ACTION-REGULATED STATE SPACES FOR ALL-IN-ONE ENDOSCOPIC IMAGE RESTORATION

Anonymous authors
Paper under double-blind review

ABSTRACT

All-in-one endoscopic image restoration faces a fundamental spatial action-allocation problem: degradations vary widely in location, extent, and severity. We identify a previously overlooked failure mode of state-space restorers under such spatial heterogeneity. Because corrective evidence is recurrently written into and propagated through shared state dynamics, an action appropriate for one region may cross its intended boundary and inadvertently alter intact tissue. We term this phenomenon restoration-action spillover. To address this issue, we propose BARS-MambaIR, a blind action-regulated state-space model that predicts dense demand, scope, and magnitude fields to encode where, how far, and how strongly restoration should act. These fields regulate state writing, propagation decay, boundary-conditioned attenuation, and local-to-global action routing, while also gating the output residual. An action-allocation consistency objective learns these fields directly from paired supervision, without requiring degradation labels or region annotations, to suppress intact-region distortion and restoration spillover. Controlled footprint–severity sweeps, localized state interventions, and field-swapping experiments further reveal how spatially regulated state propagation limits restoration-action spillover.

1 INTRODUCTION

Endoscopic vision is frequently degraded by smoke generated by surgical instruments, condensation on the lens, and contamination by blood or tissue fluid. These artifacts obscure anatomy and can impair both visual assessment and downstream perception. Recent benchmarks have made it possible to study such degradations with paired surgical data (Pei et al., 2026), and all-in-one models aim to restore several degradation types with one set of parameters (Chen et al., 2025). Yet endoscopic restoration poses a property that is easy to miss when methods are organized only by degradation category: the geometry of the required correction is highly heterogeneous. Diffuse fog may cover most of a frame, smoke may have multiple semi-transparent components, and a splash may occupy a small region with a sharp and irregular boundary. Moreover, vessels, specular highlights, and low-contrast tissue can resemble these artifacts locally.

Most all-in-one restoration methods adapt representations according to an estimated degradation condition. Contrastive encoders (Li et al., 2022), learned prompts (Potlapalli et al., 2023), frequency modulation (Cui et al., 2025), and task-discriminative experts (Zamfir et al., 2025) are effective mechanisms for deciding *what kind* of degradation is present. More recent methods provide query- or pixel-level guidance (Zhou et al., 2026; Liu et al., 2026), which improves spatial adaptation. Nevertheless, conditioning and spatial localization do not by themselves specify how a correction should travel through a deep restoration network. In particular, a state-space model transports a recurrent state along its scan order (Gu & Dao, 2024). Evidence for a strong, broad corruption can thus alter the corrective response applied to a nearby small corruption or to an intact region. Appearance similarity between artifacts and anatomy makes this failure more likely, but it is part of the same underlying problem: an incorrectly transported corrective operation crosses the boundary of its valid support.

We call the image modification induced by a restorer its *restoration action*. Operationally, it is the predicted residual $a_\theta(x)$, rather than an action in the reinforcement-learning sense. We use *action propagation* to describe how evidence at one spatial location changes this residual elsewhere

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

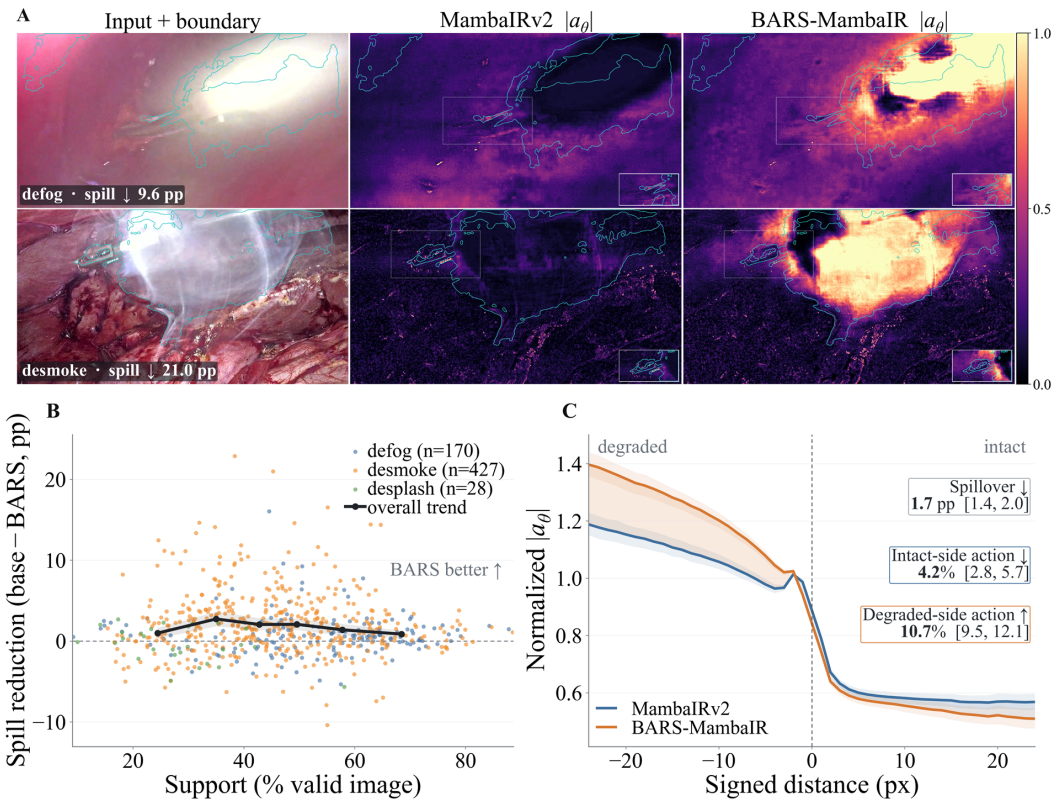

Figure 1: Diagnosing restoration-action spillover under heterogeneous endoscopic degradations. (a) Representative inputs and output-action maps $|a_\theta(x)| = |f_\theta(x) - x|$ produced by MambaIRv2 and BARS-MambaIR. Cyan contours denote evaluation-only degradation-support boundaries obtained from aligned input–reference residuals. (b) Per-image change in the spill ratio, measured in percentage points (pp), as a function of degradation support. Positive values indicate lower spillover for BARS-MambaIR. (c) Mean normalized action magnitude as a function of signed distance to the support boundary; negative and positive distances correspond to degraded and intact regions, respectively. Shaded regions denote 95% bootstrap confidence intervals.

through the network. This viewpoint exposes a failure mode that conventional PSNR or SSIM can hide: *action spillover*, where correction response extends from a degraded region into tissue that should be preserved. We operationalize this failure mode in Figure 1. Figure 1(a) visualizes the output action $|a_\theta(x)|$ for representative endoscopic degradations with markedly different spatial supports. Figure 1(b) relates the reduction in spillover to the fraction of the valid image occupied by the degradation, revealing substantial sample-wise variation rather than a uniform advantage. Figure 1(c) further measures the action magnitude as a function of signed distance to the estimated degradation boundary. Compared with the MambaIRv2 restorer, BARS-MambaIR attenuates the residual response on the intact side while retaining a stronger response inside the degraded support. These observations motivate our hypothesis that restoration quality depends not only on recognizing the degradation, but also on regulating how corrective evidence is transported across its spatial boundary.

This leads to our central question: *How can a unified state-space restorer adapt the propagation of corrective actions to endoscopic degradations with heterogeneous geometries while preserving intact tissue?* We hypothesize that a useful controller must resolve three coupled quantities: *demand*, or where a correction is warranted; *scope*, or the spatial range over which its evidence should persist; and *magnitude*, or how strongly the image should be changed. Scope without demand can propagate a wrong operation, demand without scope cannot distinguish diffuse from local corruption, and either without magnitude can still over-correct tissue.

Based on this hypothesis, we propose BARS-MambaIR, a **B**lind **A**ction-**R**egulated **S**tate-space framework. A compact estimator predicts dense demand, scope, and magnitude fields without knowing the degradation label. Together these fields parameterize a spatially varying control law over a selective state-space kernel. Demand controls whether a state is written and read; scope selects among global, intermediate, and local propagation bases; and a boundary term contracts the transition when a local action approaches the edge of its estimated support. The same action coordinates regulate the final residual, so internal transport and image-space modification are consistent. We train the controller with paired supervision. The aligned input–target residual defines a soft demand target, while its low- and high-frequency components provide a label-free proxy for the required propagation scope. No degradation labels or manually annotated artifact masks are required.

Our contributions are threefold: 1) We formulate heterogeneous all-in-one endoscopic restoration through the operational lens of spatial restoration actions, and introduce cross-boundary action spillover as a measurable failure mode complementary to global fidelity metrics. 2) We develop BARS-MambaIR, a unified control mechanism that converts blind demand, scope, and magnitude coordinates into spatial regulation of state writing, effective transition decay, boundary attenuation, multi-range routing, and image-space residual readout. 3) Through support–severity stratification, boundary-distance analysis, capacity-controlled comparisons, and localized interventions, we characterize when action regulation is effective and how its internal controls relate to reduced intact-region modification.

2 Related Work

All-in-one and spatially adaptive restoration. All-in-one restoration replaces degradation-specific networks with a shared model for unknown corruptions. Representative approaches infer a degradation representation, prompt, frequency modulation, or expert route, including AirNet, PromptIR, AdaIR, and MoCE-IR (Li et al., 2022; Potlapalli et al., 2023; Cui et al., 2025; Zamfir et al., 2025). In medical and endoscopic imaging, AMIR mitigates cross-task interference through spatial–channel routing (Yang et al., 2024), while SurgClean and EndoIR extend unified restoration to smoke, fog, and splash degradations (Pei et al., 2026; Chen et al., 2025). Recent methods further introduce query- or pixel-level guidance to localize nonuniform corruption (Zhou et al., 2026; Liu et al., 2026; Wang et al., 2026). These mechanisms mainly determine *what* degradation is present or *where* features should be adapted. Our focus is complementary: we regulate how the resulting corrective operation is transported through a shared restoration backbone.

State-space restoration and action transport. Selective state-space models provide input-dependent long-range propagation with linear sequence complexity (Gu & Dao, 2024). MambaIR adapts this formulation to image restoration through local enhancement and multidirectional scanning, and MambaIRv2 adds semantic routing and non-causal interactions (Guo et al., 2024; 2025). Subsequent work injects degradation priors, suppresses corrupted feature propagation, or learns multiple state-evolution timescales (Liu et al., 2025; Wang et al., 2026; Gao et al., 2026). BARS-MambaIR instead studies the image-space residual as a *restoration action*. It explicitly measures cross-boundary spillover and predicts demand, scope, and magnitude coordinates that jointly regulate state writing, retention, boundary attenuation, multi-range routing, and residual readout. This support–range–magnitude view connects spatial localization to the internal dynamics that determine how far a correction travels.

3 Method

3.1 Problem Formulation: Restoration as Controlled Action Transport

Let $\boldsymbol{x} \in \mathbb{R}^{H\times W\times 3}$ be a degraded endoscopic image and $\boldsymbol{y}$ its clean reference. A restorer f_θ predicts $\widehat{\boldsymbol{y}} = f_\theta(\boldsymbol{x})$. We define its image-space restoration action and the target action as

$$\boldsymbol{a}_\theta(\boldsymbol{x}) = f_\theta(\boldsymbol{x}) - \boldsymbol{x}, \qquad \boldsymbol{a}^\star = \boldsymbol{y} - \boldsymbol{x}. \tag{1}$$

This definition is deliberately operational: it describes what the model changes, independently of how its internal features are parameterized. To evaluate whether the learned restoration action extends unnecessarily beyond degradation-affected regions, we additionally use an evaluation-only

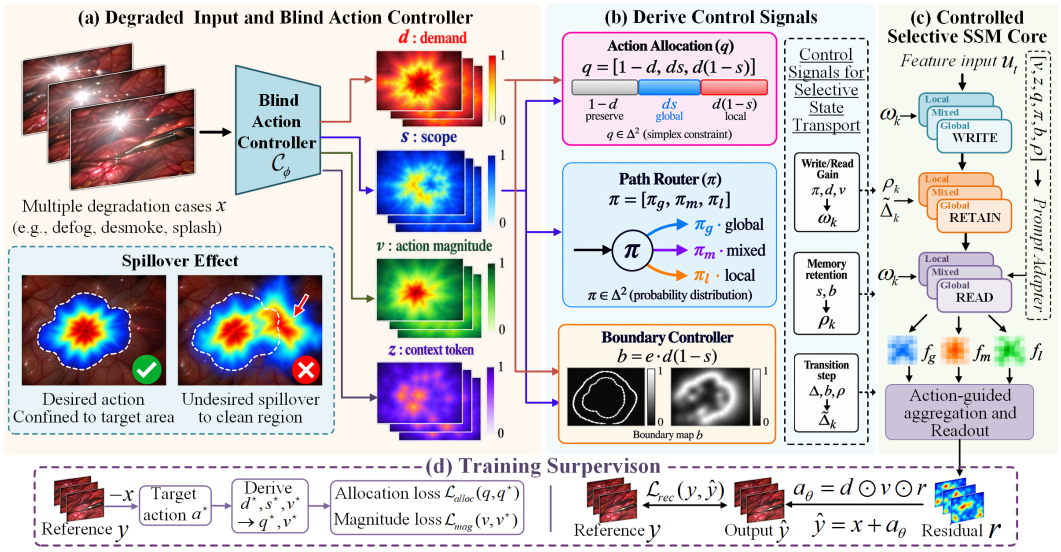

Figure 2: BARS-MambaIR uses spatial action fields to regulate state transport and residual readout.

soft degradation-support map $\boldsymbol{m} \in [0,1]^{H\times W}$, we measure action spillover by

$$\mathrm{Spill}(\boldsymbol{a}_\theta; \boldsymbol{m}) = \frac{\|(1-\boldsymbol{m}) \odot \boldsymbol{a}_\theta\|_1}{\|\boldsymbol{a}_\theta\|_1 + \varepsilon}. \tag{2}$$

Eq. (2) is an evaluation functional, not a claim that every change outside a binary mask is erroneous. Accordingly, our objective is not merely to infer a degradation identity, but to learn a spatial controller $\mathcal{C}_\phi$ that determines whether, how far, and how strongly the correction at each position may influence other positions. This controller acts on the transition kernel of a selective state-space model and on the final image residual. Figure 2 summarizes the complete computation graph: a blind controller first estimates spatial action coordinates, which are converted into write, retention, boundary, and routing controls for the selective state-space backbone, and are reused by the final residual readout.

3.2 Blind Restoration-Action Coordinates

As illustrated in Figure 2(a), a lightweight blind action controller estimates

$$(\boldsymbol{d}, \boldsymbol{s}, \boldsymbol{v}, \boldsymbol{z}) = \mathcal{C}_\phi(\boldsymbol{x}), \qquad \boldsymbol{d}, \boldsymbol{s}, \boldsymbol{v} \in [0,1]^{H\times W}, \tag{3}$$

where $\boldsymbol{d}$ is *demand*, $\boldsymbol{s}$ is *scope*, $\boldsymbol{v}$ is action *magnitude*, and $\boldsymbol{z}$ is an image-level context token. Demand should be high where modification is necessary. Scope governs the spatial extent of corrective propagation, favoring broader context for diffuse corruption and localized processing for concentrated corruption. Magnitude determines the strength of the final correction. These coordinates form a controller, not three independently interpreted segmentation maps.

For every position p, as illustrated in Figure 2(b), demand and scope induce a conservative action allocation

$$\boldsymbol{q}_p = \left[q_p^{\mathrm{pre}}, q_p^{\mathrm{g}}, q_p^{\mathrm{l}}\right] = \left[1-d_p,\ d_p s_p,\ d_p(1-s_p)\right], \qquad \mathbf{1}^\top \boldsymbol{q}_p = 1. \tag{4}$$

Here, d_p determines the total corrective mass, while s_p partitions it between global and local correction, yielding a three-way soft allocation over preservation, global correction, and local correction. The maximal component of $\boldsymbol{q}_p$ identifies the dominant action mode, while the simplex retains continuous mixtures between modes. To represent intermediate propagation, the network predicts logits $\boldsymbol{\ell}_p$ and forms the three-way path distribution

$$\boldsymbol{\pi}_p = [\pi_{p,\mathrm{g}}, \pi_{p,\mathrm{m}}, \pi_{p,\mathrm{l}}] = \mathrm{softmax}(\boldsymbol{\ell}_p + \log\left([s_p \quad 4s_p(1-s_p) \quad 1-s_p] + \varepsilon\right)), \tag{5}$$

whose entries correspond to global, mixed, and local propagation. The scope-derived terms respectively favor large, intermediate, and short propagation ranges: s_p increases toward the global

endpoint, $1 - s_p$ toward the local endpoint, while $4s_p(1 - s_p)$ peaks at $s_p = \frac{1}{2}$. Notably, $\boldsymbol{q}_p$ and $\boldsymbol{\pi}_p$ play distinct roles. $\boldsymbol{q}_p$ allocates restoration action among preservation, global correction, and local correction, whereas $\boldsymbol{\pi}_p$ governs how the resulting corrective state is routed across global, mixed, and local propagation bases. In other words, $\boldsymbol{q}_p$ specifies what type of action is allocated, while $\boldsymbol{\pi}_p$ determines how that action is propagated.

Scope controls the desired propagation extent, but not where propagation should be suppressed. For localized corruption such as a splash, corrective states should not cross into surrounding intact regions. We therefore derive a boundary barrier directly from the demand geometry

$$e_p = \mathrm{MaxPool}_{3\times3}\left(\frac{\|\nabla d_p\|_2}{\|\nabla d_p\|_2 + \overline{\|\nabla \boldsymbol{d}\|_2} + \varepsilon}\right), \qquad b_p = e_p\, d_p(1 - s_p). \tag{6}$$

where $\overline{(\cdot)}$ denotes spatial averaging. The multiplicative form makes the barrier selective: e_p localizes sharp demand transitions, d_p suppresses irrelevant boundaries where little correction is required, and $1 - s_p$ activates the barrier primarily for localized actions. Consequently, b_p becomes large only near the boundary of a locally corrected region, rather than indiscriminately suppressing long-range context.

3.3 Action-Regulated Selective State Transport

As illustrated in Figure 2(c), consider tokens $\boldsymbol{u}_t$ ordered by the semantic scan of MambaIRv2 (Guo et al., 2025), where t indexes the token position along the resulting scan sequence. A standard discretized selective state-space model can be written as

$$\boldsymbol{h}_t = \exp(\Delta_t \boldsymbol{A})\boldsymbol{h}_{t-1} + \boldsymbol{\Gamma}(\Delta_t, \boldsymbol{A}, \boldsymbol{B}_t)\boldsymbol{u}_t, \qquad \boldsymbol{o}_t = \boldsymbol{C}_t\boldsymbol{h}_t + \boldsymbol{D}\boldsymbol{u}_t, \tag{7}$$

where Δ_t, $\boldsymbol{B}_t$, and $\boldsymbol{C}_t$ are input-selective parameters. A long-range state is helpful for diffuse fog but can transport an inappropriate correction across the boundary of a small splash. We therefore treat the selective transition as a controlled transport kernel. For propagation basis $k \in \{\mathrm{g}, \mathrm{m}, \mathrm{l}\}$, define the write/read gain

$$\omega_{t,k} = d_t\, \frac{1 + v_t}{2}\, \pi_{t,k}. \tag{8}$$

Here, d_t gates corrective transport, v_t modulates its strength, and $\pi_{t,k}$ distributes it across the three propagation bases. Since $\sum_k \pi_{t,k} = 1$, the total gain satisfies $\sum_k \omega_{t,k} = d_t(1 + v_t)/2$, separating the overall amount of corrective transport from its path-wise allocation. To directly control how long the accumulated corrective state persists, we further regulate state retention and memory decay through $\rho_{t,k}$ and the effective discretization step $\widetilde{\Delta}_{t,k}$

$$\rho_{t,k} = \mathrm{clip}((a_k + c_k s_t)(1 - \beta_k b_t), \rho_{\min}, 1)\,, \tag{9}$$

$$\widetilde{\Delta}_{t,k} = \Delta_{t,k}\big[1 + \beta_k b_t + (1 - \rho_{t,k})\big]. \tag{10}$$

The triplet (a_k, c_k, β_k) jointly characterizes the propagation profile of each basis, and is chosen to establish an ordered hierarchy from long-range global transport to strongly localized, boundary-sensitive transport. In our implementation, (a_k, c_k, β_k) is set to $(0.75, 0.25, 0.25)$, $(0.40, 0.60, 0.50)$, and $(0.15, 0.45, 1.00)$ for the global, mixed, and local bases, respectively, with $\rho_{\min} = 0.05$. Since $\boldsymbol{A}_k = -\exp(\boldsymbol{A}_{\log,k})$, increasing $\widetilde{\Delta}_{t,k}$ contracts $\exp(\widetilde{\Delta}_{t,k}\boldsymbol{A}_k)$, thereby shortening the effective memory of the corresponding state transition.

The action-gated token $\boldsymbol{u}_t$ first parameterizes the selective dynamics of each propagation basis,

$$(\Delta_{t,k}, \boldsymbol{B}_{t,k}, \boldsymbol{C}_{t,k}) = g_{\theta,k}(\omega_{t,k}\boldsymbol{u}_t), \tag{11}$$

where $g_{\theta,k}$ denotes the learnable selective projection for basis k, mapping the gated token to the input-dependent SSM parameters. The controlled state propagation and readout are then given by

$$\boldsymbol{h}_{t,k} = \exp(\widetilde{\Delta}_{t,k}\boldsymbol{A}_k)\boldsymbol{h}_{t-1,k} + \boldsymbol{\Gamma}(\widetilde{\Delta}_{t,k}, \boldsymbol{A}_k, \boldsymbol{B}_{t,k})(\omega_{t,k}\boldsymbol{u}_t), \tag{12}$$

$$\widetilde{\boldsymbol{f}}_{t,k} = \omega_{t,k}(\boldsymbol{C}_{t,k} + \omega_{t,k}\boldsymbol{p}_{t,k})\boldsymbol{h}_{t,k}, \qquad \boldsymbol{f}_t = \sum_k \widetilde{\boldsymbol{f}}_{t,k}, \tag{13}$$

Here, $\boldsymbol{p}_{t,k}$ is an action-conditioned prompt generated from the full action vector $[\boldsymbol{v}, \boldsymbol{z}, \boldsymbol{q}, \boldsymbol{\pi}, b, \rho]$ through a zero-initialized adapter, and modulates the readout of basis k. The same action vector also conditions semantic routing and path features through zero-initialized adapters, preserving the pretrained backbone at initialization while enabling restoration-specific adaptation. The three paths thus form a compact basis of propagation ranges rather than independent restoration branches.

Table 1: Quantitative comparison on SurgClean and SegSTRONG-C. Spill denotes the task-macro fraction of restoration action outside the evaluation support.

SurgClean				
Method	Defog	Desmoke	Desplash	Spill↓
	PSNR↑ / SSIM↑ / LPIPS↓			
Restormer	19.22 / 0.647 / 0.477	19.74 / 0.695 / 0.391	21.19 / 0.731 / 0.364	0.3997
PromptIR	19.10 / 0.631 / 0.473	19.60 / 0.676 / 0.393	21.20 / 0.711 / 0.361	0.4019
TransWeather	19.14 / 0.645 / 0.452	19.63 / 0.696 / **0.354**	21.23 / 0.725 / 0.353	0.3964
AdaIR	19.35 / **0.651** / 0.482	19.83 / 0.701 / 0.387	21.29 / 0.732 / 0.359	0.3907
AMIR	19.24 / 0.648 / 0.488	19.67 / 0.694 / 0.393	21.16 / 0.730 / 0.358	0.4007
MambaIRv2	19.49 / 0.648 / 0.499	19.73 / 0.697 / 0.400	21.27 / 0.731 / 0.358	0.3962
MoCE-IR	19.36 / 0.641 / 0.472	19.69 / 0.697 / 0.384	21.31 / 0.731 / 0.347	0.3945
MIRAGE	19.18 / 0.645 / 0.481	19.34 / 0.685 / 0.383	21.42 / 0.731 / 0.344	0.4045
BARS-MambaIR	**19.55** / 0.649 / **0.450**	**20.24 / 0.703** / 0.386	**21.54 / 0.738 / 0.343**	**0.3863**
SegSTRONG-C				
Method	Blood	LowBrightness	Smoke	Spill↓
	PSNR↑ / SSIM↑ / LPIPS↓			
Restormer	25.12 / 0.598 / 0.638	20.20 / 0.505 / 0.684	23.04 / 0.537 / 0.655	0.4049
PromptIR	25.41 / 0.589 / 0.613	17.25 / 0.467 / 0.690	22.33 / 0.561 / 0.646	0.4480
TransWeather	25.48 / 0.590 / 0.673	20.11 / 0.535 / 0.723	23.08 / 0.533 / 0.676	0.4151
AdaIR	25.65 / 0.608 / 0.635	17.57 / 0.467 / 0.711	23.24 / 0.571 / 0.652	0.3963
AMIR	25.58 / 0.606 / 0.642	16.97 / 0.432 / 0.709	21.99 / 0.547 / 0.656	0.4608
MambaIRv2	25.43 / 0.596 / 0.687	19.95 / 0.487 / 0.715	24.48 / 0.590 / 0.679	0.4076
MoCE-IR	**25.88** / 0.586 / 0.629	17.12 / 0.418 / 0.683	22.04 / 0.552 / 0.627	0.4705
MIRAGE	25.61 / 0.600 / 0.619	17.11 / 0.423 / 0.700	22.11 / 0.549 / 0.618	0.4211
BARS-MambaIR	25.87 / **0.611 / 0.609**	**24.74 / 0.590 / 0.668**	**25.01 / 0.588 / 0.614**	**0.3904**

3.4 Action-Consistent Readout and Residual-Derived Supervision

As illustrated in Figure 2(d), to align image-space correction with the propagation decision, we apply a lightweight zero-initialized residual refinement to the coarse backbone output $\boldsymbol{c}$, with $k \in \{\mathrm{g}, \mathrm{m}, \mathrm{l}\}$ indexing the three propagation components of $\boldsymbol{\pi}$ in Eq. 5. The final restoration is written compactly as

$$\widehat{\boldsymbol{y}} = \boldsymbol{x} + \boldsymbol{d} \odot \boldsymbol{v} \odot \boldsymbol{r}. \qquad \boldsymbol{r} = \sum_{k \in \{\mathrm{g},\mathrm{m},\mathrm{l}\}} \boldsymbol{\pi}_k \odot \left[(\boldsymbol{c} - \boldsymbol{x}) + \delta_k \right], \tag{14}$$

where δ_k denotes the basis-specific residual refinement. Thus, the same propagation weights $\boldsymbol{\pi}$ used in state transport also guide the final image-space correction, while $\boldsymbol{d}$ and $\boldsymbol{v}$ determine where and how strongly the image is modified.

For each spatial position p, we learn the action coordinates without degradation labels or artifact masks. Given the paired residual r_p, we construct the target action coordinates as

$$\left(d_p^\star, s_p^\star, v_p^\star\right) = \left(1 - \exp\left(-\frac{r_p}{\tau_d + \varepsilon}\right),\ \frac{r_{\mathrm{lo},p}}{r_{\mathrm{lo},p} + r_{\mathrm{hi},p} + \varepsilon},\ \frac{r_p}{r_p + \tau_v + \varepsilon}\right), \tag{15}$$

where r_p is the noise-suppressed discrepancy between $\boldsymbol{x}$ and $\boldsymbol{y}$, while $r_{\mathrm{lo},p}$ and $r_{\mathrm{hi},p}$ denote its locally smooth and spatially concentrated components, respectively. The scale factors τ_d and τ_v control the sensitivity of the demand and magnitude targets to the residual amplitude. Following Eq. 4, the corresponding target allocation is $\boldsymbol{q}_p^\star = \left[1 - d_p^\star,\ d_p^\star s_p^\star,\ d_p^\star(1 - s_p^\star)\right]$. The overall training objective is

$$\mathcal{L} = \frac{\mathcal{L}_{\mathrm{rec}} + \lambda_{\mathrm{alloc}} \operatorname{Charb}(\boldsymbol{q}, \boldsymbol{q}^\star) + \lambda_{\mathrm{mag}} \operatorname{Charb}(\boldsymbol{v}, \boldsymbol{v}^\star)}{1 + \lambda_{\mathrm{alloc}} + \lambda_{\mathrm{mag}}}, \tag{16}$$

where $\mathcal{L}_{\mathrm{rec}}$ is the Charbonnier reconstruction loss between $\widehat{\boldsymbol{y}}$ and $\boldsymbol{y}$.

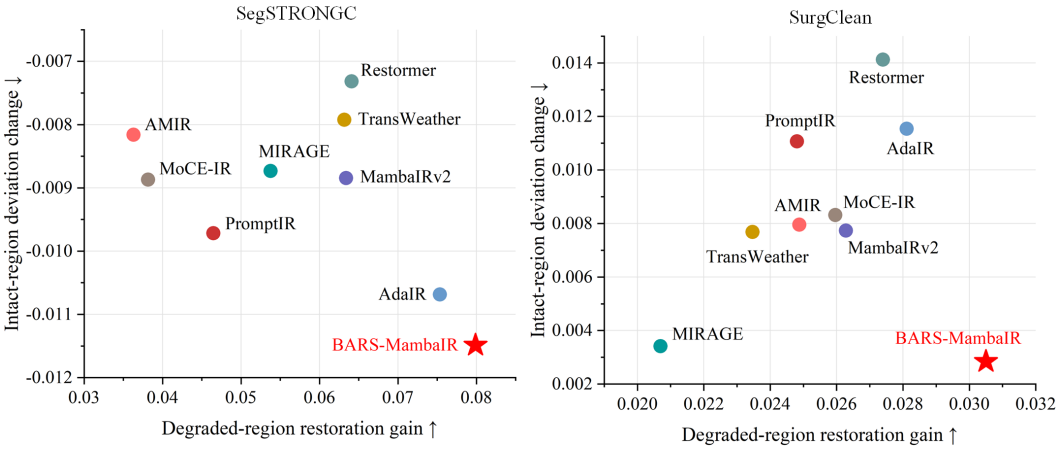

Figure 3: Restoration–preservation trade-off. Rightward indicates larger degraded-region gain and downward indicates smaller intact-region error change. BARS-MambaIR lies on the Pareto envelope on both benchmarks.

4 Experiments

4.1 Experimental Protocol

We evaluate blind all-in-one restoration on SurgClean (Pei et al., 2026), containing defog, desmoke, and desplash pairs, and SegSTRONG-C (Ding et al., 2024), using its blood, low-brightness, and smoke domains. No degradation identity is provided at inference. Alongside PSNR (Huynh-Thu & Ghanbari, 2008), SSIM (Wang et al., 2004), and LPIPS (Zhang et al., 2018), we report degraded-region gain, intact-region preservation delta, and Spill (Eq. 2). We compare BARS-MambaIR with Restormer (Zamir et al., 2022), TransWeather (Valanarasu et al., 2022), PromptIR (Potlapalli et al., 2023), AMIR (Yang et al., 2024), AdaIR (Cui et al., 2025), MambaIRv2 (Guo et al., 2025), MIRAGE (Ren et al., 2026), and MoCE-IR (Zamfir et al., 2025). For a fair comparison, all methods are trained and evaluated on identical task unions and data splits for 80,000 iterations using (128×128) crops, a batch size of 8, and Adam ($\beta_1 = 0.9$, $\beta_2 = 0.99$). The initial learning rate is (1.5×10^{-4}) and is halved after 40,000 iterations.

4.2 Restoration Quality versus Intact-Tissue Preservation

Table 1 shows that BARS-MambaIR improves restoration without merely increasing action everywhere. On SurgClean, it obtains the best PSNR for defog, desmoke, and desplash (19.55, 20.24, and 21.54 dB), exceeding the next-best result by 0.06, 0.41, and 0.12 dB, respectively. On SegSTRONG-C, it is essentially tied on blood (25.87 versus 25.88 dB), while improving low-brightness from 20.20 to 24.74 dB and smoke from 24.48 to 25.01 dB. Importantly, BARS-MambaIR also achieves the lowest Spill on both datasets, with 0.3863 on SurgClean and 0.3904 on SegSTRONG-C. Together, these results show that BARS-MambaIR couples stronger restoration with more spatially selective corrective actions. Figure 3 makes this balance explicit. BARS-MambaIR lies on the favorable Pareto envelope on both datasets, combining the largest degraded-region gain with competitive intact-region change. This directly supports the motivation in Sec. 1: recognizing the degradation is insufficient when the induced correction can cross its valid spatial support.

4.3 Where Does the Restoration Action Go?

The signed-distance profiles in Fig. 4(a) connect the aggregate result to the boundary behavior introduced in Fig. 1. BARS-MambaIR maintains a strong response on the degraded side and sharply attenuates it after crossing into intact tissue, rather than suppressing restoration globally. The intervention curves in Fig. 4(b) show a clear dose response to demand perturbations. Increasing either demand shuffling or donor replacement monotonically reduces degraded-region gain and increases Spill. At full dose, demand shuffling decreases the gain from 0.0305 to 0.0270 and increases Spill

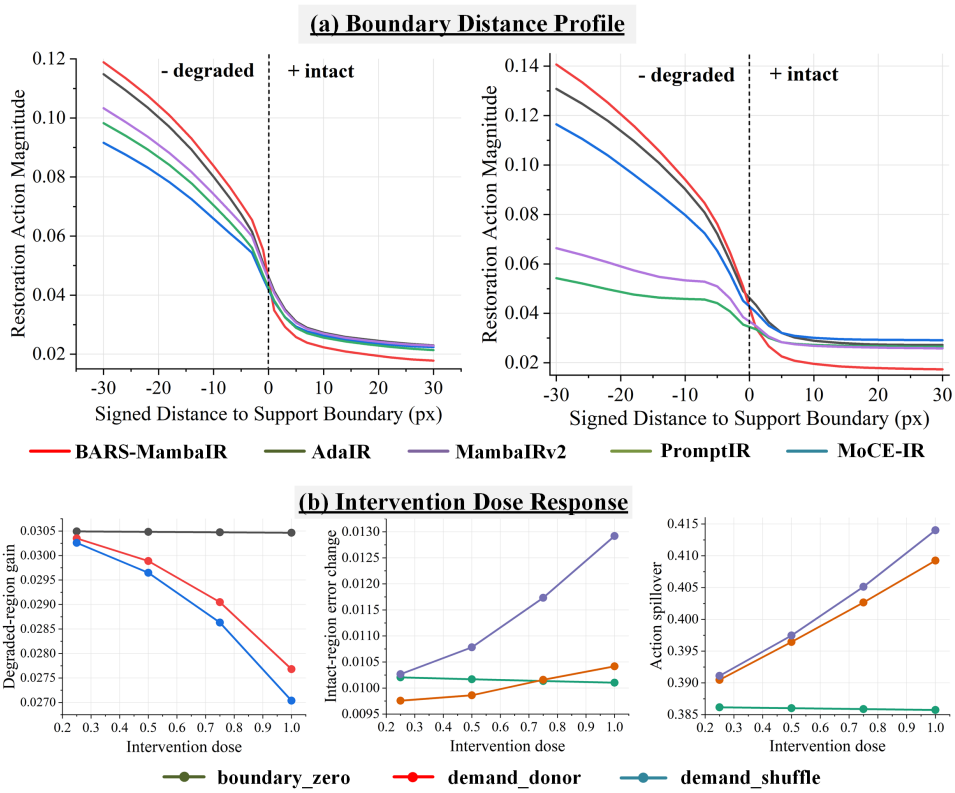

Figure 4: Mechanistic analysis of action localization. (a) Mean restoration-action magnitude versus signed distance to the evaluation-only support boundary; negative and positive distances denote degraded and intact regions. (b) Dose responses to demand-field corruption and boundary removal on SurgClean. Shuffling or donor-replacing demand progressively reduces gain and increases spillover, whereas removing the explicit boundary term has a smaller effect.

from 0.3863 to 0.4140, while donor replacement produces the same trend, yielding 0.0277 and 0.4093, respectively. Because shuffling preserves the distribution of demand values while disrupting their spatial correspondence, these results highlight the importance of assigning restoration demand to the correct locations. Boundary removal, by comparison, causes only minor changes in the aggregate metrics. Overall, spatially aligned demand is critical for effective and well-confined restoration.

4.4 Geometry-Dependent Restoration and Routing

Figure 5 reveals distinct dependencies on degradation severity and footprint. Degraded-region gain is more strongly associated with severity than with support size: fog and smoke generally improve as severity increases, whereas synthetic splash remains negative throughout the tested grid. Measured Spill is highest for compact corruptions and decreases consistently as the footprint expands. Because the intact-region complement in the Spill definition also shrinks with support size, this trend reflects both restoration-action allocation and metric geometry.

The routing statistics in Fig. 6 validate the role assigned to scope. For compact SurgClean supports, the local path receives 60% of the allocation, compared with 18% global and 21% mixed; for the largest supports, the distribution shifts to 26% local, 34% global, and 40% mixed. SegSTRONG-C exhibits an even stronger transition: the global share rises from 15% to 78%, while the local share falls from 52% to 6%. Forced-global and forced-local interventions both reduce degraded-region gain, confirming that neither endpoint is uniformly optimal. These results link the controller coordinates in Sec. 3 to an observable change in propagation range.

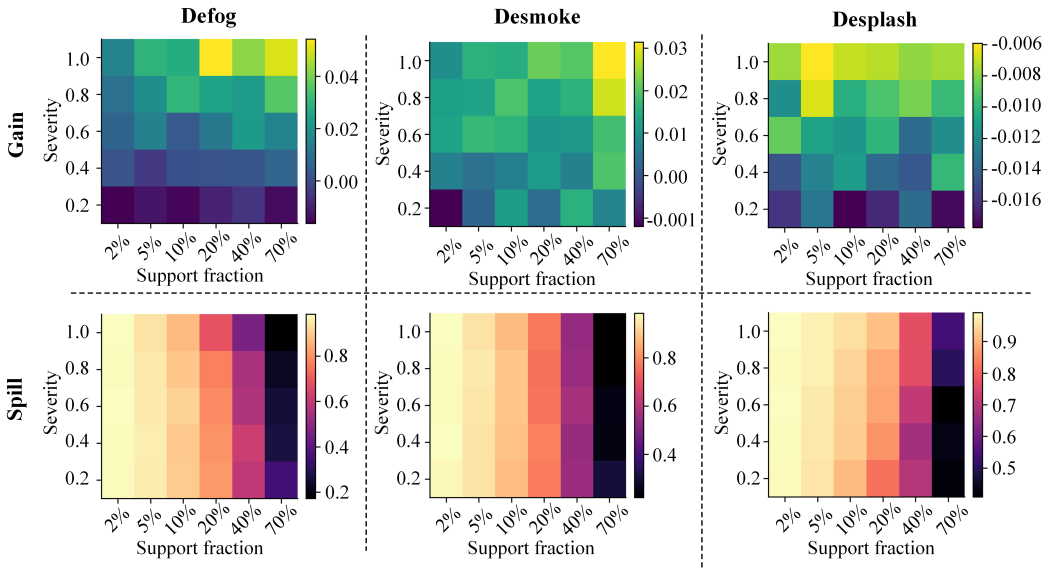

Figure 5: Controlled support–severity stress test. Degraded-region gain and Spill are reported over synthetic fog, smoke, and splash geometries. Fog and smoke improve with severity in most settings, whereas the uniformly negative splash gain exposes an unresolved failure case. Spill trends should be interpreted jointly with support fraction because the intact complement changes by construction.

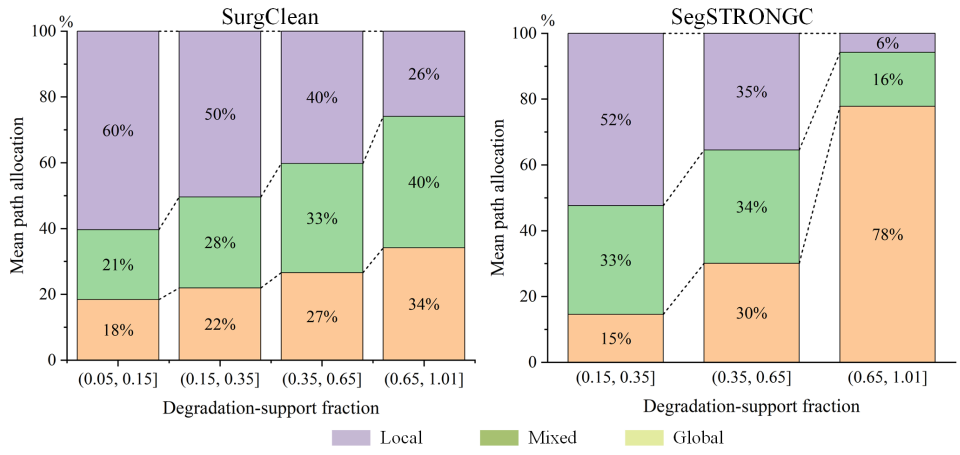

Figure 6: Support-dependent action routing. Mean global, mixed, and local path allocations are grouped by degradation-support fraction. Compact supports favor local transport, whereas larger supports shift mass toward mixed and global paths, consistent with the intended role of the predicted scope coordinate.

5 Conclusion

We introduced BARS-MambaIR, a blind controller that regulates where, how far, and how strongly restoration action propagates through a state-space restorer. Across SurgClean and SegSTRONG-C, it improves degraded-region restoration while attaining the lowest Spill. Demand interventions demonstrate the importance of spatial alignment, while routing statistics reveal support-dependent path selection. These results establish controlled action transport as a useful perspective.

AI USE STATEMENT

We used a large language model (LLM) for language polishing and manuscript editing, including sentence rewriting, grammar checking, and improving fluency. The LLM did not participate in ideation, research design, experiments, or data analysis, and was not used for other tasks requiring disclosure. All scientific content was independently developed by the authors, who take full responsibility for the paper and ensure that AI-assisted text complies with ethical guidelines.

ETHICS STATEMENT

We have read and acknowledge the ICLR Code of Ethics and comply with its guidelines. This study uses the publicly released, fully de-identified dataset. The work does not involve human-subject intervention or new data collection. We believe the presented work poses no identifiable ethical risk.

REPRODUCIBILITY STATEMENT

We have made every effort to ensure the reproducibility of the results presented in this paper. The experimental settings are described in detail in the paper to facilitate replication.

REFERENCES

Tong Chen, Xinyu Ma, Long Bai, Wenyang Wang, Yue Sun, and Luping Zhou. EndoIR: Degradation-agnostic all-in-one endoscopic image restoration via noise-aware routing diffusion. *arXiv preprint arXiv:2511.05873*, 2025. URL https://arxiv.org/abs/2511.05873.

Yuning Cui, Syed Waqas Zamir, Salman Khan, Alois Knoll, Mubarak Shah, and Fahad Shahbaz Khan. AdaIR: Adaptive all-in-one image restoration via frequency mining and modulation. In *International Conference on Learning Representations*, 2025. URL https://arxiv.org/abs/2403.14614.

Hao Ding, Yuqian Zhang, Tuxun Lu, Ruixing Liang, Hongchao Shu, Lalithkumar Seenivasan, Yonghao Long, Qi Dou, Cong Gao, Mathias Unberath, et al. SegSTRONG-C: Segmenting surgical tools robustly on non-adversarial generated corruptions—an EndoVis'24 challenge. *arXiv preprint arXiv:2407.11906*, 2024.

Hu Gao, Changshuo Wang, Yulong Chen, and Lizhuang Ma. Learning adaptive dynamical features via multi-τ liquid-mamba for all-in-one image restoration. *arXiv preprint arXiv:2606.22801*, 2026. URL https://arxiv.org/abs/2606.22801.

Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv preprint arXiv:2312.00752*, 2024. URL https://arxiv.org/abs/2312.00752.

Hang Guo, Jinmin Li, Tao Dai, Zhihao Ouyang, Xudong Ren, and Shu-Tao Xia. MambaIR: A simple baseline for image restoration with state-space model. In *European Conference on Computer Vision*, pp. 222–241. Springer, 2024. URL https://arxiv.org/abs/2402.15648.

Hang Guo, Yong Guo, Yaohua Zha, Yulun Zhang, Wenbo Li, Tao Dai, Shu-Tao Xia, and Yawei Li. MambaIRv2: Attentive state space restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 28124–28133, 2025. URL https://arxiv.org/abs/2411.15269.

Quan Huynh-Thu and Mohammed Ghanbari. Scope of validity of psnr in image/video quality assessment. *Electronics Letters*, 44(13):800–801, 2008. doi: 10.1049/el:20080522.

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. URL https://openaccess.thecvf.com/content/CVPR2022/html/Li_All-in-One_Image_Restoration_for_Unknown_Corruption_CVPR_2022_paper.html.

Cencen Liu, Wen Yin, Dongyang Zhang, Dongmin Li, Shan Zhao, Bing Su, Tao He, Jielei Wang, and Guoming Lu. What to remove, what to preserve: Dual-ambiguity rectification for all-in-one image restoration. *arXiv preprint arXiv:2607.28526*, 2026. URL https://arxiv.org/abs/2607.28526.

Zhanwen Liu, Sai Zhou, Yuchao Dai, Yang Wang, Yisheng An, and Xiangmo Zhao. DPMambaIR: All-in-one image restoration via degradation-aware prompt state space model. *arXiv preprint arXiv:2504.17732*, 2025. URL https://arxiv.org/abs/2504.17732.

Jialun Pei, Diandian Guo, Donghui Yang, Zhixi Li, Yuxin Feng, Long Ma, Bo Du, and Pheng-Ann Heng. Benchmarking endoscopic surgical image restoration and beyond. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2026. URL https://arxiv.org/abs/2505.19161.

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. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/e187897ed7780a579a0d76fd4a35d107-Abstract-Conference.html.

Bin Ren, Yawei Li, Xu Zheng, Yuqian Fu, Danda Pani Paudel, Hong Liu, Ming-Hsuan Yang, Luc Van Gool, and Nicu Sebe. Efficient degradation-agnostic image restoration via channel-wise functional decomposition and manifold regularization. In *International Conference on Learning Representations*, 2026.

Jeya Maria Jose Valanarasu, Rajeev Yasarla, and Vishal M. Patel. Transweather: Transformer-based restoration of images degraded by adverse weather conditions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2353–2363, 2022.

Shiwei Wang, Yongzhen Wang, Bingwen Hu, Liyan Zhang, Xiao-Ping Zhang, and Mingqiang Wei. M2IR: Proactive all-in-one image restoration via mamba-style modulation and mixture-of-experts. *arXiv preprint arXiv:2603.14816*, 2026. URL https://arxiv.org/abs/2603.14816.

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. doi: 10.1109/TIP.2003.819861.

Zhiwen Yang, Haowei Chen, Ziniu Qian, Yang Yi, Hui Zhang, Dan Zhao, Bingzheng Wei, and Yan Xu. All-in-one medical image restoration via task-adaptive routing. In *Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, volume 15007 of *Lecture Notes in Computer Science*, pp. 67–77. Springer Nature Switzerland, 2024. doi: 10.1007/978-3-031-72104-5_7. URL https://papers.miccai.org/miccai-2024/063-Paper0787.html.

Eduard Zamfir, Zongwei Wu, Nancy Mehta, Yuedong Tan, Danda Pani Paudel, Yulun Zhang, and Radu Timofte. Complexity experts are task-discriminative learners for any image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2025. URL https://arxiv.org/abs/2411.18466.

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.

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.

Shen Zhou, Jinghui Zhang, Wenbo Huang, Xuwei Qian, Zhen Wu, Guangwen Peng, Zhiyuan Li, Ding Ding, Dian Shen, and Fang Dong. QuReC: All-in-one image restoration with query-specific guidance and local-global response calibration. *arXiv preprint arXiv:2607.15097*, 2026. URL https://arxiv.org/abs/2607.15097.