

Vision: From Computational Theory
to
Neuronal Mechanism

January 8, 2008

Prof. Christof Koch

Prof. Pietro Perona

Prof. Shin Simojo

Class Requirements

- Attend all 19 lectures
- Hand in 6 homeworks on time
- Participate in two project presentations
- Carry out a research project and write a 4-8 pages report

Grading

- This is a 4-4-4 unit class
- Final letter grade is based on homework (6x8 points), project presentation (2x3 points), and written project report (46 points)
- Eligible students can take the class on a P/F basis (Pass > 50%)
- Grades are not normalized (we do not grade on a curve)

90% < A < 100%

80% < B < 90%

65% < C < 80%

50% < D < 65%

Fail for < 50%

(with appropriate +/- gradations)

Class Minutiae

- Any conflicts?
- We do have a strict late policy for homeworks.
- TA: Jonathan Harel (jonharel@gmail.com) and Jeffrey Elund (jedlund@its.caltech.edu)
- Class web-site www.klab.caltech.edu/cns186
- www.cns.caltech.edu/wiki/index.php/CNS186
- Email list cns186@klab.caltech.edu

Project

- A small project in vision (analytical work, computer simulation, your own psychophysical or fMRI experiment) that you can complete in 2-3 weeks. Often the starting points for thesis research
- Don't be too ambitious
- Two project presentations
 - 1/30 (Wednesday) 2 slides
 - 2/21 (Thursday) 2 slides
- Hand in 4-8 written report on 3/18 (Tuesday)

Literature

- The basic text is *Vision Science* by Stephen Palmer (1999, MIT Press)
- Other books

CNS186 Vision

Introduction to Computational Vision

Prof. Christof Koch


Why see?



Why see?



What and where



What:

object recognition
surface properties
classification of actions

Where:

spatial relationships
shape
distance
ego-motion

Cues

- Texture
- Edges
- Stereoscopic disparity
- Aerial perspective
- Occlusion
-

Occlusion,
texture
gradient




Edges

(static motion)







Texture
gradient





Images
can be
ambiguous





What is in an image?



What is in an image?



What is in an image?



Vision as a computational problem

- Vision is the *process* of discovering from images what is present in the world, and where it is.

Computational theory

Goal of the computation? Why is it appropriate and what is the logic of the strategy by which it can be carried out?

Representation and algorithm

Implementation?
Representation of the input and output?
Algorithm for the transformation?

Hardware implementation

How can representation and algorithm be realized physically?

From David Marr's *Vision* (1980)

Hierarchies of vision problems

Early vision

mid-level vision

High-level vision

Vision processes



Image(s)

The vision
black box

