

CS4670 / 5670: Computer Vision

Noah Snavely

Lecture 14: Projection




"The School of Athens," Raphael

Reading

- Szeliski 2.1.3-2.1.6

Modeling projection



- Projection equation

$$(x, y, z) \rightarrow \left(-d\frac{x}{z}, -d\frac{y}{z}\right)$$

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & -1/d & 0 \end{bmatrix} \begin{bmatrix} x \\ y \\ z \\ 1 \end{bmatrix} = \begin{bmatrix} x \\ y \\ -z/d \end{bmatrix} \Rightarrow \left(-d\frac{x}{z}, -d\frac{y}{z}\right)$$


divide by third coordinate

Projection properties



- Many-to-one: any points along same ray map to same point in image
- Points \rightarrow points
- Lines \rightarrow lines (collinearity is preserved)
 - But line through focal point projects to a point
- Planes \rightarrow planes (or half-planes)
 - But plane through focal point projects to line

Projection properties


- Parallel lines converge at a vanishing point
 - Each direction in space has its own vanishing point
 - But parallels parallel to the image plane remain parallel



Orthographic projection



Perspective projection




Questions?

Camera parameters

- How many numbers do we need to describe a camera?
- We need to describe its *pose* in the world
- We need to describe its internal parameters

A Tale of Two Coordinate Systems



Camera parameters

- To project a point (x,y,z) in *world* coordinates into a camera
- First transform (x,y,z) into *camera* coordinates
- Need to know
 - Camera position (in world coordinates)
 - Camera orientation (in world coordinates)
- Then project into the image plane
 - Need to know camera *intrinsic*s
- These can all be described with matrices


Camera parameters

A camera is described by several parameters

- Translation T of the optical center from the origin of world coords
- Rotation R of the image plane
- focal length f , principle point (x'_c, y'_c) , pixel size (s_x, s_y)
- blue parameters are called “*extrinsics*,” red are “*intrinsic*s”

Projection equation

$$\mathbf{X} = \begin{bmatrix} sx \\ sy \\ s \end{bmatrix} = \begin{bmatrix} * & * & * & * \\ * & * & * & * \\ * & * & * & * \end{bmatrix} \begin{bmatrix} X \\ Y \\ Z \\ 1 \end{bmatrix} = \mathbf{P}\mathbf{X}$$



- The projection matrix models the cumulative effect of all parameters
- Useful to decompose into a series of operations


$$\mathbf{P} = \begin{bmatrix} -fs_x & 0 & x'_c \\ 0 & -fs_y & y'_c \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix} \begin{bmatrix} \mathbf{R}_{3 \times 3} & \mathbf{0}_{3 \times 1} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix} \begin{bmatrix} \mathbf{I}_{3 \times 3} & \mathbf{T}_{3 \times 1} \\ \mathbf{0}_{1 \times 3} & 1 \end{bmatrix}$$

intrinsic projection rotation translation identity matrix

- The definitions of these parameters are **not** completely standardized
 - especially intrinsic—varies from one book to another

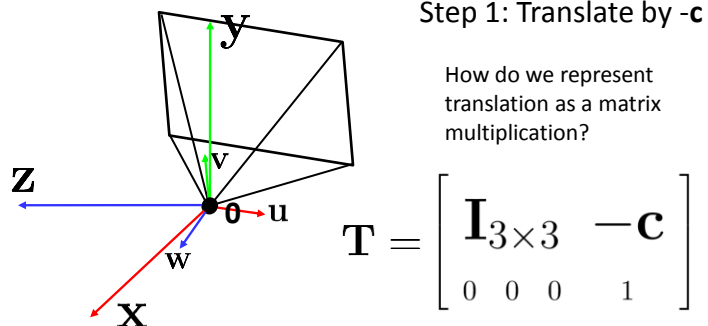
Extrinsics

- How do we get the camera to “canonical form”?
 - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)




Extrinsics

- How do we get the camera to “canonical form”?
 - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)




Extrinsics

- How do we get the camera to “canonical form”?
 - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)



Extrinsics

- How do we get the camera to “canonical form”?
 - (Center of projection at the origin, x-axis points right, y-axis points up, z-axis points backwards)



Perspective projection

$$\begin{bmatrix} -f & 0 & 0 \\ 0 & -f & 0 \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}$$

K
(intrinsic) (converts from 3D rays in camera coordinate system to pixel coordinates)

in general, $\mathbf{K} = \begin{bmatrix} -f & s & c_x \\ 0 & -\alpha f & c_y \\ 0 & 0 & 1 \end{bmatrix}$ (upper triangular matrix)

α : **aspect ratio** (1 unless pixels are not square)

S : **skew** (0 unless pixels are shaped like rhombi/parallelograms)

(c_x, c_y) : **principal point** ((0,0) unless optical axis doesn't intersect projection plane at origin)

Focal length

- Can think of as “zoom”



24mm



50mm



200mm



800mm



- Related to *field of view*

Projection matrix


$$\mathbf{\Pi} = \mathbf{K} \underbrace{\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{bmatrix}}_{\text{projection}} \underbrace{\begin{bmatrix} \mathbf{R} & \begin{matrix} 0 \\ 0 \\ 0 \end{matrix} \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\text{rotation}} \underbrace{\begin{bmatrix} \mathbf{I}_{3 \times 3} & -\mathbf{c} \\ 0 & 0 & 0 & 1 \end{bmatrix}}_{\text{translation}}$$

$$\begin{bmatrix} \mathbf{R} & | & -\mathbf{R}\mathbf{c} \end{bmatrix}$$

(t in book's notation)

$$\mathbf{\Pi} = \mathbf{K} \begin{bmatrix} \mathbf{R} & | & -\mathbf{R}\mathbf{c} \end{bmatrix}$$

Projection matrix



Questions?

Perspective distortion


- Problem for architectural photography: converging verticals




Source: F. Durand

Perspective distortion


- Problem for architectural photography: converging verticals



Tilting the camera upwards results in converging verticals





Keeping the camera level, with an ordinary lens, captures only the bottom portion of the building



Shifting the lens upwards results in a picture of the entire subject

- Solution: view camera (lens shifted w.r.t. film)



http://en.wikipedia.org/wiki/Perspective_correction_lens

Source: F. Durand

Perspective distortion

- Problem for architectural photography: converging verticals
- Result:



Source: F. Durand

Perspective distortion

- What does a sphere project to?






Image source: F. Durand

Perspective distortion

- The exterior columns appear bigger
- The distortion is not due to lens flaws
- Problem pointed out by Da Vinci




Slide by F. Durand

Perspective distortion: People



Distortion




- Radial distortion of the image
 - Caused by imperfect lenses
 - Deviations are most noticeable for rays that pass through the edge of the lens

Correcting radial distortion



from [Helmut Dersch](#)

Distortion



Modeling distortion

$$\begin{array}{l} \text{Project } (\hat{x}, \hat{y}, \hat{z}) \\ \text{to "normalized"} \\ \text{image coordinates} \end{array} \quad \begin{array}{l} x'_n = \hat{x}/\hat{z} \\ y'_n = \hat{y}/\hat{z} \end{array}$$

$$\begin{array}{l} \text{Apply radial distortion} \end{array} \quad \begin{array}{l} r^2 = x'_n{}^2 + y'_n{}^2 \\ x'_d = x'_n(1 + \kappa_1 r^2 + \kappa_2 r^4) \\ y'_d = y'_n(1 + \kappa_1 r^2 + \kappa_2 r^4) \end{array}$$

$$\begin{array}{l} \text{Apply focal length} \\ \text{translate image center} \end{array} \quad \begin{array}{l} x' = f x'_d + x_c \\ y' = f y'_d + y_c \end{array}$$

- To model lens distortion
 - Use above projection operation instead of standard projection matrix multiplication