DOES ENHANCED SHAPE BIAS IMPROVE NEURAL NETWORK ROBUSTNESS TO COMMON CORRUPTIONS?

CHAITHANYA KUMAR MUMMADI*[1,3],
RANJITHA SUBRAMANIAM*[2],
ROBIN HUTMACHER[1],
JULIEN VITAY[2],
VOLKER FISCHER[1],
JAN HENDRIK METZEN[1]

[1]BOSCH CENTER FOR AI (BCAI), GERMANY
[2]TU CHEMNITZ
[3]UNIVERSITY OF FREIBURG
* EQUAL CONTRIBUTION

Introduction

ImageNet trained CNNs are biased towards texture

- Geirhos et al., (2019) verified the performance of CNNs on various image representations.

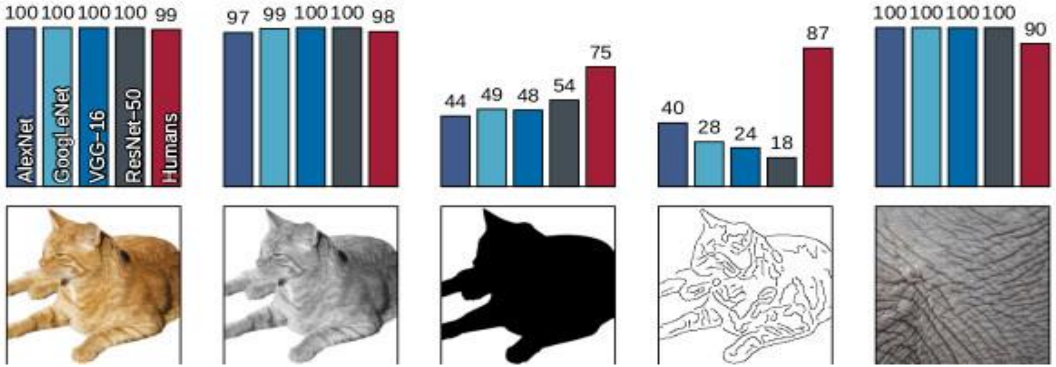

Original Greyscale Silhouette Edges Texture

- Texture-shape cue conflict

a) Texture image b) Content image c) Texture-shape cue conflict

CNN predictions: 81.4% Indian elephant 71.1% Tabby cat 63.9% Indian elephant

CNNs are texture-biased!

Geirhos, R., et al. "ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness." ICLR 2019.

© Robert Bosch GmbH 2021. Alle Rechte vorbehalten, auch bzgl. jeder Verfügung, Verwertung, Reproduktion, Bearbeitung, Weitergabe sowie für den Fall von Schutzrechtsanmeldungen.

BOSCH

Introduction
CNNs trained with stylized ImageNet (SIN)

- Avoid texture cues
- Compelled to make decisions based on shape details.

Original image (IN)

Stylized images generated using various styles

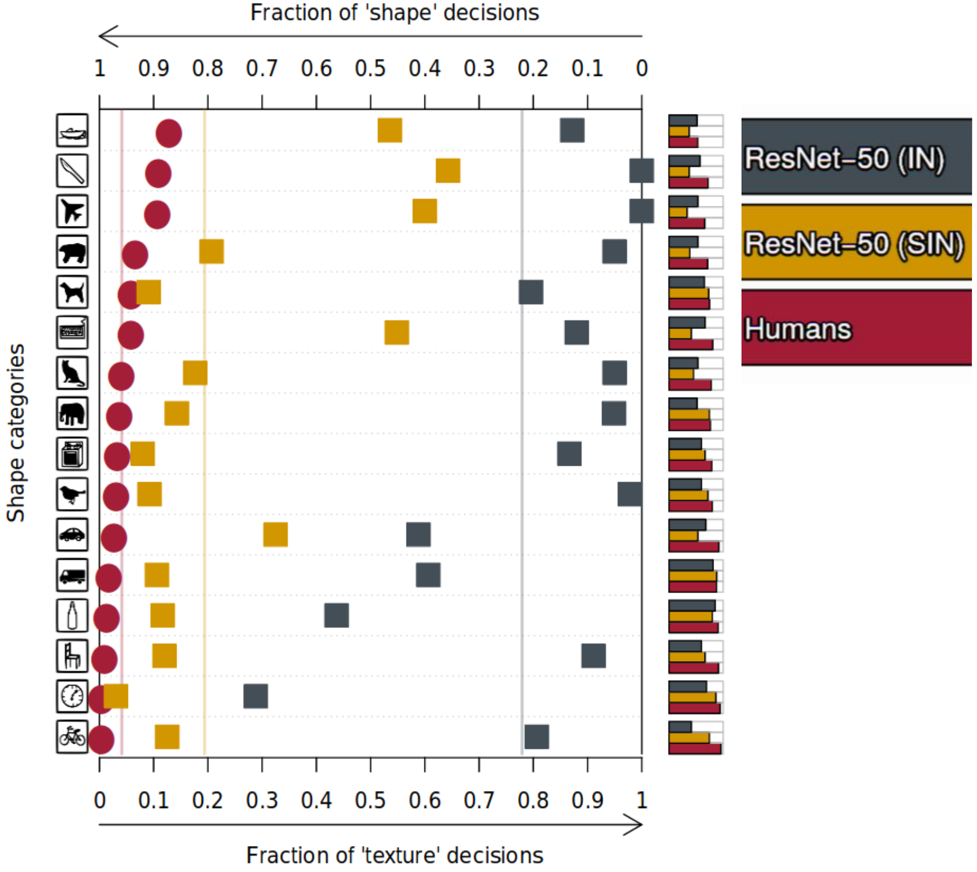

Geirhos, R., et al. "ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness." ICLR 2019.

© Robert Bosch GmbH 2021. Alle Rechte vorbehalten, auch bzgl. jeder Verfügung, Verwertung, Reproduktion, Bearbeitung, Weitergabe sowie für den Fall von Schutzrechtsanmeldungen.

BOSCH

Introduction
Motivation

▶ Robustness against common corruptions with ImageNet-C benchmark.

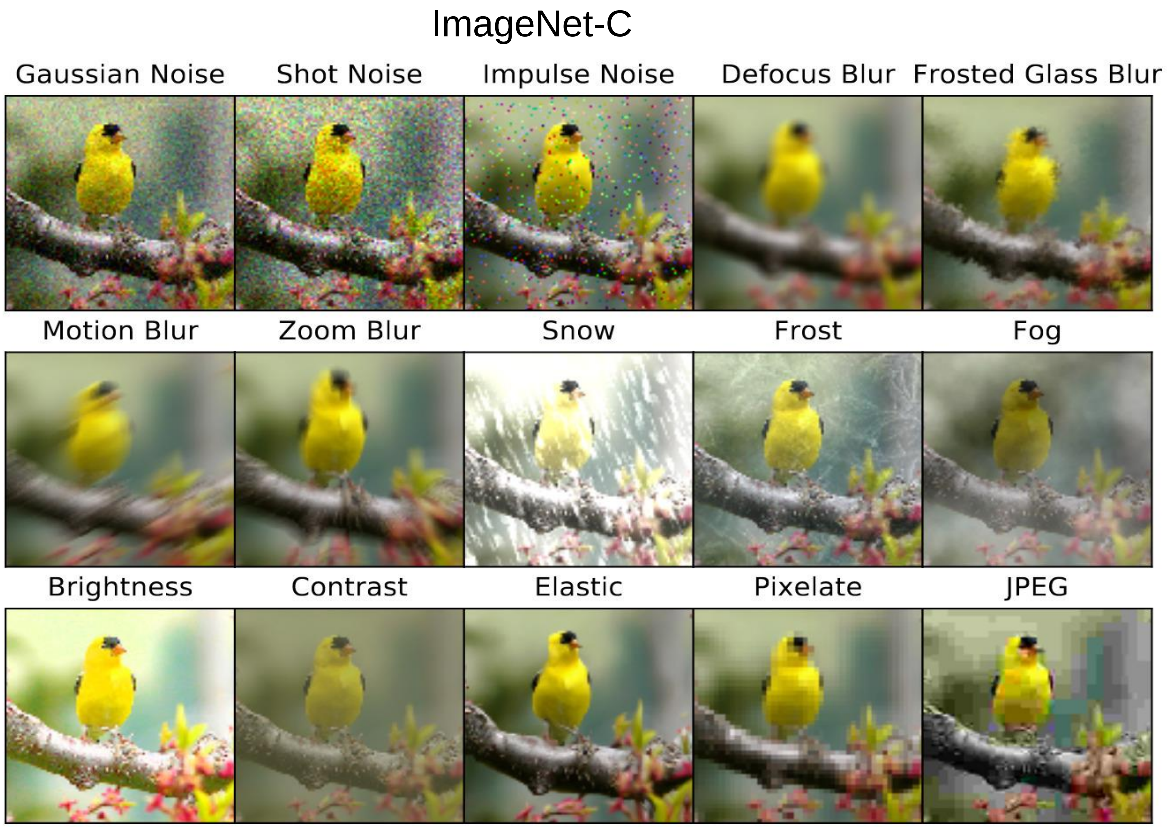

- Evaluation on ImageNet-C distortions

Network	mCE
Standard CNN	76.7
Stylized CNN	69.3

- mCE - mean Corruption Error on 15 corruptions (lower the better)

Hypothesis: Increased shape bias improves corruption robustness

Hendrycks, D., et al. " Benchmarking neural network robustness to common corruptions and perturbations" ICLR, 2019.

© Robert Bosch GmbH 2021. Alle Rechte vorbehalten, auch bzgl. jeder Verfügung, Verwertung, Reproduktion, Bearbeitung, Weitergabe sowie für den Fall von Schutzrechtsanmeldungen.

Introduction
Our contributions

- Enhance CNNs shape bias: use edge maps and randomize style information to reduce texture bias.
- Demonstrate that there is no clear correlation found between shape bias and corruption robustness.
- Study to understand the reasons for improved corruption robustness with Stylized ImageNet.

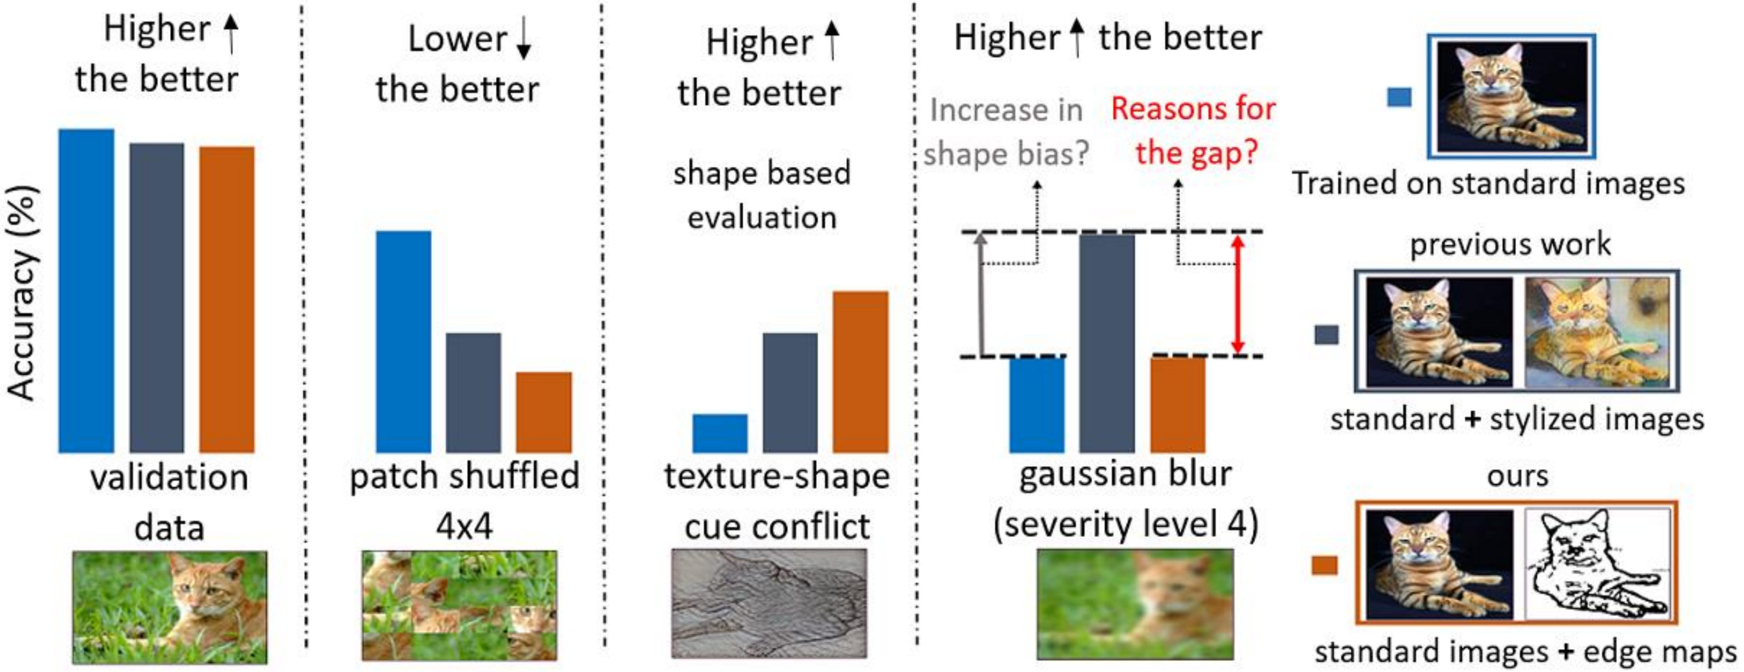

© Robert Bosch GmbH 2021. Alle Rechte vorbehalten, auch bzgl. jeder Verfügung, Verwertung, Reproduktion, Bearbeitung, Weitergabe sowie für den Fall von Schutzrechtsanmeldungen.

Learning shape based representations
Edge maps

- Train a network using standard images along with their edge maps.
  - edge maps are extracted using an existing deep network - Rich Convolutional Features(RCF)
  - enhances shape representations of the network

Standard RGB image

Canny edge map

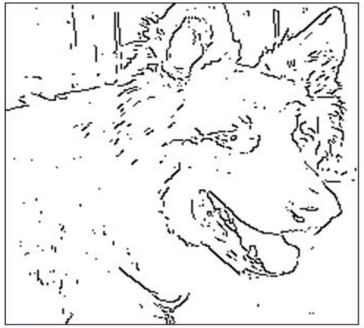

RCF edge map

Liu, Y., et al. "Richer Convolutional Features for Edge Detection." CVPR 2017.

© Robert Bosch GmbH 2021. Alle Rechte vorbehalten, auch bzgl. jeder Verfügung, Verwertung, Reproduktion, Bearbeitung, Weitergabe sowie für den Fall von Schutzrechtsanmeldungen.

Learning shape based representations
Style Randomization (SR)

- Inspiration: A Style transfer technique – Adaptive Instance Normalization (AdaIN)
  - aligns the statistics μ & σ of the content features with those of the style feature.

style y content x stylized

$$AdaIN(x, y) = \sigma(y).\left(\frac{x-\mu(x)}{\sigma(x)}\right) + \mu(x)$$

- Style Randomization: In-network component to reduce texture bias of the network.
  - samples the statistics of feature maps from uniform distribution.

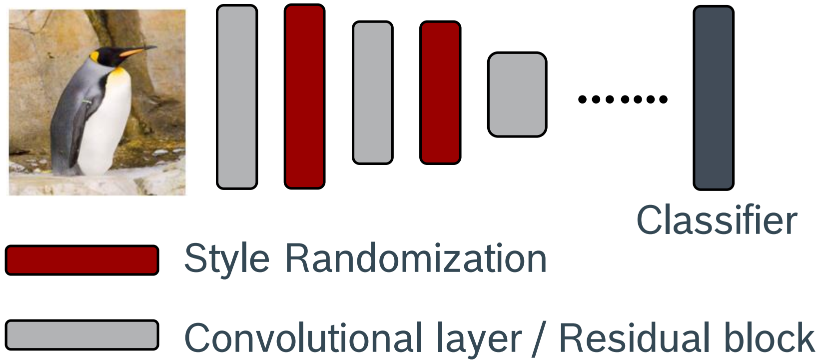

Style Randomization

$$\hat{X}_i = \hat{\sigma}_i.\left(\frac{X_i-\mu_i}{\sigma_i}\right) + \hat{\mu}_i \quad s.t.\ \ \hat{\mu}_i \sim \text{Uniform(-1, 1)},$$

$$\hat{\sigma}_i \sim \text{Uniform}(0.1, 1).$$

$X_i - ith$ feature map

μ_i, σ_i – mean & std of ith feature map

Huang & Belongie, "Arbitrary style transfer in real-time with adaptive instance normalization." ICCV 2017.

© Robert Bosch GmbH 2021. Alle Rechte vorbehalten, auch bzgl. jeder Verfügung, Verwertung, Reproduktion, Bearbeitung, Weitergabe sowie für den Fall von Schutzrechtsanmeldungen.

BOSCH

Dataset
ImageNet20

Edible Items

Mushroom

Bell Pepper

Pretzel

Automobile

Sports Car

Trolley Bus

Life Boat

Animals / Birds / Insects

African Elephant

German Shepherd

Tabby Cat

Arabian Camel

Tailed Frog

Scorpion

King Penguin

Albatross

Fly

Sulphur Butterfly

Man-made objects

Tea Pot

Stop Watch

Teddy Bear

Fur Coat

Architecture used: ResNet18

Kindly refer our paper for more results on ImageNet200, ResNet50, DenseNet121, MobileNetV2.

© Robert Bosch GmbH 2021. Alle Rechte vorbehalten, auch bzgl. jeder Verfügung, Verwertung, Reproduktion, Bearbeitung, Weitergabe sowie für den Fall von Schutzrechtsanmeldungen.

BOSCH

Evaluation of shape bias

- Shuffled patches evaluation on 653 validation images that are correctly classified by all networks.

IN – Network trained on ImageNet
SIN – Network trained on ImageNet + Stylized ImageNet (previous work)
E – Network trained on ImageNet + Edge Maps
E-SIN – Network trained on ImageNet + Edge Maps + Stylized ImageNet

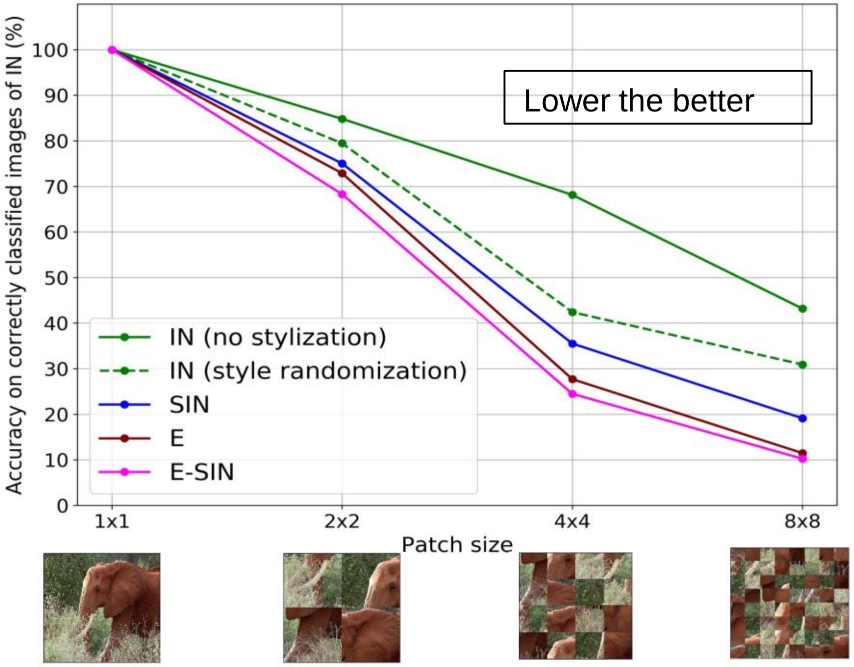

- Texture-shape cue conflict
  - #400 images – whose shape label belongs to ImageNet20 class
  - #100 images – with texture and shape labels from ImageNet20

Networks	#400 images	#100 images	
	Shape results	Texture results	Shape results
IN (no stylization)	68	34	15
IN (style randomization)	86	20	18
SIN (previous work)	156	**2**	32
E (ours)	193	15	46
E-SIN (ours)	**234**	6	**58**

Shape label : Car Cat Dog Elephant
Texture label : Cat Elephant Car Dog

© Robert Bosch GmbH 2021. Alle Rechte vorbehalten, auch bzgl. jeder Verfügung, Verwertung, Reproduktion, Bearbeitung, Weitergabe sowie für den Fall von Schutzrechtsanmeldungen.
BOSCH

Evaluation on ImageNet-C corruptions

- Enhanced shape bias doesn't improve the robustness towards common corruptions.

IN – Network trained on ImageNet
SIN – Network trained on ImageNet + Stylized ImageNet (previous work)
E – Network trained on ImageNet + Edge Maps
E-SIN – Network trained on ImageNet + Edge Maps + Stylized ImageNet

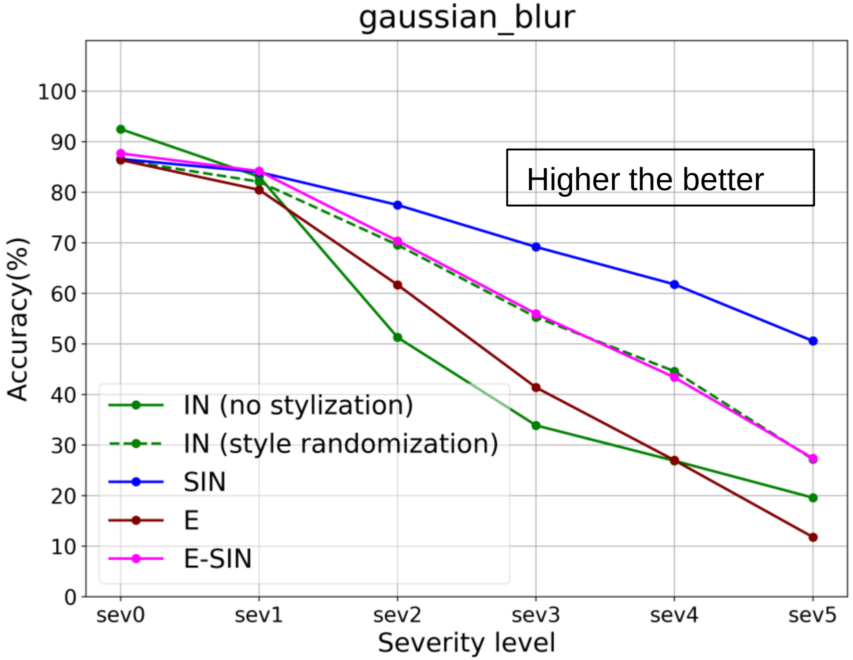

Sample ImageNet-C corruptions

Networks	Average corruption accuracies (%)				mCA (%) (Higher the better)
	Noise	Blur	Weather	Digital	
IN (no stylization)	50.4	57.81	69.25	76.03	64.69
IN (style randomization)	59.57	62.7	72.4	78.85	69.39
SIN (previous work)	**74.54**	**72.67**	**78.36**	**83.05**	**77.64**
E (ours)	60.65	48.07	66.09	71.54	62.01
E-SIN (ours)	72.03	58.31	73.51	80.67	71.55

- mCA - mean Corruption Accuracy across 15 corruptions

© Robert Bosch GmbH 2021. Alle Rechte vorbehalten, auch bzgl. jeder Verfügung, Verwertung, Reproduktion, Bearbeitung, Weitergabe sowie für den Fall von Schutzrechtsanmeldungen.

BOSCH

Study on corruption robustness
Stylization variants

- Study contribution of different factors to understand their influence on corruption robustness.
  - **Role of shape bias**
  - **Role of data augmentation via stylization**
  - **Role of style distribution**
  - **Role of preserved image statistics**

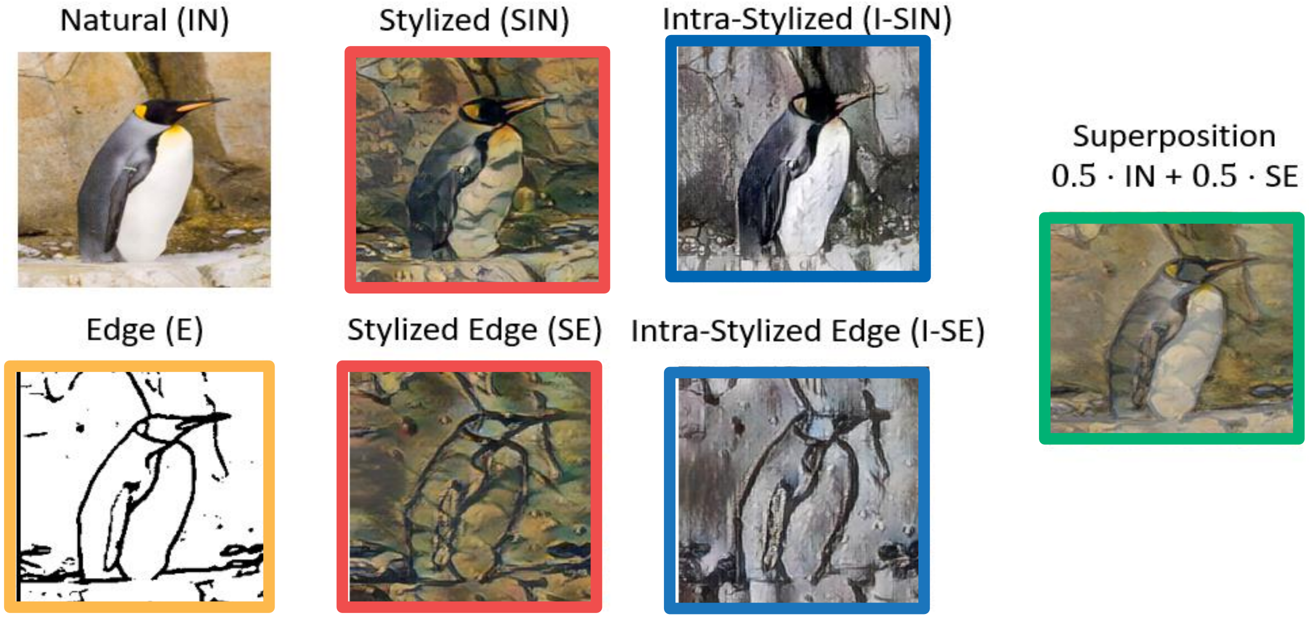

© Robert Bosch GmbH 2021. Alle Rechte vorbehalten, auch bzgl. jeder Verfügung, Verwertung, Reproduktion, Bearbeitung, Weitergabe sowie für den Fall von Schutzrechtsanmeldungen.

Study on corruption robustness
Stylization variants – Evaluation on ImageNet-C

Networks	Mean Corruption Accuracy (%)
IN	64.69
E	62.01
SIN	77.64
SE	71.81
I-SIN	77.92
I-SE	70.3
Superposition	**78.96**

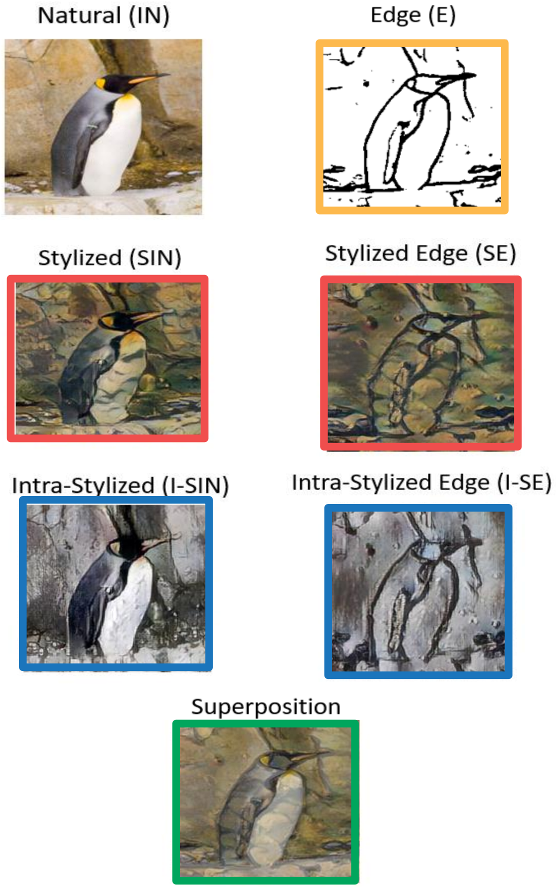

© Robert Bosch GmbH 2021. Alle Rechte vorbehalten, auch bzgl. jeder Verfügung, Verwertung, Reproduktion, Bearbeitung, Weitergabe sowie für den Fall von Schutzrechtsanmeldungen.

Study on corruption robustness
Stylization variants – Final Evaluation

Networks	Input Image Composition			#100 Cue conflict images		Mean Corruption Accuracy (%)
	Natural image	Edge Map	Style transfer	Shape results	Texture results	
IN	□	✗	✗	11	39	64.69
SIN	□	✗	□	34	**2**	77.64
E	✗	□	✗	46	15	62.01
SE	✗	□	□	**55**	6	71.81
Superposition	□	□	□	22	13	**78.96**

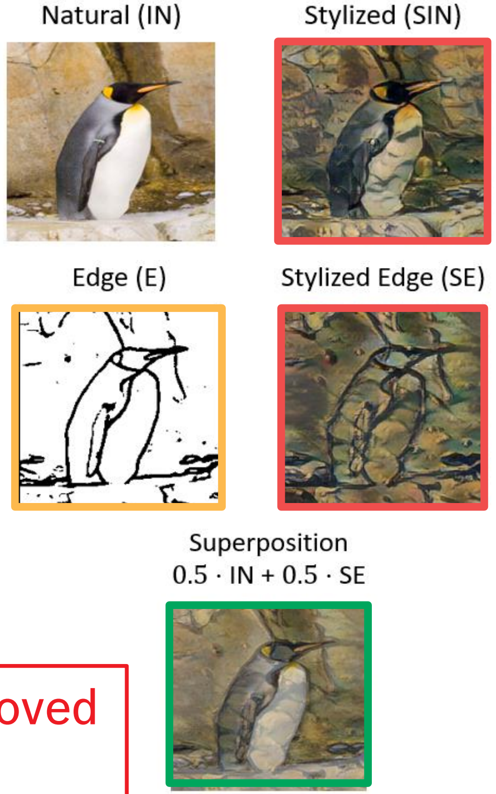

Data augmentation of natural images via stylization is the reason for improved corruption robustness and not the shape bias!

© Robert Bosch GmbH 2021. Alle Rechte vorbehalten, auch bzgl. jeder Verfügung, Verwertung, Reproduktion, Bearbeitung, Weitergabe sowie für den Fall von Schutzrechtsanmeldungen.

BOSCH

Conclusion

- Improved the shape bias of a CNN using
  - Edge maps.
  - Style randomization.

- Shown that shape bias has no influence on corruption robustness.

- Performed a detailed analysis on the improved corruption accuracies on Stylized images.

Take away message: Data augmentation of natural images via stylization is the reason for improved corruption robustness and increased shape bias is only a byproduct.

© Robert Bosch GmbH 2021. Alle Rechte vorbehalten, auch bzgl. jeder Verfügung, Verwertung, Reproduktion, Bearbeitung, Weitergabe sowie für den Fall von Schutzrechtsanmeldungen.

BOSCH