

Confidence Measure Guided Single Image De-Raining

Rajeev Yasarla¹, *Student Member, IEEE*, and Vishal M. Patel¹, *Senior Member, IEEE*

Abstract—Single image de-raining is an extremely challenging problem since the rainy images contain rain streaks which often vary in size, direction and density. This varying characteristic of rain streaks affect different parts of the image differently. Previous approaches have attempted to address this problem by leveraging some prior information to remove rain streaks from a single image. One of the major limitations of these approaches is that they do not consider the location information of rain drops in the image. We extend our previous work UMRL network, and propose Image Quality-based single image Deraining using Confidence measure (QuDeC), network addresses this issue by learning the quality or distortion level of each patch in the rainy image, and further processes this information to learn the rain content at different scales. In addition, we introduce a technique which guides the network to learn the network weights based on the confidence measure about the estimate of both quality at each location and residual rain streak information (residual map). Extensive experiments on synthetic and real datasets demonstrate that the proposed method achieves significant improvements over the recent state-of-the-art methods.

Index Terms—Deraining, convolutional neural networks, image restoration.

I. INTRODUCTION

MANY practical computer vision-based systems such as surveillance and autonomous driving often require processing and analysis of videos and images captured under adverse weather conditions such as rain, snow, haze etc. These weather-based conditions adversely affect the visual quality of images and as a result often degrade the performance of vision systems. Hence, it is important to develop algorithms that can automatically remove these artifacts before they are fed to a vision-based system for further processing.

In this paper, we address the problem of removing rain streaks from a single rainy image. Rain streak removal or image de-raining is a difficult problem since a rainy image may contain rain streaks which may vary in size, direction and density. A number of different techniques have been developed in the literature to address this problem. These algorithms can be clustered into two main groups - (i) video based algorithms [5]–[11], and (ii) single image-based algorithms [2]–[4], [12], [13]. Algorithms

Manuscript received September 9, 2019; revised January 6, 2020; accepted February 1, 2020. Date of current version February 24, 2020. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. Zhen He. (*Corresponding author: Rajeev Yasarla.*)

The authors are with the Whiting School of Engineering, Johns Hopkins University, Baltimore, MD 21218-2608 USA (e-mail: ryasar11@jhu.edu; vpatel36@jhu.edu).

Digital Object Identifier 10.1109/TIP.2020.2973802

1057-7149 © 2020 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission.
See <https://www.ieee.org/publications/rights/index.html> for more information.

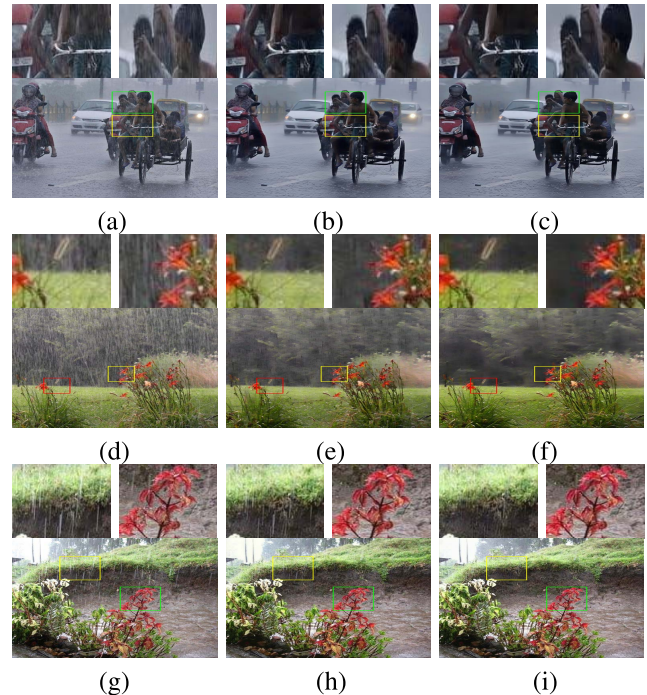


Fig. 1. Sample image deraining results. (a),(d),(g) are Rainy images. (b) De-rained using DID-MDN [2] where zoomed in parts show the rain streaks and blurry effects on faces and hands. (e) De-rained using Fu *et al.* [3] where zoomed in parts show under de-raining of the image. (h) De-rained using JORDER [4] where zoomed in part shows under de-raining near grass and blurry effects on the plant. (c),(f),(i) are de-rained images using QuDeC.

corresponding to the first category assume temporal consistency among the image frames, and use this assumption for de-raining. On the other hand, single image de-raining methods attempt to use some prior information to remove rain components from a single image [2], [12]–[14]. Priors such as sparsity [15], [16] and low-rank representation [17] have been used in the literature. In particular, the method proposed by Fu *et al.* [18] uses a priori image domain knowledge by focusing on high frequency details during training to improve the de-raining performance. However, it was shown in [2], that this method tends to remove some important parts in the de-rained image (see Figure 1(e)). Similarly, a recent work by Zhang and Patel [2] uses the image-level priors to estimate the rain density information which is then used for de-raining. Although their approach provides the state-of-the-art results, they estimate image level priors which do not consider the location information of rain drops in the image. As a result, their algorithm tends to introduce some artifacts in the final




Fig. 2. Sample location-quality-label-maps s corresponding to synthetic and real-world rainy images. Here *red color* indicates high distortion, *blue color* indicates medium distortion, *low color* indicates low distortion in the corresponding patch of the rainy image, y .

de-rained images. These artifacts can be clearly seen from the de-rained results shown in Figure 1(b).

In our previous work [1] we proposed a different approach to image de-raining where we made use of the observation that rain streak density and direction does not change drastically with different scales. Rather than relying on the rain density information (i.e. heavy, medium or light) alone, we developed a method in which the rain streak location information is also taken into consideration in a multi-scale fashion to improve the de-raining performance. In this paper, we extend our work in [1] by also incorporating the distortion level information at each location in the lower scales of the image to further improve the de-raining performance. To this end, we use the Natural Image Quality Evaluator (NIQE) no-reference image quality score [19] and formulate a joint task of estimating the level of distortion at each spatial location of the image and computing rain streak information to reconstruct the clean image. This is achieved by adding a decoder network which computes the level of distortion at each location, to our UMRL approach in [1]. This additional decoder network labels the distortion at each location into three classes and also computes the confidence scores for each patch (which indicates how confident the decoder is about those obtained labels) as shown in the Figure 3. As a result, this can be interpreted as soft labelling of the distortion levels in the image. As shown in Figure 3 different patches in image are labelled with different confident scores. This makes the proposed QuDeC network more robust by blocking the errors in level of distortion while estimating the de-rained image. In addition, we propose a novel loss function to train the network. In summary, we add the following novelties to our previous UMRL approach in [1].

- We formulate a joint task of estimating the level of distortion at each spatial location of the image and compute rain streak information by adding a new decoder to UMRL [1].
- A new loss function is proposed for estimating the confidence scores for both residual maps and distortion quality-label maps.

Fig. 1(c), (f) and (i) present sample results from our QuDeC network, where one can clearly see that QuDeC is able to remove the noise artifacts and provides better results as compared to [3], [4] and [2]. To summarize, the following are our key contributions in this paper.

- A novel method called QuDeC is proposed which jointly computes the level of distortion at each location and




Fig. 3. (a) Input rainy image, y . (b) location-quality-label-map. (c) corresponding confidence scores.

estimates rain streak information in multi-scale fashion to obtain the de-rained image.

- Confidence scores are computed for each task and judiciously combined to obtain improved results.
- We run extensive experiments to show the performance of QuDeC against several recent state-of-the-art approaches on both synthetic and real rainy images. Furthermore, an ablation study is conducted to demonstrate the effectiveness of different parts of the proposed QuDeC network.

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. Training details are provided in Section IV. Experimental results are presented in Section V, and finally, Section VI concludes the paper with a brief summary.

II. BACKGROUND AND RELATED WORK

An observed rainy image y can be modeled as the superposition of a rain component (i.e. residual map) r with a clean image x as follows

$$y = x + r. \quad (1)$$

Given y the goal of image de-raining is to estimate x . This can be done either by estimating the residual map r and then subtracting it from the observed image y or by directly estimating x from y . Various methods have been proposed in the literature for image de-raining [12], [14], [16], [20]–[23], [23]–[25] including dictionary learning-based [26], Gaussian mixture-model (GMM) based [27], and low-rank representation based [28] methods. In recent years, deep learning-based single image de-raining methods have also been proposed in the literature. Fu *et al.* [18] proposed a convolutional neural network (CNN) based approach in which they directly learn the mapping relationship between rainy and clean image detail layers from data. Zhang *et al.* [29] proposed a generative

adversarial network (GAN) based method for image de-raining. Furthermore, to minimize the artifacts introduced by GANs and ensure better visual quality, a new refined loss function was also introduced in [29]. Fu *et al.* [3] presented an end-to-end deep learning framework for removing rain from individual images using a deep detail network which directly reduces the mapping range from input to output. Zhang and Patel [2] proposed a density-aware multi-stream densely connected CNN for joint rain density estimation and de-raining. Their network automatically determines the rain-density information and then efficiently removes the corresponding rain-streaks using the estimated rain-density label. Note the methods proposed in [3], and [2] showed the benefits of using multi-scale networks for image de-raining. Recently, Wang *et al.* [30] proposed a hierarchical approach based on estimating different frequency details of an image to get the de-rained image. The method proposed by Qian *et al.* [31] generates attentive maps using the recurrent neural networks, and then uses the features from different scales to compute the loss for removing the rain drops on glasses. Note that this method was specifically designed for removing rain drops from a glass rather than removing rain streaks from an image. [32]–[34] illustrated the importance of attention based methods in low-level vision tasks. In a recent work [34], Li *et al.* proposed a convolutional and recurrent neural network-based method for single image de-raining which makes use of the contextual information for rain removal. Wei *et al.* [35] proposed a semisupervised single image deraining method using Gaussian mixture models. Yang *et al.* [36] presented an excellent literature and discussed about the various aspects like architecture and losses that beneficial to single image deraining. Hu *et al.* [37] an end-to-end deep neural network that computes a residual map to estimate the rain-free image output using depth-attentional features computed with help of depth-guided attention mechanism. It was observed in [2], that some of the recent deep learning-based methods tend to under de-rain or over de-rain the image if the rain condition present in the rainy image is not properly considered during training. References [38]–[40] show the benefits of modeling the uncertainty and use it for training the network. Inspired by their work we propose a novel confidence based loss to train our network QuDeC.

III. PROPOSED METHOD

Unlike many deep learning-based methods that directly estimate the de-rained image from the noisy observation, we take a different approach where we formulate the rain removal problem as the joint task of computing distortion level at each location and estimating the residual rain streak component \hat{r} (i.e residual map). We compute the distortion level for the patches of size 128×128 , and classify them into three classes, *green – low distortion*, *blue – medium distortion*, and *red – high distortion*. Note that, the ground truth labels for these patches are computed using the NIQE measure [19]. In this way, the distortion level at each location in the rainy image is represented by the location-quality-label-map s . Representing the distortion level at each location, and giving it as a prior information while reconstructing the image helps the

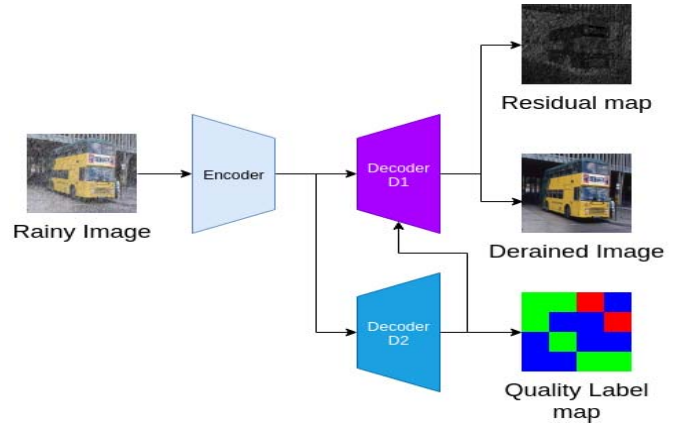


Fig. 4. An overview of the QuDeC network.

network to learn the rain streak information which may vary spatially or affect the background differently. From Figure 2 we can clearly observe that different levels of rain streak affect the background differently. In some cases the same level of rain streaks have different affects on the background. For example, as can be seen from Figure 2, rain streaks produce different distortion in bright and heavily textured regions. On the other hand the distortion level is relatively high on the dark and homogeneous regions. To handle these regions differently, we propose a method that jointly estimate the level of distortion at each location and uses it as a prior information in estimating the rain streak information (i.e residual map). In contrast to other methods [2], [4] that use either rain streak binary masks or level of rain streak present in the image, we estimate the distortion caused by rain streaks at each location based on the NIQE measure which makes our method more general and applicable for synthetic and well as real-world images.

To perform the joint task of computing the location quality label maps (s), and estimating the residual maps \hat{r} , QuDeC is constructed using one encoder and two decoders as shown in Figure 4. A rainy image is first passed through an encoder network to obtain intermediate features. These features are then further processed by the decoders. In particular, Decoder D2 estimates the location quality labels, which are then given as a prior information to Decoder D1 to estimate the residual maps (r).

A. QuDeC Network

Given a rainy image y , we estimate the residual map and its corresponding confidence maps at three different scales, $\{\hat{r}_{@1}, c_{@1}\}$ (at the original input size), $\{\hat{r}_{@2}, c_{@2}\}$ (at 0.5 scale of the input size), and $\{\hat{r}_{@4}, c_{@4}\}$ (at 0.25 scale of the input size) in a multi-scale fashion [1], as well as the location-quality-label-map \hat{s} along with the corresponding confidence scores (c_s) for label map s .

The QuDeC network aims to judiciously combine the pixel-wise rain residual information from lower scales, and patch-wise distortion level in the quality label maps while estimating the final de-rained image. To achieve this we compute the confidence maps for the estimated residual maps and

TABLE I

DETAILS OF THE QuDeC NETWORK. AVGPOOL IS THE AVERAGE POOLING LAYER, UPSAMPLE IS THE UPSAMPLING CONVOLUTION LAYER, AND CONVBlock(i, j) INDICATES CONVBlock WITH i INPUT CHANNELS AND j OUTPUT CHANNELS

Network	Network blocks
Encoder	ConvBlock(3,32)-AvgPool-ConvBlock(32,32)-AvgPool-ConvBlock(32,32)-AvgPool-ConvBlock(32,32)-AvgPool
Decoder D1	ConvBlock(32,32)-UpSample-ConvBlock(64,32)-UpSample-ConvBlock(67,32)-UpSample-ConvBlock(67,16)-ConvBlock(16,16)-Conv2d(3×3)
Decoder D2	ConvBlock(32,32)-UpSample-ConvBlock(64,32)-UpSample-ConvBlock(67,32)-UpSample-ConvBlock(67,16)-ConvBlock(16,16)-GlobalAveragePool-FullyConnectedLayer.

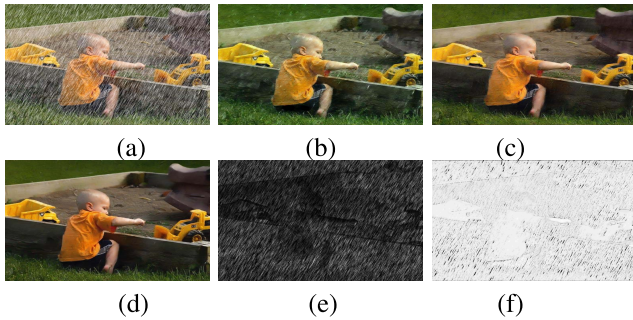


Fig. 5. (a) Input rainy image, y . (b) De-rained image using the base network. (c) De-rained using [2]. (d) De-rained using the proposed encoder and decoder D1 of the QuDeC network. (e) The residual map. (f) The confidence map at scale 1.0(@1).

location-quality-label-maps. These confidence maps block the errors in the computation by giving low confidence values to erroneous regions in the computations of \hat{r} and \hat{s} . Additionally, these confidence scores enable QuDeC to perform soft labeling of the distortion levels in the location-quality-label-map s , which makes QuDeC more robust to location quality label errors and effectively use this information as a prior while estimating \hat{r} . To compute \hat{r} and \hat{s} we start with UMRL [1] and add Decoder D2 to obtain the network architecture of QuDeC as shown in Figure 6.

1) *Encoder and Decoder D1*: Rain streaks are high frequency components and existing de-raining methods either tend to remove high frequencies that are not rain streaks or do not remove the rain near high frequency components of the clean image like edges as shown in the Figure 5. To address this issue, one can use the information about the location in the image where the network might go wrong in estimating the residual value. This can be done by estimating a confidence value corresponding to the estimated residual value and guide the network to remove the artifacts, especially near the edges. For example, we can observe clearly from Figure 5 that the residual map and its corresponding confidence map are able to capture the regions where there is high probability of incorrect estimates. In the encoder and decoder D1 networks we estimate the residual values and their corresponding confidence maps at different scales (1.0(@1), 0.5(@2) and 0.25(@4)) of the input size. This

information is then fed back to the subsequent layers so that the network can learn the residual value at each location, given the computed residual value and confidence value at lower scales.

The encoder and decoder D1 networks are similar to the the encoder and decoder networks of UMRL [1] where a convolutional block (ConvBlock as shown in Figure 7(a)) is used as the building block. The encoder network is described in Table I. The decoder D1 network contains the sequence of blocks as shown in Table I, where Residual-Confidence Network (ReCoN) is added to Decoder D1 to estimate \hat{r} at different scales and their corresponding confidence maps c . Given the feature maps as input to ReCoN network, RN (residual network) estimates the residual map \hat{r} and CN (confidence network) computes the corresponding confidence map c as shown in Figure 7(d).

Feature maps at different scales such as @2 and @4 are given as inputs to the Residual Network (RN) to estimate the residual map at the corresponding scale as shown in Figure 7(d). RN consists of the following sequence of convolutional layers, Convblock(64,32)-Convblock(32,32)-Convblock(32,3) as shown in Figure 7(b). We use the estimated residual map and the feature maps as inputs to the Confidence map Network (CN) to compute the confidence measure at every pixel, which indicates how sure the network is about the residual value at each pixel. CN consists of the following sequence of convolutional layers, Convblock(67,16)-Convblock(16,16)-Convblock(16,3) as shown in Figure 7(c). Given the estimated residual map and the corresponding feature maps as inputs to the confidence map network, it estimates $c_{@4}$ and $c_{@2}$. The element wise product of \hat{r}_i and c_i is computed, and up-sampled. This is used as an input to the subsequent layers of the Decoder D1 network in QuDeC as shown in Figure 6 for $i \in \{@2, @4\}$. Given the output residual map $r_{@1}$ and the feature maps of the final layer of the Decoder D1 network in QuDeC as input to CN, we get $c_{@1}$.

A Refinement Network (RFN) is used at the end of Decoder D1 to produce the de-rained image. It takes $y_i - \hat{r}_i$ as the input and generates \hat{x}_i (i.e. derained image) as the output. The RFN consists of the following blocks Conv2d(7×7)-Conv2d(3×3)-tanh(), where Conv2d($m \times m$) represents 2D convolution using the kernel of size $m \times m$.

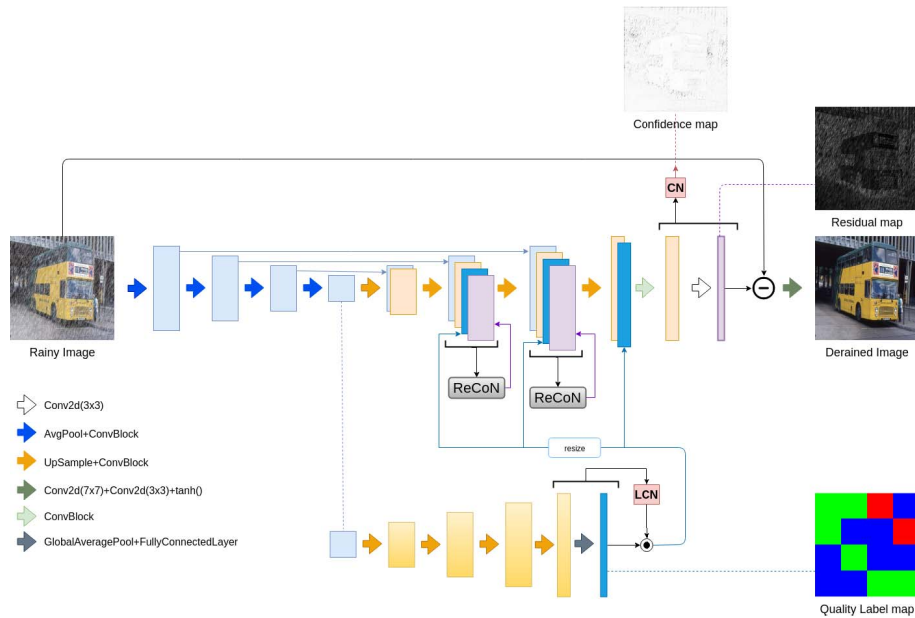


Fig. 6. An overview of the proposed QuDeC network. The aim of the QuDeC network is to estimate the clean image given the corresponding rainy image. QuDeC learns the residual maps, location-quality-label-maps, and computes the confidence maps to guide the network. To achieve this, we introduce ReCoN and LCN networks, and feed their outputs to the subsequent layers of Decoder D1.

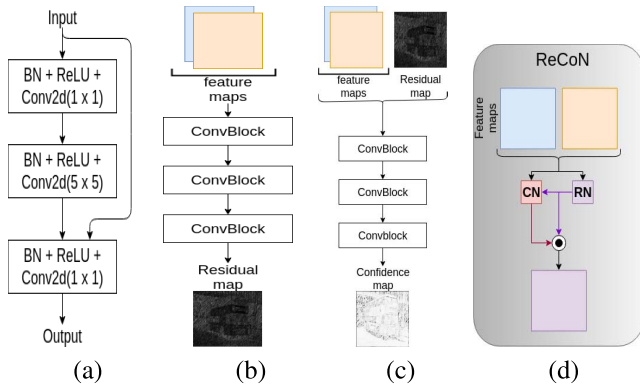


Fig. 7. (a) Convolutional block (ConvBlock). BN - batchnormalization, ReLU - Rectified Linear Units, Conv2d($m \times m$) - convolutional layer with kernel of size $m \times m$. (b) Residual Network (RN). (c) Confidence map Network (CN). (d) Residual-Confidence Network(ReCoN).




Fig. 8. (a) Input rainy image. (b) De-rained image using DIDMDN [2]. (c) De-rained using UMRL [1]. (d) De-rained using our proposed method QuDeC. (e) Location-quality-label-map s computed by QuDeC. (f) Ground truth clean image.

2) *Decoder D2*: Rain streaks have different distortion effects on different parts of the image.

As a result while reconstructing the clean image, we also estimate the prior information such as distortion caused by rain streak at every location. Not using this prior information may lead to inferior de-raining performance. As can be seen from Figure 8 that DIDMDN [2] and UMRL [1] do not perform well as they lack prior distortion level information at each location of the image. We address this by introducing a decoder D2 in our method and formulate joint task of computing the distortion level at each location and estimating the residual rain streak information. Decoder D2 is similar to Decoder D1 with the sequence of blocks shown in the Table I.

Decoder D2 takes feature maps obtained from the encoder as shown in Figure 6, to estimate the distortion levels in location-quality-label-map. Label Confidence Network (LCN)




Fig. 9. Label Confidence Network (LCN).

is used to compute the confidence scores corresponding to the location-quality-label-maps \hat{s} . Feature maps from the last layer of D2 and obtained location-quality-label-maps \hat{s} are fed to LCN to compute the confidence scores c_s as shown in Figure 9. LCN is a sequence of three ConvBlocks followed by global average pool and fully connected layers as shown in Figure 9.

B. Loss for QuDeC Network

In image restoration tasks, maximum-a-posteriori method is often used to optimize the network parameters (θ) as follows,

$$\hat{\theta} = \operatorname{argmax}_{\theta} P(f_{\theta}(y)|y; \theta) = \operatorname{argmax}_{\theta} P(\hat{x}|y; \theta), \quad (2)$$

where $P(\cdot)$ is the probability function and $f_{\theta}(\cdot)$ represents the QuDeC network, $\hat{x} = f_{\theta}(y)$. Since QuDeC performs the joint task of estimating \hat{r} and \hat{s} , the above optimization can be updated as follows,

$$\hat{\theta} = \operatorname{argmax}_{\theta} P(\hat{x}, \hat{s}|y; \theta) = \operatorname{argmax}_{\theta} P(\hat{x}|y, \hat{s}; \theta)P(\hat{s}|y; \theta).$$

To find the optimal network parameters $\hat{\theta}$, $P(\hat{x}, \hat{s}|y; \theta)$ needs to be maximized. For simplicity to solve this optimization problem, let us assume $P(\hat{x}|y, \hat{s}; \theta)$ and $P(\hat{s}|y; \theta)$ are Gaussian distributions. As our goal is to minimize the joint errors between \hat{x} and the ground-truth clean image (x), \hat{s} and the ground-truth labels (s) of rainy image y , we denote the mean of distribution $P(\hat{x}|y, \hat{s}; \theta)$ as x and variance as σ_x^2 , and mean of distribution $P(\hat{s}|y; \theta)$ as s and variance as σ_s^2 . Thus our objective becomes,

$$\begin{aligned} \hat{\theta} &= \operatorname{argmax}_{\theta} \log(P(\hat{x}|y, \hat{s}; \theta)P(\hat{s}|y; \theta)) \\ \hat{\theta} &= \operatorname{argmax}_{\theta} \log(P(\hat{x}|y, \hat{s}; \theta)) + \log(P(\hat{s}|y; \theta)) \\ \hat{\theta} &= \operatorname{argmax}_{\theta} -\frac{1}{\sigma_x^2} \|\hat{x} - x\|_2^2 - \frac{1}{\sigma_s^2} \|\hat{s} - s\|_2^2 + \log\left(\frac{1}{\sigma_x^2 \sigma_s^2}\right). \end{aligned} \quad (3)$$

Substituting, $\mathcal{L}_r = \frac{1}{\sigma_x^2} \|\hat{x} - x\|_2^2$ and $\mathcal{L}_s = \frac{1}{\sigma_s^2} \|\hat{s} - s\|_2^2$ into the above equation, we get

$$\begin{aligned} \hat{\theta} &= \operatorname{argmax}_{\theta} -\mathcal{L}_r - \mathcal{L}_s + \log\left(\frac{1}{\sigma_x^2}\right) + \log\left(\frac{1}{\sigma_s^2}\right) \\ \hat{\theta} &= \operatorname{argmax}_{\theta} -\mathcal{L}_r + \log\left(\frac{1}{\sigma_x^2}\right) - \left(\mathcal{L}_s - \log\left(\frac{1}{\sigma_s^2}\right)\right) \\ \hat{\theta} &= \operatorname{argmax}_{\theta} -\mathcal{L}_r + \mathcal{L}_c - \mathcal{L}_{cs}, \end{aligned} \quad (4)$$

where $\mathcal{L}_c = \log\left(\frac{1}{\sigma_x^2}\right)$, and $\mathcal{L}_{cs} = \mathcal{L}_s - \log\left(\frac{1}{\sigma_s^2}\right)$. In (4) variance (σ^2) and σ_s^2 can be inferred in two ways as explained in [38], [39]. (i) Epistemic uncertainty, which is explained as the model uncertainty given enough data to train, and (ii) Aleatoric uncertainty that captures noise inherent in the observations, which is data dependent. Epistemic uncertainty can be formulated as variational inference to compute variance. Aleatoric uncertainty can be formulated as MAP (maximum-a-posterior) or ML (maximum-likelihood) inference. Here, in our method we attempt to address the uncertainty caused in the outputs due to different properties of rain streaks like density, direction, and effect on background scene which are inherent in rainy images. Following the ML inference, in the above equation (4), we can view the terms $\frac{1}{\sigma_x^2}$ and $\frac{1}{\sigma_s^2}$ as the corresponding location-based confidence maps or scores. These confidence maps indicate the erroneous regions in the estimates of residual maps \hat{r} or location-quality-label-maps \hat{s} by giving low values to those regions or pixels. These errors occur at regions or pixels where the variance is high. Note that in our method, the confidence map has no ground-truths.

We compute these confidence scores using CN (Confidence Network) and LCN (Label Confidence Network) as explained in the earlier sections. Note that the values in the confidence map at every position will be in the range of [0, 1]. Additionally the L2-norm in the loss of \mathcal{L}_s in (4) for classifying the distortion levels in location-quality-map s , should be replaced with the cross entropy loss. The residual maps are estimated in a multi-scale fashion. Thus the loss is constructed as follows,

$$\begin{aligned} \mathcal{L}_r &= \sum_{i \in \{\text{@1, @2, @4}\}} \|(c_i \odot \hat{x}_i) - (c_i \odot x_i)\|_2^2, \\ \mathcal{L}_c &= \sum_{i \in \{\text{@1, @2, @4}\}} \left(\sum_j \sum_k \log(c_{ijk}) \right), \\ \mathcal{L}_{cs} &= \sum_{\forall p \in y} c_s(p) \mathcal{L}_{CE}(\hat{s}(p), s(p)) - \lambda_{cs} \log(c_s(p)) \\ \mathcal{L}_u &= \mathcal{L}_r + \mathcal{L}_{cs} - \lambda_1 \mathcal{L}_c, \end{aligned} \quad (5)$$

where

$$\mathcal{L}_{CE}(\hat{s}(p), s(p)) = -s(p) \log(\hat{s}(p)) - (1 - s(p)) \log(1 - \hat{s}(p)),$$

\odot denotes point-wise multiplication, and $\{p\}$ is a set of patches in the rainy image y . Here, λ_{cs} and λ_1 are parameters. Inspired by the importance of the perceptual loss in many image restoration tasks [15], [41], we also use it to further improve the visual quality of the de-rained images. Features from layer *relu1_2* of a pretrained network VGG-16 [42] are used to compute the perceptual loss [41], [43]. Let $F(\cdot)$ denote the features obtained using the VGG16 model [42], then the perceptual loss is defined as follows

$$\mathcal{L}_p = \frac{1}{NHW} \sum_i \sum_j \sum_k \|F(\hat{x}_1)^{i,j,k} - F(x_1)^{i,j,k}\|_2^2, \quad (6)$$

where N is the number of channels of $F(\cdot)$, H is the height and W is the width of feature maps. The overall loss used to train the QuDeC network is,

$$\mathcal{L} = \mathcal{L}_r + \mathcal{L}_{cs} - \lambda_1 \mathcal{L}_c + \lambda_2 \mathcal{L}_p, \quad (7)$$

where λ_2 is a parameter.

IV. QUDEC TRAINING

The QuDeC network is trained using the synthetic image datasets created by the authors of [2], [4], [29]. The dataset in [2] consists of 12000 images with different rain levels like low, medium and high. The dataset in [29] contains 700 training images. The dataset in [4] contains 1800 rainy images for training. We generate the ground truth location-quality-maps, s which indicate the distortion levels of the background in rainy images, using the NIQE scores. Figure 10 show the histogram of NIQE scores corresponding to the 128×128 patches of rainy images from the DIDMDN training dataset [2]. We divide the patches into three levels of distortions using the thresholds $T_1 = 6$ and $T_2 = 9$. These thresholds are chosen such that the patches are divided into three equal groups. The following pseudo code summarizes the procedure used for generating the location-quality-label-maps using the NIQE scores.




Fig. 10. Histogram of NIQE values corresponding to the 128×128 patches from the DIDMDN training images.

Input: Rainy image y , patch in y of size 128×128 .

Output: Location-quality-label-map, s^\dagger .

```

for  $\forall p \in y$  do
  if  $\text{NIQE}(y(p)) \leq T_1 - 0.2$  then
     $s^\dagger(p) = \text{green}$ 
  else if  $T_1 + 0.2 < \text{NIQE}(y(p)) < T_2 - 0.2$  then
     $s^\dagger(p) = \text{blue}$ 
  else if  $\text{NIQE}(y(p)) \geq T_2 + 0.2$  then
     $s^\dagger(p) = \text{red}$ 
  end if
end for

```

Note that the NIQE scores [19] are obtained using various features which may not be specifically useful for computing the distortion levels caused by the rain streaks. Furthermore, the thresholds T_1 and T_2 may vary from one dataset to another. In order to deal with these issues, we train a Generation-of-Labels Network (GLN) to automatically generate the location-quality-label-maps s from the input rainy images. The GLN network is trained using pairs of (y, s^\dagger) from the DIDMDN training images. We provide the network architecture and training details of GLN in the appendix. The estimated s from GLN as well as the clean image x along with the rainy image y are used to train the proposed QuDeC network.

A. Test Datasets

The QuDeC network is tested on synthetic and real-world rainy images published by the authors of [2], [4], [29]. The DIDMDN synthetic test dataset consists of two subsets *Test-1* and *Test-2* containing 1000 and 1200 images, respectively [2]. The *Rain800* dataset shared by the authors of [29] contains 100 synthetic rainy images for testing. The *Rain200H* dataset contains 200 heavy rain synthetic images provided by the authors [4]. In addition to the synthetic images, we use 100 real-world rainy images provided by the authors of [2], [4], [29] to show the qualitative performance of QuDeC.

B. Training Details

The (y, x, s) rainy-clean-distortion label image pairs are used to train QuDeC using the loss \mathcal{L} . The Adam optimizer with the batch size of 1 is used to train the network. Learning rate is set equal to 0.0002 for the first 20 epochs and 0.0001 for the remaining epochs. During training initially λ_1 , λ_{CS} and λ_2

are set equal to 0.1, 0.1 and 1.0, respectively, but when the mean of all values in the confidences maps $c_{@1}$, $c_{@2}$ and $c_{@4}$ is greater than 0.8 then λ_1 is set equal to 0.03. QuDeC is trained for 60 epochs. During inference given a rainy image y , QuDeC estimates the residual map and its corresponding confidence map at three different scales, $\{\hat{x}, \hat{r}_{@1}, c_{@1}\}$ (at the original input size), $\{\hat{x}_{@2}, \hat{r}_{@2}, c_{@2}\}$, and $\{\hat{x}_{@4}, \hat{r}_{@4}, c_{@4}\}$ (at 0.25 scale of the input size) along with the location-quality-label-map \hat{s} .

V. EXPERIMENTAL RESULTS

In this section, we evaluate the performance of our method on both synthetic and real images. Peak-Signal-to-Noise Ratio (PSNR) and Structural Similarity index (SSIM) [44] measures are used to compare the performance of different methods on synthetic images. We visually inspect the performance of different methods on real images, as we don't have the ground truth clean images. The performance of the proposed QuDeC method is compared against the following recent state-of-the-art methods:

- Fu *et al.* [18] CNN method (TIP'17),
- Joint Rain Detection and Removal (JORDER) [4] (CVPR'17),
- Deep detailed Network (DDN) [3] (CVPR'17),
- Density-aware Image De-raining method using a Multi-stream Dense Network (DID-MDN) [2] (CVPR'18),
- REcurrent SE Context Aggregation Net (RESCAN) [34] (ECCV'18)
- Uncertainty guided Multi-scale Residual Learning (UMRL) network [1] (CVPR'19).

A. Results on Synthetic Test Images

The proposed QuDeC method based on cycle spinning [1] is also compared against the state-of-the-art algorithms qualitatively and quantitatively. Table II shows the quantitative performance of our method. As it can be seen from this table, our method clearly out-performs the present state-of-the-art image de-raining algorithms. On average, QuDeC out-performs the methods like RESCAN [34] and DIDMDN [2] by approximately 1dB. Furthermore, QuDeC outperforms the state-of-the-art method, UMRL+cycle-spinning [1], by 0.3dB on average. Figure 11 shows the qualitative performance of QuDeC against other methods on synthetic rainy images. The DDN method is over de-raining on some images and on others it is slightly under de-raining as shown in the second column of Figure 11. First four rows show under de-raining performance corresponding to methods RESCAN [34] and DIDMDN [2] in the third and fourth columns of Figure 11 respectively where we can observe residue rain streaks in their outputs. The last three rows show over de-raining of methods RESCAN [34] and DIDMDN [2] in the third and the fourth columns of Figure 11, respectively where we can observe the edges of the objects are blurred and objects like cables or wires and texture of birds feather have disappeared. Visually we can see in the fifth column of Figure 11 that our method produces images without any artifacts. From Figure 11 we can see that our method is able to

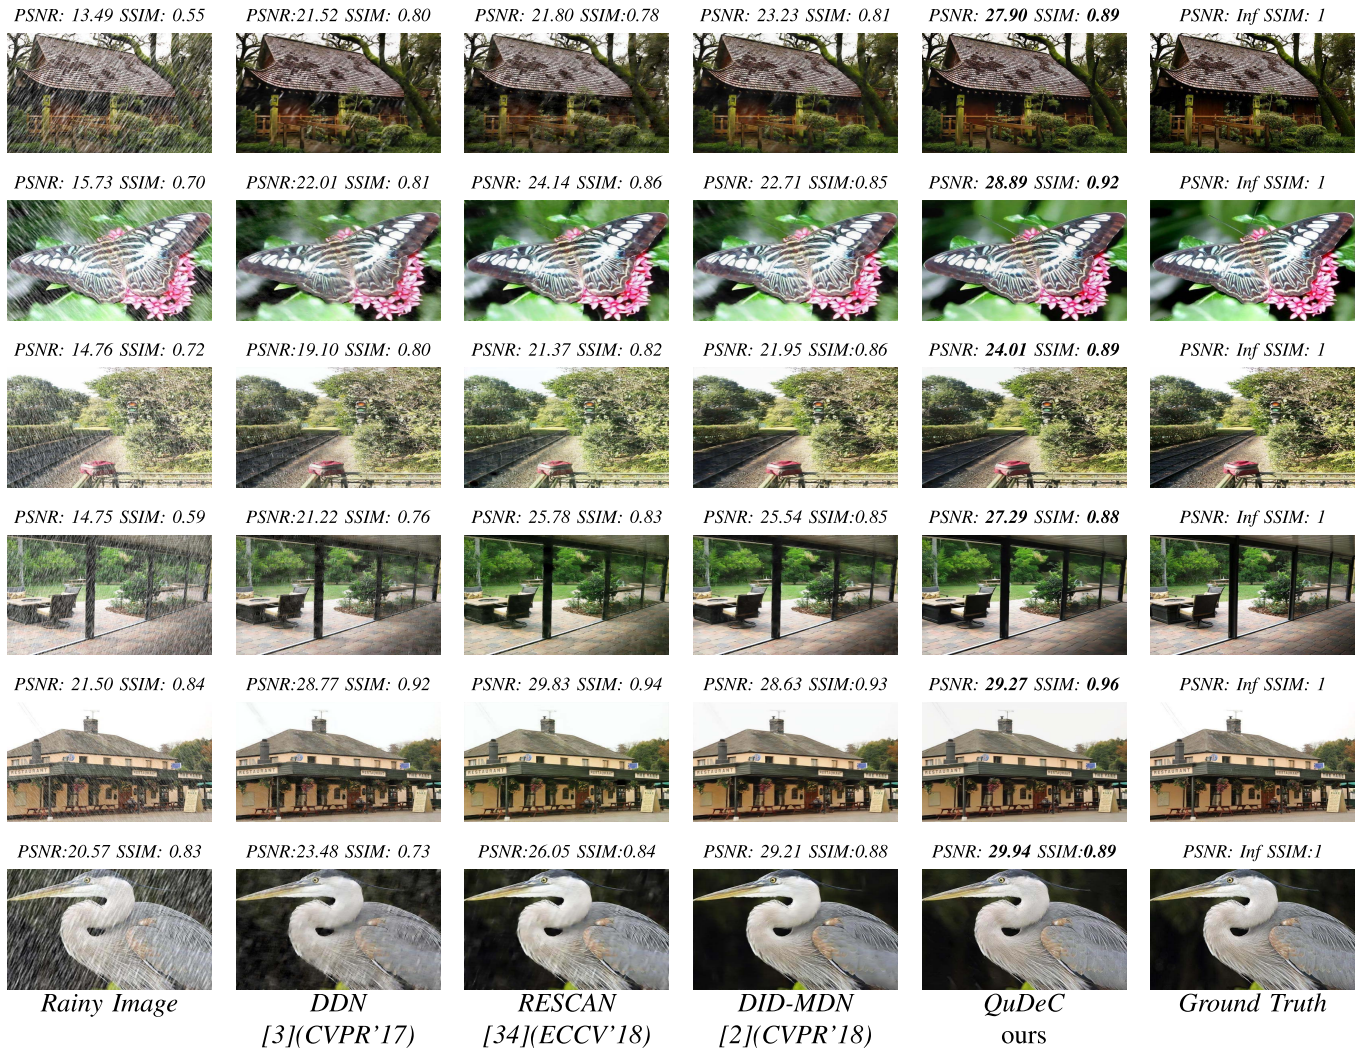


Fig. 11. Rain-streak removal results on sample images from the synthetic images consisting of different rain levels (low, medium and heavy) and directions.

TABLE II
QUANTITATIVE RESULTS EVALUATED IN TERMS OF AVERAGE SSIM AND PSNR (dB) (PSNR|SSIM)

Dataset	Testset	Fu <i>et al.</i> [18](TIP'17)	DDN [3](CVPR'17)	JORDER [4](CVPR'17)	RESCAN [34](CVPR'18)	DID-MDN [2](CVPR'18)	UMRL+ cycle spinning [1] (CVPR'19)	QuDeC ours
DIDMDN[2]	<i>Test-1</i>	22.07 0.84	27.33 0.90	24.32 0.86	27.19 0.87	27.95 0.91	29.77 0.92	30.43 0.93
	<i>Test-2</i>	19.73 0.83	25.63 0.88	22.26 0.84	25.65 0.88	26.08 0.90	26.67 0.92	26.72 0.92
Rain800 [29]	<i>Rain800</i>	18.95 0.78	21.33 0.80	21.13 0.81	24.37 0.84	23.57 0.87	24.52 0.86	24.61 0.86
JORDER [4]	<i>Rain200H</i>	21.83 0.81	23.24 0.83	24.02 0.86	25.97 0.90	23.43 0.86	26.38 0.92	26.74 0.93

Values highlighted in *red color*- indicate the best performance, *blue color* – indicate the second best performance, *green color* – indicate the third best performance among the de-raining methods on the test datasets.

- recover the texture on the wooden wall in the first row,
- produce clear objects with sharp edges in the fourth and the fifth images of the fifth column,
- remove rain streaks by maintaining the underlying textures on trees and on the feathers in the third and the sixth images of the fifth column.

B. Results on Real-World Rainy Images

We conducted experiments on the real-world images provided by the authors of [2], [4], [29]. Results are shown in

Figure 12. Similar to the results obtained on synthetic images, we observe the same trend of either over de-raining or under deraining by the other methods. On the other hand, our method is able to remove rain streaks while preserving details of objects in the resultant output images. For example, the background texture on tree is sharp when compared to other methods. Also, rain streaks are removed properly without losing background information of walls and flowers in the second, third and fourth images of the fifth column in Figure 12. All of these experiments clearly show that our method can handle




Fig. 12. De-raining results on sample real-world images.

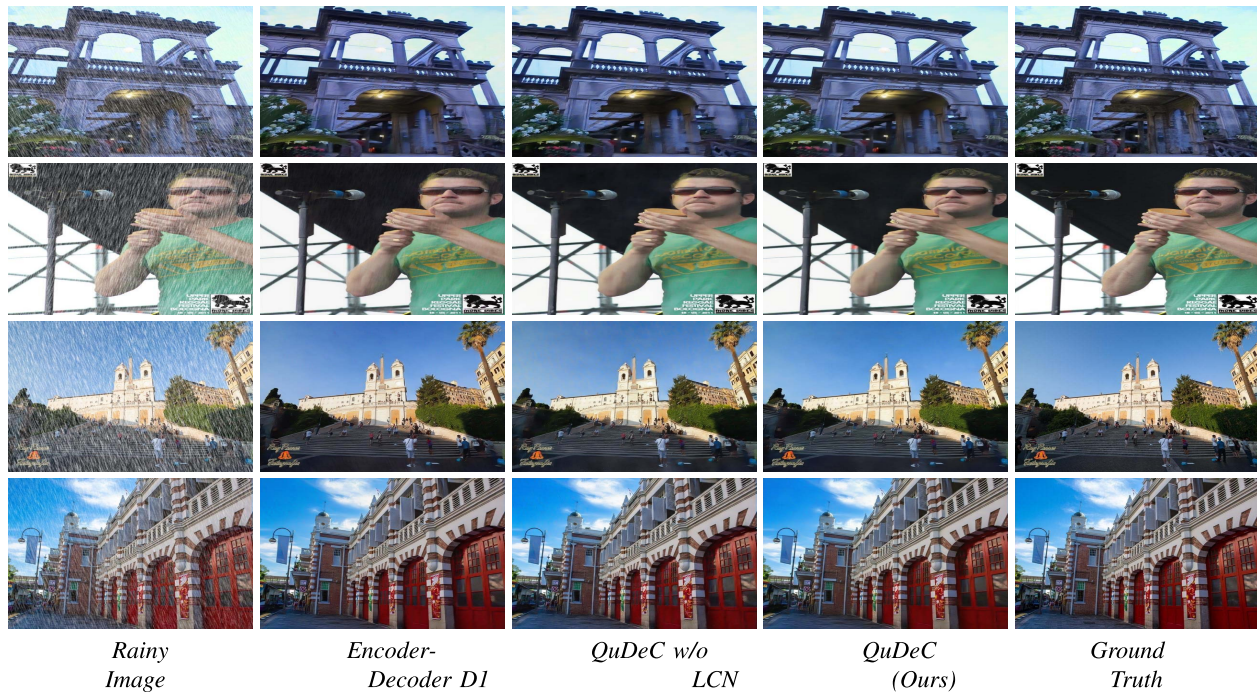


Fig. 13. Results of ablation study on synthetic images.

different levels of rain (low, medium and high) with different shapes and scales.

C. Ablation Study

We study the performance of each block's contribution to the QuDeC network by conducting extensive experiments on

the test datasets. We start with Encoder-Decoder D1 which are similar to UMRL [1] network. Encoder-Decoder D1 is trained as explained in UMRL [1] method. Now we add Decoder D2, where we call the resultant network as QuDeC w/o LCN, and output of D2 is supervised with cross-entropy loss using the ground-truth maps, s . Finally, we add LCN to Decoder

TABLE III
PSNR AND SSIM (PSNR|SSIM) RESULTS CORRESPONDING
TO THE ABLATION STUDY

Dataset	Testset	UMRL [1] (Encoder- Decoder D1)	QuDeC w/o LCN	QuDeC
DIDMDN [2]	<i>Test-1</i>	29.42 0.91	30.17 0.92	30.43 0.93
	<i>Test-2</i>	26.47 0.91	26.58 0.91	26.72 0.92
Rain800[29]	<i>Rain800</i>	24.19 0.85	24.42 0.86	24.61 0.86
JORDER [4]	<i>Rain200H</i>	26.17 0.91	26.56 0.92	26.74 0.93

D2 to construct our proposed network QuDeC. Table III shows the contribution of each block on the performance of QuDeC network. Addition of Decoder D2, i.e formulating the joint task of computing distortion level at each location and estimating the residual rain streak component improves the overall performance by approximately 0.37dB. Furthermore, introducing LCN to Decoder D2 improves the performance of QuDeC by 0.2dB. Figure 13, visually shows the improvement in performance after adding each block in constructing the QuDeC network. For example, we can clearly observe from Figure 13, QuDeC is able to reconstruct clear skies and dark backgrounds in the second and third images of the fourth column. Also QuDeC is able to reconstruct sharp objects when compared to the outputs of network with only Encoder-Decoder D1 in the first and last rows of Figure 13.

VI. CONCLUSION

We proposed a novel QuDeC to address the single image de-raining problem. In our approach, we formulate rain removal problem as a joint task of computing distortion level at each location and estimating the residual rain streak information. We judiciously combine the residual rain streak outputs at lower scales and distortion level information at each location using the corresponding confidence maps. Extensive experiments showed that QuDeC is robust enough to handle different levels of rain content for both synthetic and real-world rainy images.

APPENDIX A GENERATION-OF-LABELS NETWORK (GLN) ARCHITECTURE

Generation-of-Labels Network (GLN) is used to generate the ground-truth location-quality-label-maps s . We use residual blocks (ResBlock) as our building module for the GLN network. GLN network consists of a sequence of eight ResBlocks as shown in Figure 14. A ResBlock consists of a 1×1 convolution layer, a 3×3 convolution layer and a 3×3 convolution layer with dilation factor of 2 as shown in Figure 14. Given y , the GLN network process each patch p of size 128×128 at a time and outputs the distortion quality-label $s(p)$ for the corresponding patch p . GLN is trained on the rainy image patches and the corresponding labels, $\{y, s^\dagger\}$. s^\dagger is generated as explained in Section IV. GLN is trained for 40 epochs using the cross entropy loss with the Adam optimizer and the learning rate is set equal to 0.0002.

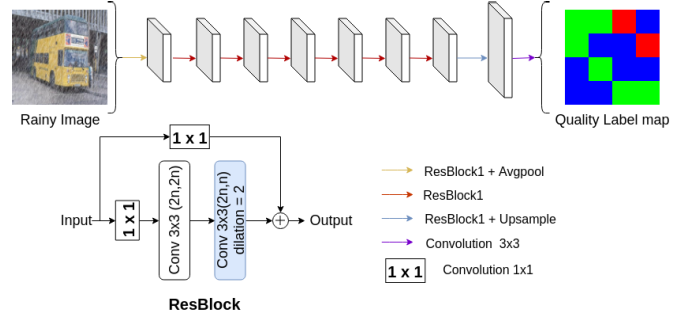


Fig. 14. Generation-of-Labels Network (GLN).

TABLE IV
QUANTITATIVE RESULTS OF QUDeC FOR DIFFERENT PATCH SIZES
IN QUALITY LABEL MAPS IN TERMS OF AVERAGE PSNR (dB)
(PSNR|SSIM) USING JORDER [4] TEST IMAGES

Patch Size	PSNR of QuDeC on JORDER [4]
256×256	26.48
128×128	26.74
64×64	26.63

Algorithm 1 Pseudo Code for Training QuDeC

Input: Pair of rainy and corresponding clean images $\{y, x\}$ and initial QuDeC network parameters θ .
Output: Derained image \hat{x} , location-quality-label-map \hat{s}
Result: $\hat{\theta}$, optimized network parameters of QuDeC
 $s = \text{GLN}(y)$;
for every epoch do
 for $\{y, x, s\}$ **do**
 $\hat{s} = \text{D2}(\text{Encoder}(y))$;
 $\hat{x} = \text{D1}(\text{Encoder}(y), \hat{s})$;
 update $\hat{\theta} \leftarrow$ using \mathcal{L}_u
 end
end

APPENDIX B QUALITY LABEL MAPS

Figure 15 shows sample \hat{s} (quality label map) and $\hat{s} \odot c_s$ computed using D2 and LCN networks. As can be seen from Figure 15, the errors in \hat{s} are corrected using c_s while multiplying element-wise with \hat{s} i.e $\hat{s} \odot c_s$. In this way, errors in image quality labels are corrected using the confidence scores c_s and given as input to Decoder D1 to compute the final de-rained image.

Table IV shows the performance of QuDeC on the experiments conducted using different patch sizes in quality label maps on the JORDER [4] test images. As shown in Table IV, QuDeC's best performance is achieved using the patch size 128×128 .

APPENDIX C PSEUDO CODE FOR QUDeC

Algorithm 1 presents the pseudo code for training QuDeC, where D1 is Decoder D1 and D2 is Decoder D2 of QuDeC.

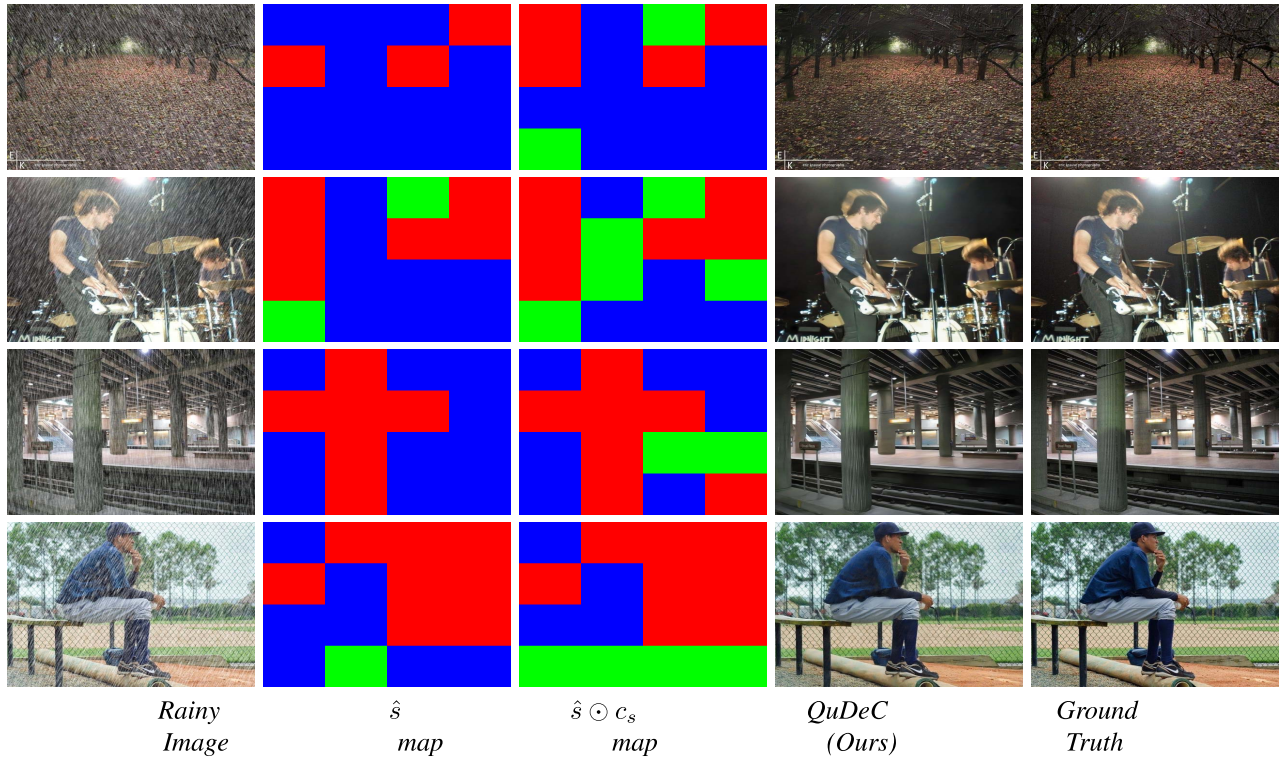


Fig. 15. Example quality label maps computed using Decoder D2 and LCN.

Algorithm 2 Pseudo Code for Testing QuDeC

Input: Input rainy image $\{y\}$, QuDeC network $f_{\hat{\theta}}(\cdot)$, $\hat{\theta}$ corresponds to the network parameters of QuDeC.
Output: Derained image \hat{x} , location-quality-label-map \hat{s}
for $\{y\}$ **do**
 $\hat{s} = \text{D2}(\text{Encoder}(y));$
 $\hat{x} = \text{D1}(\text{Encoder}(y), \hat{s});$
end

Note that QuDeC is trained for 60 epochs, and the ground truth location-quality-label-maps s are generated using GLN as explained in Section IV. Algorithm 2 presents the pseudo code for testing QuDeC.

ACKNOWLEDGMENT

This research is based upon work supported by the Office of the Director of National Intelligence (ODNI), Intelligence Advanced Research Projects Activity (IARPA), via IARPA R&D Contract No. 2019-19022600002. The views and conclusions 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.

REFERENCES

- [1] R. Yasarla and V. M. Patel, "Uncertainty guided multi-scale residual learning-using a cycle spinning cnn for single image de-raining," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 8405–8414.
- [2] H. Zhang and V. M. Patel, "Density-aware single image de-raining using a multi-stream dense network," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 695–704.
- [3] X. Fu, J. Huang, D. Zeng, Y. Huang, X. Ding, and J. Paisley, "Removing rain from single images via a deep detail network," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1715–1723.
- [4] W. Yang, R. T. Tan, J. Feng, J. Liu, Z. Guo, and S. Yan, "Deep joint rain detection and removal from a single image," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jul. 2017, pp. 1357–1366.
- [5] X. Zhang, H. Li, Y. Qi, W. Leow, and T. Ng, "Rain removal in video by combining temporal and chromatic properties," in *Proc. IEEE Int. Conf. Multimedia Expo*, Jul. 2006, pp. 461–464.
- [6] K. Garg and S. K. Nayar, "Vision and rain," *Int. J. Comput. Vis.*, vol. 75, no. 1, pp. 3–27, Jul. 2007.
- [7] V. Santhaseelan and V. K. Asari, "Utilizing local phase information to remove rain from video," *Int. J. Comput. Vis.*, vol. 112, no. 1, pp. 71–89, Mar. 2015.
- [8] J. Liu, W. Yang, S. Yang, and Z. Guo, "Erase or fill? Deep joint recurrent rain removal and reconstruction in videos," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 3233–3242.
- [9] M. Li *et al.*, "Video rain streak removal by multiscale convolutional sparse coding," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 6644–6653.
- [10] T.-X. Jiang, T.-Z. Huang, X.-L. Zhao, L.-J. Deng, and Y. Wang, "FastDeRain: A novel video rain streak removal method using directional gradient priors," *IEEE Trans. Image Process.*, vol. 28, no. 4, pp. 2089–2102, Apr. 2019.
- [11] J. Liu, W. Yang, S. Yang, and Z. Guo, "D3R-Net: Dynamic routing residue recurrent network for video rain removal," *IEEE Trans. Image Process.*, vol. 28, no. 2, pp. 699–712, Feb. 2019.
- [12] Y. Li, R. T. Tan, X. Guo, J. Lu, and M. S. Brown, "Rain streak removal using layer priors," in *Proc. IEEE Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2016, pp. 2736–2744.
- [13] L. Zhu, C.-W. Fu, D. Lischinski, and P.-A. Heng, "Joint bi-layer optimization for single-image rain streak removal," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Oct. 2017, pp. 2534–2536.
- [14] L.-W. Kang, C.-W. Lin, and Y.-H. Fu, "Automatic single-image-based rain streaks removal via image decomposition," *IEEE Trans. Image Process.*, vol. 21, no. 4, pp. 1742–1755, Apr. 2012.

- [15] H. Zhang and V. M. Patel, "Convolutional sparse and low-rank coding-based rain streak removal," in *Proc. IEEE Winter Conf. Appl. Comput. Vis. (WACV)*, Mar. 2017, pp. 1259–1267.
- [16] Y. Luo, Y. Xu, and H. Ji, "Removing rain from a single image via discriminative sparse coding," in *Proc. IEEE Int. Conf. Comput. Vis. (ICCV)*, Dec. 2015, pp. 3397–3405.
- [17] Y.-L. Chen and C.-T. Hsu, "A generalized low-rank appearance model for Spatio-temporally correlated rain streaks," in *Proc. IEEE Int. Conf. Comput. Vis.*, Dec. 2013, pp. 1968–1975.
- [18] X. Fu, J. Huang, X. Ding, Y. Liao, and J. Paisley, "Clearing the skies: A deep network architecture for single-image rain removal," *IEEE Trans. Image Process.*, vol. 26, no. 6, pp. 2944–2956, Jun. 2017.
- [19] A. Mittal, R. Soundararajan, and A. C. Bovik, "Making a 'completely blind' image quality analyzer," *IEEE Signal Process. Lett.*, vol. 20, no. 3, pp. 209–212, Mar. 2012.
- [20] D.-Y. Chen, C.-C. Chen, and L.-W. Kang, "Visual depth guided color image rain streaks removal using sparse coding," *IEEE Trans. Circuits Syst. Video Technol.*, vol. 24, no. 8, pp. 1430–1455, 2014.
- [21] D.-A. Huang, L.-W. Kang, Y.-C.-F. Wang, and C.-W. Lin, "Self-learning based image decomposition with applications to single image denoising," *IEEE Trans. Multimedia*, vol. 16, no. 1, pp. 83–93, Jan. 2014.
- [22] S. Li *et al.*, "Single image deraining: A comprehensive benchmark analysis," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 3838–3847.
- [23] W. Yang, R. T. Tan, J. Feng, J. Liu, S. Yan, and Z. Guo, "Joint rain detection and removal from a single image with contextualized deep networks," *IEEE Trans. Pattern Anal. Mach. Intell.*, to be published.
- [24] X. Fu, B. Liang, Y. Huang, X. Ding, and J. Paisley, "Lightweight pyramid networks for image deraining," *IEEE Trans. Neural Netw. Learn. Syst.*, to be published.
- [25] W. Yang, J. Liu, S. Yan, and Z. Guo, "Scale-free single image deraining via visibility-enhanced recurrent wavelet learning," *IEEE Trans. Image Process.*, vol. 28, no. 6, pp. 2948–2961, Jun. 2019.
- [26] J. Mairal, F. Bach, J. Ponce, and G. Sapiro, "Online dictionary learning for sparse coding," in *Proc. 26th Annu. Int. Conf. Mach. Learn. (ICML)*, 2009, pp. 689–696.
- [27] D. A. Reynolds, T. F. Quatieri, and R. B. Dunn, "Speaker verification using adapted Gaussian mixture models," *Digit. Signal Process.*, vol. 10, nos. 1–3, pp. 19–41, Jan. 2000.
- [28] G. Liu, Z. Lin, S. Yan, J. Sun, Y. Yu, and Y. Ma, "Robust recovery of subspace structures by low-rank representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 1, pp. 171–184, Jan. 2013.
- [29] H. Zhang, V. Sindagi, and V. M. Patel, "Image de-raining using a conditional generative adversarial network," *IEEE Trans. Circuits Syst. Video Technol.*, to be published.
- [30] Y. Wang, S. Liu, C. Chen, and B. Zeng, "A hierarchical approach for rain or snow removing in a single color image," *IEEE Trans. Image Process.*, vol. 26, no. 8, pp. 3936–3950, Aug. 2017.
- [31] R. Qian, R. T. Tan, W. Yang, J. Su, and J. Liu, "Attentive generative adversarial network for raindrop removal from a single image," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 2482–2491.
- [32] X. Wang, R. Girshick, A. Gupta, and K. He, "Non-local neural networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7794–7803.
- [33] T. Xu *et al.*, "AttnGAN: Fine-grained text to image generation with attentional generative adversarial networks," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 1316–1324.
- [34] X. Li, J. Wu, Z. Lin, H. Liu, and H. Zha, "Recurrent squeeze-and-excitation context aggregation net for single image deraining," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2018, pp. 262–277.
- [35] W. Wei, D. Meng, Q. Zhao, Z. Xu, and Y. Wu, "Semi-supervised transfer learning for image rain removal," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 3877–3886.
- [36] W. Yang, R. T. Tan, S. Wang, Y. Fang, and J. Liu, "Single image deraining: From model-based to data-driven and beyond," 2019, *arXiv:1912.07150*. [Online]. Available: <http://arxiv.org/abs/1912.07150>
- [37] X. Hu, C.-W. Fu, L. Zhu, and P.-A. Heng, "Depth-attentional features for single-image rain removal," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, Jun. 2019, pp. 8022–8031.
- [38] A. Kendall and Y. Gal, "What uncertainties do we need in Bayesian deep learning for computer vision?" in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2017, pp. 5574–5584.
- [39] R. Cipolla, Y. Gal, and A. Kendall, "Multi-task learning using uncertainty to weigh losses for scene geometry and semantics," in *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit.*, Jun. 2018, pp. 7482–7491.
- [40] R. Yasarla, F. Perazzi, and V. M. Patel, "Deblurring face images using uncertainty guided multi-stream semantic networks," 2019, *arXiv:1907.13106*. [Online]. Available: <http://arxiv.org/abs/1907.13106>
- [41] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Proc. Eur. Conf. Comput. Vis. (ECCV)*, 2016, pp. 694–711.
- [42] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2014, *arXiv:1409.1556*. [Online]. Available: <http://arxiv.org/abs/1409.1556>
- [43] H. Zhang and K. Dana, "Multi-style generative network for real-time transfer," 2017, *arXiv:1703.06953*. [Online]. Available: <https://arxiv.org/abs/1703.06953>
- [44] 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.



Rajeev Yasarla (Student Member, IEEE) received the bachelor's and master's degree 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.



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 currently 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 2016 ONR Young Investigator Award, the 2016 Jimmy Lin Award for Invention, A. Walter Tyson Assistant Professorship Award, the Best Paper Award at the IEEE AVSS 2017 & 2019, the Best Paper Award at the IEEE BTAS 2015, the Honorable Mention Paper Award at the IAPR ICB 2018, two Best Student Paper Awards at the IAPR ICPR 2018, and the Best Poster Awards at BTAS 2015 and 2016. He is currently an Associate Editor of the *IEEE Signal Processing Magazine*, the IEEE BIOMETRICS COMPENDIUM, and serves on the Machine Learning for Signal Processing (MLSP) Committee of the IEEE Signal Processing Society.