

# Graphics & Video Synthesis

## Image-based Modeling & Rendering (IBMR)

Princeton University  
COS 429 Lecture

Mar. 11, 2004

Harpreet S. Sawhney

[hsawhney@sarnoff.com](mailto:hsawhney@sarnoff.com)



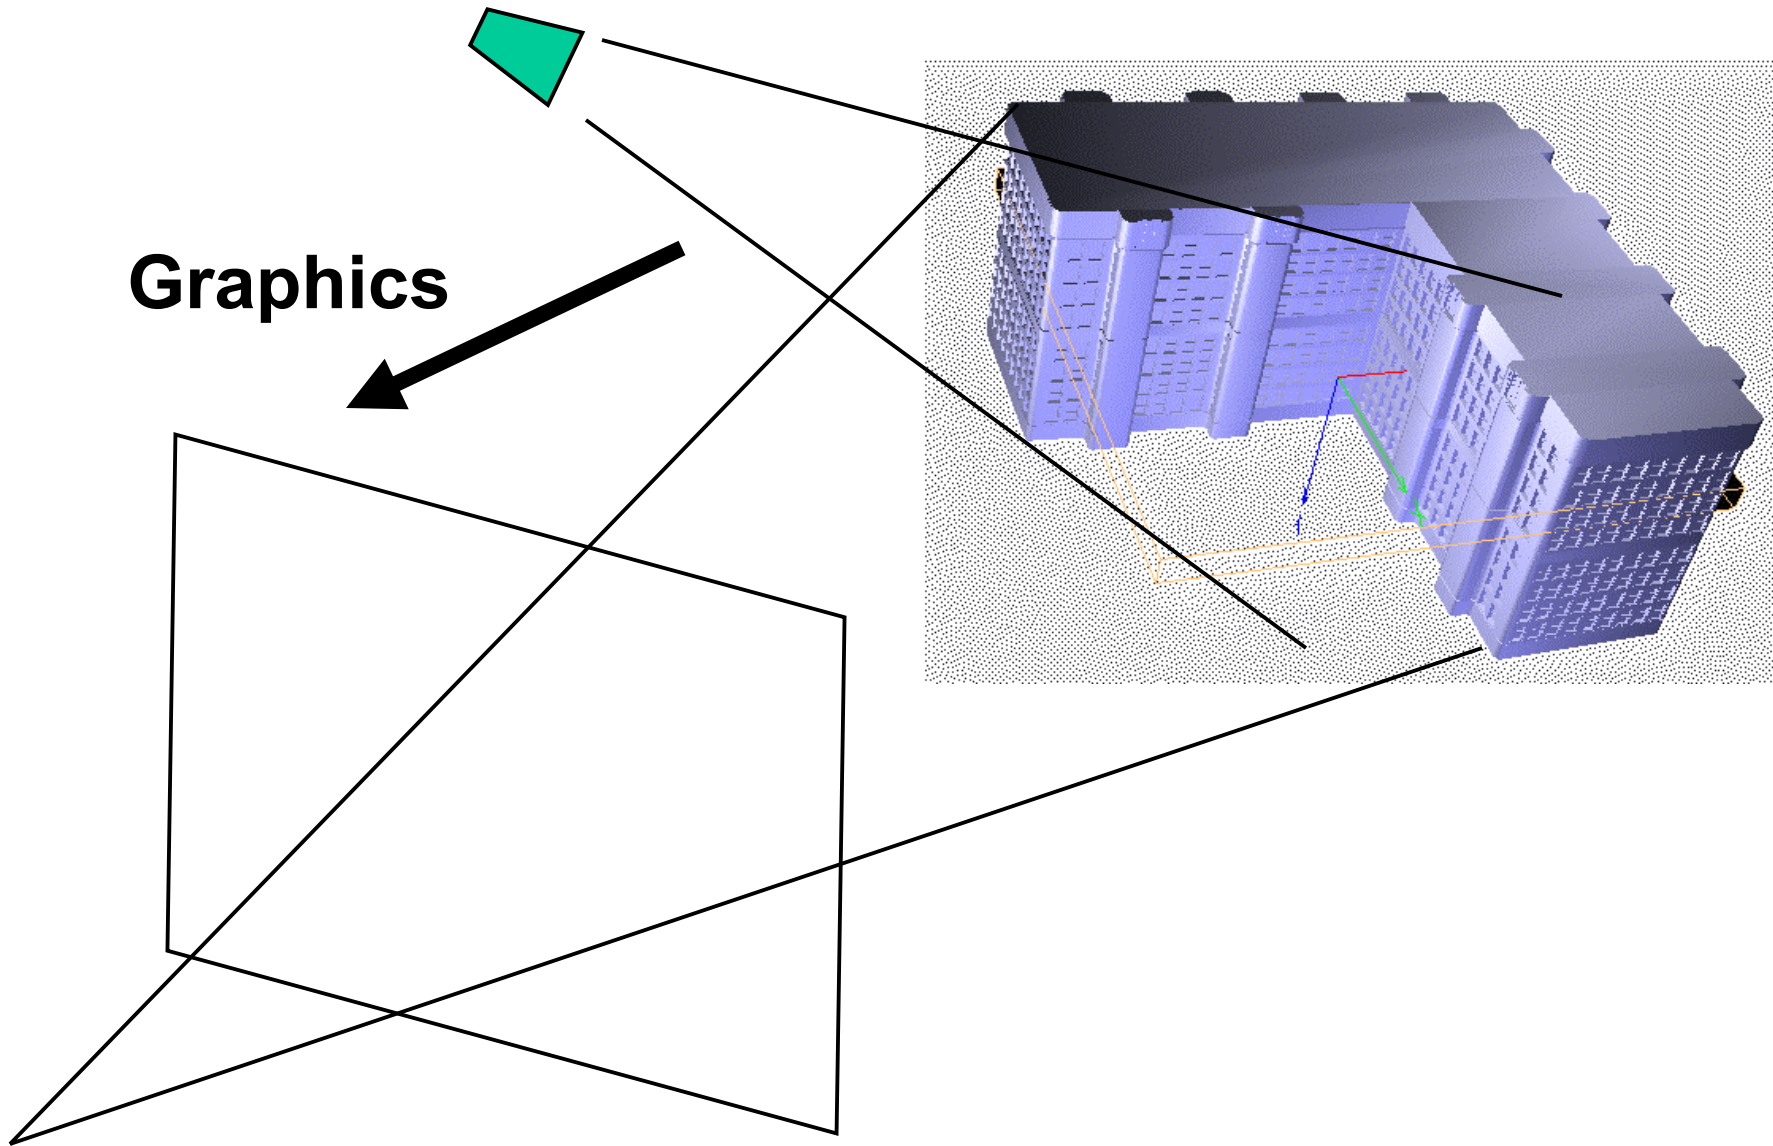
# Recapitulation

- Problem of motion estimation
- Parametric models of motion
- Direct methods for image motion estimation
- Camera models & parametric motion
- Image & video mosaicing as an application

# Plan





- Motivate Image-based Modeling & Rendering (IBMR)
  - Change in viewpoint, IMAX app
- Parameterize motion & structure for video
  - Euclidean case
  - Direct Estimation
- Plane+Parallax
  - Formulation
  - Direct Estimation
- IMAX app.
- Tweening app.
- Model-to-video pose estimation
- Video Flashlights

# Graphics Vs. Vision



# Graphics Vs. Vision

**Vision**



# Modeling Complexity vs. Realism Dilemma



Can we model and render this scene authentically ?

**A very hard vision and graphics problem !**

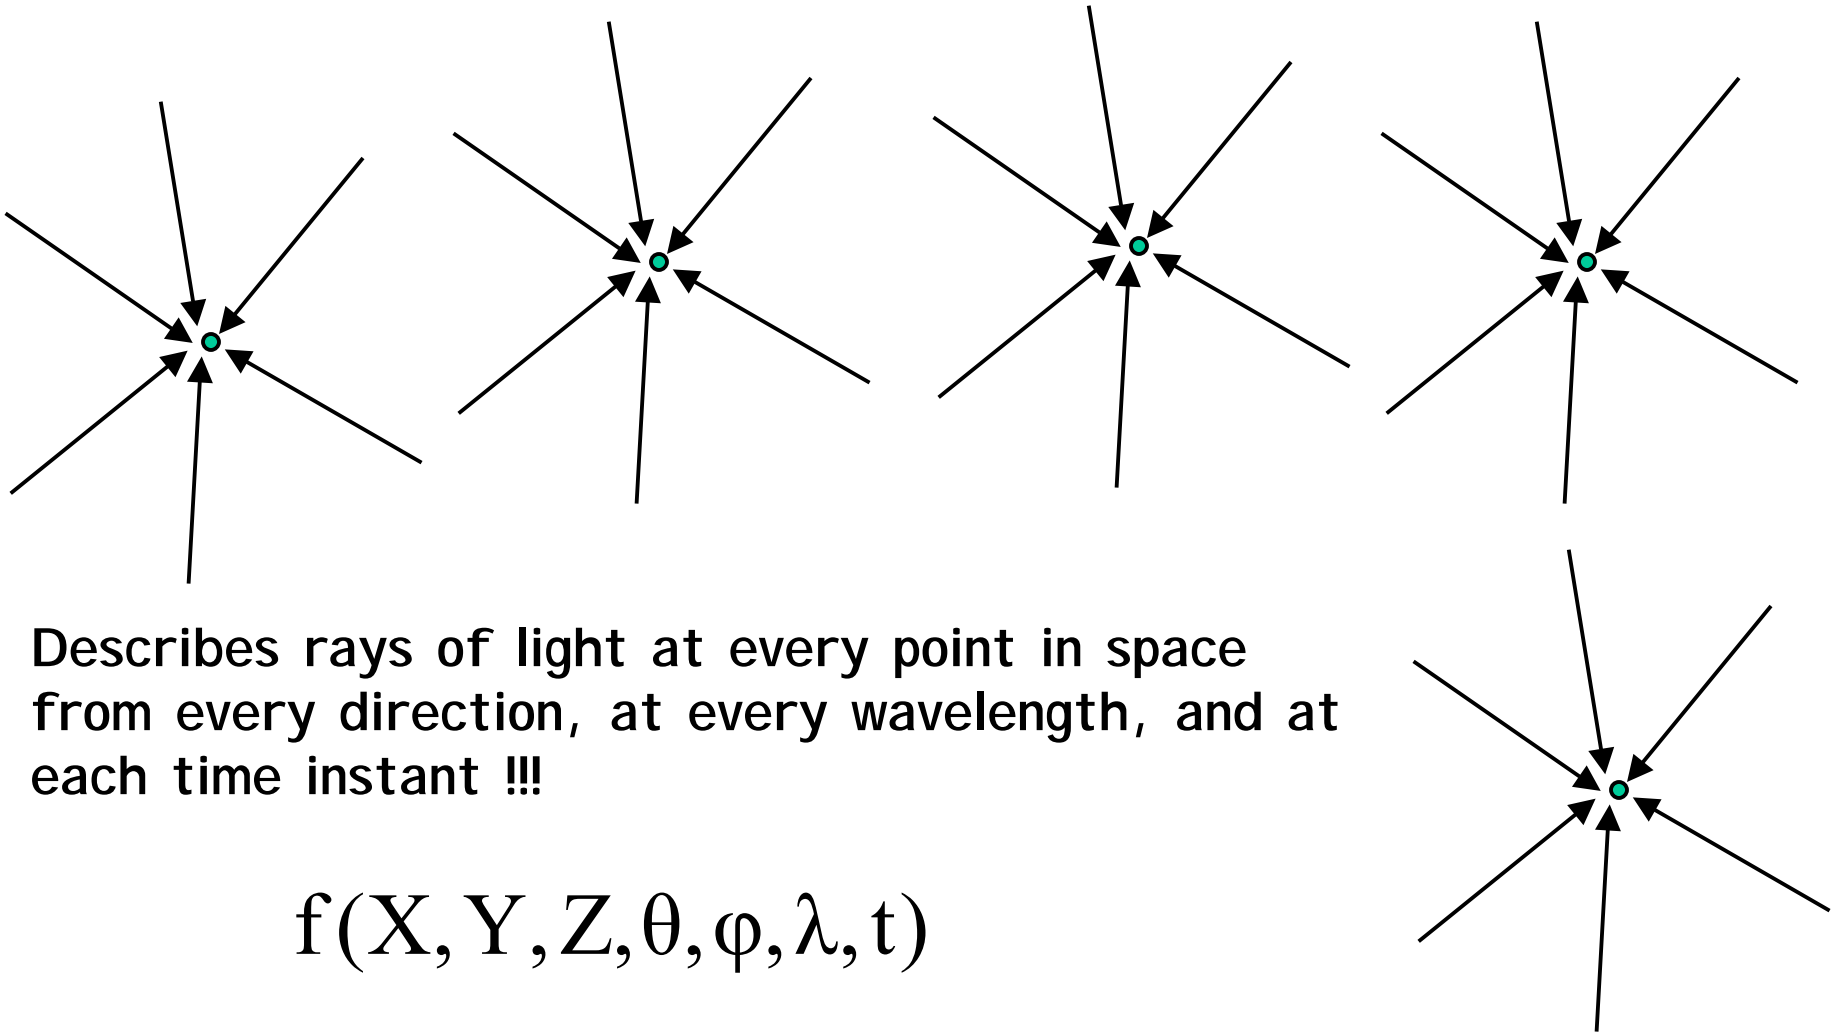
Can we model less

and

fill-in details with more and more

images ?

# Plenoptic Function



Describes rays of light at every point in space from every direction, at every wavelength, and at each time instant !!!

$$f(X, Y, Z, \theta, \varphi, \lambda, t)$$




# Pure IBR


## Use Images to Generate Novel Views

- Mosaics and Panoramas:
  - Projective, Cylindrical, Spherical
- Concentric Mosaics:
  - Restricted change in viewpoint


# An Illustrative Example Concentric Mosaics



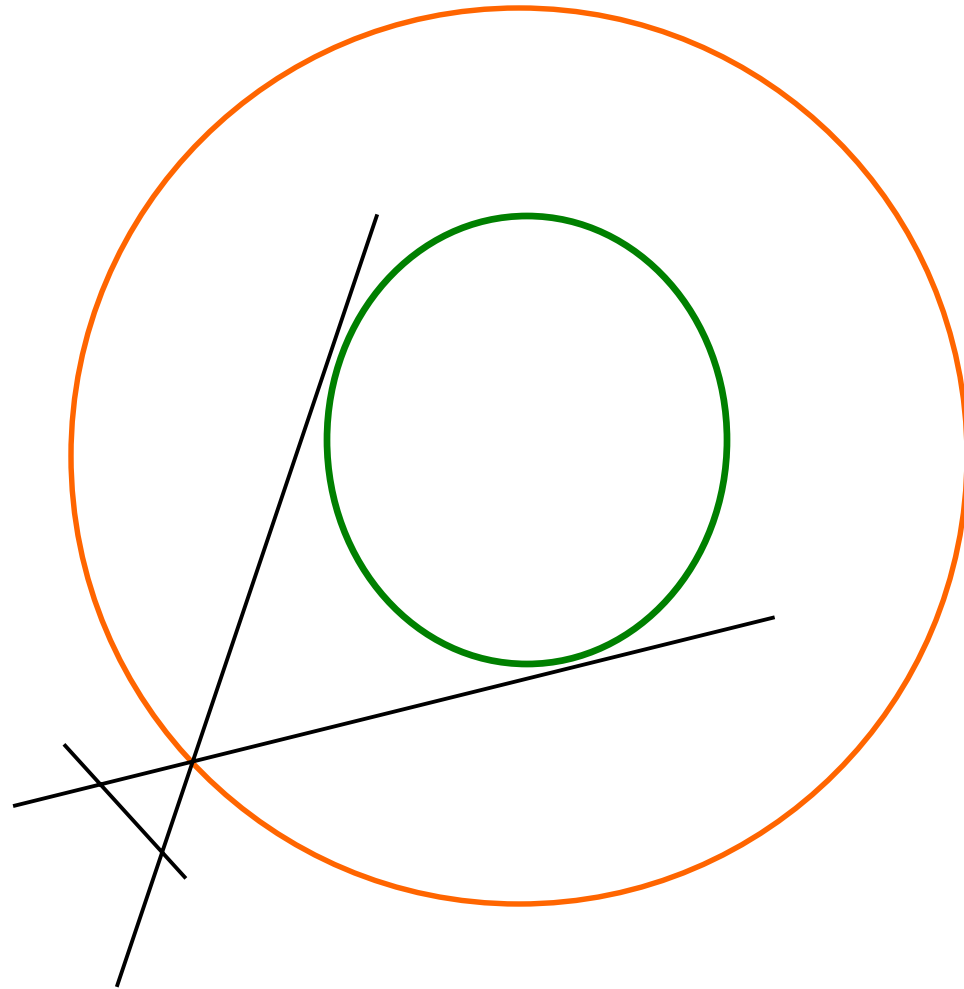
# An Illustrative Example Concentric Mosaics




# An Illustrative Example Concentric Mosaics




# An Illustrative Example Concentric Mosaics




# An Illustrative Example Concentric Mosaics



# An Illustrative Example Concentric Mosaics




# An Illustrative Example Concentric Mosaics







# An Illustrative Example Concentric Mosaics




# New View Generation



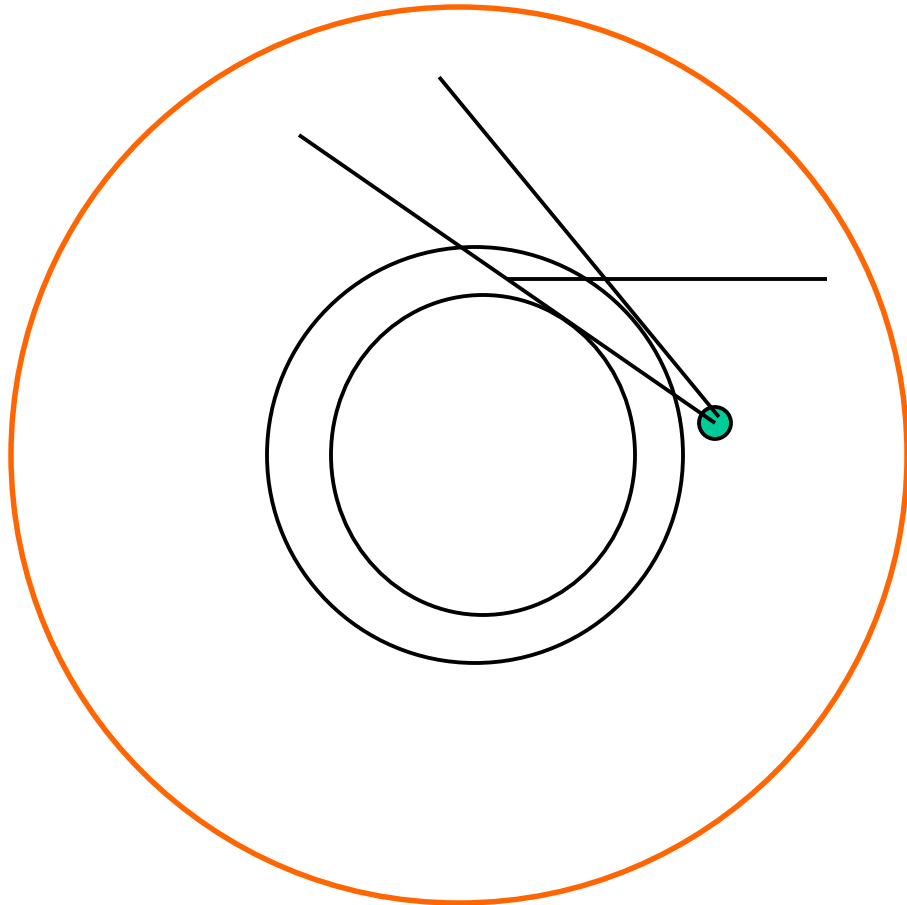
# New View Generation



# New View Generation



# New View Generation




## Details:

- New view rays will not be exactly original pixel locations
- Need to interpolate rays
- Since these rays come from cameras that have different centers of projection, interpolation ideally requires depth estimates

# Local Geometry for IBR

- “Extreme” IBR works only with images
- Need a dense collection to cover even a restricted collection of generatable viewpoints
- Extend the scope of IBR by including local geometry modeling using parallax/depth

# Image Motion : 3D Parameterization



# Image Motion : 3D Parameterization

$$\mathbf{x}' = \frac{\mathbf{X} - \omega_z \mathbf{Y} + \omega_y \mathbf{Z} + \mathbf{T}_x}{-\omega_y \mathbf{X} + \omega_x \mathbf{Y} + \mathbf{Z} + \mathbf{T}_z}$$

$$\mathbf{y}' = \frac{\omega_z \mathbf{X} + \mathbf{Y} - \omega_x \mathbf{Z} + \mathbf{T}_y}{-\omega_y \mathbf{X} + \omega_x \mathbf{Y} + \mathbf{Z} + \mathbf{T}_z}$$

$$\mathbf{x}' = \frac{\mathbf{x} - \omega_z \mathbf{y} + \omega_y + \frac{\mathbf{T}_x}{\mathbf{Z}}}{-\omega_y \mathbf{x} + \omega_x \mathbf{y} + 1 + \frac{\mathbf{T}_z}{\mathbf{Z}}}$$

$$\mathbf{y}' = \frac{\omega_z \mathbf{x} + \mathbf{y} - \omega_x + \frac{\mathbf{T}_y}{\mathbf{Z}}}{-\omega_y \mathbf{x} + \omega_x \mathbf{y} + 1 + \frac{\mathbf{T}_z}{\mathbf{Z}}}$$

Using small motion approximation

$$\mathbf{x}' - \mathbf{x} = \underbrace{\omega_y - \omega_z \mathbf{y} - \omega_x \mathbf{x} \mathbf{y} + \omega_y \mathbf{x}^2}_{\text{Rotational Component}} + \frac{\mathbf{T}_x - \mathbf{T}_z \mathbf{x}}{\mathbf{Z}}$$

Rotational  
Component

Translational  
Component

$$\mathbf{y}' - \mathbf{y} = \underbrace{-\omega_x + \omega_z \mathbf{x} + \omega_y \mathbf{x} \mathbf{y} - \omega_x \mathbf{y}^2}_{\text{Rotational Component}} + \frac{\mathbf{T}_y - \mathbf{T}_z \mathbf{y}}{\mathbf{Z}}$$



# Observations

- For a perspective camera:
  - Rotational image displacements are depth independent
  - Translational displacements contain depth information
- Each displacement vector provides two equations
  - There are 6 motion parameters common to all vectors
  - Depth parameter is typically unique to each vector
- Scale ambiguity between depth and translation

$$\mathbf{x}' - \mathbf{x} = \frac{\mathbf{T}_x - \mathbf{T}_z \mathbf{x}}{Z}$$


$(\alpha \mathbf{T}, \alpha \mathbf{Z})$  Give the same flow

$$\mathbf{y}' - \mathbf{y} = \frac{\mathbf{T}_y - \mathbf{T}_z \mathbf{y}}{Z}$$

# Observations


$$\mathbf{x}' - \mathbf{x} = \frac{\mathbf{T}_x - \mathbf{T}_z \mathbf{x}}{\mathbf{z}}$$

$$\mathbf{y}' - \mathbf{y} = \frac{\mathbf{T}_y - \mathbf{T}_z \mathbf{y}}{\mathbf{z}}$$



Focus of Expansion/Contraction

$$\left( \frac{\mathbf{T}_x}{\mathbf{T}_z}, \frac{\mathbf{T}_y}{\mathbf{T}_z} \right)$$



# Quasi-Parametric Motion Model

$$\mathbf{x}' - \mathbf{x} = \underbrace{\omega_y - \omega_z \mathbf{y} - \omega_x \mathbf{x} \mathbf{y} + \omega_y \mathbf{x}^2}_{\text{Global Parametric Component}} + \frac{\mathbf{T}_x - \mathbf{T}_z \mathbf{x}}{\mathbf{Z}}$$
$$\mathbf{y}' - \mathbf{y} = \underbrace{-\omega_x + \omega_z \mathbf{x} + \omega_y \mathbf{x} \mathbf{y} - \omega_x \mathbf{y}^2}_{\text{Global Parametric Component}} + \frac{\mathbf{T}_y - \mathbf{T}_z \mathbf{y}}{\mathbf{Z}}$$

The diagram shows two equations for motion. The first equation is  $\mathbf{x}' - \mathbf{x} = \omega_y - \omega_z \mathbf{y} - \omega_x \mathbf{x} \mathbf{y} + \omega_y \mathbf{x}^2 + \frac{\mathbf{T}_x - \mathbf{T}_z \mathbf{x}}{\mathbf{Z}}$ . A bracket under the first four terms is labeled "Global Parametric Component". A bracket under the last term is labeled "Local Parametric Component". The second equation is  $\mathbf{y}' - \mathbf{y} = -\omega_x + \omega_z \mathbf{x} + \omega_y \mathbf{x} \mathbf{y} - \omega_x \mathbf{y}^2 + \frac{\mathbf{T}_y - \mathbf{T}_z \mathbf{y}}{\mathbf{Z}}$ . A bracket under the first four terms is labeled "Global Parametric Component". A bracket under the last term is labeled "Local Parametric Component".

For rigid small motion scenario:

- 6 global model parameters
- $(\omega_x, \omega_y, \omega_z)$  Rotations       $(\mathbf{T}_x, \mathbf{T}_y, \mathbf{T}_z)$  Translations
- One local parameter per pixel : Depth Z

## Direct Model Estimation

$$E_{SSD}(\mathbf{u}) = \sum_{\mathbf{p} \in R} (\nabla I_1^T \delta \mathbf{u}(\mathbf{p}) + \delta I(\mathbf{p}))^2$$

Recall :

$$\delta \mathbf{u}(\mathbf{p}) = \mathbf{u}(\mathbf{p}) - \mathbf{u}^{(k)}(\mathbf{p})$$

$$\mathbf{u}(\mathbf{p}) = \begin{pmatrix} -xy & 1+x^2 & -y \\ -(1+y^2) & xy & x \end{pmatrix} \begin{pmatrix} \omega_x \\ \omega_y \\ \omega_z \end{pmatrix} + \frac{1}{Z} \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & y \end{pmatrix} \begin{pmatrix} T_x \\ T_y \\ T_z \end{pmatrix}$$

$$\mathbf{u}(\mathbf{p}) = \mathbf{A}\Omega + \tilde{\mathbf{Z}}\mathbf{K}\mathbf{T}$$

Observations: The error function is quadratic in inverse depth

Given the depth at each pixel, it is quadratic in rotation and translation parameters

Therefore, can be iteratively solved using depth and motion optimizations

# Solution

Simplifying assumption : Although depth varies from pixel-to-pixel we assume, that within a small window, (say 5x5) around each pixel, depth is constant

Local Optimization Step : Eliminate depth analytically

$$\mathbf{E}_{\text{local}}(\mathbf{u}) = \sum_{\mathbf{p} \in \mathbf{R}_{\text{local}}} (\nabla \mathbf{I}_1^T \delta \mathbf{u}(\mathbf{p}) + \delta \mathbf{I}(\mathbf{p}))^2$$

$$\frac{\partial \mathbf{E}_{\text{local}}}{\partial \tilde{\mathbf{Z}}} = 0 \quad \frac{\partial \mathbf{E}_{\text{local}}}{\partial \tilde{\mathbf{Z}}} = 2 \sum_{\mathbf{p} \in \mathbf{R}_{\text{local}}} \nabla \mathbf{I}_1^T \mathbf{K} \mathbf{T} (\nabla \mathbf{I}_1^T \delta \mathbf{u}(\mathbf{p}) + \delta \mathbf{I}(\mathbf{p}))$$

$$\delta \mathbf{u}(\mathbf{p}) = \mathbf{A}(\Omega - \Omega^{(k)}) + (\tilde{\mathbf{Z}} \mathbf{K} \mathbf{T} - \tilde{\mathbf{Z}}^{(k)} \mathbf{K} \mathbf{T}^{(k)})$$

$$\tilde{\mathbf{Z}}_{\text{opt}} = \frac{- \sum_{\mathbf{p} \in \mathbf{R}_{\text{local}}} \nabla \mathbf{I}_1^T \mathbf{K} \mathbf{T} (\nabla \mathbf{I}_1^T \mathbf{A}(\Omega - \Omega_0) - \nabla \mathbf{I}_1^T \tilde{\mathbf{Z}}^{(k)} \mathbf{K} \mathbf{T}^{(k)} + \delta \mathbf{I}(\mathbf{p}))}{\sum_{\mathbf{p} \in \mathbf{R}_{\text{local}}} (\nabla \mathbf{I}_1^T \mathbf{K} \mathbf{T})^2}$$

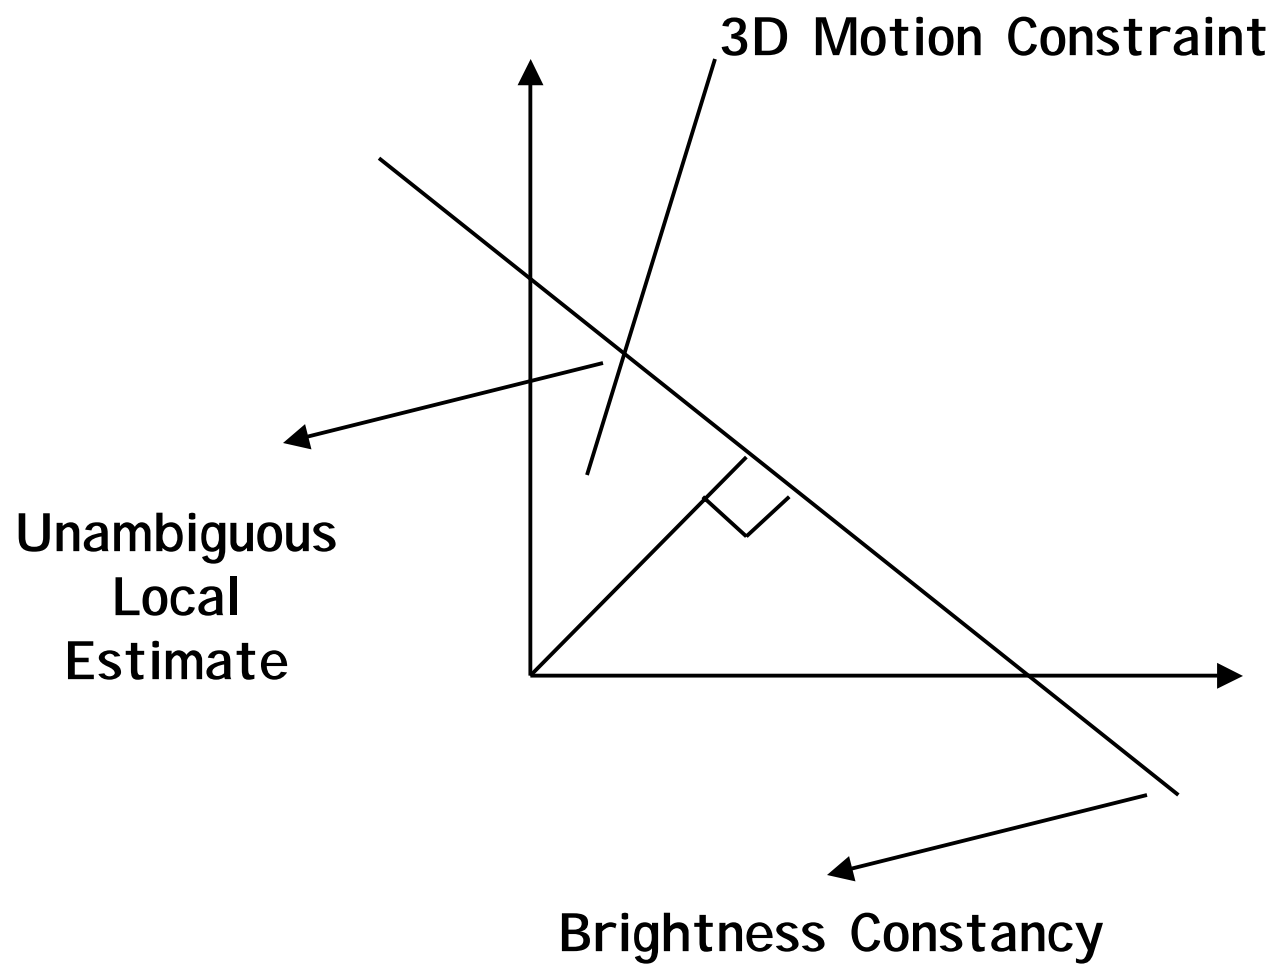
# Solution

Global Optimization Step : Substitute the analytical solution for Z in the optimization function and solve iteratively for rotation and translation

$$\mathbf{E}_{\text{global}}(\mathbf{u}) = \sum_{\mathbf{p} \in \mathbf{R}_{\text{global}}} (\nabla \mathbf{I}_1^T \mathbf{A}(\boldsymbol{\Omega} - \boldsymbol{\Omega}_0) + \nabla \mathbf{I}_1^T \tilde{\mathbf{Z}}_{\text{min}} \mathbf{K} \mathbf{T} - \nabla \mathbf{I}_1^T \tilde{\mathbf{Z}}^{(k)} \mathbf{K} \mathbf{T}^{(k)} + \delta \mathbf{I}(\mathbf{p}))^2$$

- This optimization function is quadratic in rotations
- So solve for rotations using closed form
- Solve for translations analytically
  
- Employ the older trick of warp and solve

# Global Constraints on Local Motion



# Layered Representation of 3D Geometry

- **Dense representation of geometry**
  - Planar layers capture sharp depth boundaries
  - Global matching constraints handle visibility
- **Alignment based estimation of structure**
  - Correspondences are not a pre-requisite for structure estimation
  - 3D constrained depth/parallax maps are sharper
- **Independent motion detection with 3D constraints**
  - Shape and motion constraints used to detect independent motions



# LOCAL RANGE ESTIMATION HANDLES ARBITRARY SHAPED OBJECTS



Two Views




Automatically computed  
Shape/Depth



Synthetic video rendered from  
original new views

# Multi-baseline depth estimation



Global matching method


# Depth vs. Parallax

- Euclidean depth estimation requires calibrated cameras:
  - External Calibration : Camera R & T for each position of the camera
  - Internal Calibration : Focal length, center, skew
- IBR can be accomplished usually without calibration:
  - Compute parallax instead of depth
  - A special practical version : plane + parallax

# Direct Estimation of Parallax/Depth

# Plane + Parallax 3D Model

[Sawhney '94, Kumar et al.'94, Shashua&Navab'94]



- Given a reference view, a reference plane, points in any other view are related by :

1. a plane projective transformation
2. the epipole, and
3. the “relative depth/parallax”

- That is,

$$p' \approx \underbrace{A_{\pi}}_{\text{Planar homography}} p + \underbrace{\kappa t'}_{\text{Relative depth}}$$

Planar homography

Relative depth

Epipole

$$\kappa = d_p / (Z d_{\pi})$$

## Plane+Parallax for Small Displacements

$$\mathbf{u}(\mathbf{p}) = \begin{pmatrix} -xy & 1+x^2 & -y \\ -(1+y^2) & xy & x \end{pmatrix} \begin{pmatrix} \omega_x \\ \omega_y \\ \omega_z \end{pmatrix} + \frac{1}{Z} \begin{pmatrix} 1 & 1 & 0 \\ 0 & 1 & y \end{pmatrix} \begin{pmatrix} T_x \\ T_y \\ T_z \end{pmatrix}$$

$$\mathbf{u}(\mathbf{p}) = \mathbf{A}\boldsymbol{\Omega} + \tilde{\mathbf{Z}}\mathbf{K}\mathbf{T}$$

Plane equation in image coordinates :  $\tilde{\mathbf{Z}} = \mathbf{a}\mathbf{x} + \mathbf{b}\mathbf{y} + \mathbf{c}$

$$\mathbf{u}_{\text{plane}}(\mathbf{p}) = \mathbf{A}\boldsymbol{\Omega} + (\mathbf{a}\mathbf{x} + \mathbf{b}\mathbf{y} + \mathbf{c})\mathbf{K}\mathbf{T}$$

$$\mathbf{u}(\mathbf{p}) = \mathbf{u}_{\text{plane}}(\mathbf{p}) + (\tilde{\mathbf{Z}} - \tilde{\mathbf{Z}}_{\text{plane}})\mathbf{K}\mathbf{T}$$

$$\mathbf{u}(\mathbf{p}) = \mathbf{u}_{\text{plane}}(\mathbf{p}) - \frac{\delta Z}{Z} \tilde{\mathbf{Z}}_{\text{plane}} \mathbf{K}\mathbf{T} = \mathbf{u}_{\text{plane}}(\mathbf{p}) - \frac{H}{d_\pi} \tilde{\mathbf{Z}}_{\text{plane}} \mathbf{K}\mathbf{T}$$

# Shape Estimation through 3D Constrained Image Alignment

- **Parameterize** the image correlation constraint with plane+parallax:

$$I_2(\mathbf{p}) = I_1(\mathbf{p} - \mathbf{u}(\mathbf{p}; \mathbf{A}_\pi, \mathbf{t}', \kappa(\mathbf{p})))$$

- Given an estimate of  $\mathbf{A}_\pi^{(m)}, \mathbf{t}'^{(m)}, \kappa^{(m)}(\mathbf{p})$

$I_1$  is **warped** towards  $I_2$  using the plane+parallax model:

$$I^w(\mathbf{p}) = I_1(\mathbf{p} - \mathbf{u}^{(m)}(\mathbf{p}))$$


- With  $\delta\mathbf{u}(\mathbf{p}) = \mathbf{u}(\mathbf{p}) - \mathbf{u}^{(m)}(\mathbf{p})$

compute the **increment** in the unknown parameters:


$$\min_{\mathbf{A}_\pi, \mathbf{T}', \kappa(\mathbf{p})} \sum_{\mathbf{p}} \rho(I(\mathbf{p}) - I^w(\mathbf{p}) + \nabla I^T(\mathbf{p})\delta\mathbf{u}(\mathbf{p}))$$

$\rho$  is a robust error function as before.


# Plane+parallax Demo




Input Sequence



Plane Alignment



Shape



Parallax Alignment



A Novel Real-world Application  
of IBR

IMAX 3D Movie Synthesis

**The End**