Computational Vision

- **Instructor:**
  - Thomas Serre thomas_serre@brown.edu
- **Teaching Assistant:**
  - David Mely david_mely@brown.edu
- **Web:**
  - https://canvas.brown.edu/courses/851733
- **Meet:**
  - Metcalf #107 TR 9:00–10:20 am
- **My office hours:**
  - Metcalf #343, walk-ins after class, for appointments: serre-admin@brown.edu
  - TA will add his office hours on canvas

Many excellent computer vision courses at Brown…

- ENGN1610: Image Understanding
- ENGN2500: Medical Image Analysis
- ENGN2501: Digital Geometry Processing
- ENGN2520: Pattern Recognition and Computer Vision
- ENGN2560: Computer Vision
- ENGN2911X: Video Processing
- CSCI 143: Introduction to Computer Vision
- CSCI2950Q: Topics in Computer Vision
- CSCI129: Computational Photography

But this course is different!

- Focuses on computational models of biological vision
  - Reverse-engineering the primate visual system to build an artificial one
  - An attempt to see how recent work in computer science can help us shed light on brain sciences and vice versa

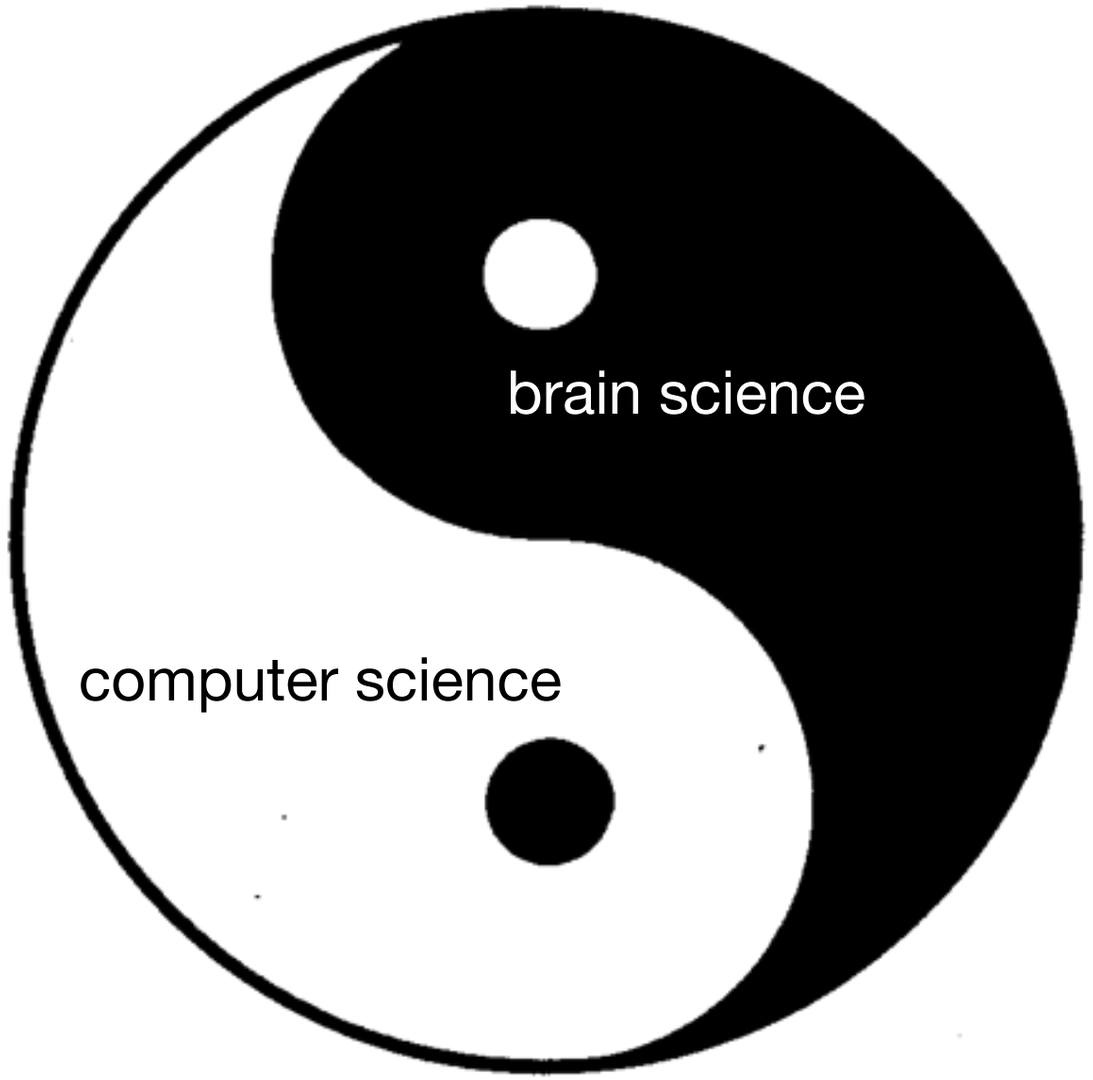

But this course is different!

- Some overlap with traditional computer vision courses…
  … But limited to very basic mathematical and computational tools (filtering, convolution, etc)

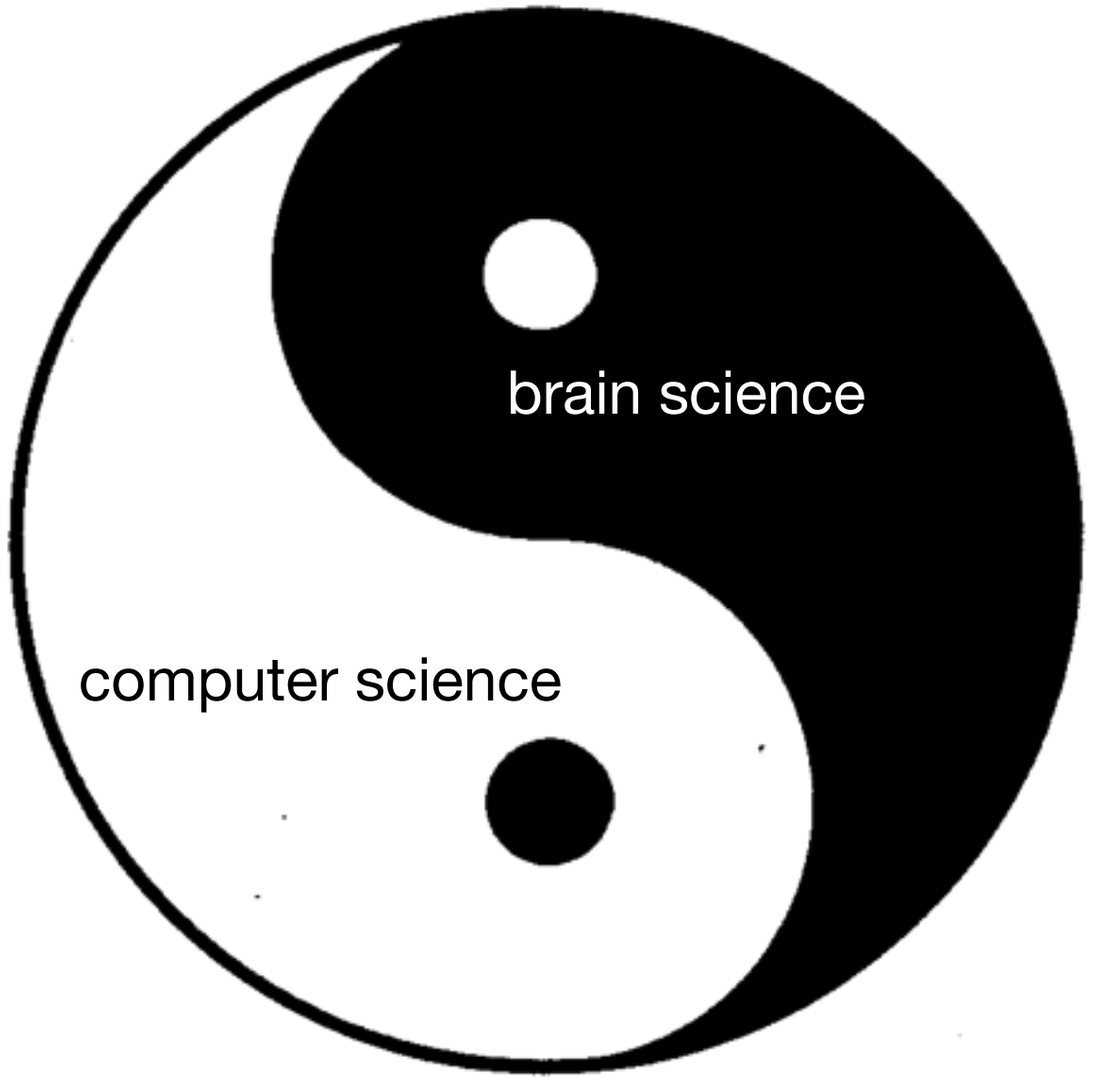

Image source: DiCarlo

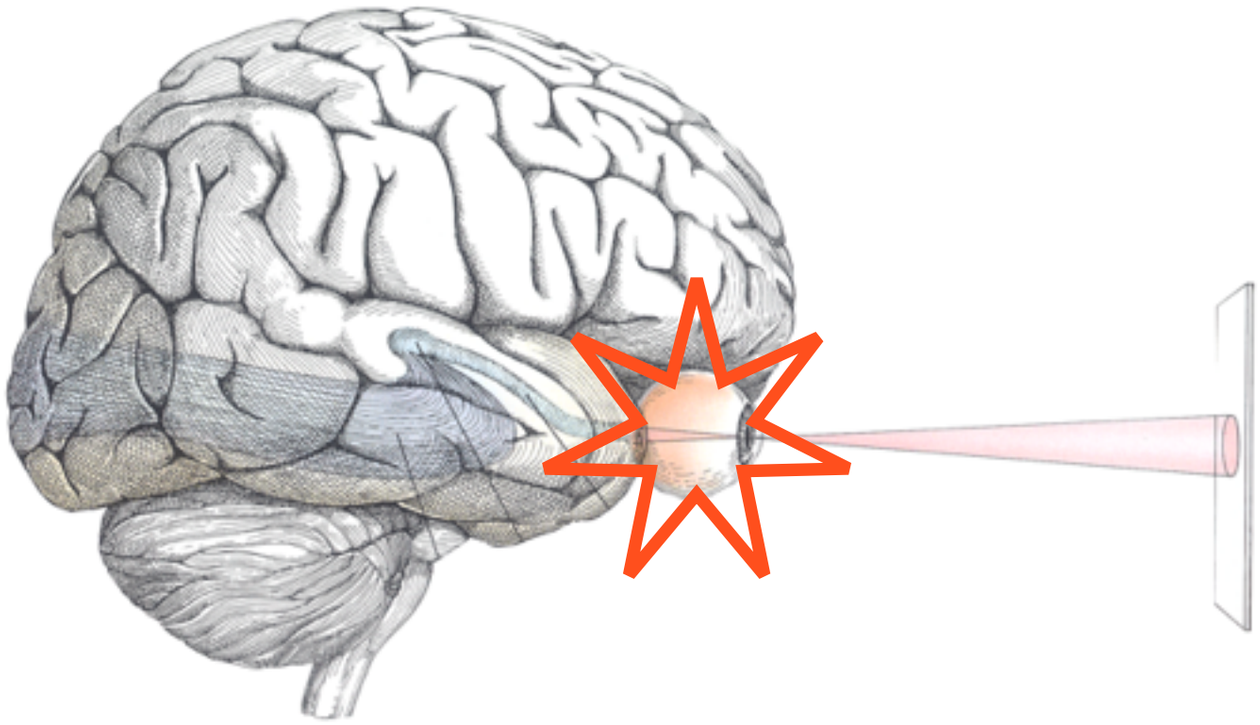

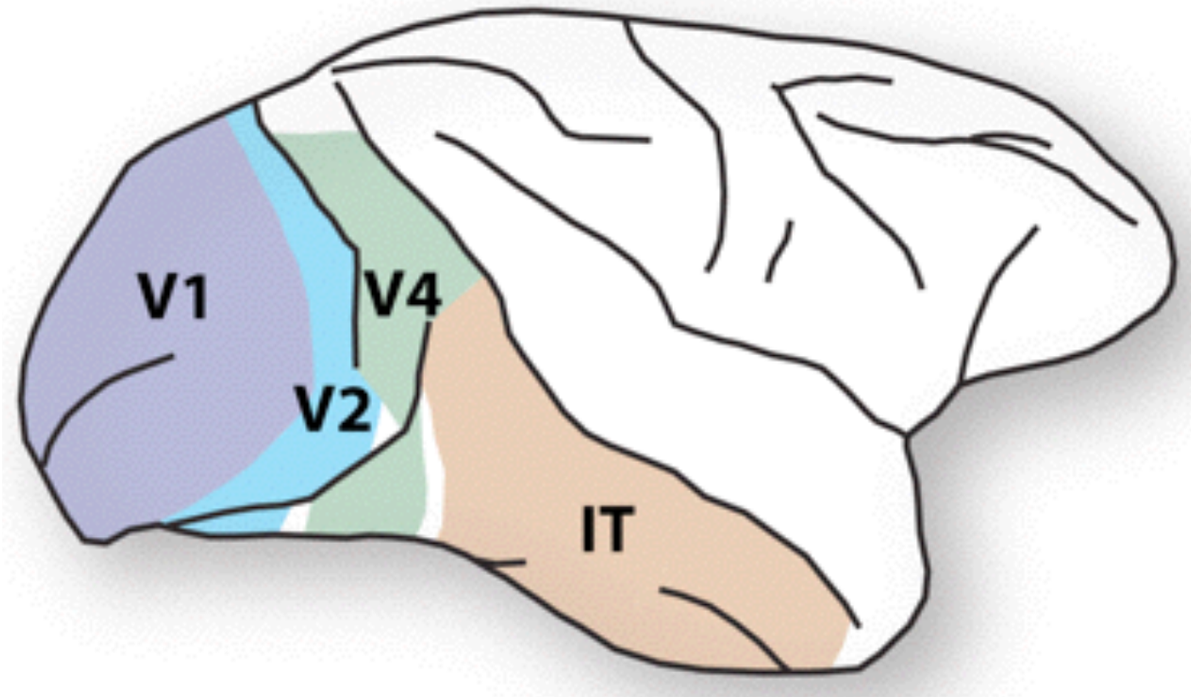

The visual system

Adapted from Motter and Mountcastle 1981

1. Foundations

- Fundamentals of image formation and digital images
- The visual system
- The computational approach to understanding vision

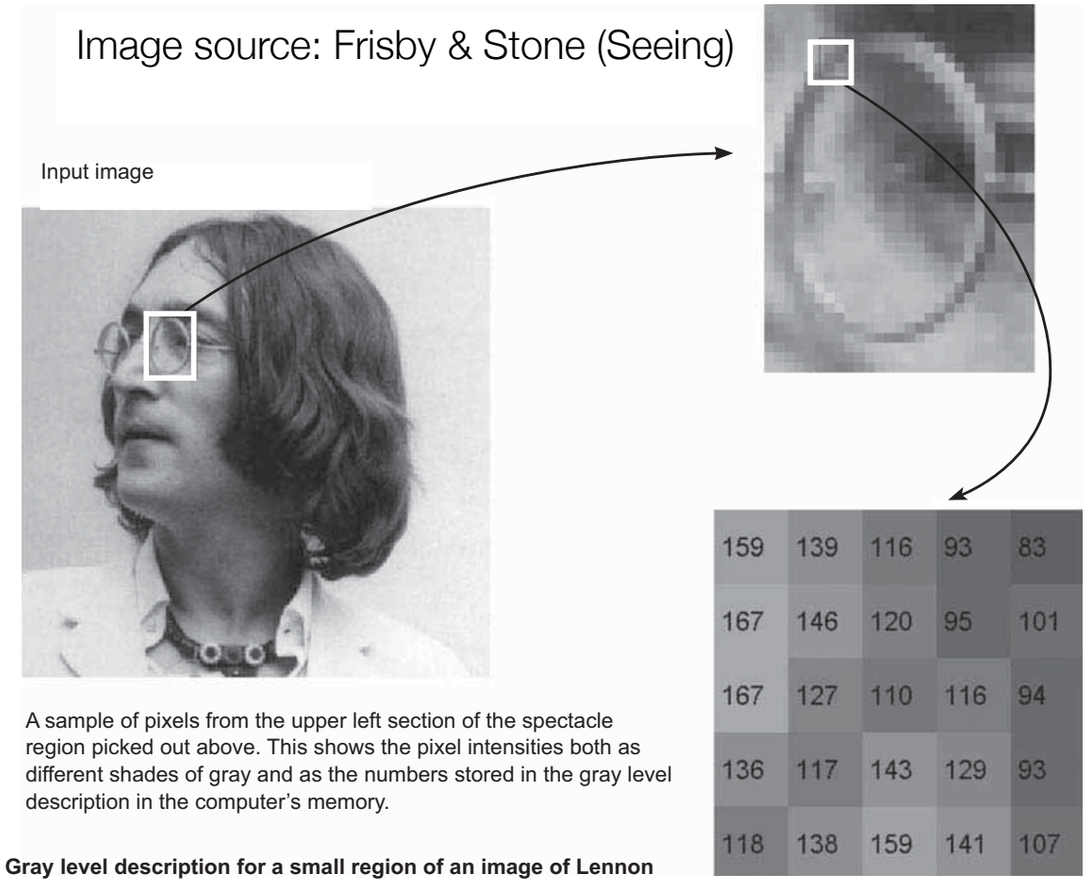

A sample of pixels from the upper left section of the spectacle region picked out above. This shows the pixel intensities both as different shades of gray and as the numbers stored in the gray level description in the computer's memory.

1.8 Gray level description for a small region of an image of Lennon

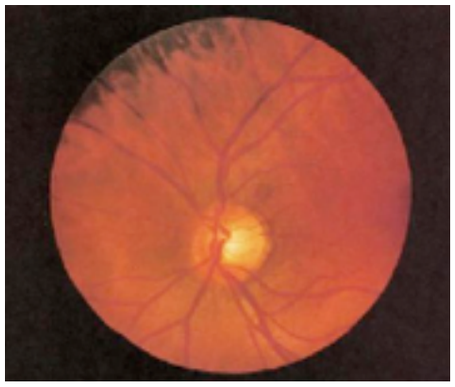

Image source: DiCarlo

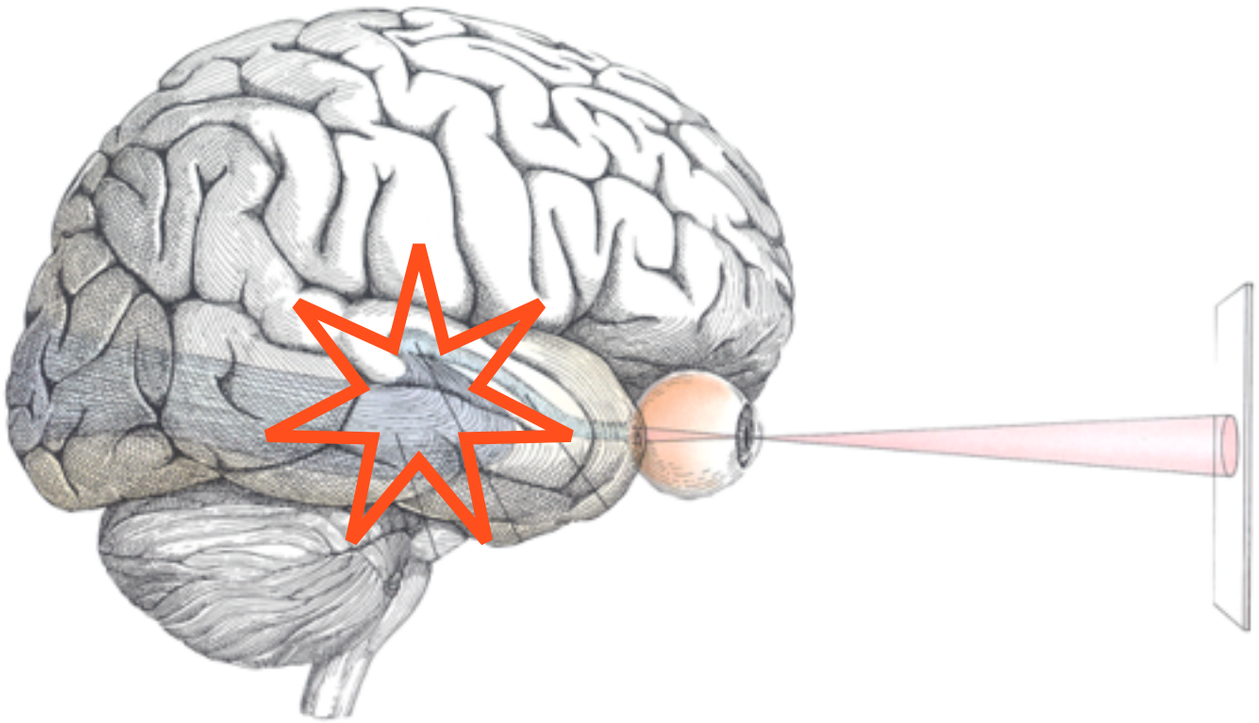

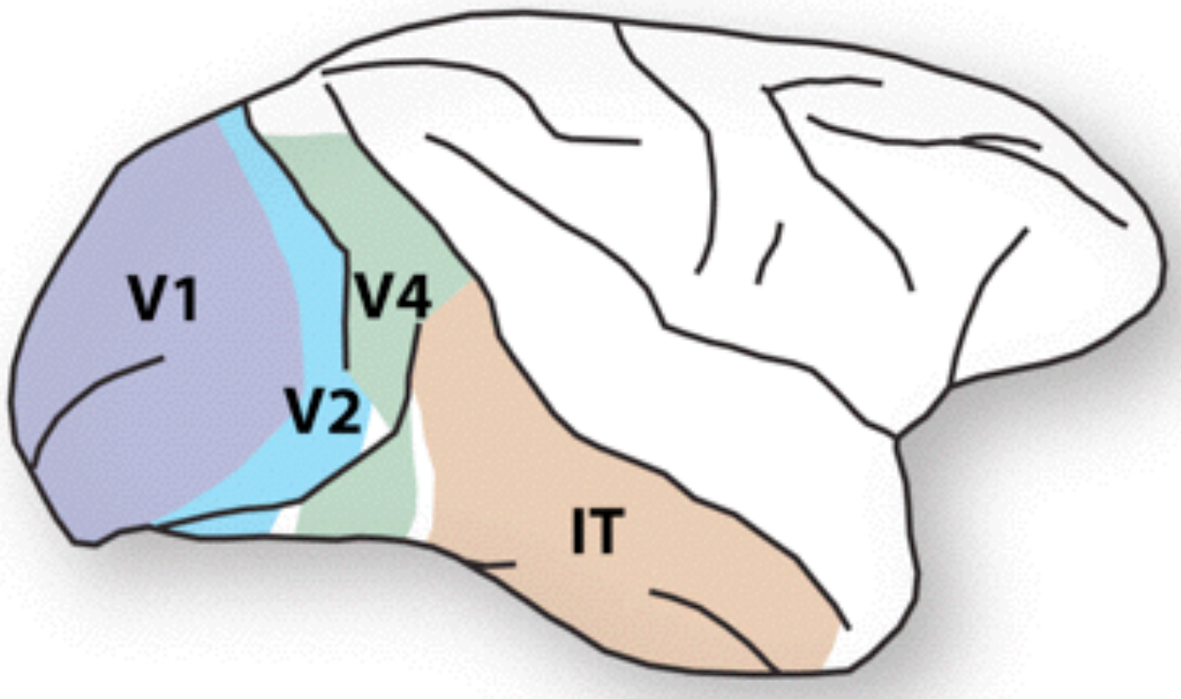

The visual system

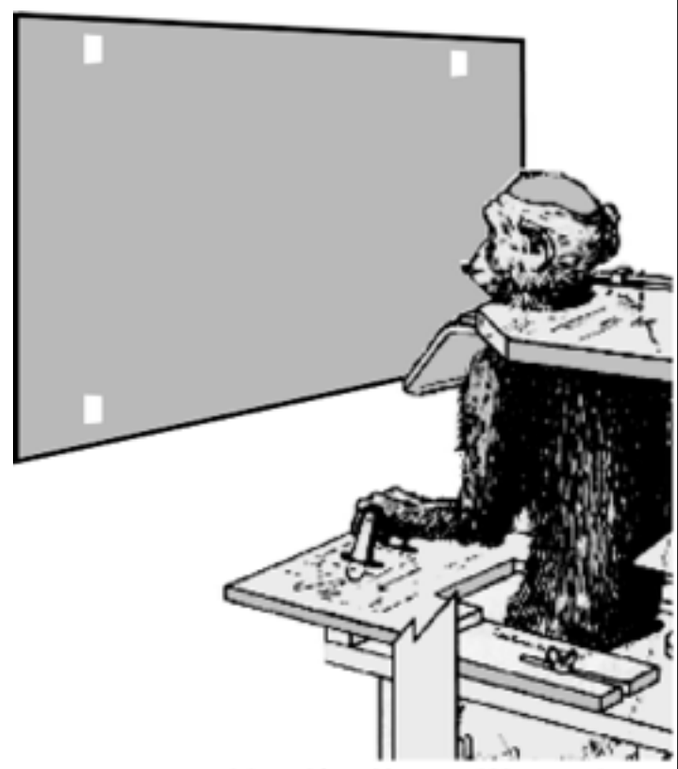

Adapted from Motter and Mountcastle 1981

2. Computing with LGN

- Basic anatomy and physiology
- Center-surround processing
- Color opponency channels
- Convolution and filtering

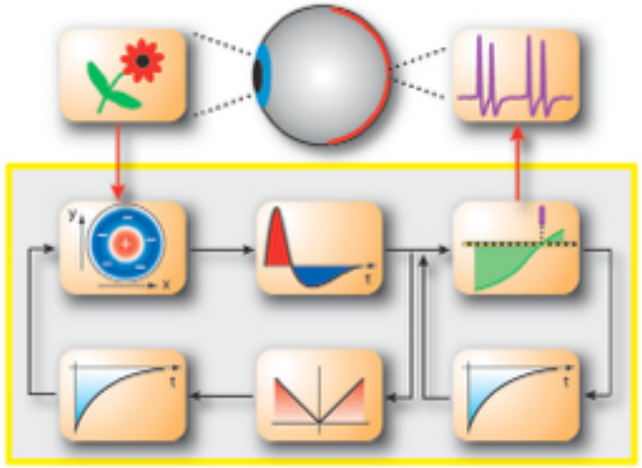

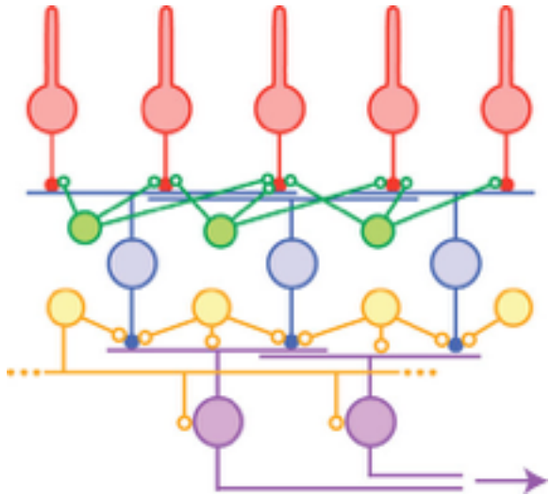

source: unknown

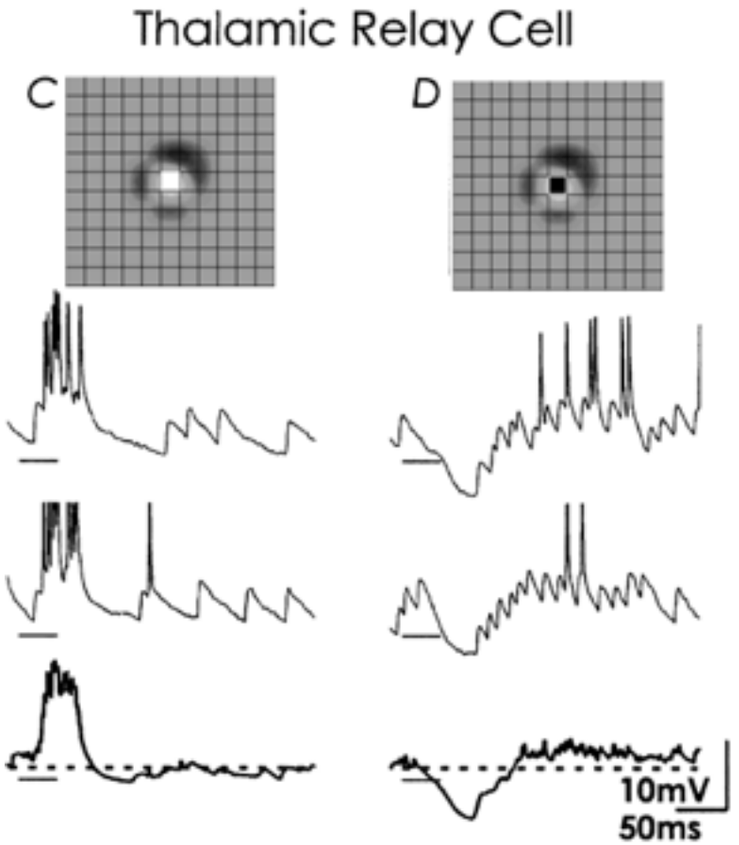

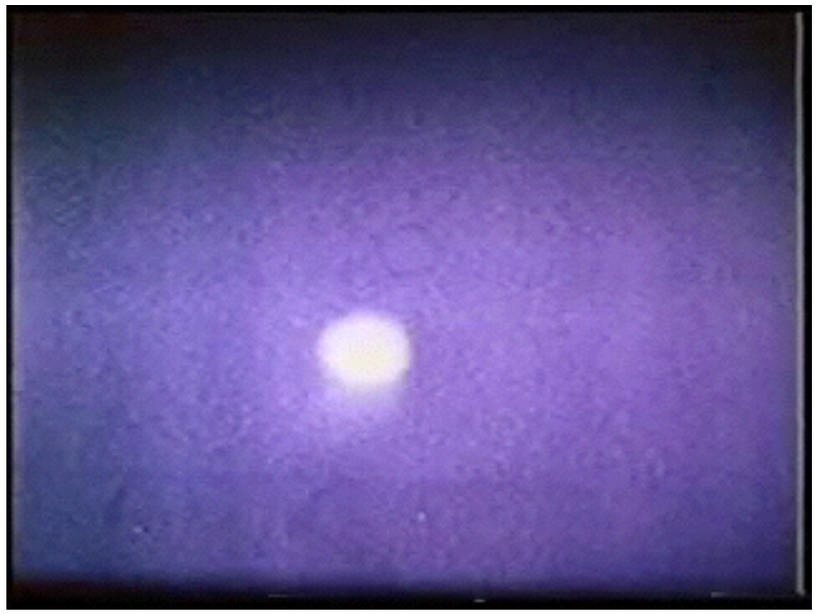

source: Hubel & Wiesel

Image source: DiCarlo

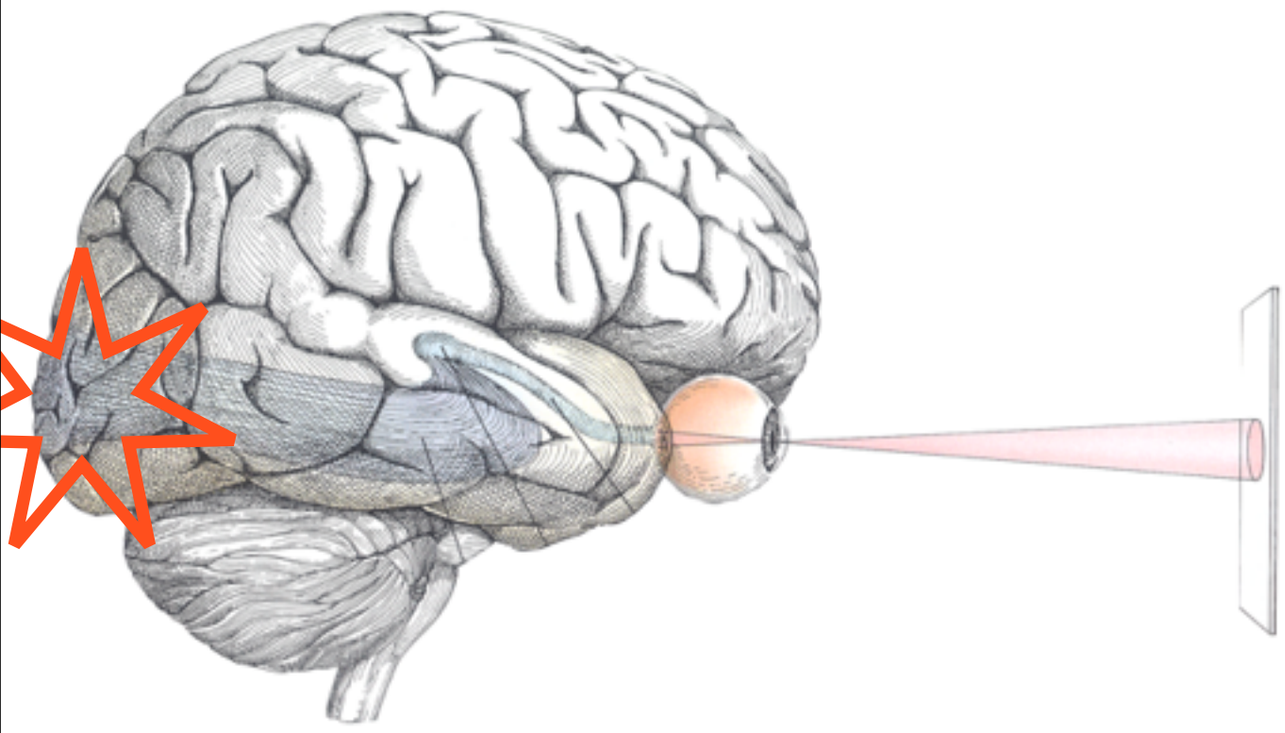

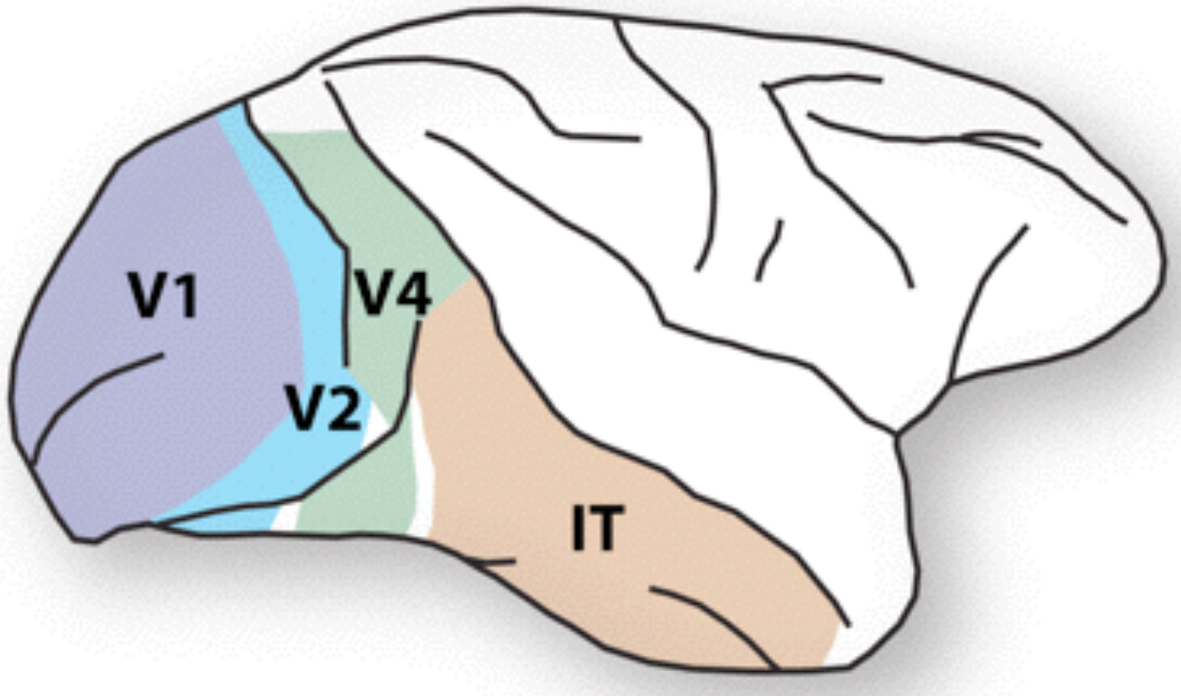

The visual system

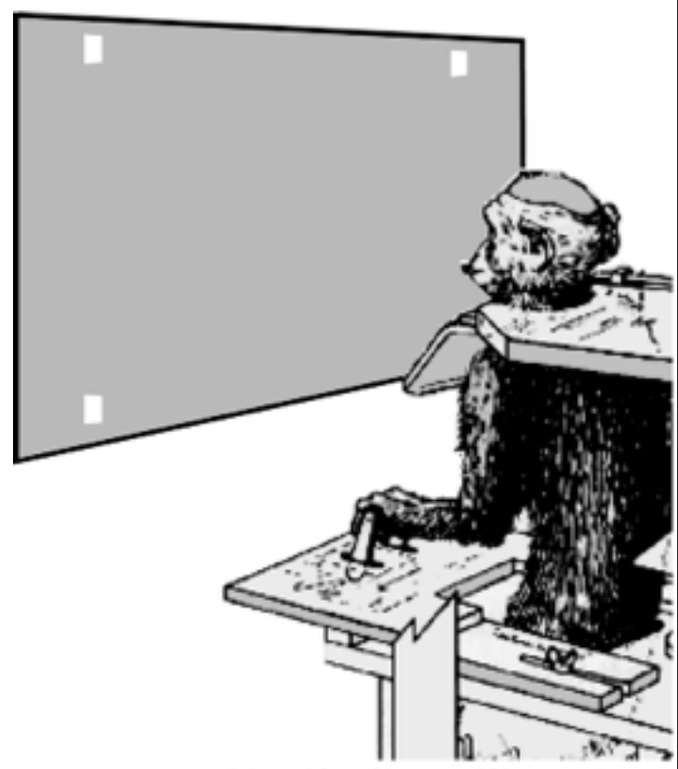

Adapted from Motter and Mountcastle 1981

3. Computing with V1

Orientation and spatial frequency

- Basic anatomy and physiology
- Simple and complex cells
- Edge detection
- Frequency analysis

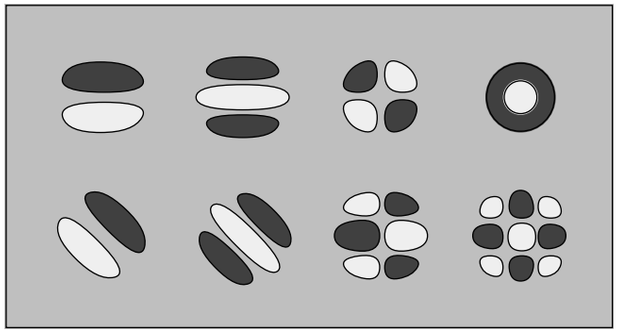

source: Aude Oliva

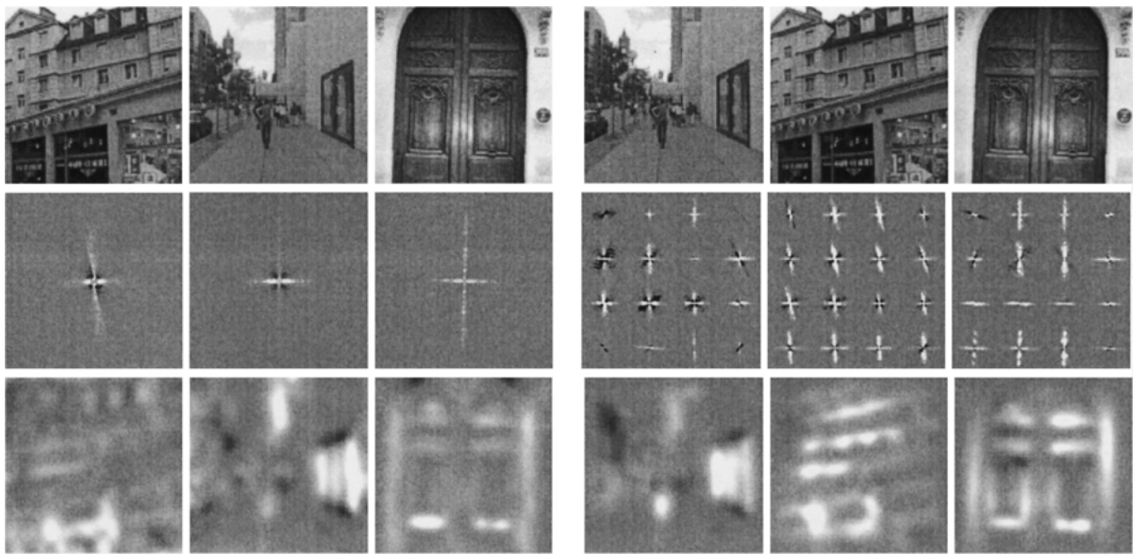

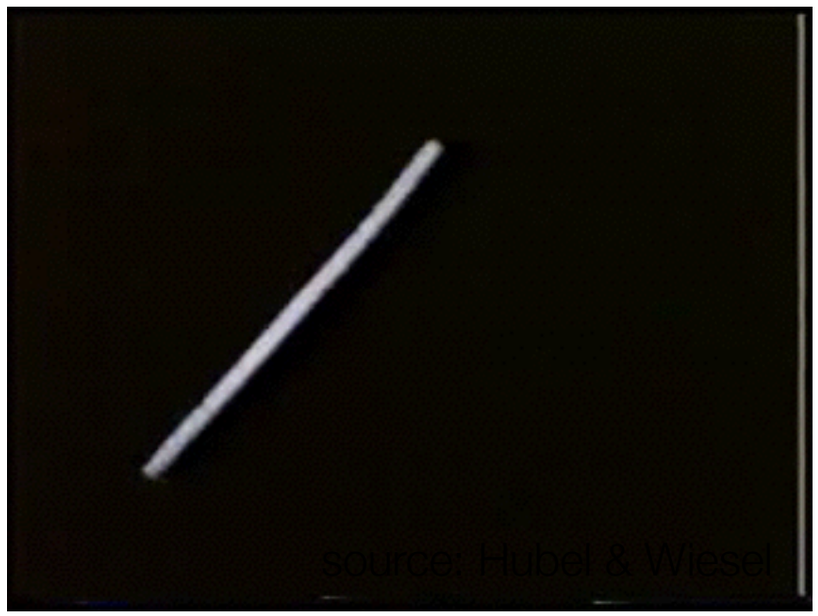

source: Hubel & Wiesel

4. Computing with V1

The coding perspective

- Natural image statistics
- Basics of dictionary learning
- Sparse coding

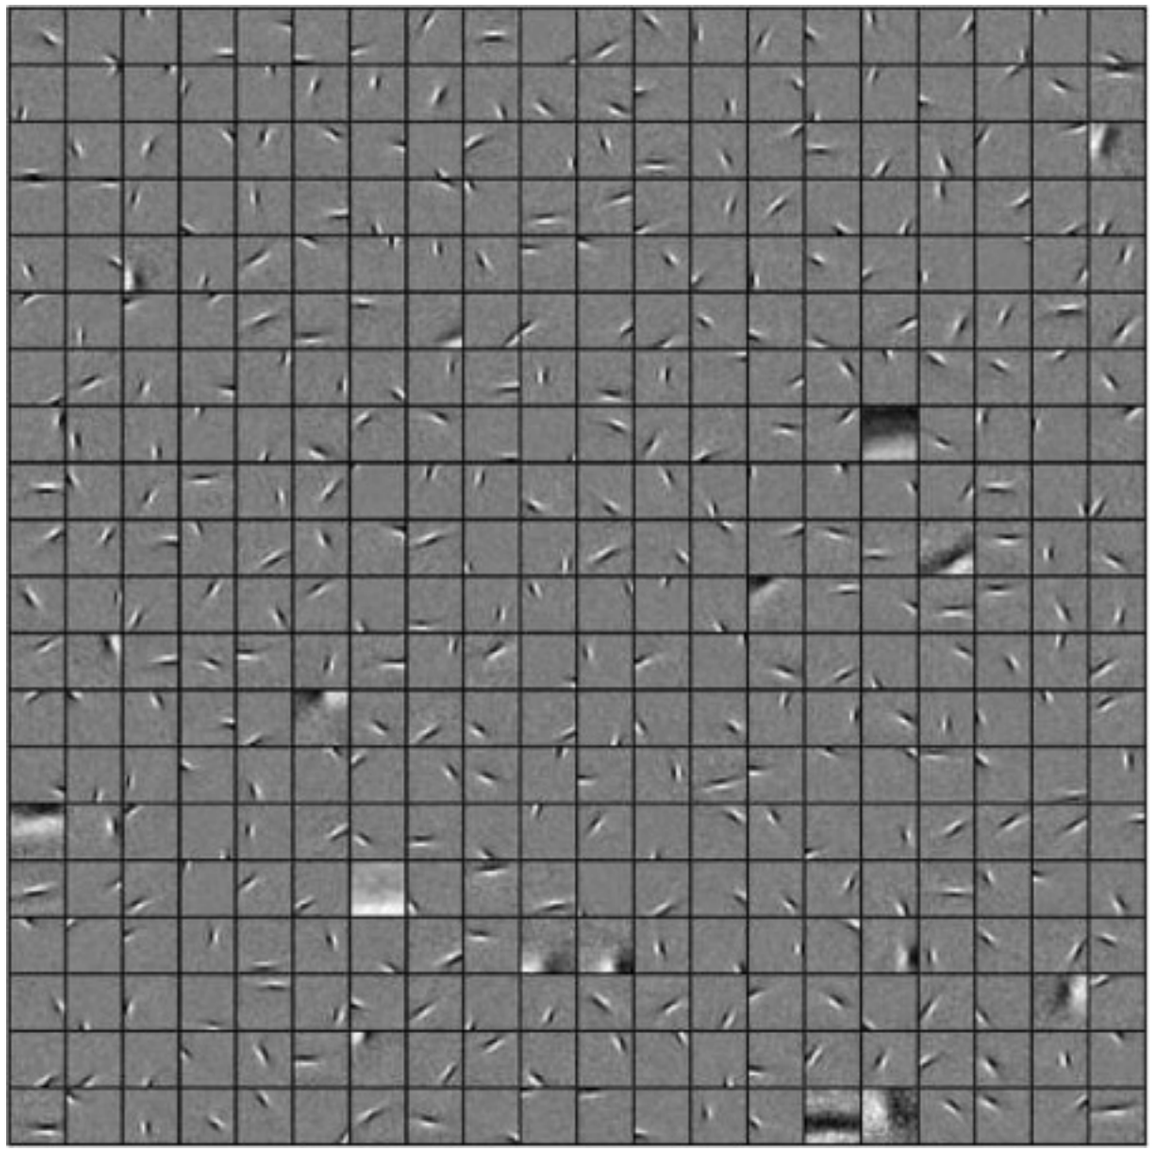

source: David Field

5. Computing with V1

Motion and stereo processing

- Motion and stereo as correspondence problems
- Depth perception
- Energy models

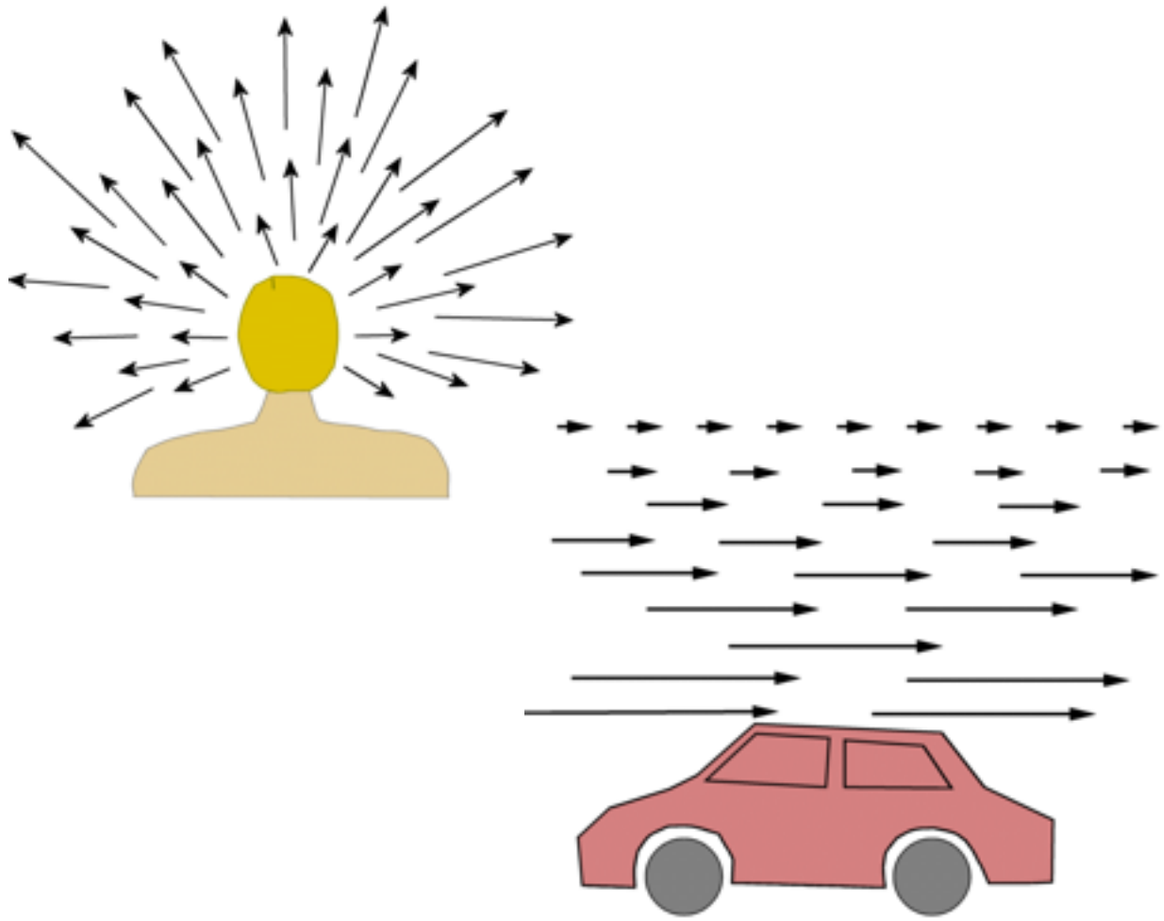

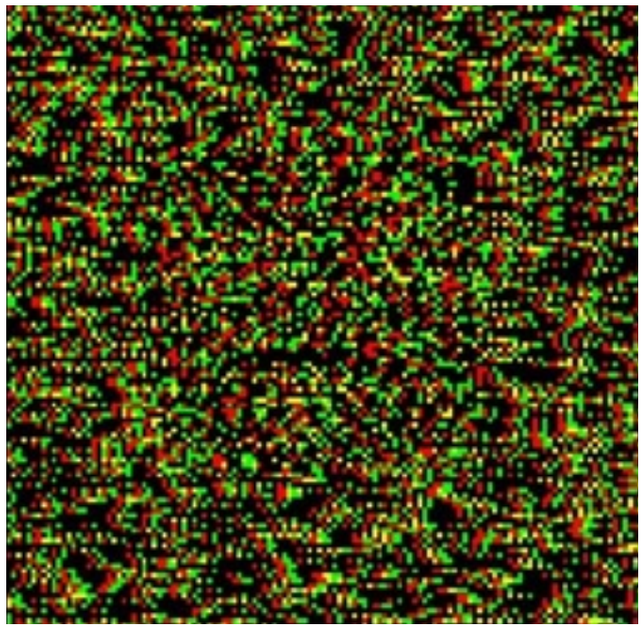

source: David Heeger

Image source: DiCarlo

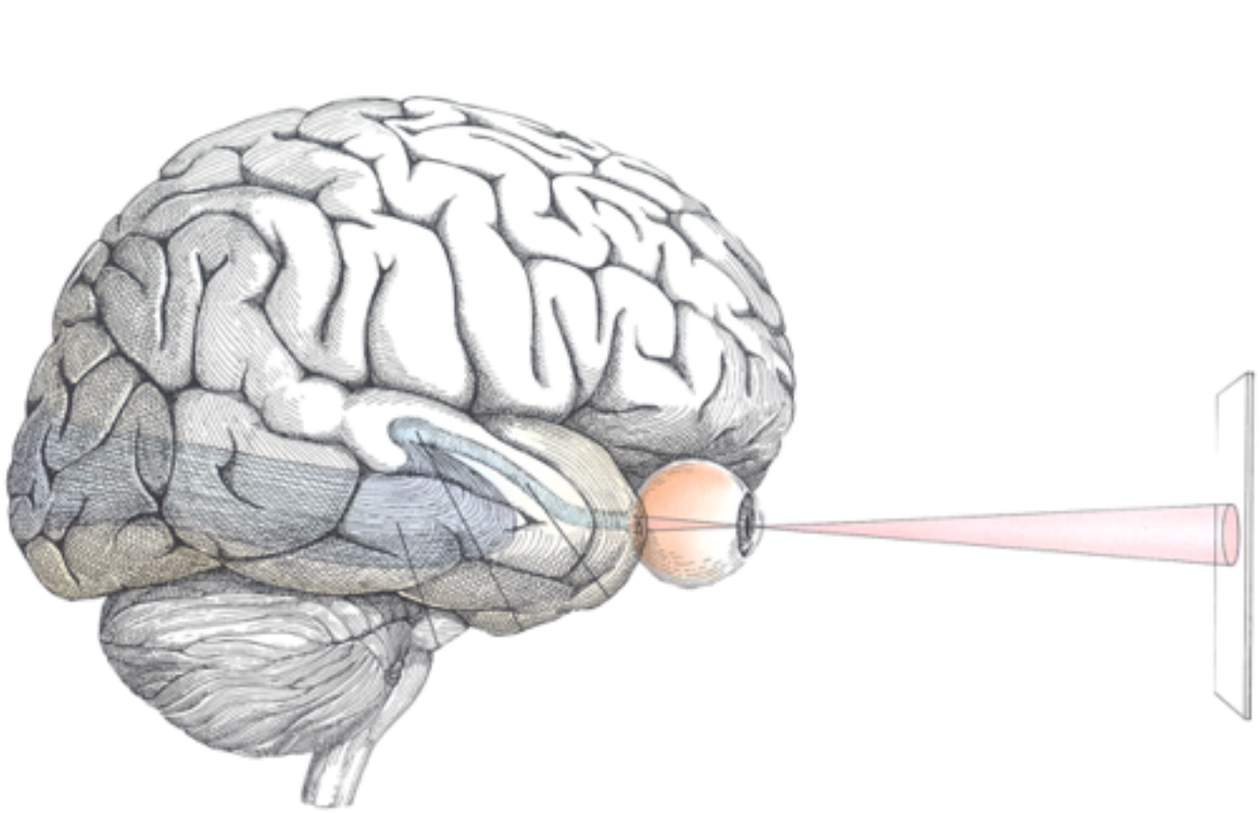

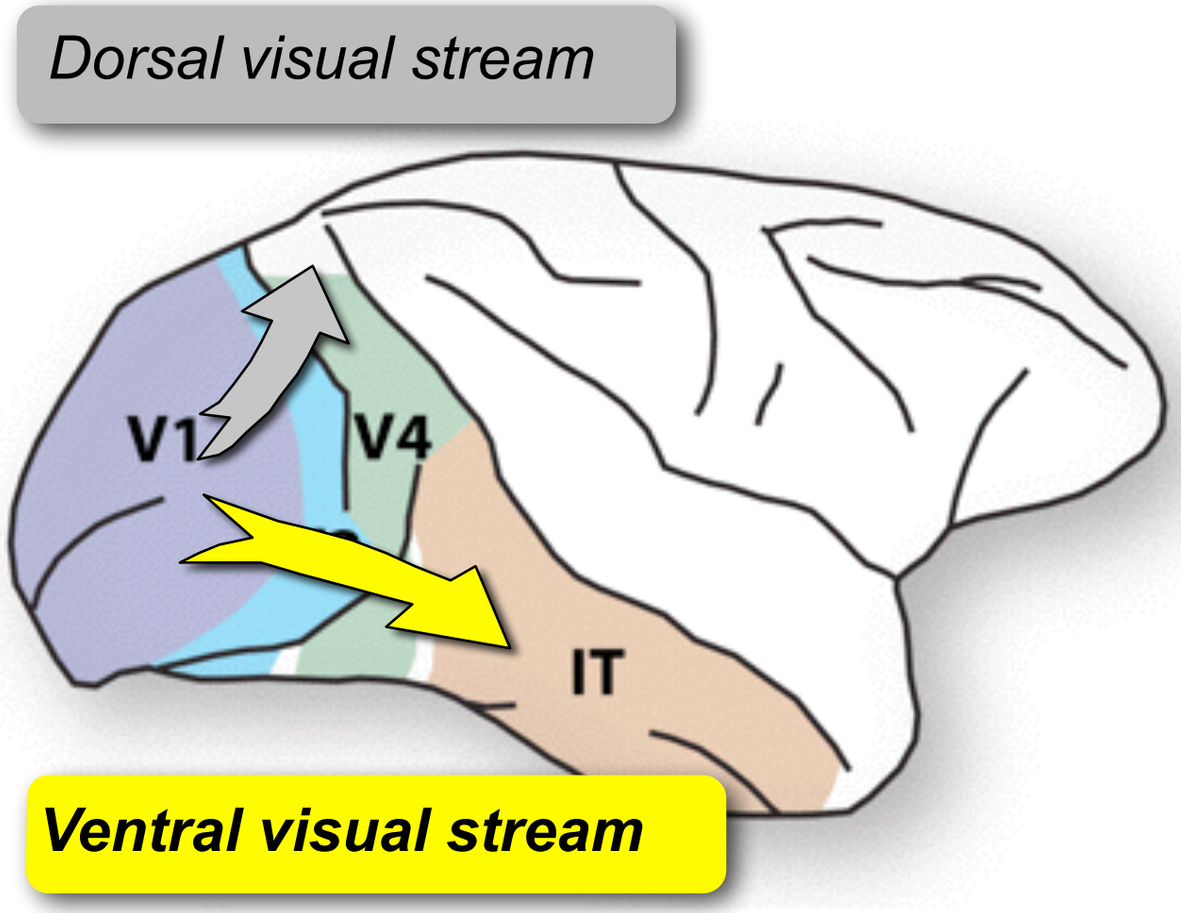

The visual system

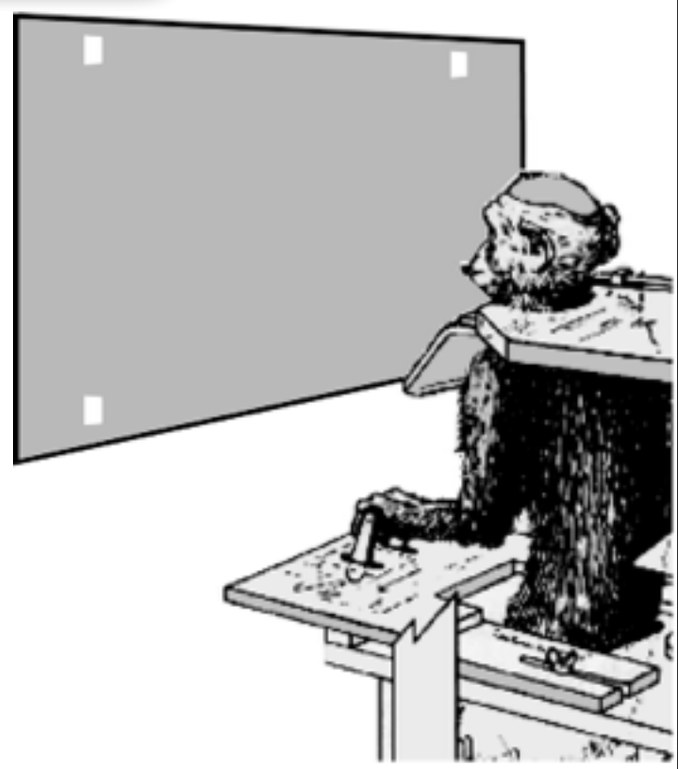

Adapted from Motter and Mountcastle 1981

6. Computing with the ventral and dorsal streams

Visual categorization

- Introduction to statistical learning theory (regression and classification)
- Levels of categorization

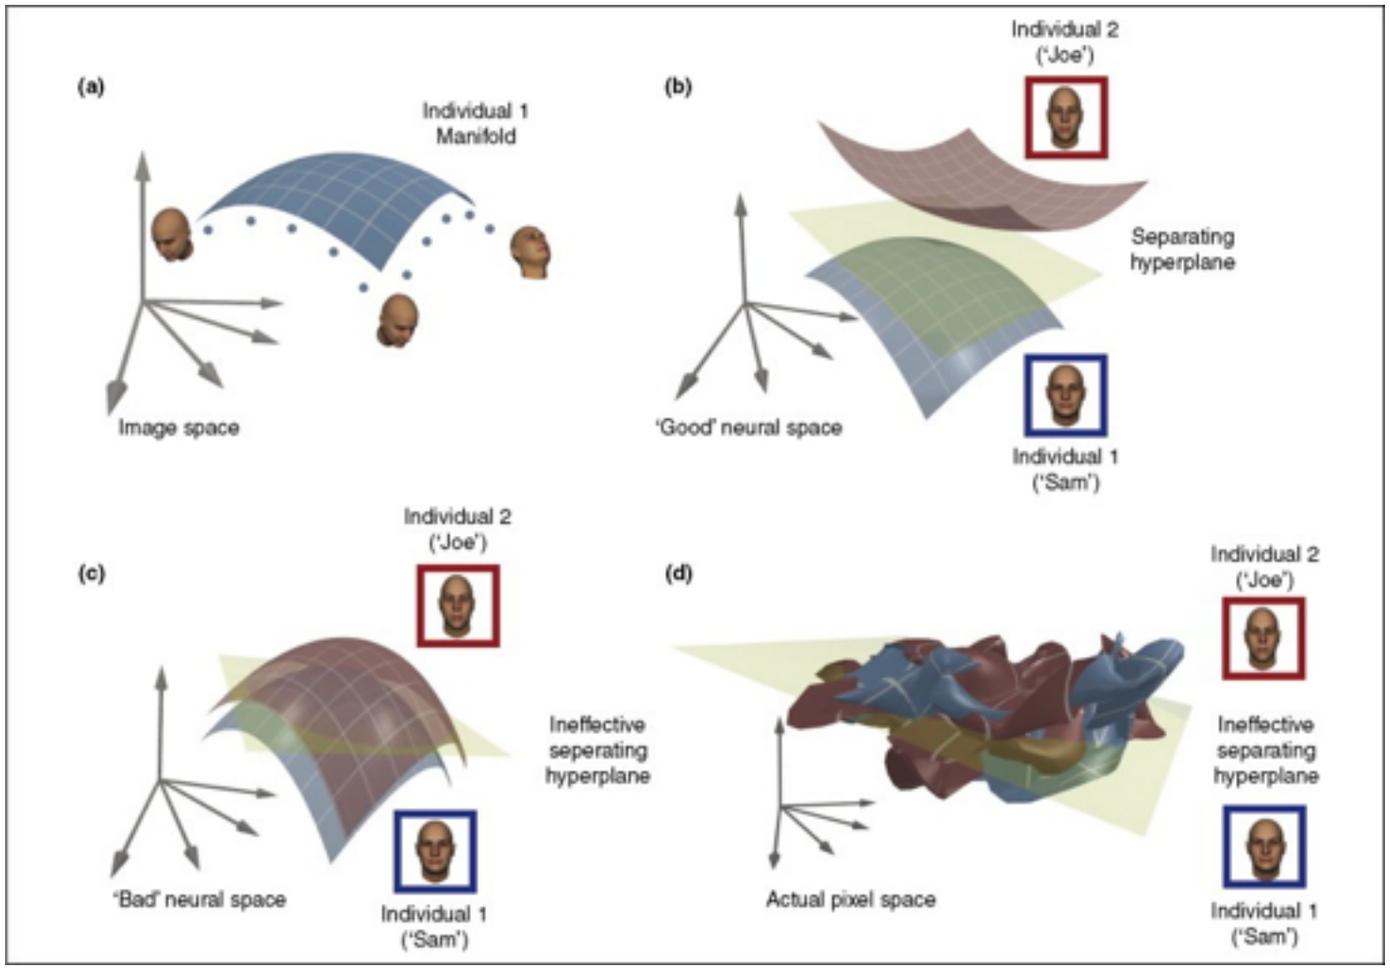

6. Computing with the ventral and dorsal streams

Invariant recognition

- Mid-level vision
- Object and action recognition
- Attention and eye movements
- Navigation

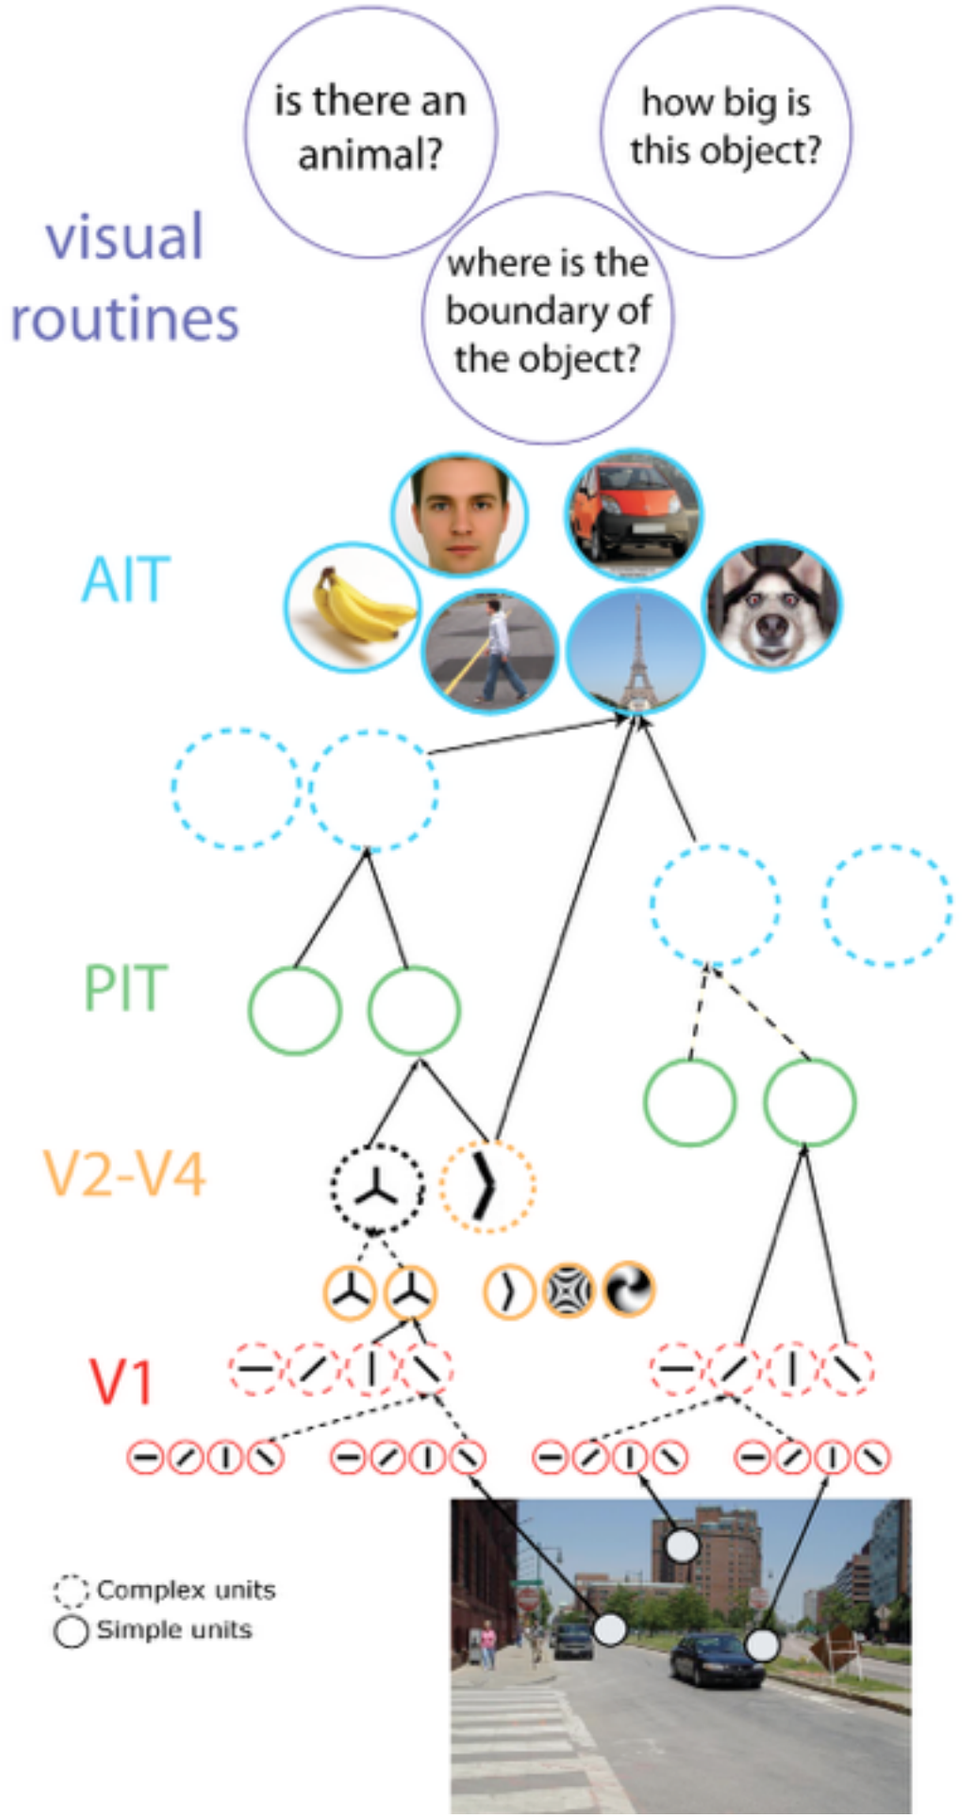

This course

- Brings together key pieces of neuroscience, computer vision, machine learning and human perception
  - Provides a toolkit (concepts, mathematical techniques, computational methods) and hands-on experience building and testing models of vision
- Exposes you to a framework for thinking about how vision works and how the brain is designed to solve specific problems
  - Helps you understand and specify levels of analysis
  - Focuses on modeling simple laboratory tasks from cognitive psychology and neuroscience

Synergies between brain and computer science

- Computer science provides:
  - Theoretical framework to study information processing in neural circuits
    - describing what nervous systems do and determining how they function
    - uncovering general principles by which they operate
  - Basic data analysis tools for neuroscience
    - decoding algorithms, information theory, etc
- Neuroscience provides novel approach to building intelligent machines by mimicking living brains!

Prerequisites

- This is an ***advanced undergrad / grad*** course!
- Assuming familiarity with:
  - ☑ Basic linear algebra (vectors, matrices)
  - ☑ Basic programing experience
  - ☑ Some prior exposure to neuroscience or cognitive science would be a plus but not mandatory
- Or willingness to do extra work!!
  - We will help you with lots of office hours and bootcamps (if needed)
- ***Please fill google form to help us plan!***

Requirements and grading

- Your grade: 90-100 ***A***, 80–90 ***B***, 70–80 ***C***, <70 ***NC***
  - 4 problem sets: 40% (up to 10% extra in bonus questions, counts towards participation)
  - 2 midterm exams: 25%
  - Final project (oral presentation): 25%
  - Participation (includes in class and on forum): 10%
- Academic integrity:
  - Discussion, collaboration, and research on the web are good
  - The work you turn in for credit must be your own
- ***Acknowledge your source or collaborators!***

MATLAB problems sets

- Content:
  - MATLAB programming
  - Experience working with real data of various kinds
- ~2 weeks to complete
- Policies for turning HWs late:
  - 5 days credit (free for everyone, use wisely)
  - After that, 20% off late penalty for each extra day

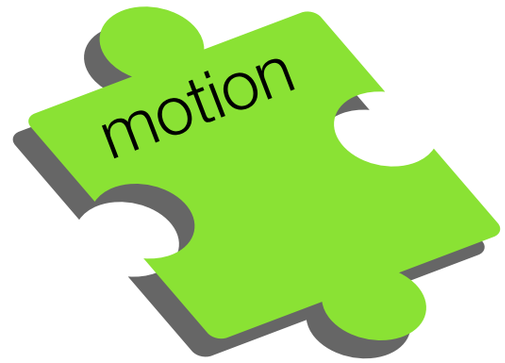

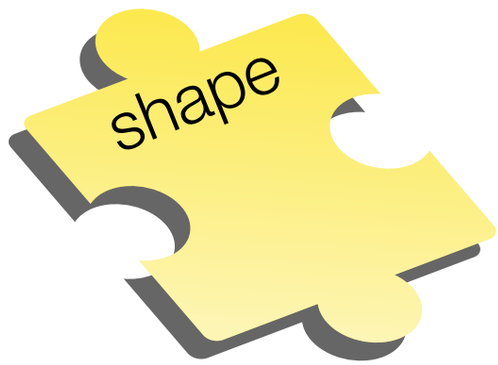

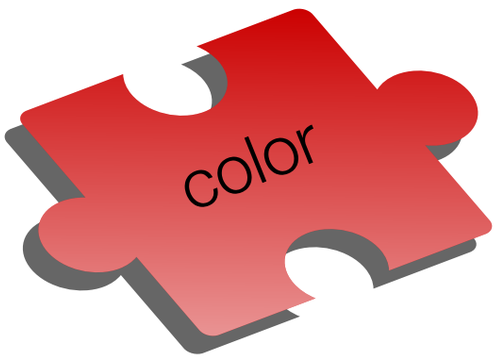

MATLAB problems sets

- Content:
  - MATLAB programming
  - Experience working with real data of various kinds
- ~2 weeks to complete
- Policies for turning HWs late:
  - 5 days credit (free for everyone, use wisely)
  - After that, 20% off late penalty for each extra day

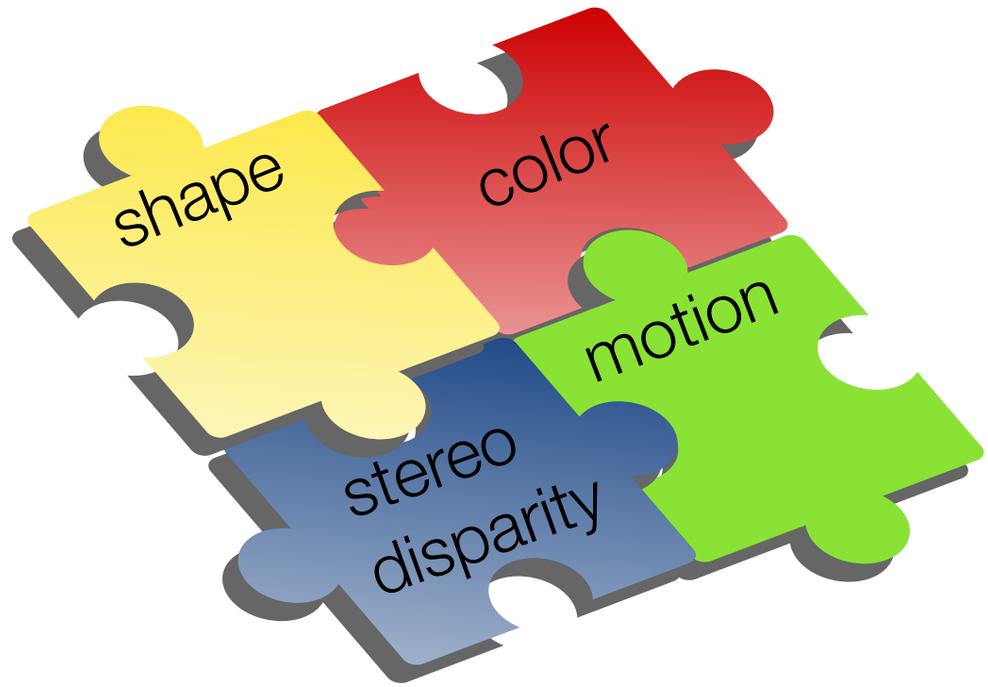

Final project

- Test artificial vision system on different vision tasks on the AR drone

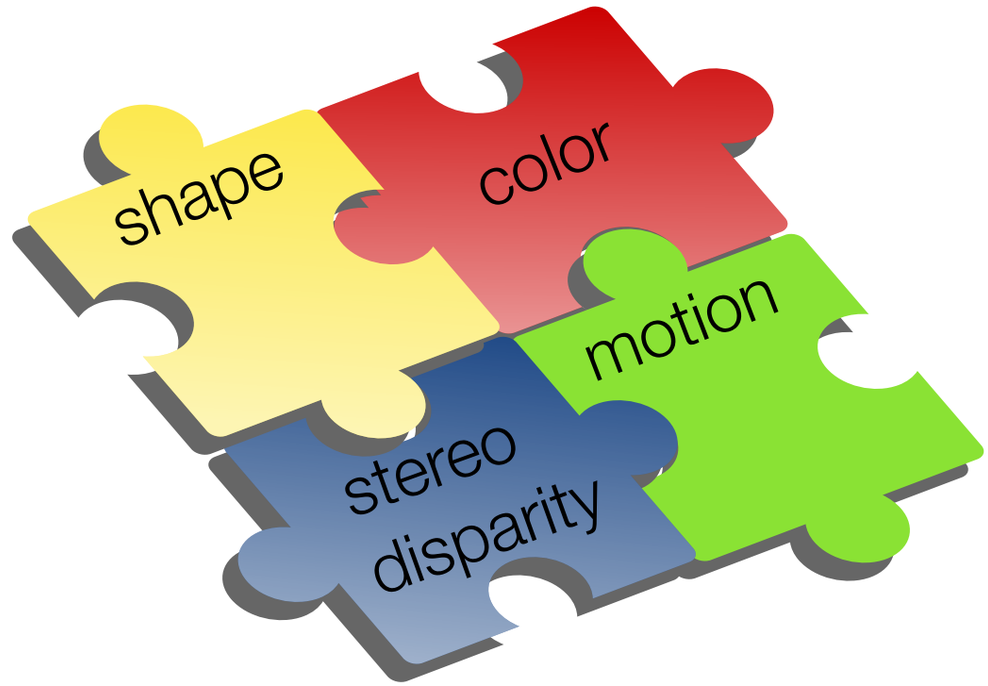

Discussion board in canvas

- Students are encouraged to use the discussion board to ask questions regarding Matlab and the material learned in class
- TAs will monitor discussion board to answer questions
- Students are encouraged to try to answer other student questions
- Answers will be considered for participation
- Discussion board should not be used to ask for the answers to your HW!

The course

- Lectures:
  - Slides and video recordings will be available online
- Readings:
  - At most 1–2 book chapters or articles per week
  - Most of the readings will be optional readings to complement the lectures

Recommended books (on reserve at the library)

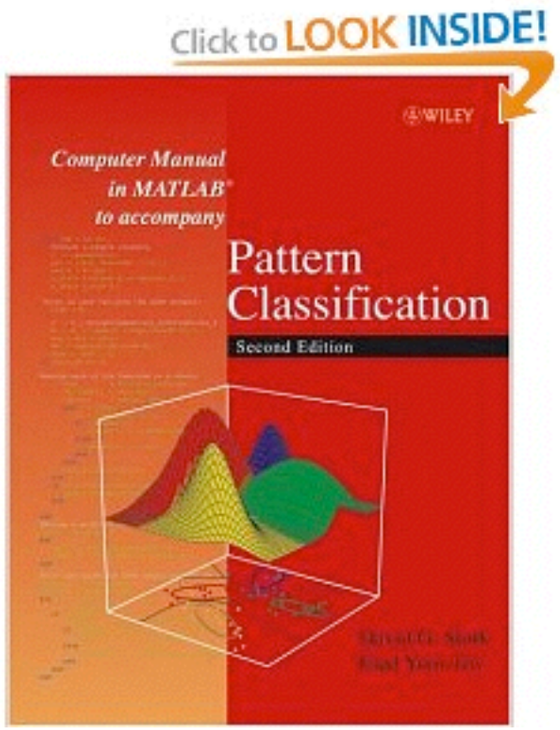

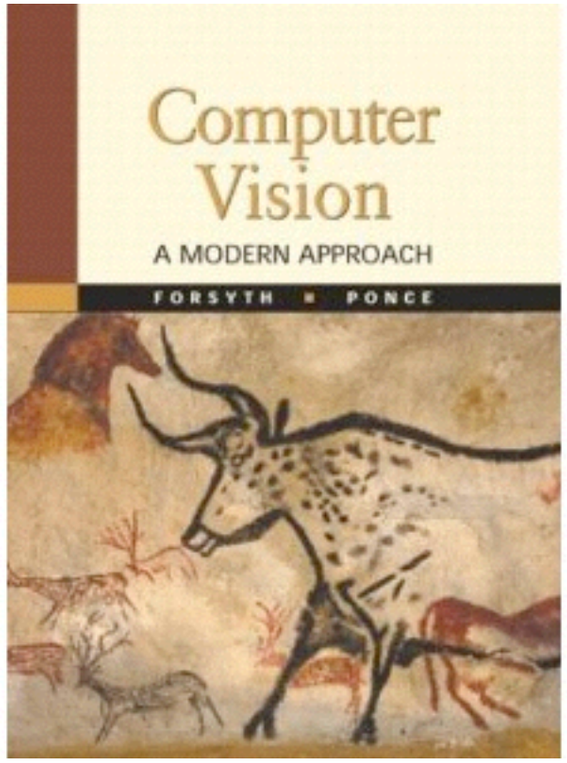

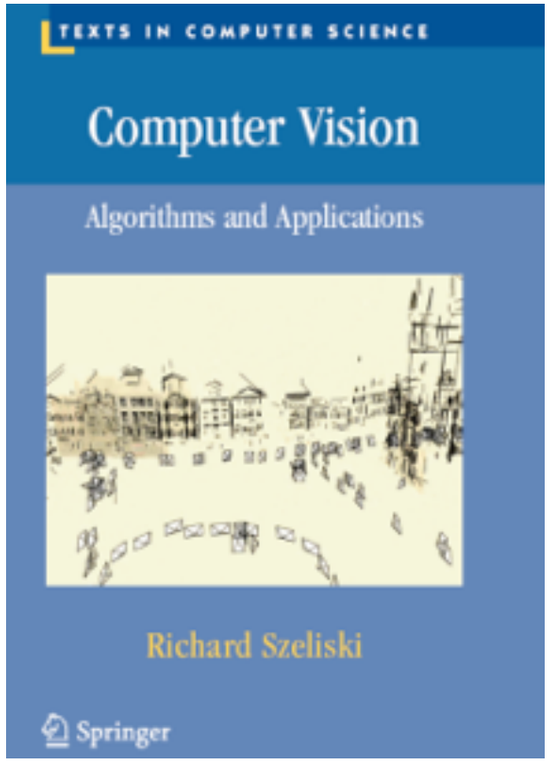

(available online – free)

Recommended books (on reserve at the library)

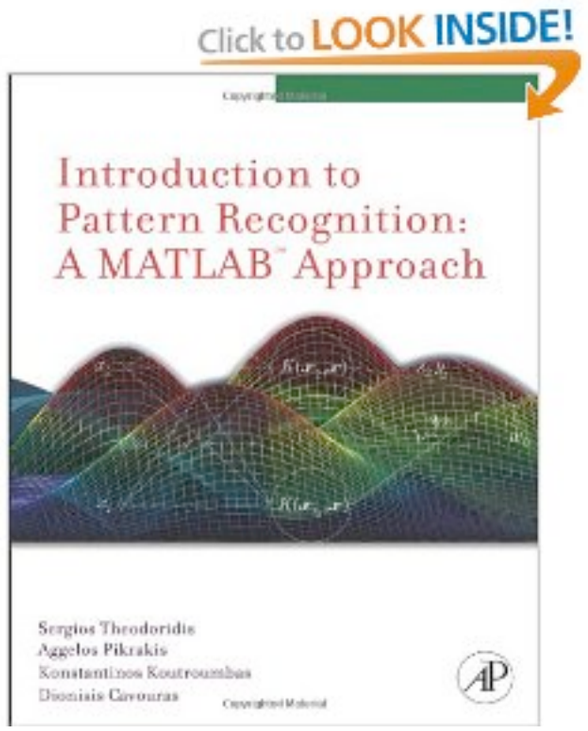

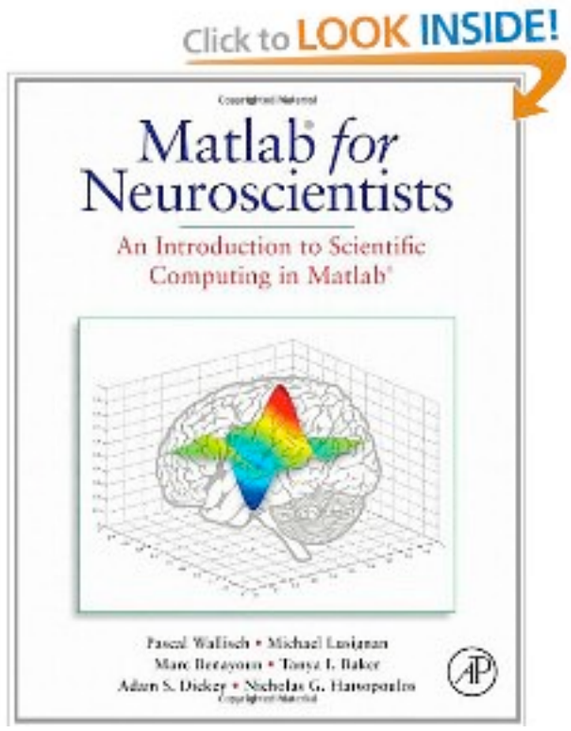

Recommended books (on reserve at the library)

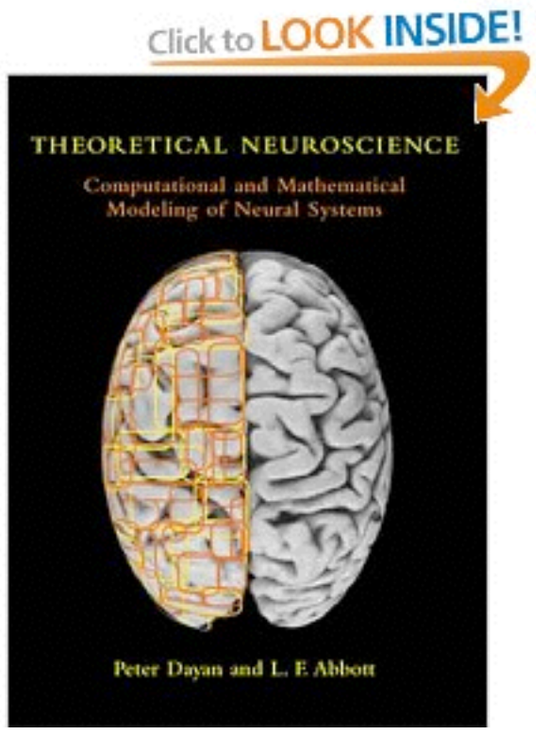

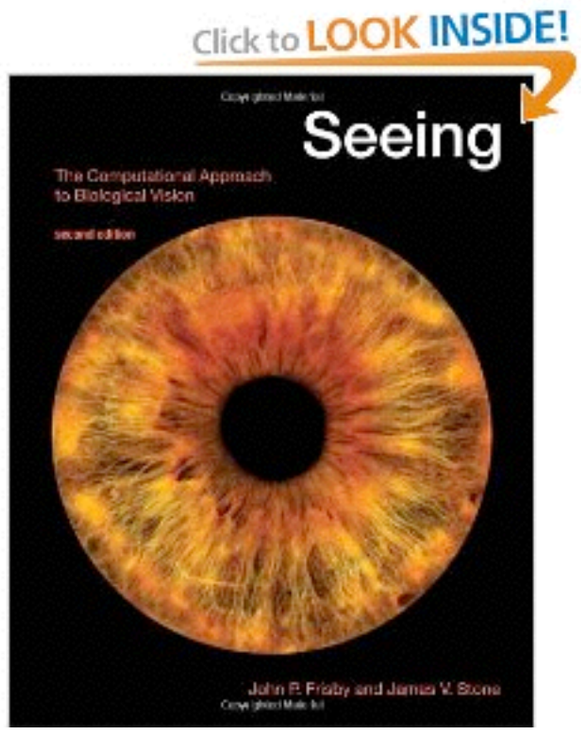

Overall

- Any problem accessing canvas?
- Questions about the course?
- If you are concerned about whether you have the proper background, please talk to me after class
- Auditing is ok but you need to be registered as an auditor