

Multi-Scale Boosted Dehazing Network with Dense Feature Fusion

Hang Dong¹ Jinshan Pan² Lei Xiang¹ Zhe Hu³ Xinyi Zhang¹
Fei Wang¹ Ming-Hsuan Yang^{4,5}

¹ College of Artificial Intelligence, Xi'an Jiaotong University

² Nanjing University of Science and Technology ³ Hikvision Research America

⁴ University of California, Merced ⁵ Google Research

Abstract

In this paper, we propose a Multi-Scale Boosted Dehazing Network with Dense Feature Fusion based on the U-Net architecture. The proposed method is designed based on two principles, boosting and error feedback, and we show that they are suitable for the dehazing problem. By incorporating the Strengthen-Operate-Subtract boosting strategy in the decoder of the proposed model, we develop a simple yet effective boosted decoder to progressively restore the haze-free image. To address the issue of preserving spatial information in the U-Net architecture, we design a dense feature fusion module using the back-projection feedback scheme. We show that the dense feature fusion module can simultaneously remedy the missing spatial information from high-resolution features and exploit the non-adjacent features. Extensive evaluations demonstrate that the proposed model performs favorably against the state-of-the-art approaches on the benchmark datasets as well as real-world hazy images.

1. Introduction

Hazy images are usually degraded by the turbid medium in the atmosphere during the imaging formation process. The goal of image dehazing is to restore a clean scene from a hazy image. This problem has received significant attention as images need to be first enhanced before applying high-level vision tasks (e.g., scene understanding [53] and detection [32]). Existing methods [5, 49, 60, 31, 58] usually model a hazy image I by:

$$I(x) = T(x)J(x) + (1 - T(x))A, \quad (1)$$

where J denotes a haze-free scene radiance, A describes the global atmospheric light indicating the intensity of ambient light, T is the transmission map, and x represents the pixel position.

To restore the haze-free scene radiance J from a hazy image I , data-driven deep learning approaches have been

demonstrated to be effective. Early approaches first use deep Convolutional Neural Networks (CNNs) to estimate transmission maps [5, 49, 60, 31, 59, 46] and then apply conventional methods (e.g., [23]) to estimate atmospheric light. However, the estimation of the transmission map or the atmospheric light from a single hazy input is not a trivial task, due to the airlight-albedo ambiguity [36] and the difficulty of obtaining ground truth data of transmission maps. In addition, inaccurate estimation of the transmission map or the atmospheric light would significantly interfere with the clear image restoration. To address this problem, several algorithms directly [35, 50, 61, 43, 9, 14, 41, 40] or iteratively [42, 10] estimate clean images based on deep CNNs. Nevertheless, these methods mainly adopt generic network architectures (e.g., DenseNet [61], U-Net [43], Dilated Network [9], Grid Network [40]), which are not well optimized for the image dehazing problem.

Different from many high-level vision tasks, inverse problems such as the image dehazing problem are highly ill-posed, where small measurement errors usually lead to dramatic changes. To solve these ill-posed problems, certain priors [15, 16, 23, 4] or careful algorithm designs are needed to make the problem well-posed. For a dehazing deep network, simply stacking more layers or using wider layers is inefficient for significant performance gain. Thus, it is of great interest and importance to tailor design network models for the dehazing problem.

In this work, we propose a dehazing network following two well-established principles for image restoration problems, i.e., boosting and error feedback mechanisms. The boosting strategy [44, 6, 51] is originally developed for image denoising by progressively refining the intermediate result from the previous iteration, and the error feedback mechanism, especially the back-projection technique [27, 12, 22], is designed for super-resolution to progressively recover details that are missed in the degradation process. We first show that the boosting strategy would facilitate the image dehazing task as well. Considering these two principles, we propose a Multi-Scale Boosted Dehazing Network (MS-

BDN) with the Dense Feature Fusion (DFF) based on the U-Net [52, 43] architecture. We interpret the decoder of the network as an image restoration module and thus incorporate the Strengthen-Operate-Subtract (SOS) boosting strategy [51] in the decoder to progressively restore the haze-free image. Due to the downsampling operations in the encoder of the U-Net, the spatial information compression may not be effectively retrieved from the decoder of the U-Net. To address this issue, we propose a DFF module based on the back-projection technique to effectively fuse features from different levels. We show that this module can simultaneously preserve spatial information from high-resolution features and exploit non-adjacent features for image dehazing. Extensive evaluations demonstrate that the proposed algorithm performs favorably against state-of-the-art dehazing methods.

The contributions of this work are summarized as follows:

- We propose a Multi-Scale Boosted Dehazing Network to incorporate the boosting strategy and the back-projection technique neatly for image dehazing.
- We show that the boosting strategy can help image dehazing algorithms under a certain axiom and show that the network design with the boosting strategy is simple but effective in practice.
- We demonstrate that the Dense Feature Fusion module based on the back-projection technique can effectively fuse and extract features from different scales for image dehazing, and help improve the performance of dehazing networks.

2. Related Work

Image dehazing. Since image dehazing is an ill-posed problem, existing methods often use strong priors or assumptions as additional constraints to restore the transmission map, global atmospheric light, and scene radiance [15, 56, 16, 23, 4]. In [15], Fattal uses the surface shading information to estimate transmission maps. By assuming that haze-free images should have higher contrast than hazy images, a method that enhances the visibility of hazy images by maximizing the local contrast is developed [56]. In [23], He et al. propose a dark channel prior on the pixel intensities of clean outdoor images and develop a dehazing method using the prior. As pixels in a given RGB space cluster are often non-local, Berman et al. [4] develop an effective non-local path prior for image dehazing. Since these priors and assumptions are introduced for specific scenes or atmospheric conditions, these dehazing methods are less effective on the scenes when the priors do not hold. For example, the dark channel prior [23] does not perform well for images without zero-intensity pixels.

To address these problems, numerous data-driven methods based on deep learning have been developed [5, 49, 60, 31, 58, 34, 42, 59, 18] to first estimate transmission

maps and then restore images. These algorithms are effective when the transmission maps and atmospheric lights are accurately estimated. However, due to the airlight-albedo ambiguity [36], they usually lead to results with significant color distortions when the estimated atmospheric lights or transmission maps are not accurate enough. On the other hand, end-to-end [35, 50, 61, 43, 9, 41, 40, 47, 13] dehazing networks have been proposed to directly restore clean radiance scenes without estimating transmission maps and atmospheric lights. Nevertheless, these methods are mainly based on some generic network architectures without significant modification, which are inefficient for the image dehazing problem.

Boosting algorithms for image restoration. Numerous boosting methods have been developed for image denoising [6, 44, 51] to progressively refine the result by feeding the enhanced previous estimation as the input. Recently, the boosting strategies are incorporated with deep CNNs for object classification [21, 45] and image denoising [7, 8]. In this work, we show that the boosting strategy can also be applied to image dehazing and incorporate it into network design for this dehazing task.

Multi-scale feature fusion. Feature fusion has been widely used in the network design for performance gain by exploiting features from different layers. Numerous image restoration methods fuse features by using dense connection [61], feature concatenation [64] or weighted element-wise summation [9, 62]. Most existing feature fusion modules reuse features of the same scale from previous layers. In [25, 37], the features at different scales are projected and concatenated by using the strided convolutional layer. Although these methods can merge multiple features from different levels, the concatenation scheme is not effective for extracting useful information. To share information among adjacent levels, the grid architectures [65, 17] are proposed by interconnecting the features from adjacent levels with convolutional and deconvolutional layers. Recently, Liu et al. propose an effective end-to-end trainable grid network [40] for image dehazing. However, these methods do not explicitly exploit features from non-adjacent levels and cannot be easily applied to other architectures. In [22], a deep feedback mechanism for projection errors [27, 12] is developed to merge features from two levels. Different from [22], we develop a DFF module which fuses features from multiple scales effectively.

3. Proposed Method

3.1. Multi-Scale Boosted Dehazing Network

The proposed network is based on the U-Net [52] architecture, and we design a multi-scale boosted decoder inspired by the SOS boosting method [51]. As shown in Figure 1, the network includes three components, an encoder module G_{Enc} , a boosted decoder module G_{Dec} , and a feature

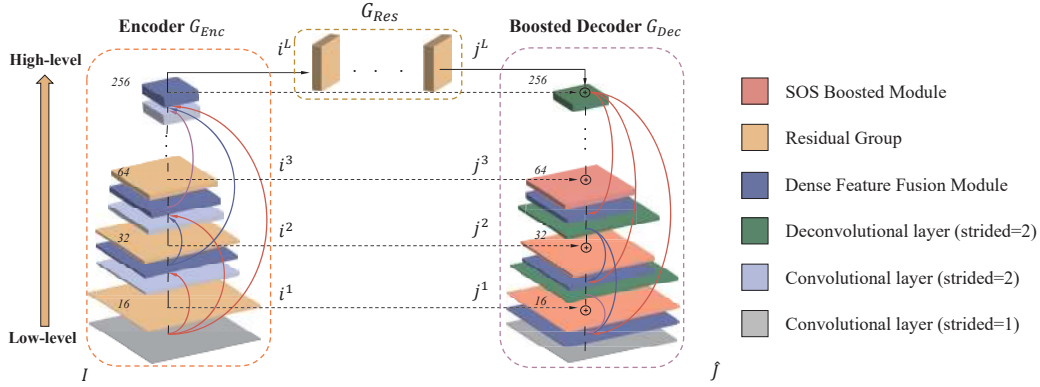


Figure 1. **Architecture of the proposed MSBDN with DFF modules.** Skip connections are employed to introduce the corresponding feature maps from the encoder module to the boosted decoder module.

restoration module G_{Res} .

Boosting in image dehazing. The boosting algorithm has been shown to be effective for image denoising [6, 44, 51]. The SOS boosting algorithm [51] operates the refinement process on the strengthened image, based on the previously estimated image. The algorithm has been shown to improve the Signal-to-Noise Ratio (SNR) under the axiom that the denoising method obtains better results in terms of SNR on the images of the same scene but less noise.

For image dehazing, the SOS boosting strategy can be formulated similarly as

$$\hat{J}^{n+1} = g(I + \hat{J}^n) - \hat{J}^n, \quad (2)$$

where \hat{J}^n denotes the estimated image at the n -th iteration, $g(\cdot)$ is the dehazing approach, and $I + \hat{J}^n$ represents the strengthened image using the hazy input I . We show that the boosting method can facilitate image dehazing performance in terms of Portion of Haze (PoH) under a similar axiom as that for denoising. Here the portion of haze of the image J in (1) is defined as $PoH(J) = (1 - T)A/J$, and it is proportional to $1 - T$ for hazy images of the same scene.

Axiom 1. *The dehazing method g obtains better results in terms of PoH on the images of the same scene but less haze. That is, if J_1 and J_2 are the images of the same scene, and $PoH(J_1) < PoH(J_2)$, then $PoH(g(J_1)) < PoH(g(J_2))$.*

Proposition 1. *Under Axiom 1, the SOS boosting strategy in (2) improves the dehazing performance, as*

$$PoH(\hat{J}^{n+1}) < PoH(\hat{J}^n). \quad (3)$$

The experimental verification of **Axiom 1** and proof of **Proposition 1** are given in the supplementary material.

According to **Proposition 1**, we develop a deep boosted network based on the SOS boosting strategy, to effectively solve image dehazing by a data-driven approach.

Deep boosted dehazing network. In a U-Net network for dehazing, we interpret the decoder as the haze-free image restoration module. To progressively refine the feature j^L

from the feature restoration module G_{Res} , we incorporate the SOS boosting strategy in the decoder of the proposed network and the structure of the SOS boosted module is illustrated in Figure 2(e). In the SOS boosted module at level n , we upsample the feature j^{n+1} from the previous level, strengthen it with the latent feature i^n from the encoder, and generate the boosted feature j^n through the refinement unit, as

$$j^n = \mathcal{G}_{\theta_n}^n(i^n + (j^{n+1}) \uparrow_2) - (j^{n+1}) \uparrow_2, \quad (4)$$

where \uparrow_2 denotes the upsampling operator with a scaling factor of 2, $(i^n + (j^{n+1}) \uparrow_2)$ represents the strengthened feature, and $\mathcal{G}_{\theta_n}^n$ denotes the trainable refinement unit at the (n) -th level parameterized by θ_n .

In this work, we implement each refinement unit with a residual group as also used in the encoder. Clearly, (4) is derivable and refines the feature $(j^{n+1}) \uparrow_2$ in a signal strengthening manner. At the end of the decoder, a convolutional layer is used for reconstructing the estimated haze-free image \hat{J} from the final features.

Alternatives to SOS boosted module. For completeness, we also list four alternatives to the proposed SOS boosted module for dehazing. The diffusion [44] and twicing [6] schemes can be applied to designing the boosted modules as shown in Figure 2(a) and Figure 2(b). They can be formulated respectively as

$$j^n = \mathcal{G}_{\theta_n}^n((j^{n+1}) \uparrow_2), \quad (5)$$

and

$$j^n = \mathcal{G}_{\theta_n}^n(i^n - (j^{n+1}) \uparrow_2) + (j^{n+1}) \uparrow_2. \quad (6)$$

We adopt the SOS boosting strategy in the proposed method, as the refinement units in (5) and (6) do not fully exploit the feature i^n , which contains more structural and spatial information compared with the upsampled feature $(j^{n+1}) \uparrow_2$.

Another related module is the pyramid module (shown in Figure 2(c)) from the Feature Pyramid Network (FPN) that has been widely used for panoptic segmentation [29], super-resolution [30], and pose estimation [11]. It can be

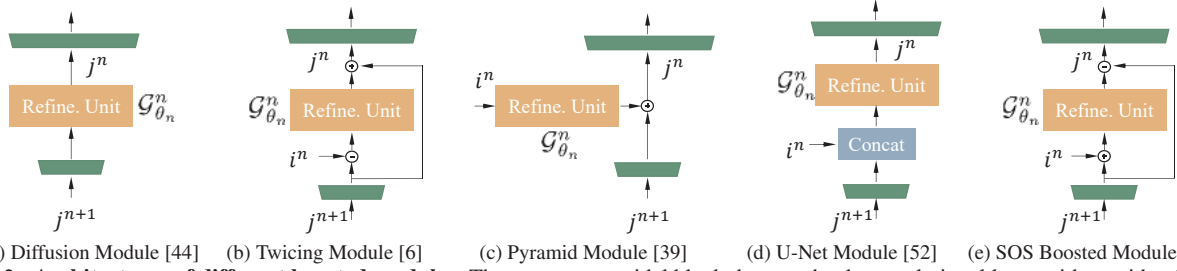


Figure 2. **Architectures of different boosted modules.** The green trapezoidal block denotes the deconvolutional layer with a stride of 2 in the decoder module.

formulated as

$$j^n = (j^{n+1}) \uparrow_2 + \mathcal{G}_{\theta_n}^n(i^n). \quad (7)$$

Here, the refinement unit is blind to the upsampled feature $(j^{n+1}) \uparrow_2$ from the previous level.

Finally, we also evaluate the decoder module of the original U-Net (shown in Figure 2(d)), which concatenates the upsampled boosted feature $(j^{n+1}) \uparrow_2$ and the latent feature i^n in the module. It can be formulated as

$$j^n = \mathcal{G}_{\theta_n}^n(i^n, (j^{n+1}) \uparrow_2). \quad (8)$$

Since the subtraction and addition operations in (4) and (6) can be absorbed by the learnable refinement unit $\mathcal{G}_{\theta_n}^n$, the decoder module of U-Net can imitate the boosting strategy with a proper training. However, this imitation is not guaranteed with an implicit and unconstrained fusing process of $(j^{n+1}) \uparrow_2$ and i^n .

We evaluate the proposed SOS boosting strategy with these alternatives in Section 4.3, and show that the network with the SOS boosting strategy obtains the best results.

3.2. Dense Feature Fusion Module

The U-Net architecture is inherently limited in several aspects, e.g., missing spatial information during the downsampling processes in the encoder and lacking sufficient connections between the features from non-adjacent levels. To remedy the missing spatial information from upper-level features and fully exploit the features from non-adjacent levels, a straightforward approach is to first resample all the features to the same scale, and then fuse them by a bottleneck layer (concatenation layer and convolutional layer) as DenseNet [26] does. However, simply using concatenation is less effective for feature fusion since the features from different levels are of different scales and dimensions.

The back-projection technique in super-resolution [27] is an efficient method designed for generating high-resolution contents by minimizing the reconstruction errors between the estimated high-resolution result \hat{H} and multiple observed low-resolution inputs. In [12], an iteratively back-projection algorithm is developed for the case with one single low-resolution input with

$$\hat{H}_{t+1} = \hat{H}_t + h(f(\hat{H}_t) - L_{ob}), \quad (9)$$

where \hat{H}_t is the estimated high-resolution output at t -th iteration, L_{ob} is the observed low-resolution image acquired by using the downsampling operator f , and h denotes the back-projection operator.

Motivated by the back-projection algorithm in (9), we propose a DFF module for effectively remedying the missing information and exploiting features from non-adjacent levels. The proposed DFF is designed to further enhance the boosted features at the current level with an error feedback mechanism, and it is used in both the encoder and the decoder. As shown in Figure 1, two DFF modules are introduced at each level, one before the residual group in the encoder and another after the SOS boosted module in the decoder. The enhanced DFF output in the encoder/decoder is directly connected to all the following DFF modules in the encoder/decoder for feature fusion.

We describe how to use the DFF in the decoder in the following, and the DFF in the encoder can be derived accordingly. The DFF for the n -th level of the decoder (\mathcal{D}_{de}^n as shown in Figure 3), is defined by

$$\tilde{j}^n = \mathcal{D}_{de}^n(j^n, \{\tilde{j}^L, \tilde{j}^{L-1}, \dots, \tilde{j}^{n+1}\}), \quad (10)$$

where j^n is the boosted feature at the n -th level of the decoder, \tilde{j}^n is the enhanced feature through feature fusion, L is the number of the network levels, and $\{\tilde{j}^L, \tilde{j}^{L-1}, \dots, \tilde{j}^{n+1}\}$ are the preceding enhanced features from all the $(L - n)$ preceding DFF modules in the decoder. Here, we adopt a progressive process to enhance the boosted feature j^n by giving one preceding enhanced feature \tilde{j}^{L-t} , $t \in \{0, 1, \dots, L - n - 1\}$ at a time. The update procedure can be defined by:

- Compute the difference e_t^n between the boosted feature at the t -th iteration j_t^n ($j_0^n = j^n$) and the t -th preceding enhanced feature \tilde{j}^{L-t} by

$$e_t^n = p_t^n(j_t^n) - \tilde{j}^{L-t}, \quad (11)$$

where p_t^n denotes the projection operator which downsamples the boosted feature j_t^n to the same dimension of \tilde{j}^{L-t} .

- Update the boosted feature j_t^n with the back-projected difference by

$$j_{t+1}^n = q_t^n(e_t^n) + j_t^n, \quad (12)$$

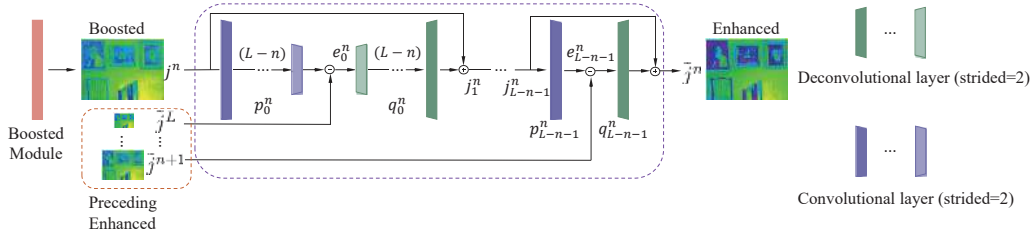


Figure 3. **Network architecture of the proposed DFF module at the n -th level of the decoder.** The dense feature fusion in the decoder module uses the back-projection technique to exploit the features from all the preceding levels. The blocks with the same color share with the same parameters.

where q_t^n denotes the back-projection operator which upsamples the difference e_t^n at the t -th iteration to the same dimension of j_t^n .

- The final enhanced feature \tilde{j}^n can be obtained after all the preceding enhanced features have been considered.

Different from the traditional back-projection techniques [27, 12], the sampling operators p_t^n and q_t^n are unknown to the network. Motivated by the recent deep back-projection network for super-resolution [22], we adopt the strided convolutional/deconvolutional layers to learn the downsampling/upsampling operators in an end-to-end manner. To avoid introducing too many parameters, we stack $(L-n-t)$ convolutional/deconvolutional layers with strides of 2 to implement the downsampling/upsampling operators in p_t^n/q_t^n .

We note that the DFF for the n -th level of the encoder (\mathcal{D}_{en}^n) can be defined by

$$\tilde{i}^n = \mathcal{D}_{en}^n(i^n | \tilde{i}^1, \tilde{i}^2, \dots, \tilde{i}^{n-1}), \quad (13)$$

where i^n is the latent feature at the n -th level of the encoder, $\{\tilde{i}^1, \tilde{i}^2, \dots, \tilde{i}^{n-1}\}$ are the preceding enhanced features from all the $(n-1)$ preceding DFF modules in the encoder, and \mathcal{D}_{en}^n shares the same architecture as the module \mathcal{D}_{de}^{L-n} at the $(L-n)$ -th level of the decoder but switches the positions of downsampling operations p_t^n and upsampling operations q_t^n .

Compared with other sampling and concatenation fusion methods, the proposed module can better extract the high-frequency information from the high-resolution features from proceeding layers due to the feedback mechanism. By progressively fusing these differences back to the downsampled latent features, the missing spatial information can be remedied. On the other hand, this module can exploit all the preceding high-level features and operate as an error-correcting feedback mechanism to refine the boosted features for obtaining better results.

3.3. Implementations

As shown in Figure 1, the proposed network contains four strided convolutional layers and four strided deconvolutional layers. The Leaky Rectified Linear Unit (LReLU) with a negative slope of 0.2 is used after each convolutional and deconvolutional layer. The residual group [38] consists of three residual blocks, and 18 residual blocks are used in G_{Res} . The filter size is set as 11×11 pixels in the first convolutional layer in the encoder module and 3×3 in all the other

convolutional and deconvolutional layers. We jointly train the MSBDN and DFF module and use the Mean Squared Error (MSE) as the loss function to constrain the network output and ground truth. The entire training process contains 100 epochs optimized by the ADAM solver [28] with $\beta_1 = 0.9$ and $\beta_2 = 0.999$ with a batch size of 16. The initial learning rate is set as 10^{-4} with a decay rate of 0.75 after every 10 epochs. All the experiments are conducted on an NVIDIA 2080Ti GPU. The source code and trained models are available at <https://github.com/BookerDeWitt/MSBDN-DFF>.

4. Experimental Results

4.1. Datasets

We evaluate the proposed algorithm on the following datasets [33, 63, 1] against the state-of-the-art methods.

RESIDE dataset. The RESIDE dataset [33] contains both synthesized and real-world hazy/clean image pairs of indoor and outdoor scenes. To learn a general dehazing model for both indoor and outdoor scenes, we select as the training set 9000 outdoor hazy/clean image pairs and 7000 indoor pairs from the RESIDE training dataset [33] by removing redundant images from the same scenes. To further augment the training data, we resize images of each pair with three random scales within the scale range of $[0.5, 1.0]$. We randomly crop 256×256 patches from hazy images and randomly flip them horizontally or vertically as the inputs. SOTS is the test subset of the RESIDE dataset, which contains 500 indoor hazy images and 500 outdoor hazy images. For comparison, all the methods are trained on the selected RESIDE training dataset and evaluated on the SOTS.

HazeRD dataset. The HazeRD dataset [63] contains 75 synthesized hazy images with realistic haze conditions. Since most evaluated methods on the HazeRD dataset [63] are trained on the synthesized NYUv2 [55] dataset, we train the proposed model on NYUv2 with the same setting as [49] for fair comparisons.

NTIRE2018-Dehazing challenge dataset. The NTIRE2018-Dehazing challenge [1] includes an indoor dataset (referred to as I-HAZE [2]) and an outdoor dataset (referred to as O-HAZE [3]). Both datasets provide training sets and test sets. For each dataset, we train the proposed model using its training set and evaluate it

Table 1. **Quantitative evaluations on the benchmark dehazing datasets.** Red texts and blue texts indicate the best and the second-best performance respectively. \uparrow and \downarrow mean the better methods should achieve higher/lower score of this metric.

Methods		DCP [23]	NLD [4]	AODNet [31]	MSCNN [49]	MsPPN [61]	DcGAN [35]	GFN [50]	GCANet [9]	PFFNet [43]	GDN [40]	DuRN [41]	Ours
SOTS [33]	PSNR \uparrow	18.75	17.27	18.80	17.57	29.94	25.37	24.11	28.13	29.22	31.51	31.92	33.79
	SSIM \uparrow	0.859	0.750	0.834	0.811	0.958	0.917	0.899	0.945	0.954	0.983	0.980	0.984
HazeRD [63]	CIEDE2000 \downarrow	14.83	16.40	13.23	13.80	15.50	12.02	14.83	14.45	14.46	13.93	12.48	10.36
	SSIM \uparrow	0.767	0.727	0.833	0.794	0.759	0.826	0.802	0.819	0.808	0.833	0.840	0.881
I-HAZE [2]	PSNR \uparrow	14.43	14.12	13.98	15.22	22.50	16.06	15.84	14.95	16.01	16.62	21.23	23.93
	SSIM \uparrow	0.752	0.654	0.732	0.755	0.871	0.733	0.751	0.719	0.740	0.787	0.842	0.891
O-HAZE [3]	PSNR \uparrow	16.78	15.98	15.03	17.56	24.24	19.34	18.16	16.28	18.76	18.92	20.45	24.36
	SSIM \uparrow	0.653	0.585	0.539	0.650	0.721	0.681	0.671	0.645	0.669	0.672	0.688	0.749

on the corresponding test set. We adopt the same data augmentation strategy as that used for the RESIDE dataset.

4.2. Performance Evaluation

We evaluate the proposed algorithm against state-of-the-art methods based on the hand-crafted priors (DCP [23] and NLD [4]) and deep convolutional neural networks (AOD [31], MSCNN [49], MsPPN[61], DcGAN [35], GFN [50], GCANet [9], PFFNet [43], GDN [40], and DuRN[41]). We use the metrics PSNR, SSIM [57], and CIEDE2000 [63] to evaluate the quality of restored images. We note that many existing dehazing methods [50, 9, 43, 41] report their results only on the SOTS indoor images with the models trained on various datasets [33, 55, 54]. Moreover, the GDN method [40] reports the results on the SOTS indoor and outdoor sets, with the models trained on the indoor scenes and outdoor scenes separately. For fair comparisons, we retrain these methods (GFN, PFFNet, GCANet, MsPPN, and DuRN) using their provided training codes on the same training dataset and evaluate them on the same test dataset, as the proposed algorithm. Other methods are evaluated with the provided pre-trained models.

The first row in Table 1 shows the quantitative results on the SOTS dataset. As expected, the methods based on hand-crafted features [23, 4] do not perform well. The methods [31, 49, 35] that use deep CNNs in a two-stage restoration (estimating the transmission maps and atmospheric lights first and then hazy-free images), are less effective when the atmospheric lights are not correctly estimated. Since the GFN [50] method applies the hand-crafted derived images as inputs, it is less effective on the scenes when these derived images cannot enhance the hazy images. On the other hand, the methods [9, 43, 61, 41, 40] that directly estimate clear images based on end-to-end trainable networks, generate better results than the other indirect ones. The proposed method outperforms these algorithms in both PSNR and SSIM, since these networks are not well optimized for the dehazing problem. In contrast, the proposed architecture based on the boosting and back-projection technique is more effective for the task.

The evaluation results on the HazeRD dataset also demonstrate the effectiveness of the proposed algorithm for dehazing realistic images. The proposed algorithm achieves better CIEDE2000 and SSIM results than other methods. On the NTIRE18-Dehazing challenge [1], our method achieves comparable results with the MsPPN [61], which is specially

designed for ultra high-resolution datasets including extra pre-processing and post-processing steps.

Figure 5 shows two examples from the SOTS dataset. The DCP algorithm generates the results with significant color distortions. The dehazed images by other deep learning frameworks still contain significant artifacts. In contrast, our algorithm restores these images well.

We further evaluate our algorithm on real images. Figure 6 shows a real hazy image and the dehazed results from state-of-the-art methods. The dehazed image by our method is sharper and brighter. More qualitative results are provided in the supplementary material.

Perceptual quality for high-level vision tasks. As the dehazing algorithms are usually used as the pre-processing step for high-level applications, it is helpful to evaluate the perceptual quality of the dehazed results. In this work, we provide evaluation results on the perceptual quality for the object detection task. To obtain the data for this evaluation, we generate hazy images using the images from the KITTI detection dataset [19]. We first estimate the depth map for each image by a single-image depth estimation method Monodepth2 [20], and then use the depth map to synthesize a haze image following the protocols of the RESIDE dataset [33]. This dataset is referred to as the KITTI Haze dataset in this work. We evaluate the proposed method with the following approaches: DCP [23], GFN [50], PFFNet [43], and DuRN [41]. For detection accuracy, we use the state-of-the-art method YOLOv3 [48] to evaluate on the dehazed images from different dehazing methods.

The detection results are shown in Table 2. The dehazed images restored from the proposed method obtain the highest detection accuracy, which shows that our method can restore images with better perceptual quality. The qualitative results in Figure 4 demonstrate that our method not only generates better dehazed image but also helps the detection algorithm to recognize cars and pedestrians.

4.3. Ablation Study and Analysis

In this section, we analyze how the proposed method performs for image dehazing. All the baseline methods mentioned below are trained using the same setting as the proposed algorithm for fair comparisons.

Study of the network configuration. To investigate the effect of the network configuration, we evaluate the PFFNet [43], MSBDN, and MSBDN-DFF on the RESIDE dataset [33] under different network configurations. The



Figure 4. **Detection results using the dehazed images from the synthesized KITTI Haze dataset.** Best viewed on a high-resolution display.



Figure 5. **Visual results on the SOTS dataset.** The results in (c)-(g) contain some color distortions and haze residual, while the dehazed image in (h) by our method is much clearer. Best viewed on a high-resolution display.




Figure 6. **Visual results on the real-world image.** The proposed method generates a clearer dehazed image with less color distortions. Best viewed on a high-resolution display.



Figure 7. **Visual results on the effect of the boosted decoder and dense feature fusion.** Compared with the baseline models, the proposed MSBDN-DFF model can generate a clearer image with more details. Best viewed on a high-resolution display.

Table 2. **Detection results on the KITTI Haze dataset.** We apply dehazing methods trained on the RESIDE dataset [33] to restore clean images and evaluate their perceptual quality for the object detection task. The mAP is the abbreviation of mean average precision. **Red texts** indicate the best detection precision.

YOLOv3		Hazy	DCP [23]	PFFNet [43]	DuRN [41]	Ours
KITTI	PSNR	10.35	13.53	11.86	16.95	17.97
	mAP	0.084	0.239	0.143	0.360	0.374

Table 3. **Effect of the number of feature levels and ResBlocks.** L denotes the number of feature levels and B denotes the number of ResBlocks [24] in G_{Res} . All the experiments are conducted on the SOTS dataset [33]. **Red texts** indicate the best performance.

Configurations		PFFNet [43]	MSBDN	MSBDN-DFF
$L = 4, B = 8$	PSNR	28.23	30.92	32.07
	Param	2.6M	3.1M	4.5M
$L = 5, B = 8$	PSNR	28.60	32.00	32.93
	Param	10.2M	12.6M	19.6M
$L = 5, B = 18$	PSNR	29.22	32.85	33.79
	Param	22M	24M	31M

PFFNet [43] adopts an encoder-decoder architecture without boosted modules and can be served as a baseline of our methods. The MSBDN is the encoder-decoder architecture with the SOS boosted modules, and the MSBDN-DFF is the

proposed method with both the SOS boosted modules and DFF modules.

Since all the above-mentioned methods are built upon the encoder-decoder architecture, we study the effect of two network parameters: the number of feature levels L and the number of ResBlocks B in the feature restoration module G_{Res} . The quantitative results are shown in Table 3. As the proposed SOS boosted modules and DFF modules can effectively extract features, larger numbers of levels (L) and ResBlocks (B) would lead to higher performance on our method. Moreover, introducing the SOS boosted modules and DFF modules can bring significant performance gains under different configurations, which demonstrates the effectiveness of the boosting and back-projection algorithms for image dehazing. For all the other experiments, we use $L = 5$ and $B = 18$ as the default network configurations.

Effectiveness of the SOS boosting strategy. The proposed MSBDN is based on an encoder-decoder architecture and the design of the boosted decoder is motivated by the SOS boosting strategy. We evaluate the following alternatives

Table 4. **Analysis on each component of the MSBDN-DFF.** All the methods are evaluated on the SOTS dataset [33] using the same training setting as the proposed algorithm. **Red texts** indicate the best performance of each part.

Baselines	Effectiveness of SOS boosting strategy						Effectiveness of the dense feature fusion				
	PFFNet	Diffusion	Twicing	Pyramid	U-Net	MSBDN	MSBDN-S	MSBDN-M	MSBDN+	PFFNet-DFF	MSBDN-DFF
boosting algorithm		✓	✓		✓	✓	✓	✓	✓		✓
FPN-like				✓							
strengthened feature						✓	✓	✓			✓
dense feature fusion							✓			✓	✓
simultaneously BP								✓			
progressively BP									✓		✓
Parameters	22M	24M	24M	24M	24M	24M	26M	31M	31M	29M	31M
PSNR	29.22	27.45	32.49	31.53	32.31	32.85	33.02	33.24	33.16	32.95	33.79

to demonstrate the effectiveness of the SOS boosting strategy. Starting from the baseline PFFNet model [43], we first remove the skip connections and add a diffusion module (shown in Figure 2(a)) to each level of the decoder to implement the Diffusion Decoder. Next, we re-introduce the skip connections and change the diffusion modules with the twicing modules in Figure 2(b) and the pyramid modules in Figure 2(c) to construct the Twicing Decoder and Pyramid Decoder. Finally, we also evaluate the original U-Net Decoder in Figure 2(d) for a thorough study.

The evaluation results are shown in Table 4. The network with the simple Diffusion Decoder performs the worst as expected. In general, the networks with the boosting modules achieve significant performance improvements over the PFFNet without using boosting strategies. Furthermore, the MSBDN with the proposed boosted decoder outperforms other boosting strategies by a margin of 0.36 dB at least. We note that the MSBDN model improves the performance without introducing any extra layers, which demonstrates that the SOS boosting strategy better fits the problem and can benefit image dehazing algorithms.

Figure 7 shows an example of the visual results. The MSBDN model with the SOS boosted decoder performs better in haze removal. Visualizations of the learned features are provided in the supplementary material.

Effectiveness of the dense feature fusion. The DFF module is designed to remedy the spatial information from high-resolution features and exploit the non-adjacent features. To demonstrate the effectiveness of the proposed DFF module, we evaluate several alternative solutions to the DFF.

Starting from the MSBDN, we first fuse all the preceding features using the sampling operator and the bottleneck layer, as stated in the first paragraph of Section 3.2. This straightforward fusion strategy is referred to as the MSBDN-S. The results in Table 4 show that the average PSNR value of the dehazed images by the MSBDN-S is 0.17 dB higher than that of the MSBDN, which demonstrate the benefits of exploiting the non-adjacent preceding features. To extract useful information from preceding features, we construct the MSBDN-M model by incorporating the back-projection technique [27], as an alternative to the proposed DFF, into the MSBDN. In the MSBDN-M model, the reconstruction errors of all the preceding features are mapped back to the estimated feature simultaneously. On the contrary, the pro-

posed DFF module adopts a progressive process to fuse one preceding feature at a time. Since extra layers are introduced by the DFF module, we construct an enhanced MSBDN (referred to as the MSBDN+) by adding two residual blocks into each residual group for fair comparisons. The MSBDN+ model has similar parameters as the MSBDN-M and MSBDN-DFF schemes. The results show that the back-projection algorithm [27] in the MSBDN-M model is less effective with a small improvement margin over the MSBDN+ scheme. It is noted that the MSBDN-DFF model outperforms the MSBDN-M scheme by a margin of 0.55 dB without introducing any extra layers, which shows the effectiveness of the proposed DFF modules. In addition, we apply the proposed DFF module to the PFFNet (referred to as the PFFNet-DFF). As shown in Table 4, the PFFNet-DFF achieves 3.73 dB performance improvement over the PFFNet, which demonstrates that the proposed DFF module can be easily deployed into other multi-scale architectures to improve the performance.

As shown in Figure 7(d), by remedying the spatial information and exploiting the preceding features, the MSBDN-DFF successfully removes the remaining haze in Figure 7(c) and recovers more details.

5. Conclusions

We propose an effective Multi-Scale Boosted Dehazing Network with Dense Feature Fusion for image dehazing. The MSBDN is constructed on an encoder-decoder architecture, where the boosted decoder is designed based on the SOS boosting strategy. The DFF module is designed on the back-projection scheme, which can preserve the spatial information and exploit the features from non-adjacent levels. The ablation studies demonstrate that the proposed modules are effective for the dehazing problem. Extensive evaluations show that the proposed model performs favorably against state-of-the-art methods on the image dehazing datasets.

Acknowledgements. H. Dong, X. Lei, X. Zhang and F. Wang are supported in part by National Major Science and Technology Projects of China (No. 2019ZX01008103), National Natural Science Foundation of China (NSFC) (No. 61603291), Natural Science Basic Research Plan in Shaanxi Province of China (No. 2018JM6057), and the Fundamental Research Funds for the Central Universities. J. Pan is supported in part by NSFC (Nos. 61872421, 61922043) and NSF of Jiangsu Province (No. BK20180471). M.-H. Yang is supported in part by NSF CAREER Grant (No. 1149783) and Gifts from Verisk, Adobe and Google.

References

- [1] Cosmin Ancuti, Codruta O Ancuti, and Radu Timofte. Ntire 2018 challenge on image dehazing: Methods and results. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 891–901, 2018.
- [2] Cosmin Ancuti, Codruta O Ancuti, Radu Timofte, and Christophe De Vleeschouwer. I-haze: a dehazing benchmark with real hazy and haze-free indoor images. In *International Conference on Advanced Concepts for Intelligent Vision Systems*, pages 620–631. Springer, 2018.
- [3] Codruta O Ancuti, Cosmin Ancuti, Radu Timofte, and Christophe De Vleeschouwer. O-haze: a dehazing benchmark with real hazy and haze-free outdoor images. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 754–762, 2018.
- [4] Dana Berman, Shai Avidan, et al. Non-local image dehazing. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1674–1682, 2016.
- [5] Bolun Cai, Xiangmin Xu, Kui Jia, Chunmei Qing, and Dacheng Tao. Dehazenet: An end-to-end system for single image haze removal. *IEEE Transactions on Image Processing*, 25(11):5187–5198, 2016.
- [6] Michael Charest, Michael Elad, and Peyman Milanfar. A general iterative regularization framework for image denoising. In *IEEE Conference on Information Sciences and Systems*, pages 452–457, 2006.
- [7] Chang Chen, Zhiwei Xiong, Xinmei Tian, and Feng Wu. Deep boosting for image denoising. In *European Conference on Computer Vision*, pages 3–18, 2018.
- [8] Chang Chen, Zhiwei Xiong, Xinmei Tian, Zheng-Jun Zha, and Feng Wu. Real-world image denoising with deep boosting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [9] Dongdong Chen, Mingming He, Qingnan Fan, Jing Liao, Liheng Zhang, Dongdong Hou, Lu Yuan, and Gang Hua. Gated context aggregation network for image dehazing and deraining. In *IEEE Winter Conference on Applications of Computer Vision*, pages 1375–1383, 2019.
- [10] Yunpeng Chen, Jianan Li, Huaxin Xiao, Xiaojie Jin, Shuicheng Yan, and Jiashi Feng. Dual path networks. In *Neural Information Processing Systems*, pages 4467–4475, 2017.
- [11] Yilun Chen, Zhicheng Wang, Yuxiang Peng, Zhiqiang Zhang, Gang Yu, and Jian Sun. Cascaded pyramid network for multi-person pose estimation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 7103–7112, 2018.
- [12] Shengyang Dai, Mei Han, Ying Wu, and Yihong Gong. Bilateral back-projection for single image super resolution. In *IEEE International Conference on Multimedia and Expo*, pages 1039–1042, 2007.
- [13] Akshay Dudhane, Harshjeet Singh Aulakh, and Subrahmanyam Murala. Ri-gan: An end-to-end network for single image haze removal. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1–10, 2019.
- [14] Deniz Engin, Anil Genç, and Hazim Kemal Ekenel. Cycle-dehaze: Enhanced cyclegan for single image dehazing. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 825–833, 2018.
- [15] Raanan Fattal. Single image dehazing. *ACM Transactions on Graphics*, 27(3):72–92, 2008.
- [16] Raanan Fattal. Dehazing using color-lines. *ACM Transactions on Graphics*, 34(1):13–31, 2014.
- [17] Damien Fourure, Rémi Emonet, Elisa Fromont, Damien Muselet, Alain Tremeau, and Christian Wolf. Residual conv-deconv grid network for semantic segmentation. In *British Machine Vision Conference*, 2017.
- [18] Yossi Gandelsman, Assaf Shocher, and Michal Irani. "double-dip": Unsupervised image decomposition via coupled deep-image-priors. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 11026–11035, 2019.
- [19] Andreas Geiger, Philip Lenz, and Raquel Urtasun. Are we ready for autonomous driving? the kitti vision benchmark suite. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3354–3361, 2012.
- [20] Clément Godard, Oisín Mac Aodha, Michael Firman, and Gabriel J. Brostow. Digging into self-supervised monocular depth prediction. In *IEEE International Conference on Computer Vision*, pages 3828–3838, 2019.
- [21] Shizhong Han, Zibo Meng, Ahmed-Shehab Khan, and Yan Tong. Incremental boosting convolutional neural network for facial action unit recognition. In *Neural Information Processing Systems*, pages 109–117, 2016.
- [22] Muhammad Haris, Greg Shakhnarovich, and Norimichi Ukita. Deep back-projection networks for super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1664–1673, 2018.
- [23] Kaiming He, Jian Sun, and Xiaoou Tang. Single image haze removal using dark channel prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12):2341–2353, 2011.
- [24] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016.
- [25] Gao Huang, Danlu Chen, Tianhong Li, Felix Wu, Laurens van der Maaten, and Kilian Q Weinberger. Multi-scale dense networks for resource efficient image classification. In *International Conference on Learning Representations*, 2018.
- [26] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4700–4708, 2017.
- [27] Michal Irani and Shmuel Peleg. Improving resolution by image registration. *Cvqip Graphical Models and Image Processing*, 53(3):231–239, 1991.
- [28] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.
- [29] Alexander Kirillov, Ross Girshick, Kaiming He, and Piotr Dollár. Panoptic feature pyramid networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 6399–6408, 2019.
- [30] Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid networks for fast and accurate super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 624–632, 2017.

- [31] Boyi Li, Xiulian Peng, Zhangyang Wang, Jizheng Xu, and Dan Feng. Aod-net: All-in-one dehazing network. In *IEEE International Conference on Computer Vision*, pages 4770–4778, 2017.
- [32] Boyi Li, Xiulian Peng, Zhangyang Wang, Jizheng Xu, and Dan Feng. End-to-end united video dehazing and detection. In *Thirty-Second AAAI Conference on Artificial Intelligence*, pages 7016–7023, 2018.
- [33] Boyi Li, Wenqi Ren, Dengpan Fu, Dacheng Tao, Dan Feng, Wenjun Zeng, and Zhangyang Wang. Reside: A benchmark for single image dehazing. *IEEE Transactions on Image Processing*, 28(1):492–505, 2018.
- [34] Jinjiang Li, Guihui Li, and Hui Fan. Image dehazing using residual-based deep cnn. *IEEE Access*, 6:26831–26842, 2018.
- [35] Runde Li, Jinshan Pan, Zechao Li, and Jinhui Tang. Single image dehazing via conditional generative adversarial network. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8202–8211, 2018.
- [36] Zhuwen Li, Ping Tan, Robby T. Tan, Danping Zou, Steven Zhiying Zhou, and Loong-Fah Cheong. Simultaneous video defogging and stereo reconstruction. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 4988–4997, 2015.
- [37] Zuoxin Li and Fuqiang Zhou. Fssd: feature fusion single shot multibox detector. *arXiv*, 2017.
- [38] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1132–1140, 2017.
- [39] Tsung-Yi Lin, Piotr Dollár, Ross Girshick, Kaiming He, Bharath Hariharan, and Serge Belongie. Feature pyramid networks for object detection. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2117–2125, 2017.
- [40] Xiaohong Liu, Yongrui Ma, Zhihao Shi, and Jun Chen. Grid-dehazenet: Attention-based multi-scale network for image dehazing. In *IEEE International Conference on Computer Vision*, pages 7314–7323, 2019.
- [41] Xing Liu, Masanori Suganuma, Zhun Sun, and Takayuki Okatani. Dual residual networks leveraging the potential of paired operations for image restoration. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 7007–7016, 2019.
- [42] Yang Liu, Jinshan Pan, Jimmy Ren, and Zhixun Su. Learning deep priors for image dehazing. In *IEEE International Conference on Computer Vision*, pages 2492–2500, 2019.
- [43] Kangfu Mei, Aiwen Jiang, Juncheng Li, and Mingwen Wang. Progressive feature fusion network for realistic image dehazing. In *Asian Conference on Computer Vision*, pages 203–215, 2018.
- [44] Peyman Milanfar. A tour of modern image filtering: New insights and methods, both practical and theoretical. *IEEE Signal Processing Magazine*, 30(1):106–128, 2013.
- [45] Mohammad Moghimi, Serge J Belongie, Mohammad J Saberian, Jian Yang, Nuno Vasconcelos, and Li-Jia Li. Boosted convolutional neural networks. In *British Machine Vision Conference*, pages 24.1–24.13, 2016.
- [46] Jinshan Pan, Sifei Liu, Deqing Sun, Jiawei Zhang, Yang Liu, Jimmy Ren, Zechao Li, Jinhui Tang, Huchuan Lu, Yu-Wing Tai, et al. Learning dual convolutional neural networks for low-level vision. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3070–3079, 2018.
- [47] Yanyun Qu, Yizi Chen, Jingying Huang, and Yuan Xie. Enhanced pix2pix dehazing network. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 8160–8168, 2019.
- [48] Joseph Redmon and Ali Farhadi. Yolov3: an incremental improvement. *arXiv*, 2018.
- [49] Wenqi Ren, Si Liu, Hua Zhang, Jinshan Pan, Xiaochun Cao, and Ming-Hsuan Yang. Single image dehazing via multi-scale convolutional neural networks. In *European Conference on Computer Vision*, pages 154–169, 2016.
- [50] Wenqi Ren, Lin Ma, Jiawei Zhang, Jinshan Pan, Xiaochun Cao, Wei Liu, and Ming-Hsuan Yang. Gated fusion network for single image dehazing. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3253–3261, 2018.
- [51] Yaniv Romano and Michael Elad. Boosting of image denoising algorithms. *Siam Journal on Imaging Sciences*, 8(2):1187–1219, 2015.
- [52] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical Image Computing and Computer-assisted Intervention*, pages 234–241, 2015.
- [53] Christos Sakaridis, Dengxin Dai, and Luc Van Gool. Semantic foggy scene understanding with synthetic data. *International Journal of Computer Vision*, pages 1–20, 2018.
- [54] Daniel Scharstein and Richard Szeliski. High-accuracy stereo depth maps using structured light. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 195–202, 2003.
- [55] Nathan Silberman, Derek Hoiem, Pushmeet Kohli, and Rob Fergus. Indoor segmentation and support inference from rgb-d images. In *European Conference on Computer Vision*, pages 746–760, 2012.
- [56] Robby T Tan. Visibility in bad weather from a single image. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.
- [57] Zhou Wang, Alan C Bovik, Hamid R Sheikh, Eero P Simoncelli, et al. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [58] Dong Yang and Jian Sun. Proximal dehaze-net: A prior learning-based deep network for single image dehazing. In *European Conference on Computer Vision*, pages 702–717, 2018.
- [59] Xitong Yang, Zheng Xu, and Jiebo Luo. Towards perceptual image dehazing by physics-based disentanglement and adversarial training. In *Thirty-second AAAI conference on Artificial Intelligence*, pages 7485–7492, 2018.
- [60] He Zhang and Vishal M Patel. Densely connected pyramid dehazing network. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3194–3203, 2018.
- [61] He Zhang, Vishwanath Sindagi, and Vishal M Patel. Multi-scale single image dehazing using perceptual pyramid deep network. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 902–911, 2018.

- [62] Xinyi Zhang, Hang Dong, Zhe Hu, Wei-Sheng Lai, Fei Wang, and Ming-Hsuan Yang. Gated fusion network for joint image deblurring and super-resolution. In *British Machine Vision Conference*, pages 153–165, 2018.
- [63] Yanfu Zhang, Li Ding, and Gaurav Sharma. Hazerd: an outdoor scene dataset and benchmark for single image dehazing. In *IEEE International Conference on Image Processing*, pages 3205–3209, 2017.
- [64] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2472–2481, 2018.
- [65] Jonathan Zung, Ignacio Tartavull, Kisuk Lee, and H Sebastian Seung. An error detection and correction framework for connectomics. In *Neural Information Processing Systems*, pages 6818–6829, 2017.