

# SCALING LAWS FOR TASK-OPTIMIZED MODELS OF THE PRIMATE VISUAL VENTRAL STREAM

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

When trained on large-scale object classification datasets, certain artificial neural network models begin to approximate core object recognition (COR) behaviors and neural response patterns in the primate visual ventral stream (VVS). While recent machine learning advances suggest that scaling model size, dataset size, and compute resources improve task performance, the impact of scaling on brain alignment remains unclear. In this study, we explore scaling laws for modeling the primate VVS by systematically evaluating over 600 models trained under controlled conditions on benchmarks spanning V1, V2, V4, IT and COR behaviors. We observe that while behavioral alignment continues to scale with larger models, neural alignment saturates. This observation remains true across model architectures and training datasets, even though models with stronger inductive bias and datasets with higher-quality images are more compute-efficient. Increased scaling is especially beneficial for higher-level visual areas, where small models trained on few samples exhibit only poor alignment. Finally, we develop a scaling recipe, indicating that a greater proportion of compute should be allocated to data samples over model size. Our results suggest that while scaling alone might suffice for alignment with human core object recognition behavior, it will not yield improved models of the brain’s visual ventral stream with current architectures and datasets, highlighting the need for novel strategies in building brain-like models. <sup>1</sup>

## 1 INTRODUCTION

The advent of neural networks has revolutionized our understanding and modeling of complex neural processes. A particularly active area of study is the ventral visual stream in primates, a key pathway in the brain responsible for processing visual information Goodale & Milner (1992); Grill-Spector et al. (2001); Malach et al. (2002); Kriegeskorte et al. (2008). Neural networks, when trained on extensive datasets, have emerged as the most accurate quantitative tools for simulating the response patterns of neurons within this stream Yamins et al. (2014); Schrimpf et al. (2018). These advanced models offer a precise computational account of how neural mechanisms in the brain give rise to visual perception.

Recent developments in machine learning have emphasized the significance of both the volume of training data and the complexity of model architectures Kaplan et al. (2020); Hoffmann et al. (2022); Zhai et al. (2022); Bahri et al. (2022); Antonello et al. (2023); Muennighoff et al. (2023); Aghajanyan et al. (2023); Isik et al. (2024). These findings raise the question: Can we build better models of the brain by scaling up model architectures and dataset sizes? Recent studies have found that in pre-trained models, the number of parameters and dataset samples respectively seem to improve predictions of fMRI and behavioral measurements (Antonello et al., 2023; Muttenthaler et al., 2023). With the numerous differences between pre-trained models however, the relative contributions of model parameters and dataset size to brain and behavioral alignment are not clear.

In this paper, we examine how scaling – of model parameters and training dataset size – impacts the alignment of artificial neural networks with the primate ventral visual stream. We systematically train models from a variety of architectural families on image classification datasets which allows us to independently control and observe the effects of model complexity and data volume. To capture

<sup>1</sup>We open-source all code, as well as checkpoints for our model zoo.

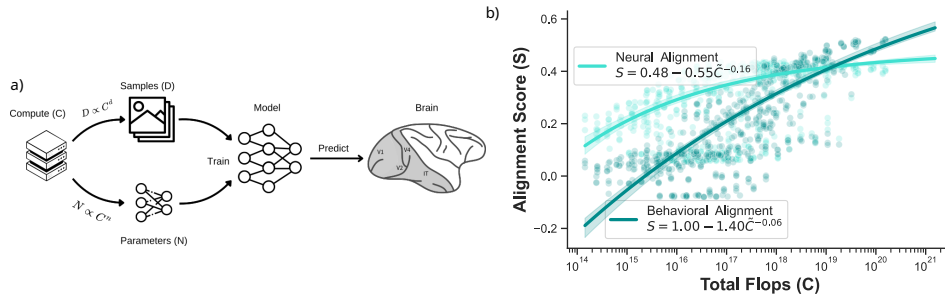


Figure 1: **a)** For a given compute budget ( $C$ ), we determine the scaling laws for maximal neural and behavioral alignment to the primate visual ventral stream. **b)** We find consistent scaling laws for brain and behavioral alignment across over 600 models. While we predict models to approach perfect behavioral alignment of 1 at large scales, the effect of scaling on brain alignment is already saturating at 0.48.

the observed trends, we introduce parametric power-law trends that describe the impact of scale on alignment with behavior and brain regions along the visual ventral stream. We summarize the contributions of this work as follows:

- While scale initially improves alignment, brain alignment saturates whereas behavioral alignment continues to improve.
- Increasing both parameter count and training dataset size improves alignment, with data providing more gains over model scaling.
- Architectures with stronger inductive bias (such as convolutions) and datasets with higher-quality images are more sample- and compute-efficient.
- Fitting parametric power-law curves, we find that model alignment with higher-level brain regions and especially behavior benefits the most from scaling.
- We publicly release our training code, evaluation pipeline, and over 600 checkpoints for models trained in a controlled manner to enable future research.

## 2 RELATED WORK

**Primate Visual Ventral Stream.** The ventral visual stream, a critical pathway in the primate brain, including humans, plays a key role in visual perception, extending from the occipital to the temporal lobes and serving as the “what pathway” for object recognition and form representation Goodale & Milner (1992); Grill-Spector et al. (2001); Malach et al. (2002); Kriegeskorte et al. (2008). Beginning in the primary visual cortex (V1), where basic visual information from retinal ganglion cells is processed, the ventral stream proceeds through areas such as V2, V3, V4, and the inferotemporal cortex (IT), each responsible for increasingly complex features of visual perception Kandel et al. (2000). Despite decades of research and a wealth of brain data, the precise neural mechanisms underlying visual perception are not well understood.

**Modeling the Primate Visual Ventral Stream.** Particular artificial neural networks (ANNs) are the most accurate models of brain responses in the visual ventral stream and associated core object recognition behaviors Schrimpf et al. (2018; 2020). Models optimized for ecologically viable tasks (Yamins & DiCarlo, 2016) in particular have demonstrated strong brain and behavioral alignment (Yamins et al., 2014; Khaligh-Razavi & Kriegeskorte, 2014; Cadena et al., 2019; Schrimpf et al., 2018; Nayebi et al., 2018; Kietzmann et al., 2019; Rajalingham et al., 2018; Zhuang et al., 2021; Geiger et al., 2022) – notably these models are trained purely on image classification datasets, without fitting to brain data.

**Scaling Laws.** Recent advancements in artificial intelligence are driven by scaling the model size and training data. Empirical evidence suggests a power-law relationship between model performance and both model parameters and dataset size, indicating that continued scaling will further

108 improve performance Kaplan et al. (2020); Cherti et al. (2023); Zhai et al. (2022); Hoffmann et al.  
 109 (2022); Dehghani et al. (2023); Henighan et al. (2020); Brown et al. (2020); Bahri et al. (2022);  
 110 Hestness et al. (2017). The power-law exponents enable the optimal allocation of compute between  
 111 model parameters and dataset samples, such that performance is maximized Kaplan et al. (2020);  
 112 Hoffmann et al. (2022).

113 While scaling laws for machine learning *performance* has been extensively studied, the scaling laws  
 114 for *brain alignment* remain unclear. Recent studies indeed suggest an involvement of both model  
 115 size and data volume in the functional alignment with brain data Azabou et al. (2023); Benchetrit  
 116 et al. (2023); Caro et al. (2024); Antonello et al. (2023). Conversely, Muttenthaler et al. Muttenthaler  
 117 et al. (2023) indicate that sample size is critical for behavioral alignment. We here unify these results,  
 118 in the realm of the primate visual ventral stream, into quantitative scaling laws for how model and  
 119 dataset sizes relate to alignment with the brain and behavior.

### 121 3 METHODS

122  
 123 **Neural & Behavioral Alignment.** To evaluate the alignment of our model with brain function, we  
 124 utilize a range of benchmarks from Brain-Score Schrimpf et al. (2018; 2020). These benchmarks  
 125 assess model performance by comparing model activations or behavior with primate neural data  
 126 using the same images. Specifically, the V1 and V2 benchmarks compare model outputs to primate  
 127 single-unit recordings from Freeman et al. (2013), using 315 texture images and data from 102  
 128 V1 and 103 V2 neurons. For the V4 and IT benchmarks, 2,560 images are used to match model  
 129 activations to primate Utah array recordings from Majaj et al. (2015), based on data from 88 V4  
 130 and 168 IT electrodes. A linear regression is trained on 90% of the images to correlate model and  
 131 neural data, with prediction accuracy for the remaining 10% evaluated using Pearson correlation,  
 132 repeated ten times for cross-validation. The behavioral benchmark assesses model predictions for  
 133 240 images against primate behavioral data from Rajalingham et al. (2018) using a logistic classifier  
 134 trained on 2,160 labeled images. Pearson correlation is used to measure the similarity in confusion  
 135 patterns between model predictions and primate responses. All benchmark scores are normalized to  
 136 their respective maximum possible values.

137 We define the model’s alignment score  $S$  (and an inverse *Misalignment Score*  $L = 1 - S$ ) as the  
 138 average across the V1, V2, V4, IT, and behavioral benchmark scores. Layers are committed to brain  
 139 regions based on models trained on a full dataset, and applied to all variants trained with subsampled  
 140 datasets. As we reused the same neural and behavioral data both to select the optimal model layer  
 141 for readout and to assess the model’s alignment, we validated the benchmark results on a private  
 142 split of each dataset on Brain-Score. We observed an almost perfect correlation between the results  
 143 on the private and public splits (Appendix C).

144 **Scaling Models and Data.** We trained an array of standard models from several architecture  
 145 families. Specifically, we used ResNet18, 34, 50, 101, 152 from He et al. (2016); EfficientNet-  
 146 B0, 1, 2 from Tan & Le (2019); Vision Transformer ViTT, S, B, L from Dosovitskiy et al. (2021);  
 147 ConvNeXtT, S, B, L from Liu et al. (2022b); CORnet-S from Kubilius et al. (2019); and AlexNet  
 148 from Krizhevsky et al. (2012). We also trained 33 modified versions of ResNet18: 22 models ob-  
 149 tained by scaling the network width from 1/16 to 4 times the original size, and 11 models derived  
 150 by adjusting the depth. Similarly, we trained four additional ConvNeXt and ViT models by scaling  
 151 the width of the ConvNeXt-T and ViT-S architectures.

152 For our experiments, we selected two image classification datasets: ImageNet Deng et al. (2009) and  
 153 EcoSet Mehrer et al. (2021). ImageNet, with millions of labeled images across 1,000 categories, has  
 154 long been a benchmark in computer vision, designed to challenge and evaluate automated visual  
 155 object recognition systems. On the other hand, EcoSet is a more recent dataset, designed to provide  
 156 an ecologically valid representation of human-relevant objects. It contains over 1.5 million images  
 157 spanning 565 basic-level categories, curated to better reflect the natural distribution of objects in the  
 158 real world, aligning with human perceptual and cognitive experiences.

159 To create subsets of ImageNet and EcoSet, we sampled  $d \in 1, 3, 10, 30, 100, 300$  images per cate-  
 160 gory. For  $d \in 1, 10, 100$ , we repeated the runs with three random seeds to ensure robustness. For  
 161 ConvNeXts (Liu et al., 2022b) and ViTs (Touvron et al., 2022), we used the training recipes de-  
 veloped by the original model authors. The remaining models were trained for 100 epochs using a

minibatch size of 512. We employed a stochastic gradient descent (SGD) optimizer with a cosine decaying learning rate schedule, starting with a peak learning rate of 0.1 and incorporating a linear warm-up phase spanning five epochs. We maintained the momentum at 0.9 and applied a weight decay of  $10^{-4}$ . Cross-entropy loss was used as the minimization objective. We utilized standard ImageNet data augmentations, specifically random resized cropping and horizontal flipping.

**Scaling Power-Law Curves.** Following previous work on scaling laws Zhai et al. (2022); Hoffmann et al. (2022); Besiroglu et al. (2024), we fit power law functions in the form

$$L = E + AX^{-\alpha} \quad (1)$$

on the data where  $L$  is the misalignment score, and  $X$  is an independent variable, such as the number of samples seen ( $D$ ), number of parameters ( $N$ ), and the total training floating point operations (FLOPs) ( $C$ ). Coefficients  $E$ ,  $A$ , and  $\alpha$  are found by minimizing

$$\min_{a,e,\alpha} \sum_{i \in [\#\text{Runs}]} \text{Huber}_\delta(\text{LSE}(a - \alpha \log X_i, e) - \log L_i) \quad (2)$$

where  $E = \exp(e)$ ,  $A = \exp(a)$  and LSE is the log-sum-exp operator. We solve Eq. 1 using BFGS minimizer with  $\delta = 1e - 3$ , and use a grid of initialiations as follows:  $e \in \{-1, -0.5, \dots, 1\}$ ,  $a \in \{0, 5, \dots, 25\}$ ,  $\alpha \in \{0, 0.5, \dots, 2\}$ .

To capture the slow initial increase in benchmark scores of modern architectures like ConvNeXt and ViT models in the low-data regime, we introduce an additional parameter  $\lambda$  to Eq. 2. This parameter allows the fitted curve to saturate at lower scales, better reflecting the observed performance of these models under limited data conditions:

$$L = E + A(X + 10^\lambda)^{-\alpha} \quad (3)$$

We minimize the modified equation as before, using  $\lambda \in 0, 0.5, 1.0, 1.5, 2.0$ . To fit the curve described by Eq. 3, we utilize all data points from the ConvNeXt and ViT models. For fitting the remaining curves, we select ConvNeXt and ViT runs that were trained on datasets with either 300 samples per class or the full dataset. This approach ensures that the fitted curves accurately represent the scaling behavior of these architectures across different data regimes.

Furthermore, we would like to describe the misalignment ( $L$ ) as a function of both the model and data size ( $N$ ,  $D$ ) and predict optimal allocations  $N^*$  and  $D^*$  by solving

$$(N^*, D^*) = \arg \min_{N, D} L(N, D) \quad \text{s.t.} \quad \text{FLOPs}(N, D) = C \quad (4)$$

In that regard, following Hoffmann et al. (2022); Besiroglu et al. (2024) we fit a parametric function of the form

$$\hat{L}(N, D) = E + \frac{A}{N^\alpha} + \frac{B}{D^\beta} \quad (5)$$

where the loss ( $\hat{L}$ ) is a function of parameter count ( $N$ ) and number of samples seen ( $D$ ). In Eq. 5, the first term represents the loss in an ideal data generation scenario (entropy), the second and the third terms reflect the under-performance of a model due to limitations in parameter and data size Hoffmann et al. (2022); Muennighoff et al. (2023). Following the example of Hoffmann et al. (2022), we learn variables  $\{E, A, \alpha, B, \beta\}$  that characterizes misalignment by solving

$$\arg \min_{e, a, \alpha, b, \beta} \sum_{i \in [\#\text{Runs}]} \text{Huber}_\delta(\text{LSE}(a - \alpha \log N_i, b - \beta \log D_i, e) - \log L_i) \quad (6)$$

with  $\delta = 10^{-3}$  and  $E = \exp(e)$ ,  $A = \exp(a)$   $B = \exp(b)$ . Initialiations of  $b$  and  $\beta$  follow  $a$  and  $\alpha$ , respectively.

Both Kaplan et al. (2020); Hoffmann et al. (2022) assume that compute follows the relationship  $C(N, D) \approx 6ND$  to predict the optimal allocation of compute ( $C$ ) to  $N$  and  $D$  using a set of equations with the learned variables mentioned above:

$$N^*(C) = G(C/6)^{a'}, \quad D^*(C) = G^{-1}(C/6)^{b'} \quad (7)$$

where  $a' = \frac{\beta}{\alpha + \beta}$ ,  $b' = \frac{\alpha}{\alpha + \beta}$ ,  $G = \left(\frac{\alpha A}{\beta B}\right)^{\frac{1}{\alpha + \beta}}$

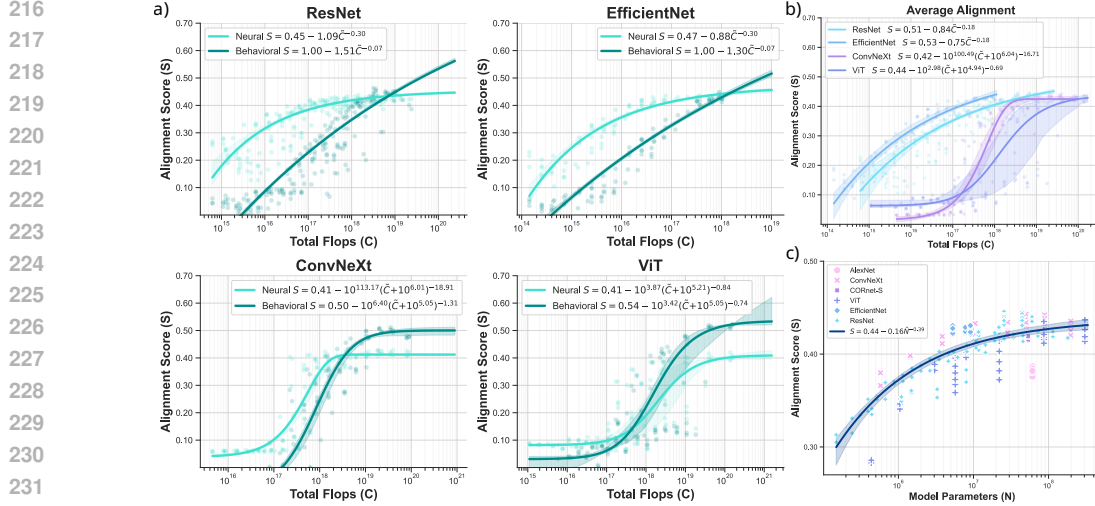


Figure 2: **Scaling Model Size.** **a)** Neural and behavioral alignments of four architecture families. Models with inductive biases (ResNet, EfficientNet) are more compute-efficient than less constrained models (ConvNeXt, ViT). **b)** Average alignment per model architecture. All models converge to similar alignments. **c)** Increasing parameters improves alignment (models trained on full datasets), but the effects saturate.

However, we observe that  $C(N, D) \approx 6ND$  does not hold with different architectures, and various CNN families have a slightly different relationship of  $C$ ,  $N$ , and  $D$ . As such, we assume a power-law relationship of the form

$$C(N, D) = m(ND)^n \quad (8)$$

where we fit  $m$  and  $n$  via linear regression of  $C$  and  $ND$  in log-log scale. Then, the updated equations governing the optimal allocation becomes

$$N^*(C) = G(C/m)^{a'/n}, \quad D^*(C) = G^{-1}(C/m)^{b'/n} \quad (9)$$

where  $a'$ ,  $b'$ , and  $G$  are calculated as before.

To evaluate the uncertainty of our model fits, we performed bootstrapping with 1,000 resamples. We compute 95% confidence intervals for each point along the fitted curves based on the variability observed across the bootstrapped estimates.

Finally, to avoid large constants during curve fitting, we rescale the variables  $C$ ,  $N$ , and  $D$  by setting  $\tilde{C} = C/10^{13}$ ,  $\tilde{N} = N/10^5$ , and  $\tilde{D} = D/10^4$ .

## 4 RESULTS

### 4.1 SCALING DRIVES BEHAVIORAL ALIGNMENT, BUT SATURATES FOR NEURAL ALIGNMENT

Our experiments show a clear and consistent improvement in behavioral alignment as both model size and training dataset size increase. Fig 1.b illustrates this trend across different architectures and scaling axes. The curve  $S = 1 - 1.4\tilde{C}^{-0.06}$  converges to perfect alignment score of 1 in the limit of  $C$ .

In contrast to behavioral alignment, neural alignment with specific brain regions demonstrated saturation as training compute scaled up in size. The curve represented by the formula  $S = 0.48 - 0.55\tilde{C}^{-0.16}$  represents a saturation at 0.48. The diminishing returns in neural alignment imply that merely scaling up models and data is insufficient to achieve better alignment with higher-level neural representations.

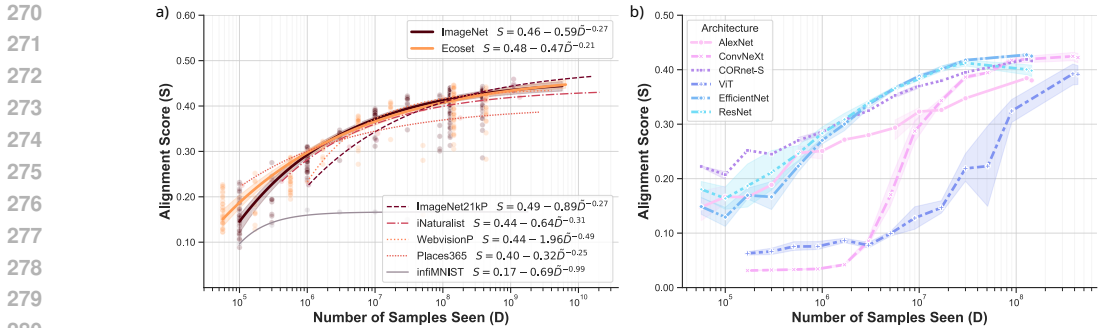


Figure 3: **Scaling Dataset Size.** **a)** Training on larger datasets enhances brain alignment. The alignment scaling curves derived from ImageNet and EcoSet closely estimate the alignment achieved when using ImageNet21k. In contrast, datasets with specialized image distributions—such as Places365—fall below the alignment scaling laws established by these generalist datasets. **b)** Model families with weaker inductive bias start at a lower alignment and require more data to improve.

#### 4.2 ARCHITECTURAL INDUCTIVE BIAS INFLUENCES ALIGNMENT AND SCALING BEHAVIOR

Experimental results indicate that modern architectures, such as ConvNeXt and Vision Transformers (ViTs), exhibit poorer neural alignment compared to models like ResNets and EfficientNets in low data regime. ResNets and EfficientNets, which have stronger inductive biases due to their fully convolutional structures, demonstrate high neural alignment even at initialization. In Fig. 2, alignment score of ResNets and EfficientNets increase steadily with additional compute in the form of training samples, however ConvNeXt and ViT requires more compute in order to start rising.

This difference in initial alignment also affects how the scaling laws evolve for each architecture. Models with weaker inductive biases require more extensive scaling—specifically in terms of training data—to achieve levels of neural alignment comparable to those with stronger inductive biases. Consequently, the scaling curves for ConvNeXt and ViT models develop differently, highlighting that architectural choices not only impact baseline alignment but also influence the efficiency of scaling strategies.

#### 4.3 MORE DATA IS BETTER THAN MORE PARAMETERS

Our analysis reveals that increasing the size of the training dataset has a more significant impact on improving brain alignment than simply enlarging the number of model parameters. While both strategies lead to performance enhancements, the benefits from data scaling exhibit less severe diminishing returns compared to model scaling. Specifically, models trained on larger datasets consistently demonstrate superior neural and behavioral alignment with the primate ventral visual stream, following a predictable power-law relationship.

In contrast, expanding the model size without proportionally increasing the training data results in steeper diminishing returns in alignment performance. Larger models rapidly reach a point where additional parameters do not translate into meaningful improvements. Fig. 2c estimates a saturation level of 0.44 by scaling model sizes with all samples of training data whereas Fig. 3a predicts maximum alignment of 0.48 and 0.50 for ImageNet and EcoSet respectively. This indicates that scaling training datasets overall improves brain alignment better than models scaling. Furthermore, Fig. 4b demonstrates that larger models of the same architecture family require much more samples to achieve the same level of alignment.

To quantitatively capture the joint interaction between data and model scaling, we fitted a parametric curve based on Eq.5, as shown in Fig.4a. This curve effectively models how compute ( $C$ ), dataset size ( $D$ ), and model size ( $N$ ) collectively influence brain alignment. Utilizing the parametric relationships described in Eq. 9, we estimate that additional compute should be allocated following the scaling laws  $D \approx C^{0.7}$  and  $N \approx C^{0.3}$ . These exponents indicate that, for optimal brain alignment, computational resources should be predominantly invested in increasing the dataset size rather than the model size.

324  
325  
326  
327  
328  
329  
330  
331  
332  
333  
334  
335  
336  
337  
338  
339  
340  
341  
342  
343  
344  
345  
346  
347  
348  
349  
350  
351  
352  
353  
354  
355  
356  
357  
358  
359  
360  
361  
362  
363  
364  
365  
366  
367  
368  
369  
370  
371  
372  
373  
374  
375  
376  
377

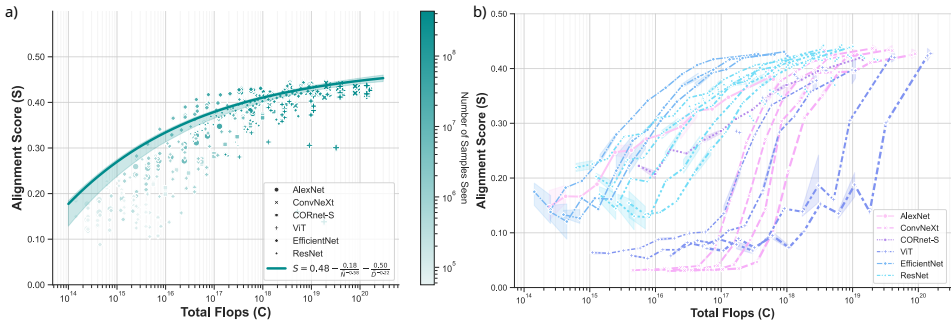


Figure 4: **Optimal Compute Allocation.** **a)** Alignment as a function of both model and training dataset sizes. Marker size is log-proportional to model size. Compute should be spent 0.3/0.7 on model/dataset size respectively. **b)** Models start out at different alignments but converge to the same saturating point.

#### 4.4 ORDERED EFFECT OF SCALE ON ALIGNMENT

Our study reveals a graded effect of scaling on alignment across the cortical hierarchy of the primate visual system. Specifically, we observe that the benefits of increased training compute—achieved through larger datasets and more complex models—vary systematically among different brain regions, reflecting their position in the visual processing pathway. Fig. 5.a illustrates the alignment as a function of training compute across various brain regions. We categorized the models into two groups based on their architectural inductive biases. Group 1 includes most models with strong inductive biases, such as ResNets and EfficientNets. These models start with higher neural alignment scores even at initialization due to their fully convolutional architectures. Group 2 consists of models with weaker inductive biases, specifically ConvNeXt and Vision Transformers (ViTs). These models exhibit lower neural alignment in the low-data regime and require more compute to achieve similar alignment levels.

To quantify the impact of scaling on each brain region, we define the alignment gain per region as  $A10^\alpha$  where  $A$  and  $\alpha$  are parameters of Eq. 2. Our findings indicate that higher regions in the cortical hierarchy show greater benefits from increased compute. Fig. 5b illustrates the alignment gain per region, highlighting how higher cortical areas benefit more from scaling efforts. This ordered effect suggests that regions higher up in the visual hierarchy, such as the Inferior Temporal (IT) cortex and behavioral outputs, gain more substantially from additional data and increased model complexity. In contrast, early visual areas like V1 and V2 exhibit smaller alignment gains with increased compute, indicating a potential saturation effect.

## 5 DISCUSSION

We establish scaling laws governing the effect of model and dataset scale on behavioral and brain alignment with the primate visual ventral stream. While scale is a necessary component for all brain-like models, model architectures with priors such as convolutions, and datasets with high-quality images are more sample efficient, leading to alignment with smaller compute requirements. Scale especially improves alignment with higher-level visual regions, but brain alignment saturates across all conditions tested here whereas behavioral alignment continuously improves with increased scale.

**Platonic representations.** Our results are consistent with the view that representations in deep neural network models are converging. This observation has been likened to Plato’s concept of an ideal reality (Huh et al., 2024). Our findings support this view in that neural network models at scale tend to converge toward similar brain and behavioral alignment, regardless of their initial architectural differences. However, we find that the representational convergence in current models is *distinct from representations in the brain’s* visual ventral stream. Intriguingly, behavioral choices on the other hand seem to converge to human-like core object recognition, suggesting that there is no unique solution to human-like behavior. Even models with differing inductive biases, such as ConvNeXt and ViT (which start with poorer neural alignment), eventually converge to the same

378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431

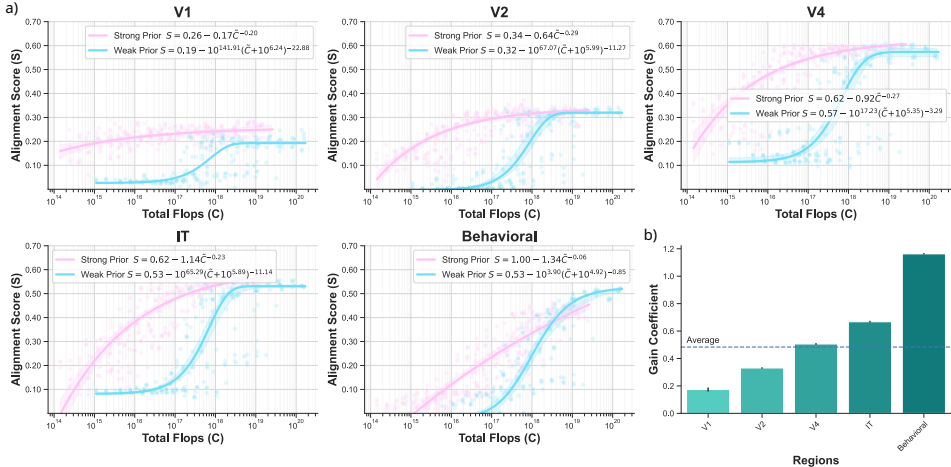


Figure 5: **Graded Effect of Scale across Cortical Hierarchy.** a) Alignment as a function of training compute across different brain regions. Group 1 contains most models except those with low inductive bias (Group 2; ConvNeXt, ViT). b) Alignment gain per region, defined as  $A10^\alpha$ . Regions higher in the cortical hierarchy show greater benefits from increased compute (Behavior > IT > V4 > V2 > V1).

behavioral, and, saturating brain alignment as models like ResNet and EfficientNet. Building models with representations that are consistent with the brain might thus require substantial changes to current architectures and training approaches.

**Dissociation of behavioral and neural alignment.** Our findings reveal a dissociation between behavioral and neural alignment as models are scaled with more parameters and larger datasets. While behavioral alignment continues to improve consistently with increased model parameters and training data – exhibiting a strong power-law relationship – neural alignment reaches a saturation point beyond which additional scaling yields minimal gains. This divergence suggests that behavioral alignment benefits more substantially from scaling efforts, whereas neural alignment may require alternative approaches beyond merely increasing model size and data volume to achieve further improvements.

This disparity is further highlighted by the correlation between task performance and alignment depicted in Figure 6. Behavioral alignment closely tracks validation accuracy, improving hand-in-hand as models become more accurate. Consistent with prior work (Schrimpf et al., 2018; Linsley et al., 2023), neural alignment eventually saturates, indicating that factors other than task performance influence neural alignment.

**Generalization Beyond Supervised Training.** We assessed whether alternative training paradigms can overcome the limitations observed in neural alignment under supervised learning. Figure 7a illustrates the scaling of alignment as a function of compute spent during self-supervised training of ResNet models using SimCLR Chen et al. (2020) on ImageNet. The results confirm the trends observed in supervised training: behavioral alignment continues to improve with increased compute, following a strong power-law relationship, while neural alignment approaches a saturation point. This consistency suggests that the saturation in neural alignment is not exclusive to supervised learning but may be inherent to the models or datasets employed.

The region-specific breakdown (as illustrated in Figure 13 of the appendix) further reinforces this observation. Even in a self-supervised learning context, higher-level visual areas like IT and behavioral outputs demonstrate more pronounced improvements with increased compute, while early visual areas like V1 and V2 show minimal gains. This suggests that the hierarchical nature of neural alignment is a fundamental characteristic that transcends specific training methodologies.






Figure 6: **Correlation between Task Performance and Alignment.** **a,b)** Correlation between validation accuracy (ImageNet & EcoSet) and brain (a) and behavioral (b) alignment. Behavioral alignment strongly correlates with task performance, whereas neural alignment shows a non-linear trend, reaching saturation. **c)** Pearson’s correlation coefficient per region, with all p-values less than  $10^{-40}$ .

Additionally, we explored the impact of adversarial fine-tuning on alignment performance. In Figure 7b, ResNet models trained on subsets of ImageNet were fine-tuned adversarially for 10 epochs using the Fast Gradient Sign Method (FGSM) (Goodfellow et al., 2015; Wong et al., 2020). Importantly, the scaling curves were estimated solely from the non-adversarial runs, yet the adversarially fine-tuned models exhibited improvements along these existing scaling curves. This indicates that adversarial training can enhance alignment without deviating from the established scaling behavior.

**Impact of Architectural Inductive Biases on Alignment Dynamics.** Our evaluation of alignment during training reveals that the alignment behavior varies significantly across different model architectures. Figure 7.c shows that while various models eventually converge to similar alignment levels with sufficient training, fully convolutional architectures—such as ResNets and EfficientNets—exhibit substantially higher alignment scores at the very beginning of training. This early advantage suggests that these architectures possess inherent features that align closely with neural data from the primate ventral visual stream even before learning from data occurs.

Further analysis in Figure 10 of the appendix confirms that this initial high alignment is due to the strong inductive biases present in fully convolutional networks. These biases enable the models to start with representations already well-suited for neural alignment. Figure 11 in the appendix reinforces this finding by demonstrating that models with strong inductive biases achieve higher initial alignment compared to architectures like ConvNeXt and ViT, which have weaker inductive biases.

**Influence of Learning Signals on Alignment Dynamics.** Our investigation reveals that the type of learning signal plays a crucial role in the dynamics of alignment during training. Figure 7d illustrates the alignment trajectories of ResNet50 and ViT-S models trained on ImageNet using supervised learning, SimCLR, and DINO (Caron et al., 2021) methods. Notably, the ViT-S model requires significantly more training steps to achieve the same level of alignment under supervised learning compared to when trained with self-supervised objectives like DINO and SimCLR. In contrast, the ResNet50 model, which possesses strong inductive biases due to its convolutional architecture, exhibits relatively consistent alignment dynamics across different learning signals. This robustness implies that models with strong inductive biases are less affected by the choice of training objective, whereas architectures like ViT-S benefit more substantially from rich, self-supervised feedback to achieve optimal alignment.

**Limitations and Future Directions.** Our study has several limitations. First, the extrapolation of our scaling functions is constrained by the specific range of model sizes and dataset volumes we examined. While we observed power-law relationships between scaling factors and brain alignment, these functions may not generalize beyond the scales tested.

Second, we evaluated a subset of models focusing primarily on standard and modern convolutional neural networks (e.g., ResNets and ConvNeXts), transformer-based architectures (e.g. ViTs) and recurrent networks (CORnet-S). While these architectures cover a range of inductive biases and complexities, they do not encompass the full spectrum of possible neural network designs, such

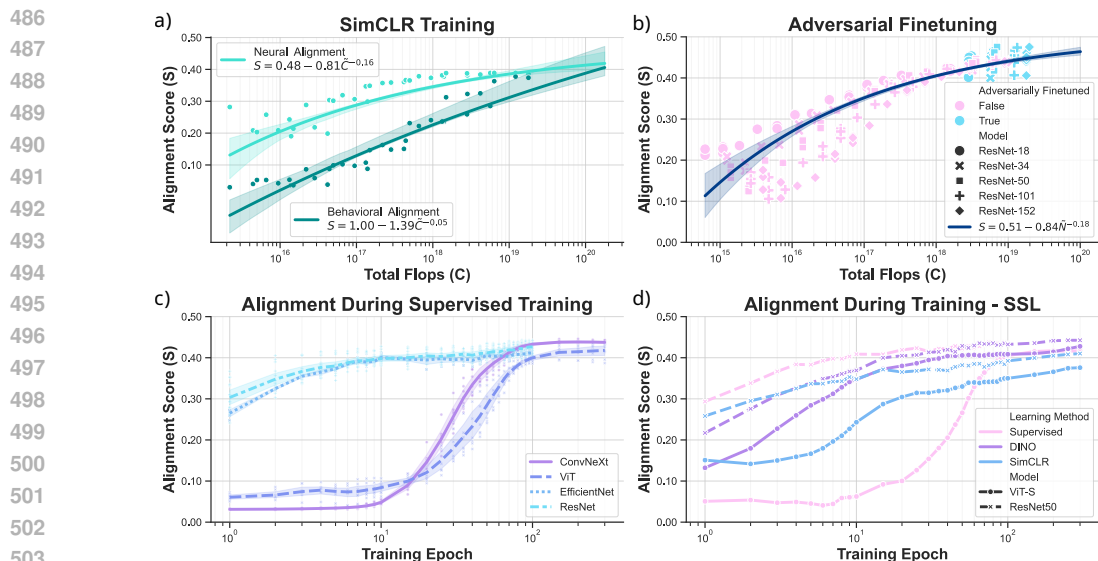


Figure 7: **Alignment Scaling of Alternative Training Strategies.** **a)** Unsupervised: ResNet models trained with SimCLR improve behavioral alignment with compute, while neural alignment saturates. **b)** Adversarial Robustness: Fine-tuning ResNets with adversarial training (FGSM) enhances alignment along the scaling curve. **c)** Architectural Prior: ResNets and EfficientNets exhibit higher initial alignment in early phases of training, due to strong inductive biases, unlike ConvNeXt and ViT. **d)** Alignment dynamics vary with training objectives, but converge to the same alignment saturation.

as more biologically plausible models. We see scaling laws as an opportunity to extrapolate the alignment of models at scale, even if their current training is compute-constrained.

Third, our experiments utilized a subset of training datasets primarily from ImageNet and EcoSet. Although these datasets are extensive and widely used, they may not capture all the nuances of visual stimuli relevant to the primate ventral visual stream. Therefore, models trained on other datasets might exhibit improved scaling properties.

Taken together, our results demonstrate that while scaling both model parameters and training data size enhances behavioral alignment with human visual perception, it leads to saturation in neural alignment with the primate ventral visual stream. Data scaling proves more effective than model scaling in improving alignment, emphasizing the critical role of extensive and diverse training datasets. We also find that architectural choices significantly influence alignment efficiency, with models possessing strong inductive biases—such as fully convolutional networks—achieving higher neural alignment even at initialization. Additionally, the impact of scaling varies across different brain regions, benefiting higher cortical areas more than early visual areas. These findings suggest that merely increasing scale is insufficient for modeling the intricate neural representations of the brain’s visual system. Future work should investigate new approaches, including alternative architectures and training strategies, to develop models that more accurately reflect the complexities of neural processing in the primate visual cortex.

To push neural alignment beyond current saturation levels, future research should explore adversarial training methods that encourage models to learn more robust, brain-like representations. Leveraging biologically inspired architectures such as VOneNets (Dapello et al., 2020) may lead to more compute-efficient models achieving higher neural alignment without extensive scaling. Additionally, investigating co-training with brain data—integrating neural recordings directly into the training process—could enhance both neural and behavioral alignment, paving the way for more accurate and efficient brain-like models.

## REFERENCES

- 540  
541  
542 Armen Aghajanyan, Lili Yu, Alexis Conneau, Wei-Ning Hsu, Karen Hambardzumyan, Susan  
543 Zhang, Stephen Roller, Naman Goyal, Omer Levy, and Luke Zettlemoyer. Scaling laws for  
544 generative mixed-modal language models. In Andreas Krause, Emma Brunskill, Kyunghyun  
545 Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Proceedings of the 40th*  
546 *International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learn-*  
547 *ing Research*, pp. 265–279. PMLR, 23–29 Jul 2023. URL [https://proceedings.mlr.](https://proceedings.mlr.press/v202/aghajanyan23a.html)  
548 [press/v202/aghajanyan23a.html](https://proceedings.mlr.press/v202/aghajanyan23a.html).
- 549 Richard J. Antonello, Aditya R. Vaidya, and Alexander G. Huth. Scaling laws for language encoding  
550 models in fmri. *Advances in Neural Information Processing Systems*, 36, 2023.
- 551 Mehdi Azabou, Vinam Arora, Venkataramana Ganesh, Ximeng Mao, Santosh B Nachimuthu,  
552 Michael Jacob Mendelson, Blake Aaron Richards, Matthew G Perich, Guillaume Lajoie, and  
553 Eva L Dyer. A unified, scalable framework for neural population decoding. In *Thirty-seventh*  
554 *Conference on Neural Information Processing Systems*, 2023. URL [https://openreview.](https://openreview.net/forum?id=sw2Y0sirtM)  
555 [net/forum?id=sw2Y0sirtM](https://openreview.net/forum?id=sw2Y0sirtM).
- 556 Yasaman Bahri, Ethan Dyer, Jared Kaplan, Jaehoon Lee, and Utkarsh Sharma. Explaining scaling  
557 laws of neural network generalization, 2022. URL [https://openreview.net/forum?](https://openreview.net/forum?id=FvfV64rovnY)  
558 [id=FvfV64rovnY](https://openreview.net/forum?id=FvfV64rovnY).
- 559 Yohann Benchetrit, Hubert Banville, and Jean-Remi King. Brain decoding: toward real-time recon-  
560 struction of visual perception. In *Submitted to The Twelfth International Conference on Learning*  
561 *Representations*, 2023. URL <https://openreview.net/forum?id=3y1K6bu08c>.
- 562 Tamay Besiroglu, Ege Erdil, Matthew Barnett, and Josh You. Chinchilla scaling: A replication  
563 attempt, 2024. URL <https://arxiv.org/abs/2404.10102>.
- 564 Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhari-  
565 wal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agar-  
566 wal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh,  
567 Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz  
568 Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec  
569 Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In  
570 H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neu-*  
571 *ral Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc.,  
572 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf)  
573 [file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf).
- 574 Alexander Buslaev, Vladimir I. Iglovikov, Eugene Khvedchenya, Alex Parinov, Mikhail Druzhinin,  
575 and Alexandr A. Kalinin. Alumentations: Fast and flexible image augmentations. *Information*,  
576 11(2), 2020. ISSN 2078-2489. doi: 10.3390/info11020125. URL [https://www.mdpi.com/](https://www.mdpi.com/2078-2489/11/2/125)  
577 [2078-2489/11/2/125](https://www.mdpi.com/2078-2489/11/2/125).
- 578 Santiago A. Cadena, George H. Denfield, Edgar Y. Walker, Leon A. Gatys, Andreas S. Tolia,  
579 Matthias Bethge, and Alexander S. Ecker. Deep convolutional models improve predictions of  
580 macaque v1 responses to natural images. *PLOS Computational Biology*, 15(4):e1006897, April  
581 2019. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1006897. URL [http://dx.doi.org/](http://dx.doi.org/10.1371/journal.pcbi.1006897)  
582 [10.1371/journal.pcbi.1006897](http://dx.doi.org/10.1371/journal.pcbi.1006897).
- 583 Josue Ortega Caro, Antonio Henrique de Oliveira Fonseca, Syed A Rizvi, Matteo Rosati, Christo-  
584 pher Averill, James L Cross, Prateek Mittal, Emanuele Zappala, Rahul Madhav Dhodapkar,  
585 Chadi Abdallah, and David van Dijk. BrainLM: A foundation model for brain activity record-  
586 ings. In *The Twelfth International Conference on Learning Representations*, 2024. URL  
587 <https://openreview.net/forum?id=RwI7ZEfR27>.
- 588 Mathilde Caron, Hugo Touvron, Ishan Misra, Hervé Jégou, Julien Mairal, Piotr Bojanowski, and  
589 Armand Joulin. Emerging properties in self-supervised vision transformers. In *Proceedings of*  
590 *the International Conference on Computer Vision (ICCV)*, 2021.

- 594 Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for  
595 contrastive learning of visual representations. In Hal Daumé III and Aarti Singh (eds.), *Pro-  
596 ceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceed-  
597 ings of Machine Learning Research*, pp. 1597–1607. PMLR, 13–18 Jul 2020. URL <https://proceedings.mlr.press/v119/chen20j.html>.
- 599 M. Cherti, R. Beaumont, R. Wightman, M. Wortsman, G. Ilharco, C. Gordon, C. Schuhmann,  
600 L. Schmidt, and J. Jitsev. Reproducible scaling laws for contrastive language-image learn-  
601 ing. In *2023 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*,  
602 pp. 2818–2829, Los Alamitos, CA, USA, jun 2023. IEEE Computer Society. doi: 10.  
603 1109/CVPR52729.2023.00276. URL [https://doi.ieeecomputersociety.org/10.  
604 1109/CVPR52729.2023.00276](https://doi.ieeecomputersociety.org/10.1109/CVPR52729.2023.00276).
- 605 Joel Dapello, Tiago Marques, Martin Schrimpf, Franziska Geiger, David Cox, and James J Di-  
606 Carlo. Simulating a primary visual cortex at the front of cnns improves robustness to im-  
607 age perturbations. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.),  
608 *Advances in Neural Information Processing Systems*, volume 33, pp. 13073–13087. Curran  
609 Associates, Inc., 2020. URL [https://proceedings.neurips.cc/paper\\_files/  
610 paper/2020/file/98b17f068d5d9b7668e19fb8ae470841-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/98b17f068d5d9b7668e19fb8ae470841-Paper.pdf).
- 611 Mostafa Dehghani, Josip Djolonga, Basil Mustafa, Piotr Padlewski, Jonathan Heek, Justin  
612 Gilmer, Andreas Peter Steiner, Mathilde Caron, Robert Geirhos, Ibrahim Alabdulmohsin,  
613 Rodolphe Jenatton, Lucas Beyer, Michael Tschannen, Anurag Arnab, Xiao Wang, Carlos  
614 Riquelme Ruiz, Matthias Minderer, Joan Puigcerver, Utku Evci, Manoj Kumar, Sjoerd Van  
615 Steenkiste, Gamaleldin Fathy Elsayed, Aravindh Mahendran, Fisher Yu, Avital Oliver, Fantine  
616 Huot, Jasmijn Bastings, Mark Collier, Alexey A. Gritsenko, Vighnesh Birodkar, Cristina Nader  
617 Vasconcelos, Yi Tay, Thomas Mensink, Alexander Kolesnikov, Filip Pavetic, Dustin Tran,  
618 Thomas Kipf, Mario Lucic, Xiaohua Zhai, Daniel Keysers, Jeremiah J. Harmsen, and Neil  
619 Houlsby. Scaling vision transformers to 22 billion parameters. In Andreas Krause, Emma  
620 Brunskill, Kyunghyun Cho, Barbara Engelhardt, Sivan Sabato, and Jonathan Scarlett (eds.), *Pro-  
621 ceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceed-  
622 ings of Machine Learning Research*, pp. 7480–7512. PMLR, 23–29 Jul 2023. URL <https://proceedings.mlr.press/v202/dehghani23a.html>.
- 624 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hi-  
625 erarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*,  
626 pp. 248–255. Ieee, 2009.
- 627 Mingyu Ding, Bin Xiao, Noel Codella, Ping Luo, Jingdong Wang, and Lu Yuan. Davit: Dual  
628 attention vision transformers. In *Computer Vision—ECCV 2022: 17th European Conference, Tel  
629 Aviv, Israel, October 23–27, 2022, Proceedings, Part XXIV*, pp. 74–92. Springer, 2022.
- 630 Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas  
631 Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszko-  
632 reit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recogni-  
633 tion at scale. In *International Conference on Learning Representations*, 2021. URL <https://openreview.net/forum?id=YicbFdNTTy>.
- 634 Jeremy Freeman, Corey M Ziemba, David J Heeger, Eero P Simoncelli, and J Anthony Movshon. A  
635 functional and perceptual signature of the second visual area in primates. *Nature Neuroscience*,  
636 16(7):974–981, 2013. doi: 10.1038/nn.3402.
- 637 Franziska Geiger, Martin Schrimpf, Tiago Marques, and James J. DiCarlo. Wiring up vision:  
638 Minimizing supervised synaptic updates needed to produce a primate ventral stream. In *In-  
639 ternational Conference on Learning Representations (ICLR) Spotlight*, 2022. URL <https://openreview.net/forum?id=g1SzIRLQXMM>.
- 640 Melvyn A Goodale and A David Milner. Separate visual pathways for perception and action. *Trends  
641 in Neurosciences*, 1992.
- 642 Ian Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial  
643 examples. In *International Conference on Learning Representations*, 2015. URL [http://  
644 arxiv.org/abs/1412.6572](http://arxiv.org/abs/1412.6572).

- 648 Benjamin Graham, Alaaeldin El-Nouby, Hugo Touvron, Pierre Stock, Armand Joulin, Herve Jegou,  
649 and Matthijs Douze. Levit: A vision transformer in convnet’s clothing for faster inference. In  
650 *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 12259–  
651 12269, October 2021.
- 652 Kalanit Grill-Spector, Zoe Kourtzi, and Nancy Kanwisher. The lateral occipital complex and its role  
653 in object recognition. *Vision Research*, 41(10-11):1409–1422, 2001.
- 654 Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Im-  
655 age Recognition. In *Proceedings of 2016 IEEE Conference on Computer Vision and Pattern  
656 Recognition, CVPR ’16*, pp. 770–778. IEEE, June 2016. doi: 10.1109/CVPR.2016.90. URL  
657 <http://ieeexplore.ieee.org/document/7780459>.
- 658 Tom Henighan, Jared Kaplan, Mor Katz, Mark Chen, Christopher Hesse, Jacob Jackson, Heewoo  
659 Jun, Tom B. Brown, Prafulla Dhariwal, Scott Gray, Chris Hallacy, Benjamin Mann, Alec Rad-  
660 ford, Aditya Ramesh, Nick Ryder, Daniel M. Ziegler, John Schulman, Dario Amodei, and Sam  
661 McCandlish. Scaling laws for autoregressive generative modeling, 2020.
- 662 Joel Hestness, Sharan Narang, Newsha Ardalani, Gregory Diamos, Heewoo Jun, Hassan Kianinejad,  
663 Md. Mostofa Ali Patwary, Yang Yang, and Yanqi Zhou. Deep learning scaling is predictable,  
664 empirically, 2017.
- 665 Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza  
666 Rutherford, Diego de las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, Tom Hen-  
667 nigan, Eric Noland, Katherine Millican, George van den Driessche, Bogdan Damoc, Aurelia  
668 Guy, Simon Osindero, Karen Simonyan, Erich Elsen, Oriol Vinyals, Jack William Rae, and  
669 Laurent Sifre. An empirical analysis of compute-optimal large language model training. In  
670 Alice H. Oh, Alekh Agarwal, Danielle Belgrave, and Kyunghyun Cho (eds.), *Advances in Neu-  
671 ral Information Processing Systems*, 2022. URL <https://openreview.net/forum?id=iBBcRUlOAPR>.
- 672 Minyoung Huh, Brian Cheung, Tongzhou Wang, and Phillip Isola. The platonic representation  
673 hypothesis. In *International Conference on Machine Learning*, 2024.
- 674 Berivan Isik, Natalia Ponomareva, Hussein Hazimeh, Dimitris Paparas, Sergei Vassilvitskii, and  
675 Sanmi Koyejo. Scaling laws for downstream task performance of large language models, 2024.  
676 URL <https://arxiv.org/abs/2402.04177>.
- 677 Eric R Kandel, James H Schwartz, and Thomas M Jessell. *Principles of Neural Science*. McGraw-  
678 Hill, 2000.
- 679 Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child,  
680 Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. Scaling laws for neural language  
681 models, 2020.
- 682 Seyed-Mahdi Khaligh-Razavi and Nikolaus Kriegeskorte. Deep supervised, but not unsupervised,  
683 models may explain it cortical representation. *PLoS Computational Biology*, 10(11):e1003915,  
684 November 2014. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1003915. URL <http://dx.doi.org/10.1371/journal.pcbi.1003915>.
- 685 Tim C. Kietzmann, Courtney J. Spoerer, Lynn K. A. Sörensen, Radoslaw M. Cichy, Olaf Hauk, and  
686 Nikolaus Kriegeskorte. Recurrence is required to capture the representational dynamics of the  
687 human visual system. *Proceedings of the National Academy of Sciences*, 116(43):21854–21863,  
688 October 2019. ISSN 1091-6490. doi: 10.1073/pnas.1905544116. URL <http://dx.doi.org/10.1073/pnas.1905544116>.
- 689 Hoki Kim. Torchattacks: A pytorch repository for adversarial attacks, 2020.
- 690 Nikolaus Kriegeskorte et al. Matching categorical object representations in inferior temporal cortex  
691 of man and monkey. *Neuron*, 60(6):1126–1141, 2008.
- 692 Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep con-  
693 volutional neural networks. In F. Pereira, C.J. Burges, L. Bottou, and K.Q. Weinberger (eds.),  
694 *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012.

- 702 Jonas Kubilius, Martin Schrimpf, Ha Hong, Najib J. Majaj, Rishi Rajalingham, Elias B. Issa, Kohitij  
703 Kar, Pouya Bashivan, Jonathan Prescott-Roy, Kailyn Schmidt, Aran Nayebi, Daniel Bear, Daniel  
704 L. K. Yamins, and James J. DiCarlo. Brain-Like Object Recognition with High-Performing  
705 Shallow Recurrent ANNs. In H. Wallach, H. Larochelle, A. Beygelzimer, F. D’Alché-  
706 Buc, E. Fox, and R. Garnett (eds.), *Neural Information Processing Systems (NeurIPS)*, pp.  
707 12785–12796. Curran Associates, Inc., 2019. URL [http://papers.nips.cc/paper/  
708 9441-brain-like-object-recognition-with-high-performing-shallow-recurrent-anns](http://papers.nips.cc/paper/9441-brain-like-object-recognition-with-high-performing-shallow-recurrent-anns).
- 709 Yann LeCun, Corinna Cortes, and Christopher J.C. Burges. The mnist database. [https://yann.  
710 lecun.com/exdb/mnist/](https://yann.lecun.com/exdb/mnist/), 1998. [Accessed 16-11-2024].
- 711 Wen Li, Limin Wang, Wei Li, Eirikur Agustsson, and Luc Van Gool. Webvision database: Vi-  
712 sual learning and understanding from web data. *CoRR*, abs/1708.02862, 2017. URL [http:  
713 //arxiv.org/abs/1708.02862](http://arxiv.org/abs/1708.02862).
- 714 Drew Linsley, Ivan F Rodriguez Rodriguez, Thomas FEL, Michael Arcaro, Saloni Sharma,  
715 Margaret Livingstone, and Thomas Serre. Performance-optimized deep neural networks  
716 are evolving into worse models of inferotemporal visual cortex. In A. Oh, T. Nau-  
717 mann, A. Globerson, K. Saenko, M. Hardt, and S. Levine (eds.), *Advances in Neural  
718 Information Processing Systems*, volume 36, pp. 28873–28891. Curran Associates, Inc.,  
719 2023. URL [https://proceedings.neurips.cc/paper\\_files/paper/2023/  
720 file/5bf234ecf83cd77bc5b77a24ba9338b0-Paper-Conference.pdf](https://proceedings.neurips.cc/paper_files/paper/2023/file/5bf234ecf83cd77bc5b77a24ba9338b0-Paper-Conference.pdf).
- 721 Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie.  
722 A convnet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and  
723 Pattern Recognition (CVPR)*, 2022a.
- 724 Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie.  
725 A convnet for the 2020s. *Proceedings of the IEEE/CVF Conference on Computer Vision and  
726 Pattern Recognition (CVPR)*, 2022b.
- 727 Gaëlle Loosli, Léon Bottou, and Patrice Simard. `projects:infimnist` [leon.bottou.org] —  
728 leon.bottou.org. <https://leon.bottou.org/projects/infimnist>, 2007. [Ac-  
729 cessed 16-11-2024].
- 730 TorchVision maintainers and contributors. Torchvision: Pytorch’s computer vision library. [https:  
731 //github.com/pytorch/vision](https://github.com/pytorch/vision), 2016.
- 732 Najib J. Majaj, Ha Hong, Ethan A. Solomon, and James J. DiCarlo. Simple learned weighted  
733 sums of inferior temporal neuronal firing rates accurately predict human core object recognition  
734 performance. *The Journal of Neuroscience*, 35(39):13402–13418, 2015. doi: 10.1523/jneurosci.  
735 5181-14.2015.
- 736 Rafael Malach, Itzhak Levy, and Uri Hasson. The topography of high-order human object areas.  
737 *Trends in Cognitive Sciences*, 6(4):176–184, 2002.
- 738 Johannes Mehrer, Courtney J. Spoerer, Emer C. Jones, Nikolaus Kriegeskorte, and Tim C. Ki-  
739 etzmann. An ecologically motivated image dataset for deep learning yields better models  
740 of human vision. *Proceedings of the National Academy of Sciences*, 118(8), 2021. doi:  
741 10.1073/pnas.2011417118.
- 742 Sachin Mehta and Mohammad Rastegari. Mobilevit: Light-weight, general-purpose, and mobile-  
743 friendly vision transformer. In *International Conference on Learning Representations*, 2022. URL  
744 <https://openreview.net/forum?id=vh-0sUt8HlG>.
- 745 Niklas Muennighoff, Alexander M Rush, Boaz Barak, Teven Le Scao, Nouamane Tazi, Aleksan-  
746 dra Piktus, Sampo Pyysalo, Thomas Wolf, and Colin Raffel. Scaling data-constrained language  
747 models. In *Thirty-seventh Conference on Neural Information Processing Systems*, 2023. URL  
748 <https://openreview.net/forum?id=j5BuTrEj35>.
- 749 Lukas Muttenthaler, Jonas Dippel, Lorenz Linhardt, Robert A. Vandermeulen, and Simon Korn-  
750 blith. Human alignment of neural network representations. In *The Eleventh International Confer-  
751 ence on Learning Representations*, 2023. URL [https://openreview.net/forum?id=  
752 ReDQ1OUQR0X](https://openreview.net/forum?id=ReDQ1OUQR0X).

- 756 Aran Nayebi, Daniel Bear, Jonas Kubilius, Kohitij Kar, Surya Ganguli, David Sussillo, James J  
757 DiCarlo, and Daniel L. Yamins. Task-driven convolutional recurrent models of the visual  
758 system. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and  
759 R. Garnett (eds.), *Advances in Neural Information Processing Systems*, volume 31. Curran  
760 Associates, Inc., 2018. URL [https://proceedings.neurips.cc/paper\\_files/  
761 paper/2018/file/6be93f7a96fed60c477d30ae1de032fd-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2018/file/6be93f7a96fed60c477d30ae1de032fd-Paper.pdf).
- 762  
763 Maxime Oquab, Timothée Darcet, Theo Moutakanni, Huy V. Vo, Marc Szafraniec, Vasil Khalidov,  
764 Pierre Fernandez, Daniel Haziza, Francisco Massa, Alaaeldin El-Nouby, Russell Howes, Po-Yao  
765 Huang, Hu Xu, Vasu Sharma, Shang-Wen Li, Wojciech Galuba, Mike Rabbat, Mido Assran,  
766 Nicolas Ballas, Gabriel Synnaeve, Ishan Misra, Herve Jegou, Julien Mairal, Patrick Labatut, Ar-  
767 mand Joulin, and Piotr Bojanowski. Dinov2: Learning robust visual features without supervision,  
768 2023.
- 769 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan,  
770 Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas  
771 Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy,  
772 Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-  
773 performance deep learning library. In *Advances in Neural Information Processing Systems* 32, pp.  
774 8024–8035. Curran Associates, Inc., 2019. URL [http://papers.neurips.cc/paper/  
775 9015-pytorch-an-imperative-style-high-performance-deep-learning-library.  
776 pdf](http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf).
- 777 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agar-  
778 wal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya  
779 Sutskever. Learning transferable visual models from natural language supervision, 2021. URL  
780 <https://arxiv.org/abs/2103.00020>.
- 781  
782 Rishi Rajalingham, Elias B. Issa, Pouya Bashivan, Kohitij Kar, Kailyn Schmidt, and James J. Di-  
783 Carlo. Large-scale, high-resolution comparison of the core visual object recognition behavior  
784 of humans, monkeys, and state-of-the-art deep artificial neural networks. *The Journal of Neuro-  
785 science*, 38(33):7255–7269, 2018. doi: 10.1523/jneurosci.0388-18.2018.
- 786  
787 Tal Ridnik, Emanuel Ben-Baruch, Asaf Noy, and Lihi Zelnik-Manor. Imagenet-21k pretraining for  
788 the masses, 2021.
- 789 Martin Schrimpf, Jonas Kubilius, Ha Hong, Najib J. Majaj, Rishi Rajalingham, Elias B. Issa, Kohitij  
790 Kar, Pouya Bashivan, Jonathan Prescott-Roy, Franziska Geiger, Kailyn Schmidt, Daniel L. K.  
791 Yamins, and James J. DiCarlo. Brain-score: Which artificial neural network for object recognition  
792 is most brain-like? *bioRxiv preprint*, 2018. URL [https://www.biorxiv.org/content/  
793 10.1101/407007v2](https://www.biorxiv.org/content/10.1101/407007v2).
- 794  
795 Martin Schrimpf, Jonas Kubilius, Michael J Lee, N Apurva Ratan Murty, Robert Ajemian, and  
796 James J DiCarlo. Integrative benchmarking to advance neurally mechanistic models of hu-  
797 man intelligence. *Neuron*, 2020. URL [https://www.cell.com/neuron/fulltext/  
798 S0896-6273\(20\)30605-X](https://www.cell.com/neuron/fulltext/S0896-6273(20)30605-X).
- 799  
800 Christoph Schuhmann, Richard Vencu, Romain Beaumont, Robert Kaczmarczyk, Clayton Mullis,  
801 Aarush Katta, Theo Coombes, Jenia Jitsev, and Aran Komatsuzaki. Laion-400m: Open dataset of  
802 clip-filtered 400 million image-text pairs, 2021.
- 803  
804 Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W Gordon, Ross Wightman,  
805 Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick  
806 Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk,  
807 and Jenia Jitsev. LAION-5b: An open large-scale dataset for training next generation image-text  
808 models. In *Thirty-sixth Conference on Neural Information Processing Systems Datasets and  
809 Benchmarks Track*, 2022. URL <https://openreview.net/forum?id=M3Y74vmsMcY>.
- Igor Susmelj, Matthias Heller, Philipp Wirth, Jeremy Prescott, and Malte Ebner. Lightly, 2020.

- 810 Mingxing Tan and Quoc Le. EfficientNet: Rethinking model scaling for convolutional neural net-  
811 works. In Kamalika Chaudhuri and Ruslan Salakhutdinov (eds.), *Proceedings of the 36th In-*  
812 *ternational Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning*  
813 *Research*, pp. 6105–6114. PMLR, 09–15 Jun 2019. URL [https://proceedings.mlr.](https://proceedings.mlr.press/v97/tan19a.html)  
814 [press/v97/tan19a.html](https://proceedings.mlr.press/v97/tan19a.html).
- 815 The Mosaic ML Team. composer. <https://github.com/mosaicml/composer/>, 2021.  
816
- 817 Hugo Touvron, Matthieu Cord, and Hervé Jégou. Deit iii: Revenge of the vit, 2022. URL <https://arxiv.org/abs/2204.07118>.  
818
- 819 Zhengzhong Tu, Hossein Talebi, Han Zhang, Feng Yang, Peyman Milanfar, Alan Bovik, and Yinxiao  
820 Li. Maxvit: Multi-axis vision transformer. *ECCV*, 2022.  
821
- 822 Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam,  
823 Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In  
824 *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June  
825 2018.
- 826 Pavan Kumar Anasosalu Vasu, James Gabriel, Jeff Zhu, Oncel Tuzel, and Anurag Ranjan. Fastvit:  
827 A fast hybrid vision transformer using structural reparameterization. In *Proceedings of the*  
828 *IEEE/CVF International Conference on Computer Vision*, 2023.  
829
- 830 Ross Wightman. Pytorch image models. [https://github.com/rwightman/](https://github.com/rwightman/pytorch-image-models)  
831 [pytorch-image-models](https://github.com/rwightman/pytorch-image-models), 2019.
- 832 Eric Wong, Leslie Rice, and J. Zico Kolter. Fast is better than free: Revisiting adversarial training. In  
833 *International Conference on Learning Representations*, 2020. URL [https://openreview.](https://openreview.net/forum?id=BJx040EFvH)  
834 [net/forum?id=BJx040EFvH](https://openreview.net/forum?id=BJx040EFvH).
- 835 Daniel L. Yamins, Ha Hong, Charles F. Cadieu, Ethan A. Solomon, Darren Seibert, and James J.  
836 DiCarlo. Performance-optimized hierarchical models predict neural responses in higher visual  
837 cortex. *Proceedings of the National Academy of Sciences*, 111(23):8619–8624, 2014. doi: 10.  
838 1073/pnas.1403112111.  
839
- 840 Daniel L K Yamins and James J DiCarlo. Using goal-driven deep learning models to understand  
841 sensory cortex. *Nature Neuroscience*, 19(3):356–365, February 2016. ISSN 1546-1726. doi:  
842 10.1038/nn.4244. URL <http://dx.doi.org/10.1038/nn.4244>.
- 843 Xiaohua Zhai, Alexander Kolesnikov, Neil Houlsby, and Lucas Beyer. Scaling vision transform-  
844 ers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*  
845 *(CVPR)*, pp. 12104–12113, June 2022.  
846
- 847 Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 mil-  
848 lion image database for scene recognition. *IEEE Transactions on Pattern Analysis and Machine*  
849 *Intelligence*, 2017.
- 850 Chengxu Zhuang, Siming Yan, Aran Nayebi, Martin Schrimpf, Michael C. Frank, James J. DiCarlo,  
851 and Daniel L. Yamins. Unsupervised neural network models of the ventral visual stream. *Pro-*  
852 *ceedings of the National Academy of Sciences*, 118(3), 2021. doi: 10.1073/pnas.2014196118.  
853  
854  
855  
856  
857  
858  
859  
860  
861  
862  
863



## A IMPLEMENTATION DETAILS

Our experiments are conducted using the PyTorch framework (Paszke et al., 2019), with Composer (Team, 2021) employed as the GPU orchestration tool to efficiently manage computational resources.

For image augmentations, we leverage the Albumentations Buslaev et al., 2020 library due to its rich set of augmentation techniques, which are crucial for enhancing model robustness and preventing overfitting. In experiments involving self-supervised learning, we use the Lightly Susmelj et al. (2020) library to facilitate the implementation of self-supervised losses, augmentations, and model heads. This library streamlines the process of setting up models for SimCLR and DINO training methods.

To generate adversarial examples for adversarial fine-tuning, we employ the Torchattacks library Kim (2020). Specifically, we use the Fast Gradient Sign Method (FGSM) to create perturbations that challenge the models, aiming to enhance their alignment with neural representations by exposing them to adversarial inputs.

## B ADDITIONAL IMAGE DATASETS

To further validate our findings across diverse image distributions and to estimate scaling curves across different sample scales, we trained ResNet18 models on subsets of several large-scale image datasets: ImageNet-21k-P, WebVision-P, iNaturalist, and Places365. Below, we provide detailed descriptions of each dataset.

### B.1 IMAGENET21K-P

ImageNet-21k-P is a processed subset of the full ImageNet-21k dataset (Ridnik et al., 2021), which originally contains over 14 million images organized into more than 21,000 categories following the WordNet hierarchy. The "P" denotes a pruned version where classes with insufficient images or noisy labels are filtered out to enhance dataset quality. This results in a refined dataset that maintains the richness of the original ImageNet-21k while improving label accuracy and image relevance. The resulting dataset contains approximately 11 million training images across 10,450 classes.

### B.2 WEBVISION-P

The WebVision dataset (Li et al., 2017) is a large-scale web image dataset designed to provide a real-world, noisy alternative to ImageNet. It originally contains over 16 million images categorized into 5,000 classes. The images are collected from the internet using queries from search engines like Google and Flickr, leading to a dataset that includes label noise, varying image resolutions, and diverse visual contexts. Due to classes with very few available samples, we processed the WebVision dataset similarly to ImageNet-21k-P to remove classes with insufficient images. The resulting dataset, which we denote as WebVision-P, contains approximately 13.5 million training images across 4,189 categories.

### B.3 INATURALIST

iNaturalist Van Horn et al. (2018) contains 2.7 million photographs of organisms in their natural environments, representing 10,000 species. The dataset features highly specialized fine-grained categories and natural backgrounds, offering insight into how domain-specific visual features influence alignment scaling.

### B.4 PLACES365

Places365 Zhou et al. (2017) is a large-scale scene-centric dataset containing approximately 1.8 million training images across 365 scene categories. Unlike object-centric datasets such as ImageNet, Places365 focuses on the recognition of environmental scenes, including natural landscapes, urban settings, and indoor environments. Each category includes a wide variety of images to capture the diversity within scene types.

918  
919  
920  
921  
922  
923  
924  
925  
926  
927  
928  
929  
930  
931  
932  
933  
934  
935  
936  
937  
938  
939  
940  
941  
942  
943  
944  
945  
946  
947  
948  
949  
950  
951  
952  
953  
954  
955  
956  
957  
958  
959  
960  
961  
962  
963  
964  
965  
966  
967  
968  
969  
970  
971

## B.5 INFIMNIST

The MNIST dataset (LeCun et al., 1998) is a classic benchmark in machine learning, comprising 70,000 grayscale images of handwritten digits (0-9), each sized 28×28 pixels. To expand this dataset for more extensive experimentation, we utilize the Infinite MNIST (Infimnist) tool Loosli et al., 2007, which generates additional MNIST-like samples through data augmentation techniques. We create an extended dataset by modifying the original training dataset 19 additional times, resulting in a total of 1.2 million images. This enlarged dataset allows for a more thorough evaluation of scaling effects on the alignment.

## C VALIDATION ON PRIVATE DATA

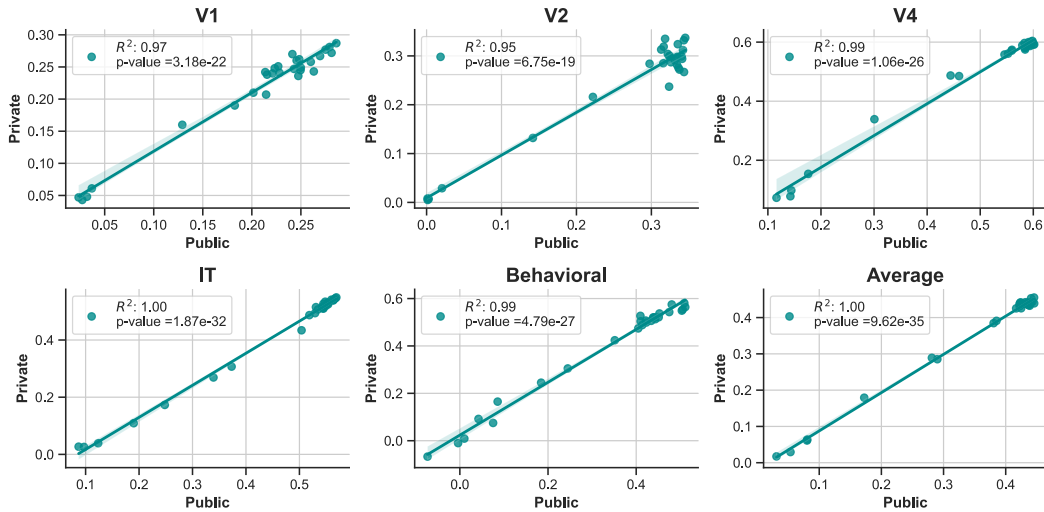


Figure 8: Public benchmarks used in this study correlates highly with private benchmarks on Brain-Score

As described in Section 3 we test a diverse set of models on private benchmarks on Brain-Score platform. All  $R^2$  values are above 0.95 with p-values less than  $10^{-18}$ .

## D PRETRAINED MODELS

As part of our comprehensive evaluation, we benchmarked a diverse set of pretrained models sourced from both `torchvision` (maintainers & contributors, 2016) and the `timm` (Wightman, 2019) libraries. We tested a total of 94, including ViT Dosovitskiy et al. (2021), DaViT Ding et al. (2022), LeViT Graham et al. (2021), ConvNeXt Liu et al. (2022a), MobileViT Mehta & Rastegari (2022), MaxViT Tu et al. (2022), FastViT Vasu et al. (2023). Each model varies in parameter count, training sample size, dataset source, and training objective, providing a broad spectrum for analysis.

To verify the generalizability of our findings, we conducted evaluations with these pretrained models, including larger networks like CLIP (Radford et al., 2021) and DINOv2 (Oquab et al., 2023), which are pretrained on richer and more diverse datasets such as LAION Schuhmann et al. (2021; 2022). We also compared variations of these models by examining base pretrained models alongside their fine-tuned counterparts on ImageNet, aiming to investigate the impact of fine-tuning on scaling behavior.

Our results indicate that models with extensive pretraining achieve enhanced behavioral alignment, likely due to their exposure to richer and more varied data. However, similar to models trained solely on ImageNet or EcoSet, these pretrained models still exhibit a saturation effect in neural alignment with the primate visual ventral stream (VVS). This suggests that while larger and more diverse datasets improve behavioral predictability, they do not substantially extend the scaling of neural alignment beyond the observed plateau.

The curves in Figure 9 closely follow the scaling patterns estimated for our trained models shown in Figure 2.c, further validating that the observed saturation is consistent across different pretraining regimes and dataset scales. This reinforces our conclusion that scaling alone is insufficient to overcome the limitations in neural alignment and highlights the need for alternative approaches to improve alignment with neural representations.

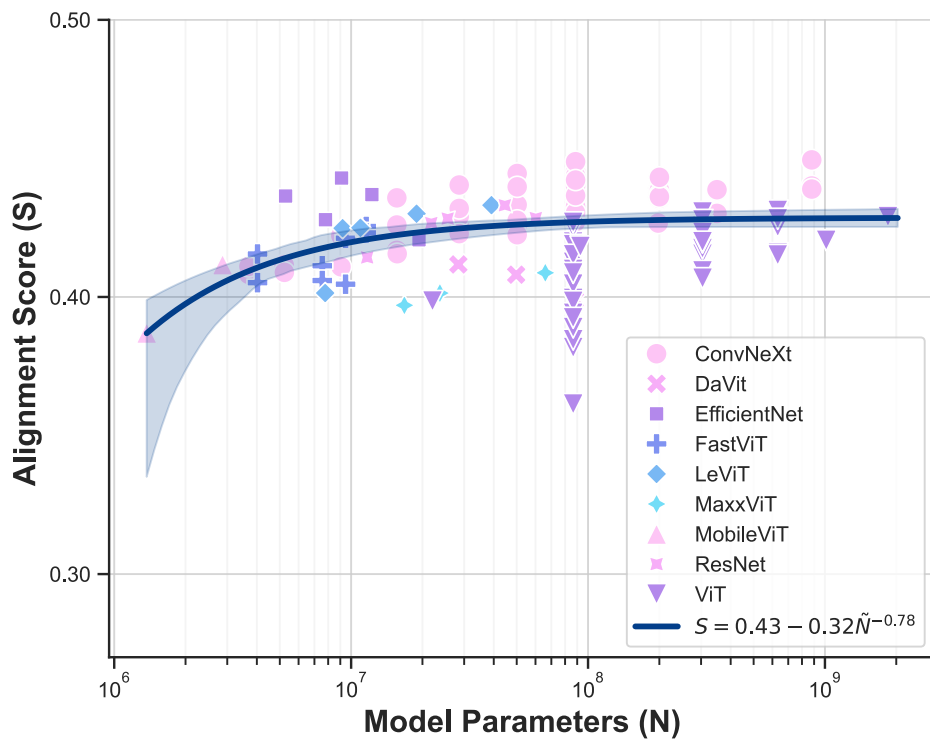


Figure 9: Alignment of pretrained vision models as a function of model parameters

## E TRAINING EVOLUTION

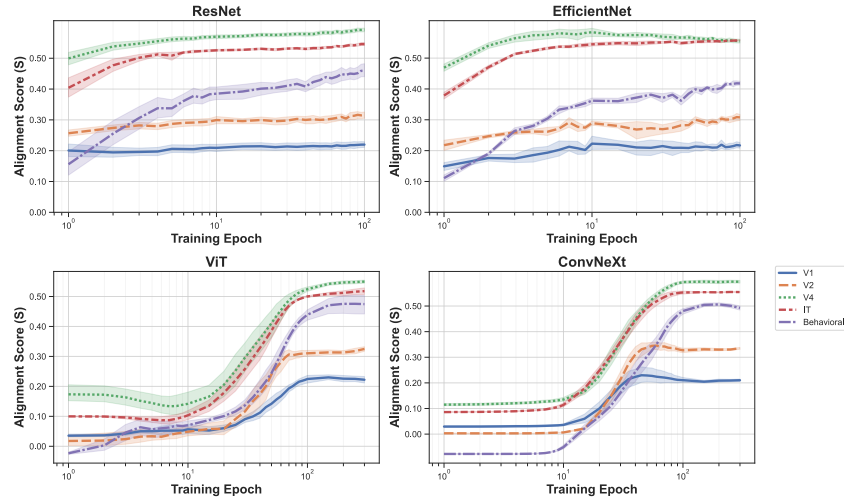


Figure 10: Evolution of per-region alignment throughout training. Models with stronger priors—such as ResNet and EfficientNet—exhibit higher neural alignment initially. However, the gap in representational power diminishes as more generalist models are trained on data.




Figure 11: Some untrained models demonstrate non-zero neural alignment at initialization. Nevertheless, all models start with almost zero behavioral alignment, indicating that initial neural alignment arises from architectural biases rather than learned behavior.

F EFFECT OF TRAINING OBJECTIVE

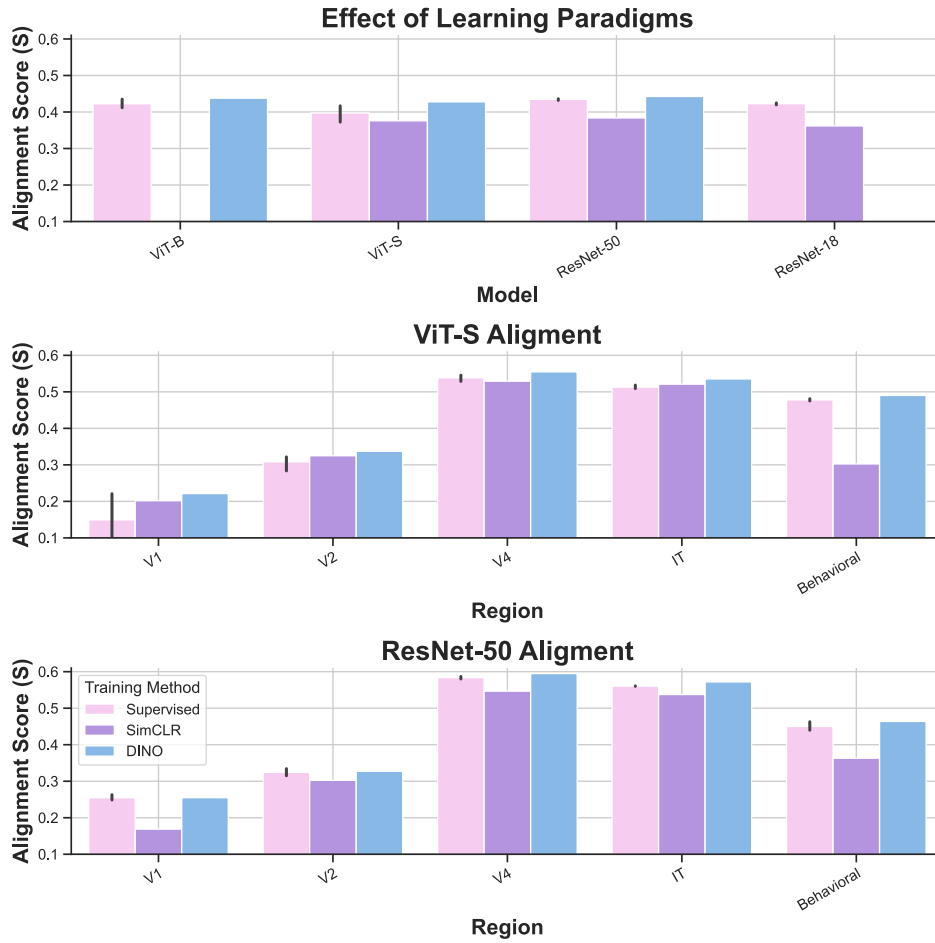


Figure 12: Effect of training objective on alignment

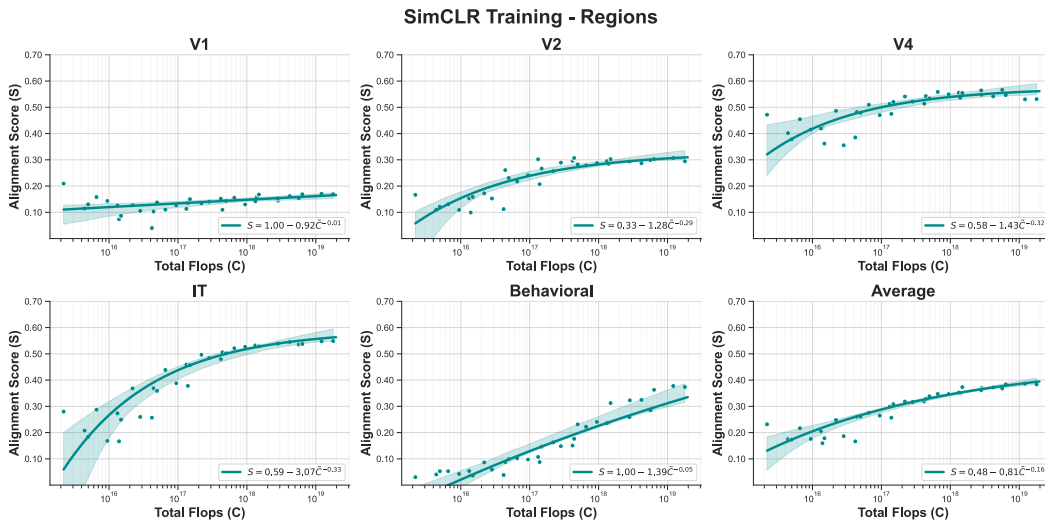


Figure 13: Scaling of SimCLR training across regions

G REGION-WISE SCATTER PLOTS

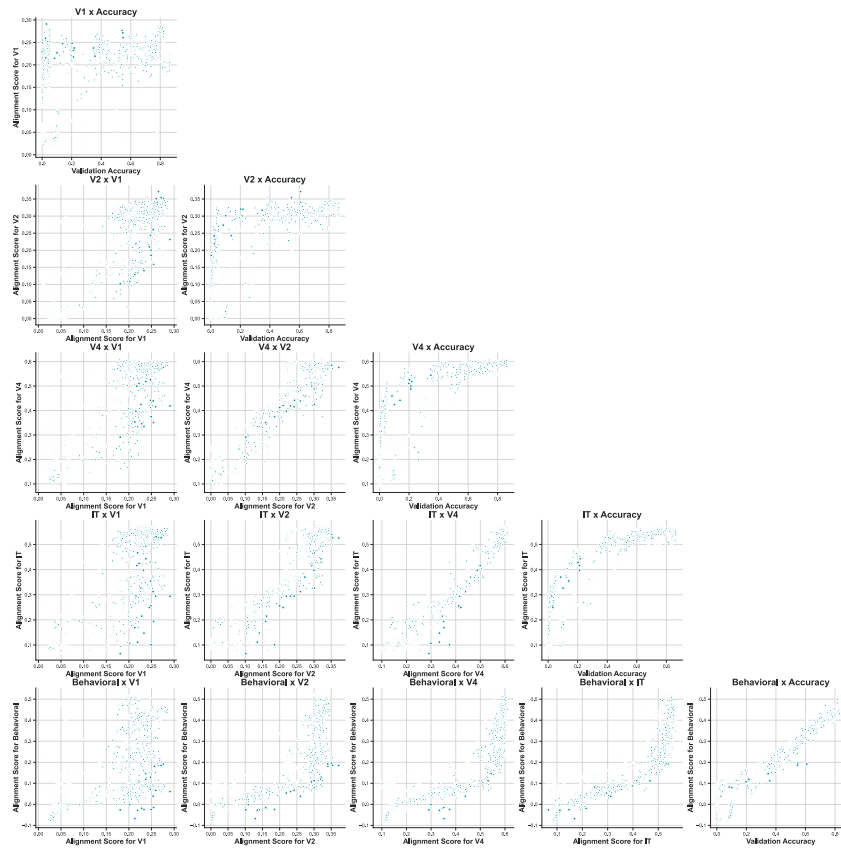


Figure 14: Region vs. Region Comparisons: This figure shows how the alignment scores for each brain region correlate with those of other regions. The diagonal plots illustrate the relationship between the alignment score of each region and the validation accuracy on ImageNet and EcoSet datasets.