Image Registration

Paul Heckbert, Sept. 1999

15-869, Image-Based Modeling and Rendering

Applications of Image Registration

- **video stabilization**
  - to find translation of image
- **image mosaicing**
  - to find affine or projective warp, most often
- **stereo** (where it is known as stereo matching)
  - to find disparity and depth at each pixel
- **structure from motion** (where it is known as optical flow or motion analysis)
  - to find motion vector at each pixel

The first two look for a global mapping or correspondence; the last two look for local ones.

Image Registration Problem

Find the mapping (correspondence) given by

x'(x,y) and *y'(x,y)*

that best registers two pictures

I(x,y) and *I'(x',y')* - *note that the pictures use different indices*

Simple example: translational image registration

Mapping is: *x'(x,y)=x+u, y'(x,y)=y+v*

Quantify misregistration error using the L_2 metric, say:

$$E = \sum_W [I'(x+u, y+v) - I(x,y)]^2$$

Find the translation *(u,v)* that minimizes *E*

Other registration problems use more general mappings and other similarity measures.

Basic Similarity Measures

Sum of Squared Differences (SSD):

$$E = \sum_W [I'(x',y') - I(x,y)]^2$$ where W is window of interest

Minimize E (zero means exact match)

Problems: sensitive to global brightness differences between images

Cross-Correlation:

simple form: $$C = \sum_W I'(x',y') I(x,y)$$

normalized form: $$C = \frac{1}{ss'} \sum_W [I'(x',y') - \bar{I'}][I(x,y) - \bar{I}]$$

where $\bar{I} = \frac{1}{|W|} \sum_W I(x,y), \quad s = \left\{ \sum_W [I(x,y) - \bar{I}]^2 \right\}^{1/2}, \quad \bar{I'} = \ldots$

Maximize C (1 means exact match)

Less sensitive to brightness differences, but still assumes images are linearly interrelated, i.e. $I' = aI + b$

Optimizing the Registration

We have transformed the registration problem into a multidimensional, nonlinear optimization problem.

Each mapping has some parameters **q**

- translation: **q**=(u,v);   x'=x+u, y'=y+v
- affine warp: **q**=(a,b,c,d,e,f);   x'=ax+by+c, y'=dx+ey+f
- projective warp: **q**=(a,b,c,d,e,f,g,h)
- range surface: **q**=(vector of many thousands of depths)

Search over the multidimensional space for the **q**=**q*** that extremizes the chosen similarity measure.

Common algorithm: Levenberg-Marquardt

- Requires computation of first and second partial derivatives of similarity measure w.r.t. mapping parameters.

To optimize: search in coarse-to-fine fashion.

Problem 1: What Window to Use?

- When mapping global, *W* can be entire overlap region.
  - No problem (except speed, possibly).

- When mapping local,
  - could *W* be a single pixel? no
  - if window too small, too many false matches, too sensitive to noise
  - if window big, it could include pixels with a mapping very different from center pixels', polluting the similarity measure.
    - This is a problem for stereo matching (foreshortening problem) or optical flow.

  - Solutions
    - Adaptive window size [Kanade-Okutomi 91]
    - Eliminate the window bias entirely [Xiong, IJCV 97]

Problem 2: Too Sensitive to Light Differences

- Both methods are sensitive to global differences such as
  - change of lighting
  - change of exposure time or aperture (common when using cameras with automatic gain control!)
  - anisotropic reflectance (methods commonly assume reflectance is diffuse, i.e. independent of viewing direction)
- SSD very sensitive to brightness differences
- cross-correlation less so, since it normalizes, but it still expects linearly-interrelated images

Partial Solutions:

- find edges with, say, Laplacian of Gaussian (Marr-Hildreth operator)
- use thresholded signed gradient of image
- passband filter (eliminate low frequencies, keep mid-range frequencies)

Problem 3: Sensor Fusion

What if images are from different sensors, e.g. CAT scan and NMR?

In this case, we won't have $I' = aI + b$ or even $I' = f(I)$, in general, but we will have that I and I' tell more about each other when properly aligned than when misaligned.

Solution:

Maximize the *mutual information* [Viola, IJCV 97], defined as

MI = (entropy of I) + (entropy of the part of I' into which I projects) - (joint entropy of I and I' as currently registered)

The third term is small when they're well registered.

This is based on weaker assumptions than most previous registration methods, so it's more general, but also perhaps less ideal for particular problems.