SynthRSF: A Novel Photorealistic Synthetic Dataset for Adverse Weather Condition Denoising

Angelos Kanlis[a], Vazgken Vanian[b], Sotiris Karvarsamis[c], Ioanna Gkika[d], Konstantinos Konstantoudakis[e] and Dimitrios Zarpalas[f]
Visual Computing Lab (VCL), Information Technologies Institute (ITI), Centre for Research and Technology - Hellas (CERTH), Thessaloniki, Greece

Keywords: Synthetic Dataset, Image Restoration, Adverse Weather Conditions, Semantic Segmentation, Depth Estimation, Benchmarking, Unreal Engine.

Abstract: This paper presents the SynthRSF dataset for training and evaluating single-image rain, snow and haze denoising algorithms, as well as evaluating object detection, semantic segmentation, and depth estimation performance in noisy or denoised images. Our dataset features 26,893 noisy images, each accompanied by its corresponding ground truth image. It further includes 13,800 noisy images accompanied by ground truth, 16-bit depth maps and pixel-accurate annotations for various object instances in each frame. The utility of SynthRSF is assessed by training unified models for rain, snow, and haze removal, achieving good objective metrics and excellent subjective results compared to existing adverse weather condition datasets. Furthermore, we demonstrate its use as a benchmark for the performance of an object detection algorithm in weather-degraded image datasets.

1 INTRODUCTION

Creating scene understanding models has become a central goal in both computer vision research and associated industrial applications. Such tasks can involve object detection, segmentation, depth estimation, as well as more complex procedures. However, adverse conditions such as rain, snow and haze, as well as variable lighting conditions, can impact the performance of such algorithms by degrading the visual data. This can affect a wide range of applications, such as autonomous driving, surveillance, robotics, computer-assisted search-and-rescue, and more.

Due to the practical constraints of collecting rain, snow and haze-specific data with an associated ground truth at real-world sites, as well as the difficulty of defining the ground truth scene at a later time, when lighting and other variables have changed, significant research effort has been devoted to generating synthetic datasets for rain, snow and haze.

Although some datasets use a synthetic noise layer superimposed on real world images, the result often appears flat and unconvincing to a human observer. Furthermore, by simply layering weather noise on top of an image, one does not account for the effect of the weather phenomenon on the landscape nor the effect of existing lighting conditions on the appearance of the weather phenomenon itself. For these reasons, deep learning models trained on such datasets often perform poorly in real-world conditions, as the domain gap between the training set and the actual input in an application is significant.

A solution to the above would be a photorealistic synthetic dataset including adverse weather effects as 3D effects fully integrated in a scene. In recent years, modern game engines are capable of producing highly realistic scenes, incorporating not only objects, weather effects and lighting, but also their interaction. Renders from such scenes can comprise multiple versions of the same view, including ones with adverse weather conditions of various types and intensities, as well as clear ground-truth images.

This paper presents SynthRSF (Synthetic with Rain, Snow and Fog), a novel, photorealistic, synthetic dataset focused on incorporating adverse weather conditions, created using the Unreal 5.2 game

[a] https://orcid.org/0009-0005-0836-9325
[b] https://orcid.org/0000-0002-2150-3446
[c] https://orcid.org/0000-0002-5302-3711
[d] https://orcid.org/0000-0001-7340-3079
[e] https://orcid.org/0000-0001-5092-8796
[f] https://orcid.org/0000-0002-9649-9306

Kanlis, A., Vanian, V., Karvarsamis, S., Gkika, I., Konstantoudakis, K. and Zarpalas, D.
SynthRSF: A Novel Photorealistic Synthetic Dataset for Adverse Weather Condition Denoising.
DOI: 10.5220/0012397700003660

In *Proceedings of the 19th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications (VISIGRAPP 2024) - Volume 3: VISAPP*, pages 567-574
ISBN: 978-989-758-679-8; ISSN: 2184-4321

engine. SynthRSF is based on 14 3D scenes, of various sizes (from indoor rooms to entire cities) set in various environments (urban day, urban night, interior, nature), within which the camera moves on a virtual rig, rendering images containing various types of noise: snow, rain, uniform and non-uniform fog. Each noisy frame is accompanied by the corresponding ground truth image, for training denoising models.

To showcase SynthRSF's added value in visual computing research, a series of experiments has been conducted, using it to train the state-of-the-art *TransWeather* (Valanarasu et al., 2022) adverse weather noise removal model. As a training dataset, SynthRSF exhibits promising performance compared to existing adverse weather datasets, and beyond state-of-the-art performance when used in combination with some existing datasets. In addition, a human subjective evaluation survey is performed, using real-world images. Results provide compelling evidence that when training models with photorealistic data, denoising results are consistently deemed preferable by human observers.

Furthermore, SynthRSF comes with an additional multi-modal expansion dataset, named SynthRSF-MM. The multi-modal dataset contains 14 scenes, with pixel-level annotations for 5 object instances per scene, and 41 object classes that are included in total. With its additional modalities, it can be used as a train and/or test dataset in a wider range of computer vision tasks, such as object detection, image segmentation, depth estimation, and scene understanding, with possible applications in autonomous driving, robotics, search-and-rescue, and more.

Hence, the main contributions of this paper can be summarised as follows:

- the **SynthRSF dataset**, a synthetic photorealistic dataset incorporating 3D weather effects and lighting, comprising 26,893 pairs of images (degraded with adverse weather and ground truth). SynthRSF, along with its expansion (see next bullet) is available on the Git repository[1].
- the **SynthRSF-MM expansion**, an additional 13,800 pairs with ground truth on additional modalities: depth map, semantic segmentation, and bounding box pixel coordinates for 39 classes.
- a novel **Dataset creation methodology** based on the Unreal 5.2 game engine, leveraging 3D models and effects and predefined virtual camera paths, used to create SynthRSF/SynthRSF-MM.
- a set of **Experiments** comparing SynthRSF to previously published adverse weather datasets in training image restoration models.

[1]https://github.com/VCL3D/SynthRSF

2 RELATED WORK

2.1 Unified Weather Denoising Models

Weather denoising research has recently undergone a shift from earlier optimization-based techniques, often requiring priors that are tailored to specific types of weather conditions, to deep-learning approaches (Yang et al., 2019) can be used to model multiple phenomena. The introduction of CNNs and GANs (Ren et al., 2020) has significantly enhanced denoising capabilities.

So far, the fraction of works on unified deraining, desnowing and dehazing is still significantly smaller than the research work on rain, snow or haze in isolation. However, very recently, multiple methods have emerged that follow a unified approach (Valanarasu et al., 2022), (Özdenizci and Legenstein, 2022), (Wang et al., 2023), (Chen et al., 2022), (Karavarsamis et al., 2022a). The authors of (Li et al., 2020) are one of the first to handle multiple weather degradations using a single network. The model is based on CNNs and consists of multiple weather-specific encoders and a single common decoder.

(Valanarasu et al., 2022) proposes a single encoder-decoder architecture based on transformers and uses weather queries to handle multiple adverse weather conditions. A novel transformer-based block is also proposed improving the networks' performance. (Wang et al., 2023) is another work based on transformers. To improve the learning capabilities and efficiency of the model, transformers-based blocks are used in a grid structure. The approach proposed by (Özdenizci and Legenstein, 2022) is based on denoising diffusion-based methods, introducing a patch based diffusive restoration architecture enabling arbitrary sized image processing.

2.2 Datasets Based on Real Images with Synthetic Weather Effects

For each single phenomenon denoising task, most methods use one or more of the following datasets, which have been documented in (Yang et al., 2020) and (Karavarsamis et al., 2022b).

For removing rain, datasets like Rain12600 (Fu et al., 2017) and Rain12000 (Zhang and Patel, 2018) have been widely used while similarly for snow, Snow-100K (Liu et al., 2018) and CSD (Chen et al., 2021) among others. When it comes to haze, notable datasets include I-HAZE (Ancuti et al., 2018a) and O-HAZE (Ancuti et al., 2018b). Recently, a novel technique was published (Ba et al., 2022) for generating ground truth for rainy images. However, similar approaches for other important weather conditions are still missing in the scientific literature.

2.3 Fully Synthetic Game Engine Datasets

There is a significant amount of fully synthetic datasets generated in game engines such as Blender and Unity 3D. Important milestones include SceneNet (A Handa, 2016) containing annotated 3D scenes that can generate unlimited ground truth training data, (Richter et al., 2016) who use game interaction with graphics hardware to generate labeled data and (Mayer et al., 2016) who provide a stereo video dataset to estimate disparity and scene flow. Furthermore, (Butler et al., 2012) create an optical flow dataset derived from a 3D animated short film. To our knowledge, weather noise has not been implemented in any fully synthetic game engine datasets.

3 THE SYNTHRSF DATASET

3.1 Design

The design goal of SynthRSF is to create a collection of photorealistic image pairs in different types of environments, of weather-degraded images and their corresponding clear ground-truth image.

This is achieved by adding 3D weather effects simulating rain, snow and fog to realistic 3D Scenes. This way, included weather noise can be parameterized into numerous combinations resulting in a wide range of visibility conditions.

Simulating real-life fog is particularly interesting, since it is both a cause of occlusion and simultaneously it interacts with existing light sources, increasing the illumination of parts of the scene. This type of simulation,[2] [3] has now been made possible with state-of-the-art game engines.

3.2 Environment

The content environment of SynthRSF is based on 14 3D scenes designed in the Unreal 5.2 game engine[4]

Scenes are sourced from the Unreal Engine's documentation, including Unreal's City Sample Project and Hillside Project, which contribute most of the images of the dataset, due to the quality and variety of the 3D assets contained. Other scenes are sourced from the documentation and the Unreal Marketplace, all detailed on the SynthRSF Git repository[1]

These scenes simulate a variety of environment and lighting characteristics: urban/rural, day/night, indoor/outdoor, wet/snowy, captured by a virtual camera moving along set paths. Aiming for high visual realism, interior scenes do not include snow or rain noise, but do include uniform or non-uniform fog, as it approximates light smoke and can be useful in emergency response applications.

Scenes 1–5 are divided into training (67%) and test sets (33%). without risking data leakage, as the camera does not revisit the same locations. Scenes 6-14 are entirely in the training set.

SynthRSF provides $26,893$ weather images of rain, snow, uniform fog and non-uniform fog. This novel addition of non-uniform fog is the reason for its name including “fog” rather than “haze”. All images are accompanied by their ground-truth pairs.

3.3 Weather Effect Implementation

Snow and rain are simulated in a virtual scene by combining elements: Particles (sprites from Unreal's Niagara System) that represent close to medium-distance occlusion and a fog component that simulates precipitation-induced light diffusion in larger distances. Particle dimensions, velocity, angle, population and fog density are assigned sinusoidal functions with different periods. Over enough time, all their possible combinations are represented. Uniform fog is created by Unreal Engine's ExponentialHeightFog module, while non-uniform fog was generated using the Unreal's Legacy particle system.

Blueprint functions for the snow and rain effects as well as the custom rendering preset are included in Git[1]

3.4 SynthRSF-MM (Multi-Modal) Set

3.4.1 Motivation and Functionality

In 3D game engines, native data on each asset is available, hence additional ground-truth modalities can be included in each sample.

SynthRSF-MM is an expansion to SynthRSF, containing fewer samples but including additional ground-truth modalities (**depth, segmentation, and object bounding boxes**). It features $13,800$ noisy images generated from 14 Unreal-Engine-sourced scenes. SynthRSF-MM's scenes have been manually populated with 39 classes of 3D objects, including persons, vehicles, animals, etc. Subsequently, 3D rain, snow and fog effects are added to the

[2] Lumen Global Illumination
[3] Volumetric Fog
[4] http://www.unrealengine.com/

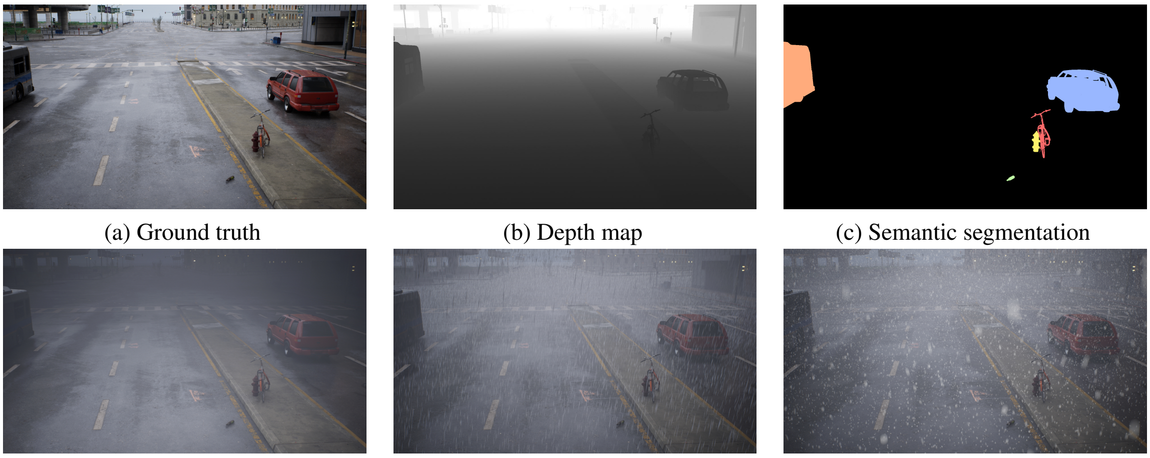

(a) Ground truth (b) Depth map (c) Semantic segmentation

(d) Fog (e) Rain (f) Snow

Figure 1: A sample scene from SynthRSF-MM. For each Ground truth image, there is one depth map, five pixel-level annotations and 8 noisy images per phenomenon.

scene. Each image is accompanied by a 16-bit depth map, pixel-accurate segmentation per object instance, YOLO-compatible bounding box .json files.

Due to the manual labour involved, 825 unique static camera views were set up, and 8 noisy images per phenomenon (rain, snow, fog), per view, were generated. Indoor scenes feature fog noise only.

Object detection bounding boxes are available for calculating the accuracy of an object detection task. SynthRSF-MM includes 39 of the YOLOv8[5] object detector classes. Occluded objects with fewer than 100 pixels appearing on an image are not being allotted a bounding box.

Having ground truth metadata, SynthRSF-MM can be used in tasks besides image restoration, such as semantic segmentation, object detection and distance estimation, both in clear and degraded conditions.

4 EXPERIMENTS

In testing SynthRSF and comparing it to previous datasets, *TransWeather*, a widely used state-of-the-art deep learning model for weather noise removal is employed. Experiments compare the results of *TransWeather* when trained by its default training dataset (AllWeather), SynthRSF, or a combination of both.

Three different experiments were conducted:

1. Objective: PSNR and SSIM metrics.
2. Subjective comparison on real adverse weather images.
3. Object detection: This experiment compares the efficiency of YoloV8 on images restored by *TransWeather*, trained on different datasets.

4.1 Training an Image Restoration Network for Adverse Weather Conditions

4.1.1 Architecture Selection

SynthRSF is suitable for unified bad weather removal architectures (i.e. "all-in-one" models that can remove multiple weather conditions). The presence of such publicly available models is rather limited. Three publicly available unified models were identified and tested: *TransWeather* (Valanarasu et al., 2022), *WeatherDiffusion* (Özdenizci and Legenstein, 2022), and *AirNet* (Li et al., 2022). Out of these, *TransWeather* features good robustness combined with a relatively fast training process; *WeatherDiffusion*, although superior in quality, is extremely slow both in training and inference, as also stated in the publication itself; and *AirNet* proved to be unstable at times due its contrastive learning approach.

Hence, the decision was made to employ *TransWeather* alone to conduct our experiments, allowing multiple training iterations and testing on tens of thousands of samples.

4.1.2 Training and Testing Datasets

In the original publication, *TransWeather* is trained on the AllWeather dataset, a combination of Snow100K (Liu et al., 2018), Outdoor-Rain (Li et al., 2019) and RainDrop (Qian et al., 2018) which contain images

[5]https://github.com/ultralytics/ultralytics

Table 1: Quantitative results based on PSNR and SSIM for *TransWeather* trained on AllWeather, SynthRSF, their combination and tested on Snow100K-L, test1 and SynthRSF test sets.

Train \ Test	Snow100K-L (AllWeather snow)		test1 (AllWeather rain and fog)		SynthRSF Snow		SynthRSF Rain		SynthRSF Fog		Average	
	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
AllWeather	28.08	0.86	**27.17**	0.87	24.83	0.79	24.33	0.73	20.07	0.73	24.90	0.80
SynthRSF	19.85	0.69	15.33	0.60	**27.68**	**0.85**	**27.89**	**0.83**	**25.17**	**0.82**	23.18	0.76
AllWeather + SynthRSF	**28.39**	**0.87**	27.16	**0.87**	27.48	**0.85**	27.82	**0.83**	24.66	**0.82**	**27.10**	**0.85**

Noisy | TransWeather (AllWeather) | TransWeather (SynthRSF) | TransWeather (AllWeather+SynthRSF)

Figure 2: Qualitative results from denoising real world images: (a) Snow (Row 1), (b) Rain (Row 2), (c) Fog (Row 3).

with snow, fog/rain and raindrop degradations respectively. In recent years AllWeather has become the go-to dataset for unified models. As such, it was chosen as the comparison dataset. For the evaluation process three instances of *TransWeather* are utilised on distinct datasets: (a) the original AllWeather dataset; (b) the SynthRSF dataset; and, (c) the combination of both datasets by using all images as input. Combining datasets often produces highly desirable results, as the literature (Liu et al., 2019; Yao et al., 2023) suggests. The three instances are evaluated on the combination of Snow-100KL (snow) and test1 (rain, fog), as well as the test set of SynthRSF (snow, rain, fog). Images with raindrop noise were not used. For training, SynthRSF images were downscaled from 1920x1080 to 720x405 pixels.

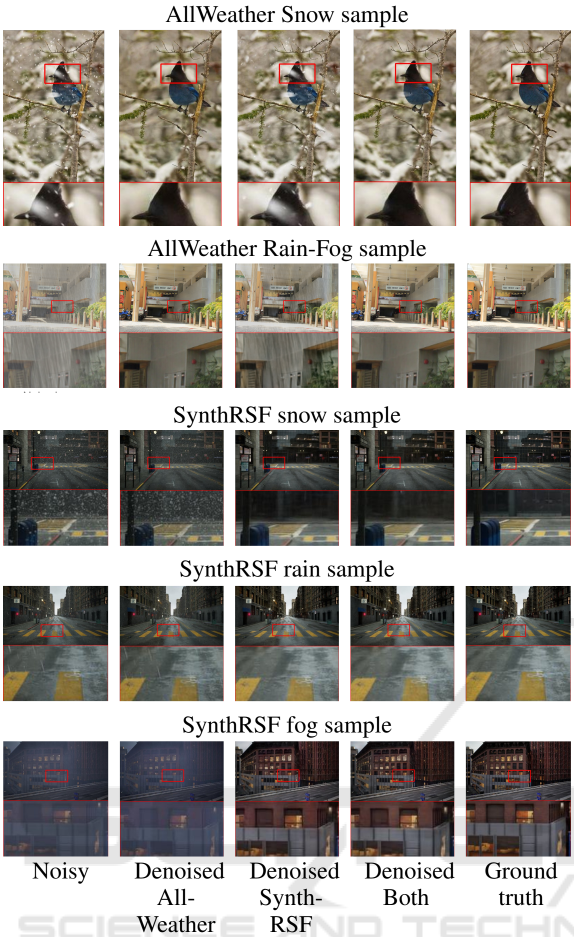

Figure 3: Denoising images from both datasets using trained TransWeather models.

4.1.3 Training Parameters

While training, *TransWeather* the default settings mentioned in (Valanarasu et al., 2022) are used without alteration to any of the hyperparameters. All models are trained on a single NVIDIA RTX 3090 GPU using the PyTorch framework (Paszke et al., 2019).

4.1.4 Quantitative Results

To evaluate performance we use the PSNR and SSIM metrics. The results for Snow-100KL, test1 and SynthRSF test set are summarized in Table 1. As expected, the model instance trained on the combination of the datasets demonstrates the best overall performance. Furthermore, using SynthRSF in combination with AllWeather improves the performance of *TransWeather* for Snow100KL while not hurting the performance for test1. On the contrary, when testing SynthRSF test set, AllWeather does not seem to improve but rather diminishes the performance of the model when used in combination with SynthRSF. This demonstrates the efficacy of SynthRSF when combined with existing datasets and as well as its comprehensiveness and efficiency as a standalone solution.

4.1.5 Qualitative Results

Synthetic Datasets. The predictions of the three model instances for images of the three test sets are illustrated in Figure 3. Similarly to quantitative results, the usage of SynthRSF appears to improve model performance, especially for the case of fog removal.

Real World Images. In this case (Figure 2), the effectiveness of SynthRSF is more apparent. All model instances showcase good denoising results, but the instances that used SynthRSF either by itself or in combination with AllWeather outperform the one that uses AllWeather only. Notably, the instance trained solely on SynthRSF removes even the farther and denser fog element, while the instance trained solely on AllWeather struggles to remove fog even at close distances.

4.2 Subjective Assessment Experiment

In order to evaluate training by SynthRSF in comparison with previous datasets, tests are performed on 75 real-world images collected from the Internet. Although a ground truth clear image for such cannot exist, and hence numerical results are not applicable, qualitative comparison can be performed on a subjective level.

Those images were fed to the three previously trained *TransWeather* models, and the results were evaluated by 70 survey participants. Participants, using their personal displays, with no time restrictions, accessed an online form to compare each noisy image with the three randomly ordered denoised versions, selecting the one they found clearest.

Survey results in Table 2 show a strong preference for SynthRSF-restored images, especially in images of rain and fog. While preferences for ”snowy” images are less pronounced, SynthRSF still leads. Despite the model trained on AllWeather often removing more individual snowflakes or rain streaks, the model trained on SynthRSF, because of its fog data, tends to clean up the distant parts of the image that are obscured by the fog-induced light scattering.

Table 2: Subjective assessment experiment - Summary of data from 70 survey participants for the total votes per model, and the number of times each model was chosen as the preferred option across 75 noisy images.

Training set	All-Weather	Synth-RSF	Both
Rain			
Total votes	149/1167	**610/1167**	408/1167
Top voted image	1/21	**12/21**	8/21
Snow			
Total votes	381/1773	**795/1773**	597/1773
Top voted image	3/30	**15/30**	12/30
Fog			
Total votes	149/1567	**911/1567**	507/1567
Top voted image	1/24	**18/24**	5/24

4.3 Benchmarking an Object Detector on Denoised Images

4.3.1 Benchmarking YoloV8 on Synthetic Noisy Images

As a test case for SynthRSF-MM, it is used to benchmark the performance of YOLOv8 on images containing adverse weather noise and their denoised counterparts. As testing data, the COCO validation dataset is used, with overlayed snow masks from CSD (Chen et al., 2021) and SRRS (Chen et al., 2020) datasets, as well as rain masks from RainTrainL (Zhang and Patel, 2018) dataset. For fog, the COCO dataset does not provide depth maps, so the RTTS dataset is used (Li et al., 2018) providing noisy images and object annotations. Results are summarised in Table 3.

4.3.2 Benchmarking YoloV8 on SynthRSF-MM

Demonstrating the utility of SynthRSF-MM's annotations, the denoising and object detection experiment is performed on the SynthRSF-MM images. The results are summarised in table 4. The combined AllWeather+SynthRSF dataset training produces better results in snow and rain images, while SynthRSF-only training was the most beneficial in fog images.

Table 3: YOLOV8 results on COCO/RTTS with synthetic noise, denoised by *TransWeather* trained on different datasets. Combined training is superior in snow and rain, SynthRSF performs better alone in fog.

Training set	mAP50	mAP50-95
Snow (COCO w/CSD/SRRS)		
noisy	0.623	0.466
AllWeather	0.629	0.468
SynthRSF	0.605	0.447
AllWeather+SynthRSF	**0.632**	**0.471**
Rain (COCO w/RainTrainL)		
noisy	0.626	0.466
AllWeather	0.615	0.455
SynthRSF	0.616	0.457
AllWeather+SynthRSF	**0.628**	**0.466**
Fog (RTTS)		
noisy	0.656	0.416
AllWeather	0.644	0.409
SynthRSF	**0.665**	**0.42**
AllWeather+SynthRSF	0.658	0.417

Table 4: YOLOV8 results on SynthRSF-MM dataset denoised by *TransWeather* trained on different datasets.

Training set	mAP50	mAP50-95
Rain		
noisy	0.293	0.218
AllWeather	0.319	0.239
SynthRSF	0.327	0.25
AllWeather+SynthRSF	**0.336**	**0.256**
Snow		
noisy	0.303	0.224
AllWeather	**0.326**	0.245
SynthRSF	0.314	0.238
AllWeather+SynthRSF	0.325	**0.247**
Fog		
noisy	0.296	0.22
AllWeather	0.288	0.217
SynthRSF	**0.307**	**0.227**
AllWeather+SynthRSF	0.302	**0.227**

5 CONCLUSIONS

In this paper, we have presented and shared SynthRSF, a novel synthetic dataset focused on adverse conditions. Its utility has been validated in multiple experiments: a) by training the *TransWeather* image denoising deep learning model in a series of both objective and subjective experiments; and b) by benchmarking a state-of-the-art object detection algorithm in its performance in the absence or presence of adverse weather conditions.

We have also presented SynthRSF-MM, a novel multi-modal dataset, which includes depth maps for all images, as well as pixel-level annotations for 39 object classes. Although its potential uses are many, the experiments highlight its functionality as a test set for measuring an object detector's performance in various inputs.

ACKNOWLEDGEMENTS

This research has been supported by the European Commission funded program RESCUER, under H2020 Grant Agreement 101021836.

REFERENCES

A Handa, V Patraucean, V. B. S. S. R. C. (2016). Understanding real world indoor scenes with synthetic data. In *CVPR 2016*.

Ancuti, C., Ancuti, C. O., Timofte, R., and De Vleeschouwer, C. (2018a). I-haze: a dehazing benchmark with real hazy and haze-free indoor images. In *ACIVS 2018, Proceedings 19*. Springer.

Ancuti, C. O., Ancuti, C., Timofte, R., and De Vleeschouwer, C. (2018b). O-haze: A dehazing benchmark with real hazy and haze-free outdoor images. In *CVPR 2018, NTIRE Workshop*, NTIRE CVPR'18, Salt Lake City, Utah, USA.

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 *ECCV*.

Butler, D. J., Wulff, J., Stanley, G. B., and Black, M. J. (2012). A naturalistic open source movie for optical flow evaluation. In *ECCV Proceedings, Part VI 12*.

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 *ECCV 2020 Proceedings, Part XXI 16*. 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 *CVPR 2021*.

Chen, W.-T., Huang, Z.-K., Tsai, C.-C., Yang, H.-H., Ding, J.-J., and Kuo, S.-Y. (2022). Learning multiple adverse weather removal via two-stage knowledge learning and multi-contrastive regularization: Toward a unified model. In *CVPR 2022*.

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 *CVPR 2017*.

Karavarsamis, S., Doumanoglou, A., Konstantoudakis, K., and Zarpalas, D. (2022a). Cross-stitched multi-task dual recursive networks for unified single image deraining and desnowing. In *WF-IoT 2022*.

Karavarsamis, S., Gkika, I., Gkitsas, V., Konstantoudakis, K., and Zarpalas, D. (2022b). A survey of deep learning-based image restoration methods for enhancing situational awareness at disaster sites: the cases of rain, snow and haze. *Sensors*, 22(13).

Li, B., Liu, X., Hu, P., Wu, Z., Lv, J., and Peng, X. (2022). All-in-one image restoration for unknown corruption. In *CVPR 2022*.

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).

Li, R., Cheong, L.-F., and Tan, R. T. (2019). Heavy rain image restoration: Integrating physics model and conditional adversarial learning. In *CVPR 2019*.

Li, R., Tan, R. T., and Cheong, L.-F. (2020). All in one bad weather removal using architectural search. In *CVPR*.

Liu, P., Zhou, X., Yang, J., El Basha Mohammad, D., and Fang, R. (2019). Image restoration using deep regulated convolutional networks.

Liu, Y.-F., Jaw, D.-W., Huang, S.-C., and Hwang, J.-N. (2018). Desnownet: Context-aware deep network for snow removal. *IEEE Trans. Image Processing*, 27(6).

Mayer, N., Ilg, E., Hausser, P., Fischer, P., Cremers, D., Dosovitskiy, A., and Brox, T. (2016). A large dataset to train convolutional networks for disparity, optical flow, and scene flow estimation. In *CVPR 2016*.

Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., Desmaison, A., Kopf, A., Yang, E., DeVito, Z., Raison, M., Tejani, A., Chilamkurthy, S., Steiner, B., Fang, L., Bai, J., and Chintala, S. (2019). PyTorch: An Imperative Style, High-Performance Deep Learning Library. In *Advances in Neural Information Processing Systems 32*.

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 *CVPR 2018*.

Ren, Y., Li, S., Nie, M., and Chuankun, L. (2020). Single image de-raining via improved generative adversarial nets. *Sensors*, 20.

Richter, S. R., Vineet, V., Roth, S., and Koltun, V. (2016). Playing for data: Ground truth from computer games. In *ECCV 2016 Proceedings, Part II 14*. Springer.

Valanarasu, J. M. J., Yasarla, R., and Patel, V. M. (2022). Transweather: Transformer-based restoration of images degraded by adverse weather conditions. In *CVPR 2022*.

Wang, T., Zhang, K., Shao, Z., Luo, W., Stenger, B., Lu, T., Kim, T.-K., Liu, W., and Li, H. (2023). Gridformer: Residual dense transformer with grid structure for image restoration in adverse weather conditions. *arXiv preprint arXiv:2305.17863*.

Yang, W., Tan, R. T., Wang, S., Fang, Y., and Liu, J. (2019). Single image deraining: From model-based to data-driven and beyond.

Yang, W., Tan, R. T., Wang, S., Fang, Y., and Liu, J. (2020). Single image deraining: From model-based to data-driven and beyond. *PAMI 2020*, 43(11).

Yao, M., Xu, R., Guan, Y., Huang, J., and Xiong, Z. (2023). Neural degradation representation learning for all-in-one image restoration.

Zhang, H. and Patel, V. M. (2018). Density-aware single image de-raining using a multi-stream dense network. In *CVPR 2018*.

Özdenizci, O. and Legenstein, R. (2022). Restoring vision in adverse weather conditions with patch-based denoising diffusion models.