Lecture 7. Coding and Representation

Reading Assignments:

None

Variability in neuronal responses

Neuronal response *f(.)* can be represented as the sum of:

- a mean response to given stimuli *s,* given by the neuron's tuning *t(.)*
- a noise term (that may also be stimulus-dependent) *n(.)*

$$f(s) = t(s) + n(s)$$

Examples of sources of variability (n(.) term):

- transduction noise at the sensors (e.g., quantum light absorption noise)
- synaptic noise
- cross-talk noise due to activity in neighboring neurons
- random fluctuations in ionic channel function

Mean response

Mean response exhibits a tuning curve, typically well approximated by a Gaussian or cosine function.

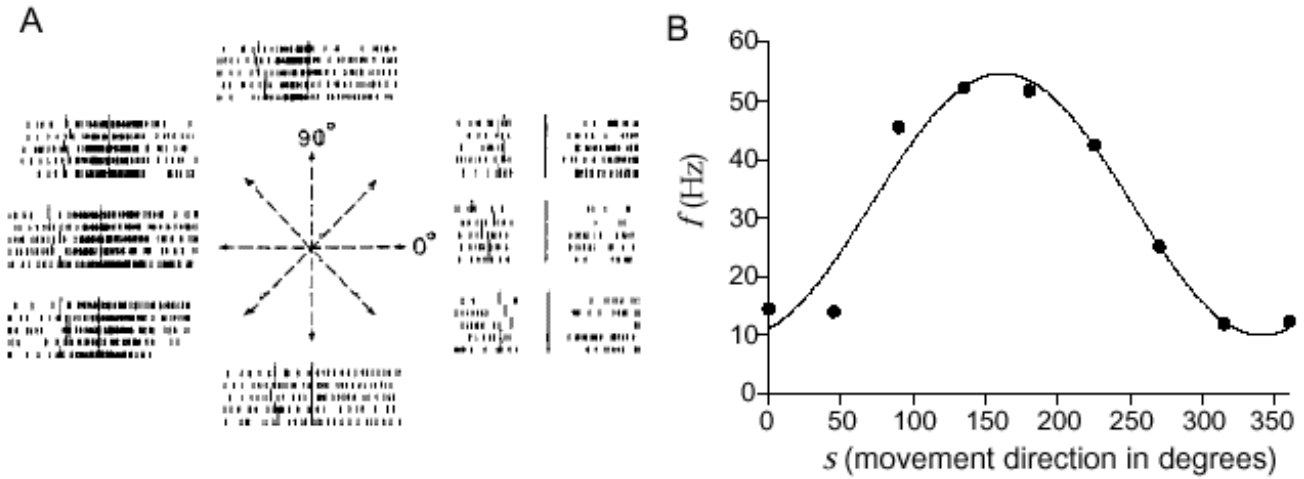

Figure 1.6: A) Recordings from the primary motor cortex of a monkey performing an arm reaching task. The hand of the monkey started from a central resting location and reaching movements were made in the directions indicated by the arrows. The rasters for each direction show action potentials fired on five trials. B) Average firing rate plotted as a function of the direction in which the monkey moved its arm. The curve is a fit using the function 1.15 with parameters $r_{max} = 54.69$ Hz, $r_0 = 32.34$ Hz, and $s_{max} = 161.25°$. (A adapted from Georgopoulos et al., 1982 which is also the source of the data points in B.)

Noise term

but responses to exactly identical stimuli can vary greatly from trial to trial…

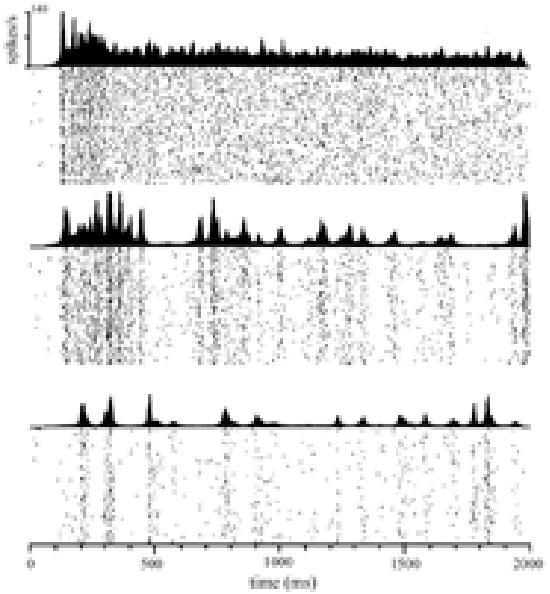

In a first approximation, we consider the variability in spike count and timing, and discard any variability in the exact shape of each action potential.

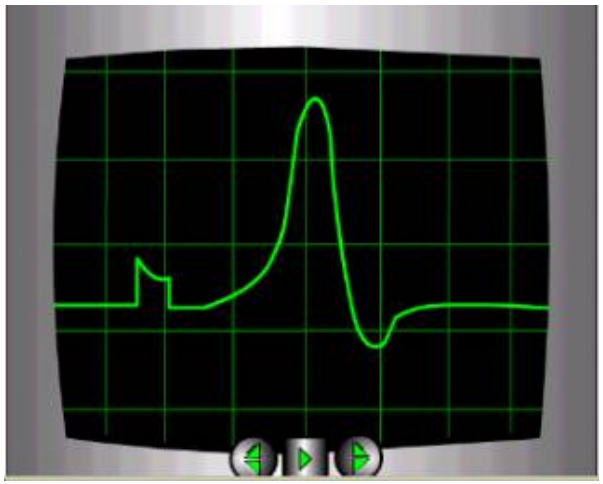

Spike Variability

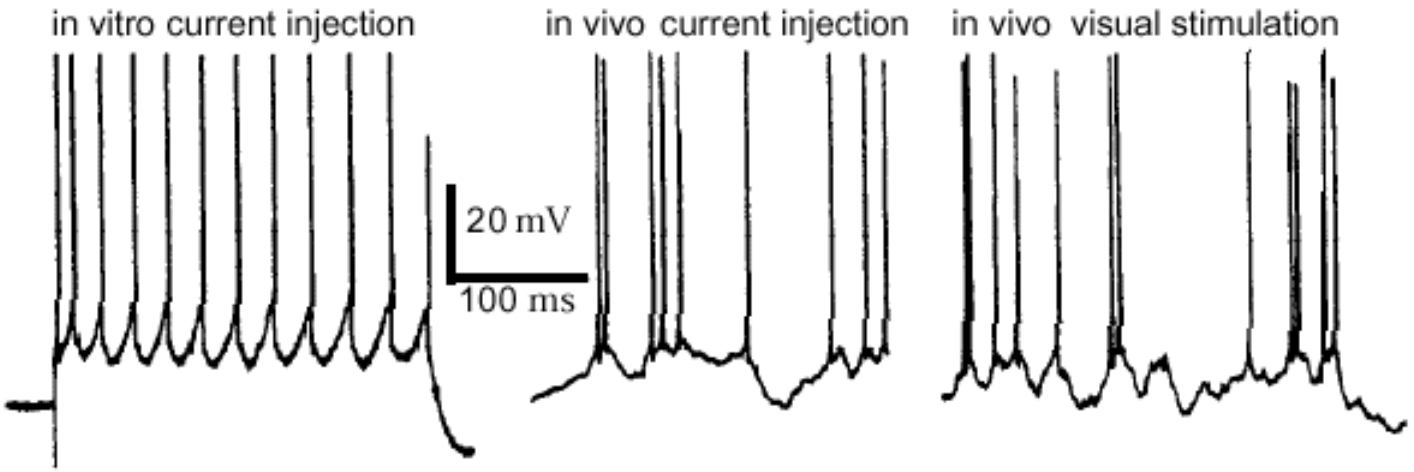

Figure 1.17: Intracellular recordings from cat V1 neurons. The left panel is the response of a neuron in an *in vitro* slice preparation to constant current injection. The center and right panels show recordings from neurons *in vivo* responding to either injected current (center), or a moving visual image (right). (Adapted from Holt et al., 1996.)

Estimating Firing Rates

Basic idea:

count spikes in
a small time interval.

Rate (frequency)=
spike count / interval.

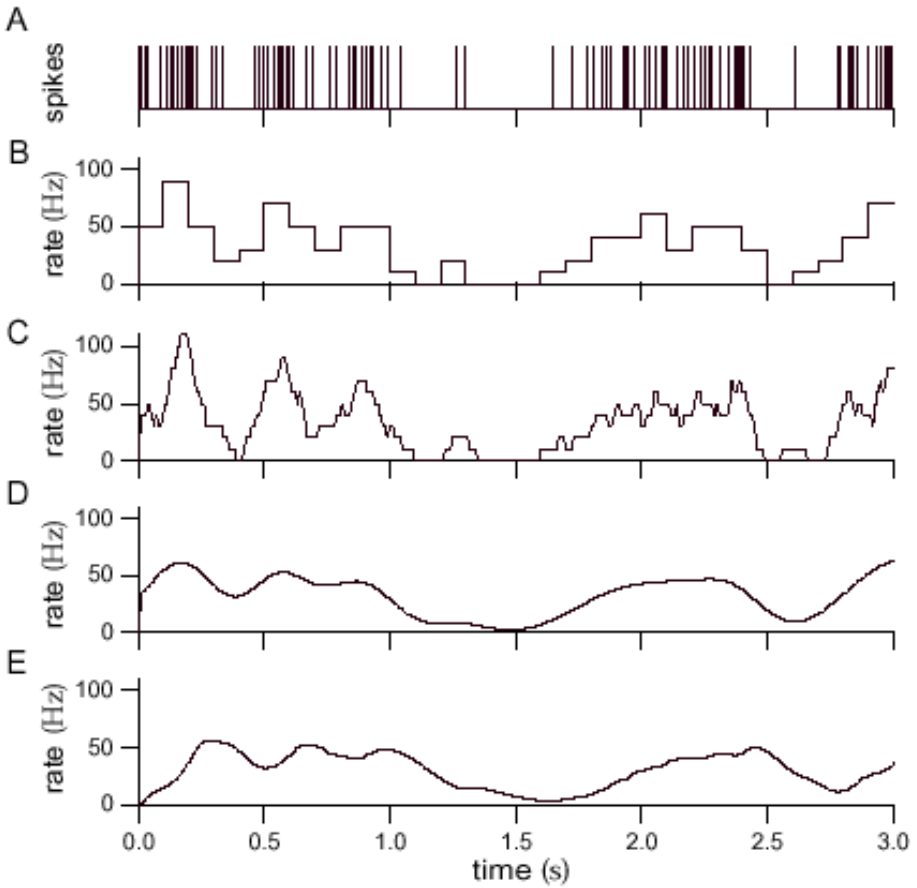

Figure 1.4: Firing rates approximated by different procedures. A) A spike train from a neuron in the inferior temporal cortex of a monkey recorded while that animal watched a video on a monitor under free viewing conditions. B) Discrete-time firing rate obtained by binning time and counting spikes with $\Delta t = 100$ ms. C) Approximate firing rate determined by sliding a rectangular window function along the spike train with $\Delta t = 100$ ms. D) Approximate firing rate computed using a Gaussian window function with $\sigma_t = 100$ ms. E) Approximate firing rate for an α function window with $1/\alpha = 100$ ms. (Data from Baddeley et al., 1997.)

Population Code

Each neuron is tuned to specific stimuli. Thus, a population of neurons with different preferred stimuli can represent all possible stimuli.

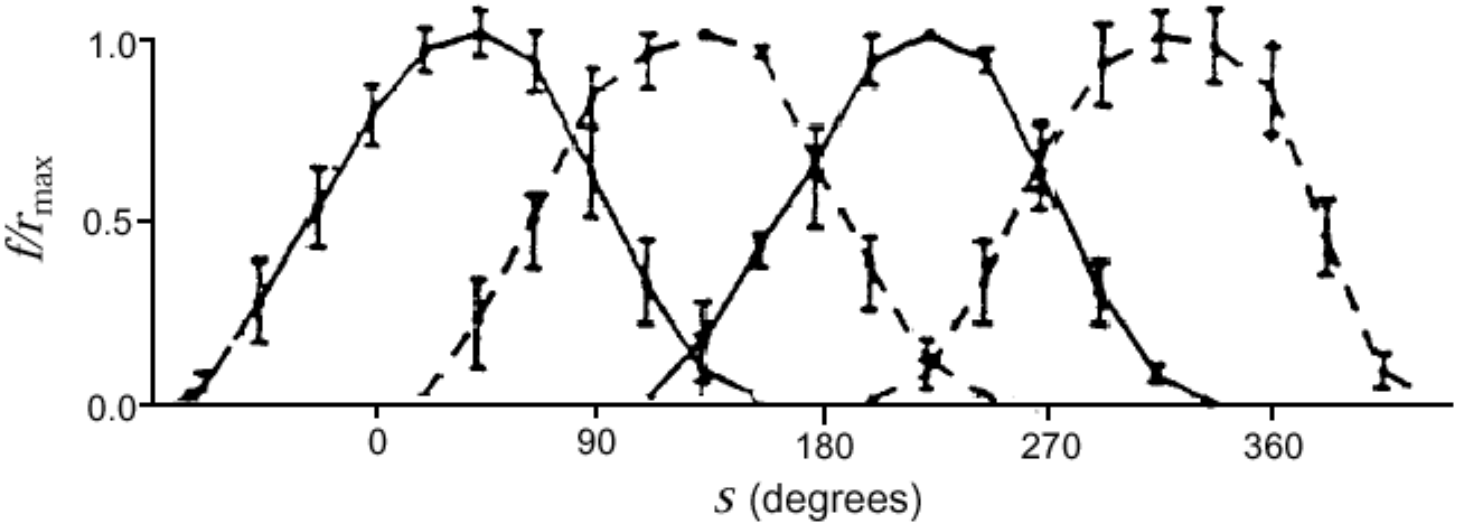

Figure 3.4: Tuning curves for the four low-velocity interneurons of the cricket cercal system plotted as a function of the wind direction s. Each neuron responds with a firing rate that closely approximated by a half-wave rectified cosine function. The preferred directions of the neurons are located 90° from each other, and r_{max} values are typically around 40 Hz. Error bars show standard deviations. (Adapted from Theunissen and Miller, 1991.)

Population Vector

- Represent each preferred stimulus by vector in stimulus space.

- Population vector: weighted sum of preferred vectors (weighted by response)

$$\vec{v}_{\text{pop}} = \sum_{a=1}^{4} \left(\frac{r}{r_{\max}} \right)_a \vec{c}_a \, .$$

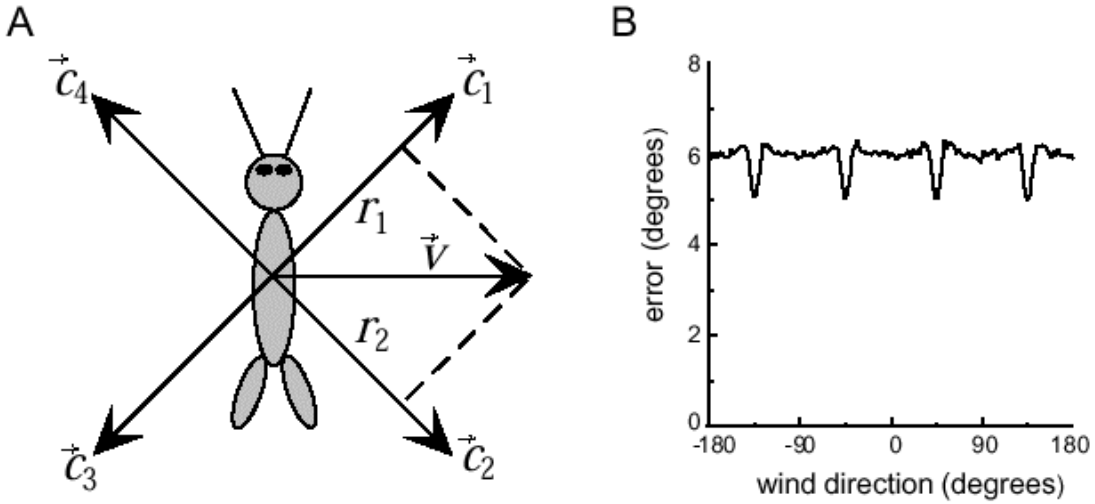

Figure 3.5: A) Preferred directions of four cercal interneurons in relation to the cricket's body. The firing rate of each neuron for a fixed wind speed is proportional to the projection of the wind velocity vector $\vec{v}$ onto the preferred direction axis of the neuron. The projection directions $\vec{c}_1$, $\vec{c}_2$, $\vec{c}_3$ and $\vec{c}_4$ for the four neurons are separated by 90°, and they collectively form a Cartesian coordinate system. B) The root-mean-square error in the wind direction determined by vector decoding of the firing rates of four cercal interneurons. These results were obtained through simulation by randomly generating interneuron responses to a variety of wind directions, with the average values and trial-to-trial variability of the firing rates matched to the experimental data. The generated rates were then decoded using equation 3.21 and compared to the wind direction used to generate them. (B adapted from Salinas and Abbott, 1994.)

Population Vector

Works very well
in many cases!

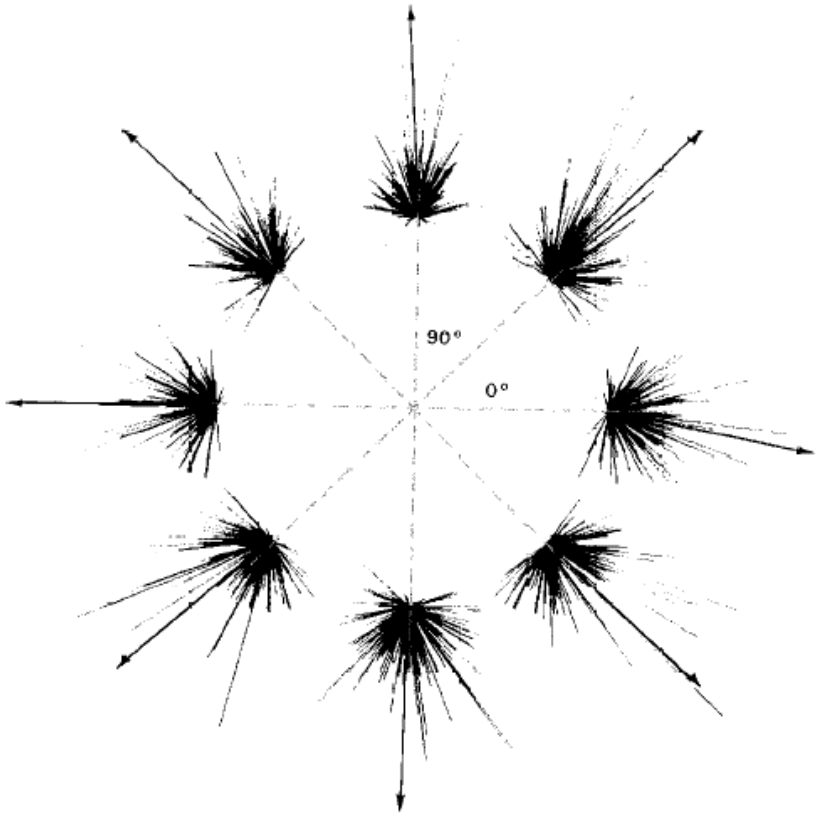

Figure 3.6: Comparison of population vectors with actual arm movement directions. Results are shown for eight different movement directions. Actual arm movement directions are radially outward at angles that are multiples of 45°. The groups of lines without arrows show the preferred direction vectors of the recorded neurons multiplied by their firing rates. Vector sums of these terms for each movement direction are indicated by the arrows. The fact that the arrows point approximately radially outward shows that the population vector reconstructs the actual movement direction fairly accurately. (Figure adapted from Kandel et al., 1991 based on data from Kalaska et al., 1983.)

Noise limits discrimination performance

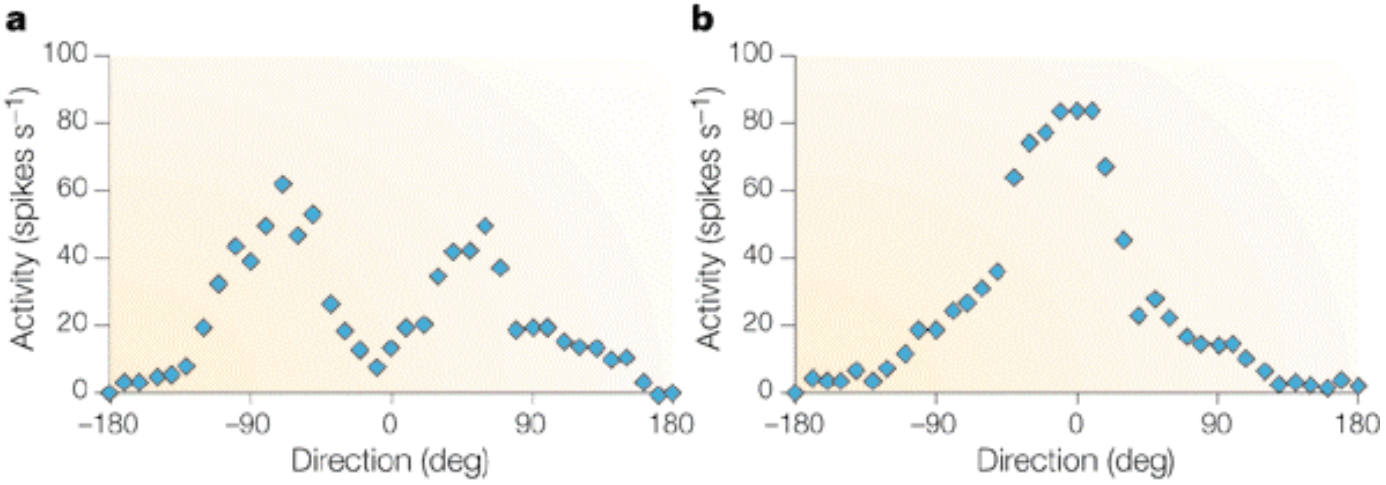

a: population response to two stimuli with different directions

b: population response to ***two*** stimuli with similar directions

Optimal Decoding

Use results from statistical estimation theory to infer a theoretical bound on how well ***any*** decoder can perform.

Key theoretical result: Cramer-Rao bound

- Consider a random variable X which depends on a parameter p

- Problem: try to recover p from the random observations of X

- Result: no matter what (unbiased) function we use to estimate p from observations $x_1 \ldots x_n$ of X, we will not be able to recover p perfectly… Rather, we will recover p up to a given accuracy bound (the Cramer-Rao bound) which we can compute formally for any unbiased estimator.

Of all possible estimators, the Maximum Likelihood estimator comes very close to this ideal bound.

Note: unbiased means that the estimator will on average yield p, not 0.9*p or 1.1*p

Population Code

a: tuning curves of neural population

b: responses of all neurons to one stimulus

c: population vector

d: maximum-likelihood estimate

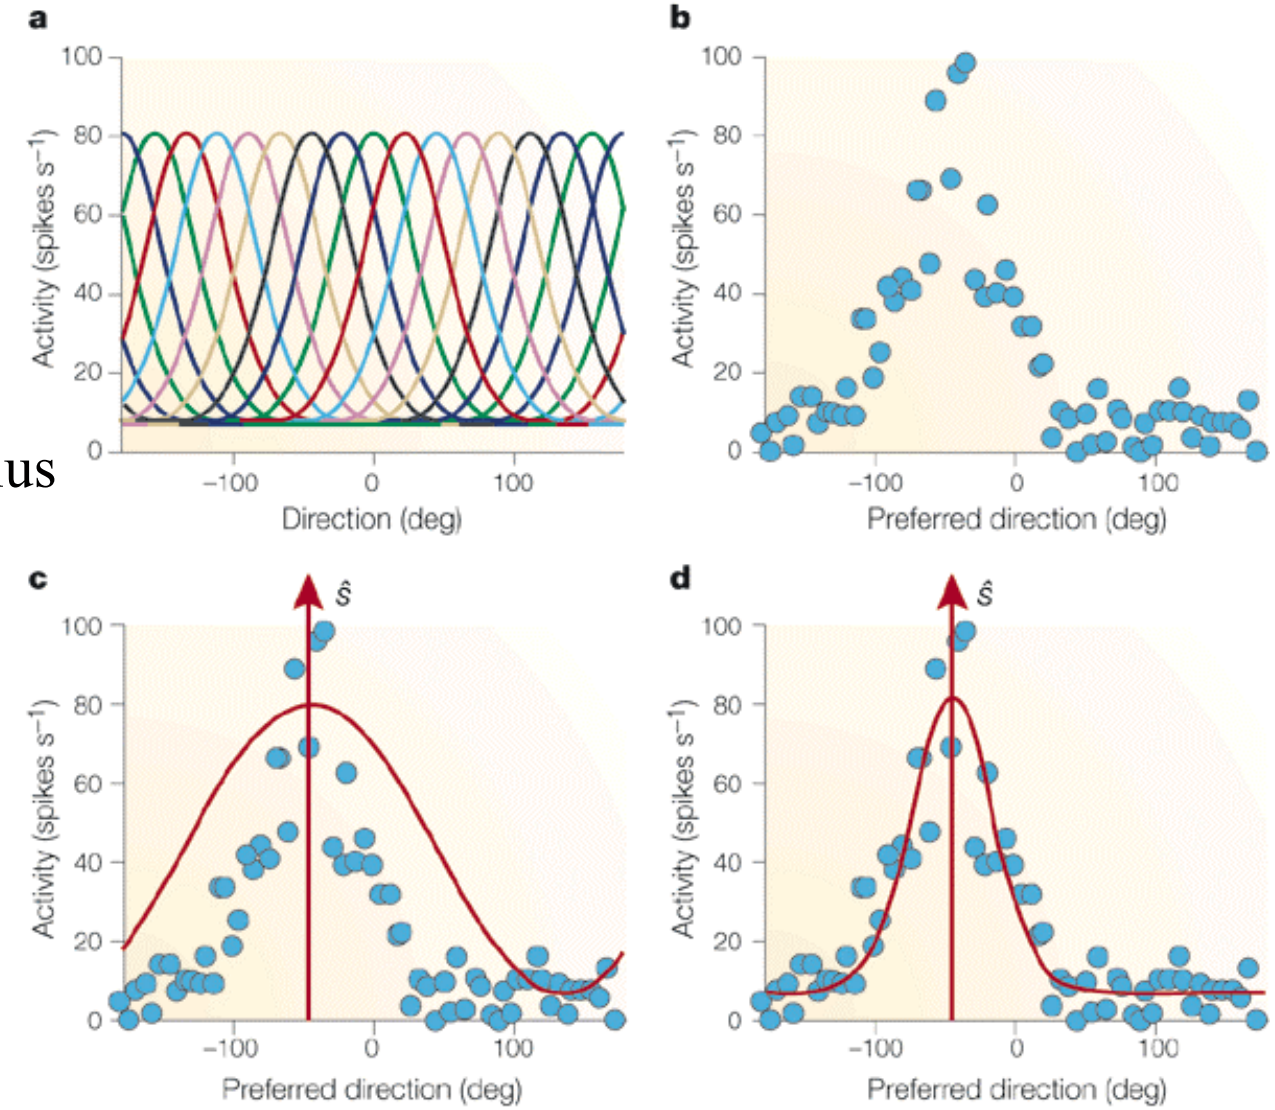

Neural implementations

A close approximation to the Maximum-Likelihood estimator can be built using a neural network (Pouget et al., Nature Reviews Neuroscience, 2000).

Input: noisy hill (bottom)

Output: clean hill whose mean is the ML estimate of stimulus orientation.

Difficulty: choice of weights depends on width of input hill…

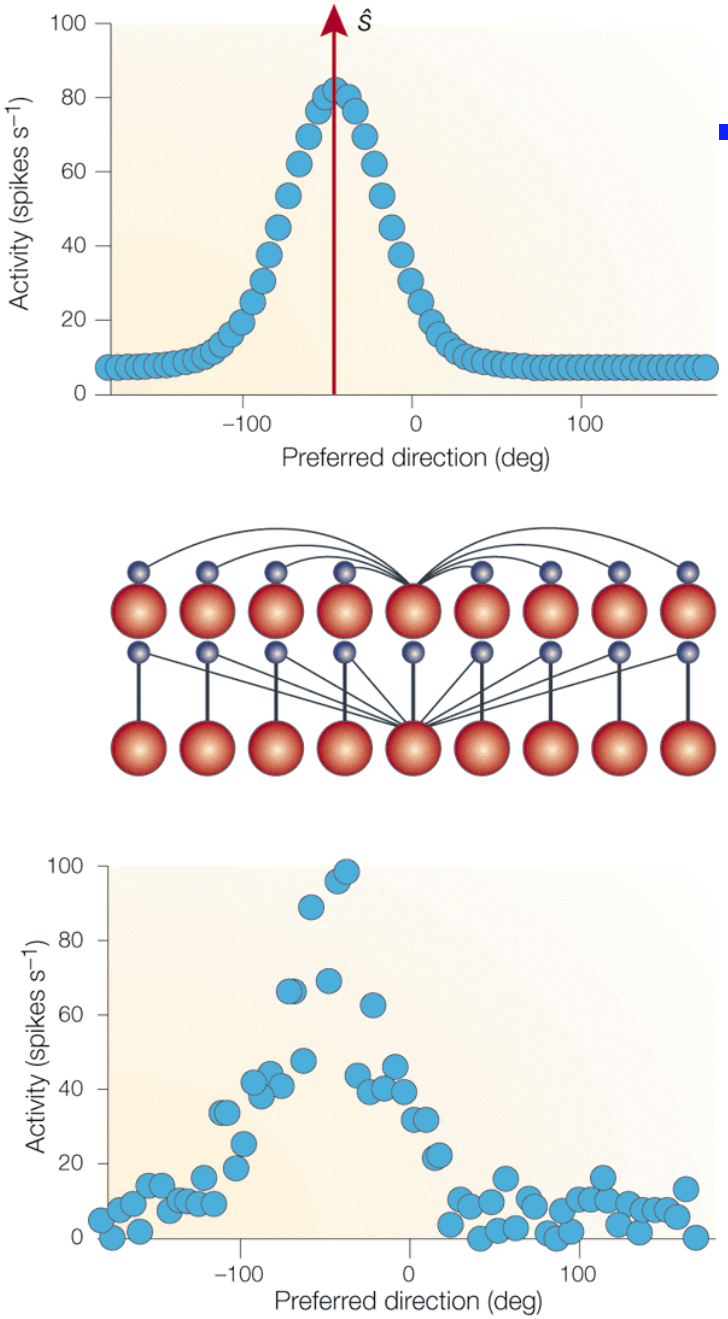

Case study: optimal estimation

Problem: record psychophysical data but simulate neurons; how can we predict performance of the whole observer based on responses of a few neurons?

Solution: train observers, and assume that they will approach an optimal decoder. Then do not worry about the implementation of the decoder itself. Rather, just use the result that it will perform close to the theoretical limit given by the Cramer-Rao bound.

Model architecture

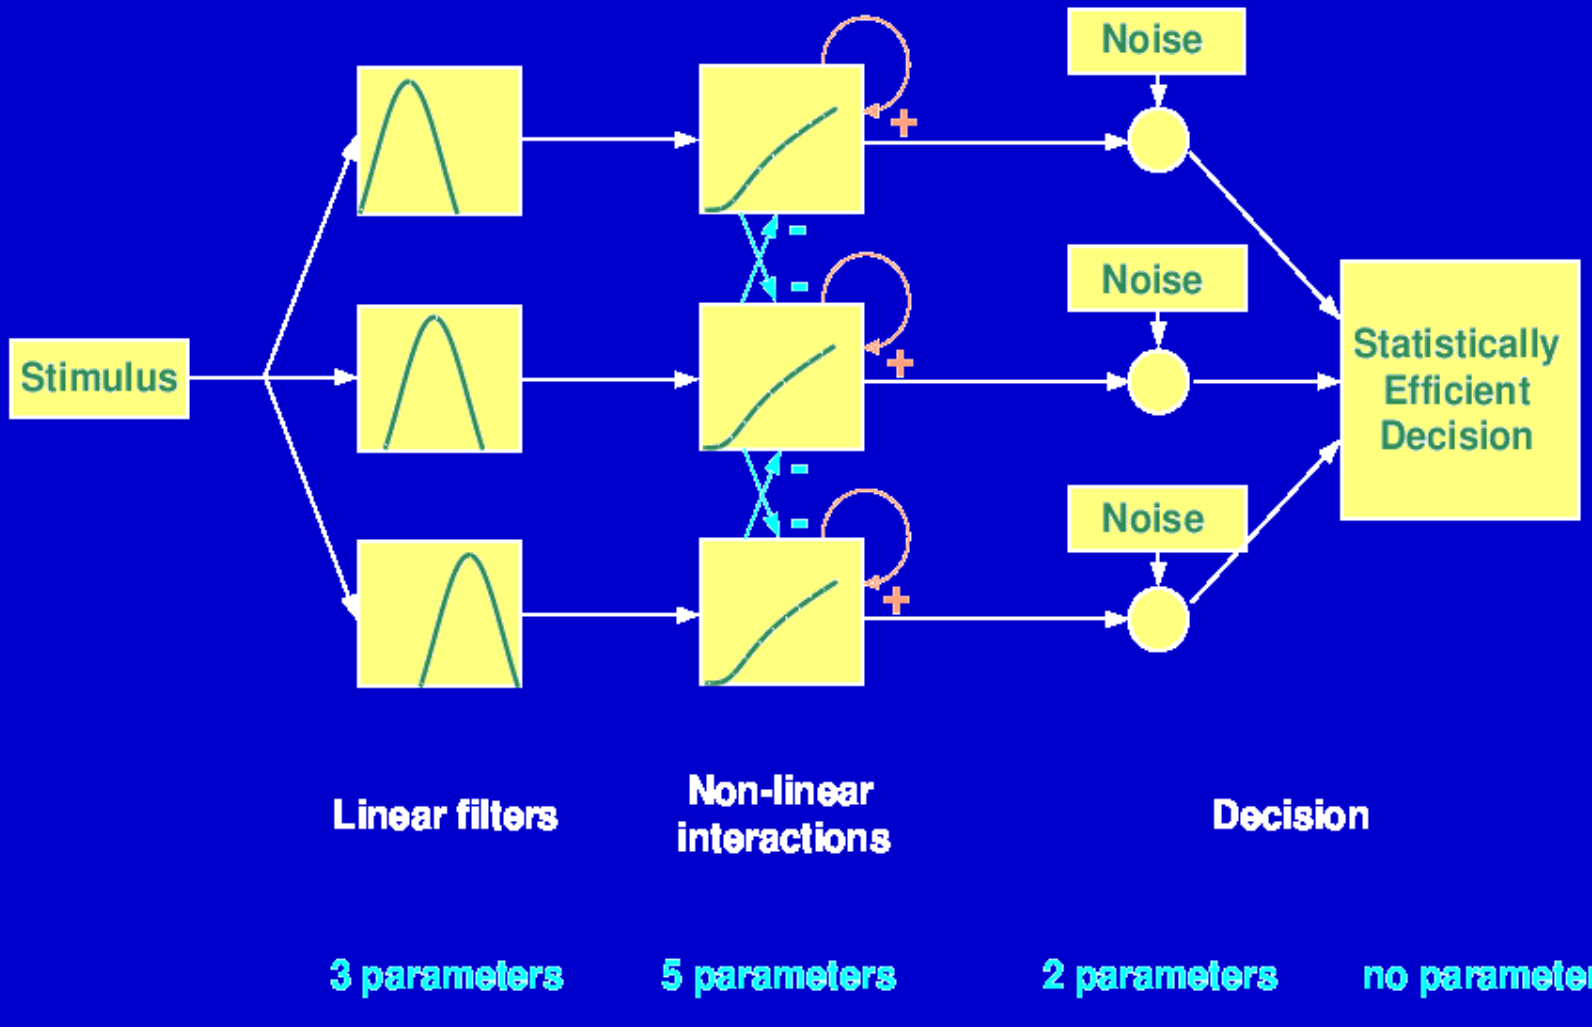

Linear filters

Separable in spatial period λ and orientation θ

Gaussian tuning curves $(\sigma_\lambda, \sigma_\theta)$

Quadrature pairs $Even_{\theta,\lambda}, Odd_{\theta,\lambda}$

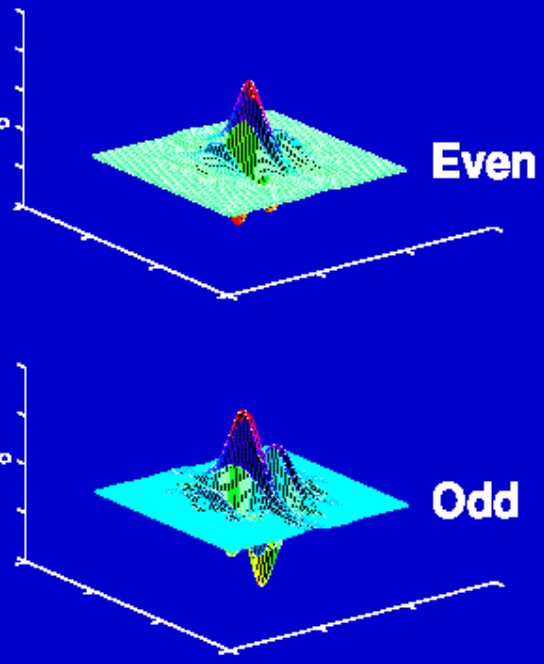

$$E_{\theta,\lambda} = \sqrt{\left(Im \star Even_{\theta,\lambda}\right)^2 + \left(Im \star Odd_{\theta,\lambda}\right)^2}$$

Non-linear interactions

Self excitation and divisive inhibition:

$$R_{\theta,\lambda} = \frac{E^{G}_{\theta,\lambda}}{S_\lambda + \sum_{\theta',\lambda'} W_{\theta',\lambda'} E^{H}_{\theta',\lambda'}}$$

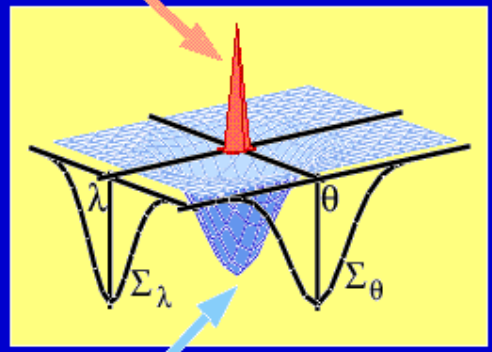

Poisson$^{\alpha}$ noise model: $V^2_{\theta,\lambda} = R^{\alpha}_{\theta,\lambda}$

variance

Properties

no pooling

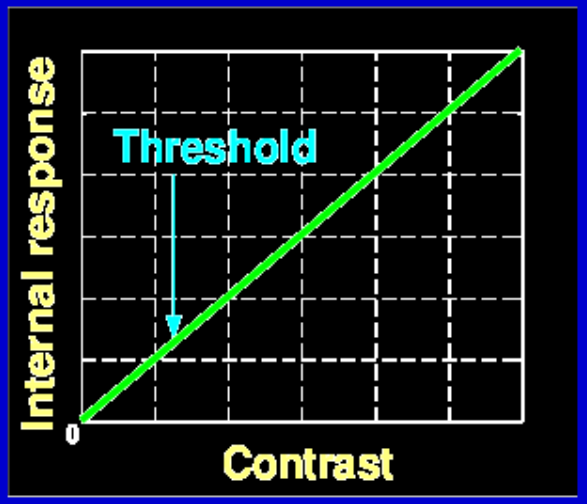

pooling

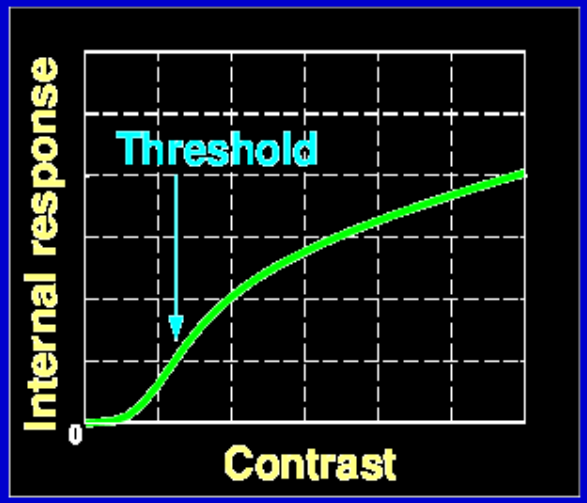

Non-linear transducer function

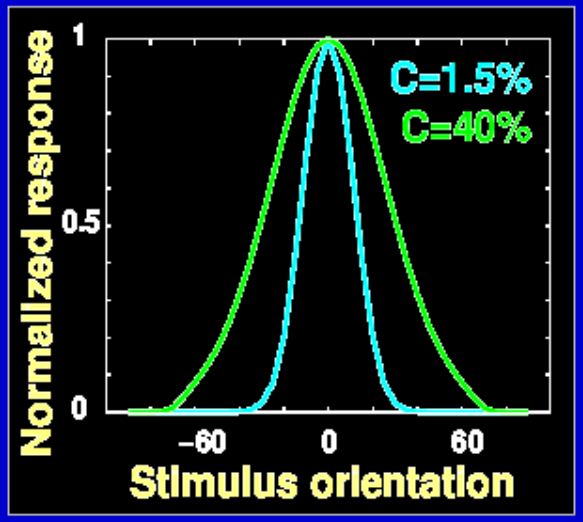

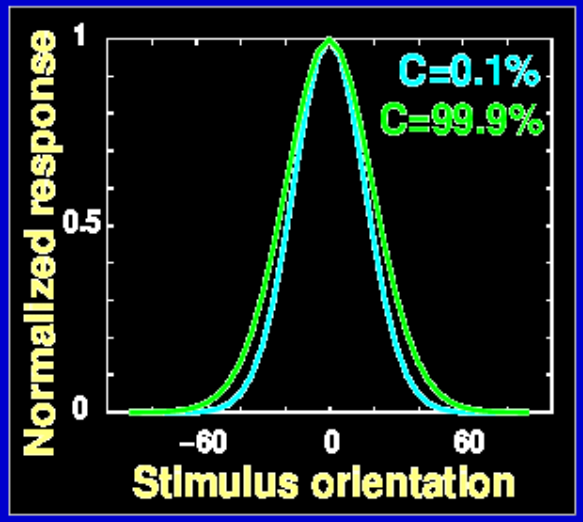

Contrast-independent tuning

Threshold prediction

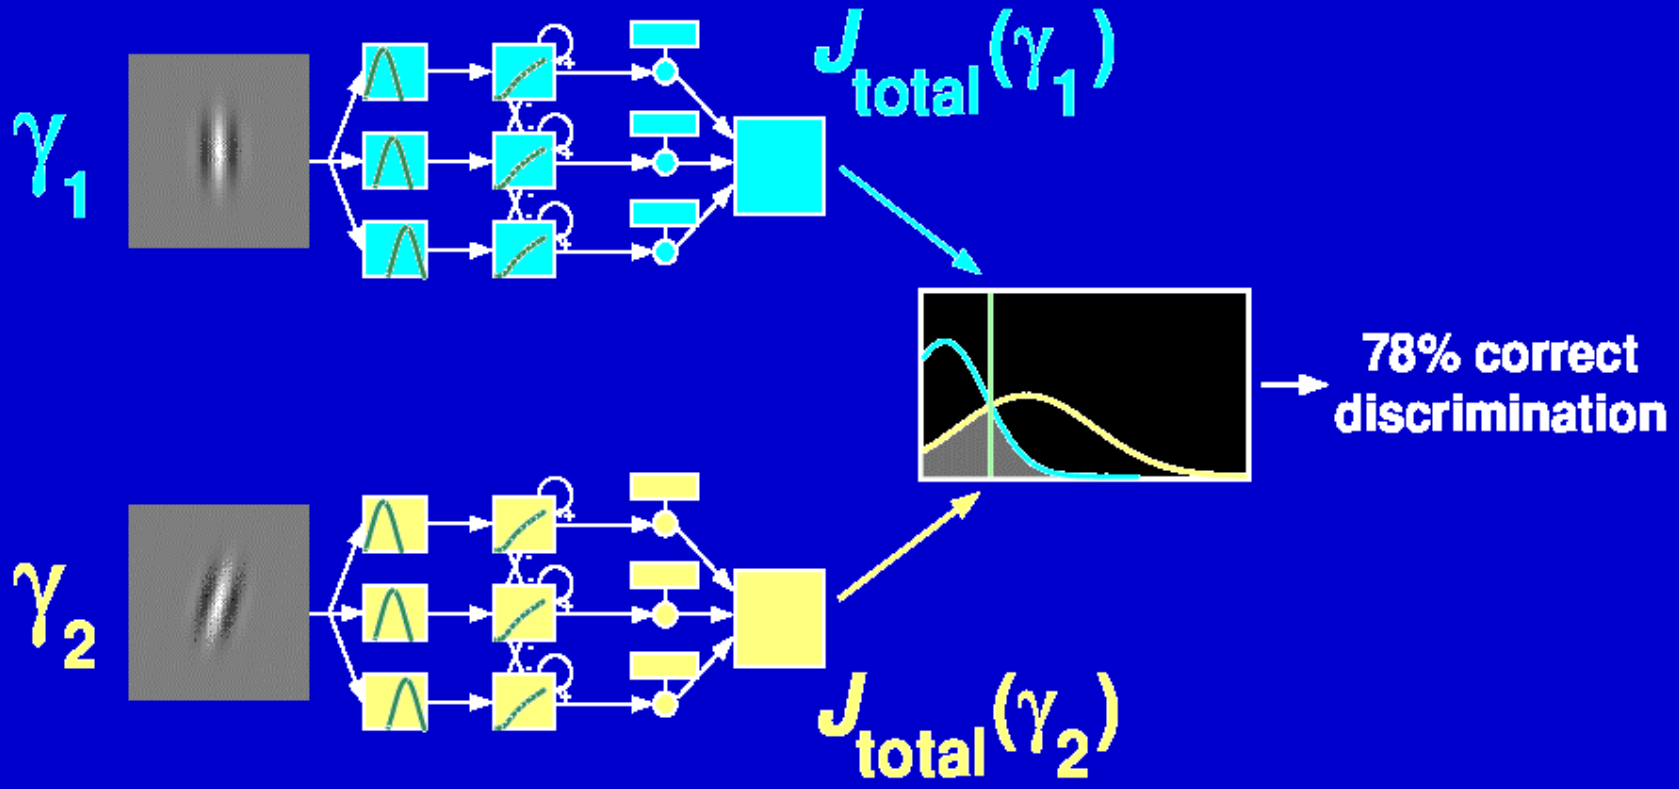

Decision strategy

One mechanism $R_{\theta,\lambda}$:

Stimulus parameter γ (contrast, orient., s.f.)

Assume existence of unbiased efficient statistic $T(R_{\theta,\lambda}, \gamma)$

then, **mean(T) = γ** and **var(T) = $1/J_{\theta,\lambda}$**

with:

$$J_{\theta,\lambda} = \frac{1}{V_{\theta,\lambda}^2}\left[\left(\frac{\delta R_{\theta,\lambda}}{\delta\gamma}\right)^2 + 2\left(\frac{\delta V_{\theta,\lambda}}{\delta\gamma}\right)^2\right]$$

Population of mechanisms:

$$J_{total} = \sum_{\theta,\lambda} J_{\theta,\lambda}$$

Fisher information

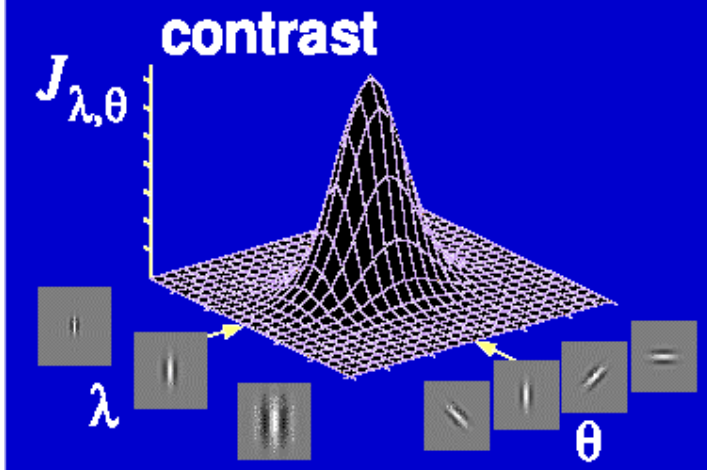

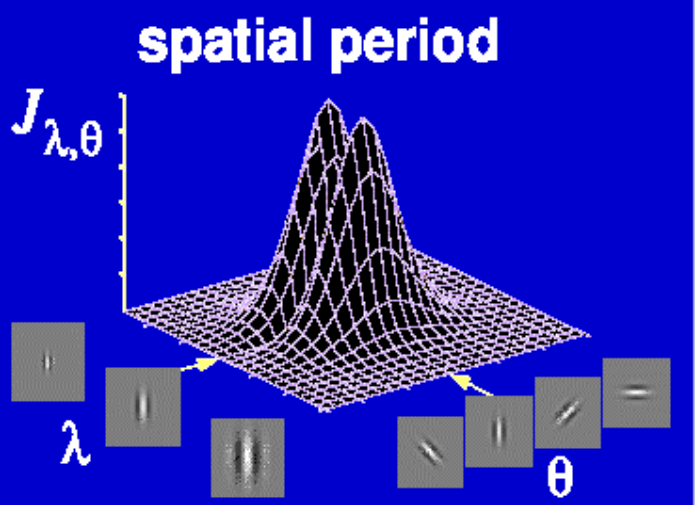

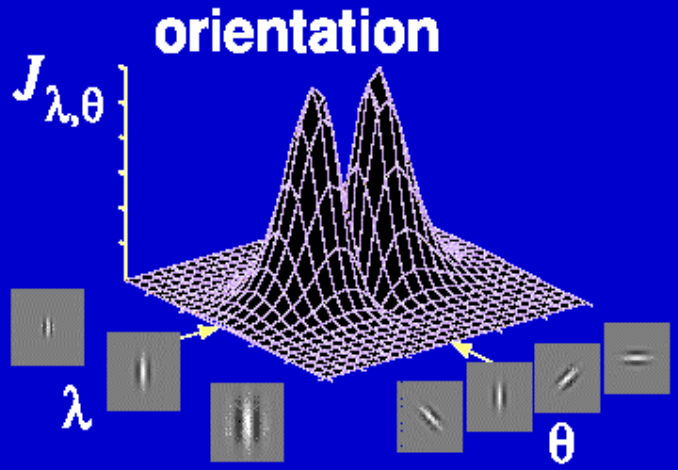

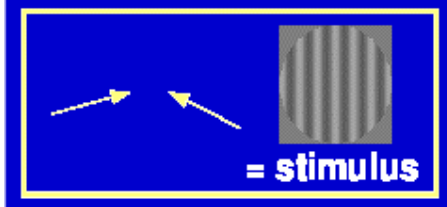

Ideal observer discrimination

Discriminate between stimulus 1 (γ_1) and stimulus 2 (γ_2)

Stimulus 1: $\mu_1 = \gamma_1$,

$\sigma_1^2 = 1/J_{total}(\gamma_1)$

Stimulus 2: $\mu_2 = \gamma_2$,

$\sigma_2^2 = 1/J_{total}(\gamma_2)$

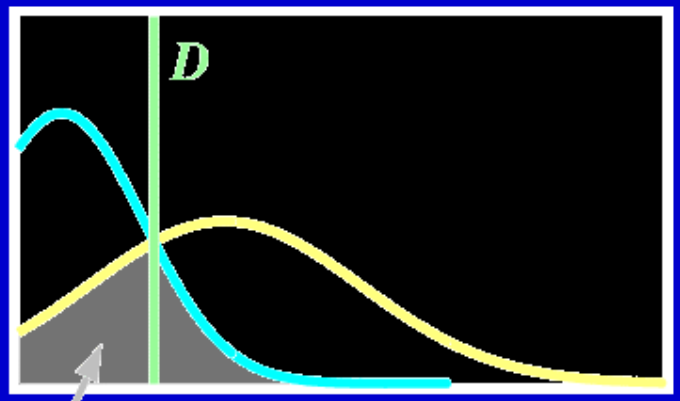

Decision criterion:

$$D = \frac{2\mu_2\sigma_1^2 - 2\mu_1\sigma_2^2 - \sqrt{(2\mu_1\sigma_2^2 - 2\mu_2\sigma_1^2)^2 - 4(\sigma_1^2 - \sigma_2^2)(\mu_2^2\sigma_1^2 - \mu_1^2\sigma_2^2 - 2\sigma_1^2\sigma_2^2 \log\frac{\sigma_1}{\sigma_2})}}{2(\sigma_1^2 - \sigma_2^2)}$$

Performance Yes/No task:

$$Performance = 1 - \frac{1}{4}erfc\left(\frac{\mu_2 - D}{\sigma_2\sqrt{2}}\right) - \frac{1}{4}erfc\left(\frac{D - \mu_1}{\sigma_1\sqrt{2}}\right)$$

How realistic is this optimal decoder?

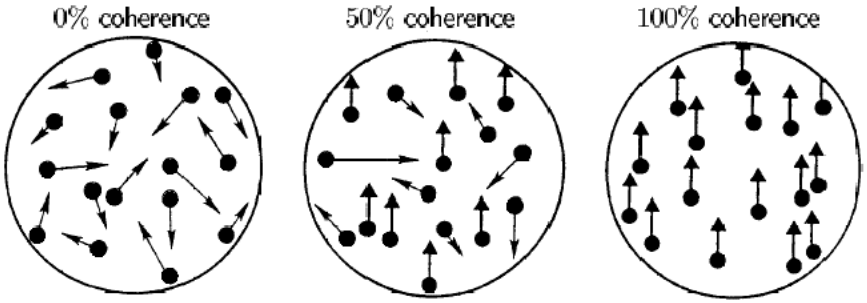

Figure 3.1: The moving random-dot stimulus for different levels of coherence. The visual image consists of randomly placed dots that jump every 45 ms according to the scheme described in the text. At 0% coherence the dots move randomly. At 50% coherence, half the dots move randomly and half move together (upwards in this example). At 100% coherence all the dots move together. (Adapted from Britten et al., 1992.)

In some studies, single-neuron firing predicts very well the behavior of the animal!

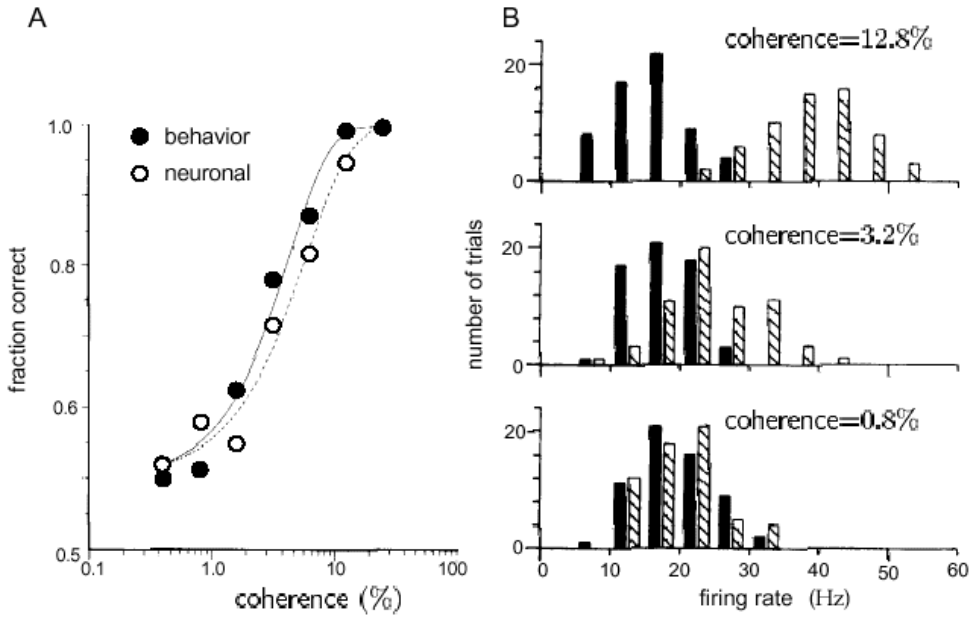

Figure 3.2: Behavioral and electrophysiological results from a random dot motion discrimination task. A) The filled circles show the fraction of correct discriminations made by the monkey as a function of the degree of coherence of the motion. The open circles show the discrimination accuracy that an ideal observer could achieve on the analogous two-alternative forced choice discrimination task given the neural responses. B) Firing rate histograms for three different levels of coherence. Hatched rectangles show the results for motion in the plus direction and solid rectangles for motion in the minus direction. The histograms have been thinned for clarity so that not all the bins are shown. (Adapted from Britten et al., 1992 .)

Threshold prediction

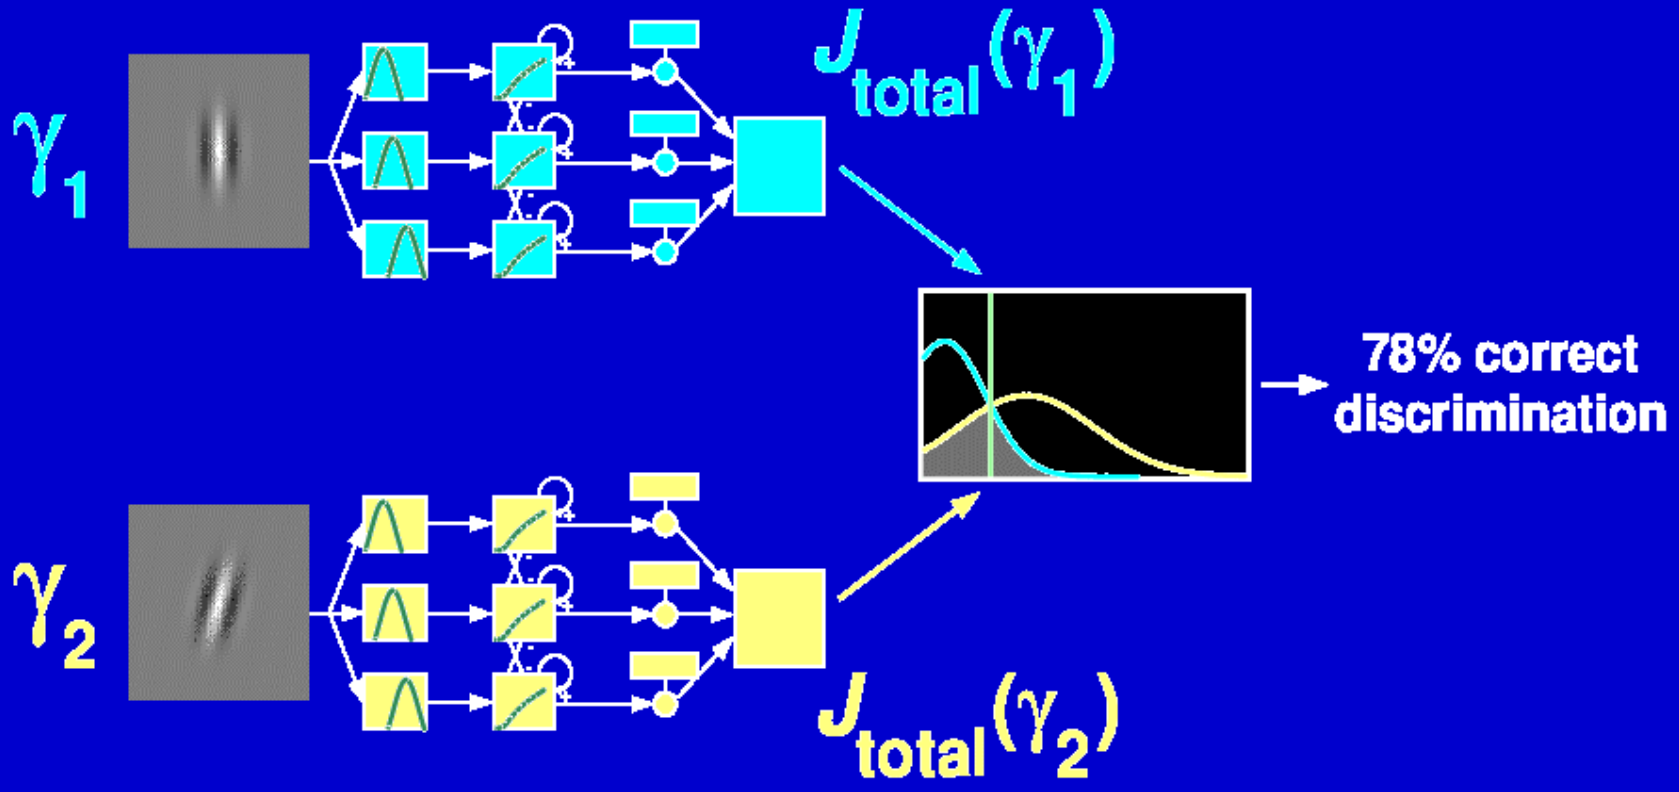

Separate fits to data

(10 parameters, 32 datapoints)

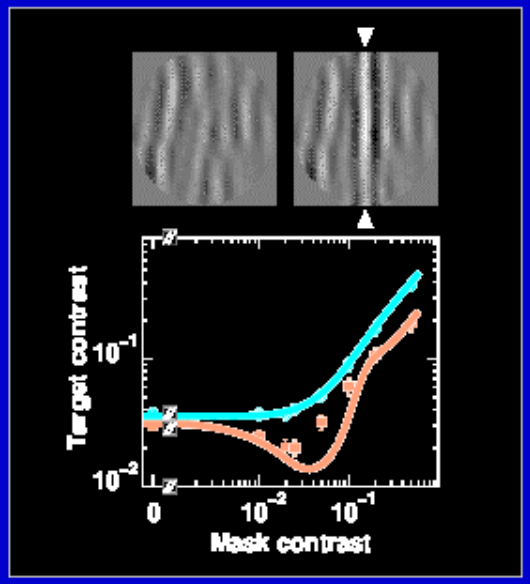

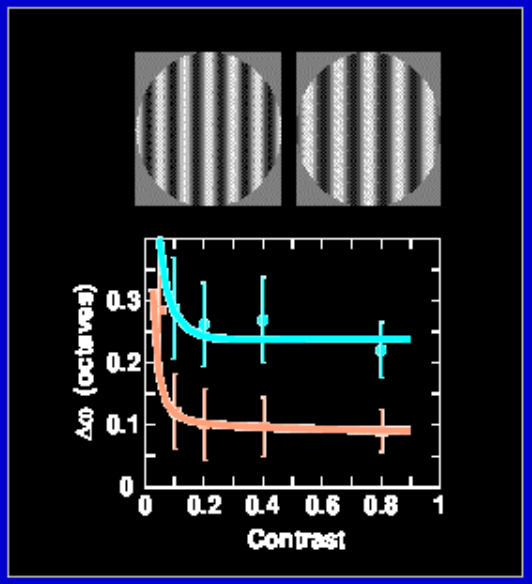

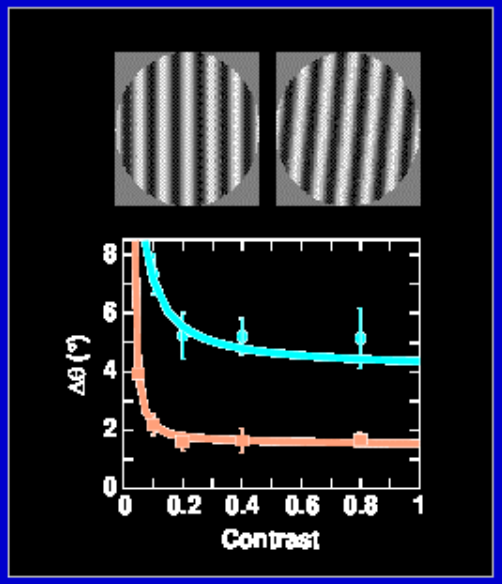

- Fully attended
- Poorly attended

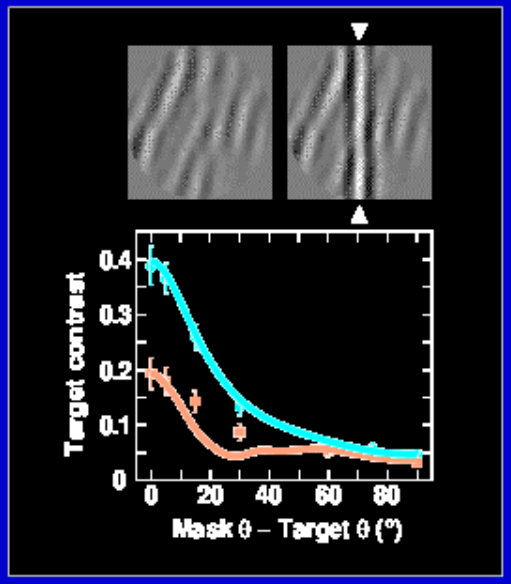

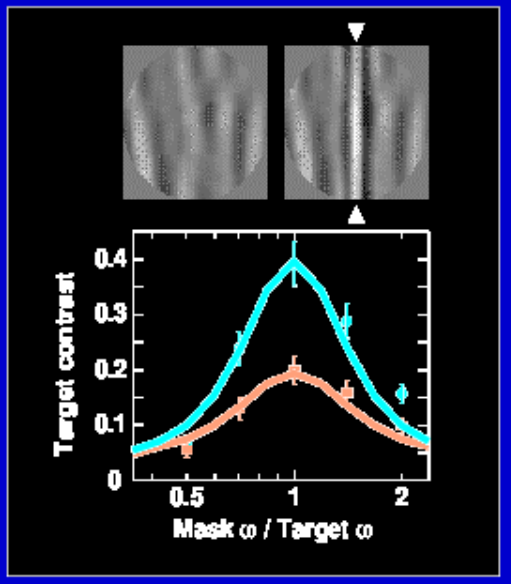

Population Codes May Not be Obvious!

Sometimes responses to different stimuli will have non-obvious characteristics:

e.g., variations in phase of oscillatory response…

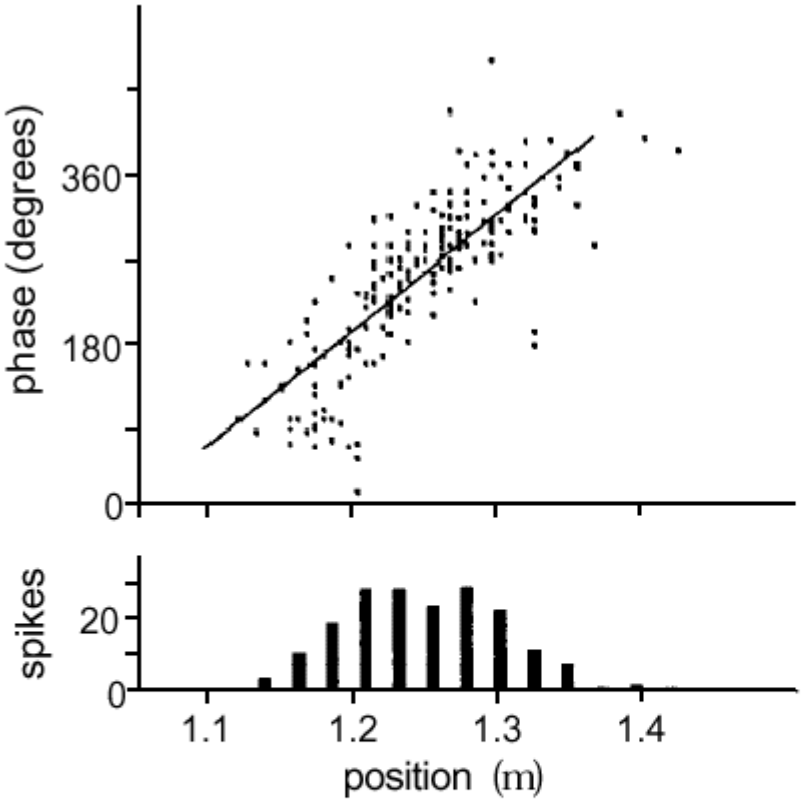

Figure 1.18: Position versus phase for a hippocampal place cell. Each dot in the upper figure shows the phase of the theta rhythm plotted against the position of the animal at the time when a spike was fired. The linear relation shows that information about position is contained in the relative phase of firing. The lower plot is a conventional place field tuning curve of spike count versus position. (Adapted from O'Keefe and Recce, 1993.)

Time-Dependent Coding

Neural responses often
are not restricted to a single
discharge event just following
a stimulus,

but may extend over time
in complex fluctuations…

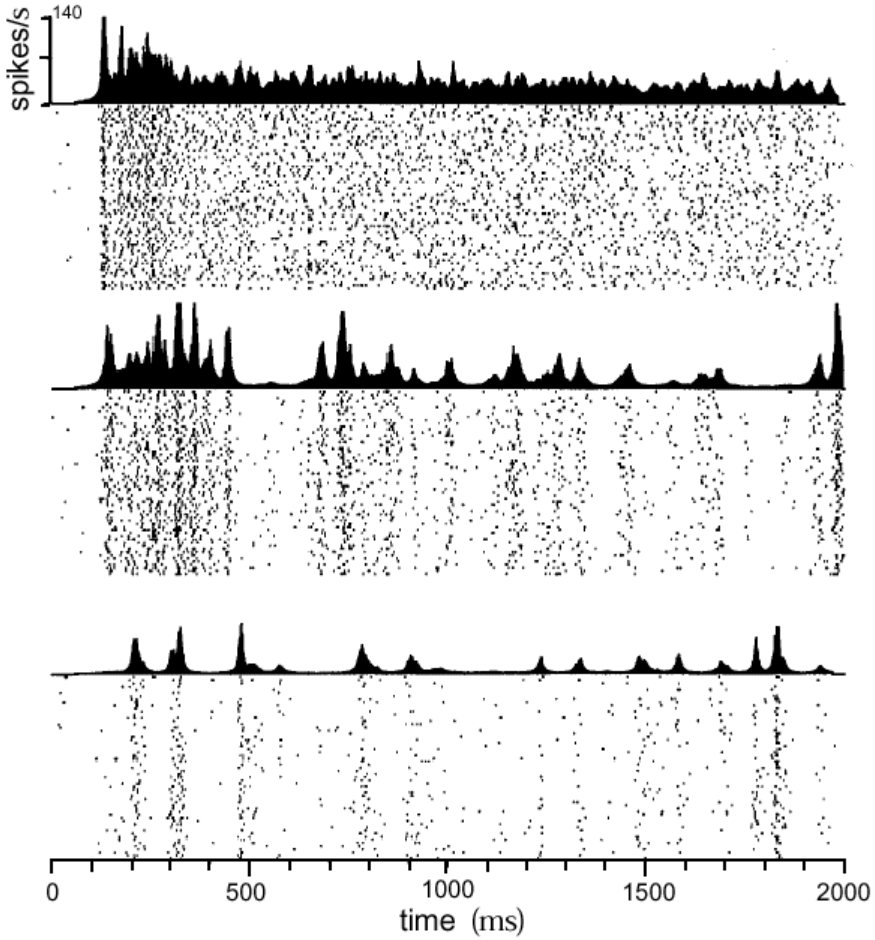

Figure 1.19: Time-dependent firing rates for different stimulus parameters. The rasters show multiple trials during which an MT neuron responded to the same moving random dot stimulus. Firing rates, shown above the raster plots, were constructed from the multiple trials by counting spikes within discrete time bins and averaging over trials. The three different results are from the same neuron but using different stimuli. The stimuli were always patterns of moving random dots but the coherence of the motion was varied (see chapter 3 for more information about this stimulus). (Adapted from Bair and Koch, 1996.)

What more can we do?

One good thing to do is to pre-process the data such as to eliminate redundancy and represent it in its most "natural" basis.

How can we do that?

Principal Component Analysis

Independent Component Analysis

Clustering techniques

Principal Component Analysis

Goal:

transforms a set of correlated variables into a smaller set of uncorrelated variables.

Approach:

find eigenvectors of covariance matrix and operate a change of basis along a subset of those vectors that correspond to non-negligible eigenvalues.

Properties:

- removes redundancy in data
- finds "true dimensionality" of dataset
- linear transformation, and hence preserves distance

PCA example

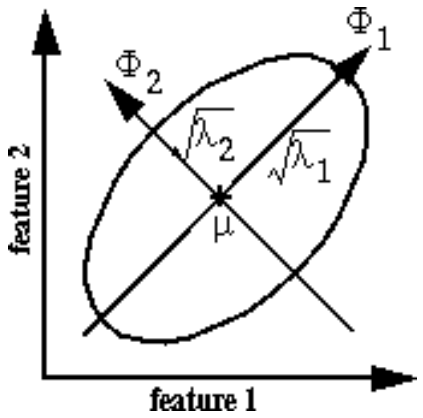

Original data is skewed along one axis.

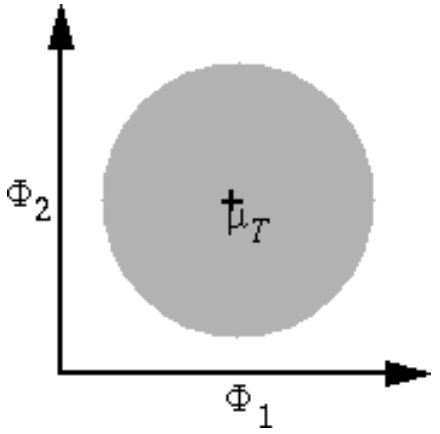

Transformed data is unskewed.

Independent Components Analysis

Goal:

extract independent sources given only observed data as a mixture of the unknown sources.

Approach:

like PCA, ICA also decorrelates the signals and reduces dimensionality. The main difference is that the new coordinate system may not be orthogonal. Algorithm works by maximizing mutual information between assumed input sources and observed output data.

Properties:

same general properties as PCA, with the additional freedom of possibly non-orthogonal principal directions.

ICA outline

Consider an ensemble of data vectors $\{\underline{x}_1 \quad , \underline{x}_r\}$

Find a new basis $\underline{u} = W \cdot \underline{x} + \underline{w}$

Lets make u zero-mean, unit variance and uncorrelated:

$$\text{i.e.} \quad \langle u \rangle = 0 \quad \langle uu^t \rangle = I$$

Solving for W in terms of x (for simplicity assume $\langle x \rangle = 0$)

$$\langle Wx(Wx)^t \rangle = I$$

$$W\langle xx^T \rangle W^T = I$$

$$W^T W = \langle xx^T \rangle^{-1}$$

Here, W is UNDER-DETERMINED.

- We can decide that W should be orthogonal => PCA

PCA vs. ICA: Example PCA Result

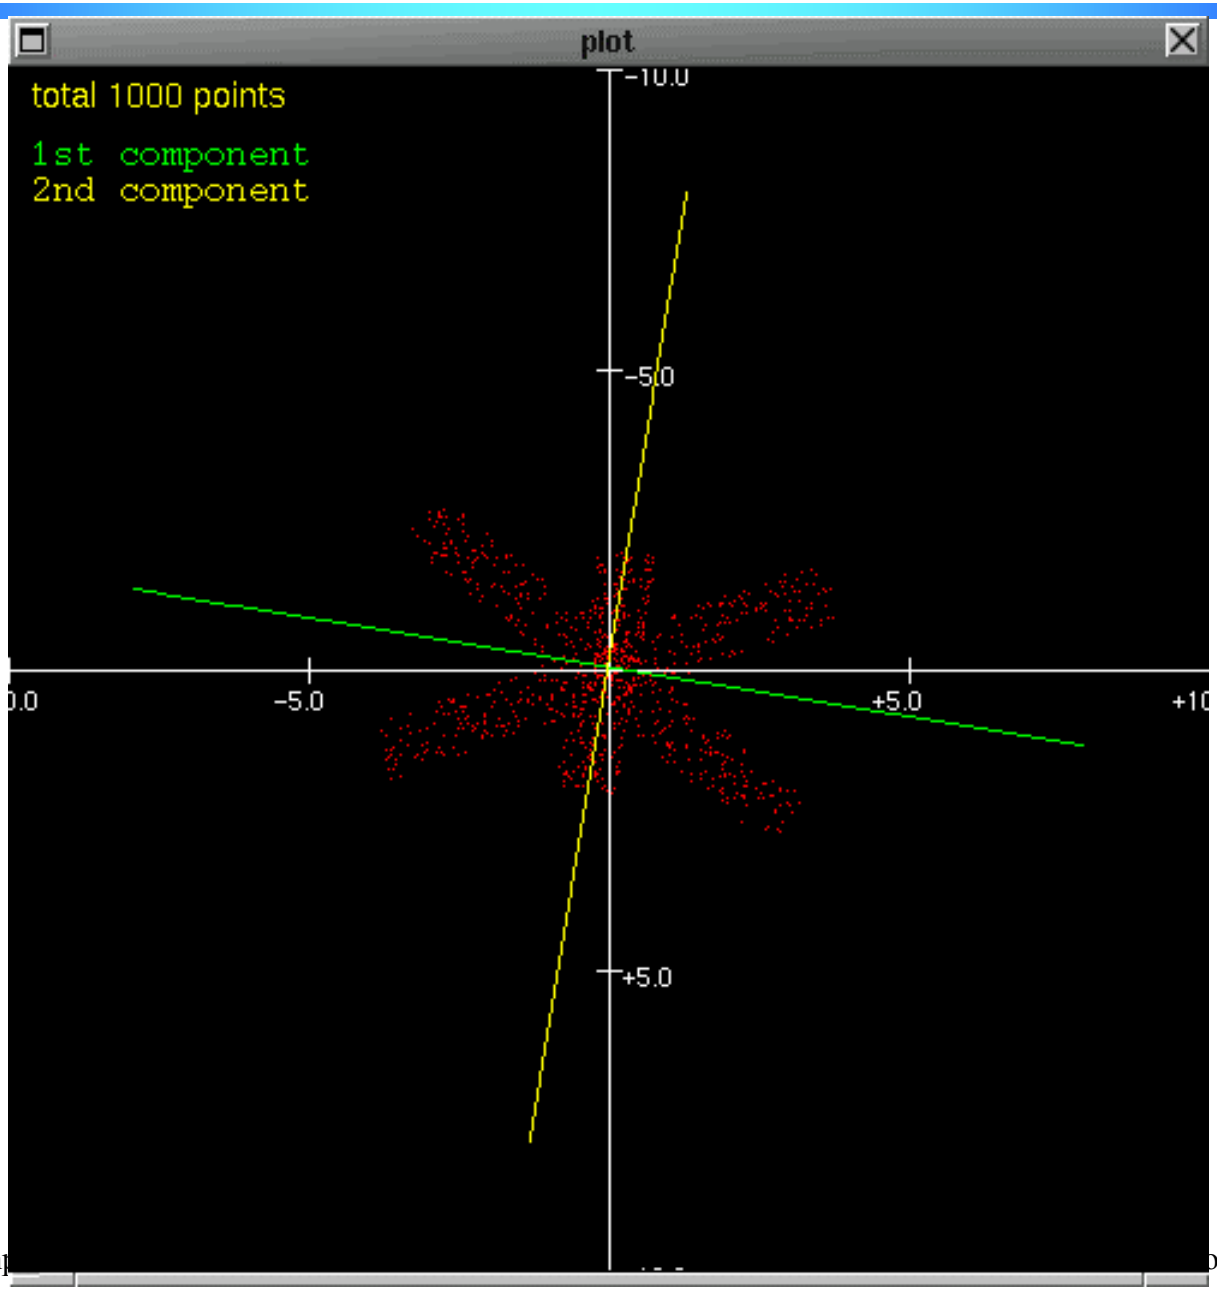

PCA vs. ICA: Example ICA Result

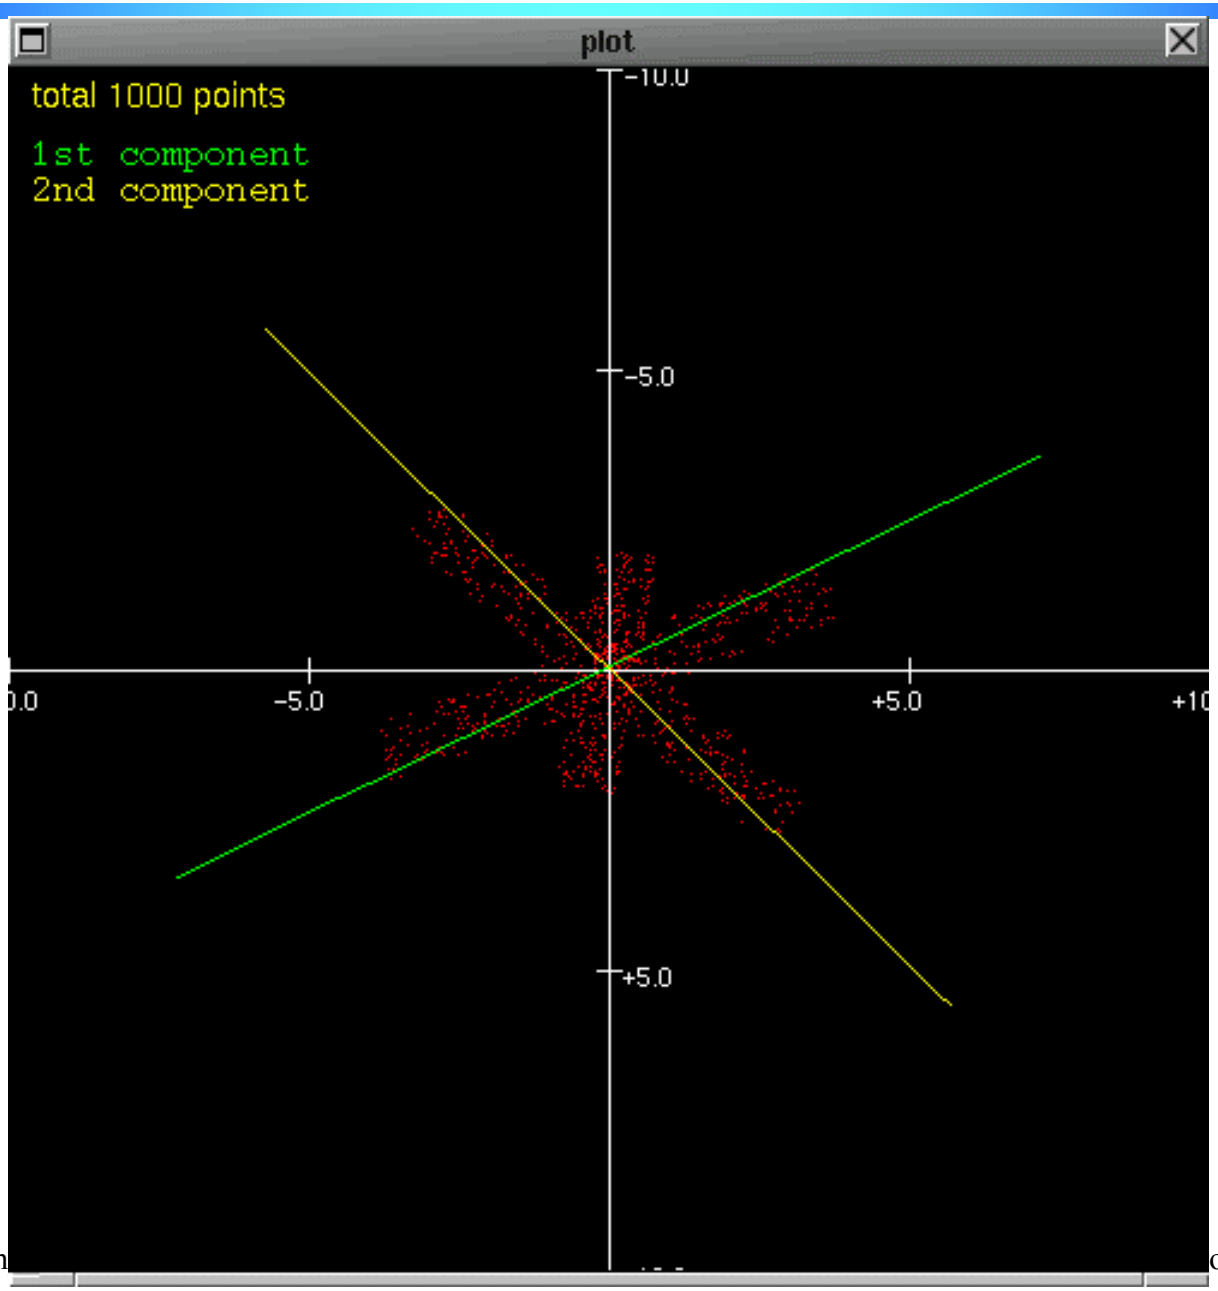

Implications for Computer Architectures

Population codes are omnipresent in cortex…

Advantages:

very robust to noisy signals

very robust to isolated failing neurons

Drawback:

need a decoder! but that may not be a real drawback if the decoder has a clear physical meaning and is fast and easy to compute: e.g., population vector.