

000 001 002 003 004 005 006 007 008 009 010 011 012 013 014 015 016 017 018 019 020 021 022 023 024 025 026 027 028 029 030 031 032 033 034 035 036 037 038 039 040 041 042 043 044 045 046 047 048 049 050 051 052 053 FEATURE SEGREGATION BY SIGNED WEIGHTS IN NEURAL NETWORKS FOR ARTIFICIAL AND BIOLOGICAL VISION

Anonymous authors

Paper under double-blind review

ABSTRACT

Signed connections are central to both artificial and biological intelligence, positive and negative weights in artificial networks, and excitatory and inhibitory synapses in the brain, yet their representational role remains unclear. Here, we investigate how signed weights shape visual representations in artificial and biological systems involved in object recognition. Using sign consistency as a proxy for biological Dale’s law, which requires neurons to send either exclusively excitatory or inhibitory outputs, we found that accuracy of ImageNet trained networks positively correlated with the Dale index of their output layer. Ablation and feature visualization reveal a functional segregation: removing positive inputs disrupts object related, low frequency structure, while removing negative inputs mainly alters background textures. This segregation is more pronounced in adversarially robust models, persists with unsupervised learning, and vanishes with non-rectified activations. In intermediate layers, the most positive Dale-like channels encoded localized, object-like features, whereas the most negative ones captured dispersed, background features. We next performed *in vivo* feature visualization in monkey ventral visual cortex (V1, V4, and IT) and fitted linear models using the input layer to the neural networks classification units. These models reproduced features similar to those preferred by the biological neurons. In the model neurons, removing positive inputs altered representations more than removing negative ones. The most Dale-like positively projecting units exhibited localized features, while the negatively projecting units showed larger, more dispersed features, suited to carrying contextual input. Consistent with this, clearing the background around each neuron’s preferred feature enhanced its response, likely by reducing inhibitory drive, supporting inhibition as a contextual modulation of the excitatory feature. Our results demonstrate that both artificial and biological vision systems segregate features by weight sign: positive weights emphasize objects and low frequencies, negative weights encode context. This shows the convergence of representational strategies in brains and machines, yielding predictions for visual neuroscience.

1 INTRODUCTION

Brains and artificial neural networks (ANNs) both rely on signed connections. In biological circuits, Dale’s law states that neurons are either excitatory or inhibitory (Dale, 1935). Excitatory neurons are thought to compute core visual features, while inhibitory neurons sharpen selectivity, modulate context, and gate information flow (Isaacson & Scanziani, 2011). The primate ventral visual stream supports object recognition (DiCarlo et al., 2012), and representations from V1 to IT increase in complexity, paralleling hierarchies in convolutional neural networks (CNNs) (Yamins et al., 2014; Güçlü & van Gerven, 2017). While early mechanisms such as lateral inhibition and center-surround receptive fields are well characterized, the excitatory–inhibitory organization of higher ventral-stream areas remains less understood (Tamura et al., 2004), motivating the question of whether information is segregated by connection sign in both artificial and biological vision.

ANNs also use positive and negative weights, loosely analogous to excitatory and inhibitory signals but unconstrained by Dale’s law. Yet it remains unclear how deep networks divide visual information across signed weights, particularly in object classification models where each output unit represents

054 a category. Prior work suggests segregation by absolute weight strength (Li et al., 2023), , but
 055 whether feature types such as foreground objects versus background context systematically separate
 056 by weight sign is unknown. Here, we hypothesize that CNNs segregate visual information into
 057 positive and negative inputs.

058 We test this hypothesis across diverse ImageNet-trained CNNs by analyzing weight sign organization,
 059 ablating positive and negative inputs, and visualizing the resulting feature selectivity. We asses
 060 feature segregation across network depth by analyzing Dale-like channels. We examine a range of
 061 architectures and training regimes, including adversarially robust models, unsupervised models, and
 062 networks with nonrectified (Tanh) activations.

063 For biological comparison, we fit linear models from ANN features to neural responses recorded
 064 across the macaque ventral stream (V1, V4, IT) and use *in vivo* feature visualization and background
 065 manipulations to probe inhibitory contributions. These analyses relate ANN-derived sign structure
 066 to biological feature preferences.

067 Our results support an emerging principle: across artificial and biological vision systems, connection
 068 sign organizes feature representation. Positive weights emphasize object-related, localized features,
 069 whereas negative weights encode contextual, more dispersed structure. This connects classic ideas
 070 rooted in Dale’s law with representational strategies in modern ANNs, generating mechanistic and
 071 testable predictions for visual neuroscience.

073 2 RELATED WORK

074 **Mechanistic interpretability of computer and biological vision** There has been progress in mechanistic
 075 interpretability in ANNs from work using perspectives adapted from neuroscience circuit
 076 dissection (Olah et al., 2020). This line of explainable AI research explains model behavior by dis-
 077 secting smaller network subgraphs, revealing how relevant features arise from input weights and are
 078 composed hierarchically. Such work has uncovered motifs involving positive and negative connec-
 079 tions between related features, reminiscent of early visual system organization. New approaches
 080 to address representations beyond single units rely on sparse dictionary learning, with early work
 081 in vision (Olshausen & Field, 1996), an approach that has recently regained popularity in language
 082 modeling (Cunningham et al., 2023), as well as in multimodal models (Pach et al., 2025). Some
 083 studies also characterized object shape and texture biases in feature visualizations by reconstructing
 084 images from sparse weight sets (Li et al., 2023). However, the systematic division between pos-
 085 itive and negative inputs across the entire range of weight strengths, and its possible role in feature
 086 segregation, remains underexplored and is a focus of this study.

087 **Feature visualization by closed-loop optimization** Characterizing learned representations is foun-
 088 dational for both biological and artificial vision research. Feature visualization, i.e. optimizing for
 089 images that strongly activate target units, was originally pioneered in the brain by hand (Hubel &
 090 Wiesel, 1959), and later *in silico* by gradient ascent on pixels of neural networks (Erhan et al., 2009;
 091 Nguyen et al., 2016a;b; Olah et al., 2017). Because gradients are unavailable when recording *in vivo*,
 092 gradient-free black-box optimization techniques were developed for synthesizing preferred images
 093 of biological neurons in real-time (Ponce et al., 2019; Xiao & Kreiman, 2020; Wang & Ponce, 2022).
 094 These approaches constrain the search space via generative adversarial networks, promoting natu-
 095 ralistic solutions (Nguyen et al., 2016a). Further methods involve first fitting a predictive network
 096 to neural data and then using *in silico* gradient ascent (Bashivan et al., 2019; Walker et al., 2019).
 097 While most prior studies use grayscale images, our study applies gradient-free visualization to color
 098 images in both CNNs and primate recordings.

099 **Robustness** Neural networks are susceptible to adversarial attacks, where noise that is nearly imper-
 100 ceptible by humans can be added to natural images, changing output classification (Szegedy et al.,
 101 2014; Salman et al., 2020; Elsayed et al., 2018). Robust training, i.e. introducing adversarial per-
 102 turbations during learning, improves resistance to such attacks and is hypothesized to align learned
 103 representations more closely with primate visual processing. Prior work does not assess how robust-
 104 ness impacts the organization of image representations after ablation of signed weights, which we
 105 systematically investigate here.

106 **Nonlinearity influence on representations** Beyond training objectives, the role of activation func-
 107 tions such as ReLU versus Tanh profoundly influences representational properties (Alleman et al.,

108 2023), with ReLU inducing representations better aligned with input features and Tanh inducing
 109 alignment with output features (labels). This prior work was done in small networks from a theo-
 110 retical perspective; thus, the impact of rectification on the potential segregation of visual features at
 111 practical scales is unknown and addressed by our study.
 112

113 3 METHODS

114 An extended methods section is in the Appendix A.1.

115 **Networks** We performed our ablation studies in CNNs pretrained on the ImageNet dataset: AlexNet
 116 (Krizhevsky et al., 2012), VGG16 (Simonyan & Zisserman, 2015), ResNet50 (He et al., 2015), and
 117 robust ResNet50 models with robustness radii specified by the L_∞ norm ($L_\infty \in \{0.5, 1, 2, 4, 8\}$);
 118 Salman et al. 2020). Here, L_∞ denotes the maximum-norm constraint used during adversarial train-
 119 ing, i.e., $\|\delta\|_\infty \leq \epsilon$, which bounds the maximum per-pixel perturbation by ϵ . To reduce computing
 120 time, we used the *Imagenette* dataset (noa, 2024) and the *ImageNet* macaque category. We also
 121 tested 100 classes sampled by k-means on the output of ResNet50 (Figs. 12,13) For all networks, we
 122 visualized the representations of the units in the fully-connected output layer (pre-softmax) matching
 123 those classes under different ablation conditions.
 124

125 **Dale index** We quantified Dale like structure with a Dale index for each outgoing channel per layer,
 126 defined as the fraction of its weights that share the majority sign, $D = \max(p_+, p_-)$, where p_+ and
 127 p_- are the proportions of positive and negative outgoing weights. The index ranges from 0.5 to 1
 128 and measures sign consistency.
 129

130 **Ablation** For each layer, we ablated positive and negative weights separately. Given a layer’s weight
 131 matrix W , we defined the sets of positive weights $P = \{w \in W : w > 0\}$ and negative weights
 132 $N = \{w \in W : w < 0\}$. For each set $S \in \{P, N\}$, we sorted its elements in decreasing order of
 133 absolute value. We then defined the ablation strength $\alpha \in [0, 1]$ as the fraction of the total magnitude
 134 of S to remove. Specifically, we identified the smallest k satisfying $\sum_{i=1}^k |w_i| / \sum_{w \in S} |w| \geq \alpha$,
 135 where w_1, w_2, \dots are the sorted weights in S , and set those k weights to zero. Because α is a
 136 normalized cumulative magnitude, it lies in $[0, 1]$, and sweeping α from 0 to 1 removes none to all
 137 of the positive (or negative) weights.
 138

139 **Feature visualization** For each ablation condition, we performed feature visualization by optimiz-
 140 ing a GAN latent code to create an activity-maximizing image. We used this closed-loop, zeroth-
 141 order-search approach to allow comparison with our neuronal experiments, where gradient ascent
 142 would not be possible. To increase the span of the stimulus space, we used two GANs: AlexNet
 143 fc6 DeePSiM (Dosovitskiy & Brox, 2016) which can render textures and objects, and BigGAN
 144 (Brock et al., 2019) that can render photo-realistic images with objects. For optimization, we used
 145 a variant of *covariance matrix adaptation evolutionary strategy* or CMAES (Wang & Ponce, 2022;
 146 Loshchilov, 2015). We optimized ten images per GAN, resulting in 20 feature visualizations per
 147 output unit and ablation condition. Diverse visualizations better capture the multifaceted high-level
 148 representations in CNNs (Nguyen et al., 2016b). For our examples, we show the best of the 20 vi-
 149 sualizations, but used all for quantitative analyses. For visualizations of neural networks predicting
 150 biological neuron responses, due to experimental time restrictions, we used five visualizations per
 151 ablation condition, via DeePSim only. Our experiments are performed in a PC with Nvidia 4090
 152 GPU, and each visualization takes about 3 mins.
 153

154 **Network training** Both ResNet18 networks were trained using the FFCV library (Leclerc et al.,
 155 2023) for 16 epochs on the ImageNet1K dataset. The top-5 classification accuracy was 0.797 for
 156 the network with Tanh activations and 0.870 for the network with ReLU activations. Note that these
 157 models were trained for only 16 epochs rather than the standard 90, so their accuracy underperforms
 158 published benchmarks. However, they are suitable for our mechanistic analyses.
 159

160 **Visual cortex electrophysiology** We recorded multi-unit (neuron microcluster) and occasional
 161 single-unit activity from chronically implanted multielectrode arrays in V1, V4, and PIT of two
 162 macaques. Animals fixated while 2–8° images were flashed briefly (100 ms ON, 150 ms inter-
 163 stimulus interval). Neurons were driven with a 160-image stimulus set (diverseSet), constructed to
 164 span a broad range of visual features derived from k-means clustering of AlexNet output layer over
 165 the IN1K validation set and typical visual neuroscience image sets. For each session, we modeled
 166 responses of a single neuron or multiunit using a one-component PLS regression between firing rates
 167

162 and AlexNet penultimate-layer activations, and used the resulting model for *in silico* ablations and
 163 feature visualizations, validating visualizations *in vivo* within the same session.
 164

165 4 RESULTS

166 4.1 PROXIMITY TO DALE'S LAW IN CNN OUTPUTS CORRELATES WITH ACCURACY

167 A key challenge in comparing artificial
 168 and biological circuits is that CNNs are
 169 not constrained by Dale's law. Therefore,
 170 we asked whether the Dale index (our
 171 measure of sign consistency) in the
 172 output layer of diverse CNNs has any relation
 173 to its performance. Dale index increased
 174 from random initialization with training
 175 (Fig 1A). Top 1 accuracy on ImageNet1K
 176 positively correlated with the mean Dale
 177 index of the output layer. Within a given
 178 architecture, the Dale index increased with
 179 network depth. And batchnorm training in
 180 VGGs produced output layers with higher
 181 Dale index. Thus, even without an ex-
 182 plicit Dale constraint, these networks nat-
 183 urally developed more sign consistent out-
 184 put channels. This motivated us to exam-
 185 ine the specific visual features carried by
 186 positive versus negative weights.
 187

188 4.2 OBJECT INFORMATION IS PREFERENTIALLY ENCODED BY POSITIVE WEIGHTS

189 **Hypothesis** Motivated by biological center-surround organization, where inhibitory surrounds
 190 convey contextual information, we hypothesized that CNN output units trained for object recognition
 191 segregate object features to positive weights and contextual features to negative weights.
 192

193 **Testing segregation by ablation and visualization.** We examined this hypothesis in ImageNet-
 194 pretrained CNNs using ablation and feature visualization. The overall ratio of positive to negative
 195 input weights per unit was close to unity (Table 2), suggesting that both polarities may encode rel-
 196 evant information. We then selectively ablated positive or negative input weights to class units and
 197 visualized features across ablation strengths. Ablating positive weights greatly reduced the max-
 198 imal achievable activation during feature visualization, whereas removing negative weights slightly
 199 increased it (appendix Fig. 11). Visually, positive-weight ablation disrupted recognizable object
 200 structure, while negative-weight ablation largely preserved object identity and instead altered back-
 201 ground or color context (Figs. 2B, C). To quantify these effects, we compared image sets generated
 202 before and after ablation using mean pairwise cosine similarity across an ensemble of readout CNNs.
 203 Positive-weight ablation yielded substantially less similar representations, whereas negative-weight
 204 ablation produced only minor shifts (Fig. 2D). These results replicated across 100 ImageNet classes
 205 and with alternative metrics such as LPIPS (Zhang et al., 2018) (appendix Fig. 13), demonstrating
 206 robustness and generality.

207 To quantify to what extent objects disappear from the preferred images under ablations, we evaluated
 208 objectness using an object-detection network (YOLOv7 Wang et al. (2022)). Relative to baseline
 209 objectness scores computed from intact visualizations, ablating positive weights reduced objectness,
 210 whereas ablating negative weights had minimal impact (Fig. 2E). Analyzing spatial frequencies re-
 211 vealed positive ablations majorly affected low frequencies (app. Fig. 14,15), consistent with the
 212 objectness reduction. Together, these results show that in ImageNet-trained CNNs, removing posi-
 213 tive input weights disrupts object features, while removing negative weights primarily alters context,
 214 indicating that positive weights preferentially encode object features.

215 4.3 SEGREGATION DEPENDS ON RELU BUT NOT ON UNSUPERVISED PRETRAINING

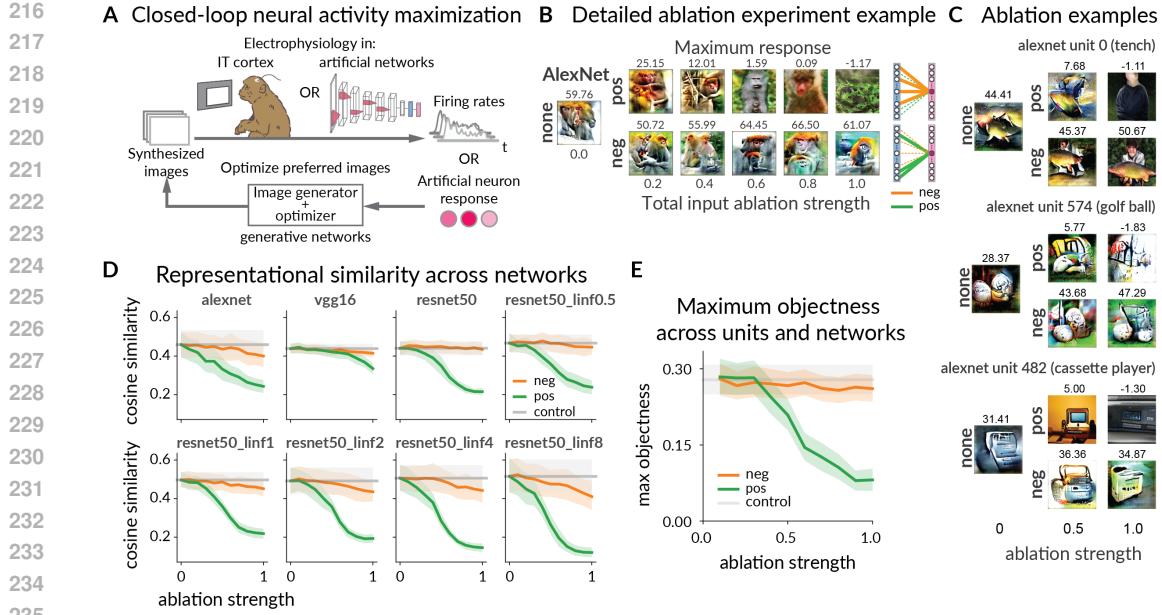


Figure 2: **A.** Schematic of feature visualization workflow in ANNs and brains. **B.** Preferred feature changes for different ablation strengths of input weights to the macaque 373 output unit of AlexNet (last fc layer of 1000 units before softmax). Images are the most activating images out of the 20 visualizations per ablation strength. Ablation strengths are below each image, and activation scores are above. **C.** Changes in preferred features to different ablations of example AlexNet output units: 0 tench, 574 golf ball, and 482 cassette player. Notice the large image changes for positive ablations. Same methods as in **B**. **D.** Representational similarity of intact vs input-ablated units across networks tested, measured by the pairwise cosine similarity of control vs ablation images over an ensemble of networks. Error bars are 95% confidence intervals over units, each unit is the mean of its 20 visualizations. The units correspond to the 10 imagenette categories ([0, 217, 482, 491, 497, 566, 569, 571, 574, 701]) plus the macaque category (373). **E.** Objectness scores across units and networks over ablation condition. As in **D**, we tested 11 units from the 1000-unit fully-connected output layer (pre-softmax) of: AlexNet, VGG16, ResNet50, and robust ResNet50 ($L_\infty \in \{0.5, 1, 2, 4, 8\}$). For each network, we averaged over the objectness scores of 20 visualizations per unit and all units. The plot shows the mean over previously described network averages. Error bars are 95% confidence interval over network averages.

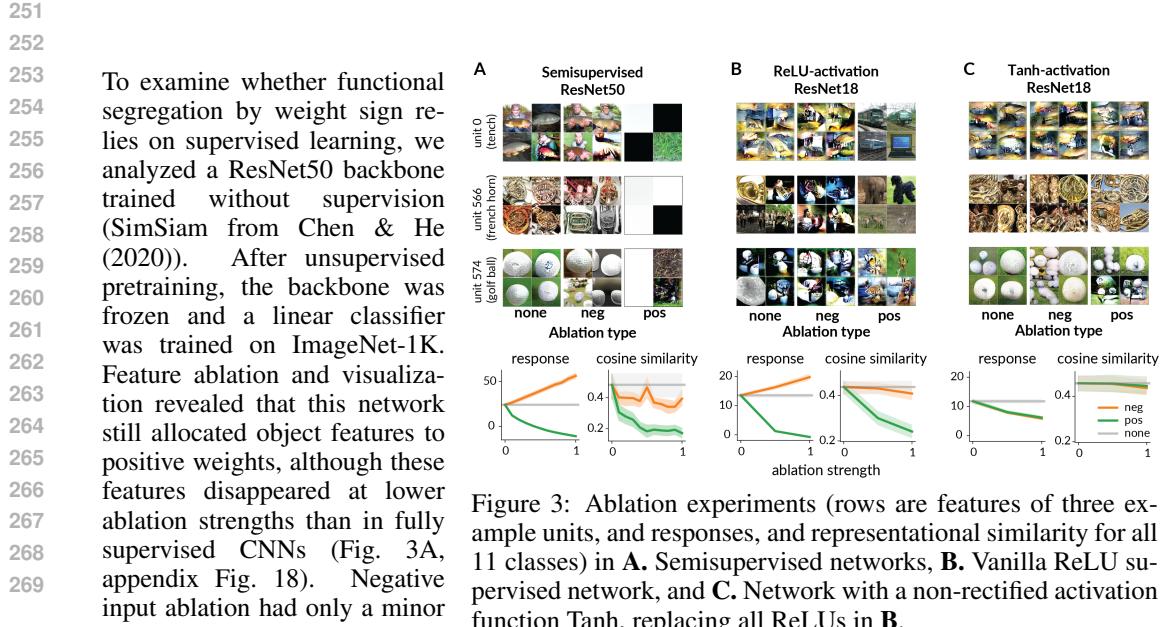
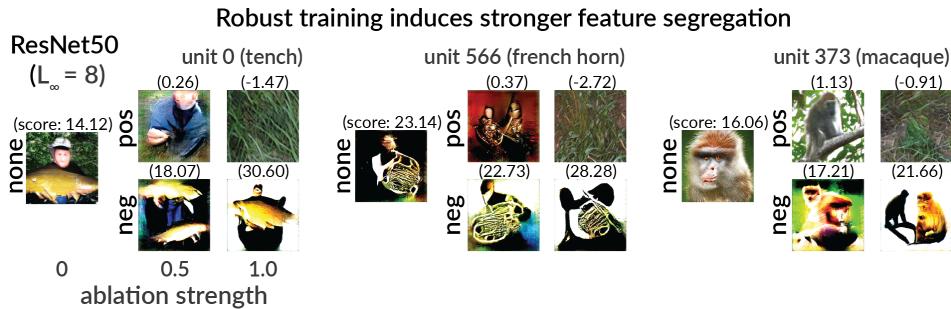


Figure 3: Ablation experiments (rows are features of three example units, and responses, and representational similarity for all 11 classes) in **A.** Semisupervised networks, **B.** Vanilla ReLU supervised network, and **C.** Network with a non-rectified activation function Tanh, replacing all ReLUs in **B**.

270 effect, suggesting that even
 271 unsupervised representations
 272 are organized so that positive weights convey most object-related features.
 273

274 We hypothesized that this segregation depends on the rectifying ReLU activation, which forces unit
 275 outputs to be non-negative and makes weight signs determine the direction of each contribution.
 276 Prior work in toy networks shows that ReLU promotes alignment to the input space, whereas Tanh
 277 promotes alignment to the output space (Alleman et al., 2023). To test the role of rectification, we
 278 trained ResNet18 models with either ReLU or Tanh. As expected, the ReLU model showed clear
 279 segregation, with the strongest disruptions arising from ablation of positive weights. In contrast,
 280 the Tanh model showed similar representational changes across ablation types and preserved key
 281 features even when either sign was removed (Fig. 3B,C). Thus, rectified activations are required for
 282 strong segregation of object information into positive weights in CNNs.
 283



293 Figure 4: Robust network ResNet50 $L_\infty = 8$ shows a large change in preferred features upon input
 294 ablation. Notice the white background in the negative-weight ablation condition.
 295

296 4.4 ADVERSARILY ROBUST NETWORKS SHOW ENHANCED FEATURE SEGREGATION 297

298 Having established that both unsupervised pretraining and rectified activations support the segre-
 299 gation of object and context information by weight sign, we next asked how this organization is
 300 affected by other salient properties of deep vision networks. In particular, adversarially robust
 301 networks, which are trained to resist small targeted image perturbations (Szegedy et al., 2014; Madry
 302 et al., 2019), are believed to better reflect aspects of biological vision and may therefore show distinc-
 303 tive patterns of feature segregation. We examined whether and how adversarial robustness influences
 304 the allocation of object and contextual information to positive and negative weights.
 305

306 In robust ResNet50 networks, intact feature visualizations appeared more object-like, and ablation
 307 of negative input weights reliably altered the background color, often rendering it white (Fig. 4).
 308 This hinted at a stronger feature segregation than in vanilla networks. Quantitative analysis con-
 309 firmed that as network robustness to adversarial attacks increased, so did the model’s vulnerability
 310 to ablation, as measured by the difference in cosine similarity between control images and ablated
 311 images (see $\Delta(\cosine \text{ similarity})$ in Fig. 5). For ablation strength of 1 (yellow/light lines), the
 312 difference increased with network robustness, and slopes in Table 1 indicate this trend holds across
 313 ablation polarities and strengths. Moreover, the robustness effects translated to higher shape bias
 314 in the benchmark by Geirhos et al. (2022), with negative ablations affecting more the texture than
 315 the shape accuracy (Fig. 16). Overall, classification CNNs segregate object information to positive
 316 weights and texture or background information to negative weights, and that adversarially robust
 317 training further sharpens this sign-based segregation.
 318

319 4.5 FEATURE SEGREGATION IS NOT EXCLUSIVE TO CLASS UNITS 320

321 To determine whether feature segregation by weight sign is present beyond the output layer, we
 322 systematically analyzed intermediate and early layers of CNNs. We first searched for channels that
 323 approximated Dale’s law, identifying those that predominantly provided positive or negative inputs
 324 to their downstream units in each layer. For each convolutional layer, we calculated the sign con-
 325 sistency of outgoing weights and ranked channels according to whether they sent mostly positive or
 326 mostly negative signals forward. We then visualized the preferred features of these channels using
 327

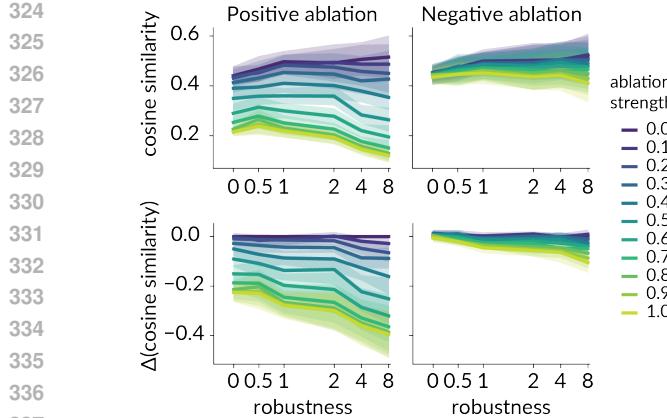


Figure 5: Representative changes under input ablation increase with robust training in ResNet50. Top: cosine similarities to control. Bottom: changes relative to control. Most ablation strengths show significant correlations with robustness.

Table 1: Spearman correlation of representational change upon ablation vs robustness (L_∞ norm)

type	Positive ρ (pvalue)	Negative ρ (pvalue)
α		
0.1	-0.17 (2e-1)	-0.10 (4e-1)
0.2	-0.39 (1e-3)	-0.21 (8e-2)
0.3	-0.34 (4e-3)	-0.14 (3e-1)
0.4	-0.38 (1e-3)	-0.34 (5e-3)
0.5	-0.47 (6e-5)	-0.46 (9e-5)
0.6	-0.48 (3e-5)	-0.34 (5e-3)
0.7	-0.51 (9e-6)	-0.52 (6e-6)
0.8	-0.50 (2e-5)	-0.49 (2e-5)
0.9	-0.48 (4e-5)	-0.62 (2e-8)
1.0	-0.47 (6e-5)	-0.57 (5e-7)

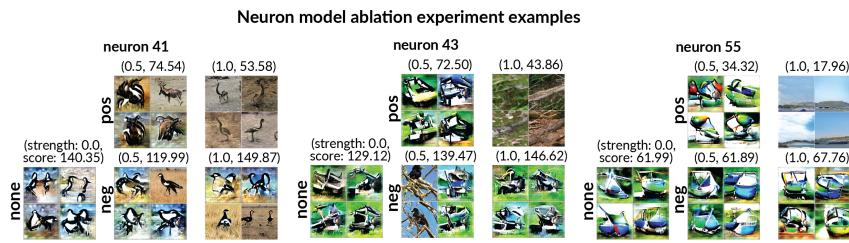


Figure 6: Preferred features of neuronal network models of visual neurons in the primate ventral stream. Pos: are positive ablations, neg are negative ablations, number indicates ablation strength. Shown are top 4 visualizations at 0, 0.5 and 1.0 ablation strengths.

gradient-based methods with the Lucent library in PyTorch. Examining all five convolutional layers of AlexNet, we found that feature segregation by sign emerged throughout the network. In the first layer, channels with mostly positive weights responded to high-frequency achromatic edges, while those with mostly negative weights responded to lower-frequency, colored edges and spots. In the middle layers, positive channels tended to emphasize edges and detailed textures, whereas negative channels often represented broader, colored, or background-like patterns. By the final convolutional layer, channels with mostly negative weights produced features that resembled background elements such as sky or grass, while channels with positive weights highlighted sharp, localized object fragments like animal snouts and eyes (appendix Fig. 19). Altogether, our results show that feature segregation by weight sign is not restricted to the output layer, but gradually develops throughout the network. This pattern is reminiscent of Dale’s law in biological circuits, suggesting that artificial neural networks can develop sign-consistent and functionally distinct representations across all layers, even in the absence of a biological constraint.

4.6 BIOLOGICAL MODELS BASED ON IMAGENET NETWORKS SEGREGATE LOCAL FEATURE INFORMATION INTO POSITIVE WEIGHTS

The ventral stream in primates is responsible for object recognition, and artificial networks are the best models of its function. We therefore wondered if the segregation of positive and negative inputs we observed in networks might also occur in the brain. However, it is not currently possible to selectively remove positive or negative synaptic inputs from real neurons the way we can in artificial networks. To address this limitation, we fit linear models mapping CNN features to macaque ventral stream neural responses, and applied feature visualization to both model units and in vivo data. This allowed us to test feature segregation in biological representations and to generate testable neuro-

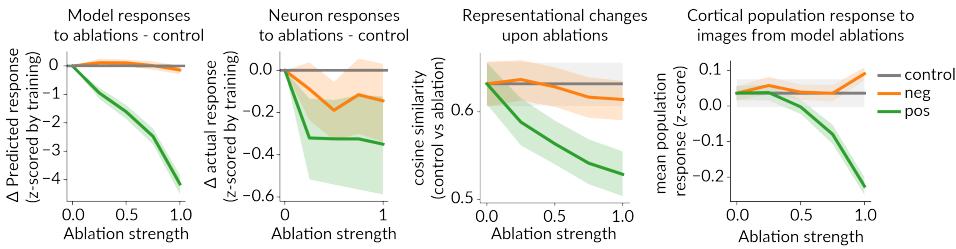


Figure 7: Left: predicted and actual neuron responses of model to ablations. Images obtained from positive ablations in the neuron models elicited a consistent activity drop on the biological neurons modeled. Right: Representational change of model to ablations measured by our cosine similarity metric on the neuron model feature visualizations upon ablation; and cortical population response to the images obtained from feature visualization from ablation of model units, neurons were z-scored before computing the population average. Plots show averages over 59 models, (35 for monkey C, and 24 for monkey D), shaded regions are the 95% C.I. of the mean. For all plots the positive ablation condition was statistically different to the control.

science predictions. We recorded neural activity from V1, V4, and IT cortex in two monkeys, using a diverse set of images and modeled each neuron’s response using partial least squares regression with activations from the penultimate layer of AlexNet (4096 units).¹ We then applied our ablation and visualization protocol to these neuron models (see appendix for validation A.4) and showed the resulting images to the monkey during the same session. First, images from intact models reliably drove biological neurons to firing rates more than one standard deviation above those observed during natural image presentations (Fig. 21, left), indicating out-of-distribution generalization. For the subset of neurons in which we performed *in vivo* closed-loop feature visualization, we found that the model’s preferred features often matched those of the neuron, providing additional validation (appendix Fig. 21, right). However, *in vivo* features were more spatially localized (procedure in A.1), whereas *in silico* features exhibited greater spatial variation (rotated, mirrored, or repeated versions). This likely reflects invariances, due to using a fully connected layer, that are not present in our recorded neurons. Moreover, unlike the images from recognition units, images from the neuron models did not resemble objects (appendix Fig. 21, Fig. 6).

Ablation experiments reveal sign-based segregation in neuron models. Ablation experiments on these neuron models showed that removing positive input weights led to a significant decrease in both predicted and observed firing rates, while removing negative weights had a smaller effect (Fig. 7). This pattern was consistent across individual neurons and at the population level in the ventral stream (Fig. 7, rightmost). The population changes suggest that sign-based functional segregation in model predictions translates to measurable changes across the ventral stream population and perhaps perception.

Dale’s law inspired analysis. To bring our models closer to Dale’s law, we applied two approaches. First, we tested whether neural responses could be predicted using only positive input weights, corresponding to receiving input exclusively from excitatory artificial neurons. This constraint reduced both training and test accuracy relative to unconstrained models (appendix Fig. 22), indicating that neuron models require both positive and negative inputs. Second, we identified Dale-like artificial units that contributed mainly positive or negative weights to all output neuron models, defining putative excitatory and inhibitory inputs. Positive-weight units corresponded to smaller-scale edges and localized spots, cleanly separated from the background, whereas negative-weight units aligned with broader textures and larger patches (appendix Fig. 23). This pattern supports the idea that inhibitory-like artificial inputs preferentially encode contextual or background structure, paralleling inhibitory modulation in biological cortex.

Experimental manipulation of background as a test for inhibition. To further test this hypothesis, we experimentally manipulated image backgrounds *in vivo*. In a subset of experiments, we presented altered images in which the background was cleared around the neuron’s preferred feature,

¹Although using other layers could improve predictive accuracy, we selected the penultimate layer to directly test if inputs optimized for classification maintain weight sign-based segregation in biological neural predictions.

432 thereby reducing the putative inhibitory drive. As predicted, this manipulation resulted in increased
 433 neuronal responses (appendix Fig. 24), providing functional support for the idea that inhibitory or
 434 negative inputs are involved in contextual modulation and that their reduction can enhance feature
 435 selectivity in high-level visual cortex. Together, these results suggest that functional segregation by
 436 input sign extends to models of ventral stream neurons, providing concrete testable predictions for
 437 future experiments targeting excitation and inhibition in visual cortex.

440 5 LIMITATIONS

441
 442 Our results hold in the last layer units of multiple networks. Due to limited computing time, we did
 443 not test all 1000 categories in as many networks as possible, our largest test consisted of 100 units.
 444 While larger scale simulations will provide exhaustive evidence, we are confident our main claims
 445 will stand. We limited our neuron recordings to a 160 image dataset for regressing neuron responses
 446 via CNNs. While we observed good fits and recovered relevant feature to the neurons, more images
 447 may improve the models, especially using larger-scale versions of our diverseSet. The neuroscience
 448 results would need to follow Dale’s law to be mapped one-to-one to excitatory and inhibitory neu-
 449 rons, but we make no claim to a perfect mapping in this work. The fundamental question of shape
 450 vs texture, foreground vs background remains to be solved. We reconciled changes in frequency
 451 structure with changes in objectness and visual representations and LPIPS image similarity. How-
 452 ever, more work remains to understand the full extent of the segregation reported here. Solving this
 453 problem for the visual cortex may provide better benchmarks for this task in AI.

454 6 DISCUSSION

455 Our study combined ablations with feature visualization guided by naturalistic image priors to re-
 456 veal the functional segregation of class-level features in the output layer of ImageNet trained CNNs:
 457 positive weights contribute object/shape/low frequency information, while negative weights con-
 458 tribute background/contextual/texture information. This effect was enhanced in robust networks, it
 459 was present in networks with unsupervised pretraining, but was absent in network trained with Tanh
 460 instead of ReLU. Our results explain how the background contribution to classification observed in
 461 (Xiao et al., 2020) emerges, backgrounds are primarily encoded by the negative inputs.

462 Importantly for neuroscience, the observed functional segregation in neuron model units in CNNs
 463 hints at a functional segregation in the brain beyond the center-surround classically studied in V1.
 464 And we crafted a diverse dataset for visual neuroscience recordings that is scalable. Neuron re-
 465 sponds to a smaller but diverse set of naturalistic, colored images, with complex foregrounds and
 466 backgrounds, led to models capturing relevant features obtained experimentally from the neuron.
 467 Thus, using both model-based and model-free approaches revealed richer neuronal representations.
 468 Preferred images from neuron models with positive input ablations elicited smaller average pop-
 469 ulation responses of cortical neurons. This suggests that ablation in networks modeling neurons
 470 holds potential as a method to control the population activity in the brain. To relate ablation-induced
 471 changes in the images to the population responses is a future direction. This ablation based on the
 472 natural division of positive and negative weights can be easily extended into arbitrary layers, e.g.,
 473 using gradients to define positive and negative contributions to any arbitrary unit. And our ablation
 474 approach proposes baselines for the functional differences between excitatory and inhibitory neu-
 475 rons in higher cortical visual areas. The functional segregation has consequences for neural coding
 476 and response selectivity. Our findings generate concrete predictions for future experiments using
 477 advanced genetic or optogenetic tools to dissect excitation and inhibition in primate cortex. Under-
 478 standing the circuit mechanism of biological vision could aid further understanding and development
 479 of computer vision models. Interpretability is thus an important field for both AI and neuroscience.

480 REFERENCES

481
 482
 483 fastai/imagenette, May 2024. URL <https://github.com/fastai/imagenette>. original-
 484 date: 2019-03-06T01:58:45Z.

486 Matteo Alleman, Jack Lindsey, and Stefano Fusi. Task structure and nonlinearity jointly deter-
 487 mine learned representational geometry. October 2023. URL [https://openreview.net/](https://openreview.net/forum?id=k9t8dQ30kU)
 488 forum?id=k9t8dQ30kU.

489 Emily J. Allen, Ghislain St-Yves, Yihan Wu, Jesse L. Breedlove, Jacob S. Prince, Logan T. Dow-
 490 dle, Matthias Nau, Brad Caron, Franco Pestilli, Ian Charest, J. Benjamin Hutchinson, Thomas
 491 Naselaris, and Kendrick Kay. A massive 7T fMRI dataset to bridge cognitive neuroscience
 492 and artificial intelligence. *Nature Neuroscience*, 25(1):116–126, January 2022. ISSN 1546-
 493 1726. doi: 10.1038/s41593-021-00962-x. URL [https://www.nature.com/articles/](https://www.nature.com/articles/s41593-021-00962-x)
 494 s41593-021-00962-x. Publisher: Nature Publishing Group.

495 Pouya Bashivan, Kohitij Kar, and James J. DiCarlo. Neural population control via deep image
 496 synthesis. *Science*, 364(6439):eaav9436, May 2019. doi: 10.1126/science.aav9436. URL
 497 <https://www.science.org/doi/full/10.1126/science.aav9436>. Publisher:
 498 American Association for the Advancement of Science.

499 Timothy F. Brady, Talia Konkle, George A. Alvarez, and Aude Oliva. Visual long-term mem-
 500 ory has a massive storage capacity for object details. *Proceedings of the National Academy
 501 of Sciences*, 105(38):14325–14329, September 2008. doi: 10.1073/pnas.0803390105. URL
 502 <https://www.pnas.org/doi/full/10.1073/pnas.0803390105>. Publisher: Pro-
 503 ceedings of the National Academy of Sciences.

504 Andrew Brock, Jeff Donahue, and Karen Simonyan. Large Scale GAN Training for High Fidelity
 505 Natural Image Synthesis, February 2019. URL <http://arxiv.org/abs/1809.11096>.
 506 arXiv:1809.11096 [cs, stat].

507 Xinlei Chen and Kaiming He. Exploring Simple Siamese Representation Learning, November 2020.
 508 URL <http://arxiv.org/abs/2011.10566>. arXiv:2011.10566.

509 Hoagy Cunningham, Aidan Ewart, Logan Riggs, Robert Huben, and Lee Sharkey. Sparse Autoen-
 510 coders Find Highly Interpretable Features in Language Models, October 2023. URL <https://arxiv.org/abs/2309.08600>. arXiv:2309.08600 [cs].

511 Henry Dale. Pharmacology and Nerve-endings (Walter Ernest Dixon Memorial Lecture). *Proceed-
 512 ings of the Royal Society of Medicine*, 28(3):319–332, January 1935. ISSN 0035-9157. doi:
 513 10.1177/003591573502800330. URL [https://pmc.ncbi.nlm.nih.gov/articles/](https://pmc.ncbi.nlm.nih.gov/articles/PMC2205701/)
 514 PMC2205701/.

515 James J. DiCarlo, Davide Zoccolan, and Nicole C. Rust. How Does the Brain Solve Visual Ob-
 516 ject Recognition? *Neuron*, 73(3):415–434, February 2012. ISSN 0896-6273. doi: 10.1016/j.
 517 neuron.2012.01.010. URL <https://www.sciencedirect.com/science/article/pii/S089662731200092X>.

518 Alexey Dosovitskiy and Thomas Brox. Generating Images with Perceptual Similarity Metrics
 519 based on Deep Networks, February 2016. URL <http://arxiv.org/abs/1602.02644>.
 520 arXiv:1602.02644 [cs].

521 Gamaleldin Elsayed, Shreya Shankar, Brian Cheung, Nicolas Papernot, Alexey Kurakin, Ian Good-
 522 fellow, and Jascha Sohl-Dickstein. Adversarial Examples that Fool both Computer Vision and
 523 Time-Limited Humans. In *Advances in Neural Information Processing Systems*, volume 31. Cur-
 524 ran Associates, Inc., 2018. URL <https://proceedings.neurips.cc/paper/2018/hash/8562ae5e286544710b2e7ebe9858833b-Abstract.html>.

525 Dumitru Erhan, Yoshua Bengio, Aaron Courville, and Pascal Vincent. Vi-
 526 sualizing higher-layer features of a deep network. *University of Montreal*,
 527 1341(3):1, 2009. URL https://www.researchgate.net/profile/Aaron-Courville/publication/265022827_Visualizing_Higher-Layer_Features_of_a_Deep_Network/links/53ff82b00cf24c81027da530/Visualizing-Higher-Layer-Features-of-a-Deep-Network.pdf.

528 Jenelle Feather, Guillaume Leclerc, Aleksander Madry, and Josh H. McDermott. Model
 529 metamers reveal divergent invariances between biological and artificial neural net-
 530 works. *Nature Neuroscience*, 26(11):2017–2034, November 2023. ISSN 1546-1726.

doi: 10.1038/s41593-023-01442-0. URL <https://www.nature.com/articles/s41593-023-01442-0>. Publisher: Nature Publishing Group.

Robert Geirhos, Patricia Rubisch, Claudio Michaelis, Matthias Bethge, Felix A. Wichmann, and Wieland Brendel. ImageNet-trained CNNs are biased towards texture; increasing shape bias improves accuracy and robustness, November 2022. URL <http://arxiv.org/abs/1811.12231>. arXiv:1811.12231 [cs].

Umut Güçlü and Marcel A. J. van Gerven. Increasingly complex representations of natural movies across the dorsal stream are shared between subjects. *NeuroImage*, 145:329–336, January 2017. ISSN 1053-8119. doi: 10.1016/j.neuroimage.2015.12.036. URL <https://www.sciencedirect.com/science/article/pii/S1053811915011490>.

Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition, December 2015. URL <https://arxiv.org/abs/1512.03385v1>.

D. H. Hubel and T. N. Wiesel. Receptive fields of single neurones in the cat's striate cortex. *The Journal of Physiology*, 148(3):574–591, October 1959. ISSN 0022-3751. URL <https://www.ncbi.nlm.nih.gov/pmc/articles/PMC1363130/>.

Chou P. Hung, Gabriel Kreiman, Tomaso Poggio, and James J. DiCarlo. Fast Readout of Object Identity from Macaque Inferior Temporal Cortex. *Science*, 310(5749):863–866, November 2005. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1117593. URL <https://www.science.org/doi/10.1126/science.1117593>.

Jaewon Hwang, Andrew R. Mitz, and Elisabeth A. Murray. NIMH MonkeyLogic: Behavioral control and data acquisition in MATLAB. *Journal of Neuroscience Methods*, 323:13–21, July 2019. ISSN 1872-678X. doi: 10.1016/j.jneumeth.2019.05.002.

Jeffry S. Isaacson and Massimo Scanziani. How Inhibition Shapes Cortical Activity. *Neuron*, 72(2):231–243, October 2011. ISSN 0896-6273. doi: 10.1016/j.neuron.2011.09.027. URL [https://www.cell.com/neuron/abstract/S0896-6273\(11\)00879-8](https://www.cell.com/neuron/abstract/S0896-6273(11)00879-8). Publisher: Elsevier.

Kohitij Kar, Jonas Kubilius, Kailyn Schmidt, Elias B. Issa, and James J. DiCarlo. Evidence that recurrent circuits are critical to the ventral stream's execution of core object recognition behavior. *Nature Neuroscience*, 22(6):974–983, June 2019. ISSN 1546-1726. doi: 10.1038/s41593-019-0392-5. URL <https://www.nature.com/articles/s41593-019-0392-5>. Number: 6 Publisher: Nature Publishing Group.

Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems*, volume 25. Curran Associates, Inc., 2012. URL <https://papers.nips.cc/paper/2012/hash/c399862d3b9d6b76c8436e924a68c45b-Abstract.html>.

Guillaume Leclerc, Andrew Ilyas, Logan Engstrom, Sung Min Park, Hadi Salman, and Aleksander Madry. FFCV: Accelerating Training by Removing Data Bottlenecks, June 2023. URL <http://arxiv.org/abs/2306.12517>. arXiv:2306.12517.

Tianqin Li, Ziqi Wen, Yangfan Li, and Tai Sing Lee. Emergence of Shape Bias in Convolutional Neural Networks through Activation Sparsity. *Advances in Neural Information Processing Systems*, 36:71755–71766, December 2023. URL https://proceedings.neurips.cc/paper_files/paper/2023/hash/e31c16c7b3e0ccee5159ae5443154fac-Abstract-Conference.html.

Ilya Loshchilov. LM-CMA: an Alternative to L-BFGS for Large Scale Black-box Optimization, November 2015. URL <http://arxiv.org/abs/1511.00221>. arXiv:1511.00221 [cs, math].

Aleksander Madry, Aleksandar Makelov, Ludwig Schmidt, Dimitris Tsipras, and Adrian Vladu. Towards Deep Learning Models Resistant to Adversarial Attacks, September 2019. URL <http://arxiv.org/abs/1706.06083>. arXiv:1706.06083.

594 Anh Nguyen, Alexey Dosovitskiy, Jason Yosinski, Thomas Brox, and Jeff Clune. Synthesizing the
 595 preferred inputs for neurons in neural networks via deep generator networks, November 2016a.
 596 URL <http://arxiv.org/abs/1605.09304>. arXiv:1605.09304 [cs].
 597

598 Anh Nguyen, Jason Yosinski, and Jeff Clune. Multifaceted Feature Visualization: Uncovering the
 599 Different Types of Features Learned By Each Neuron in Deep Neural Networks, February 2016b.
 600 URL <https://arxiv.org/abs/1602.03616v2>.

601 Chris Olah, Alexander Mordvintsev, and Ludwig Schubert. Feature Visualization. *Distill*, 2(11):e7,
 602 November 2017. ISSN 2476-0757. doi: 10.23915/distill.00007. URL <https://distill.pub/2017/feature-visualization>.

603

604 Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter.
 605 Zoom In: An Introduction to Circuits. *Distill*, 5(3):e00024.001, March 2020. ISSN 2476-
 606 0757. doi: 10.23915/distill.00024.001. URL <https://distill.pub/2020/circuits/zoom-in>.
 607

608 Bruno A. Olshausen and David J. Field. Emergence of simple-cell receptive field properties by
 609 learning a sparse code for natural images. *Nature*, 381(6583):607–609, June 1996. ISSN 1476-
 610 4687. doi: 10.1038/381607a0. URL <https://www.nature.com/articles/381607a0>.
 611 Publisher: Nature Publishing Group.

612

613 Mateusz Pach, Shyamgopal Karthik, Quentin Bouliot, Serge Belongie, and Zeynep Akata. Sparse
 614 Autoencoders Learn Monosemantic Features in Vision-Language Models, June 2025. URL
 615 <http://arxiv.org/abs/2504.02821>. arXiv:2504.02821 [cs].
 616

617 Carlos R. Ponce, Will Xiao, Peter F. Schade, Till S. Hartmann, Gabriel Kreiman, and Margaret S.
 618 Livingstone. Evolving Images for Visual Neurons Using a Deep Generative Network Reveals
 619 Coding Principles and Neuronal Preferences. *Cell*, 177(4):999–1009.e10, May 2019. ISSN 0092-
 620 8674, 1097-4172. doi: 10.1016/j.cell.2019.04.005. URL [https://www.cell.com/cell/abstract/S0092-8674\(19\)30391-5](https://www.cell.com/cell/abstract/S0092-8674(19)30391-5). Publisher: Elsevier.

621

622 Hadi Salman, Andrew Ilyas, Logan Engstrom, Ashish Kapoor, and Aleksander Madry. Do Adver-
 623 sarially Robust ImageNet Models Transfer Better?, December 2020. URL <http://arxiv.org/abs/2007.08489>. arXiv:2007.08489 [cs, stat].
 624

625 Karen Simonyan and Andrew Zisserman. Very Deep Convolutional Networks for Large-Scale Image
 626 Recognition, April 2015. URL <http://arxiv.org/abs/1409.1556>. arXiv:1409.1556
 627 [cs].
 628

629 Christian Szegedy, Wojciech Zaremba, Ilya Sutskever, Joan Bruna, Dumitru Erhan, Ian Good-
 630 fellow, and Rob Fergus. Intriguing properties of neural networks, February 2014. URL
 631 <http://arxiv.org/abs/1312.6199>. arXiv:1312.6199 [cs].
 632

633 Hiroshi Tamura, Hidekazu Kaneko, Keisuke Kawasaki, and Ichiro Fujita. Presumed Inhibitory
 634 Neurons in the Macaque Inferior Temporal Cortex: Visual Response Properties and Functional
 635 Interactions With Adjacent Neurons. *Journal of Neurophysiology*, 91(6):2782–2796, June 2004.
 636 ISSN 0022-3077. doi: 10.1152/jn.01267.2003. URL <https://journals.physiology.org/doi/full/10.1152/jn.01267.2003>. Publisher: American Physiological Society.

637

638 Edgar Y. Walker, Fabian H. Sinz, Erick Cobos, Taliah Muhammad, Emmanouil Froudarakis, Paul G.
 639 Fahey, Alexander S. Ecker, Jacob Reimer, Xaq Pitkow, and Andreas S. Tolias. Inception loops
 640 discover what excites neurons most using deep predictive models. *Nature Neuroscience*, 22(12):
 641 2060–2065, December 2019. ISSN 1546-1726. doi: 10.1038/s41593-019-0517-x. URL <https://www.nature.com/articles/s41593-019-0517-x>. Publisher: Nature Publishing
 642 Group.
 643

644 Binxu Wang and Carlos R. Ponce. High-performance evolutionary algorithms for online neuron
 645 control. In *Proceedings of the Genetic and Evolutionary Computation Conference*, GECCO '22,
 646 pp. 1308–1316, New York, NY, USA, July 2022. Association for Computing Machinery. ISBN
 647 978-1-4503-9237-2. doi: 10.1145/3512290.3528725. URL <https://doi.acm.org/doi/10.1145/3512290.3528725>.

648 Chien-Yao Wang, Alexey Bochkovskiy, and Hong-Yuan Mark Liao. YOLOv7: Trainable bag-
 649 of-freebies sets new state-of-the-art for real-time object detectors, July 2022. URL <https://arxiv.org/abs/2207.02696v1>.
 650

651

652 Kai Xiao, Logan Engstrom, Andrew Ilyas, and Aleksander Madry. Noise or Signal: The Role
 653 of Image Backgrounds in Object Recognition, June 2020. URL <http://arxiv.org/abs/2006.09994>. arXiv:2006.09994 [cs].
 654

655

656 Will Xiao and Gabriel Kreiman. XDream: Finding preferred stimuli for visual neu-
 657 rons using generative networks and gradient-free optimization. *PLOS Computational
 658 Biology*, 16(6):e1007973, June 2020. ISSN 1553-7358. doi: 10.1371/journal.pcbi.
 659 1007973. URL <https://journals.plos.org/ploscompbiol/article?id=10.1371/journal.pcbi.1007973>. Publisher: Public Library of Science.
 660

661

662 I. Zeki Yalniz, Hervé Jégou, Kan Chen, Manohar Paluri, and Dhruv Mahajan. Billion-scale semi-
 663 supervised learning for image classification, May 2019. URL <http://arxiv.org/abs/1905.00546>. arXiv:1905.00546 [cs].
 664

665

666 Daniel L. K. Yamins, Ha Hong, Charles F. Cadieu, Ethan A. Solomon, Darren Seibert, and James J.
 667 DiCarlo. Performance-optimized hierarchical models predict neural responses in higher visual
 668 cortex. *Proceedings of the National Academy of Sciences*, 111(23):8619–8624, June 2014.
 669 doi: 10.1073/pnas.1403112111. URL <https://www.pnas.org/doi/10.1073/pnas.1403112111>. Publisher: Proceedings of the National Academy of Sciences.
 670

671

672 Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang. The Unreasonable
 673 Effectiveness of Deep Features as a Perceptual Metric, January 2018. URL <https://arxiv.org/abs/1801.03924v2>.
 674

675

676 A APPENDIX

677 A.1 EXTENDED METHODS

678 **Networks** The ablation studies were performed on CNNs pretrained on the ImageNet dataset:
 679 AlexNet (Krizhevsky et al., 2012), VGG16 (Simonyan & Zisserman, 2015), ResNet50 (He et al.,
 680 2015), and robustly-trained ResNet50 ($L_\infty \in \{0.5, 1, 2, 4, 8\}$, Salman et al. (2020)). All these
 681 networks end on a 1000-unit fully connected layer, each unit corresponding to one of the 1000
 682 ImageNet categories. Neural networks were used in Pytorch.
 683

684 **ImageNet subsampling** To reduce computing time, for most of the experiments, we used a subset
 685 of ImageNet, the *imagenette* dataset (noa, 2024) and the macaque category, 11 classes in total. These
 686 classes and their corresponding output units in each network trained on the 1000-class ImageNet
 687 dataset are as follows: (0, tench), (207, English Springer), (482, cassette player), (491, chain saw),
 688 (566, church), (569, French horn), (571, garbage truck), (574, gas pump), (701, golf ball), (970,
 689 parachute), and (373, macaque). We visualized the representations of the output layer units of
 690 those classes under different ablation conditions. For Fig. 12, to sample 100 diverse classes out
 691 of the 1000 ImageNet classes, the 50k validation images were first clustered into 100 clusters via
 692 agglomerative clustering of the L2 distance matrix from the 1000-d output features of ResNet50,
 693 which was pre-trained on ImageNet. Then, one new unique class is selected from each cluster.
 694

695 **Ablation** We used two ablation conditions: we ablated weights that were (1) only positive or (2)
 696 only negative. We ablated weights cumulatively by first sorting the positive (or negative) weights
 697 by their (absolute) decreasing value. We defined the *ablation strength*, α , as a fraction of the total
 698 positive or total negative weights to a unit. We identified the top k weights necessary to reach the
 699 silencing strength, i.e., $\sum_{i=1}^k w_i \leq \alpha$, and set them to zero. We covered the range of ablations from
 700 0 to 1. For most experiments with ANNs, we used silencing strengths in 0.1 steps, from 0 (intact) to
 701 1 (complete ablation).

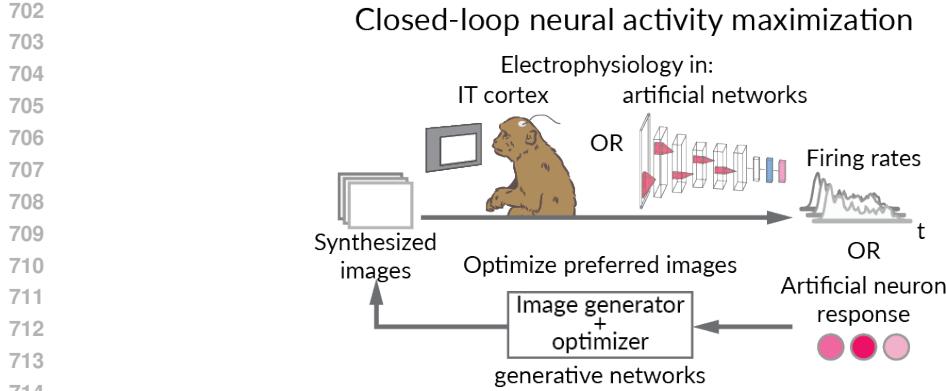


Figure 8: Schematic of feature visualization workflow in ANNs and brains. Optimizer is CMAES, image generators are DeePSim fc6 or BigGAN.

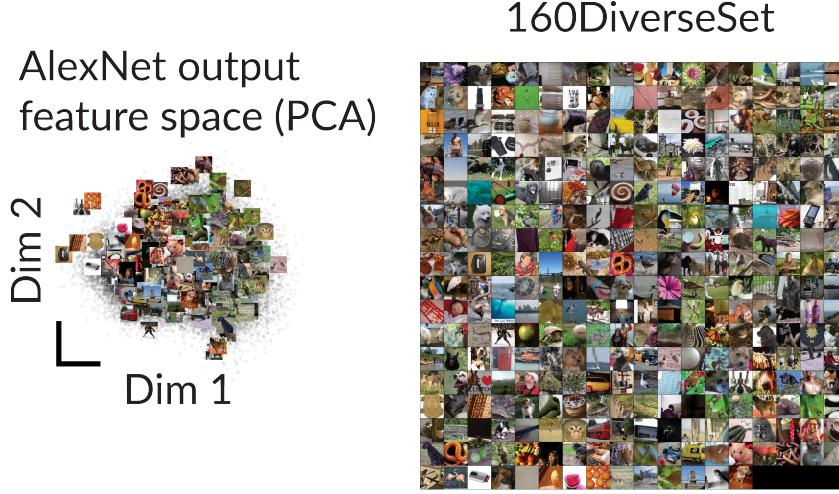


Figure 9: Illustration of a diverse dataset construction using AlexNet output feature space. The embedding is the output of the last layer before softmax of AlexNet, a vector space of 1000-dimensions. Left: PCA showing the coverage of the feature space by the diverseSet 160, only for illustration purposes. Right: images from diverseSet 160 used to fit neuron models.

Feature visualization For each ablation condition, we performed feature visualization by optimizing a GAN latent code to create an activity-maximizing image Fig. 8. We used this closed-loop, zeroth-order-search approach to allow comparison with our neuronal experiments, where gradient ascent would not be possible. To increase the span of the stimulus space, we used two GANs: AlexNet fc6 DeePSiM (Dosovitskiy & Brox, 2016) and BigGAN (Brock et al., 2019). For optimization, we used a variant of *covariance matrix adaptation evolutionary strategy* or CMAES (Wang & Ponce, 2022; Loshchilov, 2015). Initial conditions for the CMAES were given as standard deviation of 3.0 for DeePSim, and 0.2 for BigGAN. Initial images for the algorithm were small norm vectors for both GANs, close to the origin of the latent spaces. For BigGAN, we generated a fixed noise vector by scaling a 128-dimensional truncated noise sample (-1.4, 1.4), and concatenated it with a 128-dimensional zero vector of the class embedding, to form the required 256-dimensional input code. The remaining parameters are determined by the dimensionality of the search space of each GAN. We optimized ten images per GAN, resulting in 20 feature visualizations per output unit and ablation condition. Diverse visualizations better capture the multifaceted high-level representations in CNNs (Nguyen et al., 2016b). For our examples, we show the best of the 20 visualizations, but used all for quantitative analyses. For visualizations of neural networks predicting biological neuron responses, due to experimental time restrictions, we used five visualizations per ablation condition, via DeePSim only. Our experiments are performed in a PC with Nvidia 4090 GPU, and each visu-

756 alization running 100 iterations takes about 3 mins. For *in vivo* experiments, we ran from 20 to 60
 757 iterations of the AlexNet fc6 DeePSiM with the CMAES algorithm implemented in Matlab, linked
 758 to our real-time spike-sorting data acquisition. The responses fed to the CMAES algorithm were the
 759 average firing rate on the window 70-170 ms from image onset.
 760

761 **Feature analysis** We computed image similarity using an ensemble of CNNs, including AlexNet,
 762 ResNet50, and ResNet50 with robustness in $L_\infty \in \{0.5, 1, 2, 4, 8\}$, inspired by (Feather et al.,
 763 2023) And confirmed the results with LPIPS (Zhang et al., 2018) in the appendix. We computed
 764 their activations and defined similarity as the average pairwise cosine similarity (LPIPS) between
 765 control activity vs input-ablated activity. We averaged the results of the CNNs ensemble, resulting in
 766 one quantity per ablation condition. We computed *objectness* as the maximum bounding box score
 767 provided by YOLOv7 (Wang et al., 2022), this was averaged over visualizations per unit, units per
 768 network, and then across networks.
 769

770 **Visual cortex electrophysiology** We collected data from two animals (monkey C and monkey D),
 771 each implanted chronically with floating multielectrode arrays (Microprobes for Life Sciences, MD)
 772 of 32 or 16 channels (monkey C, N = 96 electrodes, monkey D, 64), in areas V1, V4 and posterior
 773 inferotemporal cortex (PIT). All institutional procedures were followed. Channels were distributed
 774 as (V1, V4, PIT): monkey C (32, 32, 32), monkey D (16, 16, 32). Some electrodes captured the
 775 activity of single units, but most showed multi-unit activity (reflecting the pooled activity of micro-
 776 clusters of neurons). The animals performed a simple fixation task, which required them to keep
 777 their eyes on a 0.25-deg diameter spot at the center of the screen, within a square fixation window
 778 measuring 0.5–1° per side. Images were presented for 100 milliseconds ON, 150-ms off, 4-5 images
 779 per trial, after which the animal received water or juice. Images were presented to monkey C were
 780 2 deg in size, and 4-8 deg for monkey D to match the receptive field centers of most channels in
 781 all cortical areas (V1, V4 and PIT). Image presentation and data acquisition (electrophysiology, eye
 782 tracking) were integrated by the MonkeyLogic2 software (Hwang et al., 2019) and OmniPlex Neu-
 783 ral Recording Data Acquisition Systems (Plexon Inc.), interfaced through custom Matlab code. We
 784 performed online spike sorting using the PlexControl client based on waveforms. We used ViewPixx
 785 EEG monitors (ViewPixx Technologies), at a resolution of 1920x1080 pixels with 120 Hz refresh
 786 rate. Eye tracking used ISCAN cameras (ISCAN Inc.). And reward was delivered using the DARIS
 787 Control Module System (Crist Instruments).
 788

789 **Feature localization *in vivo*** We conducted a perturbation-based localization to identify relevant
 790 image regions from a feature visualization performed *in vivo*, where gradient information from the
 791 animal brain is unavailable. We perturbed a circular region with a 50-pixel diameter within the
 792 256-pixel image by randomly shuffling the pixels inside this circle, effectively disrupting the local
 793 image structure while maintaining local contrast. We selected 30 such regions for perturbation at
 794 random, excluding those that extended beyond the image boundaries. The modified images were
 795 then presented to the monkey. We hypothesized that perturbing regions crucial for driving the neu-
 796 ron response would lead to a decreased firing rate. To assess local image importance, we calculated
 797 the normalized response change: the difference between the firing rate response to the intact image
 798 and the firing rate response to the perturbed image, divided by the firing rate response to the intact
 799 image. A normalized response change of 0.5 indicates the neuron response decreased by half due to
 800 perturbation. To generate the localized response mask, we averaged the circular masks correspond-
 801 ing to each perturbed region, weighted by their response change. This response mask was further
 802 smoothed using a Gaussian kernel with a 30-pixel standard deviation. We defined relevant regions
 803 as those causing a normalized response change of 0.5 or greater. Finally, we applied this mask to
 804 the original feature visualization image to highlight the local features.
 805

806 **Image dataset** We collected a reference image dataset to activate neurons in the monkey along
 807 the hierarchy of V1, V4, and PIT. Because neurons vary in their preferred features, we constructed
 808 a dataset spanning the image space as represented by the neural embedding of ImageNet-trained
 809 AlexNet. The embedding is the output of the last layer before softmax of AlexNet, a vector space
 810 of 1000-dimensions. The images from this dataset also spanned uniformly the 1000-dimensional
 811 output space of a semi-supervised trained network, trained on a billion images, ResNet50SS (Yalniz
 812 et al., 2019). To define this embedding space, we performed PCA on the output activations from
 813 AlexNet to the 50k ImageNet validation images, we kept the top 300 components (accounting for
 814

about 95% of total explained variance). Then we partitioned the space into a defined number of clusters k , according to the desired dataset size, using batched k-means to reduce computational burden. After finding the k cluster centers, we could feed arbitrary images to the network, map them to the PCA space, and then pick the nearest neighbors to the cluster centers from the desired image space. In addition to the ImageNet validation set, we added other common neuroscience datasets (Brady et al., 2008; Kar et al., 2019; Allen et al., 2022; Hung et al., 2005) to form our image space. We selected $k = 160$ images, as a set that was diverse but small enough to be used in every experimental session. We called this image dataset *diverseSet*.

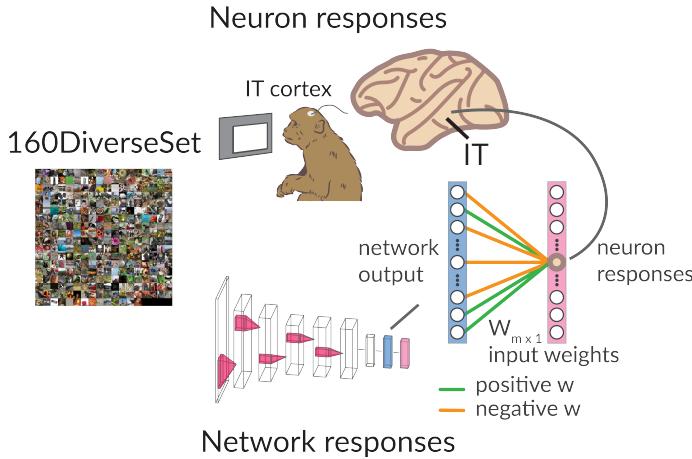


Figure 10: Schematic of model fitting using the dataset *diverseSet*. 160 images were split into train/test datasets (80/20).

Models fit on neuronal activity We recorded responses of neurons in the ventral stream to a 160 image dataset, our *diverseSet* Fig. 10. We relied on a small dataset to fit neuron responses and perform feature visualizations within the same experimental session. We performed partial least-squares linear (PLS) regression (80/20 train/test split) between the neuron responses to images and the activations of the penultimate layer of AlexNet. We used one component for the PLS regression. We selected one neuron or microcluster per experimental session, fitted a model, and performed the ablation and feature visualizations *in silico* for that model. We selected the best fitted neuron per session, based on the r^2 on the 20 % held out test set, usually in the range of 0.15 to 0.5. When time allowed, we also performed the feature visualization of the modeled neuron *in vivo* using a gradient-free approach (Ponce et al., 2019), within the same experimental session. To test whether features learned by the model were relevant to the biological neuron, we recorded the neuronal responses to the preferred images of the model. We then analyzed the representational similarity of the model features under ablations using ANNs. And analyzed the responses of the biological neuron populations from V1, V4 and IT.

A.2 SUPPORTING RESULTS

864
865
866
867
868
869
870
871
872
873
874
875
876
877
878
879
Table 2: Ratio of positive to negative weights. We divided the sum of positive weights by the sum
of the absolute values of the negative weights.

MODEL	RATIO (MEAN \pm STD)
AlexNet	1.03 \pm 0.08
VGG16	1.01 \pm 0.09
ResNet50	1.00 \pm 0.06
ResNet50 ($L_\infty = 0.5$)	1.00 \pm 0.05
ResNet50 ($L_\infty = 1$)	0.99 \pm 0.05
ResNet50 ($L_\infty = 2$)	1.00 \pm 0.04
ResNet50 ($L_\infty = 4$)	1.00 \pm 0.05
ResNet50 ($L_\infty = 8$)	1.01 \pm 0.05

880
881
882
883
884
885
886
887
888
889
890
891
892
893
894
895
896
897
898
899
900
901
902
903
904
905
906
907
908
909
910
911
912
913
914
915
916
917
Table 3: Dale index of the final classification layer before and after training, and trained Top1
accuracy.

Model	Untrained DI	Trained DI	Δ DI	Top1 Acc
alexnet	0.5126	0.5461	+0.0335	56.52
densenet121	0.5125	0.5567	+0.0442	74.43
densenet169	0.5126	0.5590	+0.0464	75.60
densenet201	0.5126	0.5601	+0.0475	76.90
resnet18	0.5128	0.5905	+0.0778	69.76
resnet34	0.5130	0.5923	+0.0793	73.31
resnet50	0.5125	0.6004	+0.0879	80.86
resnet101	0.5126	0.5998	+0.0872	81.89
resnet152	0.5127	0.6007	+0.0880	82.28
vgg11	0.5123	0.5570	+0.0447	69.02
vgg11_bn	0.5123	0.5737	+0.0614	70.37
vgg13	0.5125	0.5583	+0.0457	69.93
vgg13_bn	0.5125	0.5747	+0.0622	71.59
vgg16	0.5126	0.5606	+0.0480	71.59
vgg16_bn	0.5126	0.5761	+0.0635	73.36
vgg19	0.5129	0.5619	+0.0490	72.38
vgg19_bn	0.5129	0.5748	+0.0619	74.22
inception_v3	0.5127	0.5597	+0.0470	77.29

A.3 DALE’S LAW INSPIRED ANALYSIS OF INTERMEDIATE FEATURES

902
903
904
905
906
907
908
To determine if weight segregation of features occurs beyond the output layer, we visualized fea-
ture representations that predominantly provide negative or positive inputs to subsequent layers in
AlexNet. We calculated sign consistency by averaging spatial weights and determining the fre-
quency of positive and negative weights across output channels. The visualization of sign-consistent
input features was conducted using the Lucent library in PyTorch, leveraging gradient-descent chan-
nel activity maximization. We focused on AlexNet’s intermediate layers, examining the top and
bottom sign-consistent features for each input channel.

909
910
Layer Details:

911
912
913
914
915
Conv1: Conv2d(3, 64, kernel_size=(11, 11), stride=(4, 4), padding=(2, 2))
Conv2: Conv2d(64, 192, kernel_size=(5, 5), stride=(1, 1), padding=(2, 2))
Conv3: Conv2d(192, 384, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1))
Conv4: Conv2d(384, 256, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1))
Conv5: Conv2d(256, 256, kernel_size=(3, 3), stride=(1, 1), padding=(1, 1))

916
917
Features that contributed mostly positive weights differed from the features contributing mainly
negative weights, with object vs background arising with increasing depth. This positive vs negative

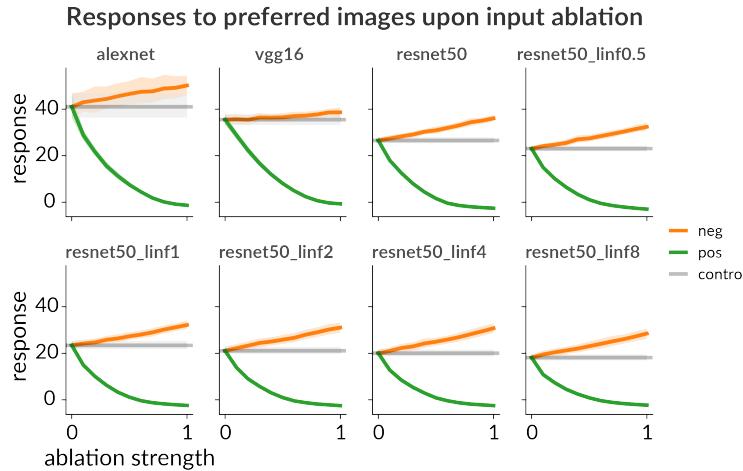


Figure 11: Mean activation scores of units used in ablation experiments. For all networks, units scores come from the last fully-connected layer, with 1000 units, before the softmax. The units correspond to the 10 imagenette categories ([0, 217, 482, 491, 497, 566, 569, 571, 574, 701]) plus the macaque category (373). Error bars are 95% confidence intervals over units (categories tested), where each unit response is the mean of its 20 visualizations. *Control* refers to the feature visualizations in the intact networks for the same units, we extended it as a horizontal line to ease visual comparisons to the different ablation strengths.

weight split is evident even in the first layer, where low-frequency color features are contrasted with high-frequency black-and-white features.

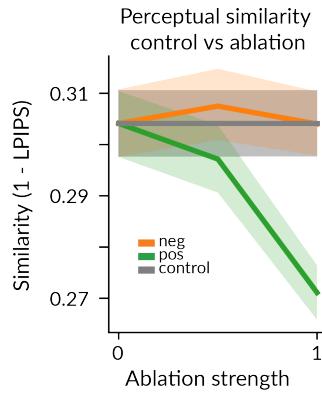


Figure 12: Functional segregation holds in a 10x larger dataset. 100 classes out of the 1000 ImageNet categories were selected by clustering the 50k validation images embedded in the 1000-d output space of ResNet50 picking one class per cluster. Thus, we now have 10x more data points that should span the representational space of the output layer we study. Consistent with the smaller dataset, the main object features degrade into more uniform background images upon positive ablation. Here we show examples from 10 of the 100 classes.

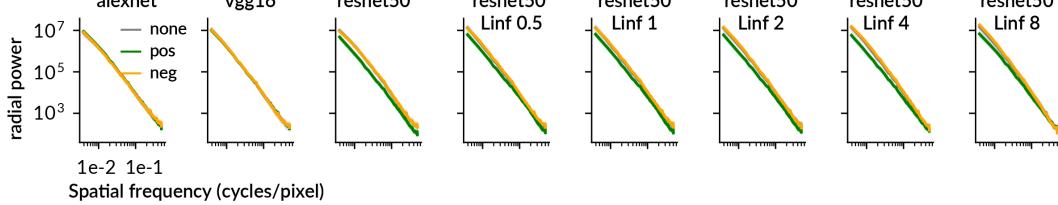
A.4 BIOLOGICAL NEURON MODELS

For each recording session, we selected the best model for further analysis, based on predictive accuracy (mean test $r^2 = 0.27 \pm 0.10$ SD across sessions). The fitted models included both positive and negative input weights, with a mean ratio of 1.17 for the sum of positive to negative absolute weights (Fig. 20). Our final dataset comprised (V1, V4, pIT): (7, 5, 23) neurons in Monkey C and (1, 5, 18) in Monkey D, with the majority of data from pIT cortex.

1026
1027
1028
1029
1030
1031
1032
1033
1034
1035
1036
1037
1038
1039
1040

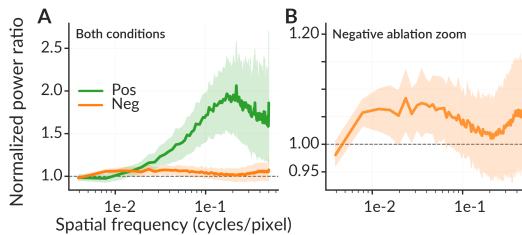


1041 Figure 13: Functional segregation holds in a 10x larger dataset with LPIPS (Zhang et al., 2018) as
1042 representational similarity measure. We measured the representational similarity of the images as $1 -$
1043 LPIPS among control images and between control images and ablation images (Fig. 12). We average
1044 results per class, and show the mean and 95% C.I. across the 100 classes. The representational
1045 similarity degrades upon positive input ablations, confirming results obtained from the imagenette
1046 dataset.



1058 Figure 14: Functional segregation of frequency content. We measured the radial power spectrum for
1059 control and ablation images of different networks. We average results per class, and show the mean
1060 and 95% C.I. across the 11 classes. Low frequencies degrade upon positive input ablations, while
1061 negative ablations overlap with control spectra.

1062
1063
1064
1065
1066
1067
1068
1069
1070
1071
1072
1073



1074 Figure 15: Functional segregation of frequency content. We measured the radial power spectrum for
1075 control and ablation images of different networks. We average results per condition and per network,
1076 and show the ratio of control to ablated spectra. The mean and 95% C.I. are across networks. Low
1077 frequencies degrade upon positive input ablations, while negative ablations overlap with control
1078 spectra but slightly enhance low frequencies.

1079

1080
 1081
 1082
 1083
 1084
 1085
 1086
 1087
 1088
 1089
 1090
 1091
 1092
 1093
 1094
 1095
 1096
 1097
 1098
 1099
 1100
 1101
 1102
 1103
 1104
 1105
 1106
 1107
 1108
 1109
 1110
 1111
 1112
 1113
 1114
 1115
 1116
 1117
 1118
 1119
 1120
 1121
 1122
 1123
 1124
 1125
 1126
 1127
 1128
 1129
 1130
 1131
 1132
 1133

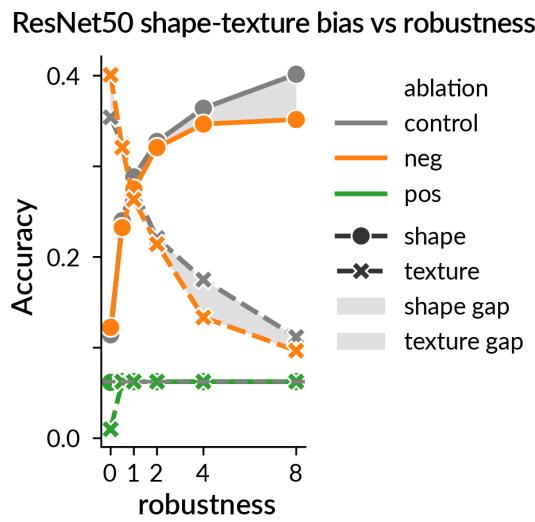


Figure 16: Functional segregation of shape vs texture. We measured robustness effect on shape vs texture encoding in ResNet50 using the shape/texture bias benchmark (Geirhos et al., 2022). Robustness induces a shift from texture to shape, and negative ablations disrupt more strongly the texture encoding of robust networks.

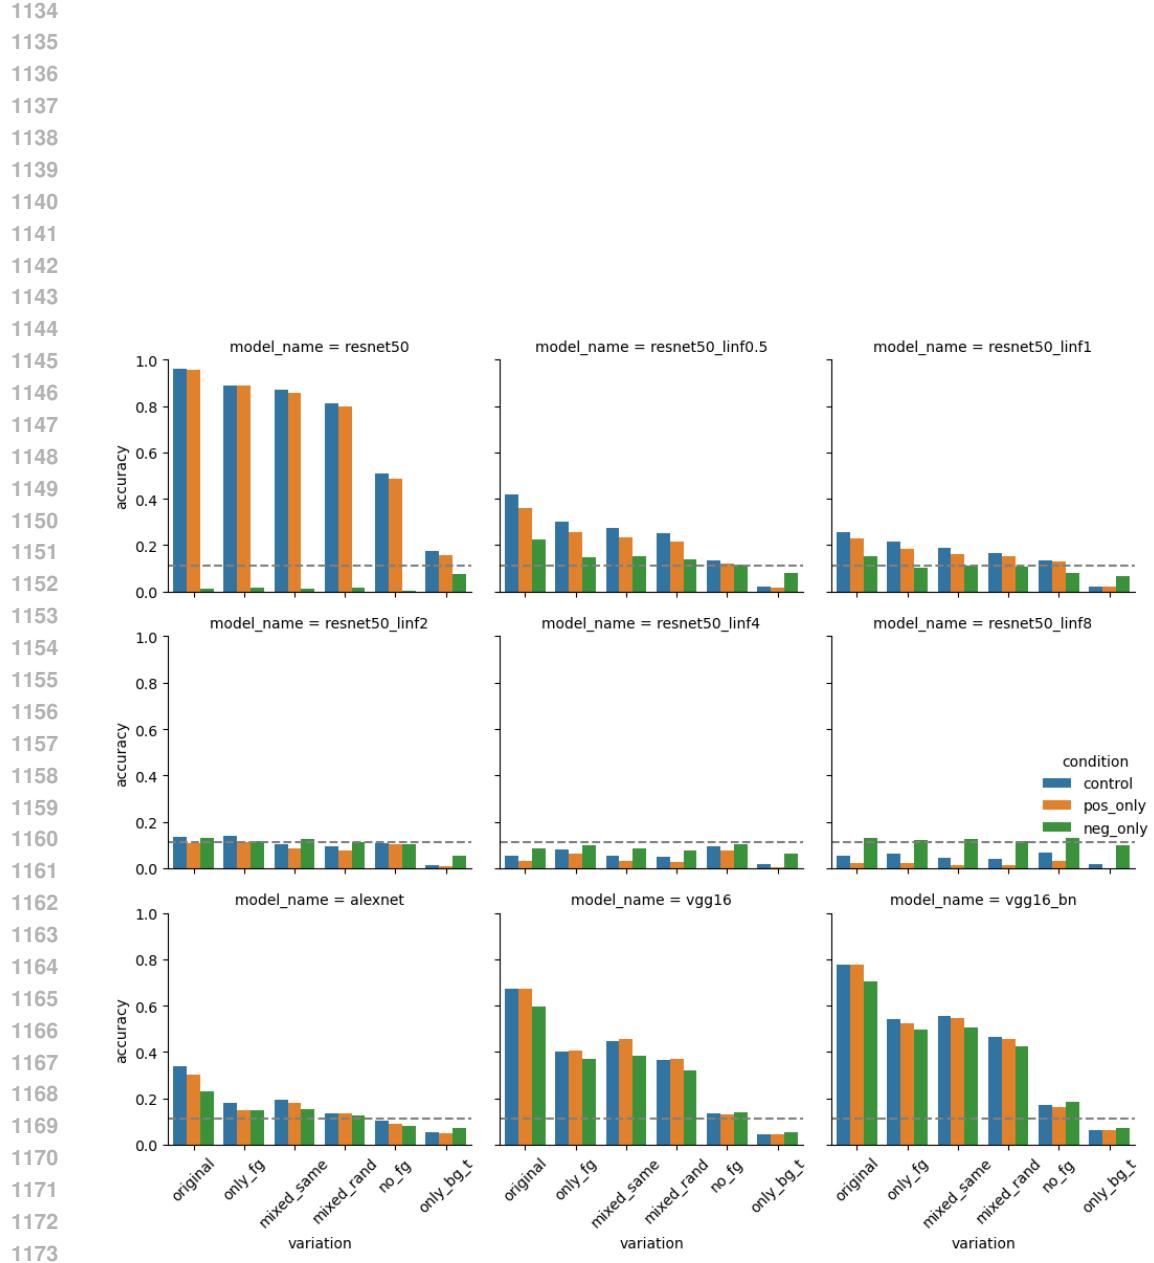


Figure 17: Functional segregation of background vs foreground. We measured the accuracy on the background challenge (Xiao et al., 2020). Robustness increases the role of negative weights in encoding backgrounds but below the chance level. Counterintuitive results show either sign can solve this challenge in AlexNet and VGG.

1188
1189
1190
1191
1192
1193
1194
1195
1196
1197
1198

resnet50-simsiam last fully connected layer

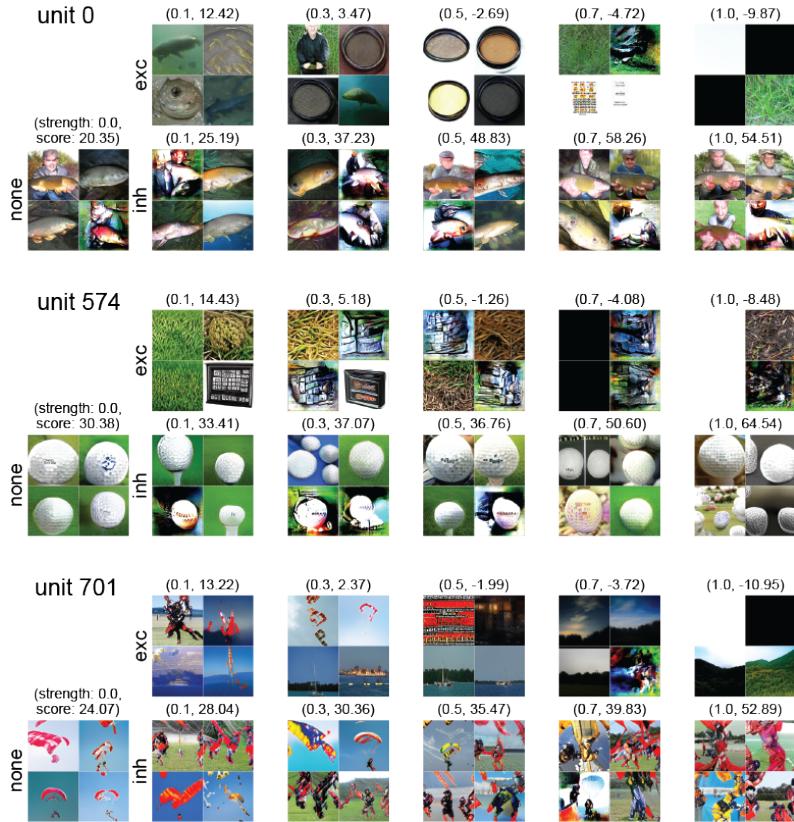


Figure 18: Feature visualizations of ablation experiments in a network pretrained with unsupervised learning. ResNet50SimSiam (Chen & He, 2020). The unsupervised network with frozen weights was coupled to a fully connected layer, only this layer was fine-tuned to classify ImageNet1000. Network units changed starting with small positive weight ablations, see unit 574 golf ball. Smaller changes are visible upon negative weight ablations, however object relevant features remain. Overall behavior is consistent with CNNs trained directly on ImageNet1000 classification.

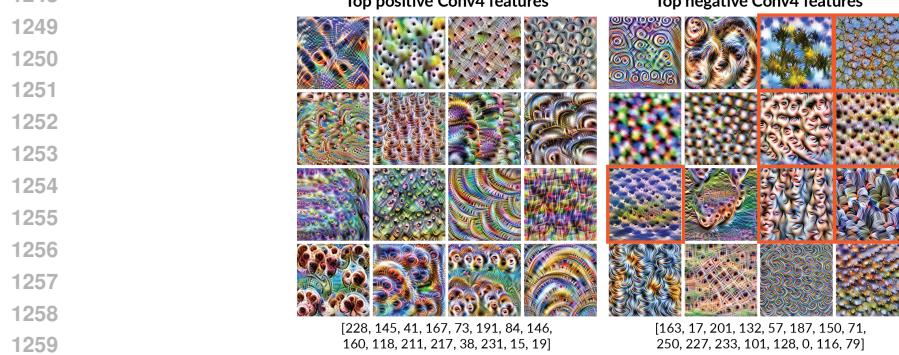
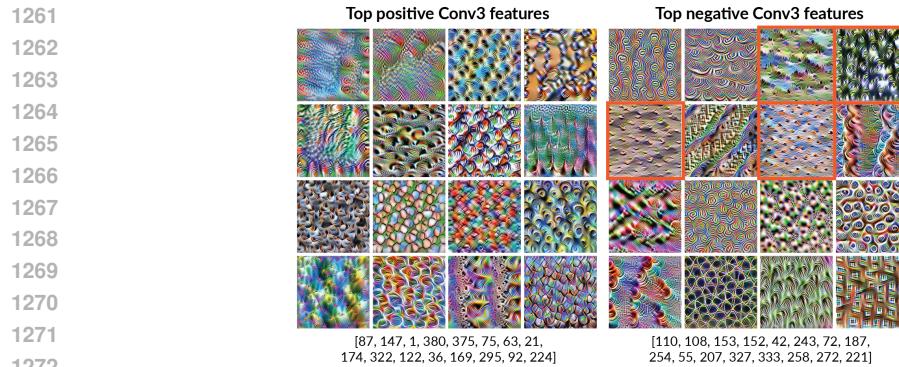
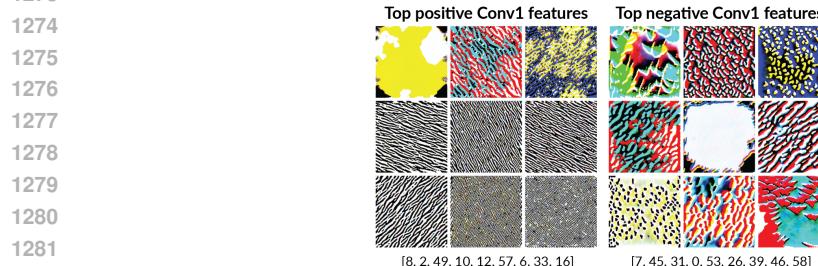
1232
1233
1234
1235
1236
1237
1238
1239
1240
1241

1242

1243

1244

1245

1246 **Visualizing top input channels with sign consistent weights
in internal layers of AlexNet**1247 **Features from Conv4 to Conv5**1260 **Features from Conv3 to Conv4**1273 **Features from Conv1 to Conv2**

1285 Figure 19: Layer Conv5 from Conv4: Features contributing mainly negative weights resemble back-
1286 grounds, such as patches of sky and grass, and sometimes face-like features (e.g., in the tench class),
1287 highlighted in orange borders. Positive weights align with localized object-like fragments, such as
1288 snouts and eyes of animals, and sharp spotted textures vs the blurry spotted textures for negative
1289 weights. Layer Conv4 from Conv3: Negative features still incorporate some background elements
1290 like ground or grass textures (orange borders), together with some spiral, square and blurry textures.
1291 Positive features exhibit more heterogeneous textures and higher frequency details, without evident
1292 background-like textures. Layer Conv2 from Conv1: Positive weights carry high-frequency edges
1293 mostly without color, while negative weights include lower frequency edge features and spotted tex-
1294 tures with color, overall more spatially coarse. Channel index from the visualized features is shown
1295 as a list below each panel.

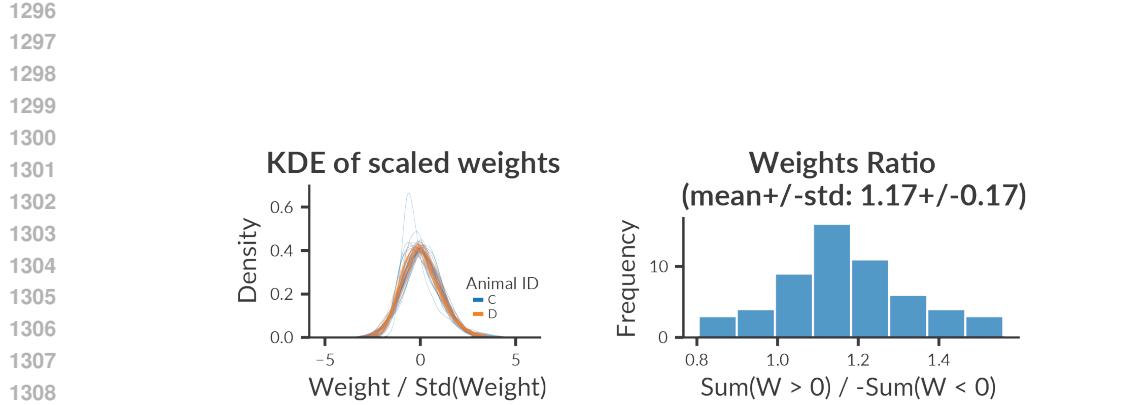


Figure 20: Left: Distribution of the model weights from neuronal fits with AlexNet penultimate layer features. Each model maps 4096 parameters from penultimate layer of AlexNet to the response of one biological neuron. Models use positive and negative weights. Model weights were normalized by their standard deviation to plot them on the same scale, for sake of visualization. Right: Ratio of total positive to total negative weights, per neuron model. Models use slightly larger positive weights with a mean of 1.17 and std of 0.17. Model numbers: 35 for monkey C, and 24 for monkey D.

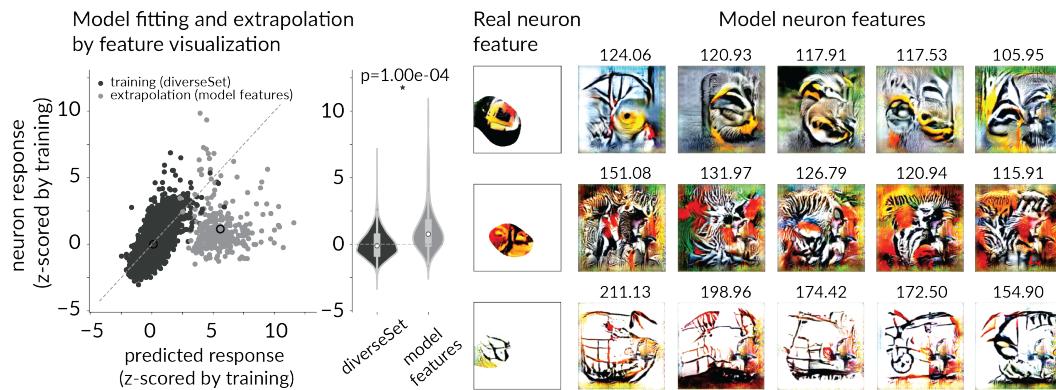
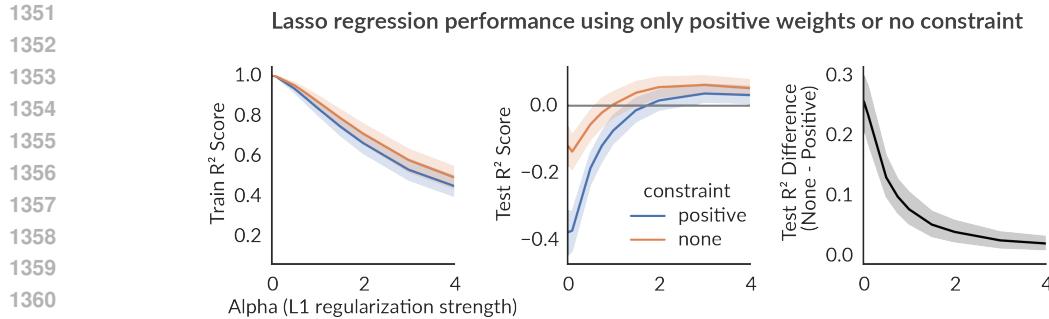


Figure 21: Neuron model units recover features relevant for the biological neurons. Left: Responses vs predicted responses of neurons to the training images, and the extrapolated features visualized from the intact models, which are extrapolations because the training data did not cover those high response ranges. Permutation t-test of neuron responses shows higher responses to images from model features than the natural images of the training dataset (diverseSet). Right: three neuron examples that show the feature visualization of the preferred feature of the neuron masked by the full-width at half-maximum obtained from perturbations to the image, and to their right the five feature visualizations of the intact model with the real neuron responses to those images on top.

1350



1362 Figure 22: Using negative weights improves neuron models obtained via Lasso regression. Lasso
 1363 regression models were fit with and without the positive constraint, over a 5-fold cross validation.
 1364 Models were a linear regression from the 4096 features to a single neuron, over all neurons modeled
 1365 from both animals. Left: performance on the training set measured by r^2 score. Middle: r^2 per-
 1366 formance on the test set. Right: Model improvement by using positive and negative weights vs using
 1367 only positive weights given by the difference in r^2 on the test set. Unconstrained models perform
 1368 better than the positively constrained model, across the range of L1 penalties (sparseness penalty)
 1369 tested, suggesting negative inputs from artificial network features are useful to predict biological
 1370 neuron responses.

1370

1371

1372

1373

1374

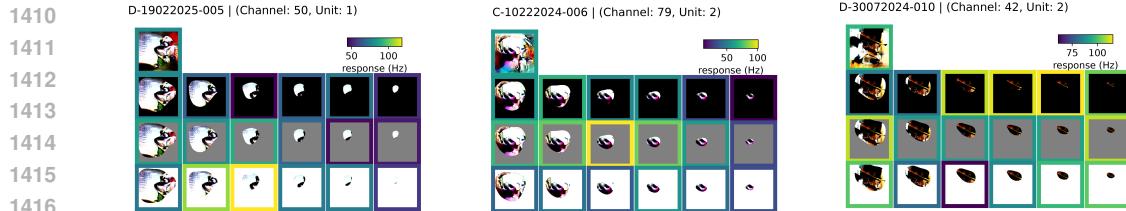
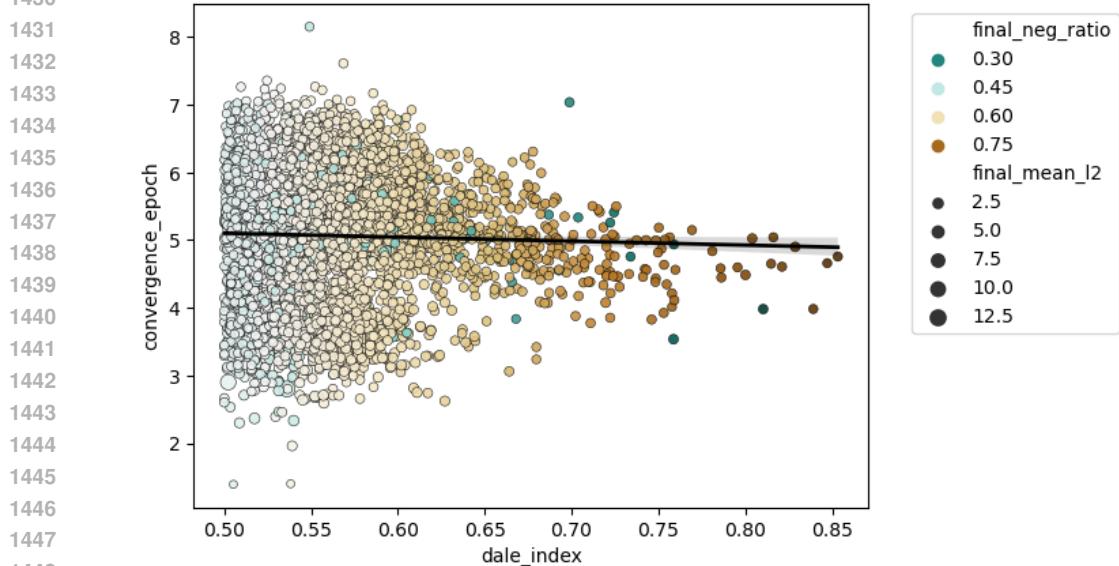
Features from neuron models, AlexNet 4096 ReLU fc layer

Positively weighted Negatively weighted for >90% neuron models



1394

1395 Figure 23: Features that had positive or negative weights in most of the neurons models (91%
 1396 of the 56 neurons, binomial test $p = 5.09e^{-10}$). These features are the closest approximation to
 1397 features respecting Dale's law from our models. Left: best of 20 feature visualizations for the
 1398 features with positive weights across neurons, feature index is on top of the image. Features are from
 1399 the penultimate fc layer post ReLU, containing 4096 units. Right: best feature visualization from
 1400 the negatively weighted features across neurons. Positively weighted features contain more local
 1401 features like curved edges, while negative features contain textures or larger image patches. Sign
 1402 consistency tested for statistical significance against the Bernoulli distribution of 0.5 probability
 1403 with Bonferroni correction for testing 4096 features.

1404
1405
1406
1407
14081409 **Background clearing can enhance neuron responses to the original feature visualization**1417
1418 Figure 24: Clearing the background around the images obtained via closed-loop visualization can
1419 further boost responses in real-time recordings. Examples of 3 neurons in 2 monkeys.1420
1421
1422
1423
1424
1425
1426
1427
14281429 **Convergence epoch vs Dale index across layers -0.03, 6.4e-02**1449
1450 Figure 25: Weight dynamics vs Dale index. Convergence rate of outgoing weights measured by
1451 cosine similarity to final weights, as time to reach 90% of the final similarity. Each dot is one output
1452 channel for all layers in ResNet18 trained over 16 epochs. Color indicates proportion of negative
1453 signs. There is no obvious correlation. However, high Dale index channels mature within narrower
1454 time windows than more mixed channels.1455
1456
1457