Image Stitching

Image Processing Lecture 2

Nicholas Dwork

Here are three images captured simply by rotating the camera.

Image Stitching

The process of combining multiple images into a single image.

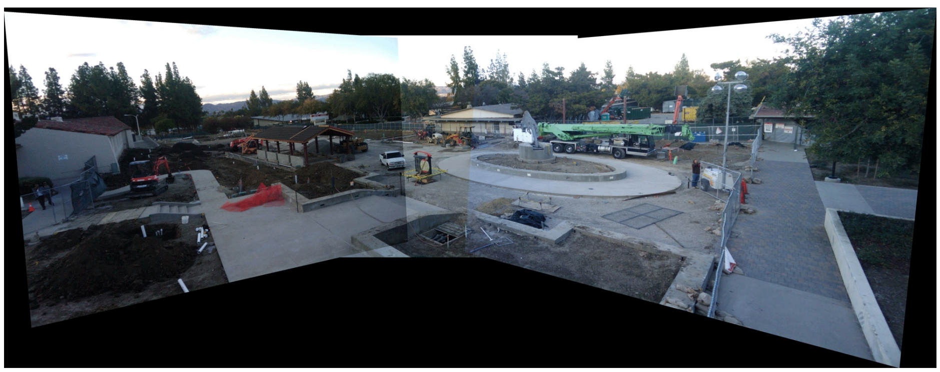

Homography

Recall that a Homography is a square matrix applied to homogeneous coordinates.

We are now going to use Homographies to stitch images together.

Theorem

Two images captured by a rotated camera are related by a Homography.

$$I_2(x) = I_1(Hx)$$

I_1, I_2 **are the first and second image.**

H **is the Homography**

Theorem

Two images of a flat plane are related by a Homography.

$$I_2(x) = I_1(Hx)$$

I_1, I_2 **are the first and second image.**

H **is the Homography**

Big Question:

How do we find the Homography that relates two images together?

Picking Matching Points

We will use matched points in the two images.

At first, we will match points by hand. Later, we'll see how to have the computer match points.

We know all the xs and the ys. We have to find the H.

$$
\begin{aligned}
y_1 &= Hx_1 \\
y_2 &= Hx_2 \\
&\vdots \\
y_N &= Hx_N
\end{aligned}
$$

Note: the equality above is in the Homogeneous sense. That is, the vector y_i is a scalar multiple of Hx_i .

Since the points are all in a two dimensional projective space, they are all three element vectors.

Since they are equal in a projective sense, the components of the vectors are scalar multiples of each other.

If we think of a pair of matching three element vectors as 3D Euclidean vectors, then they would be parallel.

This implies the following:

$$y_i \times (Hx_i) = 0$$

Recall that we can rewrite Hx_i as the following

$$H\,x_i = \begin{bmatrix} r_1^T\,x_i \\ r_2^T\,x_i \\ r_3^T\,x_i \end{bmatrix}$$

where r_i^T is the i^{th} row of H. Let $x_i' = Hx_i$.
Recall that the cross product can be rewritten as

$$y \times x_i' = \begin{bmatrix} y_{i,2}\,x_{i,3}' - x_{i,2}'\,y_{i,3} \\ x_{i,1}'\,y_{i,3} - x_{i,3}'\,y_{i,1} \\ x_{i,1}'\,y_{i,2} - y_{i,2}\,x_{i,1}' \end{bmatrix}$$

Combining the previous equations

$$y_i \times H\,x_i = \begin{bmatrix} y_{i,2}\,(r_3^T x_i) - (r_2^T x_i)\,y_{i,3} \\ (r_1^T x_i)\,y_{i,3} - (r_3^T x_i)\,y_{i,1} \\ (r_1^T x_i)\,y_{i,2} - y_{i,2}\,(r_1^T x_i) \end{bmatrix} = 0$$

We're going to use a neat trick. We're going to isolate the rows of H by rewriting the above equation as

$$\begin{bmatrix} 0^T & -y_{i,3}\,x_i^T & y_{i,2}\,x_i^T \\ y_{i,3}\,x_i^T & 0^T & -y_{i,1}\,x_i^T \\ -y_{i,2}\,x_i^T & y_{i,1}\,x_i^T & 0^T \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \\ r_3 \end{bmatrix} = 0$$

$$\begin{bmatrix} 0^T & -y_{i,3}\, x_i^T & y_{i,2}\, x_i^T \\ y_{i,3}\, x_i^T & 0^T & -y_{i,1}\, x_i^T \\ -y_{i,2}\, x_i^T & y_{i,1}\, x_i^T & 0^T \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \\ r_3 \end{bmatrix} = 0$$

This is a system of equations! Each pair of matched points gives us 3 equations with 9 unknowns.

Unfortunately, one of the equations is redundant. (The third equation does not provide any more information.) So we only have two equations with 9 unknowns.

$$\begin{bmatrix} 0^T & -y_{i,3}\, x_i^T & y_{i,2}\, x_i^T \\ y_{i,3}\, x_i^T & 0^T & -y_{i,1}\, x_i^T \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \\ r_3 \end{bmatrix} = 0$$

$$\begin{bmatrix} 0^T & -y_{i,3}\, x_i^T & y_{i,2}\, x_i^T \\ y_{i,3}\, x_i^T & 0^T & -y_{i,1}\, x_i^T \end{bmatrix} \begin{bmatrix} r_1 \\ r_2 \\ r_3 \end{bmatrix} = 0$$

We'll rewrite this as $A_i h = 0$

where $h = \begin{bmatrix} r_1 \\ r_2 \\ r_3 \end{bmatrix}$

We can combine the equations from all matched points into a single system

$$A\,h = \begin{bmatrix} A_1 \\ A_2 \\ \vdots \\ A_N \end{bmatrix} h = 0$$

We can find h**, and therefore** H **by finding a non-zero vector in the null space of** A **!**

Finding a Vector in the Null Space

Compute the SVD of the matrix A:

```
[u,s,v] = svd( A );
```

Find the smallest singular value:

```
[minS,minIndx] = min( diag(s) );
```

The null vector is the column of v **corresponding to the minimum singular value:**

```
nullVector = v(:,minIndx);
```

Question:

Now that we have the Homography, how do we stitch the images together?

Finding the Range

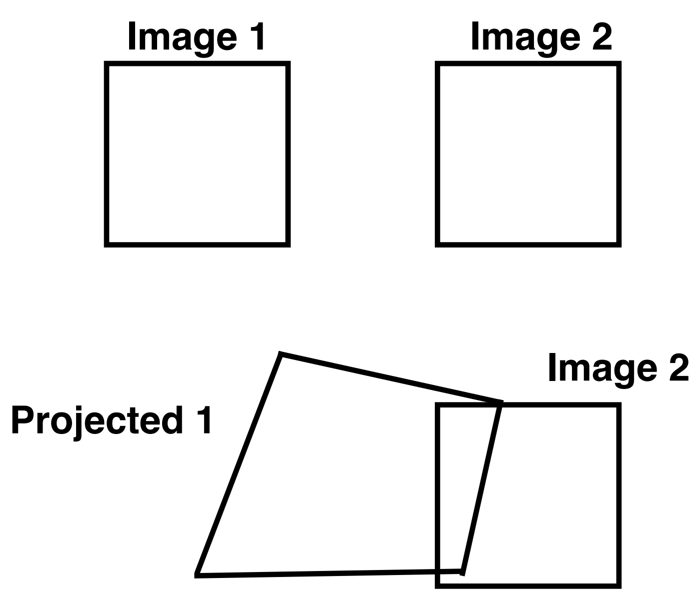

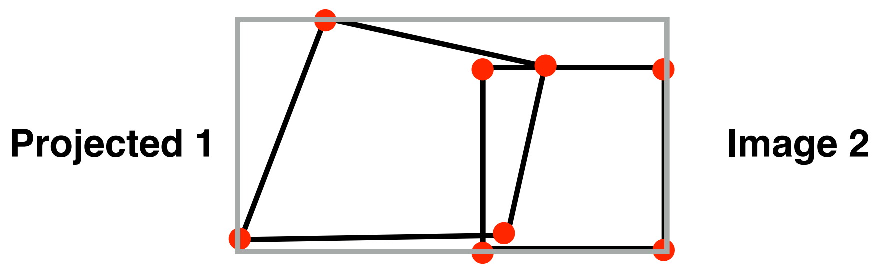

The min and max coordinates of the red dots tell you how big to make your stitched image.

They tell you the range of the projective space!

Note: You'll have to properly place Image 2.
If it comes out as shown above, you're going to have to shift Image 2 to the right.

Now that you have the range, pull Image 1 into the projected space.

Iterate over the pixels in the projected space.

For each pixel, apply H^{-1} to see which pixel in Image 1 belongs at that location.

Populate the value in the stitched image with the intensity from Image 1.

$$I_P(x) = I_1(H^{-1}x)$$

Does this picture make more sense now?