DISTILLING BOTH CLARITY AND RAIN: EQUIVARIANT SELF-COLLABORATIVE DISTILLATION FOR UNPAIRED IMAGE DERAINING

Anonymous authors
Paper under double-blind review

ABSTRACT

Unpaired image deraining methods usually jointly learn image restoration and realistic rain synthesis from independently collected rainy and clean images, leaving degradation modeling weakly constrained. We propose Equivariant Self-Collaborative Distillation (ESCD), which converts the restoration ability of a frozen derainer into complementary clarity, rain, and objective knowledge. Given a rainy image and an equivariant view, the teacher produces aligned clean predictions and dual rain residuals. The clean predictions distill clarity into the student, while the residuals provide transferable rain structure. A Structure and Appearance Factorized Rain Generator separates clean content, rain appearance, and rain structure, then uses the fused residual representation to gate affine modulation. Rain-Responsive Collaborative Regularization further converts the residual pair into bounded spatial weights for synthesis fidelity and channel consistency, emphasizing rainy regions without discarding background supervision. Iterative Self-Collaborative Distillation promotes the refined student to the next frozen teacher only at cycle boundaries and carries forward the complete collaborative system, which progressively evolves knowledge under stable supervision within each cycle. Extensive experiments on paired synthetic, paired real, and unpaired real datasets demonstrate state of the art performance among existing unpaired deraining approaches, with consistent improvements in full reference fidelity and perceptual image quality as well as strong generalization to real rain.

1 INTRODUCTION

Single image deraining seeks a clean background from an observation corrupted by spatially varying rain streaks, accumulation, and scattering. The task matters for visual quality and downstream perception, which both deteriorate in adverse weather. Supervised derainers have progressed from the detail decomposition of DDN (Fu et al., 2017), the joint detection and removal formulation of JORDER (Yang et al., 2017), density aware modeling in DID-MDN (Zhang & Patel, 2018), and spatial attention in SPANet (Wang et al., 2019), to progressive restoration in MPRNet (Zamir et al., 2021), efficient self attention in Restormer (Zamir et al., 2022), sparse attention in DRSformer (Chen et al., 2023), and implicit multiscale representations in NeRD-Rain (Chen et al., 2024b). Their performance, however, depends on paired data whose synthetic rain often differs from real weather.

Unpaired learning removes the need for aligned data but leaves the inverse problem less constrained. CycleGAN (Zhu et al., 2017) matches domain distributions without guaranteeing content preservation. Deraining methods address this limitation through rain attention in DerainCycleGAN (Wei et al., 2021), directional regularization and distillation in the method of Liu et al. (2021), nonlocal contrastive separation in NLCL (Ye et al., 2022), and dual contrastive alignment in DCD-GAN (Chen et al., 2022b). CSUD (Dong et al., 2025) combines channel consistency with self reconstruction to improve synthesis from pseudo pairs. However, degradation guidance in these methods remains tied to the rainy image rather than separating the clean and rain knowledge encoded by a trained derainer.

RGSUD (Chen et al., 2026) instead selects restorations of high quality from the training history using a pretrained image quality assessor based on a vision language model and recycles those outputs as rewards. Historical quality selection can improve optimization, but the procedure requires an

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.

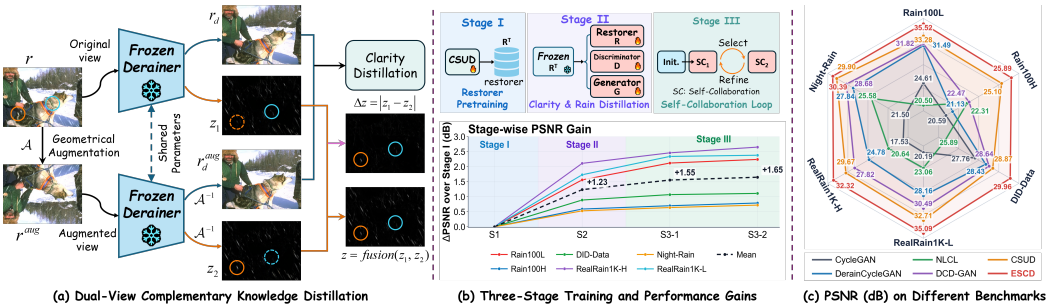

Figure 1: Motivation and empirical evidence for ESCD. (a) A frozen derainer provides aligned clean predictions and residuals from two equivariant views, which support clarity and rain distillation. (b) PSNR gains produced by the initial collaborative transfer and subsequent Iterative Self-Collaborative Distillation on six paired benchmarks. The dashed curve shows the mean across datasets. (c) Final PSNR compared with five representative unpaired baselines on the same benchmarks.

external assessor and principally reuses a selected clean output. A trained derainer, however, already encodes an implicit decomposition of the inverse problem. The clean prediction represents restoration knowledge, while the difference between the rainy input and prediction exposes degradation knowledge. A single prediction can inherit teacher bias, and existing unpaired methods neither coordinate these complementary signals nor progressively update their roles as restoration ability improves. The central problem is therefore how to extract complementary clean and rain knowledge, coordinate the resulting signals, and evolve them iteratively without paired data or external semantic supervision.

To address the above problems, we propose Equivariant Self-Collaborative Distillation (ESCD). As shown in Figure 1(a), the same frozen derainer processes a rainy image and an invertibly transformed view. Aligning the two responses produces complementary clean predictions and rain residuals in a common coordinate system. The clean responses provide stable restoration targets, whereas the residual pair reveals transferable rain structure. This decomposition allows one learned restoration ability to guide both inverse rain removal and forward rain synthesis.

Building on this dual knowledge, the proposed framework coordinates three collaborative paths. Path I performs clarity distillation by transferring both aligned clean predictions to the student. Path II introduces a Structure and Appearance Factorized Rain Generator (SAFRG), which fuses the residual pair, separates clean content, rain appearance, and rain structure, and applies affine modulation gated by the fused structure. Path III introduces Rain-Responsive Collaborative Regularization (RCR), which transforms the dual residuals into bounded spatial weights and applies them to synthesis fidelity and channel consistency. These three paths use shared teacher guidance to connect restoration, synthesis, and objective allocation beyond an isolated student constraint.

Beyond a single transfer step, Iterative Self-Collaborative Distillation (ISCD) repeatedly promotes the refined student to the next frozen teacher and carries forward the student, generator, and discriminator as one collaborative state. Freezing the teacher within each cycle prevents target drift, while replacement at cycle boundaries allows newly acquired restoration ability to generate progressively stronger clarity, rain, and objective signals. Figure 1(b) shows that the mean PSNR gain over the initial restorer reaches 1.23 dB after the first collaborative transfer and continues to 1.55 and 1.65 dB through ISCD. Figure 1(c) further demonstrates the strong overall PSNR advantage of the final model against representative unpaired baselines across six paired benchmarks. Our contributions are summarized as follows:

- We propose equivariant distillation of complementary knowledge for unpaired deraining. Aligned teacher predictions transfer clarity knowledge, while the corresponding residuals provide rain knowledge for synthesis and objective guidance.
- We develop SAFRG and RCR to coordinate rain synthesis and objective allocation. SAFRG separates content, appearance, and structure and uses the fused residual representation to gate affine modulation, while RCR converts the same residual pair into bounded spatial weights for synthesis fidelity and channel consistency.

- We introduce iterative self-collaborative distillation, which updates the frozen teacher and the complete collaborative state between cycles to evolve knowledge progressively while preserving stable supervision within each cycle.

- Extensive experiments on paired synthetic, paired real, and unpaired real datasets establish state of the art performance among the compared unpaired methods in both reference based fidelity and perceptual quality.

2 Related Work

2.1 Supervised image deraining

Supervised image deraining learns mappings from rainy to clean images using aligned pairs, progressing from convolutional and recurrent models (Fu et al., 2017; Yang et al., 2017; Li et al., 2018; Ren et al., 2019) to sparse attention and implicit multiscale representations (Chen et al., 2023; 2024b). Recent work extends this progress to more diverse rain conditions. RLP (Zhang et al., 2025) guides feature modulation with rain location priors for nighttime deraining, while CST-Net (Guan et al., 2025) combines a learnable color space transformation with implicit illumination guidance. UniRain (Yang et al., 2026) unifies daytime and nighttime removal of rain streaks and raindrops through dataset distillation, an asymmetric mixture of experts, and reweighting for multiple objectives. Despite their strong performance, supervised derainers typically rely on large paired datasets that are difficult to acquire in real scenes, while synthetic training data can limit generalization (Gu et al., 2023). These limitations motivate our method, which exploits complementary clean and rain knowledge from a frozen derainer for unpaired training.

2.2 Unpaired image deraining

To reduce reliance on paired data, unpaired deraining learns from independently collected rainy and clean images. Early studies explored image translation, rain removal without clean targets, domain transfer, and joint rain generation and removal (Zhu et al., 2017; Jin et al., 2019; Wei et al., 2019; Yasarla et al., 2020; Ye et al., 2021). Subsequent work improved cyclic synthesis, regularization, contrastive learning, and perceptual objectives (Wei et al., 2021; Liu et al., 2021; Ye et al., 2022; Chen et al., 2022b; Xu et al., 2022). CSUD (Dong et al., 2025) constrains rain synthesis through channel consistency and self reconstruction. RGSUD (Chen et al., 2026) selects historical restorations using a vision language IQA model. Our framework instead derives complementary clean predictions and rain residuals from a frozen derainer to coordinate restoration, synthesis, and spatial regularization. Appendix Table 9 compares their supervision sources and signal roles.

2.3 Knowledge distillation for image restoration

Knowledge distillation transfers predictive information from a teacher or ensemble to a student (Hinton et al., 2015). Later variants extend this principle through successive generations of teachers and students, peer learning, and consistency across perturbed views (Furlanello et al., 2018; Zhang et al., 2018; Tarvainen & Valpola, 2017). Recent restoration research applies distillation in several forms. Semantic prior distillation (Zhang et al., 2024) transfers knowledge from SAM to efficient restoration networks, IR-CM (Gong & Ma, 2024) uses consistency learning for restoration in few steps, and AdcSR (Chen et al., 2025) compresses a diffusion restorer through adversarial distillation. Despite these advances, distillation remains relatively underexplored in deraining without fully paired supervision, with only a few studies adopting data distillation, knowledge transfer between a teacher and student, or related strategies (Lin et al., 2020; Liu et al., 2021; Cui et al., 2022; Chen et al., 2024a). The proposed approach uses a frozen derainer to produce aligned clean predictions and rain residuals from two equivariant views. These signals jointly guide restoration, rain synthesis, and spatial regularization. ISCD promotes each refined student to the next frozen teacher, allowing complementary knowledge to evolve across cycles without destabilizing supervision within a cycle.

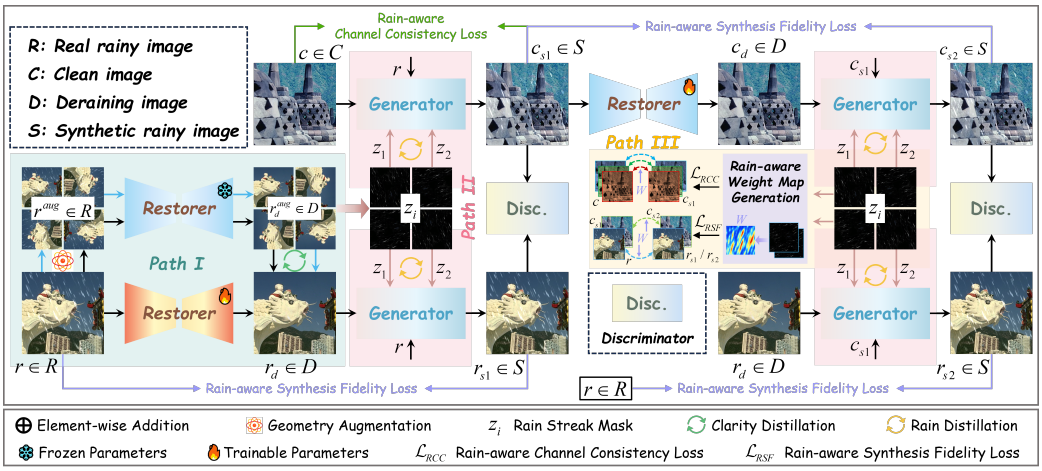

Figure 2: Overall training framework of ESCD. Path I aligns original and transformed teacher predictions to distill clarity into the student. Path II uses the dual residuals (z_1, z_2), together with clean content and a rainy reference, throughout the cyclic synthesis graph. Path III performs RCR by mapping the residual pair to W for the synthesis fidelity loss $\mathcal{L}_{\text{RSF}}$ and channel consistency loss $\mathcal{L}_{\text{RCC}}$. The teacher, generator, and discriminator are used only for training.

3 METHOD

3.1 PROBLEM SETTING AND OVERVIEW

Our Equivariant Self-Collaborative Distillation framework (ESCD) addresses unpaired deraining by reusing learned restoration ability to coordinate restoration, rain synthesis, and spatial objective allocation. Let $\mathcal{R} = \{r_i\}$ and $\mathcal{C} = \{c_j\}$ denote unpaired rainy and clean sets. The student $F_\theta : \mathcal{R} \to \mathcal{C}$ predicts clean content, while the generator $G_\phi(c, r, \mathbf{z})$ transfers rain using a rainy reference and complementary residuals $\mathbf{z} = (z_1, z_2)$. As summarized in Figure 1(b), Stage I trains CSUD (Dong et al., 2025), Stage II freezes the resulting derainer as $F_{\bar{\theta}}$ to supervise the collaborative graph, and Stage III promotes each refined student to the next teacher between cycles. In Figure 2, the three paths respectively distill aligned clean predictions, guide rain transfer with residual structure, and convert the same residuals into RCR weights. More details will be provided in the following.

3.2 EQUIVARIANT DISTILLATION OF COMPLEMENTARY KNOWLEDGE

A single frozen prediction can serve as a pseudo target but retains teacher error specific to one view and provides only one decomposition of rain and background. We therefore query the same teacher under two views that preserve the image grid. Alignment makes shared clean content a more stable restoration target. The corresponding residuals between inputs and outputs provide complementary evidence about where the teacher attributes degradation. This dual use distinguishes the views from ordinary data augmentation.

For rainy image r, we sample an invertible transform τ from horizontal flip, vertical flip, and $180°$ rotation. The complete decomposition from two views is

$$\bar{c}_1 = F_{\bar{\theta}}(r), \quad \bar{c}_2 = \tau^{-1}(F_{\bar{\theta}}(\tau(r))), \tag{1}$$

$$z_1 = |r - \bar{c}_1|, \quad z_2 = |r - \bar{c}_2|, \quad \mathbf{z} = (z_1, z_2), \quad z = \Psi(\mathbf{z}) = \sqrt{z_1^2 + z_2^2}. \tag{2}$$

The inverse transform returns both predictions to the coordinates of r, and all quantities derived from the teacher are detached from gradient updates. The generator receives the pair $\mathbf{z}$ shown in Figures 2 and 3. The structure branch first combines the two residual maps with the fixed element-wise operator Ψ and then encodes the fused map z. The individual maps z_1 and z_2 remain available to generate the spatial loss weights in Path III.

The student $\hat{c} = F_\theta(r)$ is optimized against both clean responses,

$$\mathcal{L}_{\text{clean}} = \frac{1}{2}\sum_{v=1}^{2}\left[\|\hat{c} - \bar{c}_v\|_1 + \ell_{\text{SSIM}}(\hat{c}, \bar{c}_v)\right], \tag{3}$$

where ℓ_{SSIM} follows Wang et al. (2004). Separate constraints avoid relying on a potentially biased pseudo label. Equations (1)-(3) establish a teacher that provides two forms of knowledge. Clean responses guide inverse restoration, while residual responses coordinate rain synthesis.

3.3 Structure and Appearance Factorized Rain Generator

SAFRG must extract rain appearance from r without transferring the reference scene. Directly encoding the rainy image entangles appearance and structure. Appearance specifies *how* rain looks, whereas structure specifies *where* rain should appear. Neither factor should overwrite clean content. We assign these roles to separate streams,

$$\begin{gathered} f_c = E_c(c), \quad f_a = E_a(r), \quad f_s = E_s(\Psi(\mathbf{z})) = E_s(z), \\ m = \sigma(H(f_s)), \quad q = f_a + f_s, \end{gathered} \tag{4}$$

where E_c, E_a, and E_s are the encoders for clean content, rain appearance, and prior structure labeled CBEM, RSEM, and PSEM in Figure 3. Their outputs f_c, f_a, and f_s keep the three factors separate before modulation. The mask head H converts the structure feature into gate m, so the teacher residuals determine where modulation conditioned on appearance acts.

Appearance and structure jointly parameterize an affine transform over features following FiLM (Perez et al., 2018), but the transform acts only where the gate is active:

$$\tilde{f}_c = f_c \odot (1 + \gamma(q) \odot m) + \beta(q) \odot m, \tag{5}$$

$$\rho = G_\rho(\tilde{f}_c), \quad \tilde{r} = G_\phi(c, r, \mathbf{z}) = c + \rho. \tag{6}$$

The masked residual parameterization preserves two identity cases. When $m = 0$, content features remain unchanged. When $\rho = 0$, the synthesized output equals c. The generator follows residual learning in ResNet (He et al., 2016) and image translation in the method of Johnson et al. (2016). Each encoder uses 5×5 and 3×3 convolutions to produce 128-channel features before six residual blocks and decoder G_ρ. The adversarial branch uses the PatchGAN discriminator of CycleGAN (Zhu et al., 2017), with a standard GAN realism objective (Goodfellow et al., 2014).

The synthesized image $\tilde{r}_c = G_\phi(c, r, \mathbf{z})$ pairs with known clean content c for student supervision. Setting both elements of $\mathbf{z}$ to zero provides the self reconstruction paths inherited from CSUD (Dong et al., 2025). The student must remove transferred rain, while the generator must preserve content when no rain structure is requested. Appendix D gives the path definitions and loss weights.

3.4 Rain-Responsive Collaborative Regularization

RCR uses the distilled residuals to guide both generated rain and the spatial allocation of training emphasis. Uniform reconstruction weights easy background pixels as strongly as regions dominated by rain. A hard residual mask, however, would discard useful context and amplify teacher errors. We therefore construct a bounded response from the two residual views and normalize the response within each image:

$$\begin{aligned} z_{\text{avg}} &= \tfrac{1}{2}[\text{mean}_{ch}(z_1) + \text{mean}_{ch}(z_2)], \\ W(z_{\text{avg}}) &= 1 + \alpha \ \text{sigmoid}\left[\omega\left(\frac{z_{\text{avg}} - \mu(z_{\text{avg}})}{\sigma(z_{\text{avg}}) + \epsilon} - \kappa\right)\right]. \end{aligned} \tag{7}$$

Here μ and σ are spatial statistics computed for each image. We set $(\alpha, \omega, \kappa, \epsilon) = (2, 2, 1, 10^{-5})$, so all pixels retain a unit base weight, while larger normalized residual responses receive progressively stronger emphasis. The teacher therefore allocates emphasis without suppressing the background.

The response enters the two rain aware constraints named in Figures 2 and 3. For a pair of synthesized output and target (u, v), the rain aware synthesis fidelity loss $\mathcal{L}_{\text{RSF}}$ measures spatially weighted

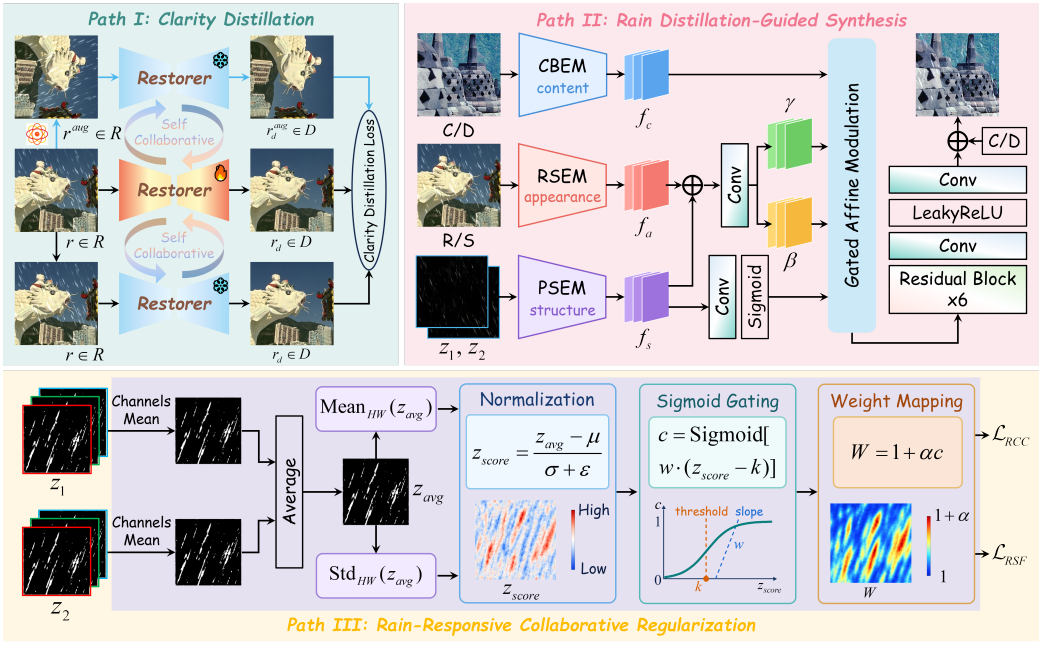

Figure 3: Details of the three collaborative paths. Path I aligns the original and augmented predictions from the frozen teacher and distills both into the student. In Path II, SAFRG uses CBEM, RSEM, and PSEM to encode clean content, rain appearance, and rain structure fused from (z_1, z_2). The structure feature gates affine modulation before residual rain decoding. Path III performs RCR by normalizing the residual response and mapping the result to bounded weight W used by $\mathcal{L}_{\mathrm{RSF}}$ and $\mathcal{L}_{\mathrm{RCC}}$.

image fidelity. For synthesized image $\tilde{r}_c$ and clean content c, the rain aware channel consistency loss $\mathcal{L}_{\mathrm{RCC}}$ encourages the weighted transferred residual to remain similar across color channels:

$$\mathcal{L}_{\mathrm{RSF}}(u, v) = \frac{1}{CHW} \|W(z_{\mathrm{avg}}) \odot (u - v)\|_1, \tag{8}$$

$$e_k = W(z_{\mathrm{avg}}) \odot (c_k - \tilde{r}_{c,k}), \quad k \in \{R, G, B\},$$

$$\mathcal{L}_{\mathrm{RCC}} = \|e_R - e_G\|_1 + \|e_G - e_B\|_1 + \|e_B - e_R\|_1. \tag{9}$$

The first term is applied to three closure pairs in the cyclic graph, whereas the second concentrates the channel consistency prior of CSUD in likely rain regions. Together, the losses make the dual residuals a source of *objective knowledge* as well as the structural input to synthesis. Appendix D lists the reconstruction pairs and complete objectives.

3.5 Iterative Self-Collaborative Distillation

The complete collaborative graph first learns from the frozen restorer obtained in Stage I. ISCD then promotes the refined student to the teacher of the next cycle and carries forward the associated student, generator, and discriminator checkpoints. Continuous teacher updates would shift clean targets and rain coordinates during student optimization. Freezing the teacher within each cycle separates *knowledge use* from *knowledge evolution*. The targets remain stationary during optimization, while replacement at the cycle boundary allows improved restoration ability to regenerate clarity, rain, and objective knowledge for the next cycle.

Appendix Algorithm 1 separates the complete collaborative graph from the gains of iterative distillation. Stage II includes every architectural and objective innovation and initializes all trainable networks from scratch. Stage III performs ISCD without changing the graph or objectives. Teacher parameters remain fixed within each cycle, making the derived targets and priors stationary under gradient updates. At the next boundary, selected student weights initialize both the new student and frozen teacher, while the generator and discriminator inherit their corresponding selected weights.

Table 1: Quantitative deraining comparisons on synthetic and real paired datasets using PSNR and SSIM. The upper block reports supervised methods, and the lower block reports methods trained without paired supervision. A dash denotes an unavailable published result. Among unpaired methods, the best and second best results are highlighted in purple and blue, respectively.

Datasets	**Venue**	*Synthetic Paired Datasets*						*Real Paired Datasets*					
		Rain100L		Rain100H		DID-Data		RealRain1K-L		RealRain1K-H		Night-Rain	
Metrics		PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑	PSNR↑	SSIM↑
DDN	CVPR'17	32.38	0.926	24.64	0.849	30.97	0.912	31.18	0.917	29.17	0.878	32.43	0.926
SPANet	CVPR'19	31.95	0.919	26.07	0.857	33.04	0.949	30.43	0.947	25.76	0.910	31.10	0.925
MPRNet	CVPR'21	34.95	0.959	28.53	0.872	33.99	0.959	36.29	0.972	34.74	0.964	36.63	0.966
Restormer	CVPR'22	37.57	0.974	29.46	0.889	35.29	0.964	40.90	0.985	39.57	0.981	36.92	0.969
DRSformer	CVPR'23	41.23	0.983	29.53	0.837	35.35	0.965	38.84	0.982	–	–	37.86	0.971
NeRD-Rain-S	CVPR'24	42.00	0.990	32.86	0.932	35.36	0.965	38.64	0.979	36.69	0.970	38.42	0.973
CycleGAN	ICCV'17	24.61	0.834	20.59	0.704	27.76	0.838	20.19	0.820	17.53	0.762	21.50	0.783
DerainCycleGAN	TIP'21	31.49	0.936	21.13	0.710	28.43	0.864	28.16	0.901	24.78	0.844	27.84	0.851
NLCL	CVPR'22	20.50	0.719	22.31	0.728	25.89	0.813	23.06	0.832	20.64	0.719	25.58	0.821
DCD-GAN	CVPR'22	31.82	0.941	22.47	0.753	28.64	0.862	30.49	0.939	27.82	0.895	28.68	0.867
CSUD	CVPR'25	33.28	0.954	25.10	0.811	28.87	0.863	32.71	0.959	29.67	0.928	29.90	0.879
RGSUD	CVPR'26	34.41	0.967	–	–	29.07	0.866	32.88	0.955	–	–	**30.54**	0.897
ESCD (Ours)	—	**35.52**	**0.968**	**25.89**	**0.832**	**29.96**	**0.867**	**35.09**	**0.970**	**32.32**	**0.954**	30.39	**0.902**

Appendix E analyzes alignment, residual fusion, bounded weighting, and identity preservation under modulation. Appendices D and A provide the objectives and checkpoint semantics.

4 Experiments

4.1 Experimental setup

Datasets. Evaluation covers paired and unpaired rain datasets. The paired benchmarks are Rain100L and Rain100H (Yang et al., 2017), DID-Data (Zhang & Patel, 2018), RealRain1K-L and RealRain1K-H (Li et al., 2022), and Night-Rain (Zhang et al., 2022; 2023). Each paired experiment uses the training and test splits of the corresponding dataset. Transfer to real scenes is additionally assessed on SIRR (Wei et al., 2019) and Real3000 (Liu et al., 2021), which lack aligned clean references.

Metrics. PSNR and SSIM (Wang et al., 2004) measure fidelity against a reference on paired benchmarks. We use the same four metrics without a reference on paired and unpaired data: CLIP-IQA (Wang et al., 2023), Q-Align (Wu et al., 2024), DeQA-Score (You et al., 2025), and NIMA (Talebi & Milanfar, 2018). Together they cover perceptual assessment based on CLIP, scoring aligned with a language model, regression over score distributions, and aesthetic quality.

Implementation details. We evaluate our method with PyTorch and train all models on a single NVIDIA GeForce RTX 4090 GPU. Training images are randomly cropped into 256×256 patches with a batch size of 2, and the initial learning rate is set to 2×10^{-4}. The first refinement cycle runs for 100k iterations, followed by 50k iterations in the second cycle for every dataset. For the rain-responsive weight map in Equation (7), the scaling factor α, slope factor ω, and response threshold κ are set to 2, 2, and 1, respectively. Further implementation details and sensitivity analyses are provided in Appendices A and A.1, respectively.

4.2 Comparison on paired datasets

We evaluate the proposed method using PSNR and SSIM on six paired benchmarks covering synthetic, daytime real, and nighttime rain. The unpaired comparison in Table 1 includes CycleGAN (Zhu et al., 2017), DerainCycleGAN (Wei et al., 2021), NLCL (Ye et al., 2022), DCD-GAN (Chen et al., 2022b), CSUD (Dong et al., 2025), and RGSUD (Chen et al., 2026). We also report results for supervised deraining models DDN (Fu et al., 2017), SPANet (Wang et al., 2019), MPRNet (Zamir et al., 2021), Restormer (Zamir et al., 2022), DRSformer (Chen et al., 2023), and NeRD-Rain-S (Chen et al., 2024b). The proposed model achieves the highest PSNR on five datasets and the highest SSIM on all six among the unpaired methods in Table 1. Compared with RGSUD, our approach improves PSNR by 1.11, 0.89, and 2.21 dB on Rain100L, DID-Data, and RealRain1K-L, respectively. On Night-Rain,

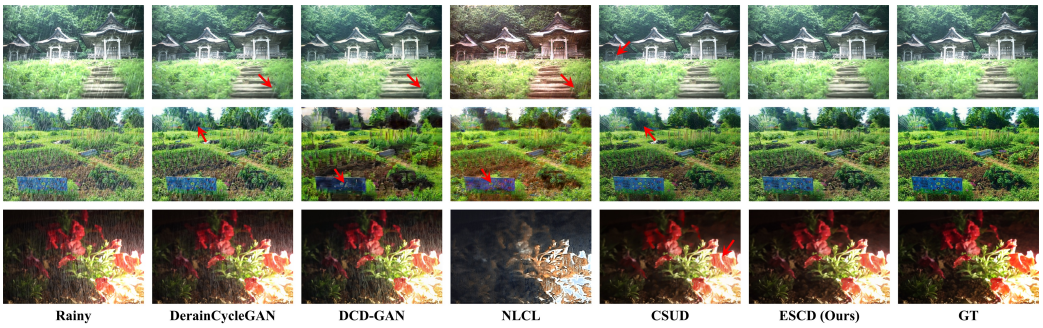

Figure 4: Visual comparisons on representative paired benchmarks. Red arrows mark residual rain, color shifts, and local structural artifacts in competing outputs. The proposed method removes rain while preserving scene appearance close to GT.

Table 2: Perceptual quality without references on the unpaired SIRR and Real3000 datasets with real rain. All methods are trained on the Rain100L dataset.

Setting	Dataset	Metric	*Supervised*	*Unpaired*				
			NeRD-Rain-S	Derain CycleGAN	NLCL	DCD-GAN	CSUD	**ESCD**
Unpaired	SIRR	CLIP-IQA↑	0.480	0.427	0.374	0.312	0.422	**0.429**
		Q-Align↑	3.625	3.462	3.053	3.107	3.437	**3.531**
		DeQA-Score↑	3.097	3.012	2.576	2.691	3.066	**3.112**
		NIMA↑	0.514	0.487	0.476	0.480	0.520	**0.521**
	Real3000	CLIP-IQA↑	0.501	0.472	0.398	0.323	0.475	**0.481**
		Q-Align↑	4.026	3.931	3.420	3.539	3.973	**4.032**
		DeQA-Score↑	3.529	3.478	2.966	3.071	3.532	**3.578**
		NIMA↑	0.585	0.563	0.534	0.529	0.588	**0.595**

the resulting restorer achieves 0.902 SSIM compared with 0.897 for RGSUD, while PSNR remains close at 30.39 versus 30.54 dB. On Rain100L, the same model also surpasses the supervised MPRNet in both metrics, reaching 35.52 dB PSNR and 0.968 SSIM. Figure 4 complements the reference based scores with visual evidence. Competing methods leave visible streaks or alter local appearance in challenging regions, whereas the proposed method produces cleaner backgrounds and preserves scene content closer to GT. Appendix Figure 5 provides further comparisons on real rainy images without aligned clean targets.

4.3 Perceptual Evaluation Without References

Full reference metrics do not completely reflect perceptual quality on Rain100L and DID-Data, while SIRR and Real3000 lack aligned clean targets. We therefore evaluate all four datasets using the same four metrics without references. Table 2 reports SIRR and Real3000, while Appendix Table 8 reports the paired datasets with GT included only as a diagnostic reference. The comparison covers supervised NeRD-Rain-S (Chen et al., 2024b) and unpaired DerainCycleGAN (Wei et al., 2021), NLCL (Ye et al., 2022), DCD-GAN (Chen et al., 2022b), and CSUD (Dong et al., 2025). Among the unpaired methods, our model ranks first in all eight comparisons on SIRR and Real3000. The paired results add five first place and three second place rankings, yielding 13 firsts and three seconds across all 16 comparisons. Relative to CSUD, the proposed approach improves Q-Align, DeQA-Score, and NIMA on every dataset and improves CLIP-IQA on DID-Data, SIRR, and Real3000, with only a 0.006 decrease on Rain100L. Although supervised NeRD-Rain-S remains stronger on DID-Data, the consistent gains over CSUD support improved perceptual quality without paired supervision or an external quality assessor.

4.4 Ablation Studies

Ablation setting. All ablations use Stage II and the same protocol unless noted. Table 3 adds components cumulatively, while Table 4 changes one design at a time. The ISCD rows evaluate later refinement, and P/S denotes PSNR/SSIM.

Table 3: Ablation of Component and refinement.

Variant	CD	SAFRG	RSF	RCC	Rain100L P/S
Base model					33.28/0.954
+ CD	✓				33.40/0.954
+ SAFRG	✓	✓			34.09/0.959
+ RSF	✓	✓	✓		34.50/0.961
+ RCC (full Stage II)	✓	✓	✓	✓	**34.84/0.964**
+ ISCD cycle 1	✓	✓	✓	✓	35.40/0.967
+ ISCD cycle 2	✓	✓	✓	✓	35.52/0.968

Table 4: Ablation of key design choices.

Variant	Rain100L P/S	RealRain1K-L P/S	Night-Rain P/S
Single view prior	33.31/0.955	32.83/0.959	29.72/0.887
Mean fusion prior	34.29/0.960	33.87/0.961	30.01/0.892
Residual fusion	**34.84/0.964**	**34.44/0.965**	**30.20/0.899**
Concatenation	34.27/0.960	34.10/0.964	30.09/0.893
SAFRG	**34.84/0.964**	**34.44/0.965**	**30.20/0.899**
Uniform weighting	34.12/0.955	34.03/0.964	29.97/0.891
RCR weighting	**34.84/0.964**	**34.44/0.965**	**30.20/0.899**

Effectiveness of collaborative paths. Table 3 shows monotonic gains as clarity distillation (CD), SAFRG, RSF, and RCC are introduced. CD raises PSNR from 33.28 to 33.40 dB, after which SAFRG provides the largest individual gain of 0.69 dB and 0.005. RSF and RCC further contribute 0.41/0.002 and 0.34/0.003, respectively. The complete Stage II model reaches 34.84 dB and 0.964, exceeding the base model by 1.56 dB and 0.010. This progression shows that clarity transfer, rain synthesis, and objective allocation provide complementary supervision.

Effectiveness of iterative refinement. Starting from the complete Stage II model, the first ISCD cycle improves PSNR/SSIM from 34.84/0.964 to 35.40/0.967, and the second cycle further reaches 35.52/0.968. Iterative refinement therefore adds 0.68 dB and 0.004 over Stage II, producing an overall gain of 2.24 dB and 0.014 over the base model. The smaller but still positive increment in the second cycle indicates diminishing returns while confirming that teacher promotion continues to improve the collaborative system without degrading the previous solution.

Effectiveness of teacher prior and rain generation. The first three rows of Table 4 show that dual teacher views consistently outperform a single view on Rain100L, RealRain1K-L, and Night-Rain. Mean fusion improves the single view prior by 0.98, 1.04, and 0.29 dB, while the residual fusion in Equation (2) adds another 0.55, 0.57, and 0.19 dB. The corresponding SSIM gains over mean fusion are 0.004, 0.004, and 0.007. This consistency supports retaining complementary residual responses rather than averaging away strong evidence. The generator comparison leads to the same conclusion across different rain domains. Relative to direct feature concatenation, SAFRG improves PSNR/SSIM by 0.57/0.004 on Rain100L, 0.34/0.001 on RealRain1K-L, and 0.11/0.006 on Night-Rain, validating the separation of content, rain appearance, and rain structure with gated modulation.

Effectiveness of objective weighting. Replacing uniform RSF and RCC with RCR weighting improves PSNR/SSIM by 0.72/0.009 on Rain100L, 0.41/0.001 on RealRain1K-L, and 0.23/0.008 on Night-Rain. The consistent PSNR gains and the larger SSIM improvements on Rain100L and Night-Rain indicate that residual guided spatial weighting emphasizes rain affected regions while preserving structural supervision over the background.

Additional ablations on the RCR parameters and the number of ISCD cycles are provided in Appendix A.1.

5 Conclusion

In this work, we presented ESCD for unpaired image deraining, which converts the restoration ability of a frozen derainer into complementary clarity, rain, and objective knowledge. Equivariant clean predictions, factorized rain generation, and rain-responsive regularization jointly coordinate restoration, synthesis, and objective allocation. Iterative self-collaborative distillation further promotes the refined student and collaborative state between cycles, allowing the transferred knowledge to evolve under stable supervision. Extensive experiments demonstrate state of the art performance among the compared unpaired methods across paired synthetic, paired real, and unpaired real datasets. The proposed method achieves the best SSIM on all six paired benchmarks and ranks first or second in all 16 evaluations without references. Ablation results verify the contribution of each collaborative path and the continued gains from iterative refinement. All auxiliary branches are confined to training, so inference retains only the student restorer. Results establish complementary knowledge transfer and iterative collaboration as an effective solution for learning image deraining without paired data.

AI USE STATEMENT

Generative AI tools (OpenAI language models) were used to aid and polish manuscript writing, including language refinement and translation, and to support the retrieval and discovery of related work. The authors verified the resulting text and all retrieved references against the original sources and take full responsibility for the manuscript.

ETHICS STATEMENT

The authors acknowledge that this work adheres to the ICLR Code of Ethics.

REFERENCES

Bin Chen, Gehui Li, Rongyuan Wu, Xindong Zhang, Jie Chen, Jian Zhang, and Lei Zhang. Adversarial diffusion compression for real-world image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 28208–28220, 2025.

Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *European Conference on Computer Vision*, pp. 17–33. Springer, 2022a.

Tingting Chen, Beibei Lin, Yeying Jin, Wending Yan, Wei Ye, Yuan Yuan, and Robby T. Tan. Dual-Rain: Video rain removal using assertive and gentle teachers. In *European Conference on Computer Vision*, pp. 127–143, 2024a.

Xiang Chen, Jinshan Pan, Kui Jiang, Yufeng Li, Yufeng Huang, Caihua Kong, Longgang Dai, and Zhentao Fan. Unpaired deep image deraining using dual contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2017–2026, 2022b.

Xiang Chen, Hao Li, Mingqiang Li, and Jinshan Pan. Learning a sparse transformer network for effective image deraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5896–5905, 2023.

Xiang Chen, Jinshan Pan, and Jiangxin Dong. Bidirectional multi-scale implicit neural representations for image deraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 25627–25636, 2024b.

Yinghao Chen, Yeying Jin, Xiang Chen, Yanyan Wei, Ziyang Yan, and Yaowen Fu. Unpaired image deraining using reward-guided self-reinforcement strategy. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 1342–1354, 2026.

Xin Cui, Cong Wang, Dongwei Ren, Yunjin Chen, and Pengfei Zhu. Semi-supervised image deraining using knowledge distillation. *IEEE Transactions on Circuits and Systems for Video Technology*, 32 (12):8327–8341, 2022.

Guanglu Dong, Tianheng Zheng, Yuanzhouhan Cao, Linbo Qing, and Chao Ren. Channel consistency prior and self-reconstruction strategy based unsupervised image deraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7469–7479, 2025.

Xueyang Fu, Jiabin Huang, Delu Zeng, Yue Huang, Xinghao Ding, and John Paisley. Removing rain from single images via a deep detail network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1715–1723, 2017.

Tommaso Furlanello, Zachary Lipton, Michael Tschannen, Laurent Itti, and Anima Anandkumar. Born again neural networks. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 1607–1616. PMLR, 2018.

Xiaoxuan Gong and Jie Ma. IR-CM: The fast and general-purpose image restoration method based on consistency model. In *Advances in Neural Information Processing Systems*, volume 37, 2024.

Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in Neural Information Processing Systems*, volume 27, pp. 75559–75584, 2014.

Jinjin Gu, Xianzheng Ma, Xiangtao Kong, Yu Qiao, and Chao Dong. Networks are slacking off: Understanding generalization problem in image deraining. In *Advances in Neural Information Processing Systems*, volume 36, pp. 28565–28584, 2023.

Qiyuan Guan, Xiang Chen, Guiyue Jin, Jiyu Jin, Shumin Fan, Tianyu Song, and Jinshan Pan. Rethinking nighttime image deraining via learnable color space transformation. In *Advances in Neural Information Processing Systems*, volume 38, pp. 3189–3225, 2025.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 770–778, 2016.

Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2015.

Xin Jin, Zhibo Chen, Jianxin Lin, Zhikai Chen, and Wei Zhou. Unsupervised single image deraining with self-supervised constraints. In *IEEE International Conference on Image Processing*, pp. 2761–2765, 2019.

Justin Johnson, Alexandre Alahi, and Li Fei-Fei. Perceptual losses for real-time style transfer and super-resolution. In *European Conference on Computer Vision*, pp. 694–711, 2016.

Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *International Conference on Learning Representations*, 2015.

Wei Li, Qiming Zhang, Jing Zhang, Zhen Huang, Xinmei Tian, and Dacheng Tao. Toward real-world single image deraining: A new benchmark and beyond. *arXiv preprint arXiv:2206.05514*, 2022.

Xia Li, Jianlong Wu, Zhouchen Lin, Hong Liu, and Hongbin Zha. Recurrent squeeze-and-excitation context aggregation net for single image deraining. In *European Conference on Computer Vision*, pp. 262–277, 2018.

Huangxing Lin, Yanlong Li, Xueyang Fu, Xinghao Ding, Yue Huang, and John W. Paisley. Rain o'er me: Synthesizing real rain to derain with data distillation. *IEEE Transactions on Image Processing*, 29:7668–7680, 2020.

Yang Liu, Ziyu Yue, Jinshan Pan, and Zhixun Su. Unpaired learning for deep image deraining with rain direction regularizer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4733–4741, 2021.

Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. FiLM: Visual reasoning with a general conditioning layer. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, pp. 3942–3951, 2018.

Dongwei Ren, Wangmeng Zuo, Qinghua Hu, Pengfei Zhu, and Deyu Meng. Progressive image deraining networks: A better and simpler baseline. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3932–3941, 2019.

Hossein Talebi and Peyman Milanfar. NIMA: Neural image assessment. *IEEE Transactions on Image Processing*, 27(8):3998–4011, 2018.

Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *Advances in Neural Information Processing Systems*, volume 30, pp. 1195–1204, 2017.

Jianyi Wang, Kelvin C. K. Chan, and Chen Change Loy. Exploring CLIP for assessing the look and feel of images. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 2555–2563, 2023.

Tianyu Wang, Xin Yang, Ke Xu, Shaozhe Chen, Qiang Zhang, and Rynson W. H. Lau. Spatial attentive single-image deraining with a high quality real rain dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12270–12279, 2019.

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.

Wei Wei, Deyu Meng, Qian Zhao, Zongben Xu, and Ying Wu. Semi-supervised transfer learning for image rain removal. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 3877–3886, 2019.

Yanyan Wei, Zhao Zhang, Yang Wang, Mingliang Xu, Yi Yang, Shuicheng Yan, and Meng Wang. DerainCycleGAN: Rain attentive CycleGAN for single image deraining and rainmaking. *IEEE Transactions on Image Processing*, 30:4788–4801, 2021.

Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Yixuan Gao, Annan Wang, Erli Zhang, Wenxiu Sun, Qiong Yan, Xiongkuo Min, Guangtao Zhai, and Weisi Lin. Q-Align: Teaching LMMs for visual scoring via discrete text-defined levels. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 54015–54029, 2024.

Wei Xu, Xinyuan Chen, Haoming Guo, Xiaolin Huang, and Wei Liu. Unsupervised image restoration with quality-task-perception loss. *IEEE Transactions on Circuits and Systems for Video Technology*, 32(9):5736–5747, 2022.

Qianfeng Yang, Qiyuan Guan, Xiang Chen, Jiyu Jin, Guiyue Jin, and Jiangxin Dong. UniRain: Unified image deraining with RAG-based dataset distillation and multi-objective reweighted optimization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12428–12437, 2026.

Wenhan Yang, Robby T. Tan, Jiashi Feng, Jiaying Liu, Zongming Guo, and Shuicheng Yan. Deep joint rain detection and removal from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 1685–1694, 2017.

Rajeev Yasarla, Vishwanath A. Sindagi, and Vishal M. Patel. Syn2Real transfer learning for image deraining using gaussian processes. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2723–2733, 2020.

Yuntong Ye, Yi Chang, Hanyu Zhou, and Luxin Yan. Closing the loop: Joint rain generation and removal via disentangled image translation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2053–2062, 2021.

Yuntong Ye, Changfeng Yu, Yi Chang, Lin Zhu, Xi-Le Zhao, Luxin Yan, and Yonghong Tian. Unsupervised deraining: Where contrastive learning meets self-similarity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5811–5820, 2022.

Zhiyuan You, Xin Cai, Jinjin Gu, Tianfan Xue, and Chao Dong. Teaching large language models to regress accurate image quality scores using score distribution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14483–14494, 2025.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 14821–14831, 2021.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 5728–5739, 2022.

Fan Zhang, Shaodi You, Yu Li, and Ying Fu. GTAV-NightRain: Photometric realistic large-scale dataset for night-time rain streak removal. *arXiv preprint arXiv:2210.04708*, 2022.

Fan Zhang, Shaodi You, Yu Li, and Ying Fu. Learning rain location prior for nighttime deraining. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13148–13157, 2023.

Fan Zhang, Shaodi You, Yu Li, and Ying Fu. Learning rain location prior for nighttime deraining and beyond. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(10):9169–9186, 2025.

He Zhang and Vishal M. Patel. Density-aware single image de-raining using a multi-stream dense network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 695–704, 2018.

Quan Zhang, Xiaoyu Liu, Wei Li, Hanting Chen, Junchao Liu, Jie Hu, Zhiwei Xiong, Chun Yuan, and Yunhe Wang. Distilling semantic priors from SAM to efficient image restoration models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 25409–25419, 2024.

Ying Zhang, Tao Xiang, Timothy M. Hospedales, and Huchuan Lu. Deep mutual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4320–4328, 2018.

Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2242–2251, 2017.

APPENDIX

A ADDITIONAL IMPLEMENTATION DETAILS AND UPDATE SCHEDULE

Training follows Figure 1(b). Stage I trains CSUD (Dong et al., 2025) to obtain a restorer with an initial ability to remove rain. Stage II freezes that restorer as the teacher and optimizes the student, rain generator, and discriminator with the three paths in Figure 2. Stage III performs ISCD through two refinement cycles. Each cycle initializes the teacher and trainable triplet from the selected preceding checkpoints, freezes the teacher, and updates the student, generator, and discriminator. The selected collaborative tuple initializes the next cycle.

The teacher and student restorers use NAFNet (Chen et al., 2022a) with width 32, encoder blocks $[2, 2, 4, 8]$, decoder blocks $[2, 2, 2, 2]$, and 12 middle blocks. The generator uses three encoders with 128 channels, followed by gated affine modulation, six residual blocks, and a decoder for rain residuals. The discriminator adopts the PatchGAN architecture of CycleGAN (Zhu et al., 2017). Optimization uses Adam (Kingma & Ba, 2015).

Table 5 makes explicit which checkpoints are loaded and which networks remain trainable. Here "scratch" means random initialization, and "previous" denotes the selected tuple from the immediately preceding stage or refinement cycle.

Table 5: Initialization and update schedule. Stage II implements the complete collaborative graph. Stage III performs ISCD through checkpoint propagation and teacher replacement.

Phase	Teacher $\bar{F}$	Student F	Generator G	Discriminator D
Stage I	–	train CSUD	train CSUD	train CSUD
Stage II init.	load F from Stage I	scratch	scratch	scratch
Stage II update	frozen	train	train	train
Stage III, cycle q init.	load previous F	previous F	previous S	previous D
Stage III, cycle q update	frozen	update	update	update

Each cycle in Stage III loads student, generator, and discriminator checkpoints from the preceding tuple. The auxiliary teacher is initialized from the same restorer checkpoint as the student and remains frozen while all three trainable networks are optimized. Cycle 1 uses 100k iterations. Cycle 2 repeats the same graph for 50k iterations on every dataset.

Iterative update procedure. Algorithm 1 specifies how checkpoints propagate across the three stages. Stage I provides the initial restoration ability, Stage II learns the complete collaborative graph, and Stage III repeats the same graph through two refinement cycles. The selected student initializes both the next student and frozen teacher, while the generator and discriminator continue from their corresponding checkpoints. Teacher replacement occurs only at cycle boundaries, which keeps the supervision fixed during each optimization cycle.

Algorithm 1 Iterative Self-Collaborative Distillation with frozen teachers

Require: Unpaired sets $(\mathcal{R}, \mathcal{C})$ and two cycles in Stage III
1: Train CSUD and retain the restorer $\theta^{(1)}$ from Stage I
2: $\bar{\theta} \leftarrow \theta^{(1)}$ and randomly initialize (θ, ϕ, ψ) for student, generator, and discriminator
3: Freeze $\bar{\theta}$ and optimize (θ, ϕ, ψ) with the collaborative graph and save $\Theta^{(2)} = (\theta^{(2)}, \phi^{(2)}, \psi^{(2)})$
4: Set $\Theta^{(3,0)} \leftarrow \Theta^{(2)}$
5: **for** $q = 1, 2$ **do**
6: Initialize $(\theta, \phi, \psi) \leftarrow \Theta^{(3,q-1)}$ and set $\bar{\theta} \leftarrow \theta$
7: Freeze $\bar{\theta}$ and update only (θ, ϕ, ψ) with the unchanged graph from Stage II
8: Select the refined tuple $\Theta^{(3,q)} \leftarrow (\theta, \phi, \psi)$
9: **end for**
10: **return** student restorer $\theta^{(3,2)}$

A.1 Sensitivity to RCR Parameters and ISCD Cycles

The weight map in Equation (7) is controlled by the amplification factor α, response slope ω, and activation threshold κ. Table 6 varies one parameter at a time while fixing the other two to the default configuration $(2, 2, 1)$. The tested ranges cover weaker and stronger amplification, smoother and sharper transitions, and lower and higher response thresholds. This protocol isolates parameter sensitivity without changing the RCR formulation or the remaining Stage II objectives.

Table 6: Sensitivity to the RCR weight map parameters. Each nondefault row changes one parameter from $(\alpha, \omega, \kappa) = (2, 2, 1)$.

Setting	α	ω	κ	Rain100L P/S	RealRain1K-L P/S	Night-Rain P/S
$\alpha = 1$	1	2	1	34.50/0.961	34.19/0.964	30.15/0.898
Default	2	2	1	**34.84/0.964**	**34.44/0.965**	**30.20/0.899**
$\alpha = 3$	3	2	1	34.52/0.962	34.26/0.963	30.18/0.899
$\omega = 1$	2	1	1	34.61/0.963	34.32/0.963	29.88/0.896
$\omega = 3$	2	3	1	34.36/0.959	34.22/0.961	30.09/0.897
$\kappa = 0.5$	2	2	0.5	34.41/0.960	34.20/0.960	29.90/0.890
$\kappa = 1.5$	2	2	1.5	34.28/0.958	34.15/0.959	29.65/0.888

The default setting gives the strongest joint performance on all three benchmarks. Changing α from 2 to 1 or 3 reduces PSNR by 0.32–0.34 dB on Rain100L and by 0.18–0.25 dB on RealRain1K-L, indicating that either insufficient or excessive amplification weakens the spatial guidance. The slope also requires a moderate value. Setting ω to 1 or 3 causes consistent drops, with the largest decrease reaching 0.48 dB and 0.005 SSIM on Rain100L. Performance is especially sensitive to the activation threshold. Increasing κ to 1.5 reduces Night-Rain performance by 0.55 dB and 0.011 SSIM, while lowering κ to 0.5 also produces a clear decline. These results support $(\alpha, \omega, \kappa) = (2, 2, 1)$ as a balanced configuration across synthetic, real daytime, and nighttime rain.

After the complete collaborative graph is learned in Stage II, Stage III can repeat self-collaborative distillation through multiple teacher promotion cycles. Table 7 compares zero to three ISCD cycles while keeping the graph and objectives unchanged. Zero cycles denotes the Stage II model before teacher promotion. This comparison isolates whether repeated knowledge evolution provides gains beyond a single self-collaborative transfer.

Table 7: Effect of the number of Stage III self-collaborative distillation cycles. Zero cycles denotes the complete Stage II model.

ISCD cycles	Rain100L P/S	RealRain1K-L P/S	Night-Rain P/S
0 (Stage II)	34.84/0.964	34.44/0.965	30.20/0.899
1	35.40/0.967	35.05/0.969	30.31/0.902
2	35.52/0.968	35.09/0.970	30.39/0.902
3	35.50/0.967	35.11/0.971	30.35/0.901

The first ISCD cycle improves PSNR over Stage II by 0.56, 0.61, and 0.11 dB on Rain100L, RealRain1K-L, and Night-Rain, respectively, together with consistent SSIM gains. A second cycle further raises PSNR by 0.12, 0.04, and 0.08 dB and preserves or improves SSIM on all three datasets. The benefit then saturates. A third cycle adds only 0.02 dB and 0.001 SSIM on RealRain1K-L, while Rain100L drops by 0.02 dB and 0.001 SSIM and Night-Rain drops by 0.04 dB and 0.001 SSIM. Repeating self-collaboration too many times therefore provides limited additional benefit and may propagate accumulated teacher bias. We consequently use two ISCD cycles as the final setting, which balances restoration accuracy and training cost while avoiding the degradation observed after further refinement.

B Training and Inference Cost

At inference, only F_θ is retained, so the teacher, generator, and discriminator add no parameters or FLOPs relative to the NAFNet student. Training requires two teacher passes without gradients for each rainy crop.

C Discussion and Limitations

Our framework and RGSUD address different learning questions. RGSUD searches the training history for clean targets of high quality. The proposed method instead decomposes the output of a frozen derainer into clean responses and rain residuals, then applies both signals during the current optimization. ISCD does not maintain an online historical reward pool. The teacher remains fixed throughout a cycle and is replaced only at the cycle boundary, where the complete collaborative tuple also initializes the next cycle. Therefore, the proposed approach evolves internally generated clarity, rain, and objective knowledge rather than recycling rewards selected by an external assessor. Combining historical selection with complementary knowledge transfer would require new experiments and lies outside the present claims.

The method has four limitations. First, the frozen teacher must provide a reasonable separation of rain and background. A biased checkpoint from Stage I could propagate errors through all three paths. Second, $|r - \bar{c}|$ may contain motion, illumination differences, or scene texture with high frequencies. The distilled residual is therefore a rain prior rather than a ground truth mask. Third, two teacher evaluations increase training cost in Stages II and III, although inference is unchanged. Fourth, Stage III increases the total training budget. The incremental refinement benefit must be separated from gains due to the complete method in Stage II.

D Implemented Training Graph

This appendix details the training graph in Stage II, which Stage III reuses without architectural or objective changes. Here r denotes a real rainy image, c an unpaired clean image, $\mathbf{z} = (z_1, z_2)$ the detached residual pair from two views, and $z = \Psi(\mathbf{z})$ the corresponding fused residual representation inside the generator. The student derainer is F_θ, and the rain generator is G_ϕ. The graph follows a fixed sequence. The teacher first produces $\mathbf{z}$, the generator transfers rain to c, and the student restores real and synthetic rain. Finally, restored content is reinserted into the generator to close the unpaired loop.

D.1 Data paths

Forward variables. The graph has two anchors. On real rain, $\hat{c}_r = F_\theta(r)$ estimates clean content. On the synthetic path, $\tilde{r}_c = G_\phi(c, r, \mathbf{z})$ transfers rain structure derived from the teacher and reference appearance onto c. Restoring $\tilde{r}_c$ gives $\hat{c}_c = F_\theta(\tilde{r}_c)$, with c as an exact target. These clean estimates then enter three paths that close the cycle:

$$\hat{c}_r = F_\theta(r), \qquad \tilde{r}_c = G_\phi(c, r, \mathbf{z}), \qquad \hat{c}_c = F_\theta(\tilde{r}_c), \tag{10}$$

$$\tilde{r}_r = G_\phi(\hat{c}_r, r, \mathbf{z}), \qquad \tilde{r}_{r\leftarrow c} = G_\phi(\hat{c}_r, \tilde{r}_c, \mathbf{z}), \qquad \tilde{r}_{c\leftarrow c} = G_\phi(\hat{c}_c, \tilde{r}_c, \mathbf{z}). \tag{11}$$

The first row defines restoration of real rain, synthesis onto known clean content, and restoration from a pseudo pair. The second row closes three synthesis paths. $\tilde{r}_r$ reconstructs observed rain from restored content of the real image. $\tilde{r}_{r\leftarrow c}$ preserves that content estimate but uses $\tilde{r}_c$ as the rainy reference. $\tilde{r}_{c\leftarrow c}$ synthesizes rain again from restored synthetic content. Arrowed subscripts distinguish these recombinations. All branches use the same residual pair $\mathbf{z}$ and therefore the same fused spatial rain coordinate z.

Cycle closure guided by rain responses. Each image synthesized again has a natural target, giving the following set of reconstruction pairs:

$$\mathcal{P} = \{(\tilde{r}_r, r),\ (\tilde{r}_{r\leftarrow c}, r),\ (\tilde{r}_c, \tilde{r}_{c\leftarrow c})\}, \qquad \overline{\mathcal{L}}_{\mathrm{RSF}} = \frac{1}{3} \sum_{(u,v)\in\mathcal{P}} \mathcal{L}_{\mathrm{RSF}}(u, v). \tag{12}$$

Here (u, v) denotes a pair of prediction and target. $(\tilde{r}_r, r)$ requires rain removal followed by synthesis to recover the real observation. $(\tilde{r}_{r\leftarrow c}, r)$ tests whether replacing the rainy reference preserves the reconstructed scene. $(\tilde{r}_c, \tilde{r}_{c\leftarrow c})$ closes the synthetic loop and limits drift under successive restoration and synthesis. The factor $1/3$ gives each closure path equal importance. Equation (8) applies the weight $W(z_{\text{avg}})$ derived from the teacher within every pair, emphasizing errors in rainy regions without discarding background supervision.

Self reconstruction with a zero prior. Cycle closure alone cannot ensure that the generator ignores scene information in the rainy reference. We therefore set the two residual inputs to zero to request no rain structure. The synthesized output should then remain c, whether the reference is clean content, restored real rain, or restored synthetic rain:

$$b_0 = G_\phi(c, c, \mathbf{0}), \qquad b_1 = G_\phi(c, \hat{c}_r, \mathbf{0}), \qquad b_2 = G_\phi(c, \hat{c}_c, \mathbf{0}). \tag{13}$$

The three outputs isolate different failure modes. b_0 is the direct identity case and tests whether the generator can leave clean content unchanged. b_1 replaces the reference with $\hat{c}_r$ and therefore penalizes leakage of content from the real scene through the appearance stream. b_2 uses $\hat{c}_c$ and checks the same property after restoration from a pseudo pair. They enter the implemented update groups as

$$\mathcal{L}^F_{\text{self}} = \|b_2 - c\|_1 + \|b_1 - c\|_1, \tag{14}$$

$$\mathcal{L}^S_{\text{self}} = \|b_0 - c\|_1 + \|b_1 - c\|_1. \tag{15}$$

The shared b_1 term couples real content estimated by the student to synthesis with a zero prior. The b_2 term detects drift on the synthetic restoration path, whereas b_0 provides a direct identity anchor. These branches test the factorization of structure and appearance. Without a request for rain structure, neither the appearance reference nor intermediate restoration should alter c.

D.2 Complete optimization objectives

Restoration from pseudo pairs. Synthesis from known clean content c makes $(\tilde{r}_c, c)$ a supervised reconstruction pair despite the unpaired source datasets. Pixel and structural terms constrain the student restoration $\hat{c}_c$:

$$\mathcal{L}_{\text{pair}} = \|\hat{c}_c - c\|_1 + \ell_{\text{SSIM}}(\hat{c}_c, c). \tag{16}$$

The ℓ_1 term constrains color and local intensity, while ℓ_{SSIM} discourages structural distortion. This exact target from a pseudo pair complements $\mathcal{L}_{\text{clean}}$ in Equation (3). The first supervises synthesized rain with known clean content, and the second distills two teacher responses on real rain.

Composition by network. Let $\mathcal{Q} = \{\tilde{r}_c, \tilde{r}_r, \tilde{r}_{r\leftarrow c}, \tilde{r}_{c\leftarrow c}\}$ collect the four synthesized rainy outputs and let $\mathcal{L}_{\text{adv}} = \sum_{x\in\mathcal{Q}} \ell_{\text{GAN}}(D(x), 1)$ denote the adversarial objective for the generator, which encourages each output to lie in the real rainy domain (Goodfellow et al., 2014). The complete student and generator objectives used in both Stages II and III are

$$\mathcal{L}_F = \mathcal{L}_{\text{pair}} + \mathcal{L}_{\text{clean}} + 0.05\overline{\mathcal{L}}_{\text{RSF}} + 0.5\mathcal{L}^F_{\text{self}}, \tag{17}$$

$$\mathcal{L}_S = 0.5\overline{\mathcal{L}}_{\text{RSF}} + 5\mathcal{L}^S_{\text{self}} + 10\mathcal{L}_{\text{RCC}} + \mathcal{L}_{\text{adv}}. \tag{18}$$

The student objective combines four sources of supervision. $\mathcal{L}_{\text{pair}}$ trains removal of synthesized rain against an exact clean target. $\mathcal{L}_{\text{clean}}$ transfers two restorations of real rain from the frozen teacher. $\overline{\mathcal{L}}_{\text{RSF}}$ weakly couples the student to cycle closure, while $\mathcal{L}^F_{\text{self}}$ limits drift under the zero prior. The generator receives stronger closure and identity weights. $\mathcal{L}_{\text{RCC}}$ additionally encourages an approximately achromatic transferred residual in likely rain regions. The implemented coefficients $(0.05, 0.5, 5, 10)$ remain fixed during the refinement in Stage III.

The PatchGAN discriminator receives input sensitive to rain without learned preprocessing. For image x, the input is the concatenation across channels $[x,\ x - G_\sigma(x)]$, where G_σ is a fixed Gaussian smoother. The second component exposes streak evidence at high frequencies alongside image appearance. During discriminator updates, D distinguishes real rainy images r from detached synthesis $\tilde{r}_c$. During joint updates of the student and generator, D remains frozen and provides $\mathcal{L}_{\text{adv}}$. Teacher evaluation uses `no_grad`, and clean predictions and residuals are detached. Teacher parameters therefore remain fixed within a cycle while gradients update the student and generator. The same teacher supports restoration targets, degradation coordinates, and objective allocation without receiving gradient updates.

E Design Properties

Three design properties support the multiple roles of the teacher signal. First, involutions that preserve the image grid align the transformed response in Equation (1) without interpolation. Equation (3) constrains the student against both views rather than forcing agreement with one potentially biased response.

Second, the fixed fusion in Equation (2) retains strong residual responses. For nonnegative responses at every pixel,

$$\max(z_1, z_2) \leq \Psi(z_1, z_2) \leq \sqrt{2}\max(z_1, z_2). \tag{19}$$

Thus a rain response visible in only one view is retained, while the amplification remains bounded. This property motivates the fusion ablation in Table 4.

Third, spatial weighting is bounded and feature modulation preserves the identity case. Equation (7) guarantees $1 < W(z_{\text{avg}}) < 1 + \alpha$, so residual weighting cannot eliminate background supervision or grow without bound. In Equation (5), $m = 0$ gives $\tilde{f}_c = f_c$ exactly. Modulation therefore acts only where the structure stream activates the gate. These properties motivate controlled tests but do not establish restoration accuracy by themselves.

F Results without Reference Metrics on Paired Datasets

Table 8 complements the unpaired evaluation on real rain in Table 2 with results on Rain100L and DID-Data. These datasets provide aligned clean targets, but the same four metrics without references are applied to the restored images so that perceptual behavior can be compared under one protocol. Ground truth images are listed only as diagnostic references and are excluded from method rankings because learned perceptual metrics need not assign them the optimal score.

Table 8: Perceptual quality without references on the paired Rain100L and DID-Data datasets. Higher is better for all metrics. Among methods trained without paired supervision, the best and second best results are highlighted in light purple and light blue, respectively. GT is included as a diagnostic reference.

Setting	Dataset	Metric	*Supervised*	*Unpaired*					*Reference*
			NeRD-Rain-S	Derain CycleGAN	NLCL	DCD-GAN	CSUD	**ESCD**	GT
Paired	Rain100L	CLIP-IQA↑	0.618	0.591	0.539	0.547	**0.600**	0.594	0.650
		Q-Align↑	3.884	3.742	3.474	3.790	3.811	**3.921**	3.956
		DeQA-Score↑	3.561	3.379	2.920	3.539	3.461	**3.712**	3.655
		NIMA↑	0.648	0.624	0.616	0.635	0.651	**0.653**	0.653
	DID-Data	CLIP-IQA↑	0.493	0.422	0.400	0.355	0.418	**0.448**	0.505
		Q-Align↑	3.851	**3.392**	2.606	2.458	3.304	3.381	3.728
		DeQA-Score↑	3.724	3.188	2.476	2.520	3.261	**3.334**	3.697
		NIMA↑	0.647	**0.573**	0.560	0.502	0.558	0.564	0.629

G Visual Results on Unpaired Real Rain

Figure 5 extends the qualitative comparison in Figure 4 to real rainy images without aligned clean references. Relative to the competing methods, the proposed model suppresses more visible streaks and rain veils while retaining natural scene appearance.



Figure 5: Visual comparisons on representative unpaired real rainy images. No aligned clean references are available. Red arrows mark residual rain or appearance distortions in competing outputs. The proposed method produces cleaner backgrounds while maintaining natural scene appearance.

H Conceptual Comparison with Prior Unpaired Methods

Table 9 clarifies how the proposed framework differs from the two methods that define the immediate training context. CSUD supplies the scaffold for pseudo pair synthesis and self reconstruction, but CSUD does not transfer clean responses from a frozen restorer or use teacher residuals to allocate loss emphasis. RGSUD introduces an external vision language IQA model to select and reuse historical restorations of high quality. By contrast, our method queries one frozen derainer under two equivariant views and reuses the resulting clean predictions and rain residuals as restoration, degradation, and objective signals. These signals are refreshed only at stage boundaries, and all teacher components are removed at inference. The table summarizes supervision pathways rather than empirical superiority. Quantitative comparisons are reported in Section 4.

Table 9: Conceptual comparison with CSUD and RGSUD in terms of supervision source, signal use, refresh strategy, and inference dependency.

Property	CSUD	RGSUD	ESCD
Core signal	channel prior	IQA reward	dual teacher views
Historical selection	no	yes	no
External VLM/IQA	no	yes	no
Clean response transfer	no	selected reward	distillation from two views
Explicit rain structure	raw guidance	degradation module	teacher residuals
Spatial loss allocation	no	reward alignment	weights from residuals
New $F/G/D$ in Stage II	–	–	trained from scratch
Reusable signal refresh	none	historical pool	stage boundaries
Teacher at test time	no	no	no