

Lifted Inference with Tree Axioms

Timothy van Bremen¹, Ondřej Kuželka²

¹KU Leuven, Belgium

²Czech Technical University in Prague, Czech Republic

timothy.vanbremen@cs.kuleuven.be, ondrej.kuzelka@fel.cvut.cz

Abstract

We consider the problem of weighted first-order model counting (WFOMC): given a first-order sentence ϕ and domain size $n \in \mathbb{N}$, determine the weighted sum of models of ϕ over the domain $\{1, \dots, n\}$. Past work has shown that any sentence using at most two logical variables admits an algorithm for WFOMC that runs in time polynomial in the given domain size (Van den Broeck 2011; Van den Broeck, Meert, and Darwiche 2014). In this paper, we extend this result to any two-variable sentence ϕ with the addition of a *tree axiom*, stating that some distinguished binary relation in ϕ forms a tree in the graph-theoretic sense.

1 Introduction

Given a first-order sentence ϕ and domain size $n \in \mathbb{N}$, the *first-order model counting* (FOMC) task is to determine the number of models of ϕ over the domain $\{1, \dots, n\}$. Its weighted variant, *weighted first-order model counting* (WFOMC), assigns each predicate in the signature of ϕ to real-valued positive and negative weights, and asks for the weighted sum of all models of ϕ .

As a natural generalization of its propositional counterpart *weighted model counting* (WMC), inference and learning problems in statistical-relational models such as Markov logic networks (Richardson and Domingos 2006) and probabilistic logic programs (Fierens et al. 2015) are reducible to WFOMC (Van den Broeck, Meert, and Darwiche 2014). In addition, several problems in enumerative combinatorics enjoy a straightforward reduction to FOMC, as long as the characterizing properties of the structure in question can be described by a first-order logic sentence. For example, setting $\phi = \forall x \exists y (E(x, y) \vee E(y, x)) \wedge \forall x \neg E(x, x)$ models directed graphs with no isolated vertices, and the FOMC of ϕ on n gives the number of such labelled graphs with n nodes. A common thread in all of these applications is that one typically has some fixed sentence ϕ , and wants to understand how the complexity of computing the (weighted) first-order model count grows with the domain size n , specified as a unary input. This corresponds broadly to the notion of *data complexity* from the databases and finite model theory literature (Vardi 1982).

In a seminal result, Van den Broeck (2011) and Van den Broeck, Meert, and Darwiche (2014) proved that the data complexity of WFOMC for any sentence in the two-variable

fragment of first-order logic, FO^2 , is in the complexity class FP. This proof, initially split across the two papers, was later consolidated by Beame et al. (2015). The authors of the latter paper also exhibited a three-variable sentence for which the data complexity of WFOMC is $\#\text{P}_1$ -complete, thereby giving a complete characterization of the data complexity for WFOMC in terms of the number of variables in the sentence.

In this paper, we extend the FP upper bound result beyond FO^2 to allow for a feature that is inexpressible in first-order logic (Libkin 2004, Chapter 3.6): a *tree axiom*, stating that some distinguished binary relation in the sentence forms a tree in the corresponding graph of any model. To do this, we use a tool from the world of graph theory: the weighted version of Kirchhoff’s matrix tree theorem (Chaiken and Kleitman 1978), which, roughly speaking, states that the weighted sum of all spanning trees of a graph can be quickly computed as the determinant of any minor of the graph’s Laplacian matrix.

This extension to tree axioms allows one to model several notable problems in enumerative combinatorics on trees, especially when combined with recent work allowing for *cardinality constraints* and *counting quantifiers* (Kuzelka 2021): for example, we may now efficiently answer questions like “how many labelled rooted trees on n nodes have exactly (resp. at most, at least) k leaves”? We show some experimental results of our approach on these types of questions near the end of this paper.

There are also interesting implications for probabilistic inference. Inference algorithms that run efficiently in the input domain size underpin the research field of *lifted inference*, which aims to exploit high-level structure for faster inference in graphical models. Using the results presented here, we are now able to efficiently deal with expressions that state, for example, that possible worlds that are “tree-like” are more likely to occur in the context of a Markov logic network. We later explore some preliminary applications of our results to Markov logic networks in this spirit. We also note that many real-world constructs are naturally modelled as trees: for example, *paths* are trees where every node has at most one child. It is not difficult to imagine Markov logic networks in which we are interested in computing the probabilities of paths along some route, subject to certain constraints modelled by the network.

2 Related Work

This work follows a long line of research in lifted inference on statistical-relational models. Many such papers have focused on operating on first-order versions of graphical models directly (Poole 2003; Milch et al. 2008; Taghipour et al. 2013; Braun and Möller 2016). On the other hand, a separate line of work has introduced WFOMC as a natural extension of its propositional counterpart WMC, and has shown how inference in the first-order models mentioned above can be reduced to WFOMC (Van den Broeck et al. 2011). Within this framework, an effort has been undertaken to understand which fragments of first-order logic allow for tractable inference. The most interesting results for us in the context of this paper are the tractability of two-variable logic (Van den Broeck 2011; Van den Broeck, Meert, and Darwiche 2014) along with its extension to functionality constraints (Kuusisto and Lutz 2018) and more recently counting quantifiers (Kuzelka 2021), as well as the fragments S^2FO^2 and S^2RU (Kazemi et al. 2016).

Notably, all existing work we are aware of in the lifted inference literature has focused on identifying tractable *fragments* of first-order logic in the strict sense, rather than considering extensions that allow for representing properties beyond what is expressible in first-order logic. In this paper, we take an important first step past this barrier by allowing for the expression of tree axioms.

However, we do note that languages going “beyond” the expressiveness of first-order logic have been investigated in other contexts. Since two-variable logic over finite trees is equivalent in expressive power to navigational XPath, a popular query language for XML documents, the satisfiability problem on this language has received attention in the logic literature (Benaim et al. 2016; Charatonik and Witkowski 2013). Similarly, in the world of probabilistic databases (closely related to *asymmetric* WFOMC, briefly discussed in the following section), results have been shown for more expressive languages. For instance, query evaluation for the class of *homomorphism-closed queries* on probabilistic graphs has been shown to admit a dichotomy between polytime and #P-hardness depending on the query in question (Amarilli and Ceylan 2020; Dalvi and Suciu 2012). In particular, this class encompasses negation-free disjunctive Datalog as well as regular path queries. Extending existing (either lower or upper bound) results for WFOMC to richer languages like these remains an open problem.

3 Background

In this section, we give some background on first-order logic and the WFOMC problem, as well as a brief review of the graph theory needed for the paper.

3.1 First-order Logic

Preliminaries We deal with the function-free, finite domain fragment of first-order logic. An *atom* of arity k takes the form $P(x_1, \dots, x_k)$, where P/k comes from a vocabulary of *predicates* (also called *relations*), and each argument

x_i is a logical variable from a vocabulary of variables. A *literal* is an atom or its negation. A *formula* is formed by connecting one or more literals together using conjunction or disjunction. A formula may optionally be surrounded by one or more quantifiers of the form $\exists x$ or $\forall x$, where x is a logical variable. A logical variable in a formula is said to be *free* if it is not bound by any quantifier. A formula with no free variables is called a *sentence*. A structure \mathcal{A} interprets each predicate in the formula over a domain; we denote the restriction of \mathcal{A} to a predicate R_i with \mathcal{A}_{R_i} . We follow the usual semantics of first-order logic for deciding whether a structure satisfies (is a model of) a formula.

WFOMC We are now ready to formally define first-order model counting along with its weighted counterpart.

Definition 1. *The first-order model count (FOMC) of a sentence ϕ over a domain of size n is defined as:*

$$FOMC(\phi, n) = |\text{models}_n(\phi)|$$

where $\text{models}_n(\phi)$ denotes the set of all models of ϕ over the domain $\Delta = \{1, \dots, n\}$.

The *weighted* first-order model count of a sentence is defined in terms of weightings on the predicates. Note that since these weightings are defined on the predicate level, all groundings of the same predicate get the same weight. This corresponds to *symmetric* WFOMC. The case in which every possible ground atom in an interpretation is instead annotated with an individual weight corresponds to *asymmetric* WFOMC and is commonly associated with probabilistic databases, which is beyond the scope of this paper. We instead refer the interested reader to (Gribkoff, Van den Broeck, and Suciu 2014) for an overview of that case.

Definition 2. *Denote the set of predicates appearing in a sentence ϕ by P_ϕ . A weighting on ϕ is a pair of mappings $w : P_\phi \rightarrow \mathbb{R}$ and $\bar{w} : P_\phi \rightarrow \mathbb{R}$.*

Definition 3. *Let (w, \bar{w}) be a weighting on a sentence ϕ . The weighted first-order model count (WFOMC) of ϕ over a domain of size n under (w, \bar{w}) is:*

$$WFOMC(\phi, n, w, \bar{w}) = \sum_{\mu \in \text{models}_n(\phi)} \prod_{L \in \mu_T} w(\text{pred}(L)) \cdot \prod_{L \in \mu_F} \bar{w}(\text{pred}(L))$$

where μ_T denotes the set of true ground atoms in the model μ , and μ_F the false ground atoms. The notation $\text{pred}(L)$ maps an atom L to its corresponding predicate name.

Cardinality Constraints Very recent work has further generalized the concept of WFOMC to *WFOMC under cardinality constraints* (Kuzelka 2021). Although this extended notion of WFOMC is not strictly needed for our core results on trees in the context of two-variable logic, it will be helpful when modelling more complex properties which often require the more expressive fragment C^2 , whose details are provided later in the paper. In particular, existing lifted inference methods for C^2 work by reducing WFOMC on a C^2 sentence to WFOMC on an FO^2 sentence under cardinality constraints.

Definition 4. Let (w, \bar{w}) be a weighting on a sentence ϕ , and let C denote a set of constraints of the form $\{|R_1| \bowtie c_1, \dots, |R_k| \bowtie c_k\}$ for predicates R_i in ϕ , $\bowtie \in \{<, >, =, \geq, \leq\}$, and $c_i \in \mathbb{N}$. The WFOMC of ϕ over a domain of size n under (w, \bar{w}) and constraints C is:

$$\text{WFOMC}(\phi, n, w, \bar{w}, C) = \sum_{\mu \in \text{models}_n(\phi), \text{sat}(\mu, C)} \prod_{L \in \mu_T} w(\text{pred}(L)) \cdot \prod_{L \in \mu_F} \bar{w}(\text{pred}(L))$$

where $\text{sat}(\mu, C)$ is true if and only if the number of true atoms in the model μ for each predicate in ϕ satisfies the relevant constraints in C , and other notation is as before.

Modularity In this paper, we will also be performing reductions between WFOMC problems. We will often want such reductions to be *modular*, a term whose definition we follow from (Van den Broeck, Meert, and Darwiche 2014).

Property 1 (Modularity). A reduction is said to be modular if, for every tuple (ϕ, w, \bar{w}) of sentence and weights passed to the reduction with output (ϕ', w', \bar{w}') , we have:

$$\text{WFOMC}(\phi \wedge \gamma, n, w, \bar{w}) = \text{WFOMC}(\phi' \wedge \gamma, n, w', \bar{w}')$$

for every domain size n and sentence γ .

Intuitively, this means that any reduction on ϕ is not invalidated by replacing γ with a different sentence.

Expressiveness Finally, we conclude with a short remark on the expressiveness of first-order logic that is useful to bear in mind when considering the results presented later in the paper.

Remark 1. Throughout this paper, we deal with finite domains; this means that any property we can think of can be expressed by grounding it out. This includes tree axioms: we can enumerate all of the possible trees that can be formed on elements of the domain. However, this induces a blow-up in the formula length by a factor exponential in the domain size, thus rendering it of limited practical utility.

3.2 Data Complexity of WFOMC

Given a fragment \mathcal{F} of first-order logic, we may consider its *data complexity* for WFOMC: fixing the input sentence as some $\phi \in \mathcal{F}$ and weights (w, \bar{w}) , what is the complexity of computing $\text{WFOMC}(\phi, n, w, \bar{w})$ when treating the domain size n as a unary input?

Data Complexity for FO^2 In Appendix C of (Beame et al. 2015), the authors show that the data complexity of WFOMC for any sentence in the syntactic fragment of first-order logic limited to two variables is in FP. We reproduce a sketch of this proof below. The term used in the statistical-relational learning literature for logical fragments with FP data complexity is *domain-liftability* (Van den Broeck 2011). We will also use this term throughout the paper, and occasionally slightly abuse notation by referring to individual sentences (rather than entire fragments) as domain-liftable as well.

Theorem 1 (Beame et al. 2015). *The fragment of first-order logic limited to two variables, FO^2 , is domain-liftable.*

Proof sketch. Suppose that we wish to compute $\text{WFOMC}(\phi, n, w, \bar{w})$ for some input sentence $\phi \in \text{FO}^2$, domain size $n \in \mathbb{N}$, and weights (w, \bar{w}) . Begin by applying the reduction in (Grädel, Kolaitis, and Vardi 1997) and eliminating existential quantifiers as shown in (Van den Broeck, Meert, and Darwiche 2014) to get a universally quantified sentence $\phi = \forall x \forall y : \psi(x, y)$ such that all atoms in ψ are either unary or binary.

Take the 2^u cells (also called *l-types*) formed as maximal consistent conjunctions of literals over the u predicates in ϕ containing only the variable x , and denote them C_1, \dots, C_{2^u} . Now, consider the possible partitions of $[n]$ into 2^u disjoint sets. Each of these partitions can be thought of as representing a series of assignments of subsets of the n domain elements to each of the cells. Then the models of ϕ over the domain $\Delta = [n]$ are precisely the models of the following sentence:

$$\eta = \bigwedge_{i, j \in [2^u], i < j} \forall x : S_i \forall y : S_j (\psi(x, y) \wedge \psi(y, x)) \wedge \bigwedge_{k \in [2^u]} \forall x : S_k \forall y : S_k \psi(x, y)$$

where S_i denotes the elements of $[n]$ assigned to the cell C_i , and the notation $\forall x : S_i$ denotes universal quantification limited to the set S_i . Since we know the truth values of the unary and reflexive binary atoms given by each cell assignment C_i , we may simplify the body of each conjunct by replacing every unary and reflexive binary atom with true or false as appropriate. Write $\psi_i(x, y)$ for the simplified version of $\psi(x, y)$ when both x and y are restricted to taking values from S_i , and $\psi_{ij}(x, y)$ for the simplified version of $\psi(x, y) \wedge \psi(y, x)$ when x and y are restricted to taking values from S_i and S_j respectively. We then have:

$$\eta = \bigwedge_{i, j \in [2^u], i < j} \forall x : S_i \forall y : S_j (\psi_{ij}(x, y)) \wedge \bigwedge_{k \in [2^u]} \forall x : S_k \forall y : S_k \psi_k(x, y)$$

Observe that each conjunct in the formula above is independent (that is, they do not share any propositional variables when grounded out). Denote $r_{ij} = \text{WMC}(\psi_{ij}(a, b), w, \bar{w})$, $s_k = \text{WMC}(\psi_k(a, b) \wedge \psi_k(b, a), w, \bar{w})$, and $w_k = \text{WMC}(\psi_k(c, c), w, \bar{w})$. Summing across the different possible configurations of cell cardinalities, and multiplying by a multinomial coefficient to account for the different possible selections of domain elements for a given configuration, we get:

$$\text{WFOMC}(\phi, n, w, \bar{w}) = \sum_{n_1 + \dots + n_{2^u} = n} \binom{n}{n_1, \dots, n_{2^u}} \prod_{i, j \in [2^u], i < j} r_{ij}^{n_i n_j} \prod_{i \in [2^u]} s_i^{n_i(n_i-1)/2} w_i^{n_i} \quad (1)$$

Clearly, evaluating this equation can be done in time polynomial in the domain size, and so we have that FO^2 is domain-liftable. \square

Data Complexity for \mathbf{C}^2 Unfortunately, in many applications, the expressiveness of \mathbf{FO}^2 as a modelling language is too limited. Recent work has shown that two-variable logic extended with *counting quantifiers* of the form $\exists^{\bowtie n} x$ for $\bowtie \in \{<, >, =, \geq, \leq\}$ (called \mathbf{C}^2), is also domain-liftable (Kuzelka 2021). Note that this fragment is strictly more expressive than \mathbf{FO}^2 : counting quantifiers allow one to express concepts like “each vertex has at most two outgoing edges”, whereas without them we could only express that each vertex has at least one outgoing edge, or no outgoing edges at all. To show the domain-liftability of \mathbf{C}^2 , we first need to understand how to efficiently deal with cardinality constraints when computing the WFOMC.

Theorem 2 (Kuzelka 2021, slightly reformulated). *Let ϕ be a first-order logic sentence and C be a set of cardinality constraints (as in Definition 4). If there is a domain-lifted oracle for $\text{WFOMC}(\phi, n, w, \bar{w})$ then $\text{WFOMC}(\phi, w, n, \bar{w}, C)$ can be computed in polynomial time using a polynomial number of queries to this oracle.*

Proof sketch. We observe that the WFOMC of ϕ can be written as a polynomial in the positive and negative weights of the predicates, given respectively by the weightings w, \bar{w} . Specifically, suppose that we have predicates $P_1/k_1, P_2/k_2, \dots, P_m/k_m$. Then we can write

$$\text{WFOMC}(\phi, n, w, \bar{w}) = \sum_{\mathbf{n} \in \mathcal{D}} A(\mathbf{n}) \mathbf{w}^{\mathbf{n}} \bar{\mathbf{w}}^{\bar{\mathbf{n}}}$$

where we used the notation:

$$\begin{aligned} \mathbf{n} &\triangleq [n_1, \dots, n_m], \\ \bar{\mathbf{n}} &\triangleq [n^{k_1} - n_1, n^{k_2} - n_2, \dots, n^{k_m} - n_m], \\ \mathbf{w}^{\mathbf{n}} &\triangleq \prod_{i=1}^m w(P_i)^{n_i}, \quad \bar{\mathbf{w}}^{\bar{\mathbf{n}}} \triangleq \prod_{i=1}^m \bar{w}(P_i)^{n^{k_i} - n_i}, \\ \mathcal{D} &\triangleq \{0, 1, 2, \dots, n^{k_1}\} \times \dots \times \{0, 1, 2, \dots, n^{k_m}\}. \end{aligned}$$

Here, $A(\mathbf{n})$ then must be the number of possible worlds in which the predicate P_1 has cardinality exactly n_1 , the predicate P_2 has cardinality exactly n_2 , etc. It follows that if we are given access to the oracle for $\text{WFOMC}(\phi, n, w, \bar{w})$ then we can use Lagrange interpolation to extract $A(\mathbf{n})$ for any \mathbf{n} , which then allows us to compute the WFOMC of ϕ with the cardinality constraints C in a straightforward way. For details, we refer to Proposition 5 in (Kuzelka 2021). \square

Importantly, the theorem above does not require the sentence ϕ to be from \mathbf{FO}^2 . It can be applied to any domain-liftable sentence. Finally, combining Theorem 1 and Theorem 2 yields the following result.

Theorem 3 (Kuzelka 2021). *The fragment of first-order logic limited to two variables with counting quantifiers, \mathbf{C}^2 , is domain-liftable.*

Proof sketch. The principle observation behind this result is that one can transform $\phi \in \mathbf{C}^2$ into a new sentence $\phi' \in \mathbf{FO}^2$, weights (w', \bar{w}') , and a set of cardinality constraints C on the predicates in ϕ' such that:

$$\text{WFOMC}(\phi, n, w, \bar{w}) = K_{\phi, n} \cdot \text{WFOMC}(\phi', n, w', \bar{w}', C)$$

for some constant $K_{\phi, n} \in \mathbb{R}$, whose value can be obtained in time polynomial in the domain size. In addition, this transformation has the helpful property that it leaves conjuncts of ϕ containing no counting quantifiers untouched. Since we know that \mathbf{FO}^2 is domain-liftable by Theorem 1, and we can handle cardinality constraints by Theorem 2, we are done. We again refer the reader to (Kuzelka 2021) for the details. \square

3.3 Graph Theory

We will also need some tools from graph theory.

Preliminaries Throughout this section, we assume that graphs are undirected and contain no self-loops. A *tree* is a connected acyclic graph. A *spanning tree* of a graph is a subgraph that is a tree and contains all vertices of the original graph. In the context of this paper, we will uniquely characterise a (weighted) graph by its symmetric *weighted adjacency matrix* W , whose element at position (a, b) denotes the weight on the edge $\{a, b\}$. Unconnected edges are given zero weight. Denote $\text{ver}(W)$ as the set of vertices in W , and $\text{edg}(W)$ as the set of edges with non-zero weight (i.e., $\text{edg}(W) = \{\{i, j\} \mid W_{(i,j)} \neq 0\}$). We associate with W a *weighted degree matrix* $D(W)$:

$$(D(W))_{(a,b)} = \begin{cases} \sum_{k=1}^{|\text{ver}(W)|} (W)_{(a,k)} & \text{if } a = b \\ 0 & \text{otherwise} \end{cases}$$

Kirchhoff’s Theorem Using the definitions above, we can define the *Laplacian matrix* of a graph.

Definition 5 (Laplacian matrix). *The Laplacian $L(W)$ of a weighted graph W is defined as:*

$$L(W) = D(W) - W$$

The Laplacian matrix is critical to the statement of Kirchhoff’s theorem below.

Theorem 4 (Kirchhoff’s matrix tree theorem). *Let W be a weighted graph, and let $L(W)$ denote the Laplacian of W . Further, let $[L(W)]_{i,i}$ denote the minor of $L(W)$ obtained by deleting row i and column i , for some $i \leq |\text{ver}(W)|$. Then the weighted sum of the spanning trees of W is the determinant of any of the minors of its Laplacian $L(W)$; that is:*

$$\sum_{P \in S(W)} \prod_{\{a,b\} \in P} (W)_{(a,b)} = \det([L(W)]_{i,i})$$

for all $i \leq |\text{ver}(W)|$, where $S(W)$ denotes the set of (weighted) spanning trees of W .

Kirchhoff’s theorem gives us a way to quickly sum the weighted spanning trees of a graph in time polynomial in the graph size. However, later in the paper we will also need a way to sum the weighted spanning trees subject to the constraint that certain edges must be present in those trees.

Definition 6. *Let*

$$\text{TreeSum}(W, \mathcal{F}) = \sum_{\substack{T \in S(W) \\ \mathcal{F} \subseteq T}} \prod_{\{a,b\} \in T} W_{(a,b)}$$

denote the weighted sum of spanning trees of W that contain all edges in $\mathcal{F} \subseteq \text{edg}(W)$.

How can we compute $\text{TreeSum}(W, \mathcal{F})$ in a manner whose runtime is still polynomial in the graph size? The opposite case—that is, imposing the constraint that certain edges are *not* present—is easy: simply set the weight of these edges to zero in W and apply Theorem 4. On the other hand, for the “must-be-present” case, we can contract edges of the graph while keeping track of their weights. Suppose we are given the set of edges $\mathcal{F} \subseteq \text{edg}(W)$ whose presence we wish to force in the spanning trees of the graph W . We may assume without loss of generality that \mathcal{F} is acyclic; otherwise we know immediately $\text{TreeSum}(W, \mathcal{F}) = 0$. With this assumption in mind, pick an edge $e_1 = \{i, j\} \in \mathcal{F}$. Then we may update W by deleting the edge $\{i, j\}$ and the vertex j , and redirecting any edge that originally connected j to a third vertex k to i . If i and k are already connected, we simply add the weight of the edge $\{j, k\}$ to the existing edge $\{i, k\}$.¹ Repeating this process for every edge $e_i \in \mathcal{F}$, we eventually end up with a new graph W' in which all of the relevant edges have been contracted. We may then compute $\text{TreeSum}(W, \mathcal{F})$ as:

$$\text{TreeSum}(W, \mathcal{F}) = w(e_1) \dots w(e_n) \det([L(W')]_{1,1})$$

where $w(e_i)$ denotes the original weight of the edge e_i in W .

4 Approach

We now come to the primary results of this paper. In this section, we will show how to efficiently compute the WFOMC of a \mathbf{C}^2 sentence containing either of the following axioms:

- $\text{Tree}(R, \text{Leaf})$, expressing that the (symmetric, antireflexive) binary relation R forms a tree with leaves defined by the unary relation Leaf .
- $\text{DirectedRootedTree}(\text{Root}, E, \text{Leaf})$, expressing that the binary relation E represents directed edges of a rooted tree, with root node given by the unary relation Root , and leaves given by the unary relation Leaf . The edges are directed from the leaves towards the root of the tree.

As we shall see, the latter axiom is expressible by conjoining the first axiom (Tree) with some appropriately-chosen sentences in \mathbf{C}^2 , so only the first axiom will require a full algorithmic treatment. The second will then follow as a consequence of Theorem 3.

4.1 Trees

In this section, we prove that the addition of a tree axiom $\text{Tree}(R, \text{Leaf})$ to a \mathbf{C}^2 sentence preserves domain-liftability. Our approach will be as follows: first, we will show that the addition of a “limited” version of the tree axiom that does not allow for quantification of the leaves (i.e. $\text{Tree}(R)$)

¹The reason why this is sound is as follows. No tree can contain the edges $\{i, j\}, \{i, k\}, \{j, k\}$ at the same time or else it would contain a cycle. Moreover there is a natural bijection between the trees which contain the edge $\{i, k\}$ and those that contain the edge $\{j, k\}$, which maps trees with weight $W \cdot w(\{i, k\})$ to trees with weight $W \cdot w(\{j, k\})$ (here W is the product of all other edges in the tree except for $\{i, k\}$ or $\{j, k\}$). The bijection simply takes a tree containing $\{i, k\}$ and replaces it by $\{j, k\}$. It is then not difficult to see the correctness of the procedure described here.

rather than $\text{Tree}(R, \text{Leaf})$) can be added to an \mathbf{FO}^2 sentence while preserving domain-liftability (Lemma 1). We will then build off of this lemma to show the same for \mathbf{C}^2 (Lemma 2). Finally, we will leverage this result to prove our final theorem, showing that we can add the “full” tree axiom ($\text{Tree}(R, \text{Leaf})$) to any \mathbf{C}^2 sentence and still compute the WFOMC in time polynomial in the domain size (Theorem 5). We begin by formalizing the semantics of the two tree axioms—both the limited version $\text{Tree}(R)$ as well as the full axiom $\text{Tree}(R, \text{Leaf})$.

Definition 7. Let ϕ be an arbitrary first-order sentence, possibly containing some binary relation R . Then a structure \mathcal{A} is a model of the sentence $\xi = \phi \wedge \text{Tree}(R)$ if and only if:

1. \mathcal{A} is a model of ϕ , and
2. the relation defined by \mathcal{A}_R is antireflexive and symmetric, and
3. the relation R forms an R -tree in \mathcal{A} : that is, \mathcal{A}_R is a tree when interpreted as an undirected graph

The concept is illustrated with the following simple example.

Example 1. Let $\xi = \text{Tree}(R)$. Then, by Cayley’s formula, $\text{FOMC}(\xi, n) = n^{n-2}$, the number of trees on n nodes (Cayley 1889).

We can now extend this definition to the full axiom that allows for expression of properties on the leaves.

Definition 8. Let ϕ be an arbitrary first-order sentence, possibly containing some binary relation R and unary relation Leaf . Then a structure \mathcal{A} is a model of the sentence $\xi = \phi \wedge \text{Tree}(R, \text{Leaf})$ if and only if:

1. \mathcal{A} is a model of $\phi \wedge \text{Tree}(R)$, and
2. the leaves of the tree defined by \mathcal{A}_R are precisely the elements in $\mathcal{A}_{\text{Leaf}}$

With these definitions in mind, we are now ready to prove our domain-liftability results, starting with the lemma for the limited tree axiom on two-variable sentences without counting quantifiers.

Lemma 1. Let:

$$\begin{aligned} \phi &= \eta \wedge \\ &\forall x : \neg R(x, x) \wedge \\ &\forall x \forall y : R(x, y) \rightarrow R(y, x) \end{aligned}$$

for some sentence $\eta \in \mathbf{FO}^2$. Further, let $\xi = \phi \wedge \text{Tree}(R)$. Then ξ is domain-liftable.

Proof. We will follow the notation and general approach from Theorem 1. Fix our domain $\Delta = \{1, \dots, n\}$. After transforming ϕ to the form $\phi = \forall x \forall y : \psi(x, y)$, recall from Equation (1) that we can write the WFOMC of ϕ as:

$$\text{WFOMC}(\phi, n, w, \bar{w}) = \sum_{n_1 + \dots + n_{2^u} = n} \binom{n}{n_1, \dots, n_{2^u}} \prod_{i, j \in [2^u], i < j} r_{ij}^{n_i n_j} \prod_{i \in [2^u]} s_i^{n_i(n_i-1)/2} w_i^{n_i}$$

for suitably defined r_{ij} , s_i , and w_i terms.

We will now need to adapt this approach for ξ so that we account only for the models of ϕ that satisfy the tree property on R . Suppose we fix some cell assignment (n_1, \dots, n_{2^u}) . All of the notation below will assume this fixed cell assignment. For $a, b \in \Delta$, denote:

$$\varphi(a, b) = \begin{cases} \psi_{ij}(a, b) & \text{if } i = C(a), j = \\ & C(b) \text{ and } i \neq j \\ \psi_i(a, b) \wedge \psi_i(b, a) & \text{if } i = C(a) = C(b) \end{cases}$$

where $C(a)$ maps a domain element a to its cell index. Define \mathcal{F} to be the set of all domain element pairs $\{a, b\} \in \Delta^2$ for which $\text{WMC}(\varphi(a, b) \wedge \neg R(a, b), w, \bar{w}) = 0$.

Consider some R -tree T over the vertices defined by Δ such that $\mathcal{F} \subseteq T$. We will compute the weighted sum w_T of every model \mathcal{A} of ξ that satisfies $T = \mathcal{A}_T$ (in other words, the WFOMC of ξ limited to models that contain exactly the tree T). We can write:

$$w_T = \prod_{i \in [2^u]} w_i^{n_i} \prod_{\{a, b\} \in \mathcal{F}} \text{WMC}(\varphi(a, b) \wedge R(a, b), w, \bar{w}) \cdot \prod_{a < b, \{a, b\} \in T \setminus \mathcal{F}} \frac{\text{WMC}(\varphi(a, b) \wedge R(a, b), w, \bar{w})}{\text{WMC}(\varphi(a, b) \wedge \neg R(a, b), w, \bar{w})} \cdot \prod_{a < b, \{a, b\} \notin \mathcal{F}} \text{WMC}(\varphi(a, b) \wedge \neg R(a, b), w, \bar{w})$$

Now, summing both sides of the equation across all trees containing \mathcal{F} :

$$\begin{aligned} & \sum_{T: \mathcal{F} \subseteq T} w_T \\ = & \sum_{T: \mathcal{F} \subseteq T} \prod_{i \in [2^u]} w_i^{n_i} \prod_{\{a, b\} \in \mathcal{F}} \text{WMC}(\varphi(a, b) \wedge R(a, b), w, \bar{w}) \cdot \prod_{a < b, \{a, b\} \in T \setminus \mathcal{F}} \frac{\text{WMC}(\varphi(a, b) \wedge R(a, b), w, \bar{w})}{\text{WMC}(\varphi(a, b) \wedge \neg R(a, b), w, \bar{w})} \cdot \prod_{a < b, \{a, b\} \notin \mathcal{F}} \text{WMC}(\varphi(a, b) \wedge \neg R(a, b), w, \bar{w}) \\ = & \prod_{i \in [2^u]} w_i^{n_i} \prod_{a < b, \{a, b\} \notin \mathcal{F}} \text{WMC}(\varphi(a, b) \wedge \neg R(a, b), w, \bar{w}) \cdot \sum_{T: \mathcal{F} \subseteq T} \prod_{\{a, b\} \in \mathcal{F}} \text{WMC}(\varphi(a, b) \wedge R(a, b), w, \bar{w}) \cdot \prod_{a < b, \{a, b\} \in T \setminus \mathcal{F}} \frac{\text{WMC}(\varphi(a, b) \wedge R(a, b), w, \bar{w})}{\text{WMC}(\varphi(a, b) \wedge \neg R(a, b), w, \bar{w})} \\ = & \text{TreeSum}(\mathbb{W}, \mathcal{F}) \prod_{i \in [2^u]} w_i^{n_i} \cdot \prod_{a < b, \{a, b\} \notin \mathcal{F}} \text{WMC}(\varphi(a, b) \wedge \neg R(a, b), w, \bar{w}) \end{aligned}$$

where \mathbb{W} is the graph with weighted adjacency matrix:

$$\mathbb{W}_{(a, b)} = \begin{cases} \text{WMC}(\varphi(a, b) \wedge R(a, b), w, \bar{w}) & \text{if } \{a, b\} \in \mathcal{F} \\ \frac{\text{WMC}(\varphi(a, b) \wedge R(a, b), w, \bar{w})}{\text{WMC}(\varphi(a, b) \wedge \neg R(a, b), w, \bar{w})} & \text{otherwise} \end{cases}$$

Recall that we saw how to compute $\text{TreeSum}(\mathbb{W}, \mathcal{F})$ in time polynomial in the size of \mathbb{W} earlier in the paper. Denote $Z(n_1, \dots, n_{2^u}) = \sum_{T: \mathcal{F} \subseteq T} w_T$. That is, $Z(n_1, \dots, n_{2^u})$ is the weighted sum of the models satisfying the tree axiom that correspond to the cell configuration (n_1, \dots, n_{2^u}) . Clearly, each $Z(n_1, \dots, n_{2^u})$ can be computed in time polynomial in the domain size n using the expression derived above. Summing across all possible cell configurations we get:

$$\text{WFOMC}(\xi, n, w, \bar{w}) = \sum_{n_1 + \dots + n_{2^u} = n} \binom{n}{n_1, \dots, n_{2^u}} \cdot Z(n_1, \dots, n_{2^u})$$

which again can be evaluated in time polynomial in the domain size.

Finally, note that this result also holds in the presence of cardinality constraints on ξ . Recall that Theorem 2 asserts that if we can efficiently compute the WFOMC of a sentence, then we can also efficiently compute the WFOMC of that same sentence extended with some cardinality constraints, where ‘‘efficiently’’ means in time polynomial in the size of the domain. We can also apply this theorem to the sentence ξ containing a tree axiom. \square

Next, we will use Lemma 1 to show that we can add the limited tree axiom to a two-variable sentence with counting quantifiers.

Lemma 2. *Let:*

$$\begin{aligned} \phi &= \eta \wedge \\ & \forall x : \neg R(x, x) \wedge \\ & \forall x \forall y : R(x, y) \rightarrow R(y, x) \end{aligned}$$

for some sentence $\eta \in \mathbf{C}^2$. Further, let $\xi = \phi \wedge \text{Tree}(R)$. Then ξ is domain-liftable.

Proof. From Lemma 1, we know that we can compute the WFOMC of any two-variable sentence with a single tree axiom and cardinality constraints in time polynomial in the size of the domain. In addition, recall that from Theorem 3, we can take an arbitrary \mathbf{C}^2 sentence ϕ and produce an \mathbf{FO}^2 sentence ϕ' along with cardinality constraints C and weights (w', \bar{w}') which have the property that $\text{WFOMC}(\phi, n, w, \bar{w}) = K_{\phi, n} \cdot \text{WFOMC}(\phi', n, w', \bar{w}', C)$. Importantly, the reduction also gives us the constant $K_{\phi, n}$ so that we can compute the WFOMC of ϕ efficiently.

Now, one problem when applying this reduction in our case could be that $\text{Tree}(R)$ is not expressible in \mathbf{C}^2 . That is true in general, but for a fixed domain we can express it in \mathbf{C}^2 as a ground formula (as described in Remark 1). Obviously, the size of such an encoding of the tree axiom would be very large (exponential in the domain size). Fortunately, we do not need to actually construct it. By inspecting the reduction in (Kuzelka 2021), one can verify that the reduction is modular (Property 1) and leaves the encoding of the tree axiom untouched, so we could replace it in the end by $\text{Tree}(R)$, which also means that we did not have to replace $\text{Tree}(R)$ by its ground encoding in the first place. Hence we see that the transformation from (Kuzelka 2021) works also in the presence of tree axioms. \square

Finally, putting Lemma 1 and 2 together gives us our final result for the full axiom $\text{Tree}(R, \text{Leaf})$.

Theorem 5. *Let:*

$$\begin{aligned}\phi &= \eta \wedge \\ &\forall x : \neg R(x, x) \wedge \\ &\forall x \forall y : R(x, y) \rightarrow R(y, x)\end{aligned}$$

for some $\eta \in \mathbf{C}^2$. Further, let $\xi = \phi \wedge \text{Tree}(R, \text{Leaf})$. Then ξ is domain-liftable.

Proof. Let

$$\gamma = \forall x : \text{Leaf}(x) \leftrightarrow (\exists^=1 y : R(x, y))$$

Clearly, $\text{Leaf}(x)$ defines precisely the leaves in the tree defined by R , and thus:

$$\text{WFOMC}(\xi, n, w, \bar{w}) = \text{WFOMC}(\phi \wedge \gamma \wedge \text{Tree}(R), n, w, \bar{w})$$

Since the right-hand of this expression is computable in time polynomial in the domain size n by Lemma 2, we have that ξ is domain-liftable. \square

4.2 Directed Rooted Trees

Of course, in practice one often wishes to model something more interesting than just the plain Tree axiom like we saw in Example 1. We will now show how to model *directed rooted trees* as described at the beginning of the section. Recall that this axiom expresses that the binary relation E represents directed edges of a tree, with root node given by the unary relation Root , and leaves given by the unary relation Leaf . Note here that the root node is never considered a leaf, even if it has degree 1.

Proposition 1. *Let $\phi \in \mathbf{C}^2$ be an arbitrary sentence, and let ψ denote the sentence below:*

$$\begin{aligned}\psi &= \exists^=1 x : \text{Root}(x) \wedge \\ &\forall x : \text{Root}(x) \rightarrow D(x, x) \wedge \\ &\forall x \exists^=1 y : D(x, y) \wedge \\ &\forall x \forall y : R(x, y) \rightarrow (D(x, y) \vee D(y, x)) \wedge \\ &\forall x \forall y : E(x, y) \leftrightarrow (D(x, y) \wedge (\neg \text{Root}(x) \vee \neg \text{Root}(y))) \wedge \\ &\text{Tree}(R) \wedge \forall x : \text{Leaf}(x) \leftrightarrow (\forall y : \neg D(y, x))\end{aligned}$$

where R and D are fresh predicate symbols not occurring in ϕ . Then:

$$\begin{aligned}\text{WFOMC}(\phi \wedge \text{DirectedRootedTree}(\text{Root}, E, \text{Leaf}), \\ w, \bar{w}, n, C) = \text{WFOMC}(\phi \wedge \psi, w', \bar{w}', n, C)\end{aligned}$$

where (w', \bar{w}') is obtained by extending (w, \bar{w}) with neutral weights on all new predicates in ψ .

Proof. We can motivate this choice of ψ as follows. We pick a single vertex as a root, which we enforce² with $\exists^=1 x : \text{Root}(x)$. We then define a new relation $D(x, y)$ that will represent auxiliary directed edges (the directed edges of the

²Equivalently, we may simply add the cardinality constraint $|\text{Root}| = 1$ instead.

tree will be represented by another relation E). The edges are oriented in such a way that the edges go from vertices further from the root, to vertices closer to it. First, we add a self-loop for the root node $\forall x : \text{Root}(x) \rightarrow D(x, x)$. We also require that D satisfies the functionality constraint and that whenever there is an original undirected edge $R(x, y)$, then there must be either a directed edge $D(x, y)$ or $D(y, x)$. Since D must satisfy the functionality constraint and the root must have a self loop, any edge that connects the root with another vertex, must be directed towards the root (if any of these edges were directed away from the root then D would not satisfy the functionality constraint). We can use induction on the depth of the tree using essentially the same argument (the fact that any vertex must have out-degree 1 because of the functionality constraint) to show that D must be oriented in such a way that there is a connected path from any leaf to the root. Note that we cannot immediately use D as our directed edge relation, since D has a self-loop at the root node, which should not be present in a directed rooted tree. As a solution, we add a ‘‘wrapper’’ predicate E that has the same edges as D except for the self-loop at the root node. Last, we enforce with the axiom $\text{Tree}(R)$ that the relation R forms an undirected tree. To encode the Leaf predicate, we can notice that leaves are precisely those vertices that have zero in-degree w.r.t. D (in particular, the root is never considered a leaf even if it has degree 1) and therefore we can encode leaves using $\forall x : \text{Leaf}(x) \leftrightarrow (\forall y : \neg D(y, x))$.

It is clear that the models of the aforementioned rules are precisely those satisfying the axiom $\text{DirectedRootedTree}(\text{Root}, E, \text{Leaf})$. It is also not difficult to show that, under the assumption that D is not used anywhere else (which we can assume w.l.o.g.), the reduction is also modular. \square

Theorem 6. *Let:*

$$\xi = \phi \wedge \text{DirectedRootedTree}(\text{Root}, E, \text{Leaf})$$

for some $\phi \in \mathbf{C}^2$. Then ξ is domain-liftable.

Proof. Follows immediately from Proposition 1 and Theorem 5. \square

5 Applications and Experiments

In this section, we will describe some applications of the tree axiom and examine the scalability of our approach in practice. To test our results experimentally, we implemented a weighted first-order model counter in Python that follows the algorithmic approach presented here with support for tree axioms. All experiments were performed on a computer with a six-core Intel i7 2.2GHz processor and 16 GB of RAM.

5.1 Combinatorics on Trees

We start by modelling some problems from enumerative combinatorics on trees.

k -coloured trees We first consider the problem of counting k -coloured trees. Recall that a graph is said to be k -colourable if every vertex can be coloured with one of k colours such that no two vertices of the same colour




Figure 1: Runtime for counting 2-, 3-, and 4-coloured trees

are adjacent to one another. A k -coloured graph is a k -colourable graph together with a valid colour assignment. We first verified that our approach gave the correct expected sequence for 2-coloured trees: namely, we should have $\text{FOMC}(\xi_2, n) = 2n^{n-2}$, because every one of the n^{n-2} trees on n nodes is 2-colourable, and any such tree has precisely two colourings (since colouring any one node uniquely determines the colour of the remainder). We then tested it on larger values of k . An indication of the runtime of our approach (with a 30 second timeout) on 2-, 3-, and 4-coloured trees is given in Figure 1.

Rooted trees with k leaves We next consider the problem posed in the introduction of the paper: how many labelled rooted trees on n nodes have exactly k leaves? This can be modelled with the following sentence:

$$\xi = \text{DirectedRootedTree}(\text{Root}, E, \text{Leaf})$$

subject to the cardinality constraint $|\text{Leaf}| = k$. Note that, although this sentence appears relatively simple at first glance, computing its FOMC takes some work. First, the axiom is expanded into a \mathcal{C}^2 sentence as shown in Proposition 1. Next, the counting quantifiers introduced by the `DirectedRootedTree` axiom are reduced to cardinality constraints (Theorem 3), for which computing the FOMC is split across several different FO^2 WFOMC oracle calls with different weights (Theorem 2). Finally, existential quantifiers in each of these calls must be eliminated as described in the proof of Theorem 1, which is a process that itself introduces new auxiliary Skolemization predicates (see (Van den Broeck, Meert, and Darwiche 2014) for the full details).

As a result of the steps above necessary to compute the FOMC, despite the fact that the runtime is still guaranteed to be polynomial in the domain size, our implementation struggled to scale to large values of n because the degree of the polynomial bounding the runtime was simply too high. We were able to scale to a domain size of $n = 9$ within a 10 minute timeout. However, we were able to check these results against *The On-Line Encyclopedia of Integer Sequences* (OEIS)³ for values $k = 1, 2$ and 3 with sequences A000142,

³<https://oeis.org>

Number of nodes n	Number of leaves k		
	1	2	3
2	2	—	—
3	6	3	—
4	24	36	4
5	120	360	140
6	720	3600	3000
7	5040	37800	54600
8	40320	423360	940800
9	362880	5080320	16087680

Table 1: Counting labelled rooted trees on n nodes with exactly k leaves

A055303, and A055304 respectively. The values we computed are given in Table 1.

5.2 Markov Logic Networks over Trees

We now turn to applications of our results to probabilistic inference in graphical models. Specifically, in this section we study the imposition of tree axioms on *Markov logic networks* (MLNs) (Richardson and Domingos 2006), a popular statistical-relational formalism for which inference and learning is reducible to WFOMC.

Preliminaries Formally, an MLN is a finite set of weighted first-order logic formulas $\{(w_1, \phi_1), \dots, (w_n, \phi_n)\}$, where each w_i is either a real-valued weight or ∞ , and ϕ_i is a quantifier-free first-order formula. MLNs can be viewed as a template for constructing Markov random fields. An MLN Φ paired with a domain Δ induces a probability distribution on possible worlds:

$$\Pr(\omega; \Phi, \Delta) = \begin{cases} \frac{1}{Z_{\Phi, \Delta}} \exp\left(\sum_{(w, \phi) \in \Phi_{\mathbb{R}}} w \cdot N(\phi, \omega)\right) & \text{if } \omega \models \Phi_{\infty} \\ 0 & \text{otherwise} \end{cases}$$

where $\Phi_{\mathbb{R}}$ and Φ_{∞} denote the real-valued and ∞ -valued formulas in Φ respectively, $N(\phi, \omega)$ denotes the number of groundings of ϕ satisfied in the world ω , and $Z_{\Phi, \Delta}$ is a normalization constant called the *partition function* of the MLN.

Reduction to WFOMC In general, computing the partition function (and by extension, marginal inference) in a given MLN is $\#P$ -hard in the domain size. However, the aforementioned reduction to WFOMC has the helpful property that the number of variables in the sentence produced by the reduction is the same as the maximum number of variables appearing in any formula of the original MLN. In particular, this means that if the number of variables used in each of the formulas in the MLN is limited to two, the inference problem for such MLNs lies in FP. We present the reduction in question below.

Definition 9 (Van den Broeck 2011). *Let $\Phi = \{(w_1, \phi_1), \dots, (w_k, \phi_k), (\infty, \phi_{k+1}), \dots, (\infty, \phi_n)\}$ be an*




Figure 2: Runtime of our tool for computing the partition function of the “friends-and-smokers” MLN for various domain sizes, subject to soft and hard tree constraints

MLN. Then the reduction of Φ to a weighted sentence is constructed as:

$$\bigwedge_{i=1}^k \forall \mathbf{x} : A_i(\mathbf{x}) \leftrightarrow \phi_i(\mathbf{x}) \wedge \bigwedge_{i=k+1}^n \forall \mathbf{x} : \phi_i(\mathbf{x})$$

where each A_i is a fresh predicate not occurring in any of the ϕ_i , \mathbf{x} denotes the collection of free variables occurring in each ϕ_i , and weights are set as follows: $w(A_i) = e^{w_i}$, $\bar{w}(A_i) = 1$, and $w(R_i) = \bar{w}(R_i) = 1$ for any other predicate R_i .

Experiments Consider the classic “friends-and-smokers” MLN described below:

$$\begin{aligned} w_1 & \text{Smokes}(x) \\ w_2 & \text{Friends}(x, y) \wedge \text{Smokes}(x) \rightarrow \text{Smokes}(y) \\ \infty & \neg \text{Friends}(x, x) \end{aligned}$$

Intuitively, this models that people who are friends with smokers are likely to smoke themselves. We consider the problem of computing the partition function of this MLN over different domain sizes, subject to one of two constraints: either, (1) the hard constraint that the *Friends* relation must form a tree, or (2) the soft constraint that models in which *Friends* forms a tree are to be preferred.

The former constraint can be dealt with by simply conjoining the axiom $\text{Tree}(\text{Friends})$ to the sentence obtained from applying the reduction in Definition 9. On the other hand, the latter constraint requires adding a new formula ($w_3, R(x, y) \leftrightarrow \text{Friends}(x, y)$) to the MLN for an appropriate weight value w_3 , and conjoining the axiom $\text{Tree}(R)$ to the reduced sentence. Experimental results for both the hard and soft constraints with fixed weight values are shown in Figure 2; the computation time of the partition function is virtually identical in both cases.

6 Conclusion and Future Work

We showed how to extend existing domain liftability results on FO^2 and C^2 to allow for the addition of tree axioms, and

showed how to extend this core notion to model directed rooted trees. We also presented some preliminary experimental results on our approach.

There are several possible directions for future work. One avenue we intend to explore involves extending the tree axiom presented here to a *forest* axiom, stating that the graph formed by some (antireflexive, symmetric) relation is (only) acyclic. This relaxes the connectedness condition imposed by the tree axiom. Efficient approaches for counting forests are known in the combinatorics literature (Takács 1990); we hope that translating them to the setting of WFOMC is not difficult and can proceed in a similar fashion to the approach taken in this paper.

More generally, we look forward to seeing how far one can go beyond first-order logic with richer fragments while maintaining domain-liftability: for instance, showing domain-liftability for fragments of Datalog could have important implications for the probabilistic logic programming community (Fierens et al. 2015). In this area, existing lifted inference techniques have struggled to see widespread practical use, due to the barriers of translating approaches geared towards first-order logic to this more expressive setting. Although initial results in this direction for probabilistic databases have been negative, showing #P-hardness for evaluation of unbounded UCQ^∞ queries on probabilistic graphs (Amarilli and Ceylan 2020), the picture may be different in the symmetric WFOMC setting considered here.

Acknowledgements

TvB was supported by the Research Foundation – Flanders (G095917N). OK was supported by Czech Science Foundation project “Generative Relational Models” (20-19104Y) and partially by the OP VVV project CZ.02.1.01/0.0/0.0/16_019/0000765 “Research Center for Informatics”.

A Directed Trees: An Alternative Transform

In this section we record a practically more efficient but less straightforward encoding for directed rooted trees:

$$\begin{aligned} \psi = & \text{Tree}(R) \wedge \forall x : \text{Root}(x) \leftrightarrow D(x, x) \wedge \\ & \forall x \exists^=1 y : D(x, y) \wedge \\ & \forall x \forall y : D(x, y) \rightarrow (R(x, y) \vee D(x, x)) \wedge \\ & \forall x \forall y : R(x, y) \rightarrow (D(x, y) \vee D(y, x)) \wedge \\ & \forall x \forall y : E(x, y) \leftrightarrow (D(x, y) \wedge (\neg \text{Root}(x) \vee \neg \text{Root}(y))) \wedge \\ & \forall x : \text{Leaf}(x) \leftrightarrow (\forall y : \neg D(y, x)) \end{aligned}$$

This transformation has one less rule with a counting quantifier than the transformation described in the main text ($\exists^=1 x : \text{Root}(x)$) and, instead of it, it contains the formula $\forall x \forall y : D(x, y) \rightarrow (R(x, y) \vee D(x, x))$ which together with the functionality constraint on D and the fact that the cardinality of R must be $n - 1$ (and together with the rest of the formula) implies that there must be exactly one x such that $D(x, x)$, which must be the root (by $\forall x : \text{Root}(x) \leftrightarrow D(x, x)$).

References

- Amarilli, A., and Ceylan, İ. İ. 2020. A dichotomy for homomorphism-closed queries on probabilistic graphs. In *ICDT*, volume 155 of *LIPICs*, 5:1–5:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- Beame, P.; Van den Broeck, G.; Gribkoff, E.; and Suciu, D. 2015. Symmetric weighted first-order model counting. In *PODS*, 313–328. ACM.
- Benaïm, S.; Benedikt, M.; Charatonik, W.; Kieronski, E.; Lenhardt, R.; Mazowiecki, F.; and Worrell, J. 2016. Complexity of two-variable logic on finite trees. *ACM Trans. Comput. Log.* 17(4):32:1–32:38.
- Braun, T., and Möller, R. 2016. Lifted junction tree algorithm. In *KI*, volume 9904 of *Lecture Notes in Computer Science*, 30–42. Springer.
- Cayley, A. 1889. A theorem on trees. *Quart. J. Pure Appl. Math.* 23:376–378.
- Chaiken, S., and Kleitman, D. J. 1978. Matrix tree theorems. *J. Comb. Theory Ser. A.* 24(3):377–381.
- Charatonik, W., and Witkowski, P. 2013. Two-variable logic with counting and trees. In *LICS*, 73–82. IEEE Computer Society.
- Dalvi, N., and Suciu, D. 2012. The dichotomy of probabilistic inference for unions of conjunctive queries. *J. ACM* 59(6):30:1–30:87.
- Fierens, D.; Van den Broeck, G.; Renkens, J.; Shterionov, D.; Gutmann, B.; Thon, I.; Janssens, G.; and De Raedt, L. 2015. Inference and learning in probabilistic logic programs using weighted boolean formulas. *TPLP* 15(3):358–401.
- Grädel, E.; Kolaitis, P. G.; and Vardi, M. Y. 1997. On the decision problem for two-variable first-order logic. *Bull. Symb. Log.* 3(1):53–69.
- Gribkoff, E.; Van den Broeck, G.; and Suciu, D. 2014. Understanding the complexity of lifted inference and asymmetric weighted model counting. In *UAI*, 280–289. AUAI Press.
- Kazemi, S. M.; Kimmig, A.; Van den Broeck, G.; and Poole, D. 2016. New liftable classes for first-order probabilistic inference. In *NIPS*, 3117–3125.
- Kuusisto, A., and Lutz, C. 2018. Weighted model counting beyond two-variable logic. In *LICS*, 619–628. ACM.
- Kuzelka, O. 2021. Weighted first-order model counting in the two-variable fragment with counting quantifiers. *J. Artif. Intell. Res.* 70:1281–1307.
- Libkin, L. 2004. *Elements of Finite Model Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer.
- Milch, B.; Zettlemoyer, L. S.; Kersting, K.; Haimes, M.; and Kaelbling, L. P. 2008. Lifted probabilistic inference with counting formulas. In *AAAI*, 1062–1068. AAAI Press.
- Poole, D. 2003. First-order probabilistic inference. In *IJCAI*, 985–991. Morgan Kaufmann.
- Richardson, M., and Domingos, P. M. 2006. Markov logic networks. *Machine Learning* 62(1-2):107–136.
- Taghipour, N.; Fierens, D.; Davis, J.; and Blockeel, H. 2013. Lifted variable elimination: Decoupling the operators from the constraint language. *J. Artif. Intell. Res.* 47:393–439.
- Takács, L. 1990. On Cayley’s formula for counting forests. *J. Comb. Theory Ser. A.* 53(2):321–323.
- Van den Broeck, G.; Taghipour, N.; Meert, W.; Davis, J.; and De Raedt, L. 2011. Lifted probabilistic inference by first-order knowledge compilation. In *IJCAI*, 2178–2185. IJCAI/AAAI.
- Van den Broeck, G.; Meert, W.; and Darwiche, A. 2014. Skolemization for weighted first-order model counting. In *KR*. AAAI Press.
- Van den Broeck, G. 2011. On the completeness of first-order knowledge compilation for lifted probabilistic inference. In *NIPS*, 1386–1394.
- Vardi, M. Y. 1982. The complexity of relational query languages (extended abstract). In *STOC*, 137–146. ACM.