



UNIVERSITY OF LEEDS

This is a repository copy of *AllWeather-Net: Unified Image Enhancement for Autonomous Driving Under Adverse Weather and Low-Light Conditions*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/216341/>

Version: Accepted Version

Proceedings Paper:

Qian, C., Rezaei, M. orcid.org/0000-0003-3892-421X, Anwar, S. et al. (4 more authors) (2024) AllWeather-Net: Unified Image Enhancement for Autonomous Driving Under Adverse Weather and Low-Light Conditions. In: Pattern Recognition. ICPR 2024, 01-05 Dec 2024, Kolkata, India. Lecture Notes in Computer Science, 15330 . Springer , Cham, Switzerland , pp. 151-166. ISBN: 978-3-031-78112-4 ISSN: 0302-9743 EISSN: 1611-3349

https://doi.org/10.1007/978-3-031-78113-1_11

This version of the article has been accepted for publication, after peer review (when applicable) and is subject to Springer Nature's AM terms of use (<https://www.springernature.com/gp/open-research/policies/accepted-manuscript-terms>), but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: https://doi.org/10.1007/978-3-031-78113-1_11 .

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

AllWeather-Net: Unified Image Enhancement for Autonomous Driving Under Adverse Weather and Low-Light Conditions

Chenghao Qian¹[0009-0008-7054-9556], Mahdi Rezaei¹[0000-0003-3892-421X],
Saeed Anwar²[0000-0002-0692-8411], Wenjing Li³[0000-0003-3201-6675], Tanveer
Hussain⁴[0000-0003-4861-8347], Mohsen Azarmi¹[0000-0003-0737-9204], and Wei
Wang⁵[0000-0001-8676-1190]

¹ Leeds University, Leeds, UK
² KFUPM, Dhahran, Saudi Arabia
³ Chinese Academy of Sciences, Beijing, China
⁴ Edge Hill University, Liverpool, UK
⁵ Sun Yat-sen University, Guangzhou, China
tscq@leeds.ac.uk



(a) Original Input Images



(b) Enhanced Output Images

Fig. 1: Given images captured under adverse conditions in (a), we propose a method that can effectively adjust color and texture, modify lighting and shadows, and remove weather effects within a unified model. This results in a visually appealing appearance that resembles normal, day-like weather conditions (b), thereby enhancing the robust performance of autonomous driving perception systems.

Abstract. Adverse conditions like snow, rain, nighttime, and fog, pose challenges for autonomous driving perception systems. Existing methods have limited effectiveness in improving essential computer vision tasks, such as semantic segmentation, and often focus on only one specific condition, such as removing rain or translating nighttime images into day-time ones. To address these limitations, we propose a method to improve the visual quality and clarity degraded by such adverse conditions. Our method, AllWeather-Net, utilizes a novel hierarchical architecture to enhance images across all adverse conditions. This architecture incorporates information at three semantic levels: scene, object, and texture, by discriminating patches at each level. Furthermore, we introduce a Scaled Illumination-aware Attention Mechanism (SIAM) that guides

the learning towards road elements critical for autonomous driving perception. SIAM exhibits robustness, remaining unaffected by changes in weather conditions or environmental scenes. AllWeather-Net effectively transforms images into normal weather and daytime scenes, demonstrating superior image enhancement results and subsequently enhancing the performance of semantic segmentation, with up to a 5.3% improvement in mIoU in the trained domain. We also show our model’s generalization ability by applying it to unseen domains without re-training, achieving up to 3.9 % mIoU improvement. Code can be accessed at: <https://github.com/Jumponthemoon/AllWeatherNet>.

Keywords: Image enhancement · Semantic segmentation · Hierarchical discrimination · Illumination-aware attention

1 Introduction

Autonomous driving systems heavily rely on clear and optimal environmental images; however, these are not guaranteed in real life due to natural conditions, like snow, rain, fog, low light at night, etc. This can significantly reduce visibility and distort the captured information within an image, which impacts the performance of autonomous driving perception systems, including but not limited to object detection and semantic segmentation.

To counter the mentioned problem, some methods remove weather artifacts via deraining [22, 24], dehazing [3, 25], and desnowing [15, 21, 27]. Moreover, some unified frameworks [4, 12, 14] handle three types of weather while mainly focusing on removing hydrometer particles, neglecting alterations in color and texture details; hence, restricting their effectiveness under adverse weather conditions for autonomous driving computer vision systems.

In contrast to weather artifacts removal, pixel-level image translation approaches transform challenging weather situations into clear, sunny-day image styles. Regardless, these methodologies mainly focus only on specific individual conditions, such as rain [13] or nighttime scenarios [2]. In addition, the model may alter irrelevant pixels or areas and introduce unwanted changes, leading to visual discrepancies and negatively impacting the performance of downstream tasks. Likewise, low-light enhancement aims to improve the visibility and quality of images captured in low-light conditions. This involves enhancing the brightness, contrast, and details of dark images due to insufficient lighting; however, this technique can mistakenly brighten already well-lit areas, leading to overexposure in weather conditions like snow, as shown in Fig. 2.

We aim to improve image quality and clarity by adjusting image attributes and enhancing texture under four distinct adverse conditions, all within a unified framework. Subsequently, we seek to improve semantic segmentation performance. To achieve this goal, we need to consider several critical factors:

Firstly, while a unified network is cost-effective, weather variability introduces instability in the learning process. Therefore, it is crucial to identify a




Fig. 2: (a) Original Image. The evaluation of image processing techniques for semantic segmentation under adverse conditions reveals the deficiencies of (b) weather effect removal [4], (c) pixel-level translation [29], and (d) low-light enhancement [16]. Images processed by these methods either fail to sufficiently enhance image quality or introduce artifacts, affecting semantic prediction accuracy. (e) Our method, AllWeather-Net, effectively enhances color and texture detail while preserving most of the original image information, achieving the best performance.

stable and invariant signal that can guide the network’s learning, ensuring consistent performance across all conditions. *Secondly*, unfavorable conditions differently impact various regions within a captured image. For example, in foggy scenes, distant objects are more blurred than nearby ones due to light scattering and attenuation. In addition, adverse weather conditions tend to preserve larger patterns in images while diminishing the clarity of finer details. So, it is essential to focus on both the overall enhancement and the intricate recovery of texture details. This motivates us to design a network architecture that is contextually aware and sensitive to variations in texture. *Lastly*, employing a pair-to-pair training strategy can improve performance, yet finding perfectly matched pairs in autonomous driving scenes is challenging due to inaccurate GPS pairing and environmental variations. Alternatively, we consider adopting a strategy that utilizes roughly aligned images for more robust discrimination during training when exactly matched pairs are unavailable.

To address these challenges, we propose a novel architecture, namely AllWeather-Net, and our contributions can be summarized as follows:

- We are the first to introduce a unified image enhancement method to address image quality degradation under adverse weather and low-light conditions, including snow, rain, fog, and nighttime.
- To achieve robust image enhancement across various adverse conditions, we introduce a Scaled Illumination-aware Attention Mechanism (SIAM) that directs a balanced learning process towards different road elements irrespective of changes in weather and scenes.
- To achieve both overall image consistency and detailed enhancement, we design a novel architecture that enhances input images by conducting discrimination tasks at three hierarchical levels of semantic patches: scene, objects, and texture.

2 Related Work

In this section, we review the image processing techniques for adverse weather conditions and low-light environments.

Weather effect removal. Current methods for removing visual artifacts, including raindrops, fog particles, and snowflakes, utilize processes such as deraining [22, 24], dehazing [3, 25, 26] and desnowing [15, 21, 27]. Recently, a unified bad weather removal network was proposed in [14]. In [12], the researcher simplifies this architecture with a single encoder-single decoder network. To reduce the computational cost, [4] proposed a knowledge transfer mechanism via teacher-student architecture.

Pixel-level translation transforms the visual representation and convert adverse weather conditions into scenes resembling sunny, daytime environments. It involves a direct modification of the image pixels, altering the fundamental appearance and context of the scene. CycleGAN [29] introduced a cycle-consistency loss for unsupervised translation between the source and target domains. CUT [18] uses contrastive learning to ensure content preservation and style transfer. Santa [23] proposes an approach to find the shortest path between source and target images without paired information.

Low-light enhancement aims to adjust attributes of an image, such as lighting and color balance, to enhance the visual appearance in low-light conditions. Traditional methods utilize histogram equalization [1] and Retinex [10] to perform low-light image enhancement. Recent deep learning approaches proposed end-to-end framework [7, 8, 16]. Compared to traditional methods, these frameworks demonstrate the capability of enhancing the quality of images captured in low-light conditions.

Limitation of Existing Works. Removing weather-related unfavourable effects typically targets minor disturbances such as snowflakes or raindrops in the image. However, merely eliminating these atmospheric particles is insufficient, as the primary cause of image quality degradation often stems from alterations in colors and texture details, which significantly contribute to domain shifts. This limitation also applies to pixel-level translation, which frequently introduces unwanted artifacts, thereby reducing the overall image quality and constraining their applicability in safety-critical scenarios. Similarly, low-light enhancement techniques, while focusing on improving visibility under low-light conditions, do not adequately address the challenges posed by adverse weather conditions.

3 Proposed Method

Our proposed method uses a generative model for generating image enhancement masks based on the original input image. We introduce a scaled illumination-aware attention mechanism (SIAM) within a unified framework to focus learning on road elements regardless of weather condition. Additionally, our hierarchical




Fig. 3: Comparison of pixel-level translation and image enhancement process.

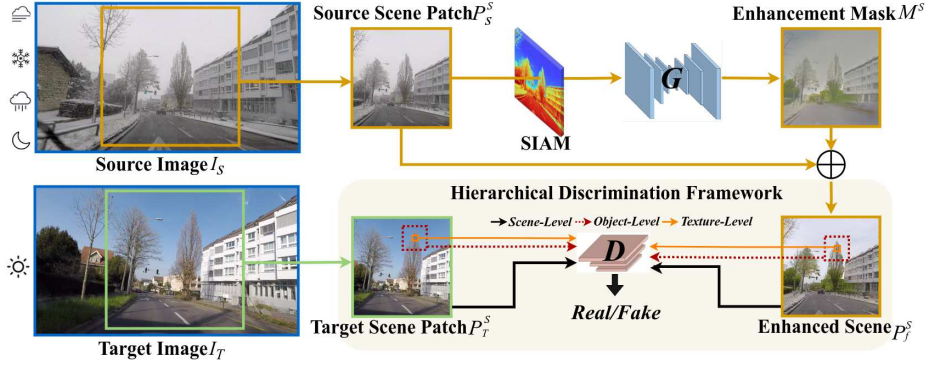


Fig. 4: Overview of AllWeather-Net architecture. SIAM: Scaled Illumination-aware Attention Mechanism. AllWeather-Net can enhance images across all adverse conditions (*e.g.*, fog, snow, rain, nighttime) with the help of the proposed SIAM and Hierarchical Discrimination Framework.

framework performs discrimination at multiple semantic levels, ensuring consistent and detailed enhancement. To further ensure precise enhancement, we utilize a ranked adaptive window pairing strategy for accurate discrimination.

3.1 Enhancement Pipeline

Our image enhancement method involves two networks: a generator and a discriminator, which are trained simultaneously through adversarial training to enhance image quality. Unlike pixel-level translation (Fig. 3a), where the generator takes the source image I_S and directly outputs translation results to mimic the style of the target image, our image enhancement generates intermediate results that are then combined with the original image I_S to produce final enhancement results. As illustrated in Fig. 3b, the process of generating enhanced image I' can be formulated as:

$$I' = G(I_S) + I_S. \quad (1)$$

Pixel-level translation often suffers from generating unwanted artifacts, which




Fig. 5: Scaled illumination-aware attention mechanism in the generator.




Fig. 6: Attention scores by illumination: naive vs. scaled attention.

can be attributed to the large search space during the training process. By conditioning the output on the input image, the image enhancement model can effectively reduce the search space for the generated result through residual learning. This method ensures that the model’s outputs are contextually relevant to the original image, thereby significantly reducing the likelihood of producing unwanted artifacts and improving the quality of the generated images.

In our model (Fig. 4), we initially cropped the same area from the paired source images I_S and target image I_T as input. The cropped source scene patch P_S^s is processed through a scaled illumination-aware attention mechanism and the generator to produce an enhancement mask M^s . This mask is then added to generate the final enhanced results P_f^s and evaluated by different discriminators according to its scale.

3.2 Scaled Illumination-aware Attention Mechanism (SIAM)

Training a unified network for image enhancement across different adverse conditions is cost-effective yet challenging. Each condition uniquely alters the scene’s visibility, color, and texture, complicating the learning of consistent information across different scenarios. This variability can significantly degrade the model’s capacity to effectively enhance images, especially when adverse conditions heavily obscure scene details. Given these complexities, the importance of a condition-invariant signal in guiding the learning process is paramount. Such a signal should guide the model to learn critical aspects of the scene regardless of weather or lighting conditions.

Drawing inspiration from previous work [9] targeting low-light conditions, we consider using illumination as a guiding cue. However, the naive approach of employing illumination intensity as attention tends to overemphasize areas of low illumination while neglecting well-lit regions. This can result in inadequate focus on pixel regions obscured by snow or fog, which often appear brighter due to higher illumination levels. This discrepancy can lead to suboptimal learning outcomes, as crucial details in these areas may not receive sufficient attention, resulting in inconsistent enhancements in the generated images.

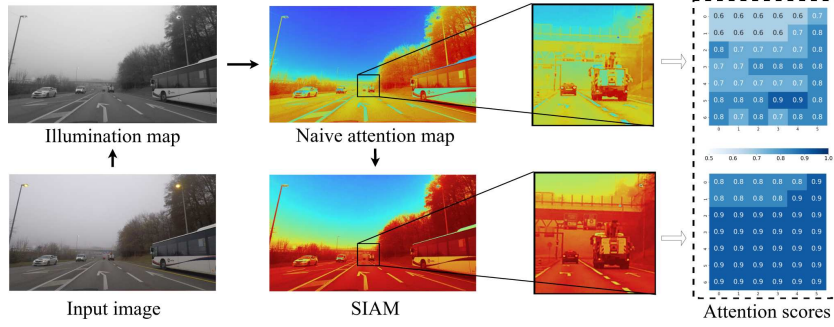


Fig. 7: The flow of generating attention map with SIAM and the comparison between naive attention and SIAM at the image and patch levels. Note that a higher attention score indicates that the model is paying more attention to such an area. This observation suggests that the presented SIAM, compared to naive attention mechanisms, is more adapt at focusing on areas containing critical road elements.

To direct a balanced learning for different road elements, rather than merely focusing on dark regions, we propose the scaled illumination-aware attention mechanism (SIAM) to allocate reasonable attention based on illumination intensity. Let I_{ij} and Att_{ij} denote the illumination intensity and the naive illumination attention value at the given pixel location (i, j) , the scaled illumination attention S_Att_{ij} can be formulated as follows:

$$Att_{ij} = 1 - I_{ij}, \quad (2)$$

$$S_Att_{ij} = -Att_{ij} \cdot (Att_{ij} - 2). \quad (3)$$

In the network, the SIAM will guide the learning throughout the generator shown in Fig. 5. As shown in Fig. 6, the scaled illumination-aware attention exhibits higher scores for high illumination and maintains consistently high attention levels across the input range for low illumination. This design demonstrates heightened sensitivity towards regions with high illumination while ensuring that low and high-illumination areas are allocated reasonable focus. With the implementation of the scaled attention, our model prioritizes objects in the distance obscured by fog particles with high illumination (Fig. 7).

3.3 Hierarchical Discrimination Framework

Hierarchical Discrimination. From the perspective of discrimination, employing a single discriminator generates unrealistic colors, while the global-local structure [9] has limited performance in providing fine-grained textural details. To address this issue, we propose our hierarchical discrimination architecture with scene-, object-, and texture-level patches/discriminators.

The scene-level discriminator D_s assesses a randomly cropped scene patch P^s from the original image I_s , as shown in Fig. 8. Its primary purpose is to

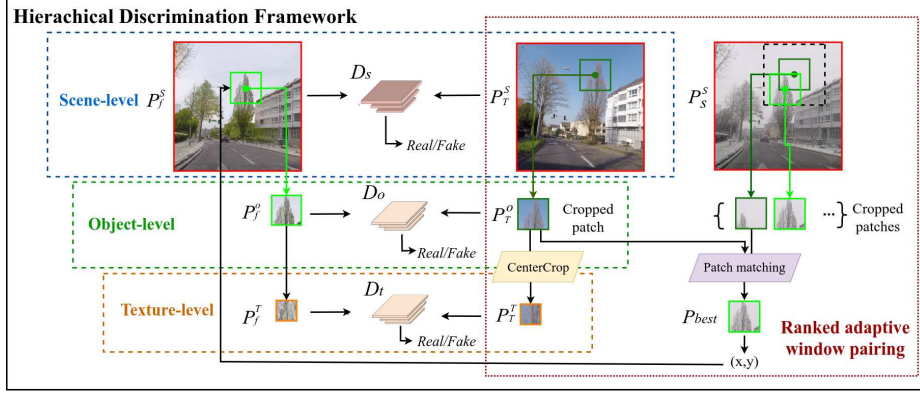


Fig. 8: The details of hierarchical discrimination framework.

evaluate the overall coherence and realism of the generated scene. Next, we crop object patches P^s with 1/4 of the size of the scene-level patch for the object-level discriminator D_o . This selective cropping allows for the exploration of intricate image details that the scene-level discriminator might overlook due to its broader perspective. Lastly, we derive the texture-level patch P^t by 1/4 center cropping from each object-level patch P^o . The texture-level discriminator D_t represents the highest level of inspection in our discriminator hierarchy, focusing on examining fine details and texture quality of the generated image.

By incorporating the mentioned three different levels of patches and discriminators, our model differentiates between the broad scales and fine levels of detail. This enables generated images that are not only realistic in overall appearance but also exhibit improved textural fidelity.

Ranked Adaptive Window Pairing. Scene-level discrimination is accurate due to similar information in paired images. However, accuracy drops at the object level, where patch pairs often show significant information shifts due to more apparent changes in perspective, leading to suboptimal outcomes.

To address the above mentioned issue, we utilize an adaptive window with a ranked score to identify the object-level patches that are most closely aligned. We begin by cropping object-level patches P^o at the same location from scene-level patches P^s . Moreover, we then define a fixed search area A with a width of w and a height of h . Within this area, we deploy a dynamic window W , of size $z \times z$, to traverse the defined area with a stride of s , thereby generating object-level patch candidates. These candidate patches are subsequently compared to the corresponding target patch using a ranked pairing score to determine the best match of P_S^o and P_T^o . Finally, we center crop the top-matched object pairs to obtain patches P_S^t and P_T^t for texture-level discrimination. Let (x, y) represent the top-left coordinate of the matched source object patch. $N = \frac{x+w-z}{s}$ repre-

sents the horizontal steps required, while $M = \frac{x+h-z}{s}$ denotes the vertical steps needed to traverse the search area with the dynamic window. The formulation of the pairing score F is defined as:

$$F(P_{Sc}^o, P_T^o) = \sum_{i=0}^N \sum_{j=0}^M |I_{P_T^o}(x, y) - I_{P_{Sc}^o}(x + i \cdot s, y + j \cdot s)|, \quad (4)$$

where P_{Sc}^o denotes the c -th candidate patch within the search window area and I represents the RGB values of the patch. The variables i and j signify the horizontal and vertical offsets within the search window, respectively. The best-matched patch, denoted as P_{best} , is then determined:

$$P_{\text{best}} = \arg \min_{P \in W} S(P_{Sc}^o, P_T^o). \quad (5)$$

By identifying the location of P_{best} in the source scene, we can locate its counterpart in the corresponding generated scene patch.

3.4 Loss Function

We utilize a relativistic approach [11] that compares the realism between real and generated images. We employ LSGAN [17] loss for direct assessment of the realism of the object- and texture-level discrimination. The scene-level losses for the discriminator and generator are given below:

$$L_D^s = \mathbb{E}_{P_r^s \sim P_{\text{real}}} [(D_s(P_r^s, P_f^s) - 1)^2] + \mathbb{E}_{P_f^s \sim P_{\text{fake}}} [D_R(P_f^s, P_r^s)^2], \quad (6)$$

$$L_G^s = \mathbb{E}_{P_f^s \sim P_{\text{fake}}} [(D_s(P_f^s, P_r^s) - 1)^2] + \mathbb{E}_{P_r^s \sim P_{\text{real}}} [D_R(P_r^s, P_f^s)^2], \quad (7)$$

where D_s represents the relativistic discriminator, P_r^s and P_f^s denote the real and generated fake scene patch. $\mathbb{E}_{P_r^s \sim P_{\text{real}}}$ and $\mathbb{E}_{P_f^s \sim P_{\text{fake}}}$ represent expectations over the real and fake data distributions. For object- and texture-level loss, The discriminator and generator for losses $P^x \in \{P^o, P^t\}$ are given by:

$$L_D^x = \mathbb{E}_{P_r^x \sim P_{\text{real}}} [(D_x(P_r^x) - 1)^2] + \mathbb{E}_{P_f^x \sim P_{\text{fake}}} [(D_x(P_f^x) - 0)^2], \quad (8)$$

$$L_G^x = \mathbb{E}_{P_f^x \sim P_{\text{fake}}} [(D_x(P_f^x) - 1)^2]. \quad (9)$$

Here, P_r^x and P_f^x denote real and fake patches of type x , respectively. D_x represents the discriminator for the patch type x , and $\mathbb{E}_{P_f^x \sim P_{\text{fake}}}$ and $\mathbb{E}_{P_r^x \sim P_{\text{real}}}$ are the same meaning as for scene patch. Consider the set $\mathcal{L} = \{s, o, t\}$ corresponding to the types of losses and use λ_1 , λ_2 , and λ_3 to control each loss contribution to balance loss; the total training loss can be written as:

$$\text{Total Loss} = \sum_{\ell \in \mathcal{L}} \Lambda_\ell \cdot (L_G^\ell + L_D^\ell),$$

where $\Lambda_s = \lambda_1$, $\Lambda_o = \lambda_2$, and $\Lambda_t = \lambda_3$.

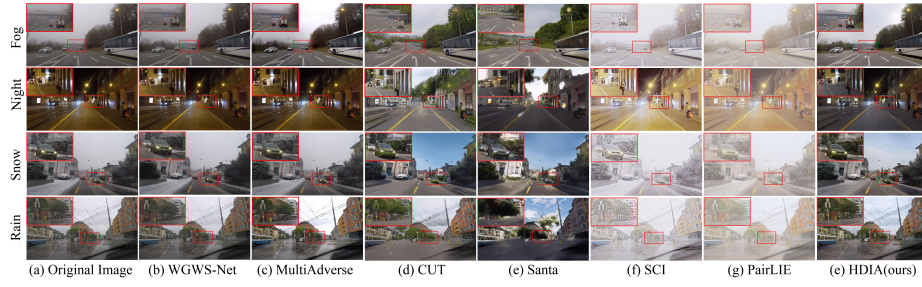


Fig. 9: Comparison with other image processing methods regarding weather effect removal, pixel-level translation, and low-light enhancement, using zoomed-in red regions to highlight visual distinctions.

4 Experiments

We conduct our experiments on image enhancement and evaluate the outcomes from two perspectives: image quality and semantic segmentation. Both aspects are assessed qualitatively and quantitatively.

4.1 Dataset

For image enhancement model training, we use 1,600 images from the ACDC [20], evenly distributed among snow, rain, night, and fog conditions, and 2,416 nighttime images from the Dark Zurich [5]. For the evaluation of semantic segmentation, our model enhances images from both the ACDC and the Dark Zurich validation set, which are subsequently tested using a pre-trained PSPNet [28] model. To demonstrate the generalization capabilities of our model, we apply it to the test datasets of Foggy Zurich [19] and Nighttime driving [6].

4.2 Comparisons

We evaluate AllWeather-Net against three distinct approaches: (a) weather effect removal, (b) pixel-level translation, and (c) low-light enhancement. We employ the Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM) to assess enhancement quality. We utilize the Natural Image Quality Evaluator (NIQE) to evaluate the image naturalness. Additionally, we consider the similarity between the improved and reference daytime images employing the Learned Perceptual Image Patch Similarity (LPIPS). We adopt the Mean Intersection over Union (mIoU) metric to evaluate semantic segmentation performance.

Image quality. In Fig. 9, the first column represents the input images captured under different adverse conditions while the subsequent columns are the enhanced results by different models. We observe that weather effect removal methods such as WGWS-Net [30] and MultiAdverse [4] excel in eliminating

Table 1: Qualitative evaluation of image quality and semantic segmentation performance on ACDC dataset.

Method Type	Models	Metrics				
		SSIM \uparrow	PSNR \uparrow	NIQE \downarrow	LPIPS \downarrow	mIoU \uparrow
Weather effect removal	WGWS-Net [30]	0.3916	11.4812	0.1480	0.4740	36.4
	MultiAdverse [4]	0.3822	10.9135	0.1782	0.4752	37.0
Pixel-level translation	CycleGAN [29]	0.3981	12.2493	0.1578	0.4655	33.6
	CUT [18]	0.3776	12.1043	0.1668	0.4833	29.3
	Santa [23]	0.3920	12.0863	0.1374	0.4770	25.7
Low-light enhancement	EnlightenGAN [9]	0.3905	11.8725	0.1651	0.4649	37.2
	Zero-DCE [8]	0.3160	10.6015	0.2726	0.4428	32.4
	SCI [16]	0.3239	8.5693	0.2600	0.5149	24.4
	PairLIE [7]	0.3577	8.9133	0.1564	0.4598	28.3
All weather enhancement	Ours	0.3983	11.6618	0.1257	0.4619	38.2

Table 2: Qualitative evaluation of image quality and semantic segmentation performance on Dark Zurich dataset.

Models	Metrics				
	SSIM \uparrow	PSNR \uparrow	NIQE \downarrow	LPIPS \downarrow	mIoU \uparrow
EnlightenGAN [9]	0.3791	10.2460	0.3186	0.4766	10.8
Zero-DCE [8]	0.3324	8.6448	0.3425	0.5109	10.6
SCI [16]	0.3519	8.5735	0.2554	0.5115	11.0
PairLIE [7]	0.3827	9.1442	0.2083	0.4844	7.1
Ours	0.3849	9.6121	0.1850	0.4589	17.6

weather-related artifacts but fall short in achieving a sunny daytime appearance. Pixel-level translation methods, including CUT [18] and Santa [23] can effectively transform images from adverse conditions to daytime-like settings, yet these models struggle with visual consistency in nighttime scenes and may inadvertently remove less prominent objects, such as a car in the distance, in foggy conditions. Low-light enhancement-based method SCI [16] and PairLIE [7] can enhance lighting and brightness in night scenes but over-expose in well-lit conditions, e.g., fog and snow. In contrast, our method demonstrates the most significant improvements in brightness and contrast across various scenes. It delivers the most realistic and clear daytime representation under diverse adverse conditions, significantly enhancing visibility without introducing artifacts or excessive noise. Additionally, it outperforms other models in color correction and detail enhancement, offering a more comprehensive solution. As shown in Table 1 and Table 2, our method achieves the most natural image enhancement outcomes with the lowest NIQE score and excels in converting images to the daytime domain, indicated by the lowest LPIPS scores for nighttime scenes.

Semantic segmentation. The effectiveness of our image enhancement model for semantic segmentation is assessed by performing a direct evaluation using the pre-trained PSPNet model [28]. We apply the model to datasets enhanced by our model as well as those enhanced by others. As shown in last column in Table 1 and Table 2, our method demonstrates superior performance in both adverse weather and nighttime scenes. This indicates that our model can enhance the performance of semantic segmentation models by significantly improving

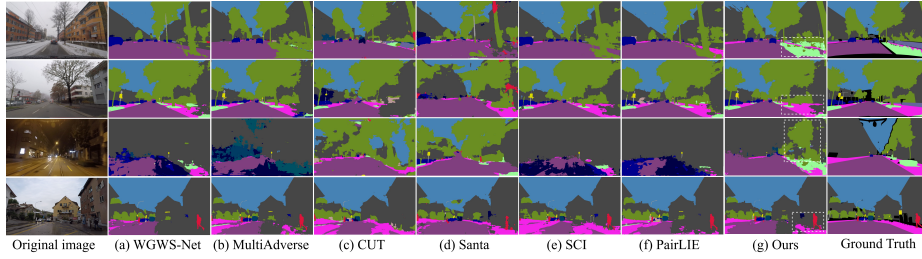


Fig. 10: Semantic segmentation results in comparison with other state-of-the-art methods of weather effect removal, pixel-level translation, and low-light enhancement, using zoomed-in white regions to highlight visual distinctions.



Fig. 11: The generalization performance of our model on the Foggy Zurich and Nighttime Driving datasets. The red and green box corresponds to the zoomed-in patches.

visual quality and visibility compared to other image processing models. In the visualization results shown in Fig. 10, our method improves the detail recognition of road elements such as trees, grass, and pedestrians in all conditions.

Generalization to unseen datasets. We evaluate our trained model’s performance in scenarios not seen during training using the Foggy Zurich and Nighttime Driving datasets. The results in Fig. 11 show that our model enhances clarity of cars, traffic lights, and road signs in the Foggy Zurich dataset, and corrects the yellowish glow on buildings and trees in the Nighttime Driving dataset, restoring their true colors and visibility. The mIoU comparison (Table 3) shows improvements of 1.8% and 3.9% respectively, highlighting the remarkable generalization capability of our model and demonstrating its ability to enhance semantic segmentation performance in unseen domains without re-training.

Table 3: mIoU scores tested on pretrained PSPNet [28] for original and enhanced versions of Foggy Zurich and Nighttime Driving test datasets.

Datasets	Original	Enhanced
Foggy Zurich [19]	26.3	28.1
Nighttime Driving [6]	23.0	26.9

Table 4: Qualitative comparison of model components using SSIM. A higher SSIM value indicates better image generation quality as it signifies greater similarity to a clear daytime image in terms of structure, luminance, and contrast.

Models	Components					SSIM \uparrow
	D_s	D_o	D_t	RWAP	SIA	
M_1	✓					0.3864
M_2	✓	✓				0.3879
M_3	✓	✓	✓			0.3910
M_4	✓	✓	✓	✓		0.3922
M_5	✓	✓	✓	✓	✓	0.3983




Fig.12: Naive attention and SIAM maps for various input adverse condition images. Darker regions indicate higher attention scores.




Fig.13: Results generated by models trained with naive attention and scaled attention.

4.3 Ablation Studies

To demonstrate the impact of each component in our method, we conduct ablation experiments with a focus on image quality improvement. These experiments examine different levels of discrimination, Ranked Adaptive Window Pairing (RAWP), and the Scaled Illumination Attention mechanism (SIAM).

In table Table 4, we observe enhanced image quality, as indicated by higher SSIM values, with the incremental inclusion of discriminators: scene discriminator D_s , object discriminator D_o , and texture discriminator D_t . The introduction of RAWP further increases SSIM, indicating that the model learns finer details in local patches through pair-to-pair training. Subsequently, incorporating attention further elevates SSIM, demonstrating the effectiveness of attention in enhancing image quality. For various adverse conditions, the scaled illumination attention mechanism effectively focuses on road elements, as demonstrated in Fig. 12. The enhanced image quality in Fig. 13 indicates that scaled attention addresses uneven lighting issues by allocating attention appropriately, particularly to areas overlaid with high illumination. This highlights the capability of the scaled attention mechanism to focus on both low and high-illumination regions within road elements and adapt to all adverse conditions.

5 Conclusions

In this work, we introduced AllWeatherNet, a unified framework designed to enhance image quality under various adverse conditions such as snow, rain, fog, and nighttime. Our objective was to develop a singular model capable of simultaneously addressing these four conditions without introducing artifacts that degrade image quality. The model can adjust lighting, brightness, and color in images in both adverse and normal weather conditions, transforming them into clear, daytime-like visuals. We implemented a hierarchical framework to recover color and texture details, along with a ranked adaptive window pair-to-pair training strategy to boost performance. We also developed a scaled-illumination attention mechanism to direct the learning process towards low and high-illumination areas, making it adaptable to different adverse scenarios. We performed semantic segmentation experiments using our enhanced dataset and observed notable improvements. Additionally, the model demonstrated exceptional generalization capability across a range of datasets without requiring re-training.

References

1. Abdullah-Al-Wadud, M., Kabir, M.H., Dewan, M.A.A., Chae, O.: A dynamic histogram equalization for image contrast enhancement. *IEEE transactions on consumer electronics* **53**(2), 593–600 (2007)
2. Anoosheh, A., Sattler, T., Timofte, R., Pollefeys, M., Van Gool, L.: Night-to-Day Image Translation for Retrieval-based Localization (Mar 2019), <http://arxiv.org/abs/1809.09767>, arXiv:1809.09767 [cs]
3. Berman, D., Avidan, S., et al.: Non-local image dehazing. In: *IEEE conference on computer vision and pattern recognition*. pp. 1674–1682 (2016)
4. Chen, W.T., Huang, Z.K., Tsai, C.C., Yang, H.H., Ding, J.J., Kuo, S.Y.: Learning multiple adverse weather removal via two-stage knowledge learning and multi-contrastive regularization: Toward a unified model. In: *CVPR*. pp. 17653–17662 (2022)
5. Dai, D., Sakaridis, C., Hecker, S., Van Gool, L.: Curriculum Model Adaptation with Synthetic and Real Data for Semantic Foggy Scene Understanding (May 2019), <http://arxiv.org/abs/1901.01415>, arXiv:1901.01415 [cs]
6. Dai, D., Van Gool, L.: Dark model adaptation: Semantic image segmentation from daytime to nighttime. In: *International Conference on Intelligent Transportation Systems (ITSC)*. pp. 3819–3824. IEEE (2018)
7. Fu, Z., Yang, Y., Tu, X., Huang, Y., Ding, X., Ma, K.K.: Learning a simple low-light image enhancer from paired low-light instances. In: *CVPR*. pp. 22252–22261 (2023)
8. Guo, C., Li, C., Guo, J., Loy, C.C., Hou, J., Kwong, S., Cong, R.: Zero-Reference Deep Curve Estimation for Low-Light Image Enhancement. In: *CVPR*. pp. 1777–1786. IEEE, Seattle, WA, USA (Jun 2020)
9. Jiang, Y., Gong, X., Liu, D., Cheng, Y., Fang, C., Shen, X., Yang, J., Zhou, P., Wang, Z.: Enlightengan: Deep light enhancement without paired supervision. *IEEE transactions on image processing* **30**, 2340–2349 (2021)
10. Jobson, D.J., Rahman, Z.u., Woodell, G.A.: Properties and performance of a center/surround retinex. *Transactions on image processing* **6**(3), 451–462 (1997)

11. Jolicœur-Martineau, A.: The relativistic discriminator: a key element missing from standard gan. arXiv preprint arXiv:1807.00734 (2018)
12. Jose Valanarasu, J.M., Yasarla, R., Patel, V.M.: TransWeather: Transformer-based Restoration of Images Degraded by Adverse Weather Conditions. In: CVPR. pp. 2343–2353. IEEE (Jun 2022)
13. Kwak, J.g., Jin, Y., Li, Y., Yoon, D., Kim, D., Ko, H.: Adverse weather image translation with asymmetric and uncertainty-aware gan. arXiv preprint arXiv:2112.04283 (2021)
14. Li, R., Tan, R.T., Cheong, L.F.: All in one bad weather removal using architectural search. In: CVPR. pp. 3175–3185 (2020)
15. Liu, Y.F., Jaw, D.W., Huang, S.C., Hwang, J.N.: Desnownet: Context-aware deep network for snow removal. IEEE Transactions on Image Processing **27**(6), 3064–3073 (2018)
16. Ma, L., Ma, T., Liu, R., Fan, X., Luo, Z.: Toward fast, flexible, and robust low-light image enhancement. In: CVPR. pp. 5637–5646 (2022)
17. Mao, X., Li, Q., Xie, H., Lau, R.Y., Wang, Z., Paul Smolley, S.: Least squares generative adversarial networks. In: ICCV (Oct 2017)
18. Park, T., Efros, A.A., Zhang, R., Zhu, J.Y.: Contrastive learning for unpaired image-to-image translation. In: ECCV (2020)
19. Sakaridis, C., Dai, D., Hecker, S., Van Gool, L.: Model adaptation with synthetic and real data for semantic dense foggy scene understanding. In: ECCV. pp. 687–704 (2018)
20. Sakaridis, C., Dai, D., Van Gool, L.: ACDC: The Adverse Conditions Dataset with Correspondences for Semantic Driving Scene Understanding (Sep 2021), <http://arxiv.org/abs/2104.13395>, arXiv:2104.13395 [cs]
21. Wang, Y., Ma, C., Liu, J.: Smartassign: Learning a smart knowledge assignment strategy for deraining and desnowing. In: CVPR. pp. 3677–3686 (2023)
22. Wang, Z., Cun, X., Bao, J., Zhou, W., Liu, J., Li, H.: Uformer: A general u-shaped transformer for image restoration. In: CVPR. pp. 17683–17693 (2022)
23. Xie, S., Xu, Y., Gong, M., Zhang, K.: Unpaired image-to-image translation with shortest path regularization. In: CVPR. pp. 10177–10187 (2023)
24. Yang, W., Tan, R.T., Feng, J., Guo, Z., Yan, S., Liu, J.: Joint rain detection and removal from a single image with contextualized deep networks. IEEE transactions on pattern analysis and machine intelligence **42**(6), 1377–1393 (2019)
25. Zhang, H., Sindagi, V., Patel, V.M.: Joint transmission map estimation and dehazing using deep networks. IEEE Transactions on Circuits and Systems for Video Technology **30**(7), 1975–1986 (2019)
26. Zhang, J., Ren, W., Zhang, S., Zhang, H., Nie, Y., Xue, Z., Cao, X.: Hierarchical density-aware dehazing network. IEEE Transactions on Cybernetics **52**(10), 11187–11199 (2021)
27. Zhang, K., Li, R., Yu, Y., Luo, W., Li, C.: Deep dense multi-scale network for snow removal using semantic and depth priors. IEEE Transactions on Image Processing **30**, 7419–7431 (2021)
28. Zhao, H., Shi, J., Qi, X., Wang, X., Jia, J.: Pyramid Scene Parsing Network (Apr 2017), <http://arxiv.org/abs/1612.01105>, arXiv:1612.01105 [cs]
29. Zhu, J.Y., Park, T., Isola, P., Efros, A.A.: Unpaired image-to-image translation using cycle-consistent adversarial networks. In: ICCV. pp. 2223–2232 (2017)
30. Zhu, Y., Wang, T., Fu, X., Yang, X., Guo, X., Dai, J., Qiao, Y., Hu, X.: Represent, compare, and learn: A similarity-aware framework for class-agnostic counting. In: CVPR (2023)