# FALCON: Frequency Adjoint Link with CONtinuous Density Mask for Fast Single Image Dehazing

Donghyun Kim Seil Kang Seong Jae Hwang\*

Yonsei University {danny0103, seil, seongjae}@yonsei.ac.kr

Abstract. Image dehazing, addressing atmospheric interference like fog and haze, remains a pervasive challenge crucial for robust vision applications such as surveillance and remote sensing under adverse visibility. While various methodologies have evolved from early works predicting transmission matrix and atmospheric light features to deep learning and dehazing networks, they innately prioritize dehazing quality metrics, neglecting the need for real-time applicability in time-sensitive domains like autonomous driving. This work introduces FALCON (Frequency Adjoint Link with CONtinuous density mask), a single-image dehazing system achieving state-of-the-art performance on both quality and speed. Particularly, we develop a novel bottleneck module, namely, Frequency Adjoint Link, operating in the frequency space to globally expand the receptive field with minimal growth in network size. Further, we leverage the underlying haze distribution based on the atmospheric scattering model via a Continuous Density Mask (CDM) which serves as a continuous-valued mask input prior and a differentiable auxiliary loss. Comprehensive experiments involving multiple state-of-the-art methods and ablation analvsis demonstrate FALCON's exceptional performance in both dehazing quality and speed (i.e., >180 frames-per-second), quantified by metrics such as FPS, PSNR, and SSIM.

Keywords: Image Dehazing  $\cdot$  Real-time Applicability  $\cdot$  Frequency Domain

## 1 Introduction

Eliminating naturally occurring visual distortions from atmospheric interference such as fog, smoke, and haze from an image, commonly termed *image dehazing*, is a ubiquitous task in vision applications. For instance, image dehazing plays a pivotal role in enhancing the visual quality of images captured in adverse weather conditions, providing clearer insights into fields such as surveillance [34, 36], autonomous driving [24, 39], and remote sensing. Despite its practical significance, image dehazing faces inherent challenges, including the variability of atmospheric conditions, the complex interaction of light with particles, and the delicate balance between preserving image details and reducing haze artifacts.

<sup>\*</sup> Corresponding author.



Fig. 1: (a) A Simplified Illustration of the FALCON Workflow. The purple arrow represents the operation of calculating the haze density map. We implement this process through an approach called Differentiable Density Pooling. (b) Analysis of dehazing performance (PSNR, Peak Signal-to-Noise Ratio) vs. dehazing speed (FPS, framesper-second) on NH-Haze2 dataset with images of size  $256 \times 256$  using RTX 3090 GPU. For each method, its circle size is proportional to the FLOPS. The goal is to achieve both high PSNR for quality and FPS for speed. Our method FALCON achieves the highest PSNR (22.41 dB) with the fastest inference FPS (182.90 frames per second), enabling a real-time single image dehazing while achieving the best dehazing quality.

In response to the aforementioned challenges in image dehazing, numerous methodologies have emerged to tackle the challenges presented by image degradation. Early research [10,11,14,35,51] predominantly concentrated on predicting features such as the transmission matrix and atmospheric light, relying on various haze-related priors and assumptions. Concurrently, the rise of deep learning has spurred the development of various dehazing networks [4,8,21,32,40]. With the progression of diverse research progress, the field of single image dehazing has expanded, witnessing a continuous influx of high-performing methods for haze removal.

While existing dehazing methods demonstrate great performance, they primarily focus on metrics like PSNR and SSIM, which measure the quality of dehazed images. However, in various practical applications such as autonomous driving, CCTV, and national defense technology where images must be dehazed quickly for responsive subsequent actions, the dehazing *speed* is another significant and perhaps necessary condition to pursue. Yet, existing state-of-the-art dehazing methods, primarily network-based, fall short in achieving *real-time applicability* where each image is dehazed within milliseconds.

This work aims to deliver a single image dehazing system that meets the demands of both quality and speed (Fig. 1(b)). Specifically, we focus on developing a network-based dehazing model, which the current top-performing methods consistently rely on. In that sense, to achieve fast inference time, our methodological efforts make minimal architectural changes to preserve the computational cost and concentrate on building a simple learning framework (Fig. 1(a)). We show how even a simple vanilla U-Net [33] outperforms existing deep image dehazing models with the following methodological contributions. The simple structure of a U-Net [33] lacks depth, posing challenges in achieving a sufficiently wide receptive field through conventional convolution operations for handling global image features. Also, real-world images often have high resolutions, such as  $1600 \times 1200$ , necessitating deeper networks at the expense of increased computational costs. To obtain larger receptive fields without proportionally growing the model, we introduced the Frequency Adjoint Link (FAL), inspired by the Fast Fourier Convolution [6]. This lightweight replacement of dense CNN-based bottlenecks facilitates comprehensive information capture from the image by performing convolution operations not only in the spatial domain but also in the frequency domain.

Although recent methods primarily leverage powerful deep models, various preceding works utilized a more fundamental understanding of haze as a type of atmospheric interference, namely, the atmospheric scattering model. The most prominent work, Dark Channel Prior [14], showed how a simple assumption about haze led to characterizing haze with a transmission map. Our work uniquely views the density map as a prior functioning as a continuous-valued mask, named Continuous Density Mask (CDM), indicating the degree of image degradation. Remarkably, a simple concatenation of this density-based prior yields potent cues about the haze distribution, significantly enhancing the learning process. Notably, we leverage CDM not only as the prior but also as an auxiliary loss by deriving a differentiable form of the density estimation.

**Contributions.** In this work, we introduce *FALCON* (Frequency Adjoint Link with **CON**tinuous density mask) for fast single image dehazing. Specifically, we make the following contributions:

- We propose a novel bottleneck module called Frequency Adjoint Link (FAL) designed to efficiently operate in the frequency space, effectively expanding the receptive field with minimal architectural growth.
- We also cleverly utilize Continuous Density Mask (CDM) characterizing the haze distribution based on the atmospheric scattering model. CDM is extremely efficiently computed, ensuring minimal computational burden, and is seamlessly integrated as a potent input prior and a differentiable auxiliary loss.
- Our study presents comprehensive results from diverse experiments, including comparisons with other state-of-the-art methods and an ablation analysis. Through metrics such as FPS, PSNR, and SSIM, we illustrate that our method achieves state-of-the-art performance, delivering exceptional speed alongside impressive dehazing quality.

We provide the code in the supplementary material which will be released upon publication.

# 2 Related Works

Atmospheric Scattering Model. In the image dehazing literature, various methodologies [26,27,35,51] have interpreted haze-like image degradation using

an atmospheric scattering model [23, 28, 29] formulated as follows:

$$\mathbf{I}(x) = \mathbf{J}(x)t(x) + \mathbf{A}(1 - t(x)), \tag{1}$$

where **I** is the hazy image, **J** is the clear or ground truth image before distortion, **A** is the global atmospheric light, and x is the index of pixel. Of particular note,  $t(x) \in (0, 1)$  is the medium transmission characterizing the proportion of light that reaches the camera after passing through the atmosphere. Conversely speaking, the proportion of light that does *not* reach the camera due to scattering is represented as 1 - t(x), which is also interpreted as the *haze density*.

**Prior-informed Dehazing.** Such finding has led early works to explicitly focus on predicting features like the transmission map and atmospheric light using haze-related priors and assumptions [10, 14, 35, 51]. A notable example is the Dark Channel Prior (DCP) [14] which often was combined with the atmospheric scattering model to effectively facilitate the transmission map computation, significantly advancing the field of dehazing. Similar efforts have also appeared including maximizing local contrast for haze removal and predicting the scene's albedo [10, 35].

Network-based Dehazing. Deep networks also began showing promising results. Starting from MSCNN [31] and DehazeNet [4], various deep models have attempted to predict the medium transmission map or the final haze-free image using CNNs, achieving significant improvements [9, 19, 25, 46, 47]. In particular, recent methods like DeHamer [13] have proposed combining CNNs and transformers to harness the strengths of both. Meanwhile, methods like C2PNet [48] have achieved impressive results using a novel learning approach called contrastive regularization.

**Frequency Domain in Dehazing.** In real-world haze images, the haze frequently appears as a *global* artifact, covering the majority of the image region. However, prior networks with convolutions have innate drawbacks of having a limited receptive field, making it challenging to efficiently derive global features. In response, inspired by prior vision techniques [6, 38, 43, 45], several dehazing methods operate in the frequency domain in which the corresponding Fourier bases span the whole image. For example, DW-GAN [12] proposed a two-branch network using wavelet transform, and many other studies [20,22,37,41] have proposed methodologies that apply the frequency domain using wavelet and Fourier transforms. Our work utilizes the Fast Fourier Convolution [6] to apply the convolution in the frequency domain, effectively expanding the receptive field in a simple convolutional framework.

# 3 Methods

In this section, we describe our method, FALCON, as shown in Fig. 2. First, we briefly outline the algorithmic process. Then, in detail, we cover each technical contribution of our approach.



**Fig. 2:** Overall pipeline of our single image dehazing method, FALCON. (1) The input hazy image is concatenated with its haze density map, namely, Continuous Density Mask, as a haze prior. (2) Our main network takes the concatenated input and leverages our Frequency Adjoint Link (FAL) to efficiently exploit the frequency domain. (3) The output dehazed image is compared against the ground-truth image in the image space  $(\mathcal{L}_{img})$ , density map space  $(\mathcal{L}_{map})$ , and VGG-16 feature space  $(\mathcal{L}_{per})$ .

#### 3.1 FALCON: Overview

Given an input 3-channel hazy image  $\mathbf{I} \in \mathbb{R}^{H \times W \times 3}$ , the goal is to generate a 3-channel dehazed image  $\hat{\mathbf{J}} \in \mathbb{R}^{H \times W \times 3}$  which is close to the ground-truth clear image  $\mathbf{J} \in \mathbb{R}^{H \times W \times 3}$ . In line with this general framework, our method brings the following novel contributions. (1) First, our model uniquely computes the haze density map, namely, Continuous Density Mask (CDM), based on the principle from the Dark Channel Prior (DCP) [14]. In particular, the original density map calculation involves a nested minimization which is seemingly non-differentiable and CPU-intensive. However, we identify that this process easily equates to a special type of pooling, turning the density map calculation to be an extremely computationally efficient process. (2) Next, we introduce our novel Frequency Adjoint Link (FAL) module which greatly improves the vanilla U-Net [33] by efficiently expanding the receptive field with minimal network size increase. (3) Lastly, our differentiable density map estimation allows us to incorporate a new kind of loss function based on the density map. Next, we describe each of the components in detail.

#### 3.2 Frequency Adjoint Link

For dense haze images which our study specifically considers, the existing networks often struggled with regions with extensive haze. When a significant area exhibits such nature, the corresponding areas showed poor dehazing results. We have identified the root of this issue as the limited receptive field of the convolution used in the network, not being able to consider the relatively clear regions

beyond the large hazy areas. In our work, to achieve dehazing with a broader receptive field, we leverage the frequency domain. In the spatial domain, each pixel of the image only has a value corresponding to its position, dealing with very local information. However, in the frequency domain, each point represented on a pixel contains information about the signal that constitutes the entire image. Therefore, performing convolution in the frequency domain allows for a wide receptive field without using many pixels [16]. For the transition to the frequency domain, we utilize the Fourier transform and adeptly apply the concept of Fast Fourier Convolution [6] to develop a module called the *Frequency Adjoint Link* (FAL).

We now describe the architecture of FAL in detail. Recall that the goal of FAL is to capture the haze pattern which is densely spread within an image. Specifically, it aims to consider both the *local* haze pattern which may not be sufficiently captured with typical CNN receptive fields, and the *global* haze pattern beyond the locally hazy area of interest.

The FAL consists primarily of the Fast Fourier Convolution Block (FFCB) and convolution block, as shown in Fig. 2(a). Let  $\mathbf{F} \in \mathbb{R}^{H \times W \times C}$  be the down-sampled feature at the bottom-most layer of U-Net, fed into the FFCB. Then, as shown in Fig. 2(b), F is first channel-wise partitioned as follows:  $\mathbf{F} \to {\mathbf{F}^l, \mathbf{F}^g}$  where  $\mathbf{F}^l \in \mathbb{R}^{H \times W \times (1-\alpha_{in})C}$  is the first  $(1 - \alpha_{in})C$  channels dedicated to the local route (*e.g.*, spatial domain), and  $\mathbf{F}^g \in \mathbb{R}^{H \times W \times \alpha_{in}C}$  is the next  $\alpha_{in}C$  channels dedicated to the global route (*e.g.*, frequency domain).  $\alpha_{in}$  is the ratio of channels between  $\mathbf{F}^l$  and  $\mathbf{F}^g$ . Specifically,  $\mathbf{F}^g$  is transformed into the Fourier feature in the frequency domain via convolution block and Fast Fourier Transform (FFT) layer:  $\mathbf{F}^g_{freq} = FFT(Conv(\mathbf{F}^g))$ , where  $Conv(\cdot)$  is the convolution block and  $FFT(\cdot)$  denotes FFT layer. Then, the subsequent convolution is applied onto  $\mathbf{F}^g_{freq}$  to derive the globally convolved receptive field. The resulting feature is then transformed back to the spatial domain via inverse Fast Fourier Transform (iFFT). More detailed information about the Fast Fourier Convolution Block can be found in the supplementary material.

By adding just this Frequency Adjoint Link to the middle of the U-Net [33], we implemented an image dehazing network, which shows both speed and impressive dehazing performance.

#### 3.3 Continuous Density Mask

Unlike some homogeneous types such as fog which almost uniformly covers the entire image with moderate visibility, we consider much more difficult cases of non-homogeneous and dense haze. That is, we observe dense haze formation resulting in very low visibility in almost all areas (*i.e.*, Dense-Haze dataset [1]; Fig. 3 first column). Also, non-homogeneous formation of haze may also result in barely visible areas (*i.e.*, NH-Haze and NH-Haze2 datasets [2,3]; Fig. 4 first column). From the methodological perspective, such dense and/or non-homogeneous haze scenarios bring a much greater challenge. While physics-informed methods such as the Dark Channel Prior (DCP) [14] can cleverly estimate the haze density, we point out that the actual dehazing process requires

a separate line of effort. Nonetheless, such haze density maps become a strong prior for the dehazing model, and we explicitly leverage this as input prior called the *Continuous Density Mask* (CDM). We next describe how CDM is derived and utilized within our framework.

Mask Mechanism In tasks like inpainting, where a mask is utilized throughout the pipeline, the mask typically assumes binary values of 0 or 1. This indicates whether a pixel retains its original value or is entirely lost. Borrowing this mechanism, we incorporated the concept of a mask in the dehazing task. Upon examining hazy images, especially real-world hazy images, it's evident that haze manifests as a complex distribution with varying densities across different locations. We interpreted the hazy image as a degraded version of the ground truth image, where each pixel is compromised to varying extents. Unlike inpainting, where compromised pixels lose all their values and assume a value of 0, in hazy images, pixels are degraded by the haze, not entirely lost. This means degradation of each pixel cannot be simply categorized into two values, 0 or 1. Instead, we assign a continuous value to each pixel based on the haze's density.

Haze Density Calculation We utilized the Dark Channel Prior (DCP) [14] to compute the haze density map. DCP is based on the observation that most local patches in outdoor haze-free images contain at least one pixel which has a very low intensity in at least one color channel. Based on the atmospheric scattering model [23, 28, 29], given a clean image J, its dark channel value  $J^{dark}(x)$  at the pixel location x is

$$J^{dark}(x) = \min_{c \in \{r,g,b\}} (\min_{y \in \Omega(x)} (J^{c}(y))),$$
(2)

where c is one of the  $\{r, g, b\}$  channels, y is a pixel from the local patch  $\Omega(x)$  centered at pixel x. Thus,  $J^c(y)$  is the value of channel c of a neighboring pixel y, and  $J^{dark}(x)$  is the minimum intensity value across all three channels within the window  $\Omega(x)$ . Then, following DCP, we assume that the dark channel value  $J^{dark}(x)$  is very low, that is,  $J^{dark}(x) = 0$ . Assuming the value of the global atmospheric light A as 1, the atmospheric scattering model [23,28,29]'s equation can be simplified using Dark Channel Prior [14] as

$$1 - t(x) = \min_{c \in \{r,g,b\}} (\min_{y \in \Omega(x)} (I^c(y))).$$
(3)

As previously mentioned, 1 - t(x) represents the haze density map. We have implemented this operation in a differentiable manner, referring to it as *Differentiable Density Pooling* (DDP). To perform the dual minimum operation present in Eq. (3), we adapted what is known as Min Pooling. Differentiable Density Pooling is as follows:

$$1 - t(x) = DDP(I(x)) = -\mathcal{M}((\mathcal{M}(-I(x)))^T).$$
(4)

 $\mathcal{M}(\cdot)$  and  $DDP(\cdot)$  represents Max Pooling and Differentiable Density Pooling, respectively. While standard Min Pooling identifies local minimum values, our

implementation goes a step further to ensure that the results are consistent with the existing density map calculations of Eq. (3). Although both operations yield the same outcome, naively implementing the dual minimum operation from Eq. (3) would result in a non-differentiable process. However, our approach, the Differentiable Density Pooling, allows for a differentiable implementation that does not break the computational graph within the overall backpropagation. Also, DDP allows us to quickly calculate the haze denisty map using a DCPbased algorithm without using a trainable network. Using Eq. (4), We swiftly and seamlessly connect the haze density map for each pixel in the hazy image as an auxiliary channel with the input hazy image. Additionally, by utilizing Differentiable Density Pooling, we have also implemented our novel loss term.

#### 3.4 Loss Formulation

Our dehazing framework optimizes a combination of Mean Squared Error (MSE) loss, Perceptual loss, and the newly introduced Density Map Loss to ensure highquality image reconstruction.

Mean Squared Error (MSE) Loss. The MSE loss computes the pixel-wise accuracy between the dehazed image  $\hat{J}$  and its ground truth J:

$$\mathcal{L}_{img} = \frac{1}{N} \sum_{i=1}^{N} (\hat{J}(i) - J(i))^2.$$
(5)

MSE provides a stable optimization landscape, aiding in model convergence and ensuring high-fidelity reconstructions.

**Perceptual Loss [15].** This loss is a combination of concept and style losses, derived using a pre-trained VGG-16 network. The concept loss captures high-level semantic differences, while the style loss ensures stylistic fidelity. Specifically, for a given dehazed image  $\hat{J}$  and its ground truth J,

$$\mathcal{L}_{concept} = ||\phi_8(\hat{J}) - \phi_8(J)||^2,$$
(6)

$$\mathcal{L}_{style} = \sum_{l \in \{3,8,15\}} ||G(\phi_l(\hat{J})) - G(\phi_l(J))||^2,$$
(7)

where  $\phi_l$  represents the feature maps from the  $l^{th}$  layer of the VGG-16 network, and G denotes the Gram matrix. The overall Perceptual loss is then

$$\mathcal{L}_{per} = \mathcal{L}_{concept} + \mathcal{L}_{style}.$$
(8)

**Density Map Loss.** To further refine the dehazing process, we introduce the Density Map Loss  $\mathcal{L}_{map}$ , which leverages the Differentiable Density Pooling  $DDP(\cdot)$  described in the previous subsection. This loss function measures the MSE between the haze density maps of the dehazed image  $DDP(\hat{J})$  and the ground truth DDP(J), providing a direct optimization target for haze density accuracy:

$$\mathcal{L}_{map} = \frac{1}{N} \sum_{i=1}^{N} (DDP(\hat{J})(i) - DDP(J)(i))^2.$$
(9)

**Combined Objective.** The final loss function that we optimize is a weighted combination of the MSE, Perceptual, and Density Map losses:

$$\mathcal{L}_{final} = \alpha \mathcal{L}_{img} + \beta \mathcal{L}_{per} + \gamma \mathcal{L}_{map} \tag{10}$$

where  $\alpha > 0$ ,  $\beta > 0$ , and  $\gamma > 0$  are the weights. In the later section, we empirically evaluate the significance of using all three losses with an ablation study as well.

## 4 Experiments

We first describe the experimental setup such as datasets, implementation details, and baselines. Then, we present the experimental results on multiple realworld image dehazing datasets against existing state-of-the-art methods along with the ablation studies. Lastly, we extensively analyze the computational speed of various dehazing networks and demonstrate the versatility of Continuous Density Mask.

#### 4.1 Experimental Setup

**Datasets.** Our experimentation involves three real-world datasets, and, consistent with previous studies [13, 40, 48], we conducted the experiments according to the official data splits of each dataset. (1) *Dense-Haze* [1] dataset is known for its dense, homogeneous haze conditions. (2) *NH-Haze* [2] dataset shows dense and non-homogeneous haze patterns. (3) Lastly, *NH-Haze2* [3] dataset is comprised of strong non-homogeneous hazy scenes. Additionally, for experiments on synthetic dataset, we employed the RESIDE [18], and more detailed information about these experimental results and the overall dataset can be found in the supplementary material.

Implementation Details. The experiments involve various GPUs to evaluate both the dehazing quality and the computational efficiency of our method. Specifically, we used an NVIDIA RTX A6000 GPU for the PSNR and SSIM evaluations. To measure the dehazing speed in frames-per-second (FPS) across different resolutions and GPU settings, we employed a diverse set of NVIDIA GPUs: RTX A6000, RTX 3090, GTX 1080 Ti, and T4. Each inference time was measured and averaged over the corresponding test set samples. We used the Adam optimizer with an initial learning rate of 0.0001 with the random flipping and random cropping augmentations. The batch size was set to 5 for training.

**Baseline Methods.** We compare our model FALCON to various traditional and recent state-of-the-art dehazing techniques. These include traditional methods like DCP [14], which relies on the dark channel prior, and a variety of deep learning approaches such as DehazeNet [4], AOD-Net [17], GCANet [5], GDNet [21], FFA-Net [30], and MSBDN [8]. Additionally, we have included recent advancements like AECR-Net [40], Dehamer [13], SFNet [7], and C2PNet [48] to cover a broad spectrum of methodologies.

**Evaluation Metrics.** Our metrics include the conventional Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) which

**Table 1:** Quantitative comparisons against other methods on real-world datasets. FPS and FLOPs are based on the inference using an RTX 3090 GPU with images of 256x256 resolution.

Method	Venue	Dense	e-Haze	NH-	Haze	NH-I	Haze2	Computatio	itational Efficiency	
monou	, onde	$\rm PSNR \uparrow$	$\uparrow$ SSIM $\uparrow$	$\overline{\mathrm{PSNR}}$ $\uparrow$	SSIM $\uparrow$	PSNR ↑	$SSIM \uparrow$	${ m FPS} \uparrow ({ m f/s})$	$\mathrm{FLOPs}\downarrow(\mathrm{G})$	
DCP [14]	TPAMI 2010	11.01	0.4165	13.28	0.4954	11.68	0.7090	-	-	
DehazeNet [4]	TIP 2016	9.48	0.4383	16.62	0.5238	11.77	0.6217	1088.14	1.162	
AOD-Net [17]	ICCV $2017$	13.14	0.4144	13.44	0.4136	12.33	0.6311	2564.10	0.230	
GCANet [5]	WACV 2019	12.62	0.4208	17.49	0.5918	18.79	0.7729	270.64	36.82	
GDNet [21]	ICCV 2019	14.96	0.5326	13.80	0.5370	19.26	0.8046	100.96	42.98	
FFA-Net [30]	AAAI 2020	16.31	0.5362	18.60	0.6374	20.45	0.8043	17.89	575.6	
MSBDN [8]	CVPR 2020	15.13	0.5551	19.23	0.7056	20.11	0.8004	75.47	83.08	
AECR-Net [40]	CVPR 2021	15.80	0.4660	19.88	0.7173	20.68	0.8282	164.07	104.4	
DeHamer [13]	CVPR 2022	16.62	0.5602	20.66	0.6844	19.18	0.7939	2.34	59.62	
SFNet [7]	ICLR 2023	17.46	0.5780	16.90	0.7052	17.81	0.8291	12.80	125.4	
C2PNet $[48]$	$\mathrm{CVPR}\ 2023$	16.88	0.5728	-	-	21.19	0.8334	0.30	461.2	
FALCON		19.51	0.5860	20.84	0.6772	22.41	0.8357	182.90	57.27	

evaluate the dehazing quality. Additionally, we measure the Frames Per Second (FPS), quantifying the number of images dehazed by the model per second, providing insights into the actual dehazing speed for real-time applications. Lastly, we also compute FLoating point Operations Per Second (FLOPs) which quantifies the number of floating point operations of the model for a single input during inference. Thus, an ideal method should achieve high PSNR, SSIM, and FPS with low FLOPs.

## 4.2 Dehazing Results

**PSNR and SSIM.** Table 1 shows the dehazing quality of all the methods on all three datasets in PSNR and SSIM. We first observe that FALCON consistently surpasses other methods with the highest PSNR and SSIM across the datasets. (1) On Dense-Haze, FALCON shows the highest PSNR of **19.51 dB** and SSIM of **0.5860**. We further note that our PSNR shows a gain of +2.05 (+11.7%) over the state-of-the-art method, which is significant considering the typical rate of improvement. (2) On NH-Haze, FALCON also shows the highest PSNR of **20.84 dB**. (3) On NH-Haze2, FALCON achieves the best PSNR of **22.41 dB** and SSIM of **0.8357**.

**FPS and FLOPs.** In Table 1, we show that FALCON achieves **182.90 FPS**, which means it can process over 180 images per second. This also equates 5 milliseconds inference time per image, which makes FALCON an extremely fast system, even applicable to standard 24 FPS videos in real-time. While there exist faster methods (i.e., rows 2 to 4 in Table 1), their PSNR and SSIM significantly underperform compared to ours, making FALCON the fastest approach within the top 5 models in each dataset. We believe the *best quality and speed* demonstrated by FALCON is a significant contribution towards standardizing practical image dehazing systems.



Fig. 3: Comparative visualization of dehazing results on Dense-Haze. The top row displays the overall results, while the bottom row shows a magnified view.



Fig. 4: Comparative visualization of dehazing results using the NH-Haze2 dataset. The top row displays the overall results, while the bottom row shows a magnified view.

**Qualitative Results.** Based on the visual assessment shown in Fig. 3 and Fig. 4, FALCON demonstrates a superior performance. On Dense-Haze, it eliminates haze while preserving the naturalness of the scene, sidestepping issues like color distortion and texture loss that plague other methods. Close-up views further underscore the details that FALCON retrieves, details that are often obscured in the outputs of other dehazing techniques. For NH-Haze2, FALCON reveals the underlying scene structures behind the dense haze, delivering results that are more true to the ground truth in both color and detail. In contrast, while AECR-Net [40] and SFNet [7] may appear to produce visually pleasing results at first glance, a closer examination reveals significant color discrepancies with the ground truth. Certain areas exhibit completely unnatural color restoration, a flaw not present in the FALCON outputs, which maintains color fidelity throughout.

## 4.3 Ablation Study

**Quantitative Results.** In Table 2, we ablate FALCON to assess its performance under different combinations of FAL, CDM, and  $\mathcal{L}_{map}$  on NH-Haze2. When none of them are included, FALCON is equivalent to U-Net. We first note that FAL alone brings a considerable improvement. Also, CDM significantly improves the performance. While  $\mathcal{L}_{map}$  does provide some gain,  $\mathcal{L}_{map}$ , including all three components brings the best result.

**Table 2:** Ablation study on FALCON assessing the effects of including the Frequency Adjoint Link (FAL), Continuous Density Mask (CDM), and  $\mathcal{L}_{map}$ 

FAL	-	$\checkmark$	-	-	$\checkmark$	-	$\checkmark$	$\checkmark$
CDM	-	-	$\checkmark$	-	$\checkmark$	$\checkmark$	-	$\checkmark$
$\mathcal{L}_{map}$	-	-	-	$\checkmark$	-	$\checkmark$	$\checkmark$	$\checkmark$
$PSNR \uparrow$	17.76	18.93	19.72	17.81	20.21	19.93	20.82	22.41
SSIM $\uparrow$	0.7644	0.7900	0.7975	0.7784	0.8012	0.8077	0.7963	0.8357



Fig. 5: An ablation study showcasing the progressive enhancement of our method. From left to right: Input image, Base setting without our proposed enhancements, Base+FAL (Frequency Adjoint Link), Base+CDM (Continuous Density Mask), Base+Density Map Loss, our final method combining all enhancements, and Ground Truth (GT) image. Each column demonstrates the visual improvements achieved by incrementally integrating our proposed techniques.

**Qualitative Results.** Fig. 5 provides a visual comparison across the different setups shown in Table 2. The progressive improvement in image clarity and detail recovery is evident with each added component. Our full model exhibits substantial dehazing effectiveness, securing between detail preservation and haze removal.

## 4.4 Computational Performance Analysis

**Comparison Across Resolutions.** Table 3 shows the inference speed in FPS using various image sizes and GPUs. We observe that across varying GPUs, FALCON maintains over 100 FPS in almost all cases. Also, on GTX 1080 Ti, FALCON processes  $256 \times 256$  images at an impressive **249.42** frames per second, and even at a resolution of  $2048 \times 2048$ , it sustains a rate of **46.51** frames per second, which is remarkable for such detailed imagery. Interestingly, using T4 which the Google Colab provides easily performs beyond 100 FPS. The results indicate that FALCON achieves a frame rate well above the 30 frames per second threshold required for real-time processing, even on high-resolution images.

**Comparison on GPUs.** Moreover, Table 4 presents a comprehensive GPU performance comparison in FPS achieved by FALCON and other leading dehazing methods. The table illustrates FALCON's exceptional speed across various GPUs, including NVIDIA RTX 3090, GTX 1080 Ti, and RTX A6000. While this table specifically addresses computational efficiency, it's important to refer to Table 1 for detailed dehazing quality metrics. Together, these results underscore

**Table 3:** Performance comparison of FPS (frames per second) across different image sizes and NVIDIA GPU models, demonstrating the scalability and efficiency of our method in various computational environments

Image Size	RTX A6000	RTX 3090	GTX 1080	Ti T4
$256 \times 256$	147.48	182.90	249.42	156.80
$512 \times 512$	120.88	178.27	195.48	130.39
$1024 \times 1024$	115.93	161.15	124.22	121.68
$2048 \times 2048$	109.98	157.94	46.51	111.53

**Table 4:** GPU performance comparison in FPS across various methods. All FPS measurements were conducted based on images with a resolution of  $1024 \times 1024$ , showing superior inference speed against other state-of-the-art methods across different NVIDIA GPUs.

GPU	DeHamer [	13] SFNet [7]	C2PNet [48	] FALCON
RTX A6000	38.96	2.45	1.14	115.93
RTX3090	0.135	2.557	0.015	161.15
GTX 1080 Ti	0.162	1.218	0.328	124.22

FALCON's capability to deliver high-speed processing without compromising on the quality of dehazing, making it a highly practical choice for real-time image dehazing.

## 4.5 Versatility of Continuous Density Mask (CDM)

To evaluate the effectiveness and versatility of the Continuous Density Mask (CDM), we designed experiments comparing the performance with and without the use of CDM across various networks. To accurately assess the effectiveness of CDM, we maintained identical settings for other hyperparameters, including learning rate and optimizer, and the experimental environment. To assess the versatility of CDM, we utilized widely used backbone networks in various fields, including U-Net [33] and U-Net++ [50], as well as networks proposed specifically for the dehazing task, such as FFA-Net [30] and MSBDN [8]. The dataset used for evaluation was NH-Haze2 [3], with PSNR and SSIM as the metrics. During training, only the Mean Squared Error (MSE) loss between images was used as the loss function. Table 5 presents the PSNR and SSIM results for each network, trained and evaluated with and without the application of CDM in the pre-network stage. It can be observed that the performance of all the various networks improved when CDM was applied before the network stage. Fig. 6 compares the dehazed images from each network with and without the application of CDM, confirming that the dehazed images with CDM applied demonstrate superior performance.

Table 5: Continuous Density Mask (CDM) on various networks

Network	U-Ne	et [33]	U-net-	++[50]	FFA-N	Vet [30]	MSB	DN [8]
CDM	-	$\checkmark$	-	$\checkmark$	-	$\checkmark$	-	$\checkmark$
$\mathrm{PSNR}\uparrow$	18.01	18.68	16.94	16.97	18.06	18.96	18.21	19.05
SSIM $\uparrow$	0.7684	0.7935	0.6202	0.6209	0.7866	0.7987	0.7679	0.7792



Fig. 6: Comparative visualization of dehazing results with and without CDM

# 5 Conclusion

In this study, we introduce the FALCON, a novel method for image dehazing that effectively balances high-quality dehazing with the demands of real-time applications. Our method, leveraging the innovative Frequency Adjoint Link, demonstrates a significant enhancement in handling high-resolution images through a wide receptive field, without the need for a complex network structure. The presentation of the Continuous Density Mask and Density Map Loss further elevates the network performance, enabling more precise and efficient dehazing by utilizing haze density information. While FALCON marks a substantial advancement in the field, it is noteworthy that the network is not the most lightweight in terms of parameters, suggesting potential areas for future optimization. Future work could focus on refining the network architecture to reduce its parameter count while maintaining, or even enhancing its current performance levels. Overall, FALCON represents a promising step forward in real-time image dehazing, offering a blend of efficiency and effectiveness that is well-suited for critical applications in autonomous driving, surveillance, and national defense technologies.

# References

- Ancuti, C.O., Ancuti, C., Sbert, M., Timofte, R.: Dense-haze: A benchmark for image dehazing with dense-haze and haze-free images. In: 2019 IEEE international conference on image processing (ICIP). pp. 1014–1018. IEEE (2019)
- Ancuti, C.O., Ancuti, C., Timofte, R.: Nh-haze: An image dehazing benchmark with non-homogeneous hazy and haze-free images. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops. pp. 444–445 (2020)
- Ancuti, C.O., Ancuti, C., Vasluianu, F.A., Timofte, R.: Ntire 2021 nonhomogeneous dehazing challenge report. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 627–646 (2021)
- Cai, B., Xu, X., Jia, K., Qing, C., Tao, D.: Dehazenet: An end-to-end system for single image haze removal. IEEE Transactions on Image Processing 25(11), 5187–5198 (2016). https://doi.org/10.1109/TIP.2016.2598681
- Chen, D., He, M., Fan, Q., Liao, J., Zhang, L., Hou, D., Yuan, L., Hua, G.: Gated context aggregation network for image dehazing and deraining. In: WACV. pp. 1375–1383. IEEE (2019)
- Chi, L., Jiang, B., Mu, Y.: Fast fourier convolution. Advances in Neural Information Processing Systems 33, 4479–4488 (2020)
- Cui, Y., Tao, Y., Bing, Z., Ren, W., Gao, X., Cao, X., Huang, K., Knoll, A.: Selective frequency network for image restoration. In: ICLR. OpenReview.net (2023)
- Dong, H., Pan, J., Xiang, L., Hu, Z., Zhang, X., Wang, F., Yang, M.: Multi-scale boosted dehazing network with dense feature fusion. In: CVPR. pp. 2154–2164. Computer Vision Foundation / IEEE (2020)
- Engin, D., Genç, A., Kemal Ekenel, H.: Cycle-dehaze: Enhanced cyclegan for single image dehazing. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 825–833 (2018)
- Fattal, R.: Single image dehazing. ACM transactions on graphics (TOG) 27(3), 1–9 (2008)
- Fattal, R.: Dehazing using color-lines. ACM transactions on graphics (TOG) 34(1), 1–14 (2014)
- Fu, M., Liu, H., Yu, Y., Chen, J., Wang, K.: Dw-gan: A discrete wavelet transform gan for nonhomogeneous dehazing. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops. pp. 203–212 (June 2021)
- Guo, C., Yan, Q., Anwar, S., Cong, R., Ren, W., Li, C.: Image dehazing transformer with transmission-aware 3d position embedding. In: CVPR. pp. 5802–5810. IEEE (2022)
- He, K., Sun, J., Tang, X.: Single image haze removal using dark channel prior. IEEE Transactions on Pattern Analysis and Machine Intelligence 33(12), 2341– 2353 (2011). https://doi.org/10.1109/TPAMI.2010.168
- Johnson, J., Alahi, A., Fei-Fei, L.: Perceptual losses for real-time style transfer and super-resolution. In: ECCV (2). Lecture Notes in Computer Science, vol. 9906, pp. 694–711. Springer (2016)
- Katznelson, Y.: An introduction to harmonic analysis. Cambridge University Press (2004)
- Li, B., Peng, X., Wang, Z., Xu, J., Feng, D.: Aod-net: All-in-one dehazing network. In: IEEE International Conference on Computer Vision, ICCV 2017, Venice, Italy, October 22-29, 2017. pp. 4780–4788. IEEE Computer Society (2017). https:// doi.org/10.1109/ICCV.2017.511, https://doi.org/10.1109/ICCV.2017.511

- 16 Kim et al.
- Li, B., Ren, W., Fu, D., Tao, D., Feng, D., Zeng, W., Wang, Z.: Benchmarking single-image dehazing and beyond. IEEE Transactions on Image Processing 28(1), 492–505 (2018)
- Li, R., Pan, J., Li, Z., Tang, J.: Single image dehazing via conditional generative adversarial network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 8202–8211 (2018)
- Liu, P., Zhang, H., Zhang, K., Lin, L., Zuo, W.: Multi-level wavelet-cnn for image restoration. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 773–782 (2018)
- Liu, X., Ma, Y., Shi, Z., Chen, J.: Griddehazenet: Attention-based multi-scale network for image dehazing. In: ICCV. pp. 7313–7322. IEEE (2019)
- Mao, X., Liu, Y., Liu, F., Li, Q., Shen, W., Wang, Y.: Intriguing findings of frequency selection for image deblurring. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 37, pp. 1905–1913 (2023)
- McCartney, E.J.: Optics of the atmosphere: scattering by molecules and particles. New York (1976)
- Mehra, A., Mandal, M., Narang, P., Chamola, V.: Reviewnet: A fast and resource optimized network for enabling safe autonomous driving in hazy weather conditions. IEEE Transactions on Intelligent Transportation Systems 22(7), 4256–4266 (2020)
- Mei, K., Jiang, A., Li, J., Wang, M.: Progressive feature fusion network for realistic image dehazing. In: Computer Vision–ACCV 2018: 14th Asian Conference on Computer Vision, Perth, Australia, December 2–6, 2018, Revised Selected Papers, Part I 14. pp. 203–215. Springer (2019)
- Narasimhan, S.G., Nayar, S.K.: Chromatic framework for vision in bad weather. In: Proceedings IEEE Conference on Computer Vision and Pattern Recognition. CVPR 2000 (Cat. No. PR00662). vol. 1, pp. 598–605. IEEE (2000)
- 27. Narasimhan, S.G., Nayar, S.K.: Vision and the atmosphere. International journal of computer vision 48, 233–254 (2002)
- Narasimhan, S.G., Nayar, S.K.: Contrast restoration of weather degraded images. IEEE transactions on pattern analysis and machine intelligence 25(6), 713–724 (2003)
- Nayar, S.K., Narasimhan, S.G.: Vision in bad weather. In: Proceedings of the seventh IEEE international conference on computer vision. vol. 2, pp. 820–827. IEEE (1999)
- Qin, X., Wang, Z., Bai, Y., Xie, X., Jia, H.: Ffa-net: Feature fusion attention network for single image dehazing. In: AAAI. pp. 11908–11915. AAAI Press (2020)
- Ren, W., Liu, S., Zhang, H., Pan, J., Cao, X., Yang, M.: Single image dehazing via multi-scale convolutional neural networks. In: Leibe, B., Matas, J., Sebe, N., Welling, M. (eds.) Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part II. Lecture Notes in Computer Science, vol. 9906, pp. 154–169. Springer (2016). https:// doi.org/10.1007/978-3-319-46475-6\_10, https://doi.org/10.1007/978-3-319-46475-6\_10
- Ren, W., Ma, L., Zhang, J., Pan, J., Cao, X., Liu, W., Yang, M.H.: Gated fusion network for single image dehazing. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3253–3261 (2018)
- Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: Medical Image Computing and Computer-Assisted Intervention-MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18. pp. 234–241. Springer (2015)

- Salazar-Colores, S., Cabal-Yepez, E., Ramos-Arreguin, J.M., Botella, G., Ledesma-Carrillo, L.M., Ledesma, S.: A fast image dehazing algorithm using morphological reconstruction. IEEE Transactions on Image Processing 28(5), 2357–2366 (2018)
- Tan, R.T.: Visibility in bad weather from a single image. In: 2008 IEEE conference on computer vision and pattern recognition. pp. 1–8. IEEE (2008)
- Ullah, H., Muhammad, K., Irfan, M., Anwar, S., Sajjad, M., Imran, A.S., de Albuquerque, V.H.C.: Light-dehazenet: a novel lightweight cnn architecture for single image dehazing. IEEE transactions on image processing **30**, 8968–8982 (2021)
- Wang, J., Li, C., Xu, S.: An ensemble multi-scale residual attention network (emranet) for image dehazing. Multimedia Tools and Applications 80(19), 29299–29319 (2021)
- Wang, Y., Tao, X., Qi, X., Shen, X., Jia, J.: Image inpainting via generative multicolumn convolutional neural networks. Advances in neural information processing systems **31** (2018)
- Wang, Y., Yan, X., Guan, D., Wei, M., Chen, Y., Zhang, X.P., Li, J.: Cyclesnspgan: Towards real-world image dehazing via cycle spectral normalized soft likelihood estimation patch gan. IEEE Transactions on Intelligent Transportation Systems 23(11), 20368–20382 (2022)
- Wu, H., Qu, Y., Lin, S., Zhou, J., Qiao, R., Zhang, Z., Xie, Y., Ma, L.: Contrastive learning for compact single image dehazing. In: CVPR. pp. 10551–10560. Computer Vision Foundation / IEEE (2021)
- Yang, H.H., Fu, Y.: Wavelet u-net and the chromatic adaptation transform for single image dehazing. In: 2019 IEEE International Conference on Image Processing (ICIP). pp. 2736–2740. IEEE (2019)
- 42. Ye, T., Jiang, M., Zhang, Y., Chen, L., Chen, E., Chen, P., Lu, Z.: Perceiving and modeling density is all you need for image dehazing. arXiv preprint arXiv:2111.09733 (2021)
- Yu, F., Koltun, V.: Multi-scale context aggregation by dilated convolutions. arXiv preprint arXiv:1511.07122 (2015)
- 44. Yu, H., Zheng, N., Zhou, M., Huang, J., Xiao, Z., Zhao, F.: Frequency and spatial dual guidance for image dehazing. In: European Conference on Computer Vision. pp. 181–198. Springer (2022)
- Yu, J., Lin, Z., Yang, J., Shen, X., Lu, X., Huang, T.S.: Generative image inpainting with contextual attention. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 5505–5514 (2018)
- Zhang, H., Patel, V.M.: Densely connected pyramid dehazing network. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 3194–3203 (2018)
- Zhang, H., Sindagi, V., Patel, V.M.: Multi-scale single image dehazing using perceptual pyramid deep network. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 902–911 (2018)
- 48. Zheng, Y., Zhan, J., He, S., Dong, J., Du, Y.: Curricular contrastive regularization for physics-aware single image dehazing. In: CVPR. pp. 5785–5794. IEEE (2023)
- Zhou, M., Huang, J., Guo, C.L., Li, C.: Fourmer: An efficient global modeling paradigm for image restoration. In: International Conference on Machine Learning. pp. 42589–42601. PMLR (2023)
- 50. Zhou, Z., Rahman Siddiquee, M.M., Tajbakhsh, N., Liang, J.: Unet++: A nested u-net architecture for medical image segmentation. In: Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: 4th International Workshop, DLMIA 2018, and 8th International Workshop, ML-CDS

2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings 4. pp. 3–11. Springer (2018)

51. Zhu, Q., Mai, J., Shao, L.: A fast single image haze removal algorithm using color attenuation prior. IEEE transactions on image processing **24**(11), 3522–3533 (2015)

# A Dataset Details

Haze generation	L	REAL-WORL	SYNTHETIC		
Dataset	Dense-Haze	NH-Haze	NH-Haze2	SOTS-indoor	SOTS-outdoor
Scene	Out/indoor	Outdoor	Outdoor	Indoor	Outdoor
Haze type	Homogeneous	Non-homogeneous	Non-homogeneous	Homogeneous	Homogeneous
Haze density	Dense	Moderate	Moderate	Low	Low
DW-GAN [12] (CVPRW'21)	✓	$\checkmark$	$\checkmark$	√	-
AECR-Net [40] (CVPR'21)	<ul> <li>✓</li> </ul>	$\checkmark$	-	$\checkmark$	-
DeHamer [13] (CVPR'22)	✓	$\checkmark$	-	$\checkmark$	$\checkmark$
PMDNet [42] (ECCV'22)	✓	$\checkmark$	-	$\checkmark$	$\checkmark$
FSDGN [44] (ECCV'22)	<ul> <li>✓</li> </ul>	$\checkmark$	-	√	-
SFNet [7] (ICLR'23)	<ul> <li>✓</li> </ul>	-	-	$\checkmark$	$\checkmark$
C2PNet [48] (CVPR'23)	✓	-	$\checkmark$	$\checkmark$	$\checkmark$
Fourmer [49] (ICML'23)	√	$\checkmark$	-	√	-
FALCON	√	$\checkmark$	$\checkmark$	$\checkmark$	-

Table 6: Details of datasets used in the image dehazing

we have utilized the best quality real-world datasets currently available in the dehazing research field. As shown in Table 6, we show the nature of dehazing datasets used by recent papers published at top venues.

# B Detailed Information about the Fast Fourier Convolution Block (FFCB)

The main component of FALCON's network, the Frequency Adjoint Link (FAL), comprises elements such as Fast Fourier Convolution Block (FFCB) and convolution block. Among these, the FFCB effectively handles both local and global features of an image through various receptive fields. The feature **F** entering the FFCB first undergoes channel-wise partitioning. After extensive experimentation to find the optimal partition ratio, FALCON uses a 1:3 ratio for channel-wise partitioning. The two partitioned features are denoted as **F**<sup>l</sup> and **F**<sup>g</sup>. **F**<sup>l</sup> generates two features **F**<sup>l</sup><sub>s1</sub> and **F**<sup>l</sup><sub>s2</sub> through two convolutions in the spatial domain, while **F**<sup>g</sup> produces the feature **F**<sup>g</sup><sub>s1</sub> and the feature **F**<sup>g</sup><sub>s1</sub> through convolution in the frequency domain using Fast Fourier Transform (FFT) and the feature **F**<sup>g</sup><sub>s</sub> through element-wise addition, followed by batch normalization and ReLU to form the feature **F**<sup>l</sup><sub>mix1</sub>. Similarly, **F**<sup>l</sup><sub>s2</sub> and **F**<sup>g</sup><sub>f</sub> are combined through element-wise addition and then passed through batch normalization and ReLU in sequence to form the feature **F**<sup>g</sup><sub>mix1</sub> from **F**<sup>l</sup> and **F**<sup>g</sup> as *local*(·) and *global*(·), respectively. The operations up to this point can be expressed as follows:

$$\mathbf{F}_{mix1}^{l} = local(\mathbf{F}^{l}, \mathbf{F}^{g}) \tag{11}$$



**Fig. 7:** Detailed Architecture of the Fast Fourier Convolution Block. The block diagram shows the detailed layer-wise composition, emphasizing the integration of Fourier transformations within the convolutional process. Each layer is annotated with its specific function, providing a comprehensive view of the block's operations. The symbol  $\bigoplus$  denotes element-wise addition in this context. This design facilitates efficient dehazing by leveraging the frequency domain for enhanced receptive fields.

$$\mathbf{F}^{g}_{mir1} = global(\mathbf{F}^{l}, \mathbf{F}^{g}). \tag{12}$$

The next step involves using  $local(\cdot)$  and  $global(\cdot)$  to compute the subsequent features from  $\mathbf{F}_{mix1}^{l}$  and  $\mathbf{F}_{mix1}^{g}$ . This process can be expressed as

$$\mathbf{F}_{mix2}^{l} = local(\mathbf{F}_{mix1}^{l}, \mathbf{F}_{mix1}^{g}) \tag{13}$$

$$\mathbf{F}_{mix2}^{g} = global(\mathbf{F}_{mix1}^{l}, \mathbf{F}_{mix1}^{g}).$$
(14)

The two resulting features,  $\mathbf{F}_{mix2}^{l}$  and  $\mathbf{F}_{mix2}^{g}$ , are then combined through channel-wise concatenation to form the final output feature of the FFCB. The detailed structure of the FFCB can be seen in Fig. 7.

## C Results on Synthetic Datasets

Just as we measure PSNR and SSIM for real-world datasets, we also evaluate the performance on synthetic dataset. We use the RESIDE [18] dataset to assess the performance of FALCON on synthetic data. Following the settings used in previous studies [21, 30, 40], we select ITS as the training dataset and

Method	Venue	SOTS		Computational Efficiency		
		$ m PSNR\uparrow$	SSIM $\uparrow$	$\overline{\mathrm{FPS}}\uparrow(\mathrm{f/s})$	$\mathrm{FLOPs}\downarrow(\mathrm{G})$	
DCP [14]	TPAMI 2010	16.62	0.8179	-	-	
DehazeNet [4]	TIP 2016	19.82	0.821	1088.14	1.162	
AOD-Net [17]	ICCV $2017$	20.51	0.816	2564.10	0.230	
GCANet [5]	WACV $2019$	30.06	0.9596	270.64	36.82	
GDNet [21]	ICCV 2019	32.16	0.9836	100.96	42.98	
FFA-Net [30]	AAAI 2020	36.39	0.9886	17.89	575.6	
MSBDN [8]	$\mathrm{CVPR}\ 2020$	33.67	0.985	75.47	83.08	
AECR-Net [40]	$\mathrm{CVPR}\ 2021$	37.17	0.9901	164.07	104.4	
DeHamer [13]	$\mathrm{CVPR}\ 2022$	36.63	0.9881	2.34	59.62	
SFNet [7]	ICLR 2023	41.24	0.996	12.80	125.4	
C2PNet $[48]$	$\mathrm{CVPR}\ 2023$	42.56	0.9954	0.30	461.2	
FALCON		37.01	0.9762	182.90	57.27	

**Table 7:** Quantitative comparisons against other methods on synthetic datasets. FPS and FLOPs are based on the inference using an RTX 3090 GPU with images of 256x256 resolution.

SOTS-indoor for testing dataset. ITS and SOTS-indoor consist of 13,990 and 500 images, respectively. All other experimental settings were identical to those used for the real-world datasets. As can be seen in the results of Table 7, FAL-CON demonstrated competitive performance with a PSNR of 37.01dB and an SSIM of 0.9762, not falling behind other methods.

# D Additional Visual Results

To accurately and transparently showcase the dehazing performance of FAL-CON, we present the results for test dataset images of every real-world dataset we experimented with. Figs. 8 to 10 display the outcomes of FALCON for images in the test datasets of Dense-Haze [1], NH-Haze [2], and NH-Haze2 [3], respectively.

## E Discussion

## 5.1 Limitation and Future Works

Designed for real-time dehazing, FALCON demonstrates a commendable balance between rapid inference speed and high-quality dehazing. To achieve swift inference, the network architecture is kept as simple as possible, resulting in remarkably low FLOPs. Despite this, due to its U-net [33] based structure, FAL-CON possesses a relatively high number of parameters. The parameter count of

the FALCON network implies a substantial memory requirement. This aspect suggests that for successful implementation of real-time dehazing in edge devices, strategies to reduce the number of parameters should be considered.

Moreover, while FALCON is designed to be optimized for real-world dehazing scenarios, it is observed that its performance on synthetic datasets is not at the forefront. This indicates a potential area for improvement. Enhancing FAL-CON's adaptability to various datasets, especially synthetic ones, could further solidify its applicability across a broader range of dehazing tasks, ensuring robust performance irrespective of the dataset's origin.

## 5.2 Broader Impact

FALCON is a remarkable dehazing technique, demonstrating its potential for significant application in fields where dehazing is crucial. Despite having room for further development, its current capabilities make it sufficiently robust for use in real-world scenarios where dehazing is essential. The field of dehazing research is driven by the need for rapid dehazing techniques in applications such as autonomous driving, CCTV, and national defense technology. These applications require the ability to continuously capture hazy environments and produce clear images promptly for subsequent high-level computer vision tasks or other necessary operations. Thus, the ultimate goal of dehazing methods extends beyond merely improving the quality of dehazing; they must also be efficient and fast for real-time applications.

Given its exceptional dehazing quality and remarkably fast inference speed on real-world datasets, FALCON emerges as a highly suitable solution in the realm of dehazing technology. FALCON's capabilities extend beyond technical excellence, offering broader societal impacts and positive contributions to various sectors. Its proficiency in delivering high-quality dehazing and rapid inference speed on real-world datasets positions it as a pivotal tool in enhancing the effectiveness of critical applications such as autonomous driving, CCTV surveillance, and national defense. By ensuring clearer visual information in challenging atmospheric conditions, FALCON can significantly improve safety and reliability in these areas. Its application in autonomous driving, for instance, could lead to safer navigation and reduced accidents in poor visibility conditions. In surveillance and defense, clearer images can enhance monitoring accuracy and situational awareness. Thus, FALCON not only represents a technological advancement in the field of dehazing but also holds the potential to bring about substantial societal benefits by improving safety, security, and operational efficiency in various real-world scenarios.

# FALCON 23



Fig. 8: Dehazing results of FALCON on Dense-Haze.



Fig. 9: Dehazing results of FALCON on NH-Haze.

# FALCON 25



Fig. 10: Dehazing results of FALCON on NH-Haze2.