

RTE-FM-Dehazer: Radiative Transfer Equation Inspired Flow Matching for Real-World Image Dehazing

Chenfeng Wei^{1,5}, Chun Wang², Boyang Zhao³, Si Zuo⁴, Shenhong Wang^{1,*}, and Chenguang Yang^{5,6,*}

¹Xi'an Jiaotong-Liverpool University, China ²Mashang Consumer Finance Co. Ltd, China ³Tsinghua University, China ⁴Hunan University, China ⁵University of Liverpool, UK ⁶The Hong Kong Polytechnic University, Hong Kong SAR, China

* Corresponding authors: Shenhong.Wang@xjtlu.edu.cn,
Chenguang.Yang@liverpool.ac.uk

Abstract. Single-image dehazing aims to recover a clear scene from a hazy image and is generally formulated as an image-to-image translation task; however, it faces two limitations. Its performance depends heavily on the haze-formation priors embedded in the model. Prevailing methods adopt the Atmospheric Scattering Model (ASM), whose assumptions of single scattering and homogeneous media are often violated, leading to residual haze and color drift. Moreover, large-scale real hazy/clear pairs are impractical to collect, and existing synthesis approaches fail to reproduce the full complexity of natural haze. To address these issues, we present RTE-FM-Dehazer, a novel dehazing approach, together with a scalable data pipeline. Unlike the ASM, the Radiative Transfer Equation (RTE) jointly accounts for both scattering and absorption, naturally accommodating the non-homogeneous, multiple-scattering media that characterize real hazy scenes. Motivated by the structural similarity between the RTE diffusion-absorption term and the ODE in flow matching, we introduce a diffusion-absorption regularizer derived from a reduced RTE, to steer the flow matching trajectory at each step. Next, leveraging modern vision-language models, we build an automated pipeline and release P-HAZE, a dataset of 50 000 realistic hazy/clear pairs. Extensive evaluations demonstrate that RTE-FM-Dehazer, trained solely on P-HAZE, effectively eliminates artifacts like residual haze and color drift, exhibits strong cross-domain generalization, and achieves leading results on five real-world dehazing benchmarks. Code and data are available at <https://github.com/vincentweikey/RTE-FM-Dehazer>.

Keywords: Image Dehazing · Flow Matching · Physics-inspired Regularization

1 Introduction

Haze severely degrades visual quality, posing challenges for real-world applications such as autonomous driving, remote sensing, and outdoor scene analysis.



Fig. 1: Our **RTE-FM-Dehazer** efficiently removes spatially-varying haze under complex weather and scenes while preserving original colors and fine textures.

In the real world, haze is spatially uneven and light is multiply scattered, leaving robust dehazing still an active research area. Dehazing is an ill-posed problem; current methods typically impose the atmospheric scattering model (ASM) [26] as a physical prior to regularize the solution.

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

where $T(\mathbf{x}) \in [0, 1]$ is the pixel-wise transmission, $J(\mathbf{x})$ is the clean radiance we wish to recover, and $A \in \mathbb{R}^3$ is the global air-light introduced by suspended particles. Estimating both T and J from a single RGB image is inherently under-constrained. Early methods using hand-crafted priors: Dark-Channel statistics [17], Color-Line distributions [13], Haze-Line clusters [6], or Color-Attenuation ratios [46] to estimate the transmission map. These cues are elegant, interpretable, and sufficient when the haze is spatially homogeneous, optically thin, and free of diffuse reflection.

In real-world hazy images, however, these assumptions are routinely violated: haze density varies with scene depth, particles exhibit anisotropic scattering, and multiple scattering dominates. Under such complex conditions, traditional priors break down, leaving residual fog, severe color casts, and over-enhanced artifacts. Recognizing these limitations, recent studies have shifted to learning direct domain-to-domain mappings via powerful deep learning architectures including CNNs [7, 10, 21, 29, 40] and Transformers [9, 16] etc, achieving impressive performance on synthetic benchmarks. However, these data-driven methods heavily rely on idealized datasets: most training pairs are synthesized under a spatially constant-haze assumption, thus limiting generalization to real-world imagery. To address this, attention has shifted toward better data-synthesis strategies [9, 32]. Building on this momentum, powerful diffusion-based generative models [11, 36] and VQGAN codebook priors [14, 38] have recently been exploited to improve dehazing performance in real-world settings. Nevertheless, these generative methods introduce a new class of artifacts: the stochastic nature of diffusion sampling may drift toward adjacent color regions, while codebook-based retrieval occasionally grafts patches lying outside the original field of view.

Single-image dehazing remains challenging due to spatially-varying haze, scarce training pairs, and generative artifacts. We address these jointly: we replace the restrictive ASM with the RTE [43] to model scattering and absorption; we regularize flow matching with physics-inspired deterministic trajectories; and we leverage VLMs with geometric alignment to synthesize scalable training data. Radiative transfer in participating media is governed by the RTE [43]:

$$\frac{dI}{ds} = -\kappa_a I + \kappa_a B(T, \nu) + \kappa_s \int_{\mathbb{S}^2} p(\hat{\mathbf{s}}, \hat{\mathbf{s}}') I(\mathbf{x}, \hat{\mathbf{s}}', \nu) d\Omega', \quad (2)$$

where I is ray specific intensity, κ_a and κ_s are absorption and scattering coefficients, and the integral captures anisotropic multiple scattering. While the RTE strictly governs radiative transfer in participating media, our latent-space formulation appropriates its diffusion-absorption structure as a soft prior: the learned terms approximate attenuation, soft halos, and varying contrast without the restrictive homogeneity assumptions of the ASM.

Rather than performing a single-step inversion, we regard the hazy observation $\mathbf{I}(\mathbf{x})$ as the initial state of a continuous light-transport trajectory whose equilibrium is the haze-free image $\mathbf{J}(\mathbf{x})$. Dehazing thus reduces to computing a physically plausible path $\mathbf{I} \rightarrow \mathbf{J}$ locally regularized by the RTE at each step—a requirement ignored by prior deep networks. Flow matching [24] supplies the ideal vehicle: it evolves the image with an ODE whose velocity field v_θ , once learned, permits single-step Euler integration at test time. We therefore introduce **RTE-FM-Dehazer**, a physics-inspired flow matching framework that regularises v_θ with an RTE-derived velocity term. To avoid scarce real pairs, we prompt a VLM [5, 35] to generate hazy counterparts and align them via dense geometry-aware descriptors and sub-pixel homography warping. Qualitative results on diverse scenes are presented in Figure. 1. Our main contributions are:

1. Reformulating single-image dehazing as RTE-regularized flow matching improves physical plausibility by aligning the generative trajectory via a physics-inspired prior, leading to more stable and realistic dehazed results across diverse real-world conditions.
2. Removing stochastic noise priors in the flow matching generation process yields a deterministic latent trajectory regularized by a diffusion-absorption term, maintaining the overall color distribution better than random-sampling baselines.
3. We propose a scalable, VLM-driven pipeline that aligns features and synthesizes realistic haze-clean pairs without requiring scarce real-world data.

2 Related Work

2.1 Single Image Dehazing

Early single-image dehazing approaches are grounded in the atmospheric scattering model and aim to invert the haze formation process by estimating the

transmission map. He et al. [17] proposed the Dark Channel Prior, which exploits the observation that at least one color channel of local patches in haze-free images tends to have very low intensity. Subsequent priors extended this philosophy: the Haze-line Prior [6] clusters colors along 1-D haze-lines, while the Color Attenuation Prior [46] models the relationship between scene depth, brightness and saturation. These handcrafted priors yield impressive results under moderate haze, yet their linear or low-order statistical assumptions break down for highly textured regions, bright-white objects, or non-uniform illumination, yielding color distortions and halo artifacts.

Driven by large-scale synthetic haze datasets [9, 22, 42] and the success of CNNs in low-level vision, recent dehazing research has predominantly adopted learning-based approaches. Early CNN-based methods such as DehazeNet [7] and AOD-Net [21] directly regress the transmission map or the haze-free image from input observations. Subsequent works have introduced architectural enhancements, including multi-scale feature fusion and attention mechanisms, to enlarge receptive fields and disentangle haze related representations. Notably, FFA-Net [29] employs channel and pixel attention to emphasize informative features, while DEA-Net [10] integrates detail enhanced convolution with content-guided attention for improved texture restoration. More recent Transformer-based models, such as Dehaze3D [16] and DehazeXL [9], leverage 3D position embeddings and patch tokenization to capture long-range dependencies and global context in high-resolution images. Despite achieving state-of-the-art PSNR and SSIM [18] on synthetic benchmarks, these supervised methods often suffer from significant performance degradation in real-world scenarios due to domain shift. This limitation stems from the oversimplified assumptions of synthetic haze generation, which typically employs homogeneous scattering models and neglects the wavelength-dependent, spatially-varying, mixed scattering-absorption characteristics of real-world haze.

2.2 Dehazing in real-world

Confronted with the scarcity of paired real hazy/clear imagery and the limited distributional coverage of synthetic datasets, recent work has sought to mitigate the generalization gap through domain adaptation [32], unsupervised priors such as the dark-channel [15], and contrastive regularization [37, 44]. a parallel line of work treats dehazing as a conditional generation problem. Li et al. [23] first brought conditional GANs to single-image dehazing, learning a deterministic pixel-to-pixel mapping via adversarial supervision; however, such models demand pixel-perfect aligned pairs and tend to yield over-smoothed predictions once the synthetic training distribution is too narrow. Diffusion probabilistic models [12] relax this determinism by gradually denoising pure noise toward the clean data manifold. Cheng et al. [11] instantiate this idea with DehazeDiff, yet the stochastic reverse process remains sensitive to the variance schedule and the number of denoising steps, making the output quality difficult to control in practice. Vector-quantized auto-encoders (VQ-VAE [31], VQ-GAN [39]) provide

a compact discrete latent space that preserves high-frequency details while enabling autoregressive or diffusion-based priors. RIDCP [38] and IPC-Dehaze [14] incorporate a pre-trained VQ-GAN as a structural prior and finetune the decoder for haze-free reconstruction. While these approaches produce visually plausible dehazing on real scenes, their generative nature is unconstrained by any physical haze-formation model and thus tends to shift global illumination, bias the original color distribution, and introduce extraneous artifacts. We retain the generative pipeline but embed approximations of the RTE at every flow-matching step, eliminating these artifacts while preserving visual realism.

3 Methodology

3.1 Core Framework

In this section we introduce RTE-FM-Dehazer, a flow matching framework regularized by the Radiative Transfer Equation (RTE). As shown in Figure 2, flow matching decomposes hazy-to-clean mapping into infinitesimal steps (unlike one-shot VQ-VAE or VQ-GAN), enabling fine-grained injection of RTE-derived absorption and scattering constraints. The network predicts a velocity field that simultaneously satisfies the flow matching objective and aligns with physical dynamics, producing trajectories that approximate valid light transport via ODE integration.




Fig. 2: RTE-FM-Dehazer pipeline. In the macro process, a frozen VAE encodes the hazy image into latent space, where flow matching learns a neural velocity field that progressively transports the haze latent toward its clean counterpart. In the micro process, each step refines this velocity so that it aligns with both the data-driven direction and the diffusion-absorption direction estimated by the RTE.

3.2 RTE-FM-Dehazer Architecture

Vanilla Flow Matching Dehazing. We formulate single-image dehazing as an *Vanilla flow matching* process in the latent space of Stable Diffusion. Given

a hazy image $x_{\text{hazy}} \in \mathbb{R}^{3 \times H \times W}$, we first encode it to $\mathbf{z}_0 = \mathcal{E}(x_{\text{hazy}}) \in \mathbb{R}^{4 \times h \times w}$ with the frozen VAE encoder $\mathcal{E}$. A U-Net is trained to match the *straight-line* velocity

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{t \sim \mathcal{U}(0,1)} \left[\left\| \mathbf{v}_\theta(\mathbf{z}_t, t) - (\mathbf{z}_1 - \mathbf{z}_0) \right\|^2 \right], \quad (3)$$

$$\mathbf{z}_t = (1 - t)\mathbf{z}_0 + t\mathbf{z}_1, \quad (4)$$

where $\mathbf{z}_1$ is the latent code of the clean target.

At inference we solve the ODE $\frac{d\mathbf{z}}{dt} = \mathbf{v}_\theta(\mathbf{z}, t)$ by *explicit Euler* with fixed step Δt :

$$\mathbf{z}_{n+1} = \mathbf{z}_n + \Delta t \cdot \mathbf{v}_\theta(\mathbf{z}_n, t_n). \quad (5)$$

Crucially, $\mathbf{v}_\theta$ outputs a *full-resolution velocity map*, so every spatial location receives its own vector. Consequently *non-uniform haze* is processed with *location-specific speeds*—bright regions, dark regions, and edges are advanced along different directions and magnitudes within the same time step. The method is purely data-driven, learning the velocity field end-to-end without transmission maps or physical priors.

Flow Matching with RTE Regularizer. While flow matching regresses to the straight-line velocity $\mathbf{z}_1 - \mathbf{z}_0$, it is agnostic to the physical propagation of light in participating media. Accordingly, we start from the diffusion approximation [43] of the RTE

$$\frac{\partial u}{\partial t} = D \nabla^2 u - \mu_a u, \quad (6)$$

$D \nabla^2 u$ governs *spatial diffusion*: it redistributes radiative energy from high to low concentration regions, thereby smoothing sharp gradients and suppressing artificial discontinuities in the recovered image. $-\mu_a u$ models *medium absorption*: it induces exponential decay of local radiative energy, directly mimicking the attenuating effect of haze or smoke and thus driving the latent state toward a cleaner image manifold. Where the right-hand side simply expresses the instantaneous rate of change of the density field $u(\mathbf{x}, t)$. In continuum mechanics such a time derivative *is* the Eulerian velocity of the field.

$$\mathbf{v}_{\text{rte}}(\mathbf{x}, t) \triangleq D \nabla^2 u(\mathbf{x}, t) - \mu_a u(\mathbf{x}, t) \quad (7)$$

Drawing an analogy between the physical radiance field u and the latent representation $\mathbf{z} \in \mathbb{R}^{4 \times 64 \times 64}$, we approximate the RTE dynamics in latent space. This approximation is justified because flow matching proceeds via infinitesimal velocity fields, restricting the regularization to the tangent space of the latent manifold where the pretrained decoder acts as a local isometry; consequently, the latent-space diffusion operator induces a transport that is structurally analogous to the physical radiative transfer in the image domain. After flipping the sign so that density decrease corresponds to haze removal, this yields:

$$\mathbf{v}_{\text{RTE}}(\mathbf{z}) \triangleq -D \nabla^2 \mathbf{z} + \mu_a \mathbf{z} \quad (8)$$

Latent-Space Discretization of RTE Dynamics. To evaluate the RTE-derived regularizer efficiently on the latent manifold, we discretize the continuous diffusion operator using the five-point stencil:

$$\nabla^2 \mathbf{z}(x, y) \approx \frac{\mathbf{z}(x+h, y) + \mathbf{z}(x-h, y) + \mathbf{z}(x, y+h) + \mathbf{z}(x, y-h) - 4\mathbf{z}(x, y)}{h^2} \quad (9)$$

where h is the grid spacing in the 64×64 latent space (typically $h = 1$), and $(x, y) \in \{1, \dots, H\} \times \{1, \dots, W\}$ denote the integer spatial coordinates indexing the location of latent feature vector $\mathbf{z}(x, y) \in \mathbb{R}^4$ on the feature map. This discrete approximation is evaluated efficiently using convolution with the kernel

$$\mathbf{K}_{\nabla^2} = \begin{bmatrix} 0 & 1 & 0 \\ 1 & -4 & 1 \\ 0 & 1 & 0 \end{bmatrix}, \quad (10)$$

D is estimated by a 1×1 convolution followed by a soft-plus activation. $\mu_a = \kappa \langle |\mathbf{z}_t| \rangle$ with $\kappa = 0.1$ (empirically set to balance the absorption term against the diffusion term in the latent domain), yielding a decoding-free proxy for local absorption strength.

Regularisation mechanism. Flow matching trains $\mathbf{v}_\theta$ to reproduce the optimal-transport vector $\mathbf{z}_1 - \mathbf{z}_0$. We now demand that, at every t , $\mathbf{v}_\theta$ is simultaneously (i) aligned with the data-driven direction and (ii) consistent with the RTE direction. The combined loss

$$\mathcal{L} = \underbrace{\|\mathbf{v}_\theta - (\mathbf{z}_1 - \mathbf{z}_0)\|^2}_{\text{transport fit(i)}} + \lambda \underbrace{\|\mathbf{v}_\theta - \mathbf{v}_{\text{RTE}}\|^2}_{\text{rte fit(ii)}} \quad (11)$$

performs an L^2 *projection*: minimising the second term encourages $\mathbf{v}_\theta$ to satisfy the diffusion-absorption balance locally, while the first term ensures the solution remains compatible with the clean target. Consequently the ODE

$$\frac{d\mathbf{z}}{dt} = \mathbf{v}_\theta(\mathbf{z}, t) \quad (12)$$

no longer follows a bare straight line but a regularised trajectory whose tangent is the closest (in L^2) vector field that fulfils the RTE at every spatial index. Mathematically, this embeds the continuous operator

$$(-D\nabla^2 + \mu_a I) \quad (13)$$

into the learning problem, yielding a physically regularised transport equation

$$\frac{d\mathbf{z}}{dt} = \Pi[(\mathbf{z}_1 - \mathbf{z}_0) \mid (-D\nabla^2 + \mu_a I)\mathbf{z}_t], \quad (14)$$

where Π denotes the L^2 projection enforced by the combined loss, and I is the identity operator acting point-wise on the latent feature map. The final latent trajectory therefore satisfies both the data distribution and the latent-space analogue of RTE, producing dehazed outputs that are sharper and physically consistent while remaining end-to-end trainable.




Fig. 3: Overview of clean-to-haze image generation: Stage 1: conditioned on a text prompt, Qwen2-VL or Gemini-2-Flash edits the clean image to produce a realistic hazy image. Stage 2: dense keypoints are extracted from both images and used to estimate a homography that warps the hazy frame into pixel-wise alignment with the clean image, yielding the final training pair.

3.3 Data Preparation

Encouraged by recent advances in Vision-Language Models (VLMs), we therefore devise a scalable, scene-agnostic, hazy-image generation pipeline, illustrated in Figure 3.

Stage 1: VLM-based hazy image generation. Given a reference image, we first sample a stochastic haze prompt—two typical examples are:

- “Inject a dense, irregular billow of heavy smoke while preserving original geometry, crisp edges, and dominant textures; zero phantom haze patches, zero repeating motifs, zero discontinuities.”
- “Add a thick, chaotic smoke plume that obeys the underlying shape and fine details; no fake fog spots, no tiled patterns, no breaks.”

The reference–prompt pair is then fed into Qwen2.5-VL [5] or Gemini [35] to produce an initial hazy frame. Advanced VLMs implicitly capture spatial and geometric relationships, enabling the generated haze to exhibit density variations, depth-dependent attenuation, and realistic color shifts without requiring handcrafted depth or weather parameters. Compared to ASM-based synthesis,

VLM generation covers arbitrarily complex scenes and better reproduces non-stationary aerosol distributions—swirling smoke, localized fog pockets, etc. Because the process is prompt-driven, we can endlessly expand the training set by sampling new prompts, whereas ASM datasets remain constrained by curated depth maps and weather priors.

Stage-2: Pixel alignment. Although the hazy images produced by VLMs are visually realistic, their pixel coordinates are not accurately aligned with the corresponding clean inputs; most of them exhibit slight distortions that closely resemble those induced by changes in camera intrinsics. To address this, we first extract dense keypoints on both clean and VLM-hazy frames and establish pixel-wise correspondences [33]. These matches are then used to estimate a planar homography $\mathbf{H} \in \mathbb{R}^{3 \times 3}$ with typically 800–1 500 inliers. Finally, the hazy frame is bilinear-warped onto the clean image via

$$I_{\text{aligned}}(\mathbf{x}) = I_{\text{hazy}}(\mathbf{H}^{-1}\mathbf{x}), \quad (15)$$

To suppress semantic artifacts (*e.g.*, phantom objects or boundary distortions) in VLM outputs, we adopt a two-stage filtering strategy: automatic segmentation verification via SAM2 [30] followed by manual human-in-the-loop inspection.

P-HAZE: A 50k Generate Dehazing Dataset. Using this scalable pipeline, we produce P-HAZE, a large-scale dehazing dataset that comprises approximately 50 000 rigorously aligned clear/haze pairs. Reference images are gathered from Mapillary Vistas [27], Flickr2K [1], Unsplash2K [20] and DIV2K [1], guaranteeing broad coverage of weather conditions, seasons, object categories, resolutions, and camera sensors; the resulting clear/hazy pairs are sufficiently diverse to support robust, real-world dehazer training.

3.4 Training and Inference

The training and inference procedures of RTE-FM-Dehazer are built on the flow matching framework to deliver an efficient, high-fidelity dehazing system.

For every hazy–clear pair (x_0, x_1) we sample $t \sim \mathcal{U}(0, 1)$, build the linear interpolation $x_t = (1 - t)x_0 + tx_1$, and minimise the flow matching loss plus an RTE-consistency term (see Algorithm 1). After training, the dehazing ODE $\frac{dx}{dt} = v_\theta(x, t)$ can be solved via T -step Euler integration:

$$x_1 \approx x_0 + \frac{1}{T} \sum_{i=0}^{T-1} v_\theta\left(x_{i/T}, \frac{i}{T}\right), \quad (16)$$

which directly estimates the clean image from the hazy input using the vector field at $t=0$. While efficient, this approach may not fully capture the complexity of the image transformation. To maximize visual quality, we recommend using more than 10 Euler steps. By subdividing the time interval $[0, 1]$, the model more accurately simulates the gradual transition from hazy to clean images,

Algorithm 1 RTE-FM-Dehazer training step for one hazy-clean pair (x_0, x_1) .**Require:** Hazy image x_0 , clean image x_1 , network parameters θ , RTE-weight λ .**Ensure:** Updated parameters θ .

- 1: Sample $t \sim \mathcal{U}(0, 1)$
- 2: Interpolate state: $x_t \leftarrow (1 - t)x_0 + tx_1$
- 3: Estimate coefficients: $(D, \mu_a) \leftarrow \phi_\theta(x_t)$
- 4: Compute Laplacian: $\nabla^2 x_t \leftarrow \text{LaplacianKernel} * x_t$
- 5: Build RTE velocity: $\hat{v}_{\text{RTE}} \leftarrow \text{RTE-velocity}(D, \mu_a, \nabla^2 x_t)$
- 6: Predict neural velocity: $v_\theta \leftarrow v_\theta(x_t, t)$
- 7: Compute losses:
- 8: Flow matching loss: $\mathcal{L}_{\text{FM}} \leftarrow \|v_\theta - (x_1 - x_0)\|_2^2$
- 9: RTE-consistency loss: $\mathcal{L}_{\text{RTE}} \leftarrow \|v_\theta - \hat{v}_{\text{RTE}}\|_F^2$
- 10: Total loss: $\mathcal{L} \leftarrow \mathcal{L}_{\text{FM}} + \lambda \mathcal{L}_{\text{RTE}}$
- 11: Back-propagate $\nabla_\theta \mathcal{L}$ and update θ

producing higher-quality dehazed results, mitigating errors inherent to single-step approximations and better preserving the detail of the image and the color fidelity.

To reduce memory usage and accelerate training, we encode images into a compact latent space via the $8\times$ autoencoder of Stable Diffusion v2.1 (SD2.1) [28]. This encoder downsamples images by a factor of 64 while preserving fine textures, enabling end-to-end latent training with 512×512 inputs on consumer-level GPUs.

4 Experiments

4.1 Experimental Settings

We train solely on our synthetic P-HAZE dataset and evaluate generalization across five benchmarks that span distinct haze-generation mechanisms. I-HAZE [4] provides indoor fog-chamber images under uniform scattering. D-HAZE [3] contains suburban scenes filled with dense glycol fog, resulting in extremely low visibility. SMOKE [19] comprises fire-smoke photographs produced with smoke generators. NH-HAZE [2] supplies outdoor 4K sequences where depth-varying, non-uniform haze coexists with bright sky regions. RESIDE-6K [22] is a synthetic set with haze parameters derived from the accompanying depth maps.

We adopt PSNR (Peak Signal-to-Noise Ratio), SSIM (Structural Similarity Index) [18], and LPIPS (Learned Perceptual Image Patch Similarity) [41] for evaluation. PSNR quantifies pixel-level fidelity, SSIM measures structural similarity (luminance, contrast, edges), and LPIPS evaluates deep perceptual distance; higher PSNR/SSIM and lower LPIPS indicate better quality.

4.2 Implementation Details

All training images are resized to 512×512 . Latent codes are pre-computed to accelerate training. The model is trained for 200 epochs on 50 000 synthetic

hazy/clear pairs, consuming ≈ 20 h on a single RTX 4090. Training adopts Adam with an initial learning rate of $1e-4$, decayed by half every 50 epochs, and a batch size of 64. We set the RTE weight to $\lambda = 0.5$. Empirically, our method is insensitive to this choice, exhibiting stable performance across the range $[0.3, 0.7]$; we select 0.5 as a balanced nominal value within this operational window.

4.3 Network Architecture

Our method builds upon a U-Net backbone adapted for latent-space RTE flow matching. The framework comprises two main components:

1. Encoder–Decoder U-Net:

- Encoder: four-level downsampling with ResNet-style blocks and spatial self-attention at the lowest resolution.
- Decoder: symmetric four-level upsampling with skip connections.
- Bottleneck: two ResBlocks with self-attention for global context.

2. Latent Generator: a pre-trained VAE encoder (SD 2.1) that maps hazy/clean images to a $64 \times 64 \times 4$ latent space, cutting computational cost while preserving perceptual quality.

4.4 Comparisons with State-of-the-art Methods

Quantitative Evaluation. We conduct quantitative comparisons on six public dehazing benchmarks that jointly cover synthetic, real indoor, real outdoor, and smoke-laden hazy scenes. As reported in Table 1, RTE-FM-Dehazer achieves the highest PSNR on four datasets and the lowest LPIPS on five out of six, demonstrating superior pixel fidelity and perceptual quality. In detail, on the challenging NH-HAZE dataset, we outperform the strongest competitor LHTD by +3.9 dB in PSNR and reduce LPIPS by 37%; on the SMOKE dataset, the gain remains +3.0 dB in PSNR and −37% in LPIPS. Although DehazeFormer and DEANet achieve higher SSIM on the synthetic RESIDE-6K, their PSNR drops by more than 9 dB on the four real-captured benchmarks, confirming the stronger generalisation of RTE-FM-Dehazer to real-world haze.

Table 1: Quantitative comparison on dehazing benchmarks. **Bold:** best; underlined: second best. RTE-FM-Dehazer is trained solely on P-HAZE. We run the official code of every mentioned methods and evaluate with IQA-PyTorch [8].

Methods	D-HAZE			I-HAZE			NH-HAZE			SMOKE			P-HAZE			RESIDE-6K		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
DCP [17] (CVPR2009)	11.38	0.438	0.704	11.56	0.542	0.355	12.04	0.515	0.505	11.75	0.488	0.405	16.87	0.762	0.303	17.89	0.867	0.097
AODNet [21] (ICCV2017)	8.17	0.316	0.873	9.85	0.525	0.487	9.31	0.328	0.700	10.26	0.394	0.570	9.43	0.470	0.505	9.35	0.599	0.324
DehazeFormer [34] (TIP2023)	11.04	0.390	0.751	16.54	0.766	0.221	12.63	0.521	0.512	13.47	0.527	0.419	17.98	0.765	0.282	<u>31.42</u>	0.981	0.012
RIDCP [38] (CVPR2023)	9.96	<u>0.474</u>	0.618	16.71	0.817	<u>0.179</u>	12.61	<u>0.631</u>	0.417	13.23	<u>0.565</u>	0.345	16.68	0.749	0.256	20.40	0.840	0.104
C2PNet [45] (CVPR2023)	11.22	0.388	0.761	16.51	0.745	0.247	12.59	0.519	0.516	13.72	0.532	0.420	16.84	0.735	0.293	24.90	0.937	0.034
DEANet [10] (TIP2024)	10.09	0.396	0.736	<u>16.96</u>	0.766	0.220	12.38	0.535	0.503	13.68	0.540	0.406	<u>18.05</u>	<u>0.786</u>	0.270	31.78	<u>0.980</u>	0.012
IPC-Dehaze [14] (CVPR2025)	8.55	0.448	0.669	15.32	0.809	0.199	11.43	0.644	0.426	13.55	0.588	<u>0.336</u>	16.02	0.771	<u>0.233</u>	18.99	0.853	0.089
LHTD [36] (CVPR2025)	12.84	0.465	<u>0.580</u>	15.81	0.749	0.221	<u>15.60</u>	0.543	<u>0.371</u>	<u>15.67</u>	0.406	0.355	13.98	0.614	0.288	17.19	0.711	0.198
RTE-FM-Dehazer (Ours)	<u>12.56</u>	0.480	0.528	18.81	<u>0.812</u>	0.145	19.53	0.600	0.233	18.67	0.554	0.222	20.29	0.792	0.159	21.77	0.858	0.068



Fig. 4: Visual comparison on challenging hazy benchmarks. Zoom in for best detail.



Fig. 5: Visual dehazing results under direct, scattered, and reflected lighting. RTE-FM-Dehazer effectively removes haze while preserving diverse lighting structures.

Qualitative Evaluation. Visual comparisons on real hazy images are shown in Figure 4. CNN-based and transformer-based methods (DehazeFormer, DEANet) leave residual haze, especially in dense regions. VQGAN priors (RIDCP, IPC-Dehaze) tend to over-expose and introduce color casts. GAN-based LHTD produces competitive dehazing but suffers from loss of detail and patch-like artifacts. In contrast, RTE-FM-Dehazer removes haze of varying density while preserving texture and color. Even under extreme haze it recovers images that are noise-suppressed and artifact-reduced, and when dust or foreground objects overlap the haze, it correctly removes turbidity without retaining floating particles. Figure 5 demonstrates that the original light paths and global color distribution are preserved, guaranteeing chromatic consistency across varying lighting conditions. Figure 6 shows that our method corrects color shifts caused by low-light or particle pollution, producing visually faithful color. These qualitative

results corroborate the quantitative rankings and verify that RTE-FM-Dehazer consistently generates high-quality, color-consistent, natural haze-free images in diverse real-world scenarios.



Fig. 6: Compared with recent SOTA generative-prior methods, RTE-FM-Dehazer provides color restoration that matches the ground-truth colors more closely.

4.5 Ablation on Regularizer Components

To isolate the necessity of RTE regularization, we train four model variants on P-HAZE: Vanilla FM without RTE regularization ($\lambda = 0$), Absorb-only FM using only the absorption term $-\mu_a \mathbf{z}$, Diffusion-only FM using only the Laplacian term $D\nabla^2 \mathbf{z}$, and DCP-regularized FM replacing RTE with dark-channel prior loss that explicitly enforces the loss gradient to align with dark channel minimization. Rectified Flow [25] learns geometrically straight-line paths in a data-driven manner, while RTE-FM constrains these trajectories with RTE to ensure optical consistency. Table 2 reveals that decoupled regularization catastrophically fails on non-uniform haze due to incomplete atmospheric modeling. Absorb-only suffers from severe performance degradation and blotchy artifacts due to the lack of spatial smoothing, while Diffusion-only leaves significant residual haze and fails to recover clear structures without attenuation guidance. Notably, DCP-regularized FM underperforms even Vanilla FM, indicating that heuristic hard constraints fundamentally misalign with flow matching’s transport cost

Table 2: Ablation study on regularizer components and training data. RTE-FM outperforms decoupled regularizer baselines.

Methods	P-HAZE			I-HAZE			NH-HAZE			D-HAZE			SMOKE			RESIDE-6K		
	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$
Absorb-only FM	18.11	0.636	0.378	11.27	0.498	0.311	12.77	0.459	0.441	9.29	0.330	0.571	11.88	0.400	0.399	16.77	0.793	0.165
Diffusion-only FM	18.91	0.552	0.356	11.19	0.577	0.398	13.85	0.505	0.480	9.11	0.324	0.601	12.21	0.421	0.411	12.03	0.693	0.201
DCP-regularizer-FM	17.22	0.621	0.344	12.21	0.455	0.444	12.22	0.523	0.547	9.99	0.369	0.621	11.19	0.398	0.412	11.79	0.777	0.189
Rectified FM	19.21	0.768	0.254	17.61	0.762	0.193	12.99	0.494	0.491	11.07	0.426	0.681	13.40	0.463	0.372	16.59	0.746	0.208
Vanilla FM	20.12	0.785	0.159	18.27	0.805	0.162	13.78	0.507	0.406	11.43	0.453	0.614	13.77	0.488	0.351	17.51	0.783	0.112
RTE-FM(Ours)	20.29	0.792	0.159	18.81	0.812	0.145	19.53	0.600	0.233	12.56	0.480	0.528	18.67	0.554	0.222	21.77	0.858	0.068

optimization. In contrast, RTE-FM’s integrated absorption-diffusion formulation provides a compatible differentiable constraint that preserves generative flexibility while enforcing radiative fidelity. This superior inductive bias enables robust cross-domain generalization (NH-HAZE, SMOKE) and prevents overfitting to synthetic patterns. Training curves (Fig. 7) confirm RTE loss dominates early training, stabilizing optimization against local optima and ensuring monotonic radiance recovery, unlike Vanilla FM’s erratic early-timestep artifacts.

5 Limitation and Conclusion

RTE-FM-Dehazer combines flow matching’s macroscopic transport with the microscopic Radiative Transfer Equation: the learned velocity field is projected onto RTE-derived energy flux, guiding a single, deterministic trajectory from hazy to clean latent space without stochastic sampling or post-processing. Extensive benchmarks demonstrate effective removal of non-uniform haze, smoke, and color casts while preserving fine textures.

Moving forward, we will pursue methods to address remaining constraints. Operating directly on pixel space will eliminate the information loss inherent in VAE latent encoding, enabling fine-grained detail preservation and improved robustness in ultra-dense haze scenarios where current latent approximations produce artifacts. Furthermore, we will investigate efficient numerical solvers, pro-




Fig. 7: Ablation study results comparing the performance of the proposed FM model with different regularization.

gressive sampling strategies, and model distillation techniques to ensure computational feasibility at full resolution without sacrificing the physical consistency of the dehazing process.

References

1. Agustsson, E., Timofte, R.: Ntire 2017 challenge on single image super-resolution: Dataset and study. In: Proceedings of the IEEE conference on computer vision and pattern recognition workshops. pp. 126–135 (2017)
2. 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)
3. Ancuti, C., Ancuti, C.O., De Vleeschouwer, C.: D-hazy: A dataset to evaluate quantitatively dehazing algorithms. In: 2016 IEEE international conference on image processing (ICIP). pp. 2226–2230. IEEE (2016)
4. Ancuti, C., Ancuti, C.O., Timofte, R., De Vleeschouwer, C.: I-haze: A dehazing benchmark with real hazy and haze-free indoor images. In: International conference on advanced concepts for intelligent vision systems. pp. 620–631. Springer (2018)
5. Bai, S., Chen, K., Liu, X., Wang, J., Ge, W., Song, S., Dang, K., Wang, P., Wang, S., Tang, J., et al.: Qwen2. 5-vl technical report. arXiv preprint arXiv:2502.13923 (2025)
6. Berman, D., Treibitz, T., Avidan, S.: Single image dehazing using haze-lines. IEEE Transactions on Pattern Analysis and Machine Intelligence **42**(3), 720–734 (2020)
7. 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)
8. Chen, C., Mo, J.: IQA-PyTorch: Pytorch toolbox for image quality assessment (2022)
9. Chen, J., Yan, X., Xu, Q., Li, K.: Tokenize image patches: Global context fusion for effective haze removal in large images. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 2258–2268 (2025)
10. Chen, Z., He, Z., Lu, Z.M.: Dea-net: Single image dehazing based on detail-enhanced convolution and content-guided attention. IEEE transactions on image processing **33**, 1002–1015 (2024)
11. Cheng, L., Ba, X., Qu, Y.: Dehazediff: When conditional guidance meets diffusion models for image dehazing. In: 2024 IEEE International Symposium on Circuits and Systems (ISCAS). pp. 1–5. IEEE (2024)
12. Croitoru, F.A., Hondru, V., Ionescu, R.T., Shah, M.: Diffusion models in vision: A survey. IEEE transactions on pattern analysis and machine intelligence **45**(9), 10850–10869 (2023)
13. Fattal, R.: Dehazing using color-lines. ACM transactions on graphics (TOG) **34**(1), 1–14 (2014)
14. Fu, J., Liu, S., Liu, Z., Guo, C.L., Park, H., Wu, R., Wang, G., Li, C.: Iterative predictor-critic code decoding for real-world image dehazing. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 12700–12709 (2025)
15. Golts, A., Freedman, D., Elad, M.: Unsupervised single image dehazing using dark channel prior loss. IEEE transactions on Image Processing **29**, 2692–2701 (2019)

16. Guo, C.L., Yan, Q., Anwar, S., Cong, R., Ren, W., Li, C.: Image dehazing transformer with transmission-aware 3d position embedding. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5812–5820 (2022)
17. 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 (2010)
18. Hore, A., Ziou, D.: Image quality metrics: Psnr vs. ssim. In: 2010 20th international conference on pattern recognition. pp. 2366–2369. IEEE (2010)
19. Jin, Y., Yan, W., Yang, W., Tan, R.T.: Structure representation network and uncertainty feedback learning for dense non-uniform fog removal. In: Asian Conference on Computer Vision. pp. 155–172. Springer (2022)
20. Kim, Y., Son, D.: Noise conditional flow model for learning the super-resolution space. 2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW) pp. 424–432 (2021)
21. Li, B., Peng, X., Wang, Z., Xu, J., Feng, D.: Aod-net: All-in-one dehazing network. In: Proceedings of the IEEE international conference on computer vision. pp. 4770–4778 (2017)
22. 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)
23. 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)
24. Lipman, Y., Chen, R.T., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. arXiv preprint arXiv:2210.02747 (2022)
25. Liu, X., Gong, C., Liu, Q.: Flow straight and fast: Learning to generate and transfer data with rectified flow. arXiv preprint arXiv:2209.03003 (2022)
26. McCartney, E.J.: Optics of the atmosphere: scattering by molecules and particles. New york (1976)
27. Neuhold, G., Ollmann, T., Rota Bulò, S., Kotschieder, P.: The mapillary vistas dataset for semantic understanding of street scenes. In: Proceedings of the IEEE international conference on computer vision. pp. 4990–4999 (2017)
28. Podell, D., English, Z., Lacey, K., Blattmann, A., Dockhorn, T., Müller, J., Penna, J., Rombach, R.: Sdxl: Improving latent diffusion models for high-resolution image synthesis. arXiv preprint arXiv:2307.01952 (2023)
29. Qin, X., Wang, Z., Bai, Y., Xie, X., Jia, H.: Ffa-net: Feature fusion attention network for single image dehazing. In: Proceedings of the AAAI conference on artificial intelligence. vol. 34, pp. 11908–11915 (2020)
30. Ravi, N., Gabeur, V., Hu, Y.T., Hu, R., Ryali, C., Ma, T., Khedr, H., Rädle, R., Rolland, C., Gustafson, L., Mintun, E., Pan, J., Alwala, K.V., Carion, N., Wu, C.Y., Girshick, R., Dollár, P., Feichtenhofer, C.: Sam 2: Segment anything in images and videos. arXiv preprint arXiv:2408.00714 (2024)
31. Razavi, A., Van den Oord, A., Vinyals, O.: Generating diverse high-fidelity images with vq-vae-2. *Advances in neural information processing systems* **32** (2019)
32. Shao, Y., Li, L., Ren, W., Gao, C., Sang, N.: Domain adaptation for image dehazing. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 2808–2817 (2020)
33. Shen, X., Cai, Z., Yin, W., Müller, M., Li, Z., Wang, K., Chen, X., Wang, C.: Gim: Learning generalizable image matcher from internet videos. arXiv preprint arXiv:2402.11095 (2024)

34. Song, Y., He, Z., Qian, H., Du, X.: Vision transformers for single image dehazing. *IEEE Transactions on Image Processing* **32**, 1927–1941 (2023)
35. Team, G., Anil, R., Borgeaud, S., Alayrac, J.B., Yu, J., Soricut, R., Schalkwyk, J., Dai, A.M., Hauth, A., Millican, K., et al.: Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805* (2023)
36. Wang, R., Zheng, Y., Zhang, Z., Li, C., Liu, S., Zhai, G., Liu, X.: Learning hazing to dehazing: Towards realistic haze generation for real-world image dehazing. In: *Proceedings of the Computer Vision and Pattern Recognition Conference*. pp. 23091–23100 (2025)
37. Wang, Y., Yan, X., Wang, F.L., Xie, H., Yang, W., Zhang, X.P., Qin, J., Wei, M.: Ucl-dehaze: Toward real-world image dehazing via unsupervised contrastive learning. *IEEE Transactions on Image Processing* **33**, 1361–1374 (2024)
38. Wu, R.Q., Duan, Z.P., Guo, C.L., Chai, Z., Li, C.: Ridcp: Revitalizing real image dehazing via high-quality codebook priors. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 22282–22291 (2023)
39. Yu, J., Li, X., Koh, J.Y., Zhang, H., Pang, R., Qin, J., Ku, A., Xu, Y., Baldridge, J., Wu, Y.: Vector-quantized image modeling with improved vqgan. *arXiv preprint arXiv:2110.04627* (2021)
40. 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)
41. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 586–595 (2018)
42. Zhang, Y., Wang, F., Yin, D.: Pixel transformer for synthetic-to-real single image dehazing. In: *2023 8th International Conference on Communication, Image and Signal Processing (CCISP)*. pp. 250–254. IEEE (2023)
43. Zhao, C.: *Numerical Methods for Solving the Radiative Transfer Equation*, pp. 54–79. Cambridge University Press (2024)
44. Zhao, S., Zhang, L., Shen, Y., Zhou, Y.: Refinednet: A weakly supervised refinement framework for single image dehazing. *IEEE Transactions on Image Processing* **30**, 3391–3404 (2021)
45. Zheng, Y., Zhan, J., He, S., Dong, J., Du, Y.: Curricular contrastive regularization for physics-aware single image dehazing. In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2023)
46. Zhu, Q., Mai, J., Shao, L.: Single image dehazing using color attenuation prior. In: *BMVC*. vol. 4, pp. 1674–1682 (2014)