Debiased All-in-one Image Restoration with Task Uncertainty Regularization

Gang Wu, Junjun Jiang, Yijun Wang, Kui Jiang*, Xianming Liu

Faculty of Computing, Harbin Institute of Technology, Harbin 150001, China
{gwu, jiangjunjun, jiangkui, csxm}@hit.edu.cn, 2021112933@stu.hit.edu.cn

Abstract

All-in-one image restoration is a fundamental low-level vision task with significant real-world applications. The primary challenge lies in addressing diverse degradations within a single model. While current methods primarily exploit task prior information to guide the restoration models, they typically employ uniform multi-task learning, overlooking the heterogeneity in model optimization across different degradation tasks. To eliminate the bias, we propose a task-aware optimization strategy, that introduces adaptive task-specific regularization for multi-task image restoration learning. Specifically, our method dynamically weights and balances losses for different restoration tasks during training, encouraging the implementation of the most reasonable optimization route. In this way, we can achieve more robust and effective model training. Notably, our approach can serve as a plug-and-play strategy to enhance existing models without requiring modifications during inference. Extensive experiments in diverse all-in-one restoration settings demonstrate the superiority and generalization of our approach. For example, AirNet retrained with TUR achieves average improvements of **1.16 dB** on three distinct tasks and **1.81 dB** on five distinct all-in-one tasks. These results underscore TUR's effectiveness in advancing the SOTAs in all-in-one image restoration, paving the way for more robust and versatile image restoration.

Introduction

Image restoration, a fundamental challenge in computer vision and image processing, seeks to recover high-quality images from their degraded counterparts. Classical image restoration relied on hand-crafted priors or domain-specific assumptions about the degradation process (Banham and Katsaggelos 1997).The advent of deep learning has ushered in a paradigm shift in this field (Su, Xu, and Yin 2022). Modern data-driven methods can now learn powerful priors directly from large-scale datasets, resulting in more robust solutions for specific image restoration tasks. These tasks include, but are not limited to, denoising (Zhang et al. 2023b), deraining (Jiang et al. 2020; Zamir et al. 2022; Chen et al. 2023), deblurring (Chen et al. 2022a; Cui et al. 2024a), and dehazing (Qin et al. 2020; Cui et al. 2022, 2023). While single-task models often struggle in real-world scenarios where multiple degradation types emerge. This limitation has given rise to the emerging field of *all-in-one* image restoration, which aims to develop unified model capable of handling multiple degradation types (Liu et al. 2022; Li et al. 2022; Potlapalli et al. 2023; Valanarasu, Yasarla, and Patel 2022; Zhang et al. 2023a; Kong, Dong, and Zhang 2024; Wu et al. 2024). For more recent achievements in the field of all-in-one image restoration, please refer to our survey paper (Jiang et al. 2024a) and continuously updated project[1].

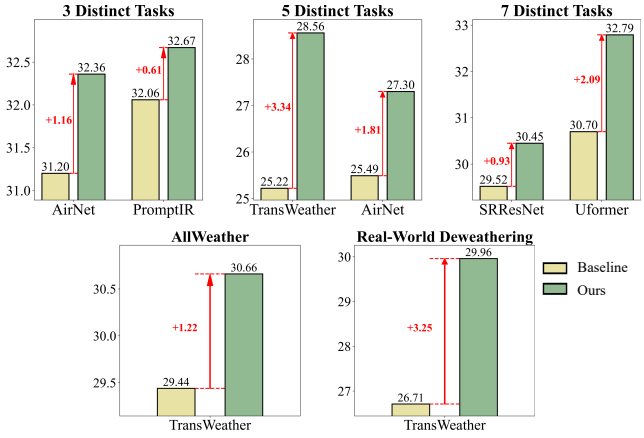

Figure 1: Performance improvements of our proposed approach across multiple all-in-one image restoration settings, demonstrating consistent and significant gains over existing methods for various tasks and models.

Existing approaches to all-in-one image restoration have primarily focused on incorporating task priors to guide model learning. These methods can be broadly categorized into two types: two-stage approaches (Li et al. 2022; Zhang et al. 2023a; Kong, Dong, and Zhang 2024) that pre-train task-specific representations, and prompt-based methods (Valanarasu, Yasarla, and Patel 2022; Potlapalli et al. 2023) that directly embed task priors into the restoration model using learnable prompts. While these approaches have improved model performance, they often overlook a critical aspect: the optimization process in multi-task learning. Cur-

*Corresponding Author

[1]https://github.com/Harbinzzy/All-in-One-Image-Restoration-Survey

rent methods typically employ uniform multi-task learning, where multiple tasks are mixed and joint trained without considering their interrelationships. However, in multi-task learning, the construction of appropriate multi-task objectives is crucial for optimal performance (Vandenhende et al. 2022). Recent work MioIR (Kong, Dong, and Zhang 2024) has made progress in this direction by proposing a sequential training strategy that groups degradation tasks based on their observed characteristics. However, it lacks the flexibility to address new degradation tasks. The key challenge remains: *How can we develop a flexible and effective approach to enhance the interpretation of task priors during the multi-task optimization process for all-in-one image restoration?*

To address this challenge, we revisit all-in-one image restoration from a Bayesian perspective. We propose a task-grouped regularization for multiple degradation tasks. It incorporates a task-related regularization that captures the distinct characteristics of multiple degradation tasks, moving beyond the uniform priors used in existing mixed multi-task training methods. More importantly, this task-dependent regularization offers a flexible and dynamic mechanism for combining objective losses across multiple degradation tasks. Drawing inspiration from Bayesian modeling of neural network uncertainty (Abdar et al. 2021), we introduce Task Uncertainty Regularization (TUR). This practical implementation of task-grouped regularization provides a principled way to balance and optimize multiple restoration tasks simultaneously, adapting to the unique characteristics of each degradation type. Furthermore, TUR offers a plug-and-play solution to enhance existing all-in-one models through retraining. As illustrated in Figure 1, we validate our approach across five distinct all-in-one image restoration settings and multiple models, demonstrating consistently superior performance. Notably, AirNet retrained with TUR achieves average improvements of **1.16 dB** on three distinct tasks and **1.81 dB** on five distinct all-in-one tasks. Furthermore, TransWeather shows remarkable gains when enhanced with TUR, with improvements of **1.22 dB** on synthetic deweathering tasks and an impressive **3.41 dB** on real-world deweathering datasets. These substantial improvements underscore the effectiveness and versatility of our TUR approach in advancing the state-of-the-art in all-in-one image restoration.

The key contributions of our work are as follows:

- We introduce a principled task uncertainty estimation for all-in-one image restoration that effectively captures the homoscedastic uncertainties associated with different degradation types.
- We derive a novel loss function that adaptively combines multi-task losses for all-in-one image restoration. This allows for automatic learning of the optimal trade-off between tasks.
- We provide detailed evaluation results across diverse all-in-one image restoration settings, demonstrating the superior performance obtained by our approach.

Related Work

Image Restoration

Image restoration aims to recover high-quality images from degraded input. Early approaches primarily relied on handcrafted priors and assumptions about the underlying degradation process (Banham and Katsaggelos 1997). With the advent of deep learning, however, researchers have shifted towards data-driven priors learned directly from large-scale datasets (Su, Xu, and Yin 2022), thereby significantly advancing performance in task-specific scenarios, including image deraining (Jiang et al. 2021; Zamir et al. 2022; Chen et al. 2023), dehazing (Qin et al. 2020; Cui et al. 2023), deblurring (Wang et al. 2022; Chen et al. 2022a; Cui et al. 2024a), and denoising (Zhang et al. 2023b).

In practice, real-world images commonly suffer from multiple degradations simultaneously (Jiang et al. 2024a). To address this complexity, "all-in-one" image restoration methods have emerged (Liu et al. 2022; Li et al. 2022; Potlapalli et al. 2023; Zhang et al. 2023a; Kong, Dong, and Zhang 2024; Wu et al. 2024). These methods present unique challenges, particularly in effectively capturing and leveraging task-specific priors to handle a variety of degradation types. For example, Li *et al.* (Li et al. 2022) apply self-supervised learning (He et al. 2020) to obtain a degradation-aware feature encoder, while Kong *et al.* (Kong, Dong, and Zhang 2024) rely on supervised classification to extract degradation-specific features. Zhang *et al.* (Zhang et al. 2023a) introduce learnable principal component analysis to generate task-oriented representations, and Cui *et al.* (Cui et al. 2024b) analyze task prompts from a spectral perspective. Potlapalli *et al.* (Potlapalli et al. 2023) propose PromptIR, which employs learnable weights as adaptive task prompts. Moreover, integrating powerful pretrained models, including multimodal priors, into all-in-one restoration frameworks is gaining traction (Conde, Geigle, and Timofte 2024; Ai et al. 2024; Luo et al. 2024; Jiang et al. 2024b). Furthermore, Wu *et al.* (Wu et al. 2024) investigate the underlying conflicts among multiple degradation tasks and introduce a robust multi-task learning strategy to enhance existing all-in-one restoration models. Guo *et al.* (Guo et al. 2024) introduce a composite-degradation all-in-one restoration setting, inspiring follow-up works (Cao, Meng, and Cao 2024; Li et al. 2024) that further push the boundaries of general-purpose restoration. A notable application domain for all-in-one restoration is adverse weather conditions, where complex and concurrent degradations such as rain streaks, haze, and snow coexist. Several studies have presented unified backbones to tackle these scenarios (Valanarasu, Yasarla, and Patel 2022; Zhu et al. 2023; Chen et al. 2022b; Sun et al. 2024). Recently, diffusion-based methods have also been explored for weather-related degradations, demonstrating strong generative capabilities (Özdenizci and Legenstein 2023; Zheng et al. 2024; Chen et al. 2024).

Uncertainty Estimation

The idea of representing uncertainties or confidence scores alongside predictions has been explored in various machine learning and computer vision problems (Abdar et al. 2021).

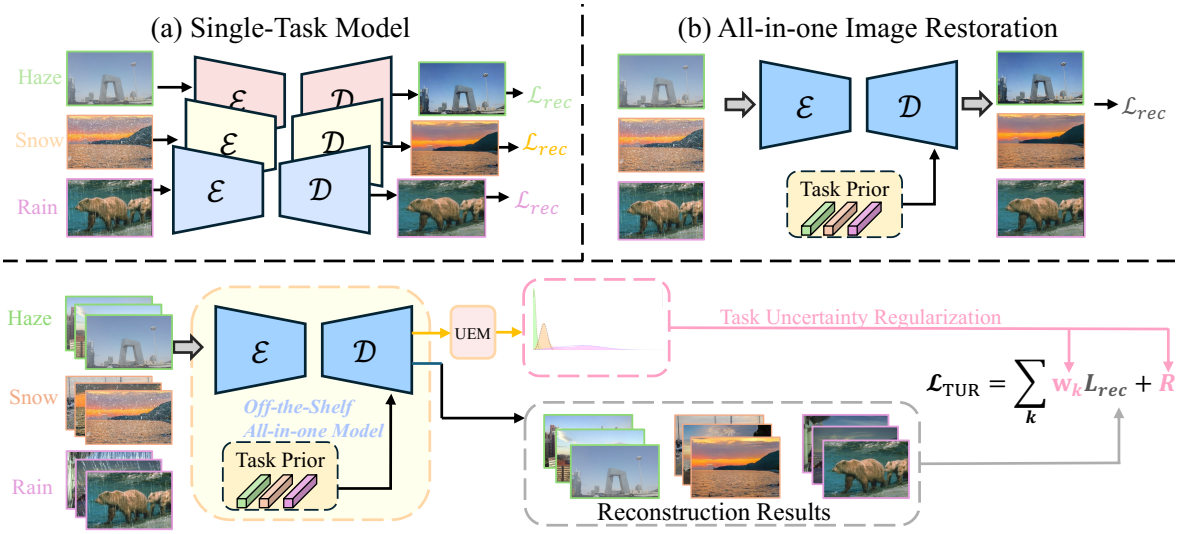

(c) Unified Model with *Task Uncertainty Regularization*

Figure 2: Evolution of image restoration frameworks. (a) Single-task models, each dedicated to a specific degradation type. (b) All-in-one models capable of handling multiple degradations by incorporating task priors. (c) Our proposed Task Uncertainty Regularization (TUR), integrated as a plug-and-play way within existing all-in-one models. By introducing an Uncertainty Estimation Module (UEM), our approach adaptively refines tasks during training.

A common categorization distinguishes between two types of uncertainties - epistemic (or model) uncertainty that accounts for uncertainties in the model itself due to limited training data; and aleatoric (or data) uncertainty that captures noise inherent in the input data (Kendall and Gal 2017). Within the context of deep learning, several works have investigated modeling epistemic uncertainties through techniques like Monte Carlo dropout(Gal and Ghahramani 2016), ensemble methods(Lakshminarayanan, Pritzel, and Blundell 2017) etc. On the other hand, aleatoric uncertainty modeling has proven beneficial to target tasks. For low-level image restoration tasks, several researches have been made (Hong et al. 2022; Fang et al. 2022; Yang et al. 2022; Ning et al. 2021, 2022; Huang, Luo, and He 2023). Ning *et al.* exploits the uncertainty map of the high-frequency regions and utilizes this to refine the super-resolution results.

Method

Preliminaries

Problem Formulation Let $\mathcal{X}$ and $\mathcal{Y}$ denote the spaces of high-quality and degraded images, respectively. All-in-one image restoration aims to learn a unified mapping $\mathcal{M} : \mathcal{Y} \to \mathcal{X}$ that restores a clean image $x \in \mathcal{X}$ from its degraded counterpart $y \in \mathcal{Y}$, where $y = \mathcal{D}(x)$ and $\mathcal{D}$ represents multiple degradation processes. Formally, we seek to learn a neural network model $\mathcal{M}_\theta : \mathcal{Y} \to \mathcal{X}$ parameterized by θ, such that $\mathcal{M}_\theta(y) \approx x$.

Conventional image restoration methods typically optimize θ using maximum likelihood estimation. The likelihood of observing the clean image x given the degraded input y is modeled as:

$$p(x|y,\theta) = \mathcal{N}(x|\mathcal{M}_\theta(y), \sigma^2 I), \tag{1}$$

where $\mathcal{N}(\cdot|\mu, \Sigma)$ denotes a Gaussian distribution with mean μ and covariance Σ, and σ^2 is a fixed scalar representing observation noise.

The corresponding negative log-likelihood loss function is:

$$\mathcal{L}_{\text{rec}}(X, Y) = \sum_{x,y \in X,Y} \|x - \mathcal{M}_\theta(y)\|_2^2. \tag{2}$$

It's worth noting that the prior σ does not affect the optimization result, leading to the commonly used MSE loss. Similarly, a Laplacian prior would result in the frequently used MAE loss.

Challenges of Multiple Degradation Types For all-in-one image restoration encompassing K degradation tasks, we combine multiple single degradations. Let $\mathcal{T} = \{t_1, \ldots, t_K\}$ denote the set of K restoration tasks. The overall loss function becomes:

$$\mathcal{L}_{\text{all-in-one}}(X, Y) = \sum_{k=1}^{K} \mathcal{L}_{\text{rec}}(X_k, Y_k), \tag{3}$$

where X_k and Y_k represent the clean and degraded subsets for degradation task k, respectively. While existing methods have focused on improving task priors in restoration models, they often overlook the optimization objective. The uniform combination of multiple task losses in current approaches fails to capture the unique distribution of each task, which is a key consideration in multi-task learning (Vandenhende et al. 2022).

Task Uncertainty Regularization

To address this challenge, we incorporate uncertainty estimation into the all-in-one image restoration framework. From a Bayesian perspective, we model the restoration process as:

$$\mathcal{L}_{\theta,\sigma}(X,Y) = \sum_{x,y\in X,Y} \frac{1}{2\sigma^2}\|x-\mathcal{M}_\theta(y)\|_2^2 + \frac{1}{2}\log\sigma^2, \quad (4)$$

where σ^2 represents model uncertainty estimated by a learnable module. In contrast to Eq. (1), this estimates the posterior data distribution $p(x|y,\theta,\sigma) = \mathcal{N}(x|\mathcal{M}_\theta(y),\sigma^2 I)$.

Building on this, we introduce Task Uncertainty Regularization to capture task-dependent uncertainty, $p(x|y,\theta,\sigma_k) = \mathcal{N}(x|\mathcal{M}_\theta(y),\sigma_k^2 I)$ for multi-task learning, where σ_k represents the task uncertainty. For task k, our loss function is formulated as:

$$\mathcal{L}_{(\theta,\sigma_k)}(X_k,Y_k) = \frac{1}{2\sigma_k^2}\mathcal{L}_{rec}(X_k,Y_k) + \frac{1}{2}\log\sigma_k^2. \quad (5)$$

To jointly optimize across all tasks, we propose the following multi-task loss function:

$$\mathcal{L}_{\text{TUR}} = \sum_{k=1}^{K} \mathcal{L}_{(\theta,\sigma_k)}(X_k,Y_k) \quad (6)$$

Remarks

Adaptive Multi-Task Combining For clarity, we reformulate Eq. (6) as follows:

$$\begin{aligned}\mathcal{L}_{\text{TUR}} =& \underbrace{\frac{1}{2\sigma_1^2}\mathcal{L}_{rec}(X_1,Y_1) + \cdots + \frac{1}{2\sigma_k^2}\mathcal{L}_{rec}(X_k,Y_k)}_{\text{Adaptive Combining Multi-Task Loss}} \\ &+ \underbrace{\frac{1}{2}\log\sigma_1^2 + \cdots + \frac{1}{2}\log\sigma_k^2}_{\text{Task Uncertainty Estimation}} \quad (7) \\ =& \sum_k w_k\mathcal{L}_{rec}(X_k,Y_k) + R\end{aligned}$$

where $w_k = 1/2\sigma_k^2$, providing flexible and dynamic multi-task reweighting, and $R = \sum_k \log\sigma_k$ represents the estimated task uncertainty. This formulation automatically learns to balance the contributions of different tasks based on their estimated uncertainties and provides an adaptive, debiased optimization objective.

Implementation Details Our TUR approach offers a flexible enhancement for all-in-one image restoration models. During training, as illustrated in Figure 2, we introduce an auxiliary projection head consisting of three stacked convolutional layers and activation layers as the uncertainty estimation module. This can be easily integrated into existing models without significant modifications. Importantly, the UEM is not required during inference, maintaining the original model structure. Thus, TUR functions as a plug-and-play module that optimizes training outcomes without altering existing restoration models.

Experiments

Experimental Settings

Here we provide comprehensive experimental settings from the literature, encompassing a wide range of degradation types and task combinations.

- ***Setting 1*: Seven Degradation Tasks**: Following the setup proposed in (Kong, Dong, and Zhang 2024), we take the DIV2K and Flickr2K as the training dataset, and evaluate our method on seven distinct degradation tasks, including super-resolution, blur, noise, JPEG compression, rain, haze, and low-light enhancement.
- ***Setting 2*: Rain-Haze-Noise**: Following (Li et al. 2022), we focus on three common degradation tasks: deraining, dehazing, and denoising. For denoising, we combine the BSD400 (Arbeláez et al. 2011) and WED (Ma et al. 2017) datasets for training. Noisy images are generated with Gaussian noise at three levels: $\sigma = 15, 25, 50$, and testing is performed on the BSD68 dataset. The Rain100L (Yang et al. 2017) dataset is used for the deraining task, and the SOTS (Li et al. 2019) dataset is employed for dehazing.
- ***Setting 3*: Rain-Haze-Noise-Blur-Dark**: Following the experimental setup in (Zhang et al. 2023a), we evaluate our method on five degradation tasks, which introduces GoPro (Nah, Kim, and Lee 2017) for deblurring, and LOL (Wei et al. 2018) for low-light enhancement beyond the 3 tasks in *Setting2*.
- ***Setting 4*: Synthetic and Real-World Deweathering**: Aligning with the setup in (Valanarasu, Yasarla, and Patel 2022), we investigate the performance of our approach on three deweather tasks: snow removal, rain streak and fog removal, and raindrop removal. The training data, termed "All-Weather," includes images from Snow100K (Liu et al. 2018), Raindrop (Qian et al. 2018), and Outdoor-Rain (Li, Cheong, and Tan 2019) datasets. Moreover, we evaluate the proposed approach on a large-scale and real-world deweathering settings following (Zhu et al. 2023), combing the multiple real-world datasets such as SPA dataset for deraining, RealSnow for desnowing, and REVIDE for dehazing.

Results and Analysis

We present and analyse the experiment results obtained by incorporating our proposed TUR into the baseline models across the four diverse all-in-one restoration settings.

Setting 1 Table 1 presents the results for the seven degradation tasks proposed in (Kong, Dong, and Zhang 2024). We evaluate our approach on two baseline models: SRResNet (Ledig et al. 2017) and Uformer (Wang et al. 2022). The results demonstrate the efficacy of our uncertainty-guided multi-task learning strategy in enhancing the performance of these baseline models across all degradation tasks. When applied to the Uformer, our method attains an average PSNR of 32.79 dB, surpassing the baseline Uformer by 2.09 dB and Uformer-S by 1.58 dB. Significant improvements are observed in multiple single tasks like Rain, Low-Light, and

Methods	SR	Blur	Noise	JPEG	Rain	Haze	Dark	Avg.
SRResNet	25.52	30.01	30.49	32.46	32.38	25.57	30.20	29.52
SRResNet-S (Kong, Dong, and Zhang 2024)	25.72	30.49	30.67	32.73	32.81	25.78	30.45	29.84
SRResNet (Ours)	**25.55**	**30.65**	**30.65**	**32.92**	**35.20**	**26.16**	**32.04**	**30.45**
Uformer	25.80	30.53	30.84	33.13	33.39	27.93	33.27	30.70
Uformer-S (Kong, Dong, and Zhang 2024)	26.07	31.11	30.96	33.27	35.96	28.29	32.80	31.21
Uformer (Ours)	**26.11**	**31.51**	**31.20**	**33.46**	**38.13**	**30.91**	**38.24**	**32.79**

Table 1: Comparison on 7 distinct degradation tasks introduced in (Kong, Dong, and Zhang 2024).

Methods	Dehazing on SOTS	Deraining on Rain100L	Denoising on BSD68			Average
			$\sigma = 15$	$\sigma = 25$	$\sigma = 50$	
BRDNet (Tian, Xu, and Zuo 2020)	23.23/0.895	27.42/0.895	32.26/0.898	29.76/0.836	26.34/0.693	27.80/0.843
LPNet (Gao et al. 2019)	20.84/0.828	24.88/0.784	26.47/0.778	24.77/0.748	21.26/0.552	23.64/0.738
FDGAN (Dong et al. 2020)	24.71/0.929	29.89/0.933	30.25/0.910	28.81/0.868	26.43/0.776	28.02/0.883
MPRNet (Zamir et al. 2021)	25.28/0.955	33.57/0.954	33.54/0.927	30.89/0.880	27.56/0.779	30.17/0.899
DL (Fan et al. 2019)	26.92/0.931	32.62/0.931	33.05/0.914	30.41/0.861	26.90/0.740	29.98/0.876
AirNet (Li et al. 2022)	27.94/0.962	34.90/0.968	33.92/0.933	31.26/0.888	28.00/0.797	31.20/0.910
AirNet (Ours)	**30.41/0.976**	**38.04/0.983**	**33.97/0.931**	**31.32/0.886**	**28.05/0.795**	**32.36/0.914**
PromptIR (Potlapalli et al. 2023)	30.58/0.974	36.37/0.972	33.98/0.933	31.31/0.888	28.06/0.799	32.06/0.913
PromptIR (Ours)	**31.17/0.978**	**38.57/0.984**	**34.06/0.932**	**31.40/0.887**	**28.13/0.797**	**32.67/0.916**

Table 2: Comparisons under the three-degradation all-in-one setting: a unified model is trained on a combined set of images obtained from all degradation types and levels.

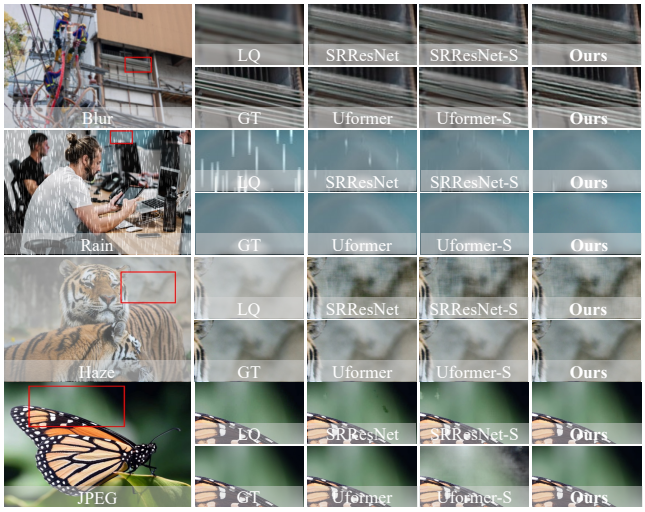

Figure 3: Visual results of *setting 1*. More visual results are available at supplementary material. Zoom in for details.

Haze compared to both original and improved ones proposed in (Kong, Dong, and Zhang 2024), highlighting the effectiveness of our uncertainty-guided multi-task learning in capturing task-specific information.

Figure 3 provides visual results for the seven distinct degradation tasks, illustrating the superior restoration quality achieved by our method across various degradation types.

Setting 2 Table 2 presents the results of our uncertainty-guided multi-task learning approach applied to AirNet (Li et al. 2022) and PromptIR (Potlapalli et al. 2023) for deraining, dehazing, and denoising tasks. Our method demonstrates significant enhancements to both models' performance across all tasks. For AirNet, we observe substantial improvements across all tasks. The average PSNR increased from 31.20 dB to 32.36 dB, a gain of 1.16 dB. Particularly notable are the improvements in dehazing and deraining. Even for the challenging denoising task, our method shows consistent gains across different noise levels, with improvements around 0.05 dB. Interestingly, our retrained AirNet now outperforms the PromptIR. This significant improvement provides strong evidence that the vanilla mixed training paradigm adopted in previous methods often leads to biased suboptimal results. Our method's ability to surpass a newer, more advanced model by retraining an older one underscores the effectiveness of our adaptive task-specific regularization approach. When applied to PromptIR, our approach yields further enhancements, pushing the boundaries of all-in-one image restoration performance.

Figure 4 provides visual results for the three distinct degradation tasks, demonstrating the superior restoration quality achieved by our retrained models. The images show clearer and more detailed results across multiple tasks, highlighting the effectiveness of our approach in handling diverse degradation types.

Setting 3 Table 3 presents the results of our approach applied to Transweather (Valanarasu, Yasarla, and Patel 2022)

Methods	Dehazing on SOTS	Deraining on Rain100L	Denoising on BSD68	Deblurring on GoPro	Low-Light on LOL	Average
NAFNet (Chen et al. 2022a)	25.23/0.939	35.56/0.967	31.02/0.883	26.53/0.808	20.49/0.809	27.76/0.881
MPRNet (Zamir et al. 2021)	24.27/0.937	38.16/0.981	31.35/0.889	26.87/0.823	20.84/0.824	28.27/0.890
SwinIR (Liang et al. 2021)	21.50/0.891	30.78/0.923	30.59/0.868	24.52/0.773	17.81/0.723	25.04/0.835
DL (Fan et al. 2019)	20.54/0.826	21.96/0.762	23.09/0.745	19.86/0.672	19.83/0.712	21.05/0.743
TAPE (Liu et al. 2022)	22.16/0.861	29.67/0.904	30.18/0.855	24.47/0.763	18.97/0.621	25.09/0.801
IDR (Zhang et al. 2023a)	25.24/0.943	35.63/0.965	31.60/0.887	27.87/0.846	21.34/0.826	28.34/0.893
Transweather (Valanarasu, Yasarla, and Patel 2022)	21.32/0.885	29.43/0.905	29.00/0.841	25.12/0.757	21.21/0.792	25.22/0.836
Transweather (Ours)	**29.68/0.966**	**33.09/0.952**	**30.40/0.869**	**26.63/0.815**	**23.02/0.838**	**28.56/0.888**
AirNet (Li et al. 2022)	21.04/0.884	32.98/0.951	30.91/0.882	24.35/0.781	18.18/0.735	25.49/0.846
AirNet (Ours)	**27.59/0.954**	**33.95/0.962**	**30.93/0.875**	**26.13/0.801**	**17.88/0.772**	**27.30/0.873**

Table 3: Comparative results on five distinct tasks in all-in-one image restoration.

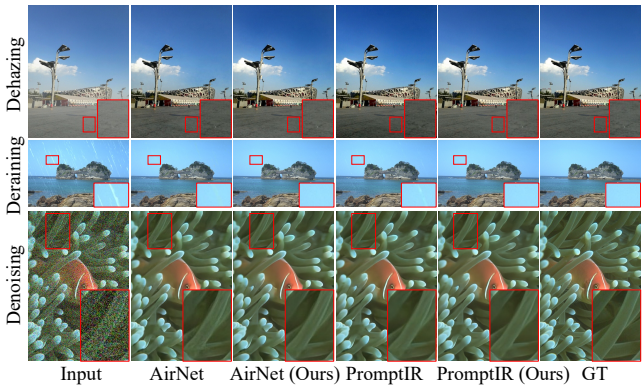

Figure 4: Visual results of 3 distinct degradation tasks. Our retrained models obtain more clear visual results across multiple tasks.

Datasets	Methods	PSNR ↑	SSIM ↑
Outdoor-Rain	All-in-One (Li, Tan, and Cheong 2020)	24.71	0.8980
	WeatherDiff$_{128}$(Özdenizci and Legenstein 2023)	29.72	0.9216
	TransWeather (Valanarasu, Yasarla, and Patel 2022)	28.83	0.9000
	TransWeather (Ours)	**29.75**	**0.9073**
Snow100K	DDMSNet (Zhang et al. 2021)	28.85	0.8772
	All-in-One (Li, Tan, and Cheong 2020)	28.33	0.8820
	WeatherDiff$_{128}$(Özdenizci and Legenstein 2023)	29.58	0.8941
	TransWeather (Valanarasu, Yasarla, and Patel 2022)	29.31	0.8879
	TransWeather (Ours)	**30.62**	**0.9086**
RainDrop	All-in-One (Li, Tan, and Cheong 2020)	31.12	0.9268
	WeatherDiff$_{128}$(Özdenizci and Legenstein 2023)	29.66	0.9225
	TransWeather (Valanarasu, Yasarla, and Patel 2022)	30.17	0.9157
	TransWeather (Ours)	**31.61**	**0.9330**
Average	All-in-One (Li, Tan, and Cheong 2020)	27.12	0.8933
	WeatherDiff$_{128}$(Özdenizci and Legenstein 2023)	29.65	0.9127
	TransWeather (Valanarasu, Yasarla, and Patel 2022)	29.44	0.9012
	TransWeather (Ours)	**30.66**	**0.9163**

Table 4: Comparison on deweathering tasks with Allweather dataset (Valanarasu, Yasarla, and Patel 2022).

and AirNet (Li et al. 2022) for five distinct tasks: dehazing, deraining, denoising, deblurring, and low-light enhancement. Our method demonstrates significant improvements for both models. For Transweather, we observe substantial enhancements across all tasks. The average PSNR increased from 25.22 dB to 28.56 dB, a remarkable gain of 3.34 dB. Particularly notable are the improvements in dehazing and deraining. It's worth noting that our method's ability to significantly improve the performance of Transweather, transforming it from one of the lower-performing models to a top performer, underscores the effectiveness of our adaptive task-specific regularization. When applied to AirNet, our approach also yields consistent improvements. The average PSNR improved from 25.49 dB to 27.30 dB, an increase of 1.81 dB. Interestingly, while the PSNR for low-light enhancement on the LoL dataset decreased slightly, we observed a significant improvement in SSIM. This suggests that our retrained AirNet achieves notable enhancements in content structure, albeit with some potential overexposure affecting overall brightness.

Setting 4 In Table 4, we apply our method to TransWeather (Valanarasu, Yasarla, and Patel 2022) on three synthetic tasks. Our approach consistently improves PSNR and SSIM across all tasks, indicating a stronger ability to handle various synthetic weather degradations and produce more faithful reconstructions.

Table 5 further demonstrates the effectiveness of our approach on real-world scenarios. Notably, on the SPA+ dataset for deraining, our method not only surpasses the original TransWeather by a substantial margin but also outperforms the previous best method, WGWSNet (Zhu et al. 2023). These results highlight our model's enhanced generalization capabilities and its capacity to handle the increased complexity and variability of real-world weather conditions. These results underscore the effectiveness and versatility of our approach in handling diverse deweathering tasks across both synthetic and real-world scenarios, consistently outperforming state-of-the-art methods.

Ablation Studies

Task Uncertainty Estimation To evaluate the effectiveness of our proposed Task Uncertainty Regularization (TUR), we conduct an ablation study on the seven distinct degradation tasks introduced in (Kong, Dong, and Zhang 2024). Table 6 presents the results of SRResNet with TUR or without it (vanilla Uncertainty Estimation). By dynamically adjusting the regularization for each degradation task,

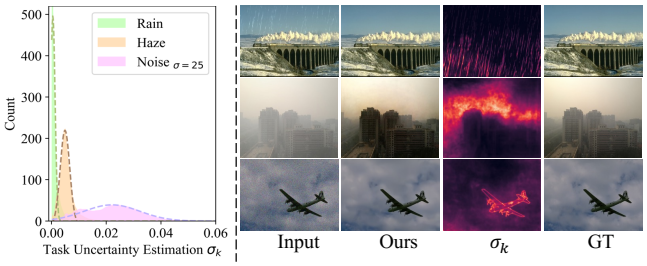

Figure 5: Illustration of task uncertainty. (a) Distribution of task uncertainty across multiple degradation types, demonstrating clear distinctions between different degradations. (b) Visualization of the sample-wise uncertainty maps, where higher uncertainty values align with more severely degraded regions.

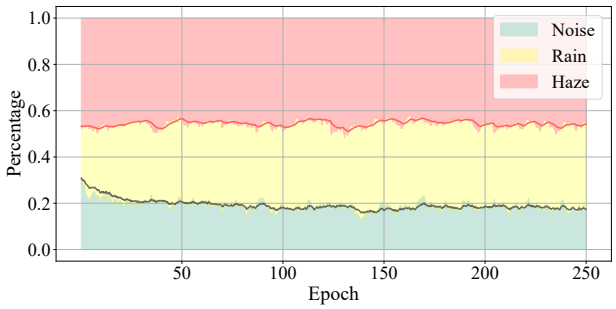

Figure 6: Dynamic multi-task balancing achieved by the proposed TUR for AirNet. Normalized mean values of task uncertainty for noise, rain, and haze degradations over 250 training epochs are presented, illustrating the adaptive balancing multiple restoration tasks.

TUR enables more adaptive and robust multi-task learning, leading to superior overall performance in all-in-one image restoration.

Distinct Task Uncertainty Distribution We analyzed the distribution of estimated uncertainties for each degradation task using AirNet in Setting 2. Figure 5(a) reveals distinct uncertainty distributions across multiple degradation tasks, highlighting a key limitation of vanilla mixed training in existing approaches and reinforcing the efficacy of our proposed TUR method. Moreover, Figure 5(b) provides a visualization of task uncertainty for individual samples. Notably, the uncertainty maps effectively highlight critical image features such as rain streaks, remote haze, and textured regions.

Dynamic Multi-Task Reweighting Figure 6 illustrates the adaptive task weighting for multiple degradation tasks during training. In contrast to the mix training strategy adopted in most existing methods, our proposed TUR demonstrates dynamic task weight adjustments throughout the training process, eventually converging to model-specific adaptive weights.

Datasets	Methods	PSNR ↑	SSIM ↑
SPA+	Chen et al. (Chen et al. 2022b)	37.32	0.97
	WGWSNet (Zhu et al. 2023)	38.94	0.98
	TransWeather (Valanarasu, Yasarla, and Patel 2022)	33.64	0.93
	TransWeather (Ours)	**39.78**	**0.98**
RealSnow	Chen et al. (Chen et al. 2022b)	29.37	0.88
	WGWSNet (Zhu et al. 2023)	29.46	0.85
	TransWeather (Valanarasu, Yasarla, and Patel 2022)	29.16	0.82
	TransWeather (Ours)	**29.72**	**0.91**
REVIDE	Chen et al. (Chen et al. 2022b)	20.10	0.85
	WGWSNet (Zhu et al. 2023)	20.44	0.87
	TransWeather (Valanarasu, Yasarla, and Patel 2022)	17.33	0.82
	TransWeather (Ours)	**20.38**	**0.88**
Average	Chen et al. (Chen et al. 2022b)	28.93	0.90
	WGWSNet (Zhu et al. 2023)	29.61	0.90
	TransWeather (Valanarasu, Yasarla, and Patel 2022)	26.71	0.86
	TransWeather (Ours)	**29.96**	**0.92**

Table 5: Comparison on de-weathering tasks on real-world datasets following (Zhu et al. 2023)

Methods	SR	Blur	Noise	JPEG	Rain	Haze	Dark	Avg.
SRResNet	25.52	30.01	30.49	32.46	32.38	25.57	30.20	29.52
UE	25.54	30.30	30.59	32.69	33.45	26.15	31.89	30.09
TUE	**25.55**	**30.65**	**30.65**	**32.92**	**35.20**	**26.16**	**32.04**	**30.45**

Table 6: Ablation study about with or without the proposed task grouped regularization on 7 distinct degradation tasks.

Discussion and Limitation

Our task uncertainty regularization approach demonstrates significant improvements across diverse all-in-one image restoration settings. It provides more stable convergence, improved generalization, and ultimately higher-quality restored images, underscoring the importance of addressing task-aware characteristics in multi-task learning process. However, the current approach is constrained by its focus on intra-task distributions, potentially overlooking valuable inter-task relationships that could further enhance performance. Typically, even distinct degradation types often share underlying patterns and interactions that, if properly leveraged, can improve the overall restoration process. Future research should therefore explore methods to capture these cross-task connections, integrating both task-specific and shared representations.

Conclusion

In this paper, we propose a novel approach for all-in-one image restoration named Task Uncertainty Regularization (TUR).By adaptively balancing multiple tasks during end-to-end training, TUR consistently achieves superior results across diverse restoration scenarios. Comprehensive experiments across various settings demonstrate significant improvements, underscoring both the importance of task-aware balancing in all-in-one image restoration and the efficacy of our proposed TUR approach. While our current formulation focuses on intra-task distribution, future work could explore inter-task relationships and more sophisticated regularization techniques to further enhance restoration tasks. We hope this work provides valuable insights and fosters continued innovation in all-in-one image restoration.

Acknowledgements

The research was supported by the National Natural Science Foundation of China (62471158, U23B2009), and partially supported by the Open Research Fund from Guangdong Laboratory of Artificial Intelligence and Digital Economy (SZ), under Grant No.GML-KF-24-09.

References

Abdar, M.; Pourpanah, F.; Hussain, S.; Rezazadegan, D.; Liu, L.; Ghavamzadeh, M.; Fieguth, P.; Cao, X.; Khosravi, A.; Acharya, U. R.; Makarenkov, V.; and Nahavandi, S. 2021. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information Fusion*, 76: 243–297.

Ai, Y.; Huang, H.; Zhou, X.; Wang, J.; and He, R. 2024. Multimodal Prompt Perceiver: Empower Adaptiveness, Generalizability and Fidelity for All-in-One Image Restoration. In *CVPR*, 25432–25444.

Arbeláez, P.; Maire, M.; Fowlkes, C.; and Malik, J. 2011. Contour Detection and Hierarchical Image Segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 33(5): 898–916.

Banham, M.; and Katsaggelos, A. 1997. Digital image restoration. *IEEE Signal Processing Magazine*, 14(2): 24–41.

Cao, J.; Meng, D.; and Cao, X. 2024. Chain-of-Restoration: Multi-Task Image Restoration Models are Zero-Shot Step-by-Step Universal Image Restorers. *CoRR*, abs/2410.08688.

Chen, L.; Chu, X.; Zhang, X.; and Sun, J. 2022a. Simple baselines for image restoration. In *ECCV*, volume 13667, 17–33.

Chen, S.; Ye, T.; Bai, J.; Chen, E.; Shi, J.; and Zhu, L. 2023. Sparse sampling transformer with uncertainty-driven ranking for unified removal of raindrops and rain streaks. In *ICCV*, 13106–13117.

Chen, S.; Ye, T.; Zhang, K.; Xing, Z.; Lin, Y.; and Zhu, L. 2024. Teaching Tailored to Talent: Adverse Weather Restoration via Prompt Pool and Depth-Anything Constraint. In *ECCV*, volume 15067, 95–115.

Chen, W.-T.; Huang, Z.-K.; Tsai, C.-C.; Yang, H.-H.; Ding, J.-J.; and Kuo, S.-Y. 2022b. Learning Multiple Adverse Weather Removal via Two-stage Knowledge Learning and Multi-contrastive Regularization: Toward a Unified Model. In *CVPR*, 17653–17662.

Conde, M. V.; Geigle, G.; and Timofte, R. 2024. InstructIR: High-Quality Image Restoration Following Human Instructions. In *ECCV*, volume 15094, 1–21.

Cui, Y.; Ren, W.; Cao, X.; and Knoll, A. 2024a. Image Restoration via Frequency Selection. *IEEE Trans. Pattern Anal. Mach. Intell.*, 46(2): 1093–1108.

Cui, Y.; Ren, W.; Yang, S.; Cao, X.; and Knoll, A. 2023. IRNeXt: Rethinking convolutional network design for image restoration. In *ICML*, 6545–6564.

Cui, Y.; Tao, Y.; Bing, Z.; Ren, W.; Gao, X.; Cao, X.; Huang, K.; and Knoll, A. 2022. Selective Frequency Network for Image Restoration. In *ICLR*.

Cui, Y.; Zamir, S. W.; Khan, S. H.; Knoll, A.; Shah, M.; and Khan, F. S. 2024b. AdaIR: Adaptive All-in-One Image Restoration via Frequency Mining and Modulation. *CoRR*, abs/2403.14614.

Dong, Y.; Liu, Y.; Zhang, H.; Chen, S.; and Qiao, Y. 2020. FD-GAN: Generative adversarial networks with fusion-discriminator for single image dehazing. In *AAAI*.

Fan, Q.; Chen, D.; Yuan, L.; Hua, G.; Yu, N.; and Chen, B. 2019. A General Decoupled Learning Framework for Parameterized Image Operators. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.

Fang, Z.; Dong, W.; Li, X.; Wu, J.; Li, L.; and Shi, G. 2022. Uncertainty Learning in Kernel Estimation for Multi-stage Blind Image Super-Resolution. In *ECCV*, volume 13678, 144–161.

Gal, Y.; and Ghahramani, Z. 2016. Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning. In *ICML*, 1050–1059.

Gao, H.; Tao, X.; Shen, X.; and Jia, J. 2019. Dynamic scene deblurring with parameter selective sharing and nested skip connections. In *CVPR*, 3848–3856.

Guo, Y.; Gao, Y.; Lu, Y.; Liu, R. W.; and He, S. 2024. OneRestore: A Universal Restoration Framework for Composite Degradation. In *ECCV*, 255–272.

He, K.; Fan, H.; Wu, Y.; Xie, S.; and Girshick, R. B. 2020. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 9729–9738.

Hong, M.; Liu, J.; Li, C.; and Qu, Y. 2022. Uncertainty-driven dehazing network. In *AAAI*, 906–913.

Huang, H.; Luo, M.; and He, R. 2023. Memory Uncertainty Learning for Real-World Single Image Deraining. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(3): 3446–3460.

Jiang, J.; Zuo, Z.; Wu, G.; Jiang, K.; and Liu, X. 2024a. A Survey on All-in-One Image Restoration: Taxonomy, Evaluation and Future Trends. *CoRR*, abs/2410.15067.

Jiang, K.; Wang, Z.; Yi, P.; Chen, C.; Huang, B.; Luo, Y.; Ma, J.; and Jiang, J. 2020. Multi-scale progressive fusion network for single image deraining. In *CVPR*, 8346–8355.

Jiang, K.; Wang, Z.; Yi, P.; Chen, C.; Wang, Z.; Wang, X.; Jiang, J.; and Lin, C.-W. 2021. Rain-free and residue hand-in-hand: a progressive coupled network for real-time image deraining. *IEEE Transactions on Image Processing*, 30: 7404–7418.

Jiang, Y.; Zhang, Z.; Xue, T.; and Gu, J. 2024b. AutoDIR: Automatic All-in-One Image Restoration with Latent Diffusion. In *ECCV*, volume 15098, 340–359.

Kendall, A.; and Gal, Y. 2017. What Uncertainties Do We Need in Bayesian Deep Learning for Computer Vision? In *NeurIPS*, 5574–5584.

Kong, X.; Dong, C.; and Zhang, L. 2024. Towards Effective Multiple-in-One Image Restoration: A Sequential and Prompt Learning Strategy. *CoRR*, abs/2401.03379.

Lakshminarayanan, B.; Pritzel, A.; and Blundell, C. 2017. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. In *NeurIPS*, 6402–6413.

Ledig, C.; Theis, L.; Huszar, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A. P.; Tejani, A.; Totz, J.; Wang, Z.; and Shi, W. 2017. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*, 105–114.

Li, B.; Liu, X.; Hu, P.; Wu, Z.; Lv, J.; and Peng, X. 2022. All-In-One Image Restoration for Unknown Corruption. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition (CVPR)*, 17431–17441.

Li, B.; Ren, W.; Fu, D.; Tao, D.; Feng, D.; Zeng, W.; and Wang, Z. 2019. Benchmarking Single-Image Dehazing and Beyond. *IEEE Trans. Image Process.*, 28(1): 492–505.

Li, H.; Chen, X.; Dong, J.; Tang, J.; and Pan, J. 2024. FoundIR: Unleashing Million-scale Training Data to Advance Foundation Models for Image Restoration. *CoRR*, abs/2412.01427.

Li, R.; Cheong, L.; and Tan, R. T. 2019. Heavy Rain Image Restoration: Integrating Physics Model and Conditional Adversarial Learning. In *CVPR*, 1633–1642.

Li, R.; Tan, R. T.; and Cheong, L. 2020. All in One Bad Weather Removal Using Architectural Search. In *CVPR*, 3172–3182.

Liang, J.; Cao, J.; Sun, G.; Zhang, K.; Gool, L. V.; and Timofte, R. 2021. SwinIR: Image restoration using swin transformer. In *ICCVW*, 1833–1844.

Liu, L.; Xie, L.; Zhang, X.; Yuan, S.; Chen, X.; Zhou, W.; Li, H.; and Tian, Q. 2022. TAPE: Task-Agnostic Prior Embedding for Image Restoration. In *ECCV*, volume 13678, 447–464.

Liu, Y.; Jaw, D.; Huang, S.; and Hwang, J. 2018. DesnowNet: Context-Aware Deep Network for Snow Removal. *IEEE Trans. Image Process.*, 27(6): 3064–3073.

Luo, Z.; Gustafsson, F. K.; Zhao, Z.; Sjölund, J.; and Schön, T. B. 2024. Controlling Vision-Language Models for Multi-Task Image Restoration. In *ICLR*.

Ma, K.; Duanmu, Z.; Wu, Q.; Wang, Z.; Yong, H.; Li, H.; and Zhang, L. 2017. Waterloo Exploration Database: New Challenges for Image Quality Assessment Models. *IEEE Trans. Image Process.*, 26(2): 1004–1016.

Nah, S.; Kim, T. H.; and Lee, K. M. 2017. Deep Multi-Scale Convolutional Neural Network for Dynamic Scene Deblurring. In *CVPR*, 257–265.

Ning, Q.; Dong, W.; Li, X.; Wu, J.; and Shi, G. 2021. Uncertainty-Driven Loss for Single Image Super-Resolution. In *NeurIPS*, 16398–16409.

Ning, Q.; Tang, J.; Wu, F.; Dong, W.; Li, X.; and Shi, G. 2022. Learning Degradation Uncertainty for Unsupervised Real-world Image Super-resolution. In *IJCAI*, 1261–1267.

Özdenizci, O.; and Legenstein, R. 2023. Restoring Vision in Adverse Weather Conditions With Patch-Based Denoising Diffusion Models. *IEEE Trans. Pattern Anal. Mach. Intell.*, 45(8): 10346–10357.

Potlapalli, V.; Zamir, S. W.; Khan, S. H.; and Khan, F. S. 2023. PromptIR: Prompting for All-in-One Image Restoration. In *NeurIPS*.

Qian, R.; Tan, R. T.; Yang, W.; Su, J.; and Liu, J. 2018. Attentive Generative Adversarial Network for Raindrop Removal From a Single Image. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2482–2491. Computer Vision Foundation / IEEE Computer Society.

Qin, X.; Wang, Z.; Bai, Y.; Xie, X.; and Jia, H. 2020. FFA-Net: Feature fusion attention network for single image dehazing. In *AAAI*, volume 34, 11908–11915.

Su, J.; Xu, B.; and Yin, H. 2022. A survey of deep learning approaches to image restoration. *Neurocomput.*, 487(C): 46–65.

Sun, S.; Ren, W.; Gao, X.; Wang, R.; and Cao, X. 2024. Restoring Images in Adverse Weather Conditions via Histogram Transformer. In *ECCV*, volume 15080, 111–129.

Tian, C.; Xu, Y.; and Zuo, W. 2020. Image denoising using deep CNN with batch renormalization. *Neural Networks*.

Valanarasu, J. M. J.; Yasarla, R.; and Patel, V. M. 2022. TransWeather: Transformer-based restoration of images degraded by adverse weather conditions. In *CVPR*, 2343–2353.

Vandenhende, S.; Georgoulis, S.; Van Gansbeke, W.; Proesmans, M.; Dai, D.; and Van Gool, L. 2022. Multi-Task Learning for Dense Prediction Tasks: A Survey. *IEEE Trans. Pattern Anal. Mach. Intell.*, 44(7): 3614–3633.

Wang, Z.; Cun, X.; Bao, J.; Zhou, W.; Liu, J.; and Li, H. 2022. Uformer: A General U-Shaped Transformer for Image Restoration. In *CVPR*, 17662–17672.

Wei, C.; Wang, W.; Yang, W.; and Liu, J. 2018. Deep Retinex Decomposition for Low-Light Enhancement. In *BMVC*, 155.

Wu, G.; Jiang, J.; Jiang, K.; and Liu, X. 2024. Harmony in Diversity: Improving All-in-One Image Restoration via Multi-Task Collaboration. In *ACM Multimedia*, 6015–6023.

Yang, W.; Tan, R. T.; Feng, J.; Liu, J.; Guo, Z.; and Yan, S. 2017. Deep Joint Rain Detection and Removal from a Single Image. In *CVPR*, 1685–1694.

Yang, Z.; Dong, W.; Li, X.; Wu, J.; Li, L.; and Shi, G. 2022. Self-feature Distillation with Uncertainty Modeling for Degraded Image Recognition. In *ECCV*, volume 13684, 552–569.

Zamir, S. W.; Arora, A.; Khan, S.; Hayat, M.; Khan, F. S.; and Yang, M.-H. 2022. Restormer: Efficient transformer for high-resolution image restoration. In *CVPR*, 5718–5729.

Zamir, S. W.; Arora, A.; Khan, S. H.; Hayat, M.; Khan, F. S.; Yang, M.; and Shao, L. 2021. Multi-Stage Progressive Image Restoration. In *CVPR*, 14821–14831.

Zhang, J.; Huang, J.; Yao, M.; Yang, Z.; Yu, H.; Zhou, M.; and Zhao, F. 2023a. Ingredient-Oriented Multi-Degradation Learning for Image Restoration. In *CVPR*, 5825–5835.

Zhang, K.; Li, R.; Yu, Y.; Luo, W.; and Li, C. 2021. Deep Dense Multi-Scale Network for Snow Removal Using Semantic and Depth Priors. *IEEE Trans. Image Process.*, 30: 7419–7431.

Zhang, K.; Li, Y.; Liang, J.; Cao, J.; Zhang, Y.; Tang, H.; Fan, D.; Timofte, R.; and Gool, L. V. 2023b. Practical Blind Image Denoising via Swin-Conv-UNet and Data Synthesis. *Mach. Intell. Res.*, 20(6): 822–836.

Zheng, D.; Wu, X.; Yang, S.; Zhang, J.; Hu, J.; and Zheng, W. 2024. Selective Hourglass Mapping for Universal Image Restoration Based on Diffusion Model. In *CVPR*, 25445–25455.

Zhu, Y.; Wang, T.; Fu, X.; Yang, X.; Guo, X.; Dai, J.; Qiao, Y.; and Hu, X. 2023. Learning Weather-General and Weather-Specific Features for Image Restoration Under Multiple Adverse Weather Conditions. In *CVPR*, 21747–21758.