


Last lecture

- Passive Stereo
- Spacetime Stereo




Today

- Structure from Motion:
Given pixel correspondences,
how to compute 3D structure and camera motion?

Epipolar geometry & fundamental matrix


The epipolar geometry



What if only C, C', x are known?


The epipolar geometry

[epipolar geometry demo](#)




C, C', x, x' and X are coplanar

The epipolar geometry



All points on π project on l and l'

The epipolar geometry



Family of planes π and lines l and l' intersect at e and e'


The epipolar geometry

epipolar pole

[epipolar geometry demo](#)

= intersection of baseline with image plane


= projection of projection center in other image



epipolar plane = plane containing baseline

epipolar line = intersection of epipolar plane with image

The fundamental matrix F



$$X \propto KX$$

$$X' \propto K'R(X - T)$$

$$p = K^{-1}X \propto X$$

$$p' = K'^{-1}X' \propto R(X - T)$$

The equation of the epipolar plane through X is

$$(X - T)^T (T \times p) = 0 \quad \rightarrow \quad (R^T p')^T (T \times p) = 0$$

The fundamental matrix F

$$(\mathbf{R}^T \mathbf{p}')^T (\mathbf{T} \times \mathbf{p}) = 0$$

$$\mathbf{T} \times \mathbf{p} = \mathbf{S} \mathbf{p}$$


$$\mathbf{S} = \begin{bmatrix} 0 & -T_z & T_y \\ T_z & 0 & -T_x \\ -T_y & T_x & 0 \end{bmatrix}$$

→ $(\mathbf{R}^T \mathbf{p}')^T (\mathbf{S} \mathbf{p}) = 0$

→ $(\mathbf{p}'^T \mathbf{R})(\mathbf{S} \mathbf{p}) = 0$

→ $\mathbf{p}'^T \mathbf{E} \mathbf{p} = 0$ essential matrix

The fundamental matrix F



$$\mathbf{p}'^T \mathbf{E} \mathbf{p} = 0$$

The fundamental matrix F

$$\mathbf{p}'^T \mathbf{E} \mathbf{p} = 0$$

Let M and M' be the intrinsic matrices, then


$$\mathbf{p} = \mathbf{K}^{-1} \mathbf{x} \quad \mathbf{p}' = \mathbf{K}'^{-1} \mathbf{x}'$$

$$\rightarrow (\mathbf{K}'^{-1} \mathbf{x}')^T \mathbf{E} (\mathbf{K}^{-1} \mathbf{x}) = 0$$

$$\rightarrow \mathbf{x}'^T \mathbf{K}'^{-T} \mathbf{E} \mathbf{K}^{-1} \mathbf{x} = 0$$

$$\rightarrow \mathbf{x}'^T \mathbf{F} \mathbf{x} = 0 \quad \text{fundamental matrix}$$

The fundamental matrix F



$$\mathbf{p}'^T \mathbf{E} \mathbf{p} = 0$$

$$\mathbf{x}'^T \mathbf{F} \mathbf{x} = 0$$

The fundamental matrix F

- The fundamental matrix is the algebraic representation of epipolar geometry
- The fundamental matrix satisfies the condition that for any pair of corresponding points $x \leftrightarrow x'$ in the two images


$$\mathbf{x}'^T \mathbf{F} \mathbf{x} = 0 \quad \left(\mathbf{x}'^T \mathbf{1}' = 0 \right)$$

The fundamental matrix F

F is the unique 3×3 rank 2 matrix that satisfies $x'^T F x = 0$
for all $x \leftrightarrow x'$

1. Transpose: if F is fundamental matrix for (P, P') , then F^T is fundamental matrix for (P', P)
2. Epipolar lines: $l' = Fx$ & $l = F^T x'$
3. Epipoles: on all epipolar lines, thus $e'^T F x = 0, \forall x \Rightarrow e'^T F = 0$, similarly $F e = 0$
4. F has 7 d.o.f. , i.e. $3 \times 3 - 1$ (homogeneous) - 1 (rank 2)
5. F maps from a point x to a line $l' = Fx$ (not invertible)

The fundamental matrix F



- It can be used for
 - Simplifies matching
 - Allows to detect wrong matches

Estimation of F — 8-point algorithm

- The fundamental matrix F is defined by

$$\mathbf{x}'^T \mathbf{F} \mathbf{x} = 0$$

for any pair of matches \mathbf{x} and \mathbf{x}' in two images.

- Let $\mathbf{x}=(u, v, 1)^T$ and $\mathbf{x}'=(u', v', 1)^T$, $\mathbf{F} = \begin{bmatrix} f_{11} & f_{12} & f_{13} \\ f_{21} & f_{22} & f_{23} \\ f_{31} & f_{32} & f_{33} \end{bmatrix}$

each match gives a linear equation

$$uu' f_{11} + vu' f_{12} + u' f_{13} + uv' f_{21} + vv' f_{22} + v' f_{23} + uf_{31} + vf_{32} + f_{33} = 0$$

8-point algorithm

$$\begin{bmatrix}
 u_1 u_1' & v_1 u_1' & u_1' & u_1 v_1' & v_1 v_1' & v_1' & u_1 & v_1 & 1 \\
 u_2 u_2' & v_2 u_2' & u_2' & u_2 v_2' & v_2 v_2' & v_2' & u_2 & v_2 & 1 \\
 \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
 u_n u_n' & v_n u_n' & u_n' & u_n v_n' & v_n v_n' & v_n' & u_n & v_n & 1
 \end{bmatrix}
 \begin{bmatrix}
 f_{11} \\
 f_{12} \\
 f_{13} \\
 f_{21} \\
 f_{22} \\
 f_{23} \\
 f_{31} \\
 f_{32} \\
 f_{33}
 \end{bmatrix}
 = \mathbf{0}$$

- In reality, instead of solving $\mathbf{A}\mathbf{f} = \mathbf{0}$, we seek \mathbf{f} to minimize $\|\mathbf{A}\mathbf{f}\|$, least eigenvector of $\mathbf{A}^T \mathbf{A}$

8-point algorithm

- To enforce that \mathbf{F} is of rank 2, \mathbf{F} is replaced by \mathbf{F}' that minimizes $\|\mathbf{F} - \mathbf{F}'\|$ subject to $\det \mathbf{F}' = 0$.
- It is achieved by SVD. Let $\mathbf{F} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T$, where

$$\mathbf{\Sigma} = \begin{bmatrix} \sigma_1 & 0 & 0 \\ 0 & \sigma_2 & 0 \\ 0 & 0 & \sigma_3 \end{bmatrix}, \text{ let } \mathbf{\Sigma}' = \begin{bmatrix} \sigma_1 & 0 & 0 \\ 0 & \sigma_2 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

then $\mathbf{F}' = \mathbf{U}\mathbf{\Sigma}'\mathbf{V}^T$ is the solution.

8-point algorithm

% Build the constraint matrix

```
A = [x2(1,:)' .* x1(1,:)'  x2(1,:)' .* x1(2,:)'  x2(1,:)' ...  
     x2(2,:)' .* x1(1,:)'  x2(2,:)' .* x1(2,:)'  x2(2,:)' ...  
     x1(1,:)'              x1(2,:)'              ones(npts,1) ];
```

```
[U,D,V] = svd(A);
```

% Extract fundamental matrix from the column of V

% corresponding to the smallest singular value.

```
F = reshape(V(:,9),3,3)';
```

% Enforce rank2 constraint

```
[U,D,V] = svd(F);
```

```
F = U*diag([D(1,1) D(2,2) 0])*V';
```

8-point algorithm

- Pros: it is linear, easy to implement and fast
- Cons: susceptible to noise

Problem with 8-point algorithm

$$\begin{bmatrix}
 u_1 u_1' & v_1 u_1' & u_1' & u_1 v_1' & v_1 v_1' & v_1' & u_1 & v_1 & 1 \\
 u_2 u_2' & v_2 u_2' & u_2' & u_2 v_2' & v_2 v_2' & v_2' & u_2 & v_2 & 1 \\
 \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots \\
 u_n u_n' & v_n u_n' & u_n' & u_n v_n' & v_n v_n' & v_n' & u_n & v_n & 1
 \end{bmatrix}
 \begin{bmatrix}
 f_{11} \\
 f_{12} \\
 f_{13} \\
 f_{21} \\
 f_{22} \\
 f_{23} \\
 f_{31} \\
 f_{32} \\
 f_{33}
 \end{bmatrix}
 = \mathbf{0}$$

~ 10000 ~ 10000 ~ 100 ~ 10000 ~ 10000 ~ 100 ~ 100 ~ 100 1




Orders of magnitude difference
 between column of data matrix
 → least-squares yields poor results

Normalized 8-point algorithm

normalized least squares yields good results

Transform image to $\sim[-1,1] \times [-1,1]$



Normalized 8-point algorithm

1. Transform input by $\hat{\mathbf{x}}_i = \mathbf{T}\mathbf{x}_i$, $\hat{\mathbf{x}}_i' = \mathbf{T}'\mathbf{x}_i'$
2. Call 8-point on $\hat{\mathbf{x}}_i$, $\hat{\mathbf{x}}_i'$ to obtain $\hat{\mathbf{F}}$
3. $\mathbf{F} = \mathbf{T}'^T \hat{\mathbf{F}} \mathbf{T}$

$$\mathbf{x}'^T \mathbf{F} \mathbf{x} = 0$$
$$\hat{\mathbf{x}}'^T \mathbf{T}'^{-T} \mathbf{F} \mathbf{T}^{-1} \hat{\mathbf{x}} = 0$$

$\hat{\mathbf{F}}$

Normalized 8-point algorithm

```
[x1, T1] = normalise2dpts(x1);
```

```
[x2, T2] = normalise2dpts(x2);
```

```
A = [x2(1,:)'.*x1(1,:)'  x2(1,:)'.*x1(2,:)'  x2(1,:)' ...  
      x2(2,:)'.*x1(1,:)'  x2(2,:)'.*x1(2,:)'  x2(2,:)' ...  
      x1(1,:)'            x1(2,:)'            ones(npts,1) ];
```

```
[U,D,V] = svd(A);
```

```
F = reshape(V(:,9),3,3)';
```

```
[U,D,V] = svd(F);
```

```
F = U*diag([D(1,1) D(2,2) 0])*V';
```

```
% Denormalise
```

```
F = T2'*F*T1;
```

Normalization

```
function [newpts, T] = normalise2dpts(pts)
```

```
    c = mean(pts(1:2,:))'; % Centroid
```

```
    newp(1,:) = pts(1,:)-c(1); % Shift origin to centroid.
```

```
    newp(2,:) = pts(2,:)-c(2);
```

```
    meandist = mean(sqrt(newp(1,:).^2 + newp(2,:).^2));
```

```
    scale = sqrt(2)/meandist;
```

```
    T = [scale    0  -scale*c(1)
```

```
         0    scale -scale*c(2)
```

```
         0    0    1    ];
```

```
    newpts = T*pts;
```

RANSAC

repeat

- select minimal sample (8 matches)

- compute solution(s) for F



- determine inliers

until $\Gamma(\#inliers, \#samples) > 95\%$ or too many times

compute F based on all inliers



Results (ground truth)

■ Ground truth with standard stereo calibration





Results (8-point algorithm)

■ 8-point algorithm



Results (normalized 8-point algorithm)

■ Normalized 8-point algorithm



From F to R, T

$$\mathbf{x}'^T \mathbf{F} \mathbf{x} = 0$$

$$\mathbf{x}'^T \mathbf{M}'^{-T} \mathbf{E} \mathbf{M}^{-1} \mathbf{x} = 0$$

$$\mathbf{E} = \mathbf{M}'^T \mathbf{F} \mathbf{M}$$

If we know camera parameters

$$\mathbf{E} = \mathbf{R}[\mathbf{T}]_{\times}$$

Hartley and Zisserman, Multiple View Geometry, 2nd edition, pp 259

Application: View morphing



Application: View morphing



Main trick

- Prewarp with a homography to rectify images
- So that the two views are parallel
 - Because linear interpolation works when views are parallel




Figure 4: View Morphing in Three Steps. (1) Original images \mathcal{I}_0 and \mathcal{I}_1 are prewarped to form parallel views $\hat{\mathcal{I}}_0$ and $\hat{\mathcal{I}}_1$. (2) $\hat{\mathcal{I}}_s$ is produced by morphing (interpolating) the prewarped images. (3) $\hat{\mathcal{I}}_s$ is postwarped to form \mathcal{I}_s .

Problem with morphing

- Without rectification




Figure 2: A Shape-Distorting Morph. Linearly interpolating two perspective views of a clock (far left and far right) causes a geometric bending effect in the in-between images. The dashed line shows the linear path of one feature during the course of the transformation. This example is indicative of the types of distortions that can arise with image morphing techniques.




Figure 10: Image Morphing Versus View Morphing. Top: image morph between two views of a helicopter toy causes the in-between images to contract and bend. Bottom: view morph between the same two views results in a physically consistent morph. In this example the image morph also results in an extraneous hole between the blade and the stick. Holes can appear in view morphs as well.

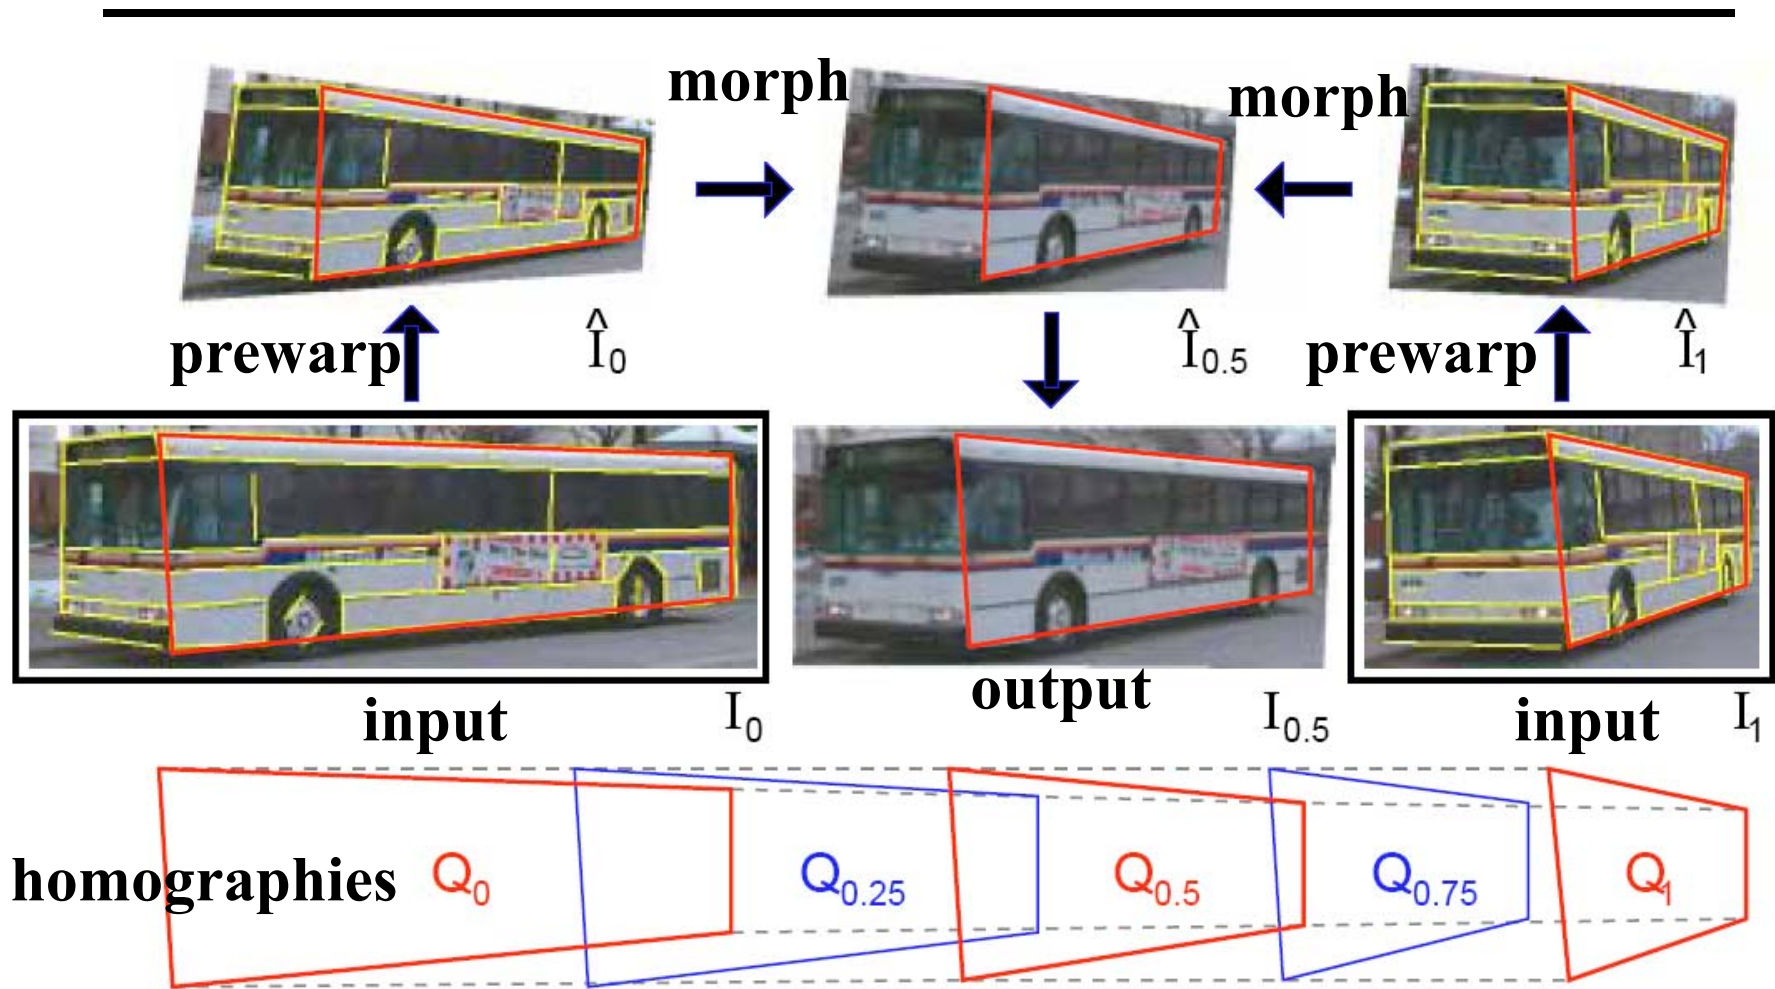


Figure 6: View Morphing Procedure: A set of features (yellow lines) is selected in original images I_0 and I_1 . Using these features, the images are automatically prewarped to produce \hat{I}_0 and \hat{I}_1 . The prewarped images are morphed to create a sequence of in-between images, the middle of which, $\hat{I}_{0.5}$, is shown at top-center. $\hat{I}_{0.5}$ is interactively postwarped by selecting a quadrilateral region (marked red) and specifying its desired configuration, $Q_{0.5}$, in $I_{0.5}$. The postwarps for other in-between images are determined by interpolating the quadrilaterals (bottom).




Figure 7: Facial View Morphs. Top: morph between two views of the same person. Bottom: morph between views of two different people. In each case, view morphing captures the change in facial pose between original images \mathcal{I}_0 and \mathcal{I}_1 , conveying a natural 3D rotation.

Video demo

Triangulation


- Problem: Given some points in *correspondence* across two or more images (taken from calibrated cameras), $\{(u_j, v_j)\}$, compute the 3D location **X**

Triangulation

- **Method I:** intersect viewing rays in 3D, minimize:

$$\arg \min_{\mathbf{X}} \sum_j \|\mathbf{C}_j + s\mathbf{V}_j - \mathbf{X}\|$$

- \mathbf{X} is the unknown 3D point
- \mathbf{C}_j is the optical center of camera j
- \mathbf{V}_j is the *viewing ray* for pixel (u_j, v_j)
- s_j is unknown distance along \mathbf{V}_j
- Advantage: geometrically intuitive



Triangulation

- **Method II:** solve linear equations in **X**

- advantage: very simple

$$u_i = \frac{m_{00}X_i + m_{01}Y_i + m_{02}Z_i + m_{03}}{m_{20}X_i + m_{21}Y_i + m_{22}Z_i + 1}$$


$$v_i = \frac{m_{10}X_i + m_{11}Y_i + m_{12}Z_i + m_{13}}{m_{20}X_i + m_{21}Y_i + m_{22}Z_i + 1}$$

- **Method III:** non-linear minimization

- advantage: most accurate (image plane error)

Structure from motion

Structure from motion




structure from motion: automatic recovery of camera motion and scene structure from two or more images. It is a self calibration technique and called *automatic camera tracking* or *matchmoving*.

Applications

- For computer vision, multiple-view shape reconstruction, novel view synthesis and autonomous vehicle navigation.
- For film production, seamless insertion of CGI into live-action backgrounds


Structure from motion



SFM pipeline

Structure from motion

- Step 1: Track Features
 - Detect good features, Shi & Tomasi, SIFT
 - Find correspondences between frames
 - Lucas & Kanade-style motion estimation
 - window-based correlation
 - SIFT matching




Structure from Motion

- Step 2: Estimate Motion and Structure
 - Simplified projection model, e.g., **[Tomasi 92]**
 - 2 or 3 views at a time **[Hartley 00]**



Structure from Motion

- Step 3: Refine estimates
 - “Bundle adjustment” in photogrammetry
 - Other iterative methods



Structure from Motion

- Step 4: Recover surfaces (image-based triangulation, silhouettes, stereo...)




Example : Photo Tourism



Photo Tourism

Exploring photo collections in 3D

Microsoft



(a)




(b)




(c)

Factorization methods

Problem statement

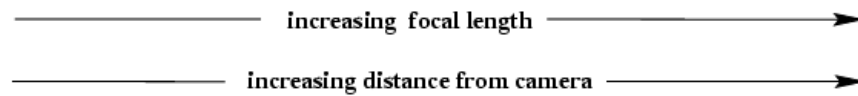


Other projection models



perspective

weak perspective



SFM under orthographic projection

2D image point orthographic projection matrix 3D scene point image offset

$$\mathbf{q} = \mathbf{\Pi} \mathbf{p} + \mathbf{t}$$

2×1 2×3 3×1 2×1

- Trick
 - Choose scene origin to be centroid of 3D points
 - Choose image origins to be centroid of 2D points
 - Allows us to drop the camera translation:

$$\mathbf{q} = \mathbf{\Pi} \mathbf{p}$$

factorization (Tomasi & Kanade)

projection of n features in one image:

$$\begin{bmatrix} \mathbf{q}_1 & \mathbf{q}_2 & \cdots & \mathbf{q}_n \end{bmatrix} = \mathbf{\Pi} \begin{bmatrix} \mathbf{p}_1 & \mathbf{p}_2 & \cdots & \mathbf{p}_n \end{bmatrix}$$

$2 \times n$ 2×3 $3 \times n$

projection of n features in m images

$$\begin{bmatrix} \mathbf{q}_{11} & \mathbf{q}_{12} & \cdots & \mathbf{q}_{1n} \\ \mathbf{q}_{21} & \mathbf{q}_{22} & \cdots & \mathbf{q}_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{q}_{m1} & \mathbf{q}_{m2} & \cdots & \mathbf{q}_{mn} \end{bmatrix} = \begin{bmatrix} \mathbf{\Pi}_1 \\ \mathbf{\Pi}_2 \\ \vdots \\ \mathbf{\Pi}_m \end{bmatrix} \begin{bmatrix} \mathbf{p}_1 & \mathbf{p}_2 & \cdots & \mathbf{p}_n \end{bmatrix}$$

$2m \times n$ $2m \times 3$ $3 \times n$

\mathbf{W} measurement

\mathbf{M} motion

\mathbf{S} shape

Key Observation: $rank(\mathbf{W}) \leq 3$

Factorization

$$\text{known} \text{ --- } \underbrace{\mathbf{W}}_{2m \times n} = \underbrace{\mathbf{M} \mathbf{S}}_{\substack{2m \times 3 & 3 \times n}} \text{ --- solve for}$$

- Factorization Technique

- W is at most rank 3 (assuming no noise)
- We can use *singular value decomposition* to factor W :

$$\underbrace{\mathbf{W}}_{2m \times n} = \underbrace{\mathbf{M}'}_{2m \times 3} \underbrace{\mathbf{S}'}_{3 \times n}$$

- S' differs from S by a linear transformation A :

$$\mathbf{W} = \mathbf{M}'\mathbf{S}' = (\mathbf{M}\mathbf{A}^{-1})(\mathbf{A}\mathbf{S})$$

- Solve for A by enforcing *metric* constraints on M

Metric constraints

- Orthographic Camera
 - Rows of Π are orthonormal: $\Pi_i \Pi_i^T = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$
- Enforcing “Metric” Constraints
 - Compute \mathbf{A} such that rows of \mathbf{M} have these properties

$$\mathbf{M}' \mathbf{A} = \mathbf{M}$$


Trick (not in original Tomasi/Kanade paper, but in followup work)

- Constraints are linear in $\mathbf{A}\mathbf{A}^T$:

$$\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} = \Pi_i \Pi_i^T = \Pi_i' \mathbf{A} (\Pi_i'^T \mathbf{A})^T = \Pi_i' \mathbf{G} \Pi_i'^T \quad \text{where } \mathbf{G} = \mathbf{A}\mathbf{A}^T$$

- Solve for \mathbf{G} first by writing equations for every Π_i in \mathbf{M}
- Then $\mathbf{G} = \mathbf{A}\mathbf{A}^T$ by SVD

Results



Extensions to factorization methods

- Paraperspective [Poelman & Kanade, PAMI 97]
- Sequential Factorization [Morita & Kanade, PAMI 97]
- Factorization under perspective [Christy & Horaud, PAMI 96] [Sturm & Triggs, ECCV 96]
- Factorization with Uncertainty [Anandan & Irani, IJCV 2002]

Bundle adjustment

Structure from motion

$$\begin{aligned}\hat{u}_{ij} &= f(\mathbf{K}, \mathbf{R}_j, \mathbf{t}_j, \mathbf{x}_i) \\ \hat{v}_{ij} &= g(\mathbf{K}, \mathbf{R}_j, \mathbf{t}_j, \mathbf{x}_i)\end{aligned}$$

- How many points do we need to match?
- 2 frames:
 (\mathbf{R}, \mathbf{t}) : 5 dof + $3n$ point locations \leq
 $4n$ point measurements \Rightarrow
 $n \geq 5$
- k frames:
 $6(k-1) - 1 + 3n \leq 2kn$
- always want to use many more

Bundle Adjustment

$$\begin{aligned}\hat{u}_{ij} &= f(\mathbf{K}, \mathbf{R}_j, \mathbf{t}_j, \mathbf{x}_i) \\ \hat{v}_{ij} &= g(\mathbf{K}, \mathbf{R}_j, \mathbf{t}_j, \mathbf{x}_i)\end{aligned}$$

- What makes this non-linear minimization hard?
 - many more parameters: potentially slow
 - poorer conditioning (high correlation)
 - potentially lots of outliers



Lots of parameters: sparsity

$$\hat{u}_{ij} = f(\mathbf{K}, \mathbf{R}_j, \mathbf{t}_j, \mathbf{x}_i)$$

$$\hat{v}_{ij} = g(\mathbf{K}, \mathbf{R}_j, \mathbf{t}_j, \mathbf{x}_i)$$


- Only a few entries in Jacobian are non-zero

$$\frac{\partial \hat{u}_{ij}}{\partial \mathbf{K}}, \quad \frac{\partial \hat{u}_{ij}}{\partial \mathbf{R}_j}, \quad \frac{\partial \hat{u}_{ij}}{\partial \mathbf{t}_j}, \quad \frac{\partial \hat{u}_{ij}}{\partial \mathbf{x}_i},$$



Robust error models

- Outlier rejection
 - use robust penalty applied to each set of joint measurements




- for extremely bad data, use random sampling [RANSAC, Fischler & Bolles, CACM'81]

$$\sum_i \sigma_i^{-2} \rho \left(\sqrt{(u_i - \hat{u}_i)^2 + (v_i - \hat{v}_i)^2} \right)$$

Structure from motion: limitations


- Very difficult to reliably estimate metric structure and motion unless:
 - large (x or y) rotation *or*
 - large field of view and depth variation
- Camera calibration important for Euclidean reconstructions
- Need good feature tracker
- Lens distortion



Issues in SFM


- Track lifetime
- Nonlinear lens distortion
- Prior knowledge and scene constraints
- Multiple motions

Track lifetime




every 50th frame of a 800-frame sequence

Track lifetime



lifetime of 3192 tracks from the previous sequence

Track lifetime





track length histogram

Nonlinear lens distortion



Nonlinear lens distortion




effect of lens distortion

Prior knowledge and scene constraints



add a constraint that several lines are parallel

Prior knowledge and scene constraints



add a constraint that it is a turntable sequence

Applications of Structure from Motion

Jurassic park



PhotoSynth



<http://labs.live.com/photosynth/>