
Computational Neuroscience

2017



Instructor: Odelia Schwartz

Computational neuroscience


- **Introductions...**

What area are you in / background?

Why Computational Neuroscience?

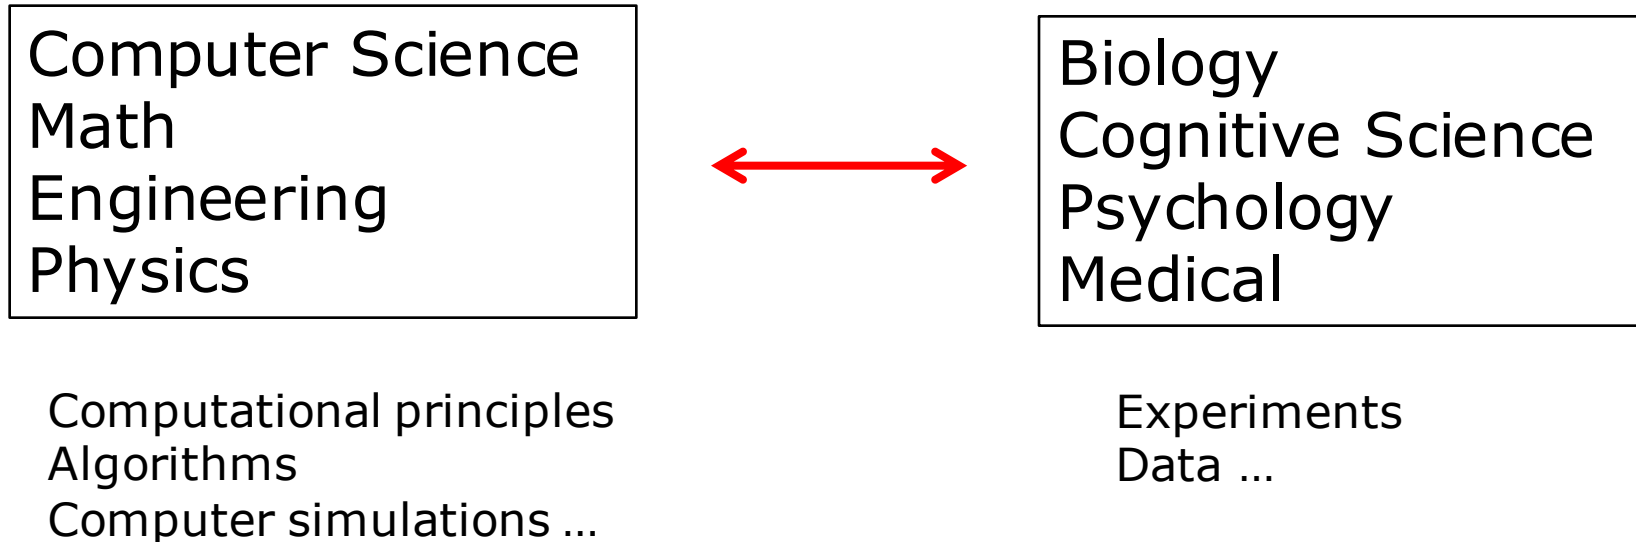
Computational neuroscience

Lots of interest from multiple fields!



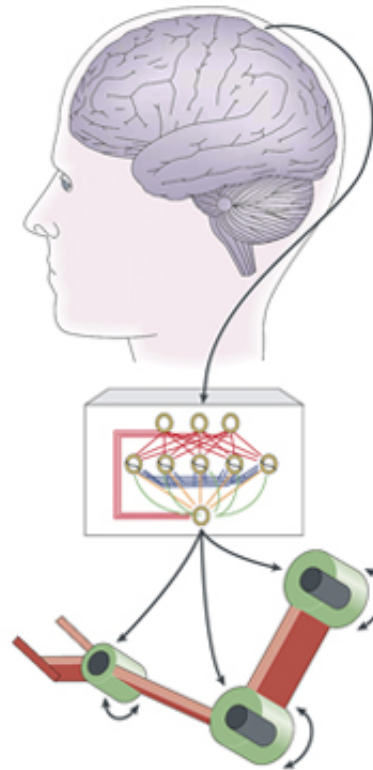
Computational neuroscience

Lots of interest from multiple fields!



Brain machine interfaces


Lots of interest from multiple fields!



Schematic from <http://www.brainfacts.org/-/media/Brainfacts/Article-Multimedia/About-Neuroscience/Technologies/Illustration-Brain-and-Arm.ashx>
Research such as L.R. Hochberg et al., Nature 2012

Computational neuroscience

Lots of interest, including from industry!



Computational neuroscience

- But what does it mean to understand how the brain works?

(discussion)

Computational neuroscience

Levels of investigation

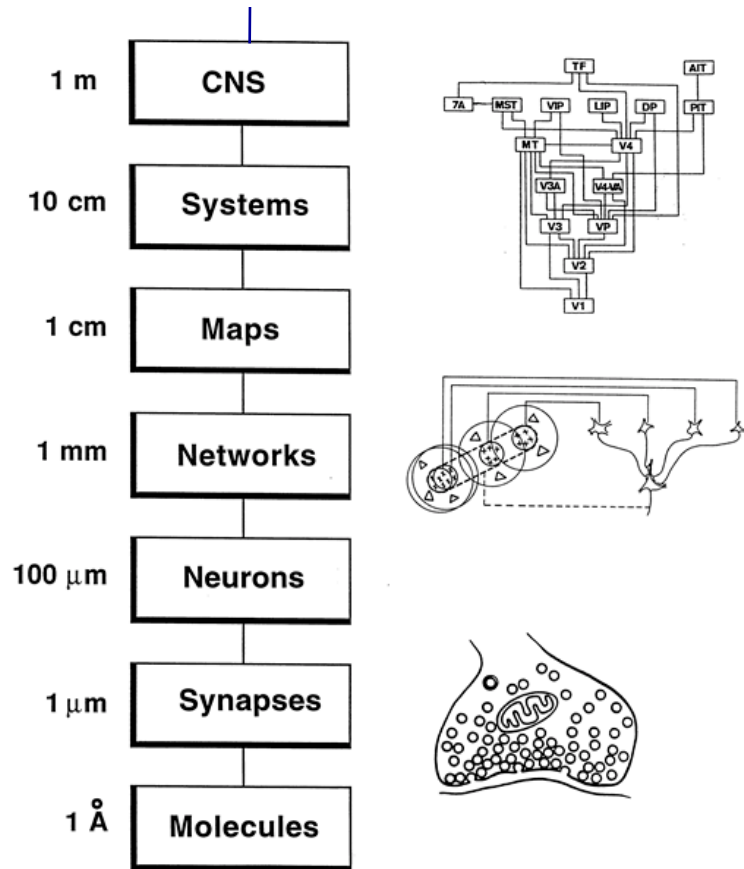


Diagram: Terrence Sejnowski

Computational neuroscience

Levels of investigation

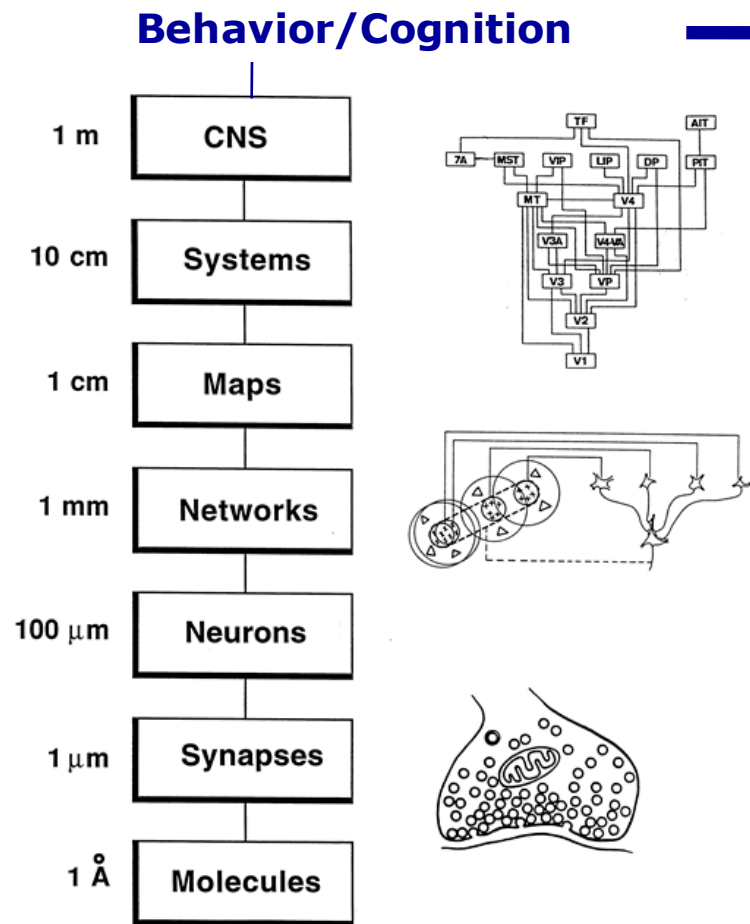


Diagram: Terrence Sejnowski

ion

Computational Neuroscience

- Finding low dimensional descriptions of high dimensional biological data
- Proposing models to explain data, making predictions, generalization
- Often close interplay between computational frameworks and experiments

Computational Neuroscience

- Finding low dimensional descriptions of high dimensional biological data
- Proposing models to explain data, making predictions, generalization
- Often close interplay between computational frameworks and experiments
- Why is this an exciting field today?

Computational Neuroscience



- Finding low dimensional descriptions of high dimensional biological data
- Proposing models to explain data, making predictions, generalization
- Often close interplay between computational frameworks and experiments
- Today: more data, stronger computers, more quantitative and advances in machine learning, advances in experimental approaches

Computational neuroscience

Types of quantitative tools

- Machine learning;
- Statistics / probability
- Information theory;
- Optimization;
- Control theory;
- Signal Processing;
- Dynamical systems;
- Statistical physics;
- Biophysics



Neurons in the brain



14

Santiago Ramón y Cajal (drawn 1899)

Neurons in the brain




15

How many neurons in the human brain?

Neurons in the brain

Around 10^{11} neurons in the human brain

Brain versus CPU



Brain versus CPU

whole brain (2 kg):

10^{11} neurons

10^{14} connections

(1000 connections per neuron)

Adapted from Gatsby Computational Neuroscience course

Brain versus CPU

whole brain (2 kg):

10^{11} neurons

10^{14} connections

(1000 connections per neuron)

whole CPU:


10^9 transistors

2×10^9 connections


(few connections per transistor)

Adapted from Gatsby Computational Neuroscience course

Man versus Machine




Neural Currency



- Spike (action potential): approximately 100 mV rise in voltage, lasting for approximately 1 msec

Neural processing



- Example: Visual neurons spike in response to features or properties of images

Computational neuroscience

David Marr, 1982

Levels of modeling

- **Computational:** What is the goal of the computation?
- **Algorithm:** What is the strategy or algorithm to achieve the computation?
- **Implementation:** How is this implemented in the brain? (mechanism in neurons; networks of neurons)

Computational neuroscience

David Marr, 1982 Levels of modeling

Computational theory	Representation and algorithm	Hardware implementation
What is the goal of the computation, why is it appropriate, and what is the logic of the strategy by which it can be carried out?	How can this computational theory be implemented? In particular, what is the representation for the input and output, and what is the algorithm for the transformation?	How can the representation and algorithm be realized physically?

Figure 1–4. The three levels at which any machine carrying out an information-processing task must be understood.

Computational neuroscience

David Marr, 1982

Levels of modeling, example:

- Computational: What is the goal of the computation?
example: maximize reward



Computational neuroscience

David Marr, 1982

Levels of modeling

- Computational: What is the goal of the computation?
example: maximize reward
- Algorithm: What is the strategy or algorithm to achieve the computation?
reinforcement learning algorithms to minimize prediction error of reward

Computational neuroscience


David Marr, 1982

Levels of modeling

- Computational: What is the goal of the computation?
example: maximize reward
- Algorithm: What is the strategy or algorithm to achieve the computation?
reinforcement learning algorithms to minimize prediction error of reward
- Implementation: How is this implemented in the brain?
dopaminergic neurons?

Computational neuroscience


Another way to parse model types



What? How? Why?

Computational neuroscience


Answer three kinds of questions about the brain



- Descriptive (**what**)
- Mechanistic (**how**)
- Interpretive/normative (**why**)

Computational neuroscience


Answer three kinds of questions about the brain



- Descriptive (**what**): What is the transform between input and output?

Computational neuroscience


Answer three kinds of questions about the brain



- Descriptive (what): What is the transform between input and output?
- Mechanistic (**how**): How does the system transform the input into the output?

Computational neuroscience

Answer three kinds of questions about the brain



- Descriptive (what): What is the transform between input and output?
- Mechanistic (how): How does the system transform the input into the output?
- Interpretive/normative (**why**): Why does the system transform the input into the output?

Computational neuroscience

Answer three kinds of questions about the brain

- Descriptive (**what**)
examples: addition or division between neural units, receptive field models
- Mechanistic (**how**)
examples: ionic channels, synaptic depression, network mechanisms
- Interpretive/normative (**why**)
examples: efficient coding, optimal estimation or decision, wiring length, metabolic cost

Computational neuroscience

Answers three kinds of questions about the brain

- descriptive (*what*)
 - mechanistic (*how*)
- } Interpretive/normative models (*why*)

Computational neuroscience

Answer three kinds of questions about the brain

So what level do we want to study?

And How can we find out how the brain works?

(discussion)

Computational neuroscience

<http://ccneuro.org/ccn-blog/>

How can we find out how the brain works? (2017)

Yoshua Bengio, U Montreal

Department of Computer Science and Operations Research

“ If there is a compact description of the computational principles which explain how the brain manages to provide us with our intelligence, this is something I would consider the core explanation for how the brain works – a little bit like the laws of physics for our physical world. Note that this is very different from the structured observation of our world in all its encyclopedic detail, which provides a useful map of our world, but not a principled explanation. Just replace ‘world’ by ‘brain’. My thesis is that those principles would also allow us to build intelligent machines and that at the heart of our intelligence is our ability to learn and make sense of the world, by observing it and interacting with it. That is why I believe in the importance of a continuous discussion between the brain researchers and AI researchers, especially those in machine learning – particularly deep learning and neural networks. This is likely to benefit AI research as well, as it has in the past.”

Computational neuroscience

<http://ccneuro.org/ccn-blog/>

How can we find out how the brain works? (2017)

Tom Griffiths, UC Berkeley

Department of Psychology and Cognitive Science

“ As a cognitive scientist I normally think about this question from the perspective of Marr’s levels of analysis. Understanding how the brain works is a question at what Marr called the “implementation” level, but I think a lot of insight can be gained by asking why the brain does what it does — a question at Marr’s “computational” level. Between those levels of analysis is the “algorithmic” level, which looks at the particular cognitive processes that are involved in solving a problem. Over the last few years a lot of progress has been made at both the implementation level and the computational level, but I think the algorithmic level gets neglected when we think about the brain. Understanding the algorithms that brains execute — and how brains learn to execute those algorithms — is going to be a critical part of finding out how the brain works.”

Computational neuroscience

<http://ccneuro.org/ccn-blog/>

How can we find out how the brain works? (2017)

Alona Fyshe, University of Victoria, BC, Canada

Department of Computer Science

“ We will need to continue to study the brain at multiple scales, both at the neuronal level, at the macro level (via brain imaging), and at the behavioral level. And we need to continue to bring these worlds together. We also need to start pushing brain imaging experiments in humans out into the real world. We can learn something about human language understanding by watching people read single words or single sentences, but we will miss out on the higher level comprehension areas that are required for larger scale understanding and reasoning. Similarly, viewing pictures or watching videos tells us something about vision, but interacting with objects in the real world will likely tell us more. There is tremendous value in tightly controlled experimental paradigms, but we also need some people doing the hard work that gets at the more holistic aspects of brain information processing.”

Computational neuroscience

<http://ccneuro.org/ccn-blog/>

How can we find out how the brain works? (2017)

Michael Shadlen, Columbia University

Department of Neuroscience

“ By testing neurobiological hypotheses that address the “how” question at a variety of levels. To me, “how the brain works” is a biological problem because I am less interested in mimicking the brain with a machine than I am in assessing what goes wrong when the brain doesn’t work, and how we might remedy the fault. To this end, functional equivalence (like airplanes to birds), which might interest the engineer, is not enough and possibly detrimental—a misguided diversion. To make progress on the “how” of cognitive function, my approach is to focus less on the representation of information and more on what the organism does with the information.”

Computational neuroscience

<http://ccneuro.org/ccn-blog/>

How can we find out how the brain works? (2017)

Birte Forstmann, University of Amsterdam

Integrative model-based cognitive neuroscience research unit

- “ I believe that formal models that make simultaneous predictions about different modalities such as behavior and the brain are powerful tools. Such tools could help to gain a better mechanistic understanding of brain function.”

Computational neuroscience

<http://ccneuro.org/ccn-blog/>

How can we find out how the brain works? (2017)

Nicole Rust, University of Pennsylvania

Department of Psychology

“ It all begins with thoughtful descriptions of the computations that the brain solves, which are often directly reflected in behavior. Ultimately, a description of “how” is formalized by a model that provides a non-trivial account of data. Crucially, while many of us have been taught that the ultimate test of understanding something is to build it, recent work in our field highlights that you can build something without deeply understanding how it works. Model interpretability is one of the biggest challenges that we currently face.”

Computational neuroscience

<http://ccneuro.org/ccn-blog/>

How can we find out how the brain works? (2017)

Odelia Schwartz, University of Miami

Department of Computer Science

“ There is a continued need for computational frameworks that interplay with experimental design and analysis at multiple levels (e.g., neurons, circuits, cognition). I have been intrigued by how neural systems represent and learn about stimuli in the natural environment, resulting in complex inferences and behavior. My main focus has been building computational neural models that push towards a more principled understanding for natural stimuli such as visual scenes. With advances in machine learning and in understanding the statistics of natural stimuli, I believe there is potential for progress in designing and interpreting experiments with naturalistic environments and tasks.”

Computational neuroscience

<http://ccneuro.org/ccn-blog/>

How can we find out how the brain works? (2017)

Wei Ji Ma, New York University

Center for Neural Science and Department of Psychology

“ ...At a more sociological level, I am old-fashioned and strongly believe in small, hypothesis-driven science. While some problems in neuroscience might be best addressed using big data, big simulations, or big collaborations, my sense is that those currently involve more hype than substance. Neuroscience and cognitive science have come far with a “letting a hundred flowers bloom” approach, and there is no evidence that this approach is bankrupt. More specifically in computational neuroscience, small science often amounts to a search for evolutionarily meaningful organizing principles, perhaps initially in a toy model – this is my favorite approach.”

This class

- Lectures
- Matlab computer tutorials and labs – most Thursdays in the Ungar computer teaching lab. Hand-in assignment mid semester.
- Several paper discussions that cover recent progress in the field, or classical papers.
- Final project (in small groups)
- No exams
- Participation

This class

Topics include...


- Neural coding
- Neural population Coding
- Brain Machine Interfaces
- Example neural system: The visual system
- Other example neural systems...
- Estimating descriptive (what) neural models from neural data
- Normative (why) models
- Relation to mechanistic (how) models
- Neural representation and processing of natural stimuli
- Unsupervised learning and neural representations
- Relation to recent advances in machine learning and deep learning
- Reinforcement learning

This class

<http://www.cs.miami.edu/home/odelia/teaching/compneuro2017/index.html>

This class

- Suggested textbook: Theoretical Neuroscience: Computational and Mathematical Modeling of Neural Systems, by Peter Dayan and L.F. Abbott.



This class

- Suggested textbook on visual processing and scene statistics:
Aapo Hyvarinen, Jarmo Hurri, Patrik O. Hoyer:
Natural Image Statistics: A probabilistic approach to early computational vision.

(particularly introduction chapters to the visual system and Interpretive (why) models)