SELF-IMPROVED PRIOR FOR ALL-IN-ONE IMAGE RESTORATION

Anonymous authors
Paper under double-blind review

ABSTRACT

Unified image restoration models for diverse and mixed degradations often suffer from unstable optimization dynamics and inter-task conflicts. This paper introduces Self-Improved Privilege Learning (SIPL), a novel paradigm that overcomes these limitations by innovatively extending the utility of privileged information (PI) beyond training into the inference stage. Unlike conventional Privilege Learning, where ground-truth-derived guidance is typically discarded after training, SIPL empowers the model to leverage its own preliminary outputs as *pseudo-privileged signals* for iterative self-refinement at test time. Central to SIPL is Proxy Fusion, a lightweight module incorporating a learnable Privileged Dictionary. During training, this dictionary distills essential high-frequency and structural priors from privileged feature representations. Critically, at inference, the *same* learned dictionary then interacts with features derived from the model's initial restoration, facilitating a self-correction loop. SIPL can be seamlessly integrated into various backbone architectures, offering substantial performance improvements with minimal computational overhead. Extensive experiments demonstrate that SIPL significantly advances the state-of-the-art on diverse all-in-one image restoration benchmarks. For instance, when integrated with the PromptIR model, SIPL achieves remarkable PSNR improvements of +4.58 dB on composite degradation tasks and +1.28 dB on diverse five-task benchmarks, underscoring its effectiveness and broad applicability.

1 INTRODUCTION

All-in-one image restoration, which aims to tackle diverse and often mixed degradations with a single, unified model, has emerged as a pivotal research area due to its immense practical value Jiang et al. (2025). However, these versatile models confront a fundamental dilemma: forcing a single network to master the distinct, often confusing or conflicting, feature representations required for tasks like denoising (local textures) and dehazing (global context) inevitably leads to optimization challenges and performance compromises Kong et al. (2024); Wu et al. (2024). This heterogeneity brings unstable training process, resulting in the suboptimal solution. To address this challenge, we argue for a paradigm shift against the end-to-end learning. We propose to explicitly learn and leverage a universal prior of what constitutes a high-quality, degradation-free image, using it as a stable guide to navigate the complex optimization landscape.

To realize this, we draw inspiration from Privilege Learning (PL) Vapnik and Vashist (2009); Vapnik et al. (2015), where auxiliary information is used to guide training progressively while can be unavailable at test time. We posit that ground-truth (GT) images can serve as the ultimate privileged information, providing a clear, degradation-agnostic supervisory signal. The incorporation of privileged information during training establishes an *inter-task comprehension bridge*, effectively harmonizing conflicting gradients and stabilizing convergence progressively. For instance, a PL-enhanced PromptIR model achieves performance comparable to the original baseline in less than half the training epochs (see Figure 1(a)), showcasing that PL therefore presents an efficient and straightforward strategy for improving existing all-in-one restoration methods. Nevertheless, vanilla PL traditionally limits its own impact, as the privileged guidance is often progressively diminished during training and is entirely unavailable at inference time.

This raises a critical question: *can the essence of privileged information be retained and repurposed to enhance performance at the inference stage?* We answer in the affirmative by proposing the novel

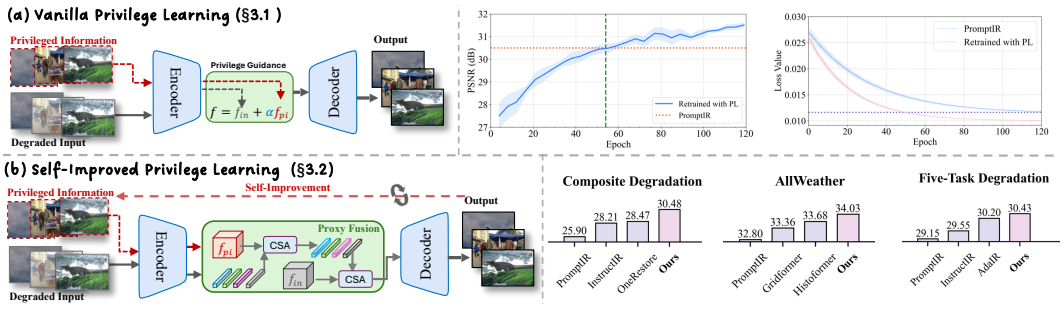

Figure 1: Conceptual comparison of learning frameworks: (a) Privilege Learning (PL) leverages privileged information during training for improved optimization. (b) Our proposed Self-Improved Privilege Learning (SIPL) framework introduces Proxy Fusion to retain privileged knowledge, enabling iterative self-refinement during inference by using intermediate restoration outputs as pseudo-privileged information. Retrained PromptIR with the proposed SIPL achieves significant improvement across diverse all-in-one tasks.

Self-Improved Privilege Learning (SIPL) framework, which extends the conventional PL paradigm in two key stages, as illustrated in Figure 1(b). First, we introduce a the Privileged Dictionary (PD) to **distill and retain universal high-quality priors** from privileged information during training. Second, we propose an iterative self-refinement cycle to **reuse these retained priors** at inference, where the model's own initial restoration serves as *pseudo-privileged information* for a powerful self-refinement loop.

In detail, this two-stage concept is realized by our proposed *Proxy Fusion* mechanism, a straightforward and efficient module centered around a learnable **Privileged Dictionary (PD)**. During training, the PD internalizes degradation-agnostic image characteristics from GT-derived features, creating a compact and persistent knowledge base. At inference, the PD's role transforms from a student to a teacher. It interacts with features extracted from the model's own preliminary restoration (the pseudo-privileged information). This interaction allows the immutable, pre-trained priors within the PD to guide a feature-level correction, progressively enhancing the output in a self-driven manner. The efficacy of this approach is evident when integrating SIPL with the PromptIR model, which yields remarkable PSNR gains of **+4.58 dB** on Composite Degradation, **+1.23 dB** on Allweather, and **+1.38 dB** on the Five-Task benchmark, as illustrated in Figure 1(b).

Our contributions can be summarized into four aspects:

- We introduce Privilege Learning (PL) to all-in-one image restoration, achieving a stronger and more stable optimization baseline by effectively mitigating inter-task conflicts.
- We extend the conventional PL framework with a novel mechanism to **retain** privileged knowledge for the inference stage, thus overcoming the ephemeral nature of guidance.
- We further propose a self-refinement strategy to **reuse** this retained knowledge. This empowers models with a new capability for iterative self-improvement, leading to promising performance boosts.
- We propose the **Proxy Fusion**, a lightweight and plug-and-play module that makes SIPL a practical reality. Its high efficiency and broad compatibility are validated by significant performance gains on multiple state-of-the-art architectures.

2 Related Work

2.1 All-in-One Image Restoration

Traditional image restoration typically targets specific degradations, such as noise or blur, using specialized models Zamir et al. (2022); Huang et al. (2023); Cai et al. (2023). However, real-world scenarios often involve unknown or mixed degradations, driving the need for *all-in-one* models that handle diverse degradation types in a unified framework Su et al. (2022-05); Jiang et al. (2025).

Early efforts toward unified models utilized powerful backbones Chen et al. (2022; 2021); Wang et al. (2022); Zamir et al. (2022); Wang et al. (2024); Guo et al. (2024a) or generative approaches like diffusion models Belhasin et al. (2024); Yue and Loy (2024); Li et al. (2025); Yue et al. (2025). A key challenge is guiding a single network to handle diverse restoration tasks, leading to a focus

on conditioning mechanisms, from early degradation encoders Li et al. (2022) to advanced learnable prompts Potlapalli et al. (2023), sometimes refined with frequency priors or dimensionality reduction Cui et al. (2025); Zhang et al. (2023). Other methods use degradation priors from classifiers Kong et al. (2024) or vision-language models like CLIP with textual prompts Luo et al. (2024); Lin et al. (2024). Additionally, multi-task learning techniques Wu et al. (2024; 2025) and pretraining strategies Qin et al. (2024) address optimization challenges. While these advances improve all-in-one capabilities, they often focus on architectural or input conditioning changes, overlooking optimization stability in multi-task learning.

2.2 Privilege Learning and Test-Time Adaptation

Privilege Learning (PL) introduced the concept of using extra, training-only information to help a model learn a more robust representation Vapnik and Vashist (2009); Vapnik et al. (2015). While explored in low-level vision Lee et al. (2020), its application has been limited. We first establish PL for all-in-one restoration to achieve a more stable training baseline. More importantly, we extend PL to *retain* and *reuse* this privileged knowledge at inference, overcoming its core limitation.

Test-Time Adaptation (TTA) aims to adapt pre-trained models to new test domains via online, unsupervised fine-tuning Liang et al. (2025). Although our self-refinement also operates at test time, it is fundamentally different. TTA *updates model parameters* using gradient-based optimization on test data to address domain shift Deng et al. (2023); Gou et al. (2024); Yang et al. (2024). In stark contrast, SIPL is a *gradient-free inference process* that operates on a *frozen model*. It does not adapt the model but instead refines the output by reusing pre-distilled knowledge, making it a distinct paradigm from TTA.

3 Method

In this section, we detail our Self-Improved Privilege Learning (SIPL) framework. We first provide an overview of how SIPL integrates into a general restoration architecture. We then elaborate on its two core stages: (1) the PL-enhanced training process that retains privileged knowledge, and (2) the inference-time interation that reuses this knowledge for iterative self-refinement.

3.1 Preliminaries

Privilege Learning Paradigm As established in Section 1, all-in-one image restoration faces fundamental optimization challenges stemming from task competition and conflicting gradient directions. Privilege Learning Vapnik and Vashist (2009); Vapnik et al. (2015) offers an elegant solution by leveraging additional information during training that enhances the learning process. Given a degraded image I_d and its corresponding ground truth I_{gt}, PL incorporates privileged information (derived from I_{gt}) into the training process. This is achieved through a simple yet effective feature fusion approach:

$$F_{fused} = (1 - \alpha) \cdot F_d + \alpha \cdot F_{PI}, \tag{1}$$

where F_d represents features extracted from the degraded image, F_{PI} denotes features derived from the ground truth (privileged information), and $\alpha \in [0, 1]$ controls the degree of privileged guidance. During training, α typically follows a decreasing schedule, ensuring the model gradually adapts to operating without privileged guidance. At inference time, α is set to 0, as privileged information is unneeded.

This straightforward mechanism provides substantial benefits for all-in-one restoration. It stabilizes training and mitigates task competition by providing clear guidance, particularly for challenging degradation types. Meanwhile, our introduced PL paradigm remains agnostic to model architecture, enhancing all-in-one image restoration generally.

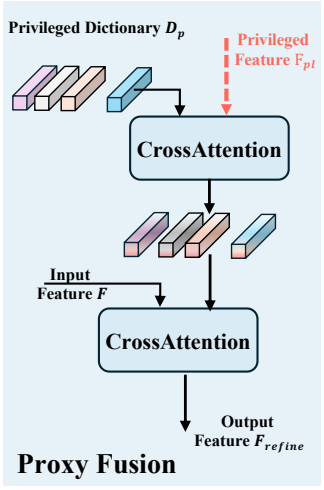

Figure 2: Implementation of the proxy fusion module.

3.2 Self-Improved Privilege Learning

Beyond the vanilla PL paradigm, we propose a novel extension to retain the privileged information at the inference stage. Our key insight is that the restoration output, though imperfect, contains valuable information closer to the ground truth than the original

degraded input. This observation motivates our proxy fusion module, which enables the model to leverage its own outputs as pseudo-privileged information during inference.

Proxy Fusion The cornerstone of SIPL is our novel Proxy Fusion mechanism, which creates a persistent bridge between training-time privileged knowledge and inference-time self-improvement. Unlike the direct feature blending in conventional PL (Eq. 1), Proxy Fusion employs a learnable Privileged Dictionary (PD) to distill and retain essential knowledge from privileged information:

$$\text{PD} \in \mathbb{R}^{N \times C}, \tag{2}$$

where N represents the number of dictionary entries and C is the feature dimension. This dictionary interacts with privileged features through cross-attention:

$$F'_{PI} = \text{CrossAttention}(Q = \text{PD}, K = F_{PI}, V = F_{PI}), \tag{3}$$

This interaction allows the PD to capture and internalize high-frequency details and statistical patterns characteristic of high-quality images. The distilled knowledge is then integrated with features from the degraded image:

$$F_{fused} = \text{CrossAttention}(Q = F_d, K = F'_{PI}, V = F'_{PI}). \tag{4}$$

During training, the entire framework, including the Privileged Dictionary, learns to extract and utilize privileged information effectively. Crucially, the learned PD persists into inference, serving as a knowledge repository that enables self-improvement without requiring actual ground truth.

Self-Improvement Mechanism The distinguishing feature of SIPL is its iterative self-refinement capability during inference:

1. **Initial Restoration**: The model produces an initial output $I^{(0)}_{restored}$ using only the degraded input:
$$I^{(0)}_{restored} = \mathcal{F}(I_d). \tag{5}$$
2. **Self-Improvement**: This initial output serves as pseudo-privileged information. Features extracted from $I^{(0)}_{restored}$ interact with the learned PD through the Proxy Fusion mechanism, guiding subsequent restoration:
$$I^{(1)}_{restored} = \mathcal{F}(I_d, I^{(0)}_{restored}). \tag{6}$$
3. **Iterative Refinement (Optional)**: This process can continue for multiple iterations, with each step potentially improving quality:
$$I^{(t)}_{restored} = \mathcal{F}(I_d, I^{(t-1)}_{restored}), \quad t \geq 1. \tag{7}$$

This elegant self-correction loop enables progressive quality enhancement without requiring actual ground truth during deployment. In practice, we find that a single refinement step ($t = 1$) typically provides substantial improvements with minimal computational overhead.

The key advantage of our Proxy Fusion approach over direct feature blending is its ability to distill and retain essential high-quality image characteristics in the learned PD parameters. This creates a persistent knowledge repository that facilitates self-improvement during inference, a capability absent in conventional PL frameworks.

3.3 Remarks

Our SIPL framework focuses exclusively on the learning paradigm rather than specific architectural choices. The Proxy Fusion module serves as a plug-and-play component that can enhance virtually any existing restoration architecture. This architectural agnosticism ensures broad applicability across diverse restoration models and tasks. For experimental validation, we integrate SIPL into multiple distinct backbone architectures, including PromptIR Potlapalli et al. (2023). As demonstrated in Section 4, SIPL consistently improves performance across all tested models, confirming its generality and effectiveness as an optimization framework for all-in-one image restoration.

4 Experiments

In this section, we evaluate our SIPL framework on diverse benchmarks (multi-task, deweathering, composite degradation). All models were retrained from scratch following original protocols for fairness. Detailed setups are in the Appendix.

Table 1: Quantitative results (PSNR/SSIM) on the Three-Task Setting. Our results are highlighted in **bold**, and best results are underlined.

Type	Method	Venue	Denoising (BSD68)			Dehazing	Deraining	Average
			$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	SOTS	Rain100L	
General	MPRNet Zamir et al. (2021)	CVPR'21	33.27/0.920	30.76/0.871	27.29/0.761	28.00/0.958	33.86/0.958	30.63/0.894
	Restormer Zamir et al. (2022)	CVPR'22	33.72/0.930	30.67/0.865	27.63/0.792	27.78/0.958	33.78/0.958	30.75/0.901
	NAFNet Chen et al. (2022)	ECCV'22	33.03/0.918	30.47/0.865	27.12/0.754	24.11/0.928	33.64/0.956	29.67/0.844
	FSNet* Cui et al. (2022)	TPAMI'23	33.81/0.930	30.84/0.872	27.69/0.792	29.14/0.968	35.61/0.969	31.42/0.906
	DRSformer* Chen et al. (2023)	CVPR'23	33.28/0.921	30.55/0.862	27.58/0.786	29.02/0.968	35.89/0.970	31.26/0.902
	MambaIR* Guo et al. (2024a)	ECCV'24	33.88/0.931	30.95/0.874	27.74/0.793	29.57/0.970	35.42/0.969	31.51/0.907
All-in-One	DL Fan et al. (2019)	TPAMI'19	33.05/0.914	30.41/0.861	26.90/0.740	26.92/0.391	32.62/0.931	29.98/0.875
	AirNet Li et al. (2022)	CVPR'22	33.92/0.932	31.26/0.888	28.00/0.797	27.94/0.962	34.90/0.967	31.20/0.910
	IDR* Zhang et al. (2023)	CVPR'23	33.89/0.931	31.32/0.884	28.04/0.798	29.87/0.970	36.03/0.971	31.83/0.911
	Gridformer* Wang et al. (2024)	IJCV'24	33.93/0.931	31.37/0.887	28.11/0.801	30.37/0.970	37.15/0.972	32.19/0.912
	NDR Yao et al. (2024)	TIP'24	34.01/0.932	31.36/0.887	28.10/0.798	28.64/0.962	35.42/0.969	31.51/0.910
	InstructIR Conde et al. (2024)	ECCV'24	34.15/0.933	31.52/0.890	28.30/0.804	30.22/0.959	37.98/0.978	32.43/0.913
	TextualDegRemoval Lin et al. (2024)	CVPR'24	34.01/0.933	31.39/0.890	28.18/0.802	31.63/0.980	37.58/0.979	32.63/0.917
	AdaIR Cui et al. (2025)	ICLR'25	34.12/0.935	31.45/0.892	28.19/0.802	31.06/0.980	38.64/0.983	32.69/0.918
	PromptIR Potlapalli et al. (2023)	NeurIPS'23	33.98/0.933	31.31/0.888	28.06/0.799	30.58/0.974	36.37/0.972	32.06/0.913
	PromptIR + SIPL	2025	**34.12/0.933**	**31.48/0.889**	**28.22/0.800**	**31.09/0.977**	**38.43/0.984**	**32.67/0.917**

Table 2: Quantitative results (PSNR/SSIM) on the Five-Task Setting. Our results are highlighted in **bold**, and best results are underlined.

Type	Method	Venue	Denoising	Dehazing	Deraining	Deblurring	Low-light	Average
			BSD68	SOTS	Rain100L	GoPro	LOL	
General	SwinIR Liang et al. (2021)	ICCVW'21	30.59/0.868	21.50/0.891	30.78/0.923	24.52/0.773	17.81/0.723	25.04/0.835
	Restormer Zamir et al. (2022)	CVPR'22	31.49/0.884	24.09/0.927	34.81/0.962	27.22/0.829	20.41/0.806	27.60/0.881
	NAFNet Chen et al. (2022)	ECCV'22	31.02/0.883	25.23/0.939	35.56/0.967	26.53/0.808	20.49/0.809	27.76/0.881
	DRSformer* Chen et al. (2023)	CVPR'23	30.97/0.881	24.66/0.931	33.45/0.953	25.56/0.780	21.77/0.821	27.28/0.873
	Retinexformer* Cai et al. (2023)	ICCV'23	30.84/0.880	24.81/0.933	32.68/0.940	25.09/0.779	22.76/0.834	27.24/0.873
	FSNet* Cui et al. (2024)	TPAMI'23	31.33/0.883	25.53/0.943	36.07/0.968	28.32/0.869	22.29/0.829	28.71/0.898
	MambaIR* Guo et al. (2024a)	ECCV'24	31.41/0.884	25.81/0.944	36.55/0.971	28.61/0.875	22.49/0.832	28.97/0.901
All-in-One	DL Fan et al. (2019)	TPAMI'19	23.09/0.745	20.54/0.826	21.96/0.762	19.86/0.672	19.83/0.712	21.05/0.743
	TAPE Liu et al. (2022)	ECCV'22	30.18/0.855	22.16/0.861	29.67/0.904	24.47/0.763	18.97/0.621	25.09/0.801
	Transweather Valanarasu et al. (2022)	CVPR'22	29.00/0.841	21.32/0.885	29.43/0.905	25.12/0.757	21.21/0.792	25.22/0.836
	AirNet Li et al. (2022)	CVPR'22	30.91/0.882	21.04/0.884	32.98/0.951	24.35/0.781	18.18/0.735	25.49/0.846
	IDR Zhang et al. (2023)	CVPR'23	31.60/0.887	25.24/0.943	35.63/0.965	27.87/0.846	21.34/0.826	28.34/0.893
	Gridformer* Wang et al. (2024)	IJCV'24	31.45/0.885	26.79/0.951	36.61/0.971	29.22/0.884	22.59/0.831	29.33/0.904
	InstructIR Conde et al. (2024)	ECCV'24	31.40/0.887	27.10/0.956	36.84/0.973	29.40/0.886	23.00/0.836	29.55/0.907
	AdaIR Cui et al. (2025)	ICLR'25	31.35/0.889	30.53/0.978	38.02/0.981	28.12/0.858	23.00/0.845	30.20/0.910
	PromptIR* Potlapalli et al. (2023)	NeurIPS'23	31.47/0.886	26.54/0.949	36.37/0.970	28.71/0.881	22.68/0.832	29.15/0.904
	PromptIR + SIPL	2025	**31.45/0.888**	**30.51/0.975**	**38.09/0.982**	**29.35/0.886**	**23.23/0.856**	**30.53/0.917**

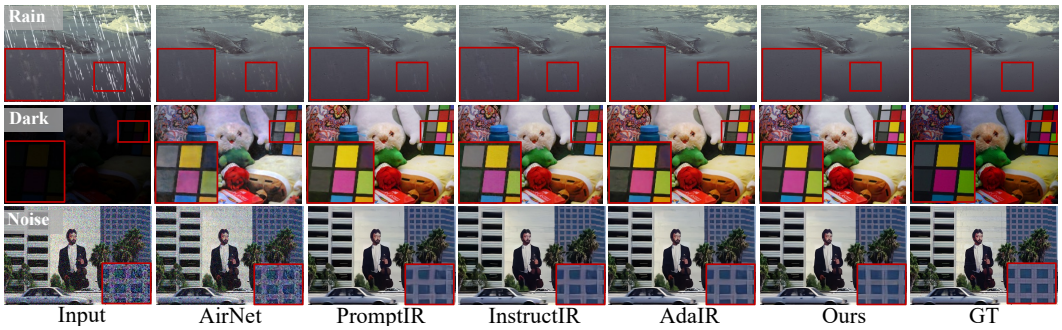

Figure 3: Visual comparison on the Five-Task benchmark. Our method demonstrates superior restoration quality across diverse degradations, effectively recovering finer details and image structures compared to other approaches.

4.1 Main Results

We now present the quantitative comparison of our SIPL framework, integrated with the PromptIR backbone (denoted as "PromptIR + SIPL"), against the original PromptIR and other state-of-the-art methods across the four benchmark settings.

Results on Three-Task Setting As detailed in Table 1, integrating our SIPL framework with the PromptIR backbone yields significant performance gains. SIPL boosts the average PSNR by a

Table 3: Quantitative results (PSNR/SSIM) on the Deweathering Setting. Our results are highlighted in **bold**, and best results are underlined.

Method	Venue	Snow100K-S		Snow100K-L		Outdoor-Rain		RainDrop		Average	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
All-in-One Li et al. (2020)	CVPR'20	–	–	28.33	0.8820	24.71	0.8980	31.12	0.9268	28.05	0.9023
Transweather Valanarasu et al. (2022)	CVPR'22	32.51	0.9341	29.31	0.8879	28.83	0.9000	30.17	0.9157	30.20	0.9094
WGWSNet Zhu et al. (2023)	CVPR'22	34.31	0.9460	30.16	0.9007	29.32	0.9207	32.38	0.9378	31.54	0.9263
WeatherDiff$_{64}$ Özdenizci and Legenstein (2023)	TPAMI'23	35.83	0.9566	30.09	0.9041	29.64	0.9312	30.71	0.9312	31.57	0.9308
WeatherDiff$_{128}$ Özdenizci and Legenstein (2023)	TPAMI'23	35.02	0.9516	29.58	0.8941	29.72	0.9216	29.66	0.9225	31.00	0.9225
AWRCP Ye et al. (2023)	ICCV'23	36.92	0.9652	31.92	0.9341	31.39	0.9329	31.93	0.9314	33.04	0.9409
GridFormer Wang et al. (2024)	IJCV'24	37.46	0.9640	31.71	0.9231	31.87	0.9335	32.39	0.9362	33.36	0.9392
MPerceiver Ai et al. (2024)	CVPR'24	36.23	0.9571	31.02	0.9164	31.25	0.9246	33.21	0.9294	32.93	0.9319
DTPM Ye et al. (2024)	CVPR'24	37.01	0.9663	30.92	0.9174	30.99	0.9340	32.72	0.9440	32.91	0.9404
Histoformer Sun et al. (2024)	ECCV'24	37.41	0.9656	32.16	0.9261	32.08	0.9389	33.06	0.9441	33.68	0.9437
PromptIR Potlapalli et al. (2023)	NeurIPS'23	36.88	0.9643	31.34	0.9200	30.80	0.9229	32.20	0.9359	32.80	0.9357
PromptIR + SIPL	2025	**37.91**	**0.9673**	**32.34**	**0.9291**	**32.91**	**0.9469**	**32.99**	**0.9462**	**34.03**	**0.9473**

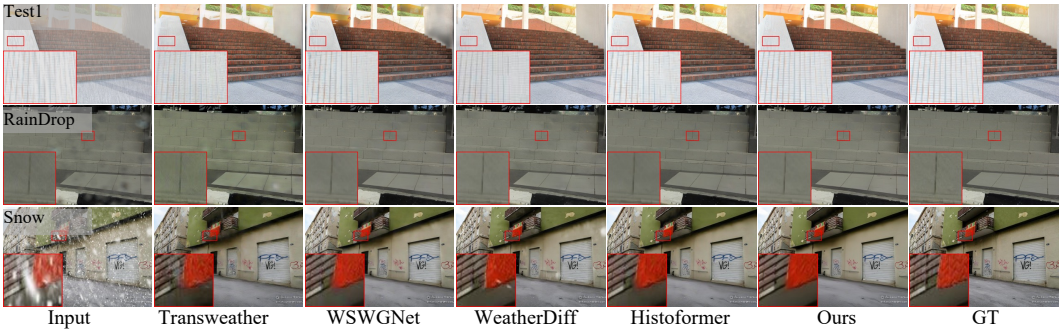

Figure 4: Qualitative examples from the AllWeather dataset. Our method exhibits robust performance in removing various challenging weather conditions. It yields visually superior results with better detail preservation and fewer artifacts.

notable **+0.61 dB** over the baseline. The benefits are particularly pronounced in challenging tasks like deraining, where SIPL achieves a substantial **+2.06 dB** improvement, and in dehazing (+0.51 dB). Our SIPL-enhanced model is highly competitive against recent state-of-the-art methods, including TextualDegRemoval Lin et al. (2024) and AdaIR Cui et al. (2025).

Results on Five-Task Setting On the more demanding five-task benchmark, the advantages of our SIPL framework become even more apparent (Table 2). SIPL significantly elevates the PromptIR baseline, boosting the average PSNR by a substantial **+1.38 dB**. The improvements are particularly striking in tasks where the baseline struggles due to task competition, such as dehazing (**+3.97 dB**) and deraining (**+1.72 dB**), alongside a solid gain in low-light enhancement (+0.55 dB). The visual results in Figure 3 corroborate these findings, showing superior detail and texture preservation across a range of degradations.

Results on Deweathering Setting The deweathering benchmark, summarized in Table 3, evaluates performance on removing diverse adverse weather conditions. Our approach again demonstrates superior capabilities, achieving the new SOTA average PSNR of 34.03 dB and SSIM of 0.9473. Consistent performance enhancements are recorded across all four test datasets. These results underscore SIPL's robustness in complex deweathering scenarios, outperforming specialized methods and recent deweathering models like Histoformer Sun et al. (2024) and GridFormer Wang et al. (2024). Qualitative results in Figure 4 corroborate these metrics, showing that our method more effectively removes severe weather artifacts while better preserving fine details and color fidelity.

Results on Composite Degradation Setting The advantages of SIPL are most pronounced on the challenging Composite Degradation benchmark (Table 4), where multiple degradations interact. Here, our approach achieves a remarkable **+4.58 dB** average PSNR improvement over the PromptIR baseline. This substantial gain, far exceeding that of PL-enhanced training alone, underscores the power of our inference-time self-refinement. The performance leap is particularly notable on severe combined degradations like haze+snow (**+9.01 dB**) and haze+rain (**+8.16 dB**). As shown in Figure 5, unlike baseline methods that typically address only one degradation, SIPL demonstrates superior generalization, effectively mitigating all co-occurring artifacts and restoring details and color fidelity.

Table 4: Quantitative results (PSNR/SSIM/LPIPS) on the Composite Degradation Setting. Our results are highlighted in **bold**, and best results are underlined.

Method	Venue	l	h	r	s	l+h	l+r	l+s	h+r	h+s	l+h+r	l+h+s	Avg.
AirNetLi et al. (2022)	CVPR'22	24.83	24.21	26.55	26.79	23.23	22.82	23.29	22.21	23.29	21.80	22.24	23.75
TransWeatherValanarasu et al. (2022)	CVPR'22	23.39	23.95	26.69	25.74	22.24	22.62	21.80	23.10	22.34	21.55	21.01	23.13
WeatherDiffÖzdenizci and Legenstein (2023)	TPAMI'23	23.58	21.99	24.85	24.80	21.83	22.69	22.12	21.25	21.99	21.23	21.04	22.49
WGWSNetZhu et al. (2023)	CVPR'23	24.39	27.90	33.15	34.43	24.27	25.06	24.60	27.23	27.65	23.90	23.97	26.96
InstructIRConde et al. (2024)	ECCV'24	26.70	32.61	33.51	34.45	24.36	25.41	25.63	28.80	29.64	24.84	24.32	28.21
OneRestoreGuo et al. (2024b)	ECCV'24	26.55	32.71	33.48	34.50	26.15	25.83	25.56	30.27	30.46	25.18	25.28	28.47
PromptIRPotlapalli et al. (2023)	NeurIPS'23	26.32	26.10	31.56	31.53	24.49	25.05	24.51	24.54	23.70	23.74	23.33	25.90
PromptIR + SIPL)	2025	**27.62**	**36.82**	**35.66**	**36.85**	**27.03**	**26.79**	**26.68**	**32.70**	**32.71**	**26.20**	**26.20**	**30.48**

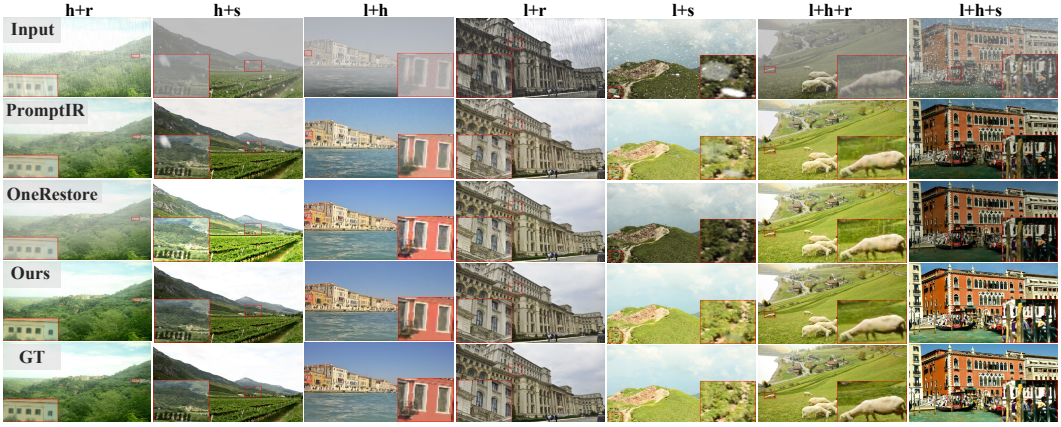

Figure 5: Visual results on the composite degradation tasks, showcasing performance on mixed degradations. Our method more effectively mitigates multiple interacting degradations, restoring clearer images with improved color fidelity and detail.

4.2 Ablation Studies

In this section, we conduct comprehensive ablation studies to meticulously validate the efficacy of our proposed SIPL framework and dissect the contributions of its core components.

Architectural Agnosticism of SIPL A core strength of our SIPL framework lies in its architectural agnosticism and ease of integration. To substantiate this plug-and-play capability, we applied SIPL to a spectrum of distinct backbone architectures, moving beyond the PromptIR Potlapalli et al. (2023). These included Restormer Zamir et al. (2022), a prominent Transformer-based network; NAFNet Chen et al. (2022), known for its high CNN efficiency; and AdaIR Cui et al. (2025), a recent state-of-the-art method notable for its frequency domain processing. All models were retrained on the five-task benchmark with and without the SIPL framework integrated. The results, presented in Figure 6, unequivocally showcase SIPL's ability to consistently elevate performance across these diverse architectural paradigms. More detailed results are available at Appendix.

Dissecting the Contributions of SIPL Components To meticulously evaluate SIPL's core components, we ablate PromptIR on a five degradation tasks, with quantitative and qualitative results presented in Figure 7. The baseline PromptIR model achieves a PSNR of 36.37 dB. Introducing privilege learning solely during training ("+PL") substantially boosts performance to 37.49 dB. This underscores PL's efficacy in stabilizing multi-task optimization by leveraging privileged information, thereby establishing a stronger foundation. Building upon this, the full SIPL framework, by incorporating the Proxy Fusion module for its initial inference application ("+SIPL"), further elevates the PSNR to 37.91 dB (+0.42 dB over "+PL"). This increment highlights the Proxy Fusion module's critical role, with its Privileged Dictionary, in effectively distilling, preserving, and transfer-

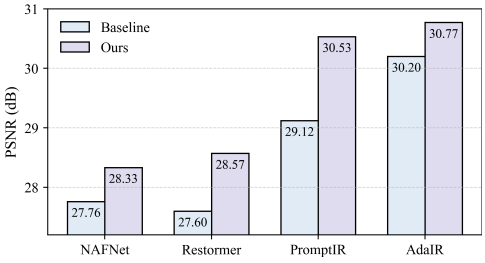

Figure 6: Ablation study on the interaction of SIPL with diverse backbone architectures on the five-task benchmark.

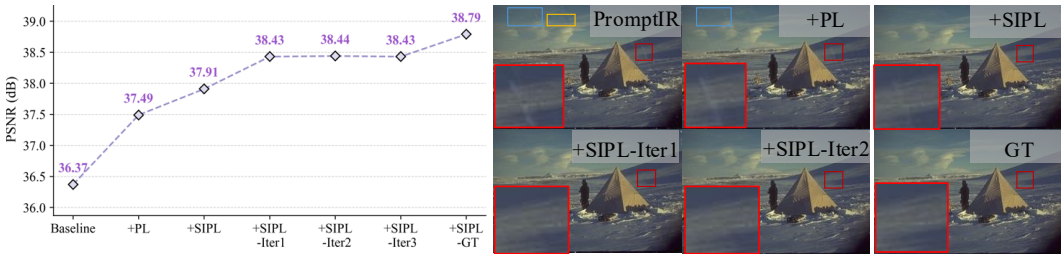

Figure 7: Ablation study dissecting the contributions of SIPL's components. The figure illustrates the progressive improvements from the baseline, through the addition of privilege learning ("+PL"), the initial application of SIPL ("+SIPL"), and subsequent iterative self-refinement stages ("+SIPL-IterX"). Performance is benchmarked against an approximate upper bound using GT-guidance ("+SIPL-GT").

ring privileged knowledge for tangible improvement at inference time. The iterative self-improvement mechanism, a key innovation of SIPL, demonstrates further significant refinement. Crucially, when the model's output from the "+SIPL" stage is *first fed back as pseudo-privileged information* (resulting in "+SIPL-Iter1"), performance impressively surges to 38.43 dB—an additional +0.52 dB gain. This substantial improvement from the initial feedback loop powerfully demonstrates the innovation and effectiveness of using self-generated outputs for refinement. While subsequent iterations show diminishing returns, they affirm the model's capacity for self-correction. This iterative process using pseudo-PI effectively narrows the gap towards the performance upper bound of 38.79 dB, achieved when utilizing true ground truth as privileged input during inference.

Table 5: Average PSNR (dB) across ten self-refinement iterations on the Three-Task and Five-Task benchmarks. The most significant performance gain occurs at the first iteration ($t = 1$), after which the performance rapidly converges to a stable plateau, demonstrating the robustness and stability of our iterative refinement mechanism. Detailed per-task metrics are available in the appendix.

Benchmark	Baseline	t=1	t=2	t=3	t=4	t=5	t=6	t=7	t=8	t=9	t=10
Three-Task	32.06	32.669	32.670	32.673	**32.673**	**32.673**	32.671	32.670	32.669	32.669	32.669
Five-Task	29.15	30.53	30.58	30.58	**30.60**	30.58	30.55	30.53	30.54	30.52	30.53

Evaluation on Multi-Step Refinement To comprehensively analyze the behavior and stability of our iterative self-refinement, we evaluated its performance over ten successive steps. Table 5 presents the average PSNR progression on both the Three-Task and Five-Task benchmarks, with peak performance highlighted in bold. The results reveal two crucial insights into our framework's dynamics. First, the most substantial performance leap universally occurs at the initial refinement step ($t = 1$). On the Three-Task benchmark, this single step yields a gain of +0.609 dB, while on the more complex Five-Task benchmark, the improvement is an even more remarkable +1.38 dB. This demonstrates that our mechanism can immediately and effectively correct the most significant errors in the initial restoration. Second, beyond this initial jump, the performance rapidly saturates and converges to a stable plateau. Subsequent iterations from $t = 2$ to $t = 10$ result in only marginal fluctuations (e.g., within $\pm$0.004 dB on the Three-Task benchmark), with performance remaining consistently high and peaking around the fourth iteration. This graceful convergence is a hallmark of our method's robustness. It stands in stark contrast to naive feedback loops, which, as our prior ablation showed, lead to catastrophic performance collapse due to error amplification. The stability of SIPL is attributed to the Privileged Dictionary (PD), which acts as a constant, reliable guide, ensuring that the feature-level corrections do not diverge. Based on this analysis, our choice of a single refinement step ($t = 1$) for the main experiments is a principled one, offering an optimal balance between substantial performance enhancement and computational efficiency.

Analysis of the Self-Improvement Mechanism To rigorously validate superiority of our proposed self-improvement mechanism, we compare it against two common inference-time enhancement strategies: naive iteration and self-ensembling. As shown in Table 6, simply feeding a model's output back as its input (Naive Iteration) results in a catastrophic performance collapse of -3.82 dB. This failure is anticipated: the initially restored image, while visually improved, constitutes an out-of-distribution (OOD) input for a model trained on heavily degraded data. This domain mismatch leads to the amplification of residual artifacts and color shifts, causing the model to diverge. This

Table 6: Ablation study on different inference-time enhancement strategies on the Three-Task benchmark. Our proposed SIPL framework not only significantly outperforms the baselines in restoration quality (PSNR/SSIM) but also demonstrates superior computational efficiency compared to standard self-ensembling.

Method	PSNR (Avg.)	SSIM (Avg.)	Forward Passes	Mechanism
Baseline (PromptIR)	32.06	0.913	1×	Standard Inference
+ Naive Iteration	28.24 (-3.82)	0.892	2×	Pixel-level Reprocessing
+ Self-Ensembling	32.42 (+0.36)	0.915	8×	Averaging Augmented Outputs
+ SIPL (Ours)	**32.67 (+0.61)**	**0.917**	**2×**	**PD-Guided Feature Refinement**

experiment unequivocally demonstrates that effective self-correction is a non-trivial task that cannot be achieved by simple looping. Unlike naive iteration's pixel-level reprocessing, SIPL performs a principled **feature-level refinement**. The initial output's features are not used to generate a new image directly, but rather to query the pre-learned Privileged Dictionary (PD). This PD, having distilled the essence of high-quality images during training, acts as a robust guardrail, providing a stable, gradient-free correction signal in the latent space.

Moreover, self-ensembling, which averages the outputs from eight geometrically augmented inputs, provides a moderate performance gain (+0.36 dB). However, our SIPL framework surpasses it with a more substantial improvement of +0.61 dB. Notably, SIPL achieves this superior result with only 2 forward passes (for one refinement step as the default), making it approximately **4× more computationally efficient** than the 8-pass self-ensembling process. The stability of this process is further corroborated by our multi-iteration experiments (as shown in Table 6), where performance gracefully saturates near the GT-guided upper bound rather than collapsing. This highlights the robustness and fundamental advantage of our PD-guided self-improvement paradigm. Further analyses are detailed in the Appendix.

4.3 Limitation and Future Work

While our experiments have demonstrated the significant success and versatility of the SIPL framework, we identify several limitations and promising avenues for future research. First, while we empirically validate that Privilege Learning (PL) stabilizes multi-degradation training, a deeper theoretical understanding of its optimization dynamics in this context remains an open question. A more formal understanding in this area is a critical challenge for the community, which could unlock even more profound improvements. Second, our iterative refinement, while more efficient than ensembling, still increases latency over single-pass baselines; optimizing this performance-cost trade-off is a key priority. Looking forward, SIPL's task- and model-agnostic design makes it a promising candidate for a wider range of restoration problems, especially in complex real-world scenarios. More broadly, our self-refinement concept resonates with the reasoning capabilities of recent large models. Future work could explore integrating principles from paradigms like Chain-of-Thought Wei et al. (2022), potentially using large vision models to guide more sophisticated, multi-step correction processes and unlock new capabilities in low-level vision.

5 Conclusion

In this work, we propose Self-Improved Privilege Learning (SIPL), a novel framework that effectively tackles critical optimization impediments in all-in-one image restoration. SIPL uniquely extends the paradigm of privilege learning to the inference stage: models are empowered to iteratively self-refine their outputs by leveraging them as pseudo-privileged information. This is realized through our proposed proxy fusion module, which employs a privileged dictionary, distilled from ground-truth priors during training, to guide this self-correction process with retrained privileged prior. Extensive evaluations across multiple challenging benchmarks, particularly those with complex composite degradations, confirm that SIPL substantially boosts the performance of diverse state-of-the-art methods, significantly enhancing their robustness and overall restoration quality. We hope that the principles and methodologies presented in SIPL will offer fresh perspectives to the community and stimulate further exploration into more effective strategies for all-in-one image restoration.

REFERENCES

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*, pages 1–20, 2025. doi: 10.1109/TPAMI.2025.3598132.

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

Gang Wu, Junjun Jiang, Kui Jiang, and Xianming Liu. Harmony in diversity: Improving all-in-one image restoration via multi-task collaboration. In *ACM Multimedia*, pages 6015–6023, 2024.

Vladimir Vapnik and Akshay Vashist. A new learning paradigm: Learning using privileged information. *Neural networks*, 22(5-6):544–557, 2009.

Vladimir Vapnik, Rauf Izmailov, et al. Learning using privileged information: similarity control and knowledge transfer. *J. Mach. Learn. Res.*, 16(1):2023–2049, 2015.

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 *IEEE conference on computer vision and pattern recognition (CVPR)*, pages 5718–5729, 2022.

Huaibo Huang, Mandi Luo, and Ran He. Memory uncertainty learning for real-world single image deraining. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 45(3):3446–3460, 2023.

Yuanhao Cai, Hao Bian, Jing Lin, Haoqian Wang, Radu Timofte, and Yulun Zhang. Retinexformer: One-stage retinex-based transformer for low-light image enhancement. In *Proceedings of the IEEE/CVF international conference on computer vision (ICCV)*, pages 12470–12479, 2023.

Jingwen Su, Boyan Xu, and Hujun Yin. A survey of deep learning approaches to image restoration. *Neurocomput.*, 487:46–65, 2022-05. ISSN 0925-2312.

Liangyu Chen, Xiaojie Chu, Xiangyu Zhang, and Jian Sun. Simple baselines for image restoration. In *Proceedings of the european conference on computer vision (ECCV)*, volume 13667, pages 17–33, 2022.

Hanting Chen, Yunhe Wang, Tianyu Guo, Chang Xu, Yiping Deng, Zhenhua Liu, Siwei Ma, Chunjing Xu, Chao Xu, and Wen Gao. Pre-trained image processing transformer. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pages 12299–12310, 2021.

Zhendong Wang, Xiaodong Cun, Jianmin Bao, Wengang Zhou, Jianzhuang Liu, and Houqiang Li. Uformer: A general u-shaped transformer for image restoration. In *Proceedings of the IEEE conference on computer vision and pattern recognition (CVPR)*, pages 17662–17672, 2022.

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

Hang Guo, Jinmin Li, Tao Dai, Zhihao Ouyang, Xudong Ren, and Shu-Tao Xia. Mambair: A simple baseline for image restoration with state-space model. In *Proceedings of the european conference on computer vision (ECCV)*, volume 15076, pages 222–241, 2024a.

Omer Belhasin, Yaniv Romano, Daniel Freedman, Ehud Rivlin, and Michael Elad. Principal uncertainty quantification with spatial correlation for image restoration problems. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(5):3321–3333, 2024.

Zongsheng Yue and Chen Change Loy. Diffface: Blind face restoration with diffused error contraction. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 46(12):9991–10004, 2024.

Miaoyu Li, Ying Fu, Tao Zhang, Ji Liu, Dejing Dou, Chenggang Yan, and Yulun Zhang. Latent diffusion enhanced rectangle transformer for hyperspectral image restoration. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(1):549–564, 2025.

Zongsheng Yue, Jianyi Wang, and Chen Change Loy. Efficient diffusion model for image restoration by residual shifting. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 47(1): 116–130, 2025.

Boyun Li, Xiao Liu, Peng Hu, Zhongqin Wu, Jiancheng Lv, and Xi Peng. All-in-one image restoration for unknown corruption. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, pages 17431–17441, 2022.

Vaishnav Potlapalli, Syed Waqas Zamir, Salman H. Khan, and Fahad Shahbaz Khan. Promptir: Prompting for all-in-one image restoration. In *Advances in neural information processing systems (NeurIPS)*, 2023.

Yuning Cui, Syed Waqas Zamir, Salman H. Khan, Alois Knoll, Mubarak Shah, and Fahad Shahbaz Khan. Adair: Adaptive all-in-one image restoration via frequency mining and modulation. In *International conference on learning representations (ICLR)*, 2025.

Jinghao Zhang, Jie Huang, Mingde Yao, Zizheng Yang, Hu Yu, Man Zhou, and Feng Zhao. Ingredient-oriented multi-degradation learning for image restoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5825–5835, 2023.

Ziwei Luo, Fredrik K. Gustafsson, Zheng Zhao, Jens Sjölund, and Thomas B. Schön. Controlling vision-language models for multi-task image restoration. In *International conference on learning representations (ICLR)*, 2024.

Jingbo Lin, Zhilu Zhang, Yuxiang Wei, Dongwei Ren, Dongsheng Jiang, Qi Tian, and Wangmeng Zuo. Improving image restoration through removing degradations in textual representations. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2866–2878, 2024.

Gang Wu, Junjun Jiang, Yijun Wang, Kui Jiang, and Xianming Liu. Debiased all-in-one image restoration with task uncertainty regularization. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2025.

Chu-Jie Qin, Ruiqi Wu, Zikun Liu, Xin Lin, Chun-Le Guo, Hyun Hee Park, and Chongyi Li. Restore anything with masks: Leveraging mask image modeling for blind all-in-one image restoration. In *Proceedings of the european conference on computer vision (ECCV)*, volume 15103, pages 364–380, 2024.

Wonkyung Lee, Junghyup Lee, Dohyung Kim, and Bumsub Ham. Learning with privileged information for efficient image super-resolution. In *ECCV*, pages 465–482, 2020.

Jian Liang, Ran He, and Tieniu Tan. A comprehensive survey on test-time adaptation under distribution shifts. *IJCV*, 2025.

Zeshuai Deng, Zhuokun Chen, Shuaicheng Niu, Thomas H. Li, Bohan Zhuang, and Mingkui Tan. Efficient test-time adaptation for super-resolution with second-order degradation and reconstruction. In Alice Oh, Tristan Naumann, Amir Globerson, Kate Saenko, Moritz Hardt, and Sergey Levine, editors, *NeurIPS*, 2023.

Yuanbiao Gou, Haiyu Zhao, Boyun Li, Xinyan Xiao, and Xi Peng. Test-time degradation adaptation for open-set image restoration. In *ICML*, 2024.

Yijun Yang, Hongtao Wu, Angelica I. Avilés-Rivero, Yulun Zhang, Jing Qin, and Lei Zhu. Genuine knowledge from practice: Diffusion test-time adaptation for video adverse weather removal. In *CVPR*, pages 25606–25616, 2024.

Syed Waqas Zamir, Aditya Arora, Salman H. 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 (CVPR)*, pages 14821–14831, 2021.

Yuning Cui, Yi Tao, Zhenshan Bing, Wenqi Ren, Xinwei Gao, Xiaochun Cao, Kai Huang, and Alois Knoll. Selective frequency network for image restoration. In *International conference on learning representations (ICLR)*, 2022.

Xiang Chen, Hao Li, Mingqiang Li, and Jinshan Pan. Learning A sparse transformer network for effective image deraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5896–5905, 2023.

Qingnan Fan, Dongdong Chen, Lu Yuan, Gang Hua, Nenghai Yu, and Baoquan Chen. A general decoupled learning framework for parameterized image operators. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.

Mingde Yao, Ruikang Xu, Yuanshen Guan, Jie Huang, and Zhiwei Xiong. Neural degradation representation learning for all-in-one image restoration. *IEEE Trans. Image Process.*, 33:5408–5423, 2024.

Marcos V Conde, Gregor Geigle, and Radu Timofte. High-quality image restoration following human instructions. In *European Conference on Computer Vision (ECCV)*, pages 1–12, 2024.

Jingyun Liang, Jiezhang Cao, Guolei Sun, Kai Zhang, Luc Van Gool, and Radu Timofte. SwinIR: Image restoration using swin transformer. In *ICCVW*, pages 1833–1844, 2021.

Yuning Cui, Wenqi Ren, Xiaochun Cao, and Alois Knoll. Image restoration via frequency selection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 46(2):1093–1108, 2024.

Lin Liu, Lingxi Xie, Xiaopeng Zhang, Shanxin Yuan, Xiangyu Chen, Wengang Zhou, Houqiang Li, and Qi Tian. TAPE: task-agnostic prior embedding for image restoration. In *Proceedings of the european conference on computer vision (ECCV)*, volume 13678, pages 447–464, 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 conference on computer vision and pattern recognition (CVPR)*, pages 2343–2353, 2022.

Ruoteng Li, Robby T Tan, and Loong-Fah Cheong. All in one bad weather removal using architectural search. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3175–3185, 2020.

Yurui Zhu, Tianyu Wang, Xueyang Fu, Xuanyu Yang, Xin Guo, Jifeng Dai, Yu Qiao, and Xiaowei Hu. Learning weather-general and weather-specific features for image restoration under multiple adverse weather conditions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 21747–21758, 2023.

Ozan Özdenizci and Robert Legenstein. Restoring vision in adverse weather conditions with patch-based denoising diffusion models. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(8):10346–10357, 2023.

Tian Ye, Sixiang Chen, Jinbin Bai, Jun Shi, Chenghao Xue, Jingxia Jiang, Junjie Yin, Erkang Chen, and Yun Liu. Adverse weather removal with codebook priors. In *IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 12619–12630, 2023.

Yuang Ai, Huaibo Huang, Xiaoqiang Zhou, Jiexiang Wang, and Ran He. Multimodal prompt perceiver: Empower adaptiveness, generalizability and fidelity for all-in-one image restoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 25432–25444, 2024.

Tian Ye, Sixiang Chen, Wenhao Chai, Zhaohu Xing, Jing Qin, Ge Lin, and Lei Zhu. Learning diffusion texture priors for image restoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2524–2534, 2024.

Shangquan Sun, Wenqi Ren, Xinwei Gao, Rui Wang, and Xiaochun Cao. Restoring images in adverse weather conditions via histogram transformer. In *European Conference on Computer Vision (ECCV)*, volume 15080, pages 111–129, 2024.

Yu Guo, Yuan Gao, Yuxu Lu, Ryan Wen Liu, and Shengfeng He. Onerestore: A universal restoration framework for composite degradation. In *Proceedings of the european conference on computer vision (ECCV)*, pages 255–272, 2024b.

Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed H. Chi, Quoc V. Le, and Denny Zhou. Chain-of-thought prompting elicits reasoning in large language models. In *NeurIPS*, 2022.

Pablo Arbeláez, Michael Maire, Charless Fowlkes, and Jitendra Malik. Contour detection and hierarchical image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(5):898–916, 2011.

Kede Ma, Zhengfang Duanmu, Qingbo Wu, Zhou Wang, Hongwei Yong, Hongliang Li, and Lei Zhang. Waterloo Exploration Database: New challenges for image quality assessment models. *IEEE Trans. Image Process.*, 26(2):1004–1016, Feb. 2017.

Wenhan Yang, Robby T. Tan, Jiashi Feng, Jiaying Liu, Zongming Guo, and Shuicheng Yan. Deep joint rain detection and removal from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1685–1694, 2017.

Boyi Li, Wenqi Ren, Dengpan Fu, Dacheng Tao, Dan Feng, Wenjun Zeng, and Zhangyang Wang. Benchmarking single-image dehazing and beyond. *IEEE Trans. Image Process.*, 28(1):492–505, 2019a.

Seungjun Nah, Tae Hyun Kim, and Kyoung Mu Lee. Deep multi-scale convolutional neural network for dynamic scene deblurring. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 257–265, July 2017.

Chen Wei, Wenjing Wang, Wenhan Yang, and Jiaying Liu. Deep retinex decomposition for low-light enhancement. In *Proceedings of the british machine vision conference (BMVC)*, page 155, 2018.

Rui Qian, Robby T. Tan, Wenhan Yang, Jiajun Su, and Jiaying Liu. Attentive generative adversarial network for raindrop removal from a single image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2482–2491, 2018.

Ruoteng Li, Loong-Fah Cheong, and Robby T. Tan. Heavy rain image restoration: Integrating physics model and conditional adversarial learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1633–1642, 2019b.

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

He Zhang and Vishal M. Patel. Density-aware single image de-raining using a multi-stream dense network. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 695–704, 2018.

Jia-Bin Huang, Abhishek Singh, and Narendra Ahuja. Single image super-resolution from transformed self-exemplars. In *CVPR*, 2015.

Rich Franzen. Kodak lossless true color image suite. `http://r0k.us/graphics/kodak/`, 1999. Online accessed 24 Oct 2021.

David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *ICCV*, 2001.

A Additional Experimental Results and Analyses

This section presents supplementary experimental results, including single-task performance benchmarks, out-of-distribution (OOD) generalization analyses, computational complexity comparisons, further iterative inference studies, and additional qualitative results.

A.1 Detailed Experimental Setup

In this section, we provide the detailed experimental setup for the benchmarks used to evaluate our SIPL framework. We evaluate SIPL across four comprehensive all-in-one settings that encompass a wide range of real-world image degradation scenarios:

1. **Three-Task Setting:** Following the established protocol in Li et al. (2022); Potlapalli et al. (2023), we address three distinct degradation tasks: image denoising (with synthetic Gaussian noise), deraining (using the Rain100L dataset), and dehazing (on the RESIDE dataset). For denoising, we use BSD400 Arbeláez et al. (2011) and WED Ma et al. (2017) for training, and test on BSD68 with noise levels of 15, 25, and 50. For deraining, we employ the Rain100L Yang et al. (2017) dataset, and for dehazing, we use the outdoor subset of RESIDE Li et al. (2019a).
2. **Five-Task Setting:** To evaluate our model's capacity to handle a broader spectrum of degradations, we utilize a five-task benchmark, including deraining (Rain100L Yang et al. (2017)), dehazing (RESIDE Indoor Training Set), denoising (BSD400 + WED), motion deblurring (GoPro Nah et al. (2017)), and low-light enhancement (LOL Wei et al. (2018)). This setting follows Zhang et al. (2023) and tests the model's ability to handle diverse degradation types in real-world scenarios.
3. **Deweathering Setting:** Based on previous work Valanarasu et al. (2022), we use the AllWeather Valanarasu et al. (2022) dataset for training, containing images from Raindrop Qian et al. (2018), Outdoor-Rain Li et al. (2019b), and Snow100K Liu et al. (2018).
4. **Composite Degradation Setting:** Following the protocol established in Guo et al. (2024b), we evaluate our method on the Composite Degradation Dataset (CDD-11), which represents a more challenging scenario with mixed degradations. CDD-11 encompasses 11 categories of image degradations including single degradations (low-light, haze, rain, snow) and their combinations (low+haze, low+rain, low+snow, haze+rain, haze+snow, low+haze+rain, low+haze+snow). The dataset is constructed using standard benchmarks: the LOw-Light dataset (LOL) Wei et al. (2018), the REalistic Single Image DEhazing Outdoor Training Set (RESIDE-OTS) Li et al. (2019a), the Rain1200 dataset Zhang and Patel (2018), and the Snow100k dataset Liu et al. (2018). This setting particularly evaluates our framework's capability to handle complex, interacting degradations that better reflect real-world scenarios.

For all settings, we adopt the same training/testing splits and protocols as in the original works to ensure fair comparisons. We integrate our proposed SIPL framework into various backbone architectures to demonstrate its versatility and effectiveness.

A.2 Single-Task Performance Evaluation

To further assess the efficacy of SIPL, we evaluated its performance on individual restoration tasks, aligning our experimental setup with that of PromptIR Potlapalli et al. (2023) and AdaIR Cui et al. (2025). These evaluations test the capability of our all-in-one model, enhanced with SIPL, on specialized degradation scenarios.

Table 7: Deraining results in the single-task setting on the Rain100L dataset. Our SIPL approach obtains a significant performance boost of 1.98 dB PSNR over baseline PromptIR and 0.12 dB over the AdaIR.

Method	DIDMDN	UMR	SIRR	MSPFN	LPNet	AirNet	Restormer	PromptIR	AdaIR	**PromptIR + SIPL (Ours)**
PSNR	23.79	32.39	32.37	33.50	33.61	34.90	36.74	37.04	38.90	**39.02**
SSIM	0.773	0.921	0.926	0.948	0.958	0.977	0.978	0.979	0.985	**0.986**

Table 8: Quantitative comparison for single-task dehazing. Our SIPL achieves significant improvement over baseline PromptIR with 0.51 dB PSNR.

Method	DehazeNet	MSCNN	AODNet	EPDN	FDGAN	AirNet	Restormer	PromptIR	AdaIR	**PromptIR + SIPL**
PSNR	22.46	22.06	20.29	22.57	23.15	23.18	30.87	31.31	31.80	**31.82**
SSIM	0.851	0.908	0.877	0.863	0.921	0.900	0.969	0.973	0.981	**0.982**

Deraining on Rain100L: The Rain100L dataset Yang et al. (2017) serves as a standard benchmark for single-image deraining. As presented in Table 7, PromptIR augmented with our SIPL framework achieves state-of-the-art performance. Specifically, it obtains a PSNR of **39.02 dB** and an SSIM of **0.986**. This represents a substantial improvement of **1.98 dB** in PSNR over the original PromptIR baseline and also surpasses the strong AdaIR model by **0.12 dB**, demonstrating the significant benefits of SIPL in effectively removing rain streaks while preserving image fidelity.

Dehazing on SOTS Outdoor: For evaluating dehazing performance, we utilize the outdoor test set from SOTS, part of the RESIDE dataset Li et al. (2019a). The results in Table 8 indicate that SIPL notably enhances PromptIR's dehazing capabilities. Our approach (PromptIR + SIPL) achieves a PSNR of **31.82 dB** and an SSIM of **0.982**. This is a gain of **0.51 dB** in PSNR compared to the PromptIR baseline. Furthermore, our method slightly outperforms AdaIR (31.80 dB PSNR / 0.981 SSIM), underscoring SIPL's efficacy in restoring clarity and detail in hazy conditions.

Table 9: Image denoising results of directly applying the pre-trained model under the five-degradation setting to the Urban100 Huang et al. (2015), Kodak24 Franzen (1999) and BSD68 Martin et al. (2001) datasets. The results are PSNR scores. Our SIPL achieves significant improvement across all test datasets compared to previous SOTA method AdaIRCui et al. (2025).

	Urban100			Kodak24			BSD68			
Method	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	Average
DL Fan et al. (2019)	21.10	21.28	20.42	22.63	22.66	21.95	23.16	23.09	22.09	22.04
Transweather Valanarasu et al. (2022)	29.64	27.97	26.08	31.67	29.64	26.74	31.16	29.00	26.08	28.66
TAPE Liu et al. (2022)	32.19	29.65	25.87	33.24	30.70	27.19	32.86	30.18	26.63	29.83
AirNet Li et al. (2022)	33.16	30.83	27.45	34.14	31.74	28.59	33.49	30.91	27.66	30.89
IDR Zhang et al. (2023)	33.82	31.29	28.07	34.78	32.42	29.13	34.11	31.60	28.14	31.48
AdaIR	34.10	31.68	28.29	34.89	32.38	29.21	34.01	31.35	28.06	31.55
PromptIR + SIPL	**35.39**	**32.74**	**29.17**	**34.98**	**32.50**	**29.36**	**34.08**	**31.45**	**28.16**	**31.98**

Denoising using Five-Task Pre-trained Model: To assess robustness and generalization for denoising, we employed the all-in-one model pre-trained on five distinct degradation tasks (including denoising) and evaluated it directly on three commonly used denoising benchmark datasets: Urban100 Huang et al. (2015), Kodak24 Franzen (1999), and BSD68 Martin et al. (2001). This setup tests the model's ability to denoise effectively without task-specific fine-tuning. As detailed in Table 9, PromptIR + SIPL demonstrates superior performance, achieving an average PSNR of **31.98 dB** across all datasets and noise levels ($\sigma \in \{15, 25, 50\}$). This is a notable improvement over AdaIR (31.55 dB average PSNR). Particularly on the Urban100 dataset, which often contains complex structures, our method shows significant gains (e.g., **+1.29 dB** for $\sigma = 15$, **+1.06 dB** for $\sigma = 25$). Consistent, positive improvements are also observed across the Kodak24 and BSD68 datasets for all noise levels. These results, especially on datasets potentially unseen during the denoising phase of the five-task training, highlight the advanced robustness and generalization capabilities endowed by the SIPL framework.

Table 10: Performance of all-in-one models on the single task of deblurring (GoPro dataset). PSNR/SSIM values are reported. Our SIPL-enhanced model demonstrates superior transferability.

Method	AirNet Li et al. (2022)	PromptIR Potlapalli et al. (2023)	Perceive-IR	**PromptIR + SIPL (Ours)**
Venue	CVPR'22	NeurIPS'23	IEEE TIP'25	2025
PSNR / SSIM	31.64 / 0.945	32.41 / 0.956	**32.83 / 0.960**	32.77 / 0.961

Enhancing All-in-One Model Transfer for Deblurring To evaluate how SIPL enhances a generalist model's capability on a specialist task, we fine-tuned our SIPL-enhanced PromptIR (originally trained on five tasks) for the single task of deblurring on the GoPro dataset. As shown in Table 10, our approach achieves a significant improvement over the baseline PromptIR and obtains a highly competitive performance against the recently proposed Perceive-IR. This result validates that the priors learned via SIPL improve the model's transferability, allowing it to effectively adapt and excel in dedicated, single-task scenarios.

Beyond validating SIPL's ability to enhance the transferability of all-in-one models, we further sought to evaluate its effectiveness in boosting a dedicated, high-performance single-task architecture. To this end, we integrated our SIPL framework into the strong NAFNet baseline and retrained it from scratch on two standard benchmarks: the GoPro dataset for deblurring and the SIDD dataset for denoising.

Table 11: Single-task state-of-the-art comparison for deblurring on the GoPro dataset. SIPL provides a clear performance boost to the strong NAFNet baseline.

Method	Restormer	UFormer	MaIR	NAFNet	NAFNet+SIPL(Iter-0)	**NAFNet+SIPL(Iter-1)**
Venue	CVPR'22	CVPR'22	ECCV'22	ECCV'22	2025	**2025**
PSNR / SSIM	32.92 / 0.961	33.06 / 0.967	33.69 / 0.969	33.69 / 0.966	33.76 / 0.968	**33.82 / 0.970**

Evaluation on Image Deblurring Task To further demonstrate SIPL's architectural agnosticism and effectiveness, we integrated it into a strong, dedicated single-task baseline, NAFNet, and retrained it from scratch on the GoPro dataset. The results in Table 11 show that our PL-enhanced training (Iter-0) already provides a better starting model than the original NAFNet. The subsequent self-refinement step (Iter-1) yields an additional performance boost, setting a new state-of-the-art result. This confirms that SIPL is not limited to enhancing all-in-one models but also serves as a general and effective framework for pushing the performance of specialized, high-performing architectures.

Table 12: Single-task denoising results on the SIDD validation set. Our SIPL framework again demonstrates its value by improving the performance of the NAFNet baseline.

Method	MAXIM	CGNet	NAFNet	NAFNet+SIPL(Iter-0)	**NAFNet+SIPL(Iter-1)**
Venue	CVPR'22	TMLR'24	ECCV'22	2025	**2025**
PSNR (dB)	40.02	40.39	39.93	40.10	**40.19**

Evaluation on Image Denoising Task To validate SIPL's versatility across different tasks, we applied it to the single-task denoising benchmark on the SIDD dataset, again using NAFNet as the backbone. As presented in Table 12, the results are consistent with our other findings. The PL-enhanced training (Iter-0) establishes a stronger baseline model compared to the original NAFNet. Subsequently, the inference-time self-refinement step (Iter-1) provides a further boost, achieving a final PSNR of 40.19 dB. These comprehensive single-task results strongly validate that SIPL is an efficient and general framework for improving image restoration models across diverse tasks, datasets, and backbone architectures.

Summary of Single-Task Evaluations. While SIPL is designed to excel in challenging all-in-one settings, these comprehensive single-task evaluations confirm its broad effectiveness and versatility. The experiments demonstrate two key strengths: first, SIPL significantly enhances the transferability of a generalist model (PromptIR) to specialized tasks. Second, it provides a substantial performance boost to a dedicated, high-performance single-task model (NAFNet) trained from scratch. The consistent improvements across diverse tasks like deraining, dehazing, and denoising validate that SIPL is a truly task- and model-agnostic framework, effectively improving restoration capabilities in both multi-task and single-task scenarios.

Table 13: OOD performance on Rain100L with Gaussian noise (σ =15, 25, and 50). Models were trained on three distinct tasks and test on this unseen dataset directly. Iter-N denotes N iterative refinement steps. Best results are underlined, our method is highlighted.

Method	Rain100L + Gaussian Noise $\sigma = 15$	$\sigma = 25$	$\sigma = 50$	Average
PromptIR Potlapalli et al. (2023)	24.92	24.50	23.79	24.40
AdaIR Cui et al. (2025)	24.91	24.50	23.77	24.39
PromptIR + SIPL (initial)	**24.95**	**24.59**	**23.86**	**24.46**
+ Iter-1	**26.90**	**25.59**	**23.96**	**25.48**
+ Iter-2	**31.74**	**29.53**	**26.41**	**29.23**

Figure 8: Visual illustration of OOD performance on a challenging Rain100L + Gaussian Noise ($\sigma = 50$) example. From left to right: Degraded Input, PromptIR, AdaIR, SIPL (Initial), SIPL (Iter-1), SIPL (Iter-2), and Ground Truth (GT). The iterative application of SIPL progressively enhances image clarity, restores fine details, and reduces artifacts, significantly outperforming baseline methods and demonstrating effective generalization to unseen composite degradations.

A.3 Out-of-Distribution (OOD) Generalization Analysis

A critical attribute of advanced image restoration models is their capacity to generalize effectively to previously unseen degradation types. This section investigates the OOD generalization capabilities of our SIPL-enhanced framework, drawing comparisons with established methods like PromptIR and AdaIR. Our analysis specifically focuses on performance when encountering complex, composite degradations not present during the training phase.

Performance on Unseen Composite Degradation (Rain100L + Noise): We evaluate models originally trained on three distinct restoration tasks (deraining, dehazing, denoising individually) on a challenging synthetic dataset: Rain100L combined with varying levels of Gaussian noise ($\sigma \in \{15, 25, 50\}$). This composite degradation scenario was deliberately excluded from the training set to rigorously test OOD performance.

The quantitative results are presented in Table 13. Baseline models, PromptIR and AdaIR, achieve average PSNR scores of 24.40 dB and 24.39 dB, respectively, on this unfamiliar task. Our PromptIR + SIPL model, in its initial single-pass inference ("PromptIR + SIPL (initial)"), yields a comparable average PSNR of 24.46 dB. However, the transformative advantage of SIPL becomes strikingly evident through its iterative self-improvement mechanism. With just one iteration ("+Iter 1"), the average PSNR significantly jumps to 25.48 dB. A second iteration ("+Iter 2") further elevates the performance dramatically to an average PSNR of **29.23 dB**. This represents a remarkable **+4.77 dB** improvement over its initial state and far surpasses the static performance of the baseline models. This progressive and substantial enhancement underscores the robust OOD generalization conferred by SIPL, particularly its ability to iteratively refine results when faced with novel degradations.

The qualitative improvements are visualized in Figure 8 using an example from the Rain100L + Noise ($\sigma = 50$) set. While the input image exhibits significant degradation, and baseline methods like PromptIR and AdaIR offer limited restoration, our SIPL demonstrates clear visual enhancements. The initial output ("SIPL-Init") shows some improvement, but subsequent iterations ("SIPL-Iter1", "SIPL-Iter2") progressively recover finer details, enhance sharpness, and reduce artifacts more effectively, approaching the ground truth quality. This visual evidence corroborates the quantitative gains and highlights the practical efficacy of iterative refinement in complex OOD scenarios.

Efficacy of Self-Improvement in OOD Contexts: The marked success of PromptIR + SIPL in handling these unseen composite degradations, especially through iteration, is attributable to its core

design featuring the Privileged Dictionary (PD) and the self-improvement learning strategy. Unlike baseline models such as PromptIR and AdaIR, which are not inherently designed to leverage their own outputs for iterative refinement on OOD tasks without a guiding mechanism, SIPL excels in this regard. Standard architectures, if naively iterated on OOD inputs, might see performance stagnate or even degrade due to the accumulation of errors or model biases when processing unfamiliar data distributions.

In stark contrast, SIPL's PD, trained to distill essential characteristics of high-quality images, provides robust guidance even when the pseudo-privileged information is derived from an imperfectly restored OOD image. The iterative process allows the model to progressively correct errors and enhance image quality by repeatedly consulting these learned priors. This capacity for effective self-correction and refinement in the face of novel, complex degradations is a key differentiator of our approach.

This OOD analysis strongly suggests that our self-improved iteration paradigm offers a novel and potent pathway for advancing all-in-one image restoration. Beyond striving for optimal performance in a single forward pass, SIPL demonstrates the significant potential of empowering models to adapt and improve their outputs dynamically at test time. This is particularly crucial for real-world scenarios where diverse and unforeseen degradations are common, showcasing a promising direction for developing more versatile and robust restoration solutions.

A.4 Analysis of Efficiency and Performance Trade-offs

A core aspect of the SIPL framework is its flexibility, offering different trade-offs between computational cost and restoration quality. In this section, we provide a comprehensive analysis of this trade-off, first by dissecting the costs of SIPL's components on a large-scale model (PromptIR), and second, by comparing SIPL's iterative refinement against the common strategy of brute-force model scaling (on NAFNet).

Cost-Benefit Analysis on PromptIR. As summarized in Table 14, our methodology provides a spectrum of enhancements. The foundational Privilege Learning (PL), as a training-only strategy, offers a zero-cost inference boost, improving the PromptIR baseline by +0.9 dB on the five-task benchmark without any additional parameters or FLOPs. Building on this, our full SIPL framework in a single-pass configuration (Iter-0) adds a marginal cost (3M parameters and 20G FLOPs) for a further performance increase to 30.17 dB. This demonstrates the high efficiency of the Proxy Fusion module in retaining and applying privileged knowledge.

The full potential of SIPL is unlocked via iterative refinement (Iter-1), which, while increasing the computational load to 434G FLOPs, delivers the highest performance at 30.53 dB. We acknowledge this increased cost. However, this configuration represents a valuable trade-off, providing a powerful mechanism for tackling the most challenging restoration scenarios (e.g., composite and OOD degradations) where single-pass models often fall short.

Table 14: Comparison of model parameters and computational complexity for the PromptIR backbone on the five-task benchmark. FLOPs are calculated for a 256×256 input.

Method	Parameters	FLOPs	Avg. PSNR (dB)
AirNet Li et al. (2022)	9M	301G	25.44
Transweather Valanarasu et al. (2022)	21.5M	115.2G	25.22
PromptIR Potlapalli et al. (2023)	36M	173G	29.15
AdaIR Cui et al. (2025)	29M	162G	30.20
PromptIR + PL	36M	173G	30.05
PromptIR + SIPL (Iter-0)	39M	193G	30.17
PromptIR + SIPL (Iter-1)	39M	434G	30.53

SIPL as an Efficient Alternative to Model Scaling. To further contextualize SIPL's efficiency, we investigate a crucial question: is it better to apply SIPL to a smaller model or to simply train a larger one (such as NAFNet-32+SIPL vs NAFNet-64)? To answer this, we compare the performance and cost of applying SIPL to NAFNet-32 against a much larger NAFNet-64 baseline on the SIDD denoising task.

Table 15: Performance vs. FLOPs comparison on the SIDD benchmark. SIPL on a smaller model (NAFNet-32) nearly matches the performance of a much larger model (NAFNet-64) with significantly less computation.

Method	PSNR (dB) on SIDD	FLOPs
NAFNet-32 (baseline)	39.96	16.08 G
NAFNet-64 (larger model)	40.30	63.28 G
NAFNet-32 + SIPL (Iter-0)	40.10	17.63 G
NAFNet-32 + SIPL (Iter-1)	**40.19**	**43.93 G**

The results in Table 15 are compelling. NAFNet-32 + SIPL (Iter-1) achieves a PSNR of 40.19 dB, closing nearly all of the performance gap to the much larger NAFNet-64 (40.30 dB). Critically, it does so while requiring approximately **31% fewer FLOPs** (43.93G vs. 63.28G). This analysis compellingly demonstrates that SIPL is not merely an add-on for performance gain; it represents a more computationally efficient strategy for achieving top-tier results than simply scaling up a model's architecture. It validates that intelligently distilling and reusing privileged priors provides a more robust and efficient path to improvement.

Concluding Remarks. In summary, our efficiency analysis positions SIPL as a highly practical and versatile framework. It offers practitioners a flexible toolkit: from a zero-cost training enhancement (PL) to a powerful iterative refinement that can serve as a more efficient alternative to training ever-larger models. This makes SIPL a valuable contribution for developing high-performance image restoration solutions within specific computational budgets.

A.5 Quantitative Results of Multi-Step Refinement

This section provides a comprehensive analysis of the behavior of our Self-Improved Privilege Learning (SIPL) framework across multiple iterative refinement steps. While the main paper establishes that a single iteration ($t = 1$) offers an optimal trade-off between performance and efficiency, a deeper examination of the iterative process reveals crucial insights into the stability and robustness of our proposed mechanism. We present detailed quantitative results for both the Three-Task and Five-Task benchmarks over ten iterations, followed by an in-depth discussion.

Tables 16 and 17 detail the performance progression from the first to the tenth refinement step ($t = 10$). We also include the performance of a "GT-guided" oracle, where the ground-truth image is used as privileged information during inference, representing a practical upper bound for our method.

Table 16: Detailed performance metrics (PSNR) across ten iterations on the **Three-Task Benchmark**. The performance rapidly improves at $t = 1$ and then converges to a stable plateau, demonstrating the robustness of the self-refinement process. The average PSNR is shown with three decimal places to highlight the subtle changes in later iterations.

Method	Average	Deraining	Dehazing	Denoising ($\sigma = 15$)	Denoising ($\sigma = 25$)	Denoising ($\sigma = 50$)
PromptIR (Baseline)	32.060	36.37	30.58	33.98	31.31	28.06
+SIPL (Iter-1)	32.669	38.431	31.092	34.119	31.481	28.220
+SIPL (Iter-2)	32.670	38.428	31.095	34.122	31.482	28.225
+SIPL (Iter-3)	32.673	38.432	31.096	34.123	31.484	28.228
+SIPL (Iter-4)	**32.673**	**38.433**	31.098	34.122	31.483	**28.230**
+SIPL (Iter-5)	**32.673**	38.432	**31.101**	34.121	31.482	28.229
+SIPL (Iter-6)	32.671	38.431	31.097	34.120	31.481	28.227
+SIPL (Iter-7)	32.670	38.430	31.096	34.120	31.480	28.226
+SIPL (Iter-8)	32.669	38.429	31.094	34.119	31.480	28.225
+SIPL (Iter-9)	32.669	38.428	31.093	34.119	31.479	28.224
+SIPL (Iter-10)	32.669	38.428	31.094	34.118	31.480	28.224
with GT (Upper Bound)	33.099	38.792	31.551	34.615	31.838	28.701

Table 17: Detailed performance metrics (PSNR) across ten iterations on the **Five-Task Benchmark**. Similar to the three-task setting, the most significant gain is achieved in the first step, followed by stable convergence.

Method	Average	Deraining	Dehazing	Denoising	Deblurring	Low-Light
PromptIR (Baseline)	29.15	36.37	26.54	31.47	28.71	22.68
+SIPL (Iter-1)	30.53	38.09	30.51	31.45	29.35	23.23
+SIPL (Iter-2)	30.58	38.22	30.52	31.50	**29.39**	23.26
+SIPL (Iter-3)	30.58	38.24	30.52	**31.53**	29.37	23.26
+SIPL (Iter-4)	**30.60**	**38.38**	**30.53**	31.50	29.37	23.21
+SIPL (Iter-5)	30.58	38.28	**30.53**	31.48	29.34	**23.28**
+SIPL (Iter-6)	30.55	38.10	30.54	31.46	29.35	23.24
+SIPL (Iter-7)	30.53	38.08	30.53	31.48	29.34	23.24
+SIPL (Iter-8)	30.54	38.15	30.52	31.47	29.33	23.23
+SIPL (Iter-9)	30.52	38.07	30.51	31.44	29.35	23.22
+SIPL (Iter-10)	30.53	38.09	30.52	31.45	29.35	23.22
with GT (Upper Bound)	30.82	38.40	30.60	31.54	29.43	24.11

Rapid Initial Improvement and Graceful Convergence A consistent pattern emerges from both benchmarks: the vast majority of the performance improvement is achieved within the very first iteration ($t = 1$). On the Five-Task benchmark, this initial step accounts for a +1.38 dB PSNR gain, which is over 95% of the total improvement observed towards the peak performance at $t = 4$ (+1.45 dB). This indicates that our Proxy Fusion mechanism is highly effective at providing a strong initial correction signal.

Beyond the first step, the performance exhibits graceful convergence, quickly settling onto a stable plateau. The fluctuations in subsequent iterations are minimal (e.g., within $\pm$0.004 dB on the Three-Task average PSNR after $t = 2$). This is a critical finding. As demonstrated in our main paper's ablation study, a naive iterative loop (feeding the output pixels back as input) leads to catastrophic performance degradation due to the amplification of artifacts from out-of-distribution inputs. The stability of SIPL empirically proves that our feature-level refinement, guided by the Privileged Dictionary (PD), is a fundamentally different and robust process. The PD acts as a constant, degradation-agnostic anchor, ensuring the iterative process does not diverge but instead converges towards the learned manifold of high-quality images.

Understanding Performance Saturation and the Role of Pseudo-PI The performance saturation observed in later iterations is an expected and desirable behavior. Our iterative process relies on the model's own output as *pseudo-privileged information* (pseudo-PI). While this output is significantly cleaner than the original degraded input, it is still imperfect. The model iteratively refines its output until the pseudo-PI is no longer informative enough to elicit further significant improvements from the PD. The remaining performance gap between the saturated performance (e.g., 30.60 dB on Five-Task) and the GT-guided upper bound (30.82 dB) quantifies the inherent limitation of using a pseudo-guide instead of a perfect one. This gap highlights a potential avenue for future research in improving the fidelity of the pseudo-PI.

Justification for Single-Step Refinement in Main Experiments This detailed analysis provides a strong empirical foundation for our choice of a single refinement step ($t = 1$) as the default configuration in the main paper. This single step captures the most substantial portion of the performance gain while being computationally efficient (requiring only two forward passes). While further iterations can yield marginal improvements, the diminishing returns suggest that a single, powerful refinement step offers the most practical and compelling trade-off between restoration quality and inference latency.

A.6 Ethics Statement

This work is in full compliance with the ICLR Code of Ethics. Our research proposes a foundational algorithm for low-level image enhancement, aiming to contribute to the scientific community by providing a more effective and efficient solution.

We uphold the principles of scientific excellence and transparency. All datasets utilized in our experiments are publicly available and open-source, and we provide a detailed description of our methodology and experimental setup to ensure reproducibility. As a computer vision algorithm, our work does not directly involve personal data or raise immediate concerns regarding privacy or societal fairness. We have, however, transparently discussed the potential limitations of our method in Section 4.3 of this paper. While we have not identified direct negative societal impacts, we acknowledge that as a foundational technology, its subsequent applications should be developed and deployed responsibly by others. We declare no competing interests in this work.

A.7 Reproducibility Statement

To ensure the reproducibility of our findings, we provide comprehensive documentation of our methodology, experimental setup, and resources. The complete source code, including scripts to reproduce all experiments and results presented in this paper, is available in the supplementary materials. A detailed description of our proposed model and algorithm is provided in Section 3 of the main paper. Further implementation details, including all hyperparameter settings, library dependencies, and the computational environment used for our experiments, are thoroughly documented in Appendix.

A.8 The Use of Large Language Models

The authors affirm that Large Language Models (LLMs) were not used for the core scientific contributions of this work. Specifically, LLMs were not utilized in the ideation phase, for the development of the proposed methodology, in the design or execution of experiments, or for the analysis of results. The conclusions presented in this paper were drawn entirely by the authors.

Following the completion of the main manuscript, an LLM-based tool was used as a general-purpose writing assistant. Its role was strictly limited to performing grammar checks and providing suggestions for improving the clarity and flow of sentences to enhance readability. The scientific content and integrity of the paper were not influenced by the use of the LLM.