General Case: Projective Reconstruction

Forsyth&Ponce: Chap. 13

Szeliski: Chap. 7

General Projection Model

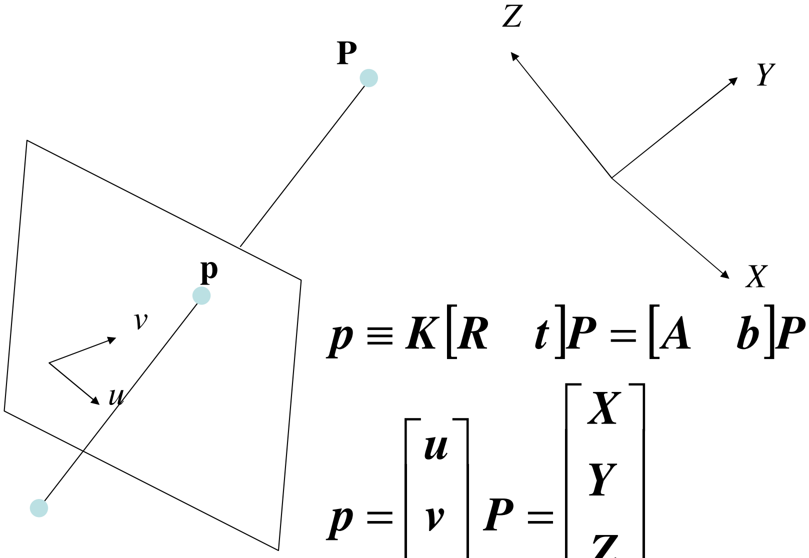

$$\boldsymbol{p} \equiv \boldsymbol{K}\left[\boldsymbol{R} \quad \boldsymbol{t}\right]\boldsymbol{P} = \left[\boldsymbol{A} \quad \boldsymbol{b}\right]\boldsymbol{P}$$

$$\boldsymbol{p} = \begin{bmatrix} \boldsymbol{u} \\ \boldsymbol{v} \\ 1 \end{bmatrix} \boldsymbol{P} = \begin{bmatrix} \boldsymbol{X} \\ \boldsymbol{Y} \\ \boldsymbol{Z} \\ 1 \end{bmatrix}$$

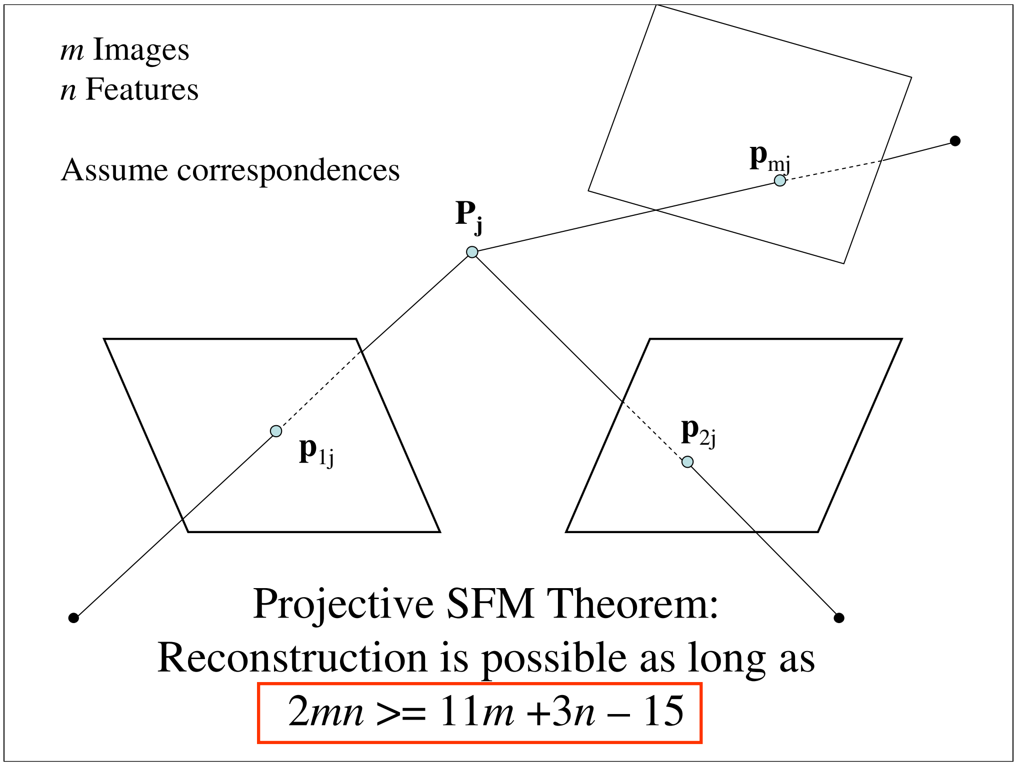

m Images
n Features

Assume correspondences

Projective SFM Theorem:
Reconstruction is possible as long as

$$\boxed{2mn >= 11m + 3n - 15}$$

Projective SFM Theorem

Given *m* images and *n* features

Each point is represented by its homogeneous coordinates $\mathbf{P}_j = [X\ Y\ Z\ 1]^T$

Each feature is represented by its homogeneous coordinates in the image plane $\mathbf{p}_{ij} = [u_{ij}\ v_{ij}\ 1]^T$

Each image is represented by its 3x4 projection matrix $\mathbf{M}_i = [\mathbf{A}_i\ \mathbf{b}_i]$

($\mathbf{A}_i$ is a 3x3 matrix, $\mathbf{b}_i$ is a 3 vector)

For each feature, we have: $\mathbf{p}_{ij} \sim \mathbf{M}_i\mathbf{P}_j$ (~ is the same as the triple-bar homogeneous equality: left hand side proportional to right hand side).

Key results:

1. The unknowns $\mathbf{M}_i$ and $\mathbf{P}_j$ can be recovered only up to a 4x4 projective transformation $\mathbf{Q}$. That is, for any 4x4 $\mathbf{Q}$, $(\mathbf{M}_i\mathbf{Q},\mathbf{Q}^{-1}\mathbf{P}_j)$ yields the same image projections as $(\mathbf{M}_i,\mathbf{P}_j)$.
2. The unknowns $\mathbf{M}_i$ and $\mathbf{P}_j$ can be recovered if $2mn >= 11m + 3n - 15$
3. In particular, for *m*=2 cameras, *at least 7 points* are needed.

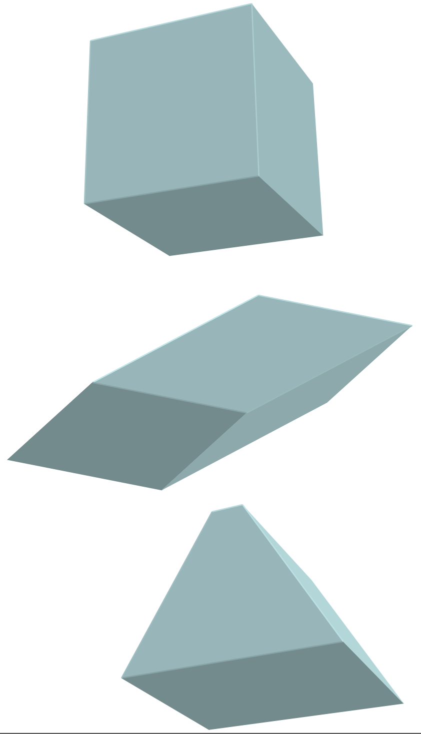

Metric

Affine

Projective

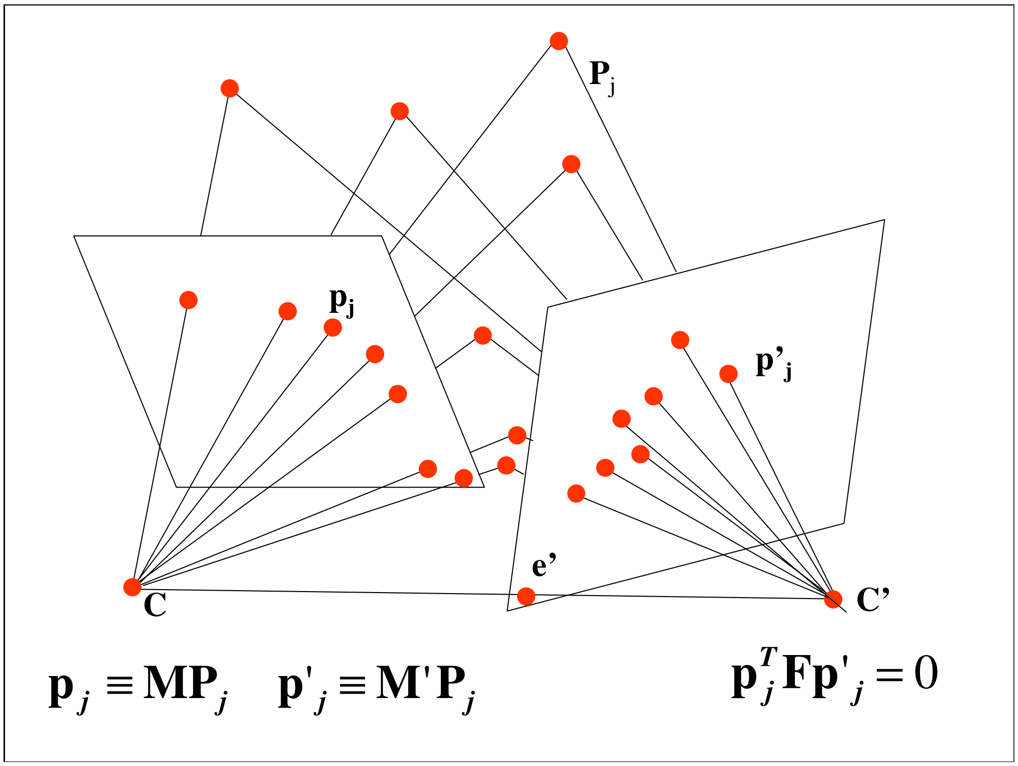

2-Camera Reconstruction

Consider the first the case of 2 cameras. The reconstruction problem is:

- Given correspondences ($\mathbf{p}_j$, $\mathbf{p}'_j$) j = 1..*n*
- Recover **M**, **M’**, $\mathbf{P}_j$

We know that the reconstruction is known only up to a projective transformation **Q** (4x4 matrix). In particular, we can choose **Q** such that: **M’Q** = [**Id 0**]. That is, we can always assume that the axis and camera parameters are such that **M’** = [**Id 0**] without loss of generality. In particular, this implies that the projection equations in that reference image are $u' = x'/z'$ and $v' = y'/z'$.

Assuming the canonical projection **M’** = [**Id 0**] in the reference image, we want to find **M** = [**A b**] that is consistent with the correspondences.

Key result:

1. If **M** is given, then the fundamental matrix is $\mathbf{F} = [\mathbf{b}_x]\mathbf{A}$
2. Conversely, if **F** is given (computed from the correspondences), then a solution to the reconstruction problem is: $\mathbf{A}_o = -[\mathbf{b}_x]\mathbf{F}$

   where **b** is the epipole: $\mathbf{F}^T\mathbf{b} = 0$
3. All the solutions to the reconstruction problem are of the form:

   $\mathbf{A} = \lambda\mathbf{A}_o + [\mu_1\mathbf{b}\ \mu_2\mathbf{b}\ \mu_3\mathbf{b}]$ and are equivalent

Note the number of parameters **A**,**b** = 11 param. = 7 for **F** + 4 free parameters

Note the analogy between $\mathbf{F} = [\mathbf{b}_x]\mathbf{A}$ in the general case and $\mathbf{E} = [\mathbf{t}_x]\mathbf{R}$ in the calibrated case.

$$\mathbf{p}_j \equiv \mathbf{M}\mathbf{P}_j \qquad \mathbf{p}'_j \equiv \mathbf{M}'\mathbf{P}_j$$

$$\mathbf{p}_j^T \mathbf{F} \mathbf{p}'_j = 0$$

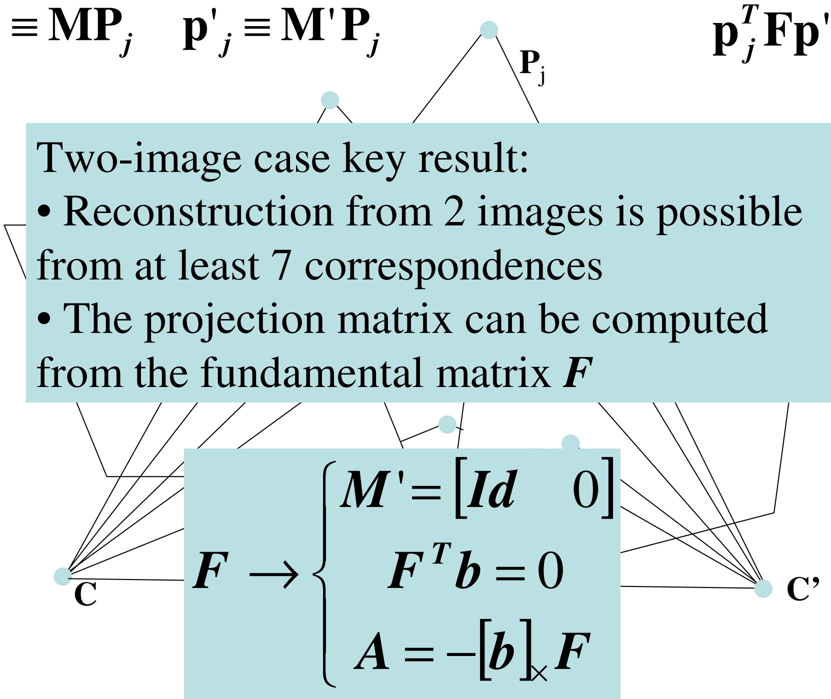

Two-image case key result:

- Reconstruction from 2 images is possible from at least 7 correspondences
- The projection matrix can be computed from the fundamental matrix ***F***

$$\boldsymbol{F} \rightarrow \begin{cases} \boldsymbol{M}' = \begin{bmatrix} \boldsymbol{Id} & 0 \end{bmatrix} \\ \boldsymbol{F}^T \boldsymbol{b} = 0 \\ \boldsymbol{A} = -[\boldsymbol{b}]_\times \boldsymbol{F} \end{cases}$$

Issues

- We know how to compute the camera geometry from 2 images*
- Remaining issues:
  1. What about $m > 2$ images?
  2. Triangulation to compute point in 3D
  3. Projective reconstruction is not unique: How can we find the metrically correct reconstruction?
  4. How can we find the "optimal" reconstruction?

*: (or 3, not shown in class)

Issue: Arbitrary number of images

- Factorization (e.g., assuming affine first)
- Sequential methods
- Hierarchical methods
- …..

Sequential Methods

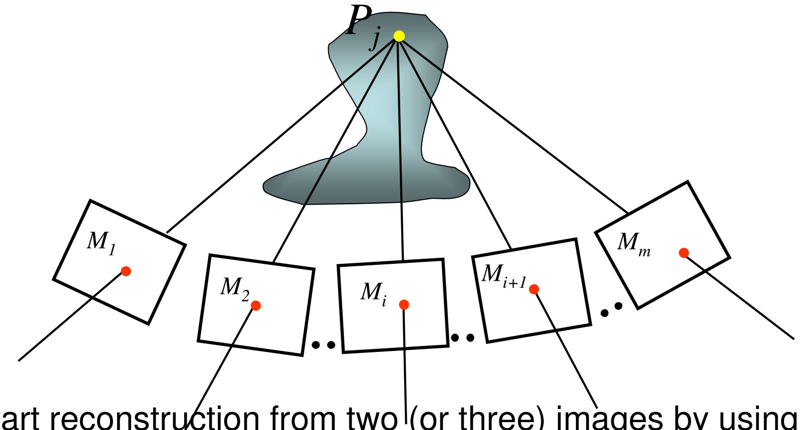

- Start reconstruction from two (or three) images by using the fundamental matrix (or the trifocal tensor).
- Given projection matrix M_i of image i, find the projection matrix of image i+1 such that:
- $M_i P_j = M_{i+1} P_j$ for the points P_j in common between the two images
- M_{i+1} can be determined from 6 points → RANSAC if correspondences not known

Hierarchical Methods

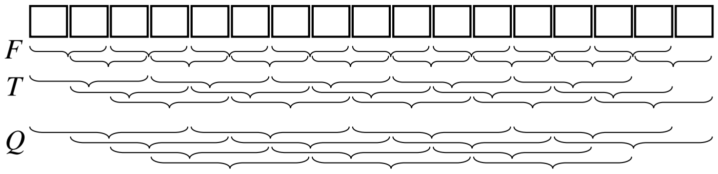

Example from Pollefeys

- Compute *F*,*T* from groups of 2 or 3 images
- Iteratively “stitch” the images into larger groups
- Stitching = Find a matrix *Q* that aligns the reconstruction from one group with the reconstruction from another group → Need some features in common between the groups

Stitching

P_j=coordinates of point j for first group
P'_j=coordinates of point j for second group

Find Q that transforms the point in the second group to the points in the first group

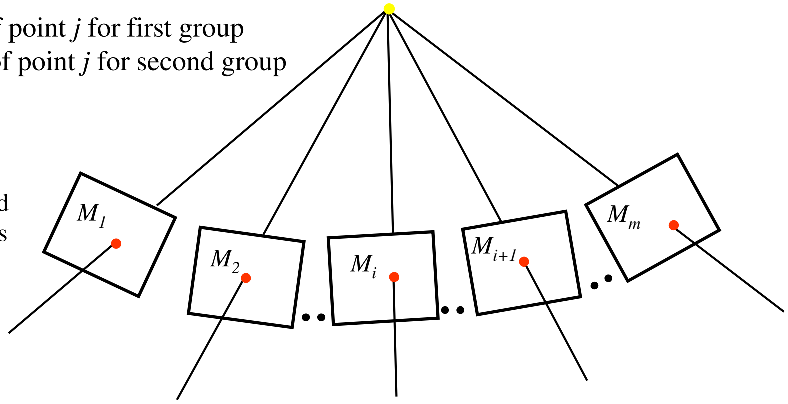

$$\min \sum d(P_j, QP'_j)$$

Distance in space

$$\min \sum d(M_i, M'_i Q^{-1})$$

Distance in camera geometry

$$\min \sum d\left(M_i Q P'_j, p_{ij}\right) + \sum d\left(M'_i Q^{-1} P_j, p_{ij}\right)$$

Reprojection error

Correspondences

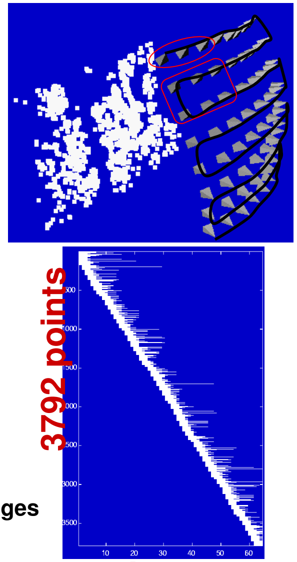

3792 points

64 images

4.8im/pt

3792 points

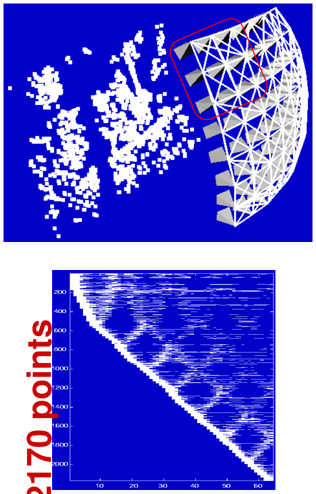

2170 points

9.8im/pt

- Problem: Need enough correspondences in common between views

Issues

- We know how to compute the camera geometry from 2 images[*]
- Remaining issues:
  1. What about $m > 2$ images?
  2. Triangulation to compute point in 3D
  3. Projective reconstruction is not unique: How can we find the metrically correct reconstruction?
  4. How can we find the "optimal" reconstruction?

*: (or 3, not shown in class)

Issue: Triangulation

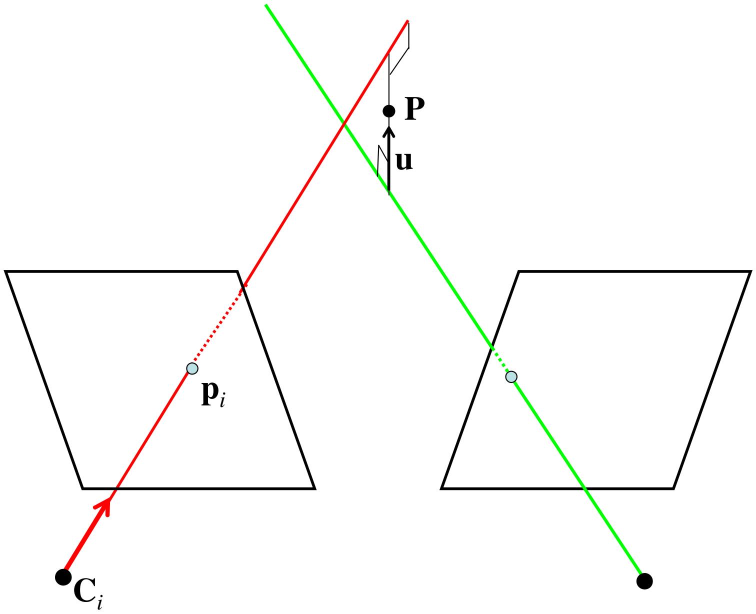

The previous discussion assumes that the viewing rays from the cameras intersect exactly. In fact, that is not usually the case because of small errors in estimating the camera geometry. In that case the two viewing rays pass close to each other but do not exactly intersect. The point **P** is reconstructed as the point that is the closest to both lines. Broadly speaking, there are 3 ways to address the problem. The point **P** has coordinates *X*, *Y*, *Z*, the corresponding feature in image *i* is $\mathbf{p}_i$ of coordinates u_i and v_i, and the corresponding ray has direction $\mathbf{W}_i$.

1. *Point closest to the rays in space (linear)*

We want to find **P** that minimizes the sum of the distances between **P** and the rays ($\mathbf{C}_i$,). Can be written as a linear least-squares problem:

$$Min \sum_i \left\| C_i + s_i W_i - \begin{bmatrix} X \\ Y \\ Z \end{bmatrix} \right\|^2$$

2. *Point with lowest algebraic reprojection error (linear)*

We can write the fact that $\mathbf{p}_i$ is the projection of **P** as $p_i \equiv M_i P$ or equivalently $p_i \times M_i P = 0$

Thus we can find **P** by minimizing the algebraic error: $Min \sum_i \|p_i \times M_i P\|$

3. *Point with lowest geometric reprojection error (non-linear)*

The algebraic error is only an approximation, by replacing $u_i = \frac{m_1^T P}{m_3^T P}$ by $u_i m_3^T P - m_1^T P = 0$

The correct error is the geometric error:

$$\sum_i \left(u_i - \frac{m_1^T P}{m_3^T P} \right)^2 + \left(v_i - \frac{m_2^T P}{m_3^T P} \right)^2$$

This leads to a non-linear optimization → Bundle adjustment presented later.

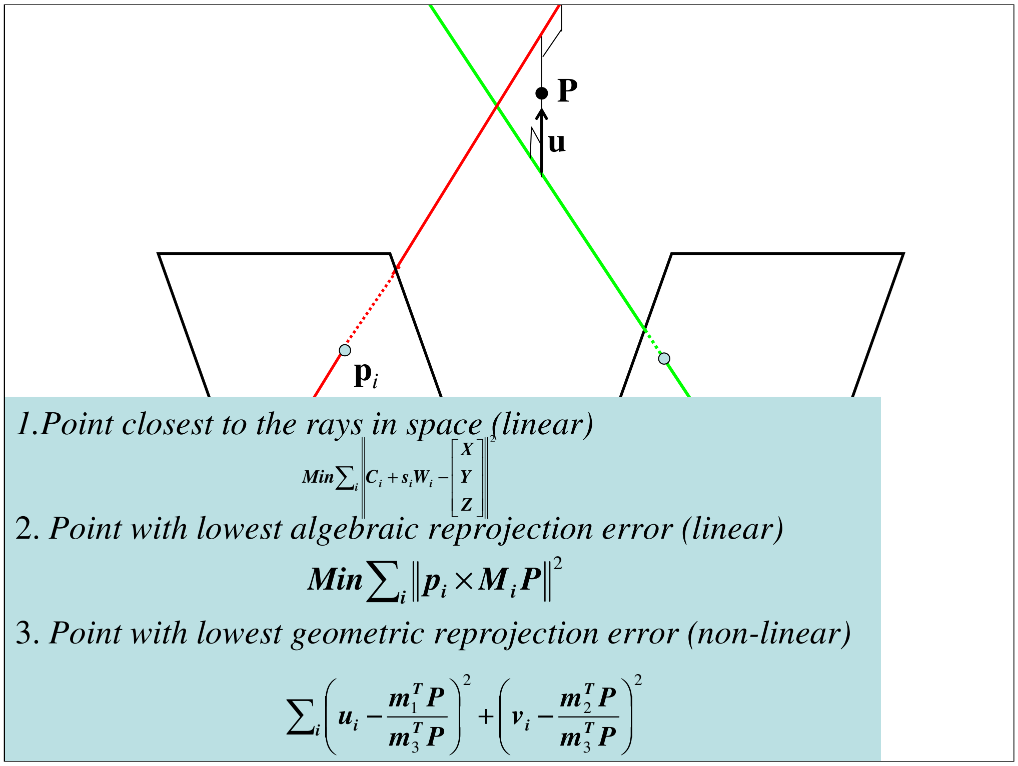

1.Point closest to the rays in space (linear)

$$Min\sum_i \left\| C_i + s_i W_i - \begin{bmatrix}X\\Y\\Z\end{bmatrix}\right\|^2$$

2. Point with lowest algebraic reprojection error (linear)

$$Min\sum_i \left\| p_i \times M_i P\right\|^2$$

3. Point with lowest geometric reprojection error (non-linear)

$$\sum_i \left(u_i - \frac{m_1^T P}{m_3^T P}\right)^2 + \left(v_i - \frac{m_2^T P}{m_3^T P}\right)^2$$

The previous discussion assumes that the viewing rays from the cameras intersect exactly. In fact, that is not usually the case because of small errors in estimating the camera geometry. In that case the two viewing rays pass close to each other but do not exactly intersect. The point **P** is reconstructed as the point that is the closest to both lines. Broadly speaking, there are 3 ways to address the problem. The point **P** has coordinates *X*, *Y*, *Z*, the corresponding feature in image *i* is $\mathbf{p}_i$ of coordinates u_i and v_i, and the corresponding ray has direction $\mathbf{W}_i$.

1. *Point closest to the rays in space (linear)*

We want to find **P** that minimizes the sum of the distances between **P** and the rays ($\mathbf{C}_i$,). Can be written as a linear least-squares problem:

$$Min\sum_i \left\| C_i + s_i W_i - \begin{bmatrix}X\\Y\\Z\end{bmatrix}\right\|^2$$

2. *Point with lowest algebraic reprojection error (linear)*

We can write the fact that $\mathbf{p}_i$ is the projection of **P** as $p_i \equiv M_i P$ or equivalently $p_i \times M_i P = 0$

Thus we can find **P** by minimizing the algebraic error: $Min\sum_i \left\| p_i \times M_i P\right\|$

3. *Point with lowest geometric reprojection error (non-linear)*

The algebraic error is only an approximation, by replacing $u_i = \frac{m_1^T P}{m_3^T P}$ by $u_i m_3^T P - m_1^T P = 0$

The correct error is the geometric error:

$$\sum_i \left(u_i - \frac{m_1^T P}{m_3^T P}\right)^2 + \left(v_i - \frac{m_2^T P}{m_3^T P}\right)^2$$

This leads to a non-linear optimization → Bundle adjustment presented later.

Issues

- We know how to compute the camera geometry from 2 images*
- Remaining issues:
  1. What about $m > 2$ images?
  2. Triangulation to compute point in 3D
  3. Projective reconstruction is not unique: How can we find the metrically correct reconstruction?
  4. How can we find the "optimal" reconstruction?

*: (or 3, not shown in class)

From Projective to Metric Reconstruction
Auto Calibration

Forsyth&Ponce: Chap. 13

Szeliski: Chap. 7

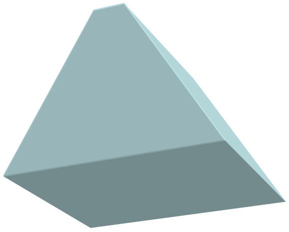

Projective

$$\mathbf{M}_i \leftarrow \mathbf{M}_i\mathbf{Q}$$

$$\mathbf{P}_j \leftarrow \mathbf{Q}^{-1}\mathbf{P}_j$$

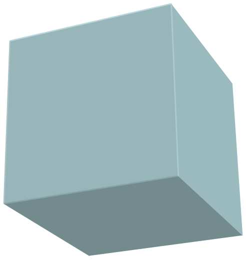

Metric

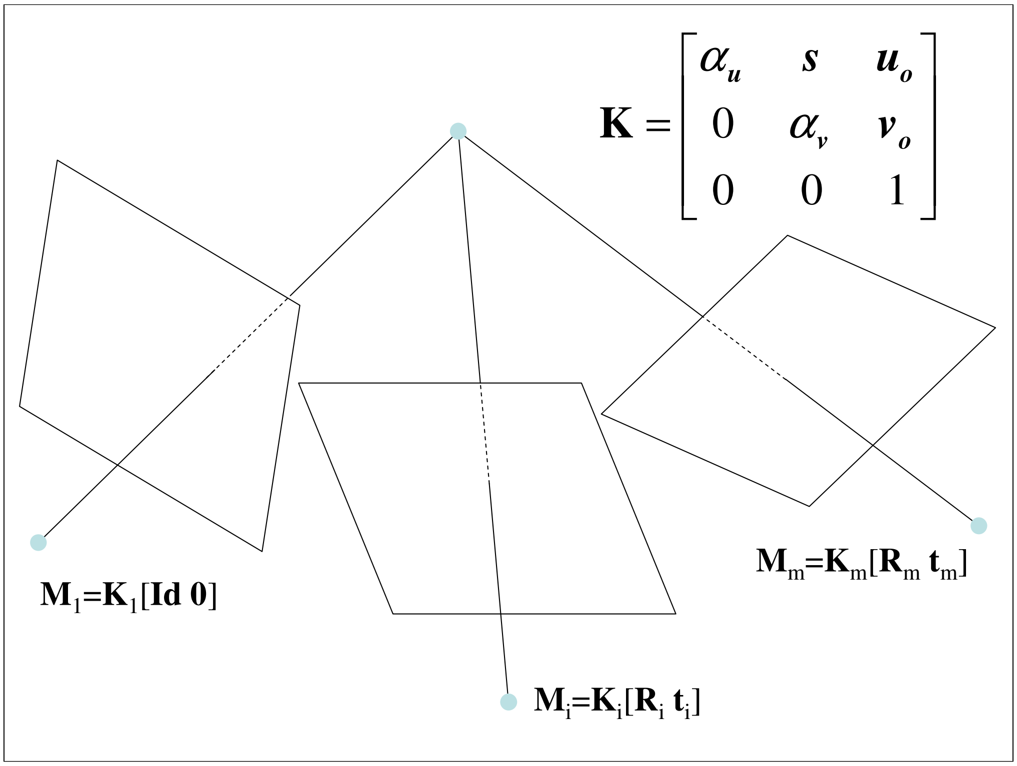

Metric Upgrade:

The projective reconstruction gives us a set of 3x4 projection matrices $\mathbf{M}_i$ for each camera i=1,..,m. The next problem is to convert this projective reconstruction to a metric reconstruction. Specifically, we want to find a 4x4 matrix **Q** such that:

$\mathbf{M}_i\mathbf{Q} \equiv \mathbf{K}_i[\mathbf{R}_i|\mathbf{t}_i]$

$\mathbf{R}_i$ and $\mathbf{t}_i$ are the rotation/translation between the coordinate system of camera i and an arbitrary coordinate system.

$\mathbf{K}_i$ is the matrix of intrinsic parameters of camera i, which is defined as:

$$\mathbf{K} = \begin{bmatrix} \alpha_u & s & u_o \\ 0 & \alpha_v & v_o \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} \alpha & -\alpha\cot\theta & u_o \\ 0 & \dfrac{\beta}{\sin\theta} & v_o \\ 0 & 0 & 1 \end{bmatrix}$$

α_x and α_y are the scales in the x and y directions, x_o and y_o are the coordinates of the center, and s is the skew of the camera (s = 0 if the axes are orthogonal.)

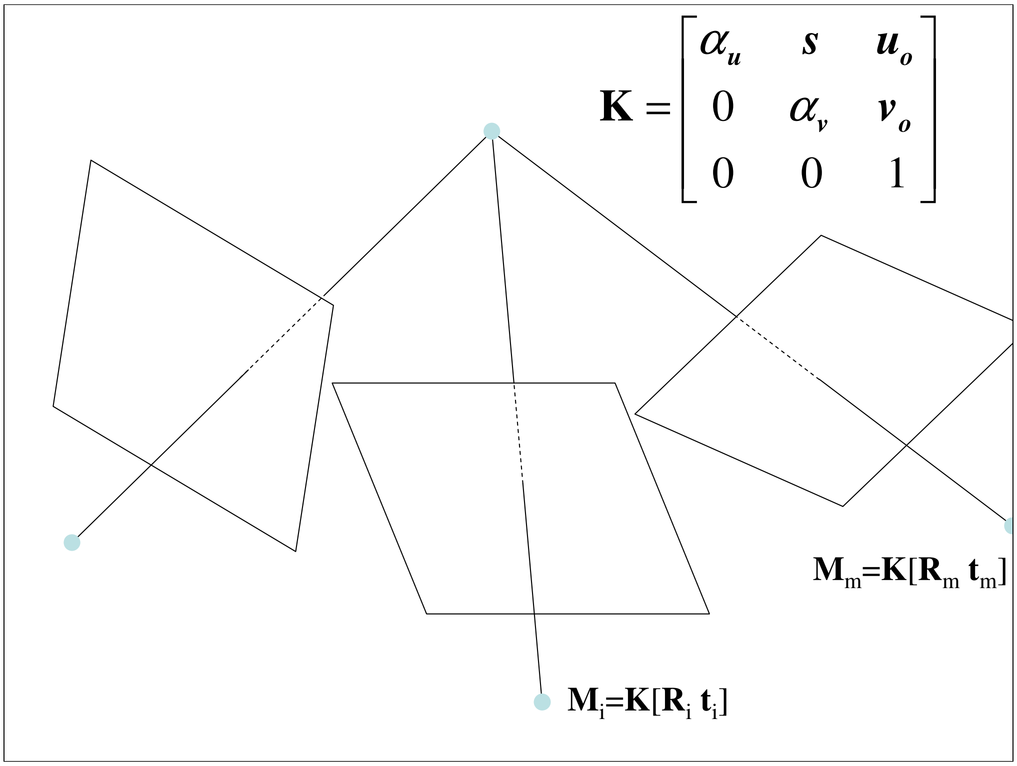

Fundamental Transformation:

Our fundamental equation is:

$$\mathbf{M}_i\mathbf{Q} \equiv \mathbf{K}[\mathbf{R}_i\ \mathbf{t}_i]$$

Denoting the matrix formed by taking the first 3 columns of **Q** by $\mathbf{Q}_3$, such that $\mathbf{Q} = [\mathbf{Q}_3\ \mathbf{q}_4]$, we have: $\mathbf{M}_i\mathbf{Q}_3 \equiv \mathbf{K}_i\mathbf{R}_i$.

Taking the first 3 columns and observing that $\mathbf{R}_i$ is a rotation matrix:

$$\mathbf{M}_i\mathbf{Q}_3\mathbf{Q}_3^T\mathbf{M}_i^T \equiv \mathbf{K}_i\mathbf{K}_i^T$$

This is the key observation: By writing that the first three columns of the product of **M** by **Q** is a rotation, we are able to eliminate the rotation from the unknowns. All that is left are the matrices of internal parameters for each of the images. This process is sometimes called auto-calibration since it amounts to calibrating the internal parameters of the cameras directly from images.

It is important to understand the number of degrees of freedom in $\mathbf{Q}_3$. The total number of entries in $\mathbf{Q}_3$ is 4x3=12. The matrix is defined up to scale since all the equalities are homogeneous. Moreover, the matrix is defined up to a rotation since for any arbitrary rotation R: $\mathbf{Q}_3\mathbf{R}\mathbf{R}^T\mathbf{Q}_3^T = \mathbf{Q}_3\mathbf{Q}_3^T$ so that if $\mathbf{Q}_3$ is a solution, so is $\mathbf{Q}_3\mathbf{R}$**.** These additional degrees of freedom simply reflect the fact that one can choose the orientation and scale of the global coordinate system arbitrarily. Therefore, $\mathbf{Q}_3$ is characterized by 12 – 1 – 3 = 8 unknowns.

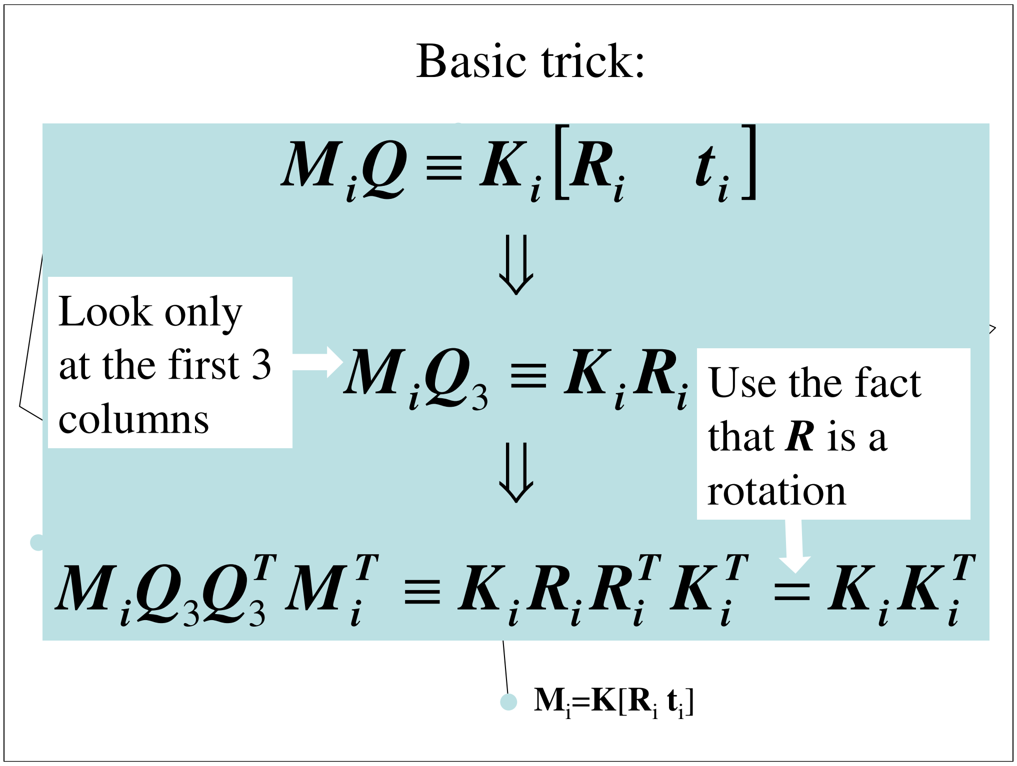

Basic trick:

$$M_i Q \equiv K_i \left[R_i \quad t_i \right]$$

$$\Downarrow$$

Look only at the first 3 columns

$$M_i Q_3 \equiv K_i R_i$$

Use the fact that *R* is a rotation

$$\Downarrow$$

$$M_i Q_3 Q_3^T M_i^T \equiv K_i R_i R_i^T K_i^T = K_i K_i^T$$

$\mathbf{M}_i = \mathbf{K}[\mathbf{R}_i \ \mathbf{t}_i]$

For convenience, we denote the matrix the matrix $\mathbf{Q}_3\mathbf{Q}_3^T$ by $\mathbf{L}$ (a 4x4 matrix) and $\mathbf{M}_i\mathbf{L}\mathbf{M}_i^T$ by $\boldsymbol{\omega}_i$. The set of equations to solve is:

$$\mathbf{M}_i\mathbf{L}\mathbf{M}_i^T \equiv \mathbf{K}_i\mathbf{K}_i^T \quad i=1,..,m$$

Each image generates 5 independent equations (the left hand side is a 3x3 symmetric matrix, but the equality is up to scale). The total number of unknowns is 8 ($\mathbf{Q}_3$) + 5m ($\mathbf{K}_i\mathbf{K}_i^T$). Therefore, the number of equations (5m) is *always* lower than the number of unknowns (8 + 5m) and we can never solve this system of equations without some constraints on the cameras. The key question is what constraints can be used. A couple of constraints are investigated below, followed with a general result.

Case 1: Identical Intrinsic Parameters:

Let us suppose now that we do not know the intrinsic parameters of the cameras, but that we do know that they are all identical, that is, $\mathbf{K}_i = \mathbf{K}$ for all cameras i. For all the cameras, we have:

$$\mathbf{M}_i\mathbf{L}\mathbf{M}_i^T \equiv \boldsymbol{\omega} \quad i=1,..,m \quad \text{(with the equality up to scale)}$$

Where $\boldsymbol{\omega}$ is computed from the first image:

$$\mathbf{M}_i\mathbf{L}\mathbf{M}_i^T \equiv \mathbf{M}_1\mathbf{L}\mathbf{M}_1^T \quad i=2,..,m$$

This gives us 5(m-1) independent equations for 8 unknowns (in $\mathbf{Q}_3$). Therefore, we can solve the reconstruction problem in this case if:

$$5(m-1) >= 8 \ \text{-}\!\!\rightarrow \ \boxed{m >= 3}$$

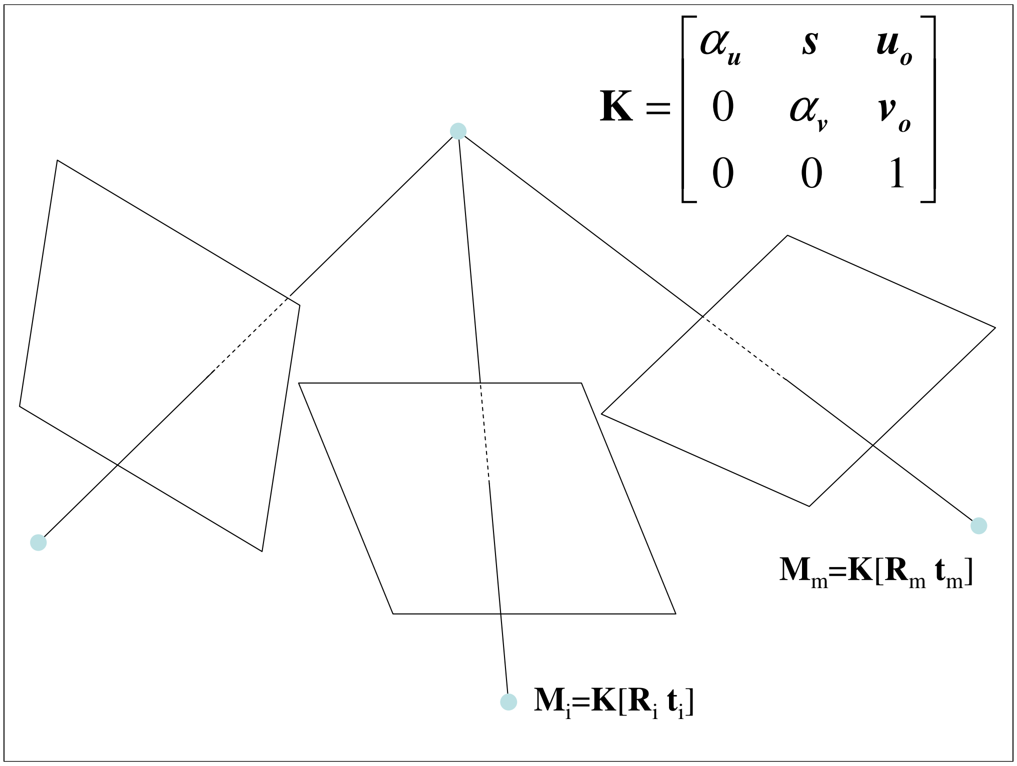

Note also that we solved for 8 unknowns for $\mathbf{Q}_3$, even though it is a 4x3 matrix → 12 elements. The difference is due to the fact that the reconstruction is defined up to a global rotation, in other words, $\mathbf{Q}_3\mathbf{R}$ for any rotation $\mathbf{R}$ is also a solution.

We can verify the parameter count:

$$\underbrace{8}_{\mathbf{Q}_3\mathbf{Q}_3^T} + \underbrace{3}_{\text{Arbitrary rotation}} + \underbrace{4}_{\mathbf{q}_4} + \underbrace{1}_{\text{Arbitrary scale}} = 16$$

It is important to note that by manipulating the equations so that the unknown becomes **L**, we have effectively eliminated the rotation and translation and reduced the problem to the recovery of the intrinsic parameters. This step is often termed *self-calibration*.

$$\mathbf{K} = \begin{bmatrix} \alpha_u & s & u_o \\ 0 & \alpha_v & v_o \\ 0 & 0 & 1 \end{bmatrix}$$

$$\downarrow$$

$$\boldsymbol{\omega} = \mathbf{K}\mathbf{K}^T = \begin{bmatrix} \alpha_u^2 + s^2 + u_o^2 & s\alpha_v + u_o v_o & u_o \\ & \alpha_v^2 + v_o^2 & v_o \\ & & 1 \end{bmatrix}$$

Case 2: Principal point at origin

In that case, $u_o = v_o = 0$, which implies that $\omega_{13} = \omega_{23} = 0$. Therefore, going back to the original equation, we can write that:

$(\mathbf{M}_i\mathbf{L}\mathbf{M}_i^T)_{13} = 0$ and $(\mathbf{M}_i\mathbf{L}\mathbf{M}_i^T)_{23} = 0$

Those are two equations in **L** that are independent of $\mathbf{K}_i$. We have $2m$ such equations for m views for 8 unknowns in **L** (meaning, 8 unknowns in $\mathbf{Q}_3$). These equations are independent of **K**. Therefore:

If the principal point is at the origin, a metric reconstruction can be obtained from a minimum of 4 views.

$m >= 4$

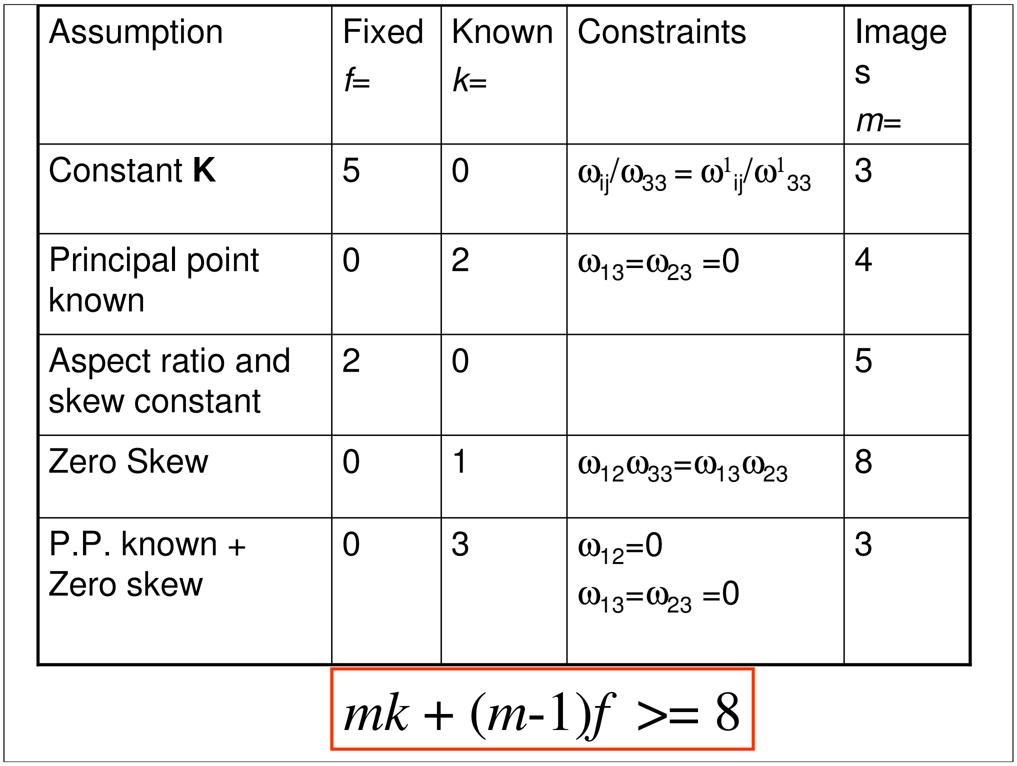

Assumption	Fixed f=	Known k=	Constraints	Images m=
Constant **K**	5	0	$\omega_{ij}/\omega_{33} = \omega^1_{ij}/\omega^1_{33}$	3
Principal point known	0	2	$\omega_{13}=\omega_{23}=0$	4
Aspect ratio and skew constant	2	0		5
Zero Skew	0	1	$\omega_{12}\omega_{33}=\omega_{13}\omega_{23}$	8
P.P. known + Zero skew	0	3	$\omega_{12}=0$ $\omega_{13}=\omega_{23}=0$	3

$$\boxed{mk + (m-1)f >= 8}$$

Case 3: Zero-Skew

If the skew is zero but all the other parameters are allowed to vary, then we have the constraint:

$\omega_{12}\omega_{33} = \omega_{13}\,\omega_{23}$

This provides m constraints. Therefore, we must have m>=8, thus 8 images are necessary.

General counting argument:

The total number of parameters to be estimated from the set of equations above is 8: $\mathbf{L} = \mathbf{Q}_3\mathbf{Q}_3^T$ is a 4x4 symmetric matrix that is defined by 10 entries, 9 of which are independent because of the scale factor. There is one more constraint that det(**L**) = 0 (it is of rank 3). Thus the number of independent parameters is 10-1-1=8.

-If we know k internal parameters, then we have km constraints

- If we know that f internal parameters are fixed (but unknown) we have $f(m-1)$ constraints

- Therefore, we can recover a metric reconstruction iff:

$$\boldsymbol{mk + (m-1)f >= 8}$$

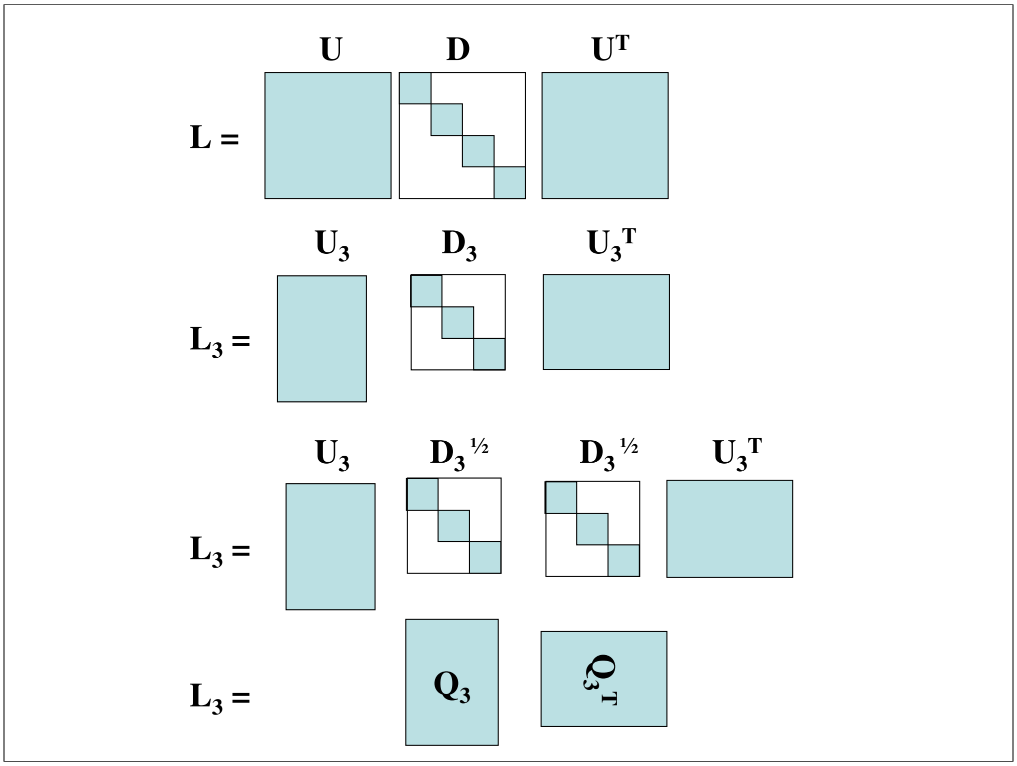

Recovering Q:

One problem in the previous result is that, while we used the matrix L for convenience, the actual unknown is $\mathbf{Q}_3$ with $\mathbf{L} = \mathbf{Q}_3\mathbf{Q}_3^T$, which is a non-linear relation. This can be inconvenient because we have to solve a set of non-linear equations. In fact, it is possible in many cases to solve the problem *linearly*. This is done by using the same trick as before with **F** and **E**, and with the factorization method: Let's pretend first that we solve the equations in **L**, which has 9 degrees of freedom (4x4 symmetric matrix = 10, but it is up to scale), and then decompose **L** into $\mathbf{L} = \mathbf{Q}_3\mathbf{Q}_3^T$.

For example, in case 2 (Principal point at origin), $u_o = v_o = 0$, which implies that $\omega_{13} = \omega_{23} = 0$. Therefore:

$(\mathbf{M}_i\mathbf{L}\mathbf{M}_i^T)_{13} = 0$ and $(\mathbf{M}_i\mathbf{L}\mathbf{M}_i^T)_{23} = 0$

Those are two equations that are linear in **L** and can be solved very easily!!

Once we have **L**, we need to decompose it back into $\mathbf{L} = \mathbf{Q}_3\mathbf{Q}_3^T$. There is a practical difficulty here: Since $\mathbf{Q}_3$ is a 4x3 matrix, for such a decomposition to exist, **L** must be of rank at most 3 which is not enforced in the linear solution.

We can find the closest matrix $\mathbf{L}_3$ of rank three as follows:

L is a symmetric matrix so $\mathbf{L} = \mathbf{U}\mathbf{D}\mathbf{U}^T$, where **D** is a 4x4 diagonal matrix of eigenvalues and **U** is a 4x4 rotation matrix. The matrix $\mathbf{L}_3$ of rank 3 that is closest to **L** can be formed by eliminating the smallest eigenvalue of **L**, that is, $\mathbf{L}_3 = \mathbf{U}_3\mathbf{D}_3\mathbf{U}_3^T$, where $\mathbf{U}_3$ is the 4x3 matrix obtained by removing the last column of **Q** and $\mathbf{D}_3$ is the 3x3 upper right block of **D**. With this decomposition, the matrix $\mathbf{Q}_3 = \mathbf{U}_3\mathbf{D}_3^{1/2}$ is a solution since $\mathbf{L}_3 = \mathbf{Q}_3\mathbf{Q}_3^T$.

K can be recovered by Cholesky decomposition $\boldsymbol{\omega} = \mathbf{K}\mathbf{K}^T$.

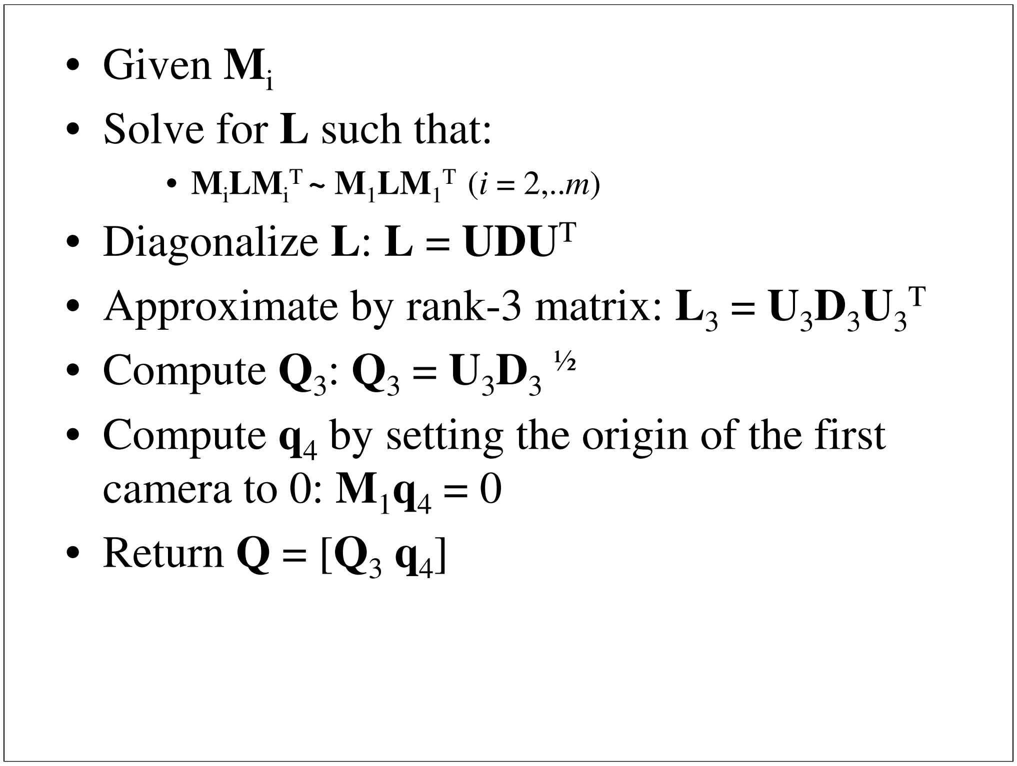

- Given $\mathbf{M}_i$
- Solve for $\mathbf{L}$ such that:
  - $\mathbf{M}_i\mathbf{L}\mathbf{M}_i^T \sim \mathbf{M}_1\mathbf{L}\mathbf{M}_1^T$ ($i = 2,..m$)
- Diagonalize $\mathbf{L}$: $\mathbf{L} = \mathbf{U}\mathbf{D}\mathbf{U}^T$
- Approximate by rank-3 matrix: $\mathbf{L}_3 = \mathbf{U}_3\mathbf{D}_3\mathbf{U}_3^T$
- Compute $\mathbf{Q}_3$: $\mathbf{Q}_3 = \mathbf{U}_3\mathbf{D}_3^{\frac{1}{2}}$
- Compute $\mathbf{q}_4$ by setting the origin of the first camera to 0: $\mathbf{M}_1\mathbf{q}_4 = 0$
- Return $\mathbf{Q} = [\mathbf{Q}_3\ \mathbf{q}_4]$

Given $\mathbf{Q}_3$, the last column of $\mathbf{Q}$ is computed by setting the origin at the origin of the first camera, that is: $\mathbf{M}_1\mathbf{q}_4 = 0$. Note that any scaled version of $\mathbf{q}_4$ is a solution. This is a consequence of the fact that it is not possible to recover the absolute scale of the translation between the cameras.

Remember that we solved for 8 unknowns for $\mathbf{Q}_3$, even though it is a 4x3 matrix → 12 elements. The difference is due to the fact that the reconstruction is defined up to a global rotation, in other words, $\mathbf{Q}_3\mathbf{R}$ for any rotation $\mathbf{R}$ is also a solution.

We can verify the parameter count:

$$8 + 3 + 4 + 1 = 16$$

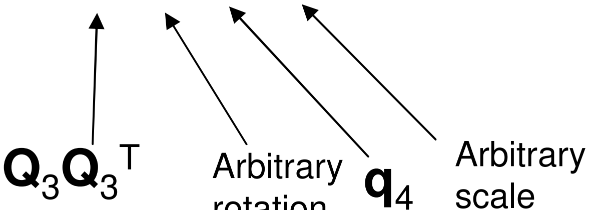

It is important to note that by manipulating the equations so that the unknown becomes $\mathbf{L}$, we have effectively eliminated the rotation and translation and reduced the problem to the recovery of the intrinsic parameters. This step is often termed *auto-calibration*.

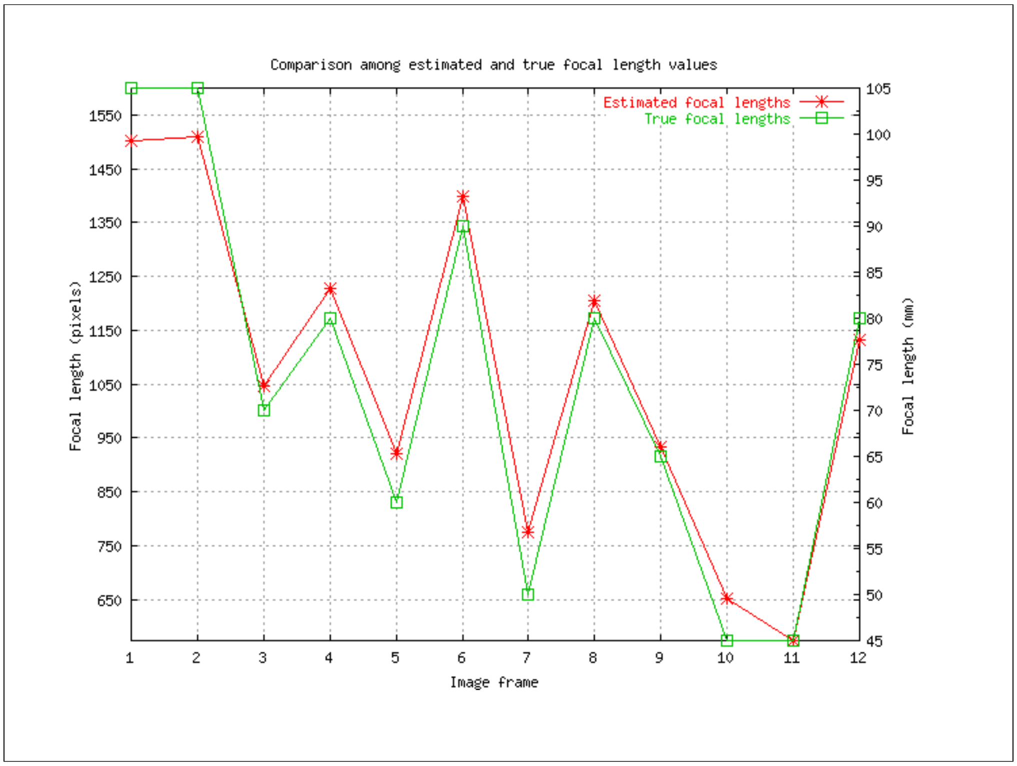
Comparison among estimated and true focal length values
Estimated focal lengths
True focal lengths
Focal length (pixels)
1550
1450
1350
1250
1150
1050
950
850
750
650
Focal length (mm)
105
100
95
90
85
80
75
70
65
60
55
50
45
1
2
3
4
5
6
7
8
9
10
11
12
Image frame

Issues

- We know how to compute the camera geometry from 2 images*
- Remaining issues:
  1. What about $m > 2$ images?
  2. Triangulation to compute point in 3D
  3. Projective reconstruction is not unique: How can we find the metrically correct reconstruction?
  4. How can we find the "optimal" reconstruction?

*: (or 3, not shown in class)