LVP-DIFFIR: DISTILLING LARGE-MODEL TRAJECTORIES FOR INFRARED IMAGE RESTORATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Infrared image enhancement faces a persistent trade-off: compact networks trained on synthetic pairs are efficient but overfit the simulator, while large generative models are accurate yet computationally expensive at inference. We propose **Large-model Visual Prior Diffusion Image Restoration (LVP-DiffIR)**, a large-to-small framework that distills an offline large-model iterative enhancement trajectory into a compact five-step student. The student copies both the first restore, which removes the dominant degradations, and the later residual refinement of edges and contrast. To save large-model inference time during distillation, we adapt the student with **Median-Ratio Adaptation (MRA)** on unlabeled real infrared: an IQA memory bank and a per-source median-ratio backward gate, built on an EMA teacher–student pair. The adapted student further outperforms pure trajectory distillation. We further introduce **InfraRS**, to our knowledge the highest-resolution public remote sensing dataset for infrared image enhancement. Distilled from offline RAR trajectories and then adapted with MRA on unlabeled real infrared, LVP-DiffIR reaches state-of-the-art no-reference perceptual scores on held-out real images, including InfraRS, without paired real-world ground truth on those captures. Our code will be released soon.

1 INTRODUCTION

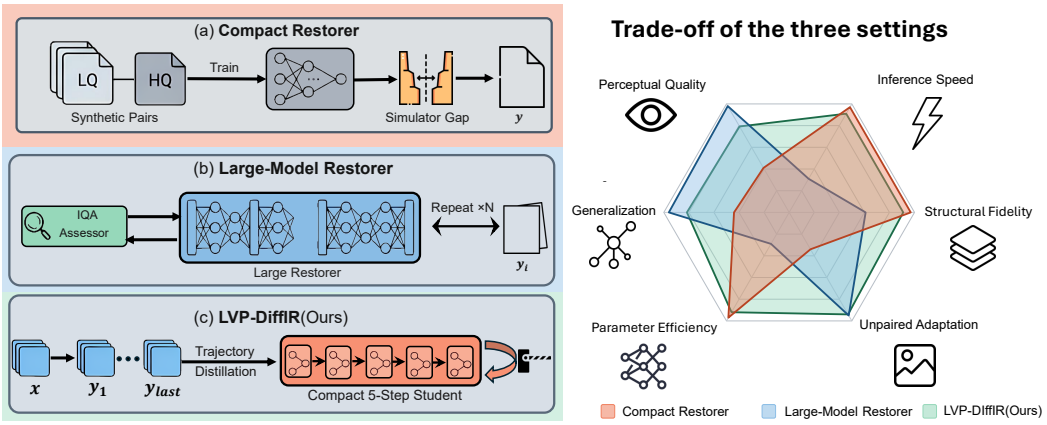

Figure 1: Three settings for infrared enhancement. (a) A compact restorer on synthetic pairs is efficient, but the simulator gap leaves artifacts on real infrared. (b) Large-model RAR is accurate, yet too costly at test time. (c) LVP-DiffIR distills the trajectory into a five-step student and adapts it with MRA, so inference uses the student only. The radar compares the three settings.

Infrared images are indispensable for mission-critical applications Zhao et al. (2023; 2024b) such as night surveillance, autonomous driving Zhao et al. (2024a); Liu et al. (2025b), military target detection Liu et al. (2025c; 2024), and industrial diagnostics Osornio-Rios et al. (2019). Nevertheless, due to the complex imaging process and susceptibility to environmental factors, raw infrared images

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.

often suffer from more severe degradations than visible images, which hinders object detection and recognition Huang et al. (2025); Downs & Vandervelde (2013). Hence, effective enhancement is essential to bridge this quality gap and make infrared imagery reliable for downstream perception.

A considerable number of deep learning-based image enhancement methods Ding et al. (2022); Liu et al. (2025d) have been proposed for infrared images. Early deep models are typically compact end-to-end networks trained on paired data. They are efficient, yet overfit to synthetic degradations and generalize poorly to real infrared noise and contrast shifts Chang et al. (2019); Kuang et al. (2019). Large generative restorers later enable strong zero-shot enhancement, but inference is costly, and unconstrained multi-step generation readily hallucinates structures that are inconsistent with the observation Jiang et al. (2024); Chen et al. (2025; 2026a).

Fig. 1 sets these choices side by side. To close the gap between a strong but computationally costly large-model loop and cheap but weak small models, we combine trajectory distillation Hinton et al. (2015); Salimans & Ho (2022) with **Median-Ratio Adaptation (MRA)**: given RAR's accepted trajectory $(x, y_1, y_2, \ldots, y_{\text{last}})$, we distill the triplet $(x, y_1, y_{\text{last}})$ into a compact student. The distilled triplet provides the Stage-1 targets, so the student inherits the teacher's restore prior without running the large loop at test time. MRA then adapts that student on unlabeled real infrared with an EMA teacher–student pair Tarvainen & Valpola (2017): as in the usual EMA pseudo-label, the student output enters the memory bank only if both quality scores improve upon the image already stored there, while the median-ratio gate decides whether a backward pass is taken. The same adaptation also applies a **Dual Canny** regularizer, which compares a sensitive edge map with an insensitive one and penalizes their gap on structures already stored in the memory bank, keeping restored texture consistent with the strong infrared structure held in that bank.

Based on the above analysis, the main contributions of this paper are summarized as follows:

- We propose **Large-model Visual Prior Diffusion Image Restoration (LVP-DiffIR)**, a large-to-small framework that transfers a large restorer's trajectory into a compact student by trajectory distillation, then adapts it with **MRA**, an IQA bank of the student output and a median-ratio gate, together with a **Dual Canny** regularizer that keeps restored texture on strong infrared structure already stored in the memory bank.
- We introduce **InfraRS**, a high-resolution infrared remote sensing dataset that, to our knowledge, is currently the highest-resolution public benchmark available for infrared image enhancement of high-resolution satellite imagery reported to date.
- We show that LVP-DiffIR, distilled from offline RAR trajectories and adapted with MRA on unlabeled real infrared, reaches state-of-the-art no-reference scores on held-out real images, including InfraRS, without paired real-world ground truth on those captures.

2 Related Work

2.1 Infrared Image Enhancement

Learning-based infrared enhancement has largely replaced hand-crafted filters Shanmugavadivu & Balasubramanian (2014); Lv et al. (2022) with compact CNNs trained on paired data Shao et al. (2021); Chang et al. (2019); Kuang et al. (2019); Ding et al. (2022); Liu et al. (2025d). Two recent methods target coupled degradations rather than a single corruption type. PPFN Liu et al. (2025a) fuses degradation- and scene-type prompts into a Restormer Zamir et al. (2022) backbone and removes stripe, blur, and contrast in reverse physical order via selective progressive training. It also introduces the HM-TIR paired-multi benchmark used here. SEGD Li et al. (2026) uses structural entropy to decouple the mixture and restores each factor with an evidence prior. Both are still trained end-to-end on synthetic degradation datasets, so they overfit the simulator and generalize poorly to real infrared images, as well as to other infrared regimes such as remote sensing. LVP-DiffIR instead distills a large-model trajectory rather than learning another end-to-end mapping from synthetic pairs in the same simulator-trained regime as those paired restorers.

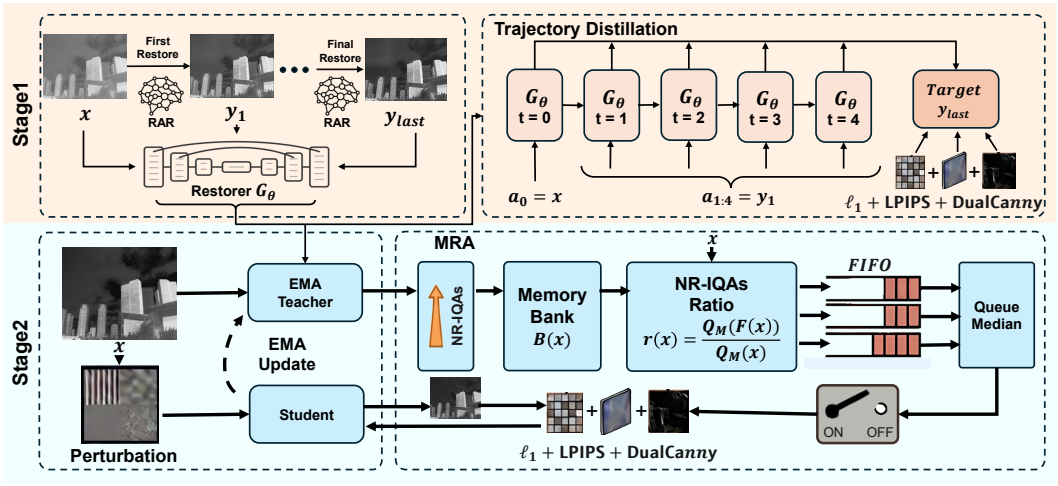

Figure 2: Overview of LVP-DiffIR. **Stage 1**, at the top, runs offline RAR to obtain the accepted trajectory $(x, y_1, y_{\text{last}})$ and distills a five-step student with centers $a_0{=}x$, $a_{1:4}{=}y_1$ and target y_{last}. **Stage 2 MRA**, at the bottom, adapts the distilled student on unlabeled infrared with an EMA teacher–student pair. The student output may update memory bank $B(x)$ by NR-IQA, and a per-source median-ratio gate decides whether a backward pass is taken.

2.2 Large-Model Image Restoration

Generative restorers on large diffusion backbones Xia et al. (2023); Lin et al. (2024); Yu et al. (2024) offer strong zero-shot perceptual quality, but unconstrained multi-step sampling can drift from the observation and is expensive. AutoDIR Jiang et al. (2024) names unknown degradations with CLIP-based Radford et al. (2021) BIQA and restores them with structural-corrected latent diffusion. UniRestore Chen et al. (2025) adapts a Stable Diffusion Rombach et al. (2022) prior for both perceptual quality and task-oriented restoration. RAR Chen et al. (2026a) iterates a large restorer with a quality assessor for several rounds at test time, which is accurate but expensive. LVP-DiffIR uses RAR only as an offline trajectory source and distills $(x, y_1, y_{\text{last}})$ into a five-step student, so inference no longer runs Restore–Assess–Repeat.

2.3 Mean Teacher

Mean Teacher Tarvainen & Valpola (2017) is a semi-supervised learning method, forming an EMA teacher and training under perturbation consistency. Restoration variants that adopt it mainly differ in how they select the teacher image. Semi-UIR Huang et al. (2023) stores a MUSIQ-scored Ke et al. (2021) best-ever final prediction as pseudo-GT. QualiTeacher Xiao et al. (2026) ensembles IQA scores and injects quality as a spatial weight and a generation condition. These methods still train on synthetic degradation datasets, even when using Mean Teacher, and thus generalize poorly to real infrared. LVP-DiffIR first distills an offline RAR trajectory into the student and then adapts it with MRA, an IQA bank of the student output and a median-ratio backward gate on an EMA teacher–student pair, which strengthens real infrared generalization without repeating the large-model loop at test time, keeping inference on the compact student alone after adaptation.

3 LVP-DiffIR

3.1 Overview

LVP-DiffIR is a distillation method that combines large-model trajectory distillation with MRA. It transfers the large model's generalization into a compact student, while MRA adapts that student on unlabeled real infrared without further large-model inference, shortening distillation and improving the small model. The method has two stages: Stage 1 distills an offline RAR trajectory into a compact student, and Stage 2 **MRA** adapts that student to images that RAR never inferred, i.e.,

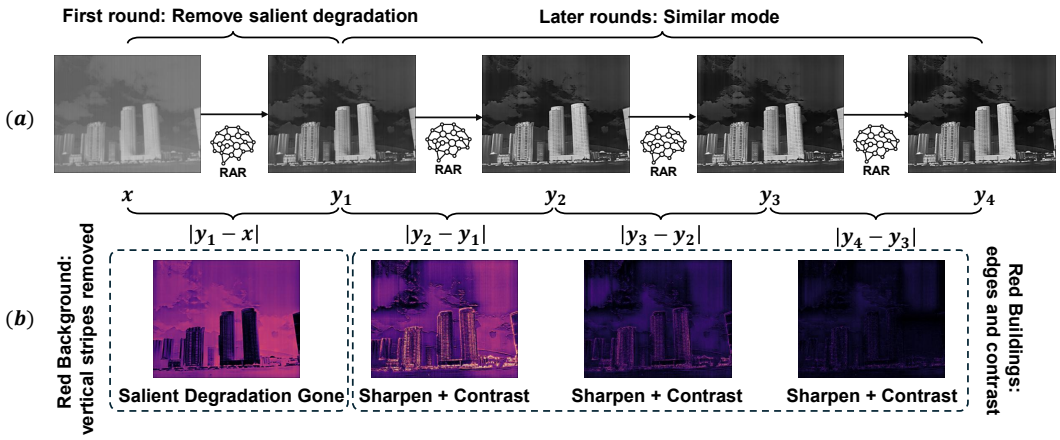

Figure 3: One RAR trajectory. Top: the observation x and later restores y_1 through y_4. Bottom: consecutive residuals. The first residual $y_1 - x$ removes the salient degradation. Later residuals stay in the same mode, sharpening edges and raising contrast. The student therefore centers its first step on x and the remaining steps on y_1, and regresses every step to y_{last}.

observations without an offline RAR trajectory. Fig. 2 summarizes the full pipeline from trajectory distillation through MRA on unlabeled real infrared images.

Stage 1: trajectory distillation. Let x be a degraded infrared observation. An offline RAR teacher Chen et al. (2026a) yields an accepted triplet $(x,\, y_1,\, y_{\text{last}})$, where y_1 is the first-step enhanced output and y_{last} is the final enhanced image. A five-step student G_θ is trained to reproduce this trajectory: the forward process is centered on x then on y_1, and every timestep regresses to y_{last}. This default supervision is called trajectory distillation in the reported forward process.

Stage 2: Median-Ratio Adaptation. MRA then fine-tunes the already-distilled student on unlabeled real infrared that has no RAR trajectory. The large-model prior is kept as a starting point, and MRA only adapts it, so the student can learn from images the large model never inferred, without repeating RAR inference on those images during adaptation.

3.2 Stage 1: Trajectory Distillation

Stage 1 only prepares a compact student that can copy an offline RAR path. For a labeled observation x, RAR Chen et al. (2026a) is run once and we store the accepted triplet $(x,\, y_1,\, y_{\text{last}})$, where y_1 is the first-step enhanced output and y_{last} is the final enhanced image. The mapping $x \mapsto y_1$ already removes the most salient degradations. Later rounds from y_1 to y_{last} largely repeat the same residual pattern, sharpening edges and raising contrast, as shown in Fig. 3. The triplet is therefore enough to form the forward process: we use a $T{=}5$ student with two centers and a DDPM-style Ho et al. (2020) Gaussian perturbation, so the first step, centered on x, learns to strip the dominant degradations, and later steps, centered on y_1, learn this residual RAR refinement of edges and contrast. Let $\{\eta_t\}_{t=0}^{T-1}$ be an exponential schedule and $\sigma_t = \kappa\sqrt{\eta_t}$. For $t \sim \text{Uniform}\{0, \ldots, T-1\}$,

$$a_t \;=\; \begin{cases} x, & t = 0, \\ y_1, & t \geq 1, \end{cases} \qquad x_t \;=\; a_t + \sigma_t\, \varepsilon, \qquad \varepsilon \sim \mathcal{N}(0, I). \tag{1}$$

The student uses the compact Swin-UNet Liu et al. (2021) of UPSR Zhang et al. (2025) as the LQ-conditioned denoiser $G_\theta(x_t, t, x)$ and is trained at every t to regress onto y_{last} with ℓ_1, LPIPS Zhang et al. (2018), and $\mathcal{L}_{\text{DC}}$. The Dual Canny loss $\mathcal{L}_{\text{DC}}$ is introduced below.

3.3 Stage 2: Median-Ratio Adaptation

Stage 1 can only see images for which RAR was run. Running RAR on every unlabeled image would make distillation prohibitively expensive. We therefore run RAR on only a subset, distill that trajectory into the student, and continue with **MRA** on unlabeled real infrared that has no triplet,

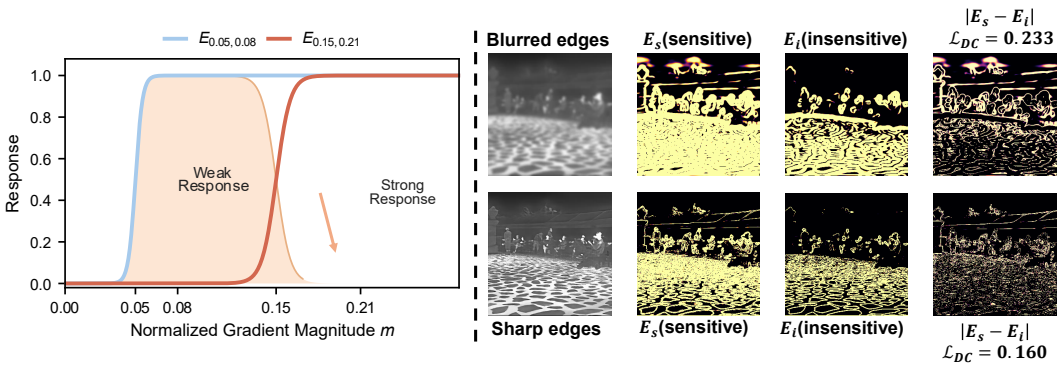

Figure 4: Dual Canny on one infrared scene. Left: the two soft-Canny responses and the gap between them. Right: a blurred observation and a sharp restore. The sensitive map E_s fires on weak texture, while the insensitive map E_i keeps strong structure. Their gap, and thus $\mathcal{L}_{\mathrm{DC}}$, is larger on the blurred example, 0.233, than on the sharp one, 0.160, on this scene.

so there is no further large-model inference and a faster path to adapting the student. MRA builds on an EMA teacher–student pair Tarvainen & Valpola (2017): the EMA copy of G_θ is the teacher, and the student is the same five-step sampler. Their inputs are not the same image: the student is trained on a geometrically and photometrically perturbed LQ, while the EMA teacher always sees the unperturbed observation. Bank scoring and the median-ratio gate likewise use this clean LQ, so the memory bank is always scored on the same unperturbed observation.

Bank initialization. A per-image bank $B(x)$ stores one pseudo-label together with MUSIQ Ke et al. (2021) Q_{M}, where higher is better, and NIQE Mittal et al. (2013) Q_{N}, where lower is better. If both a paired ground truth y_{gt} and the RAR output y_{last} are available, the bank is initialized by comparing them under the same two scores, MUSIQ and NIQE:

$$B(x) \leftarrow \begin{cases} y_{\mathrm{last}}, & Q_{\mathrm{M}}(y_{\mathrm{last}}) > Q_{\mathrm{M}}(y_{\mathrm{gt}}) \ \wedge \ Q_{\mathrm{N}}(y_{\mathrm{last}}) < Q_{\mathrm{N}}(y_{\mathrm{gt}}), \\ y_{\mathrm{gt}}, & \text{otherwise.} \end{cases} \tag{2}$$

Thus y_{last} is written only if both MUSIQ and NIQE improve upon the ground truth. Otherwise the bank starts from y_{gt}. If neither y_{gt} nor y_{last} is available, $B(x)$ is cold-started from Teacher(x), the EMA teacher's output on that unlabeled observation.

Memory bank. On each image x we run the sampler of Eq. 1 and keep the student output $\hat{y}_{T-1}$, the same choice as a standard EMA pseudo-label. That output replaces the stored image iff it jointly dominates it on both IQA scores, with MUSIQ rising and NIQE falling:

$$B(x) \leftarrow \hat{y}_{T-1} \quad \text{iff} \quad Q_{\mathrm{M}}(\hat{y}_{T-1}) > Q_{\mathrm{M}}\big(B(x)\big) \ \wedge \ Q_{\mathrm{N}}(\hat{y}_{T-1}) < Q_{\mathrm{N}}\big(B(x)\big). \tag{3}$$

Unlabeled measurements therefore become training data once this student output is admitted, without running another RAR restore on that image during Stage 2 adaptation.

Median-ratio backward gate. Admission into $B(x)$ is not yet a license to update θ. We form the relative quality

$$r(x) \ = \ \frac{Q_{\mathrm{M}}\big(B(x)\big)}{Q_{\mathrm{M}}(x)}, \tag{4}$$

and keep a per-source FIFO queue $\mathcal{W}_s$ of recent ratios. InfraRS, M3FD Liu et al. (2022), and HM-TIR are tracked separately, because absolute MUSIQ scales differ. After a short warmup, the student receives a gradient if and only if its ratio is at least the queue median:

$$\mathbb{I}_{\mathrm{bp}}(x) \ = \ \mathbf{1}[r(x) \ \geq \ \mathrm{median}(\mathcal{W}_s)]\,. \tag{5}$$

If $\mathbb{I}_{\mathrm{bp}}(x) = 0$, there is no backward pass: the sampler is still evaluated and may update $B(x)$ under Eq. 3, but θ is not changed. Weak pseudo-labels therefore cannot overwrite the Stage-1 trajectory, while a later iteration can still replace the bank if the student output improves on both scores.

Student update. When $\mathbb{I}_{\mathrm{bp}}(x) = 1$, the student uses the same ℓ_1, LPIPS, and $\mathcal{L}_{\mathrm{DC}}$ regression as Stage 1, with target $B(x)$ in place of y_{last}, on $\hat{y} = G_\theta(x_t, t, x)$:

$$\mathcal{L}_{\mathrm{S2}} \;=\; \mathbb{I}_{\mathrm{bp}}(x)\,\Big(\big\|\hat{y} - B(x)\big\|_1 \;+\; \lambda_{\mathrm{p}}\,\mathrm{LPIPS}\big(\hat{y},\, B(x)\big) \;+\; \lambda_{\mathrm{dc}}\,\mathcal{L}_{\mathrm{DC}}\big(\hat{y},\, B(x)\big)\Big). \tag{6}$$

The forward still follows the trajectory in Eq. 1 when the triplet is available. Without a triplet there is no y_1 to follow, so x_t is formed from x, as at test time. A backward pass applies geometric transforms jointly to the student LQ and $B(x)$. Stripe, blur and contrast are applied only to the student LQ, not to the bank target used as the regression GT.

Dual Canny regularizer. Weak texture is cheap for a restorer to invent in infrared, while true structure is easy to lose. We therefore compare two soft-Canny maps of $\hat{y}$: a sensitive pair that fires on weak edges and an insensitive pair that keeps only strong ones, and penalize their gap on structures already present in $B(x)$. Let $g(\cdot)$ be luminance and ∇ a Sobel operator. Write $m = \|\nabla g(\hat{y})\|_2$, min–max normalized per image before the soft threshold, and

$$E_{\ell,h}(\hat{y}) \;=\; \sigma\!\left(\tau\,\frac{m-\ell}{h-\ell}\right), \tag{7}$$

for a differentiable threshold pair (ℓ, h) and temperature τ. The Dual Canny loss is

$$\mathcal{L}_{\mathrm{DC}} \;=\; \mathrm{mean}\Big(\big|E_{\ell_1,h_1}(\hat{y}) - E_{\ell_2,h_2}(\hat{y})\big| \odot w\Big), \qquad w \;=\; \frac{\|\nabla g(B(x))\|_2^{\gamma}}{\mathrm{mean}\big(\|\nabla g(B(x))\|_2^{\gamma}\big)}, \tag{8}$$

with $B(x)$ detached. The low pair $(\ell_1, h_1) = (0.05, 0.08)$ is sensitive, and the high pair $(\ell_2, h_2) = (0.15, 0.21)$ is not. Their absolute difference is large wherever the student adds spurious texture, and w with $\gamma = 0.6$ concentrates the penalty on bank structure. Fig. 4 shows the same gap on real infrared: the blurred example leaves a larger $|E_s - E_i|$ and a higher $\mathcal{L}_{\mathrm{DC}}$ than the sharp one. We set $\lambda_{\mathrm{dc}} = 0.1$ and $\tau = 12$ for every experiment in this paper.

4 The InfraRS Dataset

High-resolution infrared remote sensing data have long been scarce Huang et al. (2025), severely hampering the development and evaluation of advanced enhancement models for real-world satellite imagery. To bridge this gap, we construct a medium-scale, ultra-high-resolution infrared remote sensing dataset, named InfraRS, which comprises 6,731 images of size 512×512 pixels at a spatial resolution of 10m. To the best of our knowledge, this is the highest-resolution infrared remote sensing dataset publicly available to date. Unlike ground-based captures, satellite imaging systems operate in space without any possibility of manual recalibration. Moreover, the complex space environment, including extreme temperature variations, cosmic radiation, and atmospheric disturbances, induces intricate and poorly modeled degradations. Consequently, achieving substantial enhancement on such data is both technically challenging and practically significant. Fig. 5 shows five typical patches. InfraRS includes these land-cover types, but is not limited to them: urban, river, mountain, industrial, and cropland scenes as well. Table 1 sets InfraRS beside other public infrared satellite datasets. Its ground sampling is 10 m, and the other listed sets span 30 m to 1 km.

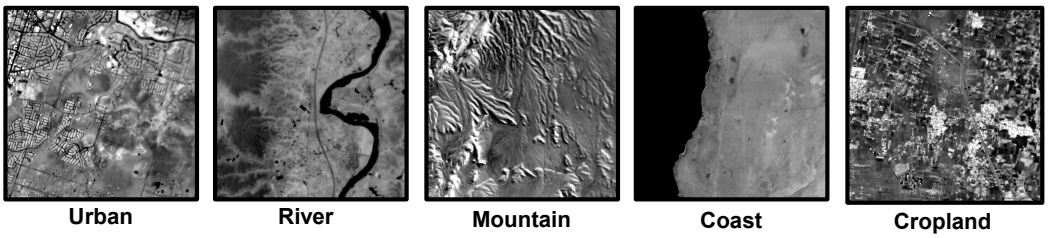

Figure 5: Typical InfraRS examples at 512×512 and 10 m. The dataset includes these land-cover types, but is not limited to them. From left to right: urban, river, mountain, industrial, and cropland scenes drawn from the collected overhead tiles here.

Table 1: Statistics of representative infrared satellite datasets. Columns give the sensor, ground sampling distance, crop size, and image count for each released set.

Dataset	Sensor	Resolution	Patch	# Images
InfraRS	XH-01	10 m	512^2	6,731
ASTER–MODIS Ait-Bachir et al. (2025)	ASTER/MODIS	1 km–90 m	256^2	8,312
EFDiff Chen et al. (2026b)	Landsat 8/9	30 m	256^2	242k
GrokLST Dai et al. (2025)	SDGSAT-1	30 m	512^2	641
STF-LST Bouaziz et al. (2026)	MODIS/Landsat 8	1 km/30 m	950^2	51
SAT Chen et al. (2024)	Landsat-8	100 m	96^2/224^2	15k/182

5 Experiments

5.1 Experimental Setup

We train and test on InfraRS, HM-TIR Liu et al. (2025a), and M3FD Liu et al. (2022). Stage 1 distills 150 RAR triplets, 50 from each source, and Stage 2 adapts on 4,906 unlabeled images: 2,000 InfraRS, 1,704 M3FD, and 1,202 HM-TIR. Held-out tests are InfraRS val300, the 301-image HM-TIR paired-multi split, and M3FD test300. We report NIQE ↓ Mittal et al. (2013), MUSIQ ↑ Ke et al. (2021), MANIQA ↑ Yang et al. (2022), and CLIP-IQA ↑ Wang et al. (2023). Compact baselines are PPFN Liu et al. (2025a), SEGD Li et al. (2026), Restormer Zamir et al. (2022), PromptIR Potlapalli et al. (2023), Mean Teacher Tarvainen & Valpola (2017), and Semi-UIR Huang et al. (2023), and RAR Chen et al. (2026a) is listed only as the offline teacher. Appendix A records the source construction, losses, schedules, inference, and baseline recipes in full.

5.2 Evaluation on Synthetic Paired HM-TIR

Table 2 reports HM-TIR results. LVP-DiffIR is trained with no paired reference, so the comparison uses only no-reference scores and omits PSNR and SSIM Wang et al. (2004). Among compact models, LVP-DiffIR attains the best MUSIQ, MANIQA and CLIP-IQA, for example MUSIQ 49.67 versus 37.15/39.00 for PPFN and SEGD. PPFN and SEGD are trained against HM-TIR HQ, whereas LVP-DiffIR is distilled from y_{last} and is judged here by no-reference quality. RAR still leads most NR scores and is reported as an upper bound for the compact student.

Table 2: Quantitative comparison on HM-TIR Liu et al. (2025a) with the dataset's composite stripe, blur and contrast degradations on 301 images. Training uses no paired reference, so PSNR and SSIM are omitted and only no-reference scores are compared. Best and second-best among compared methods, excluding RAR, are in **bold** and underline, respectively.

Method	NIQE↓	MUSIQ↑	MANIQA↑	CLIP-IQA↑
RAR (CVPR 2026, teacher)	4.17	58.92	0.447	0.454
Mean Teacher (NeurIPS 2017)	7.27	29.48	0.184	0.095
Restormer (CVPR 2022)	6.76	35.29	0.210	0.113
Semi-UIR (CVPR 2023)	6.17	31.11	0.164	0.112
PromptIR (NeurIPS 2023)	**5.88**	38.86	0.230	0.154
PPFN (NeurIPS 2025)	6.65	37.15	0.215	0.139
SEGD (CVPR Findings 2026)	<u>5.95</u>	<u>39.00</u>	<u>0.237</u>	<u>0.204</u>
LVP-DiffIR (Ours)	6.09	**49.67**	**0.296**	**0.305**

Mean Teacher and Semi-UIR, trained on the same mix, stay well below LVP-DiffIR. Semi-UIR reaches MUSIQ 31.11 and Mean Teacher 29.48. The five-step student therefore benefits from the offline RAR trajectory rather than from unlabeled EMA training alone on this mix. Fig. 8 shows the same comparison on two HM-TIR scenes alongside the other compared methods.

5.3 GENERALIZATION TO INFRARS AND M3FD

Table 3 asks whether the student overfits the HM-TIR synthetic composite by testing unpaired InfraRS and M3FD, which have no paired HQ and are scored only with no-reference metrics.

Table 3: Cross-dataset comparison on unpaired infrared tests. InfraRS val300, 300 images, and M3FD test300, 300 images. Best / second-best among compared methods exclude RAR.

	InfraRS				M3FD			
Method	NIQE↓	MUSIQ↑	MANIQA↑	CLIP-IQA↑	NIQE↓	MUSIQ↑	MANIQA↑	CLIP-IQA↑
RAR (CVPR 2026, teacher)	4.37	57.65	0.397	0.555	3.85	55.94	0.419	0.479
Mean Teacher (NeurIPS 2017)	5.96	32.71	0.277	0.328	7.18	28.01	0.236	0.170
Restormer (CVPR 2022)	6.21	41.63	0.230	0.401	6.46	39.33	0.248	0.247
Semi-UIR (CVPR 2023)	7.95	<u>49.54</u>	<u>0.278</u>	**0.477**	6.79	<u>44.46</u>	0.261	0.269
PromptIR (NeurIPS 2023)	6.10	40.29	0.221	0.390	5.99	38.45	0.247	0.231
PPFN (NeurIPS 2025)	<u>5.58</u>	46.75	0.275	0.328	**5.45**	43.31	<u>0.273</u>	0.292
SEGD (CVPR Findings 2026)	6.26	45.39	0.267	0.433	6.43	41.35	0.272	<u>0.379</u>
LVP-DiffIR (Ours)	**4.52**	**54.74**	**0.337**	<u>0.461</u>	<u>5.61</u>	**50.37**	**0.318**	**0.386**

On InfraRS, LVP-DiffIR is best among compact methods on NIQE, MUSIQ and MANIQA, and stays close to the RAR teacher on MUSIQ, 54.74 versus 57.65. Semi-UIR is higher on CLIP-IQA, 0.477 versus 0.461, while PPFN and SEGD remain lower on MUSIQ and MANIQA.

On M3FD, LVP-DiffIR leads compact methods on MUSIQ, MANIQA and CLIP-IQA. PPFN is lower on NIQE, 5.45 versus 5.61, and SEGD stays close on CLIP-IQA, 0.379 versus 0.386. The distilled student transfers to both unpaired sets better than restorers trained only on HM-TIR pairs, while remaining below the RAR teacher on both tests. Fig. 6 visualizes this comparison on one InfraRS scene and one M3FD scene, with the input on the left and the same methods zoomed on the right of both examples, shown together with the scores from Table 3.

5.4 ABLATION STUDIES

We ablate the same 150/4,906 mix and report MUSIQ ↑ / MANIQA ↑ as InfraRS / HM-TIR / M3FD. The reported run is Stage 1 trajectory distillation on 150 RAR triplets, then Stage 2 MRA with Dual Canny, last-step admission, a median gate, and Teacher(x) initialization.

Component ablations. Table 4 drops one choice at a time. Removing Dual Canny, the median gate, or Stage 2 hurts both metrics on every set, as does replacing trajectory distillation by a ResShift-style Yue et al. (2023) residual. Trajectory distillation regresses every step to y_{last}, and the ResShift residual is the control for that default. Without RAR, Stage 1 uses 150 HM-TIR synthetic pairs and MRA runs with identity HQ. HM-TIR MUSIQ can match the default, but InfraRS, M3FD, and MANIQA fall behind the reported student on both metrics across the three evaluation sets.

Table 4: Component ablations on InfraRS / HM-TIR / M3FD. Each row drops one component from the full model. Default: trajectory distillation + Dual Canny + last-step + median gate on all three datasets shown for the reported full model in this component comparison.

Method	MUSIQ↑	MANIQA↑
w/o Dual Canny	52.90 / 48.56 / 49.40	0.313 / 0.279 / 0.308
w/o median gate	53.93 / 48.48 / 49.94	0.330 / 0.294 / 0.316
w/o S2	50.02 / 45.62 / 47.45	0.297 / 0.266 / 0.293
w/o trajectory distillation	51.12 / 47.46 / 47.91	0.314 / 0.280 / 0.307
w/o RAR	51.65 / **50.26** / 47.66	0.307 / 0.282 / 0.290
LVP-DiffIR (full)	**55.03** / 50.17 / **50.78**	**0.337 / 0.299 / 0.323**

Table 5: Hard gate vs. soft IQA weights on InfraRS, HM-TIR, and M3FD.

Stage-2 filter	MUSIQ↑	MANIQA↑
ungated EMA	53.93 / 48.48 / 49.94	0.330 / 0.294 / 0.316
QualiTeacher weight	53.77 / 48.49 / 49.94	0.330 / 0.296 / 0.317
median gate	**55.03 / 50.17 / 50.78**	**0.337 / 0.299 / 0.323**

Table 6: Centers a_t under trajectory distillation, with target y_{last}. Both stages share these centers.

$a_{0:4}$	MUSIQ↑	MANIQA↑
y_1, y_1, y_1, y_1, y_1	52.60 / 47.24 / 47.83	0.322 / 0.279 / 0.309
x, y_1, y_2, y_3, y_4	52.09 / 46.55 / 49.78	0.313 / 0.286 / 0.318
x, y_1, y_1, y_1, y_1	**55.03 / 50.17 / 50.78**	**0.337 / 0.299 / 0.323**

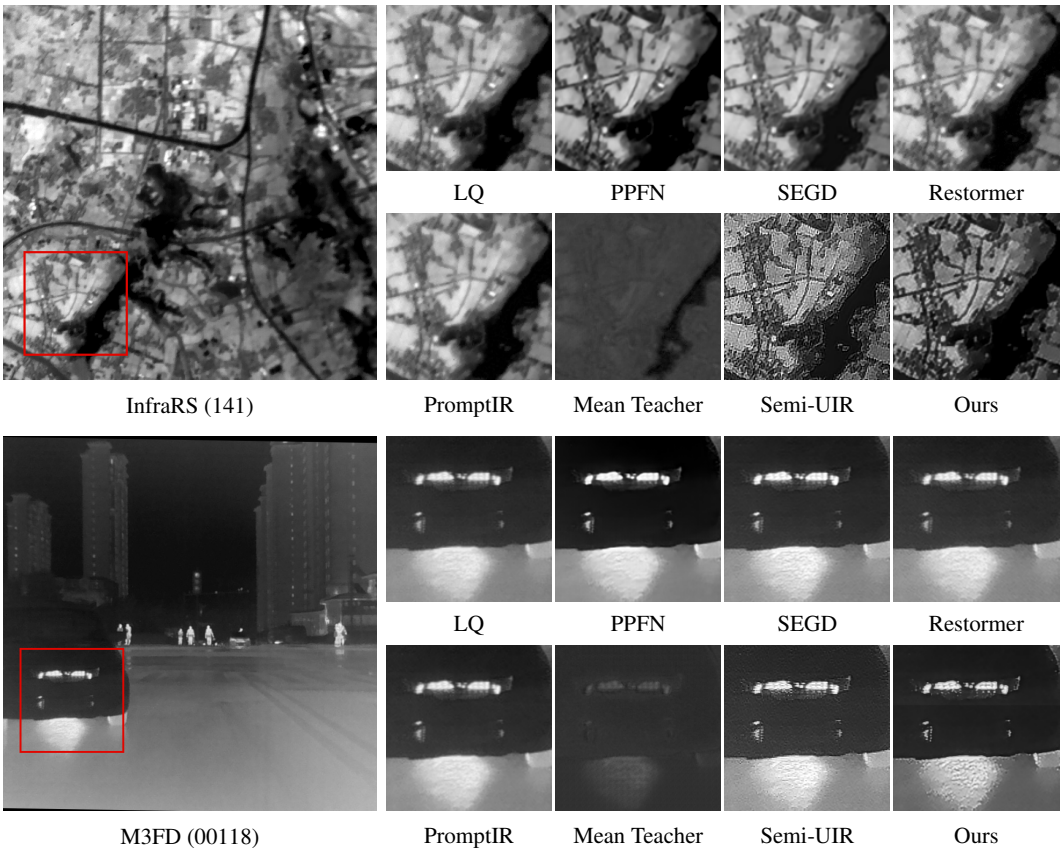

Figure 6: Generalization on InfraRS and M3FD. The input is on the left, and both examples zoom LQ, PPFN, SEGD, Restormer, PromptIR, Mean Teacher, Semi-UIR, and Ours on the right.

Hard gate vs. IQA spatial weights. Table 5 replaces the median-ratio gate with QualiTeacher Xiao et al. (2026) 2×2 spatial weights, an ensemble of MUSIQ, inverted BRISQUE, and CLIP-IQA, while keeping Dual Canny, last-step admission, and Teacher(x) init. The ungated EMA row matches w/o median gate in Table 4, and the soft weights stay near that row, below the hard median-ratio gate across all three evaluation sets.

Forward-process centers. Table 6 changes only the training centers a_t in Eq. 1. Test-time sampling still starts from x. Centering every step on y_1, or blending x, y_1, y_2, y_3, y_4 toward y_{last}, drops all three sets. We keep the two-center schedule, with the first step around the observation and later steps around the first RAR output in the reported run.

6 Conclusion

We presented LVP-DiffIR, a five-step student that distills an offline RAR trajectory and then adapts on unlabeled infrared images with MRA. On the synthetic HM-TIR split it leads compact restorers on MUSIQ, MANIQA, and CLIP-IQA under the same no-reference comparison. The same checkpoint transfers to InfraRS and M3FD, where no paired ground truth is available, and stays close to the RAR teacher without running that loop at test time.

InfraRS adds a 10 m overhead benchmark for this setting. Ablations support the two-center schedule, Dual Canny, and the hard median gate over soft IQA reweighting. The resulting student keeps a large-model restore prior, runs in five DDIM steps, and still improves when the only extra supervision is unlabeled real infrared on the three reported sets.

REFERENCES

Romuald Ait-Bachir, Carlos Granero-Belinchon, Aurélie Michel, Julien Michel, Xavier Briottet, and Lucas Drumetz. Land surface temperature super-resolution with a scale-invariance-free neural approach: Application to MODIS. *JSTARS*, 2025.

Sofiane Bouaziz, Adel Hafiane, Raphaël Canals, and Rachid Nedjai. Deep learning for spatio-temporal fusion in land surface temperature estimation: A comprehensive survey, experimental analysis, and future trends. *Remote Sens.*, 2026.

Yi Chang, Luxin Yan, Li Liu, Houzhang Fang, and Sheng Zhong. Infrared aerothermal nonuniform correction via deep multiscale residual network. *GRSL*, 2019.

I-Hsiang Chen, Wei-Ting Chen, Yu-Wei Liu, Yuan-Chun Chiang, Sy-Yen Kuo, and Ming-Hsuan Yang. UniRestore: Unified perceptual and task-oriented image restoration model using diffusion prior. In *CVPR*, 2025.

I-Hsiang Chen, Isma Hadji, Enrique Sanchez, Adrian Bulat, Sy-Yen Kuo, Radu Timofte, Georgios Tzimiropoulos, and Brais Martinez. Restore, assess, repeat: A unified framework for iterative image restoration. In *CVPR*, 2026a.

Junqi Chen, Lijuan Jia, Jinchuan Zhang, Yilong Feng, Xiaobin Zhao, and Ran Tao. Super-resolution for land surface temperature retrieval images via cross-scale diffusion model using reference images. *Remote Sens.*, 2024.

Yiheng Chen, Zihui Ma, Peishi Jiang, Yilong Dai, Qikai Hu, Xinyue Ye, Lingyao Li, Rita Sousa, and Runlong Yu. When earth foundation models meet diffusion: An application to land surface temperature super-resolution, 2026b, arXiv:2604.16841.

Qun Dai, Chunyang Yuan, Yimian Dai, Yuxuan Li, Xiang Li, Kang Ni, Jianhui Xu, Xiangbo Shu, and Jian Yang. MoCoLSK: Modality-conditioned high-resolution downscaling for land surface temperature. *TGRS*, 2025.

Dan Ding, Ye Li, Peng Zhao, Kaitai Li, Sheng Jiang, and Yanxiu Liu. Single infrared image stripe removal via residual attention network. *Sensors*, 2022.

Chandler Downs and Thomas E. Vandervelde. Progress in infrared photodetectors since 2000. *Sensors*, 2013.

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

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.

Shirui Huang, Keyan Wang, Huan Liu, Jun Chen, and Yunsong Li. Semi-supervised learning for underwater image restoration via a memory bank. In *CVPR*, 2023.

Yongsong Huang, Tomo Miyazaki, Xiaofeng Liu, and Shinichiro Omachi. Infrared image super-resolution: A systematic review and future trends. *JSTARS*, 2025.

Yitong Jiang, Zhaoyang Zhang, Tianfan Xue, and Jinwei Gu. AutoDIR: Automatic all-in-one image restoration with latent diffusion. In *ECCV*, 2024.

Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. MUSIQ: Multi-scale image quality transformer. In *ICCV*, 2021.

Xiaodong Kuang, Xiubao Sui, Yuan Liu, Qian Chen, and Guohua Gu. Single infrared image enhancement using a deep convolutional neural network. *Neurocomputing*, 2019.

Pu Li, Huafeng Li, Yafei Zhang, Yu Liu, and Wen Wang. Breaking degradation coupling: A structural entropy-guided decoupled framework and benchmark for infrared enhancement. In *CVPR Findings*, 2026.

Xinqi Lin, Jingwen He, Ziyan Chen, Zhaoyang Lyu, Bo Dai, Fanghua Yu, Yu Qiao, Wanli Ouyang, and Chao Dong. DiffBIR: Toward blind image restoration with generative diffusion prior. In *ECCV*, 2024.

Jinyuan Liu, Xin Fan, Zhanbo Huang, Guanyao Wu, Risheng Liu, Wei Zhong, and Zhongxuan Luo. Target-aware dual adversarial learning and a multi-scenario multi-modality benchmark to fuse infrared and visible for object detection. In *CVPR*, 2022.

Jinyuan Liu, Zihang Chen, Zhu Liu, Zhiying Jiang, Long Ma, Xin Fan, and Risheng Liu. Enhancing infrared vision: Progressive prompt fusion network and benchmark. In *NeurIPS*, 2025a.

Jinyuan Liu, Guanyao Wu, Zhu Liu, Di Wang, Zhiying Jiang, Long Ma, Wei Zhong, Xin Fan, and Risheng Liu. Infrared and visible image fusion: From data compatibility to task adaption. *TPAMI*, 2025b.

Jinyuan Liu, Bowei Zhang, Qingyun Mei, Xingyuan Li, Yang Zou, Zhiying Jiang, Long Ma, Risheng Liu, and Xin Fan. DCEvo: Discriminative cross-dimensional evolutionary learning for infrared and visible image fusion. In *CVPR*, 2025c.

Risheng Liu, Zhu Liu, Jinyuan Liu, Xin Fan, and Zhongxuan Luo. A task-guided, implicitly-searched and meta-initialized deep model for image fusion. *TPAMI*, 2024.

Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021.

Zhu Liu, Zijun Wang, Jinyuan Liu, Fanqi Meng, Long Ma, and Risheng Liu. DEAL: Data-efficient adversarial learning for high-quality infrared imaging. In *CVPR*, 2025d.

Hui Lv, Pengfei Shan, Hongfang Shi, and Li Zhao. An adaptive bilateral filtering method based on improved convolution kernel used for infrared image enhancement. *SIViP*, 2022.

Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a completely blind image quality analyzer. *SPL*, 2013.

Roque A. Osornio-Rios, Jose A. Antonino-Daviu, and René de Jesús Romero-Troncoso. Recent industrial applications of infrared thermography: A review. *TII*, 2019.

Vaishnav Potlapalli, Syed Waqas Zamir, Salman Khan, and Fahad Shahbaz Khan. PromptIR: Prompting for all-in-one blind image restoration. In *NeurIPS*, 2023.

Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning transferable visual models from natural language supervision. In *ICML*, 2021.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022.

Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In *ICLR*, 2022.

P. Shanmugavadivu and K. Balasubramanian. Particle swarm optimized multi-objective histogram equalization for image enhancement. *Opt. Laser Technol.*, 2014.

Yuyi Shao, Yingwei Sun, Mengmeng Zhao, Yankang Chang, Zhouzhou Zheng, Chengliang Tian, and Yan Zhang. Infrared image stripe noise removing using least squares and gradient domain guided filtering. *Infrared Phys. Technol.*, 2021.

Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consistency targets improve semi-supervised deep learning results. In *NeurIPS*, 2017.

Jianyi Wang, Kelvin C. K. Chan, and Chen Change Loy. Exploring CLIP for assessing the look and feel of images. In *AAAI*, 2023.

Zhou Wang, Alan C. Bovik, Hamid R. Sheikh, and Eero P. Simoncelli. Image quality assessment: From error visibility to structural similarity. *TIP*, 2004.

Bin Xia, Yulun Zhang, Shiyin Wang, Yitong Wang, Xinglong Wu, Yapeng Tian, Wenming Yang, and Luc Van Gool. DiffIR: Efficient diffusion model for image restoration. In *ICCV*, 2023.

Fengyang Xiao, Jingjia Feng, Peng Hu, Yuhan Chen, Dingming Zhang, Lei Xu, Guanyi Qin, Lu Li, Chunming He, and Sina Farsiu. QualiTeacher: Quality-conditioned pseudo-labeling for real-world image restoration. In *ECCV*, 2026.

Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and Yujiu Yang. MANIQA: Multi-dimension attention network for no-reference image quality assessment. In *CVPRW*, 2022.

Fanghua Yu, Jinjin Gu, Zheyuan Li, Jinfan Hu, Xiangtao Kong, Xintao Wang, Jingwen He, Yu Qiao, and Chao Dong. Scaling up to excellence: Practicing model scaling for photo-realistic image restoration in the wild. In *CVPR*, 2024.

Zongsheng Yue, Jianyi Wang, and Chen Change Loy. ResShift: Efficient diffusion model for image super-resolution by residual shifting. In *NeurIPS*, 2023.

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 *CVPR*, 2022.

Leheng Zhang, Weiyi You, Kexuan Shi, and Shuhang Gu. Uncertainty-guided perturbation for image super-resolution diffusion model. In *CVPR*, 2025.

Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *CVPR*, 2018.

Zixiang Zhao, Haowen Bai, Yuanzhi Zhu, Jiangshe Zhang, Shuang Xu, Yulun Zhang, Kai Zhang, Deyu Meng, Radu Timofte, and Luc Van Gool. DDFM: Denoising diffusion model for multi-modality image fusion. In *ICCV*, 2023.

Zixiang Zhao, Haowen Bai, Jiangshe Zhang, Yulun Zhang, Kai Zhang, Shuang Xu, Dongdong Chen, Radu Timofte, and Luc Van Gool. Equivariant multi-modality image fusion. In *CVPR*, 2024a.

Zixiang Zhao, Lilun Deng, Haowen Bai, Yukun Cui, Zhipeng Zhang, Yulun Zhang, Haotong Qin, Dongdong Chen, Jiangshe Zhang, Peng Wang, and Luc Van Gool. Image fusion via vision-language model. In *ICML*, 2024b.

Contents of the Appendix

The appendix covers three topics, listed below in the order they appear.

- Appendix A records the training data construction, LVP-DiffIR hyperparameters, inference and IQA protocol, compared-method recipes, and the ablation variants behind Table 4.
- Appendix B explains why QualiTeacher is omitted from Tables 2 and 3. The method trains against perceptual NR-IQA scores, so the reported numbers look strong while the restored images still show bright halos and doubled edges.
- Appendix C leads with the HM-TIR side-by-side, then shows more LVP-DiffIR results on HM-TIR, InfraRS, and M3FD, in that table order, and skips frames with a heavy ink-like outline.

A Implementation Details

This section expands the experimental setup of Sec. 5.1. The numbers below are those used for the reported LVP-DiffIR checkpoint (Stage 1 val-MUSIQ peak at 6k, then Stage 2 Dual Canny last-step MRA at 16k) and for the compared methods in Tables 2 and 3.

A.1 TRAINING DATA

Labeled RAR triplets (Stage 1). Stage 1 and Stage 2 use RAR on only 50 triplets per source, 150 in total, from InfraRS, M3FD-512, and HM-TIR. Each triplet is $(x, y_1, y_{\text{last}})$, stored as LQ / HQ1 / last. Stage 1 crops 128×128 patches with horizontal flip and 90° rotation.

Stage 2 mix. The Stage 2 loader uses a 4,906-slot mix: 2,000 InfraRS, 1,704 M3FD, and 1,202 HM-TIR. Each source contributes 50 labeled RAR triplets, and the remaining slots are unlabeled LQ. Unlabeled HQ1 and last are copies of LQ, so the same loader serves both roles. Median-ratio queues are kept per source. HM-TIR labeled slots additionally store the sensor GT used only for the bank contest in Eq. 2, not as a pixel-space training target in either stage of the reported run.

Test sets. HM-TIR LQ is synthesized from HQ by the benchmark's stripe, blur, and contrast composition. The official paired-multi test has 301 images at 640×512, and that HQ is not a training target for LVP-DiffIR. InfraRS val300 is 300 unpaired 512×512 patches, taken after the 50 labeled triplets and the next 2,000 unlabeled Stage 2 images. M3FD keeps the infrared channel only. Detection frames are subsampled at stride 8, and each remaining 1024×768 frame is cropped into four 512×512 tiles, which supplies the 1,704 training images. M3FD test300 is 300 unpaired fusion frames at native 1024×768. All three test sets are evaluated in full, and mid-training curves use only the first 50 images of each set while checkpoints are being selected.

A.2 LVP-DIFFIR TRAINING

The student is the UPSR Zhang et al. (2025) Swin-UNet Liu et al. (2021): T=5 steps, x_0 prediction, exponential noise schedule with power 0.2, κ=2.5, η_{end}=0.99, and minimum noise 0.2. The network uses 160 base channels, multipliers $(1, 2, 2, 4)$, Swin depth 2, embed dimension 192, window 8, and an LQ-conditioned encoder (`lq_size` = 64). Inputs are normalized to mean/std 0.5. Training uses mixed precision, AdamW with zero weight decay, batch size 4, and seed 0.

Stage 1. Supervision is trajectory distillation: every t regresses to y_{last} with ℓ_1 and LPIPS ($\lambda_{\text{p}} = 0.5$). Dual Canny is not used in Stage 1. The learning rate is $1{\times}10^{-4}$ with 1,000 warmup iterations and MultiStep decay by 0.5 at 5k and 8k, for 10,000 iterations. We select `net_g_6000.pth` by validation MUSIQ, not the last 10k iterate. EMA decay is 0.999.

Stage 2 Median-Ratio Adaptation. We initialize from the Stage 1 6k weights. The learning rate is $5{\times}10^{-5}$, decayed by 0.5 at 6k and 12k, for 16,000 iterations with no warmup. The reported model is the last checkpoint. Dual Canny uses min–max normalized Sobel magnitude, thresholds $(0.05, 0.08)$ and $(0.15, 0.21)$, τ=12, γ=0.6, and λ_{dc}=0.1. The memory bank admits only the last cascade step, and only if MUSIQ increases and NIQE decreases relative to the stored image. If several last-step candidates pass, the highest MUSIQ is written. Labeled HM-TIR slots start from y_{gt} unless y_{last} jointly beats it on both scores. Labeled InfraRS / M3FD slots, which have RAR last but no sensor GT, start from y_{last}. Unlabeled slots are cold-started from Teacher(x), the EMA teacher's output on the clean observation. The median-ratio gate uses a per-source FIFO of length 256 after a warmup of 32 ratios; a backward pass is taken iff $r(x)$ is at least the source median. Geometric augmentations (flip and 90° rotation, probability 0.5 each) are applied jointly to the student LQ and the bank target. Photometric corruptions (blur σ=1.5, stripe, and contrast; probability 0.4 each, at most three) are applied only to the student LQ. The EMA teacher always sees the unperturbed observation. Loader-level flip/rotation is off in Stage 2 so that those perturbations stay under MRA rather than the data loader during adaptation on the reported Stage 2 checkpoint.

A.3 INFERENCE AND METRICS

Test-time sampling uses five DDIM steps from the observation x, with 512 chopping, stride 448, and a minimum size of 64. There is no teacher, no IQA rollback, and no Dual Canny at test time. No-reference scores are computed with pyiqa on RGB tensors in $[0, 1]$ (a grayscale restore is repeated over three channels): NIQE ↓, MUSIQ ↑, MANIQA ↑, and CLIP-IQA ↑. The same protocol is used for every method in Tables 2 and 3. PSNR and SSIM are omitted because training and the reported comparison use no paired reference on HM-TIR or the unpaired tests.

A.4 COMPARED METHODS

RAR. RAR Chen et al. (2026a) is reported as a teacher oracle, and table entries are the accepted outputs y_{last}. The official SD3.5-Medium model is fine-tuned only on the HM-TIR paired-multi train split at a peak learning rate of 1.25×10^{-5}, with the VAE and DepictQA/LQA frozen.

Restormer and PromptIR. Both are trained on the official HM-TIR paired-multi $8{,}414$ LQ–HQ pairs (composite stripe, blur, and contrast), not on the 150-label LVP mix. Restormer follows the official progressive-patch AdamW schedule (base learning rate 3×10^{-4}, seed 100) with early stopping after iteration $8{,}000$; the reported weights are the in-memory snapshot at iteration $21{,}500$. PromptIR uses 128 patches, batch 8, and learning rate 2×10^{-4}; the table uses the epoch-60 checkpoint.

PPFN and SEGD. PPFN Liu et al. (2025a) is the official published `best.pth`, and we do not finetune it further. It is a 1-channel Restormer trained on HM-TIR at learning rate 8×10^{-5} for 300 epochs, and test keeps the order noise–blur–contrast, with HM-TIR in-domain and InfraRS and M3FD zero-shot. SEGD Li et al. (2026) is finetuned on the same $8{,}414$ HM-TIR pairs from the official pretrain, with DENet frozen and learning rate 5×10^{-5}, and training stops at best epoch 6. At test DENet predicts the degradation, and InfraRS and M3FD are zero-shot.

Mean Teacher and Semi-UIR. These two semi-supervised restorers are not given RAR last as a pixel target. Labeled data are the $1{,}000$ HM-TIR sensor-GT pairs. Unlabeled data are $2{,}000$ images from InfraRS and M3FD plus $1{,}906$ further unlabeled images, $3{,}906$ unlabeled LQ frames in total. Mean Teacher first trains a Restormer on the $1{,}000$ pairs, then runs EMA Mean Teacher (decay 0.999) for $10{,}000$ iterations with consistency on unlabeled LQ. Semi-UIR trains AIMNet from scratch for up to 200 epochs, crop 256, batch 4, with early stopping of patience 5 after epoch 15. We report `model_best.pth` as the Semi-UIR checkpoint in the comparison tables.

QualiTeacher. QualiTeacher uses the same $1{,}000{+}3{,}906$ split, an AIMNet backbone trained from scratch for $20{,}000$ iterations at 5×10^{-5}, and the official IQA-conditioned losses (DA-CLIP, spatial NR-IQA, DPO). Inference conditions on score 7. It is discussed only in Appendix B.

A.5 ABLATION VARIANTS

Table 4 starts from the Stage 2 Dual Canny last-step recipe and removes one choice. w/o Dual Canny keeps last-step MRA without $\mathcal{L}_{\text{DC}}$. w/o median gate still writes the bank but back-propagates on every sample. w/o S2 is the Stage 1 6k student with no MRA. w/o trajectory distillation replaces that x_0 regression to y_{last} by a ResShift-style residual in both stages, still with Dual Canny and last-step MRA. w/o RAR trains Stage 1 on 150 HM-TIR synthetic pairs (no RAR last) and runs MRA on the same unlabeled mix with identity HQ. Table 6 changes only the training centers a_t in both stages; test-time sampling still starts from x on every center variant in the reported run.

B WHY QUALITEACHER IS OMITTED FROM THE MAIN TABLES

QualiTeacher Xiao et al. (2026) does not treat NR-IQA as a test metric only. Each teacher pseudo-label is scored by an ensemble of the three terms below

$$S = S_{\text{MUSIQ}} + 0.4\, S'_{\text{BRISQUE}} + 0.2\, S_{\text{CLIP-IQA}},$$

with BRISQUE flipped so that larger is better, and all three terms min–max normalized. A dual-drop gate discards an augmented set when the three scores disagree, or when the same score jumps across geometric copies. Survivors are re-ranked and a memory bank keeps only the top three labels by S.

That scalar then enters the student in three places. A Fourier embedding of S is added to the U-Net bottleneck, so the network is conditioned on the perceptual score. The same ensemble is evaluated on a 2×2 grid and the resulting map weights the pixel loss, which is the spatial weight borrowed in Table 5. Finally a preference loss treats S as a reward: the output conditioned on a high score must outrank both the low-score output and the teacher's worst label. The training objective is therefore to raise MUSIQ, BRISQUE, and CLIP-IQA, not only to match a restored image.

The authors note that this reward can be increased by sharp repetitive texture that the metric likes and a person does not, and they add a cropped consistency loss to limit that hack. On infrared data the hack remains. The released model, run on InfraRS val300 and M3FD test300, still draws bright halos and doubled edges around roads, vehicles, pedestrians, and towers. Fig. 7 enlarges three such cases. The marks are not structure recovered from the observation. They are the edge pattern that an IQA model trained on natural images can mistake for sharpness, which is why a high MUSIQ or MANIQA would mis-rank this method in Tables 2 and 3.

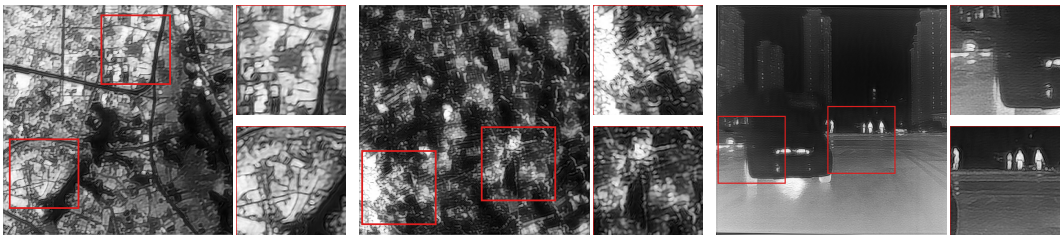

Figure 7: QualiTeacher outputs on InfraRS and M3FD. Each group is a full result with two red boxes and the enlarged boxes beside it. Halos and doubled edges follow boundaries instead of infrared structure.

C Additional Qualitative Results

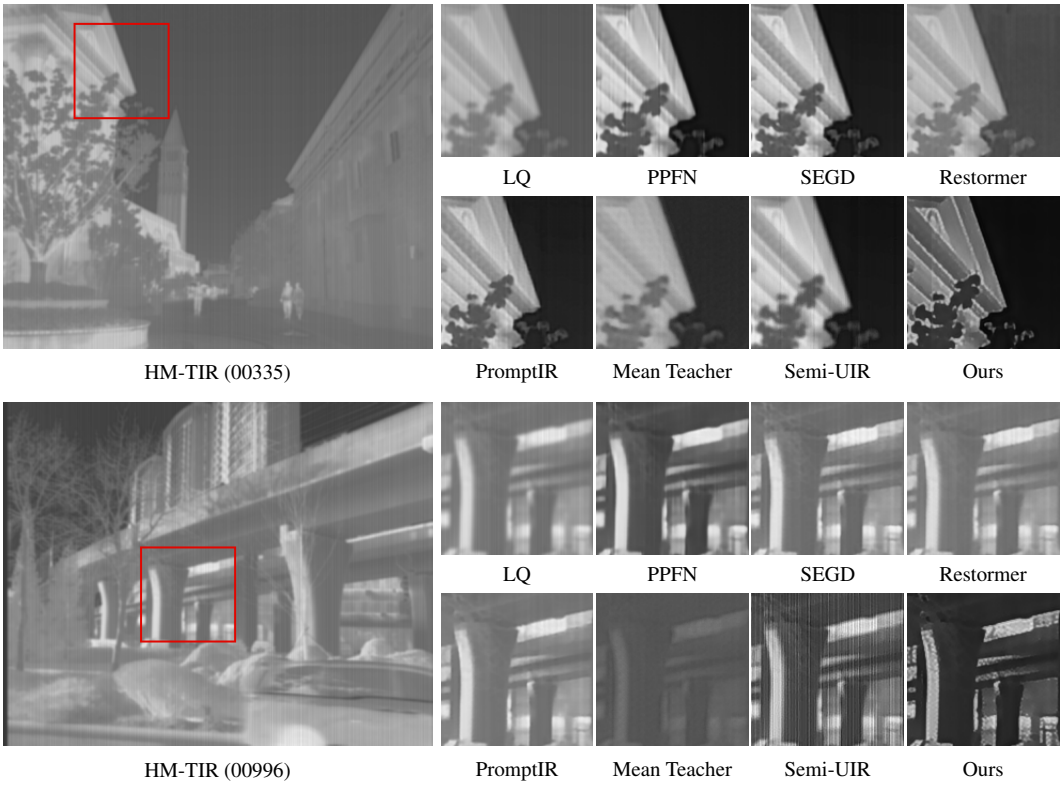

Figure 8: Visual comparison on HM-TIR. The input is on the left. Zoomed results on the right are LQ, PPFN, SEGD, Restormer, PromptIR, Mean Teacher, Semi-UIR, and Ours.



Figure 9: LVP-DiffIR on HM-TIR. Streets, bridges, and waterfronts stay readable, and frames with a crushed ink outline are not included. Each label is the source index.

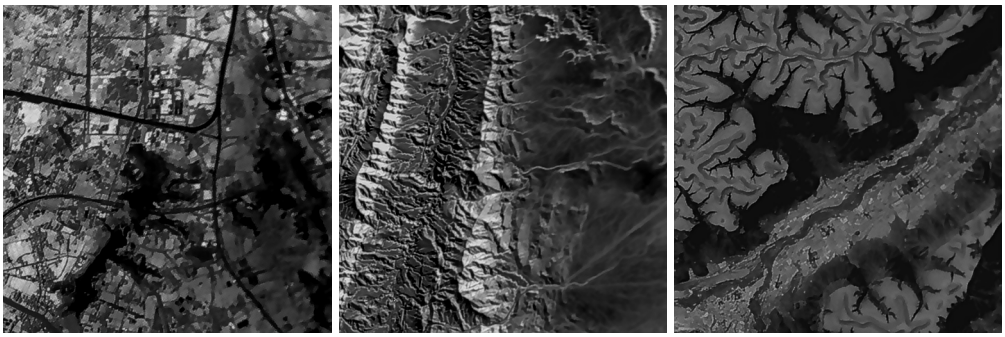

InfraRS (141) InfraRS (14) InfraRS (16)

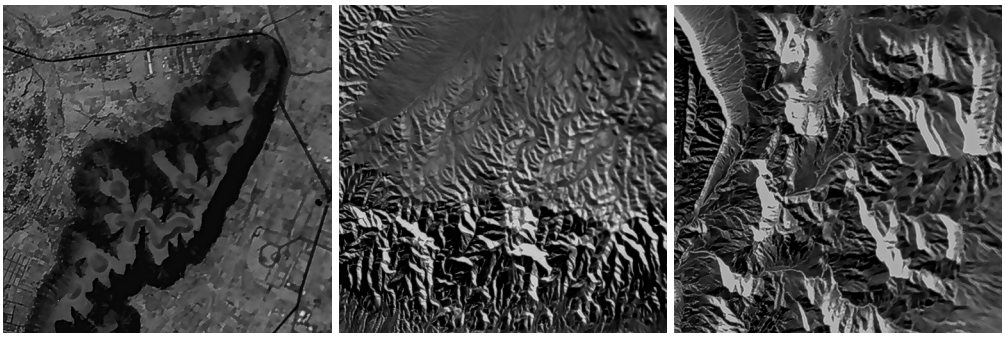

InfraRS (19) InfraRS (35) InfraRS (42)

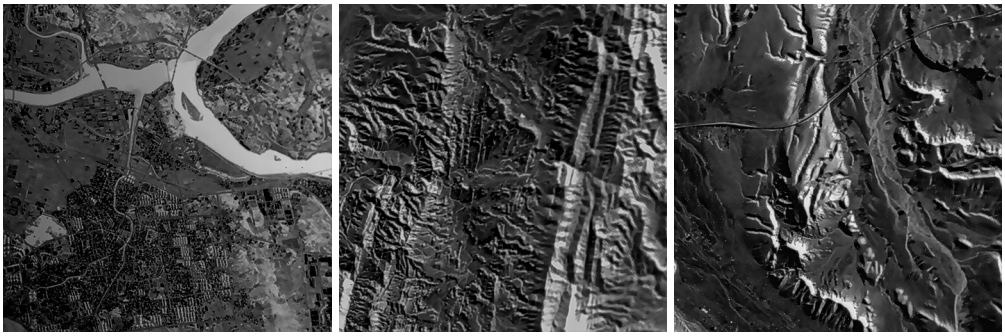

InfraRS (56) InfraRS (82) InfraRS (136)

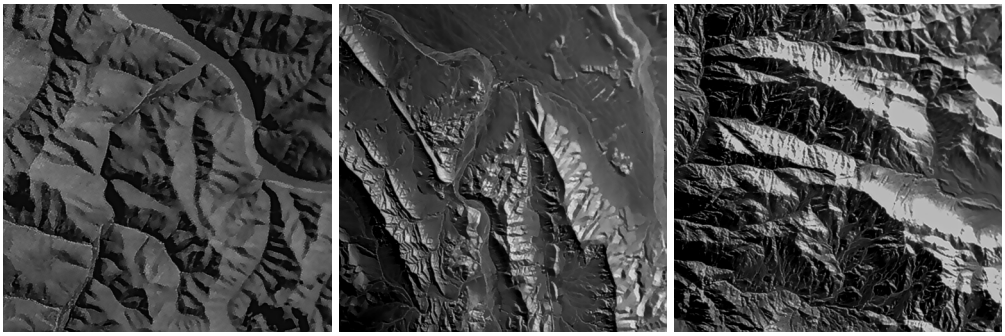

InfraRS (163) InfraRS (214) InfraRS (293)

Figure 10: LVP-DiffIR on InfraRS val300, including the frame from Fig. 6. These are the hand-picked mountain and cropland results. Each label is the validation index.



Figure 11: LVP-DiffIR on M3FD test300. Night roads, vehicles, and pedestrians are shown without the heavy ink-like stroke seen on some other frames. Each label is the test index.