

Belief propagation for networks with loops

Alec Kirkley^{1*†}, George T. Cantwell^{1,2†}, M. E. J. Newman^{1,2,3}

Belief propagation is a widely used message passing method for the solution of probabilistic models on networks such as epidemic models, spin models, and Bayesian graphical models, but it suffers from the serious shortcoming that it works poorly in the common case of networks that contain short loops. Here, we provide a solution to this long-standing problem, deriving a belief propagation method that allows for fast calculation of probability distributions in systems with short loops, potentially with high density, as well as giving expressions for the entropy and partition function, which are notoriously difficult quantities to compute. Using the Ising model as an example, we show that our approach gives excellent results on both real and synthetic networks, improving substantially on standard message passing methods. We also discuss potential applications of our method to a variety of other problems.

Copyright © 2021
The Authors, some
rights reserved;
exclusive licensee
American Association
for the Advancement
of Science. No claim to
original U.S. Government
Works. Distributed
under a Creative
Commons Attribution
NonCommercial
License 4.0 (CC BY-NC).

INTRODUCTION

Many complex phenomena can be modeled using networks, which provide powerful abstract representations of systems in terms of their components and interactions (1). Phenomena of interest are often modeled using probabilistic formulations that capture the probabilities of states of network nodes. Examples include the spread of epidemics through networks of social contacts (2), cascading failures in power grids (3), and the equilibrium behavior of spin models such as the Ising model (4). Networks are also used to represent pairwise dependencies between variables in statistical models that do not otherwise have a network component as a convenient tool for bookkeeping and visualization of model structure (5). Such “graphical models,” which allow us to represent the conditional dependencies between variables in a nonparametric manner, form the foundation for many modern machine learning techniques (6).

The solution of probabilistic models like this presents a challenge. Analytic methods such as those used for regular lattices do not generalize to the more complex topologies of networks, and mean field and other standard approximations often fail to take crucial details of network structure into account. Numerical methods can be successful but are computationally demanding on larger networks and sometimes give results of poor accuracy. Message passing or “belief propagation” methods offer an alternative and promising approach that straddles the line between analytic and numerical techniques (7, 8). Message passing works by deriving a set of self-consistent equations satisfied by the variables or probabilities of interest and then solving those equations by numerical iteration. The name “message passing” comes from the fact that the equations can be thought of as representing messages passed between neighboring nodes in the network.

Standard formulations of message passing, however, have a crucial weakness: They rely on the assumption that the states of the neighbors are uncorrelated with one another, which is only true if the network contains no loops. Unfortunately, almost all real-world networks do contain loops, and usually many of them (9), so

standard message passing can give quite poor results in practical situations. Here, we propose a solution to this problem in the form of a new class of message passing methods for probabilistic models on “loopy” networks. These methods open up a host of possibilities for novel network calculations, many of which we discuss here.

The limitations of traditional message passing have been widely noted in the past, and a number of previous attempts have been made to remedy them. The only truly loopless networks are trees, but standard message passing methods have been shown to give good results on networks that satisfy the weaker condition of being “locally tree-like,” meaning that local regions of the network take the form of trees, although the network as a whole is not a tree. In effect, this means that the network can contain long loops but not short ones (1). However, realistic networks often fail to satisfy even this weaker condition and contain many short loops. Message passing has been extended to certain classes of random graphs with short loops, such as Husimi graphs (10–12) and other tree-like agglomerations of small loopy subgraphs (13, 14), but these techniques are not generally applicable to real-world networks. Alternatively, one can incorporate the effect of loops by using a perturbative expansion around the loopless case (15, 16), although this approach becomes progressively less accurate as the number of loops increases and is therefore best suited to networks with a low loop density, which rules out a large fraction of real networks, whose loop density is often high (9, 17). Perhaps the best known extension of belief propagation, and the one most similar to our own approach, is the method known as generalized belief propagation (18), which is based on the idea of passing messages not just between pairs of nodes but between larger groups. This method is, however, quite complicated to implement and requires explicit construction of the groups, which involves nontrivial preprocessing steps (19). The method we propose requires no such steps.

In (20), we previously described message passing schemes for percolation models and spectral calculations on loopy networks. Here, we extend this approach to the solution of general probabilistic models. We derive a factorization of the probability of states for such models that allows us to write self-consistent message passing equations for the marginal probabilities on sets of nodes in a neighborhood around a given reference node. From these equations, we can then calculate a range of quantities of interest such as single-site marginals, partition functions, and entropies. To ground our discussion, we use the Ising model as an example of our approach,

¹Department of Physics, University of Michigan, Ann Arbor, MI 48109, USA. ²Santa Fe Institute, 1399 Hyde Park Road, Santa Fe, NM 87501, USA. ³Center for the Study of Complex Systems, University of Michigan, Ann Arbor, MI 48109, USA.

*Corresponding author. Email: alec.w.kirkley@gmail.com

†These authors contributed equally to this work.

showing how our improved message passing methods can produce better estimates for this model than regular message passing. We show that our methods are asymptotically exact on networks whose loop structure satisfies certain general conditions and give good approximations for networks that deviate from these conditions. We give example results for the Ising model on both real and artificial networks and also discuss applications of our method to a range of other problems, emphasizing its wide applicability.

MATERIALS AND METHODS

Our first step is to develop the general theory of message passing for probabilistic models on loopy networks. With an eye on the Ising model, our discussion will be in the language of spin models, although the methods that we describe can be applied to any probabilistic model with pairwise dependencies between variables, making it suitable for a broad range of calculations in probabilistic modeling.

Model description

Consider a general undirected, unweighted network G composed of a set V of nodes or vertices and a set E of pairwise edges. The network can be represented mathematically by its adjacency matrix A with elements $A_{ij} = 1$ when nodes i and j are connected by an edge and 0 otherwise. On each node of the network, there is a variable or spin s_i , which is restricted to some discrete set of values S . In a compartmental model of disease propagation, for instance, $s_i \in S = \{0 \text{ (susceptible)}, 1 \text{ (infected)}, 2 \text{ (removed)}\}$ could be the infection state of a node (21, 22). In a spatial model of segregation, $s_i \in S = \{0 \text{ (unoccupied)}, 1 \text{ (occupied)}\}$ could represent land occupation (23).

Spins s_i and s_j interact if and only if there is an edge between nodes i and j , a formulation sufficiently general to describe a large number of models in fields as diverse as statistical physics, machine learning, economics, psychology, epidemiology, and sociology (24–30). Interactions are represented by an interaction energy $g_{ij}(s_i, s_j | \omega_{ij})$, which controls the preference for any particular pair of states s_i and s_j to occur together. The quantity ω_{ij} represents any external parameters, such as temperature in a classical spin system or infection rate in an epidemiological model, which control the nature of the interaction. We also allow for the inclusion of an external field $f_i(s_i | \theta_i)$ with parameters θ_i , which controls the intrinsic propensity for s_i to take an particular state. This could be used, for instance, to encode individual risk of catching a disease in an epidemic model.

Given these definitions, we write the probability $P(\mathbf{s} | \omega, \theta)$ that the complete set of spins takes value \mathbf{s} in the Boltzmann form

$$P(\mathbf{s} | \omega, \theta) = \frac{e^{-H(\mathbf{s} | \omega, \theta)}}{Z(\omega, \theta)} \quad (1)$$

where the Hamiltonian

$$H(\mathbf{s} | \omega, \theta) = - \sum_{(i,j) \in E} g_{ij}(s_i, s_j | \omega_{ij}) - \sum_{i \in V} f_i(s_i | \theta_i) \quad (2)$$

is the log probability of the state to within an arbitrary additive constant, and the partition function

$$Z(\omega, \theta) = \sum_{\mathbf{s}} e^{-H(\mathbf{s} | \omega, \theta)} \quad (3)$$

is the appropriate normalizing constant, ensuring that $P(\mathbf{s} | \omega, \theta)$ sums to unity. Here, we will primarily be concerned with computing the single-site (or one-point) marginal probabilities

$$q_i(s_i) = \sum_{\mathbf{s} \setminus s_i} P(\mathbf{s} | \omega, \theta) \quad (4)$$

where $\mathbf{s} \setminus s_i$ denotes all spins with the exception of s_i . For convenience, we have dropped ω and θ from the notation on the left of the equation, but it should be clear contextually that q_i depends on both of these variables.

The one-point marginals reveal useful information about physical systems, such as the magnetization of classical spin models or the position of a phase transition. They are important for statistical inference problems, where they give the posterior probability of a variable taking a given state after averaging over contributions from all other variables (e.g., the total probability of an individual being infected with a disease at a given time). Unfortunately, direct computation of one-point marginals is difficult because the number of terms in the sum in Eq. 4 grows exponentially with the number of spins. The message passing method gives us a way to get around this difficulty and compute q_i accurately and rapidly.

Message passing can also be used to calculate other quantities. For instance, we will show how to compute the average energy (also called the internal energy), which is given by

$$U(\omega, \theta) = \sum_{\mathbf{s}} H(\mathbf{s} | \omega, \theta) P(\mathbf{s} | \omega, \theta) \quad (5)$$

The average energy is primarily of interest in thermodynamic calculations, although it may also be of interest for statistical inference, where it corresponds to the average log likelihood.

We can also compute the two-point correlation function between spins

$$P(s_i = x, s_j = y) = P(s_j = y | s_i = x) q_i(s_i = x) \quad (6)$$

This function can be computed by first calculating the one-point marginal $q_i(s_i = x)$, then fixing $s_i = x$ and repeating the calculation for s_j . The same approach can also be used to compute n -point correlation functions.

Message passing equations

Our method operates by dividing a network into neighborhoods (20). A neighborhood $N_i^{(r)}$ around node i is defined as the node i itself and all of its edges and neighboring nodes, plus all nodes and edges along paths of length r or less between the neighbors of i . See Fig. 1 for examples. The key to our approach is to focus initially on networks in which there are no paths longer than r between the neighbors of i , meaning that all paths are inside $N_i^{(r)}$. This means that all correlations between spins within $N_i^{(r)}$ are accounted for by edges that are also within $N_i^{(r)}$, which allows us to write exact message passing equations for these networks. Equivalently, we can define a primitive cycle of length r starting at node i to be a cycle (i.e., a self-avoiding loop) such that at least one edge in the cycle is not on any shorter cycle beginning and ending at i . Our methods are then exact on any network that contains no primitive cycles of length greater than $r + 2$.

This approach gives us a series of methods where the r th member of the series is exact on networks that contain primitive cycles of length $r + 2$ and less only. The calculations become progressively more complex as r gets larger: They are very tractable for smaller values but become impractical when r is large. In many real-world networks, the longest primitive loop will be relatively long, requiring an infeasible computation to reach an exact solution.

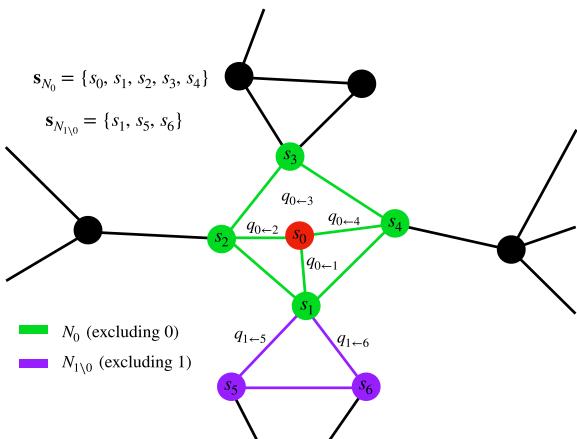


Fig. 1. Hamiltonian expansion diagram, with $r=2$. The focal node is in red, while the rest of its neighborhood N_0 is in green. Nodes and edges in purple represent the neighborhood $N_{1 \setminus 0}$ excluding node 1. We also label the corresponding spin and message variables used in Eqs. 11 and 12.

However, long loops introduce smaller correlations between variables than short ones, and moreover, the density of long loops is, in many cases, low: The network is “locally dense but globally sparse.” In this situation, we find that the message passing equations for low values of r , while not exact, give excellent results. They account correctly for the effect of the short loops in the network while making only a small approximation by omitting the long ones.

In practice, quite modest values of r can give good results. The smallest possible choice is $r=0$, which corresponds to assuming that there are no loops in the network at all and that the network is a tree. This is the assumption made by traditional message passing methods, and it gives poor results on many real-world networks. The next approximation after this, however, with $r=1$, which correctly accounts for the effect of loops of length three in the network (i.e., triangles), produces substantially better results, and the $r=2$ approximation (which includes loops of length three and four) is, in many cases, impressively accurate. In the following developments, we drop r from our notation for convenience—the same equations apply for all values of r .

Having defined the initial neighborhood N_i , we further define a neighborhood $N_{j \setminus i}$ to be node j plus all edges in N_j that are not contained in N_i and the nodes at their ends. Our method involves writing the marginal probability distribution on the spin at node i in terms of a set of messages received from nodes j that are in N_i , including nodes that are not immediate neighbors of i . (This contrasts with traditional message passing in which messages are received only from the immediate neighbors of i .) These messages are then, in turn, calculated from further messages j receives from nodes $k \in N_{j \setminus i}$ and so forth.

When written in this manner, the messages i receives are independent of one another in any network with no primitive cycles longer than $r+2$. Messages received from any two nodes j_1 and j_2 within N_i are necessarily independent, since they are calculated from the corresponding neighborhoods $N_{j_1 \setminus i}$ and $N_{j_2 \setminus i}$, which are disconnected from one another: If they were connected by any path, then that path would create a primitive cycle starting at i but passing outside of N_i , of which by hypothesis there are none. By the same

argument, we also know that $N_{j \setminus i}$ and N_i only overlap at the single node j for any $j \in N_i$.

This much is in common with our previous approach in (20), but to apply these ideas to the solution of probabilistic models, we need to go further. Specifically, we now show how this neighborhood decomposition allows us to factorize the Hamiltonian into a product of independent sums over the individual neighborhoods, with interactions that can be represented by messages passed between neighborhoods. Consider N_i as comprising a central set of nodes and edges surrounding i , then we can think of the set of neighborhoods $N_{j \setminus i}$ for all $j \in N_i$ as comprising the next “layer” in the network, the sets $N_{k \setminus j}$ for all $k \in N_{j \setminus i}$ as a third layer, and so forth until all nodes and edges in the network are accounted for. In a network with no primitive cycles longer than $r+2$, this procedure counts all interactions exactly once, allowing us to rewrite our Hamiltonian as a sum of independent contributions from the various layers, thus

$$H(\mathbf{s}) = H_{N_i}(\mathbf{s}_{N_i}) + \sum_{j \in N_i} H_{N_{j \setminus i}}(\mathbf{s}_{N_{j \setminus i}}) + \sum_{j \in N_i} \sum_{k \in N_{j \setminus i}} H_{N_{k \setminus j}}(\mathbf{s}_{N_{k \setminus j}}) + \dots, \quad (7)$$

where \mathbf{s}_{N_i} and $\mathbf{s}_{N_{j \setminus i}}$ are the sets of spins for the nodes in the neighborhoods N_i and $N_{j \setminus i}$, and we have defined the local Hamiltonians

$$H_{N_i}(\mathbf{s}_{N_i}) = - \sum_{(j,k) \in N_i} g_{jk}(s_j, s_k | \omega_{jk}) - f_i(s_i | \theta_i) \quad (8)$$

$$H_{N_{j \setminus i}}(\mathbf{s}_{N_{j \setminus i}}) = - \sum_{(k,l) \in N_{j \setminus i}} g_{kl}(s_k, s_l | \omega_{kl}) - f_j(s_j | \theta_j) \quad (9)$$

The decomposition of Eq. 7 is illustrated pictorially in Fig. 1.

The essential feature of this decomposition is that it breaks sums over spins such as those in Eqs. 3 and 4 into a product of sums over the individual neighborhoods $\{N_{j \setminus i}\}_{j \in N_i}$. Because these neighborhoods are, as we have said, independent, this means that the partition function and related quantities factorize into products of sums over a few spins each, which can easily be performed numerically. For instance, the one-point marginal of Eq. 4 takes the form

$$q_i(s_i = x) \propto \sum_{\mathbf{s}_{N_i}; s_i = x} e^{-H_{N_i}(\mathbf{s}_{N_i})} \prod_{j \in N_i} \sum_{\mathbf{s}_{N_{j \setminus i}}} e^{-H_{N_{j \setminus i}}(\mathbf{s}_{N_{j \setminus i}})} \times \prod_{k \in N_{j \setminus i}} \sum_{\mathbf{s}_{N_{k \setminus j}}} e^{-H_{N_{k \setminus j}}(\mathbf{s}_{N_{k \setminus j}})} \dots, \quad (10)$$

which can be written recursively as

$$q_i(s_i = x) = \frac{1}{Z_i} \sum_{\mathbf{s}_{N_i}; s_i = x} e^{-H_{N_i}(\mathbf{s}_{N_i})} \prod_{j \in N_i} q_{i \leftarrow j}(s_j) \quad (11)$$

with

$$q_{i \leftarrow j}(s_j = y) = \frac{1}{Z_{i \leftarrow j}} \sum_{\mathbf{s}_{N_{j \setminus i}}; s_j = y} e^{-H_{N_{j \setminus i}}(\mathbf{s}_{N_{j \setminus i}})} \prod_{k \in N_{j \setminus i} \setminus j} q_{j \leftarrow k}(s_k) \quad (12)$$

where the normalization constants Z_i and $Z_{i \leftarrow j}$ ensure that the marginals q_i and messages $q_{i \leftarrow j}$ are normalized so that they sum to 1. (In practice, we simply normalize the messages by dividing by their sum.) The quantity $q_{i \leftarrow j}(s_j)$ is equal to the marginal probability of node j having spin s_j when all the edges in N_i are removed. Alternatively,

one can think of it as a local external field on node j that influences the probability distribution of s_j . To make this more explicit, one could rewrite Eq. 11 as

$$q_i(s_i = x) = \frac{1}{Z_i} \sum_{s_{N_i} \setminus s_i = x} e^{-H_{N_i}(s_{N_i}) + \sum_{j \in N_i} \log q_{i \leftarrow j}(s_j)} \quad (13)$$

where $\log q_{i \leftarrow j}(s_j)$ plays the role of the external field.

Equations 11 and 12 define our message passing algorithm and can be solved for the messages $q_{i \leftarrow j}$ by simple iteration, starting from any suitable set of starting values and applying the equations repeatedly until convergence is reached.

With only slight modification, we can use the same approach to calculate the internal energy as well. The contribution to the internal energy from the interactions of a single node i is $\frac{1}{2} \sum_{j: A_{ij}=1} g(s_i, s_j | \omega_{ij}) + f(s_i | \theta_i)$, where the factor of $\frac{1}{2}$ compensates for double counting of interactions. Summing over all nodes i and weighting by the appropriate Boltzmann probabilities, the total internal energy is

$$U = \sum_{i \in V} \frac{1}{Z_i} \sum_{s_{N_i}} \left[\frac{1}{2} \sum_{j: A_{ij}=1} g(s_i, s_j | \omega_{ij}) + f(s_i | \theta_i) \right] e^{-H_{N_i}(s_{N_i})} \prod_{j \in N_i} q_{i \leftarrow j}(s_j) \quad (14)$$

All of the quantities appearing here are known a priori, except for the messages $q_{i \leftarrow j}(s_j)$ and the normalizing constants Z_i , which are calculated in the message passing process. Performing the message passing and then using the final converged values in Eq. 14 then gives us our internal energy.

Implementation

For less dense networks, those with node degrees up to about 20, the message passing equations of Eqs. 11 and 12 can be implemented directly and work well. The method is also easily parallelizable, as we can update all messages asynchronously based on their values from the previous iteration and compute the final marginals in parallel.

For networks with higher degrees, the calculations can become unwieldy, the huge reduction in complexity due to the factorization of the Hamiltonian notwithstanding. For a model with t distinct spin states at every node, the sum over states in the neighborhood of i has $t^{|N_i|}$ terms, which can quickly become computationally expensive to evaluate. Moreover, if just a single node has too large a neighborhood, then it can make the entire computation intractable, as that single neighborhood can consume more computational power than is available.

In such situations, therefore, we take a different approach. We note that Eq. 12 is effectively an expectation

$$q_{i \leftarrow j}(s_j = y) = \langle \delta_{s_j, y} \rangle_{N_j} \quad (15)$$

where we use the shorthand

$$\langle A \rangle_{N_j} = \sum_{s_{N_j}} A(s_{N_j}) \frac{e^{-H_{N_j}(s_{N_j})} \prod_{k \in N_j \setminus j} q_{s_k}^{j \leftarrow k}}{Z^{i \leftarrow j}} \quad (16)$$

We approximate this average using Markov chain Monte Carlo importance sampling over spin states, and, after convergence of the messages, the final estimates of the marginals q_i can also be obtained by Monte Carlo, this time on the spins in N_i . We describe the details in Results.

Calculating the partition function

The partition function Z is perhaps the most fundamental quantity in equilibrium statistical mechanics. From a knowledge of the partition function, one can calculate virtually any other thermodynamic variable of interest. Objects equivalent to Z also appear in other fields, such as Bayesian statistics, where the quantity known as the model evidence, the marginal likelihood of observed data given a hypothesized model, is mathematically analogous to the partition function and plays an important role in model fitting and selection (31–33).

Unfortunately, the partition function is difficult to calculate in practice. The calculation can be done analytically in some special cases (34, 35), but direct numerical calculations are difficult due to the need to sum over an exponentially large number of states, and Monte Carlo is challenging because of the difficulty of correctly normalizing the Boltzmann distribution.

Another concept central to statistical mechanics is the entropy

$$S = - \sum_s P(s) \ln P(s) \quad (17)$$

which has broad applications not just in physics but across the sciences (36–38). Like the partition function, entropy is difficult to calculate numerically, and for exactly the same reasons, since the two are closely related. For the canonical distribution of Eq. 1, the entropy is given in terms of Z by $S = \ln Z + \beta U$. Although we know the internal energy U therefore (which is relatively straightforward to compute), the entropy is at least as difficult to calculate as the partition function. The fundamental difficulty of normalizing the Boltzmann distribution is equivalent to establishing the zero of the entropy, a well-known hard problem (unsolvable within classical thermodynamics, requiring the additional axiom of the Third Law).

As we now show, the entropy can be calculated using our message passing formalism by appropriately factorizing the probability distribution over spin states. Since we have already developed a prescription for computing U (see Eq. 14), this also allows us to calculate the partition function. The details of the procedure are quite involved and do not follow straightforwardly from the previous discussion, so we defer the derivation to the Supplementary Materials, section S4. As shown there, the state probability $P(s)$ in Eq. 1 can be rewritten in the factorized form

$$P(s) = \frac{\prod_{i \in G} P(s_{N_i})}{\prod_{((i,j)) \in G} P(s_{\cap_{ij}})^{2/|\cap_{ij}|}} \quad (18)$$

where $P(s_{N_i})$ is the joint marginal distribution of the variables in the neighborhood of node i , $P(s_{\cap_{ij}})$ is the joint marginal distribution in the intersection $\cap_{ij} = N_i \cap N_j$ of the neighborhoods N_i and N_j , and $((i, j))$ denotes pairs of nodes that are contained in each other's neighborhoods.

By a series of manipulations, this form can be further expressed as the pure product

$$P(s) = \left[\prod_{((i,j)) \in G} P(s_{\cap_{ij}})^{1/\binom{|\cap_{ij}|}{2}} \right] \left[\prod_{(i,j) \in G} P(s_i, s_j)^{W_{ij}} \right] \left[\prod_{i \in G} P(s_i)^{C_i} \right] \quad (19)$$

where

$$W_{ij} = 1 - \sum_{((l,m)) \in G} \frac{1}{\binom{|\cap_{lm}|}{2}} \mathbf{1}_{\{(ij) \in \cap_{lm}\}} \quad (20)$$

with $\mathbf{1}_{\{\dots\}}$ being the indicator function, and

$$C_i = 1 - \left(\sum_{j \in N_i} \frac{1}{|\cap_{ij}| - 1} \right) - \left(\sum_{j \in N_i^{(0)}} W_{ij} \right) \quad (21)$$

Substituting Eq. 19 into Eq. 17, we get an expression for the entropy, thus

$$S = -\frac{1}{\binom{|\cap_{ij}|}{2}} \sum_{((i,j)) \in G} P(\mathbf{s}_{\cap_{ij}}) \ln P(\mathbf{s}_{\cap_{ij}}) \quad (22)$$

$$- \sum_{(i,j) \in G} W_{ij} P(s_i, s_j) \ln P(s_i, s_j) - \sum_{i \in G} C_i P(s_i) \ln P(s_i)$$

Note that, like the well-known Bethe approximation for the entropy (39), this expression not only has contributions from the one- and two-point marginals $P(s_i)$ and $P(s_i, s_j)$ of Eqs. 6 and 11 but also contains a term that depends on the joint marginal $P(\mathbf{s}_{\cap_{ij}})$ in the intersection \cap_{ij} , which may be nontrivial if $r > 0$. As shown in the Supplementary Materials, section S4, we can calculate this joint marginal using the message passing equation

$$P(\mathbf{s}_{\cap_{ij}}) = \frac{1}{Z_{\cap_{ij}}} e^{-\beta H(\mathbf{s}_{\cap_{ij}})} q_{i \leftarrow j}(s_j) \prod_{k \in \cap_{ij} \setminus j} q_{j \leftarrow k}(s_k) \quad (23)$$

where $H(\mathbf{s}_{\cap_{ij}})$ denotes the terms of the Hamiltonian of Eq. 2 that fall within \cap_{ij} and $Z_{\cap_{ij}}$ is the corresponding normalizing constant. For $|\cap_{ij}|$ sufficiently small, $Z_{\cap_{ij}}$ can be computed exactly. In other cases, we can calculate it using Monte Carlo methods similar to those that we used previously for the marginals $P(s_i)$.

Ising model calculations

As an archetypal application of our methods, we consider the Ising model on various example networks. The ferromagnetic Ising model in zero external field is equivalent in our notation to the choices

$$g_{ij}(s_i, s_j) = -\beta A_{ij} s_i s_j, \quad f_i(s_i) = 0 \quad (24)$$

where $\beta = 1/T$ is the inverse temperature. Note that temperature in this notation is considered a part of the Hamiltonian. It is more conventional to write temperature separately, so that the Hamiltonian has dimensions of energy rather than being dimensionless as here, but absorbing the temperature into the Hamiltonian is notationally convenient in the present case. It effectively makes the temperature a parameter ω_{ij} in Eq. 2 (and all ω_{ij} are equal).

As example calculations, we will compute the average magnetization M , which is given by

$$M = \left\langle \left\langle \frac{1}{N} \sum_{i=1}^N s_i \right\rangle \right\rangle = \frac{1}{N} \left\langle \sum_{i=1}^N [2q_i(s_i = +1) - 1] \right\rangle \quad (25)$$

and the heat capacity C , given by

$$C = \frac{dU}{dT} = -\beta^2 \frac{dU}{d\beta} \quad (26)$$

As detailed in section S1, we use an extension of the message passing equations to compute C that avoids having to use a numerical derivative to evaluate Eq. 26. Briefly, we consider the messages $q_{i \leftarrow j}$ to be a function of β , then define their derivatives with respect to β as their own set of messages

$$\eta_{i \leftarrow j} = \frac{dq_{i \leftarrow j}}{d\beta} \quad (27)$$

with their own associated message passing equations derived by differentiating Eq. 12. We then compute the heat capacity C by differentiating Eq. 14, expressing the result in terms of the $\eta_{i \leftarrow j}$, and substituting it into Eq. 26.

Behavior at the phase transition

In many geometries, the ferromagnetic Ising model has a phase transition at a nonzero critical temperature between a symmetric state with zero average magnetization and a symmetry broken state with nonzero magnetization. Substituting Eq. 24 into Eqs. 11 and 12, we can show that the message passing equations for the Ising model always have a trivial solution $q_{i \leftarrow j}(s_j) = \frac{1}{2}$ for all i, j . This choice is a fixed point of the message passing iteration: When started at this point, the iteration will remain there indefinitely. Looking at Eq. 25, we see that this fixed point corresponds to magnetization $M = 0$. If the message passing iteration converges to this trivial fixed point, therefore, it tells us that the magnetization is zero, and we are above the critical temperature; if it settles elsewhere, then the magnetization is nonzero, and we are below the critical temperature. Thus, the phase transition corresponds to the point at which the fixed point changes from being attracting to being repelling.

This behavior is well known in standard belief propagation, where it has been shown that, on networks with long loops only, there is a critical temperature T_{BP} below which the trivial fixed point becomes unstable and hence the system develops nonzero magnetization and that this temperature corresponds precisely to the conventional zero-field continuous phase transition on these networks (40). Extending the same idea to the present case, we expect the phase transition on a loopy network to fall at the corresponding transition point between stable and unstable in our message passing formulation.

Moreover, because the values of the messages at the trivial fixed point are known, we can compute an expression for the phase transition point without performing any message passing. We treat the message passing iteration as a dynamical system and perform a linear stability analysis of the trivial fixed point. Perturbing around $q = \frac{1}{2}$ (shorthand for setting all $q_{i \leftarrow j} = \frac{1}{2}$) and keeping terms to linear order, we find that the dynamics is governed by the Jacobian

$$J_{j \rightarrow i, v \rightarrow \mu} = \left. \frac{\partial q_{i \leftarrow j}}{\partial q_{\mu \leftarrow v}} \right|_{q=1/2} = \tilde{B}_{j \rightarrow i, v \rightarrow \mu} D_{j \rightarrow i, v \rightarrow \mu} \quad (28)$$

where \tilde{B} is a generalization of the so-called non-backtracking matrix (41) to our loopy message passing formulation

$$\bar{B}_{j \rightarrow i, v \rightarrow \mu} = \begin{cases} 1 & \text{if } j = \mu \text{ and } v \in N_{j \setminus i} \\ 0 & \text{otherwise,} \end{cases} \quad (29)$$

and $D_{j \rightarrow i, v \rightarrow \mu}$ is a correlation function between the spins s_μ and s_v within the neighborhood $N_{j \setminus i}$ —see section S3 for details.

When the magnitude of the leading eigenvalue λ_{\max} of this Jacobian is less than 1, the trivial fixed point is stable; when it is greater than 1, the fixed point is unstable. Hence, we can locate the phase transition temperature by numerically evaluating the Jacobian and locating the point at which $|\lambda_{\max}|$ crosses 1, for instance, by binary search.

Equation 29 is also useful in its own right. The non-backtracking matrix has numerous applications within network science, for instance in community detection (41), centrality measures (42), and percolation theory (43). The generalization defined in Eq. 29 could be used to extend these applications to loopy networks, although we will not explore such calculations here.

RESULTS

A model network

As a first example application, we examine the behavior of our method on a model network created precisely to have short loops only up to a specified maximum length. The network has short primitive cycles only of length $r + 2$ and less for a given choice of r , although it can also have long loops—it is locally dense but globally sparse in the sense discussed previously. This turns out to be a crucial point. The Ising model does not have a normal phase transition on a true tree, because, at any finite temperature, there is always a nonzero density of defects in the spin state (pairs of adjacent spins that are oppositely oriented), which, on a tree, divide the network into finite sized regions, imposing a finite correlation length and hence no critical behavior. Similarly, in the case of a network with only short loops and no long ones, there is no true phase transition. The long loops are necessary to produce criticality, a point discussed in detail in (44).

To generate networks that have short primitive cycles only up to a certain length, we first generate a random bipartite network—a network with two types of nodes and connections only between unlike kinds—then “project” down onto one type of node, producing a network composed of a set of complete subgraphs or cliques. In detail, the procedure is as follows.

1) We first specify the degrees of all the nodes, of both types, in the bipartite network.

2) We represent these degrees by “stubs” of edges emerging, in the appropriate numbers, from each node, then we match stubs at random in pairs to create our random bipartite network.

3) We project this network onto the nodes of type 1, meaning that any two such nodes that are both connected to the same neighbor of type 2 are connected directly with an edge in the projection. After all such edges have been added, the type 2 nodes are discarded.

4) Last, we remove a fraction p of the edges in the projected network at random.

If $p = 0$, then the network is composed of fully connected cliques, but when $p > 0$, some cliques will be lacking some edges; hence, the network is composed of a collection of subgraphs of size equal to the degrees of the corresponding nodes of type 2 from which they were projected. If we limit these degrees to a maximum value of $r + 2$, then there will be no short loops of length longer than this.

Figure 2 shows the magnetization per spin, entropy, and heat capacity for the ferromagnetic Ising model on an example network of 9447 nodes and 13,508 edges generated using this procedure with $r = 2$ and $p = 0.6$. We also limit the degrees of the type 1 nodes in the bipartite graph to a maximum of 5, which ensures that no neighborhood in the projection is too large to prevent a complete summation over states and hence that Monte Carlo estimation of the sums in the message passing equations is unnecessary.

Results are shown for belief propagation calculations with $r = 0, 1$, and 2 , the last of which should, in principle, be exact except for the weak correlations introduced by the presence of long loops in the network. We also show in the figure the magnitude of the leading eigenvalue of J for each value of r . The points at which this eigenvalue equals 1, which give estimates of the critical temperature for each r , are indicated by the vertical lines. Also shown in the figure for comparison are results from direct Monte Carlo simulations of the system, with the entropy calculated from values of the heat capacity computed from energy fluctuations and then numerically integrated using the identity

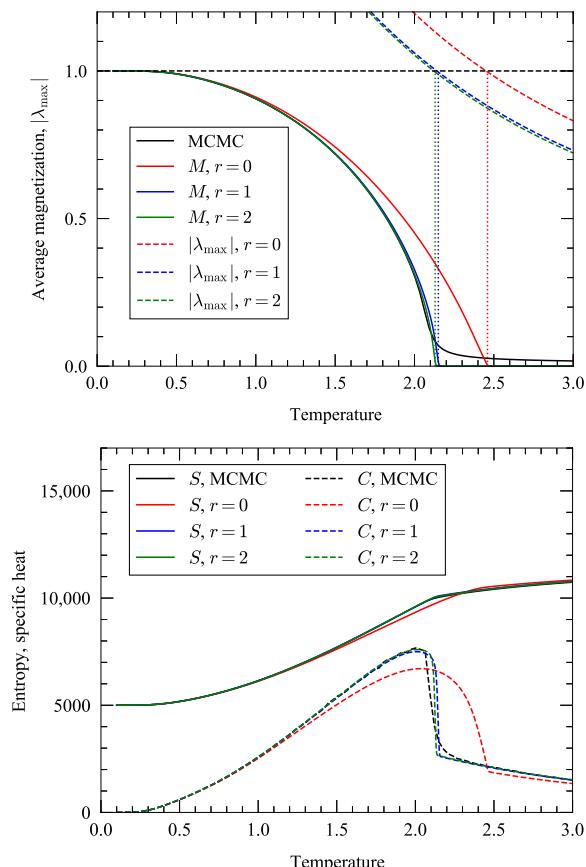


Fig. 2. Ferromagnetic Ising model critical behavior on synthetic network. The top panel shows the average magnetization, while the bottom one shows the heat capacity and the entropy (the latter shifted up for visualization purposes). Results are shown for $r = 0, 1, 2$ and direct Monte Carlo simulation (denoted MCMC). The magnitude of the leading eigenvalue for the Jacobian is also shown in the top panel for all three values of r , and we can see that the apparent positions of the phase transition, revealed by discontinuities in the physical quantities or their gradients, correspond closely to the temperatures at which the associated eigenvalues are equal to 1.

$$S = \int_0^T \frac{C(T)}{T} dT \quad (30)$$

The message passing simulations offer significantly faster results for this system: For $r = 2$, message passing was about 100 times faster than the Monte Carlo simulations.

Looking at Fig. 2, we can see that as we increase r , the message passing results approach those from the direct Monte Carlo, except close to the phase transition, where the Monte Carlo calculations suffer from finite size effects that smear the phase transition, to which the message passing approach appears largely immune. While the results for conventional belief propagation ($r = 0$) are quite far from the direct Monte Carlo results, most of the improvement in accuracy from our method is already present even at $r = 1$. Going to $r = 2$ offers only a small additional improvement in this case.

The apparent position of the phase transition aligns well with the predictions derived from the value of the Jacobian for each value of r . The transition is particularly clear in the gradient discontinuity of the magnetization. For $r = 1$ and 2, the heat capacity appears to exhibit a discontinuity at the transition, which differs from the divergence we expect on low-dimensional lattices but bears a resemblance to the behavior seen on Bethe lattices and other homogeneous tree-like networks (7, 45, 46).

Real-world networks

For our next example, we look at an application on a real-world network, where we do not expect the method to be exact, although as we will see it nonetheless performs well. The network that we examine has larger local neighborhoods than our synthetic example, which means that we are not able to sum exhaustively over all configurations of the spins $s_{N_j \setminus i}$ in Eq. 12 (and similarly s_{N_i} in Eq. 11) so, as described in Implementation, we instead make use of Monte Carlo sampling to estimate the messages $q_{i \leftarrow j}$ and marginals q_i .

The summation over local spins in Eq. 12 is equivalent to computing the expectation in Eq. 15. To calculate $q_{i \leftarrow j}(s_j = y)$, we fix the values of the incoming messages $\{q_{j \leftarrow k}\}$ and perform Monte Carlo sampling over the states of the spins in the neighborhood $N_j \setminus i$ with the Hamiltonian of Eq. 9. Then, we compute the average in Eq. 15 separately for the cases $y = 1$ and -1 and normalize to ensure that the results sum to one. The resulting values for $q_{i \leftarrow j}$ can then be used as incoming messages for calculating other messages in other neighborhoods. We perform the Monte Carlo using the Wolff cluster algorithm (47), which makes use of the Fortuin-Kasteleyn percolation representation of the Ising model to flip large clusters of spins simultaneously and can significantly reduce the time needed to obtain independent samples, particularly close to the critical point. Once the messages have converged to their final values, we compute the marginals q_i by performing a second Monte Carlo, this time over the spins s_{N_i} with the Hamiltonian of Eq. 8. More details on the procedure are given in section S2.

The Monte Carlo approach combines the best aspects of message passing and traditional Monte Carlo calculations. Message passing reduces the sums we need to perform to sets of spins much smaller than the entire network, while the Monte Carlo approach markedly reduces the number of spin states that need to be evaluated. The approach has other advantages too. For instance, because of the small neighborhood sizes, it shows improved performance in systems with substantial energy barriers that might otherwise impede

ergodicity, such as antiferromagnetic systems. But perhaps its biggest advantage is that it effectively allows us to sample very large numbers of states of the network without taking very large samples of individual neighborhoods. If we sample k configurations from one neighborhood and k configurations from another, then, in effect, we are summing over k^2 possible combinations of states in the union of the two neighborhoods. Depending on the value of r , there are at least $2m$ neighborhoods $N_j \setminus i$ in a network, where m is the number of edges, and hence we are effectively summing over at least k^{2m} states overall, a number that increases exponentially with network size. Effective sample sizes of 10^{1000} or more are easily reachable, far beyond what is possible with traditional Monte Carlo methods.

Figure 3 shows the results of applying these methods with $r = 0 \dots 4$ to a network from (48) representing the structure of an electric power grid, along with results from direct Monte Carlo simulations on the same network. As the figure shows, the magnetization is again poorly approximated by the traditional ($r = 0$) message passing algorithm but improves as r increases. In particular, the behavior in the region of the phase transition is quite poor for $r = 0$ and does not provide a good estimate of the position of the transition. For $r = 1$ and 2, however, we get much better estimates, and for $r = 3$ and 4, the method approaches the Monte Carlo results quite closely, with the critical temperature falling somewhere in the region of $T = 1.6$ in this case. We also see a much clearer phase transition in the message passing results than in the standard Monte Carlo because of finite size effects in the latter. These results all suggest that, for real systems, our method can give substantial improvements over both ordinary belief propagation and direct Monte Carlo simulation and, in some cases, show completely different behavior altogether.

DISCUSSION

Here, we have presented a new class of message passing algorithms for solving probabilistic models on networks that contain a high density of short loops. Taking the Ising model as an example, we have shown that our methods give substantially improved results in calculations of magnetization, heat capacity, entropy, marginal spin probabilities, and other quantities over standard message passing methods that do not account for the presence of loops. Our methods are exact on networks with short loops up to a fixed maximum length, which we can choose, and can give good approximations on networks with loops of any length.

Message passing methods for probabilistic models on loopy networks have been proposed in the past, the best known being the generalized belief propagation method of Yedidia *et al.* (18). Generalized belief propagation uses a region-based approximation (49), in which the free energy $\ln Z$ is approximated by a sum of independent local free energies of regions within the network. Once the regions are defined, it is straightforward to write down belief propagation equations, which can be used to calculate marginals and other quantities of interest, including approximations to the partition function and entropy. Perhaps the best known example of generalized belief propagation, at least within the statistical physics community, is the cluster variational method, in which the regions are defined so as to be closed under the intersection operation (24), and the resulting free energy is called the Kikuchi free energy (50).

The accuracy and complexity of generalized belief propagation is determined by the specific choice of regions, which has been

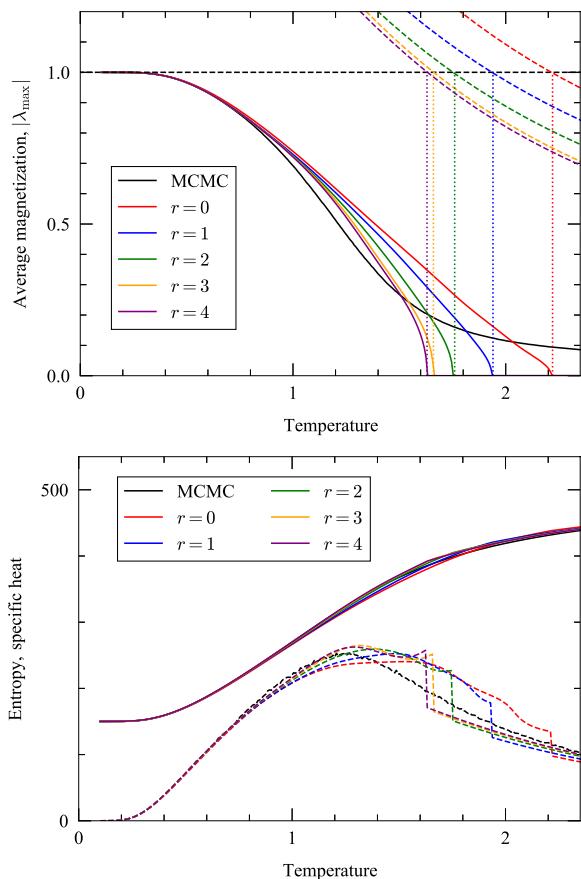


Fig. 3. Ferromagnetic Ising model critical behavior on a power grid network. Message passing and Monte Carlo calculations of the average magnetization, entropy, and specific heat on the “494 bus power system” network from (48). Results are shown for $r=0\dots 4$ and direct Monte Carlo simulation (MCMC). Again, the message passing results approximate the real solution progressively better as r grows larger.

described as being “more of an art than a science” (39). Loops contained within regions are correctly accounted for in the belief propagation, while those that span two or more regions are not and introduce error. At the same time, the computational complexity of the belief propagation calculations increases exponentially with the size of the regions (39), so choosing the right regions is a balancing act between enclosing as many loops as possible while not making the regions too large. A number of heuristics have been proposed for choosing the regions (19, 51, 52), but real-world networks can pose substantial difficulties because they often contain both high degrees and many loops (1), which effectively forces us to compromise either by leaving loops out or by using very large regions. Our method can have a substantial advantage in these systems because it can accommodate large, tightly connected neighborhoods through local Monte Carlo sampling. Our method also has the benefit that the neighborhoods are constructed automatically based on the network structure rather than being chosen by the user.

There are many ways in which the methods and results of this paper could be extended. We have studied only one application in detail, the Ising model, but the formalism that we present is a general one that could be applied to many other models. In principle, any model with sparse pairwise interactions (i.e., interactions

whose number scales subquadratically with the number of variables) could be studied using these methods. For example, there is a large class of generative models of networks in which edges appear with probabilities that depend on the properties of the adjacent nodes. Examples include the Chung-Lu model (53) and the stochastic block model and its variants (54, 55). If we assume an observed network to be drawn from such a model, then we can use statistical inference to estimate the values of hidden node attributes that influence edge probability, such as community membership. Our message passing methods could be applied to such inference calculations and could, in principle, give more accurate results in the common case where the observed network contains many short loops.

Another potential application in the realm of statistical inference is the inverse Ising model, the problem of inferring the parameters of an Ising or Ising-like model from an observed sequence of spin states, which has numerous applications including the reconstruction of neural pathways (56), the inference of protein structure (57), and correlations within financial markets (58). It can be shown that the one- and two-point correlation functions of the observed spins are sufficient statistics to reliably estimate coupling and external field parameters (59), and our method could be used to compute these statistics on loopy networks to greater accuracy than with traditional message passing and faster than standard Monte Carlo simulation. Other potential applications, further afield from traditional statistical physics, include the solution of constraint satisfaction problems, coding theory, and combinatorial optimization.

SUPPLEMENTARY MATERIALS

Supplementary material for this article is available at <http://advances.sciencemag.org/cgi/content/full/7/17/eabf1211/DC1>

REFERENCES AND NOTES

1. M. E. J. Newman, *Networks* (Oxford Univ. Press, ed. 2, 2018).
2. I. Z. Kiss, J. C. Miller, P. L. Simon, *Mathematics of Epidemics on Networks* (Springer International Publishing, 2017).
3. B. A. Carreras, V. E. Lynch, I. Dobson, D. E. Newman, Dynamical and probabilistic approaches to the study of blackout vulnerability of the power transmission grid, in *Proceedings of the 37th Annual Hawaii International Conference on System Sciences* (Institute of Electrical and Electronics Engineers, 2004), pp. 7–14.
4. S. N. Dorogovtsev, A. V. Goltsev, J. F. F. Mendes, Ising model on networks with an arbitrary distribution of connections. *Phys. Rev. E* **66**, 016104 (2002).
5. M. I. Jordan, *Learning in Graphical Models* (Kluwer Academic Publishers, 1998).
6. D. Koller, N. Friedman, *Probabilistic Graphical Models: Principles and Techniques* (MIT Press, 2009).
7. H. A. Bethe, Statistical theory of superlattices. *Proc. R. Soc. Lond. A* **150**, 552–575 (1935).
8. M. Mezard, A. Montanari, *Information, Physics, and Computation* (Oxford Univ. Press, 2009).
9. D. J. Watts, S. H. Strogatz, Collective dynamics of “small-world” networks. *Nature* **393**, 440–442 (1998).
10. M. Eckstein, M. Kollar, K. Byczuk, D. Vollhardt, Hopping on the Bethe lattice: Exact results for bondencies of states and dynamical mean-field theory. *Phys. Rev. B* **71**, 235119 (2005).
11. F. L. Metz, I. Neri, D. Bollé, Spectra of sparse regular graphs with loops. *Phys. Rev. E* **84**, 055101 (2011).
12. D. Bollé, F. L. Metz, I. Neri, On the spectra of large sparse graphs with cycles. [arXiv:1206.1512](https://arxiv.org/abs/1206.1512) (2012).
13. M. E. J. Newman, Random graphs with clustering. *Phys. Rev. Lett.* **103**, 058701 (2009).
14. S. Yoon, A. V. Goltsev, S. N. Dorogovtsev, J. F. F. Mendes, Belief-propagation algorithm and the Ising model on networks with arbitrary distributions of motifs. *Phys. Rev. E* **84**, 041144 (2011).
15. A. Montanari, T. Rizzo, How to compute loop corrections to the Bethe approximation. *J. Stat. Mech.* **2005**, 10011 (2005).
16. M. Chertkov, V. Y. Chernyak, Loop calculus in statistical physics and information science. *Phys. Rev. E* **73**, 065102 (2006).
17. M. E. J. Newman, J. Park, Why social networks are different from other types of networks. *Phys. Rev. E* **68**, 036122 (2003).

18. J. S. Yedidia, W. T. Freeman, Y. Weiss, Generalized belief propagation, in *Proceedings of the 13th International Conference on Neural Information Processing Systems* (MIT Press, 2000), pp. 668–674.
19. M. Welling, On the choice of regions for generalized belief propagation, in *Proceedings of the 20th Conference on Uncertainty in Artificial Intelligence* (AUAI Press, 2004), pp. 585–592.
20. G. T. Cantwell, M. Newman, Message passing on networks with loops. *Proc. Natl. Acad. Sci. U.S.A.* **116**, 23398–23403 (2019).
21. W. O. Kermack, A. G. McKendrick, A contribution to the mathematical theory of epidemics. *Proc. R. Soc. Lond. A* **115**, 700–721 (1927).
22. R. Durrett, *Spatial Epidemic Models: Their structure and Relation to Data* (Cambridge Univ. Press, 1995).
23. D. Stauffer, S. Solomon, Ising, Schelling and self-organising segregation. *Eur. Phys. J. B.* **57**, 473–479 (2007).
24. A. Pelizzola, Cluster variation method in statistical physics and probabilistic graphical models. *J. Phys. A* **38**, R309–R339 (2005).
25. S. Geman, C. Graffigne, Markov random field image models and their applications to computer vision, in *Proceedings of the International Congress of Mathematicians* (International Congress of Mathematicians, 1986), pp. 2.
26. M. Yasuda, T. Horiguchi, Triangular approximation for Ising model and its application to Boltzmann machine. *Physica A* **368**, 83–95 (2006).
27. A. Decelle, F. Krzakala, C. Moore, L. Zdeborová, Asymptotic analysis of the stochastic block model for modular networks and its algorithmic applications. *Phys. Rev. E* **84**, 066106 (2011).
28. W.-X. Zhou, D. Sornette, Self-organizing Ising model of financial markets. *Eur. Phys. J. B.* **55**, 175–181 (2007).
29. S. Galam, Rational group decision making: A random field Ising model at $T = 0$. *Physica A* **238**, 66–80 (1997).
30. D. Stauffer, Social applications of two-dimensional Ising models. *Am. J. Phys.* **76**, 470–473 (2008).
31. D. J. MacKay, *Information Theory, Inference and Learning Algorithms* (Cambridge Univ. Press, 2003).
32. H. S. Migon, D. Gamerman, F. Louzada, *Statistical Inference: An Integrated Approach* (CRC Press, 2014).
33. N. Friel, J. Wyse, Estimating the evidence—A review. *Statistica Neerlandica* **66**, 288–308 (2012).
34. S. Salinas, *Introduction to Statistical Physics* (Springer, 2001).
35. R. J. Baxter, *Exactly Solved Models in Statistical Mechanics* (Academic Press, 1982).
36. C. E. Shannon, Prediction and entropy of printed English. *Bell Syst. Tech. J.* **30**, 50–64 (1951).
37. W. Bialek, *Biophysics: Searching for Principles* (Princeton Univ. Press, 2012).
38. P. Cabral, G. Augusto, M. Tewolde, Y. Araya, Entropy in urban systems. *Entropy* **15**, 5223–5236 (2013).
39. J. S. Yedidia, W. T. Freeman, Y. Weiss, Understanding belief propagation and its generalizations. *Explor. Artificial Intell. New Millennium* **8**, 236–239 (2003).
40. J. Mooij, H. Kappen, On the properties of the Bethe approximation and loopy belief propagation on binary networks. *J. Stat. Mech.* **2005**, P11012 (2005).
41. F. Krzakala, C. Moore, E. Mossel, J. Neeman, A. Sly, L. Zdeborová, P. Zhang, Spectral redemption in clustering sparse networks. *Proc. Natl. Acad. Sci. U.S.A.* **110**, 20935–20940 (2013).
42. T. Martin, X. Zhang, M. E. J. Newman, Localization and centrality in networks. *Phys. Rev. E* **90**, 052808 (2014).
43. B. Karrer, M. E. J. Newman, L. Zdeborová, Percolation on sparse networks. *Phys. Rev. Lett.* **113**, 208702 (2014).
44. A. Allard, L. Hébert-Dufresne, On the accuracy of message-passing approaches to percolation in complex networks. *arXiv:1906.10377* (2019).
45. F. Mancini, A. Naddeo, Equations-of-motion approach to the spin-12 Ising model on the Bethe lattice. *Phys. Rev. E* **74**, 061108 (2006).
46. S. N. Dorogovtsev, A. V. Goltsev, J. F. Mendes, Critical phenomena in complex networks. *Rev. Mod. Phys.* **80**, 1275–1335 (2008).
47. U. Wolff, Collective Monte Carlo updating for spin systems. *Phys. Rev. Lett.* **62**, 361–364 (1989).
48. T. A. Davis, Y. Hu, The University of Florida sparse matrix collection. *ACM Trans. Math. Softw.* **38**, 1–25 (2011).
49. J. S. Yedidia, W. T. Freeman, Y. Weiss, Constructing free-energy approximations and generalized belief propagation algorithms. *IEEE Trans. Inf. Theory* **51**, 2282–2312 (2005).
50. R. Kikuchi, A theory of cooperative phenomena. *Phys. Rev.* **81**, 988–1003 (1951).
51. H. J. Kappen, W. Wiegerinck, Novel iteration schemes for the cluster variation method, in *Advances in Neural Information Processing Systems* (MIT Press, 2002), pp. 415–422.
52. P. Pakzad, V. Anantharam, Minimal graphical representation of Kikuchi regions, in *Proceedings of the Annual Allerton Conference on Communication Control and Computing* (University of Illinois, 2002), pp. 1586–1595.
53. F. Chung, L. Lu, The average distances in random graphs with given expected degrees. *Proc. Natl. Acad. Sci. U.S.A.* **99**, 15879–15882 (2002).
54. P. W. Holland, K. B. Laskey, S. Leinhardt, Stochastic blockmodels: First steps. *Soc. Networks* **5**, 109–137 (1983).
55. B. Karrer, M. E. J. Newman, Stochastic blockmodels and community structure in networks. *Phys. Rev. E* **83**, 016107 (2011).
56. E. Schneidman, M. J. Berry II, R. Segev, W. Bialek, Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* **440**, 1007–1012 (2006).
57. F. Morcos, A. Pagnani, B. Lunt, A. Bertolino, D. S. Marks, C. Sander, R. Zecchina, J. N. Onuchic, T. Hwa, M. Weigt, Direct-coupling analysis of residue coevolution captures native contacts across many protein families. *Proc. Natl. Acad. Sci. U.S.A.* **108**, E1293–E1301 (2011).
58. T. Bury, Market structure explained by pairwise interactions. *Physica A* **392**, 1375–1385 (2013).
59. H. C. Nguyen, R. Zecchina, J. Berg, Inverse statistical problems: From the inverse Ising problem to data science. *Adv. Phys.* **66**, 197–261 (2017).

Acknowledgments

Funding: This work was funded in part by the U.S. Department of Defense NDSEG fellowship program (to A.K.) and by the National Science Foundation under grants NSF IIS-1838251 (to G.T.C.), DMS-1710848 and DMS-2005899 (to M.E.J.N.). **Author contributions:** A.K., G.T.C., and M.E.J.N. designed the study; A.K. performed the simulations; and A.K. and G.T.C. performed the mathematical analysis. All authors contributed to writing the manuscript. **Competing interests:** The authors declare that they have no competing interests. **Data materials and availability:** All data needed to evaluate the conclusions in the paper are present in the paper and/or the Supplementary Materials, except for the power grid data, which are available from the original source in (48): <https://dl.acm.org/doi/10.1145/2049662.2049663>. Additional data related to this paper may be requested from the authors.

Submitted 5 October 2020

Accepted 9 March 2021

Published 23 April 2021

10.1126/sciadv.abf1211

Citation: A. Kirkley, G. T. Cantwell, M. E. J. Newman, Belief propagation for networks with loops. *Sci. Adv.* **7**, eabf1211 (2021).

Belief propagation for networks with loops

Alec Kirkley, George T. Cantwell and M. E. J. Newman

Sci Adv 7 (17), eabf1211.
DOI: 10.1126/sciadv.abf1211

ARTICLE TOOLS

<http://advances.science.org/content/7/17/eabf1211>

SUPPLEMENTARY MATERIALS

<http://advances.science.org/content/suppl/2021/04/19/7.17.eabf1211.DC1>

REFERENCES

This article cites 40 articles, 4 of which you can access for free
<http://advances.science.org/content/7/17/eabf1211#BIBL>

PERMISSIONS

<http://www.science.org/help/reprints-and-permissions>

Use of this article is subject to the [Terms of Service](#)

Science Advances (ISSN 2375-2548) is published by the American Association for the Advancement of Science, 1200 New York Avenue NW, Washington, DC 20005. The title *Science Advances* is a registered trademark of AAAS.

Copyright © 2021 The Authors, some rights reserved; exclusive licensee American Association for the Advancement of Science. No claim to original U.S. Government Works. Distributed under a Creative Commons Attribution NonCommercial License 4.0 (CC BY-NC).