

UC Merced

Proceedings of the Annual Meeting of the Cognitive Science Society

Title

Recognition Model with Narrow and Broad Extension Fields

Permalink

<https://escholarship.org/uc/item/27n4v71r>

Journal

Proceedings of the Annual Meeting of the Cognitive Science Society, 19(0)

Authors

Kalocsai, Peter

Biederman, Irving

Publication Date

1997

Peer reviewed

Recognition Model with Narrow and Broad Extension Fields

Peter Kalocsai (KALOCSAI@SELFORG.USC.EDU)

Irving Biederman (IB@RANA.USC.EDU)

University of Southern California
HNB-316, Los Angeles 90089 USA

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, b) increased the similarity of complementary feature-deleted images, but not part-deleted ones, more closely corresponding to human psychophysical results. Interestingly, performance was approximately equivalent for narrow ($\pm 15^\circ$) and broad ($\pm 90^\circ$) extension fields.

Introduction

A task that both biological and artificial vision systems have to solve is to recover boundaries of objects from many times imperfect, noisy input. The Gestalt grouping principles of co-curvilinearity, proximity, constancy of curvature can help recovering meaningful information under these circumstances. There is considerable evidence from neuroscience (e. g. Gilbert & Wiesel, 1989) and psychophysics (e. g. Field et al., 1993) 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 mainly by excitatory connections (Gilbert & Wiesel, 1989). Psychophysical results seem to suggest a broader field of connections between not only collinear units, but also curvilinear ones (Field et al. 1993). For either the narrow or the broad fields, the excitatory connections reveal smoothly decreasing strength with increasing distance and curvature differences (Field et al. 1993; Gilbert, 1992; Polat & Sagi, 1994). There is also evidence for facilitation (increase in sensitivity for detecting Gabor patches) when local and global orientations are 90 degrees offset (the virtual line connecting two segments is perpendicular to their orientation) which is not modeled here.

The smoothly decaying excitatory field around an oriented segment is referred to as an *extension field* in this paper. 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. To compare the effects of both narrow (collinear) and broad (collinear and curvilinear) connectivity patterns among processing units we decided to implement two versions of the extension field: a narrow and a broad one. In the absence of precise neurophysiological data for 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 $\pm 15^\circ$ and $\pm 90^\circ$ degrees 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 computed 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.

Brief comparison with previous work

Several previous computer vision models have used extension field type algorithms to guide the grouping process (Grossberg & Mingolla, 1985; Sha'ashua & Ullman, 1988; Parent & Zucker, 1989; Heitger & von der Heydt, 1993; Guy & Medioni, 1996; Williams & Jacobs, 1996). The main contribution in 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 allowed 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 grouping results in activity that is indistinguishable from that produced by the original image). The previously cited efforts did not result in a measure of shape similarity.

Other differences distinguishing the present effort from prior ones was that the latter studies used only one scale as opposed to our multiscale approach. Since our test images were 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., Parent & Zucker, 1989; Sha'ashua & Ullman, 1988) used an iterative relaxation algorithm as opposed to the more biologically plausible one-pass operation which was implemented here. An additional feature of the current study is that it directly compared 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 will describe two experiments on object recognition and compare human data to the performance of our baseline model.

Human experimental results and comparison of performance with a baseline model

In a psychophysical experiment (Biederman, 1987) equal amount of contour was deleted from line drawings in such a way that the parts were either *recoverable* or *nonrecoverable* as illustrated in Figure 1. Subjects were able to recognize recoverable versions, but not nonrecoverable ones. A model (Lades et. al., 1993) based on the direct output of a number of columns of multiscale and multiorientation Gabor filters (each column is roughly analogous to the simple cells in a VI hypercolumn) was tested on the same images (this recognition system was originally developed for face recognition and 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 (Konen, 1995)). The model recognized the nonrecoverable images as well as the recoverable ones, a result that does not correspond to human data (see results later).

In another task, subjects named briefly presented contour-deleted images (Biederman and Cooper, 1991). 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) (Figure 2). 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.

Members of a complementary feature-deleted image pair (Figure 2. left) were equivalent to each other for human subjects as tested with the priming paradigm (Biederman and Cooper, 1991), but not for the model since the similarity of members of a pair was markedly lower than similarity of one of the images from the pair to itself. Part-deleted complementary images (Figure 2. right) were not equivalent neither to humans nor to the model.

Figures 3 and 4 give a visual illustration of the *activation fields* (the responses of the individual kernels in the model 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 0-255 for 8-bit graphical display.



Figure 1. Examples of intact (left), recoverable (middle) and nonrecoverable (right) test images (Biederman, 1987).




Figure 2. Examples of feature-deleted complementary image pairs (left) and part-deleted complementary image pairs (right). Each member contains approximately half the parts of the object.

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 (Guy & Medioni, 1996). 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




Figure 3. 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 2D 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 degrees. The last column shows the normalized cumulative activation of the three different scales at all orientations.




Figure 4. 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 2D 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 degrees. The last column shows the normalized cumulative activation of the three different scales at all orientations.

through the edge element and that point, while its strength is proportional to the radius of that circle (Figure 5.). 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:

$$\overline{EF}(x, \rho) = e^{-Ax^2} e^{-B\rho^2} \quad (1)$$

where x is the distance along the circular arc and ρ is the curvature of the given arc. Recently, Williams and Jacobs (1996) 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 expanded to define a triangular area as shown in Fig. 6. The maximum orientation difference spanned by the broad extension field was $\pm 90^\circ$, which were at the $\pm 45^\circ$ boundaries of the extension field (Figure 6). Beyond those values, the Gaussians for orientation were set to zero so the broad

extension fields had zero values above and below the main diagonals, as illustrated in Fig. 6. The narrow extension field is a subset of the broad extension field in that it uses the same direction and strength fields except that the excitation area is limited to $\pm 15^\circ$ orientation difference. 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 convolving 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 was distributed to the activation fields in such a manner that only the corresponding orientations of the activation fields and extension fields were convolved:

$$(EFWI)(\vec{l}, \vec{x}_0) = \int EF_i(\vec{x}_0 - \vec{x}) WI_j(\vec{x}) d^2x = EF_j * WI_j \quad (2)$$

where \vec{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 ease the excitation fields were divided into 8 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.




Figure 5. 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.




Figure 6. 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 (right).

Figure 7 shows the activation fields created by convolving an image with the differently oriented and scaled kernels (altogether 24 kernels were used) and with the narrowly tuned extension field (again the 'boat' recoverable images is used as an example). Again, for the visual representation the activation values of the model are normalized to integer values between 0-255 for 8-bit graphical display.




Figure 7. 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 sizes used in the experiment. The first column shows the 2D 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 degrees. The last column shows the normalized cumulative activation of the three different scales at all orientations.

Figure 8 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 x 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 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 (similarly to the last columns of Figures 3, 4 and 7). 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).

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).

Result of the simulations

The results of the simulations are displayed on Figures 9 and 10. 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 $t(35) = .64, p = .52$ the addition of narrow extension fields significantly increased the difference between the similarity of recoverable and nonrecoverable types compared with the original images $t(35) = 4.8, p < .001$. 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 $t(35) = 9.09, p < .001$.

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 complementary part-deleted ones $t(17) = 3.04, p < .01$, but the addition of horizontal connections further improved this difference $t(17) = 8.54, p < .001$, just as did the addition of extension fields $t(17) = 9.12, p < .001$. 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.




Figure 8. The top row displays the intact, recoverable and nonrecoverable versions of an image respectively. The 3 x 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).

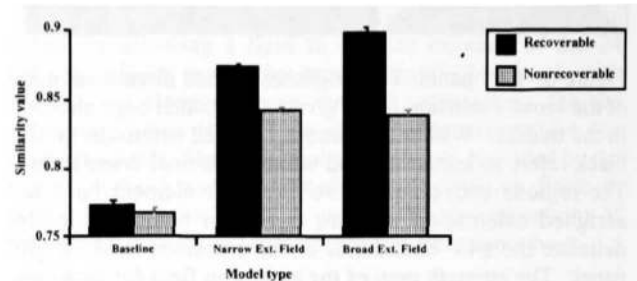


Figure 9. Average similarity values for matching the original intact images with the recoverable and nonrecoverable versions in the three model types.




Figure 10. Average similarity values for matching complementary feature-deleted pairs and complementary part-deleted pairs in the three model types.

Conclusions

The addition of extension fields to a baseline model of object recognition that operates on the output of multiscale and multiorientation Gabor filters improves its overall recognition performance (at minimum for the given set of images) and brings its performance significantly and qualitatively closer to that of human object recognition.

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

An obvious direction for future development of the model is to incorporate inhibition and endstopping into the connectivity pattern, both well known characteristics of biological low level vision systems. We might mention though that even without these additions the model's performance could significantly be improved just based on the addition of excitatory connections. The extension field model is currently being tested on 8-bit gray-scale images.

Acknowledgements

The authors wish to express their appreciation for the significant contributions of Prof. Christoph von der Malsburg, Prof. Gerard Medioni, Gideon Guy and Mi-Suen Lee in this research. This research was supported by grants ONR N00014-95-1-1108 and ARO (NVESD) grant DAAH04-94-G-0065.

References

- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115-147.
- Biederman, I. & Cooper, E. E. (1991). Priming contour-deleted images: Evidence for intermediate representations in visual object recognition. *Cognitive Psychology*, 23, 393-419.
- Field, D. J., Hayes, A. & Hess, R. F. (1993). Contour Integration by the Human Visual System: Evidence for a Local Association Field. *Vision Research*, 33(2), 173-193.
- Gilbert, C. D. & Wiesel, T. N. (1989). Columnar Specificity of Intrinsic Horizontal and Corticocortical Connections in Cat Visual Cortex. *The Journal of Neuroscience*, 9(7), 2432-2442.
- Grossberg, S. & Mingolla, E. (1985). Neural Dynamics of Perceptual Grouping: Textures, Boundaries, and Emergent Segmentations. *Perception and Psychophysics*, 38, 141-171.
- Guy, G. & Medioni G. (1996). Inferring Global Perceptual Contours from Local Features. *International Journal of Computer Vision*, 20(1/2), 113-133.
- Heitger, F. & von der Heydt, R. (1993). A Computational Model of Neural Contour Processing: Figure-Ground Segregation and Illusory Contours. In *Proceedings of the ICCV*, 32-40.
- Konen, W. (1995). ZN-Face: A system for access control using automated face recognition. In *Proceedings of the International Workshop on Automated Face- and Gesture-Recognition*, 18-23.
- Lades, M., Vortbrüggen, J. C., Buhmann, J., Lange, J., von der Malsburg, C., Würtz, R. P. & Konen, W. (1993). Distortion Invariant Object Recognition in the Dynamic Link Architecture. *IEEE Transactions on Computers*, 42, 300-311.
- Parent, P. & Zucker, S. W. (1989). Trace Inference, Curvature Consistency, and Curve Detection, *IEEE Trans. PAMI*, 11(8), 823-839.
- Polat, U. & Sagi, D. (1994). The Architecture of Perceptual Spatial Interactions. *Vision Research*, 34(1), 73-78.
- Sha'ashua, A. & Ullman, S. Structural saliency: the detection of globally salient structures using a locally connected network. In *Proceedings of the ICCV*, 321-327.
- Williams, L. R. & Jacobs, D. W. (1995). Stochastic Completion Fields: A Neural Model of Illusory Contour Shape and Saliency. In *Proceedings of ICCV*.