

PHATNet: A Physics-guided Haze Transfer Network for Domain-adaptive Real-world Image Dehazing

Fu-Jen Tsai^{1,2} Yan-Tsung Peng³ Yen-Yu Lin⁴ Chia-Wen Lin¹

National Tsing Hua University¹ MediaTek²

National Chengchi University³ National Yang Ming Chiao Tung University⁴

fjtsai@gapp.nthu.edu.tw ytpeng@cs.nccu.edu.tw lin@cs.nycu.edu.tw cwlin@ee.nthu.edu.tw

Abstract

Image dehazing aims to remove unwanted hazy artifacts in images. Although previous research has collected paired real-world hazy and haze-free images to improve dehazing models' performance in real-world scenarios, these models often experience significant performance drops when handling unseen real-world hazy images due to limited training data. This issue motivates us to develop a flexible domain adaptation method to enhance dehazing performance during testing. Observing that predicting haze patterns is generally easier than recovering clean content, we propose the Physics-guided Haze Transfer Network (PHATNet) which transfers haze patterns from unseen target domains to source-domain haze-free images, creating domain-specific fine-tuning sets to update dehazing models for effective domain adaptation. Additionally, we introduce a Haze-Transfer-Consistency loss and a Content-Leakage Loss to enhance PHATNet's disentanglement ability. Experimental results demonstrate that PHATNet significantly boosts state-of-the-art dehazing models on benchmark real-world image dehazing datasets. The source code is available at <https://github.com/pp00704831/PHATNet>.

1. Introduction

Images captured in hazy conditions often experience significant degradation, including occlusion, color distortion, reduced contrast, etc. These factors severely diminish image clarity and adversely affect downstream computer vision tasks, such as image classification [25, 38], object detection [20, 28], and semantic segmentation [29, 45, 49]. Image dehazing seeks to restore high-quality, visually enhanced images from a single hazy input; however, it remains challenging due to its inherently ill-posed nature.

Image dehazing has advanced considerably in recent years, with numerous methods developed using various deep neural network architectures, including Convolutional

Neural Networks (CNNs) [11, 12, 35, 40, 42, 43, 47, 50] and Transformers [16, 41], to improve dehazing performance. However, due to the scarcity of real-world clean-hazy image pairs for training, these methods often rely on synthetic hazy image datasets [22] generated through the following Atmospheric Scattering Model (ASM):

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

where I and $J \in \mathbb{R}^{H \times W \times 3}$ denote the hazy and corresponding haze-free images of size $H \times W$. $A \in \mathbb{R}^1$ is the global atmospheric light, and the transmission map is given by $t = e^{-\beta d(x)} \in \mathbb{R}^{H \times W}$, where β represents the haze density, $d(x)$ is the depth map, and x is the pixel index.

However, as ASM is not sufficient to accurately simulate haze distribution in real-world scenarios, some studies [2–7] have focused on collecting paired real-world hazy and clean images to improve dehazing models' performance. Although dehazing models trained on real-world pairs perform better than those trained on synthetic pairs, they often suffer from significant performance drops when applied to unseen, real-world, hazy images from different domains.

Since paired hazy and haze-free images for unseen target domains are typically unavailable during testing, some approaches [23, 46] utilize image-to-image translation techniques based on Generative Adversarial Networks (GANs) [14] to generate the paired images. However, GAN-based methods may encounter issues, such as mode collapse [1, 30] and instability during optimization [15, 31]. Moreover, they struggle to capture region-specific degradation patterns, such as non-uniform haze, commonly seen in real-world scenarios. Additionally, GANs lack effective domain adaptation capabilities, necessitating re-optimization of the generator to accommodate hazy images from unseen domains. These challenges motivate us to develop a flexible domain adaptation framework that allows dehazing models to effectively adjust to new target domains at test time.

We present a novel domain adaptation framework to update dehazing models, enhancing their performance on unseen target-domain hazy images. Since haze-free counter-




Figure 1. Illustration of the proposed Domain-Adaptive Dehazing framework. We introduce a Physics-guided Haze Transfer Network (PHATNet) which transfers haze patterns from hazy images in unseen target domains to clean images in the source domain. This process generates domain-adaptive fine-tuning sets, enabling dehazing models to be updated during testing for effective domain adaptation.

parts of target-domain hazy images are typically unavailable during testing, it is essential to develop effective methods for generating paired hazy and clean images authentically reflecting target-domain haze patterns for test-time adaptation. Our key observation is that retrieving haze patterns, especially thick haze, is generally easier than restoring haze-free content. This is because haze often forms a smooth, uniform, and semi-transparent layer that overlays the scene. ASM captures this phenomenon by simulating light scattering caused by atmospheric particles, resulting in a more predictable degradation in hazy images. Consequently, estimating haze is less challenging than reconstructing the original scene, which often contains intricate details obscured by the haze. Based on this insight, we propose a **Physics-guided HAZE Transfer Network** (PHATNet), designed to transfer haze patterns from unseen target domains to source-domain clean images. Through this process, PHATNet generates domain-adaptive fine-tuning sets with paired hazy/haze-free images, enabling effective test-time updates for dehazing models, as shown in Figure 1.

PHATNet leverages ASM guidance to transfer haze patterns extracted from target-domain hazy images. However, the standard ASM often fails to accurately represent real-world haze distributions due to its sensitivity to scene content, such as depth variations. To address this, we propose a novel Parametric Haze Disentanglement and Transfer (PHDT) module that extends ASM to the latent space to disentangle haze-related and content-related features, facilitating their recombination to generate high-quality haze-transferred images. This latent space extension also enables PHATNet to mitigate ghosting artifacts caused by scene depth variations. Additionally, we propose a Haze-Transfer-Consistency loss and a Content-Leakage Loss to further prevent PHATNet from leaking content-related features from hazy images.

PHATNet offers several advantages for domain adaptation in real-world dehazing tasks. First, by introducing ASM as an inductive bias for haze formation, PHATNet can accurately disentangle and transfer haze patterns from target-domain hazy images. This precision is achiev-

able because haze exhibits greater regularity than natural scenes within the ASM parametric domain. Additionally, the ASM-guided approach ensures that the haze transfer and subsequent augmentation processes are physically interpretable. Second, PHATNet facilitates the augmentation of ample domain-specific fine-tuning data, effectively mitigating domain gaps and promoting domain-specific adaptation for dehazing models. Third, by utilizing an offline fine-tuning process, PHATNet ensures that dehazing models do not incur increased latency during testing after adaptation.

Our key contributions are summarized as follows:

- We propose a flexible domain adaptation framework, called PHATNet, which disentangles and transfers haze patterns from target to source domains. This approach generates adaptive fine-tuning sets to update dehazing models, enhancing their performance during testing.
- Recognizing that extracting haze patterns is generally easier than restoring haze-free scene content from hazy images, we devise a Parametric Haze Disentanglement and Transfer (PHDT) module. This module, combined with a Haze-Transfer-Consistency Loss and a Content-Leakage Loss, enables effective haze pattern transfer.
- Extensive experimental results demonstrate our approach significantly boosts existing dehazing models across seven benchmark real-world image dehazing datasets.

2. Related Work

2.1. Image Dehazing

Image dehazing has made significant progress with the development of deep learning. Several studies [8, 12, 21, 34, 39] have aimed to enhance dehazing performance by leveraging prior knowledge from ASM. For instance, It was proposed in [8, 39] to estimate the transmission map of a hazy image for dehazing. Deng *et al.* [12] proposed a haze-aware representation distillation module by reformulating ASM. Other approaches [11, 17, 19, 27, 35–37, 40, 41, 43, 48, 50] attempted to directly learn a hazy-to-clean mapping. For example, Guo *et al.* [16] proposed a transformer-based dehazing model with transmission-aware position embedding.




Figure 2. Architecture of the proposed Physics-guided Haze Transfer Network (PHATNet). (a) PHATNet is a multi-scale network composed of the proposed Parametric Haze Disentanglement and Transfer (PHDT). (b) PHDT is a dual-branch network where the top branch disentangles haze-related features—specifically, atmospheric-light features F^{AL} by Atmospheric Light Encoder (ALE) and transmission-map features F^{TM} by Transmission Maps Encoder (TME)—from a hazy image. In contrast, the bottom branch extracts content-related features F^J by Content Encoder (CE) from a haze-free image. Guided by the Atmospheric Scattering Model (ASM), PHDT then combines these features in the ASM parametric domain to produce the final haze-transferred image by Rehazing Encoder (RE).

Cui *et al.* [11] devised a dual-domain selection mechanism to amplify the response of important spatial and frequency regions. Shen *et al.* [40] presented a triplet interaction network to aggregate cross-domain, cross-scale, and cross-stage features. Although many architectural improvements have been made, these methods often fail to handle real-world hazy images across various domains, limiting their applicability in practical scenarios.

2.2. Domain Adaptation for Image Restoration

In addition to enhancing restoration performance through architectural design, several studies have focused on improving the performance of restoration models via data augmentation [9, 23, 44, 46] and domain adaptation [10, 18, 24, 26] for image restoration. In real-world hazy environments, haze patterns can be unpredictable, exhibiting a variety of shapes, color distortions, and contrast attenuation. This variability negatively impacts the performance of dehazing models when tested in different domains. To bridge the domain gap in image dehazing, several methods [10, 23, 24, 46] have been proposed. Since hazy-clean image pairs are unavailable during testing, Liu *et al.* [24] proposed a helper network that reconstructs the hazy input from the dehazed output, allowing for the use of cycle-consistency loss to update dehazing models during testing. However, because the cycle reconstruction process functions as an auxiliary task, it limits the potential improvement of dehazing performance at inference. Another approach to haze synthesis involves using GANs to generate hazy-clean image pairs that mimic haze distributions in target domains. For example, Li *et al.* [23] and Yang *et al.* [46] applied ASM and GANs to transfer real-world haze patterns

to clean images. However, GAN-based methods often suffer from mode collapse, limiting their ability to accurately transfer region-specific haze patterns. To overcome these limitations, Chen *et al.* [10] proposed to generate visual prompts that align with haze distributions in target domains through patch-based image-level normalization. However, these visual prompts may contain patch artifacts, reducing the fidelity of haze-transferred images.

3. Proposed Method

To address these challenges, we introduce the Physics-guided Haze Transfer Network (PHATNet), designed to transfer haze patterns from hazy images in unseen target domains to haze-free images in the source domain. This transfer creates domain-adaptive fine-tuning sets, enabling dehazing models to adapt effectively to new domains. At the core of PHATNet is the novel Parametric Haze Disentanglement and Transfer (PHDT) module, which disentangles haze patterns from a hazy image and scene content from a haze-free image and then combines them within the ASM parametric domain. To strengthen PHDT’s capacity to separate haze features from scene content, we introduce the Haze-Transfer-Consistency and Content-Leakage loss functions, enabling precise haze transfer. This section details the design of PHDT, PHATNet’s architecture, the loss functions, and the domain adaptation process.

3.1. Parametric Haze Disentanglement & Transfer

To transfer haze patterns from a hazy image I^H to a haze-free image I^C , where $I^H, I^C \in \mathbb{R}^{H \times W \times 3}$, we introduce PHDT with ASM guidance. As illustrated in Figure 2,



Figure 3. Illustration of PHDT-based hazy image augmentation. Beyond generating the haze-transferred image, we can apply gamma correction to F^{TM} to adjust haze density, either increasing or decreasing it to further augment images A and B. Additionally, we can vertically flip F^{TM} to augment image C.

PHDT operates as a dual-branch network: the top branch extracts *haze features* from the hazy image, while the bottom branch isolates *content features* from the haze-free image. Using guidance from ASM, PHDT then fuses these features to produce the final haze-transferred image.

In the top branch, we introduce the Atmospheric Light Encoder (ALE) and Transmission Maps Encoder (TME) to disentangle atmospheric-light features $F^{\text{AL}} \in \mathbb{R}^{128}$ and transmission-map features $F^{\text{TM}} \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times 128}$ from the hazy image I^H as follows:

$$F^{\text{AL}} = \exp\{-\text{ALE}(I^H)\}, \quad (2)$$

$$F^{\text{TM}} = \exp\{-\text{TME}(I^H)\}, \quad (3)$$

where we apply the ReLU activation function at the end of the ALE and TME to clip negative values to zero, ensuring that the output features are non-negative. Following this, we use the function e^{-x} to normalize F^{AL} and F^{TM} within the range $[0, 1]$, maintaining consistency with the physical interpretation of ASM.

The bottom branch employs Content Encoder (CE) to disentangle content features $F^J \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times 128}$ from I^C :

$$F^J = \text{CE}(I^C). \quad (4)$$

Next, we fuse F^{AL} , F^{TM} , and F^J based on the ASM in Eq. (1) to generate rehazed features $F^I \in \mathbb{R}^{\frac{H}{8} \times \frac{W}{8} \times 128}$ as

$$F^I = F^J \times F^{\text{TM}} + F^{\text{AL}} \times (1 - F^{\text{TM}}). \quad (5)$$

Unlike the standard ASM in Eq. (1) that operates in the image domain, we generate transmission-map features F^{TM} in the latent space, forming haze-aware attention maps invariant to scene depth. This enables us to disentangle haze-related features across channels without introducing ghosting artifacts. Finally, we employ a Rehazing Encoder (RE) to generate the final haze-transferred image $I^O \in$




Figure 4. Illustration of the loss functions for PHATNet. The top figure depicts the cyclic haze-transfer process of PHATNet, supervised by the Haze-Transfer-Consistency Loss ($\mathcal{L}_{\text{HTC}}$). In the bottom figure, since the hazy image lacks haze content, the Content-Leakage Loss ($\mathcal{L}_{\text{CL}}$) is applied to ensure that PHATNet generates a haze-transferred image identical to the original haze-free image.

$\mathbb{R}^{H \times W \times 3}$. This image combines the haze patterns from I^H with content from I^C . The ALE, TME, CE, and RE modules are constructed from CNN-based residual blocks; additional architectural details are available in the supplementary material. By generating hazy images aligned with ASM physical principles, we can further manipulate haze-transferred images in the ASM parametric domain to augment additional target-domain hazy images. For example, applying gamma correction to F^{TM} allows for haze density adjustments, while vertically flipping F^{TM} creates additional variations, as illustrated in the augmented images A, B, and C in Figure 3.

3.2. Multi-scale Haze Transfer

Given a hazy image I^H from the target domain and a haze-free image I^C from the source domain, PHATNet can generate a haze-transferred image I^O , which incorporates target-domain haze patterns from I^H along with content information from I^C as follows: $I^O = \text{PHATNet}(I^H, I^C)$. Since haze patterns are often non-homogeneous and can vary in scales, as shown in Figure 2(a), we employ a multi-scale PHDT structure in PHATNet to facilitate the transfer of haze patterns in the following coarse-to-fine manner:

$$I^O = \text{PHATNet}(I^H, I^C) = \text{PHDT}(I^H, I^C) + \text{UP}(I^{O\downarrow}), \quad (6)$$

where

$$I^{O\downarrow} = \text{PHDT}(I^{H\downarrow}, I^{C\downarrow}) + \text{UP}(I^{O\downarrow\downarrow}), \quad (7)$$

and

$$I^{O\downarrow\downarrow} = \text{PHDT}(I^{H\downarrow\downarrow}, I^{C\downarrow\downarrow}), \quad (8)$$

where $(\cdot)^\downarrow$ and $(\cdot)^{\downarrow\downarrow}$ denote the downsampling-by-two and downsampling-by-four operations, respectively, and $\text{UP}(\cdot)$ represents the upsampling-by-two operation.

3.3. Loss Functions

We use paired real-world hazy and haze-free images $\{I_i^H, I_i^C\}_{i=1}^N$ in the source domain, consisting of N pairs, to

train PHATNet. The optimization process is driven by two key loss functions: the Haze-Transfer-Consistency Loss ($\mathcal{L}_{\text{HTC}}$) and the Content-Leakage Loss ($\mathcal{L}_{\text{CL}}$), as illustrated in Figure 4. The Haze-Transfer-Consistency Loss $\mathcal{L}_{\text{HTC}}$ assesses the difference between the haze-transferred image $I_i^O = \text{PHATNet}(I_i^H, I_i^C)$ and its corresponding hazy reference, I_i^H . To generate the haze-transferred image I_i^O , PHATNet performs a cyclic haze-transfer process: haze features from I_i^H are extracted in the ASM parametric domain, and then combined with scene content from haze-free image I_j^C , yielding I_i^O in the image domain, ideally capturing the same haze patterns as I_i^H . $\mathcal{L}_{\text{HTC}}$ is defined as

$$\mathcal{L}_{\text{HTC}} = \sum_{s=1}^3 \|I_{i,s}^O - I_{i,s}^H\|_1, \quad (9)$$

where s denotes the scale index in the multi-scale network.

To isolate haze patterns exclusively from hazy images for haze transfer, it is essential that these haze patterns, extracted from hazy images I_i^H , remain disentangled from the underlying scene content. To achieve this, we introduce the Content-Leakage Loss, which minimizes content leakage from hazy images I_i^H into the corresponding haze-transferred images. This method reconsiders the use of unpaired image translation without relying on GANs, which are often unstable in training. To prevent the atmospheric light and transmission map encoders from capturing content-related features, we use unpaired haze-free images as input to PHATNet, generating $I_{i,j}^O = \text{PHATNet}(I_i^C, I_j^C)$, where $j \neq i$. In this setup, $I_{i,j}^O$ incorporates haze patterns from I_i^C while preserving the scene content from I_j^C . Since I_i^C is haze-free, $I_{i,j}^O$ should ideally only reflect the content of I_j^C . Thus, the Content-Leakage Loss can be defined as

$$\mathcal{L}_{\text{CL}} = \sum_{s=1}^3 \|I_{i,j,s}^O - I_{j,s}^C\|_1, \quad (10)$$

Consequently, we optimize PHATNet using the total loss:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{HTC}} + \mathcal{L}_{\text{CL}}. \quad (11)$$

3.4. Domain Adaptation via Haze Transfer

After training PHATNet, we use it to transfer haze patterns from hazy images $\{\tilde{I}_i^H\}_{i=1}^M$ in the target domain to haze-free images $\{I_j^C\}_{j=1}^N$ in the source domain, producing a domain-specific fine-tuning set $S = \{\tilde{I}_{i,j}^O, I_j^C\}_{\forall i,j}$ with a total of $M \times N$ training pairs. Here, $\tilde{I}_{i,j}^O = \text{PHATNet}(\tilde{I}_i^H, I_j^C)$ represents the haze-transferred image that combines haze patterns from $\tilde{I}_i^H$ with the scene content from I_j^C , while I_j^C serves as its corresponding haze-free version. We then fine-tune dehazing models on the generated domain-specific fine-tuning set S , enabling effective domain adaptation.

Table 1. Experimental settings applied in our experiments.

Setting	Source Domain	Target Domains	
Setting1	NH-Haze20 [5]	NH-Haze21 [6]	HD-NH-Haze [7]
		I-Haze [2]	O-Haze [3]
		DenseHaze [4]	RTTS [22]
Setting2	HD-NH-Haze [7]	NH-Haze20 [5]	NH-Haze21 [6]
		I-Haze [2]	O-Haze [3]
		DenseHaze [4]	RTTS [22]

4. Experiments

4.1. Datasets

We utilize seven real-world image dehazing datasets in our experiments: NH-Haze20 [5], NH-Haze21 [6], HD-NH-Haze [7], DenseHaze [4], I-Haze [2], O-Haze [3], and RTTS [22]. Specifically, NH-Haze20, NH-Haze21, HD-NH-Haze, DenseHaze, I-Haze, and O-Haze provide paired hazy and haze-free images, with 55, 25, 40, 55, 30, and 40 pairs, respectively. In contrast, the RTTS dataset only contains 4,322 hazy images without their haze-free versions. We conduct two experimental settings to assess the effectiveness of the proposed method, as summarized in Table 1. In **Setting1**, we use NH-Haze20 as the source domain for optimizing both PHATNet and the dehazing models. Consistent with prior studies [11, 16, 40], we employ the first 50 image pairs for training and the last 5 pairs for validation. The remaining datasets are used as target domains for performance evaluation. In **Setting2**, we use HD-NH-Haze as the source domain, training on the initial 35 pairs and validating on the last 5 pairs. The other datasets then serve as target domains for testing.

4.2. Implementation Details

PHATNet PHATNet is optimized on source-domain training pairs over 1,000 epochs, with a batch size of 1 and an image resolution of 1600×1200 on an Nvidia A5000 GPU. We employ the Adam optimizer with an initial learning rate of 10^{-4} , which decays to 10^{-7} following a cosine annealing schedule.

Dehazing Models We evaluate the effectiveness of PHATNet using four SOTA dehazing models: FocalNet [11], Dehazer [16], MITNet [40], and the latest SGDNet [13]. For each model, we utilize its source-domain pre-trained weights, if available; if not, we retrain the model on the source domain’s training set using its default settings. During the testing phase, each dehazing model is fine-tuned for one epoch on the fine-tuning set generated for each target domain.

4.3. Experimental Results

Quantitative Results We compare the dehazing performance of four SOTA baseline models and their PHATNet-enhanced versions in Tables 2 and 3, where “Baseline”

Table 2. Evaluation results of dehazing performance under **setting1**. We use NH-Haze20 [5] as the source domain and evaluate the performance on target domains, including NH-Haze21 [6], HD-NH-Haze [7], DenseHaze [4], I-Haze [2], O-Haze [3], respectively. “Baseline” and “+PHATNet” denote the dehazing performance without and with PHATNet.

Model		NH-Haze21		HD-NH-Haze		DenseHaze		I-Haze		O-Haze		Average	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
FocalNet [11]	Baseline	16.45	0.635	14.76	0.509	14.80	0.434	16.13	0.588	19.10	0.706	16.25	0.574
	+PHATNet	16.90	0.654	15.87	0.570	15.60	0.423	16.96	0.602	20.01	0.711	17.07	0.592
Dehamer [16]	Baseline	16.46	0.619	14.30	0.457	14.56	0.436	17.05	0.573	19.28	0.696	16.33	0.556
	+PHATNet	17.26	0.636	15.22	0.537	15.42	0.438	17.85	0.625	20.70	0.697	17.29	0.587
MITNet [40]	Baseline	14.21	0.516	11.60	0.355	13.67	0.420	15.72	0.540	18.96	0.701	14.83	0.506
	+PHATNet	16.21	0.650	14.04	0.466	14.73	0.428	16.54	0.557	20.02	0.708	16.31	0.562
SGDN [13]	Baseline	14.79	0.560	12.40	0.424	13.71	0.438	14.74	0.568	18.55	0.701	14.84	0.538
	+PHATNet	17.64	0.677	16.67	0.609	14.21	0.438	16.03	0.599	19.31	0.705	16.77	0.606

Table 3. Evaluation results of dehazing performance on **setting2**. We use HD-NH-Haze [7] as the source domain and evaluate on target domains, including NH-Haze20 [5], NH-Haze21 [6], DenseHaze [4], I-Haze [2], O-Haze [3], respectively. “Baseline” and “+PHATNet” denote the dehazing performance without and with PHATNet.

Model		NH-Haze20		NH-Haze21		DenseHaze		I-Haze		O-Haze		Average	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
FocalNet [11]	Baseline	15.12	0.468	19.90	0.799	11.34	0.323	13.92	0.565	16.64	0.558	15.38	0.543
	+PHATNet	15.77	0.505	20.06	0.808	12.95	0.323	14.64	0.593	17.66	0.586	16.22	0.563
Dehamer [16]	Baseline	14.82	0.475	18.62	0.748	12.57	0.320	16.06	0.569	17.67	0.576	15.95	0.538
	+PHATNet	15.10	0.478	18.73	0.747	13.61	0.331	16.65	0.570	18.95	0.585	16.61	0.542
MITNet [40]	Baseline	11.61	0.385	19.30	0.775	8.6	0.285	9.99	0.548	12.70	0.502	12.44	0.499
	+PHATNet	14.02	0.477	19.64	0.783	11.22	0.292	15.14	0.550	18.94	0.561	15.79	0.533
SGDN [13]	Baseline	11.87	0.414	19.95	0.810	8.29	0.232	11.01	0.524	13.84	0.522	12.99	0.500
	+PHATNet	14.71	0.502	20.33	0.815	11.92	0.310	15.57	0.572	18.34	0.572	16.17	0.554

Table 4. Evaluation results on the RTTS [22] dataset, where we respectively use NH-Haze20 [5] and HD-NH-Haze [7] as source domains.

		FocalNet [11]		Dehamer [16]		MITNet [40]		SGDN [13]	
		Baseline	+PHATNet	Baseline	+PHATNet	Baseline	+PHATNet	Baseline	+PHATNet
NH-Haze20 → RTTS	NIQE (↓)	4.19	3.92	4.20	3.82	4.42	4.33	4.30	3.85
	BRISQUE (↓)	34.88	28.34	31.94	24.15	33.67	31.99	34.27	22.21
HD-NH-Haze → RTTS	NIQE (↓)	4.45	4.40	4.49	4.33	4.69	4.62	4.74	4.58
	BRISQUE (↓)	32.78	28.93	34.67	23.74	30.21	28.84	29.71	27.63

indicates performance without PHATNet-augmented data, while “+PHATNet” means performance with it. Table 2 compares the PSNR performance under **Setting1**, showing substantial gains across all models using PHATNet. Specifically, the average PSNR gains enhanced by PHATNet across the five datasets are 0.82 dB for FocalNet, 0.96 dB for Dehamer, 1.48 dB for MITNet, and 1.93 dB for SGDN.

Similarly, Table 3 demonstrates that under **Setting2**, PHATNet consistently enhances the average PSNR performance across all datasets, with gains of 0.84 dB for FocalNet, 0.66 dB for Dehamer, 3.35 dB for MITNet, and 3.18 dB for SGDN. Table 4 further demonstrates the effectiveness of PHATNet on RTTS, a large real-world dataset without haze-free reference images. Employing two commonly-used no-reference image quality metrics, NIQE [33] and BRISQUE [32], we show that PHATNet consistently enhances performance on RTTS.

Qualitative Results Figure 5 displays haze-transferred images generated by PHATNet, with haze-free images from

NH-Haze20 and HD-NH-Haze training sets in the top and bottom rows, respectively. PHATNet effectively transfers region-specific haze patterns from unseen target domains to source-domain haze-free images, capturing realistic haze characteristics. Figure 6 presents dehazed images produced with and without using PHATNet-augmented data under **Setting1** and **Setting2**, respectively. These figures showcase PHATNet’s ability to enhance visual quality by reducing hazy artifacts, highlighting PHATNet’s robust domain adaptation capabilities. Additional qualitative results can be found in the supplementary material.

4.4. Ablation Studies

We evaluate the impact of PHATNet-augmented data on dehazing performance after adaptation, specifically using FocalNet as the dehazing model under **Setting1**. First, we examine the contributions of the ALM and TME in the multi-scale PHDT framework within PHATNet. Next, we assess how the Content-Leakage Loss affects PHATNet’s performance. At last, we compare PHATNet with existing domain



Figure 5. Qualitative results of haze-transferred images. In the top row, we transfer haze patterns from O-Haze [3] (top-left figure) and HD-NH-Haze [7] (top-right figure) to haze-free images from NH-Haze20 [5]. In the bottom row, we transfer haze patterns from NH-Haze20 [5] (bottom-left figure) and DenseHaze [4] (bottom-right figure) to haze-free images from HD-NH-Haze [7].




Figure 6. Qualitative performance comparisons of dehazed images conducted under **Setting1** and **Setting2**. Left: We demonstrate the dehazed images on HD-NH-Haze [7], NH-Haze21 [6], and DenseHaze [4] datasets under **Setting1**. Right: We demonstrate the dehazed images on O-Haze [3], I-Haze [2], and NH-Haze20 [5] datasets under **Setting2**.

Table 5. PSNR (dB) with ALE and TME within PHDT.

	ALE	TME	PSNR
Baseline			16.25
+CNN (Concat)			16.62
+ALE	✓		16.46
+TME		✓	16.80
Ours	✓	✓	17.07

adaptation approaches [10, 24] and haze synthesis methods [23, 46] for dehazing. In the following tables, “Baseline” refers to FocalNet trained solely on the source-domain training set without PHATNet-augmented data.

Component analysis of PHATNet Table 5 evaluates the individual and combined contributions of the Atmospheric Light Encoder (“+ALE”) and Transmission Map Encoder (“+TME”). “+CNN (Concat)” denotes that features extracted from both hazy and haze-free images are concatenated for fusion. Experimental results demonstrate that integrating ALE and TME performs better than using either component alone or the feature concatenation strategy. Table 6 indicates that employing a 3-stage PHDT structure in PHATNet optimizes dehazing performance, achieving the highest PSNR.

Table 6. Impact of the stage number of PHDT on performance.

	Baseline	1-stage	2-stage	3-stage
PSNR	16.25	16.68	16.95	17.07

Effectiveness of content-leakage loss Table 7 presents the effectiveness of Content-Leakage (CL) Loss, $\mathcal{L}_{CL}$, comparing the performance of PHATNet without CL Loss (“w/o $\mathcal{L}_{CL}$ ”) and with CL Loss (“w/ $\mathcal{L}_{CL}$ ”). The $\mathcal{L}_{CL}$ mitigates content leakage from hazy images into ALE and TME, generating more accurate haze-transferred images and enhancing dehazing performance. While PHATNet without $\mathcal{L}_{CL}$ improves dehazing results on NH-Haze21, HD-NH-Haze, DenseHaze, and I-Haze datasets, incorporating $\mathcal{L}_{CL}$ yields further PSNR gains across all datasets. This effect is particularly notable on O-Haze, which contains sparse haze, making the model more susceptible to content interference during feature disentanglement. Figure 7 provides visual comparisons of haze-transferred images generated with and without $\mathcal{L}_{CL}$.

Comparison with competing methods Table 8 compares the performance and complexity of PHATNet against two domain-adaptation methods, TMD [24] and PTTD [10], and

Table 7. Impact of Content-Leakage Loss ($\mathcal{L}_{CL}$) on PHATNet’s PSNR performance (dB).

	Baseline	w/o $\mathcal{L}_{CL}$	w/ $\mathcal{L}_{CL}$
NH-Haze21	16.45	<u>16.68</u>	16.90
HD-NH-Haze	14.76	<u>15.48</u>	15.87
DenseHaze	14.80	<u>15.14</u>	15.60
I-Haze	16.13	<u>16.37</u>	16.96
O-Haze	<u>19.10</u>	18.04	20.01




Figure 7. Illustration of haze-transferred images generated without and with the content-leakage loss, denoted “w/o $\mathcal{L}_{CL}$ ” and “w/ $\mathcal{L}_{CL}$ ”, within PHATNet, where we use the hazy image from O-Haze [3] and haze-free image from NH-Haze20 [5].

two haze-synthesis methods, HTFANet [23] and D4+ [46]. For a fair comparison, we use FocalNet as the dehazing model and follow **Setting1**, using NH-Haze20 training set to optimize the competing methods and reporting the average PSNR. We also compare the average runtime needed for augmenting a hazy image for the adaptation process. For the GAN-based HTFANet and D4+, both paired and unpaired data in the training set are used for training. Table 8 shows that PHATNet outperforms all competing methods in PSNR, demonstrating its effectiveness in domain adaptation for real-world image dehazing. With 26 million parameters, PHATNet requires 0.153 seconds to generate a 1600×1200 haze-transferred image on an Nvidia A5000 GPU, which is efficient and effective compared to the others. In addition, since the adaptation process is performed offline, there is no increase in latency for dehazing models during inference.

The qualitative comparison in Figure 8 further demonstrates PHATNet’s strong ability to generate realistic haze-transferred images. In particular, since HTFANet utilizes vanilla ASM to transfer haze patterns, it inevitably leaks content-related features in the transmission map, as shown in Figure 9. In contrast, PHATNet successfully disentangles haze-related features by extending ASM to the latent space, effectively mitigating ghosting artifacts caused by scene depth variations.

Discussions Beyond haze transfer, PHATNet can further augment hazy images through gamma correction and flipping in the ASM parametric domain (see Figure 3). However, achieving comprehensive hazy image augmentation in

Table 8. Comparison of performance and complexity between our method, domain-adaptation methods, TMD [24] and PTTD [10], and haze-synthesis methods, HTFANet [23] and D4+ [46].

	Baseline	TMD	HTFANet	D4+	PTTD	Ours
PSNR (dB)	16.25	15.71	16.21	<u>16.69</u>	16.02	17.07
Params (M)	–	5	<u>0.01</u>	20	0	26
Time (s)	–	1.328	0.040	<u>0.122</u>	0.237	0.153
Offline	–	×	✓	✓	×	✓




Figure 8. Comparison of transferred hazy images among HTFANet [23], D4+ [46], PTTD [10], and our method.




Figure 9. Comparison of transmission maps between HTFANet [23] and our method.

the parametric domain remains challenging due to limited prior knowledge of real-world haze distributions. This limitation highlights a valuable avenue for future research.

5. Conclusion

We presented an effective domain adaptation framework for image dehazing that leverages a domain-specific fine-tuning set, generated via haze transfer, to adapt dehazing models during testing. At the core of our approach is the Physics-guided Haze Transfer Network (PHATNet), which effectively disentangles and transfers target-domain haze patterns onto source-domain haze-free images. PHATNet incorporates the Atmospheric Scattering Model (ASM) as an inductive bias, enabling efficient separation of haze patterns and scene content within the ASM parametric space. To facilitate accurate haze disentanglement and transfer, we proposed the Haze-Transfer-Consistency Loss to ensure accurate transfer of target-domain haze patterns and the Content-Leakage Loss to prevent scene content leakage from target-domain hazy images. Extensive experiments have demonstrated that PHATNet significantly enhances the performance of state-of-the-art dehazing models across multiple real-world haze datasets, highlighting its effectiveness in domain adaptation for image dehazing.

Acknowledgments This work was supported in part by the National Science and Technology Council (NSTC) under grants 113-2634-F-002-003, 112-2221-E-A49-090-MY3, and 113-2221-E-004-006-MY2. This work was funded in part by MediaTek and NVIDIA Academic Grant.

References

- [1] Srivastava Akash, Valkov Lazar, Russell Chris, Gutmann Michael U., and Sutton Charles. VEEGAN: Reducing mode collapse in gans using implicit variational learning. In *NeurIPS*, 2017. 1
- [2] Codruta O. Ancuti, Cosmin Ancuti, Radu Timofte, and Christophe De Vleeschouwer. I-HAZE: a dehazing benchmark with real hazy and haze-free indoor images. In *ACIVS*, 2018. 1, 5, 6, 7
- [3] Codruta O. Ancuti, Cosmin Ancuti, Radu Timofte, and Christophe De Vleeschouwer. O-HAZE: a dehazing benchmark with real hazy and haze-free outdoor images. In *CVPRW*, 2018. 5, 6, 7, 8
- [4] Codruta O. Ancuti, Cosmin Ancuti, Mateu Sbert, and Radu Timofte. Dense-Haze: A benchmark for image dehazing with dense-haze and haze-free images. In *ICIP*, 2019. 5, 6, 7
- [5] Codruta O. Ancuti, Cosmin Ancuti, and Radu Timofte. NH-HAZE: an image dehazing benchmark with non-homogeneous hazy and haze-free images. In *CVPRW*, 2020. 5, 6, 7, 8
- [6] Codruta O. Ancuti, Cosmin Ancuti, Florin-Alexandru Vasluiianu, and Radu Timofte. Ntire 2021 nonhomogeneous dehazing challenge report. In *CVPRW*, 2021. 5, 6, 7
- [7] Codruta O. Ancuti, Cosmin Ancuti, Florin-Alexandru Vasluiianu, and Radu Timofte. Ntire 2023 hr nonhomogeneous dehazing challenge report. In *CVPRW*, 2023. 1, 5, 6, 7
- [8] Bolun Cai, Xiangmin Xu, Kui Jia, Chunmei Qing, and Dacheng Tao. DehazeNet: An end-to-end system for single image haze removal. *IEEE TIP*, 2016. 2
- [9] Chih-Ling Chang, Fu-Jen Tsai, Zi-Ling Huang, Lin Gu, and Chia-Wen Lin. PANet: A physics-guided parametric augmentation net for image dehazing by hazing, 2024. 3
- [10] Zixuan Chen, Zewei He, Ziqian Lu, and Zhe-Ming Lu. Prompt-based test-time real image dehazing: a novel pipeline. In *ECCV*, 2024. 3, 7, 8
- [11] Yuning Cui, Wenqi Ren, Xiaochun Cao, and Alois Knoll. Focal network for image restoration. In *ICCV*, 2023. 1, 2, 3, 5, 6
- [12] Qili Deng, Ziling Huang, Chung-Chi Tsai, and Chia-Wen Lin. HardGAN: A haze-aware representation distillation gan for single image dehazing. In *ECCV*, 2020. 1, 2
- [13] Wenxuan Fang, JunKai Fan, Yu Zheng, Jiangwei Weng, Ying Tai, and Jun Li. Guided real image dehazing using ycbcr color space. In *AAAI*, 2025. 5, 6
- [14] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NeurIPS*, 2014. 1
- [15] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C Courville. Improved training of wasserstein gans. In *NeurIPS*, 2017. 1
- [16] Chun-Le Guo, Qixin Yan, Saeed Anwar, Runmin Cong, Wenqi Ren, and Chongyi Li. Image dehazing transformer with transmission-aware 3d position embedding. In *CVPR*, 2022. 1, 2, 5, 6
- [17] Dong Hang, Pan Jinshan, Hu Zhe, Lei Xiang, Zhang Xinyi, Wang Fei, and Yang Ming-Hsuan. Multi-scale boosted dehazing network with dense feature fusion. In *CVPR*, 2020. 2
- [18] Jin-Ting He, Fu-Jen Tsai, Jia-Hao Wu, Yan-Tsung Peng, Chung-Chi Tsai, Chia-Wen Lin, and Yen-Yu Lin. Domain-adaptive video deblurring via test-time blurring. In *ECCV*, 2024. 3
- [19] Ming Hong, Yuan Xie, Cuihua Li, and Yanyun Qu. Distilling image dehazing with heterogeneous task imitation. In *CVPR*, 2020. 2
- [20] Lukas Hoyer, Dengxin Dai, Haoran Wang, and Luc Van Gool. MIC: Masked image consistency for context-enhanced domain adaptation. In *CVPR*, 2023. 1
- [21] Boyi Li, Xiulian Peng, Zhangyang Wang, Jizheng Xu, and Dan Feng. AOD-Net: All-in-one dehazing network. In *ICCV*, 2017. 2
- [22] Boyi Li, Wenqi Ren, Dengpan Fu, Dacheng Tao, Dan Feng, Wenjun Zeng, and Zhangyang Wang. Benchmarking single-image dehazing and beyond. *IEEE TIP*, 2019. 1, 5, 6
- [23] Huafeng Li, Jirui Gao, Yafei Zhang, Minghong Xie, and Zhengtao Yu. Haze transfer and feature aggregation network for real-world single image dehazing. *Know.-Based Syst.*, 2022. 1, 3, 7, 8
- [24] Huan Liu, Zijun Wu, Liangyan Li, Sadaf Salehkalalbar, Jun Chen, and Keyan Wang. Towards multi-domain single image dehazing via test-time training. In *CVPR*, 2022. 3, 7, 8
- [25] Jiaming Liu, Ran Xu, Senqiao Yang, Renrui Zhang, Qizhe Zhang, Zehui Chen, Yandong Guo, and Shanghang Zhang. Continual-MAE: Adaptive distribution masked autoencoders for continual test-time adaptation. In *CVPR*, 2024. 1
- [26] Po-Sheng Liu, Fu-Jen Tsai, Yan-Tsung Peng, Chung-Chi Tsai, Chia-Wen Lin, and Yen-Yu Lin. Meta transferring for deblurring. In *BMVC*, 2022. 3
- [27] Xiaohong Liu, Yongrui Ma, Zhihao Shi, and Jun Chen. Grid-dehazenet: Attention-based multi-scale network for image dehazing. In *ICCV*, 2019. 2
- [28] Yajing Liu, Shijun Zhou, Xiyao Liu, Chunhui Hao, Baojie Fan, and Jiandong Tian. Unbiased faster R-CNN for single-source domain generalized object detection. In *CVPR*, 2024. 1
- [29] Xianzheng Ma, Zhixiang Wang, Yacheng Zhan, Yinqiang Zheng, Zheng Wang, Dengxin Dai, and Chia-Wen Lin. Both style and fog matter: Cumulative domain adaptation for semantic foggy scene understanding. In *CVPR*, 2022. 1
- [30] Qi Mao, Hsin-Ying Lee, Hung-Yu Tseng, Siwei Ma, and Ming-Hsuan Yang. Mode seeking generative adversarial networks for diverse image synthesis. In *CVPR*, 2019. 1
- [31] Xudong Mao, Qing Li, Haoran Xie, Raymond Y.K. Lau, Zhen Wang, and Stephen Paul Smolley. Least squares generative adversarial networks. In *ICCV*, 2017. 1

- [32] Anish Mittal, Anush Krishna Moorthy, and Alan Conrad Bovik. No-reference image quality assessment in the spatial domain. *IEEE TIP*, 2012. 6
- [33] Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a “completely blind” image quality analyzer. *IEEE Signal Processing Letters*, 2013. 6
- [34] Yanwei Pang, Jing Nie, Jin Xie, Jungong Han, and Xuelong Li. BidNet: Binocular image dehazing without explicit disparity estimation. In *CVPR*, 2020. 2
- [35] Xu Qin, Zhilin Wang, Yuanchao Bai, Xiaodong Xie, and Huizhu Jia. FFA-Net: Feature fusion attention network for single image dehazing. In *AAAI*, 2020. 1, 2
- [36] Yuwei Qiu, Kaihao Zhang, Chenxi Wang, Wenhan Luo, Hongdong Li, and Zhi Jin. MB-TaylorFormer: Multi-branch efficient transformer expanded by taylor formula for image dehazing. In *ICCV*, 2023.
- [37] Yanyun Qu, Yizi Chen, Jingying Huang, and Yuan Xie. Enhanced pix2pix dehazing network. In *CVPR*, 2019. 2
- [38] Harsh Rangwani, Sumukh K Aithal, Mayank Mishra, Arihant Jain, and R. Venkatesh Babu. A closer look at smoothness in domain adversarial training. In *ICML*, 2022. 1
- [39] Wenqi Ren, Si Liu, Hua Zhang, Jinshan Pan, Xiaochun Cao, and Ming-Hsuan Yang. Single image dehazing via multi-scale convolutional neural networks. In *ECCV*, 2016. 2
- [40] Hao Shen, Zhong-Qiu Zhao, Yulun Zhang, and Zhao Zhang. Mutual information-driven triple interaction network for efficient image dehazing. In *ACM MM*, 2023. 1, 2, 3, 5, 6
- [41] Yuda Song, Zhuqing He, Hui Qian, and Xin Du. Vision transformers for single image dehazing. *IEEE TIP*, 2023. 1, 2
- [42] Hsing-Hua Wang, Fu-Jen Tsai, Yen-Yu Lin, and Chia-Wen Lin. TANet: Triplet attention network for all-in-one adverse weather image restoration. In *ACCV*, 2024. 1
- [43] Haiyan Wu, Yanyun Qu, Shaohui Lin, Jian Zhou, Ruizhi Qiao, Zhizhong Zhang, Yuan Xie, and Lizhuang Ma. Contrastive learning for compact single image dehazing. In *CVPR*, 2021. 1, 2
- [44] Jia-Hao Wu, Fu-Jen Tsai, Yan-Tsung Peng, Chung-Chi Tsai, Chia-Wen Lin, and Yen-Yu Lin. ID-Blau: Image deblurring by implicit diffusion-based reblurring augmentation. In *CVPR*, 2024. 3
- [45] Senqiao Yang, Zhuotao Tian, Li Jiang, and Jiaya Jia. Unified language-driven zero-shot domain adaptation. In *CVPR*, 2024. 1
- [46] Yang Yang, Chaoyue Wang, Xiaojie Guo, and Dacheng Tao. Robust unpaired image dehazing via density and depth decomposition. *IJCV*, 2024. 1, 3, 7, 8
- [47] Hu Yu, Naishan Zheng, Man Zhou, Jie Huang, Zeyu Xiao, and Feng Zhao. Frequency and spatial dual guidance for image dehazing. In *ECCV*, 2022. 1
- [48] Hu Yu, Naishan Zheng, Man Zhou, Jie Huang, Zeyu Xiao, and Feng Zhao. Frequency and spatial dual guidance for image dehazing. In *ECCV*, 2022. 2
- [49] Dong Zhao, Shuang Wang, Qi Zang, Licheng Jiao, Nicu Sebe, and Zhun Zhong. Stable neighbor denoising for source-free domain adaptive segmentation. In *CVPR*, 2024. 1
- [50] Zhuoran Zheng, Wenqi Ren, Xiaochun Cao, Xiaobin Hu, Tao Wang, Fenglong Song, and Xiuyi Jia. Ultra-high-definition image dehazing via multi-guided bilateral learning. In *CVPR*, 2021. 1, 2