A connectionist account of Global Precedence: Theory and data

Elizabeth M. Olds
Department of Psychology
Stanford University
Stanford, CA 94305-2130
olds@psych.stanford.edu

Abstract

A connectionist model was developed to investigate the relationship between global and local information in visual perception, and an experiment tested a prediction generated by the model. The research focused on the fact that processing of global information is found to dominate processing of local information in many tasks ("global precedence"). The connectionist model demonstrated that global precedence can arise out of simple parallel processing. The experiment demonstrated that rotating global elements eliminates Global Precedence. This empirical result supports the possibility, raised by the model, that Global Precedence is due in part to simplicity of Input-Output mapping.

Introduction: Global Precedence

An important issue in research on visual perception is the relationship between global and local information, as they are processed by the visual system. Of particular interest is the phenomenon of "global precedence," where the processing of global information somehow tends to dominate the processing of local information. Briefly, "local" is used to refer to relatively high spatial frequency (HSF) information, or "parts", or details; and "global" is relatively low spatial frequency (LSF) information, or a "whole" object, or the arrangement of parts. The use of these terms in the literature is not consistent; however, in this paper the terms can be viewed as equivalent and thus are used interchangeably.

Navon (1977) sparked a large amount of research on this topic. In a typical experiment, subjects were shown large letters composed of smaller letters (see Figure 1). Subjects had to identify either the global letter or the local letter. When the global and local letters were different, subjects' performance was impaired only if the task was to identify the local letter. That is, only a conflicting irrelevant global aspect interfered with response. Navon concluded that global visual analysis precedes local analysis and cannot be completely ignored. Later experiments have revealed that the global precedence effect depends on many factors, including visual angle (Kinchla & Wolfe, 1979), number and spacing of the local elements (Martin, 1979), number and size of the local elements (Kimchi and Palmer, 1982), positional certainty (Grice, Canham, & Boroughs, 1983), and attentional factors (Boer & Keuss, 1982; Miller, 1981; Ward, 1982), and pattern goodness (Sebrechts & Fragala, 1985). Kimchi (1992) provides a good review of the evidence.

This research aims to specify the mechanisms relevant to global precedence more explicitly. Some researchers state or imply that in order for global precedence to occur, global analysis must start before local analysis. For example, Navon (1977) claims that "Perception proceeds from global analysis to more and more fine-grained analysis."[1] On the other hand—at the other extreme, perhaps—there exist models of perceptual processing that explain certain visual tasks in terms of synthesis of wholes (more global information) from parts (more local information). McClelland and Rumelhart's (1981) Interactive Activation model of letter and word perception is an example, in which information about local parts (e.g., letters) must be accessed before information about the global whole (e.g., a word). It does not seem likely that global precedence would arise within a system in which local information is processed first. In between these two extremes, neurophysiological research indicates that different channels exist in the visual system for the parallel processing of information at different spatial frequencies (e.g., see Livingstone & Hubel, 1988; Spillman & Werner, 1990). In addition, more precise explanations of global precedence, in terms of spatial frequency, have been proposed. Finally, Hughes, Fendrich, and Reuter-Lorentz (1990), and Lamb and Yund (1993) found that global precedence may be due to rapid processing of low spatial frequency information. How can the phenomenon of global precedence be explained within a hypothesis of parallel

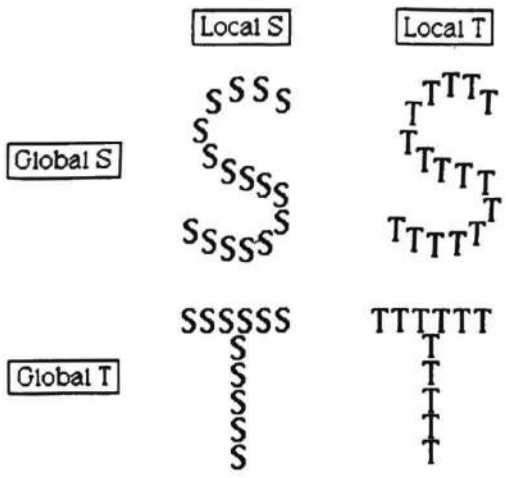

Figure 1: A schematic diagram of the type of stimuli used in a typical Global Precedence experiment.

[1] See also Sanocki, 1993 and Watt, 1987, for interesting versions of this kind of hypothesis.

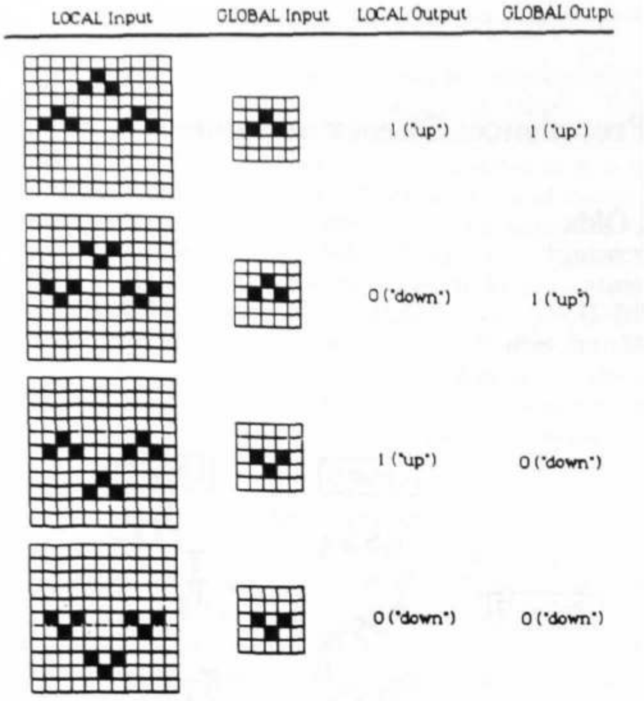

Figure 2: A sample pattern set that the model was trained on.

processing?

A Connectionist Model of Global and Local Processing

Connectionist models are useful because they allow theories about cognitive mechanisms to be tested computationally. If a particular model is able to simulate human data, then the constraints built into it may be computationally important in human cognition as well, especially if the model imitates certain "brain-style" properties of human neural architecture (e.g. see Rumelhart & McClelland, 1986, McClelland & Rumelhart, 1986). The model described here was designed to test the hypothesis that global precedence can arise quite simply out of the parallel processing of global and local information, that is, that global processing is somehow computationally easier than local processing.[2]

Task

The basic task the network learned to perform was a categorization/identification task similar to that performed by Navon's subjects. Two global patterns were created from two component patterns, analogous to Navon's letters made of letters (see Figure 2 for an example). The local input provided the whole pattern, and the component pieces ("parts") had to be categorized as pointing Up or Down; global input is a low spatial resolution summary of the local input, indicating the positions of the "parts" but not any details about them.

The input to the network was such a pattern, and the task of the network was to categorize the large overall pattern as well as its component small patterns.[3] So that the results would not depend on some spurious characteristic of one particular pattern set, the model was trained (separately) on fifteen different pattern sets. The entire group of pattern sets used is presented on the right-hand side of Figure 3.

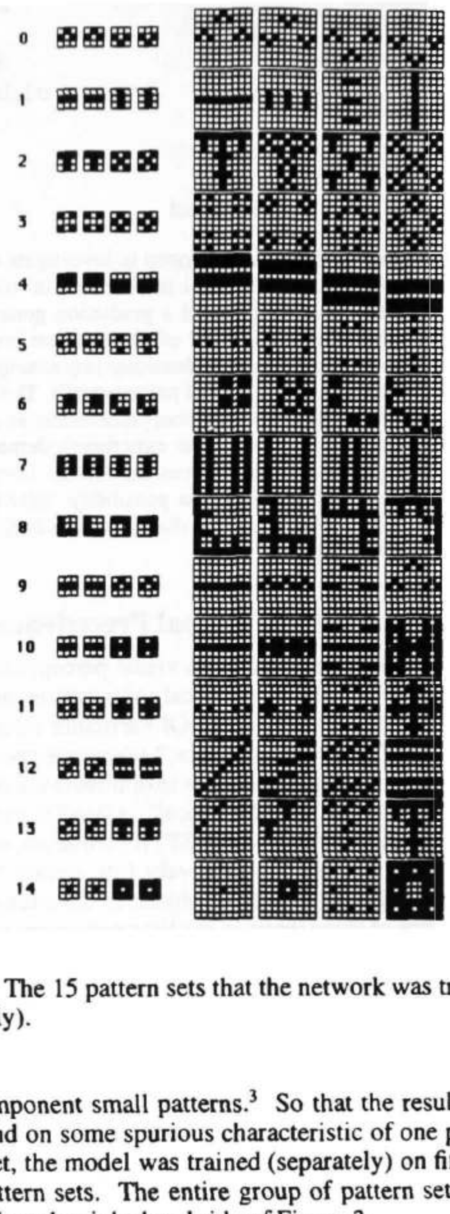

Figure 3: The 15 pattern sets that the network was trained on (separately).

Architecture

The architecture of the network is illustrated in Figure 4. The network is composed of four layers: Input, Feature-Detectors, Hidden, and Output.

There are two sections of Input. In the diagrams, empty squares represent 0's and filled squares represent 1's. The Local input is an 11x11 matrix.[4] The Global input is a 5x5

[2]See Cohen, Dunbar, & McClelland (1990) for a connectionist model of the Stroop effect, which has interesting similarities to the Global Precedence effect.

[3]For these simulations, the component parts were always homogeneous.

[4]A 9x9 pattern of 0's and 1's, surrounded by a border of 0's one

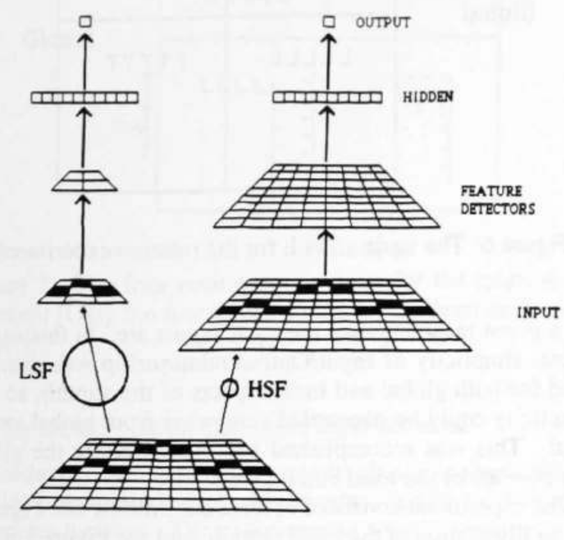

Figure 4: The neural network model of parallel processing of global and local information.

matrix[5], and it is simply a low-spatial-resolution description of the larger Local pattern. That is, the Global input has been designed by the experimenter to be a summary of the Local input, indicating the relative positions of the parts but not their fine details. In a more realistic simulation, the Global and Local inputs here would be LSF and HSF convolutions of a single original image (shown hypothetically at the bottom of Figure 4); however, the current inputs suffice as reasonable representations at some stage of processing.

One section of the Feature-Detector layer is fed by the Global input, and one section is fed by the Local input. Each Feature-Detector unit has a 2x2 receptive field. The Hidden layer has two sections, one receiving input from the Global Feature-Detector units and one receiving input from the Local Feature-Detector units. Each section of the Hidden layer projects to one of two Output units. One Output unit is taught to produce a description (0 or 1) of the Global input, the other is taught to produce a description (0 or 1) of the Local input. All the weights in the network are learned by Backpropagation (Rumelhart, Hinton, & Willams, 1986).

Results

The measures of interest were Global and Local error at each presentation (on each trial, one of the four patterns was selected at random and presented); specifically, the difference between the magnitudes of the two. The difference between Global and Local squared errors was computed at each presentation, and then averaged across chunks of 10 presentations for simplicity. The same pattern set was run four times and the results were averaged across these four runs for reliability. The mean squared error difference (between Global and Local) was added across all trials for a score for each pattern within a pattern set, and the four patterns' scores were added for an overall pattern set score.[6]

unit thick, so that the edges of the pattern could be fully analyzed.

[5]A 3x3 pattern, surrounded by a border.

Table 1: The results for the 15 pattern sets.

Patternset	Score	**Precedence**
0	-1.55	G
1	-1.12	G
2	1.10	L
3	-1.45	G
4	1.60	L
5	-7.18	G
6	-1.27	G
7	-0.96	G
8	0.30	L?
9	-1.21	G
10	1.43	L
11	0.02	?
12	-0.21	G?
13	-1.13	G
14	-2.38	G
sum	-13.07	**G**

See Table 1 for a summary of the results over all the pattern sets. Overall, Global error was lower than Local error, although this difference was not statistically significant. This was interpreted to mean that, across a wide range of stimuli, global analysis was easier overall than local analysis; however, further analyses of the individual pattern sets and the individual patterns within each pattern set, and testing of further pattern sets, are necessary for more reliability.

Discussion

To the extent that the model exhibited Global Precedence, these results are relevant to the human data in that they allow us to examine some computational explanations for Global Precedence. It is reasonable to expect that, in general, an easier task (one producing lower error) might be performed faster than a more difficult task. Further, if a task produces relatively low error and it is therefore performed relatively quickly, perhaps its results will then interfere with other, slower processes. This could explain the phenomenon of Global Precedence—global processing is easier, and thus quicker than local processing, and so it interferes with local processing in the experiments reviewed here. The model simulated the relative simplicity of global and local processing; a more complete model would include this hypothesized interference as well.

Given that the model's performance can be analyzed as potentially relevant to human performance, why did the model

[6]Scores were added at this last stage, instead of averaged—this is because the numbers were small here, due to the inclusion of many trials with zero error, after a given pattern set had been learned but training continued for a fixed number of trials.

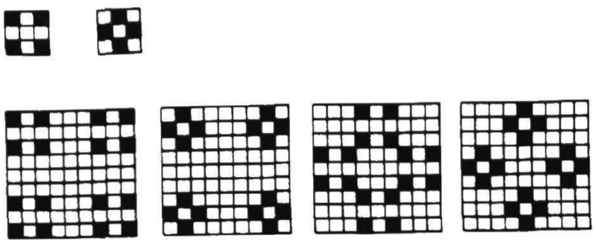

Figure 5: An illustration of how the global task might be relatively easy. The global inputs, above, can be distinguished on the basis of any one of 8 of the 9 input units, alone. However, the local inputs, below, cannot be distinguished on the basis of any one unit (because the parts shift around).

exhibit Global Precedence? Why was the global task easier? One basic answer concerns simplicity of mapping. In the model, the Global Input-Output mapping is simpler than the Local mapping: fewer distinct inputs map on to one output (global inputs are the same for one "global" response, while the local inputs corresponding to one "local" response are different; see Figure 2). In Navon's stimuli, this argument holds as well: a big S always looks the same on a large scale, no matter what its parts look like. On the other hand, the two stimuli made of little T's look quite different on either a small or large scale, even though much HSF information is the same between the two stimuli. That is, the "parts" are the same for these two stimuli (their information is even repeated), but the overall images look quite different, because the local shapes are arranged differently in the two stimuli. So, the simulation results imply that the relative simplicity of the global mapping may be responsible for its precedence in processing.

Another way to describe this simplicity is in terms of "positional certainty": it is clear where the global shape will appear, but not the local component parts. See Figure 5 for an illustration of how this positional certainty might lead a model to exhibit Global Precedence. Generally, there are fewer possible positions for a large object than for a small object, so the more local aspects of an image will tend to suffer more from this positional uncertainty. Furthermore, recent research (e.g., Grice, Canham, & Boroughs, 1983) has investigated situations in which the position of the global object was uncertain between trials as well as the positions of the local shapes, and the results are mixed. Although it might be argued that the current results would be stronger if the model were made translation invariant, the current model realistically simulates Navon's experiment. Additionally, there will always be more positional uncertainty for parts than for wholes, so this logic should apply to visual processing more generally.

Behavioral Data: The Effect of Rotation on Global Precedence

The simulation results suggest that one cause of global precedence could be the fact that the global inputs are more similar for a given response than the local inputs are. In this experiment, simplicity of Input/Output relationship was manipulated for both global and local aspects of the stimuli, so that simplicity could be decoupled somewhat from global versus local. This was accomplished by rotating either the global components or the local components.[7]

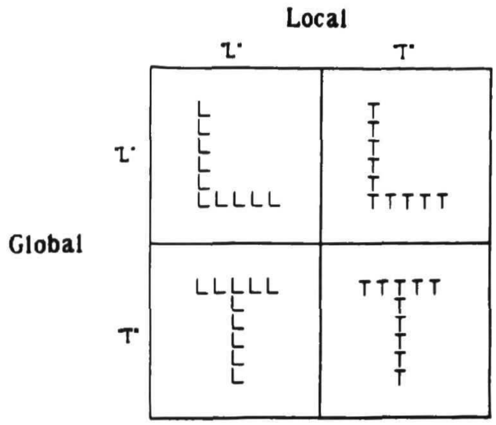

Figure 6: The basic stimuli for the rotation experiment.

The experiment consisted of three conditions. See Figure 6 for an illustration of the basic stimuli, and see Figure 7 for an illustration of the rotation conditions. In Condition 1, stimuli were L's and T's made of L's and T's presented in standard, upright form.[8] In Condition 2, the local component letters were rotated, but not the global letters. That is, an upright T or L was made up of little T's or L's rotated by 0, 90, 180, or 270 degrees.[9] In Condition 3, the global letters were rotated, but not the local letters. That is, a T rotated by 0, 90, 180, or 270 degrees was made up of *upright* L's or T's. Each subject was tested in only one of the three conditions.

Method

Subjects. 18 Stanford University undergraduates participated in order to fulfill a requirement for an introductory psychology course. 6 Stanford University graduate students participated in the experiment voluntarily.

Design and Materials. The stimuli were L's and T's composed of L's and T's (see Figure 6). For each of these four possible shapes, the subject was requested to identify either the global or the local level, on separate trials. These eight basic trials were multiplied by 10, yielding a total of 80 trials, which were then randomized for each subject. Each of these 80 stimuli was then given a random orientation (0, 90, 180, 270 degrees) for the global aspect, local aspect, or neither aspect, depending on the experimental condition (see Figure 7). Each subject was run in one experimental condition only,

[7] This might be viewed as a manipulation of "rotational certainty," comparable to the "positional certainty" manipulated in previous research.

[8] The original intent was to use H's, O's, and S's as Navon did, but the rotational symmetries of these letters made this impossible.

[9] Within a single stimulus, the little letters were all rotated the same amount.

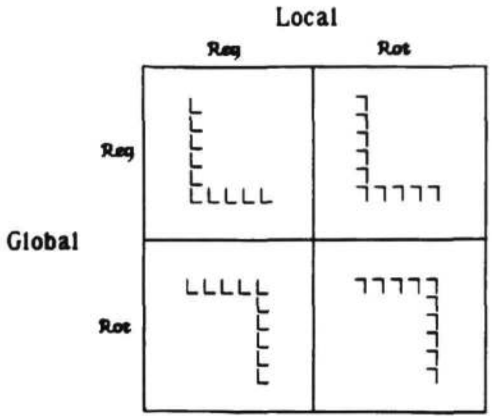

Figure 7: The four rotation conditions for the rotation experiment (Only the first three conditions are described in this report).

and for each condition there were eight subjects.

Procedure. Subjects were run individually in experimental sessions of approximately 15 minutes, seated in front of a computer monitor. The experimental stimuli were approximately 3 cm x 3 cm, which at a distance of approximately 70 cm subtended about 2.5 degrees of visual angle.

Each trial of the experiment consisted of an instruction stimulus, which was either one beep (which meant that the subject should identify the one global letter) or five beeps (which meant to identify the many local letters). Then the test stimulus was presented until the subject responded, pressing "L" if the appropriate level (global or local) of the stimulus was an L, or "T" if it was a T.

Results

The results are presented in Tables 2-4. Analyses were performed on reaction times for correct responses only (subjects made two errors, on average), and from these only those within three standard deviations of the mean for each subject were considered.[10]

Discussion

Condition 1. First, the results in the non-rotated condition replicated the Global Precedence effect reported in the literature (see Table 2). When the global and local letters were different, the local identification task was more disrupted by this conflict than the global task was. That is, the local identification task was performed slower for conflict trials than over all trials (both conflict and non-conflict trials), by an amount greater than the amount that conflict slowed down the global task. This difference was significant ($t=2.59$, $p<0.05$). Thus, this experimental situation is appropriate for the investigation of global precedence.

Condition 2. When local letters were rotated but not global letters, Global Precedence did not occur ($t=0.86$, $p=0.42$). See Table 3.

[10]So that extreme outliers could be removed, without the tail of the non-normal RT distribution being cut too short.

Table 2: Results for Condition 1 of the experiment (neither global nor local letters rotated). The left matrix shows mean RTs (in msec) when the subject was asked to identify the global letter; the right matrix shows the results for the local task.

Identify Global

Global \ Local	"L"	"T"
"L"	519	575
"T"	620	573

Identify Local

Global \ Local	"L"	"T"
"L"	581	704
"T"	671	552

Table 3: Results for Condition 2 of the experiment (local letters rotated).

Identify Global

Global \ Local	"L"	"T"
"L"	651	658
"T"	755	671

Identify Local

Global \ Local	"L"	"T"
"L"	690	820
"T"	801	846

Table 4: Results for Condition 3 of the experiment (global letters rotated).

Identify Global

Global \ Local	"L"	"T"
"L"	684	758
"T"	833	647

Identify Local

Global \ Local	"L"	"T"
"L"	761	789
"T"	799	769

Condition 3. Finally, as expected, rotating the global letters wiped out the Global Precedence effect altogether; in fact, Local Precedence was found to a significant degree ($t=3.90$, $p<0.01$). See Table 4.

Discussion

The connectionist model of Global Precedence, described here, demonstrated that global precedence can arise out of simple parallel processing. Global information may be easier to process than local information because of a greater simplicity in the global Input-Output mapping. This conclusion was supported by the empirical results, which showed that rotating global elements eliminates Global Precedence. More specific support for these claims would be gained by simulating the current rotation experiment with the connectionist model, and comparing the human data to these computational results.

Acknowledgements

David Rumelhart provided useful guidance and comments. This research was carried out under a United States Air Force Laboratory Graduate Fellowship (S.C.E.E.E.) to Elizabeth Olds.

References

[1982] Boer, L. C., & Keuss, P. J. G. (1982). Global precedence as a post-perceptual effect: An analysis of speed-accuracy tradeoff functions. *Perception & Psychophysics, 31*, 358–366.

[1990] Cohen, J. D., Dunbar, K., & McClelland, J. L. (1990). On the control of automatic processes: A parallel distributed processing account of the Stroop effect. *Psychological Review, 97*(3), 332–361.

[1983] Grice, G. R., Canham, L., & Boroughs, J. M. (1983). Forest before trees? it depends where you look. *Perception & Psychophysics, 33*, 121–128.

[1990] Hughes, H. C., Fendrich, R., & Reuter-Lorenz, P. A. (1990). Global versus local processing in the absence of low spatial frequencies. *Journal of Cognitive Neuroscience, 2*(3), 272–282.

[1992] Kimchi, R. (1992). The primacy of wholistic processing and the global/local paradigm: A critical review. *Psychological Bulletin, 112*(1), 24–38.

[1982] Kimchi, R., & Palmer, S. E. (1982). Form and texture in hierarchically constructed patterns. *Journal of Experimental Psychology: Human Perception and Performance, 8*(4), 521–535.

[1979] Kinchla, R. A., & Wolfe, J. M. (1979). The order of visual processing: 'top-down', 'bottom-up', or 'middle-out'. *Perception & Psychophysics, 25*, 225–231.

[1993] Lamb, M., & Yund, E. W. (1993). The role of spatial frequency in the processing of hierarchically organized stimuli. *Perception & Psychophysics, 54*, 773–784.

[1988] Livingstone, M., & Hubel, D. (1988). Segregation of form, color, movement, and depth: Anatomy, physiology, and perception. *Science, 240*, 740–749.

[1979] Martin, M. (1979). Local and global processing: The role of sparsity. *Memory & Cognition, 7*, 476–484.

[1981] McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: Part 1. an account of basic findings. *Psychological Review, 88*, 375–407.

[1986] McClelland, J. L., & Rumelhart, D. E. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 2. psychological and biological models*. Cambridge, MA: MIT Press.

[1981] Miller, J. (1981). Global precedence in attention and decision. *Journal of Experimental Psychology: Human Perception and Performance, 7*, 1161–1174.

[1977] Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cognitive Psychology, 9*, 353–383.

[1986] Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 1. foundations*. Cambridge, MA: MIT Press.

[1986] Rumelhart, D. E., & McClelland, J. L. (1986). *Parallel distributed processing: Explorations in the microstructure of cognition: Vol. 1. foundations*. Cambridge, MA: MIT Press.

[1993] Sanocki, T. (1993). Time course of object identification: Evidence for a global-to-local contingency. *Journal of Experimental Psychology: Human Perception and Performance, 19*(4), 878–898.

[1985] Sebrechts, M. M., & Fragala, J. J. (1985). Variations on parts and wholes: Information precedence vs. global precedence. In *Proceedings of the seventh annual conference of the cognitive science society*.

[1990] Spillman, L., & Werner, J. S. (1990). *Visual perception: The neuropsychological foundations*. San Diego: Academic Press.

[1982] Ward, L. M. (1982). Determinants of attention to local and global features of visual forms. *Journal of Experimental Psychology: Human Perception and Performance, 8*, 562–581.

[1987] Watt, R. J. (1987). Scanning from coarse to fine spatial scales in the human visual system after the onset of a stimulus. *Journal of the Optical Society of America A, 4*(10), 2006–2021.