Brain Research Bulletin, Vol. 21, pp. 505–512. © Pergamon Press plc, 1988. Printed in the U.S.A. 0361-9230/88 $3.00 + .00

The Role of the Teacher in Learning-Based Models of Parietal Area 7a

DAVID ZIPSER*[1] AND RICHARD A. ANDERSEN†

**Institute for Cognitive Science, University of California, San Diego*
†Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology

ZIPSER, D. AND R. A. ANDERSEN. *The role of the teacher in learning-based models of parietal area 7a.* BRAIN RES BULL **21**(3) 505–512, 1988.—The back-propagation learning procedure can be used to train simulated neural networks to compute arbitrary functions. We have recently shown that when such a network is trained to carry out the transformation of stimulus location to head-centered coordinates that occurs in parietal area 7a, the response properties of certain units in the network closely resemble neurons found in area 7a. The back-propagation procedure requires the use of a teacher. Here we examine the effect of using different kinds of teachers. As long as the teacher represents information about stimulus location in head-centered coordinates, the trained network contains units of the kind found in area 7a. Differences in teacher format only effect the quantitative distribution of the different unit types. When the teacher does not represent stimulus location explicitly, the network does not contain units of the required kind.

Parietal area 7a Back propagation Neural models

ONE of the main theoretical problems facing neurobiology is to explain how the brain carries out computations. For example, to determine such a simple thing as the head-centered location of a seen object, information about its position on the retina must be extracted from the retinal images and then combined with information about the direction in which the eyes are pointing. In the past theorists have often tried to develop specific neural networks to cary out each computation like this. Unfortunately, even a single computation can be implemented by an unlimited number of different networks, and even with specifically designed networks it has been exceedingly difficult to account for the actual neural firing patterns observed experimentally.

We have recently described a new approach to the problem of specifying how neural networks in the brain carry out computation (16). Our approach is based on the use of a neural network training procedure, called "back propagation," which can train simple model networks to compute a wide range of functions (12). Training a computer-simulated model of a neural network using back propagation requires knowledge of the correct output to associate with each input to the network, but it is not necessary to specifically tailor the network to do the computation because the learning process will discover an implementation. The source of the correct output values to use for training is called the "teacher." The back-propagation procedure configures the network to implement the computation by adjusting the strengths of synapses connecting neural-like units in the network. The importance of the back-propagation procedure is that it provides a way to make these changes in networks having more than one layer. This means that the networks used can have internal or "hidden" units that are free to take on the response properties that best accomplish the computation being learned. Our motivation for using back propagation was that for many problems learned by these networks, the hidden units often show a general resemblance to cortical neurons in their response properties, particularly in the way they tend to act as feature detectors.

The target of our original modeling effort was area 7a of the monkey parietal cortex. Area 7a together with 7b, MST, and LIP comprise the posterior half of the posterior parietal cortex (2–4, 7, 10, 11, 14). Lesions to the posterior parietal cortex in monkeys and humans produce profound spatial deficits in both motor behavior and perception (1, 5, 6, 9). Based on single-unit recording data and lesion studies, Andersen and others (3,6) proposed that parietal area 7a performed a spatial transformation from observation-based to head-centered coordinates by combining retinal-based and eye-position information. About 55% of the neurons in area 7a respond to both the retinal location of a visual stimulus and the position of the eyes, but are virtually insensitive to qualities of the stimulus such as size, color, shape, or intensity. By combining eye and retinal positions these 7a neurons can code the spatial location of external objects. However, these neurons do not give a constant response to spatially fixed stimuli. Their firing rates change as a function of both eye positions and retinal location of the stimulus. This means that if head-centered spatial location is represented by these

---

[1] Requests for reprints should be addressed to Dr. David Zipser, Inst. for Cognitive Science, UCSD, C-015, La Jolla, CA 92093.

neurons, it must be encoded in the changing patterns of their firing rates.

Our original back-propagation modeling studies demonstrated that a simple network can be trained to compute the transformation hypothesized to occur in area 7a and, more significantly, that the hidden units in this network have response properties that mirror those in area 7a (16). In particular, the hidden units in the model network changed their output activities with eye position and stimulus retinal location in very much the same way as neurons in area 7a. Because of the similarity between the model hidden units and the neurons in area 7a, this result also shows how the neurons in area 7a can encode head-centered spatial location in a way that can be decoded by other neurons.

Back propagation is an error-correcting procedure that requires a teacher to supply the correct answer on each training trial. An understanding of the role played by this teacher is important for two reasons. First, the properties of the hidden units depend on the nature of the teacher, so a study of the effects of using different kinds of teachers can tell us what features the teacher must have to produce units like those found in area 7a. Second, an understanding of the teacher requirements may prove useful in designing experiments to determine whether the brain uses some learning procedure analogous to back propagation to program its computations, as opposed, for example, to genetic determination or an adaptive mechanism that needs no teacher. In this paper we investigate the effects of both different teacher formats and different information on the area 7a model. It is shown that when the teacher contains head-centered location information in a variety of different formats the networks resemble 7a quite closely. When a teacher is used that does not contain explicit head-centered location information, but uses one of the same formats, the networks generated differ completely from what is observed in area 7a. This is so even though the information provided by this teacher is consistent with the required hidden unit response pattern, and a format is used that works when given head-centered information. Before describing what was done and its possible significance, it is necessary to take a closer look at the details of the area 7a model and the role of the teacher in the back-propagation paradigm.

THE BACK-PROPAGATION MODEL OF AREA 7a IN MORE DETAIL

In this section we review in some detail the relevant aspects of our previously reported results (16). Area 7a is hypothesized to take part in the computation of the transformation of stimulus location from observation-based to body-centered coordinates. The computer-simulated network we trained to model this computation was not designed to resemble the complex anatomy and physiology of a typical cortical area such as 7a. Rather, it was the simplest kind of network we felt could be trained to carry out the required coordinate transformation. The network used has only three layers of neuron-like units. These layers consist of an input layer, whose whole function is to hold the input values constant during each training cycle, a hidden layer, and an output layer. The units in the hidden and output layers are very simple models of neurons. These simple model neurons use average firing rates, not individual spikes, to represent their input and output values. The output of a neuron is computed as the synaptic strength weighted sum of the input activities passed through an S-shaped function that limits the outputs to the range 0 to 1 (12).

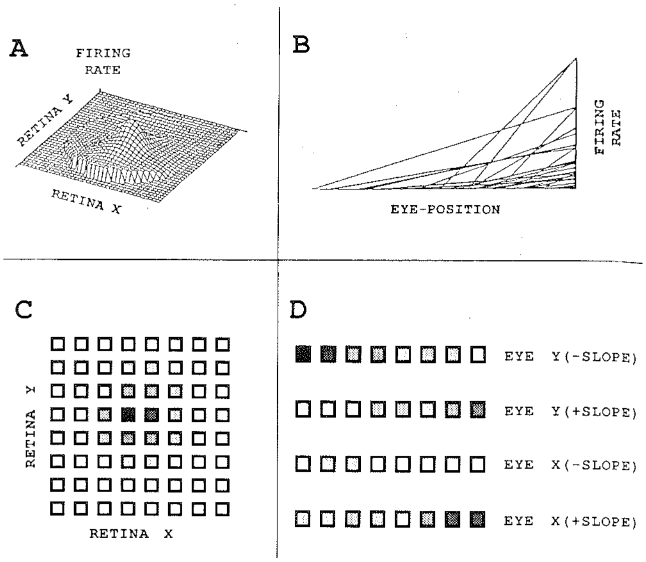

FIG. 1. The relationship between experimental data from area 7a and input to the model network. (A) The receptive field of a neuron in area 7a that responds only to visual stimulation and not to changes in eye position. The receptive field was measured and plotted as detailed in the caption to Fig. 3. The receptive field shown is 80 degrees in diameter with a single, almost Gaussian peak at the center of the retina. Fields like this were assumed to reflect the retinal stimulation input to area 7a, and were used as a basis for designing the input to the model network, as described in (C) below. (B) This is a graph of the response of 30 eye-position neurons in area 7a. These neurons changed their activity as a nearly linear function of eye position, but showed no effect of visual stimulation. The majority of such neurons fall into one of four classes: those that increase their activity with eye movement either to the right, to the left, up, or down. Neurons from all four classes are plotted in (B) as if they were all of the class that increase activity with eye movement to the right. Note that the slopes and intercepts of the responses differ greatly between neurons. (C) The visual stimulus to the model was formatted using 64 units. Each of these units has a Gaussian receptive field of the type seen in (A). The centers of these fields are uniformly distributed on a grid as shown. The response of this array to a point visual stimulus of the kind used is indicated by the relative darkening of the units. (D) The eye position input to the model consists of 32 units divided into four groups of 8. Each group represents neurons in one of the observed classes. The slopes and intercepts for each unit's response was chosen at random, but in the figure the units have been drawn in order of increasing response for their respective direction specificities.

The process of training the network consists of repeatedly giving it inputs representing the retinal location of a stimulus and the current eye position, and training it to produce the associated head-centered location of the stimulus on its output units. The teacher provides correct values of the output unit representation of head-centered location. Learning requires a large number of training cycles. The longer the training goes on, the more accurate the output becomes. We typically continued training until the error with which the output indicated stimulus location was about equal to the distance between screen points that project to neighboring retinal units. That is, we stopped training before hyperacuity set in. This required about 1000 training cycles. Further training produced more accurate outputs, but did not make the response properties of the hidden units more like area 7a.

While no attempt was made to model the physical struc-

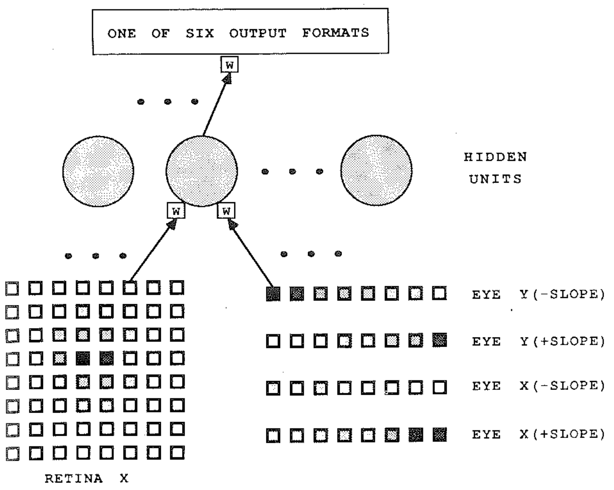

FIG. 2. Back-propagation network used to model area 7a. The input, hidden, and output layers are described in the text and in Fig. 1. The input to the network consists of retinal-position and eye-position information. The output activation of the hidden and output layer units is given by the logistic:

$$output\ activation = 1/(1 + e^{-net}),$$

where *net=weighted sum of inputs +bias*. The arrows indicate the direction of activity propagation. The *w*'s are the weights changed by learning. The network was simulated on a Symbolics 3600 LISP machine using the P3 parallel system simulator (17).

ture of area 7a, we did try to model the format of its afferent information using experimental data. The area 7a neurons relevant to the model's input are the eye-position neurons, responding to eye position only, and the visual neurons, responding to visual stimulation only. These neurons presumably represent the eye-position and retinal location information used by area 7a as input, so we modeled the network's input on their properties. The properties of the visual and eye-position neurons and how they were used to represent the input to the model network are shown in Fig. 1. Most of the eye-position neurons in area 7a responded monotonically to either the horizontal or vertical position of the eyes in the orbit. This was modeled by using 32 units with eye-position response parameters selected from the experimentally observed range. The visual neurons in area 7a have large, peaked retinal receptive fields with the peak centers distributed over the whole retina. The retinal location input was modeled using 64 units with receptive fields shaped like the simplest, most symmetrical found experimentally. The peaks of these receptive fields were distributed evenly over the modeled retinal space. Every unit in the input layer is connected to each of the hidden units, which were connected to all of the output units. Between 9 and 36 hidden units were used. The synaptic weights are initially set to small random values. The whole model network is diagrammed in Fig. 2.

A cycle of training consists of picking an eye position and retinal location at random and converting them to the format used for input to the network. This input activity pattern is then propagated through the network, first to the hidden units and then to the output units. The resulting output unit activity pattern is subtracted from the correct pattern, that is, the teacher pattern. The resulting difference, or "error," is used to modify the synaptic weights of the output and hidden units as proscribed by the back-propagation procedure. Note that the teacher pattern and the output units have matching formats. At the beginning of training, the values of the output units and the teacher pattern do not match closely, whereas after the network has been successfully trained to do a computation they are nearly equal.

In our original modeling efforts we used two different teacher formats, each representing spatial location in a head-centered frame. One format represented spatial location as the eye position at which the stimulus would be foviated. In this case the output layer and corresponding teacher pattern consisted of 32 eye-position units with properties like those of the eye-position inputs to the model. The other format represented head-centered spatial location as the retinal location of the stimulus when looking straight ahead. In this case the output layer and teacher pattern consisted of 64 units arranged in the same way as the retinal input to the network. Note that although the format of the output layer corresponded to part of the input layer in each of these cases, the actual values appearing on the input and output layers were completely different. In particular, the input layer values represent location in an observation-based reference frame, whereas the output layer values represent location in a head-centered frame.

These two different teacher formats were used to train the model in separate training sessions. With both of these teacher formats, the network learned to compute the transformation from observation-based to head-centered coordinates postulated to occur in area 7a. The really interesting aspect of these training simulations is that the network learned to do this computation in a way analogous to area 7a. This can be seen by comparing the hidden units' response patterns to the firing rates of the neurons in area 7a that respond to both a visual stimulus and eye position. The experimental data available for comparison with the model consist of measurements of retinal receptive fields at fixed eye positions and the effects of different eye positions with fixed stimulus retinal location. Both of these measurements have been compared to the model to determine the degree of similarity.

The comparison of experimental and model receptive fields is shown in Fig. 3. The experimental data on which this comparison is based was collected several years before the model was developed, and no new data has been obtained since. In these comparisons normalized firing rates are compared to normalized hidden unit activity. The comparison is difficult because there is such a wide variety of large complex fields. The only feature of the receptive fields that could be quantitatively compared between model and experiment was the eccentricity of the location of the peak activity. However, using a smoothing algorithm to visualize the overall shape of the receptive fields allows a more revealing semiquantitative comparison to be made. In Fig. 3 the fields are arranged in columns of the same eccentricity of the highest peak, and in rows with fields of about the same complexity of peak structure. Complexity of peak structure is based on an estimate of the number of major and minor peaks. At the present time this is the best method of comparison we have. The small number of available experimental fields hinders the development of meaningful statistical comparisons. The top row in Fig. 3a and 3b has the fields with one major peak and no significant minor ones. The center row has fields with a clear major peak and at most a few minor peaks. The bottom row has all the rest. While this measure of peak

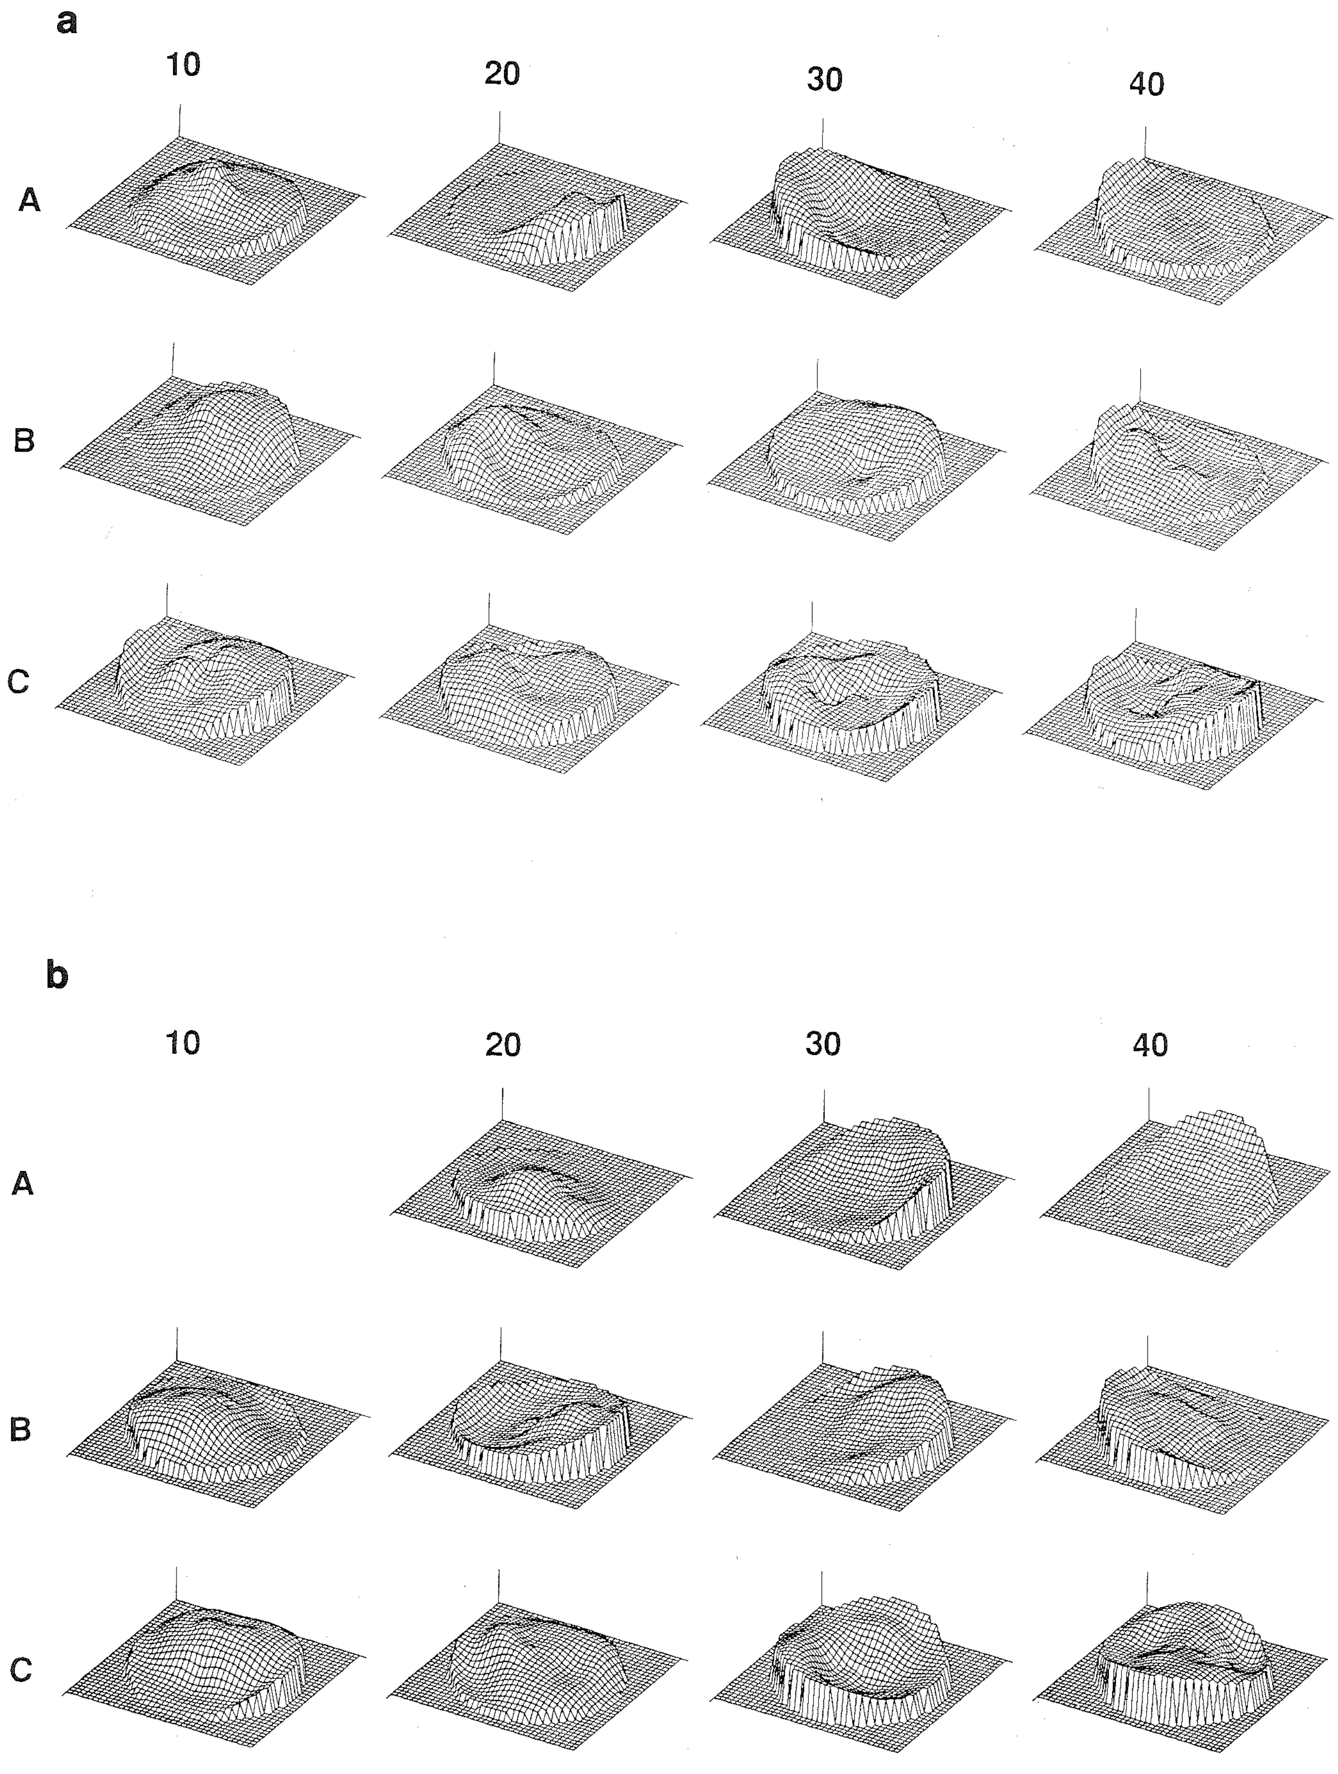
a
10
20
30
40
A
B
C
b
10
20
30
40
A
B
C

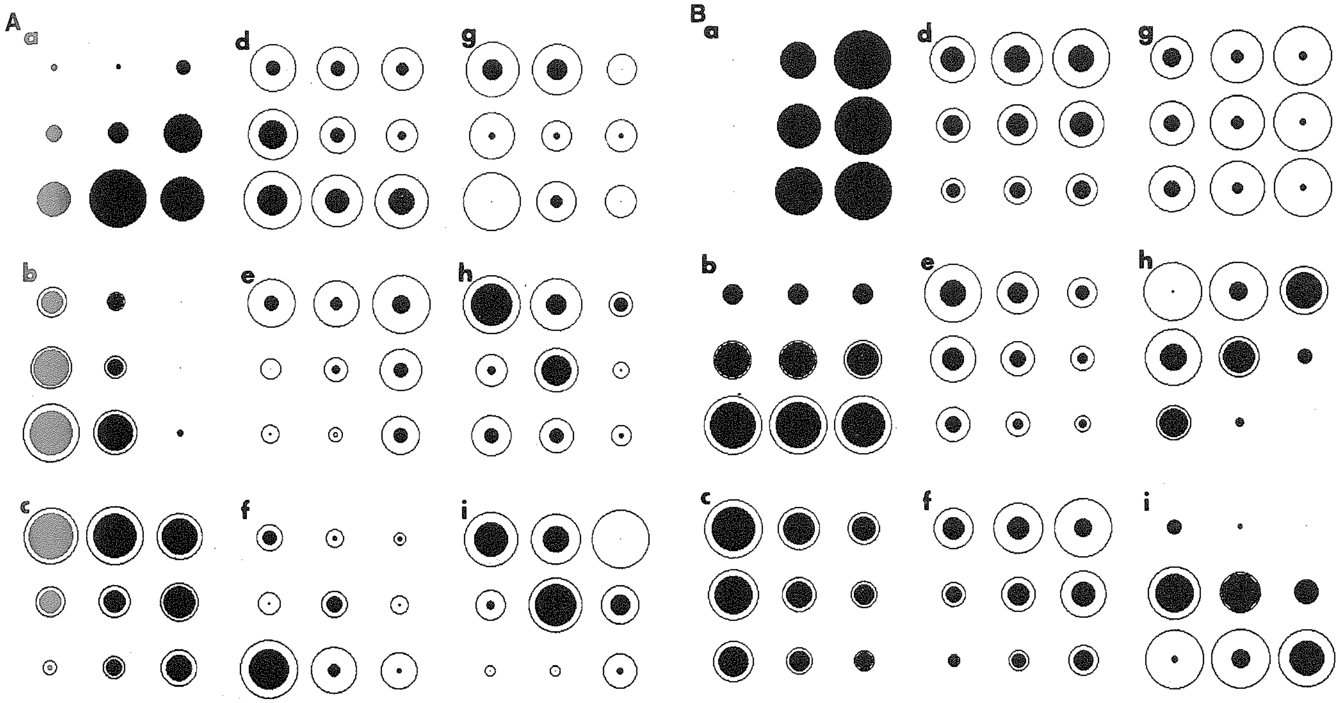

FIG. 4. (A) The spatial gain fields of nine neurons from area 7a (16). The diameter of the darkened inner circle, representing the visually evoked gain fields, is calculated by subtracting the background activity recorded for 500 msec before the stimulus onset from the total activity during the stimulus. The outer circle diameter, representing the total response gain fields, corresponds to the total activity during the stimulus. The annulus diameter corresponds to the background activity that is due to an eye position signal alone, recorded during the 500 msec prior to the stimulus presentation. (B) Hidden unit spatial gain fields generated by the model network. Fields a–f were generated using the monotonic format output; the rest used the Gaussian format output.

complexity is only semiquantitative, it does simplify the comparison of model and experimental fields. Although only 12 fields from model and experiment are compared, they represent all the complexity-eccentricity types observed. All field types observed in the model were also found in area 7a. The reverse is not true. Fields with an eccentricity of 10 degrees and a single peak are observed in area 7a but are very rarely seen in the model. Note also that three of the fields in the bottom row of Fig. 3b come from the model before it has been trained. The data in Fig. 3 make it quite clear that there is considerable similarity between the shapes of the retinal receptive fields of model and experiment.

The pattern of variation in visual response as a function of eye position, when the retinal location of the visual stimulus is kept fixed, is called the "spatial gain field." Spatial gain fields were measured by having the monkey fixate nine different positions on a frontal screen. At each of these fixation points the response was first measured in the absence of visual stimulation. Then the response was measured with a visual stimulus flashed at a screen position chosen to keep its retinal location the same for all eye positions. Measurements at these nine eye positions constitute the spatial gain field data for each area 7a neuron tested (2). The firing rate of most area 7a neurons changed in a systematic way as a function of the eye position in both the presence and absence of visual stimulation. Since the intensity and retinal location of the visual stimulus is the same at each eye position it might be expected that the effect of visual stimulation could be computed by just adding a constant representing the effect of the visual stimulus to the eye position alone response. This is not what is observed. Rather, the contribution of visual stimulation to the total response is also seen in the model's hidden units.

To facilitate the comparison of experiment and model we have visualized these nonlinear interactions by using the concentric circle pattern described in Fig. 4. The outer diameter of each circle is proportional to the normalized activity of a unit in the presence of visual stimulation. The width of the

FACING PAGE

FIG. 3. (a) Experimentally determined retinal receptive fields (2,16). The data for drawing each of these receptive field plots comes from measurements of the firing rate of a single area 7a neuron at 17 different retinal locations. These locations were at the center and at 10, 20, 30, and 40 degrees out. A neighborhood smoothing procedure was used to create the plots shown here. The receptive fields are arranged in rows with the eccentricity of the field maxima increasing to the right, and in columns with the complexity of the fields increasing downward. All the fields in row A have single peaks; those in B a major peak and a few distinguishable minor peaks. The fields in C are the most complex. The data has been normalized so the highest peak in each field is the same height. (b) Hidden unit retinal receptive fields generated by the back-propogation model. These plots were generated in the same way as those of Fig. 2a except that the data came from computer simulations of the model network. All the fields, except for the three on the left in row C, are from units that have received 1,000 learning trials. The remaining three are from untrained units and represent fields that result from the initial random assignment of synaptic weights.

annulus is proportional to the activity in the absence of a visual stimulus. The dark inner disk is the difference between these and represents the apparent affect of visual stimulation. Figure 4 shows nine spatial gain fields from area 7a compared to nine fields from the model's hidden units. Each field consists of an array of nine circle patterns with each circle placed at a position corresponding to the location of the fixation at which it was measured. As can be seen by examination of Fig. 4, the spatial gain fields of both model and experiment show a wide variety of complex patterns.

Many of the spatial gain fields could be fitted to planes appropriately tilted in the horizontal and vertical directions. All the total gain fields (outer circles) of the model's hidden units were approximately planar. This compares with 80% for the area 7a neurons. About 55% of the visually evoked gain fields observed experimentally (inner dark disks) were also planar. The fraction of planar visually evoked hidden unit gain fields depended on which teacher format was used. With the monotonic eye-position teacher format, 78% of the visual response gain fields were planar or at least monotonic. With the Gaussian retinal teacher format, only 36% fall in this class. For both model and experiment there was also a class of visually evoked gain fields (inner dark disks) that were radically nonmonotonic (see Fig. 4, fields g, h, and i).

Comparison of the gain fields of model and experiment using both the shape of the total and visually evoked fields and the details of the circular patterns that make up each field reveal the high degree of similarity. In particular, the general form of the nonlinear combination of eye-position and visual stimulation responses are similar in the model and area 7a. In the model this nonlinear interaction is a consequence of the S-shaped input-output function used to model the neural units. The source of this nonlinearity in the brain is unknown. As in the case of retinal receptive fields, a more quantitative comparison is frustrated by the wide variety of fields observed experimentally. However, the fact that every kind of hidden unit retinal receptive field and spatial gain field can be matched quite closely to one found experimentally supports the conjecture that area 7a and the model network are computing the same function in analogous ways. These results do not justify the conclusion that back propagation plays any role in training the computation occurring in area 7a, but they do raise the possibility that some form of learning is involved.

NEW TEACHER FORMATS

When using the back-propagation modeling paradigm, the nature of the teacher is a critical factor in determining which computation is learned. Recall that the teacher provides the correct output values for each input. The teacher must therefore provide the same number of values as there are output units. In the case of our area 7a model, the teacher provided information that represented the location of a visual stimulus in head-centered coordinates. There are many possible formats for this representation. Perhaps the simplest is to use just two values, one representing the horizontal and the other representing the vertical position of the stimulus on the frontal screen. Since we are concerned with monkeys with fixed head positions, this is a way to represent head-centered location. In our original model we used two different teacher formats, an eye-position and a retinal format, to represent head-centered location. Recall the distinction between teacher format and the information represented by the teacher. Although we have mentioned several different teacher formats, they have all represented the same information.

TABLE 1

DATA FROM DIFFERENT TEACHER FORMATS

Teacher Format	% at Each Eccentricity					
	0	10	20	30	40	Mono./Irreg.
7a	0	19	25	16	40	1.22
Linear	0	8	8	8	46	11.5
Squashed	0	6	12	36	46	15.7
Squared	2	6	8	30	54	11.5
Eye	0	6	20	20	54	3.55
Retina	8	2	12	12	42	0.56
Mixed	4	6	16	28	46	1.27
Untrained	2	20	18	12	36	N/A

The teacher formats are discussed in the text. 7a refers to the distribution found in monkey cortex. Monotonic/irregular is the ratio of monotonic to irregular visually evoked eye-position gain fields, that is, the inner circles of Fig. 4.

In our previous studies we used two different formats to represent head-centered stimulus location (16). While all the gain fields produced by both formats were of types found in 7a, neither format produced the exact quantitative distribution of types found in 7a. To further analyze the effect of teacher format on the kinds of hidden units generated, we have used an additional four teacher formats that represent stimulus location in a head-centered frame. The simplest of these formats contains a small number of output units that vary their activity linearly with the horizontal or vertical position of the stimulus on the frontal screen. The activity of these units is 0.0 at one extreme of possible stimulus location and 1.0 at the other extreme. Two additional simple formats used consist of units that vary monotonically but nonlinearly with stimulus position. One varies as the square of position while the other varies as an S-shaped function of position. A logistic function of the same type used to compute the output of the units in the back-propagation network was used. The most complex new output format was a combination of the retinal and eye-position formats used in the original model. In this case the network was trained to produce both of these complex output representations of stimulus location simultaneously.

All these new versions of the teacher format could train the transformation carried out by area 7a. Examination of the retinal fields of the hidden units generated showed that all the previously observed receptive field types were present for all versions of the model. The frequency distribution of the various receptive field eccentricities varies in magnitude but not in shape between the different versions (see Table 1). While all are similar to the experimentally observed distribution in area 7a, they have somewhat less of the lower eccentricity fields. An examination of the lower eccentricity experimental receptive fields shows them to be mostly of the most complex type, row C in Fig. 2. This complex type is rare among trained model receptive fields, but is common in the model receptive fields before training, presumably because the initial weights for the model are chosen at random.

All six versions of the teacher also produced eye-position gain fields of types found in area 7a (see Table 1), but there was a considerable spread in the ratio of types. The most realistic ratio was generated by the mixed format. This was to be expected since the gain field type distribution is skewed too far one way for the monotonic format and far the other way for the Gaussian format that make up this complex type. The conclusion from these simulation studies is that as long as the network is explicitly taught the location of the stimulus, the hidden units generated are all types actually found in area 7a. Differences in the teacher format affect only the distribution of the various type.

A Teacher of a Different Kind

A reasonable question to ask is whether a teacher is really needed to train the network to have the same kinds of units found in area 7a? Some learning procedures do not need explicit external teachers to train units that can respond to significant features in the input patterns. Simple Hebbian learning, competitive learning, and a form of back propagation learning called "identity mapping" are in this category (8, 13, 15). All of these procedures exploit some form of correlation between the inputs to generate feature detector units. It seems unlikely that any of these procedures could give rise to the kind of neurons found in area 7a because the neurons respond to combinations of uncorrelated inputs, that is, eye-position and retinal stimulus location. However, examination of the hidden units in the trained network indicates that, in addition to serving as an intermediate step in a coordinate transformation, they also encode sufficient information to allow the original eye position and retinal location to be reconstructed. That is, if we first train a network to do the area 7a transformation, and then fix the hidden unit weights, it would be possible to train an output layer to regenerate the input patterns. This would give us a network that could produce input patterns on its output units after passing all information through a small number of hidden units. Networks of this kind are called identity mapping networks and have been extensively studied. Their most important feature is that the hidden units often learn an efficient encoding of the input patterns. Back-propagation identity mapping, in which the input pattern and teacher are the same, is one of the more powerful of the correlational learning procedure. If back-propagation identity mapping fails to generate hidden units of the kind found in area 7a, it seems unlikely that any other correlation-based approach will succeed.

To test this we trained a network with the back-propagation identity mapping procedure. This corresponds to using the same input and output format as used in the mixed format model described earlier, but now the teacher is just the input pattern, not a representation of the head-centered location of the stimulus.

Networks trained by back-propagation identity mapping produce hidden units that are completely different from those observed in area 7a and from those produced by using teachers that represent head-centered location. The units found with identity mapping do not combine visual and eye position inputs. Rather they separate these out so that the vast majority of hidden units were either eye position only or retinal position only. This was particularly striking because the random weight distribution used at the start of training gave rise to hidden units that combine eye and retinal position to some degree. These units had to be actively destroyed by the identity mapping process. Since we believe that back-propagation identity mapping is one of the most powerful correlation-based learning procedures, it is unlikely that other procedures based on correlation, such as simple Hebbian learning or competitive learning, will generate the kind of units observed in area 7a.

DISCUSSION

The results described here demonstrate that the back-propagation model of area 7a is robust with respect to teacher format so long as the teacher explicitly represents the location of the stimulus in space. All formats with this property gave rise to hidden units of types actually seen among the space-tuned neurons in area 7a. The details of teacher format did affect the quantitative distribution of unit types. This result suggests that the back-propagation network is actually simulating the computational algorithm used by area 7a, but says nothing about the physiological mechanisms used to implement this computation. The fact that the quantitative distribution of unit types is not exactly reproduced by the network may be a consequence of the simplicity of the model network compared to the actual neural architecture of area 7a. Indeed, considering the simplicity of the model network, the degree to which it simulates the properties of 7a neurons is quite remarkable. A key question, which can only be answered empirically, is how general is this ability of back propagation to simulate neuron response properties in other parts of the brain?

Another important issue raised by this work is the question of whether the computation accomplished by area 7a is learned. Our results do not imply that it is. It could be genetically programmed to use computational procedures analogous to those found by back-propagation learning. If, however, cortical computations are learned by an error-correcting procedure analogous to back propagation they will require teacher signals to be present in the brain. The failure of identity mapping to simulate the type of neurons present in 7a is significant in this regard because it suggests that no correlation-based learning procedure will work. This raises the question of whether other learning procedures not requiring some form of teacher or back propagation of error can be found to account for the properties of 7a neurons. We know of no such learning procedure at present, but this is an active area of research and perhaps one will be found. This is an important issue because back propagation is quite difficult to implement given our current understanding of neuron function. It is not, however, impossible to implement back propagation using plausible neurophysiological concepts as has been described elsewhere (18).

ACKNOWLEDGEMENTS

This research was supported by a grant for the System Development Foundation, Contract 85-NL-059 from the Air Force Office of Scientific Research, and Contract N00014-87-K-0671 from the Office of Naval Research to the first author and by NIH grants EY05522 and EY07492, and grants from the Sloan Foundation and the Whitaker Health Sciences Foundation to the second author.

REFERENCES

1. Andersen, R. A. The neurobiological basis of spatial cognition: role of the parietal lobe. In: Stiles-Davis, J.; Kritchevsky, M.; Bellugi, U., eds. Spatial cognition: Brain bases and development. Chicago: University of Chicago Press; in press.
2. Andersen, R. A.; Essick, G. K.; Siegel, R. M. Encoding of spatial location by posterior parietal neurons. Science 230:456–458; 1985.
3. Andersen, R. A.; Mountcastle, V. B. The influence of the angle of gaze upon the excitability of the light-sensitive neurons of the posterior parietal cortex. J. Neurosci. 3:532–548; 1983.
4. Asanuma, C.; Andersen, R. A.; Cowan, W. M. The thalamic relations of the caudal inferior parietal lobule and lateral prefrontal cortex in monkeys: Divergent cortical projections from cell clusters in the medial pulvinar nucleus. J. Comp. Neurol. 241:357–381; 1985.
5. Bock, O.; Eckmiller, R.; Andersen, R. A. Goal-directed arm movements in trained monkeys following ibotenic acid lesions in the posterior parietal cortex. Brain Res., in press; 1988.
6. Critchley, M. The parietal lobes. New York:Hafner; 1953.
7. Gnadt, J. W.; Andersen, R. A. Memory related motor planning activity in posterior parietal cortex of macaque. Exp. Brain Res., in press; 1988.
8. Linsker, R. From basic neural principles to neural architecture: Emergence of orientation selective cells. PNAS 83:8390–8394; 1987.
9. Lynch, J. C. The functional organization of posterior parietal association cortex. Behav. Brain Sci. 3:485–534; 1980.
10. Lynch, J. C.; Mountcastle, V. B.; Talbot, W. H.; Yin, T. C. T. Parietal lobe mechanisms for directed visual attention. J. Neurophysiol. 40:362–389; 1977.
11. Mountcastle, V. B.; Lynch, J. C.; Georgopoulos, A.; Skata, H.; Acuna, C. Posterior parietal association cortex of the monkey: Command function for operations within extrapersonal space. J. Neurophysiol. 38:871–908; 1975.
12. Rumelhart, D. E.; Hinton, G. E.; Williams, R. J. Learning internal representations by error propagation. In: Rumelhart, D. E.; McClelland, J. L.; The PDP Research Group, eds. Parallel distributed processing: Explorations in the microstructure of cognition. vol. 1. Foundations. Cambridge, MA: Bradford Books/MIT Press; 1986:318–362.
13. Rumelhart, D. E.; Zipser, D. Feature discovery by competitive learning. Cogn. Sci. 9:75–112; 1985.
14. Van Essen, D. C.; Maunsell, J. H. R.; Bixby, J. L. The middle temporal visual area in the macaque: Myeloarchitecture, connections, functional properties, and topographic organization. J. Comp. Neurol. 199:293–326; 1981.
15. Zipser, D. Programmine neural nets to do spatial computatins. In: Sharkey, N. E., ed. Review of cognitive science. Norwood, NJ: Ablex; in press.
16. Zipser, D.; Andersen, R. A. A back-propagation programmed network that simulates response properties of a subset of posterior neurons. Nature 331:679–684; 1988.
17. Zipser, D.; Rabin, D. P3: A parallel network simulating system. In: Rumelhart, D. E.; McClelland, J. L.; The PDP Research Group, eds. Parallel distributed processing: Explorations in the microstructure of cognition. vol. 1. Foundations. Cambridge, MA: Bradford Books/MIT Press, 1986:489–506.
18. Zipser, D.; Rumelhart, D. E. Neurobiological significance of new learning models. In: Schwartz, E., ed. Computational neuroscience. Cambridge, MA:MIT Press; in press.