5a. More transformations

The pinhole camera

The first camera ever made was the "camera obscura": a dark room or chamber with a hole in the window shade.

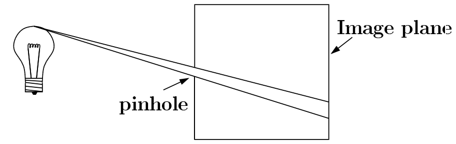

Pinhole camera

The first camera that actually recorded an image simply replaced the back wall with material sensitive to light, i.e., film.

This kind of camera is called a "pinhole camera."

Q: What does the pinhole camera have in common with a conventional camera or the human eye?

Q: How is it different?

The pinhole camera (cont'd)

Today, we typically use the pinhole camera for creating images of synthetic scenes.

However, since we are not bound by the laws of physics, we conveniently place the image plane in front of the pinhole.

Q: Why is this convenient?

Projections

"Projections" transform points in n-space to m-space, where $m < n$.

In 3-D, we map points from 3-space to the "projection plane" (PP) along "projectors" emanating from the "center of projection" (COP):

The center of projection is exactly the same as the pinhole in a pinhole camera.

There are two basic types of projections:

- "Perspective" — distance from COP to PP finite
- "Parallel" — distance from COP to PP infinite

Parallel projections

For parallel projections, we specify a "direction of projection" (DOP) instead of a COP.

There are two types of parallel projections:

- "Orthographic projection" — DOP perpendicular to PP
- "Oblique projection" — DOP not perpendicular to PP

We can write orthographic projection onto the $z = 0$ plane with a simple matrix.

Normally, we do not zero out or drop the z value right away. Why not?

Properties of parallel projection

Properties of parallel projection:

- Not realistic looking
- Good for exact measurements
- Are actually a kind of affine transformation
  - Parallel lines remain parallel
  - Angles not (in general) preserved
- Most often used in CAD, architectural drawings, etc., where taking exact measurement is important

Derivation of perspective projection

Consider the projection of a point onto the projection plane:

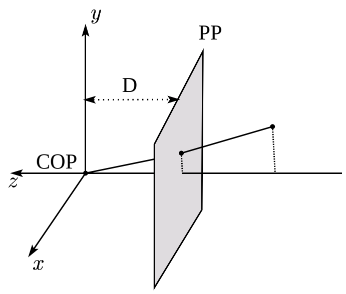

By similar triangles, we can compute how much the x and y coordinates are scaled:

Homogeneous coordinates

Remember how we said that affine transformations work with the last coordinate always set to one.

If we allow the last coordinate to vary arbitrarily, we enter the realm of *homogeneous coordinate*.

We divide all the coordinates by w:

If $w = 1$, then nothing changes.

Sometimes we call this division step the "perspective divide".

Homogeneous coordinates and perspective projection

Now we can re-write the perspective projection as a matrix equation:

After division by w, we get:

Projection implies dropping the z coordinate to give a 2D image, but we usually keep it around a little while longer. Why?

Projective goemetry

We can think of points as living in a higher dimensional space while obeying an equivalence relation:

$$(x, y, z) \to (x, y, z, 1) \Leftrightarrow (\lambda x, \lambda y, \lambda z, \lambda)$$

For points in 2D, we arrive at:

$$(x, y) \to (x, y, 1) \Leftrightarrow (\lambda x, \lambda y, \lambda)$$

An affine point is a homogeneous line

An affine point now has a geometric interpretaion of being a line in the higher dimensional homogenous space:

$$P(\lambda) = \lambda \mathbf{v}$$

where $\mathbf{v} = (x, y, z, 1)$ in 3D or $\mathbf{v} = (x, y, 1)$ in 2D.

We can sketch a point, $(x_o, y_o, 1)$:

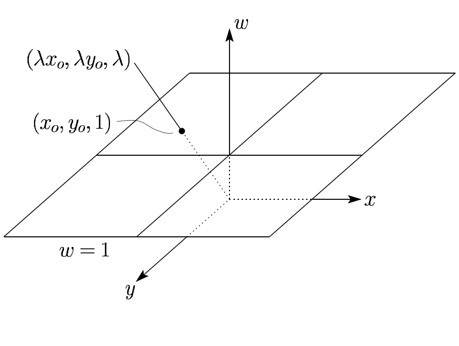

An affine line is a homogeneous plane

A 2D line describes a plane in the homogeneous coordinate system:

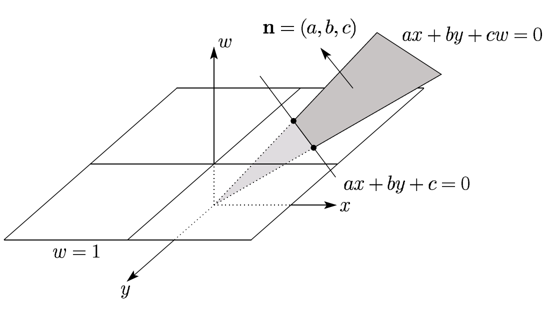

Two points make a line

Given a pair of homogeneous points:

$$p_1 = (x_1, y_1, w_1)$$
$$p_2 = (x_2, y_2, w_2)$$

We can easily compute the line that passes through them.

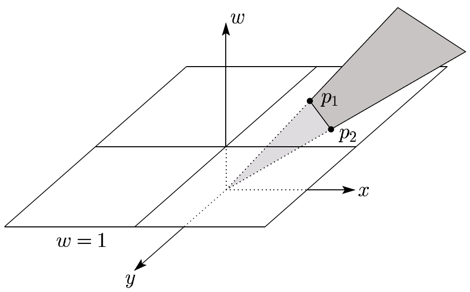

Two lines intersect in a point

Similarly, given a pair of homogeneous lines:

$$a_1 x + b_1 y + c_1 w = 0$$
$$a_2 x + b_2 y + c_2 w = 0$$

We can compute their point of intersection.

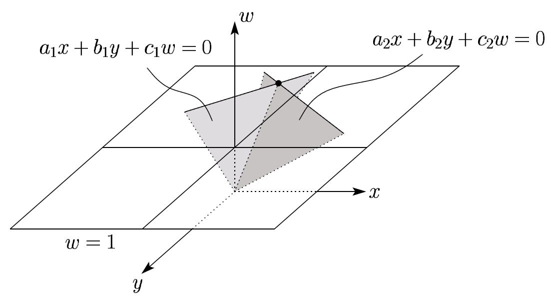

Perspective projection, revisited

An alternative formulation for perspective projection places the image plane at the x-y plane, and the center or projection at a distance D away.

The resulting projection matrix is:

$$\begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \frac{-1}{D} & 1 \end{bmatrix}$$

Which in 2D is:

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & \frac{-1}{D} & 1 \end{bmatrix}$$

Perspective projection in 2D

We can now think of 2D perspective projection as:

- A 1D projection
- of a 2D slice
- through a 3D space

For 3D perspective projection, add one to each dimension above.

Consider two parallel line segments:

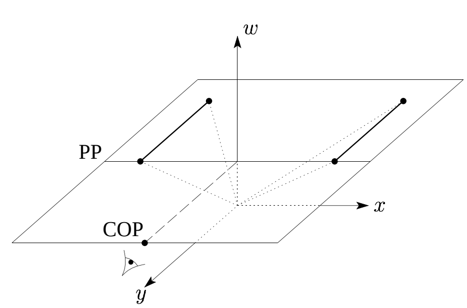

Projection in 2D

Now apply the perspective matrix. It is equivalent to a shear in w:

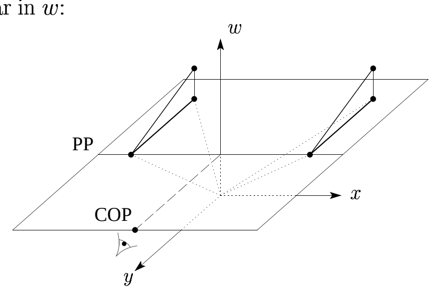

Now reproject the points through the origin to establish their new affine coordinates:

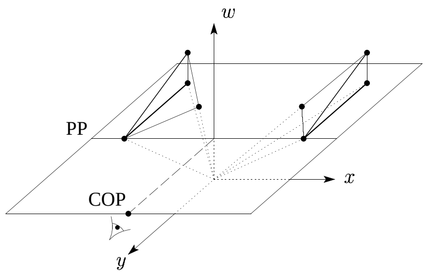

Vanishing points

What happens as the lines continue toward infinity?

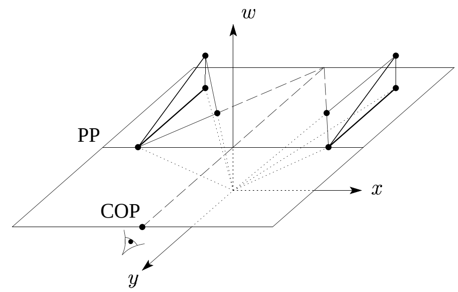

Properties of projective transformations

Here are some properties of projective transformations:

- Can be written as linear transformations in homogenous coordinates
- Lines map to lines
- Parallel lines do not necessarily remain parallel
- Ratios are not preserved
- Ratios are preserved

$$\begin{aligned} \text{cross} - \text{ratio}(P, Q, R, S) &= \frac{\text{ratio}(P, R, S)}{\text{ratio}(P, Q, S)} \\ &= \frac{\beta_1/\beta_2}{\alpha_1/\alpha_2} \end{aligned}$$

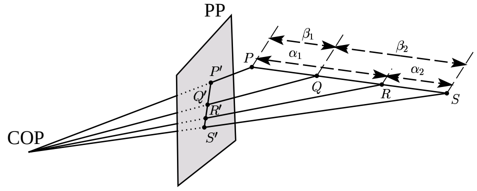

Properties of projective transformations

One of the advantages of perspective projection is that size varies inversely with distance — looks realistic.

A disadvantage is that we can't judge distances as exactly as we can with parallel projections.

Q: Why did nature give us eyes that perform perspective projections?

Q: Do our eyes "see in 3D"?