

RETHINKING EXPRESSIVITY AND DEGRADATION-AWARENESS IN ATTENTION FOR ALL-IN-ONE BLIND IMAGE RESTORATION

Anonymous authors

Paper under double-blind review

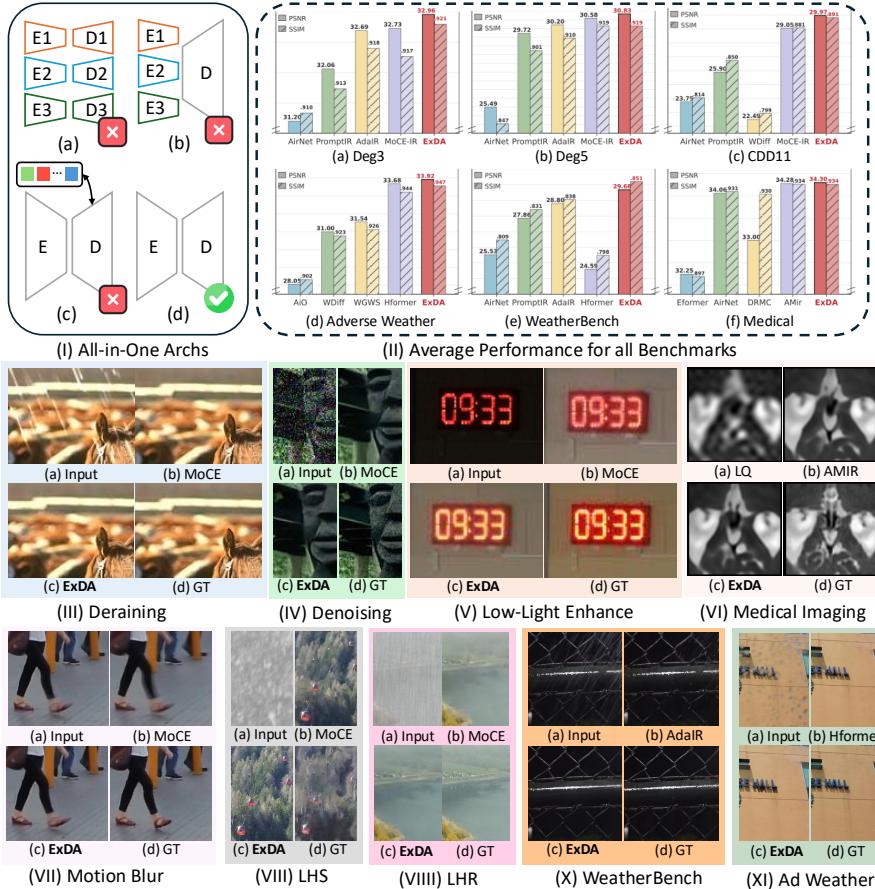


Figure 1: (I) Architectures: We propose that the simple encoder-decoder (d) pipeline is strong enough for All-in-One IR. (II) Average Performance: our ExDA consistently produced better results across six challenging All-in-One benchmarks. (III–XI) Qualitative Comparison: Please zoom in for details.

ABSTRACT

All-in-one image restoration (IR) aims to recover high-quality images from diverse degradations, which in real-world settings are often mixed and unknown. Unlike single-task IR, this problem requires a model to approximate a family of heterogeneous inverse functions, making it fundamentally more challenging and practically important. Although recent focus has shifted toward large multimodal models, their robustness still depends on faithful low-level inputs, and the principles that govern effective restoration remain underexplored. We revisit attention mechanisms through the lens of all-in-one IR and identify two overlooked bottlenecks in widely adopted Restormer-style backbones: (i) *the value path remains purely linear*, restricting outputs to the span of inputs and weakening expressivity, and (ii) *the absence of an explicit global slot* prevents attention from encoding degradation context. To address these issues, we propose two minimal, backbone-agnostic primitives: a nonlinear value transform that upgrades attention from a

054 selector to a selector–transformer, and a global spatial token that provides an ex-
 055 plicit degradation-aware slot. Together, these additions improve restoration across
 056 synthetic, mixed, underwater, and medical benchmarks, with negligible overhead
 057 and consistent performance gains. Analyses with foundation model embeddings,
 058 spectral statistics, and separability measures further clarify their roles, positioning
 059 our study as a step toward rethinking attention primitives for robust all-in-one IR.
 060
 061

1 INTRODUCTION

064 Image restoration (*i.e.*, IR) aims to recover high-quality images from degraded observations. In
 065 real-world scenarios, images often suffer from complex degradations—noise, blur, haze, rain, and
 066 their unpredictable mixtures—each corresponding to a distinct inverse function that must be learned
 067 and applied adaptively. This multiplicity of inverse mappings makes all-in-one image restoration
 068 fundamentally more challenging than single-task variants: rather than learning one specialized
 069 function, the model must approximate a *family of heterogeneous inverse transformations* while
 070 remaining robust to mixed and unseen degradations. The stakes for robust low-level vision have never
 071 been higher. As computer vision increasingly relies on large multimodal models and foundation
 072 architectures, their effectiveness hinges critically on the quality of input images. Corrupted inputs
 073 propagate errors throughout the entire pipeline, making reliable image restoration not merely a
 074 specialized tool, but a *foundational prerequisite* for robust vision systems.

075 Recent advances are dominated by Transformer-based architectures that skillfully balance global
 076 context modeling with local spatial priors. The multi-dconv head transposed attention (MDTA) and
 077 gated-dconv feed-forward network (GDFN) introduced in Restormer (Zamir et al., 2022b) exemplify
 078 this trend, replacing computationally prohibitive token-wise self-attention with efficient *channel-wise*
 079 *attention*. This design has become the de-facto standard for high-resolution IR, spawning numerous
 080 variants and establishing strong empirical baselines across diverse restoration tasks. However, when
 081 viewed through the lens of all-in-one IR, this prevalent design exposes two critical yet overlooked
 082 limitations that reflect a deeper tension between single-task efficiency and unified IR capability.

083 First, the attention mechanism’s value path remains purely linear: while queries and keys interact
 084 through nonlinear softmax operations, values are merely linearly aggregated, *constraining outputs* to lie within the convex hull of input features. This expressivity bottleneck becomes particularly
 085 severe in all-in-one settings, where the model must navigate between vastly different inverse map-
 086 pings—from high-frequency noise removal to low-frequency haze correction—yet is constrained by
 087 linear combinations of its inputs. Although (Katharopoulos et al., 2020; Shen et al., 2024; Aksenov
 088 et al., 2024; Shazeer, 2020) use of nonlinearities on queries and keys enables an efficient linearized
 089 softmax, we argue that the value space is more critical for learning robust representations. Second,
 090 unlike standard ViTs that employ class tokens for global aggregation, channel-wise attention *discards*
 091 the notion of explicit global slots entirely. This forces the degradation context to be encoded implicitly
 092 across spatial channels, making the model less capable of explicit degradation inference—a capability
 093 that proves essential when corruption types are unknown and mixed. These limitations matter less
 094 in single-task settings where the inverse function is fixed and known, but become fundamental
 095 bottlenecks in all-in-one scenarios that demand both expressivity and adaptability.

096 In this work, we address these bottlenecks through a principled rethinking of attention primitives for
 097 all-in-one IR. We introduce two minimal, backbone-agnostic extensions that transform any Restormer-
 098 style architecture into a more expressive and degradation-aware system. Our **nonlinear value**
 099 **transform** breaks the linear span constraint by augmenting values with lightweight convolutional
 100 mappings *before* aggregation, upgrading attention from a feature selector to a feature *selector-
 101 transformer*. Our **global spatial token (GST)** restores explicit global context by providing learnable
 102 slots that adaptively aggregate degradation-relevant statistics and inject them back into the attention
 103 computation. These innovations are both theoretically motivated and empirically validated: from
 104 a function approximation perspective, we show that pre-aggregation nonlinear transformations
 105 fundamentally expand the realizable function family, while diagnostic analyses involving foundation
 106 model embeddings, spectral decompositions, and cluster separability measures demonstrate that
 107 global tokens capture meaningful degradation context. Extensive experiments across synthetic, mixed,
 108 underwater, and medical benchmarks reveal consistent improvements over strong baselines, with
 109 negligible computational overhead and robust transfer across different backbone architectures.

108 Our contributions address a fundamental gap in understanding attention mechanisms for unified IR:
 109

110

- 111 • We identify and theoretically analyze two overlooked bottlenecks in prevalent Restormer-
 112 style architectures—linear value constraints that limit expressivity and absent global slots
 113 that limit degradation-awareness.
- 114 • We propose two minimal, backbone-agnostic primitives that address these bottlenecks:
 115 nonlinear value transforms that expand the realizable function family, and global spatial
 116 tokens that provide explicit degradation context.
- 117 • We provide theoretical motivation, diagnostic analyses, and extensive empirical validation
 118 across diverse benchmarks, demonstrating consistent improvements while maintaining
 119 computational efficiency.

120

121 **2 RELATED WORK**

122

124 **Image Restoration (IR).** IR aims to solve a highly ill-posed problem: reconstructing high-quality
 125 images from degraded observations. Owing to its broad importance, IR has been widely applied
 126 in numerous applications (Richardson, 1972; Xie et al., 2025; Banham & Katsaggelos, 1997; Li
 127 et al., 2023b; Zamfir et al., 2024). Early approaches were dominated by model-based solutions that
 128 searched for closed-form results under handcrafted formulations. With the advent of deep learning,
 129 learning-based IR methods have rapidly gained popularity. Recent studies include regression-based
 130 techniques (Lim et al., 2017; Lai et al., 2017; Liang et al., 2021; Chen et al., 2021; Li et al., 2023a;
 131 Zhang et al., 2024) and generative pipelines (Gao et al., 2023; Wang et al., 2023b; Luo et al., 2023;
 132 Yue et al., 2023; Zhao et al., 2024), built upon a variety of architectures such as convolutional
 133 networks (Dong et al., 2015; Zhang et al., 2017b;a; Wang et al., 2018), MLP-based designs (Tu
 134 et al., 2022), state space models (Guo et al., 2024a; Zhu et al., 2024; Gu & Dao, 2023; Dao & Gu,
 135 2024), and Vision Transformers (ViTs) (Liang et al., 2021; Ren et al., 2023a; Li et al., 2023a; Zamir
 136 et al., 2022b; Dosovitskiy et al., 2020; Liu et al., 2023). Among them, Restormer (Zamir et al.,
 137 2022b) introduced a channel-wise attention mechanism (MDTA) to achieve linear complexity while
 138 handling high-resolution inputs, and has since become a widely adopted backbone in IR due to its
 139 strong balance of efficiency and accuracy. Despite these advances, the majority of IR models are still
 140 designed for specific degradation types, such as denoising (Zhang et al., 2017b; 2019), dehazing (Ren
 141 et al., 2020; Wu et al., 2021), deraining (Jiang et al., 2020; Ren et al., 2019), and deblurring (Kong
 142 et al., 2023; Ren et al., 2023b).

143 **All-in-One Blind Image Restoration.** Training separate task-specific models for individual degradations
 144 can deliver strong results, yet maintaining one model per task is impractical and environmentally
 145 costly. In practice, images are frequently affected by mixtures of degradations, making it unrealistic to
 146 handle each corruption independently. This has motivated the study of *All-in-One image restoration*,
 147 where a single model is expected to generalize across multiple degradation types (Zamfir et al., 2025;
 148 Zeng et al., 2025; Zheng et al., 2024). Different strategies have been explored: AirNet (Li et al.,
 149 2022) learns contrastive degradation representations to guide reconstruction, whereas IDR (Zhang
 150 et al., 2023) formulates degradations as physical components and leverages a meta-learning pipeline.
 151 Prompt-based approaches (Potlapalli et al., 2024; Wang et al., 2023a; Li et al., 2023c) push this idea
 152 further by conditioning restoration on learned visual prompts, later extended to frequency-aware
 153 prompts (Cui et al., 2025) or larger-scale dynamic architectures (Dudhane et al., 2024). Although
 154 effective, these methods often incur high training costs and reduced efficiency (Cui et al., 2025),
 155 which hinders practical deployment. **More recently, besides these regression-based approaches,**
 156 **distribution-oriented models also show decent performance while with a larger model size** Tian et al.
 157 (2025); Luo et al. (2025). In this work, we take a different route. Instead of relying on auxiliary
 158 prompt modules or complex extra multi-stage strategies, we revisit the backbone design itself. **In**
 159 **parallel, recent studies have attempted to enhance attention expressivity** (Chefer et al., 2021) and
 160 **incorporate degradation-aware priors directly into the backbone design** (Tang et al., 2025). These
 161 directions motivate a closer examination of the attention operator itself rather than adding heavier
 auxiliary modules. By examining Restormer-style channel-wise attention, we identify two overlooked
 bottlenecks, linear values and the lack of an explicit global slot, and propose minimal primitives that
 directly improve expressivity and provide degradation-aware context for efficient all-in-one Blind IR.

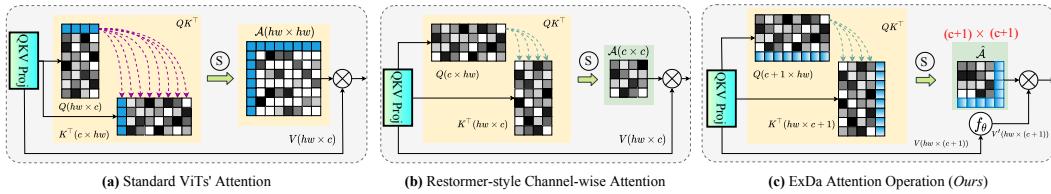


Figure 2: Comparison of (a) standard ViT attention with CLS in high-level tasks, (b) Restormer-style channel-wise attention in IR, (c) Our proposed attention design that augments values space with nonlinearity and introduces an explicit global slot.

3 THE PROPOSED METHOD

3.1 REVISITING RESTORMER-STYLE ATTENTION: HIDDEN BOTTLENECKS

Transformers have recently become central to low-level vision tasks such as image restoration (IR), where they are often re-designed to balance efficiency and spatial detail recovery. A representative example is Restormer (Zamir et al., 2022b), which replaces spatial token-wise self-attention (Fig. 2a) with *channel-wise attention* (Fig. 2b) and substitutes the standard MLP feed-forward layers with gated-dconv feed-forward networks. These design choices make the model computationally feasible for high-resolution inputs and have set a strong baseline for single-degradation IR.

However, in the broader context of low-level vision, two overlooked limitations emerge. (i), the *value (V) path remains purely linear*. While queries (Q) and keys (K) interact through a nonlinear softmax, values are only linearly aggregated, constraining outputs to lie within the span of the inputs. In the standard ViTs, this limitation is less severe because the MLP-based feed-forward networks provide strong nonlinear transformations. In contrast, IR backbones such as Restormer adopt gated-dconv feed-forward networks, where one branch is essentially linear. As a result, a portion of the information bypasses any nonlinear transformation, leaving the overall block with weaker nonlinearity. This makes the expressivity bottleneck of linear values particularly pronounced in *all-in-one restoration*, where the model must approximate a family of diverse inverse functions rather than a single degradation. (ii), *there is no explicit global slot to summarize degradation context*. Standard ViTs include a CLS token to aggregate global semantics (Fig. 2a), yet in low-level vision, this token is often discarded as “useless” for pixel-level predictions (Liang et al., 2021; Li et al., 2023a; Ren et al., 2024). Restormer’s channel-wise attention follows this practice (Fig. 2b), relying instead on local depth-wise convolutions. While sufficient for single-task restoration, this design implicitly assumes that the degradation type is fixed and known. In the all-in-one setting, however, the model must *infer* degradation type directly from the input. Without a dedicated global slot, there is no explicit mechanism to encode degradation statistics into the representation. What may appear redundant in single-task IR thus becomes indispensable in all-in-one IR, where such a slot can naturally evolve into a *degradation embedding*, capturing global statistics not only across channels but also explicitly across spatial structure (Fig. 2b). These two limitations motivate our subsequent exploration of nonlinear values and explicit global tokens.

3.2 NONLINEARITY MATTERS IN VALUES FOR EXPRESSIVITY

Building on the identified bottleneck in Sec. 3.1, we empirically validate the critical role of nonlinear values in all-in-one restoration. The linear value constraint severely limits the model’s ability to approximate diverse inverse mappings required across different degradations.

To isolate this effect, we design experiments with multi-faceted degradations that mirror all-in-one restoration complexity. Our *synthetic function* combines nonlinear sensor response, blur kernels, additive/multiplicative noise, and quantization effects. Our *MNIST restoration* applies realistic corruptions including nonlinear response curves, motion blur, spatially-varying noise, and compression artifacts. Both scenarios require modeling multiple degradation characteristics simultaneously, analogous to all-in-one restoration challenges.

Fig. 3 reveals consistent patterns across both settings. In the synthetic function approximation setting (a), linear value attention systematically fails in critical regions (orange highlights), achieving 50.4% worse convergence as it encounters inherent expressivity limitations.

216 Algorithm 1 GST: Content-Adaptive Global Spatial Token Generation

217 **Require:** Input features $X \in \mathbb{R}^{B \times C \times H \times W}$, heads h , tokens per head K , stride s
218 Ensure: Degradation-aware global tokens $G \in \mathbb{R}^{B \times h \times K \times HW}$
219 1: Efficient Spatial Compression: $\tilde{X} = \text{AvgPool}_s(X) \in \mathbb{R}^{B \times C \times \frac{H}{s} \times \frac{W}{s}}$
220 2: Multi-Head Token Projection: $\Phi = W_{\text{proj}} \star \tilde{X} \in \mathbb{R}^{B \times hK \times H_s \times W_s}$
221 3: Content-Adaptive Spatial Attention:
222 4: $G_{\text{compact}} = \text{Softmax}_{\text{spatial}}(\text{Reshape}(\Phi, [B, h, K, -1]))$
223 5: Resolution Recovery: $\hat{G} = \text{BilinearUpsample}(G_{\text{compact}}, \text{size} = (H, W))$
224 6: return $G \in \mathbb{R}^{B \times h \times K \times HW}$ ▷ Ready for attention concatenation

225
 226 MNIST restoration (b) confirms this on realistic image
 227 data, where nonlinear values achieve 5.92 dB PSNR
 228 improvement (19.2 → 25.1 dB) and produce visually su-
 229 perior reconstructions compared to blurry, artifact-laden
 230 outputs from linear value attention.

231 These findings confirm our analysis: in Restormer-style
 232 architectures where gated FFN allows information to
 233 bypass nonlinear transformations, the linear value con-
 234 straint becomes a critical expressivity bottleneck for
 235 all-in-one restoration requiring diverse degradation mod-
 236 eling. This motivates enhancing value expressivity
 237 through $V' = f_{\theta}(V)$, where f_{θ} introduces nonlinearity
 238 beyond the linear span of inputs. We adopt a residual
 239 formulation to balance preservation and transformation:

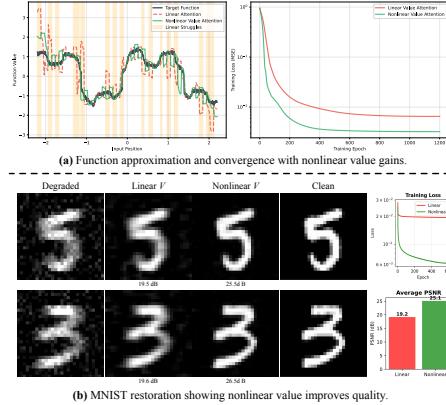
$$240 \\ 241 V' = V + g_{\theta}(V), \\ 242 g_{\theta} = \text{Conv}_{1 \times 1} \rightarrow \text{DWConv}_{3 \times 3} \rightarrow \text{GELU} \rightarrow \text{Conv}_{1 \times 1}. \\ 243 \quad (1)$$

244 Two insights guide this design. First, pre-aggregation
 245 placement is essential since attention $\text{Softmax}(QK^{\top}/\sqrt{d})V'$ constrains outputs to linear
 246 combinations—post-aggregation nonlinearity cannot escape this fundamental limitation. Second, the
 247 lightweight g_{θ} provides sufficient expressivity for diverse degradation modeling while maintaining
 248 efficiency. This transforms channel-wise attention from a linear feature selector into a nonlinear
 249 feature transformer, fundamentally addressing the expressivity gap between single-task and all-in-one
 250 restoration requirements.

252 3.3 GLOBAL CONTEXT AND THE MISSING SLOT

254 Building on the analysis in Sec. 3.1, we now address the second identified limitation: the absence
 255 of explicit global context mechanisms in Restormer-style architectures. While Sec. 3.2 tackled the
 256 expressivity bottleneck through nonlinear value transforms, the *degradation inference* challenge
 257 requires a fundamentally different approach, introducing explicit degradation-aware slots that can
 258 capture global spatial statistics. The core issue is that without dedicated global tokens, degradation
 259 context can only be distributed implicitly across local channel interactions or captured through
 260 cross-channel attention mechanisms. This becomes particularly problematic in all-in-one settings
 261 where the model must distinguish between fundamentally different corruption types. These global
 262 spatial signatures are best captured through explicit global analysis rather than relying solely on local
 263 depth-wise convolutions or implicit cross-channel attention.

264 We address this through *Global Spatial Tokens (GST)*, which provide explicit degradation-aware
 265 slots while maintaining computational efficiency. As detailed in Alg. 1, our design generates content-
 266 adaptive spatial attention maps through stride- s spatial compression and learnable projection. The
 267 key innovation lies in content-adaptive pooling: rather than fixed global averaging, GST learns spatial
 268 attention maps through learnable projections followed by spatial softmax normalization, enabling each
 269 token to naturally develop distinct spatial emphasis patterns that correspond to different degradation
 characteristics. This enables natural specialization during training—noise tokens activate scattered



251 Figure 3: Nonlinear value transforms out-
 252 perform linear value attention in both (a)
 253 function approximation and (b) MNIST
 254 restoration, achieving better convergence,
 255 higher PSNR/SSIM, and improved quality.

270 Table 1: *Comparison to state-of-the-art on three degradations.* PSNR (dB, \uparrow) and SSIM (\uparrow) metrics
 271 are reported on the full RGB images. **Best** performances is highlighted. ‘-’ means unreported results.
 272

273 Method	Venue.	Params.	Dehazing		Deraining			Denoising			Average
			SOTS	Rain100L	BSD68 $_{\sigma=15}$	BSD68 $_{\sigma=25}$	BSD68 $_{\sigma=50}$				
275 BRDNet (Tian et al., 2020)	NN’20	-	23.23	.895	27.42	.895	32.26	.898	29.76	.836	26.34
276 LPNet (Gao et al., 2019)	CVPR’19	-	20.84	.828	24.88	.784	26.47	.778	24.77	.748	21.26
277 FDGAN (Dong et al., 2020)	AAAI’20	-	24.71	.929	29.89	.933	30.25	.910	28.81	.868	26.43
278 DL (Fan et al., 2019)	TPAMI’19	2M	26.92	.931	32.62	.931	33.05	.914	30.41	.861	26.90
279 MPRNet (Zamir et al., 2021)	CVPR’21	16M	25.28	.955	33.57	.954	33.54	.927	30.89	.880	27.56
280 AirNet (Li et al., 2022)	CVPR’22	9M	27.94	.962	34.90	.967	33.92	.933	31.26	.888	28.00
281 NDR (Yao et al., 2024)	TIP’24	28M	25.01	.860	28.62	.848	28.72	.826	27.88	.798	26.18
282 PromptIR (Potlapalli et al., 2024)	NeurIPS’23	36M	30.58	.974	36.37	.972	33.98	.933	31.31	.888	28.06
283 MoCE-IR-S (Zamfir et al., 2025)	CVPR’25	11M	30.98	.979	38.22	.983	34.08	.933	31.42	.888	28.16
284 AdaIR (Cui et al., 2025)	ICLR’25	29M	31.06	.980	38.64	.983	34.12	.935	31.45	.892	28.19
285 MoCE-IR (Zamfir et al., 2025)	CVPR’25	25M	31.34	.979	38.57	.984	34.11	.932	31.45	.888	28.18
286 ExDA (<i>Ours</i>)	2025	22M	31.58	.982	39.13	.985	34.22	.937	31.56	.895	28.32
287											.808
288											.32.96
289											.921
290	Methods with the assistance of vision language, multi-task learning, natural language prompts, or multi-modal control										
291	DA-CLIP (Luo et al., 2024)	ICLR’24	125M	.29.46	.963	36.28	.968	30.02	.821	24.86	.585
292	Art _{Promp} ptIR (Wu et al., 2024)	ACM MM’24	36M	30.83	.979	37.94	.982	34.06	.934	31.42	.891
293	InstructIR-3D (Conde et al., 2024)	ECCV’24	16M	30.22	.959	37.98	.978	34.15	.933	31.52	.890
294	UniProcessor (Duan et al., 2025)	ECCV’24	1002M	31.66	.979	38.17	.982	34.08	.935	31.42	.891
295	VLU-Net (Zeng et al., 2025)	CVPR’25	35M	30.71	.980	38.93	.984	34.13	.935	31.48	.892
296	Perceive-IR (Zhang et al., 2025)	TIP’25	42M	30.87	.975	38.29	.980	34.13	.934	31.53	.890
297	DFPIR (Tian et al., 2025)	CVPR’25	31M	31.87	.980	38.65	.982	34.12	.935	31.47	.893
298	DA-RCOT (Tang et al., 2025)	TPAMI’25	50M	31.26	.977	38.36	.983	33.98	.934	31.33	.890
299											.801
300	high-frequency regions, blur tokens emphasize smooth low-frequency areas, and haze tokens respond to large-scale illumination structures, all without explicit supervision or degradation labels.										
301	The generated GST tokens $G \in \mathbb{R}^{B \times h \times K \times HW}$ are seamlessly integrated into channel-wise attention through direct concatenation with query, key, and value representations:										
302	$[Q^*, K^*, V^*] = [Q \oplus G, K \oplus G, V' \oplus G] \in \mathbb{R}^{B \times h \times (C+K) \times HW}$, (2)										
303	where \oplus denotes channel-wise concatenation and V' is the nonlinear-enhanced value from Sec. 3.2. After standard attention computation $\text{Attn} = \text{Softmax}(Q^* K^{*\top} / \tau) V^*$, we separate local channel features and global token contributions, re-injecting GST influence via learnable residual scaling:										
304	$\text{Output} = \text{Attn}_{[:, :, :, C, :]} + \alpha \cdot \text{Attn}_{[:, :, :, C, :, :]} \quad (3)$										
305	where α is initialized to 0.1 for gradual learning without overwhelming local features.										
306	This design transforms the discarded CLS token concept into a valuable primitive for all-in-one scenarios. By providing explicit degradation-aware slots that naturally specialize during training, GST enables coherent global strategies while maintaining channel-wise attention efficiency. The content-adaptive nature allows each attention head to learn specialized global patterns without requiring explicit degradation labels or supervisory signals.										
307											
308	4 EXPERIMENTS										
309											
310	We evaluate ExDA under six All-in-One IR benchmarks, <i>i.e.</i> , (i) <i>All-in-One (3Degradations)</i> , (ii) <i>All-in-One (5Degradations)</i> , (iii) <i>Mixed Degradation</i> , (iv) <i>Adverse Weather Removal</i> , (v) <i>Real-world WeatherBench</i> , and (vi) <i>Medical All-in-One</i> . The macro architecture of our model, the datasets, and implementation details are provided in our appendix.										
311											
312	4.1 MAIN RESULTS.										
313											
314	3-Degradation Setting. The results in Tab. 1 show that ExDA consistently outperforms other methods, even outperforming those with the assistance of language, multi-task, or text/visual prompts. Notably, 0.23 dB higher than the recent MoCE-IR on PSNR, while with 3M fewer parameters.										
315											
316	5-Degradation Setting. Building on the previous setting, we further include deblurring and low-light enhancement as two additional degradation types (Li et al., 2022; Zhang et al., 2023). As shown in Tab. 2, our method achieves the highest average PSNR, surpassing even substantially larger models, as well as those leveraging additional modalities or pretraining. Moreover, for dehazing and deblurring, our approach also delivers competitive results compared to existing methods.										
317											
318											
319											
320											
321											
322											
323											

324
 325
 326
 327
 328
 329
 330
 331
 332
 333
 334
 335
 336
 337
 338
 339
 340
 341
 342
 343
 344
 345
 346
 347
 348
 349
 350
 351
 352
 353
 354
 355
 356
 357
 358
 359
 360
 361
 362
 363
 364
 365
 366
 367
 368
 369
 370
 371
 372
 373
 374
 375
 376
 377
 Table 2: *Comparison to state-of-the-art on five degradations.* PSNR (dB, \uparrow) and SSIM (\uparrow) metrics are reported on the full RGB images with (*) denoting general image restorers, others are specialized all-in-one approaches. **Best** performance is highlighted.

Method	Venue	Params.	Dehazing		Deraining		Denoising		Deblurring		Low-Light		Average	
			SOTS	Rain100L	BSD68 $_{\sigma=25}$	GoPro	LOLvl							
NAFNet* (Chen et al., 2022a)	ECCV'22	17M	25.23	.939	35.56	.967	31.02	.883	26.53	.808	20.49	.809	27.76	.881
DGUINet* (Mou et al., 2022)	CVPR'22	17M	24.78	.940	36.62	.971	31.10	.883	27.25	.837	21.87	.823	28.32	.891
SwinIR* (Liang et al., 2021)	ICCVW'21	1M	21.50	.891	30.78	.923	30.59	.868	24.52	.773	17.81	.723	25.04	.835
Restormer* (Zamir et al., 2022b)	CVPR'22	26M	24.09	.927	34.81	.962	31.49	.884	27.22	.829	20.41	.806	27.60	.881
MambaIR* (Guo et al., 2024a)	ECCV'24	27M	25.81	.944	36.55	.971	31.41	.884	28.61	.875	22.49	.832	28.97	.901
DL (Fan et al., 2019)	TPAMI'19	2M	20.54	.826	21.96	.762	23.09	.745	19.86	.672	19.83	.712	21.05	.743
Transweather	CVPR'22	38M	21.32	.885	29.43	.905	29.00	.841	25.12	.757	21.21	.792	25.22	.836
TAPE (Liu et al., 2022)	ECCV'22	1M	22.16	.861	29.67	.904	30.18	.855	24.47	.763	18.97	.621	25.09	.801
AirNet (Li et al., 2022)	CVPR'22	9M	21.04	.884	32.98	.951	30.91	.882	24.35	.781	18.18	.735	25.49	.847
IDR (Zhang et al., 2023)	CVPR'23	15M	25.24	.943	35.63	.965	31.60	.887	27.87	.846	21.34	.826	28.34	.893
PromptIR (Potlapalli et al., 2024)	NeurIPS'23	36M	30.41	.972	36.17	.970	31.20	.885	27.93	.851	22.89	.829	29.72	.901
MoCE-IR-S (Zamfir et al., 2025)	CVPR'25	11M	31.33	.978	37.21	.978	31.25	.884	28.90	.877	21.68	.851	30.08	.913
AdaIR (Cui et al., 2025)	ICLR'25	29M	30.53	.978	38.02	.981	31.35	.889	28.12	.858	23.00	.845	30.20	.910
MoCE-IR (Zamfir et al., 2025)	CVPR'25	25M	30.48	.974	38.04	.982	31.34	.887	30.05	.899	23.00	.852	30.58	.919
ExDA (<i>Ours</i>)	2025	22M	31.14	.981	39.23	.985	31.52	.893	29.06	.878	23.19	.859	30.83	.919
Methods with the assistance of natural language prompts or multi-task learning														
InstructIR-5D (Conde et al., 2024)	ECCV'24	16M	36.84	.973	27.10	.956	31.40	.887	29.40	.886	23.00	.836	29.55	.908
ArtPromptIR (Wu et al., 2024)	ACM MM'24	36M	29.93	.908	22.09	.891	29.43	.843	25.61	.776	21.99	.811	25.81	.846
VLU-Net (Zeng et al., 2025)	CVPR'25	35M	30.84	.980	38.54	.982	31.43	.891	27.46	.840	22.29	.833	30.11	.905
Perceive-IR (Zhang et al., 2025)	TIP'25	42M	28.19	.964	37.25	.977	31.44	.887	29.46	.886	22.88	.833	29.84	.909
DFPIR (Tian et al., 2025)	CVPR'25	31M	31.64	.979	37.62	.978	31.29	.889	28.82	.873	23.82	.843	30.64	.913
DA-RCOT (Tang et al., 2025)	TPAMI'25	50M	30.96	.975	37.87	.980	31.23	.888	28.68	.872	23.25	.836	30.40	.911

Table 3: *Comparison on composited degradations.* PSNR (dB, \uparrow) and SSIM (\uparrow) are reported.

Method	Params.	CDD11-Single				CDD11-Double				CDD11-Triple				Avg.	
		Low (L)	Haze (H)	Rain (R)	Snow (S)	L+H	L+R	L+S	H+R	H+S	L+H+R	L+H+S			
AirNet	9M	24.83	.778	24.21	.951	26.55	.891	26.79	.919	23.23	.779	22.82	.710	23.29	.901
PromptIR	36M	26.32	.805	26.10	.969	31.56	.946	31.53	.960	24.49	.789	25.05	.771	24.51	.761
WGWSNet	26M	24.39	.774	27.90	.982	33.15	.964	34.43	.973	24.27	.780	25.06	.772	24.60	.765
WeatherDiff	83M	23.58	.763	21.99	.904	24.85	.885	24.80	.888	21.83	.756	22.69	.730	22.12	.707
OneRestore	6M	26.48	.826	32.52	.990	33.40	.964	34.31	.973	25.79	.782	25.58	.799	25.19	.789
MoCE-IR	11M	27.26	.824	32.66	.990	34.31	.970	35.91	.980	26.24	.817	26.25	.800	26.04	.793
ExDA (<i>Ours</i>)	22M	27.60	.837	34.44	.991	35.33	.973	36.85	.981	26.95	.834	26.80	.817	31.17	.968

Table 4: Comparisons for 4-task *adverse weather removal*. Missing values are denoted by '-'.

Method	Venue	Snow100K-S		Snow100K-L		Outdoor-Rain		RainDrop		Average	
		PSNR	SSIM								
All-in-One (Li et al., 2020)	CVPR'20	—	—	28.33	.882	24.71	.898	31.12	.927	28.05	.902
TransWeather (Valanarasu et al., 2022a)	CVPR'22	32.51	.934	29.31	.888	.900	.900	30.17	.916	30.20	.909
Chen et al. (Chen et al., 2022b)	CVPR'22	34.42	.947	30.22	.907	29.27	.915	31.81	.931	31.43	.925
WGWSNet (Zhu et al., 2023a)	CVPR'23	34.31	.946	30.16	.901	29.32	.921	32.38	.938	31.54	.926
WeatherDiff64 (Özdenizci & Legenstein, 2023)	TPAMI'23	35.83	.957	30.09	.904	29.64	.931	30.71	.931	31.57	.931
WeatherDiff128 (Özdenizci & Legenstein, 2023)	TPAMI'23	35.02	.952	29.58	.894	29.72	.922	29.66	.923	31.00	.923
AWRCP (Ye et al., 2023)	ICCV'23	36.92	.965	31.92	.934	31.39	.933	31.93	.931	33.04	.941
GridFormer (Wang et al., 2024)	IICV'24	37.46	.964	31.71	.923	31.87	.933	32.39	.936	33.36	.939
MPerceiver (Ai et al., 2024)	CVPR'24	36.23	.957	31.02	.916	31.25	.925	33.21	.929	32.93	.932
DTPM (Ye et al., 2024)	CVPR'24	37.01	.966	30.92	.917	30.99	.934	32.72	.944	32.91	.940
Histoformer (Sun et al., 2024)	ECCV'24	37.41	.966	32.16	.926	32.08	.939	33.06	.944	33.68	.944
ExDA (<i>Ours</i>)	2025	37.97	.968	32.49	.930	32.70	.945	32.66	.944	33.92	.947

368
 369
 370
 371
 372
 373
 374
 375
 376
 377
Composed Degradation. To better capture real-world complexities, we extend OneRestore (Guo et al., 2024b) to cover eleven scenarios, including rain, haze, snow, low-light conditions, and their composite variants. As summarized in Tab. 3, our approach consistently surpasses state-of-the-art methods such as AirNet (Li et al., 2022), PromptIR (Potlapalli et al., 2024), WGWSNet (Zhu et al., 2023a), WeatherDiff (Özdenizci & Legenstein, 2023), OneRestore (Guo et al., 2024b), and MoCE-IR (Zamfir et al., 2025). Notably, our method achieves a 0.92 dB PSNR improvement over MoCE-IR, underscoring its strength in addressing complex, mixed degradations.

Adverse Weather Removal. Following prior works (Valanarasu et al., 2022b; Zhu et al., 2023b), we evaluate ExDA on three challenging deweathering tasks: snow removal, rain streak and fog removal, and raindrop removal. Tab. 4 shows that ExDA consistently outperforms existing SOTA approaches across nearly all datasets, with the only exception being the PSNR result on RainDrop. These

Table 5: Comparisons for Real-World *WeatherBench* (Guan et al., 2025).

Method	Venue	Dehaze		Derain		Desnow		Average	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Restormer (Zamir et al., 2022b)	CVPR'22	19.30	.687	34.48	.945	27.95	.836	27.25	.823
AirNet (Li et al., 2022)	CVPR'22	20.94	.705	33.59	.942	22.05	.780	25.53	.809
TransWeather (Valanarasu et al., 2022a)	CVPR'23	19.79	.680	29.34	.903	24.96	.796	24.70	.793
PromptIR (Potlapalli et al., 2024)	NeurIPS'23	21.11	.713	34.53	.944	27.93	.836	27.86	.831
WGWS-Net (Zhu et al., 2023a)	CVPR'23	13.79	.603	37.08	.961	20.81	.909	23.89	.781
Histoformer (Sun et al., 2024)	ECCV'24	17.69	.669	30.70	.916	25.39	.808	24.59	.798
AdaIR (Cui et al., 2025)	ICLR'24	23.08	.731	34.87	.946	28.44	.837	28.80	.838
ExDA (Ours)	2025	23.74	.739	35.86	.946	29.42	.868	29.68	.851

Table 6: Comparisons for **3-task Medical Image Restoration**. Missing values are denoted by ‘-’.

Method	Venue	MRI-Super-Resolution		CT-Denoising		PET-Synthesis		Average	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
Restormer (Zamir et al., 2022a)	CVPR'22	31.72	.936	33.61	.918	37.13	.947	34.16	.934
AirNet (Li et al., 2022)	CVPR'22	31.39	.931	33.62	.917	37.17	.945	34.06	.931
Spach Transformer (Jang et al., 2023)	TMF'23	31.18	.929	33.47	.916	37.05	.945	33.90	.930
DRMC (Yang et al., 2023)	MICCAI'23	29.55	.903	33.28	.915	36.19	.945	33.00	.930
Eformer (Luthra et al.)	MedLA'23	29.19	.873	32.44	.908	35.11	.909	32.25	.897
AMIR (Yang et al., 2024)	MICCAI'24	31.86	.938	33.68	.918	37.22	.946	34.28	.934
ExDA (Ours)	2025	31.95	.938	33.67	.918	37.27	.947	34.30	.934

substantial improvements across diverse degradations highlight the robustness of ExDA. Especially, our method achieves a notable 0.24 dB PSNR gain over Histoformer (Sun et al., 2024).

Real-world WeatherBench. We further validate our approach on the challenging real-world WeatherBench benchmark (Guan et al., 2025), which involves diverse weather degradations such as haze, rain, and snow. As reported in Table 5, our method achieves the best overall performance, reaching an average PSNR/SSIM of 29.68/0.851. Notably, it delivers the highest scores on both dehazing and desnowing, and achieves competitive results in deraining compared with the strongest prior models. These consistent improvements across multiple degradations highlight the robustness and generalization ability of our method in handling complex, real-world weather conditions.

Medical All-in-One. Finally, we evaluate our method on medical image restoration tasks using the AMIR dataset (Yang et al., 2024), which unifies three restoration tasks within a single model. The results in Tab. 6 demonstrate the strong cross-domain applicability of our method, showing large gains in MRI super-resolution and competitive performance on the others.

Visual Results. Fig. 1 & 4 shows that our method restores high-quality and faithful results across diverse all-in-one settings, producing clearer boundaries and more accurate details than others.

4.2 ABLATION STUDY.

Component Analysis & Model Scaling Exploration. We start with two PromptIR-based baselines (with and without prompts). Unlike these settings, our framework does not depend on extra prompts for strong performance. Instead, the key improvements come from the proposed nonlinear value design and GST module. As shown in Fig. 5, adding a nonlinear value (“b + Nonlinear Value”) provides a clear gain, and further integrating GST (“b + GST”) yields consistent improvements. When combined (“c + d”), our full model achieves the best performance (32.96/0.921), confirming the effectiveness of these components. We then examine model scaling. Both the Small and Tiny versions (“e” and “f”) remain highly competitive despite their reduced size, with the Tiny variant (only 6M parameters) still approaching the full model’s performance. This demonstrates that once the core components are well designed, even extremely lightweight models can deliver strong results, making our framework both effective and scalable. The below ablations are all based on tiny model.

Nonlinear Value Design. To systematically investigate how nonlinear value transformations should be incorporated, we conduct ablations along three dimensions. *First*, we compare two formulations: an in-place nonlinear mapping ($V' = g_\theta(V)$) versus a residual form ($V' = V + g_\theta(V)$). Results consistently favor the residual design, which preserves the original representation while enriching it with nonlinear transformations. *Second*, we examine the choice of $g_\theta()$. While non-parametric nonlinearities (e.g., pure Sigmoid or GELU) provide limited gains, learnable parameterized mappings achieve substantially better performance, confirming the need for adaptive flexibility. *Third*, we analyze where nonlinear values matter most in all-in-one blind IR by applying them selectively to the encoder, the decoder, or both.

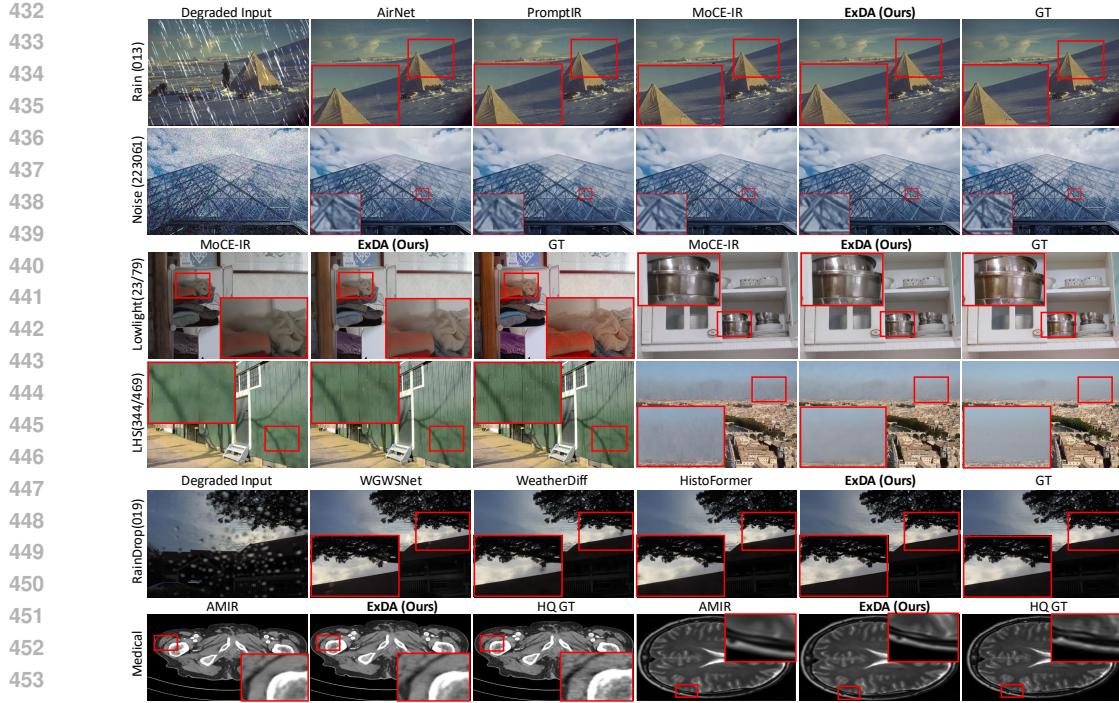


Figure 4: Visual results under diverse degradations and domains, spanning rain and noise (3-deg), low-light (5-deg), LHS composite (low-light+haze+snow), raindrop (adverse weather), and medical imaging. Please zoom in for details.

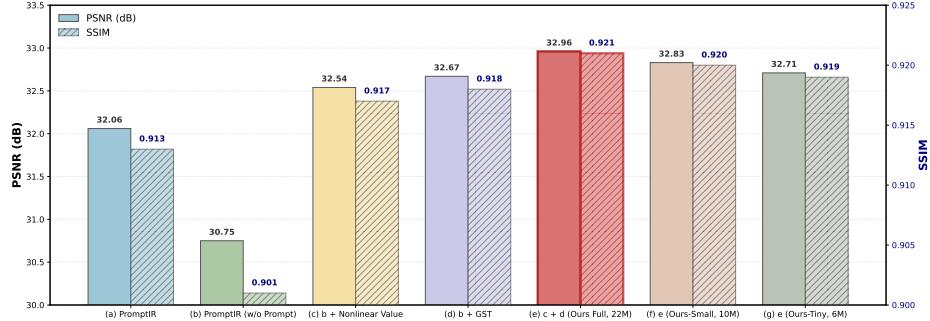


Figure 5: Qualitative results on medical restoration.

The evidence shows that deploying them in both the encoder and decoder yields the strongest improvements. Overall, these findings establish that residual, learnable, and fully integrated nonlinear value designs form the most effective strategy, significantly enhancing expressivity for diverse degradations in all-in-one IR (Tab. 7).

GST Configuration. We analyze the effect of the stride s in GST, which controls the granularity of spatial information compression. When $s = 1$, the model preserves excessive redundancy and achieves 32.45 dB PSNR. Larger strides, such as $s = 4$ and $s = 8$, oversimplify spatial structures and reduce performance to 32.52 and 32.40 dB, respectively. The best result is obtained with $s = 2$, reaching 32.71 dB PSNR, striking the optimal balance between information preservation and compactness. Therefore, we adopt $s = 2$ as the default configuration in our method.

Efficiency Analysis. Fig. 6 summarizes the engineering trade-offs of ExDA-Tiny/Small/Base. Latency grows smoothly with model size (209.6→385.5 ms), while restoration quality increases correspondingly (32.71→32.96 dB). ExDA-Small provides the best balance. To illustrate scalability

Table 7: *Ablations on Nonlinear Value Design*

Design Choice	Results	
	PSNR (dB, ↑)	SSIM (↑)
In-place Nonlinear (Sigmoid)	32.30	.913
Residual Nonlinear (Sigmoid)	32.63	.916
In-place Nonlinear (GELU)	32.45	.914
Residual Nonlinear (GELU)	32.71	.919
Non-parametric (GELU)	32.30	.914
Parametric (Learnable)	32.71	.919
Encoder-only	32.65	.916
Decoder-only	32.60	.915
Encoder + Decoder	32.71	.919

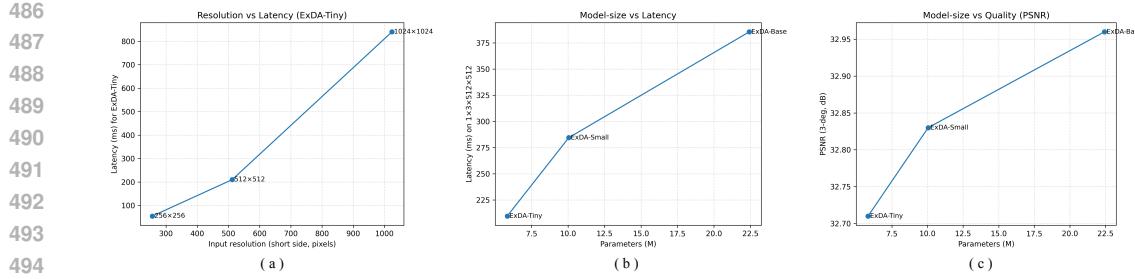


Figure 6: **Efficiency trade-offs of ExDA.** (a) Resolution–latency scaling (ExDA-Tiny). (b) Model-size–latency. (c) Model-size–quality (PSNR). ExDA-S provides the best accuracy–efficiency balance.

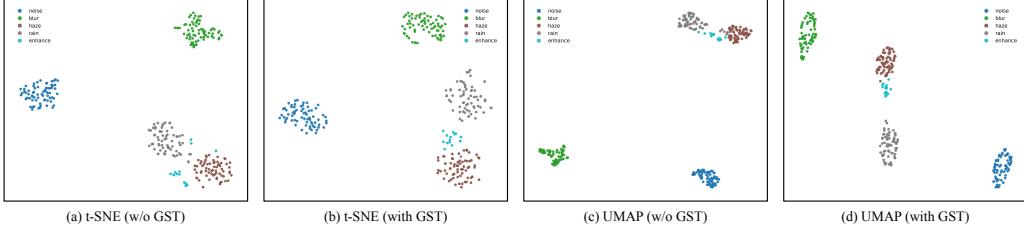


Figure 7: t-SNE and UMAP visualizations before (a,c) and after (b,d) applying GST.

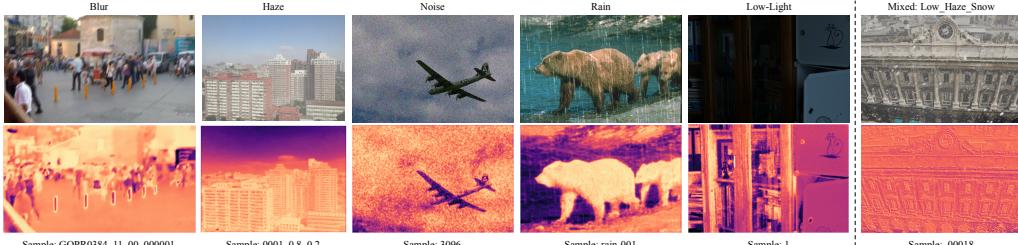


Figure 8: **GST-driven attention visualization across 5 degradations and 1 mixed case.** GST consistently highlights degradation-relevant regions, indicating that it captures meaningful degradation cues.

with input resolution, we sweep ExDA-Tiny from 256^2 to 1024^2 , observing near-linear $O(HW)$ growth ($54.5 \rightarrow 840.1$ ms). This confirms that ExDA scales predictably to high-resolution inputs.

Degradation-Aware Representation. As shown in Fig. 7, GST noticeably improves the organization of the embedding space: without GST, degradation types tend to overlap, while GST produces clearer and more compact clusters. To ensure a fair comparison given the uneven dataset sizes, we sampled 68 images for noise, blur, haze, and rain, and used all 15 samples for low-light. This balanced setting leads to substantial improvements in NMI ($0.71 \rightarrow 0.88$) and ARI ($0.56 \rightarrow 0.89$), reflecting stronger degradation-awareness. Fig. 8 further illustrates this effect by visualizing GST attention across blur, haze, noise, rain, low-light, and a mixed case (Low + Haze + Snow). Across all settings, GST consistently attends to degradation-relevant regions—such as blur boundaries, dense haze areas, noise-heavy textures, rain streaks, dark low-light structures, and mixed artifacts—showing that the module captures interpretable global cues that align with the improved embedding separation.

5 CONCLUSION

In this work, we revisited attention mechanisms for all-in-one image restoration and uncovered two central limitations in Restormer-style architectures: the linear constraint on value aggregation and the lack of explicit global context modeling. To address these issues, we introduced a nonlinear value transform and a global spatial token, two minimal yet powerful primitives that are both backbone-agnostic and lightweight. Our design consistently improves restoration quality across synthetic, real-world, underwater, and medical benchmarks, while remaining effective even in compact models with only a few million parameters. Beyond empirical gains, our analyses clarify why nonlinear value transformations enhance expressivity and why explicit global slots enable degradation-awareness. These insights highlight a broader principle: rethinking the role of values and global context is key to building robust, efficient, and general-purpose restoration models.

540
541 ETHICS STATEMENT

542 This work studies fundamental methods for all-in-one image restoration. Our research focuses on
 543 algorithmic improvements to attention mechanisms and does not involve the collection or annotation
 544 of new data. All datasets used in experiments (including synthetic, mixed, underwater, and medical
 545 benchmarks) are publicly available, widely used in the community, and employed strictly according to
 546 their licenses. The proposed methods are designed to improve robustness and generalization of image
 547 restoration models and do not raise foreseeable ethical concerns beyond the general considerations of
 548 responsible AI research.

549
550 REPRODUCIBILITY STATEMENT
551

552 We emphasize reproducibility throughout this work. Detailed descriptions of the architecture, training
 553 setups, and evaluation protocols are provided in the main text (see corresponding sections) and in
 554 the Appendix. If the paper is accepted, we will release all source code, pretrained weights, and
 555 configuration files immediately to facilitate transparent verification and future research.

556
557 REFERENCES
558

559 Yuang Ai, Huaibo Huang, Xiaoqiang Zhou, Jiexiang Wang, and Ran He. Multimodal prompt
 560 perceiver: Empower adaptiveness, generalizability and fidelity for all-in-one image restoration. In
 561 *CVPR*, pp. 25432–25444, 2024. 7

562 Yaroslav Aksenov, Nikita Balagansky, Sofia Maria Lo Cicero Vaina, Boris Shaposhnikov, Alexey
 563 Gorbatovski, and Daniil Gavrilov. Linear transformers with learnable kernel functions are better
 564 in-context models, 2024. URL <https://arxiv.org/abs/2402.10644>. 2

565 Pablo Arbelaez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and
 566 hierarchical image segmentation. *IEEE TPAMI*, 33(5):898–916, 2010. 1

567 Mark R Banham and Aggelos K Katsaggelos. Digital image restoration. *IEEE Signal Processing
 568 Magazine*, 14(2):24–41, 1997. 3

569 Hila Chefer, Shir Gur, and Lior Wolf. Transformer interpretability beyond attention visualization. In
 570 *CVPR*, pp. 782–791, 2021. 3

571 Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration.
 572 In *ECCV*, pp. 17–33, 2022a. 7

573 Wei-Ting Chen, Zhi-Kai Huang, Cheng-Che Tsai, Hao-Hsiang Yang, Jian-Jiun Ding, and Sy-Yen Kuo.
 574 Learning multiple adverse weather removal via two-stage knowledge learning and multi-contrastive
 575 regularization: Toward a unified model. In *CVPR*, pp. 17632–17641, 2022b. 7

576 Yinbo Chen, Sifei Liu, and Xiaolong Wang. Learning continuous image representation with local
 577 implicit image function. In *CVPR*, pp. 8628–8638, 2021. 3

578 Marcos V Conde, Gregor Geigle, and Radu Timofte. Instructir: High-quality image restoration
 579 following human instructions. In *ECCV*, 2024. 6, 7

580 Yuning Cui, Syed Waqas Zamir, Salman Khan, Alois Knoll, Mubarak Shah, and Fahad Shahbaz
 581 Khan. AdaIR: Adaptive all-in-one image restoration via frequency mining and modulation. In
 582 *ICLR*, 2025. 3, 6, 7, 8, 2

583 Tri Dao and Albert Gu. Transformers are SSMs: Generalized models and efficient algorithms through
 584 structured state space duality. In *ICML (ICML)*, 2024. 3

585 Chao Dong, Yubin Deng, Chen Change Loy, and Xiaoou Tang. Compression artifacts reduction by a
 586 deep convolutional network. In *ICCV*, pp. 576–584, 2015. 3

587 Yu Dong, Yihao Liu, He Zhang, Shifeng Chen, and Yu Qiao. Fd-gan: Generative adversarial networks
 588 with fusion-discriminator for single image dehazing. In *AAAI*, 2020. 6

594 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
 595 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image
 596 is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2020. 3
 597

598 Huiyu Duan, Xiongkuo Min, Sijing Wu, Wei Shen, and Guangtao Zhai. Uniprocessor: a text-induced
 599 unified low-level image processor. In *ECCV*, pp. 180–199. Springer, 2025. 6
 600

601 Akshay Dudhane, Omkar Thawakar, Syed Waqas Zamir, Salman Khan, Fahad Shahbaz Khan,
 602 and Ming-Hsuan Yang. Dynamic pre-training: Towards efficient and scalable all-in-one image
 603 restoration. *arXiv preprint arXiv:2404.02154*, 2024. 3
 604

605 Qingnan Fan, Dongdong Chen, Lu Yuan, Gang Hua, Nenghai Yu, and Baoquan Chen. A general
 606 decoupled learning framework for parameterized image operators. *TPAMI*, 43(1):33–47, 2019. 6, 7
 607

608 Hongyun Gao, Xin Tao, Xiaoyong Shen, and Jiaya Jia. Dynamic scene deblurring with parameter
 609 selective sharing and nested skip connections. In *CVPR*, 2019. 6
 610

611 Sicheng Gao, Xuhui Liu, Bohan Zeng, Sheng Xu, Yanjing Li, Xiaoyan Luo, Jianzhuang Liu, Xiantong
 612 Zhen, and Baochang Zhang. Implicit diffusion models for continuous super-resolution. In *CVPR*,
 613 pp. 10021–10030, 2023. 3
 614

615 Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv
 616 preprint arXiv:2312.00752*, 2023. 3
 617

618 Qiyuan Guan, Qianfeng Yang, Xiang Chen, Tianyu Song, Guiyue Jin, and Jiyu Jin. Weatherbench: A
 619 real-world benchmark dataset for all-in-one adverse weather image restoration. *ACM MM*, 2025.
 620 8, 1
 621

622 Hang Guo, Jinmin Li, Tao Dai, Zhihao Ouyang, Xudong Ren, and Shu-Tao Xia. Mambair: A simple
 623 baseline for image restoration with state-space model. In *ECCV*, 2024a. 3, 7
 624

625 Yu Guo, Yuan Gao, Yuxu Lu, Ryan Wen Liu, and Shengfeng He. Onerestore: A universal restoration
 626 framework for composite degradation. In *ECCV*, 2024b. 7, 1
 627

628 Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed
 629 self-exemplars. In *CVPR*, pp. 5197–5206, 2015. 1
 630

631 Se-In Jang, Tinsu Pan, Ye Li, Pedram Heidari, Junyu Chen, Quanzheng Li, and Kuang Gong. Spach
 632 transformer: Spatial and channel-wise transformer based on local and global self-attentions for pet
 633 image denoising, 2023. URL <https://arxiv.org/abs/2209.03300>. 8
 634

635 Kui Jiang, Zhongyuan Wang, Peng Yi, Chen Chen, Baojin Huang, Yimin Luo, Jiayi Ma, and Junjun
 636 Jiang. Multi-scale progressive fusion network for single image deraining. In *CVPR*, pp. 8346–8355,
 637 2020. 3
 638

639 Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are
 640 rmns: Fast autoregressive transformers with linear attention, 2020. URL <https://arxiv.org/abs/2006.16236>. 2
 641

642 Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015. 2
 643

644 Lingshun Kong, Jiangxin Dong, Jianjun Ge, Mingqiang Li, and Jinshan Pan. Efficient frequency
 645 domain-based transformers for high-quality image deblurring. In *CVPR*, pp. 5886–5895, 2023. 3
 646

647 Wei-Sheng Lai, Jia-Bin Huang, Narendra Ahuja, and Ming-Hsuan Yang. Deep laplacian pyramid
 648 networks for fast and accurate super-resolution. In *CVPR*, pp. 624–632, 2017. 3
 649

650 Boyi Li, Wenqi Ren, Dengpan Fu, Dacheng Tao, Dan Feng, Wenjun Zeng, and Zhangyang Wang.
 651 Benchmarking single-image dehazing and beyond. *TIP*, 28(1):492–505, 2018. 1
 652

653 Boyun Li, Xiao Liu, Peng Hu, Zhongqin Wu, Jiancheng Lv, and Xi Peng. All-in-one image restoration
 654 for unknown corruption. In *CVPR*, pp. 17452–17462, 2022. 3, 6, 7, 8, 1, 2
 655

656 Ruoteng Li, Loong-Fah Cheong, and Robby T Tan. Heavy rain image restoration: Integrating physics
 657 model and conditional adversarial learning. In *CVPR*, pp. 1633–1642, 2019. 1
 658

648 Ruoteng Li, Robby T Tan, and Loong-Fah Cheong. All in one bad weather removal using architectural
 649 search. In *CVPR*, pp. 3175–3185, 2020. 7
 650

651 Yawei Li, Yuchen Fan, Xiaoyu Xiang, Denis Demandolx, Rakesh Ranjan, Radu Timofte, and Luc
 652 Van Gool. Efficient and explicit modelling of image hierarchies for image restoration. In *CVPR*,
 653 pp. 18278–18289, 2023a. 3, 4

654 Yawei Li, Kai Zhang, Jingyun Liang, Jiezhang Cao, Ce Liu, Rui Gong, Yulun Zhang, Hao Tang, Yun
 655 Liu, Denis Demandolx, Rakesh Ranjan, Radu Timofte, and Luc Van Gool. LSDIR: A large scale
 656 dataset for image restoration. In *CVPRW*, pp. 1775–1787, 2023b. 3

657 Zilong Li, Yiming Lei, Chenglong Ma, Junping Zhang, and Hongming Shan. Prompt-in-prompt
 658 learning for universal image restoration. *arXiv preprint arXiv:2312.05038*, 2023c. 3

660 Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. SwinIR:
 661 Image restoration using Swin transformer. In *ICCVW*, pp. 1833–1844, 2021. 3, 4, 7

663 Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual
 664 networks for single image super-resolution. In *CVPRW*, pp. 1132–1140, 2017. 3

665 Chang Liu, Mengyi Zhao, Bin Ren, Mengyuan Liu, Nicu Sebe, et al. Spatio-temporal graph diffusion
 666 for text-driven human motion generation. In *BMVC*, pp. 722–729, 2023. 3

668 Lin Liu, Lingxi Xie, Xiaopeng Zhang, Shanxin Yuan, Xiangyu Chen, Wengang Zhou, Houqiang Li,
 669 and Qi Tian. Tape: Task-agnostic prior embedding for image restoration. In *ECCV*, 2022. 7

670 Yun-Fu Liu, Da-Wei Jaw, Shih-Chia Huang, and Jenq-Neng Hwang. Desnownet: Context-aware
 671 deep network for snow removal. *TIP*, 27(6):3064–3073, 2018. 1

673 Wenyang Luo, Haina Qin, Zewen Chen, Libin Wang, Dandan Zheng, Yuming Li, Yufan Liu, Bing
 674 Li, and Weiming Hu. Visual-instructed degradation diffusion for all-in-one image restoration. In
 675 *CVPR*, pp. 12764–12777, 2025. 3

676 Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Image restoration
 677 with mean-reverting stochastic differential equations. *arXiv preprint arXiv:2301.11699*, 2023. 3

679 Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Controlling
 680 vision-language models for universal image restoration. In *ICLR*, 2024. 6

681 A Luthra, H Sulakhe, T Mittal, A Iyer, and S Yadav. Eformer: Edge enhancement based transformer
 682 for medical image denoising. *arxiv* 2021. *arXiv preprint arXiv:2109.08044*. 8

684 Kede Ma, Zhengfang Duanmu, Qingbo Wu, Zhou Wang, Hongwei Yong, Hongliang Li, and Lei
 685 Zhang. Waterloo exploration database: New challenges for image quality assessment models.
 686 *IEEE TIP*, 26(2):1004–1016, 2016. 1

687 David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented
 688 natural images and its application to evaluating segmentation algorithms and measuring ecological
 689 statistics. In *ICCV*, pp. 416–423, 2001. 1

691 Chong Mou, Qian Wang, and Jian Zhang. Deep generalized unfolding networks for image restoration.
 692 In *CVPR*, pp. 17399–17410, 2022. 7

693 Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network
 694 for dynamic scene deblurring. In *CVPR*, pp. 3883–3891, 2017. 1

696 Ozan Özdenizci and Robert Legenstein. Restoring vision in adverse weather conditions with patch-
 697 based denoising diffusion models. *TPAMI*, 45(8):10346–10357, 2023. 7

698 Vaishnav Potlapalli, Syed Waqas Zamir, Salman H Khan, and Fahad Shahbaz Khan. Promptir:
 699 Prompting for all-in-one image restoration. *NeurIPS*, 36, 2024. 3, 6, 7, 8, 1, 2

701 Rui Qian, Robby T Tan, Wenhan Yang, Jiajun Su, and Jiaying Liu. Attentive generative adversarial
 702 network for raindrop removal from a single image. In *CVPR*, pp. 2482–2491, 2018. 1

702 Bin Ren, Yahui Liu, Yue Song, Wei Bi, Rita Cucchiara, Nicu Sebe, and Wei Wang. Masked jigsaw
 703 puzzle: A versatile position embedding for vision transformers. In *CVPR*, pp. 20382–20391, 2023a.
 704 3

705 Bin Ren, Yawei Li, Jingyun Liang, Rakesh Ranjan, Mengyuan Liu, Rita Cucchiara, Luc V Gool,
 706 Ming-Hsuan Yang, and Nicu Sebe. Sharing key semantics in transformer makes efficient image
 707 restoration. *NeurIPS*, 37:7427–7463, 2024. 4, 2

708 Dongwei Ren, Wangmeng Zuo, Qinghua Hu, Pengfei Zhu, and Deyu Meng. Progressive image
 709 deraining networks: A better and simpler baseline. In *CVPR*, pp. 3937–3946, 2019. 3

710 Mengwei Ren, Mauricio Delbracio, Hossein Talebi, Guido Gerig, and Peyman Milanfar. Multiscale
 711 structure guided diffusion for image deblurring. In *ICCV*, pp. 10721–10733, 2023b. 3

712 Wenqi Ren, Jinshan Pan, Hua Zhang, Xiaochun Cao, and Ming-Hsuan Yang. Single image dehazing
 713 via multi-scale convolutional neural networks with holistic edges. *IJCV*, 128:240–259, 2020. 3

714 William Hadley Richardson. Bayesian-based iterative method of image restoration. *Journal of the
 715 Optical Society of America*, 62(1):55–59, 1972. 3

716 Noam Shazeer. Glu variants improve transformer, 2020. URL <https://arxiv.org/abs/2002.05202>. 2

717 Zhuoran Shen, Mingyuan Zhang, Haiyu Zhao, Shuai Yi, and Hongsheng Li. Efficient attention:
 718 Attention with linear complexities, 2024. URL <https://arxiv.org/abs/1812.01243>.
 719 2

720 Shangquan Sun, Wenqi Ren, Xinwei Gao, Rui Wang, and Xiaochun Cao. Restoring images in adverse
 721 weather conditions via histogram transformer. In *ECCV*, volume 15080, pp. 111–129, 2024. 7, 8, 1

722 Xiaole Tang, Xiang Gu, Xiaoyi He, Xin Hu, and Jian Sun. Degradation-aware residual-conditioned
 723 optimal transport for unified image restoration. *TPAMI*, 2025. 3, 6, 7

724 Chunwei Tian, Yong Xu, and Wangmeng Zuo. Image denoising using deep cnn with batch renormal-
 725 ization. *Neural Networks*, 2020. 6

726 Xiangpeng Tian, Xiangyu Liao, Xiao Liu, Meng Li, and Chao Ren. Degradation-aware feature
 727 perturbation for all-in-one image restoration. In *CVPR*, pp. 28165–28175, 2025. 3, 6, 7

728 Zhengzhong Tu, Hossein Talebi, Han Zhang, Feng Yang, Peyman Milanfar, Alan Bovik, and Yinxiao
 729 Li. MAXIM: Multi-axis mlp for image processing. In *CVPR*, pp. 5769–5780, 2022. 3

730 Jeya Maria Jose Valanarasu, Rajeev Yasarla, and Vishal M. Patel. Transweather: Transformer-based
 731 restoration of images degraded by adverse weather conditions. In *CVPR*, pp. 2343–2353, 2022a. 7,
 732 8, 1

733 Jeya Maria Jose Valanarasu, Rajeev Yasarla, and Vishal M Patel. TransWeather: Transformer-based
 734 restoration of images degraded by adverse weather conditions. In *CVPR*, pp. 2353–2363, 2022b. 7

735 Cong Wang, Jinshan Pan, Wei Wang, Jiangxin Dong, Mengzhu Wang, Yakun Ju, and Junyang Chen.
 736 Promptrestorer: A prompting image restoration method with degradation perception. *NeurIPS*, 36:
 737 8898–8912, 2023a. 3

738 Tao Wang, Kaihao Zhang, Ziqian Shao, Wenhan Luo, Bjorn Stenger, Tong Lu, Tae-Kyun Kim, Wei
 739 Liu, and Hongdong Li. Gridformer: Residual dense transformer with grid structure for image
 740 restoration in adverse weather conditions. *IJCV*, 132(10):4541–4563, 2024. 7

741 Xintao Wang, Ke Yu, Chao Dong, and Chen Change Loy. Recovering realistic texture in image
 742 super-resolution by deep spatial feature transform. In *CVPR*, pp. 606–615, 2018. 3

743 Yinhuai Wang, Jiwen Yu, and Jian Zhang. Zero-shot image restoration using denoising diffusion
 744 null-space model. *ICLR*, 2023b. 3

745 Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light
 746 enhancement. *arXiv preprint arXiv:1808.04560*, 2018. 1

756 Gang Wu, Junjun Jiang, Kui Jiang, and Xianming Liu. Harmony in diversity: Improving all-in-one
 757 image restoration via multi-task collaboration. In *ACM MM*, pp. 6015–6023, 2024. 6, 7
 758

759 Haiyan Wu, Yanyun Qu, Shaohui Lin, Jian Zhou, Ruizhi Qiao, Zhizhong Zhang, Yuan Xie, and
 760 Lizhuang Ma. Contrastive learning for compact single image dehazing. In *CVPR*, pp. 10551–10560,
 761 2021. 3

762 Chengxing Xie, Xiaoming Zhang, Linze Li, Yuqian Fu, Biao Gong, Tianrui Li, and Kai Zhang. Mat:
 763 Multi-range attention transformer for efficient image super-resolution. *TCSVT*, 2025. 3
 764

765 Fuzhi Yang, Huan Yang, Jianlong Fu, Hongtao Lu, and Baining Guo. Learning texture transformer
 766 network for image super-resolution. In *CVPR*, pp. 5791–5800, 2020. 1

767 Zhiwen Yang, Yang Zhou, Hui Zhang, Bingzheng Wei, Yubo Fan, and Yan Xu. Drmc: A generalist
 768 model with dynamic routing for multi-center pet image synthesis. In *MICCAI*, pp. 36–46. Springer,
 769 2023. 8

770 Zhiwen Yang, Haowei Chen, Ziniu Qian, Yang Yi, Hui Zhang, Dan Zhao, Bingzheng Wei, and
 771 Yan Xu. All-in-one medical image restoration via task-adaptive routing. In *MICCAI*, pp. 67–77.
 772 Springer, 2024. 8, 1

773 Mingde Yao, Ruikang Xu, Yuanshen Guan, Jie Huang, and Zhiwei Xiong. Neural degradation
 774 representation learning for all-in-one image restoration. *TIP*, 2024. 6

775 Tian Ye, Sixiang Chen, Jinbin Bai, Jun Shi, Chenghao Xue, Jingxia Jiang, Junjie Yin, Erkang Chen,
 776 and Yun Liu. Adverse weather removal with codebook priors. In *ICCV*, pp. 12619–12630, 2023. 7

777 Tian Ye, Sixiang Chen, Wenhao Chai, Zhaohu Xing, Jing Qin, Ge Lin, and Lei Zhu. Learning
 778 diffusion texture priors for image restoration. In *CVPR*, pp. 2524–2534, 2024. 7

779 Zongsheng Yue, Jianyi Wang, and Chen Change Loy. ResShift: Efficient diffusion model for image
 780 super-resolution by residual shifting. *arXiv preprint arXiv:2307.12348*, 2023. 3

781 Eduard Zamfir, Zongwei Wu, Nancy Mehta, Yulun Zhang, and Radu Timofte. See more details:
 782 Efficient image super-resolution by experts mining. In *ICML*. PMLR, 2024. 3

783 Eduard Zamfir, Zongwei Wu, Nancy Mehta, Yuedong Tan, Danda Pani Paudel, Yulun Zhang, and
 784 Radu Timofte. Complexity experts are task-discriminative learners for any image restoration. In
 785 *CVPR*, 2025. 3, 6, 7, 1, 2

786 Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan
 787 Yang, and Ling Shao. Multi-stage progressive image restoration. In *CVPR*, pp. 14821–14831,
 788 2021. 6

789 Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-
 790 Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *CVPR*, pp.
 791 5728–5739, June 2022a. 8

792 Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-
 793 Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *CVPR*, pp.
 794 5728–5739, 2022b. 2, 3, 4, 7, 8

795 Haijin Zeng, Xiangming Wang, Yongyong Chen, Jingyong Su, and Jie Liu. Vision-language gradient
 796 descent-driven all-in-one deep unfolding networks. In *CVPR*, 2025. 3, 6, 7

797 Jinghao Zhang, Jie Huang, Mingde Yao, Zizheng Yang, Hu Yu, Man Zhou, and Feng Zhao. Ingredient-
 798 oriented multi-degradation learning for image restoration. In *CVPR*, pp. 5825–5835, 2023. 3, 6,
 799 7

800 Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser:
 801 Residual learning of deep cnn for image denoising. *IEEE TIP*, 26(7):3142–3155, 2017a. 3

802 Kai Zhang, Wangmeng Zuo, Shuhang Gu, and Lei Zhang. Learning deep cnn denoiser prior for
 803 image restoration. In *CVPR*, pp. 3929–3938, 2017b. 3

810 Leheng Zhang, Yawei Li, Xingyu Zhou, Xiaorui Zhao, and Shuhang Gu. Transcending the limit
 811 of local window: Advanced super-resolution transformer with adaptive token dictionary. *arXiv*
 812 *preprint arXiv:2401.08209*, 2024. [3](#)

813

814 Xu Zhang, Jiaqi Ma, Guoli Wang, Qian Zhang, Huan Zhang, and Lefei Zhang. Perceive-ir: Learning
 815 to perceive degradation better for all-in-one image restoration. *TIP*, 2025. [6](#), [7](#)

816 Yulun Zhang, Kunpeng Li, Kai Li, Bineng Zhong, and Yun Fu. Residual non-local attention networks
 817 for image restoration. *arXiv preprint arXiv:1903.10082*, 2019. [3](#)

818

819 Mengyi Zhao, Mengyuan Liu, Bin Ren, Shuling Dai, and Nicu Sebe. Denoising diffusion probabilistic
 820 models for action-conditioned 3d motion generation. In *ICASSP 2024-2024 IEEE International*
 821 *Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 4225–4229. IEEE, 2024. [3](#)

822 Xu Zheng, Yuanhuiyi Lyu, and Lin Wang. Learning modality-agnostic representation for semantic
 823 segmentation from any modalities. In *ECCV*, pp. 146–165. Springer, 2024. [3](#)

824

825 Lianghui Zhu, Bencheng Liao, Qian Zhang, Xinlong Wang, Wenyu Liu, and Xinggang Wang. Vision
 826 mamba: Efficient visual representation learning with bidirectional state space model. *arXiv preprint*
 827 *arXiv:2401.09417*, 2024. [3](#)

828 Yurui Zhu, Tianyu Wang, Xueyang Fu, Xuanyu Yang, Xin Guo, Jifeng Dai, Yu Qiao, and Xiaowei
 829 Hu. Learning weather-general and weather-specific features for image restoration under multiple
 830 adverse weather conditions. In *CVPR*, pp. 21747–21758, 2023a. [7](#), [8](#), [1](#)

831 Yurui Zhu, Tianyu Wang, Xueyang Fu, Xuanyu Yang, Xin Guo, Jifeng Dai, Yu Qiao, and Xiaowei
 832 Hu. Learning weather-general and weather-specific features for image restoration under multiple
 833 adverse weather conditions. In *CVPR*, pp. 21747–21758, 2023b. [7](#)

834

835

836

837

838

839

840

841

842

843

844

845

846

847

848

849

850

851

852

853

854

855

856

857

858

859

860

861

862

863

864 **A EXPERIMENTAL PROTOCOLS**
865
866
867868 **A.1 DATASETS**
869
870871 **3 Degradation Datasets.** For both the All-in-One and single-task settings, we follow the evaluation
872 protocols established in prior works (Li et al., 2022; Potlapalli et al., 2024; Zamfir et al., 2025),
873 utilizing the following datasets: For image denoising in the single-task setting, we combine the
874 BSD400 (Arbelaez et al., 2010) and WED (Ma et al., 2016) datasets, and corrupt the images with
875 Gaussian noise at levels $\sigma \in \{15, 25, 50\}$. BSD400 contains 400 training images, while WED
876 includes 4,744 images. We evaluate the denoising performance on BSD68 (Martin et al., 2001) and
877 Urban100 (Huang et al., 2015). For single-task deraining, we use Rain100L (Yang et al., 2020),
878 which provides 200 clean/rainy image pairs for training and 100 pairs for testing. For single-task
879 dehazing, we adopt the SOTS dataset (Li et al., 2018), consisting of 72,135 training images and 500
880 testing images. Under the All-in-One setting, we train a unified model on the combined set of the
881 aforementioned training datasets for 120 epochs and directly test it across all three restoration tasks.
882883 **5 Degradation Datasets.** The 5-degradation setting is built upon the 3-degradation setting, with
884 two additional tasks included: deblurring and low-light enhancement. For deblurring, we adopt the
885 GoPro dataset (Nah et al., 2017), which contains 2,103 training images and 1,111 testing images.
886 For low-light enhancement, we use the LOL-v1 dataset (Wei et al., 2018), consisting of 485 training
887 images and 15 testing images. Note that for the denoising task under the 5-degradation setting, we
888 report results using Gaussian noise with $\sigma = 25$. The training takes 130 epochs.
889890 **Composed Degradation Datasets.** Regarding the composite degradation setting, we use the
891 CDD11 dataset (Guo et al., 2024b). CDD11 consists of 1,183 training images for: (i) 4 kinds of
892 single-degradation types: haze (H), low-light (L), rain (R), and snow (S); (ii) 5 kinds of double-
893 degradation types: low-light + haze (l+h), low-light+rain (L+R), low-light + snow (L+S), haze +
894 rain (H+R), and haze + snow (H+S). (iii) 2 kinds of Triple-degradation type: low-light + haze + rain
895 (L+H+R), and low-light + haze + snow (L+H+S). We train our method for 150 epochs (fewer than
896 200 epochs than MoCE-IR (Zamfir et al., 2025)), and we keep all other settings unchanged.
897898 **Adverse Weather Removal Datasets.** For the deweathering tasks, we follow the experimental setups
899 used in TransWeather (Valanarasu et al., 2022a) and WGWSNet (Zhu et al., 2023a), evaluating the
900 performance of our approach on multiple synthetic datasets. We assess the capability of ExDA across
901 three challenging tasks: snow removal, rain streak and fog removal, and raindrop removal. The
902 training set, referred to as “AllWeather”, is composed of images from the Snow100K (Liu et al.,
903 2018), Raindrop (Qian et al., 2018), and Outdoor-Rain (Li et al., 2019) datasets. For testing, we
904 evaluate our model on the following subsets: Snow100K-S (16,611 images), Snow100K-L (16,801
905 images), Outdoor-Rain (750 images), and Raindrop (249 images). Same as Histoformer (Sun et al.,
906 2024), we train ExDA on “AllWeather” with 300,000 iterations.
907908 **Real-World WeatherBench.** For the real-world WeatherBench, we adopt newly released benchmark
909 from Guan et al. (2025), which contains in total 41,402 training pairs across 3 kinds of degradations
910 (*i.e.*, rain, haze, and snow). There are 600 pairs for testing and 200 pairs per degradation. Under the
911 All-in-One setting, we train a unified model on the combined set of the datasets for 120 epochs, and
912 directly test it across all three restoration tasks
913914 **Medical All-in-One Dataset.** AMIR dataset (Yang et al., 2024) include three important medical
915 image restoration tasks, include (i) MRI super-resolution dataset from public LXI MRI benchmark.
916 We use the public IXI MRI dataset, containing 578 high-quality T2-weighted MRI volumes. Low-
917 quality images are generated by 4x k-space downsampling (retaining 6.25% of central data). The
918 dataset is split 405/59/114 for training/validation/testing. (ii) CT Denoising dataset from NIH AAPM-
919 Mayo Clinic Low-Dose CT Grand Challenge, These images originate from 10 patients, with 8
920 allocated for training, 1 for validation, and 1 for testing purposes, after slicing we get 18531/128/211
921 for training/validation/testing. (iii) PET Synthesis include 159 HQ PET images acquired by (Yang
922 et al., 2024). The 3D volumes have dimensions of 192×192×400. Each volume is divided into 192
923 axial slices (192×400), excluding slices containing only air. Patient data is split into 120 training, 10
924 validation, and 29 testing cases.
925

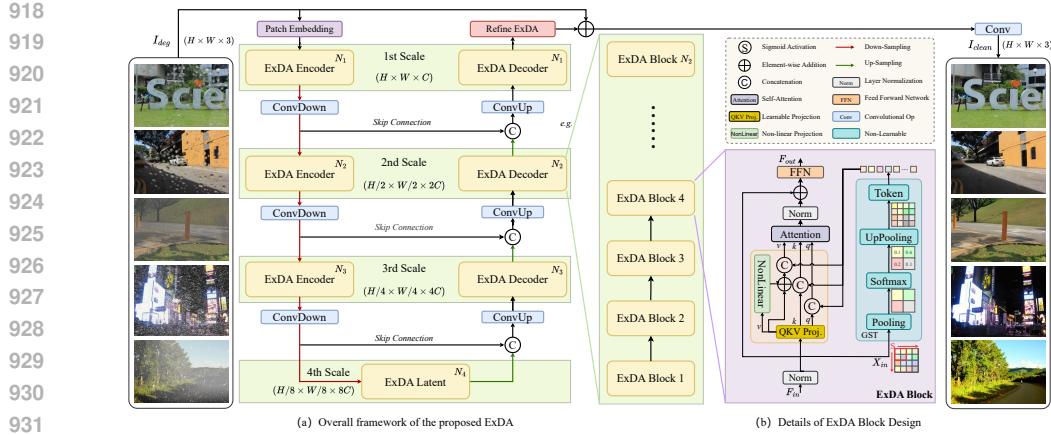


Figure A: ExDA Model Architecture Framework. (a) Demonstrates the overall pipeline for ExDA, (b) details the ExDA block design.

A.2 IMPLEMENTATION DETAILS

Implementation Details. Our ExDA framework is designed to be end-to-end trainable, removing the need for multi-stage optimization of individual components. The architecture adopts a robust 4-level encoder-decoder structure, with a varying number of Mixed Degradation Attention Blocks (MDAB) at each level—specifically [3, 5, 5, 7] from highest to lowest resolution in the Tiny variant. Following prior works (Potlapalli et al., 2024; Zamfir et al., 2025), we train the model for 120 epochs with a batch size of 32 in both the 3-Degradation All-in-One and single-task settings. The optimization uses a combination of L_1 and Fourier loss, optimized with Adam (Kingma & Ba, 2015) (initial learning rate of 2×10^{-4} , $\beta_1 = 0.9$, $\beta_2 = 0.999$) and a cosine decay schedule. During training, we apply random cropping to 128×128 patches, along with horizontal and vertical flipping as data augmentation. All experiments are conducted on a single NVIDIA H200 GPU (140 GB). Memory usage is approximately 42 GB for the Tiny (*i.e.*, ExDA -T) model and 56 GB for the Small model (*i.e.*, ExDA -S).

A.3 OPTIMIZATION OBJECTIVES

The overall optimization objective of our approach is defined as:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_1 + \lambda_{fre} \times \mathcal{L}_{\text{Fourier}} \quad (4)$$

Here, $\mathcal{L}_{\text{Fourier}}$ denotes the real-valued Fourier loss computed between the restored image and the ground-truth image, and \mathcal{L}_{SPD} represents our proposed contrastive learning objective in the SPD (Symmetric Positive Definite) space.

Specifically, we adopt an ℓ_1 loss that adopted in IR tasks (Potlapalli et al., 2024; Zamfir et al., 2025; Li et al., 2022; Cui et al., 2025; Ren et al., 2024), defined as $\mathcal{L}_1 = \|\hat{x} - x\|_1$, to enforce pixel-wise similarity between the restored image \hat{x} and the ground-truth image x . $\mathcal{L}_{\text{Fourier}}$, as utilized in MoCE-IR (Zamfir et al., 2025; Cui et al., 2025), to enhance frequency-domain consistency, the real-valued Fourier loss, is defined as:

$$\mathcal{L}_{\text{Fourier}} = \|\mathcal{F}_{\text{real}}(\hat{x}) - \mathcal{F}_{\text{real}}(x)\|_1 + \|\mathcal{F}_{\text{imag}}(\hat{x}) - \mathcal{F}_{\text{imag}}(x)\|_1, \quad (5)$$

where \hat{x} and x denote the restored and ground-truth images, respectively. $\mathcal{F}_{\text{real}}(\cdot)$ and $\mathcal{F}_{\text{imag}}(\cdot)$ represent the real and imaginary parts of the 2D real-input FFT (*i.e.*, rfft2). The final loss is computed as the ℓ_1 distance between the real and imaginary components of the predicted and target frequency spectra. Same as MoCE-IR (Zamfir et al., 2025), λ_{fre} is set to 0.1 throughout our experiments.

B MACRO ARCHITECTURE INTRODUCTION

The overall architecture of ExDA is illustrated in Fig. A. At a macro level, it adopts a U-shaped network with four hierarchical levels. Initially, a convolutional layer extracts shallow features from the

972 Table A: The details our the tiny and small version of our ExDA . FLOPs are computed on an image
 973 of size 224×224 using a NVIDIA Tesla A100 (40G) GPU.

	ExDA -Base	ExDA -S	ExDA -T
The Number of the MDAB crosses 4 scales	[4,6,6,8]	[4, 6, 6, 8]	[3, 5, 5, 7]
The Input Embedding Dimension	48	32	26
The FFN Expansion Factor	2	2	2
The Number of the Refinement Blocks	4	3	3
Params. (\downarrow)	22.42M	10.05M	5.83 M
FLOPs (\downarrow)	109G	47 G	27 G

974
 975 degraded input, creating a patch embedding of size $H \times W \times C$. As in standard U-Nets, each encoder
 976 stage doubles the embedding dimension and halves the spatial resolution, with skip connections
 977 transferring information to the corresponding decoder stage. In the decoder, features are merged with
 978 the previous decoding stage via linear projection. Finally, a global skip connection links input to
 979 output, preserving high-frequency details and producing the restored image.

980 Unlike most recent approaches that rely on additional prompts or large language models (LLMs), we
 981 argue that a simple encoder–decoder architecture, as illustrated in Fig. 1(I)(d), is sufficient to tackle
 982 the challenges of all-in-one image restoration. The key lies in identifying and addressing the latent
 983 bottlenecks.

984 **Model Scaling.** We propose two scaled variants of our ExDA , namely Tiny (ExDA -T) and Small
 985 (ExDA -S). As detailed in Tab. A, these variants differ in terms of the number of MDAB blocks across
 986 scales, the input embedding dimension, the FFN expansion factor, and the number of refinement
 987 blocks.

988 **Efficiency Comparison.** Tab. A presents a detailed comparison of PSNR, memory usage, parameter
 989 count, and FLOPs. Our Tiny model (ExDA -T) achieves the best efficiency-performance trade-off:
 990 with only 6.21M parameters and 16G FLOPs, it outperforms all prior methods, including larger
 991 models like PromptIR (Potlapalli et al., 2024) and MoCE-IR-S (Zamfir et al., 2025). Notably, ExDA
 992 -T surpasses MoCE-IR-S by +0.26 dB while requiring less than half the computational cost. Even
 993 our Small variant (ExDA -S) exceeds full MoCE-IR in both PSNR (+0.18 dB) and FLOPs (27G vs.
 994 75G). These results validate that our design achieves strong restoration quality without sacrificing
 995 computational efficiency.

1003 C ADDITIONAL EXPLANATION OF THE NONLINEAR VALUE PATH

1004 This section provides a more detailed explanation of the operator-level effect and motivation behind
 1005 introducing a nonlinear value transform into Restormer-style channel-wise attention.

1006 **Structural limitation of linear-value attention.** As discussed in Sec. 3.1 of the main paper, when
 1007 the value projection is strictly linear, the attention operator is fundamentally constrained to lie within
 1008 the linear span of the input features \mathcal{X} . For any attention head, we may write the attention operator
 1009 as:

$$1010 \text{Attn}(Q, K, V) = AV, \quad A = \text{Softmax}(QK^\top). \quad (6)$$

1011 Since V is a linear transformation of \mathcal{X} , its output is restricted by:

$$1012 \text{Range}(\text{Attn}) \subseteq \text{span}(\mathcal{X}), \quad (7)$$

1013 a limitation that cannot be removed by widening the FFN or adding post-aggregation nonlinearities,
 1014 as these operations act *after* the linear mixing in Eq. 6.

1015 **Effect of the residual nonlinear value transform.** To relax this structural constraint, the proposed
 1016 module applies a residual non-affine transform to the value branch:

$$1017 V' = V + g_\theta(V), \quad (8)$$

1018 where g_θ includes depth-wise convolutions and a GELU activation. Because g_θ is non-affine, we
 1019 have:

$$1020 g_\theta(\mathcal{X}) \not\subseteq \text{span}(\mathcal{X}), \quad (9)$$

1021 which enlarges the admissible value subspace:

$$1022 \text{span}(\mathcal{X}) \subsetneq \text{span}(\mathcal{X} \cup g_\theta(\mathcal{X})). \quad (10)$$

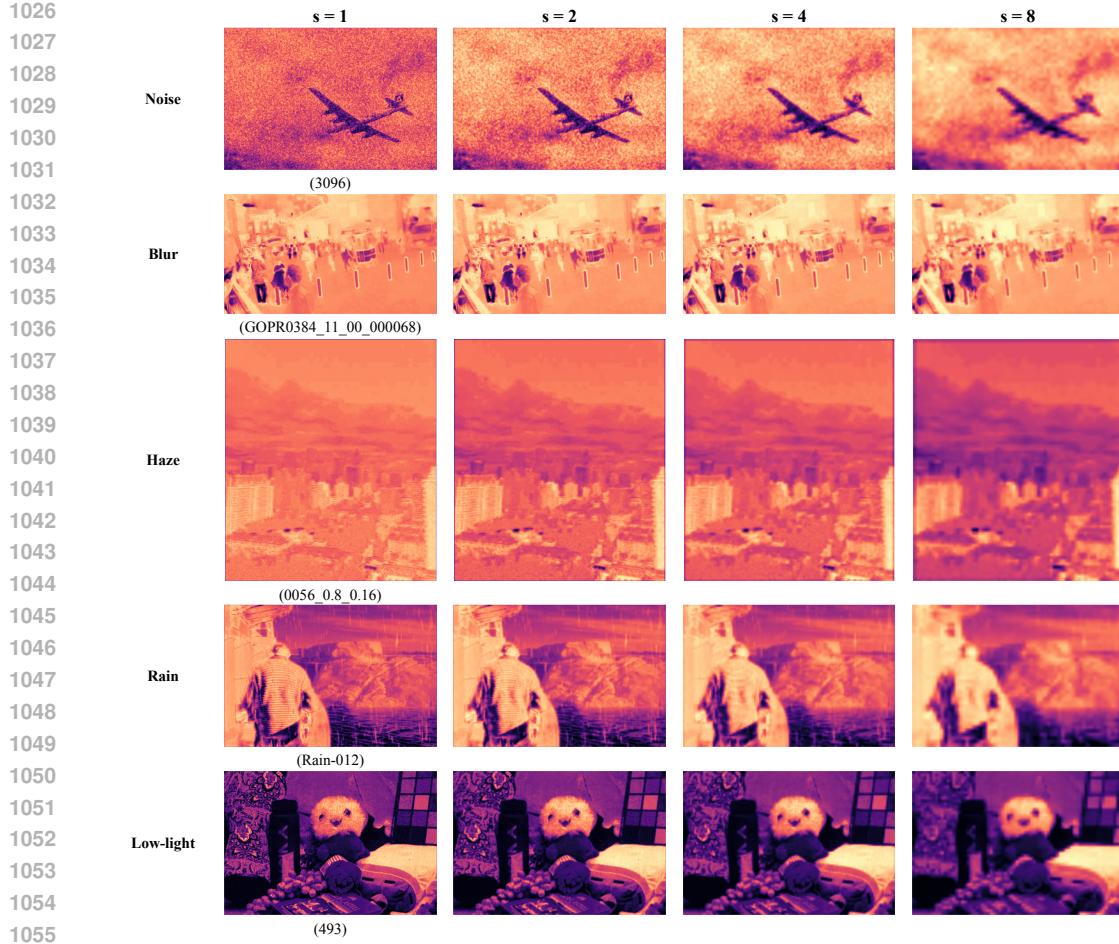


Figure B: Qualitative comparison of GST-driven attention maps across different stride values s . For five degradations (noise, blur, haze, rain, low-light), we visualize the GST heatmaps obtained with $s \in \{1, 2, 4, 8\}$. Smaller strides preserve finer spatial cues, while larger strides yield overly smoothed responses. Please zoom in for more details.

This is an operator-level property that directly increases the functional rank of the attention mechanism, independent of the global universal-approximation capacity of the entire network.

Empirical validation. Sec. 3.2 of the main paper provides evidence of the limitations of linear-value attention through synthetic function fitting and MNIST restoration experiments.

D ADDITIONAL ANALYSIS

To complement the quantitative stride ablation in the main paper, we further visualize in Fig. B how the GST-driven attention maps change with different stride values $s \in \{1, 2, 4, 8\}$ across the five degradations (noise, blur, haze, rain, low-light). A clear trend emerges: (i) $s=1$ preserves the most spatial detail but also introduces redundant high-frequency responses; (ii) $s=2$ maintains fine structures while producing clean and well-localized degradation cues; (iii) larger strides ($s=4$ and $s=8$) overly downsample the spatial statistics and yield noticeably blurred or weakened attention responses. These observations align with our quantitative findings, showing that $s=2$ provides the best balance between preserving degradation-relevant information and avoiding redundancy.

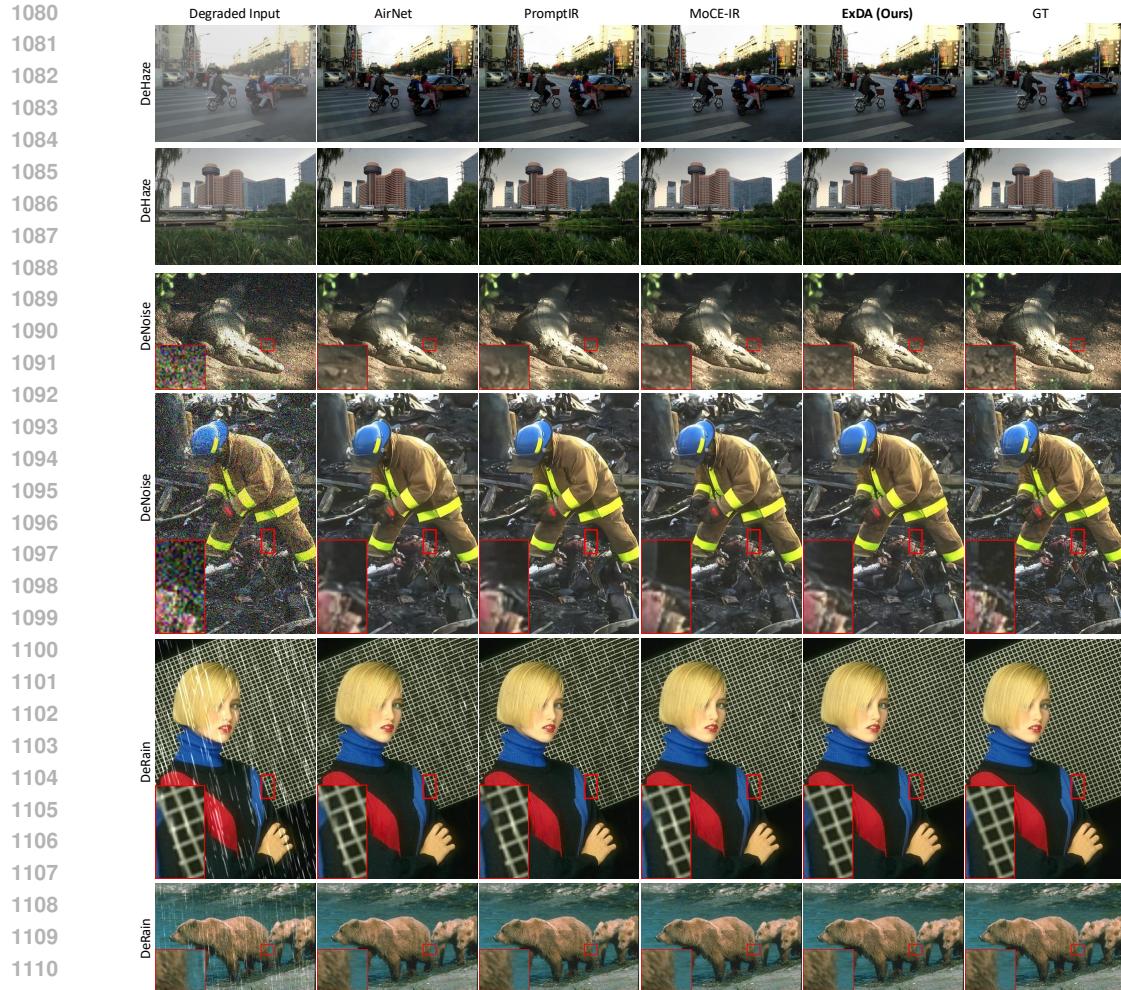


Figure C: More qualitative comparisons under 3-deg setting. Zoom in to see more details.

E ADDITIONAL VISUAL RESULTS

E.1 3-DEGRADATION

We additionally put more visual results on 3 Degradation settings, please refer to Fig. C for more detailed information.

E.2 5-DEGRADATION

We additionally put more visual results on 5 Degradation settings, please refer to Fig. D for more detailed information.

E.3 COMPOSITED DEGRADATION

We additionally put more visual results on Composed Degradation settings, please refer to Fig. E for more detailed information. Here, we selected the scenes composed by lowlight, haze and snow or lowlight, haze and rain, which are the most difficult settings. Our results indicated that our method can easily handle difficult and severe weather degradation.

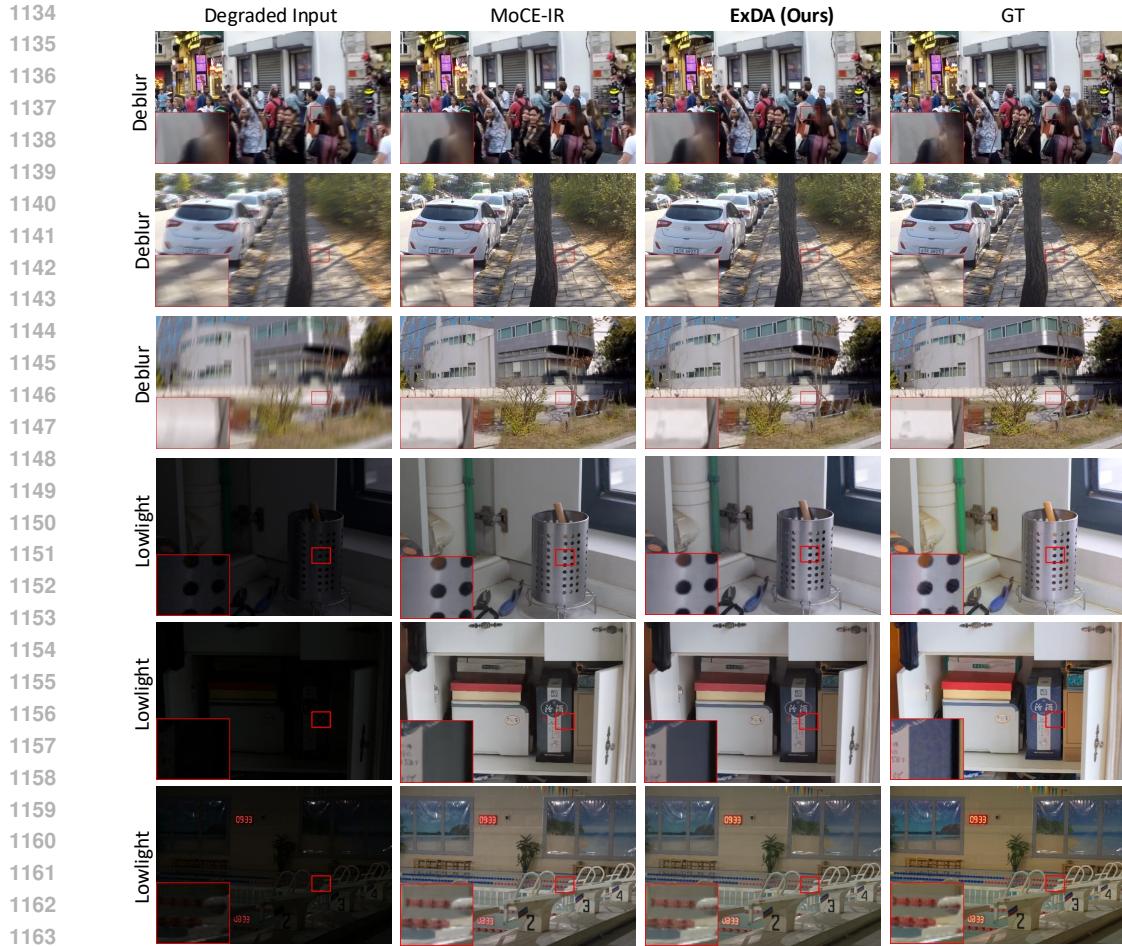


Figure D: More qualitative comparisons under 5-deg setting. Zoom in to see more details.

E.4 ADVERSE WEATHER REMOVAL

We additionally put more visual results on real-world WeatherBench data, please refer to Fig. F for more detailed information. Our results indicated that our method can handle different difficult and severe weather degradations.

E.5 REAL-WORLD WEATHERBENCH

We additionally put more visual results on real-world WeatherBench data, please refer to Fig. G for more detailed information. Our results indicated that our method can easily handle difficult and severe real-world weather degradation.

E.6 MEDICAL IMAGE RESTORATION

We provide additional visual results for medical image restoration. For better visualization, we display 64x64 crops. Both our method and the state-of-the-art (SOTA) baseline, AMIR, effectively enhance image quality. However, our approach achieves better detail, producing clearer tissue structures and sharper boundaries.

F USE OF LARGE LANGUAGE MODELS (LLMs).

Use of Large Language Models (LLMs). We used a large language model (ChatGPT) solely to aid in polishing the writing and improving the readability of the manuscript, such as refining grammar, style,



Figure E: More qualitative comparisons under composited degradation setting. Zoom in to see more details.

and clarity. The model was not involved in idea generation, experimental design, implementation, analysis, or any other research-related aspects of this work. All scientific contributions, including the problem formulation, methodology, experiments, and conclusions, are entirely the responsibility of the authors.

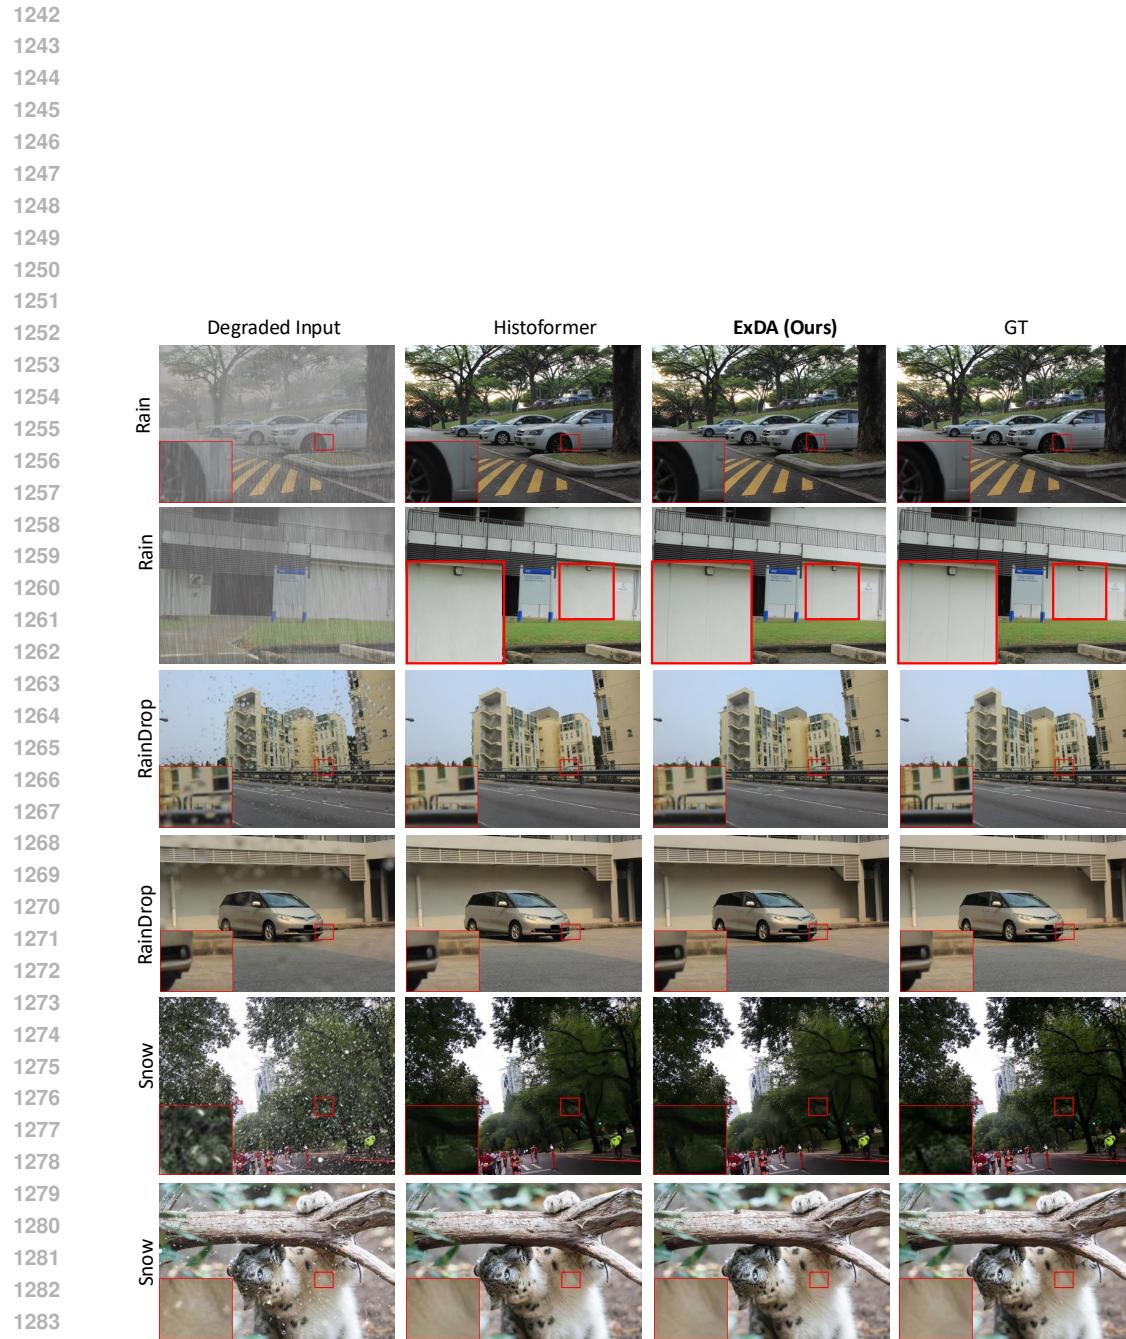


Figure F: More qualitative comparisons under adverse weather. Zoom in to see more details.

1287
1288
1289
1290
1291
1292
1293
1294
1295

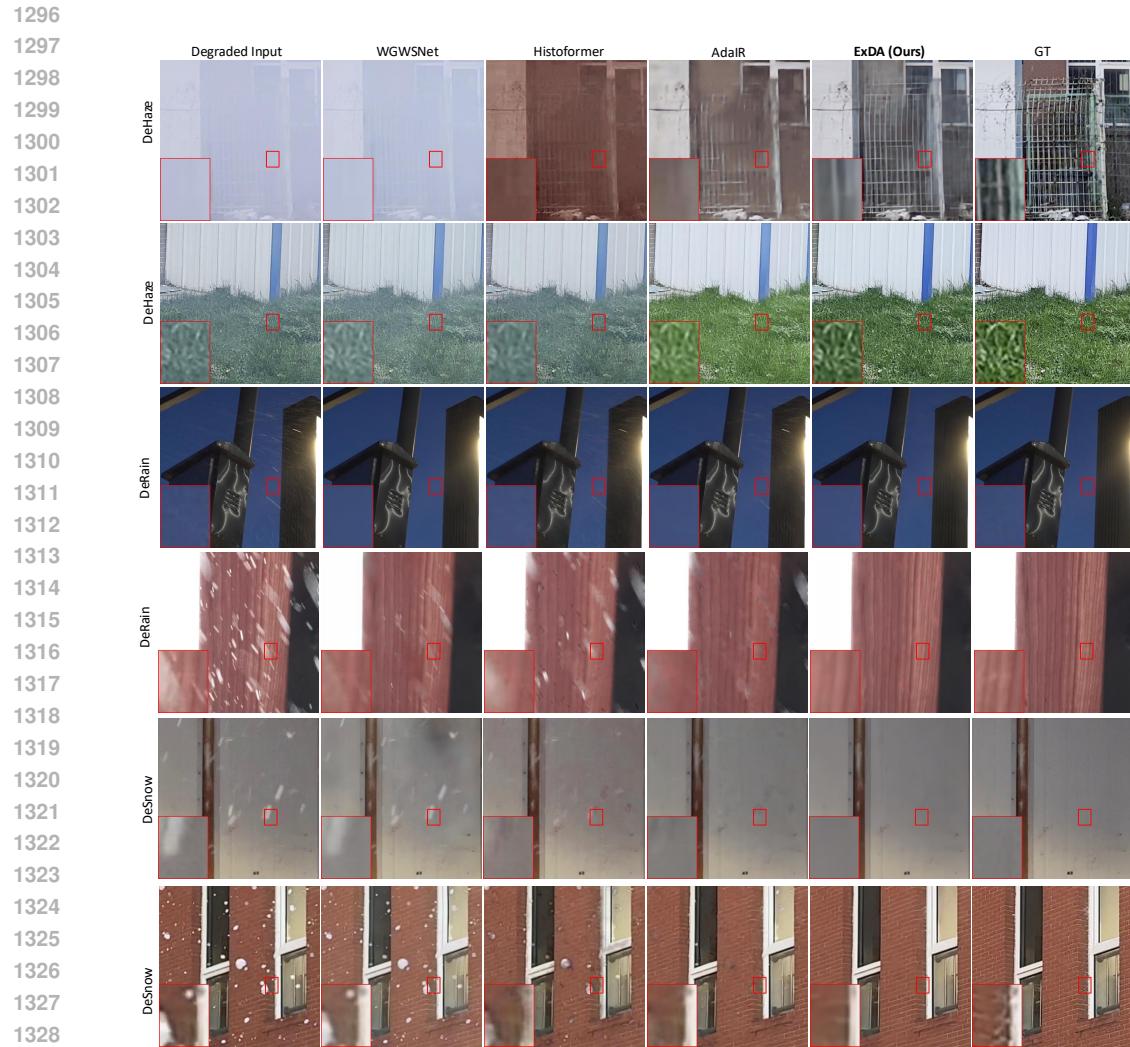


Figure G: More qualitative comparisons under real-world weatherbench data. Zoom in to see more details.

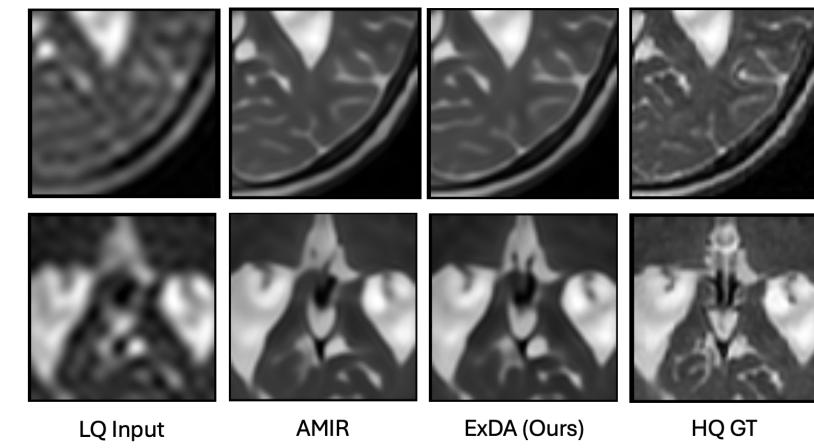


Figure H: More qualitative comparisons under medical restoration MRI super-resolution setting.

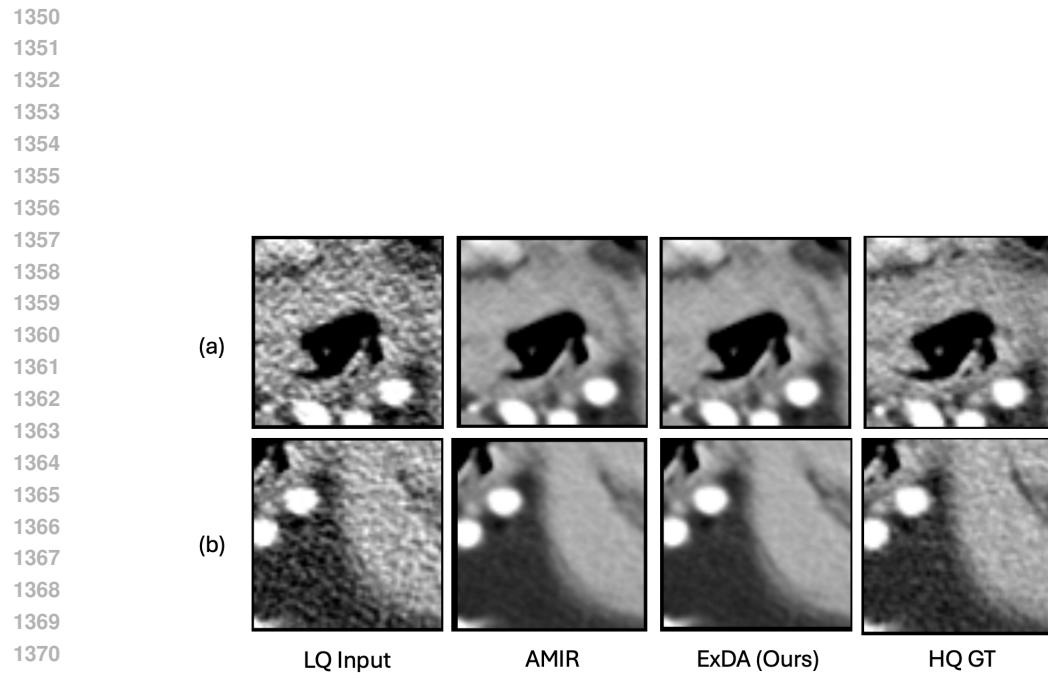


Figure I: More qualitative comparisons under medical restoration CT-Denoising setting.

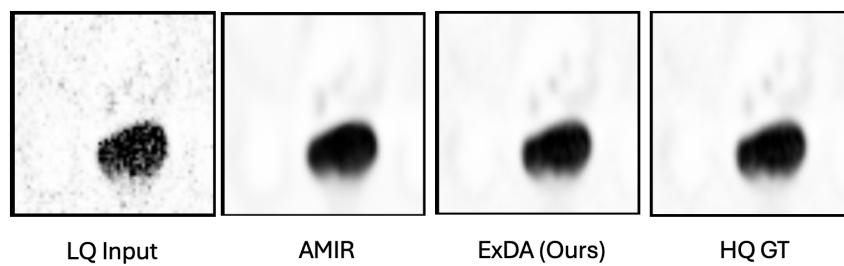


Figure I: More qualitative comparisons under medical restoration PET-Synthesis setting.