
Self-Supervised Learning of a Biologically-Inspired Visual Texture Model

Nikhil Parthasarathy
Center for Neural Science
New York University
np1742@nyu.edu

Eero P. Simoncelli
Center for Neural Science, and Courant Inst. of Mathematical Sciences
New York University, and
Howard Hughes Medical Institute
eero.simoncelli@nyu.edu

Abstract

We develop a model for representing visual texture in a low-dimensional feature space, along with a novel self-supervised learning objective that is used to train it on an unlabeled database of texture images. Inspired by the architecture of primate visual cortex, the model uses a first stage of oriented linear filters (corresponding to cortical area V1), consisting of both rectified units (simple cells) and pooled phase-invariant units (complex cells). These responses are processed by a second stage (analogous to cortical area V2) consisting of convolutional filters followed by half-wave rectification and pooling to generate V2 ‘complex cell’ responses. The second stage filters are trained on a set of unlabeled homogeneous texture images, using a novel contrastive objective that maximizes the distance between the distribution of V2 responses to individual images and the distribution of responses across all images. When evaluated on texture classification, the trained model achieves substantially greater data-efficiency than a variety of deep hierarchical model architectures. Moreover, we show that the learned model exhibits stronger representational similarity to texture responses of neural populations recorded in primate V2 than pre-trained deep CNNs.

1 Introduction

Most images contain regions of "visual texture" - comprised of repeated elements, subject to some randomization in their location, size, color, orientation, etc. Humans are adept at recognizing and differentiating materials and objects based on their texture appearance, as well as using systematic variation in texture properties to recover surface shape and depth. At the same time, we are insensitive to the details of any particular texture example - to first approximation, different instances of any given class of texture are perceived as the same, as if they were "cut from the same cloth". This invariance is usually captured through the use of statistical models. Bela Julesz initiated the endeavor to build a statistical characterization of texture, hypothesizing that a texture could be modeled using n -th order joint co-occurrence statistics of image pixels [1]. Subsequent models can be partitioned into three broad categories: 1) orderless pooling of handcrafted raw-pixel features such as local binary patterns [2, 3], 2) local statistical models using Markov random fields [4–7], and 3) statistical characterization of fixed convolutional decompositions (i.e. wavelets, Gabor filters, multi-scale pyramids) [8–10, 7, 11, 12]. More recent models are based on statistics of nonlinear features

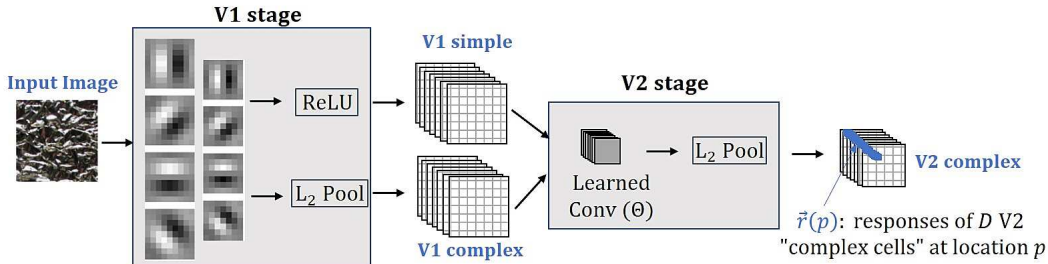


Figure 1: Proposed biologically-inspired texture model architecture. The V1 stage is built using a 5-scale 4-orientation complex steerable pyramid [7], followed by two nonlinearities to generate simple and complex cell responses. The latter uses specialized L_2 pooling. The V2 stage consists of convolution with D filters followed by spatial L_2 pooling.

extracted from pre-trained deep convolutional neural networks (CNN's) [13–17]. A comprehensive review of these is available in [18].

The fixed-filter methods are generally chosen to capture features considered fundamental for early visual processing, such as local orientation and scale. Similar filters can be learned using methods such as sparse coding [19] or independent components analysis [20]. On the other hand, deep learned methods provide great benefits in terms of extracting relevant complex features that are not so easily specified or even described.

However, recent work in understanding the representation of texture in the primate brain has shown that texture selectivity arises in Area V2 of visual cortex [21, 22], which receives primary input from Area V1. Therefore, it seems that the brain can achieve selectivity for texture in far fewer stages than are commonly used in the deep CNNs. Motivated by this fact, we construct a simple, hybrid texture model that blends the benefits of the aforementioned fixed-filter image decompositions with the power of learned representations. There are two main contributions of our work. First, the model represents textures in a relatively low-dimensional feature space (in contrast to the extremely high-dimensional representations found in CNN models). We propose that this low-dimensional representation can be used to perform texture family discrimination with small amounts of training data when it is coupled with an interpretable non-linear decoder. Moreover, we show that a novel self-supervised learning objective plays an important role in achieving this result. Finally, while pre-trained deep CNNs can achieve better texture classification accuracy, we show that our learned model exhibits much stronger representational similarity to texture responses of real neural populations recorded in primate V2.

2 Methods

2.1 V2Net Model Architecture

It is well-known that the primary inputs to V2 are feed-forward outputs from area V1 [23–25]. Inspired by these physiological results, we propose a computational texture model as a two-stage network that functionally mimics the processing in these two early visual areas.

The V1 stage is implemented using a set of fixed convolutional basis filters that serve as a functional model for V1 receptive fields [26]. The filters are localized in orientation and scale, specifically utilizing a complex-steerable derivative basis [27, 28]. We chose a specific set of 4 orientations and 5 scales (octave-spaced) with two phases (even and odd), for a total of 40 filters. The full set of V1 responses are a combination of both half-wave rectified simple cells and L_2 -pooled (square root of the sum of squares) complex cells, yielding a total of 60 feature maps.

The V1 responses provide input to a V2 stage that consists of a set of D learned convolutional filters. In the macaque, V1 and V2 are known to have similar cortical surface area and output fibers [29], so in our experiments we set $D = 60$ to match the dimensionality of the V1 and V2 stages of our model. The convolutional layer is then followed by half-wave rectification, spatial L_2 -pooling and downsampling to produce V2 'complex cell' responses (Fig. 1). Unlike standard max-pooling, L_2 -pooling is used in both stages of our model because it is more effective at capturing local energy of responses without introducing aliasing artifacts [11, 30].

2.2 Learning Objective

Consider the model in Fig. 1 as a function $f(T; \Theta)$ that takes as input a texture image T , and computes responses based on parameter vector $\Theta = [\Theta_1, \dots, \Theta_D]$, which contains the D V2 filters¹. Given a dataset of N texture images (T_n) and their corresponding model responses $\vec{r}_n = f(T_n; \Theta)$, we seek an objective function, $L(\cdot)$, for optimizing the V2 filter weights: $\Theta_{opt} = \arg \min_{\Theta} L(\{f(T_n; \Theta)\})$. We assume a curated image dataset with two properties that underlie the formulation of the objective: 1) individual images contain a single texture type (homogeneous across their spatial extent) and 2) the N images in the dataset represent a diverse set of texture types.

Our learning objective is motivated by the experimental observations in [22] suggesting that V2 represents textures such that responses within texture families (i.e. classes) are largely invariant to variability within the texture families- the responses are less variable within texture families than across families. To learn such a representation, one could simply utilize an objective function that reduces variability of responses to each family while maintaining variability across all families. This can usually be achieved by supervised methods that optimize responses to predict the class identity for an image. However, we desire an objective that has no supervisory knowledge of which images correspond to which texture families. As a result, we propose a contrastive objective that seeks to 1) Minimize the variability of model responses ($\vec{r}_n(p)$) across locations p within each individual texture image and 2) Maximize variability of these responses across neighborhoods sampled from the entire set of N images. Therefore, rather than using labels to enforce grouping of similar texture families, we utilize the natural spatial homogeneity of *individual texture images* as a form of ‘self-supervision’.

To formulate this mathematically, we first model the distribution of V2 responses over positions p within each image ($\vec{r}_n(p) \in \mathbb{R}^D$) as multivariate Gaussian, parameterized by the sample mean and covariance: $\vec{\mu}_n \in \mathbb{R}^D$ and $C_n \in \mathbb{R}^{D \times D}$. The global distribution of responses across all images is then a Gaussian mixture with mean and covariance: $\vec{\mu}_g = \frac{1}{N} \sum_{n=1}^N \vec{\mu}_n$; $C_g = \frac{1}{N} \sum_{n=1}^N C_n + (\vec{\mu}_n - \vec{\mu}_g)(\vec{\mu}_n - \vec{\mu}_g)^\top$. Under this parameterization, the two goals for the objective can be achieved by maximizing the ‘discriminability’ between the individual and global response distributions based on their covariances. A suitable measure of discriminability must capture the differences in both size (total variance) and shape of the distributions.

There has been extensive work on developing measures that approximate the discriminability between Gaussian distributions based on their mean and/or covariance statistics [31–37]. In order to choose a distance for this problem we define a set of criteria the distance must satisfy. First, the distance must be *scale invariant*: global rescaling of the image data should not change the value of the distance measure, which is meant to capture *relative* differences in variability. This is especially important for an objective function, as the responses can be arbitrarily scaled by the learned weights. Second, for maximization it is preferable that a distance *have an upper bound* as this can stabilize optimization and avoid degenerate solutions where the distance can take on extremely large, unbounded values. Third, for any given texture image, not all of the V2 dimensions may be important (i.e. the covariance is low-rank), so the distance must be stable in this regime.

Given these criteria, it is clear that many of the statistical distances and manifold-based log-Euclidean distances are problematic because the log transformation is unstable when covariances are low-rank. The work of [38] has shown that regularizing the log-Euclidean approach with standard covariance shrinkage can lead to large errors, and we have observed this in our experiments as well. A novel attempt to resolve this issue was proposed using a Riemannian optimization method [38], but this method only works for fixed low-rank matrices. As a result, we construct our distance on the form $\|C_1^{1/2} - C_2^{1/2}\|_F$ corresponding to the Bures metric² [31, 39]. We modify this to make it bounded and scale-invariant, arriving at a novel measure of distance between the global response covariance and that of image T_n :

$$d_n = \frac{\|C_g^{1/2} - C_n^{1/2}\|_F}{\|C_g^{1/2}\|_F}, \quad (1)$$

where $(\cdot)^{1/2}$ indicates matrix square-root and $\|\cdot\|_F$ is the Frobenius norm. This may be seen as a normalized variant of the log-Euclidean distances [37], in which replacement of $\log(\cdot)$ by $(\cdot)^{1/2}$

¹each Θ_d is a $60 \times 7 \times 7$ set of weights, as each V2 filter operates over the full set of 60 V1 channels

²Equivalent to the covariance term of the 2-Wasserstein distance between multivariate Gaussian distributions in the special case when the two covariance matrices commute

retains the primary benefit of the log-Euclidean framework (transforming the covariance eigenvalues with a compressive nonlinearity), while remaining stable and well-defined in low-rank conditions.

After calculating the distance in Eqn. (1) for each individual image, we then combine over all images to obtain a single scalar objective. To force all distances to be as large as possible, we maximize the minimum of these distances. For stable optimization, we use a soft-minimum function, which yields our variability-based objective:

$$\mathbf{L}_{\text{var}} = \text{softmin}(d_1, d_2, \dots, d_N) = \frac{\sum_n d_n e^{-d_n}}{\sum_n e^{-d_n}}. \quad (2)$$

To allow for robust estimation of the covariance, we make a diagonal approximation where C_n and C_g are each taken to be diagonal. Therefore, the matrix square-roots can be implemented as element-wise square roots of the individual response variances along the diagonal and the Frobenius norm becomes the standard vector L_2 norm. However, because a diagonal approximation can be poor if the covariances have strong co-variability, we use an additional orthogonal regularization term to encourage orthogonalization of the V2 filters [40]:

$$\mathbf{L}_{\text{orth}} = \|\Theta\Theta^\top - I\|_F. \quad (3)$$

Minimizing this loss forces the responses of each channel to be roughly independent and thus more amenable to the diagonal approximation. The final objective is a weighted combination of the two terms:

$$\max_{\Theta} [\mathbf{L}_{\text{var}} - \lambda\mathbf{L}_{\text{orth}}]. \quad (4)$$

2.3 Evaluation Methodology

After training the model with the self-supervised objective in Eqn. (4), we use a separate labeled dataset to train and test a texture family classifier. We first compute the spatially global-average pooled (GAP) responses for each image in the new dataset, such that each image T_n is represented by a single D -dimensional vector, $\vec{\mu}_n$. We again make a Gaussian assumption on the distribution of these mean response vectors for each texture family and fit and test a quadratic discriminant classifier (QDA) to predict the texture class labels. This process is shown in Fig. 2 (a). Although the choice of a QDA classifier is not common, state-of-the-art texture classification methods generally use some form of quadratic feature encoding (Fisher vectors, bilinear layers, etc.) before applying a trained linear classifier (i.e. SVM) [41, 16, 13]. Rather than compute all pairwise products, which can be prohibitively expensive in terms of number of parameters, we use mean pooling to produce a low-dimensional representation, followed by a bilinear readout. In our context, a quadratic discriminant is the optimal bilinear method for discrimination under the Gaussian assumption.

One issue with QDA classification is that it requires the estimation of class-covariance matrices. These matrices can only be reasonably estimated when the number of samples per class is much larger than the dimensionality of the features, so QDA is only amenable to low-dimensional feature representations. In our experiments, we compare our model to supervised methods that use popular network architectures as the base feature extractor. However, most of these networks produce very high-dimensional output feature spaces that are usually evaluated with linear classifiers. As a result, we devise an evaluation protocol for these methods such that the dimensionality of the feature representation *and* the expressivity of the classifiers is matched to that of our model (Fig. 2(b)). Specifically,

we first reduce the dimensionality of the feature representation to match that of our V2Net model ($D = 60$) using a trained 1×1 convolutional layer (as is common in the literature [42, 43]). This is followed by the same classification procedure as ours: spatial GAP followed by a bilinear classifier. However, because QDA cannot be implemented for supervised, end-to-end trained networks, we use

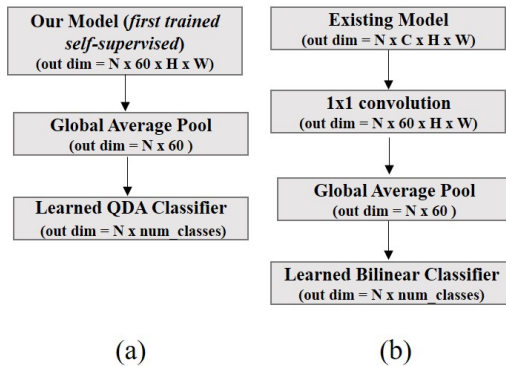


Figure 2: (a) Evaluation method for our self-supervised model. (b) Evaluation method for the supervised networks

a parameterizable bilinear layer of the form: $x^T Ax + Bx + c$. The model parameters, 1x1 conv layer, and bilinear are all trained end-to-end, in contrast with our model which is first trained separately with the self-supervised objective.

One might ask if the dimensionality reduction of the existing network architectures is too restrictive and if our comparisons will be biased because of this. In fact, a similar methodology has shown minimal loss in performance for texture retrieval with PCA reduction down to 64 dimensions [44]. Therefore, it is unlikely that we are biasing our comparisons by stifling the capacity of the network. Moreover, the 1x1 convolution approach is arguably more effective than PCA because it allows this dimensionality reduction to be optimized in the context of the classification task. Nevertheless, we additionally verified that results for all tested networks were close to those achieved using a linear classifier on the full-dimensional feature space.

The specific models we compare to are chosen to span a diverse set of methods from the literature:

ScatNet: We implement the front-end two-stage scattering model as described in [11, 12] that has 5 scales and 8 angles. The scattering model is then fixed and the 1x1 convolution layer and the bilinear classifier are learned. The number of channels before dimensionality reduction is 681.

DAWN(16-init): Recent work has performed a similar experiment using a hybrid deep adaptive wavelet network that is found to be more data-efficient than previous methods [45]. We implemented the same model and regularization, with 16 initial convolutional layers, followed by the multi-scale representation. The number of channels before dimensionality reduction is 256.

ResNet-18: Based on recent success as a feature extractor for texture recognition [43] we also included an 18-layer ResNet model. We extract features from the *layer4* level of the network, as these have been deemed as the most powerful features for texture classification in previous work [43, 46]. The number of channels before dimensionality reduction is 512.

VGG-16: VGG networks and their variants have been the most common network architectures used for feature extraction in the literature. The work of [47, 13] demonstrated that a Fisher vector decoder, and even linear classification from pooled features of the last convolutional layer, can be effective for texture classification. Based on this work, we used features from the *conv5* layer of a VGG-16 network. The number of channels before dimensionality reduction layer is 512.

3 Related Work

Model Architecture. Many fixed-filter, hierarchical image decompositions have been used in the construction of texture representations that are similar to our V1 stage [11, 27]. However, we note that our V1 responses include both rectified simple cells *and* L_2 -pooled complex cells. This formulation is motivated by physiological experiments studying the projections of V1 to V2 neurons [48], and represents a departure from the classical view of hierarchical visual modeling that assumes only pooled responses are transmitted to the downstream layers [49, 50, 11].

Recent deep learning approaches to representing texture have been heavily optimized and hand-crafted for specific tasks such as texture classification [13, 43], synthesis [15, 14], and retrieval [51, 44]. However, there are a few common themes in these methods that we highlight for their relevance to our model and the models we use for comparison. First, all SoA methods, regardless of task, rely on extraction of features or statistics from deep CNNs trained for object recognition, primarily the VGG and ResNet architectures [52, 53]. With the exception of a few studies [54, 45], performing texture classification with networks trained from scratch has been relatively understudied. Second, it has been consistently shown that "orderless" pooling of the features before classification layers results in a far better texture representation. Simple global average pooling (GAP) has been shown to be quite effective [44, 43, 55, 56] as well as methods that pool based on 2nd-order statistics [13, 14, 41].

Objective functions. In the context of texture classification, current human-labeled homogeneous texture databases are few and small, so most deep learning methods transfer features from networks trained with full supervision on an alternative task (typically, object recognition). Some authors have developed limited unsupervised methods based on vector quantization [57, 58], and non-negative matrix factorization [59]. Nevertheless, in concert with CNN models, we believe ours is the first competitive self-supervised learning objective for this problem.

Conceptually, our objective is inspired by principles of *contrastive learning* that have recently seen much success at competing with more traditional supervised methods [60–64]. However, the specific construction of our learning objective differs substantially from these methods as it relies on a diagonal Gaussian parameterization of sample distributions that provides many computational benefits such as easy generalization to incremental learning where the sufficient statistics are updated online without use of large in-memory batches.

4 Results

4.1 Data-Efficient Texture Classification

We hypothesized that our objective function enables the learning of a more powerful texture representation from small data. To test this, we used an experimental paradigm similar to [64]. We trained and tested all models on varying amounts of data from a texture dataset. We used a modified version of the challenging KTH-TIPS2-b dataset [65] for both training and evaluation. The original dataset includes 11 families of textured materials photographed with different viewpoints, illumination levels, and scales. The total dataset is relatively small (4752 images), so we augmented it with 3 rotated versions of each image (90, 180, and 270 degrees) to obtain a total of 19008 samples. As texture representations should be invariant to rotation, this is a sensible augmentation that increases the difficulty of the task. We used the original 4 splits of the KTH-TIPS2-b data (training on 3 splits and testing on the 4th). For all experiments we used a fixed validation set of 3256 images and each test set contained 4752 images. We then conducted three experiments varying the amount of training data (reducing evenly the number of images per texture family). We report results for the full training data (1000 images per family), 50 percent training (500 images per family), and 25 percent (250 images per family).

All models (ours and those listed in Sec. 2.3) were trained from scratch without any pre-trained information. For the supervised networks we varied learning rates (from 0.0001 to 0.01) and batch sizes from (50 to 200) and chose the best model for each train/test split. For our model (V2Net (self-sup)), the objective function relies on calculating the global mean and variances over the entire dataset. However, because our training is done through stochastic gradient descent, we approximated these global statistics by the global statistics over batches of 275 images. We chose the batch size heuristically so that individual batch statistics do not deviate significantly from the statistics over the whole dataset. Interestingly, the batch size does not need to be as large as is necessary in most other contrastive learning approaches [64, 66]. We used a learning rate of 0.001 and additionally included a BatchNorm layer at the output of the network to stabilize the global statistics across batches.




Figure 4: V1 and V2 model comparison based on ablating V1 simple cell contribution.

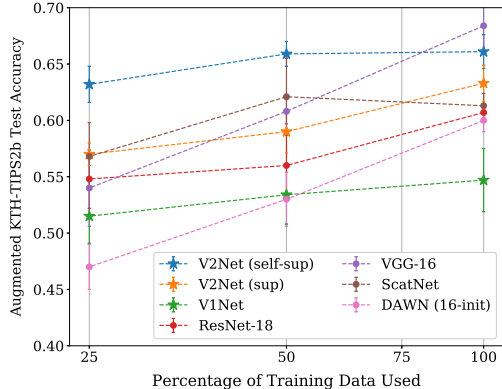


Figure 3: Mean and standard error across the 4 train/test splits as a function of the percentage of training data used.

The results for the training experiments are shown in Fig. 3. We report the mean and standard error for the 4 train/test splits within each experiment. First, we can see that just using the fixed V1 stage (V1Net) followed by QDA provides a reasonable baseline. This model has marginal performance difference across differing amounts of training data, which can be solely attributed to the estimation error of the class covariances when training the QDA classifier. Second, we find that the two-stage V2 model performs similarly to the VGG-16 network with full training data, but significantly outperforms all networks when using 50% or 25% of the training data, indicating much greater data-efficiency. To better understand the impact of our objective function,

we also report results for a network with the same architecture as V2Net, but trained with a supervised cross-entropy loss (V2Net (sup))³. As seen in Fig. 3, this network performs comparably to the other supervised networks but still seems to overfit in the small-data regimes. This suggests that even small networks can overfit with small amounts of training data, implying that it is the design of our objective function that allows our network to remain data-efficient in these cases.

To assess the impact of the inclusion of V1 simple cells in our network, we compared the V2Net classification accuracy to a model trained with V1 simple cells removed. The results are shown in Fig. 4. The performance of both V1 models is roughly the same, and in both cases the V2 model improves on the V1 model. However, the gap between the V1 and V2 performance is noticeably larger when the V1 layer contains both simple and complex cells. This result suggests that a more effective V2 representation can be learned when the inputs come from both simple and complex cells.

4.2 Transfer Learning

To verify the generalization of our learning objective, we collected an unlabeled dataset of texture photographs. Original images were manually cropped to be globally homogeneous (by eye) over their entire spatial extent. The scale, viewpoint etc. were not controlled in any particular way, although most textures are on approximately front-parallel surfaces. The types of texture in the dataset span a wide range (including leaves, grass, wood bark, brick, ceramic tile mosaics, etc) that is far more diverse than the KTH-TIPS2-b dataset. We trained our model on 11000 of these images and re-evaluated the performance on the four KTH train/test splits by retraining the QDA classifier. Performance of this pre-trained model slightly improves on the performance of the models trained from scratch (average gain of 1.4 % mean accuracy across the three experiments) and displays the same level of robustness to the reduction of training data. This demonstrates that our results are not specific to the training dataset and that our learning objective in fact generalizes across texture datasets with very different distributions of images. We additionally compared the performance of our pre-trained (but still self-supervised) network against the ResNet-18 and VGG-16 architectures pre-trained on ImageNet classification. The results of this experiment are given in Appendix A. Our network does not achieve the performance of these pre-trained networks, but the performance gap (5-10%) is surprisingly small given that our model is pre-trained without supervision, using two orders of magnitude fewer images (11k vs. 1M).

4.3 Selectivity for Natural Texture vs. Spectrally-shaped Noise

Physiological results in [21, 22] suggest that texture selectivity in the brain not only manifests as an ability to separate texture families, but also can also be used to distinguish natural textures from their phase-scrambled counterparts. We constructed a test along these lines to gain a deeper understanding of our learned model and its selectivities. We retrained our V2Net model using phase-scrambled versions of the images from our unlabeled texture dataset from Sec. 4.2. By training on phase-scrambled images, the model no longer has access to the natural statistics that define textures beyond their spectral power. As a result, if our model is truly capturing higher-order texture statistics, its performance on natural images will drop significantly when trained on the phase-scrambled images. In fact, we find that the average test accuracy of the model trained on phase-scrambled images (V2Net (PS)) is 51.5% vs. 67.4% for the model trained on natural images (V2Net (Natural)). Upon further inspection, there are certain texture classes that have high accuracy for the V2Net (PS) model, indicating that these families are readily distinguished using spectral power statistics. We verified that this is also true perceptually: phase-scrambled versions of these classes are visually similar to the original images. However, the classes where there is a large deviation between V2Net (Natural) and V2Net (PS) are those where the phase-scrambled images carry little information about the original texture. For more details, see Appendix B.

4.4 Texture Representational Similarity

Having established that our learned texture model reproduces the qualitative texture selectivities seen in populations of V2 neurons, we explored this relationship quantitatively by comparing the representational similarity between our model and recorded responses of V2 neurons to texture images. We used the dataset described in [21, 22], which provides electrophysiological recordings of

³We use the same bilinear classifier model as was used for the other networks

	V2Net	V2Net (Comp)	Res (layer1)	VGG (block1pool)
Spearman Corr.	0.658	0.597	0.410	0.405

Table 1: Spearman correlation between model RDMs and the V2 neural data RDM representing texture families. V2Net (Comp) refers to the V2Net model trained with only V1 complex cells.

103 V2 neurons responding to 15 samples of textures from 15 different texture families. To understand the representational similarity between our model and the neural data at the level of texture families, we first computed the averaged response (across samples) of both the model and neural responses to each of the 15 texture families. Next, for each representation, we constructed a dissimilarity matrix based on the pairwise correlation distance. There are many distances one could choose but the correlation distance is one of the most common and performs fairly robustly in comparison with distances such as euclidean distances [67, 68]. As has been noted in the literature [69], it is not common to assume a linear relationship between dissimilarity matrices, but it is rather more appropriate to assume the model RDM predicts the rank order of the dissimilarities [69]. Therefore, we computed the Spearman rank correlation between the dissimilarity matrices of our model and the V2 neural data. We repeated this process for our V2Net model that uses only V1 complex cells (V2Net (Comp)). Finally, we performed the same analysis for all of the major layers from pre-trained ResNet-18 and pre-trained VGG-16 networks, reporting results for the layer with the best correlation. The results are summarized in Table 1. For more details on the physiology data and image presentation see Appendix C.

We find that both V2Net representations are significantly more correlated with the V2 population representation than either deep CNN, and that the inclusion of simple cells again offers a noticeable improvement. Additionally, for both of the pre-trained deep networks we find that the best performance occurs in early layers (layer1 of ResNet and block1pool of VGG). These layers are arguably at the right level in the visual hierarchy (2 or 3 convolutional layers deep) to be matched to V2, but we note that previous work has suggested that texture classification improves when taking features from deeper layers (e.g., VGG conv5) [13]. We hypothesize that this is a result of the supervised learning objective used to train these networks. Because of this, the pre-trained deep networks are able to achieve higher overall texture classification accuracy than our current model, but are not able to capture the physiology as well. This suggests that stacking a hierarchical model on top of our learned network may lead to an improvement in SoA classification performance while maintaining consistency with biological architectures.

5 Discussion

In this work, we demonstrate successful data-efficient self-supervised learning of a simple, yet powerful computational model for representing texture. Rather than learn a very high-dimensional representation followed by linear classification, we use a simpler two-stage model whose responses are then decoded with an *interpretable* non-linear decoder (QDA). This provides the benefit that moving forward we can more easily probe the underlying learned feature space and understand explicitly how those features impact decoding of texture families (through their covariance structure). In fact, we are not the first to propose such a scheme in the context of neural decoding as QDA has been shown to provide a possible basis for a biologically-plausible non-linear decoding method that can explain quadratic transformations that have been observed between layers of processing in the visual system [70, 71]. Within this framework, we show that a modification of the common view of hierarchical visual processing (reminiscent of skip-connections [53]), that includes both V1 simple and complex cells as input to a second V2-like processing stage can provide functional benefits in the learning of the texture representation both in terms of classification accuracy and representation similarity with recording neurons in primate area V2. More importantly, we demonstrate that smaller networks do not necessarily perform much better with small training data, but that learning robustly from small numbers of training examples required the development of a novel self-supervised learning objective.

Our learning objective is inspired by recent unsupervised contrastive objectives (separating positive examples from a collection of negatives) [63, 62, 64, 61]. While these methods are general, in they are non-parametric with respect to the distribution of the data, we believe that our parameterization in terms of mean and covariance allows our method to 1) constrain learning in small data regimes

and 2) provide opportunities to explore more biologically plausible on-line learning implementations. In particular, it is implausible that the brain can store all samples of the global distribution, and our parameterization allows for on-line sequential update of the mean and covariance statistics for each observed image.

Finally, our model currently assumes a dataset of homogeneous textures as input, enabling a simple form of objective that minimizes spatial variability of the responses across each image. We are currently extending this to allow learning from whole natural scenes, by minimizing variability of responses within *local* spatial neighborhoods, while maximizing global variability. This is motivated by the local consistency of natural images - nearby spatial regions are more likely to be similar than distant ones. In fact, there have been some efforts to use spatial coherence as a learning signal [72–75], splitting the image into independent patches that are processed as inputs to the model during learning. Our objective offers an alternate methodology that can process full images while imposing the locality constraint in the response space. Because of the layer-wise nature of our objective, there is also the potential to extend the method to learn filters in multiple stages of a hierarchical model.

Broader Impact

The interplay between machine learning and neuroscience is something that has helped progress both fields throughout the previous decades. Our work lies at the intersection of these two fields and aims to provide new accounts of how the visual system processes information by utilizing computational methods and constraints from physiology. In the context of machine learning, achieving generalizable unsupervised learning is widely considered to be one of the most important open problems, and our work provides a novel unsupervised method for learning representations throughout a hierarchical model, that is additionally more biologically plausible due to its self-supervised and layer-wise nature. Our work also provides insight into understanding how the primate visual system encodes natural images. Much of the field of visual neuroscience has focused on building models of either early sensory areas through primary visual cortex, or late-sensory areas (i.e V4/IT) that can be to some extent captured by deep network models. Computational models of mid-visual areas are lacking in comparison. The model described here aims to bridge this gap, both from a perspective of understanding the feature representation in area V2, but also providing physiologists with stronger hypotheses and experimental design constraints for probing mid-visual areas.

In terms of social and ethical implications, the high-level goal of generalizable unsupervised learning has the risk of creating uninterpretable AI models that will be used as replacements for more interpretable human-level decision-making. It is generally thought that lower dimensional models such as ours can allow for greater interpretability and analysis. Finally, to the extent that it provides a good description of primate mid-visual processing, our model offers potentially important societal value in the context of medical applications such as human visual prosthesis development.

Acknowledgments

This work has been supported by the Howard Hughes Medical Institute (EPS and NP) and the NIH Training Grant in Visual Neuroscience (T32 EY007136-27)

References

- [1] Bela Julesz. Visual pattern discrimination. *IRE transactions on Information Theory*, 8(2):84–92, 1962.
- [2] Timo Ojala, Matti Pietikäinen, and Topi Mäenpää. Multiresolution gray-scale and rotation invariant texture classification with local binary patterns. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (7):971–987, 2002.
- [3] Li Liu, Paul Fieguth, Xiaogang Wang, Matti Pietikäinen, and Dewen Hu. Evaluation of lbp and deep texture descriptors with a new robustness benchmark. In *European Conference on Computer Vision*, pages 69–86. Springer, 2016.
- [4] George R Cross and Anil K Jain. Markov random field texture models. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (1):25–39, 1983.
- [5] Rama Chellappa and Shankar Chatterjee. Classification of textures using gaussian markov random fields. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 33(4):959–963, 1985.

- [6] Haluk Derin and Howard Elliott. Modeling and segmentation of noisy and textured images using gibbs random fields. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (1):39–55, 1987.
- [7] Javier Portilla and Eero P Simoncelli. A parametric texture model based on joint statistics of complex wavelet coefficients. *International journal of computer vision*, 40(1):49–70, 2000.
- [8] Alan C. Bovik, Marianna Clark, and Wilson S. Geisler. Multichannel texture analysis using localized spatial filters. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (1):55–73, 1990.
- [9] James R Bergen and Edward H Adelson. Visual texture segmentation based on energy measures (a). *J. Opt. Soc. Am. A*, vol. 3, page P99, 3, 1986.
- [10] David J Heeger and James R Bergen. Pyramid-based texture analysis/synthesis. In *Proceedings of the 22nd annual conference on Computer graphics and interactive techniques*, pages 229–238. Citeseer, 1995.
- [11] Joan Bruna and Stéphane Mallat. Invariant scattering convolution networks. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1872–1886, 2013.
- [12] Laurent Sifre and Stéphane Mallat. Rotation, scaling and deformation invariant scattering for texture discrimination. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1233–1240, 2013.
- [13] Mircea Cimpoi, Subhansu Maji, and Andrea Vedaldi. Deep filter banks for texture recognition and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3828–3836, 2015.
- [14] Leon Gatys, Alexander S Ecker, and Matthias Bethge. Texture synthesis using convolutional neural networks. In *Advances in neural information processing systems*, pages 262–270, 2015.
- [15] Dmitry Ulyanov, Andrea Vedaldi, and Victor Lempitsky. Improved texture networks: Maximizing quality and diversity in feed-forward stylization and texture synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 6924–6932, 2017.
- [16] Yang Song, Fan Zhang, Qing Li, Heng Huang, Lauren J O’Donnell, and Weidong Cai. Locally-transferred fisher vectors for texture classification. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 4912–4920, 2017.
- [17] Jia Xue, Hang Zhang, Kristin Dana, and Ko Nishino. Differential angular imaging for material recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 764–773, 2017.
- [18] Li Liu, Jie Chen, Paul Fieguth, Guoying Zhao, Rama Chellappa, and Matti Pietikäinen. From bow to cnn: Two decades of texture representation for texture classification. *International Journal of Computer Vision*, 127(1):74–109, 2019.
- [19] Bruno A Olshausen and David J Field. Emergence of simple-cell receptive field properties by learning a sparse code for natural images. *Nature*, 381(6583):607, 1996.
- [20] Anthony J Bell and Terrence J Sejnowski. The “independent components” of natural scenes are edge filters. *Vision research*, 37(23):3327–3338, 1997.
- [21] Jeremy Freeman, Corey M Ziemba, David J Heeger, Eero P Simoncelli, and J Anthony Movshon. A functional and perceptual signature of the second visual area in primates. *Nature neuroscience*, 16(7):974, 2013.
- [22] Corey M Ziemba, Jeremy Freeman, J Anthony Movshon, and Eero P Simoncelli. Selectivity and tolerance for visual texture in macaque v2. *Proceedings of the National Academy of Sciences*, 113(22):E3140–E3149, 2016.
- [23] PASCAL Girard and JEAN Bullier. Visual activity in area v2 during reversible inactivation of area 17 in the macaque monkey. *Journal of neurophysiology*, 62(6):1287–1302, 1989.
- [24] Lawrence C Sincich and Jonathan C Horton. The circuitry of v1 and v2: integration of color, form, and motion. *Annu. Rev. Neurosci.*, 28:303–326, 2005.
- [25] Peter H Schiller and Joseph G Malpeli. The effect of striate cortex cooling on area 18 cells in the monkey. *Brain research*, 1977.
- [26] Dario L Ringach. Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *Journal of neurophysiology*, 88(1):455–463, 2002.

- [27] Eero P Simoncelli and William T Freeman. The steerable pyramid: A flexible architecture for multi-scale derivative computation. In *Proceedings., International Conference on Image Processing*, volume 3, pages 444–447. IEEE, 1995.
- [28] Jorn-Henrik Jacobsen, Jan van Gemert, Zhongyu Lou, and Arnold WM Smeulders. Structured receptive fields in cnns. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2610–2619, 2016.
- [29] Pascal Wallisch and J Anthony Movshon. Structure and function come unglued in the visual cortex. *Neuron*, 60(2):195–197, 2008.
- [30] Olivier J Hénaff and Eero P Simoncelli. Geodesics of learned representations. *arXiv preprint arXiv:1511.06394*, 2015.
- [31] Donald Bures. An extension of kakutani’s theorem on infinite product measures to the tensor product of semifinite w^* -algebras. *Transactions of the American Mathematical Society*, 135:199–212, 1969.
- [32] Anil Bhattacharyya. On a measure of divergence between two multinomial populations. *Sankhyā: the indian journal of statistics*, pages 401–406, 1946.
- [33] Karim T Abou-Moustafa, Fernando De La Torre, and Frank P Ferrie. Designing a metric for the difference between gaussian densities. In *Brain, Body and Machine*, pages 57–70. Springer, 2010.
- [34] Fernando De la Torre and Takeo Kanade. Multimodal oriented discriminant analysis. In *Proceedings of the 22nd international conference on Machine learning*, pages 177–184. ACM, 2005.
- [35] Zoran Nenadic. Information discriminant analysis: Feature extraction with an information-theoretic objective. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (8), 2007.
- [36] Ian L Dryden, Alexey Koloydenko, Diwei Zhou, et al. Non-euclidean statistics for covariance matrices, with applications to diffusion tensor imaging. *The Annals of Applied Statistics*, 3(3):1102–1123, 2009.
- [37] Zhiwu Huang, Ruiping Wang, Shiguang Shan, Xianqiu Li, and Xilin Chen. Log-euclidean metric learning on symmetric positive definite manifold with application to image set classification. In *International conference on machine learning*, pages 720–729, 2015.
- [38] Masoud Faraki, Mehrtash T Harandi, and Fatih Porikli. Image set classification by symmetric positive semi-definite matrices. In *2016 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 1–8. IEEE, 2016.
- [39] Boris Muzellec and Marco Cuturi. Generalizing point embeddings using the wasserstein space of elliptical distributions. In *Advances in Neural Information Processing Systems*, pages 10237–10248, 2018.
- [40] Nitin Bansal, Xiaohan Chen, and Zhangyang Wang. Can we gain more from orthogonality regularizations in training deep cnns? In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 4266–4276. Curran Associates Inc., 2018.
- [41] Tsung-Yu Lin, Aruni RoyChowdhury, and Subhransu Maji. Bilinear cnn models for fine-grained visual recognition.
- [42] 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. *arXiv preprint arXiv:1704.04861*, 2017.
- [43] Jia Xue, Hang Zhang, and Kristin Dana. Deep texture manifold for ground terrain recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 558–567, 2018.
- [44] Augusto C Valente, Fábio VM Perez, Guilherme AS Megeto, Marcos H Cascone, Otavio Gomes, Thomas S Paula, and Qian Lin. Comparison of texture retrieval techniques using deep convolutional features. *Electronic Imaging*, 2019(8):406–1, 2019.
- [45] Maria Ximena Bastidas Rodriguez, Adrien Gruson, Luisa Polania, Shin Fujieda, Flavio Prieto, Kohei Takayama, and Toshiya Hachisuka. Deep adaptive wavelet network. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 3111–3119, 2020.
- [46] Wei Zhai, Yang Cao, Jing Zhang, and Zheng-Jun Zha. Deep multiple-attribute-perceived network for real-world texture recognition. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 3613–3622, 2019.

- [47] Mircea Cimpoi, Subhansu Maji, Iasonas Kokkinos, Sammy Mohamed, and Andrea Vedaldi. Describing textures in the wild. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3606–3613, 2014.
- [48] Yasmine El-Shamayleh, Romesh D Kumbhani, Neel T Dhruv, and J Anthony Movshon. Visual response properties of v1 neurons projecting to v2 in macaque. *Journal of Neuroscience*, 33(42):16594–16605, 2013.
- [49] Kuniyuki Fukushima. Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological cybernetics*, 36(4):193–202, 1980.
- [50] Maximilian Riesenhuber and Tomaso Poggio. Hierarchical models of object recognition in cortex. *Nature neuroscience*, 2(11):1019–1025, 1999.
- [51] Yu Qian, Eduard Vazquez, and Biswa Sengupta. Differential geometric retrieval of deep features. In *2017 IEEE International Conference on Data Mining Workshops (ICDMW)*, pages 539–544. IEEE, 2017.
- [52] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [53] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [54] Shin Fujieda, Kohei Takayama, and Toshiya Hachisuka. Wavelet convolutional neural networks. *arXiv preprint arXiv:1805.08620*, 2018.
- [55] Weixia Zhang, Kede Ma, Guangtao Zhai, and Xiaokang Yang. Uncertainty-aware blind image quality assessment in the laboratory and wild. *arXiv preprint arXiv:2005.13983*, 2020.
- [56] Vincent Dumoulin, Jonathon Shlens, and Manjunath Kudlur. A learned representation for artistic style. *arXiv preprint arXiv:1610.07629*, 2016.
- [57] H Greenspan, R Goodman, and R Chellappa. Texture analysis via unsupervised and supervised learning. In *IJCNN-91-Seattle International Joint Conference on Neural Networks*, volume 1, pages 639–644. IEEE, 1991.
- [58] PP Raghu, R Poongodi, and B Yegnanarayana. Unsupervised texture classification using vector quantization and deterministic relaxation neural network. *IEEE Transactions on Image Processing*, 6(10):1376–1387, 1997.
- [59] Lei Qin, Qingfang Zheng, Shuqiang Jiang, Qingming Huang, and Wen Gao. Unsupervised texture classification: Automatically discover and classify texture patterns. *Image and Vision Computing*, 26(5): 647–656, 2008.
- [60] Olivier J Hénaff, Johannes Ballé, Neil C Rabinowitz, and Eero P Simoncelli. The local low-dimensionality of natural images. *arXiv preprint arXiv:1412.6626*, 2014.
- [61] Zhirong Wu, Yuanjun Xiong, Stella Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance-level discrimination. *arXiv preprint arXiv:1805.01978*, 2018.
- [62] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018.
- [63] Chengxu Zhuang, Alex Lin Zhai, and Daniel Yamins. Local aggregation for unsupervised learning of visual embeddings. *arXiv preprint arXiv:1903.12355*, 2019.
- [64] Olivier J Hénaff, Ali Razavi, Carl Doersch, SM Eslami, and Aaron van den Oord. Data-efficient image recognition with contrastive predictive coding. *arXiv preprint arXiv:1905.09272*, 2019.
- [65] Barbara Caputo, Eric Hayman, and P Mallikarjuna. Class-specific material categorisation. In *Tenth IEEE International Conference on Computer Vision (ICCV’05) Volume 1*, volume 2, pages 1597–1604. IEEE, 2005.
- [66] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. *arXiv preprint arXiv:2002.05709*, 2020.
- [67] Johannes Mehrer, Courtney J Spoerer, Nikolaus Kriegeskorte, and Tim C Kietzmann. Individual differences among deep neural network models. *bioRxiv*, 2020.

- [68] Nikolaus Kriegeskorte, Marieke Mur, and Peter A Bandettini. Representational similarity analysis—connecting the branches of systems neuroscience. *Frontiers in systems neuroscience*, 2:4, 2008.
- [69] Hamed Nili, Cai Wingfield, Alexander Walther, Li Su, William Marslen-Wilson, and Nikolaus Kriegeskorte. A toolbox for representational similarity analysis. *PLoS computational biology*, 10(4), 2014.
- [70] Marino Pagan, Eero P Simoncelli, and Nicole C Rust. Neural quadratic discriminant analysis: Nonlinear decoding with v1-like computation. *Neural computation*, 28(11):2291–2319, 2016.
- [71] Qianli Yang, Edgar Walker, R James Cotton, Andreas S Tolias, and Xaq Pitkow. Revealing nonlinear neural decoding by analyzing choices. *BioRxiv*, page 332353, 2020.
- [72] Suzanna Becker and Geoffrey E Hinton. Spatial coherence as an internal teacher for a neural network. *Backpropagation: Theory, architecture and applications*, pages 313–349, 1995.
- [73] Xu Ji, João F Henriques, and Andrea Vedaldi. Invariant information distillation for unsupervised image segmentation and clustering. *arXiv preprint arXiv:1807.06653*, 2018.
- [74] Dov Danon, Hadar Averbuch-Elor, Ohad Fried, and Daniel Cohen-Or. Unsupervised natural image patch learning. *Computational Visual Media*, 5(3):229–237, 2019.
- [75] Neal Jean, Sherrie Wang, Anshul Samar, George Azzari, David Lobell, and Stefano Ermon. Tile2vec: Unsupervised representation learning for spatially distributed data. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 3967–3974, 2019.

Appendix

A Transfer Learning with Pre-trained Networks

In addition to comparing networks trained from scratch on our modified KTH dataset, we also tested the performance of features transferred from pre-trained versions of our V2Net model, VGG-16, and ResNet-18. We pre-trained our model on a dataset of 11000 unlabeled image patches using our self-supervised objective. Example images from this dataset are provided in Fig. 5:




Figure 5: Example texture images from our hand-curated dataset, comprised of a large collection of natural textures that are unlabelled, but diverse in content and homogeneous across their spatial extent.

The VGG and ResNet networks are pre-trained on the supervised task of object recognition using 1 million images from the ImageNet database.

We used these pre-trained networks as feature extractors, and retrained the respective classifiers (See Fig. 2) for texture classification. Results are shown in Fig. 6.

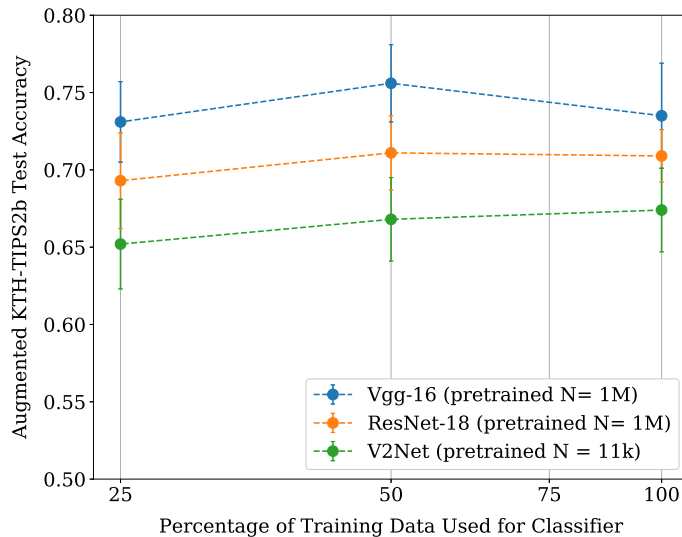


Figure 6: Mean and standard error computed across the 4 train/test splits on our KTH dataset (for each experiment where 25 %, 50 % and 100 % of training data is used to train the classifier weights). N refers to the number of images used to pre-train each model.

For all of the models, we find the classifier does not require large amounts of training data - performance is relatively constant across the different amounts of training data used. For the full (100%) classifier training set, our model achieves 67% - the performance gap (5-10%) relative to the pre-trained CNNs is surprisingly small given that our model is pre-trained without supervision, using two orders of magnitude fewer images (11k vs. 1M).

B Selectivity for Natural Texture vs. Spectrally-shaped Noise

For each of the 11 texture families in the test dataset, we plot the mean accuracy our model trained on natural images (V2Net (Natural) vs. our model trained on phase-scrambled images V2Net (PS)). Fig. 7 shows that the model trained on natural images performs better for most texture families, since it is able to capture higher-order natural statistics. If we visualize an example of one of these classes (aluminum foil), we see that this is because the scrambling of phase destroys content that is critical in defining that texture. However, for a few families, the performance of the V2Net (PS) model is about the same as the V2Net (Natural) model because certain texture families (e.g. wood) are primarily defined by their spectral content (and thus not altered significantly by phase-scrambling).

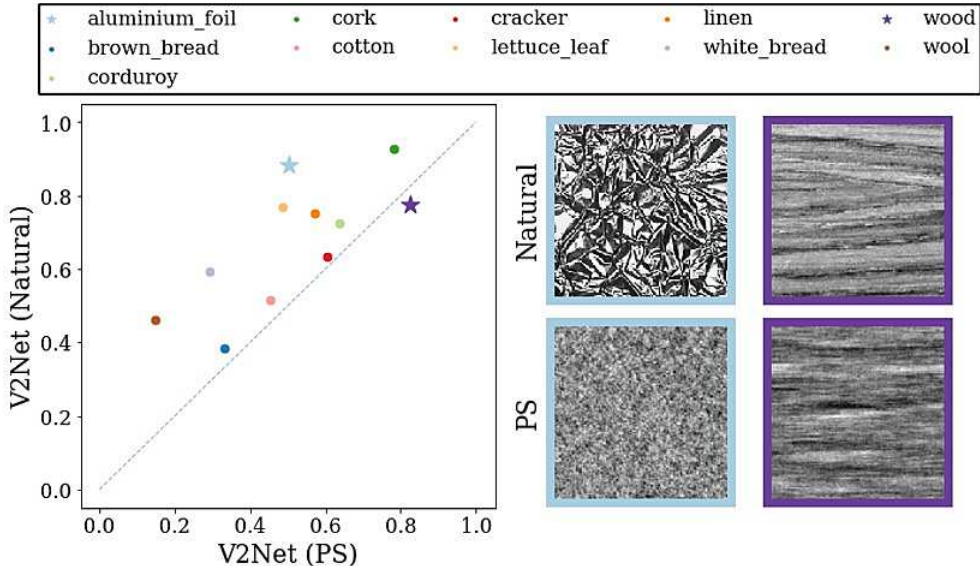


Figure 7: Left: Average KTH test accuracy (averaged over 4 splits with 100 % training data) for the V2Net (Natural) vs. V2Net (PS) models. Right: example images (both natural and phase-scrambled) for two texture families. For the ‘aluminum foil’ family, the phase-scrambled image removes the higher-order content that is necessary for identifying the texture. For the ‘wood’ family, the phase-scrambling does not alter perception of the texture significantly, because its appearance is primarily determined by spectral content.

C Representational Similarity Methods and Data

Here, we provide more details about the dataset and methods used for the representational similarity analysis presented in the main text. The neural data taken from [22] consists of electrophysiological recordings of 103 V2 neurons from anesthetized adult macaque monkeys. As is done in the original analysis, we averaged spike counts within 100-ms time windows aligned to the response onset for each single unit. To gaussianize the neural responses, we applied a variance-stabilizing transformation to the spike counts for each neuron ($r_{gauss} = \sqrt{r_{poiss}} + \sqrt{r_{poiss} + 1}$).

The visual stimuli used in the experiment are synthetic texture stimuli generated using the procedure described in [7]. A set of 15 grayscale texture photographs are used as the examples for 15 different texture families. From these seed images, 15 samples are generated for each family to provide sample variation across the family. The original stimuli have a size of 320 x 320 pixels and are presented to every V2 unit at a size of 4°, within a raised cosine aperture (this window was larger than all of the receptive fields of the neurons at the recorded eccentricities). For our representational similarity experiments, we thus pre-processed the images for input to the models such that they are resized to the appropriate pixel dimensions (224 x 224) and presented within a 4° raised cosine aperture. This ensures that the models receive inputs with content comparable to what is seen by the neurons in the V2 recordings.