

The importance of intermediate representations for the modeling of 2D shape detection: Endstopping and curvature tuned computations

Antonio J. Rodríguez-Sánchez

John K. Tsotsos

Centre for Vision Research and Dept. of Computer Science and Engineering
York University, Toronto, ON, Canada, M3J 1P3

ajrs@cse.yorku.ca

Abstract

Computational models of visual processes with biological inspiration - and even biological realism - are currently of great interest in the computer vision community. This paper provides a biologically plausible model of 2D shape which incorporates intermediate layers of visual representation that have not previously been fully explored. We propose that endstopping and curvature cells are of great importance for shape selectivity and show how their combination can lead to shape selective neurons. This shape representation model provides a highly accurate fit with neural data from [17] and provides comparable results with real-world images to current computer vision systems. The conclusion is that such intermediate representations may no longer require a learning approach as a bridge between early representations based on Gabor or Difference of Gaussian filters (that are not learned since they are well-understood) and later representations closer to object representations that still can benefit from a learning methodology.

1. Introduction

In this paper we present a biologically plausible model for shape representation. Curvature is considered an important component in order to achieve object recognition in the brain, along with corners, edges, color, texture and other important features [3]. A group of studies has shown that neurons in visual area V4 in monkeys are involved with analyzing curvature [16, 17]. In the hierarchical object recognition pathway, V4 is the area just before the inferotemporal cortex (IT), where object recognition is achieved [21].

2D shapes (closed contours) are an important feature for recognizing objects [3]. Shapes are composed of straight lines, corners, junctions and curves. Curvature has been present in many methods and algorithms that achieve object




Figure 1. The hierarchy of 2D shape processing.

recognition (e.g., [14, 7]). But until recently there was no substantial proof that the biological visual system analyzes curvature in areas that are involved in the object recognition pathway, such as V4 [16, 17] and IT [1]. In addition to this, there is evidence that endstopped neurons present in area V2 [22] are selective to contours [6, 5]. These works provide the intellectual foundation for the research described in this paper.

2. Hierarchical modeling of Shape

This section explains the details of the model. A summary of the architecture is briefly described as follows (Figure 1): V1 is composed of simple cells whose receptive fields are linear filters with the functional form of Differences of Gaussians and of complex cells that are the result of the summation of spatially displaced simple cells. V2 contains endstopped cells which respond to variations of straightness which are the input for local curvature selective neurons. Then, shape-selective neurons (V4) respond

to curvature configurations with respect to their position in the neuron's receptive field. Each type of cell is described next in detail.

2.1. Simple cells

For the modeling of simple cells Gabor filters [13] and Difference of Gaussians have been shown to provide a good fit, although a better fit to neuronal responses has been found with Difference of Gaussians [11]. We implemented both types but we obtained better results using the Difference of Gaussians formulation for our model (Figure 1a):

$$G(x, y) = \frac{1}{2\pi\sigma_{x_1}\sigma_y} e^{-\frac{1}{2}\left(\left(\frac{x'}{\sigma_{x_1}}\right)^2 + \left(\frac{y'}{\sigma_y}\right)^2\right)} - \frac{1}{2\pi\sigma_{x_2}\sigma_y} e^{-\frac{1}{2}\left(\left(\frac{x'}{\sigma_{x_2}}\right)^2 + \left(\frac{y'}{\sigma_y}\right)^2\right)} \quad (1)$$

$$\begin{aligned} x' &= x\cos(\theta) + y\sin(\theta) \\ y' &= -x\sin(\theta) + y\cos(\theta) \end{aligned}$$

where σ_y is the height and σ_{x_1} and σ_{x_2} are the width of each Gaussian function. θ is their orientation. For our experiments, we used 12 orientations and 4 different sizes, this gives a total of 48 types of V1 model simple neurons. The 4 different sizes will be used for curvature estimation as explained next.

2.2. Complex cells

In our model, a complex cell is the sum of 5 laterally displaced model simple cells (Figure 1b) as follows from [20]. The model complex cell response is given by [5]:

$$R_{CX} = \sum_{i=1}^n c_i \phi(R_i) \quad (2)$$

R_i is the response of the i th cell and c_i is its weight. Model cells are Gaussian weighted with position, with weight inversely proportional to the distance to the center. ϕ is a rectification function, where any value less than 0 is set to 0. Cells are proportionally shifted by their Difference of Gaussians parameters, so that the cell separation is related to the size (S) and aspect ratio (A) of the component model simple cells along their preferred orientation:

$$separation = \frac{S}{2A} \quad (3)$$

2.3. Endstopped cells

Endstopped cells - also known as hypercomplex - were fully described by Orban and colleagues [15]. They showed that endstopped cells had different properties from orientation-selective cells and provided a description of those properties as well as a detailed study on the end-zone inhibitory areas that were part of such cells.




Figure 2. Response of the model endstopped cells to different radius of curvatures. Simple cell sizes were a)40, b)80, c)100 and d)120 pixels. $\sigma_y=(10,20,25,30)$, $\sigma_{x_1}=(2.5,4,3,3,3)$, $\sigma_{x_2}=2.5 \times \sigma_{x_1}$, $c_{d1} = c_{d2} = (0.7,0.8,1,2)$, $c_c = 1$.

Degree of curvature. Model endstopped cells provide us with a coarse curvature estimation so that we can divide contours into curvature classes. For the design of the model endstopped cells we followed the work of [5]:

$$R_{ES} = \Phi[c_c \phi(R_c) - (c_{d1} \phi(R_{d1}) + c_{d2} \phi(R_{d2}))] \quad (4)$$

$$\Phi = \frac{1 - e^{-R/\rho}}{1 + 1/\Gamma e^{-R/\rho}}$$

c_c , c_{d1} and c_{d2} are the gains for the center and displaced cells. R_c , R_{d1} and R_{d2} are the responses of the center and the two displaced cells. The center cell is a model simple cell and the displaced cells are model complex cells ϕ and Φ are rectification functions. In our experiments, for equation 4, $\Gamma=0.01$ and ρ is the maximum response of the set of neurons for a given scale divided by 8.5, a factor that provided a good normalization approximation for this rectification. Displaced cells are shifted 1/2 of their receptive field size.

The center simple cell has an excitatory effect while the two complex cells have an inhibitory effect (Figure 1c). The response show a good selectivity to curvature, this aspect is due to the inhibitory effect of the displaced cells, which we may note are wider than the center cell, following [15]. Model endstopped cells provide us with a coarse curvature estimation as shown in Figure 2.

Direction endstopped neurons. We refer to direction as the direction of the normal to the curve. The same model endstopped cells are used adding a rotated component on each displaced complex cell with opposite directions (e.g. 45° and 135° for the 0° model endstopped neurons, Fig. 1c). These cells provide valuable information in order to compute convexity/concavity at a later stage. Two types of model direction cells are used, and we will use the term *sign* to specify if the curvature is in one direction or the opposite, positive or negative. These different directions are obtained by changing the order of the displaced subtracted neurons (e.g. Fig 3a vs Fig 3d). If the orientations of the displaced cells lie between the tangent to the curve and the normal

to the curve (in the positive direction within the coordinate system), then the sign is defined as positive, the neuron's response is R_+ , otherwise sign is defined as negative and its response is R_- :

$$R_+ = \phi[c_c \phi(R_c) - (c_{d145} \phi(R_{d145}) + c_{d2135} \phi(R_{d2135}))]$$

$$R_- = \phi[c_c \phi(R_c) - (c_{d1135} \phi(R_{d1135}) + c_{d245} \phi(R_{d245}))]$$
(5)

where c_c , c_{d1} and c_{d2} are the gains for the center and displaced cells as before. R_c , R_{d1} and R_{d2} are the responses of center and displaced cells. The difference here is that the displaced cells are at different orientations of the preferred center simple cell, for the positive sign model endstopped neuron, the displaced model complex neuron $d1$ is at 45° , while the model complex component $d2$ is at 135° . Fig.3e shows the response of direction selective cells over a circular shape for direction positive vs negative.

2.4. Local curvature cells

Curvature cells are obtained due to the neural convergence of the two types of model endstopped cells. By combining model endstopped cells and model direction endstopped cells responses (Figure 1d), we obtain twice the number of curvature classes than the number of model endstopped cells. For example, if we have four types of model endstopped cells, through the use of the direction of curvature of those cells we obtain eight curvature classes.

$$R_{curv_i} = R_{ES_i} \cap (R_+ > R_-)$$
(6)

$$R_{curv_{i+n}} = R_{ES_i} \cap (R_- > R_+)$$

where n is the number of model endstopped cell types, R_{ES_i} is the response of the model endstopped cell i and R_+ , R_- are the responses of the model direction selective endstopped neurons. Curvature is obtained directly as the pattern of different responses of local curvature neurons. Independently of the degree of curvature and direction of curvature, a local curvature neuron will always provide a response which will be higher or lower depending on its curvature selectivity (see Figure 2).

2.5. Shape-selective neurons

Our Shape cells integrate the responses from the model local curvature neurons. The proposed response of a model's Shape neuron is:

$$R_{shape} = \sum_{i=1}^n c_i R_{curv_i}(\lambda)$$

$$\lambda = \max_{j=1}^m (\lambda_j), c_i = \frac{1}{2\pi} e^{-(x^2+y^2)}$$
(7)

where R_{curv_i} is the response of the i th model curvature neuron from the set of all possible n model curvature




Figure 3. Direction selective neurons. Endstopped configuration of cells whose curve's normal direction selectivity is in the direction of the solid line (a), if the normal is in the direction of the dashed line, it is inhibited; (b), (c) and (d) show other configurations, e) An example of a circle, positive sign corresponds to the curve in blue. Negative sign is shown in red.

neurons at the preferred curvature direction (λ) inside the Shape neuron receptive field, and c_i is a Gaussian weight that would account for partial excitation depending on the selective curvature in distance - angular position (Figure 1e).

3. Experiments

We performed two sets of experiments. In the first set we study the capability of our model Shape-selective neurons to represent complete shapes as aggregates of boundary fragments in curvature \times angular position domain. In a second set of experiments we examine the performance of our model for classifying real world images.

3.1. The curvature \times angular position representation

Pasupathy and Connor [17] recorded from neurons in area V4 of the macaque monkey. The response pattern was quantified by using a 2D Gaussian tuning function. We wanted first to test the capability of our model to encode curvature in the Curvature \times angular position representation proposed by Pasupathy and colleagues and we compare our model's performance against the results from [17].

3.1.1 Methods

Experiments were run in Matlab in a Mac G5 PowerPC computer. The input to the model is a gray-value image. Images used are 400×400 pixels, a shape would span 300×300 pixels and correspond to the stimuli used in the aforementioned study.

Simple cells comprised 4 scales and 12 orientations. Sizes were 40, 60, 88 and 120 pixels, their corresponding values of σ_y were 10, 15, 22 and 30; $\sigma_{x_1} = 5.7, 4.3, 4$ and 4 ; $\sigma_{x_2} = 2.5 \times \sigma_{x_1}$. Complex cells were obtained as explained before. For the integration into model endstopped neurons, the values of gain c (Equation 4) for displaced neurons c_{d1}

$= c_{d2}$ were, from the smaller to the larger cell: 1.5, 1.25, 1, 3, $c_c = 1$ for all centre cells. For the chosen parameters, cells respond (considering 90% of their maximum value) to the following ranges of curvature radius: 6 to 11, 25 to 52, 48 to 77 and 140 to 300+ pixels. The 4 types of model endstopped neurons and the curvature direction selective neurons lead to eight curvatures.

The stimuli from [17] contains boundary elements that include sharp convex angles, and medium and high convex and concave curvatures. The combination of these boundary elements gave rise to 49 different stimuli (Figure 4, left icons). Neuronal responses from [17] are reported using convex and concave curvatures. The method to compute convexity/concavity uses our model endstopped direction neurons and the centroid of the object i.e., the center of mass of the region. Following this strategy, 8 polar regions are considered in increments of 45° . E.g. the first area (0° to 45°) is convex provided model direction neurons are to the left, any other condition means it is concave. For each area, the direction of its model neuron is considered in order to be convex or concave and a map of convexity and concavity is obtained based on the model neurons directions and the polar area they belong to relative to the centroid of the object.

The 4 types of model endstopped neurons and the curvature direction selective neurons lead to eight curvatures. In order to compare with [17] we transformed the curvature radius pixels to which our model endstopped neurons were selective into the squashed curvature values used by Pasupathy and colleagues. Pasupathy and Connor's stimuli use higher values for convexities than concavities (their range of curvatures is 1 (convex) to -0.3 (concave)), for this reason we took the higher values of the ranges for the convex contours provided before ($radius=11, 52, 77, 301$ pixels), and the lower values for the concave contours ($radius=6, 25, 48, 140$ pixels). We transformed the pixel values into curvature values considering the resolution of the screen and that the definition of curvature is $1/radius$. We finally used the *squashed* curvature equation found in [17]:

$$c' = \frac{2.0}{1 + e^{-ac}} - 1.0 \quad (8)$$

After performing the preceding transformations for comparison purposes, eight curvature classes were considered: (1) Very high curvatures/corners, $c' = 1$, (2) High convex curvatures, $c' = 0.74$, (3) Medium convex curvatures, $c' = 0.45$, (4) Low convex curvatures, $c' = 0.17$, (5) Straight lines, $c' = 0.0$, (6) Low concave curvatures, $c' = -0.1$, (7) Medium concave curvatures, $c' = -0.3$, and (8) High concave curvatures, $c' = -0.4$. To consider a curvature as class 5 (Straight lines), its response to any curvature cell is to be less than 20% the maximum curvature cells response and have a high response to a simple neuron

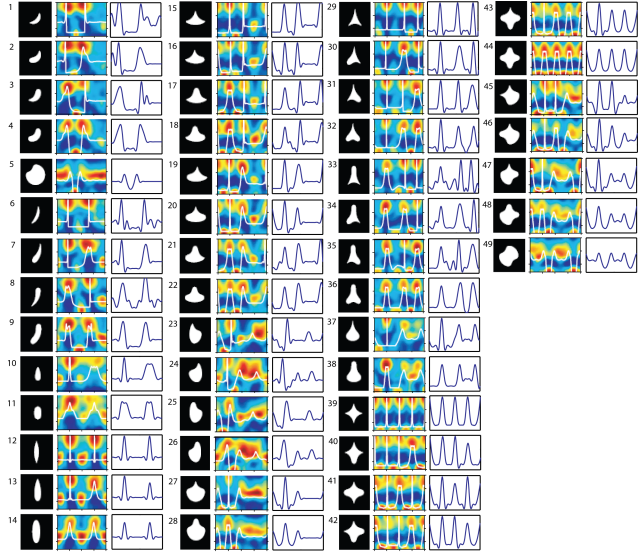


Figure 4. For every stimulus (left image within each of the 4 columns), Pasupathy and Connor's fit (white curve, center) superimposed summed population neural responses (colored values, center) was compared to our model's response (right), blue curve on white background.

(>50% maximum response).

The maximum curvature responses from model endstopped neurons were integrated at different angular positions (in 12° steps) with respect to the center of the object and data was fitted using a spline in order to obtain a function to compare with [17].

3.1.2 Results

Figure 4 shows the results to all the stimuli from [17]. For each stimulus (left column of each panel), Pasupathy and Connor's results (colored graphs with white curves, center column of each panel) are compared against our results (blue curve on white background, right column of each panel). It can be seen that the results are quite similar to those obtained for neurons in area V4 of the visual cortex.

We can see that our plots are not as smooth as the real plots corresponding to the stimuli. At some angular positions our plots show small bumps this is an effect of sampling (we use only 12 orientations), in fact those bumps occur at those orientations and provides a limitation on shape representation. Nevertheless, we observe here that the model performs closely to neuronal data, fitting very well in almost every case the Gaussian shape response pattern (colored graph). The peaks of our model are highly correlated to the cells response patterns.

We measured the difference between our model and the veridical curvature functions (white curves in middle rows). For such a task we averaged the Euclidean distance between

the four to eight angular position values provided by Pasupathy and the results from our model in curvature and angular position. Both terms were normalized to 1.

The worst case scenario would be one where the plot provided by the model would be completely the opposite from the curve in [17] (distance=1). The case where our results completely overlap Pasupathy's curve would mean that Shape neurons have the same exact response pattern as real cells (distance=0). The maximum error is 1 and a perfect match is 0. Total average error for all stimuli was 0.074 (stdev = 0.037). These results show that the proposed hierarchical representation based on endstopping to achieve curvature provides very similar data to data from neurons and curvatures, only 5 out of the 49 stimulus distances were over 10% error (or 0.10 distance), and all of them were below 19%. Most errors are on the range 0.02-0.08.

3.2. Test with real world images

Here we provide a preliminary test of our model on real images. For the task we have selected eight databases with clutter (Leaves from [8], cars back, faces, motorcycles, leopards, bottles and airplanes from Caltech256 and cars from [12]) used in previous Computer Vision studies. The task was an object present/absent classification, where the model has to detect if the object in question is present in the image or not. We used the background database as negative (absent) samples.

3.2.1 Methods

Images were resized to 240 pixels width and 160 pixels height. Simple cells comprised 4 scales and 12 orientations. Sizes were 20, 40, 60 and 80 pixels, their corresponding values of σ_y were 5, 10, 15 and 20; σ_{x_1} =1, 2, 2 and 2.7; $\sigma_{x_2}=2.5 \times \sigma_{x_1}$. Complex cells were obtained as explained before. For the integration into model endstopped neurons, the values of gain c (Equation 4) for displaced and center neurons were $c_{d1} = c_{d2} = c_c = 1$. For each image the Shape neuron pattern of responses was obtained, that is the responses from the different local curvatures neurons (8 curvature classes). Those values from local curvature neurons responses were organized in radial bins of radius=10 pixels at angular steps= 4° . Those values were used to construct a feature vector (2640 elements).

That feature vector of pattern of responses was the input to an Adaboost classifier (300 iterations). Training consisted on presenting randomly half the images containing the object (positive samples: 93 for leaves, 263 for cars back, 258 for cars-MIT, 225 for faces, 95 for leopards, 50 for bottles, 413 for motorbikes and 537 for airplanes) and half the background images (negative samples: 225 randomly chosen images). The remaining images were used to test the model (Same number as in training, but different




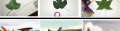
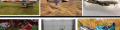


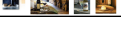
Database	Samples	Model	Benchmark	Han et al.	Serre et al.
Cars back		98.6 %	84.2 % (Fergus)	99.7 %	99.7 %
Cars		96.9 %	75.4 % (Leung)		95.1 %
Faces		89.2 %	90.3 % (Fergus)	93.4 %	98.2 %
Leaves		94.0 %	84.0 % (Weber)	98.8 %	97.0 %
Airplane		92.8 %	93.6 % (Fergus)	98.7 %	96.7 %
Leopards		96.9 %	88.0% (Fergus)	100 %	
Motorbike		96.4 %	97.3% (Fergus)	95.3 %	98.0 %
Bottle		83.3 %	76.4% (Fergus)	70.0 %	

Figure 5. Results with real images.

randomly chosen images).

3.2.2 Results

We obtained the percentage of correctly classified images (as containing objects or background). Results are shown in Figure 5. The model outperforms classical systems such as [23], [9] and [12] for most databases. Correct classification were: 98.6% for cars back (1.9% false negatives and 0.9% false positives), 96.9% for cars-MIT (5% false negatives and 0.9% false positives), 89.2% for faces (12% false negatives and 10% false positives), 94.0% for leaves (10% false negatives and 0.7% false positives), 96.9% for leopards (4% false negatives and 2.6% false positives), 83.3% for bottles (35% false negatives and 16.5% false positives) and 92.8% for airplanes (5% false negatives and 1.2% false positives). Results are similar as well to another biologically inspired model [19], and the very recent Bag-of-features approach by [10].

4. Discussion

A model of intermediate representations and computations has been shown that bridges the gap between the common very early orientation representations found in most categorization systems, and the higher level object level representations. These representational levels have been shown to not only provide very low error fit to real neuronal data but also to perform comparably than current learned representations. Because the process of learning is time-consuming, this implies that our representation may be substituted for the currently learned intermediate representations, thus reducing both learning time and the number of training samples required.

The differences between our model and other recent models, e.g. [19, 2] are several. Whereas Serre, Cadieu and colleagues define their cell types as combinations of edge maximal responses successively over a number of hierarchical layers (7 in [19]), here our neurons in each layer compute quite different quantities. Our goal was to include curvature computations directly, and not indirectly as Serre et al. through the conjunctions of edges. Similarly, [2] used

a forward selection algorithm (greedy search) to determine the subset of Gabor based subunits whose combination approximates curvature. These strategies do not explicitly include either curvature or end-stopped units, both well-known to exist in the visual cortex. Although a learning approach is valuable when there is no other option, in the case of these intermediate representations, there is sufficient knowledge to directly model these cells (in the same way no one "learns" Gabor or DOG cells, because those functions are well understood). However, except for the notable exception of [6], these intermediate cell representations have not been adequately investigated computationally. This is where our approach and that of Serre diverges and enables our true representation of curvature and 2D shape.

To conclude, we have shown that it is possible with current enhanced knowledge of the neurophysiology of curvature and shape detection, to develop new and effective intermediate representations for object categorization systems. These reduce the load on learning procedures without impacting performance but importantly, also provide a new understanding of the nature of object representation.

5. Acknowledgments

We would like to thank Allan Dobbins and Steve Zucker for all their input and advice regarding parts of the model. We would also like to thank Anitha Pasupathy for providing us with the code for constructing the stimuli as well as data corresponding to V4 neurons.

References

- [1] S. Brincat and C. Connor. Underlying principles of visual shape selectivity in posterior inferotemporal cortex. *Nature Neuroscience*, 7(8):880–886, 2004. [4321](#)
- [2] C. Cadieu, M. Kouch, C. Connor, M. Riesenhuber, and T. Poggio. A model of v4 shape selectivity and invariance. *Journal of Neurophysiology*, 3(98):1733–1750, 2007. [4325](#)
- [3] C. Connor, S. Brincatt, and A. Pasupathy. Transformation of shape information in the ventral pathway. *Current Opinion in Neurobiology*, 17(2):140–147, 2007. [4321](#)
- [4] S. Dickinson, A. Leonardis, B. Schiele, and M. J. Tarr, editors. *Object categorization, computer and human vision perspectives*. Cambridge University Press, 2009.
- [5] A. Dobbins. *Difference models of Visual Cortical neurons*. PhD thesis, Department of Electrical Engineering. McGill University, 1992. [4321](#), [4322](#)
- [6] A. Dobbins, S. Zucker, and M. Cynader. Endstopped neurons in the visual cortex as a substrate for calculating curvature. *Nature*. 1987 Oct 1-7;329(6138):438-41., 329(6138):438–441, 1987. [4321](#), [4326](#)
- [7] G. Dudek and J. Tsotsos. Shape representation and recognition from multiscale curvature. *Computer Vision and Image Understanding*, 68(2):170–189, 1997. [4321](#)
- [8] R. Fergus, P. Perona, and A. Zisserman. Object class recognition by unsupervised scale-invariant learning. *CVPR*, 2:264, 2003. [4325](#)
- [9] R. Fergus, P. Perona, and A. Zisserman. A sparse object category model for efficient learning and exhaustive recognition. *CVPR*, 2005. [4325](#)
- [10] X. Han, Y. Chen, and X. Ruan. Image recognition by learned linear subspace of combined bag-of-features and low-level features. *ICIP*, 2010. [4325](#)
- [11] M. Hawken and A. Parker. Spatial properties of neurons in the monkey striate cortex. *Proceedings of the Royal Society of London, series B, Biological Sciences*, 231:251–288, 1987. [4322](#)
- [12] B. Leung. *Component-Based Car Detection in Street Scene Images*. PhD thesis, Massachusetts Institute of Technology, Dept. of Electrical Engineering and Computer Science, 2004. [4325](#)
- [13] S. Marcelja. Mathematical description of the responses of simple cortical cells. *Journal of Optical Society of America*, 70(11):1297–1300, 1980. [4322](#)
- [14] F. Mokhtarian. Silhouette-based isolated object recognition through curvature scale space. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(5):539–544, 1995. [4321](#)
- [15] G. Orban, H. Kato, and P. Bishop. Dimensions and properties of end-zone inhibitory areas of hypercomplex cells in cat striate cortex. *Journal of Neurophysiology*, 42:833–849, 1979. [4322](#)
- [16] A. Pasupathy and C. Connor. Responses to contour features in macaque area v4. *Journal of Neurophysiology*, 82(5):2490–2502, 1999. [4321](#)
- [17] A. Pasupathy and C. Connor. Population coding of shape in area v4. *Nature Neuroscience*, 5(12):1332–1338, 2002. [4321](#), [4323](#), [4324](#), [4325](#)
- [18] J. Peissig and M. Tarr. Visual object recognition: Do we know more now than we did 20 years ago? *Annual Review of Psychology*, 58:15–96, 2007.
- [19] T. Serre, L. Wolf, S. Bileschi, and M. Riesenhuber. Robust object recognition with cortex-like mechanisms. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 29(3):411–426, 2007. [4325](#)
- [20] H. Spitzer and S. Hochstein. A complex-cell receptive-field model. *Journal of Neurophysiology*, 53(5):1266–1286, 1985. [4322](#)
- [21] K. Tanaka, H. Saito, Y. Fukada, and M. Moriya. Coding visual images of objects in the inferotemporal cortex of the macaque monkey. *Journal of Neurophysiology*, 66(1):170–189, 1991. [4321](#)
- [22] R. von der Heydt, E. Peterhans, and G. Baumgartner. Illusory contours and cortical neuron responses. *Science*, 224(4654):1260–1262, 1984. [4321](#)
- [23] M. Weber, W. Welling, and P. Perona. Unsupervised learning of models of recognition. *ECCV*, 2:1101–1108, 2000. [4325](#)