

## **Putting Gestalt Laws at Work**

**Peter Kalocsai**  
**Department of Psychology**  
**University of California, Los Angeles**  
**CA 90095, USA**

*Communication about this ms. should be addressed to:*

Peter Kalocsai  
1285 Franz Hall  
P.O. Box 951563  
Los Angeles, CA 90095-1563

email: kalocsai@psych.ucla.edu  
Phones: (310) 267-4683 3#  
FAX: (310) 774-2580

### **Abstract**

The inherent complexity and ambiguity of scene properties renders vision extremely difficult without taking advantage of the statistical regularities of natural images. The most basic forms of these image regularities are captured by the Gestalt laws of perceptual organization [5, 15, 21]. This paper reports two versions of a multiscale and multiorientation artificial recognition system in which some of these Gestalt laws: good continuation, co-curvedness, proximity, and constancy of curvature (smoothness) are explicitly implemented. The recognition performance of these two ‘Gestalt enhanced’ models, of humans and of a baseline model where no Gestalt principles were implemented are compared on several sets of line-drawn images. It was found that the performance of the two ‘Gestalt enhanced’ recognition systems not only increased significantly compared with the baseline model, but their performance also became qualitatively much closer to that of humans. There was no significant difference found between the performance of the two enhanced models with slightly different Gestalt law implementations.

## Introduction

The image features of an object could be extremely variable and ambiguous due to the effects of projection, occlusion, background clutter and illumination. An additional source of ambiguity comes from the fact that similar objects can give rise to very different images and also different objects can give rise to practically identical images [12]. Because of all these difficulties vision is considered to be a highly ill-posed problem. Yet we perceive the shapes of objects quickly and reliably despite all the complexity and objective ambiguities of natural images. How is that possible, one might ask. It seems that the best and quite possibly the only successful way of taming the complexity and ambiguity of image representation is to take advantage of the statistical regularities in natural images and scene properties. The most basic of these statistical regularities are expressed qualitatively in the Gestalt laws of perceptual organization. Hence, it seems, that a good starting point for an artificial vision system would be the explicit implementation of the Gestalt laws.

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 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 connections [6, 7, 22]. 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, 19]<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 chose the algorithmic definition of narrow and broad extension fields to be an excitatory gradients of  $\pm 15$  and  $\pm 90$  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 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.

## Brief comparison with previous work

Several previous computer vision models have used extension field type algorithms to guide the grouping process [8, 9, 10, 18, 20, 23]. 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 [13]. This representation allows a measure of shape similarity based on the combined activity produced by both the input image and the grouping process. Previous 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., [18, 20], used 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

---

<sup>1</sup> 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 [4, 19].

<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.

literature. In the following we describe two experiments on object recognition and compare human data to the performance of our baseline model.

### 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 Figure 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) (Figure 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 reaction times 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.




Figure 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.




Figure 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.

### 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. 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 (Figure 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.

### General structure of the baseline direct mapping model

The baseline model [17] 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 [14]. The following sections describe the representation and matching of the model.

### Representation

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

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

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

$$\Gamma_k(\bar{x}) = C_{k,\sigma} * \exp\left(\frac{\bar{k}^2 \bar{x}^2}{2\sigma^2}\right) * \exp(j\bar{k}\bar{x}) \quad (2)$$

Where  $C_{k,\sigma}$  is constant,  $\bar{k}$  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 x 10 in our case). The convolution of the input image with a set of Gabor kernels is modeled on Figure 3. In the current study 3 scales, 8 orientations and a 10 x 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 x 10 lattice over an image is shown in the left hand column of Figure 4. The result of the convolution with all 24 (3 scales x 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".




Figure 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 x 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".

## 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 ( $S_{vertices}$ ), and measuring the

necessary distortion of the grid in the new image in order to find similar jets ( $S_{edges}$ ). Similarity ( $S_{vertices}$ ) between an input jet ( $J_{inp}$ ) and an output jet ( $J_{out}$ ) is measured by the normalized dot product of the two 24 dimensional jets

$$S_{vertices}(J_{inp}, J_{out}) = \frac{J_{inp} * J_{out}}{\|J_{inp}\| * \|J_{out}\|} \quad (3)$$

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

$$S_{edges}(D_{inp}^{i,j}, D_{out}^{i,j}) = (D_{inp}^{i,j} - D_{out}^{i,j})^2 \quad (4)$$

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:

$$C_{total} = \lambda * \sum_{i,j} S_{edges}(D_{inp}^{i,j}, D_{out}^{i,j}) + \sum_i S_{vertices}(J_{inp}^i, J_{out}^i) \quad (5)$$

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 Figure 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. Figures 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 0-255 for 8-bit graphical display.

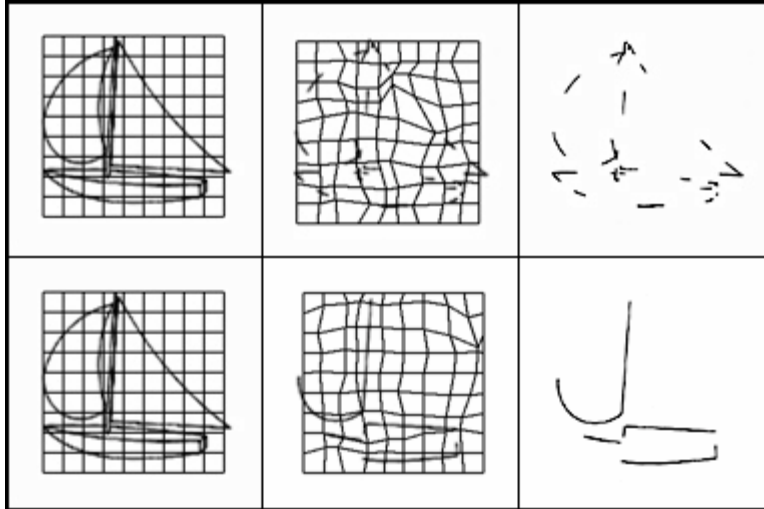


Figure 4. Illustration of the grid distortions in the recoverable-nonrecoverable experiment. The left column shows examples of the intact images that were stored in the gallery with the original grid positioned on them. The middle column shows the distortion of the grid for the recoverable and nonrecoverable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




Figure 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 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 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 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.

### 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 element and that point, while its strength is proportional to the radius of that circle (Figure 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:

$$\overline{EF}(x, \rho) = e^{-Ax^2} e^{-B\rho^2} \quad (6)$$

where  $x$  is the distance along the circular arc and  $\rho$  is the curvature of the given arc. Recently, Williams and Jacobs [23] 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 runs 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 (Figure 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 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.






Figure 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).




Figure 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.

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 is distributed to the activation fields in such a manner that only the corresponding orientations of the activation fields and extension fields are convolved:

$$(EFWI)(\bar{l}, \bar{x}_0) = EF_{\bar{l}}(\bar{x}_0 - \bar{x})WI_{\bar{l}}(\bar{x})d^2x = EF_{\bar{l}} * WI_{\bar{l}} \quad (7)$$

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 ease the excitation fields were divided into 8 sub regions 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 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 Figure 5).

Figure 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 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 Figure 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).

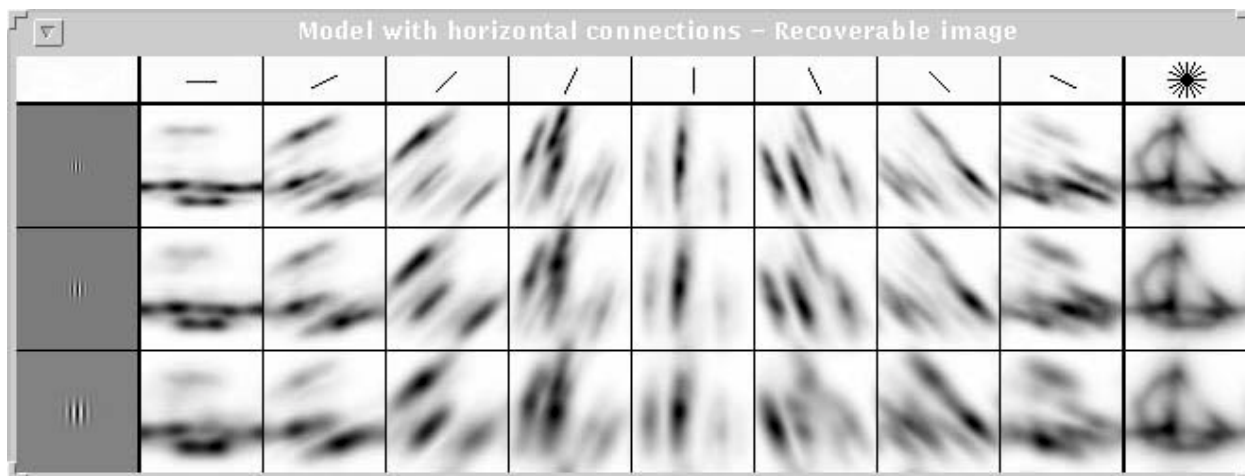


Figure 9. Activation fields produced by narrowly tuned extension fields (the one with only horizontal connections) for the recoverable 'boat' image (the model with broad extension fields gave similar results). The three rows represent the three scale sizes used in the experiment. The first column shows 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.

## 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).































		
<b>Baseline Model</b>		
		
		
		
<b>Narrow Extension Field Model</b>		
		
		
		
<b>Broad Extension Field Model</b>		
		
		
		

Figure 10. 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).

### Result of the simulations

The results of the simulations are displayed on Figures 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  $t(35) = .64, p = .52$  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 11. Average similarity values for matching the original intact images with the recoverable and nonrecoverable versions in the three model types.

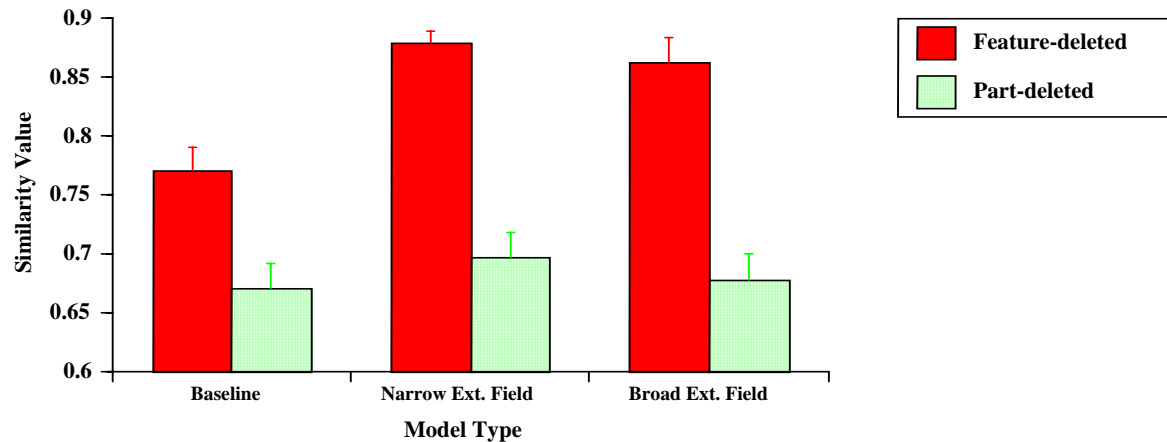


Figure 12. Average similarity values for matching complementary feature-deleted pairs and complementary part-deleted pairs in the three model types.

### Summary of the results

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 importantly, 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.

### Conclusions

Due to the inherent complexity and ambiguity of scene properties vision would be quite difficult without taking advantage of the statistical regularities of natural images. Ambiguity renders even the simplest and most basic low level visual processes such as edge detection impossible to solve correctly. Without taking into account a great deal from the statistical properties of natural scenes even the most sophisticated edge detection algorithms are doomed for failure.

It is argued that the most basic forms of image regularities are captured in the Gestalt laws of perceptual organization [5, 15, 21]. This paper reports two versions of an artificial recognition system in which some of these Gestalt laws: good continuation, co-curvilinearity, proximity, and constancy of curvature (smoothness) are explicitly implemented. Compared with a baseline model without Gestalt laws implemented it is shown on a simple set of line-drawn images that the recognition performance of the resulting new systems is not only considerably increased, but their performance is also qualitatively much closer to that of humans. Further testing of the systems on grayscale natural images is in progress.

Also note that the representation of the recognition systems described above consists of processing units that are tuned to detecting various orientations at various spatial scales in the images. These processing elements respond neither to movement nor to color. Consequently, the implemented Gestalt laws are also defined only over the orientation and scale sensitive units. There is evidence however that the multi-modal integration of Gestalt principles (over orientation, phase, movement and color sensitive elements) can significantly further reduce

ambiguity in the processing of natural images [16], and in turn that could further increase recognition performance. To take advantage of the benefits of multi-modal integration the extension of the current systems with movement and color sensitive processing units is proposed along with the application of the Gestalt laws to these additional elements as well.

## References

- [1] I. Biederman. Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115-147, 1987.
- [2] I. Biederman and E. E. Cooper. Priming contour-deleted images: Evidence for intermediate representations in visual object recognition. *Cognitive Psychology*, 23, 393-419, 1991.
- [3] R. L. DeValois and K. K. DeValois, *Spatial vision*, Oxford Press, 1988.
- [4] D. J. Field, A. Hayes, and R. F. Hess. Contour Integration by the Human Visual System: Evidence for a Local Association Field. *Vision Research*, 33(2), 173-193, 1993.
- [5] W. S. Geisler, J. S. Perry, B. J. Super and D.P. Gallogly. Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, 41, 711-724, 2001.
- [6] C. D. Gilbert and T. N. Wiesel. Columnar Specificity of Intrinsic Horizontal and Corticocortical connections in Cat Visual Cortex. *The Journal of Neuroscience*, 9(7), 2432-2442, 1989.
- [7] C. D. Gilbert. Horizontal Integration and Cortical Dynamics. *Neuron*, 9, 121-128, 1992.
- [8] S. Grossberg and E. Mingolla. Neural Dynamics of Perceptual Grouping: Textures, Boundaries, and Emergent Segmentations. *Perception and Psychophysics* 38, 141-171, 1985.
- [9] G. Guy and G. Medioni. Inferring Global Perceptual Contours from Local Features. *International Journal of Computer Vision*, 20(1/2), 113-133, 1996.
- [10] F. Heitger and R. von der Heydt. A Computational Model of Neural Contour Processing: Figure-Ground Segregation and Illusory Contours. In *Proceedings of the ICCV*, 32-40, 1993.
- [11] J. P. Jones and L. A. Palmer. The two-dimensional spatial structure of simple receptive fields in cat striate cortex. *Journal of Neurophysiology*, 58(6), 1187-1211, 1987.
- [12] D. Kersten and A. Yuille. Bayesian models of object perception. *Current Opinion in Neurobiology*, 13, 1-9, 2003.
- [13] P. Kalocsai. Recognition Model with Narrow and Broad Extension Fields. *Information Sciences*, 126, 41-56, 2000.
- [14] 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*, 18-23, 1995.
- [15] N. Krüger. Collinearity and Parallelism are Statistically Significant Second Order Relations of Complex Cell Responses. *Neural Processing Letters*, 8(2), 117-129, 1998.
- [16] N. Krüger, and F. Wörgötter. Multi-modal Estimation of Collinearity and Parallelism in Natural Image Sequences. *Network: Computation in Neural Systems*, 13(4), 553-576, 2002.
- [17] M. Lades, J. C. Vortbrüggen, J. Buhmann J., Lange, C. von der Malsburg, R. P. Würtz, and W. Konen. Distortion Invariant Object Recognition in the Dynamic Link Architecture. *IEEE Transactions on Computers*, 42, 300-311, 1993.

- [18] P. Parent and S. W. Zucker. Trace Inference, Curvature Consistency, and Curve Detection. *IEEE Trans. PAMI*, 11(8), 823-839, 1989.
- [19] U. Polat and D. Sagi. The Architecture of Perceptual Spatial Interactions. *Vision Research*, 34(1), 73-78, 1994.
- [20] A. Sha'ashua and S. Ullman. Structural saliency: the detection of globally salient structures using a locally connected network. In *Proceedings of the ICCV*, Tampa Fl., 321-327, 1988.
- [21] M. Sigman, G. A. Cecchi, C. D. Gilbert and M. O. Magnasco. On a common circle: Natural scenes and gestalt rules. *PNAS*, 98(4), 1935-1949, 2001.
- [22] D. Y. Ts'o and C. D. Gilbert. The organization of chromatic and spatial interactions in the primate striate cortex. *Journal of Neuroscience*, 8, 1712-1727, 1988.
- [23] L. R. Williams and D. W. Jacobs. Stochastic Completion Fields: A Neural Model of Illusory Contour Shape and Saliency. In *Proceedings of ICCV*, 1995.