

Learning Real-World Image De-weathering with Imperfect Supervision

Xiaohui Liu¹, Zhilu Zhang¹, Xiaohe Wu^{1*}, Chaoyu Feng², Xiaotao Wang², Lei Lei², Wangmeng Zuo¹

¹Harbin Institute of Technology

²Independent Researcher

lxh720199@gmail.com, cszlzhang@outlook.com, xhwu.cpsl.hit@gmail.com, wmzuo@hit.edu.cn

Abstract

Real-world image de-weathering aims at removing various undesirable weather-related artifacts. Owing to the impossibility of capturing image pairs concurrently, existing real-world de-weathering datasets often exhibit inconsistent illumination, position, and textures between the ground-truth images and the input degraded images, resulting in imperfect supervision. Such non-ideal supervision negatively affects the training process of learning-based de-weathering methods. In this work, we attempt to address the problem with a unified solution for various inconsistencies. Specifically, inspired by information bottleneck theory, we first develop a Consistent Label Constructor (CLC) to generate a pseudo-label as consistent as possible with the input degraded image while removing most weather-related degradation. In particular, multiple adjacent frames of the current input are also fed into CLC to enhance the pseudo-label. Then we combine the original imperfect labels and pseudo-labels to jointly supervise the de-weathering model by the proposed Information Allocation Strategy (IAS). During testing, only the de-weathering model is used for inference. Experiments on two real-world de-weathering datasets show that our method helps existing de-weathering models achieve better performance. Code is available at <https://github.com/1180300419/imperfect-deweathering>.

Introduction

The effectiveness of outdoor computer vision systems *e.g.* monitoring systems and autonomous driving is inevitably impacted by weather conditions. Image de-weathering, which aims at removing undesirable image degradation caused by adverse weather, *e.g.* rain, fog, and snow, is significant in improving the performance of the subsequent tasks. Recently, data-driven learning-based de-weathering methods have experienced notable progress. However, capturing ideal real-world pairs for training neural networks is nearly impossible since it is not feasible to capture a scene simultaneously with and without weather-related artifacts.

To circumvent this problem, some existing de-weathering methods (Valanarasu, Yasarla, and Patel 2022; Wang, Ma, and Liu 2023) initially construct paired data by simulating weather degradation and subsequently undergo super-

vised learning. However, synthesized degraded images have limited realism in modeling complex and variable weather characteristics (*e.g.* raindrop shapes and fog concentration). Consequently, models trained using synthetic pairs typically face challenges in effectively generalizing to realistic severe weather scenarios, even leading to artifacts and the loss of original details. Other methods (Wu et al. 2021; Wei et al. 2021) address the lack of real-world pairs by employing semi-supervised or unsupervised manners. While some of these methods (Rai et al. 2022; Huang, Yu, and He 2021) have demonstrated promising results, challenges remain in separating intricate overlaps of foreground and background information in real-world degraded images (Li et al. 2022).

Recognizing the aforementioned limitations, the recent works (Ba et al. 2022; Zhang et al. 2023) pin their hopes on acquiring real-world image pairs. They relax the strict requirement for simultaneous acquisition of ideal pairs and instead retrieve degraded and clean images from live streams of landscape scenes on YouTube within the shortest possible time window. And the conclusion shows that such a manner indeed bridges the domain gap better than synthetic (Hu et al. 2019; Li, Cheong, and Tan 2019) and semi-real (Wang et al. 2019) datasets. It is worth noting that, despite the adoption of strict filtering conditions during dataset collection to ensure spatial position and illumination consistency between the ground-truth (GT) and degraded images, certain inconsistencies persist, which will have a detrimental impact on the learning process of the de-weathering model.

Actually, the inconsistencies also occur in some other low-level vision (*e.g.* super-resolution (Zhang et al. 2022; Wang et al. 2023), deblurring (Li et al. 2023), learnable ISP (Zhang et al. 2021b), and etc (Feng et al. 2023)) tasks. And their works usually utilize optical flow estimation and color correction approaches (He, Sun, and Tang 2012) for aligning the network output and GT in position and color, respectively. However, these solutions are less effective when being applied to de-weathering due to the complexity and diversity of inconsistent types (see Fig. 1).

In this work, we suggest a novel perspective to alleviate the adverse impact of imperfect supervision in real-world de-weathering datasets, and propose a unified solution for a variety of inconsistencies. The main motivation is from the information bottleneck (IB) theory (Tishby and Zaslavsky 2015). Therein, during network training, it has been

*Corresponding Author

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.




Figure 1: Examples of various inconsistent types between the degraded (top) and original GT (bottom) images. These cases come from two real-world de-weathering datasets (*i.e.*, GT-Rain-Snow and WeatherStream). The pairs in (a)~(c) show inconsistencies in color and illumination. (d) and (e) show that the pairs are not aligned owing to object movement. (f) shows inconsistent textures due because of snow melting. The inconsistent parts are highlighted in red boxes.

observed that the mutual information between the learned features and target increases monotonically, while that between the learned features and input increases first and then decreases gradually. Taking advantage of it, we can develop a **Consistent Label Constructor (CLC)** for generating a pseudo-label whose non-weather content is as consistent as possible with the input degraded image while degradation are mostly removed. Considering the characteristics of the dataset itself, multiple adjacent frames of the current input can be also fed into CLC to enhance the ability of the pseudo-label. Noted that although the pseudo label is better than the original imperfect label in terms of consistency preservation, it may not be as satisfactory as the original label in terms of degradation removal. Therefore, we combine the two labels and present an **Information Allocation Strategy (IAS)** for supervising the de-weathering model, which draws on the advantages and discards the disadvantages of both. During inference, only the de-weathering model is used for removing weather-related degradation.

Experiments are conducted with RainRobust (Ba et al. 2022) and Restormer (Zamir et al. 2022) models on GT-Rain-Snow (Ba et al. 2022) and WeatherStream (Zhang et al. 2023) datasets. With the proposed method, the quantitative and qualitative results of existing de-weathering models are greatly improved, while not increasing any inference cost.

The main contributions are as follows:

- In real-world de-weathering datasets, the content inconsistencies between degraded and labeled images hurt the performance of de-weathering models, and we propose a unified and novel solution for various inconsistencies.
- We first feed multiple frames into Consistent Label Constructor (CLC) to generate a pseudo-label consistent with input content, and then combine the pseudo- and original labels for supervising the de-weathering model by the Information Allocation Strategy (IAS).
- Experiments demonstrate that our proposed method en-

hances the performance of existing de-weathering models without incurring additional inference costs.

Related Work

De-Weathering Methods

Deraining. Traditional deraining methods (Garg and Nayar 2007; Chen and Hsu 2013) rely on handcrafted image priors to remove rain streaks. but these priors heavily depend on empirical observations and fail to capture the intrinsic properties of clean images accurately. In contrast, recent advancements in deep learning-based approaches, utilizing neural convolutional networks (CNNs) (Fu et al. 2017) or Transformers (Chen et al. 2023b), have achieved remarkable rain removal performance by learning prior knowledge from paired data. Nonetheless, these methods typically rely on synthetic data and show limited generalization when applied to real-world images. To address this limitation, several studies have explored semi-supervised (Yasarla, Sindagi, and Patel 2021) or unsupervised (Ye et al. 2022) frameworks for handling real rainy images, given the scarcity of real-world deraining datasets. However, these approaches often face challenges in achieving satisfactory deraining results.

Desnowing. Early works usually employ priors derived from rainfall-driven features to model snow particles. Desnownet (Liu et al. 2018) is the first CNN-based approach introduced for snow removal. Recently, DDMSNet (Zhang et al. 2021a) further presents a deep dense multiscale network incorporated semantic and geometric priors to enhance snow removal. To address the challenges posed by diverse scenes, a multi-scale projection transformer (Chen et al. 2023a) is proposed, which employs a multi-path approach to involve various snow degradation features and utilizes self-attention operation to integrate comprehensive scene context information for clean reconstruction.

Dehazing. Previous attempts (Li et al. 2019) primarily rely on the parameter estimation of the atmospheric scatter-

ing model by the handcrafted priors on haze-free images. With the advent of deep learning techniques, learning-based methods have been developed rapidly by exploring the deployment of CNNs and Transformers effectively (Guo et al. 2022). Despite performing excellently well on synthetic data, the above approaches typically suffer from constrained generalization in handling real-world hazy images (Ancuti et al. 2021). As an alternative, some studies endeavor to employ GANs (Goodfellow et al. 2014) within the unpaired or unsupervised framework (Chen et al. 2022b) for real image dehazing. However, GANs often generate images with artifacts, which adversely affects the subsequent model training. **Unified De-Weathering.** To remove various adverse weather conditions via a single set of pre-trained parameters, All-in-One (Li, Tan, and Cheong 2020) deploys multiple encoders associated with different weather types for image restoration. However, it suffers from the model inefficiency problem. To overcome this, TransWeather (Valanarasu, Yasarla, and Patel 2022) presents a transformer-based structure comprising a single encoder-decoder network with learnable weather type queries. In contrast, (Chen et al. 2022a) draw inspiration from knowledge distillation, and leverage multiple well-trained teacher networks to distill the student model for multiple weather effects. Meanwhile, (Wang, Ma, and Liu 2023) introduce a multi-task learning way for deraining and desnowing by designing multiple heads. Although the aforementioned works achieve encouraging results in various weather types with a unified framework, they are typically trained on synthetic data, leading to degraded performance in real-world all-weather removal.

De-Weathering Datasets

Single weather removal, such as deraining (Fu et al. 2017; Zhang and Patel 2018; Li, Cheong, and Tan 2019; Hu et al. 2019), dehazing (Li et al. 2018; Ancuti et al. 2019, 2021) and desnowing (Liu et al. 2018; Chen et al. 2020, 2021) has made substantial advancements due to the numerous publicly accessible datasets. As for the de-weathering task, due to the absence of dedicated datasets, prevailing methods often explore the integration of existing weather-specific datasets. For example, some works (Valanarasu, Yasarla, and Patel 2022; Li, Tan, and Cheong 2020) sample images from Snow100K (Liu et al. 2018), Raindrop (Qian et al. 2018) and Outdoor-Rain (Li, Cheong, and Tan 2019) to facilitate the training of the model. However, those weather-specific datasets are commonly composed of synthetic paired data given certain degradation.

Recently, SPA (Wang et al. 2019) dataset introduces a semi-automatic approach to generate one high-quality clean image from each input sequence of real rain images, and builds the semi-real paired dataset. RainRobust (Ba et al. 2022) indeed establishes the real-world paired deraining dataset, *i.e.*, GT-Rain, by enforcing rigorous selection criteria to minimize the environmental variations in time multiplexed pairs. Unfortunately, it requires significant human and financial resources, which limits the size of the dataset. To reduce manpower requirements, WeatherStream (Zhang et al. 2023) then introduce an automatic pipeline capturing real-world weather effects and their clean

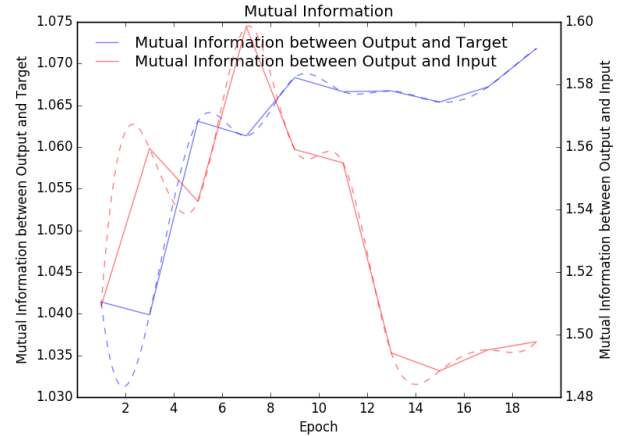


Figure 2: Comparison of mutual information in different training periods. At different training epochs, we test the RainRobust model trained on the GT-Rain-Snow dataset, and measure the mutual information between the output and the target as well as between the output and the input. The model was trained in total 20 epochs, and we calculate the mutual information every 2 epochs.

image pairs, which can be typically regarded as a real-world de-weathering dataset.

Proposed Method

Motivation and Overall Pipeline

Denote by \mathbf{I} and \mathbf{G} a degraded image caused by severe weather (*e.g.*, rain, fog, and snow) and the corresponding clean ground-truth (GT) image. The learning-based de-weathering model \mathcal{D} predicts output $\hat{\mathbf{G}}$ without weather-related degradation from \mathbf{I} , *i.e.*,

$$\hat{\mathbf{G}} = \mathcal{D}(\mathbf{I}; \Theta_{\mathcal{D}}). \quad (1)$$

$\Theta_{\mathcal{D}}$ denotes the model parameters, whose optimization objective can be defined by,

$$\Theta_{\mathcal{D}}^* = \arg \min_{\Theta_{\mathcal{D}}} \mathbb{E}_{\mathbf{I}, \mathbf{G}} [\mathcal{L}_{\mathcal{D}}(\hat{\mathbf{G}}, \mathbf{G})], \quad (2)$$

where $\mathcal{L}_{\mathcal{D}}$ denotes the loss functions.

Generally, the content of GT image \mathbf{G} should remain the same as that of input \mathbf{I} , except for the weather-related degradation to be removed. But for real-world de-weathering, \mathbf{G} and \mathbf{I} cannot be captured at the same time, and they have to be collected successively. During this collection process, changes in weather and scenes are unavoidable and uncontrollable. As a result, there are inconsistencies in illumination, spatial position, and texture between \mathbf{G} and \mathbf{I} . Although recent works (Ba et al. 2022; Zhang et al. 2023) have set strict filtering and selection conditions to alleviate this problem, some subtle inconsistencies still exist, as shown in Fig. 1. Intuitively, these irregular and diverse inconsistencies will inevitably affect the optimization of parameters $\Theta_{\mathcal{D}}$, harming the performance of the de-weathering model.

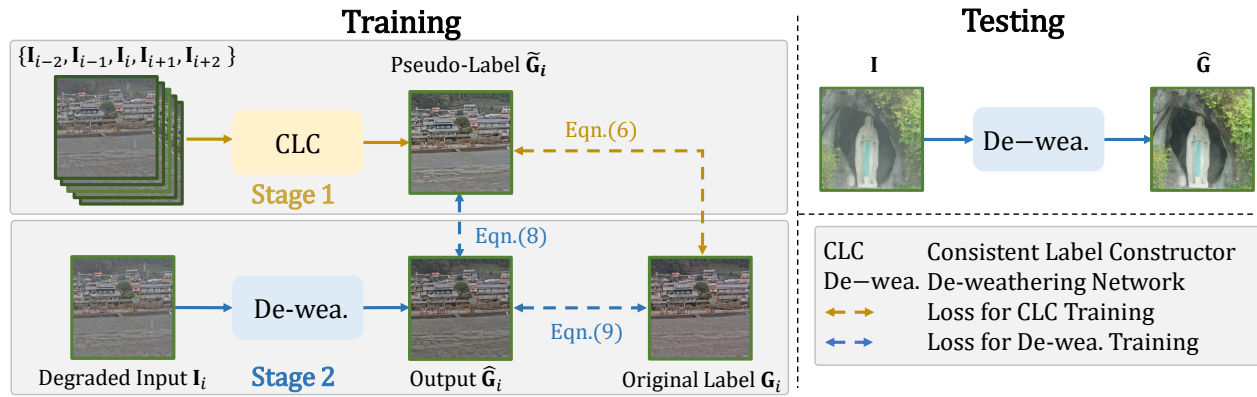


Figure 3: Overview of the proposed pipeline. Left: In training stage 1, we pre-train a Consistent Label Constructor (CLC), which utilizes multiple degraded frames to produce a pseudo-label $\tilde{\mathbf{G}}_i$ consistent with the input degraded image \mathbf{I}_i . In training stage 2, we leverage both the original label \mathbf{G}_i and pseudo-label $\tilde{\mathbf{G}}_i$ to constrain the output $\hat{\mathbf{G}}_i$ of the de-weathering model. Right: During testing, only the de-weathering model is used to remove weather-related degradation. Better viewed in color.

It is worth mentioning some other real-world low-level vision (*e.g.* super-resolution (Zhang et al. 2022; Wang et al. 2023), deblurring (Li et al. 2023), learnable ISP (Zhang et al. 2021b), and etc (Feng et al. 2023)) tasks face this problem as well, and the corresponding solutions vary. For inconsistent illumination, guided filter (He, Sun, and Tang 2012) is adopted to correct output color to target one, and then calculate the loss value (Wei et al. 2020). For spatial misalignment, optical flow estimation (Sun et al. 2018) is widely used for aligning pairs, and some works propose more robust loss functions, *e.g.*, CoBi (Kim et al. 2013) and misalignment-tolerate ℓ_1 (Xia et al. 2023) loss. Regrettably, these single and coarse-grained inconsistency mitigation methods are less effective for real-world de-weathering pairs whose inconsistencies are diverse and fine-grained, as evidenced by the experimental results presented in the subsequent section.

In this work, we aim to explore a unified approach for alleviating the adverse impact caused by a variety of inconsistencies in real-world de-weathering datasets. Inspired by information bottleneck (IB) theory (Tishby and Zaslavsky 2015), we propose a novel perspective for addressing the problem. Specifically, the IB principle has observed that the mutual information between the network features and target goes up monotonically, while that between the network features and input goes up first and then goes down during the training, as shown in Fig. 2. In other words, for the target information, the network continuously learns. For the input information, the network absorbs first and then eliminates the part that is irrelevant to the target. Thus, it gives us an opportunity to construct a pseudo-label as consistent as possible with the input. Simultaneously, for de-weathering with non-ideal supervision, it may be more difficult to fit weather-irrelevant irregular disturbance from the target than to remove weather-related degradation. So the pseudo-label is promising to remove most of the degradation while not fitting the imperfect parts of the original GT.

Furthermore, we develop a Consistent Label Constructor (CLC) for generating pseudo-labels. Multiple neighboring

frames of the current degraded image are also fed into CLC to enhance the consistency between the pseudo-label and the degraded image. Additionally, noted that it is not feasible when only taking the pseudo-label as supervision for the de-weathering model, as the pseudo-label has to make a trade-off between consistency preservation and degradation removal. Instead, we combine the pseudo-label and original imperfect labels to re-train a de-weathering model by the proposed Information Allocation Strategy (IAS), which is detailed in the following section. Please see Fig. 3 for the overview of the proposed pipeline.

Consistent Label Constructor

The CLC model \mathcal{C} learns pseudo-label $\tilde{\mathbf{G}}_i$ from input i -th degraded image \mathbf{I}_i , *i.e.*,

$$\tilde{\mathbf{G}}_i = \mathcal{C}(\mathbf{I}_i; \Theta_{\mathcal{C}}). \quad (3)$$

The parameters $\Theta_{\mathcal{C}}$ can be optimized by the constraints between the original label \mathbf{G}_i and $\tilde{\mathbf{G}}_i$, which is written as,

$$\Theta_{\mathcal{C}}^* = \arg \min_{\Theta_{\mathcal{C}}} \mathbb{E}_{\mathbf{I}_i, \mathbf{G}_i} [\mathcal{L}_{\mathcal{C}}(\tilde{\mathbf{G}}_i, \mathbf{G}_i)], \quad (4)$$

where $\mathcal{L}_{\mathcal{C}}$ denotes the loss functions of CLC model.

Moreover, given that real-world de-weathering datasets are usually collected from live streams, multiple adjacent frames of degraded images can be easily obtained. Thus, we can naturally feed the neighboring $2n$ frames into the CLC model. Eqn. (3) can be modified as,

$$\tilde{\mathbf{G}}_i = \mathcal{C}(\mathbf{I}_{i-n}, \dots, \mathbf{I}_{i-1}, \mathbf{I}_i, \mathbf{I}_{i+1}, \dots, \mathbf{I}_{i+n}; \Theta_{\mathcal{C}}). \quad (5)$$

On the one hand, it does not hinder CLC to produce a pseudo-label consistent with \mathbf{I}_i , as the weather-independent content is also consistent between frames captured in a short amount of time. On the other hand, sometimes complementary multi-frames help CLC to generate a clearer pseudo-label, as the rain streaks and snowflakes cannot always be fixed at the same position in the frame.

For the loss function in Eqn. (4), we adopt ℓ_1 as well as structural similarity (SSIM) (Wang, Simoncelli, and Bovik 2003) distance between the output $\tilde{\mathbf{G}}_i$ and the original label \mathbf{G}_i , which can be written as,

$$\mathcal{L}_C(\tilde{\mathbf{G}}_i, \mathbf{G}_i) = \|\tilde{\mathbf{G}}_i - \mathbf{G}_i\|_1 + (1 - \mathcal{L}_{SSIM}(\tilde{\mathbf{G}}_i, \mathbf{G}_i)), \quad (6)$$

where \mathcal{L}_{SSIM} denotes the multi-scale SSIM function.

The network structure of CLC depends on that of the de-weathering model used in our experiments. For example, when the de-weathering adopts Restormer (Zhang et al. 2023), we simply modified its single encoder to multiple ones as the CLC model. The details about the CLC architecture can be seen in the suppl. In addition, we set n to 2 for experiments.

Information Allocation Strategy

Although the pseudo-label $\tilde{\mathbf{G}}_i$ is more consistent with the input \mathbf{I}_i than the original label \mathbf{G}_i , it may perform worse on degradation removal. Thus, It is not appropriate to treat $\tilde{\mathbf{G}}_i$ as the only supervision of the de-weathering model. Naturally, we can also include \mathbf{G}_i as additional guidance, and combine the advantages of $\tilde{\mathbf{G}}_i$ and \mathbf{G}_i for supervising the de-weathering model. Thus, the loss function in Eqn. (2) can be written as,

$$\mathcal{L}_D(\hat{\mathbf{G}}_i, \tilde{\mathbf{G}}_i, \mathbf{G}_i) = \mathcal{L}_{pse}(\hat{\mathbf{G}}_i, \tilde{\mathbf{G}}_i) + \lambda_{ori} \mathcal{L}_{ori}(\hat{\mathbf{G}}_i, \mathbf{G}_i), \quad (7)$$

where \mathcal{L}_{pse} and \mathcal{L}_{ori} denote the optimization objectives with pseudo-label and the original label as supervision, respectively. λ_{ori} denotes a weight coefficient. The proper design of \mathcal{L}_{pse} and \mathcal{L}_{ori} is essential, and we propose the following Information Allocation Strategy (IAS) for this.

First, IAS should take advantage of the fact that the pseudo-label is more consistent with the input. Thus, for \mathcal{L}_{pse} , we directly adopt the image-level constraint between the output $\hat{\mathbf{G}}_i$ and the pseudo label $\tilde{\mathbf{G}}_i$, *i.e.*,

$$\mathcal{L}_{pse}(\hat{\mathbf{G}}_i, \tilde{\mathbf{G}}_i) = \|\hat{\mathbf{G}}_i - \tilde{\mathbf{G}}_i\|_1. \quad (8)$$

Second, IAS should efficiently leverage clean information from the original label \mathbf{G}_i while mitigating the impact of its inconsistencies as much as possible. At this time, image-level or pixel-level loss terms are not suitable. Instead, we impose feature-level and distribution-level constraints because of their robustness to spatial misalignment and color variations between image pairs. On the one hand, we adopt rain-robust loss (Ba et al. 2022), which is a variant of contrastive loss and is calculated between intermediate features of the de-weathering networks. On the other hand, we utilize Sliced Wasserstein (SW) (Zhang et al. 2022) loss, which is a probability distribution distance between VGG (Simonyan and Zisserman 2014) features. In short, the optimization objective with the original label can be written as,

$$\mathcal{L}_{ori} = \mathcal{L}_{Robust} + \lambda_{SW} \mathcal{L}_{SW}, \quad (9)$$

where \mathcal{L}_{Robust} and \mathcal{L}_{SW} denote rain-robust and SW loss terms, respectively. λ_{SW} is the weight for SW loss and set to 0.08. λ_{ori} in Eqn. (7) is set to 0.1. The details and formulas of \mathcal{L}_{Robust} and \mathcal{L}_{SW} are shown in the suppl.

Experiments

Datasets

(Zhang et al. 2023) offers datasets constructed in two ways. One of which, named GT-Rain-Snow, follows the protocol of (Ba et al. 2022) and extends to multiple weather types by manual collection. The other WeatherStream dataset is obtained by automatic collection. Specifically, GT-Rain-Snow comprises 129 scenes and WeatherStream contains a larger set of 424 scenes. Each scene consists of approximately 300 aligned degraded images and one clean image. We conduct experiments with the two training datasets, respectively. And experiments are all evaluated on the testing dataset proposed by (Zhang et al. 2023), which consists of 45 scenes and covers the sub-tasks of deraining, dehazing and desnowing.

Implementation Details

We integrate our method into the CNN-based RainRobust (Ba et al. 2022) and transformer-based Restormer (Zamir et al. 2022) to validate the effectiveness. During training, we follow (Ba et al. 2022) to employ various augmentation techniques, including random rotation, padding, cropping, Rain-Mix (Guo et al. 2021) and Snow-Mix (Zhang et al. 2023). For optimizing the models, we adopt the Adam with $\beta_1 = 0.9$ and $\beta_2 = 0.999$. A warm-up strategy is employed to gradually increase the learning rate from 5×10^{-5} to 2×10^{-4} , followed by the cosine annealing strategy to decrease the learning rate from 2×10^{-4} to 10^{-6} . Experiments based on RainRobust are conducted on a single GPU. The corresponding patch size is 256×256 , and the batch size is 8. Experiments based on Restormer are conducted on two GPUs. The corresponding patch and batch sizes are set to 168×168 and 6, respectively. All experiments are implemented with the PyTorch framework.

Comparison with State-of-the-Arts

Quantitative Comparisons. In Table 1, we report the PSNR and SSIM of the models trained with GT-Rain-Snow and WeatherStream datasets, respectively. It can be seen that our method consistently improves the performance of baseline de-weathering models regardless of whether a CNN or Transformer backbone, thereby demonstrating its satisfactory capability. In particular, when training models with GT-Rain-Snow dataset, our transformer-based Restormer obtains 0.66dB PSNR improvements while CNN-based RainRobust achieves 0.39dB PSNR gains. Although the inconsistencies in the WeatherStream are significantly reduced in comparison with GT-Rain-Snow. When training models with WeatherStream, our Restormer and RainRobust still obtain 0.21dB and 0.24dB PSNR gains, respectively. Moreover, our method universally boosts the capability of de-weathering across different weather conditions.

Qualitative Comparisons. Fig. 4 and Fig. 5 show visual comparisons on various scenarios. Our method brings visible improvements with fewer degradation and more consistent color. For instance, in the bottom-right example of Fig. 4, our method promisingly recovers the color of heavily foggy images, while Restormer only removes some streaks.

Dataset	Method	Rain		Fog		Snow		Overall	
		PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
GT-Rain-Snow (Ba et al. 2022) (Zhang et al. 2023)	Restormer (Zamir et al. 2022)	22.63	0.7940	20.14	0.7700	21.62	0.8123	21.51	0.7913
	Ours-Restormer	23.10	0.7971	21.46	0.7860	21.78	0.8178	22.17	0.7995
	RainRobust (Ba et al. 2022)	22.83	0.7887	20.95	0.7691	22.17	0.8058	22.01	0.7871
	Ours-RainRobust	23.24	0.7980	21.52	0.7860	22.31	0.8107	22.40	0.7977
WeatherStream (Zhang et al. 2023)	Restormer (Zamir et al. 2022)	23.50	0.7990	22.89	0.8034	22.29	0.8175	22.94	0.8059
	Ours-Restormer	23.64	0.8055	23.01	0.8062	22.66	0.8209	23.15	0.8102
	RainRobust (Ba et al. 2022)	23.46	0.7957	22.61	0.7962	22.11	0.8143	22.78	0.8013
	Ours-RainRobust	23.68	0.7987	22.71	0.7971	22.54	0.8163	23.02	0.8033

Table 1: Quantitative improvements of de-weathering models when applying the proposed method. Experiments are conducted with Restormer and RainRobust networks on GT-Rain-Snow and WeatherStream datasets, respectively.




Figure 4: Qualitative testing results of the de-weathering models trained with GT-Rain-Snow dataset.




Figure 5: Qualitative testing results of the de-weathering models trained with WeatherStream dataset.

In the top-left example of Fig. 5, our de-weathered image contains fewer snow streaks than Rain-Robust.

Visualization of Pseudo-Labels

As mentioned above, our pseudo-labels can have greater consistency with the degraded images compared to the original label, encompassing color, position, and other aspects. Here we give some visual examples in Fig. 6. Therein, the first two rows showcase the superior color and texture consistency of pseudo-labels with the degraded images, respec-

tively. Notably, the utilization of multiple frames in CLC enhances the visual quality in the last row, with pseudo-labels exhibiting more intricate textures, while the original label still retains rain streaks.

Ablation Study

Comparison with Other Inconsistency- Handling Methods

Here we compare our method with alternative ones from other works that can also alleviate the effect of inconsisten-




Figure 6: Visualizations of pseudo-labels.

Methods	PSNR \uparrow	SSIM \uparrow
Baseline	22.01	0.7871
Guided Filter	22.14	0.7959
Optical Flow	22.15	0.7932
Misalign-Tolerate ℓ_1	22.18	0.7951
Guided Filter + Optical Flow	22.21	0.7935
Guided Filter + Misalign-Tolerate ℓ_1	22.24	0.7976
Ours	22.40	0.7977

Table 2: Quantitative results when comparing with other inconsistency-handling methods. The baseline represents the model trained directly with the original label. The results on three sub-tasks are attached in the suppl.

cies between input and GT images. We select two types of methods that are robust to color-inconsistent supervision and misaligned supervision, respectively. Among them, guided filter (He, Sun, and Tang 2012) is adopted for aligning the color between the output and GT (Wei et al. 2020; Wang et al. 2023). The optical flow (Sun et al. 2018) is used in some works (Zhang et al. 2021b; Li et al. 2023) for aligning the content spatially. And the misalignment-tolerate ℓ_1 (Xia et al. 2023) loss is also present to tackle the misalignment.

During experiments, we train the de-weathering model with the single or combined methods by following their ways. The results are shown in Table 2. From the table, although these methods can also improve performance compared to the baseline, our method achieves better performance than them. Furthermore, even combining two complementary techniques (e.g., Guided Filter + Optical Flow), our method still performs better than them.

Frame Number	CLC		De-Weathering	
	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
1	22.01	0.7871	22.19	0.7971
3	22.84	0.8199	22.34	0.7972
5	23.30	0.8282	22.40	0.7977
7	23.20	0.8277	22.39	0.7976

Table 3: Effect of feeding the different number of frames into the CLC model. The results of the CLC and corresponding de-weathering models are both reported.

Pseudo-Label	Original Label	PSNR \uparrow	SSIM \uparrow
×	✓	22.01	0.7871
✓	×	22.18	0.7903
✓	✓	22.40	0.7977

Table 4: Quantitative results of de-weathering model trained with different supervisions.

Effect of Number of Frames in CLC

Here we investigate the effect of feeding different number of frames into CLC. Table 3 evaluates the performance of both CLC and its enhanced de-weathering model. It can be observed that their performance generally increases as the number of frames increases. When the number of frames is larger than 5, the performance tends to be saturated. See suppl. for visual examples.

Effect of Supervision Setting

Here we conduct an ablation experiment to investigate the impact of different supervisions when training the de-weathering model. As illustrated in Table 4, benefiting from CLC’s meticulous design, taking pseudo-labels for supervising the de-weathering model is more effective than taking the original labels for that. Nevertheless, combining the two labels can achieve better performance gain. This result also indicates the complementary nature of the two supervisions, and shows the effectiveness of our proposed IAS.

Conclusion

In this work, we propose a unified method to mitigate adverse effects of inconsistencies in real-world de-weathering pairs. Firstly, we develop a CLC that generates a pseudo-label consistent with the input degraded image while eliminating most of weather-related degradation. To further enhance the pseudo-label, we incorporate multiple adjacent frames of the input into the CLC. Then we propose an IAS to combine the original imperfect labels and pseudo-labels, thereby providing comprehensive and complementary supervision for the de-weathering networks. Experimental results on two real-world de-weathering datasets demonstrate that our method significantly improves the performance of existing de-weathering models.

Acknowledgements

This work was supported in part by the National Natural Science Foundation of China (NSFC) under Grant No. U19A2073 and Heilongjiang Science and Technology Project under Grant 2022ZX01A21.

References

- Ancuti, C. O.; Ancuti, C.; Sbert, M.; and Timofte, R. 2019. Dense-haze: A benchmark for image dehazing with dense-haze and haze-free images. In *2019 IEEE international conference on image processing (ICIP)*, 1014–1018. IEEE.
- Ancuti, C. O.; Ancuti, C.; Vasluiianu, F.-A.; and Timofte, R. 2021. NTIRE 2021 nonhomogeneous dehazing challenge report. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 627–646.
- Ba, Y.; Zhang, H.; Yang, E.; Suzuki, A.; Pfahnl, A.; Chandrappa, C. C.; de Melo, C. M.; You, S.; Soatto, S.; Wong, A.; et al. 2022. Not Just Streaks: Towards Ground Truth for Single Image Deraining. In *Computer Vision–ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII*, 723–740. Springer.
- Chen, S.; Ye, T.; Liu, Y.; Liao, T.; Jiang, J.; Chen, E.; and Chen, P. 2023a. MSP-Former: Multi-Scale Projection Transformer for Single Image Desnowing. In *ICASSP 2023 - 2023 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, 1–5.
- Chen, W.-T.; Fang, H.-Y.; Ding, J.-J.; Tsai, C.-C.; and Kuo, S.-Y. 2020. JSTASR: Joint size and transparency-aware snow removal algorithm based on modified partial convolution and veiling effect removal. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XXI 16*, 754–770. Springer.
- Chen, W.-T.; Fang, H.-Y.; Hsieh, C.-L.; Tsai, C.-C.; Chen, I.; Ding, J.-J.; Kuo, S.-Y.; et al. 2021. All snow removed: Single image desnowing algorithm using hierarchical dual-tree complex wavelet representation and contradict channel loss. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4196–4205.
- Chen, W.-T.; Huang, Z.-K.; Tsai, C.-C.; Yang, H.-H.; Ding, J.-J.; and Kuo, S.-Y. 2022a. Learning multiple adverse weather removal via two-stage knowledge learning and multi-contrastive regularization: Toward a unified model. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17653–17662.
- Chen, X.; Li, H.; Li, M.; and Pan, J. 2023b. Learning A Sparse Transformer Network for Effective Image Deraining. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5896–5905.
- Chen, X.; Pan, J.; Jiang, K.; Li, Y.; Huang, Y.; Kong, C.; Dai, L.; and Fan, Z. 2022b. Unpaired deep image deraining using dual contrastive learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2017–2026.
- Chen, Y.-L.; and Hsu, C.-T. 2013. A generalized low-rank appearance model for spatio-temporally correlated rain streaks. In *Proceedings of the IEEE international conference on computer vision*, 1968–1975.
- Feng, R.; Li, C.; Chen, H.; Li, S.; Gu, J.; and Loy, C. C. 2023. Generating Aligned Pseudo-Supervision from Non-Aligned Data for Image Restoration in Under-Display Camera. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5013–5022.
- Fu, X.; Huang, J.; Zeng, D.; Huang, Y.; Ding, X.; and Paisley, J. 2017. Removing rain from single images via a deep detail network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 3855–3863.
- Garg, K.; and Nayar, S. K. 2007. Vision and rain. *International Journal of Computer Vision*, 75: 3–27.
- Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *Advances in neural information processing systems*, 27.
- Guo, C.-L.; Yan, Q.; Anwar, S.; Cong, R.; Ren, W.; and Li, C. 2022. Image dehazing transformer with transmission-aware 3d position embedding. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5812–5820.
- Guo, Q.; Sun, J.; Juefei-Xu, F.; Ma, L.; Xie, X.; Feng, W.; Liu, Y.; and Zhao, J. 2021. Efficientderain: Learning pixel-wise dilation filtering for high-efficiency single-image deraining. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1487–1495.
- He, K.; Sun, J.; and Tang, X. 2012. Guided image filtering. *IEEE transactions on pattern analysis and machine intelligence*, 35(6): 1397–1409.
- Hu, X.; Fu, C.-W.; Zhu, L.; and Heng, P.-A. 2019. Depth-attentional features for single-image rain removal. In *Proceedings of the IEEE/CVF Conference on computer vision and pattern recognition*, 8022–8031.
- Huang, H.; Yu, A.; and He, R. 2021. Memory oriented transfer learning for semi-supervised image deraining. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 7732–7741.
- Kim, J.; Zhang, H.; André, P.; Chilton, L. B.; Mackay, W.; Beaudouin-Lafon, M.; Miller, R. C.; and Dow, S. P. 2013. Cobi: A community-informed conference scheduling tool. In *Proceedings of the 26th annual ACM symposium on User interface software and technology*, 173–182.
- Li, B.; Ren, W.; Fu, D.; Tao, D.; Feng, D.; Zeng, W.; and Wang, Z. 2018. Benchmarking single-image dehazing and beyond. *IEEE Transactions on Image Processing*, 28(1): 492–505.
- Li, R.; Cheong, L.-F.; and Tan, R. T. 2019. Heavy rain image restoration: Integrating physics model and conditional adversarial learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 1633–1642.
- Li, R.; Tan, R. T.; and Cheong, L.-F. 2020. All in one bad weather removal using architectural search. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 3175–3185.

- Li, W.; Zhang, Q.; Zhang, J.; Huang, Z.; Tian, X.; and Tao, D. 2022. Toward Real-world Single Image Deraining: A New Benchmark and Beyond. *arXiv preprint arXiv:2206.05514*.
- Li, Y.; Miao, Q.; Ouyang, W.; Ma, Z.; Fang, H.; Dong, C.; and Quan, Y. 2019. LAP-Net: Level-aware progressive network for image dehazing. In *Proceedings of the IEEE/CVF international conference on computer vision*, 3276–3285.
- Li, Y.; Ren, D.; Shu, X.; and Zuo, W. 2023. Learning single image defocus deblurring with misaligned training pairs. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 1495–1503.
- Liu, Y.-F.; Jaw, D.-W.; Huang, S.-C.; and Hwang, J.-N. 2018. DesnowNet: Context-aware deep network for snow removal. *IEEE Transactions on Image Processing*, 27(6): 3064–3073.
- 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*, 2482–2491.
- Rai, S. N.; Saluja, R.; Arora, C.; Balasubramanian, V. N.; Subramanian, A.; and Jawahar, C. 2022. Fluid: Few-shot self-supervised image deraining. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, 3077–3086.
- Simonyan, K.; and Zisserman, A. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*.
- Sun, D.; Yang, X.; Liu, M.-Y.; and Kautz, J. 2018. Pw-net: Cnns for optical flow using pyramid, warping, and cost volume. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 8934–8943.
- Tishby, N.; and Zaslavsky, N. 2015. Deep learning and the information bottleneck principle. In *2015 IEEE information theory workshop (itw)*, 1–5. IEEE.
- Valanarasu, J. M. J.; Yasarla, R.; and Patel, V. M. 2022. Transweather: Transformer-based restoration of images degraded by adverse weather conditions. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2353–2363.
- Wang, R.; Liu, X.; Zhang, Z.; Wu, X.; Feng, C.-M.; Zhang, L.; and Zuo, W. 2023. Benchmark Dataset and Effective Inter-Frame Alignment for Real-World Video Super-Resolution. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1168–1177.
- Wang, T.; Yang, X.; Xu, K.; Chen, S.; Zhang, Q.; and Lau, R. W. 2019. Spatial attentive single-image deraining with a high quality real rain dataset. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12270–12279.
- Wang, Y.; Ma, C.; and Liu, J. 2023. SmartAssign: Learning a Smart Knowledge Assignment Strategy for Deraining and Desnowing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3677–3686.
- Wang, Z.; Simoncelli, E. P.; and Bovik, A. C. 2003. Multi-scale structural similarity for image quality assessment. In *The Thrity-Seventh Asilomar Conference on Signals, Systems & Computers, 2003*, volume 2, 1398–1402. Ieee.
- Wei, Y.; Liu, M.; Wang, H.; Zhu, R.; Hu, G.; and Zuo, W. 2020. Learning flow-based feature warping for face frontalization with illumination inconsistent supervision. In *Computer Vision—ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XII 16*, 558–574. Springer.
- Wei, Y.; Zhang, Z.; Wang, Y.; Xu, M.; Yang, Y.; Yan, S.; and Wang, M. 2021. Deraincyclegan: Rain attentive cyclegan for single image deraining and rainmaking. *IEEE Transactions on Image Processing*, 30: 4788–4801.
- Wu, H.; Qu, Y.; Lin, S.; Zhou, J.; Qiao, R.; Zhang, Z.; Xie, Y.; and Ma, L. 2021. Contrastive learning for compact single image dehazing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 10551–10560.
- Xia, Y.; Monica, J.; Chao, W.-L.; Hariharan, B.; Weinberger, K. Q.; and Campbell, M. 2023. Image-to-Image Translation for Autonomous Driving from Coarsely-Aligned Image Pairs. In *2023 IEEE International Conference on Robotics and Automation (ICRA)*, 7756–7762. IEEE.
- Yasarla, R.; Sindagi, V. A.; and Patel, V. M. 2021. Semi-supervised image deraining using gaussian processes. *IEEE Transactions on Image Processing*, 30: 6570–6582.
- Ye, Y.; Yu, C.; Chang, Y.; Zhu, L.; Zhao, X.-l.; Yan, L.; and Tian, Y. 2022. Unsupervised Deraining: Where Contrastive Learning Meets Self-similarity. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5821–5830.
- 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 *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, 5728–5739.
- Zhang, H.; Ba, Y.; Yang, E.; Mehra, V.; Gella, B.; Suzuki, A.; Pfahnl, A.; Chandrappa, C. C.; Wong, A.; and Kadambi, A. 2023. Weatherstream: Light transport automation of single image deweathering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 13499–13509.
- Zhang, H.; and Patel, V. M. 2018. Density-aware single image de-raining using a multi-stream dense network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 695–704.
- Zhang, K.; Li, R.; Yu, Y.; Luo, W.; and Li, C. 2021a. Deep dense multi-scale network for snow removal using semantic and depth priors. *IEEE Transactions on Image Processing*, 30: 7419–7431.
- Zhang, Z.; Wang, H.; Liu, M.; Wang, R.; Zhang, J.; and Zuo, W. 2021b. Learning raw-to-srgb mappings with inaccurately aligned supervision. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 4348–4358.
- Zhang, Z.; Wang, R.; Zhang, H.; Chen, Y.; and Zuo, W. 2022. Self-supervised learning for real-world super-resolution from dual zoomed observations. In *European Conference on Computer Vision*, 610–627. Springer.