

Image-Based Rendering and Modeling

- 1 **Image-based rendering (IBR):** A scene is represented as a collection of images
- 1 **3D model-based rendering (MBR):** A scene is represented by a 3D model plus texture maps
- 1 Differences
 - u Many scene details need not be explicitly modeled in IBR
 - u IBR simplifies model acquisition process
 - u IBR processing speed independent of scene complexity
 - u 3D models (MBR) are more space efficient than storing many images (IBR)
 - u MBR uses conventional graphics “pipeline,” whereas IBR uses pixel reprojection
 - u IBR can sometimes use uncalibrated images, MBR cannot

IBR Approaches for View Synthesis

- 1 Non-physically based image mapping
 - u **Image morphing**
- 1 Geometrically-correct pixel reprojection
 - u **Image transfer** methods, e.g., in photogrammetry
- 1 Mosaics
 - u Combine two or more images into a single large image or higher resolution image
- 1 Interpolation from dense image samples
 - u Direct representation of **plenoptic function**

Image Metamorphosis (Morphing)

- 1 **Goal:** Synthesize a sequence of images that smoothly and realistically transforms objects in source image A into objects in destination image B

1 **Method 1: 3D Volume Morphing**

- Create 3D model of each object
- Transform one 3D object into another
- Render synthesized 3D object
- Hard/expensive to accurately model real 3D objects
- Expensive to accurately render surfaces such as skin, feathers, fur



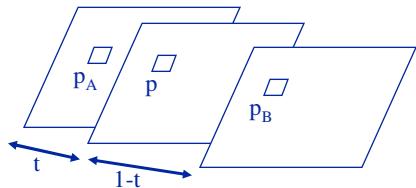
[Lerios, Garfinkle, & Levoy 95]

Image Morphing

1 Method 2: Image Cross-Dissolving

- ✉ Pixel-by-pixel color interpolation
- ✉ Each pixel p at time $t \in [0, 1]$ is computed by combining a fraction of each pixel's color at the same coordinates in images A and B:

$$p = (1 - t)p_A + tp_B$$

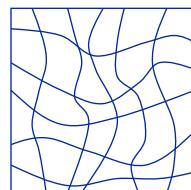
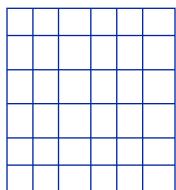


- ✉ Easy, but looks artificial, non-physical

Image Morphing

1 Method 3: Mesh-based image morphing

- ✉ G. Wolberg, *Digital Image Warping*, 1990
- ✉ Warp between corresponding grid points in source and destination images
- ✉ Interpolate between grid points, e.g., linearly using three closest grid points



- ✉ Fast, but hard to control so as to avoid unwanted distortions

Image Warping

- 1 Goal: Rearrange pixels in an image. I.e., map pixels in source image A to new coordinates in destination image B
- 1 Applications
 - u Geometric Correction (e.g., due to lens pincushion or barrel distortion)
 - u Texture mapping
 - u View synthesis
 - u Mosaics
- 1 Aka **geometric transformation, geometric correction, image distortion**
- 1 Some simple mappings: 2D translation, rotation, scale, affine, projective

Image Warping

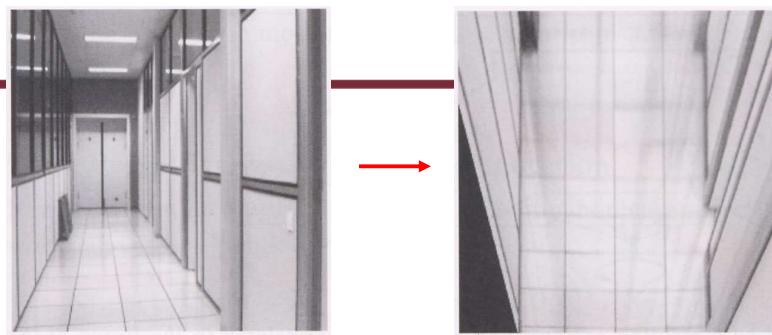


image plane in front

image plane below

black area
where no pixel
maps to

Homographies

1 Perspective projection of a plane

- u Lots of names for this:
 - u **homography**, texture-map, colineation, planar projective map
- u Modeled as a 2D warp using homogeneous coordinates

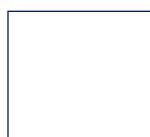
$$\begin{bmatrix} sx' \\ sy' \\ s \end{bmatrix} = \begin{bmatrix} * & * & * \\ * & * & * \\ * & * & * \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

$\mathbf{p}' \quad \mathbf{H} \quad \mathbf{p}$

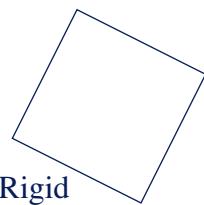
To apply a homography \mathbf{H}

- Compute $\mathbf{p}' = \mathbf{H}\mathbf{p}$ (regular matrix multiply)
- Convert \mathbf{p}' from homogeneous to image coordinates
 - divide by s (third) coordinate

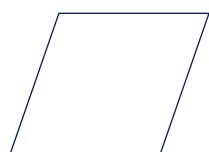
Examples of 2D Transformations



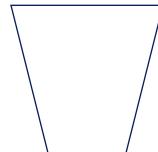
Original



Rigid



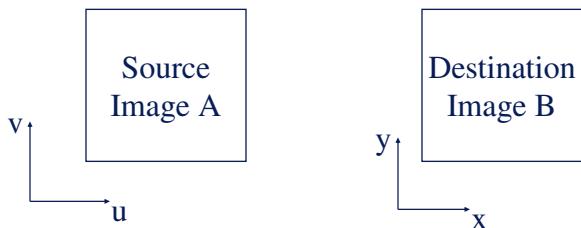
Affine



Projective

Mapping Techniques

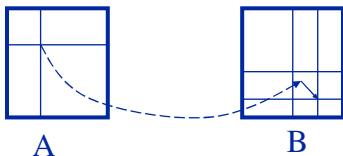
- 1 Define transformation as either
 - u **Forward:** $x = X(u, v)$, $y = Y(u, v)$
 - u **Backward:** $u = U(x, y)$, $v = V(x, y)$



Mapping Techniques

1 Forward, point-based

- u Apply forward mapping X, Y at point (u, v) to obtain real-valued point (x, y)
- u Assign (u, v) 's gray level to pixel closest to (x, y)

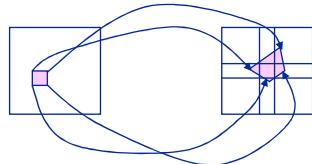


- u Problem: “**measles**,” i.e., “**holes**” (pixel in destination image that is not assigned a gray level) and “**folds**” (pixel in destination image is assigned multiple gray levels)
- u Example: Rotation, since preserving length cannot preserve number of pixels

Mapping Techniques

1 Forward, square-pixel based

- u Consider pixel at (u,v) as a unit square in source image. Map square to a quadrilateral in destination image
- u Assign (u,v) 's gray level to pixels that the quadrilateral overlaps



- u Integrate source pixels' contributions to each output pixel. Destination pixel's gray level is weighted sum of intersecting source pixels' gray levels, where weight proportional to coverage of destination pixel
- u Avoids holes, but not folds, and requires intersection test

Mapping Techniques

1 Backward, point-based

- u For each destination pixel at coordinates (x,y) , apply backward mapping, \mathbf{U} , \mathbf{V} , to determine real-valued source coordinates (u,v)
- u Interpolate gray level at (u,v) from neighboring pixels, and copy gray level to (x,y)



- u Interpolation may cause artifacts such as aliasing, blockiness, and false contours
- u Avoids holes and folds problems
- u Method of choice

Backward Mapping

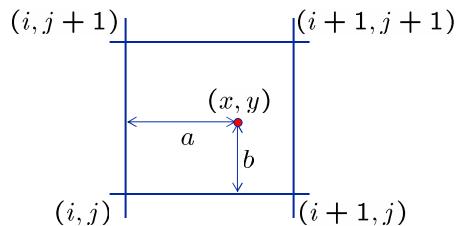
- 1 For $x = x_{\min}$ to x_{\max}
 - for $y = y_{\min}$ to y_{\max}
 - $u = \mathbf{U}(x, y)$
 - $v = \mathbf{V}(x, y)$
 - $B[x, y] = A[u, v]$
- 1 But (u, v) may not be at a pixel in A
- 1 (u, v) may be out of A 's domain
- 1 If \mathbf{U} and/or \mathbf{V} are discontinuous, A may not be connected!
- 1 Digital transformations in general don't commute

Pixel Interpolation

- 1 **Nearest-neighbor (0-order) interpolation**
 - u $g(x, y) = \text{gray level at nearest pixel}$ (i.e., round (x, y) to nearest integers)
 - u May introduce artifacts if image contains fine detail
- 1 **Bilinear (1st-order) interpolation**
 - u Given the 4 nearest neighbors, $g(0, 0), g(0, 1), g(1, 0), g(1, 1)$, of a desired point $g(x, y)$, $0 \leq x, y \leq 1$, compute gray level at $g(x, y)$:
 - u Interpolate linearly between $g(0,0)$ and $g(1,0)$ to obtain $g(x,0)$
 - u Interpolate linearly between $g(0,1)$ and $g(1,1)$ to obtain $g(x,1)$
 - u Interpolate linearly between $g(x,0)$ and $g(x,1)$ to obtain $g(x,y)$
 - u Combining all three interpolation steps into one we get:
 - u
$$g(x,y) = (1-x)(1-y) g(0,0) + (1-x)y g(0,1) + x(1-y) g(1,0) + xy g(1,1)$$
- 1 **Bicubic spline interpolation**

Bilinear Interpolation

- ① A simple method for resampling images



$$f(x, y) = \begin{aligned} & (1 - a)(1 - b) f[i, j] \\ & + a(1 - b) f[i + 1, j] \\ & + ab f[i + 1, j + 1] \\ & + (1 - a)b f[i, j + 1] \end{aligned}$$

Example of Backward Mapping

- ① **Goal:** Define a transformation that performs a **scale change**, which expands size of image by 2, i.e., $\mathbf{U}(x) = x/2$

① $A = 0 \dots 0 2 2 2 0 \dots 0$

① 0-order interpolation, I.e., $u = \lfloor x/2 \rfloor$

$B = 0 \dots 0 2 2 2 2 2 2 0 \dots 0$

① Bilinear interpolation, I.e., $u = x/2$ and average 2 nearest pixels if u is not at a pixel

$B = 0 \dots 0 1 2 2 2 2 2 1 0 \dots 0$

Image Morphing

1 Method 4: Feature-based image morphing

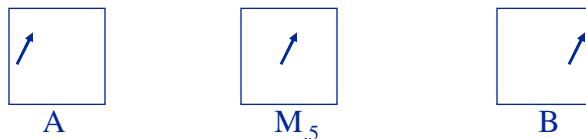
- ✉ T. Beier and S. Neely, *Proc. SIGGRAPH '92*
- ✉ Distort color *and* shape
 - ⇒ image warping + cross-dissolving
- ✉ Warping transformation partially defined by user interactively specifying corresponding pairs of line segment features in the source and destination images; only a sparse set is required (but carefully chosen)
- ✉ Compute dense pixel correspondences, defining continuous mapping function, based on weighted combination of displacement vectors of a pixel from all of the line segments
- ✉ Interpolate pixel positions and colors (2D linear interpolation)

Beier and Neely Algorithm

- 1 **Given:** 2 images, A and B, and their corresponding sets of line segments, L_A and L_B , respectively
- 1 **For**each intermediate frame time $t \in [0, 1]$ **do**
 - ✉ **Linearly interpolate** the position of each line
 - ✉ $L_t[i] = \text{Interpolate}(L_A[i], L_B[i], t)$
 - ✉ **Warp** image A to destination shape
 - ✉ $WA = \text{Warp}(A, L_A, L_t)$
 - ✉ **Warp** image B to destination shape
 - ✉ $WB = \text{Warp}(B, L_B, L_t)$
 - ✉ **Cross-dissolve** by fraction t
 - ✉ $\text{MorphImage} = \text{CrossDissolve}(WA, WB, t)$

Example: Translation

- 1 Consider images where there is one line segment pair, and it is **translated** from image A to image B:



- 1 First, linearly interpolate position of line segment in M
- 1 Second, for each pixel (x, y) in M, find corresponding pixels in A (x-a, y) and B (x+a, y), and average them

Feature-based Warping

- 1 **Goal:** Define a continuous function that warps a source image to a destination image from a sparse set of corresponding, oriented, **line segment features** - each pixel's position defined relative to these line segments

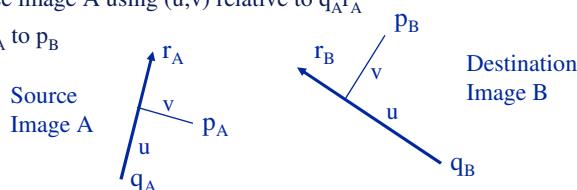
- 1 Warping with one line pair:

```
foreach pixel pB in destination image B do
```

```
    find dimensionless coordinates (u,v) relative to oriented line segment qBrB
```

```
    find pA in source image A using (u,v) relative to qArA
```

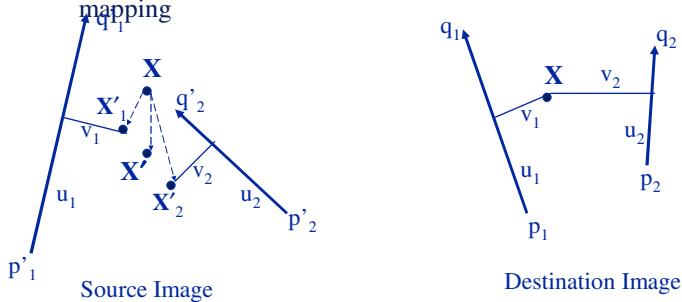
```
    copy color at pA to pB
```



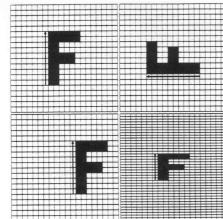
Feature-based Warping (cont.)

1 Warping with multiple line pairs

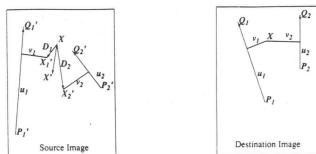
- Use a weighted combination of the points defined by the same mapping



$\mathbf{X}' = \text{weighted average of } D_1 \text{ and } D_2, \text{ where } D_i = \mathbf{X}'_i - \mathbf{X},$
and weight = $(\text{length}(p_i q_i))^c / (a + |v_i|)^b$, for constants a, b, c



- Warping with multiple line pairs:
Use a weighted combination of the points defined by the same mapping



Geometrically-Correct Pixel Reprojection

- 1 What geometric information is needed to generate virtual camera views?
 - u Dense pixel correspondences between two input views
 - u Known geometric relationship between the two cameras
 - u Epipolar geometry

View Interpolation from Range Maps

- 1 Chen and Williams, *Proc. SIGGRAPH '93* (seminal paper on image-based rendering)
- 1 **Given:** Static 3D scene with Lambertian surfaces, and two images of that scene, each with **known camera pose** and **range map**
- 1 **Algorithm:**
 1. Recover dense pixel correspondence using known camera calibration and range maps
 2. Compute forward mapping, X_F , Y_F , and backward mapping, X_B , Y_B . Each “morph map” defines an *offset vector* for each pixel

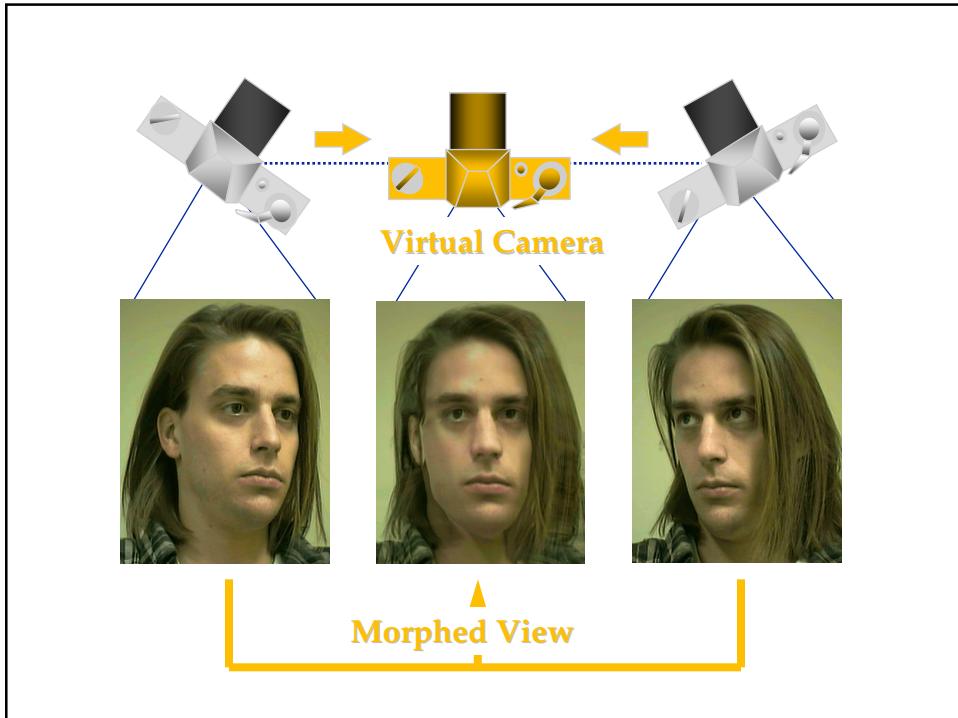
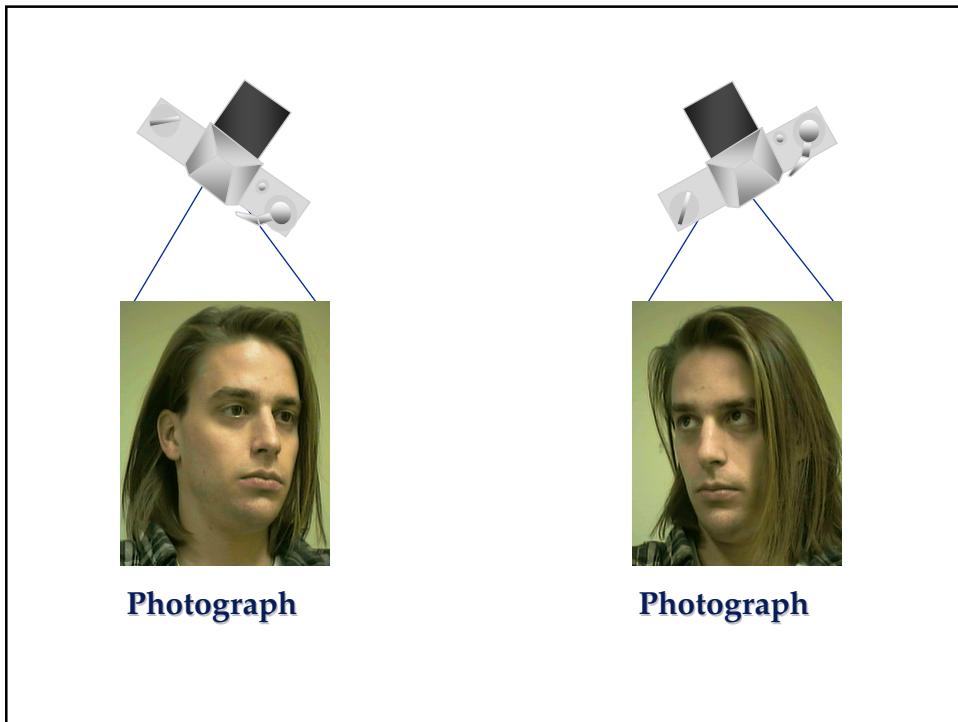
View Interpolation (cont.)

3. Compute interpolated “morph map” by linearly interpolating forward and backward offset vectors, given intermediate frame, t , $0 \leq t \leq 1$
4. Apply forward mapping from A given interpolated morph map (approximating perspective transformation for new view)
5. Use a z-buffer (and A’s range map) to keep only closest pixel (so, handles folds)
6. For each pixel in interpolated image that has no color, interpolate color from all adjacent colored pixels

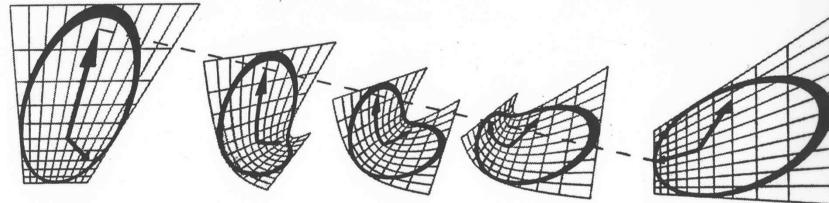
View Morphing

¹ Seitz and Dyer, *Proc. SIGGRAPH ‘95*

¹ **Given:** Two views of an unknown rigid scene, with no camera information known, compute new views from a virtual camera at viewpoints in-between two input views



Why NOT Image Morphing?



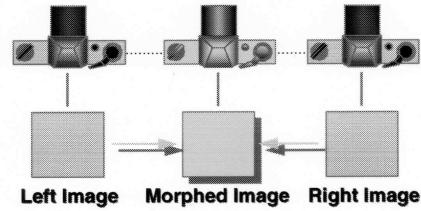
- Results are not physically consistent

11

When is View Synthesis Feasible?

- ① Given two images, I_0 and I_1 , with optical centers C_0 and C_1 , if the monotonicity constraint holds for C_0 and C_1 , then there is sufficient information to completely predict the appearance of the scene from all in-between viewpoints along line segment C_0C_1
- ① **Monotonicity Constraint:** All visible scene points appear in the *same order* along conjugate epipolar line in I_0 and I_1
- ① Any number of distinct scenes could produce I_0 and I_1 , but each one produces the **same** in-between images

View Morphing: Parallel Views



Morphing parallel views \Rightarrow new parallel views

Parallel projection matrices have special form:

$$\boldsymbol{\Pi}_0 = \begin{bmatrix} 1 & 0 & 0 & C_0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \quad \boldsymbol{\Pi}_1 = \begin{bmatrix} 1 & 0 & 0 & C_1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

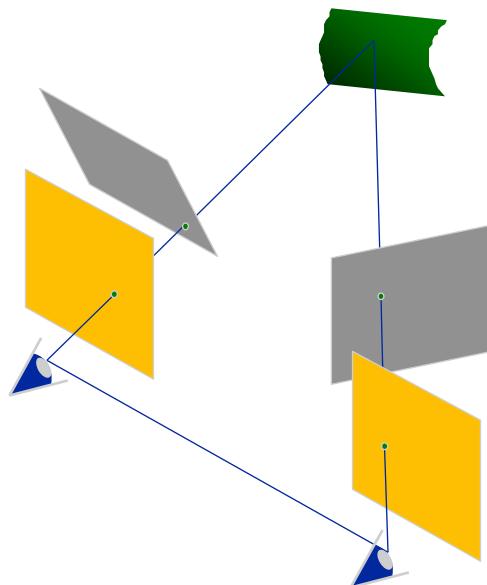
Linear combination of image point positions: $\mathbf{p}_0 + \mathbf{p}_1$

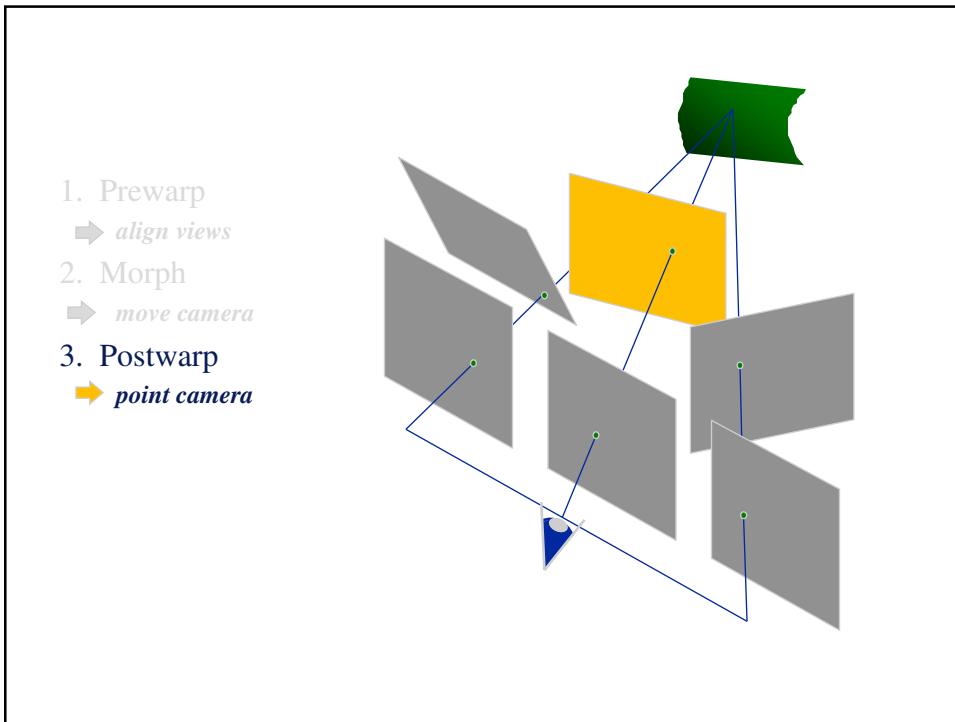
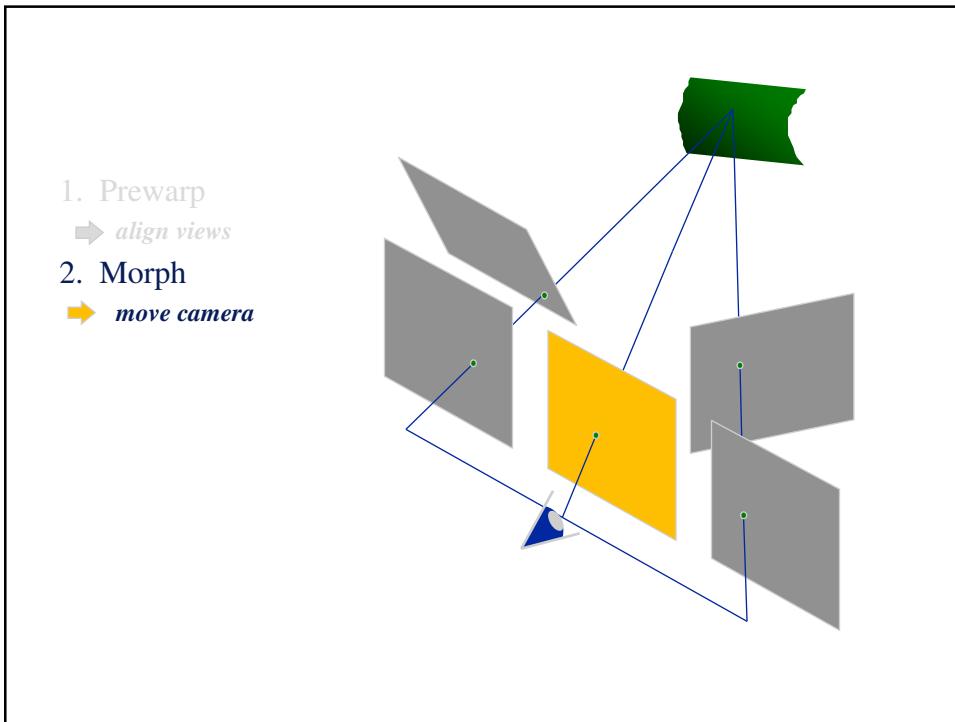
\iff

Linear combination of camera positions: $\boldsymbol{\Pi}_0 + \boldsymbol{\Pi}_1$

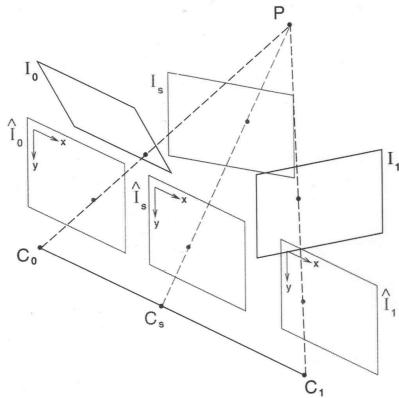
1. Prewarp

align views





Interpolating Non-Parallel Views



Three Step Algorithm

1. **Prewarp:** Reproject images to parallel configuration
2. **Morph:** Interpolate parallel images
3. **Postwarp:** Reproject images to new configuration

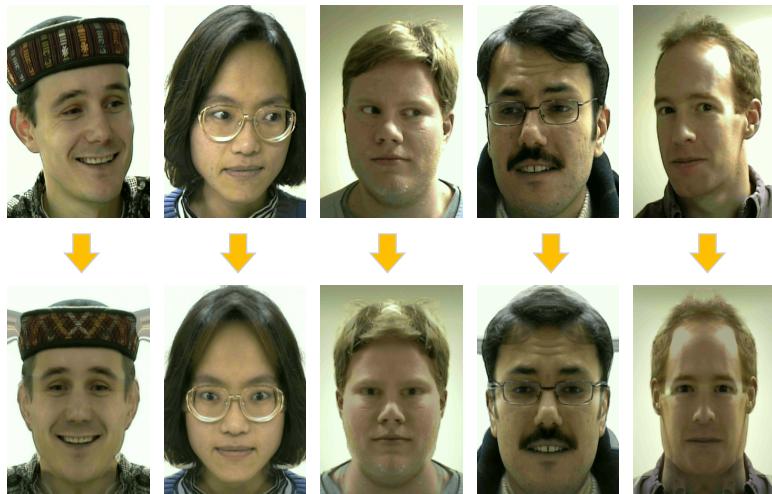
Features of View Morphing

- ① Provides
 - A mobile virtual camera
 - Better image morphs
- ① Requires no prior knowledge
 - No 3D shape information
 - No camera information
 - No training on other images

Application: Photo Correction

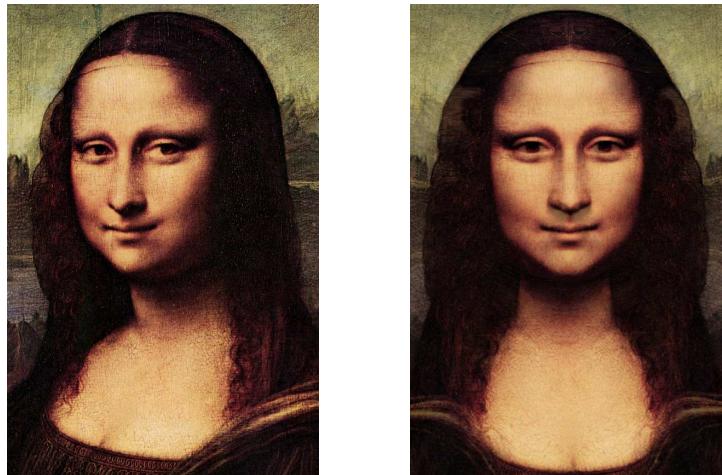
- 1 Image Postprocessing
 - ✉ Alter image perspective in the lab
- 1 Image Databases
 - ✉ Normalize images for better indexing
 - ✉ Simplify face recognition tasks

Original Photographs



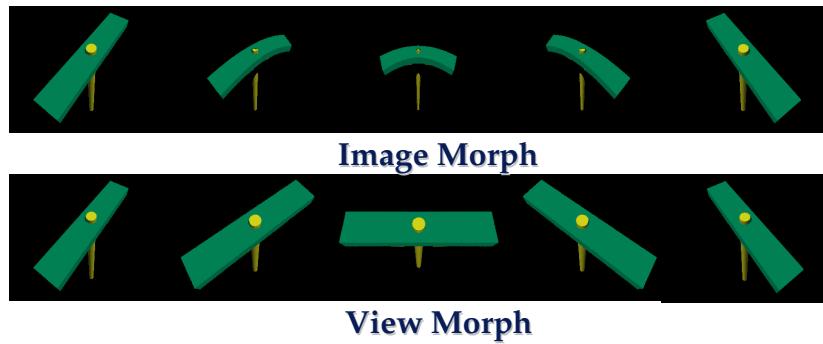
Corrected Photographs

Another Example



Application: Better Image Transitions

- ✉ Avoid bending and shearing *distortions*
- ✉ Shapes do not need to be *aligned*



Mosaics

- ① **Goal:** Given a static scene and a set of images (or video) of it, combine the images into a single “**panoramic image**” or “**panoramic mosaic**”
- ① **Motivation:** Image-based modeling of 3D scenes benefits visualization, navigation, exploration, VR walkthroughs, video compression, video stabilization, super-resolution
- ① **Example:** Apple’s Quicktime VR (Chen, *SIGGRAPH ‘95*)

Image Mosaics



① Goal

- ② Stitch together several images into a seamless composite

Mosaicing Method

- ① **Registration:** Given n input images, I_1, \dots, I_n , compute an image-to-image transformation that will map each image I_2, \dots, I_n into the coordinate frame of reference image, I_1
- ① **Warp:** Warp each image I_i , $i=2, \dots, n$, using transform
- ① **Interpolate:** Resample warped image
- ① **Composite:** Blend images together to create single output image based on the reference image's coordinate frame

When can Two Images be Aligned?

- ① Problems
 - In general, warping function depends on the **depth** of the corresponding scene point since image projection defined by $x' = fx/z$, $y' = fy/z$
 - Different views means, in general, that parts that are visible in one image may be occluded in the other
- ① Special cases where the above problems *can't* occur
 - **Panoramic mosaic:** Camera rotates (i.e., pans) about its optical center, arbitrary 3D scene
 - No motion parallax as camera rotates, so depth unimportant
 - 2D projective transformation relates any 2 images ($\Rightarrow 8$ unknowns)
 - **Planar mosaic:** Arbitrary camera views of a planar scene
 - 2D projective transformation relates any 2 images

PANORAMIC MOSAICING

- ANY 2 IMAGES OF AN ARBITRARY SCENE TAKEN FROM 2 CAMERAS WITH THE SAME CAMERA CENTER ARE RELATED BY THE PLANAR PROJECTIVE TRANSFORMATION:

$$\tilde{w}' = K R K^{-1} \tilde{w}$$

called the
"homography
induced by
the plane at
infinity"

where \tilde{w}, \tilde{w}' are homogeneous coordinates of 2 corresponding points in the two input images

K is the 3×3 upper-triangular camera calibration matrix:

$$\begin{bmatrix} \alpha_u & 0 & u_0 \\ 0 & \alpha_v & v_0 \\ 0 & 0 & 1 \end{bmatrix}$$

R is the 3×3 rotation matrix:

$$\begin{bmatrix} r_{11} & r_{12} & r_{13} \\ r_{21} & r_{22} & r_{23} \\ r_{31} & r_{32} & r_{33} \end{bmatrix}$$

4 dof + 3 dof \Rightarrow 4 point correspondences needed

Or, generalizing to the projective camera:

$$\tilde{w}' = \begin{bmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{bmatrix} \tilde{w}$$

8 degrees of freedom

$$m_{33} = 1$$

4 point correspondences needed

How to Determine Alignment Assuming a Planar Projective Transformation between Cameras?

- 1 Method 1: Find 4 point correspondences, and then solve for 8 unknowns
 - u Requires reliable detection of four corresponding features, at sub-pixel location accuracy

- 1 Method 2: Use image-based (intensity) correlation to determine best matching transformation
 - u No correspondences needed
 - u Statistically optimal (gives maximum likelihood estimate)
 - u Useful for local image registration

Example: 2D Rigid Warp Mosaics

- 1 **Assume:** Planar scene, camera motion restricted to plane parallel to scene, optical axis perpendicular to scene, intensity constancy assumption, local displacement
- 1 **3 unknowns:** 2D translation (a, b) and 2D rotation (θ)
- 1 Relation between two images, I and I' , given by:
$$I'(x', y') = I(x \cos \theta - y \sin \theta + a, x \sin \theta + y \cos \theta + b)$$
Expanding $\sin \theta$ and $\cos \theta$ to first two terms in Taylor series:
$$I'(x', y') \approx I(x + a - y\theta - x\theta^2/2, y + b + x\theta - y\theta^2/2)$$
Expanding I to first term of its Taylor series expansion:
$$I'(x', y') \approx I(x, y) + (a - y\theta - x\theta^2/2) \partial I / \partial x + (b + x\theta - y\theta^2/2) \partial I / \partial y$$

Example: 2D Rigid Warp Mosaics

- 1 Solve for a, b, θ that minimizes SSD error:

$$\begin{aligned} E &= \sum \sum (I' - I)^2 \\ &= \sum_x \sum_y (I(x, y) + (a - y\theta + x\theta^2 / 2) \partial I / \partial x + (b + x\theta - y\theta^2 / 2) \partial I / \partial y - I(x, y))^2 \end{aligned}$$

- 1 Assuming small displacement, use gradient descent to minimize E , $\nabla E = (\partial E / \partial a, \partial E / \partial b, \partial E / \partial \theta)$

- 1 Iteratively update total motion estimate, (a, b, θ) , while warping I' towards I until $E <$ threshold

2D Rigid Warp Algorithm

- 1 Because the displacement between images may not be small enough to solve directly for the motion parameters, use an iterative algorithm instead:
 1. $a^{(0)} = 0; b^{(0)} = 0; \theta^{(0)} = 0; \mathbf{m} = (0, 0, 0); t = 1$
 2. Solve for $a^{(t)}, b^{(t)}, \theta^{(t)}$ from the 3 equations
 3. Update the total motion estimate:
$$\mathbf{m}^{(t+1)} = \mathbf{m}^{(t)} + (a^{(t)}, b^{(t)}, \theta^{(t)})$$
 4. Warp I' toward I : $I' = \text{warp}(I', a^{(t)}, b^{(t)}, \theta^{(t)})$
 5. If $(a^{(t)} < \varepsilon_1 \text{ and } b^{(t)} < \varepsilon_2 \text{ and } \theta^{(t)} < \varepsilon_3)$
then halt
else $t = t + 1$; goto step 2

More generally, Szeliski paper

$$\text{minimizes } E = \sum_i [I'(x'_i, y'_i) - I(x_i, y_i)]^2 \\ = \sum_i e_i^2$$

by computing $\frac{\partial e_i}{\partial m_0}, \dots, \frac{\partial e_i}{\partial m_7}$

$$\text{E.g., } \frac{\partial e_i}{\partial m_0} = \frac{x_i}{m_0 x_i + m_1 y_i + 1} \cdot \frac{\partial I'}{\partial x'}$$

Then uses Levenberg-Marquardt iterative, nonlinear minimization algorithm to solve for m_0, \dots, m_7 .

Dealing with Noisy Data: RANSAC

- 1 How to find the best fitting data to a global model when the data are noisy – especially because of the presence of outliers, i.e., missing features and extraneous features?

- 1 **RANSAC** (Random Sample Consensus) Method
 - u Iteratively select a small subset of data and fit model to that data
 - u Then check all of data to see how many fit the model

RANSAC Algorithm for Robust Estimation

```
bestcnt := 0
until there is a good fit or  $k$  iterations do
    randomly choose a sample of  $n$  points from the dataset
    compute the best fitting model parameters to the selected subset of the  $n$ 
        data points
    cnt := 0
    foreach remaining data point do
        if the distance from the current point to the model is < T then cnt++
        if cnt  $\geq D$  then there is a good fit, so re-compute the best fitting model
            parameters using all of the good points, and halt
        else if cnt > bestcnt then bestcnt := cnt
```

Other Robust Parameter Estimation Methods

- 1 How to deal with outliers, i.e., occasional large-scale measurement errors?

1 M-Estimators

$$\arg \min_{\theta} \sum_{x_i \in X} \rho(r_i(x_i, \theta); \sigma)$$

- u Generalization of
 - u MLE (Maximum Likelihood Estimation)
 - u LMedS (Least Median of Squares)
- u Θ is the set of model parameters, ρ is a “robust loss function” whose shape is parameterized by σ , and $r_i(x_i, \theta)$ is the residual error of the model θ with the i^{th} data element

Least Median of Squares (LMedS)

- 1 Defined as

$$\arg \min_{\theta} \text{median}_{x_i \in X} r^2(x_i, \theta)$$

- 1 Up to $\frac{1}{2}$ of data points can be arbitrarily far from the optimum estimate without changing the result
- 1 Median is not differentiable, so how to search for optimum?
 - u Use randomization:
 - u Randomly select a subset of the data
 - u Compute model, θ , from the selected subset
 - u Each candidate model, θ , is tested by computing r^2 using the remaining data points and selecting the median of these values

Global Image Registration

- 1 When images are not sufficiently close, must do global registration
- 1 Method 1: **Coarse-to-fine matching using Gaussian Pyramid**
 1. Compute Gaussian pyramids
 2. $level = N$ // Set initial processing level to coarse level
 3. Solve for motion parameters at level $level$
 4. **If** $level = 0$ **then halt**
 5. Interpolate motion parameters at level $level - 1$
 6. $level = level - 1$
 7. **Goto** step 3

Global Image Registration

1 Method 2: Coarse-to-fine matching using Laplacian Pyramid

- u Use Laplacian pyramids, LA and LB
- u **for** $l = N$ **to** 0 **step -1 do**
 - Use motion estimate from previous iteration to warp level l in LA
 - for** $-1 \leq i, j \leq 1$ **do**
 - compute cross-correlation at level l :
$$CC_{i,j}(x, y) = LA_l(x, y) LB_l(x+i, y+j)$$
 - smooth using Gaussian pyramid to level S :
$$C_{i,j}(x, y) = CC_{i,j}(x, y) * w(x, y)$$
 - foreach** (x, y) interpolate 3×3 correlation surface centered at (x, y) at level S to find peak, corresponding to best local motion estimate
 - find best-fit global motion model to flow field

Planar Mosaics and Panoramic Mosaics

-
- 1 Motion model is 2D projective transformation, so 8 parameters
 - 1 Assuming small displacement, minimize SSD error
 - 1 Apply (nonlinear) minimization algorithm to solve

Panoramic Mosaics

- ① Large field of view \Rightarrow can't map all images onto a plane
- ① Instead, map onto cylinder, sphere, or cube
- ① Example: With a cylinder, first warp all images from rectilinear to cylindrical coordinates, then mosaic them
- ① “Undistort” (perspective correct) image from this representation prior to viewing

Cylindrical panoramas



- ① Steps
 - ② Reproject each image onto a cylinder
 - ② Blend
 - ② Output the resulting mosaic

Cylindrical reprojection



Image 384x300

f = 180 (pixels)

f = 280

f = 380

- Map image to cylindrical coordinates

need to know the camera focal length

Cylindrical image stitching



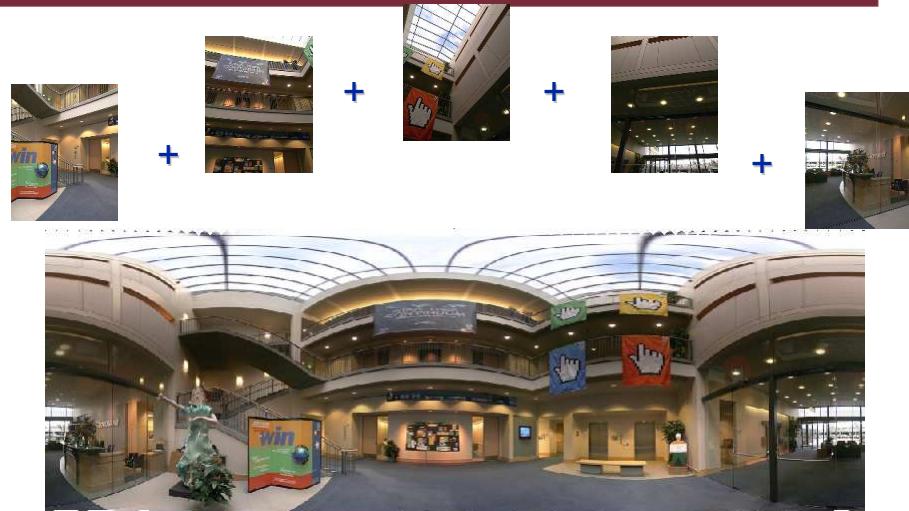
- What if you don't know the camera rotation?

Solve for the camera rotations

Note that a rotation of the camera is a **translation** of the cylinder!

Use Lucas-Kanade to solve for translations of cylindrically-warped images

Full-view Panorama

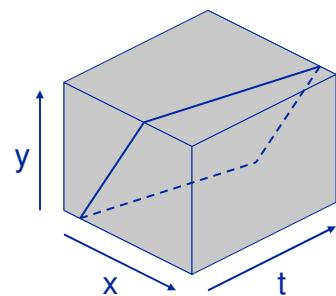
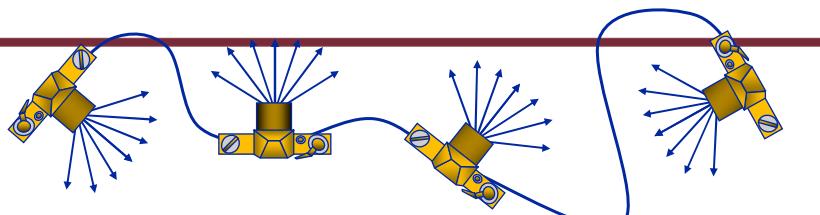


Other types of mosaics



- ¹ Can mosaic onto *any* surface if you know the geometry
- ^u See NASA's [Visible Earth project](#) for some stunning earth mosaics
 - ^u <http://earthobservatory.nasa.gov/Newsroom/BlueMarble/>

Path Images



Pushbroom image



<http://grail.cs.washington.edu/projects/stereo/>

Cyclograph



<http://grail.cs.washington.edu/projects/stereo/>

Multi-Perspective Image

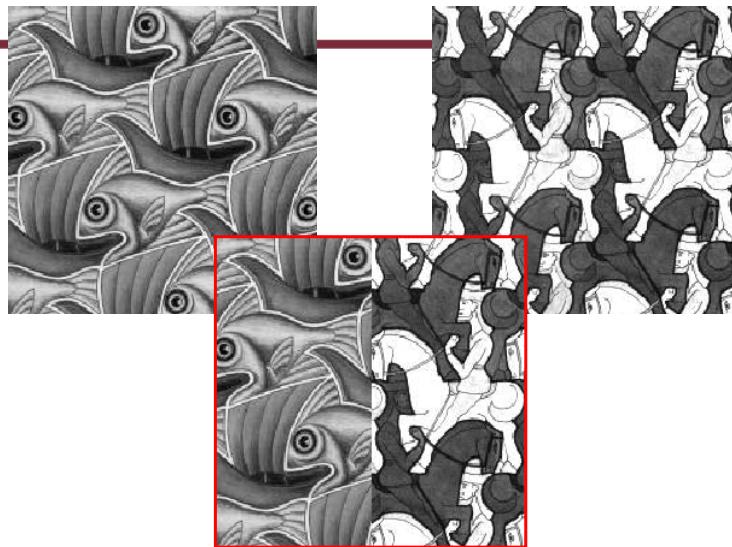


Rademacher and Bishop, 1998

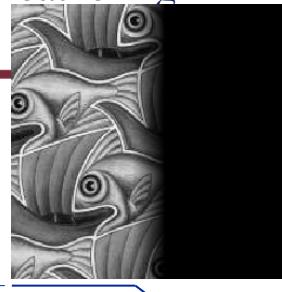
Some Blending Methods

- ① Average intensity of overlapping pixels
- ① Median intensity of overlapping pixels
- ① Newest (i.e., most recent frame) pixel's intensity
- ① Burt's pyramid blending method
- ① Bilinear blending
 - ② Weighted average with pixels near center of each image contributing more
 - ② Let w_t be a 1D triangle (hat) function of size equal to width of image, with value 0 at edges and value 1 at midpoint. Then use 2D weighting function: $w(x'_i, y'_i) = w_t(x'_i)w_t(y'_i)$

Image Blending



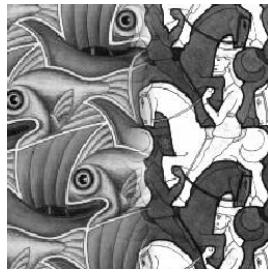
Feathering



+



=



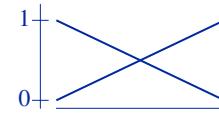
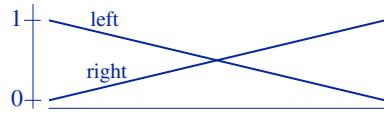
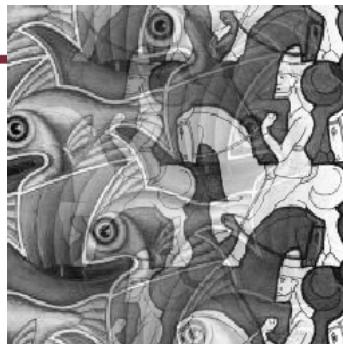
Encoding transparency

$$I(x,y) = (\alpha R, \alpha G, \alpha B, \alpha)$$

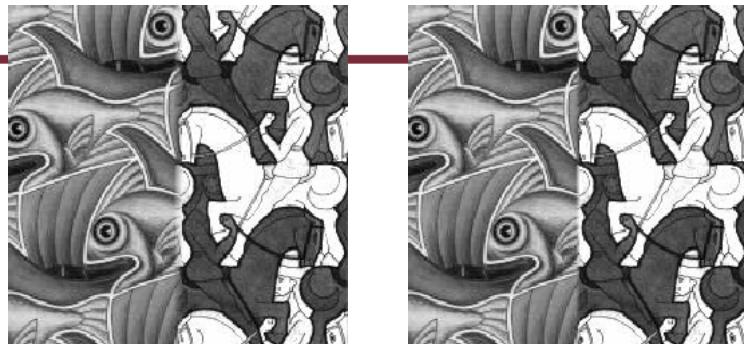
$$I_{\text{blend}} = I_{\text{left}} + I_{\text{right}}$$

(J. Blinn, CGA, 1994) for details

Effect of window size



Effect of window size



Good window size



- ① “Optimal” window: smooth but not ghosted
- ② Doesn’t always work...

Mosaicing Arbitrary Scenes

- ① For general scenes and arbitrary viewpoints, we must recover **depth**
- ① **Method 1:** Depth map given ala Chen and Williams
- ① **Method 2:** Segment image into piecewise-planar patches and use 2D projective method for each patch
- ① **Method 3:** Recover dense 3D depth map using either **stereo reconstruction** (assuming known motion between cameras) or **structure from motion** (assuming no camera motion is known)

Learning-based View Synthesis from 1 View

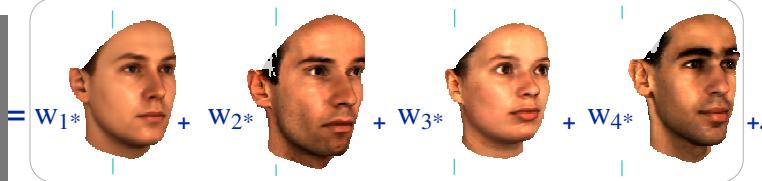
- ① Vetter and Poggio, *IEEE Trans. PAMI*, 1997
- ① **Given:** A set of views of several training objects, and a single view of a new object
- ① Single view of new object considered to be approximated by a linear combination of views of the training objects, all at the same pose as the new object
- ① Learn set of weights that specify best match between given view of new object and same view of the training objects
- ① Use learned weights with other views of training objects to synthesize novel view of new object

Approach: Example-based Modeling of Faces

2D Image



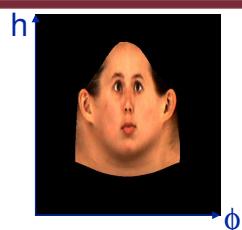
3D Face Models



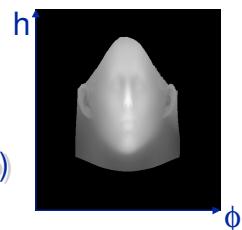
Cylindrical Coordinates



red(h, ϕ)
green(h, ϕ)
blue(h, ϕ)

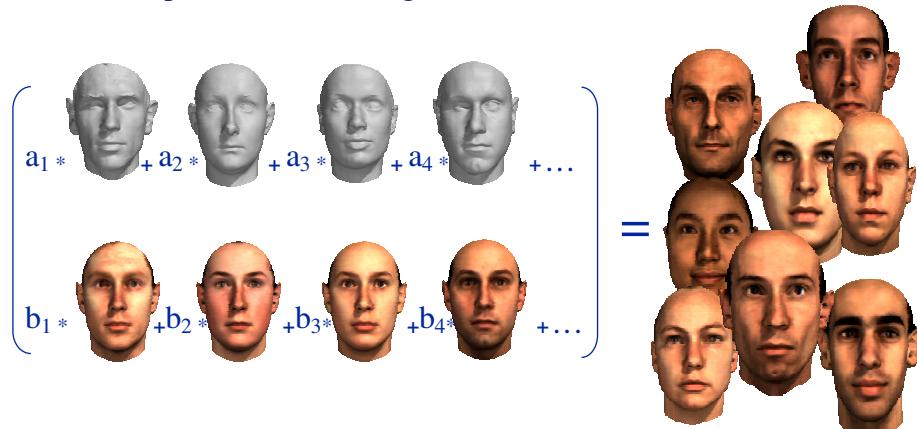


radius(h, ϕ)

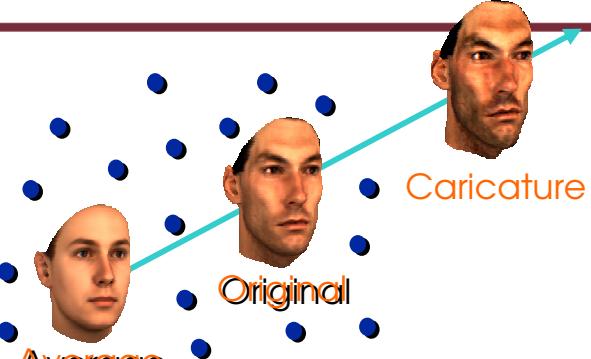


Vector Space of 3D Faces

① A Morphable Model can generate new faces



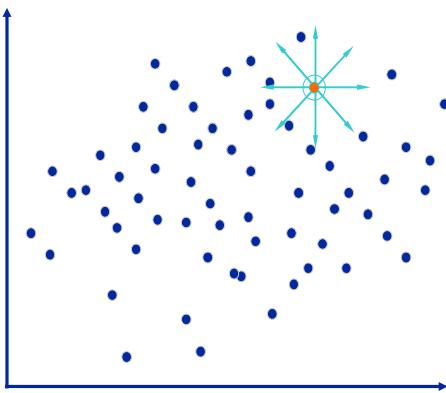
Modeling in Face Space



Modeling the Appearance of Faces

A face is represented as a point in face space

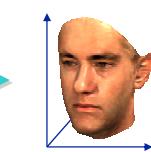
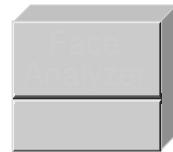
1 Which directions code for specific attributes?



3D Shape from Images



Input Image



3D Head

Matching a Morphable 3D Face Model

$$\text{Input Image} = R \left(\begin{array}{c} a_1 * \text{Face 1} + a_2 * \text{Face 2} + a_3 * \text{Face 3} + a_4 * \text{Face 4} + \dots \\ b_1 * \text{Face 5} + b_2 * \text{Face 6} + b_3 * \text{Face 7} + b_4 * \text{Face 8} + \dots \end{array} \right)$$

Optimization problem!

Vetter's Method

Given an input image of a face

1. Compute correspondence with a "reference face" in the same pose

(Uses coarse-to-fine optical flow algorithm using a Laplacian pyramid)

Result is a "shape vector"
 $s = (x_1, y_1, \dots, x_n, y_n)$

2. Compute "texture vector"
Difference map between corresponding pixels' image intensities

$t = (i_1, \dots, i_n)$

3. Linear Shape Model of Faces

Given a database of \approx 3D models of faces, any new face can be described by

$$S' = \sum_{i=1}^n \beta_i S_i$$

where $S' = (x_1, y_1, z_1, \dots, x_n, y_n, z_n)$ is the new face's 3D shape and S_i is the 3D shape of the i^{th} face in the DB.

* Arbitrarily rotating 3D shape and projecting into an image does not change the β_i 's!

$$s_r = \sum_{i=1}^n \beta_i s_i^r$$

where $s_r = PS'$, P = perspective projection

4. Compute linear shape approximation

→ Find best β_i 's (min error)
so that

$$s^r = \sum \beta_i s_i^r$$

where s^r = input image face

s_i^r = i^{th} reference face image
at same pose

5. Compute new view's shape vector

$$s^f = \sum \beta_i s_i^f$$

6. Compute linear texture approximation

$$t^r = \sum \alpha_i t_i^r$$

~~Approximate~~

7. Compute new view's
texture vector

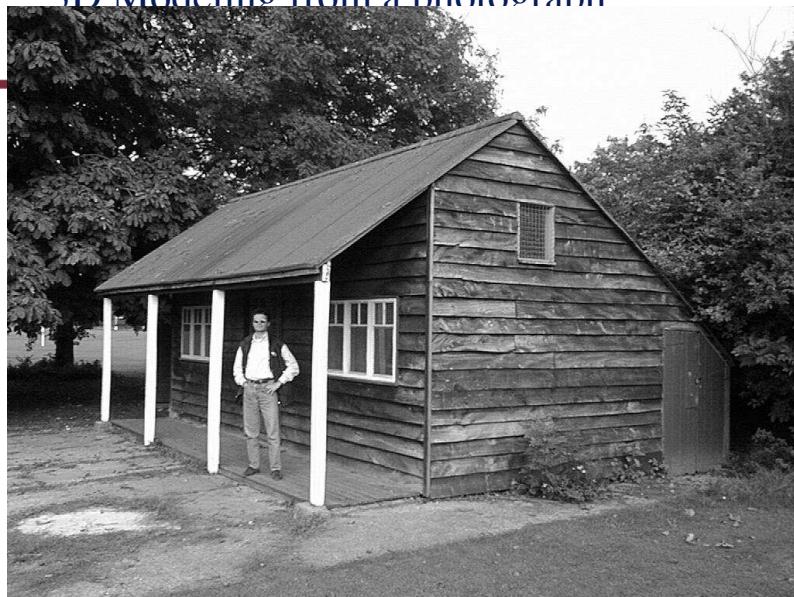
$$t^f = \sum \alpha_i t_i$$

Note: This texture mapping is
not correct except for
Lambertian surfaces.
Alternatively, use 3D head model
to remap texture

8. Warp texture onto shape

$$s^f + t^f$$

3D Modeling from a photograph

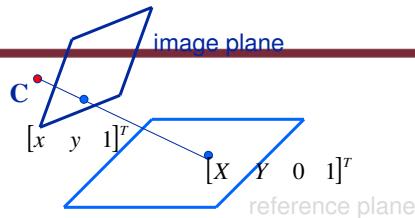


Criminisi et al., ICCV 99

1 Complete approach

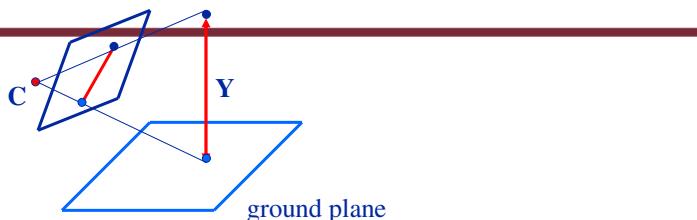
- u Load in an image
- u Click on lines parallel to X axis
 - u repeat for Y, Z axes
- u Compute vanishing points
- u Specify 3D and 2D positions of 4 points on reference plane
- u Compute homography H
- u Specify a reference height
- u Compute 3D positions of several points
- u Create a 3D model from these points
- u Extract texture maps
- u Output a VRML model

Measurements within reference plane



- 1 Solve for homography \mathbf{H} relating reference plane to image plane
 - u \mathbf{H} maps reference plane (X,Y) coords to image plane (x,y) coords
 - u Fully determined from 4 known points on ground plane
 - u Option A: physically measure 4 points on ground
 - u Option B: find a square, guess the dimensions
 - u Given (x, y), can find (X, Y) by \mathbf{H}^{-1}

Measuring height without a ruler



- 1 Compute Y from image measurements
 - u Need more than vanishing points to do this

The cross ratio

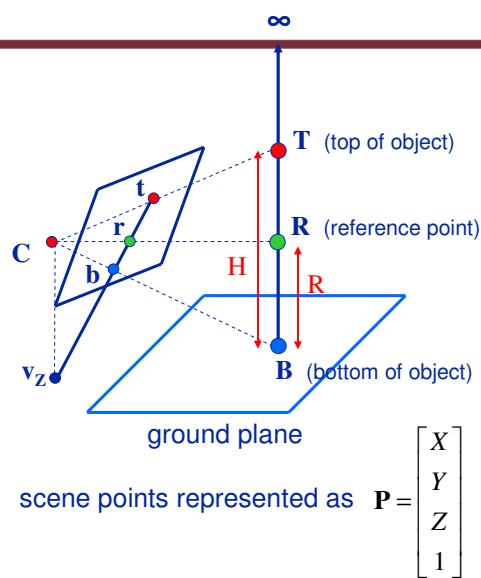
A Projective Invariant
The cross-ratio of 4 collinear points
u Something that does not change under projective transformations
(including perspective projection)

$$\frac{\|\mathbf{P}_3 - \mathbf{P}_1\| \|\mathbf{P}_4 - \mathbf{P}_2\|}{\|\mathbf{P}_3 - \mathbf{P}_2\| \|\mathbf{P}_4 - \mathbf{P}_1\|} \quad \mathbf{P}_i = \begin{bmatrix} X_i \\ Y_i \\ Z_i \\ 1 \end{bmatrix}$$

The diagram shows four points labeled $\mathbf{P}_1, \mathbf{P}_2, \mathbf{P}_3, \mathbf{P}_4$ arranged collinearly along a blue line.

- 1 Can permute the point ordering $\frac{\|\mathbf{P}_1 - \mathbf{P}_3\| \|\mathbf{P}_4 - \mathbf{P}_2\|}{\|\mathbf{P}_1 - \mathbf{P}_2\| \|\mathbf{P}_4 - \mathbf{P}_3\|}$
 - u $4! = 24$ different orders (but only 6 distinct values)
- 1 This is the fundamental invariant of projective geometry

Measuring height



$$\frac{\|T - B\|}{\|R - B\|} \frac{\|\infty - R\|}{\|\infty - T\|} = \frac{H}{R}$$

scene cross ratio

$$\frac{\|t - b\|}{\|r - b\|} \frac{\|v_z - r\|}{\|v_z - t\|} = \frac{H}{R}$$

image cross ratio

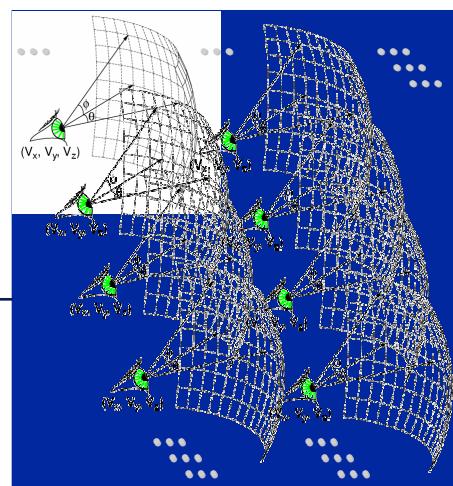
$$\text{image points as } \mathbf{p} = \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Plenoptic Function (Adelson and Bergen 1991)

- ① For a given scene, describe ALL rays through
 - u ALL pixels, of
 - u ALL cameras, at
 - u ALL wavelengths,
 - u ALL time

$$F(x, y, z, \phi, \theta, \lambda, t)$$

“eyeballs everywhere”
function (5D x 2D)



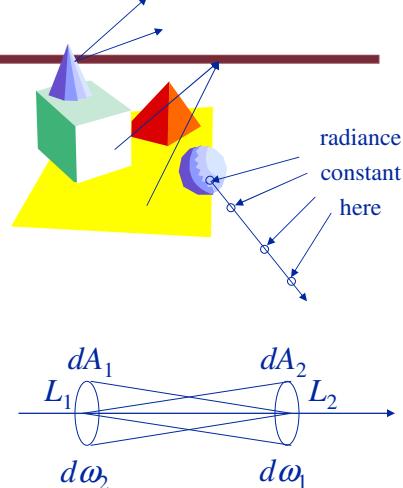
Plenoptic Array: ‘The Matrix Effect’

- 1 Brute force! Simple arc, line, or ring array of cameras
- 1 Synchronized shutter – Dayton Taylor’s TimeTrack
- 1 Warp/blend between images to change viewpoint on ‘time-frozen’ scene:



How Much Light is Really in a Scene?

- 1 Light transported throughout scene along rays
 - u Anchor
 - u Any point in 3D space
 - u 3 coordinates
 - u Direction
 - u Any 3D unit vector
 - u 2 angles
 - u Total of 5 dimensions
- 1 Radiance remains constant along ray as long as in empty space
 - u Removes one dimension
 - u Total of 4 dimensions



Light Field and Lumigraph



① Light Field

② Levoy & Hanrahan, Siggraph 1996

① Lumigraph

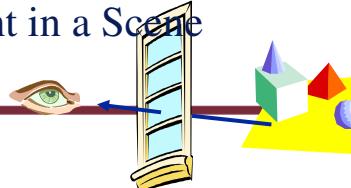
② Gortler *et al.*, Siggraph 1996

① Consider (u,v) the image plane and (s,t) the viewpoint plane

① Photographs taken from a bunch of different viewpoints

① Reconstructed photographs of scene are 2D slices of 4D light field

Representing All of the Light in a Scene



① View scene through a window

① All visible light from scene must have passed through window

① Window light is 4D

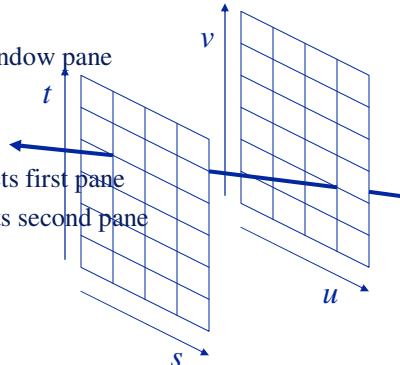
② 2 coordinates where ray intersects window pane

② 2 angles for ray direction

① Use a double-paned window

② 2 coordinates (u,v) where ray intersects first pane

② 2 coordinates (s,t) where ray intersects second pane



Object Based

- ① Can be made a cube around object

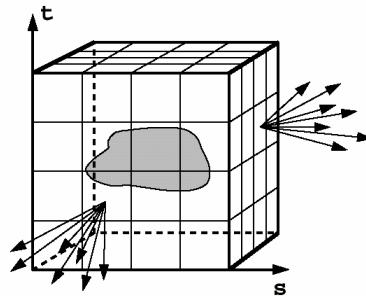
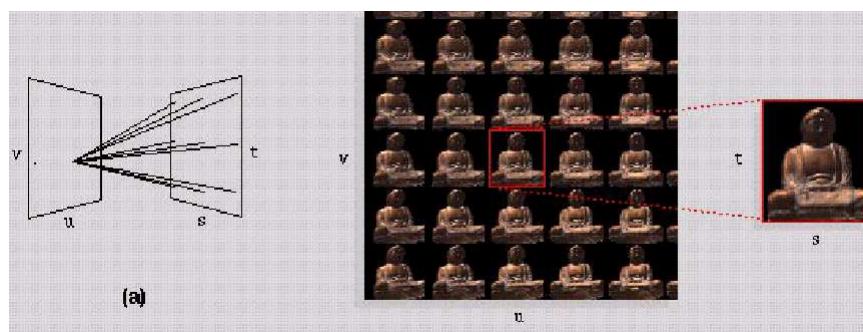
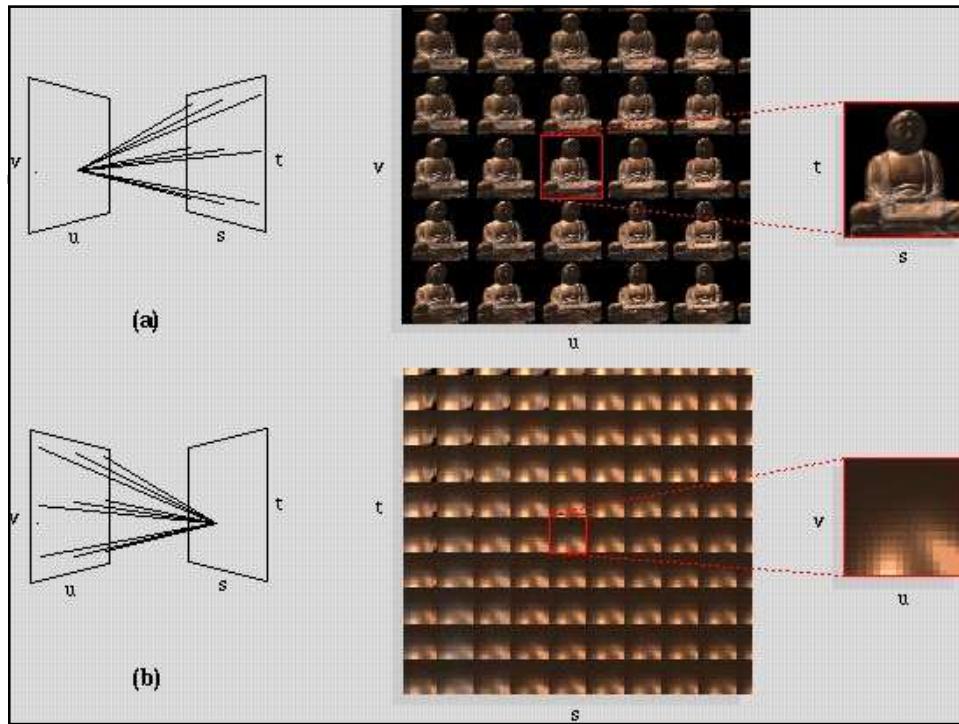


Figure 1: The surface of a cube holds all the radiance information due to the enclosed object.

Images from uv Plane





Making Light Field/Lumigraph

- ① Rendered from synthetic model
- ① Made from real world
 - ✉ With gantry
 - ✉ Handheld

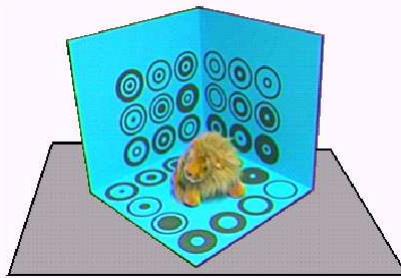
Gantry

- ① Lazy Susan
 - ② Manually rotated
- ① XY Positioner
- ① Lights turn with lazy susan
- ① Correctness by construction



Handheld Camera

- ① Blue screen on stage
 - ② Walls moveable, can turn



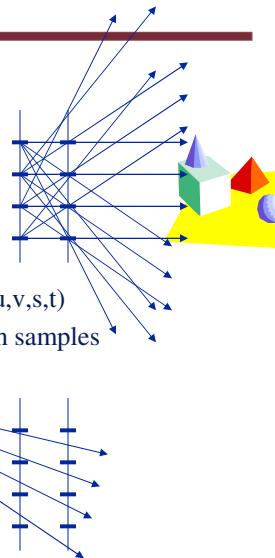
Ray Tracing and Light Fields

1 Rendering **into** a light field

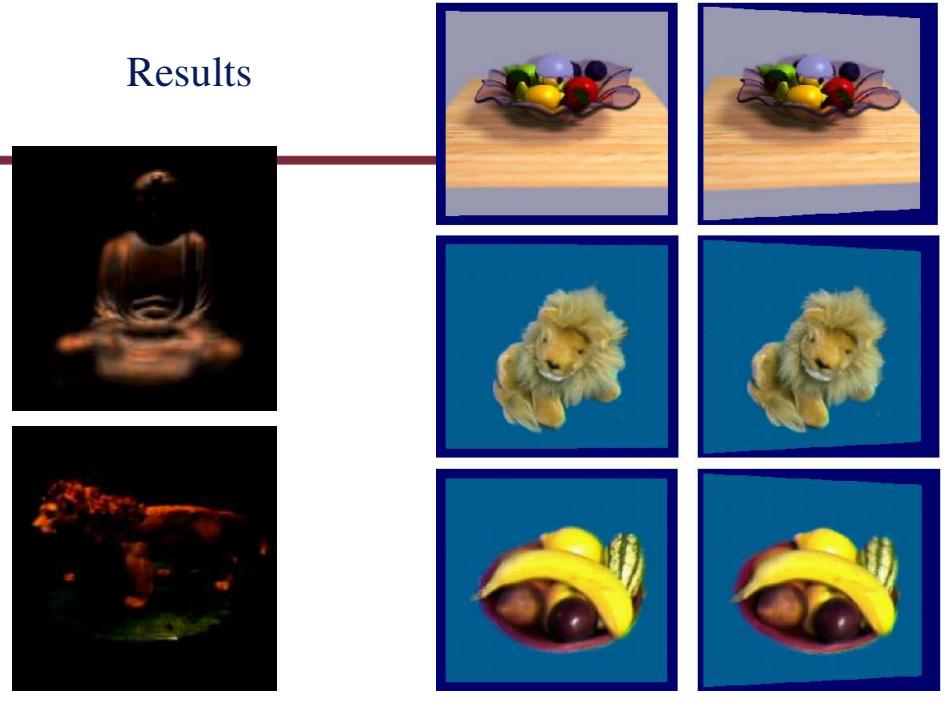
- Cast rays between all pairs of points in panes
- Store resulting radiance at (u,v,s,t)

1 Rendering **from** a light field

- Cast rays through pixels into light field
- Compute two ray-plane intersections to find (u,v,s,t)
- Interpolate u,v and s,t to find radiance between samples
- Plot radiance in pixel



Results



Results

¹ Light Field

	buddha	kidney	hallway	lion
Number of slabs	1	1	4	4
Images per slab	16x16	64x64	64x32	32x16
Total images	256	4096	8192	2048
Pixels per image	256^2	128^2	256^2	256^2
Raw size (MB)	50	201	1608	402
Prefiltering	uvst	st only	uvst	st only

Table I: Statistics of the light fields shown in figure 14.