Introduction to Computer Vision

Michael J. Black

Sept 2009

Lecture 3: Introduction

Tutorials

Matlab tutorial:

Sunlab (1st 3 rows)

Monday and Tuesday 7-8 pm

Linear Algebra tutorial:

CIT 219 Wednesday and Thursday 7-8 pm

Office hours

Michael: CIT 521

Mondays 4pm

Thursdays 3pm

Peng, CIT 271: Monday 7-9 pm

Tim: Wednesday 4-6 pm

Assignment 1

- Parts 1 (pyramids) & 2 (edges and derivative filters) out of 4 parts out today
- 1&2 Due Wed Sept 23.

For next class

Reading: Ch 3.2.1 Linear Filtering – Wednesday

Background: 2.3.1 (sampling and aliasing), 3.3 intro (Fourier transform)

Reading ahead: 3.4.1, 3.4.2 (interpolation & pyramids)

Applied Math seminar

http://www.dam.brown.edu/ptg/seminar.html

- Wednesdays at noon, 182 George Street, Room 110
- The outside vision speakers this semester are Ce Liu (10/14) and Antonio Torralba (11/6). N

Imageworld digest

- Lots of job postings, PhD positions, postdoc positions, faculty positions and conferences.
- http://lists.diku.dk/mailman/listinfo/imageworld

Help with a research project?

Play the body shape similarity game.
View 3 bodies and say which two are most alike.
Full game takes about 20 min. Have to be inside Brown.

http://cslab6g.cs.brown.edu:3000/shape_games/new?shape_game[type]=OddManOutGame&shape_game[name]=run1&shape_game[question_set_id]=4

Bringing Pictorial Space to Life

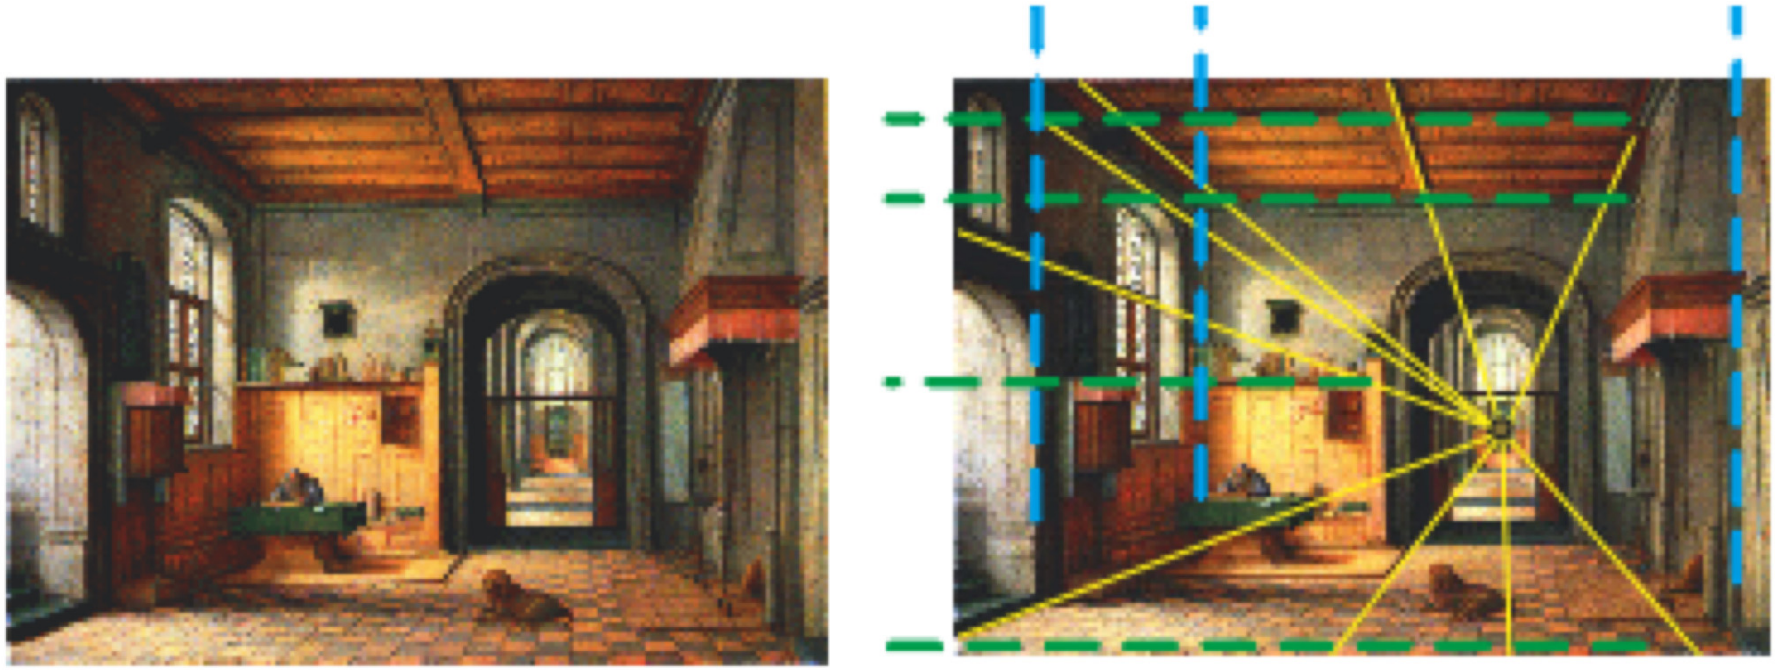

Antonio Criminisi:
http://research.microsoft.com/pubs/67260/criminisi_chart2002.pdf
http://research.microsoft.com/en-us/um/people/antcrim/acriminisi_singleviewmetrology.wmv

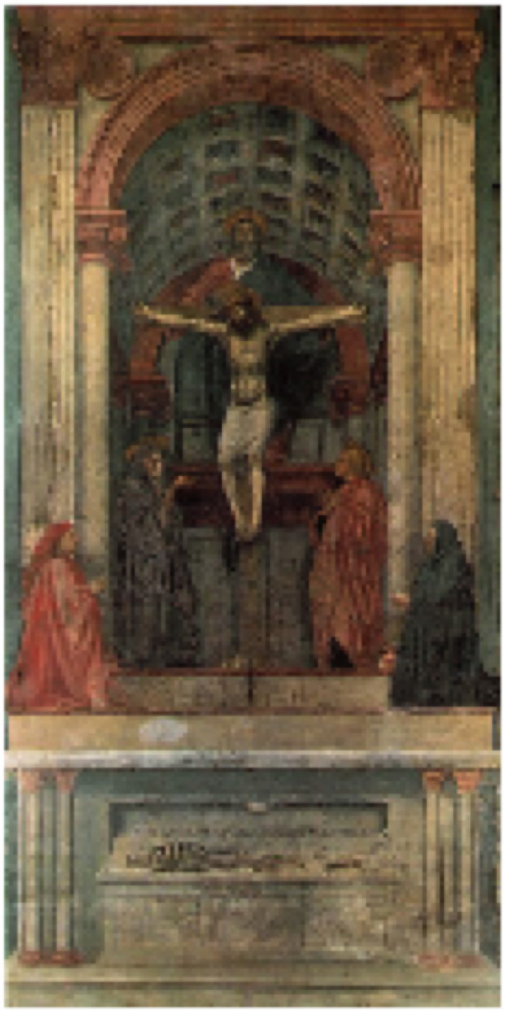

a

b

d

c

e

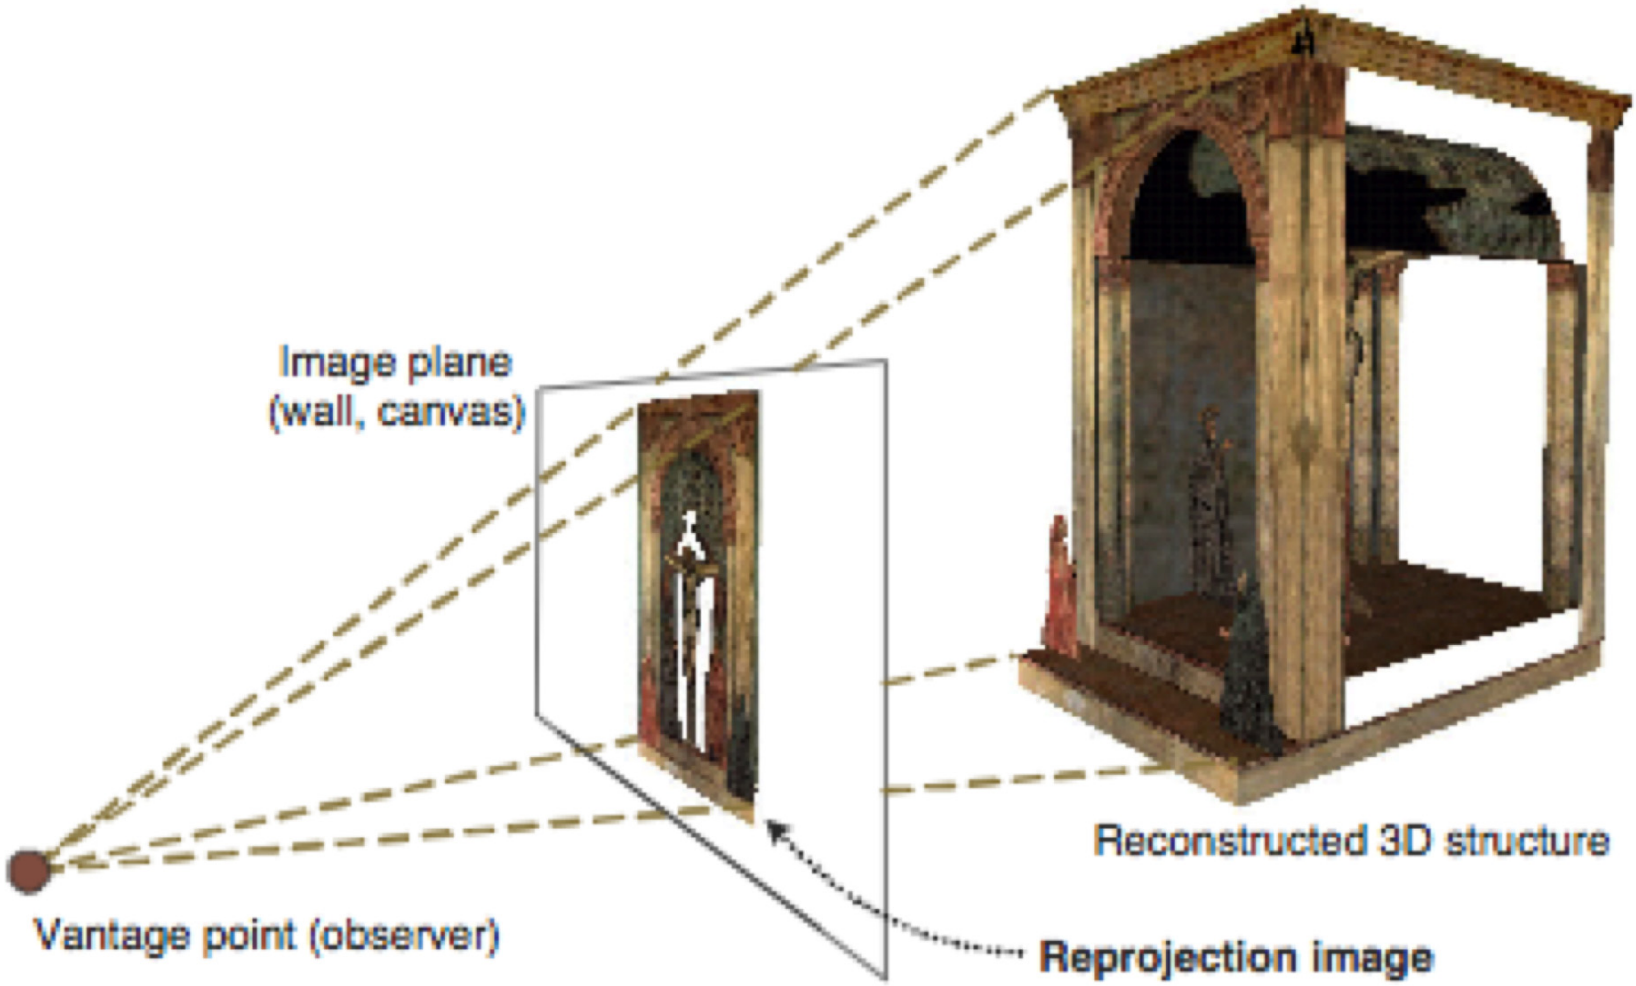
Image plane
(wall, canvas)
Reconstructed 3D structure
Vantage point (observer)
Reprojection image

a

f

e

b

c

d

Automatic Photo Pop-Up

(a) input image

(b) superpixels

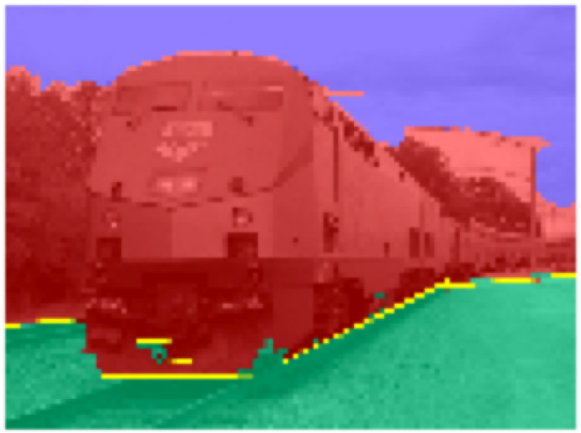

(a) Fitted Segments

(b) Cuts and Folds

Popup

- http://www.cs.uiuc.edu/homes/dhoiem/projects/popup/index.html

(e) novel view

Make3D

http://make3d.stanford.edu/

- Saxena and Ng. Submit your photo.
- Have a look: http://make3d.stanford.edu/images/showall

Goals for Today

Continue with introduction.

What does it mean to see and how do we do it?

How can we make this formal (mathematical and computational)?

Consider a case study of object recognition – in class “group” work.

Computer Vision

First pass at a definition:

- take all the cues of artists and "turn them around"

- exploit these cues to infer the structure of the world

- need mathematical and computational models of these cues

- sometimes called "inverse graphics"

Idea 1: model physics of image formation and find the best "model" that matches the image observations.

Crater illusion

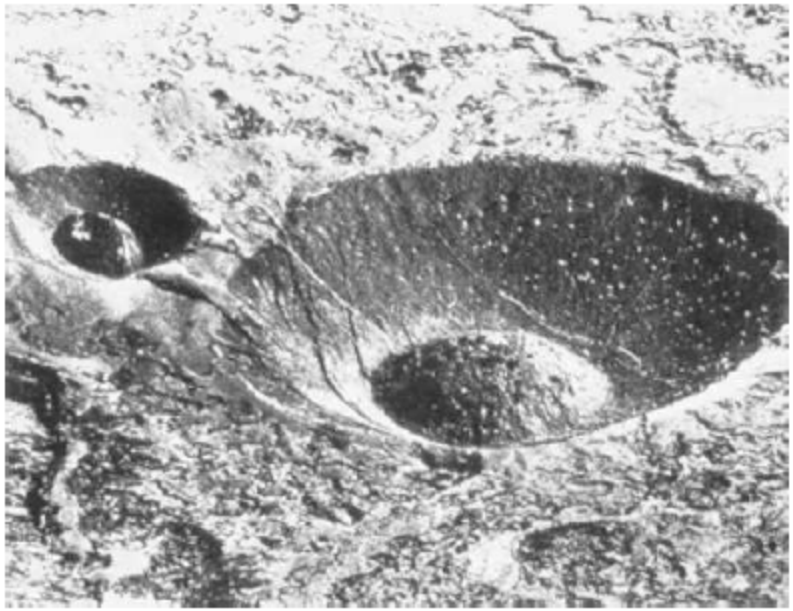

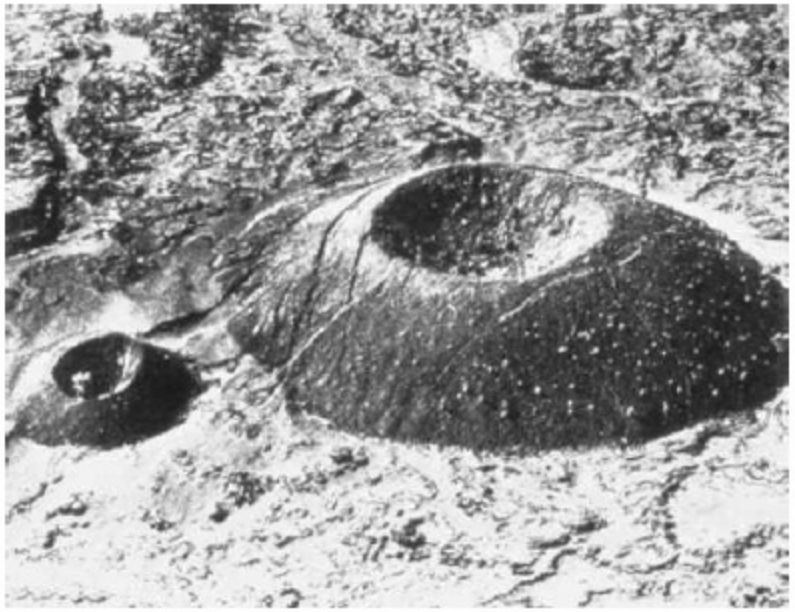

A. Pentland. Local shading analysis. *Trans. PAMI*, 6:170–187, 1984.

This picture is of an ash cone in the Hawaiian Islands (courtesy of W. Richards).

Very powerful effect

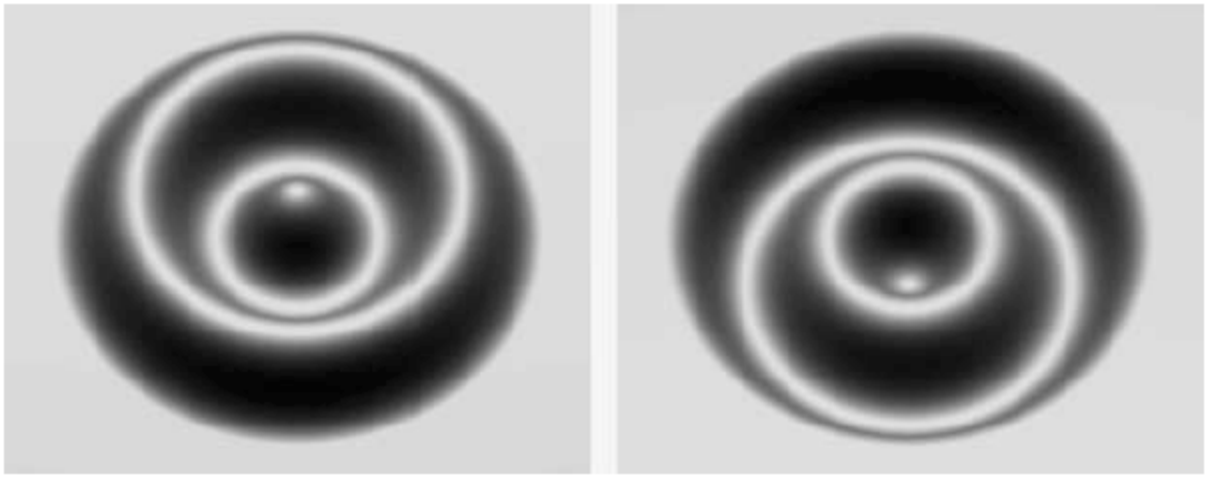

Liu and Todd, Vision Research 2004.

Idea 2: we need more than just a model of the physics of image formation. Need to model something about our prior experience with the world.

Light and shadows

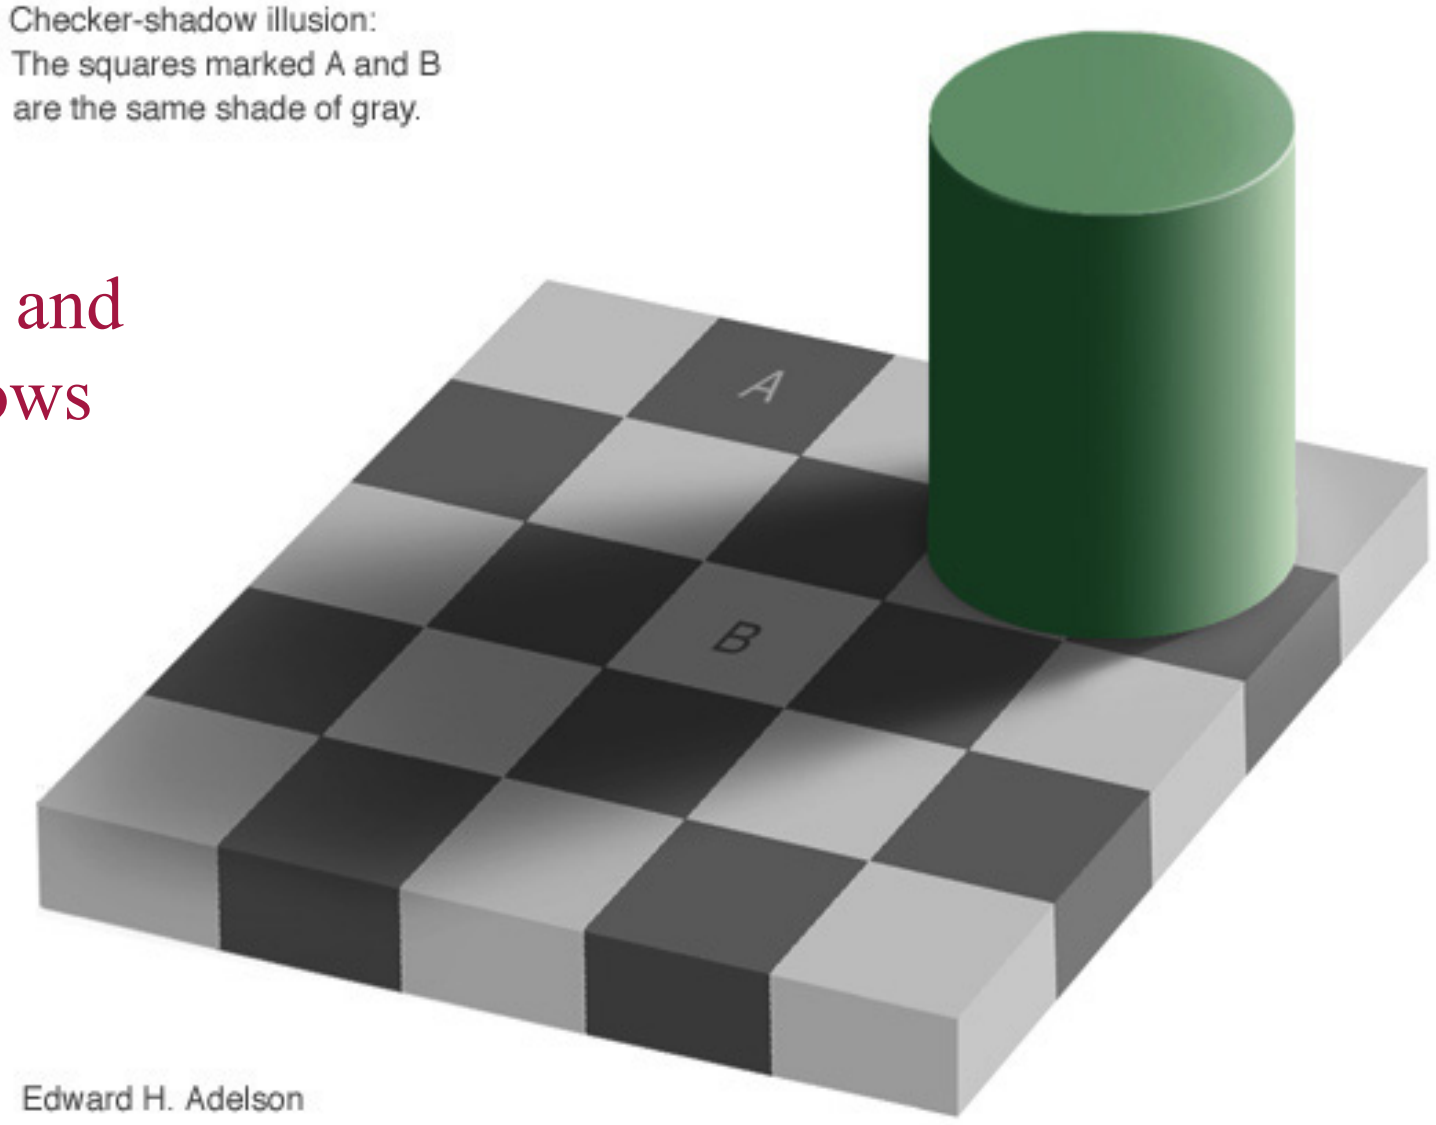

Checker-shadow illusion:
The squares marked A and B
are the same shade of gray.

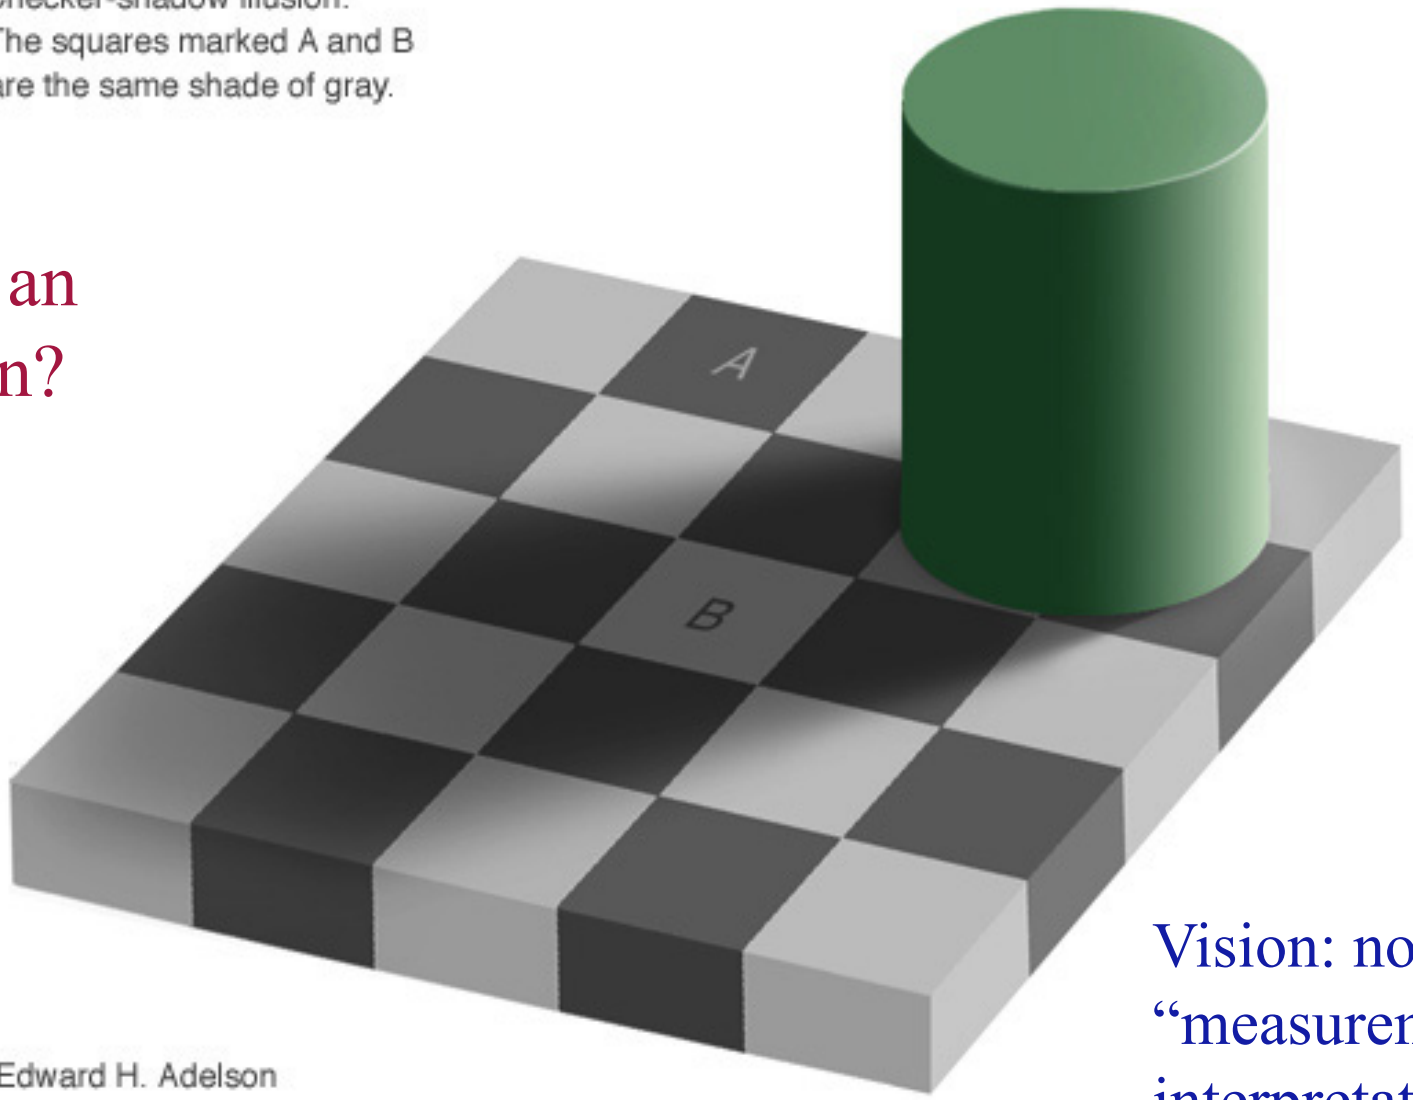

Edward H. Adelson

Is this an illusion?

Vision: not just "measurement" – interpretation.

Combination of cues

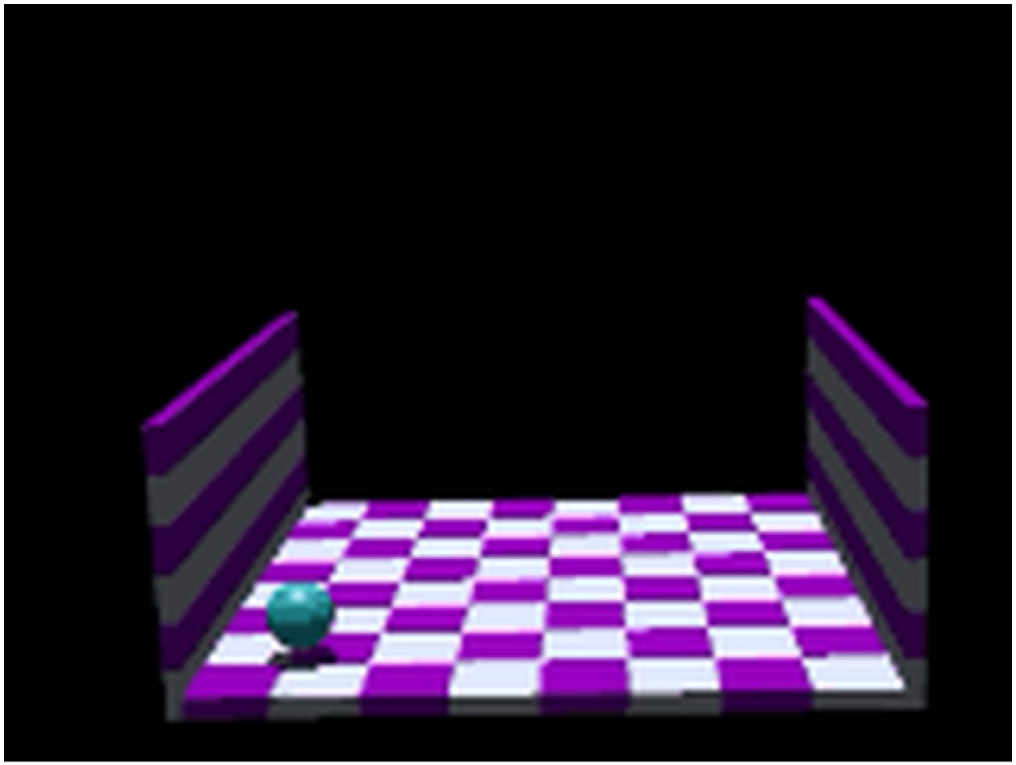

Dan Kersten, http://vision.psych.umn.edu/users/kersten/kersten-lab/shadows.html

Hartung, B., & Kersten, D. (2002). Distinguishing Shiny from Matte.

Does our brain "represent" the physical truth?

- change blindness movies

R. Rensink

Does our brain "represent" the truth?

R. Rensink

Does our brain "represent" the truth?

R. Rensink

Does our brain "represent" the truth?

R. Rensink

Does our brain "represent" the truth?

R. Rensink