EXPLORATIONS IN TEXTURE LEARNING

Blaine Hoak
University of Wisconsin-Madison
bhoak@cs.wisc.edu

Patrick McDaniel
University of Wisconsin-Madison
mcdaniel@cs.wisc.edu

ABSTRACT

In this work, we investigate *texture learning*: the identification of textures learned by object classification models, and the extent to which they rely on these textures. We build texture-object associations that uncover new insights about the relationships between texture and object classes in CNNs and find three classes of results: associations that are strong and expected, strong and not expected, and expected but not present. Our analysis demonstrates that investigations in texture learning enable new methods for interpretability and have the potential to uncover unexpected biases. Code is available at https://github.com/blainehoak/texture-learning.

1 INTRODUCTION AND BACKGROUND

Convolutional Neural Networks (CNNs) have been shown to be more biased towards texture (repeated patterns), rather than shape like human vision is Geirhos et al. (2019). Functionally, this suggests that models learn to associate object classes with textures that are present in the image, rather than shapes. This adherence to texture bias not only highlights discrepancies between human and machine vision, but may also impact model robustness and generalization Geirhos et al. (2019).

While prior works have focused on measuring, mitigating, and explaining texture bias in CNNs Geirhos et al. (2019; 2021); Hermann et al. (2020); Gatys et al. (2015), in this work we leverage the existence of texture bias to uncover what kinds of textures are learned. Specifically, we investigate *texture learning*: the identification of textures learned by models during training and the extent to which these textures are associated with objects. We do this by building a mapping of texture-object associations, which allow us to understand (a) what kind of textures models may be biased toward and (b) when this has the potential to be problematic.

2 BUILDING TEXTURE-OBJECT ASSOCIATIONS

To uncover the textures that are learned by models, we build texture-object associations. Specifically, we input texture-only images into an ImageNet trained model and measure the degree to which certain textures are classified as specific objects. Importantly, and contrasting with prior work, the textures we explore are representative of texture classes that go beyond the typical textures that may be easily associated with ImageNet objects (e.g., elephant skin texture is easily associated with elephant objects, but bumpy textures do not readily map to one object class). This is to ensure that we remain free of assumptions about what textures "should" be associated with certain objects, and that we are capturing texture learning phenomena that may not be as expected.

To this end, we use the Describable Textures Dataset (DTD) Cimpoi et al. (2014) as our texture dataset. The DTD consists of 5640 images, each of which is labeled with one of 47 texture classes (e.g., bubbly, scaly, polka-dotted). We use a pretrained ResNet50 model (see Appendix A.1 for details) to classify each of these texture images as belonging to one of the 1000 ImageNet classes (objects). **Notably, our experiments use a model trained on one dataset (ImageNet) yet are evaluated with an entirely different phenomenon (i.e., textures from DTD).**

With these texture classifications, we measure the effect size for each texture-object pairing (47 texture classes $\times$ 1000 object classes) by taking the ratio of samples belonging to the texture class that were classified as the corresponding ImageNet class. In other words, the effect size for texture class A and object class B represents how many samples belonging to texture class A were predicted to be

Texture class	Object class	Effect	Object class	Effect	Object class	Effect
honeycombed	honeycomb	0.731	chain_mail	0.071	velvet	0.027
cobwebbed	spider_web	0.655	poncho	0.046	radio_telescope	0.046
waffled	waffle_iron	0.427	honeycomb	0.117	pretzel	0.075
striped	zebra	0.381	tiger	0.169	velvet	0.093
knitted	dishrag	0.331	wool	0.239	cardigan	0.188
stratified	cliff	0.305	velvet	0.140	stone_wall	0.125
spiralled	coil	0.296	maze	0.061	chambered_nautilus	0.043
bubbly	bubble	0.286	beer_glass	0.104	Petri_dish	0.077
dotted	bib	0.248	shower_curtain	0.148	wallet	0.097
polka-dotted	bib	0.247	Windsor_tie	0.125	wallet	0.089

Table 1: First 10 rows of the texture-object associations with the top 3 most predicted objects (and their effect size) for each texture.

object class B. Thus, higher effect sizes correspond to stronger texture-object associations. For each of the 47 texture classes, the top 3 objects classes with the highest effect size are reported. Table 2 of Appendix A contains the texture-object associations for all 47 texture classes (3 object classes per texture class), sorted highest to lowest by the first effect size column. For brevity, Table 1 displays the first 10 rows of this table, corresponding to the top 10 textures with the strongest associations.

The texture-object associations yield multiple interesting results, which we divide into three types based on (a) how expected the relationship between texture and object is (i.e., if humans would naturally associate the texture with the object) and (b) the strength of the association that emerged in our results. Below we provide examples and descriptions of each type. See corresponding sections of Appendix A for images of each example.

Expected & Strongly Present. The *honeycombed* textures (which consist of repeated hexagonal patterns in objects ranging from bathroom tiles to bee honeycombs) were classified as the *honeycomb* object 73.1% of the time. This is a strong association, and is not necessarily surprising, as honeycomb objects are largely composed of honeycombed textures. Despite the "expectedness" of the association, these results are still interesting for two reasons. First, this demonstrates that models are able to generalize well on textures alone for these object classes, even for examples that are in entirely different datasets. Second, given that the DTD honeycombed texture images consisted of a variety of objects beyond honeycombs, the strength of this association suggests that the ImageNet model is predominantly relying on texture to predict the classes of these categories of images, rather than color or shape. See Appendix A.2.1 for supporting images.

Not Expected & Strongly Present. The *polka-dotted* and *dotted* texture classes were most often mapped to the *bib* object (24.8 % and 24.7% of the time, respectively). While there is not an obvious object class that the polka-dotted or dotted textures would naturally be associated with, the strength of this association suggests that the model has indeed learned to associate these textures with the bib object. This could suggest a bias in the training data that was learned by the model: a large number of training examples for the bib object may contain polka-dotted or dotted textures. In subsequent investigations of the ImageNet training data, we found this to be true; a glance through some of the *bib* images recovered multiple examples of bibs with polka-dots (shown in Figure 1) See A.2.2.

Expected & Not Present. The *scaly* texture images, while consisting of images appearing to be fish and reptile scales, were not associated with any fish or reptile objects, but rather with the honeycomb object (13.5% of the time as shown in Table 2). These types of results highlight object classes that may not have learned generalizable textures. This could be due to the fact that the textures in the training examples of these objects were not diverse enough, or that these object classes learned to build stronger associations to shapes or colors, rather than textures. See A.2.3 for supporting images.

3 Conclusions

This methodology and subsequent findings can be used to uncover learned (and potentially unexpected) associations. This not only enables greater model interpretability, but can also highlight and identify specific unwanted biases in models.

URM STATEMENT

The authors acknowledge that at least one key author of this work meets the URM criteria of ICLR 2024 Tiny Papers Track.

ACKNOWLEDGEMENTS

This material is based upon work supported by, or in part by, the National Science Foundation under Grant No. CNS1946022 and Grant No. CNS2343611. Any opinions, findings, and conclusions or recommendations expressed in this publication are those of the author(s) and do not necessarily reflect the views of the National Science Foundation, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for government purposes notwithstanding any copyright notation hereon.

REFERENCES

Mircea Cimpoi, Subhransu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing Textures in the Wild. In *2014 IEEE Conference on Computer Vision and Pattern Recognition*, pp. 3606–3613, Columbus, OH, USA, June 2014. IEEE. ISBN 978-1-4799-5118-5. doi: 10.1109/CVPR.2014.461. URL https://ieeexplore.ieee.org/document/6909856.

Leon A. Gatys, Alexander S. Ecker, and Matthias Bethge. Texture Synthesis Using Convolutional Neural Networks, November 2015. URL http://arxiv.org/abs/1505.07376. arXiv:1505.07376 [cs, q-bio].

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, January 2019. URL http://arxiv.org/abs/1811.12231. arXiv:1811.12231 [cs, q-bio, stat].

Robert Geirhos, Kantharaju Narayanappa, Benjamin Mitzkus, Tizian Thieringer, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. Partial success in closing the gap between human and machine vision, October 2021. URL http://arxiv.org/abs/2106.07411. arXiv:2106.07411 [cs, q-bio].

Katherine L. Hermann, Ting Chen, and Simon Kornblith. The Origins and Prevalence of Texture Bias in Convolutional Neural Networks, November 2020. URL http://arxiv.org/abs/1911.09071. arXiv:1911.09071 [cs, q-bio].

Sébastien Marcel and Yann Rodriguez. Torchvision the machine-vision package of torch. In *Proceedings of the 18th ACM international conference on Multimedia*, MM '10, pp. 1485–1488, New York, NY, USA, October 2010. Association for Computing Machinery. ISBN 978-1-60558-933-6. doi: 10.1145/1873951.1874254. URL https://dl.acm.org/doi/10.1145/1873951.1874254.

Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge, January 2015. URL http://arxiv.org/abs/1409.0575. arXiv:1409.0575 [cs].

A APPENDIX

A.1 EXPERIMENTAL DETAILS

The pretrained ResNet50 used in our experiments was obtained from torchvision Marcel & Rodriguez (2010) with the default model weights. The model was trained on ImageNet Russakovsky et al. (2015). The model was evaluated on the DTD dataset using the following data preprocessing steps: (1) resize the image to 256×256, (2) center crop the image to 224×224, (3) normalize the image using the mean and standard deviation of the ImageNet training dataset. All experiments

were run on a single NVIDIA 2080Ti GPU. Complete code to replicate experiments can be found at `https://github.com/blainehoak/texture-learning`.

A.2 IMAGE EXAMPLES

A.2.1 EXPECTED & STRONGLY PRESENT

Images associated with the honeycombed class of the Describable Textures Dataset can be browsed at `https://www.robots.ox.ac.uk/~vgg/data/dtd/view.html?categ=honeycombed`.

A.2.2 NOT EXPECTED & STRONGLY PRESENT

Images associated with the dotted and polka dotted classes of the Describable Textures Dataset can be browsed at `https://www.robots.ox.ac.uk/~vgg/data/dtd/view.html?categ=dotted` and `https://www.robots.ox.ac.uk/~vgg/data/dtd/view.html?categ=polka_dotted`, respectively.

Upon further inspection of a portion of the ImageNet training data, we were able to easily find multiple examples where images in the *bib* class contained dots or polka-dots. A few examples are shown in Figure 1. This supports our hypothesis that the model may have learned a bias for polka-dots in the bib class, leading to strong object-texture associations in our results. This finding demonstrates that texture learning analysis may be a fruitful direction for uncovering biases in models.

Figure 1: Examples of bibs with polka-dots in the ImageNet training data.

A.2.3 EXPECTED & NOT PRESENT

Images associated with the scaly class of the Describable Textures Dataset can be browsed at `https://www.robots.ox.ac.uk/~vgg/data/dtd/view.html?categ=scaly`.

A.3 EXTENDED RESULTS

Table 2 shows the full table of texture-object associations for all 47 texture classes on ResNet50. Results on Resnet152 can be found in Table 3

Texture class	Object class	Effect	Object class	Effect	Object class	Effect
honeycombed	honeycomb	0.731	chain mail	0.071	velvet	0.027
cobwebbed	spider web	0.655	poncho	0.046	radio telescope	0.046
waffled	waffle iron	0.427	honeycomb	0.117	pretzel	0.075
striped	zebra	0.381	tiger	0.169	velvet	0.093
knitted	dishrag	0.331	wool	0.239	cardigan	0.188
stratified	cliff	0.305	velvet	0.140	stone wall	0.125
spiralled	coil	0.296	maze	0.061	chambered nautilus	0.043
bubbly	bubble	0.286	beer glass	0.104	Petri dish	0.077
dotted	bib	0.248	shower curtain	0.148	wallet	0.097
polka-dotted	bib	0.247	Windsor tie	0.125	wallet	0.089
paisley	velvet	0.223	wool	0.112	shower curtain	0.103
wrinkled	velvet	0.219	quilt	0.153	wool	0.051
frilly	head cabbage	0.209	hoopskirt	0.105	velvet	0.069
grid	window screen	0.199	oscilloscope	0.114	shoji	0.063
crystalline	plastic bag	0.193	head cabbage	0.082	honeycomb	0.068
lacelike	handkerchief	0.191	velvet	0.119	stole	0.108
perforated	strainer	0.190	space heater	0.080	honeycomb	0.074
stained	velvet	0.184	volcano	0.040	potpie	0.035
woven	hamper	0.175	velvet	0.156	dishrag	0.100
blotchy	velvet	0.164	ant	0.058	fig	0.032
gauzy	shower curtain	0.158	velvet	0.079	window shade	0.068
cracked	stone wall	0.158	guillotine	0.074	spider web	0.074
braided	knot	0.155	hamper	0.125	dishrag	0.097
zigzagged	maze	0.153	envelope	0.131	quilt	0.115
meshed	chainlink fence	0.148	honeycomb	0.140	window screen	0.137
interlaced	maze	0.148	prayer rug	0.092	shield	0.065
veined	leaf beetle	0.143	head cabbage	0.095	sulphur butterfly	0.049
lined	shower curtain	0.142	web site	0.094	window shade	0.073
banded	shower curtain	0.142	bib	0.079	Windsor tie	0.079
marbled	velvet	0.137	cliff	0.052	spider web	0.044
flecked	wool	0.135	velvet	0.080	cardigan	0.069
scaly	honeycomb	0.135	tile roof	0.071	wool	0.061
matted	wool	0.132	komondor	0.070	wig	0.059
pleated	shower curtain	0.129	velvet	0.118	window shade	0.102
crosshatched	window screen	0.127	velvet	0.069	handkerchief	0.066
fibrous	hay	0.126	pot	0.076	matchstick	0.050
swirly	fire screen	0.116	velvet	0.103	shower curtain	0.084
grooved	radiator	0.115	velvet	0.100	doormat	0.084
porous	French loaf	0.115	honeycomb	0.049	velvet	0.044
chequered	wool	0.114	tray	0.108	crossword puzzle	0.079
studded	strainer	0.110	Windsor tie	0.105	cuirass	0.059
potholed	volcano	0.108	geyser	0.090	cliff dwelling	0.063
freckled	lipstick	0.104	seat belt	0.083	Band Aid	0.064
sprinkled	ice cream	0.075	dough	0.070	pretzel	0.052
bumpy	custard apple	0.073	jackfruit	0.049	spaghetti squash	0.047
pitted	pomegranate	0.068	doormat	0.047	switch	0.042
smeared	mask	0.057	velvet	0.054	jellyfish	0.041

Table 2: Texture-object associations with the top 3 most predicted objects (and their effect size) for each texture.

Texture class	Object class	Effect	Object class	Effect	Object class	Effect
honeycombed	honeycomb	0.753	Christmas_stocking	0.026	coil	0.026
cobwebbed	spider_web	0.691	barn_spider	0.074	shower_curtain	0.025
waffled	waffle_iron	0.533	honeycomb	0.078	tile_roof	0.044
spiralled	coil	0.471	maze	0.058	knot	0.038
striped	zebra	0.426	tiger	0.167	velvet	0.056
bubbly	bubble	0.425	beer_glass	0.085	honeycomb	0.057
knitted	dishrag	0.417	wool	0.157	cardigan	0.148
dotted	bib	0.415	shower_curtain	0.113	wallet	0.066
polka-dotted	bib	0.340	pillow	0.078	shower_curtain	0.068
stratified	cliff	0.309	cliff_dwelling	0.082	velvet	0.073
wrinkled	velvet	0.257	quilt	0.115	packet	0.044
zigzagged	pillow	0.243	maze	0.122	wool	0.113
grid	window_screen	0.243	manhole_cover	0.043	oscilloscope	0.043
paisley	velvet	0.235	shower_curtain	0.165	pillow	0.087
lacelike	handkerchief	0.221	quilt	0.142	shower_curtain	0.115
meshed	window_screen	0.198	honeycomb	0.153	chainlink_fence	0.117
potholed	volcano	0.195	manhole_cover	0.161	valley	0.068
banded	shower_curtain	0.190	web_site	0.155	bib	0.129
gauzy	shower_curtain	0.183	velvet	0.139	window_shade	0.096
chequered	wool	0.179	shower_curtain	0.103	wall_clock	0.085
perforated	window_screen	0.179	strainer	0.170	honeycomb	0.054
woven	hamper	0.177	dishrag	0.097	doormat	0.080
frilly	head_cabbage	0.169	gown	0.102	vase	0.059
matted	wool	0.169	wig	0.119	komondor	0.042
pleated	shower_curtain	0.157	window_shade	0.139	wool	0.087
braided	knot	0.155	hamper	0.121	wool	0.103
fibrous	hay	0.139	wool	0.070	pot	0.052
veined	leaf_beetle	0.139	head_cabbage	0.104	buckeye	0.087
lined	web_site	0.138	wool	0.095	shower_curtain	0.095
blotchy	velvet	0.137	ant	0.043	switch	0.043
crosshatched	window_screen	0.134	wallet	0.042	handkerchief	0.042
grooved	radiator	0.127	velvet	0.059	doormat	0.059
freckled	lipstick	0.120	seat_belt	0.068	cellular_telephone	0.051
marbled	velvet	0.119	cliff	0.068	spider_web	0.042
porous	French_loaf	0.118	stone_wall	0.042	manhole_cover	0.042
crystalline	plastic_bag	0.117	head_cabbage	0.108	shower_cap	0.058
studded	strainer	0.111	Windsor_tie	0.068	purse	0.060
flecked	cardigan	0.103	wool	0.095	velvet	0.052
cracked	volcano	0.103	stone_wall	0.077	tiger_beetle	0.077
sprinkled	ice_cream	0.102	Petri_dish	0.042	tray	0.042
bumpy	custard_apple	0.101	jackfruit	0.059	thimble	0.034
scaly	honeycomb	0.094	wool	0.068	tile_roof	0.068
stained	velvet	0.093	paper_towel	0.059	jean	0.034
interlaced	maze	0.093	prayer_rug	0.085	doormat	0.059
swirly	shower_curtain	0.087	fire_screen	0.061	pillow	0.061
smeared	paintbrush	0.050	mask	0.042	handkerchief	0.042
pitted	switch	0.050	pomegranate	0.042	ant	0.034

Table 3: Texture-object associations with the top 3 most predicted objects (and their effect size) for each texture on ResNet152.