



Generalization of graph network inferences in higher-order graphical models

Yicheng Fei¹ · Xaq Pitkow^{2,3,4,5,6} 

Received: 23 June 2022 / Revised: 5 April 2023 / Accepted: 19 August 2023

© The Author(s) 2023

Abstract

Probabilistic graphical models provide a powerful tool to describe complex statistical structure, with many real-world applications in science and engineering from controlling robotic arms to understanding neuronal computations. A major challenge for these graphical models is that inferences such as marginalization are intractable for general graphs. These inferences are often approximated by a distributed message-passing algorithm such as Belief Propagation, which does not always perform well on graphs with cycles, nor can it always be easily specified for complex continuous probability distributions. Such difficulties arise frequently in expressive graphical models that include intractable higher-order interactions. In this paper we define the Recurrent Factor Graph Neural Network (RF-GNN) to achieve fast approximate inference on graphical models that involve many-variable interactions. Experimental results on several families of graphical models demonstrate the out-of-distribution generalization capability of our method to different sized graphs, and indicate the domain in which our method outperforms Belief Propagation (BP). Moreover, we test the RF-GNN on a real-world Low-Density Parity-Check dataset as a benchmark along with other

✉ Xaq Pitkow
xaq@cmu.edu

Yicheng Fei
yf17@rice.edu

¹ Department of Physics and Astronomy, Rice University, 6100 Main St, Houston, TX 77005, USA

² Neuroscience Institute, Carnegie Mellon University, Pittsburgh, PA 15213, USA

³ Department of Machine Learning, Carnegie Mellon University, Pittsburgh, PA 15213, USA

⁴ Department of Electrical and Computer Engineering, Rice University, 6100 Main St, Houston, TX 77005, USA

⁵ Department of Neuroscience, Baylor College of Medicine, One Baylor Plaza, Houston, TX 77030, USA

⁶ Center for Neuroscience and Artificial Intelligence, Baylor College of Medicine, One Baylor Plaza, Houston, TX 77030, USA

baseline models including BP variants and other GNN methods. Overall we find that RF-GNNs outperform other methods under high noise levels.

Keywords Graph neural network · Probabilistic inference · Third-order · Probabilistic graphical model

Mathematics Subject Classification 60 · 62 · 65 · 68

1 Introduction

To draw conclusions about individual variables of interest in a task, we must marginalize out all other unobserved variables. Such exact inference computations are often infeasible in high-dimensional latent spaces, due to their exponential complexity. Conveniently, the latent space to be marginalized is often decomposable due to conditional dependencies between variables, a structure that can be described by a probabilistic graphical model (PGM) (Koller and Friedman 2009).

This decomposable structure may allow us to perform difficult global calculations using simpler computations on subsets of variables. This makes probabilistic graphical models a compelling framework for describing both machine and biological intelligence. Early graphical models used pairwise interactions, as in Hopfield networks (Hopfield 1982) and Boltzmann machines (Sherrington and Kirkpatrick 1975), to model inference, learning, and memory. However, basic pairwise interaction models may not be compact descriptions of complex data patterns seen in real data, whether in machine learning (Ranzato and Hinton 2010; Hinton 2010), neuroscience (Beggs and Plenz 2003; Ganmor et al. 2011; Shimazaki et al. 2015), biochemical networks (Ritz et al. 2014), or social networks (Centola et al. 2018; Milojević 2014; Iacopini et al. 2019). By including higher-order interactions we increase model flexibility, but the number of possible interactions grows combinatorially with the interaction order, which contributes to the difficulty of exact inference in such models. This difficulty can be reduced by reasonable prior assumptions about locality and sparsity of natural interactions, resulting in a sparse higher-order graphical model.

Message-passing algorithms take advantage of this sparse graph structure to simplify computations. Such approaches are used by algorithms like Belief Propagation (BP) (Pearl 1988) and Expectation Propagation (EP) (Minka 2001), which are widely used approaches to computing or approximating marginal probabilities using distributed computation. BP is guaranteed to yield exact results if the graph has a tree structure. However, on general graphs with loops, which are likely to be better descriptors of real data, these algorithms can make substantial approximation errors or even fail to converge.

Unfortunately, higher-order factor graph may have more loops than pairwise graphs with as many interactions, and since standard local message-passing algorithms like BP suffer in the presence of loops they are likely to perform worse on many real-world graphical models. Even when applied to higher-order trees, message updates for BP usually don't have closed-form solutions, so running exact BP on these graphs becomes impractical.


To mitigate these drawbacks of algorithms like BP and to provide an alternative on loopy graphs without analytical update formulas, in this work we present the Recurrent Factor Graph Neural Network (RF-GNN), a flexible recursive message-passing algorithm for fast approximate inference. Our method applies to a large family of higher-order graphical models with a wide range of graph structures and parameter values, including very loopy ones. In the spirit of message-passing algorithms like BP, we use a Graph Neural Network (GNN) (Scarselli et al. 2008; Li et al. 2015) to perform message-passing on factor graphs iteratively. We train this network to compute sufficient statistics of all univariate marginal probabilities simultaneously for each instance of a distribution generated from a parametrized family of PGMs.

To study the performance of RF-GNNs, we run numerical experiments on two artificial datasets where we can calculate ground truth: Gaussian Graphical Models (GGM) and a small binary spin-glass system with third-order interactions. GGMs only have pairwise interactions, but the ground truth marginals can easily be computed without message-passing for very large graphs. In addition, we construct a dataset of more complex continuous PGMs with pairwise and third-order interactions. Since closed-form marginals do not exist for this interesting model class, we train an RF-GNN to predict univariate statistics estimated by Markov Chain Monte Carlo (MCMC) sampling. We also test the performance of the RF-GNN on a Low-Density Parity-Check (LDPC) decoding dataset with a very loopy factor graph as a real-world example. We compare our recurrent approach with an alternative, the Factor Graph Neural Network (FGNN) (Zhang et al. 2020), which is a graph-structured feedforward network with a fixed number of GNN layers. In order to compare the performance when generalizing to larger unseen graphs, we also train FGNNs on our Gaussian Graphical Model dataset with diverse graph structures in addition to the LDPC dataset with a fixed structure.

Our experiments show that a trained RF-GNN has better performance than both BP and FGNN on an in-distribution test dataset, even when restricting the comparison to only those graphs for which BP dynamics converge. We also show that our model generalizes reliably out of distribution to probabilistic graphical models of different sizes than the training set. By looking at how the error distribution depends on two graph metrics—average shortest path length and cluster coefficient—we find that an RF-GNN outperforms BP particularly well on graphs with small average shortest path length and large cluster coefficients, which are common properties of many real world graphs (Watts and Strogatz 1998). This suggests there is potential for using an RF-GNN as an approximation inference method on real-world PGMs with higher-order interactions.

In Sect. 2 we provide key background material about probabilistic graphical models and Graph Neural Networks. In Sect. 3 we first define our Recurrent Factor Graph Neural Network as a message-passing algorithm to calculate approximate marginal distribution in a family of graphical models. Then we describe four types of graphical models with different attributes and complexity on which we tested our model. In Sect. 4 we show the performance of RF-GNNs on these four datasets and compare with Belief Propagation. Of particular interest, we analyze how RF-GNNs generalize to larger graphs and how different graph attributes affect generalization performance. In Sect. 5 we discuss related message-passing inference algorithms for loopy graphs,

Fig. 1 Diagram of a factor graph with 4 variable nodes and 6 factor nodes. Factors 1 and 2 are three-way factors each involving three variables. Factors 3–6 are singleton factors each only involving one variable



and discuss other types of GNNs that apply to graphs with higher-order structure. Finally, in Sect. 6 we discuss potential limitations and extensions of our framework.

2 Background

2.1 Probabilistic graphical models

Probabilistic graphical models describe multivariate probability distributions using graphs in which nodes represent variables. Nodes are connected if the corresponding variables are statistically dependent when conditioning on all other variables. Different types of graphs represent different factorization structures or conditional independence relationships within the joint probability density. Examples include Bayesian networks with directed acyclic graphs, Markov random fields using undirected graphs, and factor graphs which are undirected bipartite graphs connecting variable nodes and factor (interaction) nodes.

Here we concentrate on the latter type, the factor graph (Frey et al. 1997), which expresses a joint probability density as a product of local factors, $p(\mathbf{x}) \propto \prod_{\alpha} f_{\alpha}(\mathbf{x}_{\alpha})$, each involving only a subset of variables \mathbf{x}_{α} . These models introduce a second type of node in the graph, the factor node, one for each local factor f_{α} . Each factor node is only connected to the variable nodes involved in the interactions. An example of such a factor graph is depicted in Fig. 1.

In this work we consider probabilistic graphical models within the exponential family (Pitman 1936; Koopman 1936), a broad class of probability distributions widely used in statistical analysis. The density of such a distribution can be parameterized by a vector of natural parameters ϕ ,

$$p(\mathbf{x}) = \frac{1}{Z(\phi)} \exp(\phi^{\top} \mathbf{T}(\mathbf{x})), \quad (1)$$

where $\mathbf{T}(\mathbf{x})$ is the vector of sufficient statistics and $Z(\phi)$ is a normalization constant. In general, any single component of $\mathbf{T}(\mathbf{x})$ can be an arbitrary function of all \mathbf{x} . In a factor graph, however, these functions are restricted, with each involving only a subset of variables, \mathbf{x}_{α} , that are connected through a factor node. Each individual factor can

then be parametrized with its own vector of natural parameters ϕ_α :

$$p(\mathbf{x}) = \frac{1}{Z(\phi)} \prod_{\alpha} \exp(\phi_{\alpha}^{\top} \mathbf{T}_{\alpha}(\mathbf{x}_{\alpha})). \quad (2)$$

A PGM is then represented by sets of factor nodes, variable nodes, edges, and natural parameters as features on factor nodes, $\{\mathcal{F}, \mathcal{V}, \mathcal{E}, \{\phi_{\alpha}\}\}$.

Another, more familiar, representation of a PGM is in terms of expectations of the statistics $\mathbb{E}_{p(\mathbf{x})}[\mathbf{T}(\mathbf{x})]$. These expectations are called expectation parameters, such as means and variances for a Gaussian. When combined with a maximum entropy principle, the expectation parameters implicitly fully specify the probability distribution. Expectation parameters and natural parameters are closely related to each other: the natural parameters are the Lagrange multipliers that enforce those expectations.

Relating these parameters is hard, but often important: the process of inference can be viewed, essentially, as converting between natural parameters and expectation parameters (Wainwright et al. 2008; Koller and Friedman 2009). Marginalization is one example of computing expectations of a subset of variables given natural parameters for the full PGM. The difficulty of this computation is what motivated our efforts to approximate it by a graph network.

Message-passing algorithms like belief propagation exploit the conditional independence between variables implied by the graph structure so that the multidimensional integration could be performed separately in orthogonal subspaces. In this way, the computation complexity only depends on the maximum node degree and is independent of graph size. However, BP updates assume all neighbors are independent, which is not true in loopy graphs. Yedidia et al. (2000) shows that BP only converges to fixed points of the Bethe free energy, which is an approximation of the Gibbs free energy that only involves pairwise interactions. By generalizing the Bethe free energy to Kikuchi free energy which involves more complex clusters of nodes and corresponding messages between clusters, Yedidia et al. (2000) defines a generalized BP (GBP) algorithm that achieves better convergence and more accurate marginals on graphs with tight loops like lattices. Our method currently works on regular factor graphs, but could be extended to more complicated Kikuchi clustering approximations.

Exact inference in tree graphs can be performed by iteratively marginalizing out the leaves of the tree and propagating this information along the graph. This iterative algorithm is called belief propagation (BP), and has been applied with some success even on graphs with loops (Pearl 1988; Bishop 2006). Message updates in BP include multivariate integration, and don't always have closed-form solutions. (Fully factorized) Expectation Propagation (EP) (Minka 2001) is an approximation to mitigate this issue by projecting the outgoing message to some convenient user-chosen parametric family.

2.2 Graph neural networks

Graph Neural Networks (GNN) are artificial neural networks implementing a message-passing operation on a graph (Scarselli et al. 2008). A GNN updates each node's

representation based on aggregated messages from its neighbors. Each node i represents information as a vector \mathbf{h}_i^t that evolves over time (or layer) t , and edges are assigned a vector weight \mathbf{e}_{ij} . The updated representation at time or layer $t + 1$ can be described by:

$$\mathbf{h}_i^{t+1} = \mathcal{U}\left(\mathbf{h}_i^t, \bigsqcup_{j \in N(i) \setminus \{j\}} \mathcal{M}(\mathbf{h}_j^t, \mathbf{h}_i^t, \mathbf{e}_{ij})\right) \quad (3)$$

where every message $\mathcal{M}(\mathbf{h}_j^t, \mathbf{h}_i^t, \mathbf{e}_{ij})$ from neighbor j to node i along edge ij is first calculated using a common trainable nonlinear message function \mathcal{M} , then messages from all neighbors are combined by a permutation-invariant aggregation function \bigsqcup (e.g. summation), before being used to update each target node through another trainable update function \mathcal{U} . Update functions based on a Gated Recurrent Unit (GRU) (Cho et al. 2014) or Long Short Term Memory (Hochreiter and Schmidhuber 1997) provide a long term memory for each node state. Note that in other applications, the superscript t can index a feedforward layer in a stacked GNN, where different layers have different weights for the message and update function \mathcal{M} and \mathcal{U} (Veličković et al. 2017; Kipf and Welling 2016). Here, however, we use a recurrent GNN implementation where t represents time, although this is equivalent to tying these layers' weights as in an unrolled Recurrent Neural Network (RNN) with message and update functions and factor parameters that are shared across layers (Li et al. 2015; Gilmer et al. 2017). After passing messages throughout the graph for a certain number of time steps, a global or local readout network is used to extract the information we need, depending on the task. We recommend (Zhou et al. 2020) for a review of methods and applications of GNNs.

3 Methods

3.1 Recurrent factor graph neural network

To make more accurate inferences on probabilistic models with loopy factor graphs, we present the Recurrent Factor Graph Neural Network (RF-GNN), a network that takes in a factor graph representation of a probabilistic model from an exponential family and predicts approximate single-variate marginal distributions of every variable. The RF-GNN learns a message-passing algorithm on the factor graph in which message functions and updates are executed by neural networks rather than multivariate integration. When run, this message-passing alters its node states iteratively as a dynamical system, operating essentially as a structured RNN, that upon convergence has node states that encode approximate marginals. As a flexible message-passing algorithm, once trained, an RF-GNN can be applied to graphical models in the same parametric family with various sizes and structures.

Using the diagram in Fig. 2, we now explain the core overall structure of an RF-GNN. After this overview, we will then describe the details of our network that we chose to create a concrete implementation of an RF-GNN. Firstly, in order to define this class of message-passing algorithms on factor graphs, we define a GNN on generic

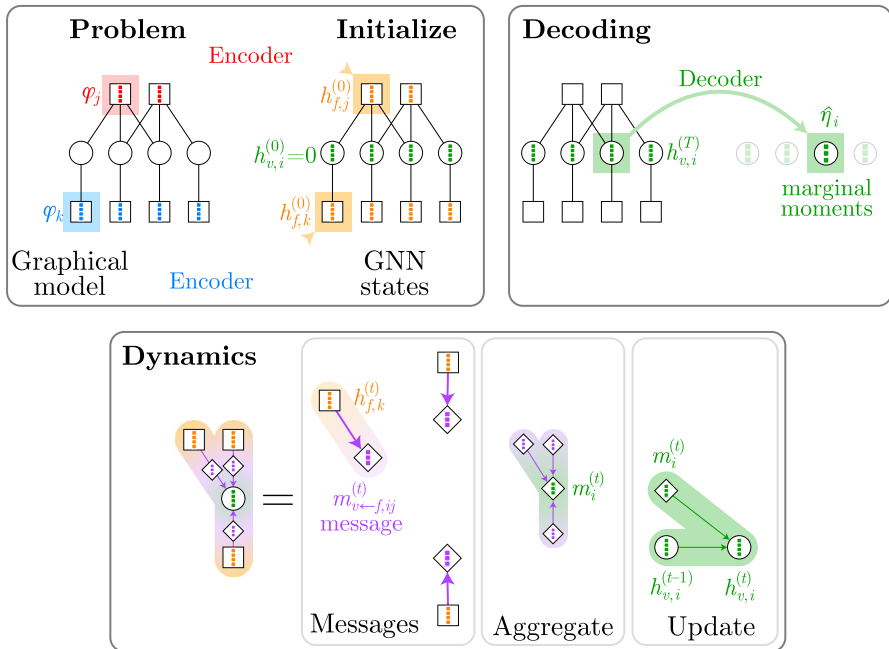


Fig. 2 Diagram of an RF-GNN. Firstly, the variable and factor node features of the input PGM ϕ_j, ϕ_k s are encoded into latent features $h_{f,j}^{(0)}, h_{v,i}^{(0)}$ using two separate encoders. Then a bidirectional message-passing network specified by the message and update neural networks recurrently update node representations. Finally, we decode the desired features of all single-variate marginal distributions $\hat{\eta}_i$ with a common decoder network

bipartite graphs with two distinct types of nodes, unlike a typical GNN that treats all nodes equally (Kipf and Welling 2016; Gori et al. 2005; Li et al. 2015). Our two node types are for variables and factors, represented by hidden vectors $\mathbf{h}_{v,i}^{(t)}$ and $\mathbf{h}_{f,i}^{(t)}$, with subscripts v and f for variable nodes and factor nodes. At $t = 0$, every node's latent state is initialized, and then the variable and factor node states $\mathbf{h}_{v,i}^{(t+1)}, \mathbf{h}_{f,i}^{(t+1)}$ are updated recursively from the values of the previous time step through message-passing. The message-passing alternates between updating variable nodes and factor nodes using their respective message and update functions. Each stage follows standard GNN dynamics, as shown in Fig. 2 (Dynamics) and described above (Eq 3). After T time steps determined by a convergence criterion, we use a common decoder network to decode the target from each variable $\mathbf{h}_{v,i}^{(T)}$ simultaneously. In this work, we train our decoder network to predict the marginal moments of our training data from node state $\mathbf{h}_{v,i}^{(T)}$.

3.2 Architectural details for an RF-GNN

Given this general algorithmic framework, we will now explain our design choices for the algorithm's components. In principle, our autonomous dynamical system could

use either continuous time, like Neural Ordinary Differential Equations (neural ODEs) (Chen et al. 2018), or discrete time, like many RNNs. We opt for a discrete-time RNN for computational efficiency. We choose GRUs (Cho et al. 2014) for the update operations as they have proven to be expressive enough while remaining more computationally efficient than LSTMs (Hochreiter and Schmidhuber 1997). We add LayerNorm (Ba et al. 2016) to the GRU state update, since that helps regularize the range of values in intermediate timesteps and avoid the common problem that activations tend to grow or shrink in a recurrent system.

We now describe the mathematical details of an RF-GNN, following the flow of computation. The inputs of an RF-GNN are graphical models $\{\mathcal{F}, \mathcal{V}, \mathcal{E}, \{\phi_\alpha\}\}$ within the exponential family, parametrized by their natural parameters $\{\phi_\alpha\}$, as described in Eq. (2).

There are two ways of incorporating local potentials into the latent representations. One way is to start from a common initialization independent of the potentials, and provide the potential's parameters as additional inputs for every step in the message and update functions, just like belief propagation. This initialization is simpler in one way, but then we need more complex message functions and update functions that take the local potentials as inputs. Another way is to embed the potentials in the *initial states* of all factor nodes, and let the dynamics run autonomously from this initial condition. The potentials then only affect the inference by how the initial state shapes the dynamics. This actually corresponds to an interpretation of belief propagation as iterative reparameterization (Wainwright et al. 2003). Here we choose a combination of both methods, as described in Eqs. (5), (6) and (8). The representation of each factor node is divided into two parts: a dynamic latent state $\mathbf{h}_{f,\alpha}^{(t)}$ and a static feature matrix \mathbf{F}_α . The first is recurrently updated through message-passing, while the second serves as a constant input in the update function to transform the aggregated messages Eq. (8). We use two separate multilayer perceptron (MLP) (Rosenblatt 1958) encoders to map the natural parameters into the initial state $\mathbf{h}_{f,\alpha}^{(0)}$ and the feature matrix \mathbf{F}_α . Different network weights are used for different factor types.

$$\mathbf{h}_{f,\alpha}^{(0)} = \text{Encoder1}(\phi_\alpha) \quad (4)$$

$$\mathbf{F}_\alpha = \text{Encoder2}(\phi_\alpha) \quad (5)$$

The hidden states for variable nodes only have a recurrently updated part, which we initialize with zero vectors:

$$\mathbf{h}_{v,j}^{(0)} = \mathbf{0} \quad (6)$$

At each iteration t , we calculate the message from variable j to factor α as a linear projection of $\mathbf{h}_{v,j}^{(t-1)}$, which is then averaged with messages from all neighbors $j \in N(\alpha)$ to form a summary message. The summary message is then transformed by a factor feature matrix \mathbf{F}_α in a similar way as (Gilmer et al. 2017) to produce the activation $\mathbf{a}_\alpha^{(t)}$. We input $\mathbf{a}_\alpha^{(t)}$ along with the old state $\mathbf{h}_{f,\alpha}^{(t-1)}$ into the GRU unit at factor

node α to calculate the new latent state $\mathbf{h}_{f,\alpha}^{(t)}$.

$$\mathbf{m}_{f,j \rightarrow \alpha}^{(t)} = \mathbf{W}_{f,\alpha} \mathbf{h}_{v,j}^{(t-1)} \quad \text{message} \quad (7)$$

$$\mathbf{a}_{f,\alpha}^{(t)} = \mathbf{F}_{\alpha} \cdot \frac{1}{\|N(\alpha)\|} \sum_{j \in N(\alpha)} \mathbf{m}_{f,j \rightarrow \alpha}^{(t)} \quad \text{aggregated message} \quad (8)$$

$$\mathbf{h}_{f,\alpha}^{(t)} = \text{GRU}_f(\mathbf{h}_{f,\alpha}^{(t-1)}, \mathbf{a}_{f,\alpha}^{(t)}) \quad \text{state update} \quad (9)$$

The updates for variable nodes are similar except that Equation (8) is replaced by $\mathbf{a}_{v,j}^{(t)} = \frac{1}{\|N(j)\|} \sum_{\alpha \in N(j)} \mathbf{m}_{v,\alpha \rightarrow j}^{(t)}$ since there is no feature associated with variables.

We also explored the use of more complicated message functions like a two-layer MLP instead of linear transformation on $\mathbf{h}_{v,j}^{(t)}$ to replace Eq. (7) and flexible message-passing mechanisms like Graph Attention Networks (Veličković et al. 2017; Brody et al. 2021). However, they do not show significant performance gain on any of the datasets described in Sect. 3.3, so we retain the simple linear message function.

The latent states of factor and variable nodes $\mathbf{h}_f^{(t)}$, $\mathbf{h}_v^{(t)}$ are updated iteratively until step T following Eqs. (7) to (9). As discussed earlier, we would like our model to learn an iterative algorithm that converges to the target, however many steps it takes. So instead of choosing the number of time steps T to be fixed, which we found often yielded RNN dynamics that passed through the target output without stopping, we randomly sample a readout time T from a range, so the network cannot rely on any particular readout time. At that point, we use a decoder MLP to read out simple statistics of the univariate marginal distributions for each variable as our target.

$$\eta_j = \text{Decoder}(\mathbf{h}_{v,j}^{(N)}). \quad (10)$$

For example, in Gaussian Graphical Models, we use the decoder to predict the inverse variance of the marginal distributions for each variable in the graph simultaneously. In continuous third-order graphical models, we use the first four central moments, which are expectation parameters, as our target. We target the expectation parameters or other simple statistics because they are a convenient way of capturing information about the local marginals, and they are easy to extract from samples from the training distributions.

3.3 Datasets

For tractability, many probabilistic models are based on pairwise interactions. Other PGMs like Bayesian Networks can capture higher-order interactions directly through dependence on multiple parents, but such models are often decomposed into additive pairwise interactions. Pairwise models would require many nonlinear auxiliary hidden variables to capture real-world data complexity such as multiplicative lighting, perspective transforms, triple synapses, or especially gating. In contrast, higher-order interactions capture some of these interactions directly. For example, third-order multiplicative interactions can be effectively viewed as soft gating operations, a crucial,

Table 1 Properties of datasets

Dataset	Variable	Complexity of marginals	BP inference	MP form	Max order
GGM-tree	Continuous	Linear	Exact	Closed-form	2
GGM-all	Continuous	Cubic or better	Approx	Closed-form	2
3-Binary	Binary	Exp	Approx	Closed-form	3
3-Continuous	Continuous	Exp. and approx.	Approx	Approx	3
LDPC	Binary	Not needed	Approx	Closed-form	6

The first column describes the type of variables (discrete or continuous). The second column indicates the computational complexity needed to obtain marginals for training. The third column describes if BP performs exact or approximate inference, and the forth column indicates whether BP's message-passing (MP) update has a closed-form formula. The last column denotes the highest factor order of PGMs in the dataset

common operation in machine learning, as seen in LSTMs (Hochreiter and Schmidhuber 1997), GRUs (Cho et al. 2014), and transformer networks (Vaswani et al. 2017). These third-order models provide more modeling power while remaining more interpretable statistically. Thus here we choose continuous graphical models with third-order interactions as a distribution family of high interest. Such distributions fall into the category of applications where traditional inference algorithms are infeasible but our GNNs still apply.

We construct four PGM datasets (first four lines in Table 1) with increasing complexity and report the performance of the RF-GNN on these datasets compared to BP. Each dataset consists of a training set and several test sets. The training set includes graphical models with a fixed number of variables, but with diverse structures and factor parameters. Each test set is constructed in the same way as the training set, but with a different graph size.

First, we construct three datasets of PGMs with known ground truth marginal distributions and closed-form BP update formulas: Gaussian Graphical Model (GGM) with tree structure, Gaussian Graphical Model with arbitrary structure, and binary Third-order Graphical Models. These datasets allow us to compare the performance of our model with Belief Propagation. Notably, since BP produces exact marginals on tree graphs, performance on the GGM-tree dataset serves a proof of concept that the RF-GNN is highly accurate on tree graphs, even when tested on larger graphs not seen during training. Second, we build a dataset of continuous PGMs with third-order interactions to test our model on a more complicated and highly interesting class of third-order PGMs. Finally, as a benchmark to compare the RF-GNN to other GNN-based models for PGMs, we test the RF-GNN on a Low-Density Parity-Check (LDPC) decoding dataset consisting of binary PGMs with a fixed structure but different parameters. Each of these datasets is described in greater detail below.

3.3.1 Testing generalization on new graphs

Flexible recurrent message-passing algorithms have the property of that we can apply them to graphs of sizes not seen during training. Here, we define how we measure an RF-GNN's ability to generalize to different sized graphs. Every PGM dataset consists

of a training set and several test sets. The training data $\mathcal{D}_{\text{train}}^N = \{(\mathcal{G}_i, \eta_i) \mid \mathcal{G}_i \sim \mathcal{P}^G(N), \eta_i \sim \mathcal{P}^\eta(N)\}$ is composed of PGMs with the same size N , but with a different graph-generating process $\mathcal{P}^G(N)$ and parameter-generating process $\mathcal{P}^\eta(N)$. We use the same processes $\mathcal{P}^G(M)$ and $\mathcal{P}^\eta(M)$ to generate test set $\mathcal{D}_{\text{test}}^M$ of size M . We define the generalization capability of the RF-GNN to a different, especially larger, size M as the performance of an RF-GNN trained on $\mathcal{D}_{\text{train}}^N$ and evaluated on $\mathcal{D}_{\text{test}}^M$.

3.3.2 Gaussian graphical model

As toy examples, we generate two datasets of random Gaussian Graphical Models (GGM): *GGM-tree*, and *GGM-all*. The first dataset consist only tree graphs as a proof of concept. The second dataset includes graphs with diverse graph structures. Since both exact marginalization and BP have closed-form solutions on GGMs, they serve as a convenient test for comparing an RF-GNN with BP, and allow us to test the generalization performance of an RF-GNN to much larger out-of-distribution graphs. For each PGM in these two datasets, we generate the graph structure and eigenvalues of the GGM independently and then find a covariance matrix that possesses these eigenvalues and complies with the graph structure, as described below.

For *GGM-tree*, we construct a random tree structure by converting from a random Prüfer sequence (Prüfer 1918), which is sampled uniformly from the symmetry group S_{n-2} where n is the number of nodes in the tree.

For *GGM-all*, we use the random graph generator algorithm WS-flex proposed in You et al. (2020) to generate diverse graph structures parametrized by their average shortest path length and cluster coefficient. To generate a random graph, we first sample the average degree k and rewiring probability p parameters for the WS-flex algorithm from a uniform distribution over $[2, n-1]$ and $[0, 1]$ where n is the number of variables, and then construct a random graph using these parameters.

We now describe how we generate a precision matrix that complies with any connected graph structure, whether it is a tree or loopy graph. Starting from a positive-definite diagonal matrix whose eigenvalues are uniformly drawn from $[0.1, 10.0]$, we first apply a random orthogonal rotation to make it dense. Then we iteratively apply Jacobi rotations to zero out elements according to the connectivity matrix until we achieve the desired structure. We focus on inferring the node variances and give all GGMs a mean of zero because BP is known to give exact means in general GGMs, but only makes errors in marginal variances (Weiss and Freeman 1999).

For a single GGM, the inverse variance of each variable's marginal Gaussian distribution is used as target for supervised learning. 10,000 random graphs with 10 variables are generated as a training dataset. Additionally, we generate various testing datasets with 10 to 50 variables, each consisting of 2000 random graphs.

3.3.3 Binary third-order factor graph

The RF-GNN works with general factor graphs, especially those with higher-order interactions, but GGMs have only pairwise interactions and can also be modeled by regular GNNs (Li et al. 2022). To test our model on graphs with higher-order factors, we

construct a dataset composed of binary graphical models with only third-order interactions. For the sake of exact solutions of marginal probabilities, we choose small binary spin glass models whose structures and connectivity strengths are randomly sampled so that exact marginal probabilities can be calculated by brute-force enumeration. The joint probability mass function is expressed as

$$p(\{\mathbf{s}\}) \propto \exp \left(\beta \left(\sum_{p=1}^3 \sum_{i \in \mathbb{I}_V} b_{i,p} s_i^p + \sum_{(ijk) \in \mathcal{N}(\mathcal{F})} J_{ijk} s_i s_j s_k \right) \right). \quad (11)$$

where $\mathcal{N}(\mathcal{F})$ is the set of neighbor variable indices for each factor. We generate these random factor graphs using a generalization of the WS-flex generator, parameterized by average variable node degree k_3 and rewiring probability p . Details of the random graph-generating procedure can be found in “Appendix A”. The singleton potential coefficients b_i^p and 3-way coupling coefficients J_{ijk} are randomly drawn from a standard Gaussian distribution. The inverse temperature, β , is set to 0.5.

3.3.4 Continuous third-order factor graphs

We are interested in continuous graphical models with third-order interactions, and we propose to use our method as an approximate inference algorithm for this model class. We construct a dataset of random continuous PGMs with third-order interactions and evaluate the accuracy of our message-passing algorithm against an expensive sampling approach. For each inference problem, a factor graph is randomly constructed such that 3-factors are generated using the aforementioned WS-flex variant, and 2-factors are generated using the regular WS-flex algorithm. Each graph structure is parametrized by three numbers: the average outdegree of variable nodes to 2-factor nodes k_2 , the average outdegree of variable nodes to 3-factor nodes k_3 , and the rewiring probability p . k_2 is uniformly sampled from $[2, n - 1]$, k_3 is uniformly sampled from $[2, \lfloor (n-1)(n-2)/6 \rfloor]$, and p is uniformly sampled from $[0, 1]$. An isotropic 4th-order base measure is added to ensure the joint density is normalizable:

$$p(\mathbf{x}) \propto \exp \left[-\beta \left(\sum_{p=1}^3 \sum_{i \in \mathbb{I}_V} b_{i,p} x_i^p + \sum_{(ij) \in \mathcal{N}(\mathcal{F}_2)} K_{ij} x_i x_j + \sum_{(ijk) \in \mathcal{N}(\mathcal{F}_3)} J_{ijk} x_i x_j x_k + \|\mathbf{x}\|_{\ell_4}^4 \right) \right] \quad (12)$$

The bias parameters $\{b_{i,p}\}$ and interaction strengths $\{K_{ij}\}$, $\{J_{ijk}\}$ are sampled from a standard normal distribution $\mathcal{N}(0, 1)$. The inverse temperature β is chosen to be 0.3. Note that this exponential family of distributions is not closed under marginalization: marginals are not in the same family as the joint.

As an approximate ground truth for our training algorithms, we run an MCMC algorithm, the No-U-Turn Sampler (NUTS) (Hoffman and Gelman 2014), using the Stan (Stan Development Team 2021) software for a large number of steps. Our read-out targets are summary statistics computed from those generated samples. For each random structure generated, we run 8 MCMC chains with 10,000 warmup steps and




Fig. 3 Example univariate marginal distributions from the continuous third-order dataset. Shaded blue curves represents empirical marginal distributions from samples. Shaded orange curves are empirical single-variate distributions by only considering singleton potentials. Green dots represent the maximum entropy fit of the sample marginal distributions by matching the first four central moments (color figure online)

10,000 sampling steps each for a random set of parameter values, and keep drawing new parameters until the potential scale reduction factor (PSRF) (Gelman et al. 1992) falls below 1.2, indicating convergence has been reached for the MCMC chains. We target the first four central moments, since the Jensen-Shannon divergence between the empirical sample distribution and corresponding moment-matched maximum-entropy distribution does not decrease substantially when including more moments. Figure 3 shows random examples of sampled univariate marginals, the corresponding singleton distributions obtained by ignoring all multivariate interactions, and their maximum entropy counterparts by matching the first four central moments. Observe that there are substantial differences between these marginals and their singleton potentials. These differences are caused by the influences of other nodes on the graph, and we would like our RF-GNNs to capture these network effects.

3.4 Training

We implement the RF-GNN¹ in PyTorch (Paszke et al. 2019) and PyTorch Geometric (Fey and Lenssen 2019) and perform all experiments on internal clusters equipped with NVIDIA GTX 1080Ti and Titan RTX GPUs. We randomly split each dataset into a training and a validation set of ratio 4:1. The testing sets are constructed separately for different graph sizes. In every experiment, we use the ADAM (Kingma and Ba 2014) optimizer with batch size 64 and initial learning rate 0.001. We multiply the learning rate by a factor of 0.2 if there is no improvement in 20 epochs, and perform early-stopping if there is no improvement in 40 epochs. We choose the dimension of hidden states for both variable and factor nodes $\mathbf{h}_v, \mathbf{h}_f$ to be 64. All MLP modules use SELU activations (Klambauer et al. 2017) and 2 hidden layers with 64 units per layer. We use mean-squared error (MSE) as loss on single-variate marginal precision values for the Gaussian dataset and on the first four central moments calculated from MCMC samples for the continuous third-order factor graph dataset respectively. For the binary third-order dataset, we use binary cross entropy as loss.

For training an FGNN (Zhang et al. 2020) on our Gaussian datasets, we use the authors' implementation for their LDPC dataset, but without the BatchNorm layers since in our hands those layers led to inferior performance on our datasets. We re-wrote the message-passing part of their code to support arbitrary graph structures since the

¹ Code is available at <https://github.com/fy0cc/factor-inference-net/>.

Table 2 Experimental results (mean squared error, MSE) on four synthetic PGM datasets

Method	GGM-tree [$\times 10^{-3}$]	GGM-all [$\times 10^{-3}$]	3-Binary [$\times 10^{-3}$]	3-Continuous [$\times 10^{-1}$]
BP	0	141.6	9.4	N/A
RF-GNN	4.4(13)	4.6(1)	5.10(1)	1.329(6)
Singleton	N/A	902.7(6)	21.5	1.930

Bold represents best performance in column

MSE is used as the training objective for datasets *GGM-tree*, *GGM-all*, and *3-continuous*, while binary cross-entropy loss is used for *3-binary*. Numbers in the parenthesis indicates the uncertainty of the last decimal

Table 3 Experimental results (coefficient of determination R^2) on four synthetic PGM datasets

Method	GGM-tree	GGM-all	3-binary	3-continuous
BP	1	0.9719	0.8290	N/A
RF-GNN	0.9995(1)	0.9991	0.9378(7)	0.8158(8)
Singleton	N/A	0.8329	0.7352	0.6543(1)

Bold represents best performance in column

Numbers in the bracket indicates the uncertainty of the last decimal

released implementation only supports graphs with fixed-degree variables. We used the ADAM (Kingma and Ba 2014) optimizer with initial learning rate 0.003, batch size 100, and the same learning rate scheduling as for the RF-GNN.

4 Experiments and results

All RF-GNNs and FGNNs are trained on graphs of size 10 only, but are tested on graphs of various sizes. We uniformly draw the number of recurrent steps T from [30, 50] during training. T is set to 30 during testing. We repeat all experiments 10 times from different random seeds, pick the best model with the lowest validation error, and report the bootstrapped mean and 95% confidence interval on held-out test sets with different graph sizes. As a baseline to show how strong the interactions are in each dataset, we also train a separate RF-GNN (singleton RF-GNN) that tries to predict the same target, but only sees a modified graphical model with all interaction terms removed. A singleton RF-GNN would have lower performance on graphs whose marginal distributions are strongly affected by multivariate interactions. We show in-distribution results of BP, the full RF-GNN, and the singleton RF-GNN in Tables 2 and 3. To illustrate the absolute and relative errors, we report Mean Squared Error (MSE) and Coefficient of determination (R^2) separately in Tables 2 and 3.

We also show in Fig. 4 the performance of RF-GNNs when generalizing to unseen graph structures with different sizes from the same parametric family. It is well known that Belief Propagation becomes non-exact on loopy graphs. Here, we empirically investigate how BP and an RF-GNN performance depends on the graph structure by examining graph metrics. We choose two independent structural features that quantify

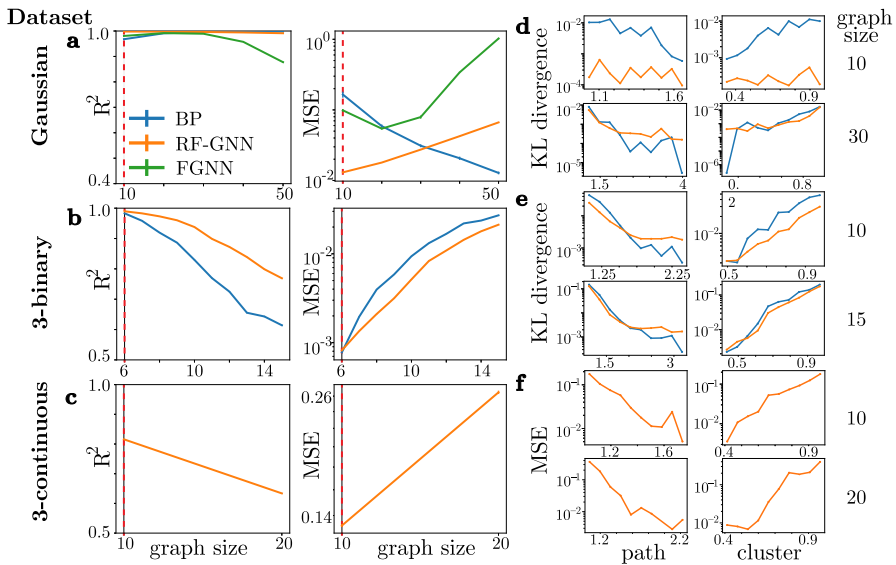


Fig. 4 Performance on Gaussian, third-order spin, and third-order continuous datasets. Rows a+d, b+e, and c+f show the results on Gaussian, binary spin, and third-order continuous datasets respectively. In subfigures a, b, and c, we compare the generalization performance of the RF-GNN, BP, and FGNN when applied to larger graphs. All RF-GNNs and FGNNs are trained on small graphs and tested on larger ones. The training graph sizes are indicated with red dashed vertical lines. Two metrics, R^2 and MSE, are shown for the Gaussian and spin datasets. Subfigures d, e, and f show how the error of different methods depends on graph structures. We plot error against the average shortest path length in the left panels, and against the cluster coefficient in the right panels. For each dataset, we show the in-sample results in the top row and generalization results on larger graphs in the bottom row (color figure online)

the loopiness of a graph: average shortest path length and cluster coefficient (Watts and Strogatz 1998). The former quantifies the average loop length between any two nodes, and the latter describes the small-world property of a graph. Graphs in each test set are binned into 10 equal bins according the graph metric, and for each bin we report the bootstrapped mean and its 95% CI for BP and RF-GNN performance metrics in Fig. 4d-f. Many real-world graphs have a small average path length and a large cluster coefficient (Watts and Strogatz 1998), and this is the region where our model outperforms BP.

4.1 Gaussian graphical models

4.1.1 Tree graphs

We trained a model on a dataset of 10^4 Gaussian tree graphs with size 10 using the eigenvalue distribution in Sect. 3.3.2. BP is exact on trees, but our model also achieves excellent generalization performance even when generalizing to graphs with size 50, giving an in-distribution R^2 score 0.9995 ± 0.0001 and out-of-sample R^2 score 0.9996 ± 0.0002 .

4.1.2 Loopy graphs

In this experiment we tested the performance of our model on Gaussian Graphical Models with zero mean, random precision matrices, and various structures as described in Sect. 3.3.2. Figure 4a, right, shows that our model trained on a dataset of GGMs with 10 variables has an average MSE that is 30-fold smaller than Belief Propagation on a test set with the same size. When generalizing to larger graphs, our model still has a smaller error than BP on graphs up to 30 variables. We also report the corresponding R^2 score as a metric for goodness of prediction in Fig. 4a, left. Since BP doesn't always converge on loopy graphs, all metrics are calculated conservatively using the subset of test graphs with convergent BP dynamics, which comprise from 73.2 to 71.2% of the test set as the graph size varies from 10 to 50. We consider BP to be non-convergent if the absolute error of beliefs between two adjacent BP updates still exceeds 10^{-5} after 1000 cycles.

BP can be viewed as a recursive algorithm that finds local fixed points of the Bethe free energy upon convergence (Heskes et al. 2003). That was indeed our original motivation to construct our recurrent method as an autonomous dynamical system. However, for comparison, we also compare with FGNN (Zhang et al. 2020), a feed-forward model with 10 layers. Figure 4a shows that although the feedforward model has more parameters, the performance is worse than the recurrent on test datasets with graph sizes varying from 10 to 50. A tradeoff is that, due to its recurrent nature, RF-GNN runs slower than FGNN.

4.2 Binary third-order graphical model

For the binary graphs with only third-order interactions and no pairwise interactions, we train an RF-GNN on a dataset of size 10 and test on graphs with sizes from 6 to 15. We didn't test on larger graphs because it becomes impractical to enumerate all spin configurations in order to compute exact marginal probabilities. For binary graphical models, unlike in Sect. 4.1, non-convergent BP dynamics won't diverge, but may oscillate. Thus, we take the beliefs from the last cycle if the dynamics does not converge in 1000 cycles. There is no qualitative difference when evaluating model performance using the whole test dataset or just BP-convergent ones, so we report the metrics on BP-convergent graphs to maintain consistency with Sect. 4.1. Within the range of testing graphs, RF-GNNs consistently outperform BP on average (Fig. 4b). In the space of graph structures, RF-GNNs perform better than BP in regions with smaller average shortest path length and larger cluster coefficient, which agrees with the result of Sect. 4.1 (Fig. 4e).

4.3 Continuous third-order graphical model

For general, continuous, non-Gaussian graphical models it is not feasible to compare the RF-GNN with BP because the BP message update is neither an explicit formula, nor calculable exactly by enumeration. Even EP message updates need to be approximated (Heess et al. 2013; Eslami et al. 2014; Jitkrittum et al. 2015). These

methods approximate local BP updates without knowing the global graph structure, thus inheriting the drawbacks of BP on loopy graphs. Instead of approximating the EP message-passing for a specific type of factors within the EP framework, we learn a new message-passing algorithm end-to-end that works with many factor types and avoids some of EP's limitations. This approach is more applicable to loopy graphs where BP and EP struggle.

For this experiment, we constructed a continuous graphical model with third-order interactions and train our model to predict the first four central moments of every univariate marginal distribution (see Sect. 3.3.4). Since BP or EP is not feasible for this dataset, we only compare the RF-GNN with the singleton RF-GNN as our baseline. An RF-GNN trained on graphs with 10 variables achieves an in-sample R^2 score of 0.816 ± 0.001 , while the baseline model without interactions has an R^2 of only 0.654 ± 0.001 . We test the generalization performance of our model with a dataset of graphs of size 20. For the MCMC algorithm to produce convergent chains for these larger graphs, we shrink the interaction strengths by a factor of 2 while keeping the singleton parameter distribution unchanged so that both input and output values fall in the training ranges. On a test set of 1000 graphs, the full model has an R^2 score of 0.654 ± 0.000 , while the baseline model has an R^2 score of 0.495 ± 0.000 . Qualitatively similar to the Gaussian and spin datasets, model performance drops when generalizing to larger graphs (Fig. 4c), and our model makes larger errors on loopy graphs (Fig. 4f) since these are harder problems.

4.4 Low-density parity-check (LDPC) codes

As one simple real-world application of our method, we tested the RF-GNN on a Low-Density Parity-Check (LDPC) decoding dataset used for the Factor Graph Neural Network (FGNN) (Zhang et al. 2020) and compare the performance of an RF-GNN and FGNN along with other baseline models used in Zhang et al. (2020). In the LDPC dataset, a noisy signal $\tilde{\mathbf{y}}$ is obtained by transmitting the original signal \mathbf{y} through a noisy channel with Gaussian noise and irregular bursting noise:

$$\tilde{y}_i = y_i + n_i + p_i z_i \quad (13)$$

$$n_i \sim \mathcal{N}(0, \sigma^2) \quad (14)$$

$$z_i \sim \mathcal{N}(0, \sigma_b^2) \quad (15)$$

$$p_i \sim \text{Bernoulli}(0.05) \quad (16)$$

The task is to decode the first 48-bit or original signal from a 96-bit noisy signal $\tilde{\mathbf{y}}$. The graphical model is fixed in this task given a specific coding scheme, with 96 variables each connected to three factors, and 48 parity-check factors each connected to six variables.

We trained an RF-GNN using the hyperparameters stated in Sect. 3.4, except in this specific task we use a fixed number of 30 iterations instead of a random number for better performance, since generalization to different graph structures is not needed in this task. The performance curves of the RF-GNN and compared models for different




Fig. 5 Performance curves of several models (MacKay and Codes 2009; Taranalli 2015; Zhang et al. 2020) on the LDPC Decoding Dataset, showing the bit error rate as a function of the independent Gaussian noise σ for different values of burst noise σ_b . SNR in decibels is calculated as $-10 \log_{10}(\sigma^2)$. All shown methods outperform the naïve baseline of decoding bits independently (dashed line), but our method exhibits the best improvement when the burst noise is high

levels of burst noise are presented in Fig. 5. Testing results for competing methods are directly taken from Zhang et al. (2020)’s public figure generation scripts on GitHub, and we did not attempt to further optimize results for competing models because they already had incentives to engineer their system for strong performance. The testing dataset is also generated with Zhang et al. (2020)’s code with fixed random seeds. Due to the recurrent nature of our method, it currently takes more clock time to process one batch compared to the feedforward FGNN. However, our model converges with $100\times$ fewer iterations, after seeing only 10^6 samples, while the FGNN converges after seeing 10^8 samples. While being more parameter-efficient, our model has similar error rates as FGNN under low SNR and smaller error rates at high SNR.

5 Related work

Previous work has explored learning to pass messages that calculate marginal probabilities in graphical models. Some examined fast and approximate calculation of Belief Propagation or Expectation Propagation updates when analytical integration is not feasible for computing on the level of single factor (Heess et al. 2013; Eslami et al. 2014; Jitkrittum et al. 2015). Others created message-passing algorithms for inference in probabilistic graphical models, trying to learn an algorithm that is more accurate than Belief Propagation when the underlying graph is loopy (Yoon et al. 2019). The most related work to ours uses stacked bidirectional GNNs on factor graphs (FGNNs) to perform maximum *a posteriori* estimation on binary graphical models (Zhang et al. 2020). One major difference is that we use a recurrent network to construct a dynamical system that runs to convergence, like Belief Propagation, instead of a feedforward network with fixed number of layers, so our method uses far fewer parameters and scales

to larger graphs without adding new layers. Another related study puts a factor graph NN layer within a recurrent algorithm to calculate marginal probabilities in binary graphical models, applying them to low density parity check codes and Ising models (Satorras and Welling 2021). Instead of using a GNN to learn a novel message-passing algorithm from scratch, the authors build their algorithm on top of Belief Propagation and let the GNN serve as an error-correction module for loopy BP.

Some work has been done to extend GNNs to such higher-order graphs (Bai et al. 2019; Zhang et al. 2021), including studies aiming to solve graph isomorphism tasks (Morris et al. 2019), and others applying GNNs to other high-order data structures like simplicial complexes (Ebli et al. 2020a).

6 Discussion

6.1 Limitations

A limitation of our method is that a separate RF-GNN needs to be learned from scratch whenever encountering a new parametric family of graphical models. One exciting but speculative possibility is to extend this framework such that whenever a new type of factor is encountered, the model only trains a dedicated new encoder/decoder, and reuses a shared core message-passing module, perhaps implementing a more universal inference engine akin to a shared language model (Denkowski and Lavie 2014). This would be useful when data is scarce, since the difficult message-passing core could be learned from multiple rich datasets. The general formulation of Belief Propagation takes such a universal form, as have some multi-factor extensions like Generalized BP (Yedidia et al. 2000). We hypothesize that the RF-GNN could potentially extend this universality while compensating for some challenges of inference in a loopy world.

As discussed and explored by Chen et al. (2019), Sato (2020), Maron et al. (2019), Keriven and Peyré (2019) and others, message-passing algorithms with only local information like BP and regular GNNs, including the RF-GNN, would fail to solve some graph isomorphism tasks on loopy graphs. These algorithms will produce the same marginals for two locally similar but globally different PGMs, even though the true marginals should be different. Some GNN variants (Murphy et al. 2019; Morris et al. 2019; Chen et al. 2019) allow nodes to see beyond their direct neighbors in order to partially mitigate this issue. We hypothesize these could be employed as a replacement of local GNNs to get better inference performance on loopy graphs, at the costs of graph locality and greater complexity.

6.2 Extensions

In this paper, we show that our recurrent factor graph neural network is a valuable alternative to traditional message-passing algorithms like Belief Propagation as an approximate inference engine on graphical models. It provides greater flexibility that accommodates inference on a far richer class of graphical models with multi-way interactions. It outperforms BP on some loopy graphs, and offers a practical way to perform fast inference for continuous non-Gaussian graphical models.

Although our method achieves promising results and does generalize to graphs of different sizes, it naturally performs better within its training distribution. However, we saw almost no performance drop on Gaussian tree graphs when generalizing to $5\times$ larger graphs (Sect. 4.1.1), and we saw promising generalization results on loopy graphs (Sect. 4.1.2, Sect. 4.2). This suggests that our methods could be used for inference on large graphs by training on smaller ones. Although our method needs training data, general MCMC sampling methods like NUTS (Hoffman and Gelman 2014) could always be used to generate training data on small graphs when efficient sampling methods for larger ones are not available. Compared to a 10-layer FGNN used in Zhang et al. (2020), our method achieves better performance with fewer parameters (Fig. 4a). Our recursive message-passing algorithm has a substantial advantage over a feedforward one when generalizing to larger graphs: the number of layers must scale with graph size for a graph-structured feedforward network in order to distribute information throughout the graph. In contrast, a recurrent algorithm can run unchanged until convergence regardless of the graph size, and this will allow information to propagate to all nodes on the graph while still respecting the graph structure.

Besides performing marginalization, RF-GNNs could also be trained to calculate the most probable state vector by altering the training objective to treat simultaneous local readout as a global target. Conditioning should also be straightforward since by conditioning on some subset of variables, we get a smaller PGM within the same family. Thus, the same RF-GNN could be directly applied to perform inference on the conditional model.

While belief propagation is exact on trees, it is only an approximation when used on loopy graphs because the update rule incorrectly assumes that messages are independent, as they are for tree graphs. The topological structure of the graph therefore has an important impact on the performance of these algorithms. Past work has attempted to mitigate this effect, creating new message-passing algorithms by training GNNs to compensate (Yoon et al. 2019). Note that it is not only the existence of loops that creates problems, but also the length of the loops and the strength of interactions along those loops. Since topology is insensitive to the size of loops, it would be blind to these effects, but topological data analyses often look for persistent homology, which is sensitive to the sizes of loops. Often, it is shorter loops that create the most problems in belief propagation. Thus, by incorporating higher-order structure, one may absorb smaller loops into cliques, potentially improving performance while preserving the large-scale topology.

Higher-order cliques are an important ingredient in other computations as well. One notable example is the simplicial complex, a foundational object in algebraic topology. Simplicial complexes are collections of fully-connected cliques that recursively contain all fully-connected subcliques among all subsets of nodes. The structure of simplicial complexes can reveal large-scale topological features of a graph, and GNNs can be applied to simplicial complexes (Bodnar et al. 2021a, b; Ebli et al. 2020b). However, using simplicial complexes as a primitive requires us to pass more messages to account for dense constraints across all neighboring orders: 4-cliques must interact with all contained 3-cliques, which interact with all contained 2-cliques, and so on. In contrast, there may be major computational advantages to using sparser messages between subsets of cliques. This could be because there are indeed sparse higher-order

interactions in the graphical model (as the PGMs in our study) without the inclusion structure implied by simplicial complexes. Alternatively, sparse higher-order messages can be helpful even when the true model has low-order interactions because it can reduce the number of messages as seen in Kikuchi clustering (Yedidia et al. 2000), Junction Tree (Lauritzen and Spiegelhalter 1988), and higher-order Weisfeiler-Leman algorithms (Morris et al. 2019).

In this paper, we focus on performing inference on known graphical models. However, in many real-world applications, it is also hard to estimate the model parameters from data. One way an approximate message-passing inference algorithm like ours could facilitate parameter estimation is through self-supervised training like in Lázaro-Gredilla et al. (2021). Given some data that is assumed to be generated by a graphical model we are fitting, we could construct a new graphical model by randomly masking a portion of observed data as hidden and use GNN on the currently fitted graphical model to infer the masked-out data. In this way, both the model parameters and the GNN could be trained to make high quality partial inferences. However, the model parameters could be biased towards partial inference instead of Maximum Likelihood Estimation, as shown in Lázaro-Gredilla et al. (2021).

Acknowledgements The authors thank KiJung Yoon for helpful conversations about graph metrics and Alex Kunin for conversations about topological spaces. This work was supported in part by NSF NeuroNex Grant 1707400 to XP, and is also based in part upon work supported by the Air Force Office of Scientific Research (AFOSR) under award number FA9550-21RT0750 to XP.

Funding Open Access funding provided by Carnegie Mellon University.

Declarations

Conflict of interest XP is a co-founder of Upload AI, LLC, a company which may have a competing interest in developing related algorithms. YF has no competing interests to declare.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Random factor graph generator for 3-factors


We use a variant of the WS-flex algorithm (You et al. 2020) to generate a random factor graph with only 3-factors. This includes two stages: generating a ring-like factor graph, and rewiring. In the first stage, given average variable node degree k and number of variable nodes n , the number of factors is $f = \lfloor nk/3 \rfloor$. On average, there are $\lfloor f/n \rfloor$ 3-factors centered at each variable node. The other two legs of these factors are connected to the closest neighbors of the center variable, while keeping the factor unique. For example, three 3-factors centered at variable 5 may be connected

Table 4 Experimental results (KL-divergence between true marginal and the predicted distribution) on three applicable synthetic PGM datasets

Method	GGM-tree [$\times 10^{-3}$]	GGM-all [$\times 10^{-3}$]	3-Binary [$\times 10^{-3}$]	3-Continuous [$\times 10^{-4}$]
BP	0	6.5	26.9	N/A
RF-GNN	1.3(4)	0.27(1)	15.9(2)	N/A
Singleton	N/A	46.0(1)	52.8	N/A

Bold represents best performance in column

We omit the KL-divergence calculation for dataset 3-continuous as we don't have exact true marginal distributions

Fig. 6 Generalization performance on two PGM datasets for which exact KL-divergence between ground truth and predicted distributions could be calculated

to variables (4, 5, 6), (4, 5, 7) and (3, 5, 6) respectively. The remaining $(f \bmod n)$ 3-factors are then formed by randomly picking the center variable and choose the closest two neighbors, while avoiding redundant factors. After this ring is created, in a second stage we may rewire each factor with probability p by randomly picking three new neighbors without duplication.

Appendix B: Additional results

As a supplement to the main results presented in Fig. 4, Tables 2 and 3, here we show some additional data and figures which support the results shown in the main text.

In addition to MSE and R^2 score, we also compute the KL divergence between the true marginal and the predicted marginal for applicable datasets. Results are shown in Table 4 and Fig. 6, which are in the same formats as Table 2 and Fig. 4.

In the main text, we show how GNN and BP perform with respect to two graph metrics: average shortest path length and cluster coefficient. Here in Figs. 7 and 8 we

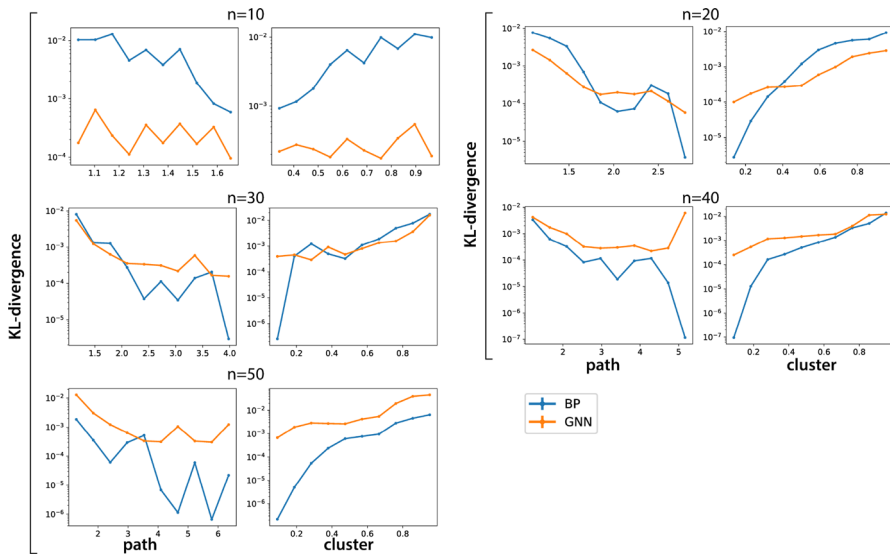


Fig. 7 Gaussian Dataset: dependence of KL divergence between true marginal distribution and model prediction on average shortest path length and cluster coefficient. We compare Belief Propagation (blue) and an RF-GNN (orange). The RF-GNN is trained on graphical models with 10 variables. Each pair of panels shows the testing performance on a separate dataset with graphs of fixed size ranging from 10 to 50 (color figure online)

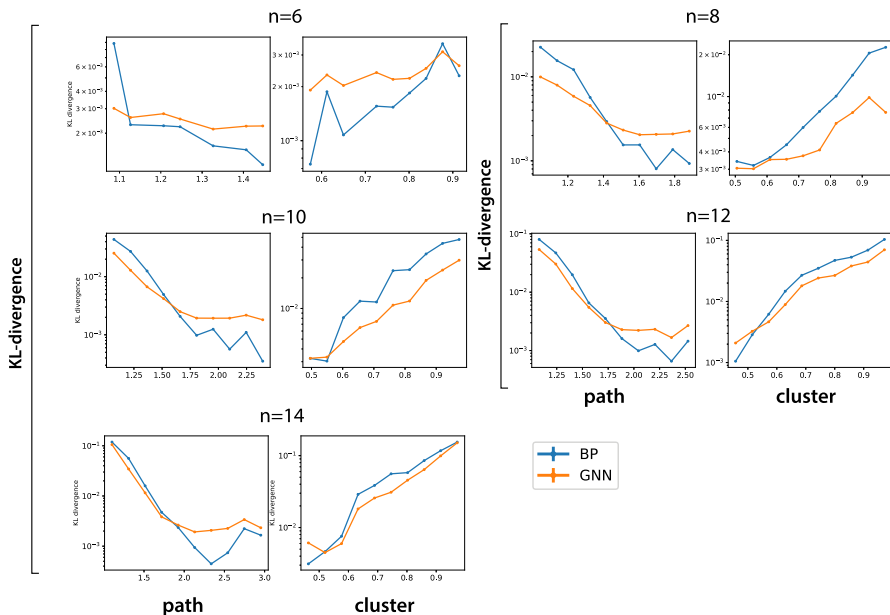


Fig. 8 Binary Dataset: dependence of KL divergence between true marginal distribution and model prediction on average shortest path length and cluster coefficient, plotted as in Fig. 7. Each pair of panels shows the testing performance on new graphical models with a fixed number of variables ranging from 6 to 14

provide two figures on testing sets with more sizes for the Gaussian dataset and the binary dataset.

References

- Ba, J.L., Kiros, J.R., Hinton, G.E.: Layer Normalization. arXiv [arXiv:1607.06450](https://arxiv.org/abs/1607.06450) (2016)
- Bai, S., Zhang, F., Torr, P.H.S.: Hypergraph Convolution and Hypergraph Attention. arXiv [arXiv:1901.08150](https://arxiv.org/abs/1901.08150) (2019)
- Beggs, J.M., Plenz, D.: Neuronal avalanches in neocortical circuits. *J. Neurosci.* **23**(35), 11,167–11,177 (2003)
- Bishop, C.M.: Pattern Recognition and Machine Learning. Springer, New York (2006)
- Bodnar, C., Frasca, F., Otter, N., et al.: Weisfeiler and Lehman go cellular: CW networks. *Adv. Neural Inf. Process. Syst.* **34**, 2625–2640 (2021a)
- Bodnar, C., Frasca, F., Wang, Y., et al.: Weisfeiler and Lehman go topological: message passing simplicial networks. In: International Conference on Machine Learning, PMLR, pp. 1026–1037 (2021b)
- Brody, S., Alon, U., Yahav, E.: How attentive are graph attention networks? arXiv preprint [arXiv:2105.14491](https://arxiv.org/abs/2105.14491) (2021)
- Centola, D., Becker, J., Brackbill, D., et al.: Experimental evidence for tipping points in social convention. *Science* **360**(6393), 1116–1119 (2018). <https://doi.org/10.1126/science.aas8827>
- Chen, Z., Villar, S., Chen, L., et al.: On the equivalence between graph isomorphism testing and function approximation with GNNs. *Adv. Neural Inf. Process. Syst.* **32** (2019)
- Chen, R.T., Rubanova, Y., Bettencourt, J., et al.: Neural ordinary differential equations. *Adv. Neural Inf. Process. Syst.* **31**, 9487 (2018)
- Cho, K., Van Merriënboer, B., Gulcehre, C., et al.: Learning phrase representations using rnn encoder-decoder for statistical machine translation. arXiv preprint [arXiv:1406.1078](https://arxiv.org/abs/1406.1078) (2014)
- Denkowski, M., Lavie, A.: Meteor universal: language specific translation evaluation for any target language. In: Proceedings of the Ninth Workshop on Statistical Machine Translation, pp. 376–380 (2014)
- Ebli, S., Defferrard, M., Spreemann, G.: Simplicial Neural Networks. arXiv [arXiv:2010.03633](https://arxiv.org/abs/2010.03633) (2020a)
- Ebli, S., Defferrard, M., Spreemann, G.: Simplicial neural networks. arXiv preprint [arXiv:2010.03633](https://arxiv.org/abs/2010.03633) (2020b)
- Eslami, S., Tarlow, D., Kohli, P., et al.: Just-in-time learning for fast and flexible inference. *Adv. Neural Inf. Process. Syst.* **27** (2014)
- Fey, M., Lenssen, J.E.: Fast graph representation learning with PyTorch Geometric. In: ICLR Workshop on Representation Learning on Graphs and Manifolds (2019)
- Frey BJ, Kschischang, F.R., Loeliger, H.A., et al.: Factor graphs and algorithms. In: Proceedings of the Annual Allerton Conference on Communication Control and Computing, Citeseer, pp. 666–680, (1997)
- Ganmor, E., Segev, R., Schneidman, E.: Sparse low-order interaction network underlies a highly correlated and learnable neural population code. *Proc. Natl. Acad. Sci.* **108**(23), 9679–9684 (2011)
- Gelman, A., Rubin, D.B., et al.: Inference from iterative simulation using multiple sequences. *Stat. Sci.* **7**(4), 457–472 (1992)
- Gilmer, J., Schoenholz, S.S., Riley, P.F., et al.: Neural message passing for quantum chemistry. In: International Conference on Machine Learning, PMLR, pp. 1263–1272 (2017)
- Gori, M., Monfardini, G., Scarselli, F.: A new model for learning in graph domains. In: 2005 IEEE International Joint Conference on Proceedings of the Neural Networks, 2005. IJCNN'05, pp. 729–734. IEEE (2005)
- Heess, N., Tarlow, D., Winn, J.: Learning to pass expectation propagation messages. *Adv. Neural Inf. Process. Syst.* **26** (2013)
- Heskes, T., et al.: Stable fixed points of loopy belief propagation are minima of the Bethe free energy. *Adv. Neural Inf. Process. Syst.* **15**, 359–366 (2003)
- Hinton, G.E.: Learning to represent visual input. *Philos. Trans. R. Soc. B Biol. Sci.* **365**(1537), 177–184 (2010)
- Hochreiter, S., Schmidhuber, J.: Long short-term memory. *Neural Comput.* **9**(8), 1735–1780 (1997)
- Hoffman, M.D., Gelman, A.: The no-u-turn sampler: adaptively setting path lengths in Hamiltonian Monte Carlo. *J. Mach. Learn. Res.* **15**(1), 1593–1623 (2014)

- Hopfield, J.J.: Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci.* **79**(8), 2554–2558 (1982)
- Iacopini, I., Petri, G., Barrat, A., et al.: Simplicial models of social contagion. *Nat. Commun.* **10**(1), 2485 (2019). <https://doi.org/10.1038/s41467-019-10431-6>. [arXiv:1810.07031](https://arxiv.org/abs/1810.07031)
- Jitkrittum, W., Gretton, A., Heess, N., et al.: Kernel-based just-in-time learning for passing expectation propagation messages. *arXiv preprint* [arXiv:1503.02551](https://arxiv.org/abs/1503.02551) (2015)
- Keriven, N., Peyré, G.: Universal invariant and equivariant graph neural networks. *Adv. Neural Inf. Process. Syst.* **32** (2019)
- Kingma, D.P., Ba, J.: Adam: a method for stochastic optimization. *arXiv preprint* [arXiv:1412.6980](https://arxiv.org/abs/1412.6980) (2014)
- Kipf, T.N., Welling, M.: Semi-supervised classification with graph convolutional networks. *arXiv preprint* [arXiv:1609.02907](https://arxiv.org/abs/1609.02907) (2016)
- Klambauer, G., Unterthiner, T., Mayr, A., et al.: Self-normalizing neural networks. *Adv. Neural Inf. Process. Syst.* **30** (2017)
- Koller, D., Friedman, N.: *Probabilistic Graphical Models: Principles and Techniques*. MIT Press, Cambridge (2009)
- Koopman, B.O.: On distributions admitting a sufficient statistic. *Trans. Am. Math. Soc.* **39**(3), 399–409 (1936). <https://doi.org/10.1090/s0002-9947-1936-1501854-3>
- Lauritzen, S.L., Spiegelhalter, D.J.: Local computations with probabilities on graphical structures and their application to expert systems. *J. R. Stat. Soc. Ser. B (Methodol.)* **50**(2), 157–194 (1988)
- Lázaro-Gredilla, M., Lech, W., Gooskar, N., et al.: Query training: learning a worse model to infer better marginals in undirected graphical models with hidden variables. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 8252–8260 (2021)
- Li, Y., Tarlow, D., Brockschmidt, M., et al.: Gated graph sequence neural networks. *arXiv preprint* [arXiv:1511.05493](https://arxiv.org/abs/1511.05493) (2015)
- Li, Z., Tolias, A.S., Pitkow, X.: Learning dynamics and structure of complex systems using graph neural networks. *arXiv preprint* [arXiv:2202.10996](https://arxiv.org/abs/2202.10996) (2022)
- MacKay, D., Codes, E.C.: David Mackay's Gallager code resources. Dostupný z <http://www.inference.phy.cam.ac.uk/mackay/CodesFiles.html> (2009)
- Maron, H., Ben-Hamu, H., Serviansky, H., et al.: Provably powerful graph networks. *Adv. Neural Inf. Process. Syst.* **32** (2019)
- Milojević, S.: Principles of scientific research team formation and evolution. *Proc. Natl. Acad. Sci.* **111**(11), 3984–3989 (2014). <https://doi.org/10.1073/pnas.1309723111>
- Minka, T.P.: Expectation propagation for approximate Bayesian inference. In: *Proceedings of the Seventeenth Conference on Uncertainty in Artificial Intelligence*, pp 362–369. Morgan Kaufmann Publishers Inc. (2001)
- Morris, C., Ritzert, M., Fey, M., et al.: Weisfeiler and leman go neural: higher-order graph neural networks. In: *Proceedings of the AAAI Conference on Artificial Intelligence*, pp. 4602–4609 (2019)
- Murphy, R.L., Srinivasan, B., Rao, V., et al.: Relational Pooling for Graph Representations. *arXiv* [arXiv:1903.02541](https://arxiv.org/abs/1903.02541) (2019)
- Paszke, A., Gross, S., Massa, F., et al.: Pytorch: an imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* **32** (2019)
- Pearl, J.: *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann, Burlington (1988)
- Pitman, E.J.G.: Sufficient statistics and intrinsic accuracy. *Math. Proc. Camb. Philos. Soc.* **32**(4), 567–579 (1936). <https://doi.org/10.1017/s0305004100019307>
- Prüfer, H.: Neuer beweis eines satzes über permutationen. *Arch. Math. Phys.* **27**(1918), 742–744 (1918)
- Ranzato, M., Hinton, G.E.: Modeling pixel means and covariances using factorized third-order Boltzmann machines. In: *2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition*, pp. 2551–2558. IEEE (2010)
- Ritz, A., Tegge, A.N., Kim, H., et al.: Signaling hypergraphs. *Trends Biotechnol.* **32**(7), 356–362 (2014). <https://doi.org/10.1016/j.tibtech.2014.04.007>
- Rosenblatt, F.: The perceptron: a probabilistic model for information storage and organization in the brain. *Psychol. Rev.* **65**(6), 386 (1958)
- Sato, R.: A survey on the expressive power of graph neural networks. *arXiv preprint* [arXiv:2003.04078](https://arxiv.org/abs/2003.04078) (2020)
- Satorras, V.G., Welling, M.: Neural enhanced belief propagation on factor graphs. In: *International Conference on Artificial Intelligence and Statistics*, pp. 685–693. PMLR (2021)

- Scarselli, F., Gori, M., Tsoi, A.C., et al.: The graph neural network model. *IEEE Trans. Neural Netw.* **20**(1), 61–80 (2008)
- Sherrington, D., Kirkpatrick, S.: Solvable model of a spin-glass. *Phys. Rev. Lett.* **35**(26), 1792 (1975)
- Shimazaki, H., Sadeghi, K., Ishikawa, T., et al.: Simultaneous silence organizes structured higher-order interactions in neural populations. *Sci. Rep.* **5**(1), 1–13 (2015)
- Stan Development Team (2021) The Stan Core Library. <http://mc-stan.org/>, version 2.26.0
- Taranalli, V.: Commpy: digital communication with python. version 03 0. <https://github.com/veeresht/Commpy> (2015)
- Vaswani, A., Shazeer, N., Parmar, N., et al.: Attention is all you need. *Adv. Neural Inf. Process. Syst.* **30** (2017)
- Veličković, P., Cucurull, G., Casanova, A., et al.: Graph attention networks. *arXiv preprint arXiv:1710.10903* (2017)
- Wainwright, M.J., Jordan, M.I., et al.: Graphical models, exponential families, and variational inference. *Found. Trends Mach. Learn.* **1**(1–2), 1–305 (2008)
- Wainwright, M.J., Jaakkola, T.S., Willsky, A.S.: Tree-based reparameterization framework for analysis of sum-product and related algorithms. *IEEE Trans. Inf. Theory* **49**(5), 1120–1146 (2003)
- Watts, D.J., Strogatz, S.H.: Collective dynamics of ‘small-world’ networks. *Nature* **393**(6684), 440–442 (1998)
- Weiss, Y., Freeman, W.: Correctness of belief propagation in Gaussian graphical models of arbitrary topology. *Adv. Neural Inf. Process. Syst.* **12** (1999)
- Yedidia, J.S., Freeman, W.T., Weiss, Y., et al.: Generalized belief propagation. In: *NIPS*, pp. 689–695 (2000)
- Yoon, K., Liao, R., Xiong, Y., et al.: Inference in probabilistic graphical models by graph neural networks. In: *2019 53rd Asilomar Conference on Signals, Systems, and Computers*, pp. 868–875. *IEEE* (2019)
- You, J., Leskovec, J., He, K., et al.: Graph Structure of Neural Networks. *arXiv arXiv:2007.06559* (2020)
- Zhang, Y., Wang, X., Shi, C., et al.: Hyperbolic graph attention network. *IEEE Trans. Big Data* (2021)
- Zhang, Z., Wu, F., Lee, W.S.: Factor graph neural networks. *Adv. Neural Inf. Process. Syst.* **33**, 8577–8587 (2020)
- Zhou, J., Cui, G., Hu, S., et al.: Graph neural networks: a review of methods and applications. *AI Open* **1**, 57–81 (2020). <https://doi.org/10.1016/j.aiopen.2021.01.001>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.