

Efficient Localized Inference for Large Graphical Models

Jinglin Chen¹, Jian Peng¹, Qiang Liu²

¹ University of Illinois at Urbana-Champaign

² University of Texas at Austin

jinglinc@illinois.edu, jianpeng@illinois.edu, lqiang@cs.utexas.edu

Abstract

We propose a new localized inference algorithm for answering marginalization queries in large graphical models with the correlation decay property. Given a query variable and a large graphical model, we define a much smaller model in a local region around the query variable in the target model so that the marginal distribution of the query variable can be accurately approximated. We introduce two approximation error bounds based on the Dobrushin's comparison theorem and apply our bounds to derive a greedy expansion algorithm that efficiently guides the selection of neighbor nodes for localized inference. We verify our theoretical bounds on various datasets and demonstrate that our localized inference algorithm can provide fast and accurate approximation for large graphical models.

1 Introduction

Probabilistic graphical models such as Bayesian networks, Markov random fields, and conditional random fields are powerful tools for modeling complex dependencies over a large number of random variables ([Koller and Friedman, 2009; Wainwright *et al.*, 2008]). Graphs are used to represent joint probability distributions, where nodes denote random variables, and edges represent dependency relationships between different nodes. With the specification of a graphical model, a fundamental problem is to calculate the marginal distributions of variables of interest. This problem is closely related to computing the partition function, or the normalization constant of a graphical model, which is known to be intractable and #P-complete. As a result, developing efficient approximation inference algorithms becomes a pressing need. The most popular algorithms include deterministic variational inference and Markov Chain Monte Carlo sampling.

However, many challenging practical problems involve very large graphs on which it is computationally expensive to use existing variational inference or Monte Carlo sampling algorithms. This happens, for example, when we use Markov random fields to represent the social network of Facebook or use a Bayesian network to model the knowledge graph that is derived from the entire Wikipedia, where in both cases the sizes of the graphical models can be prohibitively large

(e.g., millions or billions of variables). It is thus infeasible to perform traditional approximate inference such as message passing or Monte Carlo on these models because such methods need to traverse the entire model to make an inference. Despite the daunting sizes of large graphical models, in most real-world applications, users only want to make an inference on a set of query variables of interest. The distribution of a query variable is often only dependent on a small number of nearby variables in the graph. As a result, complete inference over the entire graph is not necessary and practical methods should perform inference only with the most relevant variables in local graph regions that are close to the query variables, while ignoring the variables that are weakly correlated and/or distantly located on the graph.

In this work, we develop a new localized inference method for very large graphical models. Our approach leverages the Dobrushin's comparison theorem that casts explicit bounds based on the correlation decay property in the graphs, in order to restrict the inference to a smaller local region that is sufficient for the inference of marginal distribution of the query variable. The use of the Dobrushin's comparison theorem allows us to explicitly bound the truncation error which guides the selection of localized region from the original large graph. Extensive experiments demonstrate both the effectiveness of our theoretical bounds and the accuracy of our inference algorithm on a variety of datasets.

Related Work Approximate inference algorithms of graphical models have been extensively studied in the past decades (see [Koller and Friedman, 2009; Wainwright *et al.*, 2008; Dechter, 2013]). Query-specific inference, including [Checheta and Guestrin, 2010] which proposed a focused belief propagation for query specific inference, and [Wick and McCallum, 2011; Shi *et al.*, 2015] which study query-aware sampling algorithms, have recently been introduced for large graphical models. Compared with these methods, our work is theoretically motivated by the Dobrushin's comparison theorem and enables us to efficiently construct the localized region in a principled and practically efficient manner.

2 Background on Graphical Models

Graphical models provide a flexible framework for representing relationships between random variables [Heine-

mann and Globerson, 2014]. In graph G , we use $X = (X_1, X_2, \dots, X_n)$ to denote a finite collection of n random variables and we use $\mathbf{x} = (x_1, x_2, \dots, x_n)$ to refer to an assignment. Suppose E is a set of edges and θ is a set of functions with $\theta_{ij}(x_i, x_j)$ for edge $\langle i, j \rangle \in E$ and $\theta_i(x_i)$ for node $i \in \{1, 2, \dots, n\}$. We use $\mu(\mathbf{x}; \theta)$ to represent the joint distribution of the graphical model (G, E, θ) as following,

$$\mu(\mathbf{x}; \theta) = \frac{1}{Z_\mu(\theta)} \exp \left(\sum_{\langle i, j \rangle \in E} \theta_{ij}(x_i, x_j) + \sum_i \theta_i(x_i) \right),$$

where $Z_\mu(\theta)$ is the normalization constant (also called the partition function).

In this work, we will focus on the Ising model, an extensively studied graphical model. The Ising model is a pairwise model with binary variables $x_i \in \chi = \{-1, +1\}$. The pairwise and singleton parameters are defined as follows

$$\theta_{ij}(x_i, x_j) = \begin{pmatrix} J_{ij} & -J_{ij} \\ -J_{ij} & J_{ij} \end{pmatrix}, \quad \theta_i(x_i) = \begin{pmatrix} -h_i \\ h_i \end{pmatrix}.$$

So the distribution of an Ising model is defined as,

$$\mu(\mathbf{x}; \theta) = \frac{1}{Z_\mu(\theta)} \exp \left(\sum_{\langle i, j \rangle \in E} J_{ij} x_i x_j + \sum_i h_i x_i \right). \quad (1)$$

Given a graphical model, marginal inference involves calculating the normalization constant, or the marginal probabilities of small subsets of variables. These problems require summation over an exponential number of configurations and are typically #P-hard in the worst case for loopy graphical models. However, practical problems can be often easier than the theoretically worst cases, and it is still possible to obtain efficient approximations by leveraging the special structures of given models. In this work, we focus on the query-specific inference, where the goal is to calculate the marginal distribution $\mu(x_i)$ of given individual variable x_i . For this task, it is possible to make good approximations based on a local region around x_i , thus significantly accelerates the inference in very large graphical models.

3 Localized Inference and Correlation Decay

Given a large graphical model, it is usually not feasible to compute the exact marginal of a specific variable due to the exponential time complexity. Furthermore, it is even not practical to perform the variational approximation algorithms, such as mean field and belief propagation, when the graph is very large. This is because these traditional methods need to traverse the entire graph multiple times before convergence, and thus are prohibitively slow for very large models such as these built on social networks or knowledge bases.

On the other hand, it is relatively cheap to calculate exact or approximate marginals in small or medium size graphical models. In many applications, users are only interested in certain queries of node marginals. Because users' queries of interest often have strong associations with only a small number of nearby variables in the graph, the complete inference over the full graph is not necessary. This can be formally captured by the phenomenon of correlation decay, that is, when the graph G is large and sparse, the influence of a random




Figure 1: The key goal of this work is to approximate queries in large scale graphical models using smaller models on local regions.

variable on the distribution of another random variable decreases quickly as the distance of the shortest path between the corresponding nodes in the graph G increases.

The correlation decay property has been widely studied in statistical mechanics and graphical models [Rebeschini *et al.*, 2015].

If a graphical model satisfies the property of correlation decay, it is possible that we can use only the local information in the graph to perform marginal inference, as the distant variables have little correlation with the query variable. This intuition allows us to use the information from the most relevant variables in the local region close to the queried variable to efficiently approximate its marginal distribution. Assume that μ is a large graphical model, and we want to calculate a marginal distribution $\mu(x_i)$ of variable i . Localized inference constructs a much smaller model $\nu(x_\alpha)$, defined on a small subgraph α that includes i , such that $\nu(x_i) \approx \mu(x_i)$. The challenge here, however, is how to construct a good localized model and bound its approximation error. We address this problem via the Dobrushin's comparison theorem [Föllmer, 1982], and propose an efficient algorithm to find the local graph region for a given query node and provide an error bound between its approximate and true marginals. To get started, we first introduce the Dobrushin's comparison theorem, which is used to compare two Gibbs measures.

Theorem 1 [Föllmer, 1982] *Dobrushin's comparison theorem* Let μ be a Gibbs measure on a finite product space $E = S^I$, where I is an index set. For $i, j \in I$, we define

$$C_{ij} = \frac{1}{2} \sup \{ \|\mu_i(x_i|\mathbf{x}) - \mu_i(y_i|\mathbf{y})\| : x_k = y_k, \forall k \neq j \},$$

where $\mu_i(x_i|\mathbf{x})$ is the conditional distribution of the i th coordinate with respect to the σ -field generated by the coordinates with index $j \neq i$, and $\|\cdot\|$ is the total variance distance. We compute

$$c = \max_{i \in I} \sum_{j \in I} C_{ij}, \quad (2)$$

and assume $c < 1$. Let $C = (C_{ij})_{i, j \in I}$ and $D = \sum_{n=0}^{\infty} C^n = (I - C)^{-1}$, then for any probability measure ν on the same place and any function f , we have

$$\left| \int f d\mu - \int f d\nu \right| \leq \sum_{i \in I} (Db)_i \times \delta_i(f),$$

where $b = (b_j)$ is the singleton perturbation coefficient of node x_j :

$$b_j = \frac{1}{2} \sup_{\mathbf{x}} \|\mu_j(x_j|\mathbf{x}) - \nu_j(x_j|\mathbf{x})\|, \quad (3)$$

and $\delta_i(f)$ is the oscillation of f in the i th coordinate, that is,

$$\delta_i(f) = \max_{x_i, x'_i} |f(\mathbf{x}_{-i}, x_i) - f(\mathbf{x}_{-i}, x'_i)|.$$

In Theorem 1, $\mu_i(x_i|\mathbf{x})$ is the probability of variable i conditioned on its adjacent variables whose assignments are the same as corresponding entries in \mathbf{x} . According to the Markov property, calculating $\mu_i(x_i|\mathbf{x})$ only requires information from the local star-shaped graph (the figure can be found in the full version). It is worth noting that a tighter bound can be obtained by defining b to be $b_j = \frac{1}{2} \int \|\mu_j(x_j|\mathbf{x}) - \nu_j(x_j|\mathbf{x})\| \nu(d\mathbf{x})$. Here we use the definition in (3) for lower computational complexity. The matrix C is known as the Dobrushin's interaction matrix, and the inequality $c = \max_{i \in I} \sum_{j \in I} C_{ij} < 1$ is the Dobrushin condition. If this condition holds, the theorem can give us a bound between two measures, which is the result of correlation decay.

In the following, we will apply Theorem 1 to undirected graphical models to derive an approximation bound of marginal distributions. We first denote $I = \{1, 2, \dots, n\}$ by the index set of the variables and assume that we want to query the marginal distribution of variable x_i . In order to apply Theorem 1, we set $f(\mathbf{x})$ to be the indicator function of the variable x_i , that is, $f(\mathbf{x}) = f(x_1, x_2, \dots, x_n) = \mathbb{I}[x_i = k]$. Then $|\int f d\mu - \int f d\nu|$ becomes the absolute marginal difference $|\mu(x_i = k) - \nu(x_i = k)|$ between the two measures μ and ν . In addition, the oscillation of function f is thus reduced to $\delta_i(f) = 1$ and $\forall j \neq i, \delta_j(f) = 0$. With these simplifications, we obtain a bound of the maximum difference between marginals of the queried node i for two measures:

Corollary 1 *Following the assumptions in Theorem 1 and the above text, we have*

$$\max_k |\mu(x_i = k) - \nu(x_i = k)| \leq (Db)_i. \quad (4)$$

Note that the roles of μ and ν in (4) are not symmetric because the Dobrushin coefficient D is solely defined based on μ (and independent with ν). As a result, there are two ways to use bound (4) for localized inference, depending on whether we treat μ or ν as the original model that we want to query or the localized model that we use for approximation, respectively. We will next exploit both possibilities in the next sections. In Section 4, we take μ as the global model (or measure) and ν as the localized model (or measure) and derive a simple upper bound relates the approximation error to the distance between the query node and the boundary of the local region on the graph. In Section 5, we take ν as the global model (or measure) and μ as the localized model (or measure), we derive another upper bound that only involves the localized region, and leverage it to propose a greedy expansion algorithm to construct the localized model with guaranteed approximations.

4 Distance-based Upper Bound

In this section, we assume that $\mu = \mu(x_1, x_2, \dots, x_n)$ in Theorem 1 is defined by the original graphical model that we want to query, and ν is a simpler and more tractable distribution that we use to approximate the marginal of x_i in μ .

For notational simplicity, we partition the node-set to two disjoint sets α and $\beta = \{1, 2, \dots, n\} \setminus \alpha$, where α is the local subgraph that contains query node $i \in \alpha$ and β is the rest of the graph. We use $\partial\alpha$ and $\overset{\circ}{\alpha}$ to represent the set of subscripts of nodes on the boundary and in the interior of α . Obviously, $\partial\alpha \subseteq \alpha$, $\overset{\circ}{\alpha} \subseteq \alpha$, and $\partial\alpha \cup \overset{\circ}{\alpha} = \alpha$. Similarly, $\partial\beta \subseteq \beta$, $\overset{\circ}{\beta} \subseteq \beta$, and $\partial\beta \cup \overset{\circ}{\beta} = \beta$. In addition, we use x_α to denote the variables in α and x_β to denote the variables in β . We will first apply the following lemma to obtain our first result on the relationship between the approximation error of marginals and the radius of the local subgraph α .

Lemma 1 [Rebeschini and van Handel, 2014] *Assume I is a finite set and let m be a pseudo-metric on set I . $C = (C_{ij})_{i,j \in I}$ is a non-negative matrix. Suppose that*

$$\max_{i \in I} \sum_{j \in I} e^{m(i,j)} C_{ij} \leq d < 1.$$

Then matrix $D = \sum_{n=0}^{\infty} C^n = (I - C)^{-1}$ satisfies

$$\max_{i \in I} \sum_{j \in I} e^{m(i,j)} D_{ij} \leq \frac{1}{1-d}.$$

In particular, this implies that

$$\sum_{j \in J} D_{ij} \leq \frac{e^{-m(i,J)}}{1-d}$$

for every set $J \subseteq I$, where $m(i, J) = \max_{j \in J} m(i, j)$.

This lemma indicates that if C_{ij} decays exponentially with the distance between i and j , the D_{ij} , which is used in Theorem 1 and Corollary 1, also decays exponentially with the distance between i and j . The condition of this correlation decay lemma is usually mild in practice. When we choose $m \equiv 0$, which is naturally a pseudo-metric, and use Dobrushin's interaction matrix as C , the conditions of the lemma hold once the Dobrushin condition is satisfied, because matrix C in Theorem 1 is by definition a non-negative matrix and hence D is also non-negative, and every entry in b is less than $1/2$. Applying Lemma 1, we can obtain the following result.

Theorem 2 *Suppose μ is the probability measure for a graphical model for which we want to query the marginal distribution of node i . Let ν be the another probability on the same space, whose parameters of edges on subgraph α and parameters of nodes in $\overset{\circ}{\alpha}$ are the same as μ . Assume the Dobrushin condition holds for μ ($c = \max_{i \in I} \sum_{j \in I} C_{ij} < 1$). Let $d(i, \partial\alpha)$ denote the distance between node i and node-set $\partial\alpha$ on the Markov graph G of μ . If we assume*

$$d(i, \partial\alpha) \geq \frac{\ln \frac{t}{2\varepsilon(t-1)(1-c)}}{\ln \frac{1+(t-1)c}{tc}} \text{ for some } t > 1, \quad (5)$$

then $\forall \varepsilon > 0$, we have

$$\max_k |\mu(x_i = k) - \nu(x_i = k)| \leq \varepsilon.$$

This theorem characterizes the error bound when approximating the global model μ using another model ν that matches μ locally in region α . Our result shows that in order to ensure an ε bound on the query node i , the distance $d(i, \partial\alpha)$ from the query node i to the boundary $\partial\alpha$ should be at least linear to $\ln(1/\varepsilon)$. In other words, the error ε decreases exponentially with $d(i, \partial\alpha)$. The proof of Theorem 2 can be found in the Appendix C.

As a result, given c and ε , we can get the minimum value of the lower bound of $d(i, \partial\alpha)$ by optimizing t . Theorem 2 gives a simple but general way to bound the local subgraph of variable x_i , as we only need to check the Dobrushin condition and compute c on the whole true graphical model.

5 Localized Bound and Greedy Expansion

The bound in Theorem 2 requires computing the value c as defined in (2) for a given graphical model. However, since c is the maximum C_{ij} of the entire graph, it can be very expensive to compute when the graph is large. In this section, we explore another approach of using the bound in Corollary 1, by setting $\nu = \nu(x_1, x_2, \dots, x_n)$ to be the distribution of the original graphical model and μ to be the localized model. In this way, we will derive a novel approximation approach by greedily constructing a local graph from the query variable i , with guaranteed upper bounds of the approximation error between marginal distributions of μ and ν .

To start with, we note that ν can be decomposed to

$$\nu(x) = \frac{\psi_\alpha(x_\alpha)\psi_\beta(x_\beta)\tilde{\psi}_{\partial\alpha\partial\beta}(x_{\partial\alpha}, x_{\partial\beta})}{Z_\nu},$$

where ψ_α is the exponential of a potential function of x_α , ψ_β is the exponential of potential function of x_β , and $\tilde{\psi}_{\partial\alpha\partial\beta}$ is the exponential of potential function defined on $x_{\partial\alpha}$ and $x_{\partial\beta}$.

We want to approximate ν with a simpler model μ in which the nodes in α and β are disconnected, so that the inference over $i \in \alpha$ can be performed locally within α , irrelevant to the nodes in β . Formally, we want to approximate μ by

$$\mu(x) = \frac{\psi_\alpha(x_\alpha)\psi_\beta(x_\beta)\tilde{\psi}_{\partial\alpha}(x_{\partial\alpha})\tilde{\psi}_{\partial\beta}(x_{\partial\beta})}{Z_\mu},$$

which replaces the factor $\tilde{\psi}_{\partial\alpha\partial\beta}$ with a product $\tilde{\psi}_{\partial\alpha}\tilde{\psi}_{\partial\beta}$ with approximations $\tilde{\psi}_{\partial\alpha}$ and $\tilde{\psi}_{\partial\beta}$. Therefore, the marginal distributions of x_α and x_β get decoupled in μ , that is,

$$\mu(x) = \mu(x_\alpha)\mu(x_\beta).$$

This decomposition thus allows us to approximately calculate marginal $\mu(x_i)$ efficiently within subgraph α . The challenges here are 1) how to construct the factors $\tilde{\psi}_{\partial\alpha}$ and $\tilde{\psi}_{\partial\beta}$ in μ to closely approximate ν , 2) how to decide the subgraph region and 3) how to bound the approximation error. We consider two methods for constructing $\tilde{\psi}$ in this work:

1. [Dropping out] Simply remove the $\tilde{\psi}_{\partial\alpha\partial\beta}$ in ν . To do so, we set

$$\tilde{\psi}_{\partial\alpha} = \tilde{\psi}_{\partial\beta} = 1. \quad (6)$$

This corresponds to directly remove all the edges between $\partial\alpha$ and $\partial\beta$, which is also referred as the ‘‘dropping out’’ method in our experiments.

2. [Mean field] Find $\tilde{\psi}_{\partial\alpha}\tilde{\psi}_{\partial\beta}$ to closely approximate $\tilde{\psi}_{\partial\alpha\partial\beta}$ by performing a mean field approximation, that is, we solve the following optimization problem:

$$\min_{\tilde{\psi}_{\partial\alpha}, \tilde{\psi}_{\partial\beta}} \text{KL}(\tilde{\psi}_{\partial\alpha}\tilde{\psi}_{\partial\beta} \parallel \tilde{\psi}_{\partial\alpha\partial\beta}), \quad (7)$$

where the $\text{KL}(\cdot \parallel \cdot)$ refers to the KL divergence of the corresponding normalized distributions. To apply the mean field approximation and reduce complexity, we further assume that the nodes are independent in $\tilde{\psi}_{\partial\alpha}$ and $\tilde{\psi}_{\partial\beta}$. By using the optimized approximation $\tilde{\psi}_{\partial\alpha}$, we will be able to compensate the error of marginal of x_i , which is introduced by simply removing the edges between $\partial\alpha$ and $\partial\beta$, as mentioned above.

Note that the potentials $\tilde{\psi}_{\partial\beta}$ and ψ_β in μ do not influence the calculation of $\mu(x_i)$, for $i \in \alpha$. Therefore, we remove all the edges in β . The marginal of node $i \in \alpha$ will not change.

By applying Corollary 1, we can now obtain an error bound which, remarkably, only involves the local region α . The details can be found in Appendix B.

Corollary 2 Assume $\mu(x) = \mu(x_\alpha)\mu(x_\beta)$, and the conditions in Theorem 1 holds, we have

$$\max_k |\mu(x_i = k) - \nu(x_i = k)| \leq \sum_{j \in \partial\alpha} D_{ij}b_j, \quad (8)$$

where b_j is defined in Eq 3, and D is defined by $D_{\alpha\alpha} = (I - C_{\alpha\alpha})^{-1}$; here C is defined in Theorem 1.

Note that the upper bound in (8) only involves the local region α and hence can be computed efficiently using mean field or belief propagation within the subgraph on α . The proof of Corollary 2 and the details on how to calculate C and D for Ising models in practice can be found in the full version.

Using the bound in (8), we propose a greedy algorithm to expand the local graph starting from query node i incrementally. At iteration, we add a neighboring node that yields the tightest bound using the above bound and repeat this process until the bound is tight enough or a maximum of graph size is reached. This process is summarized in Algorithm 1. After we complete the expanding phase, we can apply exact inference or on local region α to calculate the marginal of the query x_i if the size of α is small or perform approximate inference methods if the size of α is medium. The actual size of α can vary in different graphical models, which is mainly determined by the correlation decay property near the query variable x_i or the tightness of the upper bound in Eq (8).

Computational Complexity We find that both the complexity of expanding the local subgraph and the complexity of localized inference can be bounded. The details will be included in the Appendix A.

6 Experiments

We test our algorithm on both simulated and real-world datasets. The results indicate that our method provides an efficient localized inference technique.

Algorithm 1 Greedy expansion algorithm for localized inference

- 1: given a graphical model ν and a node i , approximate marginal probability $\nu(x_i)$
 - 2: **input:** $K =$ maximum number of nodes in the local subgraph α and $\delta =$ the improvement threshold
 - 3: initialize local subgraph $\alpha = \{i\}$ and $bound_{best} = 1$
 - 4: **while** $|\alpha| < K$ ($|\alpha|$ is the number of nodes in α) **do**
 - 5: set $\beta = \{1, 2, \dots, n\} \setminus \alpha$ and $\partial\beta$ the nodes in β that connects with α in ν .
 - 6: **for** node $k \in \partial\beta$ **do**
 - 7: add node k to α and get a candidate local subgraph $\alpha_k^{new} = \alpha \cup \{k\}$.
 - 8: construct local model μ by setting $\tilde{\psi}_{\partial\alpha} = 1$ (dropping out, Eq (6)) or estimating it using mean field as in Eq (7).
 - 9: calculate the bound $bound_k$ in (8) where the $bound_k$ refers to $\sum_{j \in \partial\alpha_k^{new}} D_{ij}b_j$
 - 10: **end for**
 - 11: **if** $\min_{k \in \partial\beta} bound_k < bound_{best} - \delta$ **then**
 - 12: update $bound_{best} = \min_{k \in \partial\beta} bound_k$
 - 13: update $\alpha = \alpha \cup \{\text{argmin}_{k \in \partial\beta} bound_k\}$
 - 14: **end if**
 - 15: **end while**
-

6.1 2D Ising Grid

In this section, we perform experiments on 2D-grid Ising models and regard the localized probability as ν and regard the true probability as μ . The graph is a 10×10 lattice and the coordinate of query node is $(5, 5)$. The parameters in the Ising model is generated by drawing h_i uniformly from $[-I_1, I_1]$ for all nodes i and J_{ij} uniformly from $[-I_2, I_2]$ for all edges $\langle i, j \rangle$. Here I_1 and I_2 control the locality and hardness of this Ising model.

Checking Dobrushin’s condition We start with numerically checking the Dobrushin condition $c = \max_i \sum_j C_{ij} < 1$. We find that c is smaller than one in most regions, but is larger than one when I_2 is very large and I_1 is very small, in which case the nodes are strongly coupled together (no correlation decay) and there is no significant local information. The hope, however, is that real problems tend to be easier because a large amount of information is available. The figure is shown in the full version.

Comparing Different Expansion Algorithms In this part, we compare the true approximation error $\max_k |\mu(x_i = k) - \nu(x_i = k)|$ to the bound $\sum_{j \in \partial\alpha} D_{ij}b_j$ given by our algorithm when we expand the local subgraph. The true error $\max_k |\mu(x_i = k) - \nu(x_i = k)|$ is evaluated using the brute-force algorithm. When removing the bipartite graph, we try both simply dropping edges and the mean field approximation. In all the experiments, we use the UGM Matlab package¹ for the mean field approximation.

In order to better compare the error, we also add two baselines. One baseline is that we expand the local subgraph in

¹<http://www.cs.ubc.ca/~schmidtm/Software/UGM.html>




Figure 2: The true errors and our upper bounds for the marginal approximation when we expand the local subgraphs to different sizes.

each step by randomly selecting a node in the boundary $\partial\beta$. Another baseline is that we expand the local subgraph greedily by choosing the node in $\partial\beta$ that has the maximum L_2 norm over the edge-set between such node and the subgraph α . The intuition is that when the edges’ weights are large, the node may be more related to the nodes in the subgraph.

In Figure 2, we compare our greedy expansion method stated in Algorithm 1 to the baselines stated above. For this experiment, we fix $I_1 = 1$ and $I_2 = 0.25$ and average on 100 random trials. We stop expanding the graph when the local subgraph contains 16 nodes. We calculate the mean value of the true errors and bounds in the 100 trials for a different number of nodes in the subgraph.

From Figure 2, we can see that, when combined the dropping out method for constructing $\tilde{\psi}$, our greedy expansion method significantly outperforms the two baselines. We also find that the mean field method for constructing $\tilde{\psi}$ gives about the same true error as the dropping out method, but provides a tighter upper bound. It is interesting to note that the true errors of the two baseline expansion methods are sometimes even worse than the upper bounds of our greedy expansion, indicating the strong advantage of our method.

We further investigate how the parameters of the Ising model may influence the results of the algorithms and the tightness of the bound. For this purpose, we fix $I_2 = 0.25$ and vary I_1 in the range of $\{0, 0.5, 1, \dots, 10\}$ in Figure 3. For each setting, we simulate 100 times and then calculate the mean error and bound. From Figure 3, we can find that the bound is again relatively tight, especially when the value of I_1 is large. both the bounds and the true errors decrease as I_1 increases because the correlation decay is stronger and the inference task is easier with strong local evidence on the singleton potentials (large I_1).

6.2 Cora Data Set

We perform experimental evaluations on the Cora data set². Cora consists of a large collection of machine learning papers with citation relations between the papers, in which each

²<https://people.cs.umass.edu/~mccallum/data.html>




Figure 3: Our bounds and the true errors vs. different I_1 .

paper is labeled as one of seven classes. For our experiment, we binarize the labels by taking “Neural Networks” as label 1 and the remaining classes as label -1 . We process the data by removing the hubs in the graph and truncate the graph to have a maximum degree of 15; this is done by randomly deleting edges of the nodes whose degree is larger than 15 until the whole graph is degree bounded by 15. We then experiment on the maximum connected subgraph, which consists of 2389 nodes and 4325 edges.

In order to construct an Ising model, we random draw edge potentials by $J_{ij} \sim N(0.25, 0.05)$ for each edge of the citation graph, and draw the singleton potentials by $h_i \sim N(0.1I_1, 1)$ for nodes with true label 1, and $h_i \sim N(-0.1I_1, 1)$ for nodes with true label -1 . Here I_1 is a parameter that we choose from $\{0, 1, \dots, 10\}$. When I_1 , the node potentials increases so that marginal is more dominated by the status of the query node and the querying is easier.

Comparing local inference with global inference In this part, we want to compare the performance of inference on the local graph to the inference on the global graph. Since the global graph is too large, we can only use approximate inference algorithm. Here, we use mean field to do the global inference and use it as a baseline. For the local graph, we expand the graph greedily as stated in Algorithm 1 and choose a threshold of $\delta = 0.005$ and stop expanding when the subgraph already has 16 nodes.

For $I_1 \in \{0, 1, \dots, 10\}$, we query the same 500 nodes randomly selected out of the 2389 nodes and evaluate their marginal distributions. In global inference and local inference, we have the marginal on the each query node. If the marginal is larger than 0.5, we consider our inference algorithm give it label 1, whereas if the marginal is less than 0.5, we give it label -1 .

In Figure 4, we report the accuracy of the labels given by the global and local inference evaluated w.r.t. the true labels, as well as the accuracy of the local inference evaluated w.r.t. the labels provided by global inference. We find that as I_1 increases, both the accuracies of global and local inference w.r.t. the true labels increase significantly. In addition, the local inference gives similar result as the global inference (the green curve is high) and the accuracy increases as I_1 increases




Figure 4: The accuracy of different algorithms when I_1 changes. Red and blue: the accuracy of the labels given by the global and the local inference evaluated w.r.t. the true labels. Green: the accuracy of the local inference evaluated w.r.t. the labels provided by the global inference.




Figure 5: Precision, Recall, and F-measure of the labels given by local inference vs. global inference when the value of I_1 changes.

as well. We also report in Figure 5 the precision, recall, and F-measure when comparing local inference with global inference, by treating label 1 as positives. Both figures show that when I_2 is fixed and I_1 increases, which means that the correlation decay is stronger, our local inference method achieves better results.

7 Conclusion

In this paper, we address query-specific marginal inference in large-scale graphical models using a new localized inference algorithm. We leverage the Dobrushin’s comparison theorem to derive two error bounds for localized inference, including a simple bound based on graph distance and a localized bound from which we derive an efficient greedy expansion algorithm for constructing local regions for localized inference. Our experiments have shown that our bounds are practically useful and the algorithm works efficiently on various graphical models. Future directions include theoretical investigation on tighter bounds and development of more efficient greedy expansion algorithms.

A Computational Complexity

Here we consider the computational complexity of expanding the local subgraph and the complexity of localized inference. We always suppose that the maximum degree of the graph G is d and we define the maximum distance between the query node and any node in the subgraph to be the radius of the subgraph. First, given a threshold ε , from Theorem 2, we just need a subgraph with radius

$$r = \inf_{t>1} \left\lceil \frac{\ln \frac{t}{2\varepsilon(t-1)(1-c)}}{\ln \frac{1+(t-1)c}{tc}} \right\rceil,$$

where we recall that c is the Dobrushin coefficient $c = \max_{i \in I} \sum_{j \in I} C_{ij}$. In particular, taking $t = 2$ shows that we just need $r = \lceil -\ln(\varepsilon(1-c)) / \ln \frac{1+c}{2c} \rceil$. It is worth noting that r decreases when c becomes small and/or the accuracy threshold ε becomes large. Since the size of the subgraph with radius r is no more than $1 + d + d^2 + \dots + d^r = (d^{r+1} - 1)/(d - 1)$, it can be much smaller than the whole graph. As a result, the inference over the subgraph is much more efficient.

Then, we discuss the computation complexity in each expansion step (Algorithm 1, line 6-10). We need to loop over the nodes in $\partial\beta$. In the loop, we need to calculate the vector b and the matrix D . The calculation for each element in b requires the enumeration of different assignments in the neighborhood of such node, which is bounded because it is not related to the size of the whole graph. In the calculation of matrix C , we only need to update the elements related to the new node. The number of such elements is no more than d and the calculation of each element is not related to the size of the whole graph. D can be derived from C and use historical information to calculate incrementally. The complexity is no more than computing the inverse of the whole matrix $I - C_{\alpha\alpha}$. If we use mean field approximation in the greedy expansion, the computation is also cheap because the sizes of $\partial\alpha$ and $\partial\beta$ are small.

B Proof of Corollary 2

Proof 1 Note that the Dobrushin’s interaction matrix C of μ is a block diagonal matrix. Since there are no edges between α and β , the corresponding blocks equal to zero. If the Dobrushin condition holds, D would also be a block-diagonal matrix and can be calculated easily from C . To see this, we have

$$C = \begin{bmatrix} \alpha & \beta \\ C_{\alpha\alpha} & O \\ O & C_{\beta\beta} \end{bmatrix} \begin{matrix} \alpha \\ \beta \end{matrix}$$

$$\text{and } D = \begin{bmatrix} \alpha & \beta \\ (I - C_{\alpha\alpha})^{-1} & O \\ O & (I - C_{\beta\beta})^{-1} \end{bmatrix} \begin{matrix} \alpha \\ \beta \end{matrix}.$$

Applying the bound in Corollary (1) gives the result.

C Proof of Theorem 2

Proof 2 Let $m(i, j) = d(i, j) \ln \frac{1+(t-1)c}{tc}$, where $d(i, j)$ represents the distance between node i and node j in the graph and $t > 1$. Then

$$\begin{aligned} \max_{i \in I} \sum_{j \in I} e^{m(i,j)} C_{ij} &= \frac{1 + (t-1)c}{tc} \max_{i \in I} \sum_{j \in I} e^{d(i,j)} C_{ij} \\ &\leq \frac{1 + (t-1)c}{t} < 1. \end{aligned}$$

Applying Lemma 1, we have

$$\begin{aligned} \max_{k \in \{1, -1\}} |\mu(x_i = k) - \nu(x_i = k)| \\ \leq \frac{e^{-m(i, \partial\alpha)}}{1 - \frac{1+(t-1)c}{t}} = \frac{e^{-d(i, \partial\alpha) \ln \frac{1+(t-1)c}{tc}}}{(t-1)(1-c)}. \end{aligned}$$

Substituting the inequality of $d(i, \partial\alpha)$ in the condition into right-hand side yields to the result.

References

- [Chechetka and Guestrin, 2010] Anton Chechetka and Carlos Guestrin. Focused belief propagation for query-specific inference. In *AISTATS*, pages 89–96, 2010.
- [Dechter, 2013] Rina Dechter. Reasoning with probabilistic and deterministic graphical models: Exact algorithms. *Synthesis Lectures on Artificial Intelligence and Machine Learning*, 7(3):1–191, 2013.
- [Föllmer, 1982] Hans Föllmer. A covariance estimate for gibbs measures. *Journal of Functional Analysis*, 46(3):387–395, 1982.
- [Heinemann and Globerson, 2014] Uri Heinemann and Amir Globerson. Inferring with high girth graphical models. In *ICML*, pages 1260–1268, 2014.
- [Koller and Friedman, 2009] Daphne Koller and Nir Friedman. *Probabilistic graphical models: principles and techniques*. MIT press, 2009.
- [Rebeschini and van Handel, 2014] Patrick Rebeschini and Ramon van Handel. Comparison theorems for gibbs measures. *Journal of Statistical Physics*, 157(2):234–281, 2014.
- [Rebeschini et al., 2015] Patrick Rebeschini, Ramon Van Handel, et al. Can local particle filters beat the curse of dimensionality? *The Annals of Applied Probability*, 25(5):2809–2866, 2015.
- [Shi et al., 2015] Tianlin Shi, Jacob Steinhardt, and Percy Liang. Learning where to sample in structured prediction. In *AISTATS*, 2015.
- [Wainwright et al., 2008] Martin J Wainwright, Michael I Jordan, et al. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305, 2008.
- [Wick and McCallum, 2011] Michael L Wick and Andrew McCallum. Query-aware mcmc. In *Advances in Neural Information Processing Systems*, pages 2564–2572, 2011.