More Mosaic Madness

© Jeffrey Martin (jeffrey-martin.com)

with a lot of slides stolen from Steve Seitz and Rick Szeliski

15-463: Computational Photography
Alexei Efros, CMU, Fall 2005

Homography

A: Projective – mapping between any two PPs with the same center of projection

- rectangle should map to arbitrary quadrilateral
- parallel lines aren't
- but must preserve straight lines
- same as: project, rotate, reproject

called Homography

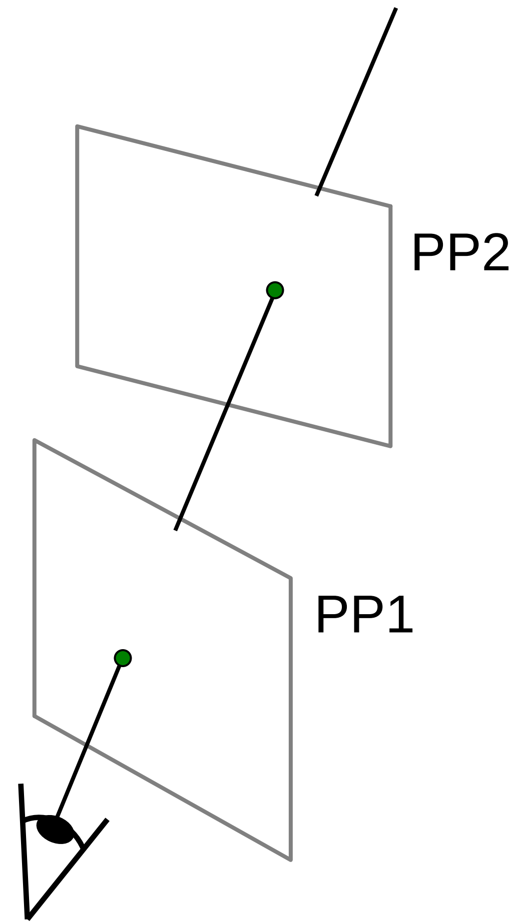

$$\underset{\mathbf{p'}}{\begin{bmatrix} wx' \\ wy' \\ w \end{bmatrix}} = \underset{\mathbf{H}}{\begin{bmatrix} * & * & * \\ * & * & * \\ * & * & * \end{bmatrix}} \underset{\mathbf{p}}{\begin{bmatrix} x \\ y \\ 1 \end{bmatrix}}$$

To apply a homography **H**

- Compute **p'** = **Hp** (regular matrix multiply)
- Convert **p'** from homogeneous to image coordinates

Rotational Mosaics

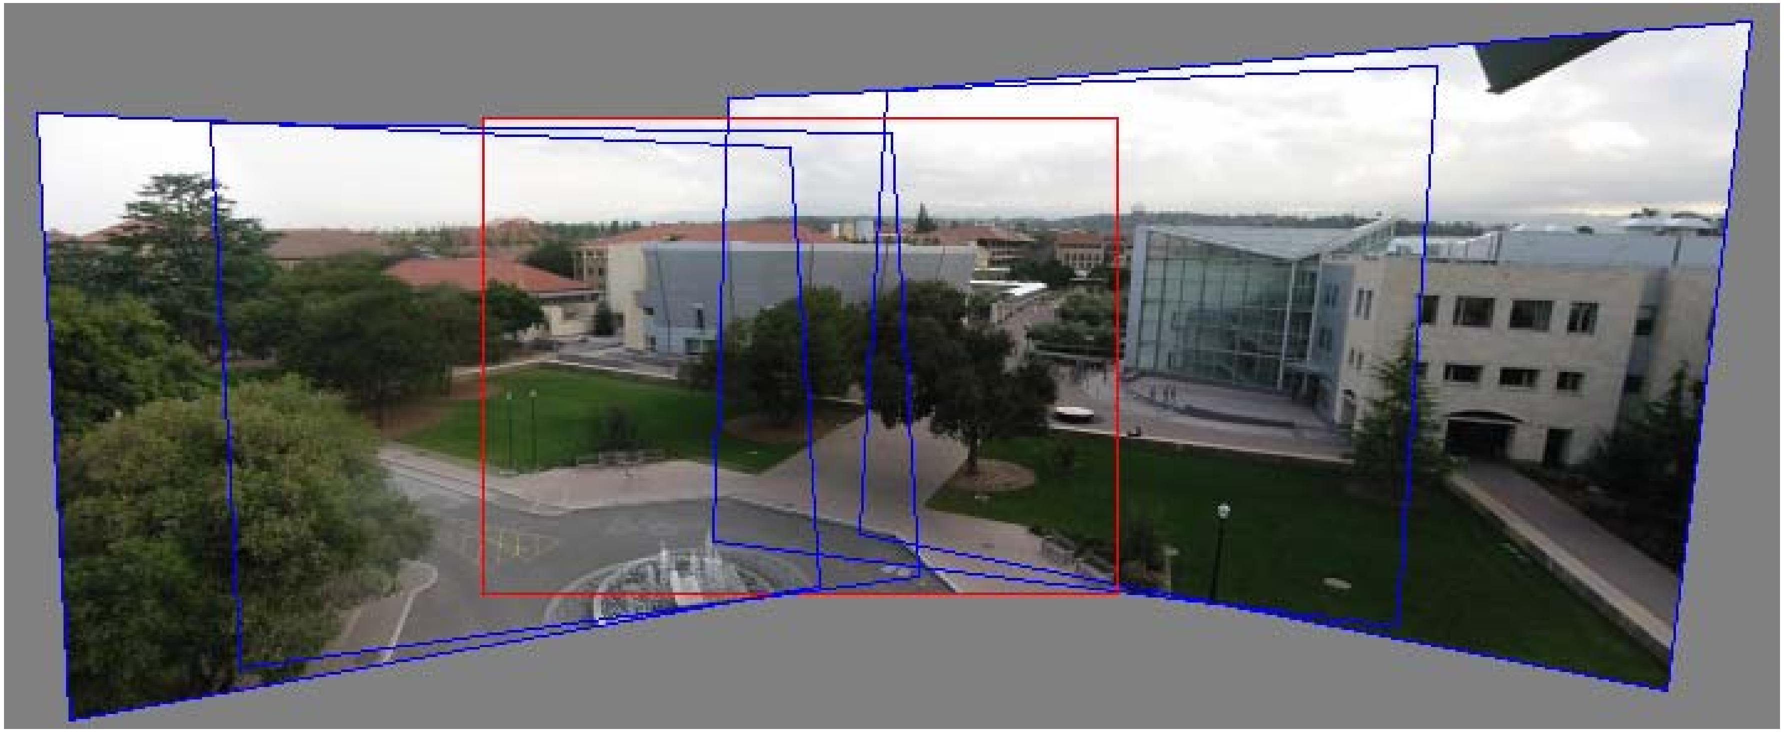

Can we say something more about rotational mosaics?

i.e. can we further constrain our H?

3D → 2D Perspective Projection

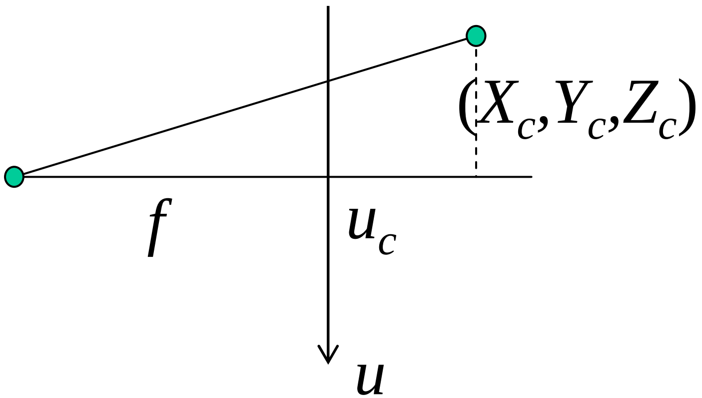

$$\begin{bmatrix} u \\ v \\ 1 \end{bmatrix} \sim \begin{bmatrix} U \\ V \\ W \end{bmatrix} = \begin{bmatrix} f & 0 & u_c \\ 0 & f & v_c \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} X_c \\ Y_c \\ Z_c \end{bmatrix}$$

K

3D Rotation Model

Projection equations

1. Project from image to 3D ray

   $(x_0,y_0,z_0) \quad = (u_0-u_c,v_0-v_c,f)$

2. Rotate the ray by camera motion

   $(x_1,y_1,z_1) \quad = \boldsymbol{R}_{01}\,(x_0,y_0,z_0)$

3. Project back into new (source) image

   $(u_1,v_1) \qquad = (fx_1/z_1+u_c,fy_1/z_1+v_c)$

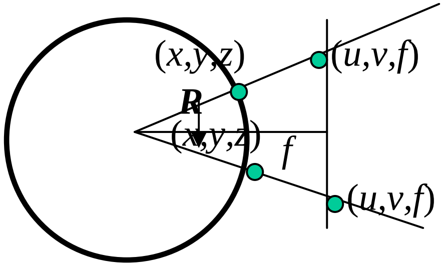

Therefore:

$$\mathbf{H} = \mathbf{K}_0 \mathbf{R}_{01} \mathbf{K}_1^{-1}$$

Our homography has only 3,4 or 5 DOF, depending if focal length is known, same, or different.

- This makes image registration much better behaved

Pairwise alignment

Procrustes Algorithm [Golub & VanLoan]

Given two sets of matching points, compute R

$p_i' = \boldsymbol{R}\, p_i$ with 3D rays

$$p_i = N(x_i, y_i, z_i) = N(u_i - u_c, v_i - v_c, f)$$

$$\boldsymbol{A} = \Sigma_{\mathbf{i}}\, p_i\, p_i'^T = \Sigma_{\mathbf{i}}\, p_i\, p_i^T \boldsymbol{R}^T = \boldsymbol{U}\,\boldsymbol{S}\,\boldsymbol{V}^T = (\boldsymbol{U}\,\boldsymbol{S}\,\boldsymbol{U}^T)\,\boldsymbol{R}^T$$

$$\boldsymbol{V}^T = \boldsymbol{U}^T \boldsymbol{R}^T$$

$$\boldsymbol{R} = \boldsymbol{V}\,\boldsymbol{U}^T$$

Rotation about vertical axis

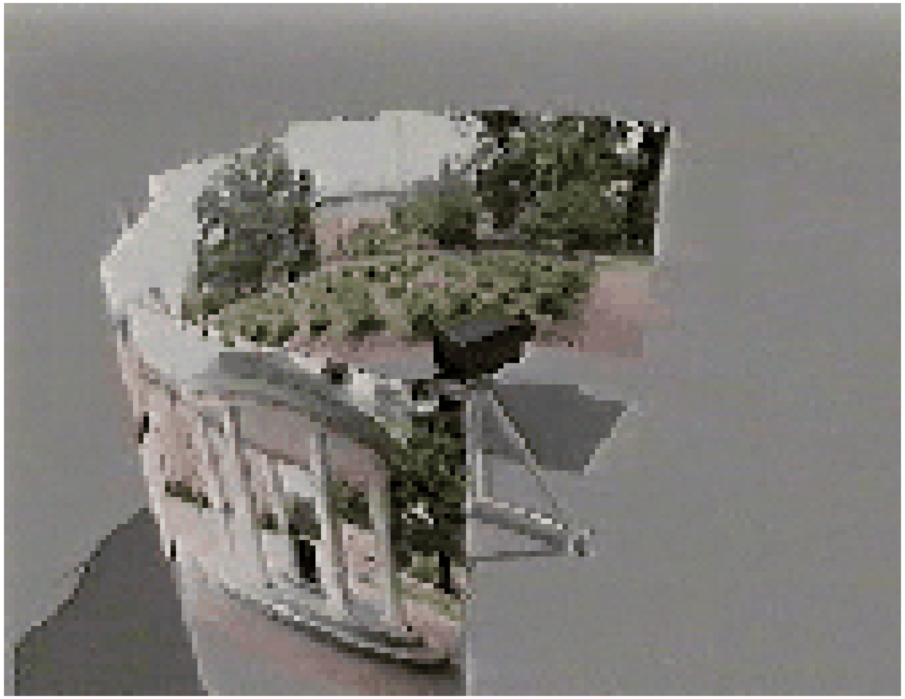

What if our camera rotates on a tripod?

What's the structure of H?

Do we have to project onto a plane?

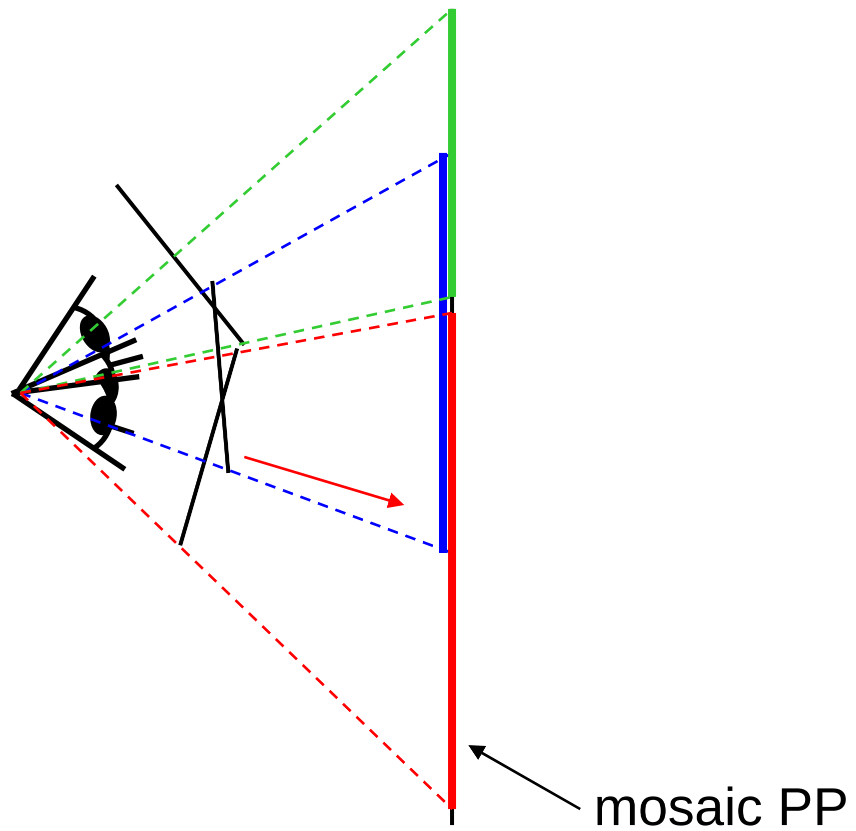

Full Panoramas

What if you want a 360° field of view?

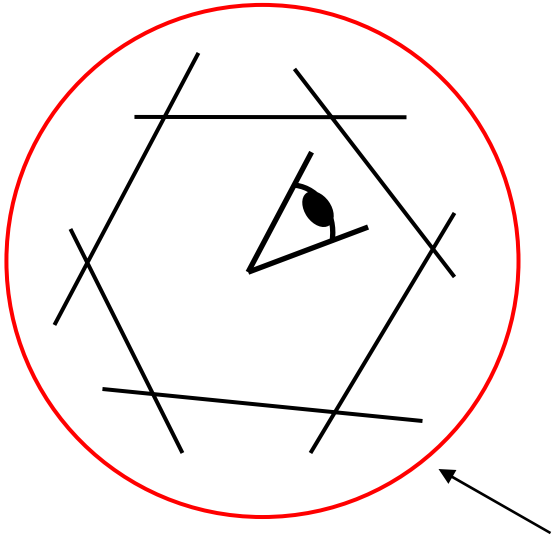

Cylindrical projection

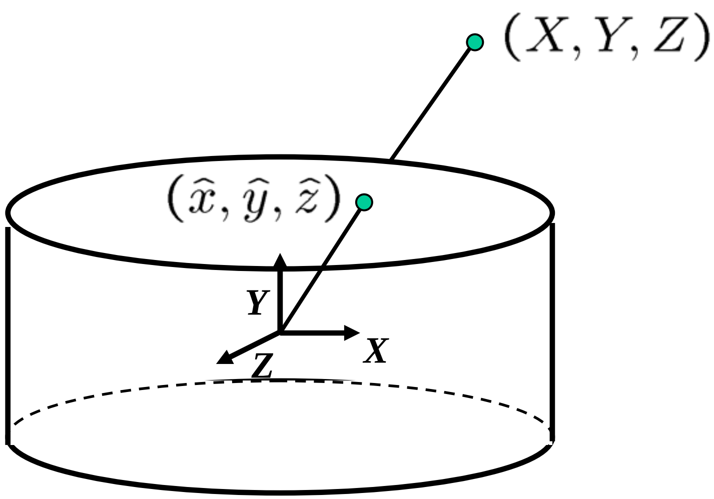

unit cylinder

- Map 3D point (X,Y,Z) onto cylinder

$$(\hat{x}, \hat{y}, \hat{z}) = \frac{1}{\sqrt{X^2+Z^2}}(X, Y, Z)$$

- Convert to cylindrical coordinates

$$(sin\theta, h, cos\theta) = (\hat{x}, \hat{y}, \hat{z})$$

- Convert to cylindrical image coordinates

$$(\tilde{x}, \tilde{y}) = (f\theta, fh) + (\tilde{x}_c, \tilde{y}_c)$$

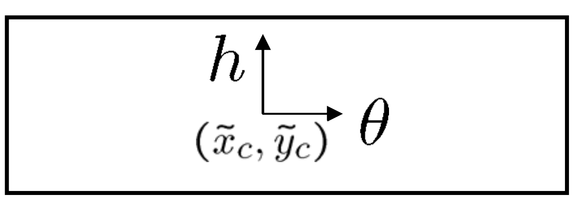

unwrapped cylinder

cylindrical image

Cylindrical Projection

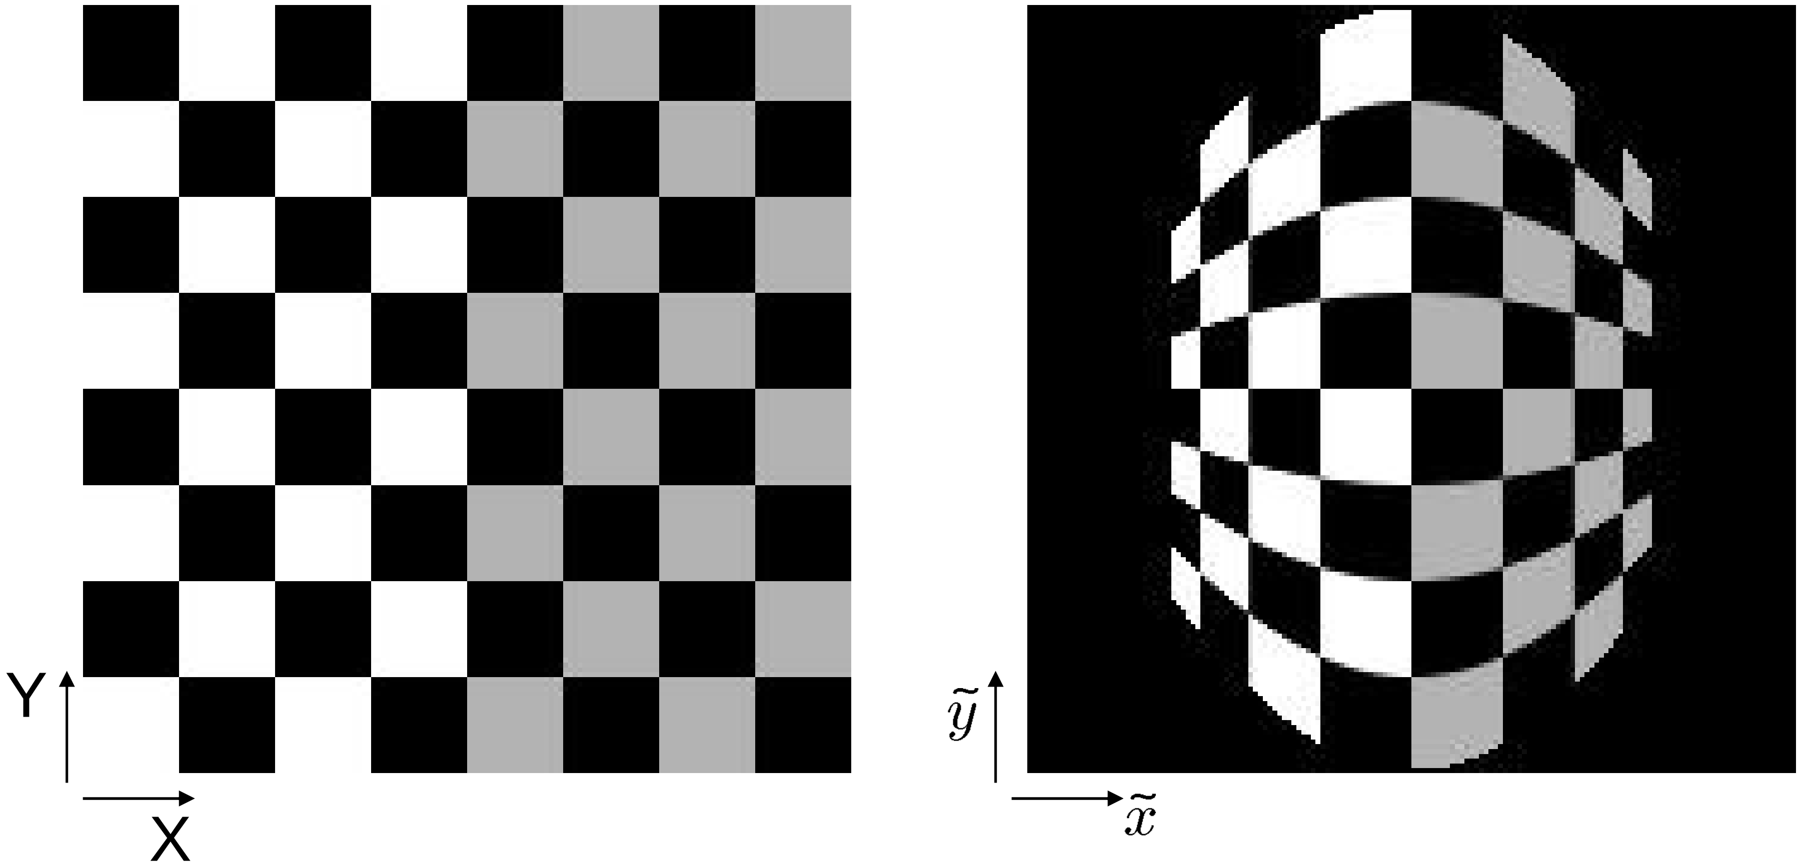

Inverse Cylindrical projection

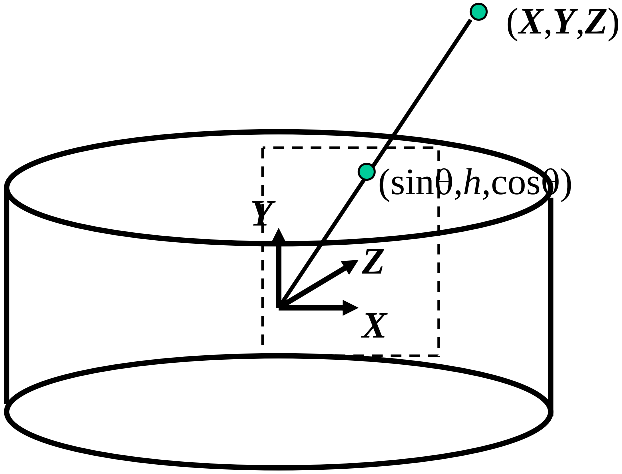

$$
\begin{aligned}
\theta &= (x_{cyl} - x_c)/f \\
h &= (y_{cyl} - y_c)/f \\
\hat{x} &= \sin\theta \\
\hat{y} &= h \\
\hat{z} &= \cos\theta \\
x &= f\hat{x}/\hat{z} + x_c \\
y &= f\hat{y}/\hat{z} + y_c
\end{aligned}
$$

Cylindrical panoramas

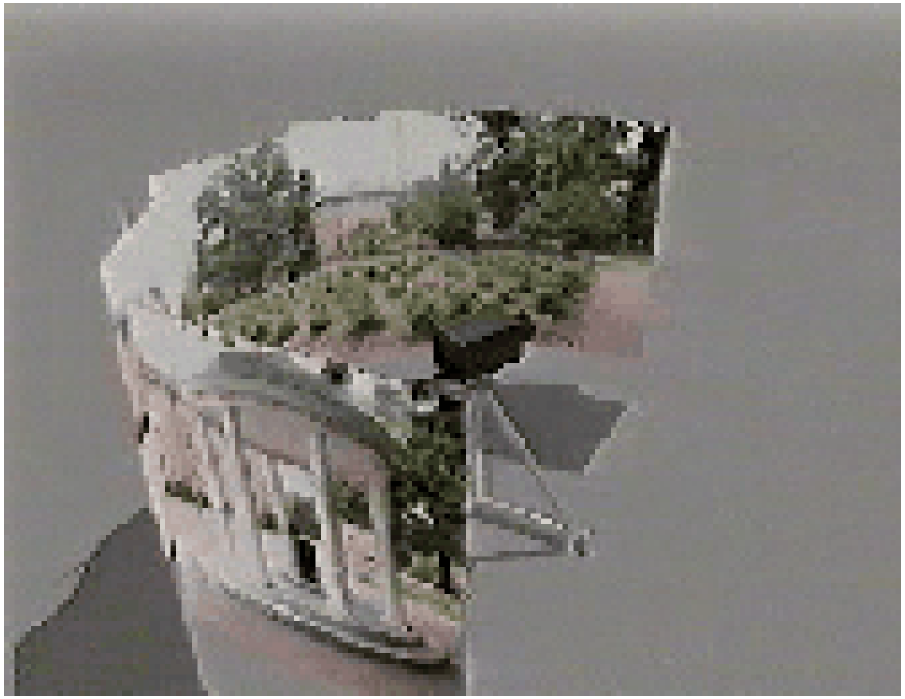

Steps

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

What are the assumptions here?

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!

Assembling the panorama

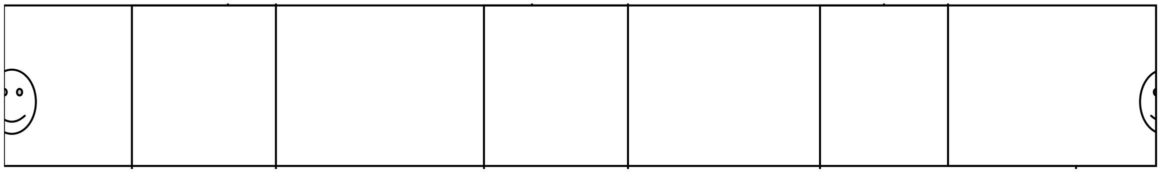

Stitch pairs together, blend, then crop

Problem: Drift

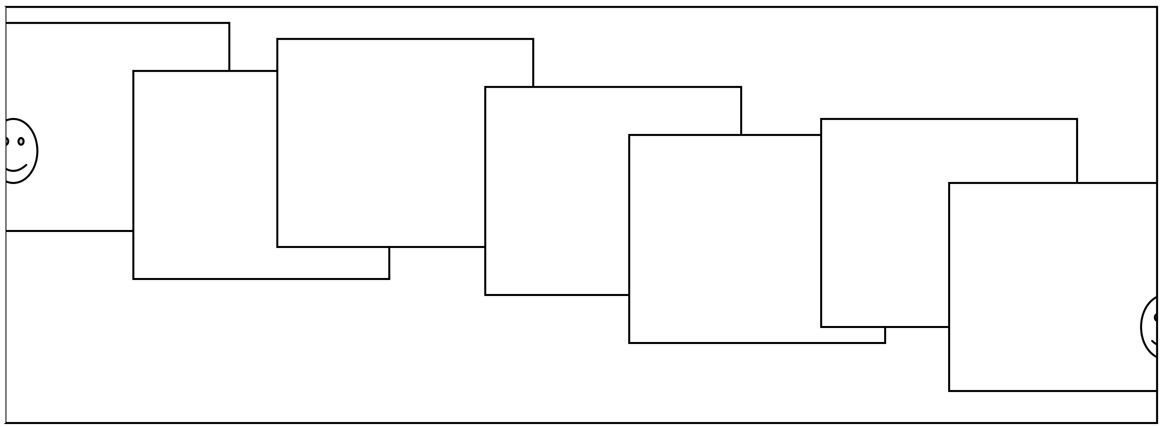

Vertical Error accumulation

- small (vertical) errors accumulate over time
- apply correction so that sum = 0 (for 360° pan.)

Horizontal Error accumulation

- can reuse first/last image to find the right panorama radius

Full-view (360°) panoramas

Spherical projection

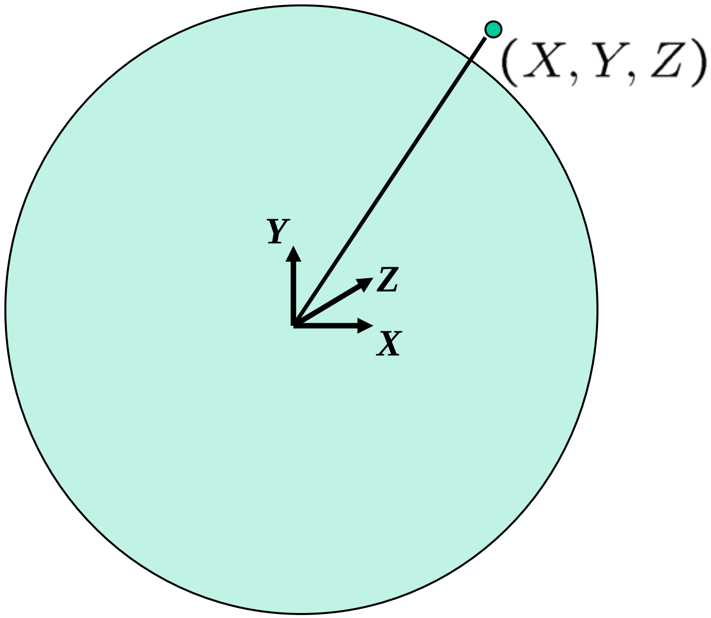

- Map 3D point (X,Y,Z) onto sphere

$$(\hat{x}, \hat{y}, \hat{z}) = \frac{1}{\sqrt{X^2 + Y^2 + Z^2}} (X, Y, Z)$$

- Convert to spherical coordinates

$$(\sin\theta\cos\phi, \sin\phi, \cos\theta\cos\phi) = (\hat{x}, \hat{y}, \hat{z})$$

- Convert to spherical image coordinates

$$(\tilde{x}, \tilde{y}) = (f\theta, fh) + (\tilde{x}_c, \tilde{y}_c)$$

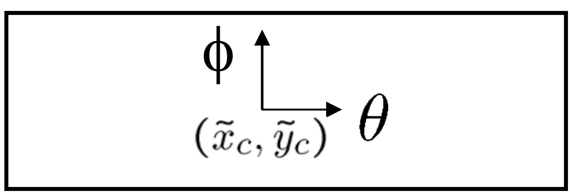

unwrapped sphere

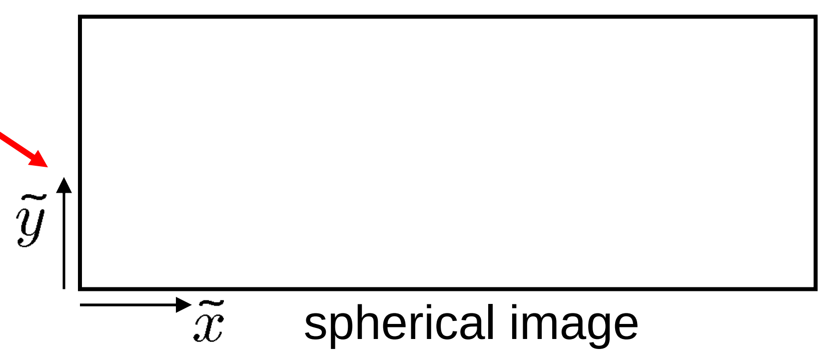

spherical image

Spherical Projection

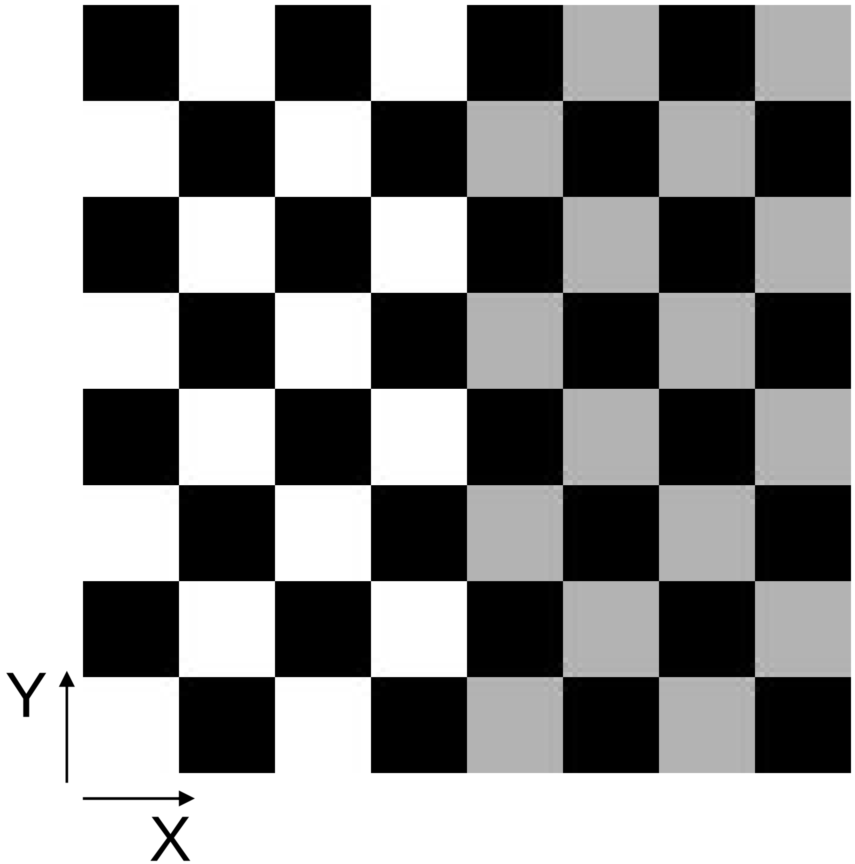

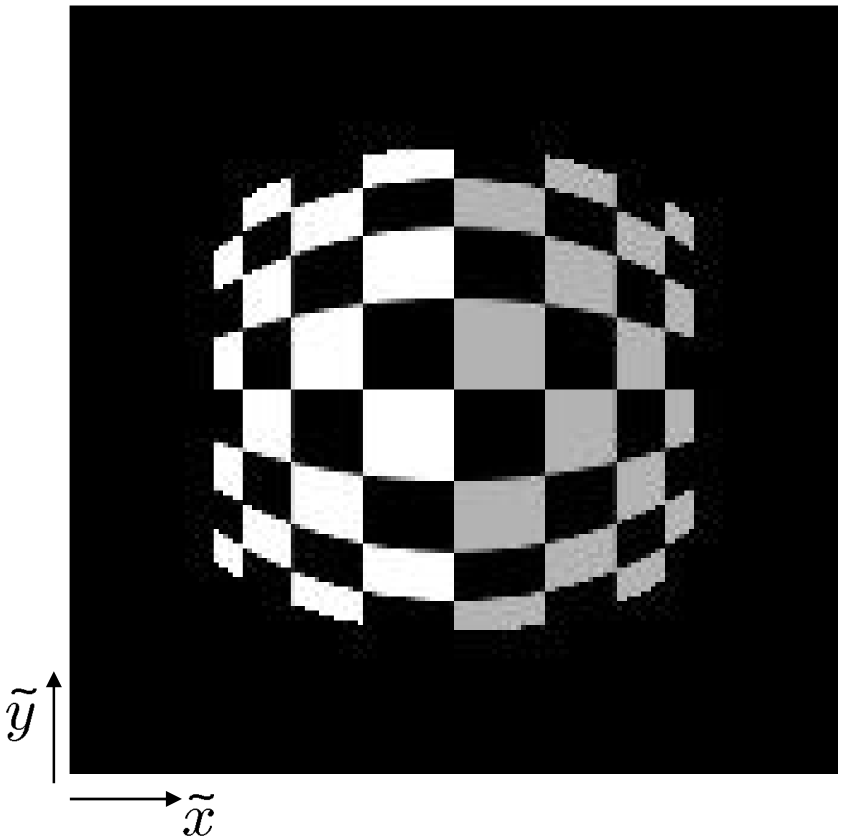

Inverse Spherical projection

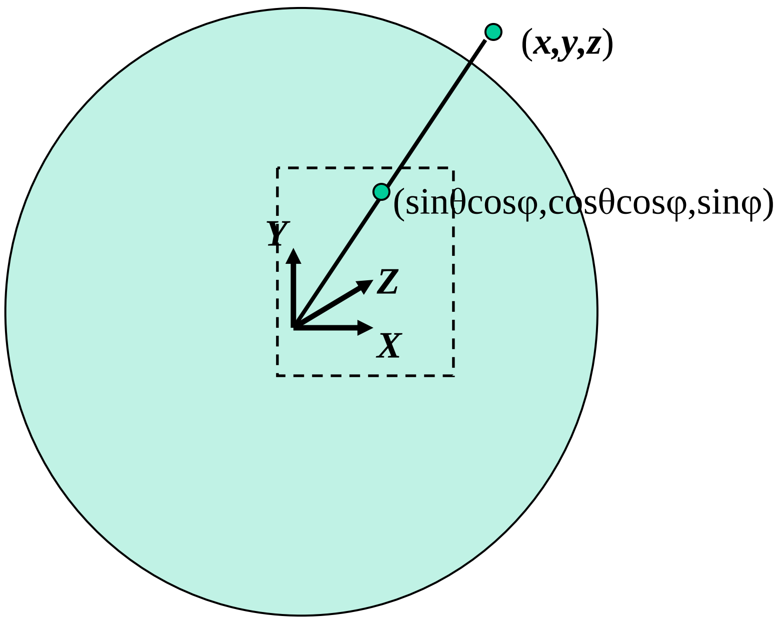

$$\theta = (x_{sph} - x_c)/f$$

$$\varphi = (y_{sph} - y_c)/f$$

$$\hat{x} = \sin\theta \cos\varphi$$

$$\hat{y} = \sin\varphi$$

$$\hat{z} = \cos\theta \cos\varphi$$

$$x = f\hat{x}/\hat{z} + x_c$$

$$y = f\hat{y}/\hat{z} + y_c$$

3D rotation

Rotate image before placing on unrolled sphere

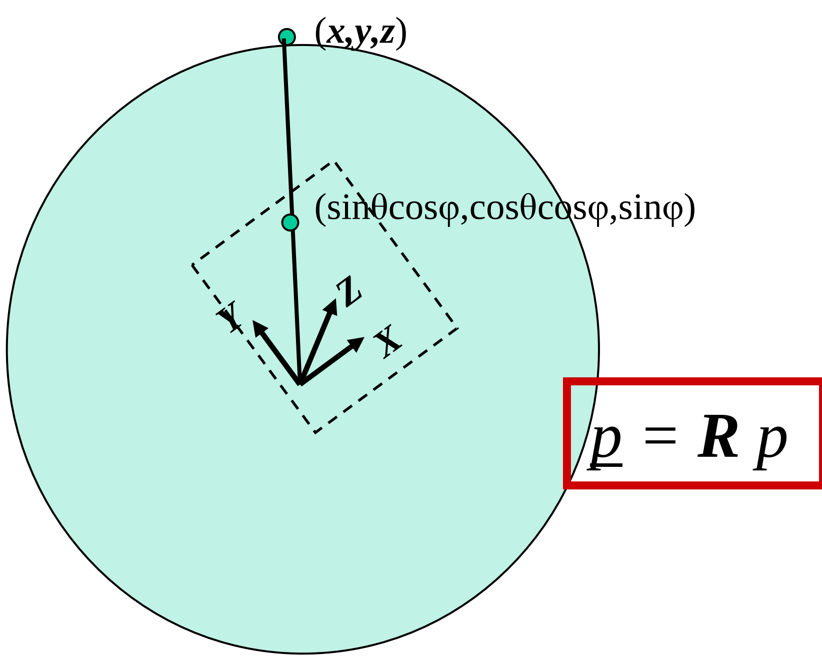

$$
\begin{aligned}
\theta &= (x_{sph} - x_c)/f \\
\varphi &= (y_{sph} - y_c)/f \\
\hat{x} &= \sin\theta \cos\varphi \\
\hat{y} &= \sin\varphi \\
\hat{z} &= \cos\theta \cos\varphi \\
x &= f\underline{\hat{x}}/\underline{\hat{z}} + x_c \\
y &= f\underline{\hat{y}}/\underline{\hat{z}} + y_c
\end{aligned}
$$

Full-view Panorama

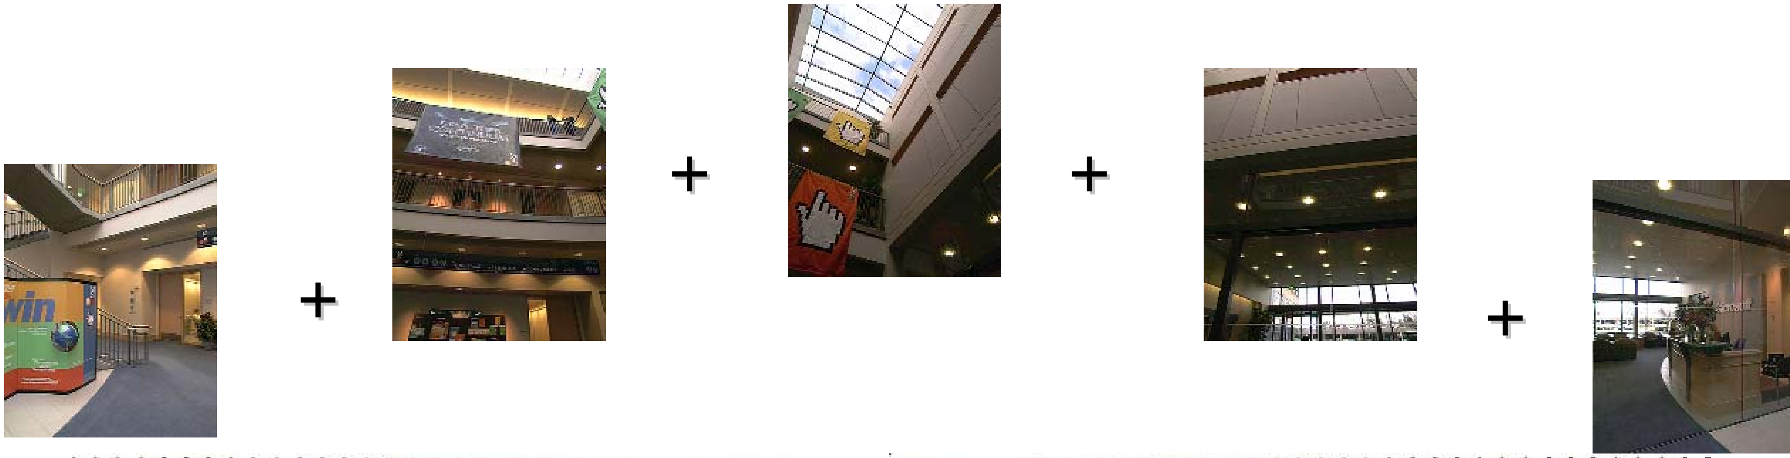

Polar Projection

Extreme "bending" in ultra-wide fields of view

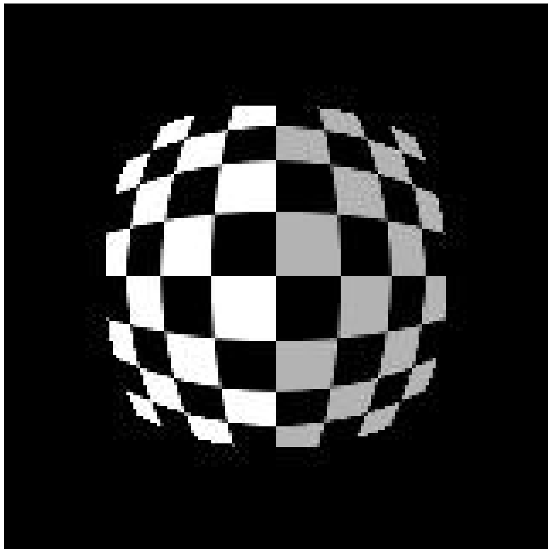

$$\hat{r}^2 = \hat{x}^2 + \hat{y}^2$$

$$(\cos\theta\sin\phi, \sin\theta\sin\phi, \cos\phi) = s\,(x, y, z)$$

uations become

$$\begin{aligned} x' &= s\phi\cos\theta = s\frac{x}{r}\tan^{-1}\frac{r}{z}, \\ y' &= s\phi\sin\theta = s\frac{y}{r}\tan^{-1}\frac{r}{z}, \end{aligned}$$

Other projections are possible

You can stitch on the plane and then warp the resulting panorama

- What's the limitation here?

Or, you can use these as stitching surfaces

- But there is a catch…

Cylindrical reprojection

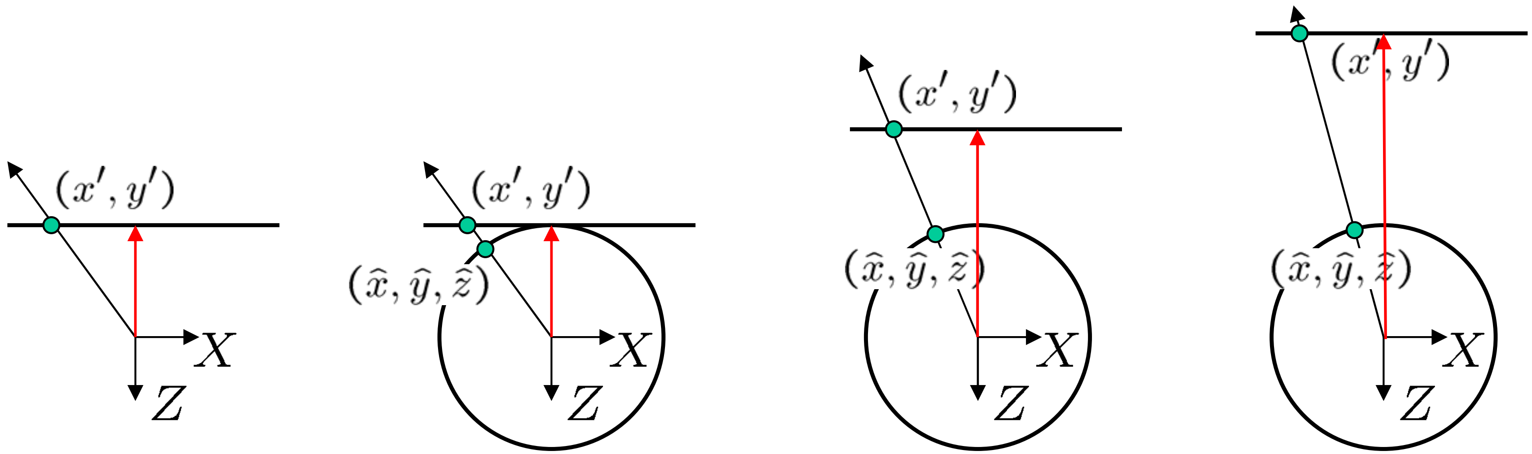

top-down view

Focal length – the dirty secret…

Image 384x300 **f = 180 (pixels)** **f = 280** **f = 380**

What's your focal length, buddy?

Focal length is (highly!) camera dependant

- Can get a rough estimate by measuring FOV:

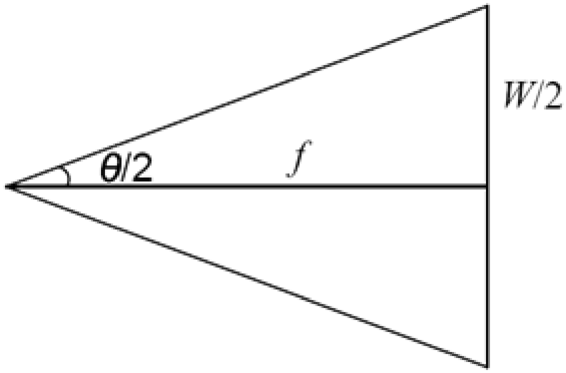

- Can use the EXIF data tag (might not give the right thing)
- Can use several images together and try to find f that would make them match
- Can use a known 3D object and its projection to solve for f
- Etc.

There are other camera parameters too:

- Optical center, non-square pixels, lens distortion, etc.

Distortion

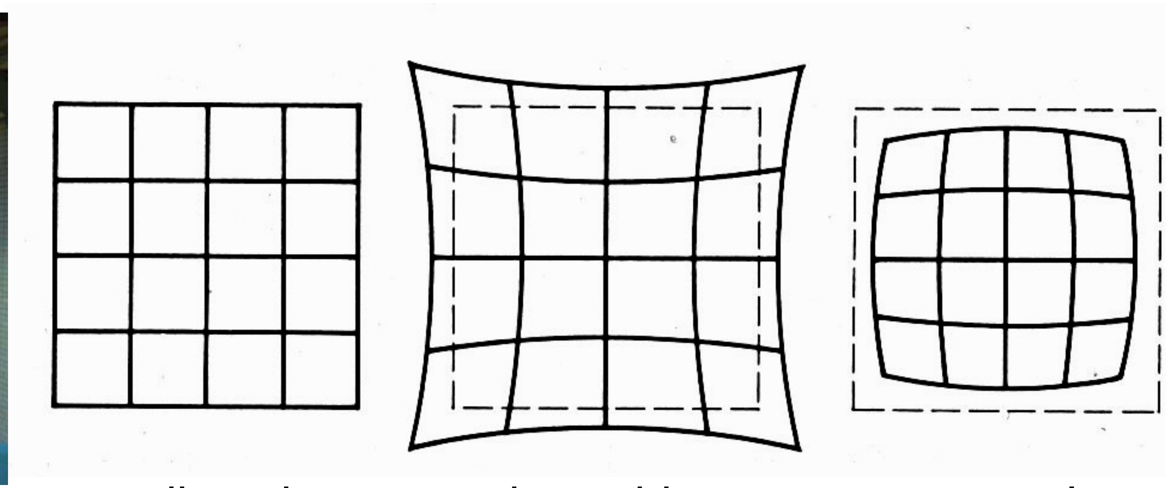

No distortion Pin cushion Barrel

Radial distortion of the image

- Caused by imperfect lenses
- Deviations are most noticeable for rays that pass through the edge of the lens

Radial distortion

Correct for "bending" in wide field of view lenses

$$
\begin{aligned}
\hat{r}^2 &= \hat{x}^2 + \hat{y}^2 \\
\hat{x}' &= \hat{x}/(1 + \kappa_1 \hat{r}^2 + \kappa_2 \hat{r}^4) \\
\hat{y}' &= \hat{y}/(1 + \kappa_1 \hat{r}^2 + \kappa_2 \hat{r}^4) \\
x &= f\hat{x}'/\hat{z} + x_c \\
y &= f\hat{y}'/\hat{z} + y_c
\end{aligned}
$$

Use this instead of normal projection

Blending the mosaic

An example of image compositing:
the art (and sometime science) of
combining images together…

Feathering

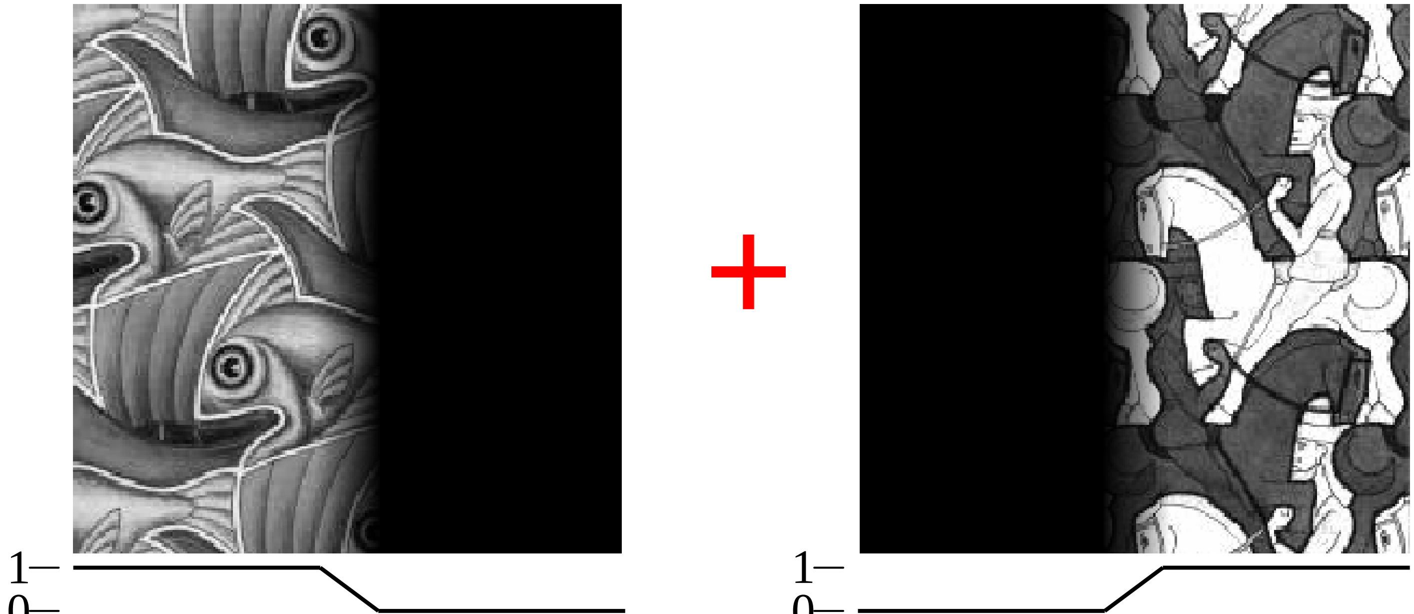

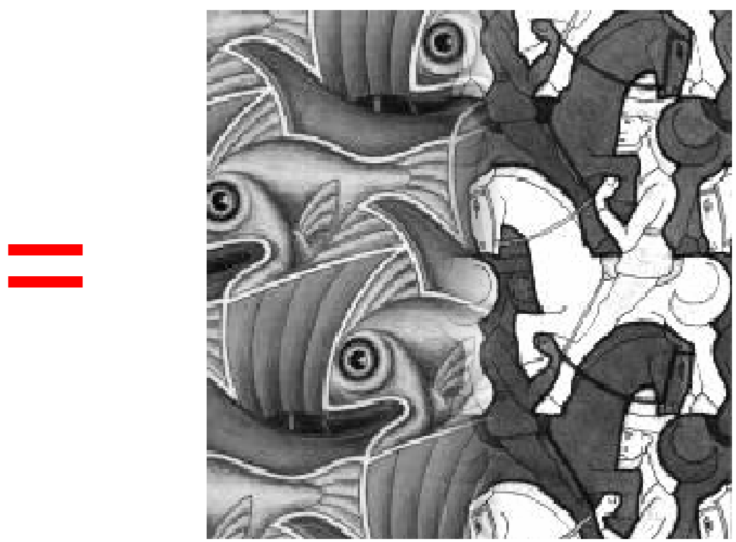

Encoding transparency

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

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

Setting alpha: simple averaging

Alpha = .5 in overlap region

Image feathering

Weight each image proportional to its distance from the edge
(distance map [Danielsson, CVGIP 1980]

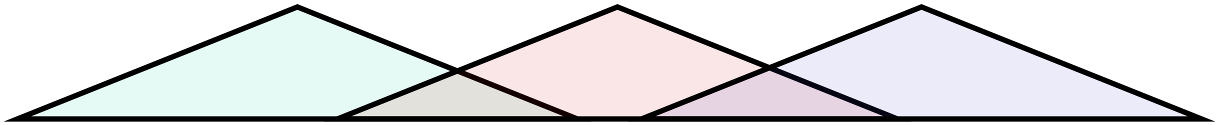

1. Generate *weight map* for each image
2. Sum up all of the weights and divide by sum: weights sum up to 1: $w_i' = w_i / (\sum_i w_i)$

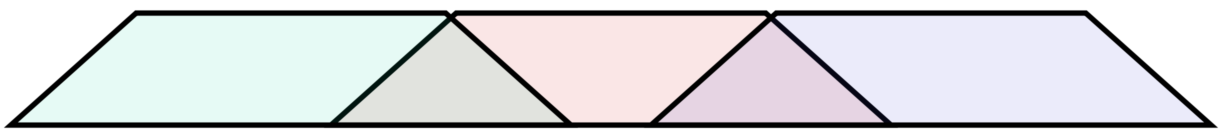

Setting alpha: center seam

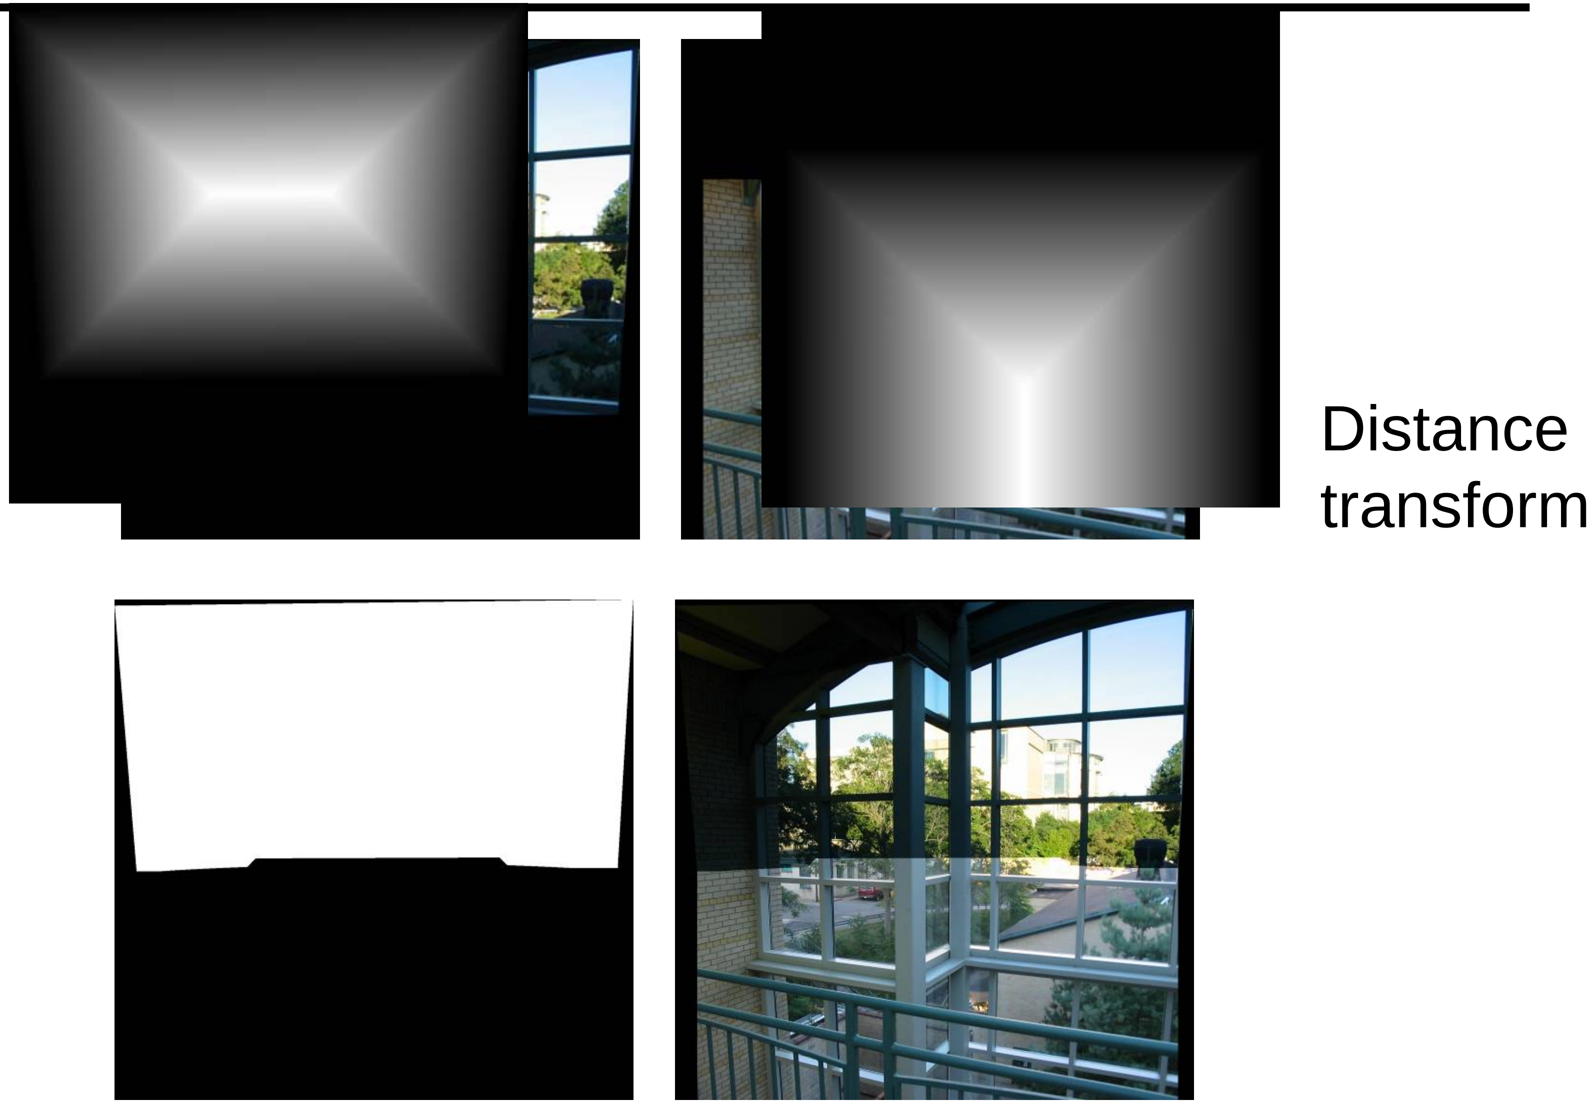

Alpha = logical(dtrans1>dtrans2)

Setting alpha: blurred seam

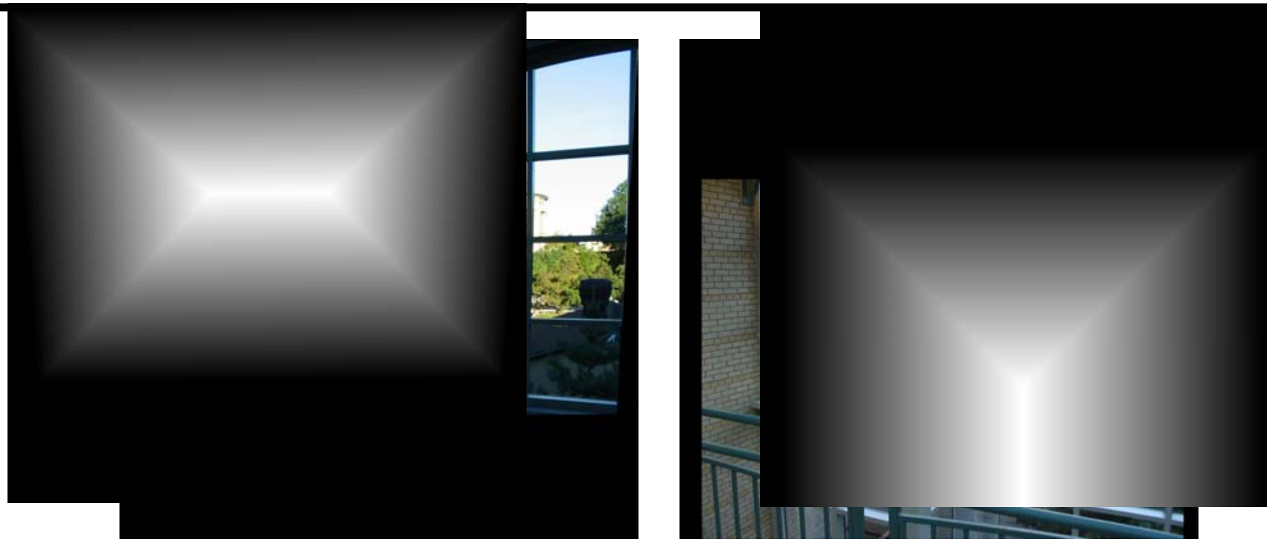

Distance transform

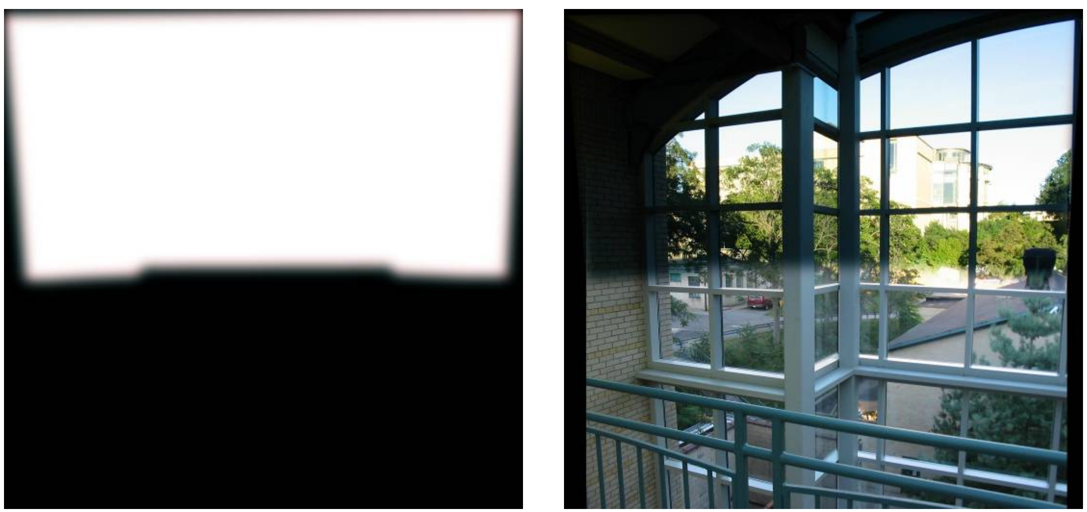

Alpha = blurred

Setting alpha: center weighting

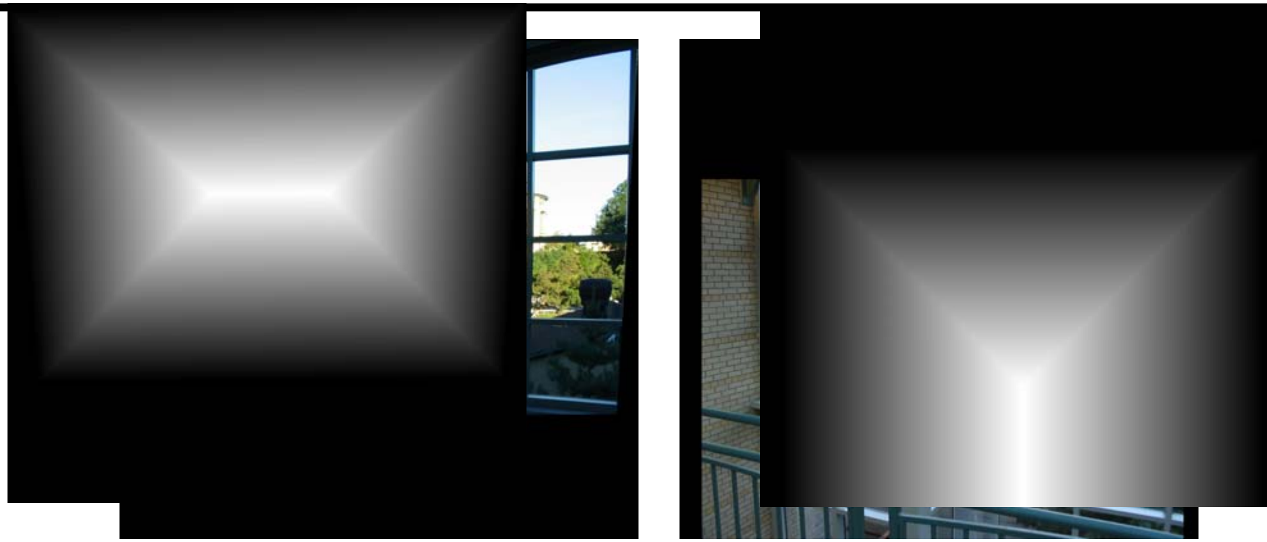

Distance transform

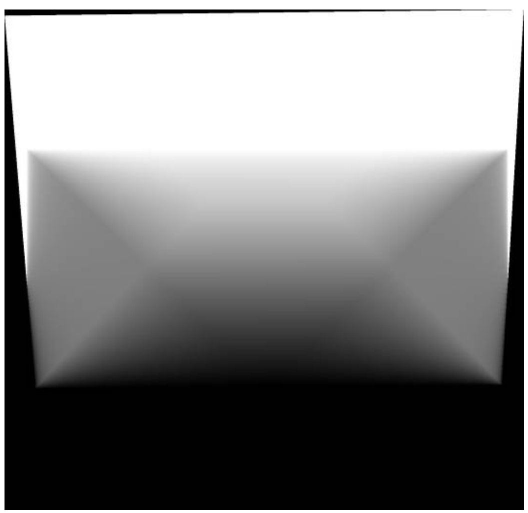

Ghost!

Alpha = dtrans1 / (dtrans1+dtrans2)

Pyramid Blending

For Alpha, use original feather weights to select *strongest* contributing image

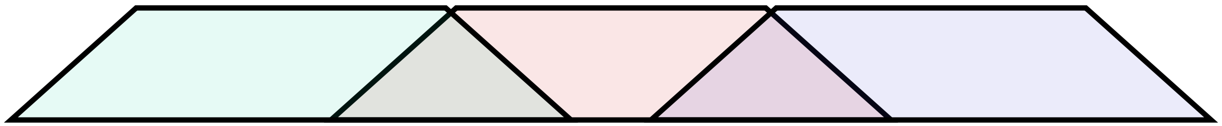

Can be implemented using L-∞ norm: ($p = 10$)

$$w_i' = [w_i^p / (\sum_i w_i^p)]^{1/p}$$

Mid-term Tuesday

Closed book, closed notes.

You are allowed one sheet of paper, both sides