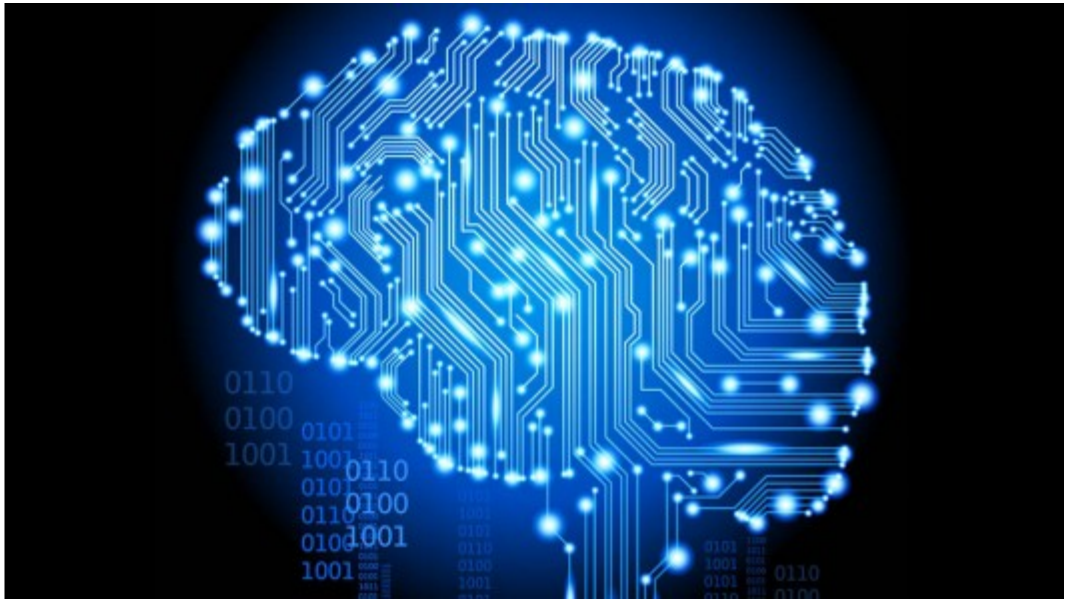

Computational modelling

Joe MacInnes

jmacinnes@hse.ru

Readings

- Some good books for those who want to dig deeper
  - Computational Modeling in Cognition - Stephan Lewandowsky, Simon Farrell
  - Computational models of cognition – Ron Sun
  - Artificial intelligence : A modern approach , Russel and Norvig
- Articles
  - We might do a few case studies of what makes a good model
  - Itti and Koch salience model
  - Wolfe ‘guided search’
  - Ratcliffe – diffusion models

Practical

- Hands on tutorials of techniques used in computational modelling
- Diffusion models
  - program in matlab
- Bayesian networks
  - Using a visual interface called Genie
- Neural networks
  - Neural net toolbox in matlab

Theoretical

- What does modelling add to psychology
- How to test models
- Advanced algorithms that may not have simple tools yet
  - Deep learning

What it is

- Models have always been a part of psychology
- Traditionally, verbal / theoretical models, but also mathematical
- Now computational
- Allow us to simulate increasingly complex brain functionality and cognitive processes

- Looking at the path of planets from earth's perspective
- What causes the retrograde motion?
  - Earth centric model explained a lot of the observable data *but not* retrograde

Model Solar system

- It was only ~500 years ago that Copernicus' model explained retrograde
- A combination of sun centric plus differing orbital velocity
- The new model also explained why planets changed brightness
  - Good models explain unexpected phenomena
  - Or make predictions about unobserved events

notes

- It took a full model of the solar system to explain one observation
- The model itself cannot be directly observed. Its an abstract representation
- Its not the only model that could have explain that data

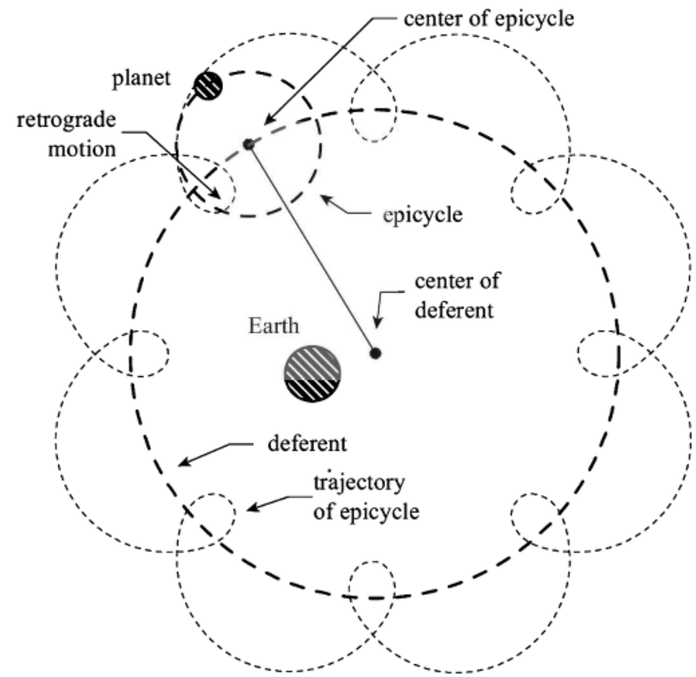

Ptolemy model, preferred model for 1300 years before Copernicus

Comparing models

- Copernicus model was slightly more accurate, and Plolomy's had an error of at most 1^o
  - Was this enough to explain its success?
- It was also simpler and more elegant
- Quantitative measurement against observable data is very important, but not everything
  - But then again, Kepler's adjustment 100 years later of elliptical orbits hit perfect accuracy
  - If two models are equally elegant, THEN go with the better fit?

Lessons learned

- Data do not ‘speak for themselves’ but require a model to be understood and explained
- Verbal theories alone cannot replace quantitative analysis
- Multiple alternative models are always possible
- Model selection depends on both quantitative analysis against the data and scholarly judgement

From Lewandowsky, computational models in cognition

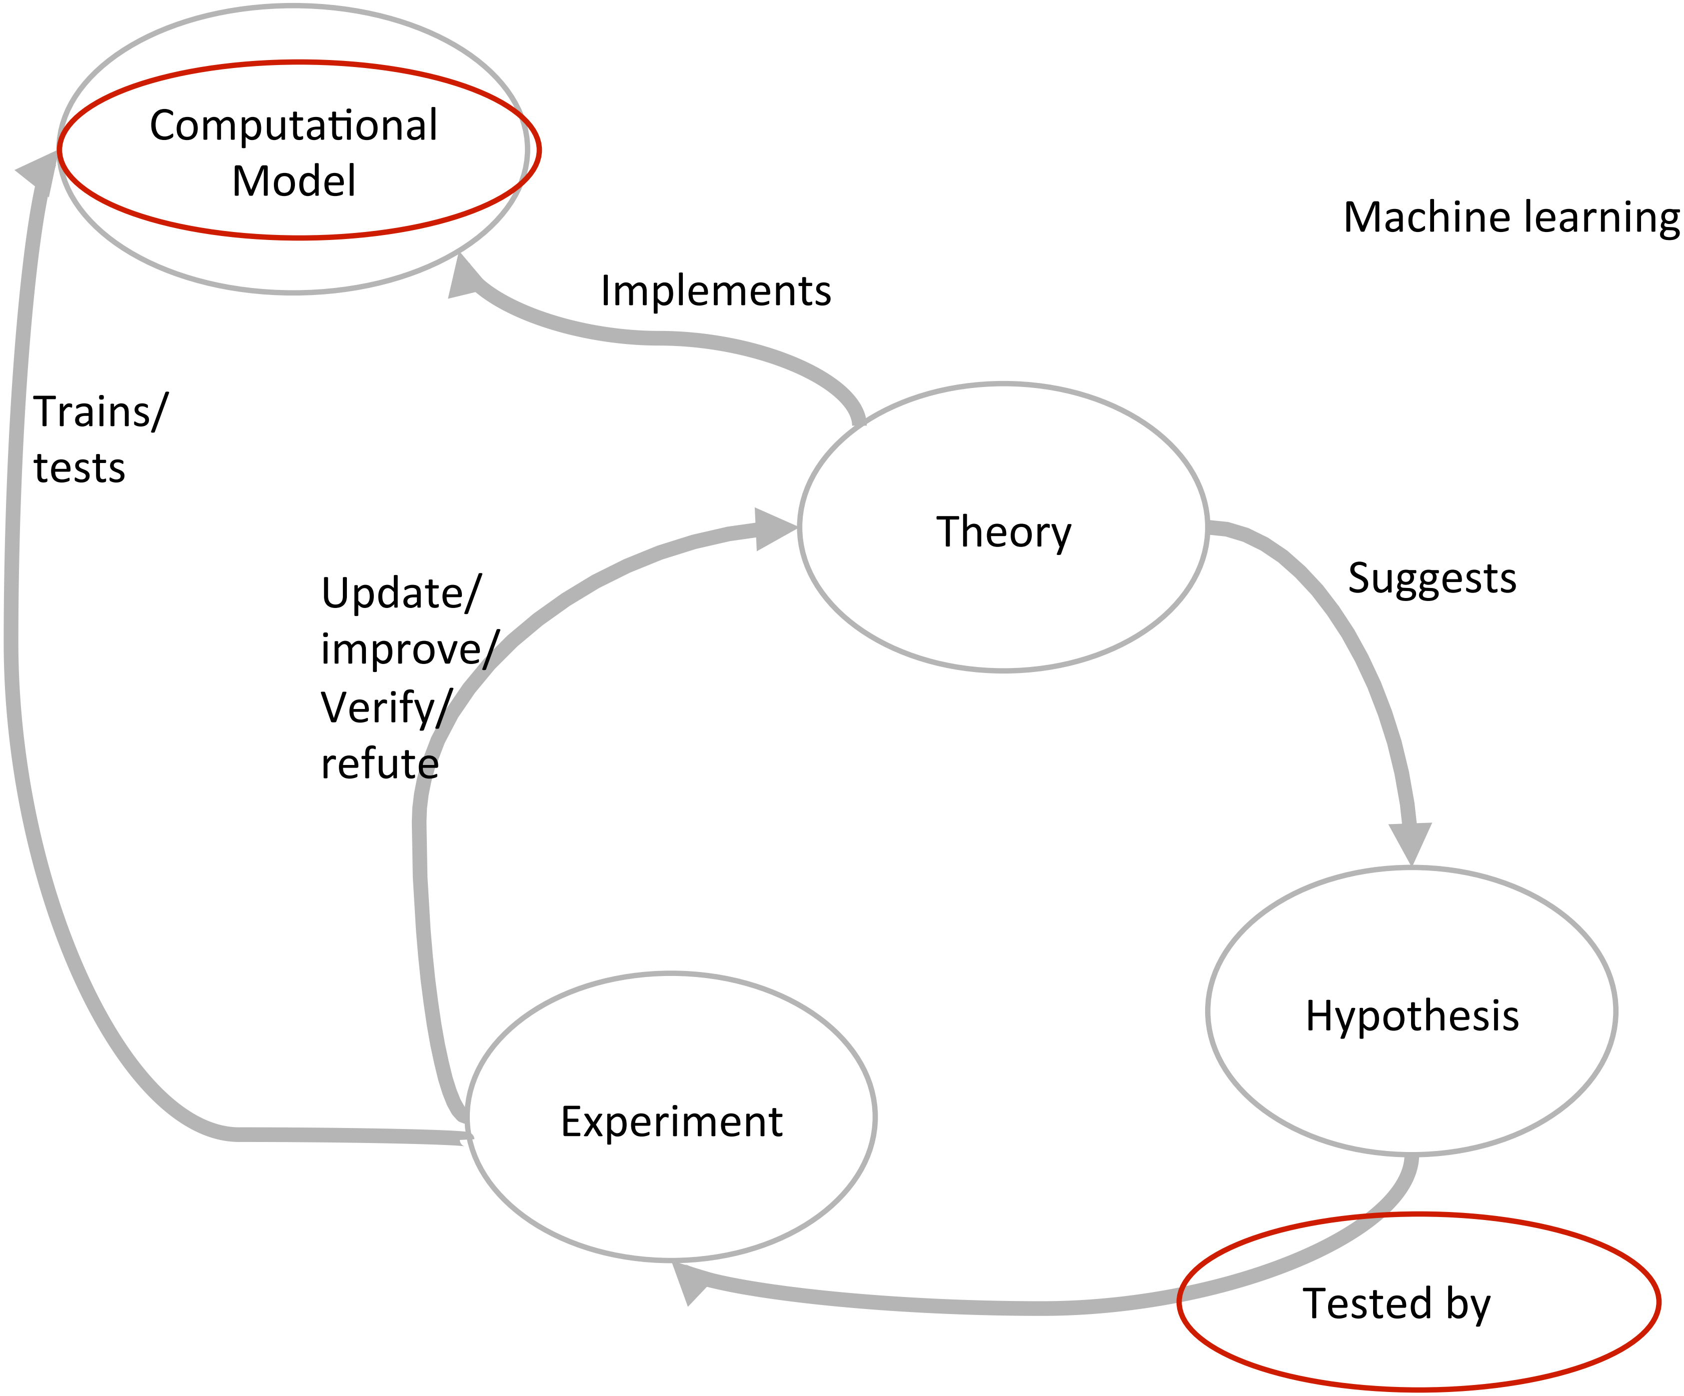
Computational Model
Machine learning
Implements
Trains/
tests
Theory
Suggests
Update/
improve/
Verify/
refute
Hypothesis
Experiment
Tested by

- Ron Sun says models aren't instantiated theory, but they are simulators that are theory building tools
  - They help you test limits and assumptions of (abstract) models
  - Also for hypothesis generation

Not directly observable

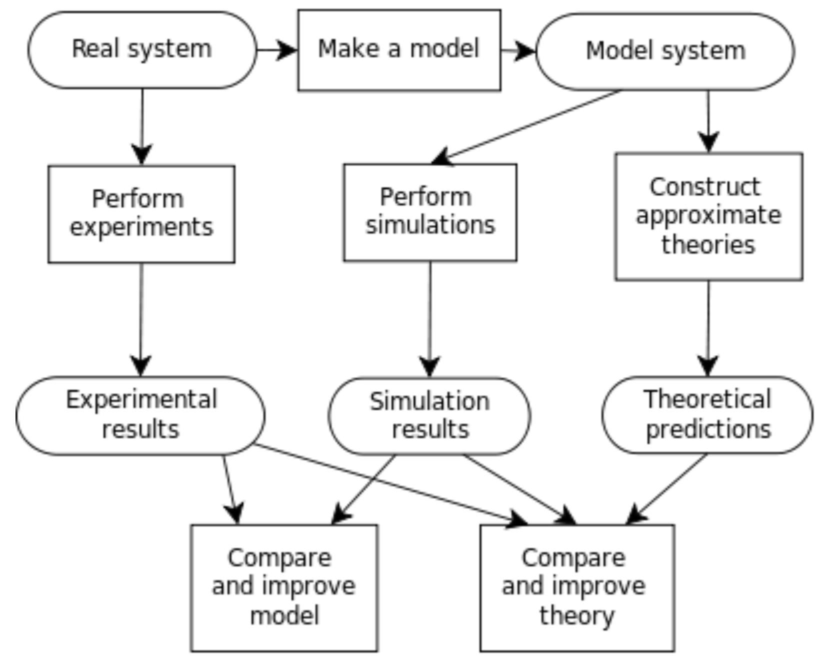

Spotlight model of attention

- Descriptive model
  - What do you mean the spotlight ‘moves’
  - How big is it?
  - How does it improve attention?
  - In what ways?

Mathematical precision of models

- When you have to state, explicitly, computationally and mathematically how each parameter of your model works, you can test it!
- Simulate experiments, lesions, individual differences

Lesions

- Patients with brain lesions have always played a huge role in psychology
- Computational models can maximize the contribution of these patients
  - Create model with theory and data from general population
  - 'Lesion' the model in the same manner as a patient
  - Compare the results
- TMS can offer the same role as lesion patients even if none are available
- Patient HM had part of hippocampus removed in a surgery that will probably never be done again
  - Modelling in now the only way this can be replicated

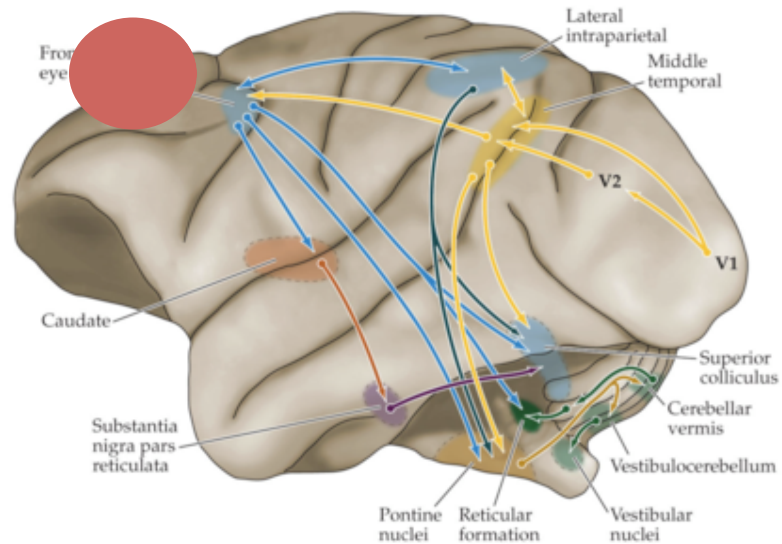

NEUROSCIENCE, Fourth Edition, Figure 20.13

What it isn't

- Analysis models vs theoretical models
  - Linear mixed effects models are becoming more common in statistics
  - But these are analysis models of experiment variables, not necessarily models of underlying cognition
- Scope
  - Single experiment vs more general
  - Multiple brain regions?
- We will not cover single experiment 'models' in the class
  - Eg linear mixed effects models
- These techniques are certainly valuable, they just aren't what we are discussing in the class
- We want more 'process' models than 'descriptive/analytic' models

Hypothesis testing vs modelling

Modelling

- Fit data from many experiments in one model
- Emphasis on HOW input matches experiment results
- Multiple hypothesis as well as generating new ones
- Uses machine learning techniques
  - To learn model parameters
  - As essential metaphor for cognitive/neural processes
- Broad scope. Interaction of multiple processes/brain regions

Hypothesis testing

- Typically single experiment
- Focus on IF experiment input variables match experiment results
- Tests validity of single hypothesis
- May use machine learning techniques as replacement for statistical analysis (classifiers)
- Limited scope. Precisely test single question/brain region

Both provide valuable (but different) tools to psychology

Lewandowsky - Three classes

- Descriptive
  - The mean of a data set is the simplest descriptive model of that data
- Process
  - State the processes and stages but agnostic to how those stages work
  - Simple mathematical or visual representation of a verbal model
- Explanatory
  - Full detail of stages and mechanisms of each stage
  - At least within the scope of the model
  - These are the models we're most interested in

Marr 1982

- For example, Marr listed 3 levels of modelling
- Computation (high)
  - Goals, logic and strategies of computation
- Algorithm
  - Input and output representation of the data used
- Hardware implementation
  - Physical realization of algorithms

- Though these days we think more along the lines of a symbolic <–> biological continuum

Prediction and classification

- What if your neighbor could predict the results of every Psych experiment (Forster, 1994)
  - Is that sufficient?
  - No, our main goal is to understand
- Classifiers in AI focus mainly on prediction
  - Is this packet a DOA attack?
  - Is this email spam?
  - Will this person click on this ad and buy something?
- They can be extremely accurate, but they aren't models

Analysis vs modelling

- Nosofowsky, 1991
  - Subjects classified faces into 2 arbitrary categories based on any features they wanted
  - Then instructed to use same categories/features with same with old plus new objects
  - Also asked 'confidence' of classification
- Did confidence of old (recognition) match confidence of new (classification)
  - Not according to analysis, only R = .36
  - Separate cognitive functions?

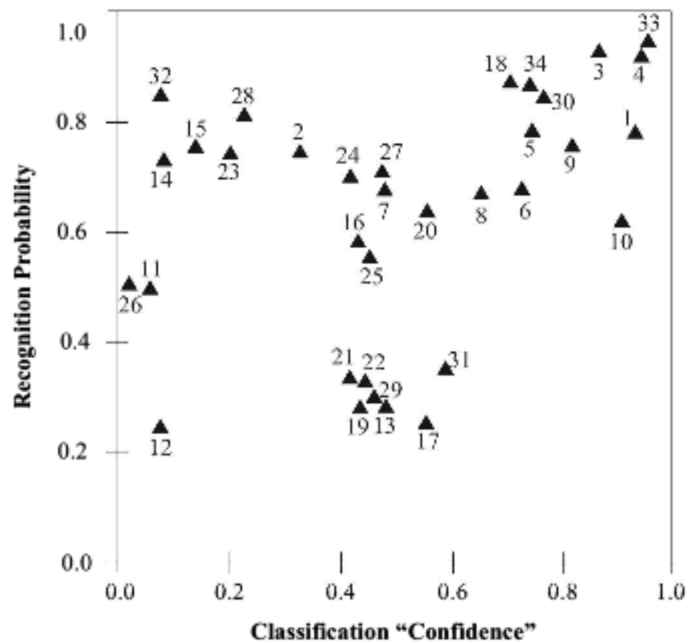

R = .36

model

- But…. A model suggested otherwise
- Successfully predicted categorization and recognition for all subjects
  - 91% and 96% accuracy
  - Using the same model and same parameters for both responses
  - This is an existence proof that both responses could be generated using the same process
- Generalized Context Model (GCM)

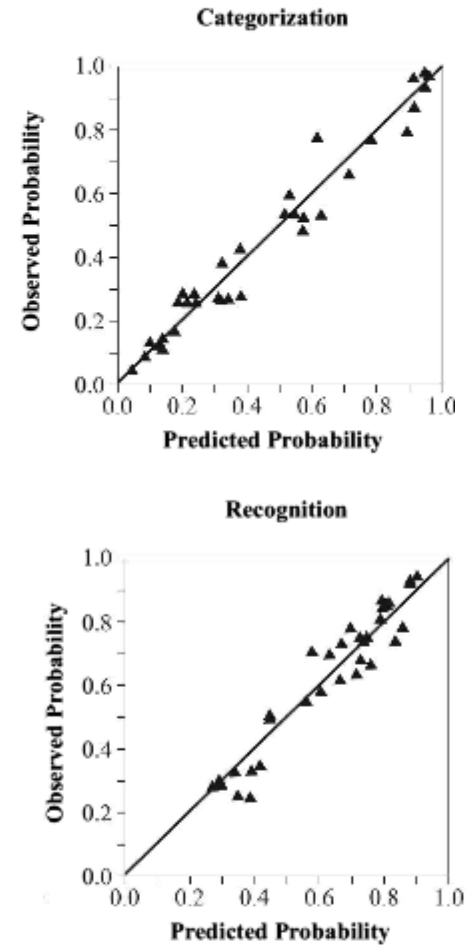

Models of models

- Johnson-Laird, Byrne, & Schaeken, 1992
- Proposes that cognition is the brains attempt to model the world
  - With purpose of predicting possible events and outcomes
  - So computational modelling is our model of the model of the world

Simplify complex systems

- What other fields use computational modelling?
  - Meteorology (weather, tornados)
  - Physics (rocket launches, black holes, big bang)
- **The goal is to understand and predict complex systems**

Limitations

- Hardware
  - The brain is massively parallel, but computers are largely sequential
- Can the brain be simulated using the language of mathematics and computers?
  - Best guess at the moment is yes
  - But.. current mathematics was created to model the physical world, not the mind
- https://aeon.co/essays/your-brain-does-not-process-information-and-it-is-not-a-computer
  - Where in the brain are the lyrics to a 'deep purple' song?
  - Where in the brain is the memory of your 5^{th} birthday?

Simplification

- The best model of a cat is another cat. Or better, the same cat
  - Neubert
- How much do you have to simplify a system to understand it?
  - Depends on your question
- **Level of granularity** (zoom/abstraction)
  - Multiple models of the same brain area/function can be equally valid at different levels of abstraction
  - Different zoom will answer different questions
- **ALL** of our current models are simplifications of the human brain
  - The only questions are how much, where, and what aspect do we chose to focus on

Falsifiability

- A perfectly accurate model that covers all possible examples is a good model, right?
- Only if it makes predictions that can be proven false
- A model that is trivially true doesn't help us understand anything

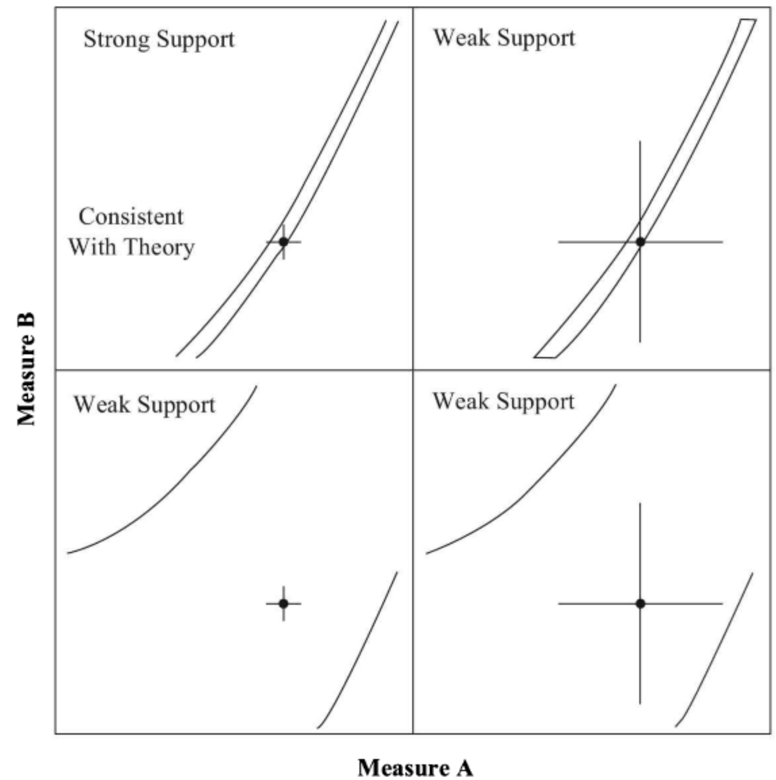

The top model can be proven false with data outside its range. The bottom model cannot. (From Roberts & Pashler, 2000)

Complexity of Scale

- Models allow looking at the complexity of scale
- Single neurons compose the brain in the same way millions of grains of sand compose a sand dune
- But studying single grains of sand does not help you understand a sand dune
- Some researchers believe large scale models are required to understand cognition
  - Emergent properties arise from interaction?

Modularity

- Experimental psychology is often accused of reductionism
- Models allow us to put those pieces back together again in a meaningful way

Types of models

- Symbolic (probably not covered in this workshop)
  - ACT-R
- Neural
  - Neural nets
  - Spreading activation
- Hybrid?
  - Bayesian

Generative vs discriminant

- Generative models allow random generation of observable data
  - Usually by specifying joint probability distribution over the variables (more on this later)
  - Over ALL relevant variables. Including hidden variables
  - Mixture models, Naïve Bayes, Hidden Markov Model
- Discriminant models
  - Target and observed variables only
  - Neural networks, Support vector machines

What is the goal

- Exploration of implications
  - We can make changes to our models and test them (simulate) in ways that are not possible with humans
  - Even animal model could replace or reduce animal research
- Emergence of understanding
  - Our models can surprise us
  - Even though they are computer rules, they will sometimes make simulated predictions that lead us to new research questions
- Point of failure
  - Given current state of the art, all models will fail on some problems and data
  - But the point of failure can also be telling

Artificial intelligence

- Sample of a few algorithms that may help

- **Neural nets (dozens of types)**
- Mixture models
- Unsupervised models
  - Clustering, K-Means, fuzzy clustering
- **Bayesian**
- Symbolic
  - Constraint satisfaction
  - Search tree algorithms (A*)
- **Genetic algorithms**
- Intelligent Agent approach (social models)
- Temporal models
  - Recurrent networks, Hidden Markov Models (HMM)
- **Diffusion models**

Thinking Humanly Cognitive Models “The exciting new effort to make computers think . . . machines with minds, in the full and literal sense.” (Haugeland, 1985) “The study of mental faculties through the use of computational models.” (Charniak and McDermott, 1985) Artificial General Intelligence	**Thinking Rationally** Problem solving “The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990) “Computational Intelligence is the study of the design of intelligent agents.” (Poole et al., 1998)
Acting Humanly Turing test Embodied Cognition (remember Gibson?) “[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . .” (Bellman, 1978) “The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)	**Acting Rationally** Software Agents, Robotics “The study of how to make computers do things at which, at the moment, people are better.” (Rich and Knight, 1991) “AI . . . is concerned with intelligent behavior in artifacts.” (Nilsson, 1998)

From AI, A modern approach

State of the art?

- There was a time that cognitive psychology and AI were closely tied
- Now? The field of AI has progressed far beyond what most people use and understand in cognitive psychology
- This means you don't have to be an AI guru to use these techniques in cognitive modelling
- There is a lot of work that can be done even with 'classic' AI techniques
- Many of these classic techniques are established enough to have very user friendly tools
  - My last paper on Bayesian modelling, I could have coded everything out by hand, but why should I? I saved weeks of programming using Genie instead with no loss to the model

Robotics

Not our focus, but there are still some research areas where cognitive modelling overlaps with robotics

Google driverless car

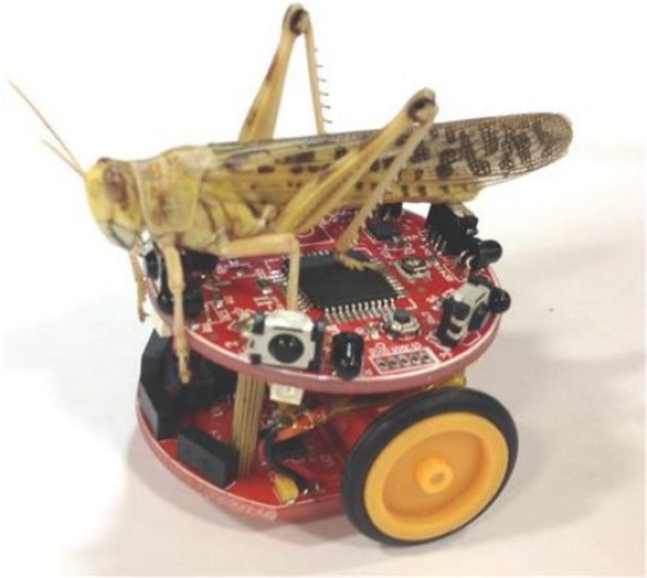

Robot insects imitate social/swarming behaviour

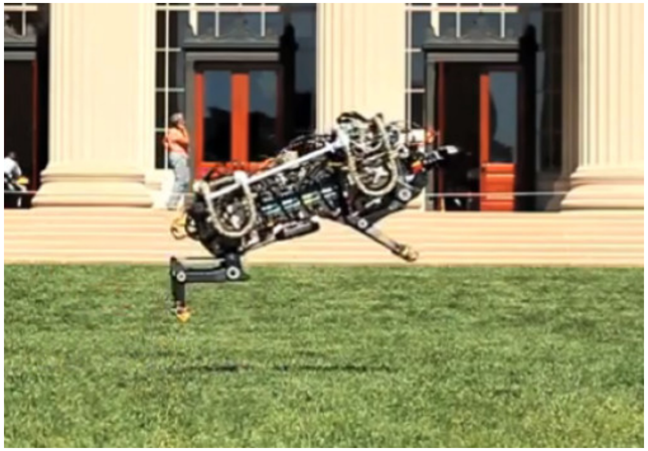

MIT robotic Cheetah running and jumping

The best model of a cat

- Can we model a full brain?
- 100,000,000,000 neurons
- Massively parallel and recurrent

- The Blue Brain Project
- at the Ecole Polytechnique Federale Lausanne
- (ted talk)
  - Biological model, not cognitive
  - Each neuron run by biologically plausible model
- ‘Neuron’ software written mostly in C and used within the GENESIS simulation environment

Euro Human Brain project

- 1 billion Euros
  - Estimated 7148 person-years of work
  - Building on successes with Blue Brain project
- Focus on hardware simulation on super computers
  - This is NOT a cognitive model, and will not provide cognitive insight, at least in this iteration
- Simulate drug treatments

- Ted markham: a brain in a supercomputer
  - http://www.ted.com/talks/henry_markram_supercomputing_the_brain_s_secrets?language=en
- Cat brain fight
  - http://spectrum.ieee.org/podcast/computing/software/cat-brain-cat-fight

Backup

Example: Diffusion models

- Overview this week
- Hands on programming next week
- Read Ratcliffe and McKoon, 2008 for detail

Ratcliff, 1978

- Diffusion model suggests that evidence builds gradually towards one of a number of outcomes

- Ratcliffe had data from a monkey saccade experiment
  - SRT behavioural data
  - SC single cell recording

- A diffusion model was built on behavioural data to predict SRT and accuracy of saccade
- And then tested on the single cell data from SC
  - Evidence accumulation from the model matched increase in activity in SC cells
  - The nearer the model was to a decision, the greater the firing rate

- Is it the behavioural data or the neural data that make this such a great paper?
- Perhaps both?
- Actually, its that you can integrate both together in a model/framework

First, a random walk

- also called drunkard's walk in
- Is a type of random search or exploration over time
- Searching physical space, but also other abstract spaces, like 'evidence'
- At every time step, randomly move some distance along possible dimensions
- Incorporated as starting point or an element of some models
  - Eg eye movements in visual search

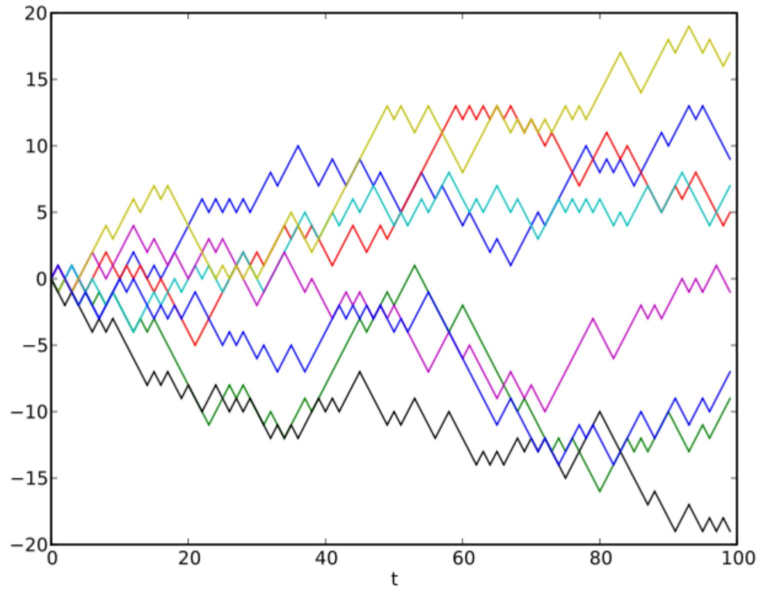

A number of random walks in one dimensional space

Diffusion model

- AKA accumulator model, threshold model
- At its simplest, we could model one or more random walks that stop when they pass one or more thresholds
- And these thresholds represent possible outcomes
  - Decisions
  - Saccade directions
  - Responses
- The model makes predictions on time, choice and accuracy

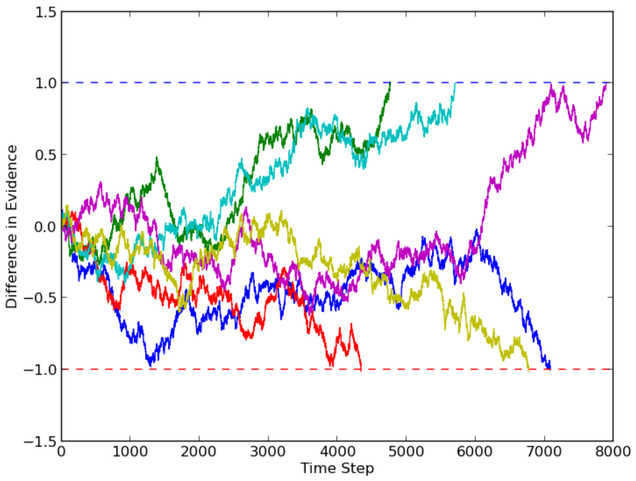

Multiple walks can represent multiple possible choices or targets. Assumes evidence and noise accumulate separately for each option

What parameters are missing?

- At every time step we can add
  - **Noise** in a random direction (random walk)
  - **Signal (or drift rate)** about the correct 'direction'
- The **threshold** (or boundary) is the value that the walk must reach for a decision or a response to be made
  - Threshold **shift** can be made in response to correct or incorrect responses to reflect an adjustment of threshold

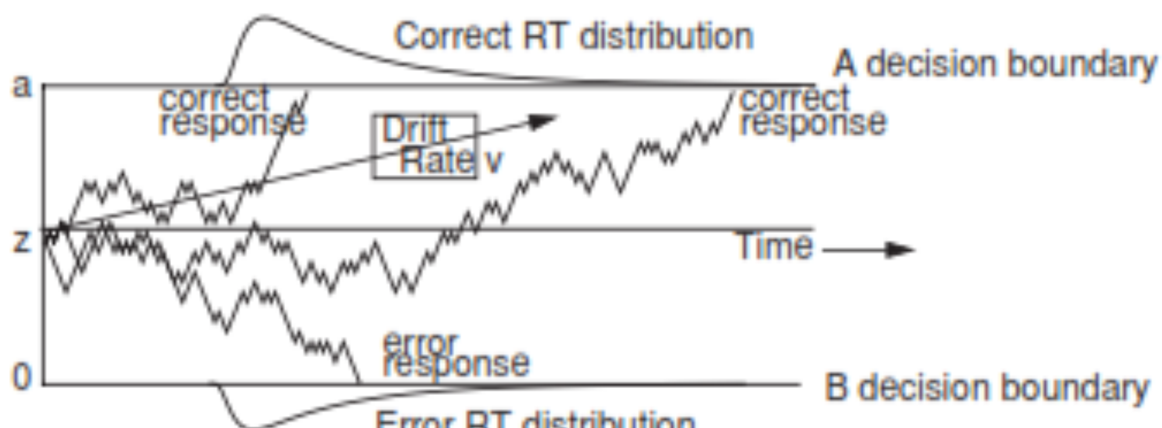

Review, Ratcliffe and McKoon, 2008

- Single model explains reaction times and error rates
- Generative: can produce response distributions
- Replicates positive skew of most RT distributions
- Drift rate can vary across trial and across subject
  - Often chosen from one or more random distributions
- Bias can be modelled as shift in starting point
- **Parameters of the model represent aspects of cognitive processing, and can change as we manipulate experiment parameters**
- Ratcliffe version: for single stage decisions less than 1500ms only

- Signal can represent
  - Size of target
  - Amount of practice on a memory task
  - Attention or cuing

- Noise can represent
  - Number of distracters
  - inhibition

tweaks

- Trial by trial adjustment of threshold criteria shift
  - Correct answers make us more liberal
  - Incorrect answers make us more conservative
- Different parameters change different quartiles of the distribution
  - For example, the same shift in drift rate (X) will have a larger impact on the tail end (Z) of the distribution than on the leading edge (Y)
- Other variants, including EZ diffusion model of RT
  - Wagenmakers, 2008
  - Linear Ballistic Accumulator

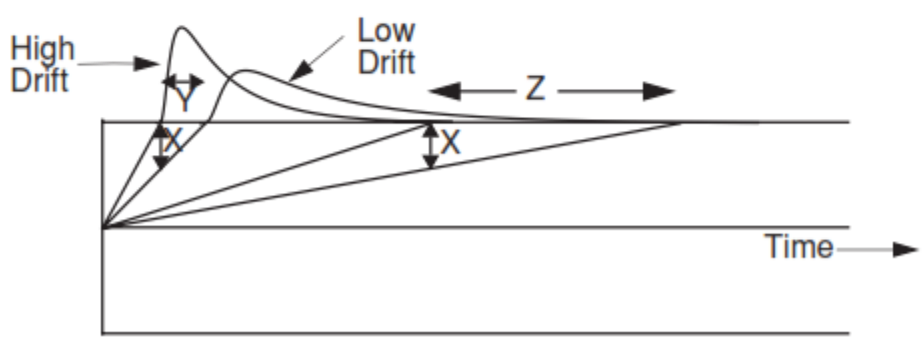

Ludwig etal 2009

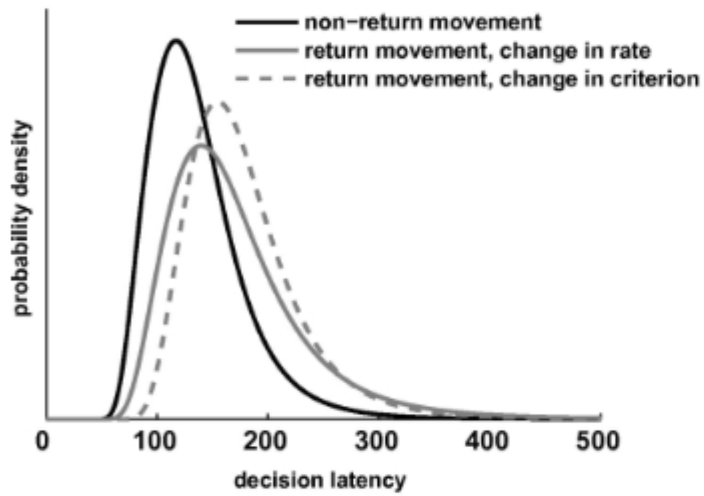

MacInnes 2016

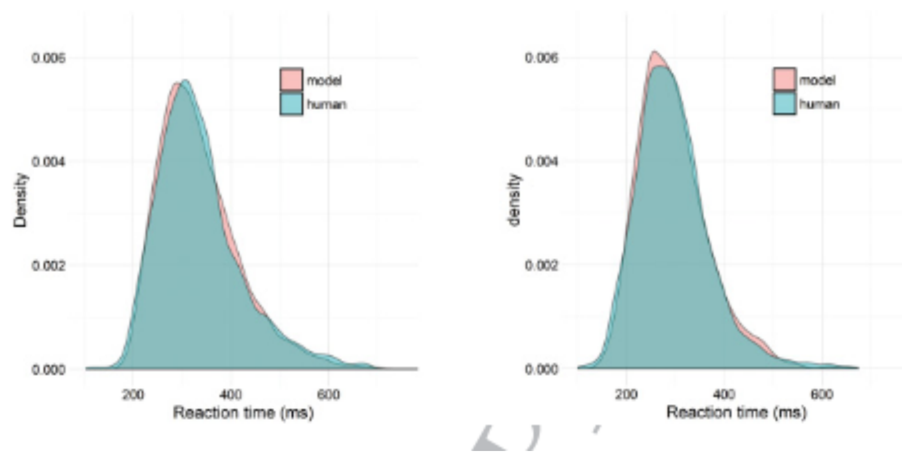

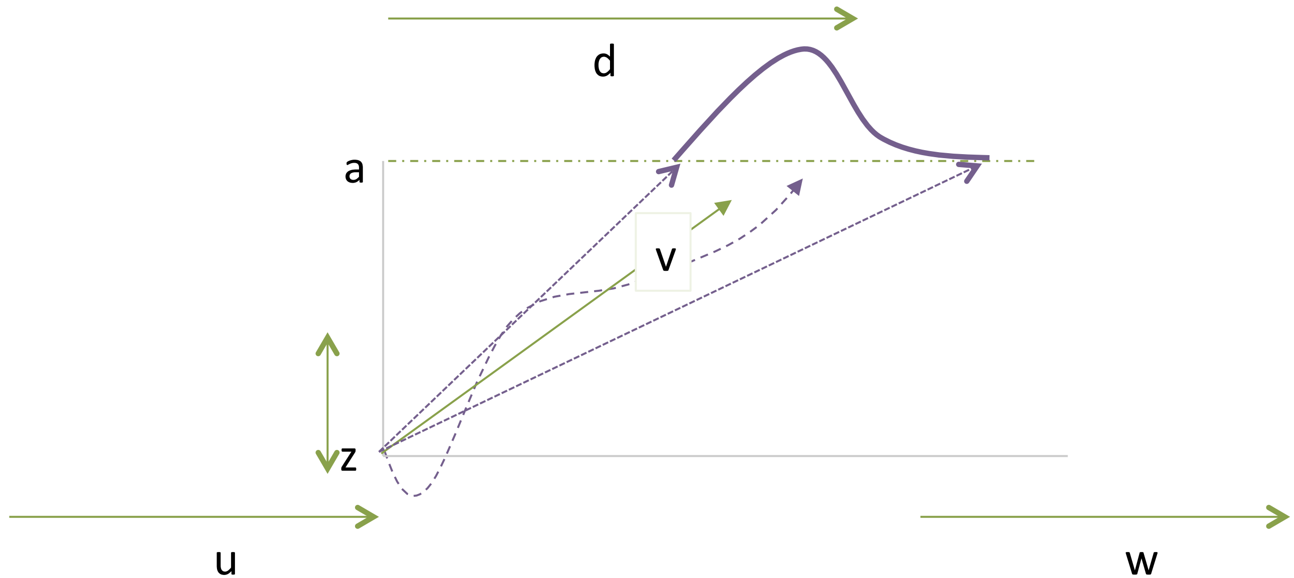

Decision components
z – starting point
a – decision boundary
v – drift rate signal
S(v) – variability (noise) in signal within trial
h – between trial variability
S(z) – flat distribution range for starting point

Non-decision components
S(t) = u+w – across trial variability of all other non decision components (Pre and post decision, possibly overlapping)

RT = (u+w) + d

Single threshold diffusion models typically do not require S(z) (Ratcliffe 2011)