Computational Cognitive Science

Lecture 1: Introduction

Frank Keller

21 September 2026

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

Slide credits: Chris Lucas, Benjamin Peters

Overview

Reading: Chapter 1 of Farrell and Lewandowsky (2018).
Optional: Sharon Goldwater's Maths Tutorials if you need a refresher.

Introduction

Cognitive Science

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

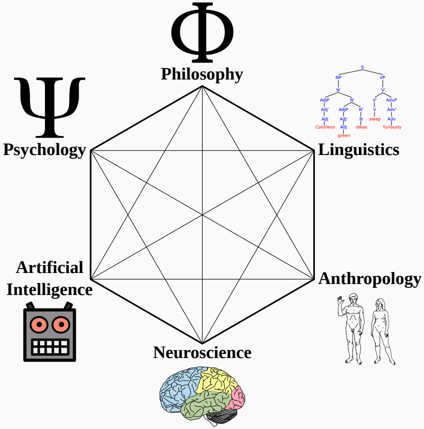

This involves describing, explaining, and predicting human behavior.

Computational cognitive science

Analyzing data and formulating verbal theories are not sufficient; we need quantitative mathematical models.

Models and Theories

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

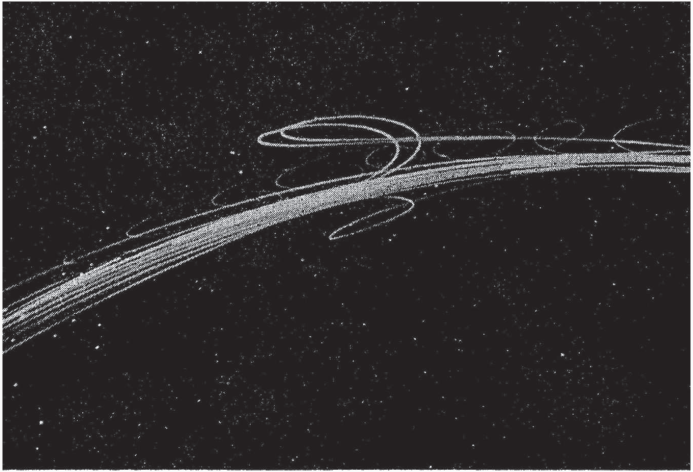

Models and Theories

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] See also this explainer video.

Ptolemaic Model of Planetary Motion

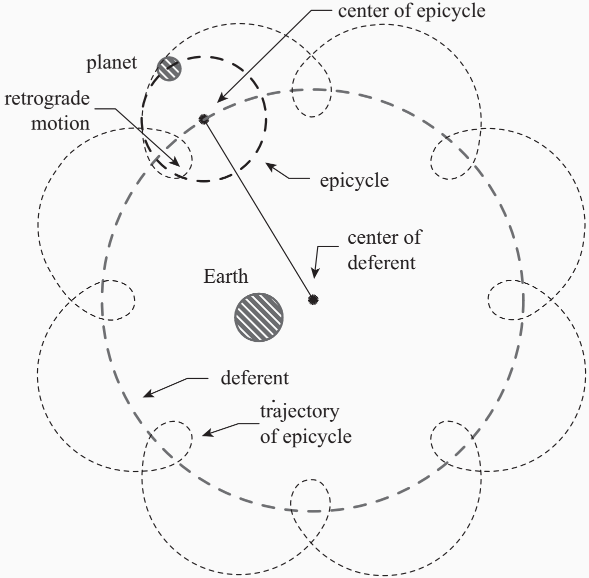

Deciding between Models

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 in the course.

Deciding between Models

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

- Kepler replaces circular orbits with elliptical ones and introduces a law governing orbital speed;
- this achieves a much better fit to the observed data.

We'll discuss model comparison in later lectures.

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):

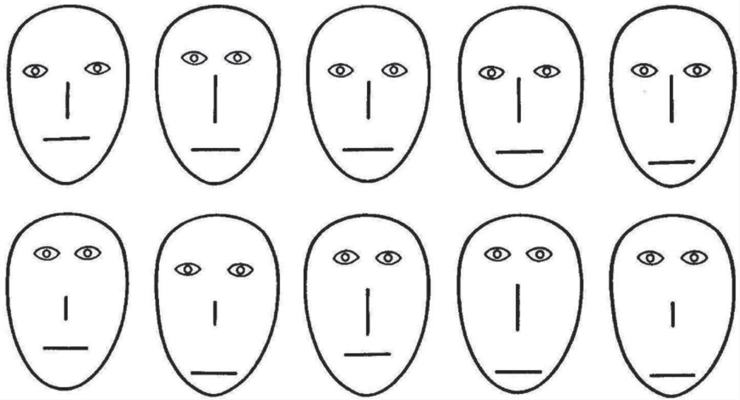

Models in Cognitive Science

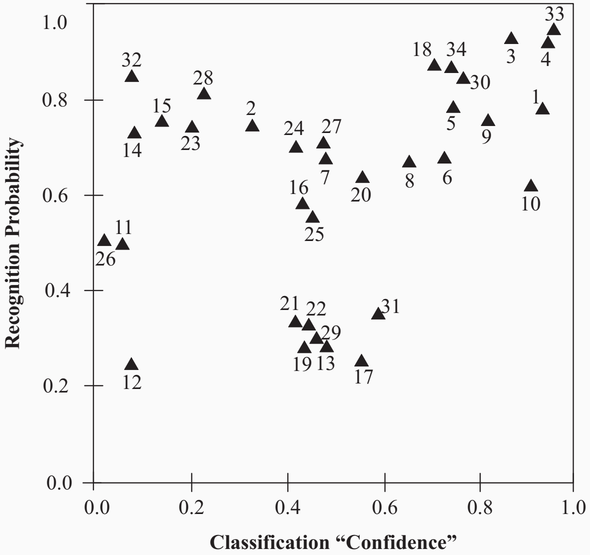

No strong relationship between classification and recognition.

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.

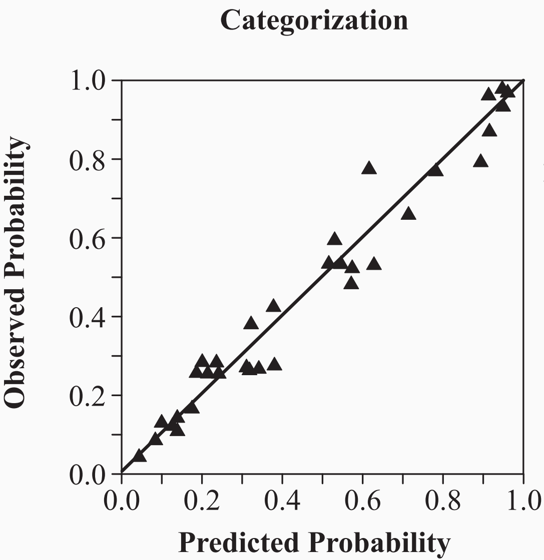
Categorization
1.0
0.8
0.6
0.4
0.2
0.0
Observed Probability
0.0
0.2
0.4
0.6
0.8
1.0
Predicted Probability

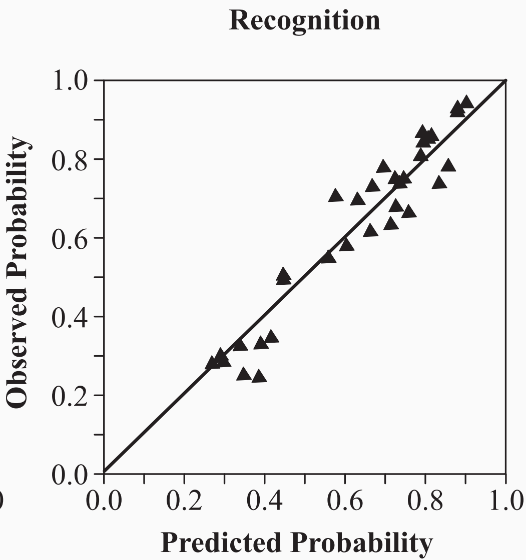
Recognition
1.0
0.8
0.6
0.4
0.2
0.0
Observed Probability
0.0
0.2
0.4
0.6
0.8
1.0
Predicted Probability

Types of Models

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 Marr (1982):

- Computational (goal) level: What problem is being solved, and why is this an appropriate computation for that problem?
- Algorithmic (process) 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 rate parameters.

Data Description

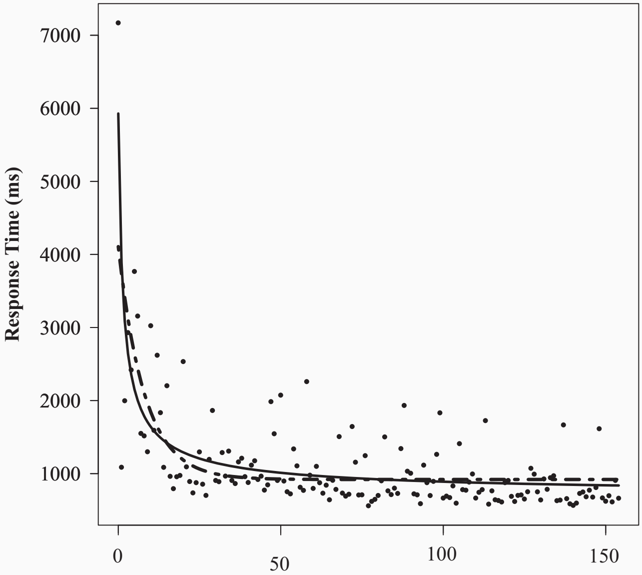

Data Description

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: analyses of individual learning curves often favour an exponential over a power function.
- empirical predictions: for the exponential model, the proportional rate of improvement is constant; for the power model, it decreases with practice.

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 training exemplar together with its category label;
- 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 each member of the category.

Generalized Context Model

Example instances (Nosofsky 1991):

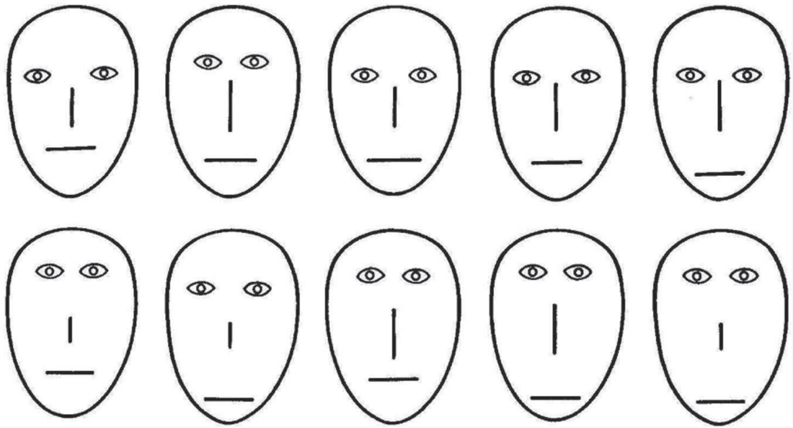

Features: eye height, eye separation, nose length, and mouth height.

A Simplified Version of the GCM

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

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 several broad areas of cognition:

- concepts and categorization;
- causality;
- active learning;
- language acquisition and processing.

As textbook, we will use Farrell and Lewandowsky's (2018) 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. If you already know Python or Java, the CCS R Cheat Sheet will be useful.

The second requirement is some mathematics:

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

If you need a refresher, Sharon Goldwater's Maths Tutorials are a good starting point.

Communication

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

- the Learn page of the course, used for the assignment and for lecture recordings;
- all other course material will appear on the OpenCourse page: `https://opencourse.inf.ed.ac.uk/ccs`;
- any course-related announcements will be posted on Learn.

There is also an EdStem Discussion Forum for the course:

- you can use it to post questions about the course content, including tutorials and the assignment;
- the main purpose is peer support: students discuss course material and help each other;
- the course staff will moderate the discussion 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 Learn page (Coursework Planner) for:

- date of assignment and how to submit it;
- policies on extensions and AI use.

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, request a group change using the MyEd Timetabling Channel.

Feedback

Feedback students will receive in this course:

- 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.

References

Farrell, Simon and Stephan Lewandowsky. 2018. *Computational Modeling of Cognition and Behavior*. Cambridge University Press, Cambridge.

Marr, David. 1982. *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. W. H. Freeman, New York.

Nosofsky, Robert M. 1986. Attention, similarity, and the identification-categorization relationship. *Journal of Experimental Psychology: General* 115(1):39–57.

Nosofsky, Robert M. 1991. Tests of an exemplar model for relating perceptual classification and recognition memory. *Journal of Experimental Psychology: Human Perception and Performance* 17(1):3–27.