

Rethinking the Role of Spatial Mixing

George Cazenavette^{1*} Joel Julin² Simon Lucey³

¹Massachusetts Institute of Technology ²Carnegie Mellon University ³University of Adelaide

Abstract

Until quite recently, the backbone of nearly every state-of-the-art computer vision model has been the 2D convolution. At its core, a 2D convolution simultaneously mixes information across both the spatial and channel dimensions of a representation. Many recent computer vision architectures consist of sequences of isotropic blocks that disentangle the spatial and channel-mixing components. This separation of the operations allows us to more closely juxtapose the effects of spatial and channel mixing in deep learning. In this paper, we take an initial step towards garnering a deeper understanding of the roles of these mixing operations. Through our experiments and analysis, we discover that on both classical (ResNet) and cutting-edge (ConvMixer) models, we can reach nearly the same level of classification performance by only learning channel mixing and leaving the spatial mixers at their random initializations. Furthermore, we show that models with random, fixed spatial mixing are naturally more robust to adversarial perturbations. Lastly, we show that this phenomenon extends past the classification regime, as such models can also decode pixel-shuffled images.

1. Introduction

For the better part of the last two decades, cascades of learned convolutions have formed the backbone of nearly every innovation in the field of computer vision and pattern recognition. From distinguishing ten digits with LeNet [20] to one thousand classes with AlexNet [19], from going very deep with VGG [24] to even deeper with InceptionNet [25], the learned 2D convolution has served as the workhorse that ushered in the new era of visual learning. Convolutional neural networks (CNNs) learned to exploit the correlations across channels and between spatially close pixels to solve a plethora of tasks.

Recently, isotropic networks (those in which the size of the representation stays fixed throughout) such as Vision Transformer [10], Image GPT [4], MLP-Mixer [26], and ConvMixer [27] have been grabbing the field’s attention. These isotropic models consist of repeated blocks wherein

each block consists of a spatial mixing operation (self-attention [4, 10], spatial MLP [26], or depthwise convolution [27]) followed by one or more *strictly* channel-mixing layers (1×1 or “pointwise” convolutions). A strictly spatial mixing operation is defined here as any operation that is applied independently across each channel. Similarly, a strictly channel-mixing layer is applied independently across spatial pixel coordinates. The fact that all these recent state-of-the-art isotropic architectures spend a significant portion of their computation budget solely on channel-mixing parameters hints toward the possibility that channel mixing may have significantly more relative importance than previously appreciated.

It has been understood for some time that impressive performance can be obtained from networks whose weights are not all completely learned [23]. Before the success of AlexNet [19] in training deep end-to-end networks, the use of random weights in early convolutional layers played an important role in training deeper CNNs. The approach was attractive as it promised faster training times, better generalisation, and the ability to learn deeper networks. Saxe et al. [23] even offered strong theoretical motivations for why random filter weights in CNNs would offer good performance in terms of: (i) frequency selection, and (ii) translation invariance. One of the most notable works with respect to using random weights within networks can be found in the Extreme Learning Machines (ELMs) of Huang et al [17]. In its simplest form an ELM takes a proposed network architecture and randomly initializes all hidden weights, leaving only the final layer to learn. ELMs are advantageous as they allow for extremely deep networks and rapid train times (as only the final layer needs to be learned). It is widely understood, however, that these training strategies have a significant performance gap in relation to their current end-to-end learned counterparts.

In this work, we leverage the separable convolution to compare the relative effects of spatial and channel mixing on the performance of deep neural networks. By leaving either the spatial-mixing or channel-mixing parameters frozen at their random initialization and only training the others, we can isolate the contributions of both types of parameters. Do-

*Correspondence: gcaz@mit.edu

ing so reveals that learning the *channel-mixing* parameters is *far* more critical to the effectiveness of a deep model, and such models that only learn channel-mixing parameters perform nearly just as well as their fully-learned counterparts.

We also show that models that only learn channel-mixing parameters are naturally more resistant to adversarial attacks than fully-learned models. This robustness can be further enhanced by then directly smoothing the random, frozen spatial-mixing filters.

Lastly, we show that this phenomenon extends beyond the regime of classification problems and find that models that only learn spatial mixing are even capable of learning to invert a permutation of pixels nearly just as well as fully-learned networks.

We hope that these gained insights as to the role and effectiveness of spatial and channel mixing in deep models will be of great value to the vision community and that they will shed light on which parts of deep models are most critical to learn while helping to inspire the intelligent design of future architectures.

2. Related Work

Isotropic Architectures First popularized by the transformer [28], isotropic architectures (those in which the size of the representation remains fixed throughout) have more recently made their way into the computer vision world. Image-GPT [4] (based on the GPT language model [4]) was able to model sequences of pixels and generate new images in raster fashion. Designed for discriminative tasks, the Vision Transformer [10] used transformer blocks to achieve state-of-the-art image classification results, but only after extensive pre-training. The MLP-Mixer [26] then built off Vision Transformer’s success, but used a simple spatial-mixing MLP instead of the expensive self-attention module. While the MLP-Mixer still required extensive pre-training, the newest addition to the line of isotropic architectures, the ConvMixer [27], replaces the self-attention or spatial-mixing MLP with a simple depthwise filter bank and achieves comparable results *without* any pre-training. The main thing all these models have in common is their *isotropic* structure. Here, “isotropic” refers to repeated blocks of operations in which the latent representation does not change in shape. Furthermore, in all of these models, the second half of each isotropic block is made up of one or more *strictly-channel-mixing* layers. Such transformer-inspired isotropic models have recently come to be known as “meta-formers” [31].

Separable Convolutions Instead of mixing across channel and spatial dimensions simultaneously, the separable convolution factorizes the operation into a depthwise and pointwise convolution. The depthwise convolution applies a disjoint set of depth-1 filters to each channel of the input independently

while the pointwise convolution is simply a 1×1 convolution (or linear projection) on the pixels of this intermediate representation. Separable convolutions are used extensively in many state-of-the-art architectures [6, 16, 32] and have built-in implementations in deep learning libraries [1]. Furthermore, a recent method [7] converts a *pre-trained* network with standard convolutions and converts them to separable ones, increasing the throughput and robustness of the model.

Lottery Ticket Hypothesis Recent work has uncovered an interesting phenomenon in deep neural networks wherein randomly initialized networks tend to contain much smaller sub-networks that can reach similar performance as the full network when both are trained to convergence. This has been dubbed the “Lottery Ticket Hypothesis” [11] with the analogy being that the optimal sub-network of randomly initialized weights has the winning “ticket.” In this work, we provide evidence of a similar phenomenon in very wide networks with fixed, randomly initialized spatial mixing layers wherein they reach the same performance as their fully-learned counterparts (Section 4.1.4).

Fixed Spatial Mixing Some recent works have explored using fixed “shift” operations [5, 29] or simple average pooling [13, 31] as spatial-mixing operations. Our work significantly differs from these in several key areas. Rather than proposing a new architecture, our work seeks to exploit existing useful mixing pattern present at the network’s random initialization (as inspired by the lottery ticket hypothesis [11]). Furthermore, these recently proposed models [5, 13, 29, 31] still *learn* varying amounts of spatial-mixing layers whereas we analyze networks that do not learn **any** spatial-mixing parameters whatsoever. Our analysis seeks to illuminate which parameters must be learned and which can be left at their random initializations in architectures already in use today.

3. Background and Setup

Before our experiments, we provide some relevant exposition regarding the operations and architectures we use to illustrate the relative importance of spatial and channel mixing in deep neural networks.

3.1. Spectral Coverage

Part of the efficacy of a spatial convolution can be attributed to the size of the *spectral envelope* of its filter bank, or how much *spectral coverage* it has. By spectral coverage, we mean the number of signal frequencies covered by the filter bank. By using multiple filters in tandem, the region of the envelope grows to be the union of each filter’s individual coverage.

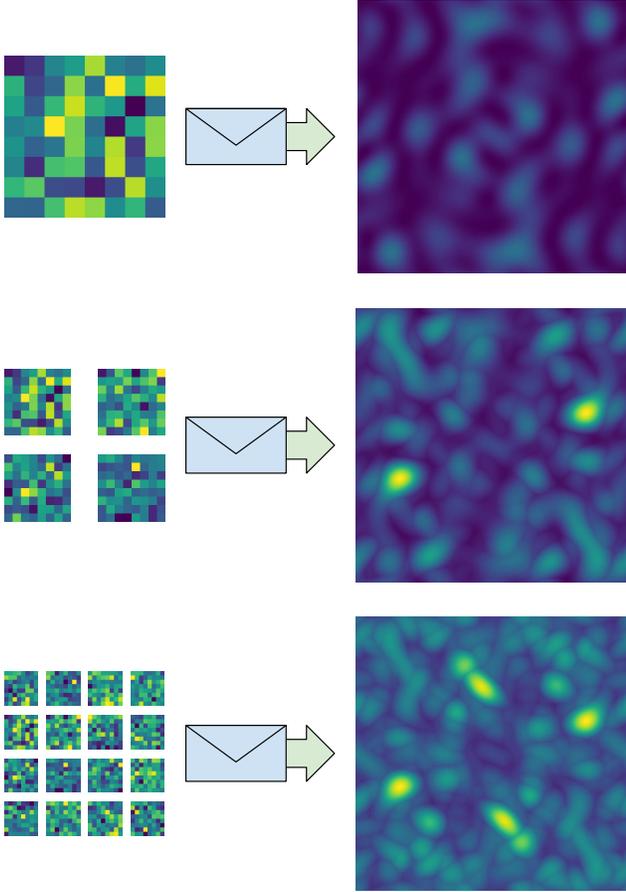


Figure 1. The spectral envelope of a bank of random filters rapidly grows as we add more uncorrelated filters to the set, allowing for quality feature extraction from natural images.

Previous work [17, 23] has shown that random filters make adequate visual feature extractors. Furthermore, motivated by decades of work in signal processing, it has been known that random filters have good spectral coverage. We hypothesize that these phenomena are related, and that the success of random filters is due in part to the size of their induced spectral envelope. As seen in Figure 1, increasing the number of filters rapidly increases the bank’s coverage in frequency space.¹

We hypothesize that random filters forming such a dense spectral envelope contributes to their ability to produce well-performing models. Intuitively, one could consider the random filters as extracting features from all different frequencies of input while the *learned* channel-mixing parameters select linear combinations of these features that are actually useful.

¹We visualize the spectral envelope by taking the max over the magnitude of the 2D fast Fourier transform (FFT) of each filter padded to be 512×512 . In our visualization the zero frequency component has been shifted to the center of the image.

3.2. The 2D Convolution

As we begin our investigation into the importance of channel mixing, we focus on the simple yet revolutionary ResNet [14] architecture. The original ResNet architecture is composed of blocks of 2D convolution operations. Given the number of in and out channels (c_{in}, c_{out}) and the size of the square kernel (k), we can represent the number of trainable parameters in the standard 2D convolution (p_{conv}) as

$$p_{conv} = c_{in} \cdot c_{out} \cdot k^2 \quad (1)$$

With the vanilla Conv2D, it is hard to analyze the respective importance of spatial and channel mixing since the two are entangled in a single linear operation.

To disentangle the spatial and channel mixing [6, 16, 32], we separate the standard 2D convolution operation into a paired spatial-mixing (depthwise) and channel-mixing (1×1 or pointwise) convolution. In a separable convolution, each filter of the depthwise convolution operates on just a single channel of the input. The pointwise convolution is simply a 1×1 convolution, or a linear projection of each pixel.

After separating the standard 2D convolution into a depthwise and pointwise convolution, we can then introduce a *depth multiplier* to increase the number of filters per input channel. Naturally, this also increases the number of channels in the pointwise filters since the intermediate representation will now have more channels. Given the number of in and out channels (c_{in}, c_{out}), the size of the square depthwise kernel (k), and the depth multiplier d , we can represent the number of trainable parameters in the separable 2D convolution (p_{sep}) as

$$\begin{aligned} p_{sep} &= p_{depth} + p_{point} \\ &= c_{in} \cdot k^2 \cdot d + c_{in} \cdot d \cdot c_{out} \end{aligned} \quad (2)$$

By exploiting the inherent dependencies between nearby pixels, the standard 2D convolution—and its separable cousin—offer vision systems a cheap, yet effective way of extracting information from images.

3.3. The ConvMixer

As one of the latest in an ever-growing line of isotropic vision models [4, 10, 26, 31], the ConvMixer [27] architecture adapts the (at this point) classical method of convolutions to the equi-sized representation (isotropic) paradigm of transformer models. Perhaps more importantly for our analysis, it serves as a state-of-the-art convolutional model in which *the convolutions are already separable*. In other words, the ConvMixer consists of strictly depthwise and pointwise convolutions. This allows us to examine the effects of learning only channel mixing in a state-of-the-art architecture without making *any* changes to the underlying architecture.

The ConvMixer [27] architecture consists of depthwise convolutions (wrapped in a residual connection) and point-

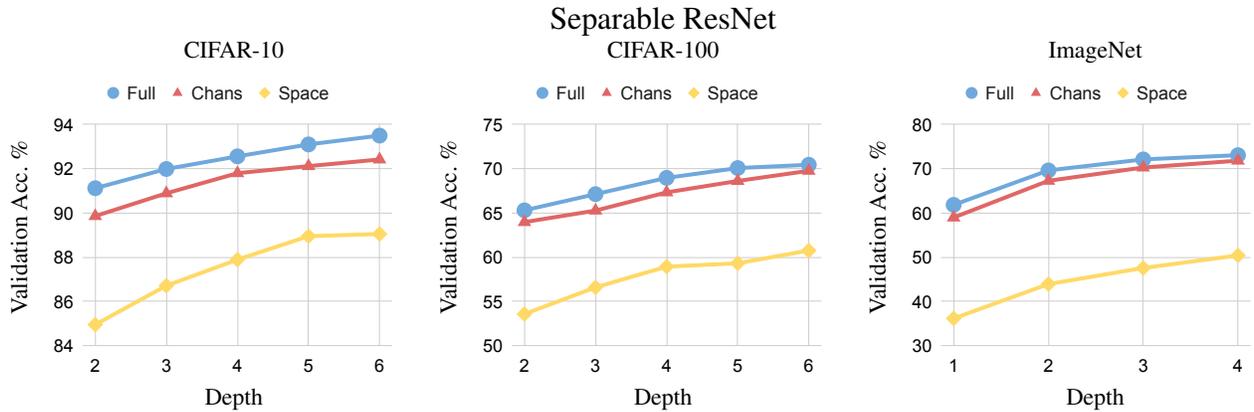


Figure 2. While the fully learned ResNets (Full) outperform all others, we see that the models that only learn channel mixing (Chans) remain quite competitive, especially so on ImageNet (right). Conversely, the models that only learn spatial mixing (Space) lag very far behind the others.

wise convolutions. Specifically, a depth- n ConvMixer consists of a one-layer patch-projection stem, n depthwise-pointwise blocks, a global average pool, and a linear classifier. All throughout, the representation maintains the size to which it is projected by the stem (hence the term *isotropic*).

3.4. Is Channel Mixing (Almost) All You Need?

By isolating the spatial and channel-mixing parts of the standard convolution into separable components, we can analyze the contribution of channel mixing alone to the success of convolutional neural networks. To do this, we now introduce a form of the separable convolution (illustrated in supplement) wherein we leave the depthwise (spatial-mixing) filters frozen at their initialized values and only learn the pointwise (channel-mixing) filters during training. In our results, this is indicated by “Chans” (or “channels-only”) whereas the fully-learned models are indicated by “Full” (or “fully-learned”).

Note that the number of trainable parameters does not depend on the size of the kernel (k) since only the pointwise parameters are learned. Furthermore, we can ensure the number of trainable parameters in the pointwise convolution is equal to that of the corresponding standard Conv2D by setting $d = k^2$:

$$\begin{aligned}
 d &= k^2 \\
 \implies c_{in} \cdot c_{out} \cdot k^2 &= c_{in} \cdot d \cdot c_{out} & (3) \\
 \implies p_{conv} &= p_{point}
 \end{aligned}$$

For ResNet architectures, the standard convolution kernel is of shape 3×3 , so we would set $d = 3^2 = 9$ to use the same number of parameters in our channel-learning-only convolution (and $d = 3$ to use one third of the parameters).

Here this is simply used as a heuristic to create models with similar parameter counts to the vanilla ResNet architecture. All of our separable ResNet models will thus have a depth-multiplier of 9, regardless of which parameters are being learned.

4. Experiments

After motivating with some relevant background, we now begin our study into the capabilities of networks that only learn channel mixing. Section 4.1 applies our hypothesis to classical ResNet [14] architectures. In Section 4.2, we move on to the naturally separable ConvMixer [27] architecture to extend our findings to a state-of-the-art isotropic model. Then, in Section 4.3, we explore a practical application in the form of architectural adversarial robustness. Lastly, in Section 4.4, we extend our findings beyond the classification regime to the pixel un-shuffling problem. Our code will be made public upon publication.

4.1. ResNet Experiments

For these experiments, we train all models under identical conditions: the same number of epochs, same batch size, same learning rate, same decay schedule, etc. Any and all differences (or similarities) in performance are due entirely to the intrinsic properties of the architectures themselves. Experiments are conducted on CIFAR-10 [18], CIFAR-100, [18] and ImageNet-1k [8].

4.1.1. Model Architectures

For our first set of experiments, we employ the ResNet [14] architecture with the identity mapping modifications introduced in [15] and convert all convolutions to the separable variant as described in Section 3.2. A depth- n CIFAR [18] ResNet contains a 1-layer stem, $2n$ layers for each of the 3 blocks, and a linear classifier, for a total of $6n + 2$ layers. Similarly an ImageNet [8] ResNet contains a 1-layer stem, $2n$ layers for each of the 4 blocks, and a linear classifier, for a total of $8n + 2$ layers. The “depth” in our plots represents this n . All separable ResNets have a depth multiplier of 9.

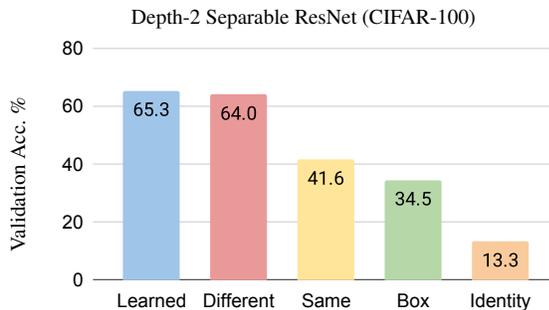


Figure 3. While *randomly* initialized filters can provide competitive results, the same is not true for *any* arbitrary, fixed filter. Random filters work best when they are uncorrelated from each other, allowing them to extract different information.

4.1.2. Classification Gap

For our ResNets that learn only channel mixing (Chans), any spatial mixing is done using depthwise filters fixed at their random initializations. Conversely, for those that learn only spatial mixing (Space), any channel mixing is done through the pointwise 1x1 convolutions frozen at their initial states. When we look at the performance of fully-learned ResNets versus those learning *only* channel mixing (Chans), we make an interesting discovery: the gap in classification accuracy is actually quite small (Figure 2). One might question the generalization of this observation, noting that CIFAR [18] is a simple dataset, and the performance may just be saturated. Yet, the trend also holds for ImageNet [8], a *much* larger and more complex dataset. Furthermore, we also see that networks only learning spatial mixing (Space) consistently perform *far* worse than those learning only channel mixing.

One might note that the number of spatial mixing parameters in a separable convolution, learned or not, is inherently smaller than that of channel mixing parameters. This is true; however, the same can also be said of the standard convolution. In a typical network, the channel mixing dimensions that contribute to the weight size ($c_{in} \cdot c_{out}$) are typically much larger than the spatial mixing dimensions ($k \cdot k$).

Ultimately, the shown fact that networks learning only channel mixing achieve competitive results to those that are fully learned comes as quite the surprise and calls into question the significance of learning the spatial versus channel-mixing operations in general.

4.1.3. Static Filter Structure

While we have shown that learning spatial mixing is not necessary to achieve competitive performance without changes to the model’s architecture, the static filters cannot be *completely* arbitrary; they must still hold some basic properties to adequately transform the input signal.

Looking at Figure 3, we see the performance of the fully-learned separable convolutions as the leftmost (blue) column. The next column (red) represents the separable convolution

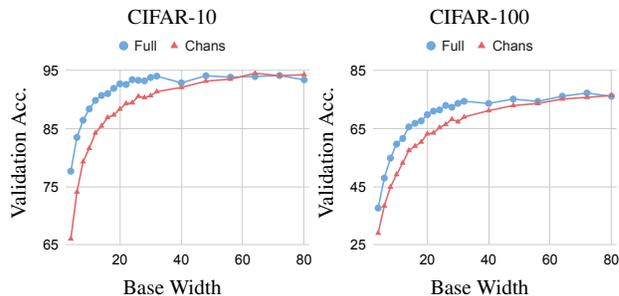


Figure 4. As we increase the network width, we see the performance of the channels-only models converges to that of the fully-learned models

that only learns channel mixing. We again note that the randomly initialized spatial mixing performs only slightly worse than the fully-learned version. We label this column “Different” to contrast with the next (yellow) column, “Same.” For this setting, we apply the *same* set of 9 randomly initialized kernels to *every input channel*. The large performance drop here shows us that highly-correlated filters do not make good feature extractors since they cover the same spectrum of signals. The next drop-off to the box filter (green) occurs because box filters are *perfectly* translationally correlated with each other. Lastly, the catastrophic failure of the identity filter (orange) tells us that although we might not need to *learn* it, spatial mixing is still an integral part of discriminative networks.

4.1.4. Extra-Wide ResNets

Our previous experiments revealed a trend of the channels-only models only lagging slightly behind the fully-learned models over a range of depths. Now, we examine the gap in performance as we greatly increase the width (number of channels) of the models.

In Figure 4, we see that the performance of the channels-only models converges to that of the fully-learned models as we continue to increase the width. While the lottery ticket hypothesis [11] states that there typically exist optimal randomly initialized sub-networks that *can be trained* to reach the same performance as the full model, our results here suggest that given enough spatial-mixing filters (or “tickets”), a channels-only model can learn to exploit the filters to **reach the same performance** as the fully-learned model while **leaving the spatial-mixing filters at their random initialization**.

4.2. ConvMixer Experiments

We use the code provided along with the ConvMixer [27] paper for these experiments and employ models with 128 channels, a kernel size of 8, and a patch size of 1 for our CIFAR [18] experiments and models with 512 channels, a kernel size of 8, and a patch size of 7 for our ImageNet [8]

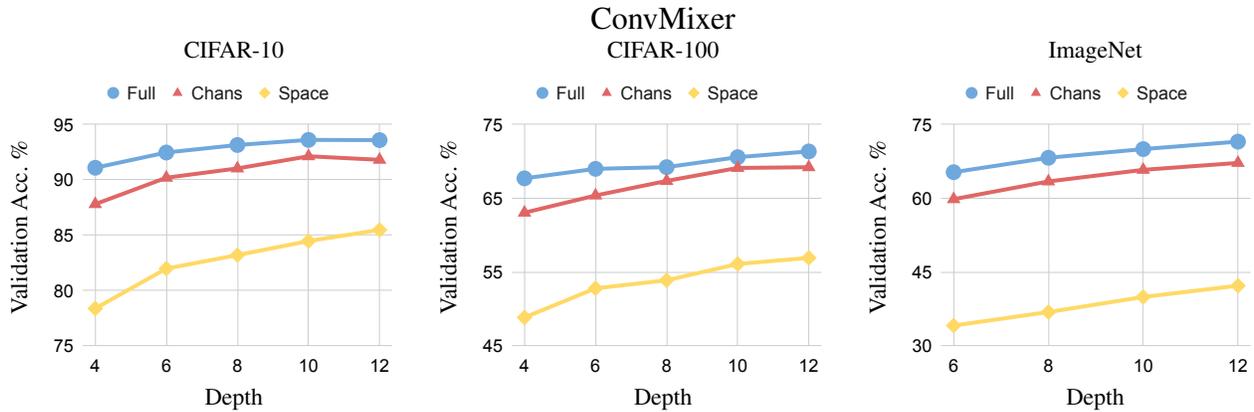


Figure 5. With the ConvMixer architecture, we can better analyze the direct contributions of spatial and channel mixing without altering the original model. We again see networks that only learn channel mixing (Chans) remain competitive with their fully-learned counterparts (Full) while completely out-classing those that only learn spatial mixing (Space).

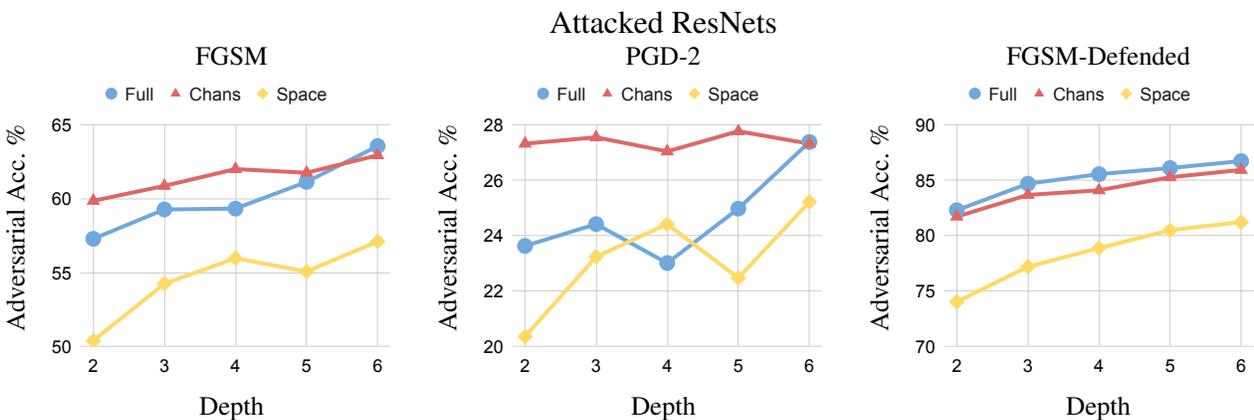


Figure 6. Adversarial accuracy in separable ResNets (CIFAR-10). Channels-only models (Chans) are clearly more *naturally* robust to adversarial perturbations than their fully-learned counterparts and those that learn only spatial mixing for both simple (left) and more advanced (middle) attacks. However, this gap disappears when adversarial training techniques are used (right). ($\epsilon = 1/255$)

experiments. Note that in channels-only models, the patch-projection layer is also left at its random initialization.

4.2.1. Naturally Separable Performance

Now that we are dealing with a naturally separable architecture, all comparisons can be made *directly* with the original model. In Figure 5, we again see that if we learn only channel mixing, we still achieve results competitive with the fully-learned architecture. Even on the much more challenging ImageNet dataset, our channels-only ConvMixers only lag slightly behind the full-learned models at all depths, showing that this phenomenon extends past toy models and small datasets to state-of-the-art architectures and real-world imagery.

4.2.2. Filter Observations

After years of networks using mostly 3×3 filters, the ConvMixer [27] takes a step back to larger filters that are more interpretable to the human eye. In Figure 9, we see the learned spectral coverage of a depth-4 ConvMixer and

note that the filters of each layer seem to follow a different distribution. Despite the structure clearly present in these learned filters, the random filters (as seen in Figure 7) still yield comparable performance. As we see in the following section, these specialized learned filters bring with them their own issues in the form of sensitivity to adversarial noise.

4.3. Adversarial Robustness

Another practical application of our discovery lies in the adversarial setting. Adversarial examples are those specifically curated to fool a neural network into making a mistake at inference time by adding targeted noise to the sample [12]. Typically, the adversarial perturbations are small enough in magnitude that the semantic meaning of the example does not change and that a human observer cannot even notice them.

4.3.1. The Fast Gradient Sign Method

First, we will quickly introduce one of the first and most simple adversarial attacks developed: the Fast Gradient Sign Method [12]. This method calculates the adversarial per-

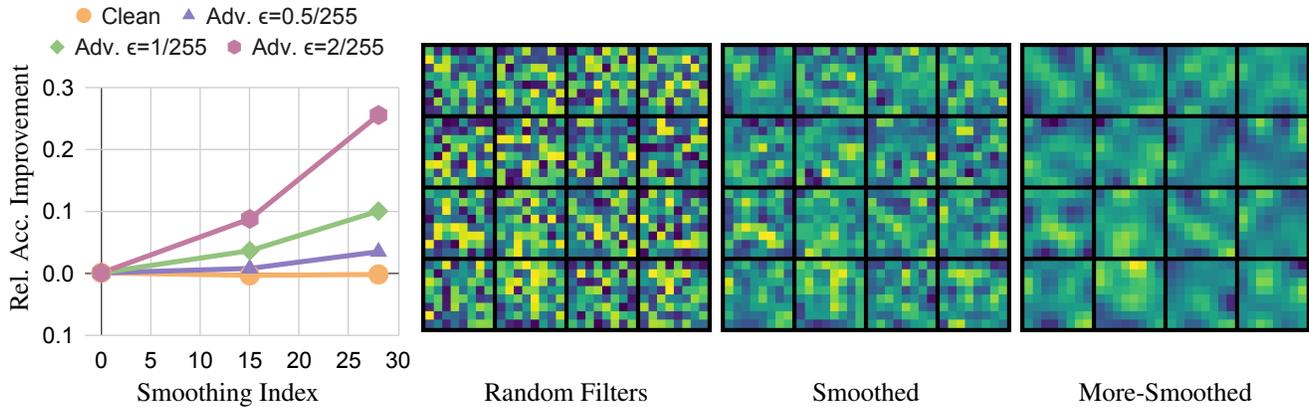


Figure 7. Smoothing the random filters causes our trained models become significantly more robust to adversarial perturbations.

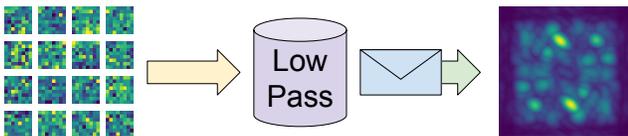


Figure 8. By filtering out the high-frequency components of our filter bank, we also eliminate the high-frequency parts of its spectral envelope, leading to a more adversarially robust model.

turbations by taking the gradient of the network’s loss with respect to the target image and mapping them to the ℓ_∞ ball.

Given a neural network (\mathcal{F}), target sample (\mathbf{x}), and label (y), the adversarial example ($\tilde{\mathbf{x}}$) is calculated as

$$\tilde{\mathbf{x}} := \mathbf{x} + \epsilon \cdot \text{sign} \left(\frac{\partial \mathcal{F}(\mathbf{x}, y)}{\partial \mathbf{x}} \right) \quad (4)$$

where ϵ is the radius of the ℓ_∞ ball.

This is a simple attack that can be defended against by a variety of methods [2, 3, 12, 21, 22], but it is still a useful tool for analyzing the natural adversarial robustness of an architecture without taking any preventative measures.

We also experiment with a more difficult version of this attack where projected gradient descent (PGD) is used to iteratively update the adversarial example using Eq. 4. Many updates caused the attack to be too strong for good comparisons, so we just used 2 iterations (PGD-2).

For the experiments where we also employed a defensive technique, FGSM was simply applied to the training samples during model training time.

4.3.2. Natural Robustness in Separable ResNets

As we see in Figure 6, ResNets with separable convolutions that only learn channel mixing (Sep-Res) are significantly more robust to the adversarial perturbations than the fully-learned separable ResNets. Without any further modification to the training process, the models that only learn channel mixing seem to be less susceptible to targeted noise.

4.3.3. Induced Robustness in ConvMixers

Looking at envelopes of the learned filters (Figure 9), one could mistakenly link the *smoothness* of learned filters to

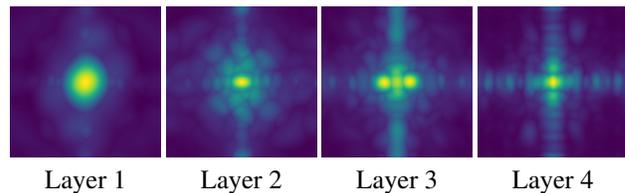


Figure 9. Spectral envelopes of learned ConvMixer filters. Naturally-learned filters have dense coverage of the lower frequencies without unnecessarily capturing the high frequencies.

their poor adversarial robustness, relative to that of random filters (Figure 6). Our next experiment disproves this.

We can artificially increase the smoothness of our *random* filters by applying a low-pass smoothing operation directly to the filters themselves. After doing so, the high-frequency components of our random filters have been removed, and the filters now focus more on the mid to low frequencies, just as the learned filters do, as seen in Figure 8.

In Figure 7, we see that instead of hurting performance, smoothing the random filters actually significantly *increases* their robustness to adversarial attacks while not hurting performance on clean data (orange line) at all. In fact, we see that smoothing results in a relative performance increase of over 25% for the highest magnitude of attack. We then conclude that there must be a different reason (statistical shortcuts, etc.) for the learned filters’ poor robustness.

4.4. Pixel Un-Shuffling

Thus far, we have only shown the power of channels-only models as they apply to classification tasks.

As evidence that such models can also prove effective for other tasks, we now apply them to the *pixel un-shuffling* problem. We define this problem as applying a deterministic random permutation of the spatial coordinates of all the pixels of the input image and tasking the model with reconstructing the un-corrupted signal. (To be clear, the same permutation pattern is applied to every input image.)

For images \mathbf{x}_i , pixel permutation σ , and network \mathcal{F} with

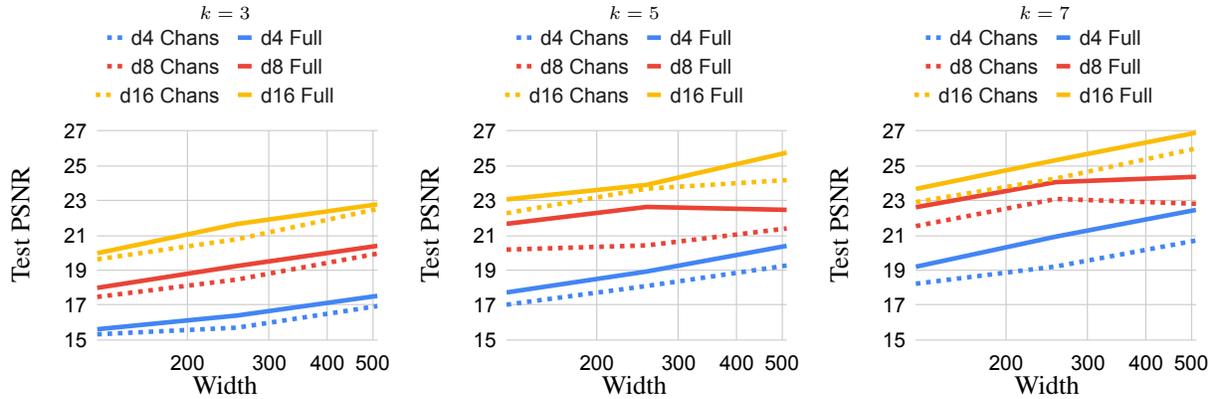


Figure 10. **CIFAR-10**: PSNR for Pixel Un-Shuffling. Once again, yet for a completely different task, we see that channels-only models only lag slightly behind their fully-learned counterparts over a wide variety of network sizes.

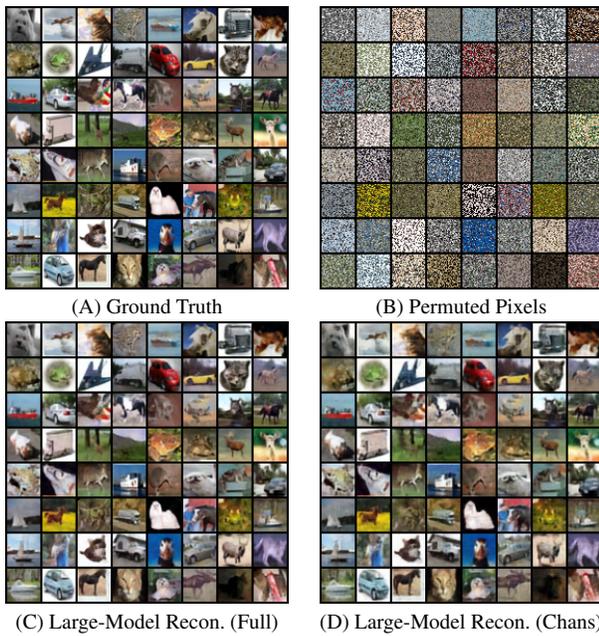


Figure 11. Channels-only un-shuffling models perform comparably to fully-learned models on the pixel un-shuffling task, indicating that the observed phenomena extends past classification.

parameters θ , our formal training objective becomes

$$\min_{\theta} \frac{1}{n} \sum_{i=1}^n \|\mathbf{x}_i - \mathcal{F}(\sigma(\mathbf{x}_i); \theta)\|^2$$

For this task, we use a patch-size 1 ConvMixer with a slight modification: instead of a global average pool and classification head, we simply apply a linear projection back up to the appropriate number of channels (1 or 3). We chose the ConvMixer because its isotropic structure makes it ideal for image-to-image translation, and its natively separable design allows for easy analysis of the effects of spatial and channel mixing.

Intuitively, one would think the learning of spatial mixing is crucial for the task of “moving” pixels back to their

correct location. However, our experiments show that, yet again, channels-only models perform nearly just as well as their fully-learned counterparts. Randomly selected examples from the test sets can be seen in Figure and 11.

On all of MNIST [9], Fashion-MNIST [30] and CIFAR-10 [18], the qualitative and quantitative test-set results for both model types are strikingly close. Unlike the classification problem, these results give us a clear, human-interpretable notion of the fact that networks can seemingly learn effective point-wise linear transforms to leverage the natural spatial mixing provided by the randomly-initialized depth-wise convolutions. In Figure 10, we include plots showing PNSR trends for various model architectures on CIFAR-10. We see that the performance increases with depth, width, and kernel size despite the parameter count of channels-only models not scaling *at all* with kernel size. (We include results for MNIST and Fashion-MNIST in the supplementary material.)

5. Conclusion

While previous works have shown the merits of fully frozen networks as feature extractors [17, 23], none (to the best of our knowledge) have explored the question of which parameters are the most important or beneficial to learn.

In this work, we have done just that; we have shown empirically that networks that only learn channel mixing can reach nearly the same performance as fully-learned networks without any further alterations. For intuition, we offer a possible explanation via the spectral envelopes of the random filter banks reaching dense coverage. We also showed that networks with random spatial-mixing weights are naturally more robust to adversarial attacks, and we also offer a method for *further* increasing the adversarial robustness of such networks by artificially altering their spectral coverage. Lastly, we showed that this phenomenon extends past the classification regime and that such restricted models are also capable of learning other tasks, such as pixel un-shuffling.

References

- [1] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dandelion Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org. [2](#)
- [2] George Cazenavette, Calvin Murdock, and Simon Lucey. Architectural adversarial robustness: The case for deep pursuit. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7150–7158, 2021. [7](#)
- [3] Anirban Chakraborty, Manaar Alam, Vishal Dey, Anupam Chattopadhyay, and Debdeep Mukhopadhyay. Adversarial attacks and defences: A survey, 2018. [7](#)
- [4] Mark Chen, Alec Radford, Rewon Child, Jeffrey Wu, Heewoo Jun, David Luan, and Ilya Sutskever. Generative pretraining from pixels. In *International Conference on Machine Learning*, pages 1691–1703. PMLR, 2020. [1](#), [2](#), [3](#)
- [5] Weijie Chen, Di Xie, Yuan Zhang, and Shiliang Pu. All you need is a few shifts: Designing efficient convolutional neural networks for image classification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7241–7250, 2019. [2](#)
- [6] Francois Chollet. Xception: Deep learning with depthwise separable convolutions. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017. [2](#), [3](#)
- [7] Hassan Dbouk and Naresh Shanbhag. Generalized depthwise-separable convolutions for adversarially robust and efficient neural networks. *Advances in Neural Information Processing Systems*, 34:12027–12039, 2021. [2](#)
- [8] 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*, pages 248–255. Ieee, 2009. [4](#), [5](#)
- [9] Li Deng. The mnist database of handwritten digit images for machine learning research. *IEEE Signal Processing Magazine*, 29(6):141–142, 2012. [8](#)
- [10] 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. In *International Conference on Learning Representations*, 2020. [1](#), [2](#), [3](#)
- [11] Jonathan Frankle and Michael Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. *arXiv preprint arXiv:1803.03635*, 2018. [2](#), [5](#)
- [12] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. *arXiv preprint arXiv:1412.6572*, 2014. [6](#), [7](#)
- [13] Dongyoon Han, YoungJoon Yoo, Beomyoung Kim, and Byeongho Heo. Learning features with parameter-free layers. In *International Conference on Learning Representations*, 2022. [2](#)
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015. [3](#), [4](#)
- [15] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016. [4](#)
- [16] Andrew G. Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications, 2017. [2](#), [3](#)
- [17] Guang-Bin Huang, Qin-Yu Zhu, and Chee-Kheong Siew. Extreme learning machine: theory and applications. *Neurocomputing*, 70(1-3):489–501, 2006. [1](#), [3](#), [8](#)
- [18] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. [4](#), [5](#), [8](#)
- [19] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012. [1](#)
- [20] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. [1](#)
- [21] Yuancheng Li and Yimeng Wang. Defense against adversarial attacks in deep learning. *Applied Sciences*, 9(1), 2019. [7](#)
- [22] Yaniv Romano, Aviad Aberdam, Jeremias Sulam, and Michael Elad. Adversarial noise attacks of deep learning architectures: Stability analysis via sparse-modeled signals. *Journal of Mathematical Imaging and Vision*, 62(3):313–327, 2020. [7](#)
- [23] Andrew M Saxe, Pang Wei Koh, Zhenghao Chen, Maneesh Bhand, Bipin Suresh, and Andrew Y Ng. On random weights and unsupervised feature learning. In *Icml*, 2011. [1](#), [3](#), [8](#)
- [24] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. [1](#)
- [25] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015. [1](#)
- [26] Ilya O. Tolstikhin, Neil Houlsby, Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Thomas Unterthiner, Jessica Yung, Andreas Steiner, Daniel Keysers, Jakob Uszkoreit, Mario Lucic, and Alexey Dosovitskiy. Mlp-mixer: An all-mlp architecture for vision. *CoRR*, abs/2105.01601, 2021. [1](#), [2](#), [3](#)
- [27] Asher Trockman and J Zico Kolter. Patches are all you need? *arXiv preprint arXiv:2201.09792*, 2022. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#)
- [28] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia

- Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017. [2](#)
- [29] Bichen Wu, Alvin Wan, Xiangyu Yue, Peter Jin, Sicheng Zhao, Noah Golmant, Amir Gholaminejad, Joseph Gonzalez, and Kurt Keutzer. Shift: A zero flop, zero parameter alternative to spatial convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 9127–9135, 2018. [2](#)
- [30] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017. [8](#)
- [31] Weihao Yu, Mi Luo, Pan Zhou, Chenyang Si, Yichen Zhou, Xinchao Wang, Jiashi Feng, and Shuicheng Yan. Metaformer is actually what you need for vision. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10819–10829, 2022. [2](#), [3](#)
- [32] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018. [2](#), [3](#)