Title: Probability Density Estimation for the Interpretation of Neural Population Codes

Author: Terence David Sanger

Institution: Massachusetts Institute of Technology

Department: Department of Brain and Cognitive Sciences

Contact Address: Terence D. Sanger

MIT, Room E25-526, Cambridge, MA 02139

tel: (617)253-8174, fax: (617)258-5342, email(preferred): tds@ai.mit.edu

Running Title: Population Code Density Estimation

Keywords: Population code, Population vector, Probability, Density Estimation, Motor Cortex, Tuning Curves, Poisson process, Cortical information processing

Submitted: Rapid Publication in *J. Neurophysiology*

1 Summary and Conclusions

1. Electrophysiological recording data from multiple cells in motor cortex and elsewhere are often interpreted using the "Population Vector" method pioneered by Georgopoulos and co-workers. This paper proposes an alternative method for interpreting coding across populations of cells that may succeed under circumstances in which the Population Vector fails.

2. Population codes are analyzed using probability theory to find the complete conditional probability density of a movement parameter given the firing pattern of a set of cells.

3. The conditional density when a single cell fires is proportional to the shape of the cell's tuning curve.

4. The conditional density when multiple cells fire is proportional to the product of their tuning curves.

5. Movement parameters can be estimated from the conditional density using statistical Maximum Likelihood or Minimum Mean-Squared Error methods.

6. Simulations show that density estimation correctly finds movement directions for non-uniform distributions of preferred directions and non-cosine cell tuning curves, while the Population Vector method fails for these cases.

7. Probability methods thus provide a statistically-based alternative to the Population Vector for interpreting electrophysiological recording data from multiple cells.

2 Introduction

The discovery of large numbers of cells in motor cortex that are tuned to specific properties of movement has led to the suggestion that collections of cells may represent a "population code" in which the aggregate activity yields more detailed and accurate movement information than could be extracted from any single cell (Georgopoulos *et al.* 1988). The "population vector" method was proposed as a technique both for the interpretation of such electrophysiological data, as well as a possible method by which the biological control system might convert cortical information into the control signals sent to the relevant muscles (Georgopoulos *et al.* 1988, Kalaska and Crammond 1992, Caminiti *et al.* 1990, Mussa-Ivaldi 1988). The population vector p is defined to be a weighted average of the cell firing rates d_i over a defined interval, where the weighting c_i is given by the vector "preferred direction" for each cell i.

$$p = \sum c_i d_i \tag{1}$$

The preferred direction is usually calculated by fitting a cosine to the cell firing rate (as a function of the direction of arm movement, for example) and choosing the peak of the resulting fit. It may also be calculated directly as the peak of a smoothed firing rate histogram. Several authors have shown sufficient conditions under which the population vector for arm movement will approximate the true direction of movement (Georgopoulos *et al.* 1988, Mussa-Ivaldi 1988) and these conditions can be generalized to (Sanger 1994)

$$\sum c_i c_i^T = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} \tag{2}$$

Despite the remarkable success of the population vector method, a number of theoretical problems have been raised (Sanger 1994, Loeb and Brown 1996). Perhaps most significant is the fact that the population vector is a linear function of the cell firing rates and thus is unable to compute any possibly nonlinear variables. As with any linear approximation it would be expected to perform well only in a limited region, and in (Caminiti *et al.* 1990) we see that the population vector gives a good approximation in a non-local region only if the individual cell tuning curves are allowed to change. By reducing the complete tuning curve to a single preferred direction, each cell is approximated by its linear component thus discarding potentially important information (Sanger 1994). If the distribution of preferred directions does not satisfy the condition (2), then the population vector may be incorrect. Finally, the population vector is not unique, since any linear transformation T can be applied to both p and c_i to yield an equally valid population vector equation:

$$Tp = \sum Tc_i d_i \ . \tag{3}$$

These concerns lead us to seek an alternative method for the interpretation of population codes. The technique proposed below uses a probabilistic interpretation of cell tuning curves to derive a statistical estimate for the true movement parameters.

3 Methods

Assume that the average firing rate of a cell in response to an externally measurable variable x can be specified by a tuning curve $\sigma(x)$. Further assume that the probability of the cell firing n times during a short time interval Δt is given by a Poisson statistic with rate $\sigma(x)$

$$P[\text{n spikes}|x] = (\sigma(x)\Delta t)^n \, e^{-\sigma(x)\Delta t}/n! \tag{4}$$

If $\sigma(x)$ is measured in spikes per second, then $\sigma(x)\Delta t$ is measured in spikes per interval Δt. x may be a vector representing a combination of many different possible experimental variables. For example, in a movement experiment it could include combinations of movement direction, initial position, load direction, and arm posture.

There are several circumstances under which the above assumptions may be violated. Trivially, a cell may not have any change in response as the variable x changes. Alternatively, the firing rate may be poorly described by a Poisson process (a derivation for a binomial process is given below). Most significantly, a cell may not encode information in its average firing rate, but might instead encode information in terms of firing time relative to another cell or to an extrinsic timing signal. The model shown does not take account of the effect of the cell refractory period, although a similar derivation could be performed for such statistics.

Bayes' law gives the distribution of x conditioned on the number of times n that the cell fired:

$$P[x|\text{n spikes}] = \sigma(x)^n P[x] e^{-\sigma(x)\Delta t}/\mathcal{N}_1 \tag{5}$$

$$P[x|\text{0 spikes}] = P[x] e^{-\sigma(x)\Delta t}/\mathcal{N}_2 \tag{6}$$

where $\mathcal{N}_1$ and $\mathcal{N}_2$ are normalization constants, and $P[x]$ is the prior probability density for x. Figure 1 shows an example of the shape of these conditional distributions for a broadly tuned cell. When the cell fires, it can be thought of as "voting" for different values of x depending on the magnitude of the tuning curve $\sigma(x)$. This is an extension of the population vector analysis in which each cell "votes" for its preferred direction (Georgopoulos *et al.* 1988). (See (Anderson 1995) for an alternative statistical interpretation of cell tuning curves.)

Figure 1 near here

(A similar derivation can be performed for statistics other than Poisson. For example, if the cell firing rate is described by a Binomial density $\sigma(x)^n(1-\sigma(x))^m$ where the cell fires n times during $n+m$ discrete time intervals, then $P[x|\text{n spikes}] = \sigma(x)^n(1-\sigma(x))^m P[x]/\mathcal{N}$.)

The conditional density for x given the spike pattern of the entire population has a particularly simple mathematical form if the Poisson processes governing different cells' firing rates are independent. This means that correlations between cells can be entirely explained by overlaps in their tuning curves, and not by direct interconnections. This situation is commonly assumed when data from single-electrode recordings on multiple trials is combined to estimate the population response. (If this assumption does not hold for some set of cells, then the joint conditional probability distribution $P[\mathrm{spike}_i, \mathrm{spike}_j, \mathrm{spike}_k, \ldots | x]$ for the related cells needs to be measured directly using multi-electrode recording. Bayes' law can then be used to calculate $P[x|\mathrm{spike}_i, \mathrm{spike}_j, \mathrm{spike}_k, \ldots]$ and this term can be multiplied into equation (7) below.) Under the independence assumption, $P[\mathrm{spike}_i, \mathrm{spike}_j|x] = P[\mathrm{spike}_i|x]P[\mathrm{spike}_j|x]$. We can then immediately write the conditional distribution for x given the state of the full population of cells:

$$P[x|\mathrm{population}] = \frac{1}{\mathcal{N}} \left(P[x] \prod_{\text{all cells}} e^{-\sigma_j(x)\Delta t} \right) \prod_{\text{firing cells}} \sigma_i(x)^{n_i} \tag{7}$$

where each cell i has tuning curve $\sigma_i(x)$ and fired n_i times during the small interval Δt, and $\mathcal{N}$ is a normalization constant. Equation (7) is the central result of this paper, and gives the mathematical formulation of the proposed algorithm for extracting density estimates from populations of spiking cells. Note that the term in parentheses is independent of which cells fired and can thus be pre-computed. This term accounts for non-firing cells as well as any possible non-uniformities in the distribution of cell tuning curves. Equation (7) can be calculated more efficiently as

$$P[x|\mathrm{population}] = \frac{1}{\mathcal{N}} P[x] \exp\left(\sum_i n_i \log \sigma_i(x) - \sigma_i(x)\Delta t \right) \tag{8}$$

Equation (7) shows that we take the product of the tuning curves of each cell that fired during the interval Δt. If a cell fires n_i times, then its tuning curve is raised to the power n_i. Multiplication of the tuning curves of all firing cells may lead to a very sharply tuned probability density estimate even for broadly-tuned cells. Figure 2 shows an example of how a product of two broadly-tuned curves can multiply to produce a more sharply-tuned conditional density estimate.

Figure 2 near here

In the discussion above, it was assumed that x is a continuous variable and that the tuning curves $\sigma(x)$ are continuous functions. However, this discussion applies equally well to categorical variables for which x takes on a small number of discrete values. (For example, x might indicate which finger moved.) In this case, equation (7) should be evaluated only at the allowed values of the categorical variable x, and the "tuning curves" $\sigma(x)$ will represent a list of firing frequencies for each possible value of x.

Once equation 7 has been evaluated for a particular experiment we have access not only to a single estimate of a movement parameter, but rather to the complete conditional probability density. We thus can make estimates of "confidence" based upon the variance in the density $P[x|\text{population}]$. In order to read out a best estimate of some desired parameter, we can use standard statistical methods such as Maximum Likelihood (ML) estimation, or Minimum Mean-Squared Error (MMSE) estimation. For ML estimation, choose the peak of the conditional density, and for MMSE estimation choose the mean.

The discussion so far has assumed that the cell tuning curves $\sigma(x)$ are known. These need to be computed from a prior experiment in which the cell behavior is systematically probed for each relevant controlled variable. (To prevent statistical overfitting, it is important that the tuning curves be estimated using a different set of spike data from that used to estimate the conditional density in equation 7.)

For categorical data, the tuning curve can be estimated directly from the average firing rate for each value of x. For continuous data, we can use Parzen windows density estimation to approximate

$$\sigma(x) \approx \sum_k \psi(\|x - x_k\|) \tag{9}$$

where $\psi()$ is a smooth function peaked at zero (frequently a narrow Gaussian), the x_k's are the values of x that occurred during each spike k, $\|x - x_k\|$ is the distance between x and x_k in a possibly multi-dimensional space, and the sum is taken over all spikes for this cell during some specified time interval. If more than one parameter is varied during the experiment, then the tuning curve will be estimated as a function of every modified variable leading to a sum of multi-dimensional "bump" functions $\psi(\|x - x_k\|)$. Note that when population estimates based on equation 7 are made from many cells the individual tuning curves need to be computed only approximately since random errors will average out. In some cases accurate results may be obtained even when the tuning curves are estimated from only a few spikes.

4 Results

As might be expected, if all cells in a population have approximate cosine-tuning and if the preferred directions are uniformly distributed, then both the population vector method (Georgopoulos *et al.* 1988) and the method presented here yield similar results. If these assumptions are violated, however, the population vector can yield substantially incorrect results. The data in (Kalaska and Crammond 1992) provides an example of a set of cells whose tuning for applied load direction is not uniformly distributed, and in this case the population vector fails to predict the true direction of applied load. In (Schwartz *et al.* 1988) and elsewhere we see that the the data is not always well fit by cosine tuning curves, and there is a large number of cells with no significant cosine-tuned component (Sanger 1994).

Figure 3 shows two simulated comparisons with the population vector method. Data are presented using a format following (Georgopoulos *et al.* 1989), and population vectors are computed as in (Georgopoulos *et al.* 1988). In the first example, all cells are primarily cosine-tuned (a cosine function predicts more than 75% of the variance), but the distribution of preferred directions is nonuniform. As in (Kalaska and Crammond 1992), the population vector yields systematic errors in estimation of the true direction. In the second example the distribution of preferred directions is uniform, but the cosine component is dominant in only 5% of cells (on average it accounts for approximately 25% of the variance). In this case the population vector generates essentially random estimates with no apparent pattern. In both cases the probability density method with ML estimation finds accurate estimates of the true direction.

Figure 3 near here

5 Discussion

The population vector method has been proposed not only as a method for interpreting experimental data, but also as a potential model for neural computation in the brain. In contrast, it is important to realize that the technique proposed here is intended only to interpret experimental data but not to model neural computation. The density estimation method is designed to extract the maximum amount of information from a neural code for the purposes of our understanding of that code. But it should be realized that the brain's purpose is often to perform a particular computation rather than extract information explicitly, and thus it may neither need nor use methods this powerful.

The technique presented here has several important advantages over current methods for interpreting population codes. First, it is based on a rigorous statistical analysis that makes use of only very general assumptions about cell firing patterns. As simulated in figure 3, there is no need for a particular shape to the tuning curves, and no dependence on the distribution of preferred directions. The density estimate will automatically indicate lower confidence (higher variance) in regions for which few cells respond, but it will still allow computation of the correct ML or MMSE estimate. Second, the density estimate uses the complete shape of the tuning curves to extract the maximum amount of information from the population. It finds the statistically best estimate of the desired parameter for the given set of spikes. Third, once the tuning curves are computed there is no requirement for spike-rate averaging, and estimates can be made rapidly in "real-time" if multi-electrode recordings are available. Fourth, the technique is easily and rapidly computable using multiplication operations. Fifth, it can be applied to categorical variables or to tuning curves computed simultaneously for many different motor or sensory variables. Sixth and perhaps most importantly, the use of this technique allows the full power of probability theory to be applied to the analysis of neural data in ways that may yield new insights into information representation in the brain.

6 Acknowledgements

I wish to thank Emilio Bizzi, Simon Giszter, Peter Dayan, and Barak Pearlmutter for comments and discussions. The author was supported by a McDonnell-Pew Post-doctoral Fellowship. Dr. Bizzi's laboratory is supported by grants from the National Institute of Health (NS09343) and the Office of Naval Research (N00014/90/J/1946).

References

Anderson C. H., 1995, Unifying perspectives on neuronal codes and processing, Proc. 19th Int'l workshop on Condensed Matter Theories, Caracas, Venezuela.

Caminiti R., Johnson P. B., Urbano A., 1990, Making arm movements within different parts of space: Dynamic aspects in the primate motor cortex, *J. Neuroscience*, 10(7):2039–2058.

Georgopoulos A. P., Kettner R. E., Schwartz A. B., 1988, Primate motor cortex and free arm movements to visual targets in three-dimensional space. II. Coding of the direction of movement by a neuronal population, *J. Neurosci*, 8(8):2928–2937.

Georgopoulos A. P., Lurito J. T., Petrides M., Schwartz A. B., Massey J. T., 1989, Mental rotation of the neuronal population vector, *Science*, 243:234–236.

Kalaska J. F., Crammond D. J., 1992, Cerebral cortical mechanisms of reaching movements, *Science*, 255:1517–1523.

Loeb G. E., Brown I. E., 1996, Directional motor control, *TINS*, 19(4):137.

Mussa-Ivaldi F. A., 1988, Do neurons in the motor cortex encode movement direction? An alternative hypothesis, *Neurosci. Lett.*, 91:106–111.

Sanger T. D., 1994, Theoretical considerations for the analysis of population coding in motor cortex, *Neural Computation*, 6(1):12–21.

Schwartz A. B., Kettner R. E., Georgopoulos A. P., 1988, Primate motor cortex and free arm movements to visual targets in three-dimensional space. I. Relations between single cell discharge and direction of movement, *J. Neurosci*, 8(8):2913–2927.

List of Figures

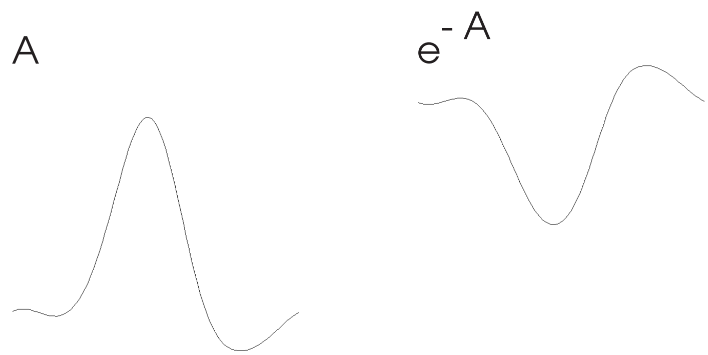

Figure 1: Sample probability densities for a movement variable x given that a simulated broadly-tuned cell fired once ("A"), or not at all ("e^{-A}").

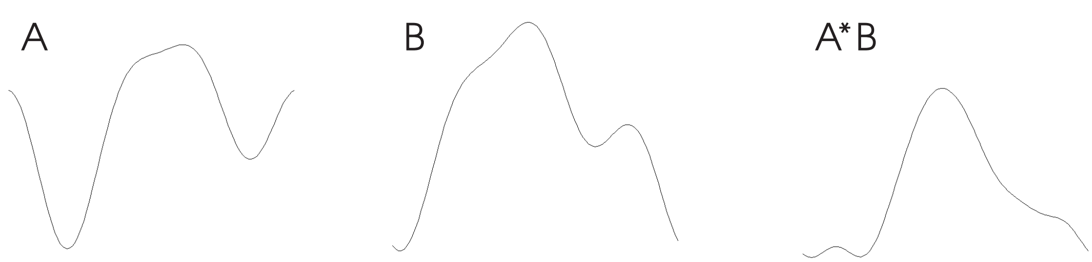

Figure 2: Example of multiplication of tuning curves. The product of curves A and B yields a more sharply tuned conditional density $A * B$.

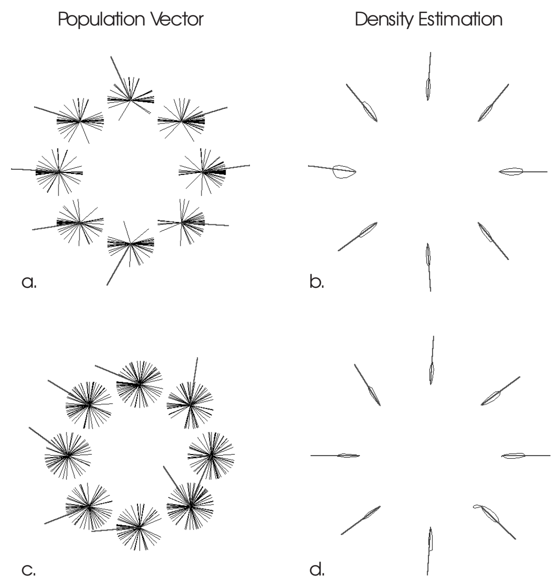

Figure 3: Comparison of Population Vector and Density Estimation methods. Long lines show the estimated direction computed by the two methods for 8 different movements. The true movement direction is radially outward from the center of each figure. (a),(b): Cells are primarily cosine-tuned, but the distribution of preferred directions favors horizontal orientations. (c),(d): There is a uniform distribution of preferred directions, but the tuning curves have significant non-cosine components in addition to the cosine components. (Tuning curves are generated for 100 cells by choosing random circular Fourier coefficients. Δt is 10 msec, the average cell firing rate is 100 spikes/sec, and maximum likelihood (ML) estimation is used to compute the estimated movement direction for the density estimation method.)