



THE UNIVERSITY *of* EDINBURGH

Edinburgh Research Explorer

Satisfiability and Model Counting in Open Universes

Citation for published version:

Belle, V 2016, Satisfiability and Model Counting in Open Universes. in *The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence Beyond NP: Technical Report WS-16-05*. AAAI Press, Thirtieth AAAI Conference on Artificial Intelligence Beyond NP, Phoenix, Arizona, United States, 12/02/16.
<<http://www.aaai.org/ocs/index.php/WS/AAAIW16/paper/view/12605>>

Link:

[Link to publication record in Edinburgh Research Explorer](#)

Document Version:

Publisher's PDF, also known as Version of record

Published In:

The Workshops of the Thirtieth AAAI Conference on Artificial Intelligence Beyond NP: Technical Report WS-16-05

General rights

Copyright for the publications made accessible via the Edinburgh Research Explorer is retained by the author(s) and / or other copyright owners and it is a condition of accessing these publications that users recognise and abide by the legal requirements associated with these rights.

Take down policy

The University of Edinburgh has made every reasonable effort to ensure that Edinburgh Research Explorer content complies with UK legislation. If you believe that the public display of this file breaches copyright please contact openaccess@ed.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.



Satisfiability and Model Counting in Open Universes

Vaishak Belle

KU Leuven

Belgium

vaishak@cs.kuleuven.be

Abstract

SAT and #SAT are at the heart of many important problem formulations in AI, the most prominent being reasoning and learning in first-order and probabilistic knowledge bases. In practice, all contemporary systems resort to domain closure: objects in the universe are all and only the ones mentioned in the knowledge base. This is in stark contrast to the natural ability of human beings to infer things about sensory inputs and unforeseen data: they infer the existence of objects from their observations; no predefined list of objects is given or known in advance. In this paper, we introduce the formal foundations for reasoning in open universes in a general way, purely based on SAT and #SAT technology.

1 Introduction

SAT is the problem of finding a satisfying assignment for a propositional formula, and model counting (or #SAT) is the problem of computing the number of satisfying assignments. SAT has innumerable applications, including reasoning in first-order knowledge bases (Mitchell and Ternovska 2005). #SAT generalizes SAT and is the canonical #P-complete problem. Although #SAT is known to be more computationally challenging than SAT, it admits a natural encoding for many counting problems in various guises, the most prominent being probabilistic reasoning and learning in graphical models and their lifted versions (Chavira and Darwiche 2008; Van den Broeck 2013). Most significantly, the simple semantics of propositional languages allows practitioners to enable and study useful features such as hard constraints, variable ordering, random restarts, component analysis and approximations – leading to solvers that are almost always competitive and sometimes outperform specialized algorithms (Bacchus, Dalmao, and Pitassi 2009; Gomes, Sabharwal, and Selman 2009). In the recent years, the maturity of SAT and #SAT solvers has encouraged the reduction of many non-propositional problems (Morgado et al. 2006; Chistikov, Dimitrova, and Majumdar 2015; Belle, Passerini, and Van den Broeck 2015) to SAT, either by an encoding or by interfacing with a higher-order solver, such as SMT.

A fundamental assumption in many contemporary AI systems is that the set of objects in the universe finite and

known. In contrast, a long-standing goal in AI has been to mimic the natural ability of human beings to infer things about sensory inputs and unforeseen data (Milch et al. 2005), often referred to as the *open universe* setting. That is to say, humans infer the existence of objects from their observations; no predefined list of objects is given or known in advance. To be concrete, consider (a categorical version of) the infamous Friends-Smokers example:

$$\forall x, y \text{ Smoker}(x) \wedge \text{Friends}(x, y) \supset \text{Smoker}(y)$$

which says that friends of smokers are also smokers, together with the following sentences:

- $\text{Smoker}(\text{john})$
- $\forall x \text{ Friends}(\text{john}, x)$

which says that John is a smoker, and perhaps describes John’s social circle and declares that John is friends with everybody. In that regard, almost all first-order (probabilistic) frameworks (Richardson and Domingos 2006) resort to *domain closure*: objects in the universe are all and only the ones mentioned in the knowledge base. In particular, if evidence is now provided that a new person Jane is rich, queries about Jane’s smoking habits are undefined because she does not appear anywhere in the knowledge base. (Of course, by logical deduction, she is a smoker because everybody, including Jane, is John’s friend from an existing declaration.) As argued in (Milch et al. 2005), many AI systems engineer away or preprocess unknown individuals in an ad hoc fashion. To alleviate this situation, the language BLOG (Milch et al. 2005), among many others (*cf.* penultimate section), have been proposed, but their semantics is specialized for the syntactical constructs they allow, making it difficult to adapt or extend them to other contexts.

Building on foundational work in the area of knowledge representation (Levesque 1998), in this paper, we propose an effective and general formulation for reasoning (satisfiability, validity, model counting) with unknown individuals in first-order knowledge bases, purely based on SAT and #SAT technology. (For simplicity, we focus on unweighted specifications: that is, no probabilities.) As the first feature, our approach will handle examples about inferred individuals and their smoking habits in a principled manner. As a second feature, the approach assumes a standard propositional reasoner (together with a grounder). To appreciate how this

latter feature is approached, consider the following alternative for unknowns: fix an extremely large set of constants in advance; since an agent has a finite life, it will only encounter finitely many individuals during its operation, and so, at any moment if it were to encounter a new individual, assign this individual’s identity to an unused constant. While such a proposal would work, it requires us to initially ground the knowledge base wrt all constants, which is very expensive. To counter this, a central property of our approach is that any given time, grounding only needs to consider the known individuals and a small number of unknowns. In the long term, our proposal can benefit from existing translations of combinatorial problems to SAT and #SAT, and in this regard, we believe this work would also provide the basis for exploring further questions in AI and machine learning over open universes in a principled way.

We are organized as follows. We begin with the logical preliminaries, and first consider satisfiability and validity in open universes. Next, we turn to how model counting should be defined in open universes. Finally, we discuss related work and conclude.

2 Preliminaries

Logical Background

We let \mathcal{L} be a classical propositional language over the vocabulary $Q = \{p, q, \dots\}$. A \mathcal{L} -model M is a $\{0, 1\}$ assignment to the elements of Q . For any $\phi \in \mathcal{L}$, we write $M \models \phi$ to mean that ϕ is *satisfied* at M , defined inductively as usual (Smullyan 1995). We write $Q(\phi)$ to refer to the propositions mentioned in ϕ . By extension, we write $\mathcal{L}(\phi)$ to mean the (finite) language built from $Q(\phi)$, and write $\mathcal{M}(\phi)$ to refer to the set of $\mathcal{L}(\phi)$ -models where ϕ is satisfied. (That is, $\mathcal{L}(\phi)$ -models are $\{0, 1\}$ assignments to the elements of $Q(\phi)$.)

We let \mathcal{L}^+ be a first-order language with equality, relational symbols $\{P(x), Q(x, y), R(x, y, z), \dots, P'(x), \dots\}$ of every arity, variables $\{x, y, z, \dots\}$, and a countably infinite set of *parameters* $\mathcal{U} = \{\#1, \#2, \dots\}$ (Levesque 1984) serving as the domain of discourse for quantification. (For presentation, we often use identifiers such as Latin alphabets (a, b, \dots) and proper names $(john, jane, \dots)$ as parameters.) Together with equality, parameters essentially realize an infinitary version of the unique name assumption.¹ The set of (ground) atoms is obtained as:

$$Q^+ = \{P(a_1, \dots, a_k) \mid P \text{ is a predicate, } a_i \text{ a parameter}\}.$$

A \mathcal{L}^+ -model M is a $\{0, 1\}$ assignment to the elements of Q^+ . The semantics for $\phi \in \mathcal{L}^+$ is defined as usual with equality as identity:

- $M \models p$ iff $p \in Q^+$ and $M[p] = 1$;
- $M \models \neg\phi$ iff $M \not\models \phi$;

¹Such a feature can also be enabled in ordinary first-order logic by following the recipe in (Levesque 1998). Note that this does not rule out capturing uncertainty about the identity of objects, as needed, for example, in (Srivastava et al. 2014), which can be enabled in a first-order language using complex terms. (Indeed, (Srivastava et al. 2014) also resort to the device of rigid designators in the underlying logic.) We ignore these concerns for simplicity.

- $M \models \phi_1 \wedge \phi_2$ iff $M \models \phi_i$ for $i \in \{1, 2\}$;
- $M \models \forall x\phi(x)$ iff $M \models \phi(a)$ for all $a \in \mathcal{U}$;
- $M \models (a = b)$ iff a and b are the same parameters.

We say that ϕ is *valid* iff for every \mathcal{L}^+ -model M , $M \models \phi$.

It is worth noting that when first-order logic is restricted to a class of models with a fixed countably infinite domain, the *compactness* property does not hold (Smullyan 1995). For example, the theory (Levesque 1998):

$$\{\exists xP(x), \neg P(\#1), \neg P(\#2), \dots\}$$

is an unsatisfiable theory whose every finite subset is indeed satisfiable.

Knowledge Bases

Since reasoning in full first-order logic is undecidable, it is natural to assume that the knowledge bases are restricted to a specified form. Our intent here is to capture (and go beyond) existing AI examples, such as the Friends-Smokers domain. To that end, we consider an expressive first-order fragment called *proper⁺* knowledge bases (Lakemeyer and Levesque 2002), equivalent to a (possibly infinite) set of ground clauses.

In the sequel, we use e to range over *ewffs*, that is, quantifier-free formulas whose only predicate is equality. For example, $x \neq john$ and $(x \neq john \wedge x \neq jane)$ are ewffs. Moreover, by a clause we simply mean a quantifier-free disjunction of (possibly non-ground) atoms. For example, $Human(john)$ and $Human(x) \vee Vehicle(x)$ are clauses. We use c to range over clauses. Finally, we write $\forall\phi$ to mean the universal closure of ϕ . For example, $\forall(Friends(x, y) \vee Brothers(x, y))$ denotes $\forall x, y. Friends(x, y) \vee Brothers(x, y)$.

Definition 1: A formula of the form $\forall(e \supset c)$ is called a \forall -clause. A knowledge base ϕ is *proper⁺* if it is a finite non-empty set of \forall -clauses.

Here are some examples for the Friends-Smokers domain:

- $Friends(john, jane) \supset \neg Smoker(john)$ says that if John is Jane’s friend, then he is not a smoker;
- $\forall(\neg Smoker(x) \vee \neg Friends(x, y) \vee Smoker(y))$ captures, as usual, that friends of smokers are smokers themselves;
- $\forall(x \neq john \supset Smoker(x))$ says that everybody other than John is a smoker;
- $\forall((x \neq john \wedge x \neq jane) \supset \neg Smoker(x))$ says that everybody other than John and Jane are not smokers.
- $\forall(e \supset Smoker(x)) \wedge \forall(\neg e \supset \neg Smoker(x))$, where $e = (x = john \vee x = jane)$ says that John and Jane are the only ones in the universe who smoke, thus enabling the closed world assumption.

In the context of logical and probabilistic reasoning, it is often the case that one provides evidence and queries. We assume that both evidence and queries are quantifier and variable-free formulas (*i.e.*, ground atoms + Boolean connectives). With this assumption, given a *proper⁺* KB ϕ , and evidence or query α , clearly $\phi \wedge (\neg)\alpha$ is also a *proper⁺* KB.²

²In the case of entailment, we check whether $\phi \wedge \neg\alpha$ is unsatisfiable. When α denotes evidence or query in probabilistic contexts, we are often interested in the model count of $\phi \wedge \alpha$.

3 Satisfiability and Validity

Before addressing model counting, we first begin with the fundamental tasks of automating satisfiability and validity checking in \mathcal{L}^+ .³ We use the following notation in our results:

- θ denotes a substitution of variables with parameters, as applicable to ewffs and clauses. For any set of parameters $C \subseteq \mathcal{U}$, we write $\theta \in C$ to mean substitutions are only allowed wrt the parameters in C .

For example, given $\theta = \langle x \mapsto \text{john}, y \mapsto \text{jane} \rangle$, ewff $e = x \neq \text{bob}$, and clause $c = \text{Friends}(x, y) \vee \text{Siblings}(x, y)$, we have $e\theta = \text{john} \neq \text{bob}$ and $c\theta = \text{Friends}(\text{john}, \text{jane}) \vee \text{Siblings}(\text{john}, \text{jane})$. If $C = \{\text{john}, \text{bob}\}$, clearly $\theta \notin C$. Instead, $\theta' = \langle x \mapsto \text{john}, y \mapsto \text{bob} \rangle$ is one such that $\theta' \in C$.

- $\text{gnd}(\phi) = \{c\theta \mid \forall (e \supset c) \in \phi \text{ and } \models e\theta\}$.

Recall that because equality is understood as identity for parameters, things $\text{john} \neq \text{jane}$ are valid, and things like $\text{john} \neq \text{john}$ are unsatisfiable. So, given $\phi = \forall (x \neq \text{john} \supset \text{Smoker}(x))$, we have $\text{gnd}(\phi) = \{\text{Smoker}(\text{jane}), \text{Smoker}(\text{bob}), \dots\}$. More precisely $\text{gnd}(\phi) = \{\text{Smoker}(a) \mid a \neq \text{john} \text{ is a parameter}\}$.

- For $k \geq 0$, $\text{gnd}(\phi, k) = \{c\theta \mid \forall (e \supset c) \in \phi, \models e\theta, \theta \in K\}$, where K is the set of parameters mentioned in ϕ plus k (arbitrary) new ones.

For example, if $\phi = \{\forall (\text{Smoker}(x)), \text{Smoker}(\text{john})\}$, then $\text{gnd}(\phi, 1) = \{\text{Smoker}(\text{john}), \text{Smoker}(\text{jane})\}$. (Note that the parameter jane is picked arbitrarily.) Also, $\text{gnd}(\phi, 0) = \{\text{Smoker}(\text{john})\}$. Instead, if $\phi' = \forall (\text{Smoker}(x))$, then $\text{gnd}(\phi', 0)$ is simply the empty set.

- Suppose ϕ is proper⁺ such that the maximum number of variables mentioned in a \forall -clause is k . We define $\text{gnd}(\phi, *)$ to mean $\text{gnd}(\phi, k)$. For example, if $\phi = \forall (\text{Smoker}(x))$ then $\text{gnd}(\phi, *) = \text{gnd}(\phi, 1)$ and if $\phi = \forall (\text{Smoker}(x)) \wedge \forall (\text{Friends}(x, z))$ then $\text{gnd}(\phi, *) = \text{gnd}(\phi, 2)$.

Proposition 2: For any proper⁺ KB ϕ , $\text{gnd}(\phi, *)$ is finite.

The first hurdle is addressing the compactness property, or the lack thereof. Such a property is useful to draw conclusions about a possibly infinite theory in terms of its finite subsets. As mentioned, \mathcal{L}^+ does not enjoy this property, but fortunately, proper⁺ knowledge bases do. For our first main result, we have:

Theorem 3: Suppose ϕ is a proper⁺ KB. Then $S = \text{gnd}(\phi)$ is satisfiable iff every finite subset of S is satisfiable.

The reason why this theorem holds is because $\text{gnd}(\phi)$ is essentially a (possibly infinite) propositional theory over a (possibly infinite) vocabulary. Compactness from propositional logic is then imported.

The second main result of this section is to then address logical consequence for proper⁺ KBs in a finitary manner:

Theorem 4: Suppose ϕ is a proper⁺ KB, and suppose α is a quantifier and variable-free formula (i.e., ground atoms + Boolean connectives). Then

$$\models \phi \supset \alpha \text{ iff } \text{gnd}(\phi \wedge \neg\alpha, *) \text{ is unsatisfiable.}$$

³The investigations in this section simplify and improve on results in (Belle and Lakemeyer 2011).

This theorem rests on an important observation made in (Levesque 1998) and elsewhere (Levesque and Lakemeyer 2001) that parameters not mentioned in the theory behave in a similar fashion. Of course, Proposition 2, Theorem 3 and our semantic setup can be used to further show:

Corollary 5: $\text{gnd}(\phi, *)$ is satisfiable (or unsatisfiable) in \mathcal{L}^+ iff it is satisfiable (or unsatisfiable resp.) in classical (finitary) propositional logic.

In sum, with these results, after grounding the knowledge base wrt a few extra parameters, an ordinary SAT solver can be used.

To see Theorem 4 in action, suppose ϕ is the union of the following:

- $\forall (\text{Smoker}(x) \wedge \text{Friends}(x, y) \supset \text{Smoker}(y))$
- $\text{Smoker}(\text{john})$
- $\forall y (\text{Friends}(\text{john}, y))$

Suppose the query $\alpha = \text{Smoker}(\text{jane})$. Then $\text{gnd}(\phi \wedge \neg\alpha, *) = \text{gnd}(\phi \wedge \neg\alpha, 2)$, i.e., we are to consider the parameters in $\phi \wedge \neg\alpha$ plus 2 new ones, say $C = \{\text{jane}, \text{john}, \text{bob}, \text{rob}\}$. On applying all possible substitutions $\theta \in C$, $\text{gnd}(\phi \wedge \neg\alpha, *)$ would include, among other things:

- $\text{Smoker}(\text{john}) \wedge \text{Friends}(\text{john}, \text{jane}) \supset \text{Smoker}(\text{jane})$
- $\text{Smoker}(\text{john})$
- $\text{Friends}(\text{john}, \text{jane})$

together with $\neg\alpha = \neg\text{Smoker}(\text{jane})$. Clearly $\text{gnd}(\phi \wedge \neg\alpha, *)$ is unsatisfiable, and so $\models \phi \supset \alpha$.

In the above example, bob and rob could have been done away with, and so one might wonder why these extra parameters are needed for grounding. To see that, consider ϕ' as the union of:

- $\forall (x \neq \text{john} \supset \text{Smoker}(x))$
- $\forall (\text{Friends}(x, y))$
- $\forall (\text{Smoker}(x) \wedge \text{Friends}(x, y) \supset \text{Smoker}(y))$

That is, all we know is that everyone other than John is a smoker, everybody is friends with everybody else, and the usual Friends-Smokers claim. In fact, by deduction, $\models \phi' \supset \text{Smoker}(\text{john})$. This is because some person other than John, say Jane, is declared to be a smoker and since she is friends with John, he must also be a smoker. However, grounding ϕ' wrt the only parameter mentioned in ϕ' , namely john , would not lead to a propositional theory where $\text{Smoker}(\text{john})$ can be deduced. Instead, consider $\text{gnd}(\phi', *) = \text{gnd}(\phi', 2)$ and parameters $\{\text{john}, \text{jane}, \text{bob}\}$ does lead to a grounding where $\text{Smoker}(\text{john})$ is entailed, as desired.

4 Model Counting

We now finally turn to model counting. The key question we attempt to answer is this: how should model counting be defined in an open universe setting?

Consider the classical definition (Gomes, Sabharwal, and Selman 2009). Let ϕ be a propositional formula. Then the model count of ϕ , written $\#\phi$, is defined as:

$$\#\phi = |\mathcal{M}(\phi)|$$

Recall that, even though \mathcal{L} is infinite, we define $\mathcal{M}(\phi)$ to be the set of satisfying interpretations of ϕ , which are $\{0, 1\}$ assignments to $\mathcal{Q}(\phi)$ where ϕ is true.⁴

In an open universe setting, given a proper⁺ KB ϕ and query α , which possibly mentions hereto unknown individuals, a reasonable definition for the model count of the proper⁺ KB $\psi = \phi \wedge \alpha$ is:

$$\begin{aligned} \#\psi &= \#gnd(\psi, *) \\ &= |\mathcal{M}(gnd(\psi, *))| \end{aligned}$$

(Note that the RHS is a propositional formula as usual.) This definition is intuitive in the sense that it accounts for all the individuals currently encountered and a few arbitrary unknowns corresponding to the quantifier rank of a \forall -clause in ψ . On the one hand, note that if ψ does not mention any \forall -clauses, *i.e.*, ψ is a ground formula, the definition is equivalent to the classical one. On the other, if ψ does mention \forall -clauses but is unsatisfiable, from Theorem 4, $\#\psi = 0$ as desired.

To see this definition in action, let ϕ be the union of:

- $\forall(\text{Smoker}(x) \vee \text{Alcoholic}(x))$
- $\forall(x \neq \text{john} \supset \neg \text{Smoker}(x))$

which says that everybody is a smoker or an alcoholic, and everybody other than John is not a smoker. Suppose we are interested in contrasting two queries in terms of their counts: let $\alpha_1 = \text{Alcoholic}(\text{jane})$ and $\alpha_2 = \text{Smoker}(\text{jane})$. Consider that $gnd(\phi \wedge \alpha_1, *)$ is the conjunction of:

- $(\text{Smoker}(\text{john}) \vee \text{Alcoholic}(\text{john})) \wedge (\text{Smoker}(\text{jane}) \vee \text{Alcoholic}(\text{jane})) \wedge (\text{Smoker}(\text{bob}) \vee \text{Alcoholic}(\text{bob}))$
- $\neg \text{Smoker}(\text{jane}) \wedge \neg \text{Smoker}(\text{bob})$
- $\text{Alcoholic}(\text{jane})$

where the grounding is wrt the parameters in $\phi \wedge \alpha_1$, namely $\{\text{john}, \text{jane}\}$ plus one extra one (because ϕ mentions one variable), namely bob . Clearly models can vary in how they interpret $\text{Smoker}(\text{john}) \vee \text{Alcoholic}(\text{john})$, and so we obtain a count of 3. However, $\phi \wedge \alpha_2$ is inconsistent, and indeed $\#gnd(\phi \wedge \alpha_2, *) = 0$.

By extension, given a proper⁺ KB ϕ and a query α , the probability of the query given the KB is defined as:

$$\text{PR}_\phi(\alpha) = \frac{\#gnd(\phi \wedge \alpha, *)}{\#gnd(\phi \wedge \alpha, *) + \#gnd(\phi \wedge \neg\alpha, *)}$$

It is easy to see that if $\models \phi \supset \alpha$, then $\phi \wedge \neg\alpha$ is not satisfiable, and so $\#gnd(\phi \wedge \neg\alpha, *)$ must be 0. Therefore, $\text{PR}_\phi(\alpha)$ is 1 and $\text{PR}_\phi(\neg\alpha)$ is 0.

For example, letting ϕ be the sentences above saying that everyone is a smoker or an alcoholic, and everybody other than John is not a smoker, we see that $\text{PR}_\phi(\text{Smoker}(\text{jane})) = 0$ and $\text{PR}_\phi(\neg \text{Smoker}(\text{jane})) = \text{PR}_\phi(\text{Alcoholic}(\text{jane})) = 1$.

More generally, for any new parameter $a \neq \text{john}$ that has not been hereto encountered, we obtain $\text{PR}_\phi(\text{Alcoholic}(a)) = 1$, as one would expect.

⁴In certain situations, given a finite set of propositions $\mathcal{Q}' \subset \mathcal{Q}$, we might be interested in the model count of a formula ϕ wrt all the propositions in \mathcal{Q}' and not just the ones in ϕ . Given our definition of $\#\phi$, this is achieved as follows: let \mathcal{Q}'' be the propositions in \mathcal{Q}' but not in ϕ . Let $\gamma = \{p \vee \neg p \mid p \in \mathcal{Q}''\}$. Clearly, $\phi \equiv \phi \wedge \gamma$, and using our construction, we would compute $\#(\phi \wedge \gamma)$.

Strong Model Counting

The above notion considers the situation of knowing some individuals, encountering unknowns via observations, and making claims thereof. It is natural to investigate a generalization wrt the *entire* domain of discourse: can anything be said about the count of one query, regardless of individuals that will be encountered in the future? To that end, we introduce the notion of *strong model counting* for a proper⁺ KB ϕ , denoted $\#\#\phi$, defined as:

$$\#\#\phi = \lim_{n \rightarrow \infty} \#gnd(\phi, n)$$

We argue for the reasonableness of this definition using mathematical intuition and examples:

- Suppose ϕ is any finite set of ground clauses. Clearly, for any $k \geq 0$, $gnd(\phi, k) = \phi$ and so $\#\#\phi = \#\phi$ as usual.
- Suppose $\phi = \forall(\text{Smoker}(x))$. Then $gnd(\phi, 1) = \{\text{Smoker}(\text{john})\}$, and for the language consisting of a single proposition $\text{Smoker}(\text{john})$ there are two interpretations, one where the proposition is false and another where it is true. So $|\mathcal{M}(\text{Smoker}(\text{john}))| = 1$. Going further, $gnd(\phi, 2) = \{\text{Smoker}(\text{john}), \text{Smoker}(\text{jane})\}$. While there are 4 interpretations for a language with 2 propositions, clearly $|\mathcal{M}(\text{Smoker}(\text{john}) \wedge \text{Smoker}(\text{jane}))| = 1$. Thus, for any k , $|\mathcal{M}(gnd(\phi, k))| = 1$, and so $\#\#\phi = 1$.
- Suppose $\phi = \forall(\text{Smoker}(x) \vee \text{Alcoholic}(x))$. Then $gnd(\phi, 1) = \{\text{Smoker}(\text{john}) \vee \text{Alcoholic}(\text{john})\}$, which has 3 models. By extension, $gnd(\phi, k) = 3^k$ models. Thus, $\#\#\phi = \lim_{n \rightarrow \infty} 3^n$, which is not bounded.

Along the lines of the Smoker-Alcoholic example, $\#\#\phi$ will be uninteresting for most interesting knowledge bases. Despite this drawback, we argue that this definition makes sense from a mathematical point of view, and the fact that counts can be unbounded is natural and not at all surprising. So where do we go from here?

Consider the generalization to PR for the case of strong model counting:

$$\text{PRLIM}_\phi(\alpha) = \lim_{n \rightarrow \infty} \frac{\#gnd(\phi \wedge \alpha, n)}{\#gnd(\phi \wedge \alpha, n) + \#gnd(\phi \wedge \neg\alpha, n)}$$

which says that the limit is considered for the ratio of the model counts for $\phi \wedge \alpha$ versus for that of $\phi \wedge \alpha$ and $\phi \wedge \neg\alpha$.

To see a simple example, consider the query $\alpha = \text{Alcoholic}(\text{jane})$, and let the KB be $\phi = \forall(\text{Smoker}(x) \vee \text{Alcoholic}(x)) \wedge \neg \text{Smoker}(\text{jane})$. Then $\#gnd(\phi \wedge \alpha, *) = 3$, and for any $n \geq 1$, we observe that $\#gnd(\phi \wedge \alpha, n) = 3^n$. Of course, $\#gnd(\phi \wedge \neg\alpha, n) = 0$. In that vein, we observe that

$$\text{PRLIM}_\phi(\alpha) = \lim_{n \rightarrow \infty} \frac{3^n}{3^n} = 1.$$

Naturally, $\text{PRLIM}_\phi(\neg\alpha) = 0$.

Thus, when strong model counting is studied in terms of PRLIM, the limit of the ratio allows us to better understand the probabilities of one query against another.

5 Discussion

Our account of $\#\phi$ was motivated by our account of satisfiability in countably infinite domains, all of which was achieved using a few parameters outside those mentioned in ϕ . Our account of $\#\#\phi$ was an attempt to understand model counting wrt the set of all parameters. Finally, when counting was considered as a ratio, we observed that PR and PRLIM are 1 for queries that are entailed, and 0 for those that are inconsistent with the knowledge base.

Although all of this was investigated in the context of grounding a proper⁺ KB, $\#\#\phi$ and PRLIM are closely related to a classical study in model theory called *0-1 laws*. Given a logical language \mathcal{L} , a class of \mathcal{L} -structures \mathcal{M} , a property of structures P , let $\mu_n(P)$ denote the fraction of \mathcal{L} -structures of domain size n that satisfy P . The logic is said to have the 0-1 law if $\mu(P) \doteq \lim_{n \rightarrow \infty} \mu_n(P)$ is either 0 or 1. Fagin (1976) proved a 0-1 law for first-order logic. In recent work, Beame et al. (2015) revisit this from the viewpoint of *first-order model counting*. The first-order model counting problem is this: given a first-order formula ϕ , and a number $n \in \mathbb{N}$, compute the number of first-order structures of size n (up to isomorphism) where ϕ holds. Basically, if a closed-form algebraic expression could be obtained for ϕ and n , then simplifications using calculus alone would yield $\mu(\phi)$. For example, for a language with two unary predicates $Smoker(x)$ and $Alcoholic(x)$, the first-order model count for $\phi = \forall(Smoker(x) \vee Alcoholic(x))$ is 3^n , and so, $\mu(\phi) = \lim_{n \rightarrow \infty} 3^n/4^n = 0$. Unfortunately, Beame et al. (2015) show that such a program is not possible in general: they obtain a first-order formula whose first-order model counting problem is $\#P_1$ -complete (Valiant 1979).

Putting it all together, we make the following observations, which we leave as directions for future work:

- The proof for the $\#P_1$ -completeness of the first-order model counting problem in (Beame et al. 2015) is based on an encoding of a Turing machine that computes a $\#P_1$ -hard function as first-order sentences. Without going into the details, the syntax of proper⁺ KBs, however, does not permit expressing some of these sentences in an obvious way (*e.g.*, existential quantifiers from the outside are not allowed). Thus, does the hardness result also apply to proper⁺ KBs?
- Extending the previous question, can computing $PR_\phi(\alpha)$, which is easy, inform us about $PRLIM_\phi(\alpha)$?

6 Related Work

This paper is concerned with extending model counting to open universes in a general way, best represented by the Friends-Smokers-Jane example in Section 1. At the outset, we remark that although there is a large body of literature on open universes, none have attempted to formulate this in terms of SAT and #SAT. We divide the literature into two broad camps: the logic camp, which includes knowledge representation and databases, and the probability camp, which includes open-universe Bayesian models.

In the logic circles, as we mentioned, our work is based on foundational work in knowledge representation (Levesque 1998), and its derivatives, including (Lakemeyer

and Levesque 2002), where proper⁺ KB were first introduced, and (Belle and Lakemeyer 2011), where a result closely related to Theorem 4 was considered. However, both (Levesque 1998) and (Lakemeyer and Levesque 2002) were concerned with sound but incomplete deductive reasoning. Model counting, SAT and #SAT formulations were not considered in any of these. We also remark that while there is extensive literature on decidable classes of first-order fragments (Lewis 1980), results such as Theorem 3 are somewhat different in flavor because they are simply establishing compactness in the presence of a countably infinite domain.

In the related area of database theory, handling incomplete information has been a longstanding goal (Van der Meyden 1998), the most prominent approach being that of *labelled null values*, which is connected to the problem of query containment (Chandra and Merlin 1977). Basically, nulls can be seen as Skolem constants, that is, these are objects in the universe whose identity is unknown. Nonetheless, the universe is assumed to be finite and fixed in database theory. A treatment of nulls in the context of (Levesque 1998) appears in (Giacomo, Lespérance, and Levesque 2011).

In the related area of automated verification, SMT (Barrett et al. 2009) investigates the satisfiability of (usually) quantifier-free decidable first-order theories, such as linear arithmetic. Although there is a lot of recent interest in model counting for SMT (Ma, Liu, and Zhang 2009; Chistikov, Dimitrova, and Majumdar 2015; Belle, Passerini, and Van den Broeck 2015), the focus is on formulas of the form $x \leq y + 3$, where \leq is a distinguished predicate, and $+$ is a distinguished function that are interpreted in the usual way. Thus, it is not clear if any of these techniques can be borrowed for examples such as Friends-Smokers with domain-specific predicates and quantification. Nonetheless, there is some research on obtaining algebraic expressions for model counts in SMT (see, for example, (Luu et al. 2014)), and so investigating these connections is worthwhile.

In the probability and machine learning circles, in practice, it is most often that probabilistic graphical models and their lifted versions make the finite domain closure assumption (Richardson and Domingos 2006; Van den Broeck 2013). Nonetheless, handling unknowns has been an important concern, perhaps best demonstrated in the early work of (Charniak and Goldman 1993) on unbounded objects for hypothesized plans. More generally, a common approach in many probabilistic models is to admit *number uncertainty*, where an explicit generating function determines (at run time) the number of objects, and the modeling language allows such a function together with statements made about these instantiated objects. This is reflected in, for example, BLOG (Milch et al. 2005) and distributional clauses (Gutmann et al. 2011). Closely related is the approach taken in (Laskey 2008), where the existence of named objects is determined by a distinguished identity function. In similar spirit, the existence (and number) of objects can be determined by explicit logical constraints in (Koller and Pfeffer 1998; Pasula et al. 2002; Getoor et al. 2001). A slightly more involved approach is to admit referential uncertainty, as in (Srivastava et al. 2014), where different logical terms might denote the same ob-

ject in the universe; see (Milch et al. 2005; Srivastava et al. 2014) for comprehensive discussions. At the other extreme, approaches such as (Singla and Domingos 2007; Doshi et al. 2011) exploit a locality property over an infinite set of variables, in the sense that all variables other than the ones in consideration are equivalent (in some appropriate formal sense). In a nutshell, the syntax and semantics for these approach vary considerably, making it difficult to adapt or extend them to other situations. To that end, by developing an open-universe account wrt SAT and #SAT, we obtained a simple and clean framework that can perhaps also be used to contrast and compare these accounts. Moreover, to the best of our knowledge, none of the other proposals have been shown to capture the entire range of examples in this paper, especially in the sense of expressivity matching that of proper⁺ KBs.

7 Conclusions

This paper was concerned with first-order knowledge bases in open universes. More precisely, for so-called proper⁺ knowledge bases that allow universally quantified clauses, where the domain of quantification is countably infinite, we showed that SAT and #SAT technology can be leveraged. Mainly, a ground version of the knowledge base wrt the known individuals and a small number of unknowns was shown to be sufficient for the tasks of satisfiability, validity and model counting. The topic was explored in a general way, and thus, we believe this work would also provide the basis for exploring further questions in AI and machine learning over open universes in a principled manner.

Outside of the foundational questions identified in Section 5, there are many interesting avenues for the future:

- Modern #SAT solvers appeal to caching components (Bacchus, Dalmao, and Pitassi 2009) or to circuit compilation (Chavira and Darwiche 2008), which perform significantly better than full backtracking in a DPLL search tree. However, if every subsequent query introduces a new parameter, the knowledge base on which we perform model counting would now include new clauses. How component caching and knowledge compilation can be adapted to handle this situation of changing knowledge bases in an effective way (*i.e.*, without a restart) is an interesting avenue for research.
- Extending and studying our results in a probabilistic context, where weights are accorded to logical sentences (Richardson and Domingos 2006), would then allow us to consider lifted graphical models in open universes.
- Finally, extending the logical language from propositions to functions would allow us to handle applications requiring referential uncertainty, as in (Srivastava et al. 2014).

Acknowledgements

The author would like to thank the anonymous reviewers for helpful suggestions and comments, and is supported by the Research Foundation-Flanders (FWO-Vlaanderen).

References

- Bacchus, F.; Dalmao, S.; and Pitassi, T. 2009. Solving #SAT and Bayesian inference with backtracking search. *J. Artif. Intell. Res. (JAIR)* 34:391–442.
- Barrett, C.; Sebastiani, R.; Seshia, S. A.; and Tinelli, C. 2009. Satisfiability modulo theories. In *Handbook of Satisfiability*. IOS Press. chapter 26, 825–885.
- Beame, P.; Van den Broeck, G.; Gribkoff, E.; and Suciuc, D. 2015. Symmetric weighted first-order model counting. In *PODS*, 313–328. ACM.
- Belle, V., and Lakemeyer, G. 2011. On progression and query evaluation in first-order knowledge bases with function symbols. In *Proc. IJCAI*, 255–260.
- Belle, V.; Passerini, A.; and Van den Broeck, G. 2015. Probabilistic inference in hybrid domains by weighted model integration. In *IJCAI*.
- Chandra, A. K., and Merlin, P. M. 1977. Optimal implementation of conjunctive queries in relational data bases. In *STOC*, 77–90.
- Charniak, E., and Goldman, R. P. 1993. A bayesian model of plan recognition. *Artif. Intell.* 64(1):53–79.
- Chavira, M., and Darwiche, A. 2008. On probabilistic inference by weighted model counting. *Artif. Intell.* 172(6-7):772–799.
- Chistikov, D.; Dimitrova, R.; and Majumdar, R. 2015. Approximate counting in smt and value estimation for probabilistic programs. In *TACAS*, volume 9035. 320–334.
- Doshi, F.; Wingate, D.; Tenenbaum, J. B.; and Roy, N. 2011. Infinite dynamic bayesian networks. In *ICML*, 913–920.
- Fagin, R. 1976. Probabilities on finite models. *The Journal of Symbolic Logic* 41(01):50–58.
- Getoor, L.; Friedman, N.; Koller, D.; and Taskar, B. 2001. Learning probabilistic models of relational structure. In *ICML*, 170–177.
- Giacomo, G. D.; Lespérance, Y.; and Levesque, H. J. 2011. Efficient reasoning in proper knowledge bases with unknown individuals. In *IJCAI*, 827–832.
- Gomes, C. P.; Sabharwal, A.; and Selman, B. 2009. Model counting. In *Handbook of Satisfiability*. IOS Press.
- Gutmann, B.; Thon, I.; Kimmig, A.; Bruynooghe, M.; and De Raedt, L. 2011. The magic of logical inference in probabilistic programming. *Theory and Practice of Logic Programming* 11:663–680.
- Koller, D., and Pfeffer, A. 1998. Probabilistic frame-based systems. In *AAAI/IAAI*, 580–587.
- Lakemeyer, G., and Levesque, H. J. 2002. Evaluation-based reasoning with disjunctive information in first-order knowledge bases. In *Proc. KR*, 73–81.
- Laskey, K. B. 2008. Mebn: A language for first-order bayesian knowledge bases. *Artif. Intell.* 172(2-3):140–178.
- Levesque, H., and Lakemeyer, G. 2001. *The logic of knowledge bases*. The MIT Press.
- Levesque, H. J. 1984. Foundations of a functional approach to knowledge representation. *Artif. Intell.* 23(2):155–212.

- Levesque, H. J. 1998. A completeness result for reasoning with incomplete first-order knowledge bases. In *Proc. KR*, 14–23.
- Lewis, H. R. 1980. Complexity results for classes of quantificational formulas. *Journal of Computer and System Sciences* 21(3):317 – 353.
- Luu, L.; Shinde, S.; Saxena, P.; and Demsky, B. 2014. A model counter for constraints over unbounded strings. In *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '14*, 565–576. New York, NY, USA: ACM.
- Ma, F.; Liu, S.; and Zhang, J. 2009. Volume computation for boolean combination of linear arithmetic constraints. In *CADE*, volume 5663 of *LNCS*. 453–468.
- Milch, B.; Marthi, B.; Russell, S. J.; Sontag, D.; Ong, D. L.; and Kolobov, A. 2005. BLOG: Probabilistic models with unknown objects. In *Proc. IJCAI*, 1352–1359.
- Mitchell, D. G., and Ternovska, E. 2005. A framework for representing and solving NP search problems. In *AAAI*, 430–435.
- Morgado, A.; Matos, P. J.; Manquinho, V. M.; and Silva, J. P. M. 2006. Counting models in integer domains. In *SAT*, 410–423.
- Pasula, H.; Marthi, B.; Milch, B.; Russell, S. J.; and Shpitser, I. 2002. Identity uncertainty and citation matching. In *NIPS*, 1401–1408.
- Richardson, M., and Domingos, P. 2006. Markov logic networks. *Machine learning* 62(1):107–136.
- Singla, P., and Domingos, P. M. 2007. Markov logic in infinite domains. In *UAI*, 368–375.
- Smullyan, R. 1995. *First-order logic*. Dover Publications.
- Srivastava, S.; Russell, S. J.; Ruan, P.; and Cheng, X. 2014. First-order open-universe pomdps. In *UAI*, 742–751.
- Valiant, L. G. 1979. The complexity of enumeration and reliability problems. *SIAM Journal on Computing* 8(3):410–421.
- Van den Broeck, G. 2013. *Lifted Inference and Learning in Statistical Relational Models*. Ph.D. Dissertation, KU Leuven.
- Van der Meyden, R. 1998. Logical approaches to incomplete information: A survey. In *Logics for databases and information systems*. Springer. 307–356.