Course notes, CS664, 10/14/04

- Administrivia: PS1 is on the web.

- Please get started on this soon! Also, please pick your paper soon.

- Next in-class quiz a week from today (10/21).

GRAPH CUTS: POTTS MODEL

- Not easy to generalize this to multiple labels. There are generalizations of min cut, but they are all NP-hard.

- Natural one is called the Multiway Cut Problem, where the goal is to separate all the terminals, and the cost of a multiway cut is the sum of the edges.

- This problem is *very* closely related to the Potts model energy minimization problem. The NP-hardness result goes in one direction (given an instance of the MWC problem, find a Potts model problem that would solve it).

- There is also a natural construction that goes in the other direction. The t-link weight between the label l and the pixel p is $K - D(p, l)$. Here, K is a large constant.

- This has the right intuition — think about what happens when $D(p, l)$ is large or small. When $D(p, l)$ is large, the t-link weight is small and it is easy to cut, hence cheap for p not to have label l.

- When $D(p, l)$ is small, the t-link weight is large and it is hard to cut, hence p likes to have label l.

- Likewise, when a pixel p is linked to α and a neighbor q is linked to β, we need to cut the n-link between them.

EXPANSION MOVE ALGORITHM

- Greedy methods: find the cheapest labeling "nearby", go there, repeat. At the end there is no cheaper labeling "nearby".

- Metropolis at 0 temperature achieves this, where a nearby labeling comes from the sampling step (i.e., change one pixel).

- What if we could find the cheapest labeling in a very large neighborhood?

- Expansion move definition: f' is within a single α-expansion move of f if $f'(p) \neq f(p) \implies f'(p) = \alpha$.

- Note that this definition is not symmetric. Also note that it is a generalization of the standard move, in that a local min w.r.t. expansion moves is also a local min w.r.t. standard moves.

- There are an exponential number of expansion moves from a given labeling. Think of there being a binary variable associated with each pixel; if that variable is 1, the pixel changes to α.

- You can get across the entire space in $|\mathcal{L}|$ expansion moves (hence, you can get to the global minimum).

- You can fill in low-texture regions in a single expansion move!

- This is actually an extremely powerful algorithm, as long as we can figure out how to find the cheapest α-expansion move starting at f, where α, f are arbitrary.

BINARY ENERGY MINIMIZATION

- Note that this is a energy minimization problem over *binary* variables v_p, where

$$v_p = 0 \implies f'(p) = f(p),$$

$$v_p = 1 \implies f'(p) = \alpha.$$

There is one variable v_p for each pixel p.

- The energy is:

$$E(v_1, \ldots, v_p, v_q, \ldots, v_{|\mathcal{P}|}) = \sum_p E^1(v_p) + \sum_{p,q \in \mathcal{N}} E^2(v_p, v_q),$$

where

$$E^1 = \begin{array}{|c|}\hline D(p, f(p)) \\ \hline D(p, \alpha) \\ \hline \end{array},$$

and

$$E^2 = \begin{array}{|c|c|}\hline V(f(p), f(q)) & V(\alpha, f(q)) \\ \hline V(f(p), \alpha) & V(\alpha, \alpha) \\ \hline \end{array}.$$

REGULARITY

- There is a property of binary energy functions called *regularity*. A function of 1 binary variable is always regular. A function of two binary variables

A	B
C	D

is regular when $A + D \leq B + C$.

- A binary energy function of no more than 3 variables can be minimized exactly using graph cuts! This is a quite powerful construction which is a subroutine at this point in time.

- All we need to do is check that our binary function, which came from the expansion move algorithm, is regular. E^1 definitely is; what about E^2?

- We need:

$$V(f(p), f(q)) + V(\alpha, \alpha) \leq V(\alpha, f(q)) + V(f(p), \alpha).$$

Note this must hold for all f.

- The specific proof is a little painful, but the general one is easy! We simply observe that if V is a metric on the space of labels, this is clearly true, since $V(\alpha, \alpha) = 0$ and it becomes the triangle inequality!

METRIC LABELING PROBLEMS

- We've actually proved something stronger than we set out to prove! It turns out that the Potts model V is a metric, as is the truncated L_1 distance.