E^{2}IR-AGENT: EXPERIENCE-DRIVEN SELF-EVOLVING IMAGE RESTORATION AGENTIC SYSTEM

Anonymous authors
Paper under double-blind review

ABSTRACT

Multimodal Large Language Model (MLLM)-driven image restoration agent demonstrates effectiveness in degradation coupling scenarios by flexibly selecting tools and determining removal orders. However, its zero-shot planning often fails without experience, necessitating severe trial-and-error overhead to achieve satisfactory outcomes. Currently, two paradigms are employed to address this issue, yet a dilemma persists: Training-based methods embed intrinsic experience into parameters, achieving high inference efficiency but lacking compatibility with new tools or degradations. In contrast, training-free methods utilize explicit experience storage for compatibility but still incur trial-and-error overhead due to naive experience. To resolve the dilemma, we propose **E^{2}IR-Agent**, which first systematically formulates the experience components of a training-free image restoration agent. Subsequently, a hierarchical experience pool is constructed, which enables coarse-to-fine guidance for diverse tools and removal orders. Furthermore, a self-evolving mechanism is introduced to update the pool from scratch using accumulated records, thereby greatly improving performance and efficiency. Extensive experiments reveal that **E^{2}IR-Agent** achieves a significant lead in the full reference metrics and yields a Pareto-optimal balance, which outperforms state-of-the-art methods in performance and is 61% more efficient.

1 INTRODUCTION

Image restoration (IR) is a low-level vision task that aims to recover a high-quality (HQ) image from a degraded low-quality (LQ) version. In recent decades, considerable progress has been made in addressing IR problems related to specific tasks (Zhang et al., 2017; Wang et al., 2018; Tao et al., 2018; Jin et al., 2025; Wang et al., 2023a). However, in real-world scenarios, multiple degradations are often coupled, making restoration more challenging.

A promising direction for addressing this challenge is all-in-one image restoration (AiOIR), which aims to handle multiple degradations with a unified model (Jiang et al., 2025a). Most methods (Li et al., 2020; Valanarasu et al., 2022; Conde et al., 2024; Jiang et al., 2024; Lin et al., 2025; Wu et al., 2026) address the degradation coupling problem by learning more general degradation feature representations. While significant strides have been made, they remain limited to a few degradations and are difficult to extend to new degradations without retraining.

Recent advances in MLLMs (Wang et al., 2024; Zhang et al., 2024) have inspired a new direction in IR, namely Image Restoration Agent (IRA). The key is to leverage a MLLM-driven agent to select restoration tools and determine removal orders. However, general-purpose MLLMs encounter difficulties with zero-shot planning (Chen et al., 2024a; Zhu et al., 2025) without experience-driven guidance, which results in severe trial-and-error overhead to reach satisfied restoration result. To resolve this issue, existing methods can be divided into two distinct paradigms. Training-based methods (Chen et al., 2024a; Lu et al., 2025; Zhang et al., 2026) internalize intrinsic experience into model parameters, offering high inference efficiency but lacking compatibility with new tools or degradations. Conversely, training-free methods (Zhu et al., 2025; Jiang et al., 2025b; Zuo et al., 2025) utilize external containers for explicit experience storage, while highly compatible, they still result in considerable trial-and-error overhead owing to being guided by naive experience. Therefore, we aim to retain the compatibility of training-free methods while endowing the model with the inference efficiency of training-based methods. We identify that the considerable trial-and-error overhead in

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

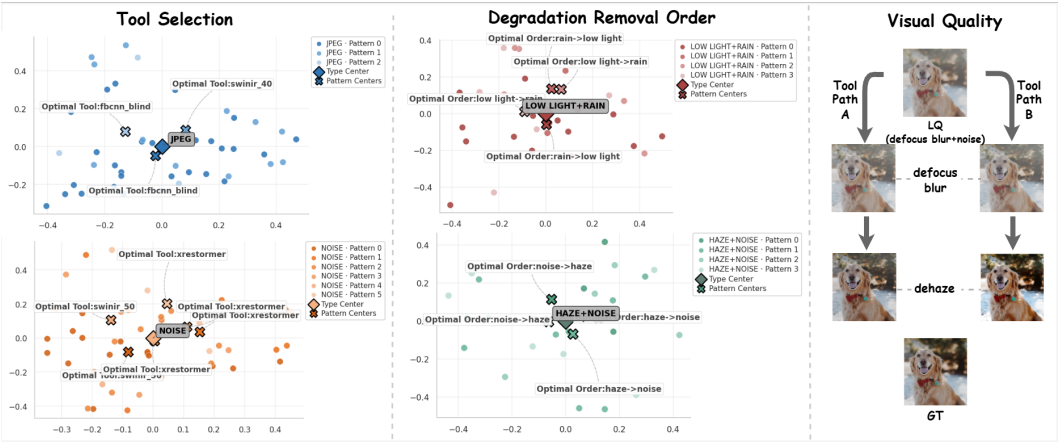


Figure 1: Illustration of the experience components in IRA. (Left&Center) We extract embeddings of various degraded images using an image encoder (Huang et al., 2026b). The scatter plots reveal that images sharing a degradation type may cluster into different degradation patterns. Crucially, the optimal *Tool Selection* and *Degradation Removal Order* are highly pattern-dependent. (Right) *Visual Quality* demonstrates that divergent tool paths can achieve restoration in distinct visual quality.

training-free IRAs stems from the underdeveloped experience mechanism. To systematically break this bottleneck, we revisit from two perspectives: (i) *Do existing IRAs lack an experience mechanism altogether?* Not entirely. Recent frameworks (Zhu et al., 2025; Jiang et al., 2025b; Zuo et al., 2025) already incorporate an experience mechanism. However, their experience components remain simplistic and static, preventing the agent from accepting comprehensive and dynamic guidance. (ii) *Is it possible to transfer experience mechanisms from general-purpose agents to IRAs?* It is possible. Recent advancements (Zhao et al., 2024; Zhang et al., 2025a; Xu et al., 2025; Fang et al., 2025) have demonstrated that explicit experience mechanisms drive self-evolving capability. However, it is unrealistic to directly apply it to IRAs without a domain-specific redesign.

Therefore, we introduce **E^{2}IR-Agent**, an image restoration agentic system explicitly designed to address the domain-specific needs. We first clearly identify the experience components in IRA as illustrated in Figure 1. (1) *Tool Selection*: The optimal tool depends not merely on broad degradation type, but on specific degradation pattern (Chen et al., 2024a). Experience should be considered to map these patterns to their optimal tool; (2) *Degradation Removal Order*: Similarly, in degradation coupling settings, the optimal order is also pattern-dependent; (3) *Visual Quality*: Due to inherent attributes in different tools, multiple paths may restore the image, but result in different visual outcomes. Overall, current experience mechanism focuses on removal order, with a few addresses on tool selection and showing an absence in focusing on visual quality. Moreover, they concentrate on the broad degradation type, while neglecting to delve into the degradation pattern at a finer granularity.

To systematically capture this coarse-to-fine information in conjunction with the aforementioned three experience components, we feature a hierarchical experience pool for storage. Inspired by (Zhang et al.; Yu et al., 2025; Pouransari et al., 2025), this pool retrieves contextually relevant experience across multiple levels, providing the IRA with robust, prioritized experience guidance. Building upon the structural foundation, the pool is equipped with a self-evolving experience mechanism. Unlike prior static experiences, ours iteratively updates using newly accumulated records. This dynamic evolution enables the system from compatibility towards adaptability in tools, degradations, and data sources, thereby promoting performance and inference efficiency. Our main contributions are summarized as follows:

(1) Systematic Formulation. We reformulate the experience components in Image Restoration Agent by introducing *visual quality* as a supplement. We also identify a limitation in current IRA experience, which solely focuses on degradation type while ignoring the degradation pattern.

(2) Self-Evolving Architecture. To effectively manage the formulated experience, we propose a hierarchical experience pool equipped with tailored retrieval methods. Building upon this structure, we introduce a self-evolving experience mechanism that updates the pool using accumulated records.

(3) Empirical Evaluation. Experiments demonstrate that our method outperforms state-of-the-art methods in performance and efficiency. Furthermore, the results validate the impact of the experience mechanism and its capacity for self-evolution, which endows the system with adaptability.

2 Related Work

All-in-one Image Restoration. All-in-one image restoration aims to handle multiple degradations within a unified model (Jiang et al., 2025a). AirNet (Li et al., 2020), PromptIR (Potlapalli et al., 2023) and MiOIR (Kong et al., 2024) focus on designing a unified backbone to model diverse degradation distributions. DA-CLIP (Luo et al., 2023) and InstructIR (Conde et al., 2024) exploit degradation-aware vision-language priors for more flexible restoration control. Some methods leverage generative priors to enhance restoration quality. AutoDIR (Jiang et al., 2024), UniRestore (Chen et al., 2025a), and Defusion (Luo et al., 2025) formulate restoration in a diffusion framework. Later methods explore more general restoration representations (Wu et al., 2026; Tang et al., 2026). However, these works are constrained to a limited number of degradations, and their capacity for extension is limited.

Image Restoration Agent. Existing image restoration agent approaches can be broadly divided into training-based and training-free methods. Training-based methods, such as RestoreAgent (Chen et al., 2024a) and Q-Agent (Zhou et al., 2025), equip the MLLM with degradation perception and planning capability through fine-tuning. Restore-R1 (Lu et al., 2025), TIR-Agent (Zhang et al., 2026), and PaAgent (Wang et al., 2026) follow in leveraging reinforcement learning (RL) for efficiency planning. In contrast, training-free methods focus on system design. AgenticIR (Zhu et al., 2025), MAIR (Jiang et al., 2025b), and 4KAgent (Zuo et al., 2025) place significant emphasis on agentic workflow orchestration. Nevertheless, these works neglect a robust experience mechanism.

Experience Mechanism in General-Purpose Agent. Recent methods have demonstrated that explicit experience can enhance the agent's capabilities (Zhang et al., 2025b; Huang et al., 2026a). Early agentic systems, such as ExpeL (Zhao et al., 2024), implement this idea by collecting and reflecting on trajectories in natural language, then retrieving distilled information at test time. Subsequent works (Xiong et al., 2025; Xia et al., 2025) examine not only how to store experience, but also how to manage it. Finally, FLEX (Cai et al., 2025) and SimpleMem (Liu et al., 2026) treat experience as the engine of continual adaptation. Overall, existing methods suggest that the experience mechanism in agents is evolving from simple episodic replay toward selective memory management, promising the system's ability to evolve and adapt.

3 Methodology

This section introduces the closed-loop framework of the **E^{2}IR-Agent**. As demonstrated in Figure 2, we begin with the outer loop by establishing the five-process inference workflow (Sec. 3.1). Subsequently, we shift to the inner loop and begin with mathematically formulating the experience. By encapsulating three core components—tool selection, degradation removal order, and visual quality—we model the experience by a priority-driven mapping function p (Sec. 3.2). To manage these components, we design a hierarchical experience pool structured across three granularities. This architecture practically instantiates p specifically for each level, providing coarse-to-fine guidance during inference (Sec. 3.3). Ultimately, we introduce a self-evolving experience mechanism that functions in a batch-wise manner. This mechanism enables the continuous updating of p from accumulated records, thereby ensuring performance efficiency and enabling adaptability in the system. (Sec. 3.4).

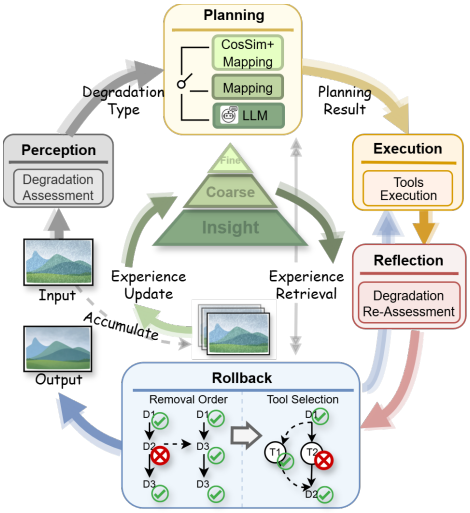


Figure 2: Overview of the experience-driven inference workflow in **E^{2}IR-Agent**.

3.1 EXPERIENCE-DRIVEN INFERENCE WORKFLOW

The five-process inference workflow adopted in prior training-free IRAs (Zhu et al., 2025) is maintained. Formally, given a degraded image I_D, the *Perception* process identifies the degradations, denoted as $D = \{d_1, d_2, \dots\}$. Based on the result, the *Planning* process performs decision-making with the degradation removal order O, which is defined as a permutation over D, i.e., $O \in \Pi(D)$. Adhering to this plan, the *Execution* process utilities a series of restoration tools for specific degradation d_i, defined as $T = \{\tau_{d_1}^1, \tau_{d_1}^2, \dots, \tau_{d_2}^1, \dots\}$ to obtain the restored image $\hat{I}_D$. The *Reflection* process is then performed to evaluate if all D have been eliminated on $\hat{I}_D$. If so, the process terminates and outputs the image. Otherwise, the *Rollback* process is triggered. The utilization of the experience pool enables the MLLM to accept robust guidance during the *Planning* and *Rollback* processes, and facilitates the management of O and T usage.

3.2 EXPERIENCE FORMULATION

We first define the aforementioned concept of visual quality. Observations from the restored images $\hat{I}_D$ from prior works (detailed in Appendix A.1) reveal a critical phenomenon: distinct tool selections can yield restored images that exhibit divergent visual quality. We conceptualize these differences as a *Preference* requirement, denoted by Q. It signifies a subjective preference interpretation that can be specified by humans and supported by different image quality assessment (IQA) metrics. In this work, *Fidelity* and *Perception* are specified as *Preference* requirements.

Subsequently, we mathematically formulate the guidance provided by the experience in the priority-driven mapping function p. The objective of *Tool Selection* is:

$$\tau^* = p_{tool}\big(T(d) \mid Q, d, H\big), \tag{1}$$

where $T(d) \subseteq T$ indicates the candidate tool set corresponds to a degradation d, and H contains the historical information during the decision-making process in the inference stage.

In cases of degradation coupling setting, the objective of determining the *Degradation Removal Order*, denoted by O, should be established in conjunction with the relevant restoration tools $T(D)$:

$$o^* = p_{order}\big(O \mid Q, T(D), D, H\big). \tag{2}$$

Experience Foundation. It should be noted that the number of permutations in O increases exponentially as D and T rise, which makes direct experience modeling inefficient and redundant. Therefore, we propose to simplify the condition on experience and form it as a foundation. The details are provided in Appendix A.1. Here's a summary:

- It is imperative to formulate p_{tool} before p_{order}, since p_{tool} can be considered as p_{order} during $|D| = 1$. Subsequent sections will denote p for unified representation.
- During the process of experience formulation, when specifying the preference Q, we decoupled p_{tool} and p_{order}, which reduces the joint optimization from the $\prod_{d_i \in D} \mathcal{T}(d_i)$ search space down to a purely D-centric problem for the priority function p_{order}.
- During rollback in inference time, O is revised first with the optimal tool τ^* derived from p_{tool}, followed by the refinement of T according to p_{tool}.

3.3 HIERARCHICAL EXPERIENCE POOL

In light of the aforementioned formulation and the findings presented in Figure 1, the optimal τ^* and o^* depend not only on the degradation type but also on the degradation pattern. Therefore, we draw inspiration from the hierarchical experience pool used in general-purpose agents, which effectively balances information across different granularities. Adopting this structure leads to a distinct instantiation of the priority-driven mapping function p:

Insight Level (Preference-oriented): At the highest level, p is instantiated by a LLM. For given Q and D, the LLM leverages distilled textual experience to guide the planning on o.

Coarse-grained Level (Degradation Type-oriented): This level instantiates p into a direct key-value mapping without the MLLM involved. Utilizing Q and D as the query, the value is considered as o.

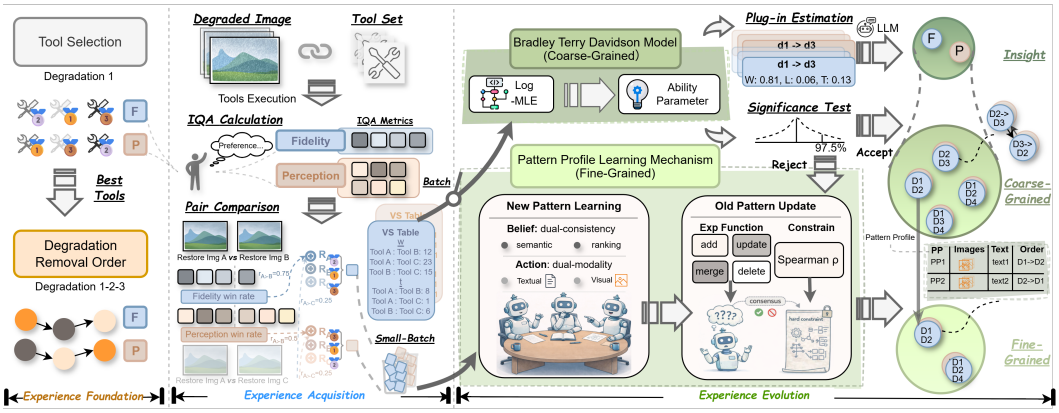


Figure 3: Overview of the inner closed-loop experience mechanism in **E²IR-Agent**. (Left) Experience pool works as a priority-driven mapping problem. (Middle) Using priority-aware modeling to generate an atomic experience record. (Right) Batch-wise experience evolving mechanism for the hierarchical pool to update the priority function.

Fine-grained Level (Degradation Pattern-oriented): According to Figure 1, a degradation type may cluster into different degradation patterns. This level still instantiates p into a key-value mapping, while combined with a retrieval strategy. The core intuition here is that a similar degradation pattern exhibits a similar priority function p. We first anchor these specific patterns using a *pattern profile*:

$$P = \left(\{I_D^{(i)}\}_{i=1}^{m},\ text \right), \tag{3}$$

where $\{I_D^{(i)}\}$ is a supporting set of degraded images characterizing the pattern, and text is its semantic description. (The construction of P is detailed in Section 3.4). During inference, given an input image I_D along with Q and D, we execute a two-stage cascade retrieval. First, we employ a CLIP encoder $e(\cdot)$ (Huang et al., 2026b) to recall the Top-K relevant profiles based on cosine similarity:

$$\mathrm{CosSim}(P, I_D) = \frac{e(P) \cdot e(I_D)}{|e(P)|\,|e(I_D)|}. \tag{4}$$

Subsequently, the MLLM evaluates the Top-K candidates $text$ against the input image to identify the most accurate profile P^*. Ultimately, P^* is utilized as a query condition, value is considered as o.

In light of the preceding experience foundation, the aforementioned three levels are characterized by the automatic selection of τ on each d with preference Q according to the coarse-grained level.

3.4 Self-Evolving Experience Mechanism

Based on the previous formulation of each experience and the structure of the experience pool, this section explains how the entire self-evolving experience mechanism is applied to the experience pool and operates in generating the latest priority-driven mapping function p.

Experience Acquisition. As illustrated in the middle panel of Figure 3, we undertake an exhaustive exploration of all potential outcomes to establish the comprehensive priority-aware model to generate an atomic experience record, which bootstraps the subsequent evolution stage.

Within an atomic experience record, a degraded image I_D with degradation types D and its corresponding tools $T(D)$ are employed, thereby yielding a set of restoration images:

$$\hat{I}_D = \{\hat{I}_D^1, \ldots, \hat{I}_D^N\},\ \text{where } N = \begin{cases} |T(D)|, & |D| = 1 \\ |D|!, & |D| > 1 \end{cases} \tag{5}$$

Subsequently, the results are evaluated using IQA metrics according to preference Q. A pairwise comparison is performed on each metric in regard to the restoration images $\hat{I}_D$:

$$r_{i \succ j} = \frac{1}{|M|} \sum_{m \in M} \mathbb{I}_m(i \succ j), \quad \mathbb{I}_m(i \succ j) = \begin{cases} 1, & m \in M\uparrow \text{ and } m(\hat{I}_D^i) > m(\hat{I}_D^j) \\ 1, & m \in M\downarrow \text{ and } m(\hat{I}_D^i) < m(\hat{I}_D^j), \\ 0, & \text{otherwise} \end{cases} \tag{6}$$

where $M \uparrow$ and $M \downarrow$ denote the higher-is-better and lower-is-better metrics, respectively. Based on these, we derive the win rate pair $r_{i \succ j}$ along with the win/loss/tie statistics w_{ij}, l_{ij}, and t_{ij}, where $r_{i \succ j} > 0.5$ indicates a vote for i, $r_{i \succ j} < 0.5$ indicates a vote for j, otherwise, considered a tie.

Following this, the average win rate of each tool is calculated, and its ranking is derived via sorting:

$$R_i = \frac{1}{|\hat{I}_D| - 1} \sum_{j \neq i} r_{i \succ j}, \qquad \text{rank}_i = \text{rank}_{\downarrow}(R_i). \tag{7}$$

Experience Evolution. As illustrated in the right panel of Figure 3, once a set of B atomic experience records has been collected for D, the process will be triggered on the batch. Ultimately, the designated level of D will be used in inference.

In the coarse-grained level, the self-evolving process commences with the utilization of the latest win/loss/tie counts w_{ij}, l_{ij}, and t_{ij} in B rounds in conjunction with the preceding counts. Leveraging these statistics, the ***Bradley–Terry–Davidson (BTD)*** model (Davidson, 1970) is adopted for the purpose of parameterizing the win/loss/tie relations in a probabilistic manner:

$$P(i \succ j) = \frac{e^{\theta_i}}{e^{\theta_i} + e^{\theta_j} + 2\nu e^{(\theta_i + \theta_j)/2}}, \quad P(i = j) = \frac{2\nu e^{(\theta_i + \theta_j)/2}}{e^{\theta_i} + e^{\theta_j} + 2\nu e^{(\theta_i + \theta_j)/2}}, \tag{8}$$

where $\nu \geq 0$ denotes the tie-intensity parameter, and θ defines as the ability parameter. θ can be estimated by maximizing the log-likelihood through a numerical process. Subsequently, θ is sorted and can be regarded as the latest priority function p in coarse-grained granularity:

$$p = \text{argsort}(\hat{\theta}), \qquad \hat{\theta}_1 > \hat{\theta}_2 > \cdots > \hat{\theta}_k. \tag{9}$$

We undertake significance testing in θ to ascertain whether fine-grained experience is necessary. A one-sided Wald test (Wald, 1943) is performed on the Top-ranked candidates θ at a significance level of α. The ability gap between two candidates is measured by $\hat{d}_{ij} = \hat{\theta}_i - \hat{\theta}_j$. We regard as significant if:

$$\hat{d}_{ij} - z_\alpha \, SE(\hat{d}_{ij}) \geq 0, \tag{10}$$

where SE denotes the standard error of θ. If the condition is met, no finer experience is demanded. Otherwise, additional fine-grained experience modeling is triggered.

In the insight level, derived from the latest estimated ability parameter θ, the pairwise win/loss/tie relations can be deduced using equation 8. Following this procedure, the deduced relations are then concatenated into textual form and combined with a guidance prompt. This enables LLM to distill textual experience, and it will be regarded as the latest priority function p in insight granularity:

$$p = \text{LLM}(\text{Guidance}, \text{Concate}(P(i \succ j), P(i = j))) \,. \tag{11}$$

In the fine-grained level, experiences necessitate a focus on the degradation pattern. We align with the previous definition of P, which anchors similar degradation patterns. However, it should be noted that the categories of P are not predefined in advance, and thus require a learning mechanism. A useful P must satisfy the following two requirements:

- **Semantic consistency** C_s: The images $\{I_D^{(i)}\}_{i=1}^m$ in P should be close in the semantic space to form a compact feature, thus facilitating retrieval by image encoder (Huang et al., 2026b).
- **Ranking consistency** C_r: The priority p mapping to P should be statistically comparable. This ensures that the priority p of a similar P in the experience pool can be reused.

In accordance with these, we design an agent-driven pattern profile learning mechanism to transform a dual-objective mathematical problem via inter-agent policy optimization.

Using cached rank_i in the latest B rounds, the system performs learning in a mini-batch manner. For each degraded image I_D^i in the small batch, we first generate a degradation description desc_i by MLLM. Subsequently, the multi-agent debate (MAD) (Chan et al.; Smit et al., 2024) paradigm is employed to partition the current small batch into a new pattern profile:

$$P_{\text{new}} = \text{MAD}(\text{role}, \text{action}, \mathcal{C}), \quad \mathcal{C} = \{C_s(\text{desc}), \; C_r(\text{rank})\}, \tag{12}$$

Table 1: Quantitative comparison under different degradation settings (Groups A–C) in two *Preference* (Fidelity and Perception). We compare methods within two distinct tool settings in accordance with the baseline methods for fair comparison: **Setting I** (AgenticIR-based) and **Setting II** (4KAgent-based). Rankings (**First** , Second , Third) are calculated independently within each group.

Setting	Method	Group A						Group B		
		Fidelity			Perception			Fidelity		
		PSNR↑	SSIM↑	LPIPS↓	MANIQA↑	CLIP-IQA↑	MUSIQ↑	PSNR↑	SSIM↑	LPIPS↓
I	AirNet	19.13	0.6019	0.4283	0.2581	0.3930	42.46	19.31	0.6567	0.3670
	PromptIR	20.06	0.6088	0.4127	0.2633	0.4013	42.62	20.47	0.6704	0.3370
	MiOIR	20.84	0.6558	0.3715	0.2451	0.3933	47.82	20.56	0.6905	0.3243
	DA-CLIP	19.58	0.6032	0.4266	0.2418	0.4139	42.51	18.56	0.5946	0.4405
	InstructIR	18.03	0.5751	0.4429	0.2660	0.3528	45.77	18.34	0.6235	0.4072
	AutoDIR	19.64	0.6286	0.3967	0.2500	0.3767	47.01	19.90	0.6643	0.3542
	AgenticIR	21.04	0.6818	0.3148	0.3071	0.4474	56.88	20.55	0.7009	0.3072
	MAIR	21.02	0.6715	0.2963	0.3330	0.4751	59.19	20.92	0.7004	0.2788
	E²IR-Agent-I	**23.49**	**0.7368**	**0.2840**	**0.3440**	**0.4889**	**60.10**	**23.41**	**0.7589**	**0.2753**
II	4KAgent	21.48	0.6720	0.3019	0.3748	**0.5544**	63.19	20.95	0.6727	0.3017
	TIR-Agent	22.07	0.6935	**0.2874**	**0.3907**	0.5454	**63.69**	22.80	0.7340	**0.2725**
	E²IR-Agent-II	**23.86**	**0.7488**	0.3085	0.3623	0.5204	61.82	**23.38**	**0.7643**	0.3117

Setting	Method	Group B			Group C					
		Perception			Fidelity			Perception		
		MANIQA↑	CLIP-IQA↑	MUSIQ↑	PSNR↑	SSIM↑	LPIPS↓	MANIQA↑	CLIP-IQA↑	MUSIQ↑
I	AirNet	0.2882	0.4274	47.88	17.95	0.5145	0.5782	0.1854	0.3113	30.12
	PromptIR	0.2893	0.4289	48.10	18.51	0.5166	0.5756	0.1906	0.3104	29.71
	MiOIR	0.2638	0.4330	51.87	15.63	0.4896	0.5376	0.1717	0.2891	37.95
	DA-CLIP	0.2435	0.4154	43.70	18.53	0.5320	0.5335	0.1916	0.3476	33.87
	InstructIR	0.3022	0.3790	50.94	17.09	0.5135	0.5582	0.1732	0.2537	33.69
	AutoDIR	0.2534	0.3986	49.64	18.61	0.5443	0.5019	0.2045	0.2939	37.86
	AgenticIR	0.3204	0.4648	57.57	18.82	0.5474	0.4493	0.2698	0.3948	48.68
	MAIR	**0.3544**	**0.5084**	**60.98**	19.42	0.5544	0.4142	0.2798	0.4239	**51.36**
	E²IR-Agent-I	0.3535	0.4984	**60.98**	**21.52**	**0.6354**	**0.4108**	**0.2954**	**0.4448**	**51.36**
II	4KAgent	0.3734	**0.5505**	62.69	19.77	0.5629	0.4271	0.3545	0.5233	55.56
	TIR-Agent	**0.3997**	0.5488	**65.24**	19.53	0.5643	**0.4120**	**0.3739**	**0.5495**	**56.93**
	E²IR-Agent-II	0.3735	0.5296	61.96	**22.12**	**0.6588**	0.4528	0.3321	0.4977	55.21

where *role* specifies the agent's designated roles included in the system, *action* encapsulates some pre-defined executable operations with dual-modality, and $\mathcal{C}(\cdot)$ denotes the dual-consistency requirements mentioned above, which are achieved as the belief (Rao et al., 1995) for each agent.

After partitioning P_{new}, we further optimize and iterate on the existing P_{old} through a iteration process between a soft updates to keep up with C_s driven by an agent and a hard constraint to keep up with C_r by Spearmanρ (spe, 1961):

$$P_{\text{new}}^{t+1} \leftarrow \text{Iter}\big(P_{\text{new}}^{t}, P_{\text{old}}^{t}; \text{Agent(func)}, \rho\big) \quad \rho(\text{rank}_i, \text{rank}_j) = 1 - \frac{6\sum_{i=1}^{n}\Delta_i^2}{n(n^2-1)}, \tag{13}$$

where *func* denotes the operation functions on P for agent, which includes the following meta operation: add, replace, update and delete. Through negotiation, the iteration process reaches a consensus on appending $\{I_D^{(i)}\}_{i=1}^{m}$ and rectifying the $text$ representation of P.

Following multiple mini-batch learning steps regularized by C_s and C_r, the information within P stabilizes. Consequently, the historically cached ranks associated with the images clustered in the specific P also converge. By leveraging these stable ranks, we derive the final priority function p bound to P, which functions as the priority in fine-grained granularity:

$$p = \text{Stabilize}\left(\text{rank}^{t}\right). \tag{14}$$

Table 2: Quantitative comparison on real-world settings from the FoundIR in two *Preference* (Fidelity and Perception). Rankings are highlighted with color (**First** , Second , and Third).

Methods	Fidelity			Perception			Invoc. ↓
	PSNR ↑	SSIM ↑	LPIPS ↓	MANIQA ↑	CLIP-IQA ↑	MUSIQ ↑	
AgenticIR	19.14	0.6283	0.3770	0.2666	0.3692	52.52	5.95
4KAgent	19.76	0.7071	0.3518	0.2969	**0.3873**	57.35	10.89
E²IR-Agent	**20.05**	**0.7101**	**0.3279**	**0.3340**	0.3832	**57.51**	**4.23**

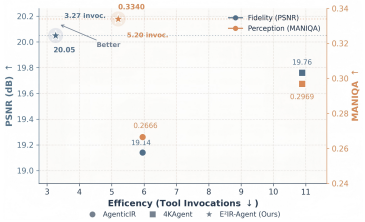


Figure 4: Performance vs. Efficiency trade-off visualization.

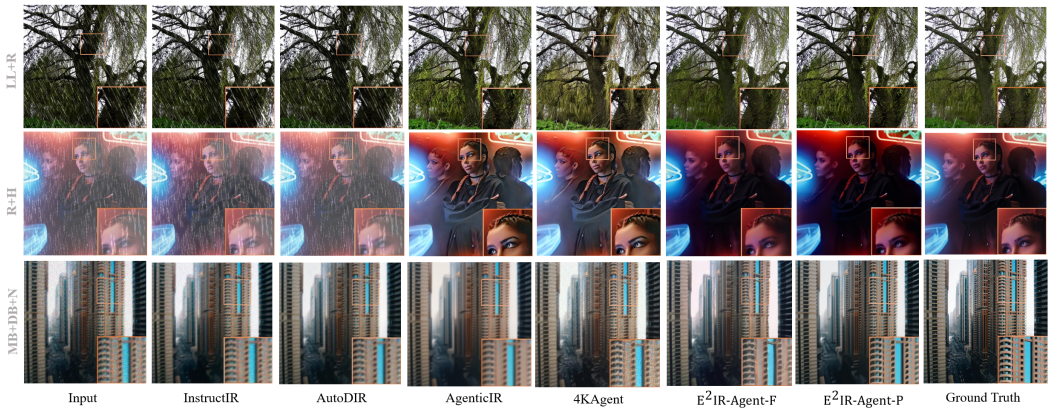


Figure 5: Visual comparison examples on MiO100 dataset. **E²IR-Agent-F** denoted the fidelity *Preference* requirement, while **E²IR-Agent-P** denoted the perception *Preference* requirement.

4 Experiments

4.1 Experimental Setup

Test Datasets and Tool Settings. To verify the applicability of **E²IR-Agent**, a comprehensive testing process is conducted using both synthetic and real-world datasets. In the case of the synthesized datasets, MiO100 is employed, which contains 1,440 low-quality images across 16 combinations of two or three types of degradations. In the case of the real-world dataset, FoundIR (Li et al., 2025) is employed for the purpose of zero-shot generalization testing in three degradation coupling scenarios. Each scenario includes 50 real-world image pairs. The detailed list of experience evolving and inference settings can be found in Appendix C.

Compared Methods. We compare our method with six AiOIR models: AirNet (Li et al., 2020), PromptIR (Potlapalli et al., 2023), MiOIR (Kong et al., 2024), DA-CLIP (Luo et al., 2023), InstructIR (Conde et al., 2024), and AutoDIR (Jiang et al., 2024). Furthermore, we compare with four agent-based IR methods: AgenticIR (Zhu et al., 2025), MAIR (Jiang et al., 2025b), 4KAgent (Zuo et al., 2025), and TIR-Agent (Zhang et al., 2026). All the methods above are without a re-training process.

Evaluation Metrics. We evaluate the quality using both full-reference (FR) and no-reference (NR) image quality assessment (IQA) metrics. The FR metrics include PSNR, SSIM (Wang et al., 2004), and LPIPS (Zhang et al., 2018), while the NR metrics include MANIQA (Yang et al., 2022), CLIP-IQA (Wang et al., 2023b), and MUSIQ (Ke et al., 2021).

4.2 Performance Comparison with Other Methods

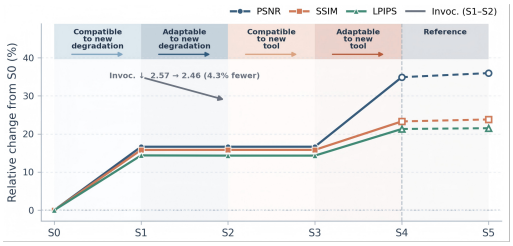
Comparison with Other Methods. As shown in Table 1, E²IR-Agent significantly outperforms state-of-the-art methods in Setting I in nearly all metrics, especially leading the baselines by an average of 2.35 dB in PSNR and 0.065 in SSIM. In Setting II, our performance is still competitive in FR metrics compared to the baselines and presents superior performance compared to Setting I, suggesting that our comparison of different tools between different baselines is fair and reasonable.

Granularity	Fidelity		Perception		Efficiency
	PSNR ↑	SSIM ↑	MANIQA ↑	MUSIQ ↑	Total Rb ↓
w/o experience	21.02	0.7923	**0.4081**	63.21	3.74
w/ insight	21.56	0.8103	0.3933	64.58	2.61
w/ coarse	22.20	0.8134	0.3913	64.55	1.46
w/ fine	**22.44**	**0.8165**	0.3941	**64.73**	**1.33**

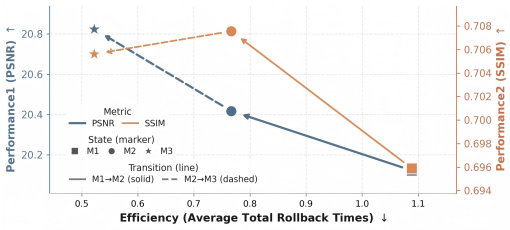
(a) Ablation on experience granularity.

Times	Fidelity		Perception		Efficiency
	PSNR ↑	SSIM ↑	MANIQA ↑	MUSIQ ↑	Total Rb ↓
0	20.59	0.6579	0.2770	56.69	2.50
1	22.21	0.7346	0.3459	61.26	1.67
2	**23.38**	0.7643	0.3735	61.96	**1.29**
3	23.03	0.7539	**0.3749**	**62.44**	1.51
4	23.04	**0.7647**	0.3714	61.79	1.47

(b) Ablation on evolving times.



(c) Adaptability on degradation and tool.



(d) Adaptability on data.

Figure 6: Ablation studies and additional experiments on core components.

Generalization to Real-World Scenarios. Table 2 and Figure 4 evaluates the zero-shot generalization testing within the tool setting II. Our method consistently surpasses existing open-source IRAs in performance and efficiency, establishing a Pareto-optimal trade-off between restoration performance and inference efficiency. Specifically, it outperforms the strongest baselines by up to 0.0239 in LPIPS and 0.0371 in MANIQA while simultaneously slashing tool invocation number by 61%.

4.3 Ablation Study

Evaluation on the hierarchical experience pool. As shown in Table 6a, when the granularity transitions from insight to fine, a substantial enhancement in performance and efficiency is observed in both *Preference* settings. The incorporation of an experience mechanism has proven to be an effective solution to address the issue of rollback in IRAs. A comprehensive examination of the hierarchy experience trade-off is furnished in Appendix B.2.

Evaluation on the evolution times. As shown in Table 6b, Times $= 2$ emerges as the optimal sweet spot between performance and efficiency when $B = 25$. The self-evolving experience mechanism not only enhances performance but also reduces inference latency during evolution on the data source. A comprehensive examination of these factors is furnished in Appendix B.4.

Demonstration of adaptability. We conduct additional verification of the adaptability of our method. The self-evolving characteristic enables the system to transition from compatibility to adaptability. Figure 6c demonstrates the ability to adapt to new degradation (from S0 to S2) and tools (from S2 to S4) in the context of continuous learning. Moreover, Figure 6d exhibits the capacity to adapt to different data sources (from M1), facilitating the transition (M2) to the optimal theoretical state (M3). A thorough exposition of the method's adaptability is furnished in Appendix B.1.

5 Conclusion

In this paper, we address a dilemma in Image Restoration Agent (IRA) for the experience design: the limited compatibility of training-based methods versus the inference inefficiency of training-free paradigms. To resolve this, we propose **E^{2}IR-Agent**, which employs a hierarchical experience pool to provide coarse-to-fine guidance for three experience components. Furthermore, we introduce a self-evolving mechanism that iteratively refines this experience pool using accumulated records. This continuous evolution enables us to achieve a balance between restoration quality and efficiency while concurrently demonstrating adaptability to degradations, tools, and data sources. Extensive experiments demonstrate that we successfully bridge the gap between existing paradigms through its self-evolving, experience-driven learning capacity.

AI USE STATEMENT

In this work, generative AI tool is used solely to aid in polishing the writing and improving grammatical clarity. All AI-assisted text is carefully reviewed, verified, and revised by the authors to ensure factual accuracy. We have not used generative AI tool for research ideation, theoretical derivation or experimental design. The authors take full responsibility for the entire work, including all text, claims, and results.

ETHICS STATEMENT

This work does not raise any direct ethical concerns.

REPRODUCIBILITY STATEMENT

To ensure reproducibility, comprehensive details regarding the datasets, hyperparameter settings, and experimental configurations are provided in the Appendix. We will make the complete source code and instructions to reproduce our results publicly available in the near future.

REFERENCES

The proof and measurement of association between two things. 1961.

Zhicheng Cai, Xinyuan Guo, Yu Pei, Jiangtao Feng, Jinsong Su, Jiangjie Chen, Ya-Qin Zhang, Wei-Ying Ma, Mingxuan Wang, and Hao Zhou. Flex: Continuous agent evolution via forward learning from experience. *arXiv preprint arXiv:2511.06449*, 2025.

Chi-Min Chan, Weize Chen, Yusheng Su, Jianxuan Yu, Wei Xue, Shanghang Zhang, Jie Fu, and Zhiyuan Liu. Chateval: Towards better llm-based evaluators through multi-agent debate. In *The Twelfth International Conference on Learning Representations*.

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, 2024a.

I Chen, Wei-Ting Chen, Yu-Wei Liu, Yuan-Chun Chiang, Sy-Yen Kuo, Ming-Hsuan Yang, et al. Unirestore: Unified perceptual and task-oriented image restoration model using diffusion prior. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 17969–17979, 2025a.

Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *European conference on computer vision*, pp. 17–33. Springer, 2022.

Xiangyu Chen, Zheyuan Li, Yuandong Pu, Yihao Liu, Jiantao Zhou, Yu Qiao, and Chao Dong. A comparative study of image restoration networks for general backbone network design. In *European Conference on Computer Vision*, pp. 74–91. Springer, 2024b.

Xiangyu Chen, Xintao Wang, Wenlong Zhang, Xiangtao Kong, Yu Qiao, Jiantao Zhou, and Chao Dong. Hat: Hybrid attention transformer for image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025b.

Shu-Chuan Chu, Zhi-Chao Dou, Jeng-Shyang Pan, Shaowei Weng, and Junbao Li. Hmanet: Hybrid multi-axis aggregation network for image super-resolution. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6257–6266, 2024.

Marcos V Conde, Gregor Geigle, and Radu Timofte. Instructir: High-quality image restoration following human instructions. In *European Conference on Computer Vision*, pp. 1–21. Springer, 2024.

Yuning Cui, Wenqi Ren, Xiaochun Cao, and Alois Knoll. Revitalizing convolutional network for image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12): 9423–9438, 2024a.

Yuning Cui, Wenqi Ren, and Alois Knoll. Omni-kernel network for image restoration. In *Proceedings of the AAAI conference on artificial intelligence*, volume 38, pp. 1426–1434, 2024b.

Roger R Davidson. On extending the bradley-terry model to accommodate ties in paired comparison experiments. *Journal of the American Statistical Association*, 65(329):317–328, 1970.

Jizhan Fang, Xinle Deng, Haoming Xu, Ziyan Jiang, Yuqi Tang, Ziwen Xu, Shumin Deng, Yunzhi Yao, Mengru Wang, Shuofei Qiao, et al. Lightmem: Lightweight and efficient memory-augmented generation. *arXiv preprint arXiv:2510.18866*, 2025.

Chih-Chung Hsu, Chia-Ming Lee, and Yi-Shiuan Chou. Drct: Saving image super-resolution away from information bottleneck. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 6133–6142, 2024.

Wei-Chieh Huang, Weizhi Zhang, Yueqing Liang, Yuanchen Bei, Yankai Chen, Tao Feng, Xinyu Pan, Zhen Tan, Yu Wang, Tianxin Wei, et al. Rethinking memory mechanisms of foundation agents in the second half. *arXiv preprint arXiv:2602.06052*, 2026a.

Weiquan Huang, Aoqi Wu, Yifan Yang, Xufang Luo, Yuqing Yang, Usman Naseem, Chunyu Wang, Qi Dai, Xiyang Dai, Dongdong Chen, et al. Llm2clip: Powerful language model unlocks richer cross-modality representation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 5131–5139, 2026b.

Jiaxi Jiang, Kai Zhang, and Radu Timofte. Towards flexible blind jpeg artifacts removal. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 4997–5006, 2021.

Junjun Jiang, Zengyuan Zuo, Gang Wu, Kui Jiang, and Xianming Liu. A survey on all-in-one image restoration: Taxonomy, evaluation and future trends. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025a.

Xu Jiang, Gehui Li, Bin Chen, and Jian Zhang. Multi-agent image restoration. *arXiv preprint arXiv:2503.09403*, 2025b.

Yitong Jiang, Zhaoyang Zhang, Tianfan Xue, and Jinwei Gu. Autodir: Automatic all-in-one image restoration with latent diffusion. In *European Conference on Computer Vision*, pp. 340–359. Springer, 2024.

Zhi Jin, Yuwei Qiu, Kaihao Zhang, Hongdong Li, and Wenhan Luo. Mb-taylorformer v2: Improved multi-branch linear transformer expanded by taylor formula for image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2025.

Junjie Ke, Qifei Wang, Yilin Wang, Peyman Milanfar, and Feng Yang. Musiq: Multi-scale image quality transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 5148–5157, 2021.

Lingshun Kong, Jiangxin Dong, Jinhui Tang, Ming-Hsuan Yang, and Jinshan Pan. Efficient visual state space model for image deblurring. In *Proceedings of the computer vision and pattern recognition conference*, pp. 12710–12719, 2025.

Xiangtao Kong, Chao Dong, and Lei Zhang. Towards effective multiple-in-one image restoration: A sequential and prompt learning strategy. *arXiv preprint arXiv:2401.03379*, 2024.

Junyong Lee, Hyeongseok Son, Jaesung Rim, Sunghyun Cho, and Seungyong Lee. Iterative filter adaptive network for single image defocus deblurring. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2034–2042, 2021.

Hao Li, Xiang Chen, Jiangxin Dong, Jinhui Tang, and Jinshan Pan. Foundir: Unleashing million-scale training data to advance foundation models for image restoration. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 12626–12636, 2025.

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*, pp. 3175–3185, 2020.

Jingyun Liang, Jiezhang 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*, pp. 1833–1844, 2021.

Xinqi Lin, Jingwen He, Ziyan Chen, Zhaoyang Lyu, Bo Dai, Fanghua Yu, Yu Qiao, Wanli Ouyang, and Chao Dong. Diffbir: Toward blind image restoration with generative diffusion prior. In *European conference on computer vision*, pp. 430–448. Springer, 2024.

Yunlong Lin, Zixu Lin, Haoyu Chen, Panwang Pan, Chenxin Li, Sixiang Chen, Kairun Wen, Yeying Jin, Wenbo Li, and Xinghao Ding. Jarvisir: Elevating autonomous driving perception with intelligent image restoration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 22369–22380, 2025.

Jiaqi Liu, Yaofeng Su, Peng Xia, Siwei Han, Zeyu Zheng, Cihang Xie, Mingyu Ding, and Huaxiu Yao. Simplemem: Efficient lifelong memory for llm agents. *arXiv preprint arXiv:2601.02553*, 2026.

Yuhao Liu, Zhanghan Ke, Fang Liu, Nanxuan Zhao, and Rynson WH Lau. Diff-plugin: Revitalizing details for diffusion-based low-level tasks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 4197–4208, 2024.

Jianglin Lu, Yuanwei Wu, Ziyi Zhao, Hongcheng Wang, Felix Jimenez, Abrar Majeedi, and Yun Fu. Simplecall: A lightweight image restoration agent in label-free environments with mllm perceptual feedback. *arXiv preprint arXiv:2512.18599*, 2025.

Wenyang Luo, Haina Qin, Zewen Chen, Libin Wang, Dandan Zheng, Yuming Li, Yufan Liu, Bing Li, and Weiming Hu. Visual-instructed degradation diffusion for all-in-one image restoration. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 12764–12777, 2025.

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*, 3(8), 2023.

Xiaoqian Lv, Shengping Zhang, Chenyang Wang, Yichen Zheng, Bineng Zhong, Chongyi Li, and Liqiang Nie. Fourier priors-guided diffusion for zero-shot joint low-light enhancement and deblurring. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 25378–25388, 2024.

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.

Hadi Pouransari, David Grangier, C Thomas, Michael Kirchhof, and Oncel Tuzel. Pretraining with hierarchical memories: separating long-tail and common knowledge. *arXiv preprint arXiv:2510.02375*, 2025.

Anand S Rao, Michael P Georgeff, et al. Bdi agents: from theory to practice. In *Icmas*, volume 95, pp. 312–319, 1995.

Ali M Reza. Realization of the contrast limited adaptive histogram equalization (clahe) for real-time image enhancement. *Journal of VLSI signal processing systems for signal, image and video technology*, 38(1):35–44, 2004.

Lingyan Ruan, Bin Chen, Jizhou Li, and Miuling Lam. 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.

Andries Petrus Smit, Nathan Grinsztajn, Paul Duckworth, Thomas D Barrett, and Arnu Pretorius. Should we be going mad? a look at multi-agent debate strategies for llms. In *International Conference on Machine Learning*, pp. 45883–45905. PMLR, 2024.

Yuda Song, Zhuqing He, Hui Qian, and Xin Du. Vision transformers for single image dehazing. *IEEE Transactions on Image Processing*, 32:1927–1941, 2023.

Lingchen Sun, Rongyuan Wu, Zhiyuan Ma, Shuaizheng Liu, Qiaosi Yi, and Lei Zhang. Pixel-level and semantic-level adjustable super-resolution: A dual-lora approach. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2333–2343, 2025.

Xiaole Tang, Xiaoyi He, Jiayi Xu, Xiang Gu, and Jian Sun. Learning continuous wasserstein barycenter space for generalized all-in-one image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2026.

Xin Tao, Hongyun Gao, Xiaoyong Shen, Jue Wang, and Jiaya Jia. Scale-recurrent network for deep image deblurring. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 8174–8182, 2018.

Zhengzhong Tu, Hossein Talebi, Han Zhang, Feng Yang, Peyman Milanfar, Alan Bovik, and Yinxiao Li. Maxim: Multi-axis mlp for image processing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5769–5780, 2022.

Jeya Maria Jose Valanarasu, Rajeev Yasarla, and Vishal M Patel. Transweather: Transformer-based restoration of images degraded by adverse weather conditions. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2353–2363, 2022.

Abraham Wald. Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Transactions of the American Mathematical society*, 54(3):426–482, 1943.

Chenxi Wang, Hongjun Wu, and Zhi Jin. Fourllie: Boosting low-light image enhancement by fourier frequency information. In *Proceedings of the 31st ACM international conference on multimedia*, pp. 7459–7469, 2023a.

Jianyi Wang, Kelvin CK Chan, and Chen Change Loy. Exploring clip for assessing the look and feel of images. In *Proceedings of the AAAI conference on artificial intelligence*, volume 37, pp. 2555–2563, 2023b.

Peng Wang, Shuai Bai, Sinan Tan, Shijie Wang, Zhihao Fan, Jinze Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, et al. Qwen2-vl: Enhancing vision-language model's perception of the world at any resolution. *arXiv preprint arXiv:2409.12191*, 2024.

Xintao Wang, Ke Yu, Shixiang Wu, Jinjin Gu, Yihao Liu, Chao Dong, Yu Qiao, and Chen Change Loy. Esrgan: Enhanced super-resolution generative adversarial networks. In *Proceedings of the European conference on computer vision (ECCV) workshops*, pp. 0–0, 2018.

Yijian Wang, Qingsen Yan, Jiantao Zhou, Duwei Dai, and Wei Dong. Paagent: Portrait-aware image restoration agent via subjective-objective reinforcement learning. *arXiv preprint arXiv:2603.17055*, 2026.

Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

Haoning Wu, Zicheng Zhang, Weixia Zhang, Chaofeng Chen, Liang Liao, Chunyi Li, Yixuan Gao, Annan Wang, Erli Zhang, Wenxiu Sun, et al. Q-align: teaching lmms for visual scoring via discrete text-defined levels. In *Proceedings of the 41st International Conference on Machine Learning*, pp. 54015–54029, 2024a.

Jiawei Wu, Zhifei Yang, Zhe Wang, and Zhi Jin. Gradient as conditions: Rethinking hog for all-in-one image restoration. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 40, pp. 10682–10690, 2026.

Rongyuan Wu, Lingchen Sun, Zhiyuan Ma, and Lei Zhang. One-step effective diffusion network for real-world image super-resolution. *Advances in Neural Information Processing Systems*, 37: 92529–92553, 2024b.

Rui-Qi Wu, Zheng-Peng Duan, Chun-Le Guo, Zhi Chai, and Chongyi Li. Ridcp: 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.

Siyu Xia, Zekun Xu, Jiajun Chai, Wentian Fan, Yan Song, Xiaohan Wang, Guojun Yin, Wei Lin, Haifeng Zhang, and Jun Wang. From experience to strategy: Empowering llm agents with trainable graph memory. *arXiv preprint arXiv:2511.07800*, 2025.

Zidi Xiong, Yuping Lin, Wenya Xie, Pengfei He, Zirui Liu, Jiliang Tang, Himabindu Lakkaraju, and Zhen Xiang. How memory management impacts llm agents: An empirical study of experience-following behavior. *arXiv preprint arXiv:2505.16067*, 2025.

Wujiang Xu, Zujie Liang, Kai Mei, Hang Gao, Juntao Tan, and Yongfeng Zhang. A-mem: Agentic memory for llm agents. *arXiv preprint arXiv:2502.12110*, 2025.

An Yang, Anfeng Li, Baosong Yang, Beichen Zhang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Gao, Chengen Huang, Chenxu Lv, et al. Qwen3 technical report. *arXiv preprint arXiv:2505.09388*, 2025.

Sidi Yang, Tianhe Wu, Shuwei Shi, Shanshan Lao, Yuan Gong, Mingdeng Cao, Jiahao Wang, and Yujiu Yang. 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.

Zhiyuan You, Zheyuan Li, Jinjin Gu, Zhenfei Yin, Tianfan Xue, and Chao Dong. Depicting beyond scores: Advancing image quality assessment through multi-modal language models. In *European Conference on Computer Vision*, pp. 259–276. Springer, 2024.

Shuo Yu, Mingyue Cheng, Daoyu Wang, Qi Liu, Zirui Liu, Ze Guo, and Xiaoyu Tao. Memweaver: A hierarchical memory from textual interactive behaviors for personalized generation. *arXiv preprint arXiv:2510.07713*, 2025.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, Ming-Hsuan Yang, and Ling Shao. Multi-stage progressive image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 14821–14831, 2021.

Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 5728–5739, 2022.

Dafeng Zhang, Feiyu Huang, Shizhuo Liu, Xiaobing Wang, and Zhezhu Jin. Swinfir: Revisiting the swinir with fast fourier convolution and improved training for image super-resolution. *arXiv preprint arXiv:2208.11247*, 2022.

Guibin Zhang, Muxin Fu, Kun Wang, Guancheng Wan, Miao Yu, and Shuicheng YAN. G-memory: Tracing hierarchical memory for multi-agent systems. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems*.

Jingyi Zhang, Jiaxing Huang, Sheng Jin, and Shijian Lu. Vision-language models for vision tasks: A survey. *IEEE transactions on pattern analysis and machine intelligence*, 46(8):5625–5644, 2024.

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.

Qizheng Zhang, Changran Hu, Shubhangi Upasani, Boyuan Ma, Fenglu Hong, Vamsidhar Kamanuru, Jay Rainton, Chen Wu, Mengmeng Ji, Hanchen Li, et al. Agentic context engineering: Evolving contexts for self-improving language models. *arXiv preprint arXiv:2510.04618*, 2025a.

Richard Zhang, Phillip Isola, Alexei A Efros, Eli Shechtman, and Oliver Wang. The unreasonable effectiveness of deep features as a perceptual metric. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 586–595, 2018.

Yisheng Zhang, Guoli Jia, Haote Hu, Shanxu Zhao, Kaikai Zhao, Long Sun, Xinwei Long, Kai Tian, Che Jiang, Zhaoxiang Liu, et al. Tir-agent: Training an explorative and efficient agent for image restoration. *arXiv preprint arXiv:2603.27742*, 2026.

Zeyu Zhang, Quanyu Dai, Xiaohe Bo, Chen Ma, Rui Li, Xu Chen, Jieming Zhu, Zhenhua Dong, and Ji-Rong Wen. A survey on the memory mechanism of large language model-based agents. *ACM Transactions on Information Systems*, 43(6):1–47, 2025b.

Andrew Zhao, Daniel Huang, Quentin Xu, Matthieu Lin, Yong-Jin Liu, and Gao Huang. Expel: Llm agents are experiential learners. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pp. 19632–19642, 2024.

Yingjie Zhou, Jiezhang 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.

Kaiwen Zhu, Jinjin Gu, Zhiyuan You, Yu Qiao, and Chao Dong. An intelligent agentic system for complex image restoration problems. In *International Conference on Learning Representations*, volume 2025, pp. 57985–58013, 2025.

Yushen Zuo, Qi Zheng, Mingyang Wu, Xinrui Jiang, Renjie Li, Jian Wang, Yide Zhang, Gengchen Mai, Lihong Wang, James Zou, Xiaoyu Wang, Ming-Hsuan Yang, and Zhengzhong Tu. 4kagent: Agentic any image to 4k super-resolution. In D. Belgrave, C. Zhang, H. Lin, R. Pascanu, P. Koniusz, M. Ghassemi, and N. Chen (eds.), *Advances in Neural Information Processing Systems*, volume 38, Main Conference, pp. 164106–164175. Curran Associates, Inc., 2025. doi: 10.52202/085713-5476. URL https://proceedings.neurips.cc/paper_files/paper/2025/file/f0075fe4e59652cf43148dcfab8d3c93-Paper-Conference.pdf.

A Additional Proofs

A.1 Experience Foundation Proving

Findings on 1. Since (Zhu et al., 2025; Zuo et al., 2025) utilizes a heuristic strategy to select a tool for each degradation, the restoration procedure is terminated once the degradation has been removed. Consequently, we observe that multiple iterations of the same process may yield different restored images, and they often exhibit noticeable differences in visual quality, as demonstrated in Figure 7. This phenomenon motivates the introduction of *Preference* as an experience component in the IRA, denoted as Q. This is a configurable feature for human utilization and should be objectively judged by IQA methods. We instantiate two *Preference* settings (Fidelity and Perception) by using score-based IQAs. More generally, it is believed that language-based IQAs (You et al., 2024; Wu et al., 2024a) would better align with human preference.

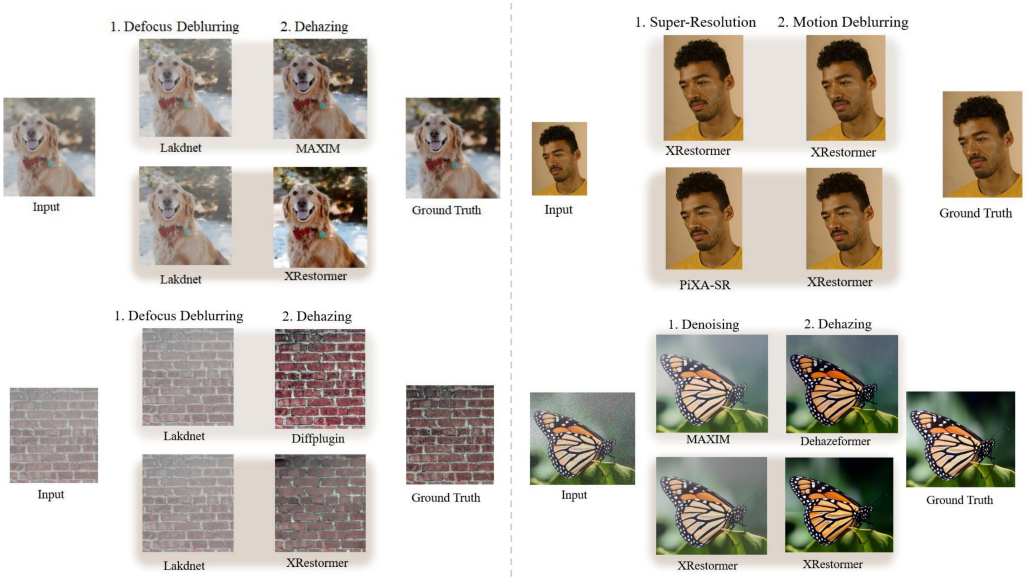


Figure 7: Visualization of *Preference* setting. It has been demonstrated that, in instances where image restoration is achieved, the visual quality of the restored image can vary within different tools.

Findings on 2. Under the degradation coupling setting, jointly exploring the tools T and degradation removal orders O leads to a prohibitive permutational space for experience modeling. We utilize the MiOIR dataset to assess three distinct degradation coupling scenarios and fix *Preference* on perception for experience evolving. This experiment is adopted to ascertain the impact of varying p_{tool} values on p_{order}.

Table 3: Influence of p_{tool} on p_{order}. $p_{tool\,i}$ denotes the priority i-th for the corresponding degradation.

Degradation pair	$\mathbf{p_{tool_1} + p_{tool_1}}$	$\mathbf{p_{tool_1} + p_{tool_2}}$	$\mathbf{p_{tool_2} + p_{tool_1}}$	$\mathbf{p_{tool_2} + p_{tool_2}}$	$\mathbf{p_{tool_3} + p_{tool_3}}$	**Agreement**
Defocus blur + Haze	D→H	D→H	D→H	D→H	D→H	5/5
Rain + Haze	H→R	R→H	R→H	R→H	H→R	3/5
Motion blur + Dark	M→B	M→B	M→B	M→B	B→M	4/5
					Overall modal-order agreement	**12/15 (80%)**

As shown in Table 3, we propose to decouple two priority functions, which reduce the joint optimization from a $\prod_{d_i \in D} \mathcal{T}(d_i)$ search space down to a purely D-centric permutation problem. The decoupling setting enables p_{tool} to be updated optionally under the degradation coupling setting.

Findings on 3. Subsequently, we conduct inference testing on the MiO100 (Kong et al., 2024) dataset. Our objective is to ascertain the optimal rollback strategy under the above findings.



(a) Defocus Blur + Haze (b) Rain + Haze (c) Low light + Motion Blur

Figure 8: Influence on the rollback mechanism between O and T. All metrics are normalized and maximized based on the optimal degradation removal order (dashed line).

As demonstrated in Figure 8, in the event that the order is correct, the system's functionality operates within a stable boundary, and the tool selection will not result in significant fluctuations. Conversely, when the order is reversed, the functionality exhibits a pronounced decline, and the utilization of tools to mitigate this decline is not feasible. Therefore, the optimal tool is anchored in preference for Q to complete the rollback in O according to p_{order} first and then T according to p_{tool}.

B Additional Experiments

B.1 Demonstration from compatibility to adaptability

The following experiments are conducted on the FoundIR (Li et al., 2025) dataset as in Table 2. We partition each degradation coupling scenario into 20 images allocated for the self-evolving process and the remaining 30 images reserved for inference testing.

B.1.1 Degradation and Tool

Our method still follows the training-free IRAs, which are compatible with new degradations. If a tool is provided that can handle the degradation, the agentic system can be compatible without training. Furthermore, IRA can be regarded as a system of tool ensembles that is compatible with multiple new tools for one degradation.

With the assistance of the self-evolving mechanism, the evolving priority-driven mapping function p allows the degradation and tool compatibility to extend to adaptability. We design six variants:

(1) **S0:** Using Tool Setting I during the evolution phase from scratch and providing only four types (total of six) of degradation tools.

(2) **S1:** Injecting two additional degradations with only one tool from Tool Setting I into S0, which presents compatibility with new degradations.

(3) **S2:** Finish the evolution process in the S1, which present the adaptability to new degradation.

(4) **S3:** Inject new tools from Tool Setting II into S2, which presents compatibility with new tools, and tools will only be added at the end of the priority function.

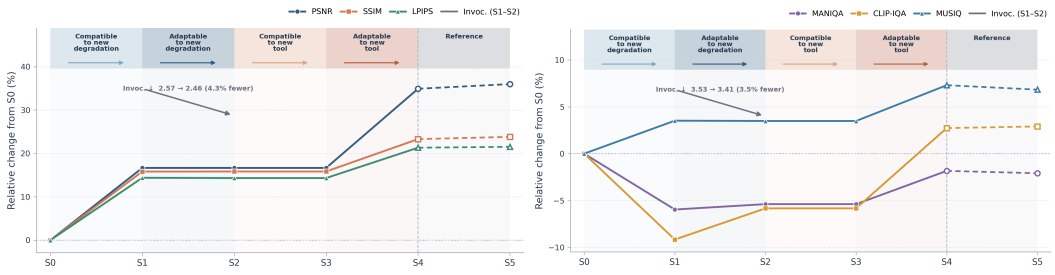
(5) **S4:** Finish the evolution process in the S3, which present the adaptability to new tool.

(6) **S5:** Use only Tool Setting II during the evolution phase, which is considered the reference standard.

As demonstrated by Table 4 from S0 to S4, our method possesses the capability to adapt to new degradations and tools in addition to the compatibility with previous IRAs. Furthermore, the adaptability enables to approach the theoretical state in terms of performance and efficiency when compared to S5.

Table 4: A study on degradation and tool adaptability with **E^{2}IR-Agent**. The reference standard is highlighted in gray.

State	Fidelity					Perception				
	Quality Metrics			Efficiency Metrics		Quality Metrics			Efficiency Metrics	
	PSNR ↑	SSIM ↑	LPIPS ↓	Invoc. ↓	Total Rb ↓	MANIQA ↑	CLIP-IQA ↑	MUSIQ ↑	Invoc. ↓	Total Rb ↓
$S0$	15.31	0.5699	0.4240	1.59	0.23	0.3084	0.3653	52.92	2.41	0.49
$S1$	17.86	0.6600	0.3630	2.57	0.39	0.2900	0.3318	54.79	3.53	0.71
$S2$	17.86	0.6601	0.3633	2.46	0.33	0.2918	0.3440	54.77	3.41	0.76
$S3$	17.86	0.6601	0.3633	2.46	0.33	0.2918	0.3440	54.77	3.41	0.76
$S4$	18.12	0.6764	0.3568	3.27	0.79	0.3029	0.3564	55.01	5.42	1.27
$S5$	20.66	0.7027	0.3336	3.10	0.83	0.3028	0.3753	56.79	3.53	0.76



(a) Fidelity (b) Perception

Figure 9: Visualization of the adaptability on degradation and tool.

Furthermore, our method inherently supports the retirement of both degradation types and tools.

Specifically, we adopt a soft-registration mechanism for different degradations, ensuring that inactive degradations are not loaded into GPU memory. Moreover, experiences of distinct degradations are stored independently within the experience pool, granting them seamless eliminability.

As for tools, benefiting from the decoupling of p_{tool} and p_{order} in the experience representation, any retired tool is simply excised from the tool priority function p, with subsequent candidates shifting forward to fill the vacancy without adverse impacts.

B.1.2 Data

Table 5 demonstrates the data source adaptability of E^{2}IR-Agent when confronted with entirely different data domains. To systematically analyze this adaptability, we design three variants:

(1) **M1**: directly applies the experience gathered from the MiO100 dataset to the out-of-domain (OOD) scenario for zero-shot testing.

(2) **M2:** leverages the prior experience and continues to self-evolve using the 20 images from the OOD dataset.

(3) **M3:** initializes the self-evolving process entirely from scratch using 20 OOD images, devoid of any prior experience.

As illustrated in the experimental results, M1 and M3 provided ideal states in their respective data domains. Considering the overall performance and efficiency, M3 outperformed M1, which is in line with expectations. Meanwhile, M2 clearly demonstrates our method's ability to adapt to data, showing an overall trend moving from M1 towards M3.

Table 5: A study on data adaptability with **E^{2}IR-Agent**. The best results are highlighted in **bold**.

Methods	Fidelity					Perception				
	Quality Metrics			Efficiency Metrics		Quality Metrics			Efficiency Metrics	
	PSNR ↑	SSIM ↑	LPIPS ↓	Invoc. ↓	Total Rb ↓	MANIQA ↑	CLIP-IQA ↑	MUSIQ ↑	Invoc. ↓	Total Rb ↓
$M1$	20.12	0.6959	0.3401	3.50	1.09	**0.3276**	**0.3761**	56.28	5.29	1.36
$M2$	20.42	**0.7076**	0.3425	3.20	0.77	0.3178	0.3713	55.31	4.93	1.19
$M3$	**20.82**	0.7056	**0.3328**	**2.76**	**0.52**	0.3020	0.3759	**56.54**	**3.60**	**0.80**

B.2 Ablation Study on Hierarchical Experience Pool

Table 6 and Figure 10 detail the ablation study of the impact of experience guidance at different granularities with $B = 25$ and $Times = 2$ on Group B of the MiO100. To ensure consistently processed images, this experiment is conducted on degradation types D, which inherently necessitates fine-grained experience.

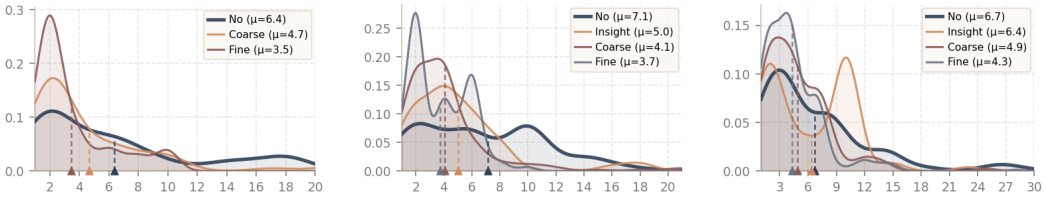
As demonstrated, the absence of prior experience forces the system into frequent rollback. The introduction of *Insight* experience provides a global fallback mechanism, substantially reducing computational overhead. Efficiency is further enhanced as the granularity refines to *Coarse* and *Fine* levels, while the quality scores remain comparable.

While *Fine-grained* experience offers the highest theoretical efficiency, relying on it may introduce three risks: (i) **Semantic Dependency:** Fine-grained experience relies heavily on the visual perception capability of MLLMs. If the comprehension of MLLM is limited, the *Pattern Profile Learning System* (Sec. 3.4) may fail to generate accurate profiles. (ii) **Visual Ambiguity:** For degradations lacking salient visual features, the image encoder (Huang et al., 2026b) struggles to extract discriminative representations. This results in erroneous experience retrieval. (iii) **API Overhead:** Relying on fine-grained experience will incur an API cost, as shown in Table 13.

Conversely, *Coarse-grained* experience provides a more robust mapping by BTD statistically. Therefore, our hierarchical design achieves an optimal equilibrium: *Coarse-grained* experience ensures generalizable stability, while *Fine-grained* experience drives peak inference efficiency.

Table 6: Detailed ablation study on experience granularity on MiO100 dataset with **E^{2}IR-Agent**.

Degrad.	Methods	Quality Metrics						Efficiency Metrics			
		PSNR ↑	SSIM ↑	LPIPS ↓	MANIQA ↑	CLIP-IQA ↑	MUSIQ ↑	Invoc. ↓	$\mathcal{O}$-Rb ↓	$\mathcal{T}$-Rb ↓	Total Rb ↓
GroupA	w/o experience	22.12	0.7451	0.2622	0.3438	0.5035	61.63	6.38	1.07	2.20	3.27
	w/ coarse	22.41	**0.7476**	**0.2474**	0.3503	0.5054	**62.86**	4.67	0.79	1.10	1.89
	w/ fine	**22.42**	0.7464	0.2495	**0.3508**	**0.5080**	62.73	**3.45**	**0.59**	**0.81**	**1.40**
GroupB	w/o experience	21.02	0.7923	0.2733	**0.4081**	**0.5858**	63.21	7.13	0.86	2.88	3.74
	w/ insight	21.56	0.8103	0.2604	0.3933	0.5408	64.58	5.00	1.00	1.61	2.61
	w/ coarse	22.20	0.8134	0.2620	0.3913	0.5317	64.55	4.08	0.78	0.68	1.46
	w/ fine	**22.44**	**0.8165**	**0.2576**	0.3941	0.5256	**64.73**	**3.72**	**0.67**	**0.66**	**1.33**
GroupC	w/o experience	21.20	0.5745	0.5364	0.1833	0.2399	35.98	6.69	1.04	1.94	2.98
	w/ insight	21.96	0.5801	0.5770	0.1968	0.2811	37.48	6.37	**0.62**	2.35	2.97
	w/ coarse	**22.15**	**0.5943**	0.5130	**0.2037**	**0.2865**	40.57	5.04	1.01	1.11	2.12
	w/ fine	**22.15**	0.5942	**0.5123**	0.2031	0.2850	**41.09**	**4.32**	0.85	**0.89**	**1.74**



(a) GroupA. (b) GroupB. (c) GroupC.

Figure 10: A visualization of the frequency distribution of tool invocation numbers on different experience granularity among GroupA-C

B.3 VISUALIZATION ON HIERARCHICAL EXPERIENCE

In this section, we will present several sample sections (insight: Table 7; coarse-grained: Table 8; fine-grained: Table 9) of the hierarchical experience pool to facilitate comprehension. Furthermore, a graph will be constructed to visualize the BTD statistical model underlying the findings, as shown in Figure 11 and Figure 12.

Table 7: Experience sample of **insight**

Preference	Experience
Fidelity	Here is the reference information from past trials: A reasonable overall elimination order that respects all strong constraints and most weaker ones is: JPEG artifacts → Defocus blur → Haze → Rain → Motion blur → Brightening → Denoising → Super-resolution.
Perception	Here is the reference information from past trials: Recommended "default" sequence, JPEG compression artifacts removal → Defocus deblurring → Dehazing → Deraining → Super-resolution → Motion deblurring → Denoising (only confidently "after JPEG" and "before brightening"

Table 8: Experience sample of **coarse-grained experience**

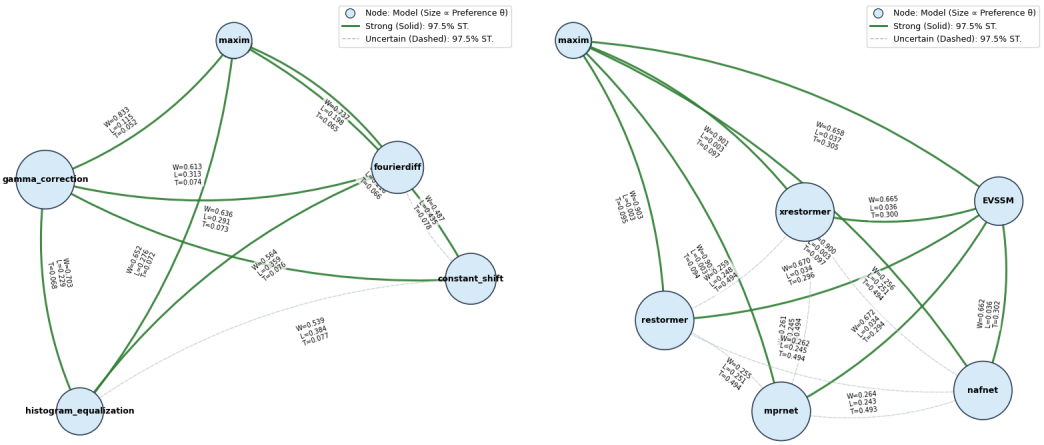
Degradation	Preference	Experience

Motion Blur	Fidelity	`"degradation_type": "motion blur",`
		`"preference": "fidelity",`
		`"ranking": {`
		`"mprnet": 1,`
		`"restormer": 2,`
		`"xrestormer": 3,`
		`...`

	Perception	`"degradation_type": "motion blur",`
		`"preference": "perception",`
		`"ranking": {`
		`"xrestormer": 1,`
		`"restormer": 2,`
		`"EVSSM": 3,`
		`...`

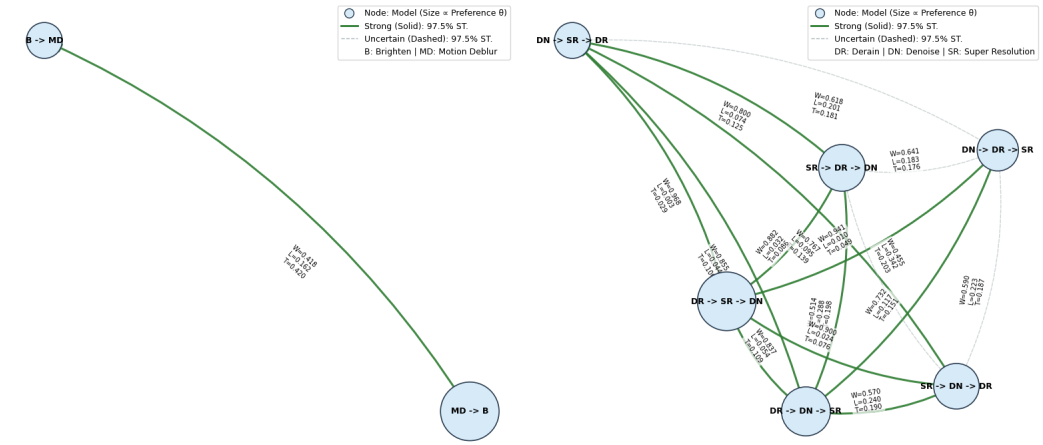
D+M	Fidelity	`"degradation_type": "dark+motion blur",`
		`"preference": "fidelity",`
		`"ranking": {`
		`"motion blur -> dark": 1,`
		`"dark -> motion blur": 2}`

	Perception	`"degradation_type": "dark+motion blur",`
		`"preference": "perception",`
		`"ranking": {`
		`"dark -> motion blur": 1,`
		`"motion blur -> dark": 2}`



(a) BTD Complete Graph on 'Dark' (b) BTD Complete Graph on 'Motion Blur'

Figure 11: Visualization of the BTD Model result in *Tool Selection*. The size of the nodes corresponds to the ability parameter as estimated by the BTD model. In instances where the coarse-grained experience is uncertain (represented by a dashed line), thereby signifying that fine-grained experience must be utilized.



(a) BTD Complete Graph on 'Dark+Motion Blur' (b) BTD Complete Graph on 'Rain+Haze+LR'

Figure 12: Visualization of the BTD Model result in *Degradation Removal Order*.

Table 9: Experience sample of **fine-grained experience**

Degradation	Preference	Experience
Motion Blur	Fidelity	{ "exp_id": 0, "degradation_type": "motion blur", "preference": "fidelity", "degradation_pattern": "Significant motion...", "ranking": { "xrestormer": 1, "mprnet": 2, ... }, "related_trajectory_ids": [0, 6, ...] }, { "exp_id": 1,...
D+M	Perception	{ "exp_id": 0, "degradation_type": "dark+motion blur", "preference": "perception", "degradation_pattern": "Significant loss...", "ranking": { "motion blur -> dark": 1, "dark -> motion blur": 2 }, "related_trajectory_ids": [0, 1, ...] }, { "exp_id": 1,...

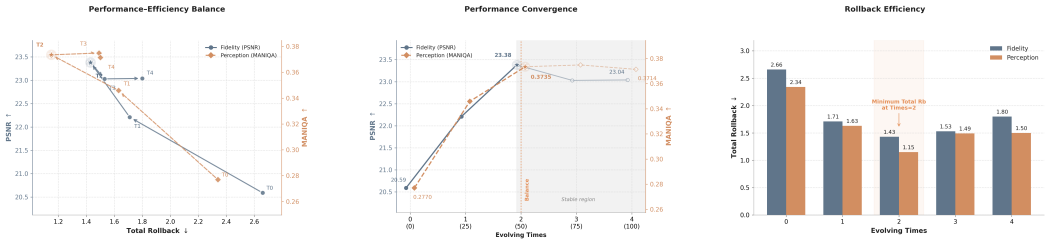
B.4 Ablation Study on Evolution Times

Table 10 and Figure 13 detail the ablation study of the self-evolving experience mechanism on Group B of the MiO100 dataset with batch size $B = 25$. By scanning the evolution times, we seek to investigate the optimal balance point.

Overall, Times $= 2$ emerges as the optimal sweet spot between performance and efficiency. Now we detail the analysis: (i) **Cost-Quality:** As illustrated in (a), the system's evolution exhibits a convergent trajectory from the initial condition of Times $= 0$ to the condition of Times $= 2$, which reaches the Pareto-optimal point in both *Preference* at Times $= 2$ (top-left corner). However, for Times > 2, the curve folds back towards the bottom-right, indicating potential experience redundancy. (ii) **Quality Progression:** The findings in (b) indicate that both *Preference* exhibit a marked increase in growth during the initial phase (Times < 2) and attains its peak level of growth at the Times $= 2$ mark. The plateau and slight fluctuations show diminishing returns in the subsequent phase (Times > 2). (iii) **Efficiency Measurement:** Furthermore, the efficiency measurement in (c) explains that the substantial decrease in total rollback times stop at Times $= 2$.

Table 10: Ablation study on evolving times with **E²IR-Agent**. The best results are highlighted in **bold**.

Times	$B = 25$	Fidelity				Perception			
		Quality Metrics			Efficiency Metrics	Quality Metrics			Efficiency Metrics
		PSNR ↑	SSIM ↑	LPIPS ↓	Total Rb ↓	MANIQA ↑	CLIP-IQA ↑	MUSIQ ↑	Total Rb ↓
0	0	20.59	0.6579	0.3287	2.66	0.2770	0.4603	56.69	2.34
1	25	22.21	0.7346	0.3430	1.71	0.3459	0.5091	61.26	1.63
2	50	**23.38**	0.7643	0.3117	**1.43**	0.3735	0.5296	61.96	**1.15**
3	75	23.03	0.7539	0.3205	1.53	**0.3749**	0.5180	**62.44**	1.49
4	100	23.04	**0.7647**	**0.3029**	1.80	0.3714	**0.5386**	61.79	1.50



(a) Quality-Efficiency Trade-off. (b) Quality Progression. (c) Efficiency Measurement.

Figure 13: Performance-Efficiency analysis of the evolution process with batch size B=25. (a) Cost-quality Pareto front between performance and efficiency in both *Preference*. (b) Quality progression over evolution times. (c) Efficiency measurement by the total rollback times.

B.5 Ablation Studies on Rollback Mechanism and Preference Setting

Evaluation on rollback mechanism. The left of Table 11 illustrates that, in nearly all instances, the metrics with the rollback process are superior. Notably, the rollback process differs from (Zhu et al., 2025; Zuo et al., 2025), as our method is structured around the experience pool.

Evaluation on preference setting. The right of Table 11 clearly demonstrates that fidelity consistently achieves superior FR metrics, while *Perception* dominates in NR metrics across three groups in the MiO100 dataset. It is noteworthy that even when prioritizing perception preference, its FR metrics demonstrate superiority over those of existing state-of-the-art methods.

Table 11: Ablation study on rollback mechanism and preference design with **E²IR-Agent**.

Degrad.	Rollback Mechanism							Preference Setting						
	Method	PSNR ↑	SSIM ↑	LPIPS ↓	MANIQA ↑	CLIP-IQA ↑	MUSIQ ↑	Dim	PSNR ↑	SSIM ↑	LPIPS ↓	MANIQA ↑	CLIP-IQA ↑	MUSIQ ↑
Group A	w/o Rb.	23.38	0.7410	0.3177	**0.3660**	**0.5214**	61.59	F	**23.86**	**0.7488**	0.3085	0.2833	0.4190	54.37
	w Rb.	**23.86**	**0.7488**	**0.3085**	0.3623	0.5204	**61.82**	P	22.31	0.7103	**0.2831**	**0.3623**	**0.5204**	**61.82**
Group B	w/o Rb.	22.93	0.7582	0.3162	**0.3814**	**0.5436**	61.62	F	**23.38**	**0.7643**	0.3117	0.2727	0.4247	53.28
	w Rb.	**23.38**	**0.7643**	**0.3117**	0.3735	0.5296	**61.96**	P	22.52	0.7392	**0.2694**	**0.3735**	**0.5296**	**61.96**
Group C	w/o Rb.	21.98	0.6576	0.4682	0.3278	0.4930	52.96	F	**22.12**	**0.6588**	0.4528	0.2032	0.3201	41.36
	w Rb.	**22.12**	**0.6588**	**0.4528**	**0.3321**	**0.4977**	**55.21**	P	20.69	0.6004	**0.4072**	**0.3321**	**0.4977**	**55.21**

B.6 Scalability Analysis on Complex Scenarios

To investigate the scalability in extreme degradation scenarios, we design an exploratory study on two challenging 4-degradation combinations (*Defocus + Haze + JPEG + Rain* and *Dark + Haze + JPEG + Motion Blur*). Specifically, during the experience evolution phase, we introduce a stochastic trajectory dropping strategy with drop rates ranging in $\{0.00, 0.25, 0.50, 0.75\}$ on equation 5.

As summarized in Table 12, a noticeable performance decay surfaces when the drop rate reaches 0.75, where the system explores merely 6 out of 24 possible permutation trajectories to update the priority function p. This demonstrates that our priority modeling exhibits strong robustness against sparse exploration trajectories in more complex cases.

Insight on Cascading Error Accumulation. Crucially, we observe that the overall quantitative scores across all drop rates and one baseline method remain relatively modest. This reveals an intrinsic challenge in long-horizon IRA tasks: in highly entangled degradation chains, the primary bottleneck is *cascading error accumulation* caused by chained sequential tool invocations. Consequently, improving search efficiency alone during the experience evolution stage is insufficient to address extreme combinations; developing error-mitigating inference mechanisms is indispensable for long-horizon agentic restoration, which we envision as our future research focus.

Table 12: Performance across four degradation combinations under varying trajectory drop rates during experience evolution.

Method	Drop Rate	Sampled Traj.	Fidelity			Perception		
			PSNR↑	SSIM↑	LPIPS↓	MANIQA↑	CLIP-IQA↑	MUSIQ↑
AgenticIR	–	–	15.86	0.5055	0.5575	0.1700	0.2635	38.62
E^2IR-Agent	0.00	24 / 24	17.58	0.5379	0.5891	0.1739	0.2586	35.03
	0.25	18 / 24	17.56	0.5397	0.5863	0.1743	0.2621	35.09
	0.50	12 / 24	17.57	0.5405	0.5858	0.1739	0.2598	35.21
	0.75	6 / 24	17.40	0.5326	0.5934	0.1779	0.2623	35.62

B.7 More Qualitative Results

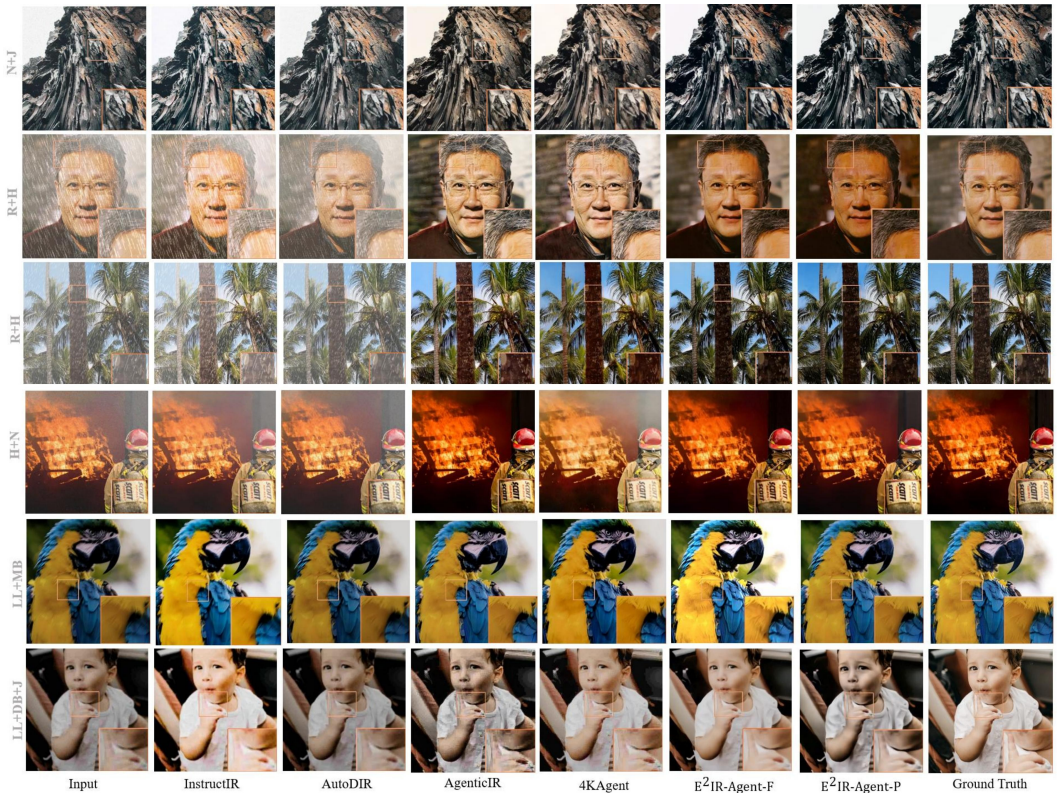


Figure 14: More visual comparison examples on MiO100 dataset. **E^2IR-Agent-F** denoted the fidelity *Preference* requirement, while **E^2IR-Agent-P** denoted the perception *Preference* requirement.

C Implementation Details

C.1 Datasets Detail

Synthetic Dataset: We adhere to the experimental protocol outlined by AgenticIR (Zhu et al., 2025). Our methodology is data-driven, and to ensure the isolation of the dataset for test-time adaptation, we require data for the purpose of experience evolution.

- Experience Evolution Set: To construct our experience pool, we utilize a separate, disjoint set of 3,450 high-quality images provided by MiOIR (Kong et al., 2024). We leverage the degradation generation pipeline from AgenticIR to generate degraded-gt image pairs for the evolution process.
- Evaluation Set: All our evaluations are conducted purely on the MiO100 dataset, which contains totally 1,440 images with 16 degradation coupling scenarios. Group A contains 8 groups with 2 degradation, Group B contains 4 groups with 2 degradation and Group C contains 4 groups with 3 degradation.

Real-world Dataset: We are interested in degradation coupling scenarios that are well-suited to the main contributions of the IRAs. To this end, we utilize three groups from a real-world dataset, FoundIR (Li et al., 2025) for evaluation: motion blur+JPEG, rain+haze, and dark+noise. Each group contains 50 degraded-gt pairs. Unless otherwise specified, this dataset is used for zero-shot generalization evaluation.

C.2 Tool Set

We leverage the methods in (Zhu et al., 2025) during the *Perception* and *Reflection* processes. The entire experiment is conducted with two NVIDIA A6000 GPUs. Since the present framework is flexible to different tools, this property is examined in two tool sets. All the ablation studies are conducted on setting II without specification.

In the **E^{2}IR-Agent**, the configuration of the tool settings in previous agents (Zhu et al., 2025; Jiang et al., 2025b; Lu et al., 2025) is adhered to.

- **Brightening:** CLAHE (Reza, 2004), Gamma correction ($\gamma = 2/3$), constant shift.
- **Defocus Deblurring:** DRBNet (Ruan et al., 2022), IFAN (Lee et al., 2021), Restormer (Zamir et al., 2022).
- **JPEG Artifact Removal:** SwinIR (Liang et al., 2021) (quality factor 40), FBCNN (Jiang et al., 2021) (quality factors 90 and 5, and blind setting).
- **Denoising:** SwinIR (Liang et al., 2021), MAXIM (Tu et al., 2022), MPRNet (Zamir et al., 2021), Restormer (Zamir et al., 2022), X-Restormer (Chen et al., 2024b).
- **Deraining:** MAXIM (Tu et al., 2022), MPRNet (Zamir et al., 2021), Restormer (Zamir et al., 2022), X-Restormer (Chen et al., 2024b).
- **Motion Deblurring:** MAXIM (Tu et al., 2022), MPRNet (Zamir et al., 2021), Restormer (Zamir et al., 2022), X-Restormer (Chen et al., 2024b).
- **Dehazing:** MAXIM (Tu et al., 2022), X-Restormer (Chen et al., 2024b), RIDCP (Wu et al., 2023), DehazeFormer (Song et al., 2023).
- **Super Resolution:** DiffBIR (Lin et al., 2024), HAT (Chen et al., 2025b), SwinIR (Liang et al., 2021), X-Restormer (Chen et al., 2024b).

In the **E^{2}IR-Agent-II**, the tool settings from prior agents (Zuo et al., 2025; Zhang et al., 2026) are utilized.

- **Brightening:** CLAHE (Reza, 2004), Constant Shift (C=40), DiffPlugin (Liu et al., 2024), FourierDiff (Lv et al., 2024), Gamma Correction ($\gamma = 2/3$), MAXIM (Tu et al., 2022).
- **Defocus Deblurring:** ConvIR (Cui et al., 2024a), DiffPlugin (Liu et al., 2024), DRBNet (Ruan et al., 2022), IFAN (Lee et al., 2021), LaKDNet (Cui et al., 2024b), Restormer (Zamir et al., 2022).

- **JPEG Artifact Removal:** FBCNN (Jiang et al., 2021) (QF=5), FBCNN (Jiang et al., 2021) (QF=90), FBCNN (Jiang et al., 2021) (BQF), SwinIR (Liang et al., 2021) (QF=40).
- **Denoising:** MAXIM (Tu et al., 2022), MPRNet (Zamir et al., 2021), NAFNet (Chen et al., 2022), Restormer (Zamir et al., 2022), X-Restormer (Chen et al., 2024b), SwinIR (Liang et al., 2021).
- **Deraining:** DiffPlugin (Liu et al., 2024), MAXIM (Tu et al., 2022), MPRNet (Zamir et al., 2021), Restormer (Zamir et al., 2022), X-Restormer (Chen et al., 2024b).
- **Dehazing:** DehazeFormer (Song et al., 2023), DiffPlugin (Liu et al., 2024), MAXIM (Tu et al., 2022), RIDCP (Wu et al., 2023), X-Restormer (Chen et al., 2024b).
- **Motion Deblurring:** EVSSM (Kong et al., 2025), LaKDNet (Cui et al., 2024b), MAXIM (Tu et al., 2022), MPRNet (Zamir et al., 2021), NAFNet (Chen et al., 2022), Restormer (Zamir et al., 2022), X-Restormer (Chen et al., 2024b).
- **Super Resolution:** DiffBIR (Lin et al., 2024), DRCT (Hsu et al., 2024), HAT-L (Chen et al., 2025b), HAT-GAN, HMA (Chu et al., 2024), OSEDiff (Wu et al., 2024b), PiSA-SR (Sun et al., 2025), SwinIR (Liang et al., 2021), SwinIR (Real-ISR), SwinFIR (Zhang et al., 2022), X-Restormer (Chen et al., 2024b).

C.3 Experience Pool Setting

Implementation Details & Hyperparameters. Our framework is implemented with the following hyperparameter configurations:

- **Retrieval & Evolution:** The Top-K recall (Sec. 3.1) is set to 3. For experience evolution (Sec. 3.4), we set the trigger threshold batch size B to 25, the significance level α to 0.975, and the mini-batch size to 12.
- **IQA Metrics:** The metrics set for *Fidelity* consist of PSNR, SSIM, LPIPS, and DISTS. The *Perception* metrics set, utilized to support the *Preference* in Sec. 3.2, includes MANIQA, MUSIQ, BRISQE, CLIPIQA, NIQE, and NIMA.

MLLM Instantiation. The estimated API cost per round for the agent-driven pattern profile learning system is reported in Table 13.

- **Base Inference:** Qwen3-VL-Flash (Yang et al., 2025) is adopted for standard, general-purpose multimodal operations.
- **Experience Evolution:** To ensure deep image interpretation and high-quality description generation during fine-grained experience evolution, we employ the more capable Qwen3-Plus (Yang et al., 2025) foundation model.

Table 13: Detailed API usage and estimated cost for formulating fine-grained experience within a single mini-batch in **E^{2}IR-Agent**.

Model Name	Calls / Batch	Avg. Tokens / Call		Total Cost per Batch ($)
		Input	Output	
Qwen3-VL-Flash	∼35	∼950	∼150	0.0189
Qwen3-Plus	∼12	∼1,300	∼400	0.0931
Total	**47**	–		**0.1120**

C.4 PROMPTS

Insight Evolving Prompt

```
Based on the following Bradley-Terry-Davison (BTD) probabilistic
    model results for multi-degradation image restoration tasks (
    Preference: {preference}), extract key insights regarding the
    degradation removal order and recommendations:

{combined_text}

Your conclusions and insights:
```

Multi-Agent Debate _Role_ Prompt

```
Your role is to act as the {role}. The theme is: "Is this a good
    enough degradation pattern?". This means you are skilled at
    processing information from both yourself and others to
    determine the next appropriate action without executing it. You
    are currently in the thinking stage.

Historical information (your current-stage reasoning & external tool
     results): {context}.

You have access to external tools represented as Python function
    prototypes, which you may invoke to support your debate
    competition:
- generate_groups() # Used to generate image degradation patterns
    solely from textual information in the `[Start Round]`.
- validate_current_group(traj_id: List[int]) # Used to validate the
    current unconvinced image degradation pattern (intra-group:
    within the same degradation pattern) using specified trajectory
    IDs. (Requirements: 1 unconvinced traj_id + 1~3 convinced
    traj_ids; maximum of 3 traj_ids in total)
- validate_other_group(traj_id: List[int]) # Used to validate
    another unconvinced image degradation pattern (inter-group:
    between two degradation patterns) using specified trajectory IDs.
     (Requirements: 1 unconvinced traj_id + 1~3 convinced traj_ids;
    maximum of 3 traj_ids in total)
- finish() # Used when you think you have obtained the final answer.

Your response must strictly follow this format. It is forbidden to
    include any output beyond these two sections:
Thought: [Your brief reasoning process (under 200 words), including
    problem analysis and determination of the next action by
    carefully reviewing the historical information from your
    opponent and your current-stage reasoning.]
Action: [The Python function call you decide to execute in the next
    step]
```

Multi-Agent Debate _Action_ Prompt

```
You want to ensure that all degradation patterns you propose are
    credible and accurate enough to withstand challenges in a debate
     competition. The theme is: `Is this a good enough degradation
    pattern?`.
```

```
Currently, there is an ongoing process defined as "semantic
    clustering".
The goal of semantic clustering is to categorize degradation modes
    according to the following rules:
1. Within each degradation pattern category, the semantic
    representations of the degradation pattern text must be as
    similar as possible.
2. Between different degradation pattern categories, the semantic
    representations of the degradation pattern text must be
    significantly distinct.
3. Under the premise of satisfying Rules 1 and 2, the tool ranking
    information of the top three (or top 2, if only two exist)
    within each category must be as close as possible.
4. Under the premise of satisfying Rules 1 and 2, the tool ranking
    information of the top three (or top 2, if only two exist)
    between different categories must be as distant as possible.
5. The number of degradation pattern categories is flexible, but it
    must avoid redundancy and satisfy Rules 1 to 4.
6. For uncertain trajectory information that cannot be confidently
    classified considering both semantic representation and ranking
    requirements, explicitly specify which trajectories are
    uncertain and require further verification inside `<
    unconvinced_info></unconvinced_info>` or `<verify_1></verify_1>`.


You are provided with the current degradation pattern categories
    derived solely from textual information: ...
Some categories contain trajectories with uncertain classifications.
     Now, you will perform a further validation step using the image
     information provided as: ...

You can provide two types of judgments:
1. Positive: Confirm that the trajectory fits the degradation
    pattern category based on the image data.
2. Negative: Determine that the trajectory does not fit the current
    degradation pattern category; instead, select a more appropriate
     category based on textual semantic similarity and specify that
    cross-category verification is needed.

Finally, you must retain all previous outputs inside `<
    convinced_info></convinced_info>` (as well as existing `<
    unconvinced_info></unconvinced_info>` and `<verify_n></verify_n
    >` tags) without modifying or deleting them. For the current
    verification, append your result inside an XML tag with an
    incremental sequence number (e.g., `<verify_1></verify_1>`).

Your output must strictly follow this format:
<convinced_info>Degradation Pattern 1: The pure semantic
    representation of the degradation pattern text information is...
     (Aspect 1: Semantic representation of the degradation pattern,
    without tool names or removal priorities, at least 15 words).
    Included trajectories are Traj1, Traj3..., reasons are: (Combine
     three aspects: Semantic representation of the degradation
    pattern & Top-3 removal priority rankings [Top-2 if only two] &
    Tool names)...</convinced_info><unconvinced_info>Among these
    trajectories, the classifications of Traj1, ... are not fully
    convincing because..., requiring further verification.</
    unconvinced_info><verify_1>After combined text and image
    validation, Traj3 indeed belongs to Degradation Pattern 1,
    because</verify_1>
```

Pattern Profile Operation *Function* Prompt

```
You are an expert in planning and knowledge management, tasked with
    updating degradation patterns stored within a vector database.
    Your objective is to analyze a given set of "new degradation
    patterns" alongside the "existing degradation patterns"
    currently in the database, and determine the precise atomic
    operations required to keep the database accurate and up-to-date.


Specifically, for each new pattern, decide which of the following
    actions applies:
- add: Introduce a completely new pattern to the database because it
     does not overlap with existing ones. Use this operation to
    balance accuracy and redundancy.
- merge: Apply when an existing pattern fully subsumes a new pattern
    . Retain the existing pattern but enrich it with additional
    information from the new one.
- replace: Apply when a new pattern subsumes an existing pattern,
    meaning the new pattern's semantics are broader or more accurate
    , and thus should replace the existing pattern.

You will receive a task description that specifies the particular
    patterns to analyze. If you receive any feedback on your output,
     revise your operations accordingly while maintaining the same
    output format.

Current task description:
Degradation type: {degradation_type}
New degradation patterns: {new_dp}

Existing degradation patterns in DB:
<none>

Previous plans:
<none>

Previous feedback:
<none>

Your output must be formatted as a Python list of strings, where
    each string represents a single atomic operation. It is strictly
     forbidden to perform multiple operations on the same existing
    degradation pattern. If an operation on a pattern is
    unsatisfactory, choose an alternative degradation pattern.

The format for each element is:
"<new pattern digit only> | <action> | <existing pattern digit only
    (optional)>"

- Include the existing pattern digit only when relevant (for merge
    or replace actions).
- For add actions, the existing pattern field must be omitted.


Example output:
["1 | merge | 2", "2 | add"]
```