
Explaining V1 Properties with a Biologically Constrained Deep Learning Architecture

Galen Pogoncheff

Department of Computer Science
University of California, Santa Barbara
Santa Barbara, CA 93106
galenpogoncheff@ucsb.edu

Jacob Granley

Department of Computer Science
University of California, Santa Barbara
Santa Barbara, CA 93106
jgranley@ucsb.edu

Michael Beyeler

Department of Computer Science,
Department of Psychological and Brain Sciences
University of California, Santa Barbara
Santa Barbara, CA 93106
mbeyeler@ucsb.edu

Abstract

Convolutional neural networks (CNNs) have recently emerged as promising models of the ventral visual stream, despite their lack of biological specificity. While current state-of-the-art models of the primary visual cortex (V1) have surfaced from training with adversarial examples and extensively augmented data, these models are still unable to explain key neural properties observed in V1 that arise from biological circuitry. To address this gap, we systematically incorporated neuroscience-derived architectural components into CNNs to identify a set of mechanisms and architectures that comprehensively explain neural activity in V1. We show drastic improvements in model-V1 alignment driven by the integration of architectural components that simulate center-surround antagonism, local receptive fields, tuned normalization, and cortical magnification. Upon enhancing task-driven CNNs with a collection of these specialized components, we uncover models with latent representations that yield state-of-the-art explanation of V1 neural activity and tuning properties. Our results highlight an important advancement in the field of NeuroAI, as we systematically establish a set of architectural components that contribute to unprecedented explanation of V1. The neuroscience insights that could be gleaned from increasingly accurate in-silico models of the brain have the potential to greatly advance the fields of both neuroscience and artificial intelligence.

1 Introduction

Advances in neuroscience have long been proposed as essential to realizing the next generation of artificial intelligence (AI). Many influential deep learning architectures and mechanisms that are widely used today (e.g. convolutional neural networks and mechanisms of attention) owe their origins to biological intelligence. Despite decades of research into computational models of the visual system, our understanding of its complexities remains far from complete. Existing neuroscientific models of the visual system are often founded upon empirical observations from relatively small datasets, and are therefore unlikely to capture the true complexity of the visual system. While these models have

successfully explained many properties of neural response to simple stimuli, their simplicity does not generalize to complex image stimuli [1].

Following their astounding success in computer vision, task-driven convolutional neural networks (CNNs) have recently been proposed as candidate models of the ventral stream in primate visual cortex [2–6], offering a path towards models that can explain hidden complexities of the visual system and generalize to complex visual stimuli. Through typical task-driven training alone, representations that resemble neural activity at multiple levels of the visual hierarchy have been observed in these models. With the emergence of such properties, CNNs are already being used to enhance our knowledge of processing in the ventral stream [7].

Despite these advancements, CNNs that achieve state-of-the-art brain alignment are still unable to explain many properties of the visual stream. Most traditional CNNs omit many well known architectural and processing hallmarks of the primate ventral stream that are likely key to the development of artificial neural nets that help us decipher the neural code. The development of these mechanisms remains an open challenge. A comprehensive understanding of the visual stream could in turn contribute to significant leaps in AI – a long-established goal of NeuroAI research.

In this work, we take a systematic approach to analyzing the hallmarks of the primate ventral stream that improve model-brain similarity of CNNs. We formulate architectural components that simulate these processing hallmarks within CNNs and analyze the population and neuron level response properties of these networks, as compared to empirical data recorded in primates. In specific:

- We introduce architectural components based on neuroscience foundations that simulate cortical magnification, center-surround antagonism, and tuned divisive normalization.
- We systematically analyze how these architectural components lead to latent representations that better explain primate V1 activity with multifaceted Brain-Score analyses. We identify center-surround antagonism, local filtering, tuned normalization, and cortical magnification as complementary ways to improve V1 alignment.
- We enrich the classic ResNet50 architecture with these architectural components and show that the resulting network achieves top V1-overall score on the integrative Brain-Score benchmark. An ablation study reveals the importance of each component and the benefits of combining multiple of these components into a single neuro-constrained CNN.

2 Background and Related Work

Model-Brain Alignment One central challenge in the field of NeuroAI is the development of computational models that can effectively explain the neural code. To achieve this goal, artificial neural networks must be capable of accurately predicting the behavior of individual neurons and neural populations in the brain. The primary visual cortex (V1) is one of the most well studied areas of the visual stream, with modeling efforts dating back to at least 1962 [8]—yet many deep learning models still fall short in explaining its neural activity.

The Brain-Score integrative benchmark [9] has recently emerged as a valuable tool for assessing the capabilities of deep learning models to explain neural activity in the visual system. This suite of benchmarks integrates neural recording and behavioral data from a collection of previous studies and provides standardized metrics for evaluating model explainability of visual areas V1, V2, V4, and IT, as well as additional behavioral and engineering benchmarks.

Although CNNs draw high-level inspiration from neuroscience, current architectures (e.g., ResNet and EfficientNet) bear little resemblance to neural circuits in the visual system. While such differences may not necessarily hinder object recognition performance, these networks still fall short in mimicking many properties of highly capable visual systems. Although there may be many paths towards next generation AI, foundational studies that have successfully merged foundations of neuroscience and AI have shown promising improvements to traditional artificial neural networks (ANNs) [10–12].

Center-Surround Antagonism As early as in the retina, lateral inhibitory connections establish a center-surround antagonism in the receptive field (RF) of many retinal cell types, which is preserved by neurons in the lateral geniculate nucleus and the visual cortex. In the primate visual stream, this center-surround antagonism is thought to facilitate edge detection, figure-ground segregation, depth




Figure 1: Design patterns of neuro-constrained architectural components. A) Difference of Gaussian implements a center-surround receptive field. B) Local receptive fields of two neurons without weight sharing. C) Tuned divisive normalization inhibits each feature map by a Gaussian-weighted average of competing features. D) Log-polar transform simulating cortical magnification

perception, and cue-invariant object perception [13–16], and is therefore a fundamental property of visual processing.

Center-surround RFs are a common component of classical neuroscience models [17–19], where they are typically implemented using a Difference of Gaussian (DoG) that produces an excitatory peak at the RF center with an inhibitory surround (Fig. 1A). Although deep CNNs have the capacity to learn center-surround antagonism, supplementing traditional convolutional kernels with fixed-weight DoG kernels has been demonstrated to improve object recognition in the context of varied lighting, occlusion, and noise [20, 21].

Local Receptive Fields The composition of convolutional operations in CNNs enables hierarchical processing and translation equivariance, both of which are fundamental to core object recognition in the primate ventral visual stream. However, the underlying mechanism through which this is achieved is biologically implausible, as kernel weights are shared among downstream neurons. Though locally connected neural network layers can theoretically learn the same operation, traditional convolutions are typically favored in practice for their computational efficiency and performance benefits. However, local connectivity is a ubiquitous pattern in the ventral stream (Fig. 1B), and visual processing phenomena (e.g., orientation preference maps [22]) have been attributed to this circuitry pattern. In artificial neural systems, Lee *et al.* [23] observed the emergence of topographic hallmarks in the inferior temporal cortex when encouraging local connectivity in CNNs. Pogodin *et al.* [24] considered the biological implausibility of CNNs and demonstrated a neuro-inspired approach to reducing the performance gap between traditional CNNs and locally-connected networks, meanwhile achieving better alignment with neural activity in primates.

Divisive Normalization Divisive normalization is wide-spread across neural systems and species [25]. In early visual cortex, it is theorized to give rise to well-documented physiological phenomena, such as response saturation, sublinear summation of stimulus responses, and cross-orientation suppression [26].

In 2021, Burg and colleagues [27] introduced an image-computable divisive normalization model in which each artificial neuron was normalized by weighted responses of neurons with the same receptive field. In comparison to a simple 3-layer CNN trained to predict the same stimulus responses, their analyses revealed that cross-orientation suppression was more prevalent in the divisive normalization model than in the CNN, suggesting that divisive normalization may not be inherently learned by task-driven CNNs. In a separate study, Ciricione *et al.* [28] showed that simulating divisive normalization within a CNN can improve object recognition robustness to image corruptions and enhance alignment with certain tuning properties of primate V1.

Tuned Normalization/Cross-Channel Inhibition While it is not entirely clear whether divisive normalization should be performed across space and/or across channels (implementations vary widely), Rust *et al.* [29] demonstrated that many response properties of motion-selective cells in the middle temporal area, such as motion-opponent suppression and response normalization, emerge from a mechanism they termed “tuned normalization”. In this scheme, a given neuron is normalized by a pool of neurons that share the same receptive field but occupy a different region in feature space. We adopt this idea in the present work (Fig. 1C), hypothesizing that enforcing feature-specific

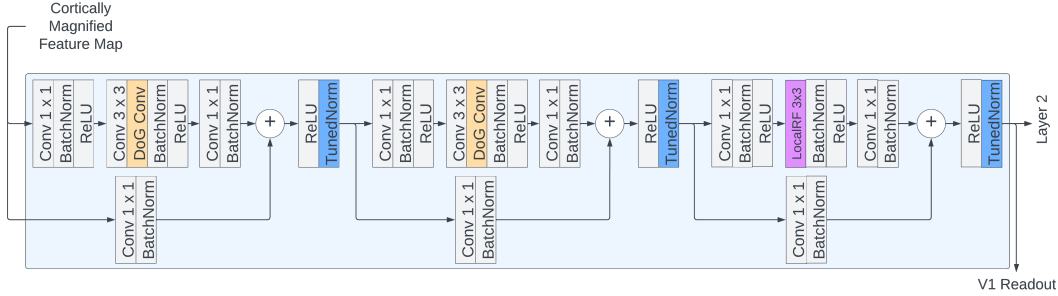


Figure 2: ResNet50 layer 1, supplemented with neuro-constrained architectural components. Throughout the modified layer 1, primary visual cortex (V1) activity is modeled with cortical magnification, center-surround convolutions, tuned normalization, and local receptive field layers. Layer 1 output units are treated as artificial V1 neurons.

weights in the pooling signal might enable a deep net to learn “opponent suppression” signals, much like cross-orientation signals found in biological V1 [30, 31].

Cortical Magnification In many sensory systems, a disproportionately large area of the cortex is dedicated to processing the most important information. This phenomenon, known as cortical magnification, reflects the degree to which the brain dedicates resources to processing sensory information accompanying a specific sense. In the primary visual cortex, a larger proportion of cortical area processes visual stimuli presented at the center of the visual field as compared to stimuli at greater spatial eccentricities [32]. The relationship between locations in the visual field and corresponding processing regions in the visual cortex has commonly been modeled with a log-polar mapping (Fig. 1D) or derivations thereof [33–36].

Layers of artificial neurons of traditional CNNs have uniform receptive field sizes and do not exhibit any sort of cortical magnification, failing to capture these distinctive properties of neuronal organization in the primary visual cortex. Recent works have demonstrated that introducing log polar-space sampling into CNNs can give rise to improved invariance and equivariance to spatial transformations [37, 38] and adversarial robustness [39].

3 Methods

3.1 Neuro-Constrained CNN Architecture

Given previous state-of-the-art V1 alignment scores achieved with ResNet50 [11, 40], we adopted this architecture as our baseline and test platform. However, the architectural components that we considered in the work are modular and can be integrated into general CNNs architectures. The remainder of this subsection details each architectural component and its integration into a neuro-constrained ResNet. In all experiments, we treated the output units from ResNet50 layer 1 as “artificial V1” neurons (refer to section 3.2 for layer selection criteria). Fig. 2 depicts ResNet50 layer 1 after enhancement with neuroscience-based architectural components.

Center-Surround Antagonism Center-surround ANN layers are composed of DoG kernels of shape $(c_i \times c_o \times k \times k)$, where c_i and c_o denote the number of input and output channels, respectively, and k reflects the height and width of each kernel. These DoG kernels (Fig. 1A) are convolved with the pre-activation output of a standard convolution. Each DoG kernel, DoG_i is of the form

$$DoG_i(x, y) = \frac{\alpha}{2\pi\sigma_{i,center}^2} \exp\left(-\frac{x^2 + y^2}{2\sigma_{i,center}^2}\right) - \frac{\alpha}{2\pi\sigma_{i,surround}^2} \exp\left(-\frac{x^2 + y^2}{2\sigma_{i,surround}^2}\right), \quad (1)$$

where $\sigma_{i,center}$ and $\sigma_{i,surround}$ were the Gaussian widths of the center and surround, respectively ($\sigma_{i,center} < \sigma_{i,surround}$), α was a scaling factor, and $(x, y) := (0, 0)$ at the kernel center. For $\alpha > 0$, the kernel will have an excitatory center and inhibitory surround while $\alpha < 0$, results in a kernel with inhibitory center and excitatory surround. Novel to this implementation, each DoG kernel has learnable parameters, better accommodating the diverse tuning properties of neurons within the

network. As in [20, 21], these DoG convolutions were only applied to a fraction of the input feature map. Specifically, we applied this center-surround convolution to one quarter of all 3×3 convolutions in layer 1 of our neuro-constrained ResNet50.

Local Receptive Fields In an effort to untangle the effects of local connectivity on brain alignment, we modified the artificial V1 layer by substituting the final 3×3 convolution of ResNet50 layer 1 with a 3×3 locally connected layer in isolation. This substitution assigns each downstream neuron its own filter while preserving its connection to upstream neurons (Fig. 1B), following the pattern in [24].

Divisive Normalization We consider the divisive normalization block proposed in [28] which performs normalization both spatially and across feature maps using learned normalization pools. Following our experimental design principle of selectively modifying the network in the vicinity of the artificial V1 neurons, we added this divisive normalization block after the non-linear activation of each residual block in ResNet50 layer 1.

Tuned Normalization We devised a novel implementation of tuned normalization inspired by models of opponent suppression [17, 29, 30]. In this scheme, a given neuron is normalized by a pool of neurons that share the same receptive field but occupy a different region in feature space (Fig. 1C), as in [27, 28]. Unlike the learned, weighted normalization proposed in [27], tuned inhibition was encouraged in our implementation by enforcing that each neuron was maximally suppressed by a neuron in a different region of feature space, and that no other neuron is maximally inhibited by activity in this feature space. Letting $x_{i,j}^c$ denote the activity of the neuron at spatial location (i, j) and channel $c \in [1, C]$ after application of a non-linear activation function. The post-divisive normalization state of this neuron, $x'_{i,j}{}^c$, is given by:

$$x'_{i,j}{}^c = \frac{x_{i,j}^c}{1 + \sum_k p_k x_{i,j}^{c_k}}, \quad (2)$$

where $p_{c,1}, \dots, p_{c,C}$ defines a Gaussian distribution with variance σ_c^2 centered at channel $(c + \frac{C}{2}) \bmod C$. By defining σ_c^2 as a trainable parameter, task-driven training would optimize whether each neuron should be normalized acutely or broadly across the feature space.

As this mechanism preserves the dimension of the input feature map, it can follow any non-linear activation function of the core network without further modification to the architecture. Similar to the divisive normalization block, tuned normalization was added after the non-linear activation of each residual block in ResNet50 layer 1 in our experiments.

Cortical Magnification Cortical magnification and non-uniform receptive field sampling was simulated in CNNs using a differentiable polar sampling module (Fig. 1D). In this module, the spatial dimension of an input feature map are divided into polar regions defined by discrete radial and angular divisions of polar space. In particular, we defined a discrete polar coordinate system partitioned in the first dimension by radial partitions r_0, r_1, \dots, r_m and along the second dimension by angular partitions $\theta_0, \theta_1, \dots, \theta_n$. Pixels of the input feature map that are located within the same polar region (i.e., are within the same radial bounds $[r_i, r_{i+1})$ and angular bounds $[\theta_j, \theta_{j+1})$) are pooled and mapped to coordinate (i, j) of the original pixel space (Fig. 1D) [41]. Pixels in the output feature map with no associated polar region were replaced with interpolated pixel values from the same radial bin. By defining the spacing between each concentric radial bin to be monotonically increasing (i.e., for all $i \in [1, m - 1]$, $(r_i - r_{i-1}) \leq (r_{i+1} - r_i)$), visual information at lower spatial eccentricities with respect to the center of the input feature map consumes a larger proportion of the transformed feature map than information at greater eccentricities.

A notable result of this transformation is that any standard 2D convolution, with a kernel of size $k \times k$, that is applied to the transformed features space is equivalent to performing a convolution in which the kernel covers a $k \times k$ contiguous region of polar space and strides along the angular and radial axes. Reflective padding was used after this transformation to enable a periodic stride along the angular axis at each radial position [37]. Furthermore, downstream artificial neurons which process information at greater spatial eccentricities obtain larger receptive fields. Treating the CNN as a model of the ventral visual stream, this polar transformation immediately preceded ResNet50 layer 1, where V1 representations were assumed to be learned.

3.2 Training and Evaluation

Training Procedure V1 alignment was evaluated for ImageNet-trained models [42]. Training and validation images were downsampled to a resolution of 64×64 . Each model of this evaluation was randomly initialized and trained for 100 epochs with an initial learning rate of 0.1 (reduced by a factor of 10 at epochs 60 and 80, where validation set performance was typically observed to plateau), and a batch size of 128.

We additionally benchmarked each neuro-constrained model on the Tiny-ImageNet-C dataset to study the effect of V1 alignment on object recognition robustness [43] (refer to appendix for evaluation details). Tiny-ImageNet-C was used as an alternate to ImageNet-C given that the models trained here expected 64×64 input images and downsampling the corrupted images of ImageNet-C would have biased our evaluations. ImageNet pre-trained models were fine-tuned on Tiny-ImageNet prior to this evaluation. As a given model will learn alternate representations when trained on different datasets (thereby resulting in V1 alignment differences), we methodologically froze all parameters of each ImageNet trained model, with the exception of the classification head, prior to 40 epochs of fine tuning with a learning rate of 0.01 and a batch size of 128.

Validation loss and accuracy were monitored during both training procedures and the model state that enabled the greatest validation accuracy during training was restored for evaluations that followed. Training data augmentations were limited to horizontal flipping (ImageNet and Tiny-ImageNet) and random cropping (ImageNet).

Training was performed using single NVIDIA 3090 and A100 GPUs. Each model took approximately 12 hours to train on ImageNet and less than 30 minutes to fine-tune on Tiny-ImageNet.

Evaluating V1 Alignment We evaluated the similarity between neuro-constrained models of V1 and the primate primary visual cortex using the Brain-Score V1 benchmark [9]. The V1 benchmark score is an average of two sub-metrics: ‘V1 FreemanZiemba2013’ and ‘V1 Marques2020’, which we refer to as V1 Predictivity and V1 Property scores in what follows. For each metric, the activity of artificial neurons in a given neural network layer is computed using in-silico neurophysiology experiments. The V1 Predictivity score reflects the degree to which the model can explain the variance in stimulus-driven responses of V1 neurons, as determined by partial least squares regression mapping. The V1 Property score measures how closely the distribution of 22 different neural properties, from 7 neural tuning categories (orientation, spatial frequency, response selectivity, receptive field size, surround modulation, texture modulation, and response magnitude), matches between the model’s artificial neural responses and empirical data from macaque V1. Together, these two scores provide a comprehensive view of stimulus response similarity between of artificial and primate V1 neurons.

Brain-Score evaluations assume a defined mapping between units of an ANN layer and a given brain region. In all analyses of V1 alignment that follow, we systematically fixed the output neurons of ResNet50 layer 1 as the artificial V1 neuron. Note that this is a stricter rule than most models submitted to the Brain-Score leaderboard, as researchers are able to choose which layer in the deep net should correspond to the V1 readout. In baseline analyses, among multiple evaluated layers, we observed highest V1 alignment between artificial units primate V1 activity from layer 1, establishing it as a strong baseline. Alternate layer V1 scores are presented in the appendix.

4 Results

4.1 Architectural Components in Isolation

Patterns of neural activity observed in the brain can be attributed to the interplay of multiple specialized processes. Through an isolated analysis, our initial investigations revealed the contribution of specialized mechanisms to explaining patterns of neural activity in V1. Tables 1 and 2 present the results of this analysis, including ImageNet validation accuracy, V1 Overall, V1 Predictivity, and V1 Property scores.

Among the four modules evaluated in this analysis, cortical magnification emerged as the most influential factor in enhancing V1 alignment. This mechanism substantially improved the ResNet’s ability to explain the variance in stimulus responses, and the artificial neurons exhibited tuning properties that were more closely aligned with those of biological neurons, particularly in terms of

	ImageNet Acc	V1 Overall	V1 Predictivity	V1 Property
ResNet50 (Baseline)	.613 ± .002	.550 ± .004	.295 ± .003	.805 ± .011
Center-surround antagonism	.610 ± .001	.545 ± .002	.304 ± .016	.786 ± .018
Local receptive fields	.609 ± .001	.550 ± .006	.300 ± .002	.799 ± .012
Divisive normalization	.606 ± .001	.543 ± .003	.271 ± .014	.815 ± .011
Tuned normalization	.608 ± .002	.547 ± .004	.274 ± .004	.820 ± .009
Cortical magnification	.548 ± .008	.587 ± .014	.370 ± .008	.805 ± .021

Table 1: ImageNet object recognition classification performance (64×64 images) and primary visual cortex (V1) alignment scores of ResNet50 augmented with each architectural component. Mean and standard deviations are reported across three runs (random initialization, training, and evaluating) of each architecture. Scores higher than baseline are presented in green and those lower are presented with in red (the more saturated the color is, the greater the difference from baseline).

	Orientation	Spatial frequency	Response selectivity	RF size	Surround modulation	Texture modulation	Response magnitude
ResNet50 (Baseline)	.893 ± .023	.826 ± .048	.684 ± .059	.832 ± .080	.820 ± .009	.786 ± .058	.790 ± .042
Center-surround	.876 ± .027	.831 ± .030	.632 ± .012	.853 ± .046	.773 ± .027	.757 ± .025	.783 ± .024
Local receptive fields	.904 ± .021	.817 ± .016	.648 ± .008	.852 ± .054	.847 ± .083	.743 ± .036	.780 ± .022
Divisive normalization	.908 ± .017	.840 ± .014	.689 ± .007	.858 ± .046	.860 ± .070	.746 ± .030	.846 ± .019
Tuned normalization	.907 ± .035	.841 ± .013	.689 ± .023	.865 ± .031	.852 ± .020	.742 ± .029	.844 ± .015
Cortical magnification	.907 ± .037	.848 ± .039	.708 ± .011	.803 ± .044	.664 ± .063	.789 ± .058	.917 ± .071

Table 2: Model alignment across the seven primary visual cortex (V1) tuning properties that constitute the V1 Property score. Mean and standard deviation of scores observed across three trials of model training and evaluation are reported.

orientation tuning, spatial frequency tuning, response selectivity, and most of all, stimulus response magnitude. However, the artificial neuronal responses of the cortical magnification network showed lower resemblance to those observed in primate V1 with regard to surround modulation, as compared to the baseline network.

Simulating neural normalization within the ResNet resulted in artificial neurons that displayed improved alignment with primate V1 in terms of response properties. Noteworthy enhancements were observed in the spatial frequency, receptive field size, surround modulation, and response magnitude properties of neurons within the modified network, leading to improvements in the V1 Property score. These results applied to both tuned and untuned forms of normalization.

In contrast, the introduction of center-surround convolutions yielded minimal improvements in neural predictivity and slight reductions in overall neuron property similarity. Surprisingly, the surround modulation properties of the artificial neurons decreased compared to the baseline model, contrary to our expectations.

Finally, replacing the final 3×3 convolution preceding the artificial V1 readout with a locally connected layer resulted in modest changes in V1 alignment. This was one of the two mechanisms that led to improvements in the surround modulation response property score (tuned normalization being the other).

These findings collectively provide valuable insights into the individual contributions of each specialized mechanism. Although mechanisms simulating center-surround antagonism (i.e., DoG convolution) and local connectivity provide little benefit to overall predictivity and property scores in isolation, we observed that they reduce the property dissimilarity gap among tuning properties that are nonetheless important and complement alignment scores where divisive normalization and cortical magnification do not.

4.2 Complementary Components Explain V1

Constraining a general-purpose deep learning model with a single architectural component is likely insufficient to explain primate V1 activity given our knowledge that a composition of known circuits play pivotal roles in visual processing. Taking inspiration from this design principle, we supplemented

Center-Surround	Local RF	Tuned Normalization	Cortical Magnification	Adversarial Training	ImageNet Acc	V1 Overall	V1 Predictivity	V1 Property
X	X	X	X		:551	:605	:357	:852
	X	X	X		:543	:605	:353	:857
X		X	X		:541	:599	:340	:858
X	X		X		:552	:592	:364	:820
X	X	X			:603	:555	:276	:834
		X	X		:541	:598	:351	:845
	X		X		:555	:593	:384	:803
	X	X			:606	:561	:287	:835
X	X	X	X	X	:448	:629	:430	:829
	X	X	X	X	:448	:625	:430	:819

Table 3: Backward component elimination results. Checkmarks denote whether or not the architectural component was included in the model. Adversarial training was performed on the two models that tied for the top V1 Overall Score.

a ResNet50 with each implemented architectural component and discern the necessary components to achieve optimal V1 alignment in an ablation study. We omit the architectural component implementing divisive normalization, however, as it cannot be integrated simultaneously with tuned normalization, which was observed to yield slightly higher V1 Predictivity and Property scores in isolated component evaluation. In this study, we employed a greedy approach reminiscent of backward elimination feature selection. In each round of this iterative approach, we selectively removed the architectural component that reduced overall V1 alignment the most until only one feature remained. This analysis allowed us to identify the subset of components that collectively yielded the most significant improvements in V1 alignment, and unraveled the intricate relationship between these specialized features and their combined explanation of V1.

The results of the ablation study are presented in Table 3. With the exception of center-surround antagonism, removing any neural mechanisms from the modified residual network reduced overall V1 alignment, suggesting that (1) each architectural component contributed to V1 alignment (the utility of center-surround antagonism is detailed in section 4.4) and (2) nontrivial interactions between these mechanisms explain V1 more than what is possible with any single mechanism. Seven of the eight models evaluated in this ablation study substantially outperformed all existing models on the Brain-Score platform in modeling V1 tuning property distributions. Furthermore, four models were observed to achieve state-of-the-art V1 Overall scores, explaining both V1 stimulus response activity and neural response properties with high fidelity.

Whether or not feed-forward, ImageNet-trained ANNs can fully approximate activity in primate V1 has stood as an open question. Previous studies have argued that no current model is capable of explaining all behavioral properties using neurons from a single readout layer. The top performing models of the current evaluation stand out as the first examples of CNNs with neural representations that accurately approximate all evaluated V1 tuning properties (see appendix), offering positive evidence for the efficacy of explaining primate V1 with neuro-inspired deep learning architectures.

4.3 Object Recognition Robustness to Corrupted Images

In contrast with the human visual system, typical CNNs generalize poorly to out-of-distribution data. Small perturbations to an image can cause a model to output drastically different predictions than it would on the in-tact image. Recent studies have demonstrated a positive correlation between model-brain similarity and robustness to image corruptions [12, 28, 45]. After freezing the brain-aligned representations of the models presented in this work and re-tuning each model's classification head on Tiny-ImageNet (see section 3.2), we evaluated the object recognition accuracy of each model from section 4.1 and the top two overall models from section 4.2 on the Tiny-ImageNet-C dataset. The results of these evaluations for each category of corruption and corruption strength are provided in the appendix.

Among the evaluated components, only tuned normalization was observed to yield improved corrupt image classification accuracy over the entire test set, albeit slight, beating the baseline accuracy (0.278) by 0.005 (i.e., an absolute improvement of 5%). More substantial improvements were observed on

‘brightness’, ‘defocus_blur’, ‘elastic_transform’, and ‘pixelate’ corruptions (improvements over the baseline of .00986, .00989, .0105, and .0133, respectively).

4.4 Adversarially Training Neuro-Constrained ResNets

Adversarial training has previously been shown to enhance the brain-similarity of artificial neural representations without any modification to the underlying network [11, 46]. Curious as to whether adversarial training would further align the neuro-constrained ResNet50s with V1 activity, we selectively trained the two networks most aligned with V1 (one model with all architectural components and the other with all components except center-surround convolution) from section 4.2 using “Free” adversarial training [47].

Despite the drop in object recognition accuracy, the artificial neural representations that emerged in each network were drastically better predictors of stimulus response variance representations. Tuning property alignment dropped in the process, but remained above previous state-of-the-art regardless. Interestingly, we found that the main difference in V1 scores between these two models can be traced to surround modulation tuning alignment. Center-surround convolution indeed contributed to improved surround modulation tuning learned while training with on corrupted images, contrasting its apparent lack of contribution to the overall network suggested in the ablation study.

In sum, both networks achieved Rank-1 V1 Overall, Predictivity, and Property scores by large margins, setting a new standard in this breed of brain-aligned CNNs. At the time of writing, the previous Rank-1 V1 Overall, Predictivity, and Property scores were .594, .409, and .816, respectively, and all achieved by separate models.

5 Discussion

Throughout this work we presented a systematic evaluation of four architectural components derived from neuroscience principles and their influence on model-V1 similarity. Specifically, we developed novel ANN layers that simulate principle processing mechanisms of the primate visual stream including center-surround antagonism, local receptive fields, tuned normalization, and cortical magnification. Through an ablation study and isolated component analyses, we found that each component contributed to the production of latent ANN representations that better resemble those of primate V1, as compared to a traditional baseline CNN. When these four components were assembled together within a neuro-constrained ResNet50, V1 tuning properties were explained better than any previous deep learning model that we are aware of. Furthermore, this neuro-constrained model exhibited state-of-the-art explanation of V1 neural activity and is the first of its kind to do so, by a large margin nonetheless, highlighting a promising direction in biologically constrained ANNs. Training this model with “free” adversarial training greatly improved its ability to predict primate neural response to image stimuli at a minor sacrifice to tuning property similarity, establishing an even larger gap between previous state of the art.

Among all architectural components examined in this work, cortical magnification was the most influential to improving V1 alignment. This mechanism on its own could not explain the neural activity as completely as the top models of this study, however. Our implementation of tuned normalization provided substantial improvement to V1 tuning property alignment, and was the only component that contributed to model robustness. The importance of center-surround antagonism seemed to be training data-dependent. In our ablation study, for which all models were trained on ImageNet, center-surround convolutional layers did not contribute to overall V1 scores. This did not surprise us, as deep CNNs have the capacity to learn similar representations without these specialized layers. When training on adversarially perturbed data, however, the center-surround antagonism provided by this layer appeared to improve surround modulation tuning properties of artificial V1 neurons. While previous attempts at improving model-brain similarity have been highly dataset dependent, our results highlight the importance of artificial network design.

A notable limitation to our work is the reduction in ImageNet classification performance that was observed upon the introduction of cortical magnification. While perfectly maintaining baseline model accuracy was not a motivation of this work, we can imagine situations in which object recognition performance needs to be preserved alongside these improvements in brain-model alignment. One scope of future work involves implementing saliency-driven polar transformations, so that the center

of the polar map assumed by the polar transform is focused on an object of interest as opposed to being fixed at the center of the image. We expect that such a mechanism would help to mitigate these reductions in ImageNet accuracy.

We additionally plan to extend this work to model architectures other than ResNet to validate the widespread application of each of these components.

This work highlights an important advancement in the field of NeuroAI, as we systematically establish a set of neuro-constrained architectural components that contribute to state-of-the-art V1 alignment. We argue that our architecture-driven approach can be further generalized to additional areas of the brain as well. The neuroscience insights that could be gleaned from increasingly accurate in-silico models of the brain have the potential to transform the fields of both neuroscience and AI.

References

- [1] Matteo Carandini, Jonathan B. Demb, Valerio Mante, David J. Tolhurst, Yang Dan, Bruno A. Olshausen, Jack L. Gallant, and Nicole C. Rust. Do We Know What the Early Visual System Does? *Journal of Neuroscience*, 25(46):10577–10597, November 2005.
- [2] Daniel L Yamins, Ha Hong, Charles Cadieu, and James J DiCarlo. Hierarchical Modular Optimization of Convolutional Networks Achieves Representations Similar to Macaque IT and Human Ventral Stream. In *Advances in Neural Information Processing Systems*, volume 26. Curran Associates, Inc., 2013.
- [3] Daniel L. K. Yamins, Ha Hong, Charles F. Cadieu, Ethan A. Solomon, Darren Seibert, and James J. DiCarlo. Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proceedings of the National Academy of Sciences*, 111(23):8619–8624, June 2014.
- [4] Daniel L. K. Yamins and James J. DiCarlo. Using goal-driven deep learning models to understand sensory cortex. *Nature Neuroscience*, 19(3):356–365, March 2016. Number: 3 Publisher: Nature Publishing Group.
- [5] Najib J. Majaj, Ha Hong, Ethan A. Solomon, and James J. DiCarlo. Simple Learned Weighted Sums of Inferior Temporal Neuronal Firing Rates Accurately Predict Human Core Object Recognition Performance. *Journal of Neuroscience*, 35(39):13402–13418, September 2015. Publisher: Society for Neuroscience Section: Articles.
- [6] Santiago A. Cadena, George H. Denfield, Edgar Y. Walker, Leon A. Gatys, Andreas S. Tolias, Matthias Bethge, and Alexander S. Ecker. Deep convolutional models improve predictions of macaque V1 responses to natural images. *PLOS Computational Biology*, 15(4):e1006897, April 2019. Publisher: Public Library of Science.
- [7] Pouya Bashivan, Kohitij Kar, and James J. DiCarlo. Neural population control via deep image synthesis. *Science*, 364(6439):eaav9436, May 2019. Publisher: American Association for the Advancement of Science.
- [8] David H Hubel and Torsten N Wiesel. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of physiology*, 160(1):106, 1962. Publisher: Wiley-Blackwell.
- [9] Martin Schrimpf, Jonas Kubilius, Ha Hong, Najib J. Majaj, Rishi Rajalingham, Elias B. Issa, Kohitij Kar, Pouya Bashivan, Jonathan Prescott-Roy, Franziska Geiger, Kailyn Schmidt, Daniel L. K. Yamins, and James J. DiCarlo. Brain-Score: Which Artificial Neural Network for Object Recognition is most Brain-Like?, January 2020. Pages: 407007 Section: New Results.
- [10] Zhe Li, Wieland Brendel, Edgar Walker, Erick Cobos, Taliah Muhammad, Jacob Reimer, Matthias Bethge, Fabian Sinz, Zachary Pitkow, and Andreas Tolia. Learning from brains how to regularize machines. In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [11] Joel Dapello, Tiago Marques, Martin Schrimpf, Franziska Geiger, David Cox, and James J DiCarlo. Simulating a Primary Visual Cortex at the Front of CNNs Improves Robustness to Image Perturbations. In *Advances in Neural Information Processing Systems*, volume 33, pages 13073–13087. Curran Associates, Inc., 2020.
- [12] Manish V. Reddy, Andrzej Banburski, Nishka Pant, and Tomaso Poggio. Biologically Inspired Mechanisms for Adversarial Robustness, June 2020. arXiv:2006.16427 [cs, stat].
- [13] J. Allman, F. Miezin, and E. McGuinness. Stimulus specific responses from beyond the classical receptive field: neurophysiological mechanisms for local-global comparisons in visual neurons. *Annual Review of Neuroscience*, 8:407–430, 1985.
- [14] J. J. Knierim and D. C. van Essen. Neuronal responses to static texture patterns in area V1 of the alert macaque monkey. *Journal of Neurophysiology*, 67(4):961–980, April 1992. Publisher: American Physiological Society.
- [15] Gary A. Walker, Izumi Ohzawa, and Ralph D. Freeman. Asymmetric Suppression Outside the Classical Receptive Field of the Visual Cortex. *Journal of Neuroscience*, 19(23):10536–10553, December 1999. Publisher: Society for Neuroscience Section: ARTICLE.
- [16] Zhi-Ming Shen, Wei-Feng Xu, and Chao-Yi Li. Cue-invariant detection of centre-surround discontinuity by V1 neurons in awake macaque monkey. *The Journal of Physiology*, 583(Pt 2):581–592, September 2007.

- [17] Gregory C DeAngelis, RALPH D Freeman, and IZUMI Ohzawa. Length and width tuning of neurons in the cat's primary visual cortex. *Journal of neurophysiology*, 71(1):347–374, 1994.
- [18] Michael P Sceniak, Dario L Ringach, Michael J Hawken, and Robert Shapley. Contrast's effect on spatial summation by macaque V1 neurons. *Nature neuroscience*, 2(8):733–739, 1999. Publisher: Nature Publishing Group.
- [19] Michael P Sceniak, Michael J Hawken, and Robert Shapley. Visual spatial characterization of macaque V1 neurons. *Journal of neurophysiology*, 85(5):1873–1887, 2001. Publisher: American Physiological Society Bethesda, MD.
- [20] Hosein Hasani, Mahdieh Soleymani, and Hamid Aghajan. Surround Modulation: A Bio-inspired Connectivity Structure for Convolutional Neural Networks. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d' Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.
- [21] Zahra Babaiee, Ramin Hasani, Mathias Lechner, Daniela Rus, and Radu Grosu. On-off center-surround receptive fields for accurate and robust image classification. In *International Conference on Machine Learning*, pages 478–489. PMLR, 2021.
- [22] Alexei A Koulakov and Dmitri B Chklovskii. Orientation preference patterns in mammalian visual cortex: a wire length minimization approach. *Neuron*, 29(2):519–527, 2001. Publisher: Elsevier.
- [23] Hyodong Lee, Eshed Margalit, Kamila M. Jozwik, Michael A. Cohen, Nancy Kanwisher, Daniel L. K. Yamins, and James J. DiCarlo. Topographic deep artificial neural networks reproduce the hallmarks of the primate inferior temporal cortex face processing network. preprint, Neuroscience, July 2020.
- [24] Roman Pogodin, Yash Mehta, Timothy Lillicrap, and Peter E Latham. Towards Biologically Plausible Convolutional Networks. In *Advances in Neural Information Processing Systems*, volume 34, pages 13924–13936. Curran Associates, Inc., 2021.
- [25] Matteo Carandini and David J. Heeger. Normalization as a canonical neural computation. *Nature Reviews Neuroscience*, 13(1):51–62, January 2012. Number: 1 Publisher: Nature Publishing Group.
- [26] David J. Heeger and Klavdia O. Zemlianova. A recurrent circuit implements normalization, simulating the dynamics of V1 activity. *Proceedings of the National Academy of Sciences*, 117(36):22494–22505, September 2020. Publisher: Proceedings of the National Academy of Sciences.
- [27] Max F Burg, Santiago A Cadena, George H Denfield, Edgar Y Walker, Andreas S Tolias, Matthias Bethge, and Alexander S Ecker. Learning divisive normalization in primary visual cortex. *PLOS Computational Biology*, 17(6):e1009028, 2021. Publisher: Public Library of Science San Francisco, CA USA.
- [28] Andrew Cirincione, Reginald Verrier, Artiom Bic, Stephanie Olaiya, James J DiCarlo, Lawrence Udeigwe, and Tiago Marques. Implementing Divisive Normalization in CNNs Improves Robustness to Common Image Corruptions.
- [29] Nicole C. Rust, Valerio Mante, Eero P. Simoncelli, and J. Anthony Movshon. How MT cells analyze the motion of visual patterns. *Nature Neuroscience*, 9(11):1421–1431, November 2006. Number: 11 Publisher: Nature Publishing Group.
- [30] M Concetta Morrone, DC Burr, and Lamberto Maffei. Functional implications of cross-orientation inhibition of cortical visual cells. I. Neurophysiological evidence. *Proceedings of the Royal Society of London. Series B. Biological Sciences*, 216(1204):335–354, 1982. Publisher: The Royal Society London.
- [31] GC DeAngelis, JG Robson, I Ohzawa, and RD Freeman. Organization of suppression in receptive fields of neurons in cat visual cortex. *Journal of Neurophysiology*, 68(1):144–163, 1992.
- [32] PM Daniel and D Whitteridge. The representation of the visual field on the cerebral cortex in monkeys. *The Journal of physiology*, 159(2):203, 1961. Publisher: Wiley-Blackwell.
- [33] Eric L Schwartz. Spatial mapping in the primate sensory projection: analytic structure and relevance to perception. *Biological cybernetics*, 25(4):181–194, 1977. Publisher: Springer.
- [34] Eric L Schwartz. Computational anatomy and functional architecture of striate cortex: a spatial mapping approach to perceptual coding. *Vision research*, 20(8):645–669, 1980. Publisher: Elsevier.
- [35] Eric L Schwartz. Computational studies of the spatial architecture of primate visual cortex: columns, maps, and protomaps. *Primary visual cortex in primates*, pages 359–411, 1994. Publisher: Springer.

- [36] Jonathan R Polimeni, Mukund Balasubramanian, and Eric L Schwartz. Multi-area visuotopic map complexes in macaque striate and extra-striate cortex. *Vision research*, 46(20):3336–3359, 2006. Publisher: Elsevier.
- [37] Carlos Esteves, Christine Allen-Blanchette, Xiaowei Zhou, and Kostas Daniilidis. Polar Transformer Networks. *International Conference on Learning Representations*, 2018.
- [38] João F. Henriques and Andrea Vedaldi. Warped Convolutions: Efficient Invariance to Spatial Transformations, 2021. [_eprint: 1609.04382](#).
- [39] Taro Kiritani and Koji Ono. Recurrent Attention Model with Log-Polar Mapping is Robust against Adversarial Attacks, 2020. [_eprint: 2002.05388](#).
- [40] 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.
- [41] M.R. Blackburn. A Simple Computational Model of Center-Surround Receptive Fields in the Retina. Technical report. Section: Technical Reports.
- [42] 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, June 2009. ISSN: 1063-6919.
- [43] Dan Hendrycks and Thomas Dietterich. Benchmarking Neural Network Robustness to Common Corruptions and Perturbations, March 2019. [arXiv:1903.12261 \[cs, stat\]](#).
- [44] Tiago Marques, Martin Schrimpf, and James J. DiCarlo. Multi-scale hierarchical neural network models that bridge from single neurons in the primate primary visual cortex to object recognition behavior, August 2021. Pages: 2021.03.01.433495 Section: New Results.
- [45] Shahd Safarani, Arne Nix, Konstantin Willeke, Santiago Cadena, Kelli Restivo, George Denfield, Andreas Tolias, and Fabian Sinz. Towards robust vision by multi-task learning on monkey visual cortex. In *Advances in Neural Information Processing Systems*, volume 34, pages 739–751. Curran Associates, Inc., 2021.
- [46] Alexander Riedel. Bag of Tricks for Training Brain-Like Deep Neural Networks. March 2022.
- [47] Ali Shafahi, Mahyar Najibi, Mohammad Amin Ghiasi, Zheng Xu, John Dickerson, Christoph Studer, Larry S Davis, Gavin Taylor, and Tom Goldstein. Adversarial training for free! In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019.