

Is Robustness To Transformations Driven by Invariant Neural Representations?

Syed Suleman Abbas Zaidi^{1,2,*}

Xavier Boix^{1,3,*†}

Neeraj Prasad^{1,*}

Sharon Gilad-Gutnick¹

Shlomit Ben-Ami¹

Pawan Sinha^{1,†}

¹ Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology (USA)

² Department of Electrical and Computer Engineering, Technical University of Munich (Germany)

³ Center of Brains, Minds and Machines, Massachusetts Institute of Technology (USA)

* Shared first authorship.

† Corresponding authors. Emails: xboix@mit.edu and psinha@mit.edu

Abstract

Deep Convolutional Neural Networks (DCNNs) have demonstrated impressive robustness to recognize objects under transformations (*e.g.* blur or noise) when these transformations are included in the training set. A hypothesis to explain such robustness is that DCNNs develop invariant neural representations that remain unaltered when the image is transformed. Yet, to what extent this hypothesis holds true is an outstanding question, as including transformations in the training set could lead to properties different from invariance, *e.g.* parts of the network could be specialized to recognize either transformed or non-transformed images. In this paper, we analyze the conditions under which invariance emerges. To do so, we leverage that invariant representations facilitate robustness to transformations for object categories that are not seen transformed during training. Our results with state-of-the-art DCNNs indicate that invariant representations strengthen as the number of transformed categories in the training set is increased. This is much more prominent with local transformations such as blurring and high-pass filtering, compared to geometric transformations such as rotation and thinning, that entail changes in the spatial arrangement of the object. Our results contribute to a better understanding of invariant representations in deep learning, and the conditions under which invariance spontaneously emerges.

1 Introduction

A widely-known strategy to gain robustness to image and object transformations in Deep Convolutional Neural Networks (DCNNs) is to include the transformation in the training set [6, 7, 18, 22]. Geirhos *et al.* [7] recently demonstrated that this strategy is not only successful in achieving robustness to local transformations, such as blur and noise, but also outperforms human-level recognition accuracy in transformed images. In a subsequent paper, Geirhos *et al.* [6] show that including some transformations in the training set that emphasize the shape of the object rather than the texture leads to robustness to transformations beyond the ones seen during training.

Despite these recent results demonstrating effective ways of achieving robustness to transformations, the field still lacks a good understanding of the underlying neural mechanisms that enable such robustness. Decades of research at the intersection of Neuroscience and Computer Vision has led to the hypothesis that neural networks may develop invariant neural representations to achieve robustness to transformations [1, 8, 15–17], *i.e.* the internal representations may remain mostly unaltered when the image is transformed¹. Yet, to what extent this hypothesis holds true is an outstanding question. It

¹Note that we use the term *robustness to a transformation*, or simply *robustness*, to denote that DNN classification accuracy does not deteriorate in transformed images. Also, we use the term *invariance to a transformation*, or *invariance*, to indicate that the internal neural representation of the DNN is not altered by the transformation.

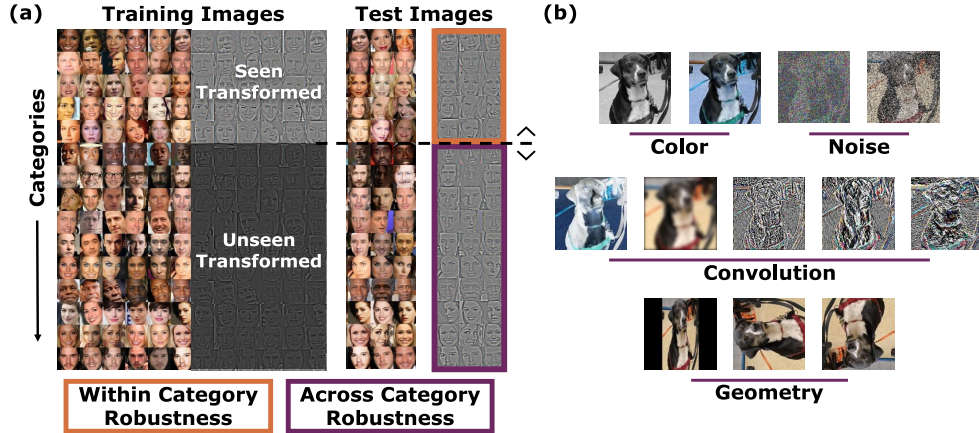


Figure 1: *Experimental Paradigm and Image Transformations* (a) The DNN is trained with the transformations in the training set restricted to the *seen-transformed* categories. The accuracy for transformed images is evaluated in categories that are either *seen* or *unseen* transformed. (b) We analyze 12 image transformations from four families: color, convolutional, noise and geometric.

is possible that DCNNs develop properties different from invariance when including transformations in the training set, *e.g.* parts of the network may specialize in recognizing either transformed or non-transformed images. Answering what representations emerge in the network could reveal new insights that help to develop more robust models.

In this paper, we investigate to what extent and under what conditions invariant representations emerge in DCNNs trained with transformations. To do so, we leverage the property of invariant representations that enable robustness to object categories not transformed in the training set [15]. We denote this form of robustness as *across-category robustness*, while we use the term *within-category robustness* to designate robustness to categories that are seen transformed during training. Note that since invariant representations result in *across-category robustness*, invariance enables reducing the amount of training examples, as not all objects need to be seen transformed to achieve robustness.

To validate that the amount of invariance in the network drives robustness and is not epiphenomenal, we evaluate across-category robustness by analyzing DCNNs trained such that the transformations in the training set are restricted to a subset of categories. A series of experiments with state-of-the-art DCNNs in FaceScrub [13] and ImageNet [4] reveal that invariant representations emerge as the number of transformed categories in the training set is increased. This trend is remarkable for local transformations, *e.g.* blurring and high-pass filtering, compared to geometric transformations such as rotations and thinning. Our analysis uncovers new evidence that invariance emerges at the individual neural level and demonstrates that it facilitates robustness to categories that the network has not seen transformed during training. These results are a first milestone to understand the neural representations that drive robustness to transformations.

2 Signatures of Invariant Neural Representations

In this section, we introduce our methodology for investigating to what extent invariant neural representations drive robustness to transformations. For simplicity, we consider DCNNs trained on a set of images with only one transformation, as well as a set of non-transformed images. This yields robustness to the trained transformation but is unlikely to result in robustness to other transformations that the DCNN was not trained on [7].

A neural network is completely invariant when the activations in the DCNN are identical for the transformed and non-transformed images. While this is ideal, it is important to note that it is not a pre-requisite for robustness: in practice, a certain amount of invariance may emerge, even without attaining identical neural activity, and result in robustness. This adds a difficulty in the analysis of invariance because it is difficult to determine whether small changes in the neural activity result in loss of robustness. Adversarial perturbations are one such example [20]. Thus, to validate that a

certain amount of invariance does indeed lead to robustness, we evaluate if that amount of invariance results in across-category robustness. Recall that across-category robustness is a consequence of invariant representations and hence, evaluating it jointly with the amount of invariance provides solid evidence about the role of invariance for robustness. In the following, we first introduce the paradigm to evaluate across-category robustness, and then, the measure of invariance to transformations.

2.1 Paradigm to Evaluate the Across-Category Robustness

To test across-category robustness we use the procedure depicted in Fig. 1. Images of the subset of categories denoted as *seen-transformed* appear both transformed and non-transformed in the training set. The remainder of the categories, denoted *unseen-transformed*, appear in the training set only in their non-transformed version. Note that the sum of the number of categories in each set is equal to the total number of categories in the dataset. The network is trained with both *seen-* and *unseen-transformed* categories by uniformly randomizing the categories and images during each training step. When an image of the *seen-transformed* set is used, it is also randomly chosen with equal probability to appear as transformed or not transformed.

The network’s task is to predict the category of the input image. We evaluate the accuracy of the network on transformed test images for the *seen-* and *unseen-transformed* categories separately, which yields the within- and across-category accuracy, respectively. We train multiple networks with different number of *seen-transformed* categories during training. In all the experiments, the network predicts a category among all categories in the dataset, *i.e.* the possible output choices of the network remain the same when evaluating within- and across-category accuracy. Thus, robustness among all the networks can be compared without bias because the only change between them is the number of *seen-transformed* categories during training.

Note that a possible alternative to our paradigm could be based on training the network in two steps, *i.e.* first training the *seen-transformed* and then the *unseen-transformed* categories. Yet, this has limitations to analyse invariance because lack of across-category robustness can be due to other reasons unrelated to invariance. If the training of the *unseen-transformed* categories is done by fine-tuning the network, there could be the so-called “catastrophic forgetting” of the transformation [11]. Otherwise, without fine-tuning, the features learnt for the *seen-transformed* categories may not be generalizing well to the *unseen-transformed*. Training both sets at the same time, as in our paradigm, alleviates those problems.

2.2 Invariance to Transformations at the Individual Neural Level

Understanding DCNNs at the individual neural level has lead to promising results that provide granular and simple explanations of the underlying mechanisms behind DCNN generalization [14, 23, 24]. Individual neurons are commonly interpreted as feature detectors that are tuned to image features. A neuron is active, *i.e.* the neuron’s output value is high, when the feature is present in the image, otherwise the neuron is not active and the output value of the neuron is low. Neurons are tuned to features that are not necessarily interpretable by humans [3]. We analyse the invariance of a neuron by evaluating the change of its activity when the image is transformed. For example, a neuron may be tuned to detect a dog nose because it is only active in the presence of a dog nose in the image. If the neuron is invariant, it will continue being active when the dog nose is transformed.

Only invariance in neurons that are active may help to achieve robustness to transformations. If a neuron is inactive in both the transformed and the non-transformed images, the activity is invariant to the transformation. Yet, this is because the feature that the neuron is tuned to is not present in the image, which does not contribute to the network’s classification accuracy. Also, neurons whose activity is constant across all images, *i.e.* neurons that are not tuned to any feature, are invariant and do not contribute to network’s accuracy. Thus, invariance is helpful to achieve robustness to transformations only when neurons that are activated by an image are also invariant to the transformation.

Following these intuitions, Goodfellow *et al.* [8] introduced a measure of invariance which analyzes the activity in images that generate the top-1% of the neuron’s output value. These images are transformed and invariance is evaluated by assessing whether the activity remains in the top-1% or not. This metric is effective to discard the cases of invariance that are not related to robustness. Yet, it is unclear to what degree the activity generated in a neuron by the rest of the 99% of images

that are not analysed relates to robustness. In the following, we introduce a procedure that controls that the majority of the neural activity related to robustness is taken into account by the invariance measurement. We first introduce the procedure to determine when a neuron is active and then, the measure of invariance of active neurons.

Active Neurons. Let \mathcal{X} be the images of the test set and let $x \in \mathcal{X}$ be one of its images. We refer to the transformed version of the image as $T(x)$. For the sake of notation simplicity, we consider $T(x)$ to be the transformation that is included in the training set without specifying the type transformation. Let $\hat{f}_k(x)$ be the activity of the neuron indexed by k when the network’s input is $x \in \mathcal{X}$. Note that $\hat{f}_k(x)$ can be multiplied by a factor without affecting the performance of the network, if the post-synaptic weights are divided by that factor. Thus, the order of magnitude of $\hat{f}_k(x)$ is not relevant and we discard it by normalizing $\hat{f}_k(x)$. We define $f_k(x)$ as the normalized activity of neuron k , such that $f_k(x)$ is equal to 1 for the image that generates the maximum activity among transformed and not-transformed images. Note that the minimum value that $f_k(x)$ can take is 0 because the neurons use ReLU activation [12].

We define the neuronal tuning of a neuron as the set of images that generate an activity above a threshold. Let τ be such threshold and let \mathcal{A}_k be the set of images that generate an output value higher than τ . Thus, \mathcal{A}_k represents the neuronal tuning of neuron k as it contains the set of images that yield the neuron active. \mathcal{A}_k is the following subset of \mathcal{X} : $\mathcal{A}_k = \{x \in \mathcal{X} \mid (f_k(x) > \tau) \vee (f_k(T(x)) > \tau)\}$, in which where \vee is the “logic or” operator. Note that the activity generated by both transformed and non-transformed images is considered.

In order to determine the value of τ , we need to take into account a compromise between the following two factors. On one hand, the value of τ should be as high as possible such that the neuronal tuning is specific and invariance due to lack of activity is discarded. On the other hand, τ should be made as low as possible in order to include in \mathcal{A}_k all the activity that is relevant for the network’s robustness. To do so, we validate that the network’s robustness is not affected by ablating all neurons except the active, given an image. Also, we validate that the network’s accuracy is at chance when all the active neurons are ablated. These tests provide reassurance that the active neurons drives the network’s robustness. In the experiments section, we show that neurons are active only on a small set of images and their activity is almost the only responsible of the network’s robustness.

Measure of Invariance in Active Neurons. In order to measure the invariance of a neuron, only the images in the set \mathcal{A}_k will be used. This is because, as previously discussed, invariance of neurons that are not active for a given image are not related to robustness. We define $h_k(x)$ as the invariance measure of neuron k given an image $x \in \mathcal{X}$. We use the following expression based on the normalized difference between the neurons’ activity of the transformed and non-transformed images:

$$h_k(x) = 1 - \frac{|f_k(T(x)) - f_k(x)|}{|f_k(T(x)) + f_k(x)|}. \quad (1)$$

We can see by analysing Eq. (1) that $h_k(x)$ is equal to 1 when the neuron is perfectly invariant, *i.e.* $f_k(T(x)) = f_k(x)$, and is equal to 0 when the neuron’s activity varies maximally, such that $f_k(T(x)) = 0$ and $f_k(x) = 1$ and vice versa. This invariance measure takes values between 0 and 1 depending on how much the activity of x and $T(x)$ varies. We define I_k as the invariance coefficient of neuron k . I_k is obtained by averaging $h_k(x)$ over the set of images that the neuron is active, \mathcal{A}_k , *i.e.* $I_k = \frac{1}{\#\mathcal{A}_k} \sum_{x \in \mathcal{A}_k} h_k(x)$, in which $\#\mathcal{A}_k$ is the cardinality of the set \mathcal{A}_k (the number of images that neuron k is active for). In the experiments we analyze the distribution of I_k among neurons in a layer.

3 Results

We evaluate robustness to 12 image transformations shown in Fig. 1 (see Supplement for details). These are grouped in four transformation types: color (gray-scale and color rotation), geometric (thinning and rotation), convolutional (blur, high-pass and horizontal, vertical filtering and contrast inversion) and noise (white and salt & pepper). Our choice of transformations was motivated by studies showing that DCNNs are robust to them when they are included in the training set [7].

To ensure that our results are not dataset- or architecture-dependent, we evaluate AlexNet [12] in the FaceScrub dataset [13] and ResNet-18 [9] in the ImageNet dataset [4] (see Supplement). In ImageNet

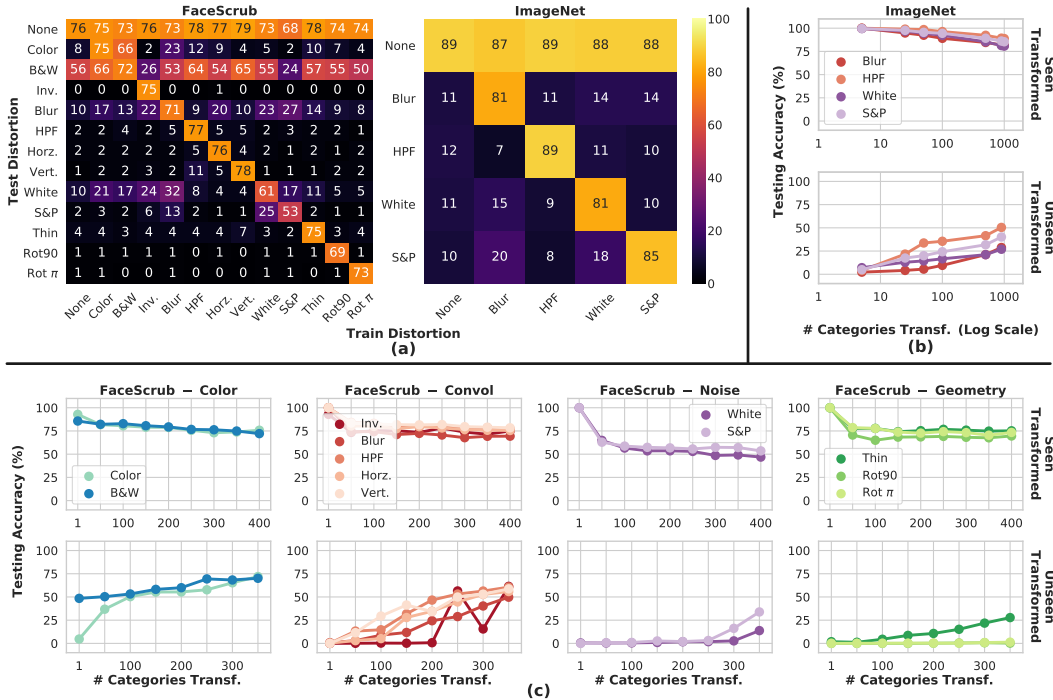


Figure 2: *Within- and Across-Category Accuracy in FaceScrub and ImageNet.* (a) Within-category accuracy when all object categories are seen transformed for FaceScrub (left) and ImageNet (right). (b) Within-category accuracy (top) and across-category accuracy (bottom) in ImageNet, for different number of transformed categories in the training set. (c) Same as (b) for FaceScrub, each family of transformations is displayed separately.

we only evaluate four transformations (blur, high-pass filtering, white noise and salt & pepper noise) as some of the other transformations are already present in the training set and can not be excluded (note that in ImageNet objects appear at different orientations, sizes, contrasts and colors). We use the loss function, regularizers, optimizer and other hyperparameter choices reported in the literature, except for the learning rates that are tuned for each experiment (see Supplement).

3.1 Results on Across-Category Robustness

In the following, we report the accuracy of the network using the paradigm introduced in section 2.1. For both FaceScrub and ImageNet, we report top-1 accuracy (%).

Reproducing Previous Results on Robustness to Transformations. Fig. 2a depicts the accuracy of networks trained with each transformation and tested on all transformations. As expected, on non-transformed images, the networks perform similarly to what is reported in the literature [9, 21] independently of which transformations being included in the training set. On transformed images, networks are robust to the transformations included in the training but not robust to other transformations, a finding consistent with Geirhos *et al.* [7]. Note that color transformations are an exception, as all networks have some degree of robustness to them. This may be because the color transformations are in part already included in the dataset as a consequence of different illumination conditions. Additionally, the network trained with blur transformation is slightly robust to the noise transformations, likely because training with blur yields large receptive fields in the first layer that may help to filter out the noise (a finding consistent with previous studies [21]).

Across-Category Robustness Improves as the Number of *Seen-transformed* Categories is Increased. In Fig. 2b and c, the within- and across-category accuracy are evaluated for different number of *seen-transformed* categories. The within-category accuracy (in the top of the figure) and the across-category accuracy (in the bottom) are reported. We can observe that the across-category accuracy increases as more categories are seen transformed. This indicates that invariant neural

representations may have emerged in the network (we can not confirm this without analysing the neural activity, as across-category robustness is a consequence of invariance but not the reverse). Also, observe that when the network is trained with few *seen-transformed* categories, the robustness for these categories can not be driven by invariance because the network is not across-category robust. Understanding the mechanisms for robustness when there is lack of invariance is an interesting open question that will be analyzed in future works.

Note also that the within-category accuracy decreases as more categories are seen transformed, and it reaches the values previously reported in Fig. 2a when all categories are seen transformed. Observe that the within-category accuracy is close to 100% when the number of *seen-transformed* categories is close to 1. This suggests that the network uses the transformation as a feature to identify the few *seen-transformed* categories. This can be validated with the confusion matrices between *seen-* and *unseen-transformed* categories (see Supplemental), which show that transformed images from *unseen-transformed* categories are mainly confused with *seen-transformed* categories.

Color, Convolutional and Noise Transformations lead to Higher Across-Category Robustness than Geometric Transformations. The increase of the across-category accuracy with the number of *seen-transformed* categories is common among transformations. Yet, Fig. 2b and c shows that the rate at which the across-category varies is different depending on the transformation. Observe that for color transformations, just few *seen-transformed* categories are sufficient to achieve a high across-category accuracy. This may be because the network is already quite robust to color transformations even training only with non-transformed images, as shown in Fig. 2a. For convolutional transformations, the increase of the across-category accuracy is not as high as for color transformations but it is remarkable. The across-category accuracy of the noise transformations is in-par with the convolutional in ImageNet but not in FaceScrub. Note that the within-category accuracy for noise transformations in FaceScrub is between 10% to 20% lower than for other transformations. Thus, in FaceScrub the amount of noise impairs both within- and across-category robustness, which suggests that the low across-category accuracy is not a characteristic of the noise transformation but of the overall accuracy in FaceScrub.

For geometric transformations, the across-category accuracy is lower than the rest of transformations, specially for rotations (the increase is only of a few percent when all categories are seen transformed). This may be because spatial re-arrangement of the object are involved in geometric transformations but not in the rest. These re-arrangements are larger for rotations of 90 or 180 degrees than for thinning, and rotations have much lower across-category accuracy than thinning. Also, to achieve across-category robustness for spatial re-arrangements requires capturing the long-range dependencies between object parts, and this has been shown to be difficult for feed-forward architectures [2, 5, 10, 19].

Across-Category Robustness depends on the Visual Homogeneity of the Object Categories. For convolutional and noise transformations, the across-category generalization in FaceScrub increases linearly with respect to the number of categories seen transformed, while in ImageNet it increases logarithmically. These difference between datasets may be because the object categories in FaceScrub are more visually homogeneous than in ImageNet, *i.e.* FaceScrub only contains faces while ImageNet is composed of a varied set of distinct objects. Thus, the visual homogeneity of the categories may facilitate across-category generalization (assuming that robustness to transformation can be learned just as well for any object category).

3.2 Results on Invariance to Transformations at the Individual Neural Level

Next, we evaluate the amount of invariance using the measure introduced in section 2.2.

Less than 20% of the Images Activate a Neuron. Recall that the threshold that determines if a neurons is active, τ , is the compromise between two factors: (i) neurons are only active for few images because they are tuned to specific patterns, *i.e.* τ is as high as possible, (ii) only the activity in active neurons matters for robustness, *i.e.* ablating all the inactive neurons in an image should not drop the accuracy whereas ablating all the active neurons in an image should bring the accuracy to chance. In practice, we implement a grid search to determine τ by evaluating values between 0 to 1 in steps of 0.05. We select the highest τ that leads to a drop of 1% of the accuracy when ablating the inactive neurons. This results in $\tau = 0.1$ for FaceScrub and $\tau = 0.05$ for ImageNet. We validate that neurons are active only for a small subset of images, *i.e.* the cardinality of \mathcal{A}_k is small for most neurons k , and also that each image activates a subset of neurons, *i.e.* the cardinality of $\{k|x \in \mathcal{A}_k\}$




Figure 3: *Active Neurons and Invariance when all categories are seen-transformed.* (a) Distribution of the number of active neurons per image and the number of images that activate a neuron. (b) and (c) Violin plots of the invariance coefficient among neurons in the penultimate layer, for FaceScrub and ImageNet, respectively. The invariance to each transformation is shown for networks trained with the transformation applied to all categories (indicated with $x + T(x)$) and for networks trained only with non-transformed images (x).

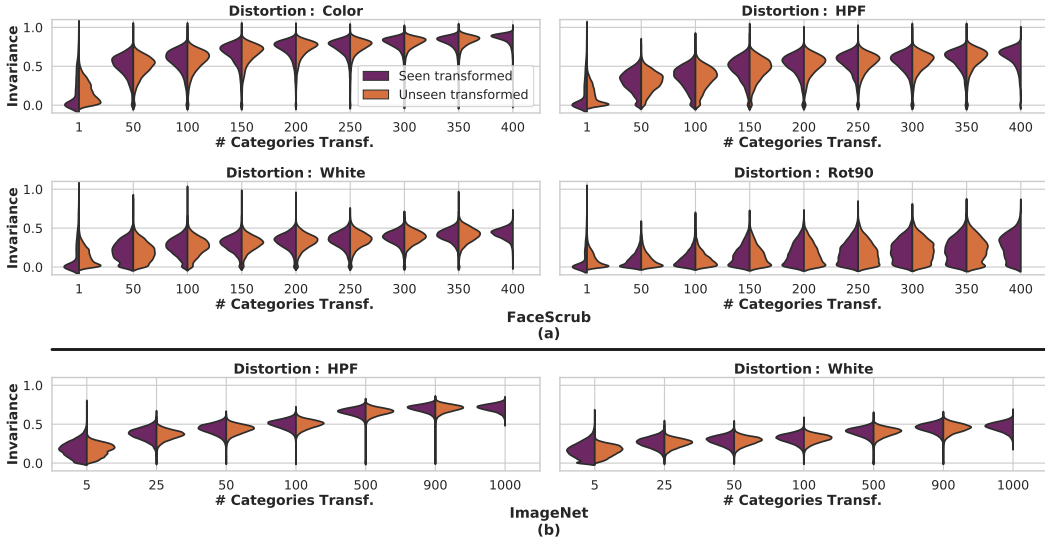


Figure 4: *Invariance for Different Number of Seen-transformed Categories.* (a) and (b) Violin plots of the invariance coefficient among neurons at the penultimate layer for FaceScrub and ImageNet, respectively. Different number of *seen-transformed* categories are displayed and the invariance coefficient is reported separately for *seen-* and *unseen-transformed* categories.

is small for most images x . In Fig. 3a we show the distributions of both cardinalities using violin plots. Notably, between 1 to 20% of the images activate a neuron, and between 10 to 20% of the neurons are active in an image. Thus, neurons are tuned to patterns that appear in few images and in most images only a subset of neurons is active, which fulfills factor (i). In the Supplement, we report the network’s accuracy obtained from ablating active and inactive neurons. We observe that when the inactive neurons are ablated (80 – 90% of the neurons per image), the network’s accuracy decreases less than 1% overall. When the active neurons are ablated (10 – 20% of the neurons per image), the network’s accuracy is very low for all transformations. This low but above chance accuracy suggests that the neural activity of the inactive neurons may play a minor but redundant role towards network’s robustness, which approximately fulfills factor (ii).

Invariance Increases when the Number of Seen-transformed Categories Increases. Recall that the invariance coefficient, I_k , is evaluated when the neurons are active. In Fig. 3b and c, we show the distribution of I_k among neurons in the penultimate layer for networks that are trained with or without transformations. All transformations lead to more invariance than when the training is only with non-transformed images. In the Supplement, we show that this invariance builds up across layers, which is in accordance with previous works [8, 15].

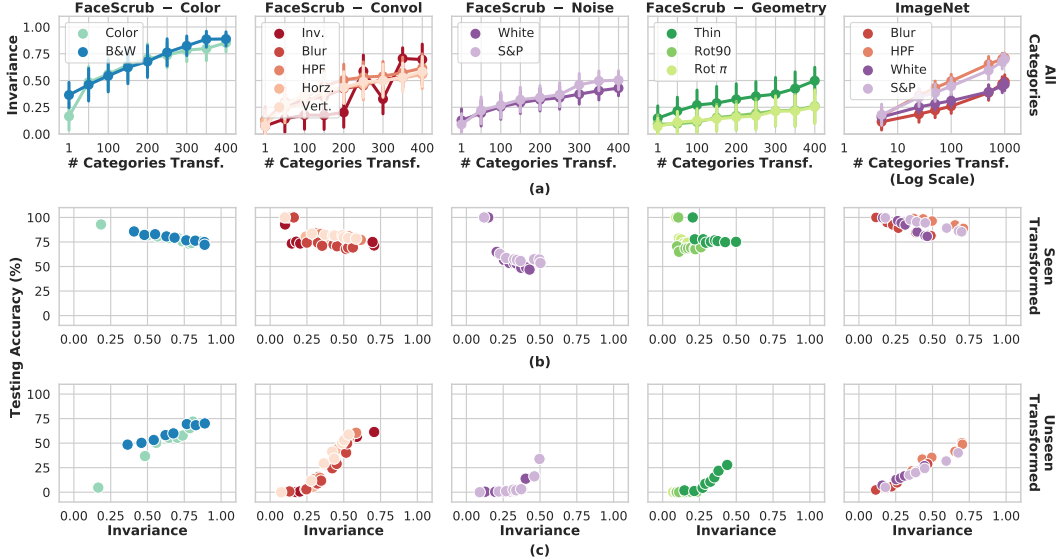


Figure 5: *Invariance and Robustness*. (a) Amount of invariance for different number of *seen-transformed* categories, transformations and datasets. (b) and (c) Relationship between amount of invariance and within- and across-category accuracy, respectively. Each dot represents a network trained with different number of *seen-transformed* categories.

In order to validate that this invariance is related to robustness and is not epiphenomenal, we investigate whether invariance drives the across-category robustness. Fig. 4 shows the distribution of I_k among neurons in the penultimate layer when the number of *seen-transformed* categories is increased. The distribution for the *seen-* and *unseen-transformed* categories is shown separately (see Supplement for all transformations). We observe that the amount of invariance remarkably increases when the number of *seen-transformed* categories is increased. Also, the amount invariance for the *seen-* and *unseen-transformed* categories is nearly identical (except when the number of *seen-transformed* categories is low because there is almost no invariance). Thus, invariance in the network is not category-dependent and emerges alongside across-category robustness. This is strong evidence that invariance drives robustness when the number of *seen-transformed* categories is large.

Invariance Explains the Across-Category Robustness. Next, we further analyse the relationship between invariance and across-category robustness by scrutinizing the differences among transformations. Fig. 5a shows the amount of invariance for all transformations and datasets as the number of *seen-transformed* categories is increased. We observe that the trends of invariance follow that of across-category robustness. To quantify whether there is a direct relationship between them, in Fig. 5b and c we depict the invariance and robustness for each network trained with a different number of *seen-transformed* categories. Observe that invariance is not needed to achieve within-category robustness (Fig. 5b), whereas the extent of across-category robustness is correlated with the amount of invariance for all transformations (Fig. 5c). Thus, the network develops properties different from invariance to achieve robustness but as we increase the number of *seen-transformed* categories the amount of invariance increases.

4 Conclusions

We have demonstrated that increasing the number of *seen-transformed* categories results in an increase of two network properties: the amount of invariance to transformations and the robustness to transformations of *unseen-transformed* categories. This suggests that invariance drives robustness to transformations depending on the number of *seen-transformed* categories. Furthermore, we have shown that invariance emerges at the individual neural level, which adds to the growing body of literature that uses individual neurons as the elemental building blocks to analyze DCNNs [14, 23, 24]. Key open questions derived from our results that will be tackled in future works are understanding why invariant representations emerge for some transformations, and what are the neural mechanisms for robustness when there is a lack of invariance (*e.g.* for geometric transformations).

Acknowledgements

We are grateful to Tomaso Poggio for his insightful advice and warm encouragement. This work is supported by the Center for Brains, Minds and Machines (funded by NSF STC award CCF-1231216), Fujitsu Laboratories Ltd. (Contract No. 40008401 and 40008819), the MIT-Sensetime Alliance on Artificial Intelligence and the R01EY020517 grant from the National Eye Institute (NIH).

Statement of Broader Impact

This work tackles the long studied but still obscure area of generalization in DCNNs on two fronts. First, it analyzes the limits of robustness to transformations by examining the extensibility of robustness to categories that are not seen transformed during training. Our results reveal weaknesses of the generalization capability of DNNs as robustness may degrade for categories that are not seen transformed during training. This phenomenon is severe for geometric transformations such as rotations. This could potentially be used to trick DCNNs in real-life applications. Nevertheless, understanding such weaknesses is necessary to build better and more robust DCNNs.

Secondly, by studying the role of invariant representations in individual neurons, this work identifies the important relationship between invariance and robustness for categories not seen transformed during training. By demonstrating the significance of invariant representations in driving robustness, this paper advocates for training methods that enforce invariance in the hidden layers. The understanding of generalization in DCNNs is a critical area that will provide strides in the advancements of modern computer vision and bring it closer to or beyond its rivaling human vision.

References

- [1] A. Achille and S. Soatto. Emergence of invariance and disentanglement in deep representations. *The Journal of Machine Learning Research*, 19(1):1947–1980, 2018.
- [2] A. Azulay and Y. Weiss. Why do deep convolutional networks generalize so poorly to small image transformations? *Journal of Machine Learning Research*, 2019.
- [3] D. Bau, B. Zhou, A. Khosla, A. Oliva, and A. Torralba. Network dissection: Quantifying interpretability of deep visual representations. In *CVPR*, pages 3319–3327. IEEE, 2017.
- [4] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [5] L. Engstrom, B. Tran, D. Tsipras, L. Schmidt, and A. Madry. Exploring the landscape of spatial robustness. In *ICML*, 2019.
- [6] R. Geirhos, P. Rubisch, C. Michaelis, M. Bethge, F. A. Wichmann, and W. Brendel. Imagenet-trained cnns are biased towards texture; increasing shape bias improves accuracy and robustness. In *ICLR*. 2019.
- [7] R. Geirhos, C. R. M. Temme, J. Rauber, H. H. Schütt, M. Bethge, and F. A. Wichmann. Generalisation in humans and deep neural networks. In *NeurIPS*. 2018.
- [8] I. Goodfellow, H. Lee, Q. V. Le, A. Saxe, and A. Y. Ng. Measuring invariances in deep networks. In *NIPS*, 2009.
- [9] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *CVPR*, 2016.
- [10] M. Jaderberg, K. Simonyan, A. Zisserman, et al. Spatial transformer networks. In *NIPS*, 2015.
- [11] J. Kirkpatrick, R. Pascanu, N. Rabinowitz, J. Veness, G. Desjardins, A. A. Rusu, K. Milan, J. Quan, T. Rammalho, A. Grabska-Barwinska, et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 2017.
- [12] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [13] H.-W. Ng and S. Winkler. A data-driven approach to cleaning large face datasets. In *ICIP*, 2014.
- [14] C. Olah, A. Satyanarayan, I. Johnson, S. Carter, L. Schubert, K. Ye, and A. Mordvintsev. The building blocks of interpretability. *Distill*, 2018.
- [15] T. Poggio and F. Anselmi. *Visual cortex and deep networks: learning invariant representations*. MIT Press, 2016.
- [16] R. Q. Quiroga, L. Reddy, G. Kreiman, C. Koch, and I. Fried. Invariant visual representation by single neurons in the human brain. *Nature*, 2005.

- [17] M. Riesenhuber and T. Poggio. Just one view: Invariances in inferotemporal cell tuning. In *NIPS*, 1998.
- [18] L. Schmidt, S. Santurkar, D. Tsipras, K. Talwar, and A. Madry. Adversarially robust generalization requires more data. In *NeurIPS*, 2018.
- [19] S. Srivastava, G. Ben-Yosef, and X. Boix. Minimal images in deep neural networks: Fragile object recognition in natural images. In *ICLR*, 2019.
- [20] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus. Intriguing properties of neural networks. *ICLR*, 2014.
- [21] L. Vogelsang, S. Gilad-Gutnick, E. Ehrenberg, A. Yonas, S. Diamond, R. Held, and P. Sinha. Potential downside of high initial visual acuity. *Proceedings of the National Academy of Sciences*, 2018.
- [22] J. Wang and L. Perez. The effectiveness of data augmentation in image classification using deep learning. *Convolutional Neural Networks Vis. Recognit*, page 11, 2017.
- [23] M. D. Zeiler and R. Fergus. Visualizing and understanding convolutional networks. In *ECCV*, 2014.
- [24] B. Zhou, Y. Sun, D. Bau, and A. Torralba. Revisiting the importance of individual units in cnns via ablation. *NeurIPS*, 2018.

A Datasets and Networks

Datasets. We use two datasets: FaceScrub [13] and ImageNet [4]. FaceScrub is a face recognition dataset containing over 50,000 images of 388 individuals. ImageNet is an object recognition dataset with around 1.4 million images of 1000 object categories.

When increasing the number of *seen-transformed* categories, we use the same *seen-transformed* categories before the increase and an additional set of categories randomly selected. For the smallest number of *seen-transformed* categories, all them are randomly selected. This procedure facilitates comparing results between different number of *seen-transformed* categories, as it minimizes the dependency of the selected *seen-transformed* categories.

Networks. For FaceScrub, we used a modified AlexNet [12] architecture with five convolutional layers followed by two fully connected layers. The modifications included removal of the max-pooling layers after the first two convolutional layers in order to preserve necessary resolution for face recognition. The rest of the network was exactly as AlexNet (*i.e.* local response normalization, dropout, initialization parameters were the same as in AlexNet). The networks were trained for a maximum of 45 epochs (sufficient for convergence in all cases) with a learning rate of $1e-4$ and a weight decay of $5e-4$ for all transformations. These values were selected via grid search over 25 possible combinations of learning rates and weight decay, evaluated on 10% of a held-out images of the training set (all transformations lead to the same hyperparameters). The optimization algorithm was stochastic gradient descent with a momentum of 0.9 and a batch size of 32.

For ImageNet, we trained Resnet18s [9] using the official tensorflow implementation and with the hyperparameters that come by default.

Preprocessing. All input images were standardized before the first hidden layer. This did not change the accuracy of the network and it was particularly important to facilitate robustness to transformations without the need of adjusting the mean and standard deviation parameters of the batch normalization to the transformation. Before the standardization, the transformations were applied to the image, whose pixel values were between 0 and 255.

B Transformations

The transformations were applied to the dataset as an input layer in the GPU, such that the transformations were computed at runtime with every forward pass.

For color transformation, the hue was rotated 180° in the hue space. The grayscale operation was performed using built-in TensorFlow functions.

For blur transformation, a gaussian filter was convolved with the input image. The standard deviation of the gaussian was selected such that the top-1 accuracy of a network trained on the untransformed FaceScrub dataset was less 10%. To do so, we increased the standard deviation in steps of 0.5 until the accuracy was less than 10%, and then we stopped. The horizontal and vertical filters have a filter size of 3×3 of form $[-1, 0, 1]$. The high-pass filter has a filter size of 5×5 of form $[-1, 2, 4, 2, -1]$ in both directions.

Noise based transformations were applied by adding white noise and salt and pepper noise to the original images. The standard deviation of noise was selected with the same procedure as the blur transformation. For the white noise the standard deviation was increase in steps of 25. The salt and pepper noise randomly sets half of the pixels to either 0 or 255.

Geometrical transformations rotated the input images anti-clockwise with angles 90° and 180° , respectively. For thinning, the image width was reduced to its half and padded with zeros on both sides.

In ImageNet, the transformations were applied with parameters such that the top-1 accuracy was under 15%, using the aforementioned procedure to set the parameters of the transformations.

C Additional Results

Confusion Matrices. In Fig. 6, we show the confusion matrix between *seen-* and *unseen-transformed* categories. The table reports the percentage of images of one set of categories classified as the same or another set.

Active Neurons. To validate that the activity of the active neurons, *i.e.* neurons with activity higher than τ , capture all the activity relevant for the network’s robustness, we performed two types of ablation experiments: set all activity below the threshold to zero and set all the activity above the threshold to zero. The performance of the DCNN in these two conditions is shown in Fig. 7.

Invariance. In Fig. 8 and 9, we display the invariance for the different transformations and number of categories *seen-transformed* in FaceScrub and ImageNet, respectively. Fig. 10 shows the amount of invariance at each layer.

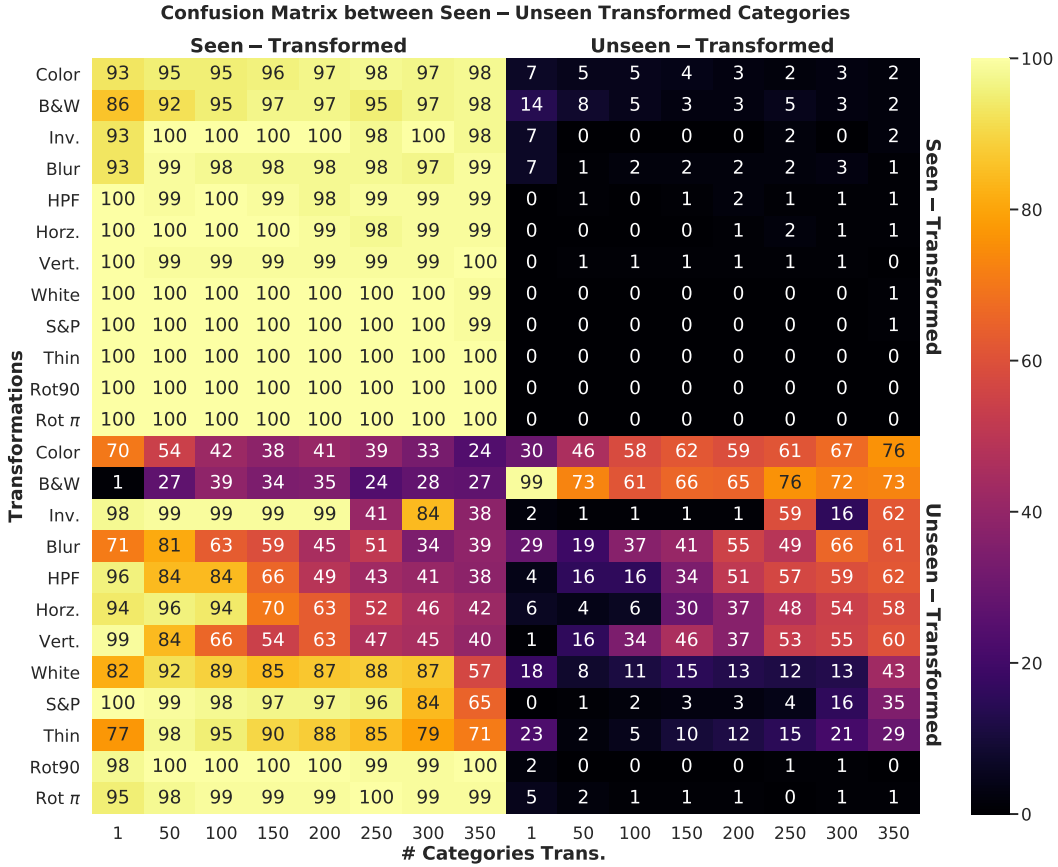


Figure 6: Confusion Matrix Between Seen- and Unseen-Transformed Categories for FaceScrub for all Transformations at different number of Categories transformed during training. Quadrant I: Percentage of Predictions for Unseen-Transformed Categories miss-classified as Seen-Transformed. Quadrant II: Percentage of Predictions for Seen-Transformed Categories correctly classified as Seen-Transformed. Quadrant III: Percentage of Predictions for Seen-Transformed Categories miss-classified as Unseen-Transformed. Quadrant IV: Percentage of Predictions for Unseen-Transformed Categories correctly classified as Unseen-Transformed.




Figure 7: Accuracy of Ablating Active and Inactive Neurons. (a) and (b) Within-category accuracy after ablating inactive (left) and active (right) neurons for FaceScrub and ImageNet, respectively.

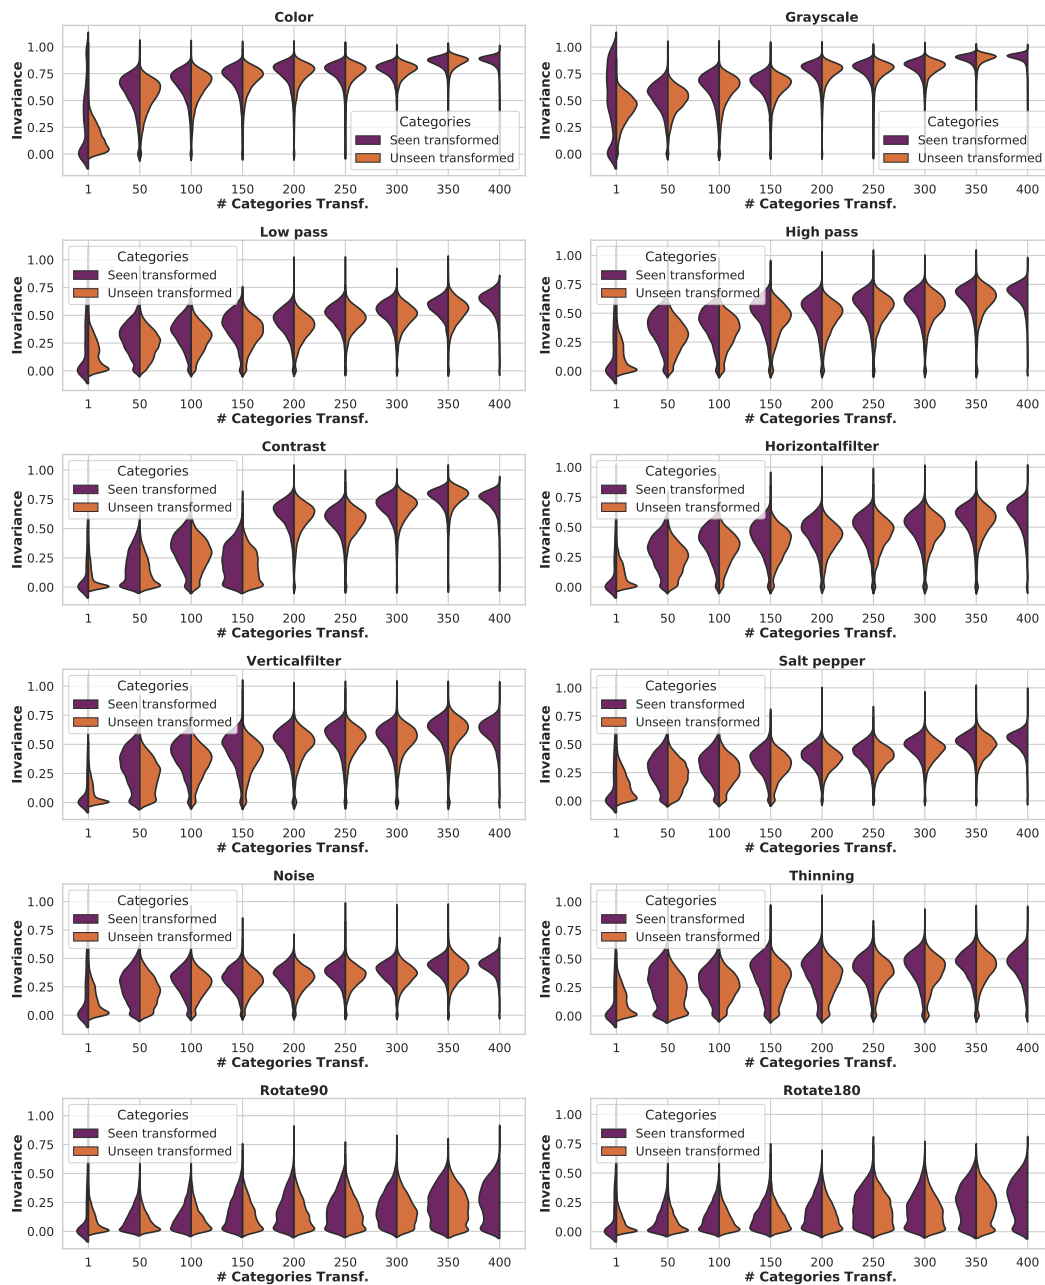


Figure 8: *Invariance for Different Number of Seen-transformed Categories.* Violin plots of the invariance coefficient among neurons at the penultimate layer for FaceScrub. Different number of *seen-transformed* categories are displayed and the invariance coefficient is reported separately for *seen-* and *unseen-transformed* categories.

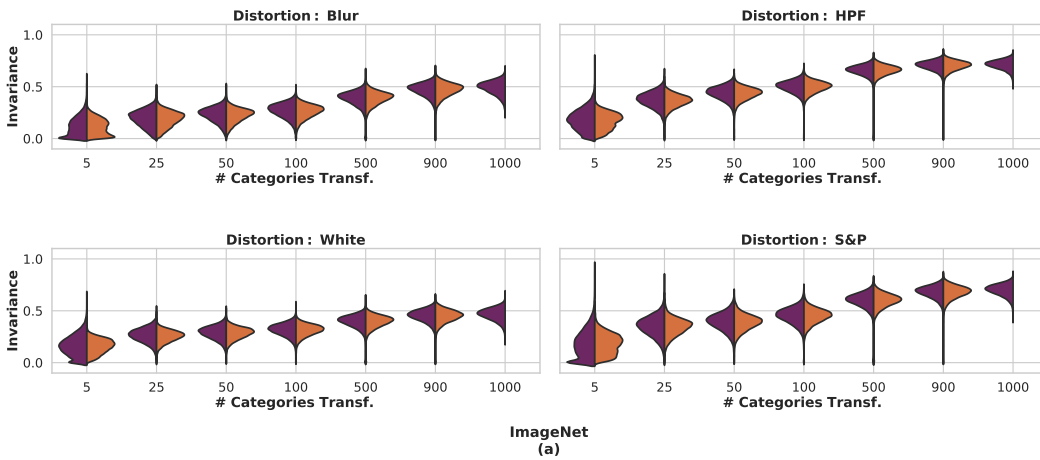


Figure 9: *Invariance for Different Number of Seen-transformed Categories.* Violin plots of the invariance coefficient among neurons at the penultimate layer for ImageNet. Different number of *seen-transformed* categories are displayed and the invariance coefficient is reported separately for *seen-* and *unseen-transformed* categories.

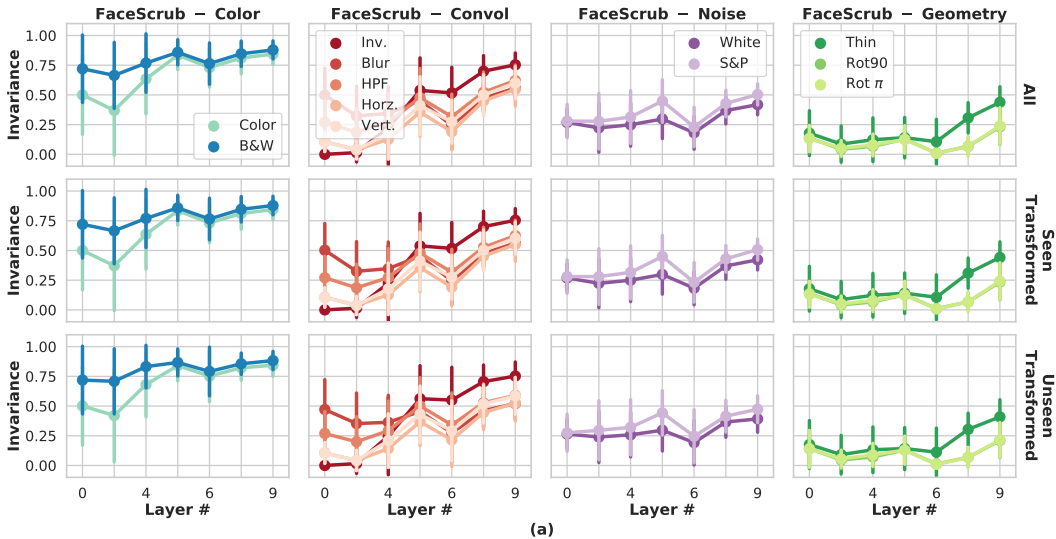


Figure 10: *Layer-wise Invariance.* Mean amount of invariance for different number of *seen-transformed* categories and transformations. Each row reports separately *seen-* and *unseen-transformed* categories, or both.