6.801/6.866 Machine Vision

Fall 2002

Tuesday and Thursday from 2:30-4:00 in 2-105

Announcements

- The first class will be held on Thursday, September 5th.

Course Information

- Syllabus
- Grading and Requirements
- Internet Resources

Contacts

All offices are on the sixth floor of 400 Technology Square (Forrester building).

Instructors	Professor Bill Freeman wtf@ai.mit.edu NE43-V601 (617) 253-8828	Professor Trevor Darrell trevor@ai.mit.edu NE43-V605 (617) 253-8966
Office Hours	TBD	TBD
Teaching Assistants	Erik Sudderth esuddert@mit.edu NE43-V609 (617) 253-3497	Louis-Philippe Morency lmorency@ai.mit.edu NE43-V603 (617) 253-4278
Office Hours	TBD	TBD

http://www.ai.mit.edu/courses/6.801/

Map showing 400 Technology Square

The building says "Forrester" on the side. Only the parking garage side building entrance is unlocked. (After normal business hours, the elevator to our floor and the building itself are both locked.) Exiting the elevator on the 6th floor, you'll see a pair of glass doors on one side. Enter the left glass door, then turn right at every opportunity to find my office, room 601.

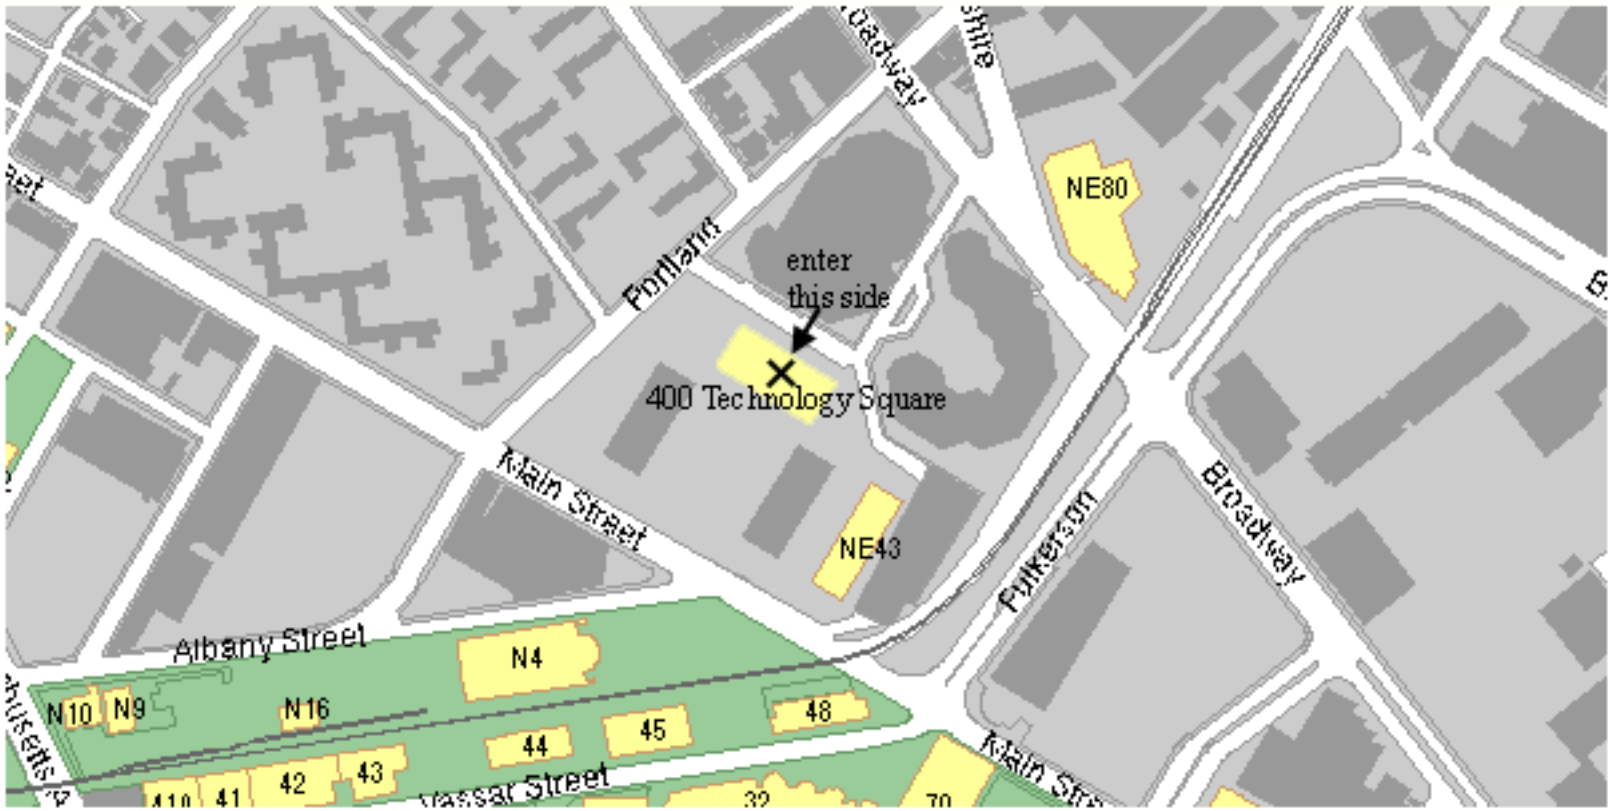

back to my home page , Sept., 2002.

6.801/6.866 Machine Vision

Course Requirements

Prerequisites

- Signals and Systems (6.003)
- Familiarity with linear algebra (18.06) and probability (6.041)
- Some Matlab experience

Textbooks

Both texts are available at Quantum Books:

- D. Forsyth and J. Ponce, *Computer Vision: A Modern Approach*, Prentice Hall, 2002.
- B. Horn, *Robot Vision*, MIT Press, 1986. (errata)

Grading

- Two take-home exams
- Four problem sets with lab exercises in Matlab
- No final exam
- Final project for graduate students (6.866)

Problem sets may be discussed, but all written work and coding must be done individually. Take-home exams may not be discussed. Individuals found submitting duplicate or substantially similar materials due to inappropriate collaboration may get an F in this class and other sanctions.

	Take-Home Exams (2)	**Problem Sets (4)**	**Final Project**
6.801	60%	40%	N/A
6.866	40%	30%	30%

The final project may be

- An original implementation of a new or published idea
- A detailed empirical evaluation of an existing implementation of one or more methods
- A paper comparing three or more papers not covered in class, or surveying recent literature in a particular area

A project proposal not longer than two pages must be submitted and approved before the end of October.

6.801/6.866 Machine Vision

Internet Resources

Matlab Links

Mathworks Online Documentation
: Online versions of the official Matlab documentation, including a general command reference and information about the image processing toolbox.

Matlab on Athena
: Guide to using Matlab at MIT. Athena consulting's stock answers may also be useful.

University of Colorado Matlab Tutorials
: A decent collection of Matlab tutorials, including one focusing on image processing.

Matlab Image Processing Tutorial
: A short introduction to the manipulation of images in Matlab, including an introduction to principal components analysis via eigenfaces.

Matlab Educational Sites
: An index of various Matlab tutorial websites.

Computer Vision Homepage

6.801/6.866 Machine Vision

Syllabus

9/5	Course Introduction	Pset #0 (not collected)
9/10	Cameras, Lenses, and Sensors	
9/12	Radiometry and Shading Models	Pset #1 Assigned
9/17	Color	
9/19	Epipolar Geometry	
9/24	Stereo	Pset #1 Due
9/26	Shape from Shading	Pset #2 Assigned
10/1	Filtering I	
10/3	Filtering II	
10/8	Intro to Bayesian Vision	Pset #2 Due
10/10	Edge Detection and Optic Flow	Exam #1 Assigned
10/15	Columbus Vacation (NO LECTURE)	
10/17	Direct SFM	
10/22	Models of Texture	Exam #1 Due
10/24	Statistical Classifiers I	Pset #3 Assigned
10/29	Statistical Classifiers II	
10/31	Multi-view Geometry	
11/5	Affine SFM	Pset #3 Due
11/7	Projective Geometry vs. Perspective Projection	PSet #4 Assigned
11/12	Clustering & Segmentation	

6.801/6.866 Machine Vision

Syllabus

11/12	Clustering & Segmentation	
11/14	EM	
11/19	Hough Transforms, RANSAC, & Voting Methods	Pset #4 Due
11/21	Tracking & Density Propagation	Exam #2 Assigned
11/26	Model-Based Vision	
11/28	Thanksgiving (NO LECTURE)	
12/3	Visual Surveillance / Activity Monitoring	
12/5	Image Databases	Exam #2 Due
12/10	Image-Based Rendering	Projects Due

Vision

- What does it mean, to see? "to know what is where by looking".
- How to discover from images what is present in the world, where things are, what actions are taking place.

from Marr, 1982

Computer vision class, fast-forward

Images and image formation

Cameras, lenses, and sensors

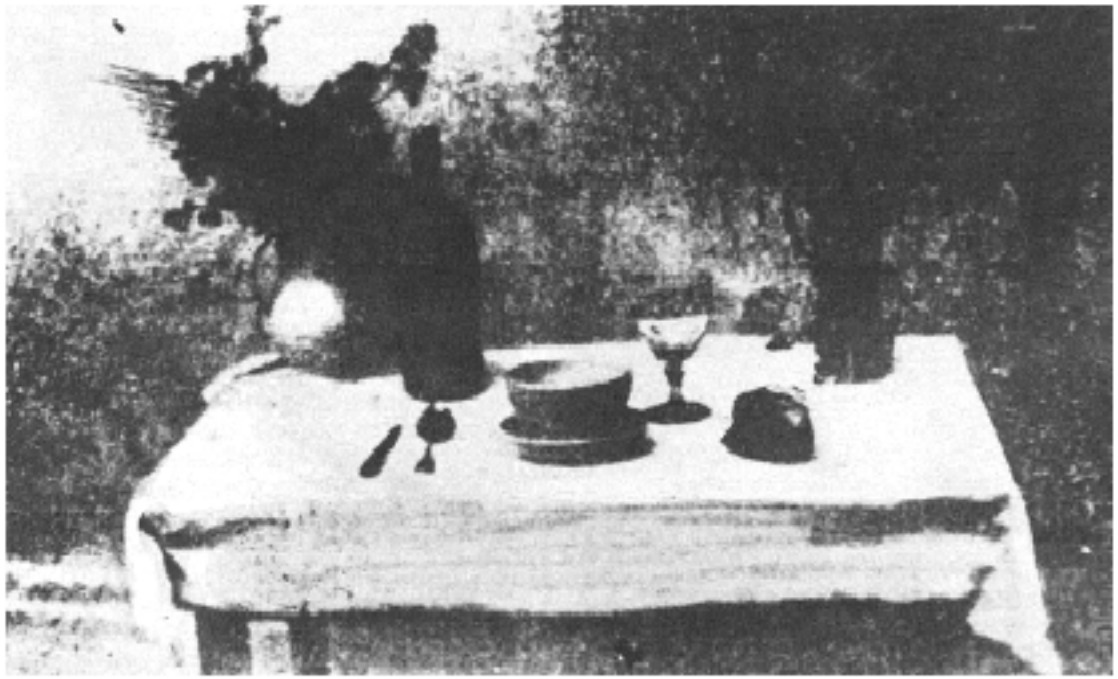

- Pinhole cameras
- Lenses
- Projection models
- Geometric camera parameters

Figure 1.16 The first photograph on record, *la table servie*, obtained by Nicéphore Niepce in 1822. *Collection Harlinge–Viollet.*

From Computer Vision, Forsyth and Ponce, Prentice-Hall, 2002.

Radiometry

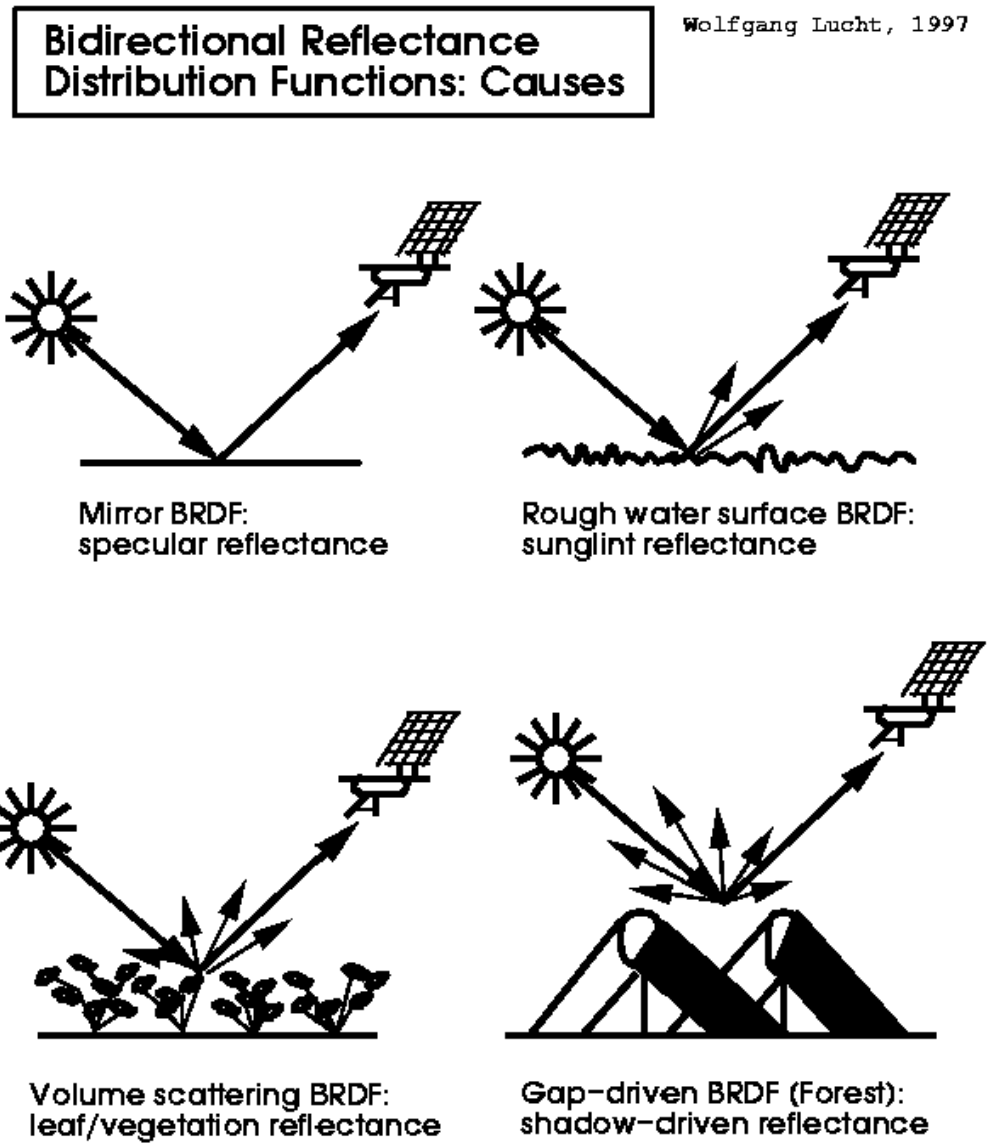

Wolfgang Lucht

http://geography.bu.edu/brdf/brdfexpl.html

Color

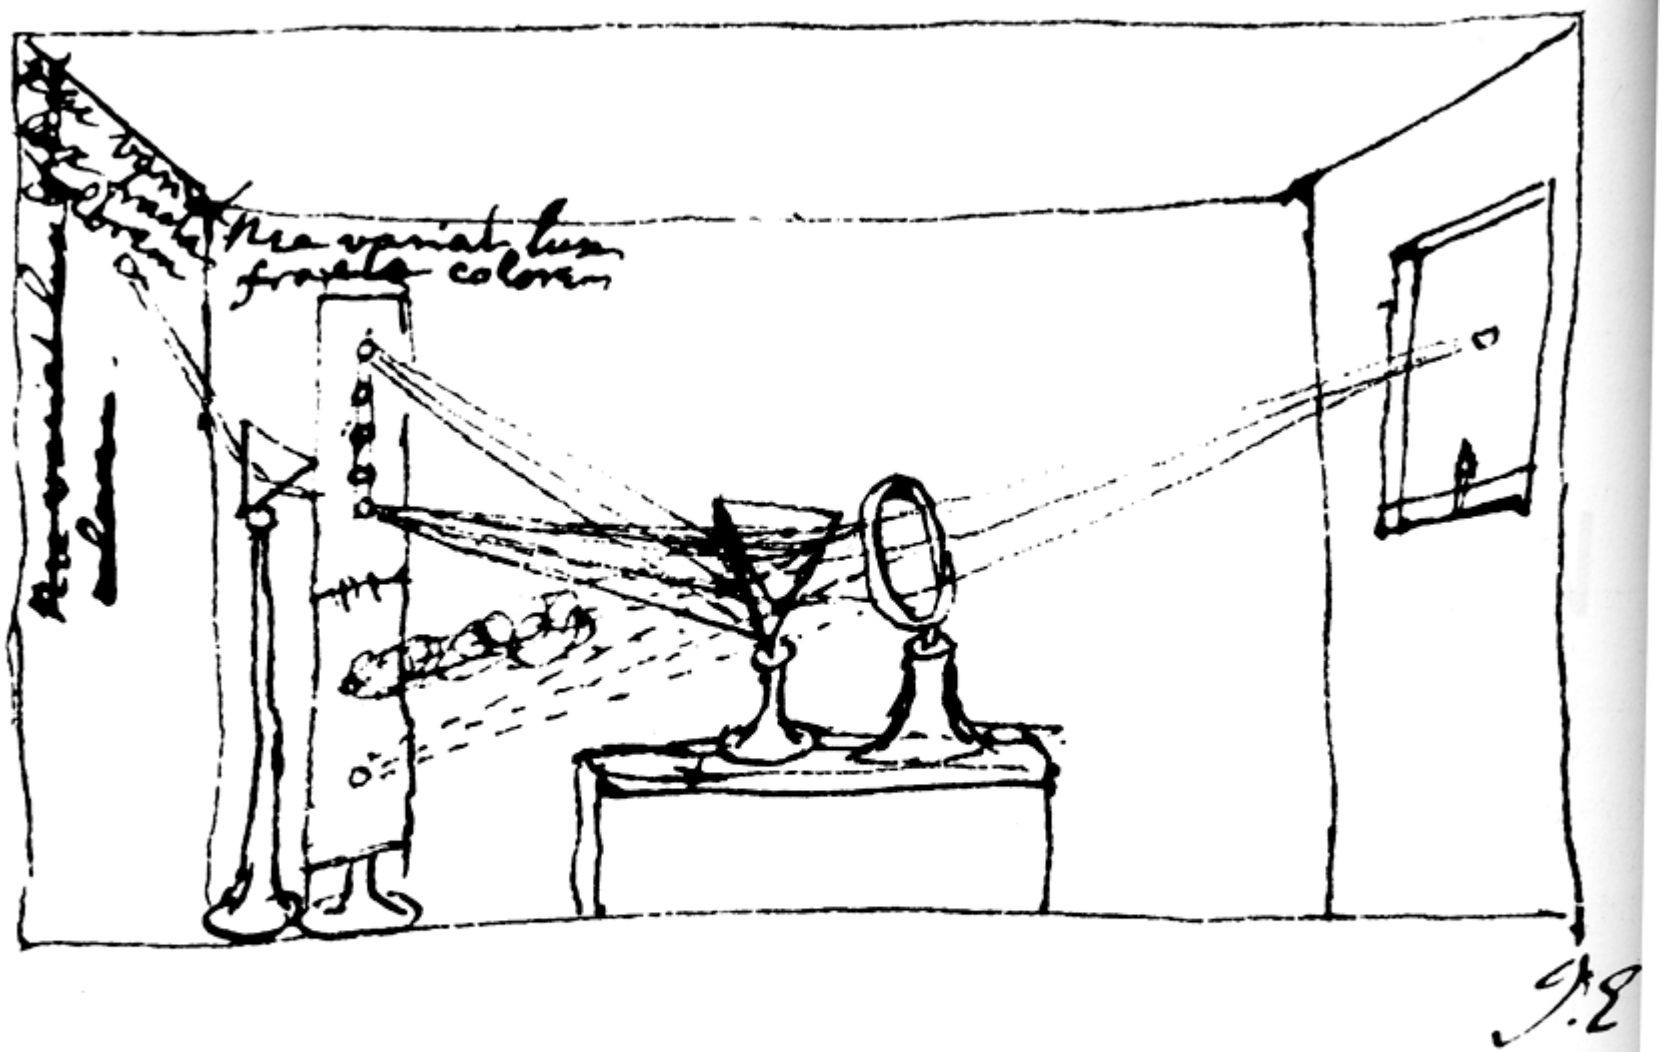

4.1 NEWTON'S SUMMARY DRAWING of his experiments with light. Using a point source of light and a prism, Newton separated sunlight into its fundamental components. By reconverging the rays, he also showed that the decomposition is reversible.

From Foundations of Vision, by Brian Wandell, Sinauer Assoc., 1995

Low-level vision

Image filtering

- Review of linear systems, convolution
- Bandpass filter-based image representations
- Probabilistic models for images

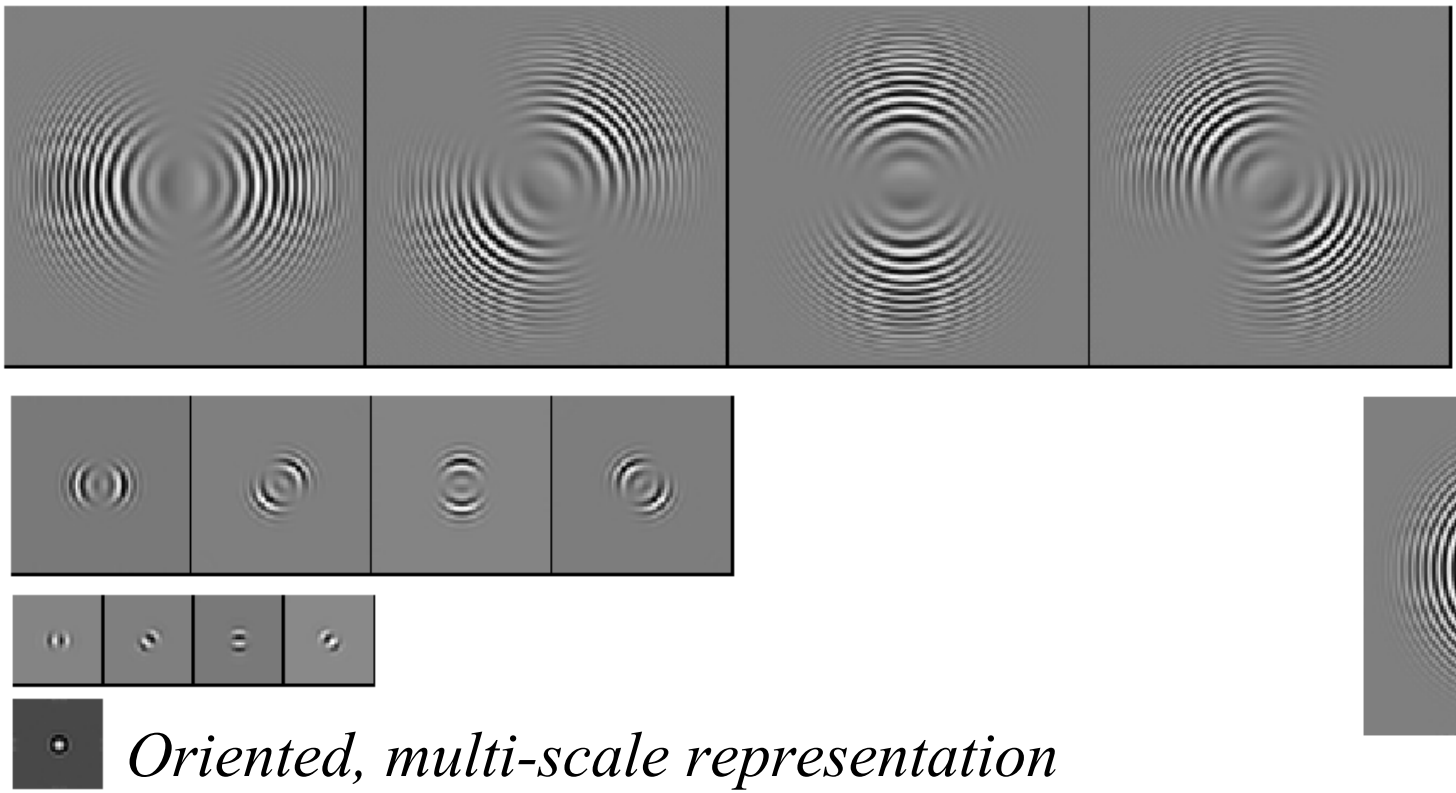

Oriented, multi-scale representation

Image

Non-linear filtering, and applications

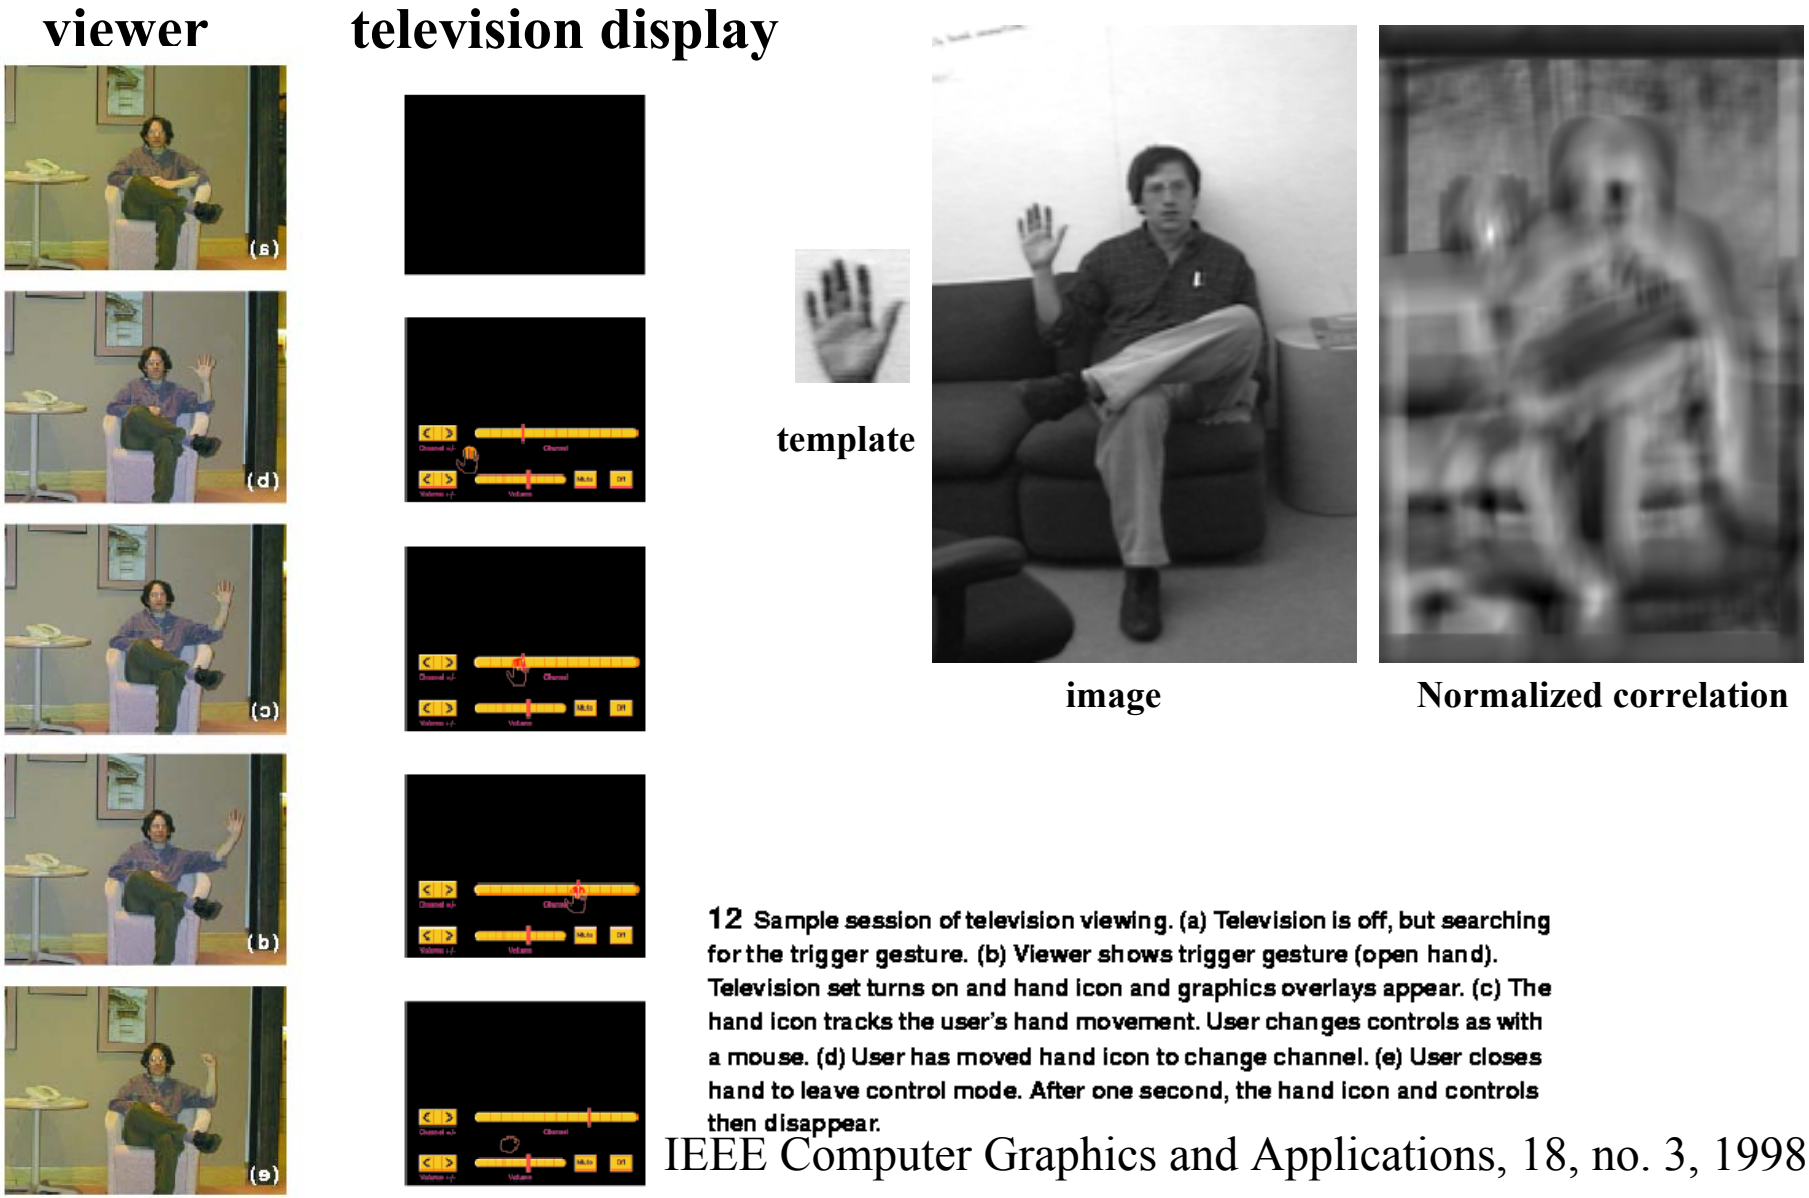

12 Sample session of television viewing. (a) Television is off, but searching for the trigger gesture. (b) Viewer shows trigger gesture (open hand). Television set turns on and hand icon and graphics overlays appear. (c) The hand icon tracks the user's hand movement. User changes controls as with a mouse. (d) User has moved hand icon to change channel. (e) User closes hand to leave control mode. After one second, the hand icon and controls then disappear.

IEEE Computer Graphics and Applications, 18, no. 3, 1998

Models of texture

Parametric model

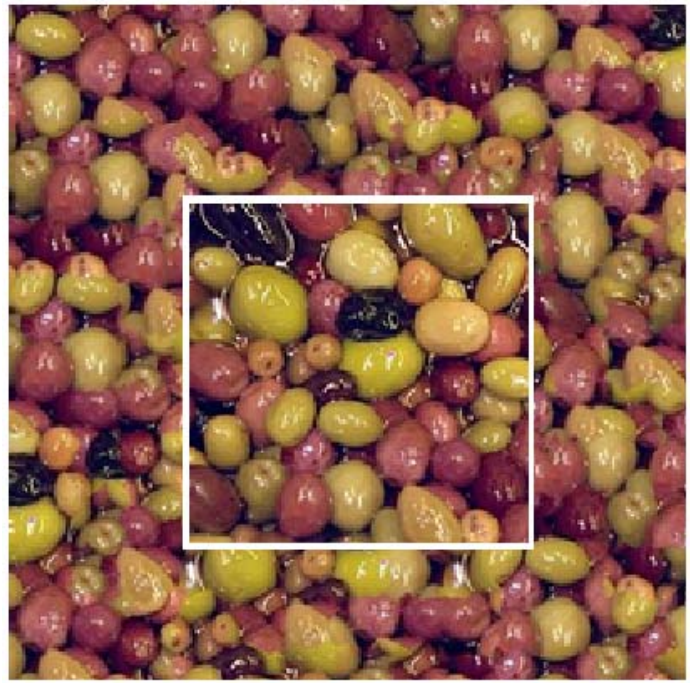

Non-parametric model

A Parametric Texture Model based on Joint Statistics of Complex Wavelet Coefficients
J. Portilla and E. Simoncelli, *International Journal of Computer Vision 40(1): 49-71, October 2000.*
© Kluwer Academic Publishers.

A. Efros and W. T Freeman, **Image quilting for texture synthesis and transfer**, SIGGRAPH 2001

Shape from shading

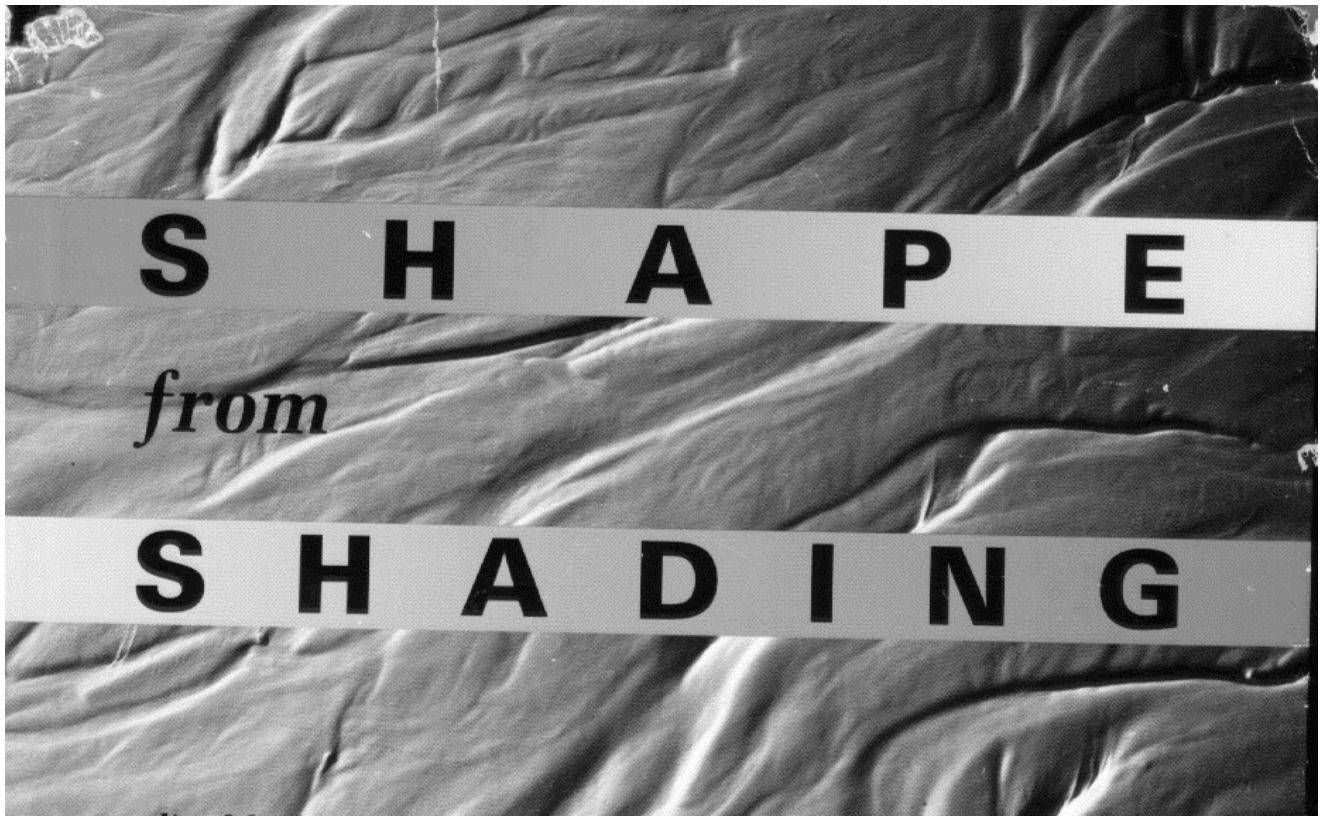

Shape from shading, Horn and Brooks, MIT Press, 1986

Bayesian framework for vision

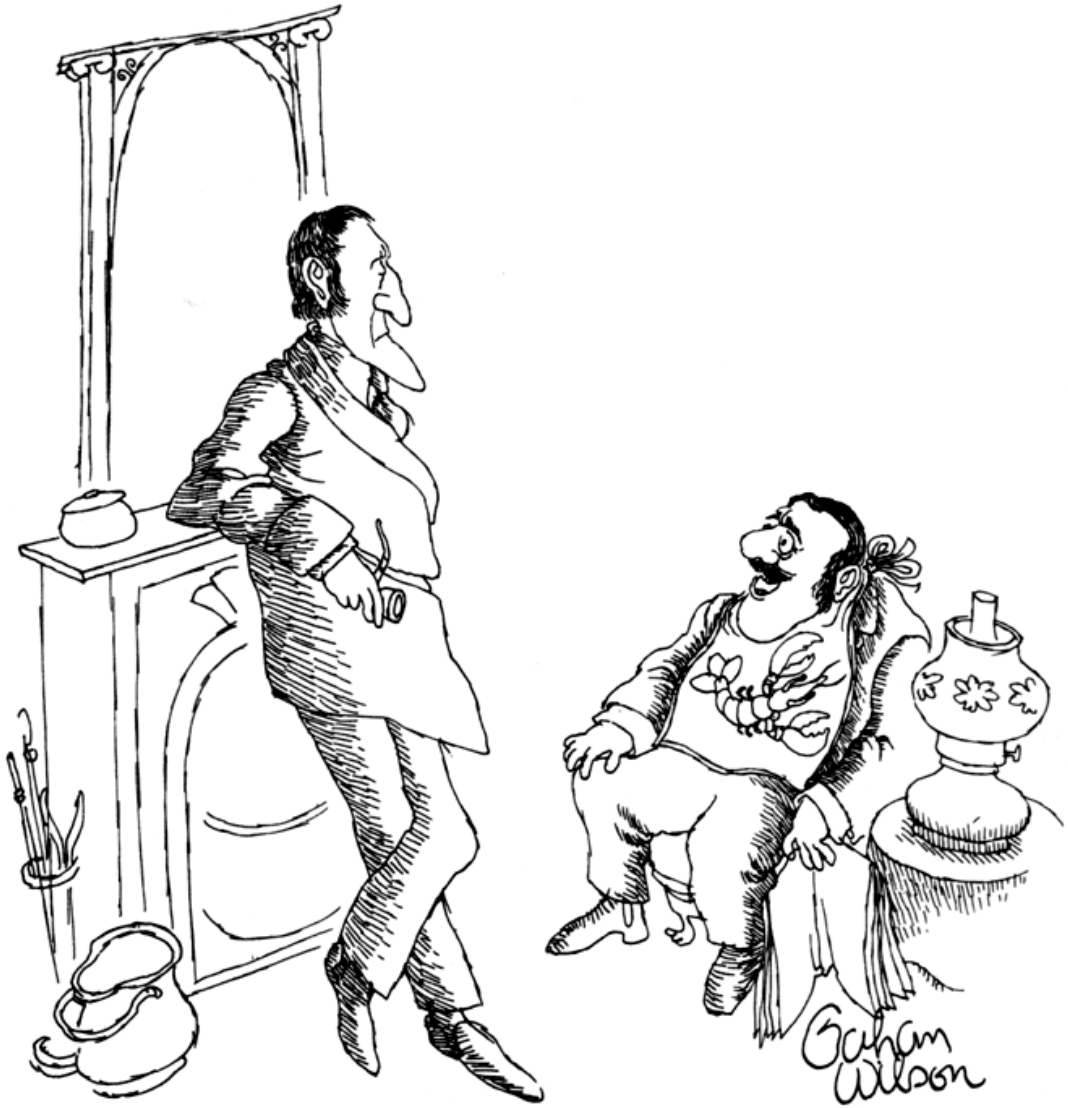

"Good lord, Holmes! How did you come to know
I'd seafood for lunch?"

Gahan Wilson's Still Weird, Forge, 1994

Bayesian framework for vision

Coincidental appearance of face profile in rock?

http://www.cs.dartmouth.edu/whites/old_man.html

Bayesian framework for vision

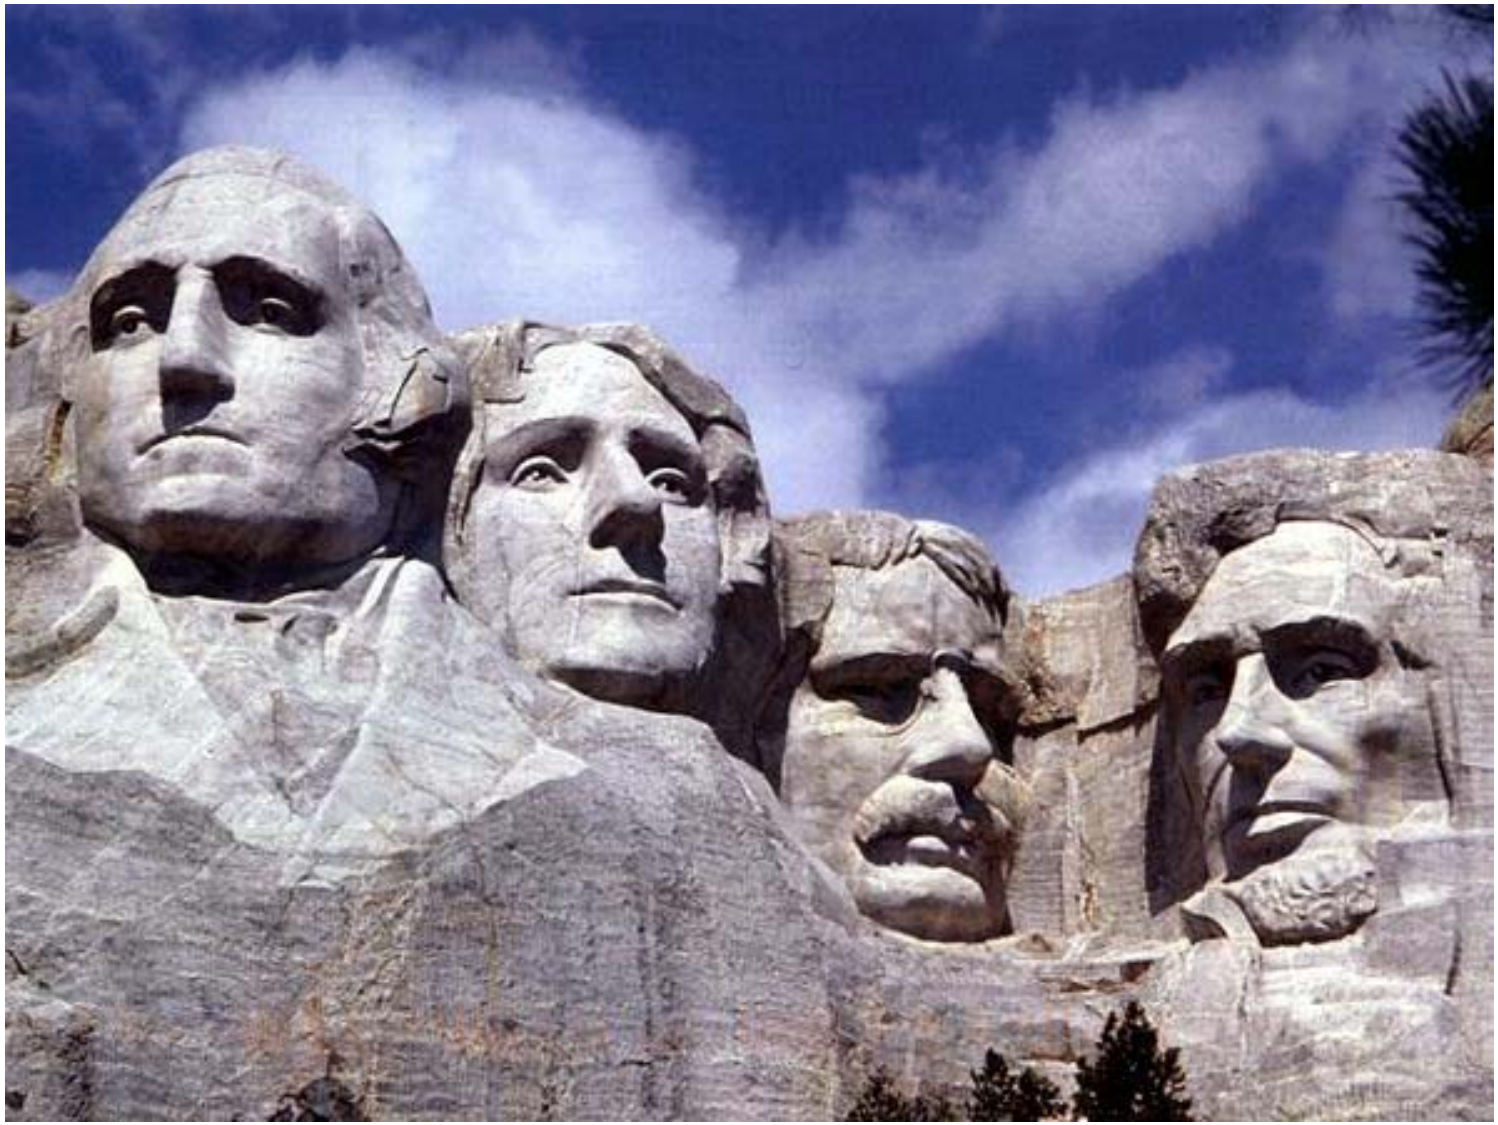

Coincidental appearance of faces in rock?

http://bensguide.gpo.gov/3-5/symbols/print/mountrushmore.html

Edge detection and optical flow

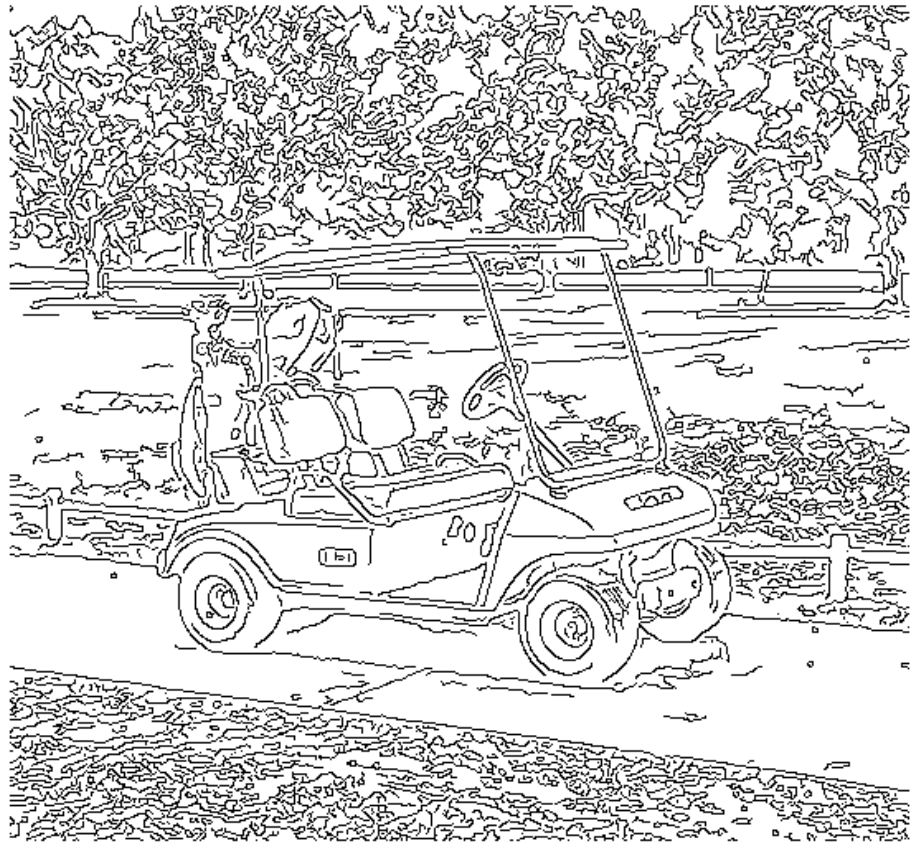

M. Heath, S. Sarkar, T. Sanocki, and K.W. Bowyer, "A Robust Visual Method for Assessing the Relative Performance of Edge-Detection Algorithms" IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 19, No. 12, December 1997, pp. 1338-1359.

Edge detection and optical flow

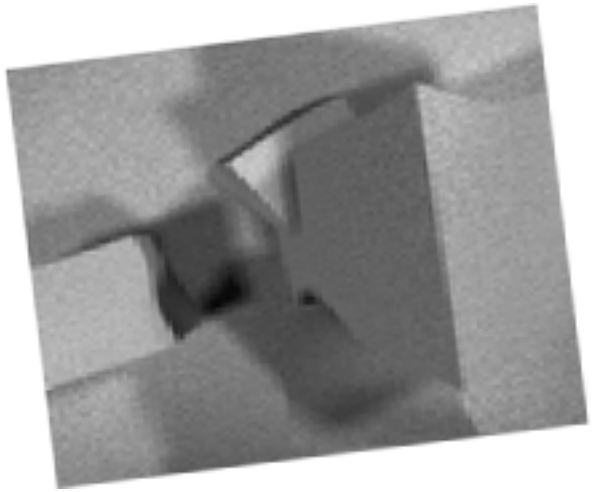

Images

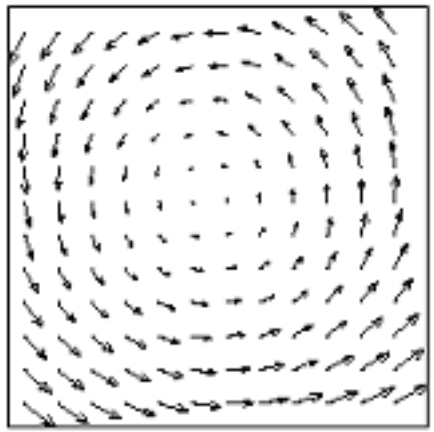

True optical flow

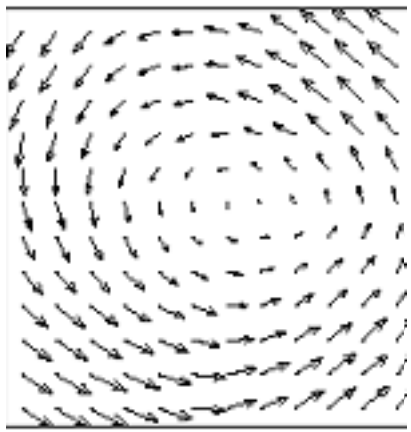

Estimated optical flow

in "Statistical Theories of the Brain", edited by R. Rao, B. Olshausen, and M. Lewicki, MIT Press, 2001

Learning and vision

Statistical classifiers

- MIT Media Lab face localization results.

- Applications: database search, human machine interaction, video conferencing.

Support vector machines and boosting

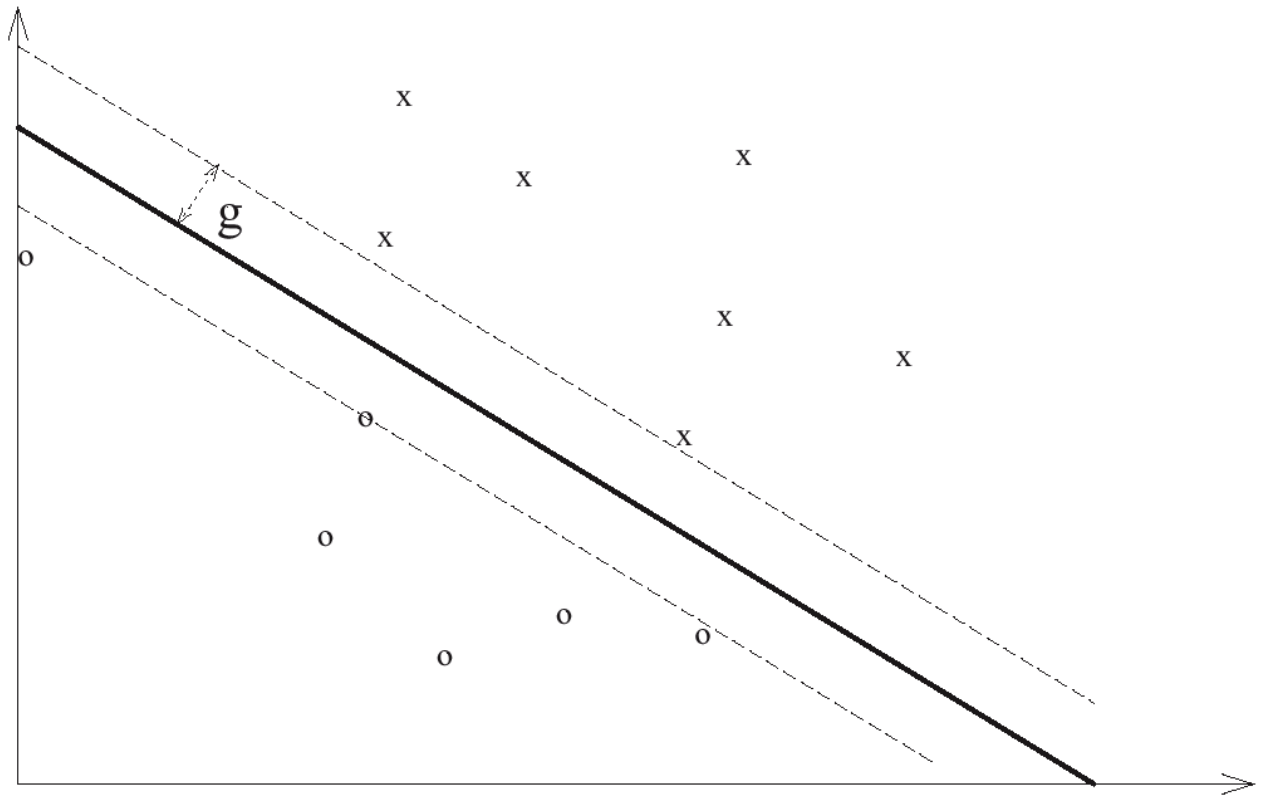

Large-margin classifier

www.support-vector.net/nello.html

Support vector machines and boosting

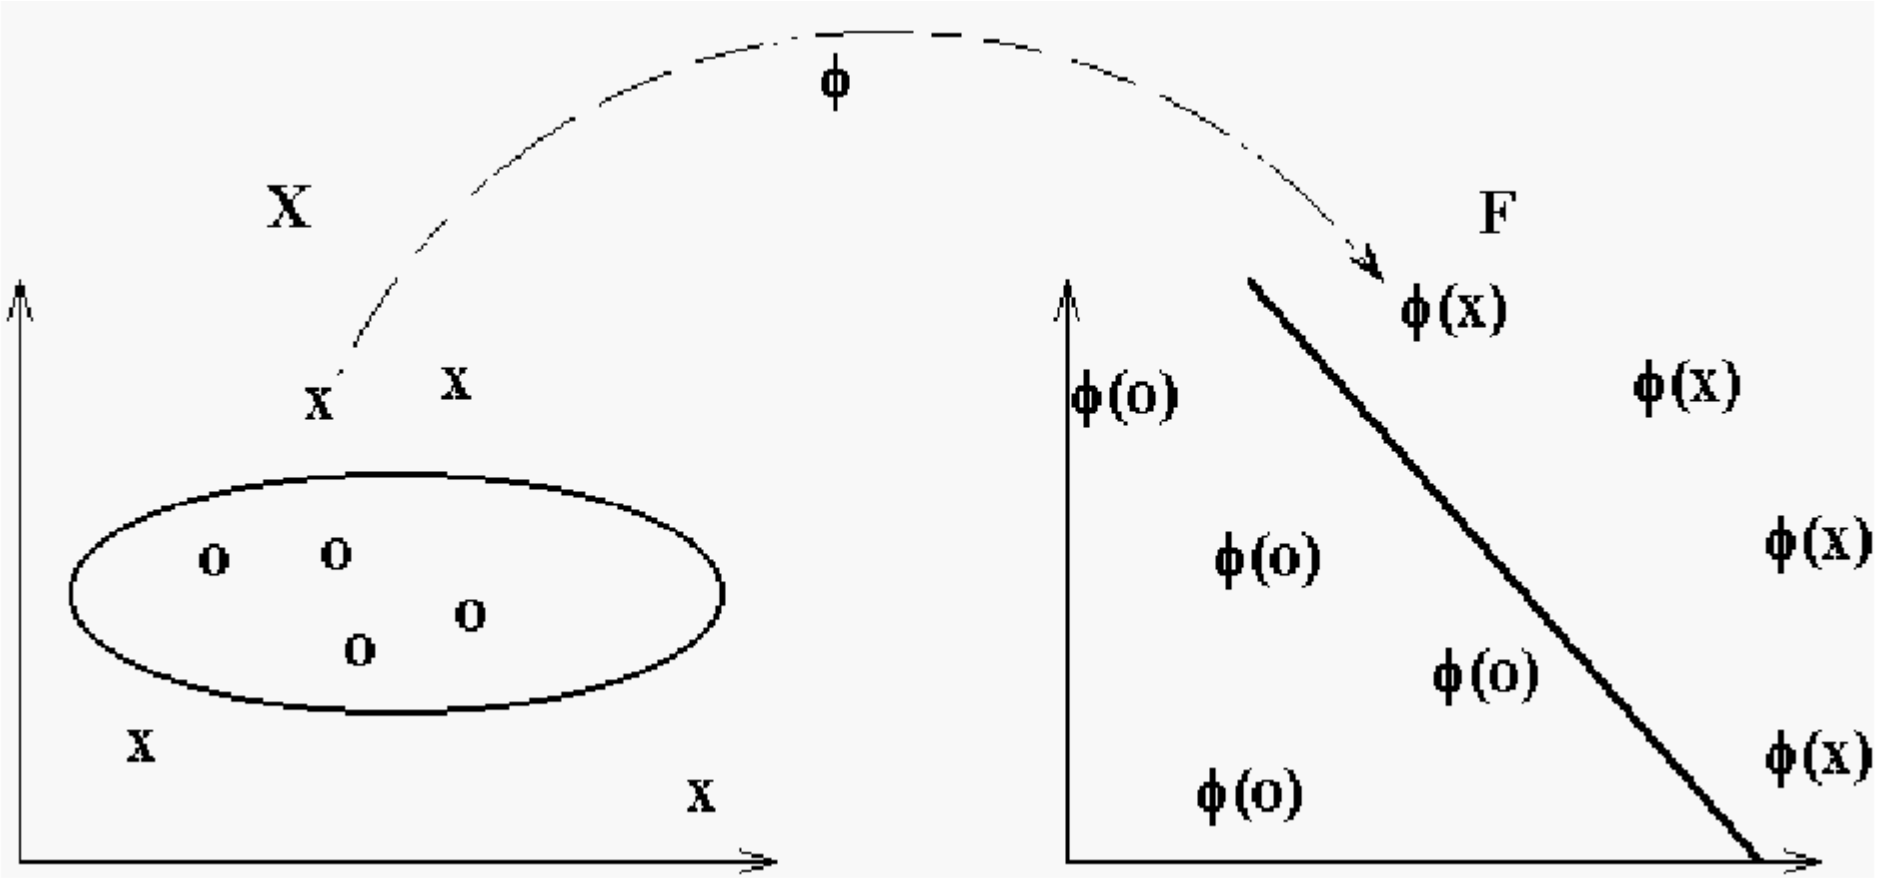

"The kernel trick"

www.support-vector.net/nello.html

Applications

- Computer vision for computer games

Computer vision applications as ocean-going vessels

this application ←

Game: Decathlete

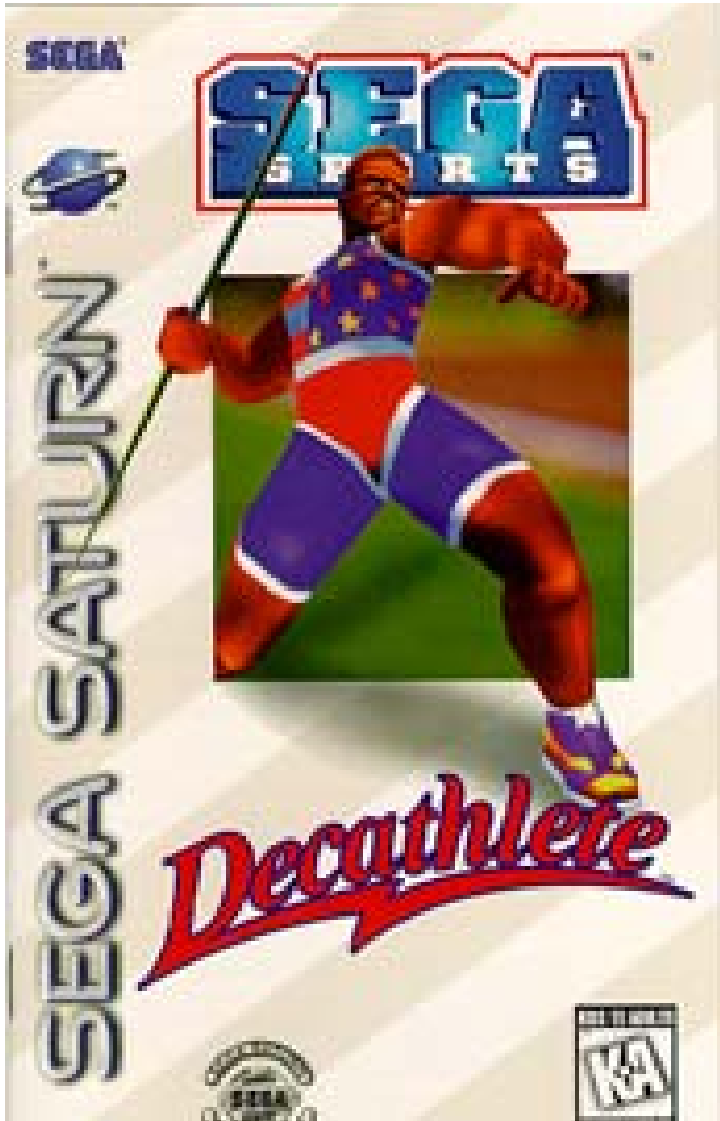

Optical-flow-based Decathlete figure motion analysis

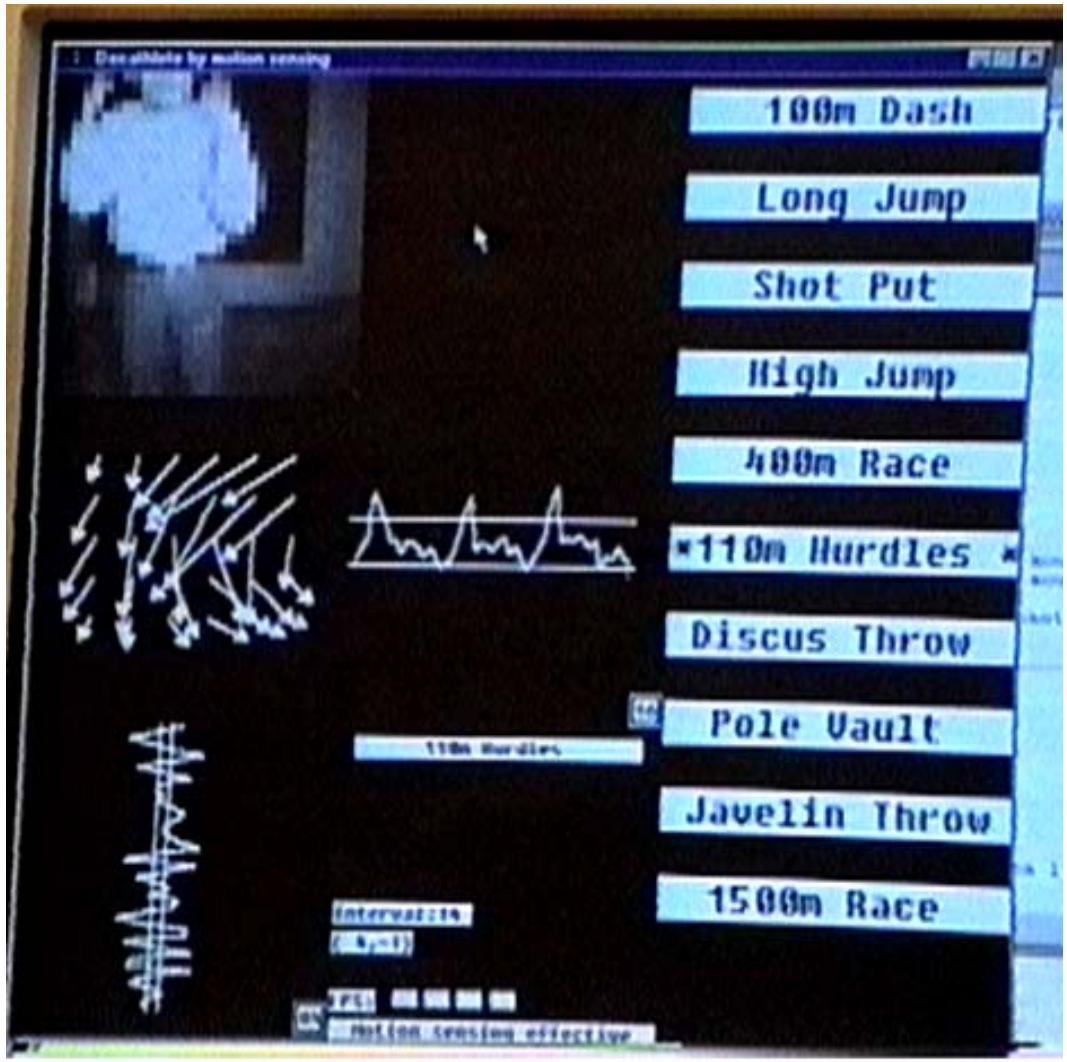

Decathlete 100m hurdles

Decathlete javelin throw

Decathlete javelin throw

video

Nintendo Game Boy

- Several million sold (most of any digital camera). Imaging chip is Mitsubishi Electric's "Artificial Retina" CMOS detector.

video

The course, in broad categories

- Images and image formation
- Geometry
- Low-level vision
- High-level vision
- Implementations and applications