Computational Cognitive Science

Lecture 1: Introduction

Chris Lucas
(Slides adapted from Frank Keller's)

School of Informatics
University of Edinburgh
clucas2@inf.ed.ac.uk

20 September, 2022

1. Introduction
   - Models and Theories
   - Models in Cognitive Science
2. Types of Models
   - Data Description
   - Explanatory models
3. Course Overview

Reading: Chapter 1 of Farrell and Lewandowsky.

Cognitive Science

The aim of cognitive science is to understand how the mind works.

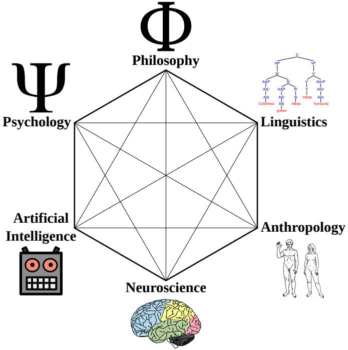

This involves *describing, explaining, and predicting* human behavior.

Computational cognitive science

Analyzing data and forming verbal theories is not sufficient, we need *quantitative mathematical models*.

Models and Theories

Example:

Planets in the night sky move back and forth in loops.

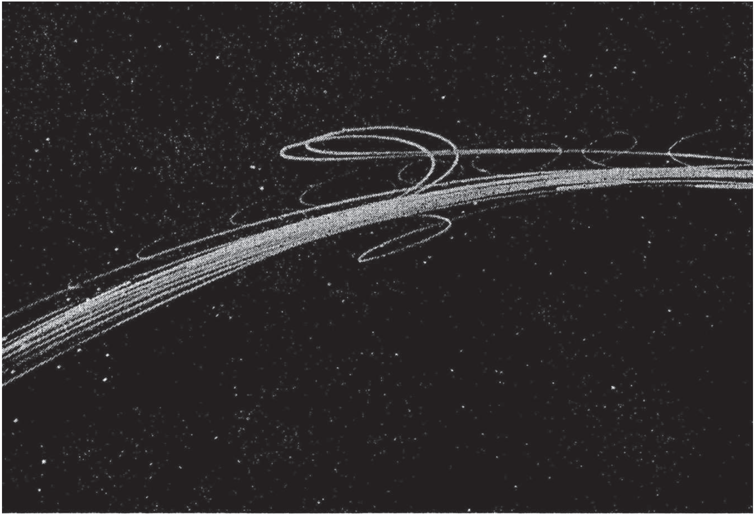

Models and Theories

Observation: *retrograde motion of planets*[1]

- this observation is hard to explain (or even to describe) without a model;
- the model itself (even though it may explain the data) is an unobservable, abstract device;
- there are always several possible models that explain the data.

Competing models of planetary motion:

- *Ptolemaic:* planets move around the earth in deferents and epicycles;
- *Copernican:* planets move around the sun in circles.

[1] Explainer video: https://www.youtube.com/watch?v=FtV0PV9MF88

Ptolemaic Model of Planetary Motion

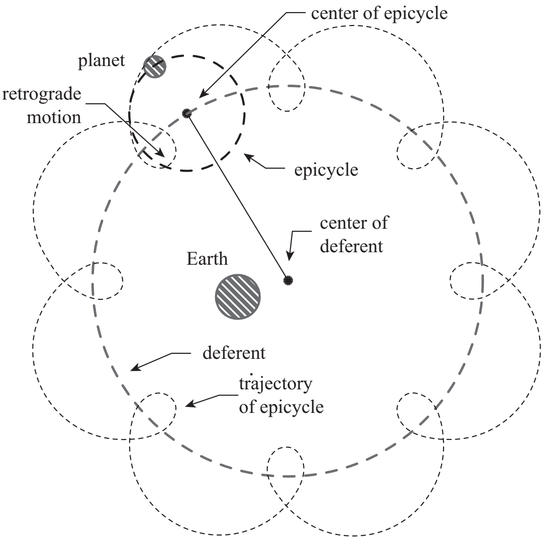

Deciding between Models

Ptolemaic (geocentric) vs. Copernican (heliocentric) model:

- both predict the position of the planets to within 1° accuracy;
- Copernican model predicts latitude slightly better;
- but its main advantage is elegance and *simplicity*, not *goodness of fit* to the data.

We will discuss ways to formalize simplicity later.

Deciding between Models

Simpler models can also be stepping stones to other theoretical advances:

- Kepler's laws of planetary motion replace the circles in the Copernican model with ellipses (of different eccentricities);
- this small modification achieves near-perfect fit with the data.

We'll discuss model comparison in later lectures.

Models in Cognitive Science

Categorization experiment (Nosofsky, 1991):

- training: participants classify cartoon faces into two categories;
- transfer: participants see a larger set, both faces they've seen before and new ones;
- they need to classify the face, say how confident they are, and whether they've seen it before.

Categorization experiment

Example instances (Nosofsky, 1991):

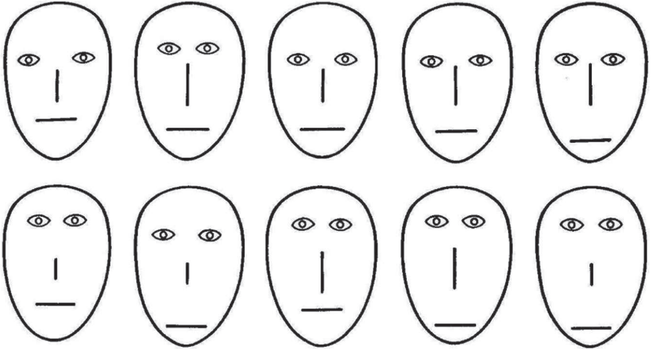

Models in Cognitive Science

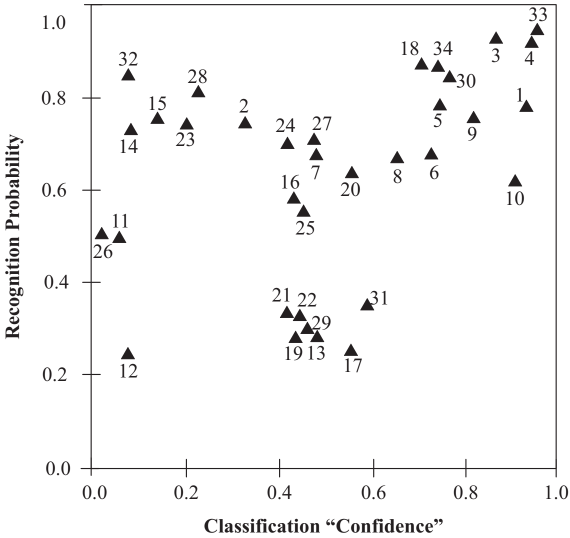

Models in Cognitive Science

No strong relationship between classification and recognition.

Can we conclude that whether you confidently classify a face doesn't depend on whether you remember it?

No, there is a cognitive model (the GCM, details below), which relates classification and recognition and predicts both accurately.

The data don't speak for themselves, but require a quantitative model to be described and explained.

Models in Cognitive Science

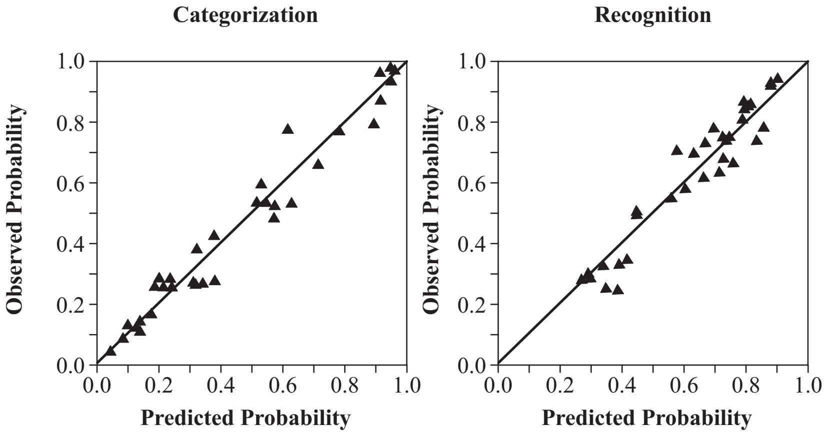

Types of Models

A model is supposed to describe existing data, predict new observations, and provide an explanation for the relevant behavior.

Farrell and Lewandowsky divide models into two kinds:

- *data descriptions:* summarize the data in mathematical form, typically involving parameters estimated from the data;
- *process models:* make commitments about the underlying processes and/or mental representations. Model parameters and features have psychological interpretations.

Another taxonomy: Marr's levels

There are other ways to classify models. One of the best-known is due to David Marr (1982):

- **Goal (computational) level**: What is the organism trying to achieve? How would an ideal or rational agent solve the problem?
- **Process (algorithmic) level**: What algorithm is implementing that solution?
- **Implementation level**: How is the algorithm implemented physically?

Data Description

Example:
The relationship between the amount of practice and the response time in a learning task can be described by a *power law* function:

$$RT = b_0 + b_1(N + 1)^{-\beta}$$

An alternative model is in terms of an *exponential function:*

$$RT = b_0 + b_1 e^{-\alpha N}$$

where RT is the response time, N is the number of trials, and α and β are learning rates.

Data Description

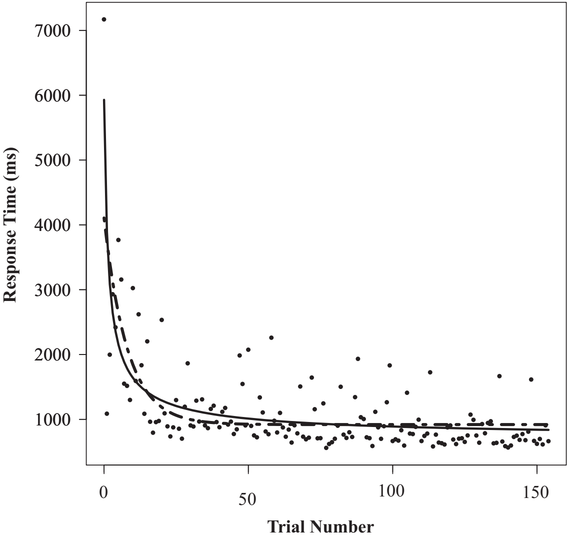

Both models provide a good fit to the data (dashed line: power law; solid line: exponential function).

Ways to decide between them:

- *goodness of fit:* recent work shows that the exponential function provides a better fit to the data on learning;
- *empirical predictions:* the mathematical form of the power law implies that the learning rate decreases with increasing practice; the exponential function implies it stays constant.

Ideally, however, we want to tie the parameters in the model to psychological processes.

Cognitive process models

We want more than a mathematical description of the data.
Proper models ("cognitive process models"):

- **explain** and **predict** cognition and behavior;
- have psychological content – their elements can be interpreted in psychological terms.

Cognitive process models

Example: *Generalized Context Model* (GCM; Nosofsky, 1986), an exemplar model of categorization:

- during training, the model stores every instance of a category;
- during testing, a new instance activates all stored exemplars depending on similarity;
- response probability depends on the sum of the similarity with each member of the category.

Generalized Context Model

Example instances (Nosofsky, 1991):

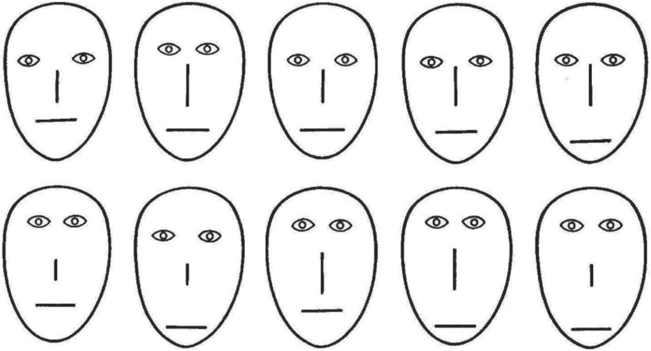

Features: eye height, eye separation, nose length, and mouth height.

Generalized Context Model

The distance d_{ij} between two instances i and j, where each has K features with values x_{ik} and x_{jk}, is:

$$d_{ij} = \left(\sum_{k=1}^{K} |x_{ik} - x_{jk}|^2 \right)^{\frac{1}{2}}$$

The similarity between i and j is (where c is a parameter):

$$s_{ij} = \exp(-c \cdot d_{ij})$$

Then the probability of classifying instance i into category A (rather than category B) is:

$$P(R_i = A|i) = \frac{\sum_{j \in A} s_{ij}}{\sum_{j \in A} s_{ij} + \sum_{j \in B} s_{ij}}$$

The Power of Models

In addition to helping us explain and predict human behavior, models can:

- help *classify phenomena* (e.g., by relating seemingly unrelated data, see categorization vs. recognition);
- help *explore the implications* of a theory (e.g., lesioning a model, scaling to larger data sets, exploring learning).

Course Overview

This course provides an introduction to computational cognitive modeling. There are two main parts:

- introduction to modeling methods;
- discussion of specific models.

We will cover three broad areas of cognition:

- concepts;
- causality;
- active learning.

The textbook is *Farrell and Lewandowsky: Computational Modeling of Cognition and Behavior*. The university has an electronic subscription. This is complemented by papers, see the course resource list.

Required Background

This course requires programming skills. We will use R throughout the course.

The second requirement is maths background:

- probability theory: random variables, distributions, expectations, Bayes theorem;
- linear algebra: basic vector and matrix operations.

If you need a refresher, use *Sharon Goldwater's maths tutorial.*

http://homepages.inf.ed.ac.uk/sgwater/math_tutorials.html

Communication

When you sign up for the course, you will have access to:

- the course mailing list: used for all essential communication;
- the Learn page of the course, used for the assignment and lecture recordings.
- all other material will appear on the course web page: https://groups.inf.ed.ac.uk/teaching/ccs/

There is also a Piazza discussion for the course:

- you can use it to post questions about the course content, including tutorials and assignment;
- the main purpose is **peer support:** students discuss course material and help each other;
- course staff will moderate and contribute

Assessment, Tutorials, Lectures

The assessment on this course will consist of:

- an assessed assignment, worth 40% of the overall mark;
- a final exam worth 60% of the overall mark.

See the course web page for:

- date of assignment and how to submit it;
- plagiarism policy;
- lecture slides, old exams.

There are weekly tutorials for this course:

- tutorials are both practical (use R) and theoretical;
- they start in **Week 3**;
- you will be automatically assigned a tutorial group; if you have a timetable clash, contact the ITO.

Feedback

Feedback students will receive in this course:

- there may one or more non-assessed quizzes – we may try to use Top Hat for these;
- tutorials will be based on non-assessed exercises; you should try to solve these before the tutorials!
- sample solutions will be released for tutorials;
- tutorials include a feed-forward session for the assignment;
- individual assignment comments will be provided by the marker;