Carney Computational Modelling Workshop (2019 Edition)

Contact: Andra Geana (andra_geana@brown.edu)
Dates: July 15 – July 19 (theoretical tutorials), July 22-July 26 (Modelling Challenge)
Locations: Metcalf Research Building #305 (The Dome Room)

Overview

What is a **computational model**? How do models help us **study the brain**? How do we **build a good model**, or **choose** the best among different models?

This intensive two-week workshop will provide tools for understanding, developing, and applying models to brain science questions. We will focus on three main classes of models – reinforcement learning, Bayesian models and Drift Diffusion Models (DDM)—and their details and applications. Week 1 will consistent of introductory sessions to each class of models, practical tutorial sessions to allow participants to write and test their own modelling code, and advanced sessions providing a deeper understanding of fundamental modelling techniques, pitfalls and concepts.

Week 2 will give everyone the opportunity to model a novel dataset using the methods taught in Week 1. Participants may work individually or together on building their models, and we will all meet once a day for a short coffee break to discuss everyone's progress and challenges.

Prizes: Modelling Challenge winner(s) will receive a $60 gift card to a local restaurant of their choice.

Goals

By the end of this workshop, you should be able to:

- **Understand** Bayesian, RL and DDM, their rationale and underlying architecture
- **Develop** and **code** your own models in MATLAB
- **Fit** models to a dataset, **compare** different models, and choose and **apply** the best one to answer your research question

Prerequisites:

- **YOU HAVE TO INSTALL THE HDDM TOOLBOX PRIOR TO THE THURSDAY SESSION.** Please find the instructions here, and contact Jae-Young Son if you have trouble with the installation.

- Basic coding experience with MATLAB preferred, but not required; you can find tutorials here and here.
- Familiarity with basic statistics preferred, but not required. No previous modelling experience or math background required.

Resources: This is the link to the workshop google drive folder. You will find there some reference papers, a couple of useful guides, and I will upload all relevant slides and MATLAB code here as well. Please email me if you can't access the drive.

Evaluation: attendance to all sessions is required to complete the workshop. No formal evaluation required for Week 1. For Week 2, you will need to submit MATLAB likelihood functions for each of your models. Files should be submitted by email to me, by the agreed-upon deadline.

Detailed WEEK 1 Schedule

Session 1: 09:00 – 10:30 a.m.
Coffee Break: 10:30 – 11:00 a.m.
Session 2: 11:00-12:30 p.m.
Session 3: 1:30 p.m. – 3:00 p.m.

Day 1: Introduction to Computational Modelling (Monday 7/15)

Session 1: What is Computational Modelling and Why Do We Care About It?

1. **What** is modeling? (oh no! equations! what do!)
   a. Problems that models can be used for (Learning! Perception! Memory! Others!)
2. **Why** are models useful?
3. Principles for **building a model**: explanatory power, simplicity, generalizability etc.
4. **Pitfalls**: overfitting, oversimplifying, reductionism, biological plausibility, interpreting parameters
5. How to **pick a model**
   a. What are you trying to model? (e.g. RT vs choice data vs. perception data)
6. Practical aspects: how to **write models** (can use different programs, e.g. Matlab, HDDM, emergent etc.)

[Coffee Break]

Session 2: PRACTICAL TUTORIAL Introduction to Writing Models

1. The lunch decision problem: a **practical example** of computational modelling
2. A simple model for computing option values: rationale and **equations** [link to tutorial1]
3. Choice likelihoods: **choice functions** vs **value functions**
4. Writing model code in **Matlab** (everyone can do this on their own, w/ help from available scripts) [Find available scripts here]

5. **Simulating data** from model, to see it in action
    a. Fitting model to actual data: parameter fitting, potential issues
6. Write function **to fit model to existing data**
    a. Interpreting parameters (i.e. what is a low learning rate vs high learning rate?)
    b. Measures of model goodness-of-fit

[Lunch Break]

Session 3: Introduction to Reinforcement Learning

1. What is **RL**? (examples, the state/action/reward diagram, how RL helps us)
2. **Value-learning** through reinforcement [Reading: Daw chapter]
3. **RL in the brain** (links to DA system) [Reading: Frank review]
4. Limitations of **RL**

Day 2: Reinforcement Learning Models (Tuesday 7/16)

Session 1: Challenges and Pitfalls of Reinforcement Learning models

Guest lecture by MJ Frank [find slides here]

Session 2: PRACTICAL TUTORIAL Designing, writing, and fitting an RL model to choice data

1. Memory-learning-perception task: introduction to available data
2. Write simple learning model (based on rules learned in Day 1 and Day 2, Session 1)
    a. Modeling choices: value initialization, selecting learning rate (static/dynamic? One rate vs different rates for win/loss? Etc.)
    b. Checking predictions: know your model mechanics!
3. Fitting model to data, interpreting parameters
4. Model iterations: change model (e.g. add second learning rate) and refit.
5. Model comparison

[coffee break]

Session 3: Advanced RL Concepts

1. More examples of underlying RL architectures: actor-critic vs Q-learning
2. Sources of Uncertainty and what to do about them
3. Partially Observable Markov Decision Processes – when do we want to use them and why?
    a. Incorporating state uncertainty into the learning problem
    b. A practical POMDP example

Day 3: Bayesian Models (Wednesday 7/17)

Session 1: Introduction to Bayes Theory and why we like it

1. Introduction to **probability theory**, Bayes theorem
2. **Optimality** and its caveats
3. An illustration of a **Bayesian inference model**: category learning
4. **Limitations** of Bayesian models:

[coffee break]

Session 2: PRACTICAL TUTORIAL Tricks and Challenges of Using a Reduced Bayesian Model for Learning Data

Guest lecture by M Nassar [find slides here] [find scripts here]

[lunch break]

Session 3: Hierarchical Bayesian Parameter Estimation [Reading: Daw chapter]

Guest lecture by MJ Frank

Day 4: Drift Diffusion Models (Thursday 7/18)

Session 1: Introduction to Drift Diffusion Models (DDM)

1. **Evidence accumulation**: the basic mechanics of DDM
2. Within-trial effects vs across-trial integration. **Why is DDM useful**?
3. **Reaction times** (RT) distributions modelled with DDM
4. DDM in **brain science**: perception, categorization examples

[coffee break]

Session 2: PRACTICAL TUTORIAL Introduction to writing and fitting DDM models

Guest lecture by Olga Lositsky [find scripts here]

[lunch break]

Session 3: **A Tutorial on HDDM toolbox (DDM with hierarchical estimation)**

Guest lecture by Jae-Young Son [find toolbox instructions here]

Day 5: Methods & Challenges in Computational Modelling (Fri 7/19)

Session 1: An Overview of different computational models

1. **Hybrid models**: RL, Bayes, DDM combined may work best
2. Discussion of **which types of models may be best fit to which kinds of data**
3. **Biological implementation** concerns
4. A summary of the **benefits of computational modelling**

Session 2 (tentative): Using Neural Networks to Approximate Likelihood Functions

Guest Lecture by Alex Fengler

Session 3: Open discussion to questions from audience

Please submit your questions no later than Thursday 5 p.m.

Day 6: Introducing the Modeling Challenge! (Monday 7/22)

Session 1 (10:00 a.m-11:00 a.m.): The Data to Be Modeled