42-590 / 18-699 Neural Signal Processing

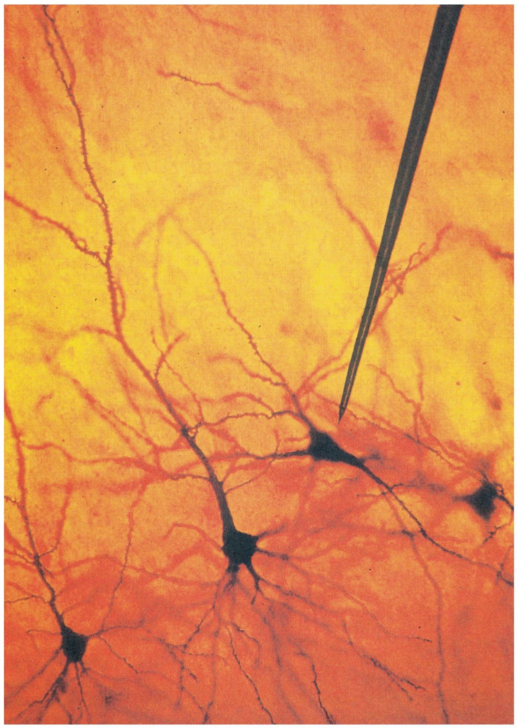

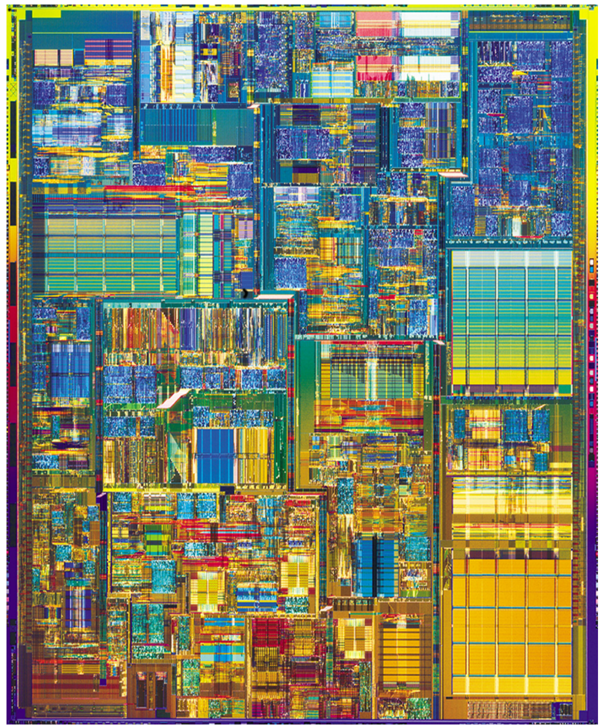

Prof. Byron Yu
Electrical & Computer Engineering / Biomedical Engineering
Carnegie Mellon University
Spring 2010

What is Neural Signal Processing?

For centuries, people have sought to understand what gives rise to our ability to perceive, to reason, and to act.

4th century BC: Aristotle identified the heart as the seat of intelligence and thought.

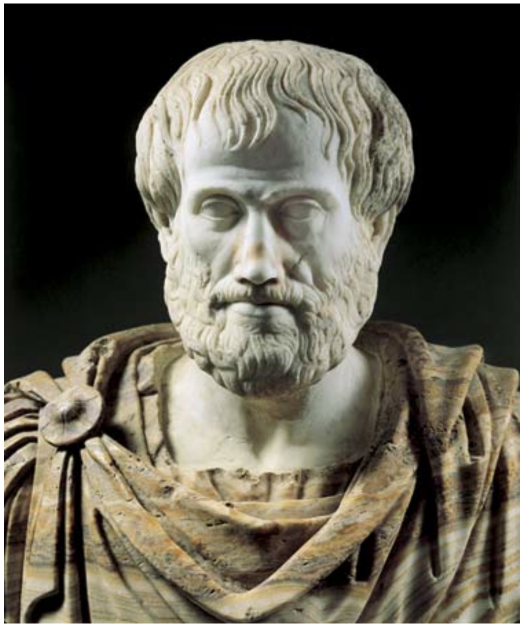

What is Neural Signal Processing?

We've learned a lot since then:

- It's the brain, not the heart.

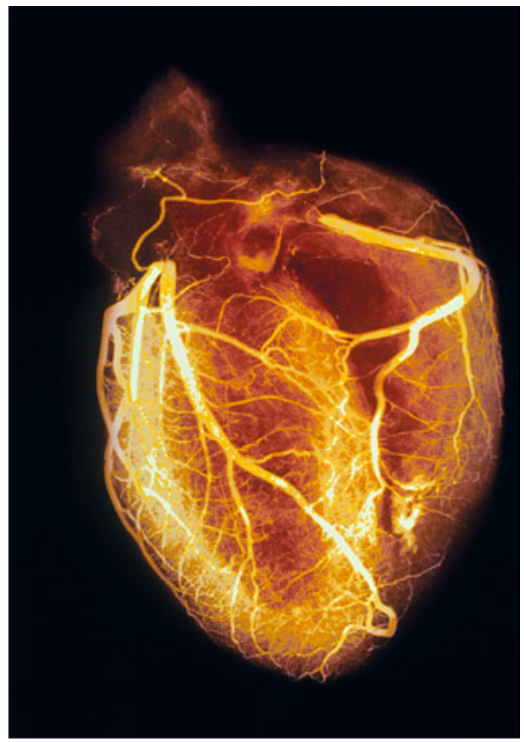

http://science.nationalgeographic.com

What is Neural Signal Processing?

- Nerves conduct electrical signals, rather than conveying fluids secreted by the brain.

http://science.nationalgeographic.com

What is Neural Signal Processing?

- Nervous system is a network of discrete cells called *neurons*, rather than a continuous web.

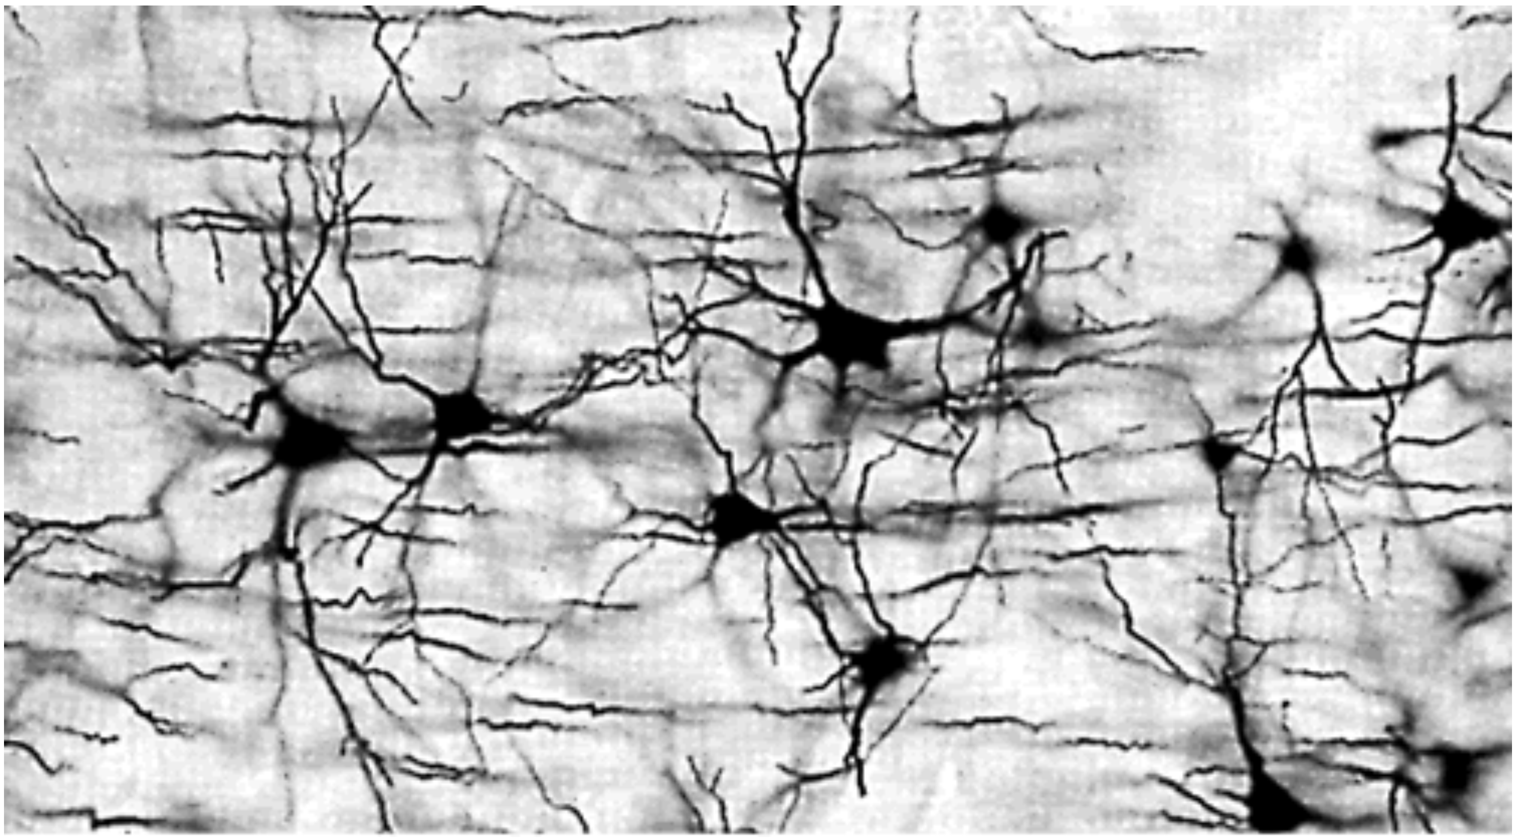

http://www.alanturing.net

What is Neural Signal Processing?

There are many important neuroscience discoveries that I haven't listed here.

However, we're still only scratching the surface today. There's *way* more that we don't know about the brain than we do know.

How are we to further our understanding of the brain?

$\Rightarrow$ We must monitor the activity of its constituent elements: neurons.

Holy grail: Monitor the activity of every neuron in the brain.

What is the best that we can do today?

One end of the spectrum…

Single-electrode recordings

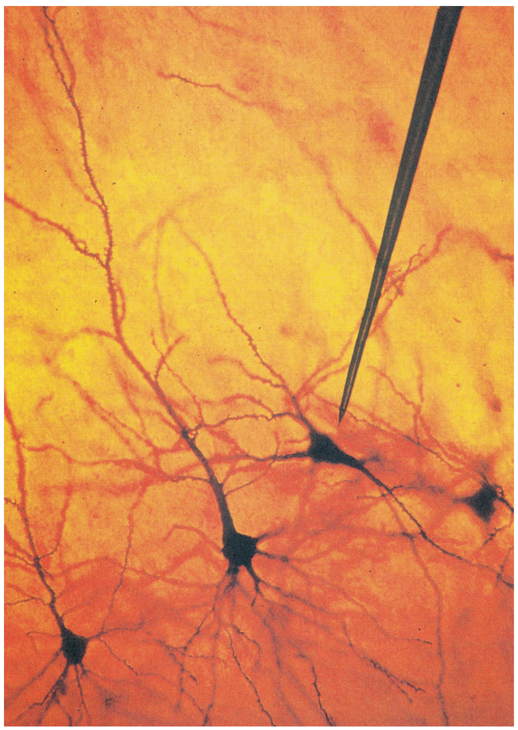

Pro: single-neuron resolution

Con: can only monitor one (or a small number of) neurons at a time

…the other end of the spectrum

Electroencephalography (EEG)

http://people.brandeis.edu/~sekuler

Functional magnetic resonance imaging (fMRI)

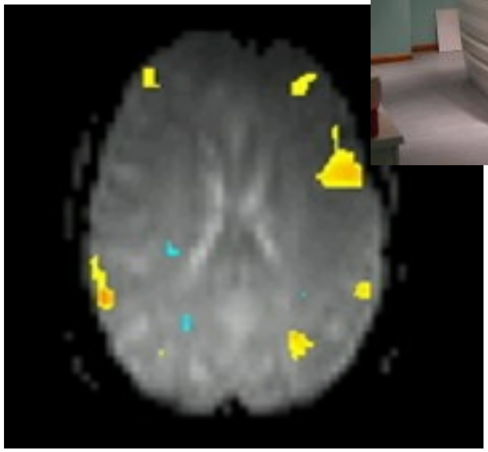

http://neurophilosophy.files.wordpress.com/

Pro: can monitor entire brain

Con: no single-neuron resolution

Different neural recording technologies: it's all about tradeoffs

Stadium analogy

Stadium is the brain.
Each person is a neuron.

Stadium analogy

Single electrode recording is like listening in to what one person is saying.

Stadium analogy

EEG and fMRI are like listening to the collective roar of the crowd.

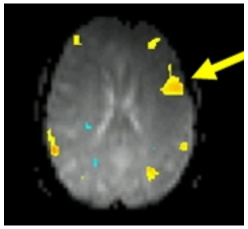

Stadium analogy

Stadium analogy

Ideally, we'd like to monitor what each individual person is saying.

More recent neural recording technologies

Only in last 15 years or so, we've been able to monitor many neurons at single-neuron resolution.

Multi-electrode arrays

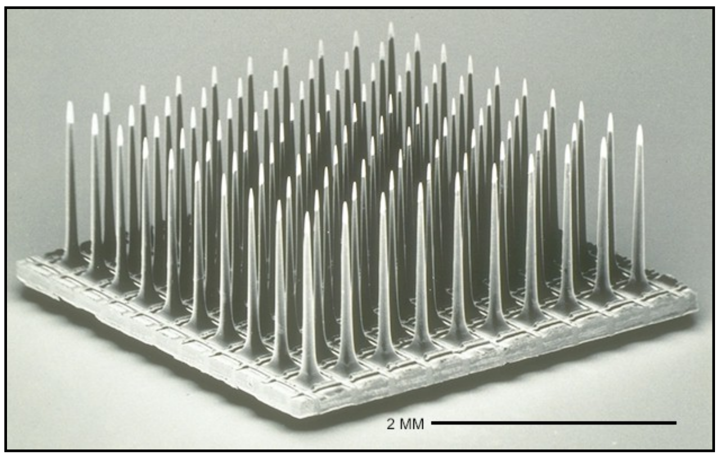

Cyberkinetics, Inc

Optical imaging

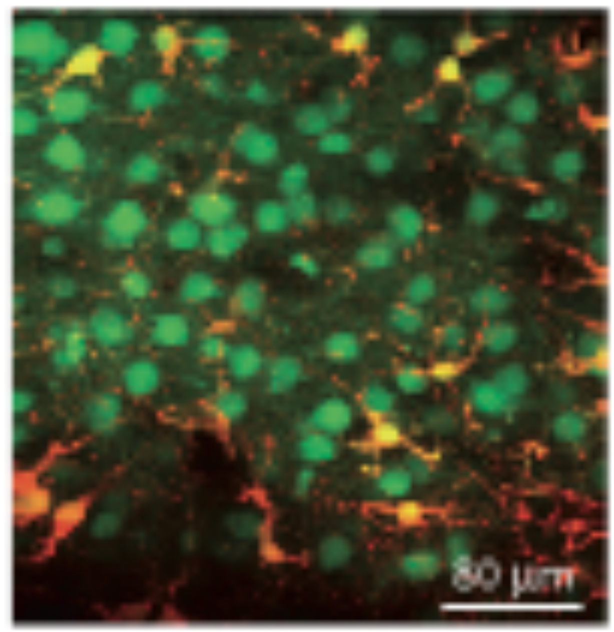

Kerr and Denk, 2008.

Stadium analogy

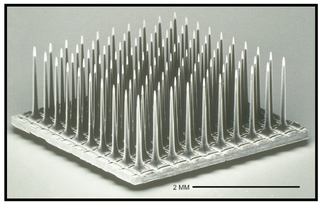

Multi-electrode array recording is like listening in on <u>multiple</u> <u>individual</u> conversations.

Stadium analogy

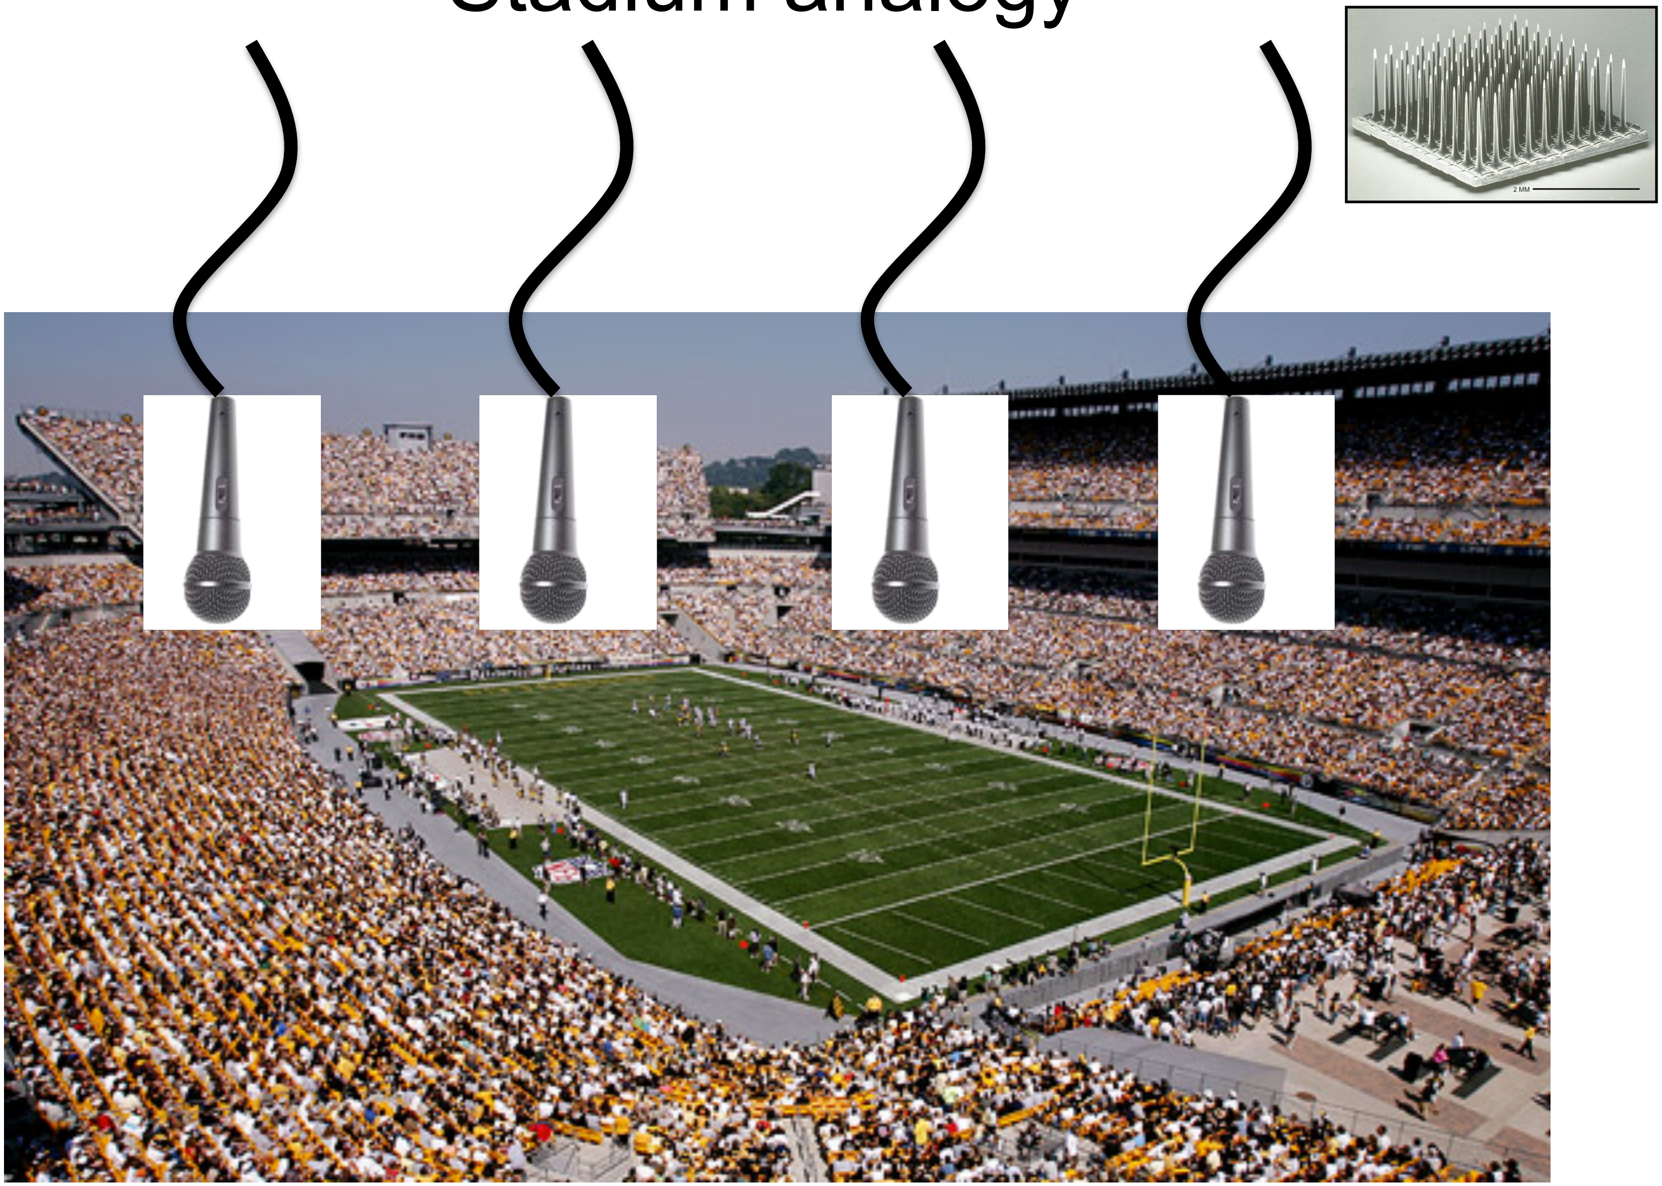

Stadium analogy

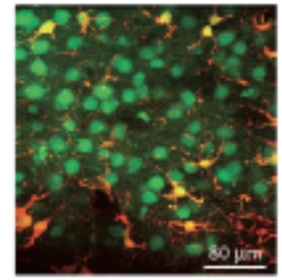

Optical imaging is like watching multiple individual mouths moving, from which we can deduce what each person is saying.

Stadium analogy

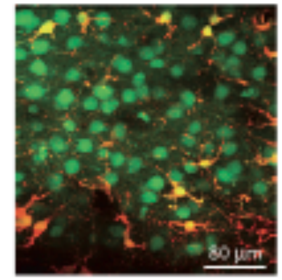

What is Neural Signal Processing?

Technologies like multi-electrode arrays and optical imaging are providing unprecedented views of the brain's activity.

Even though we're far from monitoring every neuron in the brain, we now have more data than we know what to do with it.

We need powerful statistical methods to make inferences about what the brain is doing from sparse sampling.

What is Neural Signal Processing?

<u>What we have</u>

Novel experimental paradigms

New neural recording technologies

Huge and rich data sets

<u>Goals</u>

Further our basic understanding of brain function

Develop biomedical devices that interface with the brain

<u>Missing link</u>

Signal processing methods appropriate for neural data

Creative ways of extracting meaningful features from huge data sets

Topic of this course!

Further our basic understanding of brain function

<u>Examples</u>:

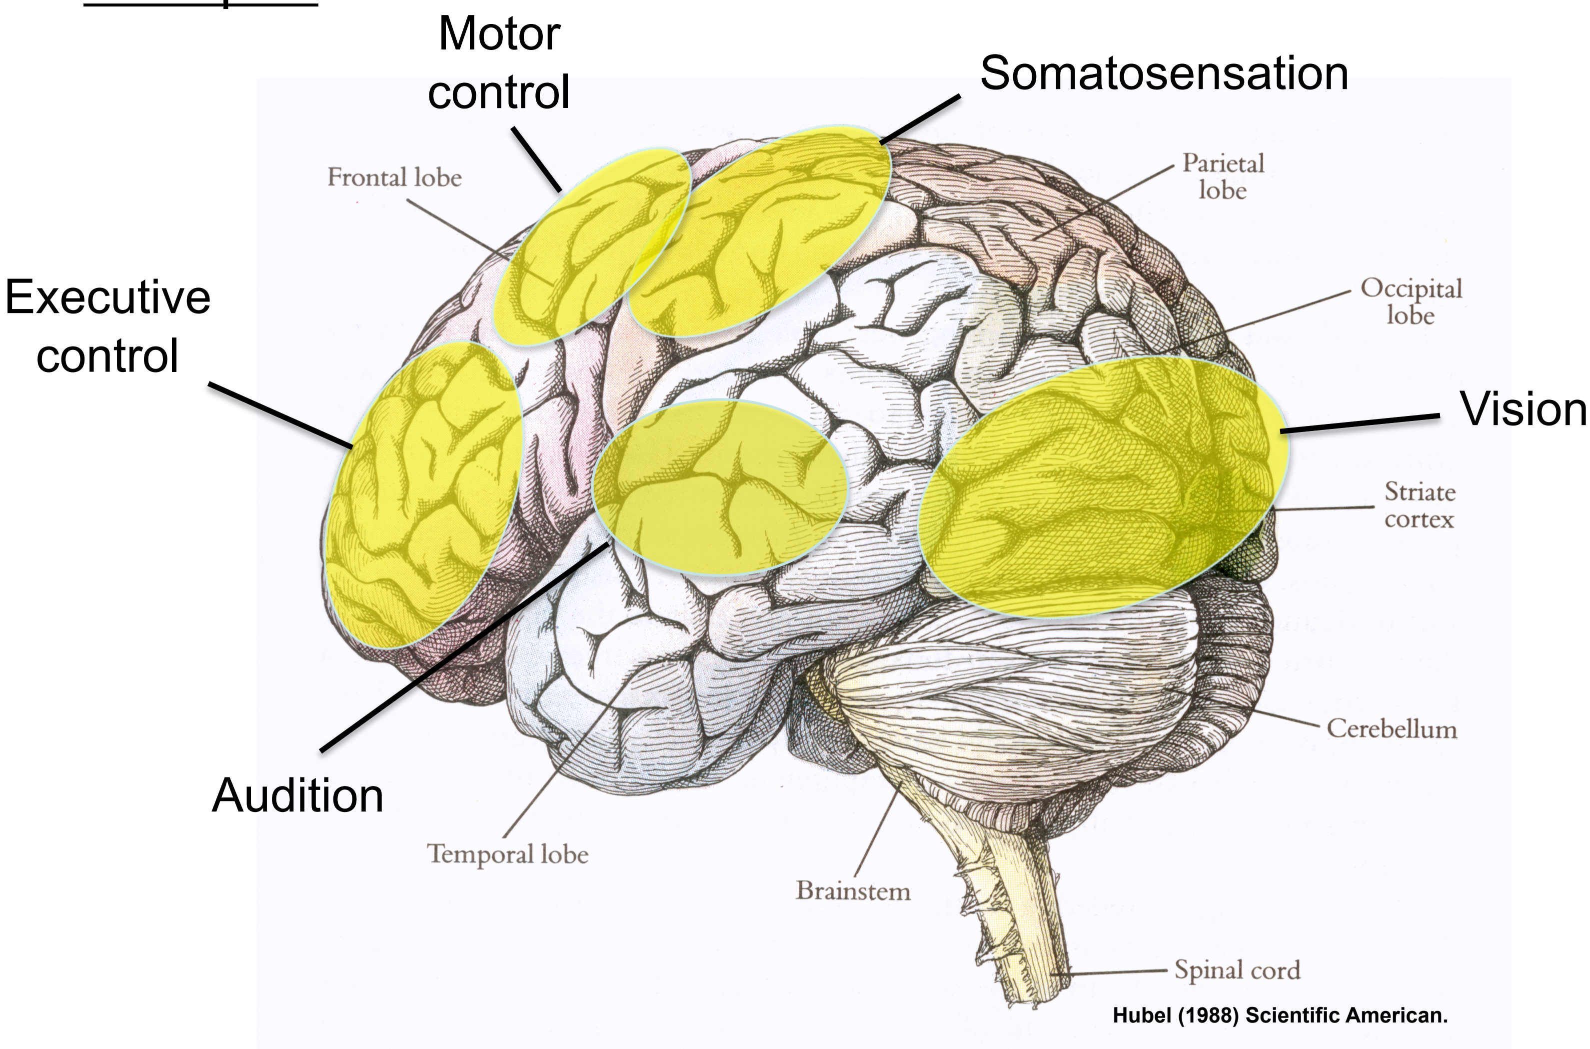

Develop biomedical devices that interface with the brain

Examples:

Cochlear implants

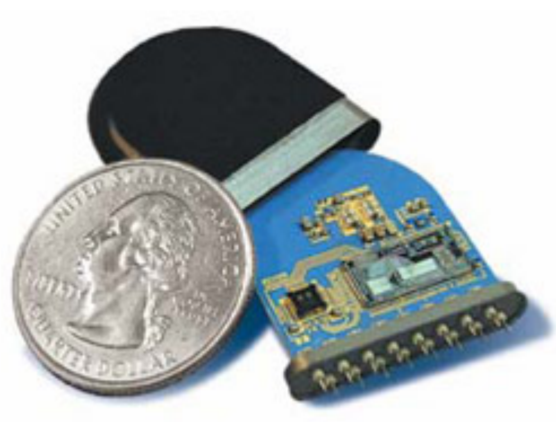

Advanced Bionics Corp.

Deep brain stimulation

Medtronic, Inc.

Motor prostheses

Cyberkinetics

Neural prosthetic systems

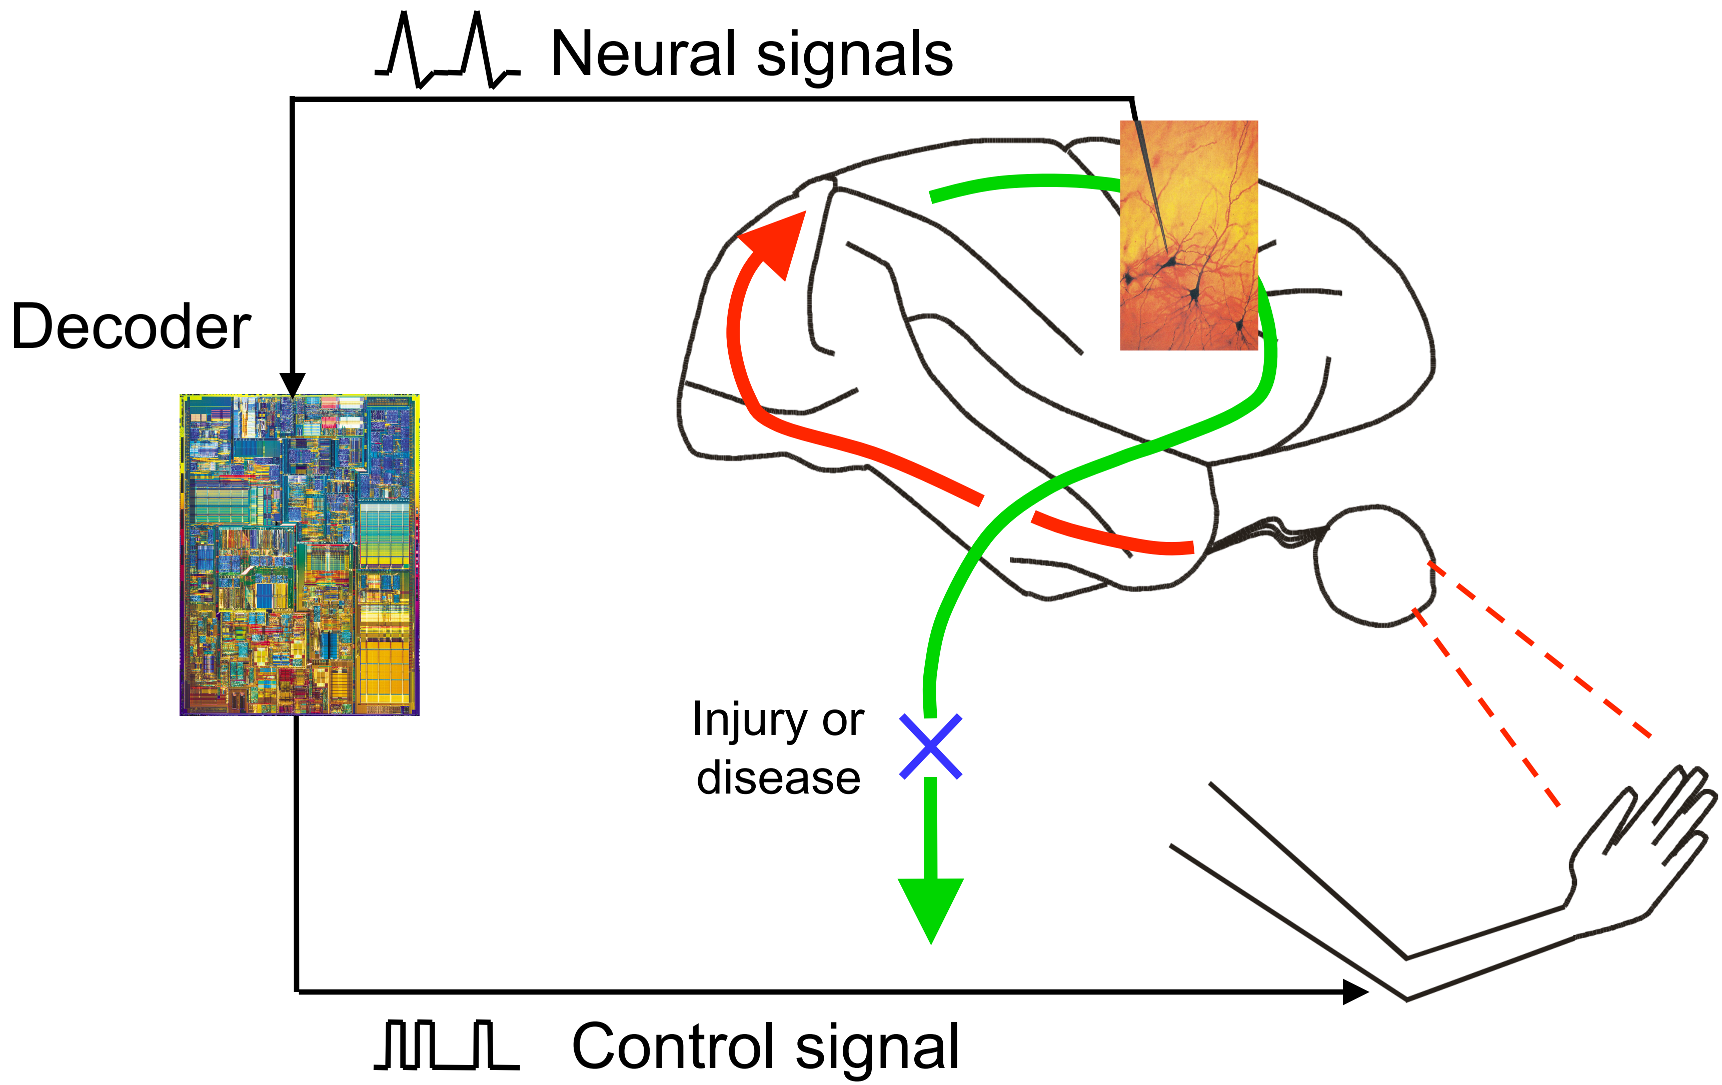

Motor prostheses

Ultimately, we want to help human patients.

Cyberkinetics, Inc.

Motor prostheses

Basic research on prosthetic systems is done with monkeys.

Schwartz Lab, U Pitt

Motor prostheses

There is still much work to be done to get decoded movements to rival natural movements.

Monkey **hand-controlling** a virtual cursor

Credit: Churchland, Kaufman, Shenoy

Motor prostheses

There is still much work to be done to get decoded movements to rival natural movements.

Monkey **hand-controlling** a virtual cursor

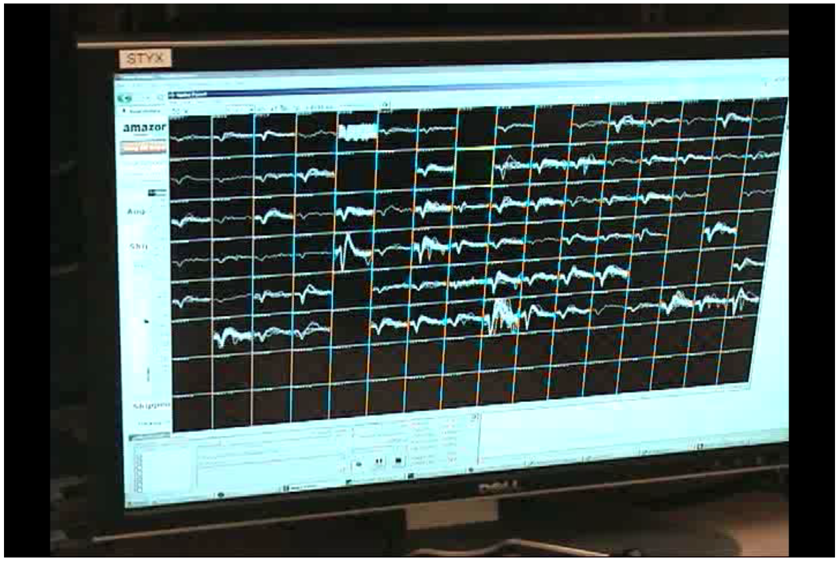

Monkey **brain-controlling** a virtual cursor

Credit: Churchland, Kaufman, Shenoy

Credit: Chestek, Gilja, Nuyujukian, Cunningham, Shenoy

Why this course is timely

- Neuroscience, as a field, used to be data-limited. Although more data is always nice, we're rapidly becoming limited by the available signal processing methods.

- CMU and Pitt together form a neuroscience hotbed. There are over 90 faculty members across the two universities that work on neuroscience-related topics. You are at the center of the action!

Syllabus

Course staff

Instructor: Byron Yu

B204 Hamerschlag Hall
byronyu@andrew.cmu.edu
(412) 268-9658
Office hours: Tues / Thurs noon--1pm

Syllabus

Course staff

Teaching assistant: Jinyin Zhang

jinyinz@andrew.cmu.edu
(412) 268-5251
Office hours: Wed / Fri noon-1pm
Location: TBD

Syllabus

Course staff

Course management assistant: Bara Ammoura

D200 Hamerschlag Hall
bammoura@ece.cmu.edu
(412) 268-6595
Office hours: Mon-Fri 8:30am-5pm

Waiting list

Syllabus

Course webpage

http://www.cmu.edu/blackboard/

Search for the course page for 18-699 (not 42-590).

Pitt students: Please let me know if you don't have access to Blackboard.

Syllabus

Course goals:

(1) to introduce the statistical tools used to study large-scale neural activity

(2) to bring out the real-world challenges of working with experimental data

By the end of the course, students should be able to ask research-level questions in neural signal processing, as well as develop new statistical tools for problems in their own research. In short, this course serves as a stepping stone to research in neural signal processing.

Syllabus

Course outline:

Motivation: doing neural signal processing requires some intuition for how neural signals are generated and how neurons communicate with each other.

First 3 weeks -- basic neuroscience from ground zero.

Rest of course – signal processing and machine learning methods as applied to neural signals.

Syllabus

Course outline:

1. What is neural signal processing?

2. Neuroscience basics. Membrane potential. Action potential. Synaptic transmission.

3. Spike train analysis. Spike histogram. Tuning curve. Poisson process.

4. Fundamentals of probabilistic machine learning. Maximum likelihood parameter estimation. Priors and likelihood functions.

5. Classification. Linear discriminant analysis. Naive Bayes.
Neuroscience application: discrete neural decoding

6. Graphical models.

Syllabus

Course outline (cont):

7. Mixture models. Expectation-maximization.
   Neuroscience application: spike sorting

8. Principal components analysis. Factor analysis.
   Neuroscience applications: dimensionality reduction, discrete neural decoding

9. Hidden Markov model. Kalman filter. Linear filter.
   Neuroscience application: continuous neural decoding

10. Cross-validation. Bayesian model selection.
    Neuroscience applications: dimensionality reduction, spike sorting

11. Point processes. Conditional intensity functions. Time-rescaling algorithm for point processes. Goodness-of-fit.
    Neuroscience application: spike train models

Syllabus

Prerequisites:

This course is ideally suited for students with a solid background in basic probability and linear algebra. Prior knowledge of neuroscience is welcome, but not required. Students with experience in neuroscience should be aware that the first 3 weeks will cover basic neuroscience.

Students should already be familiar with concepts such as:

Probability -- independence, conditional probability, Bayes rule, multivariate Gaussian distribution, Poisson distribution, Poisson process

Linear algebra -- basic matrix operations (sums and products), matrix inversion, eigenvectors and eigenvalues, singular value decomposition

For those unfamiliar with the concepts above, I would recommend Statistical Methods for Neuroscience and Psychology (36-746), and Probability Theory and Random Processes (36-217).

If you are unsure whether this class is for you, please talk with me.

Syllabus

<u>Relationship to other courses</u>:

Statistical Methods for Neuroscience and Psychology (36-746)

A natural precursor to Neural Signal Processing.

Machine Learning (10-701 and 15-781)

If you've taken Machine Learning already, Neural Signal Processing should be sufficiently different that it could be worth taking. Even though we will use the same textbook, many of the chapters are non-overlapping. Most of my examples will come from neuroscience.

Syllabus

Required textbook:

Pattern Recognition and Machine Learning
Christopher Bishop. Springer, 2006.

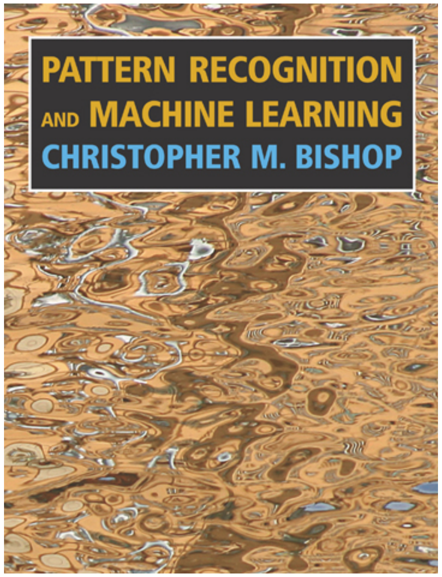

Syllabus

Optional textbooks:

Principles of Neural Science
Eric Kandel, James Schwartz, Thomas Jessell.
McGraw-Hill Medical, 2000.
Available at bookstore.

I will be photocopying sections of PNS needed for the course, but I highly recommend that you buy the book, especially if you plan to continue in neuroscience. This is the definitive neuroscience textbook and reference.

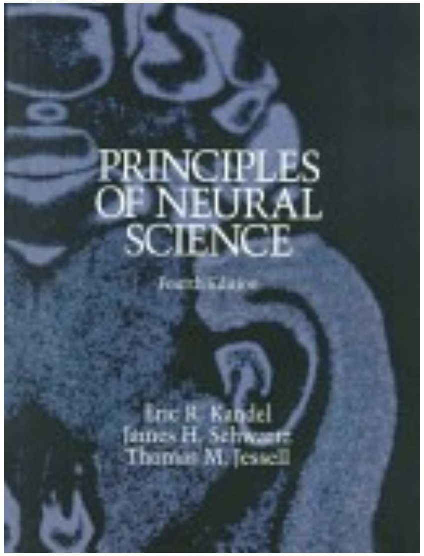

Syllabus

Optional textbooks:

Theoretical Neuroscience
Peter Dayan and L.F. Abbott. MIT Press, 2001.
Chapter 1 will be made available to you electronically.

Information Theory, Inference, and Learning Algorithms
David J.C. MacKay. Cambridge University Press, 2003.
Entire book is available online at:
http://www.inference.phy.cam.ac.uk/mackay/itprnn/ps/

Syllabus

<u>Assignments and exams:</u>

There will be approximately 8 problem sets during the semester and regular reading assignments.

<u>Midterm exam</u>: in class on **Thursday, March 4**.
<u>Final exam</u>: week of May 3, date TBD.

Most problem sets will have a Matlab component, in which students will implement various algorithms and apply them to neural data. Matlab is available on most campus machines and can also be downloaded (see course webpage).

Students may work on problem sets together, but each student must turn in his/her own solutions. *You may not simply copy another student's work*. All students are bound by the CMU Academic Integrity Code.

<u>Late policy for problem sets</u>: Each student is allowed two late problem sets during the semester (up to 24 hours after the deadline). Problem sets that are turned in outside of this grace period will receive zero credit.

Syllabus

Grading breakdown:

Problem sets	30%
Midterm exam	30%
Final exam	40%

Course Feedback

- This is the first time this course is offered.

- I welcome any feedback you might have about how this course is taught.

- Anonymous feedback can be posted on the ECE wiki: http://wiki.ece.cmu.edu/index.php

- Otherwise, feel free to send me an email.