

Belief propagation for finite networks using a symmetry-breaking source node

Seongmin Kim¹ and Alec Kirkley^{1,2,3,*}

¹*Institute of Data Science, University of Hong Kong, Hong Kong SAR, China*

²*Department of Urban Planning and Design, University of Hong Kong, Hong Kong SAR, China*

³*Urban Systems Institute, University of Hong Kong, Hong Kong SAR, China*



(Received 31 October 2025; accepted 21 January 2026; published 23 February 2026)

Belief propagation (BP) is an efficient message-passing algorithm widely used for inference in graphical models and for solving various problems in statistical physics. However, BP often yields inaccurate estimates of order parameters and their susceptibilities in finite systems, particularly in sparse networks with few loops. Here, we show for both percolation and Ising models that fixing the state of a single well-connected “source” node to break global symmetry substantially improves inference accuracy and captures finite-size effects across a broad range of networks, especially treelike ones, at no additional computational cost.

DOI: [10.1103/pxg-mhjr](https://doi.org/10.1103/pxg-mhjr)

I. INTRODUCTION

Message passing enables efficient inference in probabilistic graphical models of complex systems [1–3]. This approach has led to advances across diverse fields, including machine learning [4–6], epidemiology [7,8], social network analysis [9–12], and statistical physics [13–15]. Central to message passing is belief propagation (BP), also known as the cavity method, which efficiently computes marginal probabilities by decomposing global inference problems as the exchange of local messages between nodes and their neighbors. BP is exact on acyclic (tree) networks and often provides good approximations on cyclic (loopy) networks that are locally treelike [15,16].

In statistical physics, the application of BP to phase-transition models, such as percolation [17–22] and the Ising model [15,22–26], has yielded highly accurate predictions in many large real-world networks. However, this narrative of success holds a profound exception: BP’s ability to describe phase transitions catastrophically breaks down on the very structures where it is theoretically exact—tree networks—creating a paradox [18]. This failure persists on *almost-trees*, networks that are only slightly more connected.

This paradox arises because the system is symmetric and finite, and thus lacks true spontaneous symmetry breaking (SSB) in theory [27,28]. Although finite systems may appear to exhibit SSB and nontrivial order parameters, these are practical measurements rather than quantities defined in the thermodynamic limit. For example, percolation strength refers to the fractional size of the infinite cluster, which exists only

in infinitely large systems [29]; similarly, the average magnetization in a finite spin system vanishes if measured over an infinite time as thermal fluctuations restore symmetry [30]. Thus, theoretical order parameters are zero in finite systems.

BP is indeed exact on trees and correctly predicts the absence of SSB [18]. However, this theoretical consistency on trees makes BP incompatible with practical metrics used in finite systems, such as the fraction of the largest cluster in percolation or the absolute magnetization in the Ising model. This discrepancy highlights the need for a modified BP approach to accurately infer these practical metrics in finite systems. Reference [18] partially addressed this “treelike network catastrophe” by adding a clique to the tree to induce a nontrivial solution, but the resulting accuracy remains poor.

Here, we introduce the source-node BP (SNBP) method, a simple modification to the BP framework that explicitly breaks system symmetry by designating a single source node. This enables BP to accurately approximate order parameters and susceptibilities while maintaining computational efficiency. We show that SNBP significantly reduces finite-size errors across various networks, especially treelike ones, and closely matches ground-truth results from Monte Carlo (MC) simulations. SNBP also outperforms conventional BP and the naive mean-field approximation (MFA) in many real-world networks, offering a practical inference method for finite systems with global symmetries.

II. METHODS

We demonstrate our framework on two canonical phase-transition models, bond percolation and the Ising model, defined on an undirected graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with N nodes and M edges.

A. Bond percolation

Each bond $(i, j) \in \mathcal{E}$ is independently occupied with probability p . In the large- N limit, the model exhibits a phase

*Contact author: alec.w.kirkley@gmail.com

transition at a critical probability p_c : For $p < p_c$, only finite clusters exist, while for $p > p_c$, an infinite cluster spans a nonzero fraction of the network. The theoretical order parameter (percolation strength) is the probability P_∞ that a randomly chosen node belongs to the infinite cluster, which is strictly zero in any finite system. In practice, for a finite graph, we measure the practical order parameter as the fractional size of the largest cluster, $S_1 = |\mathcal{C}_{\max}|/N$, where $|\mathcal{C}_{\max}|$ is the number of nodes in the largest cluster, $\mathcal{C}_{\max}$ [29]. In the large- N limit, $S_1 \rightarrow P_\infty$. The theoretical susceptibility $\chi_{\text{true}} = \sum_{|\mathcal{C}| < \infty} |\mathcal{C}|^2/N$ measures the mean finite cluster size, where $|\mathcal{C}|$ is the size of cluster $\mathcal{C}$. In finite systems, the infinite cluster is absent, and the direct calculation of χ_{true} yields no peak. Therefore, the practical susceptibility is defined by excluding the largest cluster, treating it as the infinite cluster: $\chi_{\text{practical}} = \sum_{\mathcal{C} \neq \mathcal{C}_{\max}} |\mathcal{C}|^2/N$ [31].

B. Ising model

Each node $i \in \mathcal{V}$ carries a spin $\sigma_i = \pm 1$. In zero field, the Hamiltonian is $\mathcal{H}(\sigma) = -\sum_{(i,j) \in \mathcal{E}} \sigma_i \sigma_j$. The theoretical order parameter is the average magnetization, $m \equiv \sum_{i=1}^N \sigma_i/N$. In the large- N limit, the system exhibits SSB below a critical temperature T_c , and the equilibrium magnetization $\langle m \rangle$ becomes nonzero. However, for any finite system, thermal fluctuations cause global spin flips, restoring the symmetry and ensuring $\langle m \rangle = 0$ at all temperatures. Therefore, the absolute magnetization $\langle |m| \rangle$ is used as the practical order parameter in finite systems; it is nonvanishing for $T < T_c$ and converges to $|\langle m \rangle|$ as $N \rightarrow \infty$ [30]. The magnetic susceptibility for each order parameter is given by $\chi_{\text{true}} = \beta N(\langle m^2 \rangle - \langle m \rangle^2)$ and $\chi_{\text{practical}} = \beta N(\langle m^2 \rangle - \langle |m| \rangle^2)$, where $\beta \equiv 1/T$ [32]. In finite systems, $\langle m \rangle = 0$, so χ_{true} simplifies to $\beta N \langle m^2 \rangle$.

We introduce the SNBP framework for these two models. SNBP breaks global symmetry by clamping the state of a designated source node x throughout BP iterations, assuming x is always connected to the infinite cluster in percolation or has a fixed spin direction (e.g., positive) in the Ising model. The marginal probability of interest for each system is then the probability that a node is connected to x (percolation) or has a spin aligned with σ_x (Ising). Consequently, the order parameters become the fraction of nodes in x 's cluster (percolation) or magnetization aligned with σ_x (Ising). The source node x can be selected arbitrarily but different choices provide different accuracy improvements. A simple heuristic that works well in practice is to select the highest-degree node as x , since it often belongs to the largest cluster or aligns with the dominant magnetization. This facilitates reliable approximations to practical order parameters in finite systems.

Algorithmically, SNBP is distinguished from conventional BP by the inclusion of Kronecker delta terms that enforce a fixed state on the source node. Removing these terms recovers the conventional BP equations (see the Supplemental Material [33] for details). Throughout, we denote by $\mathcal{N}_i$ the set of neighbors of node i , and by $\mathcal{N}_j \setminus i$ the set of neighbors of j excluding i . The cluster containing the source node x is denoted by $\mathcal{C}(x)$.

C. SNBP for percolation

We derive the SNBP equations by extending conventional BP for percolation [17–19]. Let $\mu_{i \leftarrow j}^{(x)}$ denote the probability that node j belongs to $\mathcal{C}(x)$ when i is removed, and $\mu_i^{(x)}$ the probability that node i belongs to $\mathcal{C}(x)$. Since the source node is always connected to $\mathcal{C}(x)$, $\mu_{i \leftarrow x}^{(x)} = 1$ for $i \in \mathcal{N}_x$ and $\mu_x^{(x)} = 1$. This modification can be incorporated into the message update rule:

$$\mu_{i \leftarrow j}^{(x)} = 1 - (1 - \delta_{jx}) \prod_{k \in \mathcal{N}_j \setminus i} (1 - p \mu_{j \leftarrow k}^{(x)}). \quad (1)$$

The expected fraction of $\mathcal{C}(x)$ is then $\langle S(x) \rangle_{\text{SNBP}} \equiv \langle |\mathcal{C}(x)| \rangle_{\text{SNBP}}/N = \sum_{i=1}^N \mu_i^{(x)}/N$, where the marginal probability $\mu_i^{(x)}$ satisfies

$$\mu_i^{(x)} = 1 - (1 - \delta_{ix}) \prod_{j \in \mathcal{N}_i} (1 - p \mu_{i \leftarrow j}^{(x)}). \quad (2)$$

The message-passing equation for determining the susceptibility in the SNBP framework can be derived from linear response analysis [17,34,35] and is given by

$$\chi_{i \leftarrow j}^{(x)} = \left[1 + \sum_{k \in \mathcal{N}_j \setminus i} \frac{p \chi_{j \leftarrow k}^{(x)}}{1 - p \mu_{j \leftarrow k}^{(x)}} \right] (1 - \mu_{i \leftarrow j}^{(x)}). \quad (3)$$

The expected global susceptibility, representing the mean size of non- $\mathcal{C}(x)$ clusters, is then $\chi_{\text{SNBP}}^{(x)} = \sum_{i=1}^N \chi_i^{(x)}/N$ where

$$\chi_i^{(x)} = \left[1 + \sum_{j \in \mathcal{N}_i} \frac{p \chi_{i \leftarrow j}^{(x)}}{1 - p \mu_{i \leftarrow j}^{(x)}} \right] (1 - \mu_i^{(x)}). \quad (4)$$

D. SNBP for the Ising model

Extending the conventional BP framework [15,22,36], we derive the SNBP equations for the Ising model. Let $h_{i \leftarrow j}^{(x)}$ denote the effective field from spin j to spin i when spin x is fixed to $\sigma_x = +1$. The expected spin of x is fixed to $m_x = 1$ by setting $h_{x \leftarrow i}^{(x)} = \infty$ for $i \in \mathcal{N}_x$. This constraint can be incorporated into the message update rule:

$$\begin{aligned} \tanh(\beta h_{i \leftarrow j}^{(x)}) &= \delta_{ix} + (1 - \delta_{ix}) \tanh(\beta) \\ &\times \tanh \left(\sum_{k \in \mathcal{N}_j \setminus i} \beta h_{j \leftarrow k}^{(x)} \right). \end{aligned} \quad (5)$$

The expected global magnetization aligned with σ_x is then $\langle m \rangle_{\text{SNBP}}^{(x)} = \sum_{i=1}^N m_i^{(x)}/N$, where

$$m_i^{(x)} = \delta_{ix} + (1 - \delta_{ix}) \tanh \left(\sum_{j \in \mathcal{N}_i} \beta h_{i \leftarrow j}^{(x)} \right) \quad (6)$$

is the probability that σ_i aligns with σ_x .

Following the approach in Refs. [34,35,37], we derive the susceptibility propagation under a uniform external field:

$$q_{i \leftarrow j}^{(x)} = Q_{ij}^{(x)} \left(1 + \sum_{k \in \mathcal{N}_j \setminus i} q_{j \leftarrow k}^{(x)} \right), \quad (7)$$

where

$$Q_{ij}^{(x)} = \frac{\tanh(\beta)}{\text{sech}^2(\beta h_{i \leftarrow j}^{(x)})} \left[1 - \frac{\tanh^2(\beta h_{i \leftarrow j}^{(x)})}{\tanh^2(\beta)} \right]. \quad (8)$$

The global susceptibility is then $\chi_{\text{SNBP}}^{(x)} = \sum_{i=1}^N \chi_i^{(x)} / N$, where

$$\chi_i^{(x)} = \beta [1 - (m_i^{(x)})^2] \left(1 + \sum_{j \in \mathcal{N}_i} q_{i \leftarrow j}^{(x)} \right). \quad (9)$$

Detailed derivations of Eqs. (1)–(9) are provided in the Supplemental Material [33]. For direct comparison with percolation, Ising model results are plotted against the equivalent bond occupation probability $p = 1 - e^{-2\beta}$ from the random cluster representation [38].

We also apply this source-node concept to extend MFA and MC simulations. We call these extensions SNMFA and SNMC, respectively. In SNMFA, the state of x is fixed during mean-field updates, and in SNMC, it is fixed during Monte Carlo sampling (see the Supplemental Material [33] for details). All variants retain computational efficiency similar to their conventional forms.

III. RESULTS

A. Trees and almost-trees

We first evaluate the performance of these source-node inference methods on synthetic networks under the bond percolation model, starting with a perfect tree and its perturbations with a few added edges. For the perfect tree, we use a 3-regular Cayley tree of depth 5 ($N = 94$, $M = 93$), where internal nodes have degree 3 and leaves have degree 1.

On the perfect Cayley tree, conventional BP stably converges to a trivial solution, correctly capturing the absence of SSB: $\langle P_\infty \rangle_{\text{BP}} = 0$ [Fig. 1(a)]. The susceptibility computed by conventional BP, χ_{BP} , increases monotonically with p and exactly matches ground-truth MC results, χ_{true} , obtained by assuming $P_\infty = 0$ [Fig. 1(b)]. However, these results are useless for predicting the practical phase-transition metrics, which estimate P_∞ using $S_1 = |C_{\text{max}}|/N$.

Adding a single edge between a randomly chosen node pair in the tree causes conventional BP to exhibit a spurious critical point at $p_c^{\text{BP}} = 1$ [Figs. S1(a) and S1(b) of the Supplemental Material [33]]. This occurs because BP treats the resulting cycle as an infinite one-dimensional chain [18] with a known critical point at $p = 1$ for percolation ($T = 0$ for the Ising model). As a result, $\langle P_\infty \rangle_{\text{BP}}$ converges to zero for $p < 1$, as it does on a tree, but becomes underconstrained at $p_c^{\text{BP}} = 1$. χ_{BP} deviates from χ_{true} in MC simulations and diverges to infinity at $p_c^{\text{BP}} = 1$.

When two or more edges are added to the tree, p_c^{BP} shifts below unity. For $p > p_c^{\text{BP}}$, BP admits both the trivial and a spurious nontrivial solution; however, as the trivial solution is unstable, BP reliably converges to the spurious solution under asymmetric initialization. Consequently, $\langle P_\infty \rangle_{\text{BP}}$ becomes nonzero for $p > p_c^{\text{BP}}$ [Fig. 1(c) and Fig. S1(c) of the Supplemental Material [33]], and χ_{BP} displays a divergent peak at p_c^{BP} [Fig. 1(d) and Fig. S1(d) of the Supplemental Material [33]]. Although adding a few cycles to the tree improves BP's




FIG. 1. (a), (c) Percolation strength P_∞ and (b), (d) susceptibility χ as functions of occupation probability p for (a), (b) a 3-regular Cayley tree ($N = 94$, $M = 93$) and (c), (d) the same tree with two additional edges. Symbols indicate results from MC and SNMC. Curves show message-passing results inferred by MFA, SNMFA, BP, and SNBP. The curve “MC ($P_\infty = S_1$)” uses S_1 as the order parameter and $\chi_{\text{practical}}$ as the susceptibility, while “MC ($P_\infty = 0$)” computes χ_{true} . The red node in each network marks the source node, chosen to be of the highest degree.

estimation of practical metrics, it is still inadequate near and below p_c^{BP} .

In contrast, SNBP dramatically reduces this gap between theory and numerical experiments by explicitly breaking the symmetry. It admits only a nontrivial solution, which accurately captures the nonzero $\langle S_1 \rangle_{\text{MC}}$ and the smooth peak of χ_{MC} (Fig. 1). By construction, SNBP exactly matches SNMC whenever the removal of the source node yields a perfect tree [Figs. 1(a) and 1(b) and Figs. S1(a) and S1(b) of the Supplemental Material [33]]. This exactness is rooted in the same principle as conditioned BP (CBP), which achieves exactness by clamping nodes to remove all cycles [39–41]. However, unlike SNBP, CBP does not directly address the issue of global symmetry, and its computational complexity grows exponentially with system size due to scanning over all possible clamped values. Furthermore, as demonstrated here, sophisticated clamping schemes are often unnecessary when dealing with percolation and Ising models, which are widely studied contexts for belief propagation methods [13,17,18,20,26]. Our SNBP results indicate that fixing a single, well-chosen node is often sufficient to obtain highly accurate approximations. Moreover, even in the inexact regime when an extra cycle remains, SNBP still closely predicts SNMC [Figs. 1(c) and 1(d)]. Since the highest-degree node often lies in the largest cluster, SNMC—and thus SNBP—reliably approximates the practical metrics of conventional MC.

Although the naive MFA predicts a nonzero order parameter and a susceptibility peak even on a perfect tree [Figs. 1(a) and 1(b)], these predictions deviate significantly from MC results. This inaccuracy persists despite the addition




FIG. 2. Inference on locally treelike networks. Order parameters for the percolation model are shown in panels (a), (c), and (d), and for the Ising model in panel (b). Simulations were run for the Norwegian board directors network ($N = 179$, $M = 184$), an Erdos-Renyi random graph ($N = 94$, $M = 139$), and a scale-free network ($N = 80$, $M = 156$) generated using the Barabasi-Albert model with $k = 2$ out-edges per arriving node. The red node in each network marks the source node.

of several cycles [Figs. 1(c) and 1(d) and Figs. S1(c) and S1(d) of the Supplemental Material [33]]. Its source-node variant, SNMFA, is similarly inaccurate, except for slightly reduced finite-size errors at small p and nondiverging susceptibility. Overall, SNBP clearly outperforms both MFA and SNMFA in accuracy for practical metrics on trees and almost-trees.

B. Locally tree-like networks

The high accuracy of SNBP, demonstrated on a Cayley tree and its perturbations, generalizes to various locally tree-like networks. On the real-world Norwegian board directors network, an almost-tree, SNBP shows superior accuracy for the order parameter [Figs. 2(a) and 2(b)] and susceptibility [Figs. S2(a) and S2(b) of the Supplemental Material [33]]. A minor deviation occurs only for the Ising model at high temperature (small $p = 1 - e^{-2\beta}$), where $\langle m \rangle_{\text{SNBP}} \rightarrow 1/N$ but $\langle |m| \rangle_{\text{MC}} \rightarrow (2/(\pi N))^{1/2}$ (derived from the one-dimensional random-walk result of Ref. [42]). In contrast, SNBP is exact for percolation in this limit: As $p \rightarrow 0$, both $\langle \mathcal{C}(x) \rangle_{\text{SNBP}}$ and $\langle S_1 \rangle_{\text{MC}}$ approach $1/N$.

Even when networks contain many cycles—as quantified by a large cyclomatic number $c(\mathcal{G}) \equiv M - N + 1$ —SNBP maintains high accuracy if the network is locally treelike. This accuracy stems from the established principle that conventional BP provides highly accurate inference in such networks [15,22]. Our simulations show that SNBP outperforms the other inference methods in Erdős-Rényi random networks [Fig. 2(c) and Fig. S2(c) of the Supplemental Material [33]]




FIG. 3. Inference for the percolation strength on spatial networks: (a) 8×8 square lattice and (b) random geometric graph ($N = 100$, $\langle k \rangle = 5.9$). Both BP and SNBP fail to reproduce MC results, showing significant deviations over a wide range of p . Red node in networks marks source node.

and scale-free networks [Fig. 2(d) and Fig. S2(d) of the Supplemental Material [33]]. Although their cyclomatic numbers are not small, the sparsity of short cycles enables SNBP to remain accurate.

C. Spatial networks

In networks that are not locally treelike and contain many short cycles, such as spatial networks including lattices [Fig. 3(a)] and random geometric graphs [Fig. 3(b)], both BP and SNBP deviate significantly from MC results. Although SNBP is more accurate than BP for $p < p_c^{\text{BP}}$, its overall accuracy is still limited by errors from cycles—a well-known issue in loopy BP. These cycle-driven errors can be overcome by more computationally intensive algorithms that explicitly account for neighborhood correlations [19–21,26].

D. Real-world networks

To systematically evaluate the performance of SNBP, we benchmark its accuracy against conventional BP and MFA on 139 real-world networks (details and links to data sources are provided in the Supplemental Material [33]). For each network, we extract the largest connected component and convert it to a simple graph by removing directionality, duplicate edges, and self-loops. Accuracy is quantified by the mean absolute error in the order parameter relative to conventional MC simulations, $\Delta_{\text{Method-MC}}$, which is calculated as the area between the curves from $p = 0$ to 1 (see the Supplemental Material [33]).

For both the percolation [Fig. 4(a)] and Ising [Fig. 4(b)] models, SNBP consistently outperforms conventional BP across the real-world networks. The improvement is most pronounced in networks with a low cyclomatic number, for which conventional BP suffers from the aforementioned “tree-like network catastrophe” and can provide even less accurate results than MFA [Figs. S4(a) and S4(b) of the Supplemental Material [33]].

Unlike conventional BP, SNBP outperforms MFA across all tested networks [Figs. 4(c) and 4(d)]. The performance gap increases with decreasing $\langle k \rangle$, highlighting SNBP’s particular advantage in weakly connected networks where mean-field assumptions break down.




FIG. 4. Comparison of SNBP errors with BP and MFA errors for (a), (c) the percolation model and (b), (d) the Ising model on 139 real networks. (a), (b) SNBP outperforms BP: $\Delta_{\text{SNBP-MC}}$ is generally smaller than $\Delta_{\text{BP-MC}}$, particularly for small cyclomatic numbers. (c), (d) SNBP outperforms MFA: $\Delta_{\text{SNBP-MC}}$ is almost always smaller than $\Delta_{\text{MFA-MC}}$. The gap increases as $\langle k \rangle$ decreases.

IV. DISCUSSION

Here, we address a fundamental yet often overlooked challenge in applying BP to finite, symmetric phase-transition models: BP's exactness on trees, which correctly predicts the absence of spontaneous symmetry breaking, prevents it from effectively capturing established practical order parameters [18]. To bridge this gap between theory and practice, we propose SNBP, which breaks global symmetry by fixing a single source node. Our tests on percolation and Ising models show that SNBP accurately infers practical order parameters and their susceptibilities in a broad range of real and synthetic networks. SNBP most significantly outperforms conventional BP in networks with few cycles and naive MFA in sparse

networks. Importantly, this enhanced accuracy is achieved without sacrificing the computational efficiency of standard BP.

Opportunities remain for further improvement and broader application. Like conventional BP, SNBP's accuracy is limited by the presence of short cycles. This can be mitigated by integrating the source-node method with more sophisticated BP algorithms for loopy graphs [19–21,26], albeit at the cost of increased computational expense. Additionally, in networks lacking a clear hub, selecting the highest-degree node as the source may not be optimal. Indeed, Figs. 3 and 4 show cases where SNBP fails to improve inference accuracy over standard BP, which is a potential issue in networks with highly homogeneous degrees, such as spatially embedded networks. Future research could explore alternative strategies for choosing one or more source nodes to improve accuracy in such cases.

Moreover, SNBP shows promise for extension to other symmetric phase-transition models, including the q -state Potts model [38], XY model [43], and Heisenberg model [44]. Beyond these, the core principle of SNBP—strategically fixing node or edge values for practical message-passing approximations—could potentially be adapted to diverse inference tasks on networks, such as community detection [9,10], as well as various deep learning tasks [45–47]. However, it is unclear whether the choice of a single source node in SNBP is sufficient for these more challenging settings due to their complex symmetries—these cases may be best handled by more flexible, adaptive message-passing modifications [39–41].

DATA AVAILABILITY

All empirical networks used in this study were obtained from the Netzschleuder repository [48], and details about the networks are given in the Supplemental Material [33]. The data that support the findings of this article are openly available [48].

- [1] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference* (Morgan Kaufmann, San Francisco, CA, 1988).
- [2] M. J. Wainwright and M. I. Jordan, Graphical models, exponential families, and variational inference, *Found. Trends Mach. Learn.* **1**, 1 (2008).
- [3] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*, *Adaptive Computation and Machine Learning* (MIT Press, Cambridge, MA, 2010).
- [4] F. Kschischang, B. Frey, and H.-A. Loeliger, Factor graphs and the sum-product algorithm, *IEEE Trans. Inf. Theory* **47**, 498 (2001).
- [5] M. I. Jordan, Graphical models, *Stat. Sci.* **19**, 140 (2004).
- [6] E. Alpaydm, *Introduction to Machine Learning*, *Adaptive Computation and Machine Learning*, 2nd ed. (MIT Press, Cambridge, MA, 2010).
- [7] B. Karrer and M. E. J. Newman, Message passing approach for general epidemic models, *Phys. Rev. E* **82**, 016101 (2010).
- [8] F. Altarelli, A. Braunstein, L. Dall'Asta, A. Lage-Castellanos, and R. Zecchina, Bayesian inference of epidemics on networks via belief propagation, *Phys. Rev. Lett.* **112**, 118701 (2014).
- [9] M. B. Hastings, Community detection as an inference problem, *Phys. Rev. E* **74**, 035102(R) (2006).
- [10] A. Decelle, F. Krzakala, C. Moore, and L. Zdeborová, Inference and phase transitions in the detection of modules in sparse networks, *Phys. Rev. Lett.* **107**, 065701 (2011).
- [11] H. Nguyen and R. Zheng, On budgeted influence maximization in social networks, *IEEE J. Sel. Areas Commun.* **31**, 1084 (2013).
- [12] S. Pei, J. Wang, F. Morone, and H. A. Makse, Influencer identification in dynamical complex systems, *J. Complex Netw.* **8**, cnz029 (2020).

- [13] H. Nishimori, *Statistical Physics of Spin Glasses and Information Processing: An Introduction*, 1st ed. (Oxford University Press, Oxford, 2001).
- [14] L. Zdeborová and F. Krzakala, Statistical physics of inference: Thresholds and algorithms, *Adv. Phys.* **65**, 453 (2016).
- [15] M. Mézard and A. Montanari, *Information, Physics, and Computation* (Oxford University Press, New York, USA, 2009).
- [16] K. P. Murphy, Y. Weiss, and M. I. Jordan, Loopy belief propagation for approximate inference: An empirical study, in *Proceedings of the 15th Conference on Uncertainty in Artificial Intelligence, UAI'99* (Morgan Kaufmann, San Francisco, CA, 1999), pp. 467–475.
- [17] B. Karrer, M. E. J. Newman, and L. Zdeborová, Percolation on sparse networks, *Phys. Rev. Lett.* **113**, 208702 (2014).
- [18] A. Allard and L. Hébert-Dufresne, On the accuracy of message-passing approaches to percolation in complex networks, *arXiv:1906.10377*.
- [19] G. T. Cantwell and M. E. J. Newman, Message passing on networks with loops, *Proc. Natl. Acad. Sci. USA* **116**, 23398 (2019).
- [20] G. T. Cantwell, A. Kirkley, and F. Radicchi, Heterogeneous message passing for heterogeneous networks, *Phys. Rev. E* **108**, 034310 (2023).
- [21] P. Mann and S. Dobson, Belief propagation on networks with cliques and chordless cycles, *Phys. Rev. E* **107**, 054303 (2023).
- [22] M. E. J. Newman, Message passing methods on complex networks, *Proc. R. Soc. A* **479**, 20220774 (2023).
- [23] J. M. Mooij and H. J. Kappen, On the properties of the Bethe approximation and loopy belief propagation on binary networks, *J. Stat. Mech.* (2005) P11012.
- [24] S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes, Critical phenomena in complex networks, *Rev. Mod. Phys.* **80**, 1275 (2008).
- [25] S. Yoon, A. V. Goltsev, S. N. Dorogovtsev, and J. F. F. Mendes, Belief-propagation algorithm and the Ising model on networks with arbitrary distributions of motifs, *Phys. Rev. E* **84**, 041144 (2011).
- [26] A. Kirkley, G. T. Cantwell, and M. E. J. Newman, Belief propagation for networks with loops, *Sci. Adv.* **7**, eabf1211 (2021).
- [27] S. N. Dorogovtsev, A. V. Goltsev, and J. F. F. Mendes, Correlations in interacting systems with a network topology, *Phys. Rev. E* **72**, 066130 (2005).
- [28] J. D. Fraser, Spontaneous symmetry breaking in finite systems, *Philos. Sci.* **83**, 585 (2016).
- [29] D. Stauffer and A. Aharony, *Introduction to Percolation Theory*, 2nd ed. (Taylor & Francis, London, 1994).
- [30] M. E. J. Newman and G. T. Barkema, *Monte Carlo Methods in Statistical Physics* (Oxford University Press, New York, USA, 1999).
- [31] S. B. Lee, Universal behavior of the amplitude ratio of percolation susceptibilities for off-lattice percolation models, *Phys. Rev. E* **53**, 3319 (1996).
- [32] D. P. Landau and K. Binder, *A Guide to Monte Carlo Simulations in Statistical Physics*, 4th ed. (Cambridge University Press, Cambridge, 2014).
- [33] See Supplemental Material at <http://link.aps.org/supplemental/10.1103/pxg-mhjr> for derivations of BP and MFA equations, details of the MC algorithm, definitions of the error metrics, and details of the network datasets.
- [34] M. Welling and Y. W. Teh, Linear response algorithms for approximate inference in graphical models, *Neural Comput.* **16**, 197 (2004).
- [35] M. Mézard and T. Mora, Constraint satisfaction problems and neural networks: A statistical physics perspective, *J. Physiol. Paris* **103**, 107 (2009).
- [36] H. C. Nguyen, R. Zecchina, and J. Berg, Inverse statistical problems: From the inverse Ising problem to data science, *Adv. Phys.* **66**, 197 (2017).
- [37] E. Aurell, C. Ollion, and Y. Roudi, Dynamics and performance of susceptibility propagation on synthetic data, *Eur. Phys. J. B* **77**, 587 (2010).
- [38] F. Y. Wu, The Potts model, *Rev. Mod. Phys.* **54**, 235 (1982).
- [39] J. Pearl, Fusion, propagation, and structuring in belief networks, *Artif. Intell.* **29**, 241 (1986).
- [40] B. Bidyuk and R. Dechter, Cutset sampling for Bayesian networks, *J. Artif. Intell. Res.* **28**, 1 (2007).
- [41] F. Eaton and Z. Ghahramani, Choosing a variable to clamp, in *Proceedings of the 12th International Conference on Artificial Intelligence and Statistics*, edited by D. van Dyk and M. Welling (PMLR, Clearwater Beach, FL, 2009), Vol. 5, pp. 145–152.
- [42] W. Feller, *An Introduction to Probability Theory and its Applications* (Wiley, Oxford, England, 1950), Vol. 1.
- [43] C. Lupo and F. Ricci-Tersenghi, Approximating the XY model on a random graph with a q -state clock model, *Phys. Rev. B* **95**, 054433 (2017).
- [44] L. M. Del Bono, F. Nicoletti, and F. Ricci-Tersenghi, Analytical solution to Heisenberg spin glass models on sparse random graphs and their de Almeida–Thouless line, *Phys. Rev. B* **110**, 184205 (2024).
- [45] C. Baldassi, F. Pittorino, and R. Zecchina, Shaping the learning landscape in neural networks around wide flat minima, *Proc. Natl. Acad. Sci. USA* **117**, 161 (2020).
- [46] C. Lucibello, F. Pittorino, G. Perugini, and R. Zecchina, Deep learning via message passing algorithms based on belief propagation, *Mach. Learn.: Sci. Technol.* **3**, 035005 (2022).
- [47] H. Huang, Statistical mechanics of neural networks, in *Statistical Mechanics of Neural Networks*, edited by H. Huang (Springer Nature, Singapore, 2021), pp. 99–109.
- [48] T. P. Peixoto, The Netzscheuler network catalogue and repository, Zenodo (2020), <http://doi.org/10.5281/zenodo.7839980>.