Lecture 1. Overview and Introduction

Reading Assignments:

Textbook: "Foundationd of Vision," Brian A. Wandell, Sinauer, 1995.

Read Introduction and browse through book

Organization

Lectures: Tuesdays, 2-4:50pm, VHE 206

Textbook: "Foundations of Vision," by Brian A. Wandell, Sinauer Associates, Inc. ISBN 0-87893-853-2

Office hour: Wed 4-6pm in HNB-30A

Homepage: http://iLab.usc.edu under "classes."

Grading: 3 units; simple five-minute quizzes at the beginning of each lecture, and one project.

Motivation behind this new course

- Introduce computer science students to neuroscience methods and research
- Introduce neuroscience students to computer, mathematical and signal/image processing methods and research
- Establish a parallel between both disciplines and foster new cross-disciplinary ideas
- Topical focus: vision
- Demonstrate validity of our cross-disciplinary approach through application examples

Typical approach

[1] describe major challenges associated with a particular aspect of vision, analyze them using general mathematical, physics, and signal processing tools;

[2] Survey state of the art computer vision and image processing algorithms which give best performance at solving those vision challenges, irrespectively of their biological plausibility;

[3] Survey latest advances in neurobiology (including electrophysiology, psychophysics, fMRI and other experimental techniques, as well as theory and brain modeling) relevant to those vision challenges, and analyze these findings in computational terms;

[4] Derive a global view of the problem from a critical comparison between the computer algorithms and neurobiological findings studied.

For issues mostly studied in computational neuroscience, and for which computer vision algorithms are just emerging and inspired from neuroscience: **[1] [3] [2] [4]**.

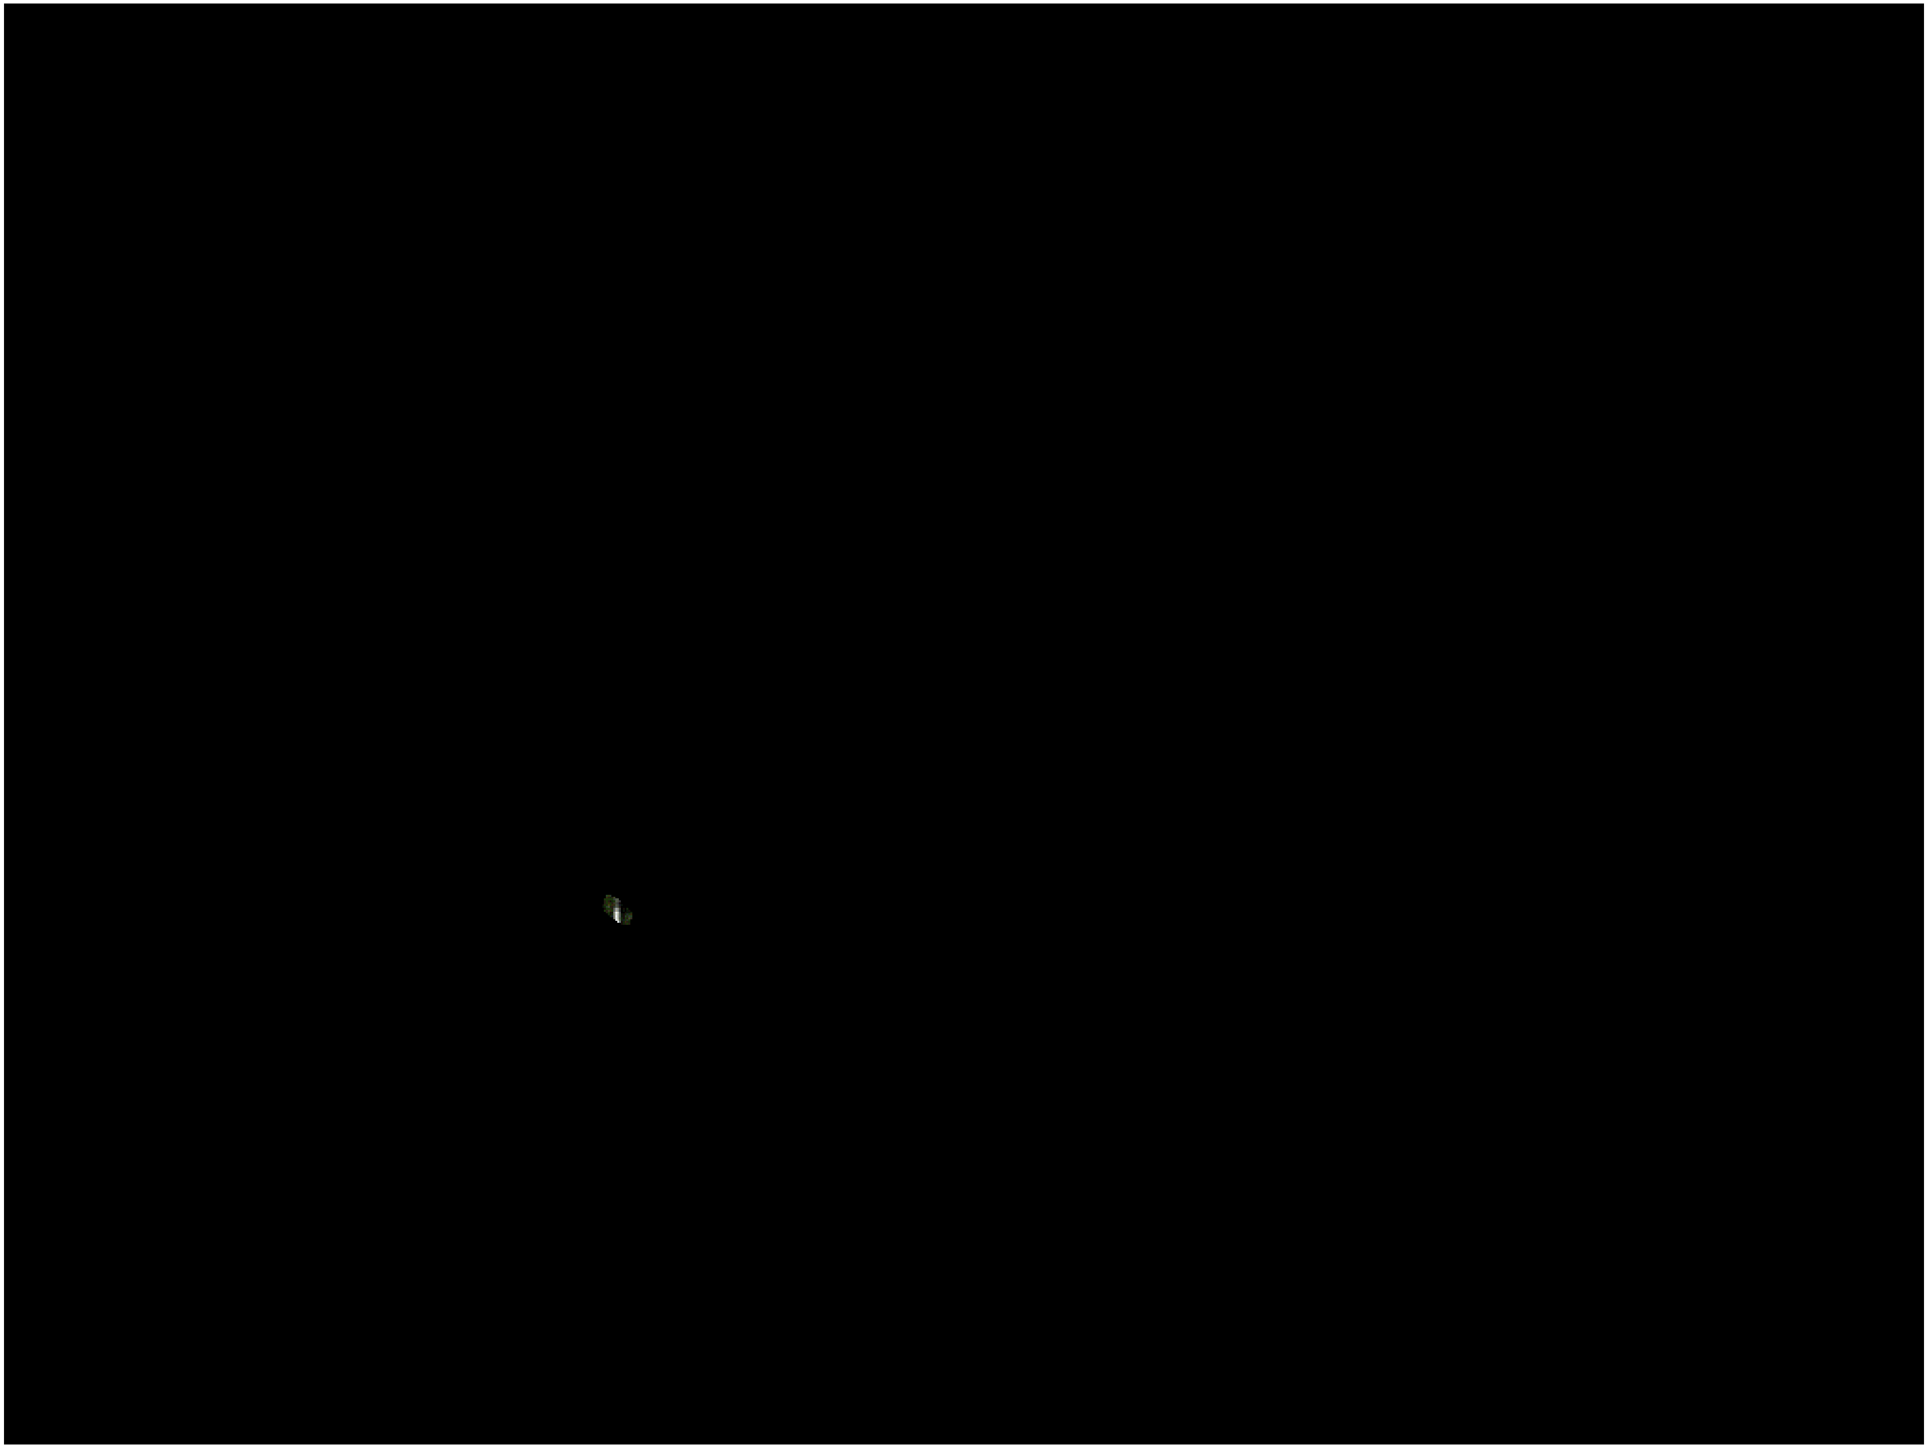

How can we "see"?

Vision as a progressive change in representation

Marr (1982): through 2 ½ D primal sketch

In the class textbook:

- Part 1: Encoding
- Part 2: Representation
- Part 3: Interpretation

Vision and the brain

Ryback et al, 1998

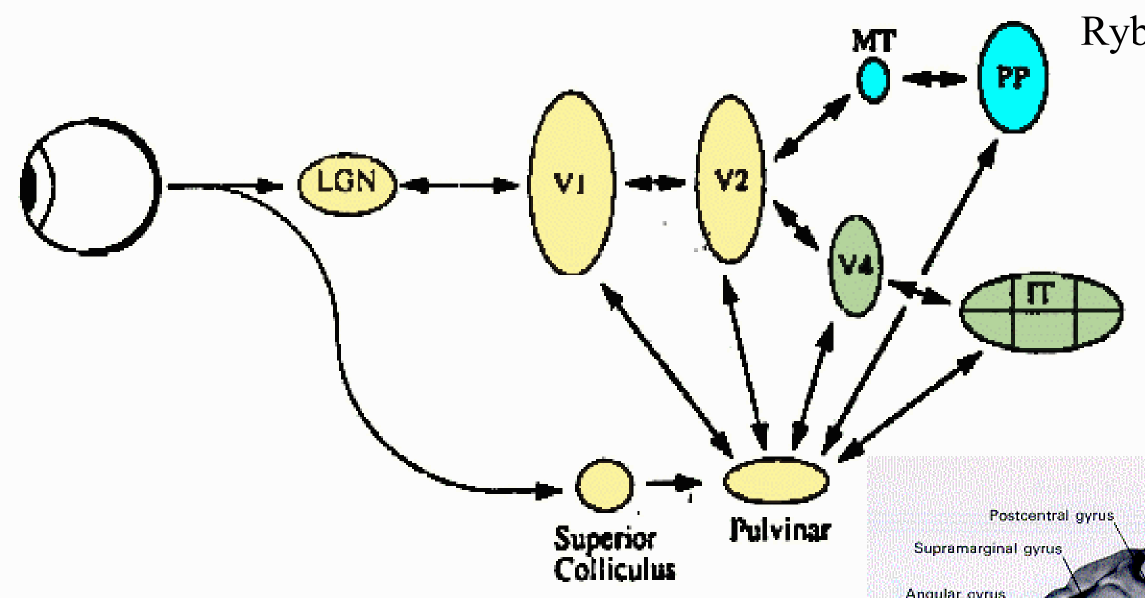

Roughly speaking, about half of the brain is concerned with vision. Although most of it is highly automated and unconscious, vision hence is a *major component of brain function.*

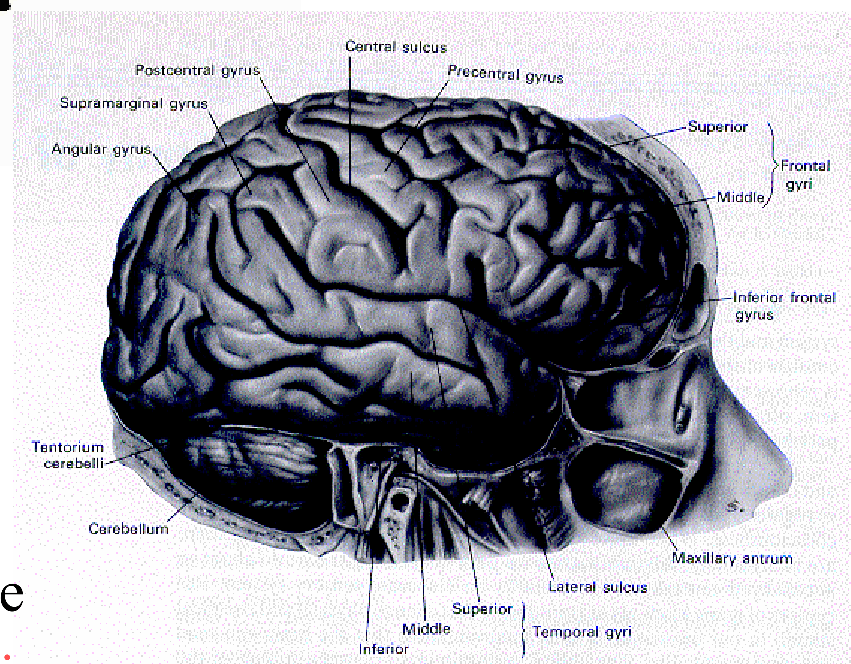

Figure 2.1. Lateral view of the brain exposed in the skull to show topographical relationships. (From Mettler's *Neuroanatomy*, 1948; courtesy of The C. V. Mosby Company.)

Vision, AI and robots

1940s: beginning of Artificial Intelligence

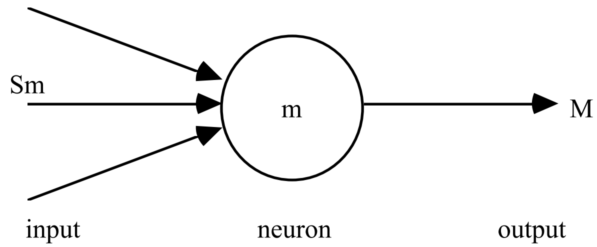

McCullogh & Pitts, 1942

$$\Sigma_i \; w_i x_i \geq \theta$$

Perceptron learning rule (Rosenblatt, 1962)

Backpropagation

Hopfield networks (1982)

Kohonen self-organizing maps

…

Vision, AI and Robots

1950s: beginning of computer vision

Aim: *give to machines same or better vision capability as ours*

Drive: AI, robotics applications and factory automation

Initially: passive, feedforward, layered and hierarchical process that was just going to provide input to higher reasoning processes (from AI)

But soon: realized that could not handle real images

1980s: Active vision: make the system more robust by allowing the vision to adapt with the ongoing recognition/interpretation

Syllabus Overview

This is tentative and still open to suggestions!

- ***Course Overview and Fundamentals of Neuroscience.***
- ***Neuroscience basics.***
- ***Experimental techniques in visual neuroscience.***
- ***Introduction to vision.***
- ***Low-level processing and feature detection.***
- ***Coding and representation.***
- ***Stereoscopic vision.***
- ***Perception of motion.***
- ***Color perception.***
- ***Visual illusions.***
- ***Visual attention.***
- ***Shape perception and scene analysis.***
- ***Object recognition.***
- ***Computer graphics, virtual reality and robotics.***

Syllabus Overview

- ***Course Overview and Fundamentals of Neuroscience.***

  - why is vision hard while it seems so naturally easy?
  - why is half of our brain primarily concerned with vision?
  - Towards domestic robots: how far are we today?
  - What can be learned from the interplay between biology and computer science?

Syllabus Overview

◆ ***Neuroscience basics.***

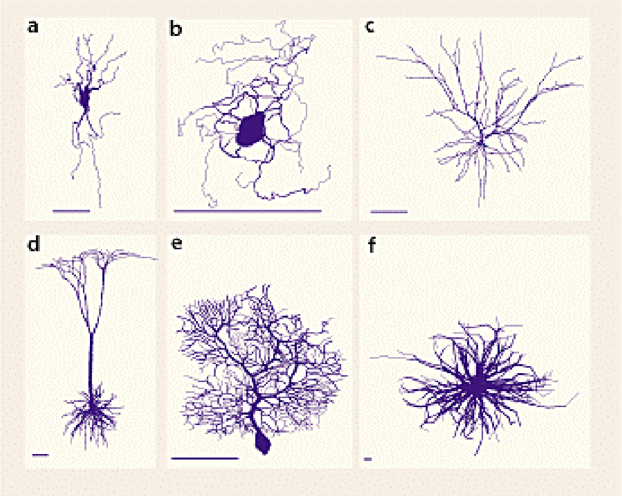

- The brain, its gross anatomy
- Major anatomical and functional areas
- The spinal cord and nerves
- Neurons, different types
- Support machinery and glial cells
- Action potentials
- Synapses and inter-neuron communication
- Neuromodulation
- Power consumption and supply
- Adaptability and learning

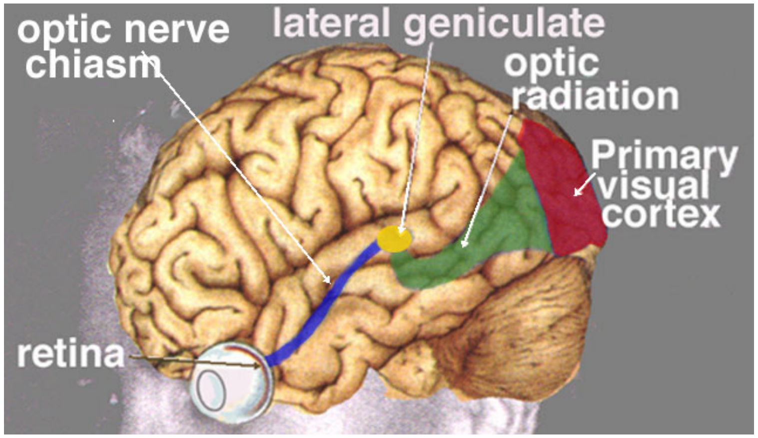

Syllabus Overview

◆ ***Experimental techniques in visual neuroscience***

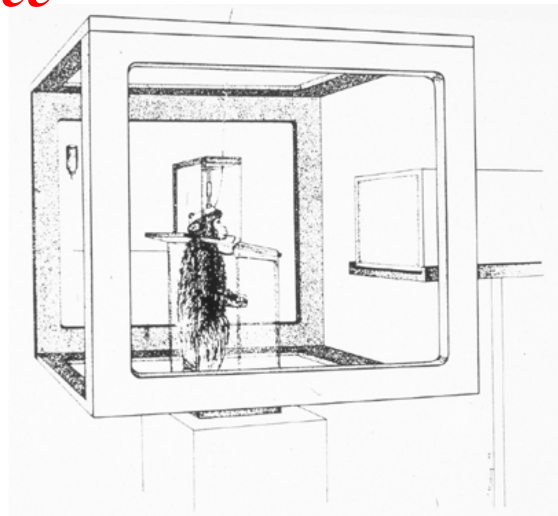

- Recording from neurons: electrophysiology
- Multi-unit recording using electrode arrays
- Stimulating while recording
- Anesthetized vs. awake animals
- Single-neuron recording in awake humans
- Probing the limits of vision: visual psychophysics
- Functional neuroimaging: Techniques
- Experimental design issues
- Optical imaging
- Transcranial magnetic stimulation

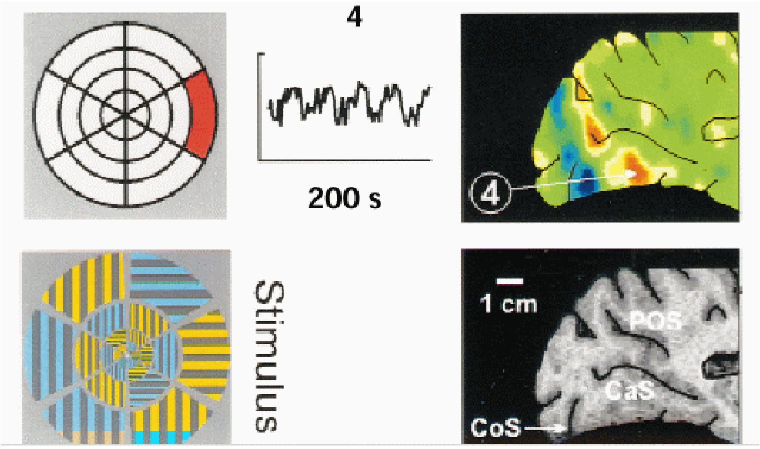

Syllabus Overview

◆*Introduction to vision.*

- Biological eyes compared to cameras and VLSI sensors
- Different types of eyes
- Optics
- Theoretical signal processing limits
- Introduction to Fourier transforms, applicability to vision
- The Sampling Theorem
- Experimental probing of theoretical limits
- Phototransduction
- Retinal organization
- Processing layers in the retina
- Adaptability and gain control.

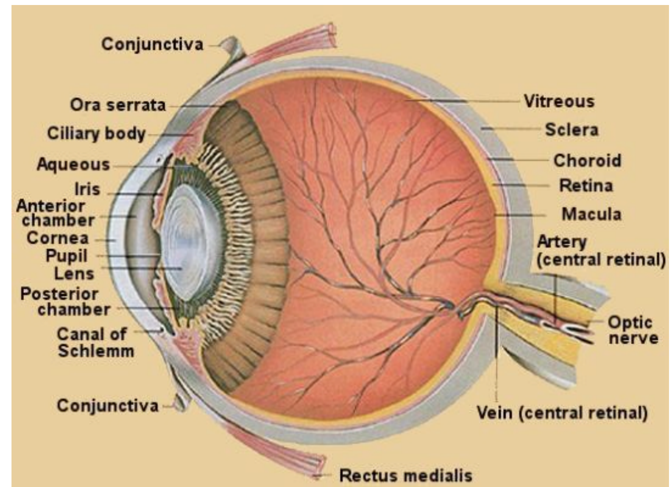

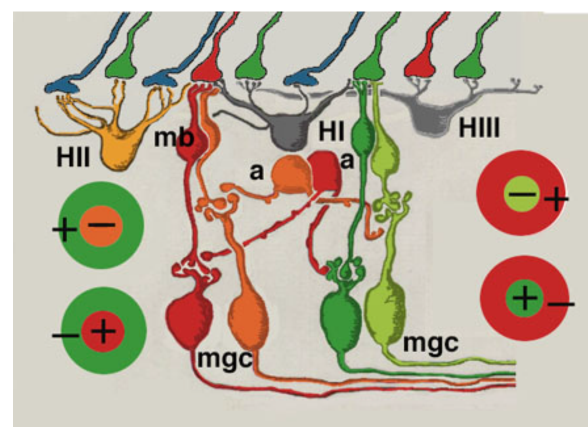

Syllabus Overview

◆ ***More Introduction to Vision.***

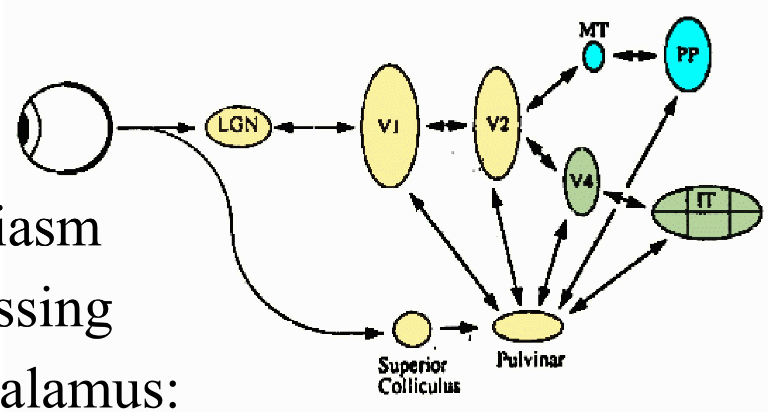

- Leaving the eyes: optic tracts, optic chiasm
- Associated pathology and signal processing
- The lateral geniculate nucleus of the thalamus: the first relay station to cortical processing
- Image processing in the LGN
- Notion of receptive field
- Primary visual cortex
- Cortical magnification
- Retinotopic mapping
- Overview of higher visual areas
- Visual processing pathways

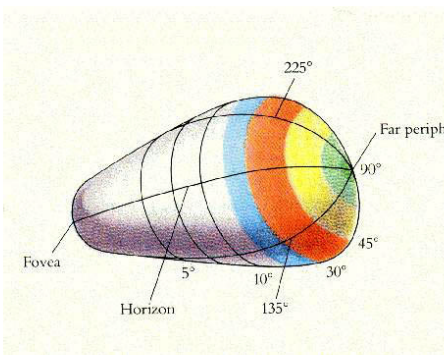

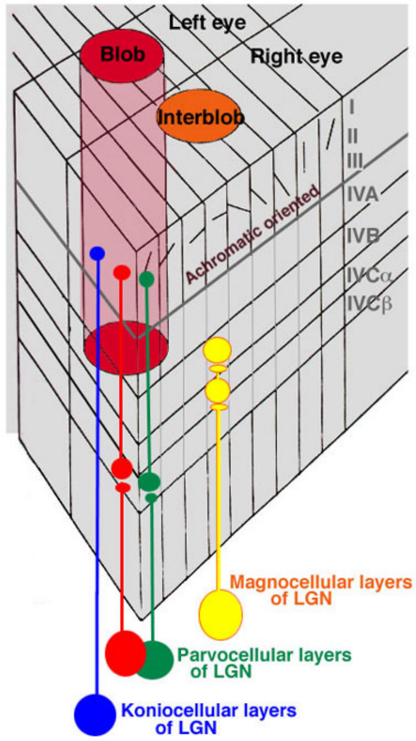

Syllabus Overview

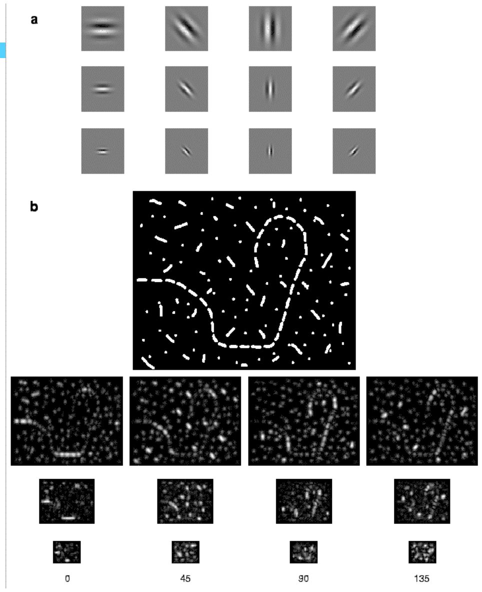

◆ ***Low-level processing and feature detection.***

- Basis transforms; wavelet transforms; jets
- Optimal coding
- Texture segregation
- Grouping
- Edges and boundaries; optimal filters for edge detection
- Random Markov fields and their relevance to biological vision
- Simple and complex cells
- Cortical gain control
- Columnar organization & short-range interactions
- Long-range horizontal connections and non-classical surround
- How can artificial vision systems benefit from these recent advances in neuroscience?

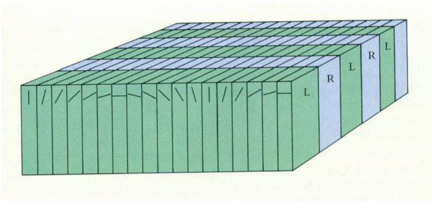

Syllabus Overview

◆ ***Coding and representation.***

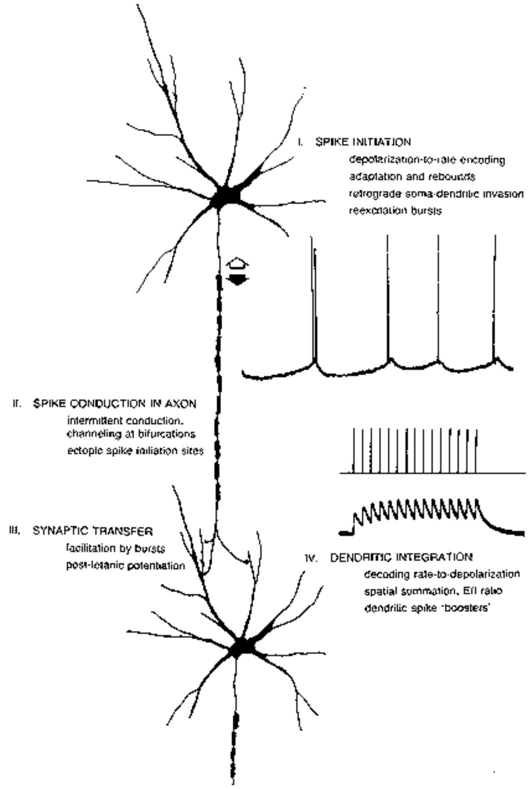

- Spiking vs. mean-rate neurons
- Spike timing analysis
- Autocorrelation and power spectrum
- Population coding; optimal readout
- Neurons as random variables
- Statistically efficient estimators
- Entropy & mutual information
- Principal component analysis (PCA)
- Independent component analysis (ICA)
- Application of these neuroscience analysis tools to engineering problems where data is inherently noisy (e.g., consumer-grade video cameras, VLSI implementations, computationally efficient approximate implementations).

Syllabus Overview

◆*Stereoscopic vision*

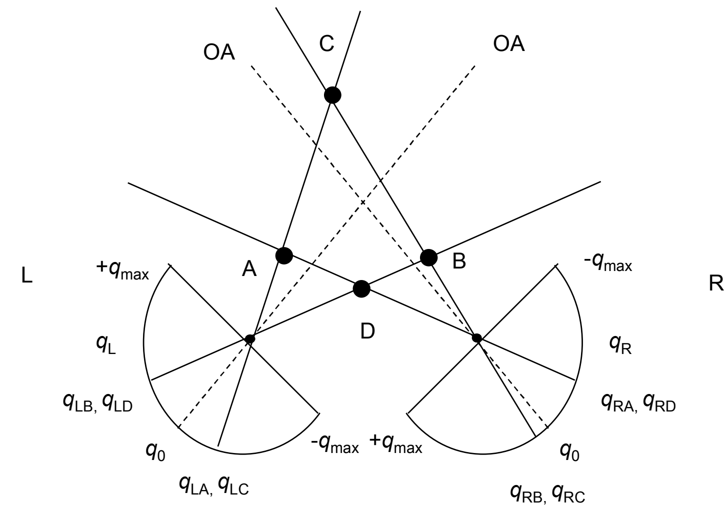

- Challenges in stereo-vision
- The Correspondence Problem
- Inferring depth from several 2D views
- Several cameras vs. one moving camera
- Brief overview of epipolar geometry and depth computation
- Neurons tuned for disparity
- Size constancy
- Do we segment objects first and then match their projections on both eyes to infer distance?
- Random-dot stereograms ("magic eye"): how do they work and what do they tell us about the brain?

Syllabus Overview

◆*Perception of motion*

- Optic flow
- Segmentation and regularization
- Efficient algorithms
- Robust algorithms
- The spatio-temporal energy model
- Computing the focus of expansion and time-to-contact
- Motion-selective neurons in cortical areas MT and MST

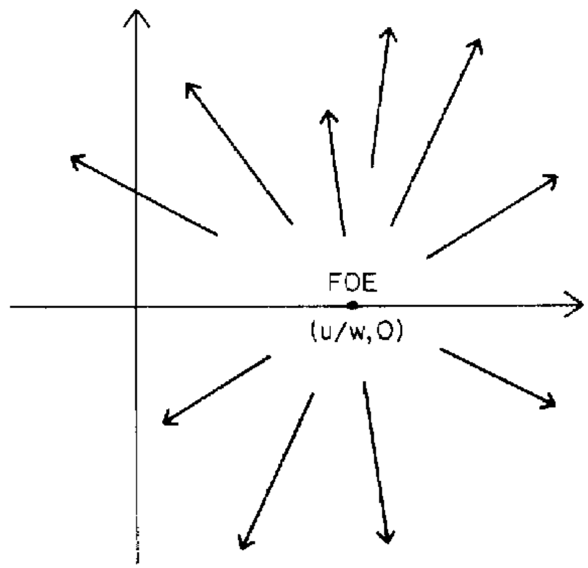

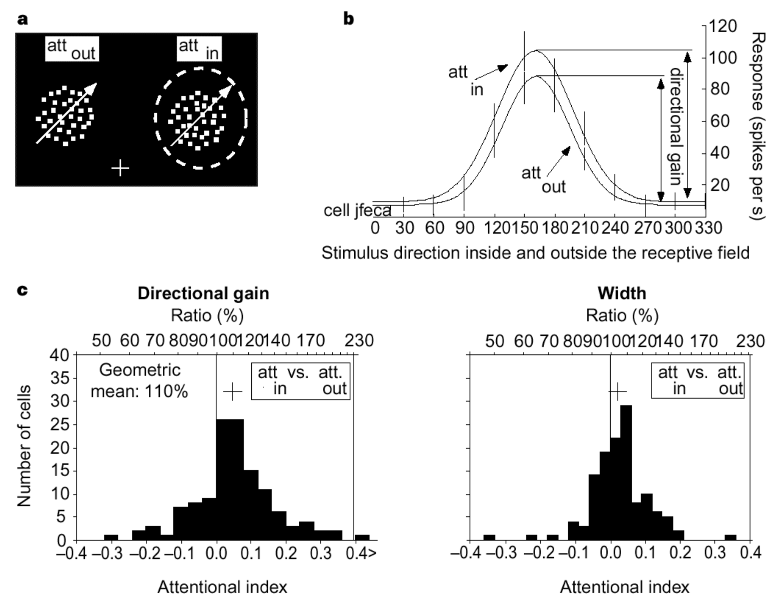

Syllabus Overview

◆ *Color perception*

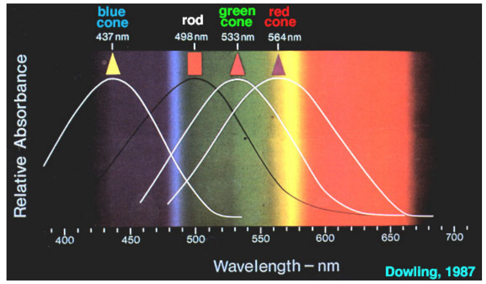

- Color-sensitive photoreceptors (cones)
- Visible wavelengths and light absorption
- The Color Constancy problem: how can we build stable percepts of colors despite variations in illumination, shadows, etc

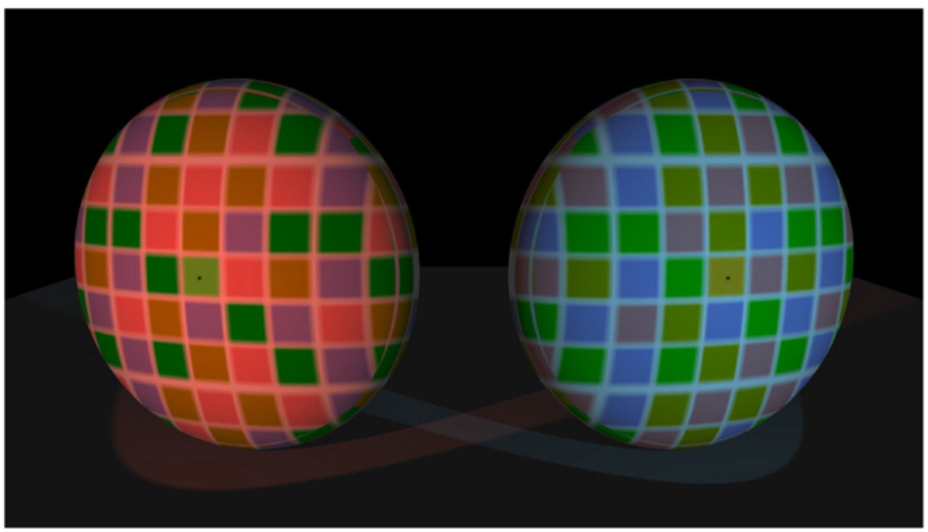

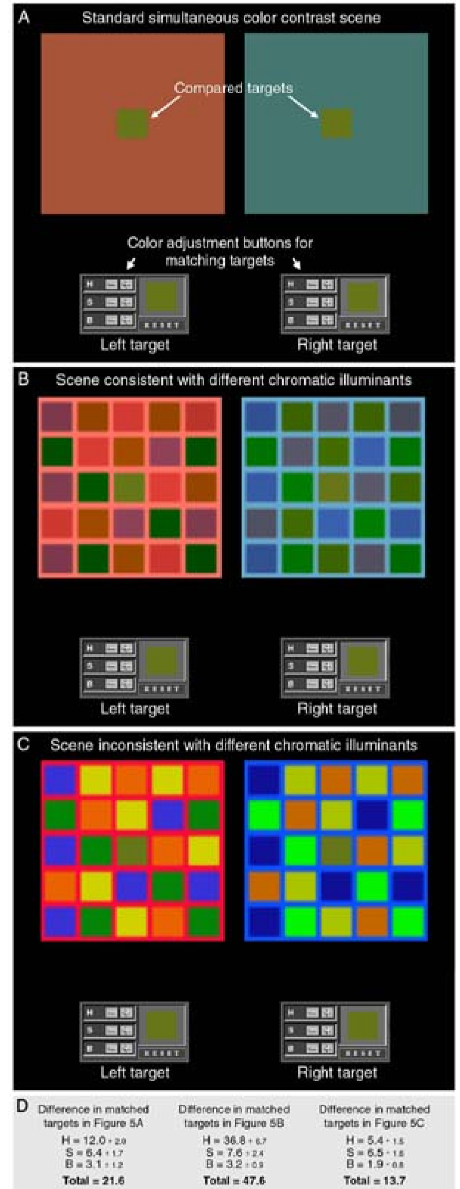

Syllabus Overview

◆ ***Visual illusions***

- What can illusions teach us about the brain?
- Examples of illusions
- Which subsystems studied so far do various illusions tell us about?
- What computational explanations can we find for many of these illusions?

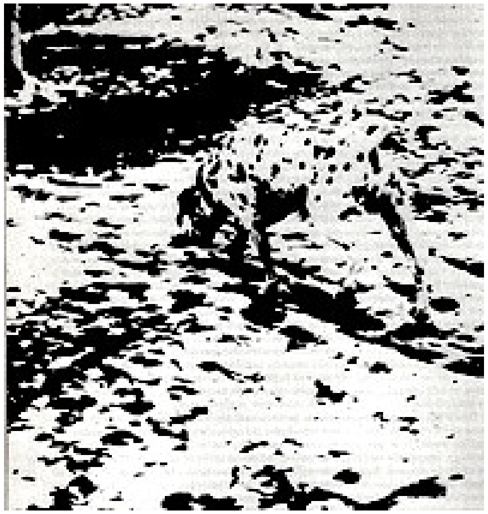

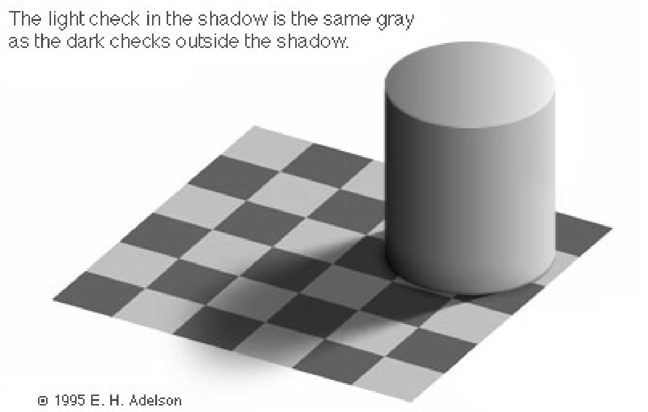

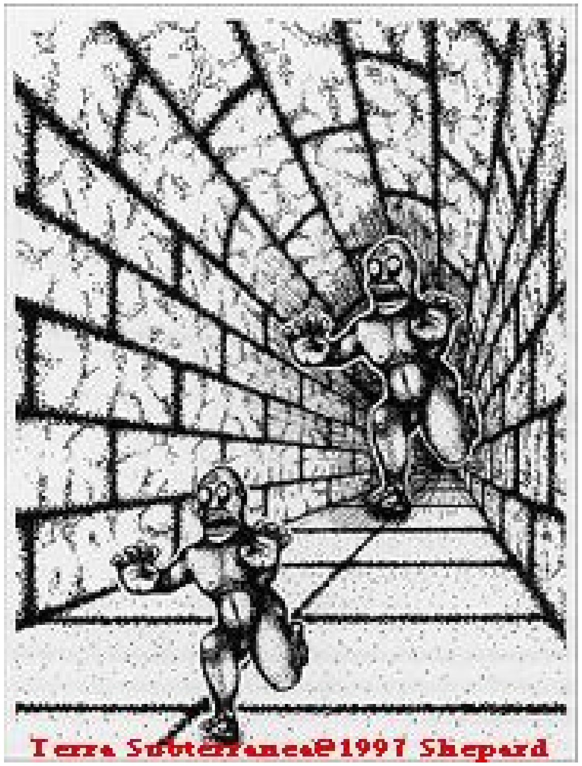

Syllabus Overview

◆ *Visual attention*

- Several kinds of attention
- Bottom-up and top-down
- Overt and covert
- Attentional modulation
- How can understanding attention contribute to computer vision systems?
- Biological models of attention
- Change blindness
- Attention and awareness
- Engineering applications of attention: image compression, target detection, evaluation of advertising, etc...

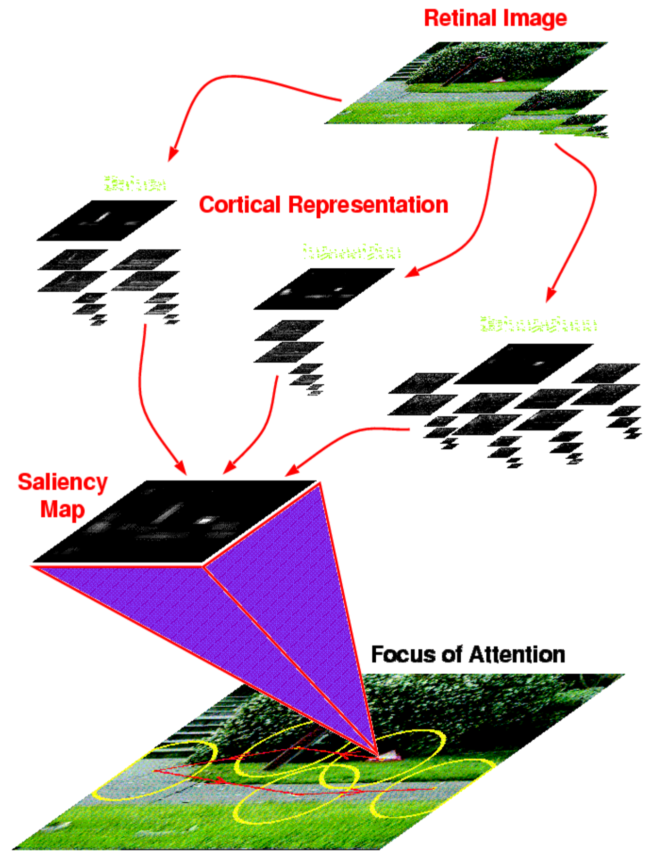

Syllabus Overview

◆*Shape perception and scene analysis*

"Face cells"

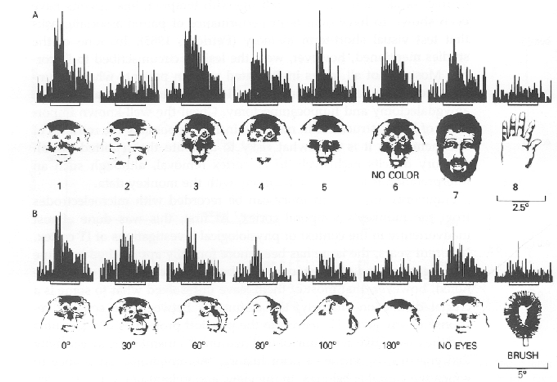

- Shape-selective neurons in cortex
- Coding: one neuron per object or population codes?
- Biologically-inspired algorithms for shape perception
- The "gist" of a scene: how can we get it in 100ms or less?
- Visual memory: how much do we remember of what we have seen?
- The world as an outside memory and our eyes as a lookup tool

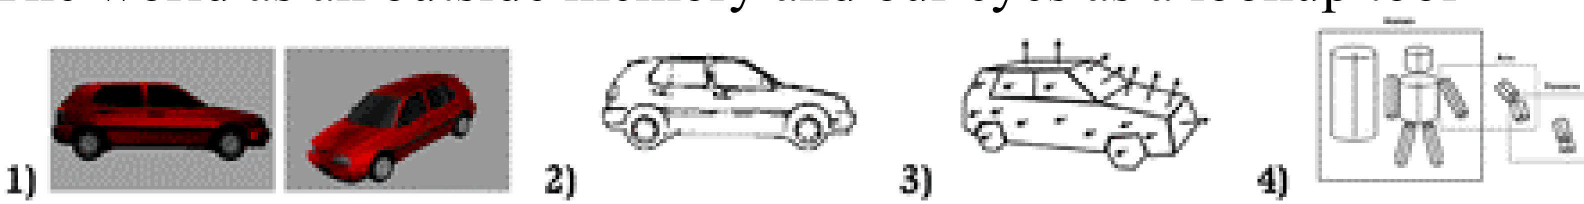

Syllabus Overview

◆ *Object recognition*

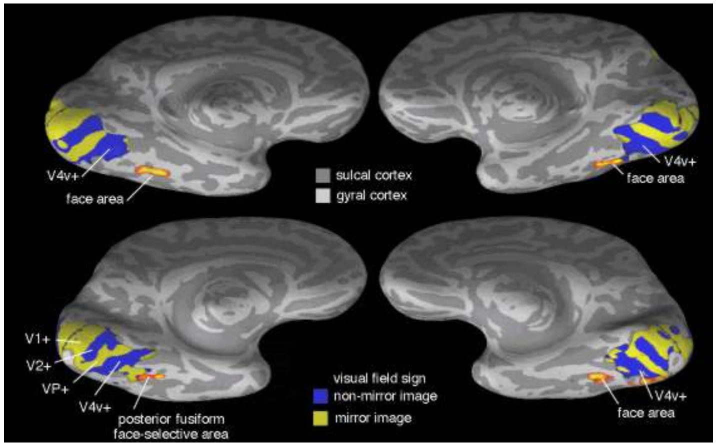

- The basic issues
- Translation and rotation invariance
- Neural models that do it
- 3D viewpoint invariance (data and models)
- Classical computer vision approaches: template matching and matched filters; wavelet transforms; correlation; etc.
- Examples: face recognition.
- More examples of biologically-inspired object recognition systems which work remarkably well

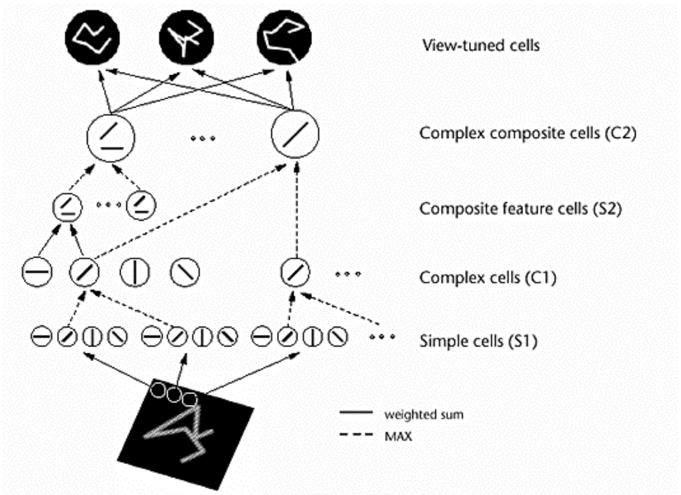

Syllabus Overview

◆ *Computer graphics, virtual reality and robotics*

- Exploiting the limitations of the human visual system when generating computer animations
- Linking vision systems to robots
- Visuo-motor interaction
- Real-time implementations
- Parallel implementations
- Towards conscious machines
- Link to artificial intelligence

Next time

We will have an introduction to the brain and neurons.

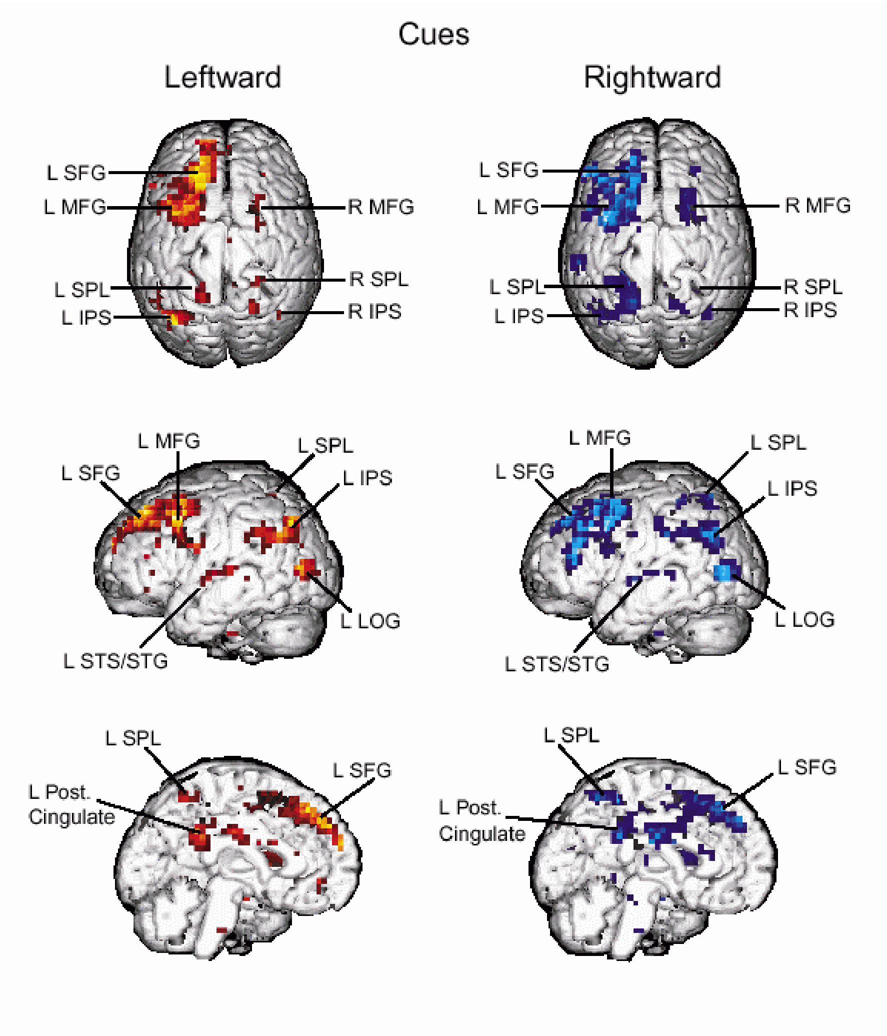

of the cues. Additionally, the cues produced bilateral activations in the region of the insula near the putamen.

In contrast, the target stimuli evoked neural activity in brain areas that were largely distinct from those activated by the attention-directing cues (**Fig. 3; Table 2**). Bilateral activations in the supplementary motor area (SMA) extending into regions of the midcingulate gyrus and activations surrounding the central sulcus were found in response to targets whether attention was directed to the right or left. Ventrolateral prefrontal areas were also bilaterally activated by the targets. None of these regions were activated by the cues. Only the superior parietal lobule was found to

Fig. 2. Activity related to attentional control. Data for brain regions significantly activated in response to the cue stimuli were overlaid onto a brain rendered in 3D. Left column, activations to cues instructing subjects to orient attention to the left visual field location. Right column, activations to cues instructing subjects to attend the right visual field location. Top panels, dorsal view of the brain (frontal pole at top); middle and bottom panels, lateral and medial views, respectively, of the left hemisphere. Labels indicate the brain regions referred to in Table 1. The *Z*-values and stereotactic coordinates for the regional maxima are given in Table 1. SFG, superior frontal gyrus; MFG, middle frontal gyrus; SPL, superior parietal lobule; IPS, intraparietal sulcus; STS, superior temporal sulcus; STG, superior temporal gyrus; LOG, lateral occipital gyrus.

Hopfinger et al, 1999