LifeNet Belief Propagation

Bo Morgan

2004/01/31

5 Node M.R.F. with 3 Cliques

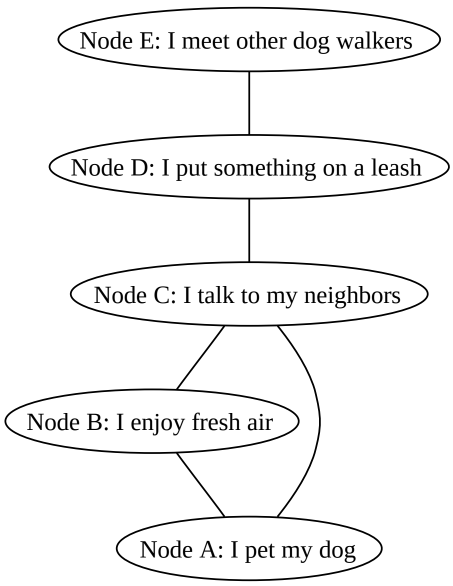

Figure 1: Sample miniature LifeNet graph.

1 Belief Propagation

Each probabilistic relationship between LifeNet phrases exists as a tabular probability distribution. These cliques relate the nodes within the markov field. We will refer to these cliques as ψ_i for $i = \{1, 2 \ldots, C\}$ when C is the number of cliques within LifeNet. ψ_i is defined in terms of the probability distribution of the set of variables within that clique, ψ_{iX}. The current version of LifeNet assumes the independence of these probability distributions, and this is what allows us to write the potential factors, ψ_i, in terms of these simple probabilities:

$$\psi_i = P(\psi_{iX}). \qquad (1)$$

LifeNet will learn from the experience gained from real world experience data such as the Diary application, which will provide a rich source of temporal event sequences and concurrences.

Probability Distribution Table

(M.R.F. Clique Potential)

A	**B**	**C**	**P(A,B,C)**
0	0	0	0.5
0	0	1	0.1
0	1	0	0.025
0	1	1	0.05
1	0	0	0.125
1	0	1	0.05
1	1	0	0.05
1	1	1	0.1

C	**D**	**P(C,D)**
0	0	0.5
0	1	0.1
1	0	0.025
1	1	0.05
0	0	0.125
0	1	0.05
1	0	0.05
1	1	0.1

D	**E**	**P(D,E)**
0	0	0.5
0	1	0.1
1	0	0.025
1	1	0.05
0	0	0.125
0	1	0.05
1	0	0.05
1	1	0.1

Figure 2: Samples of LifeNet tabular potential factors, ψ_i.

LifeNet factors, ψ_i, are tabular functions of the states of those factors, ψ_{iX}. Three sample tabular potential factors are shown in figure 2.

LifeNet's belief propagation algorithm accepts evidence, E, for the probability of a subset of LifeNet's nodes. Given this evidence, belief propagation can efficiently estimate the probabilities of the remaining nodes. Let ξ^0 be the initial estimate of $P(X|E)$, which is the initial state of the iterating belief propagation algorithm. Within LifeNet, we assume $\xi_i^0 = 0.5$ for all $i \in \{1, 2, \ldots, N\}$, where N is the total number of nodes within LifeNet. Our purpose for using the belief propagation algorithm is that it is an efficient albeit unreliable method of iteratively calculating the following limit:

$$\lim_{k \to \infty} \xi^k = P(X|E) \tag{2}$$

For each node, X_i, we find a new estimate of $P(X_i|E)$, based on the current probability estimates, ξ_i^k, which gives us ξ_i^{k+1}. At each iteration, the probabilities for the nodes within the *markov blanket* for each node is assumed to be equal to the most recent probability estimates for those nodes in the blanket.

The *markov blanket* [1] for a node, X_i, in LifeNet, or any M.R.F., is equal to the set of cliques that contain that node. The subset of all cliques, ψ, that contain a node, X_i, is the markov blanket, $X_{i\beta}$, of that node:

$$X_{i\beta} = \{\psi : X_i \in \psi\}. \tag{4}$$

The markov blanket of X_i is the minimal set of nodes that when known, effectively make $P(X_i)$ independent from any other evidence within the network.

The clique potential functions, $\psi(\Lambda)$, are tabular as shown in figure 2. For each potential function $\psi_i(\Lambda)$, the domain, Λ is not a binary space but is instead a bounded real space such that

$$\Lambda \in [0 - 1]^{|\psi_i|}, \tag{5}$$

where $|\psi_i|$ is the dimensionality of the clique, ψ_i. This function is calculated by making a weighted sum of every tabular entry in the potential. The linear weighting is equal to the probability of that entry being true, given the domain, Λ.

The belief propagation algorithm uses the potential functions by setting the domain at iteration, k, to be

$$\Lambda^k = (E, \xi^k) \tag{6}$$

The iterative algorithm for updating the probability estimates, ξ_i, for each of the nodes is

$$\xi_i^{k+1} = \prod_{\psi \in X_{i\beta}} \psi(\Lambda^k), \text{ for } i \in \{1, 2, \ldots, N\}. \tag{7}$$

[1] *markov blanket*: In this paper, we refer to the markov blanket at the set of cliques that a node belongs to because this is easier within the M.R.F. framework, but in general the markov blanket is referred to as the set of nodes that are contained within these cliques. Or more generally, the set of nodes when whose probabilities are known fully specify the probability of that node such that

$$P(X_i|X_{i\beta}) = P(X_i|X_{i\beta}, E) \tag{3}$$

for any evidence, E.

References

[1] Push Singh and William Williams "LifeNet: A Propositional Model of Human Activity" Media Lab, M.I.T.; K-CAP'03, October 23-25, 2003

[2] Daphne Coller and Nir Friedman "Probabilistic Methods" Institution; Proceeding, Date