Topics in Cognitive Modelling: Course Introduction

Jan. 14, 2014

John Lee, Chris Lucas
School of Informatics
University of Edinburgh
{J.Lee,C.Lucas}@ed.ac.uk

(with thanks to Sharon Goldwater)

The what and why

- What is cognitive modelling and why do it?

The what and why

- What is cognitive modelling and why do it?
- Why study cognitive science at all?

The what and why

- What is cognitive modelling and why do it?
- Why study cognitive science at all?
  - We want to know how the mind works:
    - How we process information and act on it.
    - How we learn and generalize.
    - How we think, reason, and make decisions.

Studying the mind

- Experiments can yield *facts* about behavior.
- If we want to *predict new behavior*, we need a theory.
    - Explains why we observed what we did.
    - Predicts what would happen in a new situation.
- A computational model is just a very explicit theory.
    - Implementation forces explicitness.
    - Often brings up issues we wouldn't have thought of otherwise.
    - Comparing the model predictions to human behavior allows us to test and refine the theory.

Levels of analysis

- Models can be explicit in different ways. Marr (1982) discussed three *levels of analysis*:
  - Computational: What is being computed?
    - Ex. Optimize a function.
  - Algorithmic: How is the computation carried out?
    - Ex. Compute derivative and use gradient ascent.
  - Implementational: What hardware is used?
    - Ex. Digital computer.
- We'll mostly focus on the first two types of model.

Other assumptions

- Models also differ in many other ways, for example assumptions about
  - Representation (symbolic or distributed).
  - Domain-specificity and modularity.
  - Need for and nature of built-in (innate) constraints.
- Studying and comparing different models can shed light on long-standing debates in cognitive science.

Goals of this course (1)

- Examine the Big Questions of cognitive science through the lens of computational modelling.
  - Is cognition a collection of separate domain-specific abilities or an interacting whole?
  - How much of cognition is innate?
  - Are mental representations symbolic or distributed?
  - Are mental processes based on rules or associations?
  - To what extent are our cognitive abilities determined by our physical body and environment (i.e., grounded/embodied)?

Goals of this course (2)

- Learn (more) about different modelling approaches and how they relate to these Big Questions.
  - Connectionist
  - Bayesian/probabilistic
  - Algorithmic/mechanistic
  - Dynamical systems
  - Cognitive architectures

Modelling approaches

- Connectionist:
  - Emphasizes distributed representations and general-purpose statistical learning mechanisms.
  - Implemented as artificial neural networks:

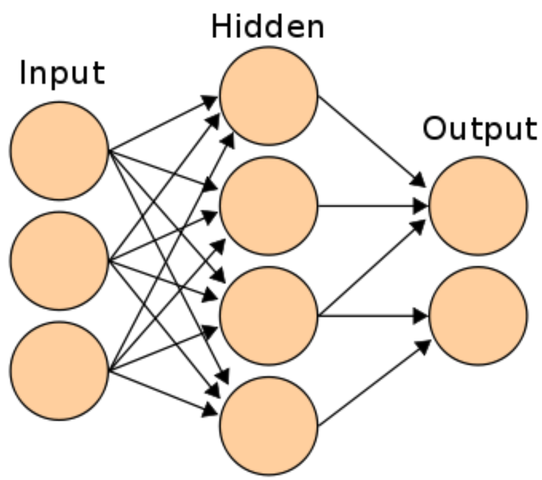

Figure: http://en.wikipedia.org/wiki/Artificial_neural_network

Modelling approaches

- Bayesian:
  - Emphasizes computational-level explanations using probability theory, optimal behavior under uncertainty.
  - Shares techniques with statistical machine learning methods.

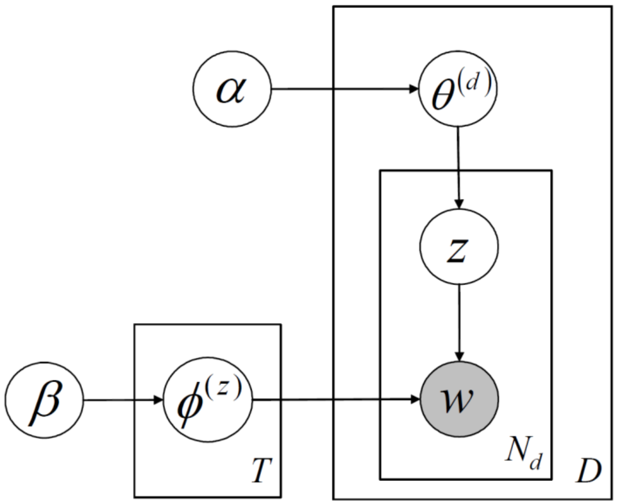

$$P(w_i) = \sum_{j=1}^{T} P(w_i \mid z_i = j) P(z_i = j)$$

Figures: Steyvers and Griffiths, 2007.

Modelling approaches

- Algorithmic/mechanistic:
  - Emphasizes procedural steps involved in processing information, usually in a specific domain.
  - Not really a single approach or philosophy, so may be symbolic/rule-based or statistical.

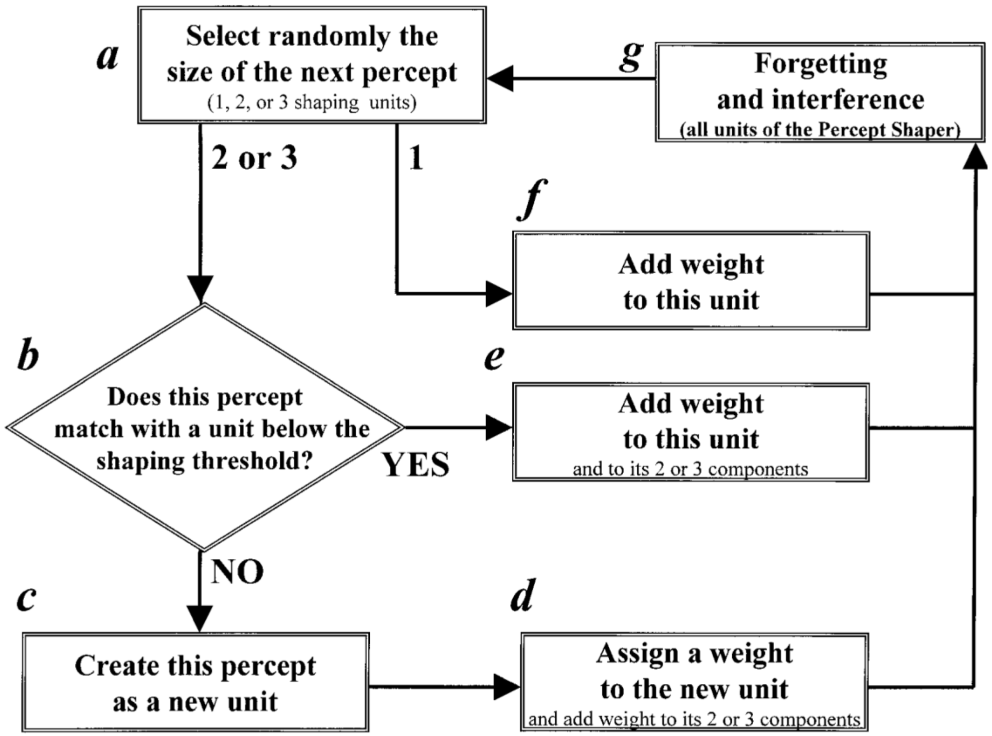

Figure: Perruchet and Vinter, 1998.

Modelling approaches

- Dynamical systems:
  - Emphasizes complex interactions between mind and environment, rather than internal representations.
  - Connections to robotics and philosophy of embodied cognition.

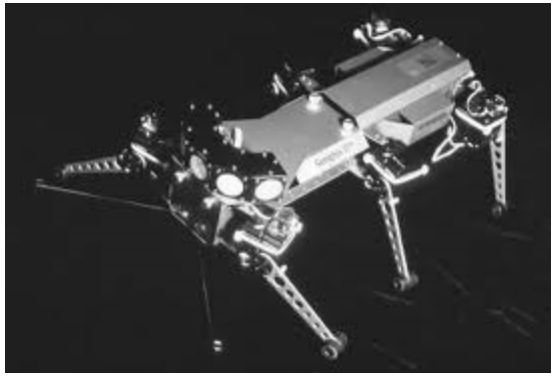

Images: scienceclarified.com (L), wikipedia (R)

Modelling approaches

- Cognitive architectures:
  - Emphasizes information flow and modularity, as well as timing. Rule-based or hybrid (rules + activation levels).
  - Also more focused on applied work than other approaches.
    - Ex. How will adding a new display to a control panel affect a pilot's reaction time and attention to a warning light?

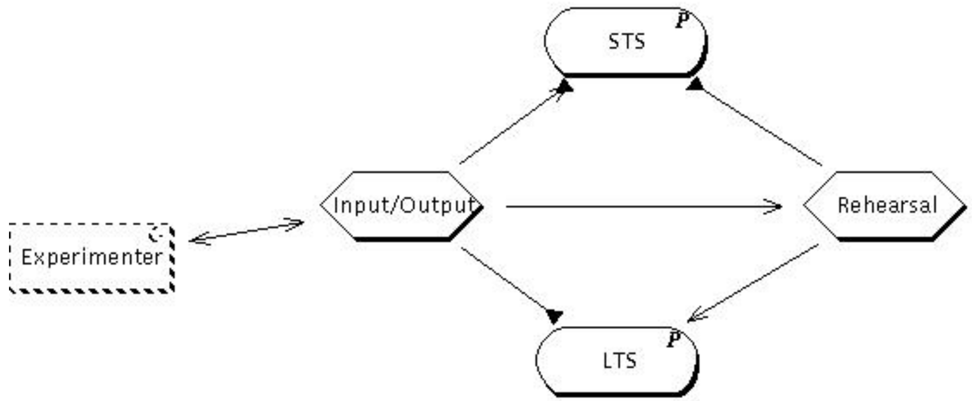

Goals of this course (3)

- Look at a few topics in relative depth.
  - Learn a bit about the phenomena in question.
  - Compare different modelling approaches.
  - What do we learn from different approaches?
  - What questions remain?
- Specific topics:
  - Various topics in language
  - Categorization
  - Infant object perception and knowledge
  - Possibly others: motor control, causal learning, etc.
  - Models of both development and adult processing.

Goals of this course (4)

- Develop students' analytical and communication skills.
  - Being able to summarize the main issues and methods in a scientific paper.
  - Critically analyzing prior work for strengths and weaknesses.
  - Comparing different approaches and techniques.
  - Presenting and discussing this information clearly in both oral and written form.

You lead this course

- Learning through reading, writing, and discussing.
- Student presentations, in-class discussion of readings.
- We will help facilitate discussions, but expect you to prepare and come with questions/comments.
- Past students have said class discussion is one of the best parts of class, but it will depend on you.

Course structure

- Weeks 1-3: 4 or 5 lectures by us.
  - Background on themes and methods.
  - How to read, analyse and present research papers.
- Weeks 4-8: ~10 presentations by you.
  - Presentations in groups of 2-3.
  - (Note: 1 week break between weeks 5 and 6 for Innovative Learning Week – www.ed.ac.uk/innovative-learning)
- Week 9: final paper due.
- No exam.

Assessment

- Oral presentation: 20%
  - In groups of 2 or 3, presenting usually 2 papers with different models of similar phenomena.
  - Students choose topics from list on course website.
  - Summarize psychological phenomena and models, discuss differences in philosophy and approach, strengths and weaknesses, relationship to other models in course. Also raise questions for further discussion with class.
  - Plan on around 35 minutes for presentation, plus 15 minutes for questions/discussion.

Assessment

- Brief paper responses: 25%
  - Each approx. 1-2 paragraphs, worth ~3-4%.
  - Due in class on each presentation day.
  - Choose one paper from that day's readings, give a brief summary and your thoughts or questions about the paper.
  - No excuses or late responses will be accepted, *but …*
    - … you may skip three responses without penalty.

Assessment

- Final essay 55%:
  - 2500-3000 words, summarizing and analysing one or more cognitive modelling papers on a single topic.
  - Topic/papers must be approved by instructor, by mid Feb. (date TBA)
  - Essay due date March 20th.

Prerequisites

- Ideally, Computational Cognitive Science.
- Some background in one or more of:
  - Cognitive psychology
  - Linguistics
  - Artificial intelligence/machine learning
- Ability/willingness to engage with mathematics.
  - Knowledge of probability helpful, but some intro provided and tutorials on website; also flexible reading list.
- Strong English skills.
  - This course requires a lot of reading and writing; if you have trouble with English, it will be extra difficult.

Course information

- website: http://www.inf.ed.ac.uk/teaching/courses/tcm/
  - Contact details, time/place of lectures, reading list, assignment requirements, etc.
  - Additional materials (lecture notes, etc) will be posted.
- course mailing list: tcm-students@inf.ed.ac.uk.
  - Will be used for important information. You will be added automatically upon registering, but this may take a few days; please register ASAP.

References

Marr, D. (1982). *Vision: A Computational Approach*. Freeman & Co., San Francisco.

Perruchet, P., and Vinter, A. (1998). PARSER: A model for word segmentation. *Journal of Memory and Language*, 39(2), 246-263.

Steyvers, M. and Griffiths, T. (2007). Probabilistic topic models. In T. Landauer, D. S. McNamara, S. Dennis, & W. Kintsch (Eds.), *Handbook of Latent Semantic Analysis.* Hillsdale, NJ: Erlbaum.