

RestoreX-AI: A Contrastive Approach towards Guiding Image Restoration via Explainable AI Systems

Aboli Marathe

Symbiosis Centre for Applied AI (SCAAI),
Symbiosis International Deemed
University (SIU), India
SCTR's Pune Institute of
Computer Technology, India

aboli.rajan.marathe@gmail.com

Pushkar Jain

Symbiosis Centre for Applied AI (SCAAI),
Symbiosis International Deemed
University (SIU), India
SCTR's Pune Institute of
Computer Technology, India

pushrjain@gmail.com

Rahee Walambe

Symbiosis Centre for Applied AI (SCAAI),
Symbiosis International Deemed
University (SIU), India
Symbiosis Institute of Technology, SIU, India

rahee.walambe@gmail.com

Ketan Kotecha

Symbiosis Centre for Applied AI (SCAAI),
Symbiosis International Deemed
University (SIU), India
Symbiosis Institute of Technology, SIU, India

drketankotecha@gmail.com

Abstract

Modern applications such as self-driving cars and drones rely heavily upon robust object detection techniques. However, weather corruptions can hinder the object detectability and pose a serious threat to their navigation and reliability. Thus, there is a need for efficient denoising, deraining, and restoration techniques. Generative adversarial networks and transformers have been widely adopted for image restoration. However, the training of these methods is often unstable and time-consuming. Furthermore, when used for object detection (OD), the output images generated by these methods may provide unsatisfactory results despite image clarity. In this work, we propose a contrastive approach towards mitigating this problem, by evaluating images generated by restoration models during and post training. This approach leverages OD scores combined with attention maps for predicting the usefulness of restored images for the OD task. We conduct experiments using two novel use-cases of conditional GANs and two transformer methods that probe the robustness of the proposed approach on multi-weather corruptions in the OD task. Our approach achieves an averaged 178 increase in mAP between the input and restored images under adverse weather conditions like dust tornadoes and snowfall. We report unique cases where greater denoising does not improve OD performance

and conversely where noisy generated images demonstrate good results. We conclude the need for explainability frameworks to bridge the gap between human and machine perception, especially in the context of robust object detection for autonomous vehicles.

1. Introduction

Sensors are employed by unmanned autonomous vehicles to navigate through their surroundings, with a substantial dependence on vision-based sensors like RGB cameras. As these sensors are impacted by bad weather conditions, perception pipelines require considerable training on diverse data to increase the robustness on downstream tasks. One specific scenario which causes distortion of images is adverse weather conditions, like heavy snowfall, haze and dust tornadoes. In these critical situations, weather corruptions can hinder the object detectability and pose a serious threat to navigation and reliability. Thus, there is a need for efficient denoising, deraining and restoration techniques.

However, denoising techniques are often evaluated using Image similarity metrics like PSNR, SSIM [8] and not by their effectiveness in achieving results for the targeted application. It is possible that the output image of these methods has high image quality but is contextually irrelevant for the object detection task. In this work we evaluate the effective-

ness of restoration techniques for denoising images with the intention of better object detection. By introducing a contrastive approach towards restoration evaluation, a method for guiding the training of GANs and restoration progress is proposed. Additionally, attention maps are leveraged for understanding why these techniques assist object detection and why certain classes are easily recognized or ignored.

Primary contributions of this work are:

1. Exploring two new data-based generative adversarial network techniques for denoising weather corrupted images.
2. Proposing a novel contrastive approach using a weighted loss for evaluating the training progress and post-training performance of the highlighted restoration techniques.
3. Explaining the training progress of the proposed generative denoising methods using attention maps and validating the results using object detection evaluation metrics.
4. Evaluating the optimal noise level of the trained Restormer denoising methods (colour and grayscale) using attention maps and validating the results using object detection evaluation metrics.

2. Related Work

The rapid development and deployment of autonomous vehicles have exposed several critical challenges in computer vision, one of which being object detection robust to weather corruptions. At the foundation of this challenge is the classic vision task of object detection, the progress of which we summarize in the following sub-section 2.2. Ultimately, when these modules are integrated in autonomous driving systems, their suggestively “black-box” nature, triggers the concerns of industry professionals, carmakers and users alike. Without human-tailored explanations for the vehicle’s behavior, even the most advanced systems fail to benefit from widespread adoption due to concerns about safety and reliability. Especially in adverse weather, when even human drivers exercise extra caution, these robust models need to provide irrefutable explanations for the detections made at every state. We discuss the progress of previous works in these 3 directions which set the background of our work in restoration (Sub-section 2.1), detection (Sub-section 2.2) and explainability (Sub-section 2.3).

2.1. Multi-weather corruption and restoration

The thermal variations accompanying weather change can adversely impact the optical, electronic, and mechanical components used in capturing visual data, thus harming

the performance of visual recognition systems [6]. Frigid temperatures, snowfall or dense fog, for example, can cause condensation on the lens; rain streaks on car windows can generate glares or act as a double lens [20]. For an autonomous car, it is critical and essential to overcome the effects of weather conditions to ensure reliability. For example, one study [4] looked at the performance of gated cameras, while another [3] expanded the research to include stereo, gated, and thermal cameras, as well as Radar and LiDAR scanners, and found considerable increases in car recognition in varying levels of fog, and other adverse weather conditions. The use of unique methodologies like domain adaptation to transform the weather conditions while keeping objects of interest intact. For example, [22] investigates the consequences of synthetic weather images on road segmentation and traffic object detection, whereas [24] shows that using synthetic time-of-day (night imagery) improves localization, and [23] proposes a deraining model to improve semantic segmentation. For their efficiency in image restoration and effectiveness in removing weather corruptions, image restoration employing Restormers [34] and image denoising algorithms [5] are becoming increasingly popular.

2.2. Object Detection

The introduction of fast detectors like SSD, Faster RCNN, and YOLO transformed the face of object detection [16,25,26]. Taking 2D detectors forward, 3D detection expanded with Stereo RCNN [15], AVOD [13], MVLidar-Net [7] and MVF algorithm [37] bringing new perspectives to the task of object detection. Specific tasks like multiscale object detection [31], pedestrian detection in crowds [17] and detection under adverse weather [32] have also been solved using ensemble methods and data augmentation.

2.3. Explainability in Object Detection

The early rise of deployable machine learning technologies was accompanied by criticism of the “black-box”-like nature of ML models. Particularly in sensitive applications like medical diagnosis, self-driving cars and algorithmic test checking, there is a need for thorough explainability in models to ensure the trust and safety of users. To meet these requirements, many techniques were proposed for explaining models [1, 2, 29, 30, 36]. In [21], explainability is demonstrated by studying what each of its neurons has learned to detect. [19] focused on individual predictions, using the technique of heatmaps to highlight important pixels. Some works also interpret classifiers by identifying representative training examples [11, 12]. [27] introduced a new perspective to this challenge by making CNN-based models more transparent by producing visual explanations.

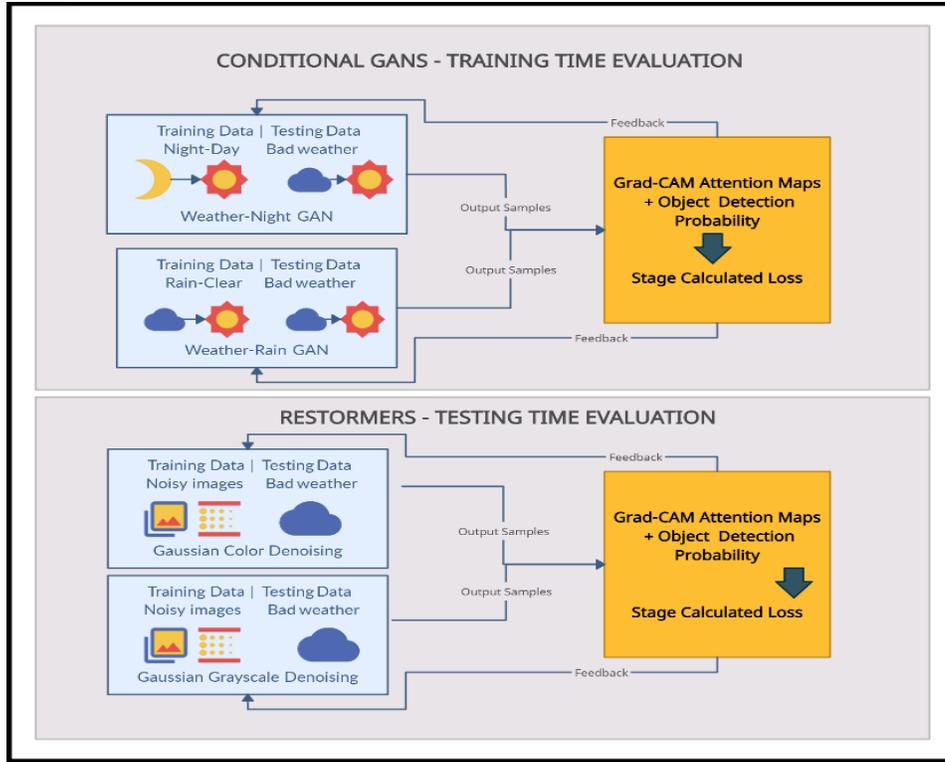


Figure 1. RestoreX-AI: Proposed contrastive approach for monitoring restoration models.

3. Methodology

This work proposes a contrastive approach (feedback) for monitoring the training progress of restoration models in the context of object detection. To put forward a diverse range of training samples, new restoration techniques using GANs have also been proposed. First, the restoration models were trained on different tasks and tested on the DAWN dataset [10]. As the models trained, their progress was monitored using the contrastive approach and simultaneously attention maps were generations to support the detection task. Finally, the OD performance of all the methods is compared using standard evaluation metrics (class AP and mAP). The 4 restoration techniques experimented with in this study are:

1. Weather-NightGAN: Conditional GANs trained on night-to-day task for multi-weather corruption restoration.
2. Weather-RainGAN: Conditional GANs trained on rain-to-clear task for multi-weather corruption restoration.
3. Restormer (Gaussian colour denoising) for multi-weather corruption restoration.

4. Restormer (Gaussian grayscale denoising) for multi-weather corruption restoration.

3.1. Conditional Generative Adversarial Networks

GANs or Generative Adversarial networks are generative models that learn mapping between noisy z and output image y , $G : z \rightarrow y$. Conditional GANs learn a mapping from observed image x and random noise vector z , to y , $G : x, z \rightarrow y$ [9]. The generator G is trained to produce images similar to the “real” images, as compared by an adversarially trained discriminator, D , which is used for detecting the “fakes”. The final objective of the conditional GAN can be expressed as:

$$G^* = \operatorname{argmin}_G \max_D \mathcal{L}_{cGAN}(G, D) + \lambda \mathcal{L}_{L1}(G)$$

where G tries to minimize this objective against an adversarial D that tries to maximize it, i.e. $G^* = \operatorname{argmin}_G \max_D \mathcal{L}_{cGAN}(G, D)$.

In this study, we propose 2 new use-cases of the conditional GAN for restoration purposes. In the first case, the GAN is trained on night-to-day images and tested for multi-weather corruption tasks. The intuition behind this is the similarity in corruptions of night and bad weather images,

like poor lighting and condensation on lens. In the second use-case, the GAN is trained on only rain-to-clear images (synthetically generated) and tested for multi-weather corruption tasks. The intuition behind this is the similarity in corruptions of rain and bad weather images, like snow and rain streaks. There is an additional challenge which notes if the single-weather trained conditional GAN can adapt to multi-weather corruptions.

3.2. Restormer

The Restormer is a highly efficient transformer that was proposed for denoising tasks in image restoration [35]. It consists of a multi-Dconv head transposed attention (MDTA) and a gated-Dconv feed-forward network (GDFN). These proposed architectural changes gave it the ability to capture long-range pixel interactions, while still remaining applicable to large images. It is both computationally efficient, and has the capacity to handle high-resolution images, a feature critical for a task like adverse weather object detection. In this work, we study the effects of trained noise levels (15, 25 and 50) on the denoising performance of colour and grayscale Restormers on the DAWN dataset. The goal is to study how the noise level affects the model performance in OD task and which level is optimal for generating explainable detections.

3.3. Grad-CAM

Gradient-weighted Class Activation Mapping (Grad-CAM) is a technique that produces visual explanations for the purpose of making CNN-based models more transparent [28]. For getting the class discriminative localization map Grad-CAM $L_{Grad-CAM}^c \in R^{u \times v}$ of width u and height v for any class c , first comes the computation of gradient of the score for class c , y^c (before the softmax), according to feature maps A^k of a convolutional layer, i.e. $\frac{\partial y^c}{\partial A^k}$. The neuron importance weights a_c^k are attained by global-average-pooling these gradients flowing back. The ‘importance’ of feature map k for a target class c is captured by this weight a_c^k . A weighted combination of forward activation maps is performed, and followed by a ReLU to obtain, This results in a coarse heat-map of the same size as the convolutional feature maps. Grad-CAM is used for the purpose of explaining object detection in the restored images of different techniques compared in this study. We additionally use the Grad-CAM model’s detection probability in calculating the contrastive metric for monitoring training progress.

3.4. Proposed Approach: RestoreX-AI

Due to the instability of training of GANs, over-or-under training does not always lead to the perfect solution images for object detection. Parallely, tuning on high noise levels does not always provide the best images from the Restormer

model. Even after producing good images by standard metrics (PSNR, SSIM), their applicability for the OD task remains uncertain, which brings the need for a new evaluation standard. We propose using a weighted sum of the explainability results (detection probability of class provided by the Grad-CAM model) and the similarity of the predicted and actual label to define this new standard. This weighted sum is calculated for every stage of training (one stage can be a user-defined set of epochs), and then used to monitor the progress of the model. We introduce this new parameter for assessing the quality of restoration which is calculated using equation 1.

$$\Delta\phi = \Delta(\Sigma(S(p, a) * d)/N) \quad (1)$$

Here S is the similarity of labels that be measured either by grouping the objects (cars, race cars and taxis have similarity 1, person, groom have 1 and so on), or by strict parameters (cars and race cars have similarity 0). p and a are the predicted and actual labels of the object under detection. N is the number of training samples generated in that stage, which are used to evaluate the current progress of that model. The explanation probability is d or the value returned by the Grad-CAM model, which is its prediction of what is present in the image. The quality of restoration between stages can be denoted as $\Delta\phi$, or the difference between qualities at consecutive stages.

3.5. Datasets

For training the Weather-RainGAN and Weather-NightGAN, corresponding images of the same scene in rain-clear and night-day conditions were required. For the Weather-RainGAN, the Rain 100L [33], a synthesized data of rain streaks with corresponding rain-free images was used. For the Weather-NightGAN, the Transient Attributes dataset [14] was used. This data was generated by a high-level image editing method which allows a user to adjust the attributes of a scene, e.g., change a scene to be “night” or “day”. The final testing of all restoration methods required a multi-weather dataset with high-resolution images, for which the DAWN dataset was selected. The DAWN dataset is a large vehicle detection dataset which has captured images of driving scenes in adverse weather conditions [10]. It consists of 1000 images from real-traffic scenes as seen in multiple adverse weather conditions including fog, snow, rain, and sandstorms. The images have been annotated with 2D annotations (boxes) with 6 object classes namely car, bus, truck, motorcycle, person and bicycle.

4. Experiments and Results

4.1. Restormer-based Grayscale Denoising

Section 1 of Table 1 shows the object detection scores of Gaussian grayscale image denoising using the Restormer

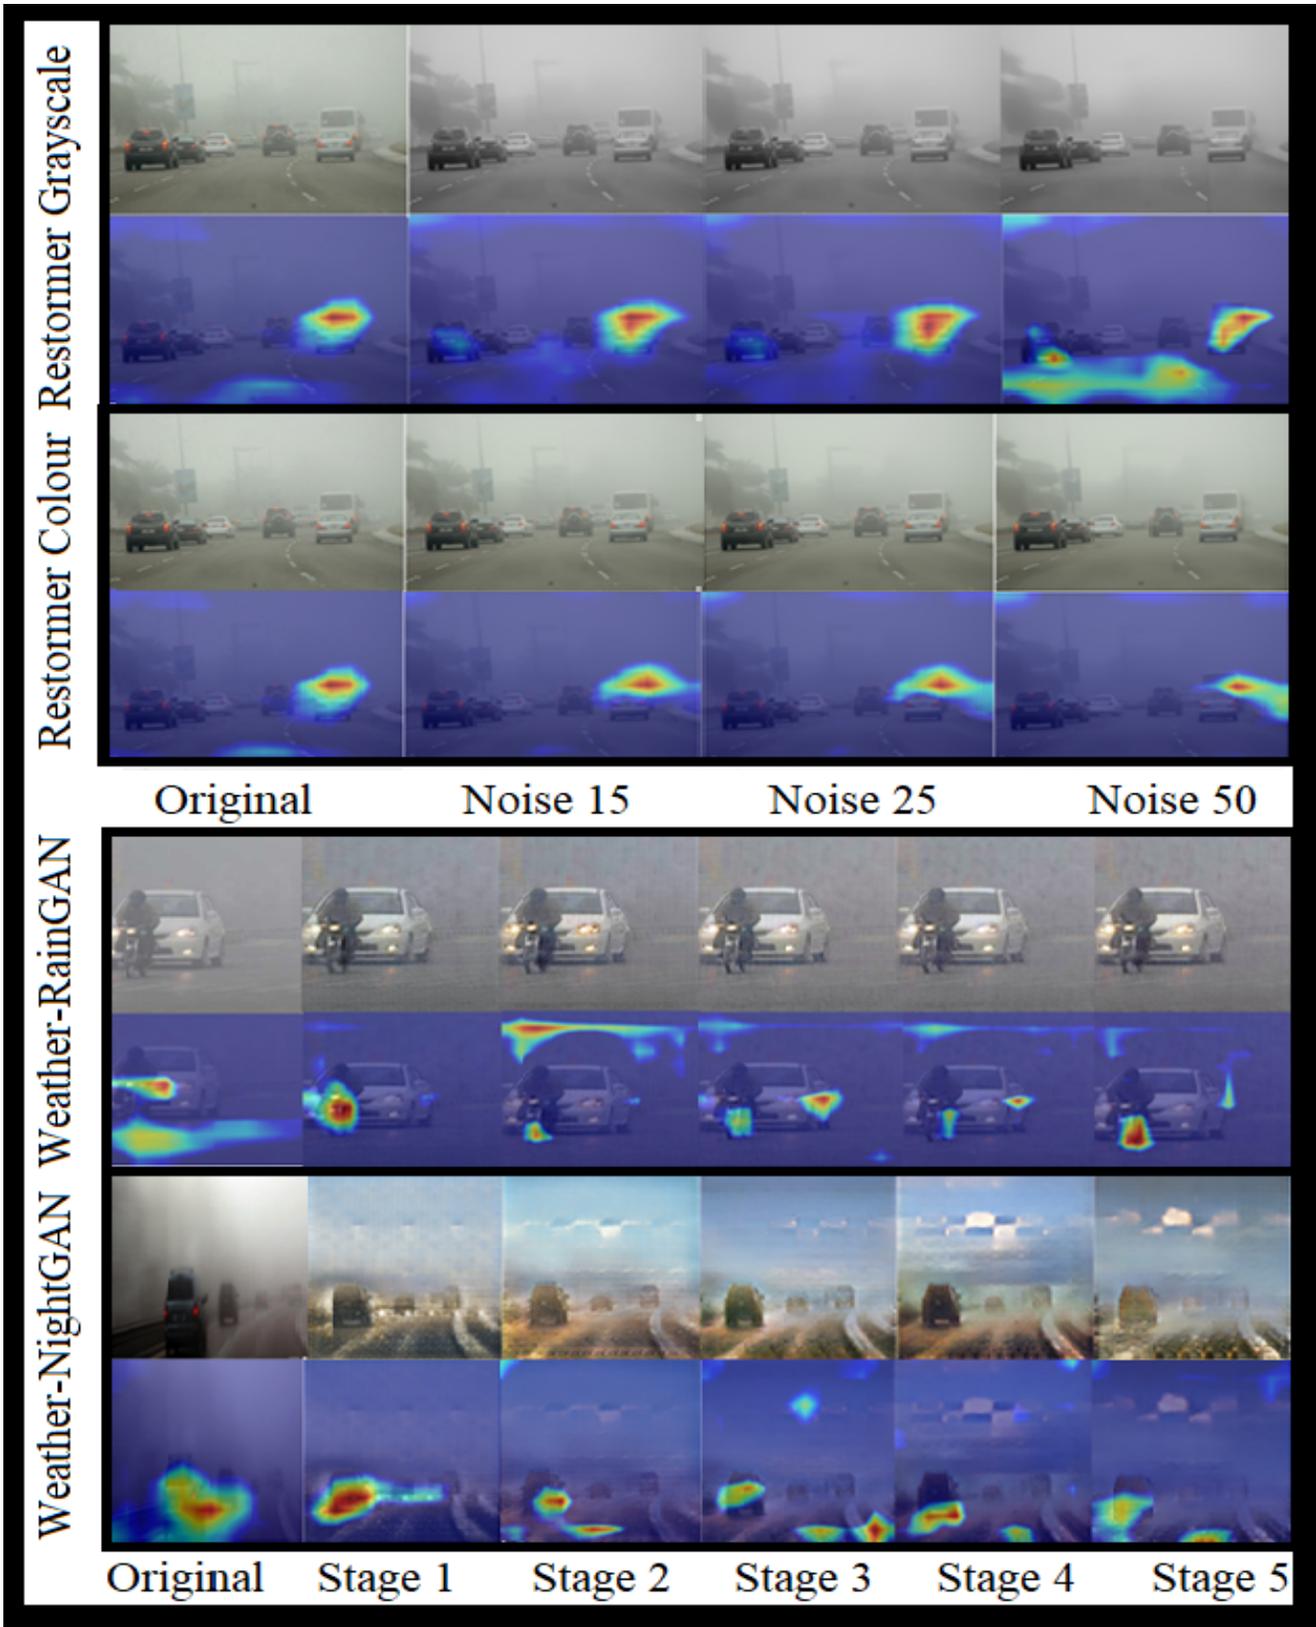


Figure 2. Foggy weather condition: Grad-CAM Attention Maps for (a) Restormer Grayscale Denoising (b) Restormer Colour Denoising (c) Weather-RainGAN and (d) Weather-NightGAN.

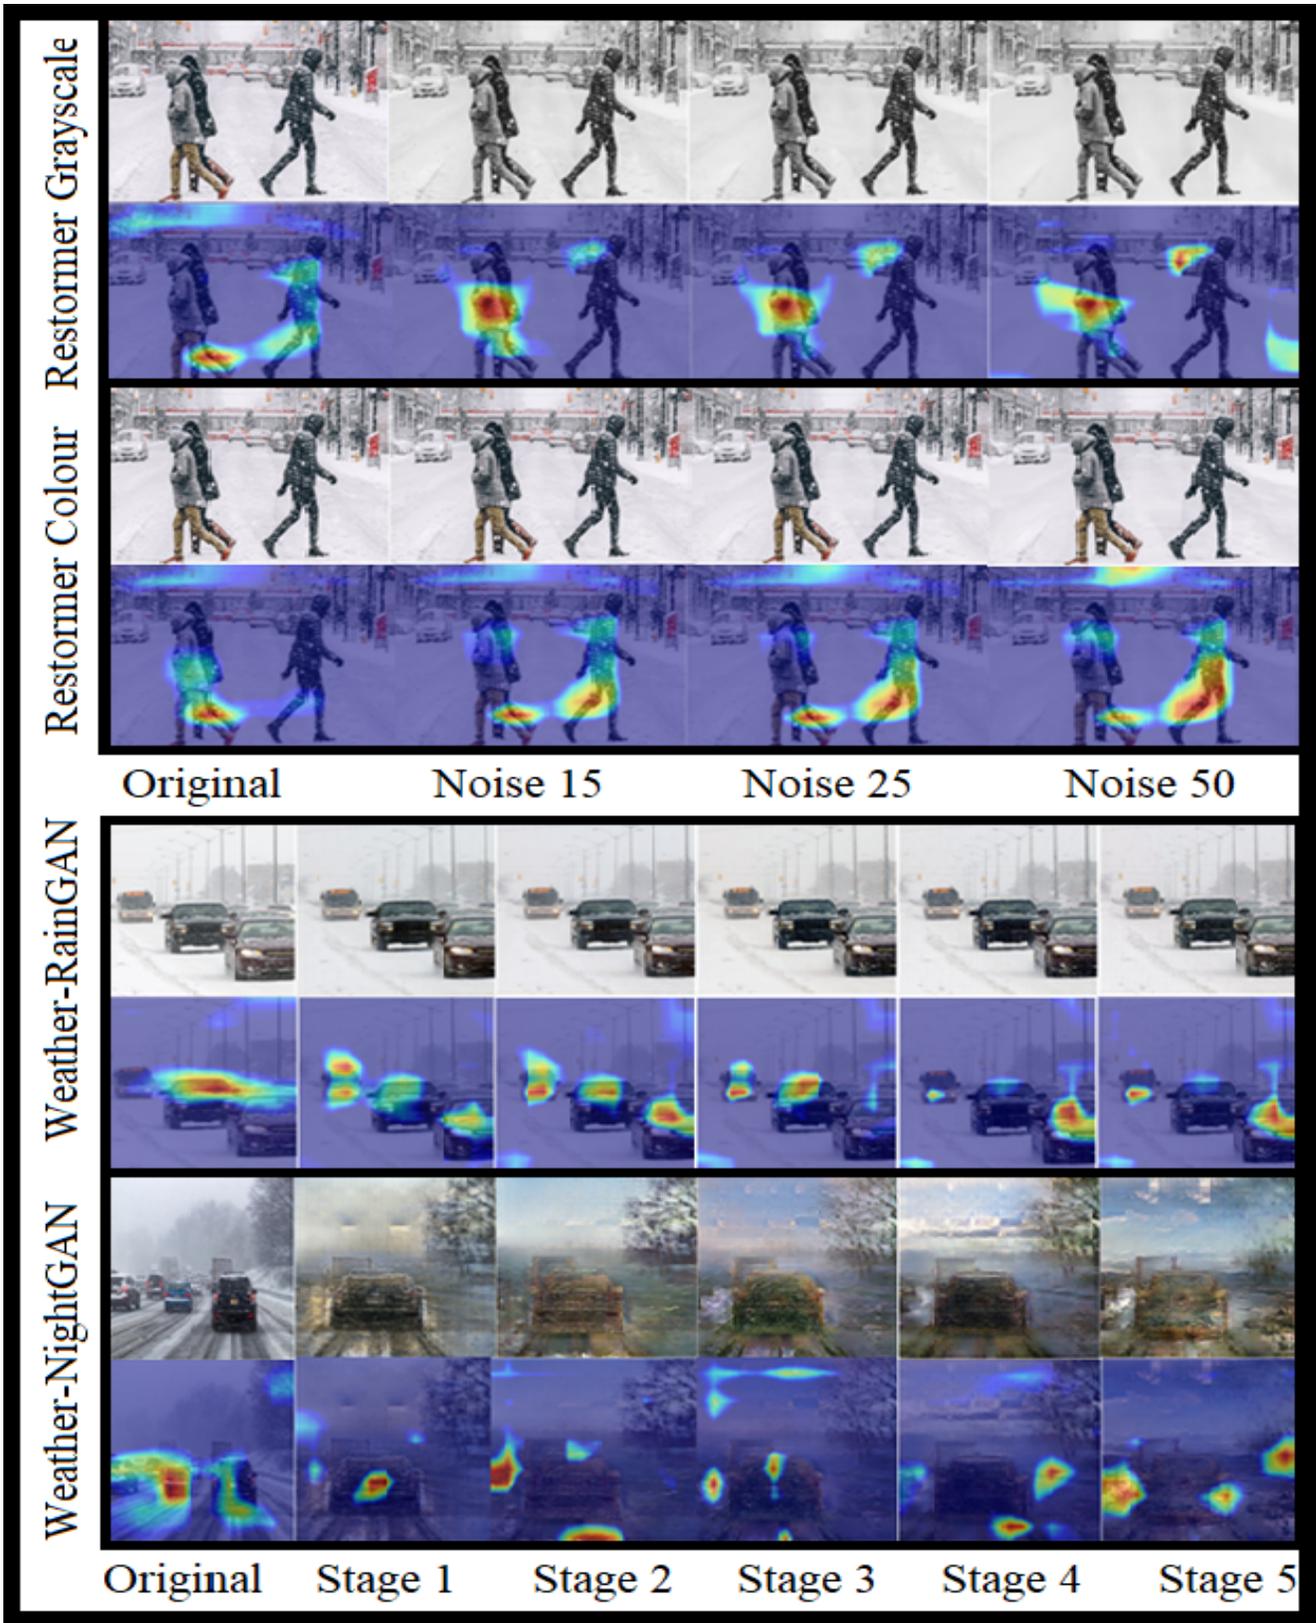


Figure 3. Snowfall weather condition: Grad-CAM Attention Maps for (a) Restormer Grayscale Denoising (b) Restormer Colour Denoising (c) Weather-RainGAN and (d) Weather-NightGAN.

Image Restoration Technique	Class 1 AP [car]	Class 2 AP [bus]	Class 3 AP [person]	Class 4 AP [motorcycle]	mAP
Gaussian Gray Denoising Restormer					
Noise 15	1	5	22	0	6
Noise 25	1	9	23	0	6
Noise 50	1	0	29	0	6
Gaussian Colour Denoising Restormer					
Noise 15	1	3	21	0	5
Noise 25	1	5	22	0	6
Noise 50	1	11	24	0	7
Weather-RainGAN					
Stage 1	1	0	0	0	0
Stage 2	1	17	0	0	4
Stage 3	1	0	0	0	0
Stage 4	1	0	0	76	15
Stage 5	1	0	0	0	0
Weather-NightGAN					
Stage 1	11	0	0	0	2
Stage 2	1	0	0	0	0
Stage 3	1	0	0	0	0
Stage 4	17	0	0	0	3
Stage 5	48	0	0	0	10

Table 1. Object detection performance of all 4 restoration methods as measured at the designed stages.

model. To compare the clarity and detection scores, models trained on different noise levels 15, 25 and 50 are included in testing. The purpose of this experimentation is to identify which noise level is optimal for the object detection task, additionally verified by the explainability setup shown in Figures 2 and 3. The overall mAP of object detection remains constant on all noise levels, however individual class scores are observed to fluctuate. The AP for bus and person class increases from noise 15 to 25 by 4 AP and 1AP respectively. However, the bus AP drops to 0 and person AP boosts to 29 when noise is set to 50. The increase in noise level would correspond to smoother images produced by Restormer, however, it is observed that the bus AP drops when noise level is set to 50. This is very interesting to note and also observable in Figure 2, where the attention maps shift from the bus to the car as the noise level increases.

4.2. Restormer-based Colour Denoising

Section 2 of Table 1 shows the object detection scores of Gaussian colour image denoising using the Restormer model. To compare the clarity and detection scores, models trained on different noise levels 15, 25 and 50 are included in testing. The purpose of this experimentation is to identify which noise level is optimal for the object detection task, additionally verified by the explainability setup. The overall mAP of object detection increases steadily by 1 with all increasing noise levels, and individual class scores are ob-

served to increase as well. The AP for bus and person class increases from noise 15 to 25 by 2 AP and 1AP respectively. The bus AP further increases to 11 and person AP boosts to 24 when noise is set to 50. The increase in noise level is improving the image quality and object detection. As visible in Figure 3, the attention maps on the people are focusing in the region of interest as the noise level increases.

4.3. Weather-RainGAN

The utilization of Pix2Pix GAN [9] for denoising in mapped rain-clear images is proposed and validated through experimentation. The dataset used is Rain 100L [33]. The images in Rain 100 L are originally from BSD 200 dataset [18]. The GAN model is trained with the rainy images as source and clear images as target, expecting this technique to produce denoised images on test images of the DAWN dataset. The intuition behind this idea was the inherent similarity between weather corruptions and synthetic rain streaks, with the goal of the GAN learning to work similarly on the 2 tasks. The training epochs are divided into stages (1,2,3,4,5) to monitor the training progress of the GANs and produce results of OD and explainability at each stage. As expected, the GANs are producing highly unstable behavior with increasing and 0 AP at most of the stages. However, the sudden boost in AP at specific epochs inspired the formulation of our proposed weighted explainability measure. It can be seen in Section 3 of Table 1, that the bus AP is 17

at Stage 2, but 0 at all other stages and the motorcycle AP is 76 at Stage 4 and 0 at other Stages. The car AP however, remains constant at 1 and overall mAP increases over time due to the individual class performance boosts. As these results are vague, we take a closer look at Figures 2 and 3, to determine possible causes for this object detection performance. It can be observed that the GAN is actually performing quite well in denoising weather conditions like fog and snow, and the attention maps are converging towards the relevant objects as the training progresses. But while the object detection results are very disjoint, the attention maps progress continuously and show improvement. The effectiveness of this GAN can be observed as a deraining solution as the image not only gets clearer, but the object detection and attention maps get more precise over training as well.

4.4. Weather-NightGAN

Using PixtoPix GAN [9] for denoising (using mapped rain-clear images) is proposed and validated using the experimental procedure. The dataset used is Transient Attributes dataset [14]. The GAN model is trained with the night images as source and day images as target and is expected to produce denoised images on test images of the DAWN dataset. The intuition behind this idea was the inherent similarity between weather corruptions and dark night images. As expected, the GANs are producing highly unstable behavior with increasing and 0 AP at most of the stages as shown in Table 1. This technique actually worked out only for 1 class (car) and boosted its AP from 11 to 48 over the training period. By taking a closer look at Figures 2 and 3, possible causes for this object detection performance can be determined. The GAN is performing randomly in denoising weather conditions like fog and snow using night training images, but the attention maps are still converging towards the relevant objects as the training progresses. This is exactly aligned with the initial analysis and intuition for the discovery i.e., cases in which object detection is clear to a CV detector, although imperceptible to the human eye. While the object detection results are very disjoint, the attention maps progress continuously and show improvement. The effectiveness of this GAN can be observed in denoising weather conditions for cars, but not for other classes. The purpose of our methodology is not just to observe the restoration progress in a positive light, but to also stop training in case the GAN goes too far. As can be seen in this particular case, the GAN distorts the images after Stage 2, which may confuse detectors when tested against it. In Stage 4, the GAN produces images which are heavily distorted but oddly easier for car detection than even the original image.

5. Conclusion

The goal of our work was to study the relationship between object detection performance and restoration in the context of novel and established restoration techniques. The scope of this study covered 4 restoration techniques, with the goal of denoising images corrupted by over 6 different weather conditions, namely fog, snowstorm, haze, dust tornadoes, rainfall and mist. All 4 techniques worked differently, with the Restormer-based methods providing stable all-rounded results, while the GANs provided class-specific boosts in performance. The overall rise in mAP before and after applying the techniques was 0%, 40%, 275%, 400% respectively. Contrary to popular beliefs, greater denoising does not always guarantee better object detection results as observed in this paper. And conversely, poor denoising does not always guarantee worse object detection results. Particularly for specific tasks like bus and car detection, it can be seen that newer approaches like our proposed method (Weather-RainGAN and Weather-NightGAN) can boost the detector's performance with its resultant images. We present conditional GANs cases that perform superbly on diverse weather conditions ranging from dust tornadoes to snowfall, in spite of having trained on limited single weather conditions.

We present 2 very interesting observations obtained through this study:

1. GANs can generate images complex to the human eye, but comparatively interpretable for vision models post processing. This opens the possibility of exploring the capabilities of vision models beyond the scope of human vision and also warding off potential attacks which can cripple modern detectors.
2. Restoration and denoising methods which produce clearer images (as measured using standard image quality metrics like PSNR), may in fact present a greater challenge for machine perception in object detection.

Understanding how differently humans and detectors perceive information in the same image demands greater exploration of explainability. Countless possibilities arise from these disparate perspectives, for example, if a model sees a pedestrian on a rainy road which a human cannot see, or conversely a pedestrian visible to a passenger's eye which a model cannot capture. Although denoising techniques are a popular choice for tackling corruptions, not all techniques are suitable for all use-cases as seen for car detection. Certain classes respond better to denoising and restoration as perceived by the object detectors. We hope to inspire further research in this direction, for overcoming complex weather corruptions using robust generalizable solutions.

References

- [1] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PloS one*, 10(7):e0130140, 2015. 2
- [2] David Baehrens, Timon Schroeter, Stefan Harmeling, Motoaki Kawanabe, Katja Hansen, and Klaus-Robert Müller. How to explain individual classification decisions. *The Journal of Machine Learning Research*, 11:1803–1831, 2010. 2
- [3] Mario Bijelic, Tobias Gruber, Fahim Mannan, Florian Kraus, Werner Ritter, Klaus Dietmayer, and Felix Heide. Seeing through fog without seeing fog: Deep multimodal sensor fusion in unseen adverse weather. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11682–11692, 2020. 2
- [4] Mario Bijelic, Tobias Gruber, and Werner Ritter. Benchmarking image sensors under adverse weather conditions for autonomous driving. In *2018 IEEE Intelligent Vehicles Symposium (IV)*, pages 1773–1779. IEEE, 2018. 2
- [5] Antoni Buades, Bartomeu Coll, and Jean-Michel Morel. A review of image denoising algorithms, with a new one. *Multiscale modeling & simulation*, 4(2):490–530, 2005. 2
- [6] Pak Hung Chan, Gunwant Dhadyalla, and Valentina Donzella. A framework to analyze noise factors of automotive perception sensors. *IEEE Sensors Letters*, 4(6):1–4, 2020. 2
- [7] Ke Chen, Ryan Oldja, Nikolai Smolyanskiy, Stan Birchfield, Alexander Popov, David Wehr, Ibrahim Eden, and Joachim Pehserl. Mvldarnet: Real-time multi-class scene understanding for autonomous driving using multiple views. In *2020 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 2288–2294. IEEE, 2020. 2
- [8] Alain Hore and Djemel Ziou. Image quality metrics: Psnr vs. ssim. In *2010 20th international conference on pattern recognition*, pages 2366–2369. IEEE, 2010. 1
- [9] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017. 3, 7, 8
- [10] Mourad A Kenk and Mahmoud Hassaballah. Dawn: vehicle detection in adverse weather nature dataset. *arXiv preprint arXiv:2008.05402*, 2020. 3, 4
- [11] Rajiv Khanna, Been Kim, Joydeep Ghosh, and Sanmi Koyejo. Interpreting black box predictions using fisher kernels. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 3382–3390. PMLR, 2019. 2
- [12] Pang Wei Koh and Percy Liang. Understanding black-box predictions via influence functions. In *International conference on machine learning*, pages 1885–1894. PMLR, 2017. 2
- [13] Jason Ku, Melissa Mozifian, Jungwook Lee, Ali Harakeh, and Steven L Waslander. Joint 3d proposal generation and object detection from view aggregation. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 1–8. IEEE, 2018. 2
- [14] Pierre-Yves Laffont, Zhile Ren, Xiaofeng Tao, Chao Qian, and James Hays. Transient attributes for high-level understanding and editing of outdoor scenes. *ACM Transactions on Graphics (proceedings of SIGGRAPH)*, 33(4), 2014. 4, 8
- [15] Peiliang Li, Xiaozhi Chen, and Shaojie Shen. Stereo r-cnn based 3d object detection for autonomous driving. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2019. 2
- [16] Wei Liu, Dragomir Anguelov, Dumitru Erhan, Christian Szegedy, Scott Reed, Cheng-Yang Fu, and Alexander C Berg. Ssd: Single shot multibox detector. In *European conference on computer vision*, pages 21–37. Springer, 2016. 2
- [17] Aboli Marathe, Rahee Walambe, and Ketan Kotecha. Evaluating the performance of ensemble methods and voting strategies for dense 2d pedestrian detection in the wild. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3575–3584, 2021. 2
- [18] David Martin, Charless Fowlkes, Doron Tal, and Jitendra Malik. A database of human segmented natural images and its application to evaluating segmentation algorithms and measuring ecological statistics. In *Proceedings Eighth IEEE International Conference on Computer Vision. ICCV 2001*, volume 2, pages 416–423. IEEE, 2001. 7
- [19] Grégoire Montavon, Wojciech Samek, and Klaus-Robert Müller. Methods for interpreting and understanding deep neural networks. *Digital Signal Processing*, 73:1–15, 2018. 2
- [20] Valentina Muşat, Ivan Fursa, Paul Newman, Fabio Cuzzolin, and Andrew Bradley. Multi-weather city: Adverse weather stacking for autonomous driving. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 2906–2915, 2021. 2
- [21] Anh Nguyen, Alexey Dosovitskiy, Jason Yosinski, Thomas Brox, and Jeff Clune. Synthesizing the preferred inputs for neurons in neural networks via deep generator networks. *Advances in neural information processing systems*, 29, 2016. 2
- [22] Vladislav Ostankovich, Rauf Yagfarov, Maksim Rassabin, and Salimzhan Gafurov. Application of cyclegan-based augmentation for autonomous driving at night. In *2020 International Conference Nonlinearity, Information and Robotics (NIR)*, pages 1–5. IEEE, 2020. 2
- [23] Horia Porav, Tom Bruls, and Paul Newman. I can see clearly now: Image restoration via de-raining. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 7087–7093. IEEE, 2019. 2
- [24] Horia Porav, Will Maddern, and Paul Newman. Adversarial training for adverse conditions: Robust metric localisation using appearance transfer. In *2018 IEEE international conference on robotics and automation (ICRA)*, pages 1011–1018. IEEE, 2018. 2
- [25] Joseph Redmon, Santosh Divvala, Ross Girshick, and Ali Farhadi. You only look once: Unified, real-time object detection. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 779–788, 2016. 2

- [26] Shaoqing Ren, Kaiming He, Ross Girshick, and Jian Sun. Faster r-cnn: Towards real-time object detection with region proposal networks. *Advances in neural information processing systems*, 28, 2015. [2](#)
- [27] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. [2](#)
- [28] Ramprasaath R Selvaraju, Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. Grad-cam: Visual explanations from deep networks via gradient-based localization. In *Proceedings of the IEEE international conference on computer vision*, pages 618–626, 2017. [4](#)
- [29] Avanti Shrikumar, Peyton Greenside, Anna Shcherbina, and Anshul Kundaje. Not just a black box: Learning important features through propagating activation differences. *arXiv preprint arXiv:1605.01713*, 2016. [2](#)
- [30] Karen Simonyan, Andrea Vedaldi, and Andrew Zisserman. Deep inside convolutional networks: Visualising image classification models and saliency maps. *arXiv preprint arXiv:1312.6034*, 2013. [2](#)
- [31] Rahee Walambe, Aboli Marathe, and Ketan Kotecha. Multiscale object detection from drone imagery using ensemble transfer learning. *Drones*, 5(3):66, 2021. [2](#)
- [32] Rahee Walambe, Aboli Marathe, Ketan Kotecha, and George Ghinea. Lightweight object detection ensemble framework for autonomous vehicles in challenging weather conditions. *Computational Intelligence and Neuroscience*, 2021, 2021. [2](#)
- [33] Wenhan Yang, Robby T Tan, Jiashi Feng, Jiaying Liu, Zongming Guo, and Shuicheng Yan. Deep joint rain detection and removal from a single image. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1357–1366, 2017. [4](#), [7](#)
- [34] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. *arXiv preprint arXiv:2111.09881*, 2021. [2](#)
- [35] Syed Waqas Zamir, Aditya Arora, Salman Khan, Munawar Hayat, Fahad Shahbaz Khan, and Ming-Hsuan Yang. Restormer: Efficient transformer for high-resolution image restoration. *arXiv preprint arXiv:2111.09881*, 2021. [4](#)
- [36] Matthew D Zeiler and Rob Fergus. Visualizing and understanding convolutional networks. In *European conference on computer vision*, pages 818–833. Springer, 2014. [2](#)
- [37] Yin Zhou, Pei Sun, Yu Zhang, Dragomir Anguelov, Jiyang Gao, Tom Ouyang, James Guo, Jiquan Ngiam, and Vijay Vasudevan. End-to-end multi-view fusion for 3d object detection in lidar point clouds. In *Conference on Robot Learning*, pages 923–932. PMLR, 2020. [2](#)