Color Equivariant Network

Felix O'Mahony [1], Yulong Yang [2], Christine Allen-Blanchette [2]
[1] Oxford University & [2] Princeton University

Equivariance for robustness to color variations

Group equivariant networks have seen limited application in the context of perceptual quantities such as hue and saturation, even though their variation can lead to significant reductions in classification performance.

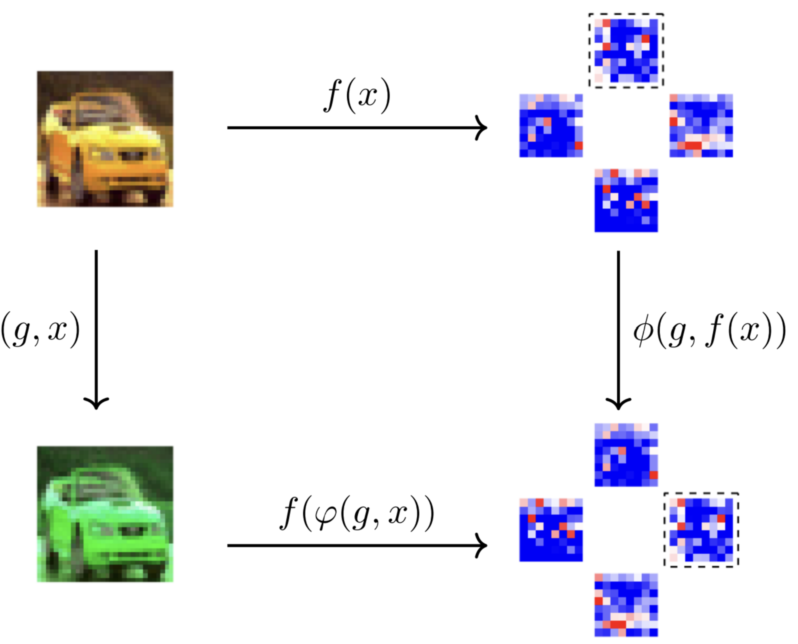

Figure 1. Equivariance of our hue equivariant model.

Our key contributions are:

- Networks that are equivariant to shifts in hue and saturation.
- Networks that perform on par or better than conventional architectures in the presence of unseen color combinations.
- Networks produce interpretable feature representations.

Hue and saturation group

Hue group and group action. In the HSL color space hue can be identified with the 2D rotation group. we identify elements of the discretized hue group, H_N, with those of the cyclic group C_N. For image $x = (x_h, x_s, x_l)$, the action of an element h_i of the hue group H_N is given by

$$\varphi_h(h_i, x) = ((x_h + h_i)(\text{mod } 255), x_s, x_l)$$

An element of the hue group acts on a function $f = (f_1, f_2, \dots, f_N)$ on the discrete hue group by the group action

$$\phi_h(h_i, f) = (f_{(1+i)(\text{mod } N)}, \dots, f_{(N+i)(\text{mod } N)})$$

Saturation group and group action. We use the group structure of $(\mathbb{R}, +)$ and consider a finite subset of the group. For image x, the action of an element s_i of the saturation group S_N is given by

$$\varphi_s(s_i, x) = (x_h, \min(x_s + s_i, 255), x_l)$$

An element of the saturation group acts on a function f on the discrete saturation group by the group action

$$\phi_s(s_i, f) = (f_{(1+i)}, \dots, f_N, 0, \dots, 0)$$

Hue-Saturation group and group action. For image x, the action of an element (h_i, s_j) of the hue-saturation group $H_N \times S_N$ is given by

$$\varphi_{hs}\left((h_i, s_j), x\right) = \varphi_h\left(h_i, \varphi_s(s_j, x)\right)$$

An element of the hue-saturation group acts on a function $f = (f_{11}, \dots, f_{1M}, f_{21}, \dots f_{NM})$ on discrete hue-saturation group by group action

$$\phi_{hs}\left((h_i, s_j), f\right) = \phi_h\left(h_i, \phi_s(s_j, f)\right)$$

Lifting layer. Lifting to the group in our implementation is achieved by transforming the images rather than the filters

$$f^1(g_{x,j}) = \varphi_h(h_j, f^0)(x)$$

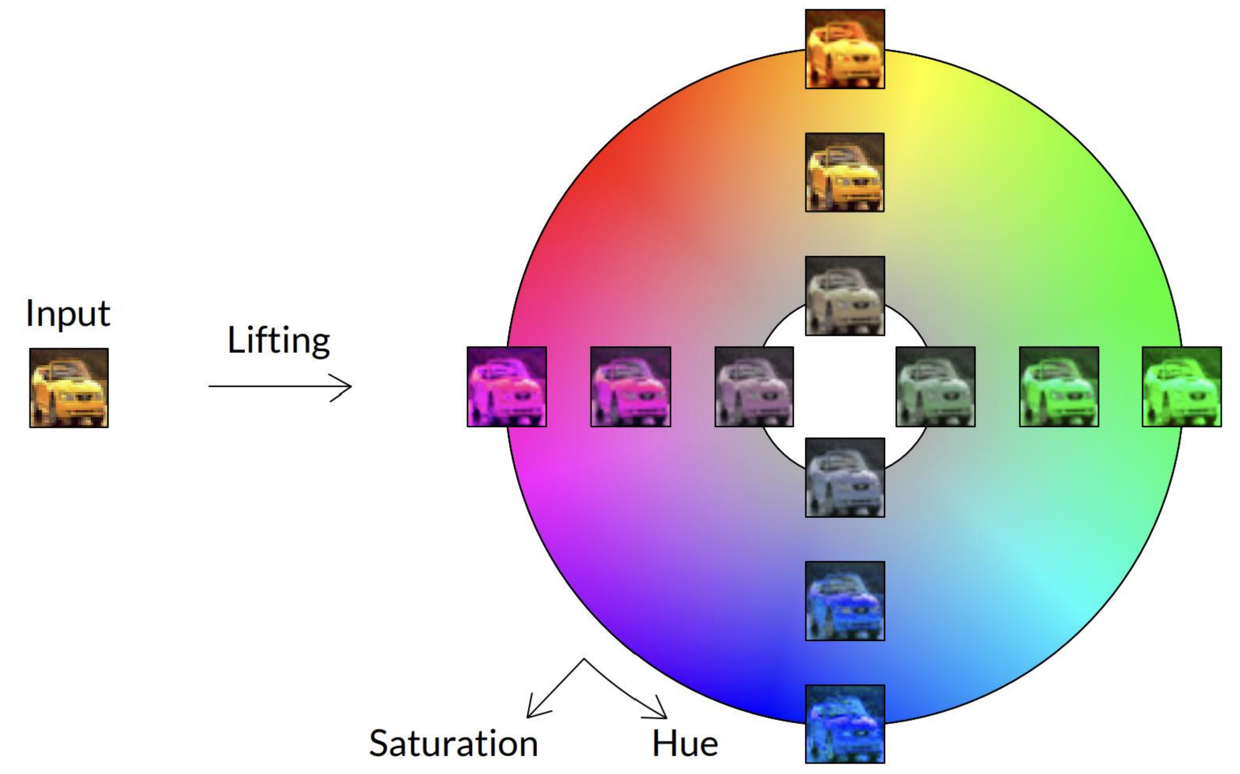

Figure 2. Lifting layer. An input image (left) is "lifted" to the hue-saturation group (right) by shifting its hue and saturation values.

Hue MNIST: Global hue equivariance

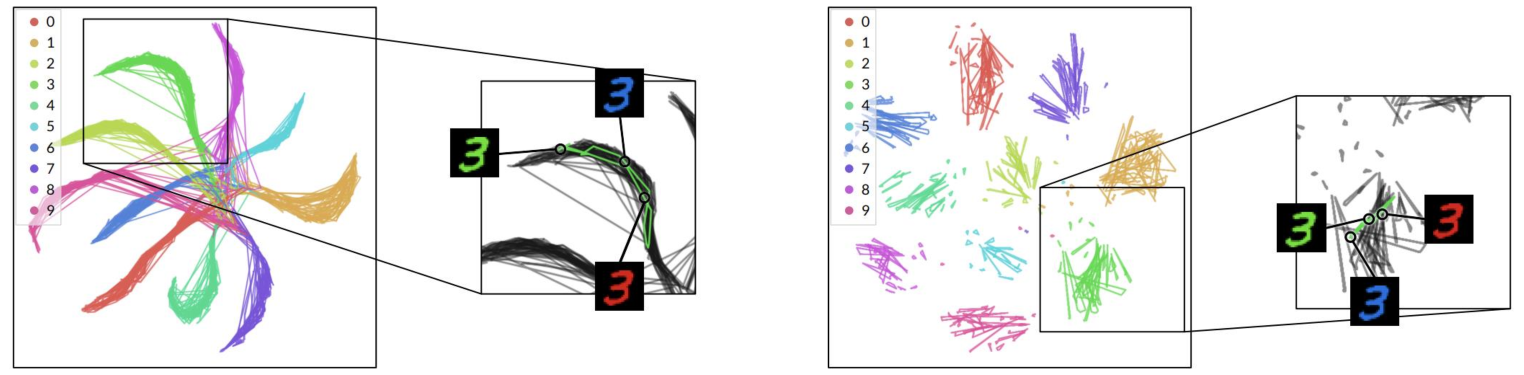

Figure 3. Hue shift MNIST feature map visualization.

	A/A	A/B	Params
Z2CNN	**1.54 (0.10)**	57.38 (22.06)	22,130
Hue-4*	1.97 (0.25)	2.08 (0.15)	25,690
Hue-3*	1.79 (0.25)	**1.81 (0.28)**	22,658
CEConv	1.79 (0.13)	1.83 (0.19)	28,739
CEConv-2	2.81 (0.44)	5.08 (0.18)	30,539

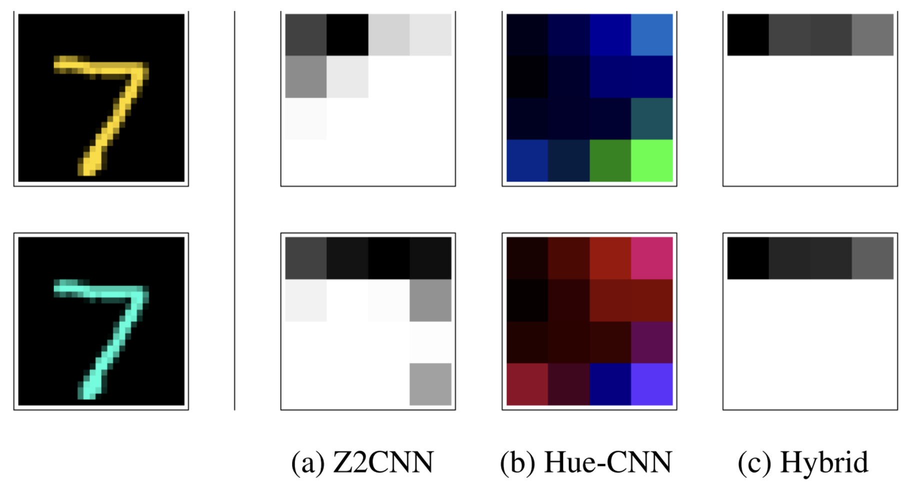

Figure 4. Comparison of feature maps for equivariant and non-equivariant architectures.

3DShapes: Local hue equivariance

Network	Group	A/A	A/B	A/C	Params
CNN	$\mathbb{Z}^2$	0.00 (0.00)	51.25 (9.59)	26.66 (19.60)	20,192
	Hue-4*	**0.00 (0.00)**	**0.00 (0.00)**	**0.00 (0.00)**	21,832
	Hue-3*	0.00 (0.00)	0.03 (0.04)	0.04 (0.03)	19,478
	CEConv	0.00 (0.00)	0.02 (0.03)	0.05 (0.04)	26,441

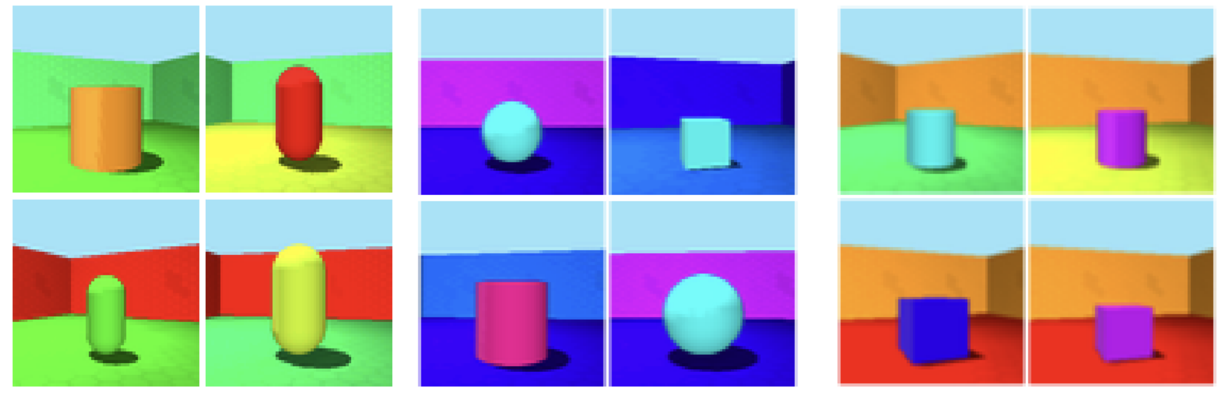

Figure 5. Hue-shifted 3DShapes dataset.

Camelyon17: Saturation shift in the wild

	Error	Params
ResNet50	28.91 (7.58)	23.5M
Sat-d_1*	19.60 (2.08)	23.3M
Sat-d_2*	**16.08 (2.68)**	23.3M
Sat-d_3*	24.53 (7.13)	23.3M
Sat-d_4*	23.57 (6.04)	23.3M

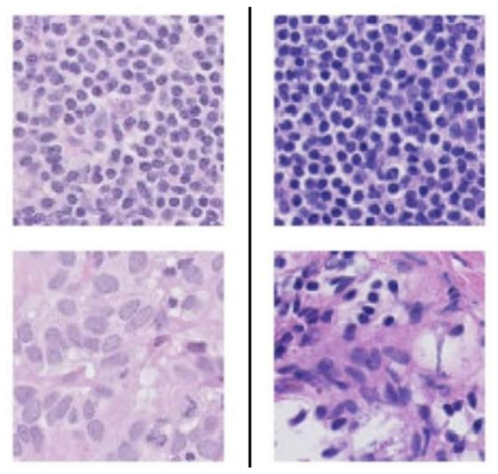

Figure 6. Camelyon17 Dataset.

CIFAR-10: Color shift in the wild

	ResNet44	Hue-4*	Sat-d_1*	Sat-d_2*	Sat-d_3*
Error	**7.86 (1.14)**	8.83 (0.64)	10.45 (1.76)	10.26 (0.43)	9.24 (0.27)
Params	2.64M	2.63M	2.55	2.55	2.55M

	Hue-Sat-d_1*	Hue-Sat-d_2*	Hue-Sat-d_3*	Hue-3*	CEConv
Error	10.68 (0.78)	11.13 (0.84)	11.16 (0.38)	8.42 (0.39)	8.86 (0.33)
Params	2.50M	2.50M	2.50M	2.55M	1.15M

Figure 7. CIFAR-10 automobile class sorted by hue.

Acknowledgements

We would like to thank members of the CAB Lab for valuable discussions and feedback, and the Princeton Department of Computer Science for providing compute resources. This work is partially supported by the Princeton School of Engineering and Applied Science.