



ELSEVIER

BioSystems 40 (1997) 189–196



# On the perceptual structure of face space

Alessandro Treves\*

*SISSA — Cognitive Neuroscience, via Beirut 2-4, 34013 Trieste, Italy*

---

## Abstract

We are able to recognize very many different faces of individuals we know, apparently using a complex and ill-understood set of identifying features; it seems natural to assume that faces are perceived as spanning the equivalent of a high-dimensional vector space. I explore ways to probe the structure of perceptual face space without making a priori hypotheses about either the space itself or the mechanisms of perception and recognition, and using solely neuronal responses recorded in the monkey, and metrics derived from their mutual similarities. Within this approach, the dimensionality of face space remains an elusive concept, but the metric content and ultrametric content of the face sets used can be quantified and compared with those of other perceptual sets.

*Keywords:* Face cells; Place cells; Information theory; Decoding; Quasi-distances; Ultrametricity

---

## 1. How do we recognize faces?

Recognizing faces is clearly a most important ability and one with high survival value, for humans and for other social species, and it is not surprising that a great deal of work has been devoted to dissecting the underlying processes, from the points of view of experimental psychology, neuropsychology and visual neurophysiology (see e.g. the papers collected in the *Phil. Trans. R. Soc. Lond. B* 335, 1992). As an information processing task, face recognition is paradigmatic of the sort of tasks at which natural nervous systems are vastly superior to artificial intelli-

gence systems, especially when considering the many sources of variability in visual inputs (from illumination to translation, scaling and rotation of the face, down to specific aspects such as haircut or removal of spectacles) with respect to which everyday recognition is relatively invariant. It is widely accepted that faces can be described as complex combinations of features, and Rolls (1992) has suggested in particular that individual cells along the primate ventral visual system might be seen as feature analyzers, tuned to low-dimensional combinations of simpler features from the preceding area, in a stepwise escalation in complexity that would culminate in specialized populations (Hasselmo et al., 1989) for face identification, the decoding of face expression, and so on.

Which are, then, the features in terms of which neurons at a given processing stage encode faces?

---

\* Corresponding author, Tel: +39 40 3787426; fax: +39 40 3787528; e-mail: ale@limbo.sissa.it

Yamane et al. (1988) have considered a set of tens of metric descriptors, including, e.g. the distance between the eyes or the width of the mouth, following a strict geometrical logic similar to that which has inspired O'Keefe (1996) and collaborators to find, among rat hippocampal place cells, neurons that code for the distance between the animal and the walls, along Cartesian coordinates. If several of these descriptors were non-redundant, one would expect the responses of a sufficiently large population of cells from the same cortical area to span a correspondingly high-dimensional space, when the animal (a monkey) is presented with face stimuli varying in terms of the descriptors. In contrast, Young and Yamane (1992), following Hasselmo et al. (1989), have applied multidimensional scaling to the responses of tens of cells from the anterior inferotemporal cortex and the anterior superior temporal polysensory area of two monkeys, and have found that 70–75% of the variance in the responses is along two dimensions alone.

A self-organizing network, in which neurons acquire their response selectivity as a result of an ongoing unsupervised learning process, can be expected to display single cell response profiles that in general do not match any simple geometrical analysis, but rather map what appear to be random regions in any given coordinate frame. Yet, the dimensionality spanned by the responses would remain a valid index of the richness of the stimulus set, as perceived by the population of cells being recorded. Faces would seem, to our intuition, to constitute a potentially rather rich set. Sirovich and Kirby (1987) have applied principal component analysis to a large database of human face pixel images, and found a Karhunen-Loeve dimension (the number of eigenvalues at least 1% of the largest) of 21; the ten largest components accounted for 82% of the variance. Although they emphasize how considering principal components lowers the effective dimensionality of the set, 10 or 20 is still a large number compared to 1 or 2, particularly since the database used included only Caucasian males with no facial hair nor spectacles (the first principal component within this rather homogeneous set, which alone accounted for nearly 40% of the variance, was

well correlated with the presence or not of hair hanging on the forehead).

The operations required of the networks involved in face recognition depend strongly on the perceived structure of face space and its effective dimensionality (perceived implies discarding extra aspects that may be there physically, but are irrelevant to perception). Just imagine the two extreme cases in which  $n$  faces, being in generic positions in a very high-dimensional space, span a simplex of dimension  $d = n - 1$ ; and in which they, instead, span a single dimension, and very high resolution  $\rho \sim 1/n$  is required along this dimension to discriminate between individuals. Infinite variations of intermediate cases are of course possible, for example a set of faces might be chosen at the vertices of a  $d$ -dimensional hypercube,  $d = \log_2(n)$ , and then  $d$ -binary neurons would conveniently identify each face by the presence or absence of any combination of features. It is therefore important to try and supplement the above analyses of physical face space with studies of how this space is represented in the brain.

Several difficulties are apparent already at this stage. First, it is difficult to extract a meaningful measure of effective dimensionality from principal component analyses (or multi-dimensional scaling), as typically the importance of successive dimensions (as measured, e.g. by the corresponding eigenvalues of the covariance matrix) decreases exponentially, and it is quite arbitrary where to set the significance level. A second difficulty, which arises when considering realistic datasets with multiple examples of each face, and is exacerbated when such analyses are applied to neuronal responses to faces, is in the rather large variability of such responses, even when exactly the same face stimulus is presented repeatedly. Moreover responses to some faces may be much more variable than to others, and in some directions more than in others. To discuss the metric of face space, and in particular its effective dimension, which are relevant to neuronal processing, one must take this variability into account, and cannot rely solely on mean responses to each stimulus to define a canonical metric. Third, and most important, analyses that rely *ab initio* on

Euclidean spaces, and reduce essentially to performing appropriate rotations of the coordinate frame, will never be able to say anything about possible non-Euclidean aspects of the stimulus set as perceived by the brain.

I consider here alternative ways to study the structure of neuronal face space, which aim at quantifying aspects beyond dimensionality, which take into account the real variability, and which do not make any a priori hypothesis as to whether the space is Euclidean. As will be clear below, this is still a preliminary formulation, to be improved in some drawbacks such as the introduction of a not mathematically valid notion of distance, and to be applied more extensively to other datasets.

## 2. Decoding the responses of face cells

Decoding the spike trains emitted by a population of neurons, when one ( $s$ ) of a set of stimuli is presented, means applying an algorithm that estimates, given the current spike trains  $\vec{r}_s$  and those previously recorded in response to each stimulus, the likelihoods for each ( $s'$ ) of the possible stimuli to be the current one,  $L(s'|\vec{r}_s)$ . The stimulus  $s' = s_p$  for which this likelihood is maximal can be said to be the stimulus predicted on the basis of the response. In general  $s_p$  will not coincide with the true  $s$  and the accuracy in the decoding can be gauged by the percent correct decoding (or the corresponding fraction  $f_{\text{cor}}$ ), or alternatively by the mutual information in the joint probability table  $Q(s, s_p)$ ,

$$I_{ml} = \sum_{s, s_p \in S} Q(s, s_p) \log_2 \frac{Q(s, s_p)}{P(s)Q(s_p)}. \quad (1)$$

This quantity measures the information in the predictions based on Maximum Likelihood, and as such it does not only reflect, like percent correct, the number of times the decoding is exact, but also, beyond percent correct, the distribution of wrong decodings. A further quantity is the mutual information

$$I_p = \sum_{s, s' \in S} P(s, s') \log_2 \frac{P(s, s')}{P(s)P(s')} \quad (2)$$

obtained from the probability  $P(s'|s)$  of confusing  $s$  with  $s'$ , which is given by averaging  $L(s'|\vec{r}_s)$  over the responses to  $s$ . This second information measure reflects, unlike the first, also the degree of certainty with which each single trial has been decoded, and it thus sheds light on a further aspect of the quality attained in decoding. Both information quantities suffer from limited sampling distortions (Treves and Panzeri, 1995; Panzeri and Treves, 1996) but the second much less than the first, in the sense that, with the limited sampling correction procedures we have developed,  $I_p$  can be estimated accurately even with few trials per stimulus, while  $I_{ml}$  requires more trials.

Decoding algorithms can be optimised to extract as much information as possible, or can be modelled on the decoding likely to be implemented by real neurons downstream of the recorded populations. Information and percent correct values in the decoding of face cells responses from the primate temporal visual cortex are reported by Rolls et al. (1996). There we show that simple, neuronally plausible decoding algorithms, based on dot product operations, perform virtually like optimal decoding algorithms in terms of  $I_{ml}$ , and are only 20–30% inferior in terms of  $I_p$ . This is because the simple dot product algorithms are poorer at quantifying likelihoods, even if they order them correctly and identify correctly the most likely stimulus that can be predicted for each trial. We emphasize measures that depend on  $P(s, s')$ , the most complete characterization of what can be extracted from the responses, and hence in the following I insist more on  $I_p$  derived with optimal decoding algorithms, and on similarity measures among stimuli that again depend on  $P(s, s')$ .

$f_{\text{cor}}$ ,  $I_{ml}$  and  $I_p$  all depend on the number of cells in the population, as recording the responses of more cells obviously allows better decoding. We have reported the important result (Rolls et al, 1996) that the information decoded from face cells appears to grow linearly with the number of

cells in the population, until it begins to saturate at the maximum allowed, which is just the entropy of the stimulus set,  $H = -\sum_s P(s) \log_2 P(s)$ . This result implies that the different cells in the sample tend to code for different aspects of the stimulus set, so that each contributes an additive term to the information provided by the population. This result appears to hold also for the data recorded in a number of other experiments, including primate hippocampal cells coding for space (Rolls et al., 1995), primate orbitofrontal cells coding for odours (Rolls, Treves and Critchley, in preparation), rat hippocampal cells coding for spatial position (Wilson and McNaughton, 1993; Treves et al., 1996), and rat somatosensory cells coding for nociceptive stimuli (Panzeri et al., 1995).

The issue I want to consider here is not, however, how the accuracy in the decoding depends on the number of cells in the population, but rather how it provides insight on the perceived structure of the stimulus set.

One way to approach this issue is to contrast the values obtained, e.g. with optimal decoding algorithms, for the three quantities mentioned above, which, again, capture increasingly detailed aspects of decoding performance. An example of this analysis has been provided in Rolls et al. (1996), where a graph shows the values of  $I_p$  plotted against the values of  $f_{\text{cor}}$  obtained with the same number of cells in the population. With  $S = 20$  stimuli in the set, presented with equal frequency (ten trials for each stimulus were used in this dataset),  $f_{\text{cor}}$  range from 0.05 ( $= 1/20$ ) to 1.0, and  $I_p$  (or  $I_{ml}$ ) from 0.0 to 4.32 ( $= \log_2(20)$ ). Datapoints cannot span these entire ranges independently, however, as there are mutual constraints between  $f_{\text{cor}}$  and  $I_p$ ,  $I_{ml}$ . The lowest information values compatible with a given  $f_{\text{cor}}$  are those attained when equal probabilities (or equal frequencies of being selected as most likely in the case of  $I_{ml}$ ) result for all wrong stimuli. In this case one finds

$$I_{\min} = \log_2 S + f_{\text{cor}} \log_2 f_{\text{cor}} + (1 - f_{\text{cor}}) \log_2 (1 - f_{\text{cor}}) - (1 - f_{\text{cor}}) \log_2 (S - 1). \quad (3)$$

Conversely, maximum information for a given  $f_{\text{cor}}$  is when stimuli are perceived as being categorized into classes of size  $1/f_{\text{cor}}$ , and on each trial the class is identified correctly, but the individual stimulus within the class is decoded on a purely random basis. It is easy to see that then

$$I_{\max} = \log_2 S + \log_2 f_{\text{cor}}. \quad (4)$$

Interpreting the probability of wrong decodings as a monotonically decreasing function of some underlying distance (e.g. as discussed below), the first situation can be taken to correspond to the limit in which the stimuli form an equilateral simplex, or equivalently the stimulus set is drawn from a space of extremely high dimensionality. In the Euclidean  $d \rightarrow \infty$  limit, points drawn at random from a finite, e.g. hyperspherical region tend to be all at the same distance from each other, and from the point of view of the metric of the set this is referred to as the trivial limit. The second situation can be taken to correspond to the ultrametric limit, instead, in which all stimuli at distance less than a critical value from each other form clusters such that all distances between members of different classes are above the critical value. This is a non-Euclidean structure (although it could be embedded in a Euclidean space of sufficiently large dimension), and it is a first example of the possible emergence of non-Euclidean aspects from a quantitative analysis that does not rely on a priori assumptions.

Intermediate situations between the two extremes are easy to imagine and can be parametrized in a number of different ways. A convenient parameter that simply quantifies the relative amount of information in excess of the minimum, without having to assume any specific distribution of wrong decodings, is

$$\lambda_m = \frac{I - I_{\min}}{I_{\max} - I_{\min}} \quad (5)$$

which range from 0 to 1 and can be interpreted as measuring the metric content of the perceived set. What is quantified by  $\lambda_m$  can be called the metric content not in the sense that it requires the introduction of a real metric, but simply be-

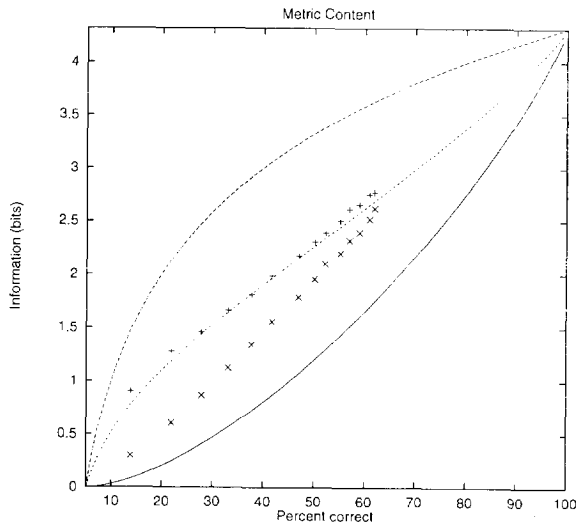


Fig. 1. Information vs. percent correct in decoding the responses of populations of up to 14 cells to 20 face stimuli. Data points are for different population sizes for  $I_{ml}$  (+) and  $I_p$  (x), and are in the region between a non-metric lower bound (full curve, Eq. (3)) and a fully categorised upper bound (long dashes, Eq. (4)). The intermediate curve (short dashes) is for metric content  $\lambda_m = 0.5$ .

cause it gives the degree to which relationships of being close or different (distant), among stimuli, are relevant to their perception. For  $\lambda_m = 0$  such relationships are irrelevant, to the point that if the correct stimulus is missed in the decoding of the response, all others have equal probabilities. For  $\lambda_m = 1$  close stimuli are so similar as to be fully confused with the correct one in the decoding, whereas other stimuli are never mistaken for the correct stimulus.

Fig. 1 reports a similar graph to that of Rolls et al. (1996) (the data is the same but the slightly different decoding algorithms produce minor differences in the experimental values).

One can see that metric content is not an invariant characterization of the responses, valid for all percent correct values (which vary in relation to the number of cells considered). Nevertheless,  $I_{ml}$  is at or above the curve for  $\lambda_m = 0.5$ , and it appears to tend to slightly higher values when more cells allow for more accurate decoding.  $I_p$  would seem to indicate very low metric content for small populations, while it tends to reach  $I_{ml}$  for larger populations. In fact one can show that

this behaviour for low percent correct (in particular, an initially quadratic dependence of  $I_p$  on  $f_{cor} - f^0$ ) derives purely from using a decoding algorithm that, when the information in the responses is low, assigns nearly equal probabilities to all stimuli (this does not happen when considering only maximum likelihoods, as even small differences in probability can result in a nonrandom ordering of likelihoods, and hence in higher values for  $I_{ml}$ ).

Although the fact that  $I_p$ , which is the more telling and safer-to-measure quantity, indicates a strongly-varying metric content, may take away significance from this index, there is something which appears very significant that can be extracted from this particular plot. It is the fact that data from other experiments, in which not faces but spatial correlates were considered as 'stimuli' (spatial position in the rat, spatial view in the monkey), appear to lie, in this plot, considerably below the face data, for any percent correct value. This is true, in particular, when considering sets with an equal number (20) of elements as the face set. Therefore, metric content may well vary with  $f_{cor}$  (strongly in the case of  $I_p$ , mildly in that of  $I_{ml}$ ), but it is likely to be significant in characterizing the representation of different correlates in the brain, and the preliminary analyses we have carried out indicate much more metric content for face representations than for space representations (results for space hippocampal data from the rat and the monkey will be reported elsewhere).

### 3. From confusion to similarity to a quasi-distance

A quantification of the metric content does not require, as we have seen, the introduction of a metric, but on the other hand it hints at another aspect in the structure of the set of face stimuli, that itself does require a metric. This is because full metricity ( $\lambda_m = 1$ ) could in particular be attained by ultrametric structures, although not necessarily (it could also be the result of the set being embedded in a standard Euclidean space, with the probability of decoding a  $\Theta$ -function differing from zero only in a region around each element of the set). Quantifying not metricity but

ultrametricity requires considering distances between triplets of points and hence the introduction of a distance among stimuli of the set.

I want to use a notion of distance among stimuli that is based on the degree of similarity of the responses and takes into account their variability, but does not assume a Euclidean structure for the underlying stimulus space. These requirements can be satisfied by first considering a symmetrized function of the probability  $P(s'|s)$  of confusing  $s$  with  $s'$ , which effectively quantifies their similarity:  $P(s'|s)P(s|s')$ ; and then normalizing it into a quasi-distance

$$D(s, s') = -\log \frac{P(s'|s)P(s|s')}{P(s|s)P(s'|s')} \quad (6)$$

The measure  $D(s, s')$  is only a quasi-distance because  $D(s, s) = 0$  and  $D(s, s') = D(s', s)$ , but the triangular inequality is not necessarily satisfied. To satisfy it,  $D(s, s') + D(s', s'') \geq D(s, s'')$ , the probabilities of confusion must verify

$$\begin{aligned} &P(s''|s)P(s|s'')P(s'|s')P(s'|s') \\ &\geq P(s''|s')P(s'|s)P(s|s')P(s'|s'') \end{aligned} \quad (7)$$

which is indeed an inequality that has an attractive logical interpretation (confusing  $s$  with  $s''$  occurs at least as often as predicted by the frequency with which both are confused with  $s'$ ), but which is not assured to hold.

For any given set of points, any notion of quasi-distance among them that does not satisfy the triangular inequality can be made to satisfy it by trivializing the measure, i.e. raising it to a sufficiently small power  $1/d^*$ :  $D \rightarrow D^{1/d^*}$ . For our particular dataset of responses to face stimuli, however, it appears that the violation of the inequality is not that serious to warrant the brutish method of trivialization. This is clear from Fig. 2, in which different points correspond to different triplets from the  $S(S-1)(S-2)/6$  possible ones (discarding those in which either  $P(s|s') = 0$  or  $D(s, s') \leq 0$  for some pair  $s \neq s'$ ). In this particular plot, which shows the ratios of the minimum and intermediate distances among those in the triplet to the maximal one ( $D_{\max}$ ), triplets that satisfy the triangular inequality are above the line




Fig. 2. Distribution of quasidistances among triplets of faces, as extracted from neuronal recordings by Rolls and Tovee, in the way specified in the text.

$D_{\min}/D_{\max} = 1 - D_{\text{med}}/D_{\max}$ , while trivialization of the metric would push all triplets towards the trivial limit  $D_{\min}/D_{\max} = D_{\text{med}}/D_{\max} = 1$ . In a fully ultrametric structure all triplets would satisfy the equality  $D_{\text{med}} = D_{\max}$ , that is all points would lie on the vertical line at the right.

For a useful comparison, one can compare with the distribution of triplets for points drawn at random from Euclidean spaces of dimension  $d$ . Since such distribution depends somewhat on the exact way in which the points are selected, this is specified here as first selecting two points really at random (with any measure of choice), and then selecting the third with a flat measure among all points whose distances from the first two are less or equal than the distance between the first two ( $D_{\max}$ ). One can easily calculate the probability distribution of triplets on the graph as

$$\begin{aligned} &P(\delta_1 = D_{\min}/D_{\max}, \\ &\delta_2 = D_{\text{med}}/D_{\max}) \propto \delta_1 \delta_2 y^{d-3}, \end{aligned} \quad (8)$$

with  $d \geq 2$  and  $y^2 = (\delta_1^2 + \delta_2^2 + \delta_1^2 \delta_2^2)/2 - (1 + \delta_1^4 + \delta_2^4)/4$  (for  $d = 1$  all points are on the line  $\delta_1 = 1 - \delta_2$ ). Two examples are given in Fig. 3. As  $d \rightarrow \infty$  all triplets converge to the trivial limit.

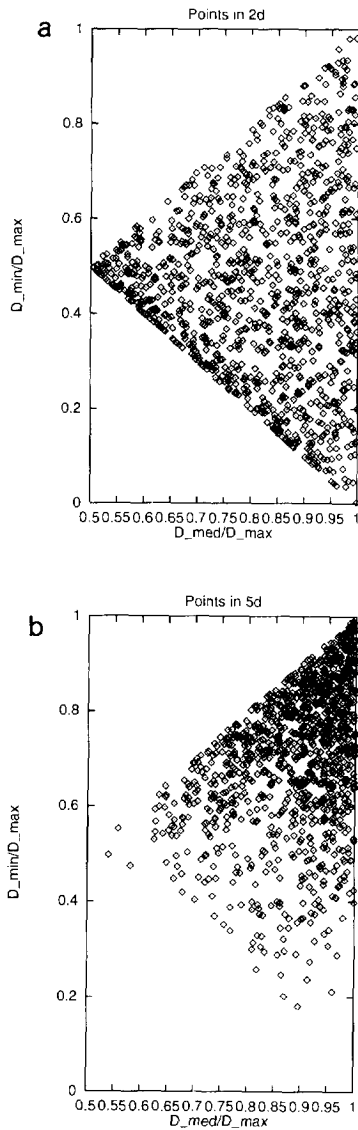


Fig. 3. Distribution of distances among triplets of points drawn at random from Euclidean spaces of dimension 2 (above) and 5 (below), as specified in the text.

Two main quantities can be extracted from each given distribution  $\langle \dots \rangle$  of triplets. The first is a measure of the dimensionality of the underlying space, given, e.g. by  $\bar{d}^2 \sim -1/\log \langle (\delta_1^2 + \delta_2^2)/2 - (\delta_1^2 \log \delta_1 + \delta_2^2 \log \delta_2) \rangle$ . For any given set of points, trivializing the metric implies  $\bar{d} \rightarrow \infty$ , and since it is a priori arbitrary to what extent to trivialize, the effective dimensionality of

the set remains ill-defined in this approach, that does not force constraints on the structure of the set. A second quantity, instead, can be defined to measure the ultrametric content of the set, i.e. the overall closeness of the triplets to the fully ultrametric limit

$$\lambda_{um} = \left\langle \frac{\log \delta_1 - \log \delta_2}{\log \delta_1 + \log \delta_2} \right\rangle. \quad (9)$$

This quantity is invariant under trivialization of the metric and range from 0 (if all triplets are such that the two furthest points are at the same distance from the third — the negation of ultrametricity) to 1 (for a fully ultrametric set). The invariance under trivialization is reflected in the fact  $\lambda_{um}$  is nearly constant in Euclidean spaces of increasing dimension: for  $d = 1, 2, 3, 4, 5, 6, \dots, \infty$  it takes the values (with the above method for selecting triplets) 0.60, 0.54, 0.52, 0.51, 0.50, 0.50,  $\dots$ , 0.5. For the face set it takes the value  $\lambda_{um} = 0.53$ . Thus the ultrametric content of the face set, as perceived by the cells analysed, is equivalent to that of an Euclidean space, from which points are drawn at random with the above procedure, of dimension  $d \cong 2.5$  (also the value  $\bar{d}$  calculated for the face set matches the value corresponding to an Euclidean set with  $d \cong 2.5$ , but trivialization would take  $\bar{d} \rightarrow \infty$  while leaving  $\lambda_{um} = 0.53$ ).

#### 4. Conclusion

I have formulated ways in which to probe the structure of perceptual representations, as they are manifested in neuronal responses recorded in a region of the brain, when stimuli drawn from a given set are presented to the animal. Application of these methods to the responses of face cells in the primate temporal visual cortex does not lead to firm quantitative conclusions at this stage, but it does indicate reliable relations in the results, with those emerging from applying the same analysis to other representations (i.e. stronger metric content than for space representations) and to known metric structures (i.e. an ultrametric content similar to that of a set drawn from an Euclidean space of dimension 2–3). This suggest the

potential relevance of applying these analyses to the systematic study of sensory representations in the brain.

### Acknowledgements

The data for these analyses were obtained by Edmund Rolls — with whom also decoding algorithms were developed — and Martin Tovee; while useful comparisons were possible with non-face datasets kindly provided by other experimenters, including Bill Skaggs in Bruce McNaughton's lab. Joseph Atick introduced to me the work of Sirovich and Kirby. Partial support was from the Human Frontier Science Program, the Human Capital and Mobility Program of the EEC, and CNR, INFM and INFN of Italy.

### References

- Hasselmo, M.E., Rolls, E.T. and Baylis, G.C., 1989, The role of expression and identity in the face-selective responses of neurons in the temporal visual cortex of the monkey. *Behav. Brain Res.* 32, 203–218.
- O'Keefe, J., 1996, Hippocampal place cells: spatially-coded neuronal oscillators, in: *Perception, Memory and Emotion: Frontier in Neuroscience*, T. Ono, B.L. McNaughton, S. Molotchnikoff, E.T. Rolls and H. Nishijo (eds.) (Elsevier, Oxford), in press.
- Panzeri, P., Biella, G., Sotgiu, M.L. and Treves, A., 1995, Information theoretical analysis of thalamocortical ensemble coding during noxious stimulation. *Soc. Neurosci. Abstr.* 21, 114.
- Panzeri, S. and Treves, A., 1996, Analytical estimates of limited sampling biases in different information measures. *Network* 7, 87–107.
- Rolls, E.T., 1992, Neurophysiological mechanisms underlying face processing within and beyond the temporal cortical visual areas. *Phil. Trans. R. Soc. Lond. B* 335, 11–21.
- Rolls, E.T., Robertson, R. and Georges-Francois, P., 1995, The representation of space in the primate hippocampus. *Soc. Neurosci. Abstr.* 21, 1494.
- Rolls, E.T., Treves, A. and Tovee, M.J., 1996, The representational capacity of the distributed encoding of information provided by populations of neurons in the primate temporal visual cortex (submitted).
- Sirovich, L. and Kirby, M., 1987, Low-dimensional procedure for the characterization of human faces. *J. Opt. Soc. Am.* 4, 519–524.
- Treves, A. and Panzeri, S., 1995, The upward bias in measures of information derived from limited data samples. *Neural Comput.* 7, 399–407.
- Treves, A., Barnes, C.A. and Rolls, E.T., 1996, Quantitative analysis of network models and of hippocampal data. in: *Perception, Memory and Emotion: Frontier in Neuroscience*, T. Ono, B.L. McNaughton, S. Molotchnikoff, E.T. Rolls and H. Nishijo (eds.) (Elsevier, Oxford), in press.
- Wilson, M. and McNaughton, B.L., 1993, Dynamics of the hippocampal ensemble code for space. *Science* 261, 1055–1058.
- Yamane, S., Kaji, S. and Kawano, K., 1988, What facial features activate face neurons in the inferotemporal cortex of the monkey? *Expl. Brain Res.* 73, 209–214.
- Young, M.P. and Yamane, S., 1992, Sparse population coding of faces in the inferotemporal cortex. *Science* 256, 1327–1331.