DIGITAL ACCESS TO SCHOLARSHIP AT HARVARD

Receptive Field Characteristics That Allow Parietal Lobe Neurons to Encode Spatial Properties of Visual Input: A Computational Analysis

The Harvard community has made this article openly available.
Please share how this access benefits you. Your story matters.

Citation	O'Reilly, Randall C., Stephen M. Kosslyn, Chad J. Marsolek, and Christopher F. Chabris. 1990. Receptive field characteristics that allow parietal lobe neurons to encode spatial properties of visual input: A computational analysis. Journal of Cognitive Neuroscience 2(2): 141-155.
Published Version	doi:10.1162/jocn.1990.2.2.141
Accessed	February 18, 2015 7:18:00 AM EST
Citable Link	http://nrs.harvard.edu/urn-3:HUL.InstRepos:3597235
Terms of Use	This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA

(Article begins on next page)

Receptive Field Characteristics That Allow Parietal Lobe Neurons to Encode Spatial Properties of Visual Input: A Computational Analysis

Randall C. O'Reilly, Stephen M. Kosslyn, Chad J. Marsolek, and Christopher F. Chabris
Department of Psychology
Harvard University

Abstract

■ A subset of visually sensitive neurons in the parietal lobe apparently can encode the locations of stimuli, whereas visually sensitive neurons in the inferotemporal cortex (area IT) cannot. This finding is puzzling because both sorts of neurons have large receptive fields, and yet location can be encoded in one case, but not in the other. The experiments reported here investigated the hypothesis that a crucial difference between the IT and parietal neurons is the spatial distribution of their response profiles. In particular, IT neurons typically respond maximally when stimuli are presented at the fovea, whereas parietal neurons do not. We found that a parallel-distributed-processing network could map a point in an array to a coordinate representation more easily when a greater proportion of its input units had response peaks off the center of the input array. Furthermore, this result did not depend on potentially implausible assumptions about the regularity of the overlap in receptive fields or the homogeneity of the response profiles of different units. Finally, the internal representations formed within the network had receptive fields resembling those found in area 7a of the parietal lobe. ■

INTRODUCTION

For many years it was the fashion in neuroscience to use common sense to analyze the functional properties of neurons. For example, if a neuron responded particularly well to faces, some inferred that it was a "face detector." But this reasoning is incomplete and misleading, in part because it ignores the role the neuron plays in the context of other neurons (e.g., see Van Essen 1985). In this article we explore an example of another way of analyzing the functional properties of neurons: Given that neurons produce specific output on receiving specific input from other neurons, they can be thought of as performing *computations*, which are systematic mappings of input to output that transform the input or operate on it in some way (cf. Marr 1982). A complete analysis of the computations performed by a set of neurons is specific enough to allow one to build a system with comparable functions. Thus, attempting to characterize neural function as computation leads one to ask much more detailed questions about neural information processing than were asked previously, and can produce insights into nonobvious properties of the specific mappings performed by neural systems.

One well-established fact about neural information processing is that separate pathways analyze different properties of visual input. In particular, spatial properties (such as location and orientation) are processed by a dorsal pathway that leads up from the occipital lobe into the parietal lobe, whereas object properties (such as shape and color) are processed by a ventral pathway that leads down into the inferior temporal lobe (e.g., Maunsell and Newsome 1987; Mishkin and Ungerleider 1982; Mishkin et al. 1983; Ungerleider and Mishkin 1982).

Rueckl, Cave, and Kosslyn (1989) explored the idea that dividing visual processing in this way is computationally more efficient than combining it in a single system that encodes both shape and location information. To test this hypothesis, Rueckl et al. trained a three layer parallel-distributed-processing network with the backpropagation algorithm (Rumelhart, Hinton, and Williams 1986). The inputs were simple patterns, which could appear in different places in an array, and the outputs were classifications of the shape and location of the stimuli. In one set of experiments, the input units were fully connected to each unit in a hidden layer, and each of these hidden units was fully connected to the output units. In another set of experiments, some of the hidden

units were connected only to the shape output units, and the rest were connected only to the location output units. The question was whether dividing the network into two processing streams would make it easier to establish the proper input/output mappings. The results were straightforward: If the hidden units were split in the right proportion, a divided system established the necessary mapping much more efficiently than did a single network that encoded both shape and location.

Gross and Mishkin (1977) speculated that the division of labor into two visual pathways is the basis of certain functional properties of the visual system. In particular, they considered how a divided visual system might achieve "stimulus equivalence across retinal translation," the ability to identify an object regardless of where it appears in the visual field and what part of the retina its image strikes. They speculated that this ability is conferred in part by the very large receptive fields of cells in area IT (in the inferior temporal lobe); the median size of these receptive fields is about 25°, but they can reach sizes of over 100° (Desimone and Gross 1979; Gross, Rocha-Miranda, and Bender 1972). For comparison, neurons in area V4 with excitatory receptive fields that include the center of gaze have much smaller receptive fields, frequently less than 1° in diameter (e.g., see Desimone, Schein, Moran, and Ungerleider 1985). IT cells will therefore respond to stimuli when they are located in a wide variety of positions, presumably allowing them to be recognized no matter where they appear.

But at the same time one is recognizing an object, one needs to know exactly where it is to reach to it, track it as it moves through space, and note the spatial relations between it and other objects. These abilities are conferred by the parietal system. In short, the idea is that the visual system "divides and conquers," with the temporal lobe system ignoring location over a wide range, which is useful for recognition, and the parietal lobe system encoding specific location, which is useful for other purposes.

Unfortunately, the receptive field data from *parietal lobe* neurons do not, at first glance, fit this model very well. Gross and Mishin speculated that stimulus equivalence was achieved in part because of the large receptive fields of neurons in IT. But neurons in the relevant areas of the parietal lobe also have very large receptive fields (Andersen, Essick, and Siegel 1985; Motter and Mountcastle 1981), yet they appear to encode location.

In this article we consider what distinctive properties of parietal neurons might allow them to register location despite their lack of fine-grained receptive fields. We will begin this effort by comparing and contrasting the receptive field properties of neurons in the temporal and parietal systems, and will consider which properties of parietal lobe receptive fields might allow them to encode location. We then will test the computational feasibility of our hypotheses with network models.

Distinctive Properties of Parietal Lobe Receptive Fields

We are concerned here with delineating properties of neurons in the "higher level" processing areas of the parietal lobes that appear to be important for spatial location encoding. In particular, we focused on area 7a within the inferior parietal lobule because of the evidence linking this area specifically with visual and visualmotor processing (Andersen, Asanuma, Essick, and Siegel in press; Andersen et al. 1985). Many properties of neurons in area 7a are not distinctive, being similar to those of neurons in area IT. Neurons in both areas respond to stimuli presented in a large range of positions in both visual hemifields, have roughly comparable receptive field sizes, tend to have peak responses to contralateral input, and have receptive fields that are not topographically organized (Andersen et al. in press, Andersen et al. 1985; Desimone et al. 1985; Desimone and Gross 1979; Gross et al. 1972).

Differences in the receptive field properties of the two types of neurons are somewhat subtle, but two reasonably clear distinctions may be drawn. First, the peak responses of neurons in areas 7a and IT tend to be distributed differently throughout their receptive fields. Virtually all IT neurons respond most vigorously when the stimulus falls on the fovea, whereas neurons in the parietal lobe rarely respond most vigorously to foveal stimuli (Andersen et al. in press; Gross et al. 1972). Indeed, Motter and Mountcastle (1981) reported a pattern of responses termed "foveal sparing," in which the receptive fields of some area 7a cells did not include the fovea at all (i.e., these neurons were excited only by stimuli presented in peripheral regions). In one study, about 40% of neurons examined in area 7a exhibited this characteristic (Motter, Steinmetz, Duffy, and Mountcastle 1987). Andersen et al. (in press) also discovered what they termed "holes" (areas in which stimuli did not elicit responses) in the receptive fields of 7a neurons, but these areas were found to be located in the periphery as well as the foveal region.

The second distinction between the two classes of neurons concerns the number of locations in the receptive field that can produce maximal responses: There typically is only a single location at which a stimulus produces a maximal response from a neuron in IT (Gross et al. 1972; Desimone and Gross 1979), but there sometimes are several locations at which a stimulus produces a maximal response from a neuron in area 7a. In other words, some neurons in area 7a have multiple-peak receptive fields (Andersen et al. in press; Andersen and Zipser 1988; Motter and Mountcastle 1981; Zipser and Andersen 1988).

In addition, the responses of some neurons in area 7a have also been found to be sensitive to eye position and oculomotor activity (Andersen et al. in press; Andersen

et al. 1985). However, over 20% of area 7a neurons respond only to visual stimulation (Zipser and Andersen 1988), and it has been argued that the inferior parietal system as a whole uses eye position to produce head-centered receptive fields, which compensate for the effects of eye position per se (Andersen 1986; Andersen et al. 1985). For these reasons, we have not emphasized effects of oculomotor behavior in our comparisons between IT and 7a neurons.

Our investigations complement those of Zipser and Andersen (1988) and Andersen and Zipser (1988), who used a neural network model to examine the role of eye position and retinal location in encoding location in head-centered coordinates. Zipser and Andersen compared the types of receptive fields developed in their network with the receptive fields of actual neurons in area 7a. They focused on the way information about eye position and retinal location is combined, whereas we are focusing on differences between area 7a and IT neurons that allow location to be encoded at all. Our model does not examine the role of eye position in computing location. However, we have adopted Zipser and Andersen's general approach to modeling computation in the parietal lobe, and will relate our findings to theirs at the end of the article.

A Coarse Coding Model

After considering the similarities and differences between neurons in areas 7a and IT, we suspected that IT neurons are impaired in their ability to encode location insofar as they respond maximally only to stimuli in the fovea. Although the common peak response location would serve to indicate whether a stimulus was on the fovea, it would not help to register other locations. Thus, we hypothesized that the variety of receptive field properties of neurons in the parietal lobes enables them to encode location. These intuitions were grounded in research on parallel distributed processing computer simulation models. Hinton, McClelland, and Rumelhart (1986) showed that units that respond coarsely to input can provide very precise information if their distributions overlap systematically. Similarly, Ballard (1986) described how, given a network with a limited number of processing units but with an abundance of connections, systematic overlapping provides high signal resolution with relatively few units. A good example of the success of such *coarse coding* processes can be found in human color vision: Although we have only three types of cones, which have relatively wide tuning functions, we see millions of colors by using the relative mixes of the outputs of the three types of cones. The wide range of response types found for neurons in area 7a might be analogous to the different tuning functions of cones, providing a sufficient range of outputs from overlapping receptive

fields to allow the system to converge on the location of the input (cf. Andersen 1986).

At first glance, however, one might suspect that the presence of multiple peaks in area 7a should *not* enhance its spatial encoding abilities in a coarse coding system. Receptive fields with multiple peaks exhibit similar responses for different stimulus locations. This ambiguity hinders the function of a coarse-coding algorithm, which employs the overlap from *separate and distinct* representations of the input (Hinton et al. 1986). However, this intuition might be faulty, and so must be tested empirically. In addition, we wondered whether the neurons with multiple peaks might play a different role in coarse coding, not being part of the input representation. Perhaps these neurons do not contribute to the initial phases of the computation, but rather are involved in the later phases of combining the inputs from different single-peak neurons. If so, then we expected such response properties to develop in the internal representations of a network that was trained to perform the proper input/output mapping. And in fact, Zipser and Andersen (1988) and Andersen and Zipser (1988) provide computational results that are consistent with these ideas, as will be discussed in more detail later.

Thus, we explored the ease of establishing different input/output mappings using backpropagation neural network simulations (Rumelhart et al. 1986). We tested the ease of mapping the location of a stimulus on a simulated "retinal" array to an explicit representation of that location when this "retinal" input has been processed by units with different receptive field properties. We reasoned that because individual neurons (in either IT or 7a) cannot encode location very well, because of their large receptive fields, the outputs from these neurons must be used together to converge on location. The receptive field properties, then, can be viewed as modulating the input to a larger system that computes precise location. Following Rueckl et al. (1989), we used the amount of error after a fixed number of training trials as a measure of the ease of making a given input/output mapping. We also analyzed the networks after successful training to discover how the hidden units used the information in the input to accomplish the mapping.

EXPERIMENT 1

In this experiment we examined how well different networks could encode the spatial location of a point of light on a simulated retina. We first examined the effect of manipulating the distribution of receptive field peaks in the input to the network. The location of peak response was systematically varied within or outside the retina's foveal region (defined as the center of the input array). We also considered the importance of tessellated versus random distributions of peaks outside the fovea. If the network shows substantial degradation when the

distributions are random, we would be suspicious of its feasibility in a biological system. Finally, we also considered the consequences of individually randomized versus fixed receptive field profiles. The "profile," in this sense, is simply the distribution of a neuron's activation over the retinal array. Again, if good performance depends on all of the receptive fields having the same drop-off function from the peak, we would not be confident that we could generalize our results to natural information processing.

In Part 1 of Experiment 1 we systematically examined the ease of registering location when different proportions of receptive fields have peak response locations on the fovea. We measured the relative difficulty of mapping from a 7 × 7 simulated retinal array with a single illuminated point to two sets of 7 output units, which indexed the horizontal and vertical coordinates of the stimulus on the retina. The input, in layer 1, was filtered through 24 "neurons," in layer 2, and we varied the number of these 24 units that had peak responses outside the fovea. The hidden layer of the network, layer 3, contained 11 units. As is illustrated in Figure 1, each of the hidden units was connected to each unit in both

output sets in layer 4, which specified horizontal position and vertical position.

The number of peaks outside the fovea was varied in increments of 2, from 0 through 24. All of the receptive fields in this case had the same profile and the peaks outside the fovea were positioned at distinct random locations. Thus, we examined the outputs of 13 networks, which differed only in the proportion of peak response locations that were outside the fovea. We ran each network 10 times to 300 epochs; each epoch consisted of an input point being presented at all 49 possible locations.

In Part 2 of Experiment 1 we examined the effects of how the peak response locations are distributed and of individual variability within the receptive field distributions. We orthogonally compared the ease of mapping input to output in networks that had tessellated versus random peak distributions and individually randomized versus fixed receptive field profiles. Thus, there were four experimental conditions. In this part of the experiment, peak responses occurred in four different sets of locations (with 24, 16, 7, and 0 peaks outside the fovea), with half corresponding to regular tessellations of the
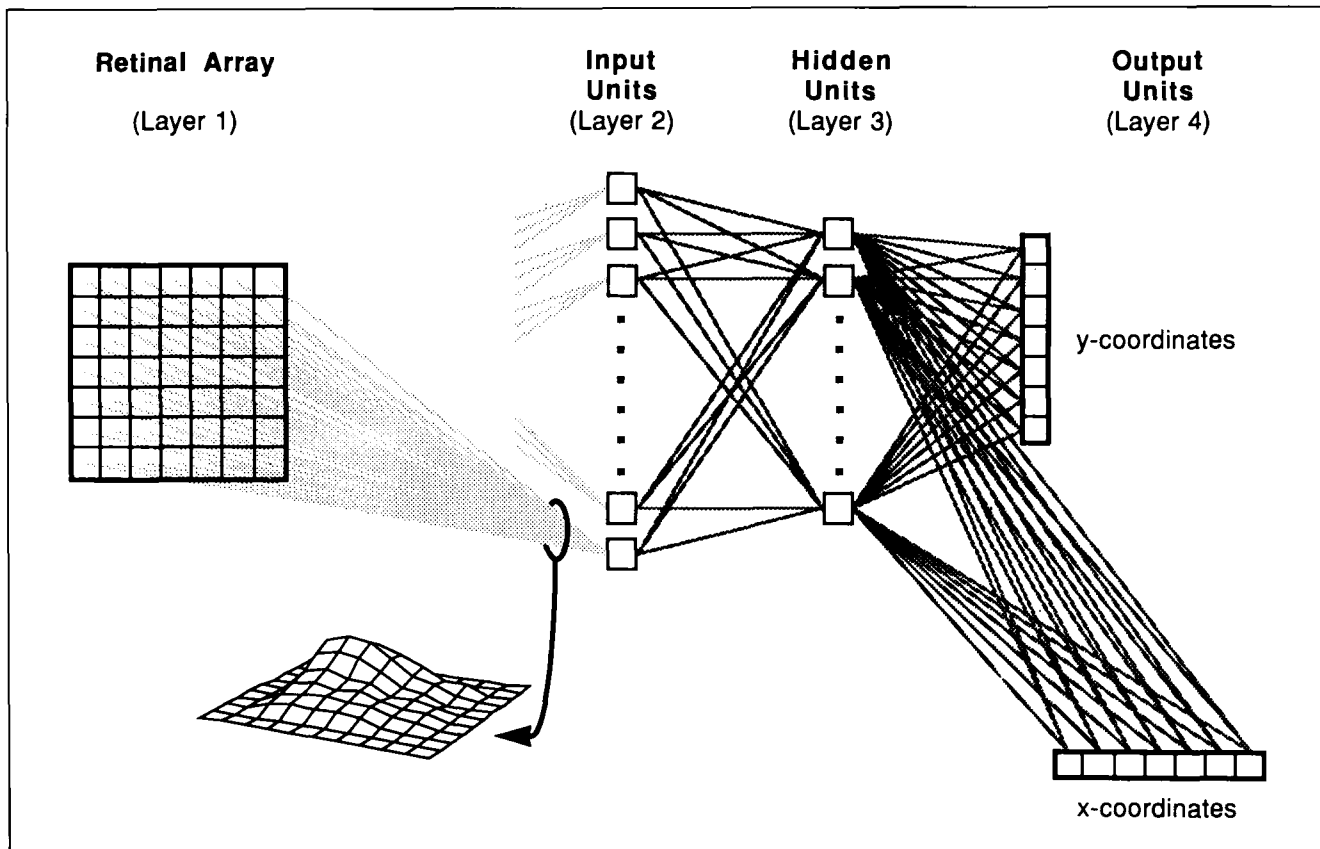

Figure 1. The feedforward (left to right) structure of the networks used in these experiments. The connection weights between points on the retinal array and the input units were fixed for each network configuration; these weights defined the properties of the receptive fields of the input units. Lightly shaded connections have fixed weights; darkly shaded connections have adjustable weights. The three-dimensional graph at the lower left illustrates how the weights to one input unit define its receptive field. Layers, 2, 3, and 4 form a standard three-layer backpropagation network. The output units, in layer 4, specify the coordinates along horizontal and vertical dimensions.

7 × 7 retinal matrix. These locations are illustrated in the top part of Figure 2. Each of the resulting 16 networks was run 10 times for 300 epochs.

Results and Discussion

Figure 3 illustrates the results from Part 1 of the experiment, in which we varied the proportion of receptive fields that had peak responses outside the fovea. Error clearly decreases in proportion to the number of peaks off the fovea, a pattern of results that supports our hypothesis that the distribution of peaks has a direct effect on the ease of mapping retinal position onto a coordinate representation. We analyzed these data using analyses of variance, with the 10 replications of each configuration being used as the random effect. The average amount of error after 300 training sets was the dependent measure. In fact, error did decrease when more peak response locations were off the "fovea," $F(12, 117) = 556.6, p < .001$. Contrast analyses revealed that there was a linear decrease in error with increasing numbers of peaks off the fovea, $F(1, 117) = 4,899, p < .001$, and that this

decrease decelerated with increasing numbers of peaks off the fovea, $F(1, 117) = 1,401, p < .001$; the linear and quadratic components together accounted for 94% of the variance. Thus, to the extent that neurons in area 7a have receptive fields whose peaks are distributed off the fovea, these results suggest one reason why the parietal lobe is better suited to encoding spatial location. The peak response locations of neurons in the inferior temporal lobe, on the other hand, are virtually always in the foveal region, which, according to our results, makes this area significantly less adept at encoding spatial information.

We next asked whether there was a critical point at which the computation became easier. Beginning with the comparison between the means for 0 and 2 off-fovea peaks, and working across the x axis in Figure 3, t tests were conducted comparing all pairs of treatment means. The critical p value was adjusted using the Bonferroni technique and the tests were continued until at least two consecutive comparisons yielded nonsignificant differences. We also examined the point at which the function no longer decreased monotonically. Both procedures indicated the presence of an "elbow" in the curve: After
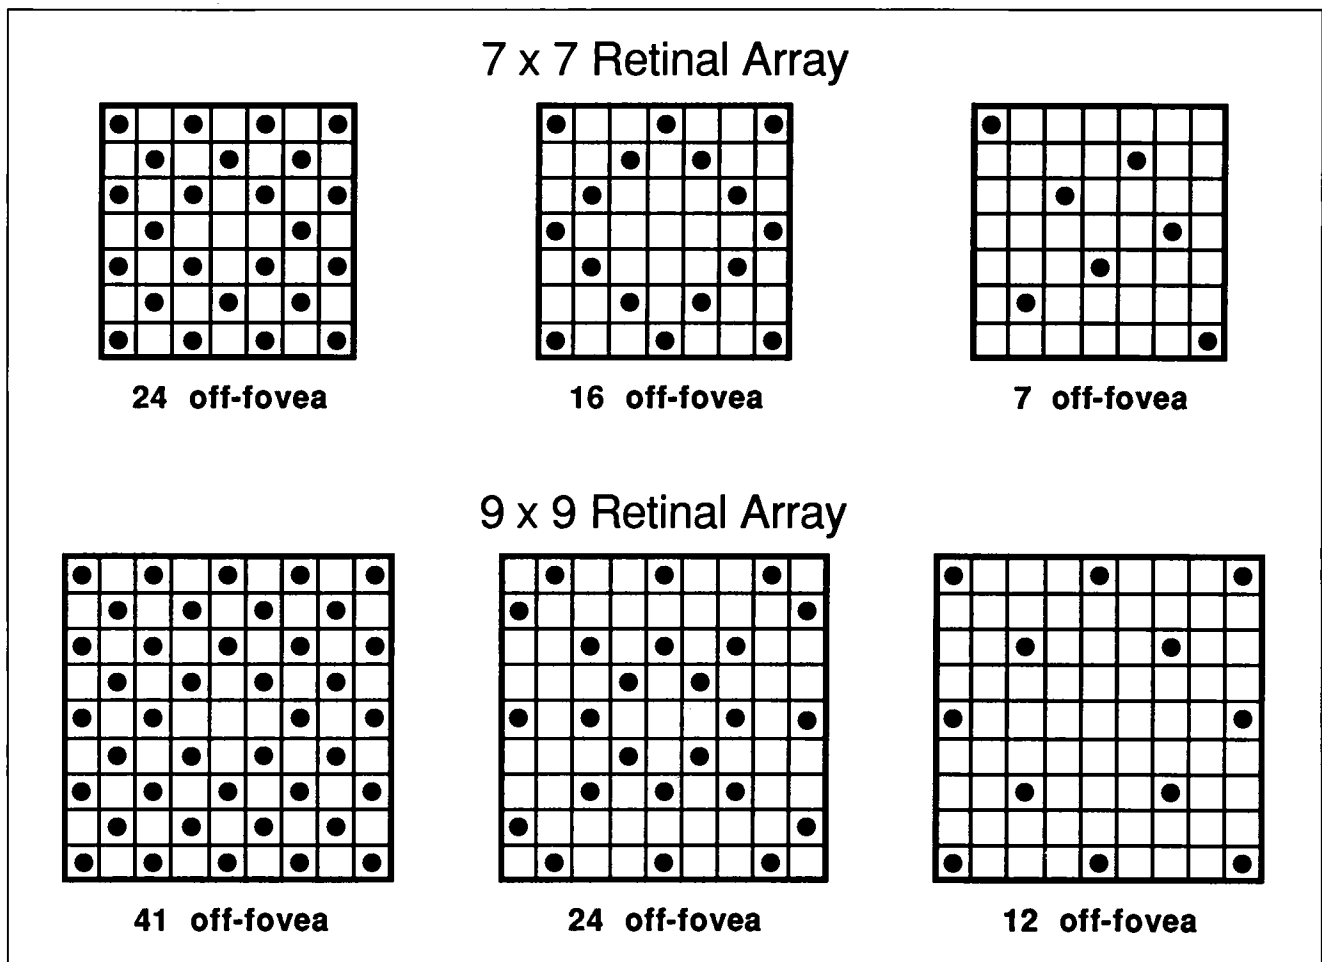

Figure 2. The regular tessellations of the 7 × 7 retinal array (above) and the 9 × 9 retinal array (below) used to derive tessellated locations of off-fovea receptive field peaks for input units. The center cell represents the fovea in all cases.
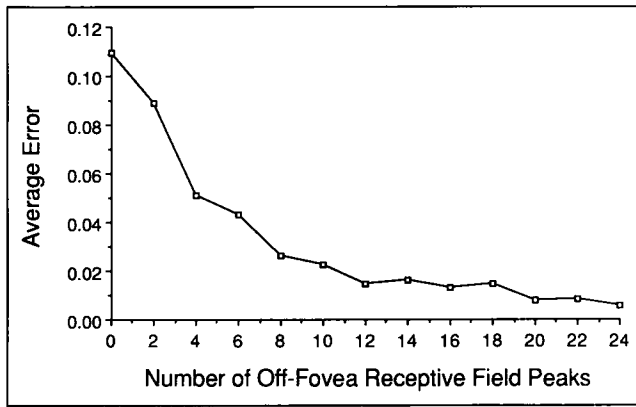
Figure 3. Results of Part 1 of Experiment 1. Mean error per output unit per stimulus for different numbers of off-fovea input unit receptive field peaks (using a 7 × 7 array network configuration).

12 of the 24 receptive field peaks were positioned off the fovea, the amount of error flattened out, as is evident in Figure 3.

This elbow at 12 off-fovea peaks provides evidence for a coarse coding model of location representation. Coarse coding works by interpolating the precise location from the variable activation of several broad representations. The degree of overlap of these representations is the crucial variable in a coarse coding system. The presence of an elbow indicates that there is sufficient lack of overlap after a certain point, before which the network's performance degrades with the decreasing proportion of off-center peaks. When the proportion of off-center peaks is above this critical point, the performance remains roughly the same with only minor improvement with decreasing overlap.

The results from Part 2 of the experiment are illustrated in Figure 4. Average error is plotted in this figure as a function of the number of off-fovea receptive field peaks at four levels: 0 and the three points that allowed for regular tessellations of off-fovea peaks across our 7 × 7 input grid. The decrease in error when more receptive fields had peak responses away from the fovea replicates the findings from the first set of networks reported above, $F(3, 132) = 2350, p < .001$. Contrast analyses revealed that there was a linear decrease in error with increasing numbers of peaks off the fovea, $F(1, 132) = 7,604, p < .001$, and that this decrease decelerated with increasing numbers of peaks off the fovea, $F(1, 132) = 1,695, p < .001$. Together, these trends accounted for 99% of the variance. There was no difference in error when the same random receptive field profile was used versus when every profile was unique, $F(1, 132) = 1.58, p > .20$. In addition, there was a trend for less average error overall when peaks were tessellated compared to when they were distributed randomly (0.0170 versus 0.0186); this difference, analyzed across the three levels allowing for regular tessellation, did not reach significance at the $p < .05$ level, $F(1, 108) = 3.26$.
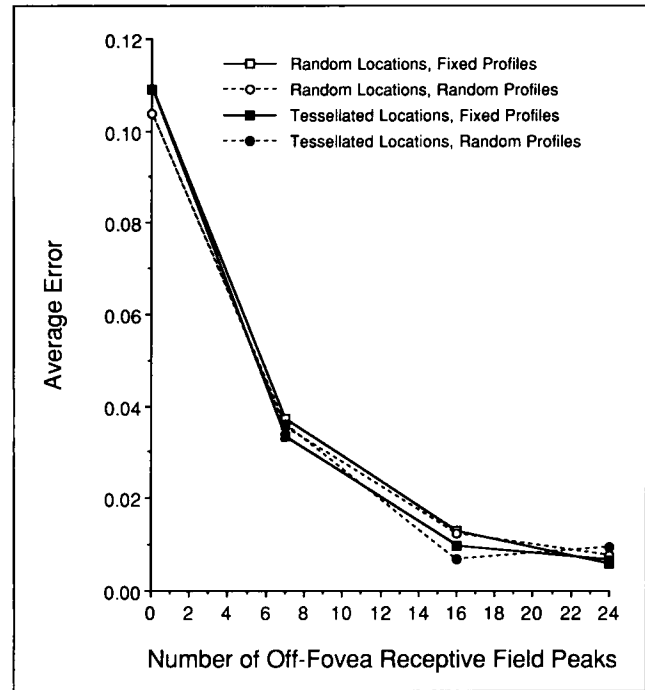

Figure 4. Results of Part 2 of Experiment 1. Mean error per output unit per stimulus for different numbers of off-fovea input unit receptive field peaks (using a 7 × 7 array network configuration). Data are plotted for four conditions created by combining the tessellated versus random distribution scheme with the randomized versus fixed receptive field profiles.

However, the network that used tessellated locations and random receptive field profiles behaved slightly differently from the others when the number of off-fovea receptive fields increased, as is evident in Figure 4, $F(3, 132) = 3.51, p < .05$ for the interaction of proportion of off-peak fields and type of profile.

In short, these results support our hypotheses that area 7a of the parietal lobe can encode location effectively, whereas IT cannot, at least in part because of the distribution of receptive field peaks outside the fovea. These findings are consistent with the fact that lesions of IT do not greatly impair an animal's ability to encode location, whereas lesions of the parietal lobe do (e.g., see Ungerleider and Mishkin 1982). Furthermore, the results do not depend critically on biologically implausible assumptions about the shapes or distributions of receptive fields.

EXPERIMENT 2

The results of Experiment 1 demonstrated that the ease of encoding location depends on the proportion of receptive fields that respond maximally when the stimulus is outside the fovea. Experiment 2 was designed to replicate this finding with a larger input array (a 9 × 9 retina) and, perhaps more importantly, to examine the representations generated by the hidden units of the

model. In the model, coarse coding is possible because of the overlapping receptive fields of units in layer 2; these units modulate the input from layer 1 (the array). Unlike actual neurons in area 7a, however, all of the units in layer 2 had only a single location at which they respond maximally. We conjectured that multiple-peaks might characterize neurons that receive input from single-peak neurons and serve to combine inputs to converge on location. If so, then we expected similar receptive fields to characterize the hidden units in layer 3 of our model, given that these units play just this role in the mapping of input to output.

The retinal array in Experiment 2 was a 9 × 9 matrix, and the output consisted of two sets of 9 units representing the horizontal and vertical coordinates, as in Experiment 1. To compensate for the increased input array size, we had 41 units in layer 2 and 24 units in layer 3. As before, we used the backpropagation algorithm to examine the ease of mapping individual locations in an array to a coordinate representation of location. In this experiment we examined the amount of error after 500 epochs, probing all of the possible 81 locations in each epoch.

This experiment had two parts. In Part 1, we constructed four networks, which had either 40, 24, 12, or 0 of the input units with receptive field peaks outside the fovea. The peaks were arranged in tessellated locations (see the bottom part of Figure 2) and the same field profile was used for each of the input units. Each network was run 10 times. In Part 2, we ran a single trial of each of the four networks until either an error level of 0.001 was reached or 5000 epochs had passed without the error criterion being met. We then recorded the receptive and projective fields of each hidden unit in the networks that reached the error criterion, and considered those that went to 5000 epochs to be incapable of achieving the proper input/output mapping.

Results and Discussion

Figure 5 illustrates the primary results of Part 1 of Experiment 2, which replicated the results of Experiment 1 with a larger network. A one-way ANOVA (with average error after 500 epochs as the dependent measure) documented the effect of varying the number of off-fovea peaks on recovering the location of the input dot, $F(3, 36) = 12.9, p < .001$. Contrast analyses revealed that there was a linear decrease in error with increasing numbers of peaks off the fovea, $F(1, 36) = 29.4, p < .001$, and that this decrease was flat in the intermediate numbers of peaks off the fovea, $F(1, 36) = 7.37, p < .05$ (for the cubic contrast); these two trends together accounted for 95% of the variance.

In Part 2 of Experiment 2, the error criterion was reached only by the network in which 40 of the receptive fields responded maximally to locations off the fovea. It took 1158 epochs for this network to achieve the 0.001
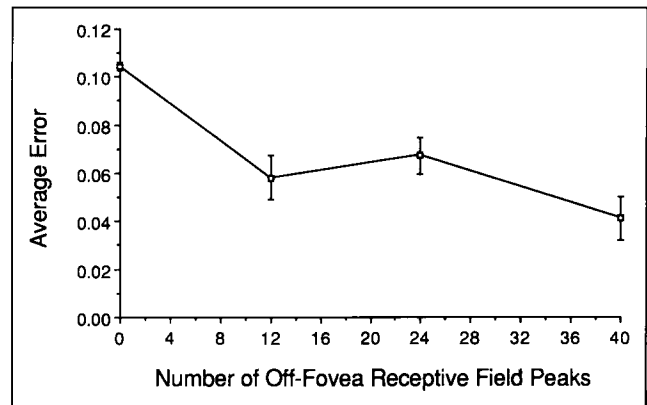

Figure 5. Results of Part 1 of Experiment 2. Mean error per output unit per stimulus for different numbers of off-fovea input unit receptive field peaks (using a 9 × 9 array network configuration).

error level; the other three networks did not attain this level of error within the 5000 epoch cutoff limit. Examination of error levels over epochs for these three networks indicated that they were oscillating around an error level of 0.020 for the last 4000 epochs, and showed no signs of improving their performance past this level.

Following Lehky and Sejnowski (1988a, 1988b), we examined the characteristics of the receptive and projective fields developed by the network to discover how the mapping was achieved. We examined each of the 24 hidden units in layer 3 of the 40 off-fovea peaks network. The three-dimensional graphs presented in Figure 6 represent the activation of each hidden unit when a single point was presented in the different locations on the retinal array, with the height (z axis value) at each location indicating the degree of activation. These receptive fields can be categorized into four groups: Group 1 fields feature a single peak in the center; Group 2 fields have a single, off-center peak; Group 3 fields have asymmetrical multiple local maxima and nonmonotonic drop-off profiles (i.e., unlike Group 2 fields, Group 3 fields included locations off the peak that "bend up," responding more strongly than the surrounding locations); and Group 4 fields have a single depression in the center surrounded by four similar-sized peaks (reminiscent of "foveal sparing").

These types of receptive fields are qualitatively similar to those found in area 7a neurons. Indeed, these results are similar to those obtained by Andersen and Zipser (1988; see also Zipser and Andersen 1988), except that this network recovered some types of multiple-peak profiles (Group 4), whereas their model did not. However, even our 9 × 9 input array is too coarse to do justice to the features of actual receptive fields (e.g., see Figure 2 of Zipser and Andersen 1988), and we must be cautious not to overstate the qualitative patterns suggested by these results.

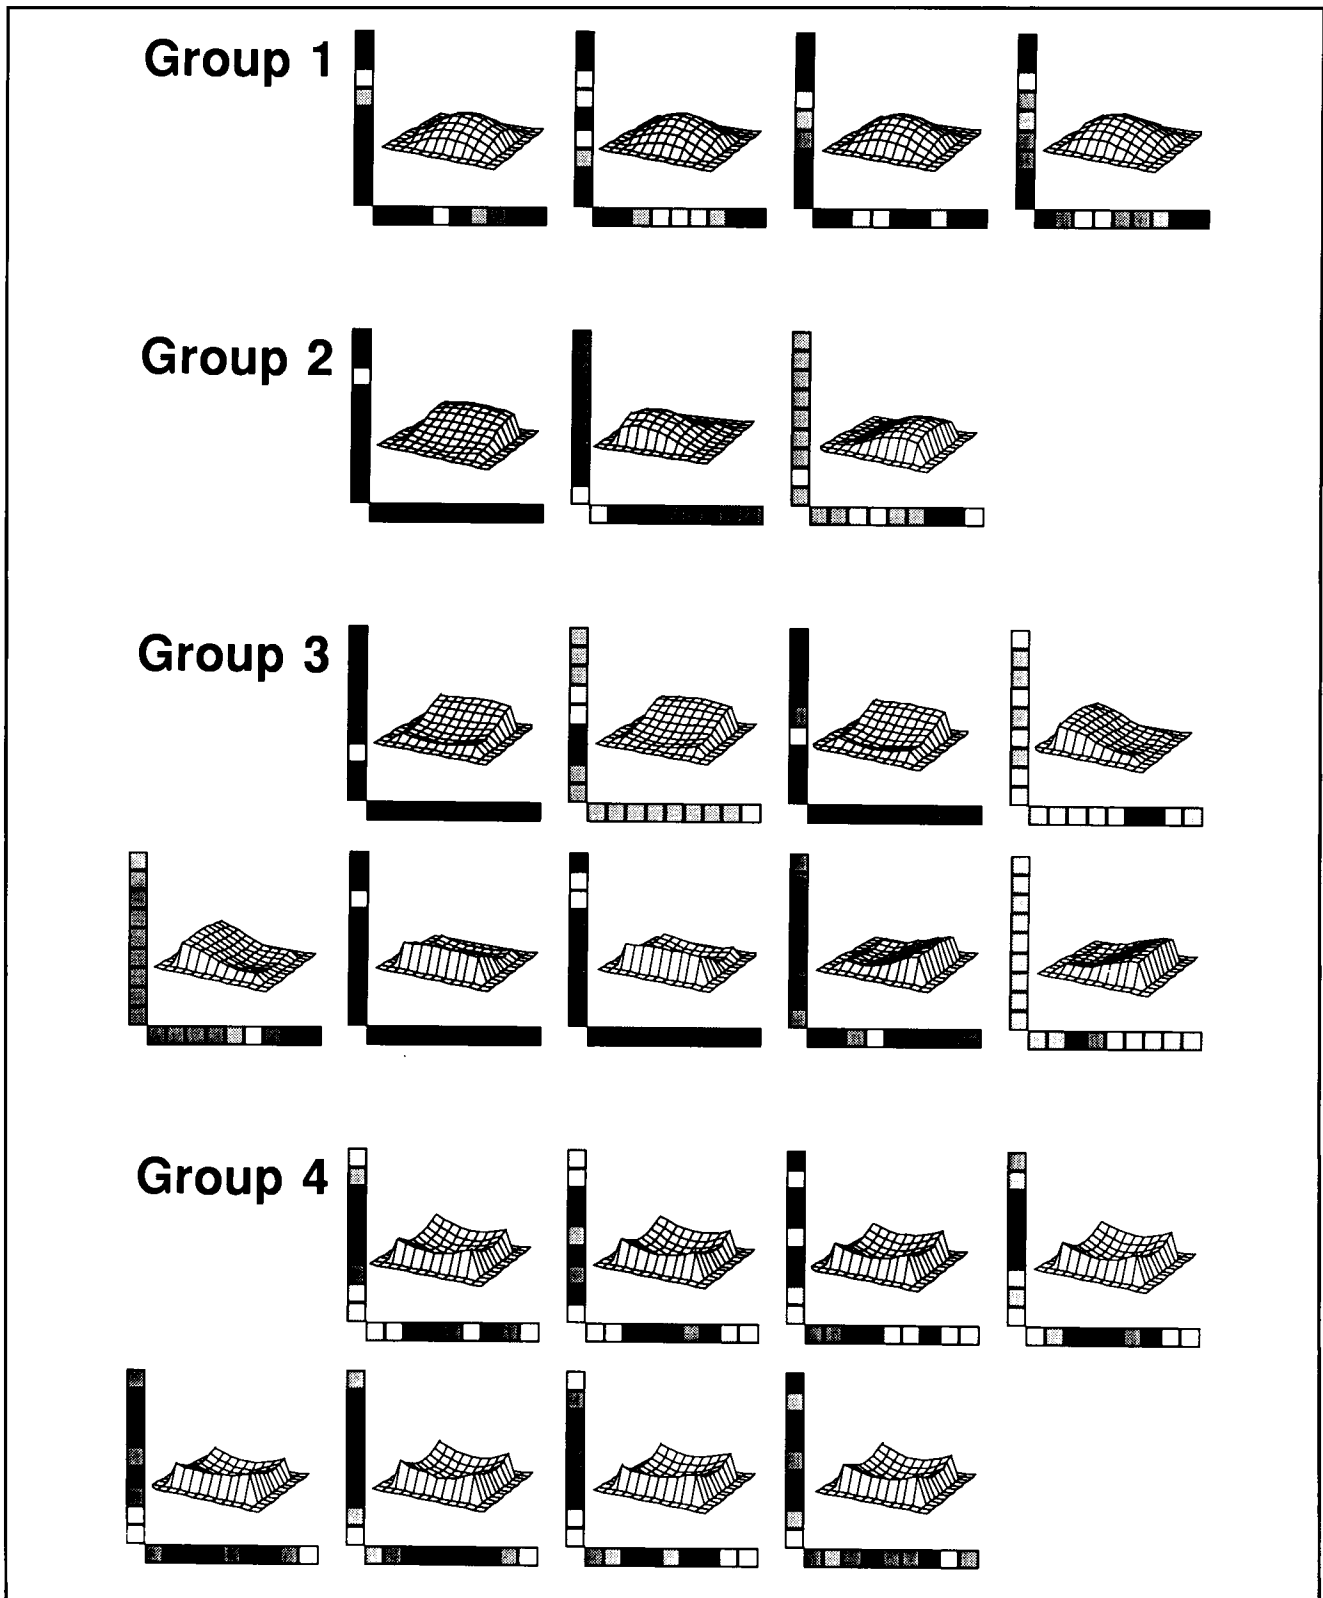In Figure 6, the horizontal and vertical bars represent the weights from each hidden unit to each of the hori-

Figure 6. Results of Part 2 of Experiment 2. Projective and receptive fields of the 24 hidden units from the 40 off-fovea receptive field peaks (9 × 9) network, after the criterion error level was reached. The units are categorized into four groups according to the receptive and projective field characteristics they share. See text for explanation.

zontal and vertical output units; these are the projective fields. Darker shades represent higher weight values; lighter shades represent lower weight values. For the Group 1 hidden units, the projective fields tend to have stronger weights at the periphery for both the horizontal and vertical indices, whereas the receptive fields have an opposite pattern. Group 1 hidden units seem to be performing a computation that differentiates between on- and off-center locations of stimuli. Group 4 hidden units, on the other hand, have receptive and projective fields complementary to those of Group 1 hidden units. The receptive fields of these units spare the foveal region and in general emphasize the periphery of the retinal input, whereas the projective fields tend to have lower values at the two extreme points on each end of both the horizontal and vertical indices. This complementary relation between Group 1 and Group 4 hidden units is particularly evident when one compares the values at the fourth horizontal projective weight, which is the strongest weight in 7 of our 8 cases for the Group 4 units but is the lowest weight in 8 of our 8 cases for the Group 1 units. Thus, it seems that these hidden units have organized into a contrasting set of representations, which specify the location of a point by the degree to which it lies in the center or the periphery.

The theme of complementary receptive fields is also evident in the Group 3 hidden units. These hidden units can be divided into four subgroups, each of which is sensitive to a set of contiguous locations along an extreme end of either the horizontal or vertical axis. Out of nine hidden units in this group, two responded maximally to stimuli along the left end of the horizontal axis, two to stimuli along the right end, two to stimuli at the top end of the vertical axis, and three to stimuli at the bottom end of the vertical axis. The correlation between the projective fields and these receptive field categories is striking. For each of these units, the projective field for the axis in which the receptive field is relatively flat contains homogeneous weights that consistently have values in the mid-range. Thus, these units do not provide much differentiation along one of the axes, indicating that they are performing important computations for a single axis. It is of interest that the maximum values of the projective fields for the active axis correspond to the lower values of the receptive field profiles, just as was found in the hidden units in Groups 1 and 4. In addition, extreme projective field weights for the subgroups of units representing the same axis are arranged in a complementary fashion, so that the maximum and minimum weights from one of the two or three members of the subgroup do not correspond to the same positions as do those in the other members of the subgroup. Furthermore, the maxima and minima are distributed in one-half of the axis, indicating that these hidden units are producing representations specific to one-half of one axis. This specificity implies that these hidden units compute specific coordinate locations, as opposed to the

more general on- and off-center representations created by the hidden units of Groups 1 and 4.

Another interesting property found in the Group 3 projective fields is that the maximal weight is consistently located one unit from the minimal weight. This arrangement would serve to define rather sharp boundaries. This configuration is also evident in the projective fields of hidden units in Group 2, but not in Groups 1 and 4. As in Group 3 receptive fields, there is a complementary distribution of peaks in the receptive fields of the three hidden units in Group 2; the extremes of each axis are represented by a peak, the remaining hidden unit has a peak in the lower left-hand corner, representing the lower horizontal and vertical axes at the same time.

To summarize, the hidden units of the network studied in Part 2 of Experiment 2 organized themselves into two general types of representations. One type, found in Groups 1 and 4, is general, differentiating the on- or off-center location of the stimulus. The other type, found in Groups 2 and 3, isolates specific subsections of the coordinate axes.

EXPERIMENT 3

The results from Experiments 1 and 2 are clear-cut: The distribution of the receptive field peaks of input units critically affects how well a network can encode spatial location. Indeed, we discovered that the peak distribution found in area IT cannot encode location well, which is consistent with the empirical findings. Furthermore, the models were not sensitive to two biologically implausible properties, specifically whether the peaks were tessellated or randomly arrayed and whether all receptive fields had the same drop-off function from the peak. In addition, when we examined the hidden units of a network that could encode location, we found that they developed complex receptive fields like those found in the parietal lobe.

The results of Experiment 2 are consistent with the idea that some neurons in area 7a have complex receptive field profiles because they integrate output from other neurons to compute location. According to the coarse coding hypothesis, there must be units that integrate the overlapping outputs from the input units. However, the previous results do not rule out another possibility, namely that multiple-peak receptive fields are themselves used in coarse coding the input.

As noted earlier, at first glance it seems implausible that multiple-peak receptive fields could play a useful role in providing input to a coarse coding system. Multiple-peak fields produce similar activation when stimuli are in a number of different locations, which intuitively seems likely to hamper the use of relative outputs from overlapping fields to converge on a location. However, it is difficult to judge intuitively what the effect of multiple peaks would be in this kind of computational system, so we decided to investigate the issue empirically.

Experiment 3 was designed to examine how easily the location of a dot can be recovered when the input is filtered through units that respond maximally to stimuli in multiple locations. We examined three networks. The first network's input units had receptive fields that mimicked the hidden unit receptive fields obtained in Part 2 of Experiment 2, which had a mix of single- and multiple-peak recepive fields. The second network's input units had approximately equal numbers of one-, two-, and three-peak receptive fields. And the third network's input units had only two- and three-peak receptive fields. In the latter two networks, the multiple-peak fields were formed by combining single-peak off-fovea fields used in the previous experiments. All networks had a 9 × 9 input array and the general structure of the networks used in Experiment 2; each network was run for 10 trials of 500 epochs each.

Results and Discussion

As is illustrated in Figure 7, the greater the proportion of single-peak receptive fields, the better the network performed. Indeed, the amount of error was predicted by the proportion of input units with single-peak receptive fields. The network that had 100% single-peak re-
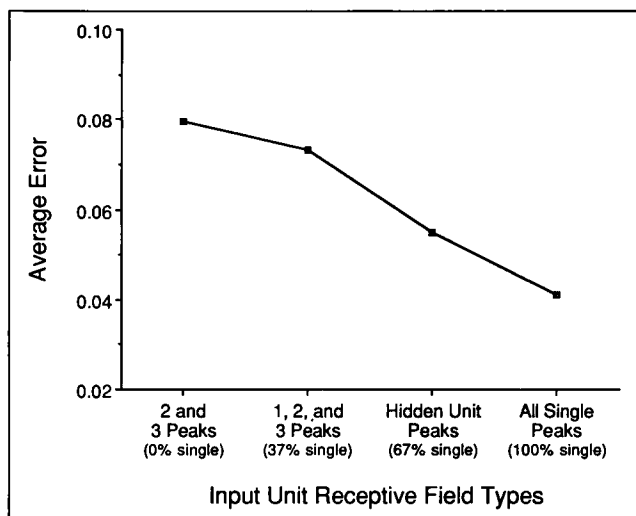

Figure 7. Results of Experiment 3. Mean error per output unit per stimulus obtained from four 9 × 9 array network configurations, each with different input unit receptive fields: 2 and 3 Peaks included roughly equal proportions of only two- and three-peak receptive fields, which were constructed from the same single-peak master field profile; 1, 2, and 3 Peaks included roughly equal proportions of one-, two-, and three-peak receptive fields, which were constructed from the same single-peak master field profile; Hidden Unit Peaks included one-, two-, and three-peak receptive fields; which were taken from those developed by the hidden units in the 40 off-fovea single-peak network of Experiment 2; All Single Peaks indicates that the receptive fields were those used for the 40 off-fovea single-peak network in Experiment 2. The numbers in parentheses indicate the percentage of single-peak receptive fields in the input units for that network.

ceptive fields was better than that with 0% single-peak receptive fields (i.e., with only two- and three-peaked receptive fields), $F(1, 36) = 11.06$, $p < .01$, and was better than the network with 37% single-peak receptive fields (i.e., with a roughly equal mix of one-, two-, and three-peak receptive fields), $F(1, 36) = 7.47$, $p < .01$; it was not better, however, than the network with 67% single-peak receptive fields (i.e., with the receptive fields of the hidden units from Experiment 2 as the receptive fields of the input units), $F(1, 36) = 1.46$, $p > .1$ (if we assume that "multiple peaks" indicates nonmonotonic orderings of activation, in which case the bars running along one side of the receptive field are taken as a single peak). The only other significant difference in Figure 7 is between the network with 67% single-peak receptive fields and that with 0% single-peak receptive fields (i.e., that with hidden unit receptive fields and that with only two- and three-peak receptive fields), $F(1, 36) = 4.49$, $p < .05$. Thus, we can conclude that single-peak receptive fields in fact provide better input to a coarse coding system than do multiple-peak receptive fields.

We next attempted to train a multiple-peak network until an error criterion of 0.001 was achieved. Even the network that performed best through 500 epochs, the one whose input units had the receptive field profiles of the hidden units in Experiment 2, could not achieve the mapping after 5800 total epochs (the average error at that point was 0.018). We concluded that it is probably impossible to train fully a network that includes a relatively large number of multiple-peak receptive fields. This finding provides additional evidence that receptive fields with multiple peaks in the input layer are not computationally efficient at encoding location.

These results, together with those from Experiment 2, are consistent with the idea that multiple "computational layers" exist within area 7a, with the layers playing different roles in registering location. According to this view, the neurons with single-peak receptive fields provide input to those with multiple-peak receptive fields, which use overlap in the input receptive fields to converge on location. The multiple-peak receptive fields, in turn, produce output that provides an explicit representation of location.

GENERAL DISCUSSION

We reasoned that the outputs from numerous receptive fields must be used to represent location in the parietal lobe, given that the individual receptive fields are neither small nor precise. In particular, we hypothesized that overlapping receptive fields could encode spatial location via coarse coding, but only if the units have different locations of maximal response. In contrast, when the locations of maximal response are clustered in one area, as tends to be true in IT neurons, we expected that overlapping fields could not easily encode location via coarse coding. These hypotheses were clearly supported

by the results of both Experiments 1 and 2: When there was a high proportion of off-fovea receptive field peaks (as is characteristic of area 7a neurons), the networks were able to accurately map a point on a retinal array to an explicit coordinate representation of location. In contrast, when there was a low proportion of off-fovea peaks (as is typical of IT neurons), this mapping was not performed accurately.

The analyses of how the successful mapping of input to output was achieved also proved illuminating. We found many properties that are reminiscent of area 7a neurons in the internal, hidden unit representations developed to perform the required mapping. For example, the proportion of on- and off-fovea peaks found in the hidden unit receptive fields was similar to the proportions found in studies of area 7a receptive fields (Andersen et al. in press). Indeed, 83% of the receptive fields of the hidden units had off-fovea peaks, which is further support for our hypothesis that this characteristic is important in encoding spatial location efficiently. In addition, we discovered examples reminiscent of fovea sparing in 33% of the receptive fields of the hidden units (those in Group 4), suggesting that this property may be important in the coarse coding mapping employed by the model. It is intriguing that the relative proportion of hidden units exhibiting "foveal sparing" is similar to that found in area 7a (40% according to Motter et al. 1987).

Andersen and Zipser (1988; Zipser and Andersen 1988) categorized spatially tuned neurons from area 7a into three types, in large part on the basis of the number of locations of maximal response in the receptive fields. In their scheme, Type 1 fields have a single, smooth peak of activity; Type 2 fields have a single, large peak of activity, but also other smaller peaks or depressions; and Type 3 fields have multiple, large peaks. Zipser and Andersen developed a backpropagation model that took eye position and retinal location as inputs and produced an explicit representation of the location of the stimulus in head-centered space. They subsequently examined the types of receptive fields developed by the hidden units. These hidden units produced receptive fields of Types 1 and 2, but did not produce the multiple-peak Type 3 fields. Our network produced fields that were in some respects similar to all three of Andersen and Zipser's types, but we did not recover the full complexities of their fields.

Our Group 1 and 2 receptive fields have clear single peaks, corresponding to Andersen and Zipser's Type 1 neurons. Group 3 receptive fields have a large single peak, a noticeable depression, and smaller peaks, as do Andersen and Zipser's Type 2 neurons. And our Group 4 fields have four similarly sized peaks, corresponding roughly to a subclass of Andersen and Zipser's Type 3 multiple-peak neurons. We did not adopt their taxonomy for our hidden unit analysis because it does not capture the features of the receptive fields we found to be important for coarse coding. They emphasized the number

of peaks and in general the complexity of the variation in response across the field, but did not differentiate between the relative locations of the peaks. Our analysis indicated that the relative locations are critical for encoding spatial location. Furthermore, we found important regularities in the number of peaks in a given location, which are organized in a complementary fashion. The differences between the results from Zipser and Andersen's network and ours could be caused by the fact that they included eye position in the input, had different numbers of hidden units, had different types of input receptive fields, and various other disparities.

The results from the hidden unit analysis of Part 2 of Experiment 2 indicated that the receptive fields of these units became organized into two general types of representations. Representations of broad sets of locations developed in Groups 1 and 4. These representations differentiated the on- or off-center location of a stimulus. In contrast, specific coordinate locations were represented in the receptive fields developed in Groups 2 and 3. This division of labor suggests one particular algorithm that is capable of solving the spatial location encoding problem, and it would be of interest to discover whether the brain in fact uses this method.

Moreover, the results of Experiment 3 demonstrated that there is a computational advantage to single-peak receptive field input to this type of network. The result suggests that there are at least two computational layers within area 7a, and that neurons with more complex fields are involved in combining input from neurons with simpler receptive fields. This hypothesis is consistent with our finding that most hidden units developed complex, multiple-peak receptive fields. To our knowledge, no relevant electrophysiology has yet been performed to test this hypothesis.

This research, then, demonstrates that certain ideas about neural function are computationally plausible. It does not demonstrate that they are in fact correct. Given the expense of neurophysiological research, it seems worth exploring the feasibility of ideas such as these in detail before attempting to test them in animals. But more than that, research such as this helps to define issues more clearly, which cannot help but be useful in furthering our understanding of neural function.

METHOD

Mapping

The mapping being established in the simulations was from a single "illuminated" point in an $N \times N$ array (corresponding to the retina) to a set of $2N$ units that indexed the X,Y coordinates on the retina. One set of N output units was a local representation of the vertical coordinate of the input point on the retina, and the other set of N similarly represented the horizontal coordinate. To represent the illumination of each point on the retina,

N^2 input patterns were required, each having a single unit on and the others off.

Network Architecture

The networks had four layers of units organized in a feedforward structure, as illustrated in Figure 1. The connection weights between layer 1 (the retinal array) and layer 2 were fixed to create receptive field shapes whose characteristics were varied to simulate properties of parietal and inferotemporal lobe neurons in the experiments. The units in layer 2, whose activations were determined by these fixed weights, served as the inputs for the backpropagation algorithm (Rumelhart et al. 1986), which was used to adjust the weights between units in layers 2 and 3 and layers 3 and 4 (the weights between layers 1 and 2 remained fixed throughout the process). We used a learning rate (epsilon) of 0.12 and a momentum factor (alpha) of 0.80 for all network runs; all initial nonfixed weights were given random values between -0.5 and 0.5.

Input Receptive Fields
(Experiments 1 and 2)

The fixed weights between layers 1 and 2 were determined as follows: First a $2N-1 \times 2N-1$ "master" receptive field was generated, and the individual receptive fields for each unit in layer 2 were generated from this by taking an $N \times N$ section from different locations of this master field, depending on where the peak was to be located. The master receptive field had a single peak with a relatively large asymmetrical component and a small random perturbation. Figure 8 illustrates the master receptive field.

This receptive field was generated in four steps. First, a $2N-1 \times 2N-1$ Gaussian surface with a σ of 1.85 was created according to the standard formula and scaled by a factor of 10,000:

$$F_g(x,y) = 10,000 \cdot \frac{e^{-(x^2+y^2)/2\sigma^2}}{2\pi\sigma^2}$$

Second, a random perturbation modulated by the natural logarithm of the Gaussian field element $[F_g(x,y)]$ was applied by the following formula, where RND$(-1,1)$ indicates a pseudorandom number in the range -1 to 1, and MIN(F_g) represents the minimum value in the Gaussian field:

$$F_r(x,y) = F_g(x,y) + 12.0 \cdot \text{RND}(-1,1) \cdot \{\ln[F_g(x,y)] - \ln[0.5 \cdot \text{MIN}(F_g)]\}$$

The natural logarithm of the minimum value of the Gaussian field was subtracted to normalize the factor to the positive-valued range of (0.69–15.01) for the minimum-valued element in the Gaussian field to the maximum element, respectively. Thus, the effect of the randomization was 21.8 times as great for the peak value

as it was for the smaller values, ensuring that the asymmetries so generated would correlate with the scale of the original field. The 12.0 scaling factor resulted in a randomization factor of 38.7% for the peak value, and 2.6×10^6% for the minimum value, which represent relatively large asymmetries in the field.

The third step was to smooth the random field generated in step 2 by convolution with the following 3×3 normalized Gaussian kernel, having a σ of 1.6:

$$\begin{bmatrix} 0.09671 & 0.11756 & 0.09671 \\ 0.11756 & 0.14292 & 0.11756 \\ 0.09671 & 0.11756 & 0.09671 \end{bmatrix}$$

After smoothing, another random factor of ±5% of the maximum F_r was added, which was not modulated by the original magnitude. The factor added a small-scale random perturbation, resulting in the source field, F_s, which needed only to be normalized to the range (0–1). The following formula was used to normalize the F_s, where MAX(F_s) is the peak value in F_s, and MIN (F_s) is the smallest value:

$$F_n(x,y) = \frac{F_s(x,y) - \text{MIN}(F_s)}{\text{MAX}(F_s) - \text{MIN}(F_s)}$$

Having generated the master receptive field, the individual $N \times N$ receptive fields were sampled from it by specifying the location of the peak in the individual $N \times N$ field. Depending on whether a random or tessellated peak location scheme was being used, the peak locations were determined in one of two ways. In the tessellated cases peaks were distributed within the square retinal matrix such that each column and row had the same number and spacing of peaks; using this scheme avoided the problem of unbalanced peak distribution that is possible with randomly located fields. In these cases, the peak locations were derived from the regular tessellations shown in Figure 2 for both the 7×7 and the 9×9 cases. The randomly distributed peaks were selected without replacement from a list of randomly generated peak coordinates.

Because each receptive field for the units in layer 2 was simply a copy of a region of the master field, all fields with the peak in the same location had exactly the same profile. This does not accurately reflect the variance between the receptive fields from single-cell recordings, so in Experiment 1 we compared the effect of the addition of a random element to each receptive field for the units in layer 2 as they were sampled from the master field. The individual randomization was produced by adding an additional random factor, which was limited to a maximum of ±5% of the maximum receptive field value, to each element in the resulting $N \times N$ field. Fixed profile receptive fields were sampled directly from the master receptive field without this additional random factor.
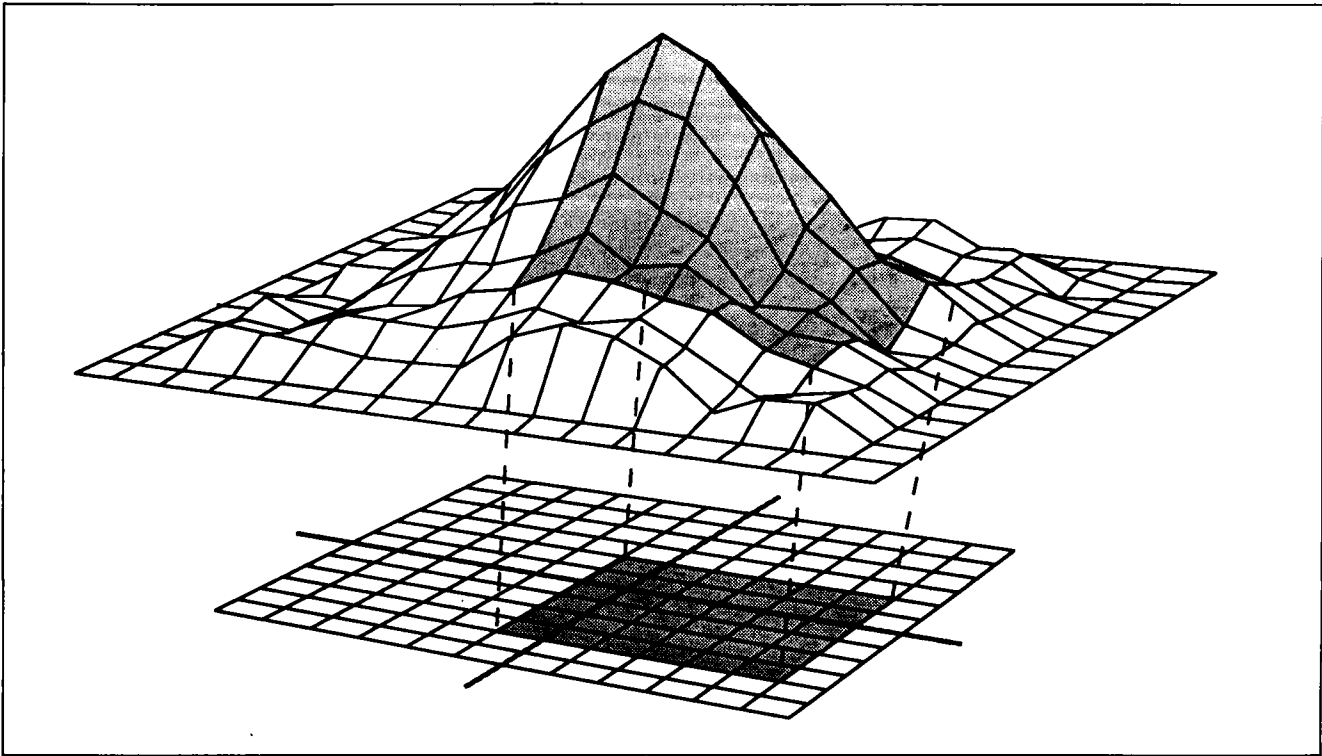
Figure 8. The selection of values from a subsection of the "master" (7 × 7) receptive field (above) to be used for individual input unit receptive field values (below). These are the fixed connection weights between points on the retina and a single input unit illustrated at the left of Figure 1.

Input Receptive Fields (Experiment 3)

In Experiment 3, we considered three cases of complex, multiple-peak input receptive fields. The first of these cases was generated by using the hidden unit receptive fields created in the 40 off-center case of Experiment 2 after training to 0.001 error. In order to generate the 41 receptive field patterns from the 24 hidden unit receptive fields, the numbers of fields from each of the four groups derived in the analysis of Experiment 2 were scaled up in roughly the correct proportions by duplicating a portion of the units as follows (units are ordered left to right, top to bottom in Figure 6): 7 receptive fields from Group 1 (the correct number should be 6.833 by proportion), using the first three twice each and the fourth once; 6 from Group 2 (5.125 by proportion), using each twice; 15 from Group 3 (15.375 by proportion), using the first, second, fourth, sixth, seventh, and eighth twice and the others once each; 13 from Group 4 (13.666 by proportion), using the first five twice and the remaining three once each.

The second multiple-peak training set was generated by using the tessellations for the 9 × 9 case (see Figure 2) to specify the peak locations for a set of single-peak stimuli, which were subsequently combined to form multiple-peak fields (two or three peaks). For the case with one, two, and three peaks, we used 15 single peaks, 12 double peaks, and 14 triple peaks. The 24 off-fovea

tessellation was used for the 12 double peaks, and the 41 off-fovea tessellation, plus an additional peak in the center, was used for the 14 triple peaks. The 15 single peaks were drawn from the 12 off-fovea tessellation with 3 additional units at coordinates $(-4,-2)$, $(4,2)$, and $(0,0)$.

For the case with double and triple peaks, 21 double peaks and 20 triple peaks were used. The same 42 peaks that were used in triple peak receptive fields in the previous case were used for the 21 double peaks. The 60 peaks necessary for the 20 triple peak receptive fields were derived by taking the 40 "holes" (unused locations) from the 41 tessellation, and adding to them 20 locations of the 24 tessellation [the four points located at coordinates $(-2,-2)$, $(2,-2)$, $(-2,2)$, $(2,2)$ were eliminated]. The peaks for the double and triple fields were selected so that the distance between peaks was not less than four grid units. The single-peak receptive field profiles were combined for the double-peak case using the following formula:

$$F_d(x,y) = \text{MAX}[F_1(x,y), F_2(x,y)] \cdot \text{MAX}[F_1(x,y), F_2(x,y)] + \{1 - \text{MAX})[(F_1(x,y), F_2(x,y)]\} \cdot \text{MIN}[F_1(x,y), F_2(x,y)]$$

A similar formula was used for the triple-peak case, as follows:

$$F_t(x,y) = \text{MAX}[F_1(x,y), F_2(x,y), F_3(x,y)]^2 +$$
$$\{1 - \text{MAX}[F_1(x,y), F_2(x,y), F_3(x,y)]\} \cdot$$
$$\text{MEDIAN}[F_1(x,y), F_2(x,y), F_3(x,y)]^2 +$$
$$\{1 - \text{MAX}[F_1(x,y), F_2(x,y), F_3(x,y)]\} \cdot$$
$$\{1 - \text{MEDIAN}[F_1(x,y), F_2(x,y), F_3(x,y)]\} \cdot$$
$$\text{MIN}[F_1(x,y), F_2(x,y), F_3(x,y)]$$

Procedure

We investigated how difficult it was for the backpropagation algorithm to train a network to map a point in an input array to an explicit representation of its location. The fixed receptive fields were used to compute the outputs from the population of input units, which served as the input to a three-layer standard backpropagation model (Rumelhart et al. 1986). The entire set of N^2 layer 2 input patterns was presented to the network in each epoch, and the weights were adjusted after the end of the epoch. The initial weights on the connections between layers two, three, and four were different randomly generated numbers between -0.5 and 0.5. Our measure of the difficulty of establishing the input/output mapping was simply the error remaining after a fixed number of epochs. The error measure used throughout was the square of the error per output unit per stimulus, which ranges between about 0.250 for a completely random set of weights to 0.001, which was used as a final error cutoff. Unlike the commonly used total sum of squared error measure, our measure of error has the same range for all network/mapping combinations, allowing comparison of different networks and different versions of the same mapping.

The receptive fields of hidden units in fully trained networks were generated by recording the activations of the hidden units (layer 3) that resulted when each point on the retina was excited. During this recording, we eliminated the sigmoidal activation function because it flattened the fine structure of the input weights that determined the receptive field of each hidden unit; as graphed, these fields were normalized in the z axis from 0 to 1. The projective field of each hidden unit onto the output layer was determined by recording the weights from that unit to each of the output units. These weights were individually normalized and gray levels assigned by dividing the total normalized range into 10 equal subdivisions.

Acknowledgments

REFERENCES

Andersen, R. A. (1986). Value, variable, and coarse coding by posterior parietal neurons. Behavioral and Brain Sciences, 9, 90–91.

Andersen, R. A. & Zipser, D. (1988). The role of the posterior parietal cortex in coordinate transformations for visual-motor integration. Canadian Journal of Physiology and Pharmacology, 66, 488–501.

Andersen, R. A., Essick, G. K., & Siegel, R. M. (1985). Encoding of spatial location by posterior parietal neurons. Science, 230, 456–458.

Andersen, R. A., Asanuma, C., Essick, G., & Siegel, R. M. (in press). Cortico-cortical connections of anatomically and physiologically defined subdivisions within the inferior parietal lobule. Journal of Comparative Neurology.

Ballard, D. H. (1986). Cortical connections and parallel processing: Structure and function. Behavioral and Brain Sciences, 9, 67–120.

Desimone, R. & Gross, C. G. (1979). Visual areas in the temporal cortex of the macaque. Brain Research, 178, 363–380.

Desimone, R., Shein, S. J., Moran, J., & Ungerleider, L. G. (1985). Contour, color and shape analysis beyond the striate cortex. Vision Research, 25, 441–452.

Gross, C. G. & Mishkin, M. (1977). The neural basis of stimulus equivalence across retinal translation. In S. Harnad, R. Doty, J. Jaynes, L. Goldstein, & G. Krauthamer (Eds.), Lateralization in the Nervous System. New York: Academic Press.

Gross, C. G., Rocha-Miranda, C. E., & Bender, D. B. (1972). Visual properties of neurons in inferotemporal cortex of the macaque. Journal of Neurophysiology, 35, 96–111.

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. Cambridge, MA: MIT Press.

Lehky, S. R. & Sejnowski, T. J. (1988a). Network model of shape-from-shading: Neural function arises from both receptive and projective fields. Nature, 333, 452–454.

Lehky, S. R. & Sejnowski, T. J. (1988b). Neural network model for the cortical representation of surface curvature from images of shaded surfaces. In J. S. Lund (Ed.), Sensory Processing in the Mammalian Brain. Oxford: Oxford University Press.

Marr, D. (1982). Vision. New York: Freeman.

Maunsell, J. H. R. & Newsome, W. T. (1987). Visual processing in monkey extrastriate cortex. Annual Review of Neuroscience, 10, 363–401.

Mishkin, M. & Ungerleider, L. G. (1982). Contribution of striate inputs to the visuospatial functions of parieto-preoccipital cortex in monkeys. Behavioural Brain Research, 6, 57–77.

Mishkin, M., Ungerleider, L. G., & Macko K. A. (1983). Object vision and spatial vision: Two cortical pathways. Trends in Neurosciences, 6, 414–417.

Motter, B. C. & Mountcastle, V. B. (1981). The functional properties of the light-sensitive neurons of the posterior parietal cortex studied in waking monkeys: Foveal sparing and opponent vector organization. The Journal of Neuroscience, 1, 3–26.

Motter, B. C., Steinmetz, M. A., Duffy, C. J., & Mountcastle, V. B. (1987). Functional properties of parietal visual neurons: Mechanisms of directionality along a single axis. The Journal of Neuroscience, 7, 154–176.

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.

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.* Cambridge, MA: MIT Press.

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.* Cambridge, MA: MIT Press.

Van Essen, D. C. (1985). Functional organization of primate visual cortex. In A. Peters & E. G. Jones (Eds.), *Cerebral Cortex*, Volume 3. New York: Plenum.

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.

This article has been cited by:

1. Padraic Monaghan, Richard Shillcock. 2004. Hemispheric Asymmetries in Cognitive Modeling: Connectionist Modeling of Unilateral Visual Neglect. *Psychological Review* **111**:2, 283-308. [CrossRef]