



ELSEVIER

Information Sciences 126 (2000) 41–56

INFORMATION  
SCIENCES

AN INTERNATIONAL JOURNAL

[www.elsevier.com/locate/ins](http://www.elsevier.com/locate/ins)

# Recognition model with narrow and broad extension fields

Peter Kalocsai \*

*U1 Consulting Group, Inc., 564 Mission Street #3, San Francisco, CA 94105, USA*

Received 1 June 1997; revised 21 June 1998; accepted 30 October 1999

---

## Abstract

A recognition model which defines a measure of shape similarity on the direct output of multiscale and multiorientation Gabor filters does not manifest qualitative aspects of human object recognition of contour-deleted images in that: (a) it recognizes recoverable and nonrecoverable contour-deleted images equally well whereas humans recognize recoverable images much better, (b) it distinguishes complementary feature-deleted images whereas humans do not. Adding some of the known connectivity pattern of the primary visual cortex to the model in the form of *extension fields* (connections between collinear and curvilinear units) among filters increased the overall recognition performance of the model and: (a) boosted the recognition rate of the recoverable images far more than the nonrecoverable ones, and (b) increased the similarity of complementary feature-deleted images, but not part-deleted ones, and thus attained a closer correspondence to human psychophysical results. Interestingly, performance was approximately equivalent for narrow ( $\pm 15^\circ$ ) and broad ( $\pm 90^\circ$ ) extension fields. © 2000 Published by Elsevier Science Inc. All rights reserved.

---

## 1. Introduction

A task that both biological and artificial vision systems need to solve is recovering boundaries of objects from input that is often noisy and imperfect. The Gestalt grouping principles of co-curvilinearity, proximity, and constancy

---

\* Fax: +1-818 774 2580.

E-mail address: [peter@u1.com](mailto:peter@u1.com); <http://www.u1.com> (P. Kalocsai).

of curvature can help the recovery of meaningful information under these circumstances. There is considerable evidence from neuroscience, e.g. [6], and psychophysics, e.g. [4], that these grouping principles are built into the mammalian visual system in the form of connectivity patterns among processing units. There is both anatomical and physiological evidence that cells with approximately collinear orientation are interconnected primarily by excitatory connection [6,7,17]. Psychophysical results seem to suggest a broader field of connections, between not only collinear units, but also curvilinear ones [4]. For either the narrow or the broad fields, the excitatory connections reveal smoothly decreasing strength with increasing distance and curvature differences [4,7,15].<sup>1</sup> The smoothly decaying excitatory field around an oriented segment is herein referred to as an *extension field*.<sup>2</sup> To compare the effects of both narrow (collinear) and broad (collinear and curvilinear) connectivity patterns among processing units, two versions of the extension field were implemented narrow and broad. In the absence of precise neurophysiological data about the strength of connections between collinear and curvilinear units we choose the algorithmic definition of narrow and broad extension fields to be an excitatory gradient of  $\pm 15$  and  $\pm 90^\circ$ , respectively centered on an oriented segment.

The goal of the present study was to investigate the consequences of adding extension fields to a recognition model that computes shape similarity based on representations of V1 hypercolumn activity. Specifically, we studied whether the extension fields would increase the resemblance of the recognition performance of the model to that shown by humans.

### 1.1. Brief comparison with previous work

Several previous computer vision models have used extension fields to guide the grouping process [8–10,14,16,18]. The main contribution of the present effort is the implementation of such a scheme on a biologically plausible multiscale and multiorientation filter representation, roughly similar to that of a lattice of V1 hypercolumns. This representation allows a measure of shape similarity based on the combined activity produced by both the input image and the grouping process (although this does not necessarily mean that

---

<sup>1</sup> There is also evidence for facilitation (increase in sensitivity for detecting Gabor patches) when local and global orientations are  $90^\circ$  offset (the virtual line connecting two segments is perpendicular to their orientation) which is not modeled here [4,15].

<sup>2</sup> The terms ‘association field’ or ‘stochastic completion field’ can be found in the literature to refer to similar constructs. These terms are generally applied to fields considered to manifest broad tuning. The term ‘horizontal connections’ has been employed to refer to the more narrowly tuned excitatory connections documented for neural units.

grouping results in activity that is indistinguishable from that produced by the original image). Previously efforts have not measured the effect of grouping on shape similarity.

Other differences distinguishing the present effort from prior ones are that the latter studies used only one scale as opposed to our multiscale approach. Since our test images are line-drawings, only one scale size – chosen to be the width (frequency) of the lines – could have very well been used, but a multiscale representation better resembles the sampling properties of biological vision systems. Many of the studies in the grouping literature, e.g., [14,16] use an iterative relaxation algorithm as opposed to the more biologically plausible one-pass operation that was implemented here. An additional feature of the current study is that it directly compares the recognition performance of a grouping model to that of humans on a large number of test images, which is relatively rare in the literature. In the following we describe two experiments on object recognition and compare human data to the performance of our baseline model.

## *1.2. Human experimental results*

In a psychophysical experiment [1], equal amounts of contour were deleted from line drawings in such a way that the parts (geons) were either recoverable or nonrecoverable, as illustrated in Fig. 1. Subjects were unable to name the nonrecoverable stimuli (median accuracy was 0%) even when given the names of the objects prior to the experiment. Given sufficient exposure duration the naming of the recoverable images was almost perfect. Thus the type of contour deletion produces an enormous effect on recognition.

In a name priming task, subjects named briefly presented contour-deleted images in two blocks of trials [2]. For each image, two sets of complementary pairs were created by deleting every other vertex and edge from each simple part in the first set (feature-deleted) and by deleting approximately half the components from each image in the second set (part-deleted) (Fig. 2). In this way each of the four images contained 50% of the original contour. If the members of the complementary feature-deleted pair or the part-deleted pair were superimposed they would provide an intact image without any overlap in contour.

For feature-deleted images, members of a complementary pair primed each other as well as they primed themselves, as evidenced by equivalent facilitation in RTs and error rates on the second block of trials [2]. However, there was no visual priming for part-deleted complementary images. Presumably in the case of feature-deleted pairs the same simple parts could be activated by either member of a complementary pair (through different image features though), but this was not possible for complementary part-deleted pairs.



Fig. 1. Examples of test images with recoverable and nonrecoverable deletions. The left column shows the intact images. The middle column shows the recoverable versions and the right column shows the nonrecoverable ones with the same amount of contour deleted [1]. The recoverable objects remain recognizable even if their left halves are occluded so they contain half the contour of the nonrecoverable images. This can be readily verified by the reader.




Fig. 2. Left panel: Examples of feature-deleted complementary image pairs. Every other vertex and edge is deleted from each part. With a very long line, as with the bottom line on the long cylinder of the flashlight, half the line was allocated to one member of a complementary pair and the other half to the other. Right panel: Examples of part-deleted complementary image pairs. Each member contains approximately half the parts of the object.

### 1.3. Comparison of performance with the baseline model

The extent to which these experimental results could be accounted for by a matching algorithm based on the direct output of a number of columns of multiscale and multiorientation Gabor filters was tested [5]. Each column is roughly analogous to the simple cells in a V1 hypercolumn (as described later). On average the model recognized the nonrecoverable images as well as the recoverable ones, a result that did not correspond to human data. Compared to humans the system performed much too well on the nonrecoverable images.

The model was also tested on the recognition of feature-deleted complementary image pairs [5] (Fig. 2). Identical images were recognized perfectly. Although the complementary images were recognized also well above chance, their recognition accuracy was markedly lower than that for the identical images. Unlike humans, who do not distinguish between members of a part-complementary pair in an object recognition task for feature-deleted images, the system clearly recognized the identical member much better than the complement.

## 2. General structure of the baseline direct mapping model

The baseline model [13] was originally developed for face recognition. It has achieved high accuracy in recognizing faces from several face databases and continues to be a success as a commercial application on the access control market [12]. The following sections describe the representation and matching of the model.

### 2.1. Representation

First we convolve the image  $I(\vec{x})$ , with a bank of filters.

$$(WI)(\vec{k}, \vec{x}_0) = \int \psi_{\vec{k}}(\vec{x}_0 - \vec{x}) I(\vec{x}) d^2x = \psi_{\vec{k}}^* I. \quad (1)$$

The filters form a self-similar family of Gabor functions which are known under the name of “Morlet wavelets” in the literature and have the general form

$$\psi_{\vec{k}}(\vec{x}) = c_{k,\sigma} \exp\left(-\frac{\vec{k}^2 \vec{x}^2}{2\sigma^2}\right) \exp(j\vec{k}\vec{x}), \quad (2)$$

where  $c_{k,\sigma}$  is a constant and  $\vec{k}$  controls the size of the Gaussian window and the frequency and orientation of the kernel (since the test images are line drawings, only the cosine part of the kernels are used). The constant parameter  $\sigma$  assures

that the ratio of the wavelength and the window size is such that in all cases the shape of the Gabor kernels are similar, and resemble the simple cell receptive field profiles found in V1 [3,11].

The result of the convolution is stored in “memory” for only certain positions on the image (at the vertices of an arbitrary lattice,  $10 \times 10$  in our case). The convolution of the input image with a set of Gabor kernels is shown on Fig. 3. In the current study 3 scales, 8 orientations and a  $10 \times 10$  lattice were used although these numbers can arbitrarily be varied. The particular choice of these parameters was based on the characteristics of the stimuli, accuracy considerations, and computational demands. Since the test images are line-drawings one scale could have been sufficient for testing except that in this case one loses the notion of sampling the frequency space which is also characteristic of biological vision systems. The positioning of the  $10 \times 10$  lattice over an image is shown in the left-hand column of Fig. 4. The result of the convolution with all 24 (3 scales  $\times$  8 orientations) kernels at the same position gives a 24 dimensional vector called a “jet”. This describes how much luminance changes of different orientation and scales are present at a given point in the image. The convolution results for all vertices of the grid along with their positions are stored for each image to form a “gallery”.




Fig. 3. Schematic representation of the baseline Gabor filter model. The model first convolves each input image with a set of Gabor kernels at three scales and eight orientations arranged in a  $10 \times 10$  lattice. The set of kernels at each node in the lattice is termed a “Gabor jet”. The activation values of the kernels in each jet along with their positions are stored for each of the images to form a “gallery”.




Fig. 4. Illustration of the grid distortions in the recoverable-nonrecoverable experiment. The left column shows examples of the intact images that were stored in the gallery with the original grid positioned on them. The middle column shows the distortion of the grid for the recoverable and nonrecoverable versions of the 'boat' image when matched against their intact versions. The right-hand column shows the recoverable and nonrecoverable images without the grids.

## 2.2. Matching

The comparison of a stored image representation with a new incoming image involves two measurements: determining the similarity between stored jets and jets computed in the new image ( $S_v$ ), and measuring the necessary distortion of the grid in the new image in order to find similar jets ( $S_e$ ). Similarity between an input jet ( $J^I$ ) and an output jet ( $J^O$ ), labeled as ( $S_v$ ), is measured by the normalized dot product of the two 24 dimensional jets

$$S_v(J^I, J^O) = - \frac{J^I \cdot J^O}{\|J^I\| \|J^O\|}, \quad (3)$$

which gives the cosine of the angle between the two vectors. Certainly, other types of similarity measure could have also been used, but after some experimentation the cosine seemed to give sufficient result. The distortion of all the edge segments between two neighboring nodes ( $S_e$ ) is measured by the quadratic difference in lengths between corresponding edges in the input and the stored image

$$S_e(D_{i,j}^I, D_{i,j}^O) = (D_{i,j}^I - D_{i,j}^O)^2, \quad (4)$$

where  $D_{i,j}$  is the distance between nodes  $i$  and  $j$ .

The graph matching occurs in the following manner. First the 24 dimensional vectors (or jets) at each pixel of the input image are computed. Next, the same grid that was used for the stored pattern is moved around rigidly on the input image searching for the best initial position for the grid. Rigidity means that the distance between two nodes of the graph does not change. The search

is performed by random walk gradient descent method with arbitrary but less than a maximum step size. At each step a cost function is evaluated which is a combined measure of jet similarities and the grid distortion.

The cost function is

$$C_t = \lambda \sum_{i,j} S_e(D_{i,j}^I, D_{i,j}^O) + \sum_i S_v(J_i^I, J_i^O), \quad (5)$$

where  $\lambda$  is a constant determining the relative importance of the two type of costs, and  $i, j$  takes each possible integer values between 1 and the number of horizontal and vertical vertices, respectively. When the grid is rigid, the first term in the cost function is zero. If the cost in the new position is lower than in the old one, the grid is repositioned. Otherwise it remains in the old position. After the optimal initial position is explored in this way, the second phase of optimization begins where the individual nodes can “diffuse” independently constrained by topographical neighborhood. The nodes take a randomly selected new position if by this step the cost function is reduced by more than a predefined threshold. In this phase the grid gets distorted, therefore the first term in the cost function also contributes to the total cost. The process stops when no improvement happens during a given number of trials. The energy landscape of a local jet is smooth enough to allow the gradient descent method to find its minimum.

The result of the diffusion over a pair of images is shown in the middle column in Fig. 4 (the test object without the distorted grids are presented in the right-hand column. To the extent that the jets move independently, the resultant positions will no longer produce a rectangular lattice, as illustrated in the figure. In general the more distorted the lattice, the less the similarity of the image to the original. The most similar match of the test image is interpreted to be the recognition response of the model. Figs. 5 and 6 give a visual illustration of the *activation fields* (the responses of the individual kernels to an image) created by convolving an image with the differently oriented and scaled kernels in the baseline model (the intact and recoverable ‘boat’ images are used as examples). In the visual representation the activation values of the model are normalized to integer values between 0 and 255 for 8-bit graphical display.

### 3. Additions to the baseline model

The extension field is essentially a probability directional vector-field describing the contribution of a single unit-length edge element to its neighborhood in terms of direction and strength [9]. In other words, it describes the preferred direction and the probability of existence of every point in space to share a curve with the original segment. The field direction at a given point in space is chosen to be tangent to the osculating circle passing through the edge






Fig. 5. Activation fields of the baseline model to the intact 'boat' image. The three rows represent the three scale sizes used in the experiment. The first column shows the two-dimensional picture of the Gabor kernels at the three different scales. From the second to the second to last column the normalized activations of the differently oriented kernels to the intact 'boat' image are displayed starting with horizontal orientation and incrementing by 22.5°. The last column shows the normalized cumulative activation of the three different scales at all orientations.




Fig. 6. Activation fields of the baseline model to the recoverable 'boat' image. The three rows represent the three scale sizes used in the experiment. The first column shows the two-dimensional picture of the Gabor kernels at the three different scales. From the second to the second to last column the normalized activations of the differently oriented kernels to the recoverable 'boat' image are displayed starting with horizontal orientation and incrementing by 22.5°. The last column shows the normalized cumulative activation of the three different scales at all orientations.

element and that point, while its strength is proportional to the radius of that circle (Fig. 7). Also, the strength decays with distance from the origin (the edge segment). The decay of extension field strength is set to be Gaussian for both the proximity and curvature constraints

$$\overline{EF}(x, \rho) = e^{-Ax^2} e^{-B\rho^2}, \quad (7)$$

where  $x$  is the distance along the circular arc and  $\rho$  is the curvature of the given arc. Recently, Williams and Jacobs [18] described a very similar type of prior probability distribution of boundary completion based on computing the probability that a particle following a random walk will pass through a given position and orientation on a path joining two edge segments.

From each end of an edge segment an extension field is defined by the triangular areas as shown in Fig. 8. Together, the two fields extending from each segment form a butterfly-shaped region. Activity that is collinear ( $=0^\circ$ ) with the segment ran through the center of each extension field with orientations that deviate from collinearity extending above and below the center. The maximum orientation difference spanned by the broad extension field is  $\pm 90^\circ$ , which is at the  $\pm 45^\circ$  boundaries of the extension field (Fig. 8). Beyond those values, the Gaussians for orientation are set to zero so the broad extension fields had zero values above and below the main diagonals, as illustrated in Fig. 8. The narrow extension field is a subset of the broad extension field in that




Fig. 7. Field direction for every point in space is chosen to be the tangent to the osculating circle passing through the edge segment and the given point. With the broadly tuned extension field, the  $90^\circ$  tangent would be at the  $\pm 45^\circ$  boundaries of the extension field on both sides of the edge element (Fig. 9 left panel).




Fig. 8. Left panel: The brightness coded directional map of the broad extension fields given a horizontal edge element in the middle. Within the butterfly shaped extension fields, black refers to horizontal and white to vertical orientations. The regions above and below the edge element have no assigned orientation and are shown in black to clearly delineate the  $\pm 45^\circ$  boundaries of the extension field. Right panel: The strength map of the extension field for locations and directions shown in the left figure. Strength declines with increasing orientation differences and distance from the edge element. There is no strength assigned above and below the diagonals.

it uses the same direction and strength fields except that the excitation area is limited to  $\pm 15^\circ$  orientation differences. The absence of grouping activity in the regions outside of the extension field merely means that additional information is needed to reconstruct curves between such pairs.

The extension fields were incorporated into the baseline model by allowing a field to operate on each of the 24 activation fields created by evolving the 24 kernels with an image. Because there were 8 orientations for the activation fields there were also 8 orientations for the extension fields. The additional excitation as provided by the extension field is distributed to the activation fields in such a manner that only the corresponding orientations of the activation fields and extension fields were convolved

$$(EFWI)(\bar{l}, \bar{x}_0) = \int EF_{\bar{l}}(\bar{x}_0 - \bar{x}) WI_{\bar{l}}(\bar{x}) d^2x = EF_{\bar{l}} * WI_{\bar{l}}, \quad (8)$$

where  $\bar{l}$  gives the orientation of both the extension and activation fields. For the broad extension field model the activation fields not only get excitation from the extension field with the same orientation, but also from all the other orientations. For computational case the excitation fields were divided into eight subregions based on orientation and only the corresponding range of orientations were applied to an activation field with a given orientation. For the broad extension field model the overall excitation applied to an activation field is then given by summing up the excitation coming from: (a) the extension field with the preferred orientation of the given activation field and (b) the excitation from all the other extension fields. In the narrow extension field model the activation field with a given orientation was only convolved with the excitation field having the same orientation.

To anticipate a point that will be made in the discussion, the grouping activity can be distinguished from the activity produced directly by the image by keeping a reference copy of the early filter activations without any connectivity.

Fig. 9 shows the activation fields created by convolving an image with the differently oriented and scaled kernels (altogether 24 kernels were used) with the narrowly tuned extension field connections (again the ‘boat’ recoverable images is used as an example). The activation pattern shows high similarity to the activation pattern achieved for the intact ‘boat’ image (as shown in Fig. 5).

Fig. 10 provides a direct visual comparison of the workings of the three different model types. The top row displays three versions of the ‘boat’ image from the set: intact, recoverable and nonrecoverable in left, middle and right columns, respectively. Below the  $3 \times 9$  blocks of images show the cumulative activation patterns induced by the three images in the three examined models: baseline, with narrowly tuned extension fields, with broadly tuned extension fields (from top nine image to bottom nine). In each of the three nine image blocks the first row represents the cumulative activation patterns of the kernels at the highest scale and at all 8 orientations. The second row represents the




Fig. 9. Activation fields produced by narrowly tuned extension fields for the recoverable ‘boat’ image (the model with broad extension fields gave similar results). The three rows represent the three scale kernels used in the experiment. The first column shows the two-dimensional picture of the Gabor kernels at the three different scales. From the second to the second to last column the normalized activations of the differently oriented kernels to the recoverable ‘boat’ image are displayed starting with horizontal orientation and incrementing by  $22.5^\circ$ . The last column shows the normalized cumulative activation of three different scales at all orientations.

cumulative response at the highest and medium scale and the last row shows the ‘total’ of the activation for all scales and all orientations as well (similarity to the last columns of Figs. 5, 6 and 9). This visualization of model activation also shows that for the second and third block of nine images (model with narrow and broad extension fields) the activation patterns for intact and recoverable images are much more similar than for the baseline model (first block of nine images).

### 3.1. Simulations

In the recoverable–nonrecoverable experiment the similarity of 36 intact images with the recoverable and nonrecoverable versions (altogether 108 images) was calculated and compared to each other.

In the feature-deleted vs. part-deleted experiment the similarity of the feature-deleted complementary image pair was compared to the similarity of the part-deleted complementary image pair for 18 images (altogether 72 used).

### 3.2. Result of the simulations

The results of the simulations are displayed in Figs. 11 and 12. The addition of narrowly tuned extension fields between similarly oriented kernels increased the similarity of both the recoverable and nonrecoverable versions to the original intact image, although it increased the similarity of the recoverable version more. Whereas for the baseline model there was no difference between the similarity of recoverable and nonrecoverable images  $t(35) = 0.64, P = 0.52$  the addition of narrow extension fields significantly increased the difference




Fig. 10. The top row displays the intact, recoverable and nonrecoverable versions of an image respectively. The  $3 \times 9$  block of images below the top row display the activation patterns of the three model types (Baseline, Narrow Extension Fields, and Broad Extension Fields) to these images. The first row in each three blocks represents the cumulative activation of the highest frequency kernels at all eight orientations to the three images. The second row in each three blocks shows the cumulative activation of the highest and medium frequency kernels at all orientations. Finally, the third row in each blocks represents the cumulative activation of all three kernel sizes at all orientations (all 24 kernels).

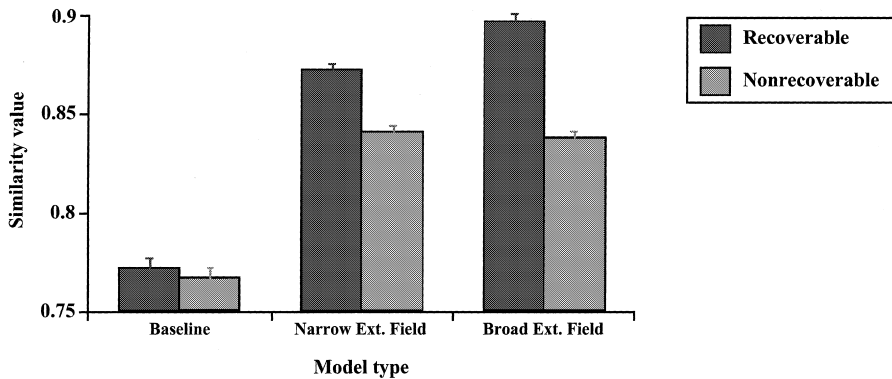


Fig. 11. Average similarity values for matching the original intact images with the recoverable and nonrecoverable versions in the three model types.




Fig. 12. Average similarity values for matching complementary feature-deleted pairs and complementary part-deleted pairs in the three model types.

between the similarity of recoverable and nonrecoverable types compared with the original images  $t(35) = 4.8, P < 0.001$ . The addition of broad extension fields further improved similarity for recoverable images, but did not improve similarity for the nonrecoverable ones compared to the narrow extension fields. Consequently, the broad extension field model further increased the difference between the similarity of recoverable and nonrecoverable images compared with the intact versions  $t(35) = 9.09, P < 0.001$ .

The addition of narrow and broad extension fields significantly increased the similarity of feature-deleted complementary images pairs, but did not improve the similarity of the part-deleted pairs. The similarity of two complementary feature-deleted images was already significantly higher than of two comple-

mentary part-deleted ones  $t(17) = 3.04, P < 0.01$ , but the addition of horizontal connections further improved this difference  $t(17) = 8.54, P < 0.001$ , just as did the addition of extension fields  $t(17) = 9.12, P < 0.001$ . The fact that similarity did not improve for part-deleted image pairs was expected considering that there was no any global knowledge provided that could relate the two different part structures in the pair to each other. However, the significant increase of similarity for the feature-deleted pairs was not an obvious outcome of the simulation. The addition of the broad extension field did not improve similarity for feature-deleted images compared to the narrow extension fields, which might be due to the large number of man-made objects in the stimuli set with mostly straight contours.

#### 4. Conclusions

The addition of extension fields to a model of object recognition that posited a representation based solely on the output of multiscale and multiorientation Gabor filters improved the model's overall recognition performance for contour-deleted images. More important, whereas the model's performance previously did not manifest qualitative aspects of human recognition of contour-deleted images in that it recognized recoverable and nonrecoverable images equally well and distinguished members of feature-deleted complementary pairs of images, the activity contributed by extension fields brought performance closer to that of human observers in that now the model had a higher recognition rate for recoverable compared to nonrecoverable images and the advantage of identical over complementary feature-deleted images was significantly reduced.

Interestingly, adding broad extension fields to the original model did not improve its performance significantly beyond the improvement already achieved by narrow extension fields.

Some of the remaining differences between model and human might be bridged by the incorporation of inhibition and endstopping into the connectivity pattern, both well-known characteristics of early cortical visual activity. The information from endstopping might be employed in making a number of nonaccidental properties explicit, such as curvature, vertices, and cusps. This information, in turn, might be employed in the activation of intermediate representations that might endow the model to express some of the orientation-in-depth robustness characteristic of human object recognition.

#### References

- [1] I. Biederman, Recognition-by-components: a theory of human image understanding, *Psychological Review* 94 (1987) 115–147.

- [2] I. Biederman, E.E. Cooper, Priming contour-deleted images: evidence for intermediate representations in visual object recognition, *Cognitive Psychology* 23 (1991) 393–419.
- [3] R.L. DeValois, K.K. DeValois, *Spatial Vision*, Oxford Press, Oxford, 1988.
- [4] D.J. Field, A. Hayes, R.F. Hess, Contour integration by the human visual system: evidence for a local association field, *Vision Research* 33 (2) (1993) 173–193.
- [5] J. Fiser, I. Biederman, E.E. Cooper, To what extent can matching algorithms based on direct outputs of low level generic descriptors account for human object recognition. *Spatial Vision*, 1996.
- [6] C.D. Gilbert, T.N. Wiesel, Columnar specificity of intrinsic horizontal and corticocortical connections in cat visual cortex, *The Journal of Neuroscience* 9 (7) (1989) 2432–2442.
- [7] C.D. Gilbert, Horizontal integration and cortical dynamics, *Neuron* 9 (1992) 121–128.
- [8] S. Grossberg, E. Mingolla, Neural dynamics of perceptual grouping: textures boundaries and emergent segmentations, *Perception and Psychophysics* 38 (1985) 141–171.
- [9] G. Guy, G. Medioni, Inferring global perceptual contours from local features, *International Journal of Computer Vision* 20 (1/2) (1996) 113–133.
- [10] F. Heitger, R. von der Heydt, A computational model of neural contour processing: figure-ground segregation and illusory contours, in: *Proceedings of the ICCV*, 1993, pp. 32–40.
- [11] J.P. Jones, L.A. Palmer, The two-dimensional spatial structure of simple receptive fields in cat striate cortex, *Journal of Neurophysiology* 58 (6) (1987) 187–1211.
- [12] W. Konen, ZN-Face: A system for access control using automated face recognition, in: *Proceedings of the International Workshop on Automated Face- and Gesture-Recognition*, 1995, pp. 18–23.
- [13] M. Lades, J.C. Vortbrüggen, J. Buhmann, J. Lange, C. vonderMalsburg, R.P. Würtz, W. Konen, Distortion invariant object recognition in the dynamic link architecture, *IEEE Transactions on Computers* 42 (1993) 300–311.
- [14] P. Parent, S.W. Zucker, Trace inference curvature consistency and curve detection, *IEEE Transactions on PAMI* 11 (8) (1989) 823–839.
- [15] U. Polat, D. Sagi, The architecture of perceptual spatial interactions, *Vision Research* 34 (1) (1994) 73–78.
- [16] A. Sha’ashua, S. Ullman, Structural saliency: the detection of globally salient structures using a locally connected network, in: *Proceedings of the ICCV*, Tampa Fl., 1988, pp. 321–327.
- [17] D.Y. Ts’o, C.D. Gilbert, The organization of chromatic and spatial interactions in the primate striate cortex, *Journal of Neuroscience* 8 (1988) 1712–1727.
- [18] L.R. Williams, D.W. Jacobs, Stochastic completion fields: a neural model of illusory contour shape and salience, in: *Proceedings of the ICCV*, 1995.