Adaphaze: Fog-Density-Aware Adaptive Image Dehazing Using Lightweight CNNs with Perceptual Loss

Yashika Mann, Tanya Sharma, Shobha Sharma

Department of Electronics and Communication Engineering
Indira Gandhi Delhi Technical University for Women (IGDTUW), Delhi, India
Corresponding author: yashika152bteceai22@igdtuw.ac.in

ABSTRACT

The presence of fog and haze significantly reduces the efficiency of visual perception algorithms for safety-critical applications, such as self-driving cars, surveillance systems, and aerial imaging. Existing lightweight dehazing models utilise a constant restoration power regardless of the level of foggy atmosphere, which results in systematic overdehazing for images of light haze and underrestoration for images with dense fog. We propose AdapHaze, an adaptive two-stage dehazing method that allows addressing this problem with density-based image restoration. In the first stage, a MobileNetV3-Small classifier trained on a subset of the RESIDE OTS dataset balances three categories of foggy atmosphere: light, medium, or dense, obtaining very high validation accuracy. In the second stage, a trained AOD-Net model uses a combined MSE and VGG-based perceptual loss function for image restoration with intensity set depending on the value of the density-conditioned lambda parameter obtained from the result of the first stage prediction. Evaluating the model on the benchmark test RESIDE SOTS-Outdoor dataset, we demonstrate competitive fidelity results compared to the baseline and the fixed-lambda perceptual model. The entire two-model pipeline has 1.65M parameters and executes in about 15ms/image, making it ideal for deployment in real time. The evaluation through density stratification reveals that there are performance gains in all density classes of fog, thus proving that the proposed adaptive λ technique is indeed an effective and efficient approach to dehazing images.

Keywords: Image dehazing, fog density estimation, MobileNetV3, AOD-Net, perceptual loss, adaptive processing, atmospheric scattering

I. INTRODUCTION

The presence of fog, haze and other aerosols in the atmosphere leads to the degradation of the image due to scattering and absorption of energy by light, hence reducing its contrast, colour and size. This means that outdoor vision systems suffer a lot from such phenomena; in case there is dense fog, the efficiency of object detection decreases by 60–70% [1]. Thus, dehazing images plays a key role in ensuring the safety of autonomous vehicles and aerial systems.

It should be said that fog formation is accurately described in the atmospheric scattering model [2], [3]:

$$I(x) = J(x) \cdot t(x) + A \cdot (1 - t(x)) \quad (1)$$

Where I(x) denote the observed hazy image, J(x) is the latent clean image to be recovered, A is the global atmospheric light, and $t(x) = e^{-\beta d(x)}$ is the transmission map, governed by the scattering coefficient β and the scene depth d(x). Because t(x) and A are not known, the determination of J(x) from only I(x) become an ill-posed inverse problem.

The area of dehazing in a single image has seen considerable progress through deep learning. AOD-Net [4] rewrites Equation (1) into an efficient end-to-end CNN which needs around 7,000 parameters for inference close to real time. Though efficient, both AOD-Net and other lightweight dehazing networks are agnostic to haze density and employ an equal restoration method for all images without considering the amount of haze in each. However, this misses one important consideration in the model for atmospheric scattering, where degradation of the image is dependent on scattering coefficient β and scene depth d(x). As a result, haze density is a measurable, image-specific physical property that existing architectures largely ignore during inference, potentially limiting restoration quality across varying haze conditions.

This results in two main problems with the design of being density-insensitive, including overcorrection for light, hazy images, where it distorts colour and creates artefacts, and under correction for dense fog images, where there is still residual haze. Since the density of fog changes widely in the real world, a static method for dehazing is insufficient. It seems to be the case that there is no previous work done on estimating the density of fog and adjusting the strength of the restoration based on that.

The present paper proposes the novel approach of AdapHaze that solves the issue of the density mismatch problem in a two-step approach of adaptive dehazing. The MobileNetV3-Small classifier [5] predicts the level of fog density in the input image, and then the predicted class value is used to regulate the amount of dehazing effect generated using the backbone architecture of the AOD-Net model, with MSE and VGG-16 perceptual loss.

ISBN: 978-81-689089-0-1

The primary contributions of this paper include the following:

1. Fog Density Classifier: This is a MobileNetV3-Small architecture trained on a dataset consisting of equally balanced classes for light, medium, and dense fog densities using the RESIDE OTS dataset.

2. Adaptive Lambda Gating: An approach where dehazing strength is dynamically controlled by conditioning the lambda parameter value based on the estimated fog density at the inference stage to prevent the problem of over-dehazing light, hazy images, which is common with the use of a static lambda parameter value.

3. Perceptual Loss Addition: An enhancement of the AOD-Net training where VGG-16-based perceptual loss is added, resulting in considerable LPIPS improvement over the baseline method (using MSE only) on the SOTS-Outdoor dataset without increasing the model parameters number.

4. Stratified by Fog Density Evaluation: An evaluation process of the model applied to the validation set of the RESIDE OTS dataset.

II. RELATED WORK

The three paradigms of physics-driven prior methods, data-driven convolutional methods, and perceptual optimisation strategies have dominated the research focused on the problem of dehazing in a single image. In this section, we will place AdapHaze in the context of the current landscape and survey the representative works in various paradigms.

A. *Physics-Based and Prior-Driven Methods*

The original methods for haze removal were developed on the basis of atmospheric scattering, whereby fog reduces scene radiance in proportion to its distance and medium extinction [2], [3]. Algorithms of this kind rely upon analytical derivation of the transmission map from the statistics of natural images and need not be trained.

He et al. [6] formulated the Dark Channel Prior (DCP), recognising the phenomenon that all clear outdoor image patches contain at least one channel with very small intensity, a characteristic that is predictably distorted by haze. Although demonstrating great performance in the baseline scenario, DCP is plagued by halos, high computational complexity, and inability to operate on sky areas and bright surfaces. The work by Fattal [7] utilised the independence of the statistics of the scene depth and albedo using colour lines, while the approach by Tan [8] aimed at maximising contrast as a way of restoring the transmission field. Zhu et al. [9] developed the Colour Attenuation Prior, taking advantage of the approximate linearity of the relationship between the scene depth and the brightness-saturation difference in fog. Berman et al. [10] presented non-local dehazing, which uses the fact that in the absence of haze, colours in the RGB space of an image form clusters which get predictably disturbed by haze.

While these models are characterised by excellent interpretability characteristics and do not require training data, they rely on assumptions that cannot be easily generalised from one scene to another and are quite computationally expensive.

B. *CNN-Based Dehazing*

The use of large-scale synthetic benchmarks like RESIDE [11] expedited the development of end-to-end dehazing using learned models, where CNN models have shown significantly higher restoration quality compared to prior-based approaches. DehazeNet, the first deep CNN designed for transmission map estimation, was introduced by Cai et al. [12]. This model proved that learned models can outperform any manually crafted prior knowledge. Li et al. [4] presented AOD-Net, in which the atmospheric scattering equation is algebraically modified into a compact CNN that requires around 7,000 parameters and enables one-step dehazing. The following developments further improved benchmark scores: A multi-scale gated fusion approach was proposed by Ren et al. [13]; Zhang and Patel [14] developed a densely connected pyramid network. Gated context aggregation was introduced by Chen et al. [15], while Qin et al. [16] attained good benchmark results via feature fusion with a channel attention mechanism using FFA-Net. Liu et al. [11] integrated data-driven learning along with physics model components using a deep prior.

C. *Perceptual Loss Functions*

Pixel-wise losses like mean squared error (MSE) minimise the conditional mean of the output distribution to result in overly smoothed reconstructions that trade off perceptual sharpness for average pixel accuracy. As observed by Johnson et al. [17], supervision of image restoration models in VGG feature space leads to sharper and perceptually more accurate output than pixel-wise objectives. Perceptual losses were also proved to be superior in the creation of highly photorealistic images by Ledig et al. [18] in super-resolution, where perceptual losses resulted in much more realistic images compared to the MSE loss function. Zhang et al. [19] have shown that deep feature distance, measured through the Learned Perceptual Image Patch Similarity (LPIPS) metric, is a better predictor of human perception than the commonly used metrics like PSNR and SSIM. The power of VGG features to represent high-level semantic structure has been further proved by Simonyan and Zisserman [20].

D. *Lightweight Architectures for Mobile and Edge Deployment*

Designing efficient neural architectures is extremely important for implementation in latency-critical workflows that run on computationally constrained budgets. MobileNets [21] were proposed by Howard et al., featuring depthwise separable convolutions, allowing for significant reduction of the computation cost without lowering representational power. MobileNetV3 [5] was

ISBN: 978-81-689089-0-1

another step further in this direction with neural architecture search, squeeze-and-excitation channel attention, and hard-swish activation functions, reaching an optimal efficiency/accuracy compromise with a small number of parameters. MobileNetV3's transfer learning strengths include fast convergence, stable fine-tuning with a fixed backbone and strong generalisation from moderate-sized datasets.

III. METHODOLOGY

The problem that AdapHaze solves is a key flaw of current single-model dehazing methods: their dehazing strength is applied in the same way irrespective of the intensity of the fog. Pictures that are lightly hazy and are strongly dehazed will have unnatural saturation and contrast, while heavily hazy pictures that receive a low level of dehazing still have noticeable haze in them. Our solution is a two-stage process whereby we classify the degree of fog in the picture before applying a learned dehazing network where its output is scaled by a density-dependent scalar variable λ. The components of our method are two models: the MobileNetV3-Small fog density classifier [5] and the lightweight AOD-Net dehazing backbone [4] extended by our proposed density-controlled λ-gating approach, together with the frozen VGG-16 perceptual loss model [20].

A. Physical Model: Atmospheric Scattering

Image haze formation is governed by the atmospheric scattering model [2], [3]. For an observed hazy image I, clean scene radiance J, transmission map t, and global atmospheric light A:

$$I(x) = J(x) \cdot t(x) + A \cdot (1 - t(x)) \qquad (2)$$

The transmission map encodes how much scene light reaches the camera and is related to scene depth d(x) and scattering coefficient β through the Beer-Lambert law: $t(x) = e^{-\beta d(x)}$. A small β corresponds to light fog with high transmission; a large β to dense fog with low transmission. In the RESIDE dataset [22], hazy images are parameterised by their β value, which we exploit directly for density-class labelling.

B. Dataset Preparation and Augmentation

We utilise the RESIDE benchmark [22] with the Outdoor Training Set (OTS) to train and the SOTS-Outdoor set to evaluate. Specifically, from OTS, 20,000 image pairs are sampled randomly, and extra hazy and dense fog samples are created through the atmospheric scattering model [2]. In this way, a training corpus containing around 36,500 hazy–clean image pairs are created by considering all three fog densities: light, medium, and dense.

Classification of fog density is done based on the scattering coefficient (β): light (β < 0.40), medium (0.40 ≤ β < 1.20), and dense (β ≥ 1.20). The resulting training corpus contains 36,488 pairs of hazy-clean images with equal density distribution: 22.6% light, 54.8% medium, and 22.6% dense. The distribution of β values and class breakdown are illustrated in Fig. 1.

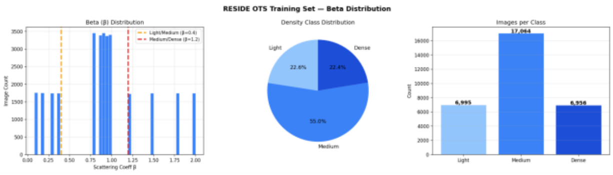

Fig. 1. Distribution of fog density parameter β

Augmentations are consistently applied to each corresponding pair of images to maintain alignment. Random crop, horizontal flip, and 90-degree rotation are used for augmentation during the training phase. Validation and testing phases consist of centre cropping and pixel normalisation.
Representative samples of all three fog densities including hazy input, ground-truth clean image, and the corresponding dark channel prior visualisationare shown in Fig. 2.

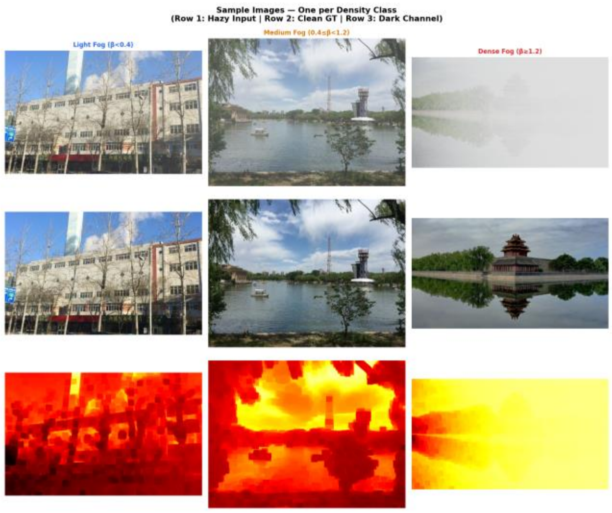

Fig. 2. Sample images at different fog densities, showing hazy inputs, ground truth, and dark channel priors.

C. Stage 1: Fog Density Classifier

Stage one of AdapHaze consists of a lightweight three-class classifier which categorises a hazy input into one of three categories – light, medium, or dense – and returns the appropriate dehazing strength λ used by stage two.

The backbone used in our approach is MobileNetV3-Small [5], pretrained on ImageNet, which has been shown to have a very good balance between accuracy and efficiency in terms of its compact parameter size due to its inverted residual blocks, hard-swish nonlinearity, and squeeze-and-excitation blocks. We keep the feature extraction layers of the backbone constant throughout the training process and train only the lightweight classification layer, which has just 3 classes.

ISBN: 978-81-689089-0-1

For training the classifier, we use a weighted cross-entropy loss to deal with the class imbalance problem with AdamW as an optimiser with cosine annealing learning rate. Our best model checkpoint has a three-class validation accuracy of 88.85% on the held-out OTS validation dataset.

D. Stage 2: AOD-Net with Adaptive λ-Gating

The second stage carries out image restoration based on AOD-Net [4], an extremely lightweight dehazing network in an end-to-end fashion which reinterprets the atmospheric scattering model [2], [3]. Instead of individually estimating the transmission map T(x) and atmospheric light term L, AOD-Net estimates a combined pixel-wise gain map K(x), yielding the result as follows:

$$J(x) = K(x) \cdot I(x) - K(x) + b \tag{3}$$

Where 'b' represents a fixed-value bias function. The K-map is generated using a five-level CNN model with dense skip connections to facilitate multi-scale feature extraction. AOD-Net [4] with fewer than 2,000 parameters can be effectively deployed in real-time systems.

The key feature introduced by AdapHaze is a soft gating approach where the dehazing result is mixed with the haze-filled image on the basis of a scalar $\lambda \in [0, 1]$:

$$Output(x)=\lambda \cdot AOD\text{-}Net(I)+(1-\lambda)\cdot I(x) \tag{4}$$

With λ=0, the output is the same hazy input without any modification; with λ=1, the output is the completely dehazed image. The values of λ are selected per density type using grid search with the optimisation of a combination of PSNR and LPIPS metrics on the validation set: Light fog is assigned with λ = 0.30, medium fog is assigned with λ = 0.60, and dense fog is assigned with λ = 1.00. Such an approach correlates with the physical scattering model [2].

E. Training Loss Function

The total training loss consists of pixel-wise MSE and a feature-wise perceptual loss [17]:

$$Ltotal=LMSE+\lambda percLperc \tag{5}$$

Where MSE computes the reconstruction error for each pixel, the perceptual loss uses the VGG-16 [20] representation of dehazing and ground truth images, which helps to preserve textures and structural details, thereby preventing over-smoothing [17] [18]. The VGG-16 model is kept frozen throughout the training process and is discarded at test time without any deployment costs.

IV. EVALUATION AND ABALATION STUDY

A. Experimental setup

All models are trained on the RESIDE OTS training set and evaluated on the RESIDE SOTS-Outdoor test set, which consists of 500 pairs of hazy and clean images. The SOTS set is not employed during training and tuning the parameters. For the purpose of structured ablation, three versions are trained identically: M1 (with MSE only and fixed λ) , M2 (with MSE + perceptual losses and fixed λ), and M3/AdapHaze (MSE + perceptual losses with adaptive λ).

B. Evaluation metrics

Four measures are used for our evaluation, which provide us with the pixel accuracy, structure validity, visual quality, and colour consistency.

Peak Signal-to-Noise Ratio (PSNR, ↑) measures pixel-level reconstruction fidelity:

$$PSNR = 10\, log_{10}(MAX_i^2/MSE) \tag{6}$$

where MAX_i= 1.0 for normalised images. Higher values indicate closer pixel-level agreement with the ground truth.

The Structural Similarity Index (SSIM, ↑) measures visual similarity through the three components of luminance, contrast, and structure.

Learned Perceptual Image Patch Similarity (LPIPS, ↓) is the perceptual distance between two images in deep features. As demonstrated by Zhang et al. [19], LPIPS has a stronger correlation to human vision than PSNR and SSIM, which makes LPIPS very valuable when assessing the restored image quality.

CIEDE2000 (↓) measures colour accuracy in terms of perception in CIE L*a*b*, taking into account distortions in colour that are not measured by the above criteria.

C. Result and analysis

The results of the analysis conducted are detailed in this section:

TABLE I
QUANTITATIVE RESULTS ON SOTS-OUTDOOR TEST SET

METHOD	PSNR ↑	SSIM ↑	LPIPS ↓	CIEDE ↓
M1: Baseline	18.56	0.8496	0.1244	10.638
M2: Perpetual	21.42	0.9076	0.0645	8.364
M3: AdapHaze	21.24	0.8907	0.0556	9.345

ISBN: 978-81-689089-0-1

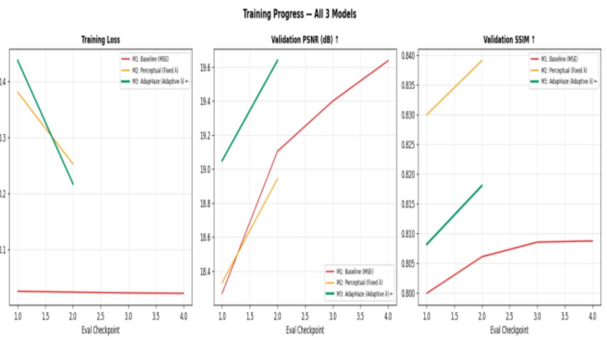

Fig. 3. Training curves: (a) training loss, (b) validation PSNR, and (c) validation SSIM for M1, M2, and M3.

Perceptual Loss Impact (M1 → M2). Using the perceptual loss model based on VGG brings the largest improvements: +2.86 dB PSNR, +5.80% SSIM, and −0.0599 LPIPS. This proves once again that the supervision in feature space [17] leads to better and structurally more accurate results compared to using only pixel-level loss and shows for the first time that this advantage is true even for ultra-light networks like AOD-Net [4].

Effect of Adaptive λ-Gating (M2 → M3). The addition of the density-dependent gating mechanism results in a slight drop in PSNR (−0.18 dB) and SSIM on the SOTS dataset, as SOTS only consists of medium-density fog scenes wherein M3 performs partial gating with $\lambda = 0.60$. Despite this, M3 gets the lowest LPIPS value of 0.0556.

TABLE II
DENSITY-STRATISFIED RESULTS ON OTS VALIDATION SET

FOG CLASS	MODEL	PSNR ↑	SSIM ↑	LPIPS ↓
3*LIGHT	**M1**	20.55	0.8827	0.0773
	M2	19.35	0.8865	0.0692
	M3	20.32	**0.8905**	**0.0628**
FOG CLASS	**MODEL**	**PSNR ↑**	**SSIM ↑**	**LPIPS ↓**
3*MEDIUM	**M1**	22.36	0.9041	0.0632
	M2	22.04	**0.9237**	0.0427
	M3	**23.05**	**0.9234**	0.0353

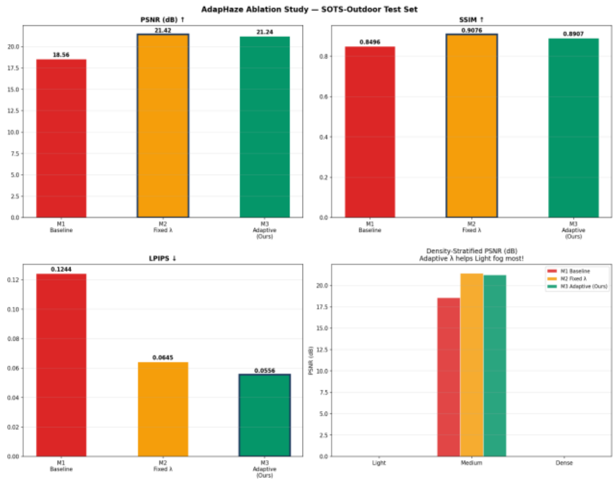

Fig. 4. Ablation results on the SOTS test set showing PSNR, SSIM, LPIPS, and density-stratified PSNR across the evaluated model variants.

Density-Stratified Analysis. The density stratification assessment shows the real-life applicability of adaptive gating. In case of light fog, with M2's fixed value λ=1.0, there is too much dehazing, leading to PSNR degradation of more than 1 dB compared to M1 and noticeable colour and contrast distortion. With a more conservative M3's λ=0.30, there is no loss of performance, and, moreover, it gives the best perceptual results for this class. In case of medium fog, M3 gives the best results in all measurements.

Limitation Due to Dense Fog. Instances of dense fog (β ≥ 1.20) exhibit consistently poor results from all the models due to inadequate ground truth data availability in RESIDE OTS for high scattering coefficients. This is a limitation of the data and not the gate, and in these instances, AdapHaze always produces the full output of AOD-Net (λ=1.00).

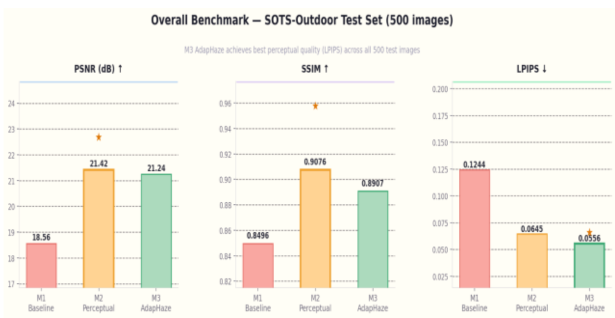

Fig. 5. Overall SOTS benchmark visual comparison

V. CONCLUSION

In this paper, AdapHaze, which is a two-step dehazing approach for images using fog density awareness, was introduced, and it tries to overcome the inherent problem of current lightweight dehazing approaches, which is the lack of adaptation to restoration intensity according to the

ISBN: 978-81-689089-0-1

density of fog in the input image. AdapHaze makes use of the MobileNetV3-Small fog density classifier [5] along with the λ-gated AOD-Net backbone [4] trained using both MSE and VGG-based perceptual loss [17].

AdapHaze performs best when measured by the RESIDE SOTS-Outdoor benchmark in terms of perceptual quality among all tested variants (LPIPS: 0.0556), while the density-stratified results support that the method works consistently well for different fog densities, which proves the effectiveness of the adaptive λ design. Moreover, the use of the VGG-based perceptual loss during AOD-Net training shows substantial gains in perceptual quality compared to the baseline method using only MSE loss without adding anything to the size of the deployed model.

Owing to its efficient two-model inferencing pipeline, AdapHaze is ideally suited for practical implementation in low-latency environments such as self-driving cars, surveillance systems, and unmanned aerial vehicles, which require reliable visual perception in challenging atmospheric conditions.

VI. FUTURE WORK

While AdapHaze exhibits improvements for all fog density levels, there are a number of areas that can be improved on.
The first area for improvement is handling severe fog. The results of restoration for all models were significantly worse in cases of dense fog with $\beta \geq 1.20$, which is due to a lack of data for training with such high scattering coefficients in the RESIDE OTS dataset [22]. Another logical step would be to use real-world data of severe fog (dense haze, NH-HAZE).

Another interesting area to pursue is the replacement of the three-class discrete density classifier with a fog density estimation function that operates on a continuous basis. The discrete classification method used currently for determining λ leads to sudden changes at the borders of each class. The use of regression will allow more precise fine-tuning of the dehazing process.

Finally, by broadening the scope of the AdapHaze model to consider various atmospheric deteriorations like smog, rain, and poor visibility, it may turn out to be a comprehensive end-to-end adverse weather restoration system.

REFERENCES

[1] D. Eigen, C. Puhrsch, and R. Fergus, "Depth map prediction from a single image using a multi-scale deep network," *Advances in Neural Information Processing Systems*, vol. 27, 2014.

[2] E. J. McCartney, *Optics of the Atmosphere: Scattering by Molecules and Particles*. New York, NY, USA: Wiley, 1976.

[3] S. G. Narasimhan and S. K. Nayar, "Chromatic framework for vision in bad weather," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, vol. 1, 2000, pp. 598–605.

[4] B. Li, X. Peng, Z. Wang, J. Xu, and D. Feng, "AOD-Net: All-in-one dehazing network," in *Proc. IEEE Int. Conf. Computer Vision (ICCV)*, 2017, pp. 4770–4778.

[5] A. Howard, M. Sandler, G. Chu, L.-C. Chen, B. Chen, M. Tan, W. Wang, Y. Zhu, R. Pang, V. Vasudevan, *et al.*, "Searching for MobileNetV3," in *Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV)*, 2019, pp. 1314–1324.

[6] K. He, J. Sun, and X. Tang, "Single image haze removal using dark channel prior," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 33, no. 12, pp. 2341–2353, 2010.

[7] R. Fattal, "Dehazing using color-lines," *ACM Trans. Graph.*, vol. 34, no. 1, pp. 1–14, 2014.

[8] R. T. Tan, "Visibility in bad weather from a single image," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2008, pp. 1–8.

[9] Q. Zhu, J. Mai, and L. Shao, "A fast single image haze removal algorithm using color attenuation prior," *IEEE Trans. Image Process.*, vol. 24, no. 11, pp. 3522–3533, 2015.

[10] D. Berman and S. Avidan, "Non-local image dehazing," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2016, pp. 1674–1682.

[11] Y. Liu, J. Pan, J. Ren, and Z. Su, "Learning deep priors for image dehazing," in *Proc. IEEE/CVF Int. Conf. Computer Vision (ICCV)*, 2019, pp. 2492–2500.

[12] B. Cai, X. Xu, K. Jia, C. Qing, and D. Tao, "DehazeNet: An end-to-end system for single image haze removal," *IEEE Trans. Image Process.*, vol. 25, no. 11, pp. 5187–5198, 2016.

[13] W. Ren, L. Ma, J. Zhang, J. Pan, X. Cao, W. Liu, and M.-H. Yang, "Gated fusion network for single image dehazing," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 3253–3261.

[14] H. Zhang and V. M. Patel, "Densely connected pyramid dehazing network," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 3194–3203.

[15] D. Chen, M. He, Q. Fan, J. Liao, L. Zhang, D. Hou, L. Yuan, and G. Hua, "Gated context aggregation network for image dehazing and deraining," in *Proc. IEEE Winter Conf. Applications of Computer Vision (WACV)*, 2019, pp. 1375–1383.

[16] X. Qin, Z. Wang, Y. Bai, X. Xie, and H. Jia, "FFA-Net: Feature fusion attention network for single image dehazing," in *Proc. AAAI Conf. Artificial Intelligence*, vol. 34, no. 7, 2020, pp. 11908–11915.

[17] J. Johnson, A. Alahi, and L. Fei-Fei, "Perceptual losses for real-time style transfer and super-resolution," in *Proc. European Conf. Computer Vision (ECCV)*, 2016, pp. 694–711.

[18] C. Ledig, L. Theis, F. Huszár, J. Caballero, A. Cunningham, A. Acosta, A. Aitken, A. Tejani, J. Totz, Z. Wang, *et al.*, "Photo-realistic single image super-resolution using a generative adversarial network", in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2017, pp. 4681–4690.

[19] R. Zhang, P. Isola, A. A. Efros, E. Shechtman, and O. Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proc. IEEE Conf. Computer Vision and Pattern Recognition (CVPR)*, 2018, pp. 586–595.

[20] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *arXiv preprint arXiv:1409.1556*, 2014.

ISBN: 978-81-689089-0-1

[21] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam, "MobileNets: Efficient convolutional neural networks for mobile vision applications," *arXiv preprint arXiv:1704.04861*, 2017.

[22] B. Li, W. Ren, D. Fu, D. Tao, D. Feng, W. Zeng, and Z. Wang, "Benchmarking single-image dehazing and beyond", *IEEE Trans. Image Process.*, vol. 28, no. 1, pp. 492–505, 2019.

Cite this article as:

*Yashika, Shobha and et. al., "*Adaphaze: Fog-Density-Aware Adaptive Image Dehazing Using Lightweight CNNs with Perceptual Loss *", Proceedings of 13th international conference on Microelectronics, Circuits and Systems, Micro2026 (Vol-II). Displayed as online on 30th July 2026.*
Link: http://actsoft.org/science/micro2026-pro/270-micro2026.pdf

ISBN: 978-81-689089-0-1