

Hybrid Agents for Image Restoration

Bingchen Li Xin Li[†] Yiting Lu Zhibo Chen[†]
University of Science and Technology of China
{lbc31415926, luyt31415}@mail.ustc.edu.cn
{xin.li, chenzhibo}@ustc.edu.cn

Abstract

Existing Image Restoration (IR) studies typically focus on task-specific or universal modes individually, relying on the mode selection of users and lacking the cooperation between multiple task-specific/universal restoration modes. This leads to insufficient interaction for unprofessional users and limits their restoration capability for complicated real-world applications. In this work, we present HybridAgent, intending to incorporate multiple restoration modes into a unified image restoration model and achieve intelligent and efficient user interaction through our proposed hybrid agents. Concretely, we propose the hybrid rule of fast, slow, and feedback restoration agents. Here, the slow restoration agent optimizes the powerful multimodal large language model (MLLM) with our proposed instruction-tuning dataset to identify degradations within images with ambiguous user prompts and invokes proper restoration tools accordingly. The fast restoration agent is designed based on a lightweight large language model (LLM) via in-context learning to understand the user prompts with simple and clear requirements, which can obviate the unnecessary time/resource costs of MLLM. Moreover, we introduce the mixed distortion removal mode for our HybridAgents, which is crucial but not concerned in previous agent-based works. It can effectively prevent the error propagation of step-by-step image restoration and largely improve the efficiency of the agent system. We validate the effectiveness of HybridAgent with both synthetic and real-world IR tasks.

1. Introduction

Image restoration (IR) has long been a popular topic in the low-level research fields [6, 26, 31, 32, 73], which aims to restore low-quality inputs (LQs) into high-quality outputs (HQs). Early IR works primarily focus on model design for diverse simple single restoration tasks, such as image denoising [12, 74, 75, 77], deblurring [41, 44, 57, 66],

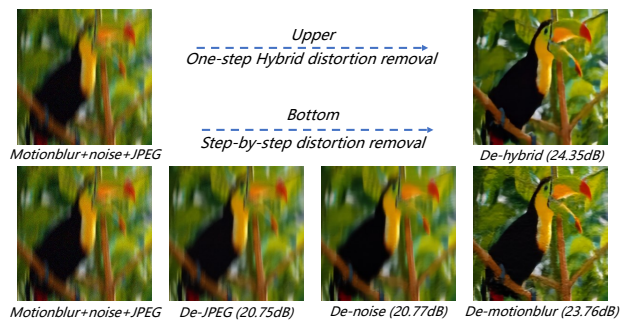


Figure 1. Removing hybrid distortions step-by-step will cause suboptimal results due to distortion entanglement.

compression artifacts removal [15, 16, 22], and super-resolution [6, 9, 23, 62, 76]. However, deploying multiple IR models simultaneously to handle diverse degradations in real-world applications incurs significant resource costs. Recently, all-in-one/universal image restoration models (AIR) [7, 10, 24, 46] have emerged to address diverse image restoration tasks within a single IR network by learning degradation descriptors through prompt learning or instruction tuning.

Despite the abundant restoration modes in existing IR studies, the restoration process still depends on professional users for mode selection when dealing with complex real-world degradations, lacking flexible cooperation between different restoration modes. For instance, an untrained user who does not understand terms like “noise” or “blur” may struggle to select the appropriate restoration tools in a step-by-step manner. To eliminate this, language-based interaction and planning strategies have been proposed to support automatic image restoration [5, 81] with multimodal large language models (MLLMs). However, existing methods suffer from two key challenges: (i) a single interactive agent lacks adaptability to varying task complexities, as even simple and clear instructions, such as “please remove the noise in the image,” are still processed using heavy-weight MLLMs, leading to unnecessary computational and time costs; (ii) error propagation occurs when step-by-step

[†]Corresponding authors.

single degradation removal tools are applied to real-world hybrid degradations, as illustrated in Fig. 1.

In this work, we propose HybridAgent, a novel interactive paradigm for universal image restoration that addresses the aforementioned challenges through two key innovations: (i) the hybrid restoration agents, consisting of fast, slow and feedback agents to balance efficiency and effectiveness during automatic restoration process; and (ii) mixed distortion removal tools integrated with single distortion removal tools to enhance adaptability and prevent error propagation during restoration process.

Specifically, we apply an MLLM as our slow restoration agent (**SlowAgent**) due to its powerful understanding, reasoning, and decision-making capabilities. To adapt SlowAgent to IR, we construct an instruction-tuning dataset designed for complex image restoration tasks. We then fine-tune the SlowAgent, equipping it with strong task understanding, distortion identification, and restoration mode/tool planning capabilities. However, user interaction requirements are not always complex, and not every task necessitates the agent to identify degradations within images. For simple interactive IR tasks (e.g., “please remove noise”), applying an MLLM would be inefficient and resource-intensive. To enable flexible and efficient interaction for different IR requirements, we introduce the fast restoration agent (**FastAgent**) by leveraging a lightweight large language model (LLM) to determine whether the user prompt is clear and simple. If so, the FastAgent will directly execute the planning for the simple IR task and invoke the corresponding restoration tool. Otherwise, the task is handed over to SlowAgent, which automatically identifies degradations and performs the restoration. Notably, to achieve an automated restoration process, SlowAgent requires external feedback to determine whether the current image still needs further restoration. To better facilitate SlowAgent in completing the restoration process, we further develop a **FeedbackAgent** to assess whether a restored image is free from degradations. The three restoration agents collaborate to enable the automatic restoration of distorted images, formulating our HybridAgent.

Another essential component of HybridAgent is the restoration tools, which are invoked by the agents during the restoration process. A straightforward approach would be to directly utilize state-of-the-art single-task restoration models as tools [5, 72, 81]. However, this approach fails to exploit shared knowledge across tasks. Moreover, distributional discrepancies between different models may cause error propagation during the restoration process. To address this, we propose a three-stage training paradigm for constructing our restoration tools, which not only enables the tools to effectively preserve common knowledge shared across different restoration tasks but also provides an intuitive approach for building both single distortion and

mixed distortion removal tools. In the first stage, following the pretraining-finetuning paradigm, we train a foundation restoration model via a multi-task learning scheme [24, 46] to learn the task-shared knowledge. Then, in the second stage, we leverage Low-Rank Adaptation (LoRA) [19] to efficiently fine-tune single distortion removal tools from the pre-trained foundation model. Consequently, in the third stage, we further tailor a mixed distortion removal tool based on the pre-trained foundation model with a new set of LoRA weights. Notably, in our three-stage training paradigm, we adopt prompt components [33, 46] to implicitly encode distortion information, while employing LoRA to efficiently adapt the pre-trained model to domain shifts caused by different distortions.

To cover the majority of application scenarios, we train all restoration tools and build the instruction dataset based on 10 degradations with various levels, including: noise [12], gaussian blur, motion blur [57], JPEG [16], HEVC [55], VVC [2], rainstreak [70], raindrop [47], haze [11], low light [61]. We incorporate two less-explored compression codecs, HEVC [55] and VVC [2], given their rising adoption in image and video compression within contemporary vision applications.

The contributions of this paper can be summarized as follows:

- We propose HybridAgent, a novel interactive paradigm for image restoration that integrates fast, slow, and feedback agents, enabling efficient and task-adaptive interactions to meet diverse user requirements and handle various degradations within a unified restoration framework.
- We identify the error propagation issue associated with step-by-step single-distortion removal strategies and propose a three-stage training paradigm to construct restoration tools. This approach enhances the reuse of shared knowledge across tasks and introduces a mixed distortion removal pattern, effectively reducing restoration steps and mitigating error propagation in image restoration.
- We build an instruction tuning dataset with over 100k image-text pairs across 10 distortion types to tailor our HybridAgent. The diversity of our dataset ensures the applicability of HybridAgent.

2. Related Works

2.1. Image Restoration

Single-task Image Restoration. Image restoration (IR) [8, 12, 16, 41, 44, 45, 74] typically focuses on addressing a specific type of image degradation. Early works leverage convolutional neural networks (CNNs) due to their effective local information processing capabilities [8, 9, 35, 79]. With the evolution of Transformers [60], which excel at uncovering long-range dependencies within the image, a series of transformer-based IR works [6, 34, 63, 73] have

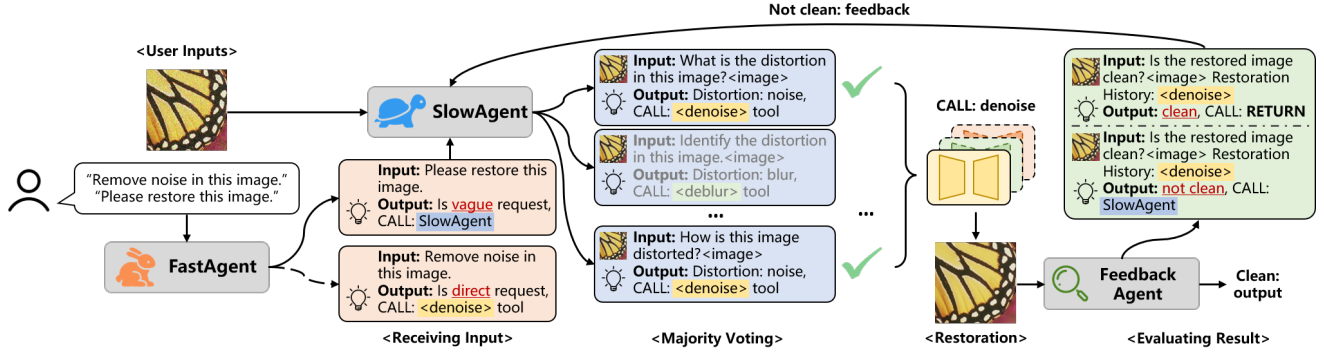


Figure 2. The overall pipeline of **HybridAgent**. We adopt a **FastAgent** to determine swiftly whether the user prompt is direct or vague. If a direct prompt is provided, **HybridAgent** will switch to the fast route (dashed lines) to invoke the corresponding restoration tool. Otherwise, **HybridAgent** will trigger the slow route (solid lines). **SlowAgent** automatically recognizes the distortion and executes the right restoration tool. To prevent incorrect tool invocation, we introduce a **FeedbackAgent** to assess whether the restored image is clean. **FeedbackAgent** and **SlowAgent** work collaboratively to generate the final clean output for the user.

emerged. However, image restoration optimized with L1 or MSE loss often results in over-smoothed outputs [23]. To mitigate this, recent works introduce diffusion techniques [18, 54] into the IR network design. Benefiting from the powerful generation abilities, diffusion-based IR models [13, 49, 56, 68, 71] can restore more realistic images with vivid texture details.

Nevertheless, these works only focus on removing one type of distortion, which lacks adaptability towards various real-world application scenarios.

All-in-one Image Restoration. To address the above-mentioned problem, some works propose all-in-one IR models, which can handle multiple distortions with a unified model weight. Pioneer work AirNet [24] leverages contrastive learning to implicitly learn mappings between various degradation distributions and clean ones. Inspired by prompt learning in natural language processing (NLP), a series of works [30, 33, 37, 38, 46, 49] have studied prompt learning-based AIR, where a set of prompt parameters are implicitly encoded to capture various distortion representations. Meanwhile, due to the inherently multi-distorted nature of weather, several all-in-one approaches [17, 28, 43] have emerged for removing weather-related distortions such as haze, rain, and snow.

Although these approaches improve efficiency in handling various distortions, they lack control over distortion removal, *i.e.*, they cannot selectively remove distortions according to user prompts. Recently, works such as InstructIR [7] and UniProcessor [10] leveraged text encoders to map user instructions as conditions of restoration models, enabling user-controllable AIR. However, these works lack generalization towards diverse user prompts due to limited encoding ability. Moreover, they adopt a step-by-step approach to handle hybrid distortions, which leads to suboptimal restoration results.

2.2. Agent

Agent typically refers to an intelligent system that can receive diverse user commands and automatically accomplish tasks accordingly [51]. Thanks to the great success of general-purposed large language models (LLMs) [3, 50, 59] and multimodal LLMs (MLLMs) [27, 36, 42], such intelligent system can be realized through the combination of MLLMs and domain-specific expert models [40, 52, 69].

Despite extensive research on agents in high-level domains, how they can be effectively applied to low-level image restoration (IR) tasks remains an open question. As a pioneer work, RL-Restore [72] proposes to solve complex restoration problems in a sequential way. At each step, RL-Restore utilizes reinforcement learning to adaptively select the most suitable restoration tool for removing a specific type of distortion. Advanced by MLLMs, Clarity ChatGPT [65] enables dynamic tool selection by understanding the user’s commands and performing reasoning. RestoreAgent [5] further optimizes the execution sequence of restoration tools and explores the most suitable model for specific degradation patterns. AgenticIR [81] improves the planning of the agent by incorporating proficient restoration experience, reflection, and rollback strategy.

Nevertheless, the aforementioned works [5, 80–82] all perform image restoration through step-by-step execution, neglecting the entanglement of distortions and potential distribution shift caused by different restoration models [4, 29]. Moreover, they do not account for efficiency during tool invocation. How to adaptively and efficiently select restoration tools based on different user prompts remains an open research question.

3. Methods

In this section, we introduce our HybridAgent, an intelligent agent system that dynamically and automatically han-

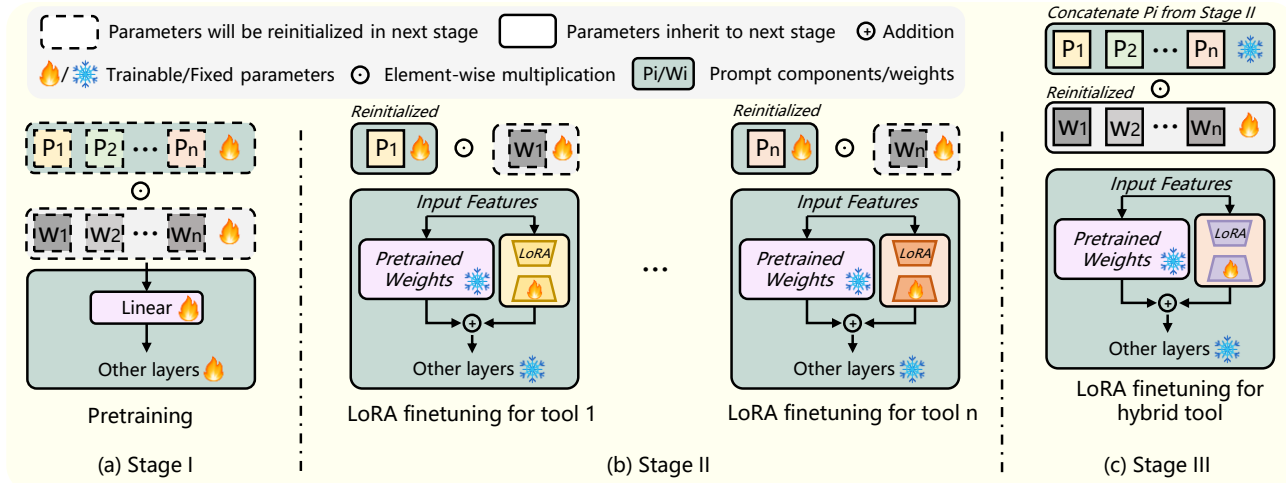


Figure 3. The illustration of three-stage training for the construction of restoration tools. We first build a well-trained base model following prompt learning-based all-in-one image restoration [25, 46] in Stage I. Subsequently, we build single-task restoration tools and hybrid restoration tool with LoRA [19] in Stage II and III. Notably, we add LoRA to weights of *Linear* layers in Attention modules and FeedForward modules of [25]. A more detailed diagram is provided in the **supplementary** due to limited space.

dles complex IR problems. We discuss the details about how we realize the automatic image restoration through MLLM agents in Section 3.1. Specifically, to address the inefficiency of MLLM-based agents when handling simple user requests, we propose a collaborative agent system composed of FastAgent, SlowAgent, and FeedbackAgent, which efficiently responds to diverse user prompts and effectively solves complex IR tasks. Subsequently, we describe how we construct the restoration tools for HybridAgent in Section 3.2. Particularly, we propose a mixed distortion removal tool to proficiently address the error propagation problem in the step-by-step restoration process.

3.1. HybridAgent

Imagine a user aiming to restore a degraded image. Typically, the user begins by identifying distortion types within the image and then searches for suitable restoration models. After applying a restoration step, the individual evaluates whether the image has been adequately restored. If degradations persist, the user continues to select and apply appropriate restoration tools iteratively until achieving a satisfactory restoration outcome. HybridAgent follows this same iterative procedure to automatically and efficiently restore degraded images, as shown in Figure 2.

Receiving Input. HybridAgent receives a distorted image along with a user-provided prompt. Typically, an agent utilizes MLLMs to perceive the degradations and plan restoration steps automatically. However, professional users may provide explicit prompts clearly specifying their restoration requirements. In such cases, using MLLMs for degradation recognition becomes unnecessary, as the intended restoration type has already been identified. To handle this

scenario efficiently, we employ the FastAgent based on a lightweight LLM. FastAgent directly analyzes the explicit user prompts, determines the required restoration tools, and proceeds immediately to the subsequent restoration step. On the contrary, if the prompt provided by the user is deemed ambiguous or unclear by FastAgent, the image will be handed over to SlowAgent and proceed to the next restoration step.

Identifying Distortion. If the image is handed over to SlowAgent, it is crucial to identify the proper distortion types and yield the suitable restoration tools. However, current MLLMs still face several challenges when handling such tasks: (i) existing MLLMs are not specifically fine-tuned for distortion type recognition, thus limiting their ability to accurately identify degradations, (ii) they cannot directly execute image restoration operations, and (iii) they lack the capability to determine whether further restoration steps are necessary.

To address the first challenge, we adopt an existing MLLM, Co-instruct [67], which is specifically tuned for image quality assessment (IQA) and possesses distortion recognition capabilities. However, the degradation types covered by Co-instruct are insufficient for handling complex real-world requirements. Therefore, we further fine-tune Co-instruct on our proposed instruction-tuning dataset (detailed in the **supplementary**) to extend its distortion-recognition capabilities, resulting in the proposed SlowAgent. However, MLLMs may hallucinate and produce incorrect judgments. To mitigate this issue, inspired by the concept of test-time scaling [53], we employ a majority voting mechanism that generates multiple candidate decisions and selects the most frequently occurring distortion

as the final recognized distortion. To address the second challenge, we construct a set of restoration tools (detailed in Section 3.2) that can be invoked by SlowAgent.

Evaluating Result. To address the third challenge, we establish a FeedbackAgent to assess the restored image. Such a FeedbackAgent is crucial because: (i) without external evaluation, SlowAgent cannot reliably determine if the restored image is already satisfactory; and (ii) it offers flexibility to terminate/continue the restoration process when a human user finds the intermediate result acceptable/unpleasant. Fortunately, determining whether a restored image is clean aligns closely with assessing image quality. However, leveraging IQA scores alone may not directly reflect whether an image meets restoration criteria. Thus, we propose the simplest way to provide feedback for SlowAgent: fine-tuning an IQA model (Co-Instruct) specifically to classify whether a restored image is clean. Nevertheless, achieving absolute perfection in restoration is nearly impossible in practice. Consequently, we further incorporate the historical information of the chosen restoration tools as the context for the FeedbackAgent, enabling it to more reliably determine whether the current image has reached a relatively clean state or still requires further restoration steps.

3.2. Restoration Tools

After determining which tool should be invoked, the distorted image will be processed by the chosen restoration tool. Current tool designs [5, 72, 81] typically suffer from two primary limitations: (i) they directly adopt multiple single-task models trained separately, thus failing to effectively leverage common knowledge across different restoration tasks, and (ii) sequential application of these single-task models leads to step-by-step processing, which struggles to resolve distortion entanglement issues [4, 29]. To simultaneously address these challenges, we propose a novel three-stage training strategy for restoration tools.

Stage I. In stage one, we aim to build a well-trained base model that shares common knowledge across various IR tasks. Inspired by the promising results of prompt learning-based methods [46], we adopt a similar network architecture. (We provide the details of this architecture in the **supplementary**.) To enhance the model’s ability to represent diverse distortions (e.g., 10 types in this paper), we follow [25] and replace the first two stages’ transformer blocks with shifted window attention blocks [6, 34].

Stage II. After obtaining the base model, we fine-tune task-specific models on different distortions to enable the agent to invoke the corresponding restoration tools. Once the base model is trained in stage I, it serves as a foundation model for distortion removal, allowing efficient fine-tuning to obtain the required restoration tools. Given LoRA’s [19] effectiveness with minimal additional parameters, we employ it

for fine-tuning in this stage, as shown in Figure 3. Additionally, to ensure that the prompt accurately conveys the degradation condition to the model, we reinitialize the prompt parameters and fine-tune them jointly with LoRA parameters. Notably, we leverage prompts to encode descriptive information regarding distortions, while employing LoRA parameters to efficiently adapt the degradation-aware semantic information within deeper network structures.

Stage III. To address the challenge of mixed distortion removal in complex IR problems, we further tailor a mixed distortion removal tool. Similar to Stage II, we employ LoRA for efficient fine-tuning. However, we initialize the prompt parameters using those obtained in Stage II, enabling the network to effectively leverage both task-specific knowledge from Stage II and common knowledge from Stage I. We follow the distortion synthesis pipeline proposed in Real-ESRGAN [62] for training our model. More details are provided in the **supplementary**.

4. Experiments

4.1. Implementation Details

Structure Details for HybridAgent. Since most MLLMs are designed for general purposes, we fine-tune existing MLLMs to enable them to effectively serve the roles of SlowAgent and FeedbackAgent. We choose to fine-tune Co-Instruct [67] because it is capable of assessing distortions in images, providing an excellent starting point for our agents. As for restoration tools, we leverage an enhanced version of PromptIR [46] from [25], where the transformer blocks in the first two stages are replaced with RHAG [6] to improve the representative abilities. For the FastAgent, we adopt Llama3.2-1B-Instruct¹.

Training Datasets for Restoration Tools. We optimize our restoration tools on 10 degradations. For noise, gaussian blur, motion blur, JPEG, HEVC, and VVC, we generate distorted samples online using 3450 images from DF2K [1, 58]. For other distortions, we adopt 1800 images from Rain100H [70] for rainstreak, 861 images from RainDrop [47] for raindrop, 6000 images from RESIDE-6k [48] for haze, and 485 images from LOL [64] for low light following previous work [37]. A more detailed explanation of degradation levels for synthesized images and the training details are given in the **supplementary**.

Instruction Tuning Datasets for Agents. Instruction tuning dataset plays an essential role in fine-tuning MLLM agents. Based on training datasets introduced above, we sample 5k images per distortion with resolution ranging between 224×224 and 784×784 . We apply a linear transformation to map 5k to the total number of images within different distortion datasets, ensuring even sampling from each dataset. For hybrid distortion, we generate 20k images

¹<https://huggingface.co/meta-llama/Llama-3.2-1B-Instruct>

Table 1. Comparison of proficiency and performance based on whether the fast route is enabled. ‘‘A.I.T.’’ represents ‘‘Average Inference Time’’, which indicates total inference time of agents. The inference time is evaluated on a RTX 4090D GPU. We use PSNR \uparrow /SSIM \uparrow to evaluate the performance of two different settings.

Setting	A.I.T. (s)	Performance	A.I.T. (s)	Performance	A.I.T. (s)	Performance	A.I.T. (s)	Performance	A.I.T. (s)	Performance
	De-noise		De-blur		De-motionblur		De-jpeg		De-HEVC	
a)	0.08	30.25/0.867	0.11	30.65/0.853	0.09	23.80/0.720	0.09	30.02/0.873	0.09	27.56/0.785
b)	0.75	30.63/0.874	0.82	30.52/0.852	0.90	23.78/0.719	0.79	30.18/0.876	0.76	27.55/0.785
	De-VVC		De-rainstreak		De-raindrop		De-haze		De-low light	
a)	0.09	27.91/0.797	0.13	30.04/0.893	0.12	30.30/0.913	0.09	29.92/0.960	0.12	22.60/0.825
b)	0.79	27.92/0.798	1.05	30.03/0.893	0.94	30.34/0.914	0.83	29.92/0.960	0.88	22.61/0.828

Table 2. Comparisons of success rate between FastAgent and SlowAgent. As both agents will invoke tools, we define success rate as the proportion of the number of correct tool invocations to the total number of tool invocations.

	De-noise	De-blur	De-motionblur	De-jpeg	De-HEVC	De-VVC	De-rainstreak	De-raindrop	De-haze	De-low light
FastAgent	72.9%	100.0%	100.0%	84.8%	96.2%	60.0%	100.0%	86.2%	100.0%	73.3%
SlowAgent	94.3%	94.3%	99.0%	96.7%	87.6%	90.9%	98.0%	96.5%	100.0%	100%

based on the combination of 10 distortions (more details are given in the **supplementary**). In total, our instruction tuning dataset for SlowAgent contains 70k image-text pairs. As for FeedbackAgent, we use images restored by the correct tools to synthesize 30k *relatively* ‘‘clean’’ images and 33k ‘‘not clean’’ images by the incorrect tools. This results in a total of 66k image-text pairs for FeedbackAgent.

Test Datasets. Following [10], we adopt the combination of CBSD68 [39], Urban100 [20], Kodak24 [14], and McMaster [78] to evaluate the performance on first six distortions. Following [10], we use 100 images from Rain100H [70], 58 images from RainDrop [47], 1000 images from RESIDE-6k [48], and 15 images from LOL [64] to evaluate the performance on last four distortions, respectively. We generate 200 images for mixed-degradation based on the mixture of 10 types of distortions. We provide more details including training settings in the **supplementary**.

User Prompts. We use GPT-4 to generate 20 direct textual prompts for each distortion type, along with an additional 20 ambiguous prompts, formulating a total of 220 diverse user prompts. We provide samples in the **supplementary**.

4.2. Effectiveness of HybridAgent

In this section, we evaluate the effectiveness of HybridAgent by addressing two key questions: i) Does FastAgent enhance the proficiency of the restoration pipeline? ii) Does mixed distortion removal outperform step-by-step distortion removal? Can they collaborate to address more complex distortions?

4.2.1. Effectiveness of FastAgent Design

We compare two settings: a) the full HybridAgent, and b) HybridAgent with the fast route disabled, meaning that all

images are processed through the SlowAgent, representing a traditional agent design. *Notice that, to provide a more intuitive comparison, we use randomly selected direct user prompt in both settings.* We evaluate the average inference time and performance on 10 distortions, as demonstrated in Table 1. With the fast route enabled, HybridAgent achieves significantly higher operational efficiency for direct prompts (requiring only about 12% of the runtime compared to SlowAgent), greatly surpassing the efficiency of SlowAgent. We further report the success rate of FastAgent and SlowAgent in Table 2, defined as the proportion of the number of correct tool invocations to the total number of tool invocations. As observed, using in-context learning with FastAgent achieves a relatively high success rate, indicating that FastAgent can adapt to diverse user prompts in real-world scenarios and make accurate tool invocations.

4.2.2. Single vs. Mixed Distortion Removal

As discussed in Section 1 and Section 3.2, removing mixed-degradation step-by-step may cause distribution shift and error propagation. We further validate this by comparing two settings: i) step-by-step restoration using only single distortion removal tools, and ii) restoration using both single and mixed distortion removal tools. As demonstrated in Table 3, mixed distortion removal significantly outperforms step-by-step single distortion removal across all metrics, particularly in handling haze and low-light distortions. This suggests that addressing mixed degradations with our proposed mixed distortion removal tools effectively mitigates error propagation and distribution shifts. We hypothesize that the poor performance of step-by-step removal for haze and low-light distortions is due to their unstable distortion modeling, which is easily disrupted by additional

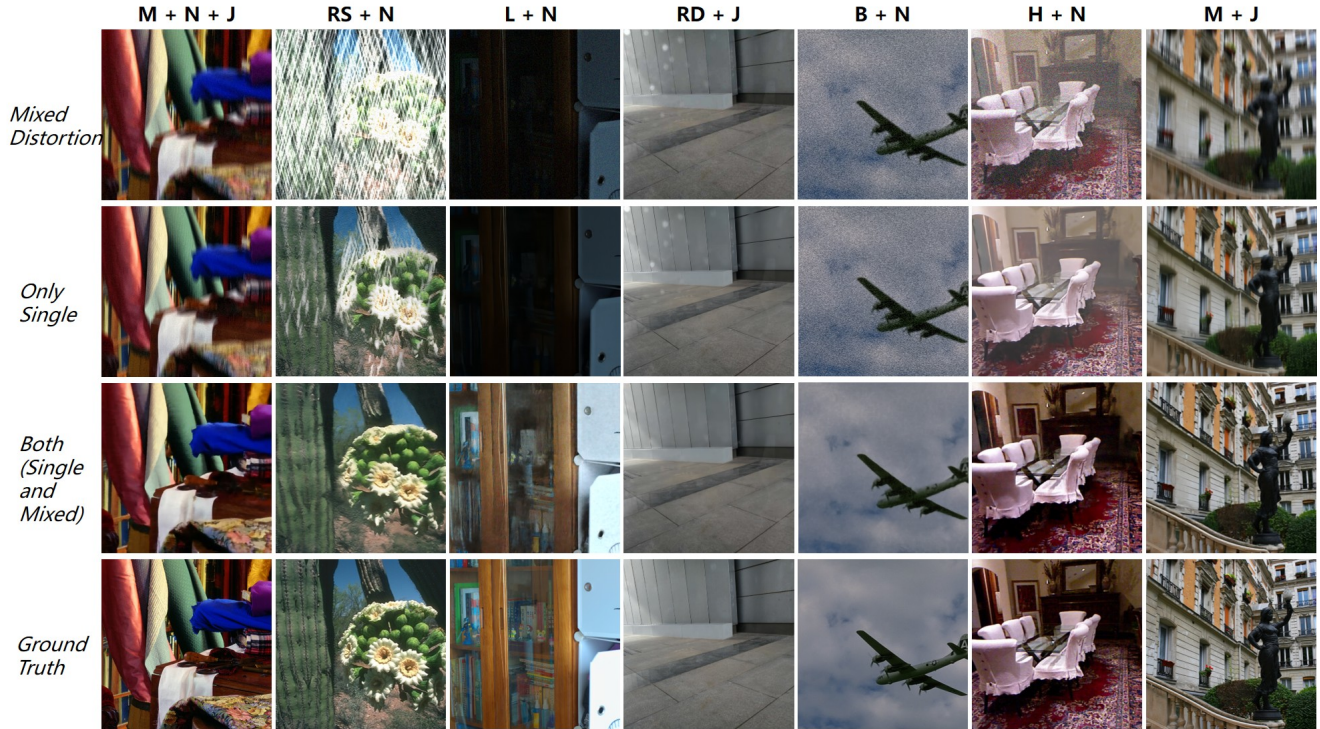


Figure 4. Qualitative comparisons of only single distortion removal tools against single and mixed distortion removal tools. **M**: Motionblur, **N**: Noise, **J**: JPEG, **RS**: Rainstreak, **L**: Low light, **RD**: Raindrop, **B**: Blur, **H**: Haze. Zoom in for a better view.

distortions. This is also evidenced in Figure 4, where step-by-step restoration fails to enhance the low-light image or remove hazy artifacts.

However, the representation ability of the mixed distortion removal tool is limited, constraining its applicability to more complex distortions or real-world scenarios. Nevertheless, the HybridAgent incorporates the FeedbackAgent, allowing the mixed distortion removal tool to be utilized in step-by-step scheduling. Consequently, the mixed distortion removal tool and single distortion removal tools can collaboratively address more complex distortion scenarios, reducing the step of tool invocations and improving overall performance. We provide a case study in Figure 5. The collaboration between the mixed distortion and single distortion removal tools reduces distribution shifts and achieves better restoration results in fewer steps.

4.3. Comparisons with Other Methods

To further demonstrate the effectiveness of HybridAgent against state-of-the-art all-in-one/agentic IR methods, we evaluate the performance on complex restoration scenarios, including both synthetic and real-world datasets. For a fair comparison, we have retrained Uformer [63], AirNet [24], Restormer [73], and PromptIR [46] on our proposed 10 degradation datasets. Following [5], we infer each of the above methods multiple times in a sequential way to



Figure 5. A case study on complex degradation removal. The image is corrupted by “Raindrop + Blur + Noise + JPEG”. Upper: step-by-step distortion removal. Bottom: tools invoked by HybridAgent: De-hybrid + De-raindrop.

achieve the best performance. For InstructIR [7], we test the method on seen distortions, ensuring a fair comparison with HybridAgent. We further compare our method with AgenticIR [81], which also employs an agentic design. As demonstrated in Table 4, our method achieves significant improvements against other all-in-one methods on mixed

Table 3. Comparison of removing mixed distortions between only using single distortion removal tools and using both single and mixed distortion removal tools.

	Blur + Noise			Blur + JPEG			Blur + Noise + JPEG			Motionblur + Noise			Motionblur + JPEG		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Only Single	23.72	0.555	0.520	26.04	0.737	0.300	22.32	0.423	0.640	22.10	0.532	0.501	23.79	0.668	0.360
Both	26.21	0.733	0.311	26.54	0.775	0.278	25.37	0.706	0.352	23.13	0.628	0.388	23.77	0.674	0.330
	Motionblur + Noise + JPEG			Rainstreak + Noise			Rainstreak + JPEG			Raindrop + Noise			Raindrop + JPEG		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Only Single	20.66	0.439	0.551	23.36	0.622	0.351	22.18	0.692	0.300	26.51	0.743	0.287	25.40	0.814	0.233
Both	22.82	0.618	0.404	26.49	0.766	0.203	25.44	0.764	0.242	27.98	0.807	0.191	28.42	0.835	0.188
	Haze + Noise			Haze + JPEG			Low light + Noise			Low light + JPEG			Average		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓
Only Single	14.42	0.536	0.497	12.57	0.668	0.302	7.69	0.188	0.761	7.05	0.174	0.616	19.84	0.557	0.444
Both	23.40	0.797	0.184	26.43	0.885	0.113	19.97	0.627	0.475	21.65	0.754	0.324	24.83	0.741	0.284

Table 4. Comparison of HybridAgent with All-in-one image restoration methods. For Uformer, Restormer, and PromptIR, we retrain them following their official code on our 10 degradation datasets. Best performances are **bolded**. Additionally, we provide averaged inference time (A.I.T.), averaged VRAM usage, and model size in the table.

	Rainstreak + Noise			Haze + Noise			Low light + Noise			Motionblur + Noise			Complexity Analysis		
	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	PSNR↑	SSIM↑	LPIPS↓	A.I.T.(s)	VRAM(G)	Size(M)
Uformer	21.37	0.565	0.409	16.27	0.587	0.452	9.28	0.289	0.781	21.17	0.494	0.549	1.18	2.55	50.88
Restormer	23.57	0.605	0.377	15.37	0.538	0.500	15.07	0.546	0.603	20.72	0.499	0.544	0.58	9.80	26.13
AirNet	22.95	0.588	0.393	14.33	0.510	0.544	15.72	0.538	0.612	20.68	0.487	0.561	0.79	4.86	8.93
PromptIR	23.17	0.591	0.388	14.54	0.513	0.532	15.87	0.541	0.610	20.79	0.489	0.558	0.58	9.70	35.59
InstructIR	14.63	0.431	0.534	18.99	0.615	0.426	19.15	0.599	0.526	20.99	0.493	0.542	9.21	7.81	15.94
AgenticIR	25.98	0.758	0.216	22.55	0.781	0.209	20.12	0.632	0.479	23.69	0.668	0.358	49.34	-	- ²
HybridAgent	26.49	0.766	0.203	23.40	0.797	0.184	19.97	0.627	0.475	23.77	0.674	0.330	11.66	32.24	400.82

degradations. Compared with AgenticIR, our method is **significantly more efficient** (11.66s vs. 49.34s), primarily because AgenticIR requires multiple GPT API calls and the invocation of tools with different scales. In contrast, our approach only switches between different LoRA modules, which greatly improves the overall efficiency of the agentic system. Additionally, HybridAgent incorporates the mixed distortion removal tool, which efficiently mitigates the impact of incorrect restoration decisions. For qualitative comparisons, we provide visual results on real-world underwater unpaired dataset EUVP [21] in Figure 6. As observed, HybridAgent first takes mixed distortion removal tool to remove entangled degradations, and then utilizes de-haze tool to further enhance the visual quality of the image. We provide more qualitative results and other ablation studies about our three-stage training design of restoration tools in the **supplementary** due to limited space.

5. Conclusion

In this work, we present HybridAgents by introducing the composition of fast and slow restoration agents, intending to solve two challenges in existing agent/instruction-based image restoration works: (i) lacking flexibility on task dif-




Figure 6. Visualization of the restoration process performed by HybridAgent on real-world image (test9003up.jpg from EUVP [21] dataset). Zoom in for best views.

ficulty; and (ii) the error propagation since the step-by-step single restoration tool invoking. To tackle the first challenge, we introduce hybrid restoration agents, where FastAgent is responsible for simple and clear user requirements while SlowAgent is optimized with our proposed large instruction tuning dataset to support ambiguous user requirements. A FeedbackAgent is also designed to collaborate with SlowAgent to provide accurate feedback and perform the termination of the restoration process. For the second challenge, we propose a three-stage training strategy which introduces the abundant mixed distortion removal tool and enhances the model reusing capability by optimizing the model with multi-task learning and task-specific prompt optimization. Extensive experiments on broad user requirements and various complicated degradations have demonstrated the effectiveness of our HybridAgent.

²We only report A.I.T. since it relies on closed-source GPT.

Acknowledgments

This work was supported in part by NSFC under Grant 62371434, U25B2010 and 623B2098, the Postdoctoral Fellowship Program of CPSF under Grant Number GZC20252293, the China Postdoctoral Science Foundation-Anhui Joint Support Program under Grant Number 2024T017AH, China Postdoctoral Science Foundation under Grant Number 2025M783529, Anhui Postdoctoral Scientific Research Program Foundation (No.2025A1015), the Fundamental Research Funds for the Central Universities (No. WK2100250064).

References

- [1] Eirikur Agustsson and Radu Timofte. Ntire 2017 challenge on single image super-resolution: Dataset and study. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 126–135, 2017. 5
- [2] Benjamin Bross, Ye-Kui Wang, Yan Ye, Shan Liu, Jianle Chen, Gary J Sullivan, and Jens-Rainer Ohm. Overview of the versatile video coding (vvc) standard and its applications. *IEEE Transactions on Circuits and Systems for Video Technology*, 31(10):3736–3764, 2021. 2
- [3] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020. 3
- [4] Jin Cao, Deyu Meng, and Xiangyong Cao. Chain-of-restoration: Multi-task image restoration models are zero-shot step-by-step universal image restorers. *arXiv preprint arXiv:2410.08688*, 2024. 3, 5
- [5] Haoyu Chen, Wenbo Li, Jinjin Gu, Jingjing Ren, Sixiang Chen, Tian Ye, Renjing Pei, Kaiwen Zhou, Fenglong Song, and Lei Zhu. Restoreagent: Autonomous image restoration agent via multimodal large language models. *Advances in Neural Information Processing Systems*, 37: 110643–110666, 2024. 1, 2, 3, 5, 7
- [6] Xiangyu Chen, Xintao Wang, Jiantao Zhou, Yu Qiao, and Chao Dong. Activating more pixels in image super-resolution transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 22367–22377, 2023. 1, 2, 5
- [7] Marcos V Conde, Gregor Geigle, and Radu Timofte. Instructir: High-quality image restoration following human instructions. In *European Conference on Computer Vision*, pages 1–21. Springer, 2024. 1, 3, 7
- [8] Chao Dong, Yubin Deng, Chen Change Loy, and Xiaoou Tang. Compression artifacts reduction by a deep convolutional network. In *Proceedings of the IEEE international conference on computer vision*, pages 576–584, 2015. 2
- [9] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015. 1, 2
- [10] Huiyu Duan, Xiongkuo Min, Sijing Wu, Wei Shen, and Guangtao Zhai. Uniprocessor: a text-induced unified low-level image processor. In *European Conference on Computer Vision*, pages 180–199. Springer, 2025. 1, 3, 6
- [11] Deniz Engin, Anil Genç, and Hazim Kemal Ekenel. Cycledehaze: Enhanced cyclegan for single image dehazing. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 825–833, 2018. 2
- [12] Linwei Fan, Fan Zhang, Hui Fan, and Caiming Zhang. Brief review of image denoising techniques. *Visual Computing for Industry, Biomedicine, and Art*, 2(1):7, 2019. 1, 2
- [13] Ben Fei, Zhaoyang Lyu, Liang Pan, Junzhe Zhang, Weidong Yang, Tianyue Luo, Bo Zhang, and Bo Dai. Generative diffusion prior for unified image restoration and enhancement. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9935–9946, 2023. 3
- [14] Rich Franzen. Kodak lossless true color image suite. <http://r0k.us/graphics/kodak/>, 1999. Online accessed 24 Oct 2021. 6
- [15] Xueyang Fu, Xi Wang, Aiping Liu, Junwei Han, and Zheng-Jun Zha. Learning dual priors for jpeg compression artifacts removal. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 4086–4095, 2021. 1
- [16] Leonardo Galteri, Lorenzo Seidenari, Marco Bertini, and Alberto Del Bimbo. Deep generative adversarial compression artifact removal. In *Proceedings of the IEEE international conference on computer vision*, pages 4826–4835, 2017. 1, 2
- [17] Yu Guo, Yuan Gao, Yuxu Lu, Huilin Zhu, Ryan Wen Liu, and Shengfeng He. Onerestore: A universal restoration framework for composite degradation. *arXiv preprint arXiv:2407.04621*, 2024. 3
- [18] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020. 3
- [19] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021. 2, 4, 5
- [20] Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5197–5206, 2015. 6
- [21] Md Jahidul Islam, Youya Xia, and Junaed Sattar. Fast underwater image enhancement for improved visual perception. *IEEE Robotics and Automation Letters*, 5(2):3227–3234, 2020. 8
- [22] Jiayi Jiang, Kai Zhang, and Radu Timofte. Towards flexible blind jpeg artifacts removal. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4997–5006, 2021. 1
- [23] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017. 1, 3

- [24] Boyun Li, Xiao Liu, Peng Hu, Zhongqin Wu, Jiancheng Lv, and Xi Peng. All-in-one image restoration for unknown corruption. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 17452–17462, 2022. [1](#), [2](#), [3](#), [7](#)
- [25] Bingchen Li, Xin Li, Yiting Lu, Ruoyu Feng, Mengxi Guo, Shijie Zhao, Li Zhang, and Zhibo Chen. Promptcir: Blind compressed image restoration with prompt learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6442–6452, 2024. [4](#), [5](#)
- [26] Bingchen Li, Xin Li, Hanxin Zhu, Yeying Jin, Ruoyu Feng, Zhizheng Zhang, and Zhibo Chen. Sed: Semantic-aware discriminator for image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 25784–25795, 2024. [1](#)
- [27] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. In *International conference on machine learning*, pages 19730–19742. PMLR, 2023. [3](#)
- [28] Ruoteng Li, Robby T Tan, and Loong-Fah Cheong. All in one bad weather removal using architectural search. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 3175–3185, 2020. [3](#)
- [29] Xin Li, Xin Jin, Jianxin Lin, Sen Liu, Yaojun Wu, Tao Yu, Wei Zhou, and Zhibo Chen. Learning disentangled feature representation for hybrid-distorted image restoration. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXIX 16*, pages 313–329. Springer, 2020. [3](#), [5](#)
- [30] Xin Li, Bingchen Li, Yeying Jin, Cuiling Lan, Hanxin Zhu, Yulin Ren, and Zhibo Chen. Ucip: A universal framework for compressed image super-resolution using dynamic prompt. In *European Conference on Computer Vision*, pages 107–125. Springer, 2024. [3](#)
- [31] Xin Li, Yeying Jin, Xin Jin, Zongwei Wu, Bingchen Li, Yufei Wang, Wenhan Yang, Yu Li, Zhibo Chen, Bihan Wen, et al. Ntire 2025 challenge on day and night raindrop removal for dual-focused images: Methods and results. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pages 1172–1183, 2025. [1](#)
- [32] Xin Li, Yulin Ren, Xin Jin, Cuiling Lan, Xingrui Wang, Wenjun Zeng, Xinchao Wang, and Zhibo Chen. Diffusion models for image restoration and enhancement: A comprehensive survey. *International Journal of Computer Vision*, 133(11):8078–8108, 2025. [1](#)
- [33] Zilong Li, Yiming Lei, Chenglong Ma, Junping Zhang, and Hongming Shan. Prompt-in-prompt learning for universal image restoration. *arXiv preprint arXiv:2312.05038*, 2023. [2](#), [3](#)
- [34] Jingyun Liang, Jiezhong Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. Swinir: Image restoration using swin transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1833–1844, 2021. [2](#), [5](#)
- [35] Bee Lim, Sanghyun Son, Heewon Kim, Seungjun Nah, and Kyoung Mu Lee. Enhanced deep residual networks for single image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 136–144, 2017. [2](#)
- [36] Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023. [3](#)
- [37] Ziwei Luo, Fredrik K Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B Schön. Controlling vision-language models for universal image restoration. *arXiv preprint arXiv:2310.01018*, 2023. [3](#), [5](#)
- [38] Jiaqi Ma, Tianheng Cheng, Guoli Wang, Qian Zhang, Xinggang Wang, and Lefei Zhang. Prores: Exploring degradation-aware visual prompt for universal image restoration. *arXiv preprint arXiv:2306.13653*, 2023. [3](#)
- [39] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings eighth IEEE international conference on computer vision. ICCV 2001*, pages 416–423. IEEE, 2001. [6](#)
- [40] Yao Mu, Qinglong Zhang, Mengkang Hu, Wenhai Wang, Mingyu Ding, Jun Jin, Bin Wang, Jifeng Dai, Yu Qiao, and Ping Luo. Embodiedgpt: Vision-language pre-training via embodied chain of thought. *Advances in Neural Information Processing Systems*, 36:25081–25094, 2023. [3](#)
- [41] Seungjun Nah, Sanghyun Son, Suyoung Lee, Radu Timofte, Kyoung Mu Lee, Liangyu Chen, Jie Zhang, Xin Lu, Xiaojie Chu, Chengpeng Chen, et al. Ntire 2021 challenge on image deblurring. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 149–165, 2021. [1](#), [2](#)
- [42] OpenAI. Gpt-4 technical report, 2023. [3](#)
- [43] Ozan Özdenizci and Robert Legenstein. Restoring vision in adverse weather conditions with patch-based denoising diffusion models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(8):10346–10357, 2023. [3](#)
- [44] Jinshan Pan, Deqing Sun, Hanspeter Pfister, and Ming-Hsuan Yang. Blind image deblurring using dark channel prior. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1628–1636, 2016. [1](#), [2](#)
- [45] Bumjun Park, Songhyun Yu, and Jechang Jeong. Densely connected hierarchical network for image denoising. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pages 0–0, 2019. [2](#)
- [46] Vaishnav Potlapalli, Syed Waqas Zamir, Salman H Khan, and Fahad Shahbaz Khan. Promptir: Prompting for all-in-one image restoration. *Advances in neural information processing systems*, 36:71275–71293, 2023. [1](#), [2](#), [3](#), [4](#), [5](#), [7](#)
- [47] Rui Qian, Robby T Tan, Wenhan Yang, Jiajun Su, and Jiaying Liu. Attentive generative adversarial network for raindrop removal from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2482–2491, 2018. [2](#), [5](#), [6](#)
- [48] Xu Qin, Zhilin Wang, Yuanchao Bai, Xiaodong Xie, and Huizhu Jia. Ffa-net: Feature fusion attention network for single image dehazing. In *Proceedings of the AAAI conference on artificial intelligence*, pages 11908–11915, 2020. [5](#), [6](#)

- [49] Yulin Ren, Xin Li, Bingchen Li, Xingrui Wang, Mengxi Guo, Shijie Zhao, Li Zhang, and Zhibo Chen. Moe-diffir: Task-customized diffusion priors for universal compressed image restoration. In *European Conference on Computer Vision*, pages 116–134. Springer, 2024. 3
- [50] Baptiste Roziere, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, et al. Code llama: Open foundation models for code. *arXiv preprint arXiv:2308.12950*, 2023. 3
- [51] Stuart J Russell and Peter Norvig. *Artificial intelligence: a modern approach*. Pearson, 2016. 3
- [52] Yongliang Shen, Kaitao Song, Xu Tan, Dongsheng Li, Weiming Lu, and Yueting Zhuang. Hugginggpt: Solving ai tasks with chatgpt and its friends in hugging face. *Advances in Neural Information Processing Systems*, 36:38154–38180, 2023. 3
- [53] Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters. *arXiv preprint arXiv:2408.03314*, 2024. 4
- [54] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020. 3
- [55] Gary J Sullivan, Jens-Rainer Ohm, Woo-Jin Han, and Thomas Wiegand. Overview of the high efficiency video coding (hevc) standard. *IEEE Transactions on circuits and systems for video technology*, 22(12):1649–1668, 2012. 2
- [56] Haoze Sun, Wenbo Li, Jianzhuang Liu, Haoyu Chen, Renjing Pei, Xueyi Zou, Youliang Yan, and Yujiu Yang. Coser: Bridging image and language for cognitive super-resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 25868–25878, 2024. 3
- [57] Xin Tao, Hongyun Gao, Xiaoyong Shen, Jue Wang, and Ji-aya Jia. Scale-recurrent network for deep image deblurring. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8174–8182, 2018. 1, 2
- [58] Radu Timofte, Eirikur Agustsson, Luc Van Gool, Ming-Hsuan Yang, and Lei Zhang. Ntire 2017 challenge on single image super-resolution: Methods and results. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 114–125, 2017. 5
- [59] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023. 3
- [60] A Vaswani. Attention is all you need. *Advances in Neural Information Processing Systems*, 2017. 2
- [61] Wenjing Wang, Chen Wei, Wenhan Yang, and Jiaying Liu. Gladnet: Low-light enhancement network with global awareness. In *2018 13th IEEE international conference on automatic face & gesture recognition (FG 2018)*, pages 751–755. IEEE, 2018. 2
- [62] Xintao Wang, Liangbin Xie, Chao Dong, and Ying Shan. Real-esrgan: Training real-world blind super-resolution with pure synthetic data. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1905–1914, 2021. 1, 5
- [63] Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general u-shaped transformer for image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 17683–17693, 2022. 2, 7
- [64] Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. *arXiv preprint arXiv:1808.04560*, 2018. 5, 6
- [65] Yanyan Wei, Zhao Zhang, Jiahuan Ren, Xiaogang Xu, Richang Hong, Yi Yang, Shuicheng Yan, and Meng Wang. Clarity chatgpt: An interactive and adaptive processing system for image restoration and enhancement. *arXiv preprint arXiv:2311.11695*, 2023. 3
- [66] Patrick Wiescholke, Michael Hirsch, Bernhard Scholkopf, and Hendrik Lensch. Learning blind motion deblurring. In *Proceedings of the IEEE international conference on computer vision*, pages 231–240, 2017. 1
- [67] Haoning Wu, Hanwei Zhu, Zicheng Zhang, Erli Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Annan Wang, Wenxiu Sun, Qiong Yan, et al. Towards open-ended visual quality comparison. In *European Conference on Computer Vision*, pages 360–377. Springer, 2024. 4, 5
- [68] Bin Xia, Yulun Zhang, Shiyin Wang, Yitong Wang, Xinglong Wu, Yapeng Tian, Wenming Yang, and Luc Van Gool. Diffir: Efficient diffusion model for image restoration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 13095–13105, 2023. 3
- [69] Jingkan Yang, Yuhao Dong, Shuai Liu, Bo Li, Ziyue Wang, Haoran Tan, Chencheng Jiang, Jiamu Kang, Yuanhan Zhang, Kaiyang Zhou, et al. Octopus: Embodied vision-language programmer from environmental feedback. In *European conference on computer vision*, pages 20–38. Springer, 2024. 3
- [70] Wenhan Yang, Robby T Tan, Jiashi Feng, Jiaying Liu, Zongming Guo, and Shuicheng Yan. Deep joint rain detection and removal from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1357–1366, 2017. 2, 5, 6
- [71] Fanghua Yu, Jinjin Gu, Zheyuan Li, Jinfan Hu, Xiangtao Kong, Xintao Wang, Jingwen He, Yu Qiao, and Chao Dong. Scaling up to excellence: Practicing model scaling for photo-realistic image restoration in the wild. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 25669–25680, 2024. 3
- [72] Ke Yu, Chao Dong, Liang Lin, and Chen Change Loy. Crafting a toolchain for image restoration by deep reinforcement learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2443–2452, 2018. 2, 3, 5
- [73] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5728–5739, 2022. 1, 2, 7

- [74] Kai Zhang, Wangmeng Zuo, Yunjin Chen, Deyu Meng, and Lei Zhang. Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing*, 26(7):3142–3155, 2017. [1](#), [2](#)
- [75] Kai Zhang, Wangmeng Zuo, and Lei Zhang. Ffdnet: Toward a fast and flexible solution for cnn-based image denoising. *IEEE Transactions on Image Processing*, 27(9):4608–4622, 2018. [1](#)
- [76] Kai Zhang, Jingyun Liang, Luc Van Gool, and Radu Timofte. Designing a practical degradation model for deep blind image super-resolution. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 4791–4800, 2021. [1](#)
- [77] Kai Zhang, Yawei Li, Jingyun Liang, Jiezhong Cao, Yulun Zhang, Hao Tang, Deng-Ping Fan, Radu Timofte, and Luc Van Gool. Practical blind image denoising via swin-conv-unet and data synthesis. *Machine Intelligence Research*, 20(6):822–836, 2023. [1](#)
- [78] Lei Zhang, Xiaolin Wu, Antoni Buades, and Xin Li. Color demosaicking by local directional interpolation and nonlocal adaptive thresholding. *Journal of Electronic imaging*, 20(2): 023016–023016, 2011. [6](#)
- [79] Yulun Zhang, Yapeng Tian, Yu Kong, Bineng Zhong, and Yun Fu. Residual dense network for image super-resolution. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2472–2481, 2018. [2](#)
- [80] Yingjie Zhou, Jiezhong Cao, Zicheng Zhang, Farong Wen, Yanwei Jiang, Jun Jia, Xiaohong Liu, Xiongkuo Min, and Guangtao Zhai. Q-agent: quality-driven chain-of-thought image restoration agent through robust multimodal large language model. *arXiv preprint arXiv:2504.07148*, 2025. [3](#)
- [81] Kaiwen Zhu, Jinjin Gu, Zhiyuan You, Yu Qiao, and Chao Dong. An intelligent agentic system for complex image restoration problems. *arXiv preprint arXiv:2410.17809*, 2024. [1](#), [2](#), [3](#), [5](#), [7](#)
- [82] Yushen Zuo, Qi Zheng, Mingyang Wu, Xinrui Jiang, Renjie Li, Jian Wang, Yide Zhang, Gengchen Mai, Lihong V Wang, James Zou, et al. 4kagent: agentic any image to 4k super-resolution. *arXiv preprint arXiv:2507.07105*, 2025. [3](#)