An Overview of Multiple View Geometry and Matching

Ricardo Fabbri

LEMS - Brown University

06/2005

Plan of Presentation

- Motivation and main problems
- Single-view calibration procedures
- Projective reconstruction and calibration
  - Obtaining camera matrices from F
  - Bundle adjustment
- Auto-calibration + Metric reconstruction
- Stereo correspondence literature survey

Main problem 1

Hottest problem in Hartley's Book:

- Given corresponding features across multiple uncalibrated views, guess:
  - Camera motion and internal parameters
  - Metric reconstruction
  - Deal with noise, mismatches, and outliers

Main problem 2

Another hot (and harder) problem

- Determine correspondences between multiple views
  - Views may be totally uncalibrated
  - Or camera structure may be known
    - Fundamental matrix
    - Or even full calibration

Problem solving

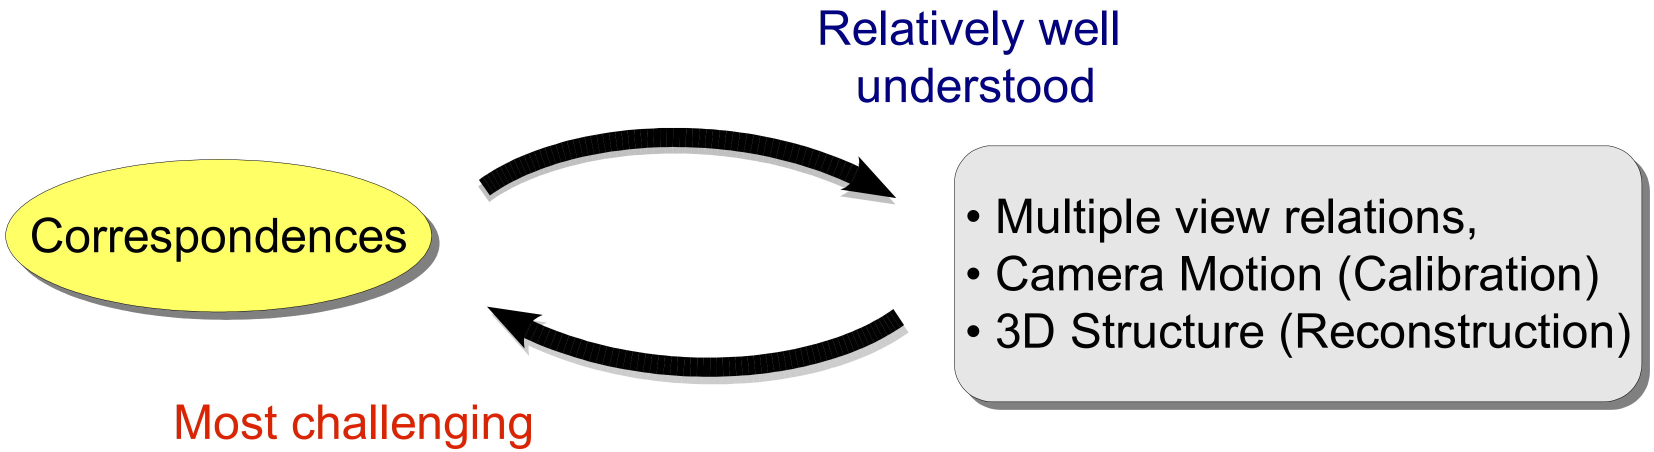

Reconstruction Process

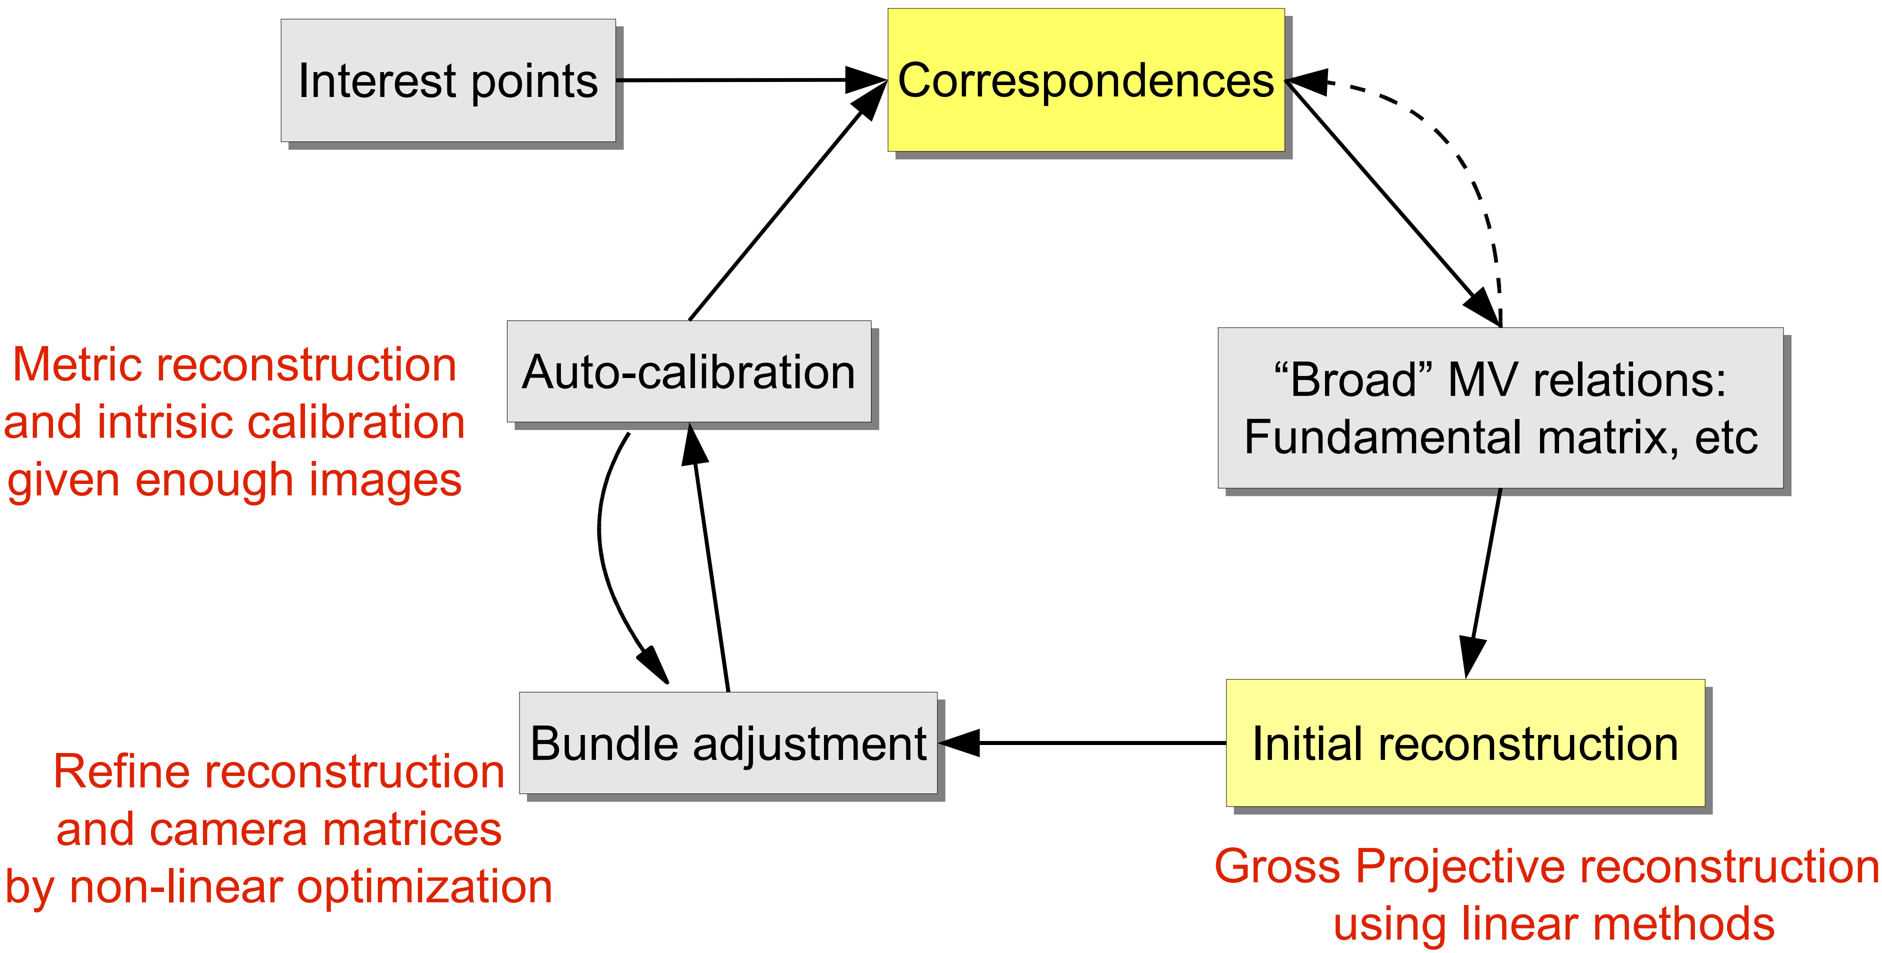

2D Projective transforms

Planar Projective transformation

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

or $\mathbf{x}' = \mathbf{H}\,\mathbf{x}$

8DOF

Any invertible linear map on homogeneous coordinates

projectivity=collineation=projective transformation=homography

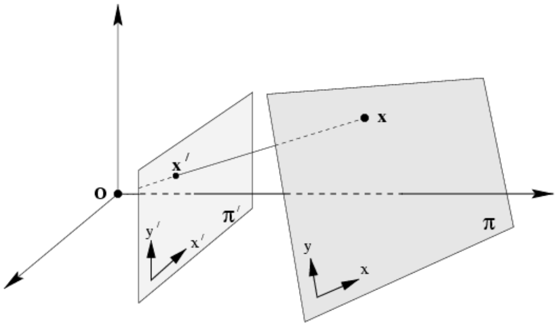

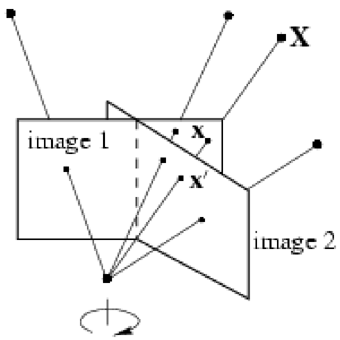

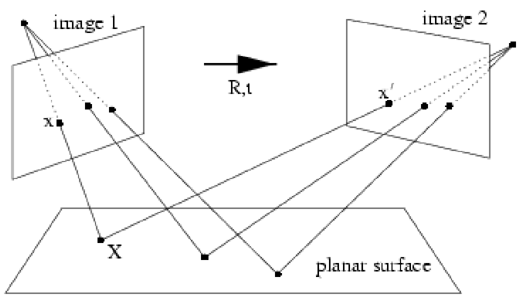

Groups of transforms

$$\mathrm{x}' = \mathbf{H}_P\,\mathrm{x} = \begin{bmatrix} \mathbf{A} & \mathrm{t} \\ \mathrm{v}^{\mathsf{T}} & v \end{bmatrix} \mathrm{x} \qquad \mathrm{v} = (v_1, v_2)^{\mathsf{T}}$$

8DOF (computable from 4 point-correspondences)

$$\mathbf{H} = \mathbf{H}_S \mathbf{H}_A \mathbf{H}_P = \begin{bmatrix} s\mathbf{R} & \mathrm{t} \\ 0^{\mathsf{T}} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{K} & 0 \\ 0^{\mathsf{T}} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I} & 0 \\ \mathrm{v}^{\mathsf{T}} & v \end{bmatrix} = \begin{bmatrix} \mathbf{A} & \mathrm{t} \\ \mathrm{v}^{\mathsf{T}} & v \end{bmatrix}$$

The line at infinity l_∞ is a fixed line under a projective transformation H if and only if H is an affinity

Affine rectification

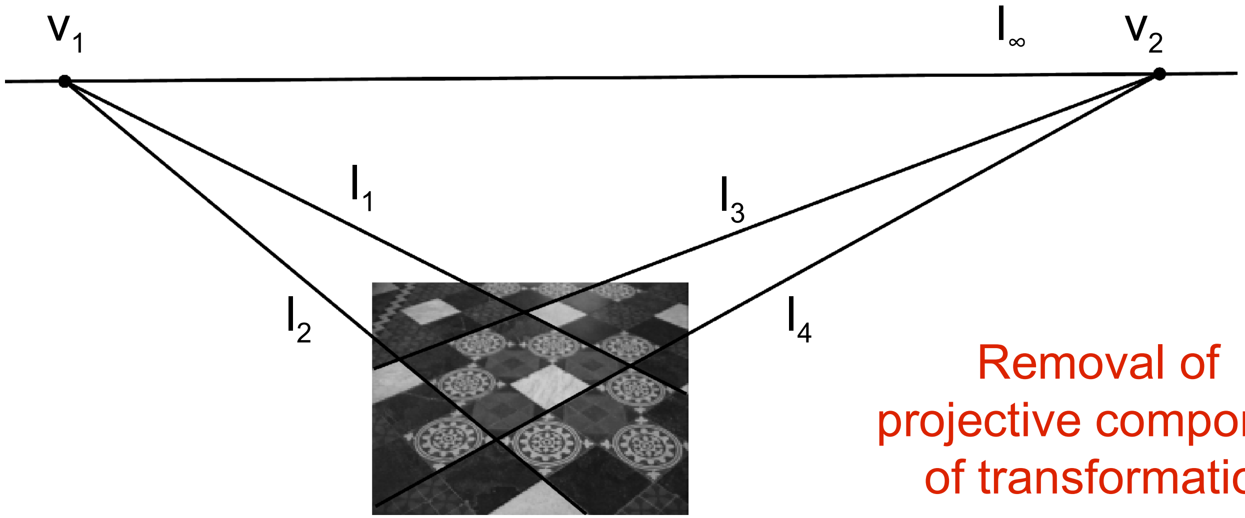

Removal of projective component of transformation

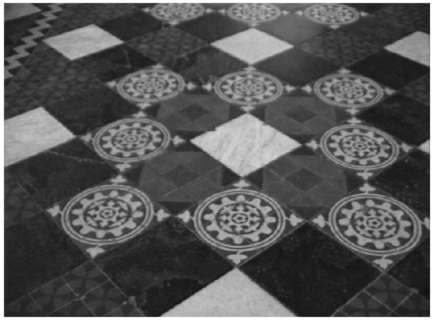

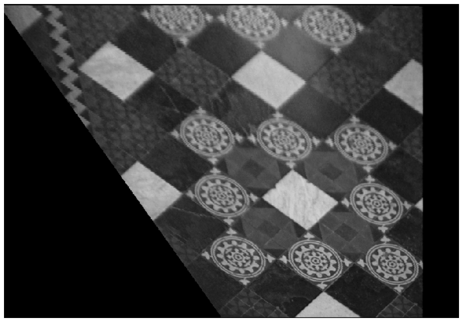

The circular points

$$\mathrm{I} = \begin{pmatrix} 1 \\ i \\ 0 \end{pmatrix} \qquad \mathrm{J} = \begin{pmatrix} 1 \\ -i \\ 0 \end{pmatrix}$$

Eigenvalues of similarities

The circular points I, J are fixed points under the projective transformation **H** iff **H** is a similarity

Conic dual to circular points

- Conic = 2nd degree homog. eq.
  - 3 x 3 symmetric matrix $\mathrm{x}^{\mathsf{T}}\mathbf{C}\,\mathrm{x}=0$
  - Dual conics = line conics: $\mathrm{l}^{\mathsf{T}}\mathbf{C}^{*}\,\mathrm{l}=0$ $\mathbf{C}^{*}=\mathbf{C}^{-1}$
- Conic dual to circular points I, J

$$\mathbf{C}_{\infty}^{*}=\mathrm{IJ}^{\mathsf{T}}+\mathrm{JI}^{\mathsf{T}}$$

  - All lines through I or J.

The dual conic $\mathbf{C}_{\infty}^{*}$ is fixed conic under the projective transformation **H** iff **H** is a similarity

Packages both circular points and l_{∞} (null vector)

Conic dual to circular points

- $\mathbf{C}_\infty^*$ Packs both circular points and l_∞ (null vector)
  - Represents information needed for determining structure up to similarity
- Enables measurement of angles

$$\cos\theta = \frac{\mathrm{l}^\mathsf{T}\mathbf{C}_\infty^*\mathrm{m}}{\sqrt{(\mathrm{l}^\mathsf{T}\mathbf{C}_\infty^*\mathrm{l})(\mathrm{m}^\mathsf{T}\mathbf{C}_\infty^*\mathrm{m})}}$$

$\boxed{\mathrm{l}^\mathsf{T}\mathbf{C}_\infty^*\mathrm{m} = 0 \quad \text{(l and m are orthogonal)}}$

Metric rectification

From affinity

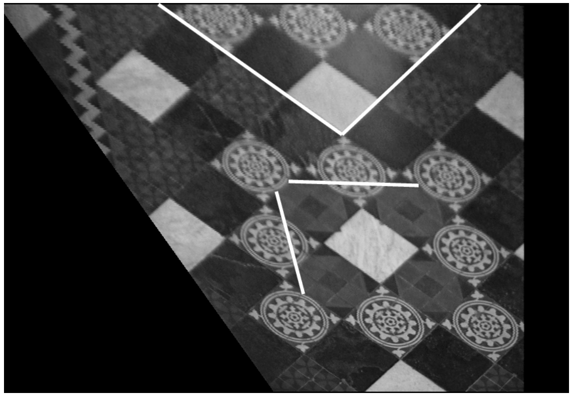

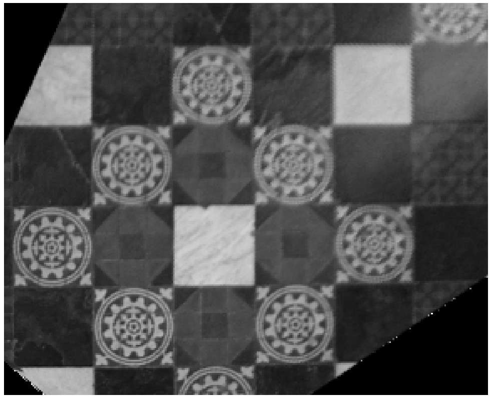

From projectivity

3D Projective transforms

3D projective transformation

$$\mathbf{X}' = \mathbf{H}\,\mathbf{X}$$

Any invertible 4x4 linear map on homogeneous coordinates

Dual: points $\leftrightarrow$ planes, lines $\leftrightarrow$ lines

3D Projective transforms

Projective 15dof	$\begin{bmatrix} A & t \\ v^T & v \end{bmatrix}$	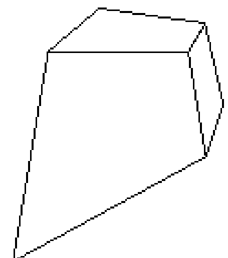	Intersection and tangency
Affine 12dof	$\begin{bmatrix} A & t \\ 0^T & 1 \end{bmatrix}$	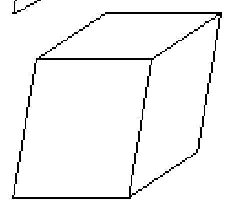	Parallellism of planes, Volume ratios, centroids, **The plane at infinity π_∞**
Similarity 7dof	$\begin{bmatrix} sR & t \\ 0^T & 1 \end{bmatrix}$	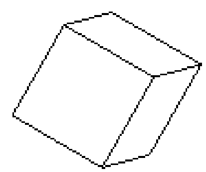	**The absolute conic Ω_∞**
Euclidean 6dof	$\begin{bmatrix} R & t \\ 0^T & 1 \end{bmatrix}$	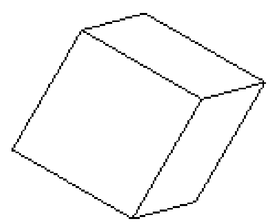	Volume

3D Projective transforms

The plane at infinity π_∞ is a fixed plane under a projective transformation H iff H is an affinity

1. canical position $\pi_\infty = (0,0,0,1)^\mathsf{T}$
2. contains directions $\mathrm{D} = (X_1, X_2, X_3, 0)^\mathsf{T}$
3. two planes are parallel $\Leftrightarrow$ line of intersection in π_∞
4. line // line (or plane) $\Leftrightarrow$ point of intersection in π_∞
5. Identifying π_∞ enables removal of projective "distortion"

The Absolute conic

- Ω_∞ is a conic with matrix I on π_∞

Canonical form: $$\left.\begin{array}{c} X_1^2 + X_2^2 + X_3^2 \\ X_4 \end{array}\right\} = 0$$

only imaginary points at infinity (!)

The absolute conic Ω_∞ is a fixed conic under the projective transformation **H** iff **H** is a similarity

- Encodes 5 DOF of affine transformation
- Identifying it enables removal of affine distortion

The Absolute conic

- Ω_∞ enables measuring angles

$$\cos\theta = \frac{\left(d_1^\mathsf{T}\Omega_\infty d_2\right)}{\sqrt{\left(d_1^\mathsf{T}\Omega_\infty d_1\right)\left(d_2^\mathsf{T}\Omega_\infty d_2\right)}}$$

- Orthogonality: $\boxed{d_1^\mathsf{T}\Omega_\infty d_2 = 0}$

The Absolute Dual Quadric

- Quadrics
  - Surfaces in P^3 defined by

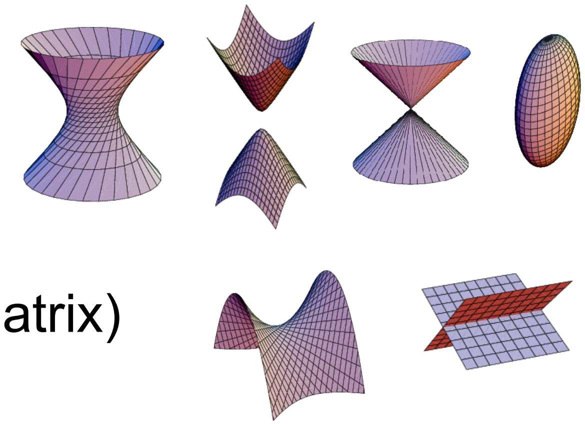

$$X^T Q X = 0 \qquad \text{(Q : 4x4 symmetric matrix)}$$

1. 9 DOF (9 points define quadric)
2. (plane ∩ quadric) = conic

- Dual quadrics
  - Equation on (tangent) planes $\pi^T Q^* \pi = 0$

1. $Q^* = Q^{-1}$ (non-degenerate)

The Absolute Dual Quadric

- Absolute dual quadric Q^*_∞
  - Set of tangent planes to absolute conic
  - Encodes both π_∞ and Ω_∞
  - 8 D.O.F. specifying projective and affine transforms, leaving only similarity

The absolute dual quadric Q^*_∞ is a fixed quadric under the projective transformation **H** iff **H** is a similarity

Estimation of multiview mappings

- 2D homography

  Given a set of (x_i,x_i'), compute H $(x_i'=Hx_i)$

- 3D to 2D camera projection

  Given a set of (X_i,x_i), compute P $(x_i=PX_i)$

- Fundamental matrix

  Given a set of (x_i,x_i'), compute F $(x_i'^{T}Fx_i=0)$

2D Homography estimation

- 4 point correspondences determine H
- In practice, there is error, so use many correspondences
- Minimize cost functions
  - Direct Linear Transformation
    - Least-squares (SVD) solution: $A\mathbf{h} \sim 0$
    - Minizes an algebraic residual, can be biased
    - Requires normalization of data
    - Advantage: fast, unique solution
    - Initial solution for iterative methods

2D Homography estimation

- Geometric cost function minimization

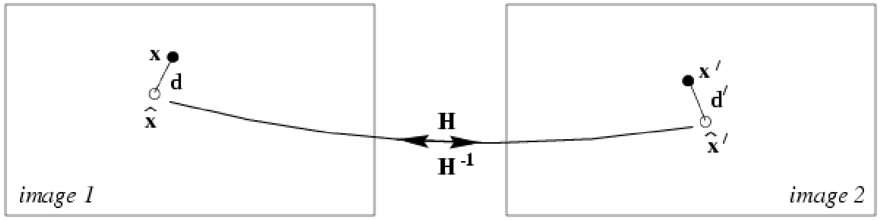

$$\left(\hat{\mathrm{H}}, \hat{\mathrm{x}}_i, \hat{\mathrm{x}}_i'\right) = \underset{\mathrm{H}, \hat{\mathrm{x}}_i, \hat{\mathrm{x}}_i'}{\operatorname{argmin}} \sum_i d\left(\mathrm{x}_i, \hat{\mathrm{x}}_i\right)^2 + d\left(\mathrm{x}_i', \hat{\mathrm{x}}_i'\right)^2$$

$$\text{subject to } \hat{\mathrm{x}}_i' = \hat{\mathrm{H}}\hat{\mathrm{x}}_i$$

- Use Levenberg-Marquadt iteration in VXL
- DLT as initial solution

2D Homography estimation

Objective

Automatically compute homography between two images

Algorithm

(iv) **Interest points:** Compute interest points in each image

(v) **Putative correspondences:** Compute a set of interest point matches based on some similarity measure

(vi) **RANSAC robust estimation:** Choose H with most inliers

(vii) **Optimal estimation:** re-estimate H from all inliers by minimizing geom. cost function with Levenberg-Marquardt

(viii) **Guided matching:** Determine more matches using prediction by computed H

Optionally iterate last two steps until stability

2D Homography estimation

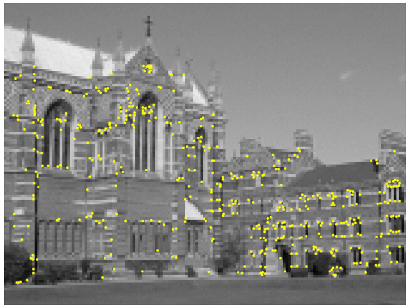
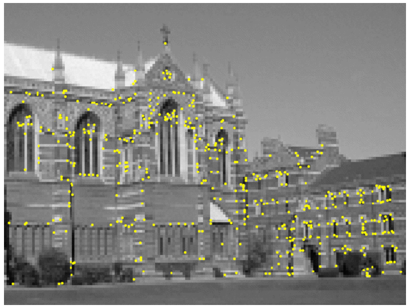

Interest points
(500/image)


Putative
correspondences (268)

Outliers (117)

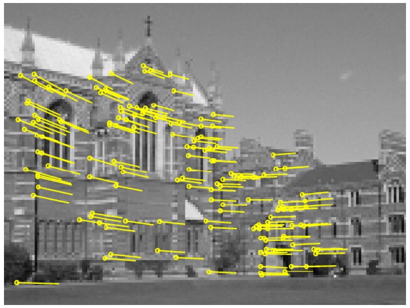
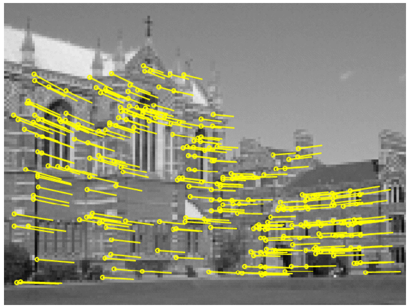

Inliers (151)

Final inliers (262)

Basic camera calibration

$$x = PX \qquad P = K[I\,|\,0]\begin{bmatrix} R & t \\ 0 & 1 \end{bmatrix}$$

$$K = \begin{bmatrix} \alpha_x & s & x_0 \\ & \alpha_y & y_0 \\ & & 1 \end{bmatrix}$$

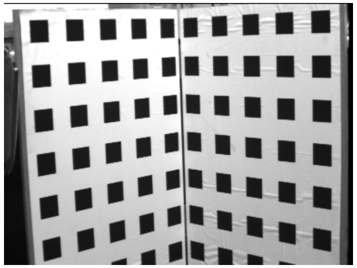

- 3x4 general homog. matrix, 11 DOF
- Minimum 6 3D to 2D point correspondences

$$Ap = 0$$

- Again, use DLT for minimizing $\|Ap\|$

Basic camera calibration

- Levenberg-Marquadt for minimizing geometric error
  - Assuming high precision in 3D
  - Geometric error:

$$\sum_i d(\mathbf{x}_i, \hat{\mathbf{x}}_i)^2$$

$$\min_{\mathrm{P}} \sum_i d(\mathbf{x}_i, \mathrm{P}\mathbf{X}_i)^2$$

- Distortion correction...

More about internal calibration

- Image of the absolute conic (**IAC**)
  - By projecting Ω_∞, one arrives at:
    - $\omega = \left(KK^T\right)^{-1} = K^{-T}K^{-1}$
    - Its dual (**DIAC**): $\omega^* = KK^T$
    - Independent of camera position or orientation!

$$\cos\theta = \frac{x_1^T \omega x_2}{\sqrt{\left(x_1^T \omega x_1\right)\left(x_2^T \omega x_2\right)}}$$

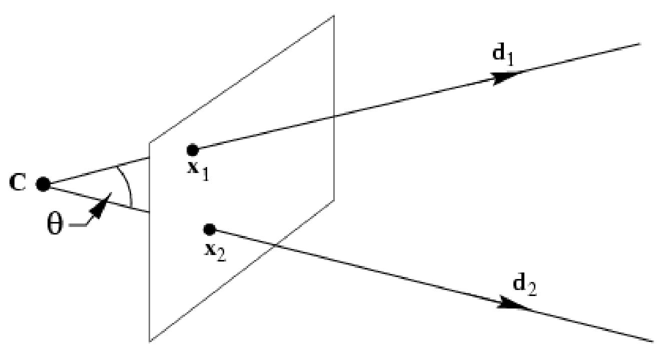

A simple calibration device

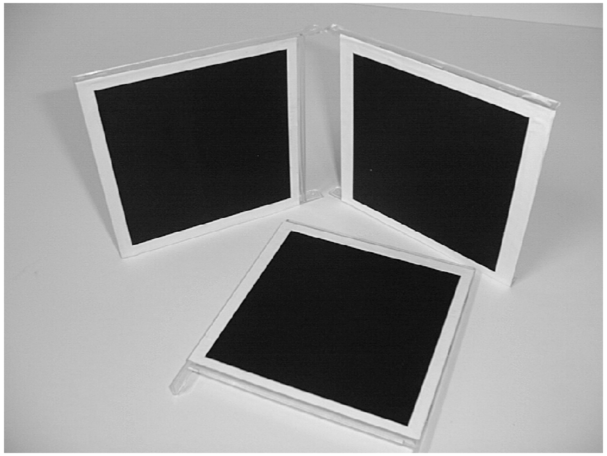

(i) compute H for each square
(corners (0,0),(1,0),(0,1),(1,1))
(iii) compute the imaged circular points $H(1,\pm i,0)^T$
(iv) fit a conic to 6 circular points
(v) compute K from ω

(= Zhang's calibration method)

Other constraints on K

- We may combine many different linear constraints on the IAC and then fit the conic and recover K
- Examples of scene constraints:
  - Planar homographies, as just seen
  - Vanishing points corresponding to orthogonal lines
- Examples of internal constraints
  - Zero skew and square pixels
- All these constraints are interpreted as known points lying on the conic or conjugate to it

The fundamental matrix

- F is the unique 3x3 rank 2 matrix that satisfies $x'^T F x = 0$ for all $x \leftrightarrow x'$
- **F** has 7 d.o.f.
  - 3x3-1(homogeneous) – 1(rank2)
  - 7-point correspondences minimum
  - Pair of camera matrices determine F uniquely
  - **F determines camera matrices up to projective ambiguity**

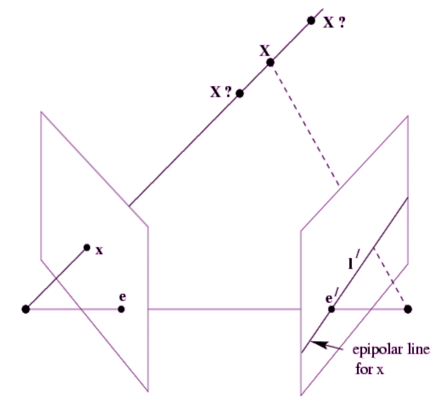

$$P = [I \mid 0] \quad P' = [[e']_\times F + e' v^T \mid \lambda e']$$

Reconstruction from 2 uncalibrated views

- given $x_i \leftrightarrow x'_i$, compute P,P' and X_i

$$x_i = PX_i \qquad x'_i = PX'_i \qquad \text{for all } i$$

- Without additional information, possible up to projective ambiguity

  (i) Compute F from correspondences
  (ii) Compute camera matrices from F
  (iii) Compute 3D point for each pair of corresponding points (triangulation)

Reconstruction from 2 uncalibrated views

- Projective reconstruction from F

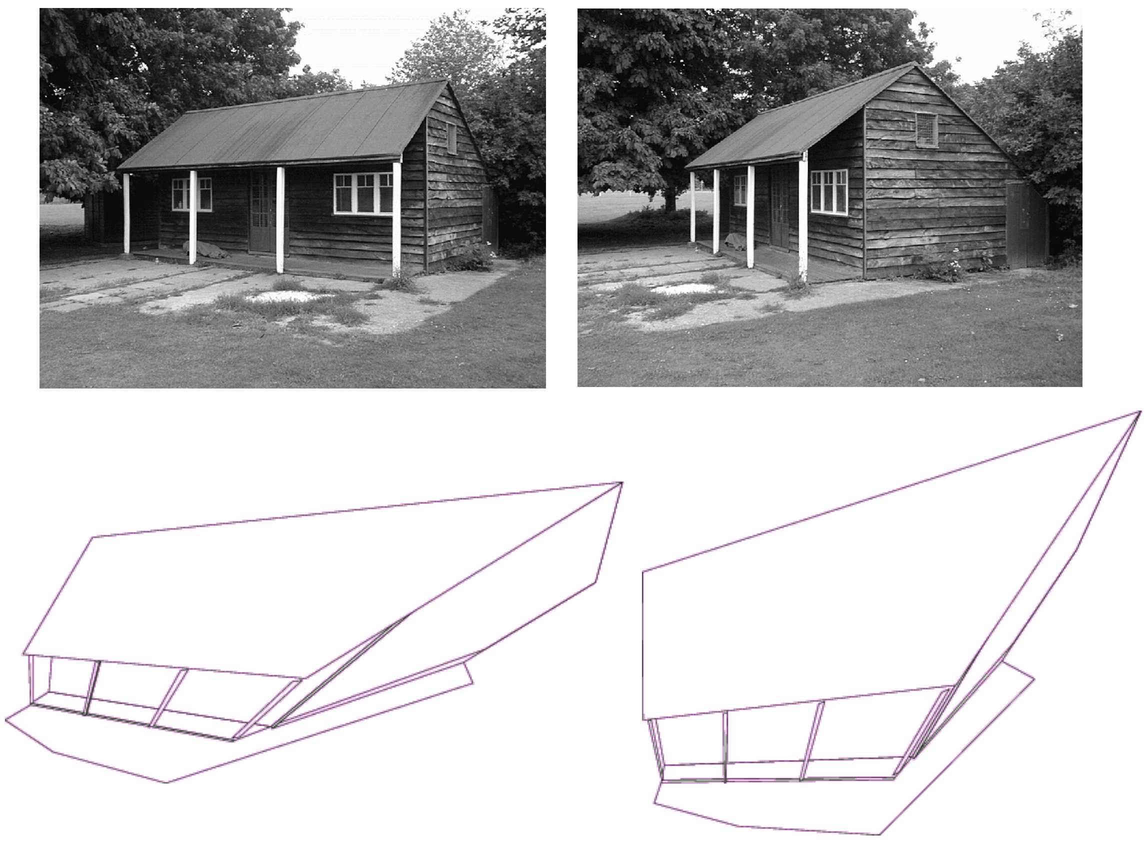

Reconstruction from 2 uncalibrated views

- Ultimate goal: metric reconstruction
  - Only similarity ambiguity

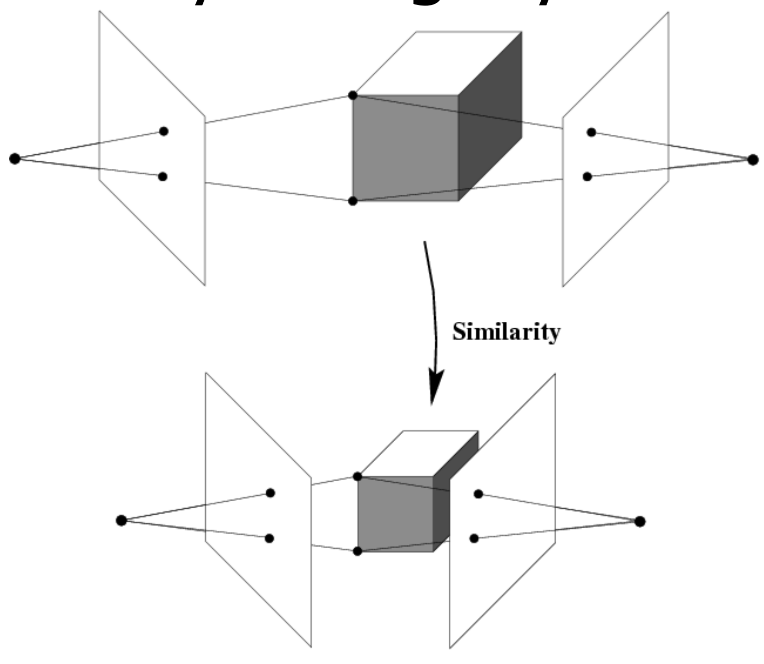

Stratified reconstruction

(i) Projective reconstruction

Hardest (ii) Affine reconstruction

(iii) Metric reconstruction

Projective to affine

- Identify π_∞ (3 points) using additional information
  - Translational camera motion

$$F = [e]_\times = [e']_\times \qquad P = [I \mid 0]$$

$$P = [I \mid e']$$

  - Scene constraints (similar to planar case)

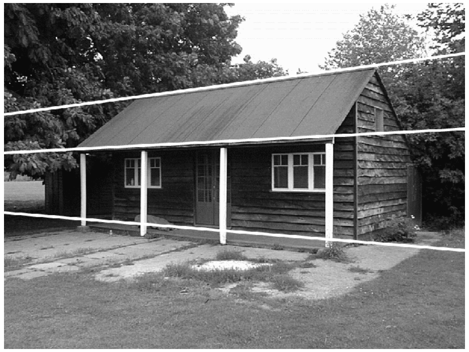

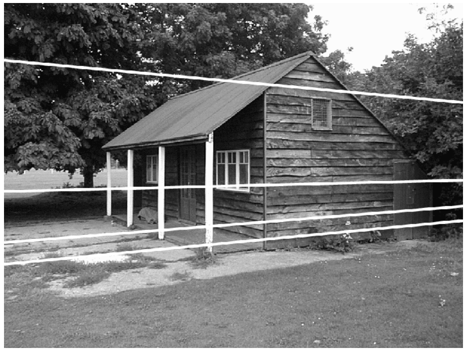

Affine to metric

- Identify absolute conic Ω_∞
  - Then apply 3D "rectification" that maps it to canonical coordinates in Euclidean world,
  - $$\Omega_\infty : X^2 + Y^2 + Z^2 = 0, \text{ on } \pi_\infty$$
- In practice, just find IAC ω in some image
  - Single view constraints as seen before:
    - Planar homographies
    - Vanishing points corresponding to orthogonal lines
    - Zero skew and square pixels

Affine to metric

- Multiple view constraints on Ω_∞
  - Idea used in auto-calibration
  - Consider same intrinsics/same ω on all cameras
    - Given sufficient images there is in general only one conic that projects to the same ω in all images:
      - The absolute conic Ω_∞
- Direct metric reconstruction
  - Ground control points (5 or more)

Bundle adjustment

- Given n correspondences across m views
  - Determine camera matrices and refine correspondences
  - minimize reprojection error

$$\min_{\hat{P}_k,\hat{M}_i} \sum_{k=1}^{m} \sum_{i=1}^{n} D\left(\mathrm{m}_{ki}, \hat{P}_k \hat{M}_i\right)^2$$

  - Levenberg-Marquadt
    - Needs specialized implementation (Matt)
- Used to refine reconstructions in many occasions

Bundle adjustment

- To many images or correspondences
  - Strategies so that not all images are optimized simultaneously
  - Partition data, bundle adjust separately, then merge
- Computation of initial structure and motion
  - According to Hartley and Zisserman:
    - “this area is still to some extend a black-art”
    - Correspondences not present in all views
      - Use overlapping subsequences
      - Stitch into final reconstruction
      - Triangulate to transfer correspondences to all views

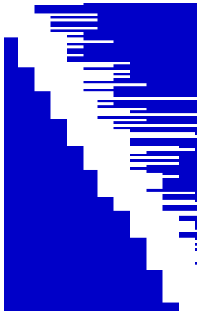

Auto-Calibration

- Metric reconstruction and intrinsics
- All we need are:
  - correspondences
  - sufficient number of views
  - assumptions on internal calibration or camera motion
- We want to find rectifying 3D homography H
  - H is completely determined by Ω_∞ and π_∞
  - Or absolute dual quadric Q^*_∞
    - K of 1st camera and π_∞ suffices: 8 parameters

Auto-Calibration

- Special imaging conditions that constrains K
  - Camera rotating about center
  - Turntable motion
- Internal constraints
  - Zero skew, fixed focal length, etc
- Strategy based on absolute dual quadric
    - Q^*_∞ is a fixed quadric under Euclidean transformations
    - DIAC $\omega^{*i} = \mathbf{K_i K_i^T}$ is its image on each view
    - So we have a relation between calibrations on each view

Auto-Calibration

- Old method based on Kruppa equations
  - Constraint based on correspondences of epipolar lines tangent to the IAC
  - Useful when only 2 views available
- Stratified strategy
  - Identify π_∞ and then K
  - π_∞ is the hardest part
    - General motion and constant parameters
    - Other ways as seen before
      - Translational motion
      - 3 vanishing points, etc

Auto-Calibration