

Self-Evolving Agentic Image Restoration via Deliberate Planning and Intuitive Execution

Shuang Cui^{1,2}, Fan Ji¹, Guanglong Sun³, Yufei Guo⁴,
Xiongxin Tang^{1†}, Jiangmeng Li^{1†}, and Fanjiang Xu¹

¹ Institute of Software, Chinese Academy of Sciences, Beijing, China

² University of Chinese Academy of Sciences, Beijing, China

³ School of Life Sciences, Tsinghua University, Beijing, China

⁴ Intelligent Science & Technology Academy of CASIC, Beijing, China
{cuishuang21, jifan22, xiongxin, jiangmeng2019, fanjiang}@iscas.ac.cn,
sgl23@mails.tsinghua.edu.cn, yfguo@pku.edu.cn
†Corresponding authors

Abstract. Real-world image restoration (IR) remains challenging due to complex and coupled degradations. While recent agentic IR frameworks leverage Large Language Models for flexible tool planning, they face two critical limitations. First, from a search scheme perspective, excessive reliance on greedy strategies fails to balance exploration and exploitation. Second, existing agentic systems underutilize information, exhibiting episodic amnesia. To address these challenges, we propose **Self-Evolving Agentic Image Restoration (SEAR)**, which formulates restoration as a sequential decision-making problem. Inspired by the dual-process theory, SEAR comprises an Intuitive Executor and a Deliberate Planner, respectively following the fast-thinking *System 1* and slow-thinking *System 2* principles. The Deliberate Planner employs Pruning-Aware Monte Carlo Tree Search for long-horizon reasoning, utilizing a hybrid no-reference reward and a Multimodal Large Language Model (MLLM)-based tournament to prevent metric exploitation. Complementarily, the Intuitive Executor leverages a self-evolving episodic memory indexed by degradation-aware state fingerprints. This mechanism distills expensive search trajectories into adaptive expertise, overcoming episodic amnesia while progressively amortizing cold-start exploration costs through memory reuse. Extensive experiments on synthetic and real-world benchmarks demonstrate its strong perceptual and quantitative performance.

Keywords: Agentic Image Restoration · Dual-Process Theory · Test-Time Search · Monte Carlo Tree Search

1 Introduction

Image Restoration (IR) is a fundamental task in computer vision that focuses on recovering high-quality images from degraded observations [21, 36, 59, 60, 65]. Task-specific networks have achieved remarkable success in handling isolated




Fig. 1: Analysis of restoration limitations. (a) Static AiO models struggle with coupled degradations. (b) Restoration quality is highly sensitive to execution sequences and tools, where greedy decisions (e.g., over-smoothed) incur irreversible failure.

degradations such as denoising [6, 42], deblurring [12, 17, 59], and super-resolution [30, 38, 52]. However, in real-world scenarios, images exhibit complex and coupled degradations [22, 49]. All-in-One (AiO) IR frameworks [23, 28, 37, 45] attempt to handle multiple tasks within a unified model, yet these static designs struggle to generalize to unknown distortion combinations [21], as shown in Fig. 1 (a).

Recently, inspired by the advanced reasoning and planning capabilities of Large Language Models (LLM) [19, 43, 48], agentic restoration frameworks [7, 64, 66] have emerged as a powerful paradigm for complex image restoration. As a representative example, AgenticIR [64] distills statistical experience from tool sequence traversals into LLM prompts to guide the agent, thereby improving restoration consistency. Despite these advances, existing agentic approaches suffer from structural limitations that hinder performance in challenging scenarios.

From a search scheme perspective, current agents [64, 66] predominantly rely on greedy heuristics that lack long-horizon foresight and fail to balance exploration and exploitation in high-dimensional decision spaces. By prioritizing immediate rewards at each step (Fig. 2 (a)), these methods often converge to sub-optimal restoration trajectories [15, 58]. As shown in Fig. 1 (b), aggressively applying denoising to maximize intermediate cleanliness may irreversibly destroy fine structural details that are essential for subsequent super-resolution, resulting in over-smoothed outputs. Moreover, from the perspective of information utility, existing systems exhibit episodic amnesia. Current agents rely primarily on static prompt-level experience, while high-reward trajectories discovered through computationally expensive search are treated as transient and discarded upon completion of each image restoration. Consequently, knowledge acquired from previous exploration cannot be transferred across diverse cases [35, 62], forcing




Fig. 2: Comparison of restoration decision strategies. (a) Myopic greedy agents prioritize local gains, yielding suboptimal trajectories. (b) SEAR synergizes an Intuitive Executor with a Deliberate Planner for long-horizon trajectory optimization.

the agent to repeatedly resolve similar degradation patterns from scratch and increasing repeated search costs in practical deployments.

To address these challenges, we reformulate real-world image restoration as a long-horizon sequential decision-making problem, where restoration operators must be applied in a carefully coordinated order to preserve image structures. This perspective moves beyond greedy, step-wise decisions to global path-level reasoning. Inspired by dual-process theory [24, 41], we propose **Self-Evolving Agentic Image Restoration (SEAR)**. SEAR employs an Intuitive Executor and a Deliberate Planner to emulate the cognitive synergy between fast *System 1* execution and slow *System 2* reasoning. As shown in Fig. 2, SEAR replaces traditional step-wise heuristics with a self-evolving architecture that continuously refines restoration trajectories by dynamically accumulating experience.

Specifically, the Deliberate Planner performs hierarchical reasoning, leveraging LLM for macro-level task scheduling and Pruning-Aware Monte Carlo Tree Search (P-MCTS) for micro-level tool coordination, balancing exploration and exploitation. P-MCTS trajectories are evaluated via a hybrid no-reference reward and safeguarded by a Multimodal Large Language Model (MLLM)-based perceptual tournament to prevent adversarial metric exploitation and ensure human-aligned fidelity. Complementarily, the Intuitive Executor addresses insufficient information utilization, i.e., episodic amnesia in existing agentic restoration systems. It maintains a self-evolving episodic memory that distills high-reward trajectories into reusable execution strategies indexed by state fingerprints. When structurally similar degradation patterns recur, the Intuitive Executor retrieves and executes these trajectories instead of replanning from scratch, effectively reusing past exploration results. Through continuous self-evolution, SEAR transforms transient search experiences into adaptive knowledge, amortizes repeated planning costs for recurring degradation patterns, and gradually shifts restoration from search-driven reasoning to memory-guided execution. Main contributions are summarized as follows:

- We propose SEAR, a dual-process restoration agent that synergizes deliberate planning with intuitive execution to tackle complex image restoration.
- We introduce P-MCTS, leveraging MLLM-guided long-horizon rollouts to coordinate tool selection while balancing exploration and exploitation.
- We develop a self-evolving episodic memory that converts search experiences into retrievable, adaptive knowledge, resolving amnesia and reducing exploration overhead.
- Extensive experiments on synthetic and real-world datasets demonstrate that SEAR achieves state-of-the-art perceptual quality, validating a robust paradigm evolving from costly deliberation to memory-guided execution.

2 Related work

Image Restoration. Image Restoration (IR) transitions from task-specific architectures for denoising [6, 42], deblurring [12, 17, 59], deraining [8, 55], super-resolution [30, 38, 52] toward All-in-One (AiO) IR frameworks that handle multiple degradations within a unified model. Representative methods include AirNet [28], which employs contrastive learning for feature discrimination, as well as PromptIR [37] and InstructIR [11], which utilize degradation-aware contexts or human instructions to guide restoration. Advancing this paradigm, AutoDIR [23] integrates semantic-agnostic VLMs to automate degradation identification and steer structure-aware diffusion processes. Furthermore, DA-CLIP [32] and MiOIR [26] leverage pre-trained foundation model priors or sequential prompting to enhance robustness. Despite these advancements, AiO methods struggle to balance generalization and reconstruction. Their efficacy remains limited on out-of-distribution degradation combinations [13, 21] due to a lack of strategic reasoning for intricate real-world scenarios.

Autonomous Agents in Image Restoration. Large Language Models (LLM) [19, 43, 48] exhibit formidable logical reasoning and task-planning capabilities, driving a paradigm shift from static restoration models toward adaptive, agentic systems [14, 18, 51]. Foundational works such as RestoreAgent [7] and AgenticIR [64] utilize in-context learning for tool dispatching and visual-linguistic feedback for iterative calibration. Multi-agent frameworks such as 4KAgent [66], MAIR [22], and HybridAgent [27] address coupled degradations through specialized sub-agents or expert integration. While training-intensive methods like JarvisIR [31] enhance MLLM via supervision or alignment, they require substantial training overhead. In contrast, SEAR follows a reasoning-centric paradigm without additional training.

Despite their adaptability, existing agentic frameworks [22, 27, 64, 66] are often constrained by myopic execution strategies that fail to balance exploration and exploitation. Furthermore, these methods suffer from diminished information utility, as they treat successful reasoning trajectories as ephemeral context rather than persistent assets. In the absence of long-horizon foresight and structural memory, such agents are trapped in a non-evolving loop, redundantly resolving recurring degradation patterns from scratch. SEAR addresses these gaps




Fig. 3: Overview of the SEAR Framework. SEAR follows a dual-process design: the Intuitive Executor performs rapid, memory-driven inference by retrieving proven strategies, while the Deliberate Planner tackles complex scenarios via long-horizon hierarchical optimization. By progressively distilling planned trajectories into the Executor, SEAR evolves from costly deliberation to memory-guided execution.

by combining hierarchical planning with a self-evolving episodic memory, systematically converting successful trajectories into reusable, long-term expertise.

3 Methods

We propose SEAR, a dual-process architecture for agentic image restoration (Fig. 3). To overcome the limitations of myopic and memory-less agents, SEAR is decoupled into an Intuitive Executor (*System 1*) for rapid, memory-driven inference, and a Deliberate Planner (*System 2*) for slow, long-horizon exploration. A self-evolving memory bank bridges the two modules by consolidating high-reward trajectories into persistent expertise, as detailed in Appendix D.

3.1 Problem Formulation

State and Action Space. Real-world image restoration is inherently order-sensitive (e.g., denoising before deblurring leads to different structural outcomes). Let $\mathcal{A} = \{f_1, \dots, f_K\}$ denote the restoration tool library and I_0 the initial degraded image. A restoration trajectory is defined as an ordered sequence $\tau = (a_0, \dots, a_{H-1})$, where each $a_t \in \mathcal{A}$, H denotes the trajectory horizon (i.e., the number of sequential operations), and $I_{t+1} = a_t(I_t)$ represents the state transition after applying a_t . By explicitly modeling the full trajectory, this formulation accounts for delayed effects across steps, ensuring that early decisions do not irreversibly compromise structures required by subsequent operations.

No-Referenced Reward Objective. Since paired ground truths are unavailable in real-world inference, evaluating long-horizon trajectories relies solely on blind quality assessment. To prevent metric overfitting that can produce misleading artifacts, we guide the search with a robust perceptual ensemble. The optimal trajectory τ^* maximizes a terminal hybrid reward $R(I_H)$, defined as a weighted combination of human-aligned aesthetics [54] and no-reference IQA metrics:

$$\tau^* = \arg \max_{\tau} R(I_H), \quad R(I_H) = \lambda_h \text{HPSv2}(I_H) + \sum_{j \in \mathcal{J}} \lambda_j \widehat{\text{IQA}}_j(I_H), \quad (1)$$

where I_H denotes the terminal restored image of trajectory τ . The IQA ensemble includes NIQE [34], MANIQA [56], MUSIQ [25], and CLIP-IQA [46]. Each IQA metric $\widehat{\text{IQA}}_j$ is normalized to zero mean and unit variance. The specific weight configurations for the reward ensemble follow the settings established in [66]. We denote $R(\tau) \equiv R(I_H(\tau))$ for the reward of trajectory τ . This approach ensures perceptual consistency over single-metric optimization (Sec. 4.4).

3.2 Intuitive Executor with Self-Evolving Memory

Unlike prior agents that suffer episodic amnesia by discarding trajectories post-inference, the Intuitive Executor leverages a self-evolving memory bank $\mathcal{M}$. To support adaptive execution, $\mathcal{M}$ is formalized as a repository of expertise tuples $\langle \mathbf{f}_i, \tau_i, \mathcal{K}_i \rangle$, where each state-trajectory pair is linked to a statistical profile $\mathcal{K}_i = \{n, \mu, R_{\max}\}$ tracking its visit count, mean reward, historical peak reward.

Degradation-Aware State Fingerprint. To facilitate efficient and content-invariant memory addressing, we first project the high-dimensional degraded image I_0 into a structured state fingerprint $\mathbf{f}$. A fine-tuned MLLM [57, 64] diagnoses the input I_0 to extract a semantic distortion profile C_I , which is quantized into a state vector $\mathbf{f} = [l_1, \dots, l_{|\mathcal{D}|}]^\top$. Here, $|\mathcal{D}| = 8$ spans a predefined taxonomy of degradation primitives (e.g., blur, noise, compression artifacts), with each element $l_i \in 0, \dots, 5$ encoding discrete intensity levels. This coarse-grained quantization acts as a regularization mechanism. By preventing the memory bank from overfitting to negligible pixel-level variances, it ensures that structurally identical degradation patterns are robustly mapped to the same execution strategy, thereby providing a low-dimensional key for efficient retrieval.

Potential-Driven Retrieval. Given the query fingerprint $\mathbf{f}_q$, the Intuitive Executor retrieves $\mathcal{M}$ to identify a candidate set $\mathcal{S}_{\text{cand}}$ composed of entries within an L2 proximity threshold. To balance prior performance with state similarity, the agent ranks these candidates using a potential-aware score:

$$\mathcal{C}(\tau_i) = R_{\max}(\tau_i) - \alpha \|\mathbf{f}_q - \mathbf{f}_i\|_2, \quad (2)$$

where $R_{\max}(\tau_i)$ is the historical peak reward of τ_i , and α penalizes state discrepancy. By focusing on $R_{\max}$ rather than the mean reward μ , the agent prioritizes trajectories with the highest demonstrated restorative potential, thereby reducing sensitivity to noisy IQA estimates and suboptimal historical averages.

Reliability-Gated Adaptive Execution. To preserve restoration fidelity, we introduce an outcome-driven verification gate for retrieved candidate trajectories. The top-ranked retrieved trajectory τ_{ret} is first executed to generate a candidate restoration $I_H(\tau_{\text{ret}})$, which is then evaluated by the hybrid reward function in Eq. 1. If the reward meets the shortcut threshold ζ , the agent accepts this trajectory and outputs $I_H(\tau_{\text{ret}})$ as the final result. Otherwise, the retrieved trajectory is used only as a warm start for the Deliberate Planner, which further performs P-MCTS search. This gate allows shortcut execution only for reliable retrieved trajectories, reducing unnecessary planning while preserving a quality-oriented fallback mechanism.

3.3 Deliberate Planner for Long-Horizon Reasoning

The Deliberate Planner is invoked when the Intuitive Executor lacks reliable precedents. To handle the combinatorial action space, it employs a hierarchical planning strategy, separating macro-level agenda reasoning from micro-level tool instantiation.

Macro-Planning: Agenda-Driven Space Pruning. Guided by the distortion profile C_I and prior restoration experience [64], an LLM generates a strategic agenda $\mathcal{T} = [T_1, \dots, T_L]$. This ordered task sequence encodes intrinsic restoration priorities (e.g., denoising before super-resolution) and enforces precedence constraints, reducing the exponential search space of joint tool selection into a stage-wise optimization sequence for subsequent tree search.

Micro-Planning: Pruning-Aware Monte Carlo Tree Search. To instantiate the strategic agenda into a concrete tool sequence, the Deliberate Planner employs Pruning-Aware Monte Carlo Tree Search (P-MCTS). This framework shifts the optimization target from isolated tool decisions to long-horizon trajectory optimization. In the search tree, each node s represents an intermediate restoration state, while each edge signifies the invocation of an expert tool $a \in \text{Tools}(T_i)$. P-MCTS identifies high-reward trajectories through an iterative process of Selection, Expansion, Simulation, and Backpropagation. The search runs for $N_{\text{total}} = \omega \cdot L$ iterations, where L is the agenda length and ω is the search budget multiplier. In practice, this process is significantly accelerated by a caching mechanism that enables the rapid reuse of previously explored results. This cycle, detailed below, balances exploration and exploitation to converge toward globally optimal restoration outcomes.

Selection: Starting from the root, the agent traverses the tree by selecting the child node s_{child} that maximizes the Upper Confidence Bound for Trees (UCT):

$$UCT(s_{\text{child}}) = \frac{Q(s_{\text{child}})}{N(s_{\text{child}})} + c \cdot \sqrt{\frac{2 \ln N(s_{\text{parent}})}{N(s_{\text{child}})}}, \quad (3)$$

where Q represents the accumulated reward, N denotes the visit count, and c is the exploration parameter. This strategy prioritizes successful paths while ensuring sufficient coverage of the decision space.

Expansion: Upon reaching a node s_n that is not yet fully expanded for the current subtask T_i , the agent expands the tree by sampling an operator $a \in \text{Tools}(T_i)$ to generate a child node s_{child} . This agenda-constrained expansion restricts the branching factor to operators associated with the current subtask, enforcing the strategic topology defined by the agenda. As a result, the search space is substantially pruned compared to unconstrained tool selection.

Simulation: To estimate the downstream impact of current decisions, the agent performs a look-ahead rollout by sequentially applying tools from the remaining subtasks $[T_{i+1}, \dots, T_L]$ to reach a terminal state s_T . The resulting terminal image I_H is then evaluated via the hybrid reward $R(I_H)$ as defined in Eq. 1.

Backpropagation: The terminal reward $R(I_H)$ is backpropagated to update the statistics of all ancestral nodes s_j along the trajectory:

$$N(s_j) \leftarrow N(s_j) + 1, Q(s_j) \leftarrow Q(s_j) + R(I_H), V_{\max}(s_j) \leftarrow \max(V_{\max}(s_j), R(I_H)). \quad (4)$$

This iterative refinement steers the selection strategy toward the high-reward sequence of expert tools. By tracking the maximum value $V_{\max}$, the framework preserves the highest potential of explored paths to guide future search iterations.

Perceptual Alignment via MLLM Tournament. To mitigate the risk of metric exploitation, where restoration operators may generate adversarial artifacts yielding deceptively high scores, we implement a pairwise tournament verification. The top four candidate trajectories $\{\tau^{(1)}, \dots, \tau^{(4)}\}$, ranked according to their terminal rewards $R(\tau^{(i)})$, undergo pairwise visual assessments by an MLLM judge. Acting as a perceptual safety net, this final check goes beyond the hybrid reward. Even if the fixed weights of Eq. 1 occasionally misalign with specific out-of-distribution degradations, the tournament selection ensures that the ultimate output I^* is anchored in human-aligned perceptual fidelity rather than spurious metric maximization.

3.4 Dual-System Synergy and Memory Evolution

When P-MCTS yields a successful trajectory τ^* ($R(I_H) \geq \eta$), its state fingerprint is evaluated. Novel trajectories are newly incorporated into $\mathcal{M}$, while familiar ones update existing statistical profiles $\mathcal{K}_i$ (i.e., n , μ , and $R_{\max}$). To maintain efficiency, pruning is triggered when $\mathcal{M}$ exceeds its capacity $M_{\max}$. We formally execute this by evicting obsolete entries with the lowest visit counts n (Least Frequently Used policy) and sub-optimal trajectories whose peak reward $R_{\max}$ falls below the memory’s population mean. As $\mathcal{M}$ expands, the agent encounters familiar states more frequently, allowing more samples to be handled by verified memory shortcuts rather than full P-MCTS. This dynamic progressively reduces repeated planning for recurring degradation patterns. Ultimately, SEAR transforms expensive search trajectories into reusable expertise, providing a quality-efficiency tradeoff by amortizing search costs through memory reuse in long-term deployment.

4 Experiments

4.1 Experimental Setups

Implementation Details. Following [64], we utilize pre-trained X-Restormer [9], HAT [10], FBCNN [20], SwinIR [29], DiffBIR [30], MPRNet [33], DRB-Net [39], DehazeFormer [40], MAXIM [44], RIDCP [53], Restormer [59], and traditional operators. To facilitate image degradation perception, we utilize a fine-tuned DepictQA [57] as the unified MLLM evaluator to assess visual quality and distortion types. For macro-task planning, GPT-4o [2] serves as the core LLM engine to orchestrate high-level restoration strategies. Key parameters are $c = 0.8, \eta = 0.55, \zeta = 0.7, \omega = 9$, and distance penalty $\alpha = 0.05$. Experiments use NVIDIA V100 GPUs (details in Appendix A).

Datasets. The synthetic benchmarks (Groups A, B, and C) [64] consist of 1,440 images derived from MiO100 [26] featuring 16 distinct combinations of mixed distortions. To assess generalization, the real-world dataset aggregates 100 pairs from I-Haze [4], NH-Haze [3], DRealSR [50], RealSR [5], T-OLED [63], SIDD [1], and LHP-Rain [16], which will be made publicly available.

Baselines. We evaluate SEAR against six AiO models (AirNet [28], PromptIR [37], MiOIR [26], DA-CLIP [32], InstructIR [11], AutoDIR [23]) and two agentic frameworks (AgenticIR [64], 4KAgent [66]). To ensure strategic parity and reproducibility, we unify toolsets and re-evaluate agentic baselines using the publicly released DepictQA module from AgenticIR [64], serving as a standardized proxy for their unreleased internal variant.

Metrics. Reference-based evaluation utilizes PSNR and SSIM [47] computed on the luminance (Y) channel following the MiO100 protocol [26]. To ensure consistent comparison, all baseline methods are re-evaluated under this unified Y-channel setting. Perceptual similarity is assessed via LPIPS [61]. To evaluate restoration quality without ground truth, we report no-reference scores including MANIQA [56], CLIP-IQA [46], and MUSIQ [25] to quantify perceptual.

4.2 Comparison with State-of-the-Arts

Quantitative Comparison. Table 1 shows that SEAR consistently outperforms AiO and agentic baselines across nearly all benchmarks. Notably, it achieves PSNR gains of 0.42 dB and 0.30 dB in Groups B and C, where out-of-domain patterns elude memory-less agents [64, 66]. SEAR also leads in perceptual metrics (LPIPS and MUSIQ), underscoring high-fidelity reconstruction in unseen scenarios. Such consistency confirms that SEAR’s self-evolving architecture bridges the gap between fixed heuristics and dynamic restoration requirements.

Qualitative Comparisons. Fig. 4 presents qualitative comparisons on the synthetic MiO100 dataset. AiO models struggle in mixed degradation scenarios, frequently failing to reconcile conflicting distortions and introducing visible artifacts. Although AgenticIR produces comparatively cleaner outputs, its greedy search strategy often results in over-smoothing, leading to the loss of characteristic structures, such as the sand ripples in Group C. 4KAgent further suffers

Table 1: Quantitative comparison of multiple-degradation image restoration tasks on three subsets (Group A, B, and C) from the MiO100 dataset. The best and second best results are highlighted in **bold** and underline, respectively.

Datasets	Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	MANIQA $\uparrow$	CLIP-IQA $\uparrow$	MUSIQ $\uparrow$
Group A	AirNet [28]	20.8902	0.6675	0.4293	0.2650	0.3938	42.5047
	PromptIR [37]	21.4075	0.6701	0.4254	0.2693	0.3957	42.8456
	MiOIR [26]	21.3563	0.6834	0.4074	0.2626	0.4120	43.8445
	DA-CLIP [32]	20.7522	0.6512	0.4225	0.2703	0.4155	44.6804
	InstructIR [11]	20.2364	0.6409	0.4626	0.2630	0.3781	42.9907
	AutoDIR [23]	20.4909	0.6543	0.4138	0.2872	0.4172	48.9844
	AgenticIR [64]	21.5321	0.6972	0.3110	0.3303	0.4842	59.8602
	4KAgent [66]	<u>21.6468</u>	0.6936	<u>0.3082</u>	<u>0.3403</u>	<u>0.5017</u>	<u>61.1519</u>
	SEAR	21.8042	<u>0.6961</u>	0.3019	0.3411	0.5039	61.2882
Group B	AirNet [28]	20.1531	0.7015	0.3753	0.2946	0.4184	47.4868
	PromptIR [37]	20.3621	0.7146	0.3550	0.2961	0.4208	47.5886
	MiOIR [26]	20.6956	<u>0.7244</u>	0.3477	0.2830	0.4325	48.5851
	DA-CLIP [32]	20.1534	0.6741	0.3968	0.2873	0.4111	48.4284
	InstructIR [11]	20.8317	0.6761	0.4133	0.3016	0.4006	48.9953
	AutoDIR [23]	20.9553	0.6897	0.3837	0.2925	0.4204	49.0215
	AgenticIR [64]	<u>21.7160</u>	0.7236	0.3025	0.3311	0.4871	59.5741
	4KAgent [66]	21.5416	0.7175	<u>0.2958</u>	<u>0.3434</u>	0.5226	61.4795
	SEAR	22.1332	0.7251	0.2890	0.3447	<u>0.5134</u>	<u>60.6520</u>
Group C	AirNet [28]	20.5462	0.5983	0.5618	0.2026	0.3221	29.9750
	PromptIR [37]	<u>20.6046</u>	<u>0.6012</u>	0.5584	0.2059	0.3189	29.8850
	MiOIR [26]	20.5308	0.6015	0.5725	0.1962	0.3171	30.1058
	DA-CLIP [32]	19.5827	0.5683	0.5948	0.1921	0.3330	29.5910
	InstructIR [11]	19.4498	0.5781	0.5756	0.1890	0.3103	29.2901
	AutoDIR [23]	19.1069	0.5716	0.5677	0.2113	0.3226	35.7965
	AgenticIR [64]	20.1974	0.5875	0.4341	0.2719	0.4325	51.7533
	4KAgent [66]	20.3078	0.5798	<u>0.4223</u>	0.2929	0.4737	<u>52.6504</u>
	SEAR	20.6121	0.5927	0.4173	<u>0.2742</u>	<u>0.4420</u>	52.7338

from noticeable residual noise and perceptual distortions. In contrast, SEAR consistently preserves fine-grained details across all categories, accurately recovering brick mortar, lizard scales, and intricate dune textures, while effectively suppressing noise and removing rain. Additional results are in Appendix B.

Efficiency Analysis. Table 2 evaluates efficiency using average inference time (min) and tool calls per image. SEAR balances fidelity and efficiency by amortizing planning costs through episodic memory. Compared to AgenticIR, SEAR improves PSNR by 0.42 dB with limited runtime and tool overhead. SEAR also outperforms 4KAgent in restoration quality and efficiency, reducing execution time by 22.4% with comparable tool usage.

Table 2: Efficiency analysis on the Group B dataset [64].

Method	PSNR	Time	Tool Calls
AgenticIR	21.72	1.09	6.11
4KAgent	21.54	2.55	8.26
SEAR	22.13	1.98	8.15





Fig. 4: Qualitative comparison on synthetic datasets [64].

Table 3: Quantitative comparison on the real-world dataset.

Method	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	MANIQA $\uparrow$	CLIP-IQA $\uparrow$	MUSIQ $\uparrow$
AirNet [28]	23.3067	0.7471	0.4484	0.2356	0.3272	35.2690
PromptIR [37]	23.8647	<u>0.7542</u>	0.4341	0.2405	0.3270	35.4811
MiOIR [26]	24.0043	0.7472	0.4239	0.2394	0.3507	36.8458
DA-CLIP [32]	23.7313	0.7437	0.4273	0.2514	0.3330	36.4113
InstructIR [11]	26.1477	0.7639	0.4237	0.2556	0.3412	37.8236
AutoDIR [23]	20.7470	0.6640	0.4792	0.2466	0.3148	42.6789
AgenticIR [64]	23.9280	0.7214	0.3730	0.3151	0.4512	52.9103
4KAgent [66]	23.6295	0.7242	<u>0.3569</u>	0.3200	<u>0.4513</u>	<u>54.4647</u>
SEAR	<u>24.4078</u>	0.7425	0.3371	<u>0.3174</u>	0.4519	54.6686

4.3 Real-World Generalization

Evaluation on the real-world dataset demonstrates the robustness of SEAR in real-world scenarios. As shown in Table 3, SEAR attains leading performance across perceptual metrics, including LPIPS and CLIP-IQA. Despite superior PSNR, pixel-fidelity models like InstructIR frequently produce oversmoothed

**Fig. 5:** Qualitative comparison on the real-world dataset.**Table 4:** Ablation study on the Group B dataset. “Tool Calls” denotes the average number of restoration tool invocations per image, reflecting inference efficiency.

Configuration	PSNR $\uparrow$	SSIM $\uparrow$	LPIPS $\downarrow$	MANIQA $\uparrow$	CLIP-IQA $\uparrow$	MUSIQ $\uparrow$	Tool Calls $\downarrow$
a. SEAR (Full)	22.1332	0.7251	0.2890	0.3447	0.5134	60.6520	8.15
b. w/o Memory	22.0222	0.7261	0.2903	0.3492	0.5124	61.0802	16.75
c. w/o P-MCTS	20.8395	0.7156	0.3046	0.3335	0.5068	60.3226	4.08
d. w/o Macro-Task	19.5378	0.6367	0.3728	0.3926	0.5488	64.4798	71.76
e. Vanilla MCTS	22.1252	0.7233	0.2893	0.3480	0.5128	60.8131	9.93
f. w/o Fingerprint	21.0774	0.6848	0.3335	0.3458	0.4885	59.9639	11.49

outputs with limited perceptual detail. Visual results in Fig. 5 confirm these observations. Particularly in “Haze+Blur” and “Noise+Blur” scenarios, SEAR restores natural textures and effectively suppresses artifacts, whereas baselines exhibit residual distortions or severe color casts. By optimizing restoration trajectories, SEAR handles complex real-world degradations with optimal perceptual quality, solving mixed patterns that challenge static models.

4.4 Ablation Study

Synergy of Dual-Process Systems. Table 4 (Configs. b, c) confirms that SEAR’s dual-process architecture is vital for balancing fidelity and efficiency. Removing the memory bank forces the agent to search without dynamic memory guidance, doubling tool calls from 8.15 to 16.75, demonstrating that memory-guided inference is essential for efficient execution. Fig. 6 illustrates the dynamic self-evolution of SEAR. While the cold-start SEAR (red) initially behaves like the deliberative search baseline (blue), it quickly distills planning into efficient




Fig. 6: Dynamic self-evolution.



Fig. 7: Ablation study of P-MCTS.




Fig. 8: Reward design ablation. (a) Radar comparison of reward components. (b) Visual analysis: tournament selection prioritizes perceptual quality over inflated scores.

execution. Cross-domain initialization (yellow) further accelerates convergence, demonstrating the transferability of distilled trajectories, while the eventual convergence of both settings highlights the robustness of the self-evolving architecture. Conversely, replacing P-MCTS with greedy heuristics reduces tool calls to 4.08 but causes a 1.29 dB drop in PSNR. As shown in Fig. 7, greedy heuristics like premature sharpening trigger irreversible artifact accumulation.

Memory Routing Analysis. To further verify that SEAR does not blindly reuse retrieved trajectories, we report memory routing statistics under cold-start deployment in Table 5. For each input, the Intuitive Executor first retrieves candidate trajectories according to degradation-aware fingerprints. The retrieved trajectory is executed only when it passes the reliability gate; otherwise, the sample is routed to the Deliberate Planner for P-MCTS search. As shown in Table 5, SEAR achieves high retrieval rates on both synthetic and real-world benchmarks, while a substantial portion of uncertain cases is still routed to System 2. This execute-then-gate mechanism confirms that the memory bank serves as a verified shortcut rather than an unconditional cache.

Efficiency of Hierarchical Planning. Table 4 (Configs. d and e) validates the necessity of hierarchical planning. Removing the macro-task forces a flat search across the entire tool set, increasing tool calls to 71.76 while degrading PSNR to 19.54 dB, clearly showing the efficiency and effectiveness benefits of hierarchical constraints. Furthermore, reverting to vanilla MCTS marginally degrades PSNR

Table 5: Memory routing statistics under cold-start deployment.

Dataset	Memory Hit	Accepted Shortcut	Routed to System 2
Group A	97%	67%	33%
Group B	98%	62%	38%
Group C	94%	43%	57%
Real-world	83%	41%	59%

Table 6: Hyperparameter sensitivity analysis on the real-world dataset.

c	PSNR	SSIM		PSNR	SSIM	ω	PSNR	SSIM
0.4	24.16	0.7373	$\zeta = 0.6$	23.89	0.7320	7	24.18	0.7349
0.6	24.08	0.7257	$\zeta = 0.7$	24.41	0.7425	8	24.39	0.7540
0.8	24.41	0.7425	$\zeta = 0.8$	24.20	0.7424	9	24.41	0.7425
1.0	24.26	0.7396	$\eta = 0.45$	24.05	0.7270	10	24.31	0.7456
1.2	24.14	0.7376	$\eta = 0.55$	24.41	0.7425	11	24.16	0.7318
			$\eta = 0.65$	24.20	0.7389	12	24.28	0.7360

while inflating tool invocations from 8.15 to 9.93, demonstrating that MLLM-guided action pruning effectively curtails redundant exploration.

Importance of State Fingerprint. Table 4 (Config. f) shows that replacing the semantic fingerprint with IQA-only descriptors reduces PSNR by 1.05 dB, indicating that simple quality metrics cannot capture heterogeneous degradations. Therefore, MLLM-based state abstraction is essential for accurate retrieval.

Perceptual Alignment Analysis. Fig. 8 shows that ablating terminal hybrid reward $R(I_H)$ components evaluates the trade-off between pixel accuracy and perceptual quality. Using the IQA-only reward preserves unsupervised metrics (e.g., NIQE, MUSIQ) but ignores aesthetic fidelity, whereas relying solely on HPSv2 favors human-aligned aesthetics at the cost of structural accuracy. Removing either component consistently degrades both quantitative and qualitative performance, confirming that the combined hybrid reward best guides restoration. To prevent trajectories that achieve high reward scores but produce visually unsatisfactory results, we apply the MLLM Tournament as a final perceptual filter. This ensures that SEAR prioritizes both metric consistency and human-aligned visual quality, achieving robust coverage across all dimensions.

Hyperparameter Sensitivity. Table 6 confirms SEAR’s structural robustness across core parameters. Specifically, the exploration constant $c = 0.8$ provides a favorable balance between search diversity and convergence, while the distillation threshold $\eta = 0.55$ effectively gates memory quality against the rate of expertise accumulation. Furthermore, the shortcut threshold $\zeta = 0.7$ supports reliable shortcut execution by bypassing deliberation for recurring patterns without performance degradation. $\omega = 9$ represents the optimal search budget multiplier in P-MCTS, as it effectively saturates restoration performance without redundant computation. Detailed tuning strategies are provided in Appendix A.

5 Conclusion

To address the greedy planning and episodic amnesia in current restoration agents, we proposed SEAR, a dual-process framework that synergizes deliberate planning with intuitive execution. By coupling P-MCTS with self-evolving episodic memory, SEAR enables long-horizon trajectory optimization while distilling expensive search experiences into adaptive, reusable expertise. This architecture allows the system to progressively shift from resource-intensive deliberation to efficient, memory-driven inference as experience accumulates. Experiments on synthetic and real-world benchmarks demonstrate that SEAR achieves competitive restoration quality, pointing to a practical pathway toward self-evolving agentic restoration systems in complex real-world environments.

Limitations. Although SEAR achieves strong restoration performance, it still has two limitations. First, SEAR mainly handles unseen degradation combinations and severities covered by a predefined restoration tool library, and its memory retrieval depends on degradation diagnosis for state fingerprint construction. Inaccurate diagnosis or missing restoration tools may affect planning under severe or out-of-distribution degradations. Second, the initial P-MCTS stage introduces cold-start overhead before sufficient high-quality trajectories are accumulated in memory. Future work will explore broader tool libraries, more discriminative fingerprints, and stronger perception-aligned evaluation for more diverse restoration scenarios.

Acknowledgements

This work is supported by the National Natural Science Foundation of China under Grant No. 62406313, the Postdoctoral Fellowship Program of the China Postdoctoral Science Foundation under Grant No. YJB20250283, and the National Key Laboratory of Space Integrated Information System under Grant No. E5GF5910.

References

1. Abdelhamed, A., Lin, S., Brown, M.S.: A high-quality denoising dataset for smart-phone cameras. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 1692–1700 (2018)
2. Achiam, J., Adler, S., Agarwal, S., Ahmad, L., Akkaya, I., Aleman, F.L., Almeida, D., Altenschmidt, J., Altman, S., Anadkat, S., et al.: Gpt-4 technical report. arXiv preprint arXiv:2303.08774 (2023)
3. Ancuti, C.O., Ancuti, C., Timofte, R.: Nh-haze: An image dehazing benchmark with non-homogeneous hazy and haze-free images. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops. pp. 444–445 (2020)

4. Ancuti, C., Ancuti, C.O., Timofte, R., De Vleeschouwer, C.: I-haze: A dehazing benchmark with real hazy and haze-free indoor images. In: International conference on advanced concepts for intelligent vision systems. pp. 620–631. Springer (2018)
5. Cai, J., Zeng, H., Yong, H., Cao, Z., Zhang, L.: Toward real-world single image super-resolution: A new benchmark and a new model. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 3086–3095 (2019)
6. Chen, H., Gu, J., Liu, Y., Magid, S.A., Dong, C., Wang, Q., Pfister, H., Zhu, L.: Masked image training for generalizable deep image denoising. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1692–1703 (2023)
7. Chen, H., Li, W., Gu, J., Ren, J., Chen, S., Ye, T., Pei, R., Zhou, K., Song, F., Zhu, L.: Restoreagent: Autonomous image restoration agent via multimodal large language models. *Advances in Neural Information Processing Systems* **37**, 110643–110666 (2024)
8. Chen, X., Li, H., Li, M., Pan, J.: Learning a sparse transformer network for effective image deraining. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5896–5905 (2023)
9. Chen, X., Li, Z., Pu, Y., Liu, Y., Zhou, J., Qiao, Y., Dong, C.: A comparative study of image restoration networks for general backbone network design. In: European Conference on Computer Vision. pp. 74–91. Springer (2024)
10. Chen, X., Wang, X., Zhang, W., Kong, X., Qiao, Y., Zhou, J., Dong, C.: Hat: Hybrid attention transformer for image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence* pp. 2676 – 2694 (2025)
11. Conde, M.V., Geigle, G., Timofte, R.: Instructir: High-quality image restoration following human instructions. In: European Conference on Computer Vision. pp. 1–21. Springer (2024)
12. Cui, S., Li, Y., Li, J., Tang, X., Su, B., Xu, F., Xiong, H.: Continual test-time adaptation for single image defocus deblurring via causal siamese networks. *International Journal of Computer Vision* **133**(7), 4134–4157 (2025)
13. Duan, J., Ai, Y., Liu, J., Huang, S., Huang, H., Cao, J., He, R.: Test-time forgery detection with spatial-frequency prompt learning. *International Journal of Computer Vision* **133**(2), 672–687 (2025)
14. Durante, Z., Huang, Q., Wake, N., Gong, R., Park, J.S., Sarkar, B., Taori, R., Noda, Y., Terzopoulos, D., Choi, Y., et al.: Agent ai: Surveying the horizons of multimodal interaction. *arXiv preprint arXiv:2401.03568* (2024)
15. Gao, N., Zhang, X., Jiang, X., You, M., Zhang, M., Deng, Y.: Rf-agent: Automated reward function design via language agent tree search. In: The Thirty-ninth Annual Conference on Neural Information Processing Systems
16. Guo, Y., Xiao, X., Chang, Y., Deng, S., Yan, L.: From sky to the ground: A large-scale benchmark and simple baseline towards real rain removal. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 12097–12107 (2023)
17. He, J.T., Tsai, F.J., Peng, Y.T., Chen, M.H., Lin, C.W., Lin, Y.Y.: Blurdm: A blur diffusion model for image deblurring. In: *Advances in Neural Information Processing Systems* (2025)
18. Hong, S., Zhuge, M., Chen, J., Zheng, X., Cheng, Y., Wang, J., Zhang, C., Wang, Z., Yau, S.K.S., Lin, Z., et al.: Metagpt: Meta programming for a multi-agent collaborative framework. In: The Twelfth International Conference on Learning Representations (2023)

19. Hui, B., Yang, J., Cui, Z., Yang, J., Liu, D., Zhang, L., Liu, T., Zhang, J., Yu, B., Lu, K., et al.: Qwen2. 5-coder technical report. arXiv preprint arXiv:2409.12186 (2024)
20. Jiang, J., Zhang, K., Timofte, R.: Towards flexible blind jpeg artifacts removal. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 4997–5006 (2021)
21. Jiang, J., Zuo, Z., Wu, G., Jiang, K., Liu, X.: A survey on all-in-one image restoration: Taxonomy, evaluation and future trends. IEEE Transactions on Pattern Analysis and Machine Intelligence (2025)
22. Jiang, X., Li, G., Chen, B., Zhang, J.: Multi-agent image restoration. arXiv preprint arXiv:2503.09403 (2025)
23. Jiang, Y., Zhang, Z., Xue, T., Gu, J.: Autodir: Automatic all-in-one image restoration with latent diffusion. In: European Conference on Computer Vision. pp. 340–359. Springer (2024)
24. Kahneman, D.: Thinking, fast and slow. macmillan (2011)
25. Ke, J., Wang, Q., Wang, Y., Milanfar, P., Yang, F.: Musiq: Multi-scale image quality transformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 5148–5157 (2021)
26. Kong, X., Dong, C., Zhang, L.: Towards effective multiple-in-one image restoration: A sequential and prompt learning strategy. arXiv preprint arXiv:2401.03379 (2024)
27. Li, B., Li, X., Lu, Y., Chen, Z.: Hybrid agents for image restoration. arXiv preprint arXiv:2503.10120 (2025)
28. Li, B., Liu, X., Hu, P., Wu, Z., Lv, J., Peng, X.: All-in-one image restoration for unknown corruption. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 17452–17462 (2022)
29. Liang, J., Cao, J., Sun, G., Zhang, K., Van Gool, L., Timofte, R.: Swinir: Image restoration using swin transformer. In: Proceedings of the IEEE/CVF international conference on computer vision. pp. 1833–1844 (2021)
30. Lin, X., He, J., Chen, Z., Lyu, Z., Dai, B., Yu, F., Qiao, Y., Ouyang, W., Dong, C.: Diffbir: Toward blind image restoration with generative diffusion prior. In: European conference on computer vision. pp. 430–448. Springer (2024)
31. Lin, Y., Lin, Z., Chen, H., Pan, P., Li, C., Chen, S., Wen, K., Jin, Y., Li, W., Ding, X.: Jarvisir: Elevating autonomous driving perception with intelligent image restoration. In: Proceedings of the Computer Vision and Pattern Recognition Conference. pp. 22369–22380 (2025)
32. Luo, Z., Gustafsson, F.K., Zhao, Z., Sjölund, J., Schön, T.B.: Controlling vision-language models for multi-task image restoration. In: International Conference on Learning Representations (2024)
33. Mehri, A., Ardakani, P.B., Sappa, A.D.: Mprnet: Multi-path residual network for lightweight image super resolution. In: Proceedings of the IEEE/CVF winter conference on applications of computer vision. pp. 2704–2713 (2021)
34. Mittal, A., Soundararajan, R., Bovik, A.C.: Making a “completely blind” image quality analyzer. IEEE Signal processing letters **20**(3), 209–212 (2012)
35. Park, J.S., O’Brien, J., Cai, C.J., Morris, M.R., Liang, P., Bernstein, M.S.: Generative agents: Interactive simulacra of human behavior. In: Proceedings of the 36th annual acm symposium on user interface software and technology. pp. 1–22 (2023)
36. Pei, Y., Huang, Y., Zou, Q., Lu, Y., Wang, S.: Does haze removal help cnn-based image classification? In: Proceedings of the European conference on computer vision (ECCV). pp. 682–697 (2018)

37. Potlapalli, V., Zamir, S.W., Khan, S.H., Shahbaz Khan, F.: Promptir: Prompting for all-in-one image restoration. *Advances in Neural Information Processing Systems* **36**, 71275–71293 (2023)
38. Qu, Y., Yuan, K., Hao, J., Zhao, K., Xie, Q., Sun, M., Zhou, C.: Visual autoregressive modeling for image super-resolution. *Proceedings of Machine Learning Research*, vol. 267, pp. 50926–50948. PMLR (2025)
39. Ruan, L., Chen, B., Li, J., Lam, M.: Learning to deblur using light field generated and real defocus images. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 16304–16313 (2022)
40. Song, Y., He, Z., Qian, H., Du, X.: Vision transformers for single image dehazing. *IEEE Transactions on Image Processing* **32**, 1927–1941 (2023)
41. Stanovich, K.E.: *Who is rational?: Studies of individual differences in reasoning*. Psychology Press (1999)
42. Tian, C., Zheng, M., Zuo, W., Zhang, B., Zhang, Y., Zhang, D.: Multi-stage image denoising with the wavelet transform. *Pattern Recognition* **134**, 109050 (2023)
43. Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al.: Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971* (2023)
44. Tu, Z., Talebi, H., Zhang, H., Yang, F., Milanfar, P., Bovik, A., Li, Y.: Maxim: Multi-axis mlp for image processing. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 5769–5780 (2022)
45. Wang, B., Qin, B., Li, J., Xu, F., Sun, F., Xiong, H.: All-in-one image restoration via causal-deconfounding wavelet-disentangled prompt network. *IEEE Transactions on Image Processing* **35**, 3421–3436 (2026)
46. Wang, J., Chan, K.C., Loy, C.C.: Exploring clip for assessing the look and feel of images. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 37, pp. 2555–2563 (2023)
47. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
48. Wei, J., Wang, X., Schuurmans, D., Bosma, M., Xia, F., Chi, E., Le, Q.V., Zhou, D., et al.: Chain-of-thought prompting elicits reasoning in large language models. *Advances in neural information processing systems* **35**, 24824–24837 (2022)
49. Wei, K., Fu, Y., Yang, J., Huang, H.: A physics-based noise formation model for extreme low-light raw denoising. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 2758–2767 (2020)
50. Wei, P., Xie, Z., Lu, H., Zhan, Z., Ye, Q., Zuo, W., Lin, L.: Component divide-and-conquer for real-world image super-resolution. In: *European conference on computer vision*. pp. 101–117. Springer (2020)
51. Wu, Q., Bansal, G., Zhang, J., Wu, Y., Li, B., Zhu, E., Jiang, L., Zhang, X., Zhang, S., Liu, J., et al.: Autogen: Enabling next-gen llm applications via multi-agent conversations. In: *First Conference on Language Modeling* (2024)
52. Wu, R., Sun, L., Ma, Z., Zhang, L.: One-step effective diffusion network for real-world image super-resolution. *Advances in Neural Information Processing Systems* **37**, 92529–92553 (2024)
53. Wu, R.Q., Duan, Z.P., Guo, C.L., Chai, Z., Li, C.: Ridep: Revitalizing real image dehazing via high-quality codebook priors. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 22282–22291 (2023)
54. Wu, X., Hao, Y., Sun, K., Chen, Y., Zhu, F., Zhao, R., Li, H.: Human preference score v2: A solid benchmark for evaluating human preferences of text-to-image synthesis. *arXiv preprint arXiv:2306.09341* (2023)

55. Xiao, J., Fu, X., Liu, A., Wu, F., Zha, Z.J.: Image de-raining transformer. *IEEE transactions on pattern analysis and machine intelligence* **45**(11), 12978–12995 (2022)
56. Yang, S., Wu, T., Shi, S., Lao, S., Gong, Y., Cao, M., Wang, J., Yang, Y.: Maniqa: Multi-dimension attention network for no-reference image quality assessment. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 1191–1200 (2022)
57. You, Z., Li, Z., Gu, J., Yin, Z., Xue, T., Dong, C.: Depicting beyond scores: Advancing image quality assessment through multi-modal language models. In: *European Conference on Computer Vision*. pp. 259–276. Springer (2024)
58. Yu, X., Peng, B., Vajipey, V., Cheng, H., Galley, M., Gao, J., Yu, Z.: Exact: Teaching ai agents to explore with reflective-mcts and exploratory learning. In: *The Thirteenth International Conference on Learning Representations* (2025)
59. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 5728–5739 (2022)
60. Zhang, K., Zuo, W., Chen, Y., Meng, D., Zhang, L.: Beyond a gaussian denoiser: Residual learning of deep cnn for image denoising. *IEEE transactions on image processing* **26**(7), 3142–3155 (2017)
61. Zhang, R., Isola, P., Efros, A.A., Shechtman, E., Wang, O.: The unreasonable effectiveness of deep features as a perceptual metric. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 586–595 (2018)
62. Zhong, W., Guo, L., Gao, Q., Ye, H., Wang, Y.: Memorybank: Enhancing large language models with long-term memory. In: *Proceedings of the AAAI conference on artificial intelligence*. vol. 38, pp. 19724–19731 (2024)
63. Zhou, Y., Ren, D., Emerton, N., Lim, S., Large, T.: Image restoration for under-display camera. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 9179–9188 (2021)
64. Zhu, K., Gu, J., You, Z., Qiao, Y., Dong, C.: An intelligent agentic system for complex image restoration problems. In: *International Conference on Learning Representations*. vol. 2025, pp. 57985–58013 (2025)
65. Zou, Z., Chen, K., Shi, Z., Guo, Y., Ye, J.: Object detection in 20 years: A survey. *Proceedings of the IEEE* **111**(3), 257–276 (2023)
66. Zuo, Y., Zheng, Q., Wu, M., Jiang, X., Li, R., Wang, J., Zhang, Y., Mai, G., Wang, L., Zou, J., et al.: 4kagent: Agentic any image to 4k super-resolution. In: *The Thirty-ninth Annual Conference on Neural Information Processing Systems* (2025)