Computer Vision and Machine Learning Winter School

ENS Lyon 2010

Camera geometry and image alignment

Josef Sivic

http://www.di.ens.fr/~josef

INRIA, WILLOW, ENS/INRIA/CNRS UMR 8548

Laboratoire d'Informatique, Ecole Normale Supérieure, Paris

With slides from: O. Chum, K. Grauman, S. Lazebnik, B. Leibe, D. Lowe, J. Philbin, J. Ponce, D. Nister, C. Schmid, N. Snavely, A. Zisserman

Outline

Part I - Camera geometry – image formation

- Perspective projection
- Affine projection
- Projection of planes

Part II - Image matching and recognition with local features

- Correspondence
- Semi-local and global geometric relations
- Robust estimation – RANSAC and Hough Transform

Motivation: Stitching panoramas

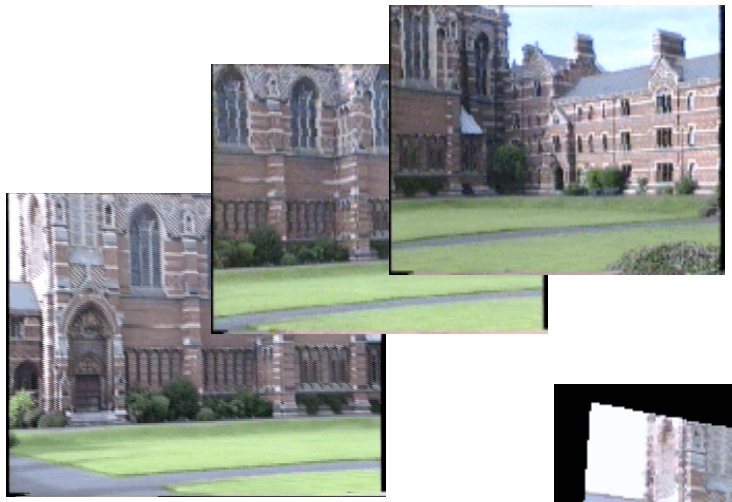

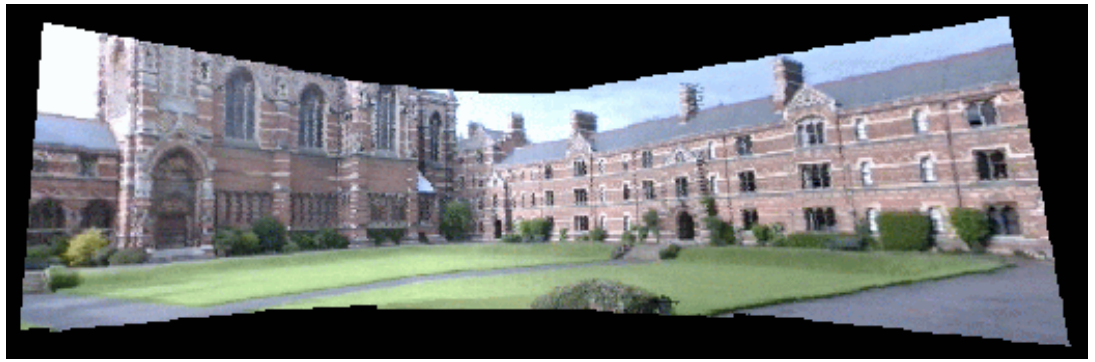

Feature-based alignment outline

Feature-based alignment outline

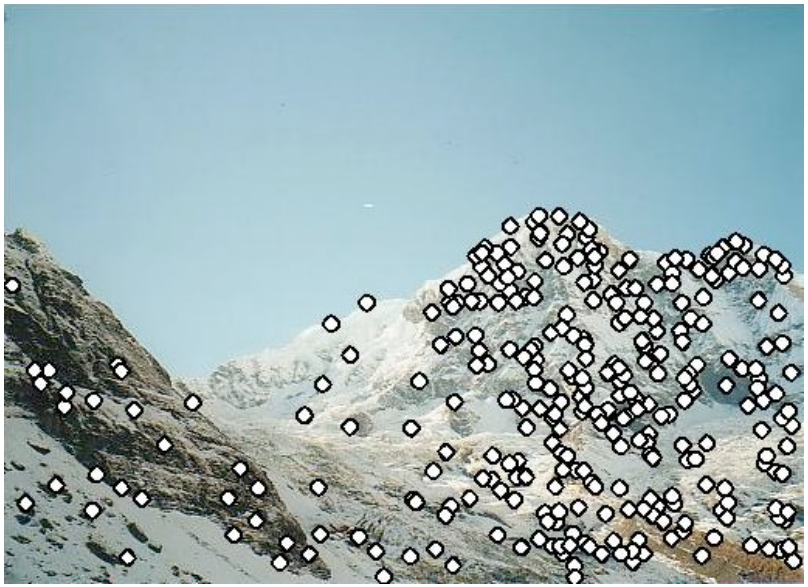

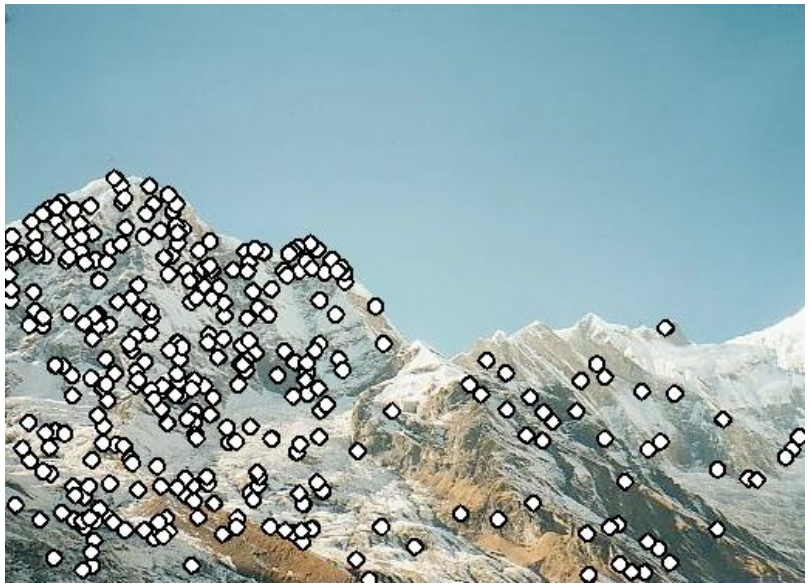

Extract features

Feature-based alignment outline

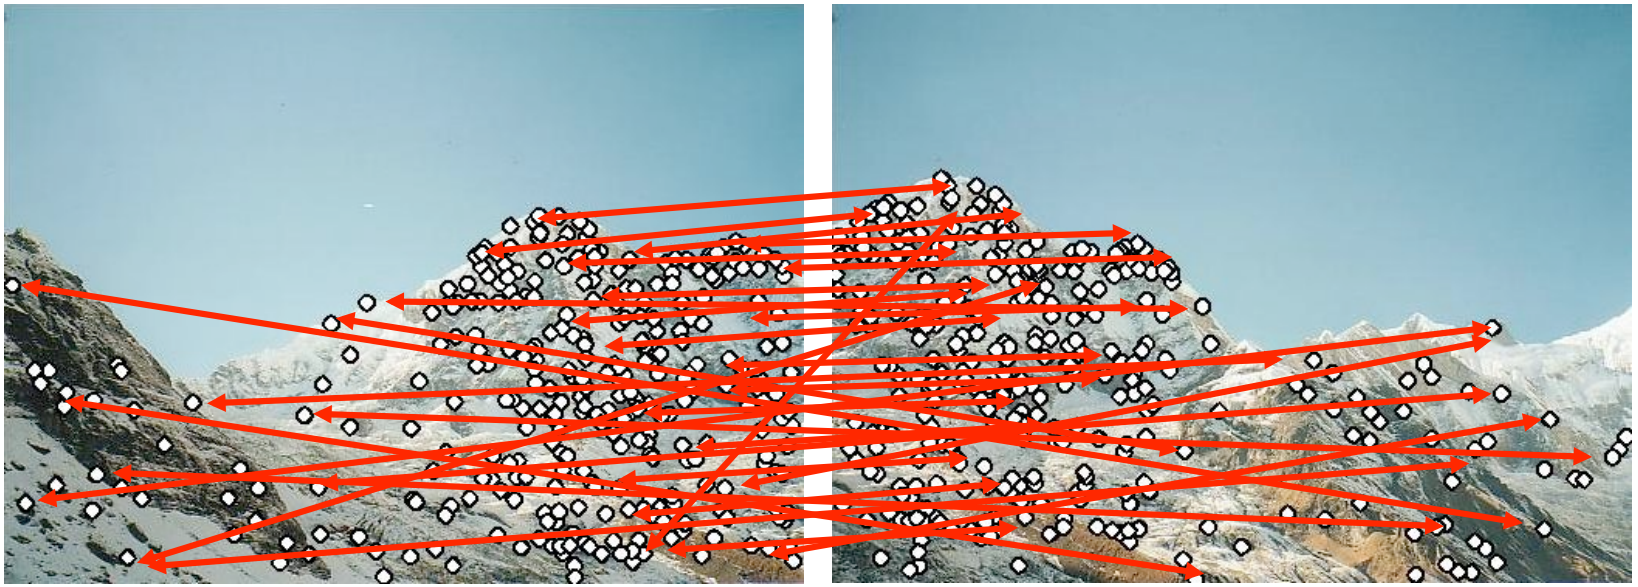

Extract features

Compute *putative matches*

Feature-based alignment outline

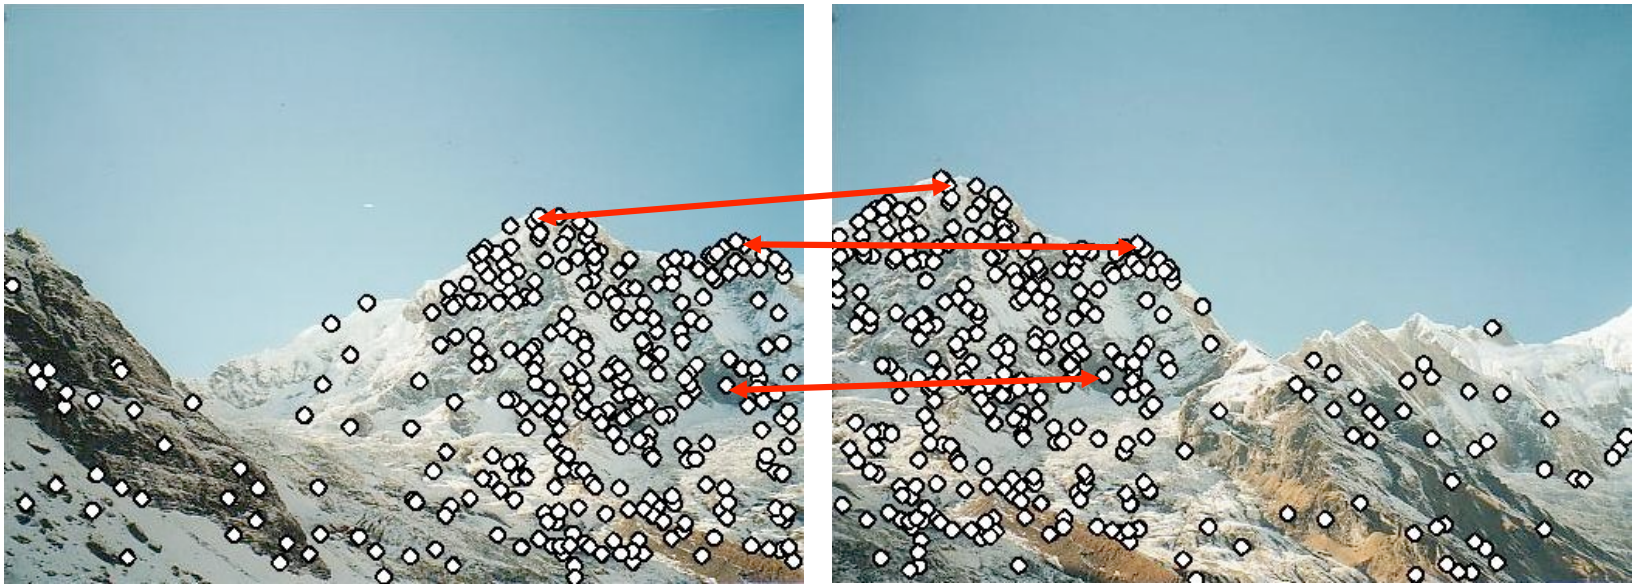

Extract features

Compute *putative matches*

Loop:

- *Hypothesize* transformation *T* (small group of putative matches that are related by *T*)

Feature-based alignment outline

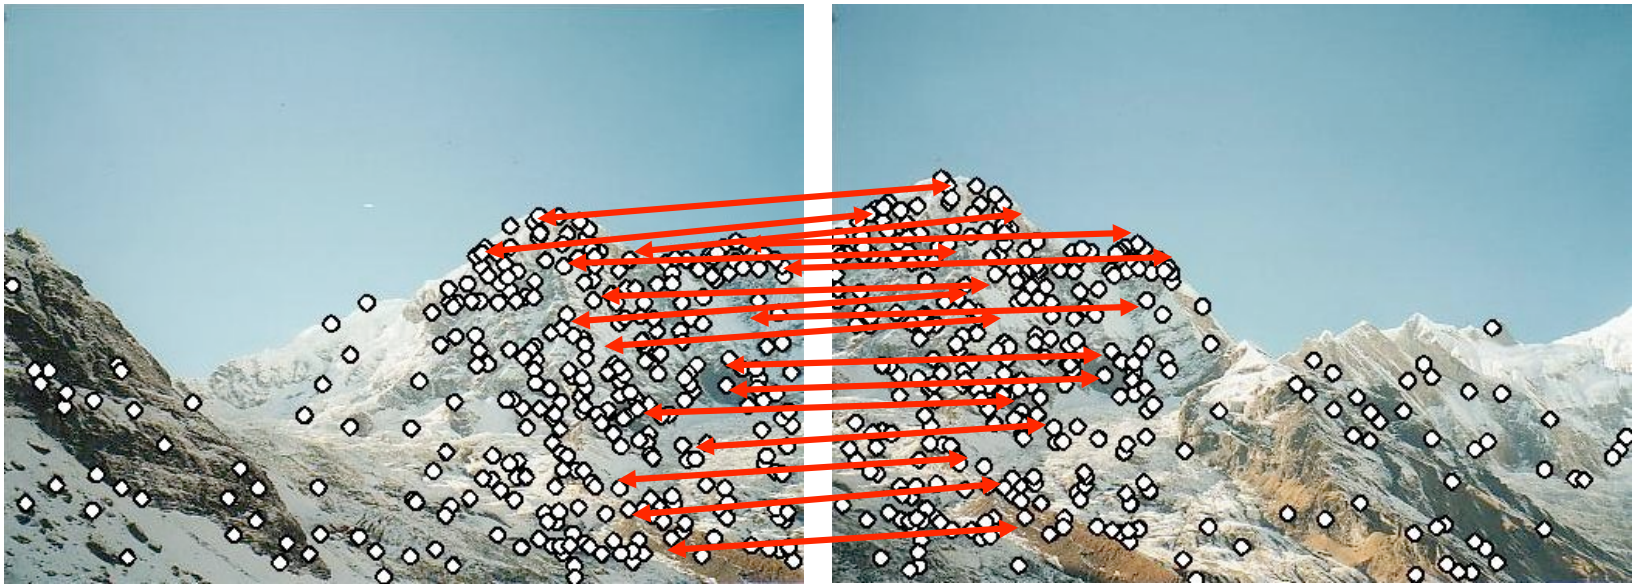

Extract features

Compute *putative matches*

Loop:

- *Hypothesize* transformation *T* (small group of putative matches that are related by *T*)
- *Verify* transformation (search for other matches consistent with *T*)

Feature-based alignment outline

Extract features

Compute *putative matches*

Loop:

- *Hypothesize* transformation *T* (small group of putative matches that are related by *T*)
- *Verify* transformation (search for other matches consistent with *T*)

2D transformation models

Similarity
(translation,
scale, rotation)

Affine

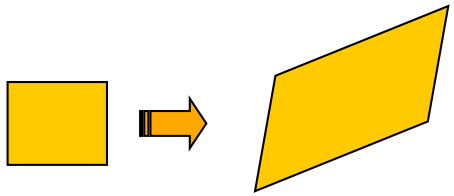

Projective
(homography)

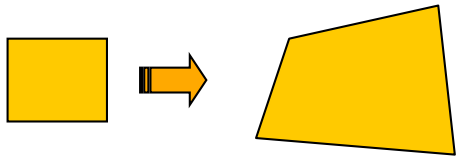

Why these transformations ???

Camera geometry

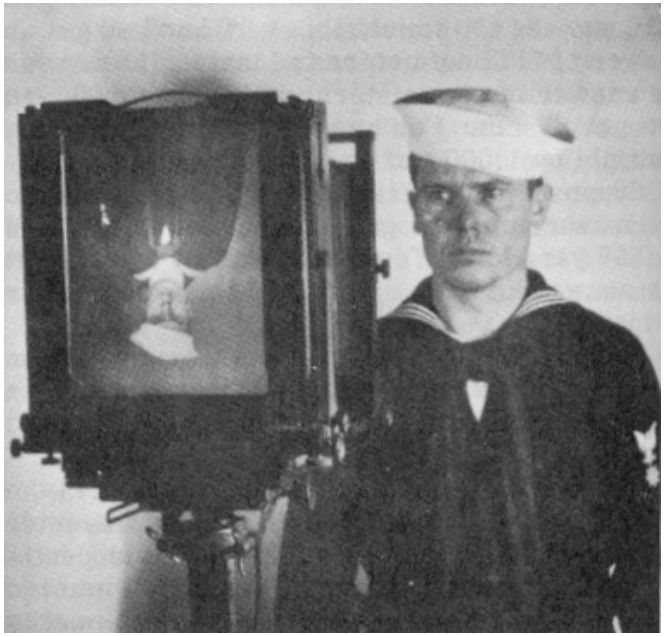

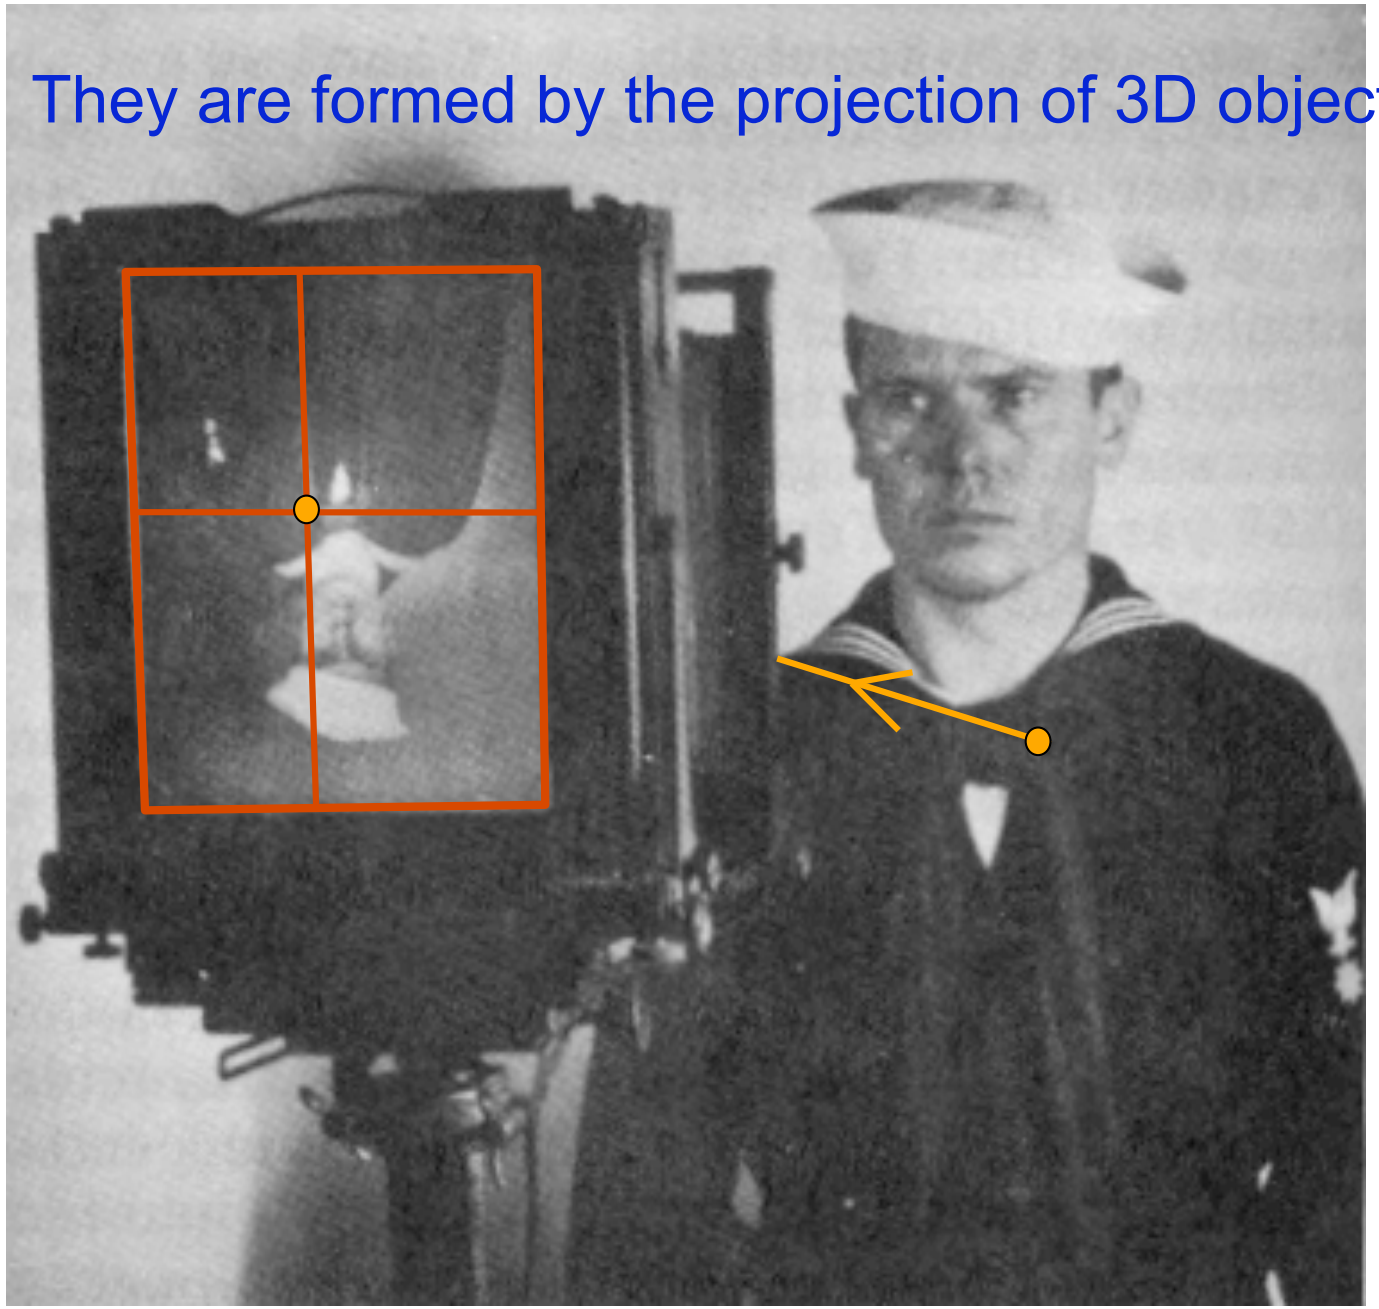

Images are two-dimensional patterns of brightness values.

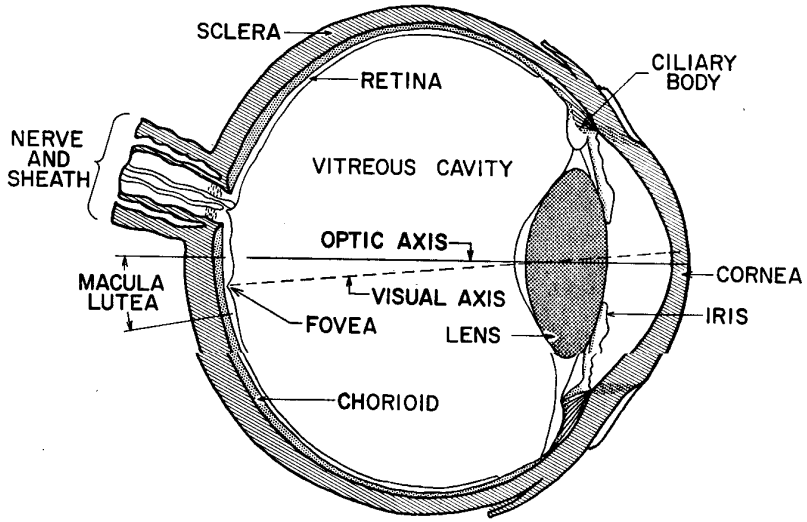

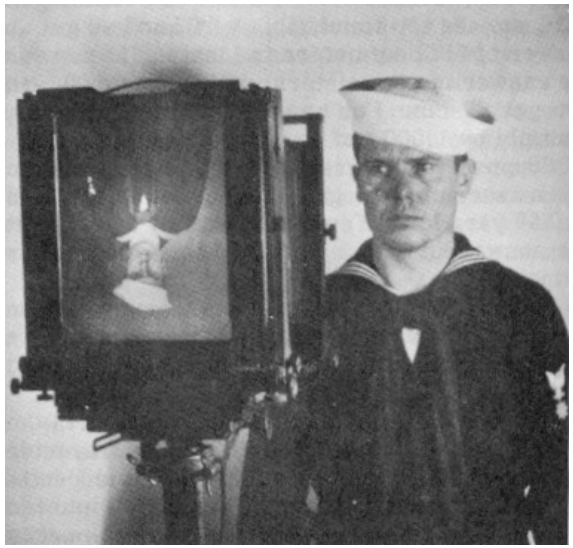

Animal eye: a looonnng time ago.

Photographic camera: Niepce, 1816.

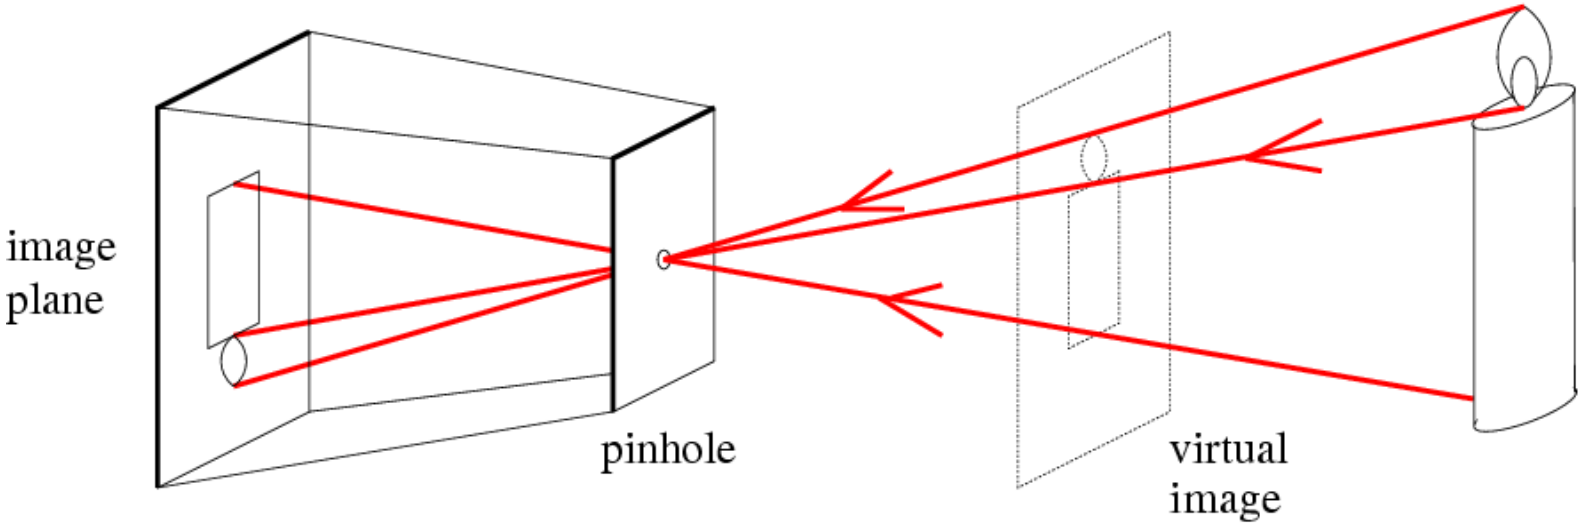

Pinhole perspective projection: Brunelleschi, XVth Century.
Camera obscura: XVIth Century.

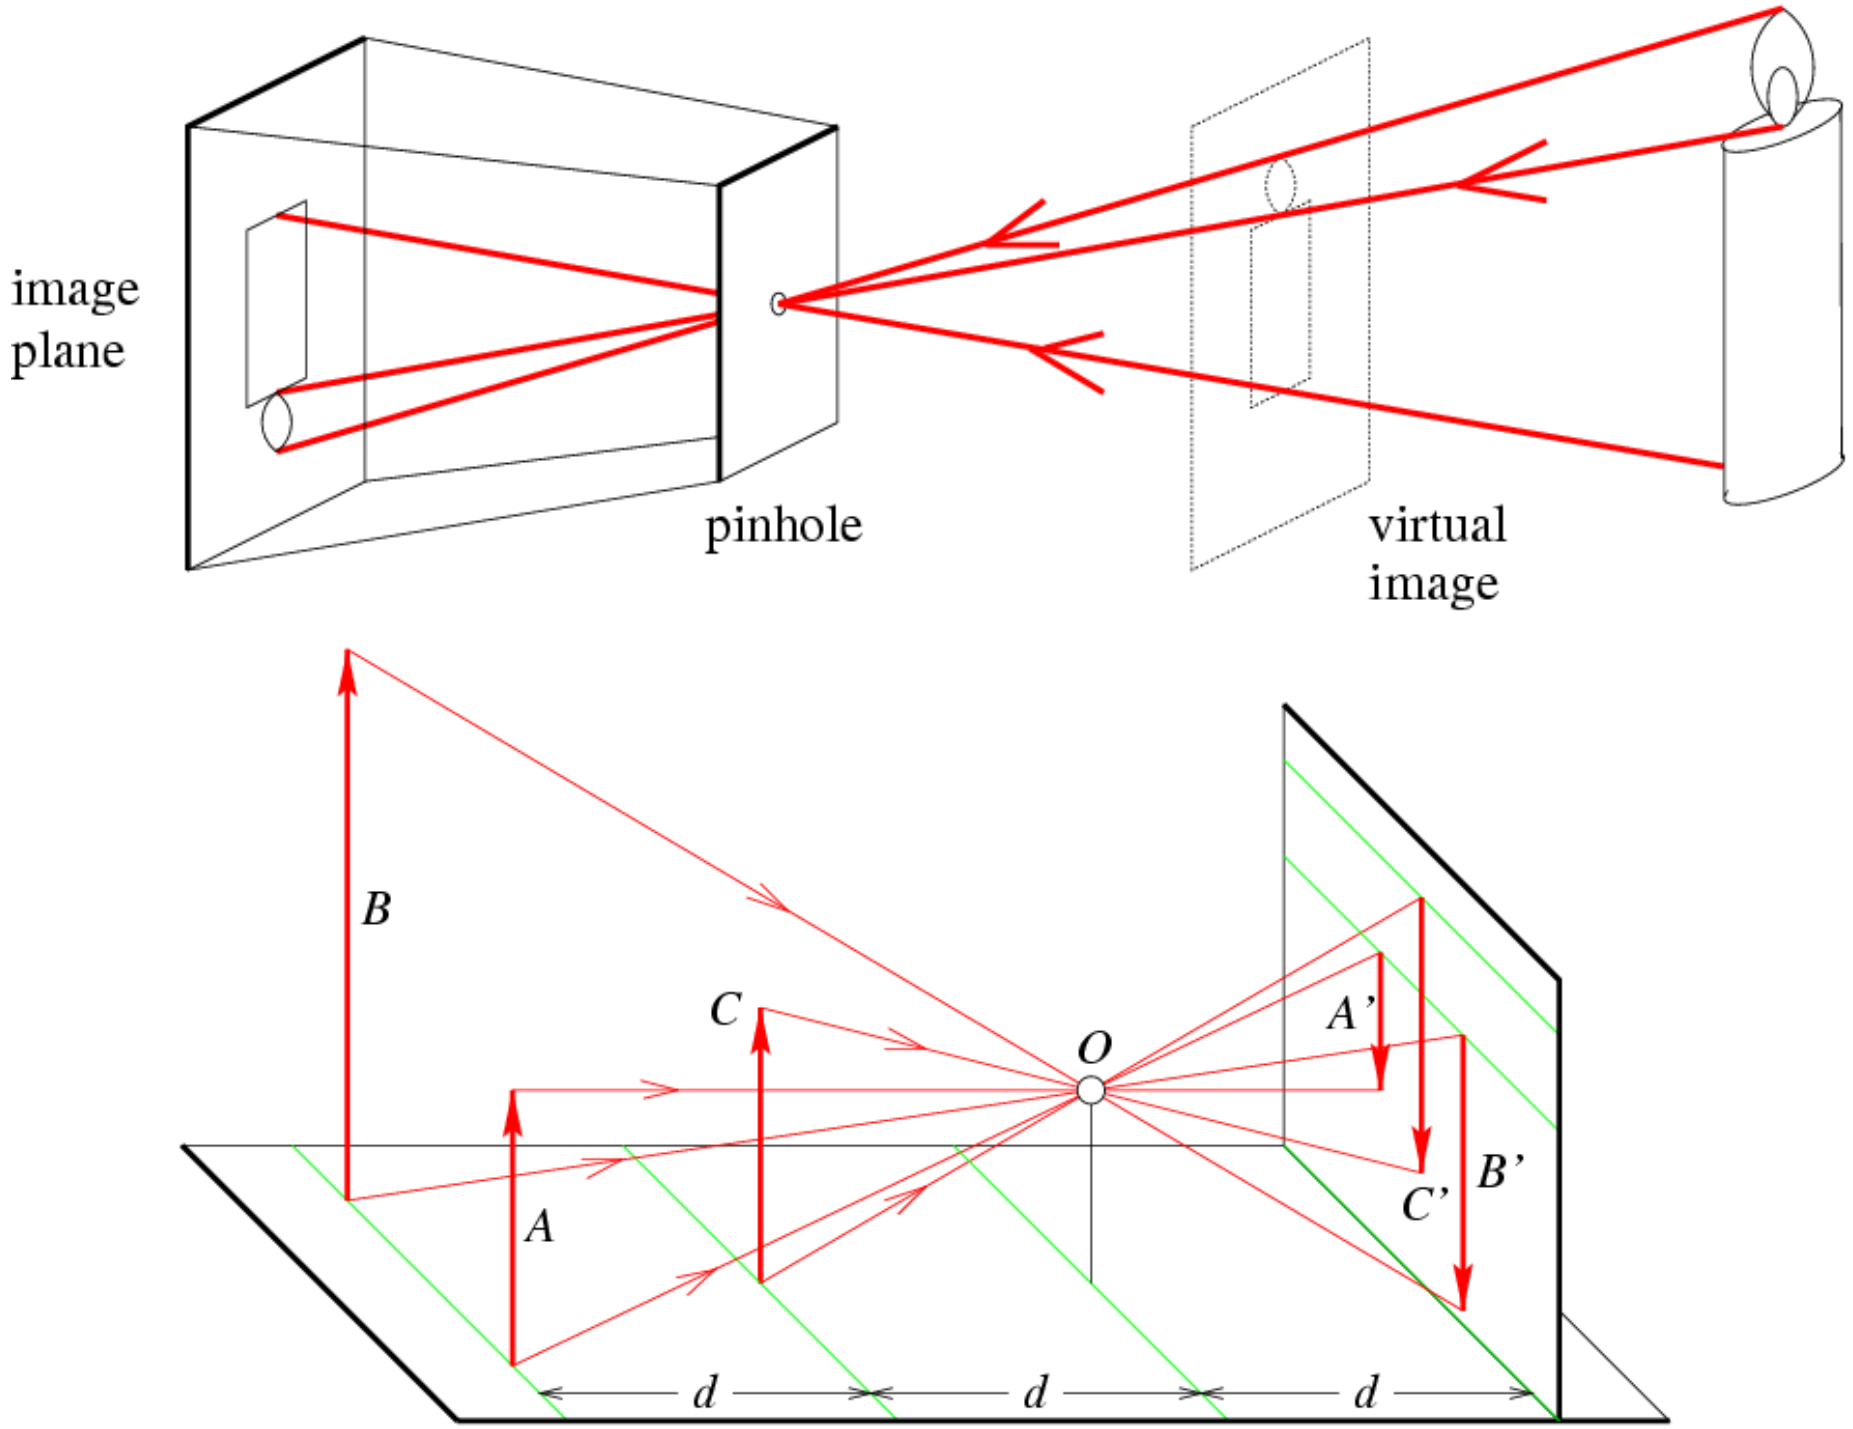
image
plane
pinhole
virtual
image
B
C
O
A'
B'
C'
A
d
d
d

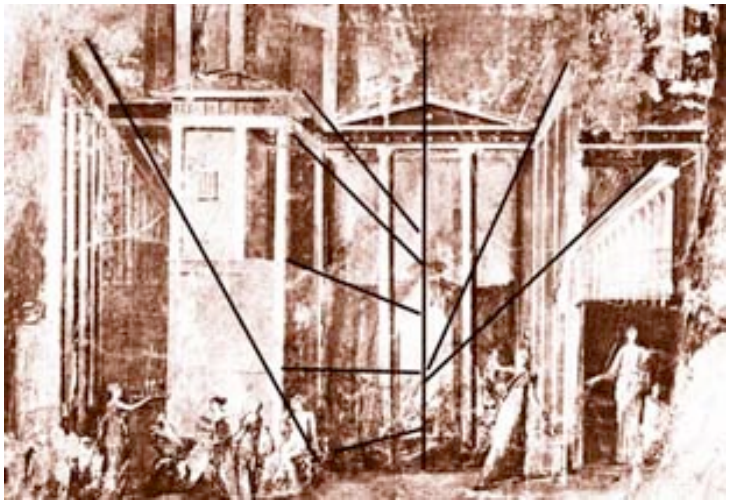

Pompei painting, 2000 years ago.

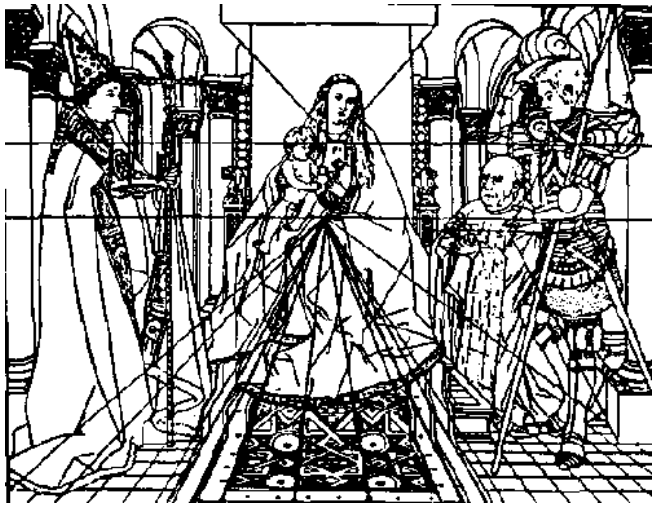

Van Eyk, XIVth Century

Brunelleschi, 1415

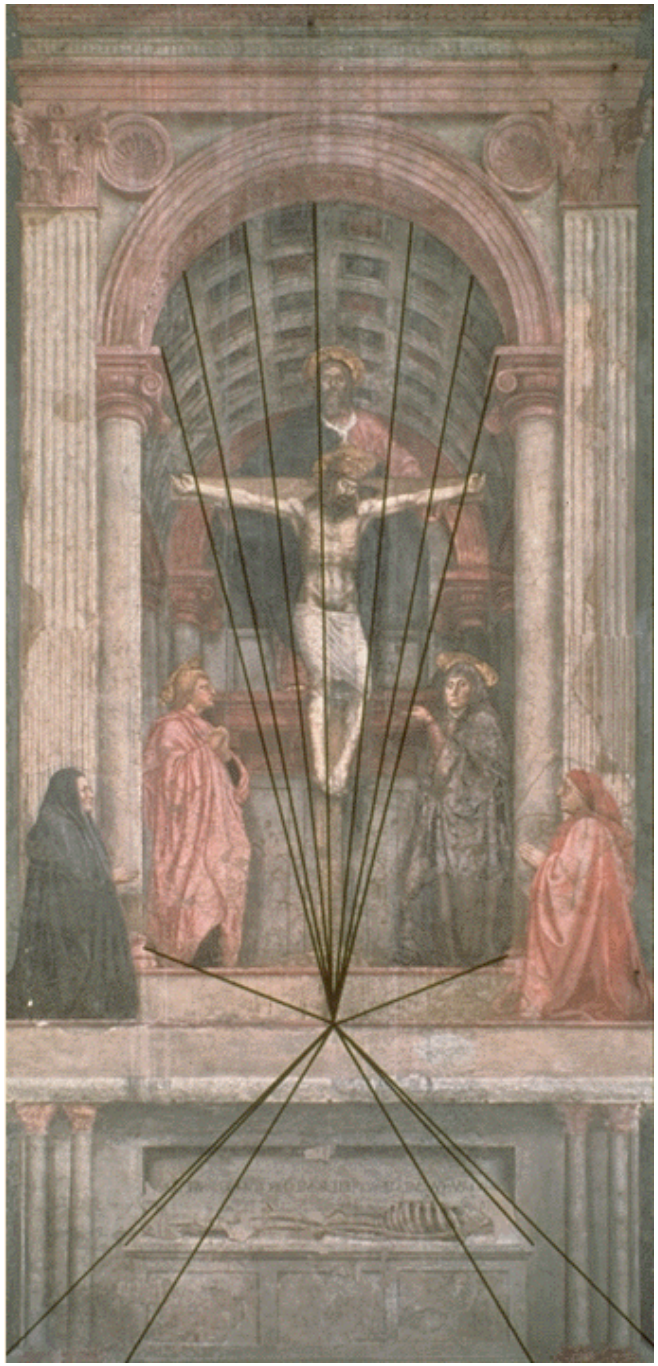

Massaccio’s Trinity, 1425

Pinhole Perspective Equation

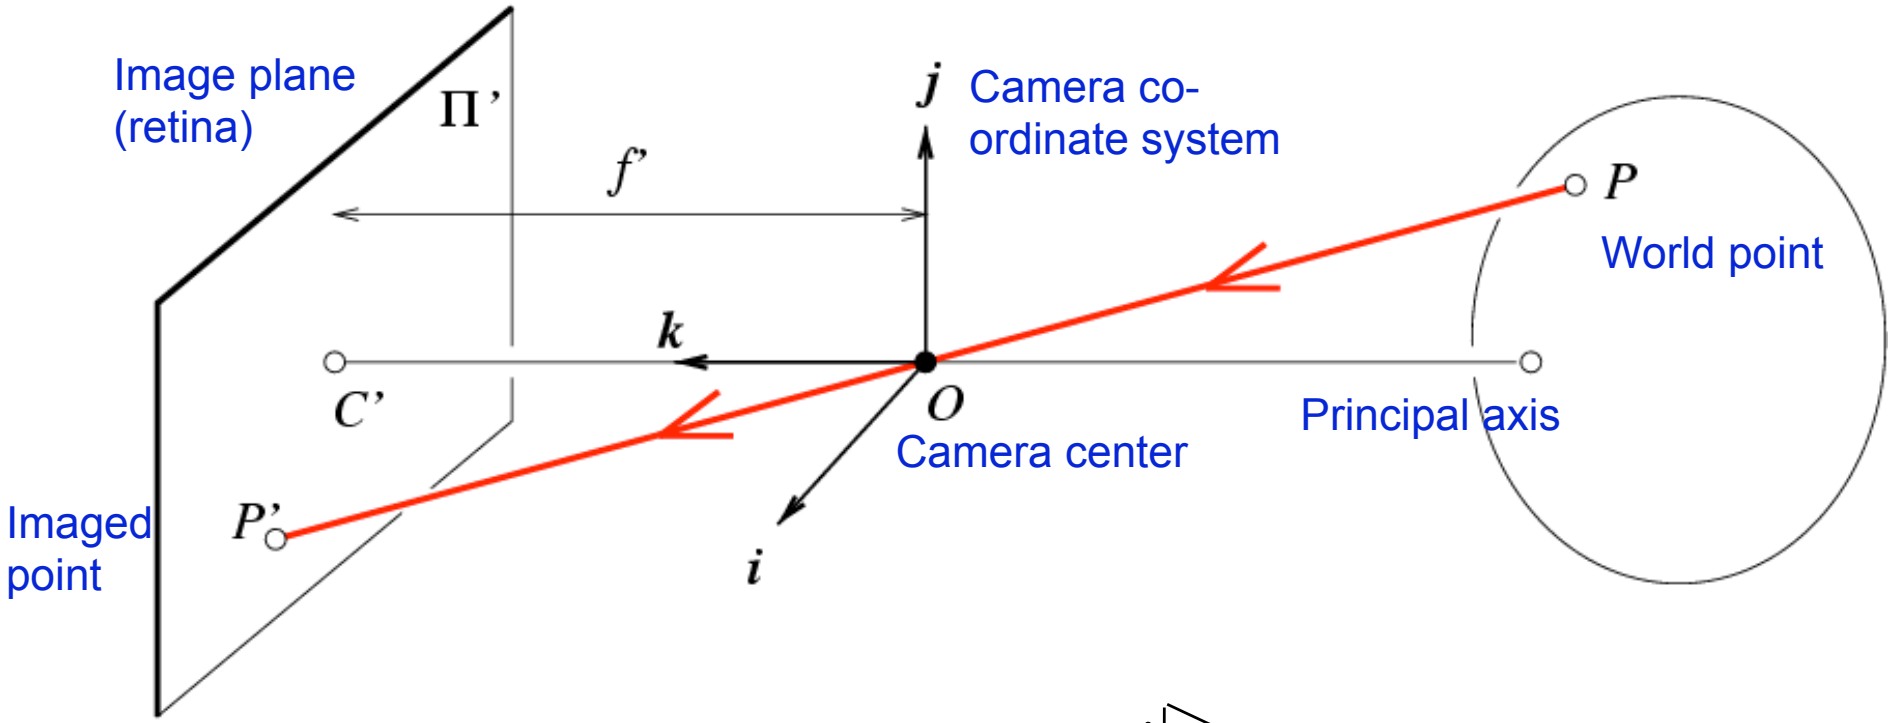

$$P = (x, y, z)^T$$

$$P' = (x', y')^T$$

$$\begin{cases} x' = f' \dfrac{x}{z} \\ y' = f' \dfrac{y}{z} \end{cases}$$

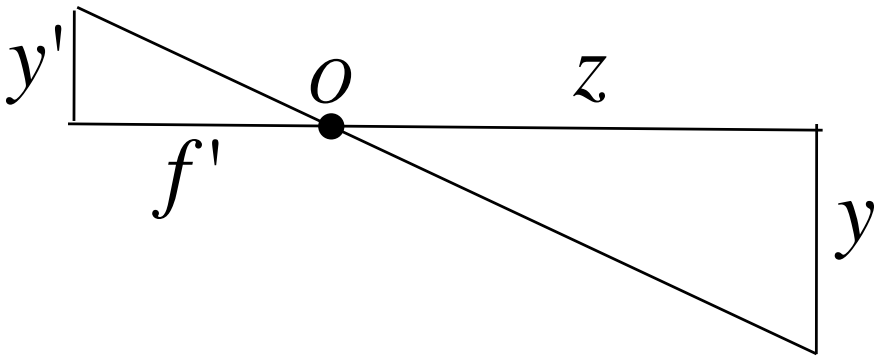

NOTE: *z* is always negative..

Affine projection models: Weak perspective projection

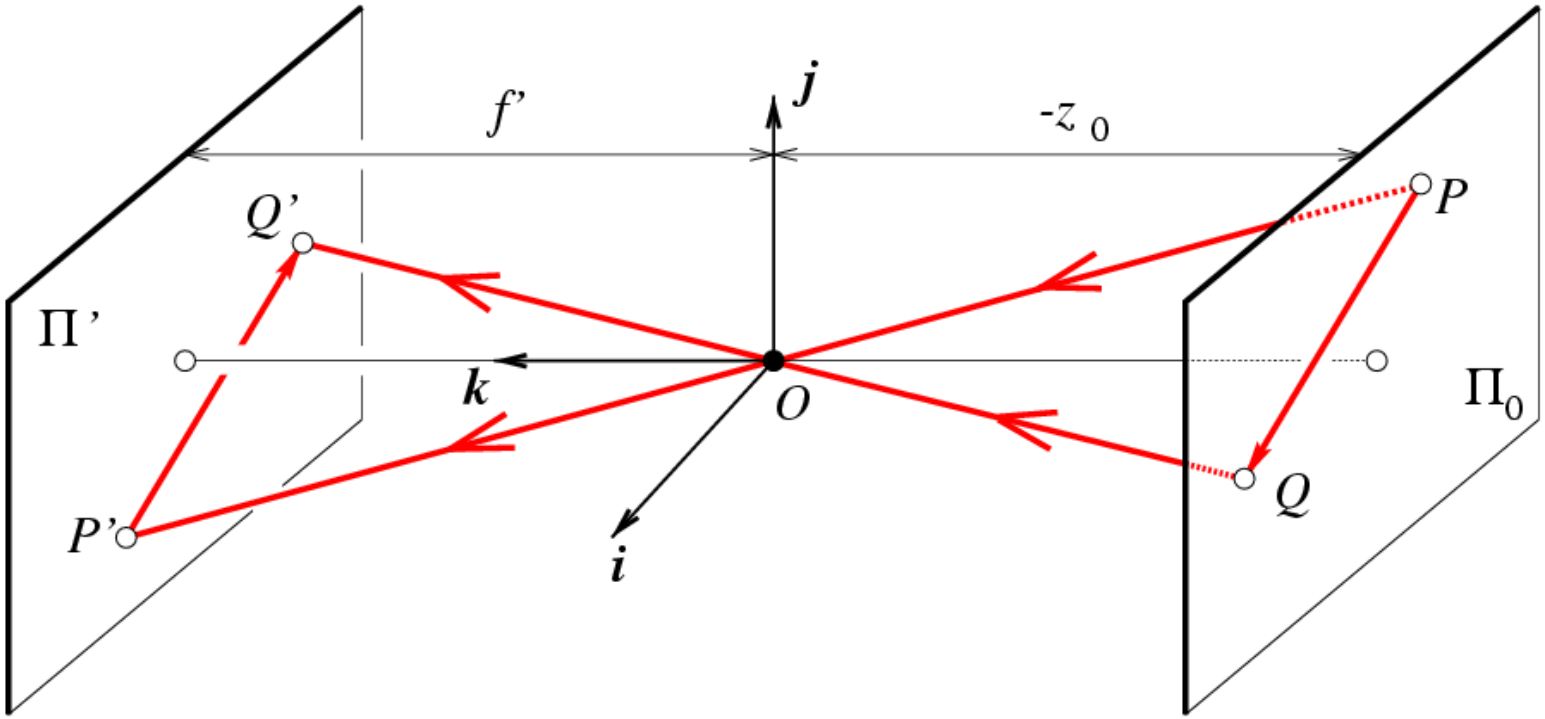

$$\begin{cases} x' = -mx \\ y' = -my \end{cases}$$ where $m = -\frac{f'}{z_0}$ is the magnification.

When the scene relief is small compared its distance from the Camera, *m* can be taken constant: weak perspective projection.

Affine projection models: Orthographic projection

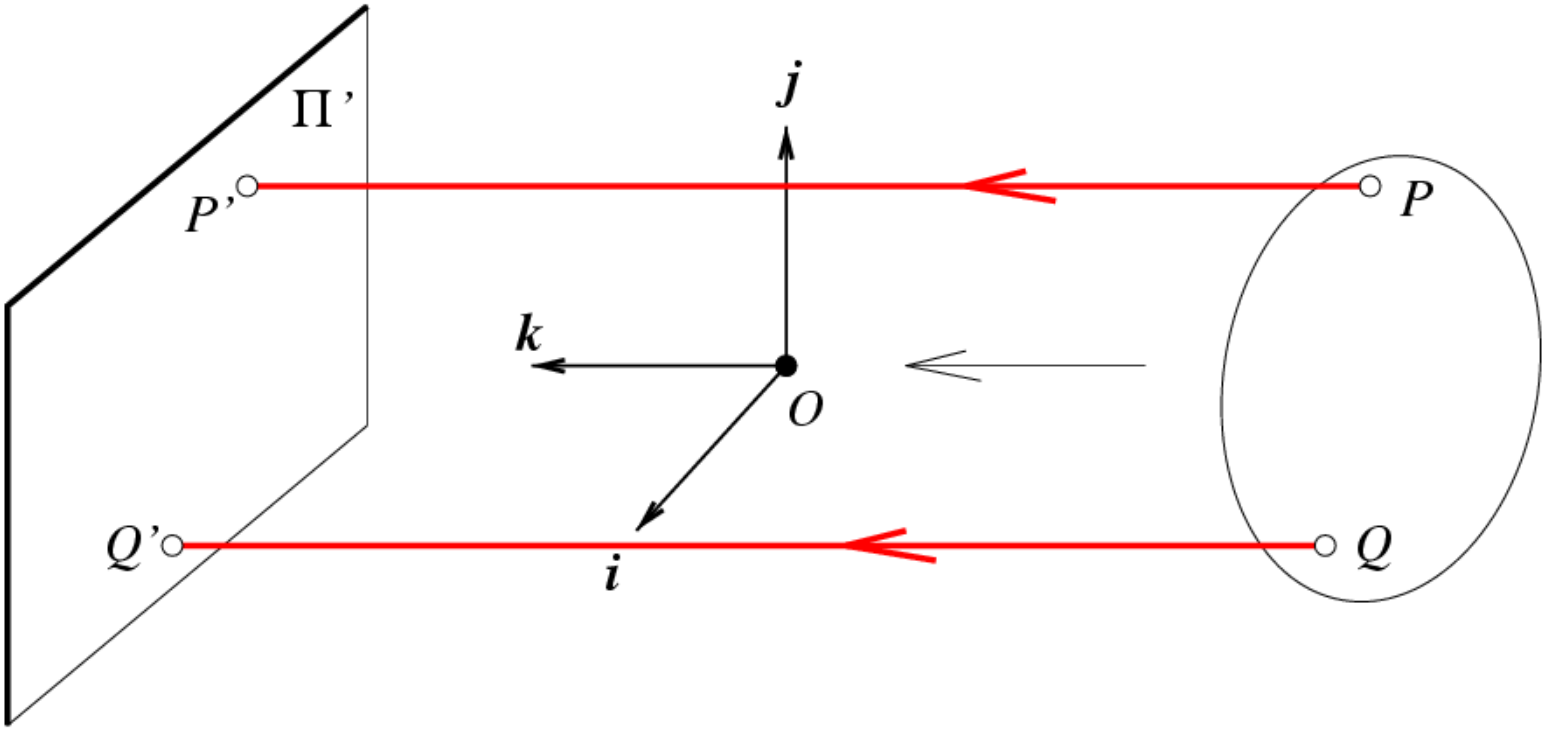

$$\begin{cases} x' = x \\ y' = y \end{cases}$$

When the camera is at a (roughly constant) distance from the scene, take m=1.

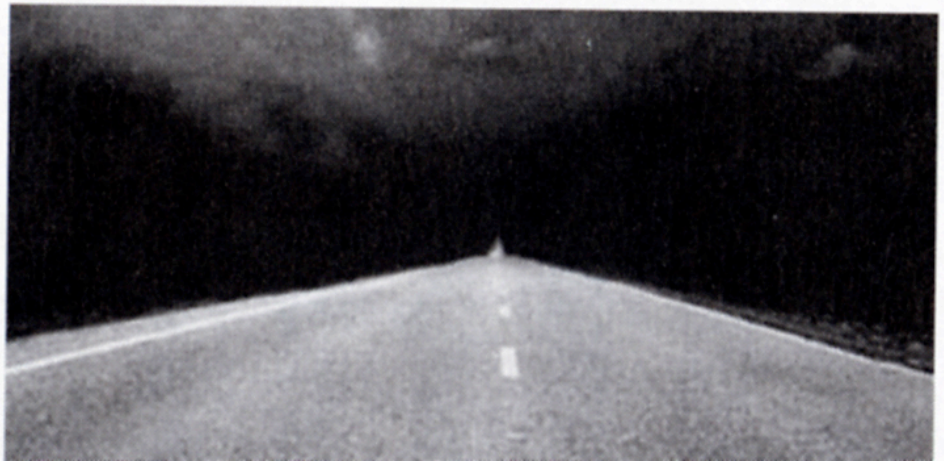

Strong perspective:
Angles are not preserved
The projections of parallel lines intersect at one point

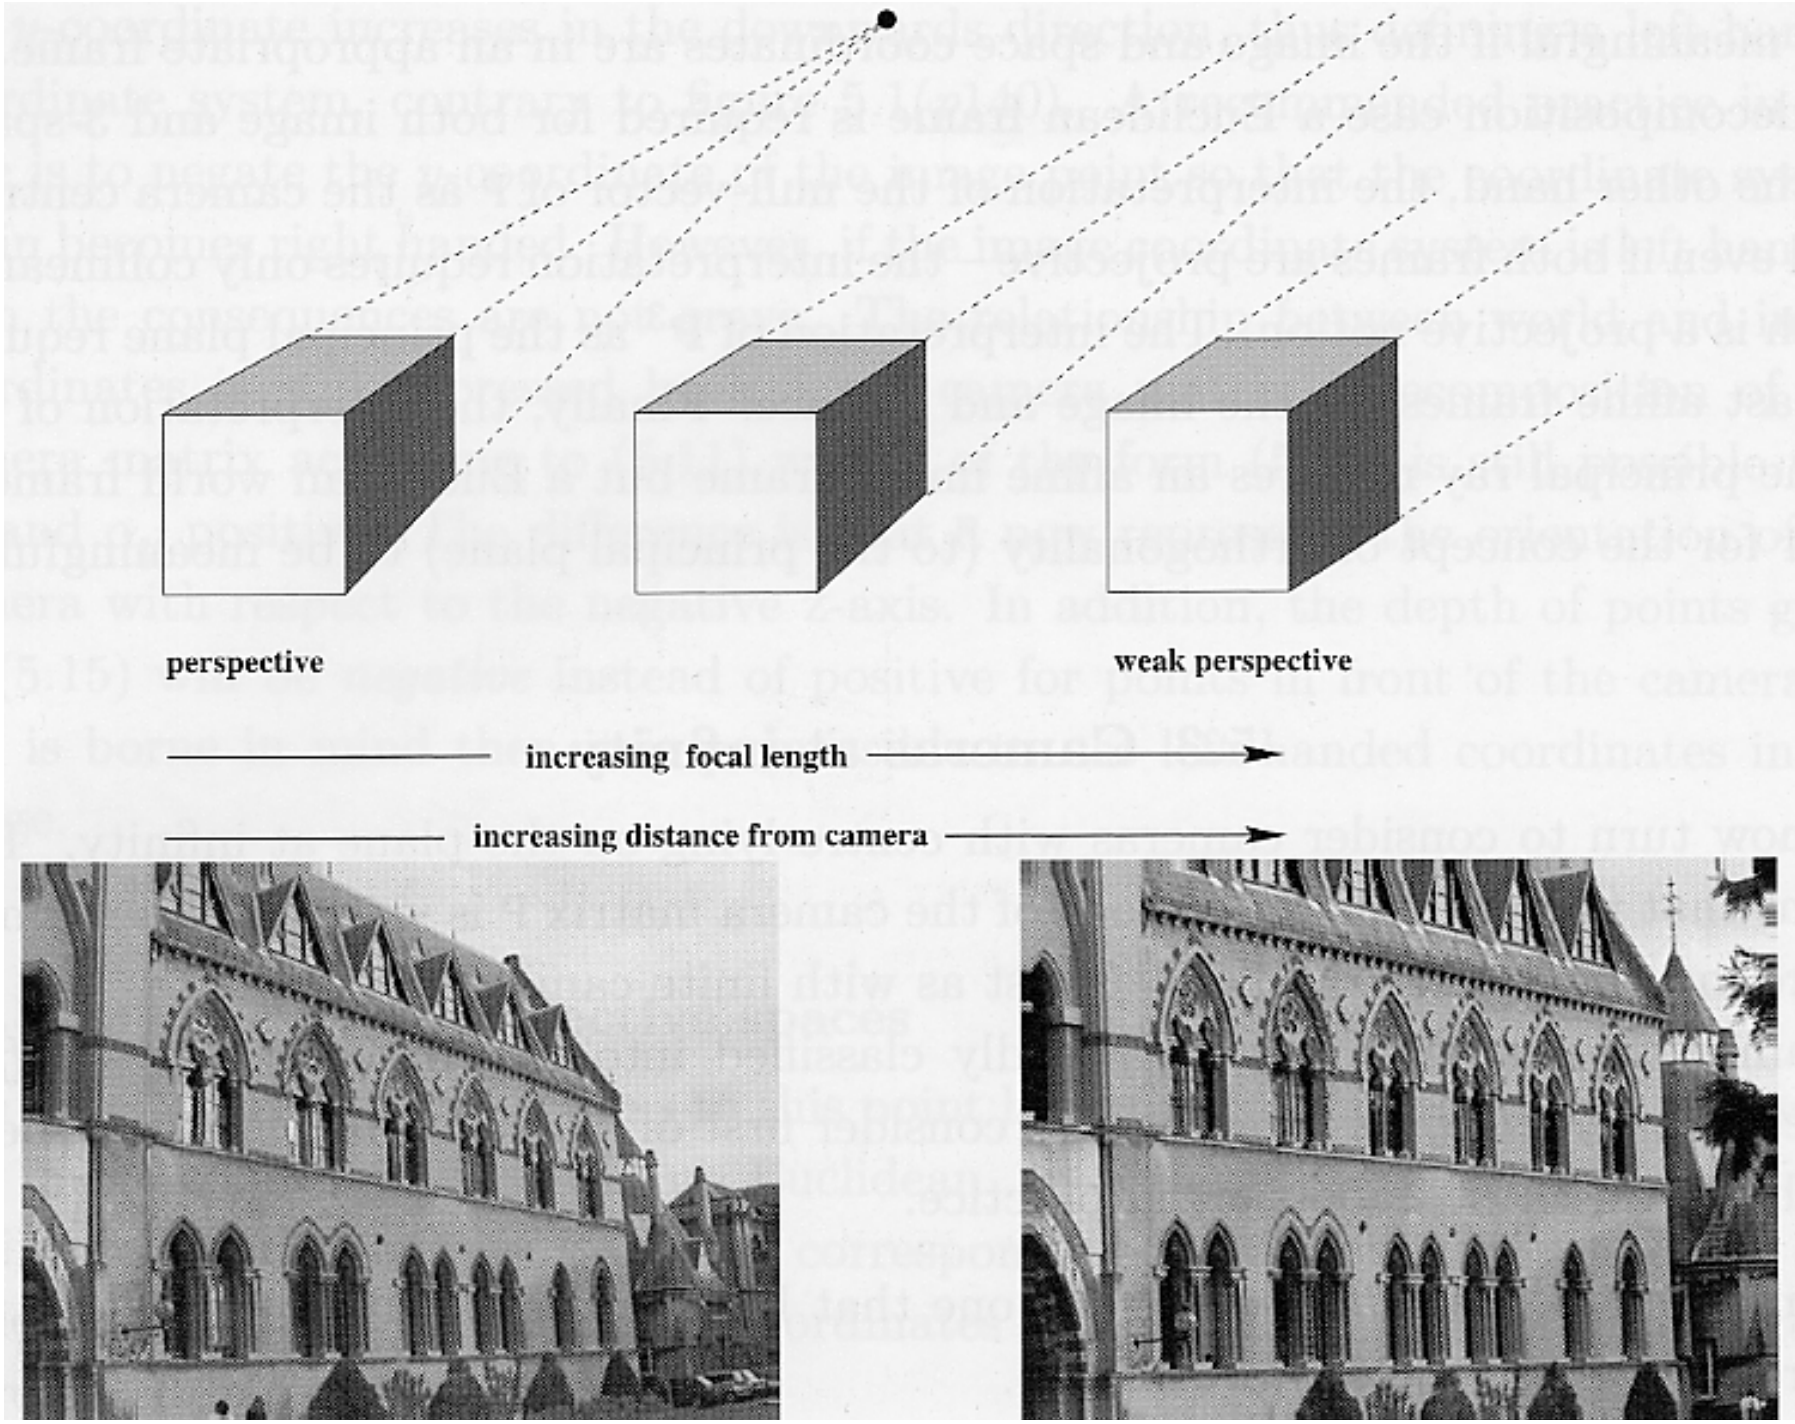

From Zisserman & Hartley

Strong perspective:
Angles are not preserved
The projections of parallel lines intersect at one point

Weak perspective:
Angles are better preserved
The projections of parallel lines are (almost) parallel

Beyond pinhole camera model: Geometric Distortion

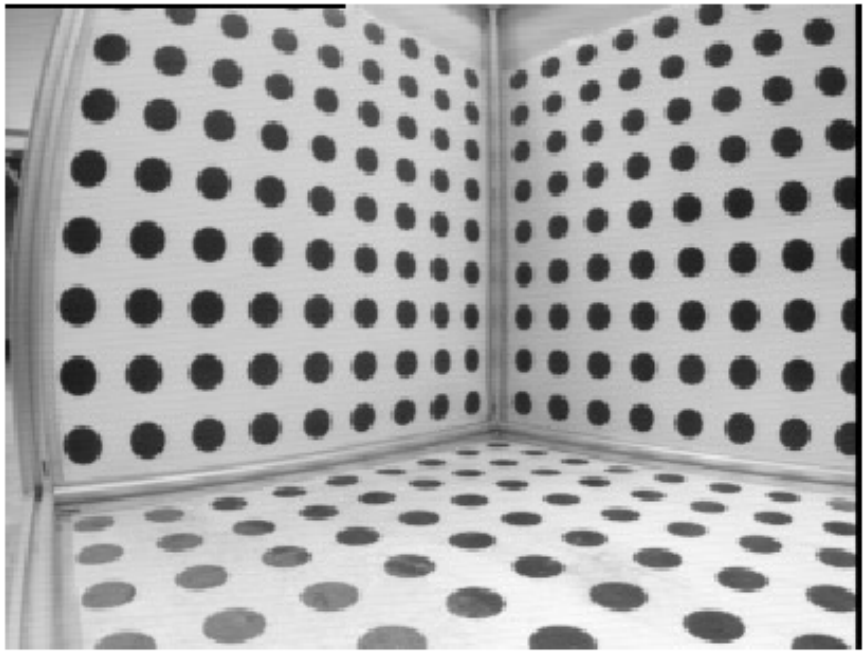

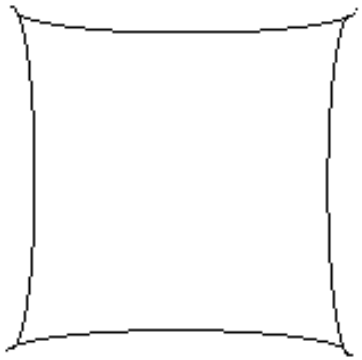

pincushion

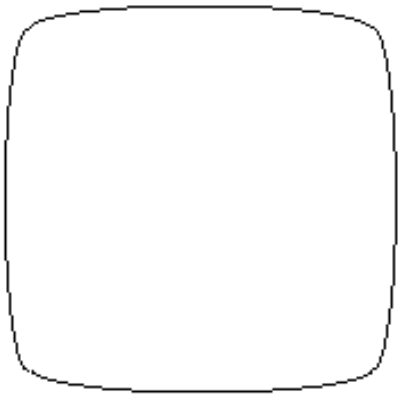

barrel

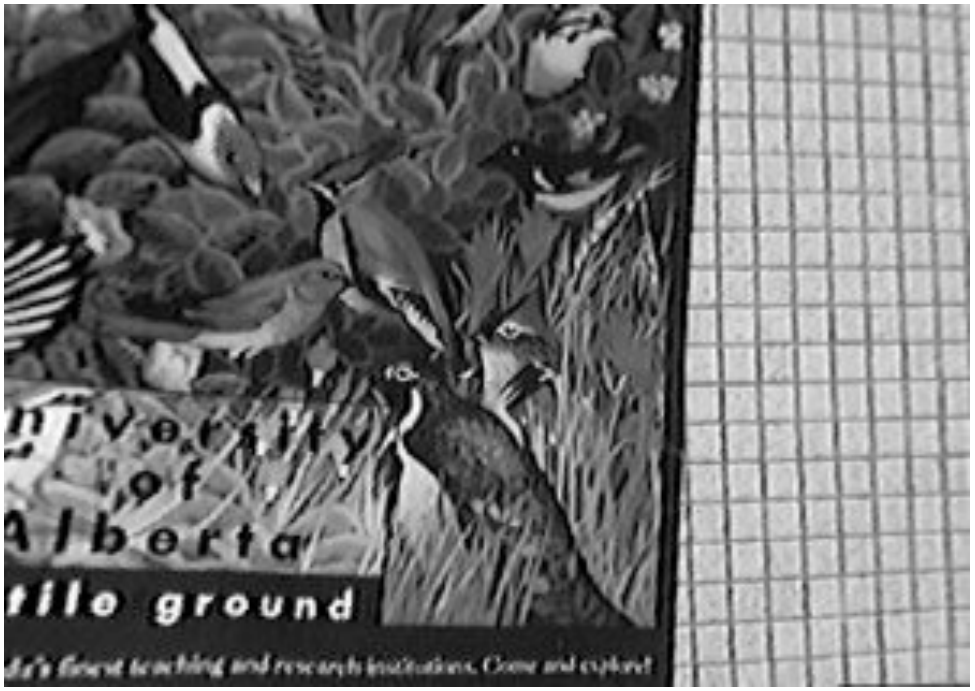

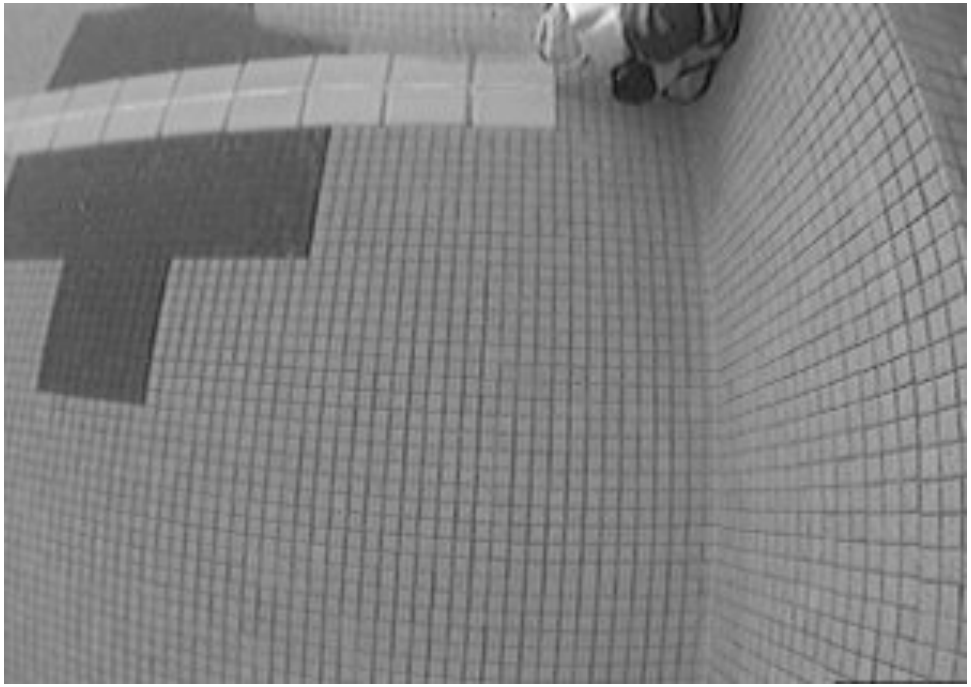

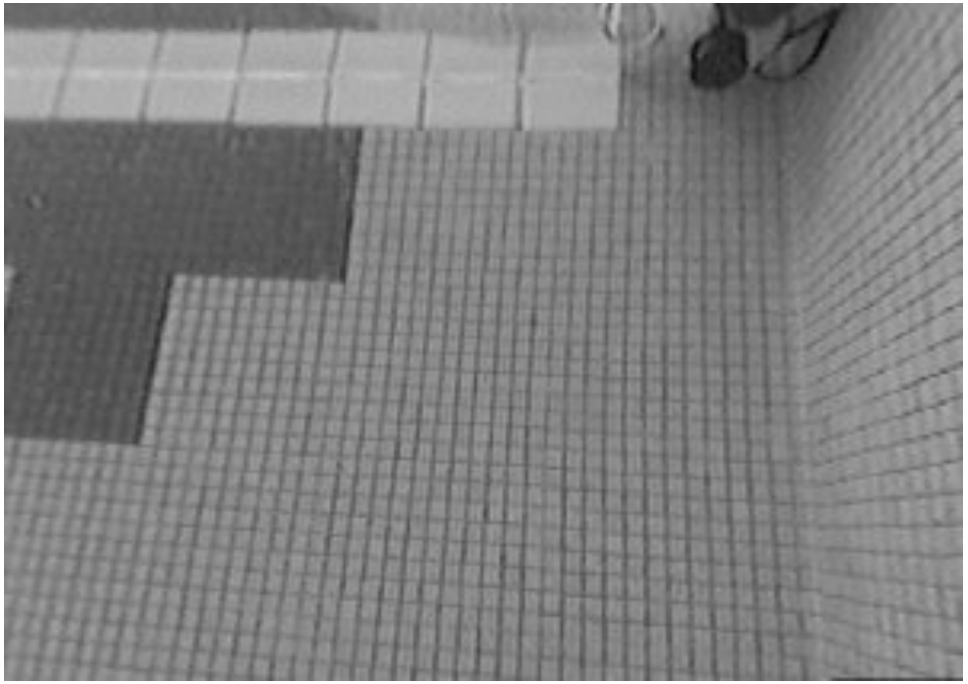

Rectification

Radial Distortion Model

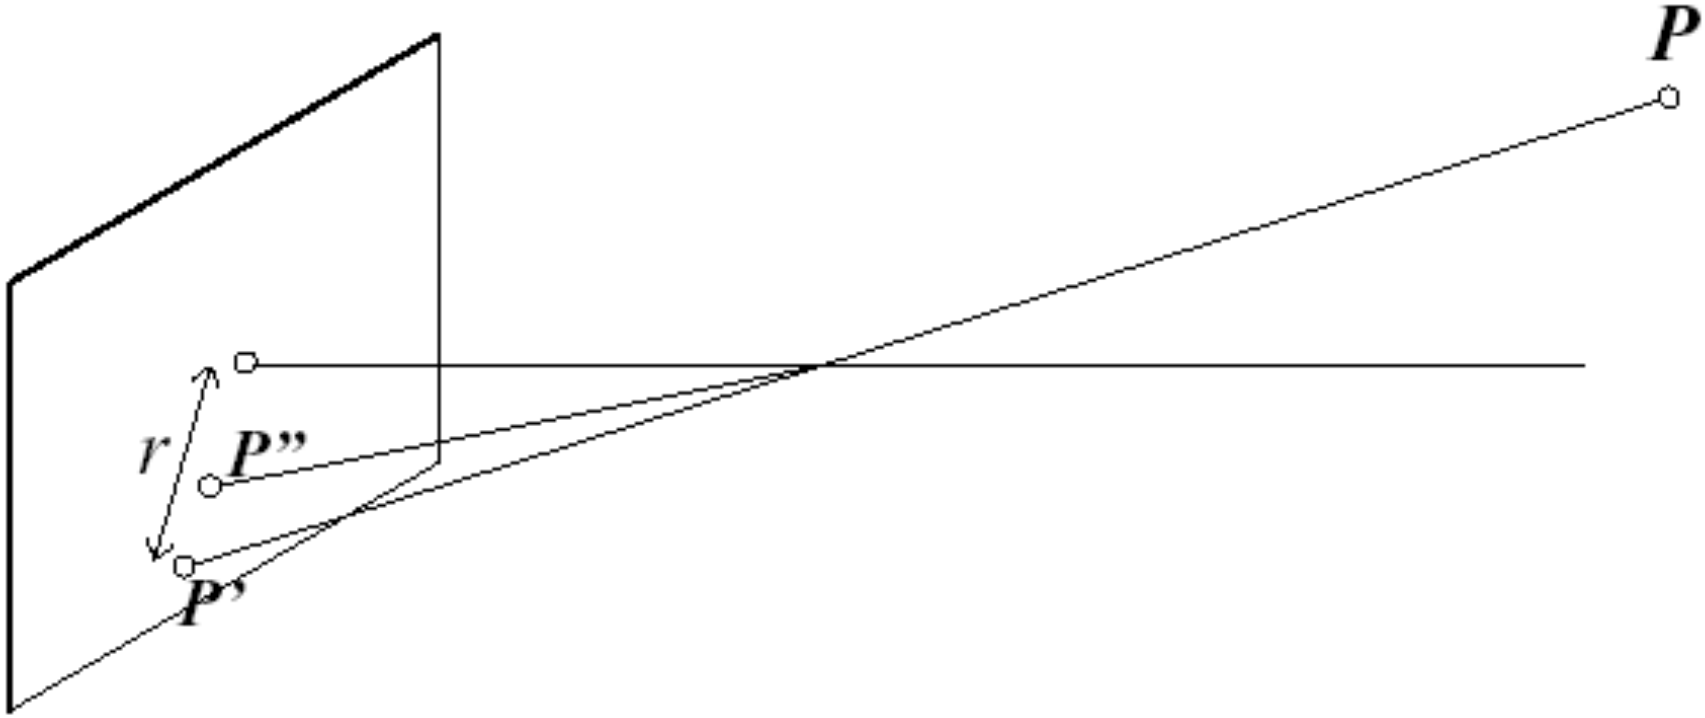

Ideal:

$$x' = f\frac{x}{z} \qquad y' = f\frac{y}{z}$$

Distorted:

$$x'' = \frac{1}{\lambda}x' \qquad y'' = \frac{1}{\lambda}y'$$

$$\lambda = 1 + k_1 r^2 + k_2 r^4 + \cdots$$

Perspective Projection	$x' = f\frac{x}{z}$ $y' = f\frac{y}{z}$	x,y: World coordinates x',y': Image coordinates f: pinhole-to-retina distance
Weak-Perspective Projection (Affine)	$x' \approx -mx$ $y' \approx -my$ $m = -\frac{f}{\bar{z}}$	x,y: World coordinates x',y': Image coordinates m: magnification
Orthographic Projection (Affine)	$x' \approx x$ $y' \approx y$	x,y: World coordinates x',y': Image coordinates
Common distortion model	$x'' = \frac{1}{\lambda}x'$ $y'' = \frac{1}{\lambda}y'$ $\lambda = 1 + k_1 r^2 + k_2 r^4 + \cdots$	x',y': Ideal image coordinates x'',y'': Actual image coordinates

Cameras and their parameters

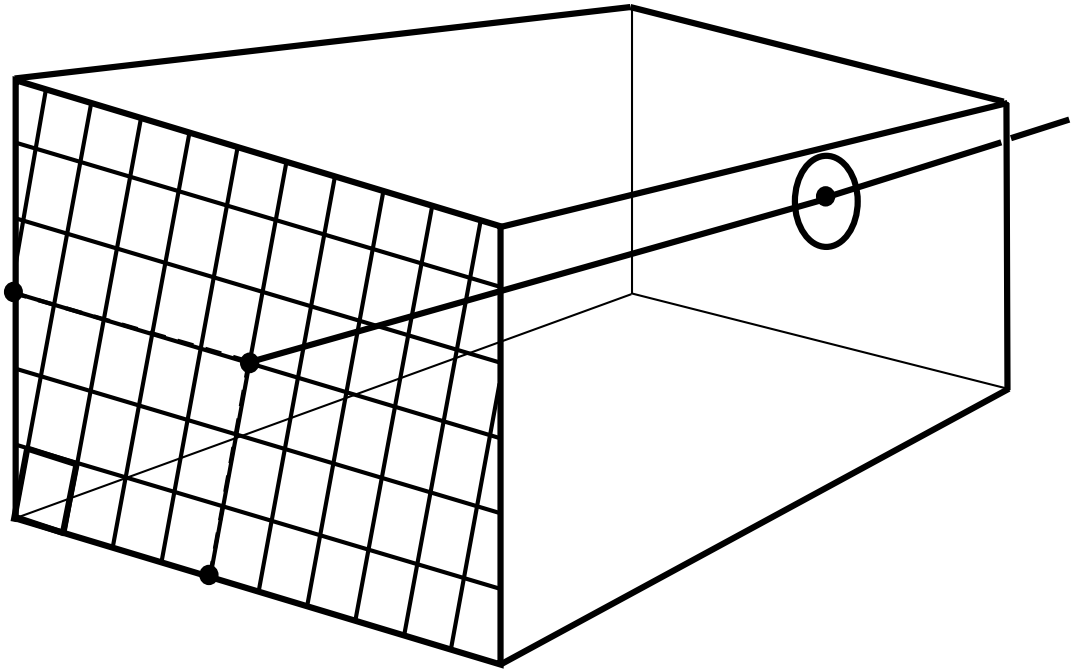

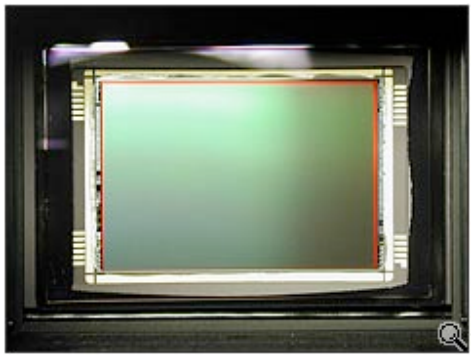

Images from M. Pollefeys

The Intrinsic Parameters of a Camera

Units:

k,l : pixel/m

f : m

α,β : pixel

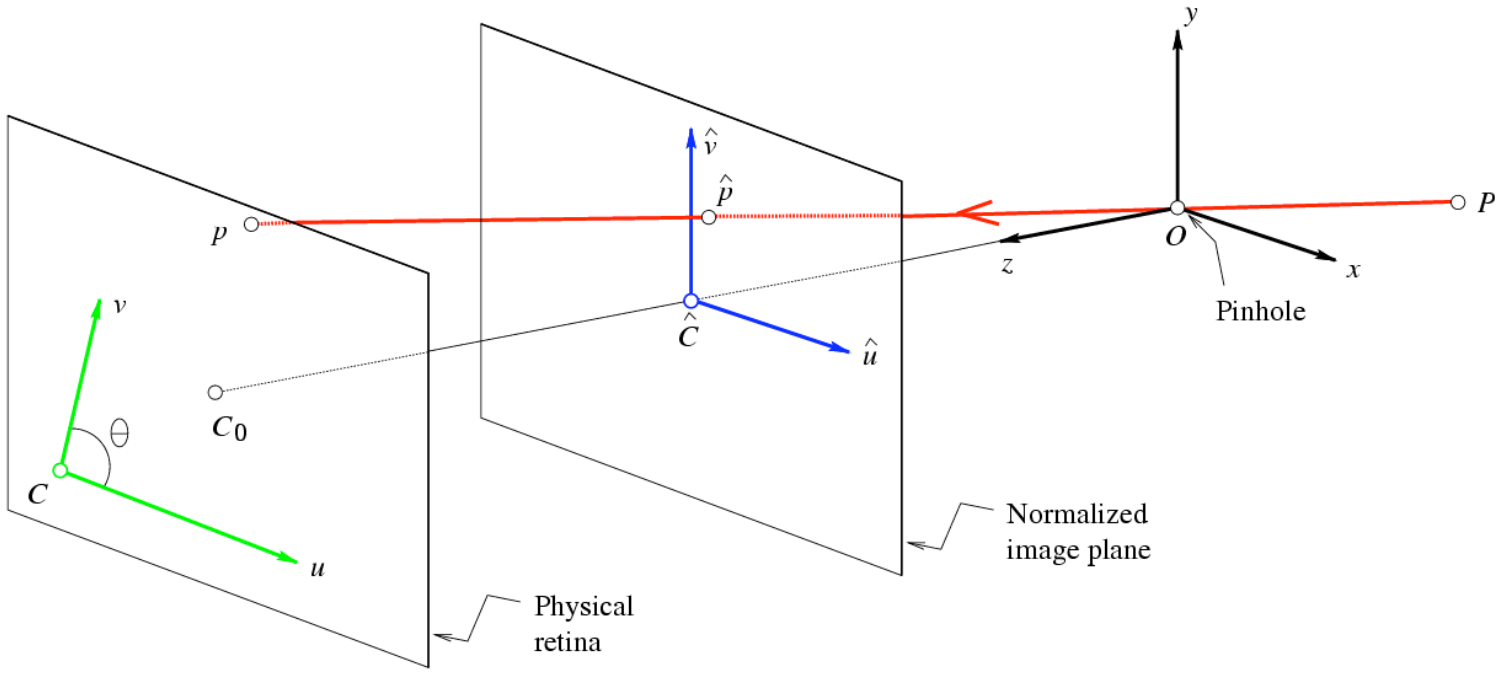

$$\begin{cases} \hat{u} = \dfrac{x}{z} \\ \hat{v} = \dfrac{y}{z} \end{cases} \iff \hat{\boldsymbol{p}} = \frac{1}{z}\begin{pmatrix} \mathrm{Id} & \mathbf{0} \end{pmatrix}\begin{pmatrix} \boldsymbol{P} \\ 1 \end{pmatrix}$$

Physical Image Coordinates

Normalized Image Coordinates

$$\begin{cases} u = kf\dfrac{x}{z} \\ v = lf\dfrac{y}{z} \end{cases} \rightarrow \begin{cases} u = \alpha\dfrac{x}{z} + u_0 \\ v = \beta\dfrac{y}{z} + v_0 \end{cases} \rightarrow \begin{cases} u = \alpha\dfrac{x}{z} - \alpha\cot\theta\dfrac{y}{z} + u_0 \\ v = \dfrac{\beta}{\sin\theta}\dfrac{y}{z} + v_0 \end{cases}$$

The Intrinsic Parameters of a Camera

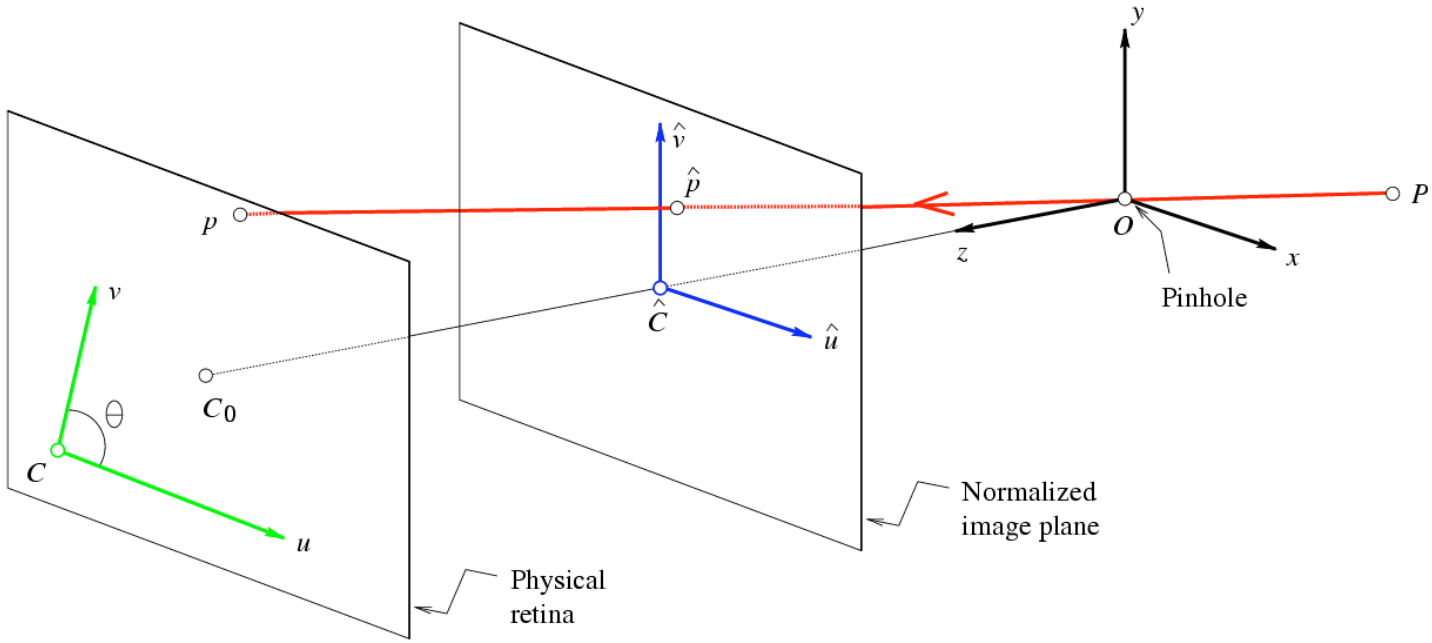

Calibration Matrix

$$\boldsymbol{p} = \mathcal{K}\hat{\boldsymbol{p}}, \quad \text{where} \quad \boldsymbol{p} = \begin{pmatrix} u \\ v \\ 1 \end{pmatrix} \quad \text{and} \quad \mathcal{K} \stackrel{\text{def}}{=} \begin{pmatrix} \alpha & -\alpha \cot\theta & u_0 \\ 0 & \dfrac{\beta}{\sin\theta} & v_0 \\ 0 & 0 & 1 \end{pmatrix}$$

The Perspective Projection Equation

$$\boldsymbol{p} = \frac{1}{z}\mathcal{M}\boldsymbol{P}, \quad \text{where} \quad \mathcal{M} \stackrel{\text{def}}{=} (\mathcal{K} \quad \mathbf{0})$$

Notation

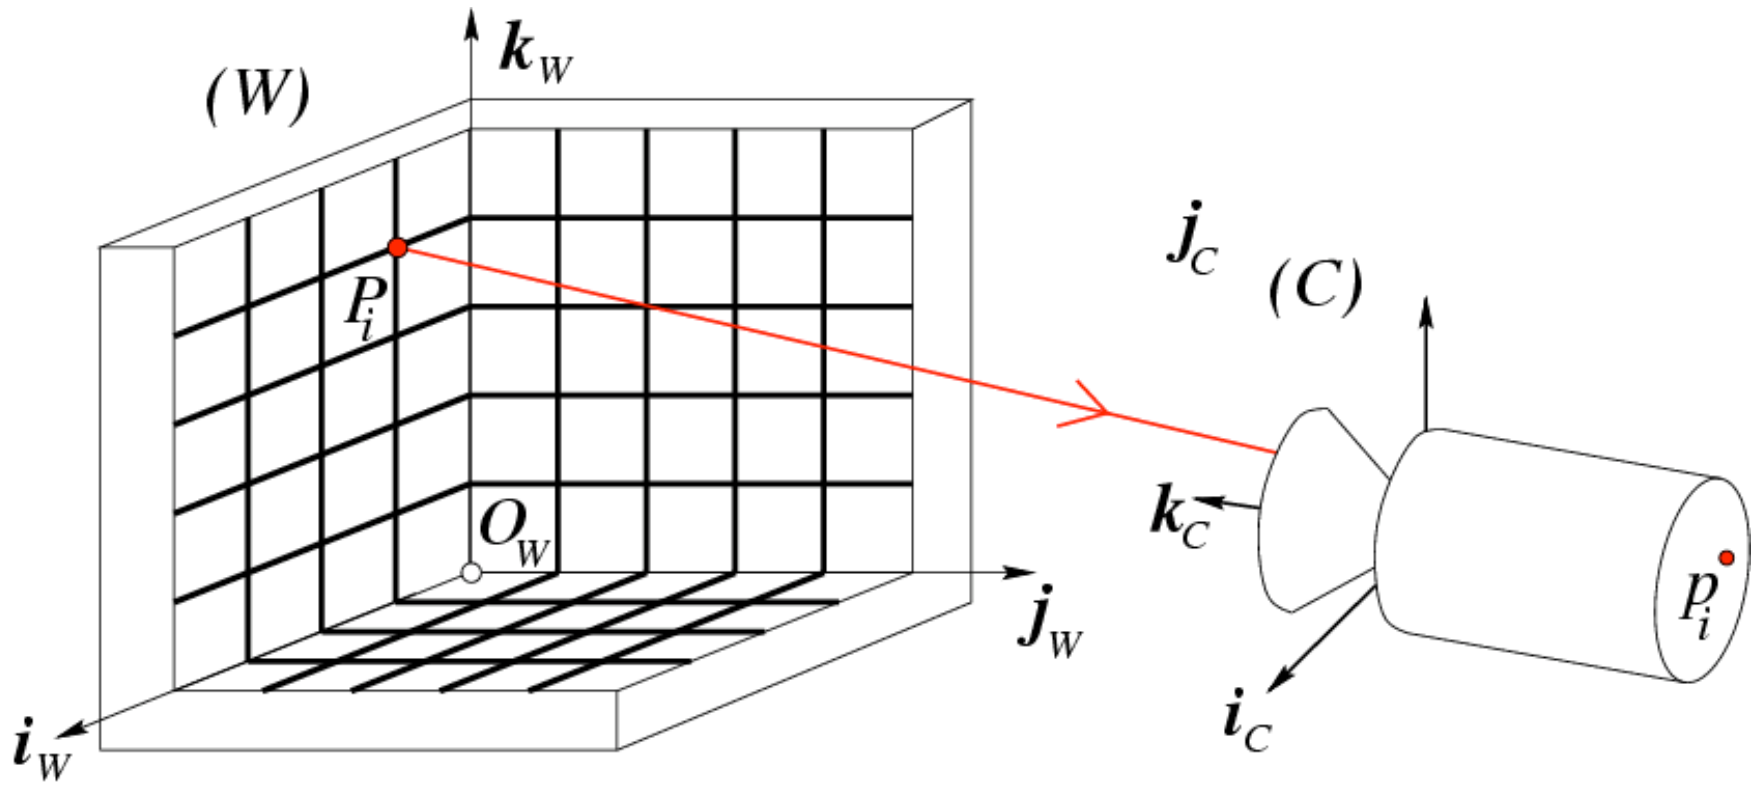

Euclidean Geometry

Recall:
Coordinate Changes and Rigid Transformations

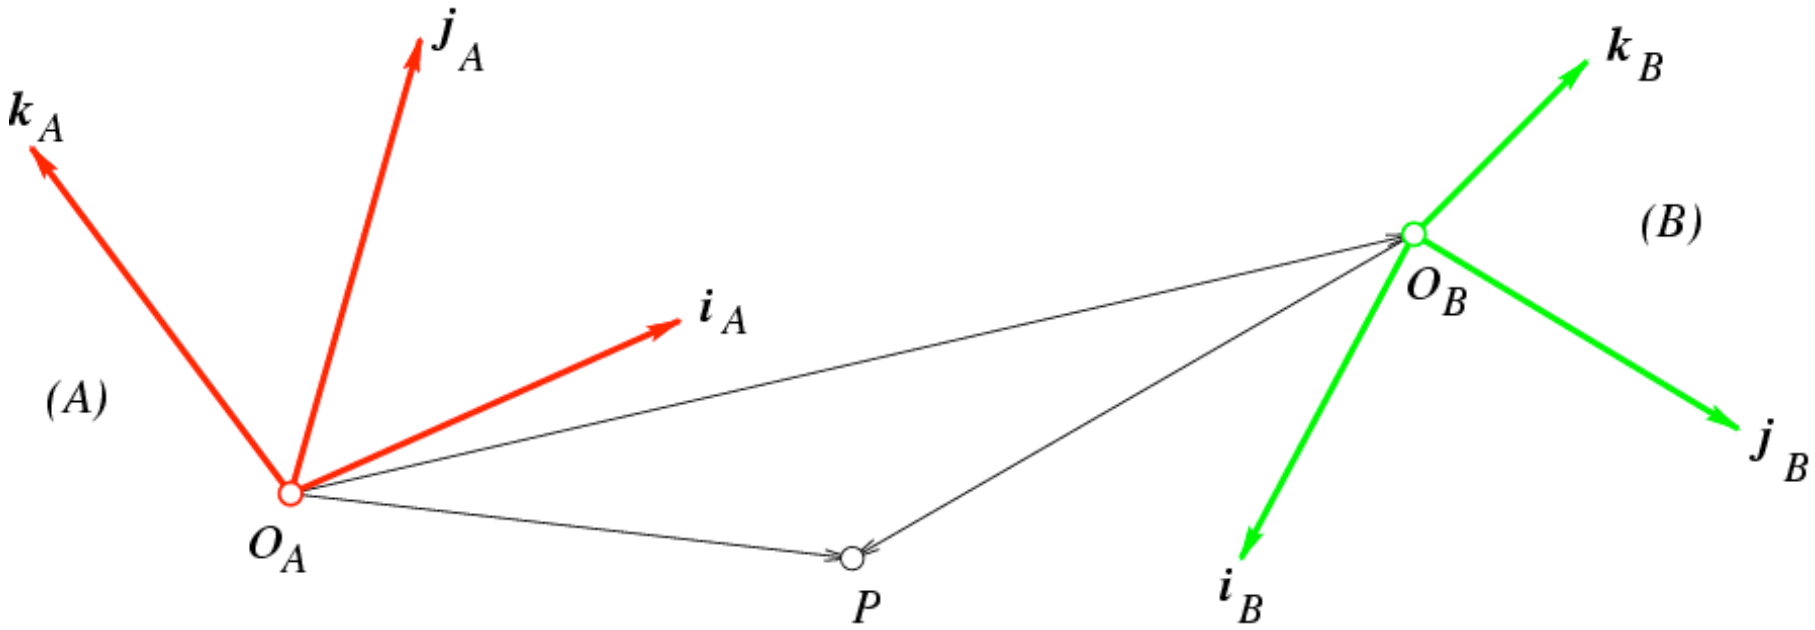

$$\begin{bmatrix} {}^{B}P \\ 1 \end{bmatrix} = \begin{bmatrix} {}^{B}_{A}R & {}^{B}O_A \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} {}^{A}P \\ 1 \end{bmatrix} = \begin{bmatrix} {}^{B}_{A}R & {}^{B}O_A \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{bmatrix} {}^{A}P \\ 1 \end{bmatrix}$$

The Extrinsic Parameters of a Camera

- When the camera frame (C) is different from the world frame (W),

$$\begin{pmatrix} {}^{C}P \\ 1 \end{pmatrix} = \begin{pmatrix} {}^{C}_{W}\mathcal{R} & {}^{C}O_{W} \\ \mathbf{0}^{T} & 1 \end{pmatrix} \begin{pmatrix} {}^{W}P \\ 1 \end{pmatrix}.$$

- Thus,

$$\boxed{\boldsymbol{p} = \frac{1}{z}\mathcal{M}\boldsymbol{P},} \quad \text{where} \quad \begin{cases} \mathcal{M} = \mathcal{K}(\mathcal{R} \quad \boldsymbol{t}), \\ \mathcal{R} = {}^{C}_{W}\mathcal{R}, \\ \boldsymbol{t} = {}^{C}O_{W}, \\ \boldsymbol{P} = \begin{pmatrix} {}^{W}P \\ 1 \end{pmatrix}. \end{cases} \quad {}^{W}P = \begin{pmatrix} {}^{W}x \\ {}^{W}y \\ {}^{W}z \end{pmatrix}$$

$$\begin{pmatrix} u \\ v \\ 1 \end{pmatrix} = \frac{1}{z}\begin{bmatrix} m_{12} & m_{12} & m_{13} & m_{14} \\ m_{21} & m_{22} & m_{23} & m_{24} \\ m_{31} & m_{32} & m_{33} & m_{34} \end{bmatrix} \begin{pmatrix} {}^{W}x \\ {}^{W}y \\ {}^{W}z \\ 1 \end{pmatrix}$$

- Note: z is *not* independent of $\mathcal{M}$ and $\boldsymbol{P}$:

$$\mathcal{M} = \begin{pmatrix} \boldsymbol{m}_1^T \\ \boldsymbol{m}_2^T \\ \boldsymbol{m}_3^T \end{pmatrix} \Longrightarrow z = \boldsymbol{m}_3 \cdot \boldsymbol{P}, \quad \text{or} \quad \begin{cases} u = \dfrac{\boldsymbol{m}_1 \cdot \boldsymbol{P}}{\boldsymbol{m}_3 \cdot \boldsymbol{P}}, \\ v = \dfrac{\boldsymbol{m}_2 \cdot \boldsymbol{P}}{\boldsymbol{m}_3 \cdot \boldsymbol{P}}. \end{cases}$$

Explicit Form of the Projection Matrix

$$\mathcal{M} = \begin{pmatrix} \alpha \boldsymbol{r}_1^T - \alpha \cot\theta \boldsymbol{r}_2^T + u_0 \boldsymbol{r}_3^T & \alpha t_x - \alpha \cot\theta t_y + u_0 t_z \\ \dfrac{\beta}{\sin\theta} \boldsymbol{r}_2^T + v_0 \boldsymbol{r}_3^T & \dfrac{\beta}{\sin\theta} t_y + v_0 t_z \\ \boldsymbol{r}_3^T & t_z \end{pmatrix}$$

Note: If $\mathcal{M} = (\mathcal{A} \quad \boldsymbol{b})$ then $|\boldsymbol{a}_3| = 1$.

Replacing $\mathcal{M}$ by $\lambda\mathcal{M}$ in

$$\begin{cases} u = \dfrac{\boldsymbol{m}_1 \cdot \boldsymbol{P}}{\boldsymbol{m}_3 \cdot \boldsymbol{P}} \\ v = \dfrac{\boldsymbol{m}_2 \cdot \boldsymbol{P}}{\boldsymbol{m}_3 \cdot \boldsymbol{P}} \end{cases}$$

does not change u and v.

M is only defined up to scale in this setting!!

Weak perspective (affine) camera

$$z_r = m_3^T P = \text{const.}$$

$$\begin{pmatrix} u \\ v \\ 1 \end{pmatrix} = \frac{1}{z_r} \begin{bmatrix} m_1^T \\ m_2^T \\ m_3^T \end{bmatrix} P = \begin{pmatrix} m_1^T P / m_3^T P \\ m_2^T P / m_3^T P \\ m_3^T P / m_3^T P \end{pmatrix}$$

$$\begin{pmatrix} u \\ v \end{pmatrix} = \frac{1}{z_r} \begin{bmatrix} m_1^T \\ m_2^T \end{bmatrix} P$$

$$\begin{pmatrix} u \\ v \end{pmatrix} = \underbrace{\begin{bmatrix} a_{11} & a_{12} & a_{13} \\ a_{21} & a_{22} & a_{23} \end{bmatrix}}_{A_{2\times 3}} \underbrace{\begin{bmatrix} b_1 \\ b_2 \end{bmatrix}}_{b_{2\times 1}} \begin{pmatrix} {}^W x \\ {}^W y \\ {}^W z \\ 1 \end{pmatrix} = A^W P + b$$

Re-cap: imaging and camera geometry (with a slight change of notation)

- perspective projection
- camera centre, image point and scene point are collinear
- an image point back projects to a ray in 3-space
- depth of the scene point is unknown

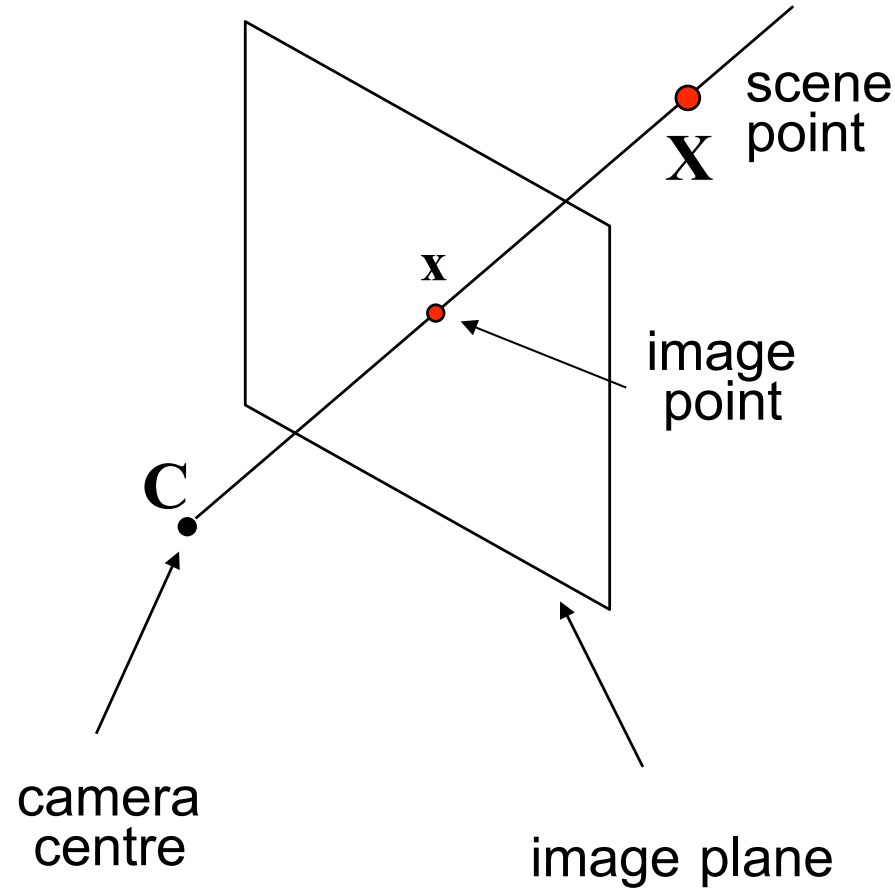

Slide credit: A. Zisserman

Notation – A Projective Camera

The camera model for perspective projection is a linear map between homogeneous point coordinates

$$\begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = \begin{bmatrix} & \mathtt{P}\ (3 \times 4) & \end{bmatrix} \begin{pmatrix} \mathrm{X} \\ \mathrm{Y} \\ \mathrm{Z} \\ 1 \end{pmatrix}$$

Image Point Scene Point

$$\lambda \mathbf{x} = \mathtt{P}\ \mathbf{X}$$

e.g. if $\mathtt{P} = [\mathtt{I}|0]$ then

$$x = \frac{\mathrm{X}}{\mathrm{Z}} \qquad y = \frac{\mathrm{Y}}{\mathrm{Z}}$$

- P has 11 degrees of freedom (essential parameters).

Slide credit: A. Zisserman

How a “scene plane” projects into an image?

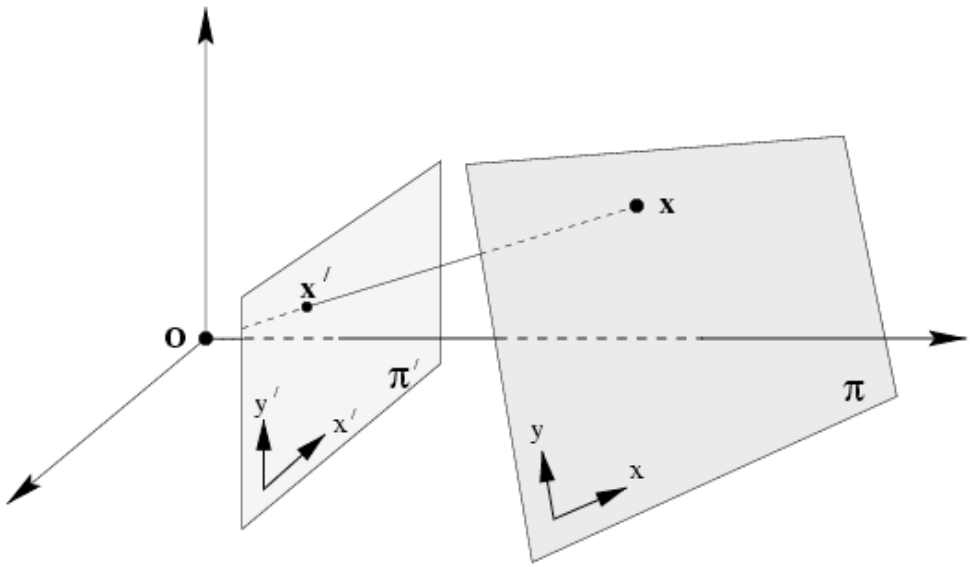

Plane projective transformations

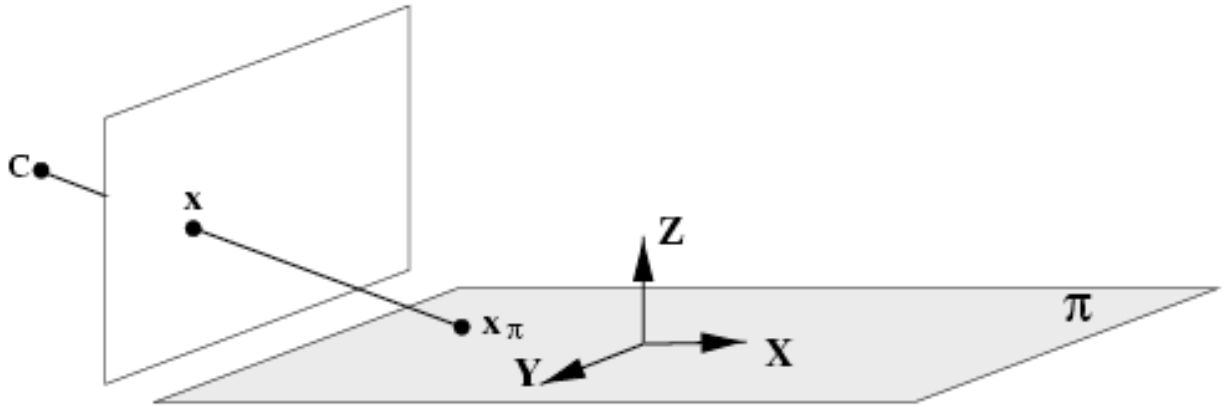

Choose the world coordinate system such that the plane of the points has zero z coordinate. Then the 3×4 matrix P reduces to

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{bmatrix} p_{11} & p_{12} & p_{13} & p_{14} \\ p_{21} & p_{22} & p_{23} & p_{24} \\ p_{31} & p_{32} & p_{33} & p_{34} \end{bmatrix} \begin{pmatrix} x \\ y \\ 0 \\ 1 \end{pmatrix} = \begin{bmatrix} p_{11} & p_{12} & p_{14} \\ p_{21} & p_{22} & p_{24} \\ p_{31} & p_{32} & p_{34} \end{bmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix}$$

which is a 3×3 matrix representing a general plane to plane projective transformation.

Slide credit: A. Zisserman

Projective transformations continued

$$\begin{pmatrix} x'_1 \\ x'_2 \\ x'_3 \end{pmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix}$$

or $\mathbf{x}' = \mathtt{H}\mathbf{x}$, where $\mathtt{H}$ is a 3×3 non-singular homogeneous matrix.

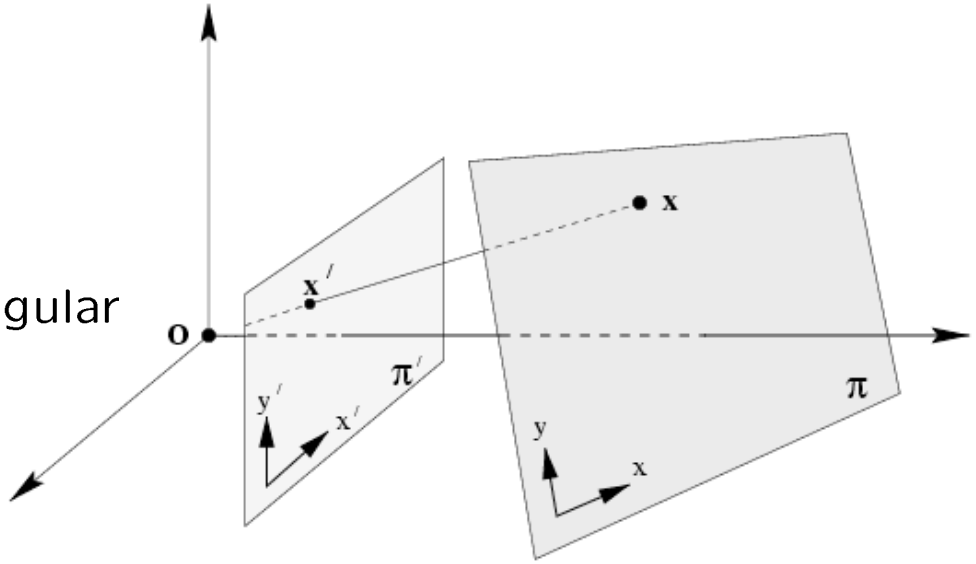

- This is the most general transformation between the world and image plane under imaging by a perspective camera.
- It is often only the 3 x 3 form of the matrix that is important in establishing properties of this transformation.
- A projective transformation is also called a ``homography'' and a ``collineation''.
- H has 8 degrees of freedom.

Slide credit: A. Zisserman

Planes under affine projection

$$\begin{pmatrix} x_1 \\ x_2 \end{pmatrix} = \begin{bmatrix} a_{11} & a_{12} & a_{13} & b_1 \\ a_{21} & a_{22} & a_{23} & b_2 \end{bmatrix} \begin{pmatrix} x \\ y \\ 0 \\ 1 \end{pmatrix} = \begin{bmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{bmatrix} \begin{pmatrix} x \\ y \end{pmatrix} + \begin{pmatrix} b_1 \\ b_2 \end{pmatrix} = A_{2\times 2} P + b_{2\times 1}$$

Points on a world plane map with a 2D affine geometric transformation (6 parameters)

Summary

• Affine projections induce affine transformations from planes onto their images.

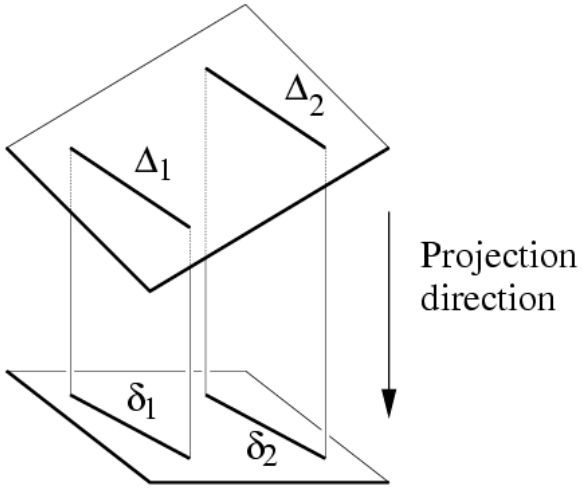

• Perspective projections induce projective transformations from planes onto their images.

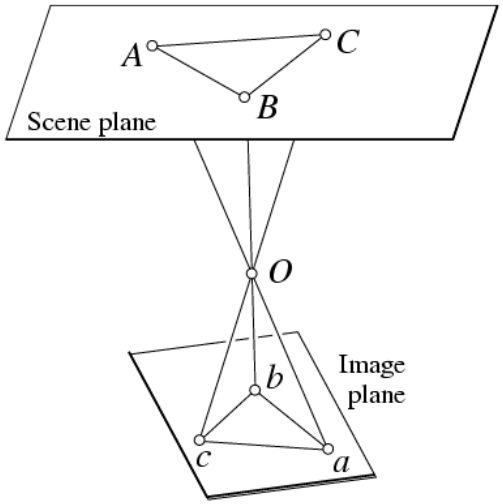

2D transformation models

Similarity
(translation,
scale, rotation)

Affine

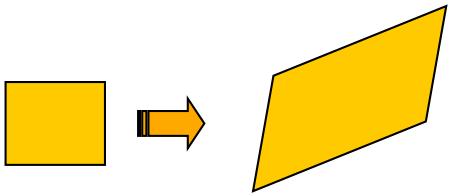

Projective
(homography)

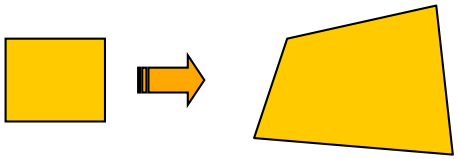

When is homography a valid transformation model?

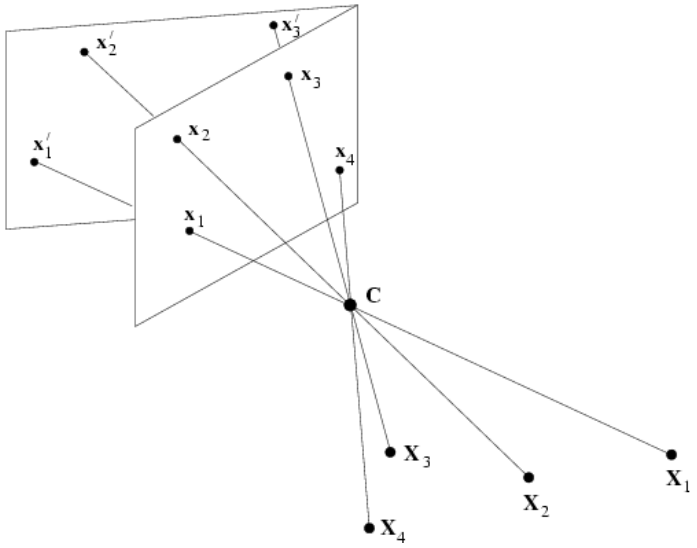

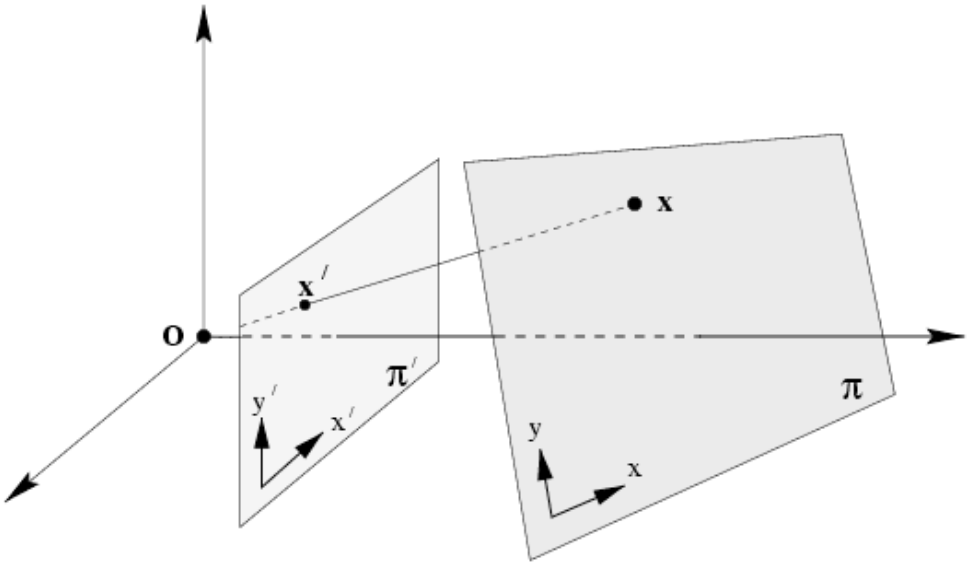

Case I: Plane projective transformations

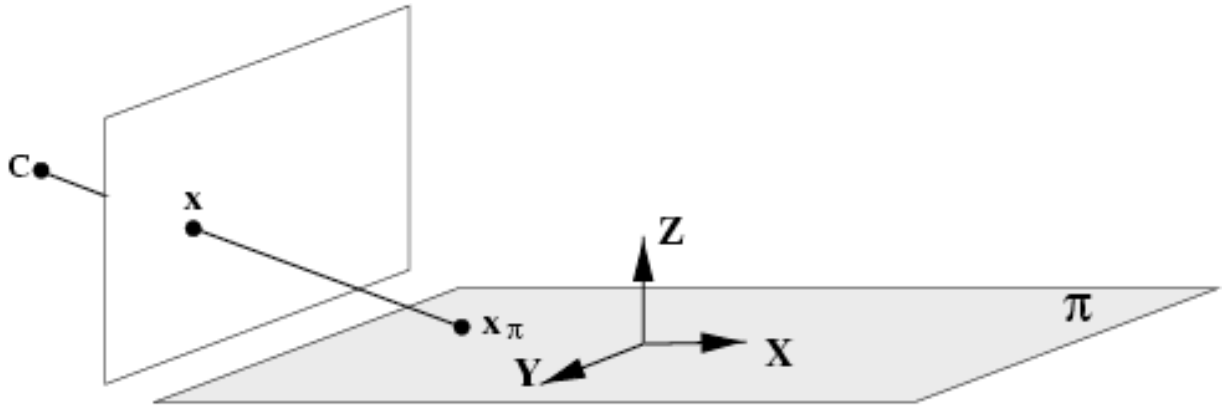

Choose the world coordinate system such that the plane of the points has zero z coordinate. Then the 3×4 matrix P reduces to

$$\begin{pmatrix} x_1 \\ x_2 \\ x_3 \end{pmatrix} = \begin{bmatrix} p_{11} & p_{12} & p_{13} & p_{14} \\ p_{21} & p_{22} & p_{23} & p_{24} \\ p_{31} & p_{32} & p_{33} & p_{34} \end{bmatrix} \begin{pmatrix} x \\ y \\ 0 \\ 1 \end{pmatrix} = \begin{bmatrix} p_{11} & p_{12} & p_{14} \\ p_{21} & p_{22} & p_{24} \\ p_{31} & p_{32} & p_{34} \end{bmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix}$$

which is a 3×3 matrix representing a general plane to plane projective transformation.

Slide credit: A. Zisserman

Case II: Cameras rotating about their centre

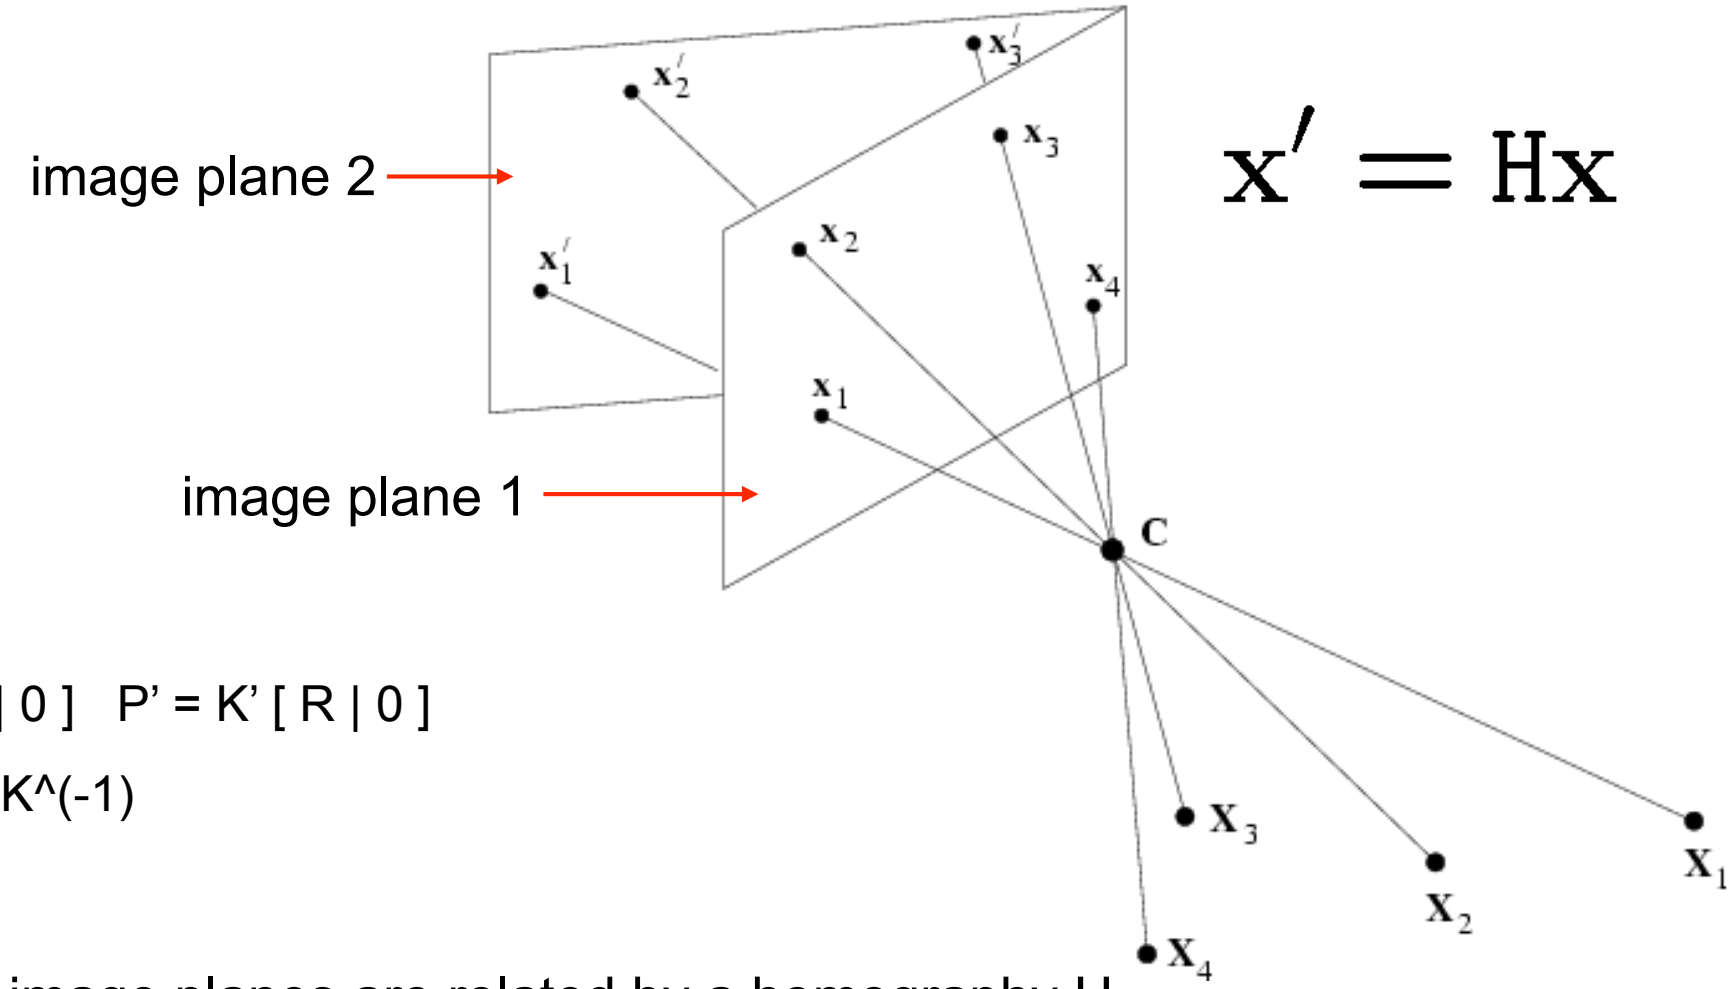

$$\mathbf{x}' = \mathrm{H}\mathbf{x}$$

P = K [I | 0] P' = K' [R | 0]

H = K' R K^(-1)

- The two image planes are related by a homography H
- H depends only on the relation between the image planes and camera centre, C, not on the 3D structure

Slide credit: A. Zisserman

Case II: Cameras rotating about their centre

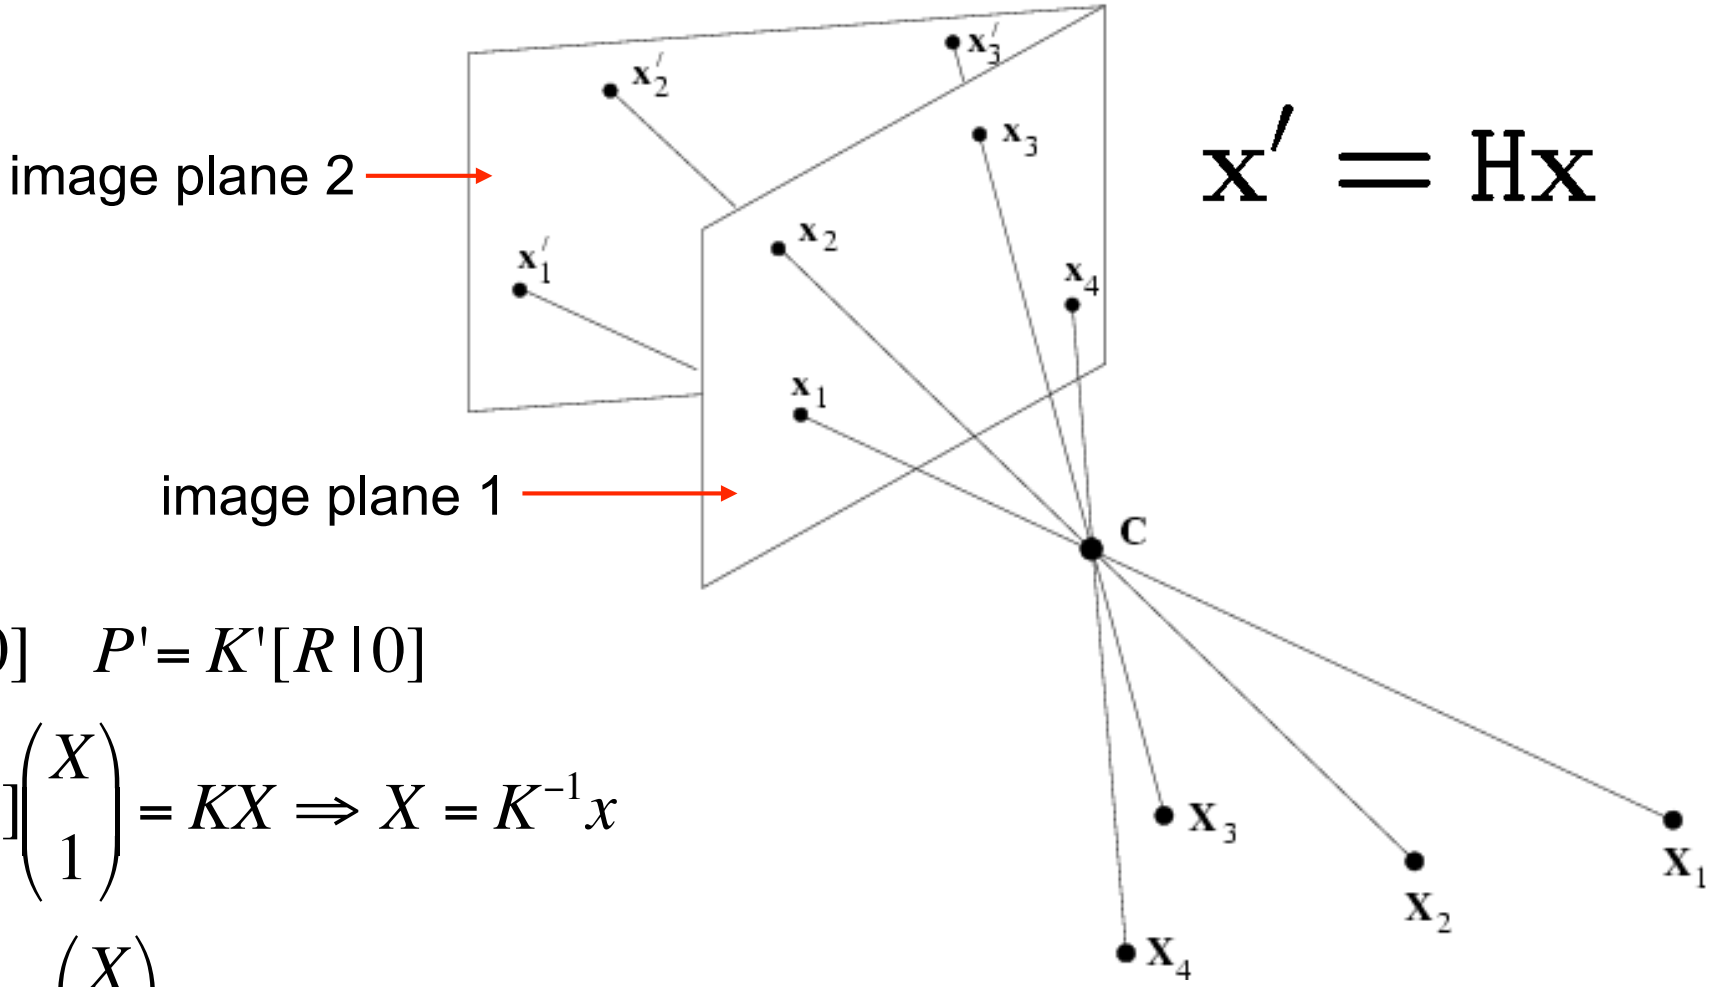

$$\mathbf{x}' = \mathrm{H}\mathbf{x}$$

$$P = K[I\,|\,0] \quad P' = K'[R\,|\,0]$$

$$x = K[I\,|\,0]\begin{pmatrix} X \\ 1 \end{pmatrix} = KX \Rightarrow X = K^{-1}x$$

$$x' = K'[R\,|\,0]\begin{pmatrix} X \\ 1 \end{pmatrix} = K'RX$$

$$x' = \underbrace{K'RK^{-1}}_{H}x$$

Slide credit: A. Zisserman

Outline – the rest of the lecture

Part 2. Image matching and recognition with local features

- **Correspondence**
- **Semi-local and global geometric relations**
- **Robust estimation – RANSAC and Hough Transform**

Image matching and recognition with local features

The goal: establish correspondence between two or more images

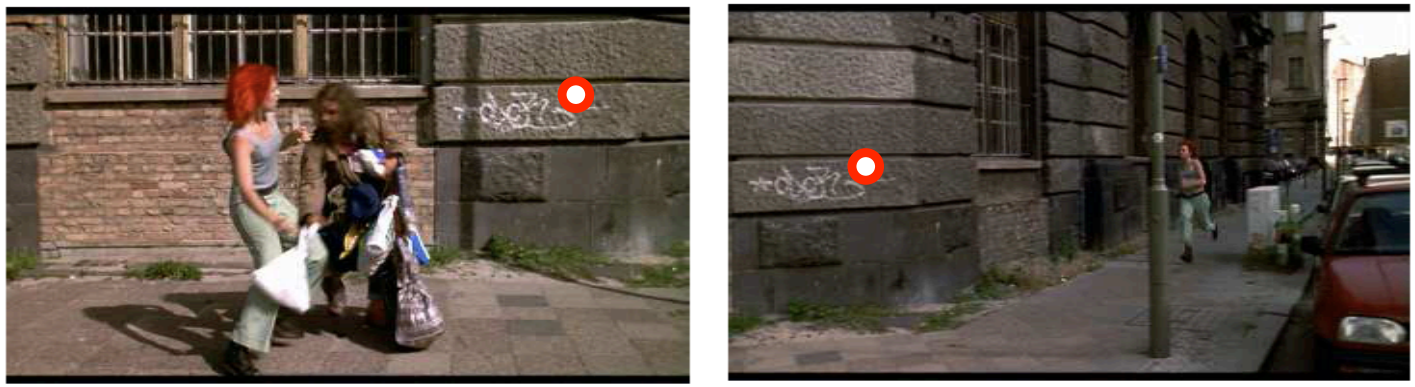

$\mathbf{x} = \mathrm{P}\mathbf{X} \quad \mathbf{x}' = \mathrm{P}'\mathbf{X}$

P : 3×4 matrix

$\mathbf{X}$: 4-vector

$\mathbf{x}$: 3-vector

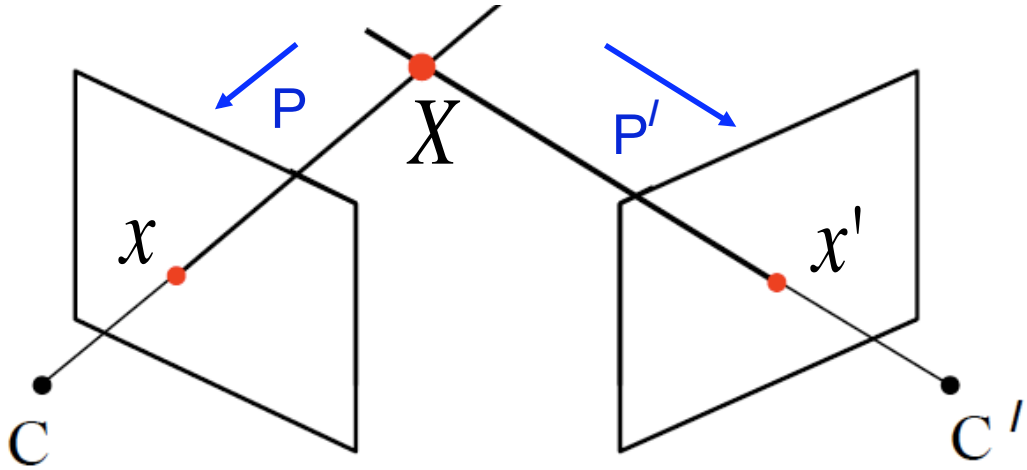

Image points x and x' are in correspondence if they are projections of the same 3D scene point X.

Images courtesy A. Zisserman

Example I: Wide baseline matching

Establish correspondence between two (or more) images.

Useful in visual geometry: Camera calibration, 3D reconstruction, Structure and motion estimation, …

Scale/affine – invariant regions: SIFT, Harris-Laplace, etc.

Example II: Object recognition

Establish correspondence between the target image and (multiple) images in the model database.

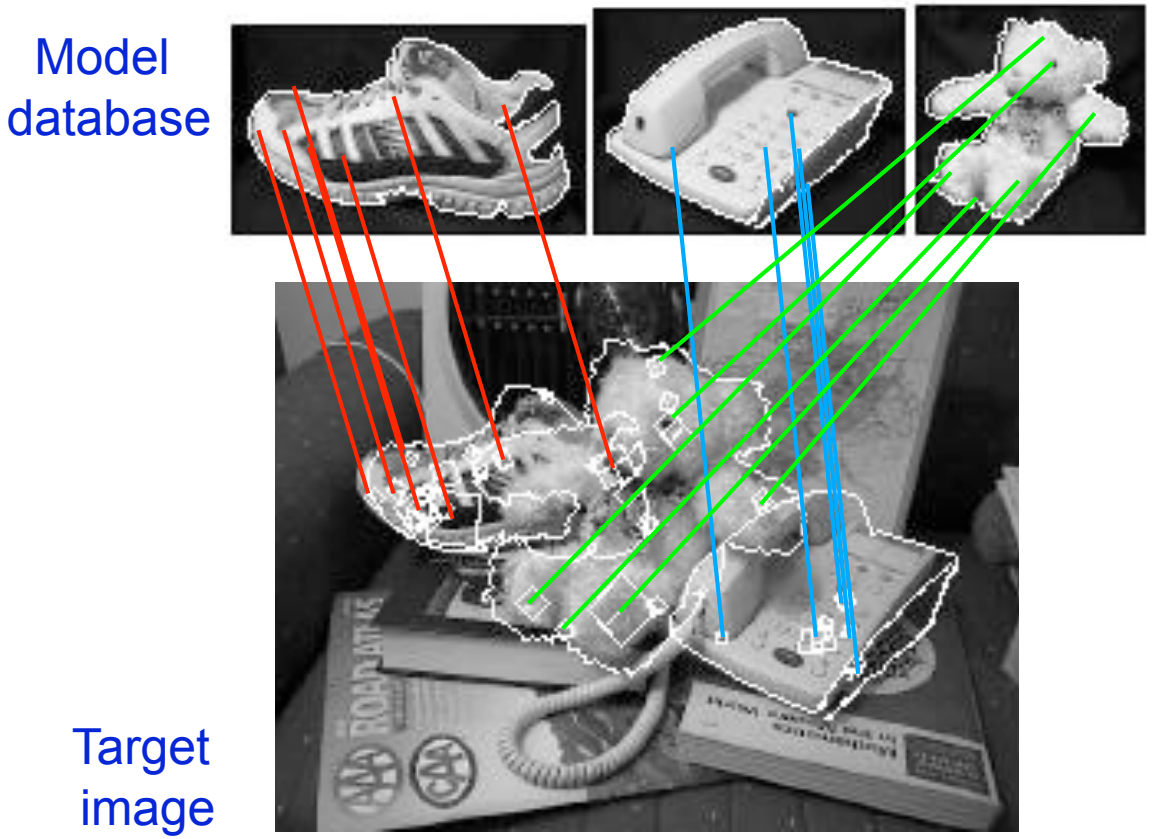

[D. Lowe, 1999]

Example III: Visual search

Given a query image, find images depicting the same place / object in a large unordered image collection.

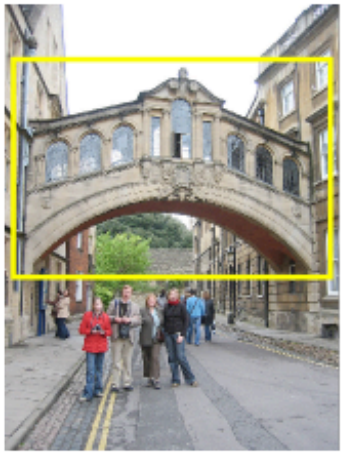

Find these landmarks

...in these images and 1M more

Establish correspondence between the query image and all images from the database depicting the same object / scene.

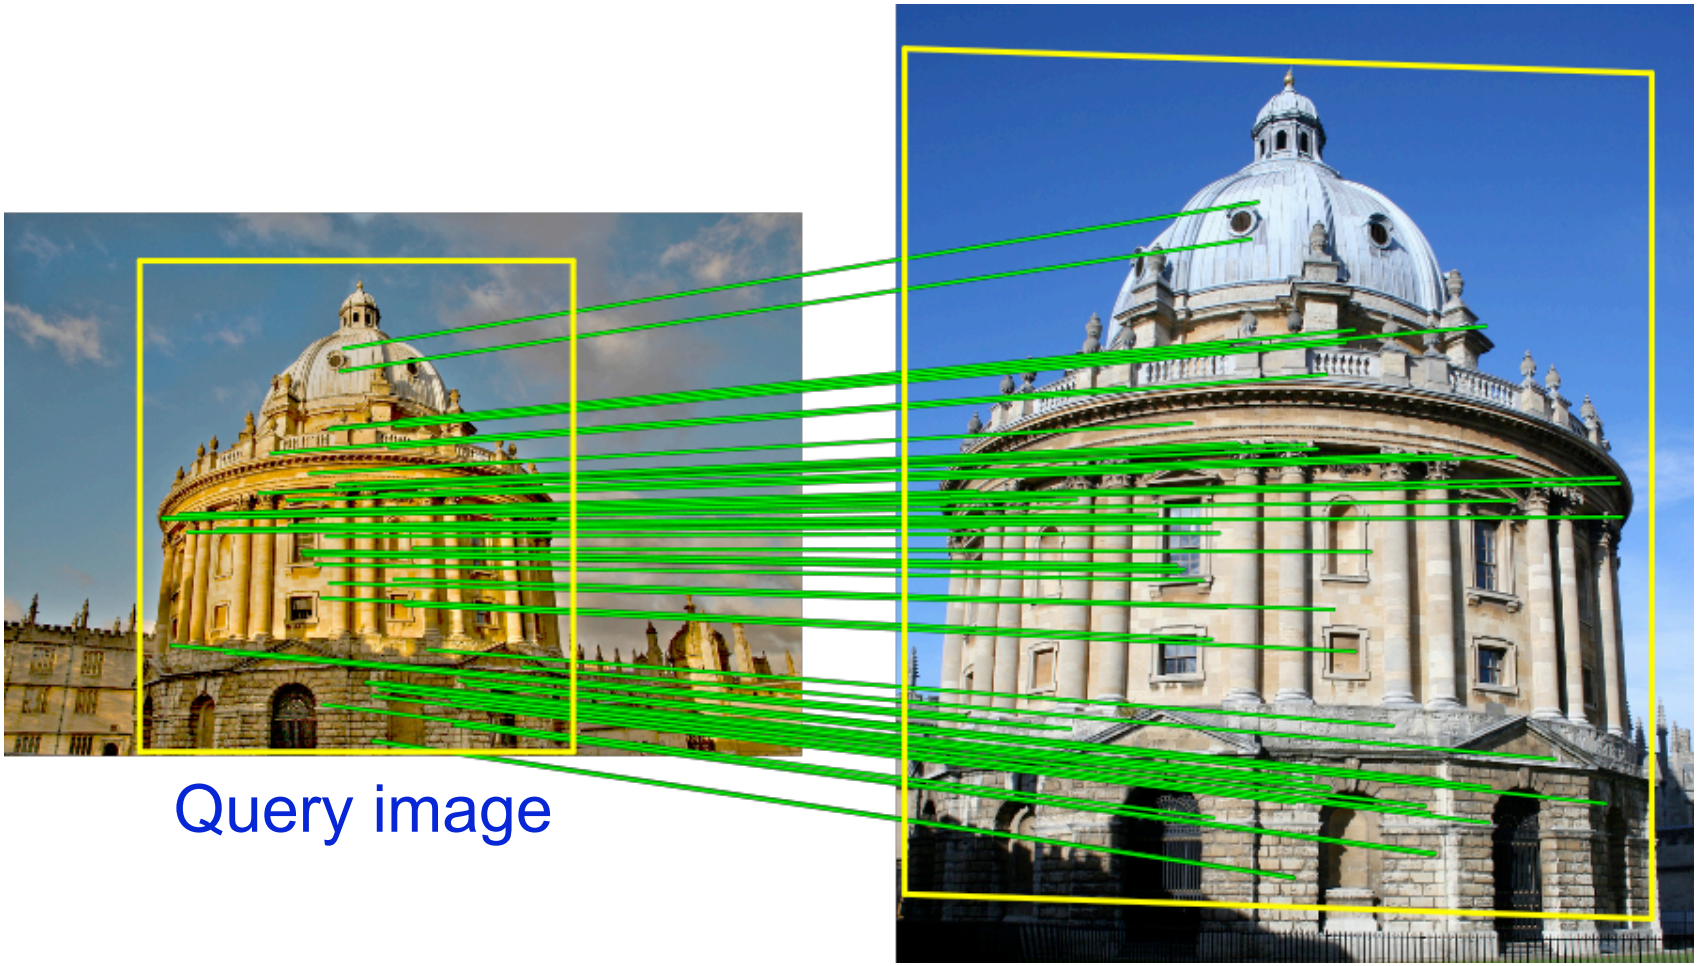

Query image

Database image(s)

Why is it difficult?

Want to establish correspondence despite possibly large changes in scale, viewpoint, lighting and partial occlusion

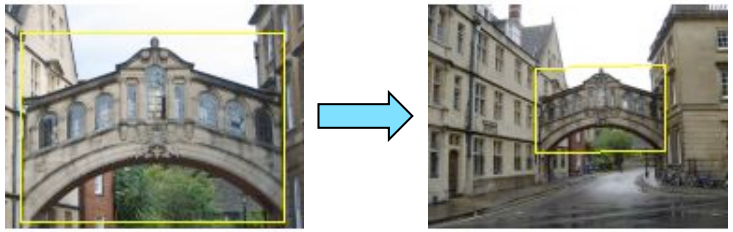

Scale

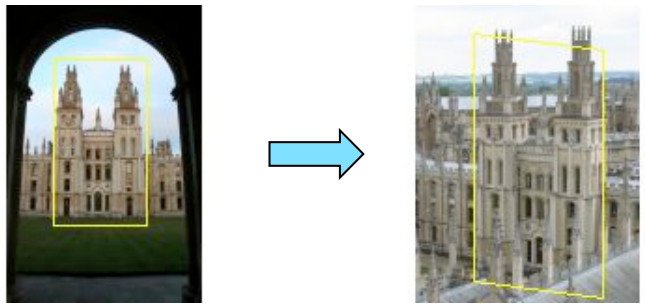

Viewpoint

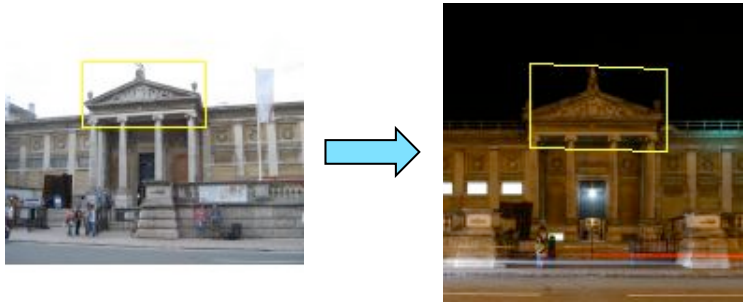

Lighting

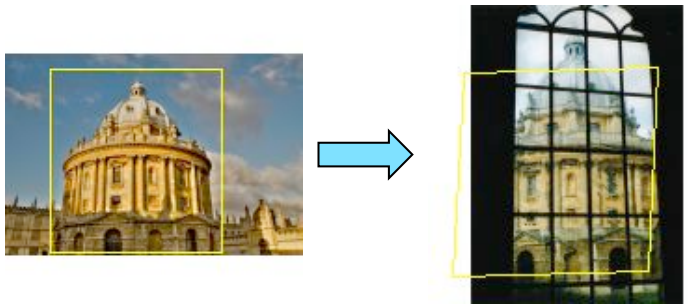

Occlusion

… and the image collection can be very large (e.g. 1M images)

Approach

Pre-processing (last lecture):

- Detect local features.
- Extract descriptor for each feature.

Matching:

1. Establish tentative (putative) correspondences based on local appearance of individual features (their descriptors).

2. Verify matches based on semi-local / global geometric relations.

Example I: Two images -“Where is the Graffiti?”

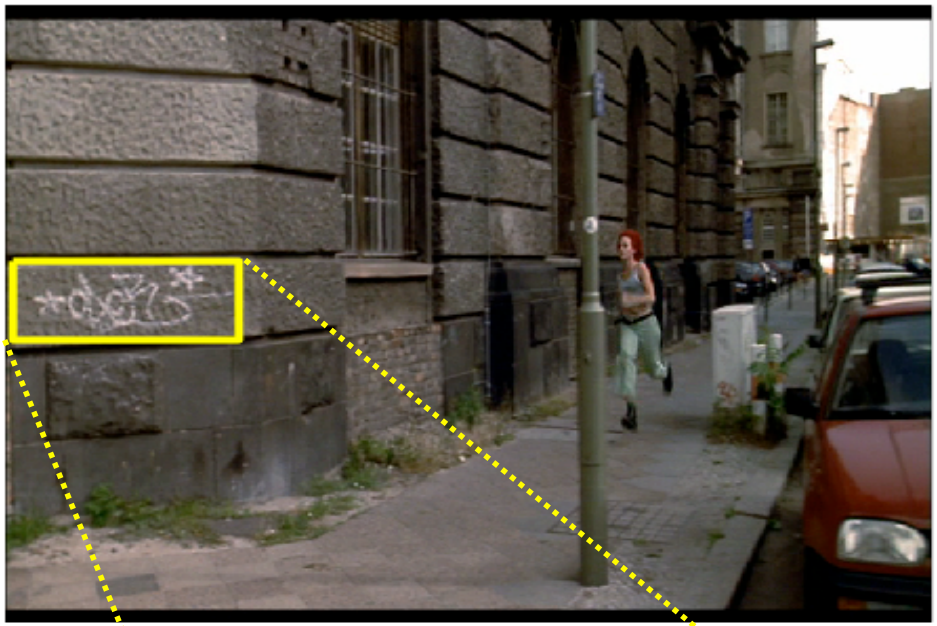

Step 1. Establish tentative correspondence

Establish tentative correspondences between object model image and target image by nearest neighbour matching on SIFT vectors

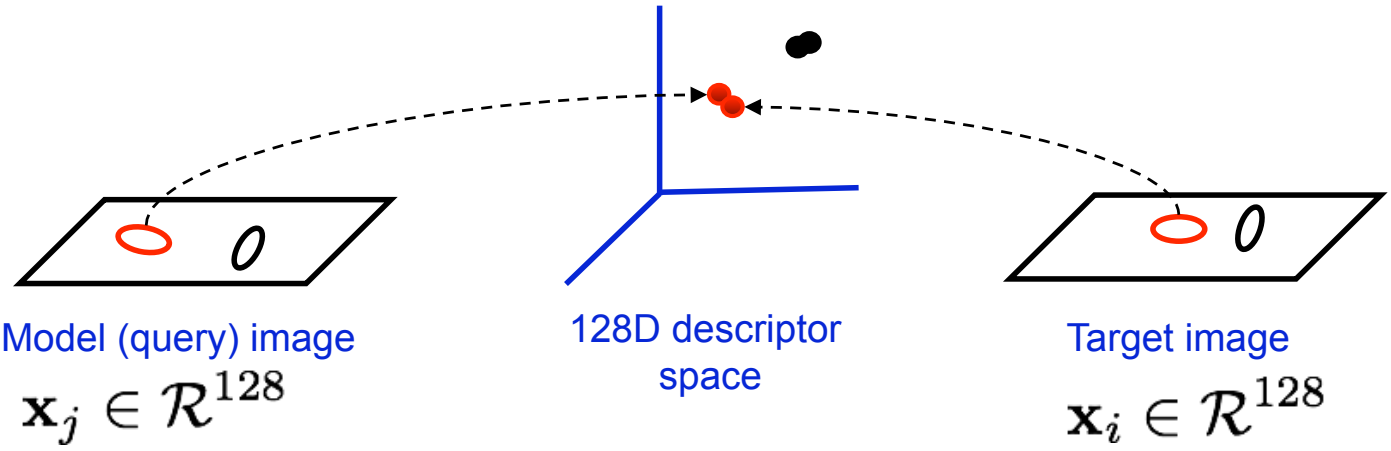

Need to solve some variant of the "nearest neighbor problem" for all feature vectors, $\mathbf{x}_j \in \mathcal{R}^{128}$, in the query image:

$$\forall j \; NN(j) = \arg\min_i ||\mathbf{x}_i - \mathbf{x}_j||,$$

where, $\mathbf{x}_i \in \mathcal{R}^{128}$, are features in the target image.

Can take a long time if many target images are considered.

Problem with matching on local descriptors alone

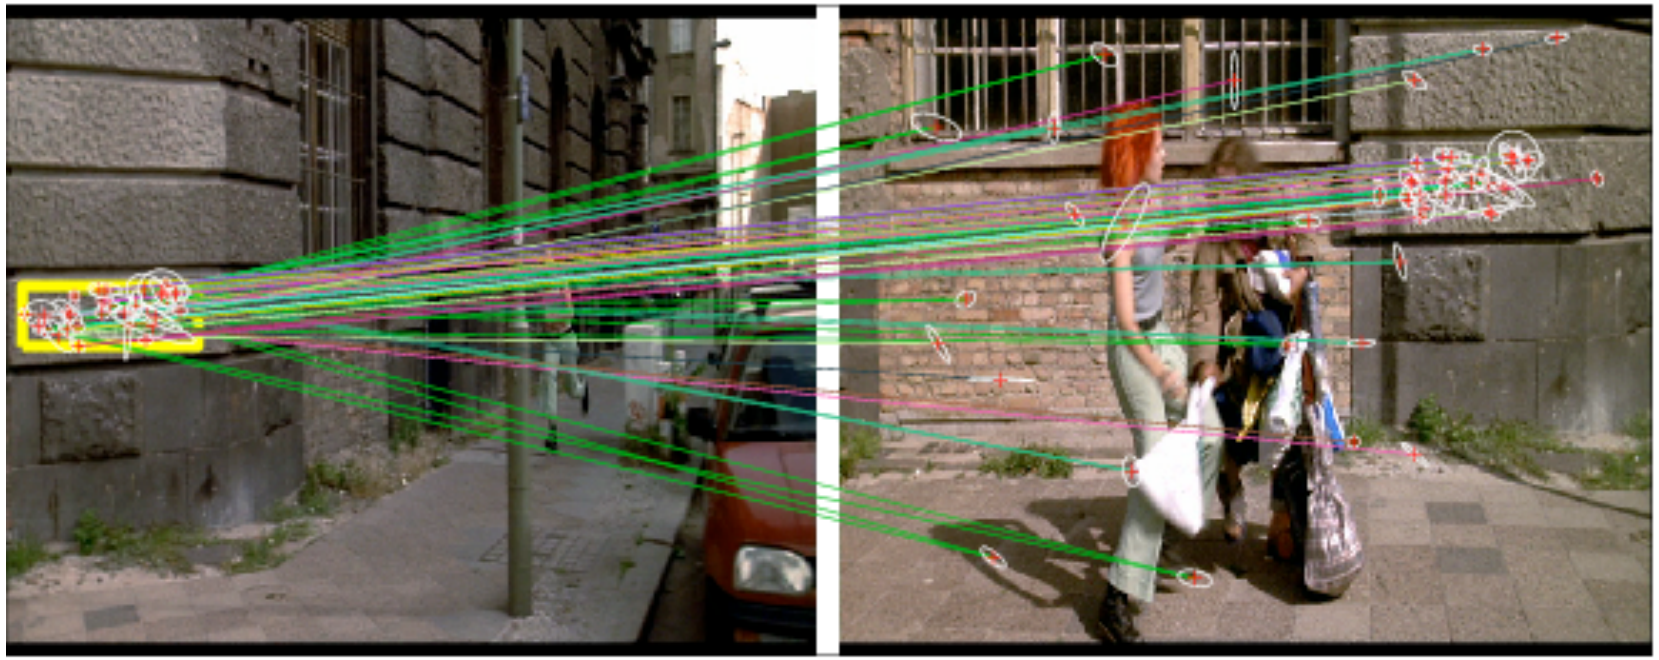

- too much individual invariance
- each region can affine deform independently (by different amounts)
- Locally appearance can be ambiguous

Solution: use semi-local and global spatial relations to verify matches.

Example I: Two images -“Where is the Graffiti?”

Initial matches

Nearest-neighbor search based on appearance descriptors alone.

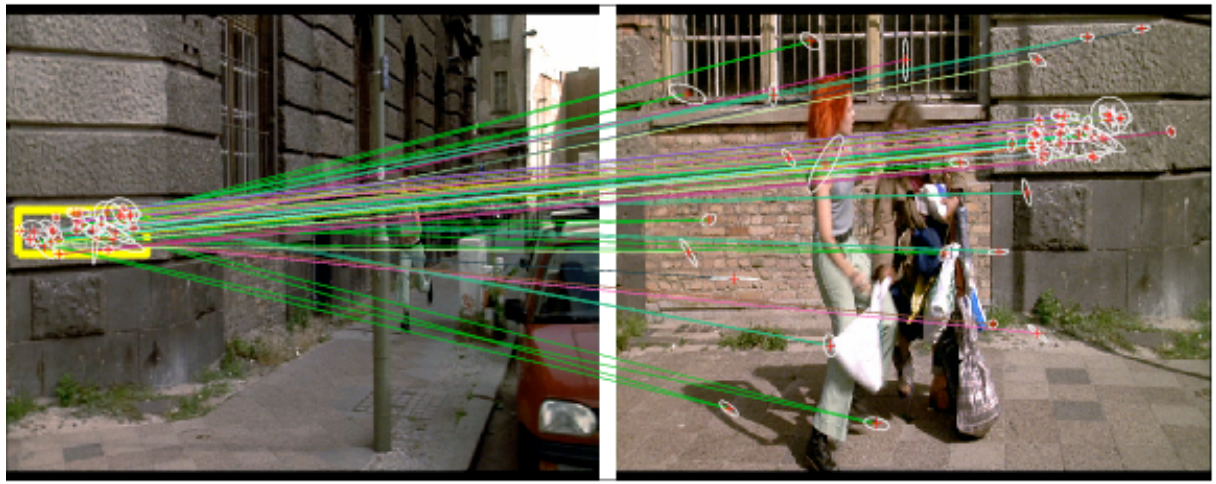

After spatial verification

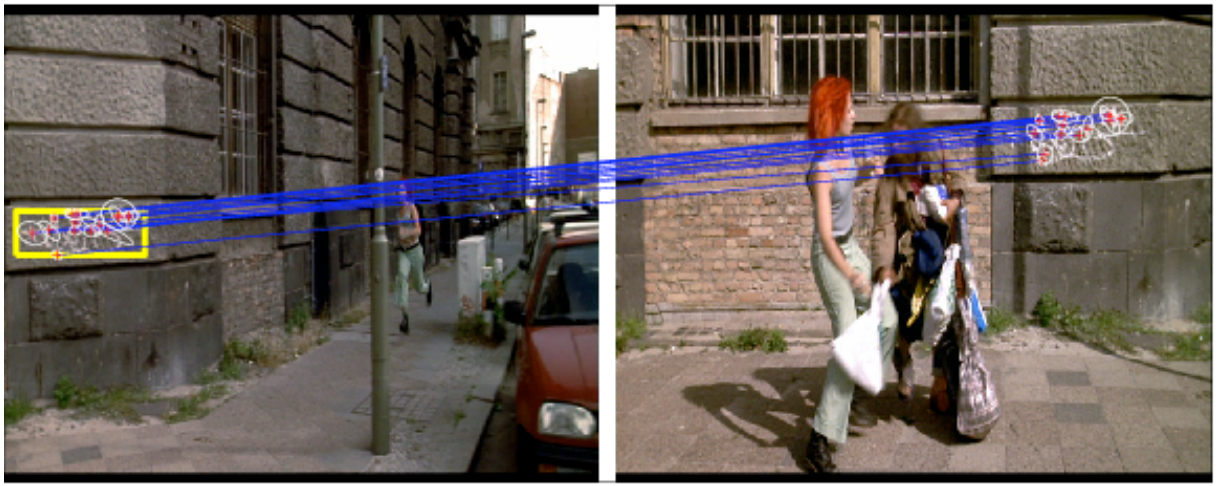

Step 2: Spatial verification (now)

a. Semi-local constraints

Constraints on spatially close-by matches

b. Global geometric relations

Require a consistent global relationship between all matches

Semi-local constraints: Example I. – neighbourhood consensus

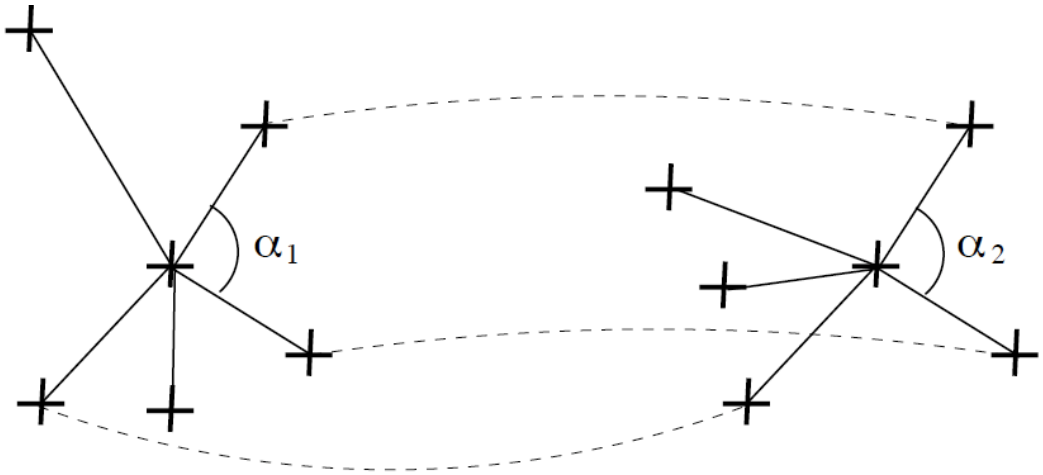

Fig. 4. Semi-local constraints : neighbours of the point have to match and angles have to correspond. Note that not all neighbours have to be matched correctly.

[Schmid&Mohr, PAMI 1997]

Semi-local constraints:

Example I. – neighbourhood consensus

Original images

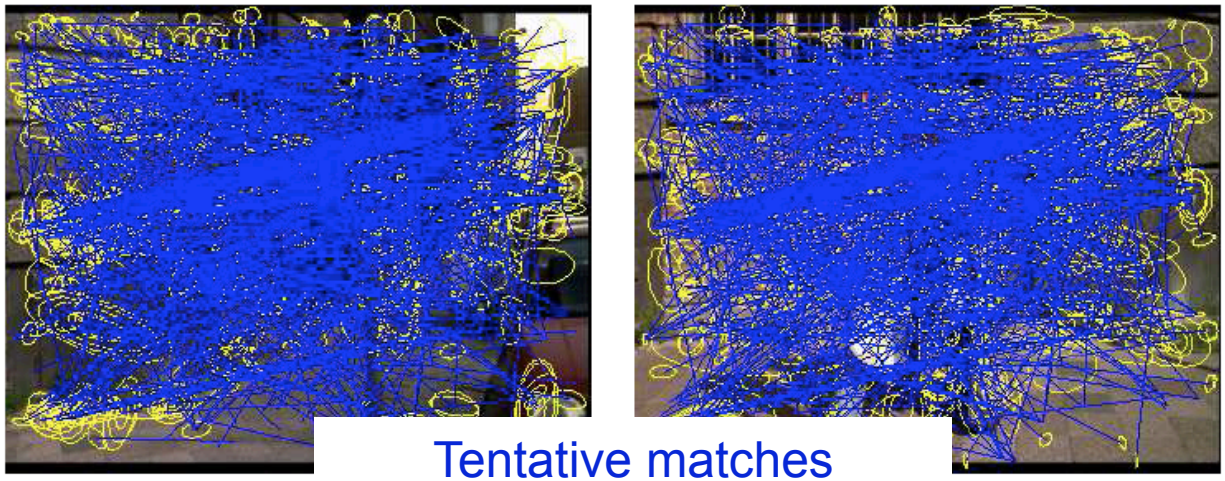

Tentative matches

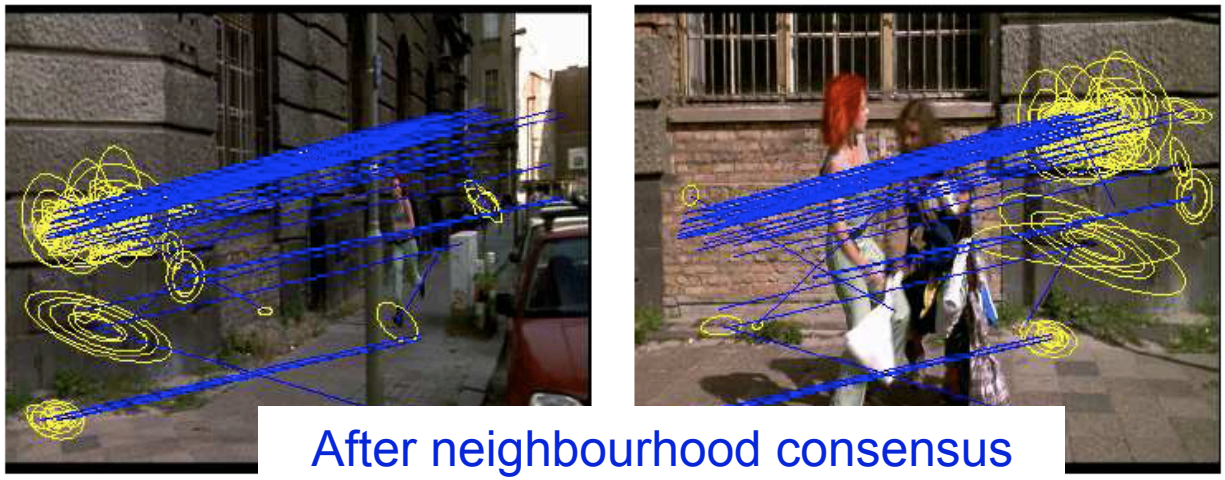

After neighbourhood consensus

[Schaffalitzky & Zisserman, CIVR 2004]

Semi-local constraints: Example II.

a

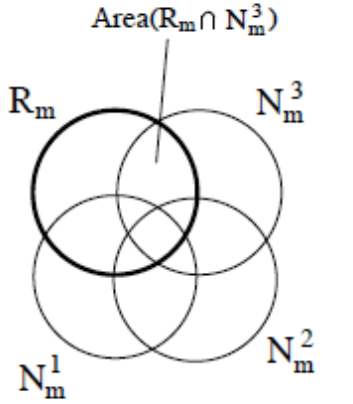

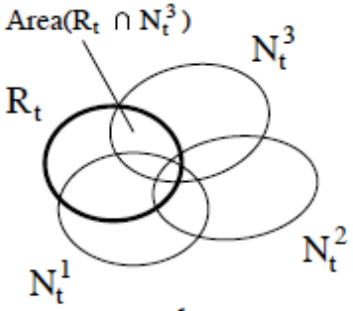

b

Figure 5: *Surface contiguity filter. a) the pattern of intersection between neighboring correct region matches is preserved by transformations between the model and the test images, because the surface is contiguous and smooth. b) the filter evaluates this property by testing the conservation of the area ratios.*

[Ferrari et al., IJCV 2005]

Model image

Matched image

Matched image

Geometric verification with global constraints

- All matches must be consistent with a global geometric relation / transformation.

- Need to simultaneously (i) estimate the geometric relation / transformation and (ii) the set of consistent matches

Tentative matches

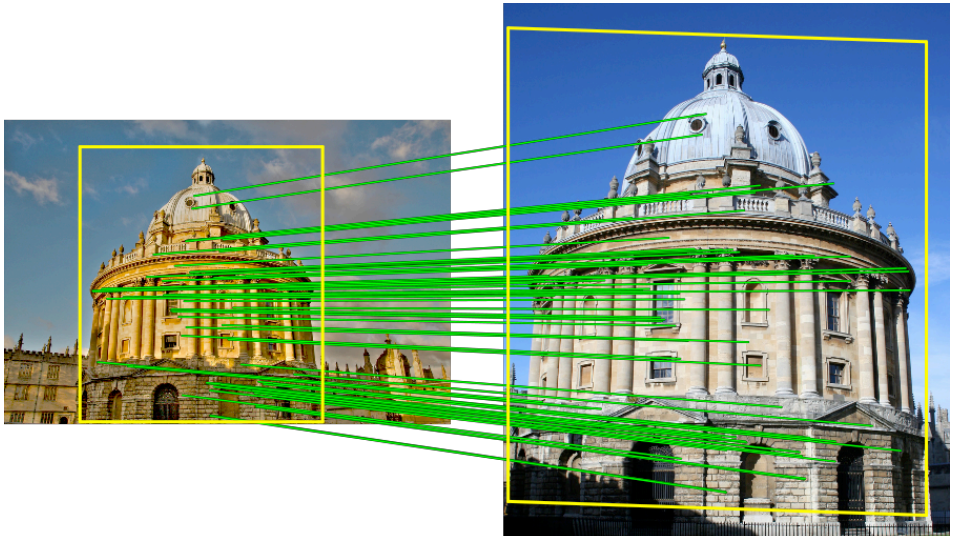

Matches consistent with an affine transformation

Epipolar geometry (not considered here)

In general, two views of a 3D scene are related by the epipolar constraint.

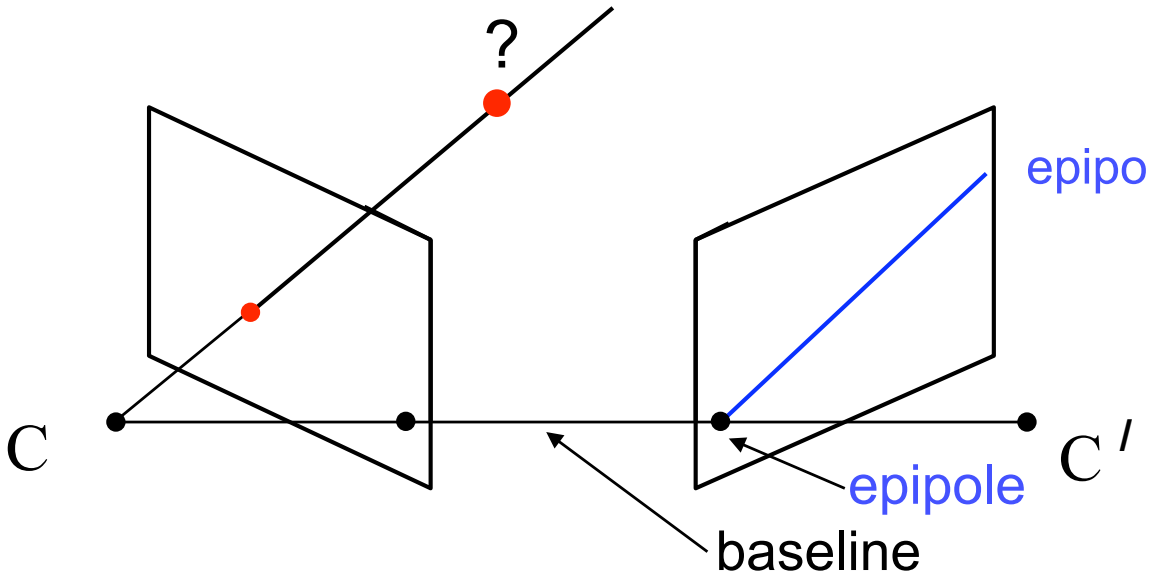

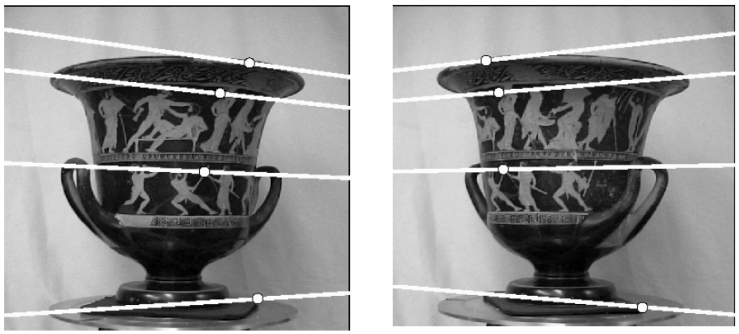

- A point in one view “generates” an epipolar line in the other view
- The corresponding point lies on this line.

Slide credit: A. Zisserman

Epipolar geometry (not considered here)

Epipolar geometry is a consequence of the coplanarity of the camera centres and scene point

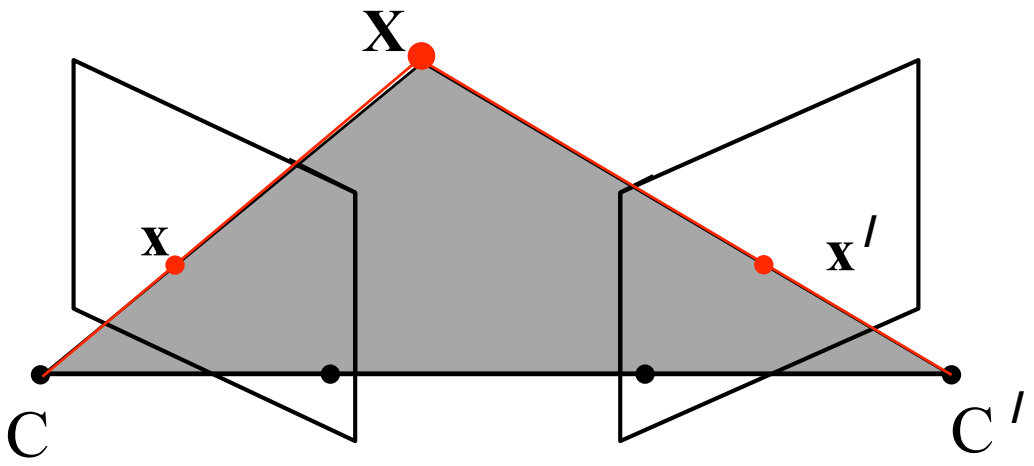

The camera centres, corresponding points and scene point lie in a single plane, known as the epipolar plane

Slide credit: A. Zisserman

Epipolar geometry (not considered here)

Algebraically, the epipolar constraint can be expressed as

$$\mathbf{x}'^{\top}\mathbf{F}\mathbf{x} = 0$$

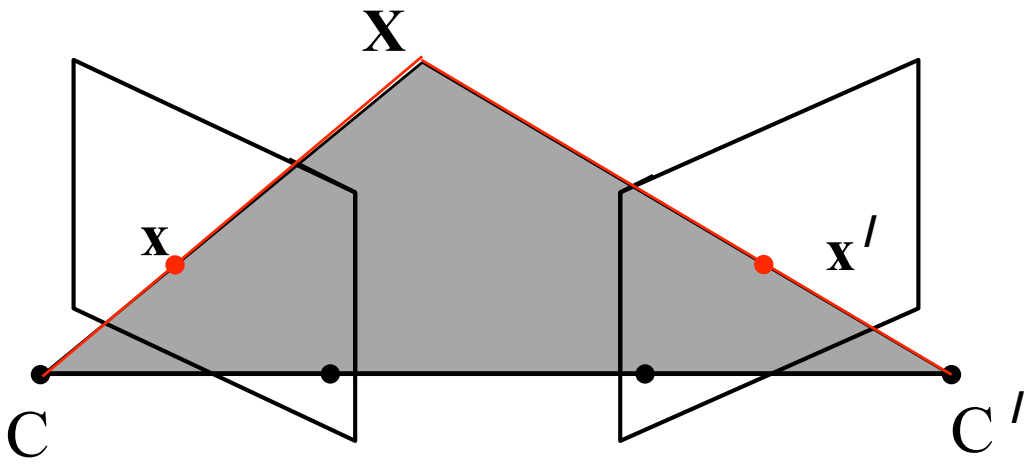

where

- x, x' are homogeneous coordinates (3-vectors) of **corresponding** image points.

- F is a 3x3, rank 2 homogeneous matrix with 7 degrees of freedom, called the **fundamental matrix**.

Slide credit: A. Zisserman

3D constraint: example (not considered here)

- Matches must be consistent with a 3D model

3 (out of 20) images used to build the 3D model

(a)

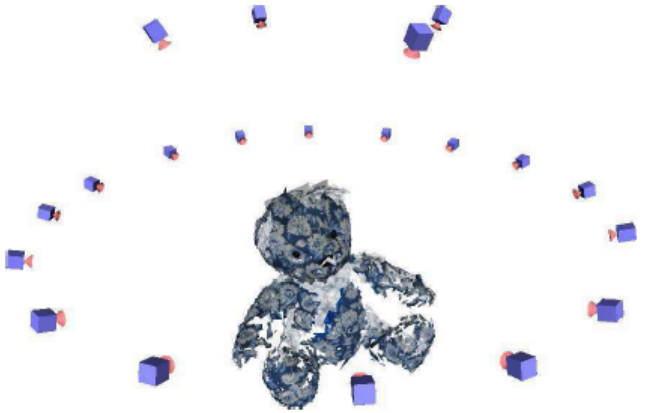

Recovered 3D model

Object recognized in a previously unseen pose

Recovered pose

(d)

[Lazebnik, Rothganger, Schmid, Ponce, CVPR'03]

3D constraint: example (not considered here)

With a given 3D model (set of known X's) and a set of measured image points x, the goal is to find find camera matrix P and a set of geometrically consistent correspondences $x \leftrightarrow X$.

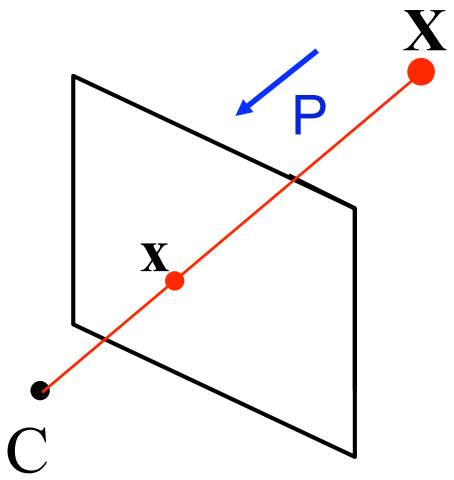

$$\mathbf{x} = \mathrm{P}\mathbf{X}$$

P : 3 × 4 matrix

x : 4-vector

x : 3-vector

2D transformation models

Similarity
(translation,
scale, rotation)

Affine

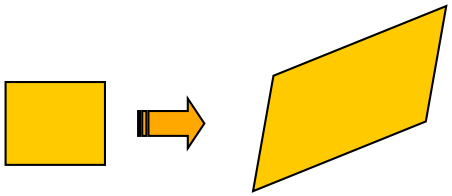

Projective
(homography)

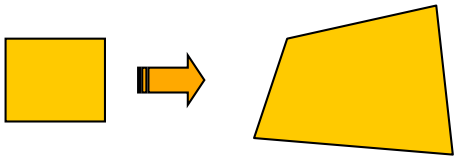

Example: estimating 2D affine transformation

- Simple fitting procedure (linear least squares)
- Approximates viewpoint changes for roughly planar objects and roughly orthographic cameras
- Can be used to initialize fitting for more complex models

Example: estimating 2D affine transformation

- Simple fitting procedure (linear least squares)
- Approximates viewpoint changes for roughly planar objects and roughly orthographic cameras
- Can be used to initialize fitting for more complex models

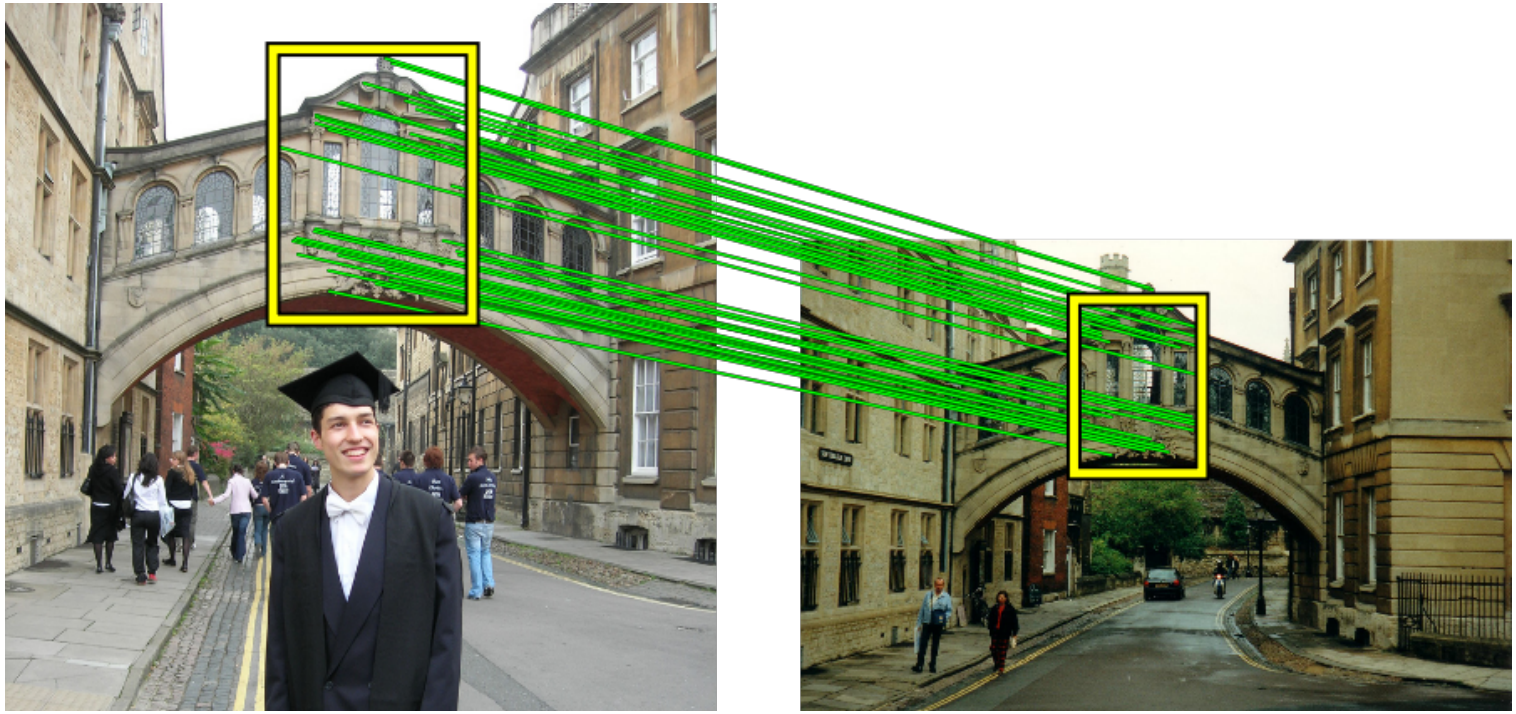

Matches consistent with an affine transformation

Fitting an affine transformation

Assume we know the correspondences, how do we get the transformation?

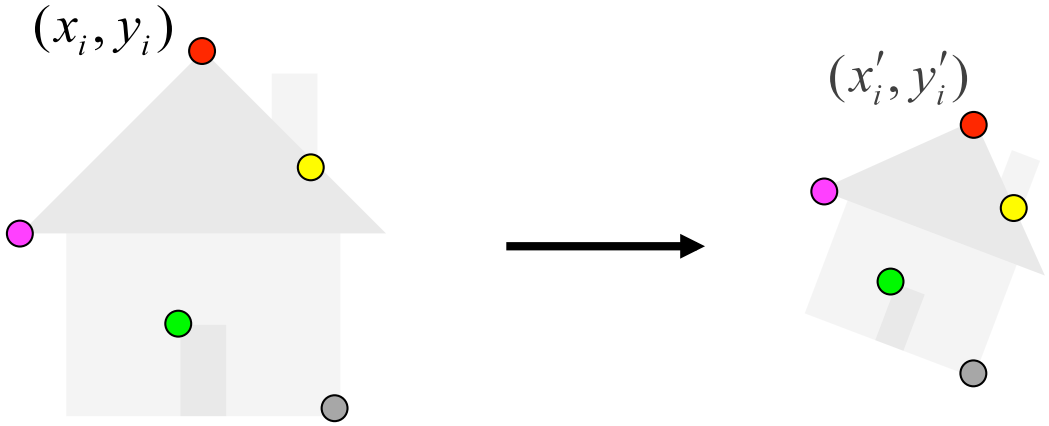

$$\begin{bmatrix} x_i' \\ y_i' \end{bmatrix} = \begin{bmatrix} m_1 & m_2 \\ m_3 & m_4 \end{bmatrix} \begin{bmatrix} x_i \\ y_i \end{bmatrix} + \begin{bmatrix} t_1 \\ t_2 \end{bmatrix}$$

$$\begin{bmatrix} & \cdots & & & & \\ x_i & y_i & 0 & 0 & 1 & 0 \\ 0 & 0 & x_i & y_i & 0 & 1 \\ & \cdots & & & & \end{bmatrix} \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \\ t_1 \\ t_2 \end{bmatrix} = \begin{bmatrix} \cdots \\ x_i' \\ y_i' \\ \cdots \end{bmatrix}$$

Fitting an affine transformation

$$\begin{bmatrix} & & \cdots & & & \\ x_i & y_i & 0 & 0 & 1 & 0 \\ 0 & 0 & x_i & y_i & 0 & 1 \\ & & \cdots & & & \end{bmatrix} \begin{bmatrix} m_1 \\ m_2 \\ m_3 \\ m_4 \\ t_1 \\ t_2 \end{bmatrix} = \begin{bmatrix} \cdots \\ x_i' \\ y_i' \\ \cdots \end{bmatrix}$$

Linear system with six unknowns

Each match gives us two linearly independent equations: need at least three to solve for the transformation parameters

Dealing with outliers

The set of putative matches may contain a high percentage (e.g. 90%) of outliers

How do we fit a geometric transformation to a small subset of all possible matches?

Possible strategies:

- RANSAC
- Hough transform

Strategy 1: RANSAC

RANSAC loop (Fischler & Bolles, 1981):

- Randomly select a *seed group* of matches
- Compute transformation from seed group
- Find *inliers* to this transformation
- If the number of inliers is sufficiently large, re-compute least-squares estimate of transformation on all of the inliers
- Keep the transformation with the largest number of inliers

Example: Robust line estimation - RANSAC

Fit a line to 2D data containing outliers

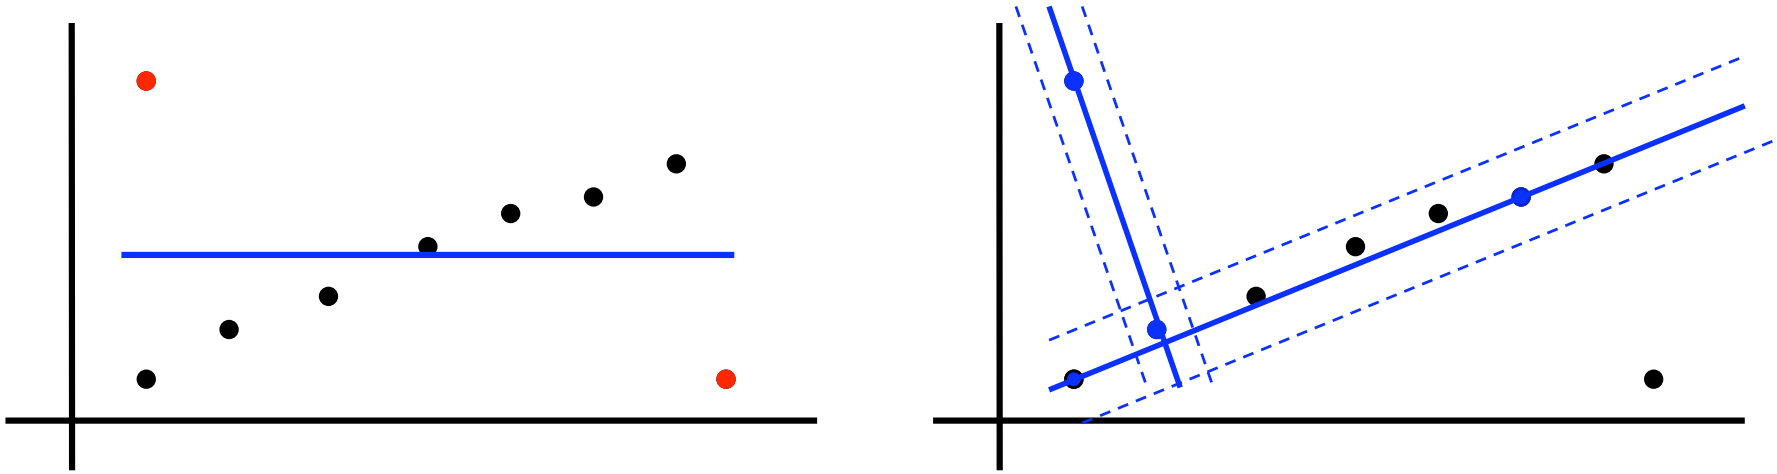

There are two problems

1. a line fit which minimizes perpendicular distance
2. a classification into inliers (valid points) and outliers

Solution: use robust statistical estimation algorithm RANSAC

(RANdom Sample Consensus) [Fishler & Bolles, 1981]

Slide credit: A. Zisserman

RANSAC robust line estimation

Repeat

1. Select random sample of 2 points
2. Compute the line through these points
3. Measure support (number of points within threshold distance of the line)

Choose the line with the largest number of inliers

- Compute least squares fit of line to inliers (regression)

Slide credit: A. Zisserman

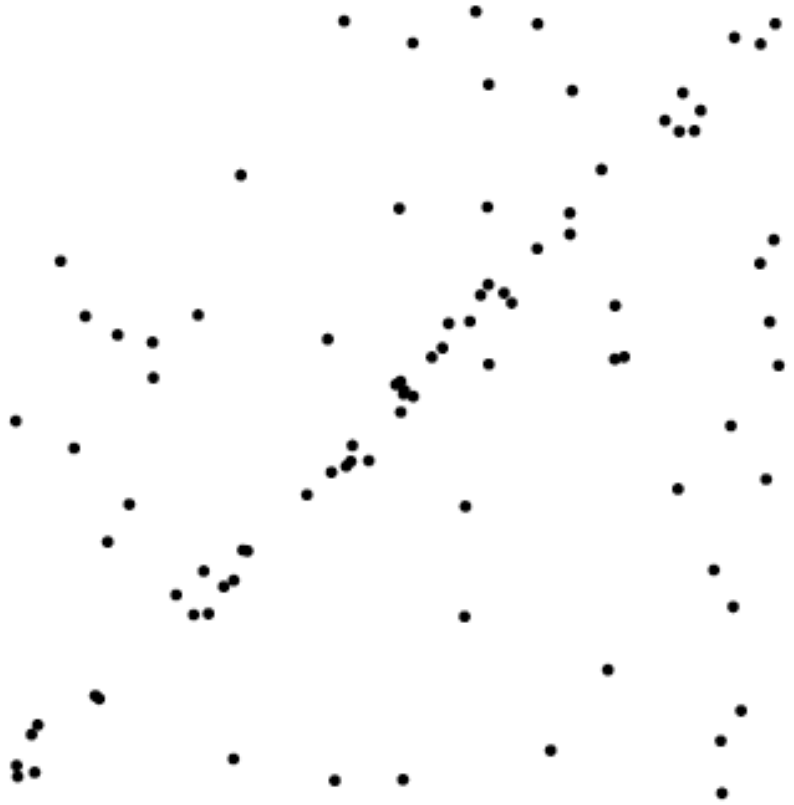

Slide credit: O. Chum

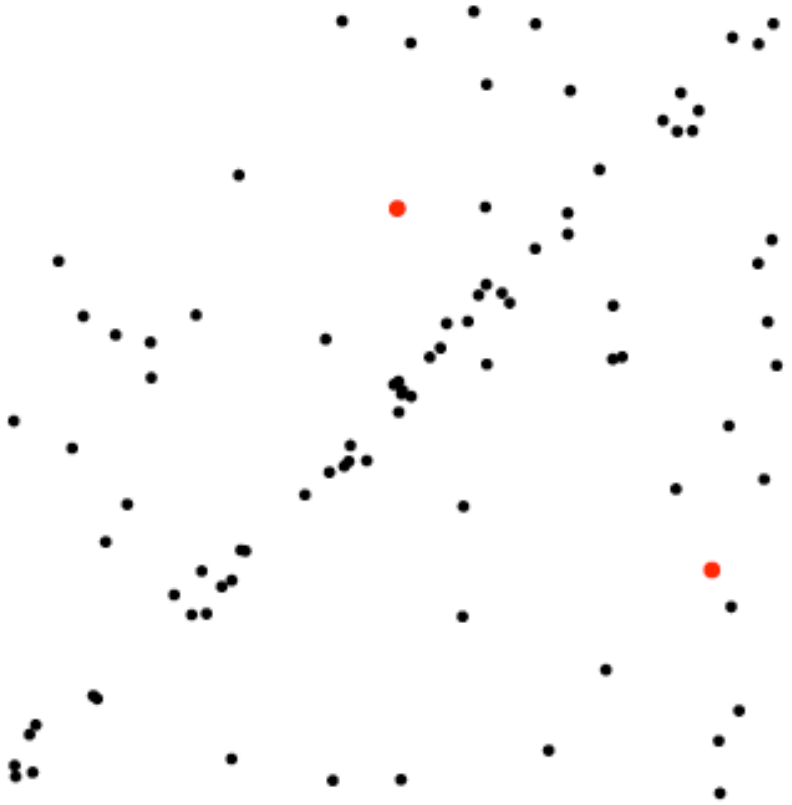

Slide credit: O. Chum

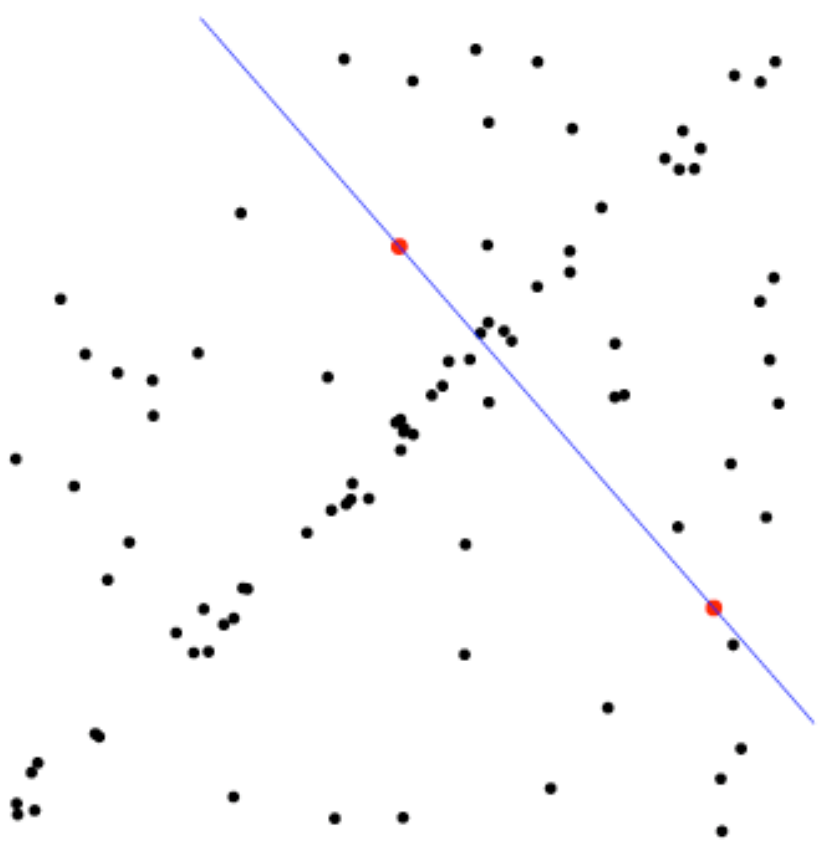

Slide credit: O. Chum

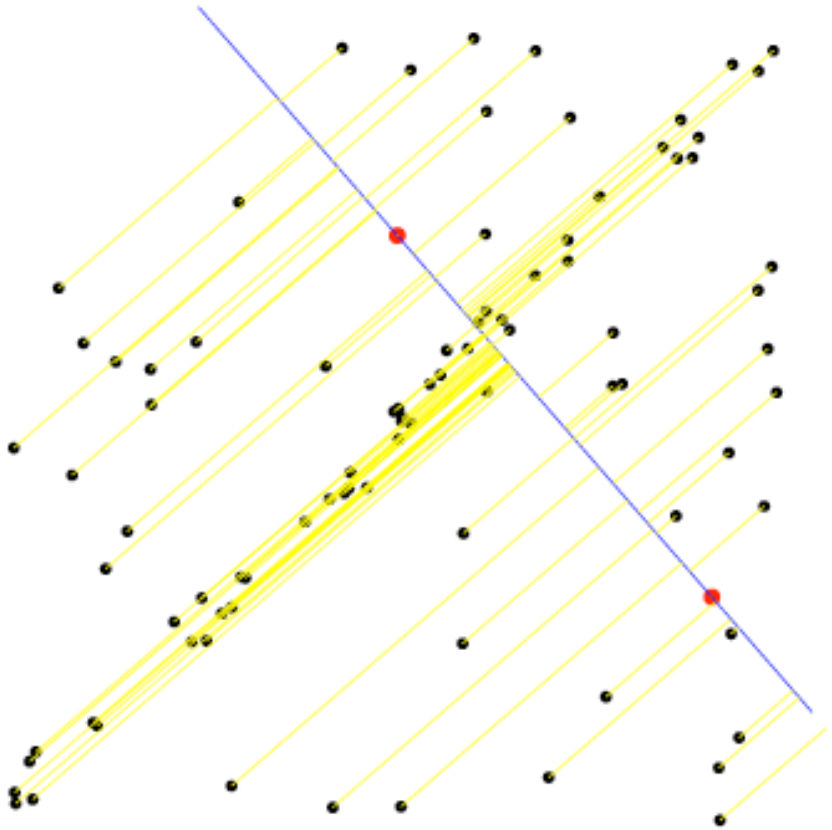

Slide credit: O. Chum

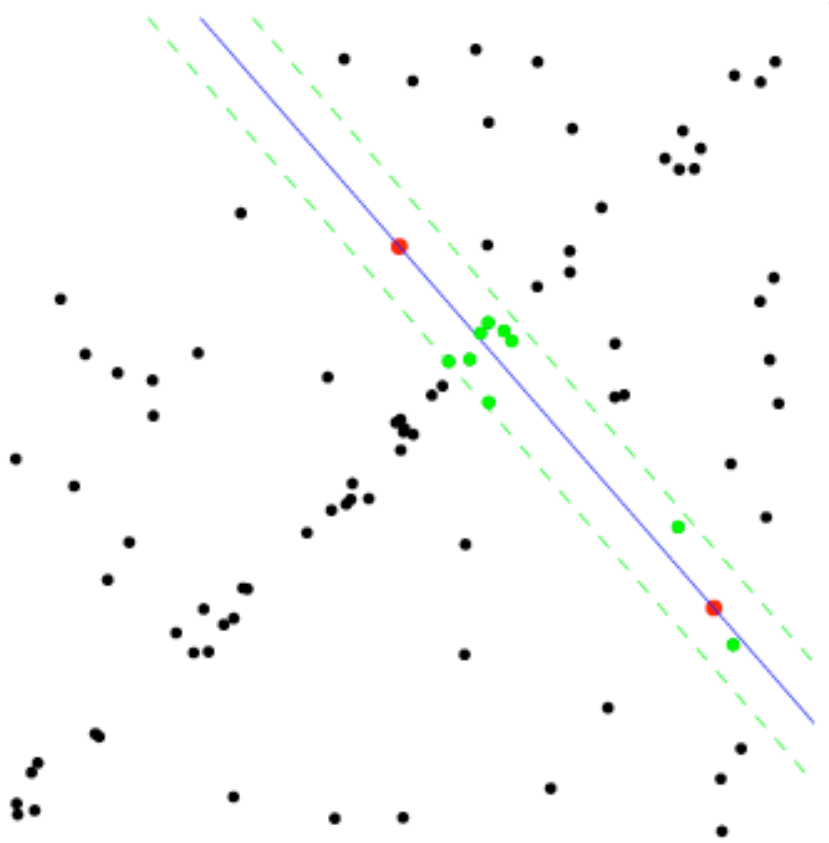

Slide credit: O. Chum

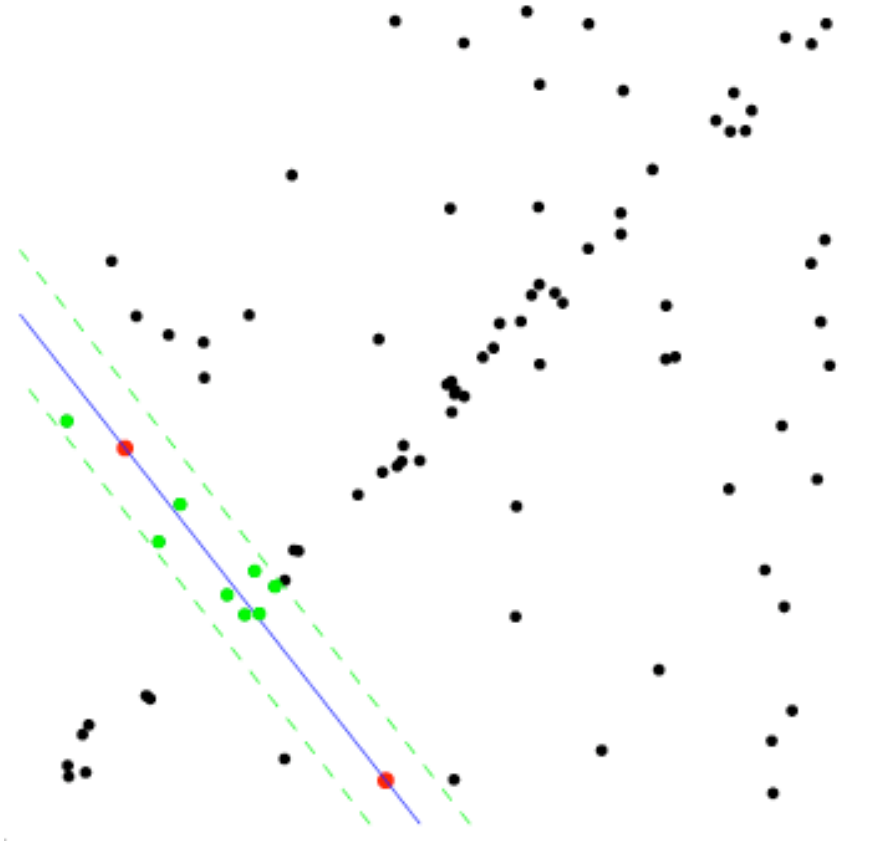

Slide credit: O. Chum

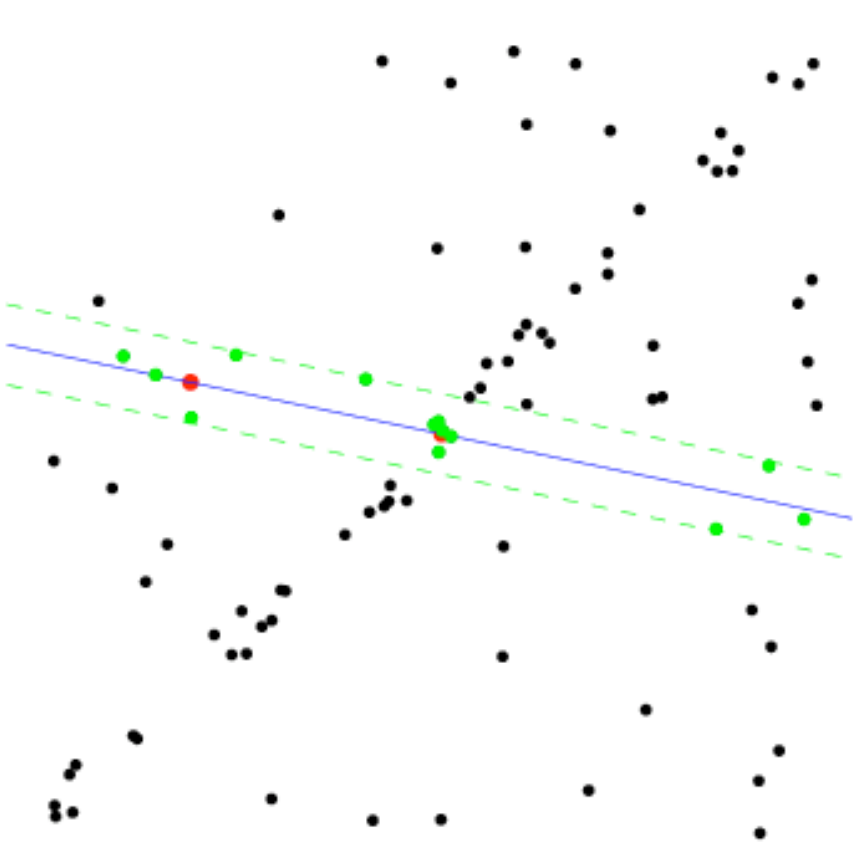

Slide credit: O. Chum

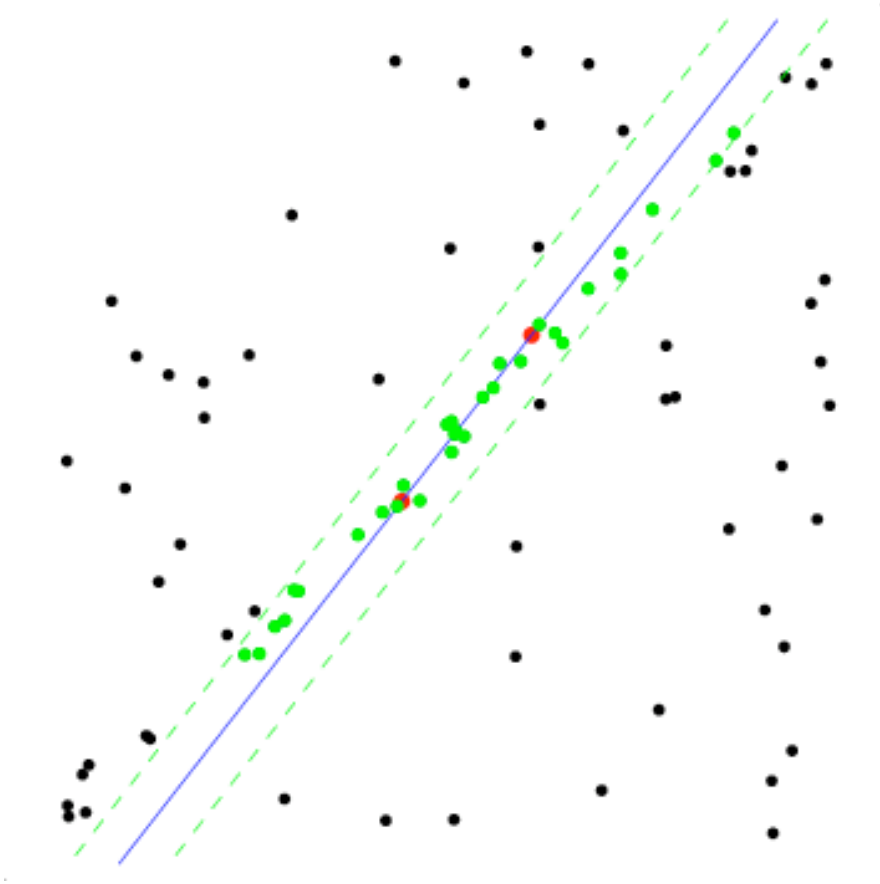

Slide credit: O. Chum

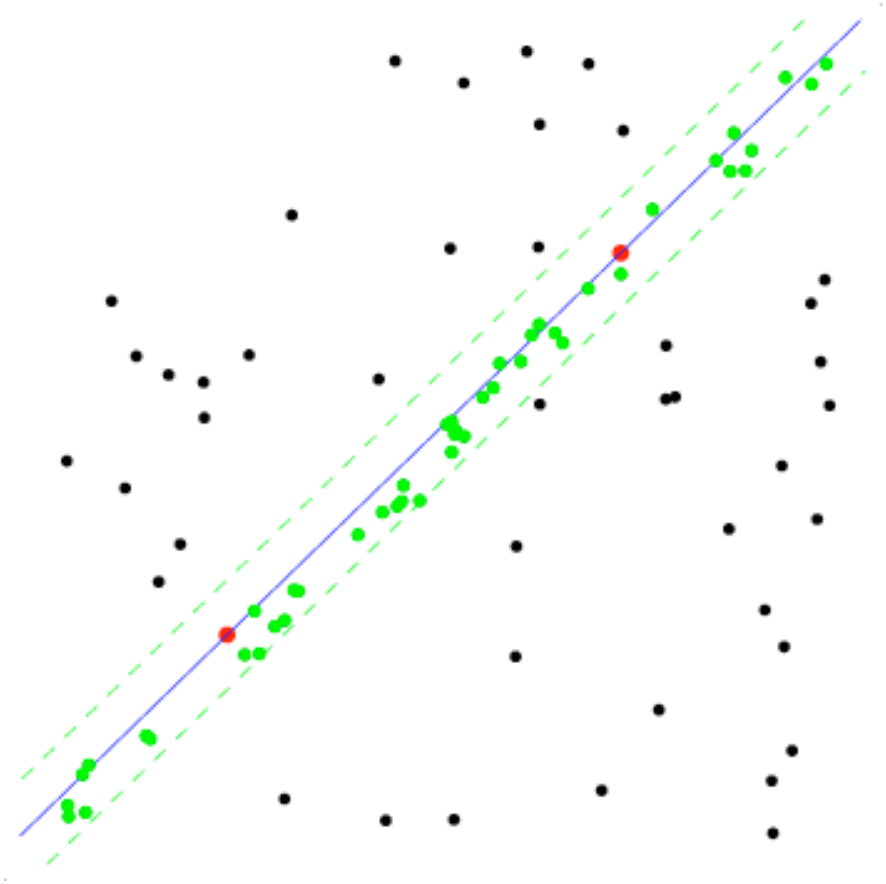

Slide credit: O. Chum

Algorithm summary – RANSAC robust estimation of 2D affine transformation

Repeat

1. Select 3 point to point correspondences
2. Compute H (2x2 matrix) + t (2x1) vector for translation
3. Measure support (number of inliers within threshold distance, i.e. $d^2_{\text{transfer}} < t$)

$$d^2_{\text{transfer}} = d(\mathbf{x}, \mathbf{H}^{-1}\mathbf{x}')^2 + d(\mathbf{x}', \mathbf{H}\mathbf{x})^2$$

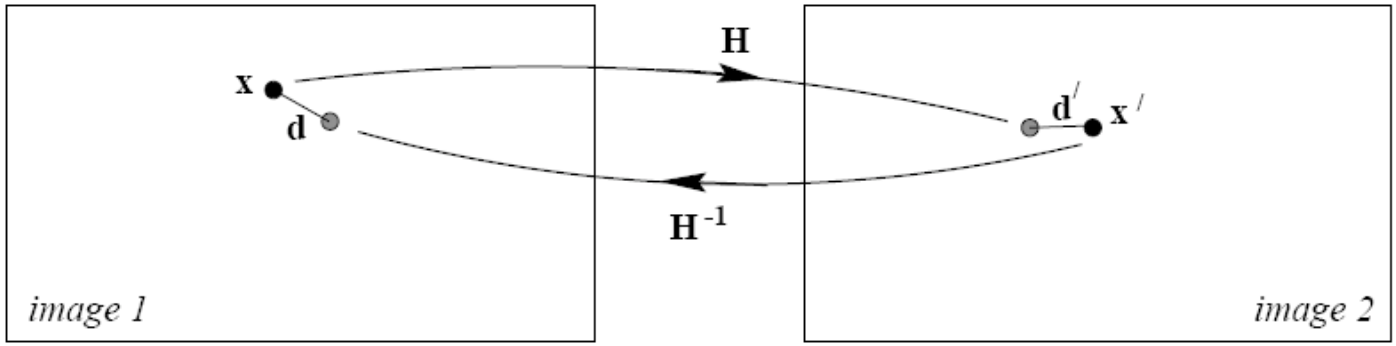

Choose the (H,t) with the largest number of inliers

(Re-estimate (H,t) from all inliers)

How many samples?

Number of samples *N*

- Choose *N* so that, with probability *p*, at least one random sample is free from outliers
- e.g.:
  - > *p*=0.99
  - > outlier ratio: *e*

Probability a randomly picked point is an inlier

$$\left(1-\underbrace{(1-e)^{s}}\right)^{N}=1-p$$

Probability of all points in a sample (of size s) are inliers

Source: M. Pollefeys

How many samples?

Number of samples *N*

- Choose *N* so that, with probability *p*, at least one random sample is free from outliers
- e.g.:
  - > *p*=0.99
  - > outlier ratio: *e*

Probability that all N samples (of size s) are corrupted (contain an outlier)

$$\left(1-\left(1-e\right)^{s}\right)^{N}=1-p$$

Probability of at least one point in a sample (of size s) is an outlier

$$N=\log\left(1-p\right)/\log\left(1-\left(1-e\right)^{s}\right)$$

	proportion of outliers *e*						
s	5%	10%	20%	30%	40%	50%	90%
1	2	2	3	4	5	6	43
2	2	3	5	7	11	17	458
3	3	4	7	11	19	35	4603
4	3	5	9	17	34	72	4.6e4
5	4	6	12	26	57	146	4.6e5
6	4	7	16	37	97	293	4.6e6
7	4	8	20	54	163	588	4.6e7
8	5	9	26	78	272	1177	4.6e8

Source: M. Pollefeys

Example: line fitting

p = 0.99

s = ?

e = ?

N = ?

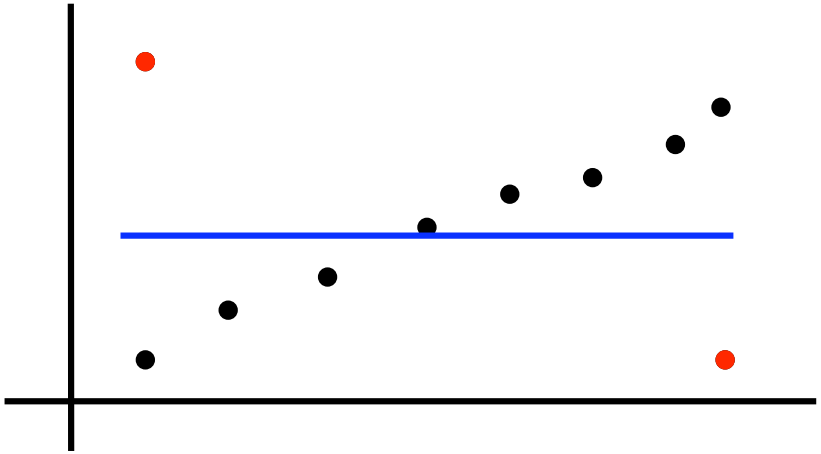

Source: M. Pollefeys

Example: line fitting

p = 0.99

s = 2

e = 2/10 = 0.2

N = 5

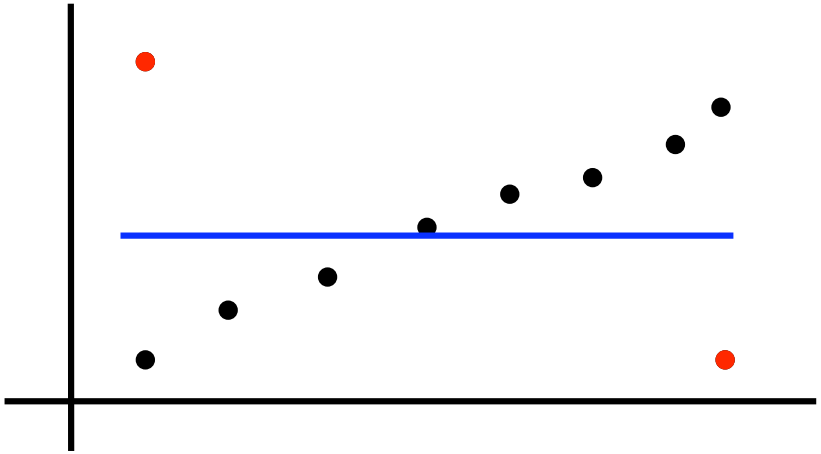

Compare with exhaustively trying all point pairs:

$$\binom{10}{2} = 10*9 / 2 = 45$$

	proportion of outliers e						
s	5%	10%	20%	30%	40%	50%	90%
1	2	2	3	4	5	6	43
2	2	3	5	7	11	17	458
3	3	4	7	11	19	35	4603
4	3	5	9	17	34	72	4.6e4
5	4	6	12	26	57	146	4.6e5
6	4	7	16	37	97	293	4.6e6
7	4	8	20	54	163	588	4.6e7
8	5	9	26	78	272	1177	4.6e8

Source: M. Pollefeys

How to reduce the number of samples needed?

1. Reduce the proportion of outliers.
2. Reduce the sample size
   - use simpler model (e.g. similarity instead of affine tnf.)
   - use local information (e.g. a region to region correspondence is equivalent to (up to) 3 point to point correspondences).

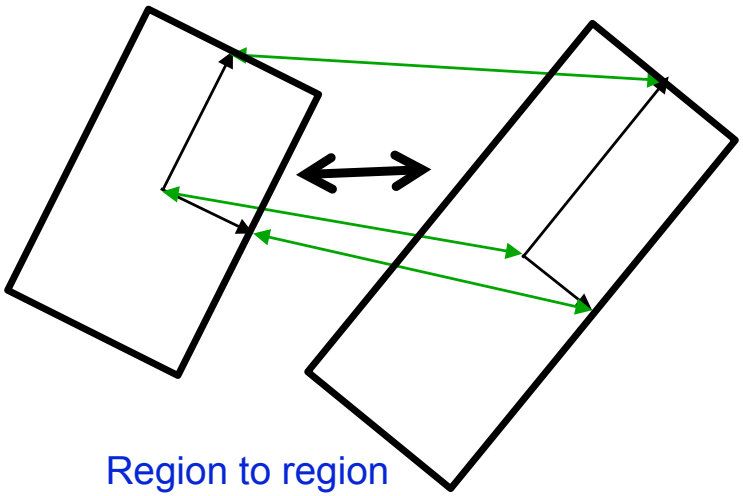

Region to region correspondence

Number of samples *N*

	proportion of outliers *e*						
s	5%	10%	20%	30%	40%	50%	90%
1	2	2	3	4	5	6	43
2	2	3	5	7	11	17	458
3	3	4	7	11	19	35	4603
4	3	5	9	17	34	72	4.6e4
5	4	6	12	26	57	146	4.6e5
6	4	7	16	37	97	293	4.6e6
7	4	8	20	54	163	588	4.6e7
8	5	9	26	78	272	1177	4.6e8

RANSAC (references)

M. Fischler and R. Bolles, “Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography,” Comm. ACM, 1981

R. Hartley and A. Zisserman, Multiple View Geometry in Computer Vision, 2nd ed., 2004.

Extensions:

B. Tordoff and D. Murray, “Guided Sampling and Consensus for Motion Estimation, ECCV’03

D. Nister, “Preemptive RANSAC for Live Structure and Motion Estimation, ICCV’03

Chum, O.; Matas, J. and Obdrzalek, S.: Enhancing RANSAC by Generalized Model Optimization, ACCV’04

Chum, O.; and Matas, J.: Matching with PROSAC - Progressive Sample Consensus , CVPR 2005

Philbin, J., Chum, O., Isard, M., Sivic, J. and Zisserman, A.: Object retrieval with large vocabularies and fast spatial matching, CVPR’07

Chum, O. and Matas. J.: Optimal Randomized RANSAC, PAMI’08

Strategy 2: Hough Transform

- Origin: Detection of straight lines in cluttered images
- Can be generalized to arbitrary shapes
- Can extract feature groupings from cluttered images in linear time.
- Illustrate on extracting sets of local features consistent with a similarity transformation.

Hough transform for object recognition

Suppose our features are scale- and rotation-covariant

- Then a single feature match provides an alignment hypothesis (translation, scale, orientation)

model

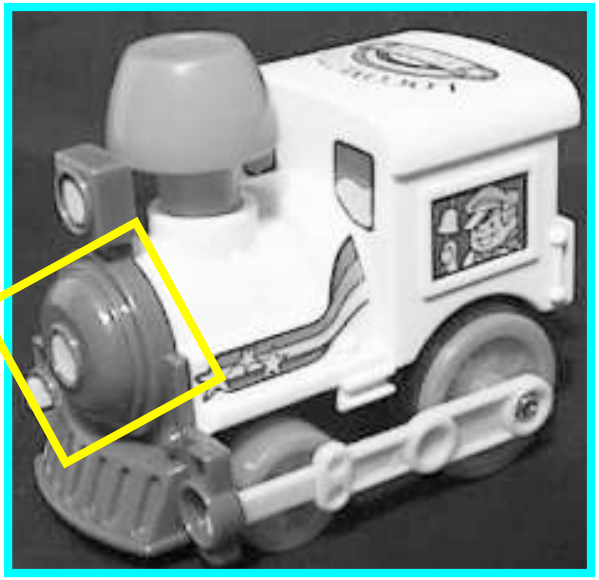

Target image

David G. Lowe. **"Distinctive image features from scale-invariant keypoints",** *IJCV* 60 (2), pp. 91-110, 2004.

Hough transform for object recognition

Suppose our features are scale- and rotation-covariant

- Then a single feature match provides an alignment hypothesis (translation, scale, orientation)
- Of course, a hypothesis obtained from a single match is unreliable
- Solution: Coarsely quantize the transformation space. Let each match vote for its hypothesis in the quantized space.

model

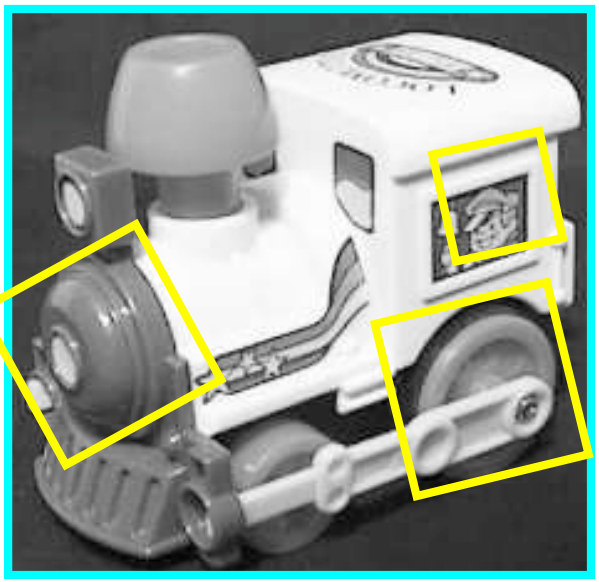

David G. Lowe. **"Distinctive image features from scale-invariant keypoints",** *IJCV* 60 (2), pp. 91-110, 2004.

Basic algorithm outline

1. Initialize accumulator H to all zeros
2. For each tentative match
   compute transformation hypothesis: tx, ty, s, θ
   H(tx,ty,s,θ) = H(tx,ty,s,θ) + 1
   end
   end
3. Find all bins (tx,ty,s,θ) where H(tx,ty,s,θ) has at least three votes

H: 4D-accumulator array
(only 2-d shown here)

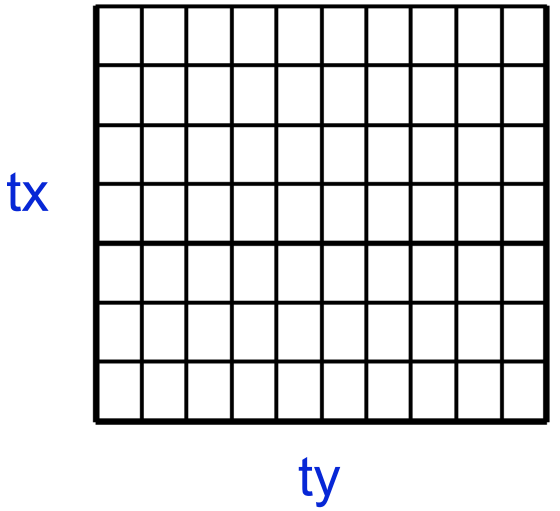

- Correct matches will consistently vote for the same transformation while mismatches will spread votes.
- Cost: Linear scan through the matches (step 2), followed by a linear scan through the accumulator (step 3).

Hough transform details (D. Lowe's system)

Training phase: For each model feature, record 2D location, scale, and orientation of model (relative to normalized feature frame)

Test phase: Let each match between a test and a model feature vote in a 4D Hough space

- Use broad bin sizes of 30 degrees for orientation, a factor of 2 for scale, and 0.25 times image size for location
- Vote for two closest bins in each dimension

Find all bins with at least three votes and perform geometric verification

- Estimate least squares *affine* transformation
- Use stricter thresholds on transformation residual
- Search for additional features that agree with the alignment

Hough transform in object recognition (references)

P.V.C. Hough, Machine Analysis of Bubble Chamber Pictures, Proc. Int. Conf. High Energy Accelerators and Instrumentation, 1959

D. Lowe, “Distinctive image features from scale-invariant keypoints”, *IJCV* 60 (2), 2004.

H. Jegou, M. Douze, C. Schmid, Hamming embedding and weak geometric consistency for large scale image search, ECCV’2008

Extensions (object category detection):

B. Leibe, A. Leonardis, and B. Schiele., Combined Object Categorization and Segmentation with an Implicit Shape Model, in ECCV'04 Workshop on Statistical Learning in Computer Vision, Prague, May 2004.

S. Maji and J. Malik, Object Detection Using a Max-Margin Hough Tranform, CVPR’2009

A. Lehmann, B. Leibe, L. Van Gool. Fast PRISM: Branch and Bound Hough Transform for Object Class Detection, IJCV (to appear), 2010.

O. Barinova, V. Lempitsky, P. Kohli, On the Detection of Multiple Object Instances using Hough Transforms, CVPR, 2010

Comparison

Hough Transform

Advantages

- Can handle high percentage of outliers (>95%)
- Extracts groupings from clutter in linear time

Disadvantages

- Quantization issues
- Only practical for small number of dimensions (up to 4)

Improvements available

- Probabilistic Extensions
- Continuous Voting Space
- Can be generalized to arbitrary shapes and objects

RANSAC

Advantages

- General method suited to large range of problems
- Easy to implement
- “Independent” of number of dimensions

Disadvantages

- Basic version only handles moderate number of outliers (<50%)

Many variants available, e.g.

- PROSAC: Progressive RANSAC [Chum05]
- Preemptive RANSAC [Nister05]

Slide credit: K. Grauman, B. Leibe

Beyond affine transformations

What is the transformation between two views of a planar surface?

What is the transformation between images from two cameras that share the same center?

Beyond affine transformations

Homography: plane projective transformation
(transformation taking a quad to another arbitrary quad)

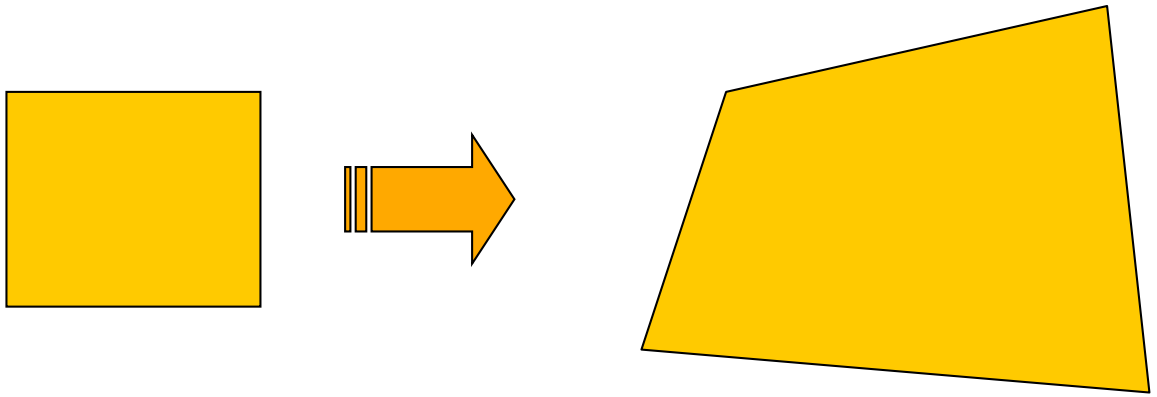

Case II: Cameras rotating about their centre

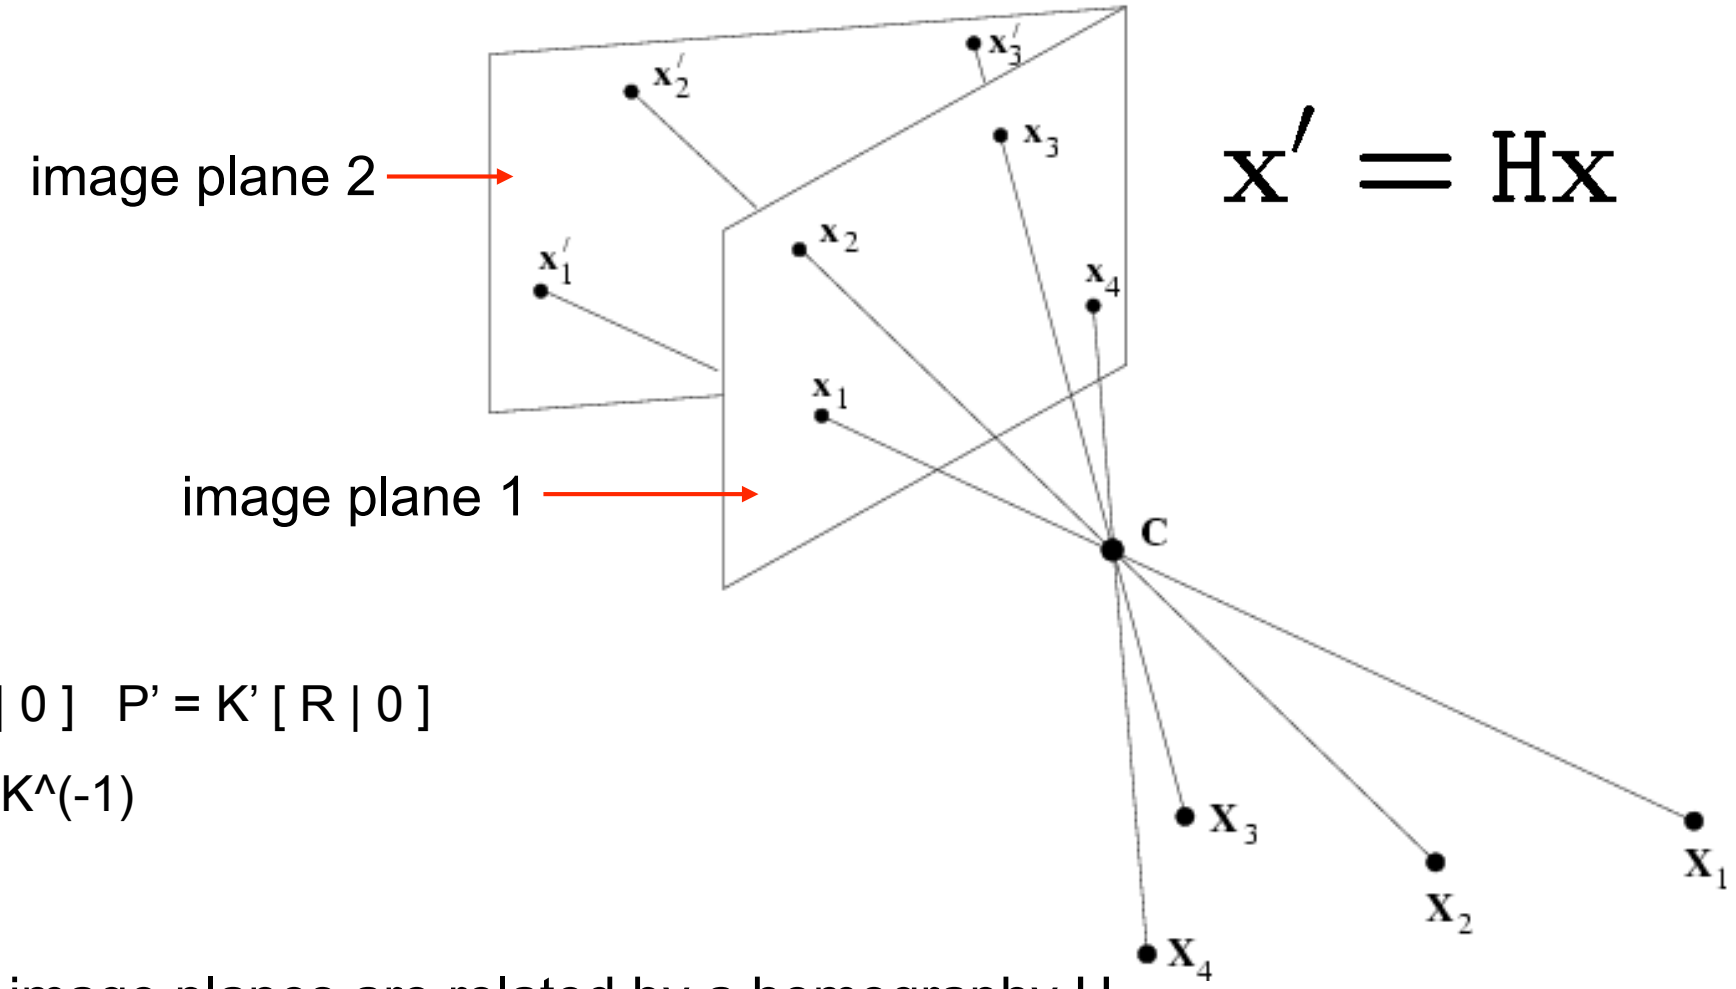

$$\mathbf{x}' = \mathrm{H}\mathbf{x}$$

P = K [I | 0] P' = K' [R | 0]

H = K' R K^(-1)

- The two image planes are related by a homography H
- H depends only on the relation between the image planes and camera centre, C, not on the 3D structure

Fitting a homography

Recall: homogenenous coordinates

$$(x, y) \Rightarrow \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Converting *to* homogenenous image coordinates

$$\begin{bmatrix} x \\ y \\ w \end{bmatrix} \Rightarrow (x/w, y/w)$$

Converting *from* homogenenous image coordinates

Fitting a homography

Recall: homogenenous coordinates

$$(x, y) \Rightarrow \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Converting *to* homogenenous image coordinates

$$\begin{bmatrix} x \\ y \\ w \end{bmatrix} \Rightarrow (x/w, y/w)$$

Converting *from* homogenenous image coordinates

Equation for homography:

$$\lambda \begin{bmatrix} x' \\ y' \\ 1 \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} x \\ y \\ 1 \end{bmatrix}$$

Fitting a homography

Equation for homography:

$$\lambda_i \begin{bmatrix} x_i' \\ y_i' \\ 1 \end{bmatrix} = \begin{bmatrix} h_{11} & h_{12} & h_{13} \\ h_{21} & h_{22} & h_{23} \\ h_{31} & h_{32} & h_{33} \end{bmatrix} \begin{bmatrix} x_i \\ y_i \\ 1 \end{bmatrix} \qquad \lambda_i \mathbf{x}_i' = \mathbf{H}\mathbf{x}_i = \begin{bmatrix} \mathbf{h}_1^T \\ \mathbf{h}_2^T \\ \mathbf{h}_3^T \end{bmatrix} \mathbf{x}_i$$

9 entries, 8 degrees of freedom
(scale is arbitrary)

$$\mathbf{x}_i' \times \mathbf{H}\mathbf{x}_i = 0 \qquad \mathbf{x}_i' \times \mathbf{H}\mathbf{x}_i = \begin{bmatrix} y_i' \mathbf{h}_3^T \mathbf{x}_i - \mathbf{h}_2^T \mathbf{x}_i \\ \mathbf{h}_1^T \mathbf{x}_i - x_i' \mathbf{h}_3^T \mathbf{x}_i \\ x_i' \mathbf{h}_2^T \mathbf{x}_i - y_i' \mathbf{h}_1^T \mathbf{x}_i \end{bmatrix}$$

$$\begin{bmatrix} 0^T & -\mathbf{x}_i^T & y_i' \mathbf{x}_i^T \\ \mathbf{x}_i^T & 0^T & -x_i' \mathbf{x}_i^T \\ -y_i' \mathbf{x}_i^T & x_i' \mathbf{x}_i^T & 0^T \end{bmatrix} \begin{pmatrix} \mathbf{h}_1 \\ \mathbf{h}_2 \\ \mathbf{h}_3 \end{pmatrix} = 0$$

3 equations, only 2 linearly independent

Direct linear transform

$$\begin{bmatrix} 0^T & \mathbf{x}_1^T & -y_1'\,\mathbf{x}_1^T \\ \mathbf{x}_1^T & 0^T & -x_1'\,\mathbf{x}_1^T \\ \cdots & \cdots & \cdots \\ 0^T & \mathbf{x}_n^T & -y_n'\,\mathbf{x}_n^T \\ \mathbf{x}_n^T & 0^T & -x_n'\,\mathbf{x}_n^T \end{bmatrix} \begin{pmatrix} \mathbf{h}_1 \\ \mathbf{h}_2 \\ \mathbf{h}_3 \end{pmatrix} = 0 \qquad \mathbf{A}\,\mathbf{h} = 0$$

- H has 8 degrees of freedom (9 parameters, but scale is arbitrary)
- One match gives us two linearly independent equations
- Four matches needed for a minimal solution (null space of 8x9 matrix)
- More than four: homogeneous least squares

Application: Panorama stitching

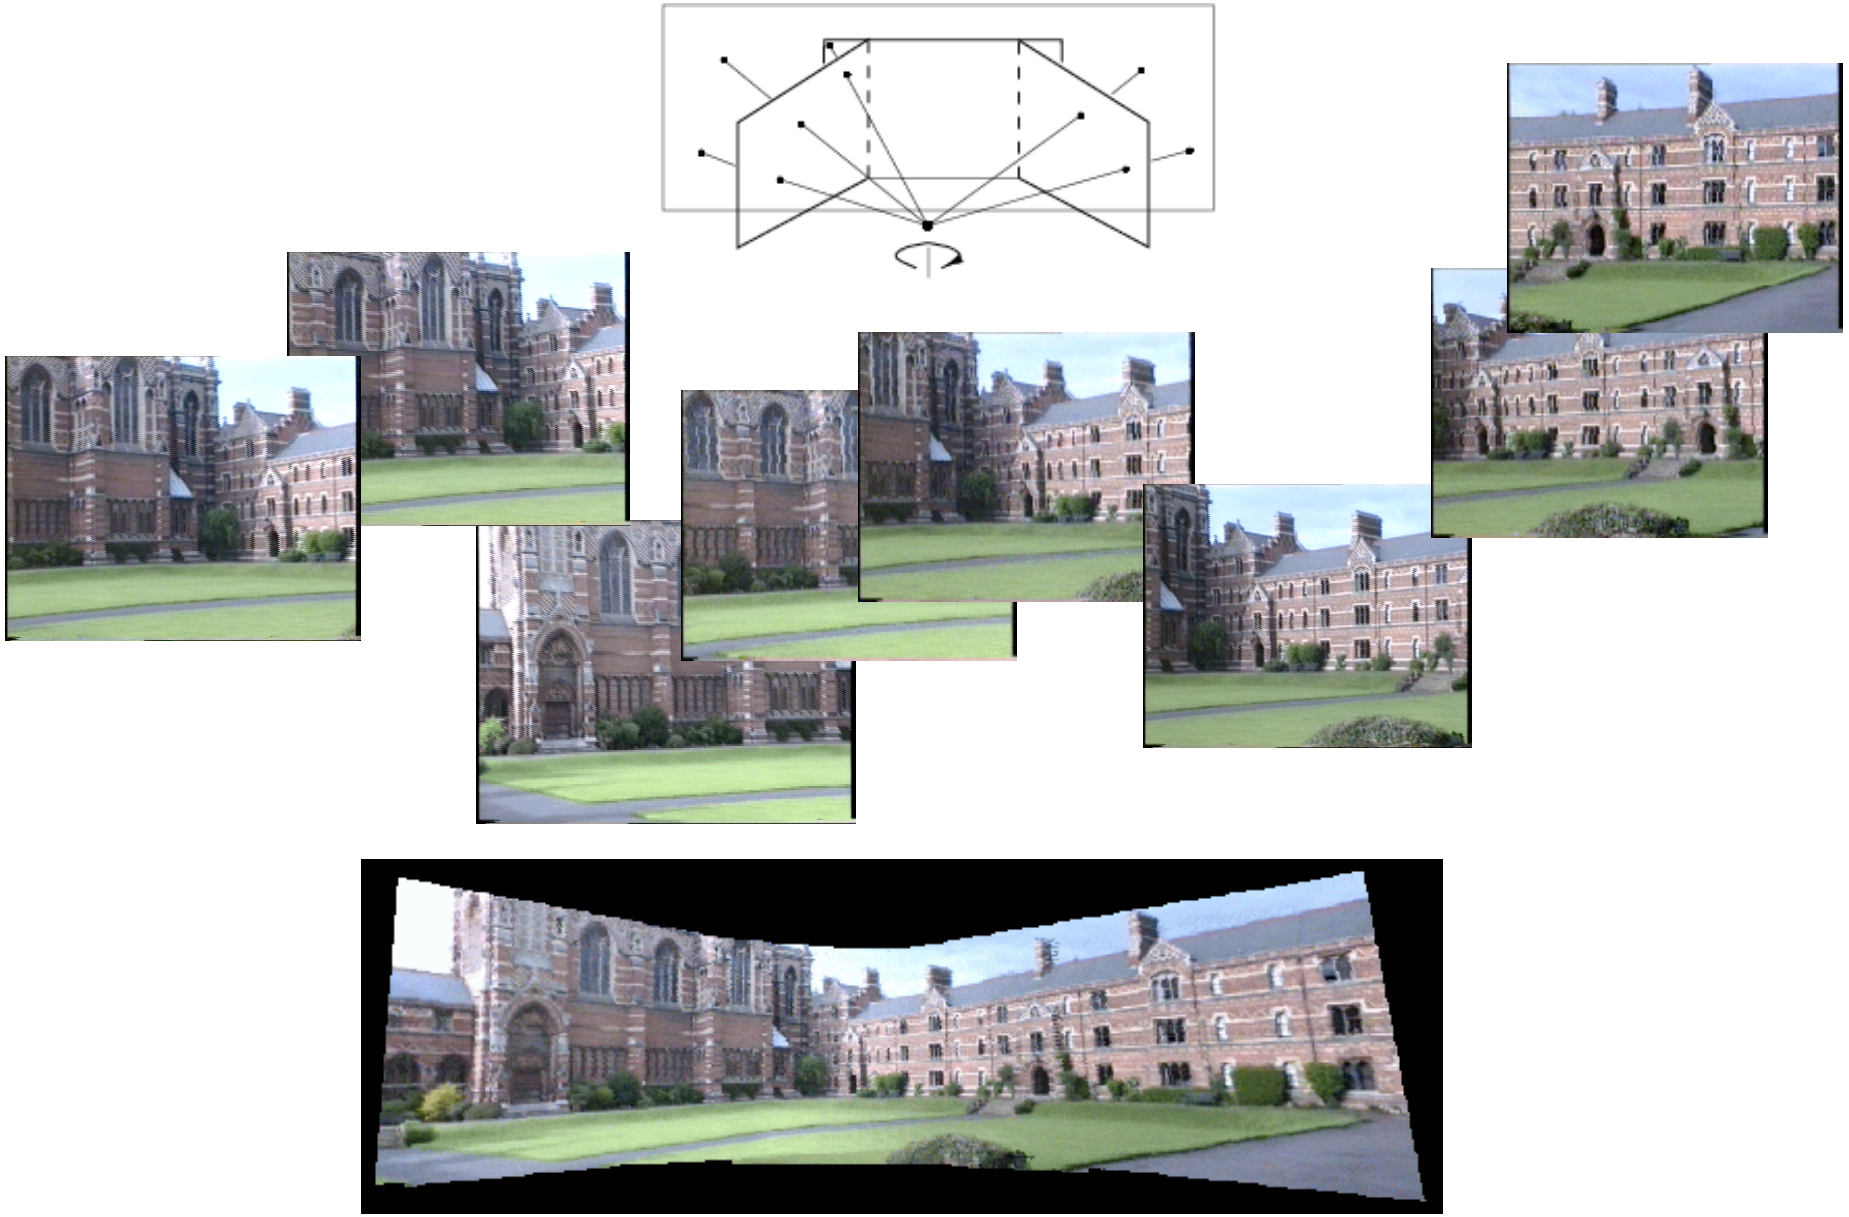

Images courtesy of A. Zisserman.

Recognizing panoramas

Given contents of a camera memory card, automatically figure out which pictures go together and stitch them together into panoramas

M. Brown and D. Lowe, "Recognizing panoramas", ICCV 2003.

1. Estimate homography (RANSAC)

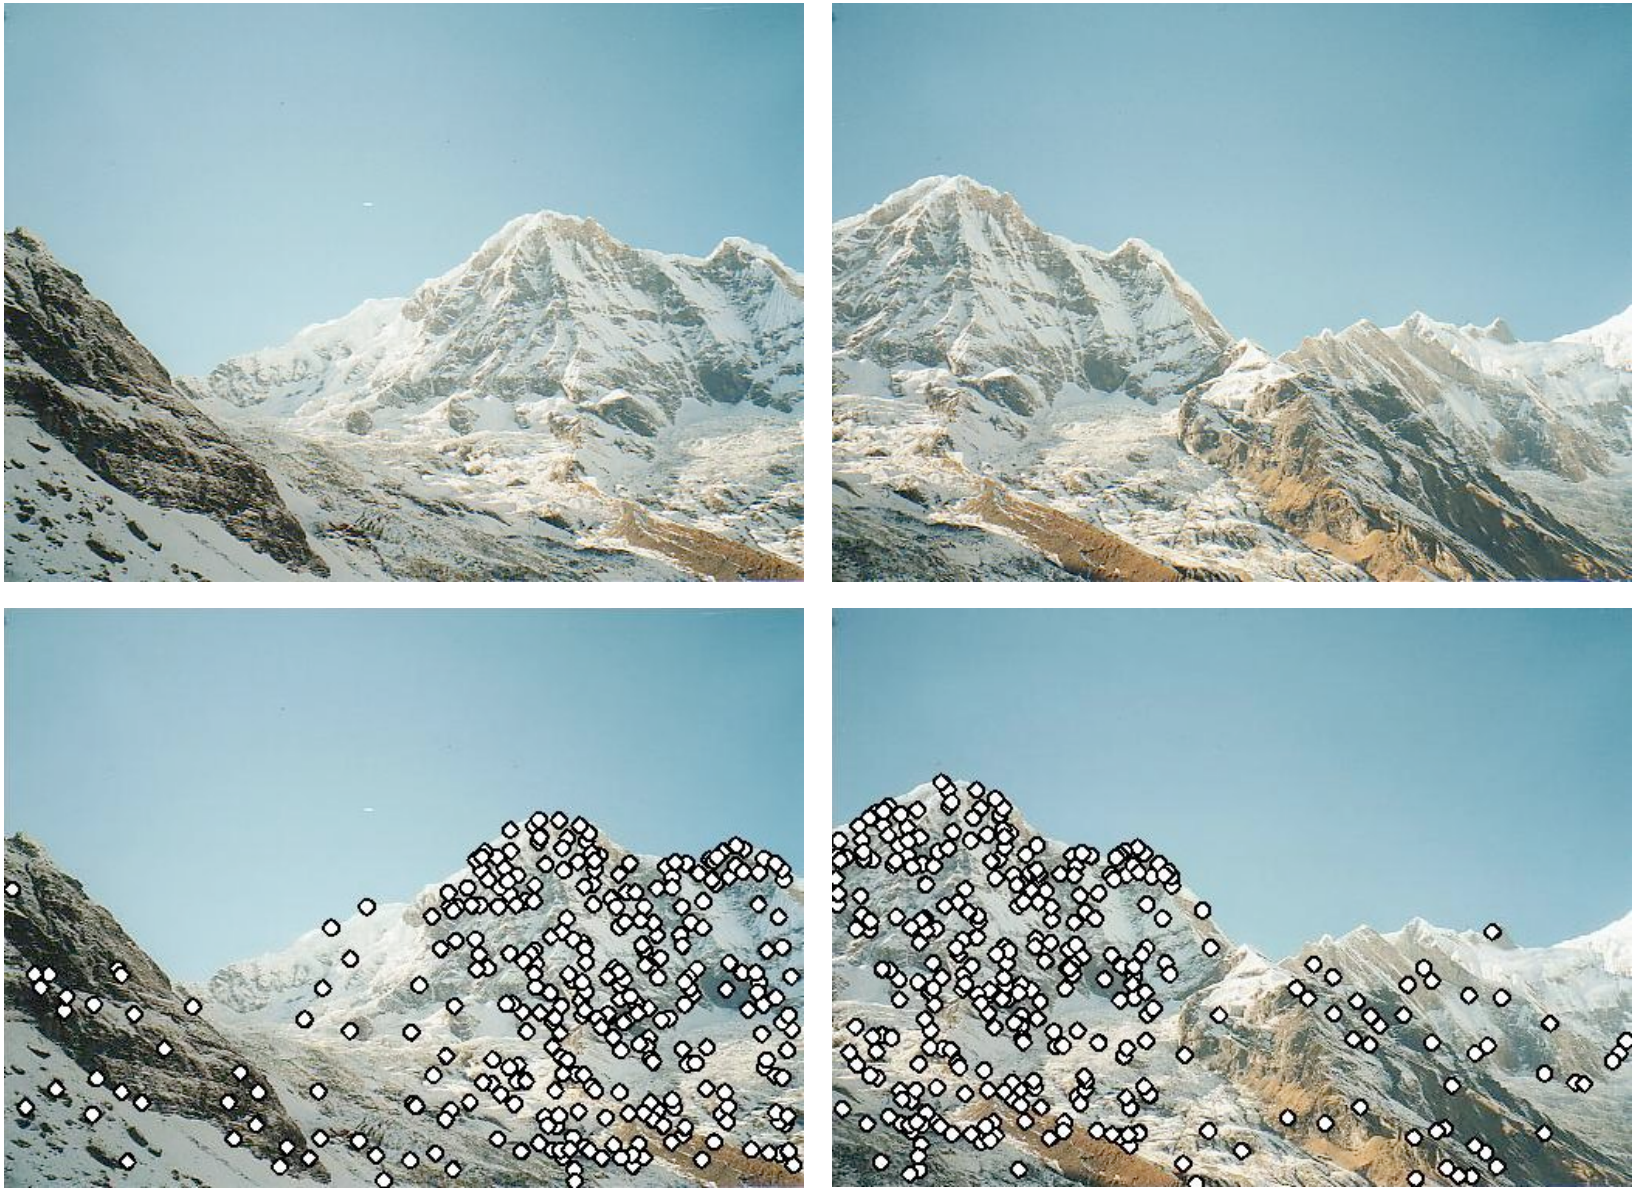

1. Estimate homography (RANSAC)

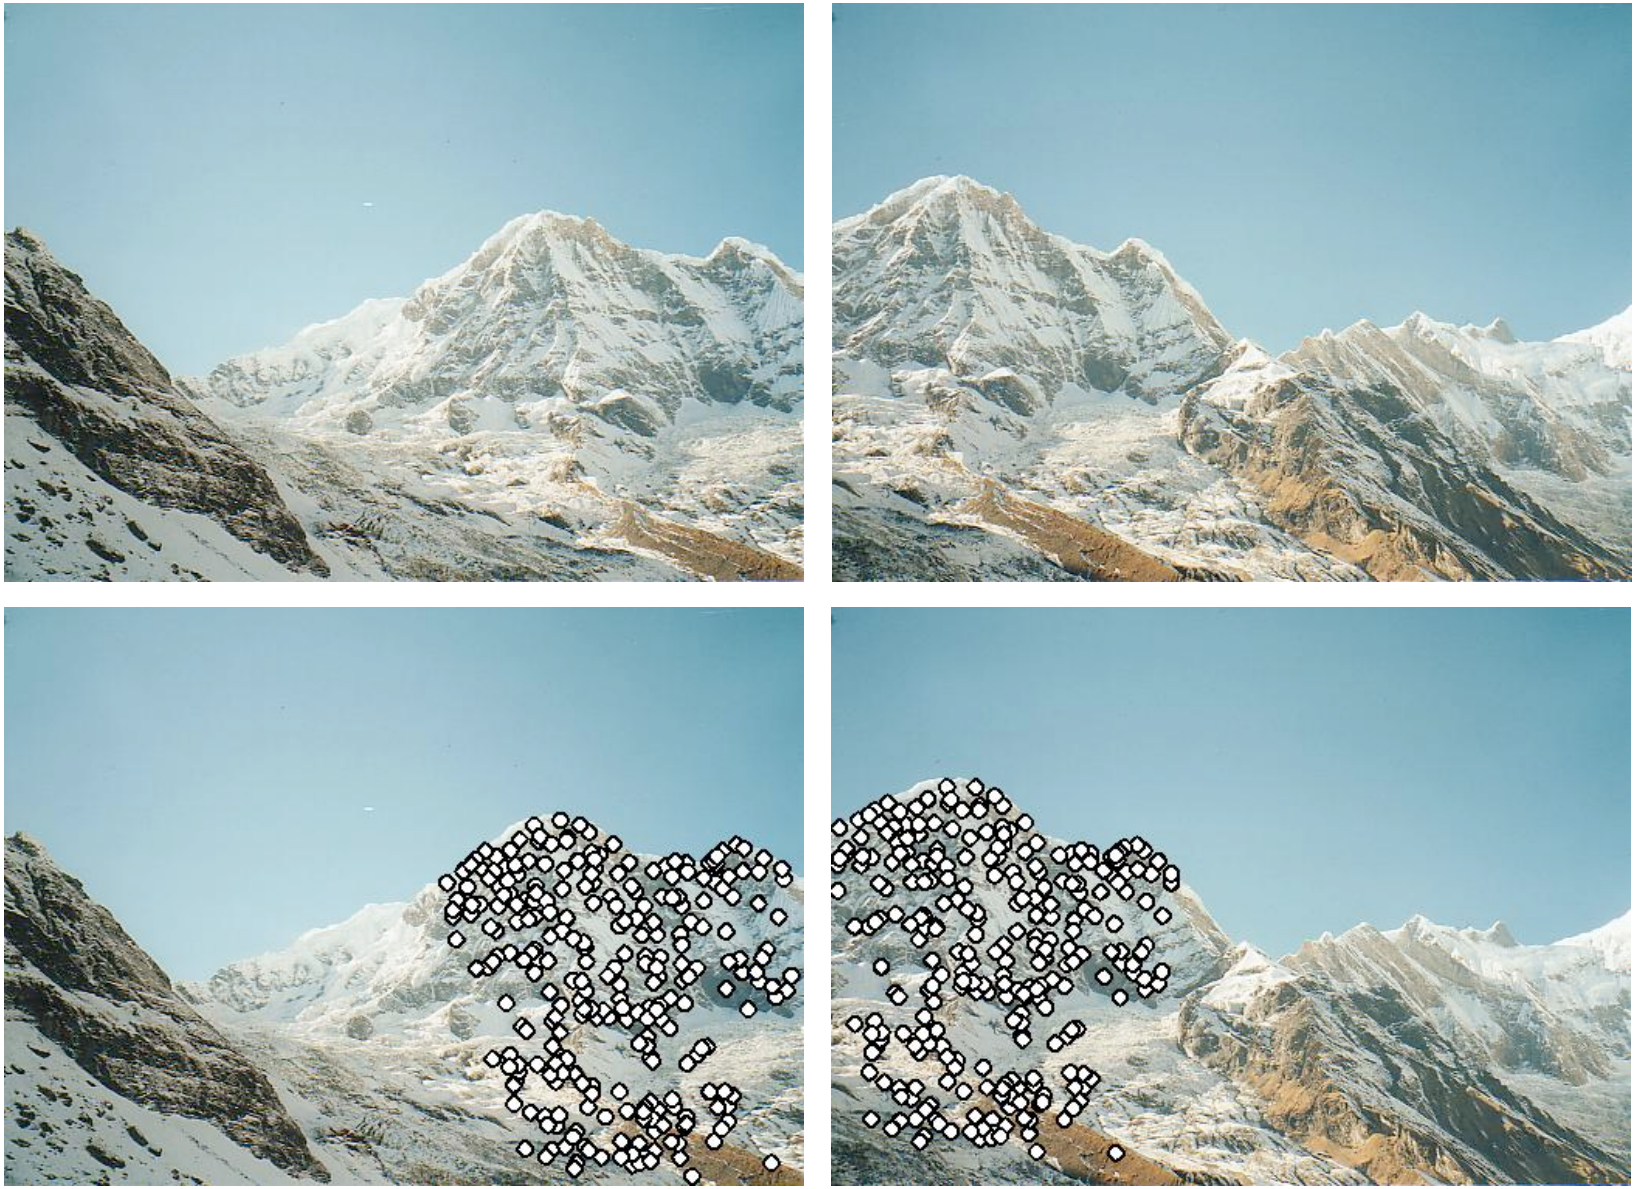

1. Estimate homography (RANSAC)

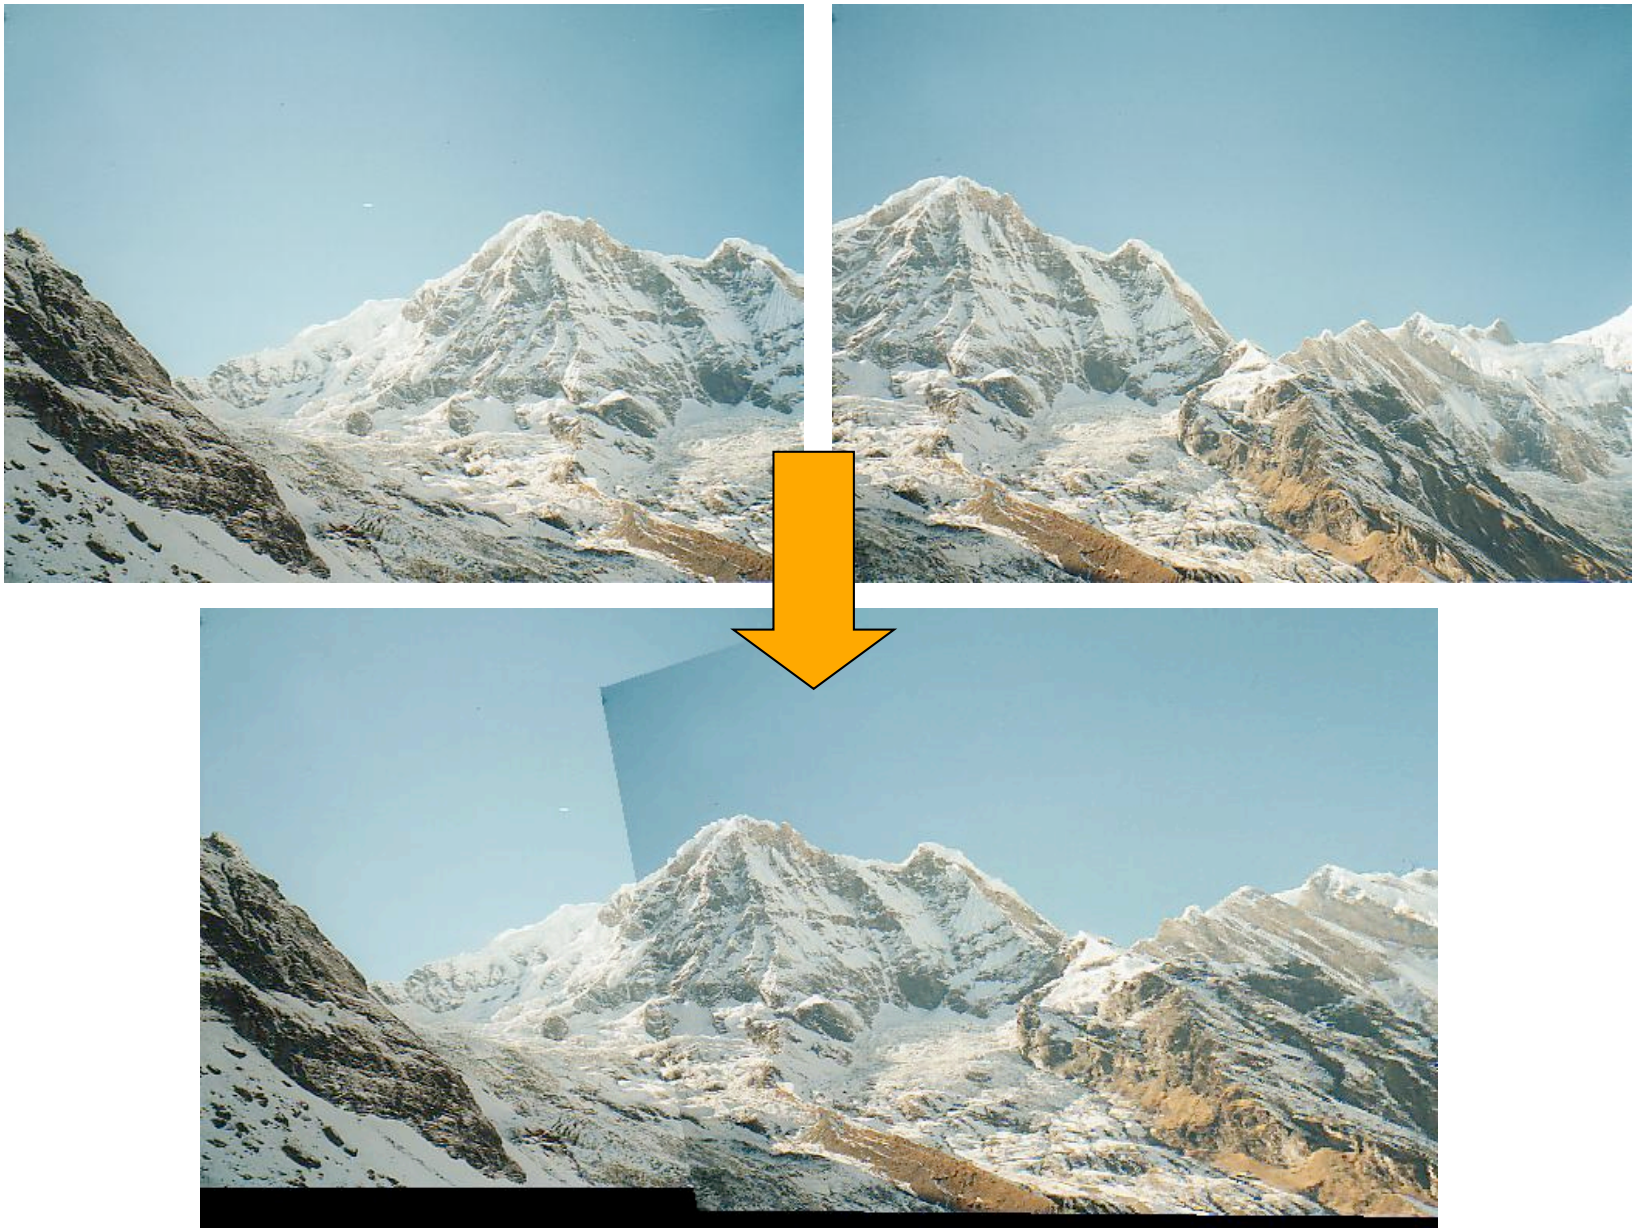

2. Find connected sets of images

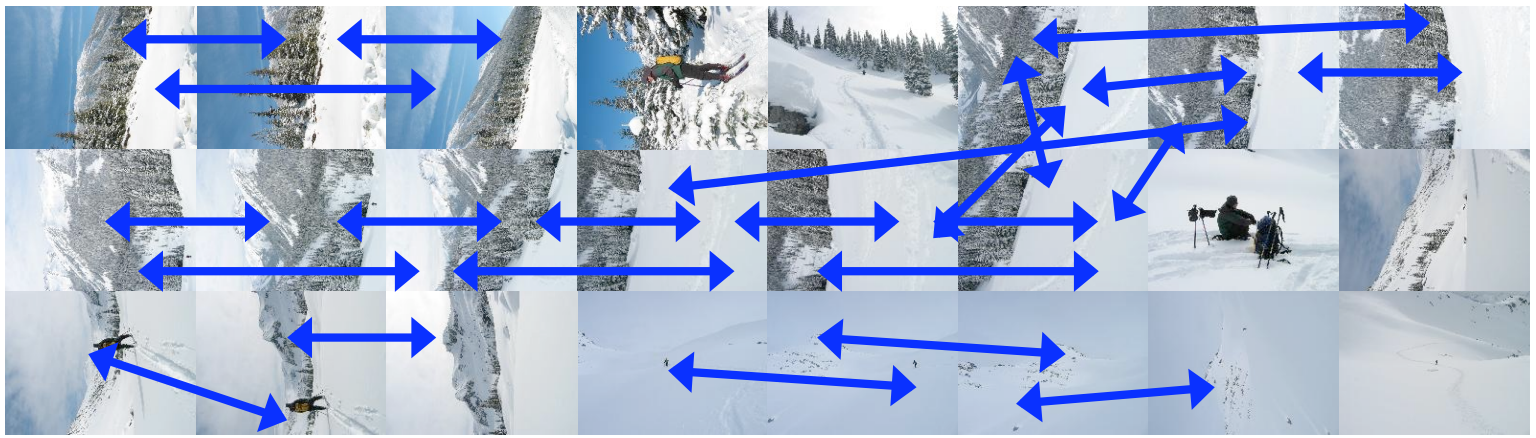

2. Find connected sets of images

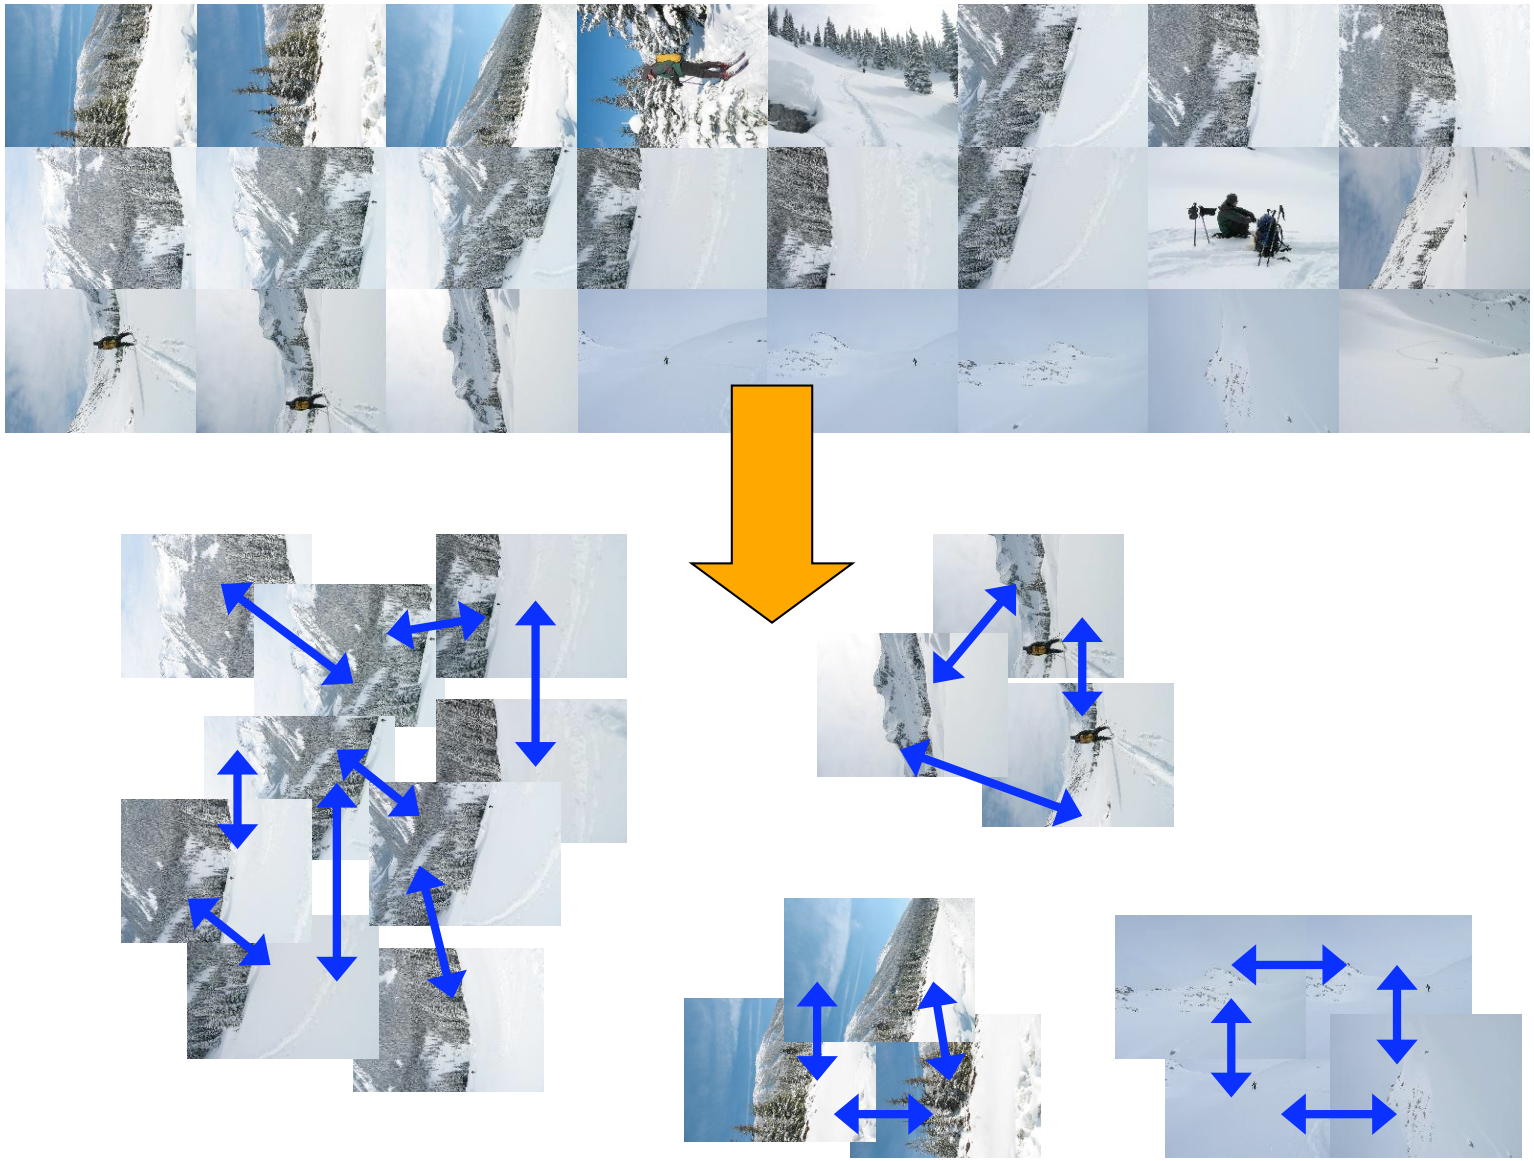

2. Find connected sets of images

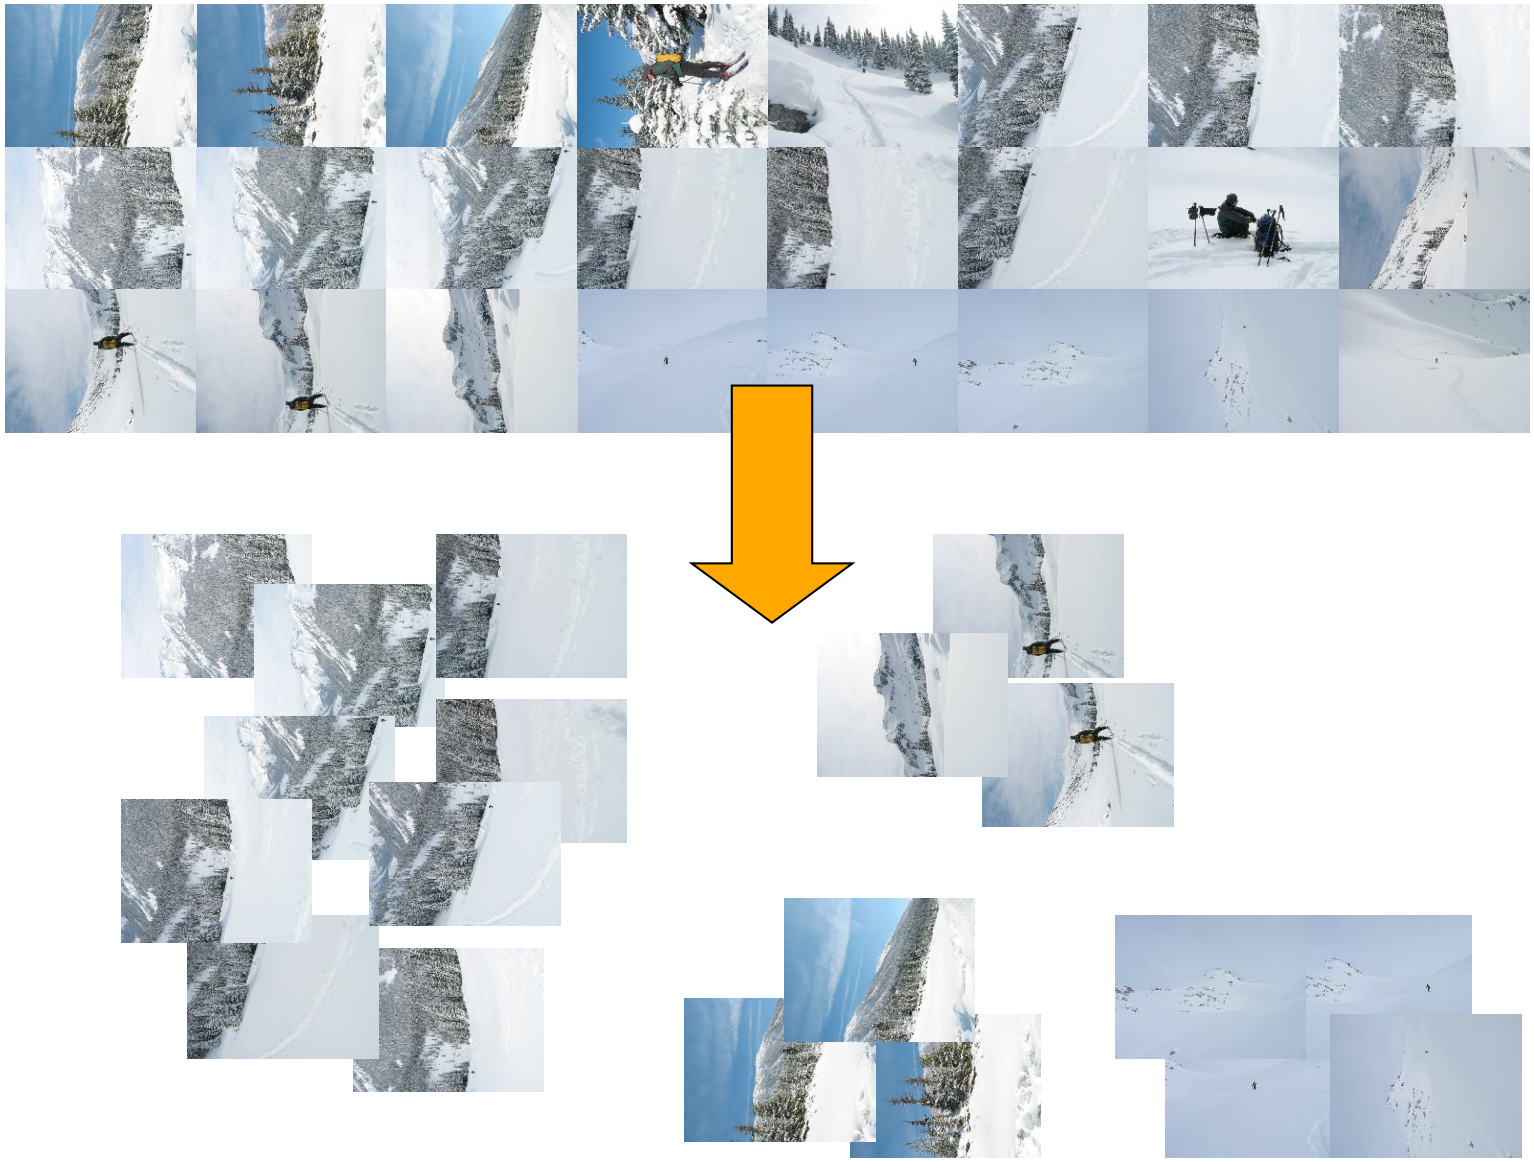

3. Stitch and blend the panoramas

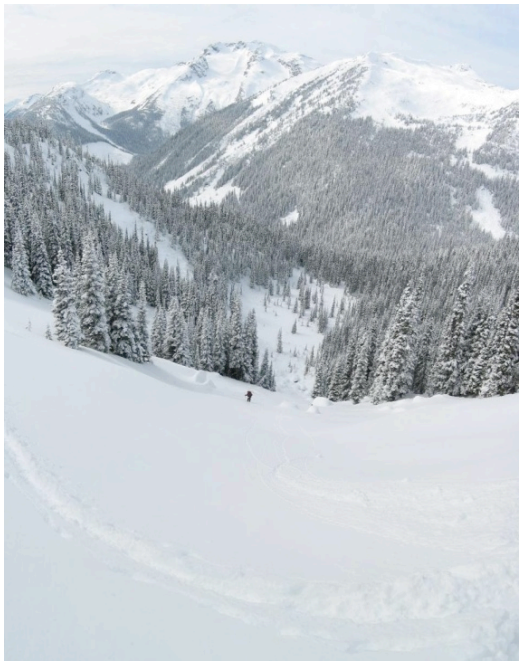

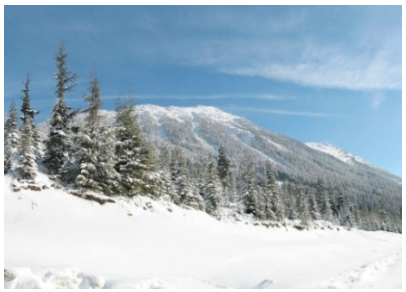

Results

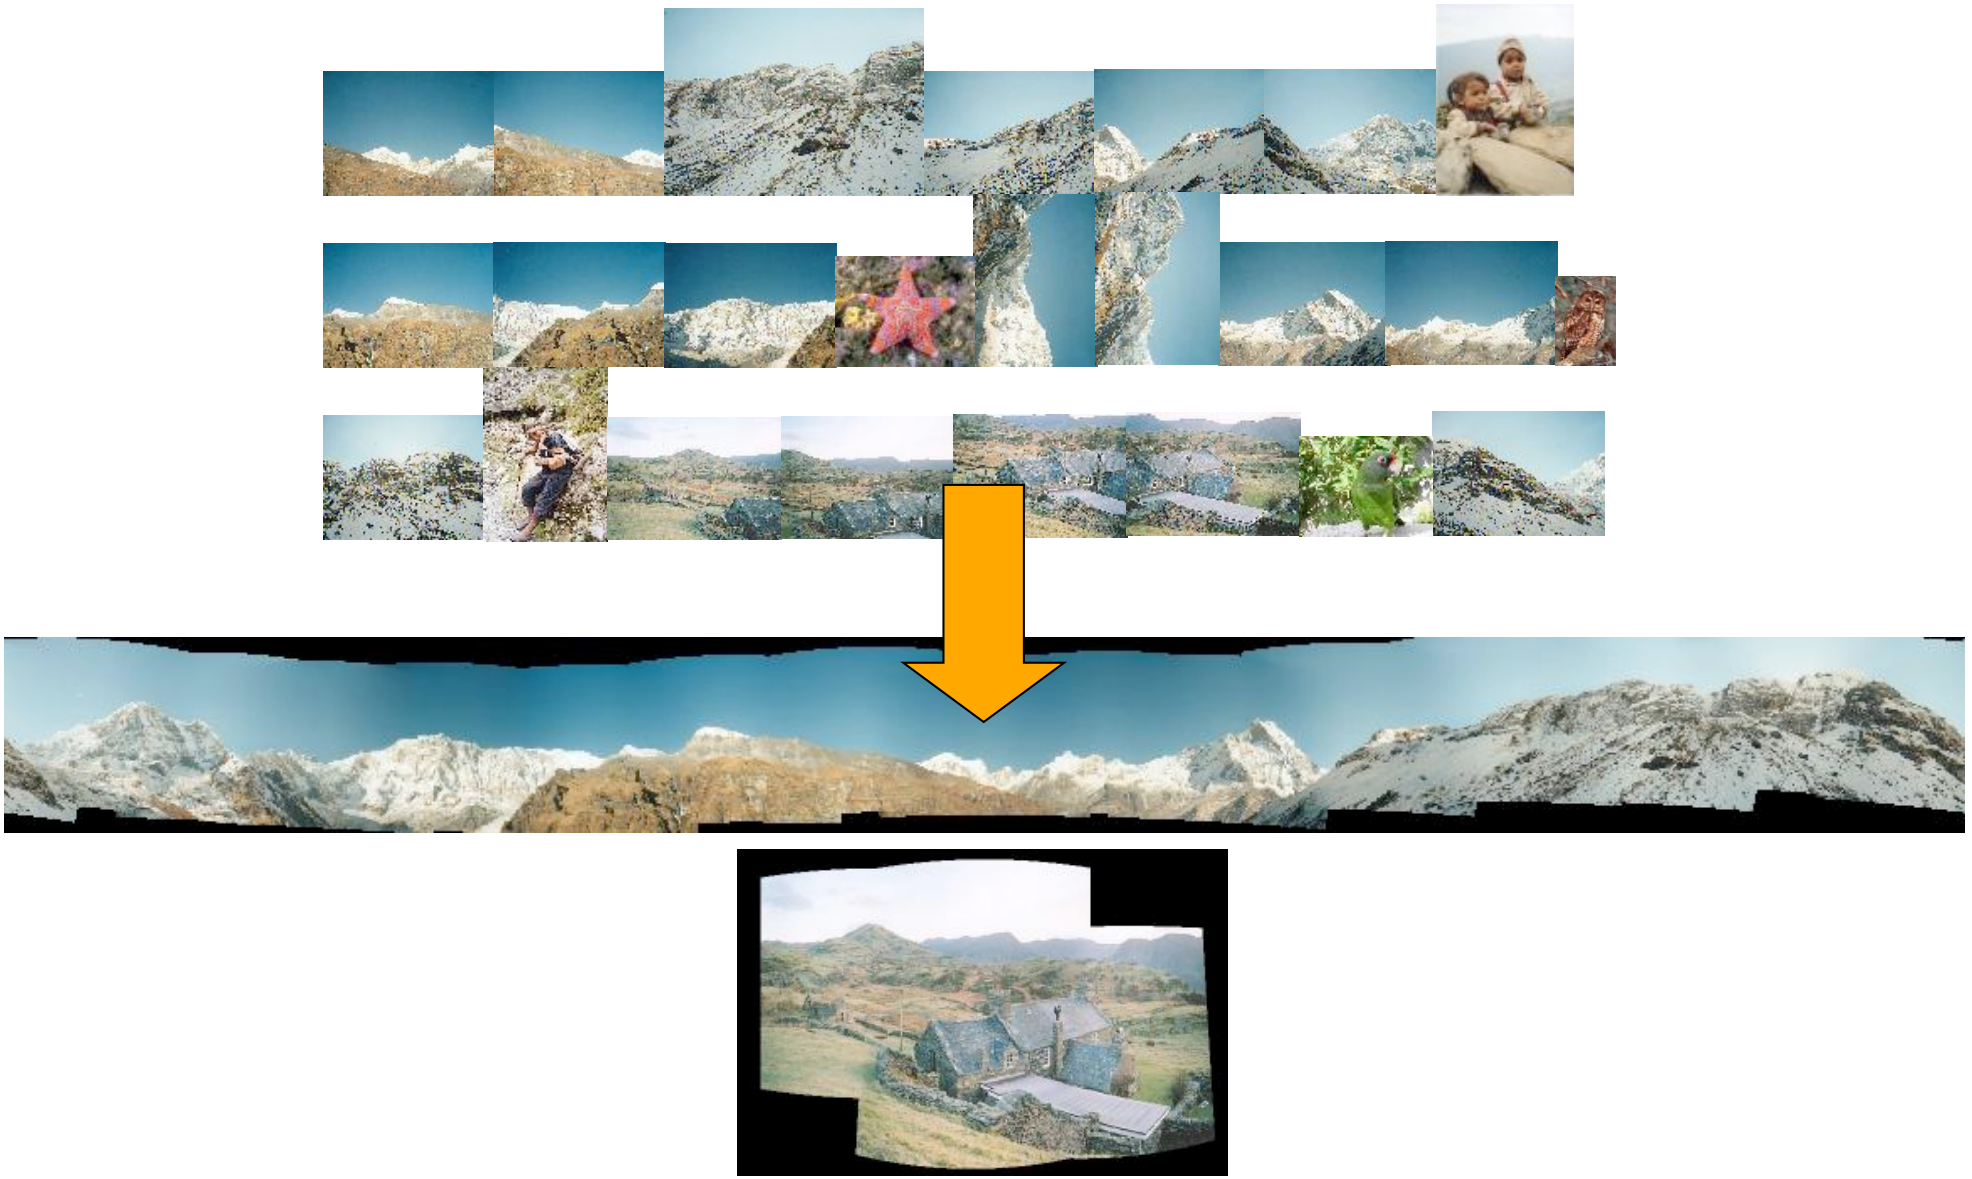

AutoStitch iPhone

Home | Usage | Gallery | FAQ | Reviews | Company | News

Automatic Image Stitching for the iPhone

AutoStitch iPhone is a fully automatic image stitcher for the iPhone. This application unleashes the power of your iPhone's camera to create wide-angle views and panoramas with any arrangment of photos.

AutoStitch uses the most advanced stitching technology available today, but it's very simple to use. To see how it works on the iPhone/iPod Touch, see our usage instructions or tutorial video.

AutoStitch iPhone brings together years of research and development experience into an amazing application that is available now on your iPhone at a very low price.

*** **Note:** If you recently upgraded to iOS4 or to version 3.0 and are having problems, please see the top of our FAQ.

M. Brown, D. Lowe, B. Hearn, J. Beis