Enhancing Object Detection Robustness: A Synthetic and Natural Perturbation Approach

Nilantha Premakumara[*], **Brian Jalaian**[†*], **Niranjan Suri**[*], **and Hooman Samani**[‡]

[*]Institute for Human & Machine Cognition (IHMC), Pensacola, Florida, USA
[†]University of West Florida, Pensacola, Florida, USA
[‡]University of Hertfordshire, Hatfield, UK
{npremakumara, bjalaian, nsuri}@ihmc.org
[†]bjalaian@uwf.edu [‡]h.samani@herts.ac.uk

Abstract

Robustness against real-world distribution shifts is crucial for the successful deployment of object detection models in practical applications. In this paper, we address the problem of assessing and enhancing the robustness of object detection models against natural perturbations, such as varying lighting conditions, blur, and brightness. We analyze four state-of-the-art deep neural network models, Detr-ResNet-101, Detr-ResNet-50, YOLOv4, and YOLOv4-tiny, using the COCO 2017 dataset and ExDark dataset. By simulating synthetic perturbations with the AugLy package, we systematically explore the optimal level of synthetic perturbation required to improve the models' robustness through data augmentation techniques. Our comprehensive ablation study meticulously evaluates the impact of synthetic perturbations on object detection models' performance against real-world distribution shifts, establishing a tangible connection between synthetic augmentation and real-world robustness. Our findings not only substantiate the effectiveness of synthetic perturbations in improving model robustness, but also provide valuable insights for researchers and practitioners in developing more robust and reliable object detection models tailored for real-world applications.

Keywords object detection · synthetic perturbation · natural perturbation · deep neural network model · data augmentation · robustness · real-world distribution shifts · ablation study

1 Introduction

Object detection, a fundamental and critical problem in computer vision, aims to identify and spatially localize objects within images or videos [1]. It is integral to various computer vision tasks, including object tracking [2], activity recognition [3], image captioning [4], image segmentation [5], and visual question answering [6]. Object detection poses a significant challenge due to high intra-class and low inter-class variance [7].

In recent years, deep learning has achieved unprecedented advances in various domains, such as generative modeling [8], computer vision [9], and natural language processing [10]. As deep learning techniques are incorporated into safety-critical application areas like autonomous vehicles, medical diagnostics, and robotics, ensuring the reliability and trustworthiness of these systems is paramount.

Although object detection models have demonstrated remarkable performance, real-world applications do not guarantee high-quality input images. Therefore, evaluating the robustness of these models against distribution shifts is essential before deployment. Numerous robustness interventions have been proposed to improve neural network robustness, including non-conventional architectures[11], adding new images to training data [12], alternative losses [13], and different optimizers [14]. These interventions target specific distribution shifts, such as noises [15] or synthetic corruptions [16].

Our work presents a novel approach to assessing the robustness of object detection models by comparing their performance against synthetic and real natural perturbations. We evaluate four state-of-the-art deep neural network models, Detr-ResNet-101, Detr-ResNet-50, YOLOv4, and YOLOv4-tiny, using the COCO 2017 dataset and the AugLy

augmentation package, which provides synthetic approximations of natural perturbations. Subsequently, we conduct a comprehensive ablation study to perform transfer learning by re-training these models on synthetic natural perturbations and evaluating their robustness against real perturbations using the ExDark dataset.

The primary contributions of this study offer a novel perspective on enhancing the robustness of object detection models against real-world distribution shifts. We meticulously unravel the intricate relationship between synthetic and real-world perturbations through the following innovative approaches:

- We delve into a systematic exploration to experimentally identify the optimal level of synthetic perturbation that effectively enhances the selected models' robustness, shedding light on the potential benefits of synthetic data augmentation for real-world deployment.

- Our comprehensive ablation study meticulously evaluates the impact of synthetic perturbations on object detection models' performance against real-world distribution shifts, establishing a tangible connection between synthetic augmentation and real-world robustness.

Our findings not only substantiate the effectiveness of synthetic perturbations in improving model robustness, but also provide valuable insights for researchers and practitioners in developing more robust and reliable object detection models tailored for real-world applications.

2   Related Work

The robustness of CNN-based object detection techniques in the presence of noise is a critical area of research in computer vision, especially for surveillance applications where image quality can be challenging. Several studies have investigated the impact of noise on object detection algorithms, exploring the effects of synthetic rain, fog, and other degradations on image quality. [17] found that noisy images significantly affect classification tasks, while [18] demonstrated that synthetic rain impacts object detection algorithms similarly to real rain.

[19] proposed a novel approach of adding realistic and varying snow and fog to existing image datasets to investigate their potential for scene reconstruction. [20] evaluated various classifiers trained on ImageNet and found a median 16% decline in classification accuracy when exposed to natural perturbations. They also discovered that natural perturbations cause classification and localization errors, resulting in reduced detection mAP for Faster R-CNN and R-FCN models. [21] assessed the effectiveness of popular deep learning techniques for object detection and recognition in thermal surveillance scenarios using a custom dataset collected in different weather conditions at night.

[22] suggested semantic adversarial editing as a technique to synthesize believable corruptions and highlighted the challenging data points on which their target model should be robust. They showed that their proposed technique can generate a wide range of corruptions and improve model robustness against natural corruptions. The performance of deep neural network-based methods is directly influenced by the availability and quality of datasets. Several challenging large-scale datasets have been introduced, consisting of images or videos captured under adverse conditions. These datasets include ExDARK, UNIRI-TID, RESIDE, UFDD, and See in the Dark [23]. Each dataset focuses on different aspects of challenging conditions, such as low light, various weather conditions, and occlusions. Table 1 summarizes the main advantages and limitations of the mentioned approaches. Recent studies have demonstrated that data augmentation can enhance model resilience [29],[16]. Various augmentations can aid resilience, including different types of noise [30], deep artificial image transformations [12], natural transformations [31], and combinations of simple image transformations such as Python Imaging Library operations [32],[16]. In this research, we used AugLy for image perturbation. AugLy is a state-of-the-art open-source Python library designed to help AI researchers apply data augmentations to assess and improve the robustness of their machine learning models. It integrates multiple modalities, including image, video, and text, which is essential in many AI research domains [33]. AugLy provides over 100 data augmentations that mimic what real users do to photos and videos on platforms like Facebook and Instagram. For our experiment, we used three image augmentations: blur, pixel degradation, and brightness.

In summary, the related work underscores the importance of developing robust deep learning models capable of handling natural data variance and corruptions. Researchers have explored various methods, such as generative models, data augmentation, and synthetic corruption benchmarks, to achieve this goal. The availability of challenging datasets and the application of data augmentation techniques, such as those provided by AugLy, play a crucial role in assessing and enhancing model robustness under adverse conditions. Our research builds upon these findings and aims to further investigate the impact of synthetic natural perturbations on model performance and the potential benefits of transfer learning and retraining on synthetic perturbations to improve robustness against real perturbations.

Table 1: A summary of advantages and limitations of methods tackling object detection in challenging conditions.

Literature	Methods	Advantages	Limitations
[24]	Image are transformed and then fed into the RFB-Net.	Context information fusion allows detection of object in low-light.	Relies on prior information about type of object, shape etc for detecting them in night-time.
[25]	Fusion of pre-trained models using Glue layer and information distillation.	Domain joining with help of glue layer reduces in computation and provides more information for models to learn from different domains.	Relies on prior domain knowledge.
[26]	Generative adversarial network with Faster R-CNN.	Networks learns both day and night-time features.	Relies on prior information of converting night-time images to day time.
[27]	Fully convolutional siamese networks with modified binary segmentation task.	Pre-frame binary segementation mask is used for low-level object representation instead of relying on feature extractor backbone.	Relies on prior information while generating binary segmentation mask. Fails when faced with motion blur and non-object pattern.
[28]	Two stream convolutional neural network with attention mechanism.	Fusion of rgb and thermal image to generate features and noise reduction with convolutional block attention modele.	Thermal images are not efficient for every envirment use.

3 Methodology

3.1 Models, Dataset, and Perturbations

We evaluated model robustness using four pretrained state-of-art neural network models, a test dataset, and synthetic perturbations. The main components of our methodology are as follows:

- **Models:** We used four pretrained neural network models: Detr-ResNet-101, Detr-ResNet-50, Yolov4, and Yolov4-tiny.

- **Dataset:** The COCO 2017 dataset was employed for training purposes and for creating synthetic perturbations. The evaluation of the models' robustness and performance was performed using the ExDARK dataset [34], containing images with real-world perturbations, as part of our ablation study.

- **Synthetic Perturbations:** We introduced three types of robustness specifications (blur, brightness, and pixel degradation) using the AugLy package. Perturbation parameters were adjusted within specific ranges to simulate natural perturbations. To simulate model robustness with natural distribution shifts on the image data, we adopted the perturbation strategies on the dataset. We used perturbations with different levels for each category. Figure 1 depicts how each model's confidence score corresponds to the perturbation level. Example results show that all the model's detection confidence scores decrease when increasing perturbation level compared to the original one.

3.2 Ablation Study Design

Ablation studies are experimental approaches used to assess the importance of individual components in complex systems or models, such as deep learning models, by removing or altering these components [35]. This methodology has been widely adopted in computer vision, particularly for object detection tasks [36], and in other domains such as natural language processing and neuroscience [37].

In this paper, we conduct an ablation study using all four models to evaluate the effectiveness of synthetic perturbations, specifically focusing on the brightness modality, in enhancing the robustness of object detection models against real-world distribution shifts. We test with varied numbers of synthetic perturbation images to determine how the size of the synthetic perturbation training set influences detection performance on natural perturbations. Our ablation study serves two primary purposes:
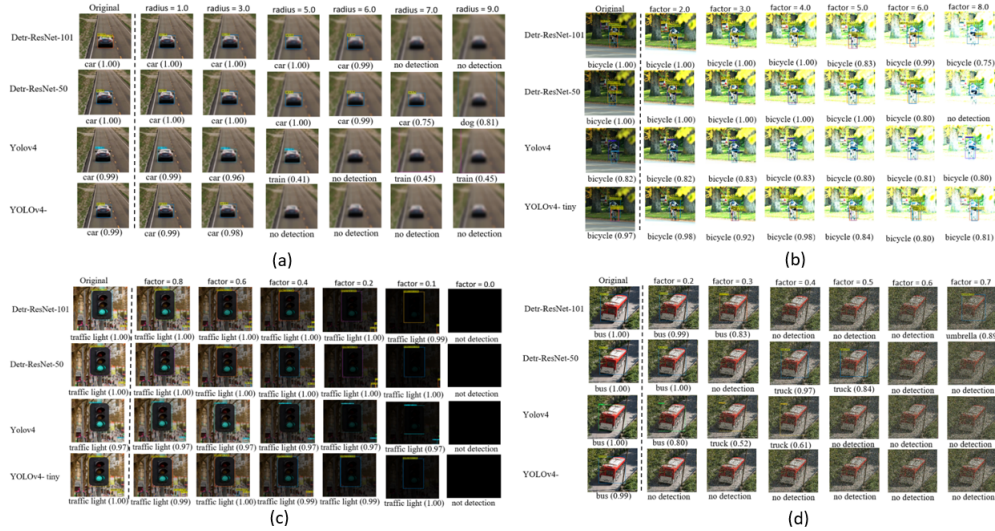
Figure 1: Illustration of how each model is robust against perturbation level (a) blur, (b) brightness (light mood), (c) brightness (dark mood), (d) degrading pixels.

1. **Quantify the impact of synthetic perturbations:** We apply varying levels of synthetic brightness perturbations to identify the optimal level that improves the models' robustness. This process helps us understand the relationship between the perturbation magnitude and the performance of the object detection models.

2. **Analyze the transferability of improvements:** We investigate whether the models' robustness improvements due to synthetic brightness perturbations translate to enhanced performance against real-world distribution shifts, thereby determining if synthetic perturbations serve as reliable proxies for natural perturbations in real-world applications.

Our findings from the ablation study provide valuable insights into developing more robust object detection models and inform future research on the role of synthetic perturbations in enhancing model resilience.

4   Results

4.1   Model Performance with Synthetic Perturbations

We first evaluated the models' performance against synthetic natural perturbations created using the AugLy augmentation package. Table 2 shows the mean average precision (mAP) scores of the four object detection models when tested with these synthetic perturbations. Most models were susceptible to strong brightness, while all models were more robust to darkness than other perturbations.

Table 2: The mean average precision (mAP) scores of various models computed with the original dataset and synthetic corruptions

Model	mAP Original	Blur(radius=5)		Brightness (factor=3)		Brightness (factor=0.2)		Degrading pixels (radius=5)	
		mAP	mAPΔ	mAP	mAPΔ	mAP	mAPΔ	mAP	mAPΔ
Detr-ResNet-101	73.56	67.26	6.3	64.39	9.17	69.02	4.54	64.53	9.03
Detr-ResNet-50	62.37	51.26	11.11	47.89	14.48	59.21	3.16	48.52	13.85
Yolov4	71.49	63.37	8.12	62.44	8.12	68.21	3.28	61.85	9.64
Yolov4-tiny	64.52	53.47	11.05	51.04	11.05	60.83	3.69	50.25	14.27

4

4.2 Comparison of Model Robustness

To provide a comprehensive comparison of the robustness of the four object detection models (Detr-ResNet-101, Detr-ResNet-50, YOLOv4, and YOLOv4-tiny), we analyzed their performance under different perturbation conditions. As shown in Figure 2, Detr-ResNet-101 and Detr-ResNet-50 exhibit similar performance patterns under various perturbations, with Detr-ResNet-101 performing slightly better than Detr-ResNet-50. YOLOv4 outperforms YOLOv4-tiny in all scenarios.

Overall, the Detr-ResNet-101 model demonstrated the best performance in handling synthetic perturbations and real-world perturbations when trained with an augmented synthetically perturbed dataset. However, it should be noted that the choice of the best model for specific scenarios may depend various factors, including the computational power available, the acceptable trade-off between accuracy and speed, and the specific perturbation types encountered in the target domain.
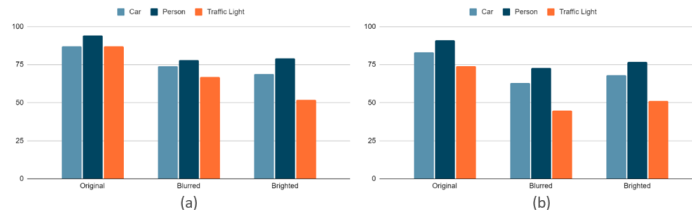

Figure 2: Illustration of performance of multiple object detection models against AugLy blur and brightness perturbation (a) Detr-ResNet-101 model and (b) Detr-ResNet-50 model

4.3 Impact of Perturbation Levels

To investigate the models' performance against varying levels of synthetic perturbation, we conducted experiments by changing the radius and factor values of each synthetic perturbation. This experiment aimed to identify the optimal level of synthetic perturbation that should be used to improve the model's robustness.

In Figure 3, we observed that the models' performance worsened when the blur radius value was greater than 5, the brightness (light mood) factor value was greater than 2, the brightness (dark mood) factor value was less than 0.2, and the degrading pixels factor value was greater than 0.2. Among the models, Detr-ResNet-101 demonstrated higher robustness compared to the other three models.
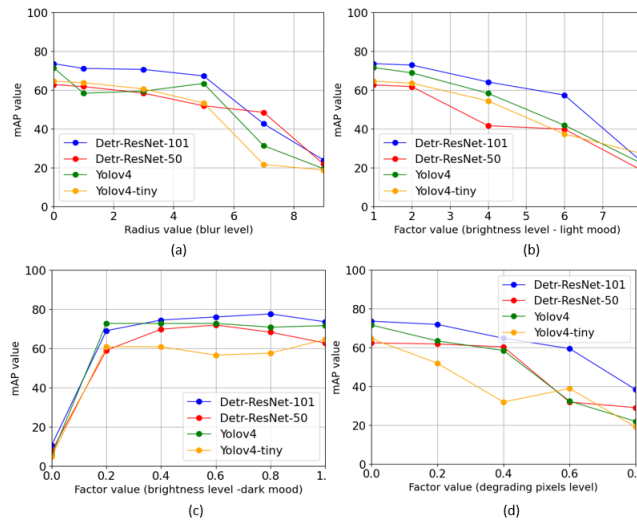

Figure 3: Illustration of model's performance against different synthetic perturbation levels.(a),(b),(c),(d) are represent different synthetic blur level, brightness (light mood) level, synthetic brightness (dark mood) level and synthetic degrading pixels level respectively

5

4.4 Ablation Study: Synthetic to Real Perturbation Robustness

In the ablation study, we examined the performance of the four object detection models (Detr-ResNet-101, Detr-ResNet-50, YOLOv4, and YOLOv4-tiny) against real perturbations from the ExDark dataset, which contains challenging lighting conditions from an out-of-domain setting. The models were re-trained with synthetic perturbations before testing. We conducted further tests with varied numbers of synthetic perturbation images to determine how the size of the synthetic perturbation training set influences detection performance on real perturbations. To test the outcomes, we added synthetic poor brightness perturbations to a subset of the COCO-2017 training dataset at random percentages of 0%, 20%, 50%, and 70% using the AugLy package.

Figure 4 shows the values of the mAP and loss functions for the models with subset ratios of 100:00, 80:20, 50:50, and 30:70, respectively. Our results demonstrate that the mAP has an increasing trend as we increase the percentage of the augmented synthetically perturbed training dataset and re-train the models. The common observation across all cases is that while there is a significant gap between the mAP for the original models, the gap narrows as the percentage of the augmented synthetically perturbed training dataset increases. This suggests that adding synthetic perturbations to the training data improves the models' robustness to real-world perturbations. However, in this study, we focused on the results with brightness adjustments only, as presented in Table 3.
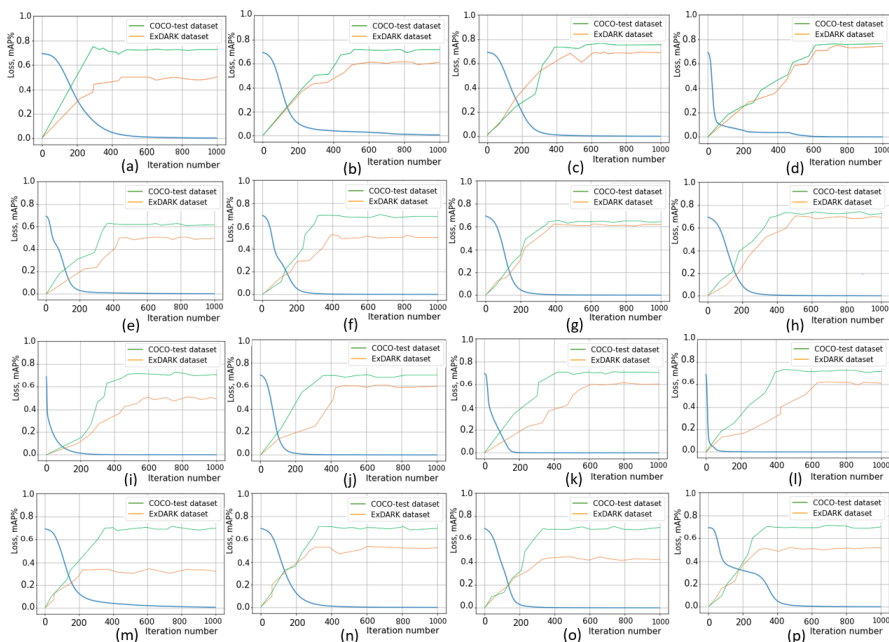

Figure 4: mAP and loss function graphs for different models and synthetic ratio combinations. (a-d) Detr-ResNet-101, (e-h) Detr-ResNet-50, (i-l) YOLOv4, and (m-p) YOLOv4-tiny. Columns represent original to synthetic ratios of 100:0, 80:20, 50:50, and 30:70, respectively.

5 Limitations and Future Work

While our study provides valuable insights into the performance of object detection models under various synthetic perturbations and their robustness against real-world perturbations, there are several limitations that should be acknowledged. First, our analysis was limited to the ExDark dataset, which primarily focuses on poor lighting conditions. To generalize our findings, it is crucial to create and test the models on datasets containing different modalities of natural perturbations, such as occlusions, weather variations, and sensor noise, among others.

Second, our experiments only considered four pre-trained object detection models. To provide a more comprehensive understanding of model robustness, future work could include a wider range of models and architectures, as well as customized models specifically designed to handle perturbations.

Table 3: The mAP scores of various models computed with the original dataset and synthetic corruptions

Model	Original:Synthetic ratio on traing dataset	Model's performance (mAP)	
		COCO-testing	ExDARK
Detr-ResNet-101	100:00	73.56	51.85
	80:20	73.24	61.43
	50:50	75.23	70.56
	30:70	77.92	76.47
Detr-ResNet-50	100:00	62.37	48.28
	80:20	64.76	58.59
	50:50	63.47	62.58
	30:70	67.58	66.89
Yolov4	100:00	71.49	49.48
	80:20	68.78	60.45
	50:50	70.37	61.57
	30:70	72.45	62.45
Yolov4-tiny	100:00	64.52	32.45
	80:20	65.43	38.58
	50:50	67.84	42.59
	30:70	69.51	56.73

Another limitation is that our study focused mainly on brightness adjustments in the ablation study. Expanding the scope of the ablation study to include other types of synthetic perturbations could provide a more complete understanding of model performance and generalization.

Lastly, the choice of augmentation techniques and their parameter settings may also impact the results. Future work could explore alternative data augmentation methods and fine-tune the parameters to find the optimal balance between model robustness and performance.

In summary, future work should aim to:

1. Develop and test object detection models on diverse datasets containing various natural perturbation modalities.

2. Investigate the performance and robustness of a broader range of object detection models and architectures.

3. Expand the ablation study to encompass different types of synthetic perturbations.

4. Explore alternative data augmentation techniques and optimize their parameters to enhance model robustness and performance.

6   Conclusion

In this study, we investigated the robustness of four state-of-the-art object detection models, namely Detr-ResNet-101, Detr-ResNet-50, YOLOv4, and YOLOv4-tiny, to natural perturbations by simulating synthetic perturbations using the AugLy package. We conducted a series of experiments to evaluate the performance of the models against synthetic and real-world perturbations, assess the impact of the size of the synthetic perturbation training set on detection performance under real perturbations, and determine the optimal level of synthetic perturbation for improving model robustness. Ablation studies were also performed to further analyze the models' performance.

Our findings indicate that the performance of object detection models can be significantly improved by augmenting the training dataset with synthetically perturbed images, particularly in scenarios involving challenging lighting conditions, such as those found in the ExDark dataset. Among the models tested, Detr-ResNet-101 exhibited the best robustness to perturbations. However, the choice of the most suitable model may depend on various factors, including computational power, the trade-off between accuracy and speed, and the specific perturbation types encountered in the target domain.

These results provide valuable insights for researchers and practitioners seeking to enhance the robustness of object detection models to natural perturbations and offer a foundation for further research into methods for improving model performance under a wide range of challenging conditions. By addressing the limitations outlined in the previous section and expanding the scope of investigation, future work can contribute to the development of more robust object detection models capable of handling diverse and complex real-world scenarios.

References

[1] P F Felzenszwalb, R B Girshick, D McAllester, and D Ramanan. Object detection with discriminatively trained part-based models. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 32:1627–1645, 9 2010.

[2] Kai Kang, Hongsheng Li, Junjie Yan, Xingyu Zeng, Bin Yang, Tong Xiao, Cong Zhang, Zhe Wang, Ruohui Wang, Xiaogang Wang, and Wanli Ouyang. T-cnn: Tubelets with convolutional neural networks for object detection from videos. *IEEE Transactions on Circuits and Systems for Video Technology*, 28:2896–2907, 10 2018.

[3] N. Vaswani, A. Roy Chowdhury, and R. Chellappa. Activity recognition using the dynamics of the configuration of interacting objects. In *2003 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, 2003. Proceedings.*, volume 2, pages II–633, 2003.

[4] Qi Wu, Chunhua Shen, Peng Wang, Anthony Dick, and Anton van den Hengel. Image captioning and visual question answering based on attributes and external knowledge. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 40:1367–1381, 6 2018.

[5] Bharath Hariharan, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Hypercolumns for object segmentation and fine-grained localization. 11 2014.

[6] Chris Alberti, Jeffrey Ling, Michael Collins, and David Reitter. Fusion of detected objects in text for visual question answering. pages 2131–2140. Association for Computational Linguistics, 2019.

[7] Mark Everingham, Luc Van Gool, Christopher K. I. Williams, John Winn, and Andrew Zisserman. The pascal visual object classes (voc) challenge. *International Journal of Computer Vision*, 88:303–338, 6 2010.

[8] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A. Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. pages 2223–2232, 2017.

[9] Sara Sabour and Nicholas Frosst. Dynamic routing between capsules. pages 3856–3866, 2017.

[10] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. 2018.

[11] Cristina Vasconcelos, Hugo Larochelle, Vincent Dumoulin, Nicolas Le Roux, and Ross Goroshin. An effective anti-aliasing approach for residual networks. 11 2020.

[12] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. 10 2017.

[13] Simon Kornblith, Ting Chen, Honglak Lee, and Mohammad Norouzi. Why do better loss functions lead to less transferable features? 10 2020.

[14] Luke Metz, Niru Maheswaranathan, Jonathon Shlens, Jascha Sohl-Dickstein, and Ekin D. Cubuk. Using learned optimizers to make models robust to input noise. 6 2019.

[15] Yiren Zhou, Sibo Song, and Ngai-Man Cheung. On classification of distorted images with deep convolutional neural networks. 1 2017.

[16] Dan Hendrycks, Norman Mu, Ekin D. Cubuk, Barret Zoph, Justin Gilmer, and Balaji Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. 12 2019.

[17] Samuel Dodge and Lina Karam. Understanding how image quality affects deep neural networks. *IEEE International Conference on Quality of Multimedia Experience, QoMEX 2016*, 6 2016.

[18] Georg Volk, Stefan Muller, Alexander Von Bernuth, Dennis Hospach, and Oliver Bringmann. Towards robust cnn-based object detection through augmentation with synthetic rain variations. 2019.

[19] Alexander von Bernuth, Georg Volk, and Oliver Bringmann. Simulating photo-realistic snow and fog on existing images for enhanced cnn training and evaluation. 2019.

[20] Vaishaal Shankar, Achal Dave, Rebecca Roelofs, Deva Ramanan, Benjamin Recht, and Ludwig Schmidt. A systematic framework for natural perturbations from videos. 7 2019.

[21] Mate Kristo, Marina Ivasic-Kos, and Miran Pobar. Thermal object detection in difficult weather conditions using yolo. *IEEE Access*, 8:125459–125476, 2020.

[22] Rakshith Shetty, Mario Fritz, and Bernt Schiele. Towards automated testing and robustification by semantic adversarial data generation. *European Conference on Computer Vision (ECCV)*, 2020.

[23] Muhammad Ahmed, Khurram Azeem Hashmi, Alain Pagani, Marcus Liwicki, Didier Stricker, and Muhammad Zeshan Afzal. Survey and performance analysis of deep learning based object detection in challenging environments, 8 2021.

[24] Yuxuan Xiao, Aiwen Jiang, Jihua Ye, and Ming-Wen Wang. Making of night vision: Object detection under low-illumination. *IEEE Access*, 8:123075–123086, 2020.

[25] Yukihiro Sasagawa and Hajime Nagahara. Yolo in the dark - domain adaptation method for merging multiple models -. 08 2020.

[26] Kun Wang and Mao Zhen Liu. Object recognition at night scene based on dcgan and faster r-cnn. *IEEE Access*, 8:193168–193182, 2020.

[27] Qiang Wang, Li Zhang, Luca Bertinetto, Weiming Hu, and Philip H. S. Torr. Fast online object tracking and segmentation: A unifying approach, 2019.

[28] Zhengzheng Tu, Yan Ma, Zhun Li, Chenglong Li, Jieming Xu, and Yongtao Liu. Rgbt salient object detection: A large-scale dataset and benchmark, 2022.

[29] Geirhos Robert, Rubisch Patricia, Michaelis Claudio, Bethge Matthias, Wichmann Felix A., and Brendel Wieland. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. 2019.

[30] Raphael Gontijo Lopes, Dong Yin, Ben Poole, Justin Gilmer, and Ekin D. Cubuk. Improving robustness without sacrificing accuracy with patch gaussian augmentation. *arXiv:1906.02611*, 6 2019.

[31] Alfred Laugros, Alfred Laugros@atos Net, Alice Caplier, and Matthieu Ospici. Using synthetic corruptions to measure robustness to natural distribution shifts. 2021.

[32] Ekin D. Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V. Le. Autoaugment: Learning augmentation strategies from data. *https://arxiv.org/abs/1805.09501v3*, 5 2018.

[33] Papakipos Zoe and Bitton Joanna. Augly: Data augmentations for robustness. 1 2022.

[34] Yuen Peng Loh and Chee Seng Chan. Getting to know low-light images with the exclusively dark dataset. *Computer Vision and Image Understanding*, 178:30–42, 2019.

[35] Sina Sheikholeslami, Moritz Meister, Tianze Wang, Amir H Payberah, Vladimir Vlassov, and Jim Dowling. Autoablation: Automated parallel ablation studies for deep learning. In *Proceedings of the 1st Workshop on Machine Learning and Systems*, page 1–7, 2021.

[36] Ross Girshick, Jeff Donahue, Trevor Darrell, and Jitendra Malik. Rich feature hierarchies for accurate object detection and semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, page 580–587, 2014.

[37] Richard Meyes, Melanie Lu, Constantin Waubert de Puiseau, and Tobias Meisen. Ablation studies in artificial neural networks. *arXiv preprint arXiv:1901.08644*, 2019.