QUERY-TO-COMPETENCE: USAGE-TIME CONTINUAL ADAPTATION FOR REAL-WORLD IMAGE RESTORATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Real-world image restoration must handle compound and continually changing degradations. We consider a deployment setting that combines an updateable self-hosted model, referred to as the Student, with a hosted image-editing service, referred to as the Teacher. In a service receiving many complex restoration requests, the Teacher can resolve difficult cases, but repeated calls incur provider cost. Our goal is to progressively reduce hosted-service dependence by turning quality-triggered usage-time Teacher calls into reusable Student competence. We propose **Query-to-Competence (Q2C)**, a usage-time adaptive restoration system for this purpose. Q2C constructs a Structured Restoration Brief (SRB) describing the degradations, restoration intent, and content-preservation constraints. The Student processes each request first, and the Teacher is queried only when the Student output fails a fixed quality criterion. Curated Teacher outputs are accumulated for Replay-Distill Adaptation (RDA), while Guarded Student Promotion (GSP) determines whether an updated candidate enters deployment. Across seven actual usage sessions, the deployed Student handles an increasing fraction of requests without Teacher fallback. On the fixed SDCR-200 evaluation set, Student-only coverage increases from 24.0% to 81.5%, showing that selected Teacher feedback can be retained as reusable Student competence.

1 INTRODUCTION

Image restoration aims to recover clear and content-faithful images from degraded observations (Zamir et al., 2022; Chen et al., 2022; Potlapalli et al., 2023). In real-world settings, degradations are rarely isolated or stationary: blur, noise, low illumination, haze, rain, compression artifacts, reflections, and old-photo damage may co-occur and change with imaging devices and environments. Although recent all-in-one methods extend restoration beyond the one-degradation–one-model paradigm (Li et al., 2022; Luo et al., 2025; Zhou et al., 2025), their generalization remains strongly constrained by available supervision. Synthetic degradations are scalable but incompletely reproduce real image formation, while collecting accurately aligned degraded–clean pairs is expensive and difficult (Wang et al., 2021; Rajagopalan et al., 2025; Yang et al., 2025). Expanding restoration competence has therefore traditionally relied on constructing increasingly large paired datasets.

General-purpose image-editing models provide an alternative source of restoration capability. We consider privacy-sensitive or domain-specific restoration services that require local data control, predictable serving, and continual model adaptation, making self-hostable models such as Qwen-Image-Edit suitable as persistent local models (Wu et al., 2025). Proprietary hosted services such as GPT-Image-2 and Nano Banana 2 provide broader open-world capability but are accessible only through external queries (OpenAI, 2026; Google DeepMind, 2026). We therefore refer to the locally deployed and updateable model as the Student and the hosted service as the Teacher.

Existing transfer methods use stronger image-editing models to construct restoration supervision before deployment (Yang et al., 2026; Kong et al., 2026). Although this offline supervision can improve an accessible model, the resulting capability remains bounded by the data collected in advance.

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.

During actual use, a service may receive a sustained stream of complex restoration requests beyond the Student's current capability. A hosted Teacher can resolve these cases, but conventional fallback does not improve the Student, so similar requests may repeatedly incur the cost of an external query. Q2C instead converts selected Teacher outputs into demand-aligned supervision for subsequent Student updates, turning one-off fallback responses into reusable restoration competence.

Realizing this query-to-competence process requires four coordinated decisions: routing identifies requests that require Teacher fallback; curation selects reliable Teacher outputs for capability transfer; retention carries acquired competence across sequential updates; and promotion prevents regressive candidates from entering deployment.

To this end, we propose Query-to-Competence (Q2C), a usage-time adaptive restoration system that progressively transfers demand-aligned Teacher supervision into an updateable self-hosted Student. Q2C uses a Structured Restoration Brief (SRB) to specify the restoration request, processes each input with the Student first, and queries the Teacher only when the Student output fails the quality criterion. Curated Teacher outputs enter Replay-Distill Adaptation (RDA), which combines LoRA training, historical replay, and previous-Student distillation (Hu et al., 2022; Buzzega et al., 2020; Yan et al., 2024). Guarded Student Promotion (GSP) then determines whether the candidate replaces the deployed Student. We instantiate Q2C with GPT-Image-2 as the hosted Teacher and a LoRA-adapted Qwen-Image-Edit-2511 model as the self-hosted Student. Across seven actual usage sessions, the deployed Student handles an increasing fraction of requests without Teacher fallback. On the unchanged SDCR-200, Student-only coverage increases from 24.0% to 81.5%, showing that the operational reduction in Teacher fallback is accompanied by reusable competence accumulated in the Student. The final candidate also achieves the strongest task-aware performance among the evaluated non-Teacher models.

Our contributions are threefold:

- We formulate real-world image restoration as an adaptive serving and learning system: quality-triggered Teacher calls resolve difficult requests and provide targeted supervision that expands the self-hosted Student's competence over time.
- We introduce SRB-guided routing and curation to connect restoration intent, content-preservation constraints, and supervision quality in the same usage-time workflow.
- We develop RDA and GSP to carry acquired competence across usage sessions and decide when an updated Student should enter deployment, combining replay and previous-Student distillation with validation on current and historical requests.

2 Related Work

2.1 Real-World Image Restoration

Image restoration recovers high-quality images from degraded inputs, with extensive research on specific tasks such as denoising, deblurring, deraining, and dehazing (Zamir et al., 2022; Chen et al., 2022). Recent advances have shifted toward unified all-in-one models based on CNNs, Transformers, multimodal prompts, diffusion models, and vision-language guidance to handle multiple degradation types (Li et al., 2022; Potlapalli et al., 2023; Ai et al., 2024; Zheng et al., 2024; Tian et al., 2025; Zeng et al., 2025; Zhou et al., 2025; Cheng et al., 2026).

Recent work also questions the optimality of fixed ground-truth supervision and introduces image-quality priors for real-world restoration (Xiao et al., 2026). While all-in-one methods improve generalization through larger models and datasets, their supervision remains fixed. Q2C acquires selected restoration supervision during deployment.

2.2 Generative Image Editing Models for Restoration

Recent methods formulate restoration as instruction-conditioned image editing and use textual, vision-language, or generative priors to handle complex degradations (Conde et al., 2024; Lin et al., 2024; Wang et al., 2025; Xia et al., 2023; Wu et al., 2025; Liu et al., 2025).

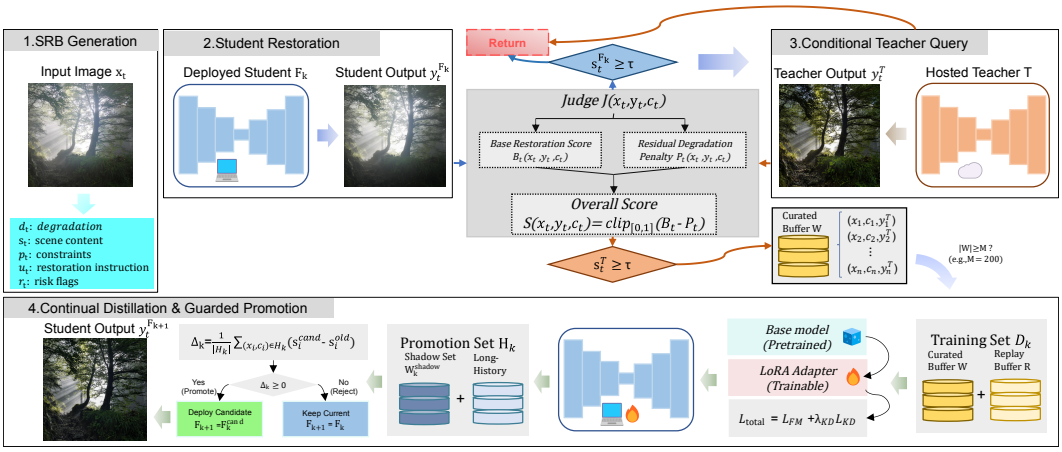

Figure 1: Overview of Q2C. The upper path serves the current request through Student-first restoration and selective Teacher fallback. Accepted Teacher outputs enter the lower path, where RDA produces a candidate Student and GSP determines whether it enters deployment.

Their strong priors have driven research on transferring capabilities to accessible restoration systems. RealRestorer constructs realistic supervision from powerful editing models to fine-tune open-source restorers (Yang et al., 2026), and GGT-100K shows that filtered generated targets improve performance across model families (Kong et al., 2026). Both construct supervision before deployment; Q2C acquires it selectively from Student failures during use.

2.3 Continual Learning and Capability Accumulation

Continual learning uses replay, regularization, and distillation to acquire new knowledge while preserving prior capabilities (De Lange et al., 2022; Buzzega et al., 2020; Smith et al., 2024; Yan et al., 2024). Replay revisits past samples, while distillation constrains changes in earlier behavior.

Recent works apply continual learning to restoration and generative tasks: MedIR investigates lifelong all-in-one medical image restoration under evolving degradation distributions (Shan et al., 2026), and ILAWR studies sequential adaptation across changing weather conditions (Lu et al., 2025). For generative models, preserving previous generation behavior has been shown to be important when updating large models (Masip et al., 2025).

Nevertheless, most existing settings assume predefined training data or task sequences. Here, supervision is generated through selective hosted-Teacher queries during deployment, coupling continual learning with query selection, quality control, and promotion. Q2C uses replay and previous-Student distillation to retain the competence acquired through these interactions.

3 Query-to-Competence

Figure 1 summarizes Q2C's serving and adaptation paths. The SRB gives the Student, Teacher, and Judge a common restoration objective and preservation constraints. The Student serves outputs that meet the quality criterion; otherwise, the Teacher provides a fallback response. Only curated Teacher outputs enter the supervision buffer. RDA trains a candidate from current and replayed examples, and GSP decides whether it replaces the deployed Student.

3.1 Usage-Time Learning Setting

Let $x_t \sim \mathcal{D}_t$ denote an image arriving at time t from a potentially nonstationary real-world restoration stream. The distribution $\mathcal{D}_t$ may change in degradation type, severity, scene content, image source, and restoration intent.

An automatic perception module P converts the open-ended restoration request into an SRB $c_t = P(x_t)$.

The currently deployed Student at usage stage k is denoted by F_k and produces $y_t^F = F_k(x_t, c_t)$.

A stronger hosted Teacher T is available through an external service and produces $y_t^T = T(x_t, c_t)$ only when the Student output does not satisfy the routing criterion.

Unlike offline distillation, the Teacher-supervision distribution is not constructed in advance. It is jointly determined by the deployment stream, the failure boundary of the current Student, and the reliability of the queried Teacher output. Denoting the routing and curation decisions by z_t and a_t, respectively, the supervision observed at stage k follows

$$\mathcal{D}_k^{\text{sup}} = \mathcal{D}_t \mid (z_t = 1,\ a_t = 1,\ F_k)\,. \tag{1}$$

Thus, as the deployed Student changes, the requests sent to the Teacher and the samples admitted for subsequent learning also change.

Under a fixed acceptance threshold τ, we describe the current quality-qualified competence of a Student by its autonomous coverage

$$A_\tau(F_k; \mathcal{X}) = \frac{1}{|\mathcal{X}|} \sum_{x_i \in \mathcal{X}} \mathbb{I}\left[S(x_i, F_k(x_i, c_i), c_i) \geq \tau\right], \tag{2}$$

where $S(\cdot)$ is the restoration score introduced below. Usage-time continual distillation aims to increase the fraction of requests satisfying this fixed quality criterion while retaining previously acquired behavior.

3.2 SRB-Guided Supervision Acquisition

Successful restoration requires both effective degradation removal and faithful content preservation. A generic instruction such as "restore this image" does not specify the observed degradation, the desired correction, or the visual content that should remain unchanged. We therefore construct an SRB shared by the Student, Teacher, and Judge:

$$c_t = (d_t,\ s_t,\ p_t,\ u_t,\ r_t)\,, \tag{3}$$

where d_t is a degradation profile containing degradation labels, severity, confidence, and localized evidence; s_t summarizes the scene content and semantic category; p_t specifies the visual structures and attributes that must be preserved; u_t describes the restoration intent and provides the image-specific restoration instruction; and r_t records potential restoration risks, such as overexposure, geometry drift, and color shift.

For an input x, an output y, and an SRB c, the Judge returns a structured output consisting of five normalized scores, residual-degradation evidence, and hard content-preservation indicators:

$$J(x, y, c) = (q_r,\ q_o,\ q_g,\ q_c,\ q_{\text{col}},\ \mathcal{E},\ \mathbf{h})\,. \tag{4}$$

The five scores, inspired by GGT-100K (Kong et al., 2026), measure restoration quality, object consistency, geometry alignment, content reasonableness, and color consistency. The set $\mathcal{E}$ contains detected residual degradations, each associated with estimated severity and confidence. The binary vector $\mathbf{h}$ records hard content-preservation failures, including identity changes, text rewriting, geometry distortion, and object insertion or removal.

The base restoration score is

$$B(x, y, c) = 0.50q_r + 0.15q_o + 0.15q_g + 0.10q_c + 0.10q_{\text{col}}. \tag{5}$$

For each residual degradation $i \in \mathcal{E}$, let ρ_i be its estimated severity multiplied by its confidence.

We aggregate residual degradation evidence as

$$P_{\text{res}} = \min\left(0.4,\ 0.25 \max_{i \in \mathcal{E}} \rho_i + 0.15\,\text{mean}_{i \in \mathcal{E}}\,\rho_i\right), \tag{6}$$

with $P_{\text{res}} = 0$ when no residual degradation is detected. The final Restoration–Preservation Score (RPS) is

$$S(x, y, c) = \text{clip}_{[0,1]}\left(B(x, y, c) - P_{\text{res}}\right). \tag{7}$$

Student-first routing. The deployed Student processes every request first. Its score is $s_t^F = S(x_t, y_t^F, c_t)$. The routing decision is

$$z_t = \mathbb{I}\left[s_t^F < \tau_{\text{route}}\right]. \tag{8}$$

When $z_t = 0$, the Student output is served without a Teacher call. When $z_t = 1$, the hosted Teacher is invoked as the fallback for the current request. We use the same routing threshold across usage stages so that a change in autonomous coverage reflects a change in Student capability rather than a relaxed acceptance rule.

Teacher-output curation. A Teacher output is not automatically a training target. Although it may remove the requested degradation, it can also alter image content or introduce plausible but unsupported details. We therefore evaluate $s_t^T = S(x_t, y_t^T, c_t)$ and admit the Teacher result only if it satisfies both the restoration-score criterion and the hard content-preservation criterion:

$$a_t = \mathbb{I}\left[z_t = 1\right] \mathbb{I}\left[s_t^T \geq \tau_{\text{cur}}\right] \mathbb{I}\left[\|\mathbf{h}_t^T\|_1 = 0\right]. \tag{9}$$

Only when $a_t = 1$ is the tuple $e_t^{\text{sup}} = (x_t, c_t, y_t^T, s_t^T)$ added to the current supervision window $\mathcal{W}_k$.

The supervision window therefore contains curated Teacher successes on requests that the Student could not serve.

3.3 Event-Triggered Replay-Distill Adaptation

To avoid updating the model after each individual Teacher output, we trigger RDA when the supervision window reaches M accepted pairs, $|\mathcal{W}_k| = M$. At each update event, the accepted window is divided into a training subset and a shadow subset:

$$\mathcal{W}_k = \mathcal{W}_k^{\text{train}} \dot{\cup} \mathcal{W}_k^{\text{shadow}}, \tag{10}$$

$$|\mathcal{W}_k^{\text{train}}| = 4 \times |\mathcal{W}_k^{\text{shadow}}|. \tag{11}$$

The shadow subset does not participate in the current gradient update and is reserved for promotion.

Let $\mathcal{M}_k$ denote a bounded replay memory constructed from previously accepted usage-time experience. Candidate training uses $\mathcal{D}_k = \mathcal{W}_k^{\text{train}} \cup \text{Sample}(\mathcal{M}_k)$. The current window concentrates on the latest Student failure boundary, whereas the replay set maintains exposure to previously acquired restoration capabilities.

The Student is an instruction-conditioned image-editing model F_{θ,ϕ_k}, where θ denotes the frozen base parameters and ϕ_k denotes the deployed LoRA parameters (Hu et al., 2022). A candidate adapter is initialized as $\phi_k^{\text{cand}} \leftarrow \phi_k$; only the LoRA parameters are updated.

For both current and replay samples, we use the native flow-matching objective of the Student backbone (Lipman et al., 2023). Given a target latent z_0, Gaussian noise ϵ, and interpolation state σ, let $z_\sigma = (1-\sigma)z_0 + \sigma\epsilon$ and $v^* = \epsilon - z_0$ denote the interpolated latent and target velocity. The adaptation loss is

$$\mathcal{L}_{\text{FM}} = \mathbb{E}_{(x,c,y^*)\sim\mathcal{D}_k}\left[w(\sigma)\left\|v_{\theta,\phi_k^{\text{cand}}}(z_\sigma, x, c, \sigma) - v^*\right\|_2^2\right], \tag{12}$$

where y^* is the curated Teacher pseudo-target used to obtain z_0.

Replay supervision alone does not directly constrain changes in the Student's historical generation function. We therefore retain a frozen copy of the previously deployed Student and perform function distillation only on replay inputs. Under the same restoration brief, noisy latent, and interpolation state, the distillation loss is

$$\begin{aligned}\mathcal{L}_{\text{KD}} = \mathbb{E}_{(x,c,\cdot)\sim\mathcal{M}_k}\Big[\Big\|&v_{\theta,\phi_k^{\text{cand}}}(z_\sigma, x, c, \sigma) \\ &- \text{sg}\left[v_{\theta,\phi_k}(z_\sigma, x, c, \sigma)\right]\Big\|_1\Big],\end{aligned} \tag{13}$$

where $\text{sg}[\cdot]$ stops gradients through the previous Student. The complete objective is $\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{FM}} + \lambda_{\text{KD}}\mathcal{L}_{\text{KD}}$. The distillation term is omitted for the first update when no previous LoRA adapter is available. In subsequent updates, it constrains functional drift on replayed historical inputs without directly suppressing adaptation to the newly accepted window.

After optimization, the resulting candidate is $F_k^{\text{cand}} = F_{\theta,\phi_k^{\text{cand}}}$. Candidate training alone does not modify the deployed Student.

3.4 GUARDED STUDENT PROMOTION AND MODEL EVOLUTION

Continual adaptation is not necessarily monotonic. A candidate can fit the latest supervision window while degrading previously learned restoration behavior. We therefore use GSP to separate model optimization from deployment.

The promotion set combines the held-out current shadow subset with a bounded long-history set:

$$\mathcal{H}_k = \mathcal{W}_k^{\text{shadow}} \cup \text{LongHistory}(\mathcal{M}_k). \tag{14}$$

For every $(x_i, c_i) \in \mathcal{H}_k$, the candidate produces $y_i^{\text{cand}} = F_k^{\text{cand}}(x_i, c_i)$. The deployed Student produces $y_i^{\text{old}} = F_k(x_i, c_i)$ for the paired comparison. Their paired score difference is

$$\delta_i = S(x_i, y_i^{\text{cand}}, c_i) - S(x_i, y_i^{\text{old}}, c_i). \tag{15}$$

The promotion statistic is the mean paired improvement

$$\Delta_k = \frac{1}{|\mathcal{H}_k|} \sum_{(x_i, c_i) \in \mathcal{H}_k} \delta_i. \tag{16}$$

The candidate is promoted only when $\Delta_k > 0$. Accordingly, the next deployed Student is

$$F_{k+1} = \begin{cases} F_k^{\text{cand}}, & \Delta_k > 0, \\ F_k, & \Delta_k \leq 0. \end{cases} \tag{17}$$

Routing decides "serve or query", curation decides "store or discard", and promotion decides "deploy or rollback". They evaluate different outputs, and promotion uses a separate shadow-plus-history set.

After promotion or rollback, the historical memory is updated. Samples on which the candidate regresses relative to the deployed Student are prioritized as regression replay. Strong candidate wins from the shadow set enter the long-history pool, while the remaining capacity is maintained across semantic groups. When promotion fails, the deployed Student remains unchanged and the rejected window is retained for later replay.

4 EXPERIMENTS

4.1 DATASETS AND EVALUATION METRICS

Usage stream. Our source image pool consists of images from LSDIR (Li et al., 2023), Pexels, and Pixabay; full details on collection, filtering, provenance, and licensing are provided in the supplementary material. The experimental corpus includes 4,454 audited input–GPT-Image-2-output pairs. Samples are partitioned by source image identity to ensure no shared-source variants appear in both continual adaptation and evaluation, yielding a 4,254-image usage stream.

The seven sessions are naturally accumulated from real deployment and vary in degradation composition. Figure 2 shows seven non-exclusive atomic factors: blur, haze, low light, noise, old-photo damage, rain, and snow. Their prevalences do not sum to 100% because degradations can co-occur.

SDCR-200. We hold out 200 source-disjoint images from the combined source pool to form the Source-Disjoint Compound Restoration Set (SDCR-200). The set is balanced across five compound restoration domains: blur with low-light detail, low light with noise, haze, rain and adverse weather, and old-photo, snow, and rare compound degradations. Each domain contains 40 images. The right panel of Figure 2 shows that these domains contain overlapping atomic degradations rather than isolated synthetic tasks.

SDCR-200 remains unchanged throughout the complete usage trajectory. It is used only for reporting and is excluded from routing, Teacher-output curation, candidate training, replay sampling, threshold selection, and GSP.

Evaluation metrics. We use the Restoration–Preservation Score (RPS) defined in Section 3.2 for task-aware evaluation. RPS measures degradation removal together with object, geometry, content, and color preservation.

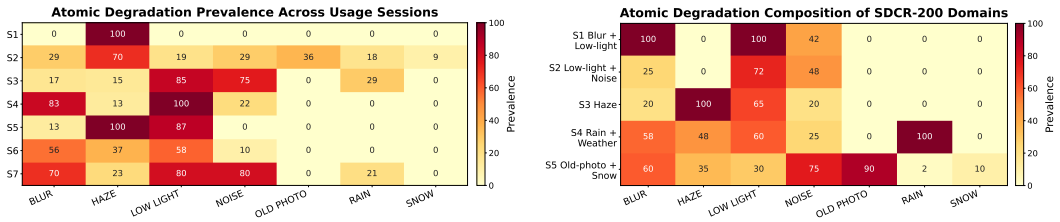

Figure 2: Atomic degradation composition of the experimental data. Left: degradation prevalence in the seven usage sessions naturally accumulated during deployment. Right: atomic composition of the five SDCR-200 domains. Individual images may contain multiple simultaneous degradations.

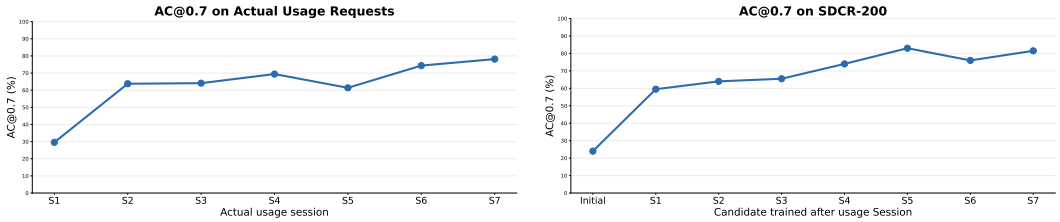

Figure 3: Teacher fallback and Student competence across seven usage sessions. Left: replacement rate on actual usage requests. Right: stage-wise candidates on the unchanged SDCR-200. Both panels use the fixed RPS ≥ 0.7 criterion.

Autonomous Coverage (AC@0.7) is the percentage of outputs whose RPS reaches 0.7 and can therefore be served without Teacher fallback. We report it on the unchanged SDCR-200 for stage-wise comparisons across domains and on each actual usage session for Student-only coverage on the requests encountered at that time.

We additionally report six no-reference image quality metrics: NIQE (Mittal et al., 2013), MUSIQ (Ke et al., 2021), MANIQA (Yang et al., 2022), TOPIQ-NR (Chen et al., 2024), BRISQUE (Mittal et al., 2012), and CLIP-IQA (Wang et al., 2023). PSNR, SSIM (Wang et al., 2004), and LPIPS (Zhang et al., 2018) computed against curated GPT-Image-2 outputs are reported in the ablation study as teacher-target fidelity measures.

4.2 Experimental Settings

We adopt Qwen-Image-Edit-2511 as the initial self-hosted Student and GPT-Image-2 as the hosted Teacher for usage-time supervision (Wu et al., 2025; OpenAI, 2026). The Student is adapted via rank-8 LoRA with a learning rate of 5×10^{-4} (Hu et al., 2022). The replay memory stores up to 200 samples, and previous-Student distillation is applied on replay inputs with $\lambda_{\mathrm{KD}} = 0.1$. Student inference uses eight generation steps.

An adaptation event is triggered once 200 curated Teacher pairs have accumulated: 160 for candidate training and 40 as the shadow subset for GSP. Standard updates run 1,500 optimization steps. If promotion fails, the subsequent update is extended to 3,000 steps with temporarily doubled replay memory capacity.

Requests in each usage session are processed by the Student deployed at that stage; collected Teacher outputs are retained for subsequent adaptation. After each session, we evaluate the resulting candidate on the unchanged SDCR-200 and report AC@0.7 on that session's requests using the Student deployed at the time. Both evaluations use the fixed RPS threshold of 0.7.

4.3 Quantitative Results

Teacher fallback diminishes as Student competence accumulates. Figure 3 tracks Student-only completion on actual usage requests alongside stage-wise candidate coverage on the unchanged SDCR-200. Both panels use the fixed RPS ≥ 0.7 criterion.

Table 1: Student-only completion across the seven usage sessions under the fixed RPS ≥ 0.7 routing criterion. The Student-only rate is computed on requests received in each session.

Session	Requests	Student-only	Rate	Teacher calls
S1	284	84	29.58%	200
S2	552	352	63.77%	200
S3	557	357	64.09%	200
S4	654	454	69.42%	200
S5	518	318	61.39%	200
S6	780	580	74.36%	200
S7	909	711	78.22%	198
Total	4,254	2,856	67.14%	1,398

Table 2: RDA ablation on SDCR-200. PSNR, SSIM, and LPIPS use curated Teacher targets; red and blue mark the best and second-best adapted variants.

Variant	RPS ↑	AC@0.7 ↑	PSNR ↑	SSIM ↑	LPIPS ↓
Initial Student	0.5575	24.0	17.6624	0.5302	0.3491
Sequential	**0.7490**	80.0	21.1114	0.6927	0.2021
Replay only	0.7453	**81.0**	**21.3202**	**0.6959**	0.2054
Distillation only	0.7471	79.0	20.7980	0.6763	**0.2019**
Full RDA	**0.7532**	**81.5**	**21.3892**	**0.7091**	**0.1983**

Table 3: Comparison on the 200 SDCR-200 images. Red and blue denote the best and second-best results, respectively. Ours denotes the Stage-7 candidate. Because the evaluation targets are curated GPT-Image-2 outputs, the Teacher's 100.0% AC@0.7 is selection-conditioned.

	Task-Aware Evaluation		No-Reference Image Quality					
Method	RPS ↑	AC@0.7 ↑	NIQE ↓	MUSIQ ↑	MANIQA ↑	TOPIQ-NR ↑	BRISQUE ↓	CLIP-IQA ↑
Qwen-Image-Edit-2511	0.5575	24.0%	5.5286	49.9790	0.3096	0.3619	49.3134	0.5087
Step1X-Edit-v1p2	0.6045	31.0%	5.3240	47.6512	0.2731	0.3331	39.7289	0.4149
Seedream 4.5	0.6665	52.8%	**4.1167**	58.4001	0.3354	**0.4941**	**18.8385**	**0.6358**
Nano Banana 2	0.7145	68.0%	**4.4602**	53.7135	0.3093	0.3784	32.8789	0.4852
GPT-Image-2	**0.8124**	**100.0%**	4.6191	**63.7704**	**0.3721**	**0.5167**	**17.2581**	**0.6337**
Ours	**0.7532**	**81.5%**	4.8084	**60.2776**	**0.3508**	0.4344	38.9432	0.5939

On the unchanged SDCR-200, AC@0.7 increases from 24.0% for the initial Student to 81.5% for the Stage-7 candidate. This 57.5-point gain under a fixed criterion shows that successive Q2C updates accumulate reusable competence across the five compound domains.

Table 1 makes the operational trend explicit: Student-only completion rises from 29.58% in S1 to 78.22% in S7, with a temporary decrease at S5. Under the fixed routing criterion, this means fewer Teacher calls per request as usage continues. Together, the actual-usage and fixed-set trajectories show that usage-time Teacher interactions become reusable Student competence, while the fluctuation highlights the value of guarded promotion under evolving request distributions.

Across all 4,254 requests, the Student completes 2,856 without Teacher fallback. The selective policy therefore selects 1,398 Teacher responses rather than 4,254 under an always-Teacher policy, corresponding to 67.14% fewer hosted calls and 67.44% fewer provider-reported tokens (Appendix B.1). Because Teacher outputs were precomputed for the study, these figures compare serving policies rather than the API calls made to prepare the dataset.

The S6 update illustrates the role of GSP. Its candidate drops from 83.0% to 76.0% AC@0.7 on the report-only SDCR-200. On the separate promotion set, its mean RPS change is $+0.0224$ on the current shadow subset but -0.0630 on long-history samples, with 34 of 40 historical comparisons regressing. The combined delta is -0.0203, so GSP rejects the candidate and leaves the Stage-5 Student deployed (Appendix B.2).

Comparison with image-editing models.

Table 3 compares the Stage-7 candidate with its Qwen-Image-Edit-2511 initialization and representative self-hosted and hosted image-editing models (Wu et al., 2025; Liu et al., 2025; ByteDance Seed, 2025; Google DeepMind, 2026; OpenAI, 2026). All methods receive the same SDCR-200 inputs and image-specific SRBs.

The Stage-7 candidate achieves an RPS of 0.7532 and AC@0.7 of 81.5%, ranking behind only GPT-Image-2 on the two task-aware metrics and first among the evaluated non-Teacher models. Relative to the initial Student, it also improves all six no-reference metrics.

The hosted editors exhibit different no-reference metric profiles: Seedream 4.5 leads several individual perceptual measures, while GPT-Image-2 leads the task-aware evaluation.



Figure 4: Qualitative comparison on representative SDCR-200 examples. Columns show the input, Qwen-Image-Edit-2511, Step1X-Edit-v1p2, Doubao Seedream 4.5, Nano Banana 2, GPT-Image-2, and ours. The two rows show a low-light and blur case and a severe compound case containing rain, low light, haze, noise, and blur.

4.4 Qualitative Results

Figure 4 compares two compound restoration cases. In the low-light and blur example, the adapted candidate improves visibility and local detail relative to the initialized Student. In the more severe example, it reduces multiple weather and imaging degradations while retaining the main road, building, and object layout.

4.5 Ablation Study

We evaluate the replay and previous-Student distillation components of RDA under the same seven-session usage trajectory. The Sequential variant uses only the current supervision window. Replay only additionally uses historical Teacher pseudo-target pairs. Distillation only uses retained historical inputs without their Teacher targets and constrains the candidate to match the previous frozen Student on those inputs. Full RDA combines both mechanisms.

All adapted variants outperform the initial Student. Sequential training on selected Teacher outputs raises AC@0.7 from 24.0% to 80.0%, while Full RDA reaches 81.5%. Thus, Teacher-supervised adaptation drives the large gain over the initial Student, and replay with previous-Student distillation adds 1.5 percentage points of coverage and 0.0042 RPS at the shared endpoint.

Full RDA achieves the best result on every reported metric. Replay only improves AC@0.7, PSNR, and SSIM over Sequential, while Distillation only obtains a lower LPIPS but weaker coverage and structural fidelity than Replay only.

5 Conclusion

We propose Q2C, an adaptive restoration system that turns selective hosted-Teacher fallback into reusable competence for a self-hosted Student. SRB-guided routing and curation identify useful supervision, while RDA and GSP adapt the Student and guard its deployment across usage sessions. On SDCR-200, Student-only coverage rises from 24.0% to 81.5%; in the usage stream, the Student completes an increasing share of requests without Teacher fallback. These results show how an updateable local model can improve through use while retaining control over when new capabilities enter service.

REFERENCES

Yuang Ai, Huaibo Huang, Xiaoqiang Zhou, Jiexiang Wang, and Ran He. Multimodal prompt perceiver: Empower adaptiveness, generalizability and fidelity for all-in-one image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 25432–25444, 2024.

Pietro Buzzega, Matteo Boschini, Angelo Porrello, Davide Abati, and Simone Calderara. Dark experience for general continual learning: A strong, simple baseline. In *Advances in Neural Information Processing Systems*, volume 33, 2020.

ByteDance Seed. Seedream 4.5. Official model documentation, 2025. URL `https://seed.bytedance.com/en/seedream4_5`. Accessed: 2026-07-26.

Chaofeng Chen, Jiadi Mo, Jingwen Hou, Haoning Wu, Liang Liao, Wenxiu Sun, Qiong Yan, and Weisi Lin. Topiq: A top-down approach from semantics to distortions for image quality assessment. *IEEE Transactions on Image Processing*, 33:2404–2418, 2024.

Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *European Conference on Computer Vision*, pp. 17–33, 2022. doi: 10.1007/978-3-031-20071-7_2.

Zihan Cheng, Liangtai Zhou, Dian Chen, Ni Tang, Xiaotong Luo, Yuan Xie, and Yanyun Qu. Unildiff: Unlocking the power of diffusion priors for all-in-one image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 37465–37475, 2026.

Marcos V. Conde, Gregor Geigle, and Radu Timofte. InstructIR: High-quality image restoration following human instructions. In *European Conference on Computer Vision*, volume 15094 of *Lecture Notes in Computer Science*, pp. 1–21, 2024. doi: 10.1007/978-3-031-72764-1_1.

Matthias De Lange, Rahaf Aljundi, Marc Masana, Sarah Parisot, Xu Jia, Ales Leonardis, Gregory Slabaugh, and Tinne Tuytelaars. A continual learning survey: Defying forgetting in classification tasks. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(7):3366–3385, 2022.

Google DeepMind. Gemini 3.1 flash image: Nano banana 2. Google DeepMind model documentation, 2026. URL `https://deepmind.google/models/gemini-image/flash/`. Accessed: 2026-07-26.

Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. In *International Conference on Learning Representations*, 2022.

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, 2026.

Boyun Li, Xiao Liu, Peng Hu, Zhongqin Wu, Jiancheng Lv, and Xi Peng. All-in-one image restoration for unknown corruption. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 17452–17462, 2022.

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.

Jingbo Lin, Zhilu Zhang, Yuxiang Wei, Dongwei Ren, Dongsheng Jiang, Qi Tian, and Wangmeng Zuo. Improving image restoration through removing degradations in textual representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2866–2878, 2024.

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. Flow matching for generative modeling. In *International Conference on Learning Representations*, 2023.

Shiyu Liu, Yucheng Han, Peng Xing, Fukun Yin, Rui Wang, Wei Cheng, et al. Step1x-edit: A practical framework for general image editing, 2025.

Xin Lu, Jie Xiao, Yurui Zhu, and Xueyang Fu. Continuous adverse weather removal via degradation-aware distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 28113–28123, 2025.

Wenyang Luo, Haina Qin, Zewen Chen, Libin Wang, Dandan Zheng, Yuming Li, Yufan Liu, Bing Li, and Weiming Hu. Visual-instructed degradation diffusion for all-in-one image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12764–12777, 2025.

Sergi Masip, Pau Rodriguez, Tinne Tuytelaars, and Gido M. van de Ven. Continual learning of diffusion models with generative distillation. In *Proceedings of The 3rd Conference on Lifelong Learning Agents*, volume 274 of *Proceedings of Machine Learning Research*, pp. 431–456, 2025.

Anish Mittal, Anush Krishna Moorthy, and Alan C. Bovik. No-reference image quality assessment in the spatial domain. *IEEE Transactions on Image Processing*, 21(12):4695–4708, 2012.

Anish Mittal, Rajiv Soundararajan, and Alan C. Bovik. Making a completely blind image quality analyzer. *IEEE Signal Processing Letters*, 20(3):209–212, 2013.

OpenAI. Gpt image 2 model. OpenAI API documentation, 2026. URL `https://developers.openai.com/api/docs/models/gpt-image-2`. Accessed: 2026-07-26.

Pexels. Pexels license. `https://www.pexels.com/legal-pages/license/`, 2026. Accessed: 2026-07-27.

Pixabay. Pixabay content license. `https://pixabay.com/service/license-summary/`, 2026. Accessed: 2026-07-27.

Vaishnav Potlapalli, Syed Waqas Zamir, Salman Khan, and Fahad Shahbaz Khan. Promptir: Prompting for all-in-one image restoration. In *Advances in Neural Information Processing Systems*, volume 36, 2023.

Sudarshan Rajagopalan, Nithin Gopalakrishnan Nair, Jay N. Paranjape, and Vishal M. Patel. Gendeg: Diffusion-based degradation synthesis for generalizable all-in-one image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 28144–28154, 2025.

Shihao Shan, Hongying Liu, Fanhua Shang, Liang Wan, and Jingjing Deng. Beyond the static-world: Lifelong learning for all-in-one medical image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13702–13711, 2026.

James Seale Smith, Lazar Valkov, Shaunak Halbe, Vyshnavi Gutta, Rogerio Feris, Zsolt Kira, and Leonid Karlinsky. Adaptive memory replay for continual learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 3605–3615, 2024.

Xiangpeng Tian, Xiangyu Liao, Xiao Liu, Meng Li, and Chao Ren. Degradation-aware feature perturbation for all-in-one image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 28165–28175, 2025.

Unsplash. Unsplash license. `https://unsplash.com/license`, n.d. Accessed July 30, 2026.

Chao Wang, Hehe Fan, Huichen Yang, Sarvnaz Karimi, Lina Yao, and Yi Yang. Adapting text-to-image generation with feature difference instruction for generic image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 23539–23550, 2025.

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

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*, pp. 1905–1914, 2021.

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.

Chenfei Wu, Jiahao Li, Jingren Zhou, Junyang Lin, Kaiyuan Gao, et al. Qwen-image technical report, 2025.

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 *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13095–13105, 2023.

Fengyang Xiao, Peng Hu, Lei Xu, XingE Guo, Guanyi Qin, Yuqi Shen, Chengyu Fang, Rihan Zhang, Chunming He, and Sina Farsiu. Beyond ground-truth: Leveraging image quality priors for real-world image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 29897–29908, 2026.

Hongwei Yan, Liyuan Wang, Kaisheng Ma, and Yi Zhong. Orchestrate latent expertise: Advancing online continual learning with multi-level supervision and reverse self-distillation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 23670–23680, 2024.

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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 1191–1200, 2022.

Wenbo Yang, Zhongling Wang, and Zhou Wang. Towards a universal image degradation model via content-degradation disentanglement. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12966–12975, 2025.

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, 2026.

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

Haijin Zeng, Xiangming Wang, Yongyong Chen, Jingyong Su, and Jie Liu. Vision-language gradient descent-driven all-in-one deep unfolding networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7524–7533, 2025.

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 Conference on Computer Vision and Pattern Recognition*, pp. 586–595, 2018.

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

Mo Zhou, Keren Ye, Mauricio Delbracio, Peyman Milanfar, Vishal M. Patel, and Hossein Talebi. Unires: Universal image restoration for complex degradations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 13237–13247, 2025.

AI USE STATEMENT

We used generative AI tools in two roles. Large language models assisted with language polishing and limited bibliographic search. GPT-Image-2(OpenAI, 2026) served as the hosted Teacher and generated audited restoration outputs used as pseudo-references for usage-time supervision and teacher-target fidelity evaluation. Generative AI tools were not used to design the methodology, analyze the experimental results or determine the conclusions. All AI-assisted text and generated artifacts were reviewed and verified by the authors.

A EXPERIMENTAL SETUP AND DATA DOCUMENTATION

A.1 USAGE-TIME PROTOCOL

We process the 4,254 non-test requests in seven ordered sessions. The Student deployed at the beginning of a session handles every request first. A request with $\mathrm{RPS} \geq 0.7$ is served locally; otherwise, Q2C follows the Teacher-fallback route. The Teacher output is then evaluated for possible inclusion in the supervision buffer. We precompute and cache hosted outputs so that all continual-learning variants use the same hosted-service version and repeated ablations avoid duplicate calls. Precomputation changes only the timing of the remote computation: a cached output remains inaccessible unless the deployed Student fails the fixed gate, and an accepted output can affect only later requests. An update is normally triggered once 200 accepted Teacher pairs have accumulated; 160 are used for candidate training and 40 form the shadow set. The final partial window contains 197 accepted pairs, of which 158 are used for training and 39 form the shadow set.

A.2 TRAINING AND PROMOTION CONFIGURATION

Each candidate is initialized from the deployed LoRA parameters, while the Qwen-Image-Edit-2511 base weights remain frozen. Stage 1 omits previous-Student distillation because no earlier adapter exists; later stages apply the distillation loss only to replay inputs. Table A.1 reports the resolved configuration. We save the final optimization checkpoint and do not select checkpoints on SDCR-200.

Table A.1: Student adaptation and inference configuration.

Item	**Value**	**Item**	**Value**
Trainable parameters	LoRA only	Rank / alpha / dropout	8 / 8 / 0
Target modules	`to_q`, `to_k`, `to_v`, `to_out.0`	Learning rate	5×10^{-4}
Training resolution	512×512	Precision	bfloat16
Batch size / accumulation	1 / 1	Random seed	42
Standard update	1,500 steps	Post-rejection retry	3,000 steps
Replay capacity	200	Temporary capacity after rejection	400
Current / replay sampling	0.5 / 0.5	Distillation weight	0.1
Student inference	8 steps	True-CFG scale	4.0

For GSP, we compare the candidate with the deployed Student on a promotion set comprising the held-out current shadow set and 40 long-history samples. The candidate is promoted only when its mean paired RPS improvement is strictly positive. Regression cases are prioritized when updating replay memory. If promotion fails, the deployed Student remains unchanged, the difficult window is retained, and the next update uses the longer schedule and temporary replay capacity in Table A.1.

A.3 Data Construction and Provenance

Our corpus contains 4,454 audited input–Teacher-output records: 3,169 from Pexels, 1,129 from Pixabay, 139 derived from LSDIR, and 17 from Unsplash. We use source-image identity rather than file path to partition the corpus, preventing degraded or restored variants of one source from crossing between continual adaptation and SDCR-200.

The constructed dataset will be made publicly available upon publication for scientific research use only. The release will include the audited data splits, source identifiers, degradation and scene annotations, restoration prompts, session assignments, and curated Teacher pseudo-references. The author-created materials will be released under a research-only license. Third-party photographs remain subject to the providers' respective terms (Li et al., 2023; Pexels, 2026; Pixabay, 2026; Unsplash, n.d.). Commercial use of the released research dataset will not be permitted.

Table A.2: Source provenance of the experimental corpus.

Original provider	**Groups**	**Fraction**
Pexels	3,169	71.15%
Pixabay	1,129	25.35%
LSDIR-derived	139	3.12%
Unsplash	17	0.38%
Total	4,454	100.00%

After holding out SDCR-200, the remaining 4,254 examples form the seven-session usage stream. The seven ordered sessions contain changing mixtures of low light, blur, haze, noise, rain, snow, and old-photo damage, causing later updates to encounter different competence boundaries. Multiple labels may occur in one image. SDCR-200 contains 40 source-disjoint examples from each of five compound domains: blur with low-light detail; low light with noise; haze; rain and adverse weather; and old-photo damage, snow, and rare compound degradations. Its curated Teacher outputs serve as pseudo-references; therefore, paired PSNR, SSIM, and LPIPS measure Teacher-target fidelity rather than fidelity to genuine clean ground truth.

A.4 Baselines and Evaluation Protocol

All image-editing models receive the same input and image-specific restoration prompt. The Qwen-Image-Edit-2511 baseline and all Q2C Student checkpoints use eight inference steps, a true-CFG scale of 4.0, and a blank negative prompt. Step1X-Edit-v1p2 uses 28 inference steps and a true-CFG scale of 6.0, with thinking and reflection enabled. Seedream 4.5 requests 2K output; Nano Banana 2 requests 1K output; GPT-Image-2 uses high quality and automatic output size. When necessary, outputs are resized to the input resolution of 1254×1254 before evaluation.

RPS and AC@0.7 are computed from the saved structured Judge records. NIQE, MUSIQ, MANIQA, TOPIQ-NR, BRISQUE, and CLIP-IQA are computed on RGB inputs using their standard implementations. For the paired Initial-Student versus Stage-7 analysis on the same 200 images, we use two-sided Wilcoxon tests for continuous metrics and the exact McNemar test for AC@0.7. We compute 95% confidence intervals from 10,000 paired bootstrap resamples with seed 20260726 and apply Holm correction jointly across the eight primary metrics.

A.5 Computing Infrastructure and Software

All local Student training, inference, promotion evaluation, and open-source metric computation were performed on one server equipped with NVIDIA A800 80GB PCIe GPUs. The server ran Ubuntu 22.04.5 LTS with Linux kernel 5.15.0-119-generic. Each LoRA candidate-training job used one GPU; independent inference and promotion evaluation workers could use up to two GPUs concurrently. No individual training job was distributed across multiple GPUs.

Table A.3: Computing infrastructure and software environment used for the reported local experiments. Commercial hosted services execute on provider-managed infrastructure.

Item	Version / specification	Item	Version / specification
GPU	NVIDIA A800 80GB PCIe	CPU	Intel Xeon Silver 4314
System memory	512 GiB (503 GiB OS-visible)	Operating system	Ubuntu 22.04.5 LTS
Linux kernel	5.15.0-119-generic	NVIDIA driver / CUDA	590.48.01 / 12.4
Python	3.10.19	PyTorch / torchvision	2.5.0+cu124 / 0.20.0
diffusers	0.39.0.dev0	transformers	4.57.1
Accelerate / PEFT	1.5.2 / 0.19.1	safetensors	0.8.0
PyIQA / LPIPS	0.1.12 / 0.1.4	NumPy / SciPy	1.26.0 / 1.15.2
scikit-image	0.25.2	OpenCV / Pillow	4.11.0.86 / 10.4.0
pandas	2.2.3	Training GPU usage	One GPU per candidate

B Additional Results

B.1 Serving Policy, Cost, and Latency

We distinguish the models from the serving policies before comparing cost. The *Local Student* is Qwen-Image-Edit-2511 with the LoRA checkpoint deployed for the corresponding session, and the *Hosted Teacher* is GPT-Image-2. Q2C is the complete framework, not a separate image-editing checkpoint. At inference time it uses a *selective-fallback policy*: the deployed Student processes each request locally, and GPT-Image-2 is invoked only when the local output has $\mathrm{RPS} < 0.7$. The *hosted-only baseline* invokes GPT-Image-2 for every request and performs no local restoration. Within Q2C, a *local completion* ends after the Student passes the criterion, whereas a *Teacher fallback* contains both the local attempt and one hosted call.

B.1.1 Overall Policy Comparison

Across 4,254 requests, Q2C runs the Local Student 4,254 times and invokes the Hosted Teacher 1,398 times; the hosted-only baseline invokes the Teacher 4,254 times. Q2C therefore avoids 2,856 hosted calls (67.14%) and 16,279,842 provider-reported tokens (67.44%). Including all 23.71 hours of Local-Student inference, Q2C requires 42.94 hours of combined serial-equivalent generation work, compared with 59.59 hours for the hosted-only baseline, a net reduction of 16.65 hours (27.93%).

Table B.1: Serving-policy comparison on the same 4,254 requests. Latencies are reported in seconds. Q2C generation work includes Local-Student inference for every request and Hosted-Teacher inference for fallback requests. Total hours are sums of per-request generation durations and therefore represent serial-equivalent work rather than parallel wall-clock makespan.

(a) Model calls, generation work, and hosted usage

Policy	Local	Hosted	Work (h)	Tokens
Hosted only	0	4,254	59.59	24,139,311
Q2C	4,254	1,398	42.94	7,859,469
Difference	$+4,254$	$-2,856$	-16.65	$-16,279,842$

(b) Per-request latency

Policy	Mean	Median	P95
Hosted only	50.43	29.00	31.00
Q2C	36.34	20.00	50.00
Difference	-14.09	-9.00	$+19.00$

B.1.2 Route-Level Recorded Restoration Time

We report *recorded restoration time* rather than complete end-to-end latency. For a local completion, the recorded interval contains only Local-Student generation from dispatch to completion. For a Teacher fallback, it contains the Local-Student generation time plus the recorded GPT-Image-2 provider turnaround. Perception, Judge evaluation, one-time model preload, and queueing outside the stored provider interval are excluded from both routes.

Table B.2 separates the two Q2C routes. Among the 4,254 requests, 2,856 (67.14%) require no Hosted-Teacher call after local evaluation. Their recorded restoration time is 20.11 seconds on average. The remaining 1,398 requests (32.86%) invoke GPT-Image-2 after the local attempt; the two recorded generation intervals sum to 69.50 seconds on average.

Because the local route accounts for more than half of all requests, the median over the complete Q2C stream is 20 seconds. The P95 is 50 seconds because fallback requests execute the local and hosted restoration components sequentially.

Table B.2: Recorded restoration time and GPT-Image-2 provider usage for the two Q2C routes. Fallback time sums Local-Student generation and Hosted-Teacher turnaround and excludes Perception and Judge evaluation. Provider tokens count GPT-Image-2 usage only.

Q2C route	Requests	Fraction	Recorded restoration components	Mean (s)	Median (s)	P95 (s)	Provider tokens/request
Local completion	2,856	67.14%	Local Student	20.11	20.00	21.00	0.0
Teacher fallback	1,398	32.86%	Local Student → Hosted Teacher	69.50	49.00	51.00	5,621.9
All Q2C requests	4,254	100.00%	Mixture of the two routes	36.34	20.00	50.00	1,847.5

B.1.3 Local-Student Inference Profile

Every request is first processed by the persistent Qwen-Image-Edit Local Student, and its dispatch records cover all 4,254 requests. Table B.3 reports two different local measurements. Panel (a) characterizes the time spent in image generation. Panel (b) characterizes the length of the SRB-derived restoration instruction encoded by the local Qwen tokenizer.

The local model does not expose a provider-style total-token usage field. Prompt-text tokens can be counted exactly. Image conditioning and image generation, however, operate through visual representations and continuous diffusion latents rather than a single provider-comparable token unit. Accordingly, the prompt-length statistics below include only local text conditioning; they exclude visual representations and diffusion latent computation and are not billing units.

Table B.3: Local-Student generation-time and text-conditioning profiles. Panel (a) measures Student dispatch-to-completion time. Its stream total is the sum of the 4,254 per-request durations and represents 23.71 serial-equivalent local generation hours. Panel (b) measures only the SRB-derived restoration-prompt length under the local Qwen tokenizer. Its stream total is the sum of prompt-text tokens across the same requests; it does not include visual representations, diffusion latent computation, or provider usage. The persistent worker uses eight diffusion steps and does not reload the base model for each request.

(a) Local generation duration (seconds per request)

Calls	Mean	Median	P95	Range	Stream total
4,254	20.06	20	21	18–29	23.71 h

(b) Restoration-prompt length (local text tokens per request)

Calls	Mean	Median	P95	Range	Stream total
4,254	132.6	131	165	79–214	564,280 tokens

B.1.4 Hosted-Provider Token Accounting

Table B.4 reports a different accounting system: provider-reported GPT-Image-2 usage. Except for the hosted-call row, every entry in the table is a multimodal token quantity returned by the provider. The always-Teacher policy invokes GPT-Image-2 on all 4,254 requests. Q2C invokes it only for the 1,398 requests routed to Teacher fallback. Consequently, the 2,856 local completions contribute zero hosted calls and zero GPT-Image-2 provider tokens.

The 564,280 local prompt-text tokens in Table B.3 are not included in this table. Local Qwen text tokens and GPT-Image-2 provider units are produced by different tokenizers and accounting rules. Moreover, the provider totals include input-text, input-image, and output-image units, whereas the local count covers prompt text only. The two systems therefore cannot be added or compared as equivalent token quantities.

Table B.4: Hosted GPT-Image-2 calls and provider-reported multimodal tokens over the same 4,254-request stream. Always Teacher invokes GPT-Image-2 for every request, while the Q2C column contains only the 1,398 fallback calls.

Accounting item	Always Teacher	Q2C fallback	Avoided by Q2C	Reduction
Hosted calls	4,254	1,398	2,856	67.14%
Input-text tokens	13,311,864	4,468,512	8,843,352	66.43%
Input-image tokens	4,671,342	1,510,357	3,160,985	67.67%
Output-image tokens	6,156,105	1,880,600	4,275,505	69.45%
All provider tokens	**24,139,311**	**7,859,469**	**16,279,842**	**67.44%**

Selective fallback avoids 2,856 hosted calls and 16,279,842 provider-reported tokens. The provider-token reduction of 67.44% differs slightly from the hosted-call reduction of 67.14% because the number of provider-reported text and image tokens varies across requests.

B.1.5 Session-Wise Policy Comparison

Table 1 in the main paper reports the session-wise routing counts. Table B.5 gives the corresponding hosted-token and generation-work comparisons. "Q2C generation hours" include the Local Student on every request and the Hosted Teacher on fallback requests. From S2 onward, every session reduces combined generation work relative to the hosted-only baseline. By S7, hosted-token use is reduced by 78.99% and combined generation work by 34.95%.

Table B.5: Session-wise hosted-token use and generation work. Routing counts appear in Table 1. Hosted-only tokens correspond to invoking GPT-Image-2 on every request. Q2C hours include 23.71 cumulative Local-Student hours and 19.24 cumulative Hosted-Teacher hours.

(a) Hosted-token use

Session	Q2C tokens	Hosted-only tokens	Reduction
S1	1,036,544	1,459,801	28.99%
S2	1,177,469	3,141,438	62.52%
S3	1,217,080	3,366,199	63.84%
S4	1,132,130	3,706,778	69.46%
S5	1,094,076	2,810,043	61.07%
S6	1,076,588	4,298,791	74.96%
S7	1,125,582	5,356,261	78.99%
Total	7,859,469	24,139,311	67.44%

(b) Serial-equivalent generation work

Session	Hosted only (h)	Q2C (h)	Reduction
S1	4.15	4.43	−6.75%
S2	7.97	5.81	27.14%
S3	7.66	6.07	20.72%
S4	9.14	6.14	32.85%
S5	7.18	5.45	24.01%
S6	10.38	6.52	37.25%
S7	13.11	8.53	34.95%
Total	59.59	42.94	27.93%

B.1.6 Accounting Scope and Limitations

Provider turnaround is measured from the stored request timestamp to the response-created timestamp; Local-Student time is measured from Student dispatch to completion. The median and P95 Teacher turnaround are 29 and 31 seconds, but the maximum is 1,725 seconds, and 147 of the 4,254 precomputed Teacher calls exceed 300 seconds. These rare service-side events explain the 50.43-second Teacher mean. All hours are sums of per-request durations rather than parallel wall-clock makespan. The comparison excludes Perception, Judge, adaptation, one-time Student preload, and queueing outside the recorded intervals. It therefore measures restoration-generation workload and provider usage, not complete user-facing latency or monetary savings.

B.2 Continual Updates and Guarded Promotion

Table B.6 reports every Full-RDA candidate on the unchanged SDCR-200 and the corresponding GSP decision. The candidate trajectory is not monotonic: S6 drops from 0.7534 to 0.7357 RPS and from 83.0% to 76.0% AC@0.7. GSP does not use SDCR-200; on its shadow-plus-history promotion set, the same candidate obtains a negative mean paired improvement and is rejected. Stage 5 therefore remains deployed until the Stage-7 retry is promoted.

The rejected S6 candidate yields a mean paired RPS change of $+0.0224$ on the current shadow subset but -0.0630 on the long-history subset; 34 of the 40 historical comparisons regress. This

observed update shows why current-window improvement alone is insufficient for deployment. The subsequent 3,000-step retry obtains a positive combined delta of $+0.0278$ and is promoted. Overall, Stage 7 improves over the Initial Student by 0.1957 RPS and 57.5 AC points, while the S6 event demonstrates that GSP can block an aggregate regression during sequential adaptation.

Table B.6: Full-RDA candidates on SDCR-200 and GSP outcomes. Candidate RPS and AC@0.7 are report-only; GSP uses the disjoint shadow-plus-history promotion set. The rejected S6 candidate is highlighted.

Boundary	Candidate RPS	Candidate AC@0.7	GSP mean Δ	Decision	Deployed afterward
Initial	0.5575	24.0%	–	–	Initial
S1	0.6770	59.5%	$+0.1773$	**Promote**	Stage 1
S2	0.7003	64.0%	$+0.0073$	**Promote**	Stage 2
S3	0.7111	65.5%	$+0.0402$	**Promote**	Stage 3
S4	0.7330	74.0%	$+0.0030$	**Promote**	Stage 4
S5	0.7534	83.0%	$+0.0402$	**Promote**	Stage 5
S6	0.7357	76.0%	-0.0203	**Reject**	Stage 5
S7	0.7532	81.5%	$+0.0278$	**Promote**	Stage 7

B.3 Additional Qualitative Comparisons

The following 16 SDCR-200 comparisons span the five compound domains.

Blur and low-light detail. Figure B.1 presents representative combinations of blur, low light, and noise.



Figure B.1: Selected examples from the blur and low-light-detail domain. Rows show low-light/noise/blur, low-light/blur, and blur cases.

Low light and noise. Figure B.2 presents representative cases dominated by low illumination and sensor noise.

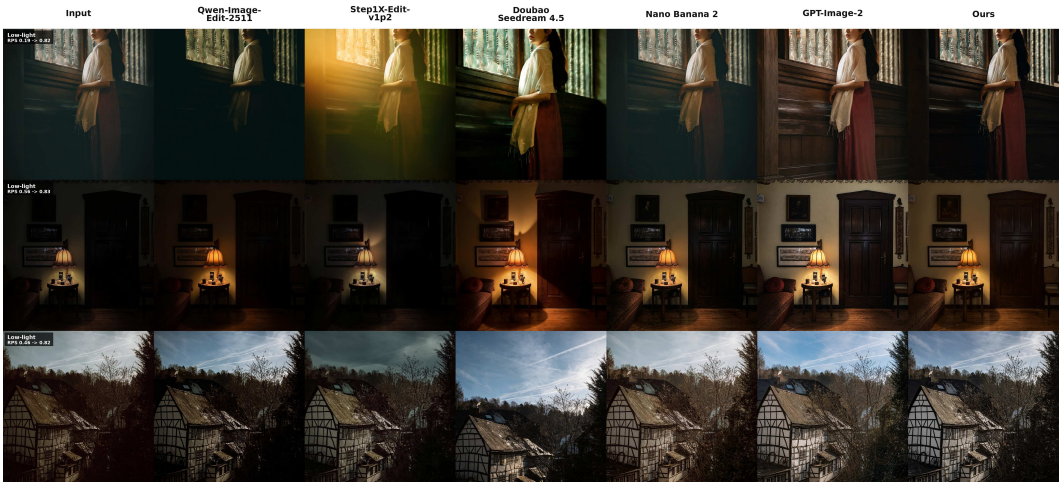

Figure B.2: Selected examples from the low-light and noise domain.

Haze and visibility. Figure B.3 presents representative haze and reduced-visibility cases.



Figure B.3: Selected examples from the haze and visibility domain.

Rain and adverse weather. Figure B.4 presents representative weather-related compound degradations.

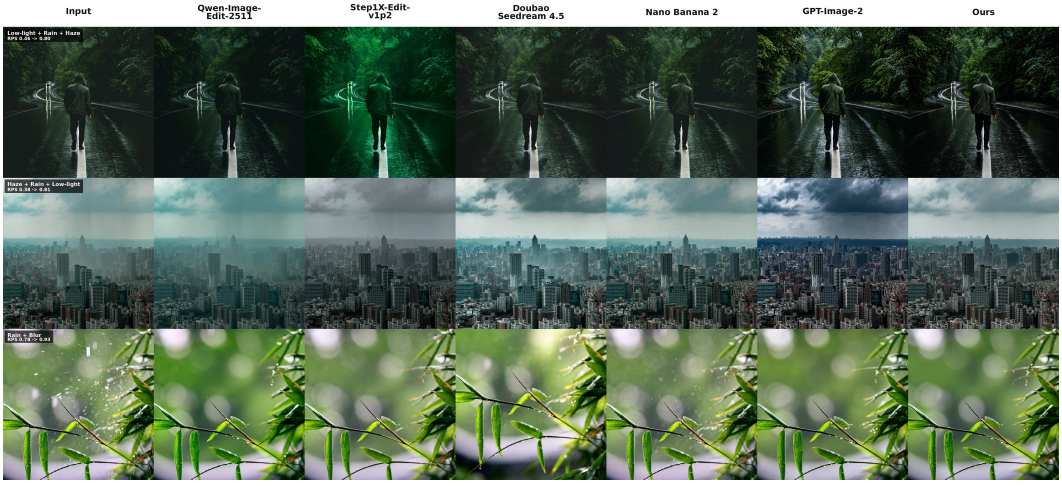

Figure B.4: Selected examples from the rain and adverse-weather domain.

Old photos, snow, and rare compounds. Figure B.5 presents the final rare-compound domain.

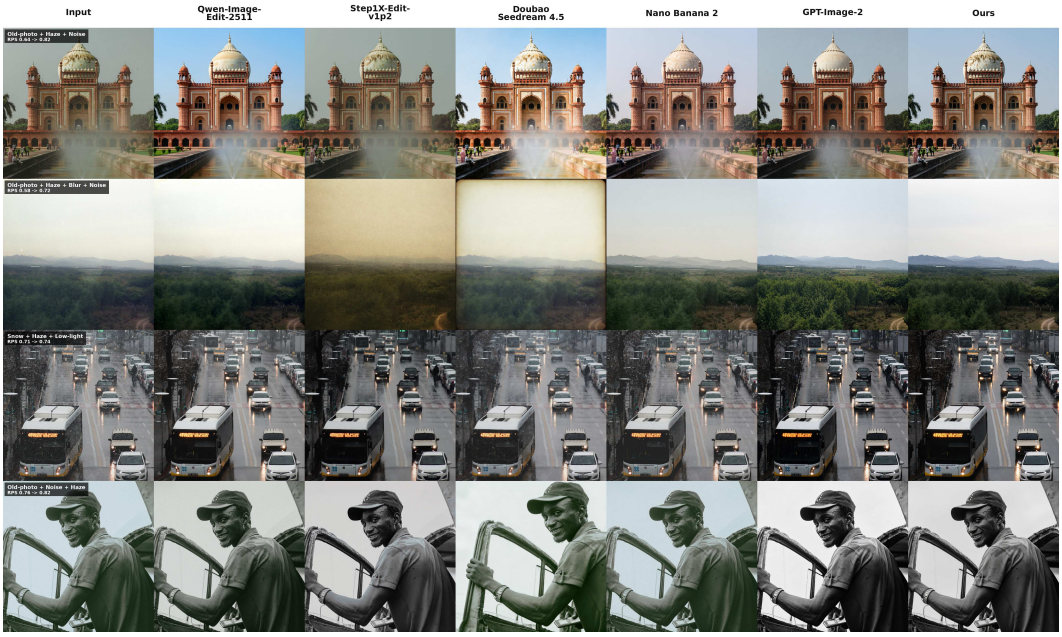

Figure B.5: Selected examples from the old-photo, snow, and rare-compound domain.

C End-to-End Q2C Example

We visualize one complete Q2C request using five saved process figures. The Initial Student obtains an RPS of 0.2053 and triggers Teacher fallback. The corresponding Teacher pseudo-reference obtains 0.8146, while the Stage-7 Student obtains 0.7675 on the same held-out request.

C.1 Perception Request and Structured SRB

Figure C.1 shows the serialized Perception request, including the degraded image, degradation vocabulary, preservation requirements, and structured response schema.

Perception-SRB Generation

developer_prompt

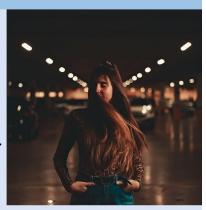

```
You are the Perception module of a real-world image restoration and pseudo-data generation agent.
Inspect the image before any image editing.
First, identify visible degradation types and describe exactly how each degradation appears.
Second, identify original image features that must be preserved by downstream GPT-Image2 calls, including
objects, faces, text, geometry, layout, colors, lighting intent, scene semantics, and fine structures.
Third, write a downstream restoration prompt with explicit restoration guidance and strong preservation constraints.
Fourth, classify the image scene into exactly one taxonomy leaf from the provided scene taxonomy.
Use only these degradation labels exactly:
          LOW_LIGHT, BLUR, NOISE, HAZE, RAIN, SNOW, NORMAL, OLD_PHOTO.
Do not invent labels. Use LOW_LIGHT, not LOW-LIGHT.
Severity and confidence must be floats in [0,1]. Evidence must name concrete visible cues or regions.
Return only valid JSON.
```

user_prompt

```
Inspect this image and return JSON exactly compatible with this schema: "
{
"degradations": [{"label": "LOW_LIGHT|BLUR|NOISE|HAZE|RAIN|SNOW|NORMAL|OLD_PHOTO", "
"severity": 0-1, "confidence": 0-1, "evidence": "visible evidence"}], "
"image_description": "concise content description", "
"preserved_features": ["features that must remain unchanged"], "
"quality_observations": ["visible quality observations"], "
"restoration_prompt": "direct restoration prompt for GPT-Image2 if the image is degraded", "
"restoration_intent": "short restoration goal", "
"risk_flags": ["hallucinated_detail|overexposure|identity_drift|text_damage|geometry_drift|color_shift"], "
"scene_classification": {"group": "taxonomy top-level key", "subcategory": "taxonomy leaf key", "
"matched_keyword": "best matching phrase", "reason": "scene evidence"}“
}.
```

Figure C.1: Serialized Perception request for the end-to-end Q2C example.

The request produces the structured SRB in Figure C.2. The SRB records the detected degradations, scene content, preserved features, restoration intent, risk flags, and scene classification.

Perception-SRB Generation

SRB Response


```
{
       "degradations": [
         {
           "label": "LOW_LIGHT",
           "severity": 0.58,
           "confidence": 0.91,
           "evidence": "Large regions of the ceiling, background, clothing, and left side of the subject are very dark with crushed shadow detail; illumination comes mainly from warm
overhead lamps and reflections on the wet-looking floor."
         },
         {
           "label": "BLUR",
           "severity": 0.32,
           "confidence": 0.72,
           "evidence": "Background vehicles, ceiling lights, and distant figures are heavily out of focus with bokeh; the subject's hair and facial edges are slightly soft, consistent with
shallow depth of field or mild focus softness."
         }
       ],
       "image_description": "Nighttime or dim indoor parking-area portrait of a young woman standing centered with long hair, hands in pockets, warm overhead lights, cars, reflective floor,
and strong background bokeh.",
       "preserved_features": [
         "Centered female subject with long brown hair partly crossing her face",
         "Face identity, pose, closed or downward-looking eyes, and hands in pockets",
         "Leopard-print long-sleeve top, teal high-waisted pants, dark belt, and small hair accessory",
         "Warm amber cinematic lighting and dark moody atmosphere",
         "Symmetrical rows of bright overhead ceiling lights receding into the background",
         "Out-of-focus cars, distant figures, reflective floor, and parking-garage/street-like layout",
         "Shallow-depth-of-field portrait composition with background bokeh",
         "No added text, logos, jewelry, vehicles, people, or facial changes"
       ],
       "quality_observations": [
         "Intentional low-key exposure creates very dark shadows and reduced visibility in clothing and background.",
         "Warm color grading dominates the scene with orange highlights and teal pants contrast.",
         "Background blur appears mostly intentional and should not be globally sharpened.",
         "Bright ceiling lights are clipped or near-clipped and should not be over-expanded.",
         "Fine hair strands across the face are important but partially lost in darkness."
       ],
       "restoration_prompt": "Restore this low-light portrait while preserving the original cinematic night mood. Gently lift shadow detail on the woman's face, hair, clothing, and hands
without making the scene look daytime or flat. Reduce excessive darkness and mild sensor-like softness only on the main subject; keep the background cars, distant figures, ceiling lights,
and bokeh intentionally out of focus. Preserve the woman's identity, facial expression, pose, hair strands across the face, leopard-print top, teal pants, belt, warm amber lighting,
reflective floor, and symmetrical parking-area geometry. Avoid adding new facial features, changing clothing patterns, altering the layout, removing bokeh, overexposing lamps, shifting
colors, or inventing background details.",
       "restoration_intent": "Improve visibility and subject clarity while keeping the warm low-light cinematic portrait intact.",
       "risk_flags": [
         "identity_drift",
         "overexposure",
         "geometry_drift",
         "color_shift",
         "hallucinated_detail"
       ],
       "scene_classification": {
         "group": "02_semantic_content",
         "subcategory": "portrait_people",
         "matched_keyword": "street portrait",
         "reason": "The image is primarily a posed portrait of one person in a real-world urban or parking-area environment, with the human subject dominating the frame."
       }
}
```


Figure C.2: Structured SRB response produced from the request in Figure C.1.

C.2 Restoration Prompt and Model Outputs

The SRB is rendered as the restoration instruction shown in Figure C.3. The same restoration goal and preservation constraints are used for the Initial Student, hosted Teacher, and Stage-7 Student.

Restoration-Image Edit

prompt

```
Restore this low-light portrait while preserving the original cinematic night mood. Gently lift shadow
detail on the woman's face, hair, clothing, and hands without making the scene look daytime or flat. Reduce
excessive darkness and mild sensor-like softness only on the main subject; keep the background cars, distant
figures, ceiling lights, and bokeh intentionally out of focus. Preserve the woman's identity, facial
expression, pose, hair strands across the face, leopard-print top, teal pants, belt, warm amber lighting,
reflective floor, and symmetrical parking-area geometry. Avoid adding new facial features, changing clothing
patterns, altering the layout, removing bokeh, overexposing lamps, shifting colors, or inventing background
details.
```

Output



(a) Initial Student



(b) Teacher Reference



(c) Stage-7 Student

Figure C.3: SRB-derived restoration instruction and the Initial Student, Teacher pseudo-reference, and Stage-7 Student outputs.

C.3 Judge Request and Structured Response

Figure C.4 shows the Judge request used to assess the Initial Student output against the degraded input under the same restoration and preservation context.

Judge-Image Evaluation

developer_prompt

```
You are a professional image restoration evaluator.
TASK: Evaluate whether the generated pseudo-GT correctly improves the low-quality input while preserving content.
Use only these project degradation labels exactly: LOW_LIGHT, BLUR, NOISE, HAZE, RAIN, SNOW, NORMAL, OLD_PHOTO.
For residual_degradations, use only non-normal labels from: LOW_LIGHT, BLUR, NOISE, HAZE, RAIN, SNOW, OLD_PHOTO.
Do not invent residual degradation categories.
IMAGES:
1. FIRST IMAGE = LQ degraded input
2. SECOND IMAGE = generated pseudo-GT restored output
Evaluate five dimensions in full detail:
1.  restoration quality (0-100)
2.  object consistency (0-100)
3.  geometry alignment (0-100)
4.  content reasonableness (0-100)
5.  color consistency (0-100)
6.  Return only valid JSON.
```



(a) LQ



(a) HQ

user_prompt

First image is LQ, second image is Generated HQ.
Evaluate in full detail.
Original LQ perception source: gpt55.
Here are the perception informations as references:
Image description: Nighttime or dim indoor parking-area portrait of a young woman standing centered with long hair, hands in pockets, warm overhead lights, cars, reflective floor, and strong background bokeh.
Detected degradations JSON:
[{"label": "LOW_LIGHT", "severity": 0.58, "confidence": 0.91, "evidence": "Large regions of the ceiling, background, clothing, and left side of the subject are very dark with crushed shadow detail; illumination comes mainly from warm overhead lamps and reflections on the wet-looking floor."},
{"label": "BLUR", "severity": 0.32, "confidence": 0.72, "evidence": "Background vehicles, ceiling lights, and distant figures are heavily out of focus with bokeh; the subject's hair and facial edges are slightly soft, consistent with shallow depth of field or mild focus softness."}].
Restoration prompt used or recommended:
Restore this low-light portrait while preserving the original cinematic night mood. Gently lift shadow detail on the woman's face, hair, clothing, and hands without making the scene look daytime or flat. Reduce excessive darkness and mild sensor-like softness only on the main subject; keep the background cars, distant figures, ceiling lights, and bokeh intentionally out of focus. Preserve the woman's identity, facial expression, pose, hair strands across the face, leopard-print top, teal pants, belt, warm amber lighting, reflective floor, and symmetrical parking-area geometry. Avoid adding new facial features, changing clothing patterns, altering the layout, removing bokeh, overexposing lamps, shifting colors, or inventing background details.
Perception quality observations:
["Intentional low-key exposure creates very dark shadows and reduced visibility in clothing and background.", "Warm color grading dominates the scene with orange highlights and teal pants contrast.", "Background blur appears mostly intentional and should not be globally sharpened.", "Bright ceiling lights are clipped or near-clipped and should not be over-expanded.", "Fine hair strands across the face are important but partially lost in darkness."]. "

Return exactly this JSON schema:
{"scores": {"restoration quality": 0-100, "object consistency": 0-100, "geometry alignment": 0-100, "content reasonableness": 0-100, "color consistency": 0-100},
"reasons": {"restoration quality": "detailed", "object consistency": "detailed", "geometry alignment": "detailed region-by-region", "content reasonableness": "detailed", "color consistency": "detailed"},
"residual_degradations": [{"label": "LOW_LIGHT|BLUR|NOISE|HAZE|RAIN|SNOW|OLD_PHOTO", "severity": 0-1, "confidence": 0-1, "evidence": "visible evidence"}],
"quality_observations": ["remaining quality observations"],
"content_fidelity_observations": ["content fidelity observations"],
"failure_modes": ["residual_blur etc."].
Residual degradations must describe visible remaining degradation in the generated HQ image using only the project labels. If the issue is hallucination, identity change, geometry error, or another non-project issue, describe it in reasons/failure_modes and reflect it in the five scores instead of residual_degradations.

Figure C.4: Serialized Judge request for evaluating the Initial Student restoration.

The structured response in Figure C.5 reports the five quality dimensions, residual degradations, content-fidelity observations, and failure modes. Its low restoration score provides the routing evidence that triggers Teacher fallback.

Judge-Image Evaluation

Judge Response


```
{
      "scores": {
        "restoration quality": 28,
        "object consistency": 72,
        "geometry alignment": 78,
        "content reasonableness": 55,
        "color consistency": 62
      },
      "reasons": {
        "restoration quality": "The generated pseudo-GT does not correctly restore the low-light input. Instead of gently lifting shadow detail, it makes the scene substantially darker and
more contrast-heavy. Large areas of the ceiling, background cars, subject torso, left arm, hair, and face are crushed into near-black, losing information that was at least partially visible
in the LQ image. The face is not meaningfully recovered; only a narrow warm highlight remains while most facial structure is obscured. The leopard-print top, hands, belt, and teal pants
retain some visible detail, but overall restoration quality is poor because the main degradation, LOW_LIGHT, is intensified rather than improved. Background bokeh is preserved, but the
output looks more like a stylized dark edit than a restored pseudo-GT.",
        "object consistency": "The central woman, long hair across the face, hands-in-pockets pose, leopard-print top, teal pants, belt, warm lights, cars, and parking-garage setting are
generally preserved. However, the subject's silhouette becomes less readable due to extreme shadow crushing, especially around the left side of the body, torso, and hair. Facial identity is
harder to verify because most of the face is hidden in darkness, and some fine hair strands appear simplified or merged into broader dark masses. No major new objects are invented, but some
background structures and cars are suppressed enough that they are effectively removed from visibility.",
        "geometry alignment": "The central composition remains broadly aligned: the woman stays centered, the overhead lights still form symmetrical perspective lines, and the floor
reflection layout remains consistent. The head, hair mass, torso, arms, hands, belt, and pants occupy approximately the same spatial locations. Region-by-region, the upper ceiling and light
rows are aligned well, though the ceiling detail is almost entirely blackened; the face and hair remain in the same position but lose internal detail; the torso and arms align but the left
side blends into the background; the pants and hands are aligned and relatively readable; the background cars and distant figures are still roughly in place but are much less visible. There
is no severe warping, but detail loss makes some boundaries ambiguous.",
        "content reasonableness": "The output remains plausible as a cinematic low-key night portrait, but it is not reasonable as a restoration target because it exaggerates darkness beyond
the source. The intended mood is preserved, yet the restoration prompt specifically asked to lift shadow detail while maintaining the nighttime look; the generated result does the opposite.
The background blur and bokeh remain visually coherent, and there are no obvious impossible objects, but the loss of subject visibility and background information makes the pseudo-GT
unsuitable as a high-quality reference.",
        "color consistency": "The warm amber lighting and teal pants are broadly consistent with the LQ image, and the cinematic orange-teal palette is retained. However, the generated image
increases contrast and saturation selectively, making hair highlights more copper-orange, skin highlights more intense, and surrounding shadows much blacker. The overall color mood is
similar, but tonal balance is inconsistent with a restoration goal because midtones are suppressed and the scene becomes overly dark and stylized. Lamps remain warm and bright without major
hue shifts, but the black crush reduces natural color variation in clothing, hair, and background."
      },
      "residual_degradations": [
        {
          "label": "LOW_LIGHT",
          "severity": 0.82,
          "confidence": 0.98,
          "evidence": "The generated image has extensive crushed shadows across the subject's face, hair, torso, left arm, ceiling, cars, and floor, with less recoverable detail than the LQ
input."
        },
        {
          "label": "BLUR",
          "severity": 0.28,
          "confidence": 0.74,
          "evidence": "The background remains heavily out of focus, which is partly intentional, but some subject hair and facial boundaries are still soft or merged by darkness rather than
cleanly restored."
        }
      ],
      "quality_observations": [
        "The output increases dramatic contrast but fails to recover shadow detail.",
        "The subject's face is less readable than in the LQ input due to stronger black crush.",
        "Hair highlights are visually striking but many fine strands and dark regions are lost.",
        "Background bokeh is preserved, but cars and garage structure become overly obscured.",
        "Teal pants and hands remain among the clearest restored regions."
      ],
      "content_fidelity_observations": [
        "Main pose, clothing type, long hair, hands in pockets, and centered portrait composition are preserved.",
        "No major hallucinated objects are introduced, but background content is suppressed by darkness.",
        "Facial identity fidelity is weakened because most facial details are hidden.",
        "Leopard-print pattern is partially preserved but less visible in shadowed areas.",
        "Parking-garage geometry and light-row symmetry remain broadly consistent."
      ],
      "failure_modes": [
        "under-restoration_of_low_light",
        "increased_shadow_crushing",
        "loss_of_facial_detail",
        "loss_of_background_detail",
        "over_stylized_contrast",
        "residual_blur"
      ],
}
```


Figure C.5: Structured Judge response and routing evidence for the Initial Student output.