Course notes, CS664, 10/28/04

- Administrivia: PS1 due soon, please choose groups (almost no one has).

- Next quiz will be a week from today, coverage through today.

- You will need to choose your final project Real Soon Now.

CONTINUOUS MINIMIZATION

- A vast number of problems can be viewed as optimization (including, e.g., least squares, which we will see in a moment).

- Convex versus not, constrained. Many interesting special cases such as linear programming.

- Generally emphasis on fast convergence to some kind of local minimum.

- Problems where the continuous method gets to the global min are rare in vision (these problems are convex).

- Example: Total Variation, $V(l_1, l_2) = |l_1 - l_2|$. For image denoising (restoration) E is convex (sum of convex functions is convex). Otherwise this is a hard continuous optimization problem (which can be solved exactly with graph cuts, for any convex V!).

- Calculus of variations is concerned with computing extrema of functionals (as calculus is with functions). Basic tool is the Euler equations which hold at extrema, but can also hold elsewhere. From a computational point of view these methods tend to compute local minima in a fairly weak sense.

MOTION

- Consider the input as a 3-D surface $I(x, y, t)$. Under the constant brightness assumption $\frac{dI}{dt} = 0$. Expanding this using the chain rule we have

$$\frac{dI}{dt} = \frac{\partial I}{\partial x}\frac{dx}{dt} + \frac{\partial I}{\partial y}\frac{dy}{dt} + \frac{\partial I}{\partial t} = 0.$$

  Now let $u(x, y) = \frac{dx}{dt}$ and $v(x, y) = \frac{dy}{dt}$.

- The partial derivatives can be measured from the image, u, v are the unknowns. So

$$\frac{\partial I}{\partial x}u + \frac{\partial I}{\partial y}v + \frac{\partial I}{\partial t} = 0.$$

OPTICAL FLOW CONSTRAINT EQUATION.

- Can write this as the dot product

$$\nabla I \cdot [u\ v] = -\frac{\partial I}{\partial t},$$

where $\nabla I = [\frac{\partial I}{\partial x}\frac{\partial I}{\partial y}] = [I_x I_y]$ is the (intensity) gradient.

- We want to compute u, v, but all we know is its projection upon the vector ∇I (the constraint line). [Draw picture]

- We can only measure flow in the direction of the gradient (steepest intensity change). APERTURE problem. This means the problem is ill-posed.

- What do we do? Use information from neighboring pixels! Why does this fail? The usual reasons.

- A good way to view the normal flow computation is as a local data cost. The main limitation is that it comes from a local differential problem formulation.

Horn and Schunk

- So-called "direct methods" (terrible term) use the normal flow vectors directly. One notion is to use local smoothness assumptions, often in a continuous minimization framework (calculus of variations).

- Horn and Schunk basic idea: minimize

$$E(u,v) = \int\int E_{data}(u,v) + E_{smooth}(u,v)$$

  where

$$E_{data}(u,v) = I_x u + I_y v + I_t$$

  which should be small if Constant Brightness holds (among other things).

- For smoothness they use

$$E_{smooth}(u,v) = |\nabla u|^2 + |\nabla v|^2$$

  which is just a local smoothness term.

- Famous paper for many reasons: OFCE, normal flow, plus energy minimization!

- Algorithm: compute local average $(\bar{u}, \bar{v})$, then move to constraint line depending on λ. Note that the current value isn't used (or is used only indirectly).