https://doi.org/10.2352/EI.2022.34.15.COLOR-156
© 2022, Society for Imaging Science and Technology

Effect of hue shift towards robustness of convolutional neural networks

Kanjar De[a, b] and Marius Pedersen[b]
[a] EISLAB, Luleå University of Technology, 971 87 Luleå, Sweden
[b] Department of Computer Science, Norwegian University of Science and Technology, Teknologiveien 22, 2802 Gjøvik, Norway

Abstract

Computer vision systems become deployed in diverse real time systems hence robustness is a major area of concern. As a vast majority of the AI enabled systems are based on convolutional neural networks based models which use 3-channel RGB images as input. It has been shown that the performance of AI systems, such as those used in classification, is impacted by distortions in the images. To date most work has been carried out on distortions such as noise, blur, compression. However, color related changes to images could also impact the performance. Therefore, the goal of this paper is to study the robustness of these models under different hue shifts.

Introduction

Since the advent of Alexnet [1] in 2012, convolutional neural networks (CNNs) have become the most extensively used tool for computer vision tasks ranging from image classification to segmentation in many different applications, such as medical imaging [2], biometrics [3], and image quality [4]. Innovative network architectures like Resnets [5] using residual blocks and skip connections, Densenets [6] using densely connected blocks, VGGNets [7], GoogleNet [8] using inception module etc. have been deployed in production level computer vision systems. For resource constrained environments like mobile devices light weight architectures like Mobilenets [9, 10] are preferred over large models where there is trade-off between model size and accuracy. With mobile device-based computer vision applications becoming more popular, robustness is also one of the key considerations. Apart from image classification, CNNs have been used successfully in other computer vision tasks like image segmentation [11], object detection [12], image super resolution [13]. CNNs typically use a 3-channel RGB image as input and the CNN architectures are composed of layers of a combination of different types of blocks like convolution blocks, pooling layers, fully connected layers to name a few. To the best of our knowledge the colour sensitivity of these CNN architectures are not deeply studied and is an open area of research. Image quality has an impact on the performance of deep convolutional networks and some of the prior research is in line with this hypothesis [14, 15, 16, 17, 18, 19, 20, 21]. Over the years, the robustness of the deep convolutional neural network has become a topic of active research among researchers [22, 23, 24, 25]. Recently, Taori et al. [26] studied in detail about measuring natural distribution shifts in images and provided a testbed for evaluating robustness. New benchmarks are proposed towards making robust deep neural network models [27]. Recent research [28] suggest that the performance of classic architectures like Resnet can be improved by revisiting augmentations, hyper-parameters, training methodology, etc.

One of the problems encountered by systems deployed using state-of–the-art convolutional neural networks is that if there is a shift in the distribution of the testing data from the trained data the performance of the CNNs drops. Robustness to distribution shift of data is an active area of research and few datasets have been published in this area like Imagenet-C [29], Imagenet-A. [30], Imagenet-R [31] etc. to name a few. Therefore, to the best of our knowledge, the effect of hue shifts on the performance of state-of–the-art convolutional neural networks has not been investigated. Therefore, the goal of this paper is to explore the effect of hue shift on the performance of state of the art convolutional neural networks. An example of effect of a hue shifted image on a CNN based classifier model is shown in Figure 1.

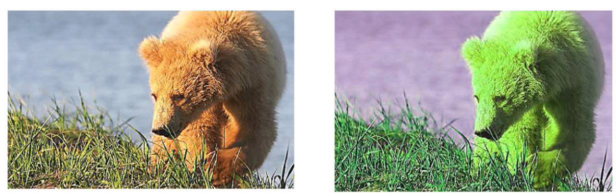

Figure 1: Example of a hue shifted image of a bear being misclassified as a tennis ball by VGG-19 classifier

Investigating this aspect is of importance also in an colour imaging aspect. Gamut mapping algorithms can create hue shifts [32], which in turn could impact CNN based computer vision systems. In addition, other applications use the hue channel such as efficient image hashing [33] and data hiding [34] etc. Recently, Generative Adversarial networks [35] have been used for image generation and image enhancement tasks like image super-resolution [36], image denoising [37], image deblurring [38] , image colourization of gray-scale images [39, 40, 41].

The paper is outlined as follows; first we present the related works, then we present the methods before we present the results and discuss them, at last we conclude.

Related works

One of the earlier works which examined the colour representation of deep convolutional neural networks like VGG-19 and Alexnet were proposed by Engilberge et al. [42] where they introduced hue specificity and colour sensitive units. Their results showed that these units have different hue-specific characteristics, dependent on the layer. The units belonging to the first layers are more sensitive to color, and the later units are more sensitive to the class. Experiments conducted by Buhrmester et al. [43] exploring the effect of colour on image classification yielded some interesting results and certain animal and landscape classes depend on the colour information present in the images. They also showed that the color information in some cases are not only increasing the performance of the CNN, and that it is color space dependent. Recently, De and Pedersen [44] investigated the impact of colour on the robustness of deep neural networks where they have synthetically generated colour distorted images using the publicly available Imagenet dataset [45] and conducted detailed experiments using state-of–the-art convolutional neural networks to study the robustness of these networks. They evaluated Densenet [46], Resnet [5], VGG19 [7], GoogleNet [8], MobileNet [9], and Alexnet [1]. The performance of these networks reduces when information from one or two colour channels in the RGB colour space is removed, and when hue or saturation changes is introduced it also impacts the performance. Geirhos et al. [47] has shown that Imagenet-trained convolutional neural networks are biased towards texture, but the impact of hue colour component in images is not explored in detail. Kantipudi et al. [48] demonstrated that colour channel perturbation attacks on the CNN architectures VGG, Resnet, and Densenet architectures posed security threats.Hosseini and Poovendran [49] have shown adversarial examples can be created by shifting hue and saturation channels in HSV colourspace and the authors have shown that VGG architecture fails terribly for these adversarial examples.

Methods

In this section we explain the details about the experiments conducted in this study.

Data Generation

In this paper, we use the colour images of the validation set of the publicly available Imagenet Database. Removing the grayscale images, we have a total of 49101 colour images and we use the Hue-Saturation-Value (HSV) colour space of these images to synthetically generate hue-shifted test images for our experiments. HSV colour space is a 3-D representation in the form of a hexacone where the intensity is represented by the central vertical axis. Hue is an angle lying in the range 0 to 2π relative to the red axis where red is at angle 0, green at angle $2\pi/3$, blue at angle $4\pi/3$ and finally red again at angle 2π. Saturation measured in the range 0 (center) to 1 (outer boundary) is the radial distance from the central axis to the outer surface. In this paper, we have generated test images where we shift the hue of the images by $\pi/6$ in the range 0 to 2π. [50] Examples of the images used for the experiments are shown in Figure 2.

CNN Architectures

Convolutional neural networks are generally a combination of convolutional blocks, pooling layers and activation functions. For this study we have considered starting from standard architectures like Alexnet, Densenet, VGG, GoogleNet, etc. to the state-of–the-art Efficient Nets, Normaliser free nets for robustness analysis for the task of image classification. Apart from this, different augmentation, training, and scaling strategies are investigated for some of the architectures. Pytorch library is used for all experiments and Imagenet1K pretrained models from Pytorch model zoo (Torchvision 0.2.0) and TIMM library [51] are used for all experiments. Top-1 Accuracy parameter was analysed for a detailed robustness analysis. Augmentation strategies during training CNN like Augmix [52] and Mixup [53] have been shown in the past to have a significant impact on the robustness of the models. EfficientNet V1 [54] are group of models which use compound scaling techniques in depth, width and resolution and inverted residual convolutions (MBconv). For our experiments we use lowest resolution (B0) to highest resolution (B7) models. We conducted further experiments on EfficientNet V1 models which were subjected to adversarial training using Adversarial prop [55] where during training a separate auxiliary batch norm was used for adversarial examples having different underlying distribution which were generated during training. Semi-supervised learning approach involving knowledge distillation using teacher and student networks [56] is known to have an effect on robustness. Teacher networks are trained on labeled images and these teacher networks are used to generate pseudo labels on unlabeled images and these are used to train a student network, JFT-300M dataset has been used some experiments to increase the performance on Imagenet images. EfficientNetV2 [57] are an advanced version of efficientnet models using the concept of Neural Architecture Search [58] (NAS) and new operations like Fused-MBConv which are smaller and can be trained faster and three versions small, medium and large are used for our experiments. Recently, Brock et al. have explored the limitations of batch normalization [59] during training of convolutional neural networks and proposed normalizer free resnets [60] and normalizer free networks [61] and these models have been explored for our experiments.

Performance metric

To assess the performance of the CNN architectures we use classification Top-1 accuracy. We make the prediction using different CNN architectures and compare it to the ground truth. The top score is then calculated as the number of times a predicted label from the CNN matched the ground truth label, and then divided by the number of data points. This is calculated for the baseline (no hue shift) and for the 11 different hue shifts.

Results and Discussion

There is a significant effect of hue shift on the robustness of these models (Tables 1, 2 and 3). The main observation is that

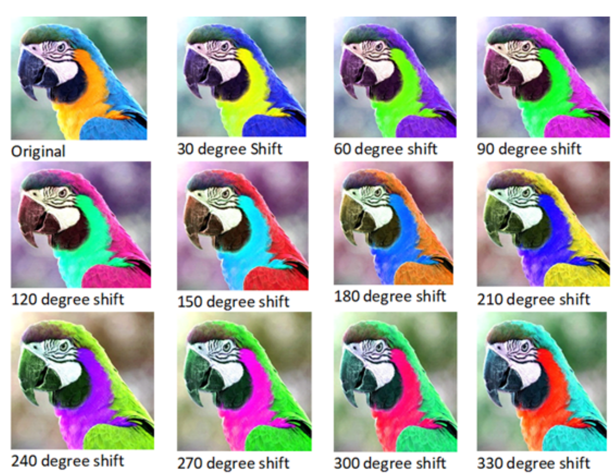

Figure 2: Examples of the generated images

Table 1: Classification Top-1 accuracy (%) of well-known CNN architectures on hue shifted images

Architecture	Hue Shift											
	None	30	60	90	120	150	180	210	240	270	300	330
Alexnet [1]	56.7	44.7	32.6	27.5	27.8	27.9	28.0	27.2	25.9	25.9	32.5	44.5
GoogleNet [8]	69.8	69.3	68.9	66.6	61.5	59.2	58.2	58.6	60.8	66.9	68.8	69.2
VGG-19 [7]	72.4	65.2	55.5	51.9	52.5	51.7	50.8	50.7	51.6	52.3	57.7	65.8
Densenet-161 [46]	77.1	72.3	66.7	64.1	63.1	61.6	61.0	61.2	61.7	62.9	67.0	71.9
Resnet-152 [5]	78.3	73.2	67.1	65.0	64.9	63.8	62.4	62.5	63.3	64.3	67.6	72.9
MobileNet-V2 [9]	71.9	65.1	56.2	53.1	53.4	52.6	51.3	51.8	53.1	54.3	58.5	65.1
EfficientNetV2S [57]	83.8	80.4	77.1	76.4	75.9	74.5	75.9	74.3	74.7	75.7	77.2	80.2
EfficientNetV2M [57]	85.0	81.8	78.9	77.8	77.6	76.9	78.2	76.4	76.8	77.8	79.2	81.7
EfficientNetV2L [57]	85.4	82.6	79.9	78.7	78.5	78.0	78.9	77.7	77.9	78.8	80.0	82.5
NF-Net F0 [61]	83.2	79.7	76.1	74.7	74.3	73.9	75.1	72.9	72.9	74.0	76.0	79.3
NF-Net F1 [61]	84.5	81.0	78.3	77.4	76.8	75.5	76.3	74.8	75.3	76.5	77.9	80.8
NF-Net F2 [61]	84.9	80.9	77.6	76.3	76.2	75.6	76.7	75.3	75.7	76.3	78.0	80.9
NF-Net F3 [61]	85.5	82.1	79.7	78.9	78.5	78.1	78.8	77.3	77.8	79.0	79.8	81.8
NF-Net F4 [61]	85.6	82.5	80.0	79.2	78.8	78.6	79.2	77.9	78.5	79.4	80.2	82.3

Table 2: Classification Top-1 accuracy (%) EfficientNet using different training strategies on hue shifted images

Training strategies	Hue Shift											
	None	30	60	90	120	150	180	210	240	270	300	330
EfficientNet-B0 [54]	76.8	71.5	66.0	65.0	65.1	64.1	65.0	63.5	63.0	64.0	66.5	71.4
EfficientNet-B0 + ap [55]	77.0	73.1	68.4	66.6	66.5	66.7	67.9	66.3	65.6	66.6	69.4	73.2
EfficientNet-B0 + ns [56]	78.6	73.0	66.7	64.0	63.8	63.2	61.8	62.3	63.7	64.5	67.6	73.5
EfficientNet-B1 [54]	78.8	74.8	70.2	68.9	68.6	67.9	68.6	67.6	68.1	68.8	70.7	74.5
EfficientNet-B1 + ap [55]	79.2	75.7	71.5	70.2	70.2	69.8	70.6	69.1	68.5	69.6	71.3	75.4
EfficientNet-B1 + ns [56]	81.4	76.7	70.8	68.7	68.6	67.7	66.2	66.7	67.9	69.0	71.6	76.5
EfficientNet-B2 [54]	80.0	75.6	71.1	70.0	69.9	69.2	69.4	67.9	68.3	69.2	71.1	75.4
EfficientNet-B2 + ap [55]	80.2	77.2	73.4	72.2	71.9	71.3	72.1	70.9	70.6	71.2	73.3	76.8
EfficientNet-B2 + ns [56]	82.4	78.0	72.9	71.0	70.7	70.1	68.5	68.5	69.6	70.4	73.0	77.8
EfficientNet-B3 [54]	81.6	77.8	73.7	72.7	72.8	71.9	72.2	71.0	71.3	72.3	74.3	77.7
EfficientNet-B3 + ap [55]	81.8	79.2	76.3	75.1	75.1	74.7	75.0	74.2	74.6	75.3	76.8	79.2
EfficientNet-B3 + ns [56]	84.0	80.1	75.4	73.9	73.8	72.9	72.0	71.8	72.4	73.0	75.8	80.2
EfficientNet-B4 [54]	83.0	79.5	75.6	74.6	75.0	74.2	74.3	73.4	73.6	74.5	76.1	79.2
EfficientNet-B4 + ap [55]	83.2	81.0	78.7	77.8	77.6	77.0	77.2	76.5	76.6	77.4	78.5	80.9
EfficientNet-B4 + ns [56]	85.1	81.0	77.6	76.6	76.6	76.3	75.8	74.9	75.5	76.3	78.0	81.1
EfficientNet-B5 [54]	83.7	80.3	77.1	76.5	76.4	75.1	75.8	74.9	75.3	76.1	77.5	80.4
EfficientNet-B5 + ap [55]	84.2	82.2	80.1	79.3	79.0	78.1	78.3	77.8	78.3	79.0	80.0	82.0
EfficientNet-B5 + ns [56]	86.0	82.4	79.6	78.9	78.7	78.1	77.6	76.7	77.3	78.1	79.7	82.3
EfficientNet-B6 [54]	84.0	81.2	78.2	77.3	77.2	76.3	76.1	75.6	75.9	76.7	78.2	81.0
EfficientNet-B6 + ap [55]	84.7	82.8	80.8	80.3	80.2	79.4	79.3	79.0	79.6	80.1	80.9	82.6
EfficientNet-B6 + ns [56]	86.4	84.1	82.1	81.7	81.4	80.5	80.6	79.3	80.1	80.9	81.8	84.0
EfficientNet-B7 [54]	84.8	81.9	79.2	78.6	78.6	77.4	78.0	77.0	77.3	78.1	79.3	81.9
EfficientNet-B7 + ap [55]	85.0	83.1	81.1	80.5	80.0	79.5	79.5	79.0	79.4	79.9	81.1	83.0
EfficientNet-B7 + ns [56]	86.8	84.3	82.3	82.1	81.6	80.9	80.1	79.9	80.7	81.3	82.2	84.0

Table 3: Classification Top-1 accuracy (%) of different architectures with Resnet-50

Res-50 Models	Hue Shift											
	None	30	60	90	120	150	180	210	240	270	300	330
Resnet 50 [5]	76.1	69.9	62.5	58.9	59.0	58.1	57.4	57.6	57.8	58.3	62.5	69.4
Resnet 50 + Augmix [52]	77.5	71.7	63.8	60.6	60.2	59.2	59.0	58.7	58.9	60.0	64.3	71.2
Resnet RS-50 [28]	79.8	75.2	70.3	68.5	68.3	67.3	69.2	67.2	66.3	67.7	70.6	75.0
NF-Resnet 50 [60]	80.7	75.9	70.9	69.5	69.4	69.8	72.5	68.6	67.2	68.6	71.5	75.9

when the hue angle is shifted in either direction in the hue circle, there is a decrease in classification accuracies of the pretrained models. This behaviour is demonstrated by all architectures included in our study, and the networks perform the poorest when the hue of the original test images are shifted by 150 to 210 degrees and with higher hue shifts we have poorer classification accuracy. Architectures with higher scale (Efficient net B6,B7, NF-Net F4) are more robust to the hue shift. Training strategies like Adversarial training, Noisy Student also improves the robustness. In addition, efficient augmentation strategies make some impact on hue shifts. GoogleNet has the smallest drop in performance for all architectures. Tables 1-3 show the summary of important results. From Table 2 we infer that the model's robustness to hue shifts increases with the increase in resolution of images (B0 has resolution of 224×224 and B7 has a resolution of 600×600). Further, it is observed that the training procedure also has an impact on the robustness of images that are impacted by hue shifts. Training the same models using Adversarial prop [55] has more robustness to hue shifts. It is also observed that noisy student training improves the robustness of the system but comes with the overhead of extra data requirements. Table 3 gives a summary of the comparison between vanilla Resnet-50 architecture, with Resnet-50 architecture trained augmix (pretrained weights from the authors), Resnet RS-50 which is trained with advanced hyper-parameters, augmentations, etc. and Normaliser free Resnet-50 models and we can see that all these have an impact on performance on hue shifted images and NF-Res50 being more robust to hue shifts. Different experiments conducted during this study suggest that the hue component has an impact on the performance of the convolutional neural networks and not only the architectures but the method in which these networks are trained and augmentations and hyper-parameters during the training stage also have an effect on the robustness of the CNN models.

Conclusion

To the best of our knowledge, very little work has been done to study the effect of hue shift on the robustness of convolutional neural networks based systems (CNNs). Previously, it was explored how different colour distortion affects CNNs, which served as a motivation to explore how something fundamental like hue shifts affects the inference of state-of-the-art deep neural networks. This can in turn motivate future researchers to incorporate colour information to build more robust systems. Colour information in general has not been deeply studied with respect to deep learning models and is an important parameter for consideration for robustness. With the robustness of CNNs to adversarial attacks becoming more important every passing day, the effects of hue must be explored in detail and the experiments presented in this paper are one of the first steps in this direction.

References

[1] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012. 1, 2, 3

[2] Ahmed Mohammed, Ivar Farup, Marius Pedersen, Sule Yildirim, and Øistein Hovde. Ps-devcem: Pathology-sensitive deep learning model for video capsule endoscopy based on weakly labeled data. *Computer Vision and Image Understanding*, 201:103062, 2020. 1

[3] Ramachandra Raghavendra, Kiran B Raja, and Christoph Busch. Contlensnet: Robust iris contact lens detection using deep convolutional neural networks. In *2017 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1160–1167. IEEE, 2017. 1

[4] Seyed Ali Amirshahi, Marius Pedersen, and Stella X Yu. Image quality assessment by comparing cnn features between images. *Journal of Imaging Science and Technology*, 60(6):60410–1, 2016. 1

[5] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016. 1, 2, 3

[6] Simon Jégou, Michal Drozdzal, David Vazquez, Adriana Romero, and Yoshua Bengio. The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition workshops*, pages 11–19, 2017. 1

[7] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 1, 2, 3

[8] C Szegedy, W Liu, Y Jia, P Sermanet, S Reed, D Anguelov, D Erhan, V Vanhoucke, and A Rabinovich. Going deeper with convolutions. In *Proc. of the IEEE conf. on computer vision and pattern recognition*, pages 1–9, 2015. 1, 2, 3

[9] M Sandler, A Howard, M Zhu, A Zhmoginov, and L Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proc. of the IEEE conf. on computer vision and pattern recognition*, pages 4510–4520, 2018. 1, 2, 3

[10] A Howard, M Sandler, G Chu, L Chen, B Chen, M Tan, W Wang, Y Zhu, R Pang, V Vasudevan, et al. Searching for mobilenetv3. In *Proc. of the IEEE/CVF Intl. Conf. on Computer Vision*, pages 1314–1324, 2019. 1

[11] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015. 1

[12] Ross Girshick. Fast r-cnn. In *Proceedings of the IEEE international conference on computer vision*, pages 1440–1448, 2015. 1

[13] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015. 1

[14] S Dodge and L Karam. Understanding how image quality affects deep neural networks. In *2016 eighth Intl. conf. on quality of multimedia experience (QoMEX)*, pages 1–6. IEEE, 2016. 1

[15] S Karahan, M K Yildirum, K Kirtac, F S Rende, G Butun, and H K Ekenel. How image degradations affect deep cnn-based face recognition? In *2016 Intl. Conf. of the Biometrics Special Interest Group (BIOSIG)*, pages 1–5. IEEE, 2016. 1

[16] S Dodge and L Karam. A study and comparison of human and deep learning recognition performance under visual distortions. In *2017 26th Intl. conf. on computer communication and networks (ICCCN)*, pages 1–7. IEEE, 2017. 1

[17] K Grm, V Štruc, A Artiges, M Caron, and H K Ekenel. Strengths and weaknesses of deep learning models for face recognition against image degradations. *IET Biometrics*, 7 (1):81–89, 2017. 1

[18] T Borkar and L Karam. Deepcorrect: Correcting dnn models against image distortions. *IEEE Transactions on Image Processing*, 28(12):6022–6034, 2019. 1

[19] P Roy, S Ghosh, S Bhattacharya, and U Pal. Effects of degradations on deep neural network architectures. *arXiv preprint arXiv:1807.10108*, 2018. 1

[20] Y Zhou, S Song, and N Cheung. On classification of distorted images with deep convolutional neural networks. In *2017 IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1213–1217. IEEE, 2017. 1

[21] S Ghosh, R Shet, P Amon, A Hutter, and A Kaup. Robustness of deep convolutional neural networks for image degradations. In *2018 IEEE Intl. Conf. on Acoustics, Speech and Signal Processing (ICASSP)*, pages 2916–2920. IEEE, 2018. 1

[22] N Papernot, P McDaniel, S Jha, M Fredrikson, Z B Celik, and A Swami. The limitations of deep learning in adversarial settings. In *2016 IEEE European symposium on security and privacy (EuroS&P)*, pages 372–387. IEEE, 2016. 1

[23] N Carlini and D Wagner. Towards evaluating the robustness of neural networks. In *2017 ieee symposium on security and privacy (sp)*, pages 39–57. IEEE, 2017. 1

[24] Hanxun Huang, Yisen Wang, Sarah Monazam Erfani, Quanquan Gu, James Bailey, and Xingjun Ma. Exploring architectural ingredients of adversarially robust deep neural networks. In *Thirty-Fifth Conference on Neural Information Processing Systems*, 2021. 1

[25] A Madry, A Makelov, L Schmidt, D Tsipras, and A Vladu. Towards deep learning models resistant to adversarial attacks. In *Intl. Conf. on Learning Representations*, 2018. 1

[26] Rohan Taori, Achal Dave, Vaishaal Shankar, Nicholas Carlini, Benjamin Recht, and Ludwig Schmidt. Measuring robustness to natural distribution shifts in image classification. *arXiv preprint arXiv:2007.00644*, 2020. 1

[27] Shiyu Tang, Ruihao Gong, Yan Wang, Aishan Liu, Jiakai Wang, Xinyun Chen, Fengwei Yu, Xianglong Liu, Dawn Song, Alan Yuille, et al. Robustart: Benchmarking robustness on architecture design and training techniques. *arXiv preprint arXiv:2109.05211*, 2021. 1

[28] I Bello, W Fedus, X Du, E D Cubuk, A Srinivas, Tsung-Yi Lin, J Shlens, and B Zoph. Revisiting resnets: Improved training and scaling strategies. *arXiv preprint arXiv:2103.07579*, 2021. 1, 3

[29] D Hendrycks and T Dietterich. Benchmarking neural network robustness to common corruptions and perturbations. *Proc. of the Intl. Conf. on Learning Representations*, 2019. 1

[30] D Hendrycks, K Zhao, S Basart, J Steinhardt, and D Song. Natural adversarial examples. *arXiv preprint arXiv:1907.07174*, 2019. 1

[31] D Hendrycks, S Basart, N Mu, S Kadavath, F Wang, E Dorundo, R Desai, T Zhu, S Parajuli, M Guo, et al. The many faces of robustness: A critical analysis of out-of-distribution generalization. *arXiv preprint arXiv:2006.16241*, 2020. 1

[32] Ivar Farup, Carlo Gatta, and Alessandro Rizzi. A multiscale framework for spatial gamut mapping. *IEEE Transactions on Image Processing*, 16(10):2423–2435, 2007. 1

[33] Yan Zhao and Shuai Liu. Robust image hashing based on cool and warm hue and space angle. *Security and Communication Networks*, 2021, 2021. 1

[34] Fatuma Saeid Hassan and Adnan Gutub. Improving data hiding within colour images using hue component of hsv colour space. *CAAI Transactions on Intelligence Technology*, 2021. 1

[35] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. 1

[36] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017. 1

[37] Zailiang Chen, Ziyang Zeng, Hailan Shen, Xianxian Zheng, Peishan Dai, and Pingbo Ouyang. Dn-gan: Denoising generative adversarial networks for speckle noise reduction in optical coherence tomography images. *Biomedical Signal Processing and Control*, 55:101632, 2020. 1

[38] Orest Kupyn, Volodymyr Budzan, Mykola Mykhailych, Dmytro Mishkin, and Jiří Matas. Deblurgan: Blind motion deblurring using conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8183–8192, 2018. 1

[39] R Zhang, P Isola, and A Efros. Colorful image colorization. In *European conf. on computer vision*, pages 649–666. Springer, 2016. 1

[40] S Halder, K De, and P Roy. Perceptual conditional generative adversarial networks for end-to-end image colourization. In *Asian Conf. on Computer Vision*, pages 269–283. Springer, 2018. 1

[41] K Nazeri, E Ng, and M Ebrahimi. Image colorization using generative adversarial networks. In *Intl. conf. on articulated motion and deformable objects*, pages 85–94. Springer, 2018. 1

[42] M Engilberge, E Collins, and S Süsstrunk. Color representation in deep neural networks. In *2017 IEEE Intl. Conf. on Image Processing (ICIP)*, pages 2786–2790. IEEE, 2017. 1

[43] V Buhrmester, D Münch, D Bulatov, and M Arens. Evaluating the impact of color information in deep neural networks. In *Iberian Conf. on Pattern Recognition and Image Analysis*, pages 302–316. Springer, 2019. 1

[44] Kanjar De and Marius Pedersen. Impact of colour on robustness of deep neural networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 21–30, 2021. 2

[45] J Deng, W Dong, R Socher, L Li, K Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conf. on computer vision and pattern recognition*, pages 248–255. Ieee, 2009. 2

[46] G Huang, Z Liu, L Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proc. of the IEEE conf. on computer vision and pattern recognition*, pages 4700–4708, 2017. 2, 3

[47] Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A Wichmann, and Wieland Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. *arXiv preprint arXiv:1811.12231*, 2018. 2

[48] Jayendra Kantipudi, Shiv Ram Dubey, and Soumendu Chakraborty. Color channel perturbation attacks for fooling convolutional neural networks and a defense against such attacks. *IEEE Transactions on Artificial Intelligence*, 1(2): 181–191, 2020. 2

[49] Hossein Hosseini and Radha Poovendran. Semantic adversarial examples. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 1614–1619, 2018. 2

[50] Shamik Sural, Gang Qian, and Sakti Pramanik. Segmentation and histogram generation using the hsv color space for image retrieval. In *Proceedings. International Conference on Image Processing*, volume 2, pages II–II. IEEE, 2002. 2

[51] Ross Wightman. Pytorch image models. `https://github.com/rwightman/pytorch-image-models`, 2019. 2

[52] D Hendrycks, N Mu, ED Cubuk, B Zoph, J Gilmer, and B Lakshminarayanan. Augmix: A simple data processing method to improve robustness and uncertainty. In *Intl. Conf. on Learning Representations*, 2019. 2, 3

[53] H Guo, Y Mao, and R Zhang. Mixup as locally linear out-of-manifold regularization. In *Proc. of the AAAI Conf. on Artificial Intelligence*, volume 33, pages 3714–3722, 2019. 2

[54] M Tan and Q Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *Intl. Conf. on Machine Learning*, pages 6105–6114. PMLR, 2019. 2, 3

[55] C Xie, M Tan, B Gong, J Wang, A Yuille, and Q Le. Adversarial examples improve image recognition. In *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, pages 819–828, 2020. 2, 3, 4

[56] Q Xie, M Luong, E Hovy, and Q Le. Self-training with noisy student improves imagenet classification. In *Proc. of the IEEE/CVF Conf. on Computer Vision and Pattern Recognition*, pages 10687–10698, 2020. 2, 3

[57] M Tan and Q Le. Efficientnetv2: Smaller models and faster training. *arXiv preprint arXiv:2104.00298*, 2021. 2, 3

[58] C Liu, B Zoph, M Neumann, J Shlens, W Hua, L Li, L Fei-Fei, A Yuille, J Huang, and K Murphy. Progressive neural architecture search. In *Proc. of the European conf. on computer vision (ECCV)*, pages 19–34, 2018. 2

[59] S Ioffe and C Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Inl. conf. on machine learning*, pages 448–456. PMLR, 2015. 2

[60] A Brock, S De, and S Smith. Characterizing signal propagation to close the performance gap in unnormalized resnets. *arXiv preprint arXiv:2101.08692*, 2021. 2, 3

[61] A Brock, S De, S Smith, and K Simonyan. High-performance large-scale image recognition without normalization. *arXiv preprint arXiv:2102.06171*, 2021. 2, 3

Author Biography

Kanjar De has completed his PhD from Indian Institute of Information Technology, Design and Manufacturing, Kancheepuram in year 2017 and Masters in Information Technology from International Institute of Information Technology, Bangalore, His research interests include Image Quality Assessment, applied machine learning and colour imaging.

Marius Pedersen received his BsC in Computer Engineering in 2006, and MiT in Media Technology in 2007, both from Gjøvik University College, Norway. He completed a PhD program in color imaging in 2011 from the University of Oslo, Norway, sponsored by Océ. He is a professor of color imaging at the Department of Computer Science at NTNU in Gjøvik, Norway. He is also the director of the Norwegian Colour and Visual Computing Laboratory (Colourlab). His work is centered on subjective and objective image quality.