
Reading out population codes with a matched filter

Mark C. W. van Rossum^{1,*}, Alfonso Renart², Sacha B. Nelson¹,
X.-J. Wang² and Gina G. Turrigiano¹

1. Dept. of Biology, MS 008
Brandeis University, Waltham, MA 02454

* email: vrossum@brandeis.edu

2. Center for Complex Systems
Brandeis University, Waltham, MA 02454

Abstract

We study the optimal way to decode information present in a population code. Using a matched filter, the performance in Gaussian additive noise is as good as the theoretical maximum. The scheme can be applied when correlations among the neurons in the population are present. We show how the read out of the matched filter can be implemented in a neurophysiological realistic manner. The method seems advantageous for computations in layered networks.

At many stages of neural information processing in the brain, information is not carried by a single neuron but shared by many neurons in parallel. Evidence for this has been found in sensory systems [7], motor systems [3, 6] and hippocampal place cells [4]. In population codes neurons have receptive fields centered at different locations, but their tuning curves are wide and overlap considerably. Thus a stimulus will activate a large population of neurons. Combining the response rates of the different neurons, the stimulus can be reconstructed. The population coding is not vulnerable to failure of a single neuron. In addition, the population signal is less noisy than the signal from a single neuron alone. This can be used to increase accuracy or limit temporal averaging to allow for a quicker response [5, 13].

One approach to read out the encoded stimulus is to average the responses of all neurons in the population in order to construct the *population vector* [3]. Due to noise in the responses of the neurons there are trial to trial fluctuations in the estimate of the encoded quantity (e.g. motion direction). A statistical construct, the *Cramer-Rao* bound, gives the minimal trial to trial error obtainable at a certain noise level. But the Cramer-Rao bound does not prescribe how to accomplish the minimal error. It is not unreasonable to assume that the nervous system in its evolutionary quest for efficiency performs close to the optimum. However, the population vector scores in most situations worse than the bound [11]. In this paper we discuss optimal readout mechanisms.

Interestingly, readout using a non-linear recurrent network approaches the optimal performance [9, 2], suggesting that possibly on a circuit level the brain has indeed implemented

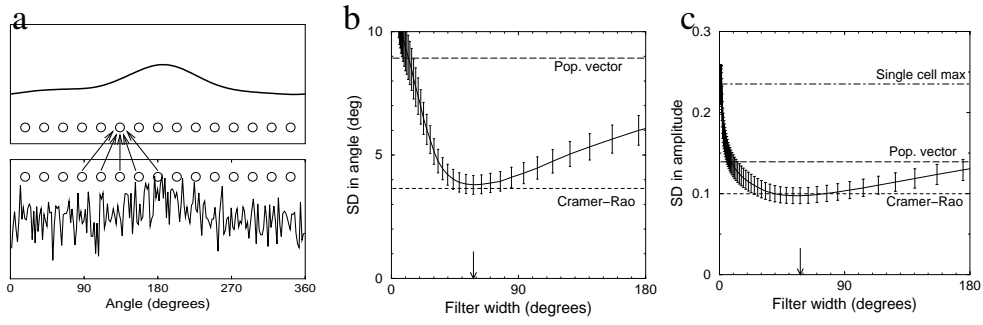


Figure 1: a) Simulated population activity of 200 units in response to a stimulus with an 180° angle. The stimulus is distorted with Gaussian additive noise, $\sigma = 0.5$ (lower panel). Upper panel: The response filtered with a matched filter. b) The error of the matched filter in the estimate of the stimulus angle as a function of the width of the filter. The standard deviation across trials in the angle estimate is plotted. The error approaches the theoretical limit (dotted line) when the filter width equals the width of the stimulus, that is, when the filter is matched (indicated by the arrow). Dashed line: error when using the population vector. c) as b) but for the amplitude estimate.

a population code reader. Yet, such a recurrent network has two possible disadvantages: 1) The output of the network becomes non-linear in the input; in an extreme case a sub-threshold input will not lead to any response in the output. 2) The network is recurrent and therefore some time passes before the final output is reached, slowing down computation.

Here we propose a simple, intuitive network to readout the population code of a direction selective population. We extend the readout mechanism to not only read the encoded angle of the stimulus but also the amplitude of the stimulus. We prove that the method equals the Cramer-Rao bound for the both the angle and the amplitude.

1 Matched filter

We consider a population N neurons coding for an angle between 0 and 360 degrees. On a given trial, the response of neuron i is $r_i = f_i(\theta) + \nu$, where $f_i(\theta) = A \exp[(\cos(\theta) - 1)/w^2]$ is the neuron's tuning curve, and ν is a Gaussian white noise with zero mean and independent across cells unless specified otherwise. The tuning width w is chosen 60° , the stimulus amplitude A is normalized to one. Consider a population response illustrated in Fig. 1a. The task is to determine the coded angle, i.e. the location of the peak, as precise as possible despite the presence of noise. In addition, we extend the task to determine the coded amplitude, i.e. the height of the peak. The amplitude of the signal is important for many psycho-physical tasks, especially when two stimuli of different amplitude need to be compared.

The question how to readout a population code optimally can be understood by considering the following two extreme approaches. One extreme is to select the neuron which gives the largest response, and to assume that the coded angle is the preferred direction of that neuron. This method is obviously very sensitive to noise and does not use that one can average over neurons to reduce the noise. The other extreme is to average over all neurons, thereby constructing the average, or population vector given by $\mathbf{p} = \frac{1}{N} \sum_i r_i \mathbf{e}_i$, where $\mathbf{e}_i = \{\cos(2\pi i/N), \sin(2\pi i/N)\}$ is a two dimensional unit vector pointing in the direction that neuron i codes for. The population vector has the advantage of averaging over many neurons. As a result it performs well, especially if the tuning curves are broad

[10]. However, if the tuning curve is narrow, most neurons only contribute noise to the population vector estimate without providing information about the coded angle. Indeed, in most cases the population vector performs under the theoretical limit [9].

The question thus seems to be: what is the optimal number of neurons one should average over and how should their responses be weighted? The answer is a matched filter. A matched filter weighs the responses according to the shape of the tuning curve itself. In other words the responses are convolved with the average tuning curve (Fig. 1a), the encoded angle is where the output of the filtered response is maximal. In addition, the stimulus amplitude is given by the height of the maximum [12]. In Fig. 1b,c the standard deviation (SD) in the angle and amplitude estimate is shown as function of the filter width. When the filter width has exactly the width of the tuning curve, that is, the filter is matched, the performance equals the theoretical limit. Note that for wide filters the result approximates the population vector result, and for narrow filters it approaches the single cell maximum result (Fig. 1c).

It is common to retrieve information from a population code using the maximum likelihood estimator. Although it performs optimally under many conditions, it has no obvious biological implementation as it is a fitting procedure. As we prove now, the maximum likelihood and the matched filter are identical for Gaussian additive noise. Using maximum likelihood the encoded angle is given by the angle which is most likely to have been encoded given the noise distribution.

$$\theta_{est}^{ML} = \arg \max_{\theta} \ln P(\mathbf{r}|\theta),$$

where $P(\mathbf{r}|\theta)$ is the probability that a stimulus with angle θ leads to response \mathbf{r} , the $\arg \max_x y$ stands for the value of x that maximizes y . For independent Gaussian additive noise this simplifies to

$$\theta_{est}^{ML} = \arg \max_{\theta} \sum_i [r_i - f_i(\theta)]^2 = \arg \max_{\theta} \sum_i r_i f_i(\theta), \quad (1)$$

where we used that the tuning curves are homogeneously distributed such that $\sum_i f_i(\theta)^2$ does not depend on θ .

On the other hand, using the matched filter we search for the output which gives the maximal response. For the matched filter the filter kernel is proportional to the tuning curve centered around zero angle (where the proportionality constant B is chosen such that the output and the input amplitudes are equal in the noiseless case), and the output of the filter is denoted o_i . The output is

$$o_i = B \sum_k r_k f_{i-k}(\theta = 0) = B \sum_k r_k f_k(\theta = 2\pi i/N)$$

We search for the output neuron i which gives the maximal response. The angle estimate is obtained by reading off to which angle neuron i corresponds,

$$\theta_{est}^{MF} = \arg \max_i o_i = \arg \max_{\theta} \sum_k r_k f_k(\theta) \quad (2)$$

where we assume that the tuning curves are densely distributed, such that the error due to discretization is much smaller than the error caused by the noise. Eqs.1 and 2 are the same, meaning that both methods give identical estimates from trial to trial and give therefore the same performance.

2 Breakdown of maximum likelihood

We have just shown that the matched filter performs as good as the maximum likelihood algorithm. But does this performance always reach the Cramer-Rao limit? It is often




Figure 2: The dependence of the reconstruction error on the strength of the noise level (defined as noise SD divided by stimulus amplitude). Upper panel: The error in the stimulus amplitude using a matched filter (dashed line) and maximum likelihood (dotted line). They approximate the Cramer-Rao bound (solid line) for all noise strengths. The matched filter estimate has a bias (triangles) but the bias is much smaller than the error.

Lower panel: The reconstruction error in the stimulus angle (dashed line) obeys the Cramer-Rao bound (solid line) for low noise. But for high noise the error worsens and deviates from the Cramer-Rao bound. The maximum likelihood estimate (dotted line, overlapping with the dashed line) performs identically to the matched filter.

assumed that the performance of the maximum likelihood fit precisely equals the value given by the Cramer-Rao bound [10, 9]. Indeed, the amplitude estimate follows the Cramer-Rao bound, Fig. 2, top. However, for the angle estimate both the maximum likelihood and the matched filter perform worse than the Cramer-Rao bound for high noise conditions, lower panel Fig. 2. The underlying reason is that the angle is encoded non-linearly in the response of the neurons. For non-linearly encoded quantities the maximum likelihood reaches the Cramer-Rao limit only for small noise amplitudes [12]. Although we did not find a superior readout method for these high noise levels, in this regime the noise is unrealistically high from a biological point of view.

One might have expected that at high noise levels the amplitude and angle estimates can confound each other. Namely, a spurious maximum can arise due to the randomness, and this maximum would lead to a wrong estimate of both angle and amplitude. Furthermore, taking the maximum seems a dangerous strategy which could lead to a systematic tendency (bias) to overestimate the amplitude. To test for this we did the maximum likelihood fit for the amplitude, using the exact location of the angle. And likewise, we did the maximum likelihood fit for the angle given the true amplitude of the stimulus. For the tested noise levels we did not see a significant change in the performance. In addition, because of symmetry, there is never a bias in the angle estimate. The bias in the amplitude estimate was always much less than the reconstruction error (triangles in Fig. 2, upper panel).

3 Effect of correlation

Until here we assumed that the noise was uncorrelated across neurons. However, when neurons share noisy inputs, the fluctuations in their responses will be correlated. Physiologically, typical correlation coefficients in firing rates are between 0.1 and 0.2 [14]. We consider all cells to have the same correlation coefficient c . The effect of these correlations on an estimation depend on the quantity to be estimated. When the average activity of a pool of cells carrying the same signal is estimated, correlation increases the error [14].




Figure 3: The error in the angle and amplitude estimate as a function of correlation level. Left: The angle estimate improves with increasing correlation. Right: The CR bound in amplitude estimate depends non-monotonically on the correlation level. After being corrected, the matched filter follows this behavior.

However, correlation has a more subtle effect if different cells carry different signals, as is our case here [1]. The angle estimate improves with increasing correlation and the amplitude estimate initially deteriorates, but, already for extremely small correlation intensities improves for increasing correlation [1]. The improvement of the estimates with increasing correlation can be understood by considering the extreme case of 100% correlation. In that case the noise is identical in all neurons, so that from trial to trial the whole response profile is shifted up or down by a random amount, but otherwise the shape of the response is noiseless. Obviously, in this case the average activity will follow the fluctuations and show large errors, but both the angle and amplitude of the signal can be read off perfectly. In essence, when different cells carry different messages, correlations make it more difficult to confound the messages of different neurons, improving the quality of the estimation.

We tested if the matched filter can still perform optimally when correlation is present. To that end we mixed in a common noise source to all neurons such that the total noise was kept constant. We measured again the performance of the estimate. Without any adjustment the matched filter, the angle estimate follows the Cramer-Rao bound, Fig. 3 left. But the amplitude reconstruction error increases with increasing correlation, Fig. 3 right.

In general, matched filters can be adjusted to perform optimally when the noise is correlated across neurons. The idea is to first filter the input such that the noise becomes white. This is called a pre-whitening filter. After pre-whitening, the task is again to detect a known signal amidst white noise, for which one uses a matched filter. Here, in the case of infinite range correlation, whitening is performed by subtracting the average signal. The performance of the matched filter after whitening is again optimal, although possibly worse than in the uncorrelated case. Thus the matched filter is improved by subtracting common fluctuations in the signal, i.e. by a center-surround filter. We set the filter to be

$$o_i = \sum_j r_j \left[f_{j-i}(0) - \frac{cN}{1 + (N-1)c} \bar{f} \right], \quad (3)$$

where $\bar{f} = \frac{1}{N} \sum_k f_k(0)$ indicates the tuning curve averaged over locations. With this expression, the matched filter and maximum likelihood are, again, identical, and thus the

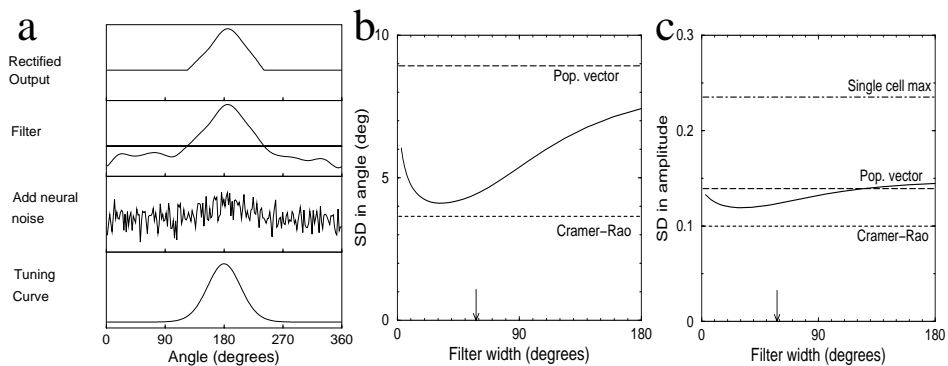


Figure 4: Reading out the maximum of the matched filter. a) Method of reading out: The tuning curve (lower panel), tuning curve with noise (next panel), the output of the filter (next panel), and the half-wave rectified filter output (upper panel). b+c) Performance of the angle and amplitude estimate (thick curves) as a function of filter width. The arrow represents the tuning curve width. Angle and amplitude are estimated by calculating the population vector of the filtered and rectified response.

corrected matched filter follows the theoretical limit, Fig. 3. The strength of the surround inhibition depends on the correlation c and the number of neurons N . It is interesting to note that the optimal balance between center and surround depends on the correlation strength. For $cN \gg 1$, the surround area equals center area and fluctuations common to all neurons will not change the output; all common noise is filtered out.

4 Reading out the matched filter

Neural implementation of a matched filter seems very plausible given the typical receptive field layout. However, the maximum still has to be located, and to reconstruct the stimulus amplitude, the value of the maximum has to be calculated. (It does not help to take the population vector of the filtered response; as the filtering is a linear operation, the performance equals the performance of the population vector of the non-filtered response.) One way is to implement a winner-take-all network. The solution we propose, is to rectify the filter output and take the population vector of this rectified response, thereby ignoring all responses which are too small and probably just contain noise, Fig 4a. Using an appropriately normalized balanced center-surround filter, a value of the rectifying threshold exists such that both the amplitude and the angle can be simultaneously accurately read out, with an error close to the Cramer-Rao limit. The tuning width of the spatial filter has to be chosen smaller than the width of the tuning curve, roughly half of that, Fig. 4b+c. As we note in the Discussion, after filtering and rectification the output curve is about as wide as the tuning curve of the input, Fig. 4a.

5 Discussion

We have shown that given a population response contaminated with Gaussian additive noise, the encoded stimulus can be extracted by using a matched filter. The matched filter resembles the template matching introduced in earlier works [8, 4], and the method also shows some similarity to a recurrent network halted after a few iterations [2]. However, the proposed method is simpler and mathematically more tractable.

The advantages of the matched filter and the proposed readout are: 1) It is easily implemented in neural circuitry. The spatial filter is provided by diverging connectivity with lateral inhibition, the rectification is inherent to a neuron's input-output relation. 2) It is purely feed-forward and is therefore fast. 3) The amplitude of the stimulus is preserved. 4) The matched filter is easily extended to cases when correlation is present, and possibly also non-Gaussian noise can be treated.

Fig. 4a shows that the output has about the same width as the input (upper and lower panel). This suggests that our readout is consistent with an architecture in which computations are done at multiple, subsequent layers. At every layer neural noise is added, but it is optimally filtered out by spatial filtering. If the filter would be too narrow, noise would accumulate in subsequent layers. If, on the other hand, the filter would be too wide the stimulus angle would be lost. With the matched filter the output has the same width as the input, and the network preserves the advantages of the population code. Computational consequences of these architectures are currently examined.

Acknowledgments

Supported by the Sloan foundation.

References

- [1] L. F. Abbott and P. Dayan. The effect of correlated variability on the accuracy of a population code. *Neural Comp.*, 11:91–101, 1999.
- [2] S. Deneve, P. E. Latham, and A. Pouget. Reading population codes: a neural implementation of ideal observers. *Nature Neurosci.*, 2:740–745, 1999.
- [3] A. P. Georgopoulos, A. B. Schwartz, and R. E. Kettner. Neuronal population coding of movement direction. *Science*, 233:1416–1419, 1986.
- [4] Zhang K, I. Ginzburg, B.L. McNaughton, and T. J. Sejnowski. Interpreting neuronal population activity by reconstruction: unified framework with application to hippocampal place cells. *J. Neurophysiol.*, 79:1017–1044, 1998.
- [5] B. W. Knight. Dynamics of encoding in a population of neurons. *J. Gen. Physiol.*, 59:734–766, 1972.
- [6] C. Lee, W. H. Rohrer, and D. L. Sparks. Population coding of saccadic eye movements by neurons in the superior colliculus. *Nature*, 332:357–360, 1988.
- [7] M. A. L. Nicolelis, A. A. Ghazanfar, C. R. Stambaugh, L. M. O. Olieviera, M. Laubach, J. K. Chapin, R. J. Nelson, and J. H. Kaas. Simultaneous encoding of tactile information by three primate cortical areas. *Nature Neurosci.*, 1:621–630, 1998.
- [8] M. A. Paradiso. A theory for the use of visual orientation information which exploits the columnar structure of striate cortex. *Biol. Cybern.*, 58:35–49, 1988.
- [9] A. Pouget, K. Zhang, S. Deneve, and P. E. Latham. Statistically efficient estimation using population coding. *Neural Comp.*, 10:373–401, 1998.
- [10] H. S. Seung and H. Sompolinsky. Simple models for reading neuronal population codes. *Proc. Natl. Acad. Sci.*, 90:10749–10753, 1993.
- [11] H. P. Snippe. Parameter extraction from population codes: a critical assessment. *Neural Comp.*, 8:511–529, 1996.
- [12] H. L. Van Trees. *Detection, Estimation, and Modulation Theory I*. Wiley, New York, 1968.
- [13] M. C. W. van Rossum, G. G. Turrigiano, and S. B. Nelson. Fast propagation of firing rates through layered networks of neurons. submitted.
- [14] E. Zohary, M. N. Shadlen, and W. T. Newsome. Correlated neuronal discharge rate and its implications for psychophysical performance. *Nature*, 370:140–144, 1994.