CS6670: Computer Vision

Noah Snavely

Lecture 20: Light, color, and reflectance

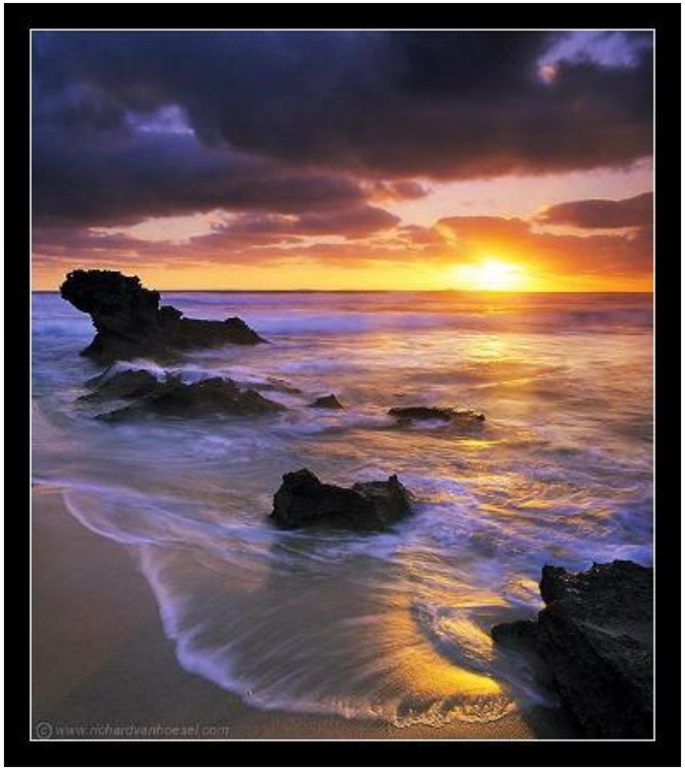

Announcements

- Final projects
  - Everyone should have received feedback by now
  - Midterm reports due November 24
  - Final presentations tentatively scheduled for the final exam period:

    Wed, December 16, 7:00 PM - 9:30 PM

Measuring height

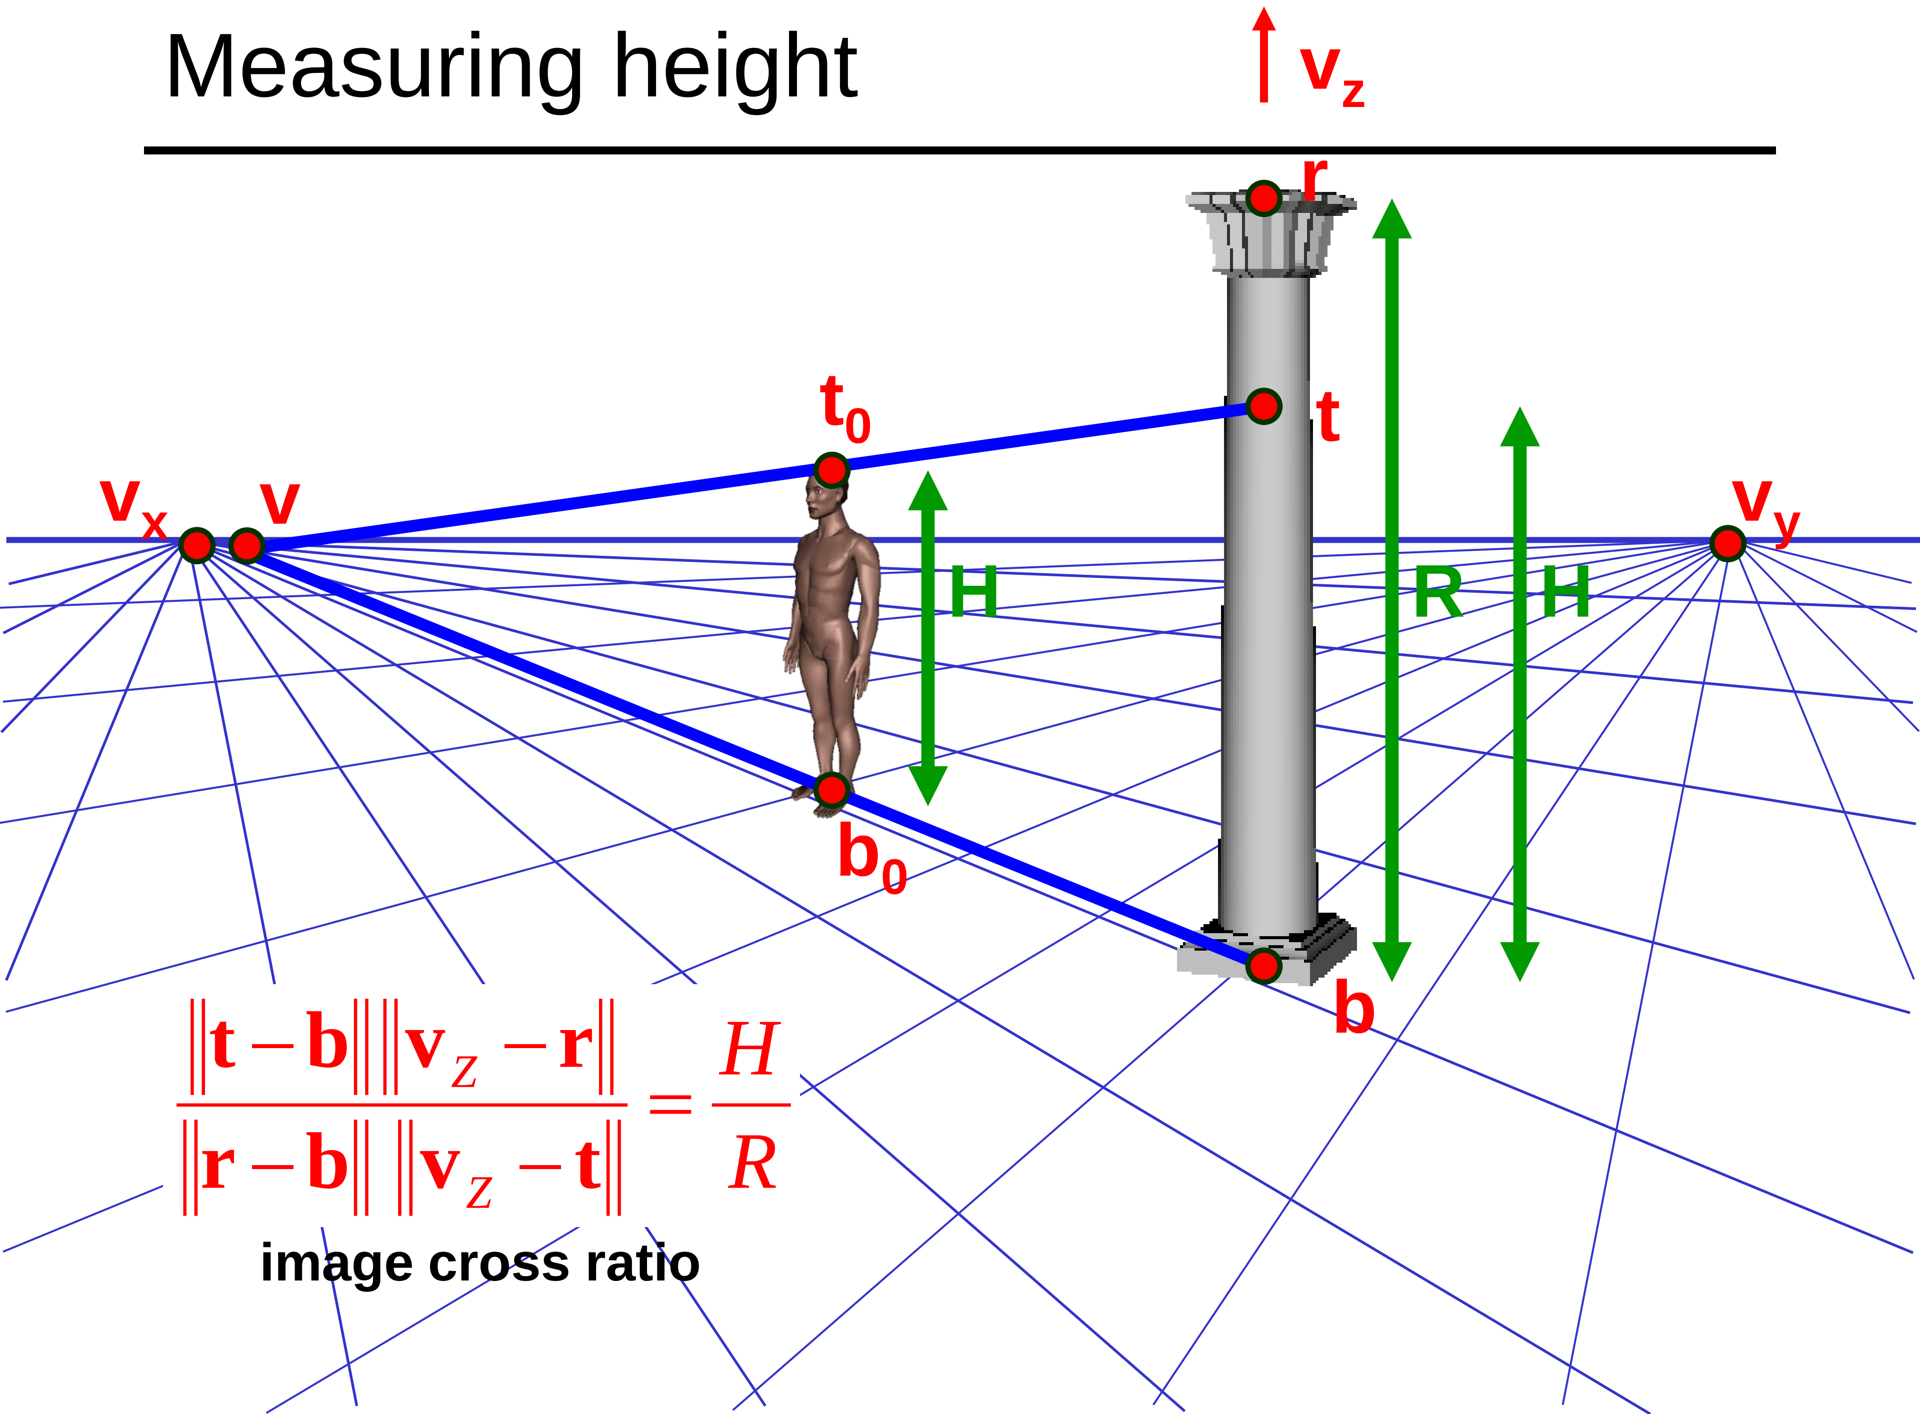

$$\frac{\|\mathbf{t}-\mathbf{b}\|\,\|\mathbf{v}_Z-\mathbf{r}\|}{\|\mathbf{r}-\mathbf{b}\|\,\|\mathbf{v}_Z-\mathbf{t}\|}=\frac{H}{R}$$

image cross ratio

Measuring height

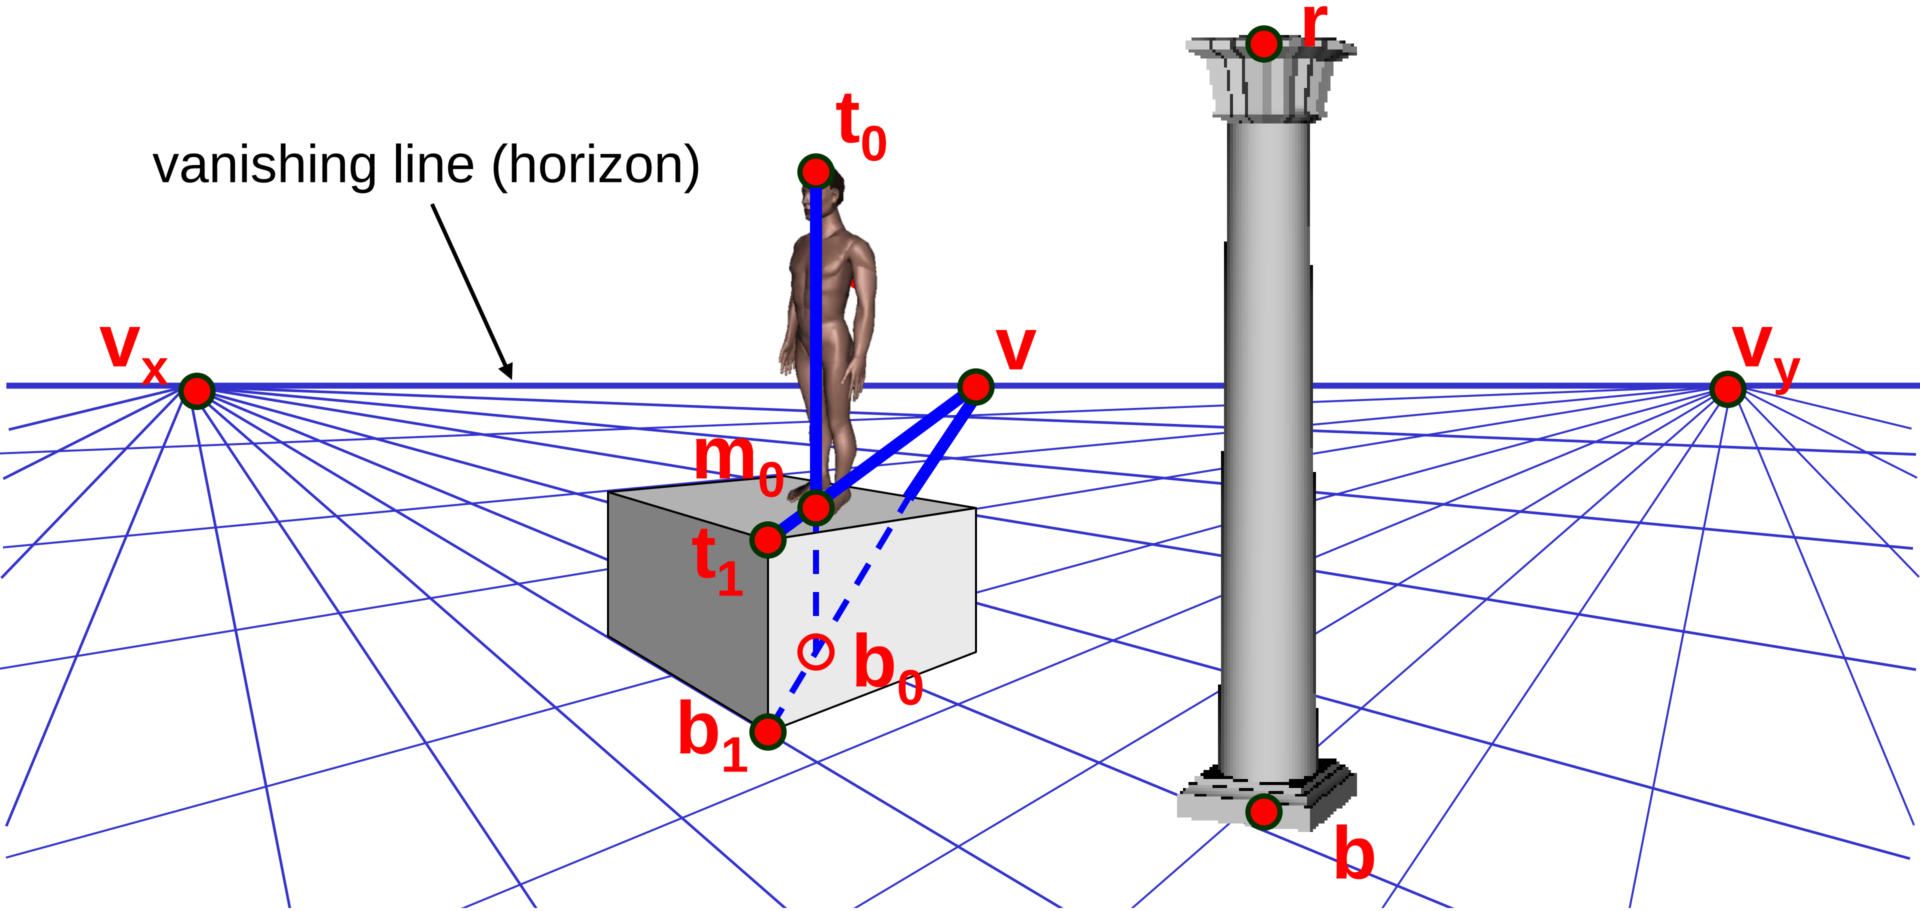

What if the point on the ground plane $\mathbf{b_0}$ is not known?

- Here the guy is standing on the box, height of box is known
- Use one side of the box to help find $\mathbf{b_0}$ as shown above

Camera calibration

Goal: estimate the camera parameters

- Version 1: solve for projection matrix

$$\mathbf{x} = \begin{bmatrix} wx \\ wy \\ w \end{bmatrix} = \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = \mathbf{\Pi X}$$

- Version 2: solve for camera parameters separately
  - intrinsics (focal length, principle point, pixel size)
  - extrinsics (rotation angles, translation)
  - radial distortion

Vanishing points and projection matrix

$$\mathbf{\Pi} = \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} = \begin{bmatrix} \boldsymbol{\pi}_1 & \boldsymbol{\pi}_2 & \boldsymbol{\pi}_3 & \boldsymbol{\pi}_4 \end{bmatrix}$$

$\boldsymbol{\pi}_1 \quad \boldsymbol{\pi}_2 \quad \boldsymbol{\pi}_3 \quad \boldsymbol{\pi}_4$

- $\boldsymbol{\pi}_1 = \mathbf{\Pi}\begin{bmatrix}1 & 0 & 0 & 0\end{bmatrix}^T$ = $\mathbf{v}_X$ (X vanishing point)
- similarly, $\boldsymbol{\pi}_2 = \mathbf{v}_Y$, $\boldsymbol{\pi}_3 = \mathbf{v}_Z$
- $\boldsymbol{\pi}_4 = \mathbf{\Pi}\begin{bmatrix}0 & 0 & 0 & 1\end{bmatrix}^T = \text{projection of world origin}$

$$\mathbf{\Pi} = \begin{bmatrix} \mathbf{v}_X & \mathbf{v}_Y & \mathbf{v}_Z & \mathbf{o} \end{bmatrix}$$

Not So Fast! We only know **v**’s up to a scale factor

$$\mathbf{\Pi} = \begin{bmatrix} a\mathbf{v}_X & b\mathbf{v}_Y & c\mathbf{v}_Z & \mathbf{o} \end{bmatrix}$$

- Can fully specify by providing 3 reference points

Calibration using a reference object

Place a known object in the scene

- identify correspondence between image and scene
- compute mapping from scene to image

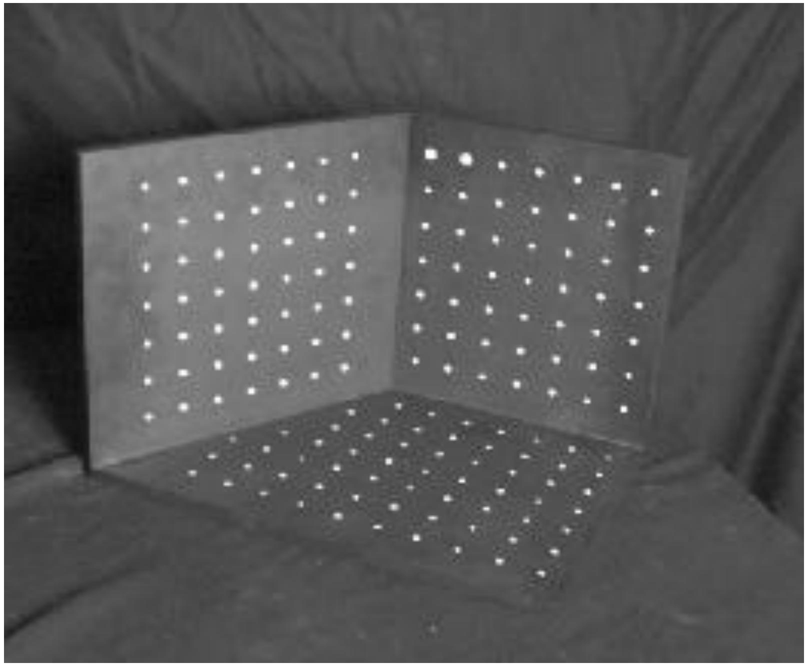

Issues

- must know geometry very accurately
- must know 3D->2D correspondence

Chromaglyphs

Courtesy of Bruce Culbertson, HP Labs
http://www.hpl.hp.com/personal/Bruce_Culbertson/ibr98/chromagl.htm

Estimating the projection matrix

Place a known object in the scene

- identify correspondence between image and scene
- compute mapping from scene to image

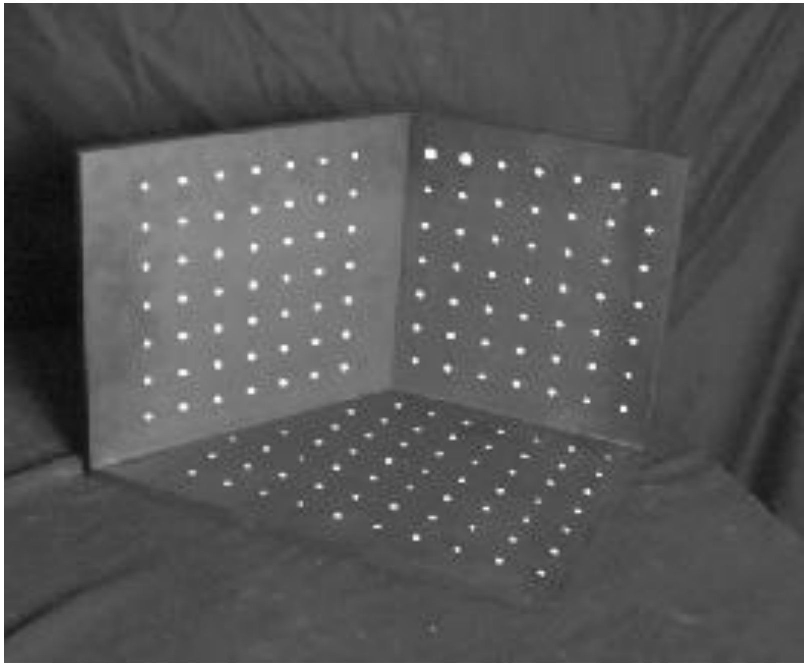

$$\begin{bmatrix} u_i \\ v_i \\ 1 \end{bmatrix} \cong \begin{bmatrix} m_{00} & m_{01} & m_{02} & m_{03} \\ m_{10} & m_{11} & m_{12} & m_{13} \\ m_{20} & m_{21} & m_{22} & m_{23} \end{bmatrix} \begin{bmatrix} X_i \\ Y_i \\ Z_i \\ 1 \end{bmatrix}$$

Direct linear calibration

$$\begin{bmatrix} u_i \\ v_i \\ 1 \end{bmatrix} \cong \begin{bmatrix} m_{00} & m_{01} & m_{02} & m_{03} \\ m_{10} & m_{11} & m_{12} & m_{13} \\ m_{20} & m_{21} & m_{22} & m_{23} \end{bmatrix} \begin{bmatrix} X_i \\ Y_i \\ Z_i \\ 1 \end{bmatrix}$$

$$\begin{aligned} u_i &= \frac{m_{00}X_i + m_{01}Y_i + m_{02}Z_i + m_{03}}{m_{20}X_i + m_{21}Y_i + m_{22}Z_i + m_{23}} \\ v_i &= \frac{m_{10}X_i + m_{11}Y_i + m_{12}Z_i + m_{13}}{m_{20}X_i + m_{21}Y_i + m_{22}Z_i + m_{23}} \end{aligned}$$

$$u_i(m_{20}X_i + m_{21}Y_i + m_{22}Z_i + m_{23}) = m_{00}X_i + m_{01}Y_i + m_{02}Z_i + m_{03}$$
$$v_i(m_{20}X_i + m_{21}Y_i + m_{22}Z_i + m_{23}) = m_{10}X_i + m_{11}Y_i + m_{12}Z_i + m_{13}$$

$$\begin{bmatrix} X_i & Y_i & Z_i & 1 & 0 & 0 & 0 & 0 & -u_iX_i & -u_iY_i & -u_iZ_i & -u_i \\ 0 & 0 & 0 & 0 & X_i & Y_i & Z_i & 1 & -v_iX_i & -v_iY_i & -v_iZ_i & -v_i \end{bmatrix} \begin{bmatrix} m_{00} \\ m_{01} \\ m_{02} \\ m_{03} \\ m_{10} \\ m_{11} \\ m_{12} \\ m_{13} \\ m_{20} \\ m_{21} \\ m_{22} \\ m_{23} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}$$

Direct linear calibration

$$\begin{bmatrix} X_1 & Y_1 & Z_1 & 1 & 0 & 0 & 0 & 0 & -u_1X_1 & -u_1Y_1 & -u_1Z_1 & -u_1 \\ 0 & 0 & 0 & 0 & X_1 & Y_1 & Z_1 & 1 & -v_1X_1 & -v_1Y_1 & -v_1Z_1 & -v_1 \\ & & & & & & \vdots & & & & & \\ X_n & Y_n & Z_n & 1 & 0 & 0 & 0 & 0 & -u_nX_n & -u_nY_n & -u_nZ_n & -u_n \\ 0 & 0 & 0 & 0 & X_n & Y_n & Z_n & 1 & -v_nX_n & -v_nY_n & -v_nZ_n & -v_n \end{bmatrix} \begin{bmatrix} m_{00} \\ m_{01} \\ m_{02} \\ m_{03} \\ m_{10} \\ m_{11} \\ m_{12} \\ m_{13} \\ m_{20} \\ m_{21} \\ m_{22} \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ 0 \end{bmatrix}$$

Can solve for m_{ij} by linear least squares

- use eigenvector trick that we used for homographies

Direct linear calibration

Advantage:

- Very simple to formulate and solve

Disadvantages:

- Doesn’t tell you the camera parameters
- Doesn’t model radial distortion
- Hard to impose constraints (e.g., known focal length)
- Doesn’t minimize the right error function

For these reasons, *nonlinear methods* are preferred

- Define error function E between projected 3D points and image positions
  - E is nonlinear function of intrinsics, extrinsics, radial distortion
- Minimize E using nonlinear optimization techniques
  - e.g., variants of Newton’s method (e.g., Levenberg Marquart)

Alternative: multi-plane calibration

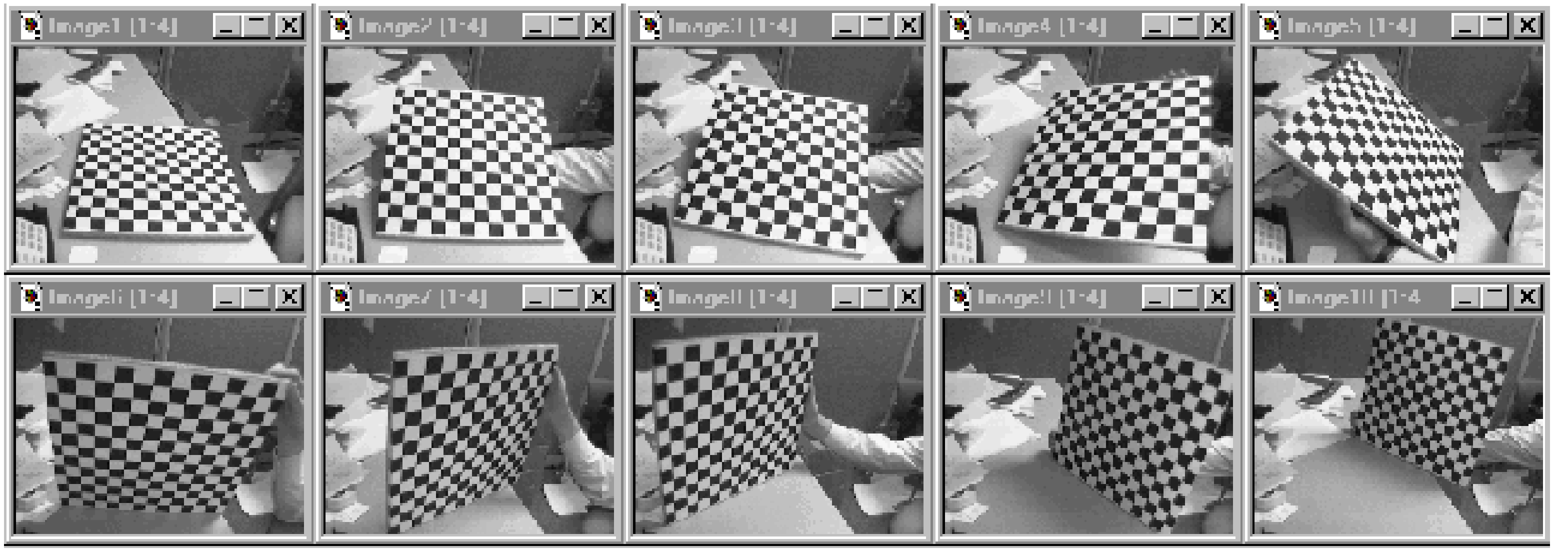

Images courtesy Jean-Yves Bouguet, Intel Corp.

Advantage

- Only requires a plane
- Don't have to know positions/orientations
- Good code available online!
  - Intel's OpenCV library: http://www.intel.com/research/mrl/research/opencv/
  - Matlab version by Jean-Yves Bouget: http://www.vision.caltech.edu/bouguetj/calib_doc/index.html
  - Zhengyou Zhang's web site: http://research.microsoft.com/~zhang/Calib/

Some Related Techniques

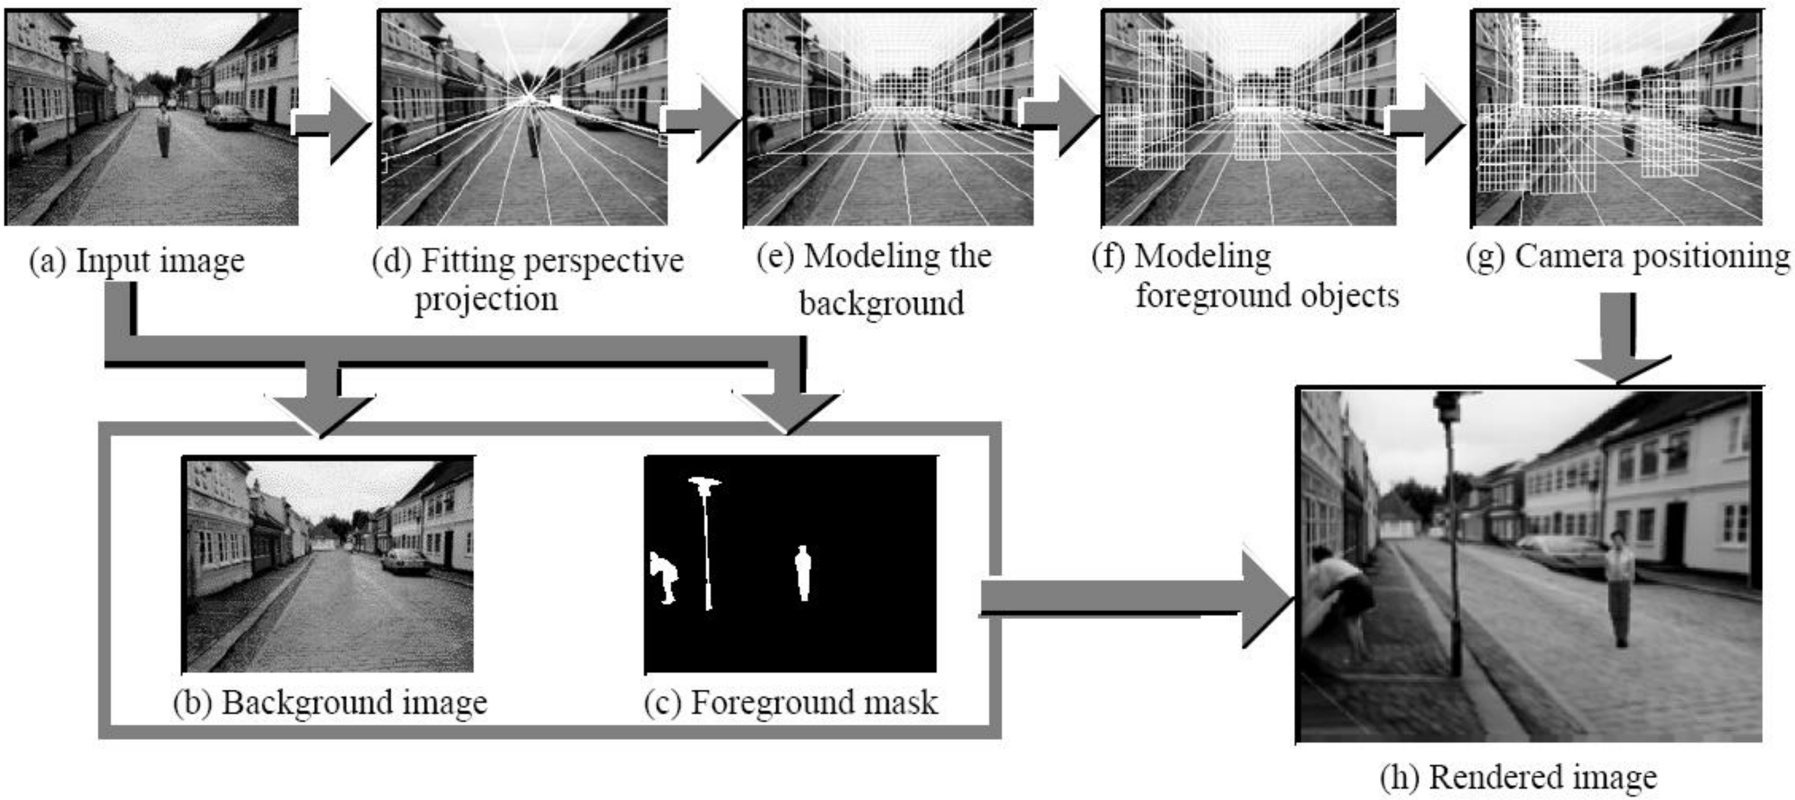

Tour Into The Picture

Anjyo et al., SIGGRAPH 1997

http://koigakubo.hitachi.co.jp/little/DL_TipE.html

Some Related Techniques


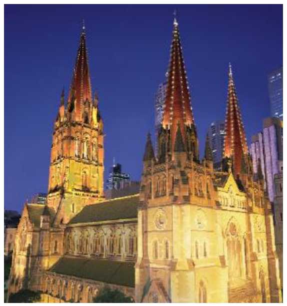

Image-Based Modeling and Photo Editing

- Mok et al., SIGGRAPH 2001
- http://graphics.csail.mit.edu/ibedit/


Single View Modeling of Free-Form Scenes

- Zhang et al., CVPR 2001
- http://grail.cs.washington.edu/projects/svm/

Automatic approaches

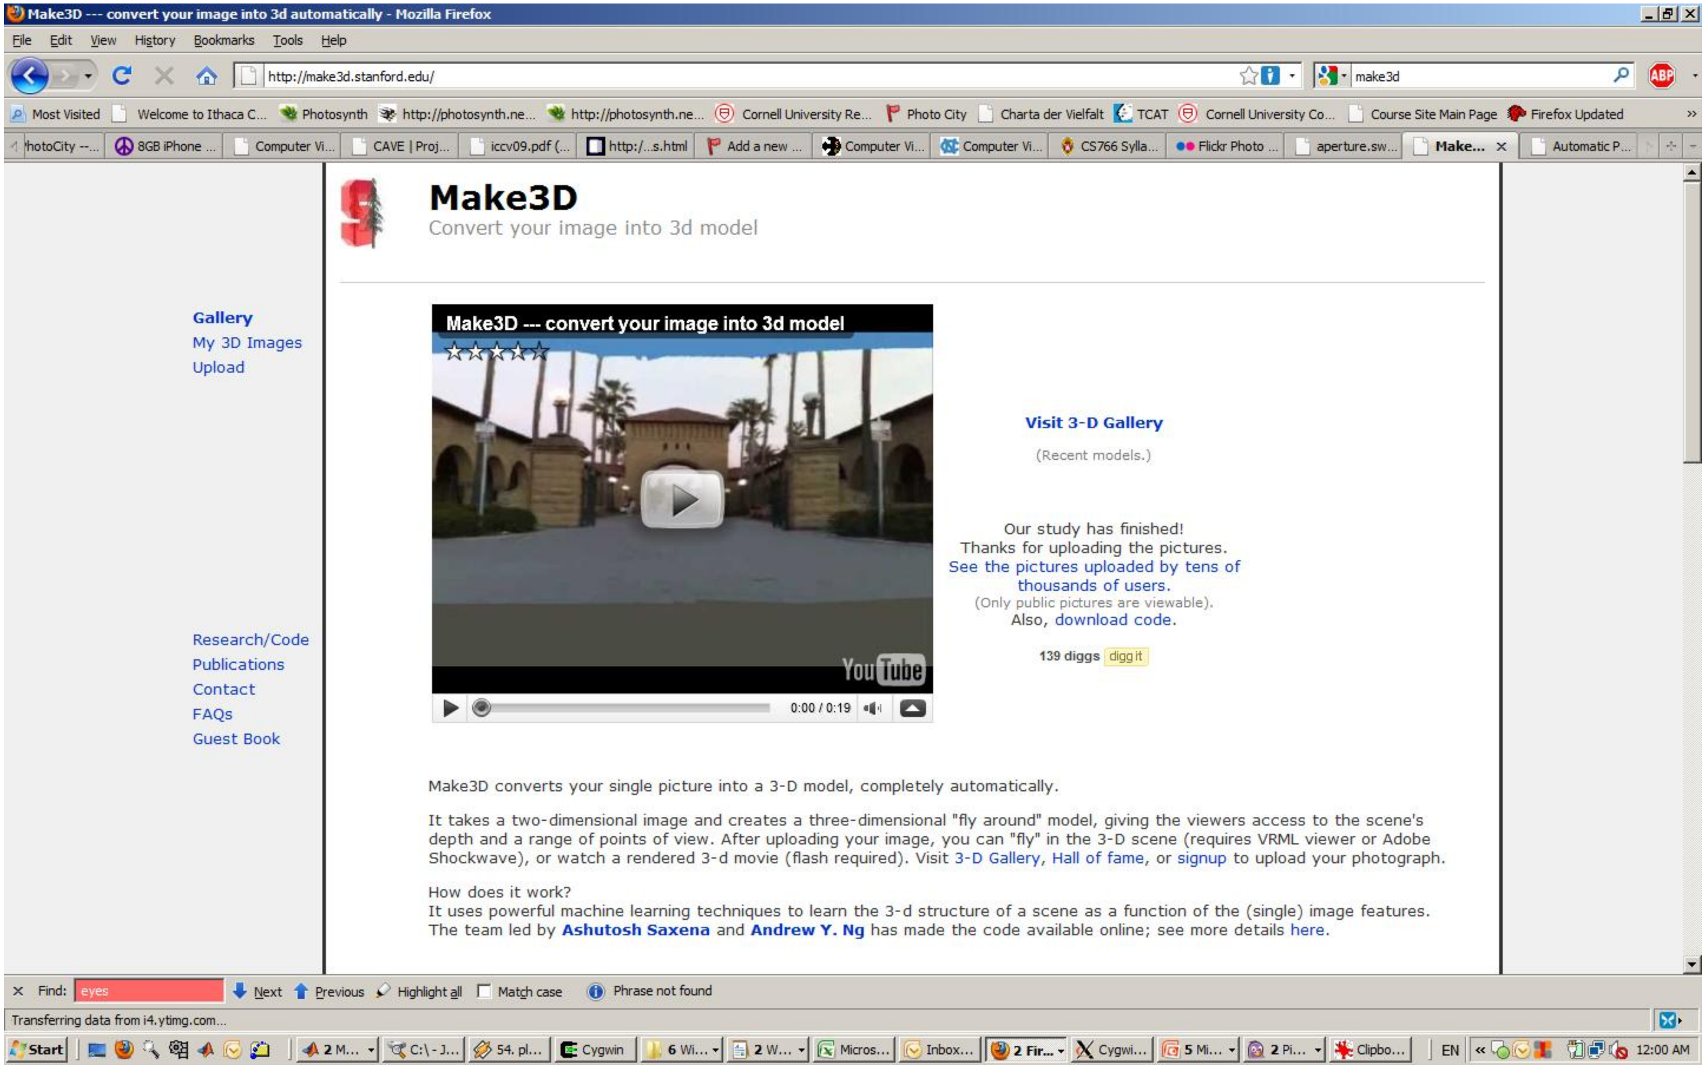

Make3D, Ashutosh Saxena

Automatic approaches

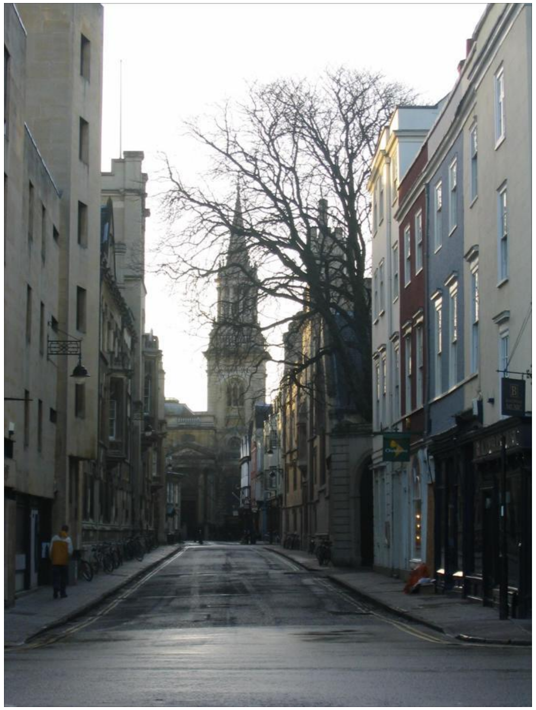

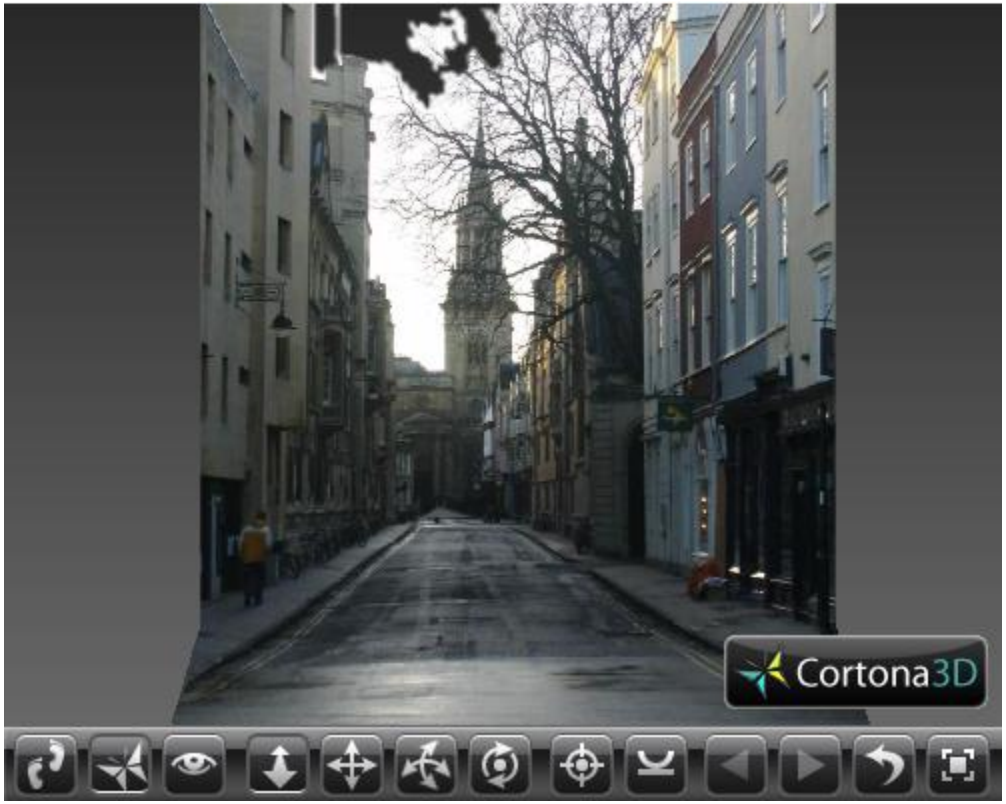

D. Hoiem, A.A. Efros, and M. Hebert, "Automatic Photo Pop-up", ACM SIGGRAPH 2005.

Light

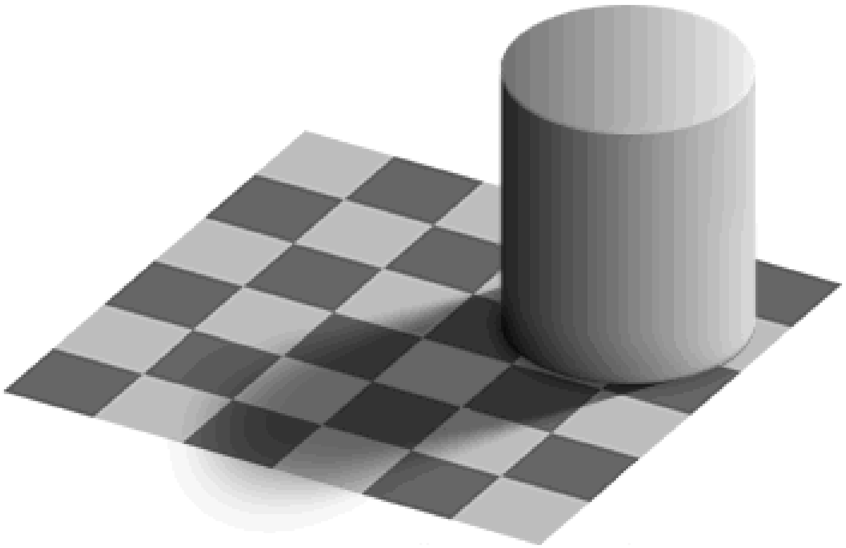

by Ted Adelson

Readings

- Szeliski, 2.2, 2.3.2

Light

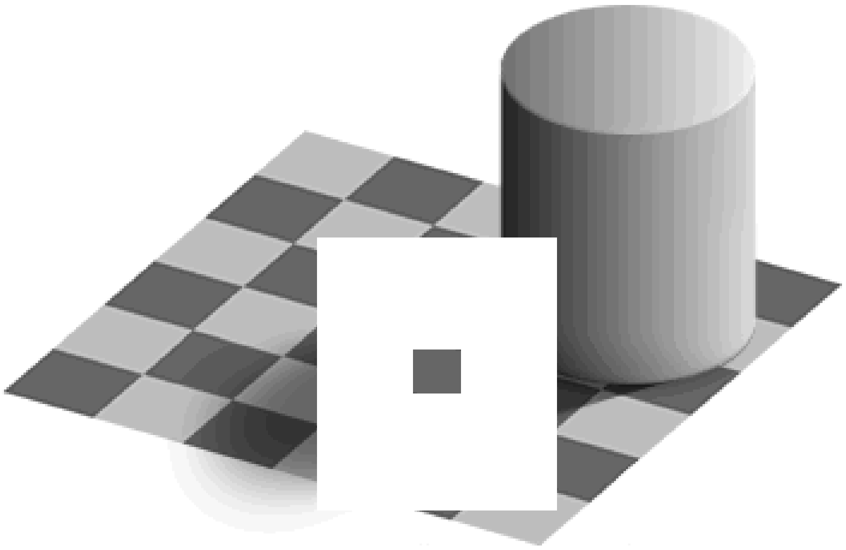

by Ted Adelson

Readings

- Szeliski, 2.2, 2.3.2

Properties of light

Today

- What is light?
- How do we measure it?
- How does light propagate?
- How does light interact with matter?

Radiometry

What determines the brightness of an image pixel?

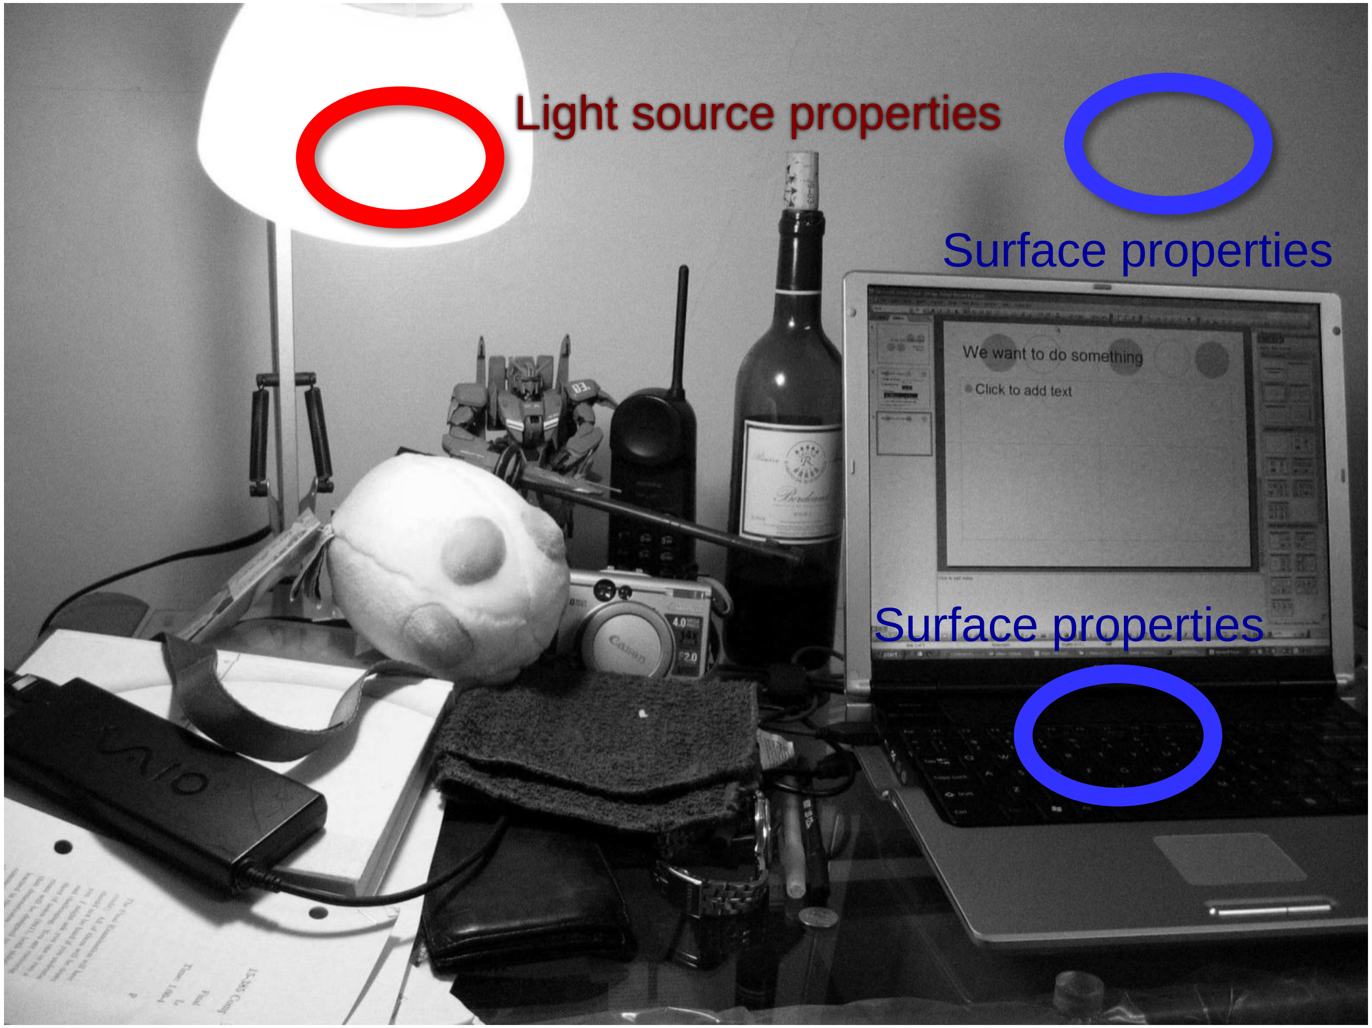

Radiometry

What determines the brightness of an image pixel?

Radiometry

What determines the brightness of an image pixel?

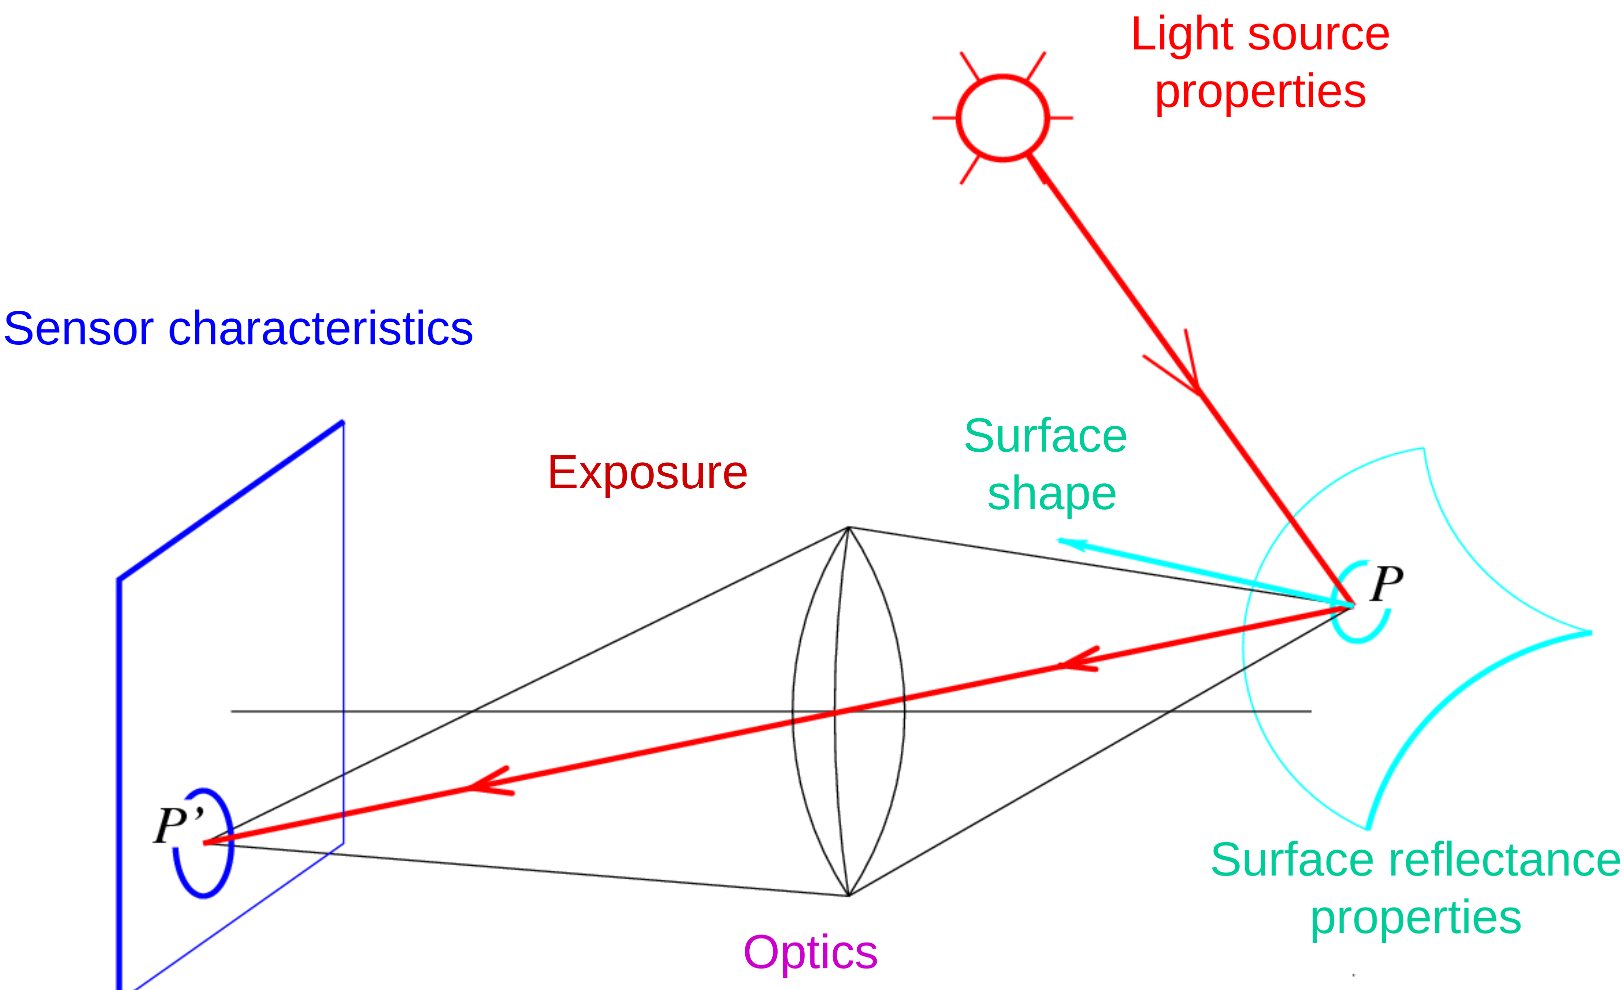

Slide by L. Fei-Fei

What is light?

Electromagnetic radiation (EMR) moving along rays in space

- $R(\lambda)$ is EMR, measured in units of power (watts)
  - λ is wavelength

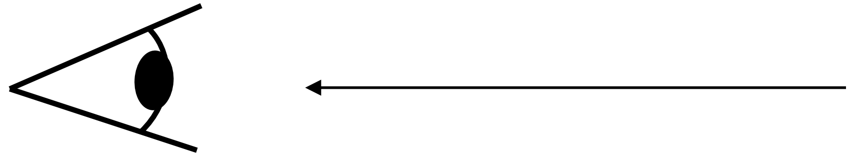

Light field

- We can describe all of the light in the scene by specifying the radiation (or **"radiance"** along all light rays) arriving at every point in space and from every direction

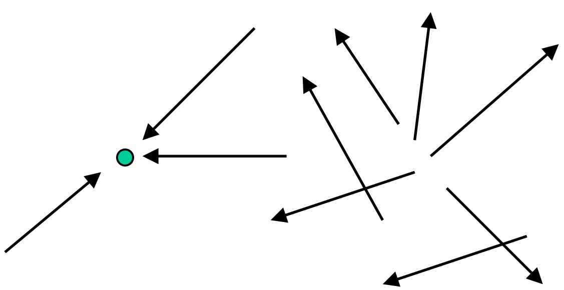

$$R(X, Y, Z, \theta, \phi, \lambda, t)$$

Radiometry

Radiometry is the science of light energy measurement

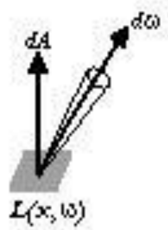

Radiance

The energy carried by a ray

energy/(area solidangle)

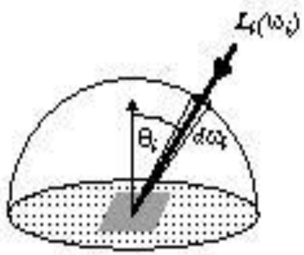

Irradiance

The energy per unit area

falling on a surface

Radiosity

The energy per unit area

leaving a surface

The light field

$$R(X, Y, Z, \theta, \phi, \lambda, t)$$

- Known as the **plenoptic function**
- If you know *R*, you can predict how the scene would appear from any viewpoint.

The **light field** $R(u, v, s, t)$ ← t is *not* time (different from above t !)

- Assume radiance does not change along a ray
  - what does this assume about the world?
- Parameterize rays by intersection with two planes:

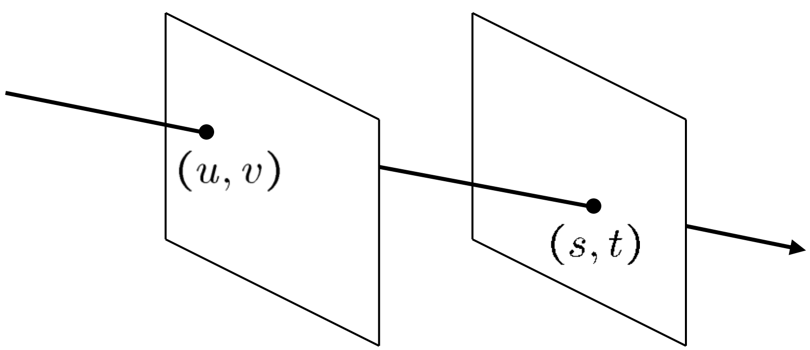

- Usually drop λ and time parameters
- How could you capture a light field?

Capturing light fields

Stanford Multi-Camera Array

Lego Mindstorms Gantry

Handheld light field camera

Light field example

More info on light fields

If you're interested to read more:

The plenoptic function

- **Original reference:** E. Adelson and J. Bergen, "The Plenoptic Function and the Elements of Early Vision," in M. Landy and J. A. Movshon, (eds) Computational Models of Visual Processing, MIT Press 1991.
- L. McMillan and G. Bishop, "Plenoptic Modeling: An Image-Based Rendering System", Proc. SIGGRAPH, 1995, pp. 39-46.

The light field

- M. Levoy and P. Hanrahan, "Light Field Rendering", Proc SIGGRAPH 96, pp. 31-42.
- S. J. Gortler, R. Grzeszczuk, R. Szeliski, and M. F. Cohen, "The lumigraph," in Proc. SIGGRAPH, 1996, pp. 43-54.

What is light?

Electromagnetic radiation (EMR) moving along rays in space

- R(λ) is EMR, measured in units of power (watts)
  - λ is wavelength

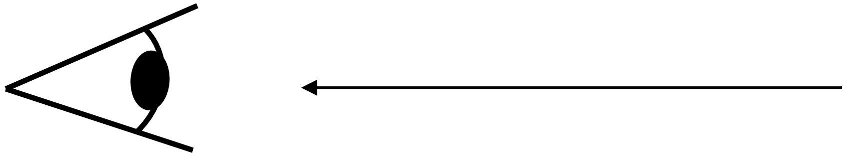

Perceiving light

- How do we convert radiation into “color”?
- What part of the spectrum do we see?

Visible light

We "see" electromagnetic radiation in a range of wavelengths