

Belief Net

Intuitively, one problem with the simple version of relaxation labeling that I've sketched is seen if we have two lines, A and B, that support each other. At one iteration, A supports B. At the next, B supports A. When this happens, A is getting support for its being figure that really originates in A itself. It would be better if we could tell the difference between B seeming to be figure because of A's support, or support from another source. Belief nets formalize probabilistic support. Belief propagation is an algorithm that performs inference according to this probabilistic support.

***** This part omitted in '06

First, a belief net is a graphical representation of a set of random variables that makes explicit some properties about causality and conditional independence. As an example, suppose we have a house alarm. It can be set off by a burglar or an earthquake. If it goes off, a neighbor, Sally or John, might call you. The arrows show causality. Conditional independence is also indicated; given that the alarm goes off, Sally and John calling are independent events. They're not independent if we don't know whether the alarm went off. The parent nodes have a different structure; the earthquake and burglary are independent if we don't know whether the alarm went off, but are not independent if we do know. That's because if the alarm goes off, knowing there was an earthquake effects my assessment of whether there was a burglary. To make this precise, we have to know the conditional probability of every variable, based on its causes.

We're going to talk about one simple algorithm for performing inference in these nets, that of belief propagation. Inference means that we are given some knowledge about the value of some variables, and want to determine the value of others. For example, in vision, we might know that some positions and orientations in the image have a contour going through them, and want to know whether others do. Though this algorithm is simple, it's very elegant and useful.

.....

We now look at how one would really perform correct probabilistic inference in an image. We first look at the 1D case, which is substantially easier than the 2D case. Suppose we have an image that we think is piecewise constant, corrupt by noise. Then the noise-free intensity at each pixel depends on the observed intensity at that pixel, and also on the rest of the image. We make a Markov assumption that the true intensity of a pixel only directly depends on its observed intensity, and the true intensity of its neighbors, and is conditionally independent of everything else. This is what we would get if we generate an image by deciding independently at each pixel, whether to copy the previous pixel or jump to a new value, and then corrupting this with independent noise.

We describe this with a chain of variables that represent the true intensities

$X_1 \rightarrow X_2 \rightarrow \dots X_n$

And corresponding variables, y_k , that indicate the observations. Collectively, we call all these e , for evidence. We want to compute the distribution of the true intensity at each pixel, conditioned on the observations. We do this by breaking the evidence into separate parts, which can all be handled independently.

Let's start with the simplest case, which will solve half the problem for us.

$$P(x_n | e) = P(e/x_n) P(x_n) / P(e).$$

$P(e)$ is the same for all values of x_n ; it's just a normalizing constant. So we can ignore it, and then normalize the values we get so they sum to one. $P(x_n)$ is the prior on x_n .

$$P(e|x_n) = P(e^- | x_n, y_n) * P(y_n|x_n)$$

e^- is defined to be all evidence to the left of x_n . $P(y_n | x_n)$ comes from our noise model, we assume we know all conditional probabilities.

$$P(e^- | x_n, y_n) = P(e^- | x_n) \lambda(x_n).$$

We compute this by taking advantage of conditional independence. If we knew x_{n-1} , x_n would then be independent of e^- . We don't know this, but we can sum over all possibilities:

$$\lambda(x_n) = \sum_{x_{n-1}} P(e^- | x_n, x_{n-1}) * P(x_{n-1} | x_n)$$

$$P(e^- | x_n, x_{n-1}) = P(e^- | x_{n-1}) = \lambda(x_{n-1}) * P(y_{n-1} | x_{n-1})$$

If we write $P(x_{n-1} | x_n)$ as a matrix, $M_{\{x_{n-1} | x_n\}}$ we get

$$\lambda(x_n) = M_{\{x_{n-1} | x_n\}} \lambda(x_{n-1}) * \text{componentwise multiplication with } P(y_{n-1} | x_{n-1}).$$

So we can compute probabilities recursively, starting at the left and working our way right.

If we want to handle a pixel that is not on the end, we can do this in much the same way, combining evidence from the left and from the right.

$$P(x_k | e) = P(e^- | x_k, y_k, e^+) P(y_k | x_k, e^+) P(x_k | e^+) * (P(e^+)/P(e))$$

The first term is just $\lambda(x_k)$. The second is the data term, $P(y_k|x_k)$. The last is the normalizing constant. So we just need to figure out how to compute:

$$P(x_k | e^+) \stackrel{\text{def}}{=} \pi(x_k)$$

This is almost the same deal as $\lambda(x_k)$. We get:

$$\pi(x_k) = \sum_{x_{k+1}} P(x_k | x_{k+1}, e^+) P(x_{k+1} | e^+) P(y_{k+1} | x_{k+1})$$

Again, the main computation is $M_{\{x_k | x_{k+1}\}} \pi(x_{k+1})$.

So our overall strategy is to pass information to the left and right, adding evidence as we go, and finally combine these by multiplying evidence from the left and right.