Neuroscience, cognitive science and machine learning

Konrad Kording (with inspiration by Peter Dayan)

May 12th 2010

Level 3 Marr: Implementation

- In our brain - there are just spikes
- Neither suggestive of generative nor discriminative algorithm of learning

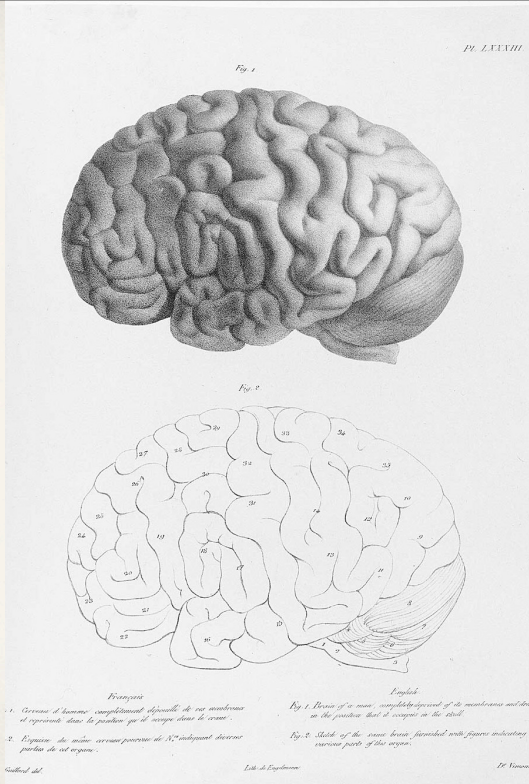

Neurons generate spikes

- Normally we do not care about generative model - just performance for decisions
- But we do want to understand the brain
- Which means understand the model which generates spikes
- Understand structure, function and objectives

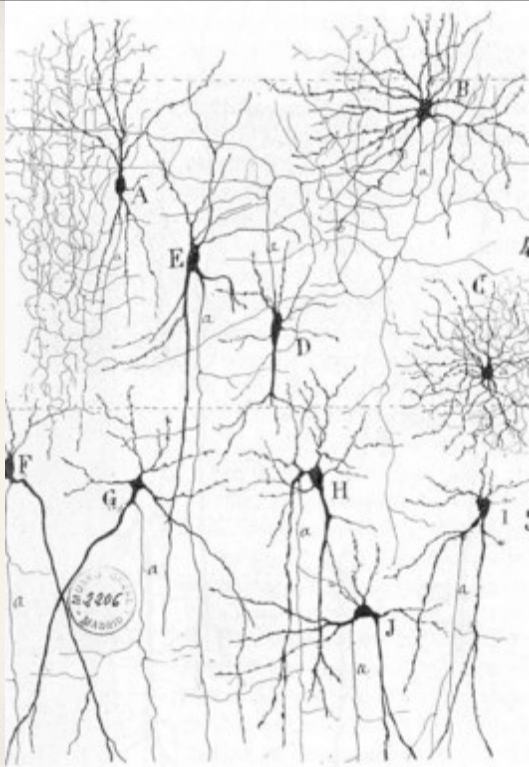

It is a big problem

- Currently we can only record from small percentage of neurons at a time
- and yet we want to understand how it works

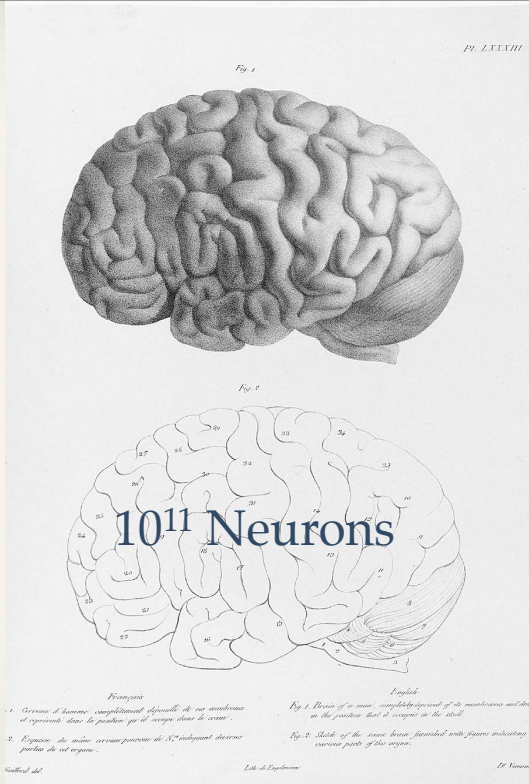

Hope #1

- Brain consists of areas
- Within each area neurons are essentially all identical
- Apart from small number of parameters
- Strong analogy to computer programs, e.g. neural networks
- Map how activity relates to outside world

Hope #2

- Brain consists of groups of similar neurons
- Neurons within each group are essentially identical
- While their relation to the outside world may be complicated their interactions may be simple
- Therefore, recording from each of the groups is enough

Hope #3

- Neurons may all be different from one another
- But their learning rules may be identical
- If we know the statistics of inputs we can predict distributions
- If we understand how they learn we might understand how it works

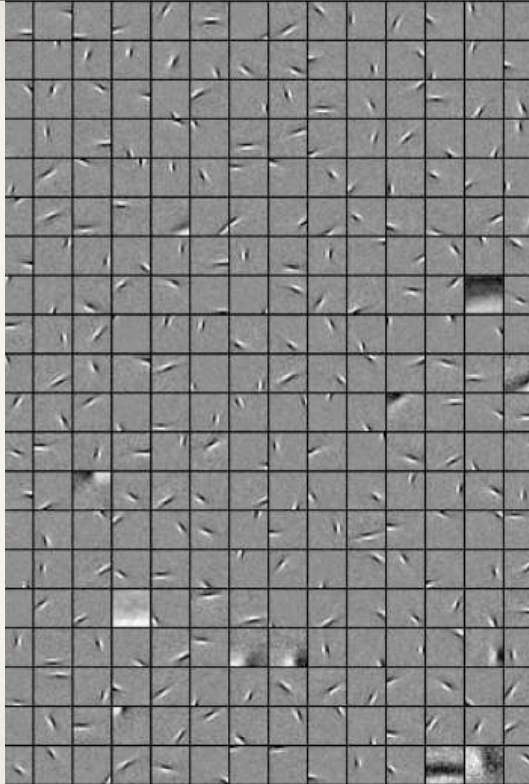

Hope #4: New technologies

- New technologies accelerate rapidly
- Record from all Neurons at the same time
- And then?

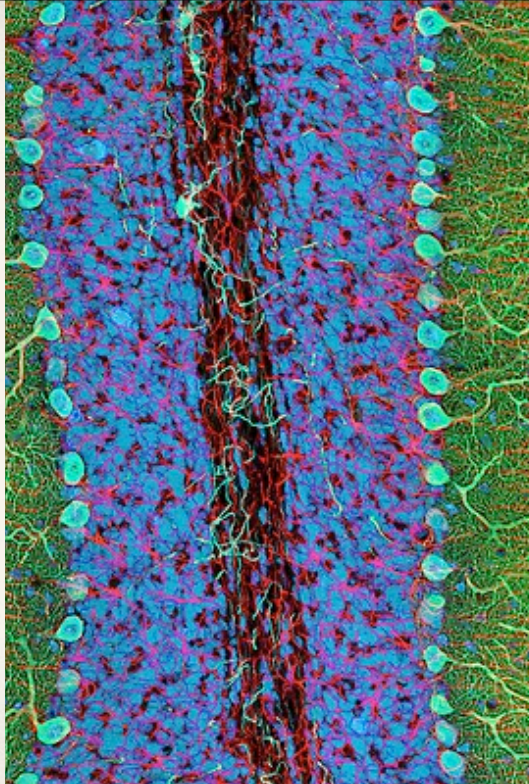

A well known Law

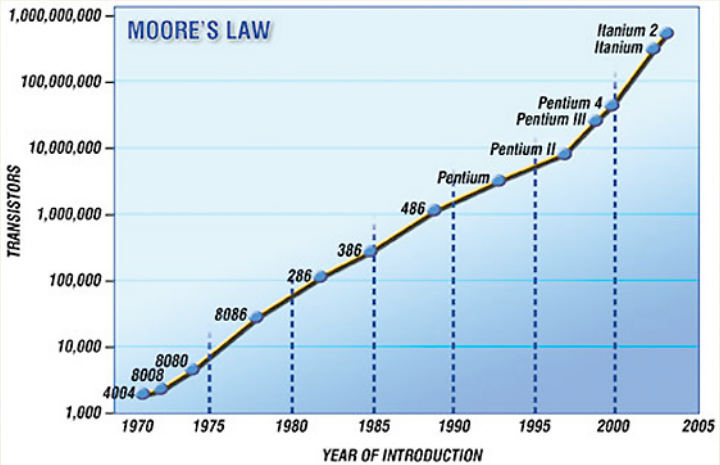

- My phone = fastest computer on the planet in 1980

Implications

- Focus not on speed but on scaling
- O(n) notation
- What is O(n), O(n log n) etc?

A Much Less known Law

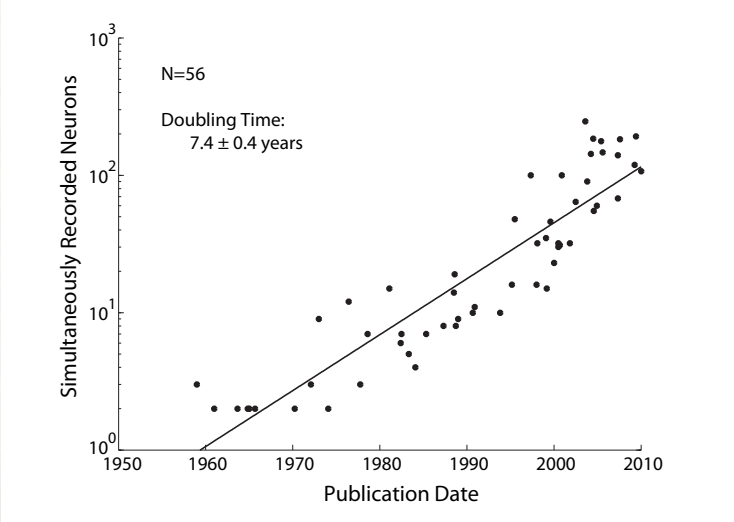

Implications

- Analyze scaling behavior
- Computer time per neuron (no worries)
- Information learned per neuron!

Outline of talk

- Understand how neurons relate to the outside world
- How neurons represent uncertainty
- How neurons relate to one another

Part 1: How neurons relate to the outside world

Experimental studies

- Vary the stimulus
- Vary the behavioral demands
- Measure behavior
- Measure neural signals

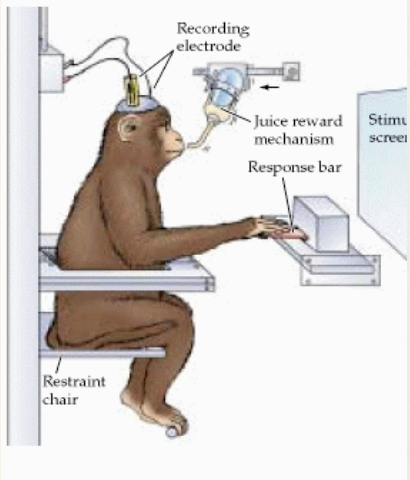

unknown image source

Motor Tuning in M1

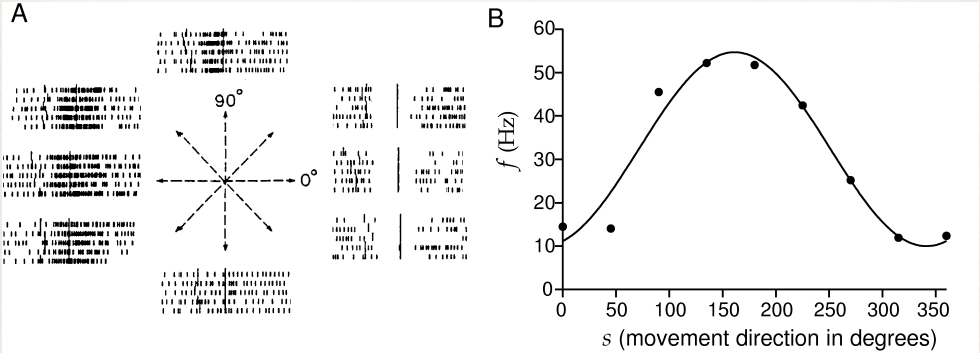

Georgopoulos et al (1982)

Slide adapted from Dayan

Orientation Tuning in V1

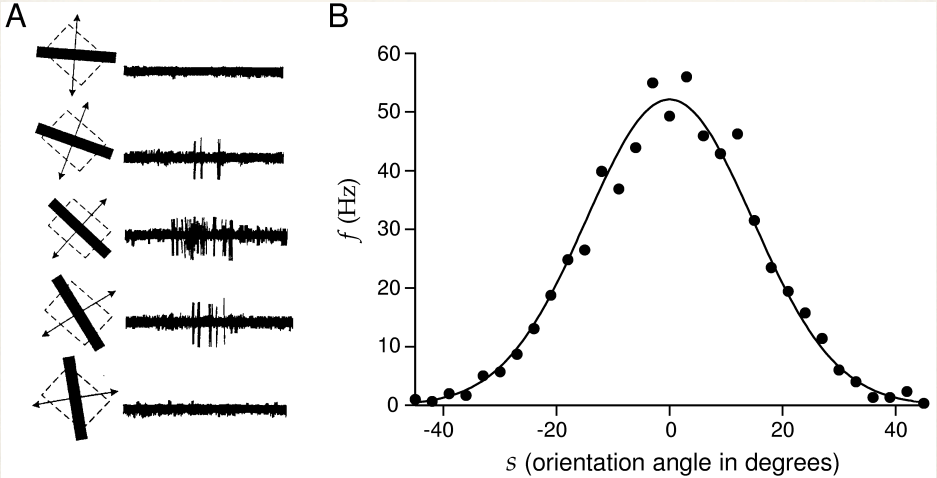

Slide adapted from Dayan

Hubel & Wiesel (1968)

Disparity Tuning in V1

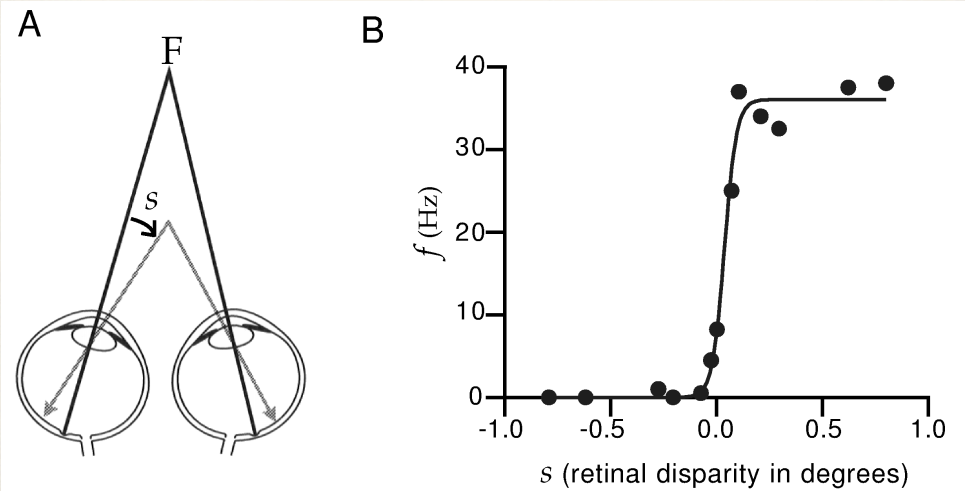

Slide adapted from Dayan

Poggio & Talbot (1981)

Spatial Tuning

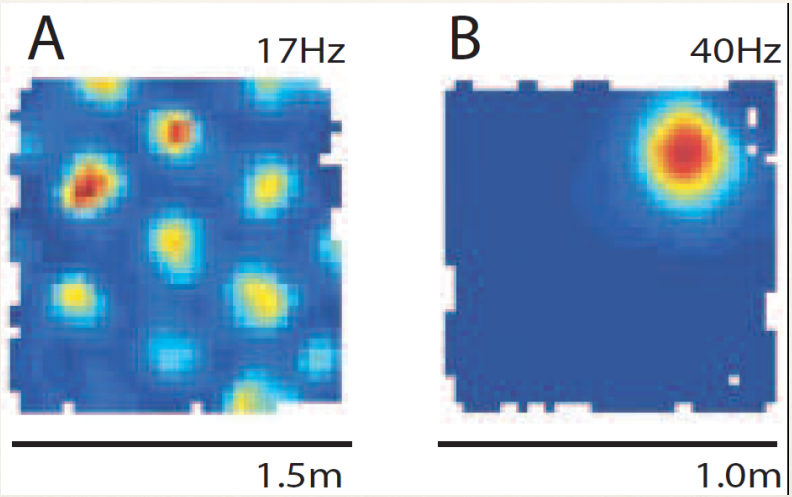

entorhinal cortex

hippocampus

Moser et al

Slide adapted from Dayan

Bygone Actress Tuning

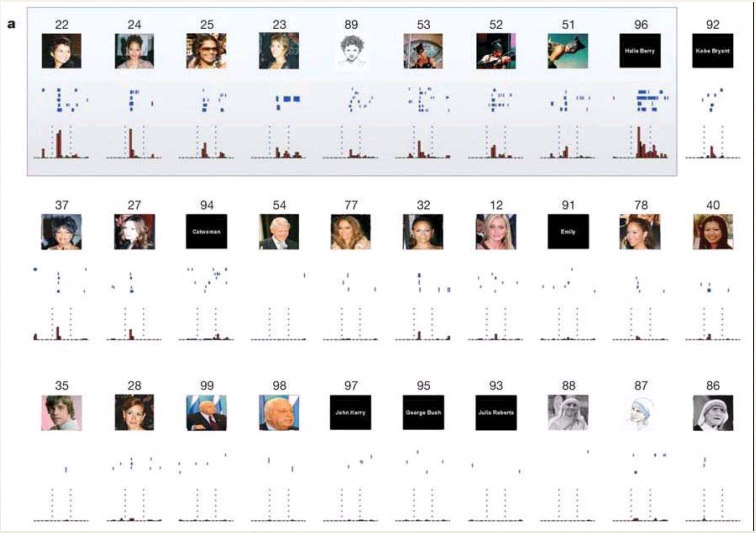

Slide adapted from Dayan

Quiroga et al (2005)

Bayesian tuning curve analysis

- We want to understand p (tuning curve | spikes)
- Markov-Chain-Monte-Carlo
- Cronin et al 2010

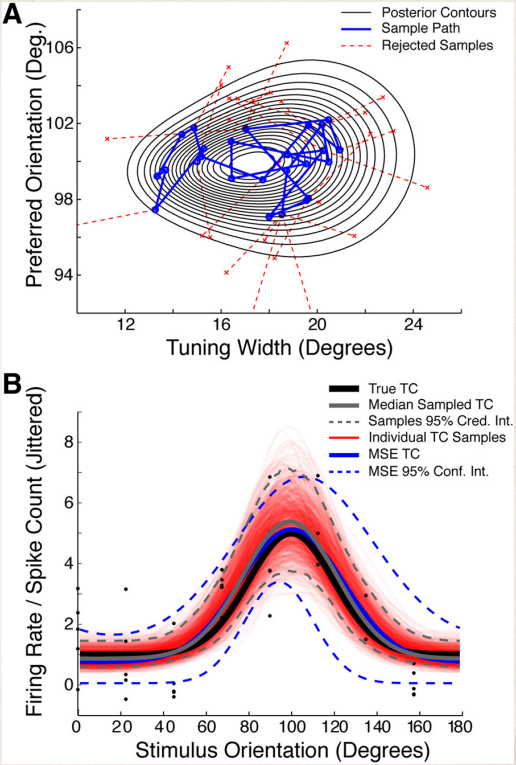

Part 2: How neurons may represent uncertainty

Many theories

- Brain has no reason to use code that Konrad can well understand

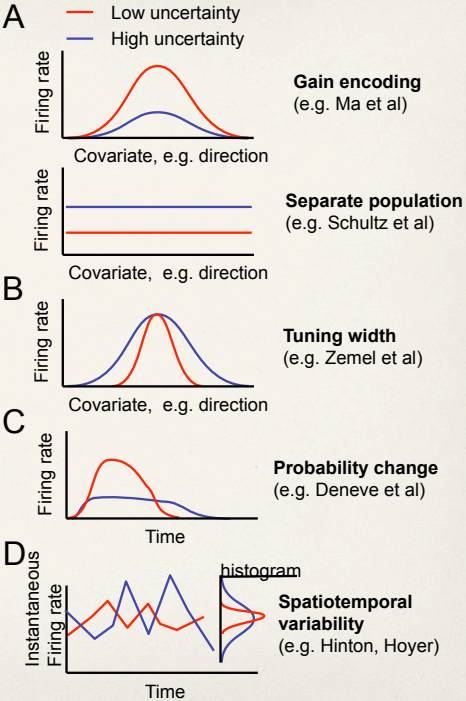

Distributed representation

- Distribution across neurons represents uncertainty
- Different variables represented by different populations
- + fast
- - requires many neurons

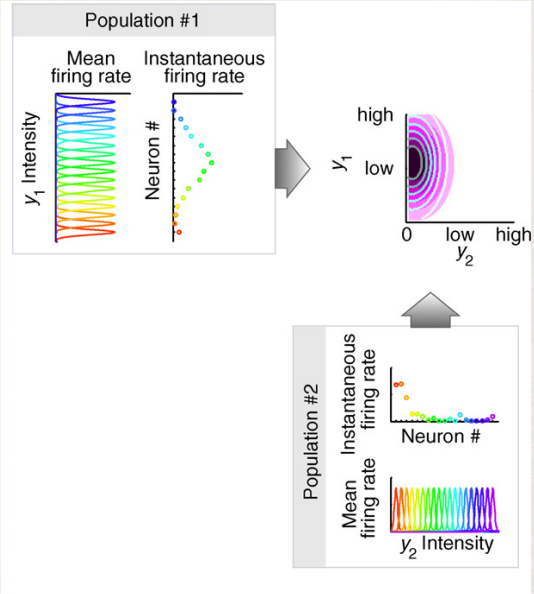

Fiser et al 2010

Representation by samples

- Distribution across neurons represents uncertainty
- Different variables represented by different populations
- + few neurons necessary
- - slow

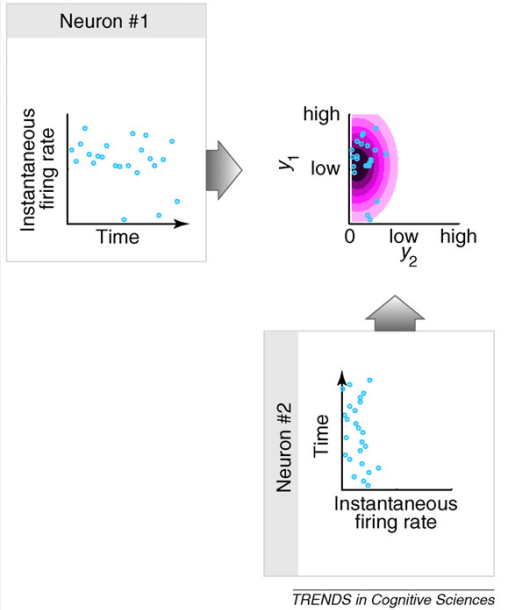

Fiser et al. TICS 2010

Part 3: Reverse engineering the way neurons interact

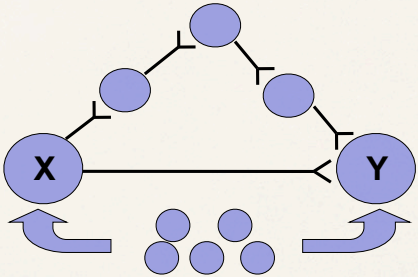

Why model interactions?

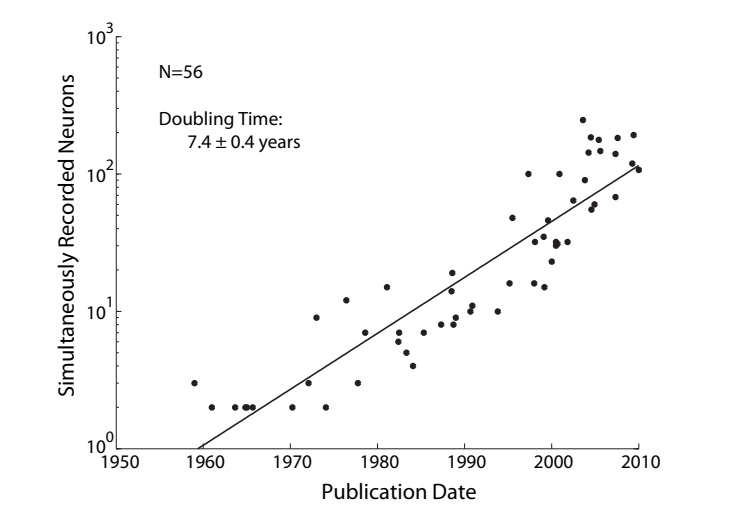

A generative model of spikes

$$\lambda_i(t) = exp\left(\sum_{i=1}^{N}\sum_{\Delta t=1}^{T} W_{i,\Delta t} s_i(t-\Delta t)\right)$$

- Biological interpretation
- Expressive power
- Real reason why we use this
- MAP estimation of weights

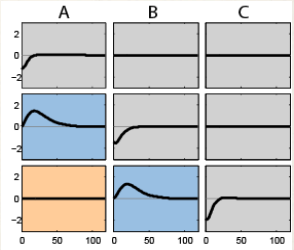

Explaining away

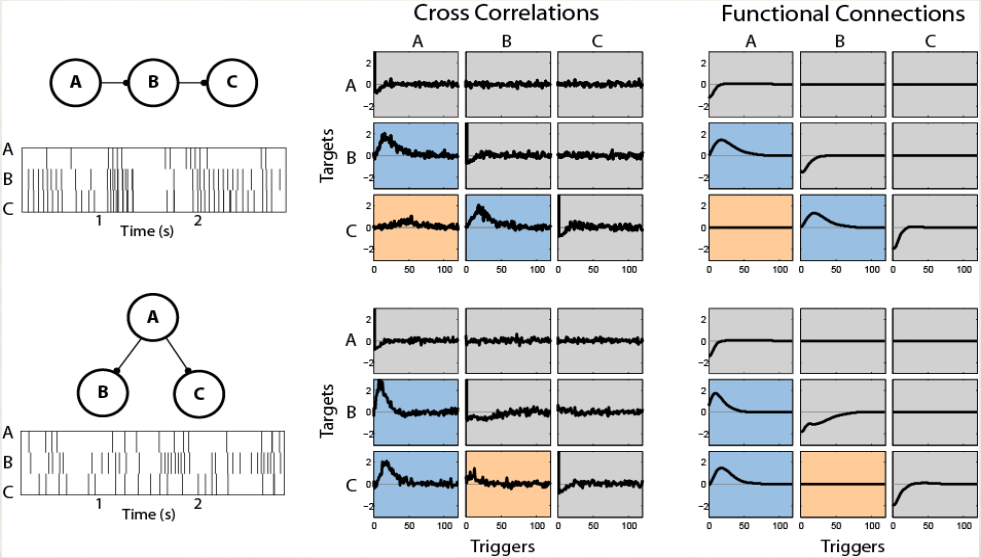

Results from real neurons

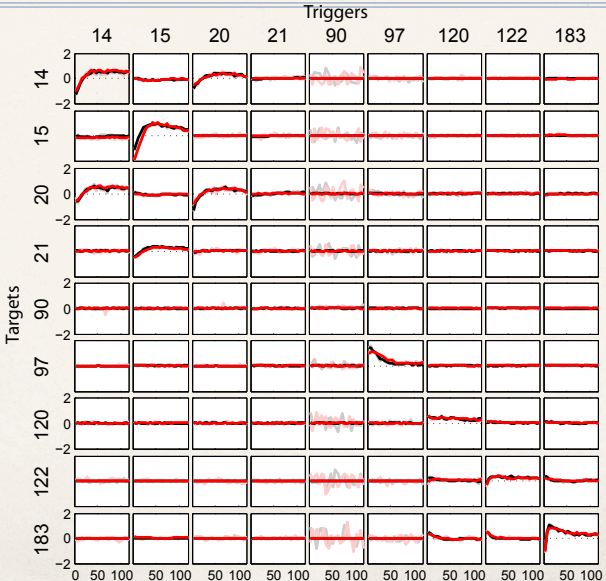

A more general model

$$\lambda_i(t) = exp\left(\sum_{i=1}^{N}\sum_{\Delta t=1}^{T} W_{i,\Delta t} s_i(t-\Delta t) + \sum_{k=1}^{K} V_{i,\Delta t} v_k(t-\Delta t)\right)$$

- Neurons are affected by other neurons
- Also affected via tuning curves by outside world

Graphical version

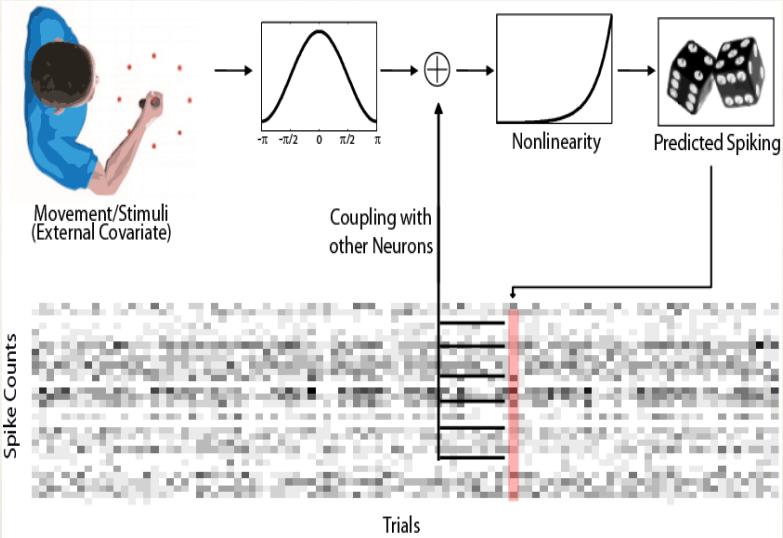

Information per Neuron

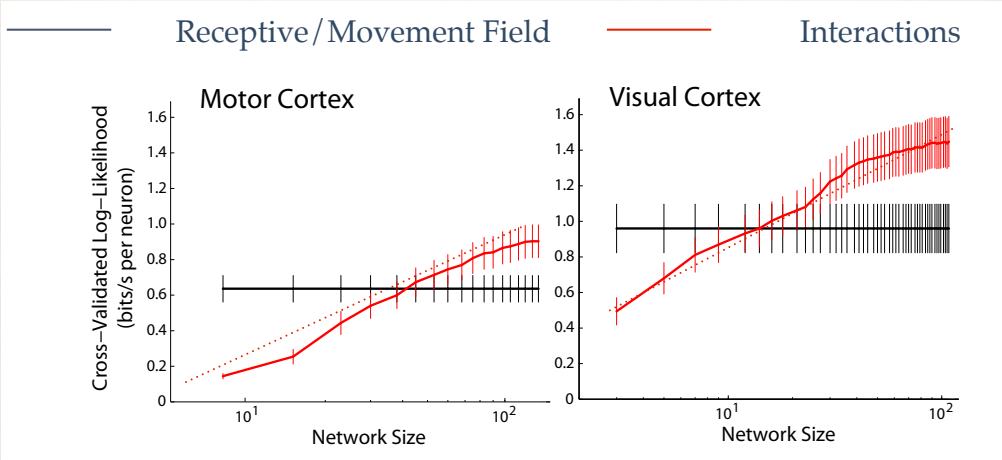

Hatsopoulos

O(n log n)

Kohn

Spike Timing Dependent Plasticity

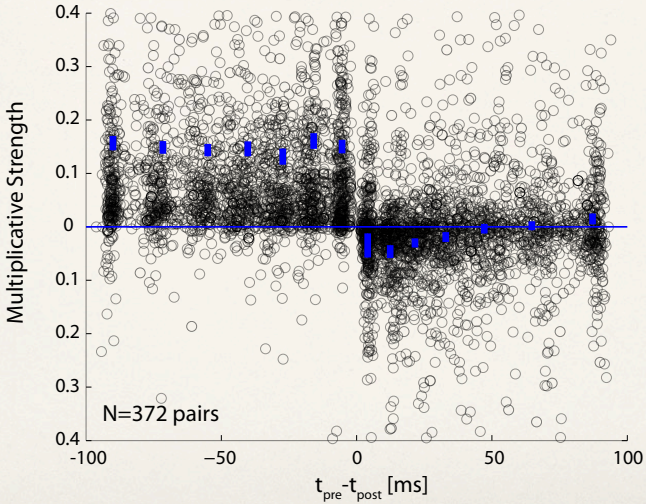

Tuning curves are explained away

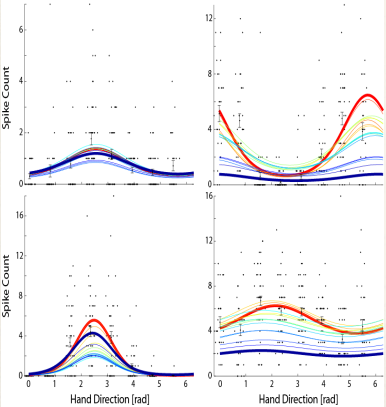

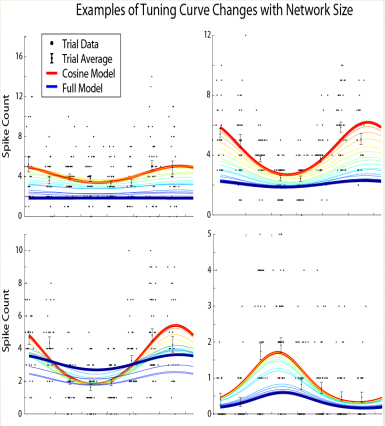

Countless future machine learning problems

- Find structure in generative model
- Meaningful priors
- Link to cognitive phenomena
- Reverse engineer learning rules

Acknowledgements

- Ian Stevenson
- Nicho Hatsopoulos
- Adam Kohn
- Lee Miller, Jim Rebesco
- Sara Solla