

Computational Cognitive Science

Lecture 20: Recap and Q&A

Frank Keller

26 November 2026

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

Slide credits: Benjamin Peters

Reading: [Wilson and Collins \(2019\)](#).

Today:

- Run through the process of developing and using computational models in cognitive science
- Point out connections to past lectures

Steps

0. What is our scientific question or hypothesis?
1. Technical design
 - Experiment
 - Model(s)
 - Analysis
2. Simulation-based checks and updates
3. Run experiments
4. Conduct analyses
5. Follow-ups and extensions
6. Sharing results

0. What is our scientific question or hypothesis?

Examples:

- “Perceptual decision-making is based on a process of sequential evidence accumulation” (Lecture 1)
- “We categorize new things by comparing them to {exemplars, prototypes} from our experience” (Lecture 1)
- “The capacity of visual working memory is limited” (Tutorial 7)

1. Design experiment/model/analysis

Your experiments, models, and analyses are interconnected — design decisions in one have implications for the others.

E.g.,

- Many parameters of interest $\Rightarrow$ many data points.
- Individual differences analysis $\Rightarrow$ many data points per participant.
- Intractable likelihoods $\Rightarrow$ no Bayes factors (probably).

1A. Design experiment(s)

- Good experiments are hard to design.
- Often relies on experience and trial-and-error.
- Bayesian optimal experimental design is one tool to help with this. Finds the design, which maximizes the information gain about the parameters of interest ([Valentin et al. 2021](#): e.g.,)
- Sometimes a simple elaboration on an earlier design.
- Sometimes inspired by watching/reflecting on everyday behavior.
- Have both the model and the analyses in mind; avoid temptation to design complex “fishing expedition” studies.


1B. Design computational model(s)

Desiderata:

- Simplicity is good ([Lecture 7](#)) but don't overdo it
 - “nuisance parameters” like left/right bias can help detect interesting patterns
- Makes precise predictions ([Lecture 2](#))
- Interpretable
- Tractable: Make sure you can fit/compare your models on the computers you have

1B. Design computational model(s)

- Simplicity, precision, interpretability, tractability good desiderata in general.
- Different researchers and domains may weigh these differently, leading to different kinds of models.



1B. Design computational model(s)

There are many kinds of models.

Recall Marr's taxonomy ([Lecture 1](#))

1. Computational-level

- Often frames cognition as Bayesian inference/decision theory
- but not necessarily probabilistic in the F&L sense
- “Rational”, “Ideal-observer” models

1B. Design computational model(s)

2. Process-level

- Often constrained/inspired by computational-level models.
- e.g., approximations of computational-level models
- e.g., Rescorla-Wagner
- e.g., Random-walk
- e.g., Nosofsky's categorization model
- ... most models

1B. Design computational model(s)

3. Implementation-level

- Often constrained/inspired by process-level models.
- “Pure” implementation-level models are rare in higher-level cognition.
- Computational (cognitive) neuroscience, mostly in sensory/motor domains.

1B. Design computational model(s)

New models + baseline/alternative models.

Alternatives are most often

1. “Lesioned” (or: “ablated”) versions of the main model (often nested within main model; [Lecture 7](#))
2. Competing proposals from the scientific literature.
3. A random-guessing baseline.

These should capture salient competing hypotheses.

1C. Plan data analysis

Decide how models will be compared and evaluated **before** collecting data.

1C. Plan model comparison/parameter estimation

- Parameter estimation (Lectures 3-5)
- Model selection (Lectures 7-8)
- Individual differences analyses (Lecture 6)

Always connect analyses to psychological questions/claims.

1C. Plan model-free analyses*

Many hypotheses/theories entail simple and distinctive predictions, e.g.,

- “ $x > y$ ”
- “ $\text{corr}(x,y) > 0$ ”

You should plan to run and report direct tests of these predictions.

- Can help convince skeptical and less technically-sophisticated audiences
- Computationally inexpensive
- Models can fit badly for many reasons; model-free analyses can suffice to make a point

2. Simulate data

Almost all cognitive models are **simulator** models; we can run them through a hypothetical experiment and what the data will look like.

1. Sanity-check our models. Do the data look implausible?
2. Run our planned analyses, using sample sizes from our planned experiment(s).

(Lecture 8)

2. Simulate data

Fake-data simulation + analysis can answer many questions:

- Can we come to meaningful conclusions, regardless of which model is actually true*?
- How often do we pick the true model or make a correct inference about θ ? (power)
- How often do we come to a spurious conclusion?
- Addresses identifiability in theory and practice.

:) → Collect data¹.

:(→ Redesign our models or experiments.

¹but consider pre-registering first!

3. Run experiment

Implementing and running experiments is an important topic in its own right.

- Often learned “on the job”, e.g., honours or MSc project.
- Tools are numerous and change often.
- Not a focus of our course.

4A. Execute analysis plan

This should be straightforward after Steps 1C and 2!

Do not be disheartened by:

- High individual variability, e.g., in what model explains a participant's behavior.
- No model “winning” decisively.
- Unexpected patterns in data.

Aspire to have something useful and true to say, not to explain everything that might be happening.

4B. Validate winning model(s)

Not discussed in the course, but increasingly recognized as important.

The best model of those considered may not be a **good** model.

One way to address this concern: Posterior predictive checks: Sample from posterior of fitted model; compare to data.

5. Follow-ups and (post-hoc) model extensions

Sometimes it is clear where a model went wrong, and a simple change can fix it.

Running additional analyses using modified models is fine — but be honest about it and report both sets of analyses.

6. Disseminating/sharing your results

Report everything that was in your plan (somewhere).

Additional analyses are fine, but should be noted as such.

- Results of fake-data simulation: Confusion matrices and parameter recovery plots.
- Overall model fits, e.g., BIC/crossVal scores, Bayes factors.
- Fits/posterior estimates for parameters.
- Number of individuals best fit by each model.
- Model-free inferential stats.

7. Beyond the current study

- Replicate results with new data (weak generalization).
- The best models have a broad scope, they apply far beyond the original experiment and provide unified explanations of many phenomena.
- Strong form of model generalization: Other domains, contexts, populations.
- Predicting hitherto unknown phenomena is often a mark of generality.

- Kording, Konrad P., Gunnar Blohm, Paul Schrater, and Kendrick Kay. 2020. Appreciating the variety of goals in computational neuroscience. *Neurons, Behavior, Data Analysis, and Theory* 3(6):1–12.
- Valentin, Simon, Steven Kleingesse, Neil R. Bramley, Michael U. Gutmann, and Christopher G. Lucas. 2021. Bayesian optimal experimental design for simulator models of cognition. In *NeurIPS 2021 Workshop on AI for Science: Mind the Gaps*.
- Wilson, Robert C. and Anne G. E. Collins. 2019. Ten simple rules for the computational modeling of behavioral data. *eLife* 8:e49547.