Course notes, CS664, 9/14/04

- Slides: good versus bad stereo algorithms; joint histograms.

REGISTRATION VERSUS CORRESPONDENCE

- Computing visual correspondence locally is usually done by solving a registration problem on a rectangular subwindow of the image. Typically T is very simple (why?)

- A disparity map can be viewed as a transformation with $|P|$ degrees of freedom. In fact, it is in some sense the most general possible, since it arbitrarily maps pixels to pixels.

- There are also intermediate representations, such as piecewise smooth disparity maps. These lead to some stereo methods that currently give the best results.

- To think of correspondence as energy minimization, consider a "data" term that comes from the joint histogram, plus some measure of smoothness of the disparity map.

- Some ways of doing correspondence involve complicated smoothness measures with more degrees of freedom than simple transformations but fewer than the number of pixels. Sometimes called "non-rigid registration", usually involves splines or similar interpolated surfaces.

REGISTRATION APPLICATION: SUPERRESOLUTION

- Superresolution: cute term, comes from astronomy (telescopes).

- Idea: how can we generate a 256x256 image from a 128x128 camera? In a static scene, by using video (plus a little algebra).

PIXEL LABELING PROBLEMS

- Image restoration problem definition: suppose true intensity is 100, we may observe 101 or 99 (from time to time). Can we invert this process?

- Motivation: pre-processing for edge detection or segmentation;

- Note that this is an ill-posed problem. Can always have output = input

- The math of image restoration is much simpler than motion or stereo, because only one image is involved, and various 3D-related issues (occlusions, transparency, etc.) don't appear.

- In fact there is a generalization of (image restoration, stereo, motion) called the *pixel labeling problem*.

- You are given a set of pixels with a neighborhood system (i.e. a graph) and a set of labels. The goal is to assign a label to each pixel.

- As a rule, there will be some "data cost" function saying how much each pixel (dis-)likes each label, and way of imposing spatial coherence (since if each pixel picks independently you'll get garbage).

- The details of how you impose spatial coherence are important, and particular to the algorithm.

- Example: for stereo or motion, the cost to assign disparity d to pixel p is $(I_l(p) - I_r(p+d))^2$. For image restoration, if we originally observed $I(p)$, the cost to assign the intensity i to the pixel p is $(I(p) - i)^2$. Note how the image restoration problem is simpler.

SOLVING PIXEL LABELING PROBLEMS

- A standard way to solve them is by energy minimization.

- Formalize the problem as finding $f : \mathcal{P} \mapsto \mathcal{L}$ that minimizes

$$E_{data}(f) + E_{smooth}(f)$$

Note that only the first term depends on the observed images!