

“© 2019 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.”

LEARNING TRANSMISSION FILTERING NETWORK FOR IMAGE-BASED PM2.5 ESTIMATION

Yinghong Liao¹, Bin Qiu¹, Zhuo Su^{1*}, Ruomei Wang¹, Xiangjian He^{2,3}

¹School of Data and Computer Science, National Engineering Research Center of Digital Life, Sun Yat-sen University, Guangzhou, China

²Fujian Provincial Key Laboratory of Information Processing and Intelligent Control, Minjiang University, Fuzhou, China

³School of Computing and Communications, University of Technology Sydney
suzhuo3@mail.sysu.edu.cn

ABSTRACT

PM2.5 is an important indicator of the severity of air pollution and its level can be predicted through hazy photographs caused by its degradation. Image-based PM2.5 estimation is thus extensively employed in various multimedia applications but is challenging because of its ill-posed property. In this paper, we convert it to the problem of estimating the PM2.5-relevant haze transmission and propose a learning model called the transmission filtering network. Different from most methods that generate a transmission map directly from a hazy image, our model takes the coarse transmission map derived from the dark channel prior as the input. To obtain a transmission map that satisfies the local smoothness constraint without regional boundary degradation, our model performs the edge-preserving smoothing filtering as the refinement on the map. Moreover, we introduce the attention mechanism to the network architecture for more efficient feature extraction and smoothing effects in the transmission estimation. Experimental results prove that our model performs favorably against the state-of-the-art dehazing methods in a variety of hazy scenes.

Index Terms— PM2.5 estimation, edge-preserving smoothing, image dehazing, deep learning

1. INTRODUCTION

In cities which suffer from heavy air pollution, large amounts of fine particles are existed in the air and severely harm the human health. PM2.5 (particulate matter whose diameter is smaller than 2.5 microns) is used to measure the concentration of fine particles in the air and is a significant indicator of the

*Zhuo Su is the corresponding author (suzhuo3@mail.sysu.edu.cn). This research is supported by Natural Science Foundation of Guangdong Province (2016A030310202), National Natural Science Foundation of China (61672547), and Open Fund Project of Fujian Provincial Key Laboratory of Information Processing and Intelligent Control (Minjiang University, MJUK-F201701).




Fig. 1. The stacking histogram of the standard deviations of pixels in 40×40 patches from 5,642 transmission maps. The pattern is skewed right and over 90% of the pixel standard deviations are distributed within $[0, 20]$.

severity of air pollution. The more severe air pollution tends to be associated with the higher level of PM2.5. The resulted heavy haze can lead to the degradation of the image quality and the decrease of the image visibility. Hence, image-based PM2.5 estimation and its relevant removal of the image degradation resulted by PM2.5 embody great usefulness in multimedia applications like self-driving and video surveillance.

The atmospheric scattering model [1] is used to describe the imaging process of the haze in photographs as a result of the light scattering resulted by PM2.5,

$$I = Jt + A(1 - t), \quad (1)$$

where I is the hazy image, J is the radiance of the haze-free scene, A is the atmospheric light. t is the transmission whose definition is $t = e^{-\beta d}$, where β is the scattering coefficient, d is the scene depth. The transmission reflects the proportion of the light that could finally reach the camera after the light scattering. Therefore, the transmission could be used to estimate the level of the PM2.5. Each hazy map has a corresponding transmission map. The dehazed map J can be computed by Eq. 1 inversely, $J = (I - A)/t + A$, and the degradation that PM2.5 exerts on the hazy image can be eliminated.

1.1. Related Work

We first review some previous models in image dehazing for PM2.5 estimation, which could be divided into two categories, including prior-based models [2, 3] and learning-based models [4, 5].

Prior-based Models. These methods tend to capture various physical or statistical cues from hazy images. Such cues are exploited as some specific priors in the estimation of the transmission and the atmospheric light in Eq. 1. Fattal [1] leveraged the local no-correlation between transmission and surface shading functions to predict the transmission. In [3], Fattal observed the local smoothness constraint, that is the small difference of pixels in a local non-edge patch. As is shown in Fig. 1, most regions in the transmission map are smooth, which ensures that the image texture could be well preserved after the haze removal. Furthermore, most standard deviations of pixels in the histogram are larger than 1, which illustrates that most pixels are not piece-wise constant and the main regional boundaries are preserved. Thus the excessively smoothed transmission map could not satisfy the constraint as well. Based on this observation, Fattal further discovered the color-lines prior. However, it is time-consuming in the transmission estimation.

The dark channel prior (DCP) [2] was proposed by He *et al.* to derive a transmission map. Its estimated transmission map is coarse, thus Soft Matting [2] or Guided Filtering [6] were developed for the transmission refinement. However, the transmission maps obtained by these methods could not satisfy the local smoothness constraint. Errors consequently occur in the transmission estimation, which could result in the color distortion and the halo artifacts in the dehazed maps. Berman *et al.* [7] proposed the non-local prior called haze-line in RGB space. Prior-based models could exert fairly thorough haze removal on hazy maps but remain the issue of over-enhancement. As a result, artifacts and color distortion are produced in local image patches.

Learning-based Models. These methods explore the mapping relationships between the image features and the variables in Eq. 1. Tang *et al.* [8] integrated four haze-relevant feature maps as the inputs of the regression model. Convolutional neural network (CNN) is employed to the image dehazing tasks with the rise of deep learning. CNN could be trained on large-scale datasets to extract useful features for the haze removal without the physical or statistical priors. DehazeNet [9] was first proposed by Cai *et al.* for the transmission estimation. Ren *et al.* [10] presented MSCNN that conducts the global estimation and the refinement of the transmission separately. AOD-Net [11] was designed by Li *et al.* to derive the dehazed map directly. Zhang and Patel [12] proposed DCPDN that jointly estimates the transmission and the atmospheric light. Li *et al.* [13] directly generated haze-free maps by cGAN. Without the aid of the empirical priors, these networks highly depend on the training dataset and might ex-

hibit the low generalization capacity within different datasets. To address this issue, some researchers attempted to exploit the hand-crafted priors in deep models. Very recently, Ren *et al.* devised GFN [4] that could produce a haze-free image via the fusion of effective dehazed patches from three feature maps. And Liu *et al.* [5] integrated the priors and the data in learning the transmission map.

1.2. Contributions

In this work, we seek to perform transmission filtering in regard to the local smoothness constraint without regional boundary degradation in a deep learning fashion. Different from previous methods that generate the transmission map from hazy images directly, our model learns the mapping from the coarse map obtained by DCP to the refined map instead. This model could effectively leverage the hand-crafted priors and the learned data, and significantly increase the accuracy of the PM2.5 estimation. The attention mechanism [14] is introduced to the residual blocks in our network, which could increase the efficiency in feature extraction and smoothing refinement. Based on the network, our dehazing algorithm has excellent performance in both outdoor and indoor hazy scenes. Our main contributions are as follows:

- Present a mapping from the coarse transmission map to the refined transmission map that fits the local smoothness constraint for accurate PM2.5 estimation.
- Propose a novel network to generate the refined transmission map that satisfies the constraint and implement the corresponding haze removal algorithm.
- Introduce the attention mechanism to the network to optimize its performance in feature extraction and smoothing filtering.

2. THE PROPOSED METHOD

This section introduces the details of our proposed model of estimating the PM2.5-relevant transmission map and its corresponding dehazing algorithm. We first illustrate the shortcomings of existing filtering methods in the refinement of the transmission map in accordance with the local smoothness constraint. We then propose our learning filtering network for the smoothing refinement on the transmission map.

2.1. Problem Formulation

The dark channel prior [2] is that at least one pixel whose dark channel value approximates zero existed in a local patch of haze-free images,

$$D_w(x) = \min_{y \in \Omega_w(x)} \min_{c \in \{R, G, B\}} \frac{I^c(y)}{A^c} \rightarrow 0, \quad (2)$$

where I^c and A^c denote one of the color channels of the hazy map I and the atmospheric light A , respectively. $\Omega_w(x)$ represents the local patch which centers at x and has the size of




Fig. 2. (a) The given hazy input and the ground truth of the transmission map. (b)-(d) The dehazed map and the transmission map obtained by DCP [2], GF [6] and WMF [15].

$w \times w$. D_w is the dark channel of a local patch $\Omega_w(x)$, which is derived from the smallest value in color channels of pixels in this patch. DCP could generate a coarse transmission map \hat{t} based on Eq. 1,

$$\hat{t}(x) = 1 - \omega \min_{y \in \Omega_w(x)} \min_{c \in \{R, G, B\}} \frac{I^c(y)}{A^c}, \quad (3)$$

where ω denotes the constant parameter ranged at (0, 1] that is introduced to control the dehazing level and preserve the depth of the dehazed map. DCP excels at seizing the distribution information of haze in images and has a relatively satisfying dehazing effects. But its directly obtained transmission map \hat{t} is coarse and has fuzzy edges. Soft Matting (SM) [2] was proposed to refine the coarse map \hat{t} and Guided Filtering (GF) [6] was further advocated to replace soft matting for computational acceleration. GF could retain the necessary edges and smooth the non-edge regions of the coarse transmission map \hat{t} on the basis of the edge information provided by the hazy map I . The refined transmission map \tilde{t} could be computed through the edge preserving GF.

We illustrate the smoothing effects of the coarse transmission map refined by GF and Weighted Median Filter (WMF) [15] in Fig. 2, respectively. Note that the transmission maps processed by these representative filters are not subject to the local smoothness constraint. The transmission maps filtered by GF and WMF contain some local details that are detrimental to the texture of the dehazed results. Furthermore, the blurred edges in the transmission maps processed by GF result in the halo artifacts in the edge regions of the dehazed map, which is versus the local smoothness constraint. Therefore, in order to obtain more accurate results of PM2.5 estimation, a more specialized model is needed to refine the coarse transmission map in consideration of the constraint without regional boundary degradation.

2.2. Learning Transmission Filtering

For the effective utilization of the physical priors and the advanced solution to the edge-preserving transmission refinement with regard to the local smoothness constraint, we propose a CNN named transmission filtering network (TFN) to learn the filtering mapping \mathcal{F} ,

$$\tilde{t} = \mathcal{F}(\hat{t}; \mathcal{W}_{\mathcal{F}}), \quad s.t. \quad \mathcal{F}(\hat{t}; \mathcal{W}_{\mathcal{F}}) \sim \mathcal{S}(\beta, d), \quad (4)$$

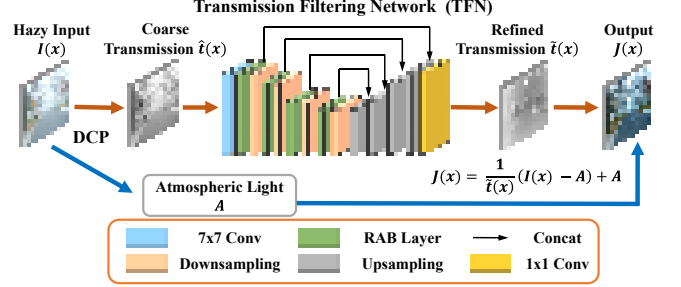


Fig. 3. The proposed dehazing procedure. Transmission Filtering Network (TFN) performs smoothing filtering in different scales on the coarse transmission map and produces the refined map that follows the local smoothness constraint.




Fig. 4. The architecture of Residual Attention Block (RAB).

where \hat{t} is the coarse transmission map, \tilde{t} is the refined transmission map and $\mathcal{W}_{\mathcal{F}}$ denotes the learnable parameters in the learning model. \mathcal{S} denotes the distribution of the real transmission map restrained by the local smoothness. β and d are the variables in the definition of the transmission t .

As the distribution \mathcal{S} lies in real images, we could approximate the dehazed map based on the refined transmission map \tilde{t} to the real haze-free map in practice so that the filtering mapping \mathcal{F} satisfy \mathcal{S} . Then we could obtain a filtered transmission map \tilde{t} that obeys the local smoothness constraint and remove the degradation caused by PM2.5 in an image.

2.3. Network Architecture

The proposed network performs the smoothing filtering to refine the coarse transmission map computed by DCP. We adopt the method in DCP to estimate the atmospheric light because of its fast speed. Our dehazing procedure is shown in Fig. 3.

TFN is a multiscale encoder-decoder network. Inspired by the attention mechanism [14], we design a Residual Attention Block (RAB) as a unit of the encoder. As is shown in Fig. 4, RAB is a residual architecture which converts the filtering mapping \mathcal{F} to an identity mapping. Each RAB learns the residual of the transmission map and smoothes the coarse map progressively. The encoder containing several RABs is named RAB Layer. RAB Layers are followed by the downsampling operations to extract features in different scales and could depict features more precisely.

Initially, two 3×3 convolutional layers preprocess the input features and extract C feature maps. In the operation of the channel attention, a global pooling operation is performed




Fig. 5. Attention mechanism in TFN. The lower training error and the more accurate estimation are obtained.

in spatial dimension to obtain a $1 \times 1 \times C$ feature vector. The size of the vector is further compressed to $1 \times 1 \times C/r$ by a 1×1 convolutional layer to aggregate more useful features, where r denotes the compression ratio. The vector is then recovered to a $1 \times 1 \times C$ channel attention vector by another 1×1 convolutional layer. The feature map before the pooling is multiplied by the channel attention vector to obtain several channel attention maps. In the next operation of the spatial attention, another max pooling operation is conducted along the channel axis of the channel attention maps, and some same-size spatial attention maps are derived. After the processing of a 3×3 convolutional layer, the spatial maps are in the multiplication with the channel maps. The final attention maps and input feature maps are in the element-wise summation.

2.4. Loss Function

We denote the output dehazed result by \mathcal{J} . ℓ_1 loss facilitates the feature selection in the model optimization, hence we adopt it as one of the loss functions. In our work, ℓ_1 loss L_{ℓ_1} is defined as below,

$$L_{\ell_1} = \frac{1}{N} \sum_{i=1}^N \|\mathcal{J}(I_i; \tilde{t}_i, A_i) - J_i\|_1, \quad (5)$$

where N is the number of training samples in every batch, I_i , \tilde{t}_i , A_i and J_i denotes the i -th hazy image, the refined transmission map, the estimated atmospheric light and the expected haze-free map, respectively.

Multiscale structural similarity index (MS-SSIM) is conceptually related to the brightness, the contrast and the structure of images. And MS-SSIM loss contributes to the edge preserving and denoising. To further enhance the dehazing performance of our model, we introduce MS-SSIM loss $L_{\text{MS-SSIM}}$ in training which is defined as follows,

$$L_{\text{MS-SSIM}} = 1 - \frac{1}{N} \sum_{i=1}^N \text{MS-SSIM}(\mathcal{J}(I_i; \tilde{t}_i, A_i), J_i), \quad (6)$$

where $\text{MS-SSIM}(\cdot)$ represents the MS-SSIM of two images.

Overall, the total loss function \mathcal{L} is defined as follows,

$$\mathcal{L} = \alpha L_{\text{MS-SSIM}} + \beta L_{\ell_1}, \quad (7)$$

where α and β are the positive weights of $L_{\text{MS-SSIM}}$ and L_{ℓ_1} , respectively.

3. EXPERIMENTAL RESULTS

In this section, we provide the experimental results to demonstrate the effectiveness in transmission filtering and the usefulness in the PM2.5-relevant dehazing task of our model.

3.1. Training and Testing Datasets

RESIDE dataset [16] contains both indoor and outdoor hazy images. 13,000 indoor images from Indoor Training Set (ITS) in RESIDE-Standard dataset and 7,665 outdoor images from Outdoor Training Set (OTS) in RESIDE- β dataset are selected to compose a mixed training dataset named MIX. ITS dataset is based on NYU2 database and its generated transmission maps are closer to the reality in depth, thus ITS images are given a larger proportion in MIX. 500 outdoor images from Synthetic Objective Testing Set (SOTS) in RESIDE-Standard are picked to form a testing dataset, Outdoor. Besides, the remaining 990 indoor images in ITS are used as another testing dataset, Indoor.

3.2. Implementation Details

Our model is implemented in PyTorch. The patch size of each image is set as 512×512 . Adam optimizer is used in training and the batch size is 5. The initial learning rate is set as 10^{-3} . The model is trained with the learning rate as 10^{-4} between 11,000 and 70,000 iterations. After that, the learning rate would decrease by 0.1 after every 20,000 iterations. The network could not search suitable parameters for optimization initially, so in the first 4,000 iterations, it is trained to approximate the transmission maps on ITS. The network is then trained with the haze-free maps in MIX as the target. The weights α , β in Eq. 7 are set as 0.2 and 0.8, respectively.

3.3. The Effectiveness of Attention Mechanism

To demonstrate the effectiveness of the attention mechanism, our model is trained with and without the attention mechanism, respectively. The relevant training error curves are displayed in the line chart above in Fig. 5, where the training errors are the numeric results of Eq. 7. The curve with attention has a smaller error after training 40 epochs, which illustrates that the attention mechanism could enhance the training convergence of the network and enable it to derive a more accurate dehazed map. A real hazy map and its corresponding transmission maps with and without the attention mechanism are also presented below the line chart in Fig. 5. Note that the transmission map refined with the attention mechanism has smoother regions and contains more clear depth information.

3.4. The Effectiveness in Transmission Filtering

A challenging hazy image is used to prove the effectiveness of our model in edge-preserving transmission filtering in regard to the local smoothness constraint. The refined and dehazed results by several representative smoothing filters are compared with ours, including SM [2], GF [6], Rolling Guidance




Fig. 6. The transmission map refined by TFN satisfies the local smoothness where the significant regional boundaries are preserved and the transitions of the smoothed regions are natural. The dehazed result achieves the best value of the SSIM.




Fig. 7. Visual comparison with state-of-the-art dehazing methods on synthetic images (SSIM/PSNR).

Table 1. Quantitative evaluation with some state-of-the-art methods on the synthetic datasets (SSIM/PSNR), where the best and the second best numeric values are marked in red and blue, respectively.

Testset	DCP [2]	DCP-G [2, 6]	MSCNN [10]	DehazeNet [9]	AOD [11]	GFN [4]	DCPDN [12]	cGAN [13]	Ours
Outdoor	0.76/15.98	0.80/16.69	0.83/19.48	0.86/22.58	0.90/20.45	0.83/21.37	0.85/23.58	0.88/ 23.91	0.93/25.16
Indoor	0.79/16.29	0.83/17.19	0.81/18.50	0.85/20.43	0.84/18.85	0.87/21.67	0.92/26.20	0.93/26.36	0.95/26.42



Fig. 8. Comparison with some state-of-the-art dehazing methods on real hazy images. Zoom in for the more clear display.

Filter (RGF) [17], WMF [15], and L_0 smoothing (L_0) [18]. As is shown in Fig. 6, the transmission maps processed by S-M, GF, RGF, and WMF contain some fine-scale details. The result computed by L_0 is shown to be over-smoothed, where the smoothed regions are piece-wise constant and some main regional boundaries are diminished. In consequence, the obvious haze remains in the dehazed result. In contrast, the refined transmission map by TFN has higher-quality smoothing effects where the significant regional edges are retained and the transitions of the smoothed regions are natural. The value of structural similarity index (SSIM) demonstrates that the generated filtered map by TFN is structurally similar to the

desired one and fits the local smoothness constraint as well.

3.5. The Evaluation on Synthetic Datasets

To effectively evaluate the haze removal effects of our model, TFN is given the comparison with some state-of-the-art dehazing methods on the synthetic testing datasets, Outdoor and Indoor. The selected methods are DCP with GF (DCP-G) [2, 6], MSCNN [10], DehazeNet [9], AOD-Net (AOD) [11], GFN [4], DCPDN [12], and cGAN [13], respectively. The numeric and visual evaluation are provided, respectively. The metrics adopted in the numeric evaluation are SSIM and peak signal-to-noise ratio (PSNR). The results of

the comparison are displayed in Table 1. The PSNR and the SSIM of our model on Outdoor are both higher than those of these methods. On Indoor, the SSIM and the PSNR of our model exceeds all the values of these models. Overall, the numeric values of the SSIM and PSNR of our model are close, which proves its superiority of the generalization capacity.

The visual comparison between state-of-the-art models and ours is conducted on an example set of hazy and haze-free images in Outdoor in Fig. 7. The results obtained by DCP-G and MSCNN have some red halo artifacts in the sky regions. The river regions in the images dehazed by GFN and cGAN are dark. The scenes on the banks in the image of DCPDN are blurred. By comparison, our result closely resembles the ground truth and is free of similar issues. The corresponding values of the SSIM and PSNR are measured. Our model has the best results in both metrics, which demonstrates its better preservation of the structure and the texture in images.

3.6. The Evaluation on Real Images

Our model is applied in real hazy scenes for the qualitative evaluation of its dehazing capacity. Our model is compared with some state-of-the-art methods on three challenging real hazy images in Fig. 8, including DCP-G [2, 6], MSCNN [10], AOD [11], GFN [4], DCPDN [12], and cGAN [13]. The color distortion is evident in dehazed images of DCP-G. The haze remains in the second images produced by MSCNN and AOD. GFN could not well remove the haze in the third one. Obvious artifacts are spotted in the glazed tiles of the palace roofs in the second map of DCPDN. Distortion occurs in the granite handrail regions of the second one of cGAN. In contrast, our results is higher-quality without the shortcomings of the remaining haze, the color distortion, and the artifacts.

4. CONCLUSIONS

In this paper, we have proposed a CNN that could refine the transmission map and solve the issue of estimating P-M2.5. Compared with the previous smoothing filters, our model could generate a transmission map that satisfies the local smoothness constraint and retains smooth transitions in local regions. The corresponding degradation caused by P-M2.5 could be eliminated by our method effectively. The attention mechanism introduced in the proposed model could enhance the performance of smoothing patches in the transmission map. The experimental results have demonstrated that our model has excellent capacity in the PM2.5-relevant transmission estimation and superior dehazing performance in hazy images of various scenes.

5. REFERENCES

- [1] R. Fattal, “Single image dehazing,” *ACM TOG*, vol. 27, no. 3, pp. 72:1–72:9, 2008.
- [2] K. He, J. Sun, and X. Tang, “Single image haze removal using dark channel prior,” *IEEE TPAMI*, vol. 33, no. 12, pp. 2341–2353, 2011.
- [3] R. Fattal, “Dehazing using color-lines,” *ACM TOG*, vol. 34, no. 1, pp. 13:1–13:14, 2014.
- [4] W. Ren, L. Ma, J. Zhang, J. Pan, X. Cao, W. Liu, and M. Yang, “Gated fusion network for single image dehazing,” in *IEEE CVPR*, 2018, pp. 3253–3261.
- [5] R. Liu, X. Fan, M. Hou, Z. Jiang, Z. Luo, and L. Zhang, “Learning aggregated transmission propagation networks for haze removal and beyond,” *IEEE TNNLS*, pp. 1–14, 2018.
- [6] K. He, J. Sun, and X. Tang, “Guided image filtering,” *IEEE TPAMI*, vol. 35, no. 6, pp. 1397–1409, 2013.
- [7] D. Berman, T. Treibitz, and S. Avidan, “Non-local image dehazing,” in *IEEE CVPR*, 2016, pp. 1674–1682.
- [8] K. Tang, J. Yang, and J. Wang, “Investigating haze-relevant features in a learning framework for image dehazing,” in *IEEE CVPR*, 2014, pp. 2995–3002.
- [9] B. Cai, X. Xu, K. Jia, C. Qing, and D. Tao, “DehazeNet: An end-to-end system for single image haze removal,” *IEEE TIP*, vol. 25, no. 11, pp. 5187–5198, 2016.
- [10] W. Ren, S. Liu, H. Zhang, J. Pan, X. Cao, and M. Yang, “Single image dehazing via multi-scale convolutional neural networks,” in *ECCV*, 2016, pp. 154–169.
- [11] B. Li, X. Peng, Z. Wang, J. Xu, and D. Feng, “AOD-Net: All-in-one dehazing network,” in *IEEE ICCV*, 2017, pp. 4780–4788.
- [12] H. Zhang and V. Patel, “Densely connected pyramid dehazing network,” in *IEEE CVPR*, 2018, pp. 3194–3203.
- [13] R. Li, J. Pan, Z. Li, and J. Tang, “Single image dehazing via conditional generative adversarial network,” in *IEEE CVPR*, 2018, pp. 8202–8211.
- [14] S. Woo, J. Park, J. Lee, and I. So Kweon, “CBAM: Convolutional block attention module,” in *ECCV*, 2018, pp. 3–19.
- [15] Q. Zhang, L. Xu, and J. Jia, “100+ times faster weighted median filter (WMF),” in *IEEE CVPR*, 2014, pp. 2830–2837.
- [16] B. Li, W. Ren, D. Fu, D. Tao, D. Feng, W. Zeng, and Z. Wang, “Benchmarking single-image dehazing and beyond,” *IEEE TIP*, vol. 28, no. 1, pp. 492–505, 2019.
- [17] Q. Zhang, X. Shen, L. Xu, and J. Jia, “Rolling guidance filter,” in *ECCV*, 2014, pp. 815–830.
- [18] L. Xu, C. Lu, Y. Xu, and J. Jia, “Image smoothing via L_0 gradient minimization,” *ACM TOG*, vol. 30, no. 6, pp. 174:1–174:12, 2011.