THE LIE DERIVATIVE FOR MEASURING LEARNED EQUIVARIANCE

Anonymous authors
Paper under double-blind review

ABSTRACT

Equivariance guarantees that a model's predictions capture key symmetries in data. When an image is translated or rotated, an equivariant model's representation of that image will translate or rotate accordingly. The success of convolutional neural networks has historically been tied to translation equivariance directly encoded in their architecture. The rising success of vision transformers, which have no explicit architectural bias towards equivariance, challenges this narrative and suggests that augmentations and training data might also play a significant role in their performance. In order to better understand the role of equivariance in recent vision models, we apply the Lie derivative, a method for measuring equivariance with strong mathematical foundations and minimal hyperparameters. Using the Lie derivative, we study the equivariance properties of hundreds of pretrained models, spanning CNNs, transformers, and Mixer architectures. The scale of our analysis allows us to separate the impact of architecture from other factors like model size or training method. Surprisingly, we find that many violations of equivariance can be linked to spatial aliasing in ubiquitous network layers, such as pointwise non-linearities, and that as models get larger and more accurate they tend to display more equivariance, regardless of architecture. For example, transformers can be more equivariant than convolutional neural networks after training.

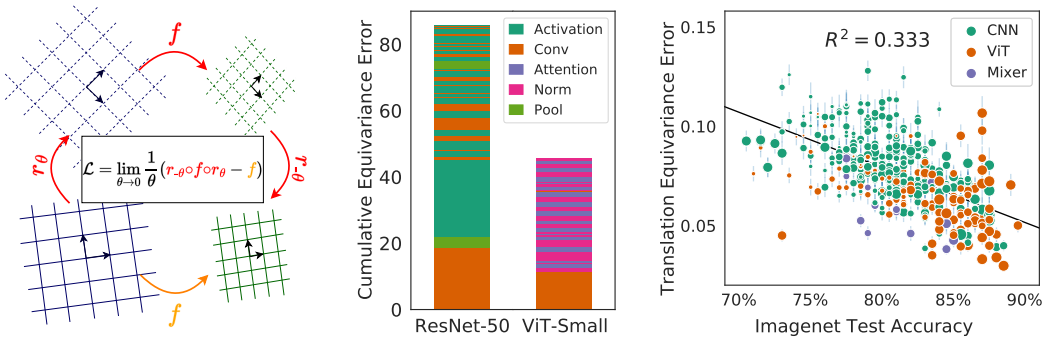

Figure 1: **(Left)**: The Lie derivative measures the equivariance of a function under continuous transformations, here rotation. **(Center)**: Using the Lie derivative, we quantify how much each layer contributes to the equivariance error of a model. Our analysis highlights surprisingly large contributions from non-linearities, which affects both CNNs and ViT architectures. **(Right)**: Translation equivariance as measured by the Lie derivative correlates with generalization in classification models, across convolutional and non-convolutional architectures. Although CNNs are often noted for their intrinsic translation equivariance, ViT and Mixer models are often more translation equivariant than CNN models after training.

1 INTRODUCTION

Symmetries allow machine learning models to generalize properties of one data point to the properties of an entire class of data points. A model that captures translational symmetry, for example, will have the same output for an image and a version of the same image shifted a half pixel to the left or right. If a classification model produces dramatically different predictions as a result of translation by half a

pixel or rotation by a few degrees it is likely misaligned with physical reality. Equivariance provides a formal notion of consistency under transformation. A function is equivariant if symmetries in the input space are preserved in the output space.

Baking equivariance into models through architecture design has led to breakthrough performance across many data modalities, including images (Cohen & Welling, 2016; Veeling et al., 2018), proteins (Jumper et al., 2021) and atom force fields (Batzner et al., 2022; Frey et al., 2022). In computer vision, translation equivariance has historically been regarded as a particularly compelling property of convolutional neural networks (CNNs) (LeCun et al., 1995). Imposing equivariance restricts the size of the hypothesis space, reducing the complexity of the learning problem and improving generalization (Goodfellow et al., 2016).

In most neural networks classifiers, however, true equivariance has been challenging to achieve, and many works have shown that model outputs can change dramatically for small changes in the input space (Azulay & Weiss, 2018; Engstrom et al., 2018; Vasconcelos et al., 2021; Ribeiro & Schön, 2021). Several authors have significantly improved the equivariance properties of CNNs with architectural changes inspired by careful signal processing (Zhang, 2019; Karras et al., 2021), but non-architectural mechanisms for encouraging equivariance, such as data augmentations, continue to be necessary for good generalization performance (Wightman et al., 2021).

The increased prominence of non-convolutional architectures, such as vision transformers (ViTs) and mixer models, simultaneously demonstrates that explicitly encoding architectural biases for equivariance is not necessary for good generalization in image classification, as ViT models perform on-par with or better than their convolutional counterparts with sufficient data and well-chosen augmentations (Dosovitskiy et al., 2020; Tolstikhin et al., 2021). Given the success of large flexible architectures and data augmentations, it is unclear what clear practical advantages are provided by explicit architectural constraints over learning equivariances from the data and augmentations. Resolving these questions systemically requires a unified equivariance metric and large-scale evaluation.

In what follows, we introduce the Lie derivative as a tool for measuring the equivariance of neural networks under continuous transformations. The *local equivariance error* (LEE), constructed with the Lie derivative, makes it possible to compare equivariance across models and to analyze the contribution of each layer of a model to its overall equivariance. Using LEE, we conduct a large-scale analysis of hundreds of image classification models. The breadth of this study allows us to uncover a novel connection between equivariance and model generalization, and the surprising result that ViTs are often more equivariant than their convolutional counterparts after training. To explain this result, we use the layer-wise decomposition of LEE to demonstrate how common building block layers shared across ViTs and CNNs, such as pointwise non-linearities, frequently give rise to aliasing and violations of equivariance.

We make our code publicly available at https://anonymous.4open.science/r/lie-deriv-6784

2 BACKGROUND

Groups and equivariance Equivariance provides a formal notion of consistency under transformation. A function $f : V_1 \to V_2$ is equivariant to transformations from a symmetry group G if applying the symmetry to the input of f is the same as applying it to the output

$$\forall g \in G : \quad f(\rho_1(g)x) = \rho_2(g)f(x), \tag{1}$$

where $\rho(g)$ is the *representation* of the group element, which is a linear map $V \to V$.

The most common example of equivariance in deep learning is the translation equivariance of convolutional layers: if we translate the input image by an integer number of pixels in x and y, the output is also translated by the same amount, ignoring the regions close to the boundary of the image. Here $x \in V_1 = V_2$ is an image and the representation $\rho_1 = \rho_2$ expresses translations of the image. The translation *invariance* of certain neural networks is also an expression of the equivariance property, but where the output vector space V_2 has the trivial $\rho_2(g) = I$ representation, such that model outputs are unaffected by translations of the inputs. Equivariance is therefore a much richer framework, in which we can reason about representations at the input and the output of a function.

Continuous signals The inputs to classification models are discrete images sampled from a *continuous* reality. Although discrete representations are necessary for computers, the goal of classification models should be learning functions that generalize in the real world. It is therefore useful to consider an image as a function $h : \mathbb{R}^2 \to \mathbb{R}^3$ rather than a discrete set of pixel values and broaden the symmetry groups we might consider, such as translations of an image by vector $\boldsymbol{b} \in \mathbb{R}^2$, rather than an integer number of pixels.

Fourier analysis is a powerful tool for understanding the relationship between continuous signals and discrete samples by way of frequency decompositions. Any $M \times M$ image, $h(\boldsymbol{x})$, can be constructed from its frequency components, H_{nm}, using a $2d$ Fourier series, $h(\boldsymbol{x}) = \frac{1}{2\pi} \sum_{n,m} H_{nm} e^{2\pi i \boldsymbol{x} \cdot [n,m]}$ where $\boldsymbol{x} \in [0,1]^2$ and $n, m \in$ [-$M/2$, -$M/2+1$, ..., $M/2$], the bandlimit defined by the image size.

Aliasing Aliasing occurs when sampling at a limited frequency f_s, for example the size of an image in pixels, causes high frequency content (above the Nyquist frequency $f_s/2$) to be converted into spurious low frequency content. Content with frequency n is observed as a lower frequency contribution at frequency

$$\mathrm{Alias}(n) = \left\{ \begin{array}{ll} n \bmod f_s & \text{if } (n \bmod f_s) < f_s/2 \\ (n \bmod f_s) - f_s & \text{if } (n \bmod f_s) > f_s/2 \end{array} \right\}. \tag{2}$$

If a discretely sampled signal such as an image is assumed to have no frequency content higher than f_s, then the continuous signal can be uniquely reconstructed using the Fourier series and have a consistent continuous representation. But if the signal contains higher frequency content which gets aliased down by the sampling, then there is an ambiguity and exact reconstruction is not possible.

Aliasing and equivariance Aliasing is critically important to our study because it breaks equivariance to continuous transformations like translation and rotation. When a continuous image is translated its Fourier components pick up a phase:

$$h(\boldsymbol{x}) \mapsto h(\boldsymbol{x} - \boldsymbol{b}) \implies H_{nm} \mapsto H_{nm} e^{-2\pi i \boldsymbol{b} \cdot [n,m]}.$$

However, when an aliased signal is translated, the aliasing operation A introduces a scaling factor:

$$H_{nm} \mapsto H_{nm} e^{-2\pi i (b_0 \mathrm{Alias}(n) + b_1 \mathrm{Alias}(m))}$$

In other words, aliasing causes a translation *by the wrong amount*: the frequency component H_{nm} will effectively be translated by $[(\mathrm{Alias}(n)/n)b_0, (\mathrm{Alias}(m)/m)b_1]$ which may point in a different direction than $\boldsymbol{b}$, and potentially even the opposite direction. Applying shifts to an aliased image will yield the correct shifts for true frequencies less than the Nyquist but incorrect shifts for frequencies higher than the Nyquist. Other continuous transformations, like rotation, create similar asymmetries.

Many common operations in CNNs can lead to aliasing in subtle ways, breaking equivariance in turn. Zhang (2019), for example, demonstrates that downsampling layers causes CNNs to have inconsistent outputs for translated images. The underlying cause of the invariance is aliasing, which occurs when downsampling alters the high frequency content of the network activations. The $M \times M$ activations at a given layer of a convolutional network have spatial Nyquist frequencies $f_s = M/2$. Downsampling halves the size of the activations and corresponding Nyquist frequencies. The result is aliasing of all nonzero content with $n \in [M/4, M/2]$. To prevent this aliasing, Zhang (2019) uses a local low pass filter (Blur-Pool) to directly remove the problematic frequency regions from the spectrum.

While studying generative image models, Karras et al. (2021) unearth a similar phenomenon in the pointwise nonlinearities of CNNs. Imagine an image at a single frequency $h(\boldsymbol{x}) = \sin(2\pi \boldsymbol{x} \cdot [n, m])$. Applying a nonlinear transformation to h creates new high frequencies in the Fourier series, as illustrated in Figure 2. These high frequencies may fall outside of the bandlimit, leading to aliasing. To counteract this effect, Karras et al. (2021) opt for smoother non-linearities and perform upsampling before calculating the activations.

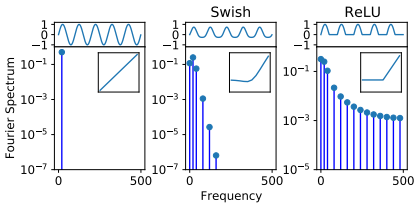

Figure 2: Non-linearities generate new high-frequency harmonics.

3 Related Work

While many papers propose architectural changes to improve the equivariance of CNNs (Zhang, 2019; Karras et al., 2021; Weiler & Cesa, 2019), others focus purely on measuring and understanding how equivariance can emerge through learning from the training data (Lenc & Vedaldi, 2015). Olah et al. (2020), for example, studies learned equivariance in CNNs using model inversions techniques. While they uncover several fascinating properties, such as rotation equivariance that emerges naturally without architectural priors, their method is limited by requiring manual inspection of neuron activations. Most relevant to our work, Bouchacourt et al. (2021) measure equivariance in CNNs and ViTs by sampling group transformations. Parallel to our findings, they conclude that data and data augmentation play a larger role than architecture in the ultimate equivariance of a model. Because their study is limited to just a single ResNet and ViT architecture, however, they do not uncover the broader relationship between equivariance and generalization that we show here.

Many papers introduce consistency metrics based on sampling group transformations (Zhang, 2019; Karras et al., 2021; Bouchacourt et al., 2021), but most come with significant drawbacks. When translations are sampled with an integer number of pixels (Zhang, 2019; Bouchacourt et al., 2021), aliasing effects will be completely overlooked. As a remedy, (Karras et al., 2021) propose a subpixel translation equivariance metric (EQ-T_{frac}) that appropriately captures aliasing effects. While this metric is a major improvement, it requires many design decisions not required by LEE, which has relatively few hyperparameters and seamlessly breaks down equivariance across layers. Relative to these other approaches, LEE offers a unifying perspective, with significant theoretical and practical benefits.

4 Measuring local equivariance error with Lie derivatives

Lie Derivatives The Lie derivative gives a general way of evaluating the degree to which a function f violates a symmetry. To define the Lie derivative, we first consider how a symmetry group element can transform a function by rearranging Equation 1:

$$\rho_{21}(g)[f](x) = \rho_2(g)^{-1} f(\rho_1(g)x)\,.$$

The resulting linear operator, $\rho_{21}(g)[\cdot]$, acts on the vector space of functions, and $\rho_{21}(g)[f] = f$ if the function is equivariant. Every continuous symmetry group (Lie group), G, has a corresponding vector space (Lie algebra) $\mathfrak{g} = \text{Span}(\{X_i\}_{i=1}^d)$, with basis elements X_i that can be interpreted as vector fields $\mathbb{R}^n \to \mathbb{R}^n$. For images, these vector fields encode infinitesimal transformations $\mathbb{R}^2 \to \mathbb{R}^2$ over the domain of continuous image signals $f : \mathbb{R}^2 \to \mathbb{R}^k$. One can represent group elements $g \in G$ (which lie in the connected component of the identity) as the *flow* along a particular vector field Φ_Y^t, where $Y = \sum_i a_i X_i$ is a linear combination of basis elements. The flow $\Phi_Y^t(x_0)$ of a point x_0 along a vector field Y by value t is defined as the solution to the ODE $\frac{dx}{dt} = Y(x)$ at time t with initial value x_0. The flow Φ_Y^t smoothly parameterizes the group elements by t so that the operator $\rho_{21}(\Phi_Y^t)[\cdot]$ connects changes in the space of group elements to changes in symmetry behavior of a function.

The Lie derivative of the function f is the derivative of the operator $\rho_{21}(g)$ at $g = \text{Identity} = \Phi_0$ along a particular vector field Y,

$$\mathcal{L}_Y(f) = \lim_{t\to 0} \frac{1}{t}(\rho_{21}(\Phi_Y^t)[f] - f) = \left.\frac{d}{dt}\right|_{t=0} \rho_{21}(\Phi_Y^t)[f]. \tag{3}$$

Intuitively, the Lie derivative measures the sensitivity of a function to infinitesimal symmetry transformations. This local definition of equivariance error is related to the typical global notion of equivariance error. As we derive in Appendix B.1, if $\forall i = 1, ..., d : \mathcal{L}_{X_i}(f) = 0$ (and the exponential map is surjective) then $\forall g \in G :\ f(\rho_1(g)x) = \rho_2(g)f(x)$ and for all x in the domain, and vice versa. In practice, the Lie derivative is only a proxy for strict global equivariance. We note global equivariance includes radical transformations like translation by 75% of an image, which is not necessarily desirable. In section 6 we show that our local formulation of the Lie derivative can capture the effects of many practically relevant transformations.

The Lie derivative of a function with multiple outputs will also have multiple outputs, so if we want to summarize the equivariance error with a single number, we can compute the norm of the Lie

```python
import torch.nn.functional as F
from torch.autograd.functional import jvp

def rotate(imgs, theta):
    """ Rotate images by angle theta and interpolate"""
    m = [[torch.cos(theta), torch.sin(theta), 0],
         [-torch.sin(theta), torch.cos(theta), 0]]
    m = torch.tensor(m)[None].expand(imgs.shape[0], -1, -1)
    return F.grid_sample(imgs, F.affine_grid(m, imgs.size(), True))

def rotation_lie_deriv(model,imgs):
    """ Lie deriv. of model w.r.t. rotation, can be scalar/image"""
    def rotated_model(theta):
        z = model(rotate(imgs,theta))
        img_like = (len(z.shape) == 4) # more complex for ViT/Mixer
        return rotate(z,-theta) if img_like else z
    return jvp(rotated_model, torch.zeros(1,requires_grad=True))[-1]

def e_lee(model,imgs):
    """ Expected equiv. error (E[|Lf|^2]/d_out) w.r.t. rotation"""
    return rotation_lie_deriv(model, imgs).pow(2).mean()
```

Figure 3: Lie derivatives can be computed using automatic differentiation. We show how a Lie derivative for continuous rotations can be implemented in PyTorch (Paszke et al., 2019). The implementation in our experiments differs slightly, for computational efficiency and to pass second-order gradients through *grid_sample*.

derivative scaled by the size of the output. Taking the average of the Lie derivative over the data distribution, we define Local Equivariance Error (LEE),

$$\mathrm{LEE}(f) = \mathbb{E}_{x\sim\mathcal{D}} \|\mathcal{L}_X f(x)\|^2 / \dim(V_2). \tag{4}$$

We provide a Python implementation of the Lie derivative calculation for rotations in Figure 3 as an example. Mathematically, LEE also has an appealing connection to consistency regularization (Athiwaratkun et al., 2018), which we discuss in Appendix C.1.

Layerwise Decomposition of Lie Derivative Unlike alternative equivariance metrics, the Lie derivative decomposes naturally over the layers of a neural network, since it satisfies the chain rule. As we show in Appendix B.2, the Lie derivative of the composition of two functions $h : V_1 \to V_2$ and $f : V_2 \to V_3$ satisfies

$$\mathcal{L}_X(f \circ h)(x) = (\mathcal{L}_X f)(h(x)) + df|_{h(x)}(\mathcal{L}_X h)(x), \tag{5}$$

where $df|_{h(x)}$ is the Jacobian of f at $h(x)$ which we will abbreviate as df. Note that this decomposition captures the fact that intermediate layers of the network may transform in their own way:

$$f(h(x)) \mapsto \rho_{31}(g)[f \circ g](x) = \rho_3(g)^{-1} f(\rho_2(g)\rho_2(g)^{-1} h(\rho_1(g)x)) = \rho_{32}(g)[f] \circ \rho_{21}(g)[h](x)$$

and the Lie derivatives split up accordingly.

Applying this property to an entire model as a composition of layers $\mathrm{NN}(x) = f_{N:1}(x) := f_N(f_{N-1}(...(f_1(x))))$, we can identify the contribution that each layer f_i makes to the equivariance error of the whole. Unrolling Equation 5, we have

$$\mathcal{L}_X(\mathrm{NN}) = \sum_{i=1}^{N} df_{N:i+1} \mathcal{L}_X f_i. \tag{6}$$

Intuitively, the equivariance error of a sequence of layers is determined by the sum of the equivariance error for each layer multiplied by the degree to which that error is attenuated or magnified by the other layers (as measured by the Jacobian). We evaluate the norm of each of the contributions $df_{N:i+1}\mathcal{L}_X f_i$ to the (vector) equivariance error $\mathcal{L}_X(\mathrm{NN})$ which we compute using autograd and stochastic trace estimation, as we describe in Appendix B.3. Importantly, the sum of norms may differ from the norm of the sum, but this analysis allows us to identify patterns across layers and pinpoint operations that contribute most to equivariance error.

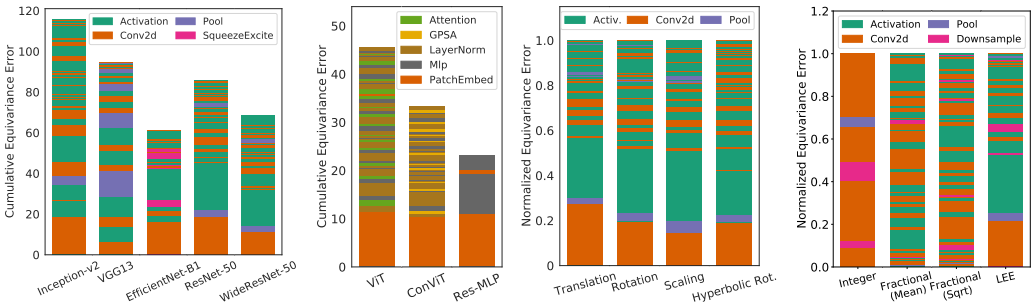

Figure 4: Contributions to equivariance shown cumulatively by layer, in the order the layers occur in the network. **Left**: Convolutional architectures. In all the CNNs, much of the equivariance error comes from downsampling and non-linearities. **Middle-Left**: Non-convolutional architectures. The initial patch embedding, a strided convolution, is the largest contributor for the ViTs and Mixers. The rest of the error is distributed uniformly across other nonlinear operations. **Middle-Right**: ResNet-50 across different transformations as a percentage. Despite being designed for translation equivariance, the fraction of equivariance error produced by each layer is almost identical for other affine transformations, suggesting that aliasing is the primary source of equivariance error. **Right**: Comparing LEE with alternative metrics for translation equivariance. Using integer translations misses key contributors to equivariance errors, such as activations, while using fractional translations can lead to radically different outcomes depending on choice of normalization (N or $\sqrt{N}$). LEE captures aliasing effects and has minimal design decisions.

5 Layerwise Equivariance Error

As described in section 3, subtle architectural details often prevent models from being perfectly equivariant. Aliasing can result from careless downsampling or from an activation function with a wide spectrum. In this section, we explore how the Lie derivative uncovers these types of effects automatically, across several popular architectures. We evaluate the equivariance of pretrained models on 100 images from the ImageNet (Deng et al., 2009) test set.

Using the layerwise analysis, we can dissect the sources of translation equivariance error in convolutional and non-convolutional networks as shown in Figure 4 (left) and (middle-left). For the Vision Transformer and Mixer models, we see that the initial conversion from image to patches produces a significant portion of the error, and the remaining error is split uniformly between the other nonlinear layers: LayerNorm, tokenwise MLP, and self-attention. The contribution from these nonlinear layers is seldom recognized and potentially counterintuitive, until we fully grasp the deep connection between equivariance and aliasing. In Figure 4 (middle-right), we show that this breakdown is strikingly similar for other image transformations like rotation, scaling, and hyperbolic rotations, providing evidence that the cause of equivariance error is not specific to translations but is instead a general culprit across a whole host of continuous transformations that can lead to aliasing.

We can make the relationship between aliasing and equivariance error precise by considering the aliasing operation Alias defined in Equation 2.

Theorem 1. *For translations along the vector $v = [v_x, v_y]$, the aliasing operation A introduces a translation equivariance error of*

$$\|\mathcal{L}_v(A)(h)\|^2 = (2\pi)^2 \sum_{n,m} H_{nm}^2 \big(v_x^2(\mathrm{Alias}(n) - n)^2 + v_y^2(\mathrm{Alias}(m) - m)^2\big),$$

where $h(\boldsymbol{x}) = \frac{1}{2\pi}\sum_{n,m} H_{nm} e^{2\pi i \boldsymbol{x}\cdot[n,m]}$ is the Fourier series for the input image h.

We provide the proof in Appendix C.2. The connection between aliasing and LEE is important because aliasing is often challenging to identify despite being ubiquitous (Zhang, 2019; Karras et al., 2021). Aliasing in non-linear layers impacts all vision models and is thus a key factor in any fair comparison of equivariance.

As alternative equivariance metrics exist, it is natural to wonder whether they can also be used for layerwise analysis. In Figure 4 (right), we show how two equivariance metrics from Karras et al. (2021)

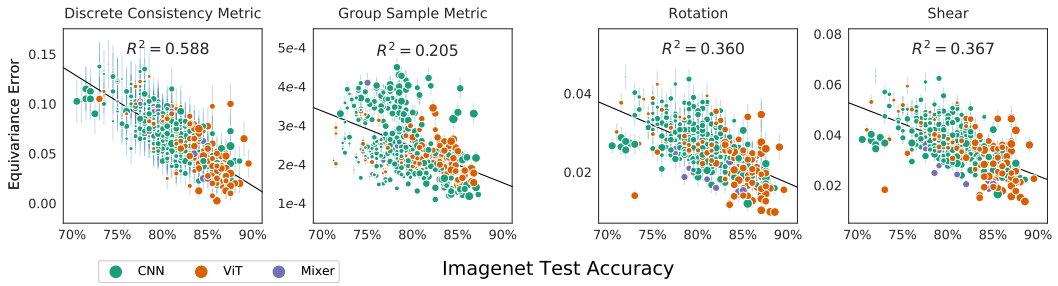

Figure 5: Equivariance metrics evaluated on the ImageNet test set. **Left:** Non-LEE equivariance metrics display similar trends to Figure 1, despite using larger, multi-pixel transformations. **Right:** Norm of rotation and shear Lie derivatives. Across all architectures, models with strong generalization become more equivariant to many common affine transformations. Marker size indicates model size. Error bars show one standard error over test set images used in the equivariance calculation.

compare with LEE, highlighting notable drawbacks. (1) Integer translation equivariance completely ignores aliasing effects, which are captured by both LEE and fractional translations. (2) Though fractional translation metrics (EQ-T_{frac}) correctly capture aliasing, comparing the equivariance of layers with different resolutions ($C \times H \times W$) requires an arbitrary choice of normalization. This choice can lead to radically different outcomes in the perceived contribution of each layer and is not required when using LEE, which decomposes across layers as described in section 4. We provide a detailed description of the baselines in Appendix D.1.

6 Trends in Learned Equivariance

Methodology We evaluate the Lie derivative of many popular classification models under transformations including $2d$ translation, rotation, and shearing. We define continuous transformations on images using bilinear interpolation with reflection padding. In total, we evaluate 410 classification models, a collection comprising popular CNNs, vision transformers, and MLP-based architectures (Wightman, 2019). Beyond diversity in architectures, there is also substantial diversity in model size, training recipe, and the amount of data used for training or pretraining. This collection of models therefore covers many of the relevant axes of variance one is likely to consider in designing a system for classification. We include an exhaustive list of models in the Appendix D.3.

Equivariance across architectures As shown in Figure 1 (right), the translation equivariance error (Lie derivative norm) is strongly correlated with the ultimate test accuracy that the model achieves. Surprisingly, despite convolutional architectures being motivated and designed for their translation equivariance, we find no significant difference in the equivariance achieved by convolutional architectures and the equivariance of their more flexible ViT and Mixer counterparts when conditioning on test accuracy. This trend also extends to rotation and shearing transformations, which are common in data augmentation pipelines (Cubuk et al., 2020) (in Figure 5 (right)). Additional transformation results included in Appendix D.5.

For comparison, we also evaluate the same set of models using two alternative equivariance metrics: prediction consistency under discrete translation (Zhang, 2019) and expected equivariance under group samples (Finzi et al., 2020; Hutchinson et al., 2021), which is similar in spirit to EQ-T_{frac} (Karras et al., 2021) (exact calculations in Appendix D.4). Crucially, these metrics are slightly less *local* than LEE, as they evaluate equivariance under transformations of up to 10 pixels at a time. The fact that we obtain similar trends highlights LEE's relevance beyond subpixel transformations.

Effects of Training and Scale In section 3 we described many architectural design choices that have been used to improve the equivariance of vision models, for example Zhang (2019)'s Blur-Pool low-pass filter. We now investigate how equivariance error can be reduced with non-architectural design decisions, such as increasing model size, dataset size, or training method. Surprisingly, we show that equivariance error can often be significantly reduced without any changes in architecture.

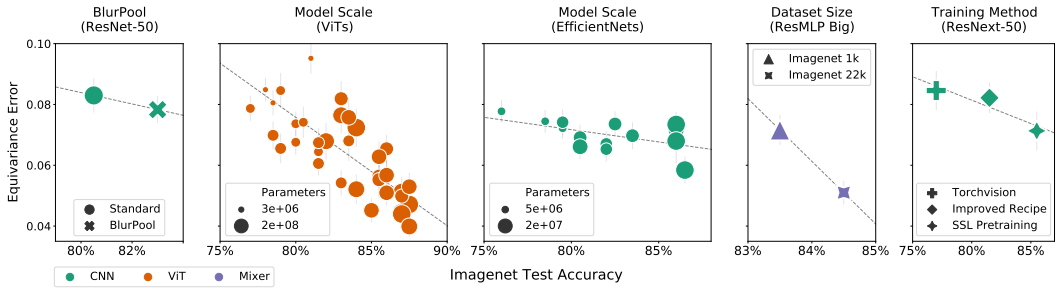

Figure 6: Case studies in decreasing translational equivariance error, numbered left-to-right. **1**: Blur-Pool (Zhang, 2019), an architectural change to improve equivariance, decreases the equivariance error but by less than can be accomplished by improving the training recipe or increasing the scale of model or dataset. **2-3**: Increasing the number of parameters for a fixed model family (here ViTs (El-Nouby et al., 2021) and EfficientNets (Tan & Le, 2019a)). **4**: Increasing the training dataset size for a ResMLP Big (Touvron et al., 2021a) model. **5**: Changing the training recipe for ResNeXt-50 (Xie et al., 2017) with improved augmentations (Wightman et al., 2021) or SSL pretraining (Yalniz et al., 2019). Error bars show one standard error over images in the Lie derivative calculation.

In Figure 6, we show slices of the data from Figure 1 along a shared axis for equivariance error. As a point of comparison, in Figure 6 (left), we show the impact of the Blur-Pool operation discussed above on a ResNet-50 (Zhang, 2019). In the accompanying four plots, we show the effects of increasing model scale (for both ViTs and CNNs), increasing dataset size, and finally different training procedures. Although Zhang (2019)'s architectural adjustment does have a noticeable effect, factors such as dataset size, model scale, and use of modern training methods, have a much greater impact on learned equivariance.

As a prime example, in Figure 6 (right), we show a comparison of three training strategies for ResNeXt-50 – an architecture almost identical to ResNet-50. We use Wightman et al. (2021)'s pretrained model to illustrate the role of an improved training recipe and Mahajan et al. (2018b)'s semi-supervised model as an example of scaling training data. Notably, for a fixed architecture and model size, these changes lead to decreases in equivariance error on par with architectural interventions (BlurPool). This result is surprising when we consider that Wightman et al. (2021)'s improved training recipe benefits significantly from Mixup (Zhang et al., 2017) and CutMix (Yun et al., 2019), which have no obvious connection to equivariance. Similarly, Mahajan et al. (2018b)'s semi-supervised method has no explicit incentive for equivariance.

Equivariance out of distribution From our analysis above, large models appear to learn equivariances that rival architecture engineering in the classification setting. When learning equivariances through data augmentation, however, there is no guarantee that the equivariance will generalize to data that is far from the training distribution. Indeed, Engstrom et al. (2019) shows that carefully chosen translations or rotations can be as devastating to model performance as adversarial examples. We find that vision models do indeed have an *equivariance gap*: models are less equivariant on test data than train, and this gap grows for OOD inputs as shown in Figure 7. Notably, however, architectural biases do not have a strong effect on the equivariance gap, as both CNN and ViT models have comparable gaps for OOD inputs.

Why aren't CNNs more equivariant than ViTs? Given the deep historical connection between CNNs and equivariance, the results in Figure 5 and Figure 7 might appear counterintuitive. ViTs, CNNs, and Mixer have quite different inductive biases and therefore often learn very different representations of data (Raghu et al., 2021). Despite their differences, however, all of these architectures are fundamentally constructed from similar building blocks–such as convolutions, normalization layers, and non-linear activations which can all contribute to aliasing and equivariance error. Given this shared foundation, vision models with high capacity and effective training recipes are more capable of fitting equivariances already present in augmented training data.

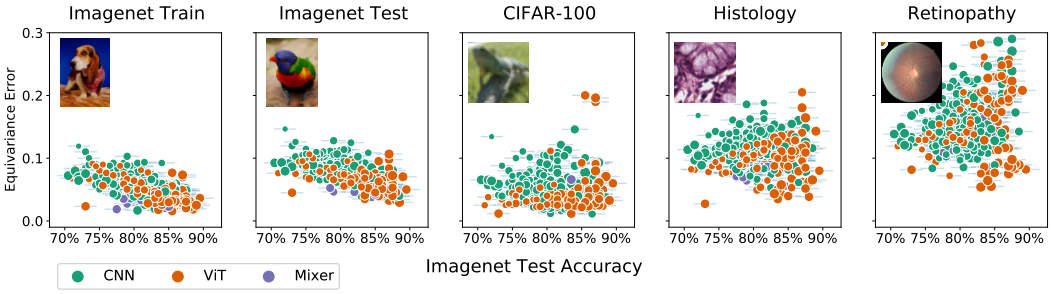

Figure 7: Models are less equivariant on test data and becoming decreasingly equivariant as the data moves away from the training manifold. As examples of data with similar distributions, we show equivariance error on the ImageNet train and test sets as well as CIFAR-100. As examples of out-of-distribution data, we use two medical datasets (which often use Imagenet pretraining), one for Histology (Kather et al., 2016) and one for Retinopathy (Kaggle & EyePacs, 2015).

Learning rotation equivariance We finally consider the extent to which large-scale pretraining can match strong architectural priors in a case where equivariance is obviously desirable. We fine-tune a state-of-the-art vision transformer model pretrained with masked autoencoding (He et al., 2021) for 100 epochs on rotated MNIST (Weiler & Cesa, 2019) (details in Appendix D.6). This dataset, which contains MNIST digits rotated uniformly between -180 and 180 degrees, is a common benchmark for papers that design equivariant architectures. In Table 1 we show the test errors for many popular architectures with strict equivariance constrainets alongside the error for our finetuned model. Surprisingly, the finetuned model achieves competitive test accuracy, in this case a strong proxy for rotation invariance. Despite having relatively weak architectural biases, transformers are capable of learning and generalizing on well on symmetric data.

Model	Test Error (%)
G-CNN [13]	2.28
H-NET [76]	1.69
ORN [86]	1.54
TI-Pooling [43]	1.2
Finetuned MAE	1.14
RotEqNet [51]	1.09
E(2)-CNN [73]	0.68

Table 1: Our finetuned MAE is competitive with several architectures explicitly engineered to encode rotation invariance on RotMNIST, where rotation invariance is clearly crucial to generalization.

7 Conclusion

We introduced a new metric for measuring equivariance which enables a nuanced investigation of how architecture design and training procedures affect representations discovered by neural networks. Using this metric we are able to pinpoint equivariance violation to individual layers, finding that pointwise nonlinearities contribute substantially even in networks that have been designed for equivariance. We argue that aliasing is the primary mechanism for how equivariance to continuous transformations are broken, which we support theoretically and empirically. We use our measure to study equivariance learned from data and augmentations, showing model scale, data scale, or training recipe can have a greater effect on the ability to learn equivariances than architecture.

Many of these results are contrary to the conventional wisdom. For example, transformers can be more equivariant than convolutional neural networks after training, and can learn equivariances needed to match the performance of specially designed architectures on benchmarks like rotated MNIST, despite a lack of explicit architectural constraints. These results suggest we can be more judicious in deciding when explicit interventions for equivariance are required, especially in many real world problems where we only desire approximate and local equivariances.On the other hand, explicit constraints will continue to have immense value when exact equivariances and extrapolation are required — such as rotation invariance for molecules. Moreover, despite the ability to learn equivariances on training data, we find that there is an equivariance *gap* on test and OOD data which persists regardless of the model class. Thus other ways of combating aliasing outside of architectural interventions may be the path forward for improving the equivariance and invariance properties of models.

REFERENCES

Ben Athiwaratkun, Marc Finzi, Pavel Izmailov, and Andrew Gordon Wilson. There are many consistent explanations of unlabeled data: Why you should average. *arXiv preprint arXiv:1806.05594*, 2018.

Haim Avron and Sivan Toledo. Randomized algorithms for estimating the trace of an implicit symmetric positive semi-definite matrix. *Journal of the ACM (JACM)*, 58(2):1–34, 2011.

Aharon Azulay and Yair Weiss. Why do deep convolutional networks generalize so poorly to small image transformations? *arXiv preprint arXiv:1805.12177*, 2018.

Hangbo Bao, Li Dong, and Furu Wei. Beit: Bert pre-training of image transformers. *arXiv preprint arXiv:2106.08254*, 2021.

Simon Batzner, Albert Musaelian, Lixin Sun, Mario Geiger, Jonathan P Mailoa, Mordechai Kornbluth, Nicola Molinari, Tess E Smidt, and Boris Kozinsky. E (3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature communications*, 13(1):1–11, 2022.

Irwan Bello, William Fedus, Xianzhi Du, Ekin D Cubuk, Aravind Srinivas, Tsung-Yi Lin, Jonathon Shlens, and Barret Zoph. Revisiting resnets: Improved training and scaling strategies. *arXiv preprint arXiv:2103.07579*, 2021.

Diane Bouchacourt, Mark Ibrahim, and Ari Morcos. Grounding inductive biases in natural images: invariance stems from variations in data. *Advances in Neural Information Processing Systems*, 34: 19566–19579, 2021.

Andrew Brock, Soham De, Samuel L Smith, and Karen Simonyan. High-performance large-scale image recognition without normalization. *arXiv preprint arXiv:2102.06171*, 2021.

Chun-Fu Chen, Quanfu Fan, and Rameswar Panda. Crossvit: Cross-attention multi-scale vision transformer for image classification. *arXiv preprint arXiv:2103.14899*, 2021.

Yunpeng Chen, Jianan Li, Huaxin Xiao, Xiaojie Jin, Shuicheng Yan, and Jiashi Feng. Dual path networks, 2017.

François Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1251–1258, 2017.

Xiangxiang Chu, Zhi Tian, Yuqing Wang, Bo Zhang, Haibing Ren, Xiaolin Wei, Huaxia Xia, and Chunhua Shen. Twins: Revisiting the design of spatial attention in vision transformers. *arXiv preprint arXiv:2104.13840*, 2021.

Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pp. 2990–2999. PMLR, 2016.

Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops*, pp. 702–703, 2020.

Zihang Dai, Hanxiao Liu, Quoc V Le, and Mingxing Tan. Coatnet: Marrying convolution and attention for all data sizes. *arXiv preprint arXiv:2106.04803*, 2021.

Stéphane d'Ascoli, Hugo Touvron, Matthew Leavitt, Ari Morcos, Giulio Biroli, and Levent Sagun. Convit: Improving vision transformers with soft convolutional inductive biases. *arXiv preprint arXiv:2103.10697*, 2021.

Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pp. 248–255. Ieee, 2009.

Xiaohan Ding, Xiangyu Zhang, Ningning Ma, Jungong Han, Guiguang Ding, and Jian Sun. Repvgg: Making vgg-style convnets great again. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 13733–13742, 2021.

Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.

Alaaeldin El-Nouby, Hugo Touvron, Mathilde Caron, Piotr Bojanowski, Matthijs Douze, Armand Joulin, Ivan Laptev, Natalia Neverova, Gabriel Synnaeve, Jakob Verbeek, et al. Xcit: Cross-covariance image transformers. *arXiv preprint arXiv:2106.09681*, 2021.

Logan Engstrom, Brandon Tran, Dimitris Tsipras, Ludwig Schmidt, and Aleksander Madry. A rotation and a translation suffice: Fooling cnns with simple transformations, 2018.

Logan Engstrom, Brandon Tran, Dimitris Tsipras, Ludwig Schmidt, and Aleksander Madry. Exploring the landscape of spatial robustness. In *International conference on machine learning*, pp. 1802–1811. PMLR, 2019.

Marc Finzi, Samuel Stanton, Pavel Izmailov, and Andrew Gordon Wilson. Generalizing convolutional neural networks for equivariance to lie groups on arbitrary continuous data. In *International Conference on Machine Learning*, pp. 3165–3176. PMLR, 2020.

Nathan Frey, Ryan Soklaski, Simon Axelrod, Siddharth Samsi, Rafael Gomez-Bombarelli, Connor Coley, and Vijay Gadepally. Neural scaling of deep chemical models, 2022.

Shanghua Gao, Ming-Ming Cheng, Kai Zhao, Xin-Yu Zhang, Ming-Hsuan Yang, and Philip HS Torr. Res2net: A new multi-scale backbone architecture. *IEEE transactions on pattern analysis and machine intelligence*, 2019.

Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016.

Jian Guo, He He, Tong He, Leonard Lausen, Mu Li, Haibin Lin, Xingjian Shi, Chenguang Wang, Junyuan Xie, Sheng Zha, Aston Zhang, Hang Zhang, Zhi Zhang, Zhongyue Zhang, Shuai Zheng, and Yi Zhu. Gluoncv and gluonnlp: Deep learning in computer vision and natural language processing. *Journal of Machine Learning Research*, 21(23):1–7, 2020. URL http://jmlr.org/papers/v21/19-429.html.

Brian C Hall. Lie groups, lie algebras, and representations. In *Quantum Theory for Mathematicians*, pp. 333–366. Springer, 2013.

Dongyoon Han, Sangdoo Yun, Byeongho Heo, and YoungJoon Yoo. Rethinking channel dimensions for efficient model design. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 732–741, 2021a.

Kai Han, An Xiao, Enhua Wu, Jianyuan Guo, Chunjing Xu, and Yunhe Wang. Transformer in transformer. *arXiv preprint arXiv:2103.00112*, 2021b.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.

Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. *arXiv:2111.06377*, 2021.

Tong He, Zhi Zhang, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of tricks for image classification with convolutional neural networks. *arXiv preprint arXiv:1812.01187*, 2018.

Byeongho Heo, Sangdoo Yun, Dongyoon Han, Sanghyuk Chun, Junsuk Choe, and Seong Joon Oh. Rethinking spatial dimensions of vision transformers. *arXiv preprint arXiv:2103.16302*, 2021.

Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 7132–7141, 2018.

Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4700–4708, 2017.

Michael J Hutchinson, Charline Le Lan, Sheheryar Zaidi, Emilien Dupont, Yee Whye Teh, and Hyunjik Kim. Lietransformer: Equivariant self-attention for lie groups. In *International Conference on Machine Learning*, pp. 4533–4543. PMLR, 2021.

John Jumper, Richard Evans, Alexander Pritzel, Tim Green, Michael Figurnov, Olaf Ronneberger, Kathryn Tunyasuvunakool, Russ Bates, Augustin Žídek, Anna Potapenko, et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873):583–589, 2021.

Kaggle and EyePacs. Kaggle diabetic retinopathy detection, July 2015. URL https://www.kaggle.com/c/diabetic-retinopathy-detection/data.

Tero Karras, Miika Aittala, Samuli Laine, Erik Härkönen, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Alias-free generative adversarial networks. *Advances in Neural Information Processing Systems*, 34, 2021.

Jakob Nikolas Kather, Cleo-Aron Weis, Francesco Bianconi, Susanne M Melchers, Lothar R Schad, Timo Gaiser, Alexander Marx, and Frank Gerrit Z"ollner. Multi-class texture analysis in colorectal cancer histology. *Scientific reports*, 6:27988, 2016.

Samuli Laine and Timo Aila. Temporal ensembling for semi-supervised learning. *arXiv preprint arXiv:1610.02242*, 2016.

Dmitry Laptev, Nikolay Savinov, Joachim M Buhmann, and Marc Pollefeys. Ti-pooling: transformation-invariant pooling for feature learning in convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 289–297, 2016.

Yann LeCun, Yoshua Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995.

Karel Lenc and Andrea Vedaldi. Understanding image representations by measuring their equivariance and equivalence. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 991–999, 2015.

Xiang Li, Wenhai Wang, Xiaolin Hu, and Jian Yang. Selective kernel networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 510–519, 2019.

Hanxiao Liu, Zihang Dai, David R So, and Quoc V Le. Pay attention to mlps. *arXiv preprint arXiv:2105.08050*, 2021a.

Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin transformer: Hierarchical vision transformer using shifted windows. *arXiv preprint arXiv:2103.14030*, 2021b.

Dhruv Mahajan, Ross Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens Van Der Maaten. Exploring the limits of weakly supervised pretraining. In *Proceedings of the European conference on computer vision (ECCV)*, pp. 181–196, 2018a.

Dhruv Kumar Mahajan, Ross B. Girshick, Vignesh Ramanathan, Kaiming He, Manohar Paluri, Yixuan Li, Ashwin Bharambe, and Laurens van der Maaten. Exploring the limits of weakly supervised pretraining. In *ECCV*, 2018b.

Diego Marcos, Michele Volpi, Nikos Komodakis, and Devis Tuia. Rotation equivariant vector field networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5048–5057, 2017.

Dushyant Mehta, Oleksandr Sotnychenko, Franziska Mueller, Weipeng Xu, Mohamed Elgharib, Pascal Fua, Hans-Peter Seidel, Helge Rhodin, Gerard Pons-Moll, and Christian Theobalt. Xnect. *ACM Transactions on Graphics*, 39(4), Jul 2020. ISSN 1557-7368. doi: 10.1145/3386569.3392410. URL http://dx.doi.org/10.1145/3386569.3392410.

Chris Olah, Nick Cammarata, Chelsea Voss, Ludwig Schubert, and Gabriel Goh. Naturally occurring equivariance in neural networks. *Distill*, 5(12):e00024–004, 2020.

Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (eds.), *Advances in Neural Information Processing Systems 32*, pp. 8024–8035. Curran Associates, Inc., 2019. URL http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf.

Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10428–10436, 2020.

Maithra Raghu, Thomas Unterthiner, Simon Kornblith, Chiyuan Zhang, and Alexey Dosovitskiy. Do vision transformers see like convolutional neural networks?, 2021.

Antônio H Ribeiro and Thomas B Schön. How convolutional neural networks deal with aliasing. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 2755–2759. IEEE, 2021.

Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 4510–4520, 2018.

Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.

Ke Sun, Yang Zhao, Borui Jiang, Tianheng Cheng, Bin Xiao, Dong Liu, Yadong Mu, Xinggang Wang, Wenyu Liu, and Jingdong Wang. High-resolution representations for labeling pixels and regions, 2019.

Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alex Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning, 2016.

Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pp. 6105–6114. PMLR, 2019a.

Mingxing Tan and Quoc V. Le. Mixconv: Mixed depthwise convolutional kernels, 2019b.

Mingxing Tan and Quoc V Le. Efficientnetv2: Smaller models and faster training. *arXiv preprint arXiv:2104.00298*, 2021.

Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V. Le. Mnasnet: Platform-aware neural architecture search for mobile, 2019.

Ilya Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, et al. Mlp-mixer: An all-mlp architecture for vision. *arXiv preprint arXiv:2105.01601*, 2021.

Hugo Touvron, Piotr Bojanowski, Mathilde Caron, Matthieu Cord, Alaaeldin El-Nouby, Edouard Grave, Gautier Izacard, Armand Joulin, Gabriel Synnaeve, Jakob Verbeek, et al. Resmlp: Feedforward networks for image classification with data-efficient training. *arXiv preprint arXiv:2105.03404*, 2021a.

Hugo Touvron, Matthieu Cord, Matthijs Douze, Francisco Massa, Alexandre Sablayrolles, and Hervé Jégou. Training data-efficient image transformers & distillation through attention. In *International Conference on Machine Learning*, pp. 10347–10357. PMLR, 2021b.

Hugo Touvron, Matthieu Cord, Alexandre Sablayrolles, Gabriel Synnaeve, and Hervé Jégou. Going deeper with image transformers. *arXiv preprint arXiv:2103.17239*, 2021c.

Cristina Vasconcelos, Hugo Larochelle, Vincent Dumoulin, Rob Romijnders, Nicolas Le Roux, and Ross Goroshin. Impact of aliasing on generalization in deep convolutional networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 10529–10538, 2021.

Bastiaan S Veeling, Jasper Linmans, Jim Winkens, Taco Cohen, and Max Welling. Rotation equivariant cnns for digital pathology. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 210–218. Springer, 2018.

Chien-Yao Wang, Hong-Yuan Mark Liao, I-Hau Yeh, Yueh-Hua Wu, Ping-Yang Chen, and Jun-Wei Hsieh. Cspnet: A new backbone that can enhance learning capability of cnn, 2019.

Maurice Weiler and Gabriele Cesa. General E(2)-Equivariant Steerable CNNs. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.

Ross Wightman. Pytorch image models. https://github.com/rwightman/pytorch-image-models, 2019.

Ross Wightman, Hugo Touvron, and Hervé Jégou. Resnet strikes back: An improved training procedure in timm. *arXiv preprint arXiv:2110.00476*, 2021.

Daniel E Worrall, Stephan J Garbin, Daniyar Turmukhambetov, and Gabriel J Brostow. Harmonic networks: Deep translation and rotation equivariance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 5028–5037, 2017.

Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 1492–1500, 2017.

I. Zeki Yalniz, Hervé Jégou, Kan Chen, Manohar Paluri, and Dhruv Mahajan. Billion-scale semi-supervised learning for image classification. *CoRR*, abs/1905.00546, 2019. URL http://arxiv.org/abs/1905.00546.

Fisher Yu, Dequan Wang, Evan Shelhamer, and Trevor Darrell. Deep layer aggregation, 2019.

Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 6023–6032, 2019.

Hang Zhang, Chongruo Wu, Zhongyue Zhang, Yi Zhu, Zhi Zhang, Haibin Lin, Yue Sun, Tong He, Jonas Muller, R. Manmatha, Mu Li, and Alexander Smola. Resnest: Split-attention networks. *arXiv preprint arXiv:2004.08955*, 2020.

Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.

Richard Zhang. Making convolutional networks shift-invariant again. In *International conference on machine learning*, pp. 7324–7334. PMLR, 2019.

Zhi Zhang, Tong He, Hang Zhang, Zhongyue Zhang, Junyuan Xie, and Mu Li. Bag of freebies for training object detection neural networks. *arXiv preprint arXiv:1902.04103*, 2019.

Zizhao Zhang, Han Zhang, Long Zhao, Ting Chen, Sercan Ö Arik, and Tomas Pfister. Nested hierarchical transformer: Towards accurate, data-efficient and interpretable visual understanding. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 36, pp. 3417–3425, 2022.

Yanzhao Zhou, Qixiang Ye, Qiang Qiu, and Jianbin Jiao. Oriented response networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 519–528, 2017.

Appendix

Table of Contents

A Aliasing Extended Discussion **15**

B Lie Groups, Lie Derivatives, and LEE **16**
B.1 Lie Groups and Local/Global Notions of Equivariance 16
B.2 Lie Derivative Chain Rule 16
B.3 Stochastic Trace Estimator for Layerwise Metric 17

C LEE Theorems **17**
C.1 LEE and consistency regularization 17
C.2 Translation LEE and aliasing 18

D Learned Equivariance Experiments **19**
D.1 Layer-wise Equivariance Baselines 19
D.2 Subnetwork Equivariance Analysis 19
D.3 Model List 20
D.4 Alternative End-to-End Equivariance Metrics 21
D.5 LEE for Additional Transformations 21
D.6 Rotated MNIST Finetuning 21

A Aliasing Extended Discussion

When working with signal f we can write the fourier series

$$f(k) = \sum_{n=-\infty}^{\infty} F(n) e^{i2\pi kn/N}$$

where $F(n)$ are the frequency components and the values n are called *harmonics*. Take the f to be discretely sampled at uniformly spaced points in a bounded interval $[-1, -1+\Delta, ..., 1-\Delta, 1]$, with $\Delta = 1/K$. Because the sampling rate is limited, it is impossible to correctly measure components $F(n)$ where $n > K/2$. Evaluated only at the grid points, such content could have identical values to components with lower frequencies, causing fundamental ambiguities:

$$\sin(2\pi(k+nK)t+\phi) = \begin{cases} +\sin(2\pi(k+nK)t+\phi) & k+nK \geq 0 \\ -\sin(2\pi|k+nK|t+\phi) & k+nK < 0 \end{cases}$$

The default mechanism for resolving these ambiguities in the reconstruction is to choose the lowest frequency component for the corresponding observations, leading to the aliasing operation given in Equation 2.

This operation can also be considered a translation in the frequency domain. Crucially, operations in frequency domain have corresponding operations in the spatial domain, and thus aliasing can give rise to recognizable patterns in images with poorly chosen resolutions, for example moire patterns. This relationship also means aliasing's effects on translations in frequency space, for example, can effect translational spatial symmetries.

To make this relationship explicit, let us consider the translational symmetries of the first set of feature maps in a CNN in two scenarios. In both scenarios the transformation is downward translation of the input by 10% of its height. First, let us consider the case where this transformation happens to result in a translation by a discrete numbers of pixels, p in the feature maps. Obviously the pixels at the bottom of the image become lost to the boundary and thus cannot be recovered from the corresponding

feature maps, as would be required for equivariance, as illustrated in Figure 1. As the amount of this translation gets smaller and smaller, however, the effect of the boundary should decrease, and yet the ability to recover the image can still be strongly affected by innate signal processing properties.

Consider the case where the CNN has a stride of 2. The feature maps will have half the width of the original image. Therefore the Nyquist frequency will also be half that of the Nyquist frequency of the image, and there will be aliasing of all the frequencies in between the original Nyquist frequency and the new value. When we try to reverse the transformation by translating p pixels upwards, the resulting translation will no longer be the inverse of the translation on the image. Therefore we cannot achieve perfect equivariance.

As another important subcase, let's also consider the non-linear activation in the CNN layer by itself. If we apply the non-linearity to a translated input, we can simply use the fact the result was a discrete translation in the output space to map the values at the grid points to values at different grid points under the reverse transformation. In this case there is clearly no issue introduced from the frequency domain properties of non-linearities on their own.

Now let's consider a translation of $1/p$ pixels. In this case, reconstructing the image after the translation is non-trivial, and we need to perform interpolation to calculate the values of the corresponding continuous image at the points that will become translated to the evaluation points. In order to perform this interpolation we must actually consider the full frequency spectrum of the image. Now the effects of pointwise non-linearities can become apparent. Because non-linearities can introduce high frequency content, these high frequencies become important when reconstructing the signal using interpolation. Aliasing makes this reconstruction fundamentally challenging and thus equivariance is impossible to achieve.

B Lie Groups, Lie Derivatives, and LEE

B.1 Lie Groups and Local/Global Notions of Equivariance

The key to understanding why the local - global equivalence holds is that $(\exp(X) - 1) = \sum_{k=1}^{\infty} X^k/k!$ has the same nullspace as X (here repeated application of X on a function f is just the repeated directional derivative, and this is the definition of a vector field used in differential geometry). Since they have the same nullspace, the space of functions for which $\exp(X)f = f$ is the same as the space $Xf = 0$. The same principle holds for $\rho(\exp(X))f = f$ and $d\rho(X)f = 0$ since $\rho(\exp(X)) = \exp(d\rho(X))$ (a basic result in representation theory, which can be found in (Hall, 2013)) where $d\rho$ is the corresponding Lie algebra representation of ρ, which for vector fields is the Lie derivative $d\rho(X) = L_X$. Hence carrying over the constraint for each element $\forall X \in \mathfrak{g} : L_X f = 0$ is equivalent to $\forall X \in \mathfrak{g} : \rho(\exp(X))f = f$ which is the same as $\forall g \in G : \rho(g)f = f$. Unpacking the representation ρ_{12} of f, this is just the global equivariance constraint $\forall g \in G : \rho_2(g)^{-1} f(\rho_1(g)x) = f(x)$.

B.2 Lie Derivative Chain Rule

Suppose we have two functions $h : V_1 \to V_2$ and $f : V_2 \to V_3$, and corresponding representations ρ_1, ρ_2, ρ_3 for the vector spaces V_1, V_2, V_3. Expanding out the definition of ρ_{31},

$$\begin{aligned}
\rho_{31}(g)[f \circ h](x) &= \rho_3(g)^{-1} f(h(\rho_1(g)x)) \\
&= \rho_3(g)^{-1} f(\rho_2(g)\rho_2(g)^{-1} h(\rho_1(g)x)) \\
&= \rho_{32}(g)[f] \circ \rho_{21}(g)[h](x).
\end{aligned}$$

From the definition of the Lie derivative, and using the chain rule that holds for the derivative with respect to the scalar t, and noting that $g_0 = \mathrm{Id}$ so $\rho(g_0) = \mathrm{Id}$, we have

$$\begin{aligned}
\mathcal{L}_X(f \circ h)(x) &= \frac{d}{dt}\left(\rho_{31}(g_t)[f \circ h](x)\right)\Big|_0 \\
&= \frac{d}{dt}\left(\rho_{32}(g_t)[f] \circ \rho_{21}(g_t)[h](x)\right)\Big|_0 \\
&= \left(\frac{d}{dt}\rho_{32}(g_t)[f]\big|_{t=0}\right) \circ \rho_{21}(g_0)[h](x) + \left[d(\rho_{32}(g_0)[f])\Big|_{h(x)}\right]\left(\frac{d}{dt}\rho_{21}(g_t)[h]\big|_{t=0}\right)(x) \\
&= \left(\frac{d}{dt}\rho_{32}(g_t)[f]\big|_{t=0}\right) \circ h(x) + df|_{h(x)}\left(\frac{d}{dt}\rho_{21}(g_t)[h]\big|_{t=0}\right)(x) \\
&= (\mathcal{L}_X f) \circ h(x) + df|_{h(x)}(\mathcal{L}_X h)(x),
\end{aligned}$$

where $df|_{h(x)}$ is the Jacobian of f at $h(x)$ and $df|_{h(x)}(\mathcal{L}_X h)(x)$ is understood to be the Jacobian vector product of $df|_{h(x)}$ with $(\mathcal{L}_X h)(x)$, equivalent to the directional derivative of f along $(\mathcal{L}_X h)(x)$. Therefore the Lie derivative satisfies a chain rule

B.3 Stochastic Trace Estimator for Layerwise Metric

Unrolling this chain rule for a sequence of layers $\mathrm{NN}(x) = f_{N:1}(x) := f_N(f_{N-1}(...(f_1(x))))$, or even an autograd DAG, we can identify the contribution that each layer f_i makes to the equivariance error of the whole as the sum of terms $C_i = df_{N:i+1}\mathcal{L}_X f_i$, $\mathcal{L}_X(\mathrm{NN}) = \sum_{i=1}^N C_i$.

Each of these C_i, like $\mathcal{L}_X(\mathrm{NN})$ measure the equivariance error for all of the outputs (which we define to be the softmax probabilities), and are hence vectors of size K where K is the number of classes. In order to summarize the C_i as a single number for plotting, we compute their norm $\|C_i\|$ which satisfy $\|\mathcal{L}_X(\mathrm{NN})\| \leq \sum_i \|C_i\|$.

To compute $df_{N:i+1}\mathcal{L}_X f_i$, one can use autograd to perform Jacobian vector products (as opposed to typical vector Jacobian products) and build up $df_{N:i+1}$ in a backwards pass. Unfortunately doing so is quite cumbersome in the PyTorch framework where the large number of available models are implemented and pretrained. A trick which can be used to speed up this computation is to use stochastic trace estimation (Avron & Toledo, 2011). Since vector Jacobian products are cheap and easy, we can compute $\|C_i\|^2 = \mathbb{E}[\hat{A}]$ as the expectation of the estimator $\hat{A} = (1/N)\sum_n^N (z_n^\top C_i)^2 = (1/N)\sum_n^N (z_n^\top df_{N:i+1}\mathcal{L}_X f_i)^2$ with iid. Normal probe vectors $z_n \sim \mathcal{N}(0, I)$, and the quantity $z_n^\top df_{N:i+1}$ which is a standard vector Jacobian product.

One can see that $\mathbb{E}[\hat{A}] = C_i^\top \mathbb{E}[zz^\top]C_i = C_i^\top I C_i = \|C_i\|^2$. We can then measure the variance of this estimator to control for the error and increase N until this error is at an acceptable tolerance (we use $N = 100$ probes). The convergence of this trace estimator is shown in Figure 9 (right) for several different layers of a ResNet-50. In producing the final layerwise attribution plots, we average the computed quantity $\|C_i\|$ over 20 images from the ImageNet test set.

C LEE Theorems

C.1 LEE and consistency regularization

As shown in Athiwaratkun et al. (2018), consistency regularization with Gaussian input perturbations can be viewed as an estimator for the norm of the Jacobian of the network, but in fact when the perturbations are not Gaussian but from small spatial transformations, consistency regularization actually penalizes the Lie derivative norm. In the Π-model (Laine & Aila, 2016) (the most basic form of consistency regularization), the consistency regularization minimizes the norm of the difference of the outputs of the network when two randomly sampled transformations T^a and T^b are applied to the input,

$$\mathrm{L}_{\mathrm{cons}} = \|f(T^a(x)) - f(T^b(x))\|^2. \tag{7}$$

Suppose that the two transformations are representations of a given symmetry group and can be written as $T^a = \rho(g_a)$ and $T^b = \rho(g_b)$, and the group elements can be expressed as the flow

generated by a linear combination of the vector fields which form the Lie Algebra: $g_a = \Phi_{\sum_i a_i X_i}$ for some coefficients $\{a_i\}_{i=1}^d$ and likewise for g_b. We can define the log map, mapping group elements top their generator values in this basis: $\log(g_a) = a$. Then, assuming a_i are small (and therefore the transformations are small), Taylor expansion yields $\mathrm{L}_{\mathrm{cons}} = \|f(x) + \sum_i a_i \mathcal{L}_{X_i} f(x) + O(a^2) - [f(x) + \sum_j b_j \mathcal{L}_{X_j} f(x) + O(b^2)]\|^2$. Therefore, taking the expectation over the distribution which a and b are sampled over (which is assumed to be centered with $\mathbb{E}[a_i] = \mathbb{E}[b_i] = 0$ as well as the input distribution x, we get that

$$\mathbb{E}_{a,b,x}[\mathrm{L}_{\mathrm{cons}}] = 2\mathbb{E}[\| \sum_i \mathcal{L}_{X_i} f(x)\|_\Sigma^2] + \text{higher order terms}, \tag{8}$$

where $\|\|_\Sigma^2$ denotes the norm with respect to the covariance matrix $\Sigma = \mathrm{Cov}(a) = \mathrm{Cov}(b)$.

When the transformations are not parameter space perturbations such as dropout, but input space perturbations like translations (which have been found to be far more important to the overall performance of the method (Athiwaratkun et al., 2018)), we can show that consistency regularization coincides with minimizing the expected Lie derivative norm. In this sense, consistency regularization can be viewed as an intervention for reducing the equivariance error on unlabeled data.

C.2 Translation LEE and aliasing

Below we show that spatial aliasing directly introduces translation equivariance error as measured by the Lie derivative, where the aliasing operation $A[\cdot]$ is given by Equation 2. The Fourier series representation of an image $h(x, y)$ with pixel locations (x, y) is H_{nm} with spatial frequencies (n, m), where the band limited reconstruction

$$h(x, y) = \tfrac{1}{2\pi} \sum_{nm} H_{nm} e^{2\pi i(xn+ym)} = F^{-1}[H]$$

and F^{-1} is the inverse Fourier transform, and the sums range over frequencies of $-M/2$ to $+M/2$ for both n and m where M is the image height and width (assumed to be square for convenience).

Applying a continuous translation by $t\mathbf{v}$ along vector $\mathbf{v} = (v_x, v_y)$ to the input means resampling the translated band limited continuous reconstruction $h(x, y)$ at the grid points.

$$T_{t\mathbf{v}}[h](x, y) = h(x - tv_x, y - tv_y) = \tfrac{1}{2\pi} \sum_{n,m=-M/2}^{M/2} H_{nm} e^{2\pi i[(x-tv_x)n+(y-tv_y)m]}$$

To simplify the notation, we will consider translations along only x and suppress the m index of H_{nm}, effectively deriving the result for the translations of a 1d sequence, but that extends straightforwardly to the 2 dimensional case.

$$T_{t\mathbf{v}}[h](x) = h(x - tv_x) = \tfrac{1}{2\pi} \sum_{n=-M/2}^{M/2} [H_n e^{-2\pi i t v_x n}] e^{2\pi i x n}$$

Applying the aliasing operation, sampling the image to a new size M' (with Nyquist frequency $M'/2$), we have

$$\begin{aligned} A[T_{t\mathbf{v}}[h]](x) &= \tfrac{1}{2\pi} \sum_{n=-M/2}^{M/2} [H_n e^{-2\pi i t v_x n}] e^{2\pi i x \mathrm{Alias}(n)} \\ &= \tfrac{1}{2\pi} \sum_{n'=-M'/2}^{M'/2} \left[\sum_{n=\mathrm{Alias}^{-1}(n')} H_n e^{-2\pi i t v_x n} \right] e^{2\pi i x n'} \end{aligned}$$

where the last line follows from applying a change of variables $n' = \mathrm{Alias}(n)$.

Applying the final inverse translation (which acts on the M' sampling rate band limited continuous reconstruction), we have

$$T_{-t\mathbf{v}}[A[T_{t\mathbf{v}}[h]]](x) = \tfrac{1}{2\pi} \sum_{n'=-M'/2}^{M'/2} \left[\sum_{n=\mathrm{Alias}^{-1}(n')} H_n e^{-2\pi i t v_x (n-n')} \right] e^{2\pi i x n'}.$$

Taking the derivative with respect to t, we have

$$\begin{aligned}\mathcal{L}_{\mathbf{v}}(A)(h) &= \tfrac{d}{dt}\big|_0 T_{-t\mathbf{v}}[A[T_{t\mathbf{v}}[h]]] \\ &= \tfrac{1}{2\pi}\sum_{n'=-M'/2}^{M'/2}\Bigg[\sum_{n=\mathrm{Alias}^{-1}(n')} 2\pi i v_x(n'-n)H_n\Bigg]e^{2\pi i x n'}.\end{aligned}$$

Notably, for aliasing when the frequency is reduced by a factor of 2 from downsampling, there are only two values of n that satisfy $\mathrm{Alias}(n) = n'$: the value $n = n'$ and the one that gets aliased down, therefore when multiplied by $n - n'$ the sum

$$\Bigg[\sum_{n=\mathrm{Alias}^{-1}(n')} 2\pi i v_x(n'-n)H_n\Bigg]$$

consists only of a single term.

According to Parseval's theorem, the Fourier transform F is unitary, and therefore the norm of the function as a vector evaluated at the discrete sampling points $x = 1/M', 2/M', \ldots$ is the same as as the norm of the Fourier transform:

$$\begin{aligned}\|\mathcal{L}_{\mathbf{v}}(A)(h)\|^2 &= \|F[\mathcal{L}_{\mathbf{v}}(A)(h)]\|^2 \\ \|\mathcal{L}_{\mathbf{v}}(A)(h)\|^2 &= \sum_{n'=-M'/2}^{M'/2}\Bigg|\sum_{n=\mathrm{Alias}^{-1}(n')} 2\pi i v_x(n'-n)H_n\Bigg|^2 \\ \|\mathcal{L}_{\mathbf{v}}(A)(h)\|^2 &= \sum_{n=-M/2}^{M/2}(2\pi)^2 v_x^2(\mathrm{Alias}(n)-n)^2 H_n^2,\end{aligned}$$

using the fact that only one element is nonzero in the sum. Finally, generalizing to the 2d case, we have

$$\|\mathcal{L}_{\mathbf{v}}(A)(h)\|^2 = (2\pi)^2\sum_{nm} H_{nm}^2\big(v_x^2(\mathrm{Alias}(n)-n)^2 + v_y^2(\mathrm{Alias}(m)-m)^2\big), \tag{9}$$

showing how the translation Lie derivative norm is determined by the higher frequency components which are aliased down.

D Learned Equivariance Experiments

D.1 Layer-wise Equivariance Baselines

We use EQ-T and EQ-T_{frac} (Karras et al., 2021) to calculate layer-wise equivariance by caching intermediate representations from the forward pass of the model. For image-shaped intermediate representations, EQ-T samples integer translations in pixels between -12.5% and 12.5% of the image dimensions in pixels. EQ-T_{frac} is identical but with continuous translation vectors. The individual layer is applied to the transformed input and then the inverse group action is applied to the output, which is compared with the original cached output. Many different normalization could be chosen to compare equivariance errors across layers. The most obvious are $\frac{1}{N}$, $\frac{1}{\sqrt{N}}$, and $\frac{1}{1}$ (no normalization), where $N = C \times H \times W$. As we show in section 5, the normalization method can have a large effect of the relative contribution of a layer, despite the decision being relatively arbitrary (in contrast to LEE, which removes the need for doing so as the scale is automatically measured relative to the contribution to the output).

D.2 Subnetwork Equivariance Analysis

Another way one might use LEE to study the effects of different layers that make up a network is to break the network in question down into its constituent subnetworks (networks starting at the input and ending at every intermediate representation in the network) and calculate the LEE of the corresponding function. We show the result of this calculation on a ResNet50 in Figure 8.

As an alternative to our layerwise analysis, this method has a significant drawback that makes analysis challenging: the functions under consideration have different outputs. In our calculation, we applied batch normalization over the outputs in order to make their scales comparable. Despite this rescaling, comparing activations and preactivations, for example, remains challenging. By contrast, our layerwise breakdown specifically targets a layer's contribution to a shared output.

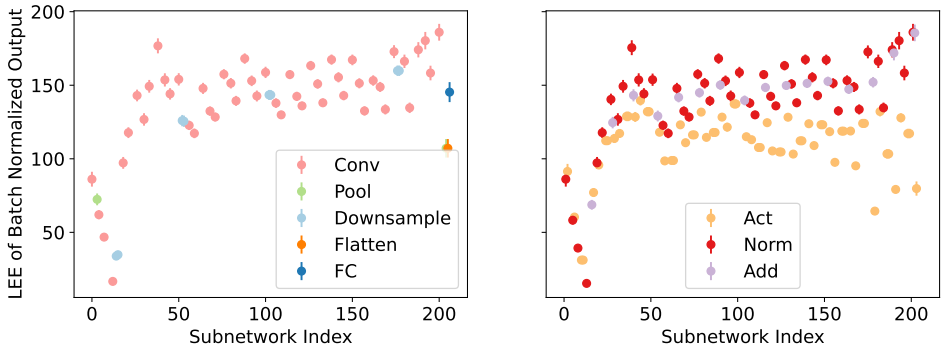

Figure 8: LEE calculated over the subnetworks of a ResNet50. Specifically a subnetwork is constructed between the input and every intermediate representation in the network's computation graph. We use batch normalization of the outputs to make the output scale of different subnetwork comparable. For visual clarity, layer types are broken across the left and right plots, which share the same axes. Similar to the pattern observed in Figure 4, we see a rapid increase in equivariance error in the early layers of the network, followed by many smaller increases later in the network. Unlike in our layerwise decomposition, comparison across layer types is challenging in this setting because layers have significantly different outputs. For example, comparing activations with preactivations is complicated by the ReLUs acting as contractions of the input and having potentially many zeroed values.

D.3 Model List

The models included in Figure 1 are

- Early CNNs: ResNets (He et al., 2015), ResNeXts (Xie et al., 2017), VGG (Simonyan & Zisserman, 2014), Inception (Szegedy et al., 2016), Xception (Chollet, 2017), DenseNet (Huang et al., 2017), MobileNet (Sandler et al., 2018), Blur-Pool Resnets and Densenets (Zhang, 2019), ResNeXt-IG (Mahajan et al., 2018a), SeResNe*ts (Hu et al., 2018), ResNet-D (He et al., 2018), Gluon ResNets (Guo et al., 2020; Zhang et al., 2019; 2020), SKResNets (Li et al., 2019), DPNs (Chen et al., 2017)

- Modern CNNs: EfficientNet (Tan & Le, 2019a; 2021), ConvMixer (**?**), RegNets (Radosavovic et al., 2020), ResNet-RS, (Bello et al., 2021), ResNets with new training recipes (Wightman et al., 2021), ResNeSts (Zhang et al., 2020), RexNet (Han et al., 2021a), Res2Net (Gao et al., 2019), RepVGG (Ding et al., 2021), NFNets (Brock et al., 2021), XNect (Mehta et al., 2020), MixNets (Tan & Le, 2019b), ResNeXts with SSL pretraining (Yalniz et al., 2019), DLA (Yu et al., 2019), CSPNets (Wang et al., 2019), ECA NFNets and ResNets (Brock et al., 2021), HRNet (Sun et al., 2019), MnasNet (Tan et al., 2019)

- Vision transformers: ViT (Dosovitskiy et al., 2020), CoaT (Dai et al., 2021), SwinViT (Liu et al., 2021b), (Bao et al., 2021), CaiT (Touvron et al., 2021c), ConViT (d'Ascoli et al., 2021), CrossViT (Chen et al., 2021), TwinsViT (Chu et al., 2021), TnT (Han et al., 2021b), XCiT (El-Nouby et al., 2021), PiT (Heo et al., 2021), Nested Transformers (Zhang et al., 2022)

- MLP-based architectures: MLPMixer (Touvron et al., 2021b), ResMLP (Touvron et al., 2021a), gMLP (Liu et al., 2021a), MLP-Mixers with (Si)GLU (Wightman, 2019)

D.4 Alternative End-to-End Equivariance Metrics

Discrete Consistency We adopt the consistency metric from Zhang (2019), which simply measures the fraction of top-1 predictions that match after applying an integer translation to the input (in our case by 10 pixels). Instead of reporting consistency numbers, we report $(1 - \%$ matching$)$, so that consistency because a measure of equivariance error. Equivariant models should exhibit end-to-end invariance, high consistency, and low equivariance error.

Expected Group Sample Equivariance Inspired by work in equivariant architecture design (Finzi et al., 2020; Hutchinson et al., 2021), we provide an additional equivariance metric for comparison against the Lie derivative. Following (Hutchinson et al., 2021), we sample k group elements in the neighborhood of the identity group element, with sampling distribution $\mathcal{D}(G)$, and calculate the sample equivariance error for model f as $\frac{1}{k}||\rho_2^{-1}(g)f(\rho_1(g)x) - f(x)||$. For translations we take $\mathcal{D}(G)$ to be Uniform$(-5, 5)$ in pixels.

Versus LEE There are several reasons why the continuous lie derivative metric is preferable over discrete and group sample metrics. Firstly, it allows us to break down the equivariance error layerwise enabling more fine grained analysis in a way not possible with the discrete analog. Secondly, the metric is less dependent on architectural details like the input resolution of the network. For example, for discrete translations by 1 pixel, these translations have a different meaning depending on the resolution of the input, whereas our lie derivatives are defined as the derivative of translations as a fraction of the input size, which is consistently defined regardless of the resolution. Working with the vector space forming the Lie algebra rather than the group also removes some unnecessary freedom in how one constructs the metric. Rather than having to choose an arbitrary distribution over group elements, if we compute the Lie derivatives for a set of basis vectors of the lie algebra, we have completely characterized the space, and all lie derivatives are simply linear combinations of the computed values. Finally, paying attention to continuous transformations reveals the problems caused by aliasing which are far less apparent when considering discrete transformations, and ultimately the relevant transformations are continuous and we should study them directly.

D.5 LEE for Additional Transformations

Beyond the 3 continuous transformations that we study with Lie derivatives above, there are many more that might reveal important properties of the network. Here we include an three additional transformations–hyperbolic rotation, brightening, and stretch.

Figure 9 (left) shows that, perhaps surprisingly, models with high accuracy become more equivariant to hyperbolic rotations. We suspect this surprisingly general equivariance to diverse set of continuous transformations is probably the result of improved anti-aliasing learned implicitly by more accurate models. LEE does not identify any significant correlation between brightening or stretch transformations and generalization ability.

D.6 Rotated MNIST Finetuning

In order to test the ability of SOTA imagenet pre-trained models to learn equivariance competitive with specialized architectures, we adapted the example rotated MNIST notebook available in E2CNN repository (Weiler & Cesa, 2019). We use the base model and default finetuning procedure from (He et al., 2021), finetuning for 100 epochs, halving the learning rate on loss plateaus.

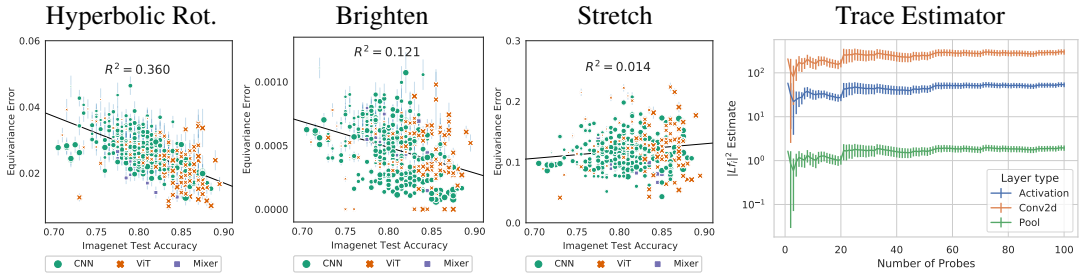

Figure 9: **(Left)**: Extending Figure 5 we show the Lie derivate norm for hyperbolic rotation, brightening, and stretch transformations. We observe that more accurate models are also more equivariant to hyperbolic rotations and to brighten transformation, to a more limited extent. In the case of hyperbolic rotations, this result is surprising, as nothing has directly encouraged this equivariance. One possible explanation is decreased aliasing in models with higher accuracy. Marker size indicates model size. Error bars show one standard error over the images use to evaluate the Lie derivative. **(Right)**: Cumulative mean and standard error of the estimator (computed for translations on a ResNet-50).