COSC 421/521

Computational Cognitive Neuroscience

Bruce MacLennan

Spring 2020

Course Information

- Instructor: Bruce MacLennan [he/his/him]
  - Teaching Assistant: Michael Price (mprice35@vols.utk.edu)
- Course website: web.eecs.utk.edu/~bmaclenn/Classes/421-521
- Email: maclennan@utk.edu
- Office hours: 2:30–3:30 WF, most M (or make appt.)
- Prereqs: no specific prereqs, but will be taught at senior/graduate level
- Grading: weekly homework, occasional pop quizzes
- Piazza for discussions

About the Course

- A course in <u>computational</u> <u>cognitive</u> <u>neuroscience</u>
- Intended for computer science and neuroscience majors
- Focus on cognitive processes (including perception, categorization, memory, language, action, and executive control)
- Understanding neural implementation of these processes
- Using computer simulations to model processes and to test hypotheses

Value for Computer Science Students

- Important if interested in artificial intelligence, neural networks, or neuromorphic computing
- Will help you understand how brains do things that are still difficult for computers
- You will be able to take the concepts and theories of neural information processing and use them to develop better AI systems
- You will learn about neuroscience applications of computer modeling

Tricking AI to think a banana is a toaster

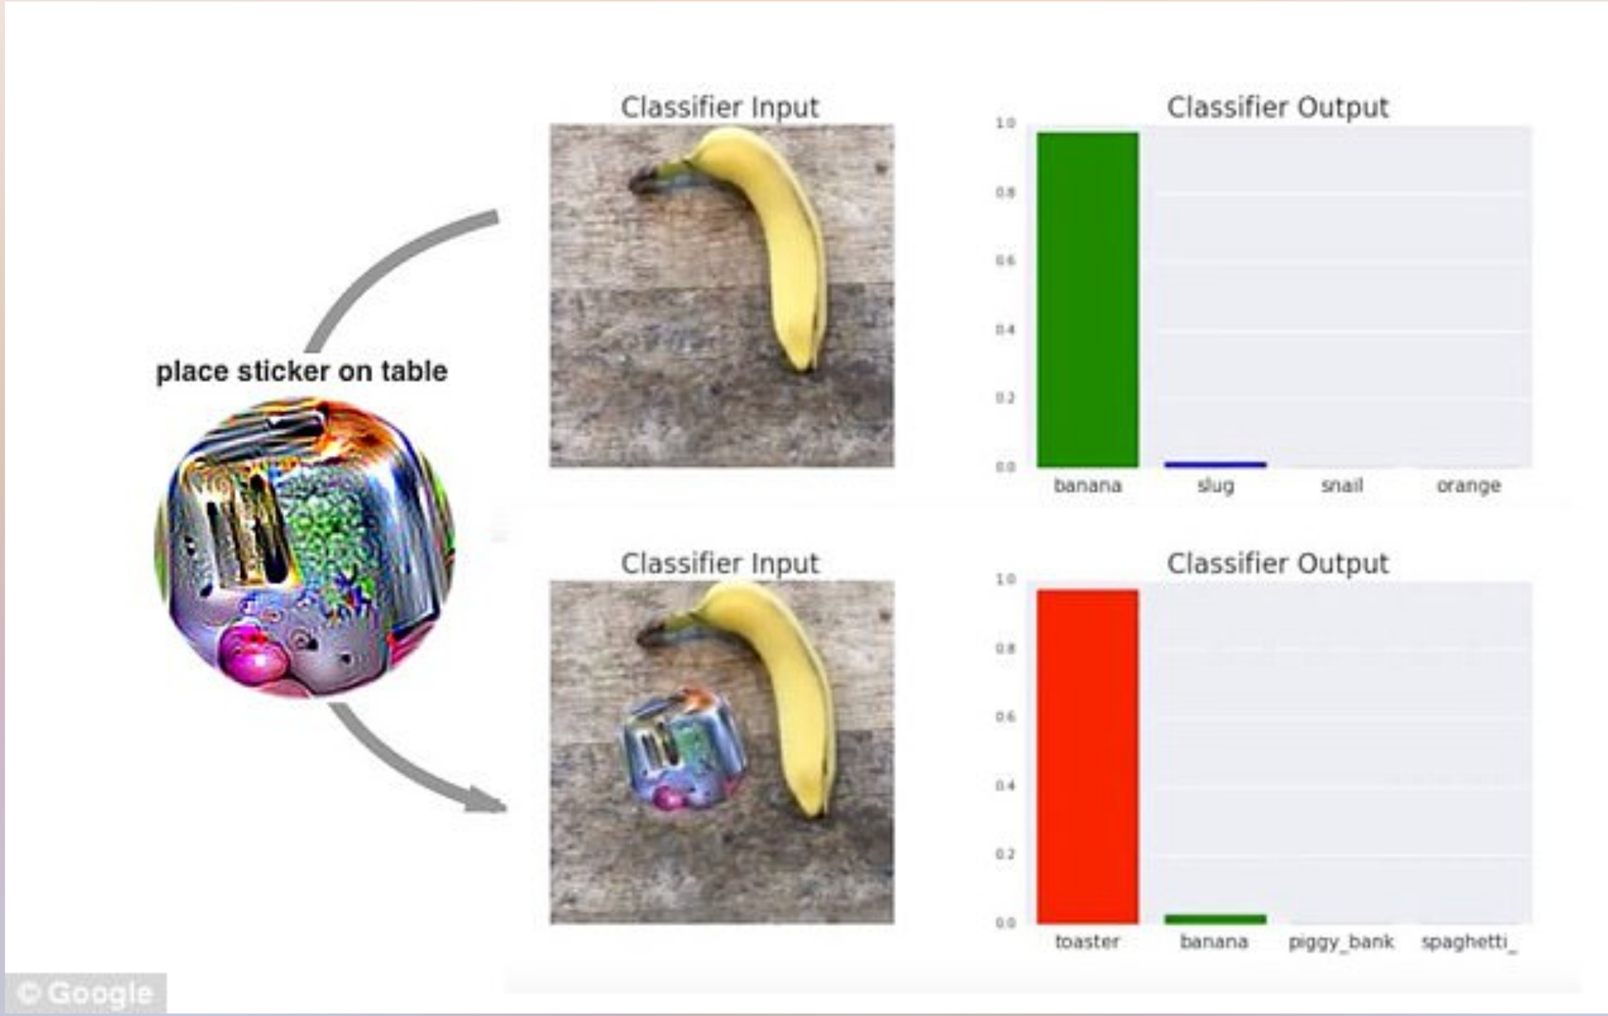

Value for Neuroscience Students

- You will learn how computers can be applied to modeling the neural processes underlying cognition
- You will get hands-on experience using these tools
  - reinforce your neuroscience knowledge
  - give you a deeper understanding of neural information processing in the brain
- You will learn how these processes can be implemented on computers in order to achieve artificial intelligence

Prerequisites

- This course is intended for:
  - Computer science students with no experience in neuroscience
  - Neuroscience students with no experience in computer modeling
- It is intended to be interdisciplinary and self-contained
- Therefore, no specific prerequisites
- No mathematics beyond elementary calculus
- Course will be taught at a level appropriate for seniors and graduate students

Practical Course Prerequisite

- You will be required to run models on the (free) emergent software system
- You can install it on your own computer (Mac, Windows, or linux)
- You can use a friend's installation
- However, if you cannot find a way to run it within a week, *you should drop the course!*

Suggested Installation Strategy

- Before the next class, attempt to install emergent on your computer of choice
- Make sure the neuron and detector projects run
  - I will demo them shortly
- If you have had difficulties, we will try to help you out on Friday
- If you do not have a workable solution by next Wednesday, we will need to consider your options

Assignments

- Readings in free, online text: *Computational Cognitive Neuroscience* by Randall O'Reilly et al.
- Weekly observations from experiments run on emergent system
- Weekly "reading reflections": paragraph on most interesting ideas from readings
- Occasional pop quizzes
- Term papers and presentations by students in COSC 521

Demonstration of emergent

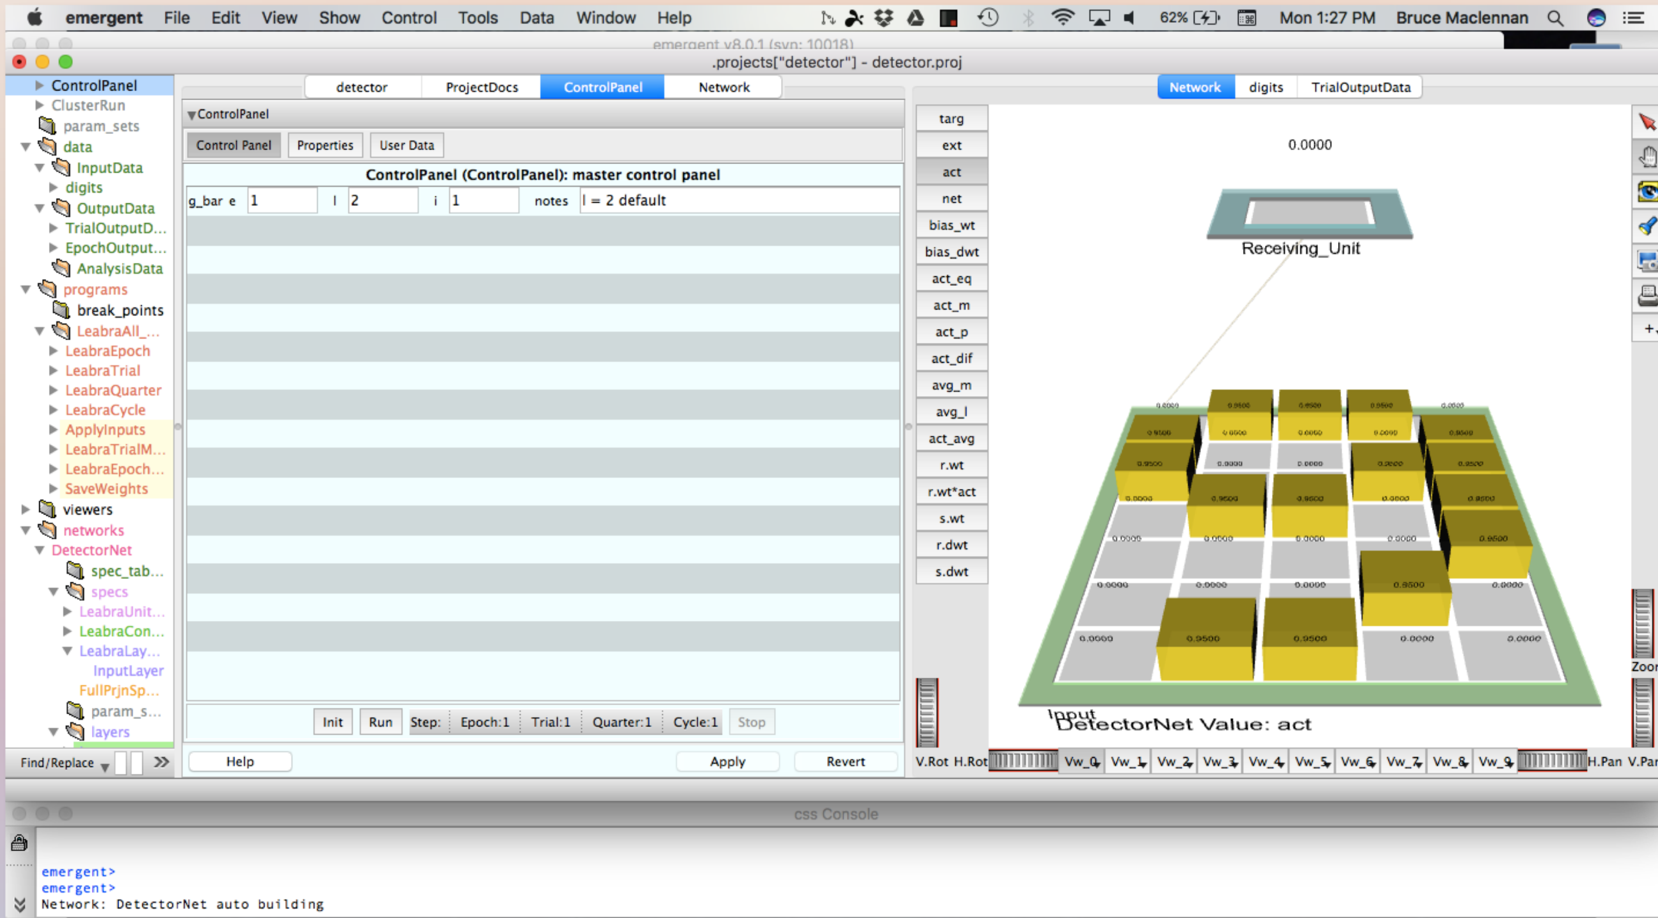

Summary of Steps to Check emergent

1. Go to text and click on chapter 2. Neuron
2. Scroll to bottom and click on Detector (not New Detector!)
3. In the right-hand frame, click File:detector.proj
4. Control- or right-click detector.proj to download it to some convenient place
5. Launch emergent
6. Under File menu, click Open Project and open detector.proj
7. In upper border, click ControlPanel
8. In lower border, click Init
9. In lower border, click Run and you should see neurons updating in RH frame
10. Quit emergent. You're up and running!

From Neurons to Behavior and Back Again

The Challenge

- The brain is massively parallel
  - 86 billion neurons in human brain
  - 20 billion neurons in neocortex
- The brain is massively interconnected
  - each neuron gets inputs from thousands of others
  - 100–1000 trillion connections
- Neurons are slow
  - take milliseconds to respond
- Yet brain responds in real time
  - 100 step rule

The brain is organized over sizes that span 6 orders of magnitude

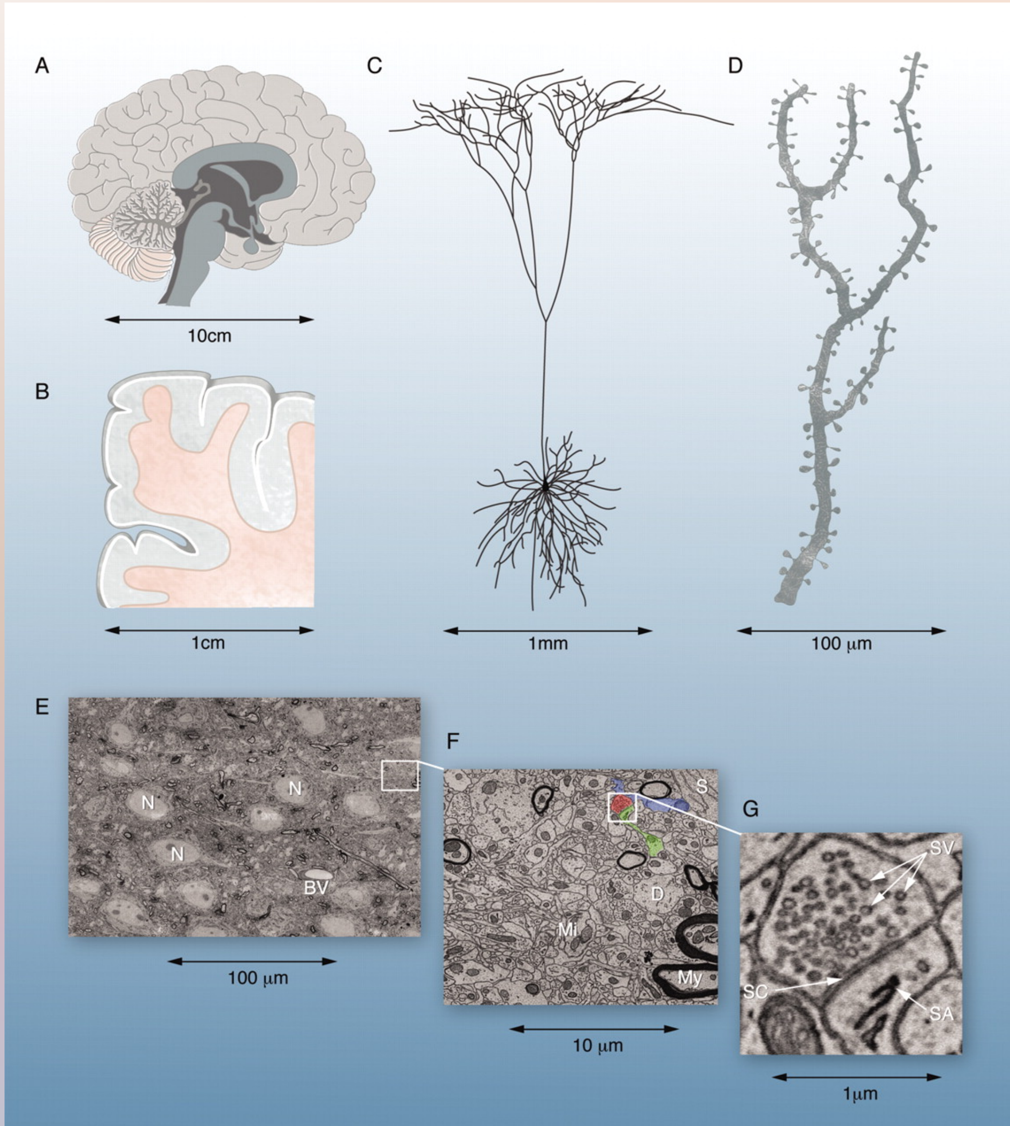

Levels of Investigation

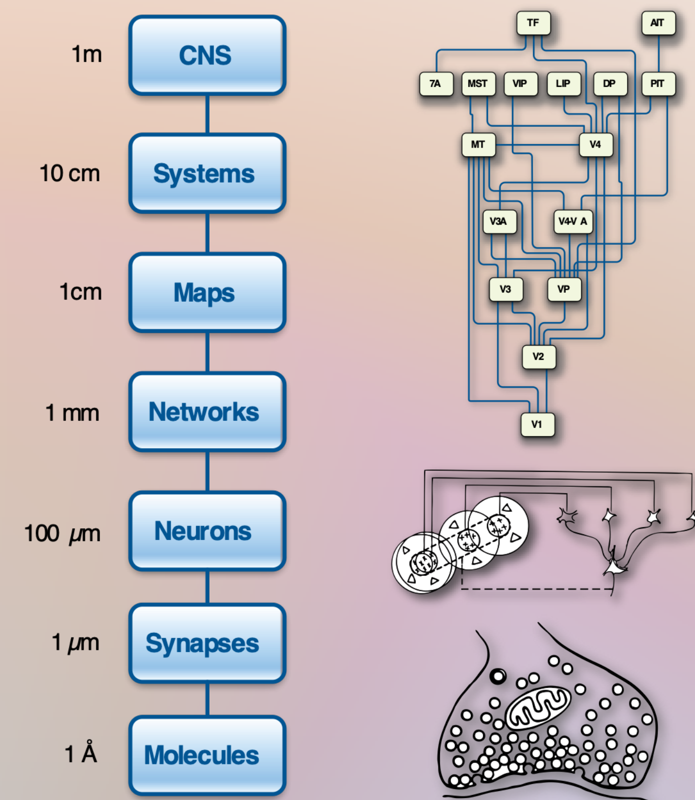

(figure < Sejnowski, Salk)

J W Lichtman, W Denk
Science 2011;334:618-623

Published by AAAS

Overview of Brain to Neurons

<http://www.youtube.com/watch?v=DF04XPBj5uc>

(play flash video)

Neural Density in Cortex

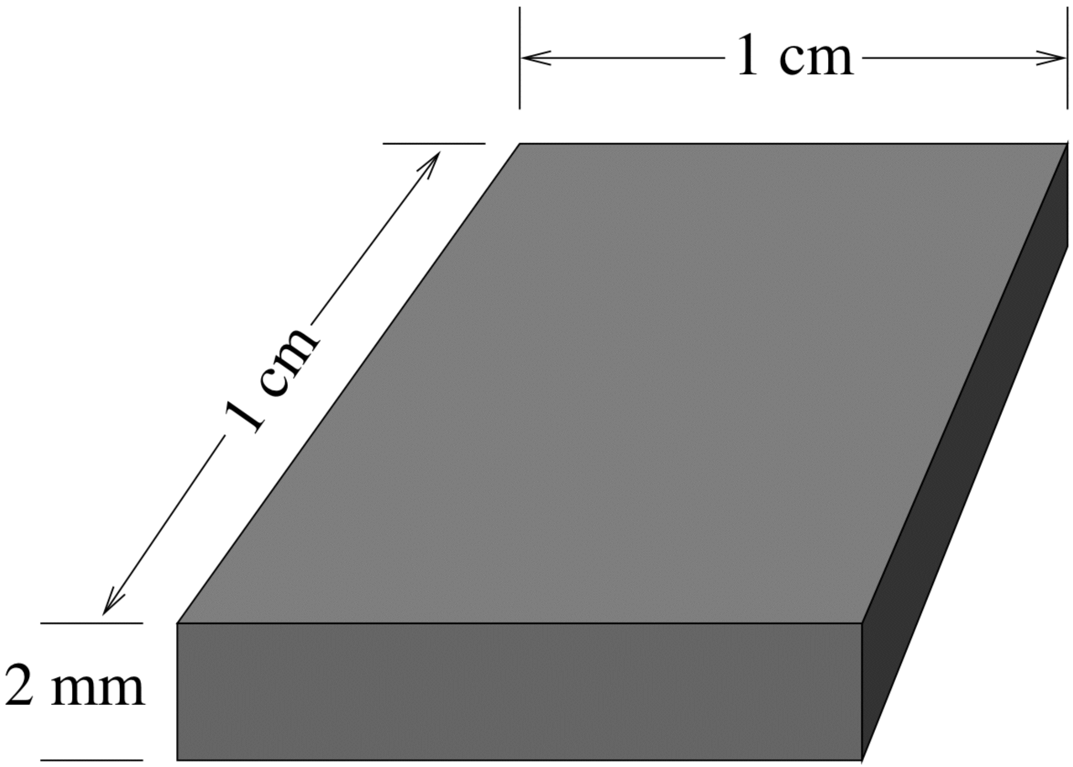

127 000 neurons / sq. mm

Hence, about 13 million / sq. cm

Cortical Areas

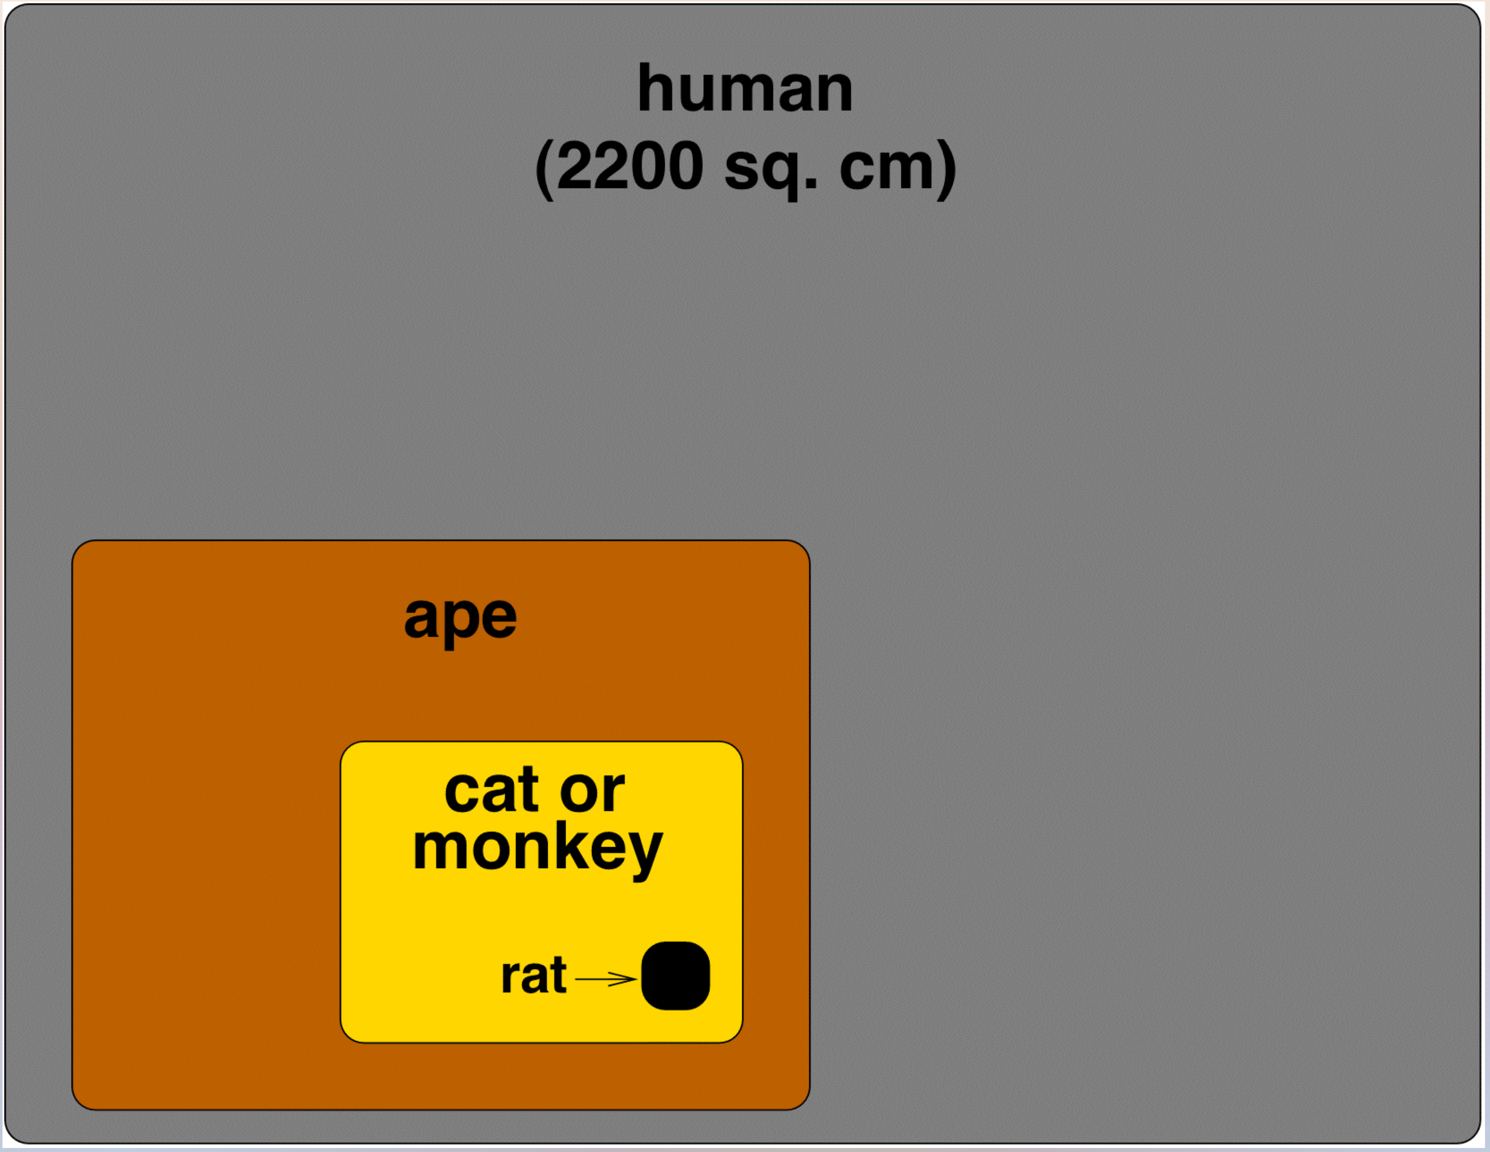

Reductionism and Reconstructionism

- Reductionism
  - Understanding something in terms of its parts
  - Understanding high-level processes in terms of lower-level processes
  - Understanding behavior in terms of neurons
- Reconstructionism
  - Complementary: putting the pieces back together
  - “What I cannot create, I do not understand” (Feynman)

Feynman’s blackboard when he died

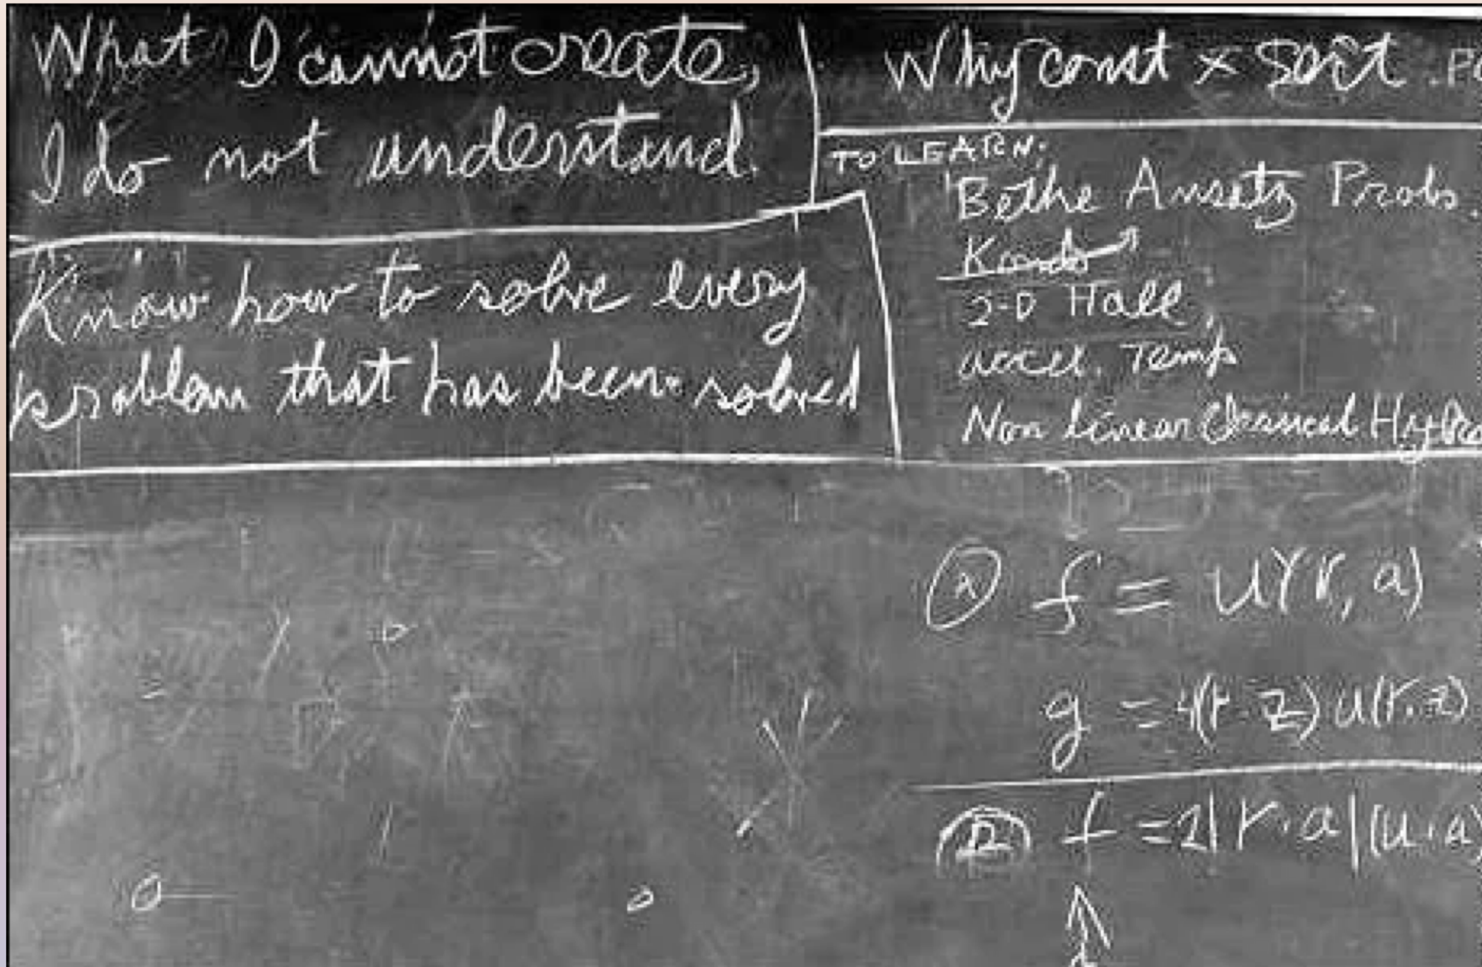

Emergence and Learning

- With 86 billion neurons, you can't build it by hand
- It basically has to build itself (through development & learning)
- Complexity must emerge from simplicity (not that many genes control brain development)

(based on O'Reilly slide)

Emergence

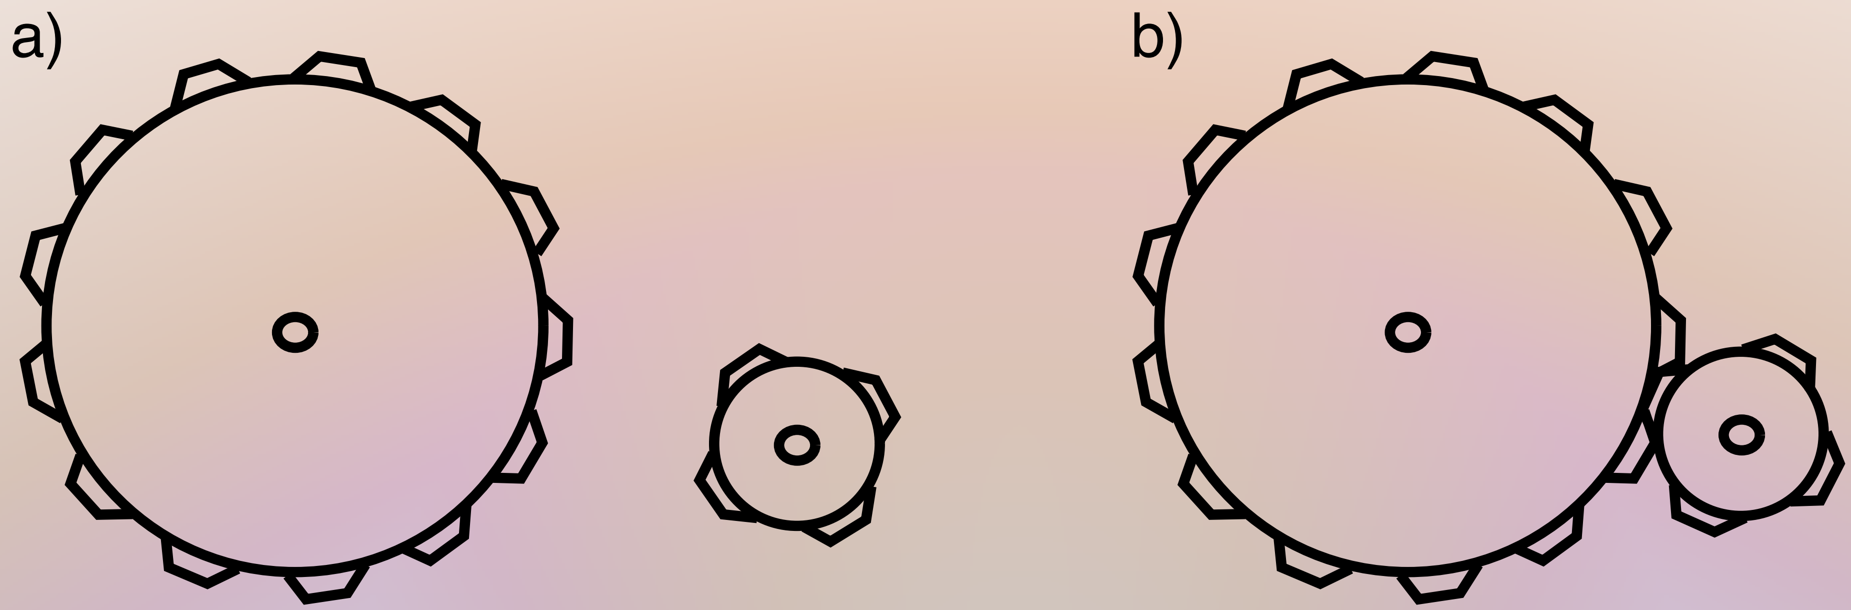

(Now Imagine 10,000,000,000 gears, each interacting with 10,000 others)

Why Computational Models?

- Brain is a computing device (information processor)
- Computational models allow us to described models in a precise way

  — "What I understand, I can program"
- Abstract and formal theory can help us organize and interpret data
- Computational models are a path to AI

Marr's Levels of Abstraction

- Computational (*goal*)
  - What computations are being performed? What information is being processed?
- Algorithmic (*strategy*)
  - How are these computations being performed, in terms of a sequence of information processing steps?
- Implementational (*representation*)
  - How does the hardware actually implement these algorithms?

Why We Can't Ignore Implementation

- Traditional view: we can ignore implementation level because all computers are functionally the same
  — Von Neumann architecture
  — Universal Turing machine
- But the brain has a radically different architecture
  — Low-precision analog devices
  — Massively parallel
  — Massively interconnected

Some Levels of Computational Model

- High-level/symbolic cognitive processes
- Bayesian inference
- Artificial neural networks (BP, deep learning)
- Biologically realistic rate-based models (Leabra)
- Spiking neuron models (integrate and fire)
- Compartment models
- Microphysiological

Course Overview

From neurons to networks to the brain/mind...

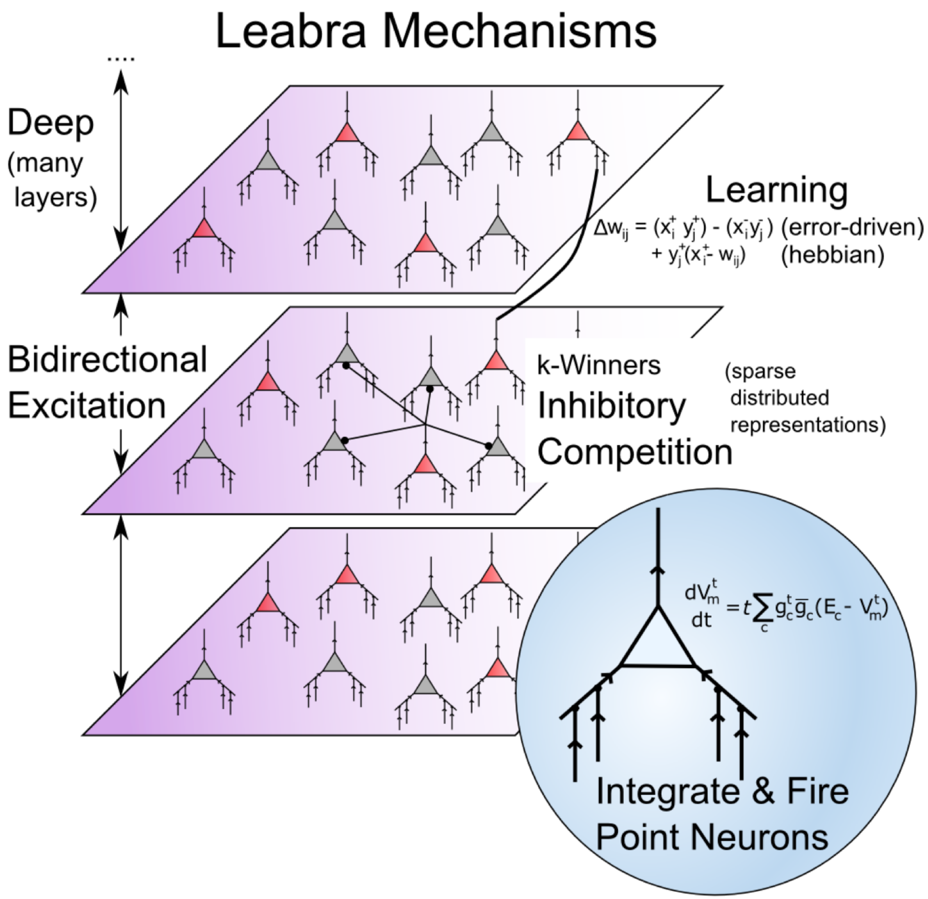

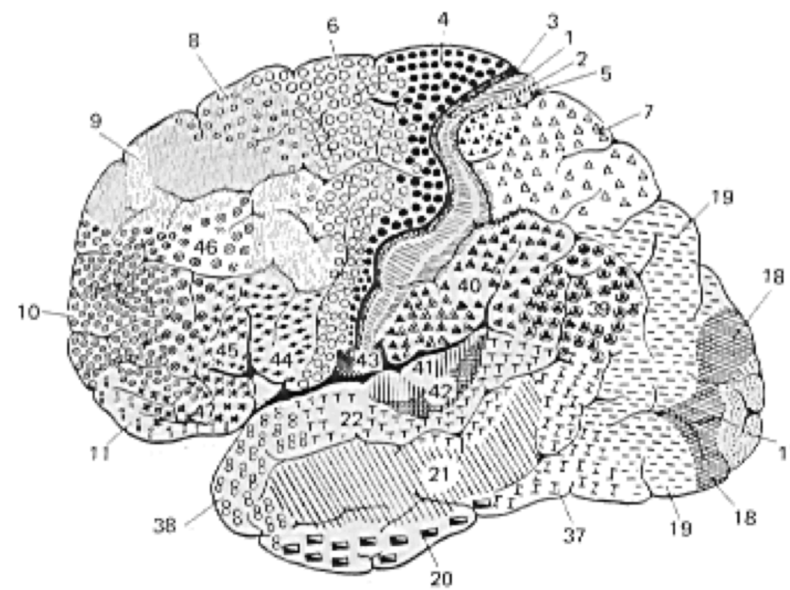

Basic Computational Mechanisms

- Neurons
  — serve as *detectors*, signal with *activity*
- Networks
  — *link*, *coordinate*, *amplify*, and *select* patterns of activity over neurons
- Learning
  — *organizes networks* to perform tasks and develop models of environment

Cognitive Phenomena (1)

- **Visual encoding:** A network views natural scenes (mountains, trees, etc.), and develops brain-like ways of encoding them using principles of learning.
- **Spatial attention:** Taking advantage of interactions between two different streams of visual processing, a model focuses its attention in different locations in space, and simulates normal and brain-damaged people.
- **Episodic memory:** Replicating the structure of the hippocampus, a model forms new episodic memories and solves human memory tasks.
- **Working memory:** A neural network with specialized biological mechanisms simulates our working memory capacities (e.g., the ability to mentally juggle a bunch of numbers while trying to multiply multidigit values).

(slide < O’Reilly)

Cognitive Phenomena (2)

- **Word reading:** A network learns to read and pronounce nearly 3,000 English words, and generalizes to novel non-words (e.g., “mave” or “nust”) just as people do. Damaging a reading model simulates various forms of dyslexia.
- **Semantic representation:** A network “reads” every paragraph in a textbook, acquiring a surprisingly good semantic understanding by noting which words tend to be used together or in similar contexts.
- **Task directed behavior:** A network simulates the “executive” part of the brain, the prefrontal cortex, which keeps us focused on performing the task at hand and protects us from distraction.

(slide < O’Reilly)