Counting and Sampling Models in First-Order Logic

Ondřej Kuželka
Faculty of Electrical Engineering, Czech Technical University in Prague
ondrej.kuzelka@fel.cvut.cz

Abstract

First-order model counting (FOMC) is the task of counting models of a first-order logic sentence over a given set of domain elements. Its weighted variant, WFOMC, generalizes FOMC by assigning weights to the models and has many applications in statistical relational learning. More than ten years of research by various authors has led to identification of non-trivial classes of WFOMC problems that can be solved in time polynomial in the number of domain elements. In this paper, we describe recent works on WFOMC and the related problem of weighted first-order model sampling (WFOMS). We also discuss possible applications of WFOMC and WFOMS within statistical relational learning and beyond, e.g., automated solving of problems from enumerative combinatorics and elementary probability theory. Finally, we mention research problems that still need to be tackled in order to make applications of these methods really practical more broadly.

1 Introduction

First-order logic is a formal framework that allows reasoning about statements such as *all humans are mortal*, which can be expressed by the first-order logic sentence $\forall x : Human(x) \Rightarrow Mortal(x)$. It is a useful tool for reasoning about information that is certain. *Weighted first-order model counting* (WFOMC) [Van den Broeck *et al.*, 2011] is the task of computing the weighted sum of models of a given first-order logic sentence (how the weights are computed will be explained in more detail in Section 2). Crucially, WFOMC extends the reach of what can be done with first-order logic to reasoning under uncertainty—it is used among others for inference tasks in statistical relational learning [Getoor and Taskar, 2007]. Another related problem is *weighted first-order model sampling* (WFOMS) [Wang *et al.*, 2022; Wang *et al.*, 2023] which, given a first-order logic sentence, asks to sample a model of the sentence with probability proportional to its weight and can be used in generative statistical relational learning models.

In general, both WFOMC and WFOMS are intractable [Beame *et al.*, 2015; Wang *et al.*, 2023], however, there are non-trivial fragments of first-order logic which have been identified as tractable for WFOMC [Van den Broeck, 2011; Van den Broeck *et al.*, 2014; Kazemi *et al.*, 2016; Kuusisto and Lutz, 2018; Kuželka, 2021; Van Bremen and Kuželka, 2021; Tóth and Kuželka, 2023] and for WFOMS [Wang *et al.*, 2022; Wang *et al.*, 2023]. These fragments are already interesting enough for applications, which we discuss in more depth in Section 4.

Before delving into the technical details and before describing the recent results on tractable fragments, which we do in the next sections, let us use the the remainder of this section to give two toy examples illustrating the use of WFOMC and WFOMS.

1.1 Model Counting: A Toy Example

We start with an example illustrating the use of first-order model counting for high-level probabilistic reasoning.

Example 1. *Alice wants to set up a Secret Santa event for students of her combinatorics class. For that she needs to randomly assign who will be giving presents to whom, with the constraint that no one should be assigned to give presents to themselves. So she puts the names of her students written on pieces of paper in a hat and lets everyone draw one name from the hat. If someone draws their own name, the whole procedure is restarted from scratch. She knows that with high probability she will not need to repeat the process for too long. Why? Assuming the process by which the students are drawing names from the hat is really random, the process can be thought of as sampling a permutation uniformly at random. The event that no one is assigned to themselves corresponds to drawing a permutation without fixed points, also known as a derangement. Having taught combinatorics, Alice knows how to compute the number of derangements of a set of* n *elements, usually denoted* $!n$*. Say, Alice's class has 30 students. Then the probability of success is* $!30/30! \approx 0.37$*. From that she can easily compute, for instance, an upper bound on the expected number of repetitions she would need.*

The problem of counting derangements was solved already in 18th century and is a standard topic in introductory discrete math courses, nonetheless, let us now pretend that we do not know how to solve it. So instead of computing the formula for the number of derangements by hand, we now show how this problem could be solved using first-order model counting.

Example 2 (Continued)**.** *We can start by writing a first-order logic sentence that encodes permutations:*

$$\Psi_{perm} = \forall x \exists^{=1} y\ R(x, y) \wedge \forall y \exists^{=1} x\ R(x, y),$$

which can be read as follows: First, $\forall x \exists^{=1} y\ R(x, y)$ expresses that R is a binary relation such that for every x there is exactly one[1] y for which $R(x, y)$ holds (i.e., R is a function). Second, $\forall y \exists^{=1} x\ R(x, y)$ means that for every y there is exactly one x such that $R(x, y)$. Having already established that R is a function, we can read this together as: R is a bijection, or in other words, R is a permutation. So the next step is to encode derangements but that is easy—we only need to forbid the permutations to have fixed points and that can be done by appending $\forall x\ \neg R(x, x)$ to Ψ_{perm}. That is:

$$\Psi_{dera} = \Psi_{perm} \wedge \forall x\ \neg R(x, x).$$

Now, if we enumerated the models of the sentence Ψ_{perm} on a set of n domain elements, they would all correspond to permutations. Likewise, if we enumerated the models of the sentence Ψ_{dera} they would all correspond to derangements. Hence, to compute the probability that a randomly sampled permutation of 30 elements is a derangement, we can just compute $\mathrm{FOMC}(\Psi_{dera}, 30)/\,\mathrm{FOMC}(\Psi_{perm}, 30) \approx 0.37$, where we used $\mathrm{FOMC}(\Psi, n)$ to denote the (unweighted) model count of the sentence Ψ on a domain of size n. Importantly, as it turns out both Ψ_{perm} and Ψ_{dera} are tractable for WFOMC, which means we can compute the model counts in time polynomial in n.

1.2 Model Sampling: A Toy Example

The strategy used by Alice in Example 1 to sample derangements is known as *rejection sampling*. It works fine in this case because the probability of success, which here corresponds to obtaining a derangement when sampling permutations, is relatively high. However, that is not always the case as illustrated by the next example.

Example 3. *Suppose that Alice does not teach just combinatorics but also discrete mathematics, some students take both courses while others take only one, and she wants to organize a joint Secret Santa for all her students, satisfying the following constraints: (i) no one is assigned to give a present to themselves, (ii) if A is assigned to give a present to B then both A and B must be taking at least one of the two courses together. In this case, depending on the size of the overlap of the two courses, if Alice wanted to use the same rejection sampling strategy as in Example 1, the number of required repetitions might grow very fast (exponentially in the number of students).[2] Therefore in this case, Alice needs a different solution.*

[1] The "exists exactly one" part is what the *counting* quantifier $\exists^{=1}$ expresses. Counting quantifiers are introduced in Section 2.

[2] To illustrate this, denote the set of students in the combinatorics course $\mathcal{C}$ and the set of students in the discrete mathematics course $\mathcal{D}$. Suppose that there are $3n$ students in total and that $|\mathcal{C} \setminus \mathcal{D}| = n$, $|\mathcal{C} \cap \mathcal{D}| = n$, and $|\mathcal{D} \setminus \mathcal{C}| = n$. Then it can be shown that the probability that the constraint (ii) is satisfied in a randomly sampled permutation is bounded from above by $1/2^n$. That also means that the expected runtime of the rejection sampling scheme would be exponential in n.

Fortunately, as we discuss in Section 2, one can use WFOMS to solve Alice's problem from the above example in polynomial time.[3]

Example 4 (Example 3 continued)**.** *To solve her sampling problem using WFOMS, Alice can construct the first-order logic sentence:*

$$\begin{aligned}&\Psi_{dera} \wedge \\ &\forall x \forall y\ (R(x, y) \Rightarrow ((C(x) \wedge C(y)) \vee (D(x) \wedge D(y)))) \\ &\wedge \bigwedge_{t \in \mathcal{C}} C(t) \wedge \bigwedge_{t \in \Delta \setminus \mathcal{C}} \neg C(t) \wedge \bigwedge_{t \in \mathcal{D}} D(t) \wedge \bigwedge_{t \in \Delta \setminus \mathcal{D}} \neg D(t),\end{aligned}$$

where $\mathcal{C}$ contains the students taking combinatorics, $\mathcal{D}$ the students taking discrete math and $\Delta = \mathcal{C} \cup \mathcal{D}$ is the domain. We can read this sentence as: R must be a derangement (that is the Ψ_{dera} part), if x is assigned to give a present to y then they both take the combinatorics course $((C(x) \wedge C(y)))$ or they both take the discrete math course $((D(x) \wedge D(y)))$, the last part specifies which students take which classes. What remains is just to call WFOMS to generate a sample.

2 Tractable Counting

The search for fragments of first-order logic that are tractable for WFOMC has roots in the area of artificial intelligence known as statistical relational learning [Getoor and Taskar, 2007], where people had noticed that the first-order nature of statistical relational learning models allows more efficient inference than what can be achieved by standard inference algorithms on ground instances of the same models, e.g., [Poole, 2003]. The overarching term used to describe the class of methods that exploit first-order nature and symmetries in such models is *lifted inference* (see, e.g., the book [Van den Broeck *et al.*, 2021] for a recent overview). The actual WFOMC task then appeared in the literature a bit later [Van den Broeck *et al.*, 2011].

In this paper, we focus on tractability of the so-called *symmetric* weighted first-order model counting problem [Beame *et al.*, 2015] which is defined as follows.

Definition 1. *Let Δ be a set of domain elements and let $w(P)$ and $\overline{w}(P)$ be functions from predicate symbols to real numbers. Let HB denote the Herbrand base over the domain Δ w.r.t. a given first-order logic language $\mathcal{L}$. Let $\mathsf{pred} : \mathsf{HB} \mapsto \mathcal{P}$ map each atom to its predicate symbol. We define*

$$\mathrm{WFOMC}(\Psi, \Delta, w, \overline{w}) = \sum_{\omega \subseteq \mathsf{HB} : \omega \models \Psi} W(\omega, w, \overline{w}),$$

where the weight of ω is computed as

$$W(\omega, w, \overline{w}) = \prod_{l \in \omega} w(\mathsf{pred}(l)) \prod_{l \in \mathsf{HB} \setminus \omega} \overline{w}(\mathsf{pred}(l)).$$

That is, WFOMC is the sum of weights of all models of the sentence Ψ.

[3] Obviously, it is also possible to construct a fast sampler by hand without relying on WFOMS, but that would require effort and expertise, whereas one of the goals we wish to achieve with WFOMS is to free programmers/users from such tedious tasks by giving them a declarative framework for combinatorial sampling.

We also define first-order model count (FOMC) as a special case: $\text{FOMC}(\Psi, \Delta) = \text{WFOMC}(\Psi, \Delta, w_{\mathbf{1}}, w_{\mathbf{1}})$, *where* $w_{\mathbf{1}}(P) = 1$ *for all* P.

Example 5. *Let* $\Delta = \{A, B\}$, *let* $\mathcal{L}$ *consist of predicates heads, tails,* $w(\mathit{heads}) = 2$, $w(\mathit{tails}) = \overline{w}(\mathit{heads}) = \overline{w}(\mathit{tails}) = 1$, *and* $\Gamma = \forall x : (\mathit{heads}(x) \lor \mathit{tails}(x)) \land (\neg \mathit{heads}(x) \lor \neg \mathit{tails}(x))$. *There are four models of* Γ *on the domain* Δ: $\omega_1 = \{\mathit{heads}(A), \mathit{heads}(B)\}$, $\omega_2 = \{\mathit{heads}(A), \mathit{tails}(B)\}$, $\omega_3 = \{\mathit{tails}(A), \mathit{heads}(A)\}$ *and* $\omega_4 = \{\mathit{tails}(A), \mathit{tails}(B)\}$. *The resulting first-order model count is* $\text{FOMC}(\Gamma, \Delta) = 4$ *and the weighted model count is* $\text{WFOMC}(\Gamma, w, \overline{w}, \Delta) = 4 + 2 + 2 + 1 = 9$.

So far in this paper, we have talked about *tractability* of WFOMC without precisely defining what we mean by it. Since there is no hope to find an algorithm that would scale polynomially with the size of the given first-order logic sentences even for fragments such as $\mathbf{FO}^1$, which contains sentences with at most one logical variable (unles $\mathbf{P} = \mathbf{NP}$), most of the focus in this area has been on identifying fragments containing first-order logic sentences for which WFOMC can be computed in time polynomial in the size of the domain. The term coined for this kind of tractability by Van den Broeck [2011] is *domain liftability* and it is an analogue to data complexity used in database theory.

From Zero to $\mathbf{FO}^2$. One of the first breakthroughs in lifted inference for WFOMC came from the two seminal papers [Van den Broeck, 2011; Van den Broeck *et al.*, 2014] which together established domain liftability of $\mathbf{FO}^2$, which is the two-variable fragment of first-order logic. This was quickly complemented by a hardness result [Beame *et al.*, 2015], showing that WFOMC for the three-variable fragment $\mathbf{FO}^3$ is not domain-liftable (under plausible complexity-theoretic assumptions).

From $\mathbf{FO}^2$ to $\mathbf{C}^2$. While $\mathbf{FO}^3$ is not a domain-liftable fragment (under plausible assumptions) due to the negative results from [Beame *et al.*, 2015], it does not mean that the frontiers of tractability cannot be pushed beyond $\mathbf{FO}^2$. Recently, we showed that another fragment, called $\mathbf{C}^2$, is domain-liftable [Kuželka, 2021]. $\mathbf{C}^2$ is an interesting fragment, which had been previously studied in theoretical computer science [Graedel *et al.*, 1997]. It extends $\mathbf{FO}^2$ by allowing counting quantifiers of the form $\exists^{=k}$ ("exist exactly k"), $\exists^{\geq k}$ ("exist at least k") and $\exists^{\leq k}$ ("exist at most k") and is strictly more expressive than $\mathbf{FO}^2$. For instance, it is possible to encode permutations, derangements, k-regular graphs and many other structures in $\mathbf{C}^2$ (and then count them).

Beyond $\mathbf{C}^2$. Another strategy how to extend tractable fragments of first-order logic is to add extra axioms. This strategy was first pursued by Kuusisto and Lutz [2018] who showed that the fragment obtained by adding a single functionality axiom to $\mathbf{FO}^2$ is domain-liftable.[4] Later works extended domain liftability to the fragments $\mathbf{C}^2$ + $\mathbf{Tree}$ [Van Bremen and Kuželka, 2021] and $\mathbf{C}^2$ + $\mathbf{LinOrder}$ [Tóth and Kuželka, 2023], which are obtained by adding to $\mathbf{C}^2$ an axiom that specifies that a distinguished relation R should correspond to an undirected tree or to a linear order, respectively. Very recently, a new result [Malhotra and Serafini, 2023] added another such class $\mathbf{C}^2$ + $\mathbf{DAG}$, which allows restricting a distinguished binary relation to be a directed acyclic graph.

Other Fragments. Several other fragments, orthogonal to $\mathbf{C}^2$, were identified as tractable for WFOMC: $\mathbf{S}^2\mathbf{FO}^2$, $\mathbf{S}^2\mathbf{RU}$ [Kazemi *et al.*, 2016] and $\mathbf{U}_1$ [Kuusisto and Lutz, 2018]. Any of these fragments can also be extended by constraints on cardinalities of relations using techniques from [Kuželka, 2021].

Handling Evidence. We mentioned that there is, in general, no hope to obtain a WFOMC algorithm running in polynomial time in the size of the formula even for simple fragments such as $\mathbf{FO}^1$ or $\mathbf{FO}^2$. In Example 4, we constructed a formula whose size grows with the size of the domain and claimed that it will allow us to perform sampling in polynomial time. So what is going on here? It turns out that, e.g., for $\mathbf{FO}^2$, one can add arbitrary long conjunctions of ground unary atoms while still guaranteeing runtime polynomial in the domain size [Van den Broeck and Davis, 2012] and, in fact, one can also show that this holds for any tractable fragment that contains $\mathbf{C}^2$ (both for counting and sampling).

3 Tractable Sampling

While non-exact sampling methods such as Gibbs sampling have been studied and used a lot in statistical relational learning, exact first-order logic sampling has received significantly less attention. In the two recent papers [Wang *et al.*, 2022; Wang *et al.*, 2023], we introduced the exact weighted first-order model sampling problem (WFOMS): Given a sentence Ψ and two non-negative weighting functions w and $\overline{w}$, the WFOMS task is to sample a model ω of Ψ with probability proportional to its weight $W(\omega, w, \overline{w})$, which is defined as in Definition 1. We also introduced the notion of *domain-liftability under sampling* as a natural generalization of domain-liftability for WFOMC and showed that $\mathbf{FO}^2$ is domain-liftable under sampling and that this remains to hold even when one adds counting quantifiers and cardinality constraints [Wang *et al.*, 2023].

One might naively expect that WFOMS could be solved efficiently for all sentences with domain-liftable WFOMC by applying the classical reduction from *counting to sampling* [Jerrum *et al.*, 1986]. However, it turns out that such a strategy does not work in the lifted setting [Wang *et al.*, 2022] because of the complexity of conditioning in lifted inference [Van den Broeck and Davis, 2012]. Thus, even though we managed to prove domain-liftability under sampling for the two-variable fragment $\mathbf{FO}^2$, including counting quantifiers and cardinality constraints, which required a rather elaborate argument, it is currently not known whether there exists a general reduction showing that any fragment that is domain-liftable is also domain-liftable under sampling.

4 Applications

While lifted inference was originally conceived for inference problems in statistical relational learning, the range of pos-

[4] This fragment is contained in $\mathbf{C}^2$, which can encode an arbitrary number of functionality constraints.

sible applications of lifted-inference algorithms for WFOMC and WFOMS is broader. In this section we sketch some of them.

Combinatorics. Lifted-inference techniques are well-suited for solving problems from enumerative combinatorics. For instance, [Totis *et al.*, 2023] recently introduced a declarative domain-specific language which allows users to specify textbook-style combinatorial problems and then also solve them. What they focused on in this work are problems inspired by the so-called "Twelvefold way", promoted by [Stanley, 1986]. As they point out, there are also simple problems that cannot be solved by their approach. This includes problems with circular permutations, e.g., arranging people around a round table, or problems that involve permutations with relative position constraints. An example given by [Totis *et al.*, 2023] of this kind, which cannot be solved by their techniques, is: *"Nine chairs in a row are to be occupied by six students and Professors Alpha, Beta and Gamma. These three professors arrive before the six students and decide to choose their chairs so that each professor will be between two students. In how many ways can Professors Alpha, Beta and Gamma choose their chairs?"* Incidentally, this problem and similar ones, including those with circular permutations, can be solved by WFOMC in the fragment $\mathbf{C}^2 + \mathbf{LinOrder}$ [Tóth and Kuželka, 2023]. Moreover, the various tractable fragments of WFOMC can be used to count many other non-trivial structures as well, e.g., k-regular graphs, trees with k leaves etc. [Barvínek *et al.*, 2021; Van Bremen and Kuželka, 2021]. At the same time, the question whether all problems that can be solved by the approach of [Totis *et al.*, 2023] can also be solved by WFOMC is currently open. Finally, the close connection between WFOMC and enumerative combinatorics was used for developing a method capable of generating a database of integer sequences with combinatorial interpretation [Svatoš *et al.*, 2023] that we hope could complement the well known OEIS database [OEIS Foundation Inc., 2023].

Samplers. Standard programming language libraries typically provide some limited support for sampling problems. For instance, the NumPy package provides support for sampling simple structures such as permutations and combinations. However, when we need to sample just a bit more complex structure, e.g., an undirected graph without isolated vertices, we typically need to develop a sampler from scratch. We believe that having a declarative framework based on WFOMS would be a useful tool for programmers making their work easier (recall also Examples 1-4). What stops us from already building such a sampling framework based on the current WFOMS algorithms [Wang *et al.*, 2023] is their practical performance, which is the focus of ongoing works (however, it is worth recalling that these algorithms already outperform state-of-the-art propositional model samplers on first-order sampling problems, as shown experimentally in [Wang *et al.*, 2023], but there is still room for improvement).

Machine Learning. Lifted inference has primarily been studied with machine learning applications in mind. In particular, it has been shown that lifted-inference algorithms for WFOMC can be used to obtain polynomial-time algorithms for maximum-likelihood learning of Markov logic networks [Van Haaren *et al.*, 2016; Kuželka and Kungurtsev, 2019; Kuželka *et al.*, 2020].[5] The results on tractability of WFOMS are quite new, so there have not been any real applications yet, but we expect WFOMS to find use as a component in deep generative models such as variational autoencoders [Kingma and Welling, 2013] or as a proposal distribution in MCMC algorithms.

5 Other Related Works

Lifted inference is a broader area than what it may seem from the discussion in this paper, in particular, it does not focus only on the WFOMC problem. Many works in lifted inference focus on developing faster algorithms for factor graphs, mostly building on the variable elimination algorithm [de Salvo Braz *et al.*, 2005; Taghipour *et al.*, 2013; Braun and Möller, 2016]. Also, regarding WFOMC, a very closely related early work was [Gogate and Domingos, 2011]. Besides symmetric WFOMC that we discussed in this paper, there is also its asymmetric variant [Gribkoff *et al.*, 2014], which is less tractable than the symmetric one (even $\mathbf{FO}^2$ is intractable for it).

As for (W)FOMC, an interesting recent development appeared in the PhD thesis [Dilkas, 2023], which shows that *domain recursion* [Van den Broeck, 2011] can be extended in a way that allows finding symbolic solutions for some FOMC problems. Another interesting direction is the extension of lifted-inference techniques for *weighted first-order model integration* [Feldstein and Belle, 2021], which can be used for inference in probabilistic models that contain both discrete and continuous random variables.

There is also a sizeable body of literature on propositional model counting and sampling [Gomes *et al.*, 2021; Meel, 2022]. Even though it is theoretically possible to use propositional counters and samplers for tractable first-order sampling and counting problems, quite naturally, propositional counters and samplers do not match the performance of lifted algorithms on these problems (e.g. [Van den Broeck, 2011; Wang *et al.*, 2023]).

6 Conclusions and What Is Next?

In this paper we discussed recent works on lifted algorithms for weighted first-order model counting and sampling. We also described possible applications of these methods which seem to be on the horizon. So what should be next? We mention just a few things here. First, there is likely still scope for extending the tractability frontiers. Second, most works on lifted inference have so far focused on identifying fragments of first-order logic for which WFOMC (or WFOMS) can be computed in time polynomial in the domain size, but less attention has been given to the degree of these polynomials, which is as important if we want to use lifted inference for practical applications. Finally, is it the case that all (reasonable) fragments tractable for WFOMC are also tractable for WFOMS?

[5] Apart from lifted-inference, this relies on fundamental results about maximum-entropy distributions obtained surprisingly only recently [Singh and Vishnoi, 2014; Straszak and Vishnoi, 2019].

Acknowledgments

The author is grateful to the many collaborators who worked on the problems described in this paper (alphabetically): Jáchym Barvínek, Timothy van Bremen, Peter Jung, Vyacheslav Kungurtsev, Martin Svatoš, Jan Tóth, Yuanhong Wang and Yuyi Wang. The work of the author was supported by the Czech Science Foundation projects 20-19104Y and 23-07299S, and the OP VVV project *CZ.02.1.01/0.0/0.0/16_019/0000765* "Research Center for Informatics".

References

[Barvínek *et al.*, 2021] Jáchym Barvínek, Timothy van Bremen, Yuyi Wang, Filip Zelezný, and Ondrej Kuzelka. Automatic conjecturing of p-recursions using lifted inference. In *Inductive Logic Programming - 30th International Conference, ILP 2021, Proceedings*, volume 13191 of *Lecture Notes in Computer Science*, pages 17–25. Springer, 2021.

[Beame *et al.*, 2015] Paul Beame, Guy Van den Broeck, Eric Gribkoff, and Dan Suciu. Symmetric weighted first-order model counting. In *Proceedings of the 34th ACM SIGMOD-SIGACT-SIGAI Symposium on Principles of Database Systems*, pages 313–328. ACM, 2015.

[Braun and Möller, 2016] Tanya Braun and Ralf Möller. Lifted junction tree algorithm. In *Joint German/Austrian Conference on Artificial Intelligence (Künstliche Intelligenz)*, pages 30–42. Springer, 2016.

[de Salvo Braz *et al.*, 2005] Rodrigo de Salvo Braz, Eyal Amir, and Dan Roth. Lifted first-order probabilistic inference. In *IJCAI-05, Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, pages 1319–1325, 2005.

[Dilkas, 2023] Paulius Dilkas. *Generalising weighted model counting*. PhD thesis, University of Edinburgh, 2023.

[Feldstein and Belle, 2021] Jonathan Feldstein and Vaishak Belle. Lifted reasoning meets weighted model integration. In *Uncertainty in Artificial Intelligence*, pages 322–332. PMLR, 2021.

[Getoor and Taskar, 2007] Lise Getoor and Ben Taskar. *Introduction to statistical relational learning*, volume 1. MIT press Cambridge, 2007.

[Gogate and Domingos, 2011] Vibhav Gogate and Pedro M. Domingos. Probabilistic theorem proving. In *UAI 2011, Proceedings of the Twenty-Seventh Conference on Uncertainty in Artificial Intelligence, Barcelona, Spain, July 14-17, 2011*, pages 256–265, 2011.

[Gomes *et al.*, 2021] Carla P Gomes, Ashish Sabharwal, and Bart Selman. Model counting. In *Handbook of satisfiability*, pages 993–1014. IOS press, 2021.

[Graedel *et al.*, 1997] Erich Graedel, Martin Otto, and Eric Rosen. Two-variable logic with counting is decidable. In *Proceedings of Twelfth Annual IEEE Symposium on Logic in Computer Science*, pages 306–317. IEEE, 1997.

[Gribkoff *et al.*, 2014] Eric Gribkoff, Guy Van den Broeck, and Dan Suciu. Understanding the complexity of lifted inference and asymmetric weighted model counting. In *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence, UAI*, pages 280–289. AUAI Press, 2014.

[Jerrum *et al.*, 1986] Mark R Jerrum, Leslie G Valiant, and Vijay V Vazirani. Random generation of combinatorial structures from a uniform distribution. *Theoretical computer science*, 43:169–188, 1986.

[Kazemi *et al.*, 2016] Seyed Mehran Kazemi, Angelika Kimmig, Guy Van den Broeck, and David Poole. New liftable classes for first-order probabilistic inference. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems*, pages 3117–3125, 2016.

[Kingma and Welling, 2013] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

[Kuusisto and Lutz, 2018] Antti Kuusisto and Carsten Lutz. Weighted model counting beyond two-variable logic. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018*, pages 619–628, 2018.

[Kuželka and Kungurtsev, 2019] Ondřej Kuželka and Vyacheslav Kungurtsev. Lifted weight learning of Markov logic networks revisited. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pages 1753–1761. PMLR, 2019.

[Kuželka *et al.*, 2020] Ondřej Kuželka, Vyacheslav Kungurtsev, and Yuyi Wang. Lifted weight learning of Markov logic networks (revisited one more time). In *Int. Conf. on Probabilistic Graphical Models*, pages 269–280. PMLR, 2020.

[Kuželka, 2021] Ondřej Kuželka. Weighted first-order model counting in the two-variable fragment with counting quantifiers. *Journal of Artificial Intelligence Research*, 70:1281–1307, 2021.

[Malhotra and Serafini, 2023] Sagar Malhotra and Luciano Serafini. Weighted first order model counting with directed acyclic graph axioms. https://arxiv.org/abs/2302.09830, 2023.

[Meel, 2022] Kuldeep S. Meel. Counting, sampling, and synthesis: The quest for scalability. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI*, pages 5816–5820. ijcai.org, 2022.

[OEIS Foundation Inc., 2023] OEIS Foundation Inc. The on-line encyclopedia of integer sequences. http://oeis.org, 2023. Accessed: 2023-05-01.

[Poole, 2003] David Poole. First-order probabilistic inference. In *IJCAI-03, Proceedings of the Eighteenth International Joint Conference on Artificial Intelligence*, pages 985–991, 2003.

[Singh and Vishnoi, 2014] Mohit Singh and Nisheeth K Vishnoi. Entropy, optimization and counting. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing (STOC)*, pages 50–59. ACM, 2014.

[Stanley, 1986] Richard P Stanley. What is enumerative combinatorics? In *Enumerative combinatorics*, pages 1–63. Springer, 1986.

[Straszak and Vishnoi, 2019] Damian Straszak and Nisheeth K Vishnoi. Maximum entropy distributions: Bit complexity and stability. In *Conference on Learning Theory*, pages 2861–2891. PMLR, 2019.

[Svatoš *et al.*, 2023] Martin Svatoš, Peter Jung, Jan Tóth, Yuyi Wang, and Ondřej Kuželka. On discovering interesting combinatorial integer sequences. In *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence*, 2023.

[Taghipour *et al.*, 2013] Nima Taghipour, Daan Fierens, Jesse Davis, and Hendrik Blockeel. Lifted variable elimination: Decoupling the operators from the constraint language. *Journal of Artificial Intelligence Research*, 47:393–439, 2013.

[Totis *et al.*, 2023] Pietro Totis, Jesse Davis, Luc De Raedt, and Angelika Kimmig. Lifted reasoning for combinatorial counting. *Journal of Artificial Intelligence Research*, 76:1–58, 2023.

[Tóth and Kuželka, 2023] Jan Tóth and Ondřej Kuželka. Lifted inference with linear order axiom. In *Thirty-Seventh AAAI Conference on Artificial Intelligence*. AAAI Press, 2023.

[Van Bremen and Kuželka, 2021] Timothy Van Bremen and Ondřej Kuželka. Lifted inference with tree axioms. In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, pages 599–608, 2021.

[Van den Broeck and Davis, 2012] Guy Van den Broeck and Jesse Davis. Conditioning in first-order knowledge compilation and lifted probabilistic inference. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*. AAAI Press, 2012.

[Van den Broeck *et al.*, 2011] Guy Van den Broeck, Nima Taghipour, Wannes Meert, Jesse Davis, and Luc De Raedt. Lifted probabilistic inference by first-order knowledge compilation. In *Proceedings of the Twenty-Second international joint conference on Artificial Intelligence*, pages 2178–2185, 2011.

[Van den Broeck *et al.*, 2014] Guy Van den Broeck, Wannes Meert, and Adnan Darwiche. Skolemization for weighted first-order model counting. In *Proceedings of the 14th International Conference on Principles of Knowledge Representation and Reasoning (KR)*, pages 1–10, 2014.

[Van den Broeck *et al.*, 2021] Guy Van den Broeck, Kristian Kersting, Sriraam Natarajan, and David Poole. *An Introduction to Lifted Probabilistic Inference*. MIT Press, aug 2021.

[Van den Broeck, 2011] Guy Van den Broeck. On the completeness of first-order knowledge compilation for lifted probabilistic inference. In *Advances in Neural Information Processing Systems*, pages 1386–1394, 2011.

[Van Haaren *et al.*, 2016] Jan Van Haaren, Guy Van den Broeck, Wannes Meert, and Jesse Davis. Lifted generative learning of markov logic networks. *Machine Learning*, 103(1):27–55, 2016.

[Wang *et al.*, 2022] Yuanhong Wang, Timothy van Bremen, Yuyi Wang, and Ondřej Kuželka. Domain-lifted sampling for universal two-variable logic and extensions. In *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence (AAAI-22)*. AAAI Press, 2022.

[Wang *et al.*, 2023] Yuanhong Wang, Juhua Pu, Yuyi Wang, and Ondřej Kuželka. On exact sampling in the two-variable fragment of first-order logic. In *LICS '23: 37th Annual ACM/IEEE Symposium on Logic in Computer Science*. IEEE, 2023.