BSCS 2019 - Neural Computation

I - Introduction

Mihály Bányai
banyai.mihaly@wigner.mta.hu
http://golab.wigner.mta.hu/people/mihaly-banyai/

Course details

- Schedule
  - 9:00 - 10:45
  - 11:00 - 12:45
- Exam
  - essay, before Monday noon
- Resources
  - http://www.rmki.kfki.hu/~banmi/bscs/

Course outline

- **Monday**
  - Introduction
    - Models in neuroscience
    - Functions of the brain
  - Knowledge representation
    - How to formalise knowledge
    - How to handle uncertainty
- **Tuesday**
  - Probabilistic models
    - Generative models of observations
    - Bayesian inference
  - Models of cognition
    - Mental models of the world
    - Prediction of behaviour with probabilistic models
- **Wednesday**
  - Neural code
    - Tying algorithms to biology
    - Sampling hypothesis
  - Prediction of neural activity
    - Probabilistic models of vision
    - Deep learning models
    - *Topic assignments*
- **Thursday**
  - Selected topics in neural computation
    - Using deep learning to predict neural responses
    - Decision making and strategy learning
  - *Discussion of essay topics*
- **Friday: essay writing**

How does the brain work

according to wired.com

CHELSEA LEU SCIENCE 12.12.16 3:00 PM

WATCH A RESTING BR UP WITH ACTIVITY

VIBRANT N REVEAL W YOU

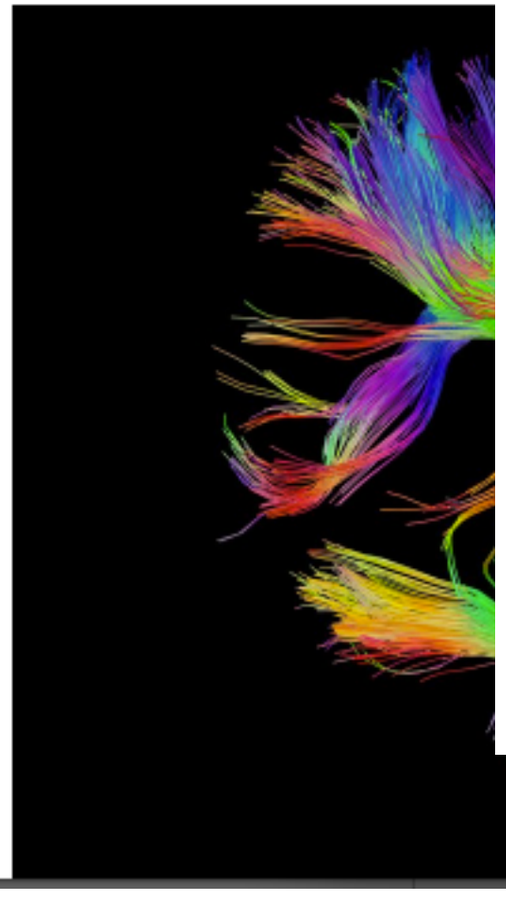

YING MA AND

CHELSEA LEU SCIENCE 07.13.16 1:00 PM

NEUROSCIENTISTS STILL DON'T KNOW WHY MUSIC SOUNDS GOOD

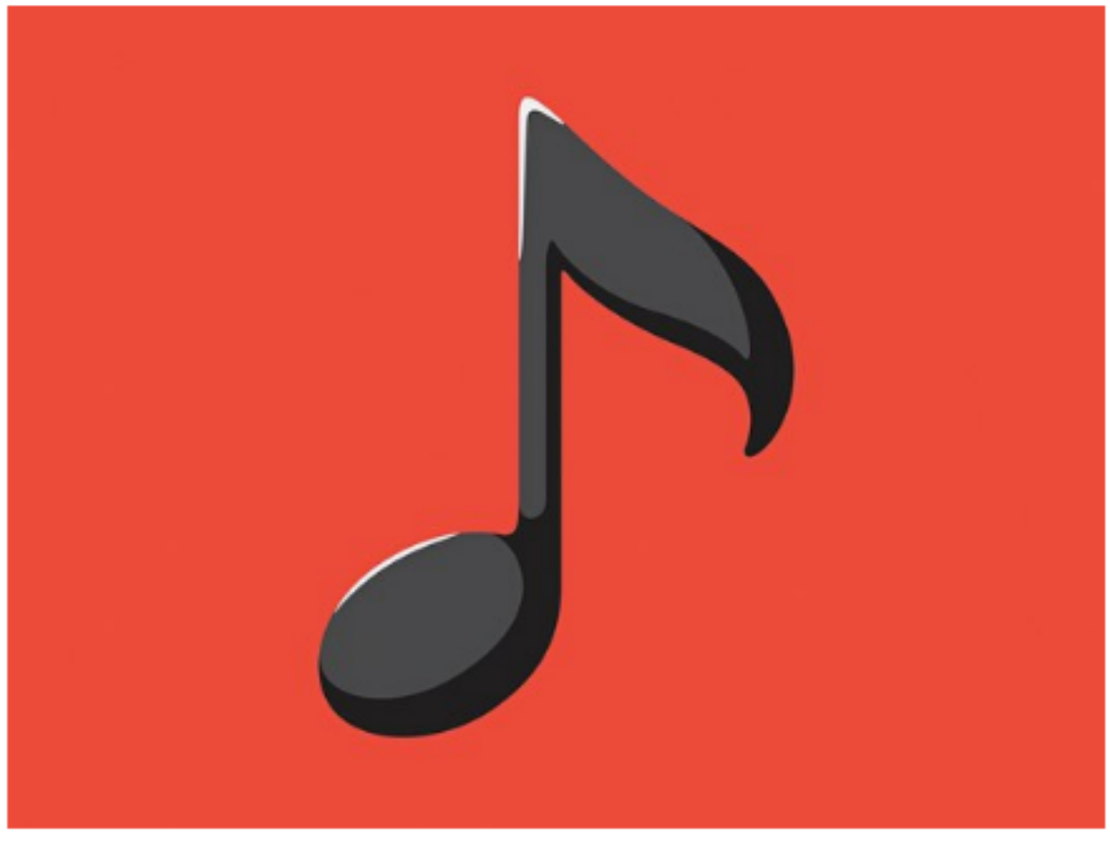

GETTY IMAGES

BRAIN NDARIES OF

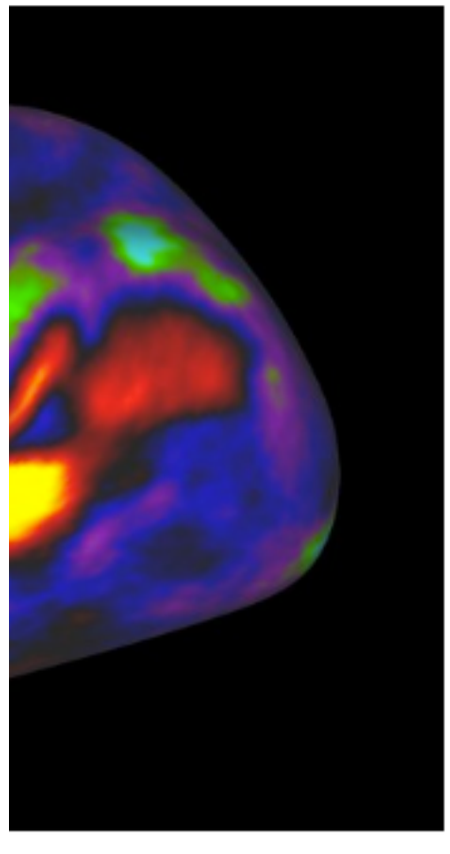

activation (red, yellow) and deactivation (blue, green) in
tories in the MRI scanner. MATTHEW F.

How does the brain work?

- You can find many answers by a Google search and there are multiple videos on YouTube that tell you the answer to this question

- But what does this question even mean? What kind of answer would be satisfactory?

- Do we have an answer to this question?

- Mathematical modelling of brain functionality
- Functions of the brain
- Brief history of computational intelligence
- Localising the structures that implement computation

- Mathematical modelling of brain functionality
- Functions of the brain
- Brief history of computational intelligence
- Localising the structures that implement computation

Taking a natural scientific approach towards the brain

- This means building models (theories) and using them to make predictions about observations

- Computational neuroscience
    - makes predictions about biophysical quantities, coming from **physiological** measurements

- Computational cognitive science
    - makes predictions about **behavioural** experiments

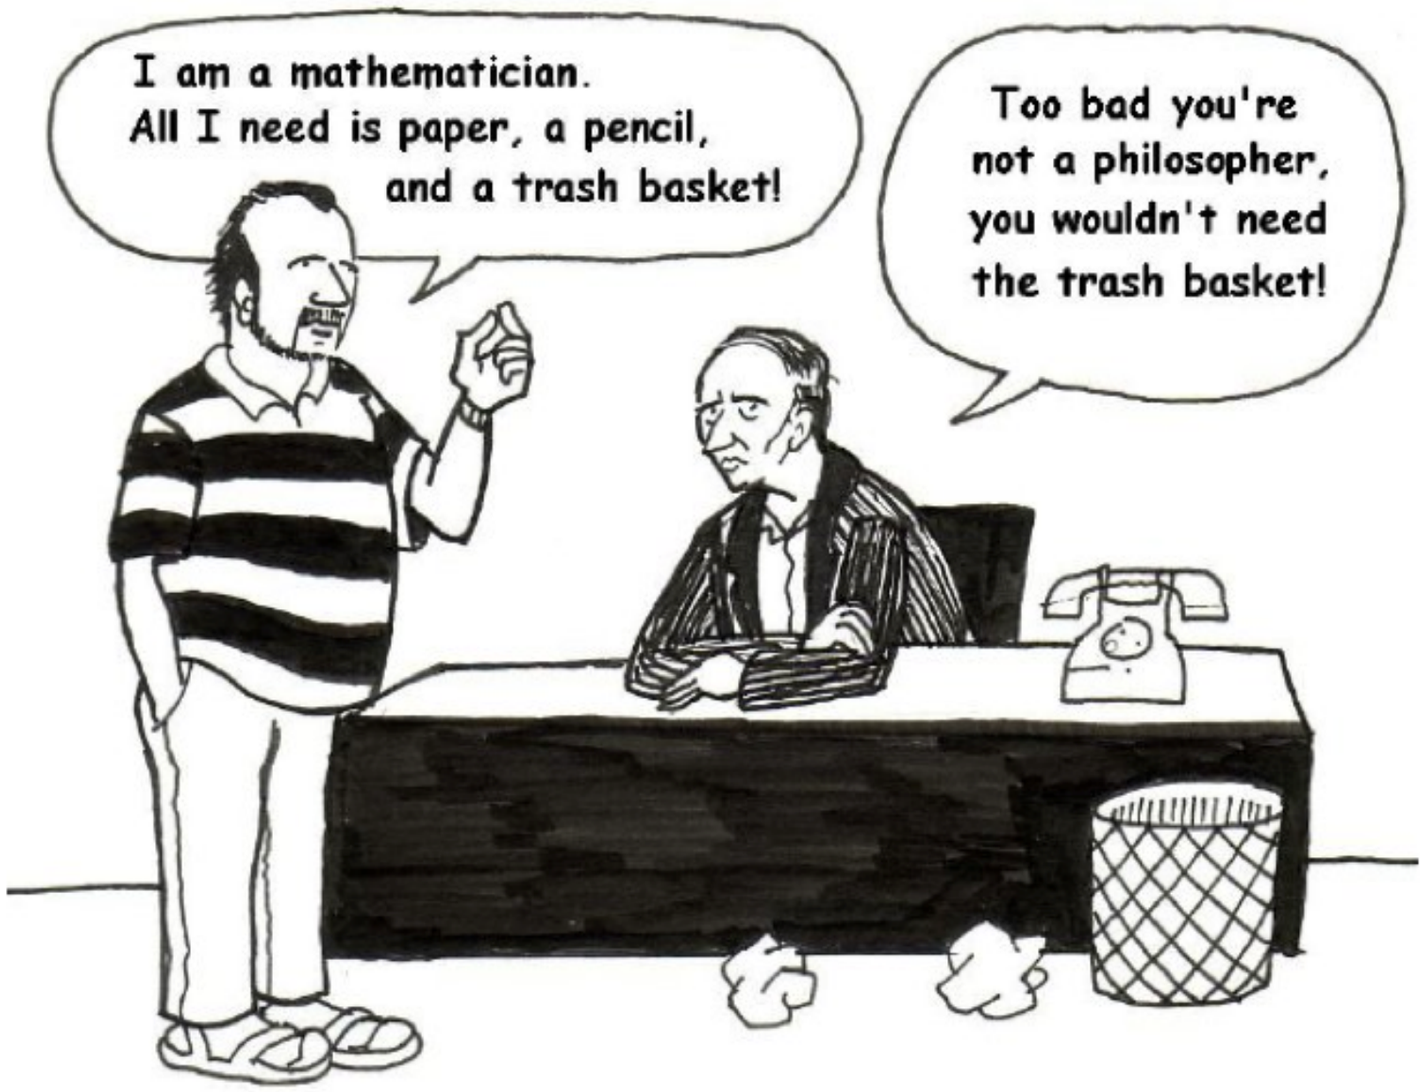

– Eugene M. Izhikevich: *Dynamical Systems in Neuroscience: The Geometry of Excitability and Bursting*

"Half of what we are going to teach you is wrong, and half of it is right. Our problem is that we don't know which half is which."

Charles Sidney Burwell

Recommended reading
http://www.smbc-comics.com/index.php?id=3905

What quantities do we want to predict?

- That is, what kind of properties of the brain do we deem essential?
- What would you be happier with?
  - A model that describes the structure of the brain well, but tells little about behaviour
  - A model that reproduces behaviour well, but does not resemble the structure of the brain too much

> “Prediction is very difficult, especially about the future.”
>
> *Niels Bohr*

Levels of abstraction

- **Computation** - specification of the brain function as an input-output mapping
- **Algorithm** - a step-by-step mathematical description of how to calculate the mapping
- **Implementation** - realisation of the algorithm by biological structures and their dynamical properties

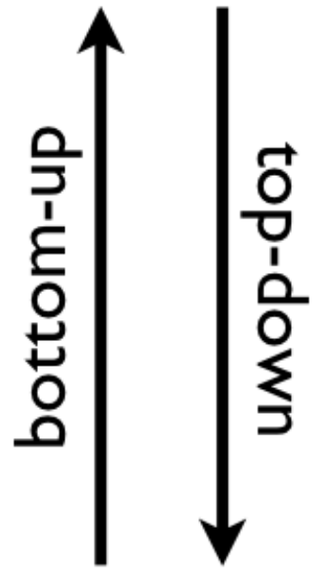

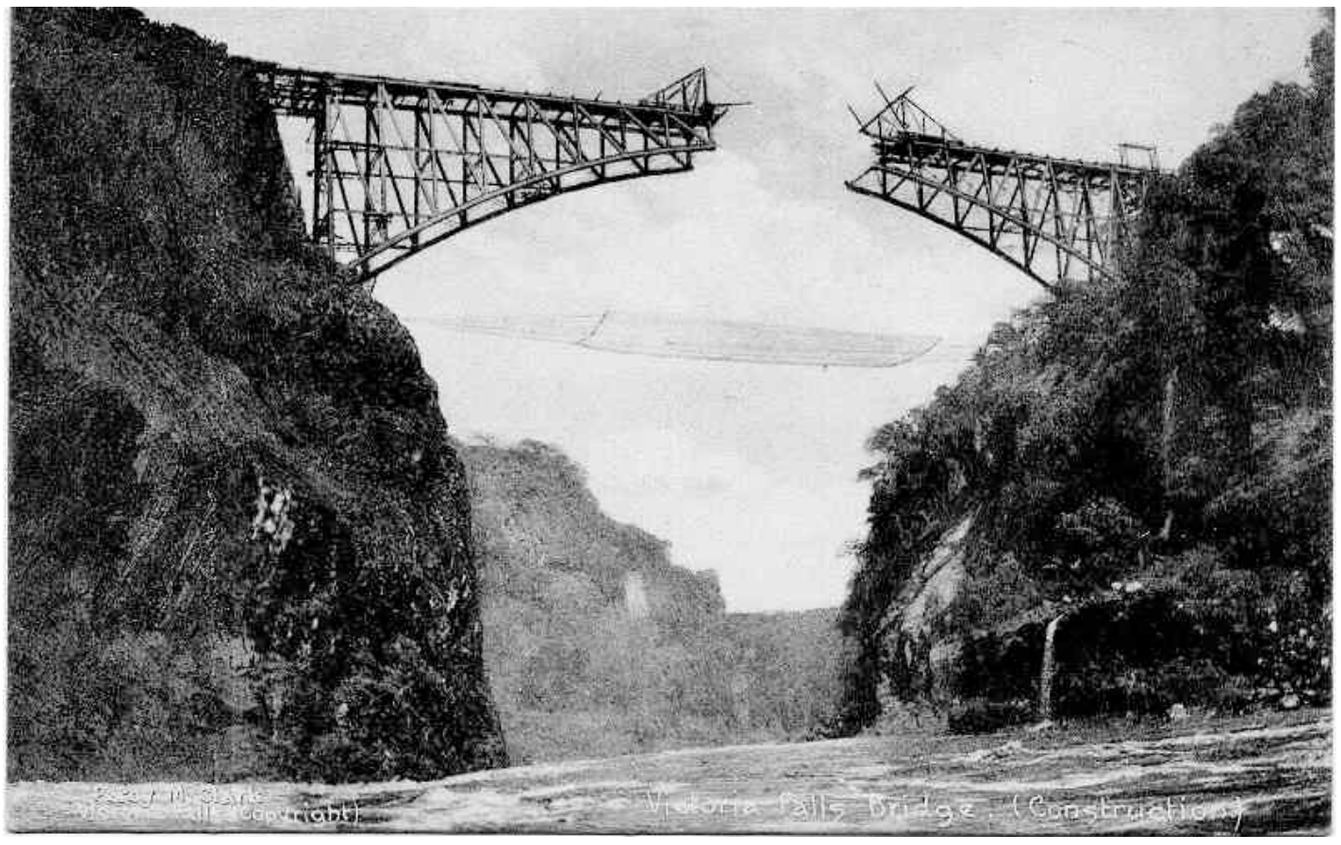

David Marr, 1976

Normative modelling

- We start by trying to reproduce the high-level properties of the object under study
  - And then move towards making the model more structurally similar to the object
  - also called top-down modelling
- The opposite approach is bottom-up
  - or descriptive modelling
  - trying to reproduce structure (and dynamical properties) as well as possible
  - function should be an emergent property of the structurally accurate model

"A wing would be a most mystifying structure if one did not know that birds flew."

Horace Barlow

Pointer
http://biorxiv.org/content/early/2016/05/26/055624

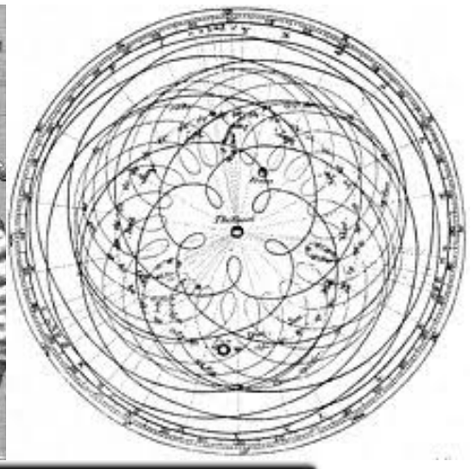

predictive accuracy

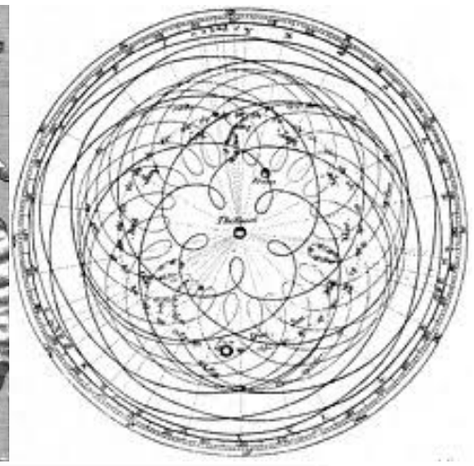

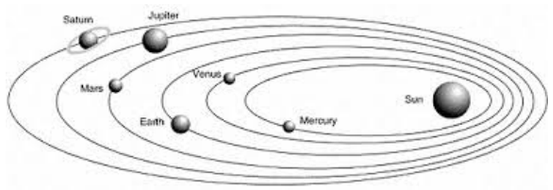

predictive accuracy

Looking for a unifying perspective

- The difference between alchemy and physics: specificity vs. generality
- A central aim in computational neuroscience is to find a theory that explains the different aspects of the brain and behaviour in similar terms, building on a small set of fundamental principles
- Examples of such unifying perspectives from other disciplines
  - Newtonian forces in physics
  - Evolutionary principle in biology
  - Computational complexity, Turing-completeness
  - String theory / holographic principle (?)

The role of theory

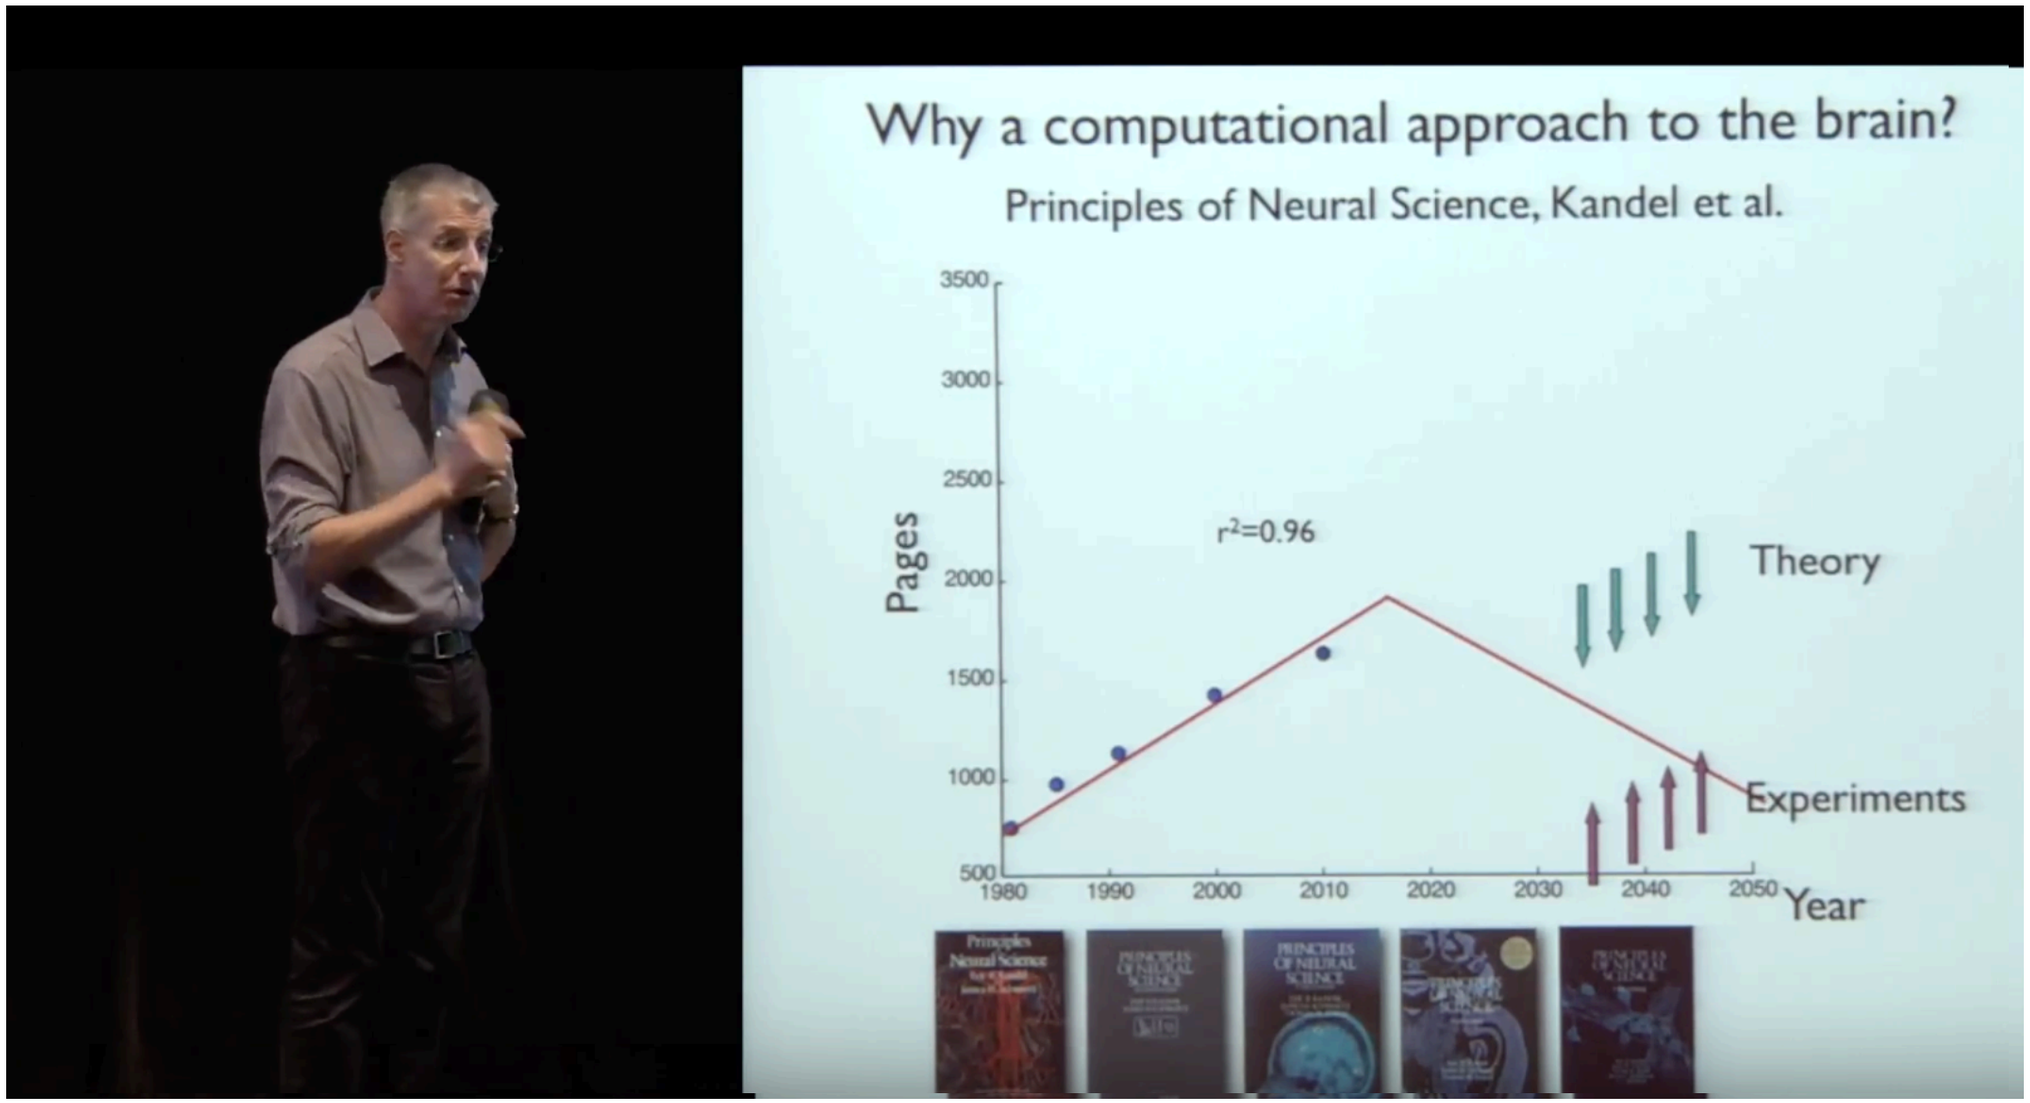

https://www.youtube.com/watch?v=wTYHF4LAKQI

Why should theory inform experimental design?

What if the brain…

…disentangles the input space for classification

Hierarchical models of object recognition in cortex

Maximilian Riesenhuber & Tomaso Poggio

Nature Neuroscience **2**, 1019–1025 (1999) | Dov

…implements probabilistic inference

Hierarchical Bayesian inference in the visual cortex

Tai Sing Lee and David Mumford

Journal of the Optical Society of America A

Vol. 20, Issue 7, pp. 1434-1448 (2003)

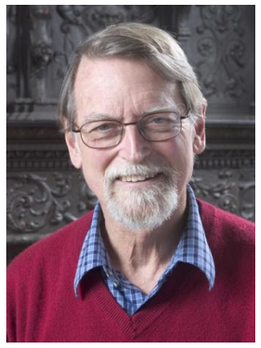

…maximises mutual information with future stimuli

Predictability, Complexity, and Learning

William Bialek, Ilya Nemenman and Naftali Tishby

Posted Online March 13, 2006
https://doi.org/10.1162/089976601753195969

Neural Computation
Volume 13 | Issue 11 | November 2001
p.2409-2463

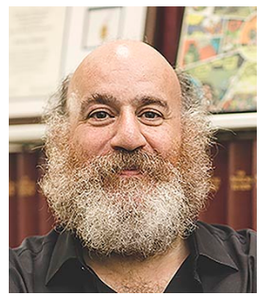

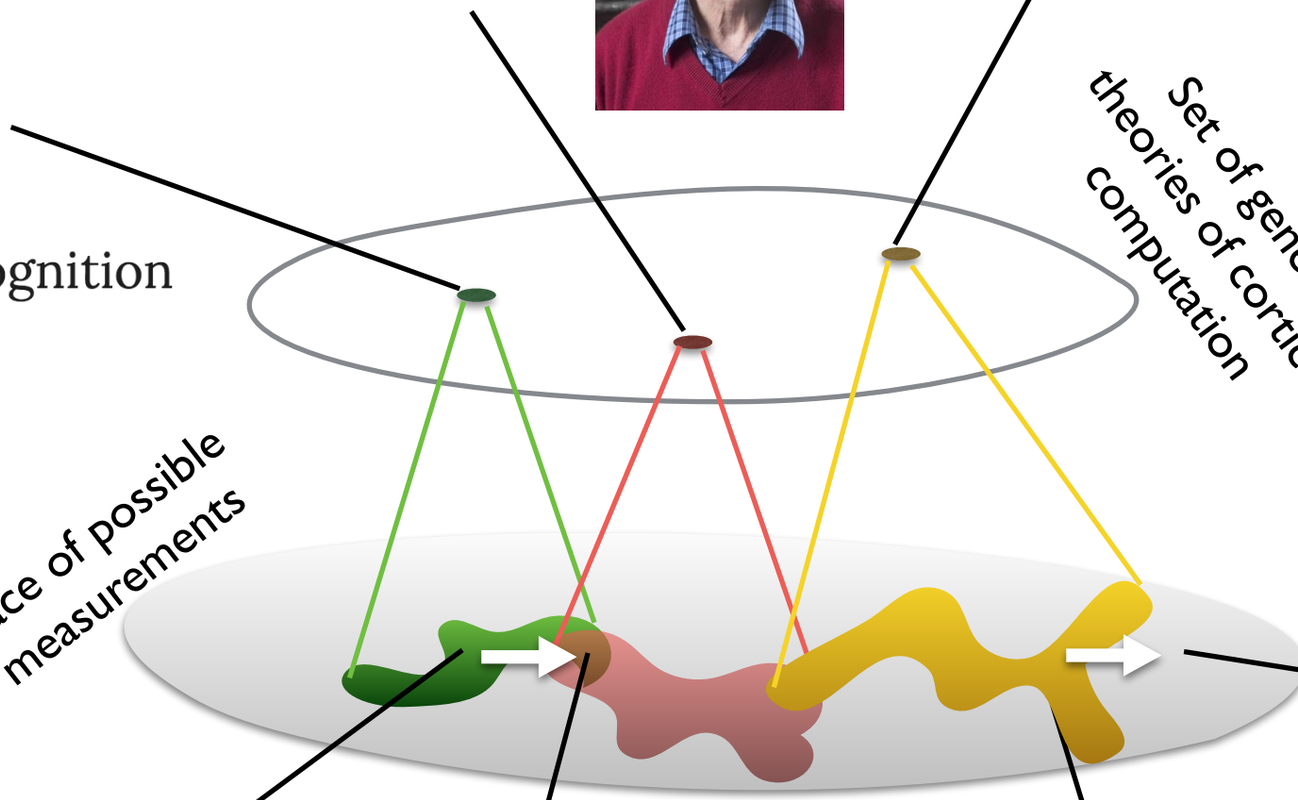

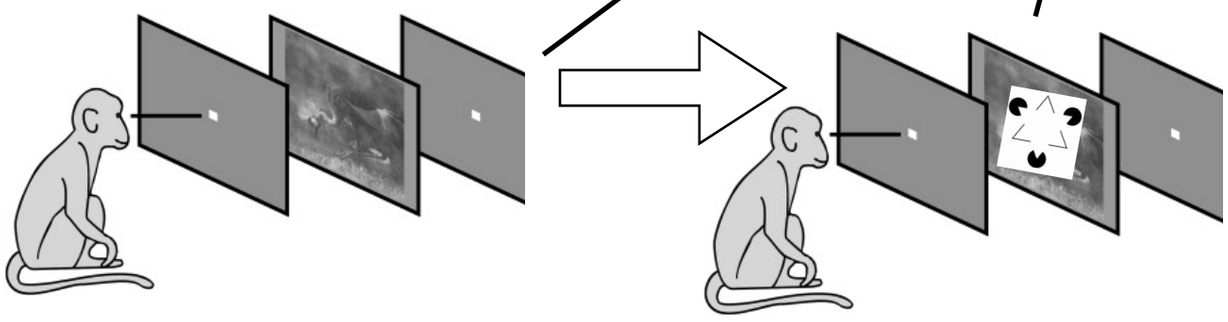

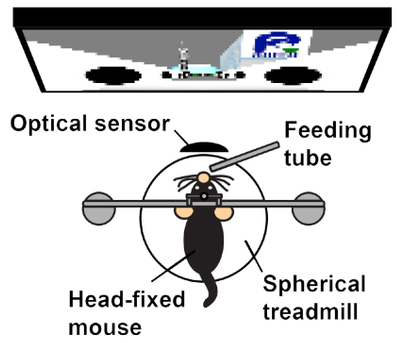

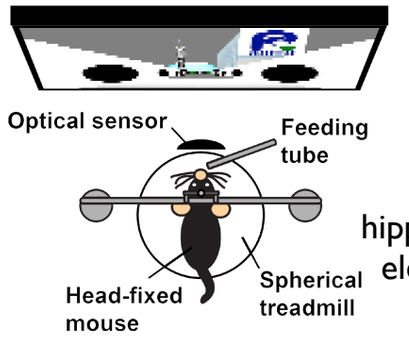

prefrontal electrode

hippocampal electrode

- Mathematical modelling of brain functionality
- Functions of the brain
- Brief history of computational intelligence
- Localising the structures that implement computation

What does the brain do?

- Moves the muscles

Knowledge

Past events

Learned regularities of the world

Inferred information from actual observation

Consciousness

?

Decision making

Motor control

Perception

What does the brain do?

- Moves the muscles
  - that is to make decisions
  - that's easy as long as you only want reflexes
  - otherwise you also have to remember stuff - create a **representation** of knowledge in the neural tissue
    - we have to use the knowledge to make sense observations - **inference**
    - we have to continually update the knowledge base with new information - **learning**
- The brain uses the knowledge to predict outcomes of decisions - so it builds models of the world (yes, we'll be modelling procedures that build models of other procedures)
  - we may build models to interpret data, and talk about models that the brain builds - never confuse the two (even if it's not always easy)

What is the task of the visual system?

- From the uninteresting sensory input (wavelength distribution of incoming light) acquire interesting information (objects)
- But what are the useful objects and how do we recognise them?

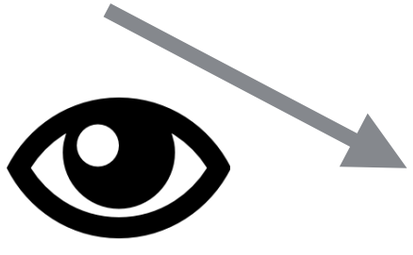

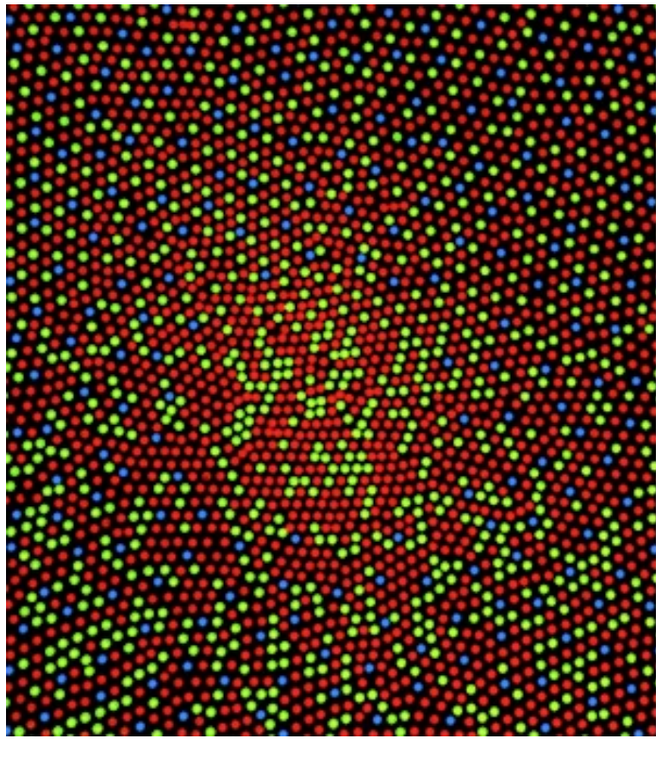

?

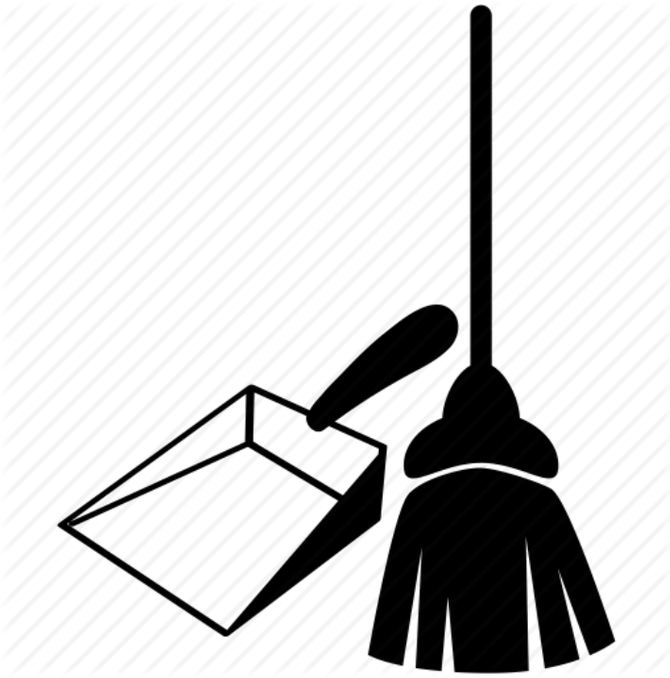

Ambiguity of observations

- A defining property of the environment is that it's ambiguous
  - this is the rule, not the exception
- Sensors also introduce uncertainty, but the main source of it is the lack of complete information
- Thus, perception is an inference problem, from the observations we have to reconstruct the content of the environment

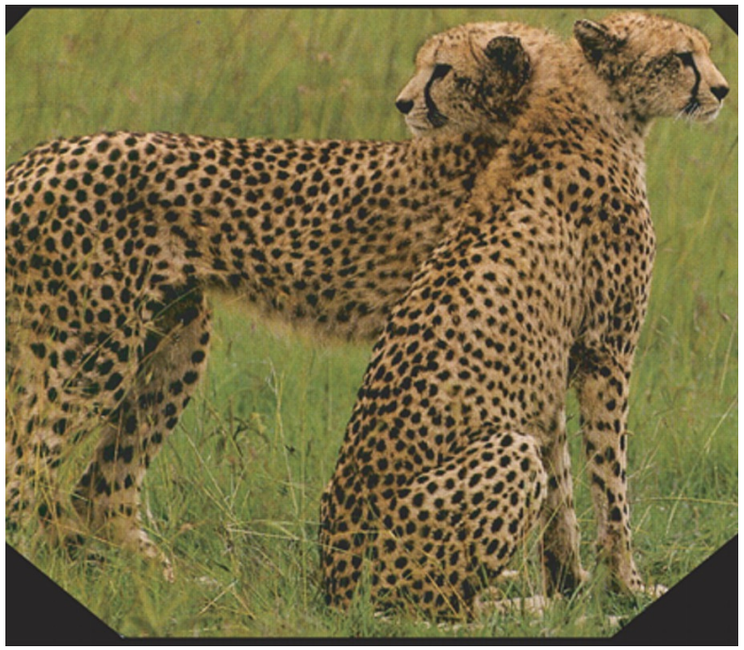

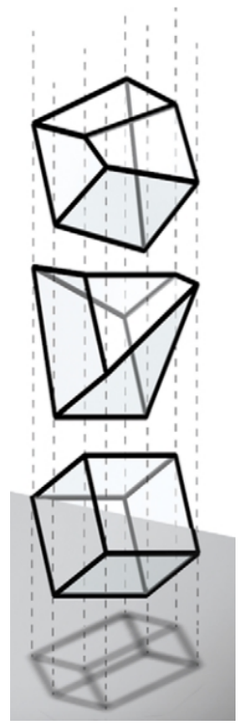

Integration of learned knowledge and observations

- Illusions can be explained by using our knowledge about regularities of the world to interpret observations

- Gestalt principles describe the rules of shape perception - continuity in space and time

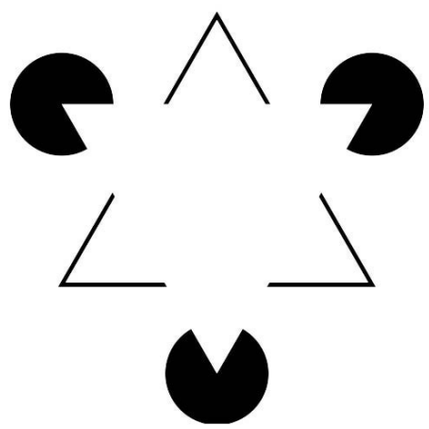

Where is the sun?

Jennifer Sun[1] and Pietro Perona[1,2]

[1] *California Institute of Technology 136-93, Pasadena, California 91125, USA*

[2] *Universita di Padova, Via Ognissanti 72, 35131 Padova, Italy*

Correspondence should be addressed to P.P. (perona@vision.caltech.edu)

nature *neuroscience* • volume 1 no 3 • july 1998

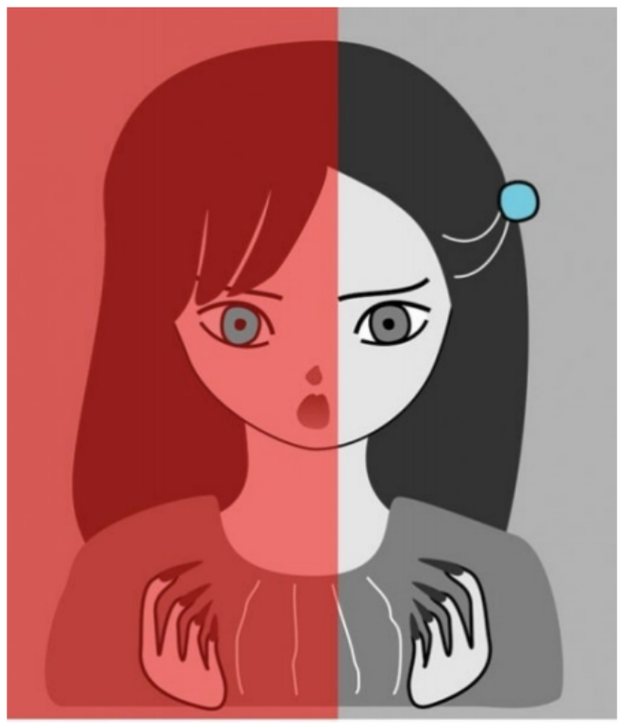

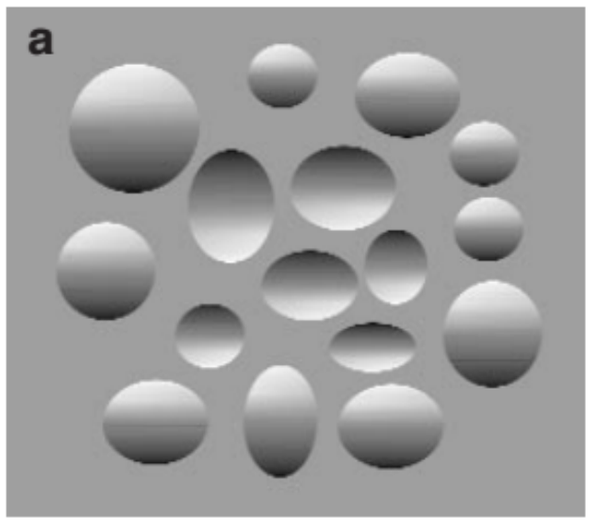

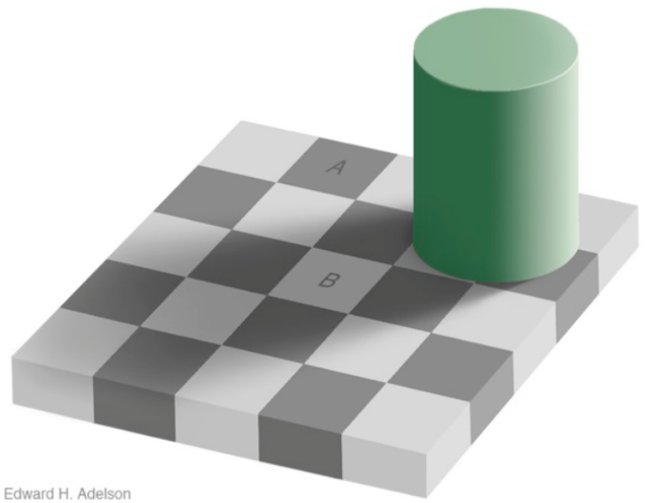

Edward H. Adelson

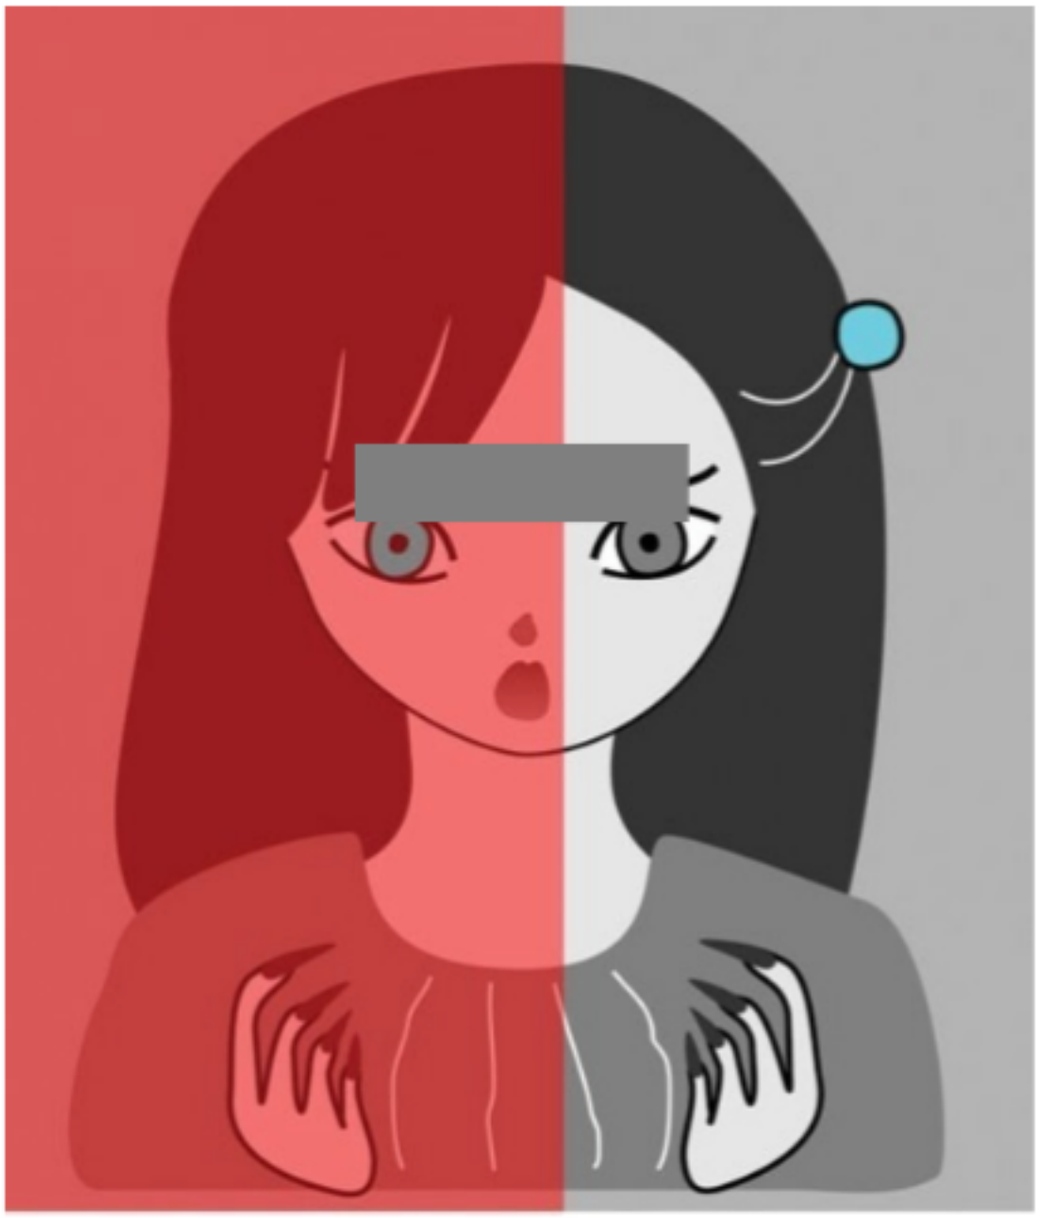

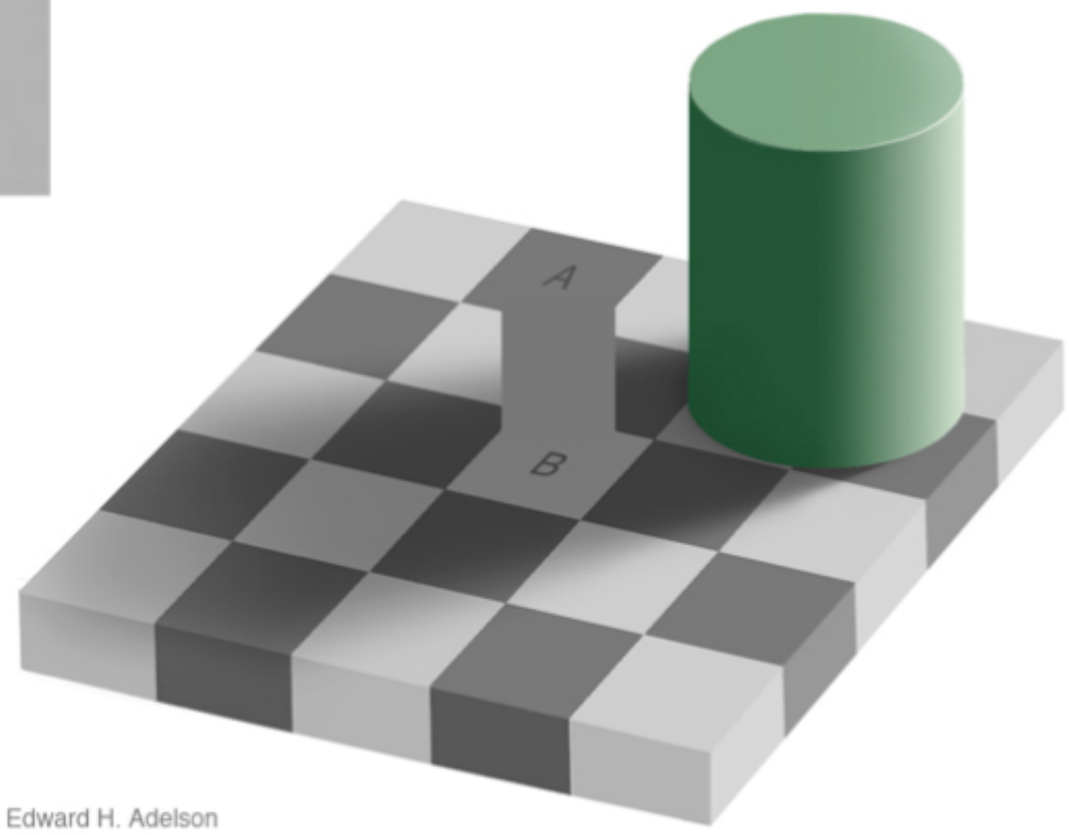

Edward H. Adelson

“Illusions” in artificial neural networks

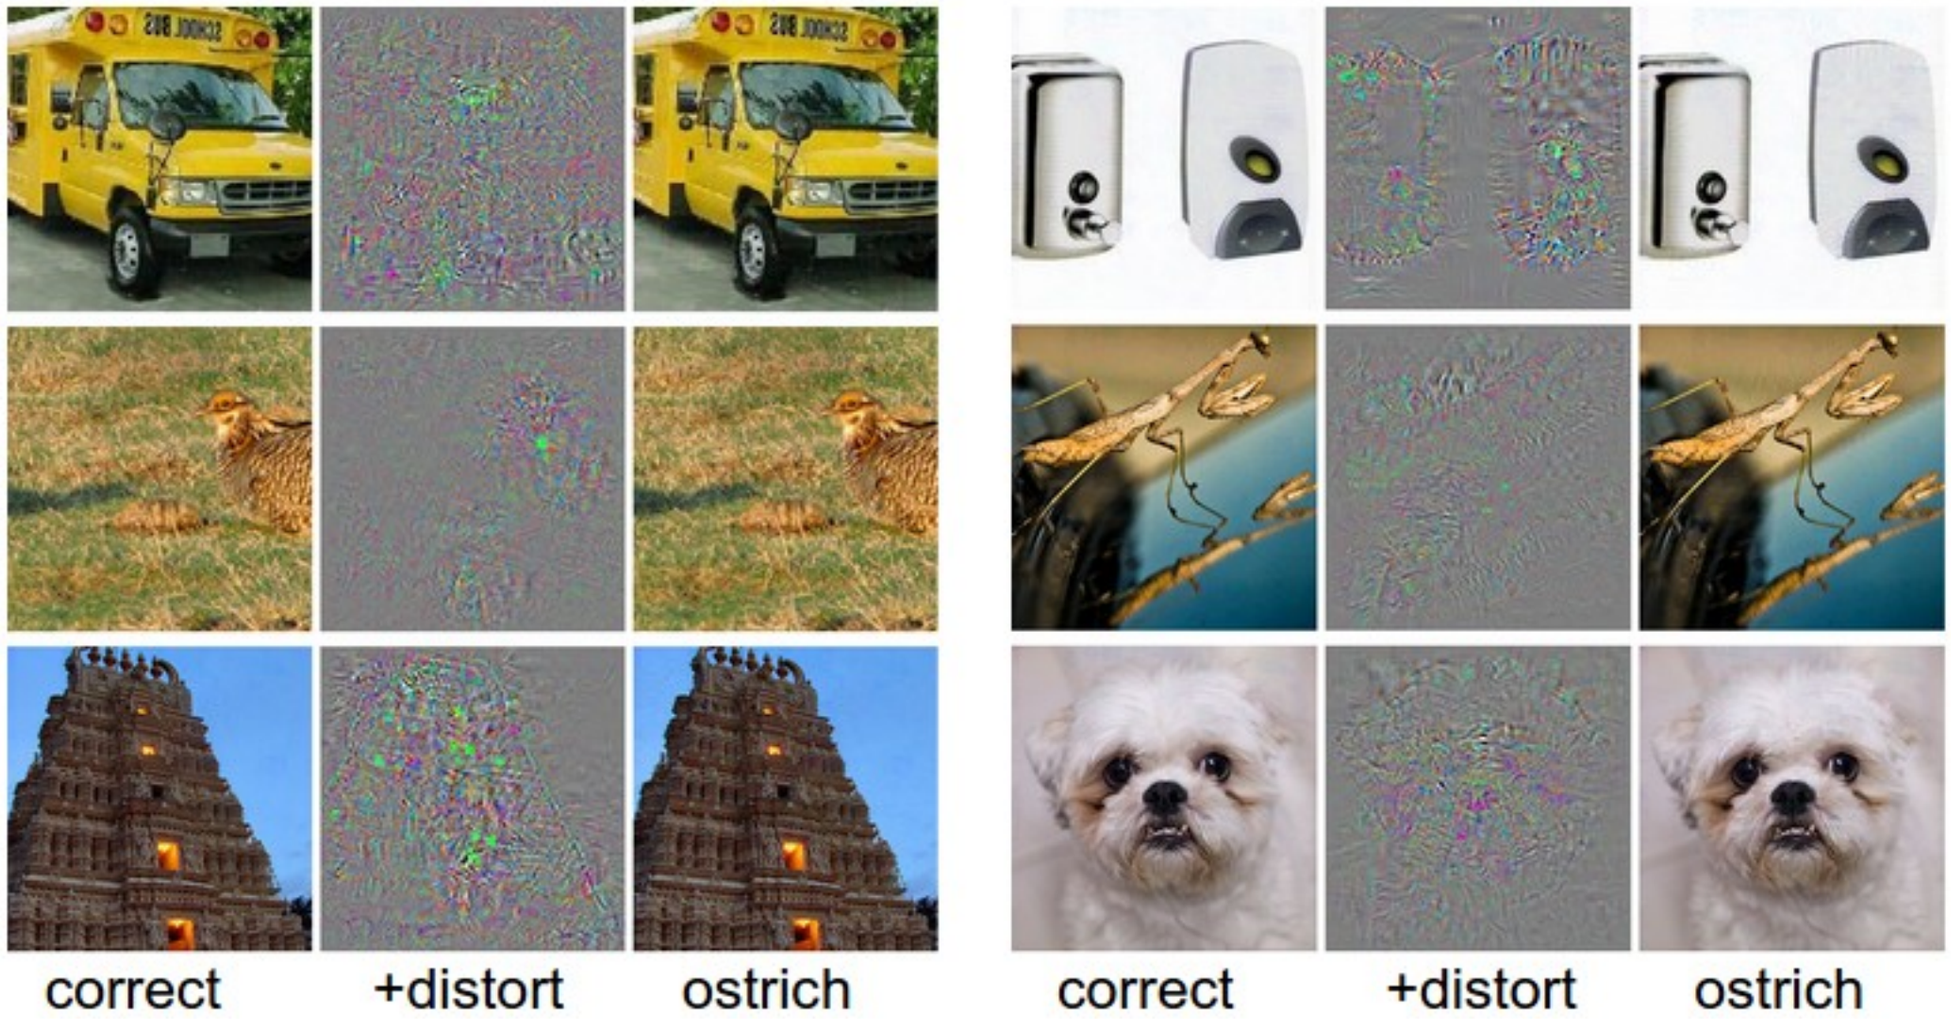

Learning the regularities of the world

- How many wheels does a car have?
  - you typically see 3, but you'd still guess 4
  - you learn this, as if you walk around it, you see all 4
    - but to realise this, you need to know that it is the same car all along -> objects are invariant in time (on short scales)
  - it sounds trivial, but newborns don't know it

Questions of representation

- What are the quantities that we (or an animal) need to extract from the environment to make decisions?
  - not necessarily the same question as what are the ones we *do* extract, but related (hopefully)
- What are the intermediate quantities between these and the observables that are useful to compute first?
  - If we know the answer to these two, we have the structure of a mental model - we mostly aim only for smaller parts of the whole
- What are the algorithms that compute these quantities?
- Do we have a task-independent model of the world, or a toolbox of models that we select from based on context?
- How do we generate stories - or how to use the mental model to test counterfactual hypotheses?

The neural code

- The way representation is tied to biophysical quantities
- Do we really need to know this?
  - If we had the complete description of a mental model and the algorithms operating on it, and we only wanted to predict behaviour - no
  - practically, figuring out the model only seems to be possible by discovering the biophysical constraints on it
  - of course, medicine needs the neural level descriptions too

The brain-computer analogy

- We can try to think about the brain as something similar to a computer
  - the hardware is the tissue, and the software is the behaviour
- Actually, nowadays there are a lot of software applications that try to do the job of a human brain
  - extracting object identities from images
  - building useful categorisations of objects, gradually building up abstract concepts
  - language processing and speech production
- The analogy is far from perfect, but it helps to define the level of abstraction on which we want to think

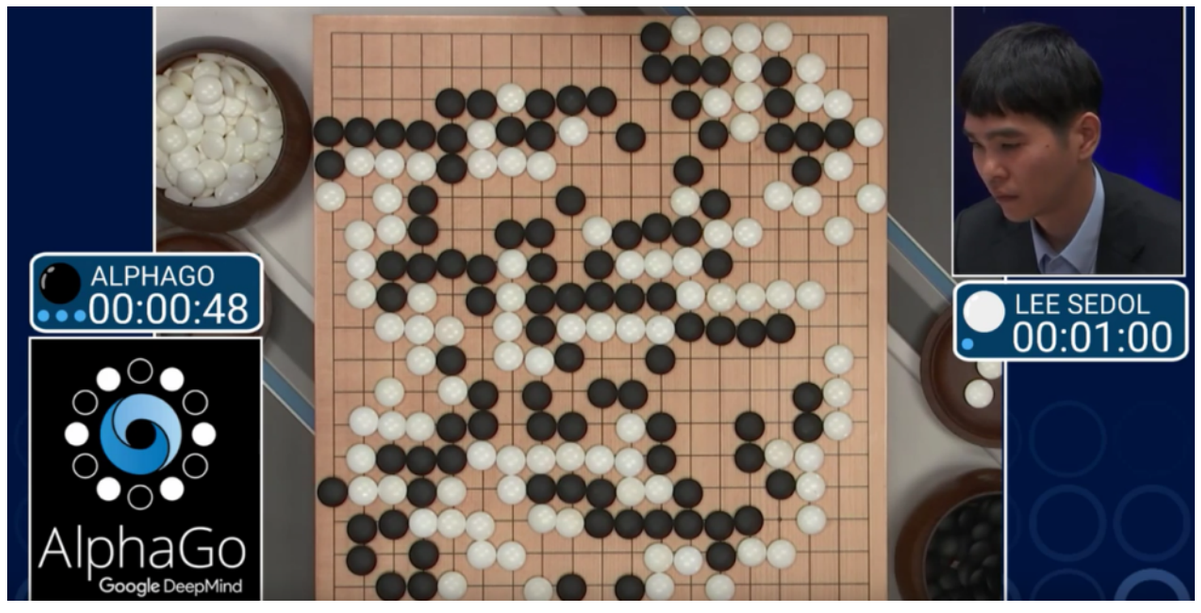

The brain-computer analogy

In computational neuroscience, we try to address the problem of figuring out how a web browser works by measuring the electric potentials of a hundred transistors of a CPU

PLOS | COMPUTATIONAL BIOLOGY

Browse | Publish

OPEN ACCESS PEER-REVIEWED

RESEARCH ARTICLE

Could a Neuroscientist Understand a Microprocessor?

Eric Jonas, Konrad Paul Kording

Published: January 12, 2017 • https://doi.org/10.1371/journal.pcbi.1005268

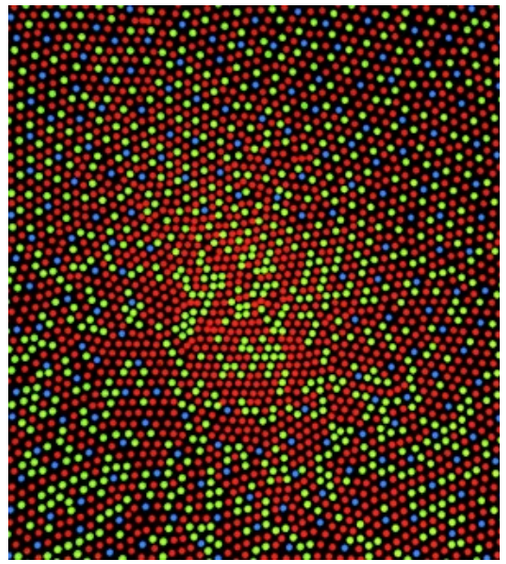

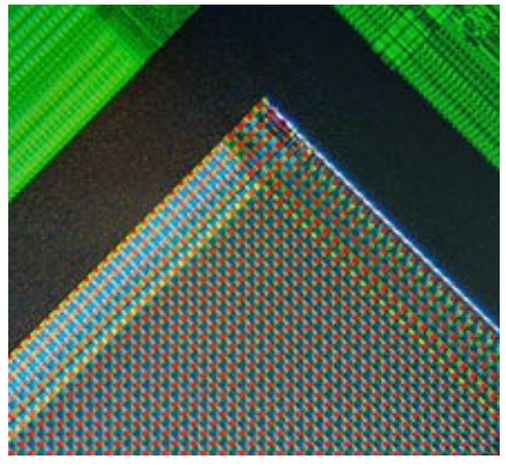

- Mathematical modelling of brain functionality
- Functions of the brain
- Brief history of computational intelligence
- Localising the structures that implement computation

History of computational neuroscience - roots

- Perception as inference about latent causes
  - Ibn al-Haytham, ~ 1020, Basra

  - Hermann von Helmholtz, 1867, Heidelberg

- In the 20th century, electrophysiological measurements made it possible to formulate cognitive theories on a neural level

The beginning

results in logic

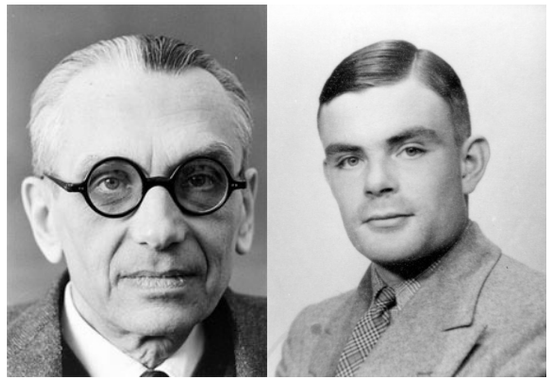

digital computer

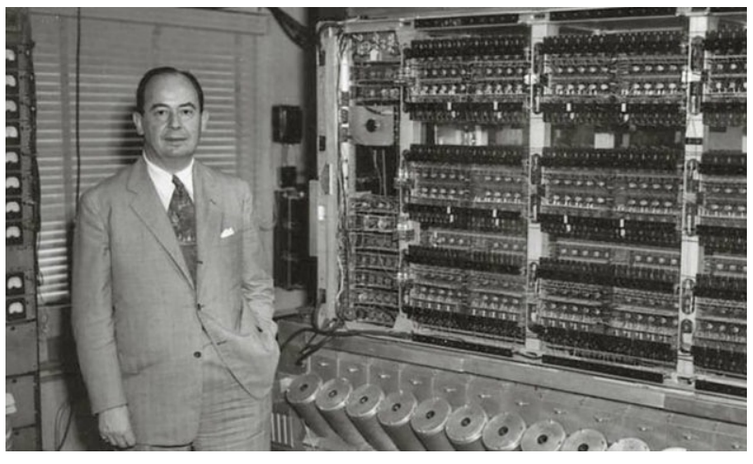

cybernetics

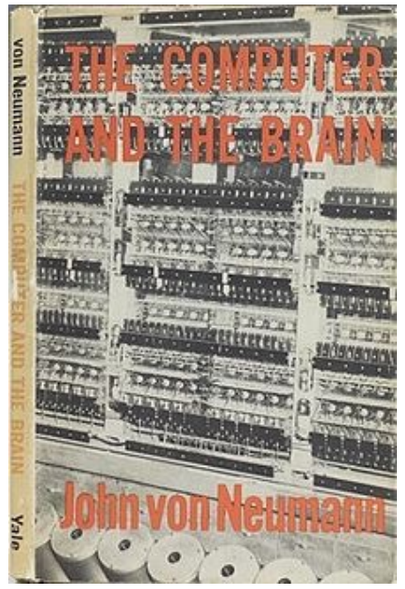

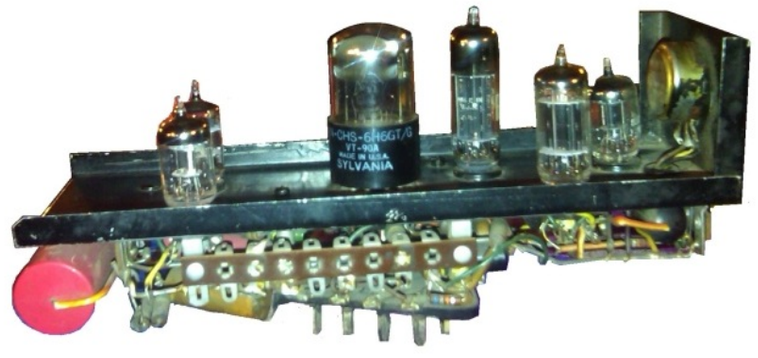

artificial neuron

neural recordings

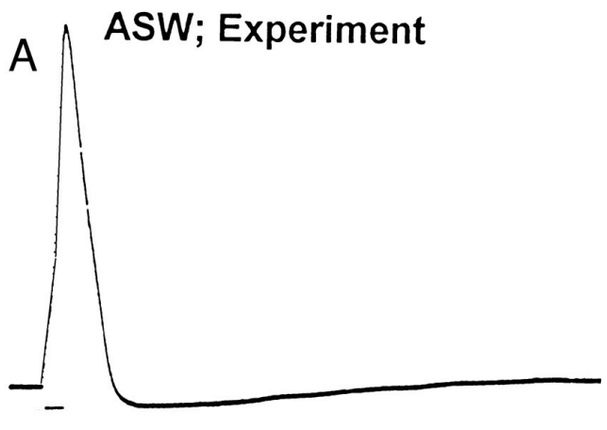

Recommended reading
http://nautil.us/issue/21/information/the-man-who-tried-to-redeem-the-world-with-logic

Coevolution with AI

- Most CN problems have a parallel AI problem
  - AI: how to do it best?
  - CN: How is it actually done in the brain?
- Artificial intelligence, **machine learning**, data mining/science, statistical modelling, adaptive control, robotics, bioinformatics
- Ideas spread both ways
  - early AI models inspired by neural computation
  - advanced high-level CN models inspired by machine learning solutions

Trends in computational neuroscience

- Biophysically detailed modelling
- Spiking networks
  - Balanced networks
- Connectionist networks
  - Deep learning
- Systems neuroscience
- Cognitive neuroscience

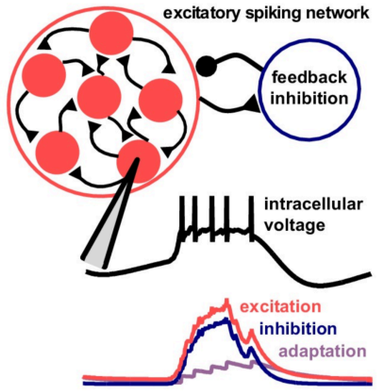

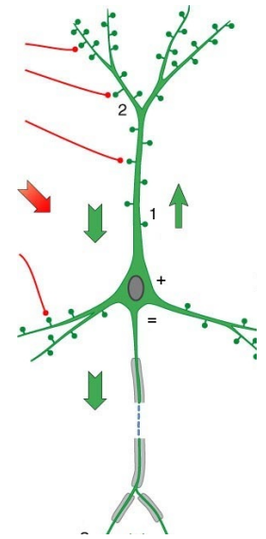

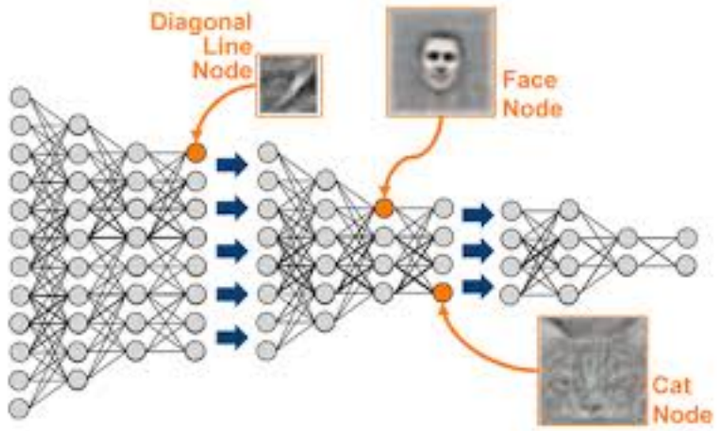

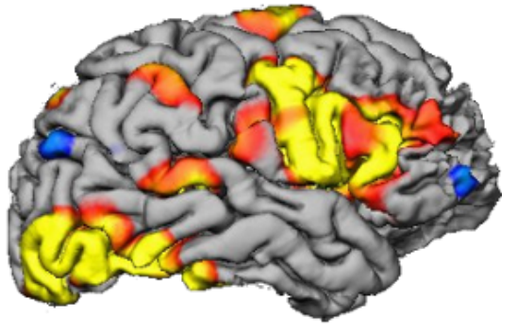

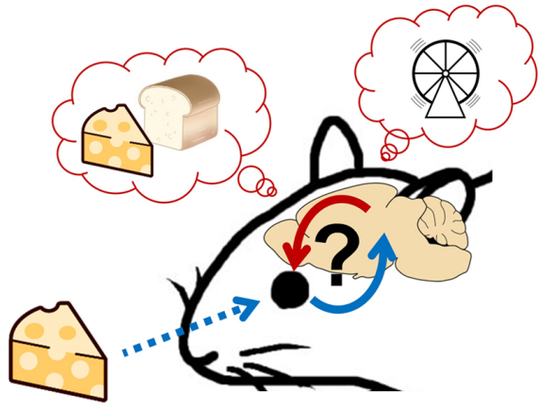

Big neuroscience

- Human Brain Project - EU
  - build a complete model of the cortex with biophysical detail
- Human Connectome Project - USA
  - map connections on multiple scales
- BRAIN Initiative - USA
  - improve measurability
  - develop new theories

Recommended reading
http://www.nature.com/news/a-better-way-to-crack-the-brain-1.20935

- Mathematical modelling of brain functionality
- Functions of the brain
- Brief history of computational intelligence
- Localising the structures that implement computation

Localisation of functions

- a prerequisite of thinking about how computations are realised is to find out by which part of the brain are they realised
- Measure their averaged activity or how much power they consume when a certain task is performed
- The search can be conducted on multiple levels
    - regions of the cortex that have specific functions, e.g. visual cortex, etc.
    - subpopulations of neurons dealing with a specific set of phenomena, e.g. motor cortex subpopulations for certain muscles, face selective visual area
    - neurons responding to a certain feature of the stimulus
- after this we can try to figure out the actual computation and algorithm that they implement
    - this will require much more than just looking at average activity or power consumption, but we have to know where to look

Large scale localisation

- Crude measure: lesions
  - if the back of your head is damaged, you'll go blind

- A nicer way: fMRI, EEG

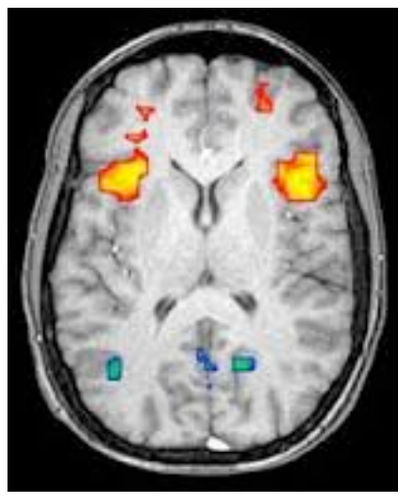

Recommended reading
http://prefrontal.org/files/posters/Bennett-Salmon-2009.pdf

The receptive field of a cell

- Looking for the variables of the mental model
- How the average activity of a neuron varies when certain properties of the stimulus change
    - we are looking for the quantities that actually make a difference for the cell
- There are potentially infinitely many quantities that could define the RF of a specific cell or subpopulation
- Functionality provides intuition

Looking for receptive fields - visual cortex

- lesions say dorsal parts of cortex are involved in vision
- what's the goal of vision? - to realise what's in front of me, e.g. what kind of objects
- objects are best defined by their contours (?)
- contours can be defined as sets of oriented line segments
- what if we look for cells that respond to those

Primary visual cortex - V1

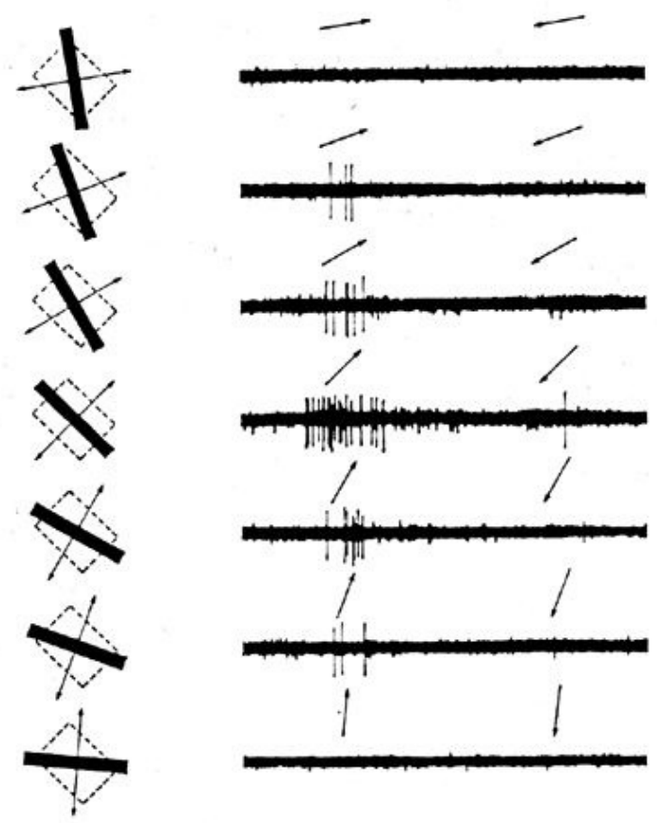

Looking for receptive fields - auditory cortex

- lesions: certain parts of the temporal cortex are involved in hearing

- the goal of hearing - discrimination of sources

- a good way to do this is to decompose input to frequency bands

- look for cells that are responsive to frequencies

Primary auditory cortex

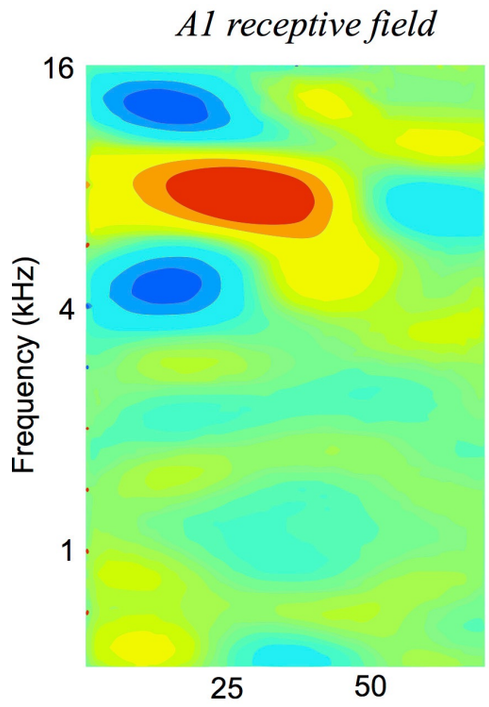

Looking for receptive fields - hippocampus

- lesions: hippocampus is involved in episodic memory
- what simple (well testable) task do you need that for? - navigation (even better if you're a rat, because than that's about all you care about)
- for navigation, you need to have a map
- what if cells are responsive to specific places?

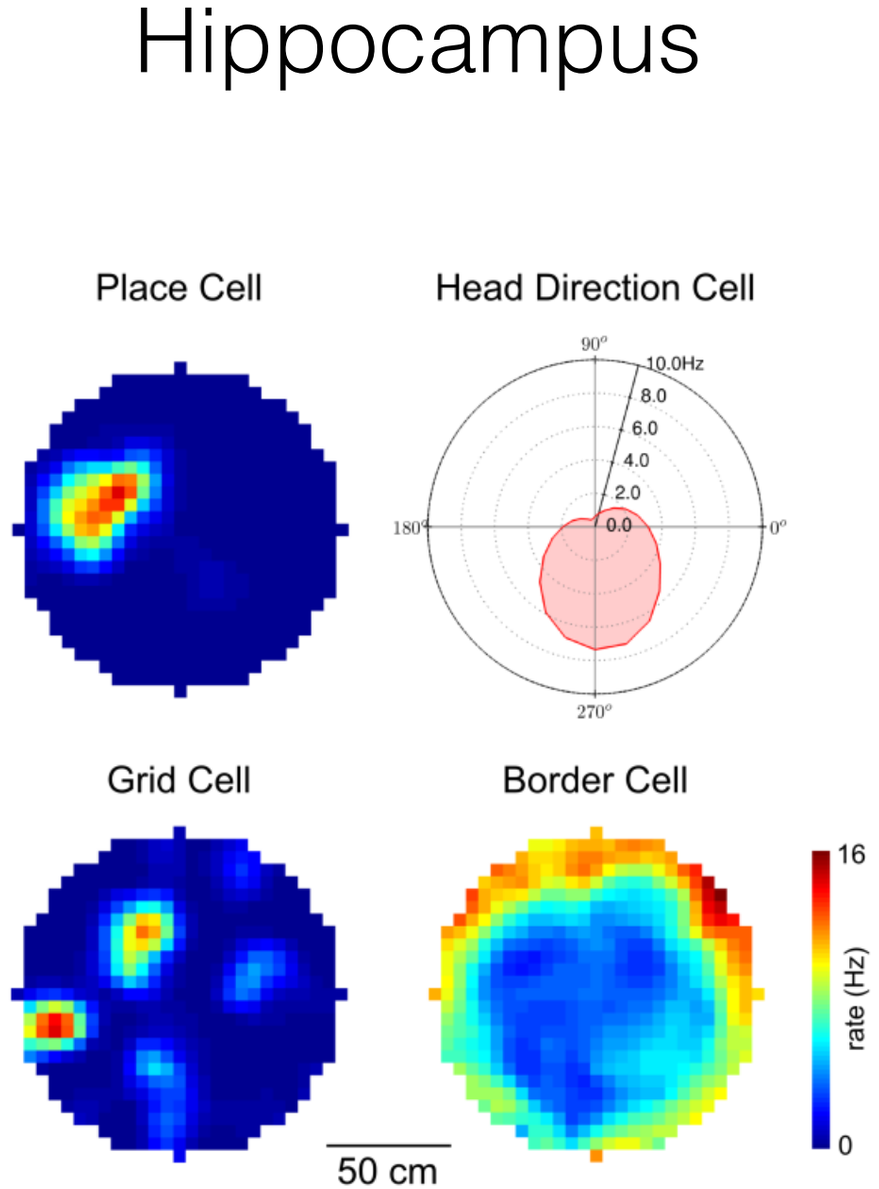

Hierarchy of visual processing

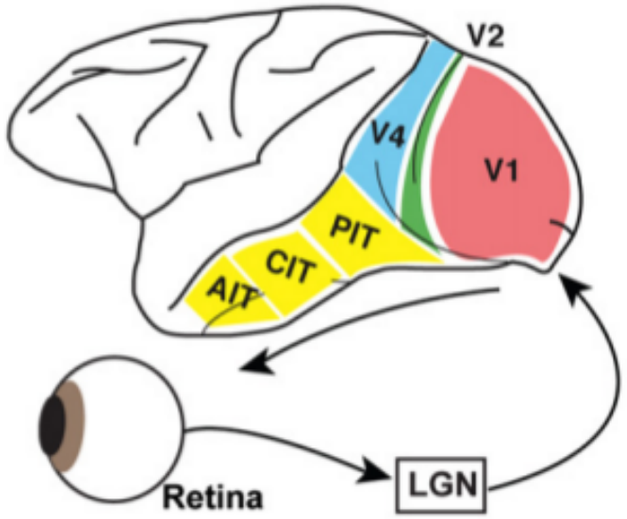

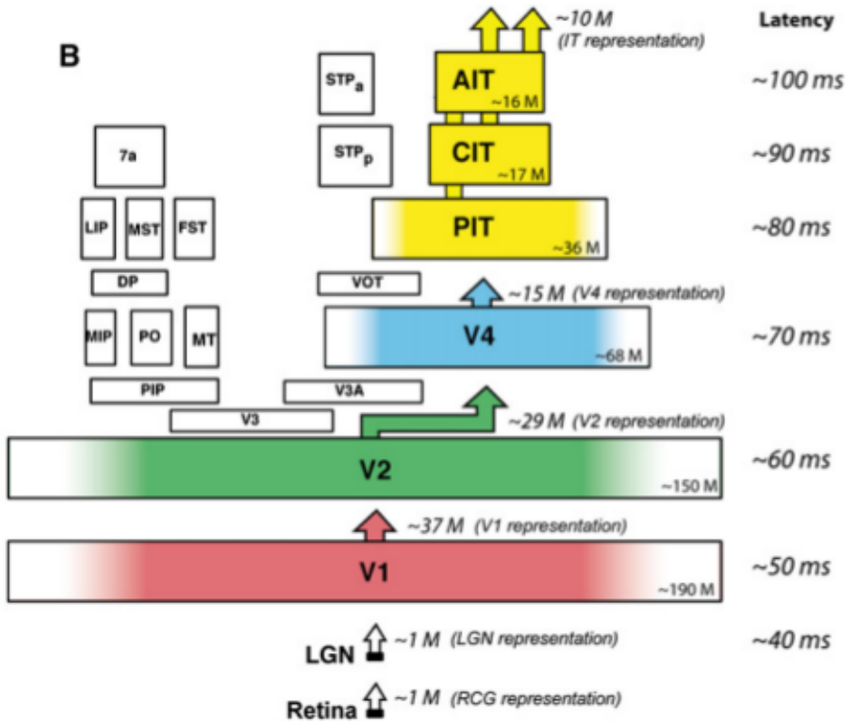

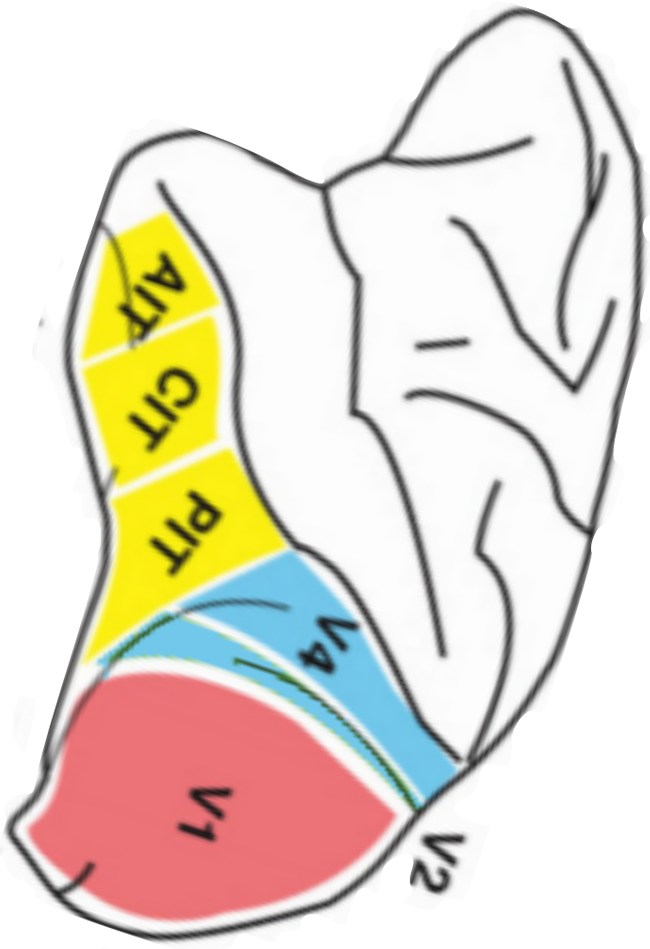

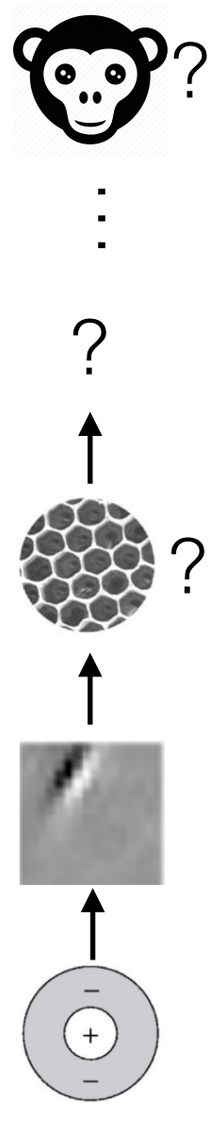

Thalamus ← Retina

Recognition

Secondary visual cortex

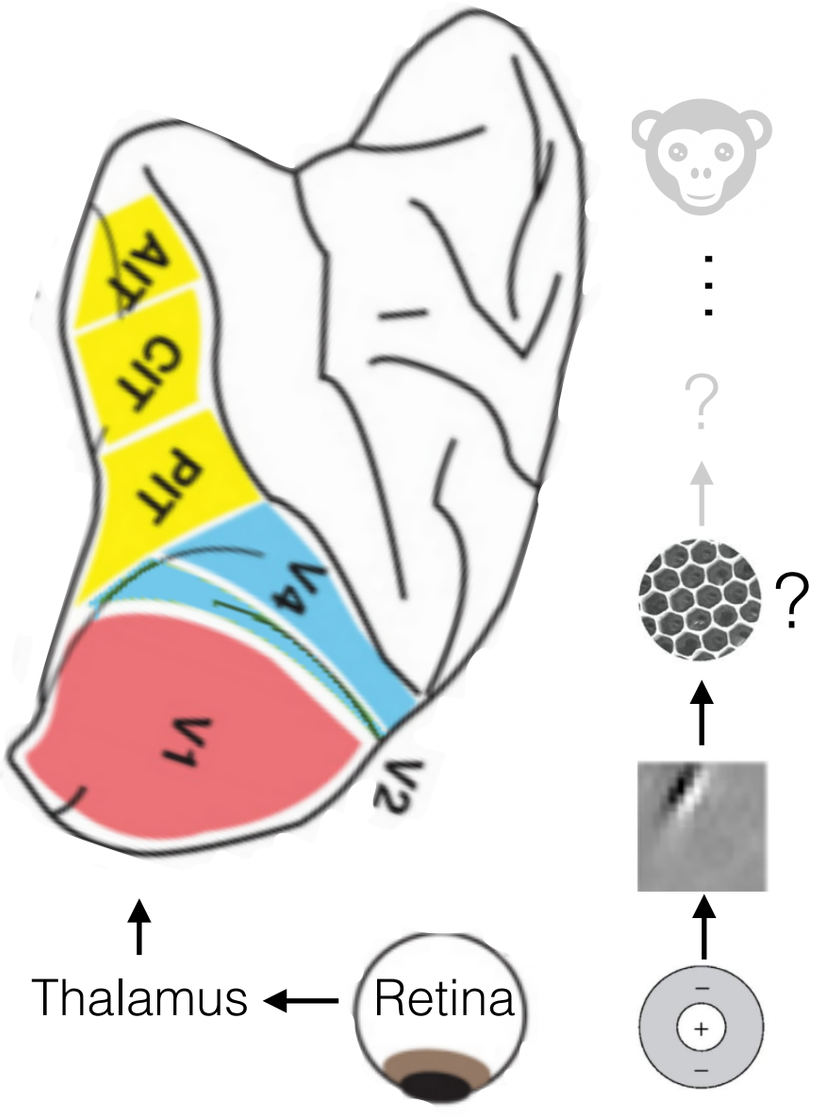

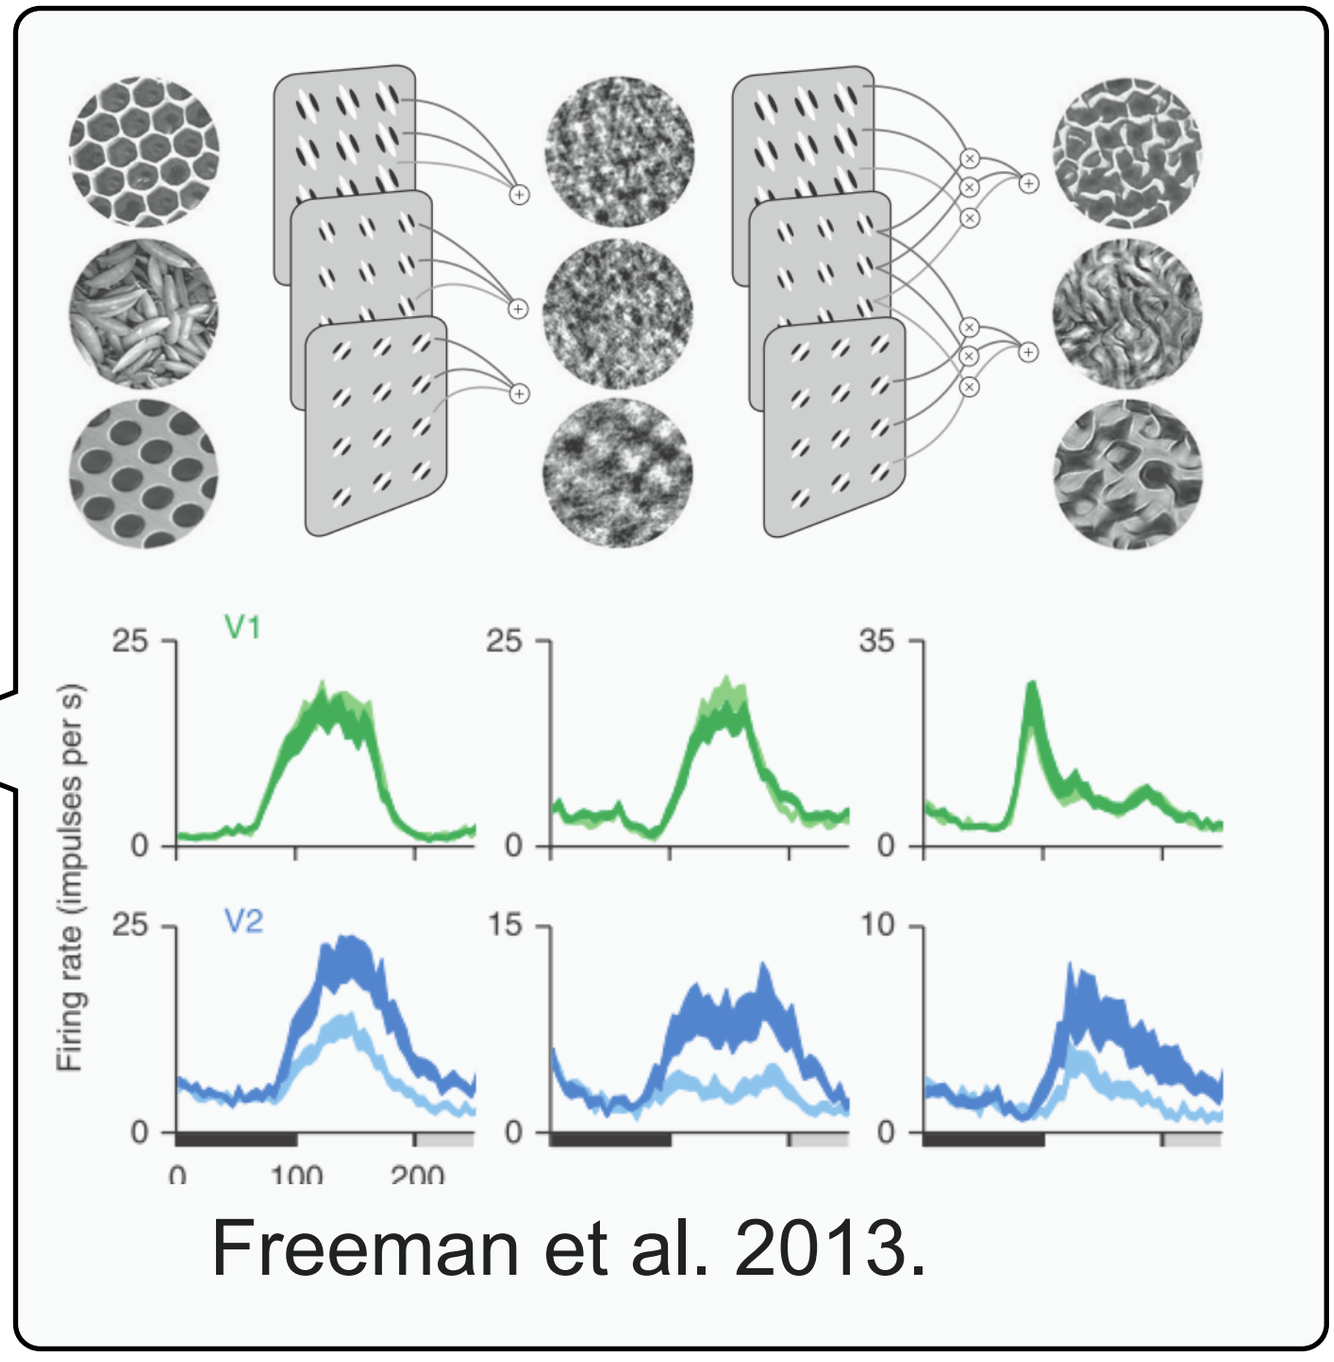

Freeman et al. 2013.

Selectivity in higher-order visual areas

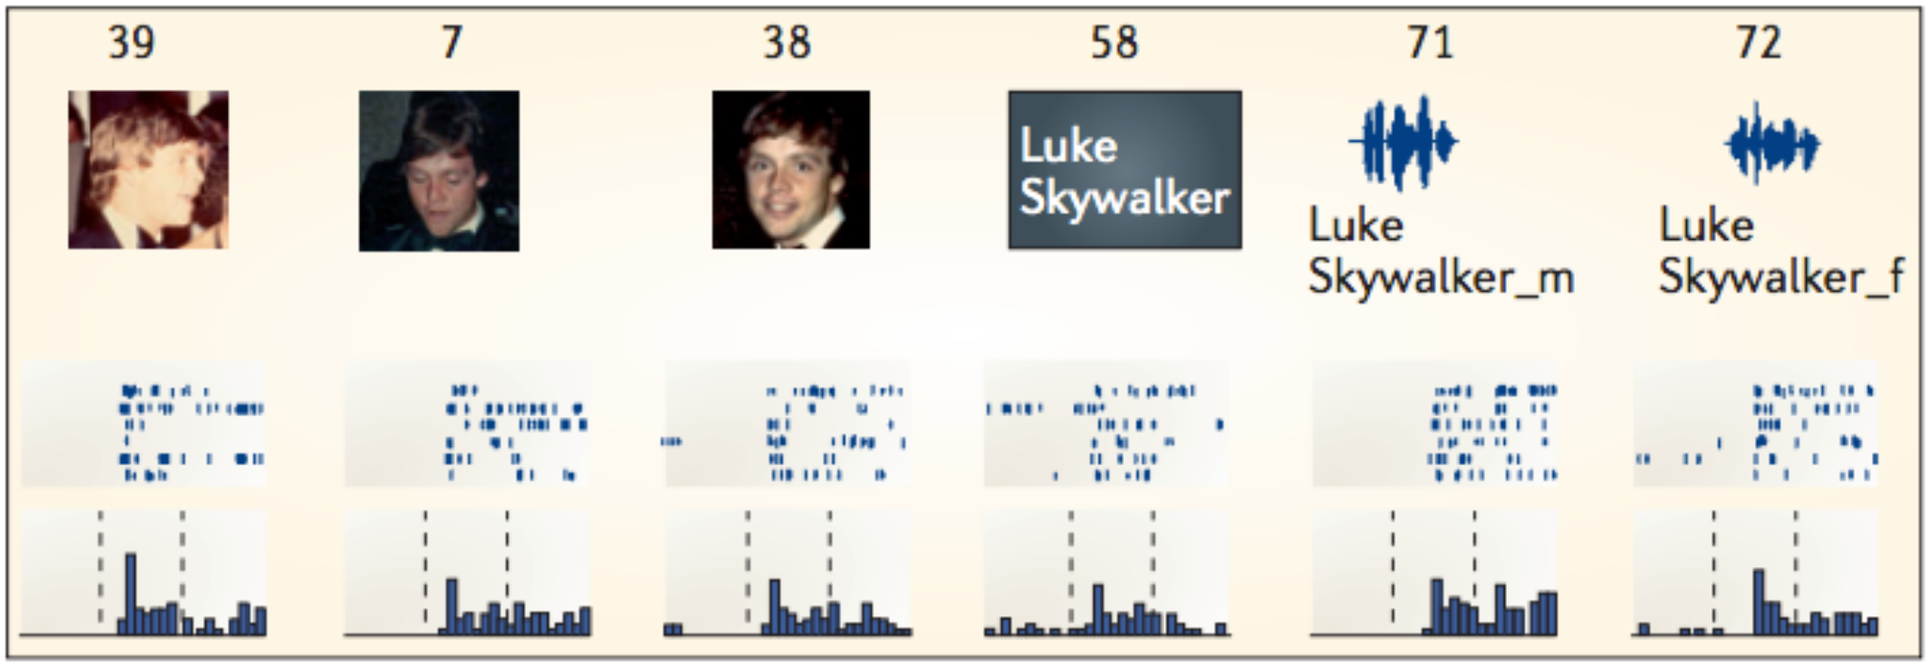

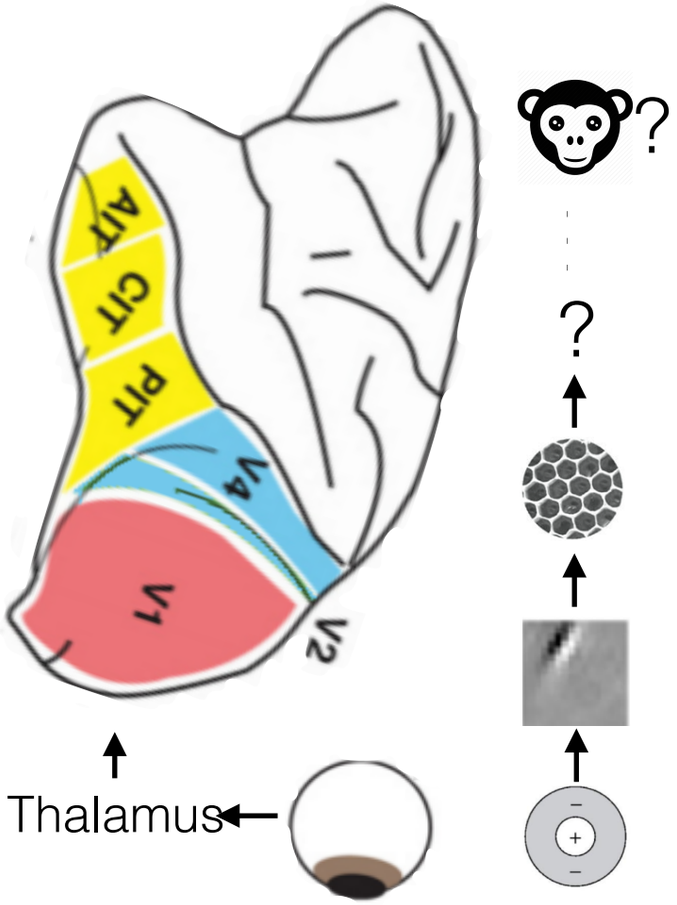

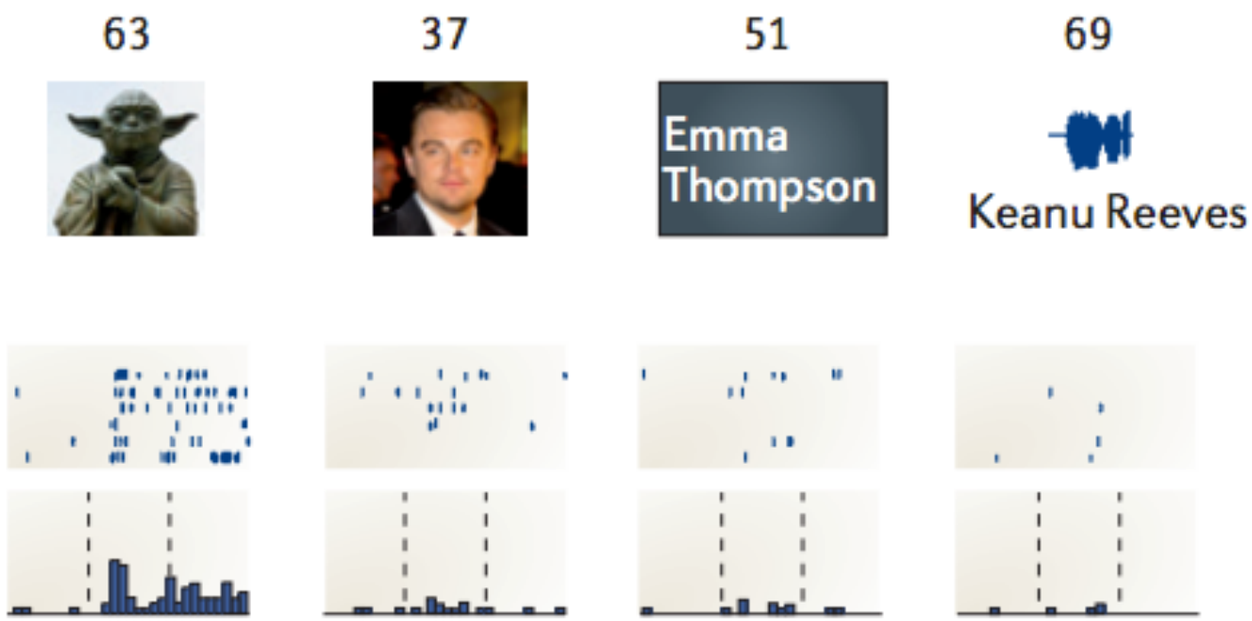

The way forward

- Now we have some well-grounded ideas about which parts of the brain are involved in certain tasks

- We can try to establish how exactly they implement computations

- In this class we will mainly look at examples regarding perception

    - this is a very large part of brain functionality: everything that maps sensory input to knowledge, including sensory processing, memory formation, learning, language processing, etc.

    - due to time restrictions, we will defer non-visual perception, motor function, decision making and higher-level phenomena (such as consciousness) to the very end, if we'll have time for them at all

    - In the spirit of normative modelling, now we have to look at how to define knowledge and its acquisition in mathematical terms

- and when we have a nice formal framework to handle information processing, we have to do two things

    - check whether our framework provides a sensible description of what animals and people do in different situations, i.e. whether it predicts behavioural data from psychological experiments well

    - then we can try to relate it to biophysical quantities measured from the brain