



## OPEN ACCESS

## EDITED BY

Denise Manahan-Vaughan,  
Ruhr University Bochum, Germany

## REVIEWED BY

Satoshi Eifuku,  
Fukushima Medical University, Japan  
Hiroshi Tamura,  
Osaka University, Japan

## \*CORRESPONDENCE

Christian Quaia  
✉ quaiaac@nei.nih.gov

RECEIVED 29 April 2024

ACCEPTED 01 July 2024

PUBLISHED 12 July 2024

## CITATION

Quaia C and Krauzlis RJ (2024) Object  
recognition in primates: what can early visual  
areas contribute?

*Front. Behav. Neurosci.* 18:1425496.  
doi: 10.3389/fnbeh.2024.1425496

## COPYRIGHT

© 2024 Quaia and Krauzlis. This is an  
open-access article distributed under the  
terms of the [Creative Commons Attribution  
License \(CC BY\)](#). The use, distribution or  
reproduction in other forums is permitted,  
provided the original author(s) and the  
copyright owner(s) are credited and that the  
original publication in this journal is cited, in  
accordance with accepted academic  
practice. No use, distribution or reproduction  
is permitted which does not comply with  
these terms.

# Object recognition in primates: what can early visual areas contribute?

Christian Quaia\* and Richard J. Krauzlis

Laboratory of Sensorimotor Research, National Eye Institute, NIH, Bethesda, MD, United States

**Introduction:** If neuroscientists were asked which brain area is responsible for object recognition in primates, most would probably answer infero-temporal (IT) cortex. While IT is likely responsible for fine discriminations, and it is accordingly dominated by foveal visual inputs, there is more to object recognition than fine discrimination. Importantly, foveation of an object of interest usually requires recognizing, with reasonable confidence, its presence in the periphery. Arguably, IT plays a secondary role in such peripheral recognition, and other visual areas might instead be more critical.

**Methods:** To investigate how signals carried by early visual processing areas (such as LGN and V1) could be used for object recognition in the periphery, we focused here on the task of distinguishing faces from non-faces. We tested how sensitive various models were to nuisance parameters, such as changes in scale and orientation of the image, and the type of image background.

**Results:** We found that a model of V1 simple or complex cells could provide quite reliable information, resulting in performance better than 80% in realistic scenarios. An LGN model performed considerably worse.

**Discussion:** Because peripheral recognition is both crucial to enable fine recognition (by bringing an object of interest on the fovea), and probably sufficient to account for a considerable fraction of our daily recognition-guided behavior, we think that the current focus on area IT and foveal processing is too narrow. We propose that rather than a hierarchical system with IT-like properties as its primary aim, object recognition should be seen as a parallel process, with high-accuracy foveal modules operating in parallel with lower-accuracy and faster modules that can operate across the visual field.

## KEYWORDS

visual recognition, modeling, visual cortex, face processing, neural networks

## Introduction

It is commonly accepted that in primates visual recognition of an object is strongly dependent on the activity of neurons in the inferior temporal (IT) cortex (Logothetis and Sheinberg, 1996; Tanaka, 1996), an area at the end of the ventral visual stream (Ungerleider and Mishkin, 1982; Goodale and Milner, 1992). This hypothesis was first prompted by the discovery in monkey IT of neurons that respond vigorously to complex object shapes and faces (Gross et al., 1972; Desimone et al., 1984; Gross, 2008). Comparisons between the activity of artificial neural networks (ANNs) trained on object classification and the activity of large populations of IT neurons have lent further support to this hypothesis (Khaligh-Razavi and Kriegeskorte, 2014; Yamins et al., 2014; Kriegeskorte, 2015; Yamins and DiCarlo, 2016;

Rajalingham et al., 2018). Additional support came from the observation that large bilateral lesions of IT in monkeys result in considerable deficits in performing fine discriminations (Eldridge et al., 2018; Setogawa et al., 2021; Matsumoto et al., 2022). Similarly, causal manipulations of monkey IT activity while the animal is performing difficult discriminations impairs performance, further indicating a causal involvement of IT (Afraz et al., 2006, 2015; Moeller et al., 2017).

Visual object recognition is however not limited to such fine discriminations. It is also important for quickly identifying and locating a potential prey or predator, thus providing information critical for survival. These needs are shared by animals that do not have the high visual acuity of diurnal primates, or anything resembling the ventral visual pathway (Kirk and Kay, 2004; Vinken et al., 2016; Kaas, 2020; Leopold et al., 2020), and can usually be satisfied by a coarser type of recognition (enabling, for example, to tell apart a mouse from a cat). Just as importantly, fine discriminations are supported almost exclusively by foveal vision. Accordingly, IT has a strong foveal bias (De Beeck and Vogels, 2000; Arcaro et al., 2017; Op de Beeck et al., 2019), meaning that its neurons respond preferentially to images placed on or near the fovea. Initial reports seemed to indicate that IT had large receptive fields covering most of the contralateral visual field, but that was confounded by the use of large stimuli. When small stimuli are used, moving the stimulus by as little as 2° away from the fovea results in a dramatic reduction in activity, even for preferred objects (DiCarlo and Maunsell, 2003; Hung et al., 2005). Taken together these observations suggest that IT in primates might be specialized for fine, mostly foveal, discriminations. This would readily explain the rather mild deficits in coarse object classification (e.g., dog vs. cat) observed following bilateral ablations of IT in monkeys (Matsumoto et al., 2016).

Setting aside the needs of foveated animals, we now end up with a chicken-and-egg problem: If IT is crucial for fine discriminations, and it relies predominantly on foveal inputs, how can an object of interest be foveated in the first place? Since our visual field covers approximately 180° × 120°, and the fovea is approximately 2° across, it would take 5,400 refixations—18 min at a rate of 5 saccades/s—to scan the entire visual field. Since this is not our daily experience (Potter, 1975; Najemnik and Geisler, 2008; Drewes et al., 2011), there must exist a system that is capable of locating objects in the environment and prioritizing those of likely behavioral interest. Importantly, this system must provide accurate information about the position of the target object, so that foveation can be achieved in a single saccade. This would require specificity for retinotopic location, thus eschewing the position invariance sought by ANNs.

The need for such a system—often involving a saliency or priority map—has been long recognized (Treisman, 1986; Fecteau and Munoz, 2006; Serences and Yantis, 2006). The saliency map is usually envisioned as a topographically organized neural map of visual space, in which activity at any given location represents the priority to be given to that location for covert or overt shifts of attention (Itti et al., 1998). Originally, bottom-up saliency (e.g., local contrast, orientation, etc.) was considered the main determinant of this map (Itti and Koch, 2000), but it has become clear that the identification of objects (Einhäuser et al., 2008), and even more complex cognitive processes (Schütz et al., 2011), play a major role. Its neural substrate has been hypothesized to lie in a network of interconnected areas involving the superior colliculus (SC) in the midbrain (McPeck and

Keller, 2002; White et al., 2017), the pulvinar nuclei in thalamus (Robinson and Petersen, 1992), the frontal eye fields (FEF) in prefrontal cortex (Thompson and Bichot, 2005), and the lateral intraparietal (LIP) area in parietal cortex (Gottlieb et al., 1998). The SC is likely to play a pivotal role, as it has a finer topographical organization than FEF and LIP, and it is more directly involved in controlling gaze and attention shifts (Krauzlis, 2005; Krauzlis et al., 2013). It also receives retinotopic inputs from most of the visual cortex, including areas V1, V2, V3, V4, and IT (Fries, 1984; Cusick, 1988; May, 2006). As noted, IT has a strong foveal bias, and presumably contributes less to the processing of peripheral objects and events than areas V1–V4. Topographically organized inputs from these areas to the salience network might then form the basis of a spatially localized system for peripheral (and coarse) object detection.

We thus sought to evaluate the ability of the first stages of cortical visual processing, simple and complex cells in area V1, to convey information about object category. We focused on information that can be extracted with a simple linear classifier (i.e., a weighted sum followed by a threshold operation, a good approximation of the operation performed by a neuron downstream of V1, Rosenblatt, 1962). We restricted our model to using only relatively low spatial frequencies, thus approximating processes in peripheral vision (Robson and Graham, 1981) or in infancy (Sokol, 1978). Finally, we focused on the classification task of identifying a face among non-face images. This is a coarse recognition task of strong behavioral significance for primates, and at which they are highly proficient (Rosenfeld and Van Hoesen, 1979; Diamond and Carey, 1986; Carey, 1992; Tanaka and Gauthier, 1997; Pascalis and Bachevalier, 1998; Pascalis et al., 1999), even when very young (Nelson and Ludemann, 1989; Pascalis and de Schonen, 1994; Nelson, 2001), and even in the (near) periphery (Mäkelä et al., 2001; McKone, 2004). To investigate the plausibility of V1 being directly involved in the detection of faces in the periphery, we assessed how variations in scale, orientation, and background affected the ability of a V1-like model to provide information useful to classification.

## Methods

### Image sets

For our face vs. non-face image classification task, we assembled three sets of images: A set of 1,756 frontal images of human faces, a set of 1,502 images of non-face stimuli (which includes a variety of human-made objects, animals, fruits, and vegetables), and a set of 1,525 images of urban and natural landscapes.

We derived the face dataset from a larger (2,000 images) set of frontal images of faces that was kindly shared with us by Prof. Doris Tsao. The set contained images of subjects from both sexes and from all races, under varying levels of illumination. We modified this set in a few ways. First, we dropped 244 images that either had a lower resolution compared to the others or in which a large fraction of the face was occluded by hair or face coverings. Next, we resized each image from the original 256 × 360 (width × height) pixels (px) to our desired square size (256 × 256 px). This was done by preserving the aspect ratio and adding white space on the left and right sides. The image was then saved as a grayscale image with a transparency (alpha) channel (an RGBA image in which the red, green, and blue channels all have the same value). Because most images

also contained hair, neck and shoulder, we used the dlib library (King, 2009) and the MTCNN face segmentation neural network (Zhang et al., 2016) to isolate the face region, removing completely neck and shoulders and leaving only minimum hair (while retaining an oval shape). Anything outside the face region was then set to transparent. Finally, the face region was rescaled (preserving its aspect ratio) to make its largest side equal to 246 px, and it was centered on the  $256 \times 256$  px canvas. 8 samples from this set, centered on a  $340 \times 340$  px mid gray background, are shown in Figure 1 (top row). Note that to make the subjects not identifiable we have blurred their images in Figure 1 and Figure 4. The images presented to the models were however not blurred.

The images for the non-face dataset were selected by us from images freely available on the web for non-commercial use. We selected images of a wide range of animals with many different poses, fruits and vegetables, and man-made objects. We intentionally did not select any “unnatural” images, such as cartoons, flags, letters, and numbers, and we tried to add as many oval/circular shapes as we found. The images were of varying sizes and formats, and almost all were in color with a transparent background (we segmented those that did not have a transparent background to create one). We converted them all to grayscale and resized them (preserving their aspect ratio) so that their longest side was 246 px, and then centered them on a  $256 \times 256$  px transparent canvas.

We then computed the distribution of mean luminance and contrast of both sets of images. We found that their contrast levels overlapped, but non-face images were on average darker. We thus increased their luminance by 15% to approximately match the distributions. 8 samples from this set, centered on a  $340 \times 340$  px mid gray background, are shown in Figure 1 (second row). The distribution of mean luminance, contrast, and two measures of shape (the fraction

of foreground pixels within the original  $256 \times 256$  px canvas, and the elongation ratio, defined as the ratio of the largest and smallest dimension of the rectangle that encompasses the foreground image) for the images we fed to the models are shown in Figure 2. Not surprisingly, the shape variation across the non-face set is much larger than across the face set.

Landscape images were selected from MIT’s Places dataset of 10 million images (Zhou et al., 2018). We went through the dataset one image at a time until we had collected approximately 1,500 images that included neither humans nor animals, and depicted natural and urban scenes (for the latter mostly exteriors, with some interior scenes). Since all the images were in color, we then converted them to grayscale, resized them to  $340 \times 340$  px, and corrected over and under exposed images by equalizing their luminance histogram using the OpenCV library (Bradski and Kaehler, 2008). 8 samples from this set of scenes are shown in Figure 1 (third row).

A second set of landscape images was obtained from this set by reducing the contrast of each image by 50% and blurring it with a Gaussian filter with a standard deviation of 3 px. We refer to this set as blurred scenes. 8 samples from this set of blurred scenes are shown in Figure 1 (bottom row).

Images from the face and non-face sets were provided as inputs to the models (see below) one at a time, superimposed over a background canvas having a size of  $340 \times 340$  px. Besides the high contrast and blurred scenes discussed above, the background could be a uniform mid gray (128), a uniform gray with the same luminance as the image presented, 2-D pink (1/f) noise, pixelated ( $4 \times 4$  px blocks) white noise (i.e., each block of pixels could have any luminance value from 0 to 255, randomly selected), pixelated binary noise (i.e., each block of pixels was set to 0 or 255, randomly), or pixelated noise with



FIGURE 1

Random samples of images from sets. Top row: 8 images from our set of 1,756 face images. Variations in sex, age, gender, and race, as well as contrast and mean luminance, seen above are representative of the entire set. Here, but not in the simulations, the images have been blurred to prevent identification of the subject. Second row: 8 images from our set of 1,502 non-face images. Variations in shape, contrast and mean luminance seen above are representative of the entire set. Third row: 8 images from our set of 1,525 high contrast background scenes. Variations in content, contrast and mean luminance seen above are representative of the entire set. Bottom row: Same 8 background images, but blurred and with their contrast reduced.

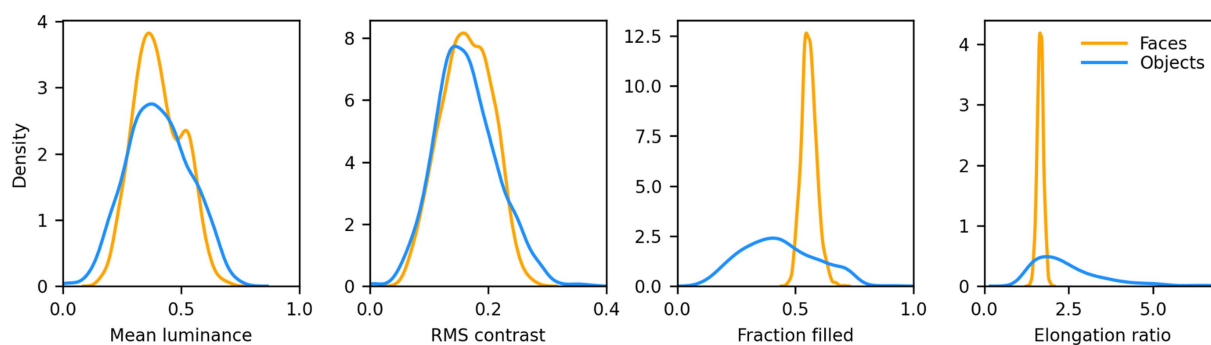


FIGURE 2

Distributions of low-level features for the faces and non-faces image sets. Mean luminance and RMS contrast of the images were well matched (Average mean luminance: face = 102.96, non-face = 103.07; Average RMS contrast: face = 41.55, non-face = 41.54). The fraction filled represents the fraction of pixels within the  $256 \times 256$  area of the image proper (i.e., no background) that belong to the foreground (i.e., the face or non-face itself). Faces had a much narrower distribution, and a larger average (0.56 vs. 0.43), than non-faces. The elongation ratio is the largest of the width/height or height/width ratios for the foreground. Faces had a much narrower distribution, and a smaller average (1.66 vs. 2.51).

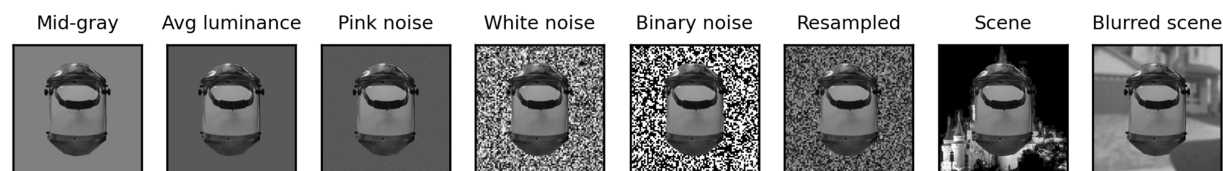


FIGURE 3

Superposition of a single non-face image on the backgrounds that we considered.

luminance values sampled from the image content. Samples of a single non-face image superimposed on these eight backgrounds are shown in Figure 3.

Besides being superimposed on the background, each image could be rescaled (always reduced in size, up to 50% in each dimension) or rotated (up to  $180^\circ$  in each direction). When in the text we refer to a scaling of images by  $x\%$  we mean that each image was scaled by a random factor that varied between 100% and  $x\%$ . Similarly, when we refer to a rotation of  $x^\circ$  we mean that each image was rotated by a random angle between  $x^\circ$  clockwise and  $x^\circ$  counterclockwise. In Figure 4, we show 4 sample faces and 4 sample non-faces presented over four different types of background at a scaling of up to 50% and rotation of up to  $\pm 90^\circ$ .

Finally, the location of each image on the background canvas was jittered (relative to the center) by a random number of pixels in the horizontal and vertical direction, up to  $\pm 42$  px (because the original image was always centered in a  $256 \times 256$  px window, even with the shift the image was always fully contained in the  $340 \times 340$  px background canvas).

In our last experiment, we tested our models on a gender classification task. For this task we manually classified the images from our face set into male or female, based on subjective appearance. This step was carried out on the original images, but the cropped versions of these images (same as for the face/non-face classification) were used as inputs to the models. A small fraction of the faces for which we felt unsure about the gender were excluded. We ended up with 924 images of males and 805 images of females, which we fed to our V1-like and AlexNet model (see below). For the FaceNet model we also created a

set of images that were cropped to a section of each image that was identified by MTCNN (Zhang et al., 2016) as the bounding box for the face (the full horizontal extent of the face is usually kept, but the image is cropped vertically across the chin and the forehead). This region was then stretched (without preserving the aspect ratio) to match the image size required by FaceNet ( $160 \times 160$  px). In a final test we also used these images for AlexNet (resizing them to  $224 \times 224$  px) or our V1-like model (resizing them to  $256 \times 256$  px, centered on a  $340 \times 340$  px mid gray background).

## V1-simple cells model

We implemented a two-layer model for binary classification in which the first layer was composed of a fixed set of oriented Gabor filters that resemble the receptive field properties of V1 simple cells (Marcelja, 1980; Jones and Palmer, 1987a). We used four sets of Gabor filters, each tuned to a different spatial frequency, spaced in one octave (i.e., a factor of two) steps. Since our images were scaled to fit in  $256 \times 256$  px, we used for our largest scale 2-D Gabor functions with a carrier having a wavelength of 256 px. The wavelengths for the other three sets were then 128, 64, and 32 px. We set the standard deviation of the Gaussian envelope of the Gabor to 40% of the carrier wavelength, a value representative of orientation tuned cells in macaque V1 (Ringach, 2002) and cat area 17 (Jones and Palmer, 1987b). We truncated the size of each filter to 1.5 times the wavelength, again matching average properties of macaque V1 cells. At each scale we defined 16 different filters, covering 4 orientations ( $0^\circ$ , i.e., vertical,

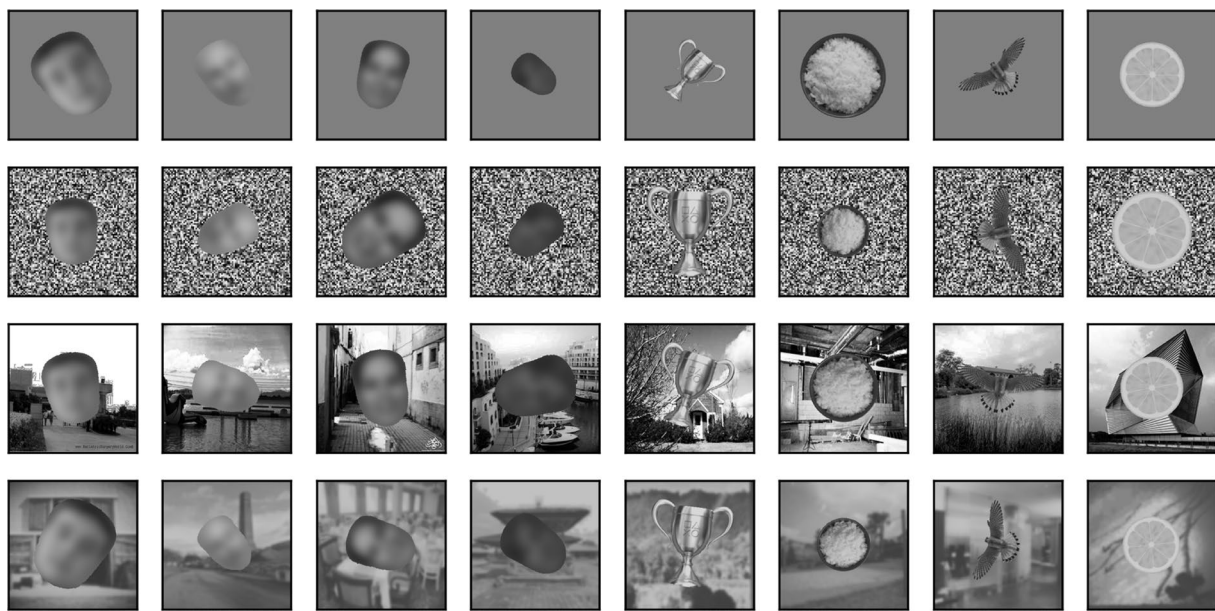


FIGURE 4

Samples of 4 faces and 4 non-faces on different backgrounds (from top to bottom: mid-gray, pixelated white noise, high contrast scene, blurred scene). Each image was randomly scaled down by up to 50% of its original size and rotated by up to 90° in either direction before being superimposed on the background. Here, but not in the simulations, face images have been blurred to prevent identification of the subject.

45°, 90°, and 135°) and four spatial phases (0°, 90°, 180°, and 270°) for the sinusoidal carrier. The filters for one scale are shown in Figure 5A (the top row shows the amplitude profile of the filters shown in the second row, those tuned to vertical orientation). Each set of 16 filters then tiled the entire space of the image (convolutional network), with a stride (distance in the horizontal or vertical direction between two nearby filters in image space) equal to half their wavelength. Because higher frequency filters have smaller sizes, they were more numerous than low spatial frequency filters. Overall, the model had 7,088 linear filters (distributed across the four spatial frequency channels as 144, 144, 1,024, and 5,776). Each filter was followed by a half-rectifying nonlinearity (negative values were set to zero, positive values were kept unchanged). The half-rectified outputs of the filters were then fed into a linear classifier, whose 7,088 trainable weights were initialized with the Kaiming uniform initialization method (He et al., 2015; the bias term was initially set to zero). This process of summing the rectified outputs of the filters with learned weights is depicted in Figure 5B.

Since we randomly jittered the location of the images by  $\pm 42$  px, units with shorter wavelengths/higher spatial frequencies than those we used (e.g., 16 px) cannot be expected to extract any reliable information from the images, and thus would mostly contribute noise to the classifier.

## V1-complex cells model

We implemented a model of V1 complex cells by simply summing together the outputs of the V1 simple cells in the above model that shared the same spatial location, spatial frequency, and orientation (i.e., we summed the output of single cells that differed only in spatial phases, Figure 5C). This is a standard way of simulating V1 complex

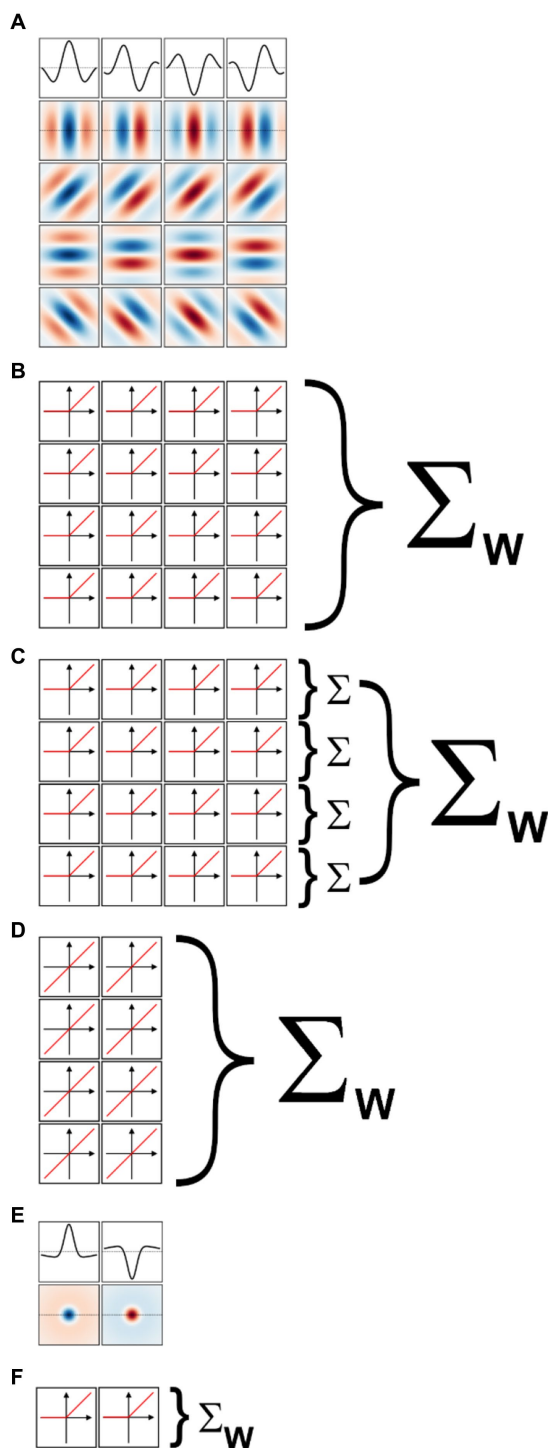
cells, analogous to the energy model of V1 complex cells (Carandini et al., 2005; Lian et al., 2021) but more biologically realistic (Hubel and Wiesel, 1962; Movshon et al., 1978; Martinez and Alonso, 2003). Because we considered four spatial phases at each scale and orientation, our V1-complex cells model had 1,772 units, whose outputs were fed to the linear classifier to learn the summing weights.

## V1 linear receptive field model

We also implemented a model of V1 simple cells linear receptive fields. Essentially, we took the model of V1-simple cells described above and removed the half-rectifying nonlinearity. Because with this modification the outputs of units with a spatial phase of 180° (270°) were simply the opposite of units with a spatial phase of 0° (90°), and thus entirely redundant (the weights of the classification layers can be positive or negative), we removed them. Thus, our V1 linear receptive field (RF) model had 3,544 units, whose outputs were fed to the linear classifier (Figure 5D).

## LGN-like model

We also implemented a two-layer model for binary classification in which the first layer was composed of a fixed set of difference of Gaussians (or Mexican-hat) filters that resemble the receptive field properties of neurons in the lateral geniculate nucleus (LGN; Rodieck, 1965; Enroth-Cugell and Robson, 1966). Just as for the V1-like model, we used four sets of filters, each tuned to a different spatial frequency, spaced in one octave (i.e., a factor of two) steps. To match the preferred spatial frequency of the V1 filters, we set the standard deviation of the center Gaussian to 1/12th of the desired spatial frequency, the standard



**FIGURE 5**  
 Models of early visual processing. **(A)** Examples of oriented Gabor filters used for V1 model cells at one spatial location, and for one spatial scale. Four spatial phases (columns) and four orientations (rows) were used. Positive values are shown in shades of blue, and negative values in shades of red. The top row shows the amplitude profile of the filters in the second row. **(B)** In a model of V1 simple cells the output of each filter is half-rectified, and then summed with a weight ( $\Sigma w$ ) that is learned during model training. **(C)** In a model of V1 complex cells the outputs of half-rectified filters associated with the same orientation, but different phases, are first summed ( $\Sigma$ ). The values for the various orientations are then summed with learned weights. **(D)** In a linear model of V1 cells, only two phases are considered (the others are redundant), and their outputs are directly summed with learned weights. **(E)** Difference of Gaussians filters used for LGN model cells at one spatial location, and for one spatial scale. On-center (left) and off-center (right) cells were used. **(F)** In a model of LGN the output of each filter is half-rectified and then summed with a learned weight.

deviation of the surround Gaussian to 5 times that value, and the gain of the surround Gaussian to 0.2 (the gain of the center Gaussian was 1). The size of the filters was the same as for the V1 filters. At each scale

we defined 2 different filters, one mimicking on-center cells (i.e., cells that prefer a bright center and a dark surround) and the other mimicking off-center cells (i.e., cells that prefer a dark center and a

bright surround). These filters are shown, for one spatial scale, in [Figure 5E](#) (the first row shows the amplitude profile of the filters shown in the second row). As with the V1 model, each set of 2 filters then tiled the entire space of the image (convolutional network), with a stride (distance in the horizontal or vertical direction between two nearby filters in image space) equal to half their preferred wavelength. Because higher frequency filters have smaller sizes, they were more numerous than low spatial frequency filters. Overall, the model had 886 linear filters (distributed across the four spatial frequency channels as 18, 18, 128, and 722). Each filter was followed by a half-rectifying nonlinearity (negative values were set to zero, positive values were kept unchanged). The half-rectified outputs of the filters were then fed into a linear classifier ([Figure 5F](#)), whose weights were initialized with the Kaiming uniform initialization method ([He et al., 2015](#); the bias term was initially set to zero). As with the V1 model, using filters tuned to higher spatial frequencies would be unwarranted given our image jittering.

## AlexNet model

In some simulations we also used the model developed by [Krizhevsky et al. \(2012\)](#) for image classification. It is not the most sophisticated model available now, but it is the one that opened the ANN revolution in image recognition and it is still widely used as a benchmark. We used the model pretrained on the ImageNet classification task, which involves classifying images in one of 1,000 categories. We retained its stack of feature layers (which has 9,216 outputs), stripped it of its fully connected and classification layers, and replaced those layers with a single binary classification layer (identical to the one we used for our V1 model). Since AlexNet expects images of size  $224 \times 224$  px, the  $340 \times 340$  px images we used for our V1 model were rescaled to this size before feeding them to the model.

## FaceNet model

For the gender classification task we also used, as a benchmark, the model developed at Google by [Schroff et al. \(2015\)](#) for face recognition. We used the Inception Resnet V1 model pretrained on  $160 \times 160$  px images from the VGGFace2 face recognition task, which involves identifying a subject (out of a pool of 9,131 subjects, for which 3.31 million images are available). We retained its stack of feature layers (which in this model had 1,792 outputs), stripped it of its fully connected and classification layers, and replaced those layers with a single binary classification layer (identical to the one we used for our V1 model). Since FaceNet expects images of size  $160 \times 160$  px and limited to the face bounding box extracted by MTCNN, the images we used for this model from our gender data set were pre-processed accordingly.

## Multi-tail AlexNet model

To quantify how the ability of AlexNet to extract useful features from our images varies as their processing progresses, we created a variant of AlexNet in which the classification is not based only on the final layer of features (as typically done and as we did above), but also on the output of intermediate layers. More precisely, we added a

classification layer after the third, sixth, eighth, and tenth layers in the feature stack (in addition to the classifier that follows the 13th and last layer in the feature stack). Each binary classifier thus had a different number of inputs: classifier C1 had 46,656, C2 had 32,448, C3 had 64,896, C4 had 43,264, and C5 had 9,216 as before. These classification layers were trained concurrently but independently, and thus each maximized its ability to best predict the image based on the features it extracted.

## Model training and performance evaluation

For the face vs. non-face task, we used the same procedure for all models. First, we randomly split each set of images (faces and non-faces) into two distinct sets, each with half the images (50/50 split). We then used one half for training the model, and the other half to evaluate its performance, which is thus fully cross-validated (the model was tested on images that were never seen by the model during the training phase). Unlike in classic model training tasks we did not feed the model identical images multiple times. Instead, we initially sampled, with replacement, 1,500 images from each of the faces and non-faces training subsets. Before feeding each of these images to the model, they were jittered by a random amount and, when required, also rescaled and/or rotated by a random amount. Finally, a background was selected for the image (which, except for the mid gray background, would also be different for every image). Thus, if an image was present in the resampled training set twice, its location, scale, rotation, and background could/would be different in the two presentations. And in different epochs (training runs through the entire training sets) the images would thus also always be different. This can be seen as an extreme form of data augmentation, as the model was effectively never fed the same input twice. In practice this minimizes the risk of overfitting and maximizes model generalization. We empirically verified that performance on the testing set increased with the number of epochs, but quickly saturated, plateauing after 50 epochs, but improving only marginally after 10 epochs. Accordingly, we did not feel the need to use a validation set to prevent training from overfitting, and simply trained each model for 10 epochs.

After the model was trained, we sampled, with replacement, 1,500 images from each of the faces and non-faces testing subsets. We then fed these images through the model once and evaluated its performance (fraction of correctly classified images).

This entire process was repeated 100 times (100 different training/testing 50/50 splits), and we collected the performance for each round, thus obtaining a distribution of performance values. To summarize the performance of a model for a given condition (amount of scaling and rotation, and type of background) we report the median of the distribution of these 100 performance values. Another measure often used to quantify performance for binary categorization is  $d'$ , which is computed as  $d' = Z(\text{hit rate}) - Z(\text{false alarm rate})$ , where  $Z$  is the inverse of the cumulative distribution of the standard normal distribution. Because in our study classification was on average symmetric (i.e., there were as many faces identified as non-faces as there were non-faces identified as faces) this can be simplified to  $d' = 2 \times Z(\text{hit rate})$ .

For the gender classification task, since we had a considerably smaller image set, and some of the simulations (particularly those associated with the FaceNet model) were restricted to no or minimal

data augmentation and thus a high risk of overfitting to the training set, we split the image set in non-overlapping training (50%), validation (10%), and testing (40%) sets. Training was then executed over up to 40 epochs and stopped based on the performance on the validation set (when performance failed to improve over 5 consecutive epochs). For the training and testing sets we sampled, with replacement, 400 images from each of the male and female training subsets. For the validation set we used 80 male and 80 female images (sampled without replacement) from the validation set. Since model performance was measured on the test set, whose images were never used during the training phase, it was fully cross-validated.

The models were implemented in Python 3.10, using the PyTorch 2.0 library. We trained the models to minimize the binary cross entropy with logits loss, the default choice for binary classification problems (Hopfield, 1987; Solla et al., 1988; Bishop, 1996). We trained the model using Stochastic Gradient Descent (SGD) with a learning rate of 0.01 and momentum of 0.9. Batch size was set to 64, and shuffling was applied during training.

## Results

### Face vs. non-face classification

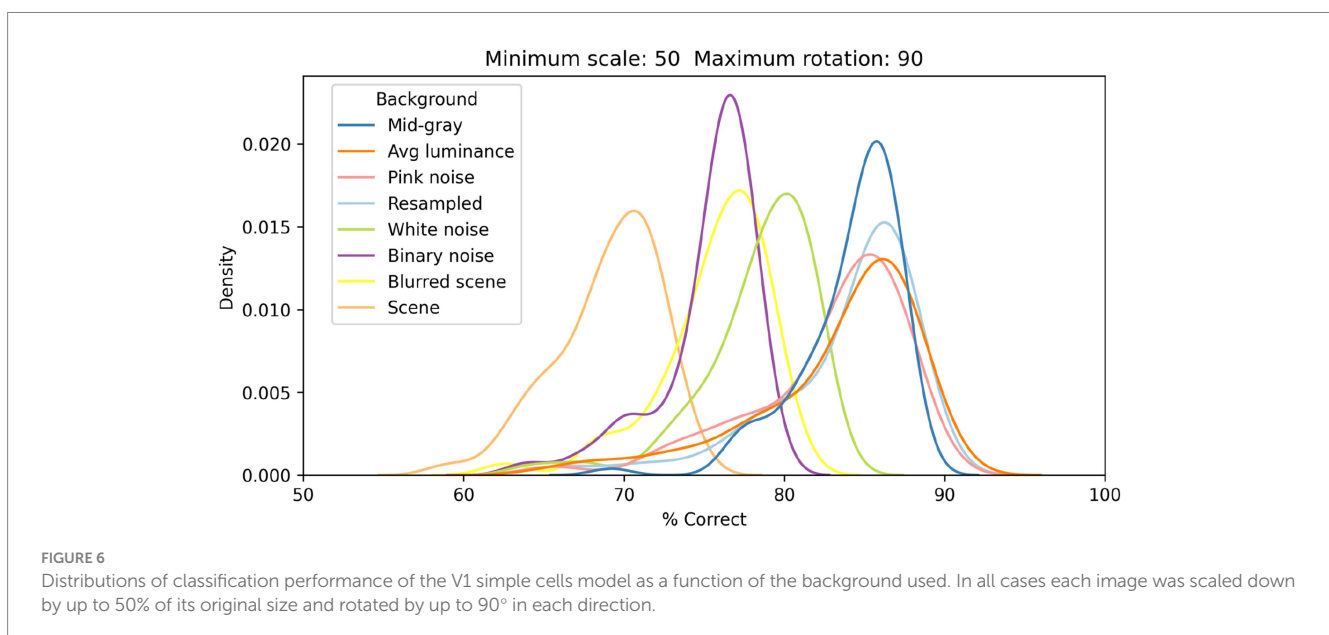
Our primary goal was to assess the impact of various factors (extent of scale variation, extent of orientation variation, and type of background) on the classification performance of our model of V1 simple cells. We started by fixing scale and rotation variation (at up to 50% and 90°, respectively) and evaluated the impact of 8 different types of backgrounds. Figure 4 shows samples of both faces (left four columns) and non-faces (right four columns) at these levels of scale and rotation variation, on four different backgrounds (each on a different row): mid-gray, pixelated white noise, high-contrast scenes, and blurred scenes.

For each type of background, we split the face and non-face image sets into two non-overlapping halves, and used one half (of each image set) for training the model, and the other half for testing its

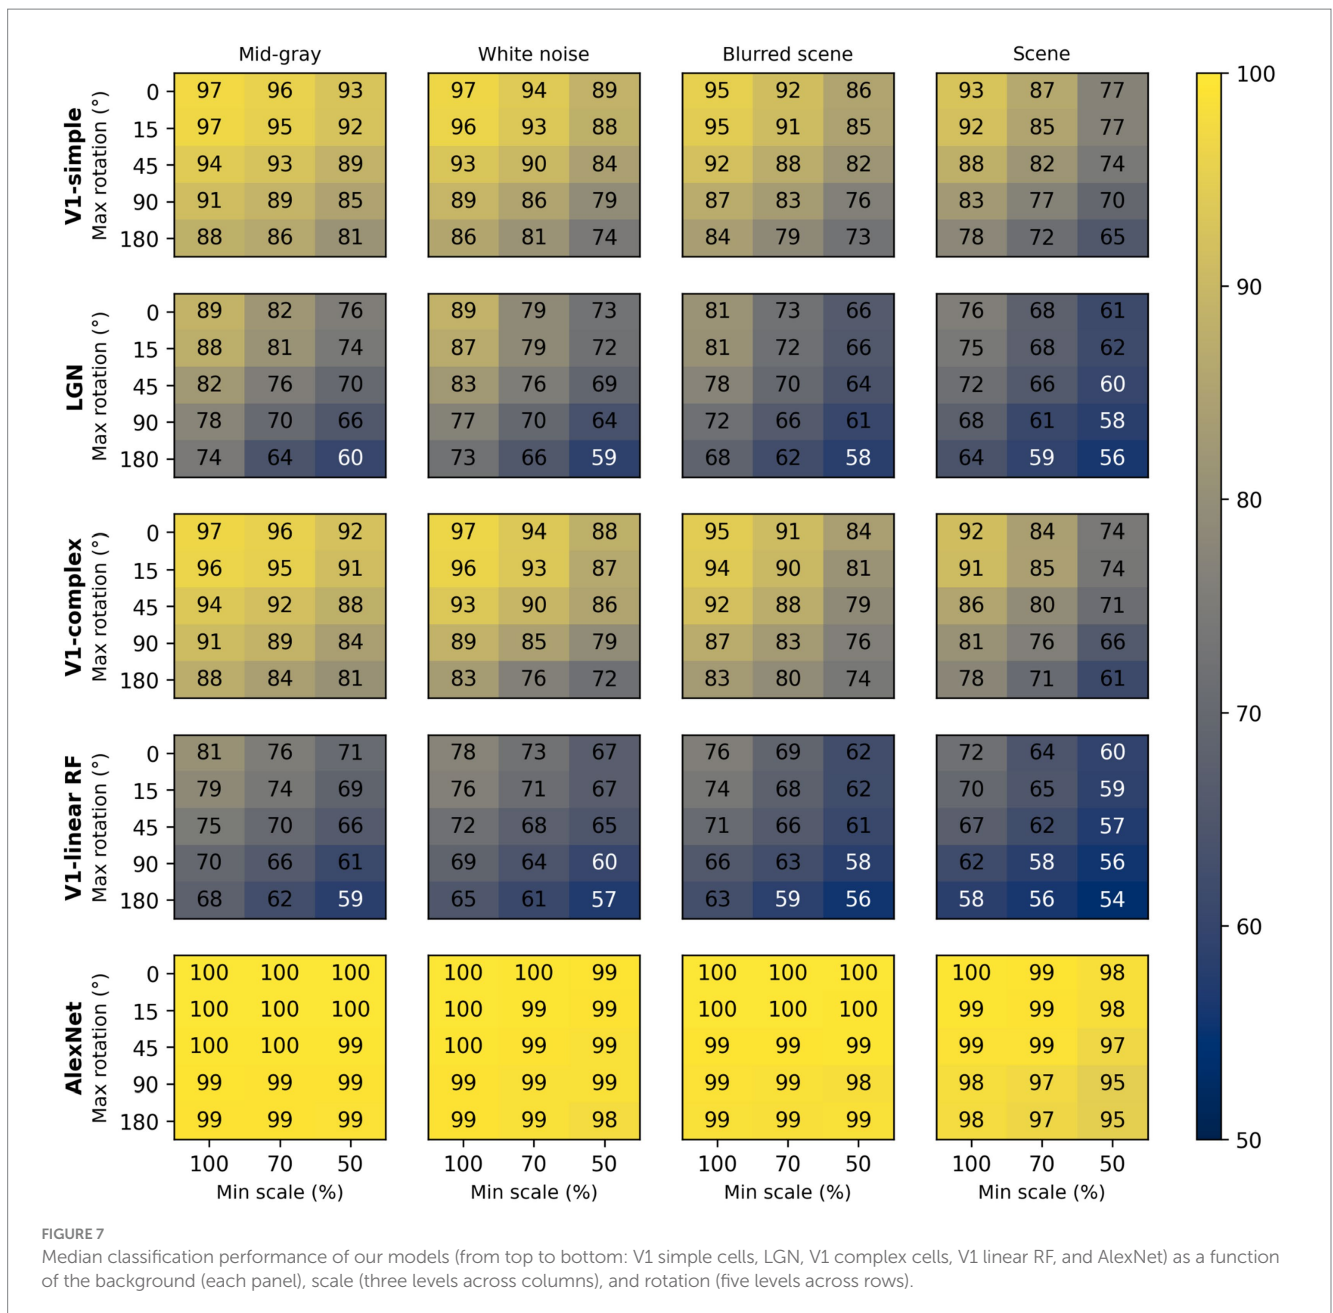
performance. The results are thus fully cross-validated, meaning that no images that were used to train the model were used to evaluate its performance. This entire process was repeated 100 times for each background type, so that we could assess the performance of the model over many different training/testing splits of the image sets. For each background we thus ended up with a distribution of performances.

In Figure 6, we show the distributions of classification performance for the 100 training/testing sets used for each background. Gray backgrounds (either at a fixed mid value, Figure 4 top row, or at the average luminance of the image content), pixelated pink noise, and pixelated noise with luminance values matching the image content, were all associated with similar high performance (median performance 85%, 85%, 84%, and 85%, respectively). Using pixelated white noise as a background (Figure 4, second row) reduced performance considerably (79%). A blurred visual scene (Figure 4, bottom row) or pixelated binary noise (which had higher contrast than all background types we tested, Figure 3) made performance even worse (76% in both cases). Finally, high-contrast scenes (Figure 4, third row) had the most deleterious impact (69%).

To assess the impact of scale and rotation variability, we then selected four of these backgrounds (those shown in Figure 4), and for each of them varied independently scale across three levels (100%, i.e., no variation, 70%, and 50%) and rotation across five levels (0, i.e., no variation, 15°, 45°, 90°, and 180°). Performance degraded as variations in scale and rotation increased (Figure 7, top row), but the degradation was severe only for the largest variations and for the high contrast scene background. Because both faces and objects have a natural orientation, and are rarely tilted more than 45° from it, larger rotations are likely not representative of our daily experience. Similarly, pixelated and high-contrast backgrounds are not common, in part because backgrounds are often far from our plane of fixation and therefore defocused (i.e., have lower contrast) due to the limited depth-of-field of the human eye (Marcos et al., 1999). If we consider scale variations of up to 50%, and rotation variations of up to 45°, we see that with a gray background the median performance is 89%, and even with a







blurred scene it is still a respectable 82%. While these levels of performance would not win any image classification contests, they are probably high enough to account for the preferential looking toward faces observed in human adults (Cerf et al., 2009; Crouzet et al., 2010) and infants (Fantz, 1961; Goren et al., 1975; Johnson et al., 1991), and in monkeys (Gothard et al., 2004; Sugita, 2008; Taubert et al., 2017).

These simulations show that a model with oriented filters like those observed in V1 simple cells is capable of extracting information that is useful, and potentially sufficient, to distinguish a face-like image from something that is unlikely to be a face (a coarse recognition task). We thus wondered if an even simpler model, with circular symmetric filters like those found in retinal ganglion cells and in the lateral geniculate nucleus (LGN), might achieve similar performance. However, when we trained and tested a model with LGN-like filters as we did for the V1-like model, we found (Figure 7, second row) that its

performance was considerably worse, by as much as 20%. Thus, LGN-like filters seem to be much less useful for detecting face-like patterns, although they are sufficient for detecting areas of high contrast.

Next, we wondered how much of the performance of the V1 model could be accounted for by its sensitivity to the spatial phase of its filters, a characteristic of V1 simple cells. We thus implemented a model of V1 complex cells, which are insensitive to spatial phase. We found (Figure 7, third row) that its performance closely matched that of the model of V1 simple cells, underperforming by a small margin across the set of conditions tested, indicating that spatial phase plays only a minor role.

Finally, to estimate the contribution of the rectifying nonlinearity to classification performance, we evaluated a model of V1 simple cells linear receptive fields. This is essentially our model of V1 simple cells, but with only two spatial phases represented at each location (because the other two become redundant) and without the rectifying

nonlinearity. We found that this model underperformed the V1 simple and complex cell models by large margins (Figure 7, fourth row), and was even worse than the LGN model. This clearly highlights the critical importance of nonlinearities, even in models as simple as those presented here, and counters our tendency to intuitively attribute the lion's share of effects to the linear filter.

Some might be surprised by how well these models performed. To provide an additional reference, and a reality check, we also tested a model that has been compared in the past to the primate ventral visual pathway (Cadieu et al., 2014). This model was developed by Krizhevsky et al. (2012) and won the ImageNet image classification contest in 2012, and is usually referred to as AlexNet (after Krizhevsky's first name). We selected a version of this model that was already trained on the ImageNet classification task, in which an image must be categorized as belonging to one of 1,000 categories. We then froze the weights of its feature maps (which can be loosely seen as a cascade of several of the filter layers that we used in our V1- and LGN-like models, but with many more, and more complex, filters), and replaced its classification layers with a linear classifier, identical to the one we used for the V1- and LGN-like models (although with a different number of inputs, corresponding to the number of filters in the last feature layer of AlexNet). Not surprisingly, we found that AlexNet does a much better job compared to our simple V1-like model, and it is much less sensitive to changes in scale, orientation, or background (Figure 7, bottom row). Even with our most challenging condition (up to 50% scaling, up to 180° rotation, high-contrast scene background) its median cross-validated performance is 95% (compared to 65% for the V1-like model); with up to 50% scaling, up to 45° rotation, and a blurred scene background it performed at 99% (compared to 82% for the V1-like model).

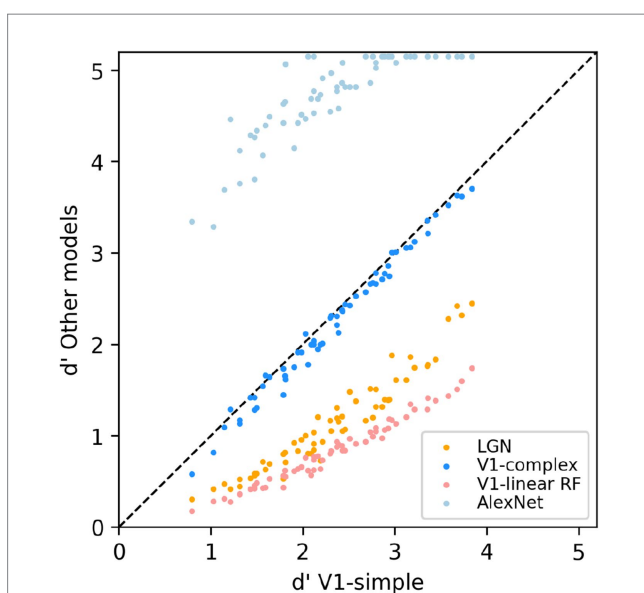
To allow for a more direct comparison between the performance of the various models, in Figure 8 we directly compare the performance of the V1-simple cells model to the that of the other models. Instead of using the percent correct measure that we have use throughout the

paper, here we used the sensitivity index  $d'$  (see Methods), another commonly used metric to characterize binary discrimination (which we clipped at 5.15, corresponding to 99.5% correct). Easy (hard) tasks are associated with high (low) values of  $d'$ . Each data point is associated with one condition from Figure 7 (3 scaling factors, 5 rotation factors, 4 backgrounds, thus 60 points for each model), and we plot on the x-axis the value of  $d'$  observed with the V1-simple cells model and on the y-axis the value of  $d'$  observed with one of the other models (see legend) for the same condition. Points below (above) the diagonal indicate that the V1-simple cells models performed better (worse) than the other model. Disregarding values at or near saturation at high and low values of  $d'$ , when expressed in units of  $d'$  the relative performance of the various model is relatively constant as difficulty varies (i.e., the points for each model are distributed along lines parallel to the main diagonal).

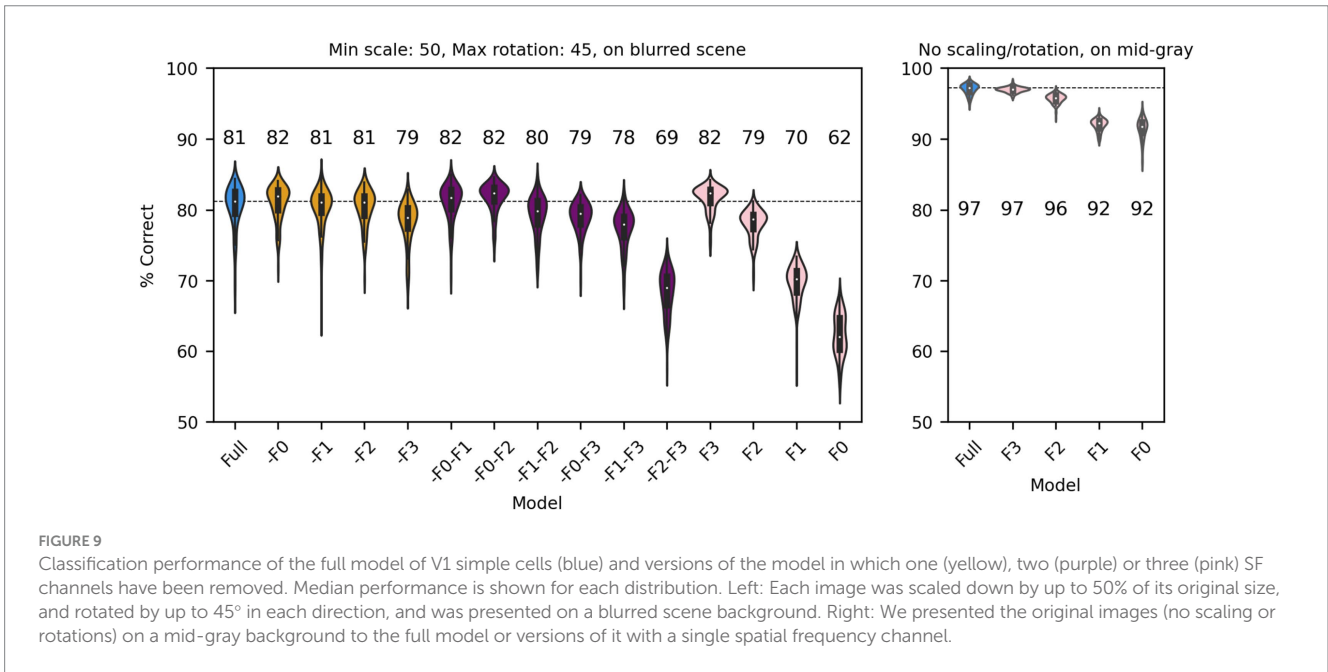
## Probing the V1-like model through selective lesions

So far, we have shown that a V1-like model (whether based on responses of simple or complex cells) contains more (linearly separable) information about distinguishing a face from a non-face object than a LGN-like model, but much less than a modern ANN could extract from the same images. We also showed that the linear receptive field and the output nonlinearity are both critical determinants of performance. We next sought to gain further understanding about which units of the V1-like models are most useful to the classifier. In a manner that is akin to what is done in neuropsychology, we thus simulated "lesions" of the model by removing subsets of model units, and evaluated their impact on the ability of the model to classify our stimuli. Since our V1-like models have four spatial frequency channels, we first eliminated from our model of V1 simple cells one or more of them and refit it to the data, for the same amount of scaling (up to 50%), rotation (up to 45°), and type of background (blurred scene). We found (Figure 9, left) that, compared to the full model (blue), dropping any single channel (yellow) had a minimal impact on performance. Dropping two channels (purple) also had a small impact, unless the two channels tuned to the higher frequencies were both dropped. Finally, when a single channel was kept (pink), we found that the two highest channels by themselves can perform essentially as well as the single model, whereas the two tuned to lower SFs are significantly impaired. The lowest spatial frequencies, which might have been thought to be associated with the overall shape of the face, thus do not play a major role in determining performance under these conditions (although there is the potential confound that the number of filters increases with spatial frequencies, see Methods). When images are presented with no random changes in scale and orientation, and on a gray background, the two lowest SF channels can effectively classify our dataset (Figure 9, right), indicating that their usefulness decreases as image variation increases.

Next, we considered a set of lesions that targeted the most consistent inputs to the classifier. We kept or removed only those filters that in the trained full model of V1 simple cells were associated with a classifier weight that was positive or negative on at least: 90/100 model runs (CS90, 1,686 filters); 95/100 model runs (CS95, 1,162 filters); 99/100 model runs (CS99, 612 filters); or all 100 model runs (CS100, 406 filters). Note that filters were kept (or removed) across all



**FIGURE 8**  
Scatter-plot of the median sensitivity index  $d'$  of our models. The V1 simple cells model is used as a reference and is plotted on the x-axis, and all other models we tested (see legend) are plotted against it (y-axis).



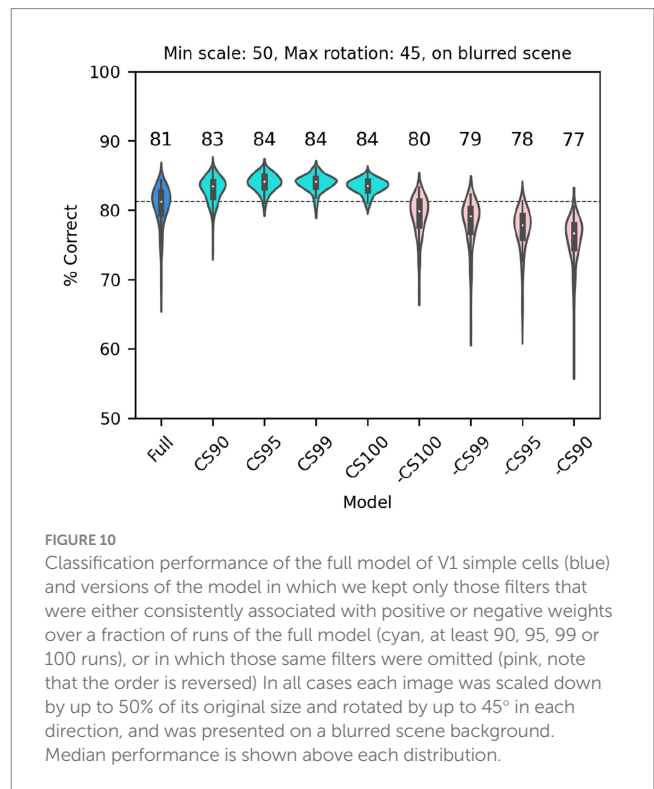
**FIGURE 9** Classification performance of the full model of V1 simple cells (blue) and versions of the model in which one (yellow), two (purple) or three (pink) SF channels have been removed. Median performance is shown for each distribution. Left: Each image was scaled down by up to 50% of its original size, and rotated by up to 45° in each direction, and was presented on a blurred scene background. Right: We presented the original images (no scaling or rotations) on a mid-gray background to the full model or versions of it with a single spatial frequency channel.

SF channels (e.g., for CS90, 22% of filters in F0, 38% of filters in F1, 32% of filters in F2, and 20% of filters in F3 were kept; for CS100 the percentages were 2%, 9%, 10%, 5%, respectively) although in absolute numbers many more filters were kept (or removed) from the higher SF channels. In all cases in which we kept only the most consistent weights (Figure 10, cyan), the model performed as well as the full model (or slightly better than it, probably because filters that contributed mostly noise were omitted). What was more surprising, removing the filters associated with these weights (Figure 10, pink) resulted in only a relatively small drop in performance, indicating that there is a fair amount of redundancy in the model, which allowed the remaining filters to compensate for the loss. However, an inspection of the scatter of model performance across runs (which can be visually estimated by the length of each violin in Figure 10) reveals that keeping only the most relevant weights has a strong impact on the reliability of the classification, presumably by making the classifier easier to train (i.e., less prone to settle on a local minimum).

Applying these same analyses to our model of V1 complex cells yielded similar results (not shown).

### A fine recognition task: gender classification

While it is difficult to understand what exactly the V1-like model responded to (a possibly ill-posed question), the large within-group shape similarity across faces is an obvious candidate (Rosch et al., 1976). To evaluate the V1-like model under more challenging conditions, requiring finer discrimination abilities, we tested it on a gender classification task. We subdivided our face set into two groups (male and female), and then trained the V1-like and the AlexNet models as we did before for the face vs. non-face classification (see Methods). We only tested two backgrounds, mid-gray and blurred scenes. Not surprisingly, cross-validated performance was considerably worse than before (Figure 11), dropping by approximately 25% points for the V1-like model and by 20% point for AlexNet. The



**FIGURE 10** Classification performance of the full model of V1 simple cells (blue) and versions of the model in which we kept only those filters that were either consistently associated with positive or negative weights over a fraction of runs of the full model (cyan, at least 90, 95, 99 or 100 runs), or in which those same filters were omitted (pink, note that the order is reversed) In all cases each image was scaled down by up to 50% of its original size and rotated by up to 45° in each direction, and was presented on a blurred scene background. Median performance is shown above each distribution.

performance of the V1-like model stayed above 70% only for small variations in scale and rotation; at the reference condition we considered before (up to 50% scaling, up to 45° rotation, blurred scene background) the performance was essentially at chance level, compared to the 82% that we observed in the face vs. non-face task.

We also tested FaceNet (Schroff et al., 2015) in this task, a model developed explicitly for face recognition, as opposed to image classification. We expected it to perform much better than the other two models in this task, and indeed its cross-validated classification

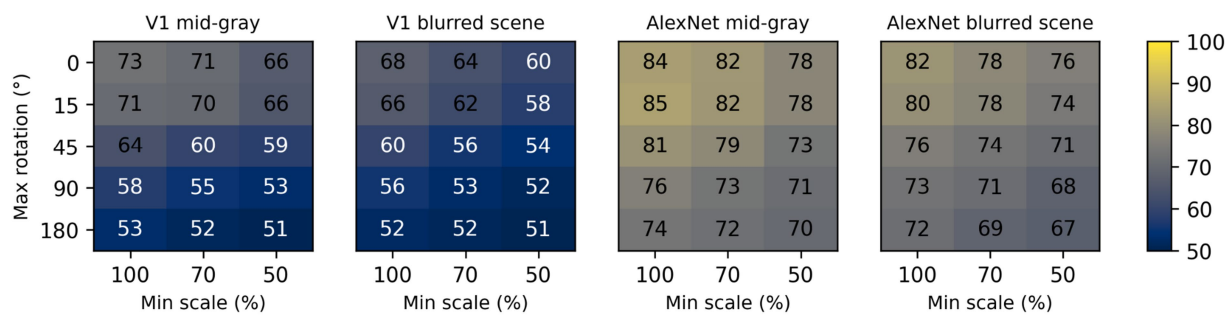


FIGURE 11

Median classification performance of the V1-like (simple cells) and AlexNet models in the gender classification task. Two different backgrounds were tested for each model, varying scale (three levels across the x axis) and rotation (five levels across the y axis) as before.

performance on our data set was 93%. However, there is a catch: This network is not trained to work directly on the images, but instead it first processes the images with another ANN, called MTCNN, which extracts a bounding box containing a face, and then this part of the image is stretched to the size preferred by FaceNet (160 by 160 px) and fed to it. There is thus a single scale, no rotation, and very little background. A fairer comparison would thus involve feeding to our two other models the same images that we fed to FaceNet. When we did that the V1-like model performed at 91% correct, and AlexNet matched FaceNet at 93% correct. Note that these performances are much better than those we obtained with no scaling or rotation on our image set on a mid-gray background (73% and 84%, respectively). What accounts for this difference? One possibility is the stretching of the images operated by MTCNN. Another possibility is the lack of jitter in the images—recall that in all previous tasks the images were always randomly jittered by up to 42 px in all directions. To find the answer we again classified the original images with the V1-like and AlexNet models, this time omitting jittering, for the no scaling, no rotation, and mid-gray background condition. The cross-validated performance of both models was now 90%, indicating that the lack of jitter accounts for the vast majority of the difference observed. Obviously the presence of jitter limits, for both our V1-like model and for AlexNet, the ability to perform fine (but not coarse) discriminations.

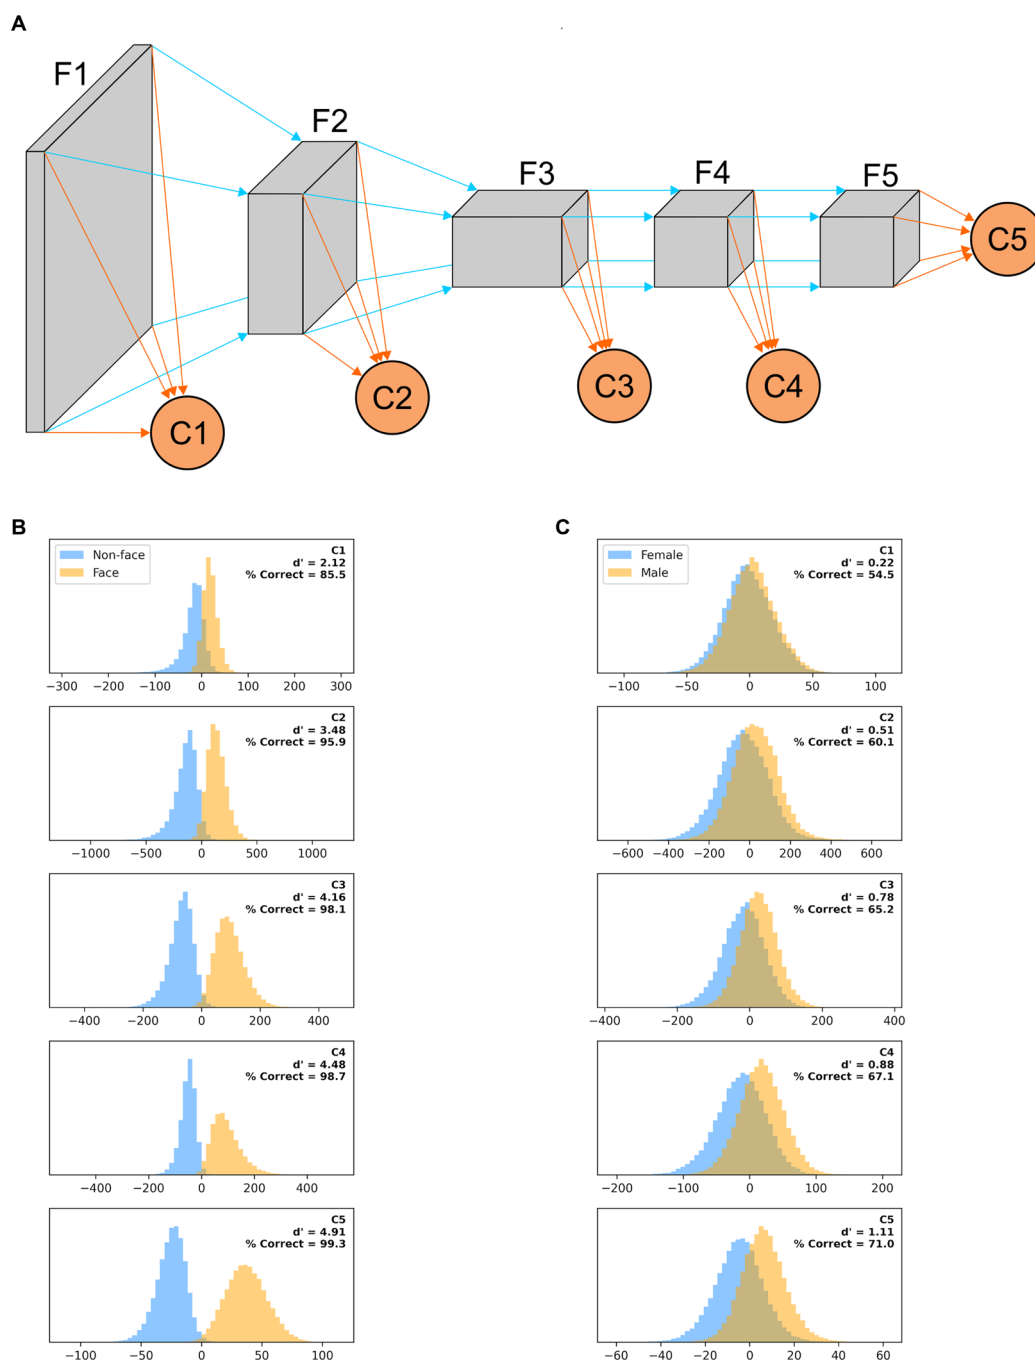
## Object recognition along the visual hierarchy

We have shown here how recognition performance differs, for two tasks, in models of the earliest stages of visual processing (LGN, V1 simple and complex cells) and in a widely used ANN, whose final stage has been likened to IT processing. It would of course be interesting to understand how performance varies along the entire ventral visual pathway. Unfortunately, there is not yet consensus on how the properties of V2, V3, V4, and IT emerge from their inputs. However, since it has been argued that there is a hierarchical similarity, in terms of feature selectivity of individual units, between layers of ANNs and areas of the primate visual system, we evaluated how classifiers trained on the features extracted by intermediate layers of AlexNet fared in our tasks (Figure 12A). While direct comparisons

have not been carried out, in our multi-tail version of AlexNet one can roughly consider the stack of layers F1 as akin to area V1, F2 to V2, F3 to V3, F4 to V4, and F5 to IT. In Figure 12B we plot, for a model trained on a face vs. non-face classification task (with images presented on a blurred background, scaled by up to 50% and rotated by up to 45°), the distribution of classifier outputs when presented with non-faces (blue) or faces (orange) from the testing image set. While performance improved as we moved deeper in the feature stack, classifier C1 already provided actionable information, and going from C2 to C5 provided a very limited improvement in performance for the additional computational load (for computers) and time (for the brain and computers) required. In Figure 12C, we plot the results for a model trained on a gender discrimination task (with images presented on a blurred background, scaled by up to 50% and rotated by up to 45°). This is a much harder task, and shortcuts incur significant costs, making waiting for the output of C5 a sensible strategy. For all their limitations, these simulations again highlight the ability of earlier stages of processing to perform adequately in tasks requiring coarse, but not fine, categorization.

## Discussion

The last 10 years have seen a resurgence in the study of object recognition, both in computer science and in primate neuroscience. What often goes unacknowledged is that almost all the focus is on fine discrimination, or what in primates could be termed foveal object recognition. We have shown here that for simpler discrimination tasks signals in early visual areas are sufficient to achieve reasonable accuracy. In itself this is not a new finding (Serre et al., 2007), but it is one that has been sidelined once the limitations of early visual areas to support fine object recognition under large variations in orientation, illumination, and background (such as recognizing the image of an elephant upside-down over a high contrast skyline of New York City) were revealed (Pinto et al., 2008; Rajalingham et al., 2018). Our main contribution here is to more accurately describe these limitations for a task that is behaviorally important and that had not been considered in previous studies. Furthermore, by comparing the performance of different variations of our models we were able to highlight the relative contributions of their various parts.



**FIGURE 12** (A) A “multi-tail” (our nomenclature) version of AlexNet, in which the feature stack has been divided into five stages, and a binary classification layer is appended to each stage (see Methods). (B) Distribution of each classifier outputs when presented with non-faces (blue) or faces (orange) from the testing image set for a model trained on a face vs. non-face classification task with images presented on a blurred background, scaled by up to 50% and rotated by up to 45°. The outputs are collated across 100 different training/testing splits of the data. (C) Same as panel (B), but for a model trained on a gender discrimination task (with images presented on a blurred background, scaled by up to 50% and rotated by up to 45°).

Unlike computer science models, the brain of primates must contend with visual inputs that do not uniformly sample the visual field, making object recognition in the periphery an altogether different problem from object recognition in the fovea. There is no doubt that foveal vision is critical to primates’ ability to survive and thrive, as clearly demonstrated by the fact that up to 30% of area V1 is

devoted to this region (Perry and Cowey, 1985; Azzopardi and Cowey, 1996), and by the devastating effects of macular degeneration (Fleckenstein et al., 2021). However, some fundamental aspects of object recognition are overlooked when the focus is on foveally presented objects in an unusual pose against an out-of-context high contrast background (what has been termed core object recognition,

DiCarlo and Cox, 2007; DiCarlo et al., 2012). Most of our daily experience is quite different, and involves quickly localizing familiar objects in familiar environments, across the visual field (Eckstein, 2017). Because the fovea covers less than one part in 5,000 of our visual field, without a fairly sophisticated system to detect objects in the periphery, our ability to make good use of IT's capacity for fine object discrimination would be severely hindered.

Recognizing that an object of interest is somewhere in the periphery is not sufficient: its precise localization is critical. Hence, peripheral recognition should be based on neurons that carry position information. It has been argued that, in tasks involving difficult discriminations of objects on high contrast backgrounds, the available position information increases along the visual stream (Hong et al., 2016), and is thus higher in IT than in V1. However, that study used a rate code to extract such information (where a larger value of a single output indicates a more eccentric position); such a code would be difficult to use for controlling behavior, as areas of the salience network (such as LIP, FEF, and SC) are retinotopically organized and interconnected. In contrast, a retinotopic code (where different cells are active to indicate the presence of an object at different retinal locations) would be much more readily usable, but this information is progressively diluted deeper (i.e., more anteriorly) in the hierarchy of visual areas. Object recognition based on signals in retinotopically organized early visual areas is thus not only possible (as we have shown here) but would also be eminently useful in primates. Importantly, early visual areas do project in a retinotopically organized manner to areas of the salience networks, such as the SC (Cerkevich et al., 2014). In afoveated animals, which devote to vision a smaller fraction of their brain, this kind of rudimentary object recognition might very well comprise the entirety of visual object recognition, as they might rely more on motion signals and other sensory modalities to direct their behavior (motion is of course an important visual cue also for primates, but it normally plays a secondary role relative to form vision, with some important exceptions, e.g., breaking camouflage, biological motion).

Here we focused on face recognition, because it is one of the most widely studied visual classification tasks, and it is one at which primates excel (Rosenfeld and Van Hoesen, 1979; Diamond and Carey, 1986; Carey, 1992; Pascalis and Bachevalier, 1998; Pascalis et al., 1999). It is also notable that a preference for looking at faces is already present in newborns (Goren et al., 1975; Pascalis and de Schonen, 1994), well before the development of the cortical machinery that supports face identification in older children and adults, the so-called face-patch network (Livingstone et al., 2017). This observation had led to the suggestion that preferentially looking toward faces might initially be mediated by direct projections from the retina to the SC, with cortical projections to the SC taking over during development (Morton and Johnson, 1991). Indeed, SC cells have been found to respond more strongly to images of faces than non-faces in adult monkeys (Nguyen et al., 2014; Yu et al., 2024). Our finding that circular center-surround filters like those found in retinal ganglion cells and LGN are not well-suited for this task argues against a retinal source for this information. However, because it is difficult to calibrate the classification performance of our models against the preferential looking methods used when studying infants, it remains possible that even the limited performance of our LGN model might be able to account for the preferences of infants.

We tested our models also on gender classification. We chose this task because it requires discrimination of fine features, and yet it has

been shown that it requires less time than face identification (Dobs et al., 2019), and it is preserved in prosopagnosic patients (Sergent et al., 1992), whose cortical face processing network (Tsao et al., 2003; Kanwisher and Yovel, 2006) is known to be disrupted (Rosenthal et al., 2017). These findings point to gender classification as possibly relying on visual cortical signals that precede the face-network in IT. Our results indicate that V1 is an unlikely source of such signals. However, since both texture and curvature have been proposed to play important roles in gender discrimination (Brown and Perrett, 1993; Bruce et al., 1993; Hole and Bourne, 2010), area V2, which is specialized for processing naturalistic textures (Freeman et al., 2013; Movshon and Simoncelli, 2014), and area V4, which is specialized for processing curvature (Pasupathy and Connor, 2002; Nandy et al., 2013), might carry useful signals. V2 and V4, like V1, project directly to the SC (May, 2006), and would thus also be well-positioned to bias refixations.

## A different architecture for object recognition

The ideas about object classification and recognition that dominate the field, with their emphasis on foveal inputs and context-independent processing, are certainly important, but they omit some crucial functional aspects of how our specialized primate visual system identifies objects. A unitary hierarchical architecture of ever more complex feature detectors with a read-out/classification stage at its end, similar in structure to ANNs for object recognition (with only C5 in Figure 12), might work well in computer science, but it is a brittle architecture, at odds with evolutionary principles and experimental evidence. We envision a parallel architecture, in which classification is based on signals in multiple visual areas (as in Figure 12), but, unlike in ANNs, with learning for coarse recognition separately at each retinal location. Such a system would be more flexible, has strong redundancy, and it is easier to see how it could have evolved from simpler systems. Information from some stages might be available earlier than that from others, but might also be less reliable, yielding a speed/accuracy trade-off that could be easily arbitrated, as we saw in our multi-tail version of AlexNet (Figure 12). The information coming from each area could be weighted differently in the center and in the periphery, and also by pathways involved in action (e.g., controlling refixation through the saliency network) vs. perception. Such a system would be more flexible than a classic ANN-like architecture, in which the loss of any part/stage would be devastating, requiring retraining of the entire model (as in our simulated lesions). It would be akin to the “wisdom of the crowd” (Aristotle, 2013; Saha Roy et al., 2021; Madirolas et al., 2022) or “ensemble learning” (Hansen and Salamon, 1990; Schapire, 1990; Polikar, 2006) architectures, in which a highly complex decision process is replaced by many simpler ones, each evaluating the evidence independently and reaching its own conclusion. If any of the processes becomes unavailable, the others carry on, as is observed in real life (Matsumoto et al., 2016).

## Data availability statement

The original contributions presented in the study are included in the article/supplementary material, further inquiries can be directed to the corresponding author.

## Author contributions

CQ: Conceptualization, Data curation, Formal analysis, Methodology, Resources, Software, Visualization, Writing – original draft. RK: Conceptualization, Funding acquisition, Project administration, Supervision, Writing – review & editing.

## Funding

The author(s) declare that financial support was received for the research, authorship, and/or publication of this article. This work was supported by National Eye Institute Intramural Research Program.

## Acknowledgments

The authors thank Doris Tsao for kindly sharing with us a carefully curated diverse set of frontal images of human faces.

## References

- Afraz, A., Boyden, E. S., and DiCarlo, J. J. (2015). Optogenetic and pharmacological suppression of spatial clusters of face neurons reveal their causal role in face gender discrimination. *Proc. Natl. Acad. Sci. U. S. A.* 112, 6730–6735. doi: 10.1073/pnas.1423281112
- Afraz, S.-R., Kiani, R., and Esteky, H. (2006). Microstimulation of inferotemporal cortex influences face categorization. *Nature* 442, 692–695. doi: 10.1038/nature04982
- Arcaro, M. J., Schade, P. F., Vincent, J. L., Ponce, C. R., and Livingstone, M. S. (2017). Seeing faces is necessary for face-domain formation. *Nat. Neurosci.* 20, 1404–1412. doi: 10.1038/nn.4635
- Aristotle, A. (2013). Aristotle's politics. Chicago, IL: University of Chicago Press.
- Azzopardi, P., and Cowey, A. (1996). The overrepresentation of the fovea and adjacent retina in the striate cortex and dorsal lateral geniculate nucleus of the macaque monkey. *Neuroscience* 72, 627–639. doi: 10.1016/0306-4522(95)00589-7
- Bishop, C. M. (1996). Neural networks for pattern recognition. 1st Edn. Oxford: Oxford University Press, USA.
- Bradski, G., and Kaehler, A. (2008). *Learning OpenCV*. Sebastopol, CA: O'Reilly Media.
- Brown, E., and Perrett, D. I. (1993). What gives a face its gender? *Perception* 22, 829–840. doi: 10.1068/p220829
- Bruce, V., Burton, A. M., Hanna, E., Healey, P., Mason, O., Coombes, A., et al. (1993). Sex discrimination: how do we tell the difference between male and female faces? *Perception* 22, 131–152. doi: 10.1068/p220131
- Cadiou, C. F., Hong, H., Yamins, D. L. K., Pinto, N., Ardila, D., Solomon, E. A., et al. (2014). Deep neural networks rival the representation of primate IT cortex for core visual object recognition. *PLoS Comput. Biol.* 10:e1003963. doi: 10.1371/journal.pcbi.1003963
- Carandini, M., Demb, J. B., Mante, V., Tolhurst, D. J., Dan, Y., Olshausen, B. A., et al. (2005). Do we know what the early visual system does? *J. Neurosci.* 25, 10577–10597. doi: 10.1523/JNEUROSCI.3726-05.2005
- Carey, S. (1992). Becoming a face expert. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 335, 95–102. doi: 10.1098/rstb.1992.0012
- Cerf, M., Frady, E. P., and Koch, C. (2009). Faces and text attract gaze independent of the task: experimental data and computer model. *J. Vis.* 9:10.1-15. doi: 10.1167/9.12.10
- Cerkevich, C. M., Lyon, D. C., Balaram, P., and Kaas, J. H. (2014). Distribution of cortical neurons projecting to the superior colliculus in macaque monkeys. *Eye Brain* 2014, 121–137. doi: 10.2147/EB.S53613
- Crouzet, S. M., Kirchner, H., and Thorpe, S. J. (2010). Fast saccades toward faces: face detection in just 100 ms. *J. Vis.* 10:16. doi: 10.1167/10.4.16
- Cusick, C. G. (1988). "Chapter 1: anatomical organization of the superior colliculus in monkeys: corticotectal pathways for visual and visuomotor functions" in *Vision within Extrageniculate-striate systems Progress in brain research*. eds. T. P. Hicks and G. Benedek (Amsterdam, Netherlands: Elsevier), 1–15.
- De Beeck, H. O., and Vogels, R. (2000). Spatial sensitivity of macaque inferior temporal neurons. *J. Comp. Neurol.* 426, 505–518. doi: 10.1002/1096-9861(20001030)426:4<505::aid-cne1>3.0.co;2-m
- Desimone, R., Albright, T. D., Gross, C. G., and Bruce, C. (1984). Stimulus-selective properties of inferior temporal neurons in the macaque. *J. Neurosci.* 4, 2051–2062. doi: 10.1523/JNEUROSCI.04-08-02051.1984
- Diamond, R., and Carey, S. (1986). Why faces are and are not special: an effect of expertise. *J. Exp. Psychol. Gen.* 115, 107–117. doi: 10.1037/0096-3445.115.2.107
- DiCarlo, J. J., and Cox, D. D. (2007). Untangling invariant object recognition. *Trends Cogn. Sci.* 11, 333–341. doi: 10.1016/j.tics.2007.06.010
- DiCarlo, J. J., and Maunsell, J. H. R. (2003). Anterior inferotemporal neurons of monkeys engaged in object recognition can be highly sensitive to object retinal position. *J. Neurophysiol.* 89, 3264–3278. doi: 10.1152/jn.00358.2002
- DiCarlo, J. J., Zoccolan, D., and Rust, N. C. (2012). How does the brain solve visual object recognition? *Neuron* 73, 415–434. doi: 10.1016/j.neuron.2012.01.010
- Dobs, K., Isik, L., Pantazis, D., and Kanwisher, N. (2019). How face perception unfolds over time. *Nat. Commun.* 10:1258. doi: 10.1038/s41467-019-09239-1
- Drewes, J., Trommershäuser, J., and Gegenfurtner, K. R. (2011). Parallel visual search and rapid animal detection in natural scenes. *J. Vis.* 11:20. doi: 10.1167/11.2.20
- Eckstein, M. P. (2017). Probabilistic computations for attention, eye movements, and search. *Annu. Rev. Vis. Sci.* 3, 319–342. doi: 10.1146/annurev-vision-102016-061220
- Einhäuser, W., Spain, M., and Perona, P. (2008). Objects predict fixations better than early saliency. *J. Vis.* 8:18. doi: 10.1167/8.14.18
- Eldridge, M. A., Matsumoto, N., Wittig, J. H., Masseau, E. C., Saunders, R. C., and Richmond, B. J. (2018). Perceptual processing in the ventral visual stream requires area TE but not rhinal cortex. *eLife* 7:310. doi: 10.7554/eLife.36310
- Enroth-Cugell, C., and Robson, J. G. (1966). The contrast sensitivity of retinal ganglion cells of the cat. *J. Physiol. (Lond.)* 187, 517–552. doi: 10.1113/jphysiol.1966.sp008107
- Fant, R. L. (1961). The origin of form perception. *Sci. Am.* 204, 66–73. doi: 10.1038/scientificamerican0561-66
- Fecteau, J. H., and Munoz, D. P. (2006). Saliency, relevance, and firing: a priority map for target selection. *Trends Cogn. Sci.* 10, 382–390. doi: 10.1016/j.tics.2006.06.011
- Fleckenstein, M., Keenan, T. D. L., Guymer, R. H., Chakravarthy, U., Schmitz-Valckenberg, S., Klaver, C. C., et al. (2021). Age-related macular degeneration. *Nat. Rev. Dis. Primers.* 7:31. doi: 10.1038/s41572-021-00265-2
- Freeman, J., Ziemba, C. M., Heeger, D. J., Simoncelli, E. P., and Movshon, J. A. (2013). A functional and perceptual signature of the second visual area in primates. *Nat. Neurosci.* 16, 974–981. doi: 10.1038/nn.3402
- Fries, W. (1984). Cortical projections to the superior colliculus in the macaque monkey: a retrograde study using horseradish peroxidase. *J. Comp. Neurol.* 230, 55–76. doi: 10.1002/cne.902300106
- Goodale, M. A., and Milner, A. D. (1992). Separate visual pathways for perception and action. *Trends Neurosci.* 15, 20–25. doi: 10.1016/0166-2236(92)90344-8
- Goren, C. C., Sarty, M., and Wu, P. Y. (1975). Visual following and pattern discrimination of face-like stimuli by newborn infants. *Pediatrics* 56, 544–549. doi: 10.1542/peds.56.4.544

We thank Leor Katz, Gongchen Yu, and Corey Ziemba for many insightful discussions and for comments on a previous version of the manuscript.

## Conflict of interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

## Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

- Gothard, K. M., Erickson, C. A., and Amaral, D. G. (2004). How do rhesus monkeys (*Macaca mulatta*) scan faces in a visual paired comparison task? *Anim. Cogn.* 7, 25–36. doi: 10.1007/s10071-003-0179-6
- Gottlieb, J. P., Kusunoki, M., and Goldberg, M. E. (1998). The representation of visual salience in monkey parietal cortex. *Nature* 391, 481–484. doi: 10.1038/35135
- Gross, C. G. (2008). Single neuron studies of inferior temporal cortex. *Neuropsychologia* 46, 841–852. doi: 10.1016/j.neuropsychologia.2007.11.009
- Gross, C. G., Rocha-Miranda, C. E., and Bender, D. B. (1972). Visual properties of neurons in inferotemporal cortex of the macaque. *J. Neurophysiol.* 35, 96–111. doi: 10.1152/jn.1972.35.1.96
- Hansen, L. K., and Salamon, P. (1990). Neural network ensembles. *IEEE Trans. Pattern Anal. Mach. Intell.* 12, 993–1001. doi: 10.1109/34.58871
- He, K., Zhang, X., Ren, S., and Sun, J. (2015). Delving deep into rectifiers: surpassing human-level performance on ImageNet classification. arXiv [Preprint].
- Hole, G., and Bourne, V. (2010). Face processing: Psychological, neuropsychological, and applied perspectives. 1st Edn. New York: Oxford University Press.
- Hong, H., Yamins, D. L. K., Majaj, N. J., and DiCarlo, J. J. (2016). Explicit information for category-orthogonal object properties increases along the ventral stream. *Nat. Neurosci.* 19, 613–622. doi: 10.1038/nn.4247
- Hopfield, J. J. (1987). Learning algorithms and probability distributions in feed-forward and feed-back networks. *Proc. Natl. Acad. Sci. U. S. A.* 84, 8429–8433. doi: 10.1073/pnas.84.23.8429
- Hubel, D. H., and Wiesel, T. N. (1962). Receptive fields, binocular interaction and functional architecture in the cat's visual cortex. *J. Physiol. (Lond.)* 160, 106–154. doi: 10.1113/jphysiol.1962.sp006837
- Hung, C. P., Kreiman, G., Poggio, T., and DiCarlo, J. J. (2005). Fast readout of object identity from macaque inferior temporal cortex. *Science* 310, 863–866. doi: 10.1126/science.1117593
- Itti, L., and Koch, C. (2000). A saliency-based search mechanism for overt and covert shifts of visual attention. *Vision Res.* 40, 1489–1506. doi: 10.1016/s0042-6989(99)00163-7
- Itti, L., Koch, C., and Niebur, E. (1998). A model of saliency-based visual attention for rapid scene analysis. *IEEE Trans. Pattern Anal. Mach. Intell.* 20, 1254–1259. doi: 10.1109/34.730558
- Johnson, M. H., Dziurawiec, S., Ellis, H., and Morton, J. (1991). Newborns' preferential tracking of face-like stimuli and its subsequent decline. *Cognition* 40, 1–19. doi: 10.1016/0010-0277(91)90045-6
- Jones, J. P., and Palmer, L. A. (1987a). An evaluation of the two-dimensional Gabor filter model of simple receptive fields in cat striate cortex. *J. Neurophysiol.* 58, 1233–1258. doi: 10.1152/jn.1987.58.6.1233
- Jones, J. P., and Palmer, L. A. (1987b). The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *J. Neurophysiol.* 58, 1187–1211. doi: 10.1152/jn.1987.58.6.1187
- Kaas, J. H. (2020). "Evolution of visual cortex in primates" in *Evolutionary neuroscience* ed. J. H. Kass (London, UK: Academic Press) 547–564.
- Kanwisher, N., and Yovel, G. (2006). The fusiform face area: a cortical region specialized for the perception of faces. *Philos. Trans. R. Soc. Lond. B Biol. Sci.* 361, 2109–2128. doi: 10.1098/rstb.2006.1934
- Khaligh-Razavi, S.-M., and Kriegeskorte, N. (2014). Deep supervised, but not unsupervised, models may explain IT cortical representation. *PLoS Comput. Biol.* 10:e1003915. doi: 10.1371/journal.pcbi.1003915
- King, D. E. (2009). Dlib-ml: a machine learning toolkit. *J. Machine Learn Res* 10, 1755–1758. doi: 10.5555/1577069.1755843
- Kirk, E. C., and Kay, R. F. (2004). "The evolution of high visual acuity in the anthropoidea" in *Anthropoid origins*. eds. C. F. Ross and R. F. Kay (Boston, MA: Springer US), 539–602.
- Krauzlis, R. J. (2005). The control of voluntary eye movements: new perspectives. *Neuroscientist* 11, 124–137. doi: 10.1177/1073858404271196
- Krauzlis, R. J., Lovejoy, L. P., and Zénon, A. (2013). Superior colliculus and visual spatial attention. *Annu. Rev. Neurosci.* 36, 165–182. doi: 10.1146/annurev-neuro-062012-170249
- Kriegeskorte, N. (2015). Deep neural networks: a new framework for modeling biological vision and brain information processing. *Annu. Rev. Vis. Sci.* 1, 417–446. doi: 10.1146/annurev-vision-082114-035447
- Krizhevsky, A., Sutskever, I., and Hinton, G. E. (2012). ImageNet classification with deep convolutional neural networks. *Commun. ACM* 60, 84–90. doi: 10.1145/3065386
- Leopold, D. A., Mitchell, J. F., and Freiwald, W. A. (2020). "Evolved mechanisms of high-level visual perception in Primates," in *Evolutionary Neuroscience*, ed. J. H. Kass (London, UK: Academic Press) 589–625.
- Lian, Y., Almasi, A., Grayden, D. B., Kameneva, T., Burkitt, A. N., and Meffin, H. (2021). Learning receptive field properties of complex cells in V1. *PLoS Comput. Biol.* 17:e1007957. doi: 10.1371/journal.pcbi.1007957
- Livingstone, M. S., Vincent, J. L., Arcaro, M. J., Srihasam, K., Schade, P. F., and Savage, T. (2017). Development of the macaque face-patch system. *Nat. Commun.* 8:14897. doi: 10.1038/ncomms14897
- Logothetis, N. K., and Sheinberg, D. L. (1996). Visual object recognition. *Annu. Rev. Neurosci.* 19, 577–621. doi: 10.1146/annurev.ne.19.030196.003045
- Madirolas, G., Zaghi-Lara, R., Gomez-Marin, A., and Pérez-Escudero, A. (2022). The motor wisdom of the crowd. *J. R. Soc. Interface* 19:20220480. doi: 10.1098/rsif.2022.0480
- Mäkelä, P., Näsänen, R., Rovamo, J., and Melmoth, D. (2001). Identification of facial images in peripheral vision. *Vision Res.* 41, 599–610. doi: 10.1016/s0042-6989(00)00259-5
- Marcelja, S. (1980). Mathematical description of the responses of simple cortical cells. *J. Opt. Soc. Am.* 70, 1297–1300. doi: 10.1364/josa.70.001297
- Marcos, S., Moreno, E., and Navarro, R. (1999). The depth-of-field of the human eye from objective and subjective measurements. *Vision Res.* 39, 2039–2049. doi: 10.1016/s0042-6989(98)00317-4
- Martinez, L. M., and Alonso, J.-M. (2003). Complex receptive fields in primary visual cortex. *Neuroscientist* 9, 317–331. doi: 10.1177/1073858403252732
- Matsumoto, N., Eldridge, M. A. G., Fredericks, J. M., Lowe, K. A., and Richmond, B. J. (2022). Comparing performance between a deep neural network and monkeys with bilateral removals of visual area TE in categorizing feature-ambiguous stimuli. *J. Comput. Neurosci.* 51, 381–387. doi: 10.1007/s10827-023-00854-y
- Matsumoto, N., Eldridge, M. A. G., Saunders, R. C., Reoli, R., and Richmond, B. J. (2016). Mild perceptual categorization deficits follow bilateral removal of anterior inferior temporal cortex in rhesus monkeys. *J. Neurosci.* 36, 43–53. doi: 10.1523/JNEUROSCI.2058-15.2016
- May, P. J. (2006). The mammalian superior colliculus: laminar structure and connections. *Prog. Brain Res.* 151, 321–378. doi: 10.1016/S0079-6123(05)51011-2
- McKone, E. (2004). Isolating the special component of face recognition: peripheral identification and a Mooney face. *J. Exp. Psychol. Learn. Mem. Cogn.* 30, 181–197. doi: 10.1037/0278-7393.30.1.181
- McPeck, R. M., and Keller, E. L. (2002). Saccade target selection in the superior colliculus during a visual search task. *J. Neurophysiol.* 88, 2019–2034. doi: 10.1152/jn.2002.88.4.2019
- Moeller, S., Crapse, T., Chang, L., and Tsao, D. Y. (2017). The effect of face patch microstimulation on perception of faces and objects. *Nat. Neurosci.* 20, 743–752. doi: 10.1038/nn.4527
- Morton, J., and Johnson, M. H. (1991). CONSPEC and CONLERN: a two-process theory of infant face recognition. *Psychol. Rev.* 98, 164–181. doi: 10.1037/0033-295x.98.2.164
- Movshon, J. A., and Simoncelli, E. P. (2014). Representation of naturalistic image structure in the primate visual cortex. *Cold Spring Harb. Symp. Quant. Biol.* 79, 115–122. doi: 10.1101/sqb.2014.79.024844
- Movshon, J. A., Thompson, I. D., and Tolhurst, D. J. (1978). Receptive field organization of complex cells in the cat's striate cortex. *J. Physiol. (Lond.)* 283, 79–99. doi: 10.1113/jphysiol.1978.sp012489
- Najemnik, J., and Geisler, W. S. (2008). Eye movement statistics in humans are consistent with an optimal search strategy. *J. Vis.* 8:4. doi: 10.1167/8.3.4
- Nandy, A. S., Sharpee, T. O., Reynolds, J. H., and Mitchell, J. F. (2013). The fine structure of shape tuning in area V4. *Neuron* 78, 1102–1115. doi: 10.1016/j.neuron.2013.04.016
- Nelson, C. A. (2001). The development and neural bases of face recognition. *Infant Child Dev.* 10, 3–18. doi: 10.1002/icd.239
- Nelson, C. A., and Ludemann, P. M. (1989). Past, current, and future trends in infant face perception research. *Can. J. Psychol.* 43, 183–198. doi: 10.1037/h0084221
- Nguyen, M. N., Matsumoto, J., Hori, E., Maior, R. S., Tomaz, C., Tran, A. H., et al. (2014). Neuronal responses to face-like and facial stimuli in the monkey superior colliculus. *Front. Behav. Neurosci.* 8:85. doi: 10.3389/fnbeh.2014.00085
- Op de Beeck, H. P., Pillot, I., and Ritchie, J. B. (2019). Factors determining where category-selective areas emerge in visual cortex. *Trends Cogn. Sci.* 23, 784–797. doi: 10.1016/j.tics.2019.06.006
- Pascalis, O., and Bachevalier, J. (1998). Face recognition in primates: a cross-species study. *Behav. Processes* 43, 87–96. doi: 10.1016/s0376-6357(97)00090-9
- Pascalis, O., and de Schonen, S. (1994). Recognition memory in 3- to 4-day-old human neonates. *Neuroreport* 5, 1721–1724. doi: 10.1097/00001756-199409080-00008
- Pascalis, O., Petit, O., Kim, J. H., and Campbell, R. (1999). Picture perception in primates: the case of face recognition. *Curr. Psychol. Cogn.* 18, 901–933.
- Pasupathy, A., and Connor, C. E. (2002). Population coding of shape in area V4. *Nat. Neurosci.* 5, 1332–1338. doi: 10.1038/972
- Perry, V. H., and Cowey, A. (1985). The ganglion cell and cone distributions in the monkey's retina: implications for central magnification factors. *Vision Res.* 25, 1795–1810. doi: 10.1016/0042-6989(85)90004-5
- Pinto, N., Cox, D. D., and DiCarlo, J. J. (2008). Why is real-world visual object recognition hard? *PLoS Comput. Biol.* 4:e27. doi: 10.1371/journal.pcbi.0040027
- Polikar, R. (2006). Ensemble based systems in decision making. *IEEE Circuits Syst. Mag.* 6, 21–45. doi: 10.1109/MCAS.2006.1688199
- Potter, M. C. (1975). Meaning in visual search. *Science* 187, 965–966. doi: 10.1126/science.1145183



- Rajalingham, R., Issa, E. B., Bashivan, P., Kar, K., Schmidt, K., and DiCarlo, J. J. (2018). Large-scale, high-resolution comparison of the Core visual object recognition behavior of humans, monkeys, and state-of-the-art deep artificial neural networks. *J. Neurosci.* 38, 7255–7269. doi: 10.1523/JNEUROSCI.0388-18.2018
- Ringach, D. L. (2002). Spatial structure and symmetry of simple-cell receptive fields in macaque primary visual cortex. *J. Neurophysiol.* 88, 455–463. doi: 10.1152/jn.2002.88.1.455
- Robinson, D. L., and Petersen, S. E. (1992). The pulvinar and visual salience. *Trends Neurosci.* 15, 127–132. doi: 10.1016/0166-2236(92)90354-b
- Robson, J. G., and Graham, N. (1981). Probability summation and regional variation in contrast sensitivity across the visual field. *Vision Res.* 21, 409–418. doi: 10.1016/0042-6989(81)90169-3
- Rodieck, R. W. (1965). Quantitative analysis of cat retinal ganglion cell response to visual stimuli. *Vision Res.* 5, 583–601. doi: 10.1016/0042-6989(65)90033-7
- Rosch, E., Mervis, C. B., Gray, W. D., Johnson, D. M., and Boyes-Braem, P. (1976). Basic objects in natural categories. *Cogn. Psychol.* 8, 382–439. doi: 10.1016/0010-0285(76)90013-X
- Rosenblatt, F. (1962). Principles of neurodynamics: Perceptrons and the theory of brain mechanisms. Washington, DC: Spartan books.
- Rosenfeld, S. A., and Van Hoesen, G. W. (1979). Face recognition in the rhesus monkey. *Neuropsychologia* 17, 503–509. doi: 10.1016/0028-3932(79)90057-5
- Rosenthal, G., Tanzer, M., Simony, E., Hasson, U., Behrmann, M., and Avidan, G. (2017). Altered topology of neural circuits in congenital prosopagnosia. *eLife* 6:69. doi: 10.7554/eLife.25069
- Saha Roy, T., Mazumder, S., and Das, K. (2021). Wisdom of crowds benefits perceptual decision making across difficulty levels. *Sci. Rep.* 11:538. doi: 10.1038/s41598-020-80500-0
- Schapiro, R. E. (1990). The strength of weak learnability. *Mach. Learn.* 5, 197–227. doi: 10.1007/BF00116037
- Schroff, F., Kalenichenko, D., and Philbin, J. (2015). FaceNet: A unified embedding for face recognition and clustering. In 2015 In: *IEEE conference on computer vision and pattern recognition (CVPR)* (IEEE), 815–823.
- Schütz, A. C., Braun, D. I., and Gegenfurtner, K. R. (2011). Eye movements and perception: a selective review. *J. Vis.* 11:9. doi: 10.1167/11.5.9
- Serences, J. T., and Yantis, S. (2006). Selective visual attention and perceptual coherence. *Trends Cogn. Sci.* 10, 38–45. doi: 10.1016/j.tics.2005.11.008
- Sergent, J., Ohta, S., and MacDonald, B. (1992). Functional neuroanatomy of face and object processing. A positron emission tomography study. *Brain* 115, 15–36. doi: 10.1093/brain/115.1.15
- Serre, T., Oliva, A., and Poggio, T. (2007). A feedforward architecture accounts for rapid categorization. *Proc. Natl. Acad. Sci. U. S. A.* 104, 6424–6429. doi: 10.1073/pnas.0700622104
- Setogawa, T., Eldridge, M. A. G., Fomani, G. P., Saunders, R. C., and Richmond, B. J. (2021). Contributions of the monkey inferior temporal areas TE and TEO to visual categorization. *Cereb. Cortex* 31, 4891–4900. doi: 10.1093/cercor/bhab129
- Sokol, S. (1978). Measurement of infant visual acuity from pattern reversal evoked potentials. *Vision Res.* 18, 33–39. doi: 10.1016/0042-6989(78)90074-3
- Solla, S. A., Levin, E., and Fleisher, M. (1988). Accelerated learning in layered neural networks. *Complex Syst.* 2, 625–639.
- Sugita, Y. (2008). Face perception in monkeys reared with no exposure to faces. *Proc. Natl. Acad. Sci. U. S. A.* 105, 394–398. doi: 10.1073/pnas.0706079105
- Tanaka, K. (1996). Inferotemporal cortex and object vision. *Annu. Rev. Neurosci.* 19, 109–139. doi: 10.1146/annurev.ne.19.030196.000545
- Tanaka, J., and Gauthier, I. (1997). “Expertise in object and face recognition” in Perceptual learning: the psychology of learning and motivation. eds. R. L. Goldstone, D. L. Medin and P. G. Schyns (San Diego, CA: Academic Press), 83–125.
- Taubert, J., Wardle, S. G., Flessert, M., Leopold, D. A., and Ungerleider, L. G. (2017). Face pareidolia in the rhesus monkey. *Curr. Biol.* 27, 2505–2509.e2. doi: 10.1016/j.cub.2017.06.075
- Thompson, K. G., and Bichot, N. P. (2005). A visual salience map in the primate frontal eye field. *Prog. Brain Res.* 147, 251–262. doi: 10.1016/S0079-6123(04)47019-8
- Treisman, A. (1986). Features and objects in visual processing. *Sci. Am.* 255, 114–125. doi: 10.1038/scientificamerican1186-114B
- Tsao, D. Y., Freiwald, W. A., Knutsen, T. A., Mandeville, J. B., and Tootell, R. B. H. (2003). Faces and objects in macaque cerebral cortex. *Nat. Neurosci.* 6, 989–995. doi: 10.1038/nn1111
- Ungerleider, L. G., and Mishkin, M. (1982). “Two cortical visual systems” in Analysis of visual behavior. eds. D. J. Ingle, M. A. Goodale and R. J. W. Mansfield (Cambridge: MIT Press), 549–586.
- Vinken, K., Van den Bergh, G., Vermaercke, B., and Op de Beeck, H. P. (2016). Neural representations of natural and scrambled movies progressively change from rat striate to temporal cortex. *Cereb. Cortex* 26, 3310–3322. doi: 10.1093/cercor/bhw111
- White, B. J., Berg, D. J., Kan, J. Y., Marino, R. A., Itti, L., and Munoz, D. P. (2017). Superior colliculus neurons encode a visual saliency map during free viewing of natural dynamic video. *Nat. Commun.* 8:14263. doi: 10.1038/ncomms14263
- Yamins, D. L. K., and DiCarlo, J. J. (2016). Using goal-driven deep learning models to understand sensory cortex. *Nat. Neurosci.* 19, 356–365. doi: 10.1038/nn.4244
- Yamins, D. L. K., Hong, H., Cadieu, C. F., Solomon, E. A., Seibert, D., and DiCarlo, J. J. (2014). Performance-optimized hierarchical models predict neural responses in higher visual cortex. *Proc. Natl. Acad. Sci. U. S. A.* 111, 8619–8624. doi: 10.1073/pnas.1403112111
- Yu, G., Katz, L. N., Quaia, C., Messinger, A., and Krauzlis, R. J. (2024). Short-latency preference for faces in primate superior colliculus depends on visual cortex. *Neuron*. doi: 10.1016/j.neuron.2024.06.005
- Zhang, K., Zhang, Z., Li, Z., and Qiao, Y. (2016). Joint face detection and alignment using multitask cascaded convolutional networks. *IEEE Signal Process. Lett.* 23, 1499–1503. doi: 10.1109/LSP.2016.2603342
- Zhou, B., Lapedriza, A., Khosla, A., Oliva, A., and Torralba, A. (2018). Places: a 10 million image database for scene recognition. *IEEE Trans. Pattern Anal. Mach. Intell.* 40, 1452–1464. doi: 10.1109/TPAMI.2017.2723009