Computational Cognitive Science

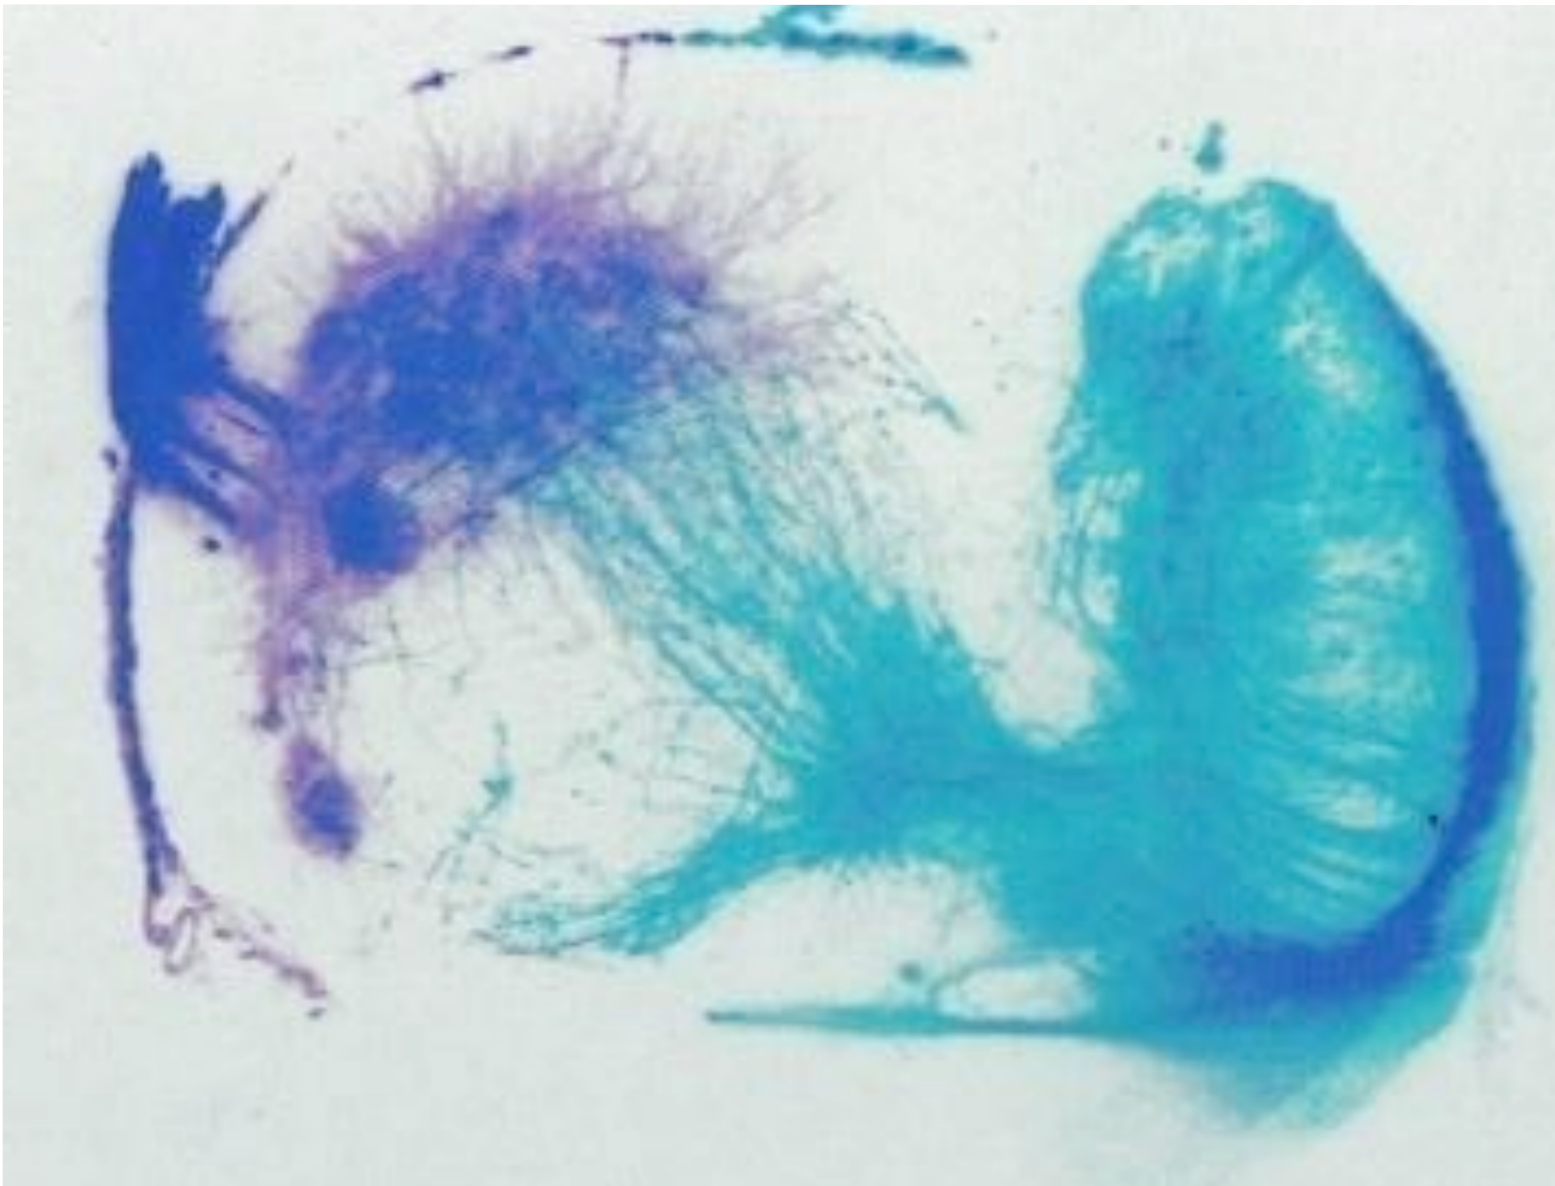

Lecture 1: Introduction

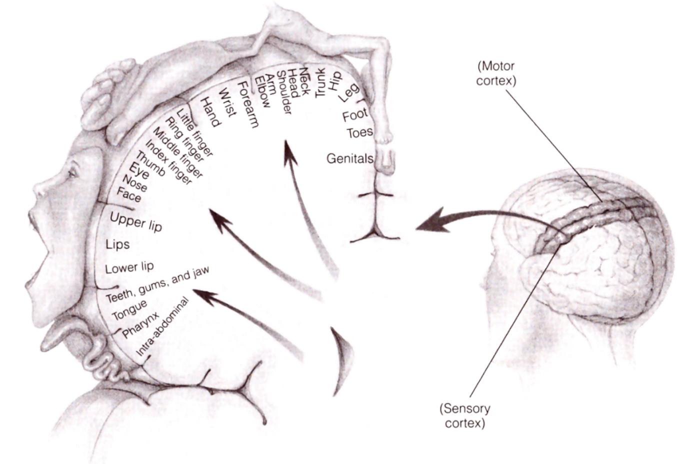

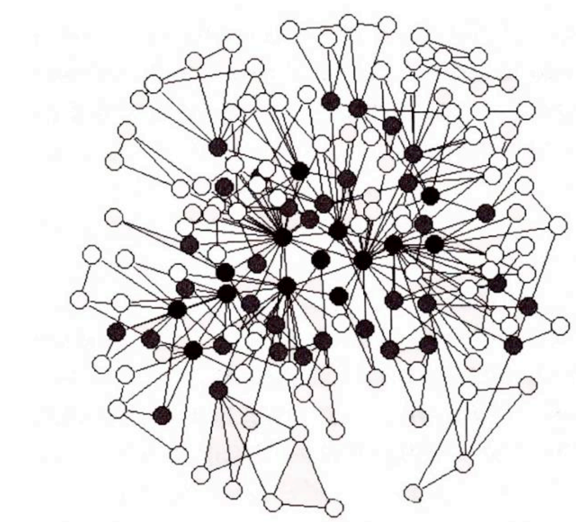

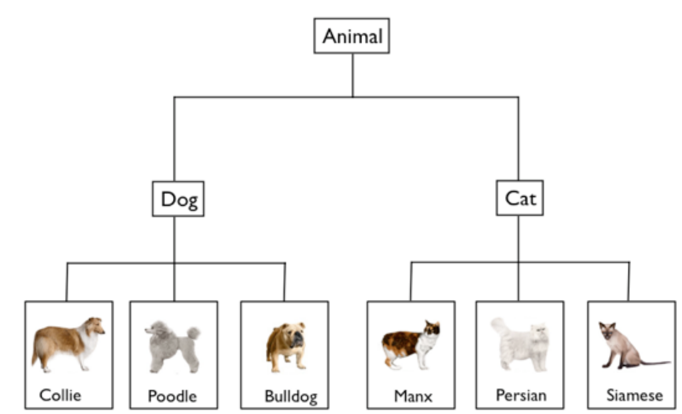

Lecture outline

- Boring logistical details
- What is computational cognitive science?
  - Why is human cognition a puzzle?
  - What kinds of questions can we address with computers and computer models?
  - What makes a good model?
- Plan for the semester

Who are we?

Dr Dan Navarro
daniel.navarro@adelaide.edu.au

Dr Amy Perfors
amy.perfors@adelaide.edu.au

Language

Concepts

Decision making

Lecturers in psychology, but our research is primarily concerned with building computational models of the human mind

Goals of this class

- Learn about key issues and theories in cognitive science
- Explore how mathematical and computational modelling illuminates our understanding of these topics

Both elements matter!

The models without the theories are irrelevant and unimportant (though sometimes have applications to AI or machine learning)

Theories without models are hard to test or define

Tutorials

Tutor: Wai Keen Vong
Weekly: Tuesdays 3.10-4.00pm
Ingkarni Wardli, B21
waikeen.vong@adelaide.edu.au

Tute participation: 4% of final mark

First tute is next week!

Assessments

- Problem sets (three, worth 7% each)
  - Programming, simple math, or conceptual practice problems
  - See information sheet for due dates
- You can program in R, Matlab/Octave, C, or Java.
  - However, our example code (in class and on some problem sets) will be in R so you are highly encouraged to use it.
  - R is open-source and free, with syntax similar to Matlab.
- R Resources are available on MyUni as well as (for now) our website: ua.edu.au/ccs/
  - We plan on gently introducing you to it via problem sets!

Assessments

- Problem sets (three, worth 7% each)
  - Programming, simple math, or conceptual practice problems
  - See information sheet for due dates
- Final project (15%)
  - Literature review or implemented project (preferred)
  - Postgrads must be slightly more complicated
  - Details in information sheet on webpage
  - Proposal due 2 May, final project due 13 June
- Exam (60%)
  - Problems are conceptual or mathematical
- Tute attendance and participation (4%)

Resources

- All available on MyUni
- We will rarely if ever have required readings
- We will often put (optional but potentially helpful) readings up in the appropriate folder. Use if you wish!
  - Technical notes
  - Theory articles
  - Chapters from textbooks
- Useful generic background
  - Russell & Norvig, *Artificial Intelligence.* Any edition
  - Goldstein (2011). *Cognitive Psychology*. Or any general text.

Lecture outline

- Boring logistical details
- ➡ What is computational cognitive science?
  - ➡ Why is human cognition a puzzle?
  - What kinds of questions can we address with computers and computer models?
  - What makes a good model?
- Plan for the semester

Humans are smart.

Our encephalisation quotient (brain to body mass ratio) is off the charts compared to other animals

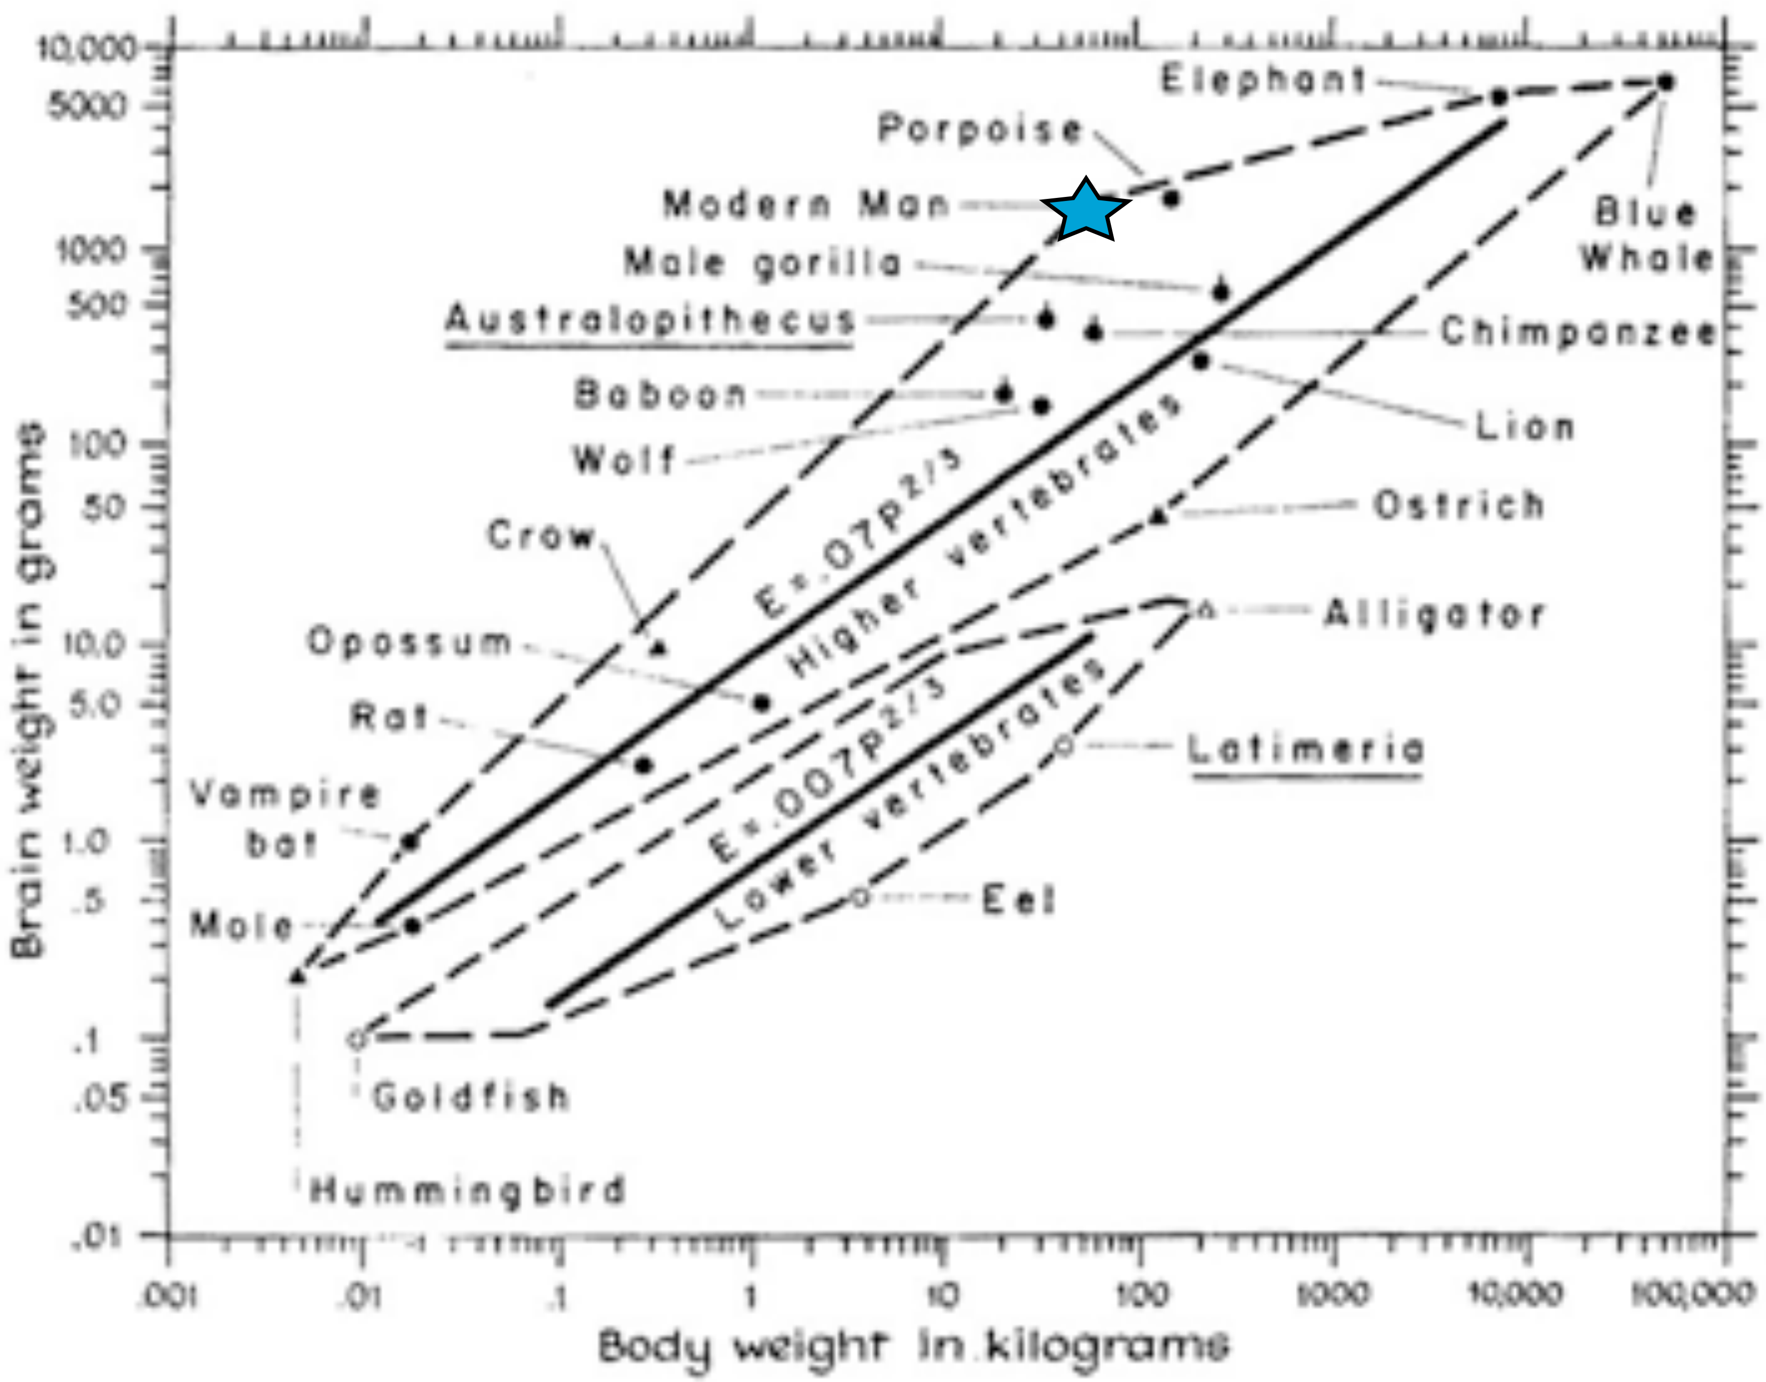

Humans are smart. Like, really smart.

Even compared to other primates!

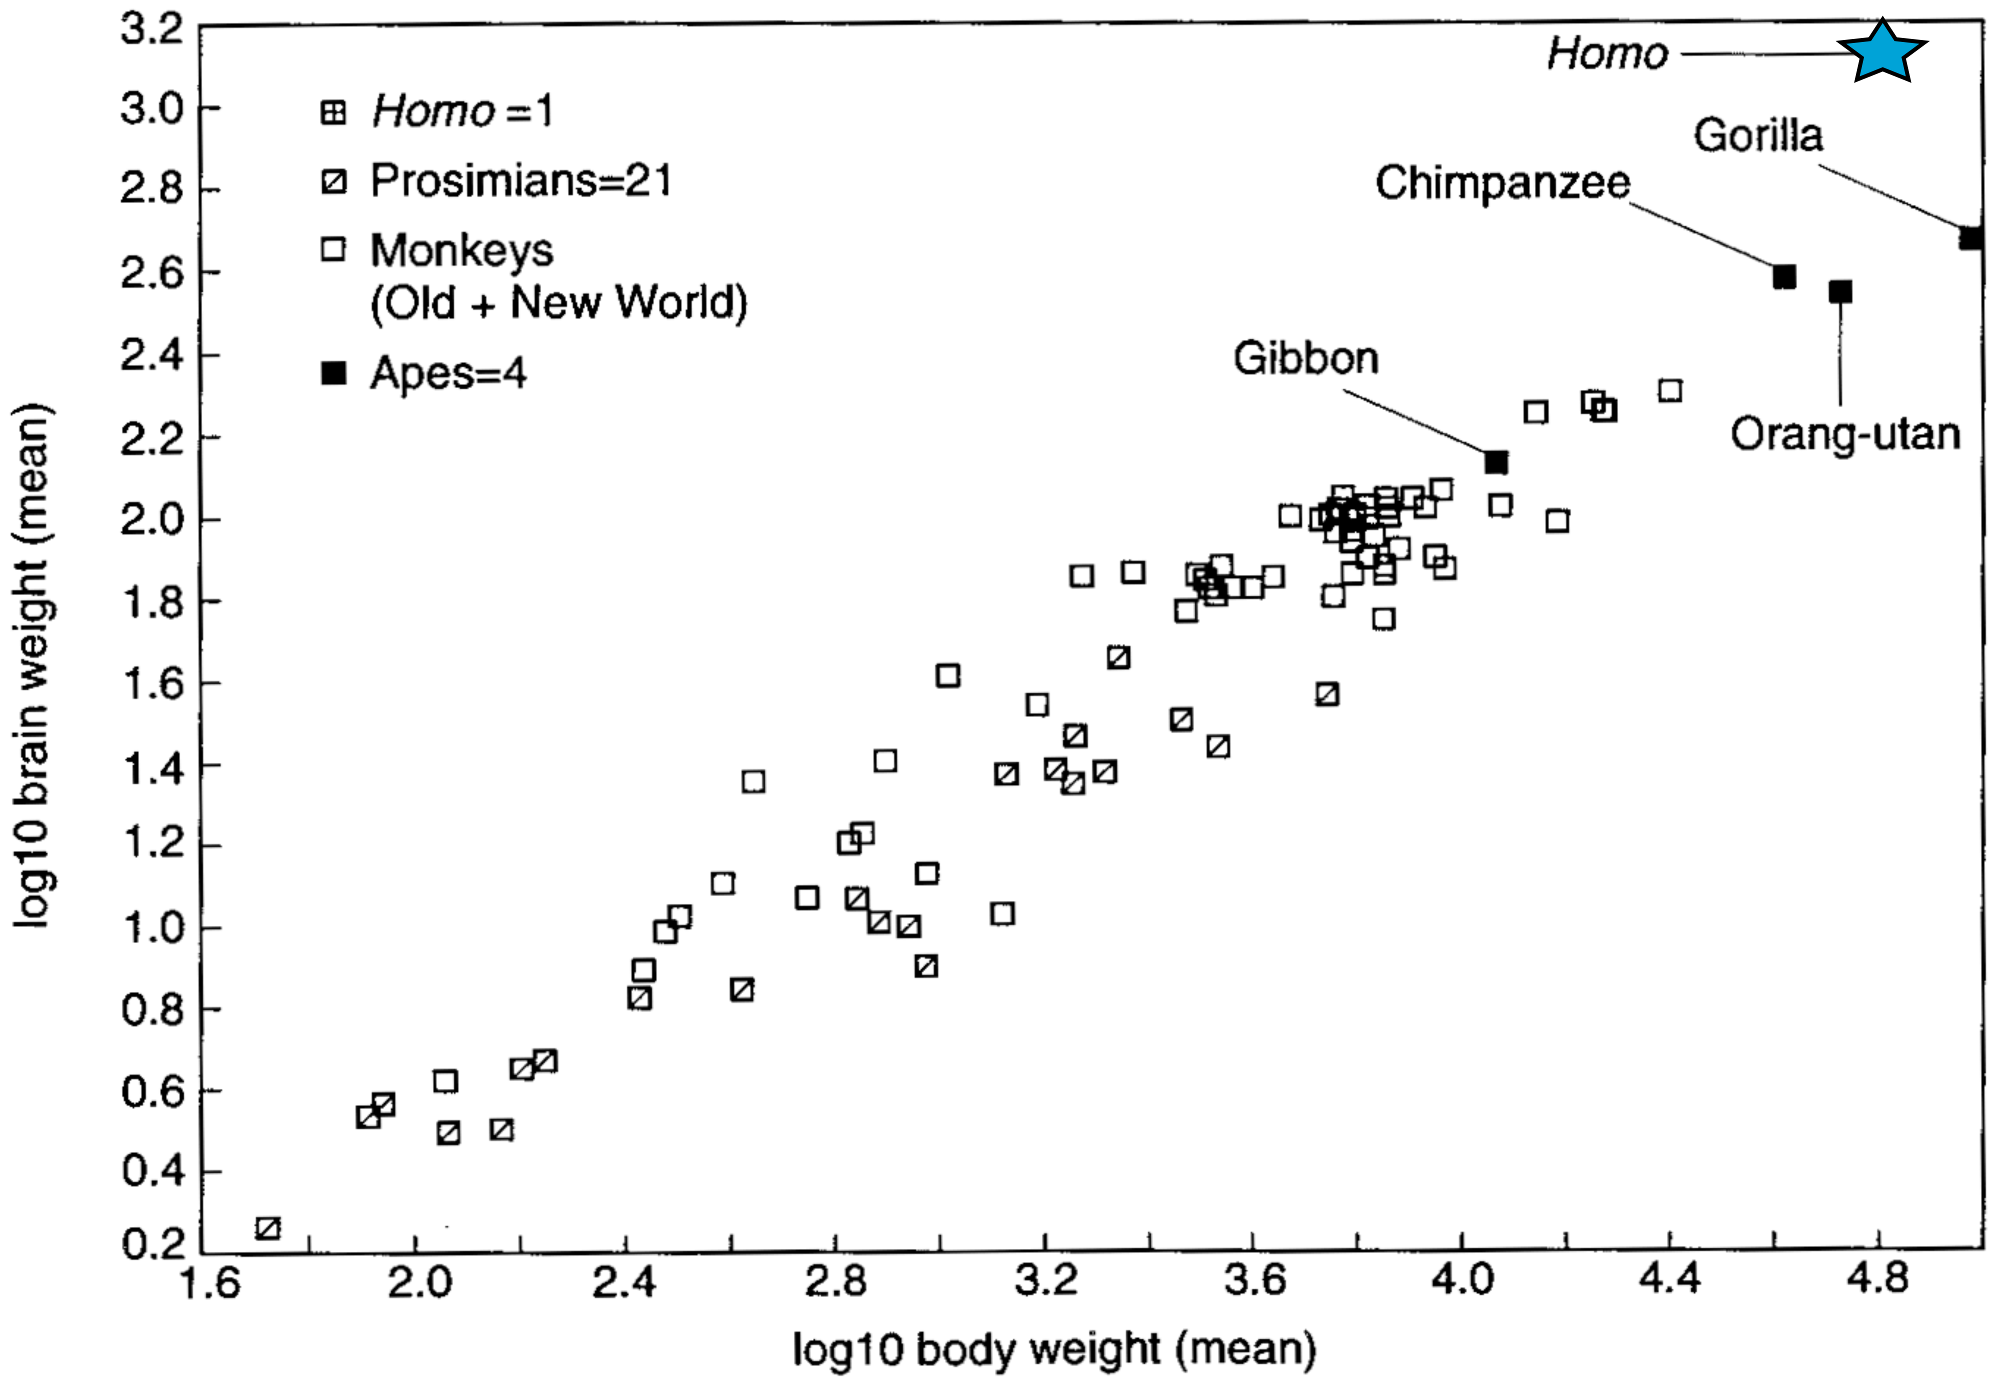

What is that intelligence for?

The world has inherent structure -- lots of it.

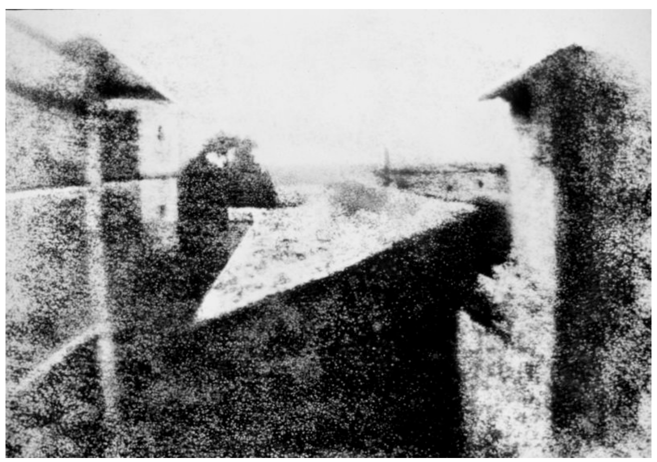

Spatial distribution of light intensity uncovered in an early photographic image (1836)

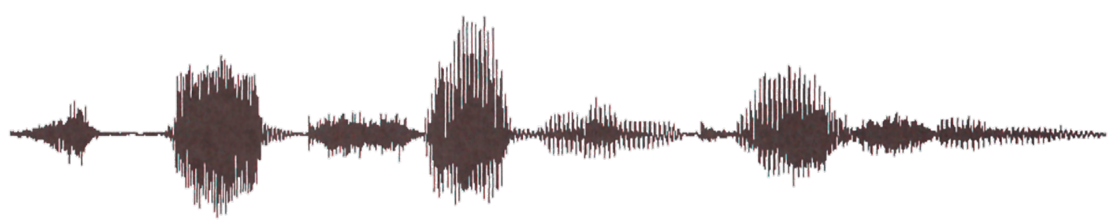

Variable gamma-ray sensitivity of the interior of a truck

Atmospheric vibrations corresponding to the English phrase “speech segmentation”

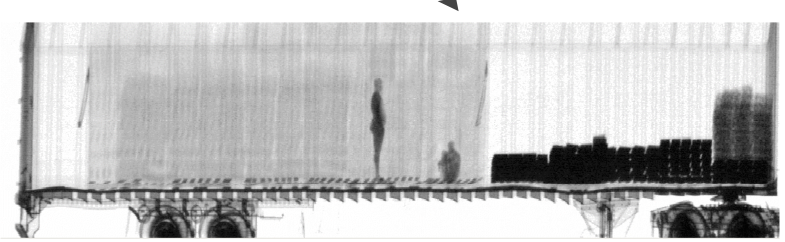

What is that intelligence for?

Sensory systems detect some of that structure

Our sensory equipment is sensitive to only a slice of the world (no x-ray vision!).... but it's still a massive torrent of data

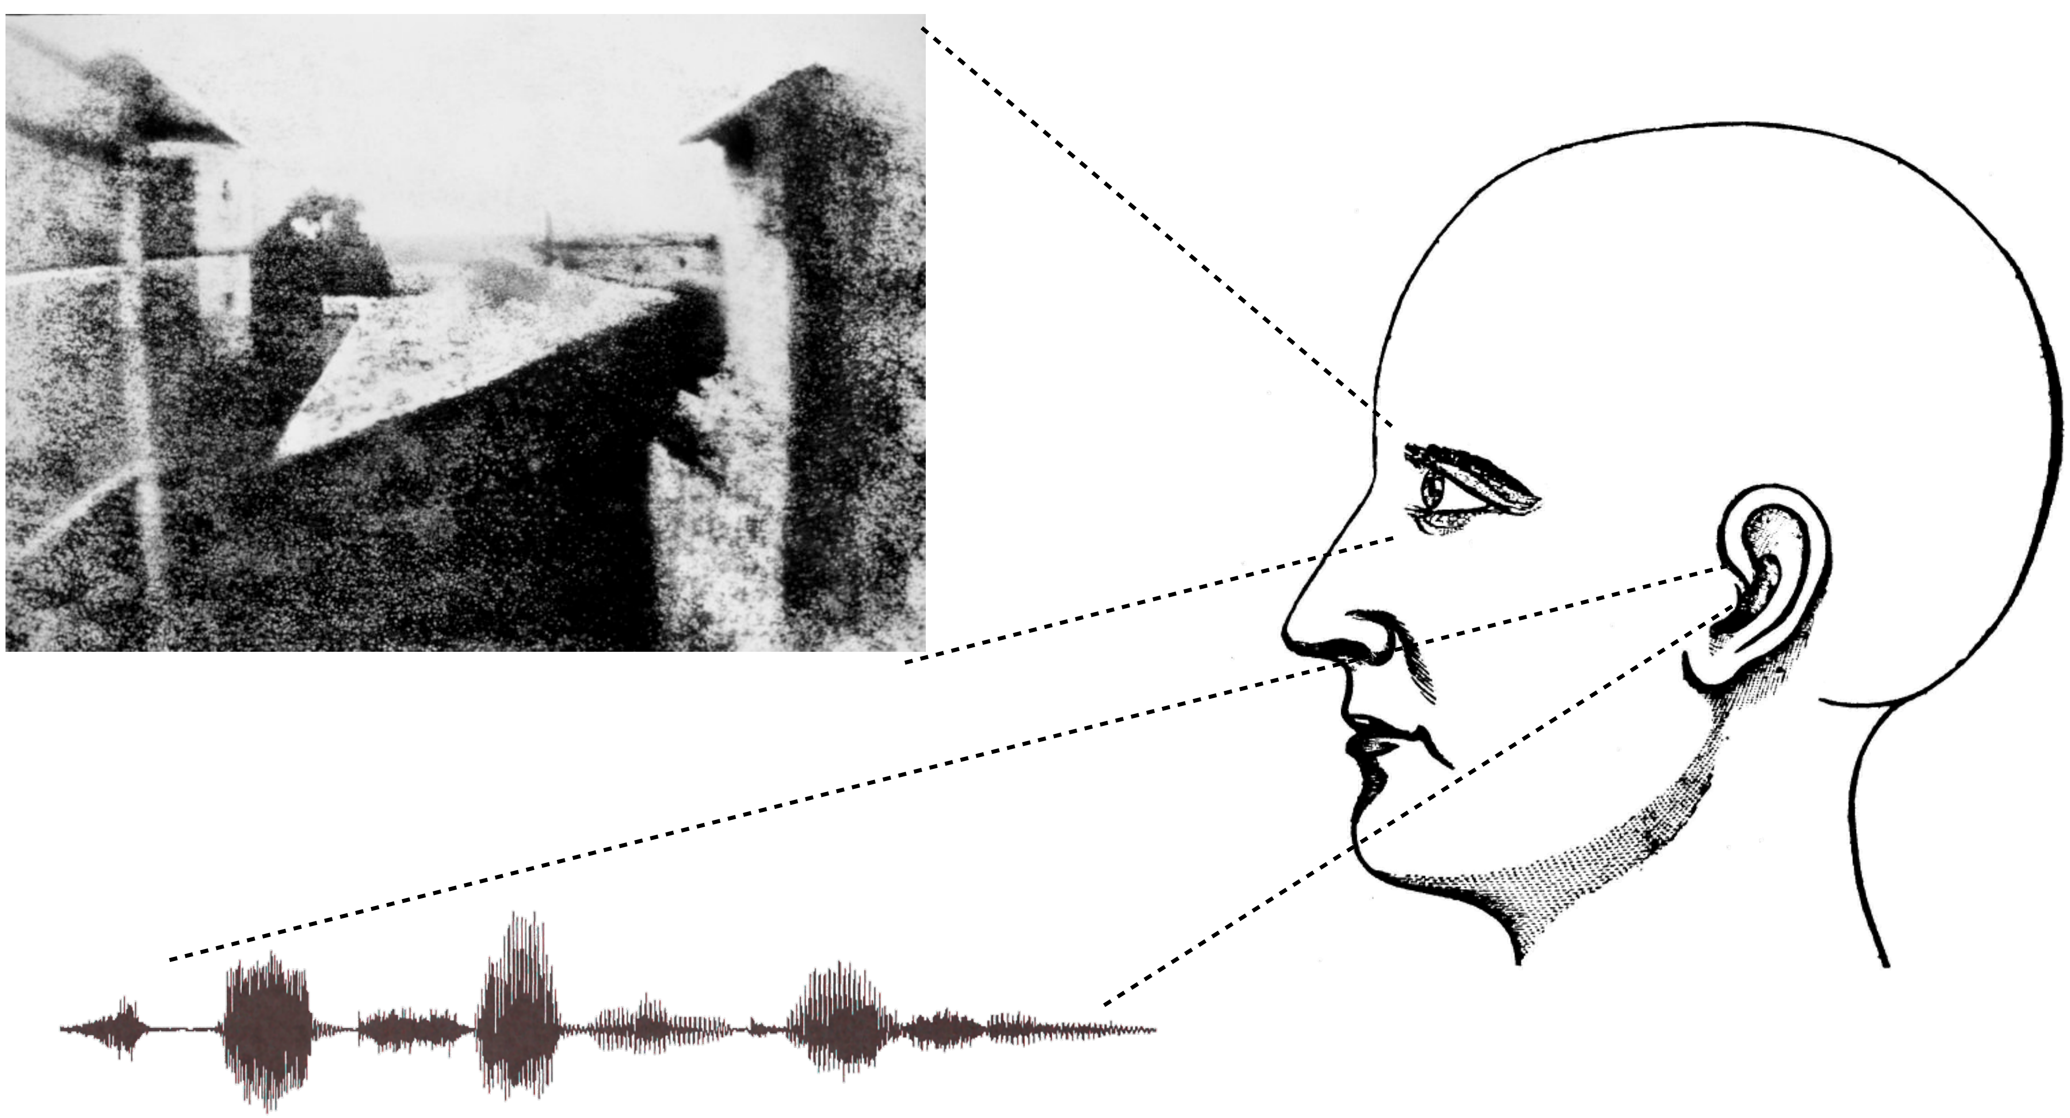

What is that intelligence for?

We have a brain under there that uses all of this information
.... and it is *very* complex

A human brain has about 10^{15} synapses
which operate at about 10^{2} per second
implying about 10^{17} bit ops per second

So 1 second of brain activity would fill up about
40,000 ordinary 300GB hard drives

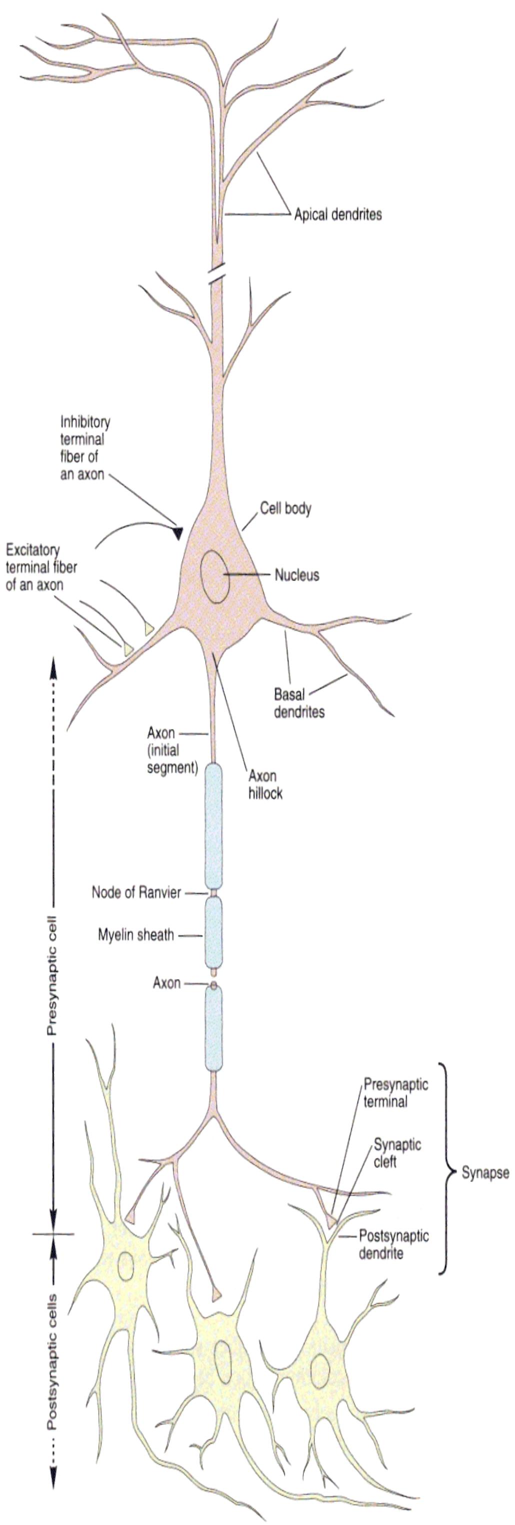

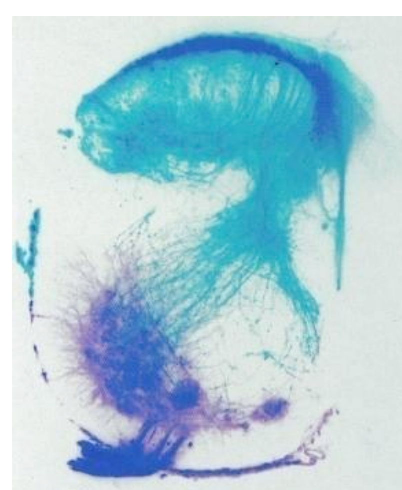

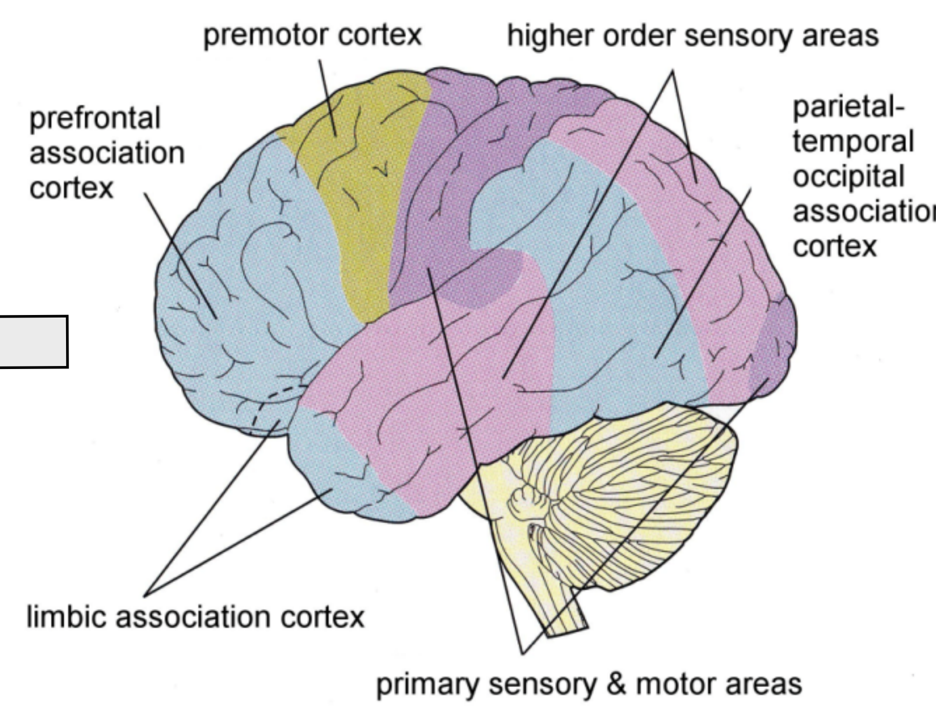

What is that intelligence for?

We have a brain under there that uses all of this information
.... and it is *very* complex

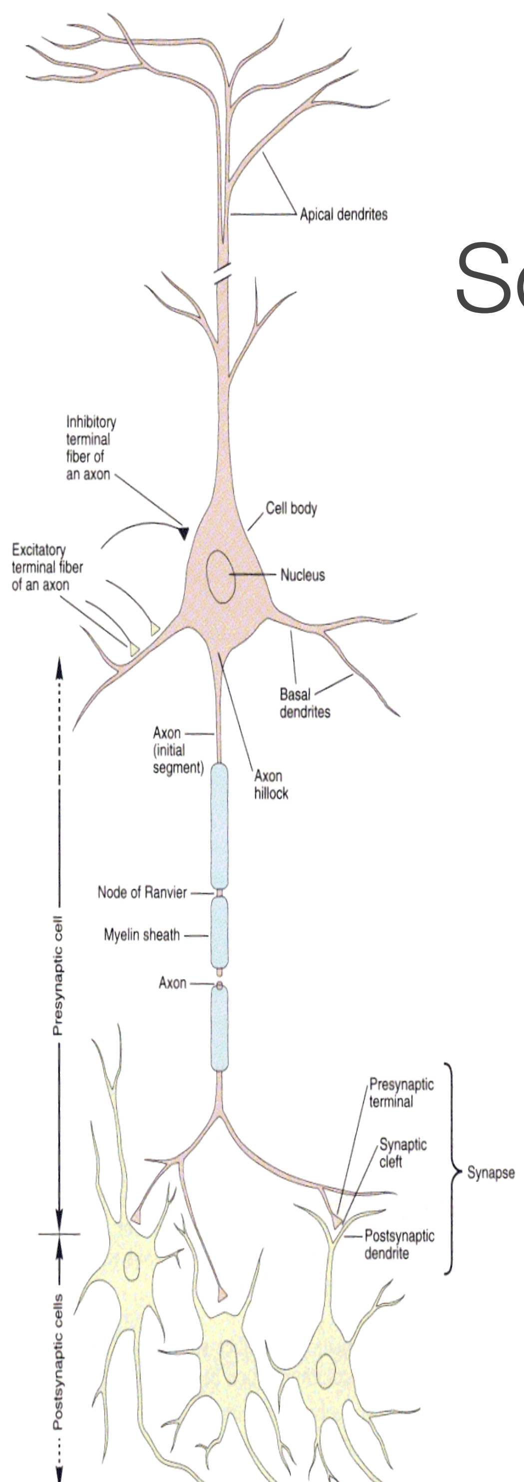

So it's probably doing something very important!

Accounts for 20% of body's oxygen consumption,
even though it is only 1.3kg. Very expensive.

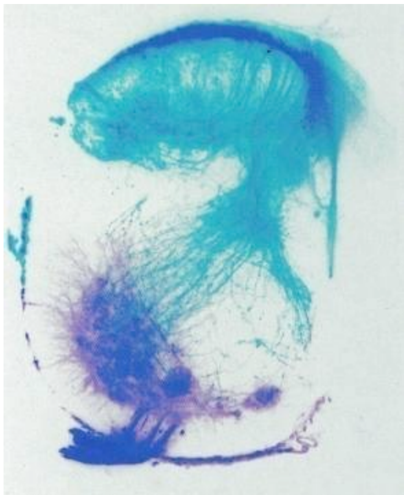

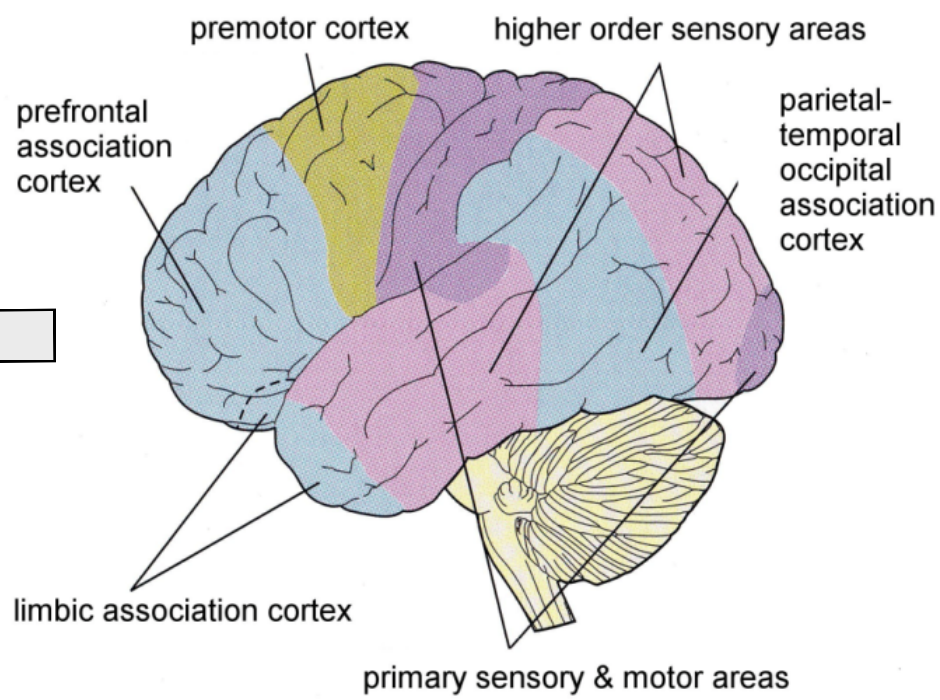

Creating internal representations

We don’t perceive the world “as is”; we perceive it in terms of a set of structures we extrapolate from sense data

Optic information from the environment is mapped out in the visual cortex

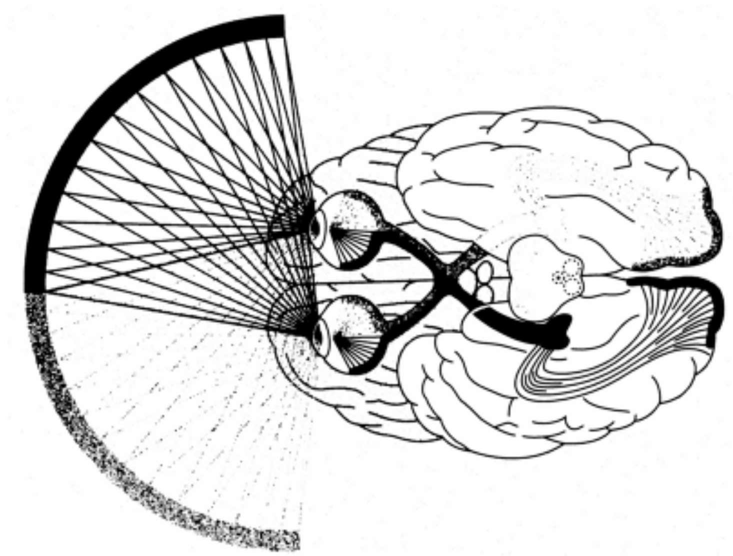

Touch (“somatosensory”) data is also mapped internally, with different parts of the body mapped onto different parts of the cortex

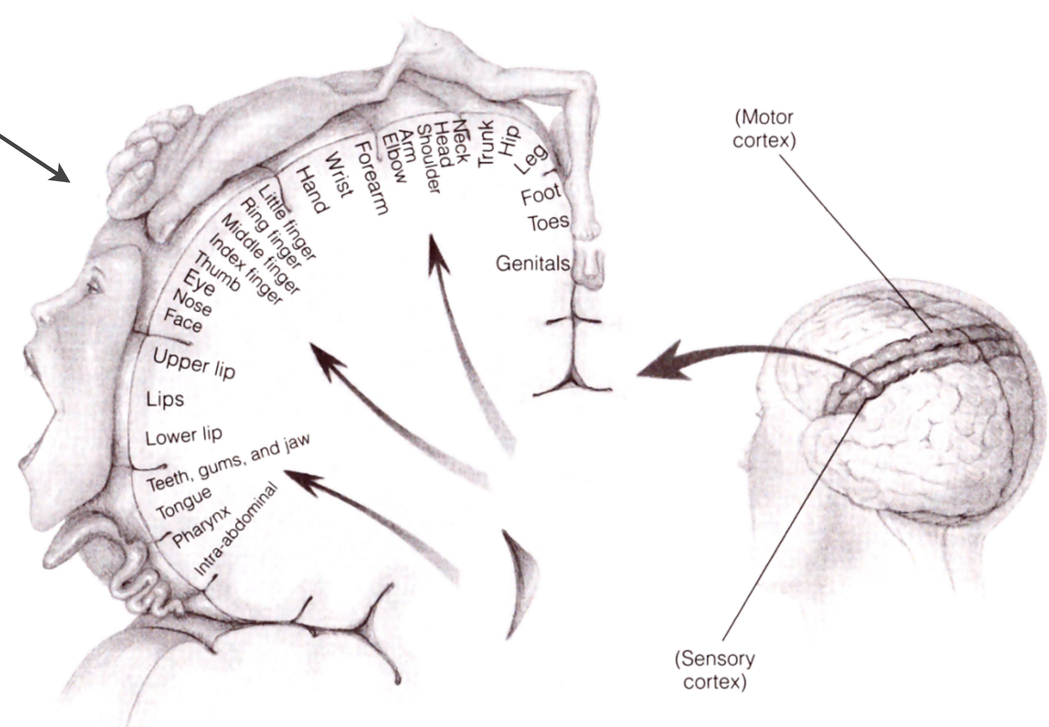

These are basically the “data structures” of the mind

Internal representations vary

Compare how humans map our body in somatosensory cortex with how other species do it

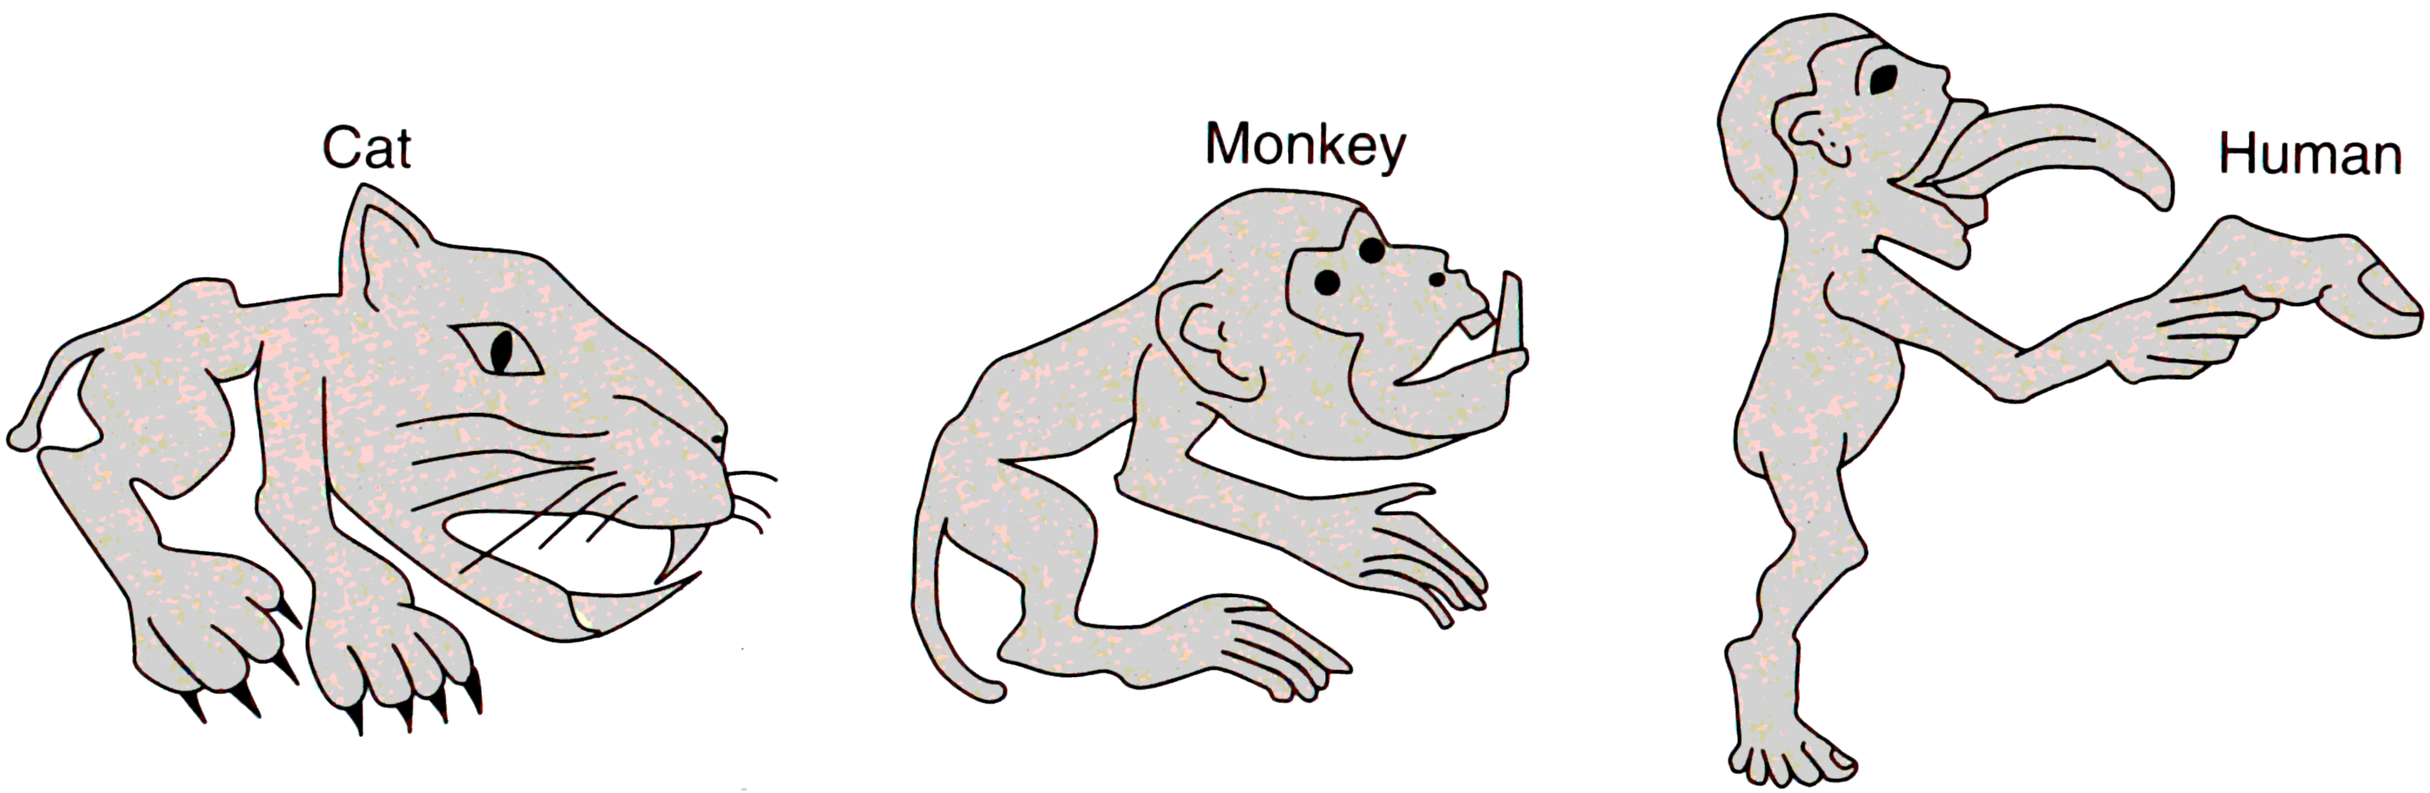

If it is important, more of the brain is devoted to it

Representation guides action

Consider the problem faced by a bird guiding a swooping descent into water. It needs to choose different actions at each stage of descent.

The movement of the different points in the optic array (i.e., the visual data) provide the cues needed to control this descent

For humans, this can get quite complex

Cognition is the study of the more complex representations we construct, and how we use them to guide actions and behaviour.

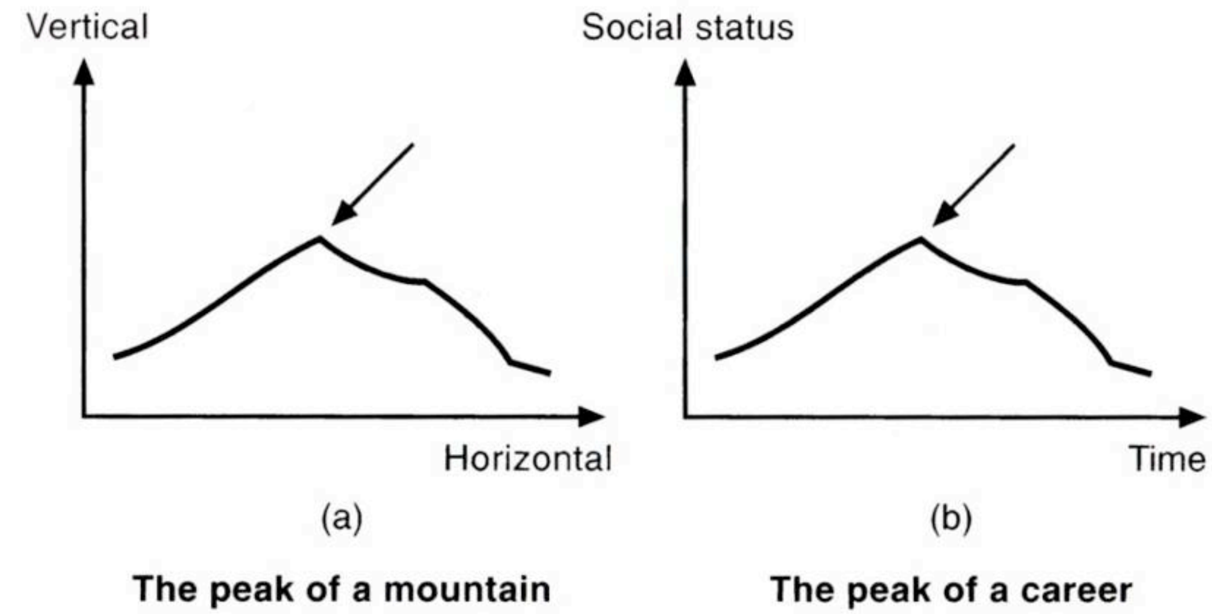

Constructing the abstract *concept* for "peak", using perceptual data and analogy

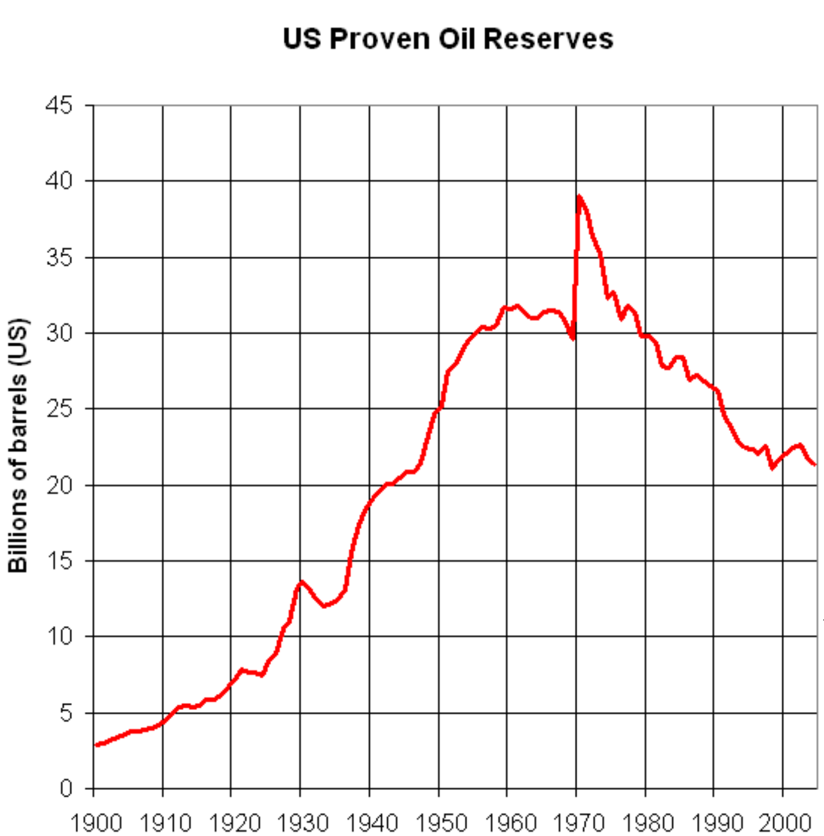

Using a "peak oil" concept to motivate a decision to drive a hybrid vehicle

How do we make sense of all of this?

- Boring logistical details
- ➡ What is computational cognitive science?
  - Why is human cognition a puzzle?
  - ➡ What kinds of questions can we address with computers and computer models?
  - What makes a good model?
- Plan for the semester

Questions of (Computational) Cognitive Science

Essentially, we're trying to reverse-engineer the human mind

Data (x) → 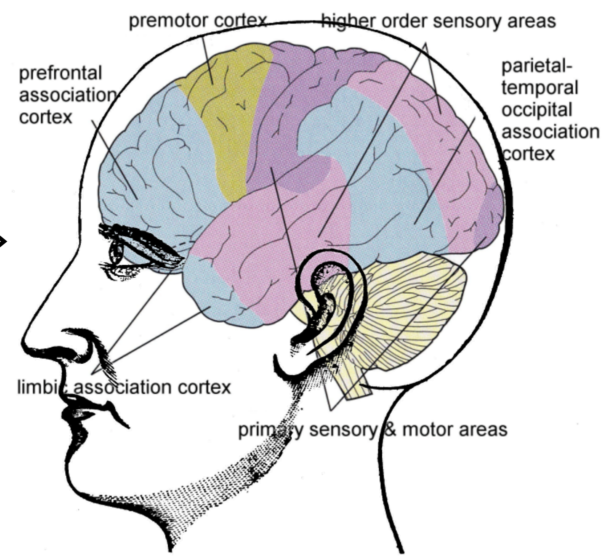

→ Behaviour (y)

Questions of (Computational) Cognitive Science

Essentially, we're trying to reverse-engineer the human mind

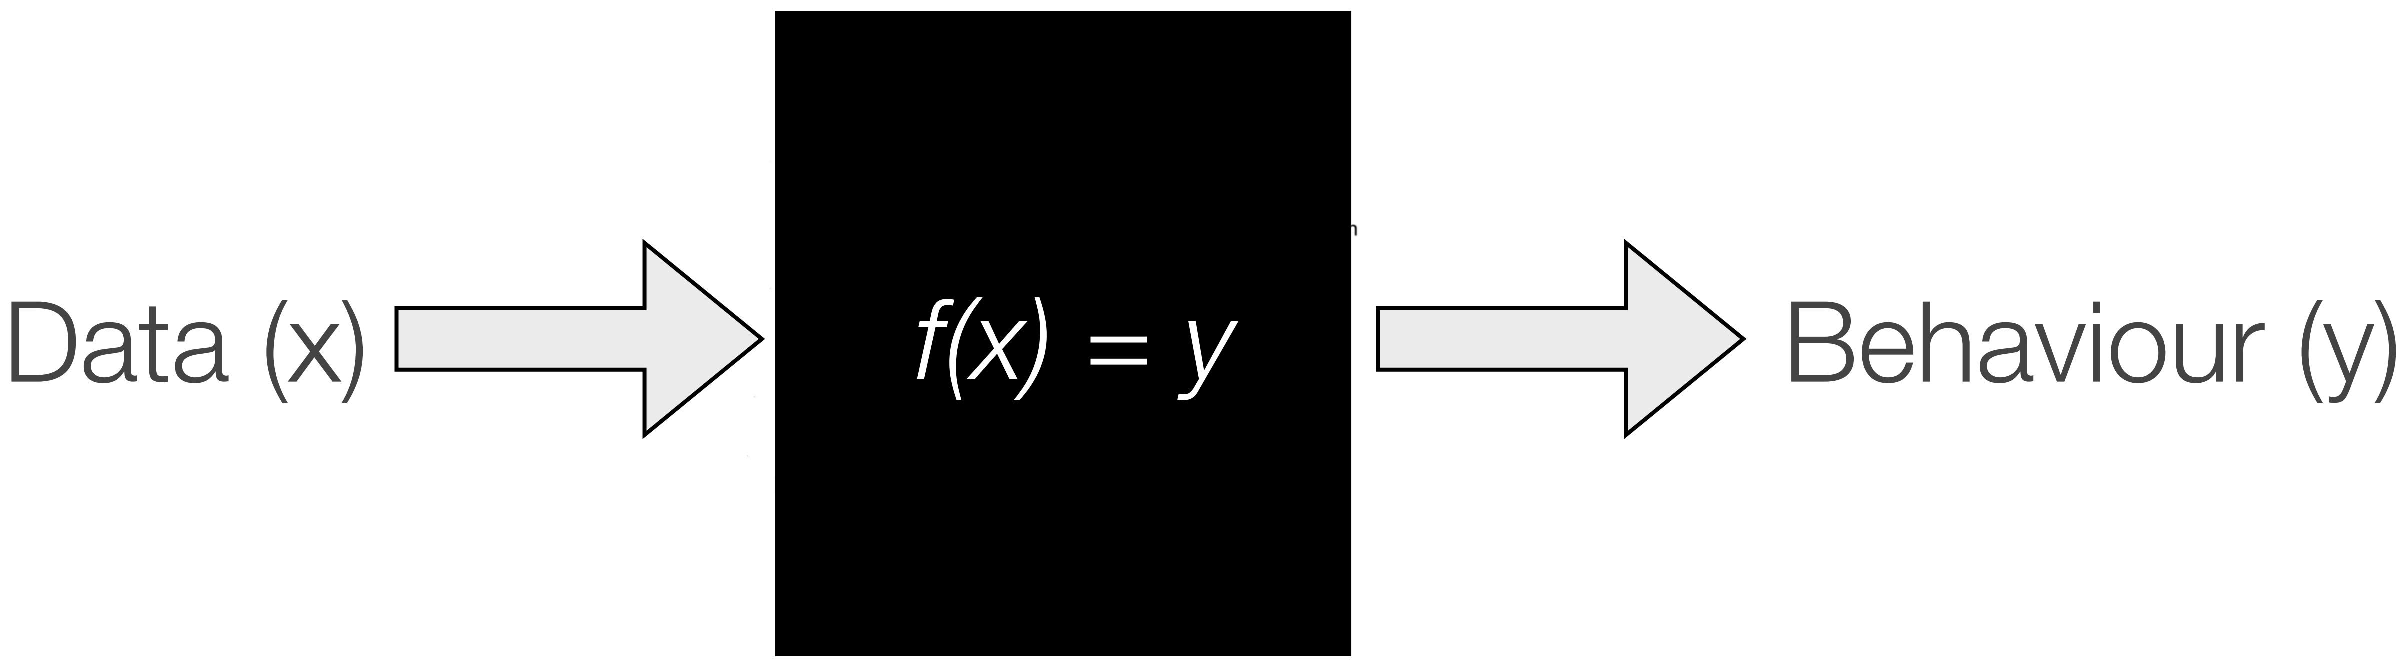

Questions of (Computational) Cognitive Science

1. What are the representations ("data structures") of the mind?

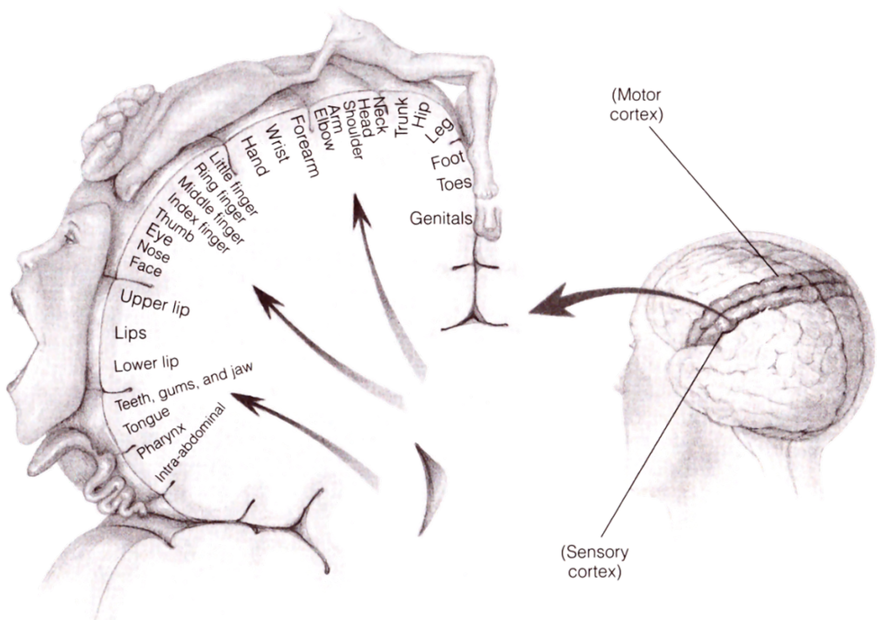

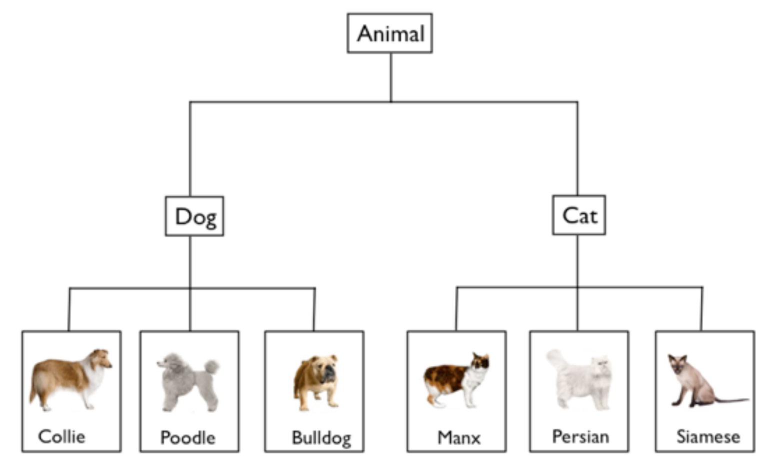

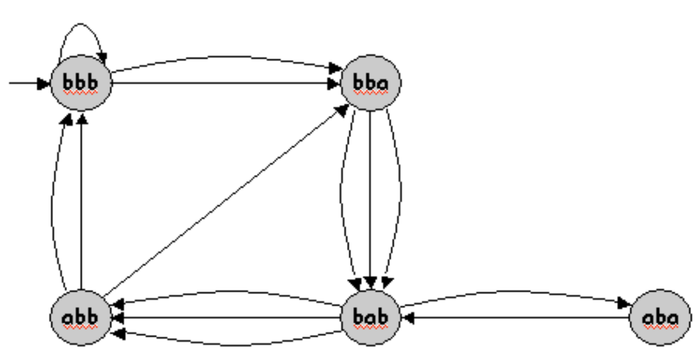

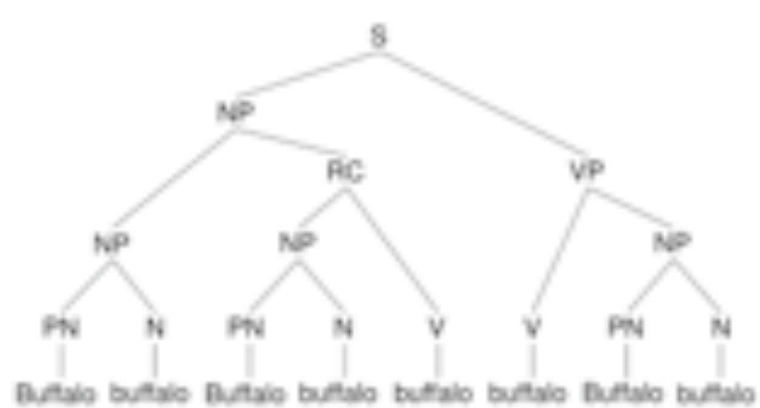

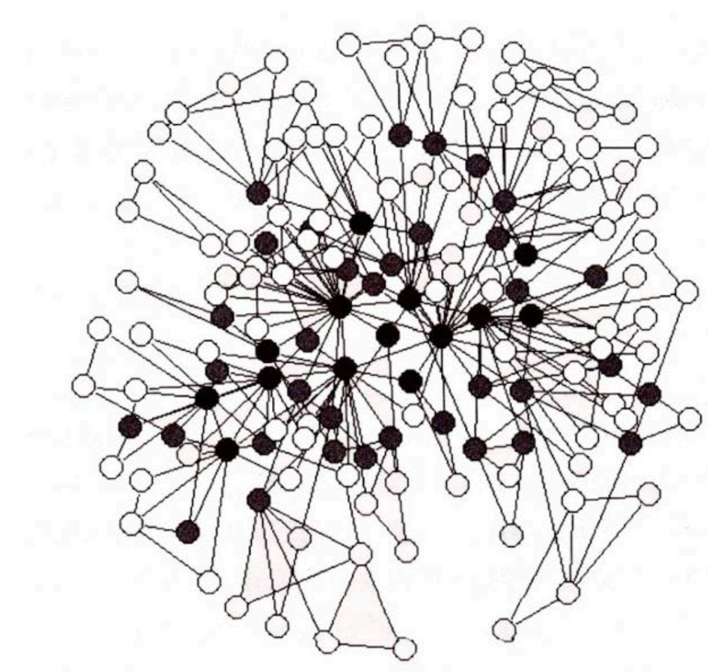

Questions of (Computational) Cognitive Science

1. What are the representations (“data structures”) of the mind?
2. How does the mind (“program”) learn (“update”) in response to environmental input (“data”)?

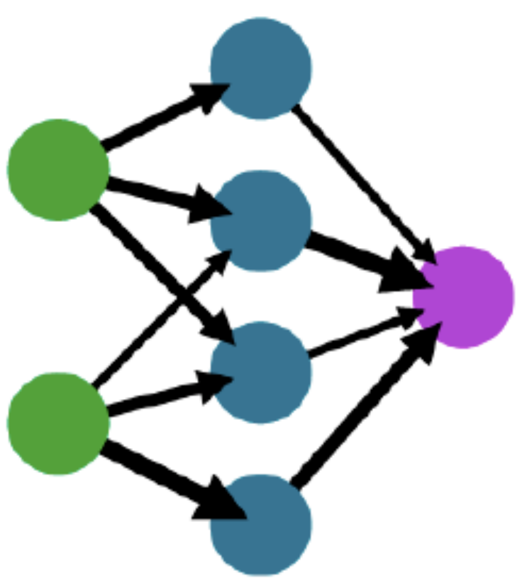

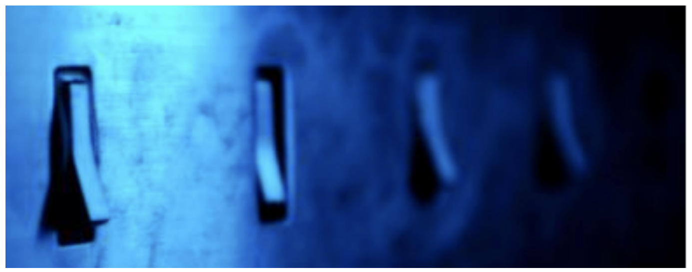

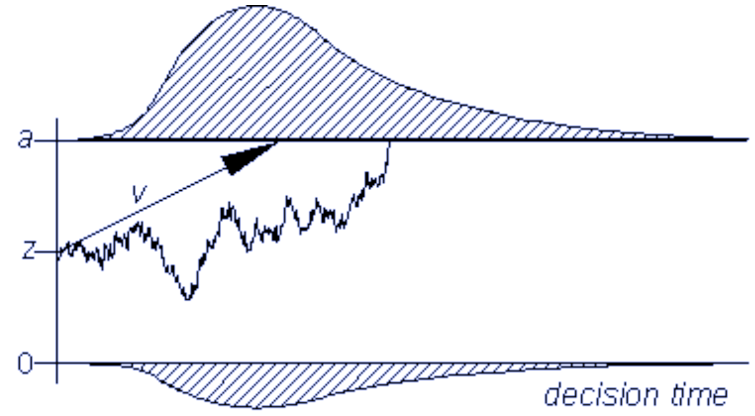

$$P(H|D) = \frac{P(D|H)P(H)}{\sum_{H_i} P(D|H_i)P(H_i)}$$

Questions of (Computational) Cognitive Science

1. What are the representations ("data structures") of the mind?
2. How does the mind ("program") learn ("update") in response to environmental input ("data")?
3. What innate biases do people have? (What is "built in" by the programmer?)

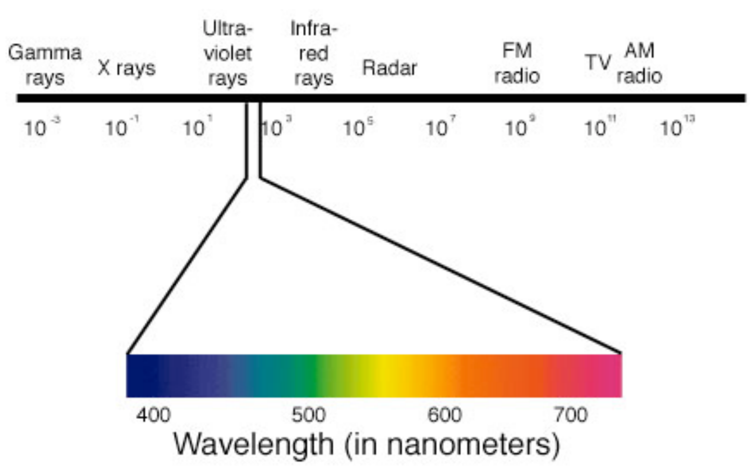

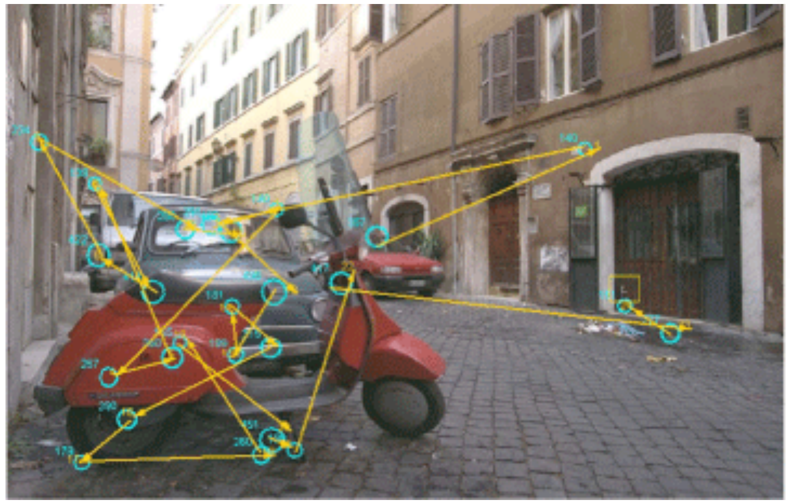

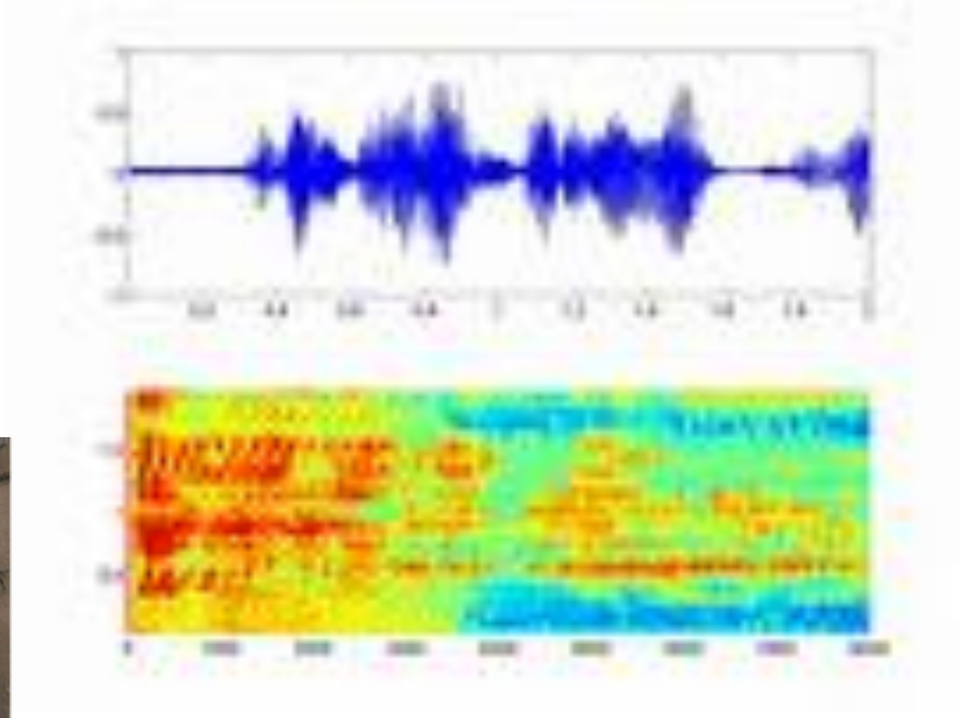

Questions of (Computational) Cognitive Science

1. What are the representations (“data structures”) of the mind?
2. How does the mind (“program”) learn (“update”) in response to environmental input (“data”)?
3. What innate biases do people have? (What is “built in” by the programmer?)
4. What is the nature of the input (“data”) received?

How do we do this?

- Boring logistical details
- ➡ What is computational cognitive science?
  - Why is human cognition a puzzle?
  - What kinds of questions can we address with computers and computer models?
  - ➡ What makes a good model?
- Plan for the semester

Desiderata

- Must be good statistical model
  - Fit the data parsimoniously (not too many parameters, etc)

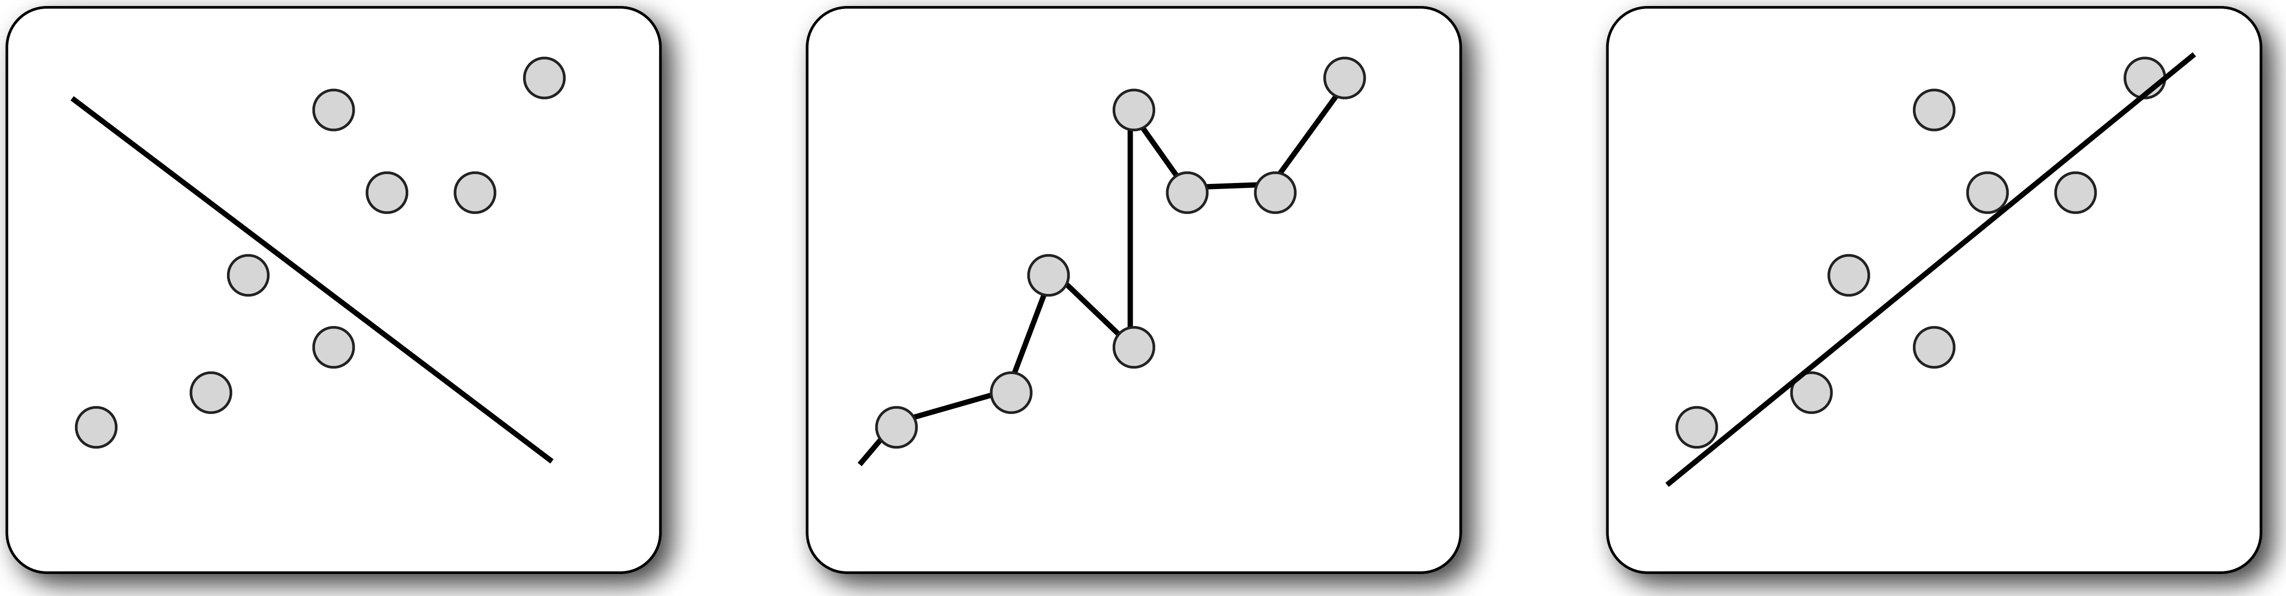

Bad (doesn't fit!) | Bad (not parsimonious!) | Good model

Desiderata

- Must be good statistical model
  - Fit the data parsimoniously (not too many parameters, etc)
- Must be interpretable
  - Model parameters should map onto something in cognition
  - Model must relate to some psychological idea

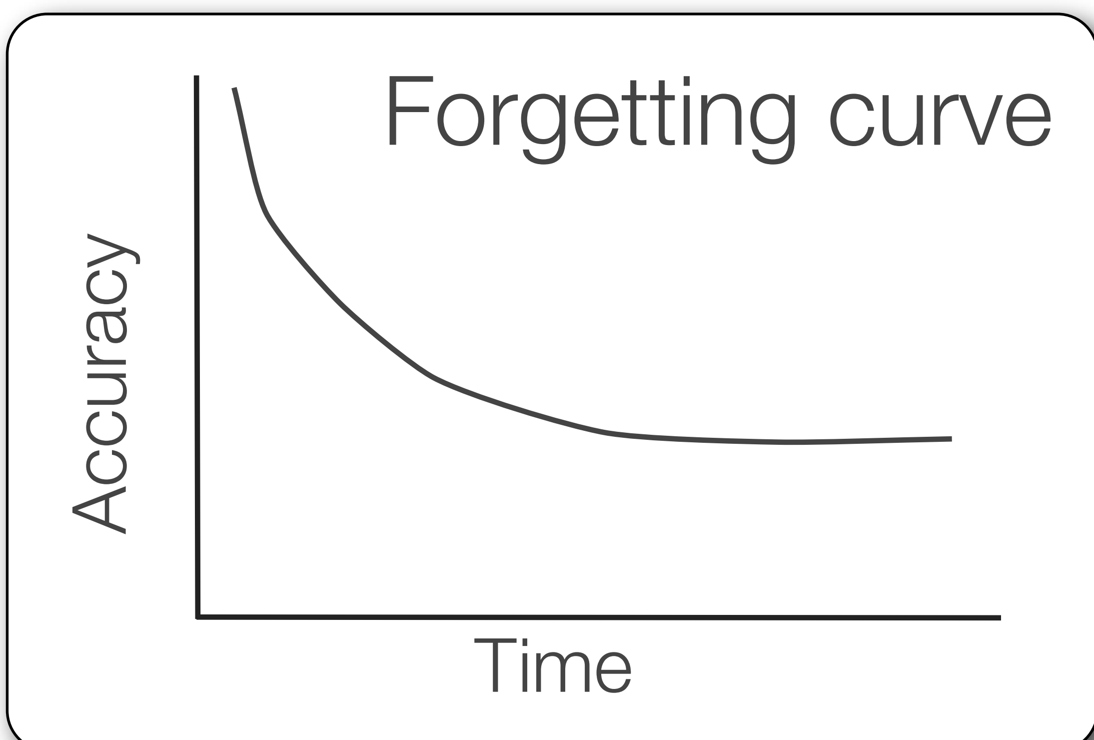

$$R = e^{-t/S}$$

R = retention
t = time
S = strength of memory

Desiderata

- Must be good statistical model
  - Fit the data parsimoniously (not too many parameters, etc)
- Must be interpretable
  - Model parameters should map onto something in cognition
  - Model must relate to some psychological idea
- Should teach us something
  - Something about the human mind should become clearer by looking at the model's behaviour...
  - The point is to learn about the mind, not make a better computer or show of how smart we are

Desiderata

How do we know if a model has taught us something?

What counts as "something?"

- Should teach us something
  - Something about the human mind should become clearer by looking at the model's behaviour...
  - The point is to learn about the mind, not make a better computer or show of how smart we are

Different levels of explanation (Marr)

- Three levels of abstraction at which different models describe and explain cognition

Computational level: What are the abstract problems that the mind needs to solve, and what would a solution look like?

Algorithmic level: What informational processing steps are followed to arrive at the solutions?

Implementation level: How does the brain carry out these operations?

Marr's Levels: Example (Memory)

- Implementation: How does the brain carry it out?

Memory formation has a clear biological mechanism

- Connections between neurons are strengthened by a process that is mediated by a neurotransmitter called NMDA

- Inhibiting NMDA operation prevents the formation of new memories (in mice, at least)

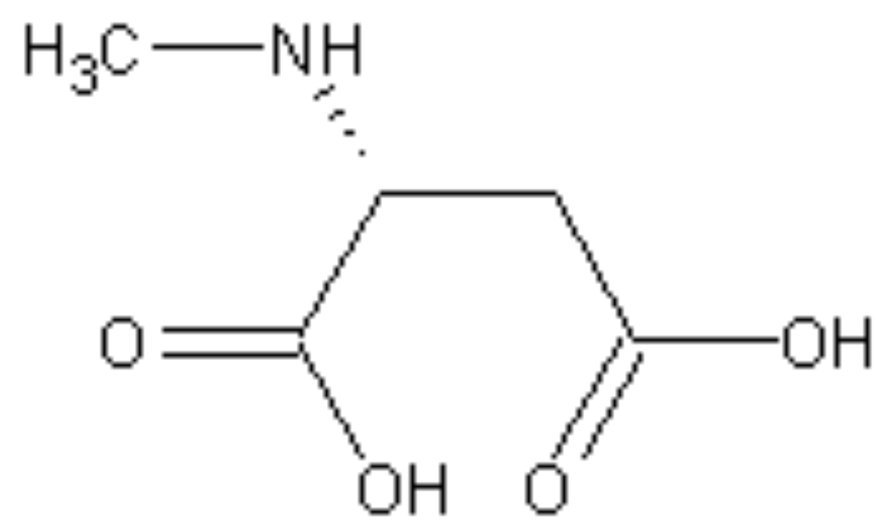

N-methyl D-aspartate

Marr’s Levels: Example (Memory)

- Algorithmic: What steps are carried out?

Memories get systematically edited over time, without us noticing

> One day after the Challenger explosion: “**I was in my religion class** and some people walked in and started talking about [it]. I didn’t know any details except that it had exploded and the schoolteacher’s students had all been watching, which I thought was so sad. **Then I went to my room** and watched the TV program talking about it, and I got all the details from that.”

> 2.5 years later, same person: “When I first heard about the explosion **I was sitting in my freshman dorm** room with my roommate, and we were watching TV. It came on a newsflash, and we were both totally shocked. I was really upset, and **I went upstairs to talk to a friend** of mine, and then I called my parents.”

Marr's Levels: Example (Memory)

- Computational: What is the abstract problem being solved?

The probability that people recall a particular piece of information is closely matched to the probability that it is needed by the person.
(Anderson & Schooler, 1991)

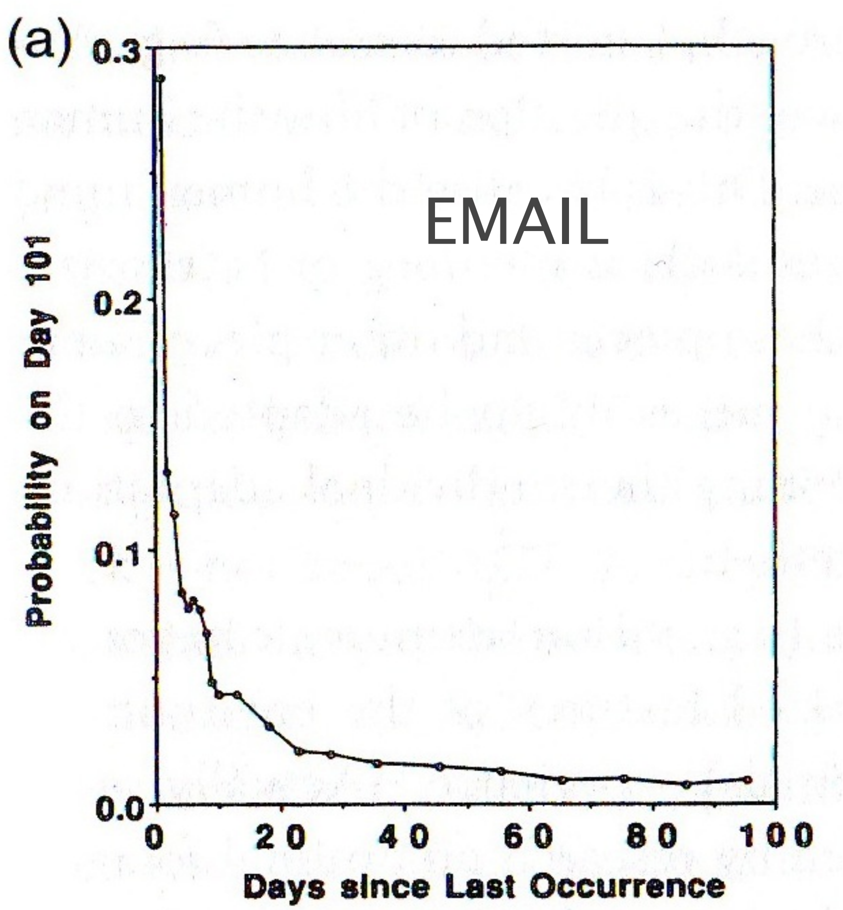

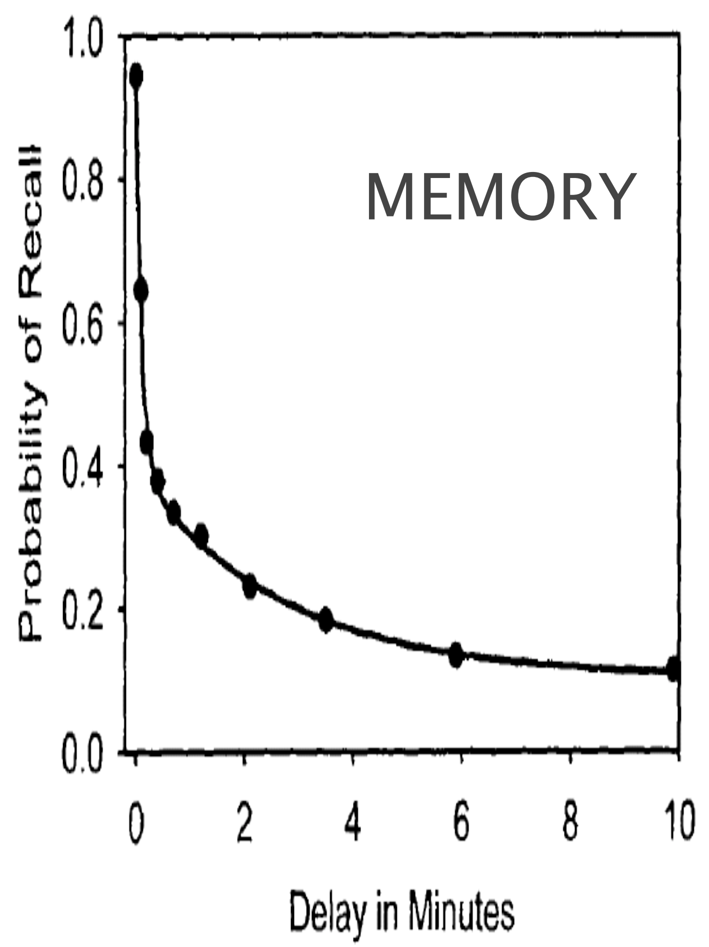

The point

- A full understanding of memory includes an explanation at all three levels
- Any single explanation in cognitive science generally occurs on one (or at most two)
- Most of the models we discuss will be on the highest, "computational" level

Our question: Do they illuminate the *nature of the problem* facing a human and what a *solution to that problem* would look like?

Why bother with modelling at all?

- Formal models are harder to build than verbally specified theories. Why go to all of that work?

- It's fun! Also several other reasons
  - Precision: Solves the horrible "no, no, what I meant to say was X, not Y" problems that plague verbal theories
  - Openness: Formal models, once specified, are usable by anyone with training - not stuck in the head of the person who invented it
  - Counterintuitive predictions: Often, when you run the model, it makes different predictions than what you thought. This is the best way to learn!
  - Interoperability: Speak the "same" language across disciplines: AI, machine learning, neuroscience, cognitive science

Plan for the semester

- Introduction
  - Basic Bayesian inference
- What can (and do) people learn?
  - Inductive generalisations; simple categorisation and applications; more complex categorisation; structure in time
- What kind of information do people get?
  - Different kinds of sampling; learning from teachers; informational utility
- How do people find and use that information?
  - Sensitivity to sampling and utility; cognitive biases; stopping decisions
- How do people act on that information?
  - Explore/exploit dilemmas, rewards, sensitivity to environment
- Bonus! (computational statistics)

Additional references (not required)

Levels of explanation

- Marr, D. (1982). *Vision: A computational investigation into the human representation and processing of visual information.* Freeman. Ch 1 & 2.

Modelling in cognitive science

- McClelland, J. (2009). The place of modelling in cognitive science. *Topics in Cognitive Science 1*: 11-38

- McClelland, J., Botvinick, M., Noelle, D., Plaut, D., Rogers, T., Seidenberg, M., & Smith, L. (2010). Letting structure emerge: Connectionist and dynamical systems approaches to cognition. *Trends in Cognitive Sciences 14*(8): 348-356.

- Griffiths, T., Chater, N., Kemp, C., Perfors, A., & Tenenbaum, J. (2010). Probabilistic models of cognition: Exploring representations and inductive biases. *Trends in Cognitive Sciences 14*(8): 357-364.