R2AS: RELATIVE-TO-ABSOLUTE SUPERVISION FOR REAL-WORLD DEHAZING

Anonymous authors
Paper under double-blind review

ABSTRACT

Adapting image dehazing models to unlabeled real images requires supervision beyond physical reconstruction: different combinations of transmission and clean radiance can explain the same hazy image. Under a model with shared atmospheric light, controlled haze interventions prescribe exact relative optical depth shifts between views. This separates the inverse problem into a known relative structure and one unknown absolute offset per pixel. We propose Relative-to-Absolute Supervision (R2AS) for real-world dehazing. First, Relative-State Calibration (RSC) constructs views with known transmission scaling and trains a teacher to match their prescribed optical depth differences while enforcing reconstruction compatibility with the observations. Second, Absolute-State Anchoring (ASA) projects the teacher's aligned transmission and radiance predictions onto a shared, observation-compatible family, producing joint targets through a closed-form scalar solution at each pixel. The calibrated teacher remains frozen while these targets supervise a student that retains the backbone architecture and requires a single dehazer evaluation at inference. We establish a bound on relative optical depth error and prove that anchoring does not increase weighted joint target error when the true state lies in the feasible family. Experiments on six real-world benchmarks demonstrate that R2AS achieves the highest MUSIQ on all four unpaired datasets and the highest PSNR on both paired datasets among the compared methods.

1 INTRODUCTION

Learning to restore images from unlabeled observations is central to bridging the persistent domain gap between synthetic training data and complex real-world degradations (Feng et al., 2024; Shin et al., 2025). In image dehazing, atmospheric scattering severely attenuates visibility and alters contrast, making clean radiance difficult to recover from a single observation (He et al., 2011; Li et al., 2019). While supervised learning on synthetic pairs has driven significant progress, the diversity of natural haze limits real-world generalization. Consequently, recent approaches increasingly turn to unlabeled real images to construct target domain supervision (Liu et al., 2022; Fang et al., 2024; Ma et al., 2025). The challenge is to separate supervision determined by the observation from state estimates that require prior evidence.

Existing approaches broadly follow three routes to address this challenge: adaptation, restoration priors, and physics-based supervision. Adaptation transfers source knowledge through test-time learning or prompt-guided updates (Liu et al., 2022; Zhang et al., 2026a;b); restoration priors supply clean image evidence through unpaired images or pseudo labels (Wang et al., 2024; Fang et al., 2024; Wu et al., 2023); and physics-based methods construct signals through haze transfer or nonhomogeneous generation (Tsai et al., 2025; Shin et al., 2025). Although these routes supply useful constraints, they do not explicitly distinguish which components of the restoration are physically identifiable from the observation alone and which require prior evidence.

This limitation stems from the inherent physical ambiguity of the inverse problem: *a single hazy observation can correspond to multiple physically plausible combinations of transmission and scene radiance*. Under the atmospheric scattering model $I = A + e^{-q}(J - A)$, where I is the observed hazy image, J is the scene radiance, A is the atmospheric light, and q is the optical depth, the variables q and J can covary while leaving I unchanged. Specifically, under this model, adding a feasible shift to the optical depth and correspondingly scaling the radiance difference $J - A$ exactly preserves the

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

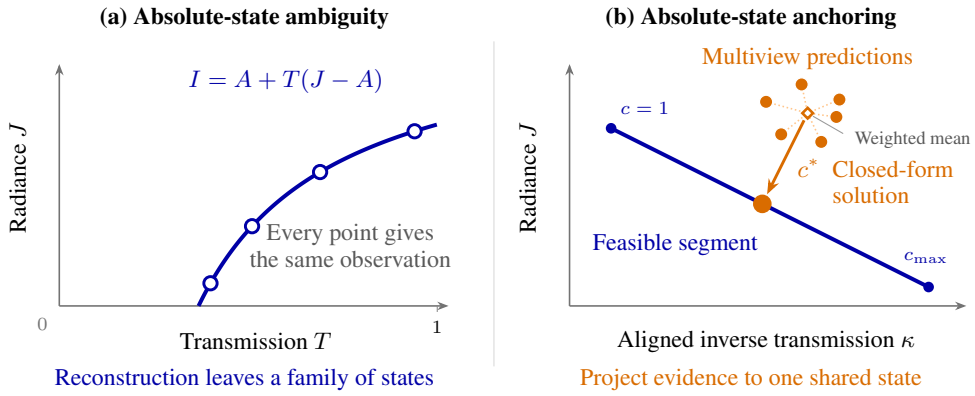


Figure 1: **From reconstruction ambiguity to geometric anchoring.** (a) Multiple transmission–radiance pairs explain the same observation at fixed atmospheric light. (b) ASA projects the aligned teacher mean onto the feasible segment to select c^*.

observation. Therefore, exact physical reconstruction merely constrains the prediction to a family of observation-compatible states, without specifying which member represents the true scene. Enforcing relative consistency across constructed views still leaves a common optical depth offset unrestricted at each pixel. The core challenge is how to exploit the deterministic information that target observations do provide, while assigning prior evidence a precise role in resolving the remaining ambiguity.

To address this challenge, we propose **Relative-to-Absolute Supervision (R2AS)** based on HazeFlow, which separates relative calibration from prior-guided absolute state selection. The key insight is that a controlled haze intervention prescribes exact relative optical depth shifts between views, even when their absolute values remain unknown. Under shared atmospheric light, multiplying the transmission by a chosen factor ρ_i induces a known shift $s_i = -\log \rho_i$, separating the optical depths across views into a known relative structure $q_i = a + s_i$ and a single unknown scalar offset a per pixel. R2AS decomposes the ill-posed inverse problem into two tractable tasks: *Relative-State Calibration (RSC)* learns these prescribed optical depth differences and enforces reconstruction compatibility across views, while *Absolute-State Anchoring (ASA)* projects calibrated teacher evidence onto a one-dimensional convex feasible segment, selecting a shared absolute state in closed form. Figure 1 illustrates this division of supervision. Our main contributions are summarized as follows:

- We formalize the absolute-state ambiguity in unsupervised dehazing and propose R2AS, a framework that separates identifiable relative calibration from prior-guided absolute anchoring.
- We introduce RSC and ASA with a graph-based bound for relative optical depth error (Prop. 1) and a closed-form projection that does not increase joint target error when the true state is feasible (Prop. 2).
- Experiments on RTTS, Fattal, NHRW, RealFOGGY, Dense-Haze, and O-Haze demonstrate that R2AS achieves the highest MUSIQ on the four unpaired benchmarks and the highest PSNR on the two paired benchmarks among the compared methods. Ablations and sensitivity analyses characterize the contributions of the supervision and its dependence on atmospheric light estimation.

2 Related Work

Adaptation to Real Haze Domain adaptation transfers restoration knowledge from synthetic pairs to unlabeled real images (Shao et al., 2020; Chen et al., 2021). Liu et al. (2022) use a helper network for test-time supervision, while CoA (Ma et al., 2025) combines compression and bilevel adaptation. PDDA (Zhang et al., 2026a) learns haze prompts, and BiLaLoRA (Zhang et al., 2026b) combines text guidance with bilevel selection of LoRA layers. Self-supervised variants instead exploit consistency between a hazy input and its reconstruction (Liang et al., 2022; Cheng et al., 2022). These methods accept whichever compatible state the adapted model reaches; R2AS uses prescribed optical depth shifts to calibrate the model before selecting absolute states from its predictions.

Unpaired Supervision and Restoration Priors UCL-Dehaze (Wang et al., 2024) uses contrastive and adversarial objectives on unpaired images; DiffUIR (Zheng et al., 2024) and DehazeSB (Lan

et al., 2025) use diffusion and Schrödinger bridge models, respectively, to restore images toward the clean distribution. RIDCP (Wu et al., 2023) supplies codebook priors, IPC-Dehaze (Fu et al., 2025) iteratively refines code predictions, and CORUN (Fang et al., 2024) combines physical unfolding with coherence-based pseudo label selection. Language-guided supervision (Wang et al., 2025a) and generated hazy data (Wang et al., 2025b) extend this line by supplying textual or synthesized evidence. Prior evidence in these works enters as a data or feature constraint, so the state it selects still lies somewhere in the reconstruction-compatible family. ASA instead constrains that selection explicitly, coupling attenuation and RGB through a scalar inverse transmission.

Physical Modeling and Haze Generation Physical methods use prior-based inversion (He et al., 2011), decomposition consistency (Liu et al., 2021), and density-based self-augmentation (Yang et al., 2022). KA_Net (Feng et al., 2024) models multiple degradation factors, PHATNet (Tsai et al., 2025) transfers target haze onto source clean images, while PRISM (Fang et al., 2026) unifies physical realization with restoration in a single model and RTE-FM-Dehazer (Wei et al., 2026) derives a flow matching solver from the radiative transfer equation. Recent flow matching solvers for inverse problems (Pourya et al., 2026; Tauberschmidt et al., 2026) inject physical constraints into the generative trajectory, and HazeFlow (Shin et al., 2025) reformulates scattering as an ODE with nonhomogeneous haze generation. These approaches shape the reconstruction trajectory but leave the absolute state underdetermined; we use HazeFlow's transmission and radiance outputs to calibrate prescribed optical depth differences and construct joint restoration targets.

3 Problem Formulation

We consider adaptation from paired source data $\mathcal{D}_s = \{(I_n^s, J_n^s)\}_{n=1}^{N_s}$ to unlabeled real hazy images $\mathcal{D}_r = \{I_m^r\}_{m=1}^{N_r}$. Starting from a source-pretrained dehazer, our goal is to construct target domain supervision without clean references.

3.1 Ambiguity of Absolute Supervision

For one pixel, let $I, J, A \in \mathbb{R}^3$ denote the observed RGB vector, clean radiance, and atmospheric light, respectively. The atmospheric scattering model relates them through transmission $T \in (0, 1]$, or equivalently optical depth $q = -\log T$:

$$I = A + T(J - A) = A + e^{-q}(J - A). \tag{1}$$

Our physical analysis is pixelwise and assumes linear radiance, shared atmospheric light, and no noise or clipping.

Under these assumptions and with known A, reconstructing I does not uniquely determine the restoration state $z = (q, J)$. For any feasible scalar shift β,

$$q' = q + \beta, \qquad T' = e^{-\beta}T, \qquad J' = A + e^{\beta}(J - A) \tag{2}$$

preserves $A + T'(J' - A) = I$. Reconstruction therefore defines a family of compatible states parameterized by β.

3.2 Relative Supervision from Controlled Haze Changes

We prescribe optical depth changes by constructing additional views. Given a real observation I_0, we estimate atmospheric light $\widehat{A}$ and construct views using chosen factors $\rho_i \in (0, 1]$, with $\rho_0 = 1$:

$$I_i = \widehat{A} + \rho_i(I_0 - \widehat{A}), \qquad s_i = -\log \rho_i. \tag{3}$$

When $\widehat{A} = A$, these views share the same clean radiance and satisfy $T_i = \rho_i T_0$. Their optical depth differences are thus known: $q_i - q_j = s_i - s_j$. Writing $s = (s_i)_i$ and $q = (q_i)_i$, we obtain

$$q = a\mathbf{1} + s, \qquad a = q_0. \tag{4}$$

The intervention determines s, leaving one shared offset a per pixel.

For a fixed $\widehat{A}$, a common restoration of the constructed views must satisfy

$$q_i = a + s_i, \qquad J = \widehat{A} + e^{a}(I_0 - \widehat{A}). \tag{5}$$

R2AS learns the prescribed differences, then uses teacher evidence to select the remaining scalar within this family. If $\widehat{A} \neq A$, the family may exclude the true restoration.

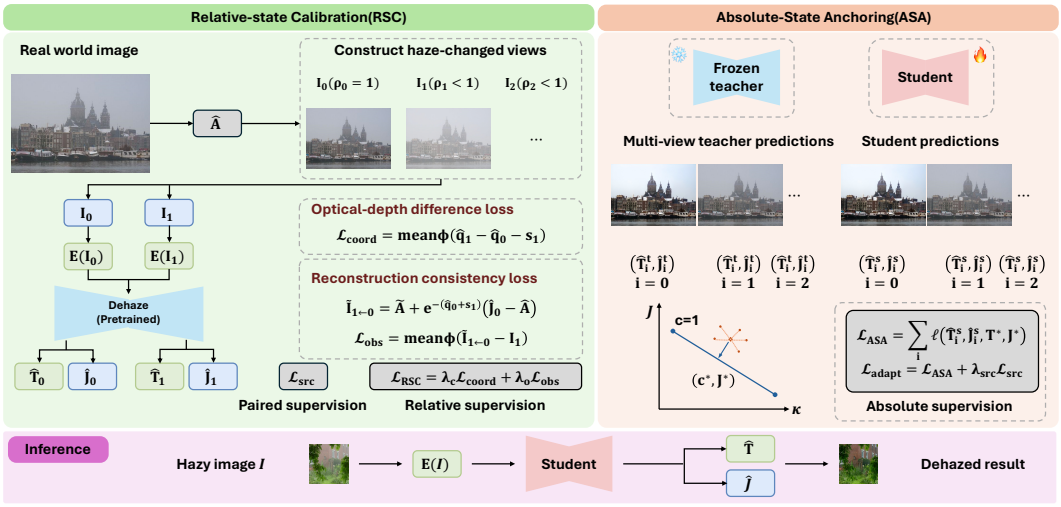


Figure 2: **Training and inference of** R2AS.

4 Relative-to-Absolute Supervision

4.1 Overview

Figure 2 shows the full workflow: source pretraining followed by RSC and ASA with student distillation. We start from the official pretrained HazeFlow checkpoint and retain paired source supervision during both adaptation stages.

HazeFlow provides transmission $\widehat{T}_i = \tau_\theta(I_i)$ and clean radiance $\widehat{J}_i = F_\theta(I_i)$. RSC uses $\widehat{q}_i = -\log \widehat{T}_i$ to match optical depth differences and penalizes reconstruction residuals. After calibration, we freeze the model as a teacher and initialize the student from the same parameters. ASA projects the teacher's multiview predictions to construct joint targets, which supervise the student on individual views. The adaptation schedule is detailed in Section 5.1, and source pretraining is described in Appendix B.1.

4.2 Relative-State Calibration

Learning prescribed differences. RSC uses the original image I_0 and one constructed view I_1 from Equation (3). The selected haze factor specifies $q_1 - q_0 = s_1$, providing a training target without estimating either absolute optical depth. We match this difference with

$$\mathcal{L}_{\text{coord}} = \text{mean}\,\text{Charb}\big(\widehat{q}_1 - \widehat{q}_0 - s_1\big), \tag{6}$$

where $\text{Charb}(x) = \sqrt{x^2 + \epsilon^2}$ and the mean averages pixels.

Enforcing reconstruction compatibility. Matching optical depth differences alone does not constrain the predicted clean radiance. We therefore apply the same haze change to the prediction for I_0 and require it to reproduce I_1:

$$\widetilde{I}_{1\leftarrow 0} = \widehat{A} + e^{-(\widehat{q}_0 + s_1)}(\widehat{J}_0 - \widehat{A}), \tag{7}$$

$$\mathcal{L}_{\text{obs}} = \text{mean}\,\text{Charb}(\widetilde{I}_{1\leftarrow 0} - I_1). \tag{8}$$

This is a scaled physical reconstruction residual that couples transmission and radiance while leaving the absolute ambiguity unresolved (Appendix A.1).

Calibrating the teacher. We combine the two real-image losses and retain HazeFlow's original supervised source objective, denoted by $\mathcal{L}_{\text{src}}$ (Shin et al., 2025):

$$\mathcal{L}_{\text{RSC}} = \lambda_c \mathcal{L}_{\text{coord}} + \lambda_o \mathcal{L}_{\text{obs}}, \tag{9}$$

$$\mathcal{L}_{\text{adapt}} = \mathcal{L}_{\text{RSC}} + \lambda_{\text{src}} \mathcal{L}_{\text{src}}. \tag{10}$$

After calibration, we freeze the model as the ASA teacher. The following result characterizes relative supervision; Appendix A.1 gives the graph formulation and proof.

Proposition 1 (Relative identifiability)**.** *Under the exact intervention model, suppose the supervised view pairs form a connected graph. If all predicted differences match their prescribed values, then $\widehat{q}_i = q_i + \beta$ for a common scalar β at each pixel. More generally, the squared variation of optical depth prediction errors across views is bounded by a graph-dependent multiple of the squared difference residuals, as specified in Equation* (22)*.*

RSC controls relative optical depth error while leaving the common offset for ASA to select.

4.3 Absolute-State Anchoring

ASA combines the calibrated teacher's predictions for K views, including I_0, into one compatible state. It selects transmission and radiance jointly through the shared scalar in Equation (5).

A scalar parameterization. Let $r = I_0 - \widehat{A}$ and write the base inverse transmission as $c = e^a = 1/T_0$. Every common restoration in the estimated family then has the form

$$T_i(c) = \rho_i / c, \qquad J(c) = \widehat{A} + cr, \qquad 1 \le c \le c_{\max}. \tag{11}$$

The interval enforces nonnegative optical depth, valid RGB values, and a numerical transmission floor. Its construction is given in Appendix A.2, Equation (25). In aligned inverse transmission and RGB coordinates, the family is affine in c, allowing a scalar quadratic objective.

Aligning teacher predictions. Because the views exhibit different attenuation levels, their transmission predictions must first be expressed on a common scale. We define

$$\kappa_i^t = \frac{\rho_i}{\widehat{T}_i} = \rho_i e^{\widehat{q}_i}. \tag{12}$$

For exact predictions, all κ_i^t equal the base inverse transmission c, so $(\kappa_i^t, \widehat{J}_i)$ are aligned estimates of the same state.

Joint target selection. We choose the feasible restoration closest to these estimates:

$$c^* = \operatorname*{arg\,min}_{c \in [1, c_{\max}]} \sum_i \pi_i \left[w_\kappa (\kappa_i^t - c)^2 + w_J \|\widehat{J}_i - \widehat{A} - cr\|_2^2 \right]. \tag{13}$$

The positive view weights π_i sum to one. Define the detached teacher means $\bar{\kappa} = \sum_i \pi_i \kappa_i^t$ and $\bar{J} = \sum_i \pi_i \widehat{J}_i$. We use $w_J = 1/3$ to average RGB channels and $w_\kappa = 1/\max(1, \bar{\kappa})^2$ to normalize inverse transmission scale. These weights are fixed within each pixel's projection. The scalar quadratic has the closed-form solution

$$c^* = \mathrm{clip}_{[1, c_{\max}]} \frac{w_\kappa \bar{\kappa} + w_J r^\top (\bar{J} - \widehat{A})}{w_\kappa + w_J \|r\|_2^2}. \tag{14}$$

The resulting joint targets are

$$\kappa_i^* = c^*, \qquad J^* = \widehat{A} + c^* r, \qquad \tau_i^* = \rho_i / c^*. \tag{15}$$

All views share J^*, recover the prescribed transmission ratios, and reconstruct their observations exactly under $\widehat{A}$.

Projection guarantee. In the stacked coordinates $\zeta = (\kappa_i, J_i)_i$, the family in Equation (11) defines a convex segment $\mathcal{C}$. Equation (14) projects the teacher state $\zeta_t = (\kappa_i^t, \widehat{J}_i)_i$ onto this segment in the metric W induced by Equation (13). Let $p = (c^*, J^*)_i$ be the projected state. The formal set and metric are specified in Appendix A.2.

Proposition 2 (Error reduction by absolute-state anchoring)**.** *Suppose $\mathcal{C}$ is nonempty and contains the true joint state $\zeta_{\mathrm{gt}} = (1/T_0, J)_i$. Then*

$$\|p - \zeta_{\mathrm{gt}}\|_W^2 \le \|\zeta_t - \zeta_{\mathrm{gt}}\|_W^2 - \|\zeta_t - p\|_W^2. \tag{16}$$

If the true state is feasible, projection strictly reduces weighted joint target error whenever it changes the teacher state. Appendix A.2 gives the proof.

4.4 Student Training and Inference

The student fits targets constructed from the frozen teacher. To control anchoring strength, we interpolate in aligned inverse transmission and RGB coordinates:

$$\kappa_i^\eta = (1-\eta)\kappa_i^t + \eta c^*, \qquad J_i^\eta = (1-\eta)\widehat{J}_i + \eta J^*, \tag{17}$$

where $\eta \in [0,1]$ controls anchoring strength: $\eta = 1$ uses the projected state, while smaller values retain teacher evidence and may lie outside the feasible family. In stacked coordinates, $\zeta_\eta = (1-\eta)\zeta_t + \eta p$.

For a view sampled as $i \sim \pi$, the student predicts $\kappa_i^\phi = \rho_i/\tau_\phi(I_i)$ and $J_i^\phi = F_\phi(I_i)$. Its training objective is

$$\mathcal{L}_{\text{joint}} = \mathbb{E}_{i\sim\pi}\,\text{mean}\left[w_\kappa(\kappa_i^\phi - \kappa_i^\eta)^2 + w_J\|J_i^\phi - J_i^\eta\|_2^2\right], \tag{18}$$

$$\mathcal{L}_{\text{distill}} = \lambda_{\text{KD}}\mathcal{L}_{\text{joint}} + \lambda_{\text{src}}\mathcal{L}_{\text{src}}. \tag{19}$$

Only the student is updated; teacher predictions and targets are detached. Inference uses the student with a single dehazer evaluation. Algorithms 1 and 2 summarize training.

5 Experiments

5.1 Experimental Setup

Datasets. Source training data come exclusively from RIDCP (Wu et al., 2023), while URHI (Li et al., 2019) supplies unlabeled real images for RSC and anchored distillation. The same trained model is evaluated on all six datasets without dataset-specific adaptation. RTTS (Li et al., 2019) is used only for testing; it, Fattal (Fattal, 2008), nighttime NHRW (Zhang et al., 2020), and RealFOGGY use no-reference metrics. Dense-Haze (Ancuti et al., 2019) and O-Haze (Ancuti et al., 2018) use their official test sets for full-reference evaluation, and their clean reference images are not used in training.

Comparison methods. We compare R2AS with eleven methods: DEA-Net (Chen et al., 2024), DiffUIR (Zheng et al., 2024), DehazeSB (Lan et al., 2025), CoA (Ma et al., 2025), BiLaLoRA (Zhang et al., 2026b), CORUN (Fang et al., 2024), HazeFlow (Shin et al., 2025), IPC-Dehaze (Fu et al., 2025), KA_Net (Feng et al., 2024), PDDA (Zhang et al., 2026a), and RIDCP (Wu et al., 2023). We evaluate all outputs on the same test images with the same metric implementations.

Implementation details. All module comparisons start from the same official pretrained checkpoint, retaining the backbone, velocity output, transmission refinement branch, and input conditioning. After RSC, we freeze a copy as the teacher and initialize the student from the same parameters. Both adaptation stages include paired source replay. We run adaptation on two NVIDIA A800 GPUs. Following HazeFlow's adaptation schedule (Shin et al., 2025), each adaptation stage uses Adam for 10K iterations, with learning rates 10^{-6} for RSC and 5×10^{-5} for distillation. Calibration uses $\rho \sim \mathcal{U}[0.65, 0.95]$, $\lambda_c = 0.1$, $\lambda_o = 1$, and Charbonnier $\epsilon = 10^{-3}$. Real and source loaders contribute at a 1:1 original-image ratio, with each loss averaged separately. The main distillation configuration uses $\rho = (1, 0.9, 0.75)$, $\pi = (0.5, 0.25, 0.25)$, $\eta = 1.0$, $w_J = 1/3$, $w_\kappa = 1/\max(1,\bar{\kappa})^2$, and $\lambda_{\text{KD}} = \lambda_{\text{src}} = 1$. The transmission floor is $\tau_{\min} = 10^{-3}$. Real and source data use independent shuffled iterators. Source replay reuses the MCBM-generated source pairs from pretraining and the original supervised loss, including its all-one atmospheric light label convention. Source pretraining, gradient accumulation, and target construction are detailed in Appendix B.

Evaluation metrics. For RTTS, Fattal, NHRW, and RealFOGGY, we report four higher-is-better no-reference quality scores: MUSIQ (Ke et al., 2021), MANIQA, TOPIQ, and Q-Align. We additionally report FADE (Choi et al., 2015), for which a lower value indicates less residual fog. Dense-Haze and O-Haze are evaluated with full-reference PSNR, SSIM (Wang et al., 2004), and LPIPS (Zhang et al., 2018); higher PSNR and SSIM and lower LPIPS indicate better reconstruction.

Table 1: Quantitative results on unpaired datasets RTTS Li et al. (2019) and Fattal Fattal (2008). Best and second-best results are shown in bold and underlined, respectively.

	RTTS					Fattal				
Method	MUSIQ↑	MANIQA↑	TOPIQ↑	FADE↓	Q-Align↑	MUSIQ↑	MANIQA↑	TOPIQ↑	FADE↓	Q-Align↑
DEA-Net	54.287	0.304	0.403	1.886	3.066	63.807	0.383	0.573	0.508	3.988
DiffUIR	54.442	0.315	0.407	2.080	3.111	61.547	0.366	0.533	0.783	3.819
DehazeSB	59.635	0.312	0.461	0.846	3.047	68.546	0.412	0.611	0.439	4.041
CoA	55.588	0.263	0.405	0.841	3.140	64.146	0.322	0.510	0.294	3.902
BiLaLoRA	63.545	0.331	0.511	0.728	3.533	69.335	0.393	0.623	0.330	4.228
CORUN	63.281	<u>0.354</u>	0.538	0.810	<u>3.556</u>	<u>70.327</u>	<u>0.446</u>	**0.681**	0.342	<u>4.305</u>
HazeFlow	<u>63.869</u>	0.345	**0.592**	<u>0.581</u>	3.373	68.320	0.410	0.628	<u>0.264</u>	4.155
IPC-Dehaze	59.604	0.327	0.500	1.094	3.469	65.672	0.420	0.617	0.435	4.183
KA_Net	50.770	0.269	0.370	3.445	2.583	60.313	0.332	0.480	1.766	2.908
PDDA	59.907	0.316	0.462	1.161	3.457	69.146	0.372	0.570	0.431	4.152
RIDCP	59.357	0.275	0.453	0.917	3.395	66.538	0.337	0.571	0.410	4.150
R2AS	**68.312**	**0.362**	<u>0.586</u>	**0.543**	**3.755**	**73.208**	**0.490**	<u>0.635</u>	**0.242**	**4.752**

Table 2: Quantitative results on paired real-world datasets: Dense-Haze Ancuti et al. (2019) and O-Haze Ancuti et al. (2018). Best and second-best results are shown in bold and underlined, respectively.

	Dense-Haze			O-Haze		
Method	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
DEA-Net	9.97	0.419	0.746	15.44	0.607	0.602
DiffUIR	9.91	0.405	0.789	15.04	0.552	0.696
DehazeSB	11.14	0.390	0.668	15.74	<u>0.643</u>	0.575
CoA	9.59	0.452	0.717	14.38	0.617	0.603
BiLaLoRA	<u>11.54</u>	<u>0.473</u>	<u>0.649</u>	15.78	0.636	0.535
CORUN	10.28	0.452	0.671	15.08	**0.647**	<u>0.517</u>
HazeFlow	11.52	0.412	0.679	<u>15.85</u>	0.499	0.612
IPC-Dehaze	8.47	0.420	0.713	14.02	0.639	0.538
KA_Net	7.40	0.376	0.801	7.84	0.434	0.755
PDDA	9.33	0.451	0.697	14.21	0.623	0.566
RIDCP	9.98	0.444	0.693	<u>15.85</u>	0.613	0.598
R2AS	**12.08**	**0.493**	**0.632**	**16.51**	0.639	**0.513**

5.2 Main Results

Unpaired real-world benchmarks. Table 1 reports RTTS and Fattal results. On RTTS, R2AS increases MUSIQ from HazeFlow 63.869 to 68.312 (6.96%); on Fattal, it increases MUSIQ from CORUN's 70.327 to 73.208 (4.10%). It also obtains the highest MANIQA and Q-Align on both datasets. Relative to HazeFlow, which has the lowest competing FADE on both datasets, FADE decreases from 0.581 to 0.543 on RTTS and from 0.264 to 0.242 on Fattal. These gains are accompanied by lower TOPIQ scores: 0.586 versus 0.592 for HazeFlow on RTTS, and 0.635 versus 0.681 for CORUN on Fattal. NHRW and RealFOGGY results appear in Appendix C, Table 7.

Paired real-world benchmarks. On Dense-Haze, R2AS obtains 12.08 dB PSNR, compared with BiLaLoRA's 11.54 dB (+0.54 dB). On O-Haze, PSNR increases from 15.85 dB for HazeFlow and RIDCP to 16.51 dB (+0.66 dB). Its O-Haze LPIPS is 0.513 versus CORUN's 0.517, while SSIM is 0.639 versus 0.647 (Table 2).

5.3 Ablation Study

Table 3 compares the baseline, RSC alone, ASA alone, and their combination on the RTTS test set. Removing ASA lowers MUSIQ by 0.724 points (68.312 to 67.588); removing RSC lowers it by 0.465 points (to 67.847). Combining the modules also improves MANIQA, FADE, and Q-Align over either module alone, while TOPIQ decreases from 0.596 with RSC alone to 0.586.

Relative supervision and physical regularization. Table 4(a) separates the optical depth difference loss $\mathcal{L}_{\text{coord}}$ from the reconstruction loss $\mathcal{L}_{\text{obs}}$. All variants use the same source initialization, source

Table 3: Ablation study on RTTS. Best and second-best scores are shown in bold and underlined, respectively.

Method	MUSIQ↑	MANIQA↑	TOPIQ↑	FADE↓	Q-Align↑
R2AS (RSC + ASA)	**68.312**	**0.362**	0.586	**0.543**	**3.755**
w/o ASA (RSC only)	67.588	<u>0.357</u>	**0.596**	0.557	3.621
w/o RSC (ASA only)	<u>67.847</u>	0.354	0.589	<u>0.551</u>	<u>3.638</u>
Baseline	63.873	0.348	<u>0.590</u>	0.578	3.376



Figure 3: Qualitative dehazing comparison on the RTTS dataset.

replay, view pairs, and optimizer update budget; scores are measured after calibration. Each loss alone improves both diagnostic errors and MUSIQ over the baseline. Combining them yields the lowest normalized optical depth difference error (0.09) and observable error (0.0095), while increasing MUSIQ to 67.588, compared with 66.562 for optical depth difference supervision alone and 66.934 for reconstruction alone.

Table 4: Relative-state calibration on RTTS: loss ablations and held-out intervention response.

(a) RSC loss ablations ($\rho = 0.75$).

Real-image objective	$E^{\mathrm{rel}}_{\Delta q}\downarrow$	$E_{\mathrm{obs}}\downarrow$	MUSIQ↑
Baseline	0.17	0.0156	63.873
$\mathcal{L}_{\mathrm{coord}}$	0.12	0.0128	66.562
$\mathcal{L}_{\mathrm{obs}}$	0.11	0.0114	66.934
$\mathcal{L}_{\mathrm{coord}} + \mathcal{L}_{\mathrm{obs}}$ (RSC)	0.09	0.0095	67.588

(b) Held-out intervention response.

	$\rho = 0.60$		$\rho = 0.75$	
Method	$E^{\mathrm{rel}}_{\Delta q}\downarrow$	$E_{\mathrm{obs}}\downarrow$	$E^{\mathrm{rel}}_{\Delta q}\downarrow$	$E_{\mathrm{obs}}\downarrow$
Baseline	0.10	0.0132	0.17	0.0156
+ RSC	0.07	0.0082	0.09	0.0095

5.4 Analysis and Visualization

Held-out intervention response. On RTTS test images, we measure $E_{\mathrm{obs}} = \mathcal{L}_{\mathrm{obs}}$ and normalize the optical depth difference error by the prescribed shift $-\log \rho$:

$$E^{\mathrm{rel}}_{\Delta q}(\rho) = \frac{\mathrm{mean}\,|[\widehat{q}(I_\rho) - \widehat{q}(I_0)] - [-\log \rho]|}{-\log \rho}. \tag{20}$$

RSC reduces normalized optical depth difference error by 30.0% at $\rho = 0.60$ and 47.1% at $\rho = 0.75$, and lowers both observable errors (Table 4(b)). The $\rho = 0.60$ intervention lies outside the training interval $[0.65, 0.95]$.

Qualitative comparison. Figure 3 compares two RTTS scenes, showing differences in residual haze, visibility, and color rendition. Additional Fattal, NHRW, and RealFOGGY examples are provided in Appendix C, Figures 4 and 5.

5.5 Sensitivity to Atmospheric Light Estimation

Both modules use the estimated atmospheric light to define the constructed views and feasible family. We perturb the DCP estimate as $\widehat{A}' = \mathrm{clip}_{[0,1]}(\widehat{A} + \delta)$ with $\delta \in \{0, 0.01, 0.03, 0.05\}$, applying the same $\widehat{A}'$ to view construction and projection. We keep RSC and the calibrated teacher fixed and repeat ASA student training for each perturbation and anchoring strength $\eta \in \{0.5, 0.75, 0.9, 1.0\}$. The $\delta = 0$ row uses the original DCP estimate.

Table 5(a) reports restoration quality with full anchoring ($\eta = 1.0$). At $\delta = 0.01$, MUSIQ changes from 68.312 to 68.316, while MANIQA and Q-Align decrease and FADE increases. At $\delta = 0.05$, MUSIQ falls to 67.780, a drop of 0.532 points (0.78%); MANIQA decreases from 0.362 to 0.331, Q-Align from 3.755 to 3.482, and FADE increases from 0.543 to 0.587. MUSIQ remains relatively stable across the tested offsets, whereas MANIQA, Q-Align, and FADE exhibit larger relative changes under atmospheric light perturbations. Table 5(b) shows how anchoring strength interacts with the perturbation. Full anchoring has the highest MUSIQ for $\delta \leq 0.03$. At $\delta = 0.05$, partial anchoring with $\eta = 0.75$ obtains 67.932, exceeding full anchoring by 0.152 points.

Table 5: Sensitivity to atmospheric light estimation on RTTS. δ is the additive perturbation to the DCP estimate. Panel (a) reports restoration metrics at $\eta = 1.0$, and panel (b) reports MUSIQ for each (δ, η) pair.

(a) Restoration quality.

	Restoration quality			
Airlight offset	MUSIQ↑	MANIQA↑	FADE↓	Q-Align↑
$\delta = 0$	68.312	0.362	0.543	3.755
$\delta = 0.01$	68.316	0.357	0.549	3.708
$\delta = 0.03$	67.984	0.352	0.565	3.653
$\delta = 0.05$	67.780	0.331	0.587	3.482

(b) Anchoring strength.

	Anchoring strength η			
Airlight offset	0.5	0.75	0.9	1.0
$\delta = 0$	68.275	68.287	68.306	68.312
$\delta = 0.01$	68.231	68.277	68.301	68.316
$\delta = 0.03$	67.893	67.946	67.975	67.984
$\delta = 0.05$	67.845	67.932	67.773	67.780

5.6 Contribution of Multiple Views

We hold the teacher and $\eta = 1.0$ fixed and vary the ASA view set as listed in Table 6. All four metrics improve as K increases from 1 to 3, with MUSIQ rising from 67.764 to 68.312. The gain is not confined to quality metrics: FADE falls from 0.554 to 0.543 and Q-Align rises from 3.685 to 3.755, indicating that extra views add consistent evidence rather than conflicting residuals.

Table 6: Ablation of multiview evidence on RTTS with a fixed calibrated teacher and $\eta = 1.0$.

Views	MUSIQ↑	MANIQA↑	FADE↓	Q-Align↑
$K = 1$ (original only)	67.764	0.358	0.554	3.685
$K = 2$ ($\rho = 1, 0.75$)	67.928	0.360	0.549	3.708
$K = 3$ ($\rho = 1, 0.9, 0.75$)	68.312	0.362	0.543	3.755

6 Conclusion

R2AS uses controlled haze interventions to supervise optical depth differences and projects calibrated teacher predictions to select joint restoration targets. Student distillation retains one dehazer evaluation at inference. The method obtains the highest MUSIQ on four unpaired benchmarks and the highest PSNR on two paired benchmarks among the compared methods. Ablations show gains from both calibration losses and multiple views; partial anchoring improves MUSIQ at the largest tested atmospheric light offset.

Limitations. The implementation operates in sRGB, which departs from the assumption of linear radiance of the scattering model. Multiple scattering, saturation, and spatially varying illumination introduce further model errors. Atmospheric light errors can shift the feasible family away from the true state, affecting the anchored targets. The method requires an explicit transmission output and incurs additional training cost from storing the frozen teacher and evaluating multiple views.

AI USE STATEMENT

We used generative AI tools solely for language polishing to improve the grammar, clarity, and readability of the manuscript. No other content, including research ideas, methodology, theoretical analysis, code, data, figures, or experimental results, was generated using these tools. We take full responsibility for the accuracy and integrity of the final manuscript.

REFERENCES

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 *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 754–762, 2018.

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 *IEEE International Conference on Image Processing*, pages 1014–1018, 2019.

Zeyuan Chen, Yangchao Wang, Yang Yang, and Dong Liu. Psd: Principled synthetic-to-real dehazing guided by physical priors. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 7176–7185, 2021.

Zixuan Chen, Zhiwei He, and Zhe-Ming Lu. DEA-Net: Single image dehazing based on detail-enhanced convolution and content-guided attention. *IEEE Transactions on Image Processing*, 33: 1002–1015, 2024.

De Cheng, Yan Li, Dingwen Zhang, Nannan Wang, Xinbo Gao, and Jiande Sun. Robust single image dehazing based on consistent and contrast-assisted reconstruction. In *International Joint Conference on Artificial Intelligence*, pages 848–854, 2022.

Lark Kwon Choi, Jaehee You, and Alan Conrad Bovik. Referenceless prediction of perceptual fog density and perceptual image defogging. *IEEE Transactions on Image Processing*, 24(11): 3888–3901, 2015.

Chengyu Fang, Chunming He, Fengyang Xiao, Yulun Zhang, Longxiang Tang, Yuelin Zhang, Kai Li, and Xiu Li. Real-world image dehazing with coherence-based pseudo labeling and cooperative unfolding network. In *Advances in Neural Information Processing Systems*, volume 37, pages 97859–97883, 2024.

Chengyu Fang, Chunming He, Yuelin Zhang, Chubin Chen, Chenyang Zhu, Hongqiu Wang, Longxiang Tang, Xiu Li, and Sina Farsiu. PRISM: Rethinking atmospheric scattering reconstruction as a unified understanding and restoration model for real-world dehazing. *arXiv preprint arXiv:2604.07048*, 2026.

Raanan Fattal. Single image dehazing. *ACM Transactions on Graphics*, 27(3):1–9, 2008.

Yuxin Feng, Long Ma, Xiaozhe Meng, Fan Zhou, Risheng Liu, and Zhuo Su. Advancing real-world image dehazing: Perspective, modules, and training. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):9303–9320, 2024.

Jiayi Fu, Siyu Liu, Zikun Liu, Chun-Le Guo, Hyunhee Park, Ruiqi Wu, Guoqing Wang, and Chongyi Li. Iterative predictor-critic code decoding for real-world image dehazing. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 12700–12709, 2025.

Kaiming He, Jian Sun, and Xiaoou Tang. Single image haze removal using dark channel prior. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 33(12):2341–2353, 2011.

Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. MUSIQ: Multi-scale image quality transformer. In *IEEE International Conference on Computer Vision*, pages 5128–5137, 2021.

Yunwei Lan, Zhigao Cui, Xin Luo, Chang Liu, Nian Wang, Menglin Zhang, Yanzhao Su, and Dong Liu. When Schrödinger bridge meets real-world image dehazing with unpaired training. In *IEEE International Conference on Computer Vision*, 2025.

Boyi Li, Wenqi Ren, Dengpan Fu, Dacheng Tao, Dan Feng, Wenjun Zeng, and Zhangyang Wang. Benchmarking single-image dehazing and beyond. *IEEE Transactions on Image Processing*, 28 (1):492–505, 2019.

Yudong Liang, Bin Wang, Wangmeng Zuo, Jiaying Liu, and Wenqi Ren. Self-supervised learning and adaptation for single image dehazing. In *International Joint Conference on Artificial Intelligence*, pages 1137–1143. International Joint Conferences on Artificial Intelligence Organization, 2022.

Huan Liu, Zijun Wu, Liangyan Li, Sadaf Salehkalaibar, Jun Chen, and Keyan Wang. Towards multi-domain single image dehazing via test-time training. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 5821–5830, 2022.

Ye Liu, Lei Zhu, Shunda Pei, Huazhu Fu, Jing Qin, Qing Zhang, Liang Wan, and Wei Feng. From synthetic to real: Image dehazing collaborating with unlabeled real data. In *ACM International Conference on Multimedia*, pages 50–58, 2021.

Long Ma, Yuxin Feng, Yan Zhang, Jinyuan Liu, Weimin Wang, Guang-Yong Chen, Chengpei Xu, and Zhuo Su. CoA: Towards real image dehazing via compression-and-adaptation. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 11197–11206, 2025.

Mehrsa Pourya, Bassam El Rawas, and Michael Unser. Flower: A flow-matching solver for inverse problems. In *International Conference on Learning Representations*, 2026.

Yuanjie Shao, Lerenhan Li, Wenqi Ren, Changxin Gao, and Nong Sang. Domain adaptation for image dehazing. In *IEEE Conference on Computer Vision and Pattern Recognition*, page 2805–2814, 2020.

Junseong Shin, Seungwoo Chung, Yunjeong Yang, and Tae Hyun Kim. HazeFlow: Revisit haze physical model as ODE and non-homogeneous haze generation for real-world dehazing. In *IEEE International Conference on Computer Vision*, pages 6263–6272, 2025.

Jan Tauberschmidt, Sophie Fellenz, Sebastian Vollmer, and Andrew Duncan. Physics-constrained fine-tuning of flow-matching models for generation and inverse problems. In *International Conference on Learning Representations*, pages 143992–144037, 2026.

Fu-Jen Tsai, Yan-Tsung Peng, Yen-Yu Lin, and Chia-Wen Lin. PHATNet: A physics-guided haze transfer network for domain-adaptive real-world image dehazing. In *IEEE International Conference on Computer Vision*, pages 5591–5600, 2025.

Ruiyi Wang, Wenhao Li, Xiaohong Liu, Chunyi Li, Zicheng Zhang, Xiongkuo Min, and Guangtao Zhai. HazeCLIP: Towards language guided real-world image dehazing. In *IEEE International Conference on Acoustics, Speech and Signal Processing*, pages 1–5, 2025a.

Ruiyi Wang, Yushuo Zheng, Zicheng Zhang, Chunyi Li, Shuaicheng Liu, Guangtao Zhai, and Xiaohong Liu. Learning hazing to dehazing: Towards realistic haze generation for real-world image dehazing. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 23091–23100, 2025b.

Yongzhen Wang, Xuefeng Yan, Fu Lee Wang, Haoran Xie, Wenhan Yang, Xiao-Ping Zhang, Jing Qin, and Mingqiang Wei. UCL-Dehaze: Toward real-world image dehazing via unsupervised contrastive learning. *IEEE Transactions on Image Processing*, 33:1361–1374, 2024.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4): 600–612, 2004.

Chenfeng Wei, Chun Wang, Boyang Zhao, Si Zuo, Shenhong Wang, and Chenguang Yang. Rte-fm-dehazer: Radiative transfer equation inspired flow matching for real-world image dehazing, 2026.

Rui-Qi Wu, Zheng-Peng Duan, Chun-Le Guo, Zhi Chai, and Chongyi Li. Ridcp: Revitalizing real image dehazing via high-quality codebook priors. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 22282–22291, 2023.

Yang Yang, Chaoyue Wang, Risheng Liu, Lin Zhang, Xiaojie Guo, and Dacheng Tao. Self-augmented unpaired image dehazing via density and depth decomposition. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2027–2036, 2022.

Jing Zhang, Yang Cao, Zheng-Jun Zha, and Dacheng Tao. Nighttime dehazing with a synthetic benchmark. In *ACM International Conference on Multimedia*, pages 2355–2363, New York, NY, USA, 2020. Association for Computing Machinery.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 586–595, 2018.

Yan Zhang, Yuxin Feng, Xin Li, Fan Zhou, and Zhuo Su. PDDA: Prompt-driven domain adaptation for real-world image dehazing. *IEEE Transactions on Image Processing*, 35:4280–4294, 2026a.

Yan Zhang, Long Ma, Yuxin Feng, Zhe Huang, Fan Zhou, and Zhuo Su. Bilevel layer-positioning LoRA for real image dehazing. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 29969–29978, 2026b.

Dian Zheng, Xiao-Ming Wu, Shuzhou Yang, Jian Zhang, Jian-Fang Hu, and Wei-Shi Zheng. Selective hourglass mapping for universal image restoration based on diffusion model. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 25445–25455, 2024.

A THEORETICAL ANALYSIS

The analysis follows the two supervision stages: relative calibration controls optical depth differences and enforces reconstruction compatibility, while absolute anchoring selects a joint state within the resulting feasible family.

A.1 RELATIVE CALIBRATION

Relative identifiability. For K constructed views, let $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ be a connected intervention graph with positive edge weights w_{ij}. The general optical depth difference loss is

$$\mathcal{L}_{\text{coord}} = \frac{1}{|\mathcal{E}|} \sum_{(i,j)\in\mathcal{E}} w_{ij} \,\text{mean}\,\text{Charb}\big((\widehat{q}_i - \widehat{q}_j) - (s_i - s_j)\big). \tag{21}$$

The practical choice $\mathcal{E} = \{(0,1)\}$ and $w_{01} = 1$ gives Equation (6), since Charb is even.

For Proposition 1, write $\widehat{q}_i = q_0 + s_i + b_i$, and let $\mathbf{D}$ be an oriented incidence matrix of a connected graph over the K interventions. An edge residual is $\widehat{q}_i - \widehat{q}_j - (s_i - s_j) = b_i - b_j$. Let $L = \mathbf{D}^\top\mathbf{D}$ and $\bar{b} = K^{-1}\sum_i b_i$. Since $\ker\mathbf{D} = \text{span}\{\mathbf{1}\}$, zero edge residuals imply a common optical depth prediction error. Also, $\mathbf{D}\mathbf{1} = 0$ gives $\mathbf{D}b = \mathbf{D}(b - \bar{b}\mathbf{1})$. The vector $b - \bar{b}\mathbf{1}$ is orthogonal to the nullspace of L. The variational characterization of its smallest positive eigenvalue gives

$$\|b - \bar{b}\mathbf{1}\|_2^2 \le \frac{\|\mathbf{D}b\|_2^2}{\lambda_2(L)}. \tag{22}$$

Thus the squared relative residual controls variation across views, whereas $\bar{b}$ remains unrestricted.

Reconstruction compatibility. For any constructed pair, the transition used in Equation (7) generalizes to

$$\widetilde{I}_{j\leftarrow i} = \widehat{A} + e^{-(\widehat{q}_i + s_j - s_i)}(\widehat{J}_i - \widehat{A}). \tag{23}$$

Let $e_i^{\text{rec}} = \widehat{A} + e^{-\widehat{q}_i}(\widehat{J}_i - \widehat{A}) - I_i$ denote its physical reconstruction residual. Using $I_j - \widehat{A} = (\rho_j/\rho_i)(I_i - \widehat{A})$ gives

$$\widetilde{I}_{j\leftarrow i} - I_j = \frac{\rho_j}{\rho_i} e_i^{\text{rec}}. \tag{24}$$

Thus the observable loss is a scaled reconstruction constraint. It couples transmission and radiance predictions, but supplies no independent information about their shared absolute offset. The forward pair $(0,1)$ used in training increases haze and avoids extrapolation to negative optical depth.

A.2 ABSOLUTE ANCHORING

Feasible family and metric. For $\widehat{A}, I_0 \in [0,1]^3$, the interval contains $c = 1$. The constraint $J_k^* \in [0,1]$ for RGB channel k gives upper bounds $c_{\max,k} = (1 - \widehat{A}_k)/r_k$ if $r_k > 0$ and $c_{\max,k} = -\widehat{A}_k/r_k$ if $r_k < 0$. Channels with $r_k = 0$ impose no additional bound. Requiring every projected view transmission $\rho_i/c \ge \tau_{\min}$ gives

$$c_{\max} = \min\left\{\min_i \frac{\rho_i}{\tau_{\min}},\ \min_{k:r_k\neq 0} c_{\max,k}\right\}. \tag{25}$$

Here $\tau_{\min} > 0$ is the transmission floor, and the minimum over an empty channel set is $+\infty$.

The aligned inverse transmission and RGB family used by ASA is

$$\mathcal{C} = \left\{(\kappa_i, J_i)_{i=0}^{K-1} : \kappa_i = c,\ J_i = \widehat{A} + cr,\ 1 \le c \le c_{\max}\right\}. \tag{26}$$

For $c_{\max} \ge 1$, this is a nonempty closed convex segment. For $\zeta = (\kappa_i, J_i)_i$, its weighted norm is

$$\|\zeta\|_W^2 = \sum_i \pi_i \left(w_\kappa|\kappa_i|^2 + w_J\|J_i\|_2^2\right). \tag{27}$$

The positive weights are fixed for each projection, so Equation (13) minimizes $\|\zeta_t - \zeta\|_W^2$ over $\zeta \in \mathcal{C}$.

Closed-form solution. Dropping terms independent of c in (13), its objective has the form

$$f(c) = (w_\kappa + w_J\|r\|_2^2)c^2 - 2\left[w_\kappa\bar{\kappa} + w_J r^\top(\bar{J} - \widehat{A})\right]c + \text{constant}. \tag{28}$$

The quadratic coefficient is strictly positive even when $r = 0$. Its unconstrained minimizer is the ratio in (14); restricting a scalar convex quadratic to an interval clips that minimizer to the interval. When $r = 0$, radiance is fixed at $\widehat{A}$ and the inverse transmission term still selects a unique c^*. For nonempty closed $\mathcal{C}$, this is exactly its projection in the positive definite metric W.

Projected targets and error reduction. The projected optical depths are $q_i^* = \log c^* + s_i$. Together with Equation (15), they satisfy

$$q_i^* - q_0^* = s_i, \qquad I_i + (1 - \tau_i^*)(J^* - \widehat{A}) = J^*. \tag{29}$$

The projected state has $J^* - \widehat{A} = c^* r$ and $I_i = \widehat{A} + \rho_i r$. Consequently,

$$I_i + (1 - \rho_i/c^*)(J^* - \widehat{A}) = \widehat{A} + \rho_i r + (c^* - \rho_i)r = J^*, \tag{30}$$

$$-\log(\rho_i/c^*) - \log c^* = -\log\rho_i = s_i. \tag{31}$$

These equalities follow directly from the constructed group definition.

For any $\zeta \in \mathcal{C}$, the derivative of the squared projection objective in the feasible direction $\zeta - p$ is nonnegative. Hence $\langle \zeta_t - p, \zeta - p\rangle_W \le 0$. For feasible ζ_{gt},

$$\|\zeta_t - \zeta_{\text{gt}}\|_W^2 = \|\zeta_t - p\|_W^2 + \|p - \zeta_{\text{gt}}\|_W^2 + 2\langle \zeta_t - p, p - \zeta_{\text{gt}}\rangle_W \tag{32}$$

$$\ge \|\zeta_t - p\|_W^2 + \|p - \zeta_{\text{gt}}\|_W^2. \tag{33}$$

B Training and Implementation

B.1 Model Initialization and Training Protocol

The backbone is the NCSN++-based U-Net and transmission refinement branch of HazeFlow (Shin et al., 2025). Its source pretraining uses 100K iterations on MCBM-generated hazy images from RIDCP500, Adam, and cosine annealing from 2×10^{-4} to 10^{-6}. The source transmission refinement loss weight is 0.5, with unit atmospheric light labels.

At each accumulation microstep, one batch from each stream contributes to the same trainable model. Each branch loss is averaged over its own samples and pixels, and the accumulated gradients contribute to one optimizer, scheduler, and exponential moving average update.

Algorithms 1 and 2 show one optimizer update without gradient accumulation. With M accumulation microsteps, each branch loss was divided by M before backpropagation, and the optimizer stepped after all M microsteps. Mixed precision used the same gradient scaler for both branches.

B.2 Target Construction and Execution

We used the existing DCP routine with the following call:

```
A_rgb, t_dcp = get_dcp_t(rgb01, return_A=True, A1=True)
```

Here `rgb01` is a contiguous NumPy RGB array of shape $H \times W \times 3$ in $[0, 1]$. The return order is atmospheric light first, then transmission: `A_rgb` has three components and `t_dcp` is an $H \times W$ map. In the public implementation, `A1=True` makes transmission preprocessing use unit atmospheric light, while the returned atmospheric light remains the DCP estimate. Retaining this flag preserves the baseline conditioning convention.

After geometric preprocessing of the original image, we estimated its $\widehat{A}$ once and broadcast it over the crop. All physical interventions and the projection of that group shared this estimate. The DCP condition was recomputed separately for each intervention; the newly returned atmospheric light estimate was discarded. The estimator was fixed and detached, with crop-specific estimates. Cached predictions were associated with the actual crop and preprocessing.

With the group atmospheric light fixed, ASA constructs the feasible interval in Equation (25) and uses the detached teacher mean to set the joint metric. The weights remain fixed during projection and student fitting. Finiteness was checked before logarithms, inverse transmission, and interval construction. Pixels with nonfinite or empty intervals used raw teacher targets.

Training and inference cost. Algorithms 1 and 2 combine this preprocessing with the training protocol in Section 5.1 and the accumulation details in Appendix B.1. RSC evaluates two real views per original crop; ASA uses three frozen teacher predictions and one sampled student view, with source replay in both adaptation stages. Inference retains the HazeFlow architecture and parameter count, with one evaluation of the baseline condition $E(I)$ and one dehazer evaluation per image.

Algorithm 1 Relative-State Calibration with source retention

Require: Source-initialized h_θ; real and source iterators
for each optimizer update **do**
  Real crop batch I_0 and paired source batch S
  Estimate $\widehat{A}$ from I_0; sample ρ
  $I_1 \leftarrow \widehat{A} + \rho(I_0 - \widehat{A})$
  Compute $E(I_0)$ and $E(I_1)$ independently
  Predict $(\widehat{q}_i, \widehat{J}_i)$ for $i = 0, 1$ from the transmission and radiance heads
  Evaluate $\mathcal{L}_{\mathrm{coord}}, \mathcal{L}_{\mathrm{obs}}$ by (6)–(8)
  Evaluate $\mathcal{L}_{\mathrm{src}}(S)$ using the supervised branch
  Update θ once using (10)
end for

Algorithm 2 Absolute-State Anchoring and student training

Require: Frozen calibrated teacher h_θ; student $\phi \leftarrow \theta$
for each optimizer update **do**
  Real crop batch I_0 and paired source batch S
  Estimate $\widehat{A}$; construct $\{I_i\}$ including I_0
  Independently compute all input conditions $E(I_i)$
  With gradients disabled, predict $\widehat{J}_i$ and align κ_i^t by (12)
  Compute fixed W, feasible intervals, and c^* by (14)
  Form $p = (c^*, J^*)_i$ and $\zeta_\eta = (1-\eta)\zeta_t + \eta p$
  Draw $i \sim \pi$ and predict κ_i^ϕ, J_i^ϕ
  Evaluate joint distillation and supervised source losses
  Update ϕ once using (19)
end for

C Additional Experimental Results

C.1 Additional Experiments on Unpaired Datasets

Table 7 reports additional comparisons on NHRW and RealFOGGY. Relative to HazeFlow, R2AS increases MUSIQ from 58.869 to 63.468 on NHRW and from 68.355 to 72.697 on RealFOGGY. NHRW Q-Align decreases from 3.377 to 3.352, while RealFOGGY FADE changes from 0.537 to 0.538.

Table 7: Quantitative results on unpaired datasets (NHRW Zhang et al. (2020) and RealFOGGY). Best and second-best results are shown in bold and underlined, respectively.

	NHRW					RealFOGGY				
Method	MUSIQ↑	MANIQA↑	TOPIQ↑	FADE↓	Q-Align↑	MUSIQ↑	MANIQA↑	TOPIQ↑	FADE↓	Q-Align↑
DEA-Net	48.212	0.300	0.407	0.575	2.961	64.119	0.439	0.576	1.108	3.720
DiffUIR	48.507	0.310	0.414	0.563	2.988	64.128	0.443	0.580	1.351	3.694
DehazeSB	56.980	0.352	0.478	<u>0.323</u>	2.936	69.169	0.477	0.629	0.621	3.735
CoA	51.174	0.242	0.383	0.587	2.916	67.723	0.367	0.560	0.645	3.867
BiLaLoRA	57.256	0.323	0.480	0.506	3.340	70.685	0.448	0.649	0.539	<u>4.098</u>
CORUN	<u>62.102</u>	<u>0.396</u>	<u>0.571</u>	0.567	3.314	<u>71.542</u>	0.472	**0.683**	0.648	4.039
HazeFlow	58.869	0.375	0.515	0.367	**3.377**	68.355	0.454	0.645	**0.537**	3.865
IPC-Dehaze	54.170	0.320	0.453	0.652	3.153	68.196	**0.489**	0.649	0.837	4.009
KA_Net	49.589	0.292	0.388	2.981	2.322	60.877	0.372	0.509	2.735	2.986
PDDA	57.829	0.317	0.461	0.701	3.180	71.371	0.434	0.624	0.884	4.083
RIDCP	56.618	0.287	0.449	0.425	3.155	68.711	0.393	0.621	0.788	3.953
R2AS	**63.468**	**0.411**	**0.598**	**0.280**	<u>3.352</u>	**72.697**	<u>0.484</u>	<u>0.651</u>	<u>0.538</u>	**4.265**

Figure 4 shows two Fattal scenes: vegetation and brickwork, and a foggy pumpkin field. Figure 5 compares nighttime NHRW and daytime RealFOGGY examples.

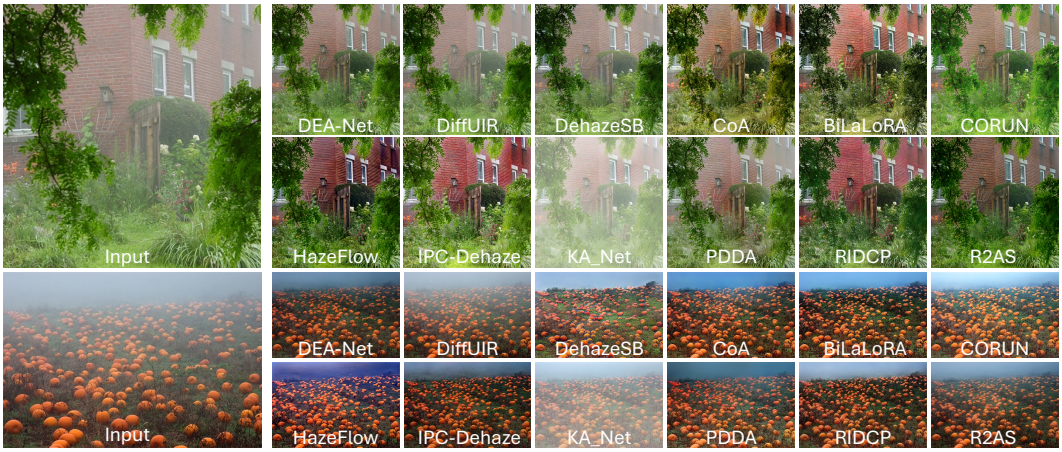


Figure 4: Qualitative dehazing comparison on the Fattal dataset.

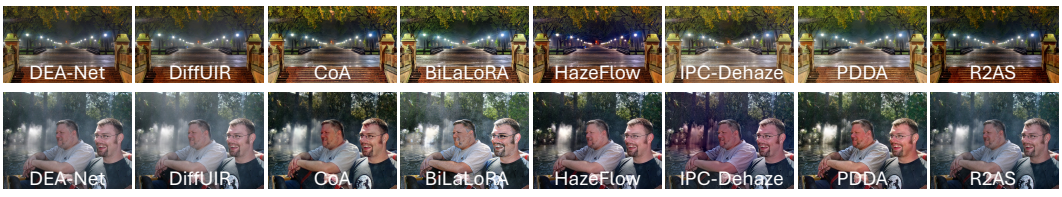


Figure 5: Qualitative dehazing comparison on NHRW (top) and RealFOGGY (bottom).