


Camera calibration




Odilon Redon, *Cyclops*, 1914

Overview


- Camera calibration
 - Intrinsic camera parameters
 - Extrinsic camera parameters
 - Estimation
- First taste of 3D reconstruction: triangulation

Perspective projection in normalized coordinates



- **Normalized (camera) coordinate system:** camera center is at the origin, the *principal axis* is the z -axis, x and y axes of the image plane are parallel to x and y axes of the world

Perspective projection in normalized coordinates



$$x = f \frac{X}{Z}, y = f \frac{Y}{Z}$$

$$\begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \cong \begin{pmatrix} fX \\ fY \\ Z \end{pmatrix} = \begin{bmatrix} f & 0 & 0 & 0 \\ 0 & f & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix}$$


x
Homogeneous
coord. vec. of
image point

P
Camera
projection
matrix

X
Homogeneous
coord. vec. of 3D
point

$x \cong PX$
↑
Equality up to scale

Camera calibration: Overview




- **Camera calibration:** figuring out transformation from *world* coordinate system to *image* coordinate system

$$\begin{pmatrix} \text{2D point} \\ (3 \times 1) \end{pmatrix} \cong \begin{pmatrix} \text{Camera to pixel coord.} \\ \text{trans. matrix} \\ \mathbf{K} (3 \times 3) \end{pmatrix} \begin{pmatrix} \text{Canonical} \\ \text{projection matrix} \\ \mathbf{[I | 0]} (3 \times 4) \end{pmatrix} \begin{pmatrix} \text{World to} \\ \text{camera coord.} \\ \text{trans. matrix} \\ \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} (4 \times 4) \end{pmatrix} \begin{pmatrix} \text{3D} \\ \text{point} \\ (4 \times 1) \end{pmatrix}$$

\mathbf{x} **Intrinsic camera parameters:** principal point, scaling factors **Extrinsic camera parameters:** rotation, translation \mathbf{X}


Camera calibration: Overview



- Camera calibration:** figuring out transformation from *world* coordinate system to *image* coordinate system


$$\begin{array}{c}
 \begin{pmatrix} \text{2D point} \\ \text{(3x1)} \end{pmatrix} \\
 \mathbf{x}
 \end{array}
 \cong
 \underbrace{
 \begin{pmatrix} \text{Camera to pixel coord. trans. matrix} \\ \mathbf{K} \text{ (3x3)} \end{pmatrix}
 \begin{pmatrix} \text{Canonical projection matrix} \\ \mathbf{[I | 0]} \text{ (3x4)} \end{pmatrix}
 \begin{pmatrix} \text{World to camera coord. trans. matrix} \\ \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} \text{ (4x4)} \end{pmatrix}
 }_{\mathbf{P} = \mathbf{K}[\mathbf{R}|\mathbf{t}]}
 \begin{pmatrix} \text{3D point} \\ \text{(4x1)} \end{pmatrix} \\
 \mathbf{X} \\
 \text{General camera projection matrix}
 \end{array}$$

Intrinsic parameters: Principal point



- **Principal point (p)**: point where principal axis intersects the image plane
- In the *normalized* coordinate system, the **origin** of the image is at the **principal point**
- In the *image* coordinate system: the **origin** is in the **corner**

Intrinsic parameters: Principal point




We want the principal point to map to (p_x, p_y) instead of $(0,0)$

$$x = f \frac{X}{Z} + p_x, \quad y = f \frac{Y}{Z} + p_y$$

$$\begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \cong \begin{pmatrix} fX + Zp_x \\ fY + Zp_y \\ Z \end{pmatrix} = \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & Z \end{bmatrix} \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix}$$

Intrinsic parameters: Principal point



Principal point: (p_x, p_y)

$$\begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} = \begin{bmatrix} f & 0 & p_x & 0 \\ 0 & f & p_y & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

calibration
matrix \mathbf{K}

Canonical
projection matrix
 $[\mathbf{I} \mid \mathbf{0}]$

$$\mathbf{P} = \mathbf{K}[\mathbf{I} \mid \mathbf{0}]$$

Intrinsic parameters: Principal point

- What are the units of the focal length f and principal point coordinates (p_x, p_y) ?
 - Same as world units – presumably metric units
- What units do we want for measuring image coordinates?
 - Pixel units
- Thus, we need to introduce *scaling factors* for mapping from world to pixel units


$$\begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} = \begin{bmatrix} f & 0 & p_x & 0 \\ 0 & f & p_y & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

calibration
matrix \mathbf{K}

Canonical
projection matrix
 $[\mathbf{I} \mid \mathbf{0}]$

$$\mathbf{P} = \mathbf{K}[\mathbf{I} \mid \mathbf{0}]$$

Intrinsic parameters: Scaling factors




m_x pixels/m in horizontal direction,
 m_y pixels/m in vertical direction

Pixel size (m): $\frac{1}{m_x} \times \frac{1}{m_y}$

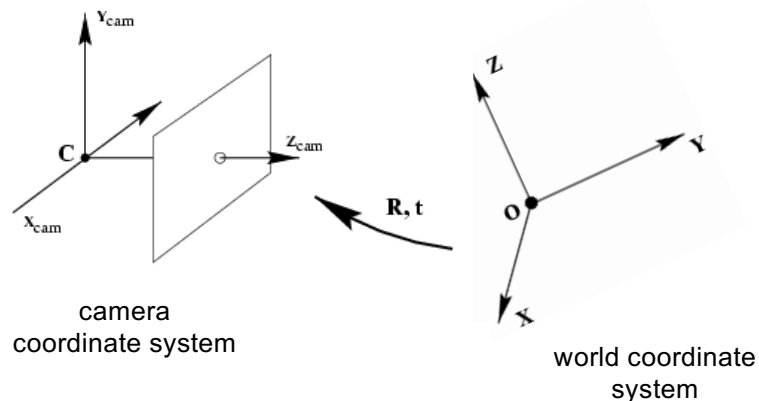
Scaling factors	Calibration matrix K in metric units	Calibration matrix K in pixel units
$\begin{bmatrix} m_x & 0 & 0 \\ 0 & m_y & 0 \\ 0 & 0 & 1 \end{bmatrix}$ <p>pixels/m</p>	$\begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix}$ <p>m</p>	$\begin{bmatrix} \alpha_x & 0 & \beta_x \\ 0 & \alpha_y & \beta_y \\ 0 & 0 & 1 \end{bmatrix}$ <p>pixels</p>
	=	

Extrinsic parameters: Rotation and translation



- In general, the camera coordinate frame will be related to the world coordinate frame by a rotation and a translation

Extrinsic parameters: Rotation and translation




- In general, the camera coordinate frame will be related to the world coordinate frame by a rotation and a translation

- In non-homogeneous coordinates, the transformation from **world** to normalized **camera** coordinate system is given by:

$$\tilde{X}_{cam} = R(\tilde{X} - \tilde{C}) = R\tilde{X} + t$$

coords. of point in normalized camera frame \rightarrow \tilde{X}_{cam}
 \tilde{X} \rightarrow 3x3 rotation matrix R
 \tilde{X} \rightarrow coords. of a point in world frame
 \tilde{C} \rightarrow coords. of camera center in world frame

Extrinsic parameters: Rotation and translation



In *non-homogeneous*
coordinates:


$$\tilde{\mathbf{X}}_{\text{cam}} = \mathbf{R}\tilde{\mathbf{X}} + \mathbf{t}$$

In *homogeneous*
coordinates:

$$\begin{pmatrix} \tilde{\mathbf{X}}_{\text{cam}} \\ 1 \end{pmatrix} = \begin{bmatrix} \mathbf{R} & \mathbf{t} \\ \mathbf{0}^T & 1 \end{bmatrix} \begin{pmatrix} \tilde{\mathbf{X}} \\ 1 \end{pmatrix}$$

3D transformation
matrix (4 x 4)

Extrinsic parameters: Rotation and translation



In *non-homogeneous*
coordinates:


$$\tilde{X}_{cam} = R\tilde{X} + t$$

In *homogeneous*
coordinates:

$$X_{cam} = \begin{bmatrix} R & t \\ \mathbf{0}^T & 1 \end{bmatrix} X$$

3D transformation
matrix (4 x 4)

Extrinsic parameters: Rotation and translation



In *non-homogeneous*
coordinates:

$$\tilde{X}_{\text{cam}} = R\tilde{X} + t$$


In *homogeneous*
coordinates:

$$\mathbf{X}_{\text{cam}} = \begin{bmatrix} R & t \\ \mathbf{0}^T & 1 \end{bmatrix} X$$

Transformation from normalized 3D
coordinates to pixel image coordinates:

$$x \cong K[I|0]X_{\text{cam}}$$

Extrinsic parameters: Rotation and translation




In *homogeneous*
coordinates:

$$x \cong K[I|\mathbf{0}] \begin{bmatrix} R & t \\ \mathbf{0}^T & 1 \end{bmatrix} X$$

Finally:

$$x \cong K[R|t]X \quad t = -R\tilde{C}$$

Extrinsic parameters: Rotation and translation



$$x \cong K[R|t]X$$

$$t = -R\tilde{C}$$

coords. of
camera center
in world frame

- What is the projection of the camera center in world coordinates?

$$PC = K[R \mid -R\tilde{C}] \begin{pmatrix} \tilde{C} \\ 1 \end{pmatrix} = K(R\tilde{C} - R\tilde{C}) = 0$$

- The camera center is the **null space** of the projection matrix!

Camera parameters: Summary


$$P = K[R|t]$$

- Intrinsic parameters
 - Principal point coordinates
 - Focal length
 - Pixel magnification factors
 - *Skew (non-rectangular pixels) – not important in practice*
 - *Radial distortion – important in practice!*

$$K = \begin{bmatrix} m_x & 0 & 0 \\ 0 & m_y & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} f & 0 & p_x \\ 0 & f & p_y \\ 0 & 0 & 1 \end{bmatrix} = \begin{bmatrix} \alpha_x & 0 & \beta_x \\ 0 & \alpha_y & \beta_y \\ 0 & 0 & 1 \end{bmatrix}$$




radial distortion



correction →



linear image



Overview

- Camera calibration
 - Intrinsic camera parameters
 - Extrinsic camera parameters
 - Estimation



Camera calibration

$$\mathbf{x} \cong \mathbf{K}[\mathbf{R} \ \mathbf{t}]\mathbf{X}$$

$$\begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \cong \begin{bmatrix} p_{11} & p_{12} & p_{13} & p_{14} \\ p_{21} & p_{22} & p_{23} & p_{24} \\ p_{31} & p_{32} & p_{33} & p_{34} \end{bmatrix} \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix}$$

Camera calibration

- Given n points with known 3D coordinates X_i and known image projections x_i , estimate the camera parameters



Camera calibration

- Given n points with known 3D coordinates X_i and known image projections x_i , estimate the camera parameters



Image credit: J. Hays

Known 2D
image coords

Known 3D
locations

880	214	312.747	309.140	30.086
43	203	305.796	311.649	30.356
270	197	307.694	312.358	30.418
886	347	310.149	307.186	29.298
745	302	311.937	310.105	29.216
943	128	311.202	307.572	30.682
476	590	307.106	306.876	28.660
419	214	309.317	312.490	30.230
317	335	307.435	310.151	29.318
783	521	308.253	306.300	28.881
235	427	306.650	309.301	28.905
665	429	308.069	306.831	29.189
655	362	309.671	308.834	29.029
427	333	308.255	309.955	29.267
412	415	307.546	308.613	28.963
746	351	311.036	309.206	28.913
434	415	307.518	308.175	29.069
525	234	309.950	311.262	29.990
716	308	312.160	310.772	29.080
602	187	311.988	312.709	30.514

Camera calibration: Linear method

- Recall homography fitting:

$$x_i \cong \mathbf{P}X_i \quad x_i \times \mathbf{P}X_i = \mathbf{0} \quad \begin{pmatrix} x_i \\ y_i \\ 1 \end{pmatrix} \times \begin{pmatrix} \mathbf{P}_1^T X_i \\ \mathbf{P}_2^T X_i \\ \mathbf{P}_3^T X_i \end{pmatrix} = \mathbf{0}$$

$$\begin{bmatrix} \mathbf{0}^T & -X_i^T & y_i X_i^T \\ X_i^T & \mathbf{0}^T & -x_i X_i^T \\ -y_i X_i^T & x_i X_i^T & \mathbf{0}^T \end{bmatrix} \begin{pmatrix} \mathbf{P}_1 \\ \mathbf{P}_2 \\ \mathbf{P}_3 \end{pmatrix} = \mathbf{0}$$

One match gives **two** linearly independent constraints

Camera calibration: Linear method

- Final linear system:

$$\begin{bmatrix} \mathbf{0}^T & \mathbf{X}_1^T & -y_1 \mathbf{X}_1^T \\ \mathbf{X}_1^T & \mathbf{0}^T & -x_1 \mathbf{X}_1^T \\ \dots & \dots & \dots \\ \mathbf{0}^T & \mathbf{X}_n^T & -y_n \mathbf{X}_n^T \\ \mathbf{X}_n^T & \mathbf{0}^T & -x_n \mathbf{X}_n^T \end{bmatrix} \begin{pmatrix} \mathbf{P}_1 \\ \mathbf{P}_2 \\ \mathbf{P}_3 \end{pmatrix} = 0 \quad \mathbf{A}\mathbf{p} = 0$$

- One 2D/3D correspondence gives two linearly independent equations
 - The projection matrix has 11 degrees of freedom
 - 6 correspondences needed for a minimal solution
- Homogeneous least squares: find \mathbf{p} minimizing $\|\mathbf{A}\mathbf{p}\|^2$
 - Solution is eigenvector of $\mathbf{A}^T\mathbf{A}$ corresponding to smallest eigenvalue

Camera calibration: Linear method

- Final linear system:

$$\begin{bmatrix} \mathbf{0}^T & \mathbf{X}_1^T & -y_1 \mathbf{X}_1^T \\ \mathbf{X}_1^T & \mathbf{0}^T & -x_1 \mathbf{X}_1^T \\ \dots & \dots & \dots \\ \mathbf{0}^T & \mathbf{X}_n^T & -y_n \mathbf{X}_n^T \\ \mathbf{X}_n^T & \mathbf{0}^T & -x_n \mathbf{X}_n^T \end{bmatrix} \begin{pmatrix} \mathbf{P}_1 \\ \mathbf{P}_2 \\ \mathbf{P}_3 \end{pmatrix} = 0 \quad \mathbf{A}\mathbf{p} = 0$$

- What if all the n 3D points are *coplanar*, i.e., there exists a set of line parameters $\mathbf{\Pi}^T = (a, b, c, d)^T$ such that $\mathbf{\Pi}^T \mathbf{X}_i = 0$ for all i ?
 - Then we will get *degenerate solutions* $(\mathbf{\Pi}, \mathbf{0}, \mathbf{0})$, $(\mathbf{0}, \mathbf{\Pi}, \mathbf{0})$, or $(\mathbf{0}, \mathbf{0}, \mathbf{\Pi})$

Camera calibration: Linear vs. nonlinear

- Linear calibration is easy to formulate and solve, but it doesn't directly tell us the camera parameters

$$\begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \cong \begin{bmatrix} p_{11} & p_{12} & p_{13} & p_{14} \\ p_{21} & p_{22} & p_{23} & p_{24} \\ p_{31} & p_{32} & p_{33} & p_{34} \end{bmatrix} \begin{pmatrix} X \\ Y \\ Z \\ 1 \end{pmatrix} \quad \text{vs.} \quad \mathbf{x} \cong \mathbf{K}[\mathbf{R} \mid \mathbf{t}]\mathbf{X}$$

Camera calibration: Linear vs. nonlinear

- Linear calibration is easy to formulate and solve, but it doesn't directly tell us the camera parameters
- In practice, *non-linear* methods are preferred
 - Write down objective function in terms of intrinsic and extrinsic parameters, as sum of squared distances between measured 2D points and estimated projections of 3D points:

$$\sum_i \|\text{proj}(\mathbf{K}[\mathbf{R} \mid \mathbf{t}]\mathbf{X}_i) - \mathbf{x}_i\|_2^2$$

- Can include radial distortion and constraints such as known focal length, orthogonality, visibility of points
- Minimize error using non-linear optimization package
- Can initialize solution with output of linear method (perform QR decomposition to get \mathbf{K} and \mathbf{R} from \mathbf{P})

Overview

- Camera calibration
 - Intrinsic camera parameters
 - Extrinsic camera parameters
 - Estimation
- First taste of 3D reconstruction: triangulation


Triangulation

- Given projections of a 3D point in two or more images (with known camera matrices), find the coordinates of the point




Triangulation

- Given projections of a 3D point in two or more images (with known camera matrices), find the coordinates of the point




Triangulation

- Given projections of a 3D point in two or more images (with known camera matrices), find the coordinates of the point




Triangulation

- We want to intersect the two visual rays corresponding to x_1 and x_2 , but because of noise and numerical errors, they don't meet exactly



Triangulation: Geometric approach


- Find shortest segment connecting the two viewing rays and let X be the midpoint of that segment



Triangulation: Nonlinear approach

- Find \mathbf{X} that minimizes

$$\|\text{proj}(\mathbf{P}_1\mathbf{X}) - \mathbf{x}_1\|_2^2 + \|\text{proj}(\mathbf{P}_2\mathbf{X}) - \mathbf{x}_2\|_2^2$$



Triangulation: Linear approach

$$\begin{array}{lll} \mathbf{x}_1 \cong \mathbf{P}_1 \mathbf{X} & \mathbf{x}_1 \times \mathbf{P}_1 \mathbf{X} = \mathbf{0} & [\mathbf{x}_{1 \times}] \mathbf{P}_1 \mathbf{X} = \mathbf{0} \\ \mathbf{x}_2 \cong \mathbf{P}_2 \mathbf{X} & \mathbf{x}_2 \times \mathbf{P}_2 \mathbf{X} = \mathbf{0} & [\mathbf{x}_{2 \times}] \mathbf{P}_2 \mathbf{X} = \mathbf{0} \end{array}$$

- Rewrite cross product as matrix multiplication:

$$\mathbf{a} \times \mathbf{b} = \begin{bmatrix} 0 & -a_3 & a_2 \\ a_3 & 0 & -a_1 \\ -a_2 & a_1 & 0 \end{bmatrix} \begin{pmatrix} b_1 \\ b_2 \\ b_3 \end{pmatrix} = [\mathbf{a}_{\times}] \mathbf{b}$$

Triangulation: Linear approach

$$\mathbf{x}_1 \cong \mathbf{P}_1 \mathbf{X}$$


$$\mathbf{x}_2 \cong \mathbf{P}_2 \mathbf{X}$$

$$\mathbf{x}_1 \times \mathbf{P}_1 \mathbf{X} = \mathbf{0}$$

$$\mathbf{x}_2 \times \mathbf{P}_2 \mathbf{X} = \mathbf{0}$$

$$[\mathbf{x}_{1 \times}] \mathbf{P}_1 \mathbf{X} = \mathbf{0}$$


$$[\mathbf{x}_{2 \times}] \mathbf{P}_2 \mathbf{X} = \mathbf{0}$$



Two independent equations each in terms of three unknown entries of \mathbf{X}


Preview: Structure from motion

Figure credit:
Noah Snavely




- Given 2D point correspondences between multiple images, compute the camera parameters and the 3D points

Preview: Structure from motion



- **Structure:** Given *known cameras* and projections of the same 3D point in two or more images, compute the 3D coordinates of that point
 - Triangulation!

Preview: Structure from motion



- **Motion:** Given a set of *known* 3D points seen by a camera, compute the camera parameters
 - Calibration!

Multi-view geometry “Bible”

