Population Coding and Neuronal Variability

AMATH 342
Monica Liu

From encoding to decoding: predicting the stimulus from neural responses

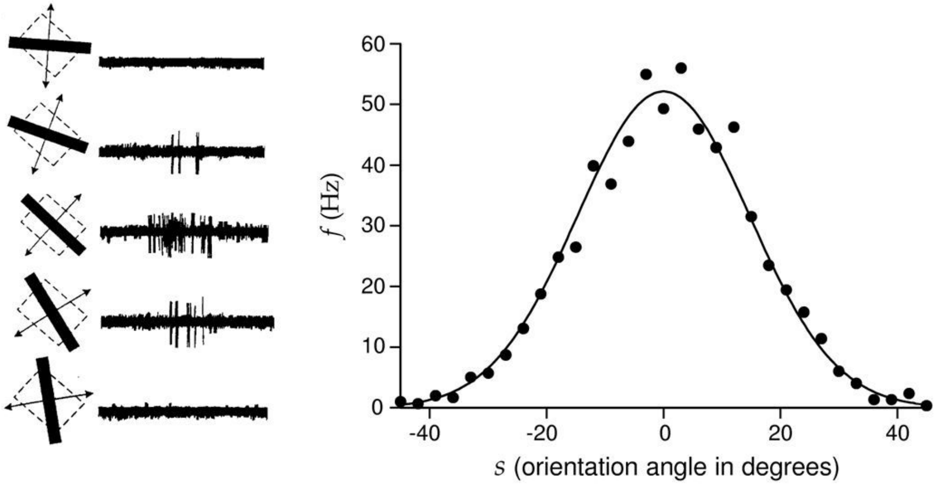

Gaussian tuning of a cortical (V1) neuron

If we observe the cell firing at 52Hz, what is the most likely orientation of the stimulus being shown?

From encoding to decoding: predicting the stimulus from neural responses

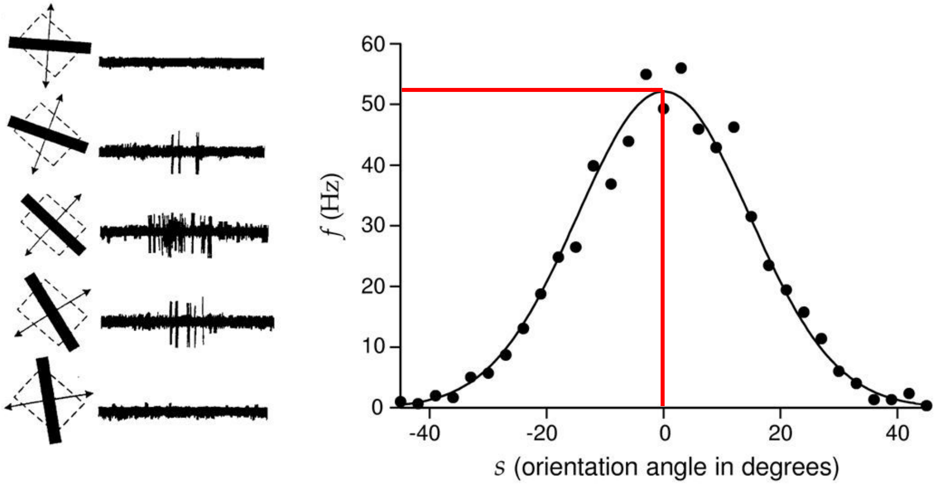

Gaussian tuning of a cortical (V1) neuron

If we observe the cell firing at 52Hz, what is the most likely orientation of the stimulus being shown?

From encoding to decoding: predicting the stimulus from neural responses

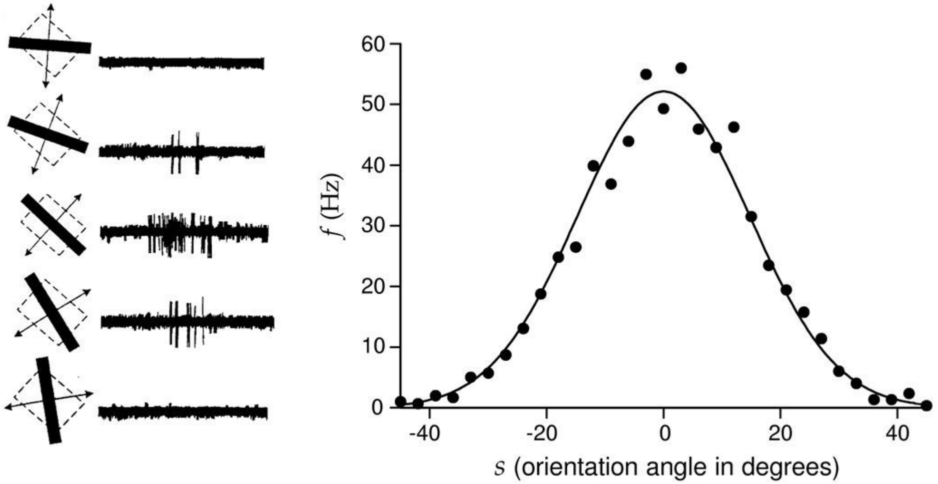

Gaussian tuning of a cortical (V1) neuron

If we observe the cell firing at 52Hz, what is the most likely orientation of the stimulus being shown?

How certain are we?

Variability of neural responses influences decodability of the stimulus

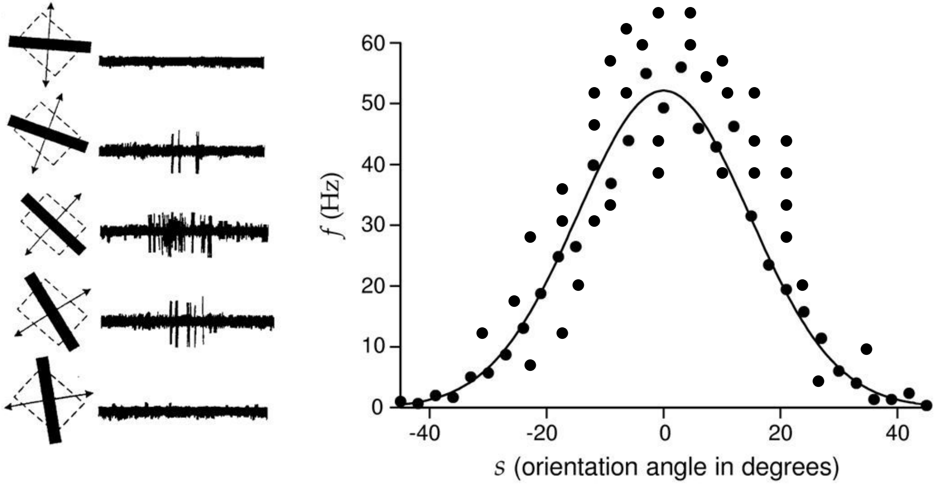

Gaussian tuning of a cortical (V1) neuron

If we observe the cell firing at 52Hz, what is the most likely orientation of the stimulus being shown?

How certain are we?

Measuring variability of neural responses: Fano factor

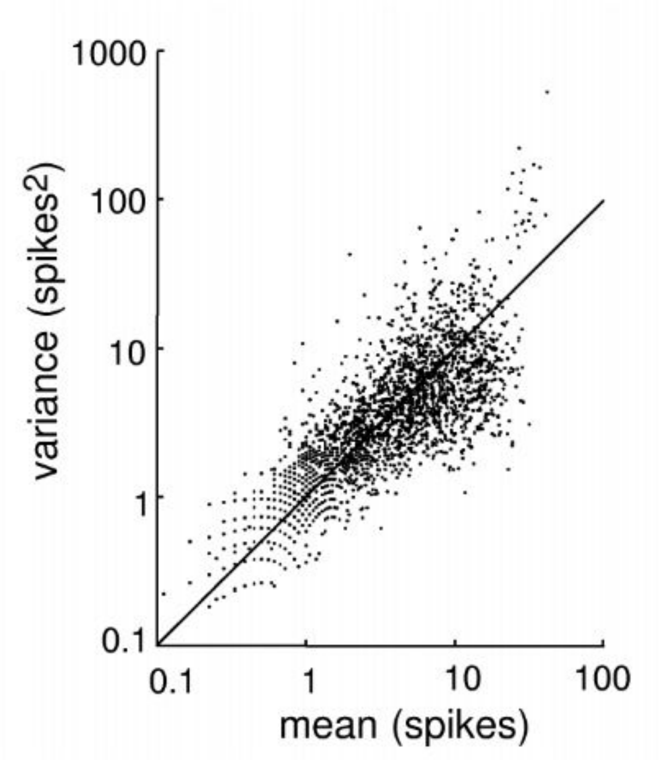

$$\text{Fano factor} = \frac{\text{Variance}}{\text{Mean}}$$

= 1 for a Poisson process

Boardwork on Fano factor

Measuring variability of neural responses: Fano factor

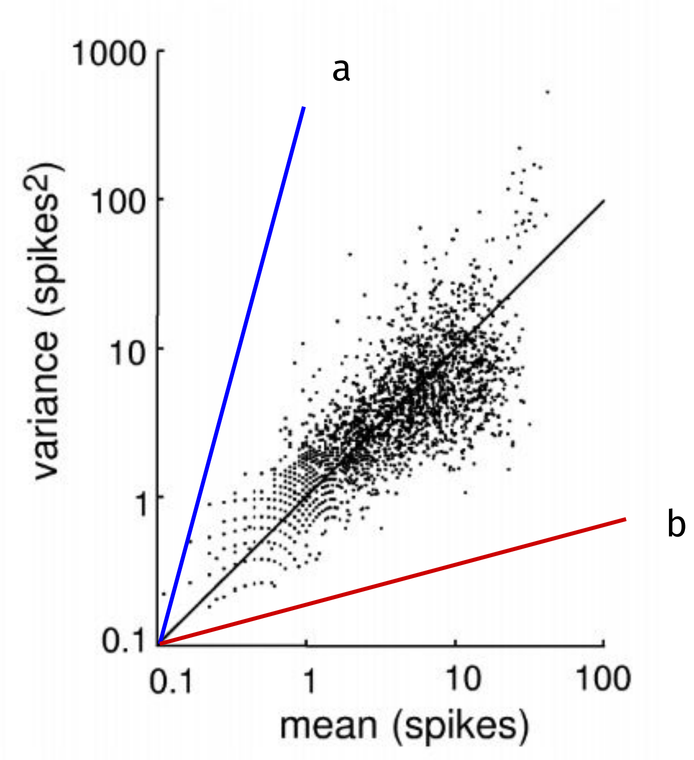

$$\text{Fano factor} = \frac{\text{Variance}}{\text{Mean}}$$

= 1 for a Poisson process

From which neuron, a or b, could we decode more reliably?

Coding breakout 1

From encoding to decoding: predicting the stimulus from neural responses

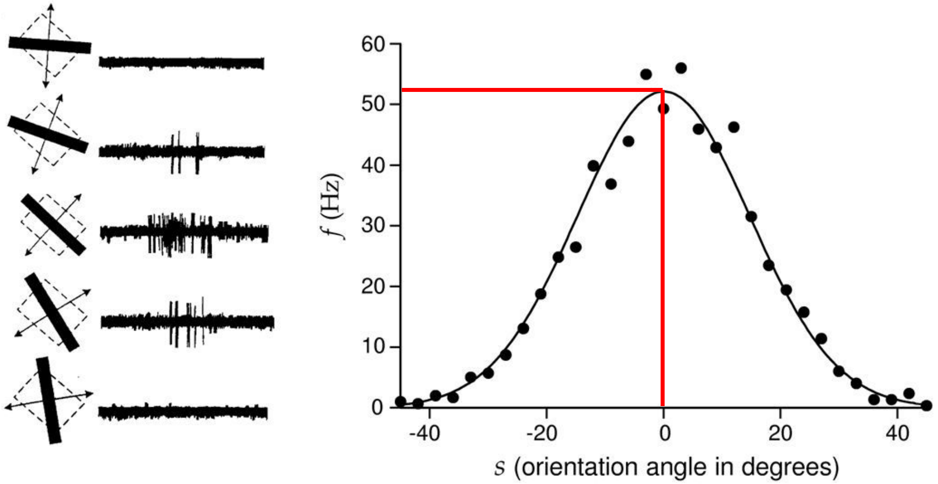

Gaussian tuning of a cortical (V1) neuron

If we observe the cell firing at 52Hz, what is the most likely orientation of the stimulus being shown?

From encoding to decoding: predicting the stimulus from neural responses

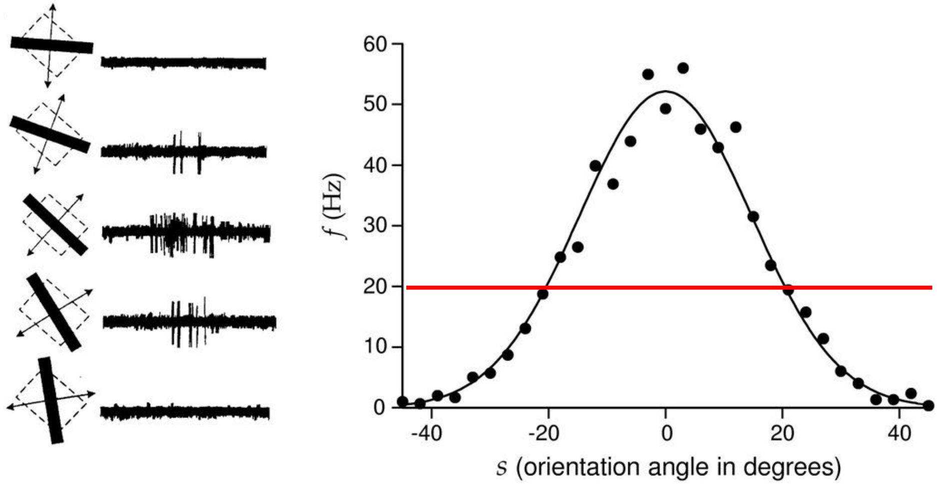

Gaussian tuning of a cortical (V1) neuron

If we observe the cell firing at 20Hz, did the animal see -20º or + 20º?

What can a population of neurons tell us that individual neurons cannot?

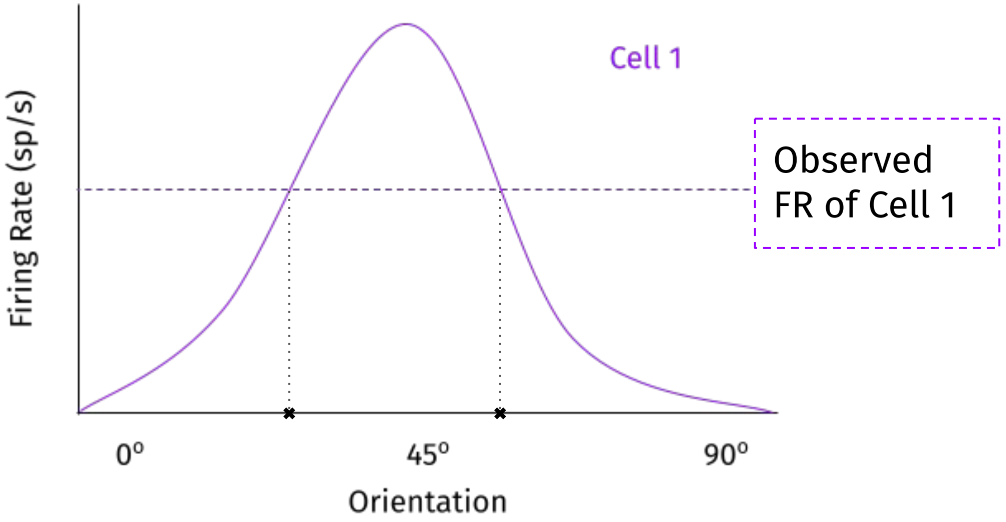

With the tuning curve of a single neuron, some stimuli can be hard to distinguish

What can a population of neurons tell us that individual neurons cannot?

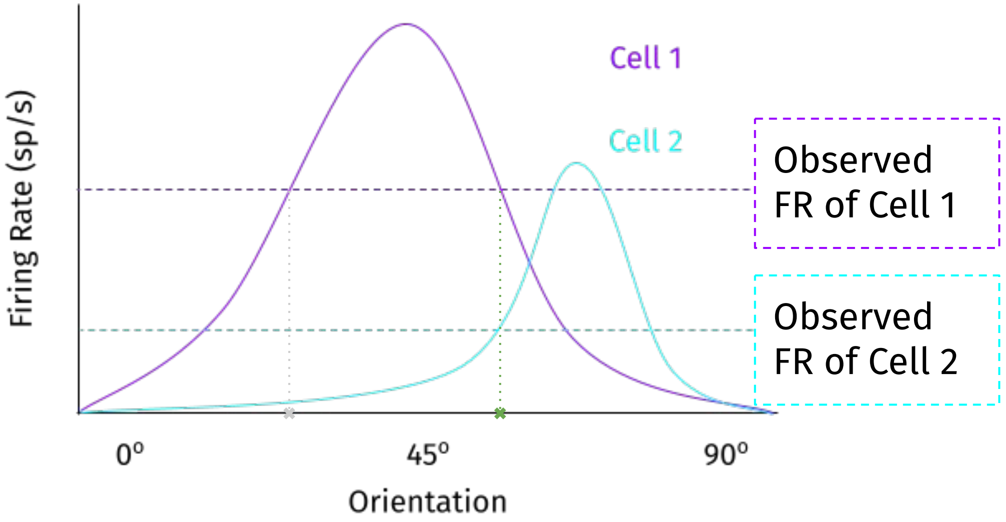

Observing the firing rate of a second neuron provides unique encoding for different stimuli at the population level

Recording from lots of neurons: Population recording technologies

- 1970s-’80s: Patch-clamp recordings of isolated single neurons by Sakmann & Neher
- 1980s-’90s: Michigan Probes and Utah arrays (~100s of neurons simultaneously)
- Currently: Neuropixels and more (~1000+ neurons)
- Calcium imaging

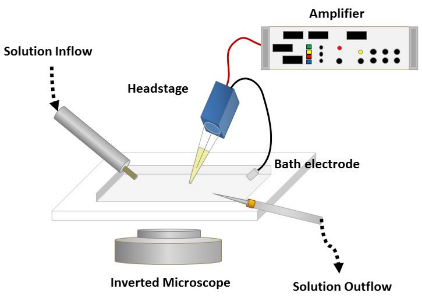

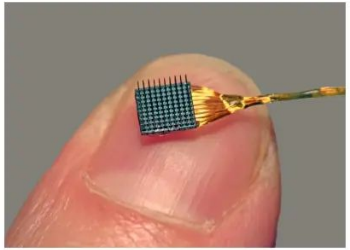

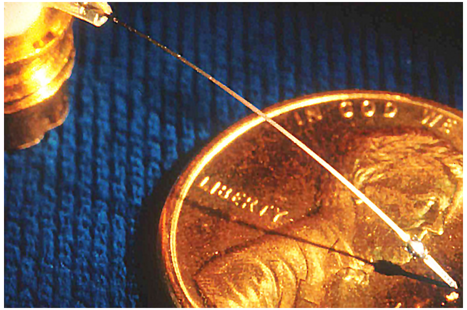

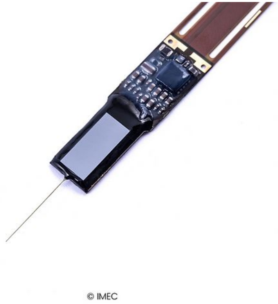

What can a population of neurons tell us that individual neurons cannot?

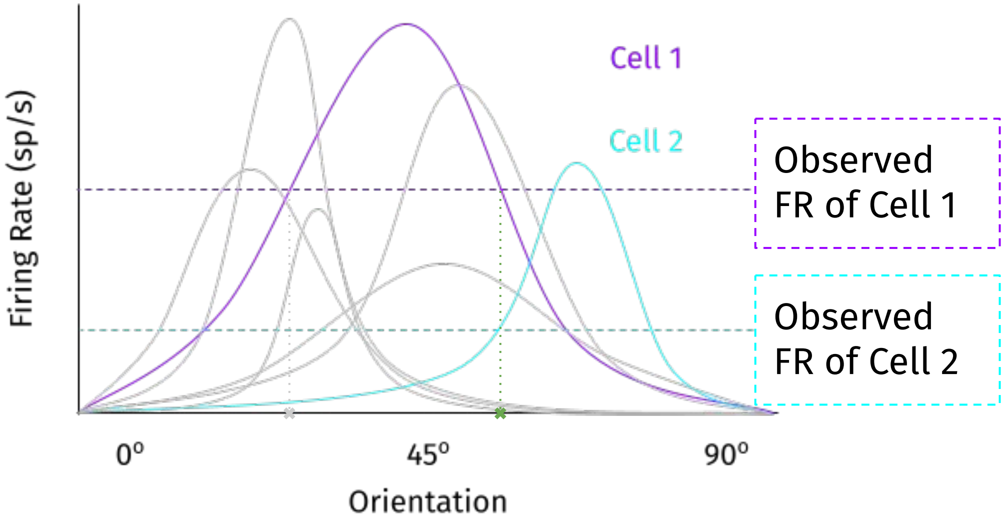

With many neurons, looking at individual tuning curves becomes intractable.

We can look at population activity in a high-dimensional “neural space” instead.

Representing population activity in high-dimensional “neural” space

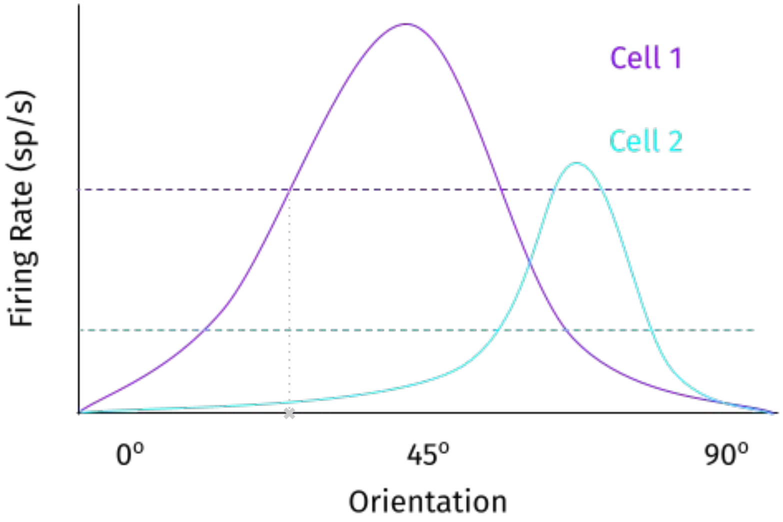

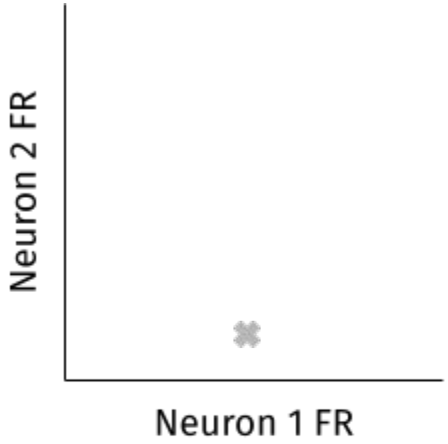

Representing population activity in high-dimensional "neural" space

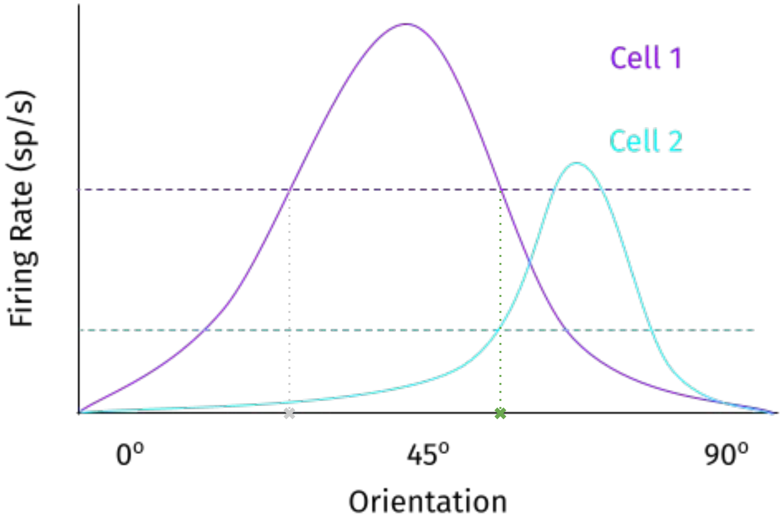

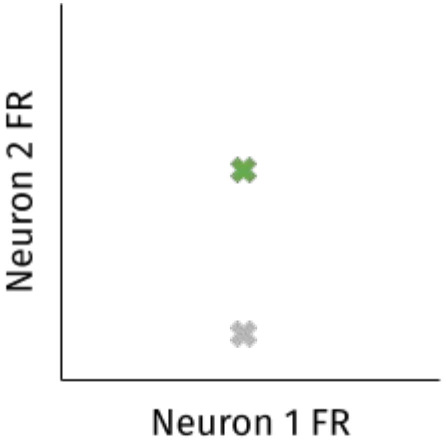

Stimulus classification for populations of neurons

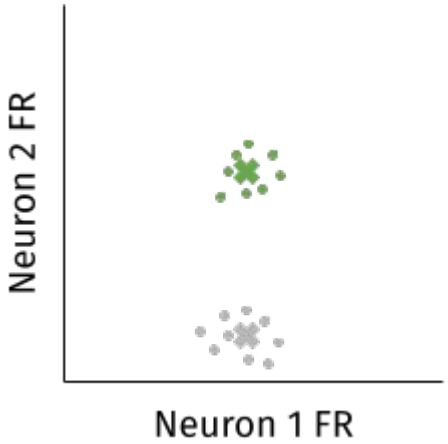

As with single neurons, we can record the population neural response over many trials.

However, we encounter the same problem: organisms don't have time to trial average in real-life situations.

Stimulus classification for populations of neurons

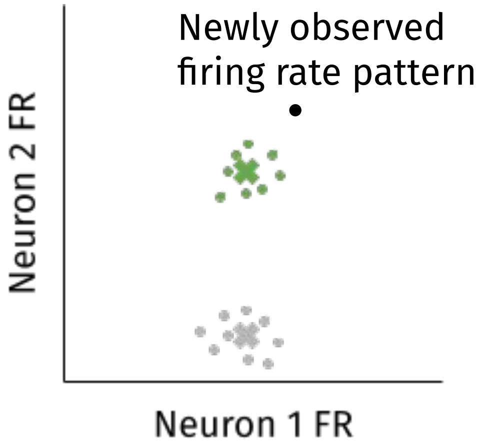

As with single neurons, we can record the population neural response over many trials.

However, we encounter the same problem: organisms don't have time to trial average in real-life situations.

Given a new firing rate pattern, how does the organism decide what stimulus is being shown?

Boardwork and Coding Exercise: Maximum likelihood estimation with multivariate Gaussians

Linear classification/decision boundaries

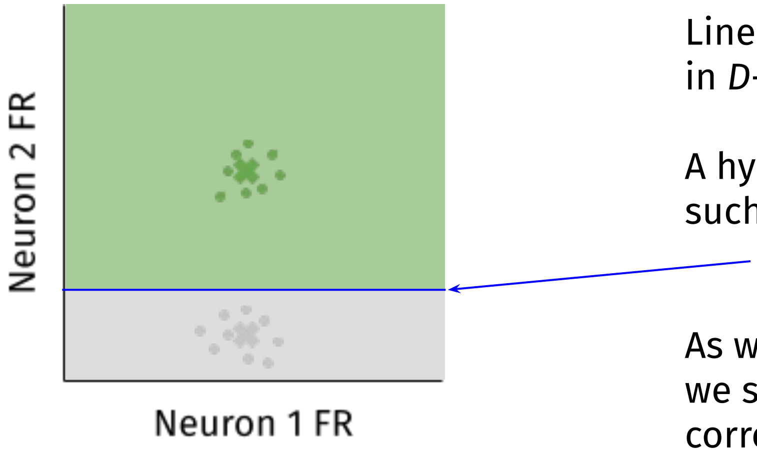

Classification with multivariate Gaussians results in a linear decision boundary.

Linear decision boundaries are hyperplanes in *D*-dimensional space.

A hyperplane is defined as the set of all *x* such that

$$y(x) = w^T x + w_o = 0$$

As we go from 2 classes to *N* classes, how do we scale the classification problem correspondingly?

From classification to continuous decoding

Brain-computer interfaces can be used to turn neural activity into movement. We want to be able to do more than classification here–we want to decode a continuous variable to enable smooth control of the end effector (robotic arm).

How do we get smooth decoding of a continuous variable like velocity of the robot arm with classification?

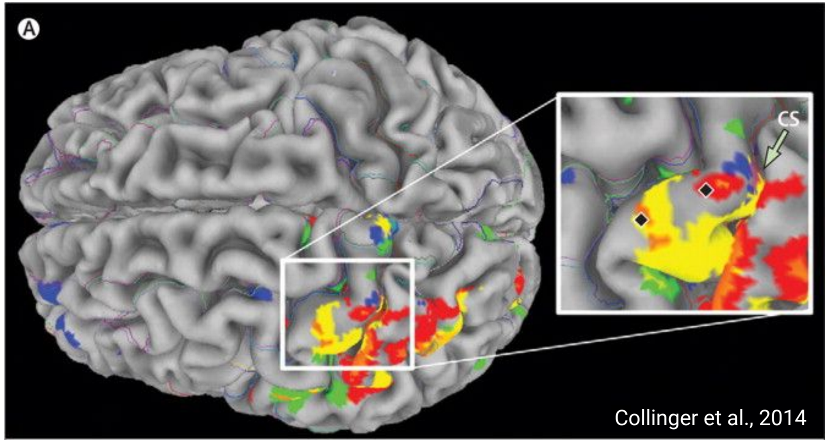

Collinger et al., 2014

Continuous decoding with linear regression

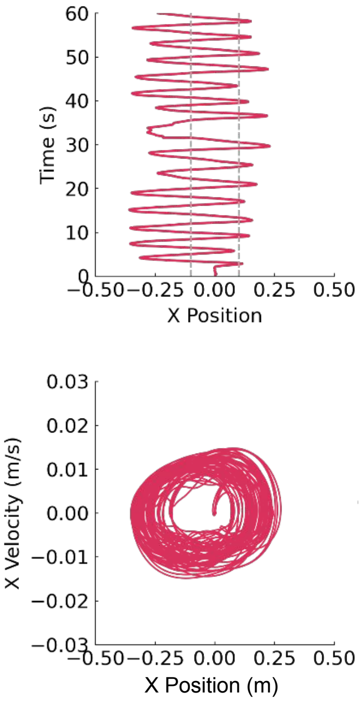

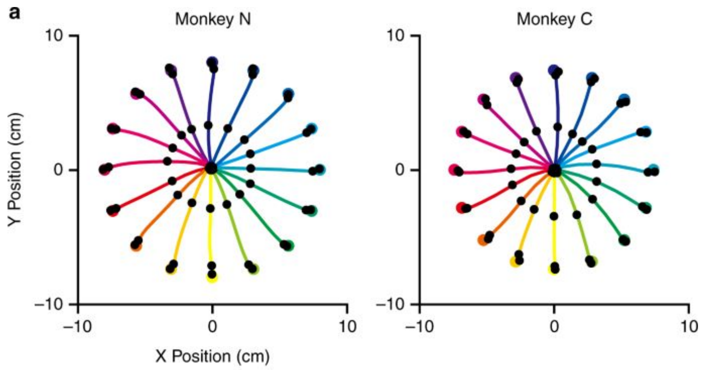

Classifier for every reach direction?
Discretization of velocity into very small bins?

Continuous decoding with linear regression

Our decoder is linear, so it takes the same form as a hyperplane:

$$y = w^T x + w_0$$

But instead of separating classes on either side of the line, it tells us how velocity maps to neural activity

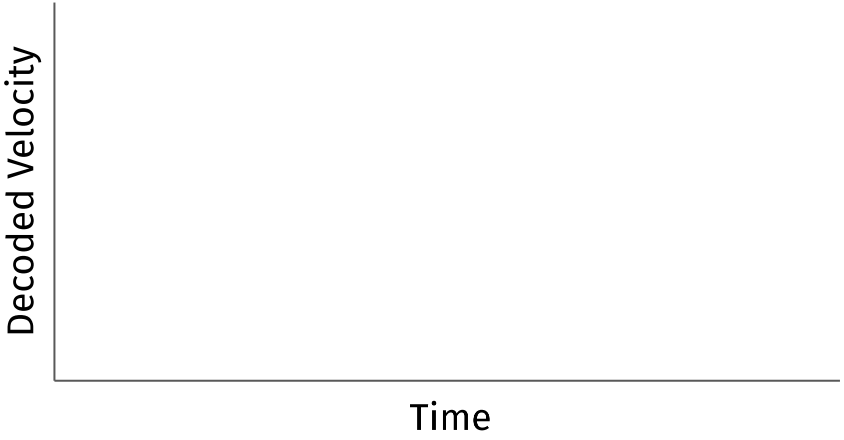

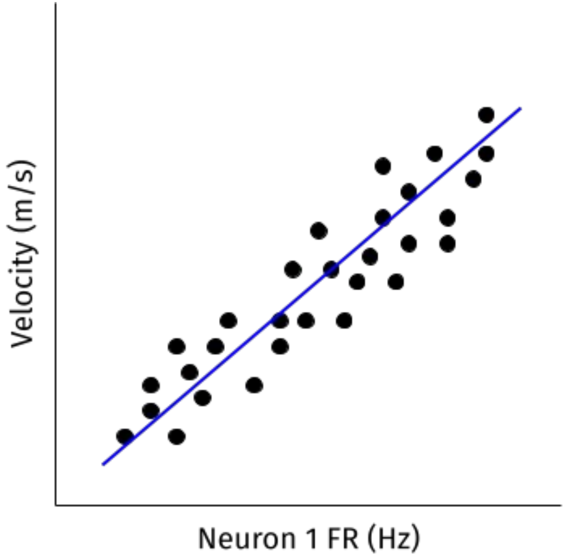

Continuous decoding with linear regression

Our decoder is linear, so it takes the same form as a hyperplane:

$$y = w^T x + w_0$$

But instead of separating classes on either side of the line, it tells us how velocity maps to neural activity

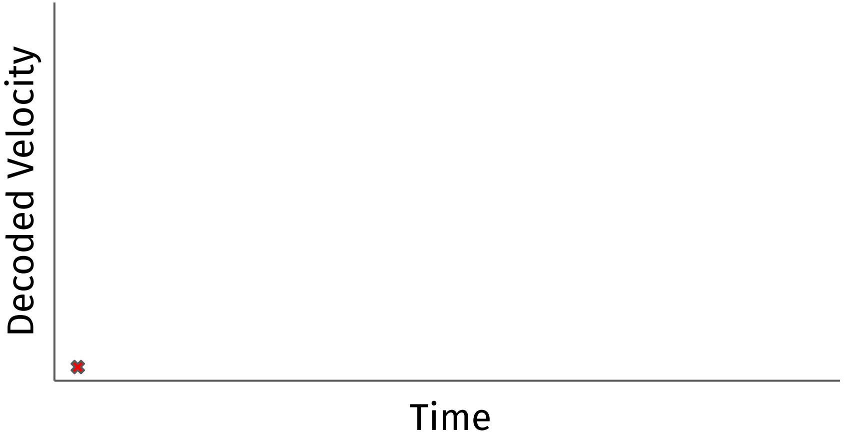

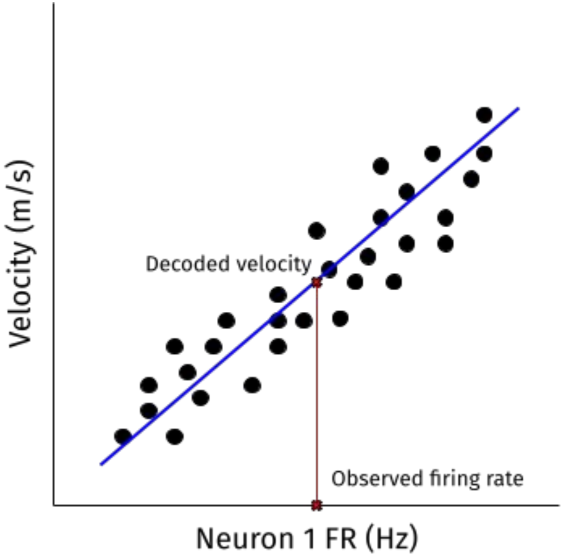

Continuous decoding with linear regression

Our decoder is linear, so it takes the same form as a hyperplane:

$$y = w^T x + w_0$$

But instead of separating classes on either side of the line, it tells us how velocity maps to neural activity

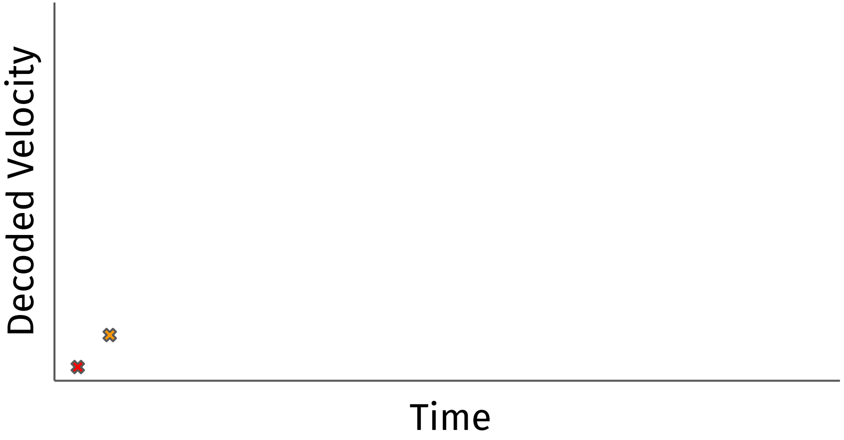

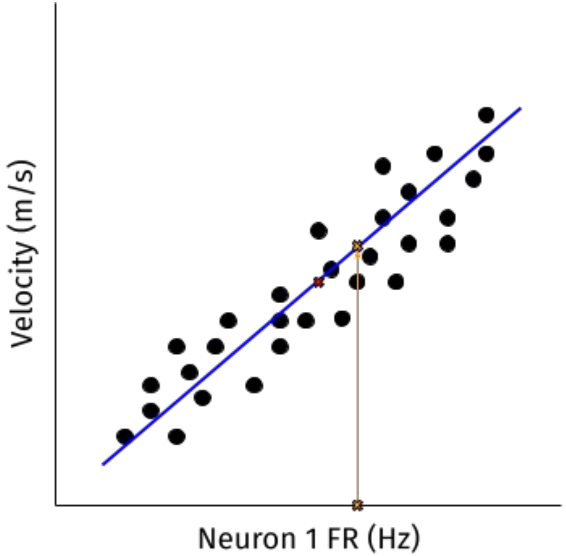

Continuous decoding with linear regression

Our decoder is linear, so it takes the same form as a hyperplane:

$$y = w^T x + w_0$$

But instead of separating classes on either side of the line, it tells us how velocity maps to neural activity

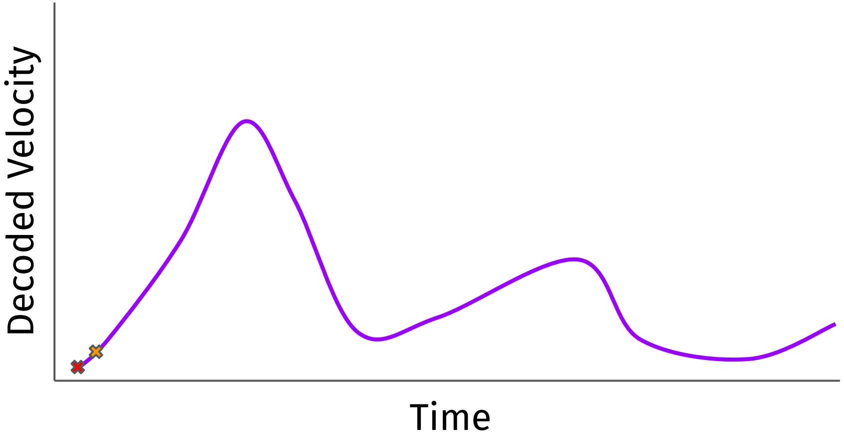

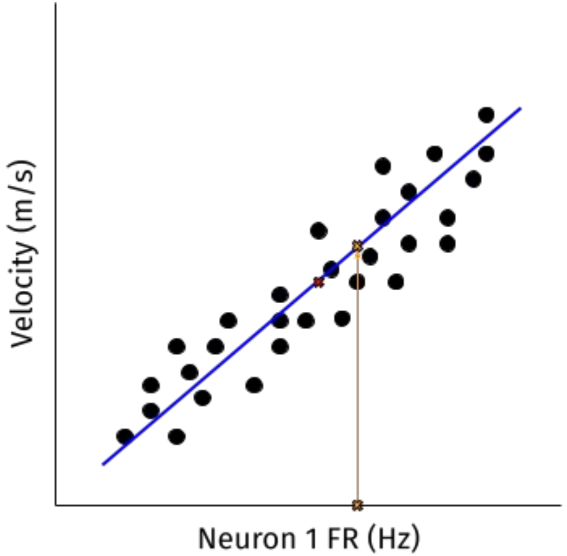

Coding breakout: linear regression

Summary: from encoding to decoding neural responses

- Neural responses vary across trials, and this variability influences how reliably we can decode from a neuron. Metrics like the Fano factor allow us to measure neural variability, but they're not perfect.
- Populations of neurons gives us improved decodablility by embedding our stimuli in a higher-dimensional space. Tools that work with one neuron also work with populations of neurons, but we need to think about "neural space" as opposed to single units.
- Continuous decoding extends upon the decoding through classification to decoding through regression. Many modern methods for motor control focus on continuous decoding techniques.