Motivation

"...abstract representations of light have not been systematically studied in the same way as modeling and display primitives..."
- Levoy and Hanrahan, 1996

How applicable are
object-centered versus image-centered
representations of
radiance versus reflectance
for common
graphics and vision
operations?

Outline of talk

- representations
  - reflectance
  - radiance
- operations
  - changing the viewpoint
  - changing the illumination
  - changing the reflectance
  - changing the shape
- determining reflectance from radiance
  - diffuse, dichromatic, general BRDF
  - non-convex or unknown shape
  - unknown illumination

Object-centered representations

- reflectance
  - 2D surface $\times$ 4D function = 6D
  - Gaussian (Ward, 1992)
  - spherical harmonics (Sillion et al., 1991)
  - wavelets (Schroeder and Sweldens, 1995)
  - functions on the disk (Koenderink et al., 1996)
- radiance
  - 2D surface $\times$ 2D function = 4D
  - hemicube (Immel, 1986)
  - spherical harmonics (Sillion et al., 1991)
  - wavelets, functions on the disk, etc.

Image-centered representations

- radiance
  - 3D space $\times$ 2D function = 5D, or
  - 4D line space $\times$ scalar function = 4D
  - light field (Levoy and Hanrahan, 1996)
  - Lumigraph (Gortler et al., 1996)
- reflectance ?
  - 3D space $\times$ 4D function = 7D, or
  - 4D line space $\times$ 2D function = 6D

Hybrid representations

- radiance
  - 3D scanned surface + projected light field (Levoy)
  - image-based-modeling + projected photographs (Debevec et al., 1996)
  - adaptive ray tracing + light field per object (Teller et al., 1996)
- reflectance ?

Object-based versus image-based rendering

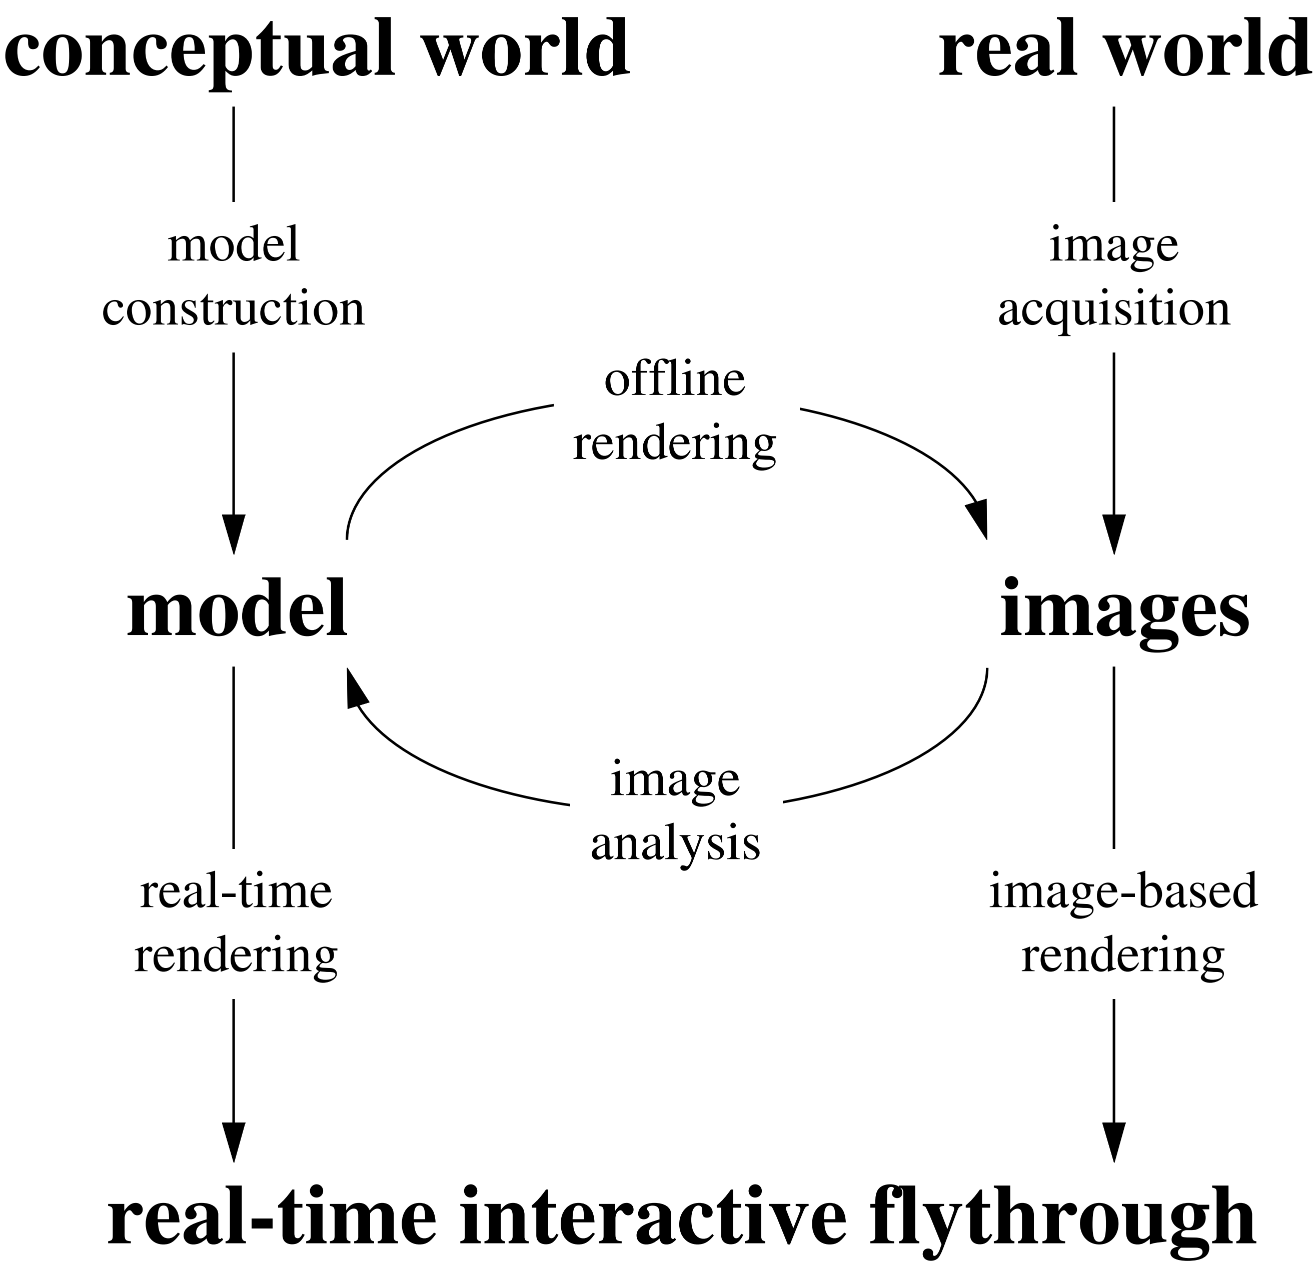

Changing the viewpoint

- object-centered radiance
  - radiance interpolation
    (Teller et al., 1996)
- image-centered radiance
  - light field rendering
    (Levoy and Hanrahan, 1996, Gortler et al., 1996)
- object-centered reflectance
  - classic lighting calculation
    (including global illumination)
- image-centered reflectance ?
  - with **N** per ray
    (no shadows or interreflections)
  - with Z per ray segment (“ZZ light field”)
    (sufficient for global illumination calculations)

Changing the illumination

- 4D lights $\times$ 4D views = 8D
- object-centered radiance
  - illumination basis functions
  - steerable angle (2D)
    (Nimeroff et al., 1994)
  - polynomial position $\times$ steerable angle (4D)
    (Teo & Heeger)
  - position $\times$ position (4D “illumination fields”) **?**
    (Ashdown, 1993, for near-field photometry)
    (untried as basis functions)
- image-centered radiance **?**
  - view interpolation from range images
    (Chen & Williams, 1993; Nimeroff et al., 1995)
- reflectance
  - same as “Changing the viewpoint”

Changing the reflectance

- using a paint program
- object-centered
  - classic 3D painting
  - followed by rerendering
- image-centered **?**
  - light field “painting”
  - requires image-to-image correspondences
  - followed by rerendering
  - global illumination requires knowing 3D shape
  - knowing Z per ray segment might suffice**?**

Changing the shape

- requires knowing reflectance
- object-centered
  - classic 3D modeling
  - followed by rerendering
- image-centered **?**
  - light field “modeling”
  - requires image-to-image correspondences
  - followed by rerendering
  - requires knowing 3D shape**?**

Determining reflectance from radiance

- multiple views of a diffuse convex surface of known shape and albedo under known illumination
  - calculate irradiance on surface
  - reflectance = observed radiance / irradiance

$$\rho_d(2\pi \to 2\pi) = \frac{\pi\, L_{r,d}}{E}$$

  - use multiple views to improve robustness

- dichromatic reflection
  - specular surface reflection + diffuse body reflection (Shafer, 1984)
  - segmentation from single view to compute $\mathbf{C_b}$, $\mathbf{C_s}$ (Klinker et al., 1990)
  - multiple views with moving light or camera (Nayer et al., 1989; Sato & Ikeuchi, 1993)

- general BRDF

Dichromatic reflection model

- colorant particles in bulk medium
  - paints, plastics, ceramics, etc.
- interface reflection
  - Fresnel term, a.k.a. specular term
  - depends on incoming & outgoing directions
  - approximately the color of the illuminant
- body reflection
  - independent of viewing direction (i.e. diffuse)
  - color mainly determined by properties of particles

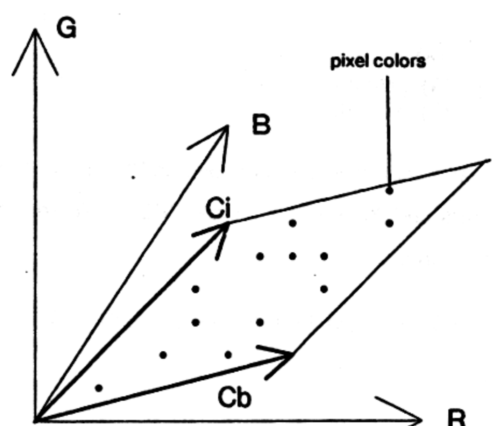

pixel values lie on a parallelogram in color space

(from Shafer)

General BRDFs of 3D objects

- assume collimated illumination
- homogeneous sphere
  - one view captures hemisphere of exiting directions
  - need only move light through incident directions
  - size of pixel and degree of collimation perform the 4D antialiasing integral
  - can use illumination basis functions
- known arbitrary shape
  - each view captures parts of hemisphere
  - use redundant data to improve robustness
  - non-convex is a harder problem
- spatially varying reflectance
  - must build separate BRDFs everywhere

Harder problems

- non-convex: handling shadows
  - shadowless scanning

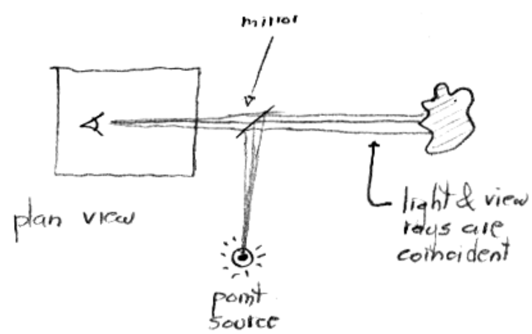

  - by outlier rejection
  - by shadow calculation
- non-convex: handling interreflections
  - diffuse: inverse radiosity
  - general: inverse radiance

Inverse radiosity

- single complete observation of known shape under calibrated illumination
  - radiosity system with known E's and F's, observed B's, and unknown ρ 's

$$\rho_i = \frac{B_i - E_i}{\sum_{j \neq i} B_j F_{i,j}}$$

- multiple partial observations under different calibrated illuminations
  - for each observation, some B's are missing
  - system of non-linear equations
  - iterative solutions converge quickly
  - initialize unknown radiosities and reflectances to zero
    do until converged
        for each patch
            for each observation
                calculate reflectance as above
            estimated reflectance = average of observations
        use classic radiosity solver to compute radiosities
        discard computed radiosities of observed patches

Much harder problems

- unknown illumination
  - light source calibration
    (gazing ball)
- unknown shape
  - homogeneous surface: photometric stereo
    (Horn, 1977; Woodham, 1978)
  - spatially varying: shape from light fields
    (Freyburger & Levoy)
- single view
  - only works for homogeneous surfaces
  - shape from shading
    (Horn, 1970)
  - shape from interreflections
    (Nayer, 1990)

Graphics and vision computations

more info ◄······► less info

forward methods

inverse methods

output (columns) / **input** (rows)

	obj refl	obj rad	LF +z refl	LF +z rad	LF refl	LF rad	im+z refl	im+z rad	im refl	im rad
obj refl	3D paint	lighting	LF build	0	LF build	0	render	0	0	0
obj rad	inverse lighting **?**	1	0	LF build	0	LF build	0	render	0	0
LF +z refl	occup-ancy grid	0	LF zpaint**?**	LF lighting **?**	drop z	0	LF view	0	0	0
LF +z rad	0	occup-ancy grid	LF inverse lighting **?**	1	0	drop z	0	LF view	0	0
LF refl	shape from LF **?**	0	find z from LF **?**	0	LF paint**?**	0	photo-metric stereo	0	LF view	0
LF rad	0	shape from LF **?**	0	find z from LF **?**	pick refl	1	0	0	0	LF view
im +z refl	0	0	0	0	0	0	zpaint	0	drop z	0
im +z rad	0	0	0	0	0	0	pick refl	1	0	drop z
im refl	0	0	0	0	0	0	shape from shading	0	paint	
im rad	0	0	0	0	0	0	0	0	pick refl	1

– elements along diagonal are editing operations

– 0 means no direct path from input to output

– powers of matrix give indirect paths

– proper representation is a 6D tensor