Spikes and the Neural Code

Jonathan D. Victor

Department of Neurology and Neuroscience Weill Medical College of Cornell University

> Computational Neuroscience: Vision Cold Spring Harbor, June 2006

Spikes and the Neural Code

- Spikes and variability
- Information: what is it, why is it difficult to calculate, why bother trying?
- Two case studies in V1
 - The Direct Method, and variability of RF properties within a local cluster
 - The Metric Space Method, and representation of spatial phase
- A new view of spatial processing in V1

Why do neurons fire spikes?

Neurons are trying to send continuous signals, but there is no other good way to send signals quickly, reliably, and for long distances

BUT PERHAPS ALSO

Spikes may play an intrinsic role in neural coding and computation, and hence behavior

Variability: A striking feature of cortical neural activity

What physiologists say:

- Neurons have definite selectivities ("tuning")
- Tuning properties can account for behavior

What physiologists also know:

- Responses depend on multiple stimulus parameters
- Response variability (number of spikes, and firing pattern) is substantial



Kinds of Variability

default model

- Poisson ("completely random") -
 - probability of firing is independent of history
 - no correlation among neurons



- correlations across neurons

Implications of Poisson Variability

- Instantaneous firing probability determines everything about firing patterns
- Other aspects of firing patterns (bursts, correlations) are coincidental and therefore not informative
- Spikes are a necessary evil

How to determine instantaneous firing probability?

- In the laboratory: average across replicate trials
- In the brain
 - average across time?
 - average across (many) similar neurons?

Implications of Non-Poisson Variability

- Firing rate is only one aspect of firing pattern: bursts, correlations, etc. are not determined by rate, and may be (independent or primary) carriers of information
- Such firing patterns are
 - difficult for neurons to create (not maximally random)
 difficult for neurons to read out (can't just count spikes)
- If such firing features are used, then spikes (and variability) are not nuisances -- they are critical to neural coding
- There is always a better way to determine firing rate than merely averaging

A Simple Consequence of Variability



Flawed Arguments that Firing Patterns are a Nuisance

Neurons must fire irregularly.

No, sensory neurons have low variability, and may be clock-like. *Cortical* neurons are intrinsically variable.

No, if their inputs are precisely controlled, they work like machines.

Controlling firing patterns is not worth the effort.

No, specific channels lead to particular firing patterns (e.g., thalamocortical neurons).

Reading firing patterns is not worth the effort or not biologically plausible.

No, postsynaptic mechanisms are highly specialized.

- coincidence detection
- "facilitating" and "depressing" synapses (sensitivity to intervals)
- dendritic processing (not just global summation of inputs)

Arguments for Role of Firing Patterns: Not Convincing to All

EEG patterns reflect state of arousal.

But that's clinical. (Or that's an epiphenomenon).

Burst and tonic modes in the thalamus reflect state of arousal.

That's clinical too.

Spike time dependent plasticity!

That's just learning.

Direct evidence: locust olfactory system (Laurent)

But that's a locust. What about Newsome's cortical microstimulation experiments?

Microstimulation doesn't only change rate, but also pattern. And who knows what the animal experiences?

A purely experimental test of a neural coding hypothesis is difficult

Rate, pattern, timing are *interdependent** and difficult to manipulate independently

*"Time is that great gift of nature which keeps everything from happening at once." (C.J. Overbeck, 1978)

- It is unclear what are the "words" of the neural code (firing rate, firing pattern, etc.)
- So, we need to use a theory

A Win-Win Situation

- If firing pattern and variability turns out to be important, then we'll understand why the primate brain is so complicated.
- If variability turns out to be just a nuisance to be averaged out, we'll understand why the primate brain is so big.

Information = Reduction in Uncertainty (Shannon)



- Reduction in uncertainty from 4 possibilities to 2
- Information = log(4/2)

Entropy: Simple definition

Suppose a system may be in any one of some set of states. The entropy of the system is the minimum number of yes-no questions, on average, that must be asked to determine the state of the system.

Say the probability that the system is in state *j* is p_j . Then the entropy *H* is given by

 $H = -\sum_{j} p_{j} \log p_{j} \qquad \text{(use } \log_{2} \text{ to get "bits")}$

Higher entropy means more uncertainty.

Entropy: Properties Mixing property $H\{(1-\lambda)p_J + \lambda q_J\} \ge (1-\lambda)H\{p_J\} + \lambda H\{q_J\}$

Chain rule, general form



Chain rule, table form



H(Z) = H(X) + H(Y)

Information: Difference of Entropies



Information ={Entropy of the a priori distribution of input symbols} {Entropy of a posteriori distribution of input symbols, given the
observation k, averaged over all k}

Information: Symmetric Difference of Entropies



Information =

{entropy of output}+{entropy of input} -{entropy of table}

 $I = H\{p(j,\bullet)\} + H\{p(\bullet,k)\} - H\{p(j,k)\}$

Information: Properties

I symmetric in input and output

I independent of labeling within input and output

 $l \ge 0$, and l = 0 if and only if $p(j,k) = p(j,\bullet)p(\bullet,k)$

$$I \leq H\{p(j,\bullet)\}$$
 and $I \leq H\{p(\bullet,k)\}$

Data Processing Inequality: if Y determines Z, then $I(X,Y) \ge I(X,Z)$

Information: Why Calculate It?

- An interesting, natural quantity
- A constraint on models
- Biological insight
 - Compare across systems (e.g., "one spike per bit")
 - Determine the constraints on a system (e.g., metabolic cost of information)
 - See where it is lost (e.g., within a neuron)
 - Determine what the system is "designed" to do
- Nonparametric measure of association
 - Test statistical models (e.g., triplets)
 - Determine candidates for neural codes
 - Rule out candidates for neural codes

Investigating neural coding: not Shannon's paradigm

- Shannon
 - symbols and codes are known
 - joint (input/output) probabilities are known
 - what are the limits of performance?
- Neural coding
 - symbols and codes are not known
 - joint probabilities must be measured
 - ultimate performance often known (behavior)
 - what are the codes?

Information estimates depend on partitioning of stimulus domain



Information estimates depend on partitioning of response domain



Partition as finely as possible? (entropy, but also information)

- $H=-\Sigma p_i \log p_i$
- The probabilities p_i must be estimated empirically (say, as p_i^{est})
- $H^{est} = -\Sigma p_i^{est} \log p_i^{est}$ is a biased estimator of H
- Bias can be estimated
- Bias roughly proportional to number of partitions
- Once the number of partitions exceeds the number of observations, estimates of bias are useless

The Basic Difficulty

We want to determine *p* log *p*, but we only have an estimate of *p*, not its exact value.

p log *p* is a nonlinear function of *p*, and the nonlinearity is greatest near *p*=0. Therefore, the error in replacing $by \log$ is greatest near *p*=0.

Refining the "bins" (stimuli or responses) pushes *p* to near 0.

Multineuronal analysis: motivation

- Cortical neurons are organized in *columns* in which nearby neurons have similar stimulus selectivities: orientation, spatial frequency, etc.
- Individual cortical neurons appear to be "noisy"
- Perhaps, local population activity should be read out by summing the activity of neighboring neurons

Questions

- Is the information transmitted by neighboring neurons *independent* or *redundant*?
- Is information fully captured by simply summing local activity?

Information in spike trains: "Direct Method" Strong et al. 1998; Ruyter van Steveninck et al. 1997

- Compare number of different firing patterns that the neuron can produce to reliability of firing pattern in response to the same stimulus
- Information = entropy of all responses to all stimuli average entropy of repeated responses to the same stimulus

(entropy ~ log of number of different kinds of responses)

Extension to multiple neurons

Neuron 1

Neuron 2





summed population

average responses neuron of origin irrelevant



labeled line

keep responses separate specific neuron relevant



Preparation

 Recordings from primary visual cortex (V1) of macaque monkey



- ensures neurons are neighbors (ca. 100 microns)
- requires spike sorting



Channel

5 msec

Online spike sorting with tetrodes



Offline spike sorting via tetrodes



M-sequence analysis: cross-correlation



Cross-correlate the spike trains with the stimulus, pixel by pixel, to determine the average stimulus at each instant prior to a spike (here, 72 ms before a spike)



Receptive field map in space and time



This receptive field is oriented and directional.



- Tabulate the occurrence of each interspike interval
- Numerical experiment: compare with the expectation based on firing rate by randomizing spikes across trials

Interspike interval distribution



Interspike interval distributions: often multimodal


Subset boundaries: consistent across neurons



Receptive field characterization depends on interspike intervals



Direct Method: Multiple Neurons

Neuron 1

Neuron 2





summed population

average responses neuron of origin irrelevant



labeled line

keep responses separate specific neuron relevant



M-sequence responses



Responses and information rates

info. rate (bits/sec)



Another pair



Quantifying redundancy



Redundancy indices



Population summary $\frac{3}{4}$ $\frac{60}{0}$ $\frac{1}{2}$ $\frac{1}{4}$ $\frac{60}{6}$



8

neurons

Pitfalls of the "Direct Method"

- Stimulus must elicit full repertoire of neural responses – m-sequences
- Estimates of response probabilities must be accurate – consider only information rates over brief periods (8 ms)
- Biases in entropy estimates (<log p> ≠log)
 asymptotics
 - simulations
 - split-set analyses

The "Direct Method": A sampling problem (Strong, Bialek, et al.)

All calculations based on discretized responses



- T_{letter} must be small to capture temporal detail – timing precision of spikes: <1 ms

- T_{word} must be large to include slow events

 inhibitory potentials, macroscopic rhythms: >100 ms
- T_{word} / T_{letter} >100. Up to 2⁽T_{word}/T_{letter}) probabilities must be estimated.

Multiple neurons: a severe sampling problem

• One dimension for each bin and each neuron



• $2^{L(T_{word} / T_{letter})]$ probabilities must be estimated.

"Time does not chop itself up for our convenience."

Itamar Moses in "Bach at Leipzig"

Acknowledgements

Weill Medical College Department of Neurology and NeuroscienceKeith PurpuraFerenc MechlerDanny ReichMichael RepucciDmitriy Aronov

Weill Medical College Laboratory of Neuroinformatics Dan Gardner David Goldberg

The Rockefeller University Laboratory of Biophysics Bruce Knight

Support

Human Brain Project/Neuroinformatics (NIMH, NIA, NIBIB, NINDS, and NSF)

and NEI

What are we missing?

- The response space is a continuous domain
- There are natural notions of "similar" responses



Exploiting this topology is important

Metric-Space Analysis: Overview

- Formulate a variety of plausible candidate notions of distance
- Determine which do a good job of classifying responses to stimuli via the mutual information of the "confusion matrix"
- Ask if the distances between responses correspond to a perceptual space

In what ways can spike trains be similar?



Cost-based distances

- Defined by a list of allowed transformations of spike trains and an assignment of costs to each of these transformations
- Distance between two spike trains A and B is the least total cost of *any* sequence of transformations from A to B
- With very minor fine print, this assignment of numbers to pairs of spike trains has all the required properties of a distance (metric)

Distance based on spike times

- Allowed transformations:
 - insert a spike: unit cost
 - delete a spike: unit cost
 - shift a spike in time by ΔT : cost is q ΔT
- Spike trains are similar only if spikes occur at similar times (i.e., within 1/q sec).
- Formalizes the hypothesis that processing of coincidences between spike trains (by mechanisms such as thresholds) is key to neural integration.

Transforming a spike train in elementary steps: Spike time metric



Transforming a spike train in elementary steps: Spike interval metric



Distances between all pairs of responses determine a response space



Configuration of the response space depends on the choice of distance



Weak clustering: responses to the four stimuli are interspersed



Strong clustering: responses to the four stimuli are grouped

Information quantifies strength of clustering



Information = row entropy + column entropy - table entropy

Distance based on interspike intervals

• Allowed transformations:

- insert or delete an interval: unit cost

–change an interspike interval by ΔT : cost is q ΔT

- Spike trains are similar only if corresponding interspike intervals match (within 1/q sec)
- Formalizes the hypothesis that processing of patterns of intervals (by mechanisms such as facilitating and depressing synapses) is a key element of neural integration

Visual cortex: contrast responses



a (1000) 8 () (1 9 () (1 1 (1) 1 (1) 1 (1) 1 (1) 1 (1) 1 (1) 1 (1)	1001100 111200 0 1 10000 1 1 0011 0 1 0011 0 1 10 1 1 1 0111 1 1 1		01 011 10 0 0 1110100 0 01 00 0 11 0 11	 1 8 9 1 1<
Trial number				

256 ms

Visual cortex: contrast coding



Multiple visual sub-modalities



Comparison of types of codes



Attributes are coded in distinct ways





Joint representation of contrast and spatial phase



Distances calculated by spike time metric, with q=64 (informative precision ~15 ms)

Colors indicate spatial phase (22.5 deg intervals) Sphere sizes indicate contrast (0.25 and 1.0)

Multineuronal codes

A two-parameter family analyzes importance of neuron of origin, and importance of timing



Candidate codes are evaluated via an informationtheoretic measure of response clustering

Multineuronal spike time distance

- Allowed transformations:
 - insert a spike: unit cost
 - delete a spike: unit cost
 - shift a spike in time by ΔT : cost is q ΔT
 - change the neuron of origin of a spike: cost is k



neuron of origin ignored ("population code") each neuron's response considered individually ("labeled line")

Transforming a multineuronal spike train in elementary steps



The key to the multineuronal algorithm



forbidden and allowed crossings



summarized by a 3-d geometry

Aronov, 2003
Spatial phase coding in two simple cells



Multineuronal representation of spatial phase



Distances calculated by spike time metric, with q=32 (informative resolution, ~ 30 ms)

Colors indicate spatial phase (22.5 deg intervals)

Conclusions from Single-Neuron Recordings

- Cortical neuronal variability differs qualitatively from that of a Poisson process
- Spike times and spike intervals, and not just the firing rate, convey information about spatial attributes of visual stimuli.
- The coding strategy depends on the visual attribute (contrast, pattern), especially beyond V1.
- Multiple visual attributes are simultaneously represented in a single neuron's responses.

Conclusions from Multi-Neuron Recordings

- The neuron of origin is important for signaling, even within local neuronal clusters.
- Reducing the activity within a local cluster of neurons to a population firing rate discards substantial information.
 - Neurons within a population appear more redundant.
 - The representation of a stimulus space (e.g., spatial phase) is less faithful, both quantitatively and qualitatively.

Many information-theoretic approaches to neural data

- Spike train considered as a symbol sequence
 - Direct method (Strong)
 - Birthday paradox (Nemenman)
 - Optimized polynomial estimators (Paninski)
 - "Coverage" (Chao and Shen)
 - Information bottleneck (Tishby), Codebook (Miller and Dimitrov)
 - LZW (Wyner; Farach; Levy)

Suffix tree method (Shlens)

spike train model

Spike train considered in continuous time

Stimulus reconstruction method (Bialek)

Metric space method

- Binless embedding method
- Power series expansion (Schultz and Panzeri)

stimulus-response model

code model

neurodatabase.org

