BEYOND SYNTHETIC: REAL-WORLD POST-TRAINING FOR IMAGE RESTORATION WITH IMPERFECT GENERATIVE GROUND TRUTH

Anonymous authors
Paper under double-blind review

ABSTRACT

Generative image restoration has achieved impressive performance, but models trained solely on synthetic data struggle to generalize to real-world scenarios. We propose a degradation-based pre-training and post-training paradigm, in which post-training data are constructed by generating ground-truth for real low-quality (LQ) images using image editing models. However, generated targets usually contain hallucinations, making them unreliable for direct supervision. Our key insight is that effective degradation adaptation does not require every region of a generated target to be fully reliable. Based on this observation, we develop a simple yet effective approach that employs a vision-language model to localize hallucinated regions and masks their gradients during post-training. A localized refinement and verification loop further refines these regions, thereby recovering additional usable supervision. Experiments on old-photo enhancement, user-generated content enhancement, and real-world super-resolution across three representative generative restoration models demonstrate that the proposed framework consistently improves restoration performance in specific real-world scenarios.

1 INTRODUCTION

Generative image restoration (Wang et al., 2024; Lin et al., 2024; Yang et al., 2024) aims to recover high-quality images from degraded observations, leveraging generative priors. Training a generative restoration model requires pairs of low-quality (LQ) images and clean ground truth (GT) (Dong et al., 2016). Current practice is to synthesize large-scale training pairs with manually designed degradation pipelines (Wang et al., 2021b; Zhang et al., 2021; Wang et al., 2021a). However, generalizing to real LQ images by training on synthetic degradation data remains challenging.

A broad synthetic degradation distribution improves generalization, but sacrifices accuracy on a specific degradation distribution. (Zhang et al., 2023). A narrow synthetic degradation distribution will lead to artifacts on real-world images (Wang et al., 2021b). Training on real degradations is the ideal way to achieve the best performance, but such data are difficult to obtain (Cai et al., 2019).

We argue that generative image restoration can follow a degradation-based pre-training and post-training paradigm. In pre-training, a generative restoration model learns general restoration capabilities from large-scale data synthesized with diverse degradation models. Post-training then adapts the pretrained model to the degradation distribution of a specific application using limited target-domain data. Our experiments show that such post-training can substantially improve restoration performance under specific degradations. For real-world scenarios, we use generative image editing to produce ground truth for real LQ images. These generated targets are referred to as generative ground truth (GGT) (Kong et al., 2026).

Combining the degradation-based pre-training and post-training paradigm with GGT post-training data, we obtain generative restoration models tailored to specific application scenarios, such as old-photo enhancement and user-generated content enhancement.

Constructing such GGT pairs, however, is not straightforward. The generated targets usually contain hallucinations, making them unreliable as direct supervision. Our key observation is that degradation adaptation does not require every region of a generated target to be perfectly faithful. Hallucinations

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.

are localized, while the remaining regions can still provide useful supervision. Based on this observation, we propose to enable post-training with imperfect GGT. Specifically, we divide each image into patches and employ a vision-language model (VLM) to identify hallucinated regions. During post-training, gradients from these regions are masked out, allowing reliable regions to contribute to degradation adaptation.

To further exploit imperfect samples, we introduce a localized refinement and verification loop. Given an LQ image and its imperfect GGT, the VLM localizes problematic regions and generates corresponding correction instructions. An image editing model then performs targeted refinement, followed by a verifier that determines whether the refined result should be accepted. This refinement loop further improves data utilization by recovering supervision from regions that would otherwise be discarded.

We evaluate our approach on three real-world scenarios, including old-photo enhancement, user-generated content enhancement, and real-world super-resolution. Extensive experiments demonstrate that our framework consistently improves pretrained generative restoration models in specific real-world scenarios and that imperfect GGT can serve as effective supervision.

Our main contributions are summarized as follows:

- We introduce a degradation-based pre-training and post-training paradigm for generative image restoration, where general restoration capabilities learned from diverse synthetic degradations are adapted to specific application domains through GGT post-training data.
- We propose a simple and effective framework that enables post-training with imperfect GGT containing hallucinations. Hallucinated regions are identified and masked during training to avoid inaccurate supervision, while a VLM-guided localized refinement and verification loop further improves data utilization.
- We validate the proposed framework across three real-world scenarios and three representative generative restoration models, demonstrating consistent improvements and broad applicability.

2 Related Work

Generative Image Restoration Generative image restoration leverages the strong priors of pretrained text-to-image diffusion models. Early methods introduced degraded observations to guide the diffusion process for image restoration (Wang et al., 2024; Lin et al., 2024; Yang et al., 2024), followed by semantic guidance to further improve semantic fidelity (Wu et al., 2024c). More recent work has shifted toward efficient one-step restoration, using either distribution matching or adversarial training (Wu et al., 2024b; Zhang et al., 2024a; Lin et al., 2025). Despite these advances, these models are trained on synthetic degradations, leaving a gap in adapting to the degradation distributions encountered in specific real-world scenarios.

Degradation Adaptation Several studies have explored adapting SR models to specific degradation distributions. One line of work adjusts the training degradation distribution according to target-domain observations, improving accuracy on specific real-world degradations (Zhang et al., 2023). Degradation-aware pre-training has also been investigated to facilitate transfer to downstream restoration tasks (Liu et al., 2023). More recent approaches adapt pretrained SR models using real target-domain LQ images through self-supervised learning, knowledge distillation, or test-time adaptation, without requiring paired HR supervision (Chen et al., 2024; Zhang et al., 2024b; Deng et al., 2023). These studies demonstrate the importance of degradation adaptation. In contrast, our work formulates it as a real-world degradation post-training stage and focuses on constructing effective paired supervision directly from real target-domain LQ images using generative image editing, even when the generated targets contain hallucinations.

Image Editing for Restoration Data Construction Image editing models provide a promising way to construct paired restoration data directly from real degraded images. Recent works generate clean counterparts for real LQ images and improve data quality through automatic filtering, VLM-assisted refinement, and manual verification (Yang et al., 2026; Kong et al., 2026). These approaches primarily aim to obtain sufficiently reliable generated targets for restoration training. In contrast, our

work enables post-training with imperfect GGT that may still contain hallucinations. We localize unreliable regions and mask their gradients during training, while retaining the remaining reliable regions as supervision. This allows pretrained restoration models to adapt to specific application domains without requiring every generated target to be fully faithful. A localized refinement and verification loop further recovers useful supervision from imperfect samples.

3 METHOD

In this section, we first discuss our degradation-based pre-training and post-training paradigm. Building on this paradigm, we present our generated-target post-training strategy, which leverages image editing models to construct target pairs and introduces masked supervision to enable effective post-training with imperfect targets.

3.1 DEGRADATION-BASED PRE-TRAINING AND POST-TRAINING

Experimental setup. We study the proposed paradigm using synthetic pairs with controlled degradation settings. We adopt HYPIR (Lin et al., 2025) built upon Stable Diffusion 2.1 (Rombach et al., 2022) as the generative restoration backbone. We crop LSDIR (Li et al., 2023) images into non-overlapping 512×512 patches, with 10k patches used for post-training and 500 patches reserved for evaluation. The synthetic degradation is defined as $x = [((y * k_\sigma)_{\downarrow_s} + n)]_{\uparrow_s}$, where y denotes the clean image, k_σ is an isotropic Gaussian blur kernel, $n \sim \mathcal{N}(0, \sigma_n^2 I)$ represents additive Gaussian noise, and $\downarrow_s$ and $\uparrow_s$ denote bilinear downsampling and upsampling with a scale factor of s, respectively. For pre-training, we employ a Wide degradation distribution with blur standard deviation $\sigma \in [0.5, 9.5]$, downsampling factor $s \in [1, 9]$, and noise standard deviation $\sigma_n \in [0, 20]$. For post-training, we adopt a Target degradation distribution with narrower ranges, including $\sigma \in [4, 6]$, $s \in [4, 6]$, and $\sigma_n \in [8, 12]$. The model is pre-trained for 30k steps and subsequently post-trained for 5k steps.

Effectiveness of post-training. We first evaluate the effectiveness of post-training. Starting from the Wide pretrained model, we perform post-training on the Target degradation distribution and track the metric changes throughout the adaptation process. As shown in Figure 1(a), post-training consistently improves both fidelity-oriented metrics (SSIM (Wang et al., 2004) and LPIPS (Zhang et al., 2018)) and perceptual quality metrics (DeQA (You et al., 2025) and QAlign (Wu et al., 2024a)). Although the training process exhibits certain fluctuations due to the optimization instability of generative restoration models, the overall gains remain positive, demonstrating that post-training enhances both fidelity and perceptual quality. The qualitative comparison in Figure 1(top) further illustrates the visual improvements brought by post-training. Compared with the pretrained model, the post-trained model recovers finer details on the bird feathers and produces sharper textures, resulting in a reconstruction closer to the ground truth.

We further study the data efficiency of post-training by varying the amount of target-domain data. With the pretrained initialization fixed, we perform post-training using 100, 1k, and 10k target patches (Figure 1(b)). When using only 100 patches, DeQA improves at the early stage but gradually decreases with further optimization, indicating overfitting to the limited target data. With 1k patches, the model achieves stable improvements over the pretrained initialization throughout the post-training process, suggesting that degradation adaptation can be achieved with a moderate amount of target-domain data. Using 10k patches further improves the adaptation performance, demonstrating the benefit of additional target-domain supervision.

The necessity of wide degradation pre-training. We first investigate whether pre-training is necessary for effective adaptation to the Target distribution when only limited target-domain data are available. We compare models post-trained on the Target distribution with and without pre-training, and track their performance throughout the adaptation process. As shown in Figure 2(a), the model with pre-training consistently outperforms the model without pre-training, demonstrating the importance of pre-training as an effective initialization for subsequent adaptation.

We further study how the coverage of the pre-training degradation distribution influences post-training performance. Specifically, we compare the Wide pre-training distribution with a Narrow distribution, defined by $\sigma \in [2, 8]$, $s \in [2, 8]$, and $\sigma_n \in [2, 18]$, and evaluate their adaptation to the

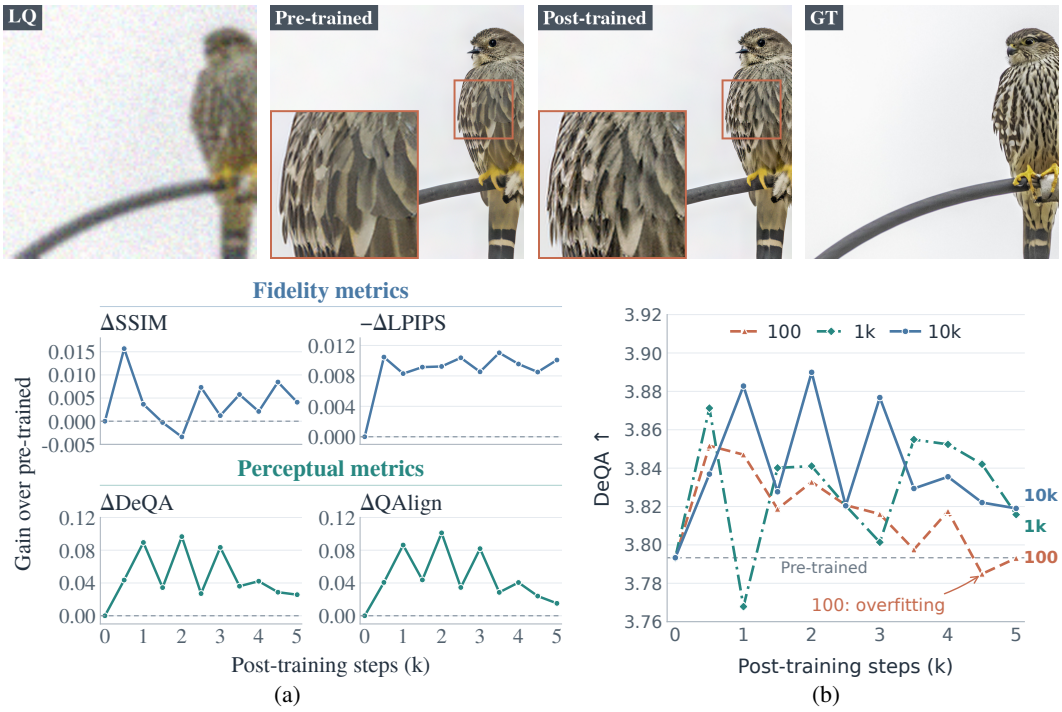


Figure 1: Effectiveness of post-training on the Target degradation. Top: qualitative comparison of the LQ input, pretrained output, post-trained output, and ground truth. (a) Gains over the pretrained model on SSIM, LPIPS, DeQA, and QAlign during post-training; signs are flipped for LPIPS so that higher values indicate improvement. (b) DeQA during post-training with 100, 1k, and 10k target patches; the dashed line marks the pretrained model.

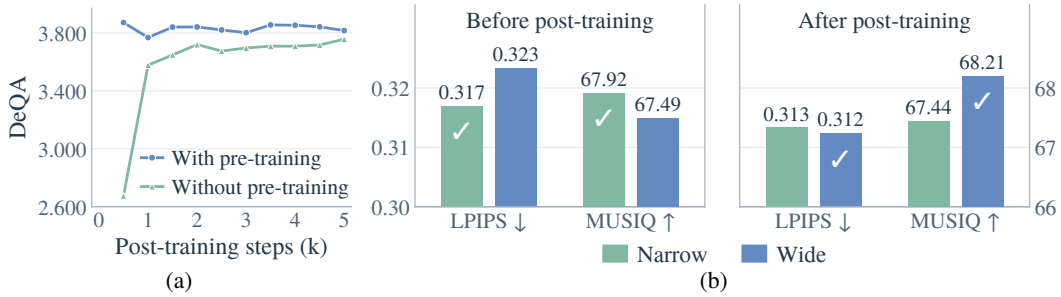


Figure 2: Analysis of pre-training. (a) DeQA during post-training on the Target degradation, with and without pre-training. (b) LPIPS and MUSIQ of the Narrow and Wide pre-trained models before and after post-training on the Target degradation.

Target distribution. As shown in Figure 2(b), Narrow achieves better LPIPS and MUSIQ (Ke et al., 2021) before post-training, which is expected due to its closer alignment with the Target distribution. After post-training, however, Wide surpasses Narrow on both metrics. These results suggest that broader degradation coverage during pre-training improves adaptation capability.

Having established the paradigm, we next introduce how generative image editing models are used to construct post-training data for real-world target domains.

3.2 Data Generation and Masked Post-training

Initial target generation. For each real LQ image x, we use an image editing model $\mathcal{G}$ to generate an initial generative ground truth: $y^{(0)} = \mathcal{G}(x)$. The generated targets typically contain hallucina-

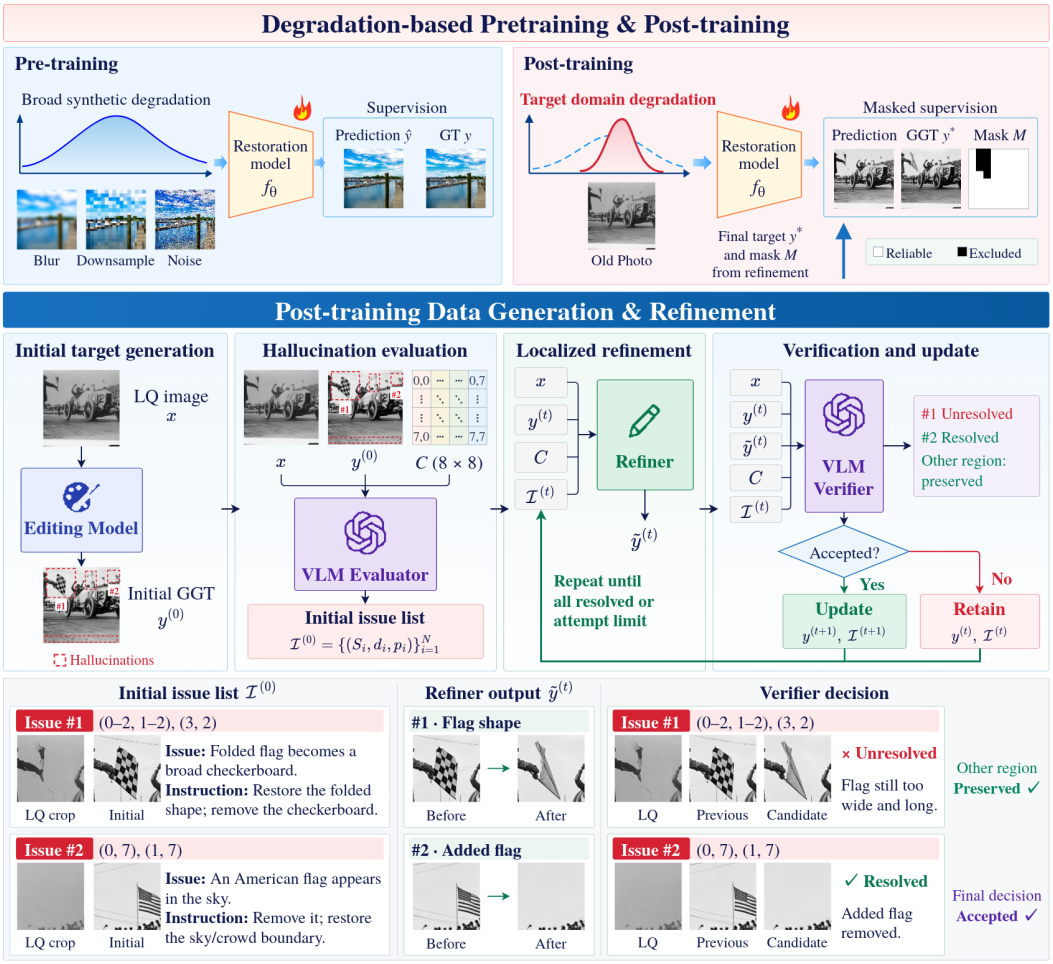


Figure 3: Overview of the framework. Top: degradation-based pre-training and post-training. Pre-training learns a restoration model from a broad synthetic degradation distribution, and post-training adapts it to the target-domain degradation distribution with masked supervision from generated targets. Bottom: target generation and refinement. An image editing model generates an initial GGT from the LQ image, a VLM evaluator localizes hallucinations, a refiner performs localized corrections, and a verifier determines whether the refined result is acceptable. The lowest row illustrates the initial issue list, the refiner output, and the verifier decisions for one refinement step.

tions, resulting in unreliable supervision. Our key observation is that degradation adaptation does not require a generated target to be fully reliable across all regions. Since hallucinations are localized, reliable regions can still provide useful supervision for adaptation. We therefore mask unreliable regions and retain supervision only from reliable regions during post-training.

Masked post-training. Given a binary supervision mask M, where $M(u) = 1$ indicates reliable regions and $M(u) = 0$ indicates unreliable regions, we apply masked supervision to spatially aligned restoration objectives. For the restoration output $\hat{y} = f_\theta(x)$ and generated target y, the reconstruction loss is $\mathcal{L}_{\text{MSE}} = \|M \odot (\hat{y} - y)\|_2^2$; perceptual supervision uses $\mathcal{L}_{\text{LPIPS}} = \text{LPIPS}(M \odot \hat{y}, M \odot y)$. For diffusion-based restoration models, we downsample M to the latent resolution as $M_\downarrow$ and mask the denoising loss, $\mathcal{L}_{\text{diff}} = \mathbb{E}_{t,\epsilon}\|M_\downarrow \odot (\epsilon_\theta(z_t, t, x) - \epsilon)\|_2^2$. Distribution-level objectives, such as adversarial (Goodfellow et al., 2014) and DMD (Yin et al., 2024) losses, remain unchanged.

VLM-based hallucination evaluation. To identify unreliable regions, a VLM-based evaluator is introduced to localize hallucinated regions in the generated targets. To enable precise spatial

localization, we divide each image into 8×8 patches. We additionally construct a patch-coordinate map C, consisting of a white canvas divided into the same 8×8 patches, with each patch annotated by its spatial coordinate. Given (x, y, C), the evaluator $\mathcal{V}_{\text{eval}}$ identifies hallucinated patches and outputs $\mathcal{S} = \mathcal{V}_{\text{eval}}(x, y, C) = \{S_i\}_{i=1}^{N}$, where S_i denotes the affected patches. The identified patches are converted into a binary supervision mask $M(u) = \mathbf{1}[u \notin \bigcup_{i=1}^{N} S_i]$, where u indexes spatial locations.

VLM-guided localized refinement. Although masking hallucinated regions prevents unreliable supervision, it also discards potentially useful information from these regions. To recover additional supervision, we further refine the hallucinated regions through a localized refinement loop. The Evaluator is first used to provide additional region-level feedback for the detected hallucinations. Specifically, given the LQ image x, the initial GGT $y^{(0)}$, and the patch-coordinate map C, the Evaluator augments the detected hallucination regions with mismatch descriptions and correction instructions, forming the initial issue set:

$$\mathcal{I}^{(0)} = \mathcal{V}_{\text{eval}}(x, y^{(0)}, C) = \{(S_i, d_i, p_i)\}_{i=1}^{N}, \tag{1}$$

where S_i denotes the hallucinated region, d_i describes the mismatch between the LQ image and the generated target, and p_i specifies the corresponding correction instruction. The Evaluator is invoked only once at initialization. The subsequent refinement process consists of a Refiner and a Verifier, with the issue set maintained by the Verifier.

Refiner. At iteration t, the Refiner, implemented by the image editing model $\mathcal{G}$, takes the LQ image x, the current generated target $y^{(t)}$, the patch-coordinate map C, and the current issue set $\mathcal{I}^{(t)}$ as inputs, and produces a refined candidate $\widetilde{y}^{(t)} = \mathcal{G}(x, y^{(t)}, C, \mathcal{I}^{(t)})$. The Refiner is instructed to correct the specified hallucinated regions while preserving all remaining regions.

Verifier. The Verifier $\mathcal{V}_{\text{ver}}$ assesses the refined candidate by checking whether the hallucinations are effectively corrected while ensuring that the remaining regions preserve their original content and visual quality. Based on the verification result, the Verifier updates the issue set by removing resolved issues and revising the descriptions of remaining issues to reflect their current refinement states:

$$(\widetilde{\mathcal{I}}^{(t)}, a^{(t)}) = \mathcal{V}_{\text{ver}}(x, y^{(t)}, \widetilde{y}^{(t)}, C, \mathcal{I}^{(t)}), \tag{2}$$

where $\widetilde{\mathcal{I}}^{(t)}$ is the updated issue set. The acceptance flag $a^{(t)} = 1$ only when the refinement resolves at least one existing issue without introducing unintended modifications or degrading the visual quality of unaffected regions. The refinement state is updated as

$$(y^{(t+1)}, \mathcal{I}^{(t+1)}) = \begin{cases} (\widetilde{y}^{(t)}, \widetilde{\mathcal{I}}^{(t)}), & a^{(t)} = 1, \\ (y^{(t)}, \mathcal{I}^{(t)}), & a^{(t)} = 0. \end{cases} \tag{3}$$

This update rule ensures monotonic refinement: $|\mathcal{I}^{(t+1)}| \leq |\mathcal{I}^{(t)}|$. The refinement process terminates when all initial issues are resolved or the maximum number of refinement attempts is reached. The remaining unresolved regions are masked during post-training. In our implementation, we use Nano Banana 2 (Google, 2026) as the image editing model and GPT-5.6 Sol (OpenAI, 2026) as the VLM. Figure 3 provides an overview of the data generation and masked post-training framework. The prompts for the Evaluator, Refiner, and Verifier are provided in the supplementary file.

4 Experiments

4.1 Experimental Setup

Data collection and preparation. We evaluate our framework in three real-world scenarios: old-photo enhancement, user-generated content (UGC) enhancement, and real-world super-resolution. For old-photo enhancement, we collect historical images from the Library of Congress Flickr Commons collection[1], using 5K images for training-data construction and 100 images for testing, which

[1] https://www.flickr.com/photos/library_of_congress/

Table 1: Quantitative comparisons on old-photo and UGC enhancement. *+pt* denotes post-trained models. Bold values indicate better results within each pretrained/post-trained pair.

Old-photo enhancement

Method	LF-PSNR ↑	LF-SSIM ↑	LF-LPIPS ↓	MUSIQ ↑	DeQA ↑	QAlign ↑	VQ-R1 ↑
SeeSR	**33.324**	0.901	**0.162**	65.099	3.339	3.403	3.227
SeeSR +pt	30.538	**0.911**	0.184	**67.367**	**3.509**	**3.697**	**3.470**
OSEDiff	**30.400**	0.897	0.191	65.364	3.334	3.523	3.494
OSEDiff +pt	29.578	**0.919**	**0.169**	**65.921**	**3.418**	**3.609**	**3.518**
HYPIR	29.530	0.854	0.209	65.580	3.134	3.499	3.500
HYPIR +pt	**29.703**	**0.872**	**0.201**	**65.836**	**3.365**	**3.642**	**3.605**

UGC enhancement

Method	LF-PSNR ↑	LF-SSIM ↑	LF-LPIPS ↓	MUSIQ ↑	DeQA ↑	QAlign ↑	VQ-R1 ↑
SeeSR	34.354	0.945	**0.112**	66.984	3.582	3.689	3.286
SeeSR +pt	**34.653**	**0.954**	0.129	**69.296**	**3.763**	**3.900**	**3.608**
OSEDiff	31.125	0.938	0.130	**66.114**	3.447	3.626	3.391
OSEDiff +pt	**33.571**	**0.962**	**0.102**	65.893	**3.497**	**3.638**	**3.494**
HYPIR	30.930	0.907	0.164	**68.746**	3.491	3.708	3.503
HYPIR +pt	**35.205**	**0.948**	**0.129**	66.816	**3.626**	**3.755**	**3.585**

yield about 22K training patches after processing. For UGC enhancement, we collect about 41K real LQ images from Kuaishou, yielding about 172K training patches after processing, with a separate set of 100 images for testing. More details on data preprocessing and construction are provided in the supplementary file. For real-world super-resolution, we evaluate on the RealSR (Cai et al., 2019) and DRealSR (Wei et al., 2020) benchmarks, with the corresponding results provided in the supplementary file.

Models and evaluation. Across all three scenarios, we evaluate SeeSR (Wu et al., 2024c), OSEDiff (Wu et al., 2024b), and HYPIR (Lin et al., 2025), representing multi-step diffusion, one-step distribution-matching, and one-step adversarial restoration, respectively. All models are initialized from their pretrained weights and post-trained for 50K, 10K, and 10K steps with effective batch sizes of 16, 16, and 32, respectively. We assess perceptual quality using four no-reference metrics: MUSIQ (Ke et al., 2021), DeQA (You et al., 2025), QAlign (Wu et al., 2024a), and VisualQuality-R1 (VQ-R1) (Wu et al., 2025). We further assess low-frequency (LF) input consistency (Wald et al., 1997; Scarpa & Ciotola, 2022). Specifically, we apply the same 5×5 Gaussian blur to both the LQ input and the restored image, and compute LF-PSNR (Horé & Ziou, 2010), LF-SSIM (Wang et al., 2004), and LF-LPIPS (Zhang et al., 2018) between the filtered images.

4.2 Experimental Results

Quantitative results. Tables 1 report the quantitative results on old-photo and UGC enhancement. Although the evaluated models adopt substantially different restoration formulations, including multi-step diffusion with denoising objectives, one-step restoration based on distribution matching, and adversarially trained one-step restoration, they all achieve substantial improvements after post-training. More importantly, rather than moving along the quality–fidelity trade-off, the post-trained models achieve a Pareto improvement, improving both perceptual quality and fidelity. This observation is consistent with our synthetic experiments, where post-training simultaneously improves perceptual and fidelity-oriented metrics. These results validate large-scale synthetic degradation pre-training followed by target-domain post-training as an effective paradigm for real-world image restoration. Our data construction strategy further makes such post-training feasible on real target domains by leveraging an image editing model to construct target-domain supervision, while masked supervision restricts optimization to reliable regions of the generated targets.

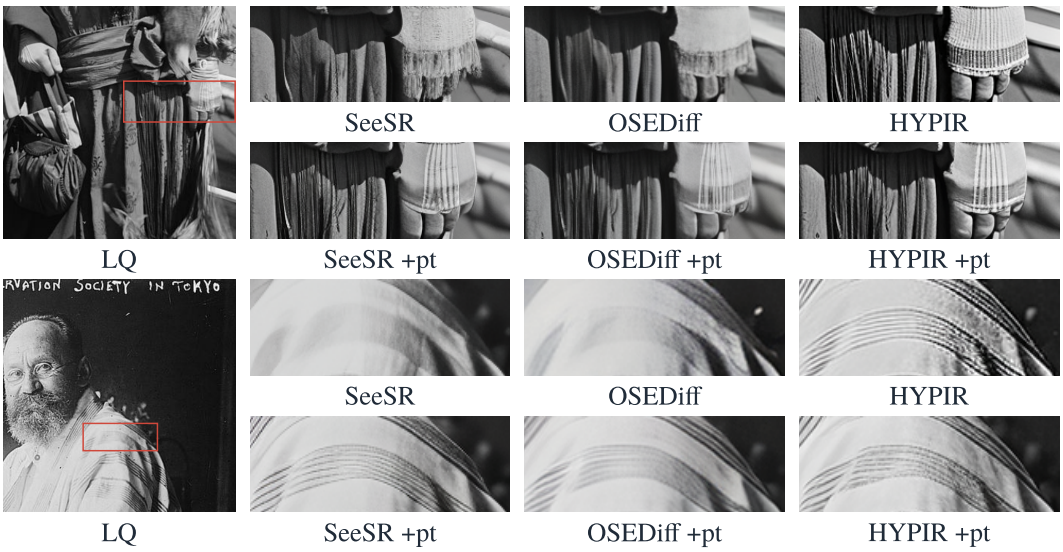


Figure 4: Qualitative comparisons on old-photo enhancement. *+pt* denotes post-trained models.

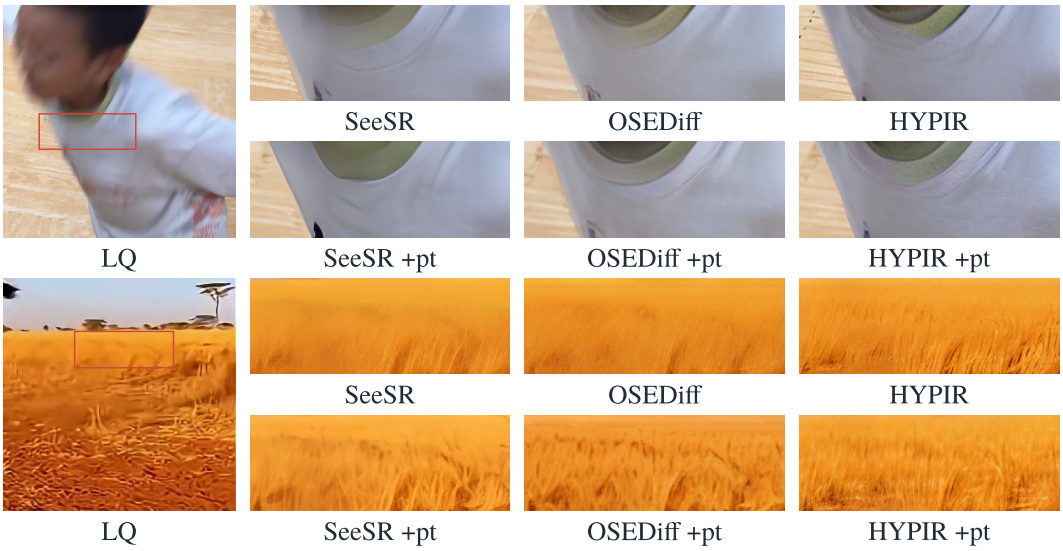


Figure 5: Qualitative comparisons on UGC enhancement. *+pt* denotes post-trained models.

Table 2: Ablation results of masking and refinement. Parentheses show stepwise gains.

Training setting	MUSIQ ↑	DeQA ↑	QAlign ↑	VQ-R1 ↑
Drop	58.076	3.112	3.243	3.143
+ Mask	58.757 (+0.681)	3.135 (+0.023)	3.294 (+0.050)	3.211 (+0.067)
+ Refinement	59.153 (+0.396)	3.185 (+0.050)	3.299 (+0.006)	3.221 (+0.010)

Table 3: Ablation results of the refinement loop. Bold values indicate the best result.

Loop	Location	Verification	PSNR ↑
×	✓	✓	32.600
✓	×	✓	31.360
✓	✓	×	31.385
✓	✓	✓	**32.871**

Table 4: Ablation results of degradation alignment. Bold values indicate better results.

Training data	LF-PSNR ↑	LF-SSIM ↑	LF-LPIPS ↓	DeQA ↑	VQ-R1 ↑
GGT + Real-ESRGAN LQ	31.877	0.940	0.129	**3.509**	3.456
GGT + Real LQ	**33.571**	**0.962**	**0.102**	3.497	**3.494**

Qualitative results. Figures 4 and 5 show qualitative comparisons on old-photo and UGC enhancement. For old-photo enhancement, post-training produces substantially more faithful hand, garment, and stripe structures than the pretrained models. This improvement can be attributed to degradation adaptation: pre-training with the broad Real-ESRGAN degradation space requires the model to accommodate a wide range of degradations, while post-training allows it to specialize in old-photo degradations, improving restoration accuracy. In the second example, post-training recovers clearer stripe patterns and richer fabric textures on the subject's clothing, demonstrating improved perceptual quality. For UGC enhancement, the first example is dominated by severe motion blur, while the second contains strong compression artifacts. Such degradations are not well covered by the Real-ESRGAN degradation distribution. As a result, the pretrained models struggle to recover garment structures and fabric details under motion blur, as well as meaningful field textures under compression. After target-domain post-training, the models recover substantially clearer structures and richer fine-grained textures in both cases. These results demonstrate the effectiveness of post-training in adapting pretrained restoration models to real-world degradations.

5 Ablation Study

5.1 Effectiveness of Masked Supervision and Refinement

We perform an ablation study to evaluate masked supervision and target refinement. We consider three settings: *Drop* discards samples whose generated targets contain detected hallucinations; + *Mask* retains these samples while masking unreliable regions during training; and + *Refinement* further refines the detected regions before post-training. As shown in Table 2, + *Mask* consistently improves all evaluated metrics over *Drop*, indicating that imperfect generated targets still provide valuable supervision through their reliable regions. + *Refinement* further improves performance, showing that correcting unreliable regions can recover additional useful supervision.

5.2 Effectiveness of the Refinement Loop

We refine hallucinated generated targets for 20 real LQ images and compute PSNR against their corresponding hallucination-free references. We ablate three components: *Loop*, *Location*, and *Verification*. Without *Loop*, only one refinement round is performed. Without *Location*, patch-level localization is replaced with coarse spatial descriptions (e.g., left or right). Without *Verification*, the verifier is removed, and the evaluator is rerun after each refinement to update the remaining hallucinations. As shown in Table 3, the full pipeline achieves the highest PSNR, with large drops when disabling *Location* or *Verification*.

5.3 Effectiveness of Degradation Adaptation

During post-training, both the image content and degradation are aligned with the target domain. To isolate the effect of degradation adaptation, we additionally construct synthetic LQ inputs by applying the Real-ESRGAN degradation model to the GGTs. We perform post-training on both the synthetic and real LQ inputs and compare the resulting performance. As shown in Table 4, using real LQ inputs consistently improves input consistency. These results confirm degradation alignment with the target domain as a key source of the performance gains from post-training.

6 Conclusion

We proposed a degradation-based pre-training and post-training paradigm for generative image restoration, which adapts pretrained models to real-world target domains through post-training on GGT data with masked supervision. Experiments on old-photo enhancement, UGC enhancement, and real-world super-resolution demonstrate consistent improvements over pretrained models. However, constructing GGT data remains costly, and the VLM evaluator may miss hallucinated regions. Reducing the cost of data construction and improving the reliability of hallucination localization are important directions for future work.

AI USE STATEMENT

In this work, we used generative AI tools to assist with writing research code, designing and creating the method framework figure, and polishing the language of the manuscript. The authors take full responsibility for the final code, figure, and manuscript, including any content produced with AI assistance.

ETHICS STATEMENT

This work studies post-training with real low-quality images collected from online sources, including historical photographs and user-generated content. Such images may contain identifiable individuals or copyrighted material, and online availability does not imply permission for unrestricted reuse or redistribution. Before any public release of the collected images, their copyright status, source-specific terms of use, and privacy implications would need to be carefully reviewed.

REPRODUCIBILITY STATEMENT

The Method section describes generated-target construction, hallucination localization, masked post-training, and localized refinement. Dataset sources, model configurations, training schedules, and evaluation metrics are reported in Section 4.1. Additional data preparation details and real-world super-resolution benchmark results are provided in the supplementary file.

REFERENCES

Jianrui Cai, Hui Zeng, Hongwei Yong, Zisheng Cao, and Lei Zhang. Toward real-world single image super-resolution: A new benchmark and a new model. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3086–3095, 2019.

Haoyu Chen, Wenbo Li, Jinjin Gu, Jingjing Ren, Haoze Sun, Xueyi Zou, Zhensong Zhang, Youliang Yan, and Lei Zhu. Low-res leads the way: Improving generalization for super-resolution by self-supervised learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024.

Zeshuai Deng, Zhuokun Chen, Shuaicheng Niu, Thomas H. Li, Bohan Zhuang, and Mingkui Tan. Efficient test-time adaptation for super-resolution with second-order degradation and reconstruction. In *Advances in Neural Information Processing Systems*, 2023.

Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 38(2): 295–307, 2016.

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, 2014.

Google. Gemini 3.1 Flash Image (Nano Banana 2). Google AI for Developers, 2026. URL `https://ai.google.dev/gemini-api/docs/models/gemini-3.1-flash-image`.

Alain Horé and Djemel Ziou. Image quality metrics: PSNR vs. SSIM. In *2010 20th International Conference on Pattern Recognition*, pp. 2366–2369, 2010. doi: 10.1109/ICPR.2010.579.

Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. MUSIQ: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 5148–5157, 2021.

Xiangtao Kong, Jixin Zhao, Lingchen Sun, Rongyuan Wu, and Lei Zhang. GGT-100K: Generative ground truth for generalizable real-world image restoration. *arXiv preprint arXiv:2605.31039*, 2026.

Yawei Li, Kai Zhang, Jingyun Liang, Jiezhang Cao, Ce Liu, Rui Gong, Yulun Zhang, Hao Tang, Yun Liu, Denis Demandolx, Rakesh Ranjan, Radu Timofte, and Luc Van Gool. LSDIR: A large scale dataset for image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1775–1787, 2023.

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

Xinqi Lin, Fanghua Yu, Jinfan Hu, Zhiyuan You, Wu Shi, Jimmy S. Ren, Jinjin Gu, and Chao Dong. Harnessing diffusion-yielded score priors for image restoration. *ACM Transactions on Graphics*, 2025. doi: 10.1145/3763346.

Yihao Liu, Jingwen He, Jinjin Gu, Xiangtao Kong, Yu Qiao, and Chao Dong. DegAE: A new pretraining paradigm for low-level vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 23292–23303, 2023.

OpenAI. GPT-5.6 Sol Model. OpenAI API Documentation, 2026. URL `https://developers.openai.com/api/docs/models/gpt-5.6-sol`.

Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Bjorn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10684–10695, 2022.

Giuseppe Scarpa and Matteo Ciotola. Full-resolution quality assessment for pansharpening. *Remote Sensing*, 14(8):1808, 2022. doi: 10.3390/rs14081808.

Lucien Wald, Thierry Ranchin, and Marc Mangolini. Fusion of satellite images of different spatial resolutions: Assessing the quality of resulting images. *Photogrammetric Engineering & Remote Sensing*, 63(6):691–699, 1997. URL `https://www.asprs.org/wp-content/uploads/pers/97journal/june/1997_jun_691-699.pdf`.

Jianyi Wang, Zongsheng Yue, Shangchen Zhou, Kelvin C. K. Chan, and Chen Change Loy. Exploiting diffusion prior for real-world image super-resolution. *International Journal of Computer Vision*, 2024.

Xintao Wang, Yu Li, Honglun Zhang, and Ying Shan. Towards real-world blind face restoration with generative facial prior. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2021a.

Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-ESRGAN: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision Workshops*, 2021b.

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.

Pengxu Wei, Ziwei Xie, Hannan Lu, Zongyuan Zhan, Qixiang Ye, Wangmeng Zuo, and Liang Lin. Component divide-and-conquer for real-world image super-resolution. In *European Conference on Computer Vision*, pp. 101–117, 2020.

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, 2024a.

Rongyuan Wu, Lingchen Sun, Zhiyuan Ma, and Lei Zhang. One-step effective diffusion network for real-world image super-resolution. In *Advances in Neural Information Processing Systems*, 2024b.

Rongyuan Wu, Tao Yang, Lingchen Sun, Zhengqiang Zhang, Shuai Li, and Lei Zhang. SeeSR: Towards semantics-aware real-world image super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2024c.

Tianhe Wu, Jian Zou, Jie Liang, Lei Zhang, and Kede Ma. VisualQuality-R1: Reasoning-induced image quality assessment via reinforcement learning to rank. In *Advances in Neural Information Processing Systems*, 2025. URL https://proceedings.neurips.cc/paper_files/paper/2025/file/7f8f7bf2d357a09bdffff04b2e0f5c4e-Paper-Conference.pdf.

Tao Yang, Rongyuan Wu, Peiran Ren, Xuansong Xie, and Lei Zhang. Pixel-aware stable diffusion for realistic image super-resolution and personalized stylization. In *European Conference on Computer Vision*, 2024.

Yufeng Yang, Xianfang Zeng, Zhangqi Jiang, Fukun Yin, Jianzhuang Liu, Wei Cheng, Jinghong Lan, Shiyu Liu, Yuqi Peng, Gang Yu, and Shifeng Chen. RealRestorer: Towards generalizable real-world image restoration with large-scale image editing models. *arXiv preprint arXiv:2603.25502*, 2026.

Tianwei Yin, Michaël Gharbi, Richard Zhang, Eli Shechtman, Frédo Durand, William T. Freeman, and Taesung Park. One-step diffusion with distribution matching distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 6613–6623, 2024.

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.

Aiping Zhang, Zongsheng Yue, Renjing Pei, Wenqi Ren, and Xiaochun Cao. Degradation-guided one-step image super-resolution with diffusion priors. *arXiv preprint arXiv:2409.17058*, 2024a.

Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 2021.

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

Ruofan Zhang, Jinjin Gu, Haoyu Chen, Chao Dong, Yulun Zhang, and Wenming Yang. Crafting training degradation distribution for the accuracy-generalization trade-off in real-world super-resolution. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 41078–41091, 2023.

Yuehan Zhang, Seungjun Lee, and Angela Yao. Pairwise distance distillation for unsupervised real-world image super-resolution. In *European Conference on Computer Vision*, 2024b.