

Supplementary Material for the Paper “Model Vulnerability to Distributional Shifts over Image Transformation Sets”

Riccardo Volpi¹, Vittorio Murino^{1,2}

¹Istituto Italiano di Tecnologia ²Università di Verona

{riccardo.volpi,vittorio.murino}@iit.it

1. Transformation Sets

Transformation Sets \mathbb{T} for the Experiments						
Transformations	Range	No. Levels	Experiment			
			MNIST [7]	CIFAR-10 [6]	CamVid [5]	Faces [1]
<i>Autocontrast</i>	[0.0, 0.3]	20	✓	✓	✓	✓
<i>Brightness</i>	[0.6, 1.4]	20	✓			
	[0.8, 1.2]	20		✓	✓	✓
<i>Color</i>	[0.6, 1.4]	20	✓	✓	✓	✓
<i>Contrast</i>	[0.6, 1.4]	20	✓	✓	✓	✓
<i>Sharpness</i>	[0.6, 1.4]	20	✓	✓	✓	✓
<i>Solarize</i>	[0.0, 20.0]	20	✓			
<i>Grayscale</i>	—	1	✓			✓
<i>R-channel enhancer</i>	[−120, 120]	30	✓		✓	✓
	[−30, 30]	30		✓		
<i>G-channel enhancer</i>	[−120, 120]	30	✓		✓	✓
	[−30, 30]	30		✓		
<i>B-channel enhancer</i>	[−120, 120]	30	✓		✓	✓
	[−30, 30]	30		✓		
Size of \mathbb{T}_N			211^N	190^N	190^N	191^N

Table 1. List of transformations used.

In this Section, we report the image transformations briefly introduced in Section 2.1, and used throughout Sections 3.3 and 5. Table 1 reports them (column 1), with the range of magnitude values (column 2) and the number of values in which the ranges have been discretized (column 3). Columns 4 – 7 indicate whether a transformation is used (✓) in the experiments of Sections 5.1, 5.2, 5.3 and 5.4, respectively. Grayscale conversion (“*Grayscale*” row) has only one magnitude value. For all the transformations, excluding the *R/G/B-channel enhancement* operations, the reader can refer to the PIL library [4], and in particular to the modules [3, 2]. We summarize them in the following, reporting the core descriptions.

- **Autocontrast:** “Maximize (normalize) image contrast. This function calculates a histogram of the input image, removes cutoff percent of the lightest and darkest pixels from the histogram, and remaps the image so that the darkest pixel becomes black (0), and the lightest becomes white (255)” [3].
- **Brightness:** “Adjust image brightness. [...] An enhancement factor of 0.0 gives a black image. A factor of 1.0 gives the original image” [2].
- **Color:** “Adjust image color balance. [...] An enhancement factor of 0.0 gives a black and white image. A factor of 1.0 gives the original image” [2].

- **Contrast:** “Adjust image contrast. [...] An enhancement factor of 0.0 gives a solid grey image. A factor of 1.0 gives the original image” [2].
- **Sharpness:** “ Adjust image sharpness. [...] An enhancement factor of 0.0 gives a blurred image, a factor of 1.0 gives the original image, and a factor of 2.0 gives a sharpened image” [2].
- **Solarize:** “Invert all pixel values above a threshold” [3].
- **Grayscale:** “Convert the image to grayscale” [3]. We treat the output as an RGB image by replicating it in three different channels.
- **R-channel enhancer:** add a value to the R-channel of all pixels.
- **G-channel enhancer:** add a value to the G-channel of all pixels
- **B-channel enhancer:** add a value to the B-channel of all pixels

References

- [1] Face recognition. https://github.com/ageitgey/face_recognition. 1
- [2] Pillow imageenhance module. <https://pillow.readthedocs.io/en/3.0.x/reference/ImageEnhance.html>. 1, 2
- [3] Pillow imageops module. <https://pillow.readthedocs.io/en/3.0.x/reference/ImageOps.html>. 1, 2
- [4] Python imaging library. <https://github.com/python-pillow/Pillow>. 1
- [5] G. J. Brostow, J. Shotton, J. Fauqueur, and R. Cipolla. Segmentation and recognition using structure from motion point clouds. In *ECCV (I)*, pages 44–57, 2008. 1
- [6] A. Krizhevsky and G. Hinton. Learning multiple layers of features from tiny images. *Master’s thesis, Department of Computer Science, University of Toronto*, 2009. 1
- [7] Y. Lecun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. In *Proceedings of the IEEE*, pages 2278–2324, 1998. 1