

# Categorical Versus Coordinate Spatial Relations: Computational Analyses and Computer Simulations

Stephen M. Kosslyn, Christopher F. Chabris, and  
Chad J. Marsolek  
Harvard University

Olivier Koenig  
University of Geneva, Geneva, Switzerland

Results of 4 sets of neural network simulations support the distinction between categorical and coordinate spatial relations representations: (a) Networks that were split so that different hidden units contributed to each type of judgment performed better than unsplit networks; the reverse was observed when they made 2 coordinate judgments. (b) Both computations were more difficult when finer discriminations were required; this result mirrored findings with human Ss. (c) Networks with large, overlapping "receptive fields" performed the coordinate task better than did networks with small, less overlapping receptive fields, but vice versa for the categorical task; this suggests a possible basis for observed cerebral lateralization of the 2 kinds of processing. (d) The previously observed effect of stimulus contrast on this hemispheric asymmetry could reflect contributions of more neuronal input in high-contrast conditions.

Vision, like all other complex mental functions, is accomplished by a method of divide and conquer. Many relatively simple component systems work together to process information (cf. Maunsell & Newsome, 1987; Van Essen, 1985). A major division of labor is achieved by systems in the temporal and parietal lobes, whereby the former encodes object properties (such as shape and color) and the latter encodes spatial properties (such as location and size; see Kosslyn, Flynn, Amsterdam, & Wang, 1990; Maunsell & Newsome, 1987; Ungerleider & Mishkin, 1982). Both of these major systems can be further divided into component subsystems. In this article, we focus on the system that encodes spatial properties and argue that this system is divided into at least two subsystems that compute different kinds of representations of spatial relations.

The research reported here builds on the analyses and findings of Kosslyn (1987) and Kosslyn, Koenig, et al. (1989). We further develop their conception of the two kinds of spatial relations representations by considering implications of the new results reported by Sergent (1991), additional analyses of what is required to build a system that behaves in particular ways (Marr, 1982), and additional facts about the neural substrate. Sergent's findings have led us both to characterize

the distinction more rigorously and to link the distinction more tightly to properties of the brain.

Computational considerations suggest that different kinds of representations of spatial relations would be useful for different purposes. Consider two contexts in which people use spatial information. First, people must use spatial information to guide actions, ranging from moving the eyes to reaching and navigating. In fact, many cells in the posterior parietal lobes appear to have some role in movement control, firing either before or after a movement or registering the position of an effector (see Andersen, 1987; Hyvarinen, 1982). For guiding action, metric spatial information must be specified; simply knowing that a table is "next to" a wall does not help one walk right up to it without bumping into the edge.

Second, people often need to encode spatial relations to identify an object or scene. For this purpose, the brain does not need to represent metric information precisely; differences in the precise positions of two objects or parts often are not relevant (and in fact are potentially harmful) for distinguishing them from other objects or parts (cf. Biederman, 1987). Rather, spatial relations are assigned to a category, such as "connected to," "left of," or "above." For some purposes, it may also be useful to assign a spatial relation to a distance category, such as "one inch away," but this sort of category must be distinguished from the kind of analog encoding of metric distance that is necessary to guide action (e.g., see Kosslyn & Koenig, in press; Osherson, Kosslyn, & Hollerbach, 1989).

These considerations lead us to the hypothesis that the brain represents spatial relations in two ways. First, *coordinate* representations specify precise spatial locations in a way that is useful for guiding action. The units of these representations are not equivalence classes; rather, they delineate the finest possible division of space (subject to the resolution limitations of the visual system). These representations do not correspond to particular movements; rather, they specify spatial coordinates in a way that can be used to guide a variety of movements (see Chapter 7 of Kosslyn & Koenig, in press). Second,

---

This work was supported by Grant 88-0012 from U.S. Air Force Office of Scientific Research, Grant BNS 90-09619 from the National Science Foundation, and the National Institute of Neurological Diseases and Stroke Grant 2 P01-17778-09 awarded to Stephen M. Kosslyn. Chad J. Marsolek was supported by a Jacob K. Javits Graduate Fellowship, and Olivier Koenig was supported by Fellowship 81.357.0.86 from the Swiss Science Foundation.

We thank Kris Kirby and Randy O'Reilly for technical assistance; Mark Beeman, Joseph Hellige, and Justine Sergent for helpful comments on an earlier draft; and Robert Rosenthal for valuable statistical advice.

Correspondence concerning this article should be addressed to Stephen M. Kosslyn, Department of Psychology, Harvard University, Cambridge, Massachusetts 02138.

*categorical* representations assign a range of positions to an equivalence class (such as connected/unconnected, above/below, left/right). For many objects, parts retain the same categorical spatial relations, no matter how the object contorts; thus, the specification of categorical spatial relations is a critical aspect of a robust representation of an object's shape (cf. Marr, 1982). For example, even though its position in space varies widely, a cat's paw remains connected to (a categorical spatial relation) its foreleg regardless of whether the cat is curled up asleep, running, or batting an insect.

We distinguish coordinate representations from those used in recognition partly because the spatial information used to guide action appears to be "encapsulated" (e.g., McLeod, McLaughlin, & Nimmo-Smith, 1985); that is, the information used to guide action is not readily accessible to the systems used to categorize stimuli.

Neuropsychological findings have supported the contention that the two kinds of spatial representations are encoded by separate processing subsystems. Hellige and Michimata (1989), Koenig, Reiss, and Kosslyn (1990), and Kosslyn, Koenig, et al. (1989) all found that subjects judge distances relatively faster when the stimuli are presented initially to the right cerebral hemisphere (i.e., in the left visual field), whereas they evaluate some categorical spatial relations relatively faster when stimuli are presented initially to the left cerebral hemisphere (i.e., in the right visual field) or equally well when stimuli are presented to the left or right hemispheres. Although the left-hemisphere advantage is rarely significant in a single experiment, a trend toward a left-hemisphere advantage for categorical relations was evident in six experiments in which low-contrast stimuli were presented to normal adult subjects (Hellige & Michimata, 1989; Experiments 1, 2, 3, 4 of Kosslyn, Koenig, et al., 1989; Koenig et al., 1990), whereas a trend toward a right-hemisphere advantage was evident only once, in a difference of less than 1 ms (Experiment 4 of Sergent, 1991). According to the binomial distribution (sign test), the probability that this pattern of results is attributable to chance is .06.<sup>1</sup>

More recently, however, Sergent (1991) reported that this dissociation occurs only when stimuli are relatively degraded, and she inferred that this result does not reflect a distinction between two qualitatively different ways of representing spatial relations. Sergent assumed only that when the stimuli are degraded, the right hemisphere can more effectively encode precise spatial location. She offered as one piece of evidence against the distinction the fact that more difficult discriminations (defined by relative distance) affect categorical judgments as well as metric judgments. In addition, Sergent offered arguments that the distinction itself is conceptually flawed.

In this article, we report the results of computer simulations that support three assertions about the distinction between categorical and coordinate spatial relations representation. First, there is a clear conceptual distinction between categorical and coordinate spatial relations representations; second, both sorts of representations are more difficult to compute when fine discriminations must be made, just as was found previously with human subjects; and, third, the effects of stimulus quality can be accounted for easily by reference to a

simple computational mechanism. However, Sergent's (1991) findings have led us to reconsider Kosslyn's (1987) original motivation for the distinction between categorical and coordinate spatial representations, which was based on the idea that the left hemisphere is specialized for language and the right is involved in navigation.

### Study 1

We have argued that categorical and coordinate spatial relations are qualitatively distinct, and we hypothesized that they are encoded by different processing subsystems. Kosslyn et al. (1990) assumed that these subsystems correspond to separate neural networks, each of which maps an input (in this case, a representation of a pair of locations) to an output (in this case, a representation of a spatial relation). We also hypothesize that separate networks are used to perform qualitatively different types of input/output mappings, such as when different types of spatial relations representations are computed.

Accordingly, it is appropriate to use computer simulations of "neural networks" to investigate whether the two types of computations are in fact qualitatively distinct. These models establish mappings—from sets of input stimuli to correct responses—that appear to share critical features with the corresponding mappings in the brain. For example, Lehky and Sejnowski (1988), O'Reilly, Kosslyn, Marsolek, and Chabris (1990), and Zipser and Andersen (1988) all found that their networks developed an internal organization that mimicked properties of neurons that are thought to be involved in performing the relevant tasks. These networks apparently extracted specific aspects of the input to achieve the mapping, and the brain also extracts those properties when performing the corresponding mapping (for further discussion of this point, see Rueckl & Kosslyn, 1992). It is possible that a simple principle, such as gradient descent or distributed representation, is responsible for such correspondences between the models and the brain, but we need not press this issue further here; for our purposes, all that is important is that these models can be used appropriately to study properties of input/output mappings performed by the brain.

Although we cannot guarantee that the results of studying network simulation models necessarily generalize to the brain, at the very least such results enable us to (a) discover whether

<sup>1</sup> Sergent (1991) reported data from normal subjects for four categorical tasks, three of which involved high stimulus contrast (two tasks from Experiment 1 and one in Experiment 2) and one of which involved low stimulus contrast (Experiment 4). She obtained a 4.5-ms hemispheric difference in one task and differences of less than 1 ms in the other three. Using the Fisher method for combining independent probabilities discussed by Rosenthal (1984, p. 96), we estimated the probability of obtaining differences this small or smaller according to chance. Because some of the *F* values necessary for this calculation were missing, we estimated them on the basis of the available data; the obtained probability of finding differences this small or smaller was only 0.013,  $\chi^2(8) = 19.39$ . This finding leads us to suspect that Sergent's findings with high contrast reflect a floor effect for this sort of processing.

certain conclusions or findings can follow from specific assumptions and (b) rule out the possibility that certain conclusions or findings *cannot* follow from specific assumptions.

In the first set of simulations we obtained computational evidence that at least one categorical relation, above/below, is conceptually distinct from the specification of metric location, which is a critical component of coordinate representations.<sup>2</sup> We ran simulations of the bar-and-dot tasks developed by Hellige and Michimata (1989) and also used by Kosslyn, Koenig, et al. (1989, Experiment 3) and Koenig et al. (1990). In these tasks, subjects saw a short horizontal bar and a dot and were asked to determine either whether the dot was above or below the bar (the categorical task) or whether the dot was within a fixed distance from the bar (the coordinate task). In our simulations, the categorical task required the network to judge whether an activated input unit (the dot) was "above" or "below" a landmark (the bar), which consisted of two activated input units flanked by one inactivated input unit to either side. In contrast, the coordinate task required the network to judge whether or not the activated input unit was within four units of the landmark. We considered this a coordinate task because the network had to encode the finest possible distinctions among locations; in contrast, the above/below task required grouping the locations into categories.<sup>3</sup>

We studied these tasks by using the "partition" paradigm developed by Rueckl, Cave, and Kosslyn (1989), in which the efficacies of two types of networks on a pair of tasks are compared; each type performs both tasks simultaneously, using separate sets of output units. In one type of network, all of the hidden units are connected to all of the output units, whereas in the other type, the hidden units are split into two groups. In a "split" network, one group of hidden units is connected exclusively to the output units for one task, whereas the other group is connected exclusively to the output units for the other task. Consequently, the representations developed by the hidden units for one input/output mapping cannot be used for the other, and vice versa. In the "unsplit" networks, in contrast, the hidden units form a single group that is fully connected to all the output units for both tasks.<sup>4</sup>

Following the reasoning of Rueckl et al. (1989), we expected that if two tasks rely on distinct computations, a split network would perform the necessary mapping better than would an unsplit network. The segregation of processing prevents patterns of weights that are useful for accomplishing one input/output mapping from interfering with those that are useful for accomplishing the other mapping. However, this effect may not be evident until the networks have enough hidden units, because a split network has the inherent disadvantage of having fewer connections (and consequently fewer weight space dimensions) than the corresponding unsplit network. With sufficient numbers of hidden units, the advantage of separating distinct representations should overcome the disadvantage of having fewer resources. Accordingly, we systematically varied the number of hidden units.

To establish the input/output mappings, we used the backward error propagation algorithm of Rumelhart, Hinton, and Williams (1986), as modified by Stornetta and Huberman (1987). This algorithm is sometimes characterized as a "learning" procedure, and its behavior is often compared with that

of biological systems that learn (e.g., McClelland & Rumelhart, 1986; Rumelhart & McClelland, 1986; Seidenberg & McClelland, 1989). We do not assume that the kind of learning performed by the networks has a direct relation to learning in actual neural networks in the brain. Instead, we use the difficulty of learning in the models solely as a measure of how difficult it is to establish a specific input/output mapping. We treat the amount of error after a fixed number of training trials as a measure of the difficulty of establishing the mapping (cf. Rueckl & Kosslyn, 1992). As noted earlier, we are interested in this type of mapping because it seems to reflect properties of mappings performed by the brain, but we do not regard the *training process* itself as necessarily having any direct correspondence to neural events. If the mappings are distinct, we will have evidence that the conceptual distinction between categorical and coordinate representations is sound, and will be able to rule out the possibility that the two types of representations are logically the same and must be produced by the same computation.

In Part 1 of this study, we compared the ease of establishing categorical and coordinate (metric) mappings in split and unsplit networks with various numbers of hidden units. Each network was trained to establish both types of mappings, and we observed the amount of error after a fixed number of trials. In Part 2, we considered the possibility that any advantage shown by split networks might have nothing to do with the distinct types of representations: Perhaps dividing resources between two tasks is always beneficial, regardless of the nature of the tasks and the degree of similarity between the mappings. To address this question, we compared split and unsplit networks that performed two different variants of the coordinate task, using the same input patterns and network architectures as in Part 1.

## Part 1

### Method

*Materials.* We created standard three-level networks that had 28 input units and 4 output units (2 for each task); we parametrically varied the number of hidden units, examining the range from 8 to 17. For each size, we created one unsplit network and split networks


<sup>2</sup> We do not mean to imply that a single subsystem necessarily computes all categorical spatial relations. For example, "inside/outside" and "connected/unconnected" may rely on different sorts of mappings than does "above/below." However, we have yet to develop the theory further in this direction.

<sup>3</sup> Although both judgments require a categorization at some level of processing, the above/below judgment requires only approximate localization of the dot, and so the right hemisphere's superior ability to compute precise location is not relevant, whereas the left hemisphere's superior ability to categorize should affect performance. In contrast, although the coordinate task requires categorization to produce the output, the difficult part of the judgment is to evaluate the precise distance. Thus the right hemisphere's abilities should be more important for good performance in this task.

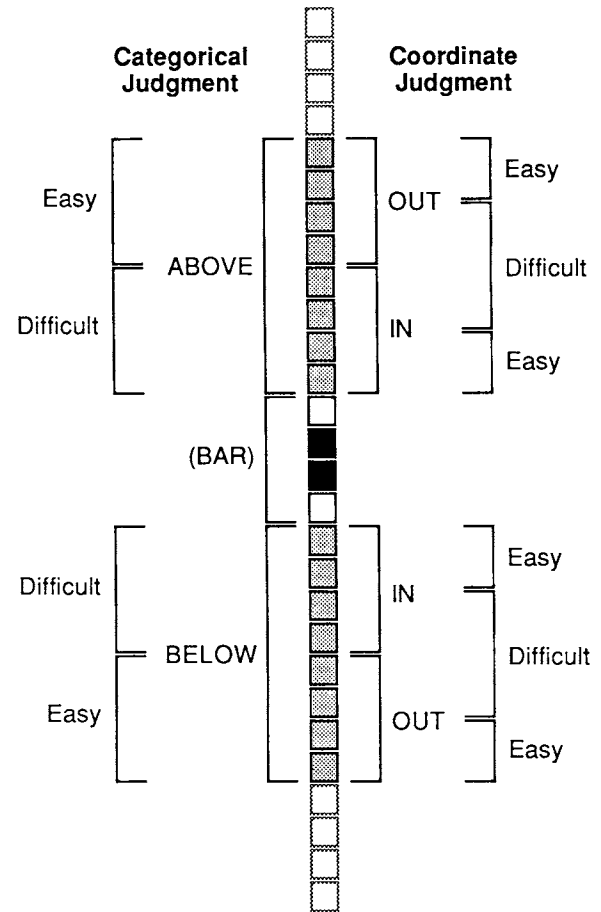
<sup>4</sup> Note that a split network with *a* hidden units in one partition and *b* hidden units in the other partition is computationally identical to two separate networks with *a* and *b* hidden units.

with 3–6 different ratios of hidden units allocated to the categorical and coordinate tasks. As illustrated in Figure 1, the split networks were identical to the unsplit networks in all respects except for the missing connections across the partition.

As illustrated in Figure 2, the input units were conceptualized as simulating a vertical spatial array. A bar was represented by a four-unit set in which the top and bottom units were always “off” and the middle two units were always “on.” This set could appear at any of five positions: at the center of the array, two units up, two units down, four units up, or four units down. A dot was represented by a single activated unit, which could appear at any of eight positions above and eight positions below the bar. Thus there were  $5 \times (8 + 8)$ , or 80, input patterns in all. Each input pattern was associated with



*Figure 1.* The architecture of the unsplit (top) and split (bottom) networks. (Numbers in parentheses indicate the size of the layers; note that a range of sizes of hidden layers was tested. CAT = categorical; COO = coordinate.)



*Figure 2.* The input to the networks, indicating the bar and possible locations for the dot. (The judgments were based on the relative location of the dot and bar, as noted in the figure. The bar could appear in five distinct locations: at the center and shifted two or four elements up or down from the center.)

an output pattern, which represented the target values for the four output units, two of which indicated the categorical judgment and two of which indicated the coordinate judgment. For the categorical judgment, one output unit indicated *ABOVE* and the other *BELOW*. For the coordinate judgment, one output unit indicated *IN* and the other *OUT*; the *IN* response was correct if the dot was within four elements of the bar, otherwise the *OUT* response was correct.

*Procedure.* We tested 18 networks of each type (unsplit and each ratio of split). Testing consisted of training with the backward error propagation algorithm for 50 epochs, at which point we measured the average squared error per output unit per input pattern in the network. An epoch was defined as a complete run through all 80 stimuli in the training set, followed by a single backward error propagation pass. The output signals of the units ranged from  $-0.5$  to  $+0.5$  (in accordance with Stornetta & Huberman's 1987 procedure); for each output unit, a threshold of  $\pm 0.4$  was used, so that once the output of the unit was within 0.1 of its target, it was considered to have no error at all. Before each new test, the weights on the connections in the network were reset to random values between  $-0.5$  and  $+0.5$ . The parameters epsilon ("learning rate") and alpha ("momentum factor") were set to 0.25 and 0.90, respectively.

## Results

The results were analyzed in an analysis of variance (ANOVA), with replication (different tests, each using new initial random weights) as the random effect and average squared error per output unit per input pattern as the dependent measure. For each number of hidden units, we compared the unsplit case to the best performer among the split networks; different ratios of categorical to coordinate task hidden units sometimes enabled different-sized networks to perform best. Figure 3 illustrates the results of this analysis and also indicates which ratio was optimal at each network size. As expected, if categorical and coordinate spatial relations computations are qualitatively distinct, the split networks generally produced significantly less error ( $M = 0.039$  averaged squared error), than did the unsplit ones ( $M = 0.057$ ),  $F(1, 340) = 7.64$ ,  $p < .01$ . Furthermore, error decreased linearly with more hidden units.  $F(1, 340) = 5.01$ ,  $p < .05$ , for the appropriate contrast, but there was no interaction between network type and the number of hidden units ( $F < 1$ ).

## Part 2

### Method

**Materials.** To examine the possibility that split networks generally perform better when two judgments of any type must be made by a single network, we created split and unsplit networks that performed two coordinate judgments. The networks evaluated both whether the dot was within two elements and whether it was within six elements of the bar; thus the only major change from Part 1 was that the output units coded for two distance judgments, rather than the categorical and original distance judgment. The split networks were created by partitioning the hidden units in the same way as in Part 1.

The networks had 10, 12, 14, 16, 18, and 20 hidden units. In order to reduce the number of models that were tested, half of the hidden units in all but one size of the split networks were allocated to each of two tasks; for the 20-hidden-unit size, we also created split networks with ratios of 5:15, 8:12, 12:8, and 15:5 hidden units. On the basis of Rueckl et al.'s (1989) results, we expected that allocating insufficient hidden units to a partition would severely impair the mapping, and thus we chose to manipulate the ratios of the largest split networks to consider this potential problem.

**Procedure.** Ten networks of each type and size were tested initially; for the 20-hidden-unit size, we tested an additional 6 networks of each type as well as 16 networks of each of the supplementary partition ratios. Testing was conducted as in Part 1, except that the parameter epsilon was set to 0.10 (this modification was necessary because pilot data indicated that the networks would often reach high local minima early in testing with the 0.25 epsilon value).

### Results

The data were analyzed as in Part 1. For each network size, we compared the unsplit networks with the even-ratio split networks in an ANOVA with replication as the random effect. Figure 4 illustrates the results of this analysis. In sharp contrast to the results from Part 1, the unsplit networks consistently produced less error ( $M = 0.058$ ) than did the split networks ( $M = 0.092$ ),  $F(1, 108) = 63.9$ ,  $p < .001$ . Error also decreased




Figure 3. Results from Part 1 of Study 1, illustrating that the split networks established the categorical (CAT) and coordinate (COO) mappings more effectively than the unsplit networks, provided that a sufficient number of hidden units was available. (The ratios under the  $x$  axis show the numbers of hidden units allocated to the categorical and coordinate tasks, respectively, in the split networks that performed best in each case; in only one case—the 10 hidden unit networks—even the most effective split networks performed worse than the unsplit networks.)

linearly with more hidden units,  $F(1, 108) = 58.1$ ,  $p < .001$ , for the appropriate contrast. However, the interaction between network type and the number of hidden units was not significant,  $F(5, 108) = 1.73$ ,  $p > .10$ .

One could argue that the poor performance of the split networks was caused by the division of the hidden units into two equal groups. We ruled out this possibility by examining the other ratios tested in the 20-hidden-unit networks. The split networks with 20 hidden units performed worse than the unsplit networks, regardless of the ratio. The mean error was

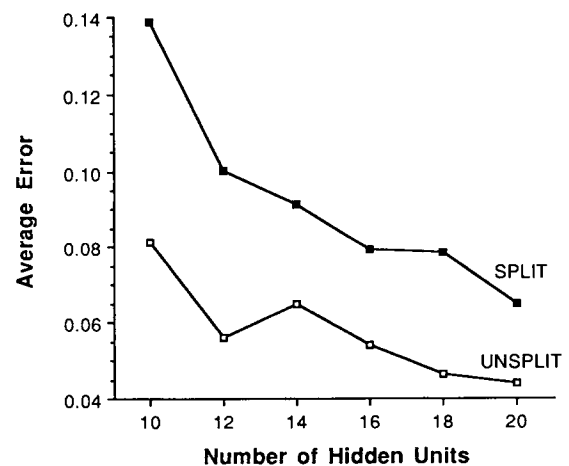


Figure 4. Results of Part 2 of Study 1, illustrating that the unsplit networks established the mappings more effectively when two coordinate judgments were required.

0.042 for unsplit networks, in comparison with mean error of 0.062 for a 12:8 split, 0.068 for a 10:10 split, 0.083 for a 8:12 split, 0.085 for a 15:5 split, and 0.089 for a 5:15 split. Four of the comparisons revealed significant differences according to *t* tests with the Bonferroni adjustment (adjusted  $p < .05$ ), and the other comparison (between unsplit and 12:8 ratio split networks) approached significance,  $t(30) = 2.49$ , adjusted  $p < .10$ .

### Discussion

The results of these simulations provide evidence that categorical and coordinate spatial relations are conceptually distinct. In Part 1 we found that networks in which the two types of processing are segregated perform both mappings better than do otherwise equivalent networks in which the representations are intermingled in a single set of units and weights. In Part 2 we found that this result was not merely a consequence of a general advantage for segregating the processing of any two tasks.

For the split networks in Part 1, the optimal ratio of dividing the hidden units was usually close to an even split between the two mappings. This finding may, however, have been an accident of the particular way in which we set up the tasks (e.g., the size of the input array, the amount that the bar moved, and so on), and we do not wish to draw strong inferences from this result; indeed, Kosslyn, Koenig, et al. (1989) manipulated the difficulty of discrimination to make the two kinds of tasks equivalent or either one more difficult than the other. However, it is fortunate that the two mappings were about equally difficult in these networks because it suggests that our results do not somehow reflect an effect of having a difficult task mixed with an easier one.

In addition, the advantage of the unsplit model in Part 2 cannot be ascribed to the "overlap" in outputs between the two tasks; whenever the response for the two-element distance judgment was *IN*, it must also have been *IN* for the six-element judgment and vice versa for the *OUT* judgment. The categorical and coordinate tasks in Part 1 overlapped to the same degree: For half of the input patterns, the output unit targets had the same pattern for both judgments.

These results do not imply that the human brain necessarily has evolved to use separate networks to compute the different types of spatial relations; even though a split network would be better, the brain may not be optimized in this way (we return to this issue in the General Discussion section). Thus these results simply support the contention that the two kinds of representations are not logically intertwined and need not rely on the same computation. Testing a simple prediction of this result would be of interest. If different networks are in fact used for the two tasks, then repetition priming should not transfer well between them. If a subject practices making metric judgments, this practice should prime other judgments in which the same network is used but not judgments in which a different network is used.

### Study 2

Sergent (1991) suggested that our theory makes a straightforward prediction: If separate subsystems encode categorical

and coordinate spatial relations, categorical judgments should not be influenced by the distance separating the two objects. However, our theory does not imply that distance affects only coordinate judgments; many categorical spatial relations rely on dividing space into discrete bins, and this process may be more difficult when the boundaries of these regions must be delineated more precisely. For example, a dot can be classified as above or below a landmark by observing whether it falls into one of two pockets of space; although any location within each bin is treated as equivalent, it may be more difficult to assign a dot to a category if the bins must be delineated carefully. Furthermore, even after the regions of space are delineated, it may be more difficult to assign a dot to a category if it appears near the boundary. Such center-versus-periphery effects are found in a wide range of categorization tasks (e.g., see Smith & Medin, 1983). Thus we do not assume a pure distinction between "encoding" and "judgment" processes; interactions between the two may influence the speed of response.

In this study, we used networks to discover whether our theory always implies that discriminability has different effects on the two types of computations or whether distance may in fact have similar effects on both computations. We compared the effect of discriminability in individual networks that performed either the categorical or the coordinate mapping (not both simultaneously, as in Study 1). We tested these specialized networks on two complementary subsets of the complete categorical and coordinate tasks: those stimulus patterns that Sergent's (1991) results suggest should be relatively easy and those that her results suggest should be relatively difficult.

In the categorical task, the easy discriminations were those in which the dot was far from the bar, and the difficult discriminations were those in which the dot was near the bar. In the coordinate task, the easy discriminations were those in which the dot was far from the criterion distance (defined in relation to the bar), and difficult judgments were those in which the dot was close to the criterion distance. This simulation enabled us to address directly the prediction that Sergent inferred.

### Method

#### Materials

The networks were identical to those of Study 1, except that we examined only standard unsplit networks with two output units and varied the number of hidden units from 6 to 12. Figure 5 illustrates this network architecture. We chose this range of hidden layer sizes because it was sufficient to enable the mappings to be achieved but still show differences if they are not equally easy; when there are too many hidden units, a "floor effect" may obscure differences in computational difficulty among mappings.

The categorical task was divided into easy and difficult conditions. In the easy condition, the dot was farther than four elements above or below the bar; in the difficult condition, it was within four elements of the bar. Separate networks were trained in the easy and difficult conditions, each receiving a set of 40 input patterns during training. The coordinate judgment task was similarly divided into easy and difficult conditions. In the easy condition, the dot was farther than two elements from the criterion distance (which

was four elements from the bar); in the difficult condition, the dot was within two elements of the criterion distance. Again, separate networks were trained to map the easy and difficult conditions; each received 40 input patterns during training.

### Procedure

We tested 25 networks of each size in each categorical condition and 25 networks of each size in each coordinate condition. Testing proceeded as in Part 1 of Study 1, except that error was measured after 30 epochs rather than 50 because the amount of error in the easy conditions was too low to compare after 50 epochs.

### Results

The data were analyzed as in the previous studies. As is evident in Figure 6, the networks generally performed better when given the easier discriminations ( $M = 0.025$  error) than when given the difficult discriminations ( $M = 0.078$ ),  $F(1, 532) = 190.9, p < .001$ , and there was no hint of an interaction between the difficulty of the discrimination and the type of task ( $F < 1$ ). In addition, in this study the categorical mappings were easier than the coordinate mappings ( $M$ s = 0.043 and 0.060 error),  $F(1, 532) = 20.9, p < .001$ , and overall error decreased linearly with more hidden units,  $F(1, 532) = 21.1, p < .001$ , for the appropriate contrast. No other effects or interactions approached significance.

### Discussion

These findings reveal that subtle discriminations can impair both kinds of judgments in networks. This result is important because it is not intuitively clear that a categorical spatial relation should be harder to establish for stimuli that appear in a smaller range of positions. The findings of Study 3 will support our characterization of this sort of categorical spatial relation as delineating regions of space; hence we can interpret the results of Study 2 as suggesting that the more difficult it

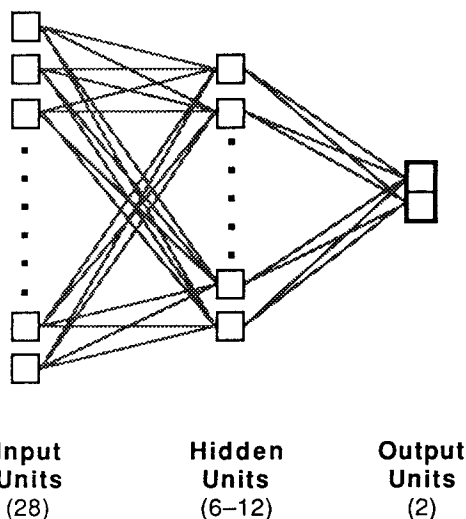


Figure 5. The architecture of the networks used in Study 2. (Only one type of judgment was performed in these models.)

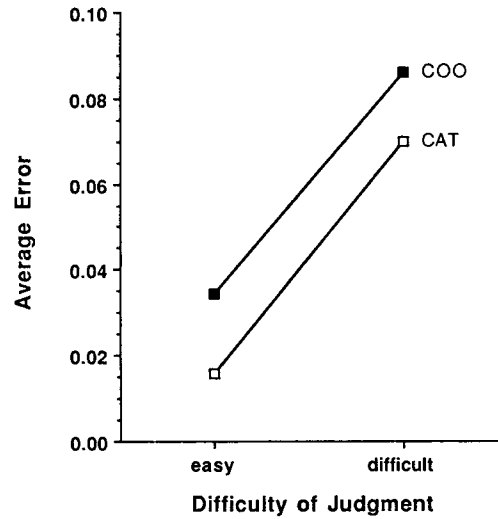


Figure 6. The results of Study 2, showing that both categorical (CAT) and coordinate (COO) judgments are affected by the difficulty of the discrimination (the range of distances).

is to delineate the regions to be related, the more difficult it is to establish the mapping.

In short, the simulation results of Study 2 show that Sergeant's (1991) finding that humans display a similar effect does not contradict the distinction between categorical and coordinate representations of spatial relations. It would be of interest to discover whether the effects of the difficulty of discrimination on the categorical and coordinate judgments can dissociate after brain damage; if our interpretation is correct, some brain-damaged patients should show selective deficits for difficult categorical discriminations but not difficult coordinate discriminations, and vice versa for other brain-damaged patients.

### Study 3

The distinction between categorical and coordinate spatial relations representations was formulated on the basis of an analysis of the purposes of spatial representations. We hypothesized that coordinate representations play a special role in action control, whereas categorical representations play a special role in recognition and identification. In this study we considered some implications of these ideas in more detail.

Action control depends on precise representation of spatial location. One way to represent spatial location precisely depends on overlap among rather coarse representations of location (Hinton, McClelland, & Rumelhart, 1986). This sort of *coarse coding* underlies color vision, for example, in which the three types of cones in the retina have overlapping distributions of sensitivity to different wavelengths of light; it is this overlap that enables the three types of cones to encode a wide range of colors. O'Reilly et al. (1990) used network models to show that this mechanism was indeed an effective way of encoding metric spatial location.

It is possible that differences in the use of coarse coding can account for the hemispheric differences in computing the two types of spatial relations. For instance, the right hemisphere

may use more input from low-level visual neurons that have relatively large receptive fields (i.e., that receive input from relatively large regions of space), which have a large degree of overlap. These broadly tuned receptive fields would enable effective coarse coding and could be responsible for the right hemisphere's superior ability to encode precise location. In contrast, the left hemisphere may use more input from low-level visual neurons that have relatively small receptive fields, which do not overlap as much. Sets of these receptive fields would define particular areas, which could be used to specify regions that are above or below a reference point, left or right of a reference point, and so on. In the limit, if the receptive fields did not overlap at all and the categories corresponded to discrete regions of space, such mappings would be "linearly separable"—so straightforward that they could be accomplished by direct connections from the input units to the output units, without a hidden layer (see Minsky & Papert, 1969; Rumelhart et al., 1986).

The idea that differences in receptive field properties may be at the root of differences in hemispheric specialization for spatial encoding is intriguing for a number of reasons. First, it is compatible with Sergent's (1982) finding that the left hemisphere encodes smaller, high-spatial-frequency patterns better than the right, and the right hemisphere encodes larger, low-spatial-frequency patterns better than the left (see also Delis, Robertson, & Efron, 1986; Van Kleeck, 1989). Higher spatial frequencies may be encoded more effectively by smaller receptive fields; these fields register smaller variations in space than larger ones. By the same token, lower spatial frequencies may be encoded more effectively by larger receptive fields (cf. De Valois & De Valois, 1988).

Second, neurons with large receptive fields would be useful in preattentive processing, which by definition must monitor large regions of space. Such processing plays a critical role in controlling actions; one often moves one's eyes, head, and body toward a stimulus that suddenly appears, moves, or changes in some other way. Furthermore, one wants to look at or reach toward an object with reasonable accuracy, even if it is seen out of the corner of one's eye. M. Livingstone (personal communication, May 1990) has suggested that the magnocellular ganglia may project preferentially to the right hemisphere (see Livingstone & Hubel, 1988); these neurons have relatively large receptive fields and are thought to be involved in preattentive processing (see also de Schonen & Mathivet, 1989).

Thus it is of interest that overlapping large receptive fields not only enable the system to monitor a large area but also can produce the necessary precision to guide an initial movement, even if a target is seen only out of the corner of one's eye. In keeping with this idea, Fisk and Goodale (1988; see also Goodale, 1988) reported that patients with right hemisphere damage are impaired in the initial phases of reaching toward a visual target.

These findings must be evaluated in the context of Kitterle, Christman, and Hellige's (1990) failure to find any difference in the sensitivities of the two hemispheres to different spatial frequencies in a simple detection task. However, Kitterle et al. did find that high-spatial-frequency gratings were *identified* faster and more accurately when they were presented in the

right visual field (and hence were processed initially by the left hemisphere), whereas in some conditions low-spatial-frequency gratings were identified more readily when they were presented in the left visual field (and hence were processed initially by the right hemisphere). Thus, in keeping with Sergent's (1982) ideas, the hemispheric asymmetry cannot be ascribed to low-level processing; rather, it depends on high-level encoding and comparison processes.

In short, we hypothesized that the right hemisphere uses more input than the left from low-level visual neurons that have large receptive fields. These large receptive fields overlap, which enables the right hemisphere to encode coordinate spatial relations better than the left. In contrast, the left hemisphere uses more input from low-level neurons with small receptive fields. These receptive fields have less overlap than do the larger ones, which enables the left hemisphere to specify some categorical relations by delineating discrete sets of locations. For example, if an  $X$  is "left of" a  $Y$ , one may be able to represent the relation by defining two regions: one for the left and one for the right.

On this view, if a categorical spatial relation cannot be computed by defining discrete pockets of space, then the left hemisphere will not encode that relation better than the right. This hypothesis is consistent with Sergent's (1991) failure to find left-hemisphere superiority in tasks involving spatially complex stimuli (which had target shapes that could appear in several noncontiguous regions); these stimuli are difficult to delineate into specific regions of space that can be used to categorize stimuli.

These hypotheses rely on subtle distinctions and several steps of reasoning. Thus, they are ideal candidates for computer simulation modeling, which can show whether we have merely engaged in so much hand waving or whether our assumptions can have the consequences that we infer. In this study, we used network models to investigate these hypotheses in two ways. In Part 1, we examined the mappings performed by categorical and coordinate networks by analyzing the "receptive fields" developed by the hidden units of different networks; that is, for each hidden unit, we examined which regions of the input array most strongly influenced its level of activation. The larger the weight on the connection from an input unit to the hidden unit, the more strongly a dot in that location will affect the hidden unit; thus the pattern of weights on the connections to a hidden unit defined its receptive field (Rueckl et al., 1989; see also Lehky & Sejnowski, 1988; O'Reilly et al., 1990; Zipser & Andersen, 1988). Specifically, we tested the possibility that after training proceeds until the network makes no errors, networks that encode coordinate spatial relations will have larger receptive fields than will networks that encode categorical spatial relations. Recall that when a unit's output value was within 0.1 of its target value, it was assigned zero error; otherwise, it was assigned the square of the deviation from the target value. These error scores were then summed over output units and training patterns and divided by the number of each to yield the average error value. In Part 1, we trained the networks until this value was zero.

In Part 2 of this study, adapting the method of O'Reilly et al. (1990), we constructed networks that were "hard wired"



to have large or small receptive fields and considered how effectively they performed the two kinds of mappings.

### Part 1

#### Method

*Materials.* Two network models were constructed in this study: one that performed the categorical mapping and one that performed the coordinate mapping. These networks were identical to those in Study 2, except that they had 10 hidden units each, and they mapped the entire set of stimuli rather than only the easy or the difficult stimuli (as in Study 2).

*Procedure.* Testing was conducted as in Study 2, except that we continued training on each network until the average error had decreased to zero.

#### Results

We examined the receptive fields developed by the hidden units of each network. To test the hypothesis that the coordinate network developed relatively large receptive fields, we first normalized the weights on the connections between the input units and the hidden units in the two networks. This was necessary because the two networks developed different patterns of weights and had different maximal and minimal values. We then calculated the radius of a receptive field by determining the average number of contiguous weights from each peak weight value down to the peak value of  $1/e$  (cf. Andersen, Asanuma, Essick, & Siegel, 1990; Zipser & Andersen, 1988). For hidden units with multiple-peak receptive fields, we calculated the average crest size. The receptive fields developed by the coordinate network model (average radius 9.7 units) were larger than those developed by the categorical network model (average radius 4.8 units),  $t(18) = 2.61$ ,  $p < .05$ .

#### Discussion

As expected, the networks that performed the coordinate task spontaneously developed larger receptive fields than did the networks that performed the categorical task. However, as O'Reilly et al. (1990) found, these receptive fields often tended to have complex shapes (see also Zipser & Andersen, 1988). It is possible that the size differences are somehow related to the various sets of shapes that developed. To examine the effects of size in isolation, we manipulated this variable directly in Part 2 of this Study.

### Part 2

#### Method

*Materials.* We constructed a new set of network models that were the same as those used in Study 2 except for the following changes. First, as shown in Figure 7, each network had four layers of units. The first layer was conceptualized as a "retinal" array, and the 28 units in this layer received the same input patterns as did the 28 input units in Study 1. The weights on the connections between Layers 1 and 2 were not modified during the training procedure; these weights

defined fixed receptive fields of the units in Layer 2. Layers 2, 3, and 4 and the connections among them constituted a three-layer network that was trained with backward error propagation. Thus, only the connections between Layers 2 and 3 and between Layers 3 and 4 were modified during the course of training. There were 14 units in Layer 2, which functioned as the input units for the backward error propagation procedure, and there were 10 hidden units (Layer 3 units). (Layer 2 units can also be referred to as "input units" because these were the input units for the backward error propagation procedure.) As in Study 2, the two output units specified *ABOVE* or *BELOW* for the categorical task and *IN* or *OUT* for the coordinate task.

For half of the networks tested in this study, the Layer 2 units had small fixed receptive fields; for the other half, the Layer 2 units had large fixed receptive fields. For each network, 14 receptive fields were created, one for each Layer 2 unit. All receptive fields were defined by normalized Gaussian distributions that were determined as follows. First, standard Gaussian fields were created according to the formula

$$F_g(y) = \frac{e^{-(y^2/2\sigma^2)}}{2\pi\sigma^2},$$

in which  $y$  corresponds to the "vertical" position in Layer 1 and  $\sigma$  is a constant that determines the size of the receptive field. All fields were then normalized to the range of 0–1. We created small receptive fields by setting the value of  $\sigma$  to 0.71, whereas we created large fields by setting the value of  $\sigma$  to 1.42. Figure 8 illustrates the shape and the scope of the two types of receptive fields; as is evident, all receptive field peaks covered two retinal elements. We varied the locations of the receptive field peaks in Layer 1 by modifying the  $y$  value in the Gaussian formula; as illustrated in the left panel of Figure 9, the 14 receptive field peaks for each network were tessellated across the retinal array. In all other respects, the networks were constructed like those used in Part 1 of this study.

*Procedure.* The networks were trained with the stimulus set used in Part 1. Input patterns were presented to the "retinal" units of Layer 1 in the networks, and the receptive fields modulated the input sent to Layer 2. The connection weights between Layer 1 and Layer 2 units (Layer 2 receptive fields) were not modified during training. Thus because the sigmoidal activation function was used to facilitate the modification of connection weights during training, it was not used to modulate the flow of activation from Layer 1 to Layer 2 in this network. Because the sigmoidal activation function was not used, activation values for Layer 2 units were set to the 0–1 range through a linear normalization procedure: Activation was first computed as the sum of the three active connection weights, two from the bar and one from the dot in the retinal array. The maximal possible activation for a Layer 2 unit, given the receptive field size and the constraints of line and dot positioning in the retinal array, was calculated for each network. This value served as the maximal value in the normalization procedure. A minimal possible activation was calculated in an analogous manner, and this value served as the minimum in the normalization procedure.

We tested 10 networks in each of the four conditions, defined by large or small Layer 2 receptive field sizes and the type of judgment. In all other respects, testing was the same as in Part 1.

#### Results

The results are illustrated in Figure 10. As expected, larger receptive field sizes greatly facilitated the coordinate judgment, whereas smaller sizes were more useful for the categorical judgment,  $F(1, 36) = 62.4$ ,  $p < .001$ , for the interaction between size of input layer receptive fields and task. Separate tests revealed that the difference for the coordinate judgment




Figure 7. The architecture of the network used in Part 2 of Study 3. (Receptive fields were defined over the input units so that only some elements in the input array fed into units in Layer 2. The connection weights between Layers 1 and 2 were fixed throughout training, whereas the connection weights between Layers 2 and 3 and between Layers 3 and 4 were modified during training according to the standard three-layer back-propagation procedure. Fourteen Layer 2 units were used in Study 3, and 28 were used in Study 4.)

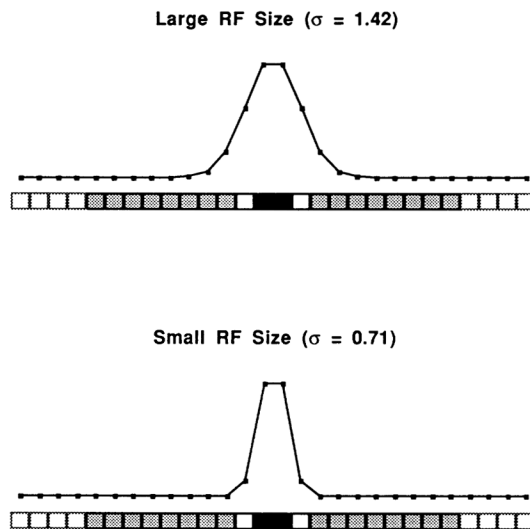


Figure 8. The sizes and profiles of the receptive fields used in Part 2 of Study 3.

( $M_s = 0.057$  error for smaller and  $0.013$  error for larger receptive field networks) was significant,  $F(1, 36) = 97.7$ ,  $p < .001$ , whereas the corresponding difference for the categorical judgment ( $M_s = 0.028$  and  $0.034$  error) was not,  $F(1, 36) = 2.22$ ,  $p > .10$ . Overall, networks with larger receptive fields performed better than those with smaller receptive fields ( $M_s = 0.024$  and  $0.042$  error),  $F(1, 36) = 34.0$ ,  $p < .001$ , but under these conditions the two types of judgments did not differ overall ( $M_s = 0.031$  for the categorical task and  $0.035$  for the coordinate task),  $F(1, 36) = 1.60$ ,  $p > .20$ .

### Discussion

Dovetailing with the results from Part 1 of this study, networks with fixed large receptive fields performed the coordinate task better than did networks with fixed small receptive field sizes, and there was a tendency for the opposite pattern in the categorical task. These findings are like those from the corresponding experiments with human subjects. Indeed, as noted in the introduction, the left hemisphere has only a weak (but consistent) advantage for encoding categorical spatial relations.




Figure 9. The distribution of receptive field peaks in the networks tested in Part 2 of Study 3 (left) and in Study 4 (right). (For mixed-size networks in Study 4, “medium” peak locations refer to the 14 peak locations for intermediate-sized receptive fields, “smaller” to the 7 peak locations for receptive fields that had relatively smaller sizes, and “larger” to the 7 peak locations for receptive fields that had relatively larger sizes. In the homogeneous networks, peak locations for receptive fields—all of which were the same size—were also distributed across these 28 locations.)

Thus these results provide additional support for the contention that categorical and coordinate spatial relations representations are conceptually distinct. Furthermore, they hint at one possible reason why the hemispheres are specialized for encoding the different types of representations. A prediction of these results is that the left-hemisphere advantage for categorical spatial relations should be present only when the relation can be computed by delineating regions of space. Thus any manipulation that makes this computation more difficult should eliminate this effect.

Study 4

The results of Study 3 may provide an insight into Sergent’s (1991) finding that hemispheric differences do not arise when the stimuli have high contrast. In Study 4, we examined networks with narrow or wide ranges of receptive field sizes in the input units. If the hemispheres differ in their sensitivities to input from units with different-sized receptive fields, then any differences in the sizes of the receptive fields of input units may be obliterated if high contrast enables a more varied set of input units to contribute to the computation.

Kosslyn (1987) and Kosslyn, Koenig, et al. (1989) assumed that the hemispheres differ in their relative efficacy at encoding the two types of spatial relations; they did not intend to claim that the hemispheres were exclusively specialized for the different types of encoding. Indeed, Kosslyn, Sokolov, and Chen (1989) simulated Kosslyn’s (1987) “snowball process,” in which the relative hemispheric specialization of various subsystems develops gradually over time; at the heart of this model is the idea that the two hemispheres differ in the relative efficacy of individual subsystems. It is unfortunate that we did not state this assumption clearly, as noted by Sergent (1991); this assumption is important because it leads us to expect that the difference in inputs to the two hemispheres is one of degree. Given the previous simulation results, we might expect the left hemisphere to use more input from neurons with relatively small receptive fields and the right hemisphere to use more input from neurons with relatively large receptive fields, but we expect a distribution of inputs from neurons with different receptive field sizes in both hemispheres.

This idea suggests that the “modulation transfer functions” (see Kaufman, 1974) of high-level visual areas may differ in the two hemispheres. We are not considering the modulation transfer functions of low-level areas involved in detection, but rather those of higher areas involved in memory and comparison. Figure 11 illustrates this hypothesis: The modulation transfer functions of the high-level areas are slightly shifted, so that the peak sensitivity for the right hemisphere is at a lower spatial frequency, which could reflect its use of more input from larger receptive fields (and vice versa for the left hemisphere). At low contrast, the performances of the two hemispheres would be well separated, as indicated by the relatively small amount of overlap at the top horizontal dotted line. This difference could arise because the hemispheres differ

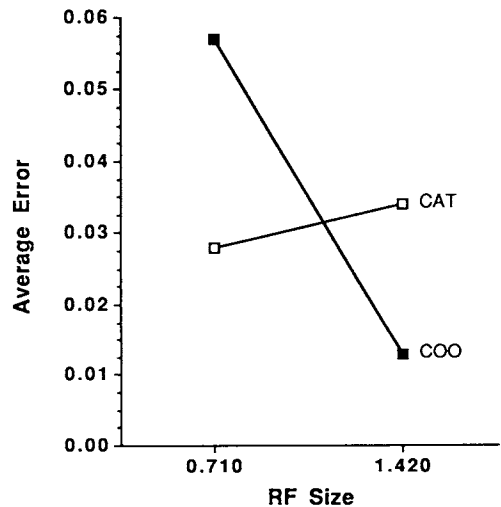


Figure 10. Results of Part 2 of Study 3, showing that the difficulty of establishing the two types of mappings depended on the relative sizes of the receptive fields (RFs). (Receptive field sizes are indicated by  $\sigma$  values for the Gaussian field equation. CAT = categorical; COO = coordinate.)

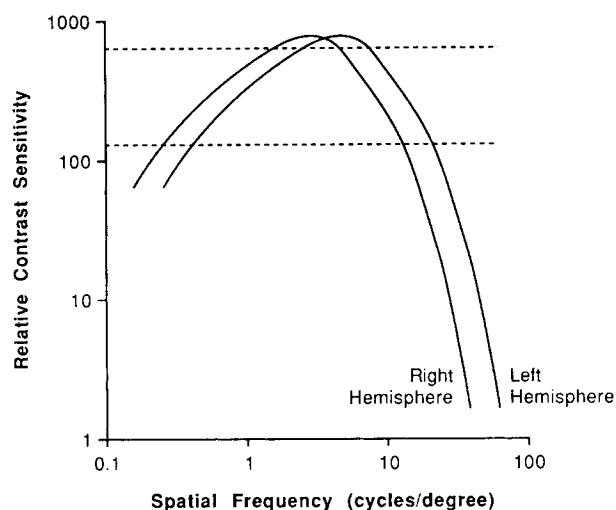


Figure 11. A hypothesized difference between identification modulation transfer functions of high-level visual areas in the left and right hemispheres. (Scales are logarithmic.)

in the degree to which they use inputs from units with different-sized receptive fields. However, at high contrast (when less sensitivity is required, as indicated by the bottom dotted line), both hemispheres use input from units with a wide range of receptive field sizes; hence the two distributions have a large amount of overlap, and the performance of the two hemispheres would not be well separated.

This idea might explain Sergent's (1991) finding that hemispheric differences between categorical and coordinate processing occur only when the stimuli are relatively degraded (i.e., are presented with low contrast). Accordingly, we decided to explore this hypothesis with another set of simulated networks. In these models, we again varied the sizes of the receptive fields of the inputs but now compared networks in which a narrow range of receptive field sizes was used (corresponding to low contrast) with those in which a wider range of receptive field sizes was used (corresponding to high contrast).

In addition, we compared this hypothesis with a simpler one: Perhaps increased contrast does not recruit additional neurons that increase the range of receptive field sizes; rather, it eliminates hemispheric differences simply because more of the same type of low-level neurons are stimulated over threshold. Even if neurons have small receptive fields, enough overlapping outputs would enable coarse coding to be used effectively even in the left hemisphere.

Thus we compared two sets of networks: ones in which greater contrast was assumed to produce outputs from a wide range of sizes of receptive fields ("mixed" receptive-field-size networks) and ones in which greater contrast was assumed simply to produce more input from the lower visual areas ("homogeneous" receptive-field-size networks). If the effect of receptive field size found in Study 3 were eliminated in mixed networks only, the first hypothesis would be supported; if this effect were eliminated in homogeneous networks, the second hypothesis would be supported.

## Method

### Materials

These networks differed from those used in Study 3 in only two ways: First, Layer 2 was composed of 28 units for all networks, instead of the 14 used before. Second, half of the networks had Layer 2 units with receptive fields of mixed sizes within a single network, whereas the other half had Layer 2 units with receptive fields of one size.

Half of the networks with homogeneous Layer 2 receptive field sizes were created so that all the units in Layer 2 had small receptive fields ( $\sigma = 0.71$ ) and half so that all of these units had large receptive fields ( $\sigma = 1.42$ ). These networks were designated *homogeneous/small* and *homogeneous/large* networks, respectively. Half of the networks with mixed Layer 2 receptive field sizes were created so that 14 of the Layer 2 units had small receptive fields ( $\sigma = 0.71$ ), 7 had even smaller receptive fields ( $\sigma = 0.36$ ), and the other 7 had large receptive fields ( $\sigma = 1.42$ ). These networks were designated *mixed/small* networks. The other half of the networks with mixed Layer 2 receptive field sizes were created so that 14 of the Layer 2 units had large receptive fields ( $\sigma = 1.42$ ), 7 had even larger receptive fields ( $\sigma = 2.84$ ), and 7 had small receptive fields ( $\sigma = 0.71$ ). These networks were designated *mixed/large* networks.

The receptive field peak locations for the 28 Layer 2 units in these networks were tessellated across the retinal array, as illustrated in the right panel of Figure 9. For mixed/small and mixed/large networks, all three receptive field sizes were evenly distributed across the retinal array; Figure 9 also indicates the peak locations for the three receptive field sizes.

### Procedure

The procedure was the same as in Study 3, except that eight conditions were tested; we produced these conditions by orthogonally combining categorical and coordinate mappings with homogeneous/small, homogeneous/large, mixed/small, and mixed/large networks. Ten networks were tested in each condition.

## Results

The results from this study are shown in Figure 12. To investigate whether mixed or homogeneous networks more accurately paralleled the behavioral findings, we analyzed the results from the two types of networks in separate ANOVAs. The mixed networks performed the mappings differently when they had different-sized receptive fields, as indicated by an interaction between receptive field size and task,  $F(1, 36) = 5.37, p < .05$ . Specifically, mixed networks performed the categorical task better when they had a majority of small receptive fields ( $M = 0.0116$  error) than when they had a majority of large fields ( $M = 0.0179$  error),  $F(1, 36) = 8.98, p < .01$ , but performed the coordinate tasks equally well with both receptive field sizes ( $M_s = 0.0051$  and  $0.0045$  for small and large receptive field networks),  $F < 1$ . In addition, the mixed networks generally performed the coordinate task better than the categorical one ( $M_s = 0.0048$  and  $0.0148$ ),  $F(1, 36) = 45.1, p < .001$ . These findings thus do not mirror Sergent's (1991) behavioral results, which did not reveal a difference in the categorical task with high contrast.

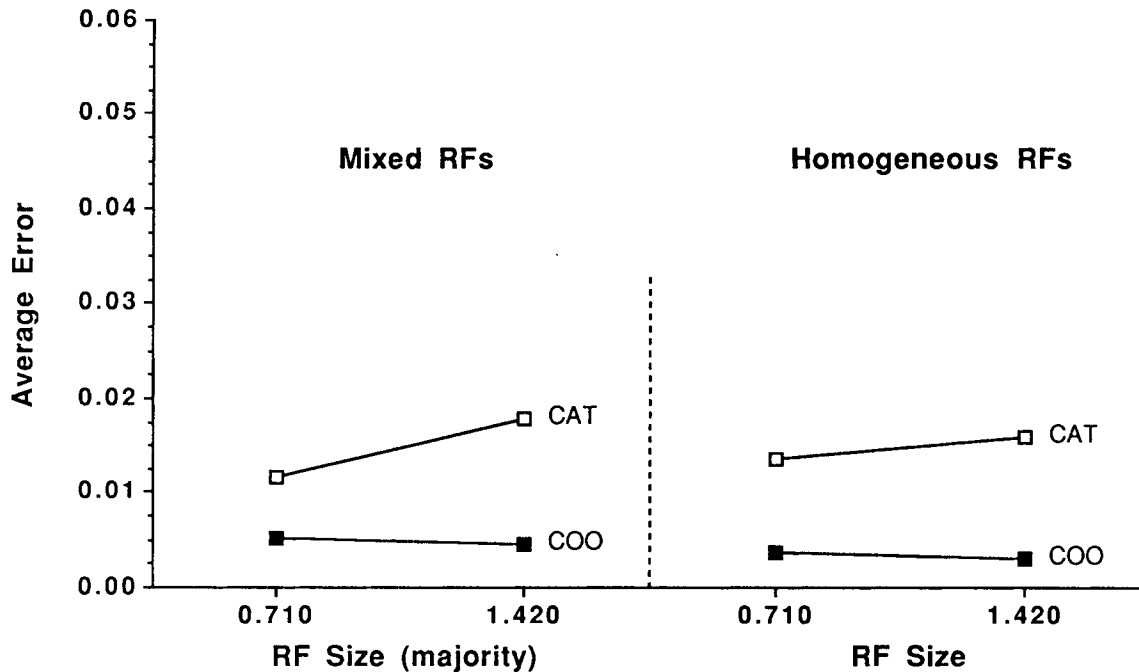


Figure 12. Results of Study 4, showing that adding inputs eliminated the advantage of large receptive fields (RFs) for accomplishing the coordinate mapping. (Moreover, if additional inputs with same-sized receptive fields were added, small receptive fields no longer allowed the network to establish the categorical mapping more easily than did large receptive fields. CAT = categorical; COO = coordinate.)

In contrast, homogeneous networks did produce the same pattern of results found in Sergent's experiments with high contrast. In these networks, there was no interaction between receptive field size and task,  $F(1, 36) = 1.14, p > .25$ . In addition, the categorical task was performed equally well with small and large receptive fields ( $M_s = 0.0135$  and  $0.0159$ ),  $F(1, 36) = 1.55, p > .20$ , as was the coordinate task ( $M_s = 0.0036$  and  $0.0030$ ),  $F < 1$ . Finally, these networks performed the coordinate task better than the categorical task ( $M_s = 0.0033$  and  $0.0147$ ),  $F(1, 36) = 68.5, p < .001$ .

### Discussion

We were best able to model Sergent's (1991) findings if we simply assumed that more contrast leads to more output from low-level neurons. Adding inputs to the network effectively eliminated the advantage of large receptive fields for encoding precise location; the networks apparently used coarse coding effectively, even if the receptive fields were relatively small. This observation makes sense when one considers that coarse coding is an effective strategy only when both sufficient and systematic (but not excessive) overlap exists in the distribution of response profiles. (It has previously been demonstrated that overlap in receptive fields must be staggered in order to encode precise locations; O'Reilly et al., 1990; see also Ballard, 1986.)

Presenting stimuli with higher levels of contrast probably caused more units with overlapping receptive fields to enter the distribution. Therefore, if the added receptive fields represent areas distributed fairly evenly across the input space, a

higher and more effective degree of overlap is obtained even among relatively small receptive fields.

### General Discussion

Our results provide additional support for the conceptual distinction between categorical and coordinate representations of spatial relations. We not only found that at least in some circumstances different computations encode the two kinds of spatial relations, and hence they are not logically bound together, but also described a simple mechanism that can explain why coordinate representations are computed more effectively in the right cerebral hemisphere and categorical representations are computed more effectively in the left cerebral hemisphere, as well as why these hemispheric differences are not evident when stimuli are presented with high contrast.

Specifically, our networks evaluated whether a dot was above or below a line or whether a dot was within four elements of a line; thus they mimicked the categorical and coordinate tasks developed by Hellige and Michimata (1989). We found that the two judgments were performed more effectively in a neural network model when the hidden units were segregated into two separate subsystems, whereby one provided input to the categorical judgment output units and the other provided input to the coordinate judgment output units. When a single, unsplit network was used to perform both mappings, there apparently was interference between the different types of internal representations needed to accom-

plish the two mappings. In contrast, the reverse pattern was found when a network made two metric judgments; there, common representations could be used for both mappings, and dividing the networks made it more difficult to establish the mappings.

In addition, we found that both sorts of mappings were more difficult when finer discriminations were required. Sergent (1991) took the corresponding finding with human subjects as evidence against the distinction; our results suggest that her findings do not contradict the contention that the two mappings are qualitatively distinct. Our most general point is that a subset of shared properties does not imply that entities are exactly the same; different species of mammals, for example, share many properties (such as warm blood, hair, and so on), but other characteristics individuate them.

Furthermore, we found that the coordinate mapping was easier if the input was filtered through larger overlapping receptive fields, whereas the categorical mapping was easier if the input was filtered through smaller receptive fields. The idea that the right hemisphere monitors large receptive fields is consistent with its possible role in encoding locations to initiate an action.

Finally, the network models suggested an account of Sergent's (1991) finding that contrast alters the observed pattern of lateralization. When a relatively large number of units contributed input to a network, as may occur with high-contrast stimuli, the sizes of their receptive fields did not affect performance.

These findings suggest that the distinction between categorical and coordinate spatial relations encoding is both conceptually sound and computationally plausible. In some ways, our conceptualization of this distinction is now closer to Sergent's (1991) interpretation of her results than to the original formulation of Kosslyn (1987), and we are grateful to Sergent for leading us to characterize the distinction more rigorously. However, Sergent accounted for the right-hemisphere advantage in the distance task by assuming that it is generally superior at making efficient use of lower quality information and apparently assumed that there is no hemispheric difference in the above/below task because precise location need not be encoded or because high-quality information is not needed to compute this relation. This interpretation fails to account for the fact that people judge categorical relations better when the stimulus is presented initially to the left hemisphere than when it is presented initially to the right hemisphere (as we reviewed in the introduction). Furthermore, patients with Gerstmann's syndrome cannot judge left from right, and these patients typically have damage to the left parietal lobe (e.g., Levine, Mani, & Calvanio, 1988).

Sergent (1991) argued against the idea that each hemisphere computes only one type of representation (either categorical or coordinate). Although we apparently have not been clear about our position previously, we have always agreed with this view; we assume that both hemispheres can compute both types of spatial relations, but not equally effectively. Indeed, one of the reasons why we find neural network simulation models attractive is that they are consistent with five principles that Kosslyn and Koenig (in press) inferred

about actual neural computation in the brain, one of which posits that the brain has only "weak modularity." We exploited these five principles in the following way.

First, *division of labor*: It is more efficient for separate networks to perform different types of mappings. Because the same patterns of weights on connections are used to accomplish different input/output mappings, different types of mappings interfere with each other. This principle does not, however, imply that the brain always does things in the most efficient way. Nevertheless, we assume that basic perceptual/motor processing, such as we studied, has become relatively efficient through the course of evolution; hence it is plausible that these sorts of processing often are performed efficiently. Second, *weak modularity*: The subsystems of the brain are not like electronic parts, with completely discrete functions. Rather, there is some overlap in the operation of the different components. For example, neurons in many visual areas are sensitive to more than one stimulus dimension (e.g., Van Essen, 1985). The fact that both categorical and coordinate mappings are sensitive to the distances between objects being judged may suggest that the mappings include a common underlying component (but this does not imply that the mappings are the same any more than the fact that dogs and bears have fur means that they are the same). Third, *constraint satisfaction*: Precise information is computed by satisfying a variety of weak constraints simultaneously. Coarse coding is one mechanism that carries out such computations. Furthermore, we assume that categorical spatial relations are encoded in conjunction with information about the identity of parts during object identification, and it is the combination of the two sorts of information that places strong constraints on what object is being viewed (see Kosslyn et al., 1990). Fourth, *concurrent processing*: We do not expect the system to "decide" which kind of representation to compute; rather, we expect it simply to compute whatever it can on the basis of the input. If so, separate subsystems operating in parallel may often be used to compute independent mappings; our experiments suggest that this would be a feasible arrangement for computing the two kinds of spatial relations representations. Finally, *opportunism*: Mechanisms that originally evolved for one purpose may be recruited later for another (cf. Gould & Lewontin, 1978). We have suggested that hemispheric differences in spatial relations encoding may have arisen because of mechanisms used in preattentive processing. Once the receptive field differences were present for this purpose, they could be exploited by processes that perform other tasks.

It is important to note that we have used network models in an unusual way in these studies; we are interested in them solely in terms of how well they perform specific input/output mappings. We were careful to design the tasks given to the networks to parallel the important features of the tasks given to humans, and we used the models to study the difficulty of establishing the necessary input/output mappings under different conditions. We made minimal assumptions about the psychological reality of the details of the models, and we did not use weights as parameter estimates or the like (see Mas-saro, 1988, for problems with some more common uses of such models). Rather, we argue that differences in the ease of

establishing input/output mappings in these networks are of interest because these mappings capture—at a rather abstract level—certain aspects of corresponding mappings in the brain. The results provide evidence that the two kinds of mappings are distinct within these kinds of networks and enable us to formulate a hypothesis that is sufficient to account for how the hemispheric differences arise.

We have not shown that the same patterns of results would occur with all possible parameter values in the models. For our purposes, however, this is not a very interesting question; it is likely that some parameter settings would allow networks to perform all of the tasks very easily (displaying a floor effect, of the sort seen in Study 4), whereas others (e.g., including only a couple of hidden units) would hamstring them in all of the tasks. The fact that we found selective differences in performance under any conditions is evidence that the mappings are distinct. A critical part of this logic is that we were able to demonstrate not only simple main effects, in which one network configuration was better than another, but also interactions: Depending on the precise task, one network performed better or worse than another. Hence the performance of the network could not be ascribed to a combination of parameter settings that simply made it effective or ineffective in general.

Intuitively, it is easy to question whether the distinction between categorical and coordinate spatial relations is coherent, let alone plausible. In fact, Kosslyn (1987) and Kosslyn, Koenig, et al. (1989) characterized the distinction slightly differently than we did, partly because they had not yet had to grapple with the problems of implementing simulation models and interpreting the results. It is clear that models such as these are a useful supplement to intuition, providing further bases for formulating and evaluating hypotheses about human information processing.

## References

- Andersen, R. A. (1987). Inferior parietal lobule function in spatial perception and visuomotor integration. In F. Plum (Vol. Ed.) & V. B. Mountcastle (Sec. Ed.), *Handbook of physiology, Section 1: The nervous system, Volume 5. Higher functions of the brain* (pp. 483–518). Bethesda, MD: American Physiological Society.
- Andersen, R. A., Asanuma, C., Essick, G., & Siegel, R. M. (1990). Cortico-cortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. *Journal of Comparative Neurology*, 296, 65–113.
- Ballard, D. H. (1986). Cortical connections and parallel processing: Structure and function. *Behavioral and Brain Sciences*, 9, 67–120.
- Biederman, I. (1987). Recognition-by-components: A theory of human image understanding. *Psychological Review*, 94, 115–147.
- Delis, D. C., Robertson, L. C., & Efron, R. (1986). Hemispheric specialization of memory for visual hierarchical stimuli. *Neuropsychologia*, 24, 205–214.
- De Schonen, S., & Mathivet, E. (1989). First come, first served: A scenario about the development of hemispheric specialization in face recognition during infancy. *European Bulletin of Cognitive Psychology*, 1, 3–44.
- De Valois, R. L., & De Valois, K. K. (1988). *Spatial vision*. New York: Oxford University Press.
- Fisk, J. D., & Goodale, M. A. (1988). The effects of unilateral brain damage on visually guided reaching: Hemispheric differences in the nature of the deficit. *Experimental Brain Research*, 72, 425–435.
- Goodale, M. A. (1988). Hemispheric differences in motor control. *Behavioral Brain Research*, 30, 203–214.
- Gould, S. J., & Lewontin, R. C. (1978). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society of London B*, 205, 581–598.
- Hellige, J. B., & Michimata, C. (1989). Categorization versus distance: Hemispheric differences for processing spatial information. *Memory & Cognition*, 17, 770–776.
- Hinton, G. E., McClelland, J. L., & Rumelhart, D. E. (1986). Distributed representations. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations* (pp. 77–109). Cambridge, MA: MIT Press.
- Hyvarinen, J. (1982). Posterior parietal lobe of the primate brain. *Physiological Review*, 62, 1060–1129.
- Kaufman, L. (1974). *Sight and mind*. New York: Oxford University Press.
- Kitterle, F. L., Christman, S., & Hellige, J. B. (1990). Hemispheric differences are found in the identification, but not the detection, of low versus high spatial frequencies. *Perception & Psychophysics*, 48, 297–306.
- Koenig, O., Reiss, L. P., & Kosslyn, S. M. (1990). The development of spatial relations representations: Evidence from studies of cerebral lateralization. *Journal of Experimental Child Psychology*, 50, 119–130.
- Kosslyn, S. M. (1987). Seeing and imagining in the cerebral hemispheres: A computational approach. *Psychological Review*, 94, 148–175.
- Kosslyn, S. M., Flynn, R. A., Amsterdam, J. B., & Wang, G. (1990). Components of high-level vision: A cognitive neuroscience analysis and accounts of neurological syndromes. *Cognition*, 34, 203–277.
- Kosslyn, S. M., & Koenig, O. (in press). *Wet mind: The new cognitive neuroscience*. New York: Free Press.
- Kosslyn, S. M., Koenig, O., Barrett, A., Cave, C. B., Tang, J., & Gabrieli, J. D. E. (1989). Evidence for two types of spatial representations: Hemispheric specialization for categorical and coordinate relations. *Journal of Experimental Psychology: Human Perception and Performance*, 15, 723–735.
- Kosslyn, S. M., Sokolov, M. A., & Chen, J. C. (1989). The lateralization of BRIAN: A computational theory and model of visual hemispheric specialization. In D. Klahr & K. Kotovsky (Eds.), *Complex information processing: The impact of Herbert H. Simon* (pp. 3–29). Hillsdale, NJ: Erlbaum.
- Lehky, S. R., & Sejnowski, T. J. (1988). Network model of shape-from-shading: Neural function arises from both receptive and projective fields. *Nature*, 333, 452–454.
- Levine, D. N., Mani, R. B., & Calvanio, R. (1988). Pure agraphia and Gerstmann's syndrome as a visuospatial-language dissociation: An experimental case study. *Brain and Language*, 35, 172–196.
- Livingstone, M., & Hubel, D. (1988). Segregation of form, color, movement, and depth: Anatomy, physiology, and perception. *Science*, 240, 740–749.
- Marr, D. (1982). *Vision*. New York: W. H. Freeman.
- Massaro, D. W. (1988). Some criticisms of connectionist models of human performance. *Journal of Memory and Language*, 27, 213–234.
- Maunsell, J. H. R., & Newsome, W. T. (1987). Visual processing in monkey extrastriate cortex. *Annual Review of Neuroscience*, 10, 363–401.
- McClelland, J. L., & Rumelhart, D. E. (Eds.) (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*.

- Volume 2: Psychological and biological models.* Cambridge, MA: MIT Press.
- McLeod, P., McLaughlin, C., & Nimmo-Smith, I. (1985). Information encapsulation and automaticity: Evidence from the visual control of finely timed actions. In M. I. Posner & O. S. M. Marin (Eds.), *Attention and performance XI* (pp. 391–406). Hillsdale, NJ: Erlbaum.
- Minsky, M. L., & Papert, S. A. (1969). *Perceptrons*. Cambridge, MA: MIT Press.
- O'Reilly, R. C., Kosslyn, S. M., Marsolek, C. J., & Chabris, C. F. (1990). Receptive field characteristics that allow parietal lobe neurons to encode spatial properties of visual input: A computational analysis. *Journal of Cognitive Neuroscience*, 2, 141–155.
- Osherson, D. N., Kosslyn, S. M., & Hollerbach, J. M. (Eds.) (1989). *An invitation to cognitive science, Volume II: Visual cognition and action*. Cambridge, MA: MIT Press.
- Rosenthal, R. (1984). *Meta-analytic procedures for social research*. Beverly Hills, CA: Sage.
- Rueckl, J. G., Cave, K. R., & Kosslyn, S. M. (1989). Why are “what” and “where” processed by separate cortical visual systems? A computational investigation. *Journal of Cognitive Neuroscience*, 1, 171–186.
- Rueckl, J. G., & Kosslyn, S. M. (1992). What good is connectionist modelling? In A. Healy, S. M. Kosslyn, & R. M. Shiffrin, (Eds.), *Essays in honor of W. K. Estes* (pp. 249–266). Hillsdale, NJ: Erlbaum.
- Rumelhart, D. E., Hinton, G., & Williams, R. (1986). Learning internal representations by error propagation. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations* (pp. 318–362). Cambridge, MA: MIT Press.
- Rumelhart, D. E., & McClelland, J. L. (Eds.) (1986). *Parallel distributed processing: Explorations in the microstructure of cognition. Volume 1: Foundations*. Cambridge, MA: MIT Press.
- Seidenberg, M. S., & McClelland, J. L. (1989). A distributed, developmental model of word recognition and naming. *Psychological Review*, 96, 523–568.
- Sergent, J. (1982). The cerebral balance of power: Confrontation or cooperation? *Journal of Experimental Psychology: Human Perception and Performance*, 8, 253–272.
- Sergent, J. (1991). Judgments of relative position and distance on representations of spatial relations. *Journal of Experimental Psychology: Human Perception and Performance*, 17, 762–780.
- Smith, E. E., & Medin, D. (1983). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- Stornetta, W. S., & Huberman, B. A. (1987). *An improved three-layer back propagation algorithm*. Paper presented at the Proceedings of the IEEE First International Conference on Neural Networks, San Diego, CA.
- Ungerleider, L. G., & Mishkin, M. (1982). Two cortical visual systems. In D. J. Ingle, M. A. Goodale, & R. J. W. Mansfield (Eds.), *The analysis of visual behavior* (pp. 549–586). Cambridge, MA: MIT Press.
- Van Essen, D. (1985). Functional organization of primate visual cortex. In A. Peters & E. G. Jones (Eds.), *Cerebral cortex* (Vol. 3, pp. 259–329). New York: Plenum Press.
- Van Kleeck, M. H. (1989). Hemispheric differences in global versus local processing of hierarchical visual stimuli by normal subjects: New data and a meta-analysis of previous studies. *Neuropsychologia*, 27, 1165–1178.
- Zipser, D., & Andersen, R. A. (1988). A back-propagation programmed network that simulates response properties of a subset of posterior parietal neurons. *Nature*, 331, 679–684.

Received November 6, 1990

Revision received June 7, 1991

Accepted June 10, 1991 ■