COMS30017
Computational Neuroscience

Week 3 / Video 5 / Neural decoding

Dr. Cian O'Donnell
cian.odonnell@bristol.ac.uk

Intended learning outcomes

- Understand a common approach behind decoding activity from a single neuron.
- Gain an intuition for how to scale this method up to multiple neurons.

Decoding from a single neuron

- Decoding is the process of inferring a stimulus from a neuron's spiking.
- This can give us insight into the neural code: how the brain represents information.
- As an example we can use signal detection theory to compute decoding quality for a moving dots task.

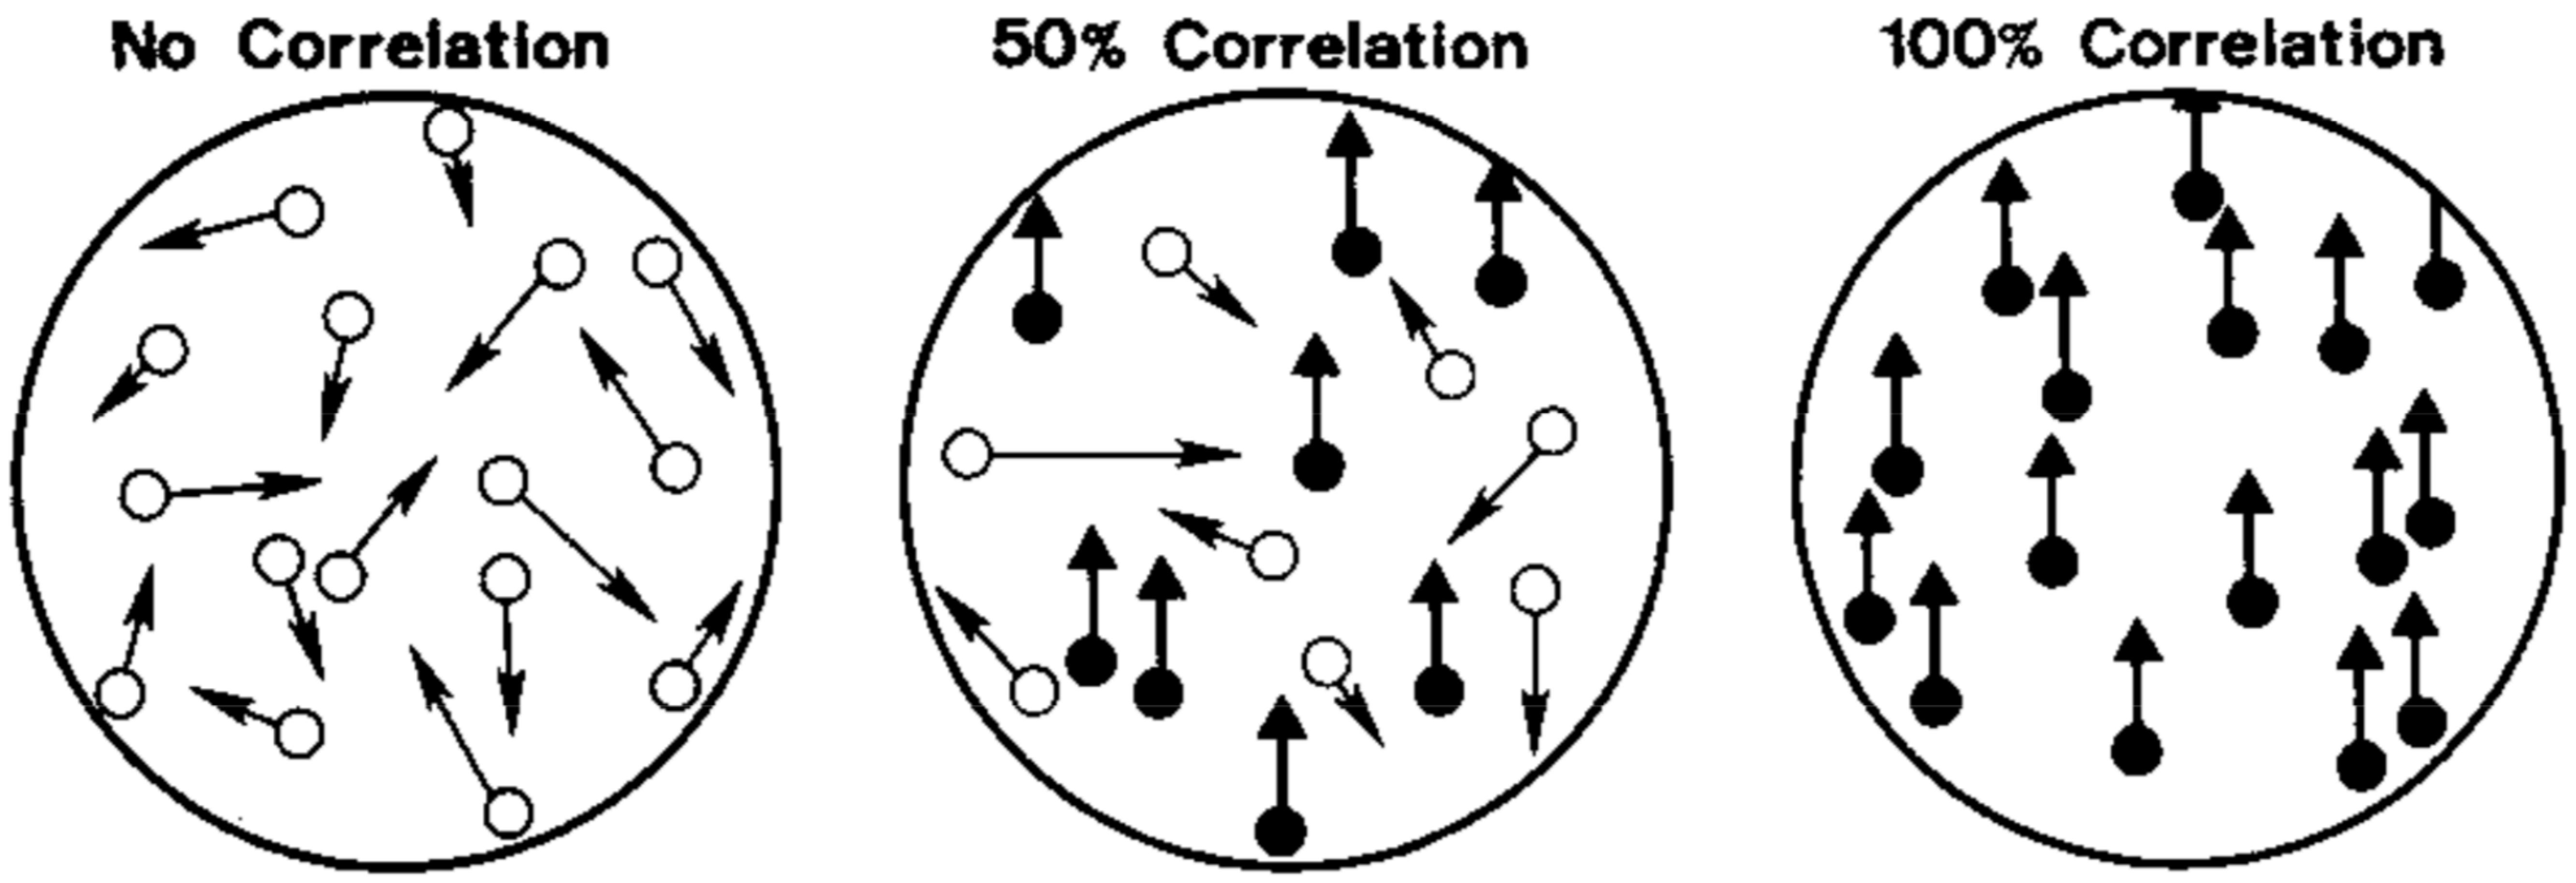

Example video: https://youtu.be/xUcwbjaGGNM?t=48

Britten, K.H., Shadlen, M.N., Newsome, W.T., Movshon, J.A., (1992). Journal of Neuroscience 12, 4745–4765.

Decoding from a single neuron

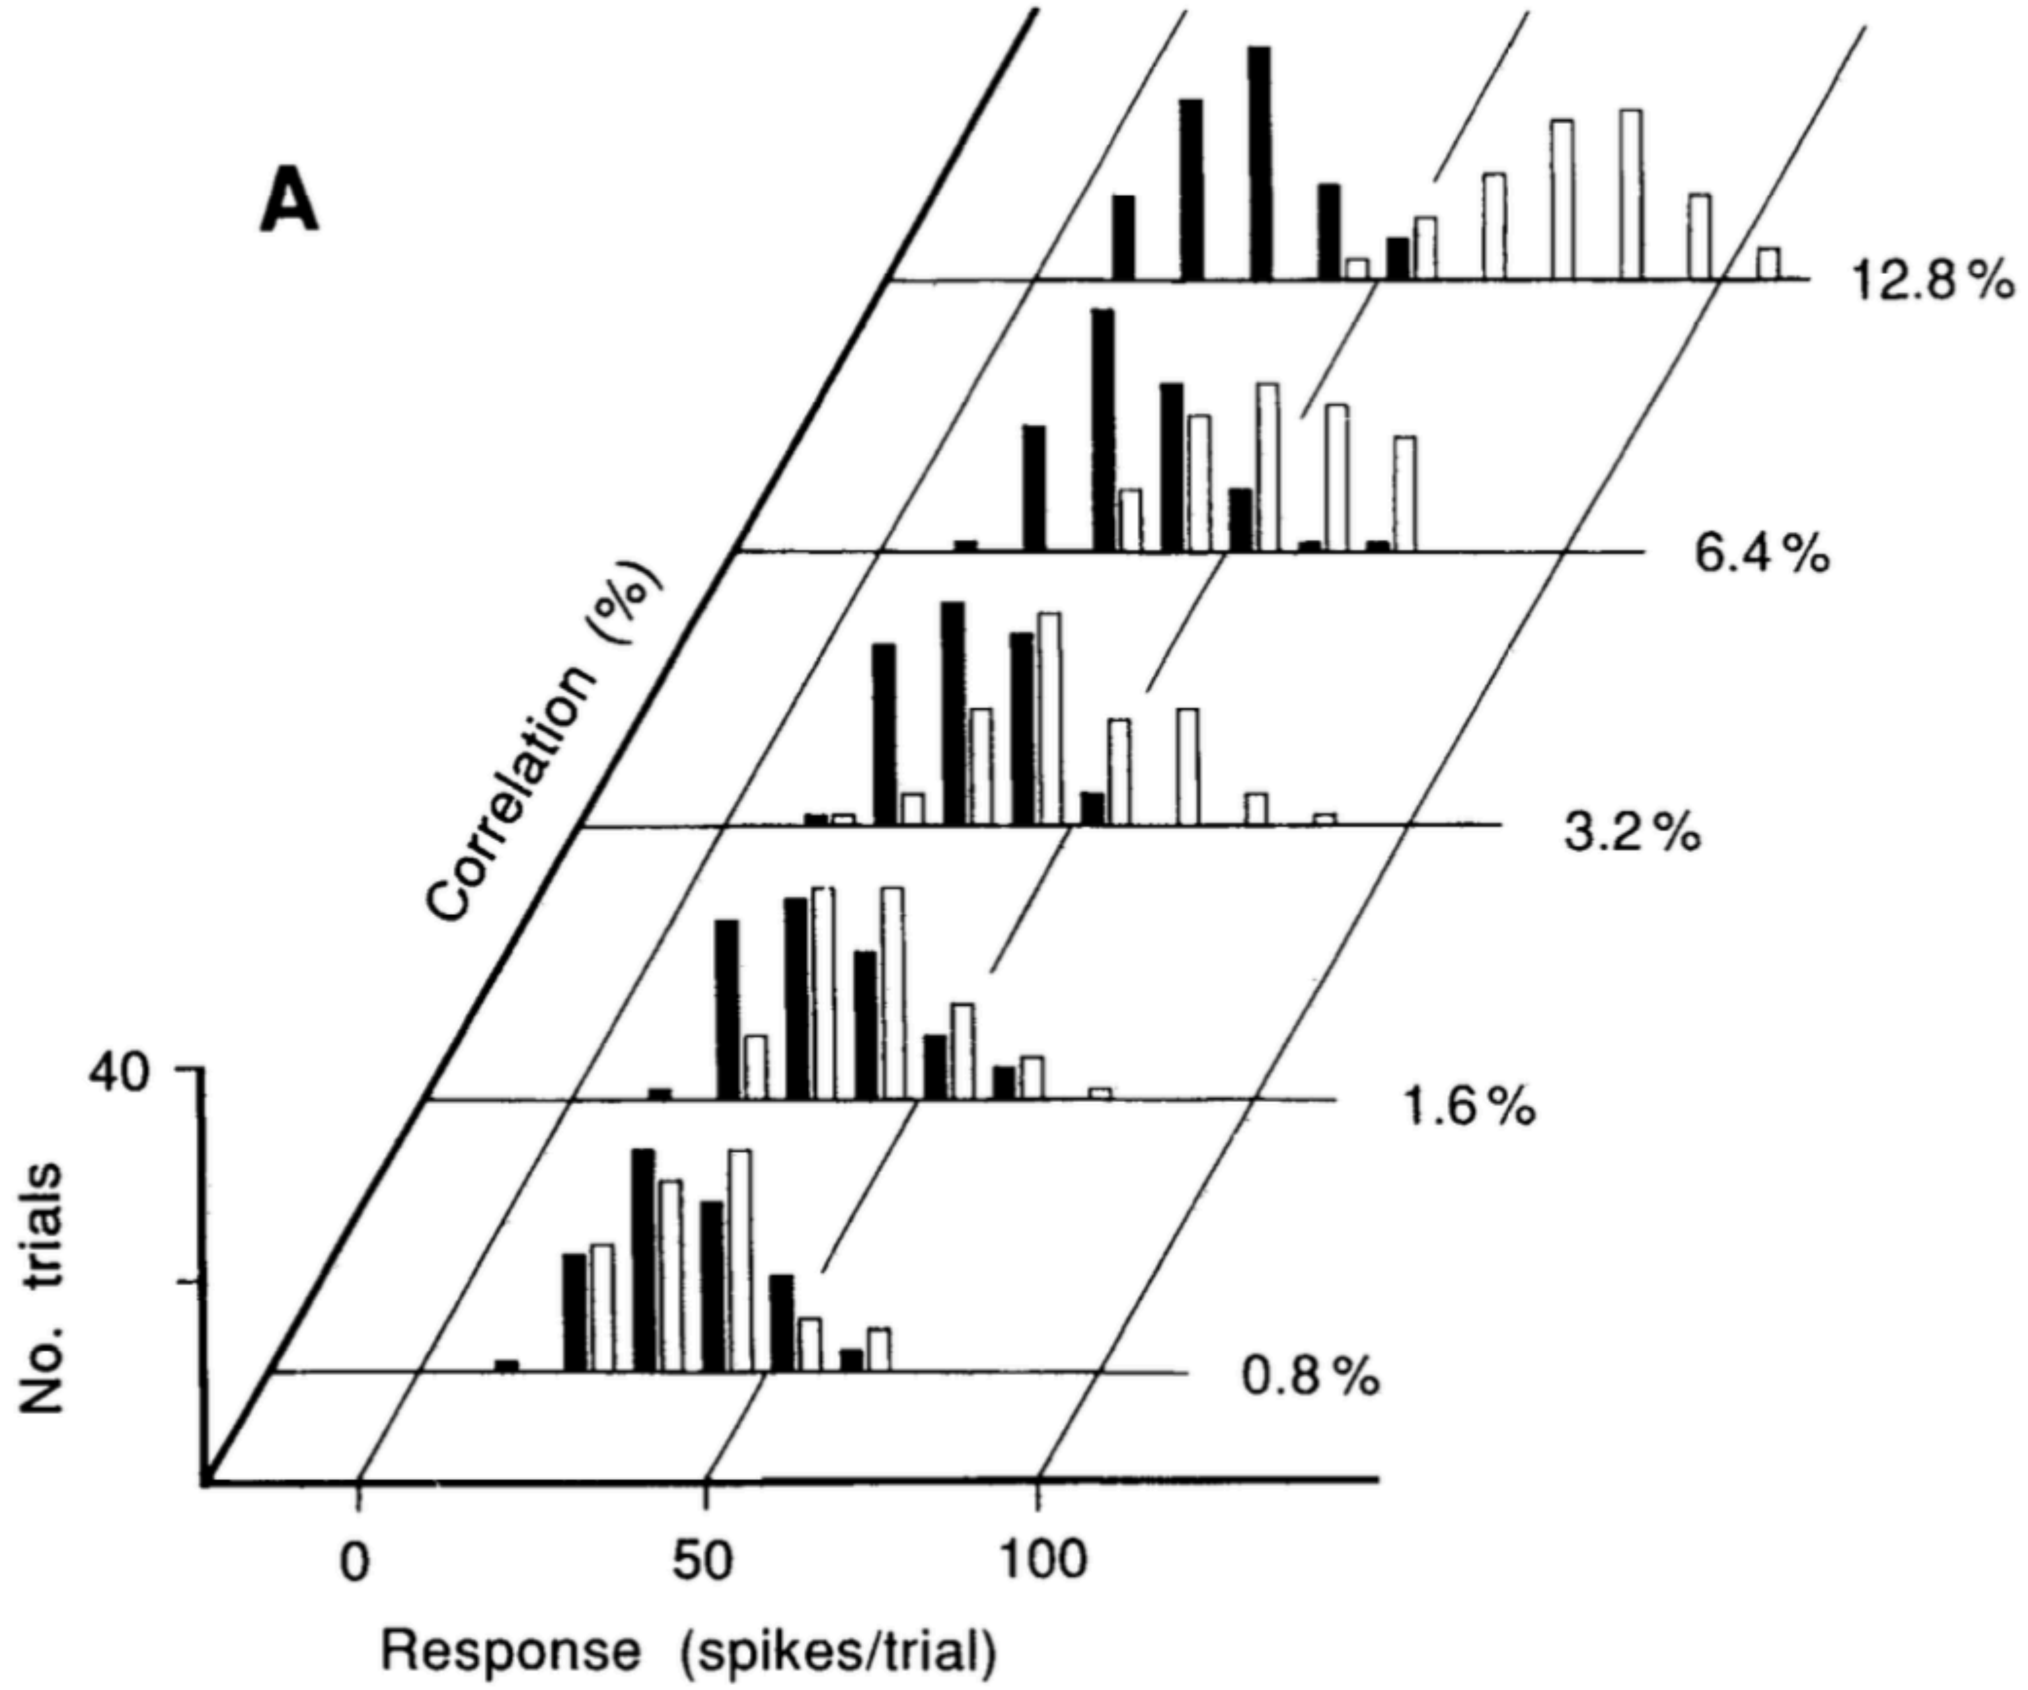

The firing rate distributions become more separated for higher moving dot image correlation levels

Britten, K.H., Shadlen, M.N., Newsome, W.T., Movshon, J.A., (1992). Journal of Neuroscience 12, 4745–4765.

Decoding from a single neuron

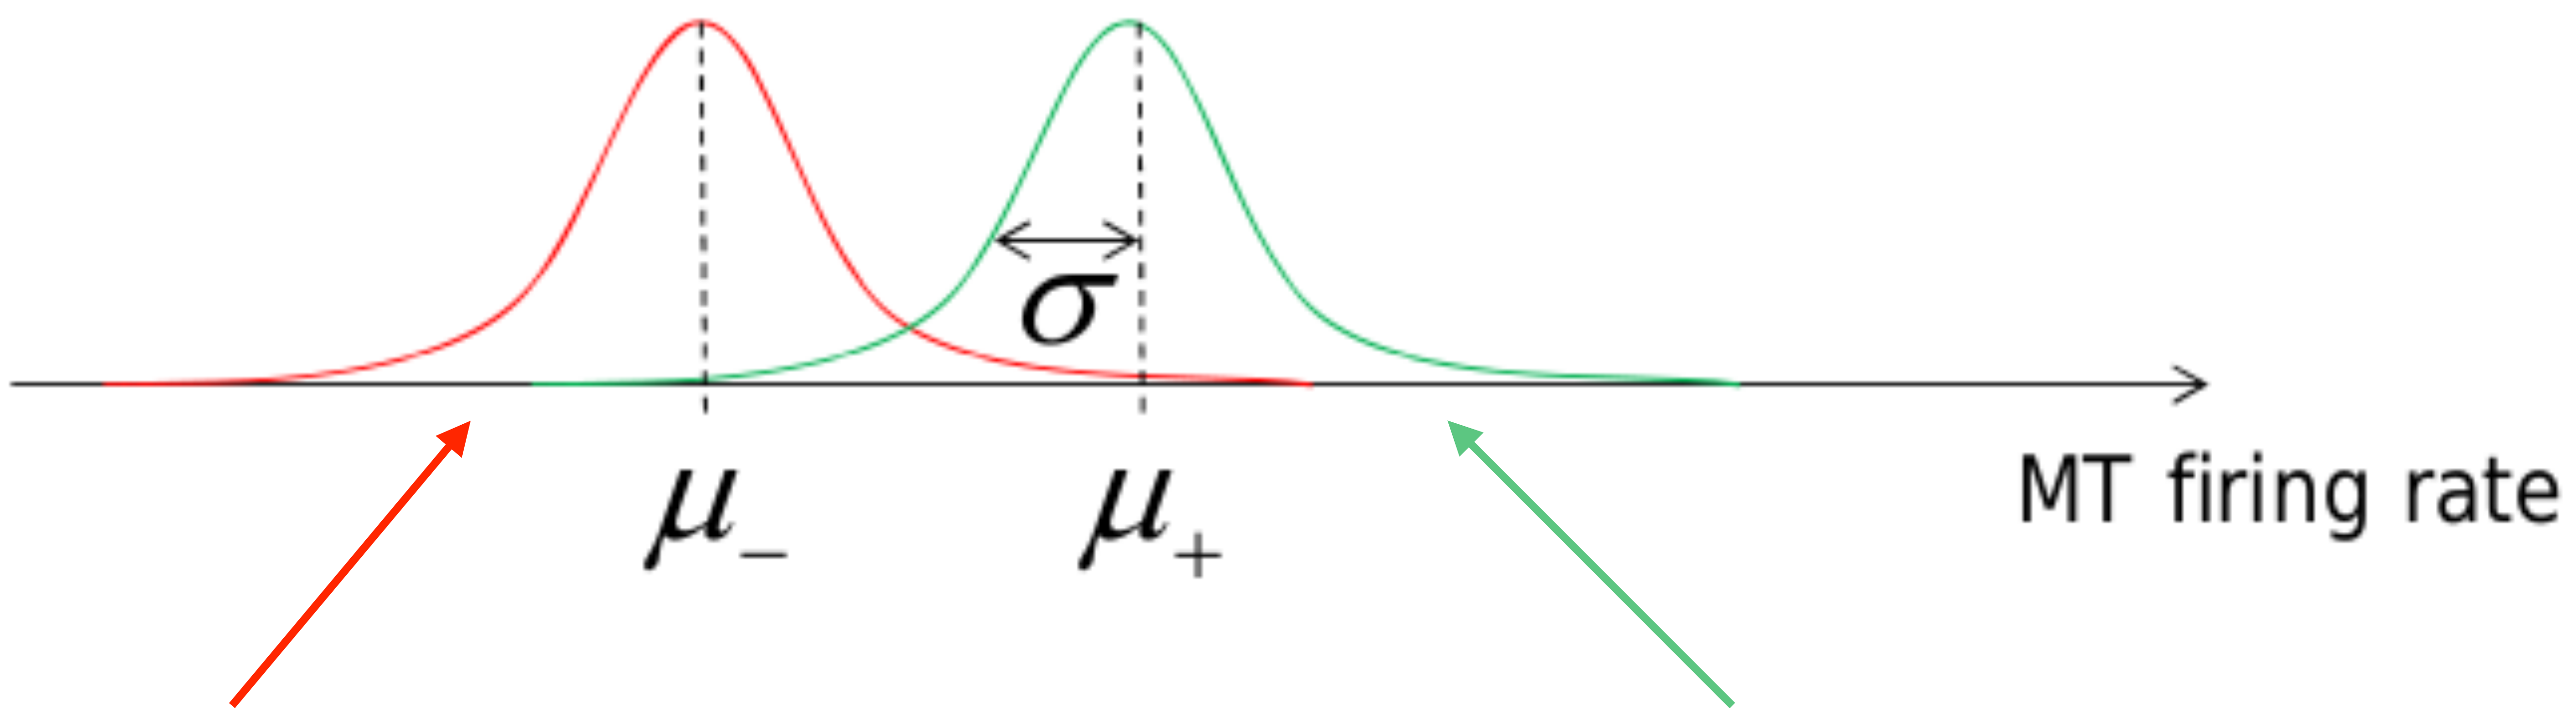

$$P(\text{rate}\,|\,\text{corr}=0) = \mathcal{N}(\mu_{-}, \sigma) \qquad P(\text{rate}\,|\,\text{corr}>0) = \mathcal{N}(\mu_{+}, \sigma)$$

- Fit a probability distribution to each response per stimulus condition: $P(\text{response}\,|\,\text{stimulus})$
- Now, if we are given a value of the response for a trial where we don't know the stimulus, e.g. rate = 52 Hz, we ask the question:
  "Under which stimulus condition was that response more probable?"
- In this 1D, 2-choice task, the decision making rule simplifies to putting a decision boundary at the crossover point in the two probability distributions (if your goal is to minimize overall errors, as opposed to caring differently about false positives vs false negatives.)

Figure from Rafal Bogacz

Decoding from a single neuron

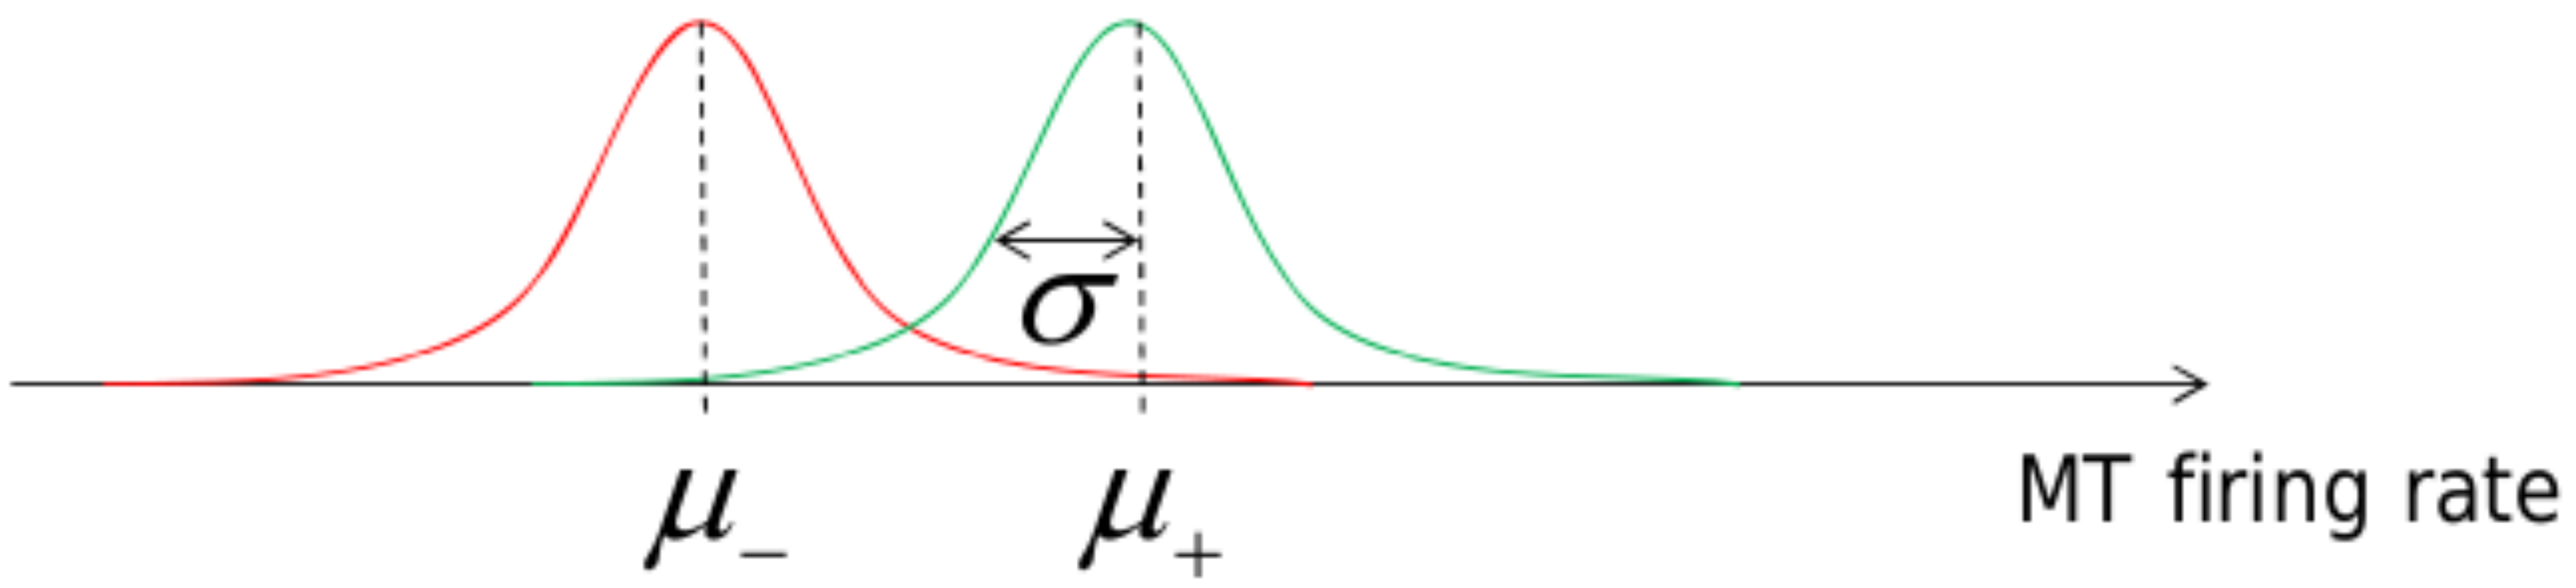

d' (pronounced "dee-prime") is a measure of discriminability of two normal distributions.

$$d' = \frac{\mu_+ - \mu_-}{\sigma}$$

$$P(\text{correct}) = P(\mathcal{N}(\mu_+, \sigma) > \mathcal{N}(\mu_-, \sigma))$$
$$= \Phi(d'/\sqrt{2})$$

Where Φ is the cumulative distribution for the normal distribution.

Figure from Rafal Bogacz

Decoding from a single neuron

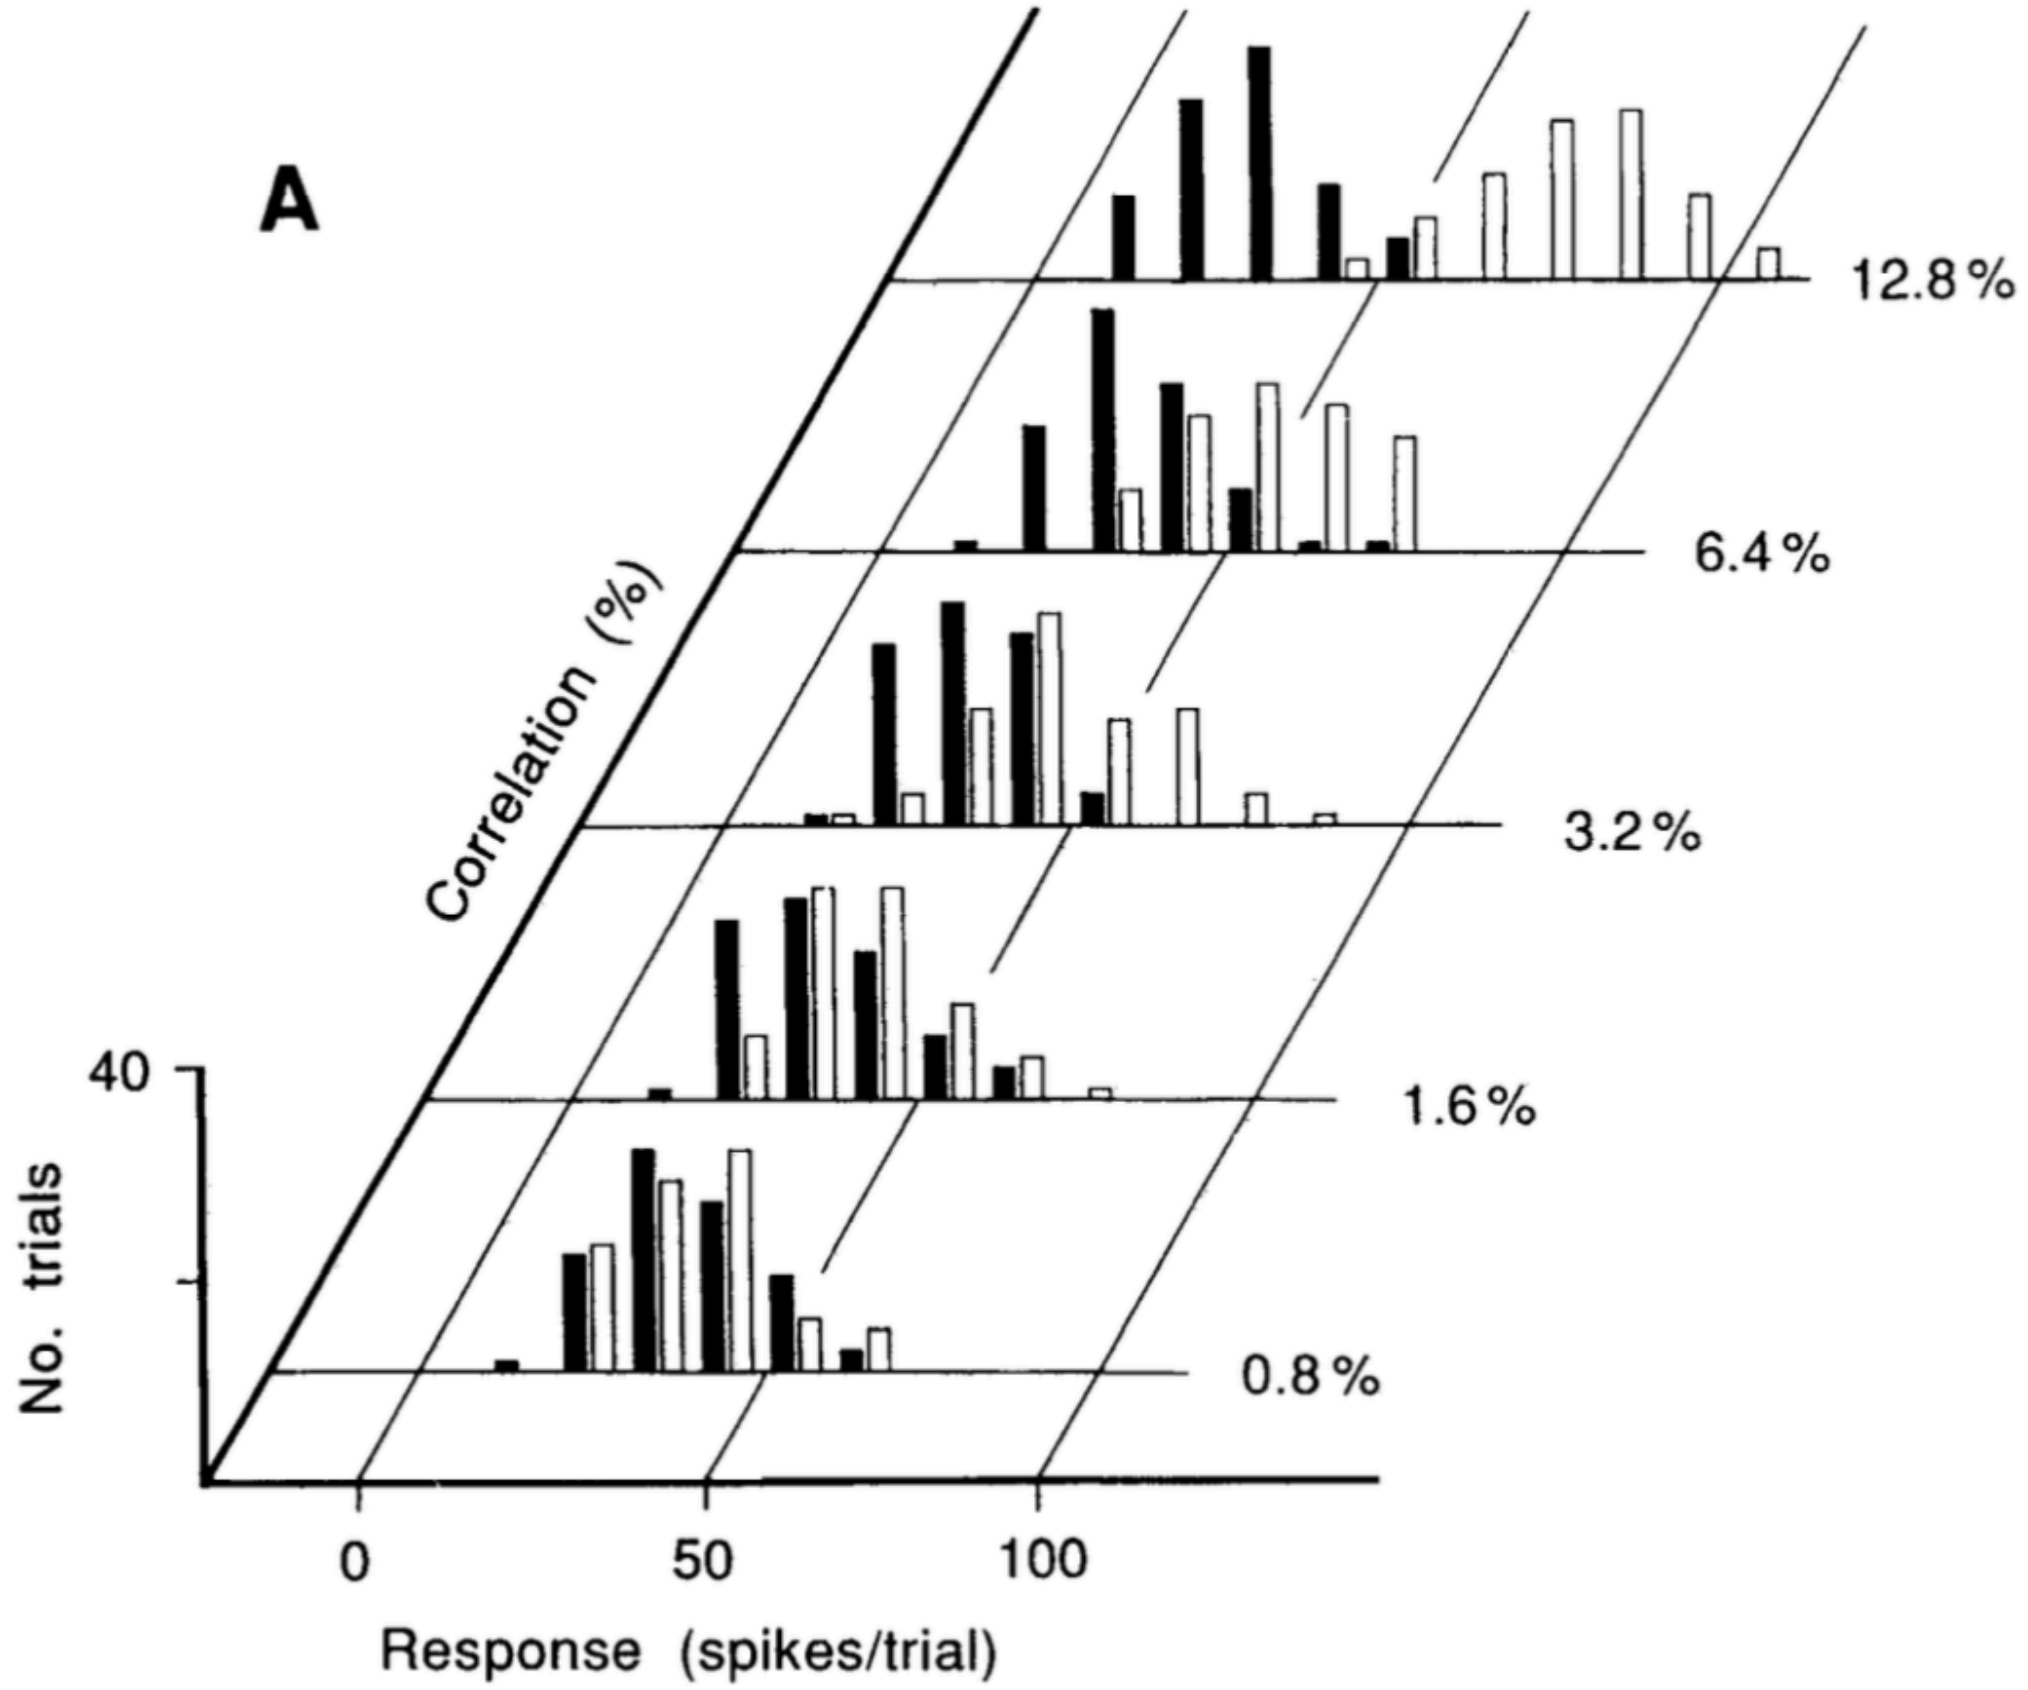

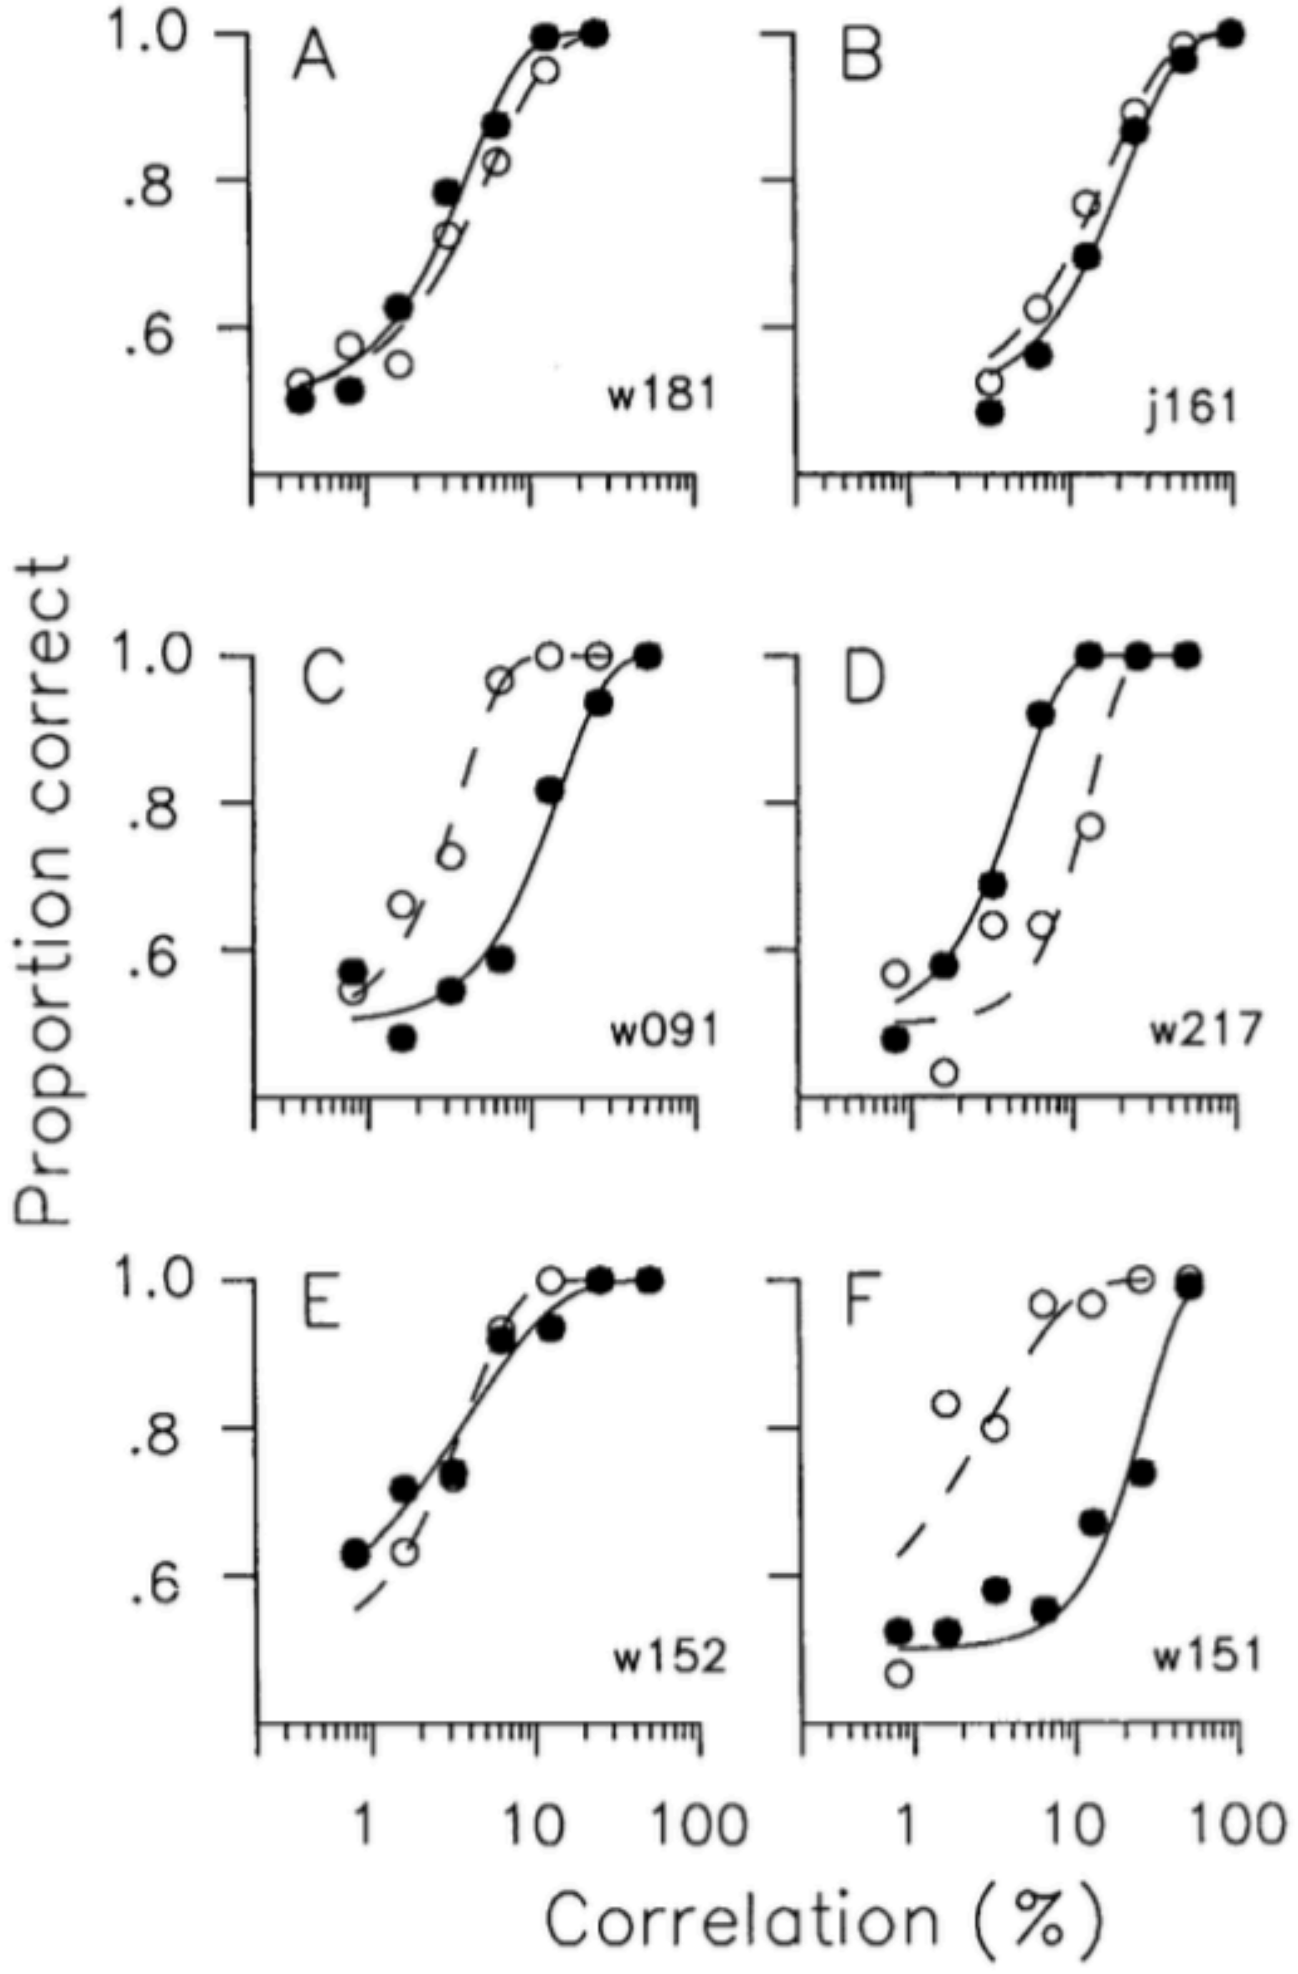

Figure 6. Psychometric and neurometric functions obtained in six experiments. The *open symbols* and *broken lines* depict psychometric data, while the *solid symbols* and *solid lines* represent neurometric data. The six examples illustrate the range of relationships present in our data. *A*, Results of the experiment illustrated in Figures 4 and 5. Psychophysical and neuronal data were statistically indistinguishable in this experiment. Thresholds and slope parameters are given in the captions for Figures 4 and 5. *B*, A second experiment in which psychometric and neurometric data were statistically indistinguishable. Psychometric α = 17.8% correlation, $\beta = 1.20$; neurometric α = 23.0% correlation, β = 1.31. *C*, An experiment in which psychophysical threshold was substantially lower than neuronal threshold. Psychometric α = 3.7% correlation, $\beta = 1.68$; neurometric α = 14.8% correlation, $\beta = 1.49$. *D*, An experiment in which neuronal threshold was substantially lower than psychophysical threshold. Psychometric α = 13.0% correlation, β = 2.15; neurometric α = 4.7% correlation, $\beta = 1.58$. *E*, An experiment in which thresholds were similar but slopes were dissimilar. Psychometric α = 3.9% correlation, $\beta = 1.36$; neurometric α = 4.0% correlation, β = 0.79. *F*, An experiment in which threshold and slope were dissimilar. Psychometric α = 3.1% correlation, $\beta = 0.91$; neurometric α = 27.0% correlation, $\beta = 1.81$.

Britten, K.H., Shadlen, M.N., Newsome, W.T., Movshon, J.A., (1992). Journal of Neuroscience 12, 4745–4765.

Multiple neurons

- What if we have data from multiple neurons recorded simultaneously?
  With calcium imaging or new "neuropixels" electrical probes it is now routine for many labs to record from ~100s of neurons in animals.

- In this case there are two main routes for decoding:

  1. Extend the same probabilistic approach to calculate the joint conditional response to stimuli:
     $P(\text{rate}_1, \text{rate}_2, \ldots, \text{rate}_N \,|\, \text{stimulus})$.
     You need a good model for the joint activity, which can be tricky. But powerful because you can compare the performance of various models.

  2. Turn it into a standard regression/classification statistical problem: trying to infer an independent variable (the stimulus value) from a dependent variable (neural firing rates).
     Then there are several common choices:
     - Generalised linear models (linear regression, logistic regression, etc)
     - Support vector machines
     - Neural networks
     - Gradient boosting

     These tend to perform well, but are usually pretty black-box so can't tell us much about how the brain works.

Summary

- Neural decoding is the process of inferring the value of an external variable from brain activity: “mind reading”.

- Modern methods performed on hundreds of neurons can do very well on low-dimensional tasks (e.g. predicting an animal’s 2D location or which of a small set of visual stimuli were shown).

- But not yet at the point where it could infer an entire visual scene, or someone’s “thoughts” — we will need some breakthroughs in understanding how the brain works before that’s possible.

- If you are interested in reading more here is an excellent tutorial paper:
  “Machine Learning for Neural Decoding"
  JI Glaser, AS Benjamin, RH Chowdhury, MG Perich, LE Miller and KP Kording
  eNeuro 31 July 2020, 7 (4)
  https://www.eneuro.org/content/eneuro/7/4/ENEURO.0506-19.2020.full.pdf