Principles of Cognition

Lecture 1: The Cognitive Science Approach

Paula Parpart
Experimental Psychology

a little bit about myself..

Computational approaches to cognition (e.g., "Bayesian models") ⟷ Heuristics in Judgement and Decision Making

Medical Decision Making

Active Learning: How do people learn/acquire information?

COLLABORATORS:
The Berliners

Gerd Gigerenzer (MPI)

Leal Schooler (Syracuse)

Edward Cokeley (Oklahoma)

Mirta Galesic (Santa Fe)

Max Planck Institute for Human Development

COLLABORATORS:
The London crew

Maarten Speekenbrink

Eric Schulz (LJDM)

Takao Noguchi

Prof Brad Love

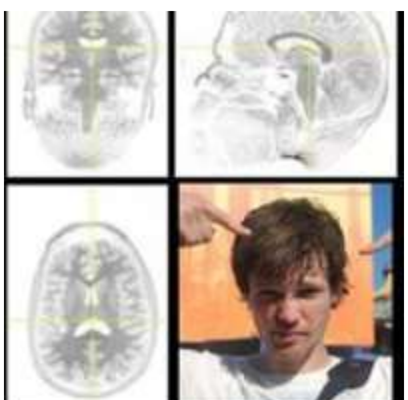

Neil Bramley

Problems, problems, problems….

- Field is at early stages
- Only few established “Principles of Cognition”
- Lack of agreement on a number of central issues – active debates
- Multidisciplinary nature of the field
  - You all have different backgrounds!

What is cognitive science?

- According to MIT Press' exhaustive *A Companion to Cognitive Science*: "Cognitive science is the multidisciplinary scientific study of cognition and its role in intelligent agency. It examines what cognition is, what it does, and how it works."

-William Bechtel and George Graham

What is cognitive science?

- According to MIT Press' exhaustive *A Companion to Cognitive Science*: "Cognitive science is the multidisciplinary scientific study of cognition and its role in intelligent agency. It examines what cognition is, what it does, and how it works."

-William Bechtel and George Graham

Multidisciplinary: Crossing traditional disciplinary boundaries or using the methods of more than one area of study.

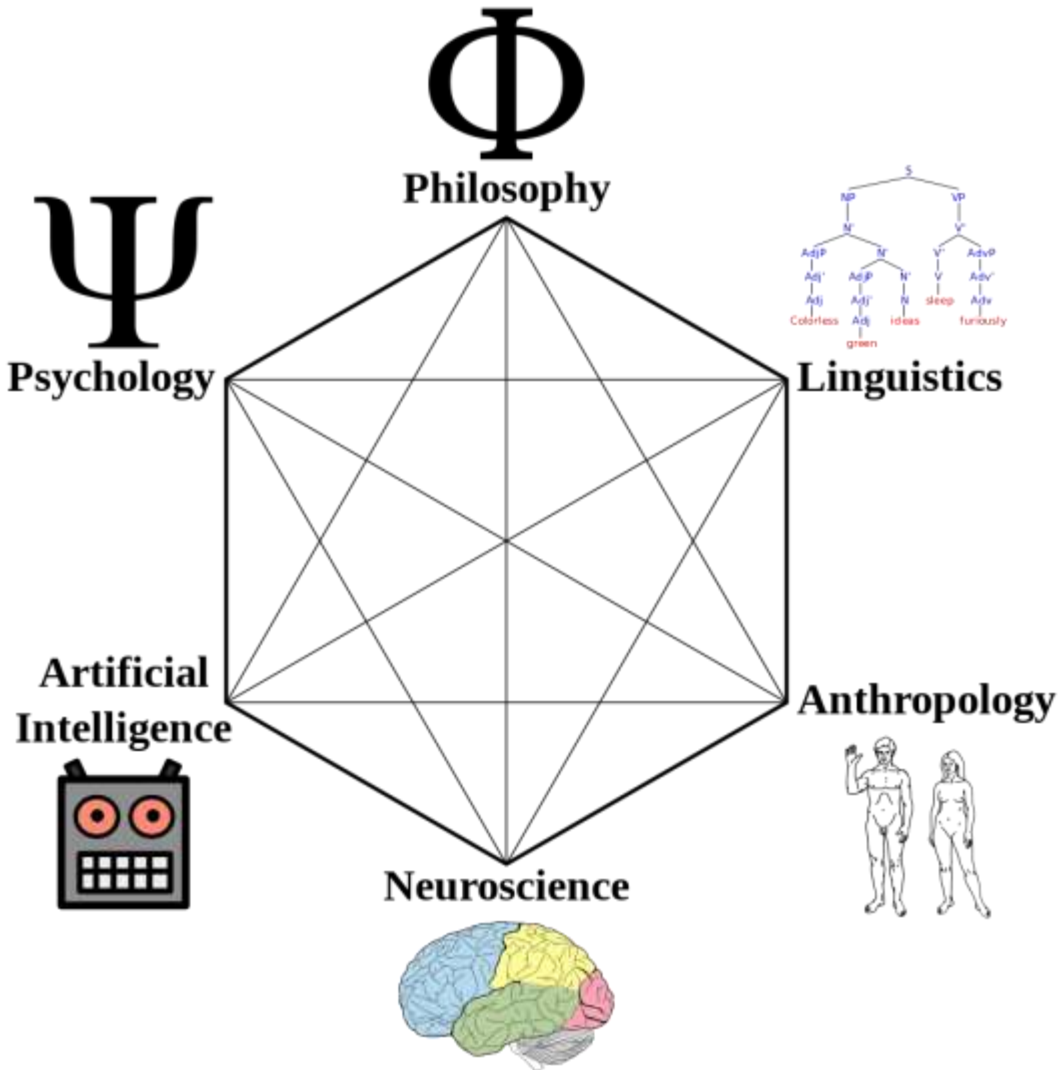
Philosophy
Psychology
Linguistics
Artificial Intelligence
Anthropology
Neuroscience

Approach

- introductory **lecture** + course outline
- Then:
  – Every week a different expert will discuss:

A different subject
- Perception
- Mental Representation
- Language
- Decision Making
- Embodied Cognition
- Social Cognition

and/ or

A (somewhat) different approach
- Experimental Psychology
- Neuroscience
- Computational modelling

COURSE OVERVIEW

- **03/10/2016 – Paula Parpart: The Cognitive Science Approach**
- 10/10/2016 – Sam Schwartzkopf: Scientific reasoning & Perceptual Bias
- 17/10/2016 – David Shanks: Mental construction
- 24/10/2016 – Constantin Rezlescu**:** Face Perception
- 31/10/2016 – Lasana Harris**:** Social cognition

-- READING WEEK –

- 14/11/2016 – Jeremy Skipper: Language cognition
- 21/11/2016 – David Vinson**:** Embodied cognition
- 28/11/2016 – Tali Sharot: Positivity Bias in Belief Formation
- 05/12/2016 – Peter Dayan: Computational Decision and Learning models
- 12/12/2016 – David Tuckett: The Role of Feelings in Judgement and Decision Making

Available on Moodle + Readings + Info

COURSE OVERVIEW

- **03/10/2016 – The Cognitive Science Approach**

Introduction to Cognitive Science

- 10/10/2016 – Sam Schwartzkopf: **Perceptual** Biases
- 17/10/2016 – David Shanks: Mental **Construction**
- 24/10/2016 – Constantin Rezlescu**:** Face **Perception**

Perception

- 31/10/2016 – Lasana Harris**: Social Cognition**

Social Cognition

-- READING WEEK –

- 14/11/2016 – Jeremy Skipper: **Language** Cognition

Language Cognition

- 21/11/2016 – David Vinson**: Embodied** Cognition

Embodied Cognition

- 28/11/2016 – Tali Sharot: Positivity Bias in **Belief** Forma
- 05/12/2016 – Peter Dayan: Computational **Decision**/ Learning **models**
- 12/12/2016 – David Tuckett: Feelings in **JDM**

Decision Making

UCL

- **The Cognitive Science Approach**

- 10/10/2016 –**Perceptual** Biases
- 17/10/2016 - Mental **Construction**
- 24/10/2016 – Face **Perception**

- 31/10/2016 – **Social Cognition**

-- READING WEEK –

- 14/11/2016 – **Language** Cognition
- 21/11/2016 – **Embodied** Cognition

- 28/11/2016 – Positivity Bias in **Belief** Formation
- 05/12/2016 – Computational **Decision**/ Learning **models**
- 12/12/2016 – Feelings in **JDM**

Approaches and Methods

1 **Perception**

2 **Social Cognition**

3 **Language**

4 **Embodied Cognition**

5 **Decision Making**

1. Perceptual Biases

1. The mind
is able to
construct things…
HOW?

1. Face **perception**/
recognition

2. **Social** Cognition
Facial Feedback → Emotions, Attribution, Attitudes, Affect and the Self

5. Decision Making

4. Embodied Cognition
Embodied account of language comprehension, abstract concepts, emotions

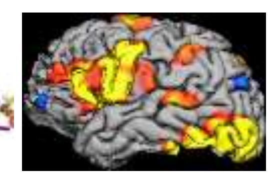

3. **Language** Cognition, multimodal language incl. gestures, neuroscientific basis of language

Objectives

- Give you a broad overview of the field & understanding of big theoretical questions and challenges
- Introduce you to (some of) the active research areas, experimental methods and computational approaches.
- Orient you towards the subject(s) more suited to your interests
- Get you thinking about your final project

Logistics

- Readings for all lectures are already uploaded on Moodle.
  - Read at least one of each before the lecture
- Lecture slides will be on Moodle before or slightly after each lecture.
- Lectures are recorded – recordings will be appearing on Moodle 2-3 days after each lecture
- Assessment: 100% coursework
  - Essay topics will be announced soon on Moodle
  - Deadline: 30/01/2017
- For any lecture related question email the lecturer
- For any general issue about the module email Paula (paula.parpart@ucl.ac.uk)

TOPIC OVERVIEW

1. **Cognitive science as reverse engineering**
2. **Historical background**
3. **Methods in Cognitive Science**
4. **Levels of Explanation**

Note: Thanks to Nick Chater (WBS) and Brad Love (UCL) for sharing some of their introductory materials for 1. and 2.

1. COGNITIVE SCIENCE AS REVERSE ENGINEERING

COGNITIVE SCIENCE vs ARTIFICIAL INTELLIGENCE

- Cognitive science: towards computational models of human intelligence

- Artificial intelligence: aims to build computer systems to do things normally requiring intelligence in humans

BUT THE TWO ARE CLOSELY RELATED

- The biggest single empirical constraint on any model of a cognitive process is that it solves the task successfully
- And without powerful computational tools, it is impossible even to understand the problems the brain solves
  - Cognitive science needs engineering (Artificial intelligence)
- The *only* working example of an intelligent learning system is the brain
- And many engineering problems are defined in terms of human cognition (from object recognition, to machine translation)
  - Engineering needs cognitive science (cf Biomimetics)

burrs  → 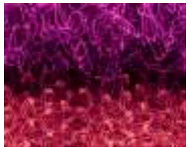 velcro

2. HISTORICAL BACKGROUND:

INTRODUCTION, BEHAVIORISM, THE COMPUTER METAPHOR

WUNDT AND THE BEGINNING OF EXPERIMENTAL PSYCHOLOGY

- Founded first experimental psychology laboratory Leipzig, in 1879

- **Structuralism**: “Human mental experience, no matter how complex, can be viewed as blends or combinations of simple processes or elements.”

- But rather than **computational components**, building blocks are **subjective experience** (qualia)

PIONEERS OF PSYCHOLOGY

Wilhelm Wundt (1832–1920). *[Archives of the History of American Psychology, University of Akron].*

THE INTROSPECTIVE METHOD

- **Experimental Psychology** "…the investigation of conscious processes in the modes of connection peculiar to them"
- **Method** – Systematic introspection, under experimental control, replicability

THE METHOD IN ACTION

- very simple stimuli → verbal report

O listens to a metronome. After a time the beats form rhythmic groupings and various conscious experiences may be reported, such as, at the end of a group there is an impression of an "agreeable whole". He then tries to describe the qualities of this experience, such as feelings of pleasure or displeasure, tension or excitement

- **Attempt** to isolate the "elements of consciousness" out of which more complex mental events are made.

- **Metronome**: single beat = a sensation
  Combination into rhythms = an idea.

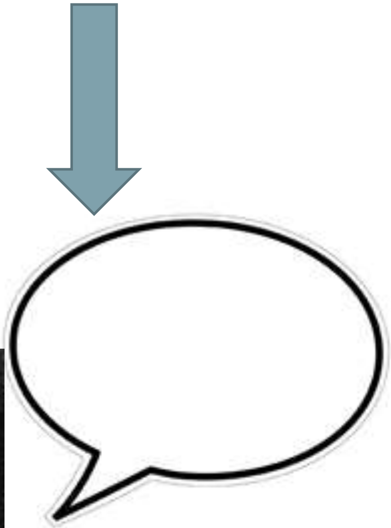

TROUBLE FOR INTROSPECTION

- the imageless thought controversy: Wundt vs Külpe
- Different labs produced very different results
  - e.g., Leipzig vs. Cornell
- Introspection can change the phenomenon observed

THE UNRELIABILITY OF INTROSPECTION IN PERCEPTION

- How much can we introspect about the retinal image?

- Endless visual illusions

THE UNRELIABILITY OF INTROSPECTION IN PERCEPTION

- How much do can we introspect about the retinal image?
- Endless visual illusions

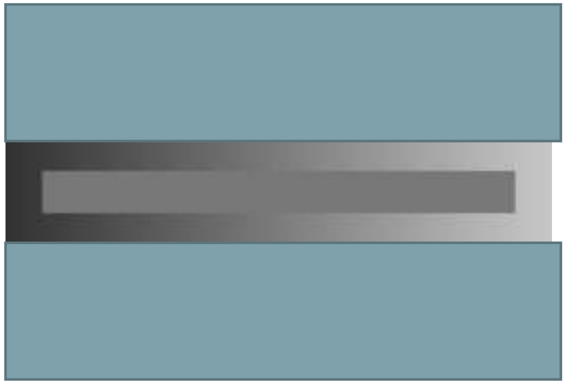

THE UNRELIABILITY OF INTROSPECTION IN PERCEPTION

- How much do can we introspect about the retinal image?
- Endless visual illusions

LITTLE BETTER FOR KNOWLEDGE OR DECISION

- People often can't explain their behaviour (Johansson & Hall's choice blindness, e.g., Science 2005)

BEHAVIORISM

- As a reaction to the subjectivity of introspection: Perhaps psychology is not about subjective experience but objective behaviour?

- Emergence of behaviorism: very strong claim that psychological laws should be framed over direct relationships between physically characterised stimuli **(S)** and/or responses **(R)**

- Main method: looking at animal learning: where stimuli and learning can be carefully **controlled** and **measured**

SETTING THE STAGE: PAVLOV (1849 – 1936)

Physiologist studying dog digestion, found, by chance, that dog salivation was triggered by a bell that usually preceded food

Thus, the dog had learned an CS-US (bell-food) association

Classical or Pavlovian conditioning

J.B. WATSON (1878 – 1959)
FOUNDER OF BEHAVIORISM

- **Watson** dismissed introspection as hopelessly unscientific. To be replaced with:
  - Psychology should restrict itself to examining the relation between **observable stimuli** and observable behavioural responses.
  - **Explained** via Stimulus-Stimulus and Stimulus-Response links
  - **Language** is movement of the larynx
  - **Thought** is movement of the larynx (hidden behaviour)

Psychology as the Science of Behavior

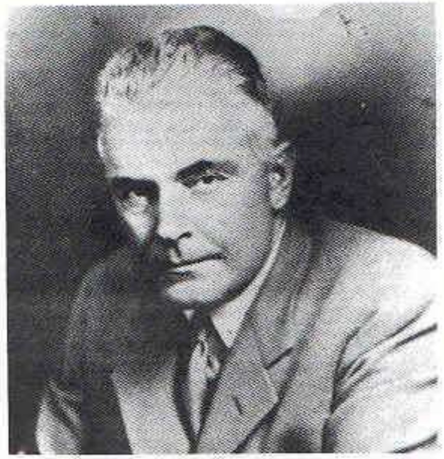

John Broadus Watson (1878–1958). *(Archives of the History of American Psychology, University of Akron).*

J.B. WATSON (1878 – 1959)
FOUNDER OF BEHAVIORISM

- Psychology is the science of behaviour.
- Psychology is not the science of mind.
- Behaviour can be described and explained without making ultimate reference to mental events or to internal psychological processes.

Psychology as the Science of Behavior

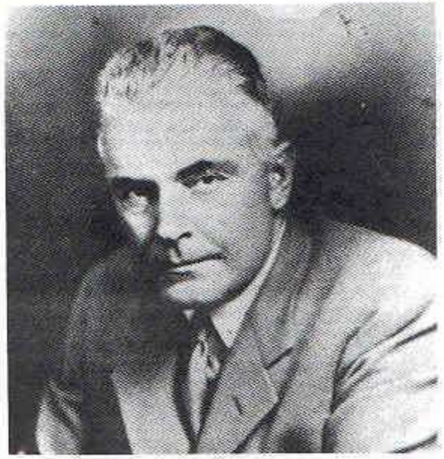

John Broadus Watson (1878–1958). *(Archives of the History of American Psychology, University of Akron).*

B. F. SKINNER (1904 – 1990)
RADICAL BEHAVIORISM

Further developed behavorist research, working on **operant or Skinnerian conditioning** (building relationships between Responses and Rewards/Punishments)

Produced a vast research programme on learning in pigeons in the "Skinner box"

Schedules of food reward determined by, e.g., lever pressing

Starting point for ideas of reinforcement learning in neuroscience and machine learning

PIONEERS OF PSYCHOLOGY

B. F. Skinner with one of his subjects. *[Ken Heyman].*

AND IT *WAS* RADICAL

Aimed to explain *all* human behaviour, including language

No real theory of perception or motor control

Typically no attempt to link with the brain

Reinforcement history explicitly viewed as the correct alternative to our view of ourselves as reasoning beings.

Only innate structure is principles of association

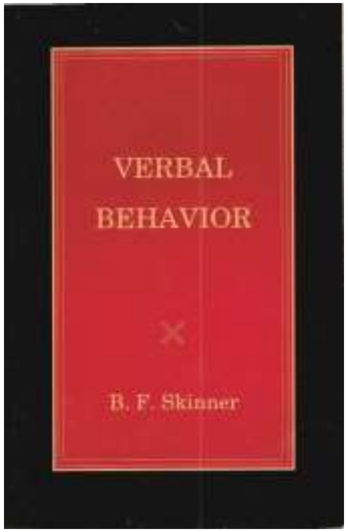

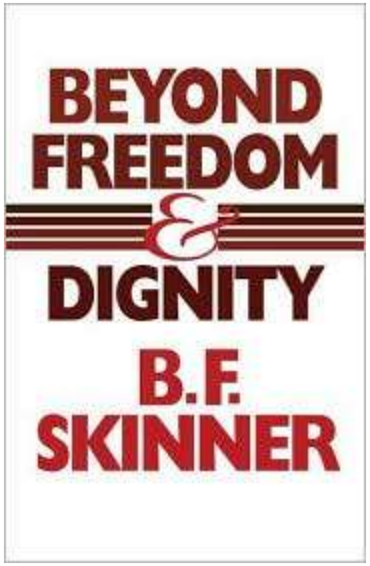

PROBLEMS FOR BEHAVIORISM

- Perception and motor control: Associations between categories "lever," "press," - but this is circular – these categories must be explained.

So: circularity

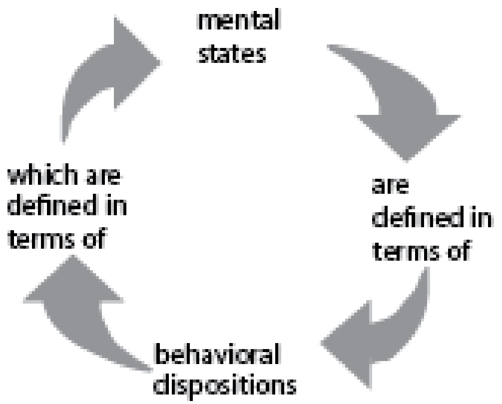

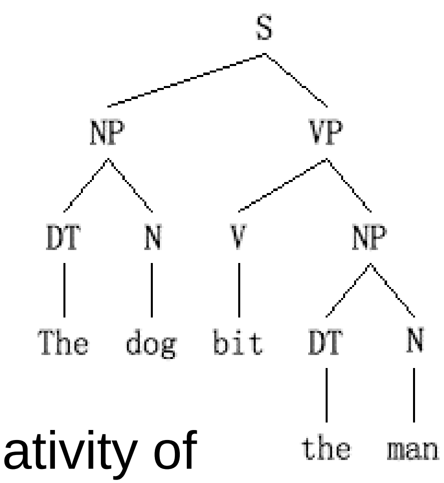

- Language : Chomsky showed that the infinite creativity of language cannot be explained in S-R terms
  (we learn rules for language, not S-R associations)

PROBLEMS FOR BEHAVIORISM

- Chomsky: Language comprehension with mental grammars consisting of rules.

- Behaviorist models cannot explain the rapid acquisition of language by young children, i.e., the phenomenon of “lexical explosion.”

- When put to the test of uttering a grammatical sentence, a person has a virtually infinite number of possible responses available, and the only way in which to understand this virtually infinite generative capacity is to suppose that a person possesses a powerful and abstract innate grammar.

- It appears to be a fundamental fact about human beings that our behavioral capacities often surpass the limitations of individual reinforcement histories.

- **Flexibility of behaviour** i.e., behaviour guided by reasoning to the solution of a novel problem
- **Example** Rats tend to take the shortest route through a maze rather than the one that has been most reinforced **(Hull and Tolman)**

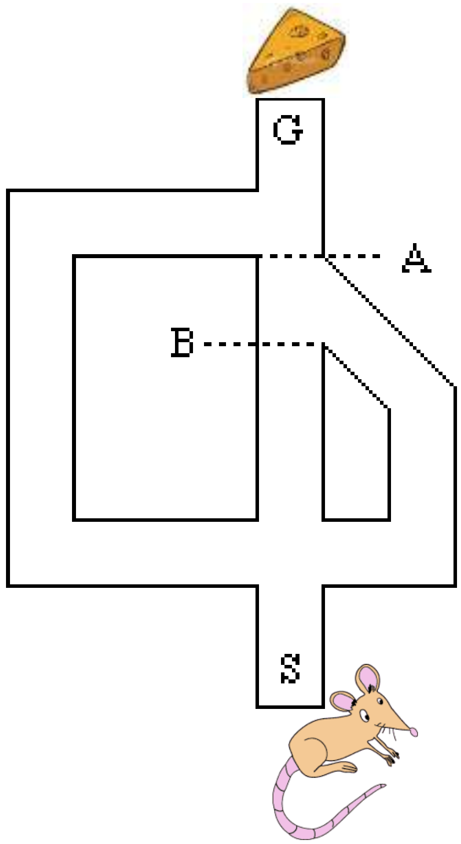

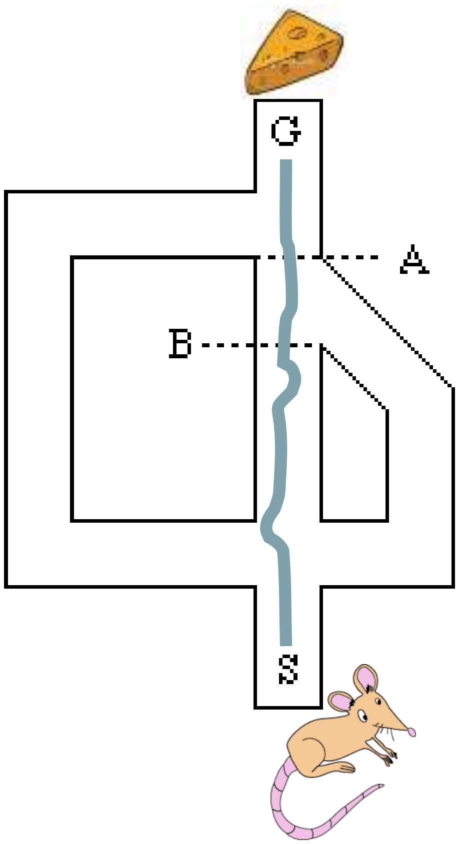
G
A
B
S

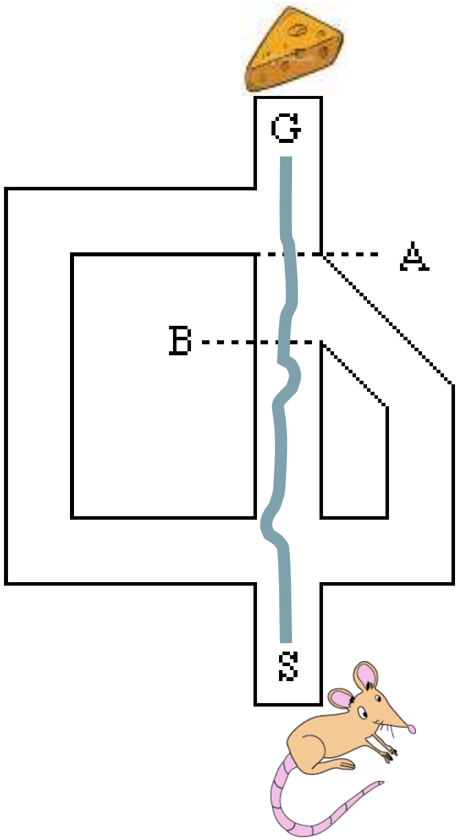
G
A
B
S

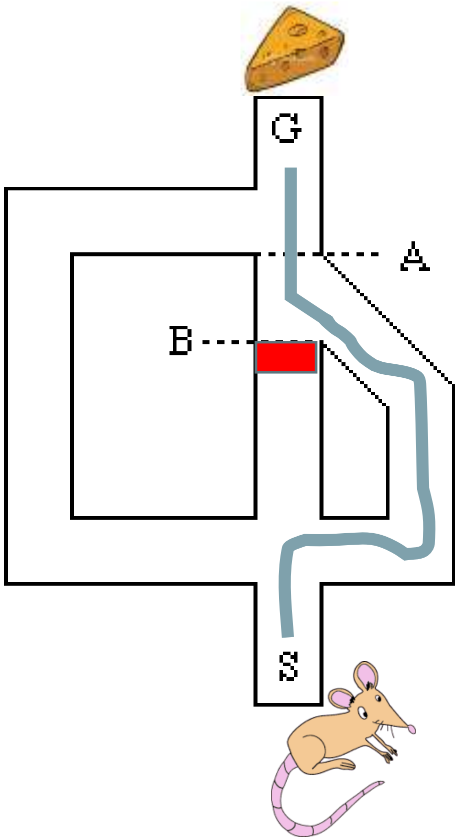
G
A
B
S

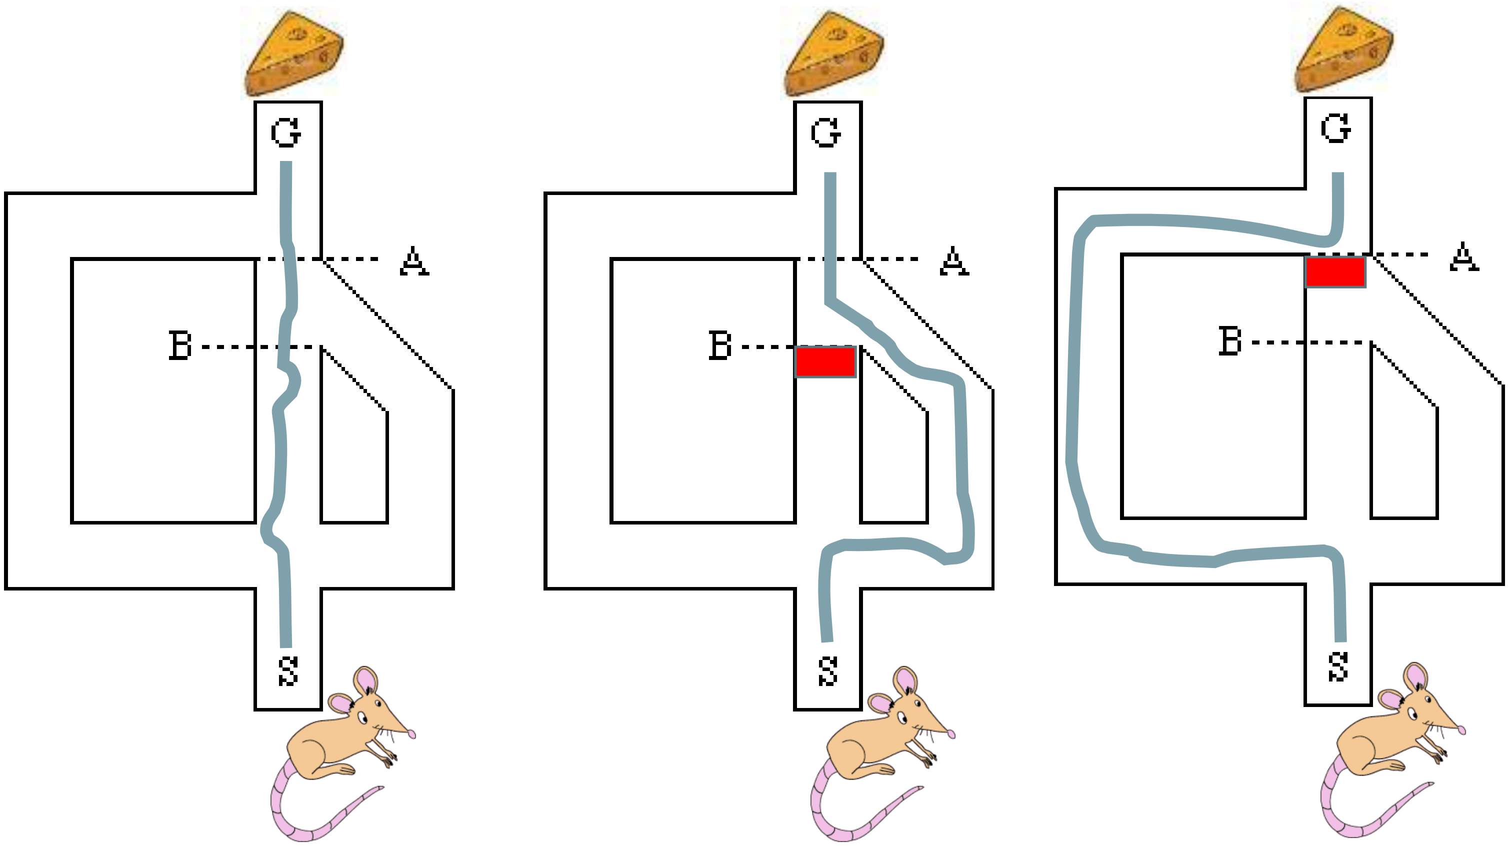
G
A
B
S
G
A
B
S
G
A
B
S

- **Flexibility of behaviour** i.e., behaviour guided by reasoning to the solution of a novel problem
- **Example** Rats tend to take the shortest route through a maze rather than the one that has been most reinforced **(Hull and Tolman)**

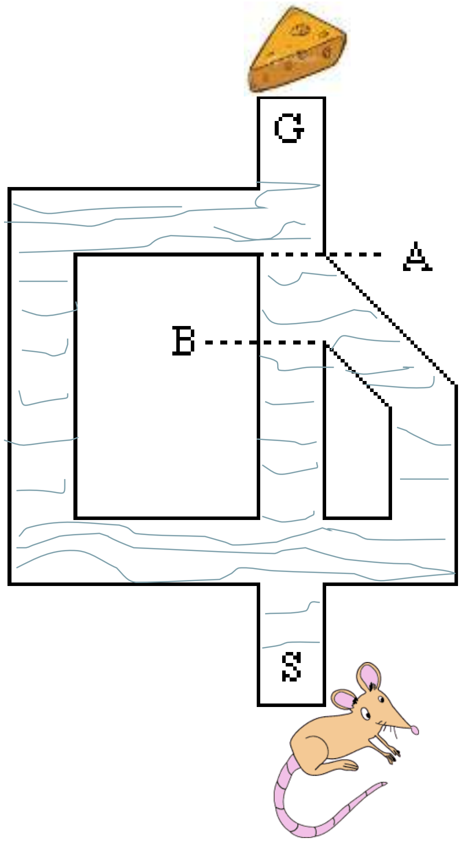

AND THE RAT CAN NAVIGATE AFTER A MAZE IS FLOODED – S-R LINKS ARE RADICALLY DIFFERENT


THE FAR SIDE
By GARY LARSON
8-28
© 1986 FarWorks, Inc./Dist. by Universal Press Syndicate
"Stimulus, response! Stimulus, response! Don't you ever *think*?"

SUGGESTS THE NEED TO PEER INSIDE THE BLACK BOX

COMPUTATION AS A FRAMEWORK

SUGGESTS THE NEED TO PEER INSIDE THE BLACK BOX

- Rats with maps
- Abstract goals
- Human Language
- Beliefs about causality
- Other minds
- Memory

COMPUTATION AS A FRAMEWORK

Internal models

Memory

- **Cognitive Science** gradually emerged through the late 40s, 50s, and 60s in the work of Kenneth Craik, George Miller, Jerome Bruner, Herbert Simon, Alan Newell, Noam Chomsky...

Perception /education

ion as computation

Allowed the mind back in to mediate between S and R, in causing intelligent behaviour

THE COGNITIVE REVOLUTION

- **Cognitive Science** gradually emerged through the late 40s, 50s, and 60s in the work of Kenneth Craik, George Miller, Jerome Bruner, Herbert Simon, Alan Newell, Noam Chomsky...

- At this time, primitive computers had been around for only a few years, but pioneers such as John McCarthy, Marvin Minsky, Allen Newell, and Herbert Simon were founding the field of artificial intelligence

Information Processing Assumption

"The mind is a complex system that receives, stores, retrieves, transforms and transmits information. These operations are called computations or information processes."

-Stillman

Brain as Computer

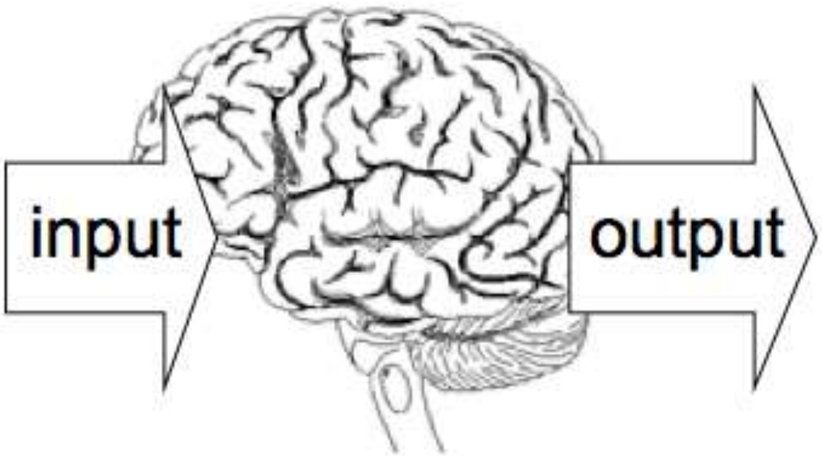

Brains vs. Computers

- 1000 operations/sec
- 100,000,000,000 units
- 10,000 connections/
- graded, stochastic
- embodied
- fault tolerant
- evolves
- learns

- 1,000,000,000 ops/sec
- 1-100 processors
- ~ 4 connections
- binary, deterministic
- abstract, disembodied
- crashes frequently
- explicitly designed
- is programmed

Brain as Computer

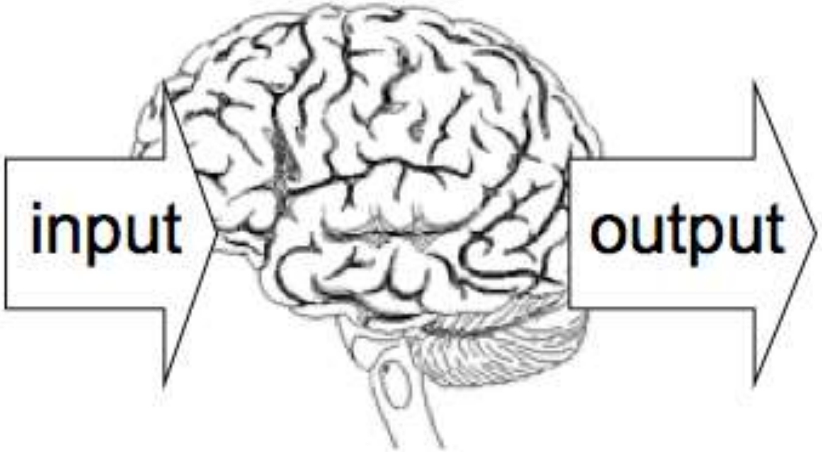

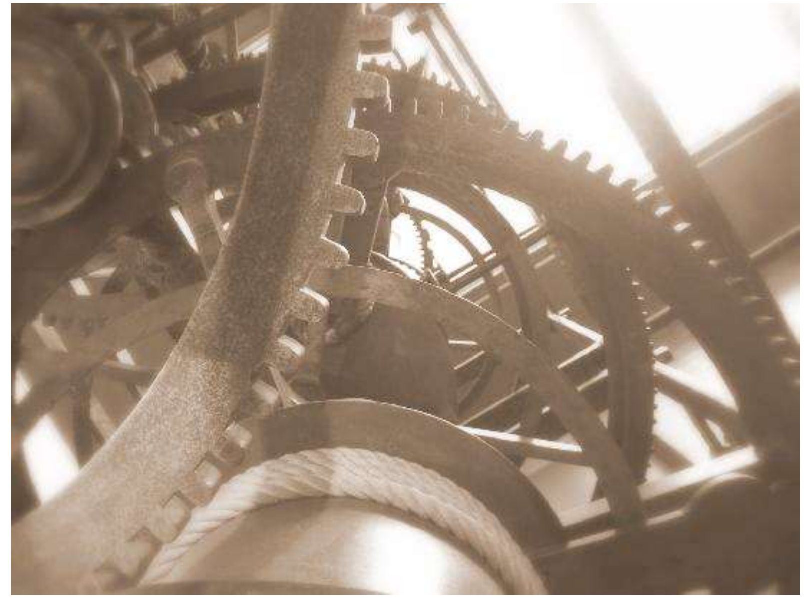

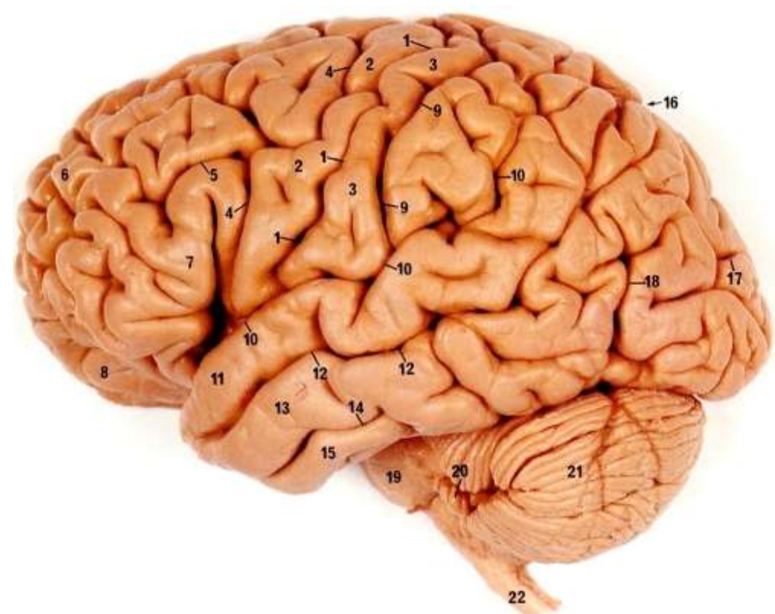

ition

THE DAWN OF COMPUTERS

- Alonzo Church (1936 thesis): everything that can be computed can be computed with recursive functions
- Alan Turing (same time): Turing machine: An abstract machine capable of calculating all recursive functions -> a machine that can compute anything.

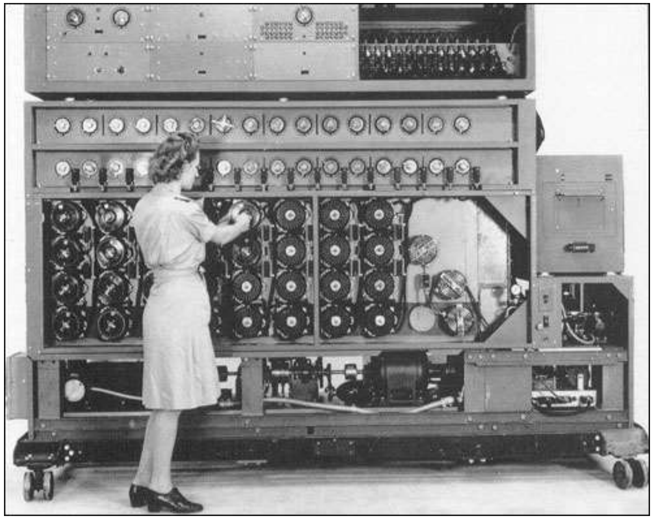

(Turing's) **Definition of computation**

- "A function is said to be computable if it can be implemented on a Turing Machine."
- **Roughly speaking**, a function or task is computable if its solution can be found in "finite" time.
- A problem in which the time required to solve grows exponentially as the problem size grows may be practically uncomputable (i.e., unsolvable) → NP-hard problem
  - (e.g., Traveling Salesman Problem)

THE DAWN OF COMPUTERS

- Alonzo Church (1936 thesis): everything that can be computed can be computed with recursive functions
- Alan Turing (same time): Turing machine: An abstract machine capable of calculating all recursive functions -> a machine that can compute anything.
- The first machines: early 1940s
- McCulloch and Pitts (1943): "A Logical Calculus of the Ideas Immanent in Nervous Activity": Neuron-binary digit analogy

Natural versus Artificial Neuron

- Natural Neuron McCullough Pitts Neuron

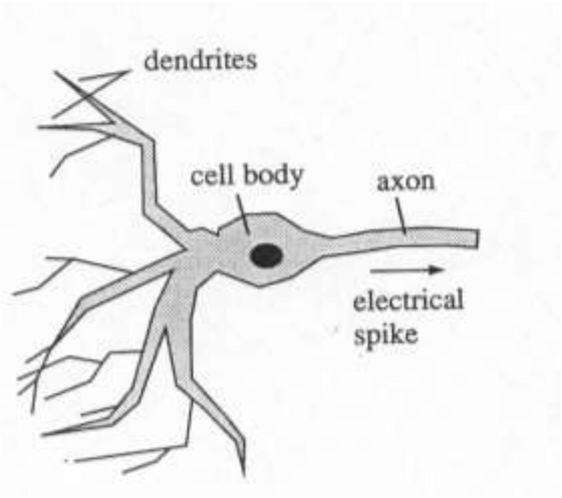

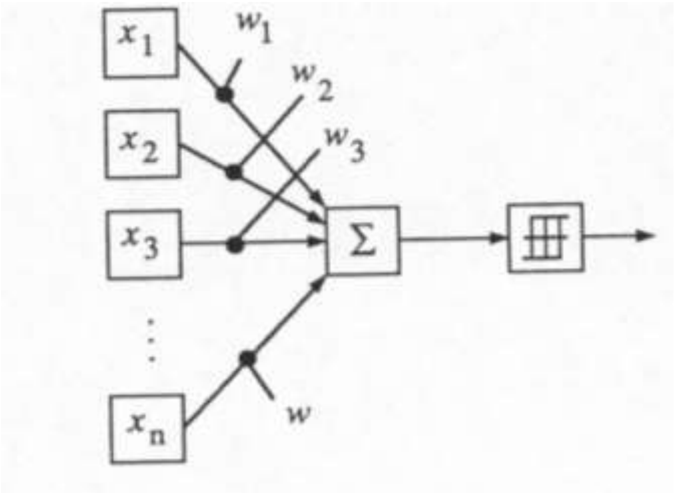

Representability

- What functions can be represented by a network of Mccullough-Pitts neurons?

- Theorem: Every logic function of an arbitrary number of variables can be represented by a three level network of neurons.
  - Using logical functions and, or and not.

THE BIRTH OF COGNITIVE SCIENCE

- The first AI conference (1956): Dartmouth College
  - Newell & Simon: The first computer programme: The Logic Theorist
  - “Logic Theory Machine” (1956): "In this paper we describe a complex information processing system, which we call the logic theory machine, that is capable of discovering proofs for theorems in symbolic logic. “
  - 1st draft of Marvin Minsky's "Steps toward AI"

THE BIRTH OF COGNITIVE SCIENCE

- Concensusal birthday: Symposium on Information Theory at MIT in 1956

  (Revolution against behaviourism)

THEME: *Is cognition 'information processing' (data+ algorithms)?*

- Newell & Simon (AI)
  The first computer program
- McCarthy, Minsky (AI)
  Modelling intelligence
- Miller (Experimental psychology)
  "Human Memory and the Storage of Information": magic number 7
- Chomsky (Linguistics)
  Transformational grammar

SYMBOLIC KNOWLEDGE REPRESENTATION

- Knowledge organized in semantic networks (Collins & Quillian, 1969)
- Or other logic-based representations, frames, scripts, situation calculus, etc.

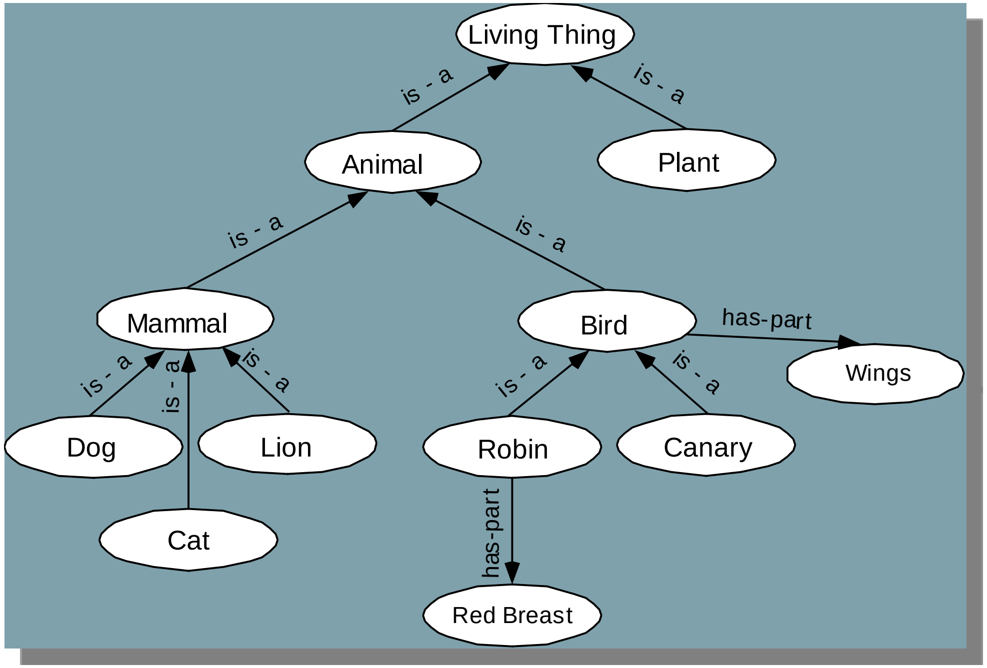

PROBLEM SOLVING AS SYMBOLIC SEARCH

- Problem solving as search in a symbolic problem space (Newell & Simon; here in the Towers of Hanoi problem)

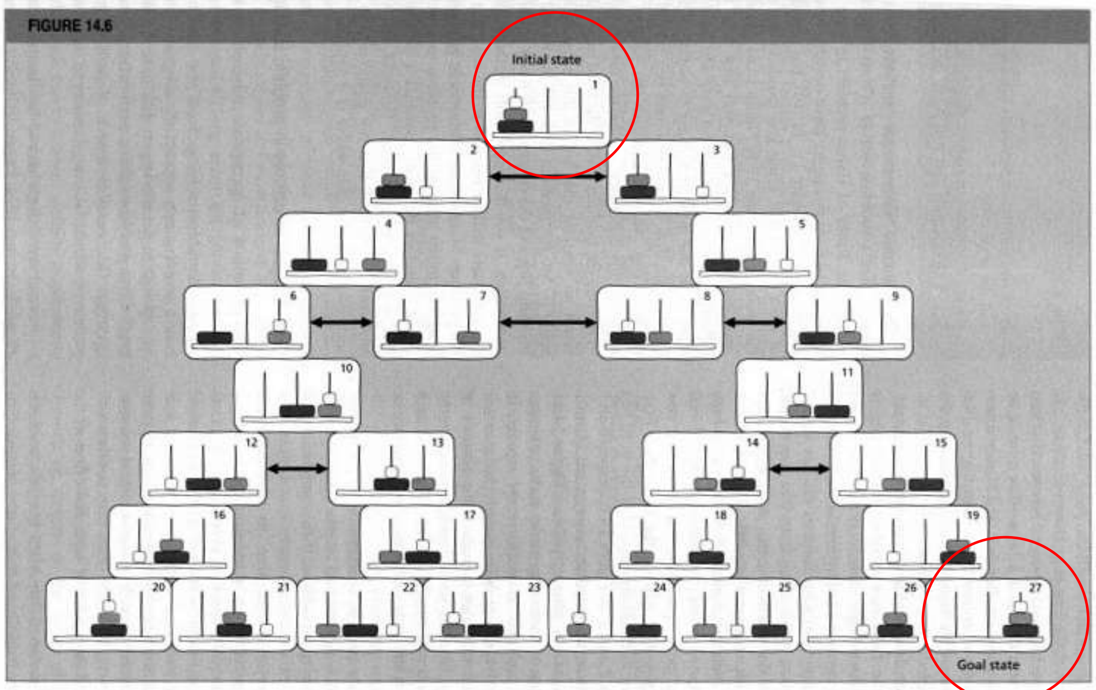

The problem space of legal moves in the Tower of Hanoi problem. If boxes touch each other, or are joined by arrows, this indicates that one can move from one state to the other using a legal operator.

LANGUAGE PROCESSING AS SYMBOL MANIPULATION

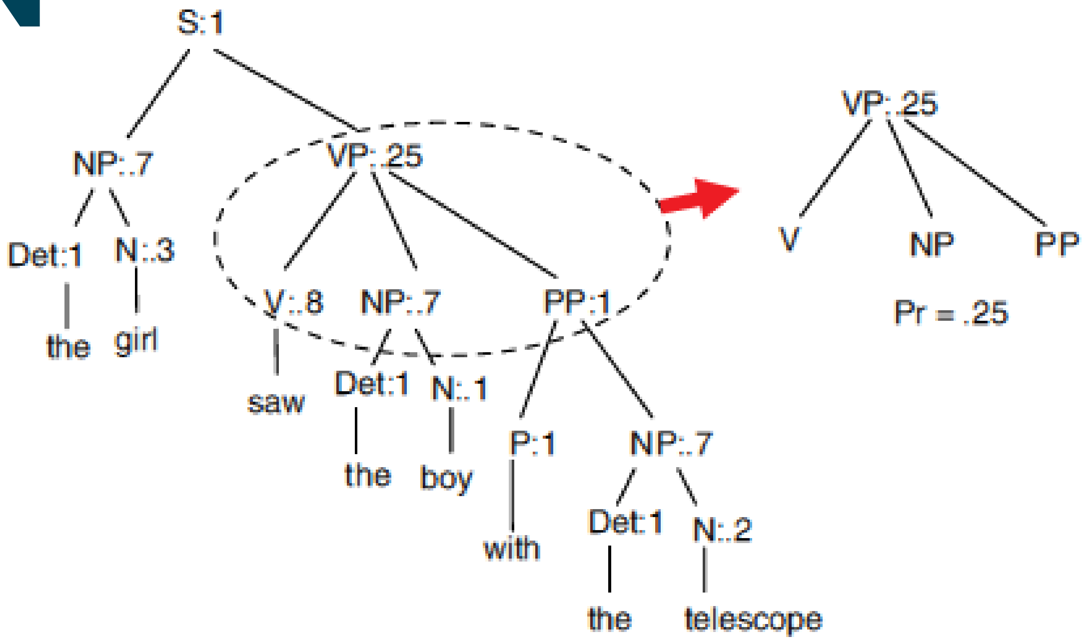

(The girl) (saw) (the boy with the telescope)

vs

(The girl) (saw) (the boy) (with the telescope)

Two challenges for the symbolic approach

Challenge of Learning

- Knowledge typically hand-coded
- Grammar specified by the linguist
- Inference via logic (or equivalent)—but does not go *beyond the information given*

- Parallel computation
- How to implement symbolic representations?

Challenge of Neural Plausibility

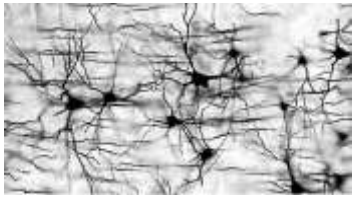

3. METHODS IN COGNITIVE SCIENCE

Various Methods

- Observation
- Experiment
- Modelling
- Brain Imaging
- Combined Approaches
  - E.g., modelling, experiment, neuroscience measures.

Various Methods

1. Experiment

   people, often undergraduates satisfying course requirements, are brought into the laboratory so that different kinds of thinking can be studied under controlled conditions. → **Amazon Mechanical Turk**

- Although theory without experiment is empty, experiment without theory is blind.

- One of the best ways of developing theoretical frameworks is by forming and testing computational models intended to be analogous to mental operations.

Various Methods

2. Modelling

To complement psychological experiments on deductive reasoning, concept formation, mental imagery, and analogical problem solving, researchers have developed computational models that simulate aspects of human performance.

- Designing, building, and experimenting with computational models is the central method of artificial intelligence (AI)

- Ideally in cognitive science, computational models and psychological experimentation go hand in hand…But..

Various Methods

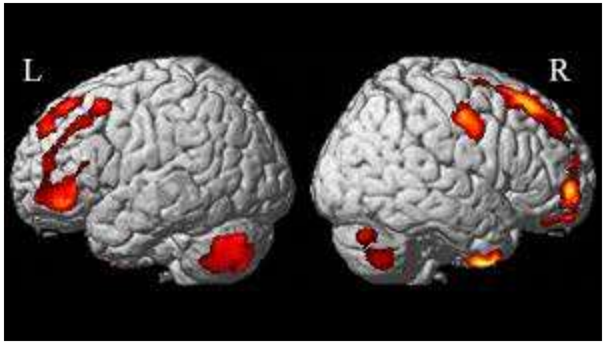

3. Neuroscience

4. Combined Approaches: Model-based fMRI Analysis

– Fit a cognitive model to behavioral data.

– Use model measures to interrogate brain correlates of model's operations.

4. LEVELS OF EXPLANATION

MARR'S (1982) LEVELS

- **Computational**
  - **What problem is the brain solving? What information is required? What is the structure of the environment?**

- Algorithmic
  - What processes does the mind execute to produce the solution?
  - What algorithms are computed?
    - In general, *not* a direct implementation of calculations the theorist employs at the computational level
- Implementational
  - Hardware: How are those algorithms implemented in the brain?

MARR'S (1982) LEVELS

- Cognitive scientists disagree over whether explanations at all levels are useful, and on the order in which levels should be explored.
- Connectionists: bottom-up or 'mechanism-first' strategy, starting by exploring the problems that neural processes can solve.
- →This often goes with a philosophy of 'emergentism': higher level explanations are at best approximations to the mechanistic truth; they describe emergent phenomena produced by lower-level mechanisms

CONNECTIONISM

- Connectionism models mental or behavioural phenomena as the **emergent processes** of interconnected networks of simple units.

- Units in the network could represent neurons and the connections could represent synapses like in the brain

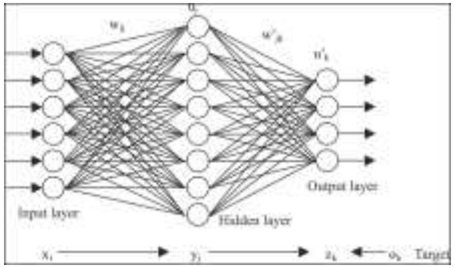

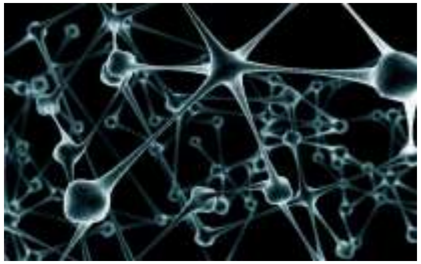

- Neural networks are by far the most commonly used connectionist model today

PROBABILISTIC MODELS OF COGNITION

Griffiths et al., 2010, *Trends in Cognitive Sciences*

- By contrast, probabilistic models of cognition pursue a top-down or ‘function-first’ strategy
- Beginning with abstract principles that allow agents to solve problems posed by the world – the functions that minds perform – and then attempting to reduce these principles to psychological and neural processes

PROBABILISTIC MODELS OF COGNITION Griffiths et al., 2010, *Trends in Cognitive Sciences*

- Top-down approach: to explore a broad range of different assumptions about how people might solve inductive problems, and what representations might be involved.
- Representations and inductive biases are selected by considering what is needed to account for the functions the brain performs, by assuming that those functions (perception, learning, reasoning, and decision) can be described as forms of probabilistic inference.

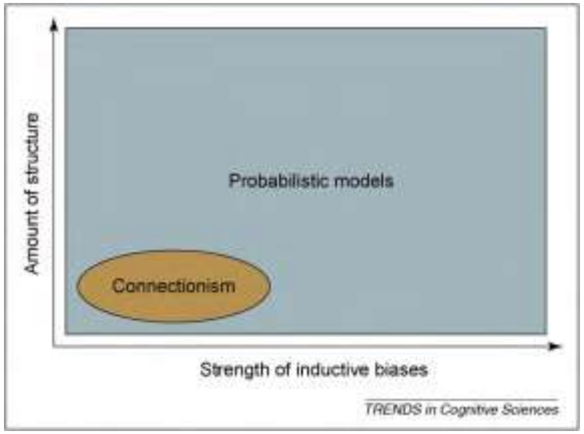

Probabilistic Inference: Bayes Theorem

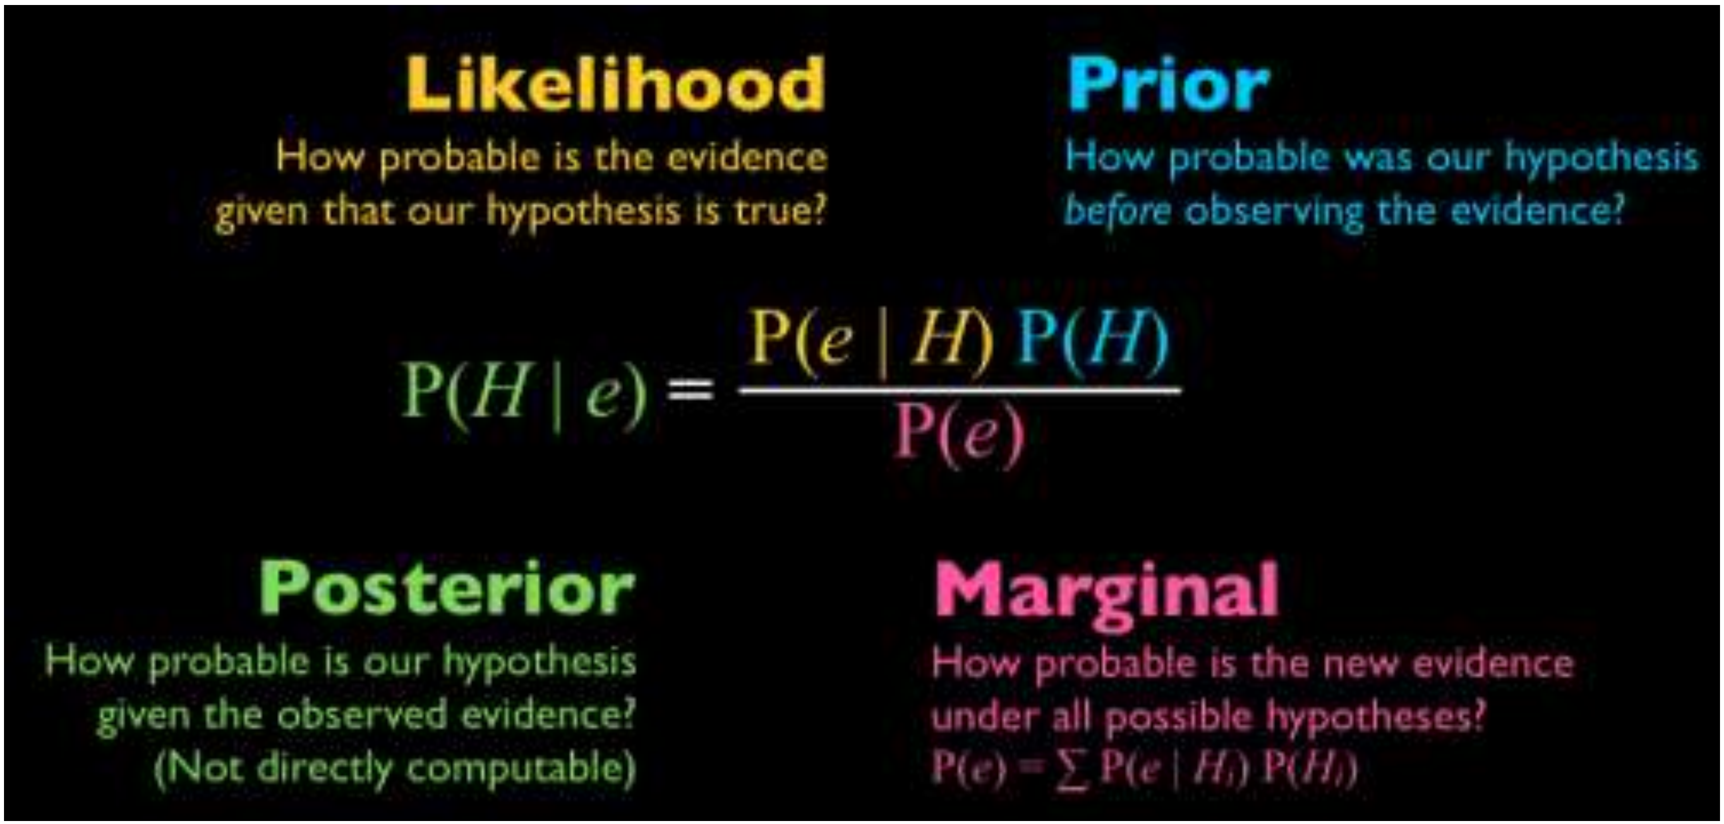

PROBABILISTIC MODELS OF COGNITION

Griffiths et al., 2010, *Trends in Cognitive Sciences*

Box 1. Probabilistic inference

Probability theory provides a solution to the problem of induction, indicating how a learner should revise her degrees of belief in a set of hypotheses in light of the information provided by observed data. This solution is encapsulated in Bayes' rule: if a learner considers a set of hypotheses H that might explain observed data d, and assigns each hypothesis $h \in H$ a probability $p(h)$ before observing d (known as the 'prior' probability), then Bayes' rule indicates that the probability $p(h|d)$ assigned to h after seeing d (known as the 'posterior' probability) should be

$$p(h|d) = \frac{p(d|h)\,p(h)}{\sum_{h \in H} p(d|h)\,p(h)} \quad (1)$$

where $p(d|h)$ is the 'likelihood', indicating the probability of observing d if h were true, and the sum in the denominator simply ensures that the posterior probabilities sum to one. Bayes' rule thus indicates that the conclusions reached by the learner will be determined by how well hypotheses cohere with prior knowledge, and how well they explain the data.

PROBABILISTIC MODELS OF COGNITION

- **Examples of their success:** see Griffith et al. 2010 paper
- **Limitations?**

Criticism: Bayesian Fundamentalism

- Significantly under constrained.
  - Excludes entire fields
  - Revisit assumptions as needed.
- Unfavorable comparisons to Behaviorism and Evolutionary Psychology.
- Optimality incomplete without mechanism
  - Time, energy, history, etc.

Jones, M. & Love, B.C. (2011). Bayesian Fundamentalism or Enlightenment? On the Explanatory Status and Theoretical Contributions of Bayesian Models of Cognition. Behavioral and Brain Sciences, 34, 169-231.

Bayesian Enlightenment

- Connect to mechanism. (Marr's Level 2)
- Evaluate assumptions of convenience

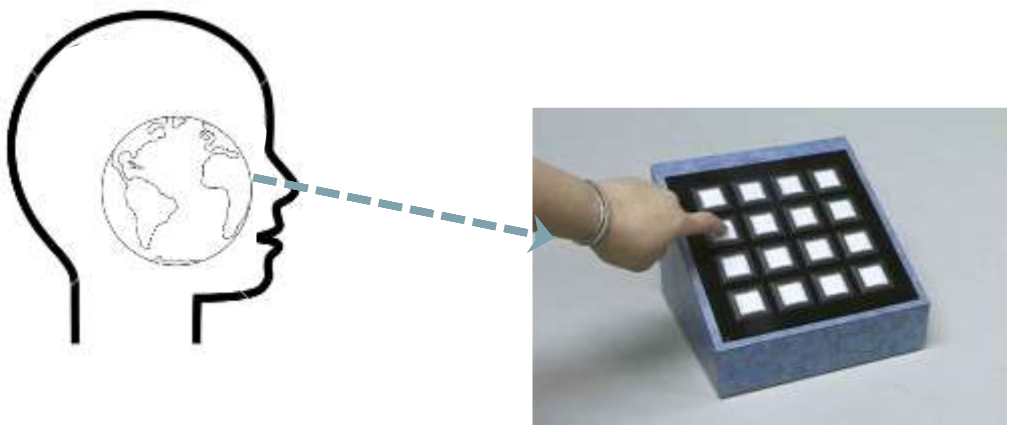

Why models?

Model Evaluation

- Capturing Qualitative Patterns
- Goodness-of-fit
  - overfitting
- Model Selection Statistics

Overfitting

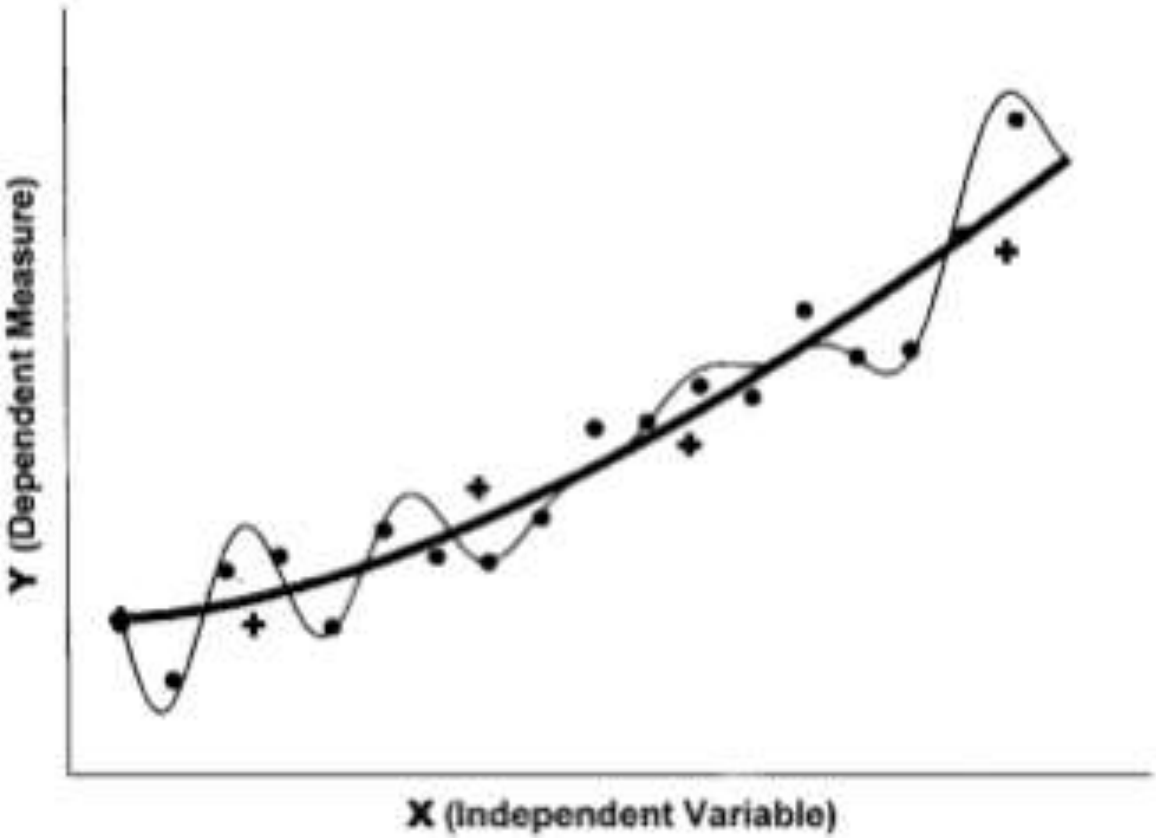

Overfitting

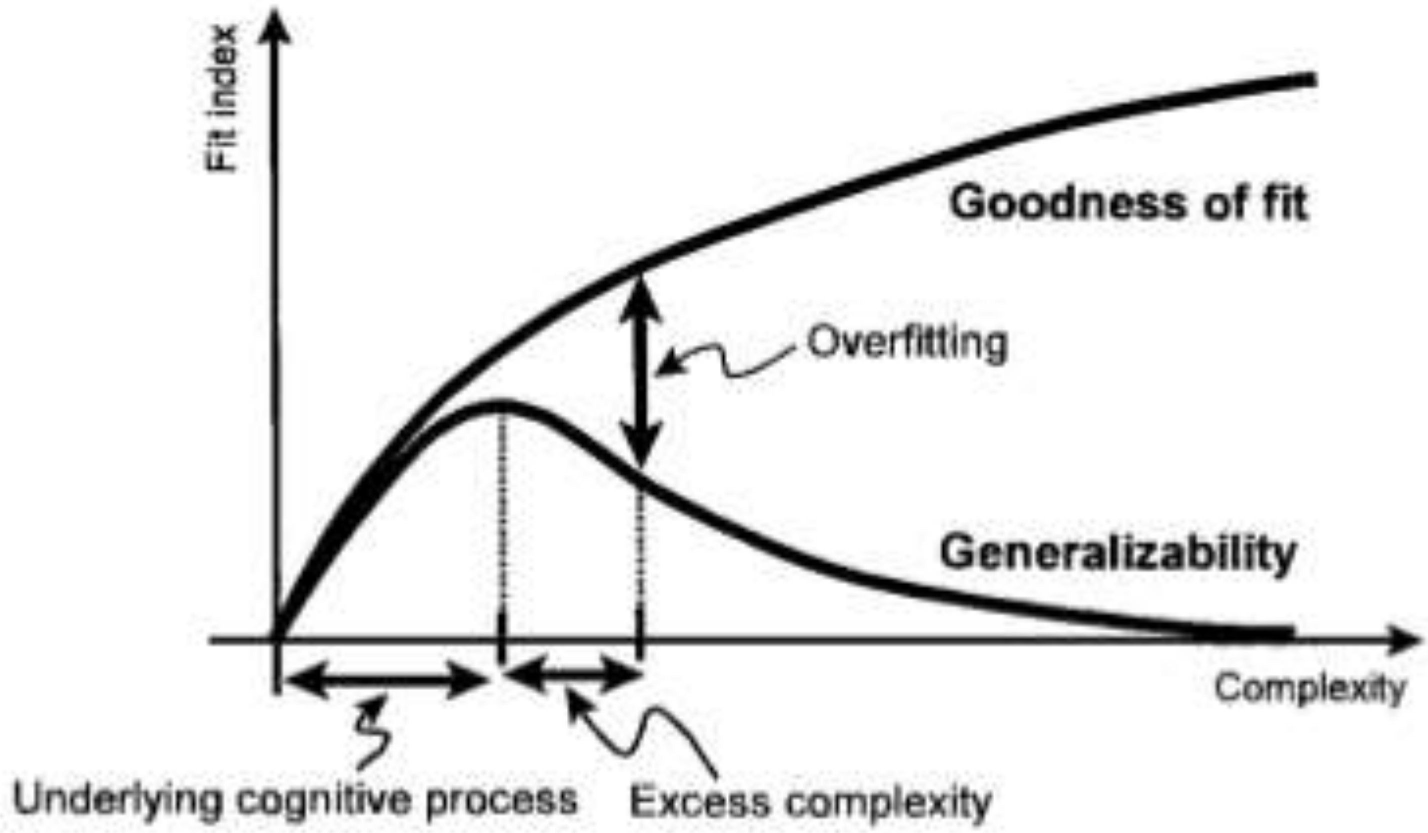

model selection

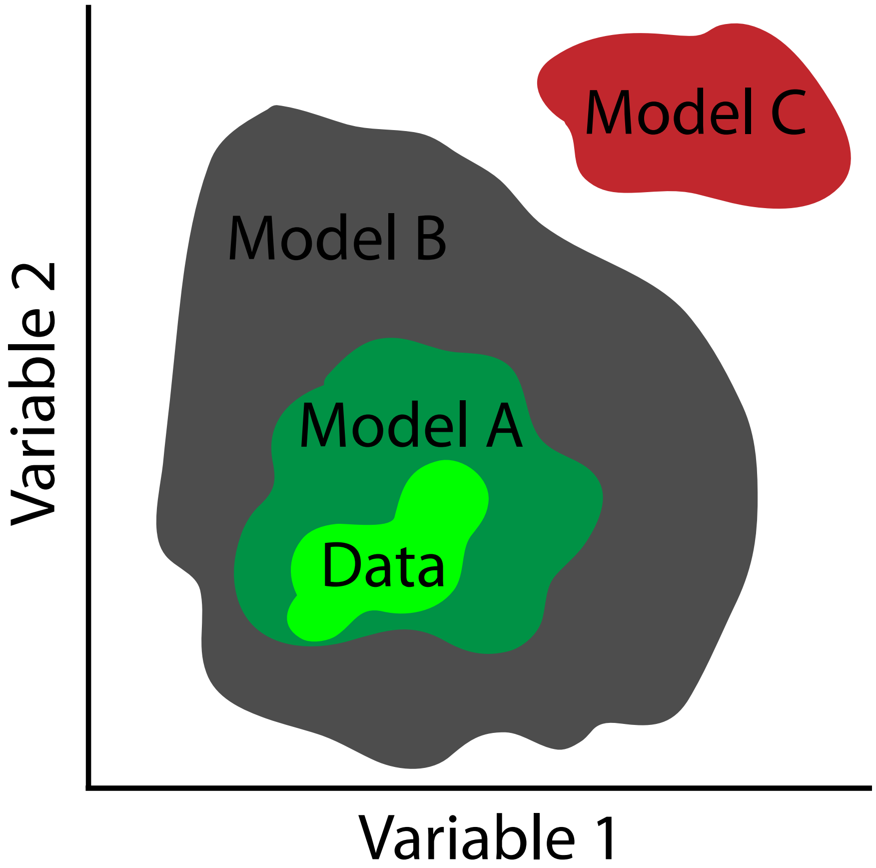

SUMMARY

1. Cognitive science as reverse engineering
2. Historical background: introspection, behaviourism, the computer metaphor
3. Methods in Cognitive Science
4. Levels of explanation

REFERENCES

- Nadel, L., & Piattelli-Palmarini, M. (2003). What is cognitive science. *Encyclopedia of cognitive science, London: Macmillan*. http://dingo.sbs.arizona.edu/~massimo/publications/PDF/LN&MPPIntro.pdf
- http://plato.stanford.edu/entries/cognitive-science/
- Griffiths et al (2010). Trends in Cognitive Sciences. http://cocosci.berkeley.edu/tom/papers/probmodelsofcognition.pdf
- Marr, D. (1982). Vision. Freeman: San Francisco.
- Jones, M. & Love, B.C. (2011). Bayesian Fundamentalism or Enlightenment? On the Explanatory Status and Theoretical Contributions of Bayesian Models of Cognition*. Behavioral and Brain Sciences*, 34, 169-231.

THANK YOU!