

Correctness of Belief Propagation in Gaussian Graphical Models of Arbitrary Topology

Yair Weiss

Computer Science Division, University of California, Berkeley, CA 94720-1776, U.S.A.

William T. Freeman

Mitsubishi Electric Research Labs, Cambridge, MA 02139, U.S.A.

Graphical models, such as Bayesian networks and Markov random fields, represent statistical dependencies of variables by a graph. Local “belief propagation” rules of the sort proposed by Pearl (1988) are guaranteed to converge to the correct posterior probabilities in singly connected graphs. Recently, good performance has been obtained by using these same rules on graphs with loops, a method we refer to as *loopy belief propagation*. Perhaps the most dramatic instance is the near Shannon-limit performance of “Turbo codes,” whose decoding algorithm is equivalent to loopy propagation.

Except for the case of graphs with a single loop, there has been little theoretical understanding of loopy propagation. Here we analyze belief propagation in networks with arbitrary topologies when the nodes in the graph describe jointly gaussian random variables. We give an analytical formula relating the true posterior probabilities with those calculated using loopy propagation. We give sufficient conditions for convergence and show that when belief propagation converges, it gives the correct posterior means for all graph topologies, not just networks with a single loop.

These results motivate using the powerful belief propagation algorithm in a broader class of networks and help clarify the empirical performance results.

1 Introduction ---

Problems involving probabilistic belief propagation arise in a wide variety of applications, including error correcting codes, speech recognition, and image understanding. Typically, a probability distribution is assumed over a set of variables, and the task is to infer the values of the unobserved variables given the observed ones. The assumed probability distribution is described using a graphical model (Lauritzen, 1996); the qualitative aspects of the distribution are specified by a graph structure. The graph may be either directed, as in a Bayesian network (Pearl, 1988; Jensen, 1996), or undirected, as in a Markov random field (Pearl, 1988; Geman & Geman, 1984).

If the graph is singly connected (there is only one path between any two given nodes), then there exist efficient local message-passing schemes to calculate the posterior probability of an unobserved variable given the observed variables. Pearl (1988) derived such a scheme for singly connected Bayesian networks and showed that this “belief propagation” algorithm is guaranteed to converge to the correct posterior probabilities (or “beliefs”). However, as Pearl noted, the same algorithm is not guaranteed to converge in multiply connected networks, and even if it does, it will not calculate the correct beliefs.

Several groups have recently reported excellent experimental results by running algorithms equivalent to Pearl’s algorithm on networks with loops (Frey, 1998; Murphy, Weiss, & Jordan, 1999; Freeman, Pasztor, & Carmichael, 2000). Perhaps the most dramatic instance of this performance is in an error-correcting code scheme known as Turbo codes (Berrou, Glavieux, & Thitimajshima, 1993). These codes have been described as “the most exciting and potentially important development in coding theory in many years” (McEliece, Rodemich, & Cheng, 1995) and have recently been shown (Kschischang & Frey, 2000; McEliece & MacKay, 1998) to use an algorithm equivalent to belief propagation in a network with loops. Although there is widespread agreement in the coding community that these codes “represent a genuine, and perhaps historic, breakthrough” (McEliece et al., 1995), a theoretical understanding of their performance has yet to be achieved.

Progress in the analysis of loopy belief propagation has been made for the case of networks with a single loop (Weiss, 1997, 2000; Forney, Kschischang, Marcus, & Tuncel, 2000; Aji, Horn, & McEliece, 1998). For these networks, it can be shown that:

- Unless all the compatibilities are deterministic, loopy belief propagation will converge.
- An analytic expression relates the correct marginals to the loopy marginals. The approximation error is related to the convergence rate of the messages; the faster the convergence, the more exact the approximation.
- If the hidden nodes are binary, then both the loopy beliefs and the true beliefs are maximized by the same assignments, although the confidence in that assignment is wrong for the loopy beliefs.

In this article we analyze belief propagation (BP) in graphs of arbitrary topology where the nodes describe jointly gaussian random variables. We show that if BP converges, then it will give the correct posterior means for all graph topologies, not just networks with a single loop. This result holds for vector-valued or scalar-valued nodes and under arbitrary message scheduling. In particular, it holds for Turbo decoding when the densities are assumed to be gaussian.

We also analyze the dynamics of BP under a synchronous message-passing protocol. We derive an exact relationship between the means and




Figure 1: Any Bayesian network can be converted into an undirected graph with pairwise cliques by adding cluster nodes for all parents that share a common child. (a) A Bayesian network. (b) The corresponding undirected graph with pairwise cliques. A cluster node for (x_2, x_3) has been added. The potentials can be set so that the joint probability in the undirected network is identical to that in the Bayesian network. In this case, the update rules presented in this article reduce to Pearl's propagation rules in the original Bayesian network (Weiss, 2000).

covariances estimated using BP at every iteration and the correct means and covariances. The covariance estimates will generally be incorrect, but we present a relationship between the error in the covariance estimates and the convergence speed.

2 Belief Propagation in Undirected Graphical Models

Pearl's original algorithm was described for directed graphs, but in this article, we focus on undirected graphs. Every directed graphical model can be transformed into an undirected graphical model before doing inference (see Figure 1). An undirected graphical model (or a Markov random field, MRF) is a graph in which the nodes represent variables and arcs represent compatibility constraints between them. Assuming all probabilities are nonzero, the Hammersley-Clifford theorem (e.g., Pearl, 1988) guarantees that the probability distribution will factorize into a product of functions of the maximal cliques of the graph.

Denoting by x the values of all unobserved variables in the graph, the factorization has the form

$$P(x) = \prod_c \Psi_c(x_c), \quad (2.1)$$

where x_c is a subset of x that form a clique in the graph and Ψ_c is the potential function for the clique.

We will assume, without loss of generality, that each x_i node has a corresponding y_i node that is connected only to x_i . Thus,

$$P(x, y) = \prod_c \Psi_c(x_c) \prod_i \Psi_{ii}(x_i, y_i). \quad (2.2)$$

The restriction that all the y_i variables are observed and none of the x_i variables are is just to make the notation simple; $\Psi_{ii}(x_i, y_i)$ may be independent of y_i (equivalent to y_i being unobserved) or $\Psi_{ii}(x_i, y_i)$ may be $\delta(x_i - x_o)$ (equivalent to x_i being observed, with value x_o).

In describing and analyzing belief propagation, we assume the graphical model has been preprocessed so it becomes a pairwise MRF—an MRF in which the joint probability factorizes into a product of bivariate potentials (potentials involving only two variables). Any graphical model can be converted into this form before doing inference through a suitable clustering of nodes into large nodes (Weiss, 2000). Figure 1 shows an example; a Bayesian network is converted into an MRF in which all the cliques are pairs of units and the probability factorizes into a product of bivariate potentials. The formalism of “factor graphs” (Frey, Koetter, & Vardy, 1998) is more explicit than the MRF formalism in representing the factorization of the joint distribution. In Weiss and Freeman (2001) we show how any factor graph can be converted into a pairwise MRF and that when this conversion is performed, running BP on the pairwise MRF is equivalent to running probability propagation on the factor graph.

Equation 2.2 becomes

$$P(x, y) = \prod_{i,j} \Psi_{ij}(x_i, x_j) \prod_i \Psi_{ii}(x_i, y_i), \quad (2.3)$$

where the first product is over connected pairs of nodes.

By preprocessing the graph into one with pairwise cliques, the description and the analysis of belief propagation become simpler. For completeness, we review the belief propagation scheme used in Weiss (2000).

At every iteration, each node sends a (different) message to each of its neighbors and receives a message from each neighbor. Let x_i and x_j be two neighboring nodes in the graph. We denote by $m_{ij}(x_j)$ the message that node x_i sends to node x_j , and by $b_i(x_i)$ the belief at node x_i .

The belief update (or “sum-product” update) rules are:

$$m_{ij}(x_j) \leftarrow \alpha \int_{x_i} \Psi_{ij}(x_i, x_j) \Psi_{ii}(x_i, y_i) \prod_{x_k \in N(x_i) \setminus x_j} m_{ki}(x_i) \quad (2.4)$$

$$b_i(x_i) \leftarrow \alpha \Psi_{ii}(x_i, y_i) \prod_{x_k \in N(x_i)} m_{ki}(x_i), \quad (2.5)$$

where α denotes a normalization constant and $N(x_i) \setminus x_j$ means all nodes neighboring x_i , except x_j .

The procedure is initialized with all message vectors set to constant functions. The normalization of m_{ij} in equation 2.4 is not necessary; whether or not the message are normalized, the belief b_i will be identical. However, normalizing the messages avoids numerical underflow and adds to the stability of the algorithm. We assume throughout this article that all nodes simultaneously update their messages in parallel.

It is easy to show that for singly connected graphs, these updates will converge in a number of iterations equal to the diameter of the graph and the beliefs are guaranteed to give the correct posterior marginals: $b_i(x_i) = \text{Prob}(X_i = x_i \mid Y)$ where Y denotes the set of observed variables.

We emphasize that this transformation of a graphical model into a pairwise MRF is simply to make the analysis simpler. The belief propagation scheme described here includes as special cases Pearl’s (1988) BP algorithm in directed graphs, the generalized distributive law algorithm of Aji and McEliece (1999), and the factor graph propagation algorithm of Frey et al. (1998). These three algorithms correspond to the algorithm presented here with a particular way of preprocessing the graph in order to obtain pairwise potentials (Weiss & Freeman, 2001). Thus, our analysis of BP holds for these alternative formulations of BP as well (see section 5.1 for an explicit discussion of vector-valued nodes that will generally arise as a result of clustering).

For particular graphs with particular settings of the potentials, equations 2.4 and 2.5 yield other well-known Bayesian inference algorithms, such as the forward-backward algorithm in hidden Markov models, the Kalman filter, and even the fast Fourier transform (Aji & McEliece, 1999; Kschischang & Frey, 2000).

A related algorithm, max-product, changes the integration in equation 2.4 to a maximization. This message passing is equivalent to Pearl’s “belief revision” algorithm in directed graphs. For particular graphs with particular settings of the potentials, the max-product algorithm is equivalent to the Viterbi algorithm for hidden Markov models and concurrent dynamic programming. We define the max-product assignment at each node to be the value that maximizes its belief (assuming a unique maximizing value exists). For singly connected graphs, the max-product assignment is guaranteed to give the maximum a posteriori (MAP) assignment.

2.1 Gaussian Markov Random Fields. A gaussian MRF (GMRF) is an MRF in which the joint distribution is gaussian. We assume, without loss of generality, that the joint mean is zero (the means can be added-in later), so the joint distribution of $z = \begin{pmatrix} x \\ y \end{pmatrix}$ is given by:

$$Prob(z) = \alpha e^{-\frac{1}{2}z^T V z}, \quad (2.6)$$

where α is a normalization constant and V has the following structure:

$$V = \begin{pmatrix} V_{xx} & V_{xy} \\ V_{yx} & V_{yy} \end{pmatrix}. \quad (2.7)$$

It is straightforward to write the inverse covariance matrix describing the GMRF which respects the statistical dependencies within the graphical model (Cowell, 1998).

Again we assume that the graph has been preprocessed so that the joint distribution factorizes as a product of pairwise potentials. Thus, there exist matrices V_{ij} , one corresponding to each connected pair of nodes and matrices V_{ii} corresponding to each node such that

$$Prob(z) = \alpha \prod_{i,j} e^{-\frac{1}{2}(x_i x_j) V_{ij} (x_i x_j)^T} \prod_i e^{-\frac{1}{2}(x_i y_i) V_{ii} (x_i y_i)^T}, \quad (2.8)$$

where the first product is again over connected pairs of nodes. We will assume that the preprocessing is such that V_{ij} are all positive semidefinite matrices.

Note that the decomposition of V into V_{ij} , V_{ii} is not unique. For scalar nodes, any set of V_{ij} , V_{ii} that satisfies the following constraints are valid:

$$V_{xy}(i, i) = V_{ii}(1, 2) \quad (2.9)$$

$$V_{xx}(i, j) = V_{ij}(1, 2) \quad (2.10)$$

$$V_{xx}(i, i) = V_{ii}(1, 1) + \sum_{x_j \in N(x_i)} V_{ij}(1, 1). \quad (2.11)$$

2.2 Exact Inference in Gaussian MRFs. Writing out the exponent of 2.6 and completing the square shows that the mean μ of x , given y , is a solution to:

$$V_{xx}\mu = -V_{xy}y, \quad (2.12)$$

and the covariance matrix $C_{x|y}$ of x given y is

$$C_{x|y} = V_{xx}^{-1}. \quad (2.13)$$

We will denote by $C_{x_i|y}$ the i th row of $C_{x|y}$, so the marginal posterior variance of x_i , given the data, is $C_{x_i|y}(i)$.

2.3 Belief Propagation for MRFs. Belief propagation in MRFs gives simpler update formulas than the general case (see eqs. 2.4 and 2.5). The messages and the beliefs are all gaussians, and the updates can be written directly in terms of the means and inverse covariance matrices. Each node sends and receives a mean vector and inverse covariance matrix to and from each neighbor, in general, each different.

To write the updates explicitly in terms of means and covariances, we denote by μ_{ij} the mean of the message that node x_i sends to node x_j and by P_{ij} the precision or inverse covariance matrix that node x_i sends to node x_j . Similarly, we denote by μ_i the mean of the belief at node x_i and by P_i the inverse covariance matrix of the belief. As before, we use μ_{ii}, P_{ii} for the mean and inverse covariance of $\Psi_{ii}(x_i, y_i)$. We partition the matrix V_{ij} into

$$V_{ij} = \begin{pmatrix} a & b \\ b^T & c \end{pmatrix}. \quad (2.14)$$

The message update rules are

$$P_{ij} \leftarrow c - b(a + P_0)^{-1}b^T \quad (2.15)$$

$$\mu_{ij} \leftarrow -P_{ij}^{-1}b(a + P_0)^{-1}P_0\mu_0, \quad (2.16)$$

with

$$P_0 = P_{ii} + \sum_{x_k \in N(x_i) \setminus x_j} P_{ki} \quad (2.17)$$

$$\mu_0 = P_0^{-1} \left(P_{ii}\mu_{ii} + \sum_{x_k \in N(x_i) \setminus x_j} P_{ki}\mu_{ki} \right). \quad (2.18)$$

The beliefs are given by

$$P_i \leftarrow P_{ii} + \sum_{x_k \in N(x_i)} P_{ki} \quad (2.19)$$

$$\mu_i \leftarrow P_i^{-1} \left(P_{ii}\mu_{ii} + \sum_{x_k \in N(x_i)} P_{ki}\mu_{ki} \right). \quad (2.20)$$

We can now state the main question of this article. What is the relationship between the true posterior means and covariances (calculated using equation 2.12) and the belief propagation means and covariances (calculated using the belief propagation rules equations 2.15–2.20)?

3 Dynamics of Belief Propagation

In this section we analyze the dynamics of BP in a gaussian pairwise MRF. To simplify the notation, we assume throughout this section that all nodes are scalar valued. In section 5.1 we generalize the analysis to vector-valued nodes.

To compare the correct posteriors and the loopy beliefs, we construct an unwrapped tree. The unwrapped tree is the graphical model that the loopy belief propagation is solving exactly when applying the belief propagation rules in a loopy network (Gallager, 1963; Wiberg, 1996; Weiss, 2000). In error-correcting codes, the unwrapped tree is referred to as the computation tree; it is based on the idea that the computation of a message sent by a node at time t depends on messages it received from its neighbors at time $t - 1$ and those messages depend on the messages the neighbors received at time $t - 2$, and so forth.

To construct the topology of the unwrapped tree, set an arbitrary node, say x_1 , to be the root node, and then iterate the following procedure t times:

1. Find all leaves of the tree (start with the root).
2. For each leaf, find all k nodes in the loopy graph that neighbor the node corresponding to this leaf.
3. Add $k - 1$ nodes as children to each leaf, corresponding to all neighbors except the parent node.

The potential matrices and observations for each node in the unwrapped network are copied from the corresponding nodes in the loopy graph. Each node in the loopy graph will have a different unwrapped tree with that node at the root.

Figure 2 shows an unwrapped tree around node x_1 for the diamond-shaped graph on the left. Each node has a shaded observed node attached to it that is not shown for clarity. Since belief propagation is exact for the unwrapped tree, we can calculate the beliefs in the unwrapped tree by using the marginalization formulas for gaussians.

We use $\tilde{\cdot}$ for unwrapped quantities. We scan the tree in breadth-first order and denote by \tilde{x} the vector of values in the hidden nodes of the tree when scanned in this fashion. Similarly, we denote by \tilde{y} the observed nodes scanned in the same order. As before, $\tilde{z} = \begin{pmatrix} \tilde{x} \\ \tilde{y} \end{pmatrix}$.

The basic idea behind our analysis is to relate the wrapped and unwrapped inverse covariance matrices. From equation 2.11, all elements $\tilde{V}_{xy}(i', j')$ and $\tilde{y}(i')$ are copies of the corresponding elements $V_{xy}(i, j)$ and $y(i)$ (where \tilde{x}_i, \tilde{x}_j are replicas of x_i, x_j). Also, all elements $\tilde{V}_{xx}(i', j')$ are copies of $V_{xx}(i, j)$, and elements $\tilde{V}_{xx}(i', i')$ for non-leaf nodes are replicas of $V_{xx}(i, i)$. However, the elements $\tilde{V}_{xx}(i', i')$ for the leaf nodes are not copies of $V_{xx}(i, i)$ because the leaf nodes are missing some neighbors.




Figure 2: (Left) A Markov network with multiple loops. (Right) The unwrapped network corresponding to this structure. The unwrapped networks are constructed by replicating the potentials $\Psi_{ij}(x_i, x_j)$ and observations y_i while preserving the local connectivity of the loopy network. They are constructed so that the messages received by node A after t iterations in the loopy network are equivalent to those that would be received by A in the unwrapped network. An observed node, y_i , not shown, is connected to each depicted node.

Intuitively, we might expect that if all the equations that $\tilde{\mu}$ satisfies are copies of the equations that μ satisfies, then simply creating $\tilde{\mu}$ by many copies of μ would give a valid solution in the unwrapped network. However, because some of the equations are not copies, this intuition does not explain why the means are exact in gaussian networks.

An additional intuition, which we formalize below, is that the influence of the noncopied equations (those at the leaf nodes) decreases with additional iterations. As the number of iterations is increased, the distance between the leaf nodes and the root node increases and their influence on the root node decreases. When their influence goes to zero, the mean at the root node is exact.

Although the elements $\tilde{V}_{xx}(i', j')$ are copies of $V_{xx}(i, j)$ for the nonleaf nodes, the matrix \tilde{V}_{xx} is not simply a block replication of V_{xx} . The system of equations that defines $\tilde{\mu}$ is a coupled system of equations. Hence the variance at the root node $\tilde{V}_{xx}^{-1}(1, 1)$ differs from the correct variance $V_{xx}^{-1}(1, 1)$.

In the following section, we prove the following four claims regarding the dynamics of loopy belief propagation in gaussian graphical models.

Assume, without loss of generality, that the root node is x_1 . Let $\tilde{\mu}(1)$ and $\tilde{\sigma}^2(1)$ be the conditional mean and variance at node 1 after t iterations of loopy propagation. Let $\mu(1)$ and $\sigma^2(1)$ be the correct conditional mean and variance of node 1. Let $\tilde{C}_{x_1|y}$ be the conditional correlation of the root node

with all other nodes in the unwrapped tree; then:

Claim 1:

$$\tilde{\mu}(1) = \mu(1) + \tilde{C}_{x_1|y}r, \quad (3.1)$$

where r is a vector that is zero everywhere but the last L components (corresponding to the leaf nodes).

Claim 2:

$$\tilde{\sigma}^2(1) = \sigma^2(1) + \tilde{C}_{x_1|y}r_1 - \tilde{C}_{x_1|y}r_2, \quad (3.2)$$

where r_1 is a vector that is zero everywhere but the last L components and r_2 are equal to 1 for all components corresponding to nonroot nodes in the unwrapped tree that reference x_1 . All other components of r_2 are zero.

Claim 3: If the conditional correlation between the root node and the leaf nodes decreases rapidly enough, then (1) belief propagation converges, (2) the belief propagation means are exact, and (3) the belief propagation variances are equal to the correct variances minus the summed conditional correlations between \tilde{x}_1 and all \tilde{x}_j that are replicas of x_1 .

Claim 4: Assume all V_{ij} are diagonally dominant. Then (1) belief propagation converges, (2) the belief propagation means are exact, and (3) the belief propagation variances are equal to the correct variances minus the summed conditional correlations between \tilde{x}_1 and all \tilde{x}_j that are replicas of x_1 .

To obtain intuition, Figure 3 shows $\tilde{C}_{x_1|y}$ for the diamond figure in Figure 2. We generated random potential functions and observations for the loopy diamond figure and calculated the conditional correlations in the unwrapped network. Note that the conditional correlation decreases with distance in the tree; we are scanning in breadth-first order so the last L components correspond to the leaf nodes. As the number of iterations of loopy propagation is increased, the size of the unwrapped tree increases, and the conditional correlation between the leaf nodes and the root node decreases.

From equations 3.1 and 3.2, it is clear that if the conditional correlations between the leaf nodes and the root nodes are zero for all sufficiently large unwrappings, then (1) belief propagation converges, (2) the means are exact, and (3) the belief propagation variances are equal to the correct variances minus the summed conditional correlations between \tilde{x}_1 and all \tilde{x}_j that are replicas of x_1 . In practice, the conditional correlations will not actually be equal to zero for any finite unwrapping, so claim 3 states this more precisely. Claim 4 gives sufficient conditions, in terms of the V_{ij} matrices, for the conditional correlation to decrease rapidly enough.




Figure 3: The conditional correlation between the root node and all other nodes in the unwrapped tree for the diamond figure after seven iterations. Potentials were chosen randomly. Nodes are presented in breadth-first order, so the last elements are the correlations between the root node and the leaf nodes. It can be proven that if this correlation goes to zero, then (1) belief propagation converges, (2) the loopy means are exact, and (3) the loopy variances equal the correct variances minus the summed conditional correlation of the root node and all other nodes that are replicas of the same loopy node. Symbols plotted with a star denote correlations with nodes that correspond to the node x_1 in the loopy graph. We prove that the sum of these correlations gives the correct variance of node x_1 while loopy propagation uses only the first correlation.

How wrong will the variances be? The term $\tilde{C}_{x_1|y}r_2$ in equation 3.2 is simply the sum of many components of $\tilde{C}_{x_1|y}$. Figure 3 shows these components. The correct variance is the sum of all the components, while the loopy variance approximates this sum with the first (and dominant) term.

Note that when the conditional correlation decreases rapidly to zero, two things happen: the convergence is faster (because $\tilde{C}_{x_1|y}r_1$ approaches zero faster), and the approximation error of the variances is smaller (because $\tilde{C}_{x_1|y}r_2$ is smaller). Thus, as in the single-loop case, we find that quick convergence is correlated with good approximation.

3.1 Relation of Loopy and Unwrapped Quantities. The proof of the claims relies on the relationship between the elements of y , V_{xy} and V_{xx} with their unwrapped quantities, described below.

Each node in \tilde{x} corresponds to a node in the original loopy network. Let O be a matrix that defines this correspondence. $O(i, j) = 1$ if \tilde{x}_i corresponds to x_j and zero otherwise. Thus, in Figure 2, ordering the nodes alphabetically,

the first rows of O are:

$$O = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ \dots & \dots & \dots & \dots & \dots \end{pmatrix}. \quad (3.3)$$

Using O we can formalize the relationship between the unwrapped quantities and the original ones. The simplest one is \tilde{y} , which contains only replicas of the original y :

$$\tilde{y} = Oy. \quad (3.4)$$

Since every x_i is connected to a y_i , V_{xy} and \tilde{V}_{xy} are zero everywhere but along their diagonals (the block diagonals, for vector-valued variables). The diagonal elements of \tilde{V}_{xy} are simply replications of V_{xy} ; hence:

$$\tilde{V}_{xy}O = OV_{xy}. \quad (3.5)$$

\tilde{V}_{xx} also contains the elements of the original V_{xx} , but here special care needs to be taken. Note that by construction, every node in the interior of the unwrapped tree has exactly the same statistical relationship with its neighbors as with the corresponding node in the loopy graph. If a node in the loopy graph has k neighbors, then a node in the unwrapped tree will have, by construction, one parent and $k - 1$ children. The leaf nodes in the unwrapped tree, however, will be missing the $k - 1$ children and hence will not have the same number of neighbors. Thus, for all nodes \tilde{x}_i, \tilde{x}_j that are not leaf nodes, $\tilde{V}_{xx}(i, j)$ is a copy of the corresponding $V_{xx}(k, l)$, where unwrapped nodes i and j refer to loopy nodes k and l , respectively.

Therefore,

$$\tilde{V}_{xx}O + E = OV_{xx}, \quad (3.6)$$

where E is an error matrix. E is zero for all nonleaf nodes, so the first $N - L$ rows of E are zero.

3.2 Proof of Claim 1. The marginalization equation for the unwrapped problem gives

$$\tilde{V}_{xx}\tilde{\mu} = -\tilde{V}_{xy}\tilde{y}. \quad (3.7)$$

Substituting equations 3.4 and 3.5, relating loopy and unwrapped network quantities, into equation 3.7, for the unwrapped posterior mean, gives

$$\tilde{V}_{xx}\tilde{\mu} = -OV_{xy}y. \quad (3.8)$$

For the true means, μ , of the loopy network, we have

$$V_{xx}\mu = -V_{xy}y. \quad (3.9)$$

To relate that to the means of the unwrapped network, we left-multiply by O :

$$OV_{xx}\mu = -OV_{xy}y. \quad (3.10)$$

Using equation 3.6, relating V_{xx} to \tilde{V}_{xx} , we have

$$\tilde{V}_{xx}O\mu + E\mu = -OV_{xy}y. \quad (3.11)$$

Comparing equations 3.11 and 3.8 gives

$$\tilde{V}_{xx}O\mu + E\mu = \tilde{V}_{xx}\tilde{\mu} \quad (3.12)$$

or

$$\tilde{\mu} = O\mu + \tilde{V}_{xx}^{-1}E\mu. \quad (3.13)$$

Using equation 2.13,

$$\tilde{\mu} = O\mu + \tilde{C}_{x|y}E\mu. \quad (3.14)$$

The left- and right-hand sides of equation 3.14 are column vectors. We take the first component of both sides and get

$$\tilde{\mu}(1) = \mu(1) + \tilde{C}_{x_1|y}E\mu. \quad (3.15)$$

Since E is zero in the first $N - L$ rows, $E\mu$ is zero in the first $N - L$ components.

3.3 Proof of Claim 2. From equation 2.13,

$$V_{xx}C_{x|y} = I. \quad (3.16)$$

Taking the first column of this equation gives

$$V_{xx}C_{x_1|y}^T = e_1, \quad (3.17)$$

where $e_1(1) = 1$, $e_1(j > 1) = 0$.

Using the same strategy as in the previous proof, we left-multiply by O ,

$$OV_{xx}C_{x_1|y}^T = Oe_1, \quad (3.18)$$

and similarly we substitute equation 3.6:

$$\tilde{V}_{xx}OC_{x_1|y}^T + EC_{x_1|y}^T = Oe_1. \quad (3.19)$$

The analog of equation 3.17 in the unwrapped problem is

$$\tilde{V}_{xx}\tilde{C}_{x_1|y}^T = \tilde{e}_1, \quad (3.20)$$

where $\tilde{e}_1(1) = 1$, $\tilde{e}_1(j > 1) = 0$.

Subtracting equations 3.19 and 3.20 and rearranging terms gives

$$\tilde{C}_{x_1|y} = OC_{x_1|y}^T + \tilde{V}_{xx}^{-1}EC_{x_1|y}^T + \tilde{V}_{xx}^{-1}(\tilde{e}_1 - Oe_1). \quad (3.21)$$

Again, we take the first row of both sides of equation 3.21 and use the fact that the first row of \tilde{V}_{xx}^{-1} is $\tilde{C}_{x_1|y}$ to obtain

$$\tilde{\sigma}^2(1) = \sigma^2(1) + \tilde{C}_{x_1|y}EC_{x_1|y}^T + \tilde{C}_{x_1|y}(\tilde{e}_1 - Oe_1). \quad (3.22)$$

Again, since E is zero in the first N_L rows, $EC_{x_1|y}$ is zero in the first $N - L$ components.

3.4 Proof of Claim 3. Here we need to define what we mean by “rapidly enough.” We restate the claim precisely:

Claim 3: Suppose for every ϵ there exists a t_ϵ such that for all $t > t_\epsilon$ $|\tilde{C}_{x_1|y}r| < \epsilon \max_i |r(i)|$ for any vector r that is nonzero only in the last L components (those corresponding to the leaf nodes). In this case, (1) belief propagation converges, (2) the means are exact, and (3) the variances are equal to the correct variances minus the summed conditional correlations between \tilde{x}_1 and all nonroot \tilde{x}_j that are replicas of x_1 .

This claim follows from the first two claims. The only thing to show is that $E\mu$ and $EC_{x_1|y}$ are bounded for all iterations. This is true because the rows of E are bounded and μ , $C_{x_1|y}$ do not depend on the iteration.

3.5 Proof of Claim 4. The proof is based on the following lemma:¹

Conditional correlation lemma. Assume $P(x, y) = \alpha e^{-\frac{1}{2}z^T Vz}$ with z, V as in equations 2.6 and 2.7. Let r be an arbitrary vector. Then $C_{x|y}r$ can be calculated by (1) modifying the joint probability so that $V_{xy} = -I$, (2) setting the observations $y = r$, and (3) calculating the posterior means of x given this y in the modified joint.

Proof. This follows directly from equations 2.12 and 2.13.

Using this lemma, we can give sufficient conditions for convergence in terms of the potentials of the loopy network.

Claim 4: Assume all V_{ij} are diagonally dominant ($|V_{ij}(k, l)| < V_{ij}(k, k)$) and all $V_{ii}(1, 1) > 0$. Then (1) belief propagation converges, (2) the belief propagation means are exact, and (3) the belief propagation variances are equal to the correct variances minus the summed conditional correlations between \tilde{x}_1 and all \tilde{x}_j that are replicas of x_1 .

Proof. The proof is based on claim 3 and the conditional correlation lemma. From claim 3, we know that it is sufficient to show that for any ϵ and r that is nonzero only at the leaves, $|\tilde{C}_{x_1|y}r| < \epsilon \max_i |r(i)|$ for sufficiently large unwrappings. By the conditional correlation lemma, we know we can calculate $|\tilde{C}_{x_1|y}r|$ by constructing an unwrapped tree with observations r and computing the conditional mean at the root node. Since the unwrapped tree is singly connected, we can calculate the mean at the root node exactly by running belief propagation. We start by sending messages from the leaf nodes to the layer above and continue sending messages upward until we reach the root node.

We use M_l to denote the maximum absolute value of means of all messages sent upward by layer l . By equation 2.18, the value μ_0 calculated by all nodes at layer $l - 1$ is a weighted average of means from the previous layer and the means from the observations at layer $l - 1$. Since r is zero for all layers but the bottom one, the means of the messages sent by observations at the $l - 1$ layer are zero. Thus:

$$|\mu_0| \leq M_l. \quad (3.23)$$

If we rewrite equation 2.16 taking advantage of the fact that x_i, x_j are scalars, we have

$$\mu_{ij} = \frac{-b(a + P_0)^{-1}P_0\mu_0}{c - b^2(a + P_0)^{-1}}. \quad (3.24)$$

¹ We thank Andrew Ng for suggesting this proof.

Multiplying top and bottom by $\frac{(a+P_0)}{P_0}$,

$$\mu_{ij} = \frac{-b}{(ac - b^2)/P_0 + c} \mu_0. \quad (3.25)$$

Similarly, we can rewrite equation 2.15 taking advantage of x_i, x_j being scalar to give

$$P_{ij} = \frac{ca - b^2 + cP_0}{a + P_0}. \quad (3.26)$$

Note that for diagonally dominant matrices, P_{ij} is nonnegative if P_0 is nonnegative. Note also that if $V_{ii}(1, 1) > 0$ for all the observation potentials, the observations will send a positive precision to the unobserved nodes. Thus, all precisions will be nonnegative.

Since V_{ij} is diagonally dominant and P_0 is nonnegative, both terms in the denominator of equation 3.25 are nonnegative, so

$$|\mu_{ij}| \leq \frac{|b|}{|c|} |\mu_0|. \quad (3.27)$$

We now denote by $\beta = \max_{i,j} |V_{ij}(1, 2)/V_{ij}(2, 2)|$. Since all V_{ij} are diagonally dominant, $\beta < 1$. Combining equations 3.23 and 3.27 gives

$$M_{l-1} \leq \beta M_l. \quad (3.28)$$

So for any ϵ if we choose $t_\epsilon > \frac{\log \epsilon}{\log \beta}$ then $|\tilde{C}_{x_i|y} r| < \epsilon \max |r(i)|$.

The conditional correlation lemma can also be used to give bounds on the loopy variances—for example:

Corollary. *If V_{ij} are diagonally dominant and the off-diagonal elements are negative, then the loopy beliefs are overconfident $\tilde{\sigma}^2(1) < \sigma^2(1)$.*

Proof. By the conditional correlation lemma, we can calculate $\tilde{C}_{x_i|y} r_2$ by setting the observations to be one at all copies of the root node and zero elsewhere. From equation 3.25 it is clear that when $b < 0$ and the observations are zero or one, the mean messages are weighted averages of positive values; hence, the mean at the root will be positive.

The decomposition of a GMRF into V_{ij}, V_{ii} is not unique. The following lemma shows that even when we have a loopy graph where V_{ij} are not diagonally dominant, belief propagation will still converge when there exists a reparameterization using diagonally dominant matrices.

Reparameterization Lemma. *If the unwrapped tree for any iteration of loopy propagation can be parameterized so that \tilde{V}_{ij} are diagonally dominant and $\tilde{V}_{ii}(1, 1) > 0$, then (1) belief propagation converges, (2) the belief propagation means are exact, and (3) the belief propagation variances are equal to the correct variances minus the summed conditional correlations between \tilde{x}_1 and all \tilde{x}_j that are replicas of x_1 .*

Proof. This follows from the proof of claim 4. Since the unwrapped tree is singly connected, belief propagation using any parameterization is exact. We can calculate $\tilde{C}_{x_1|y}r$ by running belief propagation on the reparameterized tree in which all the matrices are diagonally dominant.

We emphasize that claim 4 gives only sufficient conditions for convergence. It is easy to construct networks in which V_{ij} are not all diagonally dominant but loopy belief propagation still converges. In section 6 we show an example.

4 Fixed Points of Loopy Propagation

Each iteration of belief propagation can be thought of as an operator F that inputs a list of messages $m^{(t)}$ and outputs a list of messages $m^{(t+1)} = Fm^{(t)}$. Thus, belief propagation can be thought of as an iterative way of finding a solution to the fixed-point equations $Fm = m$ with an initial guess m_0 in which all messages are constant functions.

Note that this is not the only way of finding fixed points. McEliece et al. (1995) have shown a simple example for which F contains multiple fixed points and belief propagation finds only one. They also showed an example where a fixed point exists but the iterations $m \leftarrow Fm$ do not converge. Murphy et al. (1999) describe an alternative method for finding fixed points of F .

An alternative way of finding solutions to $Fm = m$ is to use an asynchronous update order, that is, update equations 2.4 in some fixed update order rather than the synchronous order we use here.

Update Order Lemma. *Let $F^A m$ be an updating of the messages in some fixed order so that all messages m_{ij} are updated sequentially using equation 2.4. Let Fm be an updating of the messages in parallel. A set of messages m satisfies $Fm = m$ if and only if they satisfy $F^A m = m$.*

Proof. By the definition of a fixed point, m is a fixed point of F^A or of F if and only if the left-hand side of equation 2.4 equals the right-hand side.

In this section we ask, suppose a fixed-point $m^* = Fm^*$ has been found by some method. How are the beliefs calculated based on these messages related to the correct beliefs?

Claim 5: For a gaussian graphical model of arbitrary topology, if m^* is a fixed point of the message-passing dynamics, then the means based on that fixed point are exact. The proof is based on the following lemma.

Periodic Beliefs Lemma. *If m^* is a fixed point of the message-passing dynamics in a graphical model G , then one can construct a modified unwrapped tree T of arbitrarily large depth such that (1) all nonleaf nodes in T have the same statistical relationship with their neighbors as the corresponding nodes in G and (2) all nodes in T will have the same belief as the beliefs in G derived from m^* .*

Proof. The proof is by construction. We first construct an unwrapped tree T of the desired depth. We then modify the potentials and the observations in the leaf nodes in the following manner. For each leaf node \tilde{x}_i , find the $k - 1$ nodes in G that neighbor $x_{\tilde{y}}$ (where \tilde{x}_i is a replica of $x_{\tilde{y}}$) excluding the parent of \tilde{x}_i . Calculate the product of the $k - 1$ messages that these neighbors send to the corresponding node in G under the fixed-point messages m^* and the message that $y_{\tilde{y}}$ sends to $x_{\tilde{y}}$. Set \tilde{y}_i and $\Psi(\tilde{y}_i, \tilde{x}_i)$ such that the message \tilde{y}_i sends to \tilde{x}_i is equal to this product.

By this construction, all leaf nodes in T will send their neighbors a message from m^* . Since all nonleaf nodes in T have the same statistical relationship to their neighbors as the corresponding nodes in G , the local message passing updates in T are identical to those in G . Thus, all messages in T will be replicas of messages in m^* .

Proof of Claim 5. Using this lemma, we can prove claim 5. Let $\tilde{\mu}$ be the conditional mean in the modified unwrapped tree. Then, by the periodic beliefs lemma,

$$\tilde{\mu} = O\mu_0, \quad (4.1)$$

where $\mu_0(i)$ is the posterior mean at node i under m^* .

We also know that $\tilde{\mu}$ is a solution to

$$\tilde{V}_{xx}\tilde{\mu} = -\tilde{V}_{xy}\tilde{y}, \quad (4.2)$$

where $\tilde{V}_{xx}, \tilde{V}_{xy}, \tilde{y}$ refers to quantities in the modified unwrapped tree. So

$$\tilde{V}_{xx}O\mu_0 = -\tilde{V}_{xy}\tilde{y}. \quad (4.3)$$

We use the notation $[A]_m$ to indicate taking the m first rows of a matrix A . Note that for any two matrices, $[AB]_m = [A]_m B$. Taking the first m rows of equation 4.3 gives

$$[\tilde{V}_{xx}O]_m\mu_0 = -[\tilde{V}_{xy}\tilde{y}]_m. \quad (4.4)$$

As in the previous proofs, the key idea is to relate the inverse covariance matrix of the modified unwrapped tree to that of the original loopy graph. Since all nonleaf nodes in the modified unwrapped tree have the same neighborhood relationships with their neighbors as the corresponding nodes in the loopy graph, we have, for any $m < N - L$,

$$[\tilde{V}_{xx}O]_m = [OV_{xx}]_m \quad (4.5)$$

and

$$[\tilde{V}_{xy}\tilde{y}]_m = [OV_{xy}y]_m. \quad (4.6)$$

Substituting these relationships into equation 4.4 gives

$$[O]_m V_{xx}\mu_0 = -[O]_m V_{xy}y. \quad (4.7)$$

This equation holds for any $m < N - L$. Since we can unwrap the tree to arbitrarily large size, we can choose m such that $[O]_m$ has n independent rows (this happens once all nodes in the loopy graph appear at least once in the modified unwrapped tree). Thus:

$$V_{xx}\mu_0 = -V_{xy}y. \quad (4.8)$$

Hence, the means derived from the fixed-point messages are exact.

5 Extensions

5.1 Vector-Valued Nodes. Most of the results we have derived so far hold for vector-valued nodes as well, but the indexing notation is rather more cumbersome. We use a stacking convention, in which we define the vector x by

$$x = \begin{pmatrix} x_1 \\ x_2 \\ \dots \end{pmatrix}. \quad (5.1)$$

Thus, supposing x_1 is a vector of length 2, then $x(1)$ is the first component of x_1 and $x(2)$ is the second component of x_1 (not x_2). We define y in a similar fashion.

Using this stacking notation, the equations for exact inference in gaussians remain unchanged, but we need to be careful in reading out the posterior means and covariances from the stacked vectors. Thus, we can still complete the square in stacked notation to obtain

$$V_{xx}\mu = -V_{xy}y \quad (5.2)$$

and $C_{x|y} = V_{xx}^{-1}$. Assuming x_1 is of length 2, μ_1 the posterior mean of x_1 is given by

$$\mu_1 = \begin{pmatrix} \mu(1) \\ \mu(2) \end{pmatrix}, \quad (5.3)$$

and the posterior covariance matrix Σ_1 is given by

$$\Sigma_1 = \begin{pmatrix} C_{x|y}(1, 1) & C_{x|y}(1, 2) \\ C_{x|y}(2, 1) & C_{x|y}(2, 2) \end{pmatrix}. \quad (5.4)$$

We use the same stacked notation for \tilde{x} and define the matrix O such that $O(i, j) = 1$ if $\tilde{x}(i)$ is a replica of $x(j)$ and zero otherwise. Using this notation, the relationships between unwrapped and loopy quantities (e.g., $[\tilde{V}_{xx}O]_m = [OV_{xx}]_m$) still hold. Thus, all the analysis done in the previous sections holds. The only difference is the semantics of quantities such as $\mu(1)$, which need to be understood as a scalar component of a (possibly) larger vector μ_1 . For explicitness, we restate the five claims for vector-valued nodes.

For any i, j less than or equal to the number of components in x_1 , we have:

Claim 1a:

$$\tilde{\mu}(i) = \mu(i) + \tilde{C}_{x|y}r, \quad (5.5)$$

where r is a vector that is zero everywhere but the last L components (corresponding to the leaf nodes).

Claim 2a:

$$\tilde{C}_{x|y}(i, j) = C_{x|y}(i, j) + \tilde{C}_{x|y}r_1 - \tilde{C}_{x|y}r_2, \quad (5.6)$$

where r_1 is a vector that is zero everywhere but the last L components (corresponding to the leaf nodes) and r_2 is equal to 1 for all components corresponding to nonroot nodes in the unwrapped tree that reference $x(i)$. All other components of r_2 are zero.

Claim 3a: If the conditional correlation between all components of the root node and the leaf nodes decreases rapidly enough, then (1) belief propagation converges, (2) the belief propagation means are exact, and (3) the i, j component of the belief propagation covariance matrices is equal to the i, j component of the true covariance matrices minus the summed conditional correlations between $\tilde{x}(j)$ and all nonroot $\tilde{x}(k)$ that are replicas of $x(i)$.

Claim 5a: For a (possibly vector-valued) gaussian graphical model of arbitrary topology, if m^* is a fixed point of the message-passing dynamics, then the means based on that fixed point are exact.

Claim 6a: For a (possibly vector-valued) graphical model of arbitrary topology with continuous potential functions, if m^* is a fixed point of the max-product message-passing dynamics, then the assignment based on that fixed point is a local maximum of the posterior probability.

We emphasize that these claims do not need to be reproven. All the equations used in proving the scalar-valued case still hold; only the semantics we place on the individual components are different.

We end this analysis with two simple corollaries:

Corollary 1. *Let m^* be a fixed point of Pearl's belief propagation algorithm on a gaussian Bayesian network of arbitrary topology and arbitrary clique size. Then the means based on m^* are exact.*

Corollary 2. *Let m^* be a fixed point of Pearl's belief revision (max-product) algorithm on a Bayesian network with continuous joint probability, arbitrary topology, and arbitrary clique size. The assignment based on m^* is at least a local maximum of the posterior probability.*

These corollaries follow from claims 5a and 6a, along with the equivalence between Pearl's propagation rules and the propagation rules for pairwise undirected graphical models analyzed here (Weiss, 2000). Note that even if the Bayesian network contained only scalar nodes, the conversion to pairwise cliques may necessitate using vector-valued nodes.

5.2 Non-gaussian Variables. In section 2, we described the max-product belief propagation algorithm that finds the MAP estimate for each node (Pearl, 1988; Weiss, 2000) of a network without loops. As with max-product, iterating this algorithm is a method of finding a fixed point of the message-passing dynamics. How does the assignment derived from this fixed point compare the MAP assignment?

Claim 6: For a graphical model of arbitrary topology with continuous potential functions, if m^* is a fixed point of the max-product message-passing dynamics, then the assignment based on that fixed point is a local maximum of the posterior probability.

Proof. Since the posterior probability factorizes into a product of pairwise potentials, the log posterior will have the form

$$\log P(x | y) = \sum_{ij} J_{ij}(x_i, x_j) + J_{ii}(x_i, y_i). \quad (5.7)$$

Assuming the clique potential functions are differentiable and finite, the MAP solution, u , will satisfy

$$\frac{\partial}{\partial x_i} \log P(x | y)|_{x=u} = 0. \quad (5.8)$$

We will write this as

$$Vu = 0, \quad (5.9)$$

where V is a nonlinear operator.

As in the previous section, we can use the periodic belief lemma to construct a modified unwrapped tree of arbitrary size based on m^* . If we denote by \tilde{V} the nonlinear set of equations that the solution to the modified unwrapped problem must satisfy, we have

$$\tilde{V}\tilde{u} = 0. \quad (5.10)$$

Because of the periodic belief lemma,

$$\tilde{u} = Ou_0. \quad (5.11)$$

Similarly, as in the previous section, all the nonleaf nodes will have the same statistical relationship with their neighbors as do the corresponding nodes in the loopy network, so

$$[\tilde{V}O]_m = [OV]_m, \quad (5.12)$$

where the left- and right-hand sides are nonlinear operators.

Substituting equations 5.11 and 5.12 into 5.10 gives

$$Vu_0 = 0. \quad (5.13)$$

A similar substitution can be made with the second derivative equations to show that the Hessian at u_0 is positive definite. Thus, the assignment based on m^* is at least a local maximum of the posterior.

Claim 6 can be extended to discrete-valued nodes as well. In Weiss and Freeman (2001) we prove that an assignment based on a fixed point of the max-product BP algorithm is a “neighborhood maximum” of the posterior probability. The posterior probability of the max-product assignment is guaranteed to be greater than all other assignments in a large neighborhood around that assignment. The neighborhood includes all assignments that differ from the max-product assignment in any subset of nodes that form no more than a single loop in the graph.

6 Simulations

To illustrate the analysis, we ran belief propagation on a 25×25 two-dimensional grid. The joint probability was

$$P(x, y) = \exp \left(- \sum_{ij} w_{ij}(x_i - x_j)^2 - \sum_i w_{ii}(x_i - y_i)^2 \right), \quad (6.1)$$

where $w_{ij} = 0$ if nodes x_i, x_j are not neighbors and 0.01 otherwise and w_{ii} was randomly selected to be 10^{-6} or 1 for all i with probability of 1 set to 0.2. When $w_{ii} = 1$, we set the observations y_i to be samples from the surface $z(x, y) = x + y$. When $w_{ii} = 10^{-6}$, we set $y_i = 0$. This problem corresponds to an approximation problem from sparse data where only 20% of the points are visible and there is a weak prior on the unobserved nodes pulling them toward zero.

We found the exact posterior by solving equation 2.12. We also ran loopy belief propagation and found that when it converged, the loopy means were identical to the true means up to machine precision. Also, as predicted by the theory, the loopy variances were too small; the loopy estimate was overconfident.

How is this predicted from our analysis? This illustrates the reparameterization lemma in section 3.5. The parameterization we used in the loopy network was

$$V_{ij} = \begin{pmatrix} w_{ij} & -w_{ij} \\ -w_{ij} & w_{ij} \end{pmatrix}, \quad V_{ii} = \begin{pmatrix} w_{ii} & -w_{ii} \\ -w_{ii} & w_{ii} \end{pmatrix}.$$

Thus, V_{ij} are not diagonally dominant. However, we can also use a different parameterization in which we shift the diagonal component from V_{ii} to V_{ij} :

$$\tilde{V}_{ij} = \begin{pmatrix} w_{ij} + \frac{w_{ii} - \epsilon}{\deg(\tilde{x}_i)} & -w_{ij} \\ -w_{ij} & w_{ij} + \frac{w_{ii} - \epsilon}{\deg(\tilde{x}_j)} \end{pmatrix}, \quad \tilde{V}_{ii} = \begin{pmatrix} \epsilon & -w_{ii} \\ -w_{ii} & w_{ii} \end{pmatrix}.$$

In this parameterization, all \tilde{V}_{ij} are diagonally dominant. Thus, claim 4 guarantees that belief propagation will converge and the variances will be overconfident.

We also ran belief propagation on this problem with $w_{ii} = 0$ for the unobserved nodes. Note that under this case, we cannot reparameterize the unwrapped tree so that all V_{ij} are diagonally dominant, so claim 4 does not hold. We found that belief propagation still converged with a similar convergence rate and that the means were exact. This illustrates the fact that claim 4 gives only sufficient (but not necessary) conditions for the conditional correlation to decrease rapidly enough. Claim 5 guarantees that the means will be exact.

In many applications, the solution of equation 2.12 by matrix inversion is intractable and iterative methods are used. Figure 4 compares the error in the means as a function of iterations for loopy propagation, successive overrelaxation (SOR), and conjugate gradient (CG). We used an overrelaxation constant of 1.9. Note that after five iterations, loopy propagation has a mean squared error of order 10^{-1} , while SOR and CG require many more (although the asymptotic performance of CG is the best). As expected by the

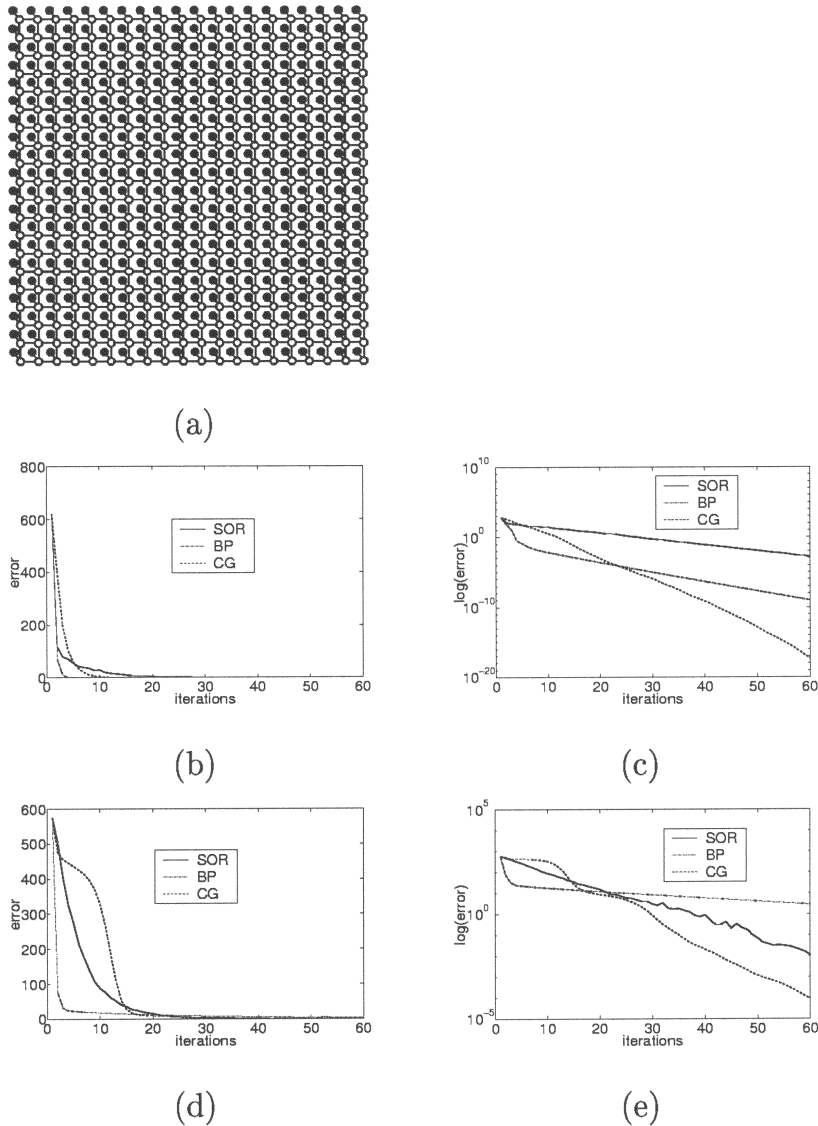


Figure 4: (a) 25×25 graphical model for simulation. (b) The error of the estimates of belief propagation (BP), conjugate gradient (CG), and successive overrelaxation (SOR) as a function of iteration. (c) The same plot as in *b* but with log scaling on the y axis. (d–e) Similar plots when the connections between unobserved nodes were increased by four orders of magnitude. BP converges very quickly to a good solution, although SOR and CG may have better asymptotic properties.

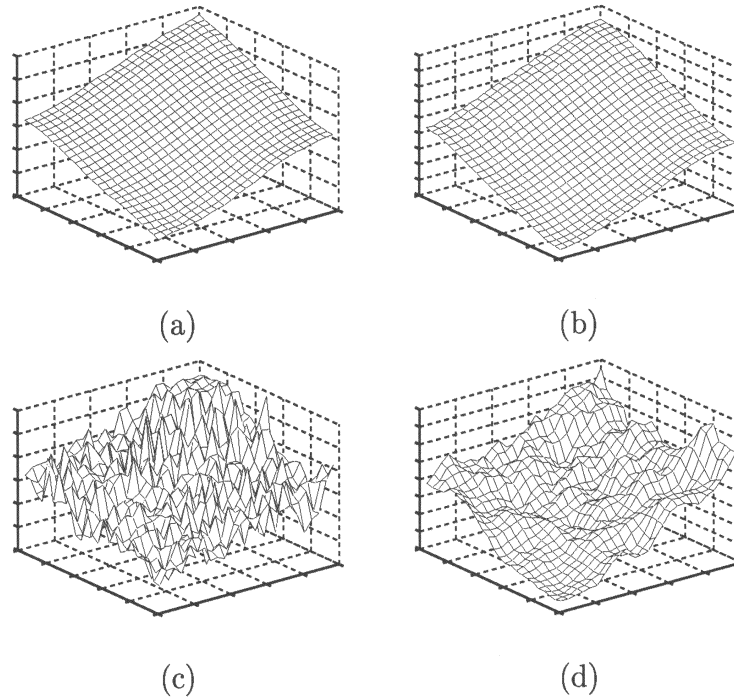


Figure 5: (a) The correct interpolation. (b) Results of BP after 10 iterations. (c) Results of SOR after 10 iterations. (d) Results of CG after 10 iterations.

fast convergence, the approximation error in the variances was quite small. The median error was 0.018. For comparison, the true variances ranged from 0.01 to 0.94, with a mean of 0.322. Also, the nodes for which the approximation error was worse were indeed the nodes that converged more slowly.

The analysis in section 3.5 suggests that the convergence rate is related to the ratio of diagonal to off-diagonal elements in \tilde{V}_{ij} . To illustrate this, we redid the simulations with the same network but set $w_{ij} = 10.0$ between neighboring nodes. Indeed convergence is much slower in this case (see figures 4d and 4e). As expected from the slower convergence, the error in the variance estimates is larger here with the median error 0.299. For visual comparison, Figure 5 shows the optimal answer and the results of BP, CG, and SOR after 10 iterations.

The slow convergence of SOR on problems such as these leads to the development of multiresolution models in which the MRF is approximated by a tree (Luetgen, Karl, & Willsky, 1994; Daniel & Willsky, 1999) and an algorithm equivalent to belief propagation is then run on the tree. Although the multiresolution models are much more efficient for inference, the tree

structure often introduces block artifacts in the estimate. Our results suggest that one can simply run belief propagation on the original MRF and get the exact posterior means (in theory, BP may fail to converge, but in many simulation runs, we never observed a failure to converge of BP in these two-dimensional interpolation problems). It would be interesting to see whether one can use the convergence rate of the beliefs to improve the variance estimate at every node.

7 Discussion

Independent of our results, two groups have recently analyzed special cases of gaussian graphical models. Frey (2000) analyzed the graphical model corresponding to factor analysis and gave conditions for the existence of a stable fixed point. Rusmevichientong and Van Roy (2000) analyzed a graphical model with the topology of turbo decoding but a gaussian joint density. They showed that for this specific case, belief propagation converges and the means are exact.

Our main interest in analyzing the gaussian case was to understand the performance of belief propagation in networks with multiple loops. Although there are many special properties of gaussians, we are struck by the similarity of the analytical results reported here for multiloop gaussians and the analytical results for single loops and general distributions reported in Weiss (2000). The most salient similarities are these:

- In single-loop networks with binary nodes, the mode at each node is guaranteed to be correct, but the confidence in the mode may be incorrect. In gaussian networks with multiple loops, the mean at each node is guaranteed to be correct, but the confidence around that mean will in general be incorrect.
- In single-loop networks, fast convergence is correlated with good approximation of the beliefs. This is also true for gaussian networks with multiple loops.
- In single-loop networks, the convergence rate and the approximation error were determined by a ratio of eigenvalues λ_1/λ_2 . This ratio determines the extent of the statistical dependencies between the root and the leaf nodes in the unwrapped network for a single loop. In gaussian networks, the convergence rate and the approximation error are determined by the off-diagonal terms of $\tilde{C}_{x|y}$. These terms quantify the extent of conditional dependencies between the root nodes and the leaf nodes of the unwrapped network.

These similarities are even more intriguing when one considers how different gaussian graphical models are from discrete models with arbitrary potentials and a single loop. In gaussian, the conditional mean is equal to the conditional mode, and there is only one maximum in the posterior probabil-

ity, while the single-loop discrete models may have multiple maxima, none of which will be equal to the mean. Furthermore, in terms of approximate inference, the two classes behave quite differently. For example, mean-field approximations give the exact means for gaussian MRFs, while they work poorly in discrete networks with a single loop in which the connectivity is sparse (Weiss, 1996). The resemblance of the results for gaussian graphical models and for single loops leads us to believe that similar results may hold for a larger class of networks.

The sum-product and max-product belief propagation algorithms are appealing, fast, and easily parallelizable algorithms. Due to the well-known hardness of probabilistic inference in graphical models, belief propagation will obviously not work for arbitrary networks and distributions. Nevertheless, there is a growing body of empirical evidence showing its success in many loopy networks. Our results give a theoretical justification for applying belief propagation in networks with multiple loops. This may enable fast, approximate probabilistic inference in a range of new applications.

Acknowledgments

We thank A. Ng, K. Murphy, P. Pakzad, B. Frey, and M. I. Jordan for comments on previous versions of this article. Y. W. is supported by MURI-ARO-DAAH04-96-1-0341

References

- Aji, S., Horn, G., & McEliece, R. (1998). On the convergence of iterative decoding on graphs with a single cycle. In *Proc. 1998 ISIT*.
- Aji, S. M., & McEliece, R. (1999). The generalized distributive law. *IEEE Transactions on Information Theory*, 46, 325–343.
- Berrou, C., Glavieux, A., & Thitimajshima, P. (1993). Near Shannon limit error-correcting coding and decoding: Turbo codes. In *Proc. IEEE International Communications Conference '93*.
- Cowell, R. (1998). Advanced inference in Bayesian networks. In M. Jordan (Ed.), *Learning in graphical models*. Cambridge, MA: MIT Press.
- Daniel, M. M., & Willsky, A. S. (1999). The modeling and estimation of statistically self-similar processes in a multiresolution framework. *IEEE Trans. Info. Theory*, 45(3), 955–970.
- Forney, G., Kschischang, F., Marcus, B., & Tuncel, S. (2000). Iterative decoding of tail-biting trellises and connections with symbolic dynamics. In *Proc. IMA Workshop on Codes, Systems, and Graphical Models*. New York: Springer-Verlag.
- Freeman, W., Pasztor, E., & Carmichael, O. (2000). Learning low-level vision. *Intl. J. of Comp. Vision*, 40(1), 25–47.
- Frey, B. J. (1998). *Graphical models for pattern classification, data compression and channel coding*. Cambridge, MA: MIT Press.

- Frey, B. (2000). Local probability propagation for factor analysis. In S. A. Solla, T. K. Leen, & K.-R. Müller (Eds.), *Advances in neural information processing systems*, 12. Cambridge, MA: MIT Press.
- Frey, B. J., Koetter, R., & Vardy, A. (1998). Skewness and pseudocode-words in iterative decoding. In *Proc. 1998 ISIT*.
- Gallager, R. (1963). *Low density parity check codes*. Cambridge, MA: MIT Press.
- Geman, S., & Geman, D. (1984). Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images. *IEEE Trans. PAMI*, 6(6), 721–741.
- Jensen, F. (1996). *An introduction to Bayesian networks*. New York: Springer-Verlag.
- Kschischang, F. R., & Frey, B. J. (2000). Iterative decoding of compound codes by probability propagation in graphical models. *IEEE Journal on Selected Areas in Communication*, 16(2), 219–230.
- Lauritzen, S. (1996). *Graphical models*. New York: Oxford University Press.
- Luetzgen, M. R., Karl, W. C., & Willsky, A. S. (1994). Efficient multiscale regularization with application to the computation of optical flow. *IEEE Transactions on Image Processing*, 3(1), 41–64.
- McEliece, R., MacKay, D., & Cheng, J. (1998). Turbo decoding as an instance of Pearl's "belief propagation" algorithm. *IEEE Journal on Selected Areas in Communication*, 16(2), 140–152.
- McEliece, R., Rodemich, E., & Cheng, J. (1995). The Turbo decision algorithm. In *Proc. 33rd Allerton Conference on Communications, Control and Computing* (pp. 366–379). Monticello, IL.
- Murphy, K., Weiss, Y., & Jordan, M. (1999). Loopy belief propagation for approximate inference: An empirical study. In *Proceedings of the Fifteenth Conference on Uncertainty in Artificial Intelligence*. San Mateo, CA: Morgan Kaufmann.
- Pearl, J. (1988). *Probabilistic reasoning in intelligent systems: Networks of plausible inference*. San Mateo, CA: Morgan Kaufmann.
- Rusmevichientong, P., & Van Roy, B. (2000). An analysis of Turbo decoding with gaussian densities. In S. A. Solla, T. K. Leen, & K.-R. Müller (Eds.), *Advances in neural information processing systems*, 12. Cambridge, MA: MIT Press.
- Weiss, Y. (1996). Interpreting images by propagating Bayesian beliefs. In M. Mozer, M. Jordan, & T. Petsche (Eds.), *Advances in neural information processing systems*, 9. Cambridge, MA: MIT Press.
- Weiss, Y. (1997). *Belief propagation and revision in networks with loops* (Tech. Rep. No. 1616). Cambridge, MA: MIT AI Lab.
- Weiss, Y. (2000). Correctness of local probability propagation in graphical models with loops. *Neural Computation*, 12, 1–41.
- Weiss, Y., & Freeman, W. (2001). On the optimality of solutions of the max-product belief propagation algorithm. *IEEE Transactions on Information Theory*, 37(2), 736–744.
- Wiberg, N. (1996). *Codes and decoding on general graphs*. Unpublished doctoral dissertation, University of Linköping, Sweden.