

Computational Models of Neural Systems: 15-883

Fall 2013

Instructor:
David S. Touretzky
Computer Science Department

dst@cs.cmu.edu
Office: 9013 GHC
tel. 412-268-7561

Course Info

Time: Mon/Wed 4:30 to 5:50

Place: 4101 GHC (G. Hillman Center)

Credit: 12 units

Current syllabus: on the class web site

Textbook: none

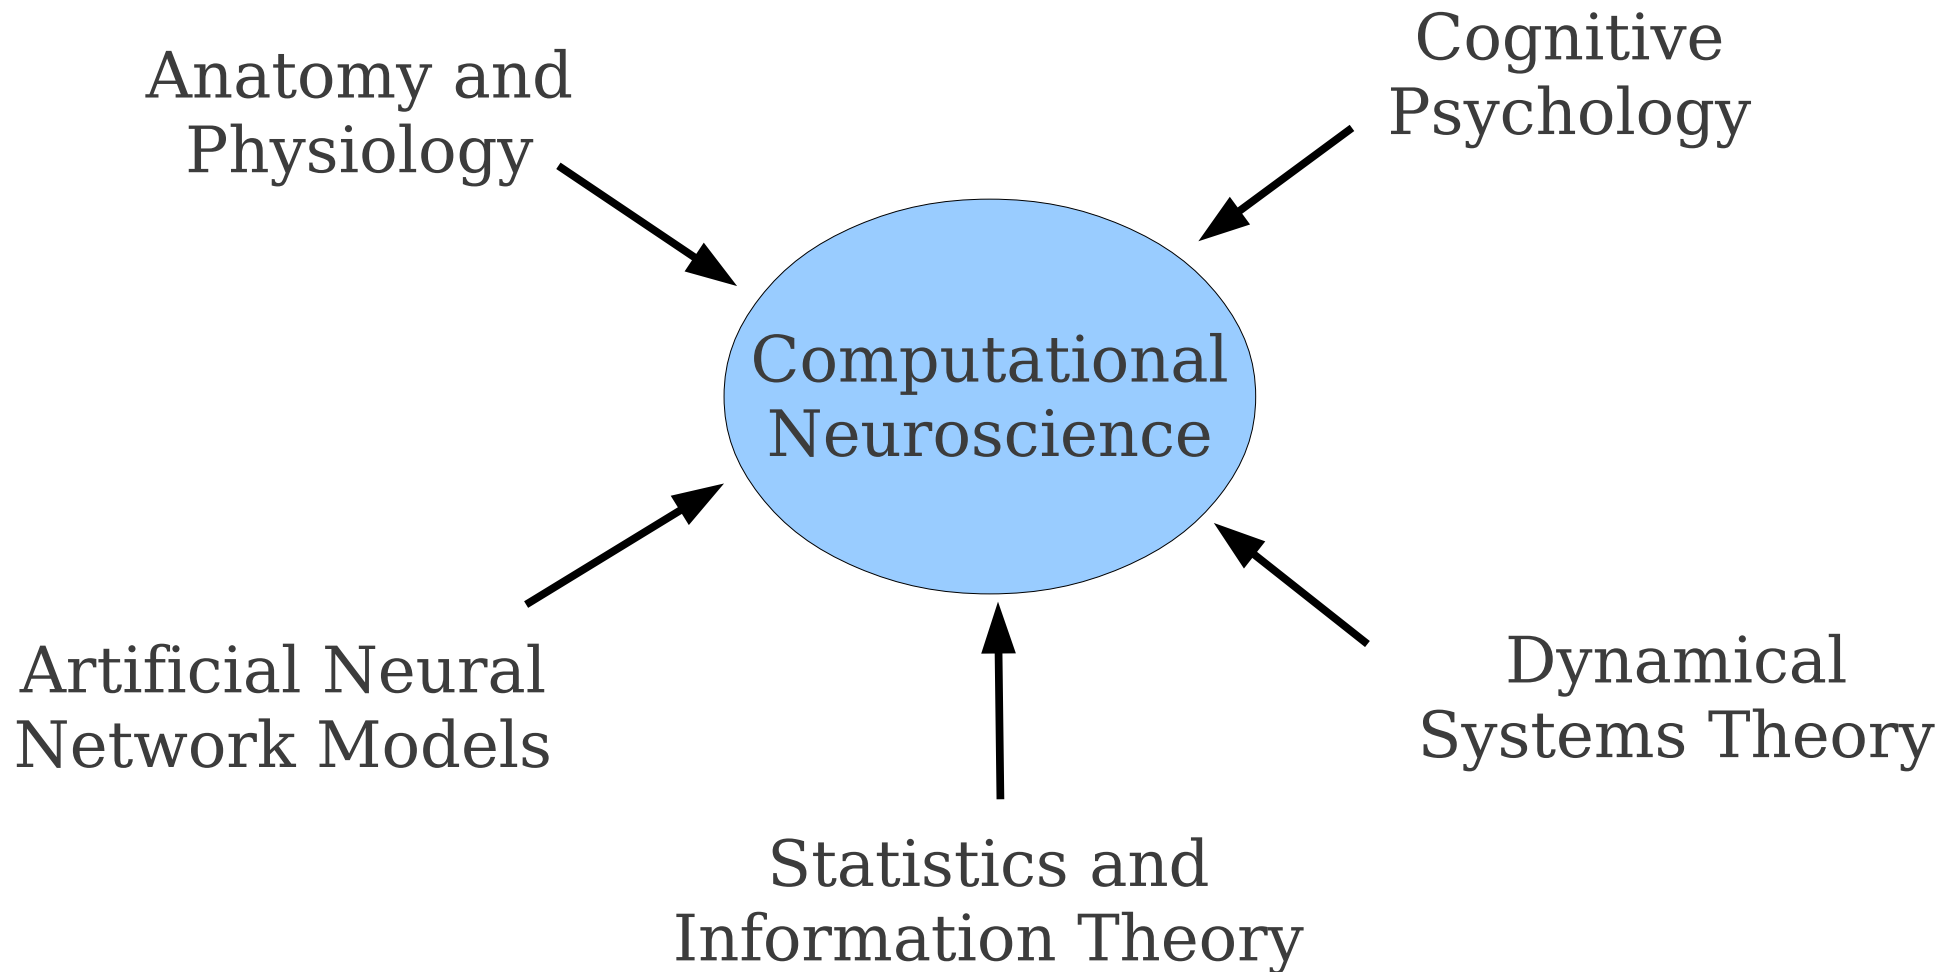
Readings:

Web repository (linked from syllabus)

Who Should Take This Course?

- Computer scientists who want to learn about the brain.
 - No prior neuroscience background required.
- Neuroscientists who want a computational perspective on brain function.
 - Focus is on representations and algorithms, rather than anatomy and physiology.
- Cognitive scientists who want to study brains as computing devices.
 - Taking the “brain as computer” metaphor seriously requires learning as much as possible about both.

Computational Neuroscience Intellectual Landscape



Varieties of “Neural Network” Research

- 1) Neuronal Modeling
- 2) Computational Neuroscience
- 3) Connectionist (PDP) Models
- 4) Artificial Neural Networks (ANNs)

Each area asks a different kind of question.

Some investigators work in more than one area.

Courses in all four areas are available at CMU or Pitt.




1: Neuronal Modeling

Understand the operation of single neurons or small neural circuits.

Detailed biophysical models of nerve cells, and collections of cells.

What makes a neuron spike?

Comp. neuro. course at Pitt
(Bard Ermentrout, Jon Robin,
Math Dept.)




Churchland & Sejnowski 1988

2: Computational Neuroscience

Model information processing in actual brain systems.

The models refer to specific anatomical structures, but their operation may be abstract.

How does the hippocampus retrieve memories?



Churchland & Sejnowski 1988

15-883 Computational Models of Neural Systems course (Touretzky)


3: Connectionist (PDP) Modeling

Modeling human cognition in a brain-like way: parallel constraint satisfaction; distributed activity patterns instead of symbols.



Models are fairly abstract.

How do priming effects act to influence reading?

85-719 PDP models course (Dave Plaut)



McClelland & Rumelhart 1981




4: Artificial Neural Nets

Pattern recognition, adaptive control, time series prediction. (This is where the money gets made.)

Simple, “neuron-like” computing elements; local computation.


How can a machine learn to efficiently recognize patterns?

15-782 (Artificial Neural Nets)



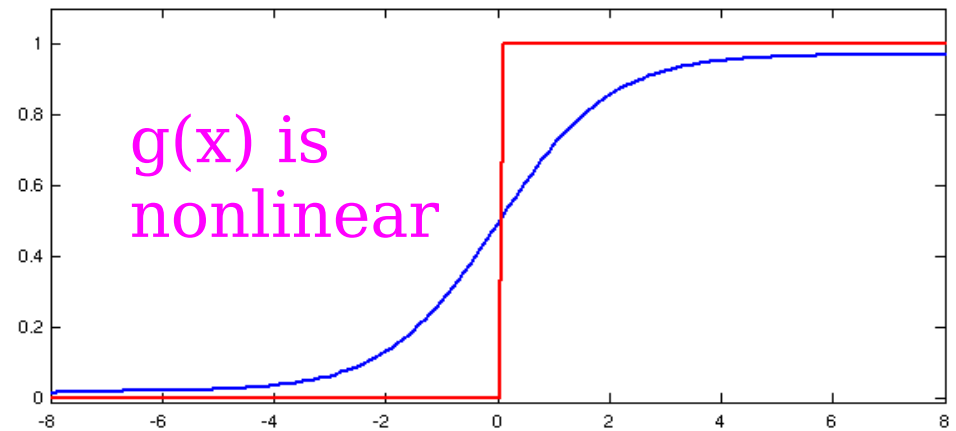
Pomerleau 1993: ALVINN

What's a "Neuron"?



$$netact = \sum_i w_i x_i = \vec{w} \cdot \vec{x}$$

$$y = g(netact)$$



Organization of this Course

- Specific domain (e.g., the hippocampus)
 - Background lecture: anatomy and physiology
 - Family of models (e.g., associative memory models)
 - One or more papers in each family
 - Class discussion
 - Occasionally: experimentation in MATLAB
- Occasional problem sets
- Modeling project (or term paper)
- Mid-term exam
- Final exam


Grading (Approximate Weightings)

Problem sets	20%	
Modeling project	20%	(or term paper)
Midterm exam	30%	
Final exam	30%	

MATLAB

You need to learn MATLAB. It's fun!


Type “matlab” on Andrew to run it.
“peaks” will display this graph;
“doc peaks” will tell you about it



Student Version of MATLAB: available for
Windows/Linux/Mac for \$99. Purchase from
mathworks.com or CMU bookstore.

Pitt students can purchase a license for \$10.

Tutorials are available online:
see the class homepage.



What You Should Do Today

- Hand in your student survey questionnaire.
- Read Churchland intro.
- Start learning MATLAB.
 - Type “demo” for a list of demos, and scroll down to the “Graphics” section. Play around a bit.
 - We'll have more formal MATLAB instruction later.
- Get started on Wednesday's reading.