

Sparse Population Coding of Faces in the Inferotemporal Cortex

Malcolm P. Young* and Shigeru Yamane

How does the brain represent objects in the world? A proportion of cells in the temporal cortex of monkeys responds specifically to objects, such as faces, but the type of coding used by these cells is not known. Population analysis of two sets of such cells showed that information is carried at the level of the population and that this information relates, in the anterior inferotemporal cortex, to the physical properties of face stimuli and, in the superior temporal polysensory area, to other aspects of the faces, such as their familiarity. There was often sufficient information in small populations of neurons to identify particular faces. These results suggest that representations of complex stimuli in the higher visual areas may take the form of a sparse population code.

An unresolved issue in cortical neurophysiology is whether the sensory hierarchies venture in small numbers of single cells tuned to complex patterns or in large populations of broadly tuned cells. Sparse coding theories suppose that individual cells would show specificity for behaviorally relevant stimuli (1), whereas population theories suppose that distributed patterns of activity in neuronal populations underlie perception and behavior and, correspondingly, expect cells to exhibit broadly graded responses (2, 3). In the context of this

issue, the specificity of neurons responsive to faces in the inferotemporal cortex (4-7) has been interpreted as strong support for sparse coding theories. On the other hand, although neurons responsive to faces may be sharply tuned to a class of stimuli, they modulate their firing to more than one stimulus and tend to be broadly tuned to stimuli within the category (5).

To address the question of the type of coding evidenced by neurons responsive to faces, we examined 850 unit recordings in the anterior inferotemporal cortex (AIT) and in the anterior superior temporal polysensory area (STP) of macaque monkeys (*Macaca fuscata*) while the monkeys performed a face discrimination task (6). The face discrimination task involved differential response to 3 of the faces from 27 other faces. The monkeys responded at greater than 90% correct performance. We analyzed

responses to the set of 27 faces that met two selection criteria (Fig. 1). The cells were divided into two groups according to whether they were recorded from AIT (41 cells, 26 from monkey A and 15 from monkey B) or in the STP (30 cells, 25 from monkey A and 5 from monkey B). This sample of cells represented 8% of the total number of recorded cells.

To represent quantitatively the population responses to the faces, we applied multidimensional scaling (MDS) to the two populations, a technique that has been used for qualitative analysis of population encoding of complex stimuli (7). MDS produces a configuration of points in a small number of dimensions that represent the population responses to the faces. The distances between the points of the MDS configuration are as close as possible to the Euclidean distances between points corresponding to each face response in a high dimensional space whose dimensions are defined by the cells (8). Two-dimensional configurations were derived (Fig. 1), and these explained 70% and 75% of the variance in the AIT and STP data, respectively. These variance-explained statistics were surprisingly high given the sharp dimensional reduction from tens of dimensions (the number of cells in the two populations analyzed) to only two dimensions, and this suggested that the coding of faces in the two populations was redundant. The configurations of points that represented the face stimuli for the AIT and STP populations are shown in Fig. 1, A and B. Faces plotted close together evoked a similar pattern of response across the population, whereas faces plotted far apart evoked very different population responses.

We investigated what characteristics of the faces the AIT and STP cells were coding, en masse. To do this we exploited the fact that the physical properties of the set of faces on which the monkeys performed the face discrimination task had been extensively quantified (Fig. 2A) (6). In addition to the 29 variables that quantified distances between the facial elements, we derived two further models. The first was a "general physical similarity" model, which was computed from the face measurement variables by MDS. A two-dimensional solution (Fig. 2B) explained 94% of the variability in the face measurements, reflecting their redundancy (all the width measurements tend to covary with the general width of the face, for example). The second additional model represented the "familiarity" of the faces, some of which belonged to humans known in varying degrees to the monkeys. Faces of people unknown to the monkeys were coded with a 1, those of people occasionally seen by the monkeys with a 2, those of people more

* P. Young, University Laboratory of Physiology, Parks Road, Oxford, OX1 3PT, United Kingdom, and Laboratory for Neural Information Processing, Frontier Research Program, The RIKEN Institute, Wako-shi, Saitama, Japan.

Shigeru Yamane, Division of Neuroscience, Electrotechnical Laboratory, Tsukuba-shi, Ibaraki-ken, 305, Japan.

† To whom correspondence should be addressed.

Take conclusion from data

often seen but not involved in the experiment with a 3, those of people who participated in the experiment with a 4, and the face of the person most familiar to the monkeys with a 5.


We quantitatively compared the population responses with the models that represented the properties of the faces. We made the comparisons using PROCRUSTES rotation (9, 10), which finds the optimal reflection, rotation, and scaling of each face distance variable (H1 and E2, for example) with each population configuration. At the optimal comparison, the procedure yields a statistic, which reflects the goodness-of-fit between the two com-

pared models. The statistical rarity of each comparison was assessed by an approximate randomization test (11), which repeated the PROCRUSTES rotation with one of the models shuffled randomly on each of 600 iterations. For the AIT population, the general physical similarity model was significantly related to the population response model ($r = 0.36, P < 0.05$), so the responses of these cells, en masse, were more similar when the faces were more physically similar. In addition, measurement variables encoding the relations between the eyes and the hairline (H3, H4, H6, and H7; each $r > 0.35$, each $P < 0.05$) were significantly related. No other models were

significantly related to the population response model. Hence, these exploratory analyses suggest that the AIT population may have been coding the general physical properties of the presented faces, with particular emphasis on the upper part of the face.

For the STP population, W1 was significantly related to the population response model ($r = 0.31, P < 0.05$). The variance explained statistic for this variable, however, was lower than that for any of the variables that were related to the AIT population model, and it is not, therefore, very persuasive that the STP population was coding this particular variable. Th-

Fig. 1. Diagrams of the population responses to faces. Single unit activity was recorded during exposure to photographs of the heads of Japanese men in full face. These faces were presented on a screen for 600 ms, and responses were quantified as the difference between the mean firing rate in the prestimulus baseline and that in the 500-ms period beginning 100 ms after stimulus onset. Selection of responses was made according to two criteria. The first criterion was that the firing rate of each cell to any face stimulus was significantly different from that during the prestimulus baseline, and the second was that each cell's responses to the faces showed a standard deviation greater than five spikes per second, to include only cells that clearly discriminated the face stimuli. Proximity data suitable for MDS were derived from each matrix, in which cells were represented by columns and the response to each face by rows, by computing Euclidean distances between the face responses (10). Nonmetric MDS was applied to the proximity data and produced configurations in one to six dimensions so that analyses in different dimensions could be compared in a "scree" test (10, 13). The scree test revealed that two-dimensional solutions were satisfactory for both the AIT and STP populations. (A) MDS-derived configuration representing the population response to the face stimuli for the AIT cells. Faces that evoked a similar pattern of response across the population are close together, whereas faces that evoked a very different population response are far apart. Example faces are shown at the position of the point corresponding to the population response to that face in the diagram or, for faces whose population response was represented near the center, at the side of the diagram. This configuration, derived from the physiological selectivities of cells in AIT, shows a configuration of faces similar to the general physical similarity model in Fig. 2B. (B) MDS-derived configuration representing the population response to the face stimuli for the STP cells. Faces belonging to people known to the monkeys are shown. This configuration, derived from the physiological selectivities of a population of cells in STP, shows most of the known faces distributed toward the bottom right, with the face of the person most closely associated with the monkeys at the furthest point of this tangential dimension.




et two
s were
ng to
T (41
from
ls, 25
y B).
of the

pula-
plied
e two
used
ncod-
ces a
er of
ation
s be-
ation
dean
ng to
ional
y the
tions
ained
AIT
var-
ingly
tion
er of
l) to
sted
pula-
ns of
i for
n in
ech-
onse
cted
tion

s of
were
ited
the
per-
had
6).
nti-
ts,
irst
lel,
ea-
en-
of
ts,
lth
en-
he
he
ch
le-
n-
l,
he
re

Angle



familiarity model, however, despite a 3-year period of exposure during which all the faces are likely to have become familiar, was also significantly related to the STP population model but was associated with a higher variance-explained statistic ($r = 0.36$, $P < 0.025$). This correspondence is illustrated in Fig. 1B, in which the faces of people known to the monkeys are shown. There is a greater density of known faces toward the bottom right of the figure, and the face of the person most closely concerned with the monkeys in these experiments is represented at the extremity of this dimension of the STP population model (Fig. 1B, bottom right). We cannot rule out, however, the possibility that this dimension could have been related to the perception of the elements of a human hierarchy, because the person at the extremity of this dimension was also the head of the laboratory in which the experiments were undertaken. In the context of the coding of familiarity or social status though, macaques are capable both of identifying photographs of familiar individuals and of recalling the social relations of the individuals in the photographs (12).

These results suggest that neurons responsive to faces share dimensions of specificity and that these shared dimensions correspond, in the case of AIT cells, to physical properties of the faces and, in the

case of STP cells, to some other possibly social properties of the faces. We addressed the further question of whether, in these shared dimensions of specificity, the responses of the neurons were sufficiently systematic to allow a readable population code that could identify particular faces.

If the population code were readable, the responses of cells should be systematically graded across stimuli, that is, responses to suboptimal stimuli should not be noise. To investigate whether these neurons responsive to faces responded in a systematically graded way, we used an analysis developed by Georgopoulos *et al.* (13). This analysis has been applied to cells in motor cortex, which are broadly tuned to movement direction, in such a way as to suggest that their broad tuning is a reflection of a systematically graded preference for movement direction, and the degree to which population vectors derived by these means recover the appropriate directions is a useful measure of the systematic nature of the responses of broadly tuned cells. Our hypotheses in applying this test to the AIT and STP populations of neurons responsive to faces were that, if the cells were better understood as elements from a population coding mechanism, then population vectors derived from the assumption of systematically graded responses would tend to recover their appropriate stimulus vectors

and that, if they were grandmother cells signaling only one face and responding randomly to others, then the population vectors would not.

We began by calculating the angle of each face stimulus in each of the AIT and STP configurations using north as 0° and proceeding clockwise. The responses of each cell were plotted against the angle of each stimulus in the derived space and a sine function fitted to the data. Single cell vectors for each stimulus were derived from the peak angle of the sine function for each cell (for direction) and from the response of the cell to each stimulus (for length), according to the methodology of Georgopoulos *et al.* (13). A population vector for each stimulus was then computed by summation. Figure 3, A and B, shows plots of stimulus vectors, single cell vectors, and population vectors for arbitrarily chosen stimuli. The direction of the population vectors was, in general, close to the direction of the stimulus vectors.

Figure 4 illustrates the distribution of angular discrepancies between the population vectors and their corresponding stimulus vectors. The distribution is skewed toward the left, indicating that discrepancies were typically small. Under the assumption that the cells' responses were not systematic to suboptimal stimuli, it would be expected that the distribution of angular





Fig. 2. Physical properties of the face stimuli. (A) The presented faces were extensively quantified according to the measurement variables shown here. The measurements of each face were used to identify what elements in the configuration of faces were coded at the single cell level (5) and, in the present analyses, at the level of the neuronal population. (B) The configuration of the 27 faces based on their "general physical similarity." We derived this configuration from the measurements of the faces made according to the variables noted in (A) by first computing a proximity matrix by taking Euclidean distances between the faces across the face measurement variables, and then submitting this proximity matrix to MDS. The two dimensions of this configuration preserved 94% of the variability in the original 29 dimensions, suggesting that



the measurement variables tend to covary and are consequently redundant. A selection of faces were plotted at the position of their corresponding point on the diagram to illustrate that the dimensions of the configuration appear to correspond to the amount of hair possessed by the face (left to right) and an elongated to round face dimension (top to bottom). Hence, the model represents round faces with a lot of hair at bottom right, round faces with less hair at bottom left, and long faces toward the top of the figure.

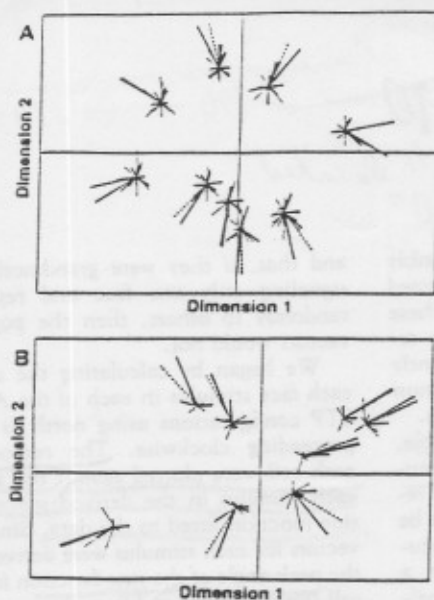


Fig. 3. Population vector plots, (A) from the AIT data and (B) from the STP data, showing single cell vectors (light lines), population vectors (heavy lines), and stimulus vectors (dashed lines). Each set of vectors is plotted at the position of its corresponding face stimulus in the space. Both population vectors and stimulus vectors have been arbitrarily elongated to show more clearly the angular discrepancies between them. The cross-hairs pass through the centroid of each space. Despite the small samples of cells, the population vectors typically matched their corresponding stimulus vectors closely, a fact which suggests both that population encoding is present and that the population code is sparse.

errors would be flat, indicative of randomly pointing population vectors. The derived distribution was clearly not flat, and this suggests that the responses of these neurons responsive to faces were systematically graded across the face stimuli. Because each cell would therefore participate in a principled way in the representation of more than one face stimulus, these results suggest that these AIT and STP neurons were engaged in the population processing of faces.

However, the strong skew to the left of the distribution in Fig. 4 suggests that these populations approached the acuity necessary to identify individual faces. This, taken with the fact that only a few tens of cells were included in the analyzed populations, suggests that only slightly larger populations would be able to identify particular faces and that consequently the population code evidenced by these cells was sparse.

The type of population code derived above may be a useful way to approach the question of whether responses to suboptimal stimuli are noise or are systematic and to represent activity distributed over en-

sembles of neurons. It may also be useful for identifying what new stimuli are being coded by a particular pattern of activity over the neurons because the physical or psychological dimensions to which the population response corresponds were identified. Hence, inferences, derived from a pattern of response, about the likely characteristics of a face stimulus not included in the faces analyzed here are possible for the AIT population if it is assumed that general physical similarity was being coded, and for the STP population if it is assumed that familiarity was being coded. Tests of predictions of this nature with new stimuli may be a useful way of investigating the robustness of a population code and the likelihood that the nature of the physical or psychological factors encoded by a population has been correctly identified.

These results suggest, for three reasons, that neurons responsive to faces in the inferotemporal cortex evidence a sparse population code. First, the redundancy of the coding characteristics of the neurons, which was illustrated by the fact that only two dimensions explained most of the variance in both populations, suggests that individual cells were not responding independently but exhibited shared dimensions of specificity. It was evident from the breadth-of-tuning of the cells that it was not the case that each cell responded to only a limited range of these shared dimensions, as would be required for a punctate code (3). Second, population responses exhibited a statistically significant relation to identifiable dimensions of the face stimuli. This implied that there was information encoded at the population level, which again would not be the expected result if all information transfer were at the level of the

single cell. Third, neurons responsive to faces exhibited systematically graded responses with respect to the face stimuli. Hence, each cell would systematically participate in the representation of many faces, which straightforwardly implies a population code. We think it a sparse population code because only a few tens of cells generated so precise a code that a relatively small increase in population size would likely be sufficient to generate a code as precise as behavior.




Fig. 4. Distribution of angular discrepancies between population vectors and their corresponding stimulus vectors. A strong assumption in the population vector procedure is that the graded nature of the responses of each cell with respect to direction is sufficiently systematic. A single parameter, the cell's direction preference, represents the "tuning curve" of the cell. This will only be a good descriptor if the cell's responses to other stimuli are well fitted by the sine function. In the ideal case, where the responses are perfectly systematically graded according to a sine function, the summated population vector would perfectly represent the summation of the cell's preferences and would point with no error in the same direction as the associated face. In the least ideal case, where the cell's responses to suboptimal faces are entirely random noise, the cell's direction preference would be a good estimate only for the optimal face for the cell and would be very bad for suboptimal faces. In this case the population vector would summate a few "good" descriptors and a large number of random ones and consequently point in random directions for each face. The distribution derived here is markedly skewed to the left for both AIT and STP cells, in contrast to the flat distribution expected if responses to suboptimal stimuli were not systematic.

REFERENCES AND NOTES

- H. B. Barlow, *Perception* 1, 371 (1972); *Q. J. Exp. Psychol.* 37A, 121 (1985); J. Konorski, *Integrative Activity of the Brain: An Interdisciplinary Approach* (Univ. of Chicago Press, Chicago, 1967).
- W. J. Freeman, *Mass Action in the Nervous System* (Academic Press, New York, 1975); A. P. Georgopoulos et al., in *Dynamic Aspects of Neocortical Function*, G. M. Edelman, W. E. Gall, W. M. Cowan, Eds. (Wiley, New York, 1984).
- S. R. Lahnky and T. J. Sejnowski, *J. Neurosci.* 10, 2281 (1990).
- D. I. Perrett, E. T. Rolls, W. Caan, *Exp. Brain Res.* 47, 329 (1982); D. I. Perrett et al., *Hum. Neurobiol.* 3, 197 (1984); D. I. Perrett, A. J. Mistlin, A. J. Chitty, *Trends Neurosci.* 10, 358 (1987).
- G. C. Baylis, E. T. Rolls, C. M. Leonard, *Brain Res.* 342, 91 (1985).
- S. Yamane, S. Kaji, K. Kawano, *Exp. Brain Res.* 73, 209 (1988). MDS has also been used to study the psychological dimensions of faces (G. Rhodes, *Perception* 17, 43 (1988)).
- M. E. Hasselmo, E. T. Rolls, G. C. Baylis, *Behav. Brain Res.* 32, 203 (1989).
- R. N. Shepard, *Science* 210, 390 (1980).
- P. Schonemann and R. M. Carroll, *Psychometrika* 35, 245 (1970); J. C. Gower, in *Mathematics in the Archeological and Historical Sciences* (Edinburgh Univ. Press, Edinburgh, 1971), pp. 138-149.
- M. P. Young, thesis, University of St. Andrews, St. Andrews, Fife, Scotland (1990).
- E. S. Edgington, *Randomization Tests* (Dekker, New York, 1980).
- V. Dasser, in *Machiavellian Intelligence. Social Expertise, and the Evolution of Neglect*, R. Byrne and A. Whiten, Eds. (Oxford Univ. Press, Lon-

- don, 1988), pp. 85-93.
- A. P. Georgopoulos, J. F. Kalaska, R. Caminiti, J. T. Massey, *J. Neurosci.* 2, 1527 (1982); A. P. Georgopoulos, A. B. Schwartz, R. E. Kattner, *Science* 233, 1416 (1986); A. P. Georgopoulos, J. T. Lurito, M. Petrides, A. B. Schwartz, J. T. Massey, *Science* 243, 234 (1989).
 - We are grateful for the support of a Royal Society STA research fellowship to M.P.Y., the Frontier Re-

search Program, the RIKEN Institute during 1990, and for the support of an Agency of Industrial Science and Technology visiting research fellowship to M.P.Y., the U.K. Medical Research Council, and the Oxford McDonnell-Pew Centre for Cognitive Neuroscience during 1991. We thank C. Blakemore, K. Tanaka, and S. J. Judge for comments.