Functional models of neural computation

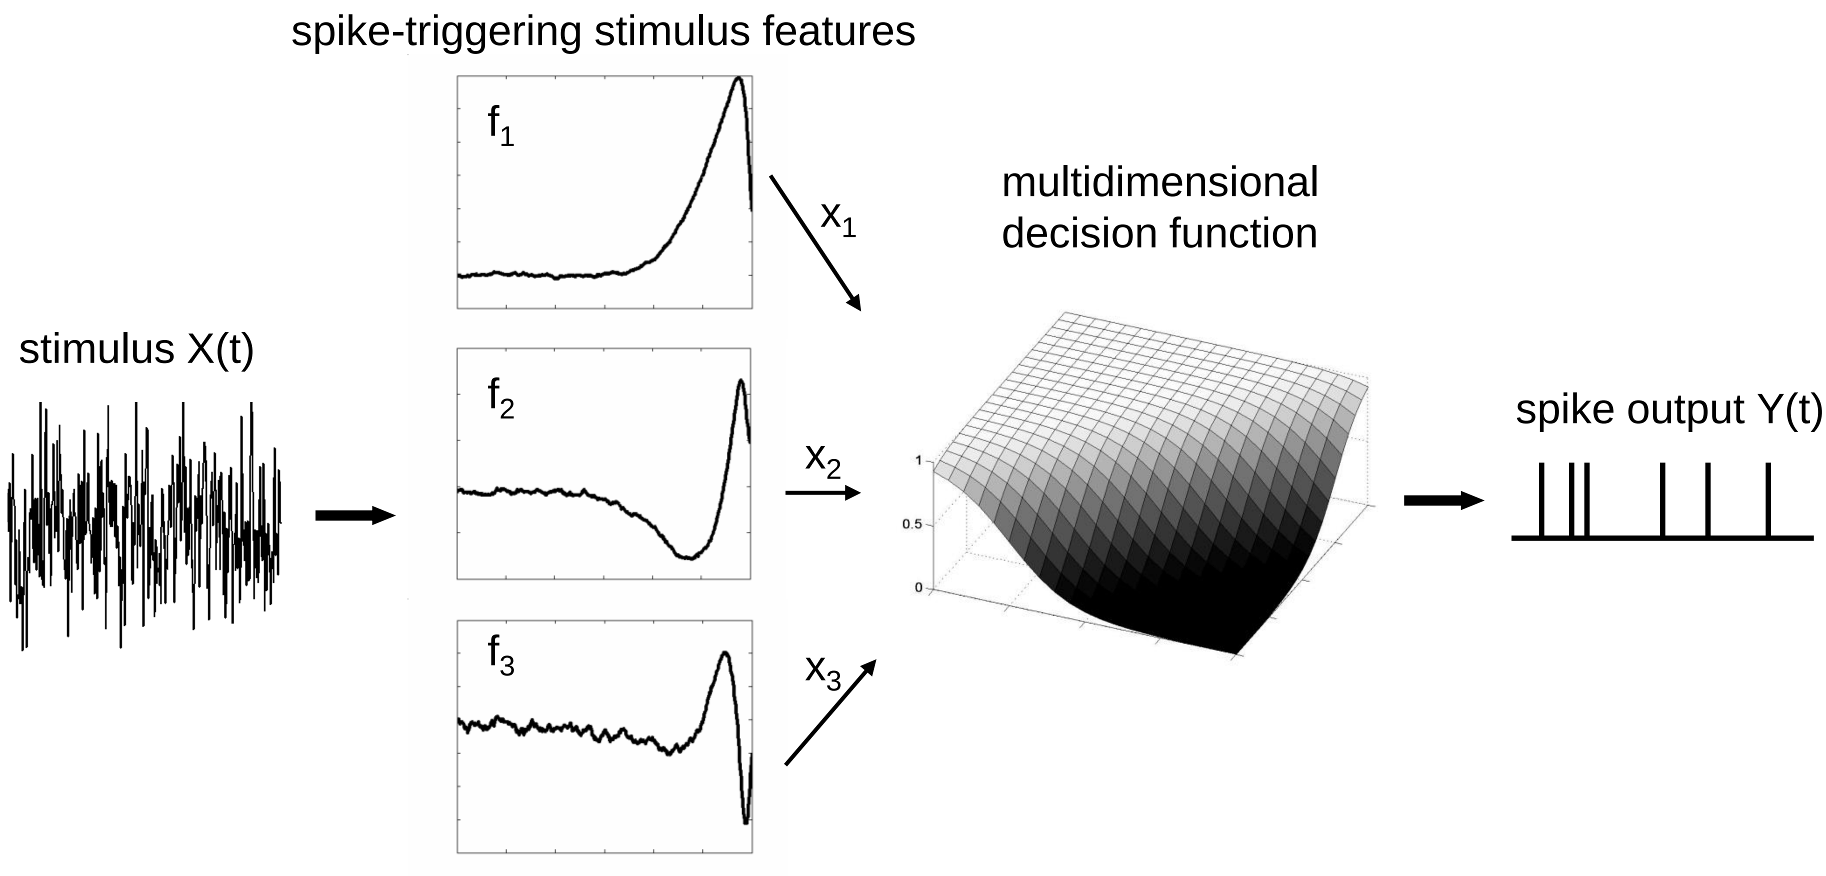

Covariance analysis

Let's develop some intuition for how this works: the Keat model

Keat, Reinagel, Reid and Meister, Predicting every spike. Neuron (2001)

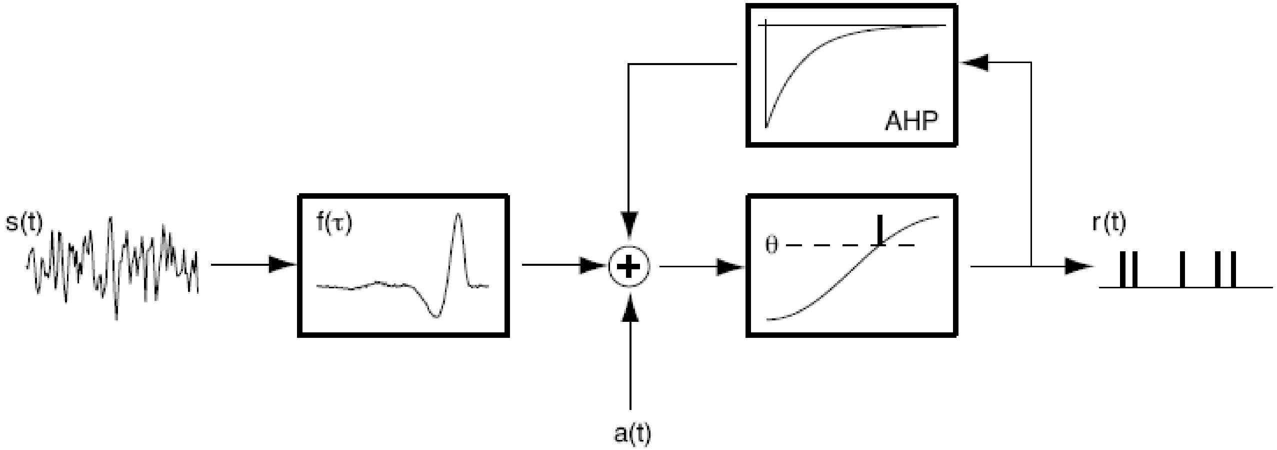

- Spiking is controlled by a single filter
- Spikes happen generally on an upward threshold crossing of the filtered stimulus

➔ expect 2 modes, the filter F(t) and its time derivative F'(t)

Given a set of data, want to find the best reduced dimensional description.

The data are the set of stimuli that lead up to a spike, $S_n(t)$, where t = 1, 2, 3, …., D

Variance of a random variable = $< (X\text{-mean}(X))^2 >$

Covariance = $< (\mathbf{X} - \text{mean}(\mathbf{X}))^T (\mathbf{X} - \text{mean}(\mathbf{X})) >$

Compute the *difference matrix* between covariance matrix of the spike-triggered stimuli and that of all stimuli

Find its *eigensystem* to define the dimensions of interest

Eigensystem:

any matrix M can be decomposed as

$$M = U\ V\ U^T,$$

where U is an orthogonal matrix;

V is a diagonal matrix, $\mathrm{diag}([\lambda_1, \lambda_2, .., \lambda_D])$.

Each eigenvalue has a corresponding eigenvector, the orthogonal columns of U.

The value of the eigenvalue classifies the eigenvectors as belonging to

column space = orthogonal basis for relevant dimensions

null space = orthogonal basis for irrelevant dimensions

We will project the stimuli into the column space.

This method finds an orthogonal basis in which to describe the data, and ranks each “axis” according to its importance in capturing the data.

Related to principal component analysis.

Functional basis set.

Example:

An auditory neuron is responsible for detecting sound at a certain frequency ω. Phase is not important.

The appropriate "directions" describing this neuron's relevant feature space are

$$\text{Cos}(\omega t) \text{ and } \text{Sin}(\omega t).$$

This will describe any signal at that frequency, independent of phase:

$\cos(A+B) = \cos(A)\cos(B) - \sin(A)\sin(B)$

→ $\cos(\omega t + \phi) = a\cos(\omega t) + b\sin(\omega t)$,

$$a = \cos(\phi),\ b = -\sin(\phi).$$

Note that $a^2 + b^2 = 1$; all such stimuli lie on a ring.

Covariance analysis

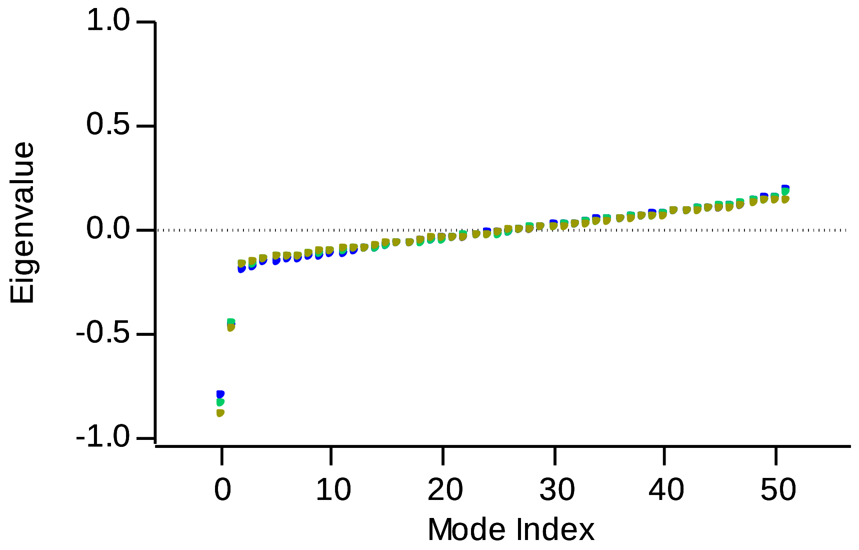

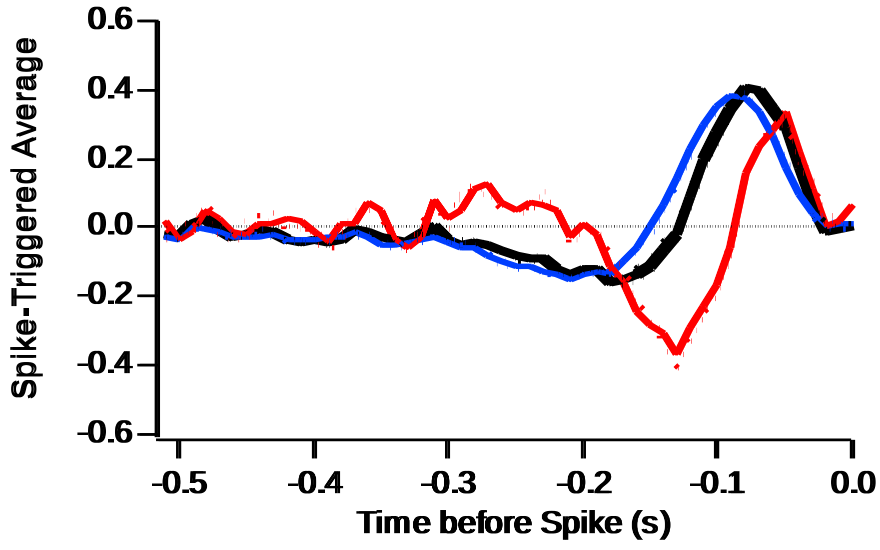

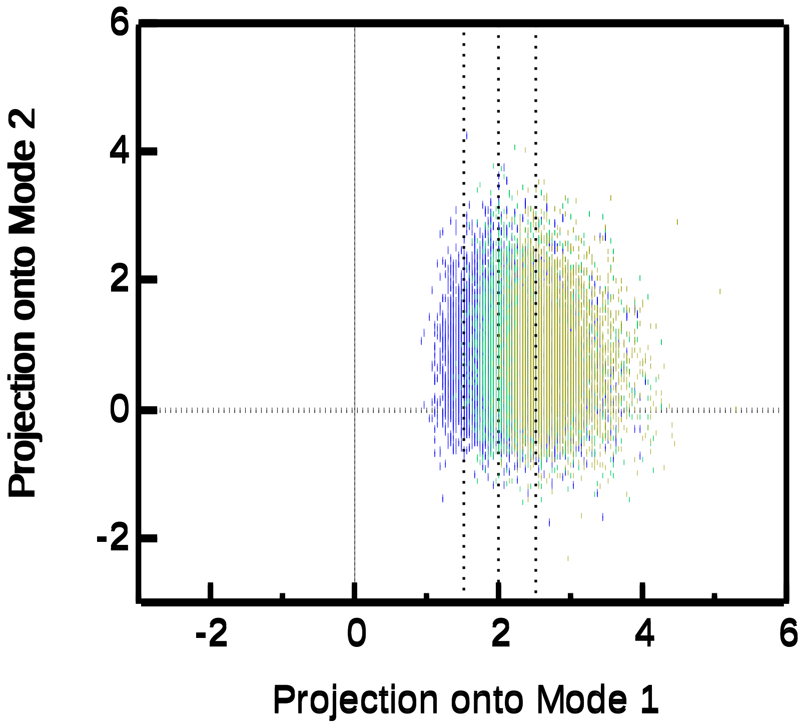

Covariance analysis

Let's try a real neuron: rat somatosensory cortex

(Ras Petersen, Mathew Diamond, SISSA)

Record from single units in barrel cortex

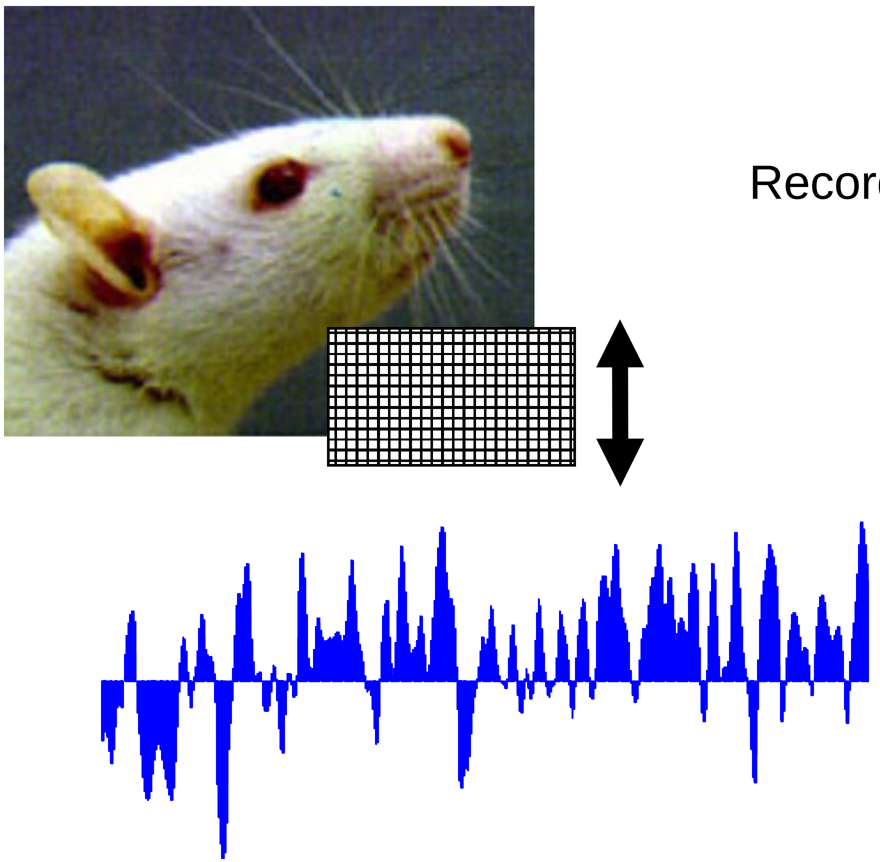

Covariance analysis

Spike-triggered average:

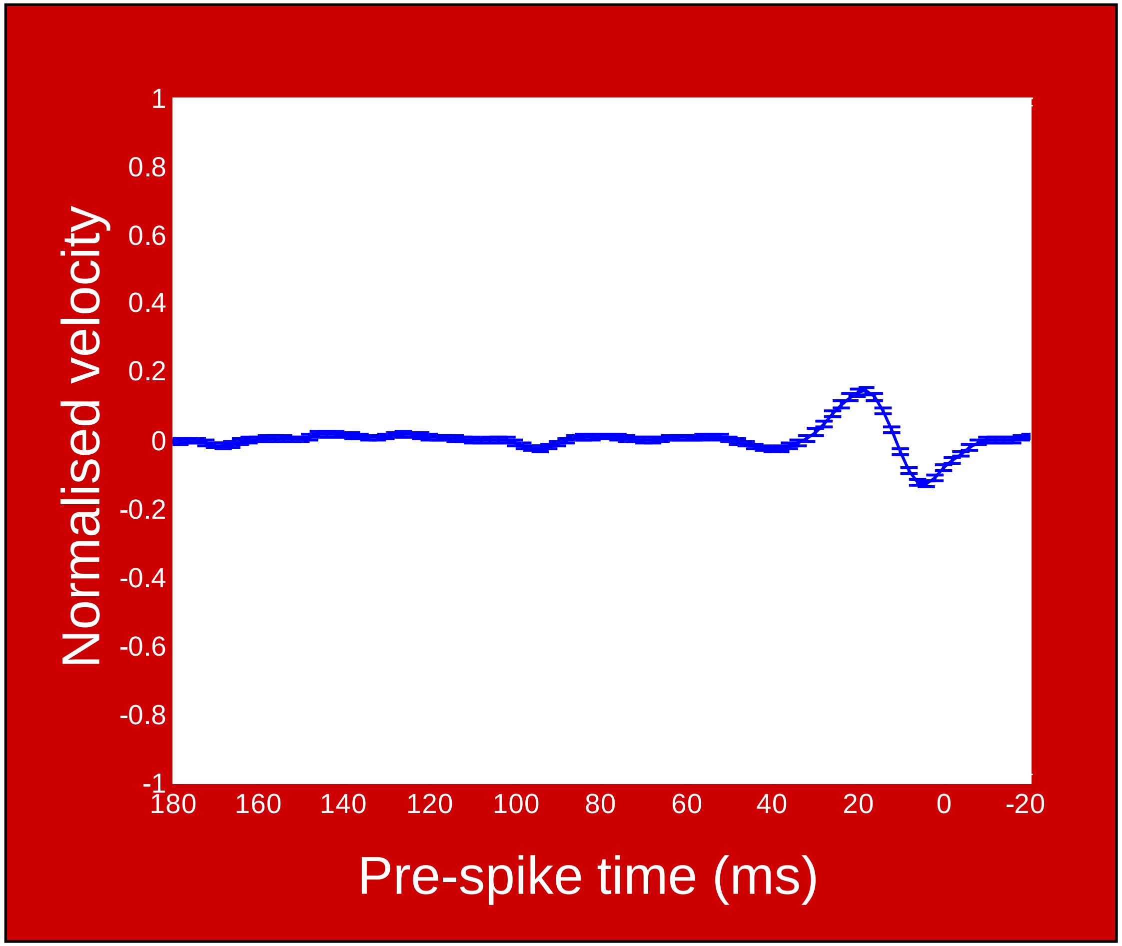

Covariance analysis

Is the neuron simply not very responsive to a white noise stimulus?

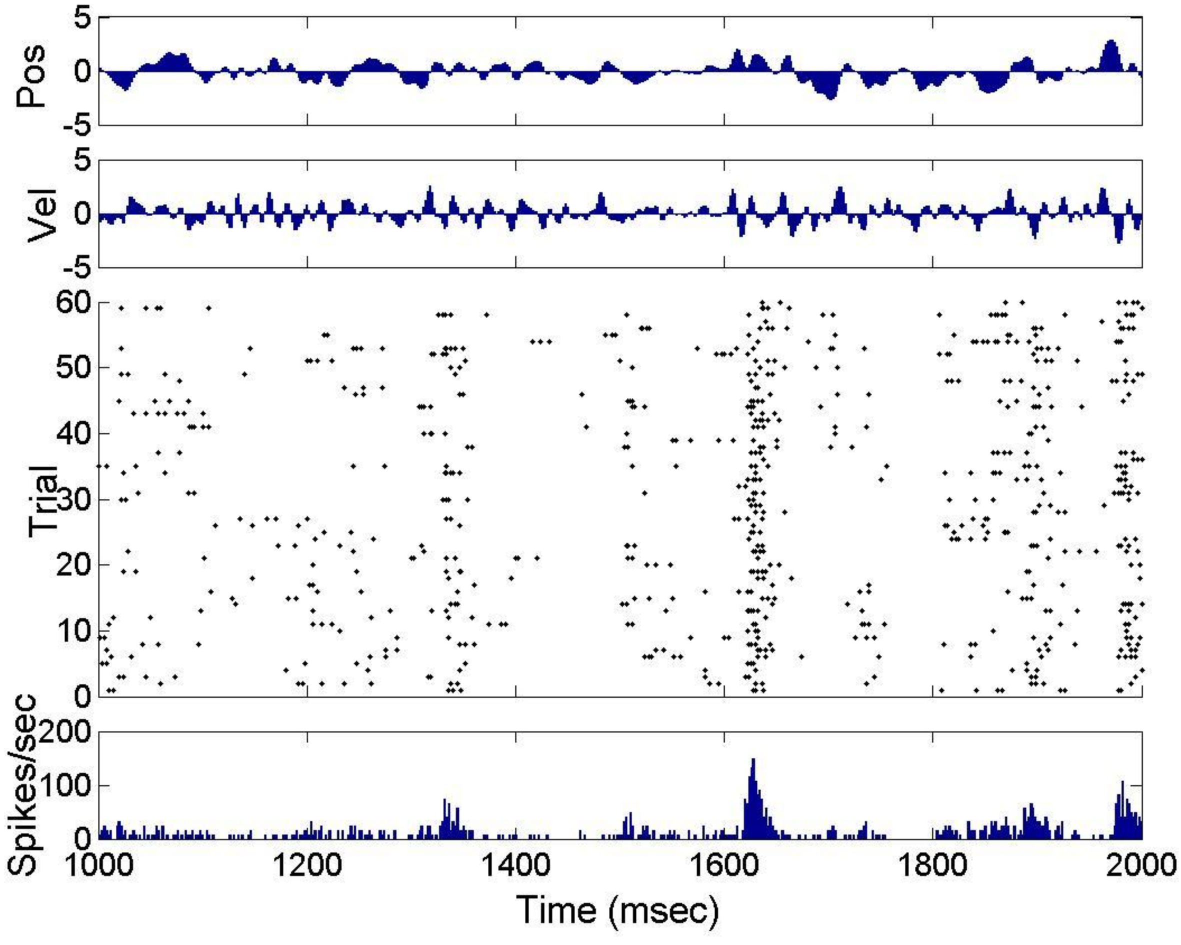

Covariance analysis

Prior

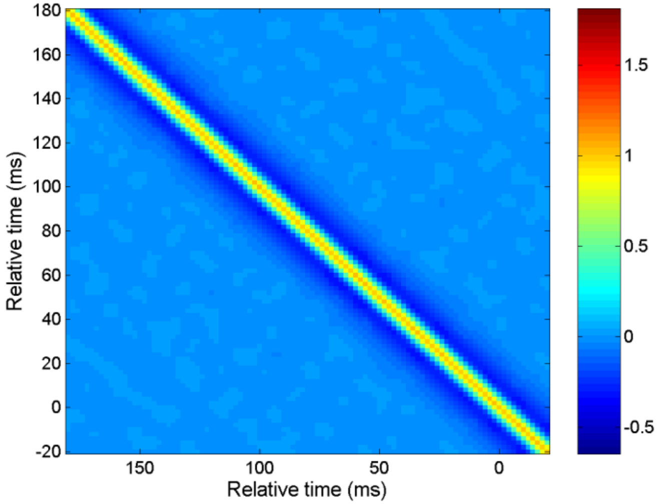

Spike-triggered

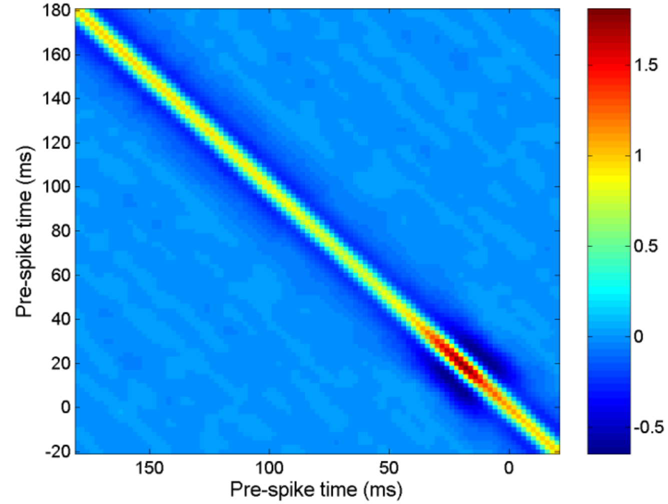

Difference

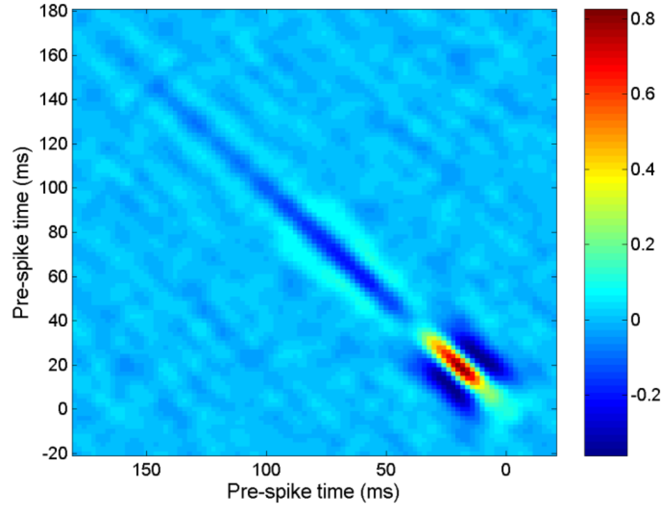

Covariance analysis

Eigenspectrum

Leading modes

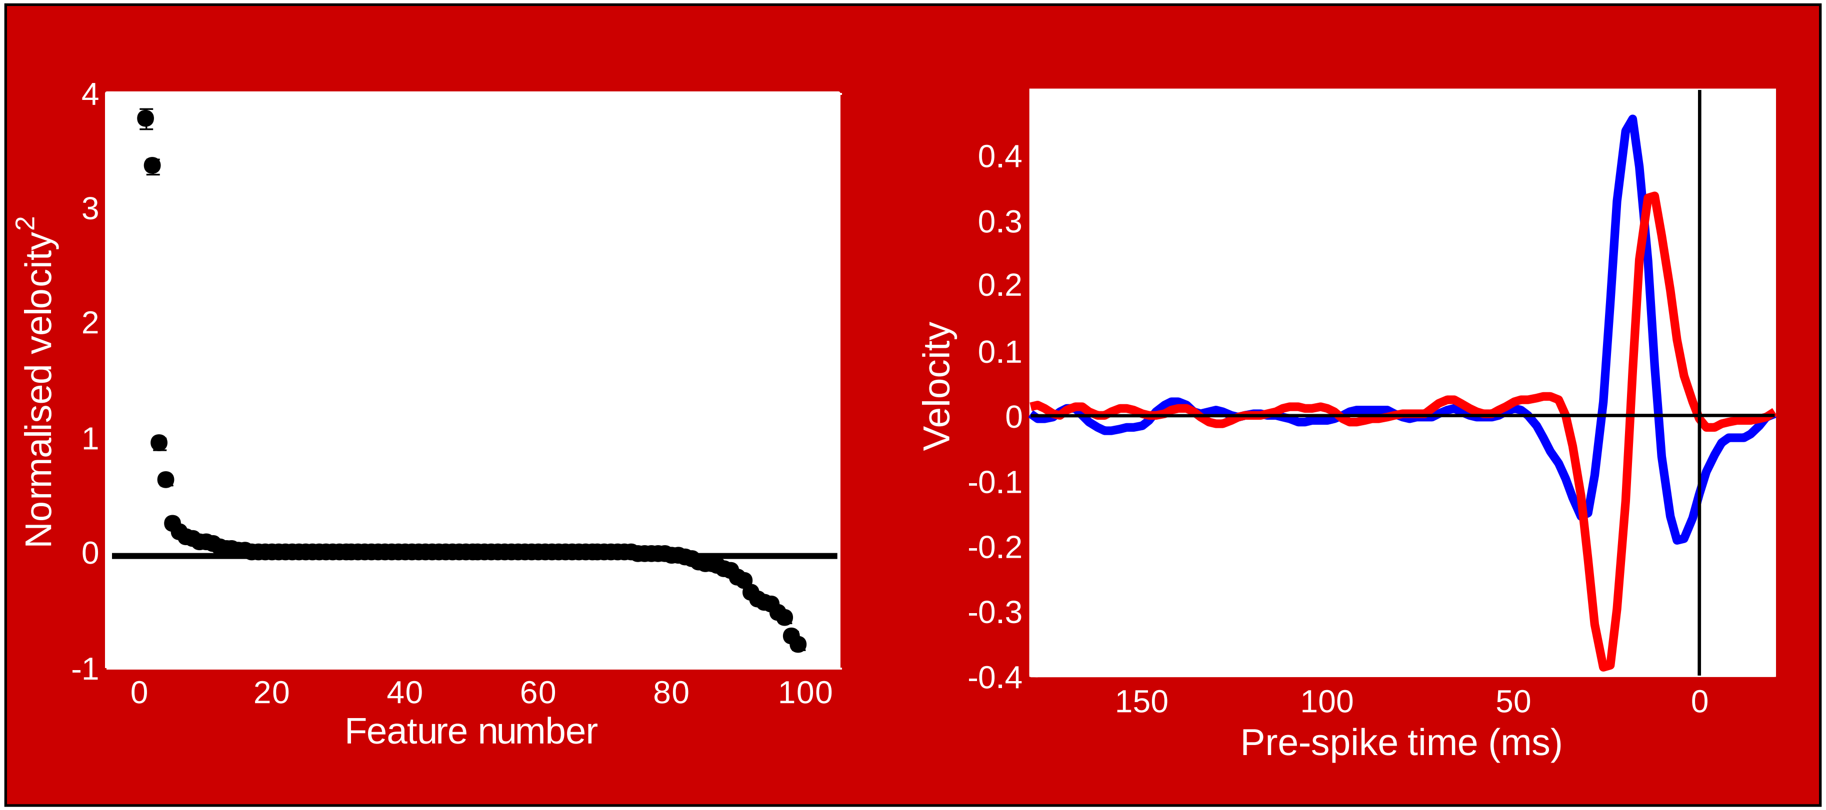

Covariance analysis

Input/output relations wrt first two filters, alone:

and in quadrature:

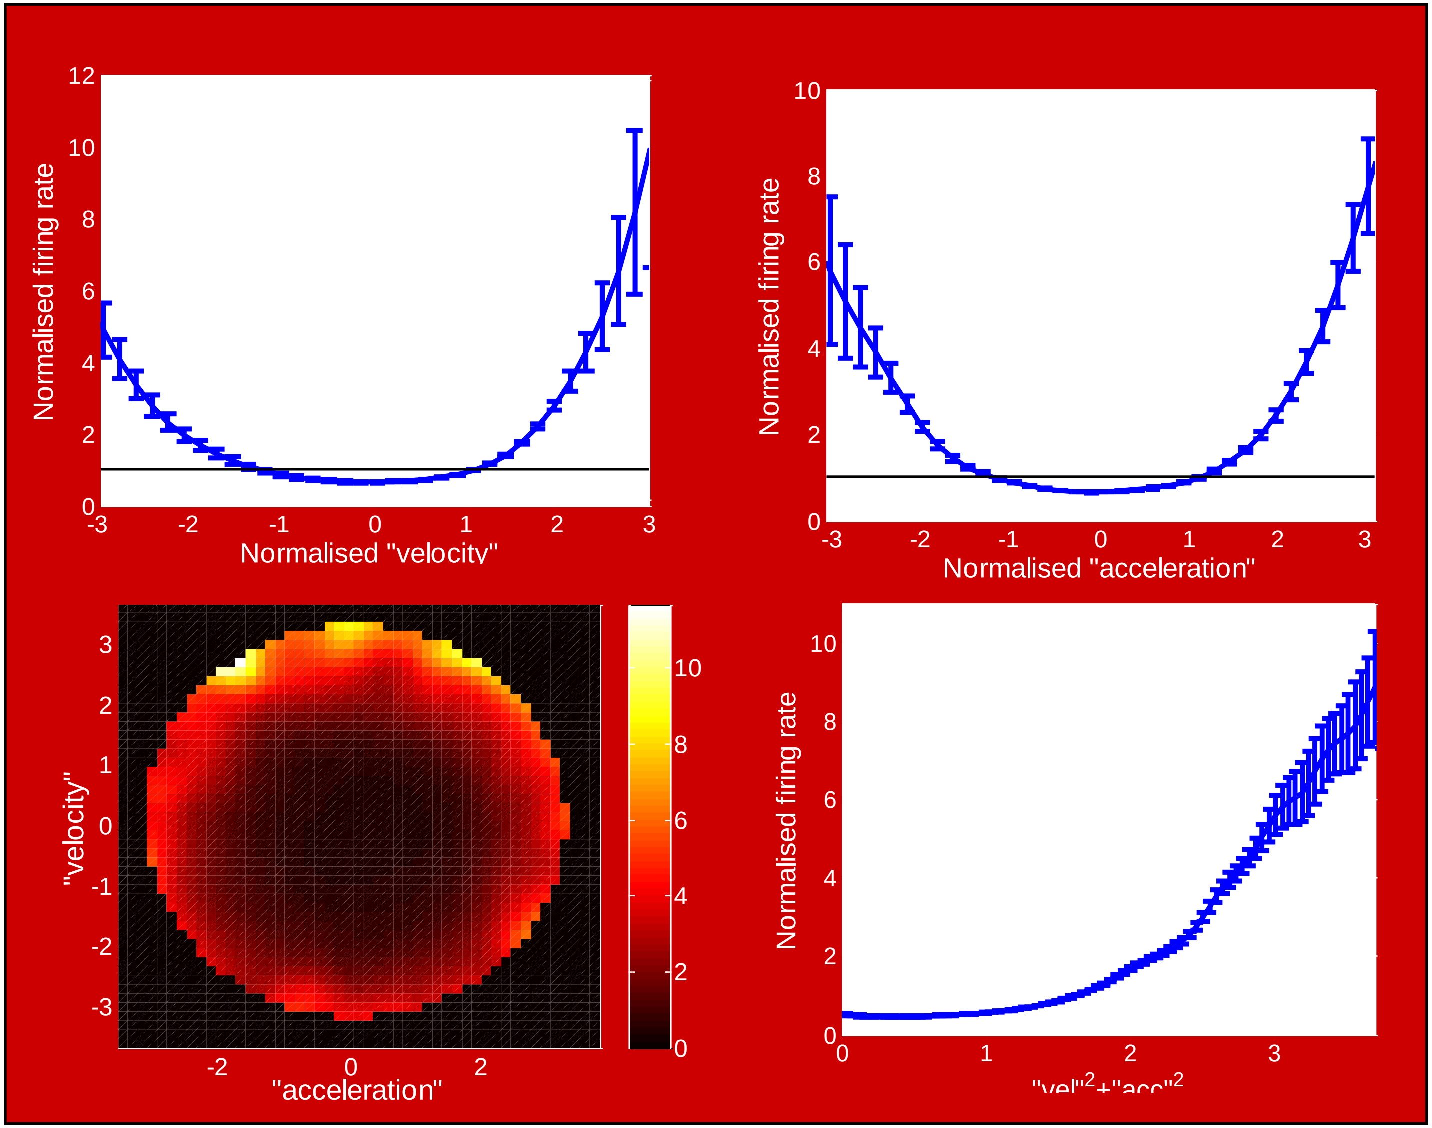

Covariance analysis

How about the other modes?

Next pair with +ve eigenvalues

Pair with -ve eigenvalues

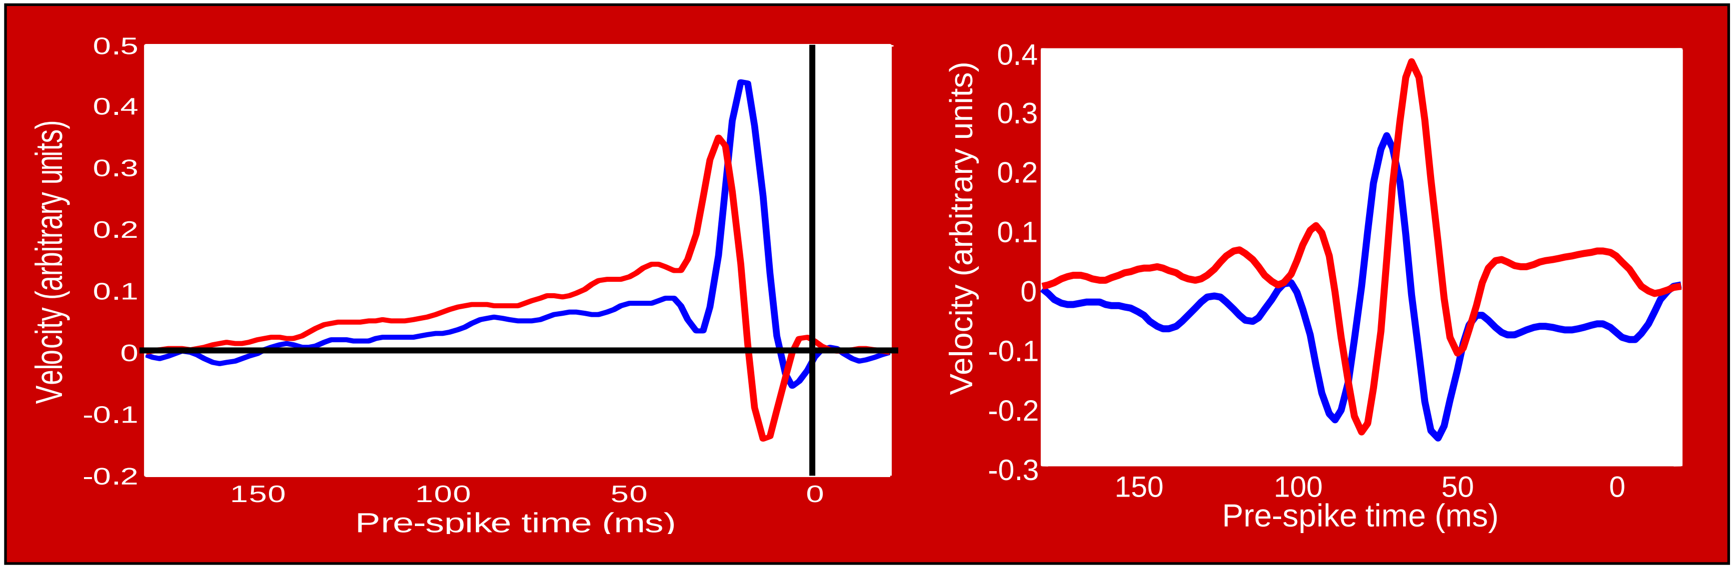

Covariance analysis

Input/output relations for negative pair

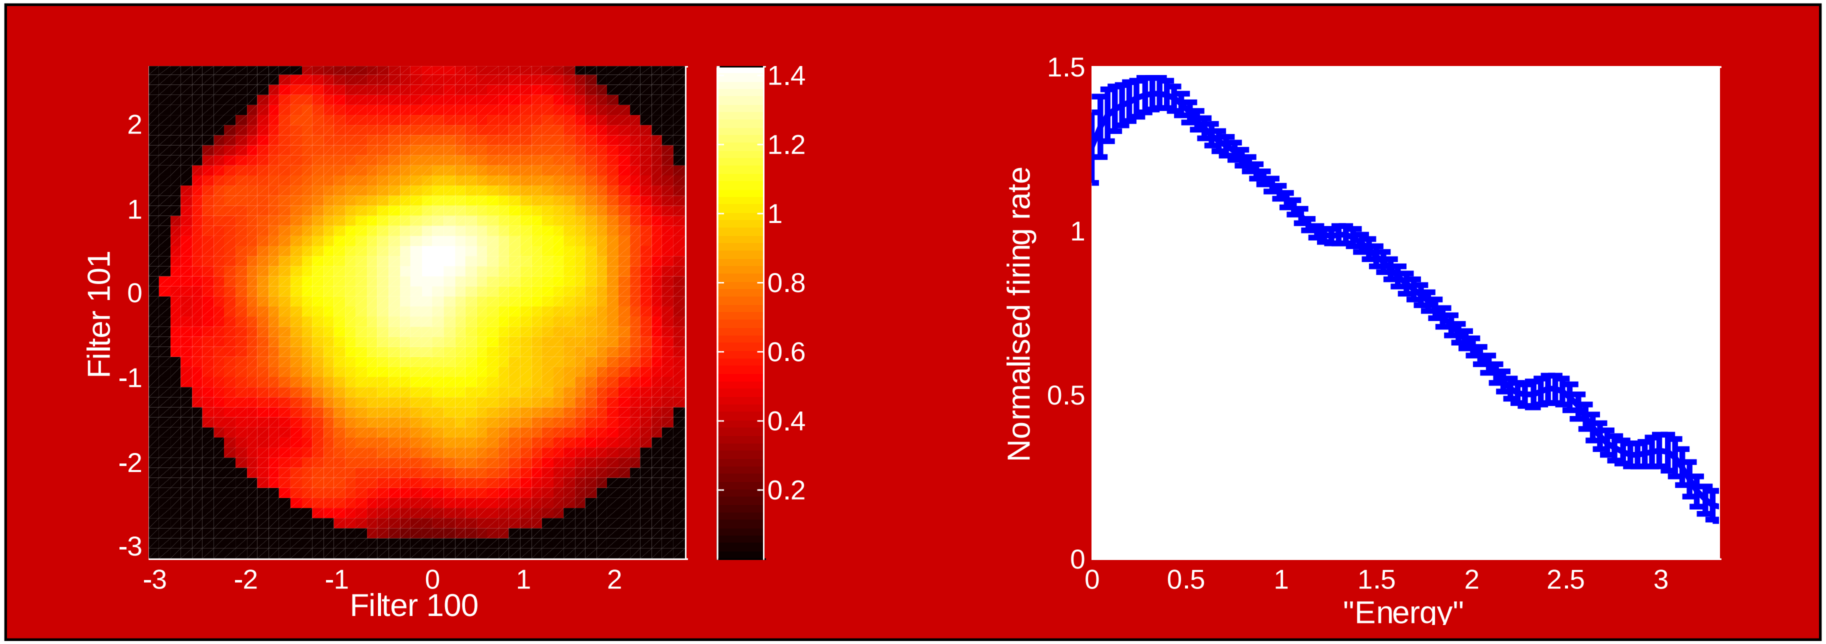

Firing rate *decreases* with increasing projection:
suppressive modes

Basic types of computation:

- integrators (H1)
- differentiators (retina, simple cells, single neurons)
- frequency-power detectors
(complex cells, somatosensory, auditory, retina)

Beyond covariance analysis

1. Single, best filter determined by the first moment
2. A family of filters derived using the second moment
3. Use the entire distribution: information theoretic methods

→ Find the dimensions that maximize the ***mutual information*** between stimulus and spike

Removes requirement for Gaussian stimuli

Limitations

Not a completely “blind” procedure:

have to have some idea of the appropriate stimulus space

Very complex stimuli:

does a geometrical picture work or make sense?

Rates vs spikes:

what is our model trying to do? What do we want to recover?

Adaptation:

stimulus representations change with experience!

Functional models of neural computation

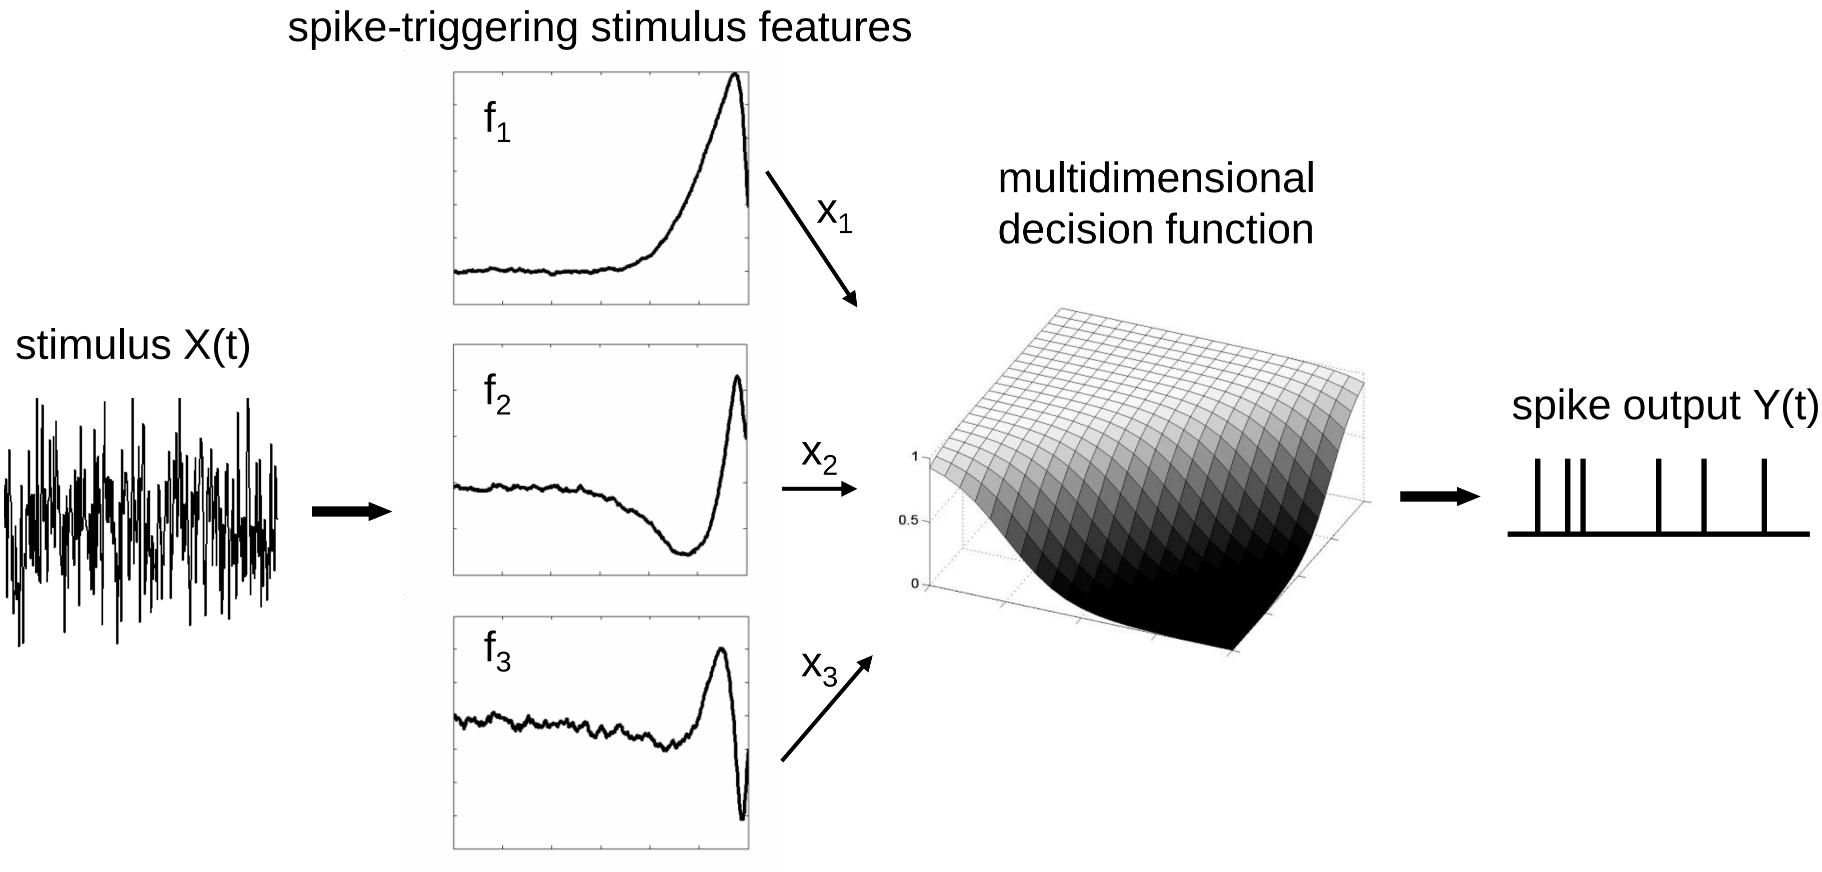

Spike statistics

Stochastic process that generates a sequence of events: ***point process***

Probability of an event at time t depends only on preceding event: ***renewal process***

All events are statistically independent: ***Poisson process***

Homogeneous Poisson: $r(t) = r$ independent of time
probability to see a spike only depends on the time you watch.

$P_T[n] = (rT)^n \exp(-rT)/n!$

Exercise: the *mean* of this distribution is rT
the *variance* of this distribution is also rT.

The *Fano factor* = variance/mean = 1 for Poisson processes.
The CV = coefficient of variation = STD/mean = 1 for Poisson

Interspike interval distribution $P(T) = r \exp(-rT)$

The Poisson model (homogeneous)

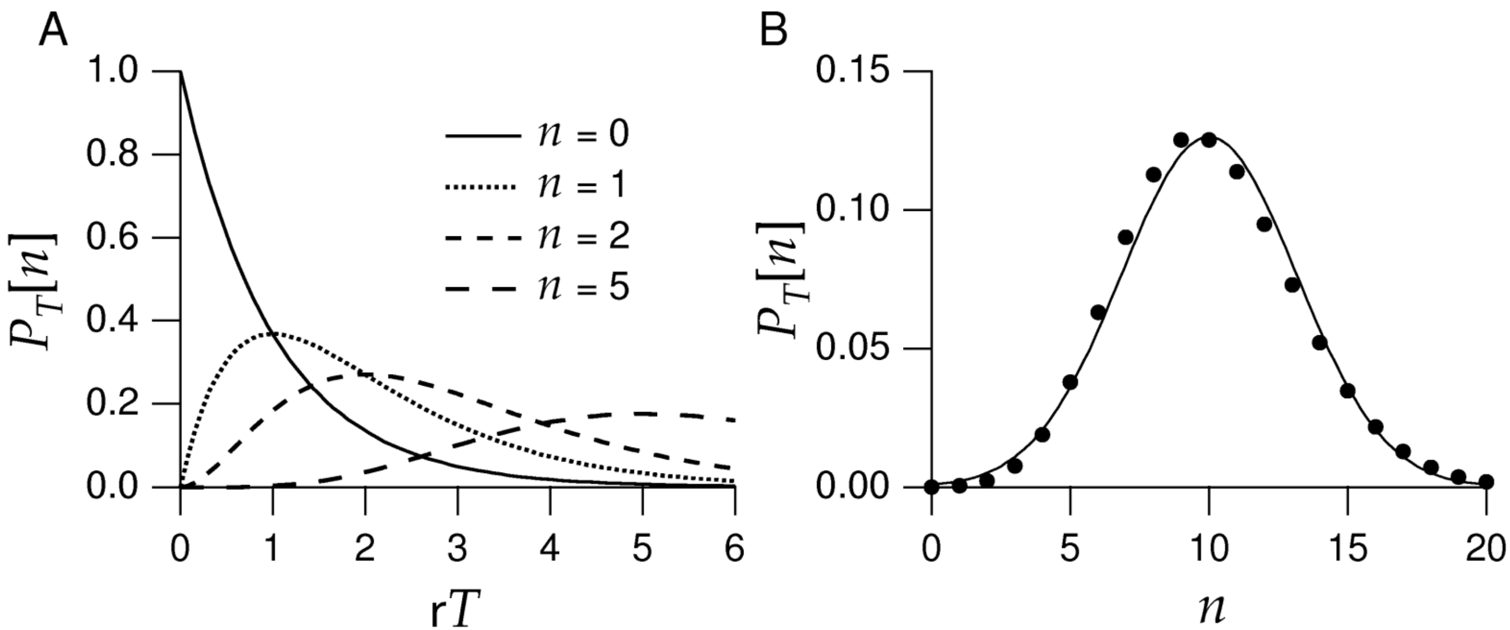

Probability of n spikes in time T as function of (rate × T)

Poisson approaches Gaussian for large rT (here = 10)

How good is the Poisson model? Fano Factor

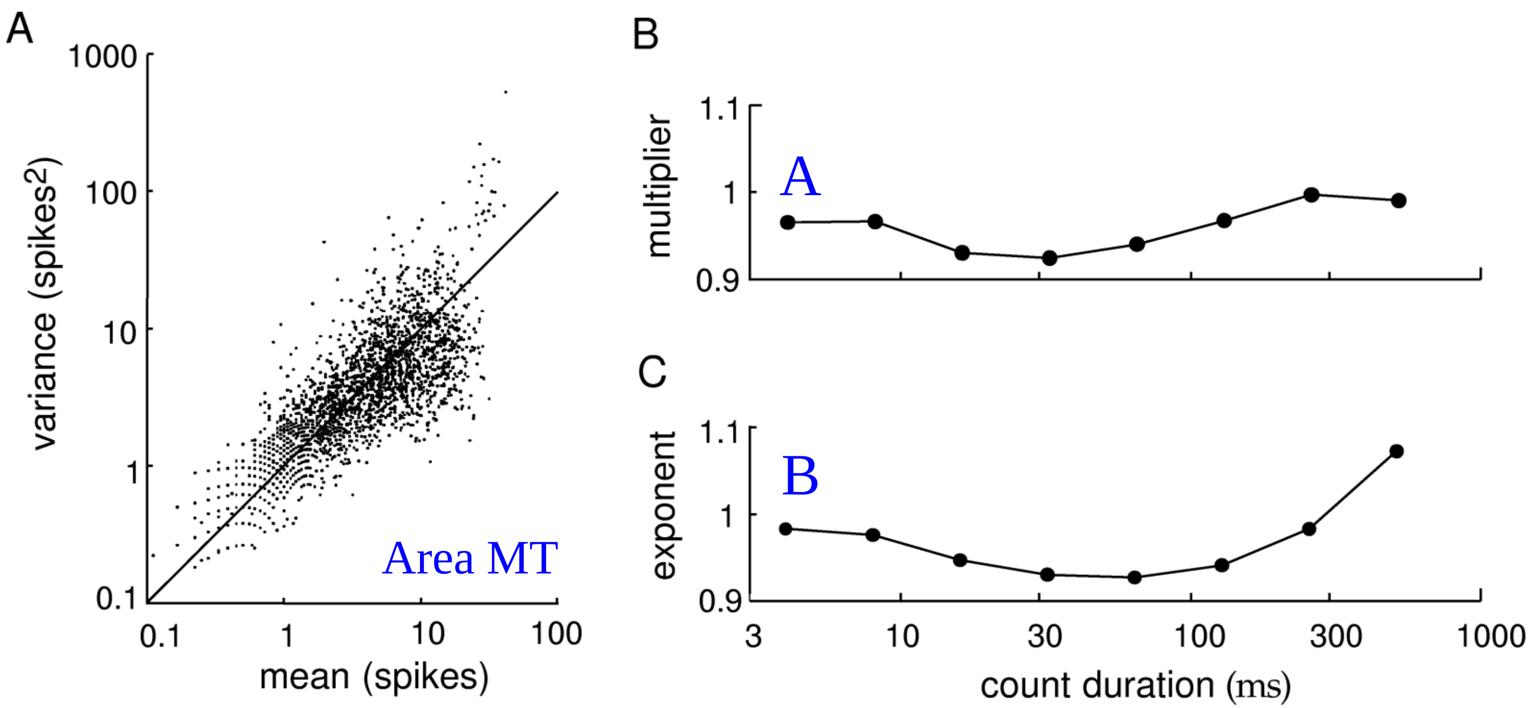

Fano factor

Data fit to:
variance = A $\times$ mean^{B}

How good is the Poisson model? ISI analysis

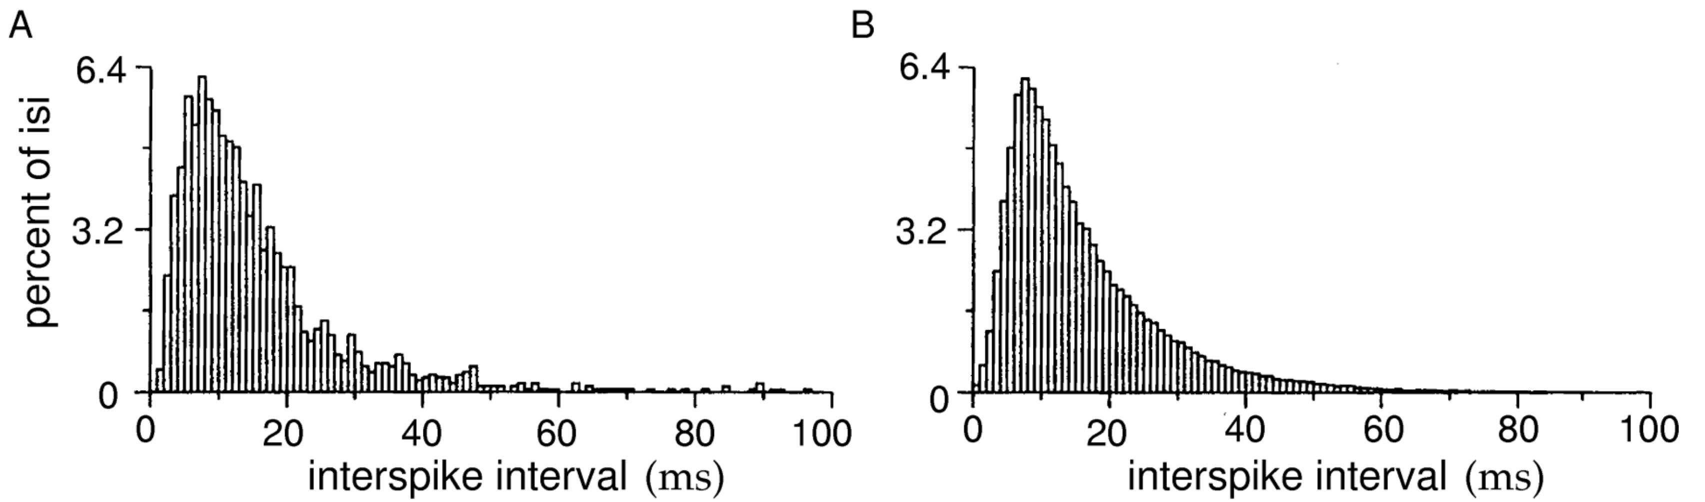

ISI Distribution from an area MT Neuron

ISI distribution generated from a Poisson model with a Gaussian refractory period

How good is the Poisson Model? C_V analysis

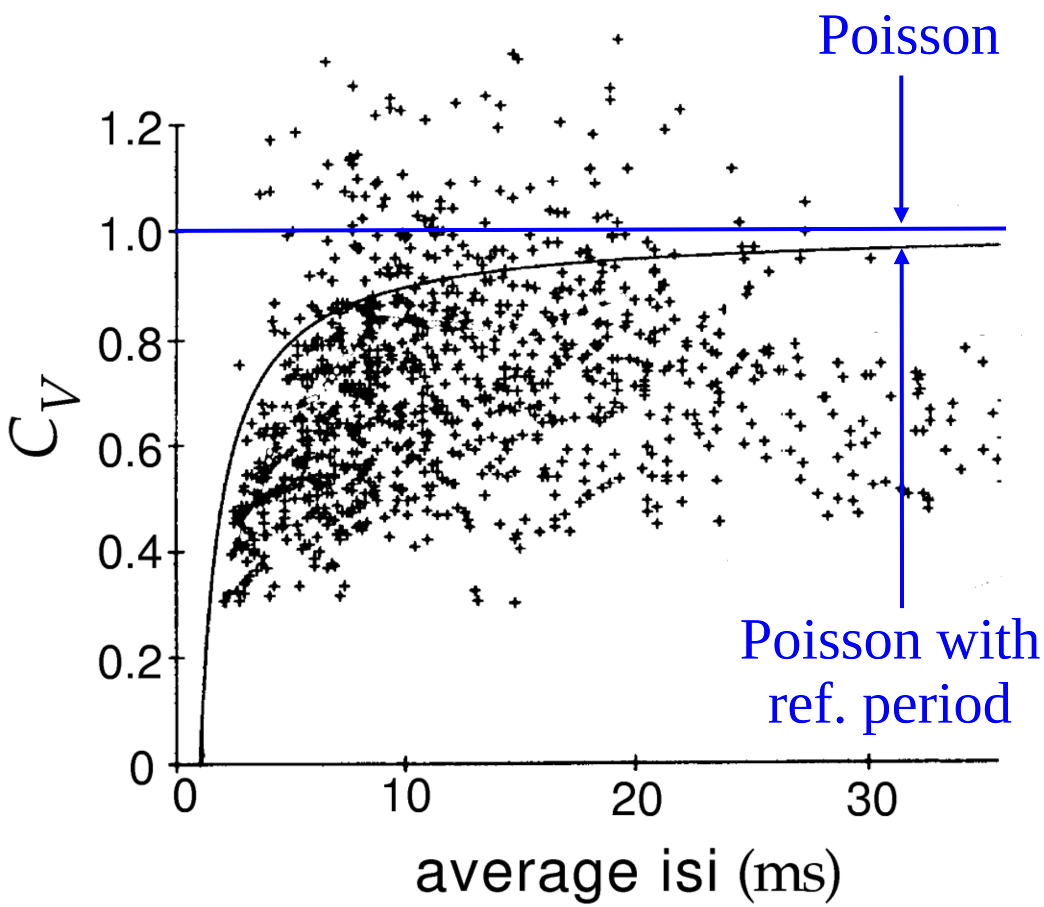

Coefficients of Variation for a set of V1 and MT Neurons

Interval distribution of Hodgkin-Huxley neuron driven by noise

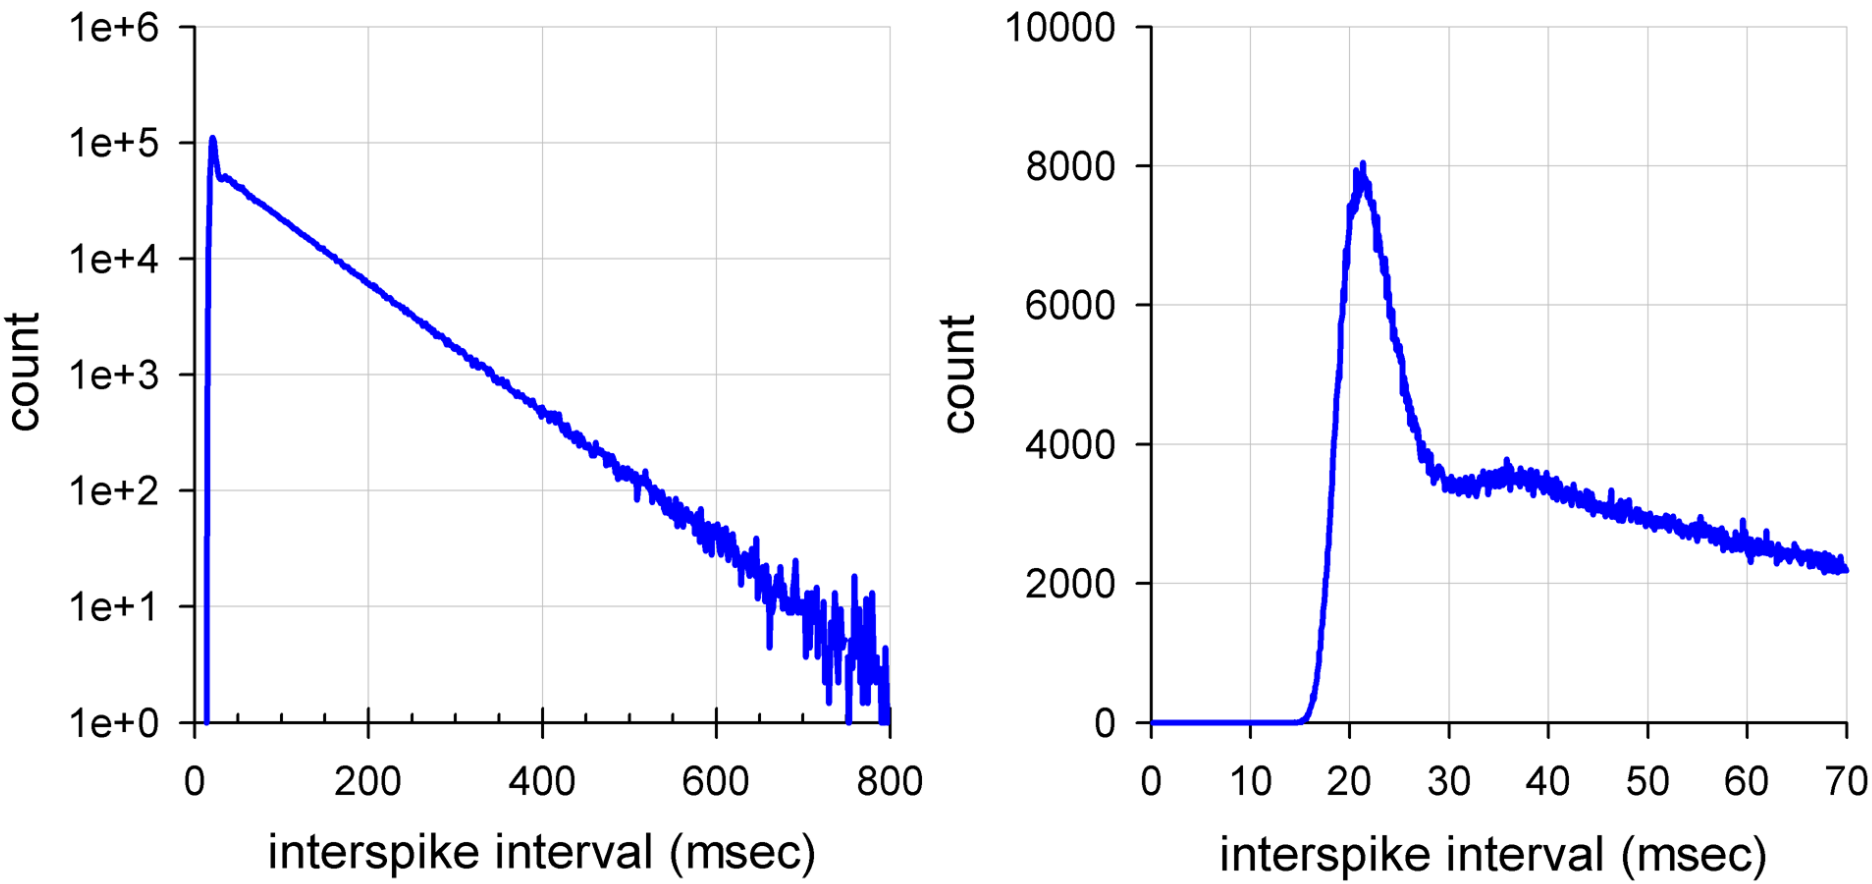

What is the language of single cells?

What are the elementary symbols of the code?

Most typically, we think about the response as a firing rate, $r(t)$, or a modulated spiking probability, $P(r = \text{spike}|s(t))$.

We measure spike times.

Implicit: a Poisson model, where spikes are generated randomly with local rate $r(t)$.

However, most spike trains are not Poisson (refractoriness, internal dynamics). Fine temporal structure might be meaningful.

→Consider spike patterns or “words”, e.g.

- symbols including multiple spikes and the interval between
- retinal ganglion cells: “when” and “how much”

Multiple spike symbols from the fly motion sensitive neuron

Spike Triggered Average

2-Spike Triggered Average
(10 ms separation)

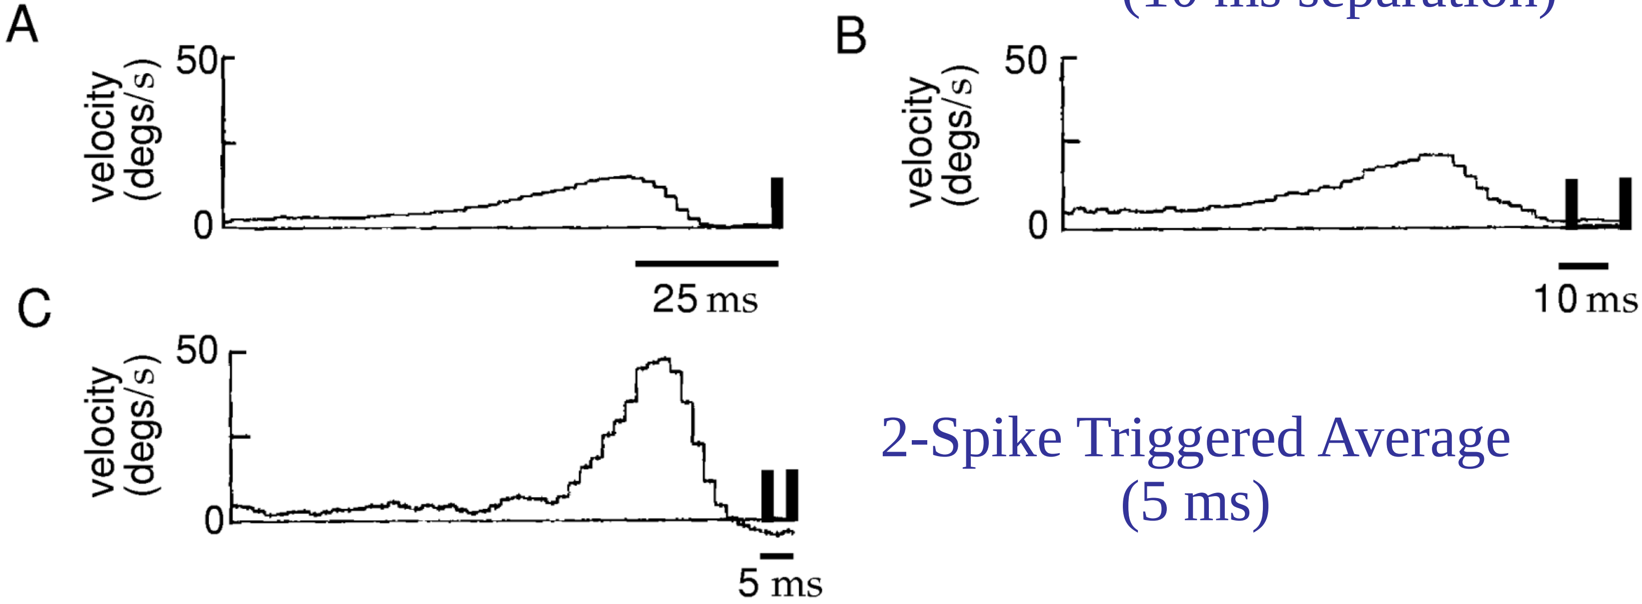

2-Spike Triggered Average
(5 ms)

Decoding

How well can we learn what the stimulus is by looking at the neural responses?

Two approaches:

- devise explicit algorithms for extracting a stimulus estimate
- directly quantify the relationship between stimulus and response using information theory

Predicting the firing rate

Starting with a rate response, r(t) and a stimulus, s(t),

the optimal linear estimator finds the best kernel K such that:

$$r_{\text{est}}(t) = \bar{r} + \int d\tau \, s(t-\tau) K(\tau)$$

is close to r(t), in the least squares sense.

Solving for K(t),

$$K(t) = \frac{1}{2\pi} \int d\omega \, e^{-i\omega t} \frac{\tilde{C}_{rs}(-\omega)}{\tilde{C}_{ss}(\omega)}$$

Stimulus reconstruction

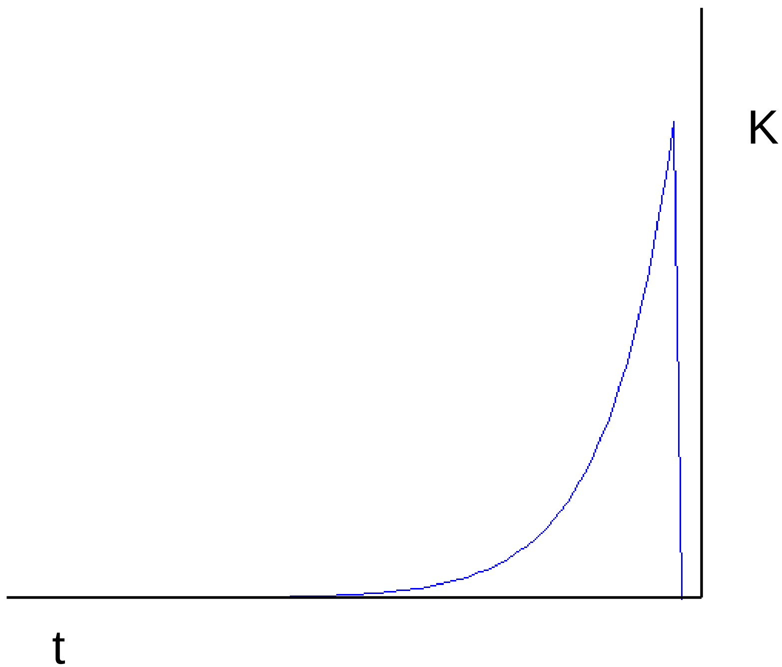

Stimulus reconstruction

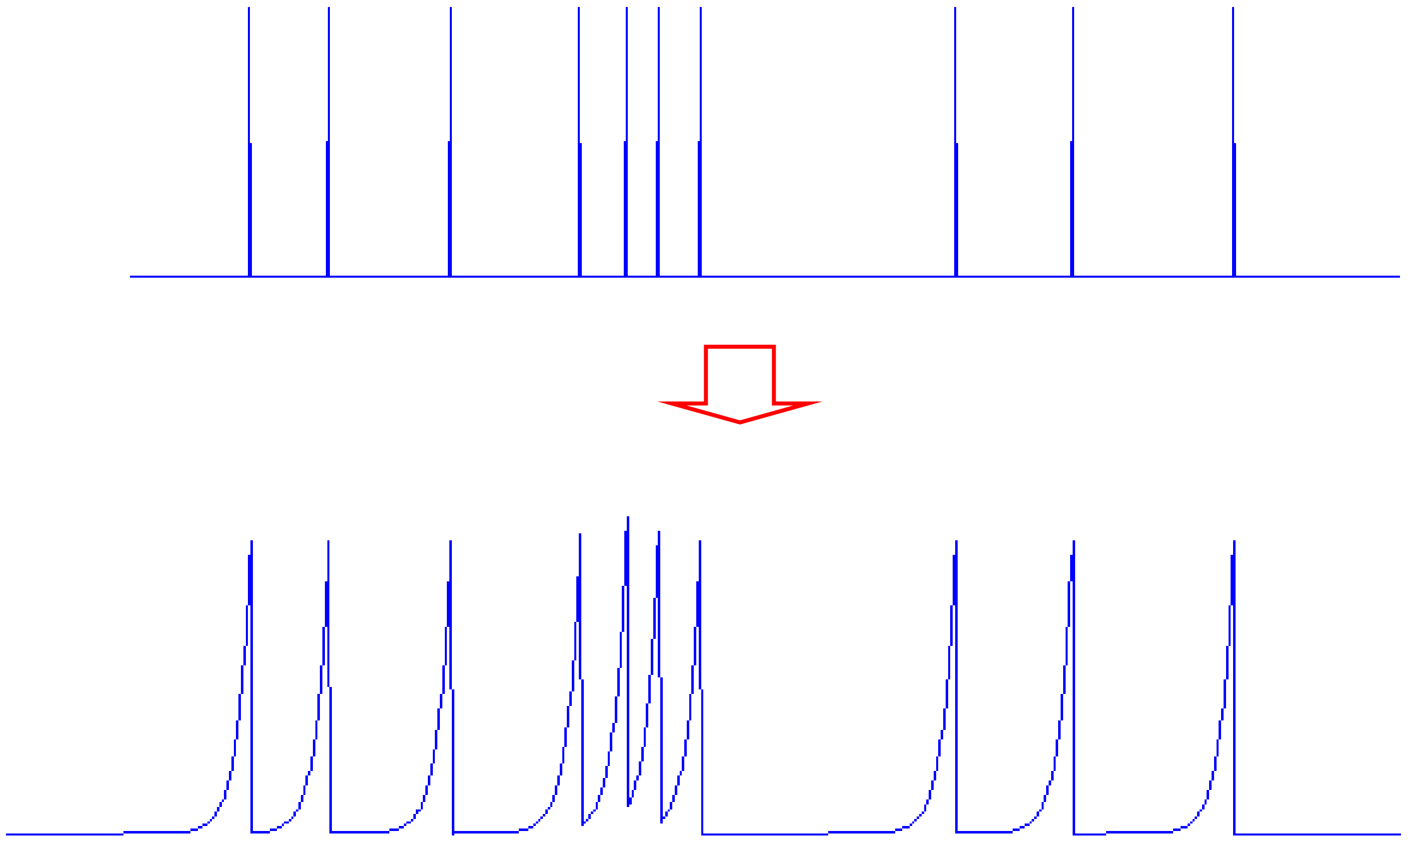

Stimulus reconstruction

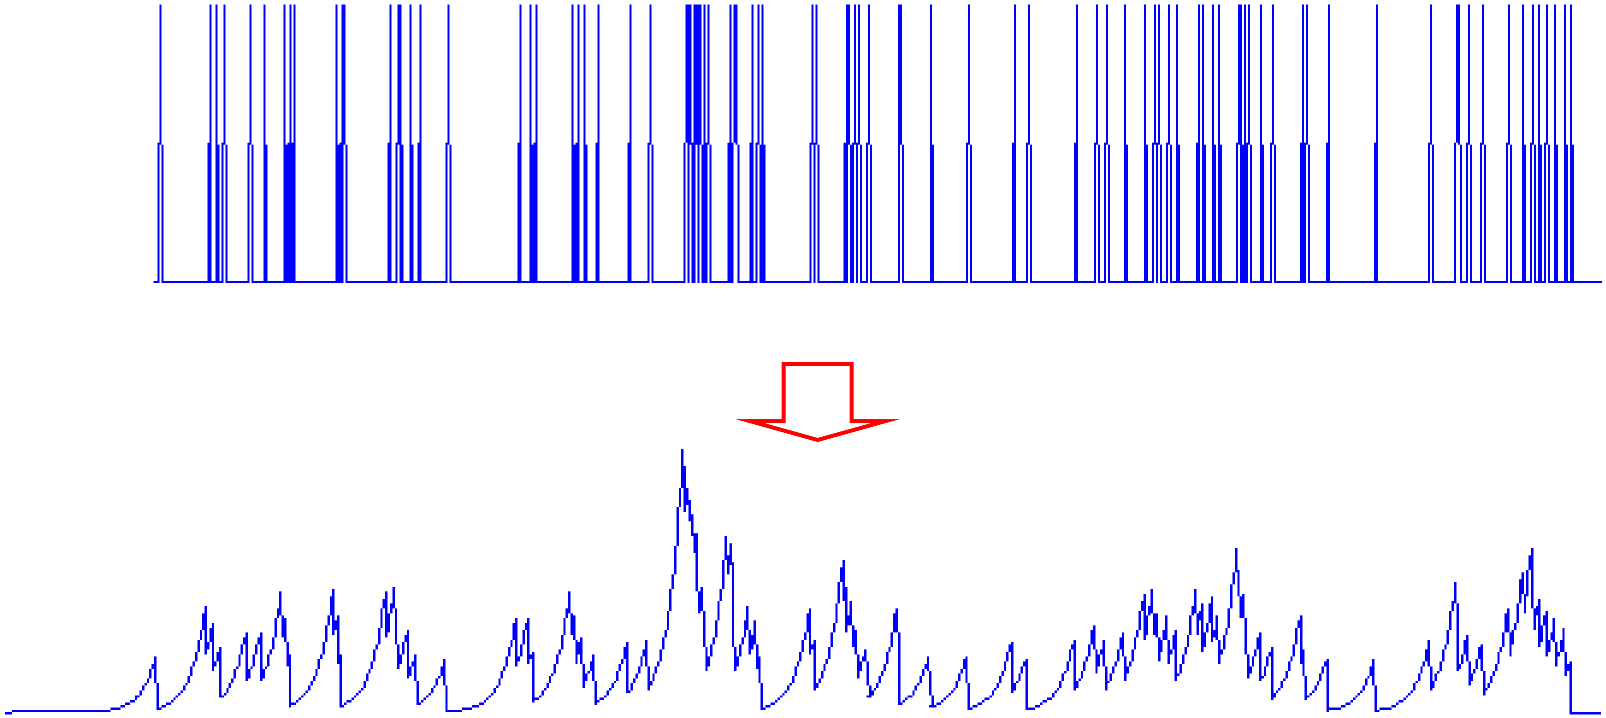

Reading minds: the LGN

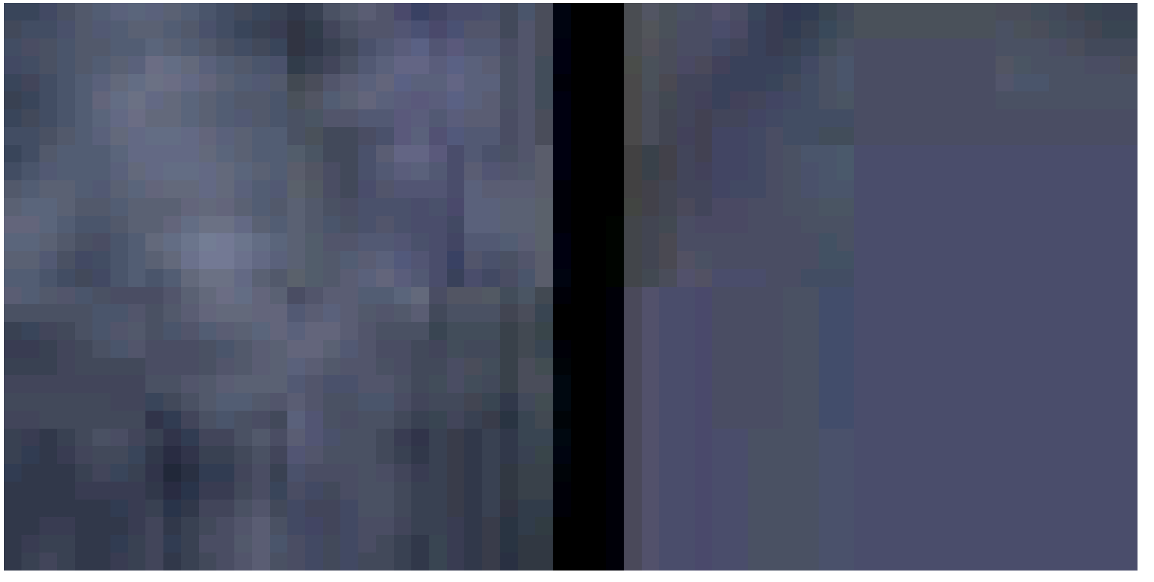

Yang Dan, UC Berkeley