

000 LEARNING TO COMPOSE DEGRADATIONS: 001 002 DEGRADATION-ADAPTIVE CODEBOOK FOR ALL- 003 IN-ONE IMAGE RESTORATIO 004

006 **Anonymous authors**

007 Paper under double-blind review

011 ABSTRACT

013 All-in-one image restoration aims to develop a single model for diverse degra-
014 dations, a challenge whose success critically hinges on the precise representation
015 of the underlying degradation process. Existing methods simplify this challenge
016 by mapping each degradation to a coarse-grained, monolithic representation—
017 effectively treating them as discrete categories (e.g., “haze,” “noise”). This
018 paradigm, even in prompt-learning variants, fundamentally fails to capture the
019 continuous and fine-grained nature of real-world corruptions, such as varying in-
020 tensities, leading to suboptimal performance. To address this, we argue that degra-
021 dations are better represented as a composition of a finite set of learnable, elemen-
022 tary degradation primitives. We introduce DACode, a novel framework built upon
023 a global, learnable codebook embodying these primitives. The core of DACode
024 is a two-stage, dual cross-attention mechanism. First, in the Context-Aware Code
025 Adaptation stage, the codebook primitives act as queries to attend to the input
026 image features, generating a contextually-adapted codebook. Subsequently, in
027 the Code-based Feature Modulation stage, the image features query this adapted
028 codebook, aggregating relevant primitive information to perform targeted feature
029 restoration. This dynamic process allows DACode to construct highly specific
030 restorative features for each input. Notably, our analysis reveals that DACode
031 learns to activate distinct code combinations in response to both varying degra-
032 dation types (e.g., haze vs. rain) and severities (e.g., light vs. heavy haze), providing
033 direct evidence for its fine-grained modeling capability and interpretability. Ex-
034 tensive experiments show that DACode significantly outperforms state-of-the-art
035 methods across all-in-one restoration benchmarks. Code are availale in an anony-
036 mous repository <https://anonymous.4open.science/r/DAcode-847A/>

037 1 INTRODUCTION

038 Image restoration (Banham & Katsaggelos, 1997) is a fundamental low-level vision task that aims
039 to recover high-quality, clean images from observations degraded by a multitude of factors. While
040 deep learning models have yielded impressive results on specific restoration tasks such as denois-
041 ing (Zhang et al., 2017), deraining (Li et al., 2018b), and dehazing (Wu et al., 2021), these spe-
042 cialized networks embody a “one-model-per-task” paradigm. This approach is not only resource-
043 intensive but also lacks the flexibility to handle the diverse and often unpredictable degradations
044 encountered in real-world applications. Consequently, developing a single, unified model for All-
045 in-One image restoration has become a key and highly-watched research direction.

046 Significant progress (Chen et al., 2021) has been made in this area (Li et al., 2022; Park et al., 2023;
047 Guo et al., 2024). A common thread in existing methods, however, is the reliance on a **coarse-**
048 **grained, categorical representation** of degradation. They simplify the complex problem by map-
049 ping each degradation type to a monolithic, high-level label (e.g., “haze”), using mechanisms like
050 task-specific prompts (Potlapalli et al., 2023; Kong et al., 2024; Chen et al., 2024) or classifier
051 heads (Conde et al., 2024). This paradigm is fundamentally ill-equipped to capture the fine-grained,
052 continuous spectrum of real-world corruptions. For instance, it struggles to differentiate between a
053 light mist and a heavy fog (intensity differences), a limitation conceptually illustrated in Figure 1(a).
Moreover, attempting to address this continuous spectrum by discretizing it—for instance, requiring

054 a distinct prompt for each specific noise level (*e.g.*, $\sigma = 15, 25, 50$) as seen in methods like Prompt-
 055 tIR (Potlapalli et al., 2023)—inevitably leads to a combinatorial explosion. This inherent limitation
 056 in modeling rich details severely constrains the performance and generalization of existing models.
 057



070 Figure 1: Conceptual illustration of our proposed Compositional Paradigm versus the prior Categorical
 071 Paradigm. **(a)** Prior models based on the **Categorical Paradigm**, represented here by the
 072 state-of-the-art InstructIR (Conde et al., 2024). This coarse, one-size-fits-all representation leads to
 073 suboptimal performance, failing on the challenging case (PSNR: 27.3dB). **(b)** In contrast, our DA-
 074 Code operates under a **Compositional Paradigm**. It intelligently assembles different combinations
 075 of fine-grained “degradation primitives” to form bespoke, adaptive cues for each scenario, enabling
 076 consistently restoration quality.

077
 078 To address this fundamental representational bottleneck, we advocate for a paradigm shift towards
 079 **fine-grained, compositional modeling**. We posit that any complex degradation can be more effec-
 080 tively represented by adaptively combining a set of learnable, elemental “degradation primitives”.
 081 Building on this insight, we propose **DACode** Degradation-Adaptive Codebook, a novel framework
 082 that materializes this new paradigm, as depicted in Figure 1(b). The core of DACode is a learnable
 083 Codebook trained to capture these underlying primitives. Its key innovation lies in a two-stage, dual
 084 cross-attention mechanism that forms a bespoke degradation representation for each input. First, the
 085 codebook queries the image features to become context-aware. Subsequently, the image features
 086 query this adapted codebook to aggregate targeted restorative information. This design empowers
 087 our model to transcend discrete categorical constraints and naturally handle continuously varying
 088 and mixed degradations.

089 Our extensive experiments validate the superiority of this approach, with DACode establishing a
 090 new state-of-the-art across multiple challenging benchmarks. The main contributions of this work
 091 are summarized as follows:

- 092 • We identify the limitation of coarse-grained representation in existing methods and propose
 093 a new fine-grained, compositional modeling paradigm for all-in-one image restoration.
- 094 • We design DACode, a novel framework centered around a degradation-adaptive codebook
 095 and a dual cross-attention mechanism, which materializes the proposed paradigm.
- 096 • We achieve state-of-the-art performance and provide analyses that validate the model’s
 097 efficacy and its unique ability for interpretable, fine-grained degradation modeling.

100 2 RELATED WORK

101 2.1 SINGLE-TASK IMAGE RESTORATION

102 Deep learning-based image restoration has evolved rapidly, beginning with foundational CNN ar-
 103 chitectures (Dong et al., 2014; Zhang et al., 2017). The performance of these CNN-based mod-
 104 els was progressively advanced through the integration of more sophisticated mechanisms, such
 105 as channel and spatial attention (Zhang et al., 2018; Zamir et al., 2020), multi-stage refinement
 106 schemes (Zamir et al., 2021), and long-range feature interactions (Liu et al., 2018). The advent of
 107

108 the Vision Transformer (ViT) (Dosovitskiy et al., 2021) marked a significant paradigm shift. Re-
 109 searchers began to tailor Transformer architectures for image restoration to better model long-range
 110 dependencies (Yang et al., 2020; Chen et al., 2023b). However, the quadratic computational cost of
 111 the self-attention mechanism posed a major bottleneck for high-resolution images. This challenge
 112 spurred a wave of innovation focused on efficiency, leading to the development of window-based
 113 attention (Wang et al., 2022; Chen et al., 2022), linearized attention mechanisms (Deng et al., 2022),
 114 and sparse attention patterns (Chen et al., 2023a). This line of research culminated in powerful and
 115 versatile architectures like Restormer (Zamir et al., 2022) and Uformer (Wang et al., 2022), which
 116 can achieve state-of-the-art results on various individual restoration tasks. However, despite their
 117 architectural unity, these models must be independently trained and stored for each specific degra-
 118 dation. This “one-model-per-task” constraint severely limits their practicality in real-world scenarios
 119 that demand handling of multiple, unpredictable degradations, thereby motivating the need for the
 120 true All-in-One models discussed in the following section.

122 2.2 ALL-IN-ONE IMAGE RESTORATION

124 Early efforts in all-in-one image restoration, such as IPT (Chen et al., 2021), established the vi-
 125 ability of using a single, large Transformer-based backbone for multiple tasks. Following this, a
 126 dominant trend has emerged: developing explicit mechanisms to make models degradation-aware.
 127 This has been approached through various strategies, including learning degradation representations
 128 via contrastive learning (Li et al., 2022), employing explicit classifiers to guide task-specific fil-
 129 ters (Park et al., 2023; Hu et al., 2025), and learning distinct task-oriented centers (Zhang et al.,
 130 2023). More recently, prompt-based learning has become a popular paradigm. In this approach,
 131 a unique prompt is learned for each degradation type to enhance the feature representation of the
 132 restoration network, as exemplified by PromptIR (Potlapalli et al., 2023; Kong et al., 2024). This
 133 concept has been extended by leveraging textual information for more flexible control (Conde et al.,
 134 2024; Luo et al., 2024) or by optimizing the multi-task learning process to resolve conflicts between
 135 degradation tasks (Wu et al., 2024). A common denominator across these diverse approaches is their
 136 emphasis on distinguishing degradations at a category level. By focusing on high-level categorical
 137 distinctions, these methods largely neglect crucial details like varying degradation intensity or tex-
 138 ture. Our work, DACode, is fundamentally different. It directly addresses this gap by proposing a
 139 compositional approach that models degradations at a much finer granularity, assembling bespoke
 140 representations from a learned set of elementary primitives.

141 3 METHODOLOGY

144 Our overall network architecture is a hierarchical U-Net (Ronneberger et al., 2015), as illustrated in
 145 Figure 2. The encoder takes a degraded image $I_d \in \mathbb{R}^{H \times W \times 3}$ and passes it through four hierarchi-
 146 cal stages. The first three stages each contain a series of Transformer Blocks (TBs) followed by a
 147 downsampling layer, which halves the spatial resolution while doubling the feature channel dimen-
 148 sion to extract progressively abstract features. The symmetric decoder mirrors this structure with
 149 three corresponding upsampling stages, where features from the encoder are re-introduced via skip
 150 connections to preserve fine-grained details, ultimately restoring the clean image $I_c \in \mathbb{R}^{H \times W \times 3}$.
 151 In line with contemporary all-in-one restorers like AdaIR (Cui et al., 2025) and PromptIR (Potla-
 152 palli et al., 2023), the backbone is constructed with Transformer Blocks, whose design is based on
 153 the powerful Restormer (Zamir et al., 2022), comprising a Multi-Dconv Head Transposed Atten-
 154 tion (MDTA) module and a Gated-Dconv Feed-Forward Network (GDFN), which we refer to as the
 155 FFN. The detailed architectures of the MDTA and FFN modules are provided in the supplementary
 156 material.

157 While this powerful backbone excels at modeling spatial context, it lacks a specialized mechanism
 158 to adapt to the fine-grained characteristics of diverse degradations, a critical capability for the all-
 159 in-one challenge. To address this gap, our key innovation is the strategic insertion of our proposed
 160 **Degradation-Adaptive Codebook (DAC) Block**. Following the architectural patterns of recent
 161 leading methods (Potlapalli et al., 2023; Cui et al., 2025), we insert a DAC Block within each stage
 162 of the decoder. As detailed in Figure 2, each DAC Block is a complete processing unit, comprising
 163 our novel **DACode module**—which serves as a degradation-aware attention mechanism—followed

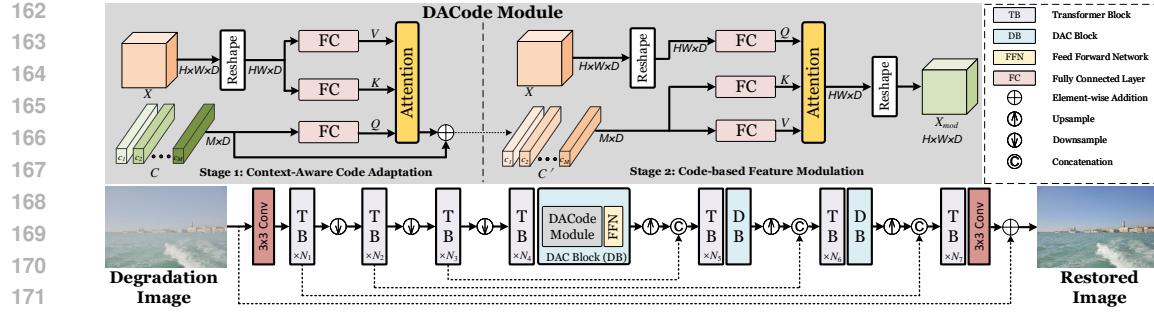


Figure 2: The overall architecture of our proposed network. Our key innovation, the Degradation-Adaptive Codebook (DAC) block, is placed in the decoder stages to guide restoration with fine-grained degradation knowledge.

by the same FFN used in the standard TBs. This design allows our model to inject degradation-specific knowledge at multiple scales during the reconstruction process.

3.1 THE DACODE MODULE

The DACode module, detailed in the top panel of Figure 2, is the technical core of our framework and the materialization of our proposed compositional paradigm. It is designed to dynamically generate a bespoke representation for any given degradation through a two-stage, dual cross-attention process: (a) Context-Aware Code Adaptation and (b) Code-based Feature Modulation. Let the input feature map to the module be $X \in \mathbb{R}^{H \times W \times D}$. For processing, we flatten its spatial dimensions to obtain $X_{\text{flat}} \in \mathbb{R}^{N \times D}$, where $N = H \times W$.

At the heart of our framework lies a globally shared and learnable codebook, denoted as $C \in \mathbb{R}^{N_c \times D}$. This codebook consists of N_c code vectors, which we conceptualize as “degradation primitives”, each with dimension D .

3.1.1 STAGE 1: CONTEXT-AWARE CODE ADAPTATION

For the universal primitives to be effective, they must first be tailored to the context of a specific degraded image. We achieve this by having the primitives *query* the image content via a cross-attention mechanism, where the codebook vectors C serve as queries (Q) and the input image features X_{flat} act as keys (K) and values (V). This allows the codebook to “read” the image and produce an update vector ΔC :

$$\Delta C = \text{Attention}(\text{LN}(C), X_{\text{flat}}, X_{\text{flat}}), \quad (1)$$

where $\text{LN}(\cdot)$ denotes Layer Normalization. The original codebook is then refined through a residual connection, controlled by a learnable, per-primitive scaling factor $\alpha_c \in \mathbb{R}^{N_c \times 1}$:

$$C' = C + \alpha_c \odot \Delta C. \quad (2)$$

Here, \odot denotes broadcasted element-wise multiplication. This process yields a contextually-adapted codebook $C' \in \mathbb{R}^{N_c \times D}$ that is now conditioned on the specific content and degradation style of the input.

3.1.2 STAGE 2: CODE-BASED FEATURE MODULATION

With the context-aware codebook C' obtained, the second stage performs feature modulation by reversing the roles in the attention mechanism. This time, the image features X_{flat} serve as queries (Q) to the adapted codebook C' , which acts as both keys (K) and values (V). This allows each image feature location to “look up” and aggregate the most relevant restorative information from the entire set of adapted primitives. The modulated feature X_{mod} is computed as:

$$X_{\text{mod}} = \text{Attention}(\text{LN}(X_{\text{flat}}), \text{LN}(C'), \text{LN}(C')). \quad (3)$$

Intuitively, this step materializes the compositional principle: the restoration of each pixel is guided by a bespoke representation, synthesized on-the-fly by combining the most relevant degradation

primitives from the adapted codebook. The final output X_{mod} is then passed to the subsequent FFN within the DAC Block.

4 EXPERIMENT

To rigorously evaluate our proposed DACode framework, we conduct a series of comprehensive experiments on the challenging task of all-in-one image restoration. In this section, we first detail our experimental setup, including the datasets and implementation specifics. Subsequently, we present the main quantitative and qualitative results, where we compare DACode against a range of state-of-the-art methods on both 3-task and 5-task benchmarks. Across all tasks, we quantify restoration quality using two standard metrics: the Peak Signal-to-Noise Ratio (PSNR) and the Structural Similarity Index (SSIM). For both metrics, higher scores signify superior restoration performance.

Table 1: Quantitative comparison (PSNR / SSIM) for all-in-one restoration on three tasks. The best results are in **bold**, and the second-best are underlined.

Method	Dehazing	Deraining	Denoising on BSD68			Average	Params
	SOTS	Rain100L	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$		
AirNet (Li et al., 2022)	27.94 / 0.962	34.90 / 0.967	33.92 / 0.933	31.26 / 0.888	28.00 / 0.797	31.20 / 0.910	9M
PromptIR (Potlapalli et al., 2023)	30.58 / 0.974	36.37 / 0.972	33.98 / 0.933	31.31 / 0.888	28.06 / 0.799	32.06 / 0.913	36M
Art-PromptIR (Wu et al., 2024)	30.83 / 0.979	37.94 / 0.982	34.06 / 0.934	31.42 / 0.891	28.14 / 0.801	32.49 / 0.917	33M
InstructIR (Conde et al., 2024)	30.22 / 0.959	37.98 / 0.978	34.15 / 0.933	<u>31.52</u> / 0.890	<u>28.30</u> / 0.804	32.43 / 0.913	16M
PromptIR-TUR (Wu et al., 2025)	31.17 / 0.978	38.57 / 0.984	34.06 / 0.932	31.40 / 0.887	28.13 / 0.797	32.67 / 0.916	33M
AdaIR (Cui et al., 2025)	31.06 / 0.980	38.64 / 0.983	34.12 / 0.935	31.46 / 0.892	28.19 / 0.802	32.69 / 0.918	29M
VLU-Net (Zeng et al., 2025)	30.71 / 0.980	<u>38.93</u> / 0.984	34.13 / 0.935	31.48 / <u>0.892</u>	28.23 / 0.804	32.70 / 0.919	35M
MoCE-IR (Zamfir et al., 2025)	<u>31.34</u> / 0.979	38.57 / 0.984	34.11 / 0.932	31.45 / 0.888	28.18 / 0.800	<u>32.73</u> / 0.917	25M
Ours (DACode)	31.50 / 0.982	39.10 / 0.985	34.24 / 0.937	31.60 / 0.895	28.34 / 0.809	32.96 / 0.922	29M

Table 2: Quantitative comparison (PSNR / SSIM) for all-in-one restoration on five tasks. Best results are in **bold**, second-best are underlined. Note that for denoising, we report results for $\sigma = 25$ following standard practice in this setting.

Method	Dehazing	Deraining	Denoising	Deblurring	Low-Light	Average	Params
	SOTS	Rain100L	$BSD68_{\sigma=25}$	GoPro	LOL		
AirNet (Li et al., 2022)	21.04 / 0.884	32.98 / 0.951	30.91 / 0.882	24.35 / 0.781	18.18 / 0.735	25.49 / 0.847	9M
PromptIR (Potlapalli et al., 2023)	25.20 / 0.931	35.94 / 0.964	31.17 / 0.882	27.32 / 0.842	20.94 / 0.799	28.11 / 0.883	33M
Gridformer (Wang et al., 2024)	26.79 / 0.951	36.61 / 0.971	<u>31.45</u> / 0.885	29.22 / 0.884	22.59 / 0.831	29.33 / 0.904	34M
InstructIR (Conde et al., 2024)	27.10 / 0.956	36.84 / 0.973	31.40 / 0.873	29.40 / <u>0.886</u>	23.00 / 0.836	29.55 / 0.908	17M
Transweather-TUR (Wu et al., 2025)	29.68 / 0.966	33.09 / 0.952	30.40 / 0.869	26.63 / 0.815	<u>23.02</u> / 0.838	28.56 / 0.888	38M
AdaIR (Cui et al., 2025)	30.53 / 0.978	38.02 / 0.981	31.35 / 0.889	28.12 / 0.858	23.00 / 0.845	30.20 / 0.910	29M
VLU-Net (Zeng et al., 2025)	<u>30.84</u> / 0.980	<u>38.54</u> / 0.982	31.43 / 0.891	27.46 / 0.840	22.29 / 0.833	30.11 / 0.905	35M
MoCE-IR (Zamfir et al., 2025)	30.48 / 0.974	38.04 / <u>0.982</u>	31.34 / 0.887	30.05 / 0.899	23.00 / <u>0.852</u>	<u>30.58</u> / 0.919	25M
Ours (DACode)	31.13 / 0.981	39.27 / 0.986	31.54 / 0.894	29.46 / 0.886	23.21 / 0.860	30.92 / 0.921	29M

4.1 EXPERIMENTAL SETTINGS

4.2 IMPLEMENTATION DETAILS

We propose two model variants built upon a 4-level encoder and 3-level decoder U-Net architecture. Our standard model, **DACode**, has a base channel dimension of 48 and uses [4, 6, 6, 8, 6, 6, 12] Transformer blocks across its seven stages. The lightweight **DACode-s** variant reduces the channel dimension to 32 with a shallower block configuration of [4, 6, 6, 8, 2, 4, 8]. The number of primitives in the DACode module is set to $N_c = 64$ for both variants. Models are trained a total of **150 epochs** on two NVIDIA L40 GPUs using the AdamW optimizer (Loshchilov & Hutter, 2017) to minimize the L1 loss. We use a batch size of 12. The learning rate is initialized to 2×10^{-4} and decayed to zero via a cosine annealing schedule. Input patches of size 128×128 are randomly cropped from training images.

Datasets and Benchmarks. To comprehensively evaluate DACode, our experiments are structured across three distinct benchmarks that progressively increase in complexity. First, we use a standard **3-task setting** to assess performance on common, isolated degradations. This includes: **(i)** **Deraining** on the Rain100L dataset Wenhan Yang & Yan (2017); **(ii)** **Dehazing** on the SOTS indoor

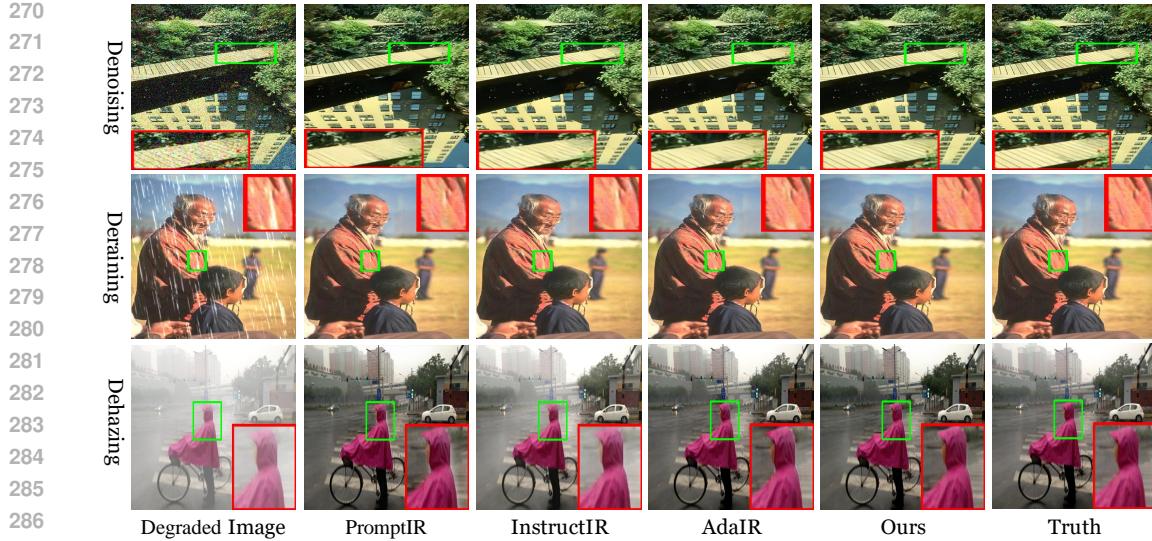


Figure 3: Qualitative comparison for all-in-one restoration on three tasks.

dataset Li et al. (2018a); and (iii) **Denoising**, for which we train on a composite of BSD400 Arbelaez et al. (2010) and WED Ma et al. (2016) with additive white Gaussian noise ($\sigma \in \{15, 25, 50\}$) and test on the BSD68 benchmark Martin et al. (2001). Second, we expand to a more extensive **5-task setting** to demonstrate the model’s versatility. This setup supplements the above tasks with (iv) **De-blurring** on the GoPro dataset Nah et al. (2017) and (v) **Low-Light Enhancement** on the LOL-v1 dataset Wei et al. (2018). Finally, to directly validate our model’s core capability in handling complex scenarios, we conduct experiments on the **CDD11** dataset Guo et al. (2024) with **Composited Degradations**. This benchmark is specifically designed to test performance on mixed corruptions, containing 11 degradation types that include both single modalities (e.g., rain, snow, haze, low-light) and their combinations (e.g., haze-rain, low-light-snow).

4.3 QUANTITATIVE AND QUALITATIVE COMPARISONS

To comprehensively evaluate our framework, we first benchmark DACode against state-of-the-art (SOTA) methods on two standard all-in-one settings: a foundational 3-task benchmark and a more extensive 5-task benchmark. The quantitative results are presented in Table 1 and Table 2. Our DACode framework demonstrates clear superiority, achieving the best overall performance in both settings. On the demanding 5-task benchmark, for instance, DACode surpasses the powerful MoCE-IR (Zamfir et al., 2025) baseline by a notable **0.34 dB** in average PSNR. These quantitative improvements are visually substantiated by our qualitative results in Figure 3. The comparisons reveal our model’s enhanced ability to restore fine-grained details and vibrant colors, such as the intricate skin texture in the deraining example and the faithful color rendition of the cyclist’s raincoat in the dehazing case.

Beyond the standard benchmarks of isolated degradations, we further probe our model’s capabilities in more realistic and complex scenarios using the **CDD11** benchmark for composite degradations. For a fair comparison against the predominantly lightweight methods evaluated on this dataset, we utilize our smaller variant, **DACode-S**. As shown in Table 3, our approach demonstrates overwhelming superiority. Notably, DACode-S achieves an average PSNR of 29.81 dB, surpassing the next-best method, Moce-IR-S, by a substantial margin of **0.76 dB**, despite having a comparable model size. This commanding performance is consistent across all 11 degradation types, spanning single, double, and even the most challenging triple-composite corruptions. This result provides strong, direct evidence for the effectiveness of our fine-grained, compositional paradigm in disentangling and restoring complex, real-world image degradations.

Collectively, these strong quantitative and qualitative results across all benchmarks validate the effectiveness and robustness of our proposed framework.

324
 325 **Table 3: Quantitative comparison on the CDD11 dataset. Best results are in **bold**, second-best are
 326 underlined. **(Top)** PSNR results. **(Bottom)** SSIM results.**

Method	Haze(H)	Low(L)	Rain(R)	Snow(S)	H+R	H+S	L+H	L+R	L+S	L+H+R	L+H+S	Average	Params.
AirNet (Li et al., 2022)	24.21	24.83	26.55	26.79	22.21	23.29	23.23	22.82	23.29	21.80	22.24	23.75	8.9M
PromptIR (Potlapalli et al., 2023)	<u>26.10</u>	26.32	31.56	31.53	24.54	23.70	24.49	25.05	24.51	24.49	23.33	25.97	38.5M
OneRestore (Guo et al., 2024)	32.52	26.48	33.40	34.31	29.99	30.21	25.79	25.58	25.19	24.78	24.90	28.47	6.0M
Moce-IR-S (Zamfir et al., 2025)	32.66	27.26	34.31	35.91	<u>29.93</u>	<u>30.19</u>	26.24	26.25	26.04	25.41	<u>25.39</u>	<u>29.05</u>	11.0M
DACode-S (Ours)	34.30	27.32	35.22	36.83	31.24	31.35	26.59	26.64	26.58	25.97	25.97	29.81	12.5M
Method	Haze(H)	Low(L)	Rain(R)	Snow(S)	H+R	H+S	L+H	L+R	L+S	L+H+R	L+H+S	Average	Params.
AirNet Li et al. (2022)	0.951	0.778	0.891	0.919	0.868	0.901	0.779	0.710	0.723	0.708	0.725	0.814	8.9M
PromptIR Potlapalli et al. (2023)	0.969	0.805	0.946	0.960	0.924	0.925	0.789	0.771	0.761	0.789	0.747	0.853	38.5M
OneRestore Guo et al. (2024)	0.990	0.826	0.964	0.973	0.957	0.964	0.822	0.799	0.789	0.788	0.791	0.878	6.0M
Moce-IR-S Zamfir et al. (2025)	0.990	0.824	0.970	0.980	0.964	0.970	0.817	0.800	0.793	0.789	0.790	0.881	11.0M
DACode-S (Ours)	0.991	0.834	0.974	0.981	0.968	0.972	0.831	0.814	0.808	0.806	0.805	0.890	12.5M

336 **Table 4: Ablation study on the number of code primitives
 337 (N_c). The case $N_c = 0$ represents our backbone without
 338 the DACode module. Best results are in **bold**.**

N_c	0	32	48	64 (Ours)	80	96	128
PSNR	31.98	32.51	32.72	32.96	32.74	32.93	32.96
SSIM	0.909	0.914	0.918	0.922	0.918	0.920	0.921

345 5 ABLATION STUDIES

346
 347 In this section, we conduct a series of targeted ablation studies to rigorously validate our proposed
 348 DACode framework. Our investigation is structured to first analyze the impact of the codebook size
 349 (N_c), which demonstrates the overall efficacy of our module and determines its optimal configura-
 350 tion. We then specifically investigate the criticality of the first stage in our dual-attention mechanism:
 351 the Context-Aware Code Adaptation. Finally, we provide in-depth qualitative visualizations that of-
 352 fer direct evidence of our model’s fine-grained and adaptive modeling capabilities. Unless otherwise
 353 specified, all ablation experiments are conducted on the **3-task** setting.

354 5.1 ANALYSIS ON THE NUMBER OF CODE PRIMITIVES

355
 356 The number of learnable primitives, N_c , in our codebook is a critical hyperparameter that directly
 357 influences the model’s representational capacity and parameter overhead. To quantify the overall
 358 effectiveness of our DACode module and find an optimal configuration, we evaluate the model’s
 359 performance while varying N_c in the set $\{0, 32, 48, 64, 80, 96, 128\}$. The $N_c = 0$ setting serves as a
 360 crucial baseline, as it effectively removes the DACode module.

361
 362 The quantitative results, summarized in Table 4, lead to two primary conclusions. First, the results
 363 reveal the substantial impact of our DACode module. Removing it entirely ($N_c = 0$) causes a
 364 drastic performance drop of 1.83 dB in average PSNR (from 32.96 dB to 31.13 dB). This clearly
 365 demonstrates that our proposed module is essential for high-quality restoration. Second, we observe
 366 a consistent trend of performance improvement as N_c increases from 32 to 64, with both PSNR and
 367 SSIM metrics reaching their peak at $N_c = 64$. Interestingly, further increasing the codebook size
 368 beyond 64 leads to performance saturation. While the PSNR at $N_c = 128$ matches our peak result,
 369 it offers no additional benefits, incurs a higher parameter cost, and results in a slightly lower SSIM
 370 score. This analysis indicates that $N_c = 64$ provides the optimal trade-off between representational
 371 capacity and model efficiency. Consequently, we adopt $N_c = 64$ as the default configuration for all
 372 our experiments.

373 5.2 EFFICACY OF CONTEXT-AWARE CODE ADAPTATION

374
 375 We hypothesize that a static codebook of universal primitives is suboptimal for the diverse na-
 376 ture of image degradations. To this end, our DACode module incorporates a crucial first stage:
 377 **Context-Aware Code Adaptation**, designed to dynamically refine these primitives based on the
 378 input image’s context. To validate the importance of this stage, we design an ablation variant named

336 **Table 5: Ablation on the Context-Aware
 337 Code Adaptation stage. Best results are
 338 in **bold**.**

Method	PSNR	SSIM
DACode w/o Adaptation	32.71	0.914
Full DACode (Ours)	32.96	0.922

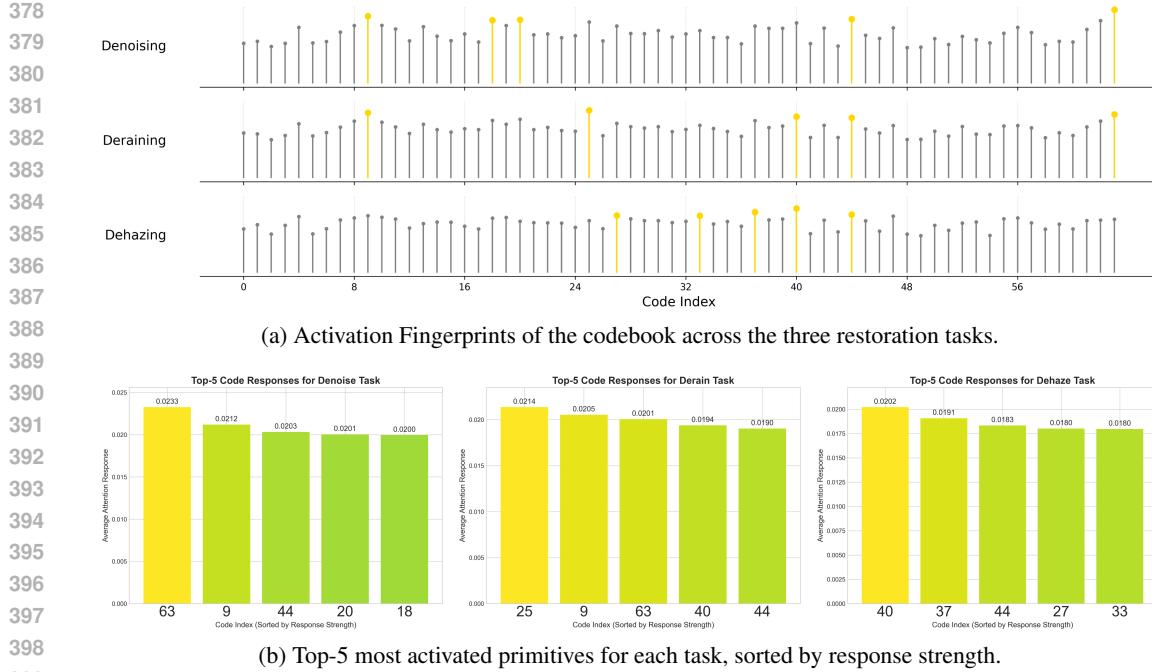


Figure 4: Functional specialization of primitives across different tasks. The visualizations in (a) the Activation Fingerprints and (b) the Top-5 primitives reveal a clear spectrum of specialization. We can identify three primary roles: (1) **Task-Specific Primitives**, which are dominantly activated for a single task (e.g., primitives #40 and #37 for Dehazing). (2) **Property-Specific Primitives**, which are co-activated for tasks with shared underlying properties, such as the additive, high-frequency nature of Denoising and Deraining (e.g., primitives #9 and #63). (3) **Universal Primitives**, such as #44, which are consistently activated across all three tasks, likely modeling fundamental restoration properties. This clear division of labor validates the richness and interpretability of our compositional approach.

“**DACode w/o Adaptation**”, where this mechanism is disabled. Specifically, we bypass the code update step in Eq. 2, forcing the model to use the same static, global codebook for all inputs.

The results, presented in Table 5, confirm our hypothesis. Disabling the adaptation mechanism leads to a substantial performance drop of **1.79 dB** in average PSNR and **0.013** in average SSIM. This significant gap demonstrates that while a global codebook can learn generic primitives, the ability to dynamically specialize them for each unique degradation instance is critical for achieving high-fidelity restoration. Without this context-aware adaptation, the model is constrained to a less effective, one-size-fits-all approach.

5.3 QUALITATIVE ANALYSIS: VERIFYING ADAPTIVE MODELING

Beyond quantitative metrics, it is crucial to qualitatively verify that our DACode module operates according to our core motivation. To this end, we visualize the codebook’s activation patterns in response to different degradation tasks on their respective standard benchmarks, as presented in Figure 4 and Figure 5. The complete response figures for different tasks are available in the supplementary material.

Functional Specialization of Primitives. Figure 4 illustrates how the primitives in our codebook have specialized for different roles when tested on distinct tasks. We observe a clear spectrum of specialization: First, some primitives are highly task-specific. For instance, primitives like #40 and #37 are dominantly activated only for Dehazing, demonstrating that the model learns representations dedicated to a single degradation family. Second, other primitives specialize in shared properties across tasks. A notable overlap exists between Denoising and Deraining, which share

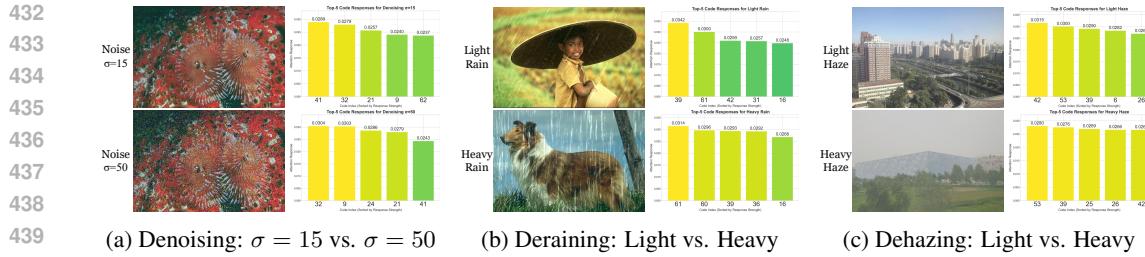


Figure 5: **Severity-aware activation of the DACode module.** The model assembles different combinations of Top-5 primitives in response to varying degradation intensity. This is demonstrated across three tasks: (a) **Denoising**, where the dominant primitive shifts from #41 ($\sigma = 15$) to #32 ($\sigma = 50$); (b) **Deraining**, with a shift from #39 (light) to #61 (heavy); and (c) **Dehazing**, shifting from #42 (light) to #53 (heavy). This behavior validates our model’s ability to perform fine-grained modeling by composing different “recipes” of primitives to match the input’s specific characteristics. Zoom in for best view.

key primitives such as **#9** and **#63**. We attribute this to the shared nature of their data synthesis as **additive degradations** that introduce high-frequency artifacts, suggesting these primitives function as general-purpose high-frequency artifact removers. Finally, we find evidence of universal primitives. **Code #44**, for example, appears as a Top-5 activated primitive across all three tasks. We hypothesize that such universal codes learn to model fundamental, task-agnostic properties essential for general image reconstruction, such as restoring fine textures. This clear division of labor—from universal, to property-specific, to task-specific functions—powerfully demonstrates the richness of the representations learned by our compositional paradigm.

Severity-Aware Activation. Furthermore, our model demonstrates a remarkable capability for severity-awareness. As illustrated in Figure 5, the model astutely adapts its response even within the same task. It composes different “recipes” of primitives, with visibly different activation strengths and participating codes, to precisely match the intensity of each input image. This, combined with the hierarchical activation, validates that DACode learns a rich, internal language of degradation rather than relying on monolithic labels.

6 CONCLUSION

In this paper, we identified a fundamental limitation in existing all-in-one image restoration methods: their reliance on coarse-grained, categorical representations of degradation. To address this, we proposed a paradigm shift towards fine-grained, compositional modeling. Our framework, DA-Code, materializes this new paradigm through a novel degradation-adaptive codebook and a dual cross-attention mechanism. The superiority of our compositional approach was validated through extensive experiments, where DACode established a new state-of-the-art on multiple challenging benchmarks. More than just performance, our in-depth analyses provided direct evidence for our core hypothesis. We revealed that the learned codebook exhibits a remarkable degree of functional specialization, with primitives ranging from universal to task-specific. This confirms that our model does not merely memorize tasks, but learns a rich, internal language of degradation, assembling different “recipes” of primitives to precisely match the type and severity of each unique corruption.

7 THE USE OF LARGE LANGUAGE MODELS (LLMs)

During the preparation of this paper, we employed a Large Language Model (LLM) to assist with improving the language and readability of the text. The primary use of the LLM was for proofreading, including correcting grammatical errors and refining sentence structure to enhance clarity. We confirm that the LLM was not used for research ideation, developing the methodology, conducting experiments, analyzing results, or drawing conclusions. All intellectual contributions and scientific claims are solely those of the authors.

486 REFERENCES
487

- 488 Pablo Arbelaez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and hi-
489 erarchical image segmentation. *IEEE transactions on pattern analysis and machine intelligence*,
490 33(5):898–916, 2010.
- 491 Mark R Banham and Aggelos K Katsaggelos. Digital image restoration. *IEEE signal processing
492 magazine*, 14(2):24–41, 1997.
- 493 Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chun-
494 jing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. In *Proceedings of the
495 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 12299–12310,
496 June 2021.
- 497 Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration.
498 In *European conference on computer vision*, pp. 17–33. Springer, 2022.
- 499 Sixiang Chen, Tian Ye, Kai Zhang, Zhaohu Xing, Yunlong Lin, and Lei Zhu. Teaching tailored to
500 talent: Adverse weather restoration via prompt pool and depth-anything constraint. In *European
501 Conference on Computer Vision*, pp. 95–115. Springer, 2024.
- 502 503 Xiang Chen, Hao Li, Mingqiang Li, and Jinshan Pan. Learning a sparse transformer network for
504 effective image deraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and
505 Pattern Recognition (CVPR)*, pp. 5896–5905, June 2023a.
- 506 507 Zheng Chen, Yulun Zhang, Jinjin Gu, Linghe Kong, Xiaokang Yang, and Fisher Yu. Dual aggre-
508 gation transformer for image super-resolution. In *Proceedings of the IEEE/CVF International
509 Conference on Computer Vision (ICCV)*, pp. 12312–12321, October 2023b.
- 510 Marcos V Conde, Gregor Geigle, and Radu Timofte. Instructir: High-quality image restoration
511 following human instructions. In *European Conference on Computer Vision*, pp. 1–21. Springer,
512 2024.
- 513 Yuning Cui, Syed Waqas Zamir, Salman Khan, Alois Knoll, Mubarak Shah, and Fahad Shahbaz
514 Khan. AdaIR: Adaptive all-in-one image restoration via frequency mining and modulation. In
515 *The Thirteenth International Conference on Learning Representations*, 2025.
- 516 517 Ye Deng, Siqi Hui, Sanping Zhou, Deyu Meng, and Jinjun Wang. T-former: An efficient transformer
518 for image inpainting. In *Proceedings of the 30th ACM international conference on multimedia*,
519 pp. 6559–6568, 2022.
- 520 521 Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional net-
522 work for image super-resolution. In *Computer Vision–ECCV 2014: 13th European Conference,
523 Zurich, Switzerland, September 6–12, 2014, Proceedings, Part IV 13*, pp. 184–199. Springer,
524 2014.
- 525 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas
526 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit,
527 and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at
528 scale. In *International Conference on Learning Representations*, 2021.
- 529 Yu Guo, Yuan Gao, Yuxu Lu, Huilin Zhu, Ryan Wen Liu, and Shengfeng He. Onerestore: A
530 universal restoration framework for composite degradation. In *European Conference on Computer
531 Vision*, pp. 255–272. Springer, 2024.
- 532 533 JiaKui Hu, Lujia Jin, Zhengjian Yao, and Yanye Lu. Universal image restoration pre-training via
534 degradation classification. In *International Conference on Learning Representations*, 2025.
- 535 536 Xiangtao Kong, Chao Dong, and Lei Zhang. Towards effective multiple-in-one image restoration:
537 A sequential and prompt learning strategy. *arXiv preprint arXiv:2401.03379*, 2024.
- 538 539 Boyi Li, Wenqi Ren, Dengpan Fu, Dacheng Tao, Dan Feng, Wenjun Zeng, and Zhangyang Wang.
Benchmarking single-image dehazing and beyond. *IEEE Transactions on Image Processing*, 28
(1):492–505, 2018a.

- 540 Boyun Li, Xiao Liu, Peng Hu, Zhongqin Wu, Jiancheng Lv, and Xi Peng. All-in-one image restora-
 541 tion for unknown corruption. In *Proceedings of the IEEE/CVF conference on computer vision*
 542 and *pattern recognition*, pp. 17452–17462, 2022.
- 543
- 544 Xia Li, Jianlong Wu, Zhouchen Lin, Hong Liu, and Hongbin Zha. Recurrent squeeze-and-excitation
 545 context aggregation net for single image deraining. In *European conference on computer vision*
 546 (*ECCV*), pp. 254–269, 2018b.
- 547 Ding Liu, Bihang Wen, Yuchen Fan, Chen Change Loy, and Thomas S Huang. Non-local recurrent
 548 network for image restoration. volume 31, 2018.
- 549
- 550 Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint*
 551 *arXiv:1711.05101*, 2017.
- 552
- 553 Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Controlling
 554 vision-language models for multi-task image restoration. In *The Twelfth International Conference*
 555 *on Learning Representations*, 2024.
- 556
- 557 Kede Ma, Zhengfang Duanmu, Qingbo Wu, Zhou Wang, Hongwei Yong, Hongliang Li, and Lei
 558 Zhang. Waterloo exploration database: New challenges for image quality assessment models.
 559 *IEEE Transactions on Image Processing*, 26(2):1004–1016, 2016.
- 560
- 561 David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented
 562 natural images and its application to evaluating segmentation algorithms and measuring ecological
 563 statistics. In *Proceedings of the IEEE international conference on computer vision*, 2001.
- 564
- 565 Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network
 566 for dynamic scene deblurring. In *IEEE Conference on Computer Vision and Pattern Recognition*,
 567 2017.
- 568
- 569 Dongwon Park, Byung Hyun Lee, and Se Young Chun. All-in-one image restoration for unknown
 570 degradations using adaptive discriminative filters for specific degradations. In *2023 IEEE/CVF*
 571 *Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 5815–5824. IEEE, 2023.
- 572
- 573 Vaishnav Potlapalli, Syed Waqas Zamir, Salman H Khan, and Fahad Shahbaz Khan. Promptir:
 574 Prompting for all-in-one image restoration. volume 36, pp. 71275–71293, 2023.
- 575
- 576 Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical
 577 image segmentation. In *Medical image computing and computer-assisted intervention-MICCAI 2015: 18th international conference, Munich, Germany, October 5-9, 2015, proceedings, part III 18*, pp. 234–241. Springer, 2015.
- 578
- 579 Tao Wang, Kaihao Zhang, Ziqian Shao, Wenhan Luo, Bjorn Stenger, Tong Lu, Tae-Kyun Kim, Wei
 580 Liu, and Hongdong Li. Gridformer: Residual dense transformer with grid structure for image
 581 restoration in adverse weather conditions, 2024.
- 582
- 583 Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li.
 584 Uformer: A general u-shaped transformer for image restoration. In *Proceedings of the IEEE/CVF*
 585 *conference on computer vision and pattern recognition*, pp. 17683–17693, 2022.
- 586
- 587 Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light
 588 enhancement. *arXiv preprint arXiv:1808.04560*, 2018.
- 589
- 590 Jiashi Feng Jiaying Liu Zongming Guo Wenhan Yang, Robby T. Tan and Shuicheng Yan. Joint rain
 591 detection and removal from a single image. 2017.
- 592
- 593 Gang Wu, Junjun Jiang, Kui Jiang, and Xianming Liu. Harmony in diversity: Improving all-in-one
 594 image restoration via multi-task collaboration. In *Proceedings of the 32nd ACM International*
 595 *Conference on Multimedia*, pp. 6015–6023, 2024.
- 596
- 597 Gang Wu, Junjun Jiang, Yijun Wang, Kui Jiang, and Xianming Liu. Debiased all-in-one image
 598 restoration with task uncertainty regularization. In *Proceedings of the AAAI Conference on Artificial*
 599 *Intelligence*, volume 39, pp. 8386–8394, 2025.

- 594 Haiyan Wu, Yanyun Qu, Shaohui Lin, Jian Zhou, Ruizhi Qiao, Zhizhong Zhang, Yuan Xie, and
 595 Lizhuang Ma. Contrastive learning for compact single image dehazing. In *Proceedings of the*
 596 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 10551–10560, 2021.
 597
- 598 Fuzhi Yang, Huan Yang, Jianlong Fu, Hongtao Lu, and Baining Guo. Learning texture transformer
 599 network for image super-resolution. In *Proceedings of the IEEE/CVF conference on computer*
 600 *vision and pattern recognition*, pp. 5791–5800, 2020.
- 601 Eduard Zamfir, Zongwei Wu, Nancy Mehta, Yuedong Tan, Danda Pani Paudel, Yulun Zhang, and
 602 Radu Timofte. Complexity experts are task-discriminative learners for any image restoration.
 603 In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 12753–12763,
 604 2025.
- 605 Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-
 606 Hsuan Yang, and Ling Shao. Learning enriched features for real image restoration and enhance-
 607 ment. In *European Conference on Computer Vision*, pp. 492–511. Springer, 2020.
- 608 Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-
 609 Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *Proceedings of the*
 610 *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 14821–14831,
 611 June 2021.
- 612 Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-
 613 Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceed-
 614 ings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5728–5739,
 615 2022.
- 616 Haijin Zeng, Xiangming Wang, Yongyong Chen, Jingyong Su, and Jie Liu. Vision-language gradient
 617 descent-driven all-in-one deep unfolding networks. In *Proceedings of the Computer Vision and*
 618 *Pattern Recognition Conference*, pp. 7524–7533, 2025.
- 619 Jinghao Zhang, Jie Huang, Mingde Yao, Zizheng Yang, Hu Yu, Man Zhou, and Feng Zhao.
 620 Ingredient-oriented multi-degradation learning for image restoration. In *Proceedings of the*
 621 *IEEE/CVF conference on computer vision and pattern recognition*, pp. 5825–5835, 2023.
- 622 Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser:
 623 Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26
 624 (7):3142–3155, 2017.
- 625 Yulun Zhang, Kunpeng Li, Kai Li, Lichen Wang, Bineng Zhong, and Yun Fu. Image super-
 626 resolution using very deep residual channel attention networks. In *Proceedings of the European*
 627 *conference on computer vision (ECCV)*, pp. 286–301, 2018.

632 A APPENDIX

633 A.1 INTRODUCTION

637 This supplementary document provides additional details, experiments, and visualizations to com-
 638 plement our main paper. The contents are organized as follows:

- 640 • **Section A.2:** A detailed architectural breakdown of the standard Transformer Block used
 641 in our backbone, including the specific structures of the Multi-Dconv Head Transposed
 642 Attention (MDTA) and the Feed-Forward Network (FFN).
- 643 • **Section A.3:** Complete, unabridged visualizations of the codebook activation responses
 644 (i.e., the full 64-primitive bar charts), providing a comprehensive view of the results sum-
 645 marized in the main paper.
- 646 • **Section A.4:** Additional qualitative results, providing more visual comparisons of our
 647 method against state-of-the-art approaches on various restoration tasks.

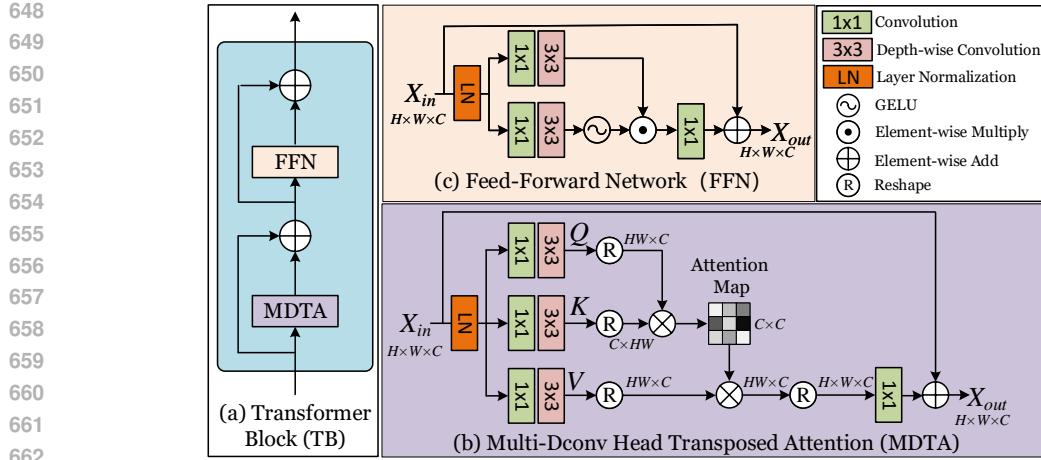


Figure 6: (a) Detailed architecture of the core components within a standard Transformer Block (TB). (b) The Multi-Dconv Head Transposed Attention (MDTA) module. Unlike standard self-attention which computes spatial relationships, MDTA computes attention across feature channels, making it more efficient for high-resolution images. (c) The Gated-Dconv Feed-Forward Network (GDFN), which we refer to as the FFN in our main paper. It employs a gating mechanism to control feature flow and enhance representational power.

A.2 DETAILED NETWORK ARCHITECTURES

In the main paper, we mentioned that the Transformer Blocks (TBs) in our backbone are based on the Restormer Zamir et al. (2022) design. Each TB is composed of two core components: a Multi-Dconv Head Transposed Attention (MDTA) module for feature aggregation, and a Gated-Dconv Feed-Forward Network (GDFN) for feature transformation, which we refer to as the FFN. Here, we provide a detailed architectural breakdown of these two modules, as illustrated in Figure 6.

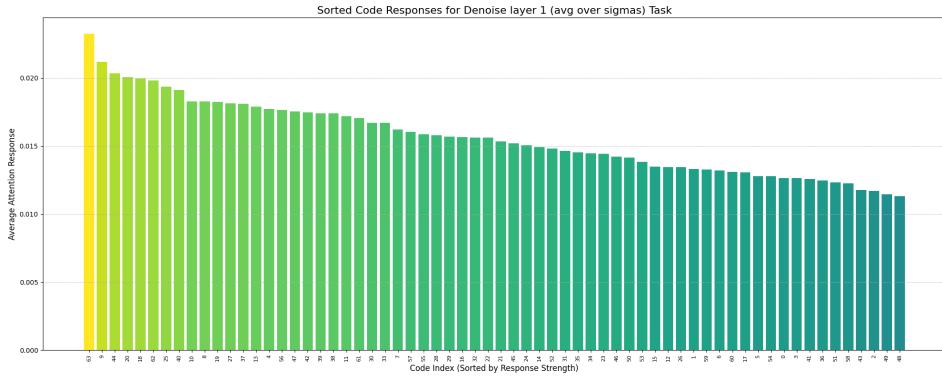
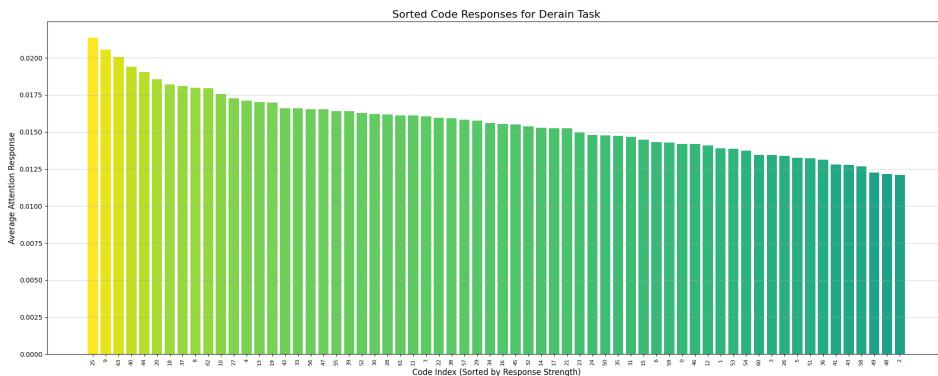
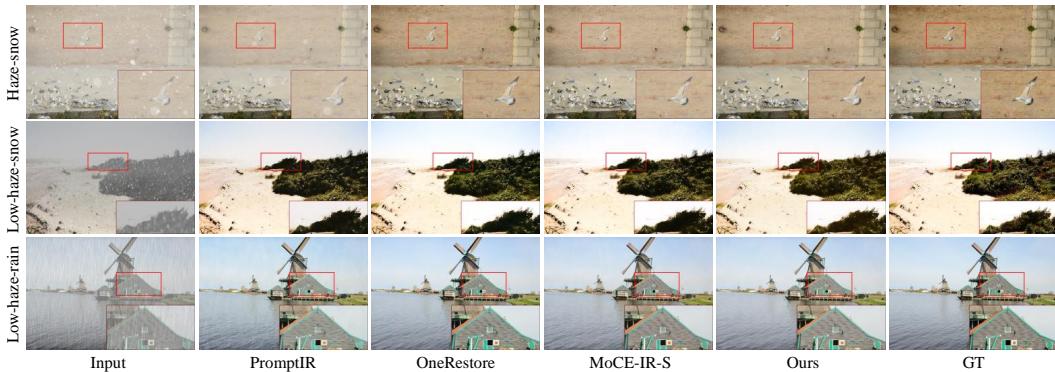


Figure 7: The complete distribution of the 64 code primitives’ responses for the **Denoising** task, averaged over all noise levels ($\sigma \in \{15, 25, 50\}$) on the BSD68 dataset.

Multi-Dconv Head Transposed Attention (MDTA) The MDTA module, shown in Figure 6(b), is designed to efficiently model long-range dependencies while maintaining a low computational cost. Given an input feature map $X_{in} \in \mathbb{R}^{H \times W \times C}$, it first passes through a Layer Normalization (LN). Then, to generate the query (Q), key (K), and value (V) projections, the normalized features are processed by three parallel branches. Each branch consists of a 1×1 convolution to adjust the channel dimension, followed by a 3×3 depth-wise convolution to aggregate local spatial context.

The key innovation of MDTA is its use of transposed attention. Instead of computing attention across the spatial dimension ($N \times N$, where $N = H \times W$), it computes attention across the channel dimension. To achieve this, the key projection K is transposed and multiplied with the query projection Q .





756
 757
 758
 759
 760
 761
 762
 763
 764
 765
 766
 767
 768
 769 **Figure 10: Qualitative comparison on challenging composite degradations from the CDD11**
 770 **dataset.** The examples shown are (from top to bottom): haze-snow, low-light-haze-snow, and low-
 771 light-haze-rain.

772
 773 clarity, these full-distribution plots allow for a detailed inspection of the entire 64-primitive code-
 774 book’s behavior. The results shown here correspond to the activations from the DACode module at
 775 Layer 1.

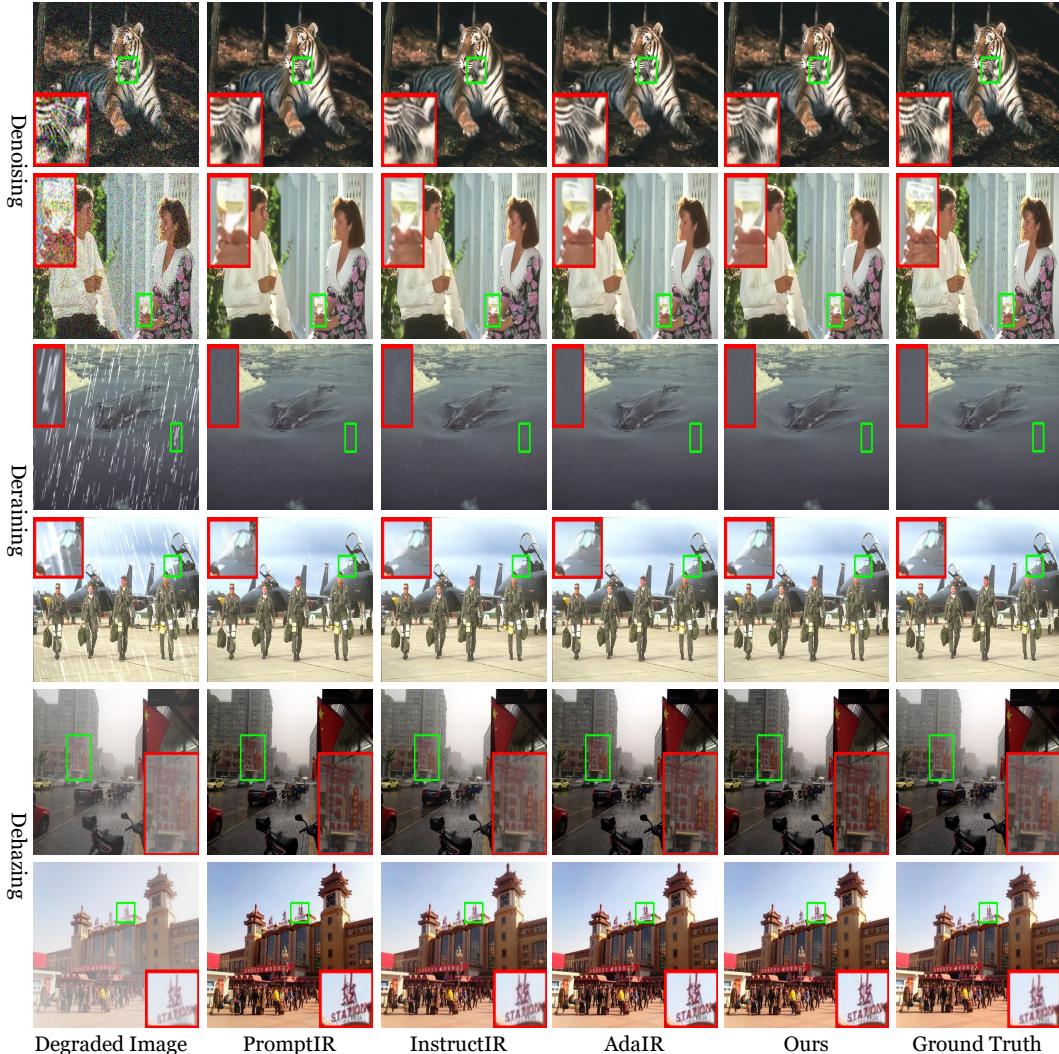
776 The activation patterns shown in Figure 7, 8, and 9 provide a comprehensive view of the functional
 777 specialization within our codebook. While each task has a unique activation signature, the full dis-
 778 tributions allow for deeper analysis. For instance, one can observe the long tail of less-activated
 779 primitives and how their relative importance shifts between tasks. These plots serve as the foun-
 780 dational data from which the more condensed visualizations and analyses in the main paper are
 781 derived.

782
 783 **A.4 ADDITIONAL QUALITATIVE RESULTS**

784 To further demonstrate the superior performance and robustness of our proposed DACode frame-
 785 work, this section provides additional qualitative results. We present a wider variety of challenging
 786 visual examples for the Denoising, Deraining, and Dehazing tasks, comparing our method against
 787 several state-of-the-art approaches. As shown in Figure 11, the columns compare results from differ-
 788 ent methods against our proposed DACode and the Ground Truth. In the **Denoising** examples, our
 789 method excels at restoring intricate textures and natural colors (e.g., the tiger’s fur and the pattern
 790 on the book cover), avoiding the color shifts or blurriness present in other methods. For **Deraining**,
 791 DACode effectively removes heavy rain streaks while preserving challenging background details
 792 (e.g., the jets on the tarmac and the soldiers’ uniforms). In the **Dehazing** scenarios, our approach
 793 recovers significantly more vibrant colors and sharper text details (e.g., the “STATION” sign), out-
 794 performing competitors that leave a residual hazy appearance.

795 To visually substantiate these strong quantitative results on composite degradations, we present
 796 challenging qualitative comparisons in Figure 10. A consistent pattern emerges across scenarios
 797 like haze-snow and low-light-haze-rain: while competing methods often struggle to handle the co-
 798 occurrence of multiple degradations—resulting in residual haze, color casts, or detail loss—our
 799 DACode-S demonstrates a superior ability to disentangle them. For instance, in the windmill scene
 800 (bottom row), DACode-S is the only method to effectively remove the thick atmospheric inter-
 801 ference and restore the fine structural details, closely matching the ground truth. Similarly, in the beach
 802 scene (middle row), our approach excels at recovering the vibrant green hues of the foliage and cor-
 803 recting the severe color cast, whereas other methods yield washed-out or color-biased results. These
 804 visual results provide compelling, intuitive evidence for our compositional paradigm, highlighting
 805 its effectiveness in generating high-fidelity restorations for complex degradation mixtures.

806
 807
 808
 809



853 **Figure 11: Additional qualitative comparisons for all-in-one restoration on the 3-task bench-**
 854 **mark.** The figure presents two additional challenging examples for each of the three tasks: **Denois-**
 855 **ing** (rows 1-2), **Deraining** (rows 3-4), and **Dehazing** (rows 5-6).