those class specific features by estimating low confidence values and higher gradients for those classes through the CN network. Additionally, by re-weighting the contribution of loss from each class, it counters for the imbalances in the error estimation from different classes. The loss function can be written as,

$$\mathcal{L}_c(\hat{x}, x) = \sum_{i=1}^M C_i \mathcal{L}(m_i \odot \hat{x}, m_i \odot x) - \lambda \log(C_i), \quad (4)$$

where $\log(C_i)$ acts as a regularizer that prevents the value of C_i going to zero and λ is a constant. We estimate the confidence measure C_i for each class by passing $m_i \odot \hat{x}$, $m_i \odot x$ as inputs to CN as shown in Fig. 6. C_i represents how confident UMSN is in deblurring the i^{th} class components of the face image. Note that, $C_i \in [0, 1]$, confidence measure is used only in the loss function while training the weights of UMSN, and it is not used (or estimated) during inference.

Inspired by the benefits of the perceptual loss in style transfer [41], [42] and image super-resolution [43], we use it to train our network. Let $\Phi(\cdot)$ denote the features obtained using the VGG-Face model [44], then the perceptual loss is defined as follows,

$$\mathcal{L}_p = \|\Phi(\hat{x}) - \Phi(x)\|_2^2. \quad (5)$$

The features from layer *relu1_2* of a pretrained VGG-Face network [44] are used to compute the perceptual loss. The total loss used to train UMSN is as follows,

$$\mathcal{L}_{total} = \mathcal{L}_c + \lambda_1 \mathcal{L}_p, \quad (6)$$

where λ_1 is a constant.




Fig. 6. An overview of the Confidence Network (CN). x is ground truth image. \hat{x} is deblurred image obtained from UMSN. m semantic maps of x .

TABLE I

PSNR AND SSIM VALUES FOR EACH DEBLURRED SEMANTIC CLASS ON THE CelebA DATASET (C) AND HELLEN DATASET (H)

	Method	class-1		class-2		class-3		class-4	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
C	Shen <i>et al.</i>	27.86	0.88	30.27	0.93	33.33	0.92	31.79	0.91
	UMSN	29.27	0.91	30.75	0.95	33.83	0.96	33.10	0.94
H	Shen <i>et al.</i>	29.13	0.85	31.20	0.88	35.13	0.92	33.96	0.88
	UMSN	30.83	0.87	31.76	0.90	35.67	0.93	34.81	0.91

IV. EXPERIMENTAL RESULTS

We evaluate the results using the images provided by the authors of [11] which consists of 8000 blurry images generated using the Helen dataset [13], and 8000 blurry images generated using the CelebA dataset [14].¹ Furthermore, we test our network on a test dataset called PubFig, provided by the authors of [12] which contains 192 blurry images. We quantitatively evaluate the results using the Peak-Signal-to-Noise Ratio (PSNR) and the Structural Similarity index (SSIM) [45].

A. Implementation Details

1) *Segmentation Network*: We initially train the S-Net network using the Helen dataset [13] which contains 2000 clean images with corresponding semantic labels provided by Smith *et al.* [46]. S-Net is trained using the cross entropy loss with the Adam optimizer and the learning rate is set equal to 0.0002. We train S-Net for 60000 iterations on clean images. As can be seen from Fig. 7 and Table II, S-Net trained on clean images only is not able to produce good results on the blurry images. Thus, we fine-tune S-Net using the blurry images and the corresponding semantic maps generated using the Helen dataset [13]. In this case, S-Net is fine-tuned for 30000 iterations with the learning rate of 0.00001. We test our S-Net network on 330 clean Test images provided by the Helen dataset as well as 8000 blurry images. Results are shown in Fig. 7 and Table II in terms of F-scores. As can be seen from these results, our method is able to produce reasonable segmentation results from blurry images.

2) *UMSN Network*: Training images for UMSN are generated using 2000 images from the Helen dataset [13], and randomly selected 25000 images from the CelebA dataset [14]. We generate 25000 blur kernels sizes ranging from 13×13 to 29×29 , using 3D camera trajectories [47]. Patches of size 128×128 are extracted from those images and convolved with

¹We follow the same training and testing protocols defined by [11].

TABLE II
F-SCORE VALUES OF SEMANTIC MASKS FOR CLEAN AND BLURRY IMAGES FROM THE HELEN DATASET

Class	Clean Images	Blurry Images	
	S-Net trained on clean images	S-Net trained on clean images	S-Net fine-tuned w/ blurry images
class-1	0.910	0.889	0.901
class-2	0.903	0.877	0.884
class-3	0.856	0.775	0.814
class-4	0.625	0.538	0.566

25000 blur kernels randomly to generated about 1.7 million pairs of clean-blurry data. We added Gaussian noise with $\sigma = 0.03$ to the blurry images.

a) *First stage semantic network*: Every first stage semantic network is trained with clean-blurry paired data where only the corresponding class components of the face images are blurred. The first stage semantic network is trained with a combination of the L1-norm and the perceptual losses using the Adam optimizer with learning rate of 0.0002. The first stage network is trained for 50,000 iterations. To illustrate what kind of outputs are observed from the first stage semantic networks, in Fig. 8 we display semantic masks (first row), the corresponding blurry images (second row) and the final estimated deblurred outputs from first stage semantic networks (last row). As it can be seen from the last row, first stage semantic networks are able to remove most of the blur from corresponding class in the image. Note that, F-Net- i 's from these first stage semantic networks are used along with B-Net to construct UMSN, and trained to perform deblurring face image.

b) *Overall UMSN network*: UMSN is trained using \mathcal{L}_{total} with the Adam optimizer and batchsize of 16. The learning rate is set equal to 0.0002. Note that, semantic maps $\{m_i\}_{i=1}^4$ of the ground truth clean images in \mathcal{L}_{total} are generated using S-Net, since we do not have the actual semantic maps for them. λ and λ_1 values are set equal to 0.01 and 0.0002, respectively. UMSN is trained for 0.1 million iterations.

B. Evaluation

1) *PSNR and SSIM*: We compare the performance of our method with the following state-of-the-art algorithms: MAP-based methods [15], [16], [48]–[50], face deblurring exemplar-based method [30], and CNN-based methods [9], [11], [24]. Results are shown in Table III. As can be seen from this table, UMSN outperforms state-of-the-art methods including the methods that make use of semantic maps for image deblurring [11]. Furthermore, we evaluate UMSN's performance in reconstructing individual semantic classes against [11]. As it can be seen from the Table I, our method's performance in reconstructing individual classes is better than the state-of-the-art method [11]. In addition, we compare the performance of our method against CNN-based face deblurring method that uses post processing [12] on the PubFig dataset consisting of 192 images released by the authors of [12]. The PSNR and SSIM values achieved by [12] are 30.21 and 0.84, whereas our method achieved 30.95 and 0.87, respectively.




Fig. 7. Semantic maps generated by S-Net on blurry images from the Helen dataset. First row contains Blurry images. Second rows consists of the corresponding Semantic masks obtained from S-Net trained on the clean images. Third row consists of the corresponding Semantic masks obtained from S-Net fine tuned with the blurry images.




Fig. 8. First row: semantic masks. Second row: blurry images. Third row: deblurred images using first stage semantic network for classes (a) m_1 , (b) m_2 , (c) m_3 (d) m_4 . Note that first stage semantic network is combination of F-Net- i , ResBlock and Conv3 \times 3.

2) *VGG-Face Distance*: We compare the performance of our method with the following state-of-the-art algorithms using feature distance (i.e, L_2 -norm distance between output feature map from network) between deblurred image and the ground truth image. We use the outputs of *pool5* layer from the VGG-Face [44] to compute the VGG-Face distance (d_{VGG}). d_{VGG} is a better perceptual metric to compare as it computes the L_2 -norm distance in the feature space using VGG-Face [44] network. Table IV clearly shows that our method UMSN outperforms the existing state of the art methods using the perceptual measures like d_{VGG} .

TABLE III
PSNR AND SSIM COMPARISON OF UMSN AGAINST STATE-OF-THE-ART METHODS

Deblurring Method	Helen		CelebA	
	PSNR	SSIM	PSNR	SSIM
Krishnan <i>et al.</i> [16] (CVPR'11)	19.30	0.670	18.38	0.672
Pan <i>et al.</i> [30] (ECCV 2014)	20.93	0.727	18.59	0.677
Shan <i>et al.</i> [48] (SIGGRAPH'08)	19.57	0.670	18.43	0.644
Xu <i>et al.</i> [15] (CVPR'13)	20.11	0.711	18.93	0.685
Cho <i>et al.</i> [49] (SIGGRAPH'09)	16.82	0.574	13.03	0.445
Zhong <i>et al.</i> [50] (CVPR'13)	16.41	0.614	17.26	0.695
Nah <i>et al.</i> [9] (CVPR'17)	24.12	0.823	22.43	0.832
Shen <i>et al.</i> [11] (CVPR'18) w/GAN	25.58	0.861	24.34	0.860
Shen <i>et al.</i> [11] (CVPR'18)	25.99	0.871	25.05	0.879
Kupyn <i>et al.</i> [24] (ICCV'19)	26.45	0.880	25.42	0.884
UMSN (ours w/ \mathcal{L}_{total})	27.75	0.897	26.62	0.908

TABLE IV
COMPARISON OF UMSN AGAINST STATE-OF-THE-ART METHODS USING DISTANCE OF FEATURE FROM VGG-FACE (d_{VGG}), LOWER IS BETTER

Deblurring Method	Helen	CelebA
	d_{VGG}	d_{VGG}
Xu <i>et al.</i> [15] (CVPR'13)	10.15	11.38
Zhong <i>et al.</i> [50] (CVPR'13)	12.28	14.72
Nah <i>et al.</i> [9] (CVPR'17)	7.46	8.92
Shen <i>et al.</i> [11] (CVPR'18)	5.87	6.49
Kupyn <i>et al.</i> [24] (ICCV'19)	4.91	5.84
UMSN (ours w/ \mathcal{L}_{total})	2.69	2.85

3) *Face Recognition*: In order to show the significance of different face deblurring methods, we perform face recognition on the deblurred images. We use the CelebA dataset to perform



Fig. 9. Sample results from the Helen and CelebA datasets. As can be seen from this figure, MAP-based methods [15], [50] are adding some artifacts on the final deblurred images especially near the eyes, nose and mouth regions. Deblurring [24] without prior produced smooth output face images with artifacts, this can be seen in the fifth column corresponding to [24]. Global prior-based method [11] produces very smooth outputs and some parts of the deblurred images are deformed. This can be clearly seen by looking at the reconstructions in the fourth column corresponding to [11], where the right eye is added in the face image, the fingers are removed, the mustache is removed, and the mouth is deformed. On the other hand by using class-specific multi-stream networks and \mathcal{L}_{total} , our method is able to reconstruct all the classes of a face perfectly and produces sharper face images. For example, UMSN is able to reconstruct eyes, nose, mustache and mouth clearly. Additionally UMSN does not add parts like eyes when they are occluded (last row). Note that the method proposed in [11] uses a generative adversarial network (GAN) to reconstruct sharper images. In contrast, our method is able to reconstruct sharper images even without using a GAN.

this comparison, where we use 100 different identities as the probe set, and for each identity we select 9 additional clean face images as the gallery set. 8000 blurry images are generated using the 100 images in the probe set. Given a blurry image, the deblurred image is computed using the corresponding deblurring algorithm. The most similar face for this deblurred image is selected from the gallery set to check whether they belong to same identity or not.

To perform face recognition we use OpenFace toolbox [51] to compute the similarity distance between the deblurred image and all the gallery images. We select the Top-K nearest matches for each deblurred image to compute the accuracy

of the deblurring method. Table V clearly shows that the deblurred images by our method have better recognition accuracies than the other state-of-the-art methods. This experiment clearly shows that our method is able to retain the important parts of a face while performing deblurring. This in turn helps in achieving better face recognition compared to the other methods.

4) *Qualitative Results*: The qualitative performance of different methods on several images from the Helen and CelebA test images are shown in Fig. 9. We can clearly observe that UMSN produces sharp images without any artifacts when compared to the other methods [11], [15], [50].




Fig. 10. Sample results on real blurry images.

In addition, we conducted experiments on 100 real blurry images provided by the authors of [11], [15], [50]. Results corresponding to different methods are shown in Fig. 10. As can be seen from this figure, UMSN produces sharper and clear images when compared to the state-of-the-art methods. For example, methods [15], [50] produce results which contain artifacts or blurry images. Shen *et al.* [11], Kupyn *et al.* [24] are not able reconstruct eyes, nose and mouth as shown in the fourth and fifth columns of Fig. 10, respectively. On the other hand, eyes, nose, and mouth regions are clearly visible in the images corresponding to the UMSN method.

C. Ablation Study

We conduct experiments on two datasets to study the performance contribution of each component in UMSN. We start with the base network (B-Net) and progressively add

components to establish their significance to estimate the final deblurred image. The corresponding results are shown in Table VI. It is important to note that B-Net and UMSN are configured to have the same number of learnable parameters by setting appropriate number of output channels in the intermediate convolutional layers of the ResBlock. The ablation study demonstrates that each of the components improve the quality of the deblurred image. In particular, the introduction of the multi-stream architecture improved the performance by 0.6dB. When UMSN is trained using \mathcal{L}_{total} the performance improved by an additional 0.75dB. Finally, the combination of all the components of UMSN produce the best results with 2.5dB gain compared to B-Net.

Sample reconstructions corresponding to the ablation study using the Helen and CelebA datasets are shown in Fig. 11. As can be seen from this figure, B-Net and B-Net+semantic maps produce reconstructions that are still blurry and they




Fig. 11. Ablation study. (a) Blurry image, (b) B-Net, (c) B-Net + semantic maps, (d) B-Net + semantic maps + NRL, (e) UMSN without \mathcal{L}_c , (f) UMSN (g) ground Truth.

TABLE V
TOP-1, TOP-3 AND TOP-5 FACE RECOGNITION ACCURACIES
ON THE CelebA DATASET

Deblurring Method	Top-1	Top-3	Top-5
Clear images	71%	84%	89%
Blurred images	31%	46%	53%
Pan <i>et al.</i> [30] (ECCV 2014)	44%	57%	64%
Xu <i>et al.</i> [15] (CVPR'13)	43%	57%	64%
Zhong <i>et al.</i> [50] (CVPR'13)	30%	44%	51%
Nah <i>et al.</i> [9] (CVPR'17)	42%	59%	65%
Shen <i>et al.</i> [11] (CVPR'18)	54%	68%	74%
Kupyn <i>et al.</i> [24] (ICCV'19)	58%	70%	77%
UMSN (ours w/ \mathcal{L}_{total})	64%	75%	83%

TABLE VI
PSNR AND SSIM RESULTS OF THE ABLATION STUDY

Deblurring Method	Helen		CelebA	
	PSNR	SSIM	PSNR	SSIM
baseline (B-Net)	25.12	0.82	24.38	0.81
+ semantic maps	25.55	0.84	24.72	0.82
+ nested residuals	26.08	0.85	25.26	0.84
UMSN without \mathcal{L}_c	26.96	0.87	25.90	0.87
UMSN	27.75	0.90	26.62	0.91

fail to reconstruct the facial parts like eyes, nose and mouth well. On the other hand, B-Net+semantic maps+ NRL is able to improve the reconstruction quality but visually it is still not comparable to the state-of-the-art method. As shown in Fig. 12, NRL is able to estimate the class-specific residual feature maps, which are further used for estimating the final deblurred image. From Fig. 12 we can see what different residual feature maps are learning for different classes. In particular, the use of F-Nets helps UMSN to produce qualitatively good results, where all the parts of a face are clearly visible. Furthermore, training UMSN using \mathcal{L}_{total} results in sharper images.




Fig. 12. (a) Blurry image. (b) UMSN. (c) Intermediate residual feature map learned for the eyes and the hair. (d) Intermediate residual feature map learned for the face skin in NRL.




Fig. 13. Sample results from the PubFig dataset. (a) Blurry image. (b) Ground Truth. (c) Song *et al.* [12](ICCV'19). (d) UMSN.

We also compare our method qualitatively with post processing-based method [12] on the PubFig dataset. As can be seen from the results shown in Fig. 13, even after applying some post processing, [12] tends to produce results that are smooth. In comparison, UMSN is able to produce sharper and high-quality images without any post-processing.

V. CONCLUSION

We proposed a new method, called Uncertainty guided Multi-stream Semantic Network (UMSN), to address the




Fig. 14. Semantic maps generated by S-Net on a blurry image from the Helen dataset. First row contains Blurry images. Second rows consists of corresponding Semantic masks obtained from S-Net. Third consists of deblurred images using UMSN. Fourth row contains ground-truth images.

single image blind deblurring of face image problem that entails the use of facial semantic information. In our approach, we introduced a different way to use the output features from the sub-networks which are trained for the individual semantic class-wise deblurring. Additionally, we introduced novel techniques such as nested residual learning, and class-based confidence loss to improve the performance of UMSN. In comparison to the state-of-the-art single image blind deblurring methods, the proposed approach is able to achieve significant improvements when evaluated on three popular face datasets.

APPENDIX A UMSN NETWORK SEQUENCE

The BaseNetwork is similar to the first stage semantic network as shown in Fig. 5 with the following sequence of layers,

ResBlock(3,64)-Avgpool-ResBlock(64,64)-ResBlock(64,64)-ResBlock(64,64)-ResBlock(64,64)-ResBlock(64,64)-ResBlock(64,64)-Upsample-ResBlock(64,16)-Conv $3 \times 3(16,3)$.

F-Net- i is a sequence of five ResBlocks with dense connections,

ResBlock(3,16)-Avgpool-ResBlock(16,16)-ResBlock(16,16)-ResBlock(16,16)-ResBlock(16,8)

where ResBlock(m, n) indicates m input channels and n output channels for ResBlock.

F-Net- i 's are joined to form the first stage of UMSN, further these outputs are concatenated to the output of first layer of Base Network(B-Net) for constructing the UMSN network. Thus UMSN is constructed by combining F-Net- i 's and B-Net, where F-Net- i 's acts as first stage of UMSN, and B-Net acts as second stage of UMSN.

APPENDIX B CONFIDENCE SCORES

In this section, we show some example test outputs of UMSN (trained with \mathcal{L}_{total}) at different time instances of training, and their corresponding confidence values to show how confidence measure is helping the network. Confidence measure helps the UMSN in deblurring different parts of the face. For example, as shown in the Fig. 15 and Table VII, the outputs of UMSN after 50,000 iterations are blurry around eyes, nose, and mouth regions which is reflected in low confidence scores for C_3 when compared to the other classes. Eventually, UMSN learns to reconstruct eyes, nose and mouth regions which is reflected in high confidence scores for C_3 . This can be seen from Figure 15 where the output deblurred images are sharp with eyes, nose and mouth regions reconstructed perfectly. Table VIII and Fig. 16 show the confidence scores for different classes and outputs of UMSN for two different identities from the test dataset blurred with kernel sizes of 17×17 and 27×27 .




Fig. 15. Sample deblurring results corresponding to UMSN at different time instances of training. (a) Blurry image. (b) Trained for 3000 iterations, (c) 50,000 iterations, and (d) 0.1 million iterations.

TABLE VII

CONFIDENCE VALUES CORRESPONDING TO IMAGE1 AND IMAGE2 IN FIRST ROW AND SECOND ROW IMAGES IN FIG. 15, RESPECTIVELY

Blurry Image	Iteration	C_1	C_2	C_3	C_4
Image1	3000	0.573	0.576	0.610	0.584
	50000	0.676	0.902	0.789	0.847
	100000	0.858	0.917	0.989	0.962
Image2	3000	0.519	0.528	0.493	0.611
	50000	0.728	0.883	0.817	0.872
	100000	0.896	0.919	0.966	0.944




Fig. 16. Sample deblurring results of UMSN for different kernel sizes 17×17 and 27×27 for two different identities. (a) Blurry images with kernel size 17×17 . (b) Deblurred images using UMSN. (c) Blurry image with kernel size 27×27 . (d) Deblurred image using UMSN.

APPENDIX C
ADDITIONAL RESULTS

Extreme blur and occlusions by object like sun glasses or other body parts such as hands can cause S-Net to produce poor quality segmentations as shown in the second row of Fig. 14. Even when the S-Net fails to produce accurate segmentation masks, UMSN is able to produce visually pleasing results as shown in the third row of Fig. 14. This is mainly due to the concatenation of features from F-Net- i with the feature from the first layer of B-Net. By doing this, features extracted and reconstructed for each class in corresponding F-Net helps the UMSN to deblur the face image even if the face parsing does not work well. Note that the first stage semantic networks are initially trained on the images where

TABLE VIII

CONFIDENCE VALUES FOR THE BLURRY IMAGES GENERATED USING IDENTITY1 AND IDENTITY2 CORRESPONDING TO FIRST ROW AND SECOND ROW IMAGES IN FIG. 16, RESPECTIVELY

Clean image Identity	kernel	C_1	C_2	C_3	C_4
Identity1	17×17	0.923	0.948	0.913	0.973
	27×27	0.871	0.832	0.848	0.945
Identity2	17×17	0.874	0.931	0.942	0.945
	27×27	0.817	0.844	0.869	0.880




Fig. 17. First row corresponds to the images without occlusion and second row corresponds to the images with occlusion. (a) Blurry images. (b) Ground-truth clean images. (c) Deblurred images using the proposed UMSN network.




Fig. 18. (a) Blurry image. (b) Ground-truth clean image. (c) Deblurred images using the proposed UMSN network.

only the corresponding class is blurred and remaining are clean. Thus even if the class mask is extremely corrupted, the F-Net can extract the corresponding class features and helps the UMSN to produce visually pleasing images as shown in Fig. 14.

Figure 17 and Figure 18 show cases in which the face images are occluded and taken in low-light conditions, respectively. In such conditions, the proposed method fails to produce better quality images.

ACKNOWLEDGMENT

The views and conclusion contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon.

REFERENCES

- [1] R. Fergus, B. Singh, A. Hertzmann, S. T. Roweis, and W. T. Freeman, "Removing camera shake from a single photograph," *ACM Trans. Graph.*, vol. 25, no. 3, pp. 787–794, Jul. 2006.
- [2] C. J. Schuler, H. C. Burger, S. Harmeling, and B. Scholkopf, "A machine learning approach for non-blind image deconvolution," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2013, pp. 1067–1074.
- [3] V. M. Patel, G. R. Easley, and D. M. Healy, "Shearlet-based deconvolution," *IEEE Trans. Image Process.*, vol. 18, no. 12, pp. 2673–2685, Dec. 2009.
- [4] J. Ni, P. Turaga, V. M. Patel, and R. Chellappa, "Example-driven manifold priors for image deconvolution," *IEEE Trans. Image Process.*, vol. 20, no. 11, pp. 3086–3096, Nov. 2011.
- [5] W. Ren, X. Cao, J. Pan, X. Guo, W. Zuo, and M.-H. Yang, "Image deblurring via enhanced low-rank prior," *IEEE Trans. Image Process.*, vol. 25, no. 7, pp. 3426–3437, Jul. 2016.
- [6] L. Sun, S. Cho, J. Wang, and J. Hays, "Edge-based blur kernel estimation using patch priors," in *Proc. IEEE Int. Conf. Comput. Photography (ICCP)*, Apr. 2013, pp. 1–8.
- [7] T. M. Nimisha, A. K. Singh, and A. N. Rajagopalan, "Blur-invariant deep learning for blind-deblurring," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 4752–4760.
- [8] L. Xu, J. S. Ren, C. Liu, and J. Jia, "Deep convolutional neural network for image deconvolution," in *Proc. Adv. Neural Inf. Process. Syst.*, 2014, pp. 1790–1798.
- [9] S. Nah, T. H. Kim, and K. M. Lee, "Deep multi-scale convolutional neural network for dynamic scene deblurring," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 3883–3891.
- [10] S. Zhang, X. Shen, Z. Lin, R. Mech, J. P. Costeira, and J. M. F. Moura, "Learning to understand image blur," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 6586–6595.
- [11] Z. Shen, W.-S. Lai, T. Xu, J. Kautz, and M.-H. Yang, "Deep semantic face deblurring," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 8260–8269.
- [12] Y. Song *et al.*, "Joint face hallucination and deblurring via structure generation and detail enhancement," *Int. J. Comput. Vis.*, vol. 127, nos. 6–7, pp. 785–800, Jun. 2019.
- [13] V. Le, J. Brandt, Z. Lin, L. Bourdev, and T. S. Huang, "Interactive facial feature localization," in *Computer Vision—ECCV (Lecture Notes in Computer Science)*, vol. 7574, A. Fitzgibbon, S. Lazebnik, P. Perona, Y. Sato, and C. Schmid, Eds. Berlin, Germany: Springer, 2012, pp. 679–692.
- [14] Z. Liu, P. Luo, X. Wang, and X. Tang, "Deep learning face attributes in the wild," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 3730–3738.
- [15] L. Xu, S. Zheng, and J. Jia, "Unnatural l0 sparse representation for natural image deblurring," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2013, pp. 1107–1114.
- [16] D. Krishnan, T. Tay, and R. Fergus, "Blind deconvolution using a normalized sparsity measure," in *Proc. CVPR*, Jun. 2011, pp. 233–240.
- [17] J. Pan, D. Sun, H. Pfister, and M.-H. Yang, "Blind image deblurring using dark channel prior," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 1628–1636.
- [18] J. Kruse, C. Rother, and U. Schmidt, "Learning to push the limits of efficient FFT-based image deconvolution," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 4586–4594.
- [19] U. Schmidt, C. Rother, S. Nowozin, J. Jancsary, and S. Roth, "Discriminative non-blind deblurring," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2013, pp. 604–611.
- [20] U. Schmidt and S. Roth, "Shrinkage fields for effective image restoration," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 2774–2781.
- [21] S. Vasu, V. R. Maligireddy, and A. N. Rajagopalan, "Non-blind deblurring: Handling kernel uncertainty with CNNs," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 3272–3281.
- [22] C. J. Schuler, M. Hirsch, S. Harmeling, and B. Scholkopf, "Learning to deblur," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 38, no. 7, pp. 1439–1451, Jul. 2016.
- [23] A. Chakrabarti, "A neural approach to blind motion deblurring," in *Proc. Eur. Conf. Comput. Vis.* Springer, 2016, pp. 221–235.
- [24] O. Kupyn, T. Martyniuk, J. Wu, and Z. Wang, "DeblurGAN-v2: Deblurring (Orders-of-Magnitude) faster and better," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 8878–8887.
- [25] Y.-H. Tsai, X. Shen, Z. Lin, K. Sunkavalli, and M.-H. Yang, "Sky is not the limit: Semantic-aware sky replacement," *ACM Trans. Graph.*, vol. 35, no. 4, pp. 1–11, Jul. 2016.
- [26] D. Eigen, D. Krishnan, and R. Fergus, "Restoring an image taken through a window covered with dirt or rain," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 633–640.
- [27] P. Svoboda, M. Hradis, L. Marsik, and P. Zemcik, "CNN for license plate motion deblurring," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2016, pp. 3832–3836.
- [28] J. Pan, Z. Hu, Z. Su, and M.-H. Yang, "Deblurring text images via L₀-regularized intensity and gradient prior," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2014, pp. 2901–2908.
- [29] S. Baker and T. Kanade, "Limits on super-resolution and how to break them," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 24, no. 9, pp. 1167–1183, Sep. 2002.
- [30] J. Pan, Z. Hu, Z. Su, and M.-H. Yang, "Deblurring face images with exemplars," in *Proc. Eur. Conf. Comput. Vis.*, 2014, pp. 47–62.
- [31] N. Joshi, W. Matusik, E. H. Adelson, and D. J. Kriegman, "Personal photo enhancement using example images," *ACM Trans. Graph.*, vol. 29, no. 2, pp. 1–15, Mar. 2010.
- [32] A. Bulat and G. Tzimiropoulos, "Super-FAN: Integrated facial landmark localization and super-resolution of real-world low resolution faces in arbitrary poses with GANs," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 109–117.
- [33] D. Chen *et al.*, "Gated context aggregation network for image dehazing and deraining," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Jan. 2019, pp. 1375–1383.
- [34] X. Yu, B. Fernando, B. Ghanem, F. Porikli, and R. Hartley, "Face super-resolution guided by facial component heatmaps," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 217–233.
- [35] W. Ren, J. Yang, S. Deng, D. Wipf, X. Cao, and X. Tong, "Face video deblurring using 3D facial priors," in *Proc. IEEE/CVF Int. Conf. Comput. Vis. (ICCV)*, Oct. 2019, pp. 9388–9397.
- [36] B. Lu, J.-C. Chen, and R. Chellappa, "Unsupervised domain-specific deblurring via disentangled representations," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, p. 10.
- [37] D. Ulyanov, A. Vedaldi, and V. Lempitsky, "Instance normalization: The missing ingredient for fast stylization," 2016, *arXiv:1607.08022*. [Online]. Available: <http://arxiv.org/abs/1607.08022>
- [38] O. Ronneberger, P. Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Proc. Med. Image Comput. Comput. Assist. Intervent. (MICCAI)*, vol. 351, 2015, p. 241. [Online]. Available: <http://lmb.informatik.uni-freiburg.de/Publications/2015/RFB15a>
- [39] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 4700–4708.
- [40] Z. Wang and S. Ji, "Smoothed dilated convolutions for improved dense prediction," in *Proc. 24th ACM SIGKDD Int. Conf. Knowl. Discovery Data Mining*, Jul. 2018, pp. 2486–2495.
- [41] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Proc. ECCV*, 2016, pp. 694–711.
- [42] H. Zhang and K. Dana, "Multi-style generative network for real-time transfer," 2017, *arXiv:1703.06953*. [Online]. Available: <http://arxiv.org/abs/1703.06953>
- [43] C. Ledig *et al.*, "Photo-realistic single image super-resolution using a generative adversarial network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 4681–4690.
- [44] O. M. Parkhi, A. Vedaldi, and A. Zisserman, "Deep face recognition," in *Proc. Brit. Mach. Vis. Conf.*, 2015, pp. 1–10.
- [45] Z. Wang, A. C. Bovik, H. R. Sheikh, and E. P. Simoncelli, "Image quality assessment: From error visibility to structural similarity," *IEEE Trans. Image Process.*, vol. 13, no. 4, pp. 600–612, Apr. 2004.
- [46] B. M. Smith, L. Zhang, J. Brandt, Z. Lin, and J. Yang, "Exemplar-based face parsing," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2013, pp. 3484–3491.
- [47] G. Boracchi and A. Foi, "Modeling the performance of image restoration from motion blur," *IEEE Trans. Image Process.*, vol. 21, no. 8, pp. 3502–3517, Aug. 2012.
- [48] Q. Shan, J. Jia, and A. Agarwala, "High-quality motion deblurring from a single image," *ACM Trans. Graph.*, vol. 27, no. 3, pp. 1–10, Aug. 2008.
- [49] S. Cho and S. Lee, "Fast motion deblurring," *ACM Trans. Graph.*, vol. 28, no. 5, p. 145, 2009.
- [50] L. Zhong, S. Cho, D. Metaxas, S. Paris, and J. Wang, "Handling noise in single image deblurring using directional filters," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit.*, Jun. 2013, pp. 612–619.
- [51] B. Amos, B. Ludwiczuk, and M. Satyanarayanan, "Openface: A general-purpose face recognition library with mobile applications," CMU School of Computer Science, Pittsburgh PA, USA, Tech. Rep. CMU-CS-16-118, 2016.



Rajeev Yasarla (Graduate Student Member, IEEE) received the bachelor's and master's degrees in electrical engineering from IIT Madras. He is currently pursuing the Ph.D. degree with the Department of Electrical and Computer Engineering (ECE), Johns Hopkins University. His research interests include deep learning-based image restoration (like de-raining, deblurring, and atmospheric turbulence distortion removal), object detection, and medical image segmentation.



Federico Perazzi (Member, IEEE) received the degree in computer science in 2008, the degree in electrical engineering in 2008, the M.Sc. degree in entertainment technology from Carnegie Mellon University in 2010, and the M.Sc. degree in computer science and the Ph.D. degree from the Swiss Federal Institute of Technology (ETHZ). He is a Research Scientist with Adobe Inc. Prior to that, he was a Postdoctoral Researcher at Disney Research Zurich, Switzerland.



Vishal M. Patel (Senior Member, IEEE) received the Ph.D. degree in electrical engineering from the University of Maryland, College Park, MD, USA, in 2010. He is an Assistant Professor with the Department of Electrical and Computer Engineering (ECE), Johns Hopkins University. Prior to joining Hopkins, he was an A. Walter Tyson Assistant Professor with the Department of ECE, Rutgers University, and a member of the research faculty at the University of Maryland Institute for Advanced Computer Studies (UMIACS). His current research interests include signal processing, computer vision, and pattern recognition with applications in biometrics and imaging. He is a member of Eta Kappa Nu, Pi Mu Epsilon, and Phi Beta Kappa. He has received a number of awards, including the A. Walter Tyson Assistant Professorship Award, the Best Paper Award from the IEEE BTAS in 2015, the Best Poster Awards from BTAS in 2015 and 2016, the 2016 ONR Young Investigator Award, the 2016 Jimmy Lin Award for Invention, the Best Paper Award from the IEEE AVSS in 2017 and 2019, the Honorable Mention Paper Award from IAPR ICB in 2018, and two Best Student Paper Awards from IAPR ICPR in 2018. He is an Associate Editor of the *IEEE Signal Processing Magazine* and the IEEE BIOMETRICS COMPENDIUM, and serves on the Machine Learning for Signal Processing (MLSP) Committee of the IEEE Signal Processing Society.