Computational Neuroscience of Vision

Dr. James A. Bednar

jbednar@inf.ed.ac.uk

http://homepages.inf.ed.ac.uk/jbednar

Welcome to CNV!

CNV is about using computational simulations to understand how biological visual systems work.

The course is geared towards people who are interested in and may be considering doing further research in computational neuroscience.

It may also be valuable as background for work in computer vision or machine learning, but those more–practical areas are not the focus.

Why study vision?

- Early stages are relatively well understood
- Easy to control stimuli
- Large percentage of brain
- Standard test case for understanding the brain

We will focus on animals whose visual systems are similar to a human's (typically monkey, ferret, and cat).

Questions about the visual system

- How does it work?
  I.e., what algorithm(s) does it perform?
- What physical parts implement the algorithm(s)?
- What does each part do?
  I.e., what do they contribute to the overall algorithm(s)?
- How are the parts and their connections constructed?
  (From a blueprint? By learning?)
- How specific is the algorithm to vision in particular?

Computational models

A computational model represents a concrete implementation of a theory of how (part of) the visual system works.

Computational models integrate information from many previous experiments of different types.

Ideally, computational models will be designed to answer specific questions that arise from either existing experiments or theoretical considerations.

What to include in models

Models should be constructed to answer specific scientific questions.

The models should then include the features of the nervous system and its context that are most relevant for the questions being asked, modeled at an appropriate level of abstraction and detail.

Everything else should be omitted.

Suitable vision models

Vision in mammals requires large areas of brain tissue, composed of many billions of neurons and many trillions of connections.

Representations of the visual field cover a large fraction of the brain (nearly half in monkeys).

Disabling any single neuron, and probably a very large fraction of scattered neurons, has no significant effect.

Thus, for understanding vision I believe it is crucial to model large numbers of neurons, rather than modeling single neurons or parts of neurons extremely well.

Types of models in this course

This course will thus focus on modeling the visual system at a level large enough to see interesting behavior, simulating regions at least several square mm in size (large enough to represent recognizable visual features, and matching data available from optical imaging).

I will generally draw examples from my own research and related research, because those are the ones that I can cover most authoritatively (and provide working code for!). At the same time, the course readings, the tools we will use, and the overall approaches are very general, and cover a very large range of other possible models.

Example model

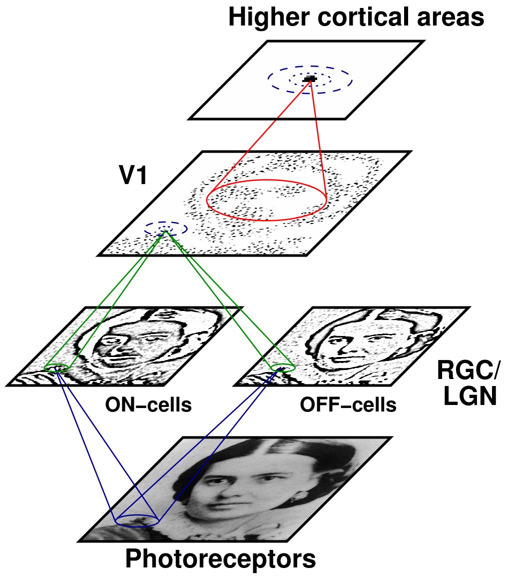

GCAL/LISSOM: Model photoreceptor responses, subcortical processing, and the development and function of the visual cortex

Course topics

Biological background: Visual systems in animals and humans

Computational modeling levels: Single-unit models, topographic map models

Modeling early visual system development: How brain areas organize

Modeling adult low-level processing: Edge detection, orientation processing, etc.

Higher-level processing: Face and object processing

Course components

The course will consist of lectures, readings, coursework, and exams.

The coursework and exams assume a basic level of mathematical and programming competence, but the focus is on the biology throughout.

Prior experience with Python is helpful, but not in any way required.

Topographica simulator

Course exercises will use the Topographica simulator, which we have built from general-purpose open-source Python packages for scientific computing.

These packages provide comprehensive support for visualizing and analysing your results in a reproducible way, and learning about these will be very helpful for future projects in any area of science or research.

Contributions to any of our packages are welcome, in the form of pull requests submitted at github.com/ioam.

Text

The text for the course is our book *Computational Maps in the Visual Cortex* (Springer, 2005), obtainable by:

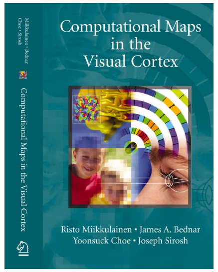

- Ordering from Springer or Amazon (115 GBP!)
- George Sq. library (including 1 on reserve)
- DTC student offices (2 copies)

The book is a synthesis of research from four authors using closely related models to study the visual system. All of the simulations in the book are my own except for those in chapters 3 (SOM) and 11-14 (perceptual grouping).

Ideal model? (1)

To address a possible confusion head-on: One might imagine that the ideal computational model would include all of the properties of each part of the visual system and its environment, in full detail. This would require simulating:

- many billions of nerve cells and connections
- morphology of each cell
- chemical context surrounding cells
- ionic currents
- specific ion channels

Ideal model? (2)

And also:

- quantum effects
- developing from a single cell
- sensory environment during development
- all previous visual experience
- sensory environment during testing
- behavioral task being performed
- all parts of the brain that connect to the visual system

Ideal model? (3)

And more! Obviously, not all of these can be included in any single model:

- We don't have sufficient computing power
- Only a tiny fraction of data necessary to model all these things has been measured
- We don't understand how the components that are connected to the visual system work (i.e., the rest of the nervous system)

Ideal model? (4)

Most importantly, such an "ideal" model would be entirely incomprehensible, and would not actually show what causes any particular observed behaviour.

One could just as well say that the animal and its environment are their own ideal model, because nothing has been gained or lost by using the model rather than the animal.

We will instead look at *useful* models, which have to leave out most of the properties of the "ideal" model, simplifying and abstracting them to focus on a few important properties.

Summary

- Vision is an important test case for the brain
- Computational models at the right level can be used to understand vision
- We will study the basic circuitry and function of early vision, and how they can be investigated using simulations
- Few of the techniques are specific to vision, and thus they can also be applied to other systems
- Required reading: *CMVC book* chapter 1, and sections 2–4 of *Report of the 1st INCF Workshop on Large-scale Modeling of the Nervous System*