

SCIENTIFIC REPORTS



OPEN

Quantum Enhanced Inference in Markov Logic Networks

Peter Wittek^{1,2} & Christian Gogolin¹

Received: 06 December 2016

Accepted: 02 March 2017

Published: 19 April 2017

Markov logic networks (MLNs) reconcile two opposing schools in machine learning and artificial intelligence: causal networks, which account for uncertainty extremely well, and first-order logic, which allows for formal deduction. An MLN is essentially a first-order logic template to generate Markov networks. Inference in MLNs is probabilistic and it is often performed by approximate methods such as Markov chain Monte Carlo (MCMC) Gibbs sampling. An MLN has many regular, symmetric structures that can be exploited at both first-order level and in the generated Markov network. We analyze the graph structures that are produced by various lifting methods and investigate the extent to which quantum protocols can be used to speed up Gibbs sampling with state preparation and measurement schemes. We review different such approaches, discuss their advantages, theoretical limitations, and their appeal to implementations. We find that a straightforward application of a recent result yields exponential speedup compared to classical heuristics in approximate probabilistic inference, thereby demonstrating another example where advanced quantum resources can potentially prove useful in machine learning.

Graphical models combine uncertainty and logical structure in an intuitive representation. Examples include Bayesian networks, Markov networks, conditional random fields, and hidden Markov models, but also Ising models and Kalman filters. Their main advantage is the compactness of representation, stemming from capturing the sparsity structure of the model and independence conditions among the variables reflected in the correlations. The graph structure encompasses the qualitative properties of the distribution. Exact probabilistic inference in a general Bayesian or Markov network is #P-complete¹, which is why one often resorts to Markov chain Monte Carlo (MCMC) Gibbs sampling to approximate exact probabilistic inference. However, the task remains computationally intensive even with MCMC.

Graphical models belong to a school of machine learning that emphasizes the importance of probability theory. First-order logic on the contrary comes from the symbolist tradition of artificial intelligence and it relies on inverse deduction to perform inference. Markov logic networks reconcile the two schools, and in one limit, they recover first-order logic². A Markov logic network is essentially a template for generating Markov networks based on a knowledge base of first-order logic. MCMC Gibbs sampling can be used in the same way as in ordinary Markov networks to perform approximate probabilistic inference, but it suffers from the enormous number of nodes that are generated by the template.

There has been a recent surge of interest in using quantum resources to improve the computational complexity of various tasks in machine learning^{3–6}, similar to what one aims to achieve more generally in the fields of quantum communication⁷ and quantum computation^{8–11}. This approach has been successful in training Boltzmann machines, which are simple generative neural networks of a bipartite structure—a set of hidden and a set of visible nodes—where the connectivity is full between the two layers. Edges carry weights and these are adjusted during training. We can view Boltzmann machines as Markov networks with a special topology, in which the largest clique has size two. One method employed for training Boltzmann machines^{12–14} is quantum annealing. It is a global optimization method that relies on actual physical phenomena and it can be used to generate a Gibbs distribution. For all current quantum annealing approaches to Gibbs sampling, restrictions on the topology of the physical hardware remain the main obstacle, which is why the limited clique size of the Boltzmann machines is attractive. An alternative approach of training Boltzmann machines is by using Gibbs state preparation and sampling protocols, which can also exploit the structure of the graph and achieve polynomial improvements in computational complexity relative to its classical analogue¹⁵.

Here, we go beyond the training of Boltzmann machines and consider more general Markov logic networks, keeping the expressiveness of first-order logic and concentrate on inference, rather than training. We analyze

¹ICFO-The Institute of Photonic Sciences, 08860 Castelldefels (Barcelona), Spain. ²University of Borås, 50190 Borås, Sweden. Correspondence and requests for materials should be addressed to P.W. (email: srep2017@peterwittek.com)

First-order formula	Graph characteristic
Number of atoms in formulas	Clique size
Domain size and number of atoms in formula	Total number of nodes
Maximum shared variables	Largest degree

Table 1. Brief summary of how the structure of the first-order formulas in the knowledge base underlying a Markov logic network influences the generated Markov network. Shared variables are variables that appear in more than one formula.

the usefulness of quantum Gibbs sampling methods to outperform MCMC methods. The runtime of quantum Gibbs sampling algorithms is sensitive to both the connectivity structure and the overall number of subsystems. Methods of lifted inference can be used to address these issues.

Probabilistic Inference and Lifting

Markov networks are undirected graphical models that offer a simple perspective on the independence structure of a joint probability distribution of random variables, and the task of probabilistic inference based on this structure¹. Nodes of the network are random variables and edges between nodes imply influence or direct correlation, that is, lack of conditional independence. Instead of conditional probabilities on parent nodes, as in Bayesian networks, Markov networks operate with unnormalized *factors* f_j , that is, functions that map from subsets of the random variables to nonnegative reals. The factors are defined over the cliques of the graph. To obtain a valid joint probability distribution over the random variables from the factors, a partition function normalizes the unnormalized measure, so that the probability distribution takes the form $P(\mathbf{X} = \mathbf{x}) = \frac{1}{Z} \prod_j f_j(\mathbf{x}_j)$, where \mathbf{x}_j are subsets of \mathbf{x} corresponding to the cliques and Z is the partition function. If P is a positive distribution over the random variables $\mathbf{X} := (X_1, \dots, X_n)$, we can associate a Gibbs distribution to the Markov network as $P(\mathbf{X} = \mathbf{x}) = \frac{1}{Z} \exp(\sum_j w_j g_j(\mathbf{x}))$, where the *features* g_j are functions of a subset of the state, and w_j are real weights.

In first-order logic, *constants* are objects over some *domain* (e.g., Alice, Bob, ... in the domain of people), and *variables* range over the set of constants in the domain. A *predicate* is a symbol that represents an attribute of an object (e.g., Smokes), or a relation among objects (e.g., Friends). An *atom* is a predicate applied to a tuple of variables or constants. A *ground atom* only has constants as arguments. These definitions apply to a function free language with finite size domains—technically, this is a strict subset of first-order logic. A *formula* is constructed of atoms, logical connectives, and quantifiers over variables. A *knowledge base* is a set of formulas connected by conjunction. A *world* is an assignment of a truth value to each possible grounding of all atoms in a knowledge base. An essential task in a first-order knowledge base is to check whether a formula is *satisfiable*, that is, there exists at least one world in which it is true.

To relax the rigid true-or-false nature of first-order logic, Markov logic networks (MLNs) introduce a real weight w_j for each formula f_j in a knowledge base². A Markov logic network \mathcal{M} is a set of pairs (f_j, w_j) , representing a probability distribution over worlds as

$$P_{\mathcal{M}}(\omega) = \frac{1}{Z(\mathcal{M})} \exp\left(\sum_j w_j N(f_j, \omega)\right), \quad (1)$$

where $N(f_j, \omega)$ is the number of groundings of f_j that are True in the world ω . An MLN can be thought of as a graph over the set of all possible groundings of the atoms appearing in the knowledge base. The size of this graph is $n \in \mathcal{O}(D^c)$, where D is the maximum domain size, and c is the highest number of atoms in any of the formulas in the knowledge base¹⁶. Groundings are viewed as connected if they can jointly appear in a grounding of some formula of the knowledge base. The ground network thus contains cliques, i.e., fully connected sub-graphs, consisting of grounded atoms that jointly appear in the grounding of some formula. The maximum clique size k is given by the maximum number of atoms per formula. Table 1 summarizes how the structure of the first-order knowledge base influences the characteristics of the generated Markov network.

MLNs belong to the class of methods known as statistical relational learning, which combine relational structures and uncertainty¹⁷. An MLN essentially uses a first-order logic knowledge base as a template to generate a Markov network by grounding out all formulas. An MLN can always be converted to a *normal* MLN, which has the following two properties: (i) there are no constants in any formula; (ii) given two distinct atoms with the same predicate symbol with two variables x and y in the same argument, then the domain of the two variables is identical. In the rest of this work we assume all MLNs to be given in this normal form. We further assume that skolemization is applied to convert existential quantifiers to universal quantifier, which can be done in polynomial time in the size of a formula with no unquantified variables¹⁸.

A main task in graphical models and in MLNs is probabilistic inference. One aspect of it is computing the partition function. The other aspect deals with the problem of assigning probabilities to or finding (the most) likely assignment of variables given evidence, that is, given a fixed assignment for a subset of its variables. This is a hard problem in general: the worst-case complexity of exact probabilistic inference of a graphical model is $\#\mathcal{P}$ -complete and that of approximate inference is $\#\mathcal{P}$ -hard¹.

For some common graphical models with a special topology, efficient exact probabilistic inference methods are known. Examples include belief propagation¹⁹ and the junction tree algorithm²⁰. In other cases, MCMC Gibbs sampling is often used for approximate inference to escape the worst-case complexity of exact inference. MCMC is hereby used to approximately sample from the distribution $P_{\mathcal{M}}(\omega)$ given in (1) or from a suitable conditional probability distribution $P_{\mathcal{M}}(\omega|\mathcal{E})$ conditioned on the evidence \mathcal{E} .

Graphical models often have symmetries that reduce the overall complexity of both exact and approximate inference. For instance, counting belief propagation exploits symmetries for exact inference¹⁶, and orbital Markov chains do the same for approximate inference²¹. Some of these methods have special extensions for MLNs, for instance, one can detect a subset of components in the ground network that would behave identically during belief propagation²². It is worth exploiting the symmetries that emerge from first-order logic and they are best exploited before grounding out, that is, symmetries should be addressed at the propositional level.

Approximate and exact probabilistic inference for first-order probabilistic languages predates MLNs^{23–25}. The core idea is a form of coarse graining by grouping similar variables together. This idea was exploited in lifted first-order probabilistic inference for MLNs²⁶. For hierarchically typed MLNs, one can move from coarse-graining over the highest level in a type hierarchy to more refined types²⁷.

Exploiting symmetries in the presence of evidence must be done with great care. Given evidence, the symmetries can become skewed, as random variables do not appear symmetrically in the formulas of the knowledge base²⁸. In this case, importance sampling helps^{29,30}, which clusters similar network components together given the evidence³¹, and approximates the correct probabilities by an easier probability distribution and an estimated importance or weight of the error.

For most practical applications, either belief propagation or MCMC, augmented with some of the described techniques as appropriate for the problem at hand, is the method of choice for approximate probabilistic inference with MLNs. While often yielding useful results with an effort far smaller than the worst case complexity, they remain very expensive computationally and so more efficient alternatives are desirable.

Quantum Gibbs Sampling

The distribution (1) we would like to sample from can be thought of as the Gibbs distribution of a suitably constructed physical system. According to the rules of statistical mechanics, the probability to find a system in a certain state of configuration when it is in thermal equilibrium follows a Gibbs distribution. The distribution can thus be sampled by preparing a suitable physical system in a thermal equilibrium Gibbs state and then measuring its configuration. This is generally rather easy to do at high temperatures, but cooling to low temperatures typically becomes increasingly difficult. Thereby methods of quantum information processing can offer advantages over classical strategies. They open up fundamentally new ways to of preparing systems approximately in Gibbs states in a well-controlled way.

Going from the abstract definition of the probability distribution in (1) to a physical model can be done in the following way: We can think of $-\sum_j w_j N(f_j, \omega) / \max_j |w_j|$ as the “energy” of a system of n spin 1/2 “particles” in a quantum state $|\omega\rangle$. The states $|\omega\rangle$ are then product state vectors in the Hilbert space $\text{span}(\{|True\rangle, |False\rangle\})^{\otimes n}$ with span the complex linear span. We can think of $\beta \max_j |w_j|$ as the inverse $\beta = 1/(k_B T)$ of the “temperature” T of the system, times the Boltzmann constant k_B (other decompositions of the features are also possible). We can try to find a Hamiltonian H such that we can rewrite the probability distribution from (1) as follows

$$P_{\mathcal{M}}(\omega) = \langle \omega | \exp(-\beta H) / Z(\mathcal{M}) | \omega \rangle. \quad (2)$$

Thereby $\langle \omega |$ is the Hermitian conjugate of the state vector and $Z(\mathcal{M}) := \text{tr}(\exp(-\beta H))$ is the partition function, where \exp is the matrix exponential and tr the matrix trace.

In the concrete case of an MLN, the number of particles $n \in \mathcal{O}(D^c)$ is equal to the number of all possible groundings of the atoms in the knowledge base underlying the MLN. The Hamiltonian H inherits the locality structure of the MLN: it can be written as a sum $H = \sum_l h_l$ of local terms h_l , one for each clique of the MLN. More precisely, for each j the expression $N(f_j, \omega)$ translates to a sum over local terms each acting on one of the cliques produced by grounding out f_j and acting on this clique like $-w_j / \max_j |w_j|$ times the projector on the subspace of assignments to the atoms in the clique for which f_j evaluates to True. The local terms h_l of the Hamiltonian can be constructed from the truth tables of the the f_j and the sum over l in the decomposition of H collects all such terms for the different values of j . Figure 1 illustrates the matching concepts in MLNs and this description.

The number k of subsystems on which each such term acts non-trivially is bounded by the maximum number of atoms per formula and its operator norm is bounded by one $\|h_l\| \leq 1$. Hence (1) is the thermal Gibbs distribution of a system of n spin 1/2 particles with a so-called k -local Hamiltonian H . To prepare the system in a state that is suitable to sample from (1) it is sufficient to reach a high effective temperature if all weights are of moderate magnitude (no assignments are strongly suppressed), but it is necessary to cool to a low temperature if weights have a high magnitude (at least one assignment is strongly suppressed).

Computational Complexity

Quantum Gibbs sampling methods can be used to obtain samples from the Gibbs distributions of the type of systems described in the previous section. Typically these methods consist of two phases:

1. A preparation phase in which a quantum system is prepared in (a state close to) a state encoding information about the Gibbs state or such a state itself; and
2. A measurement phase, in which, by performing measurements on this state, samples from the Gibbs distribution are obtained.

The measurement phase is trivial, consisting only of local measurements and has complexity $\mathcal{O}(n)$. The known quantum methods for Gibbs sampling differ in the kind of resources they require during the preparation, their expected improvement in runtime over classical methods, and the extent to and effort with which their performance for a concrete Gibbs distribution can be predicted.

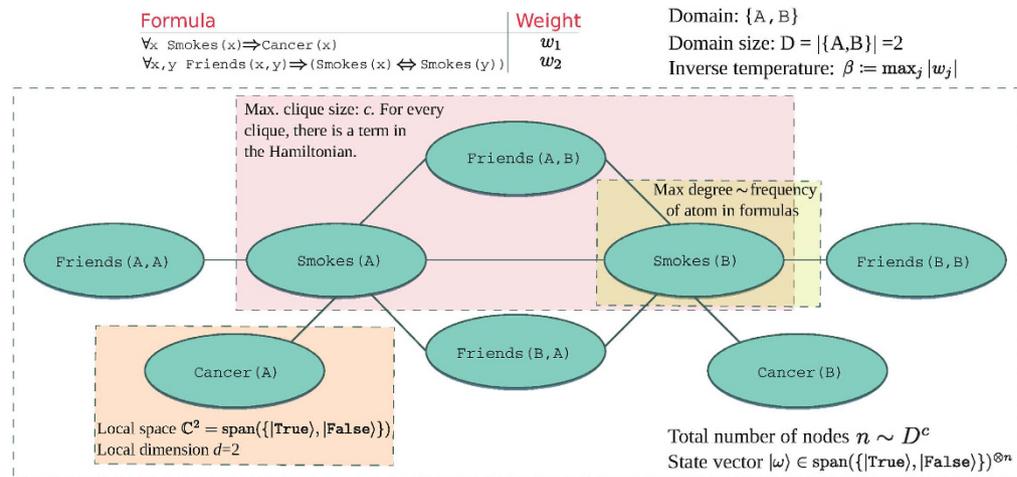


Figure 1. An example of a first-order knowledge base, a matching MLN, and the corresponding concepts of a thermal state and a local Hamiltonian. The knowledge base has only two formulas, and the variables range over a finite domain of two elements, $\{A, B\}$. Grounding out all formulas in all possible way, we obtain the MLN of maximal size (i.e., lifted inference is not used). The maximum of absolute value of the weights w_1 and w_2 defines the inverse temperature β in the thermal state. Since all ground atoms are binary valued, the local space is $\mathbb{C}^2 = \text{span}(|\text{True}\rangle, |\text{False}\rangle)$, and thus the thermal state $|\omega\rangle$ is in $\text{span}(|\text{True}\rangle, |\text{False}\rangle)^{\otimes n}$, where n is the total number of nodes.

The state prepared in the preparation phase is usually either close to a thermal Gibbs state^{32–35} at inverse temperature β of a given Hamiltonian H , or to a so-called pure thermal state^{36,37}, i.e., a pure state whose overlap with any energy eigenstates of H with energy E is proportional to the square root of the Gibbs weight $\exp(-\beta E)$. Recently, an algorithm for the computation based approximate preparation of thermal states of arbitrary k -local Hamiltonians has been proposed in ref. 35. For this algorithm a particularly favorable upper bound on the gate complexity—the scaling of the number of elementary operations in the preparation step—is known. This bound can be expressed in terms of the inverse temperature β , the local dimension d ($d = 2$ for Gibbs states corresponding to MLNs), the number of local terms in H , the gate complexity of time evolution under these terms (or the size of their support) and their strength, as well as the value of the partition function $Z = \text{tr}(\exp(-\beta H))$, and the final distance to the thermal state ϵ .

Proposition 1 Assuming that the maximum size of the support of the local terms of the Hamiltonian H is constant and that for some constant α the number of terms in H is in $\mathcal{O}(n^\alpha)$, the overall complexity of the Gibbs sampling method from ref. 35 is in $\mathcal{O}(\sqrt{d^n \beta / Z} \text{polylog}(\sqrt{d^n \beta / Z} / \epsilon))$.

More generally it is enough if the gate complexity of time evolution under the local terms of H scales at most linear with n . When applying this to the graph structure generated by an MLN, α can be taken to be the maximum number of atoms in any formula, the maximum size of the supports of the local terms of H is equal to the maximum clique size, and the number of terms is the number of cliques in the MLN. So long as the maximum number of atoms in any formula is constant, the above scaling of complexity is achieved. It is important to note that the complexity does not directly depend on the maximal degree of the MLN.

This result improves upon the previously known methods in several respects, but in particular it improves the scaling of the runtime with $1/\epsilon$ and β . In the natural parameters, the problem size n and the precision ϵ , this method yields an exponential improvement over the runtime of classical simulated annealing, which scales like $1/(\delta \epsilon^2)$, where δ is the gap of the Markov process, which in interesting cases typically is in $\mathcal{O}(1/d^n)$ ³⁶. However, the exponential dependence on n remains. The possibility of a logarithmic scaling with $1/\epsilon$ was anticipated in refs 38–40. This scaling is particularly relevant when small probabilities are to be estimated with small relative error.

Following early works⁴¹, several previous methods for quantum Gibbs sampling with improved scaling of complexity had been proposed^{32,33,36–40,42,43}. This in particular concerns the dependence of the runtime on the dimension of the Hilbert space d^n or the inverse gap of a Markov chain $1/\delta$, which was reduced from linear to square root by using techniques such as Szegedy’s quantum walks, the Grover/Long algorithm^{44,45}, phase estimation⁴⁶, or amplitude amplification⁴⁷. Algorithms that speed up the convergence of Markov Chains with quantum techniques^{36–40,42,43} often offer more flexibility than such more specific to the problem of preparing thermal states^{32,33,35}. In cases in which the gap of a Markov chain is much larger than $1/d^n$, they combine their quantum speedup with the advantage inherent in MCMC. However, the interesting cases are usually those in which $1/\delta \approx d^n$ and then both types of algorithms perform essentially equally well. A different method, based on the preparation of microcanonical states, was developed in ref. 34, but has an at least exponential scaling in $\beta \|H\|$.

If the Hamiltonian H has more structure and/or the effective temperature is high, more efficient special purpose procedures are available^{15,48,49}, which however are of limited relevance for inference in MLNs. In addition to that, there exists a quantum generalization of the Metropolis sampling algorithm⁵⁰, that however does not aim at achieving a speedup, but rather works around the sign problem in fermionic systems and makes MCMC techniques available for general local quantum Hamiltonians with non-commuting terms.

Can We Hope for Something Better?

As we have seen, quantum methods reduce the complexity of approximate Gibbs sampling quite drastically. Still, an exponential scaling with the number of all possible groundings of all atoms n remains, and the complexity diverges in the low temperature limit as β goes to infinity. A valid question is: Can we hope that future advances will remedy this? After all, the quantum Gibbs sampling methods presented above are able to do Gibbs sampling from Hamiltonians much more general than those that can arise from MLNs, like ones that have non-commuting terms, for example. Yet, the answer is probably negative. It is highly unlikely that any general purpose quantum algorithm for inference in MLNs exists that is efficient in cases with high weights (i.e., at low temperatures), as this would imply the existence of an efficient algorithm for solving satisfiability problems more general than 3-SAT, which is known to be NP-complete by the Cook–Levin theorem⁵¹. Further, the $\log(1/\varepsilon)$ scaling of complexity is known to be optimal for Hamiltonian simulation⁵² and hence for any Gibbs sampling method based on it (the algorithm of ref. 52, like any other algorithm whose operations are written as linear combinations of unitaries can be thought of⁵³ as a duality quantum computing algorithm^{54–56}; this is always possible as the unitary Pauli Matrices form a complete basis). The situation is different in the high temperature regime, where more efficient Gibbs sampling methods exist⁴⁸.

Computing the Partition Function with First-Order Lifting

The great advantage of working with lifting at the first-order level is that we save potentially exponentially many groundings given the compact representation when we count the models in Eq. (1). There are trivial cases: for instance, when there are no shared variables between the atoms, then there is a closed form to calculate the number of satisfied groundings⁵⁷. Here we follow the outlines of lifted importance sampling^{29,58}, but without reference to an importance or proposal distribution: our aim is to reduce the complexity of the generated Markov network and potentially split it into disconnected graphs when computing the partition function. We run quantum Gibbs sampling on the smaller network and post-process the result with some book-keeping values to return the value of the partition function. Algorithm 1 summarizes the steps. Since the sampling is not based on a proposal distribution, the actual variance will depend on the error term that estimates the accuracy of the quantum Gibbs sampler. We follow the simplification steps from lifted importance sampling to cater to the critical parts of quantum thermal state preparation, but in principle, the sampling part of the algorithm can also use classical MCMC Gibbs sampling. For this reason, Algorithm 1 does not specify what kind of Gibbs sampling protocol we use.

If we have a normal-form network as the input, that is, all domains have size one, we can run the Gibbs sampler and return the value of the partition function.

The first interesting case is if we detect a *decomposer*—this can be done in linear time—that is, a set of logical variables \mathbf{x} such that (i) every atom in \mathcal{M} contains exactly one variable from \mathbf{x} , and (ii) for every predicate R there exists a position such that variables from \mathbf{x} only appear at that position. If we have a decomposer, \mathcal{M} can be simplified to $\mathcal{M}[X/\mathbf{x}]$ that is obtained by substituting all variables in \mathbf{x} by the same constant X in D_x , $x \in \mathbf{x}$, then converting the result to normal form. The partition function is calculated as $Z(\mathcal{M}) = [(Z(\mathcal{M}[X/\mathbf{x}]))]^{|\mathcal{D}_x|}$.

The next structural simplification comes from *isolated variables*—one such variable in a predicate R at position m is exclusive to R in all formulas containing R . Let \mathbf{x} denote all isolated variables of R and \mathbf{y} the rest of the variables, and $\mathbf{Y}_i \in D_y$. We obtain a simplified MLN $\mathcal{M}[R, \mathbf{x}]$ by generating the groundings of $R(x, \mathbf{Y}_i)$ for $i = 1, \dots, |D_y|$, deleting the formulas that evaluate True or False, deleting all groundings of R , and normalizing the result. We get a combinatorial multiplier term to adjust the value of the partition function.

Algorithm 1 Lifted Sampling (LS) of an MLN

Require: A normal MLN \mathcal{M}

Ensure: The value of the partition function $Z(\mathcal{M})$

if \mathcal{M} is fully ground out **then**

 Run Gibbs sampler to obtain $Z(\mathcal{M})$

return $Z(\mathcal{M})$

end if

if there exists a decomposer \mathbf{x} **then**

 Let $x \in \mathbf{x}$ and $X \in D_x$

return $LS(\mathcal{M}[X/\mathbf{x}])^{|\mathcal{D}_x|}$

end if

if there exists an isolated variable x **then**

return $w(R)2^{p(R)} \prod_{i=1}^{|\mathcal{D}_y|} \binom{|\mathcal{D}_x|}{j_i} LS(\mathcal{M}[R, x])$

end if

if exists singleton atom R that does not appear more than once in the same formula **then**

return $\sum_{i=0}^{|\mathcal{D}_x|} \binom{|\mathcal{D}_x|}{i} w(i)2^{p(i)} LS(\mathcal{M}|\bar{R}^i)$

end if

Choose an atom A

return $\sum_{\bar{A} \text{ groundingsof } A} LS(\mathcal{M}|\bar{A})$

The final simplification is known as the *generalized binomial rule*, which relies on singleton atoms that do not appear more than once in the same formula. Given such an atom $R(x)$, we can simplify the MLN as $\mathcal{M}|\bar{R}^i$, where \bar{R}^i is a truth assignment to all groundings of R such that exactly i groundings are set to True. The simplified network is obtained by grounding all $R(x)$ and setting all its groundings to match the assignment given by \bar{R}^i , deleting the formulas that evaluate True or False, deleting all groundings of R , and normalizing the result. We can compute the partition function by $Z(\mathcal{M}|\bar{R}^i) = \sum_{i=0}^{D_x} \binom{D_x}{i} Z(\mathcal{M}|\bar{R}^i) w(i) 2^{p(i)}$, where $w(i)$ is the exponentiated sum of the formulas that evaluate to True, and $p(i)$ is the number of ground atoms that are removed when removing the formulas.

If we cannot find any heuristics, we have to resort to fully grounding out an atom, normalizing the result, and continuing with the remaining expressions.

Probabilistic Inference Given Evidence

If we look at probabilistic inference given evidence, at the level of the quantum protocol, this can be done in at least two ways: First, to the Hamiltonian H one can add some strong local “clamping” terms, effectively forcing some of the assignments to the desired values. This is convenient from an implementation point of view, as it only requires few local changes in the Hamiltonian simulation procedure⁵² underlying the algorithm of ref. 35. However, it can be difficult to quantify the additional error due to the finite clamping strength and adding very strong clamping terms unfavorably affect the runtime of the algorithm. Second, one can construct the local terms h_i not from the full truth tables of the f_j , but instead use reduced truth tables given the evidence, to construct local terms h_i that act non-trivially only on the grounded atoms for which no evidence exists. This can only decrease the maximal weight (i.e., increase the temperature $1/\beta$), decrease the number of terms (in case some of them become completely trivial), and reduce the number of sites n . Gibbs sampling with the new Hamiltonian is hence always at most as computationally costly as with the original Hamiltonian.

We can also use classical heuristics before employing the quantum protocol, as in the algorithm described in the previous section. For first-order lifting methods, the presence of evidence is a problem, as it skews symmetries and potentially leads to a complete grounding out. To avoid this, ref. 31 proposed a distance function on the partially clamped network, and suggested a clustering to find clusters of similar groundings. All groundings in a cluster are replaced by their cluster center, reducing the overall network size to $\mathcal{O}(r^c)$, where r is maximum cluster size, compared to the original $\mathcal{O}(D^c)$. This in turn reduces n in the overall complexity of the quantum Gibbs sampling protocol, as stated in Proposition 1.

Conclusions and Future Work

We hope that by fostering knowledge exchange between communities, for example concerning the typical properties of Gibbs distributions relevant for machine learning, progress towards more realistic and useful quantum algorithms can be made. In summary, we addressed the following aspects of probabilistic inference in MLNs:

- We analyzed the computational complexity of the state-of-the-art quantum Gibbs sampling protocol given the structural properties of MLNs and we argued the theoretical limits of the approach. A term in the computational complexity reduces exponentially, albeit the overall complexity remains exponential in the number of nodes.
- Understanding the impact of the properties of the graph generated by an MLN on the computational complexity of quantum Gibbs sampling, we adapted a classical first-order lifting algorithm to reduce the complexity of the network. The algorithm mirrors lifted importance sampling, but instead of using a proposal distribution, it uses either classical MCMC or quantum Gibbs sampling.
- We studied the effects of evidence on quantum Gibbs sampling.

The protocols we considered rely on a universal quantum computer, which, given the hurdles in implementation, is still mainly of academic interest. We can, however, turn to methods that use current or near future quantum annealing devices, for instance, technology using quantum annealing with manufactured spins^{59,60}. In this technology, the distribution of excited states after annealing follows approximately a Boltzmann distribution¹², albeit one has to pay attention to estimating persistent biases and the effective temperature estimation^{13,14}. This technology was used, for instance, for learning the structure of a Bayesian network⁶¹, but the restricted connectivity between the spins causes difficulties for arbitrary graph structures, in contrast to the methods discussed here. Recent progress allows embedding arbitrary graphs, albeit at a quadratic cost in the number of spins in the worst-case scenario^{62,63}, and there is also a proposal for a quantum annealing architecture with all-to-all connectivity⁶⁴. Given the techniques described in this paper, it would be interesting to see whether we can achieve a scalable implementation with contemporary quantum annealing technologies, since machine learning demonstrations with this paradigm mainly focused on Boltzmann machines so far: MLNs have different topological features than Boltzmann machines, but they also have regularities that might allow an efficient embedding and subsequent inference.

References

1. Koller, D., Friedman, N., Getoor, L. & Taskar, B. Graphical models in a nutshell. In Getoor, L. & Taskar, B. (eds.) *Introduction to Statistical Relational Learning* (MIT Press, 2007).
2. Richardson, M. & Domingos, P. Markov logic networks. *Machine Learning* **62**, 107–136, doi: 10.1007/s10994-006-5833-1 (2006).
3. Schuld, M., Sinayskiy, I. & Petruccione, F. An introduction to quantum machine learning. *Contemporary Physics* **56**, 1–14, doi: 10.1080/00107514.2014.964942 (2014).

4. Wittek, P. *Quantum Machine Learning: What Quantum Computing Means to Data Mining* (Academic Press, New York, NY, USA, 2014).
5. Adcock, J. *et al.* Advances in quantum machine learning. *arXiv:1512.02900* (2015).
6. Biamonte, J. *et al.* Quantum machine learning. *arXiv:1611.09347* (2016).
7. Pfaff, W. *et al.* Unconditional quantum teleportation between distant solid-state quantum bits. *Science* **345**, 532–535, doi: 10.1126/science.1253512 (2014).
8. Tiecke, T. G. *et al.* Nanophotonic quantum phase switch with a single atom. *Nature* **508**, 241–244, doi: 10.1038/nature13188 (2014).
9. Reiserer, A., Kalb, N., Rempe, G. & Ritter, S. A quantum gate between a flying optical photon and a single trapped atom. *Nature* **508**, 237–240, doi: 10.1038/nature13177 (2014).
10. Ren, B.-C., Wang, G.-Y. & Deng, F.-G. Universal hyperparallel hybrid photonic quantum gates with dipole-induced transparency in the weak-coupling regime. *Physical Review A* **91**, 032328, doi: 10.1103/PhysRevA.91.032328 (2015).
11. Wei, H.-R., Deng, F.-G. & Long, G. L. Hyper-parallel Toffoli gate on three-photon system with two degrees of freedom assisted by single-sided optical microcavities. *Optics Express* **24**, 18619, doi: 10.1364/OE.24.018619 (2016).
12. Adachi, S. H. & Henderson, M. P. Application of quantum annealing to training of deep neural networks. *arXiv:1510.06356* (2015).
13. Benedetti, M., Realpe-Gómez, J., Biswas, R. & Perdomo-Ortiz, A. Estimation of effective temperatures in a quantum annealer and its impact in sampling applications: A case study towards deep learning applications. *Physical Review A* **94**, 022308, doi: 10.1103/physrev.94.022308 (2015).
14. Perdomo-Ortiz, A., O’Gorman, B., Fluegemann, J., Biswas, R. & Smelyanskiy, V. N. Determination and correction of persistent biases in quantum annealers. *arXiv:1503.05679* (2015).
15. Wiebe, N., Kapoor, A. & Svore, K. M. Quantum deep learning. *arXiv:1412.3489* (2014).
16. Kersting, K., Ahmadi, B. & Natarajan, S. Counting belief propagation. In *Proceedings of UAI-09, 25th Conference on Uncertainty in Artificial Intelligence*, 277–284 (2009).
17. Getoor, L. & Taskar, B. (eds.) *Introduction to Statistical Relational Learning* (MIT Press, 2007).
18. Van den Broeck, G., Meert, W. & Darwiche, A. Skolemization for weighted first-order model counting. In *Proceedings of KR-14, 14th International Conference on Principles of Knowledge Representation and Reasoning*, 1–10 (2014).
19. Pearl, J. Reverend Bayes on inference engines: A distributed hierarchical approach. In *Proceedings of AAAI-82, 2nd National Conference on Artificial Intelligence*, 133–136 (1982).
20. Lauritzen, S. L. & Spiegelhalter, D. J. Local computations with probabilities on graphical structures and their application to expert systems. *Journal of the Royal Statistical Society. Series B (Methodological)* **50**, 157–224 (1988).
21. Niepert, M. Markov chains on orbits of permutation groups. In *Proceedings of UAI-12, 28th Conference on Uncertainty in Artificial Intelligence*, 624–633 (2012).
22. Singla, P. & Domingos, P. M. Lifted first-order belief propagation. In *Proceedings of AAAI-08, 23rd Conference on Artificial Intelligence*, vol. **8**, 1094–1099 (2008).
23. Pasula, H. & Russell, S. Approximate inference for first-order probabilistic languages. In *Proceedings of ICJAI-01, 17th International Joint Conference on Artificial Intelligence*, 741–748 (2001).
24. Poole, D. First-order probabilistic inference. In *Proceedings of IJCAI-03, 18th International Joint Conference on Artificial Intelligence*, 985–991 (2003).
25. De Salvo Braz, R., Amir, E. & Roth, D. Lifted first-order probabilistic inference. In *Proceedings of IJCAI-05, 19th International Joint Conference on Artificial Intelligence*, 1319–1325 (2005).
26. Kersting, K. Lifted probabilistic inference. In *Proceedings of ECAI-12, 20th European Conference on Artificial Intelligence*, 33–38 (2012).
27. Kiddon, C. & Domingos, P. Coarse-to-fine inference and learning for first-order probabilistic models. In *Proceedings of AAAI-11, 25th AAAI Conference on Artificial Intelligence*, 1049–1056 (2011).
28. Ahmadi, B., Kersting, K., Mladenov, M. & Natarajan, S. Exploiting symmetries for scaling loopy belief propagation and relational training. *Machine Learning* **92**, 91–132, doi: 10.1007/s10994-013-5385-0 (2013).
29. Gogate, V., Jha, A. & Venugopal, D. Advances in lifted importance sampling. In *Proceedings of AAAI-12, 26th AAAI Conference on Artificial Intelligence*, 1910–1916 (2012).
30. Venugopal, D. & Gogate, V. Scaling-up importance sampling for Markov logic networks. In *Advances in Neural Information Processing Systems*, vol. **27**, 2978–2986 (2014).
31. Venugopal, D. & Gogate, V. Evidence-based clustering for scalable inference in Markov logic. In *Machine Learning and Knowledge Discovery in Databases*, 258–273 (2014).
32. Poulin, D. & Wocjan, P. Sampling from the thermal quantum Gibbs state and evaluating partition functions with a quantum computer. *Physical Review Letters* **103**, 220502, doi: 10.1103/PhysRevLett.103.220502 (2009).
33. Chiang, C.-F. & Wocjan, P. Quantum algorithm for preparing thermal Gibbs states - detailed analysis. In *Quantum Cryptography and Computing*, vol. 26 of *NATO Science for Peace and Security Series - D: Information and Communication Security*, 138–147 (2010).
34. Riera, A., Gogolin, C. & Eisert, J. Thermalization in nature and on a quantum computer. *Physical Review Letters* **108**, 080402, doi: 10.1103/PhysRevLett.108.080402 (2012).
35. Chowdhury, A. N. & Somma, R. D. Quantum algorithms for Gibbs sampling and hitting-time estimation. *arXiv:1603.02940* (2016).
36. Wocjan, P., Chiang, C.-F., Nagaj, D. & Abeyesinghe, A. Quantum algorithm for approximating partition functions. *Physical Review A* **80**, 022340, doi: 10.1103/PhysRevA.80.022340 (2009).
37. Boixo, S., Knill, E. & Somma, R. D. Quantum state preparation by phase randomization. *Quantum Information & Computation* **9**, 833–855 (2009).
38. Richter, P. C. Quantum speedup of classical mixing processes. *Physical Review A* **76**, 042306, doi: 10.1103/PhysRevA.76.042306 (2007).
39. Somma, R. D., Boixo, S., Barnum, H. & Knill, E. Quantum simulations of classical annealing processes. *Physical Review Letters* **101**, 130504, doi: 10.1103/PhysRevLett.101.130504 (2008).
40. Tucci, R. R. Quantum Gibbs sampling using Szegedy operators. *arXiv:0910.1647* (2009).
41. Terhal, B. M. & DiVincenzo, D. P. Problem of equilibration and the computation of correlation functions on a quantum computer. *Physical Review A* **61**, 022301, doi: 10.1103/PhysRevA.61.022301 (2000).
42. Somma, R. D., Boixo, S. & Barnum, H. Quantum simulated annealing. *arXiv:0712.1008* (2007).
43. Wocjan, P. & Abeyesinghe, A. Speedup via quantum sampling. *Physical Review A* **78**, 1–8, doi: 10.1103/PhysRevA.78.042336 (2008).
44. Grover, L. K. A fast quantum mechanical algorithm for database search. In *Proceedings of STOC-96, 28th Annual Symposium on Theory of Computing*, 212–219 (1996).
45. Long, G. L. Grover algorithm with zero theoretical failure rate. *Physical Review A* **64**, 22307, doi: 10.1103/PhysRevA.64.022307 (2001).
46. Luis, A. & Peřina, J. Optimum phase-shift estimation and the quantum description of the phase difference. *Physical Review A* **54**, 4564–4570, doi: 10.1103/PhysRevA.54.4564 (1996).
47. Brassard, G., Hoyer, P., Mosca, M. & Tapp, A. Quantum amplitude amplification and estimation. *arXiv:quant-ph/0005055* (2000).
48. Kastoryano, M. J. & Brandão, F. G. S. L. Quantum Gibbs samplers: The commuting case. *Communications in Mathematical Physics* **344**, 915–957, doi: 10.1007/s00220-016-2641-8 (2016).

49. Bilgin, E. & Boixo, S. Preparing thermal states of quantum systems by dimension reduction. *Physical Review Letters* **105**, 170405, doi: 10.1103/PhysRevLett.105.170405 (2010).
50. Temme, K., Osborne, T. J., Vollbrecht, K. G., Poulin, D. & Verstraete, F. Quantum Metropolis sampling. *Nature* **471**, 87–90, doi: 10.1038/nature09770 (2011).
51. Cook, S. A. The complexity of theorem-proving procedures. In *Proceedings of STOC-71, 3rd Annual Symposium on Theory of Computing*, 151–158 (1971).
52. Berry, D. W., Childs, A. M., Cleve, R., Kothari, R. & Somma, R. D. Simulating Hamiltonian dynamics with a truncated Taylor series. *Physical Review Letters* **114**, 090502, doi: 10.1103/PhysRevLett.114.090502 (2015).
53. Wei, S. J. & Long, G. L. Duality quantum computer and the efficient quantum simulations. *Quantum Information Processing* **15**, 1189–1212, doi: 10.1007/s11128-016-1263-6 (2016).
54. Long, G. L. General quantum interference principle and duality computer. *Communications in Theoretical Physics* **45**, 825–844, doi: 10.1088/0253-6102/45/5/013 (2006).
55. Long, G. L. Duality quantum computing and duality quantum information processing. *International Journal of Theoretical Physics* **50**, 1305–1318, doi: 10.1007/s10773-010-0603-z (2011).
56. Wei, S.-J., Ruan, D. & Long, G.-L. Duality quantum algorithm efficiently simulates open quantum systems. *Scientific Reports* **6**, 30727, doi: 10.1038/srep30727 (2016).
57. Sarkhel, S., Venugopal, D., Singla, P. & Gogate, V. Lifted MAP inference for Markov logic networks. In *Proceedings of AISTATS-14, 17th International Conference on Artificial Intelligence and Statistics*, 859–867 (2014).
58. Gogate, V. & Domingos, P. Probabilistic theorem proving. In *Proceedings of UAI-11, 27th Conference on Uncertainty in Artificial Intelligence*, 256–265 (2011).
59. Johnson, M. W. *et al.* Quantum annealing with manufactured spins. *Nature* **473**, 194–198, doi: 10.1038/nature10012 (2011).
60. Boixo, S. *et al.* Evidence for quantum annealing with more than one hundred qubits. *Nature Physics* **10**, 218–224, doi: 10.1038/nphys2900 (2014).
61. O’Gorman, B. A., Perdomo-Ortiz, A., Babbush, R., Aspuru-Guzik, A. & Smelyanskiy, V. Bayesian network structure learning using quantum annealing. *European Physics Journal Special Topics* **224**, 163–188, doi: 10.1140/epjst/e2015-02349-9 (2015).
62. Zaribafiyani, A., Marchand, D. J. J. & Rezaei, S. S. C. Systematic and deterministic graph-minor embedding for Cartesian products of graphs. *arXiv:1602.04274* (2016).
63. Benedetti, M., Realpe-Gómez, J., Biswas, R. & Perdomo-Ortiz, A. Quantum-assisted learning of graphical models with arbitrary pairwise connectivity. *arXiv:1609.02542* (2016).
64. Lechner, W., Hauke, P. & Zoller, P. A quantum annealing architecture with all-to-all connectivity from local interactions. *Science Advances* **1**, e1500838–e1500838, doi: 10.1126/sciadv.1500838 (2015).

Acknowledgements

We would like to thank the anonymous referee 1 for suggesting to us refs 45, and 53–56 and the anonymous referee 2 for suggesting to us refs 7–11. P.W. and C.G. acknowledge financial support from the European Research Council (CoG QITBOX and AdG OSYRIS), the Axa Chair in Quantum Information Science, Spanish MINECO (FOQUS FIS2013-46768, QIBEQI FIS2016-80773-P and Severo Ochoa Grant No. SEV-2015-0522), Fundació Privada Cellex, and Generalitat de Catalunya (Grant No. SGR 874 and 875, and CERCA Programme). C.G. acknowledges support by the European Union’s Marie Skłodowska-Curie Individual Fellowships (IF-EF) programme under GA: 700140.

Author Contributions

P.W. and C.G. have contributed equally and wrote the main manuscript text. P.W. prepared Fig. 1. All authors reviewed the manuscript.

Additional Information

Competing Interests: The authors declare no competing financial interests.

How to cite this article: Wittek, P. and Gogolin, C. Quantum Enhanced Inference in Markov Logic Networks. *Sci. Rep.* **7**, 45672; doi: 10.1038/srep45672 (2017).

Publisher’s note: Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



This work is licensed under a Creative Commons Attribution 4.0 International License. The images or other third party material in this article are included in the article’s Creative Commons license, unless indicated otherwise in the credit line; if the material is not included under the Creative Commons license, users will need to obtain permission from the license holder to reproduce the material. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>

© The Author(s) 2017