Computational Photography

Prof. Feng Liu

Spring 2022

http://www.cs.pdx.edu/~fliu/courses/cs510/

04/28/2022

With slides by C. Dyer, Y. Chuang, R. Szeliski, S. Seitz, M. Brown and V. Hlavac

Last Time

- Panorama
  - Feature detection
  - Feature matching

Today

- Panorama
  - Homography estimation
  - Cylindrical panorama
  - Blending

Stitching Recipe

- Align pairs of images
  - Feature Detection
  - Feature Matching
  - Homography Estimation

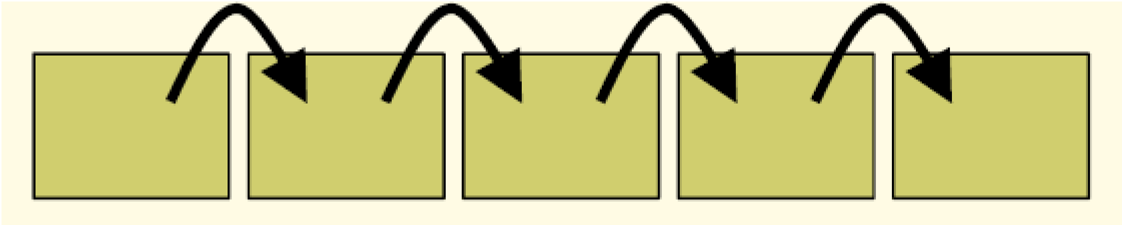

- Align all to a common surface
- Adjust (Global) & Blend

What can be globally aligned?

- In image stitching, we seek for a model to globally warp one image into another. Are any two images of the same scene can be aligned this way?
  - Images captured with the same center of projection
  - A planar scene or far-away scene

Credit: Y.Y. Chuang

A pencil of rays contains all views

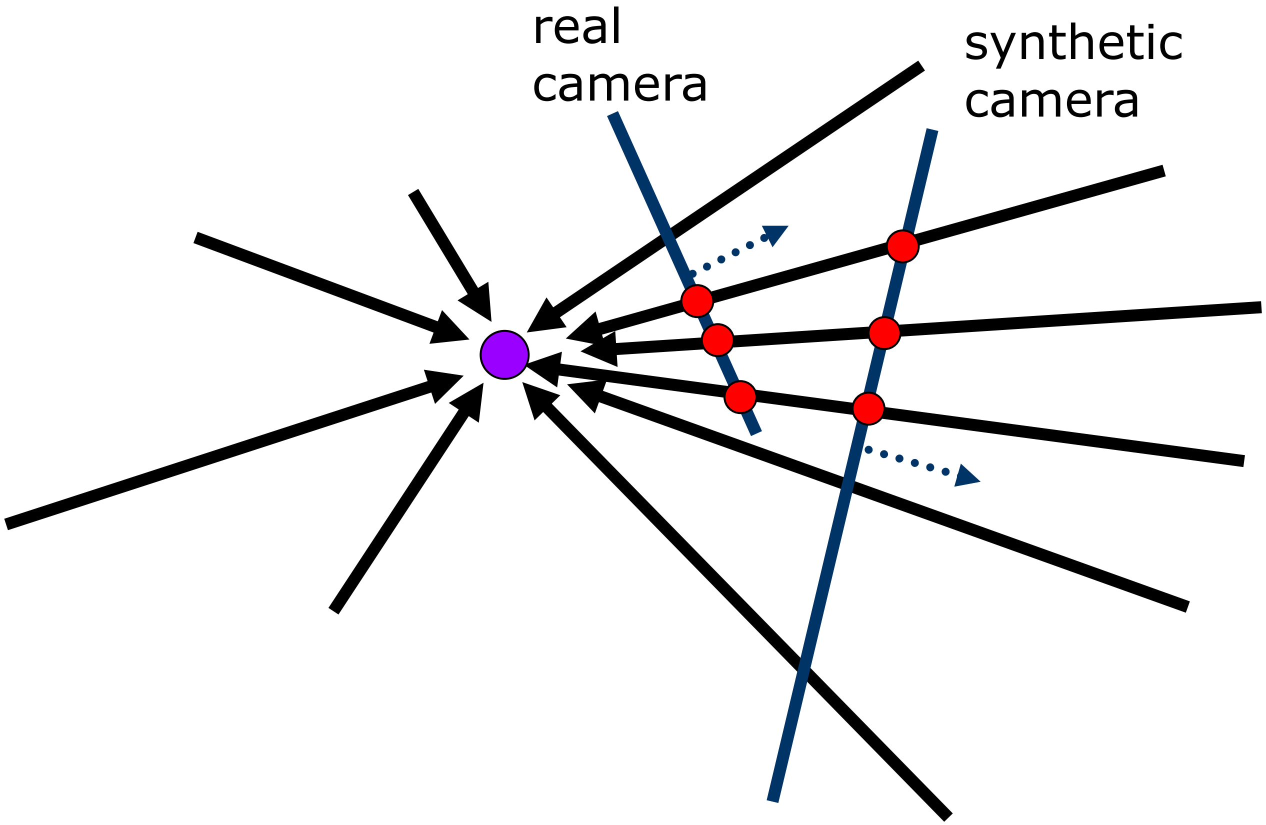

Can generate any synthetic camera view
as long as it has the same center of projection!

Credit: Y.Y. Chuang

Mosaic as an image reprojection

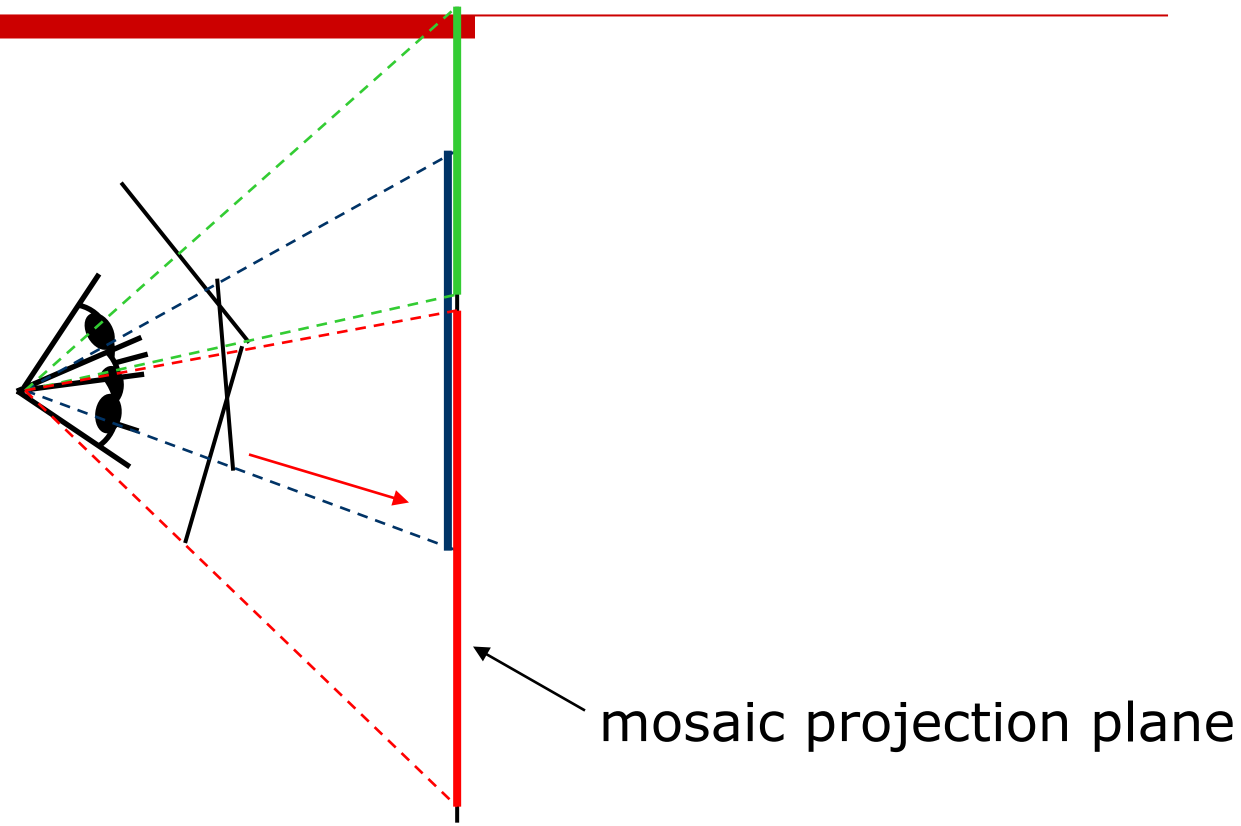

- The images are reprojected onto a common plane
- The mosaic is formed on this plane
- Mosaic is a *synthetic wide-angle camera*

Credit: Y.Y. Chuang

Changing camera center

Does it still work?

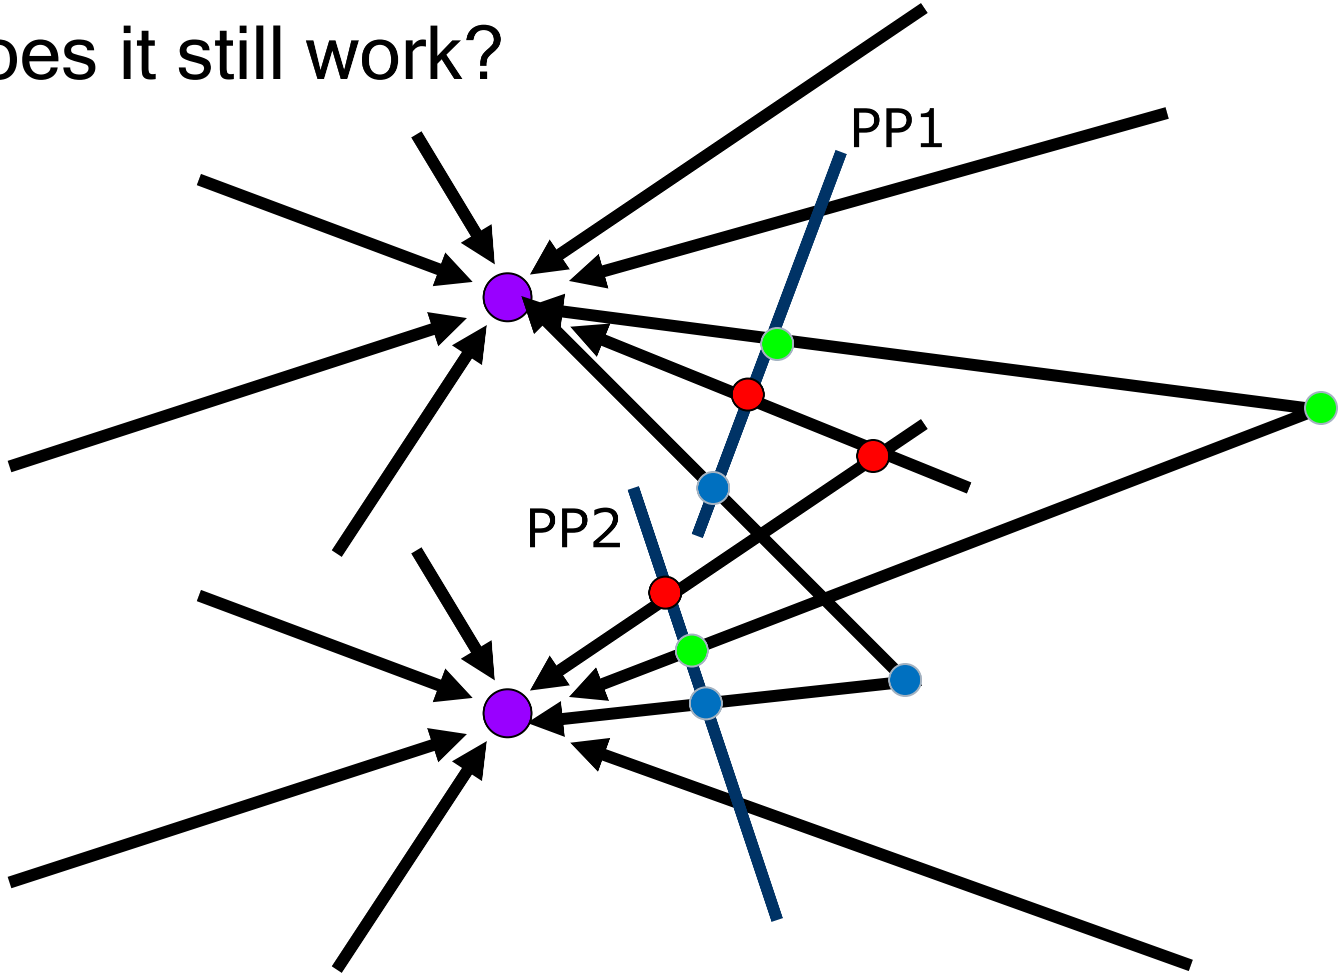

Credit: Y.Y. Chuang

Planar scene (or a faraway one)

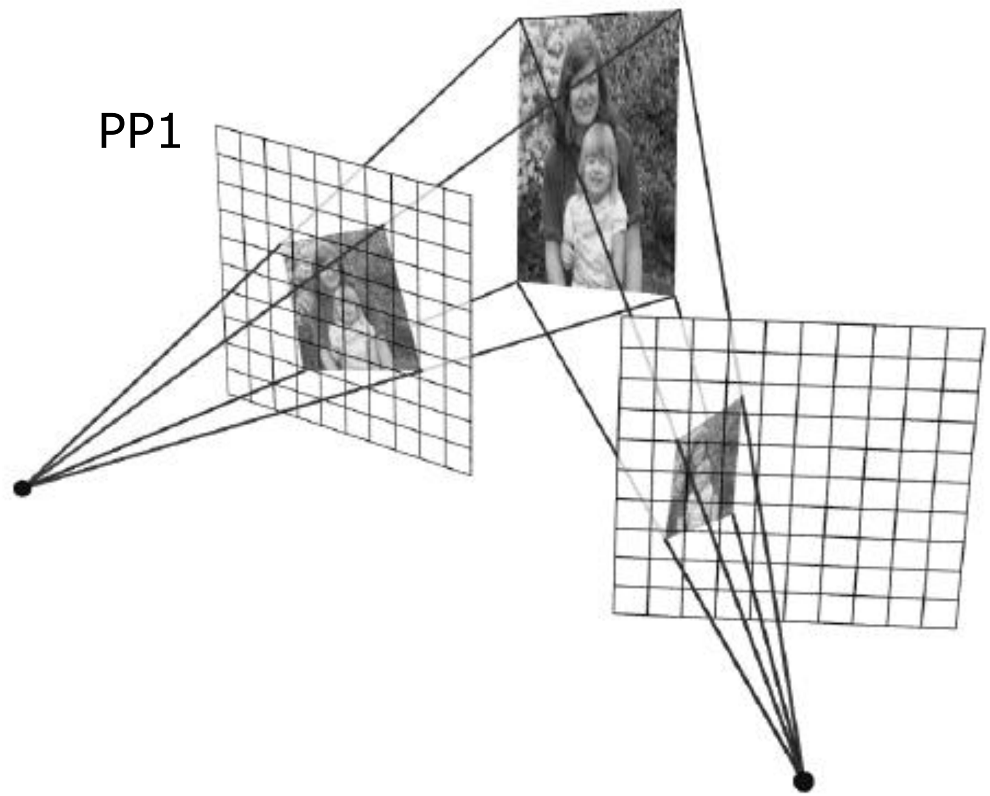

- PP3 is a projection plane of both centers of projection, so we are OK!
- This is how big aerial photographs are made

Credit: Y.Y. Chuang

Motion models

- Parametric models as the assumptions on the relation between two images.

Credit: Y.Y. Chuang

2D Motion models

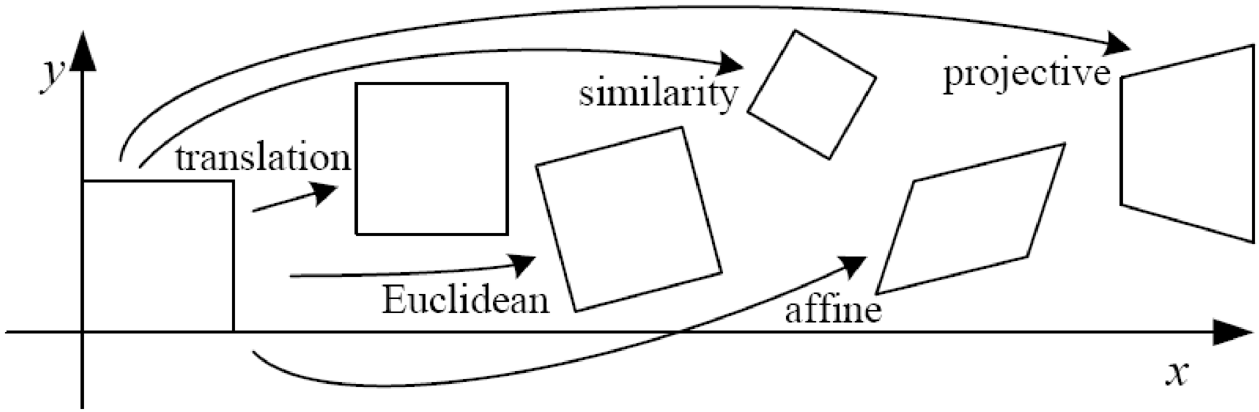

Name	Matrix	# D.O.F.	Preserves:	Icon
translation	$\begin{bmatrix} I \mid t \end{bmatrix}_{2\times 3}$	2	orientation $+\cdots$	
rigid (Euclidean)	$\begin{bmatrix} R \mid t \end{bmatrix}_{2\times 3}$	3	lengths $+\cdots$	
similarity	$\begin{bmatrix} sR \mid t \end{bmatrix}_{2\times 3}$	4	angles $+\cdots$	
affine	$\begin{bmatrix} A \end{bmatrix}_{2\times 3}$	6	parallelism $+\cdots$	
projective	$\begin{bmatrix} \tilde{H} \end{bmatrix}_{3\times 3}$	8	straight lines	

Credit: Y.Y. Chuang

Motion models

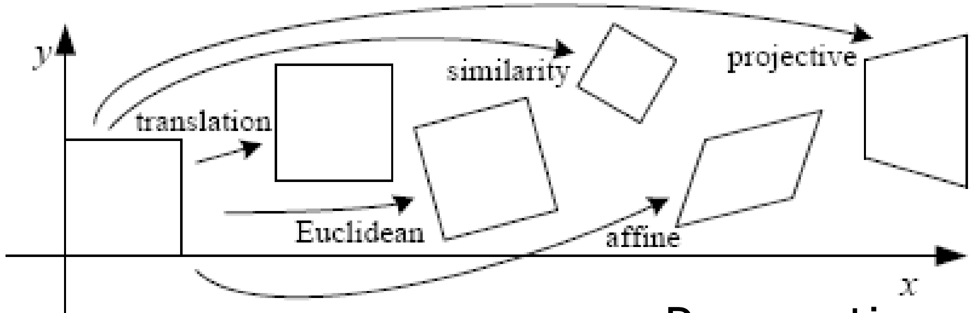

Translation

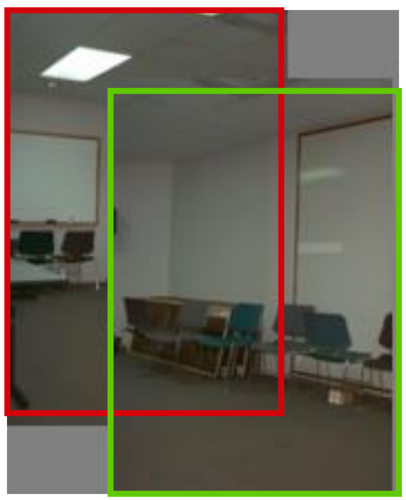

2 unknowns

Affine

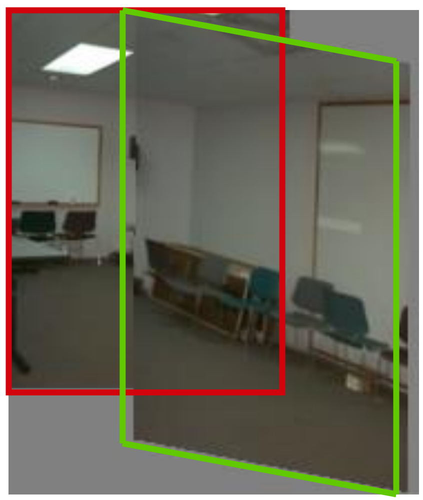

6 unknowns

Perspective

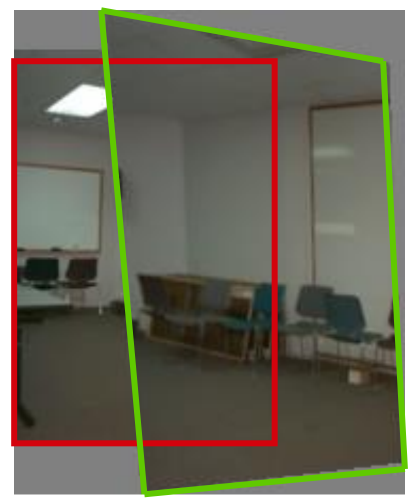

8 unknowns

3D rotation

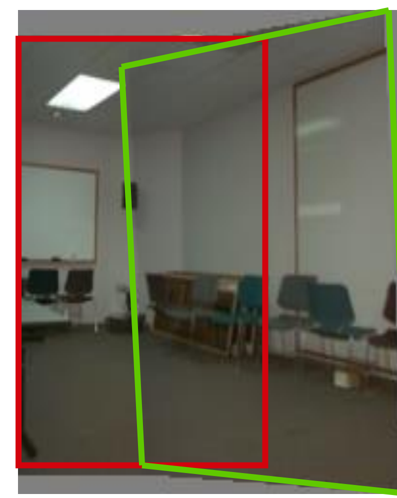

3 unknowns

Credit: Y.Y. Chuang

Determine pairwise alignment?

- Feature-based methods: only use feature points to estimate parameters
- We will study the “Recognising panorama” paper published in ICCV 2003
- Run SIFT (or other feature algorithms) for each image, find feature matches.

Credit: Y.Y. Chuang

Determine pairwise alignment

- p'=Mp, where M is a transformation matrix, p and p' are feature matches
- It is possible to use more complicated models such as affine or perspective
- For example, assume M is a 2x2 matrix

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}$$

- Find M with the least square error

$$\sum_{i=1}^{n} \left(Mp - p' \right)^2$$

Credit: Y.Y. Chuang

Determine pairwise alignment

$$\begin{pmatrix} x' \\ y' \end{pmatrix} = \begin{pmatrix} m_{11} & m_{12} \\ m_{21} & m_{22} \end{pmatrix} \begin{pmatrix} x \\ y \end{pmatrix}$$

$$x_1 m_{11} + y_1 m_{12} = x_1'$$
$$x_1 m_{21} + y_1 m_{22} = y_1'$$

$$\begin{pmatrix} x_1 & y_1 & 0 & 0 \\ 0 & 0 & x_1 & y_1 \\ x_2 & y_2 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots \\ x_n & y_n & 0 & 0 \\ 0 & 0 & x_n & y_n \end{pmatrix} \begin{pmatrix} m_{11} \\ m_{12} \\ m_{21} \\ m_{22} \end{pmatrix} = \begin{pmatrix} x_1' \\ y_1' \\ x_2' \\ \vdots \\ x_n' \\ y_n' \end{pmatrix}$$

Credit: Y.Y. Chuang

Over-determined system

Normal equation

Given an over-determined system

$$\mathbf{A}\mathbf{x} = \mathbf{b}$$

the normal equation is that which minimizes the sum of the square differences between left and right sides

$$\mathbf{A}^{\mathrm{T}}\mathbf{A}\mathbf{x} = \mathbf{A}^{\mathrm{T}}\mathbf{b}$$

Why?

Credit: Y.Y. Chuang

Normal equation

$$E = (\mathbf{A}\mathbf{x} - \mathbf{b})^2$$

$$= (\mathbf{A}\mathbf{x} - \mathbf{b})^T (\mathbf{A}\mathbf{x} - \mathbf{b})$$

$$= \left((\mathbf{A}\mathbf{x})^T - \mathbf{b}^T\right)(\mathbf{A}\mathbf{x} - \mathbf{b})$$

$$= \left(\mathbf{x}^{\mathrm{T}}\mathbf{A}^{\mathrm{T}} - \mathbf{b}^{\mathrm{T}}\right)(\mathbf{A}\mathbf{x} - \mathbf{b})$$

$$= \mathbf{x}^{\mathrm{T}}\mathbf{A}^{\mathrm{T}}\mathbf{A}\mathbf{x} - \mathbf{b}^{\mathrm{T}}\mathbf{A}\mathbf{x} - \mathbf{x}^{\mathrm{T}}\mathbf{A}^{\mathrm{T}}\mathbf{b} + \mathbf{b}^{\mathrm{T}}\mathbf{b}$$

$$= \mathbf{x}^{\mathrm{T}}\mathbf{A}^{\mathrm{T}}\mathbf{A}\mathbf{x} - (\mathbf{A}^{\mathrm{T}}\mathbf{b})^{\mathrm{T}}\mathbf{x} - (\mathbf{A}^{\mathrm{T}}\mathbf{b})^{\mathrm{T}}\mathbf{x} + \mathbf{b}^{\mathrm{T}}\mathbf{b}$$

$$\frac{\partial E}{\partial \mathbf{x}} = 2\mathbf{A}^{\mathrm{T}}\mathbf{A}\mathbf{x} - 2\mathbf{A}^{\mathrm{T}}\mathbf{b}$$

Credit: Y.Y. Chuang

Determine pairwise alignment

- p'=Mp, where M is a transformation matrix, p and p' are feature matches
- For translation model, it is easier.

$$E = \sum_{i=1}^{n} \left[\left(m_1 + x_i - x_i^{'} \right)^2 + \left(m_2 + y_i - y_i^{'} \right)^2 \right]$$

$$0 = \frac{\partial E}{\partial m_1}$$

Credit: Y.Y. Chuang

Stitch with Homography

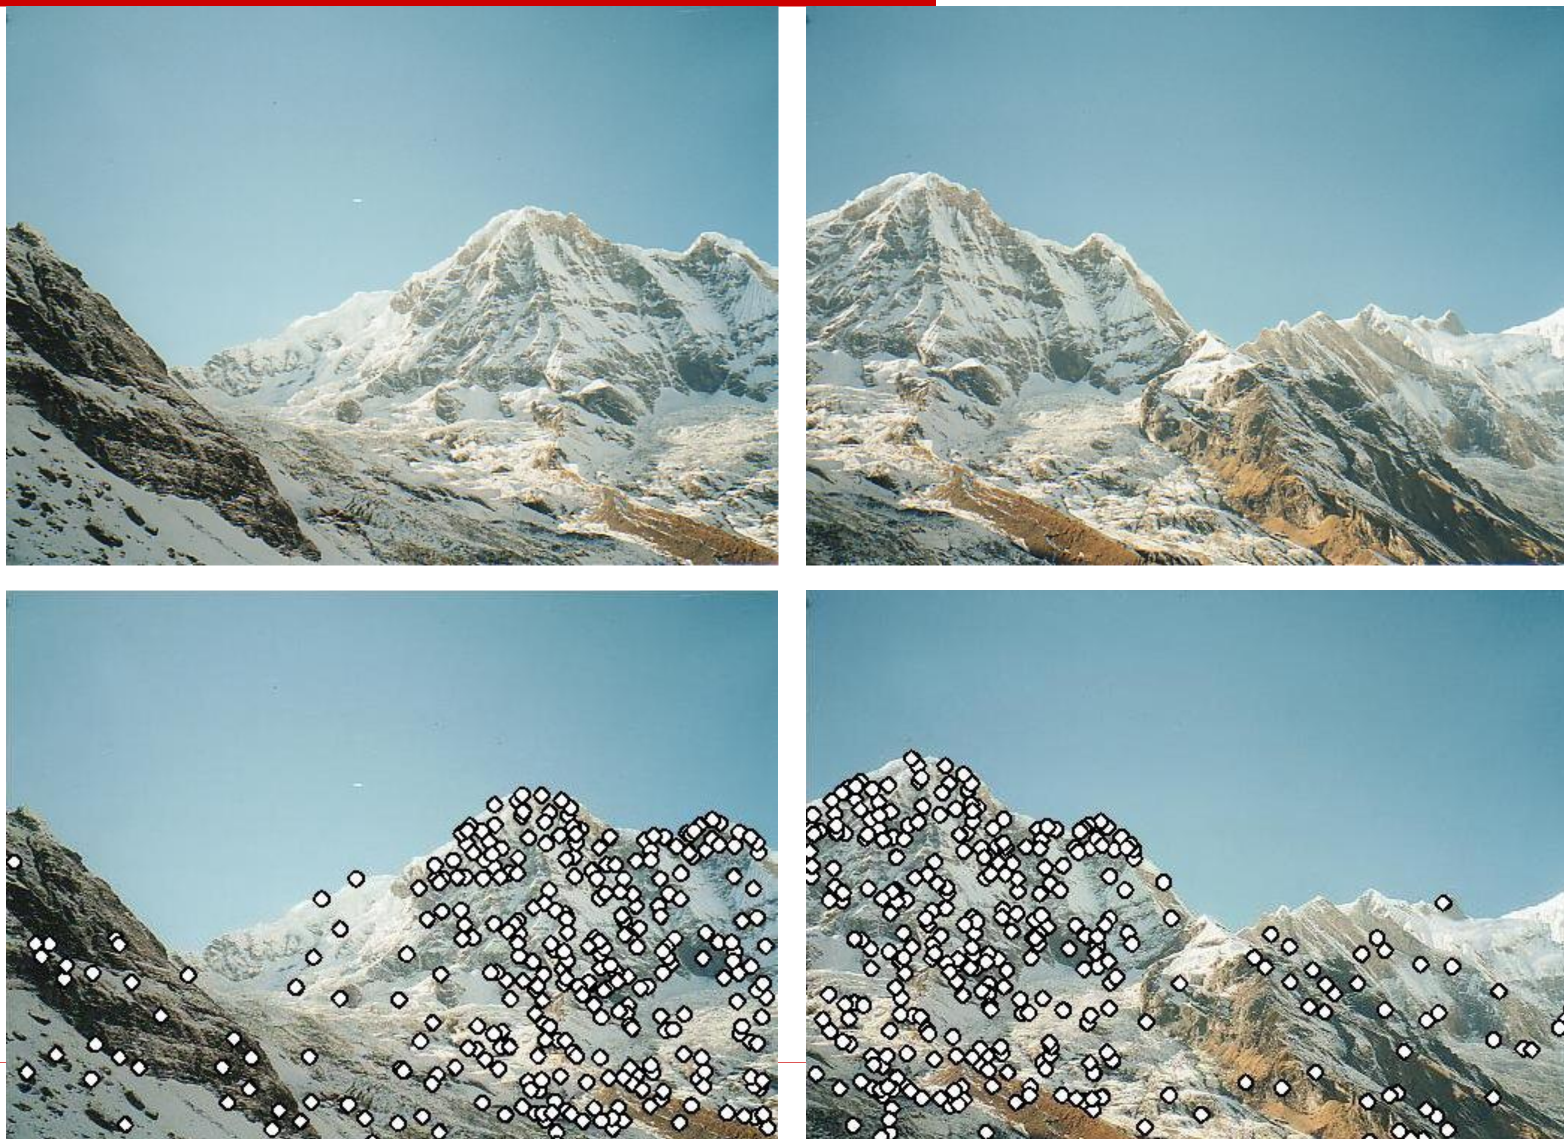

Credit: Y.Y. Chuang

Stitch with Homography

Credit: Y.Y. Chuang

Stitching Recipe

- Align pairs of images
  - Feature Detection
  - Feature Matching
  - Homography Estimation

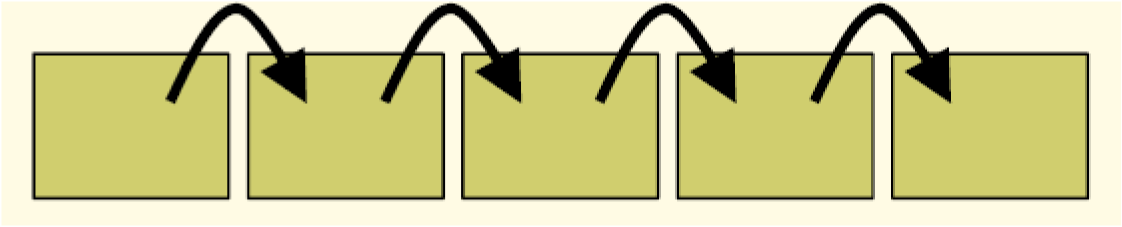

- Align all to a common surface
- Adjust (Global) & Blend

A case study: cylindrical panorama

- What if you want a 360° field of view?

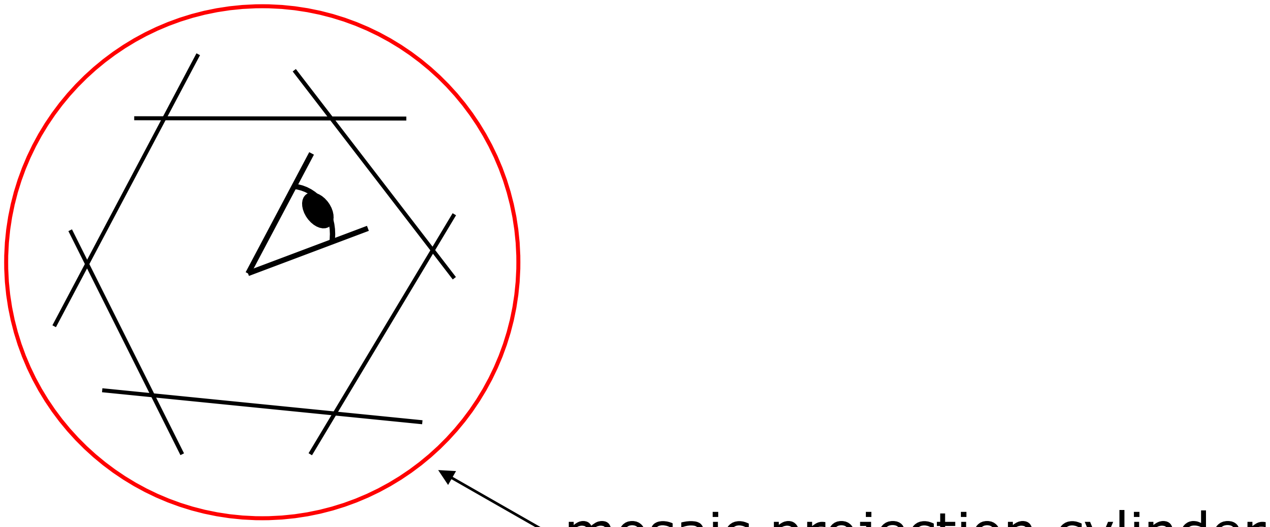

Credit: Y.Y. Chuang

Cylindrical panoramas

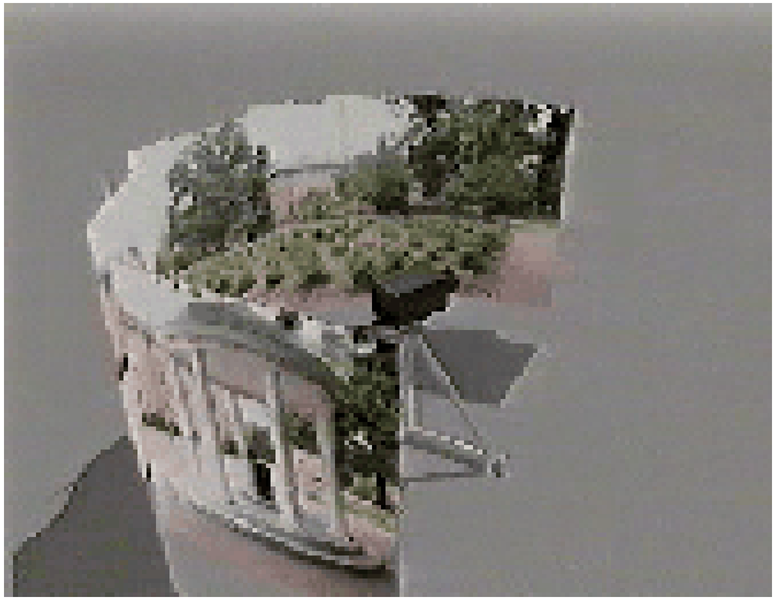

- Steps
  - Reproject each image onto a cylinder
  - Blend
  - Output the resulting mosaic

Credit: Y.Y. Chuang

Cylindrical panorama

1. Take pictures on a tripod (or handheld)
2. Warp to cylindrical coordinate
3. Compute pair-wise alignments
4. Fix up the end-to-end alignment
5. Blending
6. Crop the result and import into a viewer

It is required to do radial distortion correction for better stitching results!

Credit: Y.Y. Chuang

Taking pictures

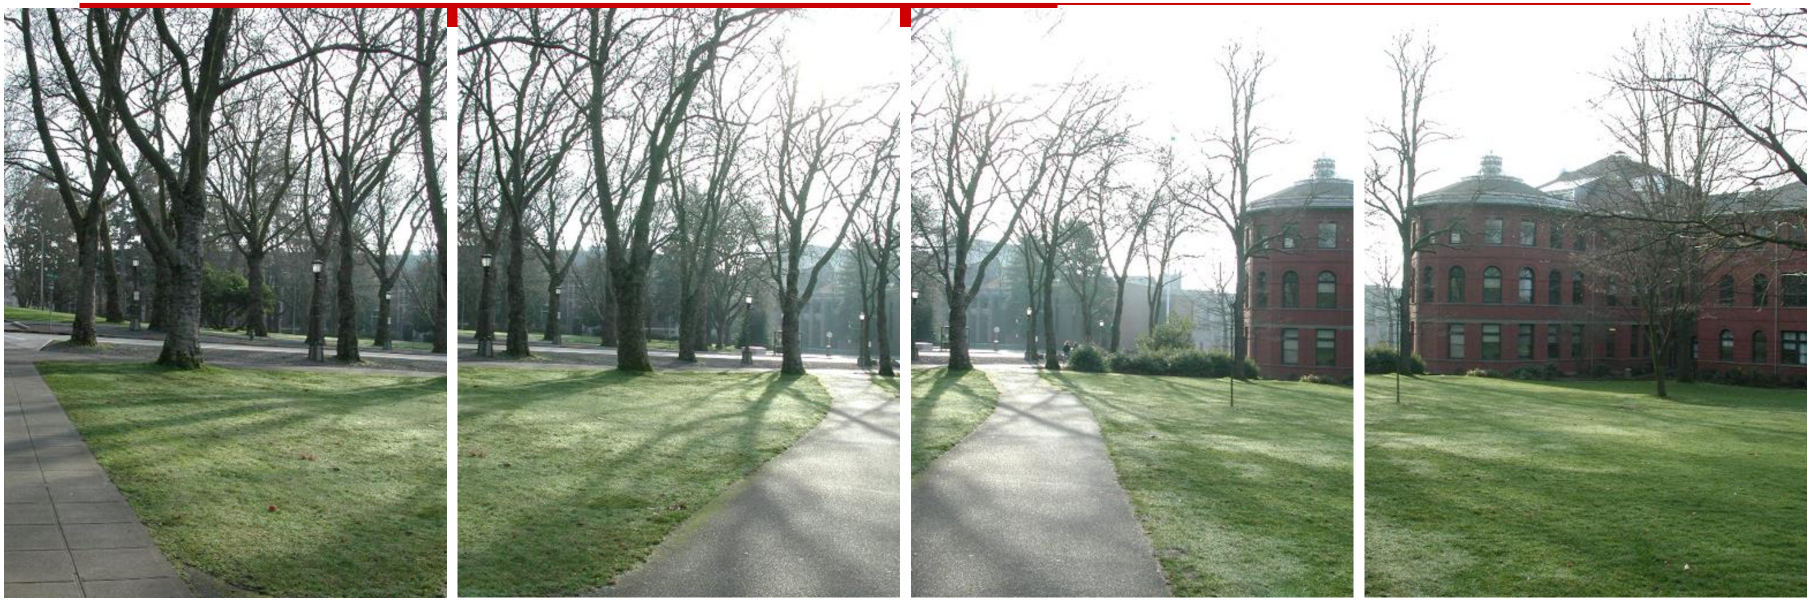

Kaidan panoramic tripod head

Credit: Y.Y. Chuang

Where should the synthetic camera be

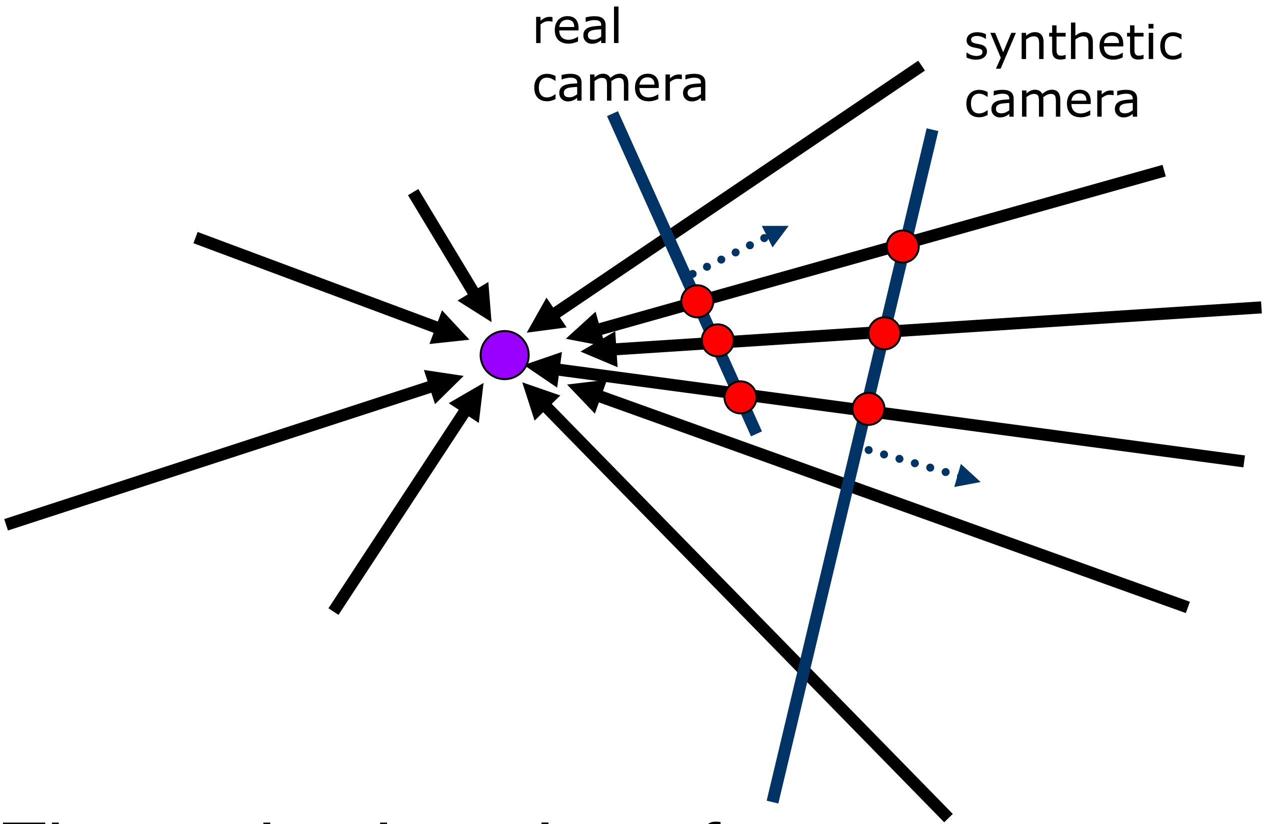

- The projection plan of some camera
- Onto a cylinder

Credit: Y.Y. Chuang

Rectlinear projection

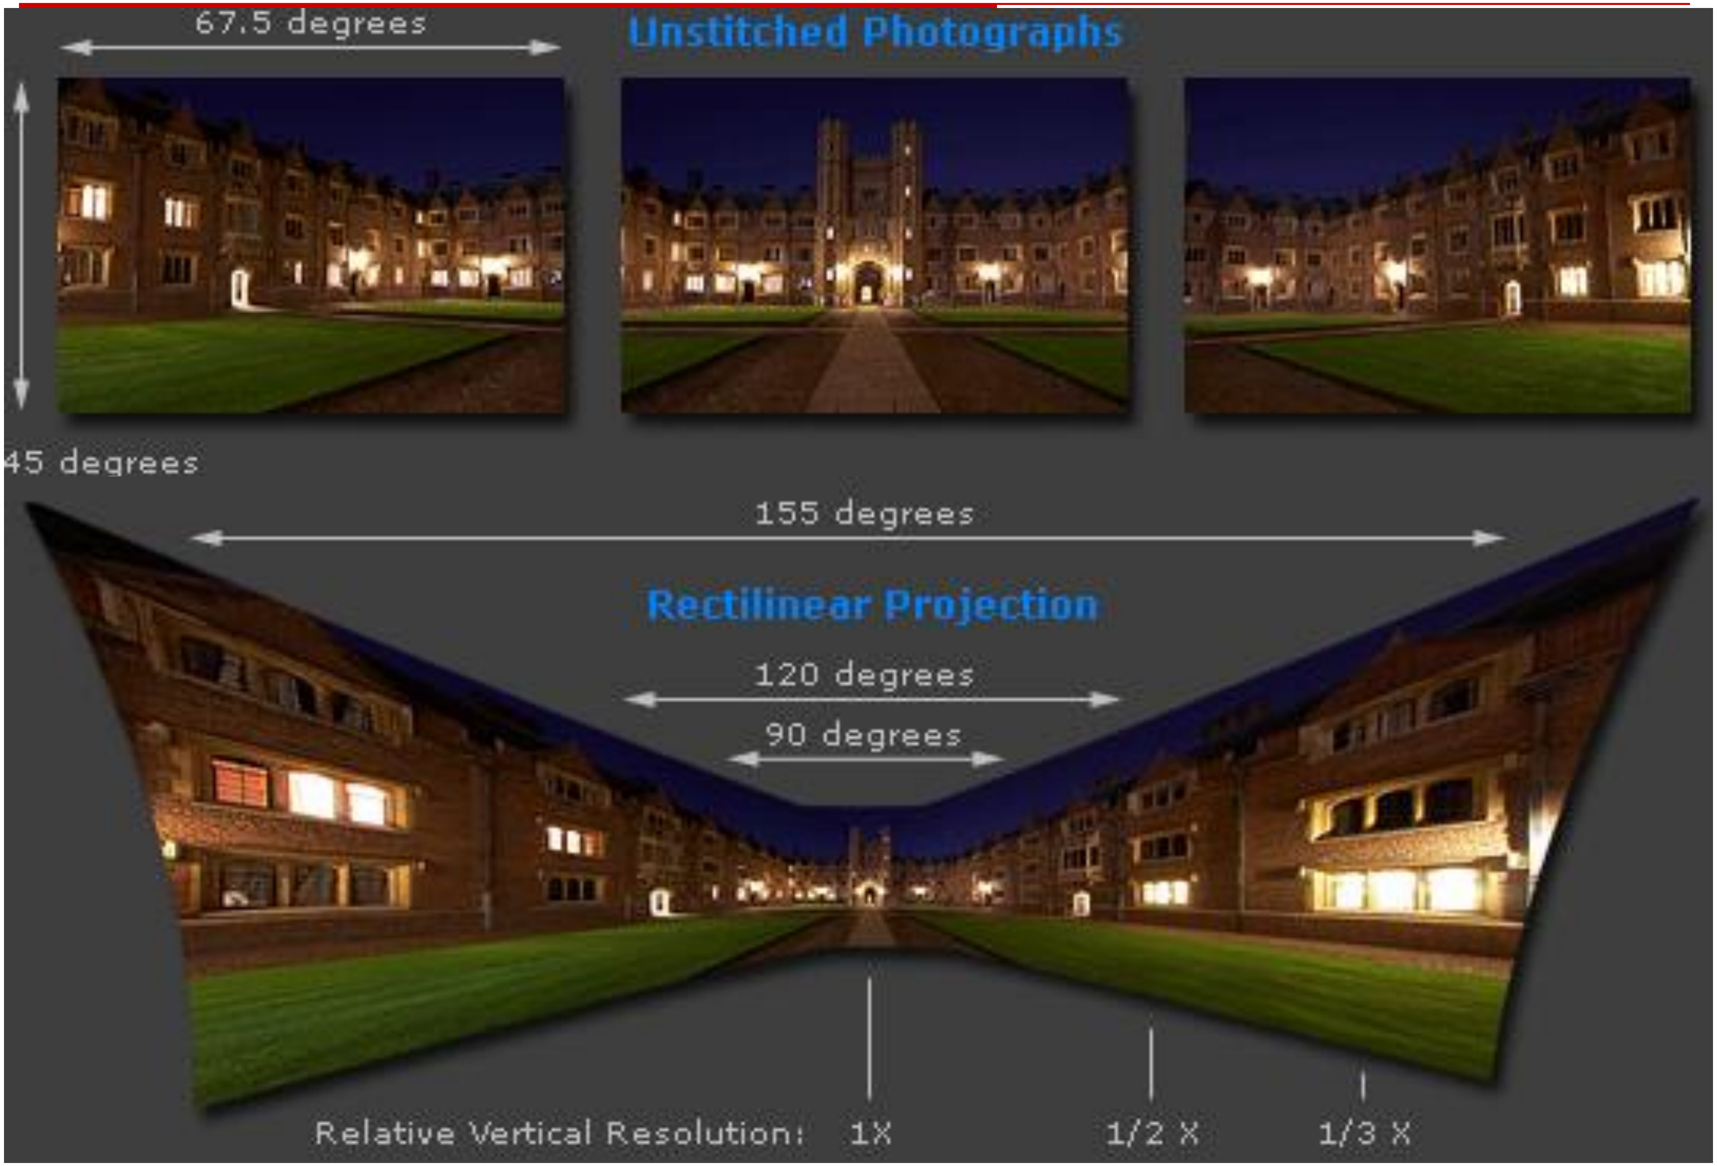

Credit: Y.Y. Chuang

Adopted from http://www.cambridgeincolour.com/tutorials/image-projections.htm

Cylindrical projection

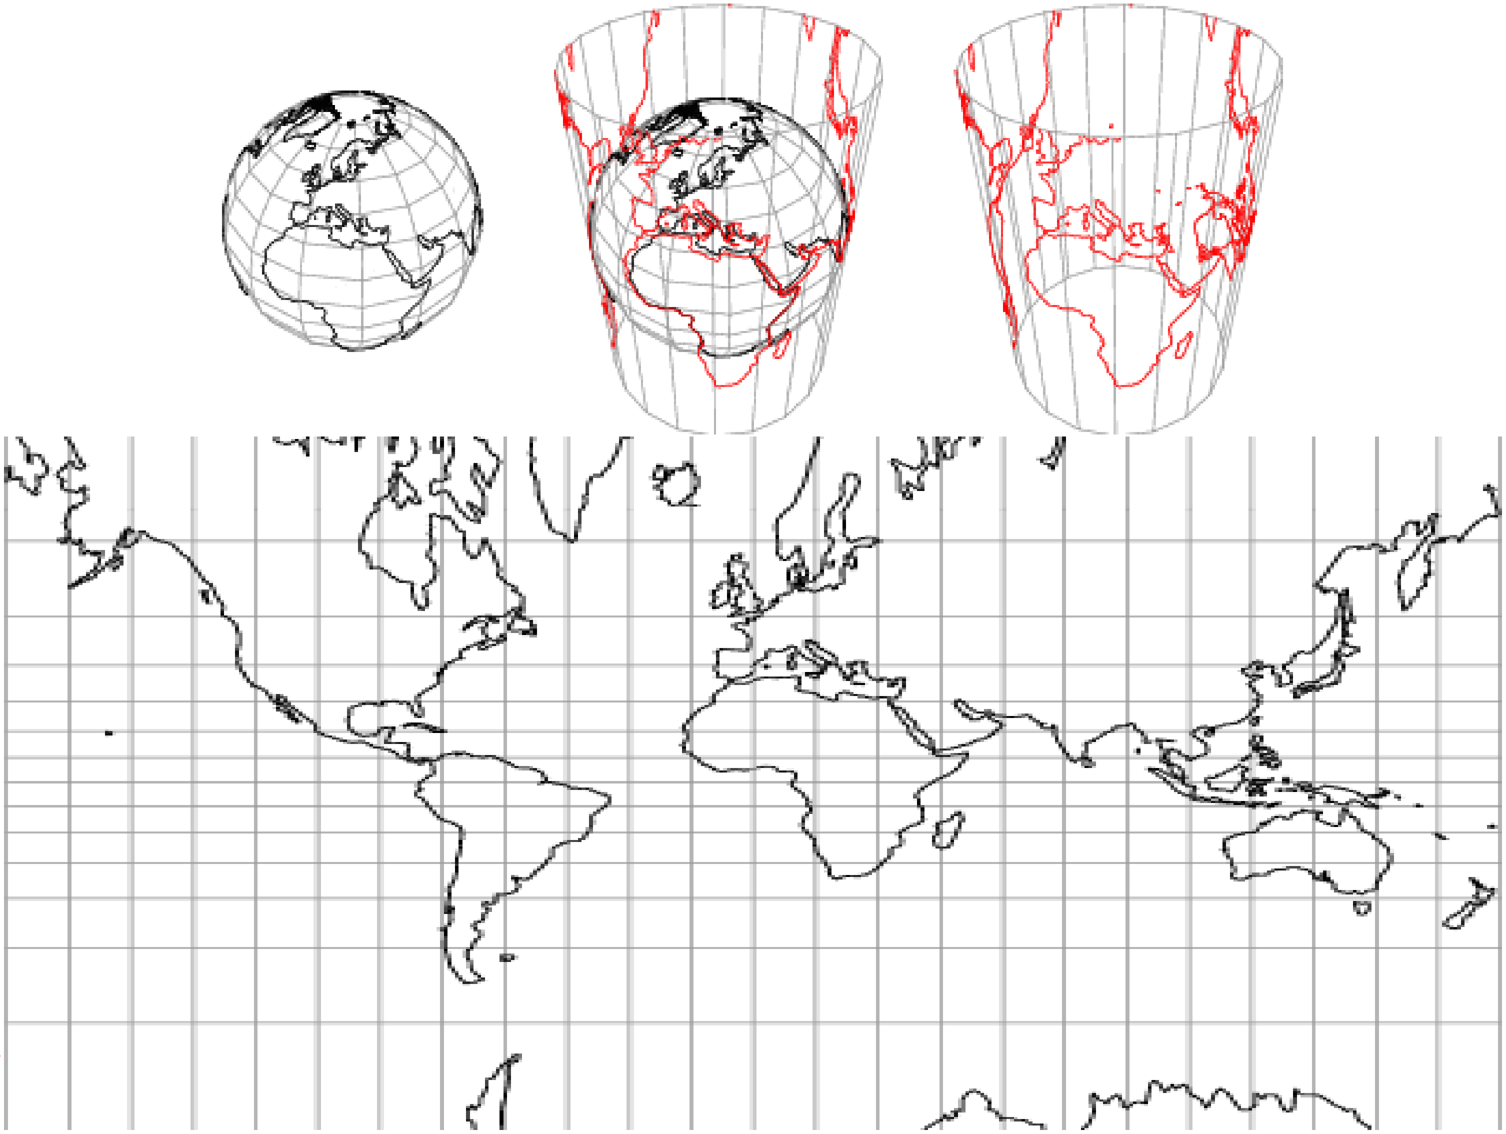

Credit: Y.Y. Chuang

Cylindrical projection

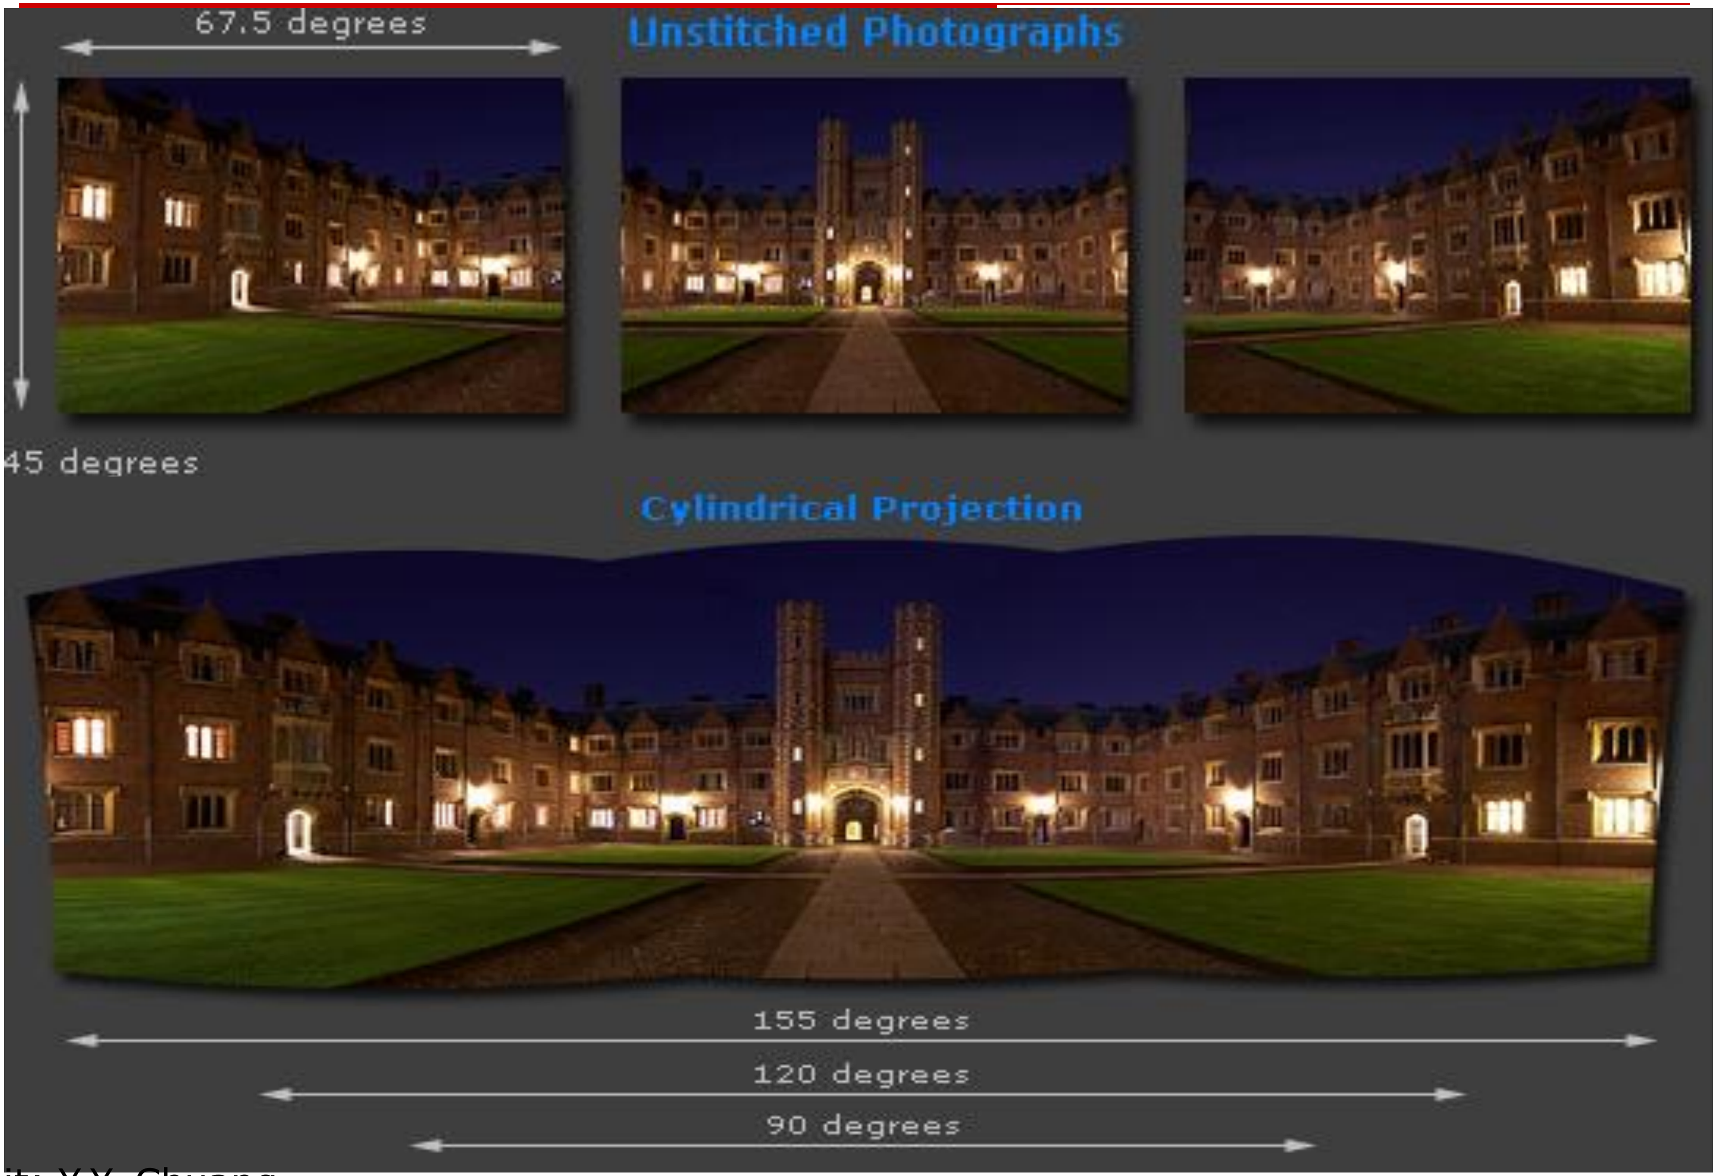

Credit: Y.Y. Chuang

Adopted from http://www.cambridgeincolour.com/tutorials/image-projections.htm

Cylindrical projection

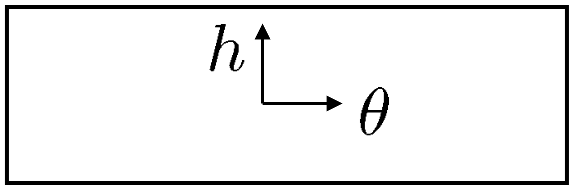

unwrapped cylinder

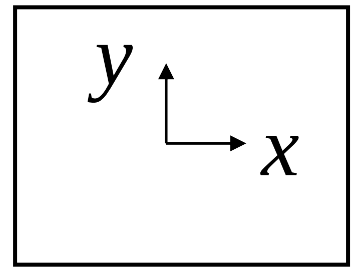

$$(\sin\theta, h, \cos\theta) \propto (x, y, f)$$

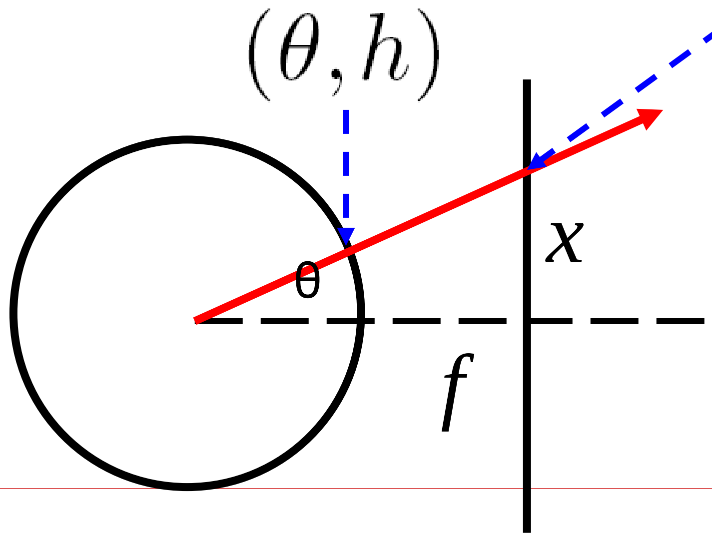

$$\theta = \tan^{-1}\frac{x}{f}$$

Credit: Y.Y. Chuang

Cylindrical projection

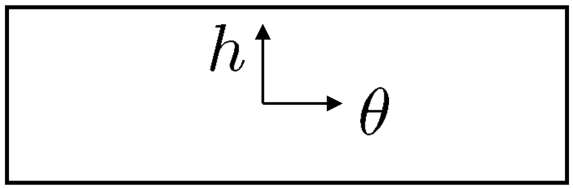

unwrapped

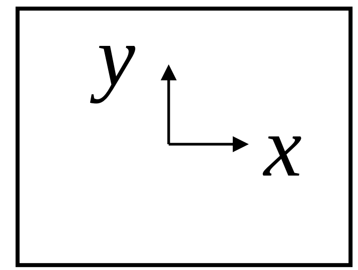

$$(\sin\theta, h, \cos\theta) \propto (x, y, f)$$

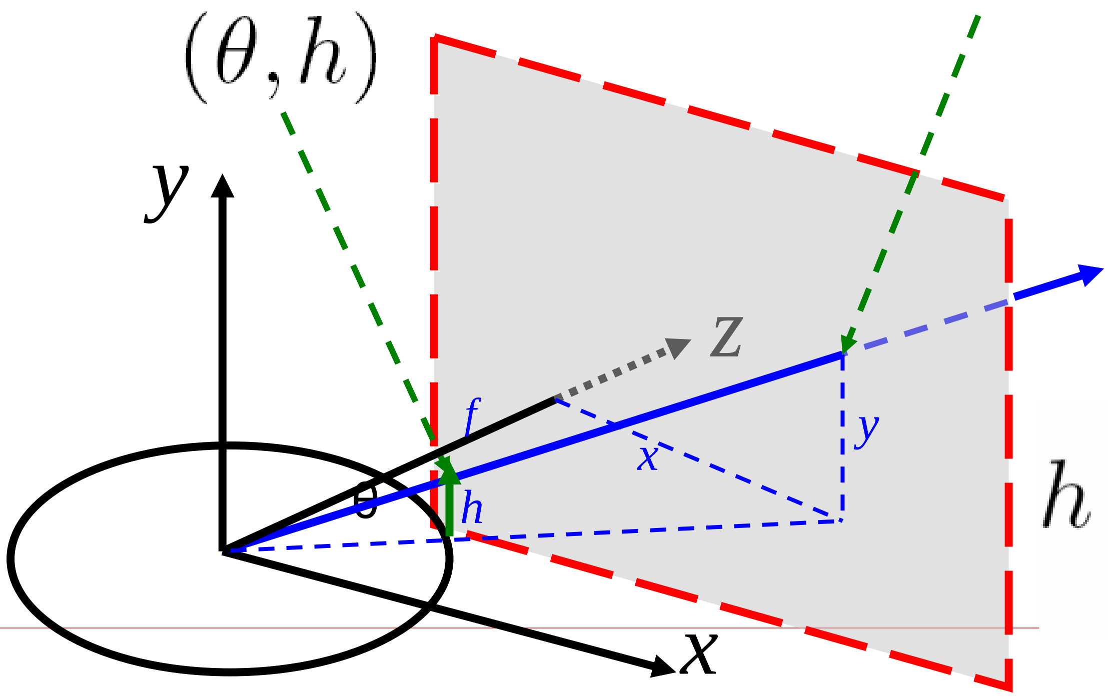

$$h = \frac{y}{\sqrt{x^2 + f^2}}$$

Credit: Y.Y. Chuang

Cylindrical projection

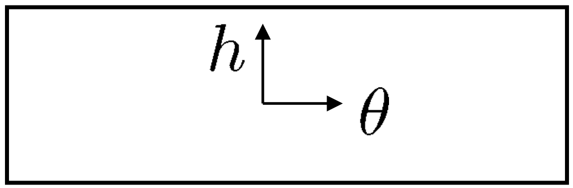

unwrapped
cylinder

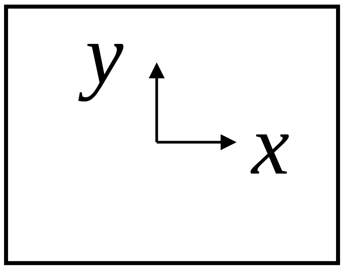

$$x' = s\theta = s \tan^{-1} \frac{x}{f}$$

$$y' = sh = s \frac{y}{\sqrt{x^2 + f^2}}$$

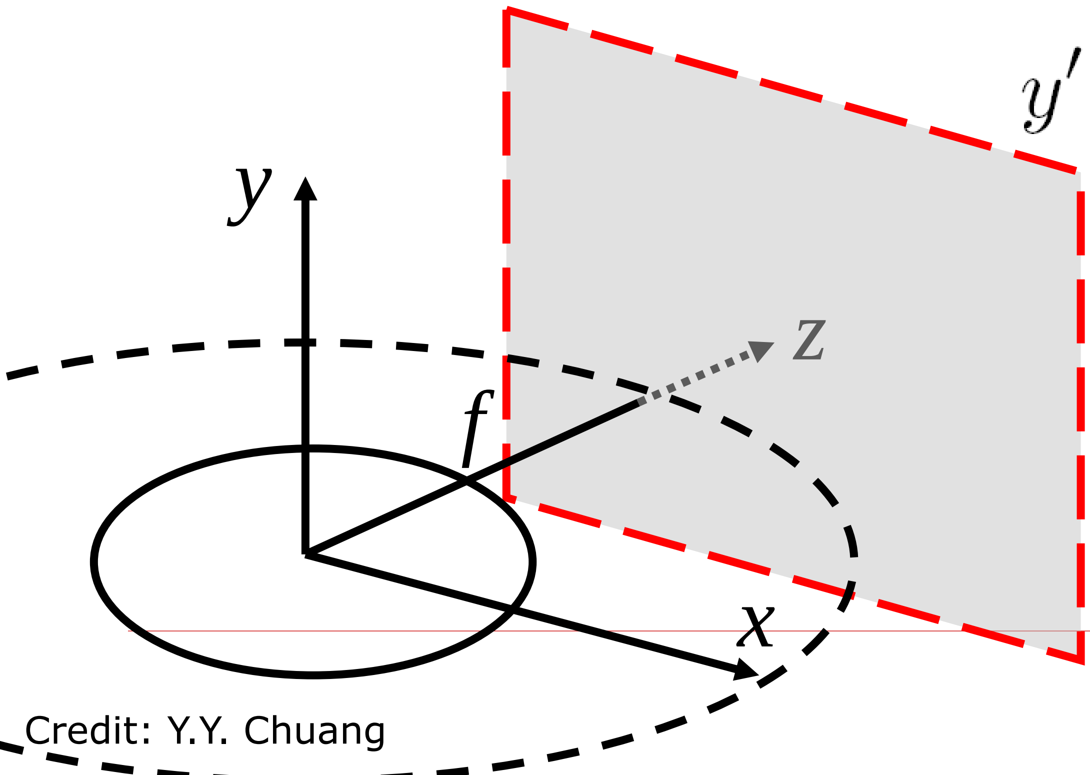

s defines size of the final image, often convenient to set s = f

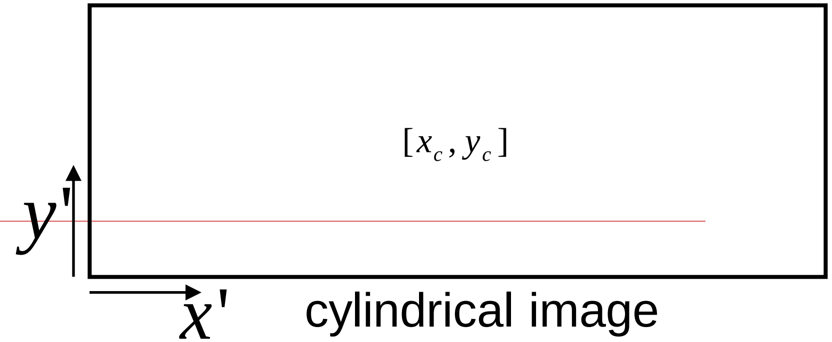

cylindrical image

Credit: Y.Y. Chuang

Cylindrical reprojection

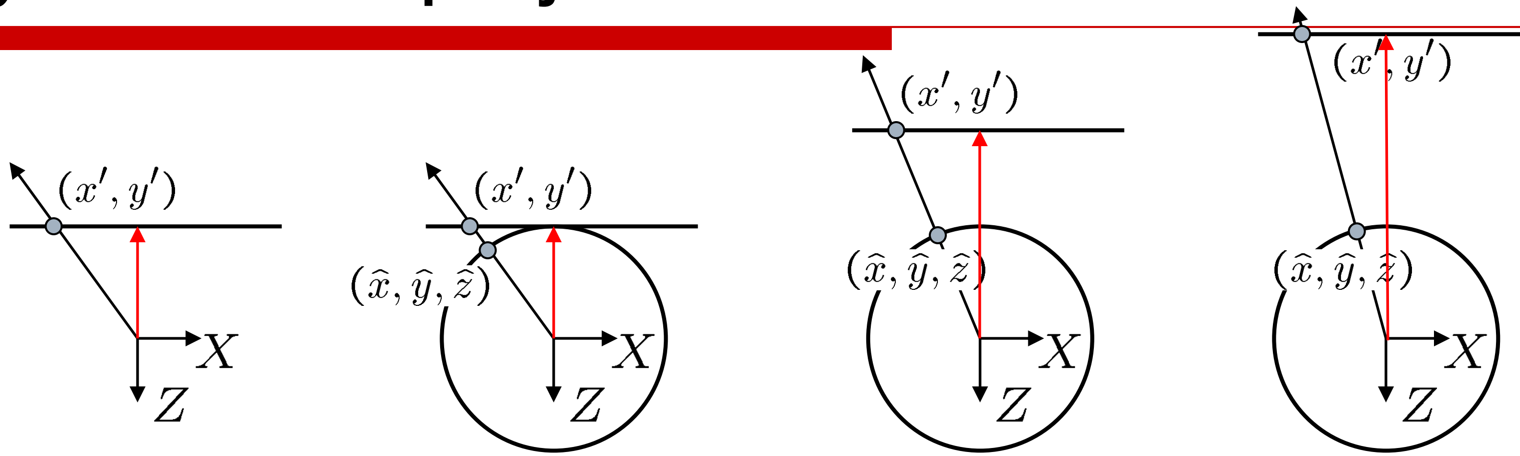

top-down view

Focal length – the dirty secret…

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

Credit: Y.Y. Chuang

A simple method for estimating f

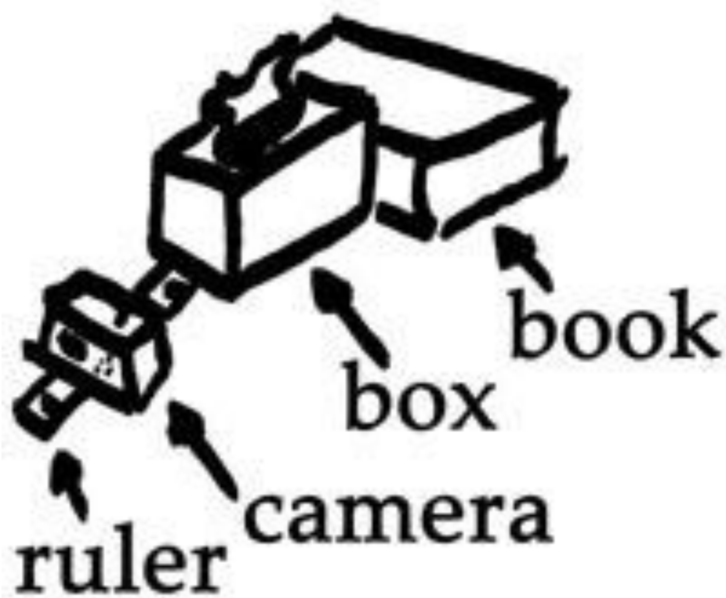

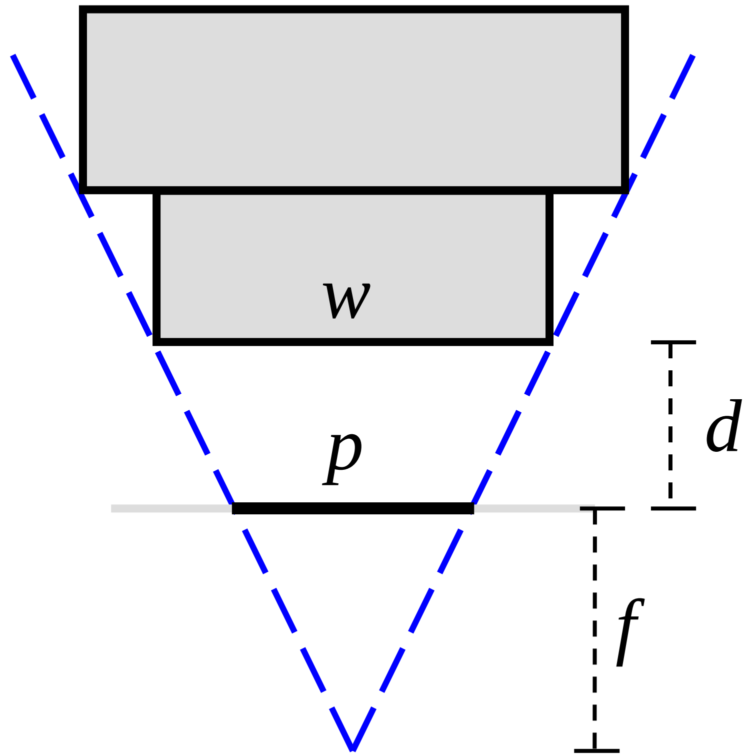

Or, you can use other software, such as AutoStich, to help.

Credit: Y.Y. Chuang

Input images

Credit: Y.Y. Chuang

Cylindrical warping

Credit: Y.Y. Chuang

Alignment

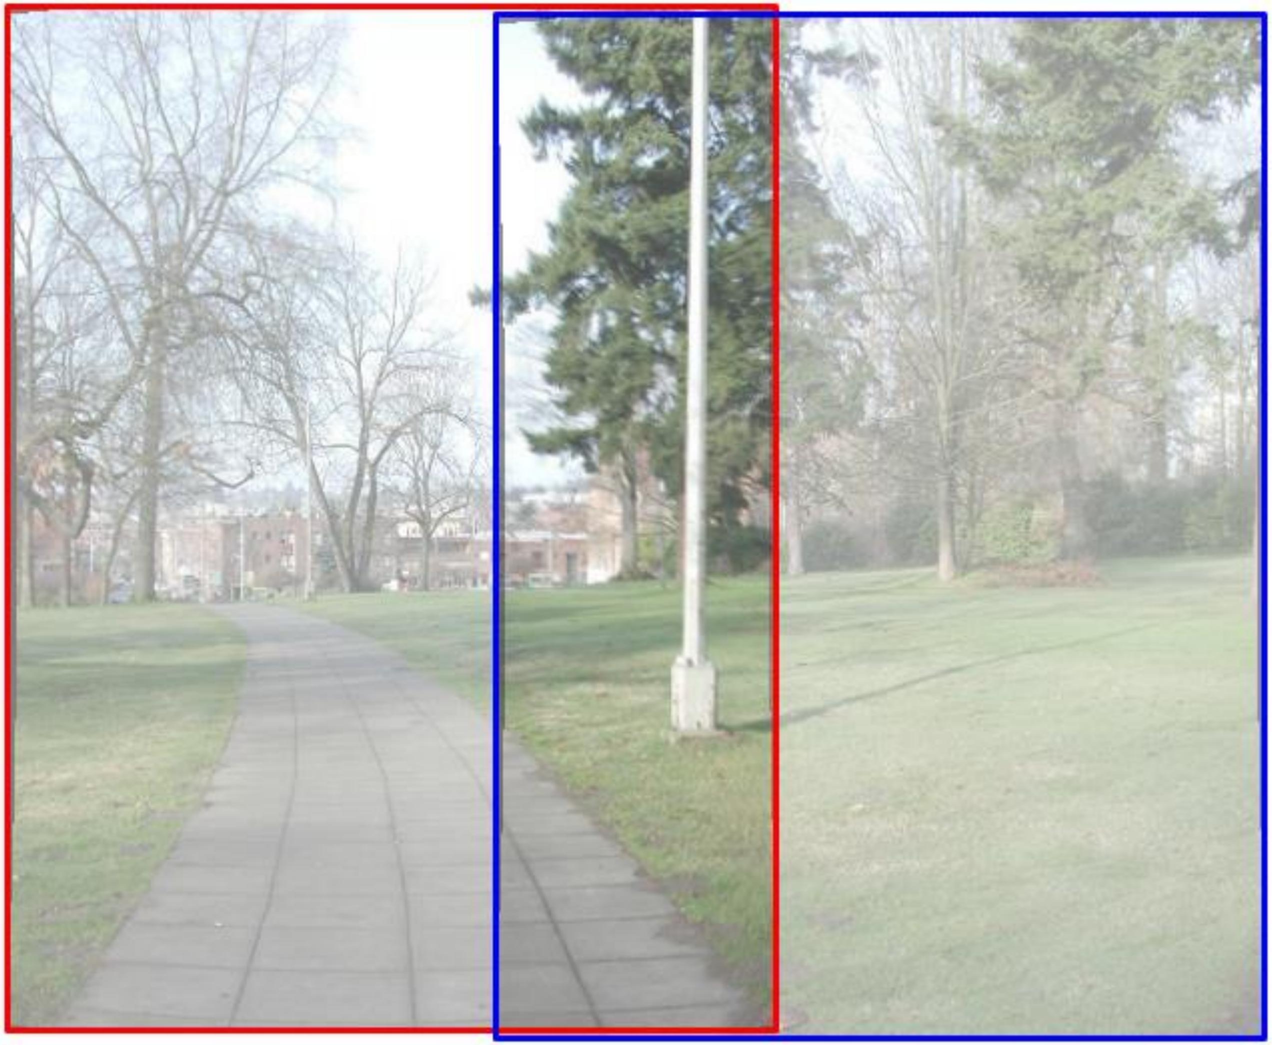

a rotation of the camera is a **translation** of the cylinder!

Credit: Y.Y. Chuang

Blending

- Why blending: parallax, lens distortion, scene motion, exposure difference
- Alpha-blending
- Poisson blending
- Adelson's pyramid blending

Blending

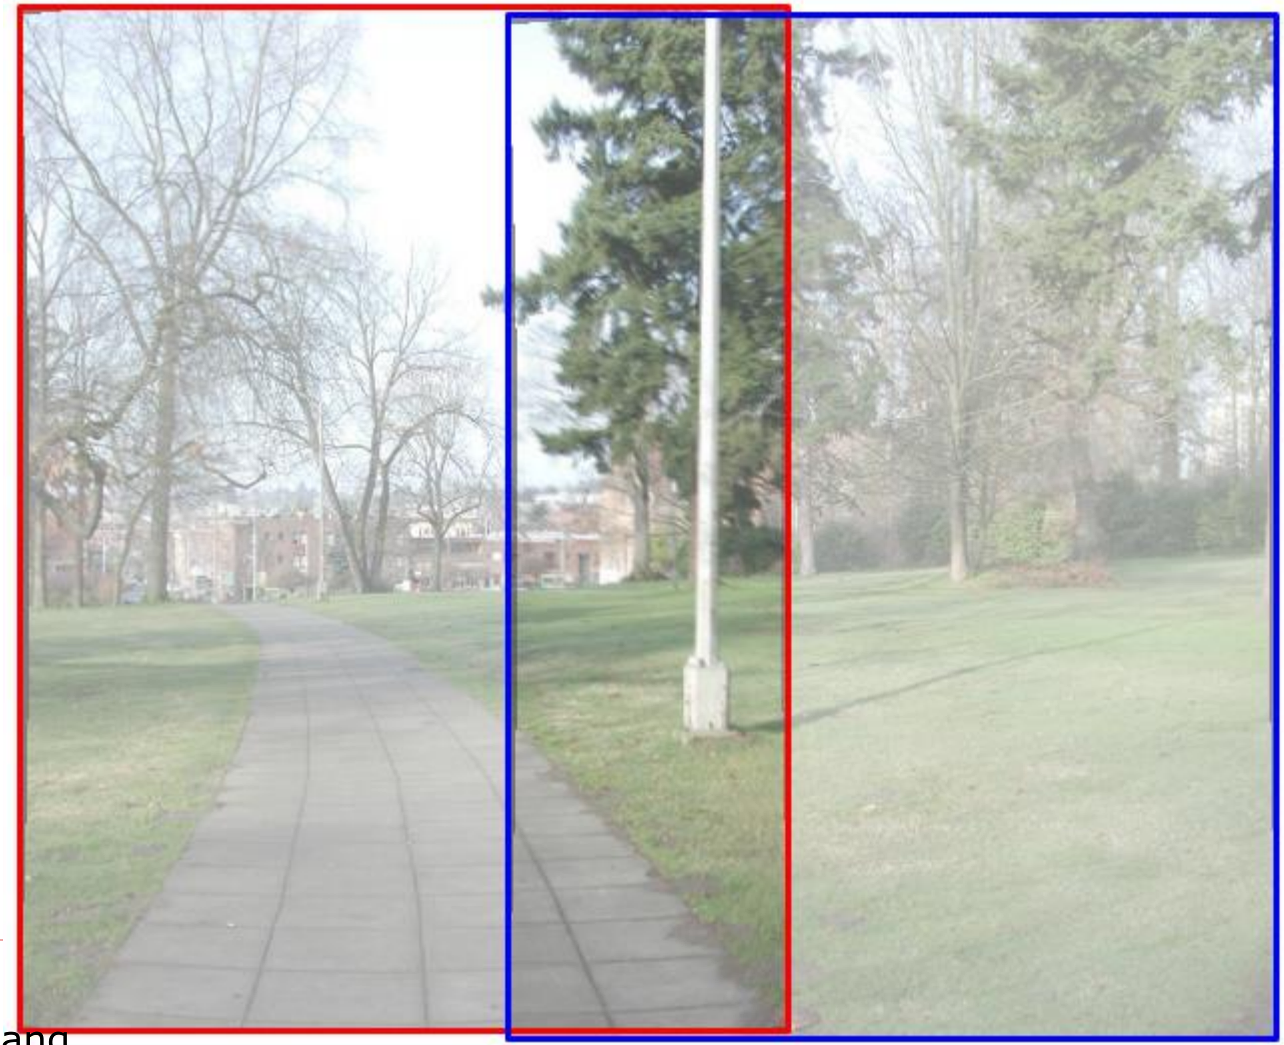

Credit: Y.Y. Chuang

Linear Blending

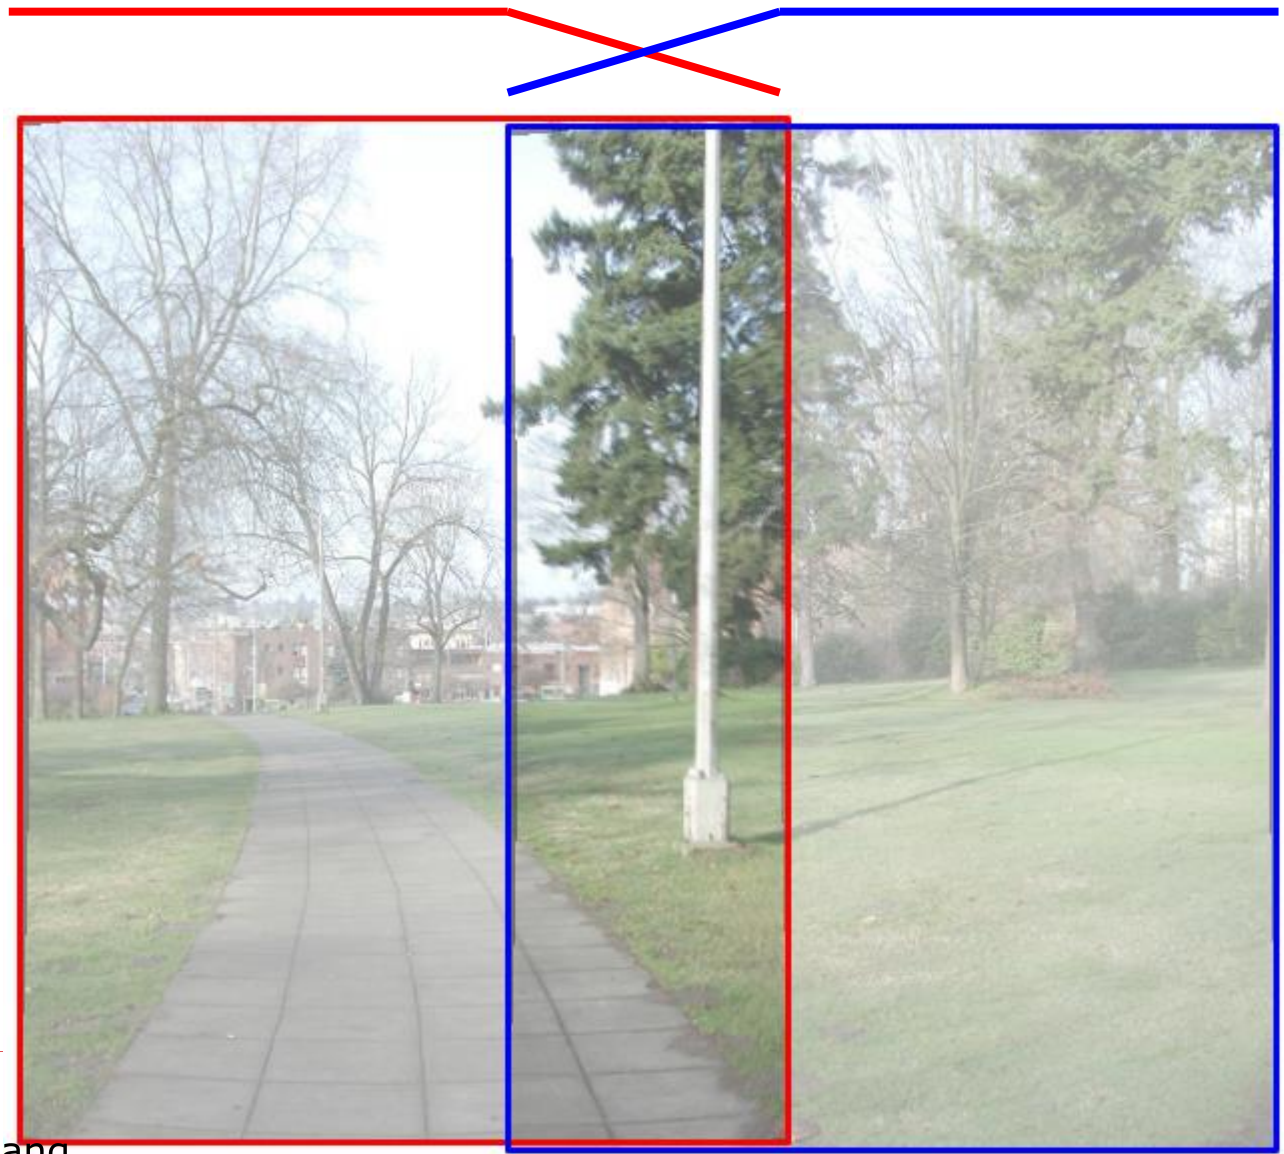

Credit: Y.Y. Chuang

Linear Blending

Credit: Y.Y. Chuang

Linear Blending

Image from http://www.cs.ubc.ca/~lowe/425/slides/11-PanoramasAR.pdf

Multi-band Blending [BURT and ADELSON 83]

Linear blending

Multi-band blending

A multi-resolution spline with application to image mosaics.
Peter J. Burt and Edward Adelson. ACM Transactions on Graphics, 1983.

Multi-band Blending

1. Laplacian pyramids LA and LB are constructed for images A and B respectively.
2. A third Laplacian pyramid LS is constructed by copying nodes from the left half of LA to the corresponding nodes of LS, and nodes in the right half of LB to the right half of LS
3. The final image S is obtained by expanding and summing the levels of LS.

Multi-band Blending

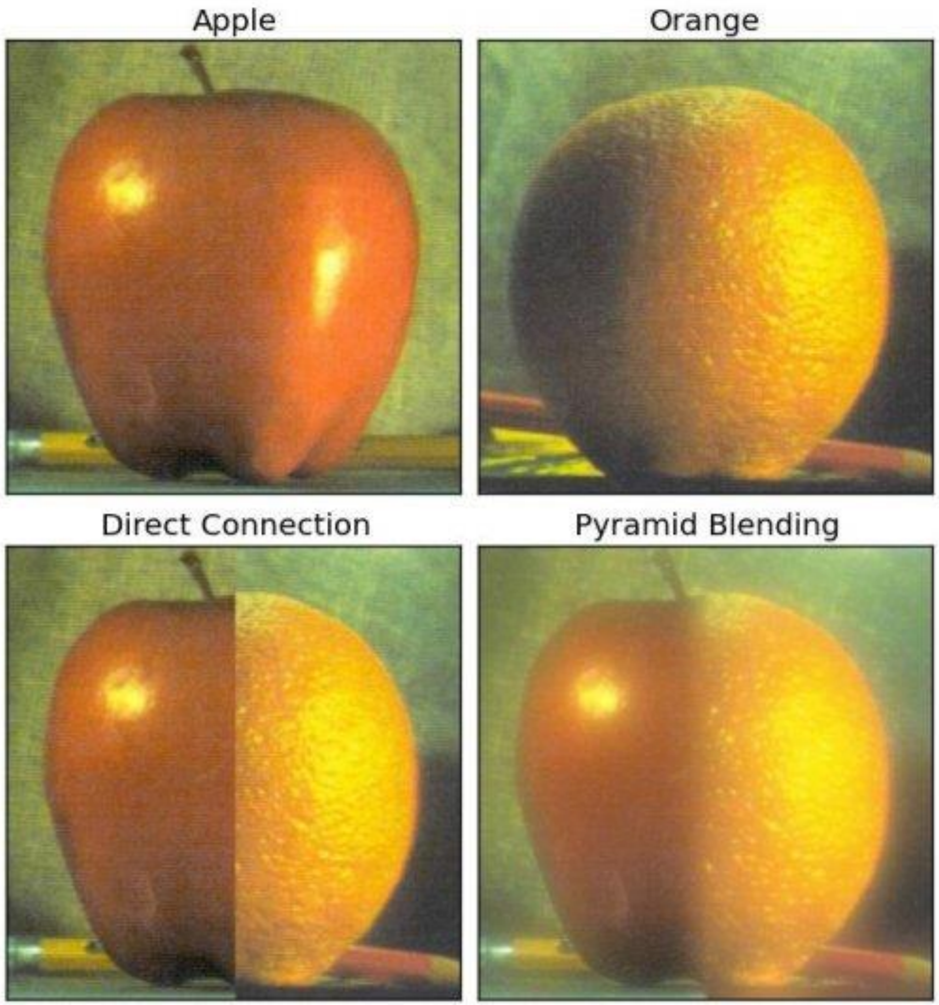

Source: https://docs.opencv.org/3.1.0/dc/dff/tutorial_py_pyramids.html

2-band Blending

Low frequency (l > 2 pixels)

High frequency (l < 2 pixels)

Credit: Y.Y. Chuang

Linear Blending

2-band Blending

Assembling the panorama

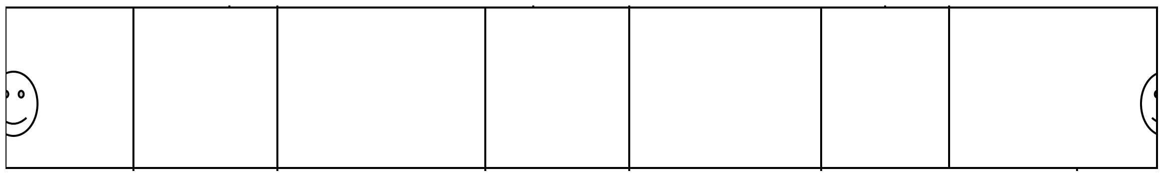

- Stitch pairs together, blend, then crop

Credit: Y.Y. Chuang

Problem: Drift

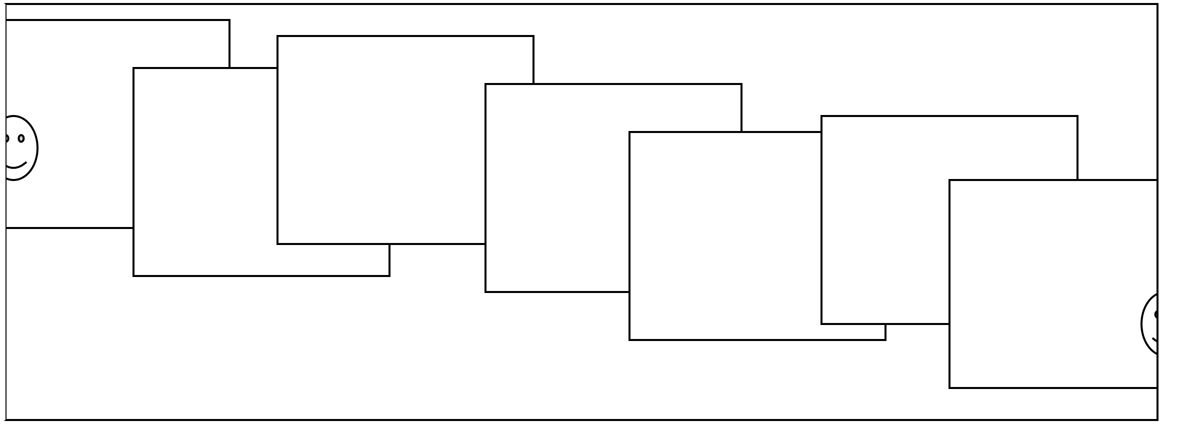

- Error accumulation
  - small errors accumulate over time

Credit: Y.Y. Chuang

Problem: Drift

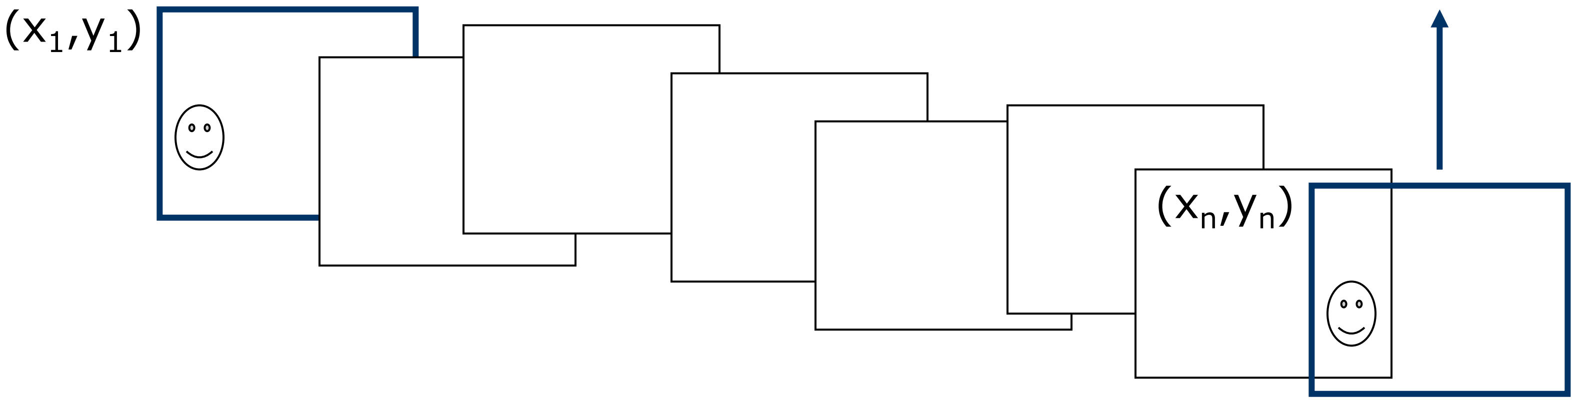

- copy of first image

- Solution
  - add another copy of first image at the end
  - there are a bunch of ways to solve this problem
    - add displacement of $(y_1 – y_n)/(n -1)$ to each image after the first
    - compute a global warp: $y' = y + ax$
    - run a big optimization problem, incorporating this constraint
      - best solution, but more complicated
      - known as “bundle adjustment”

Credit: Y.Y. Chuang

End-to-end alignment and crop

Credit: Y.Y. Chuang

Viewer: panorama

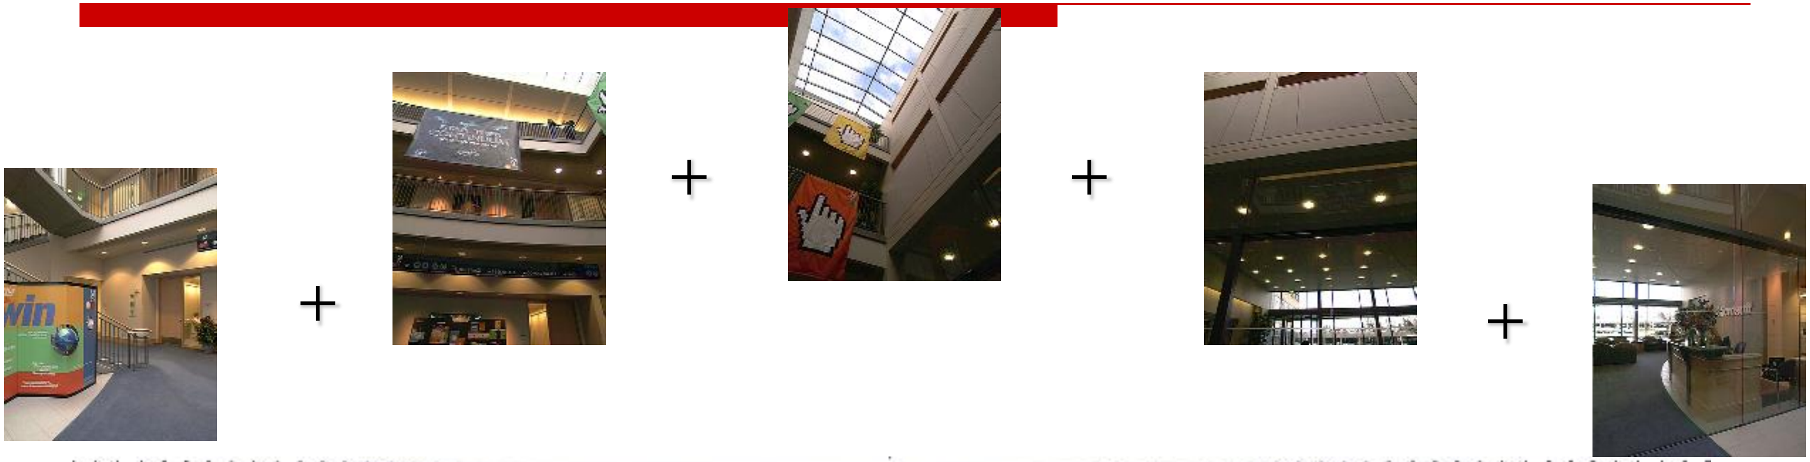

example:

http://www.cs.washington.edu/education/courses/cse590ss/01wi/projects/project1/students/dougz/index.html

Credit: Y.Y. Chuang

Student paper presentation

Color Image Colorization

R. Zhang, P. Isola, and A. Efros
ECCV 2016

Presenter: Seward, Garrett

Student paper presentation

Burst photography for high dynamic range and low-light imaging on mobile cameras

Samuel W. Hasinoff, Dillon Sharlet, Ryan Geiss, Andrew Adams, Jonathan T. Barron, Florian Kainz, Jiawen Chen, and Marc Levoy
SIGGRAPH Asia 2016

Presenter: Peters, Emerson

Next Time

- Image segmentation
- Student paper presentations
  - 05/03: Little, Samuel
    - Deep High Dynamic Range Imaging of Dynamic Scenes. N. K. Kalantari and R. Ramamoorthi, SIGGRAPH 2017
  - 05/03: Joshi, Vijay
    - Night Sight: Seeing in the Dark on Pixel Phones
      - https://ai.googleblog.com/2018/11/night-sight-seeing-in-dark-on-pixel.html
      - https://www.blog.google/products/pixel/see-light-night-sight/