Open-Universe Weighted Model Counting: Extended Abstract

Open-Universe Weighted Model Counting: Extended Abstract

Vaishak Belle
University of Edinburgh
vaishak@ed.ac.uk

This is an extended abstract for a paper to appear in the proceedings of Thirty-First AAAI Conference on Artificial Intelligence (Belle 2017).

Weighted model counting (WMC) has recently emerged as an effective and general approach to probabilistic inference, offering a computational framework for encoding a variety of formalisms, including factor graphs and Bayesian networks (Chavira and Darwiche 2008). WMC generalizes #SAT, where we are to count the models of a propositional formula, in that models can be additionally accorded numeric weights, whose sum we are to compute. In particular, the encoding of graphical models to a propositional theory allows us to leverage context-specific dependencies, express hard constraints, and reason about logical equivalence. Exact solvers are based on knowledge compilation or exhaustive DPLL search; approximate ones use local search or sampling.

The advent of large-scale probabilistic knowledge bases, such as Google's Knowledge Vault and Microsoft's Probase, often containing billions of tuples and structured data extracted from the Web and other text corpora has generated enormous interest in relational probabilistic representations. In Markov logic networks, for example, the formula:

$$1.2 \quad \forall x, y\ Smoker(x) \land Friends(x, y) \supset Smoker(y)$$

is indicative of a fairly involved Markov network obtained by grounding the formula wrt all possible values for $\{x, y\}$ and assigning a potential of 1.2 to the edges of the network. Such *template models* can be seen as modest relatives of general-purpose probabilistic logics in succinctly characterizing large finite graphical models, making them appropriate for reasoning and learning with big uncertain data. In practice, the semantics of template models is given in terms of the ground (propositional) theory, for which WMC and its extensions suffice.

By construction, and indeed by design, a fundamental limitation of these template models is its *finite domain closure assumption*: the range of the quantifiers is assumed to be finite, typically the set of constants appearing in the logical theory. As argued in (Russell 2015), such a closed-world system is at odds with areas such as vision, language understanding, and Web mining, where the existence of objects must be inferred from raw data. Various languages have been proposed to alleviate this with specialized machinery, including probabilistic programming and infinite-state Bayesian networks.

In this paper, we revisit *open-universe* (OU) template models, and develop an account of probabilistic inference by weighted model counting. Formally, our proposal will allow quantifiers to range over an *infinitary set of rigid designators* – constants that exist in all possible worlds – a common technical device from modal logic, also used elsewhere in OU languages (Srivastava et al. 2014). Such an extension is powerful, and can be used for features such as:

- *unknown atoms*: $\{\forall x\ (x \neq john \supset Smoker(x))\}$ says that infinitely many individuals other than John are smokers, while also leaving open whether John is a smoker;
- *unknown values*: $\{\forall x(Canary(x) \supset \neg Color(x, black))\}$ says the color of canaries is anything but black;
- *varying sets of objects*: $\{\forall x(Thing(x) \supset \phi(x))\}$ says that property ϕ is true of all things, but the set of (possibly infinite) things can vary depending on the world;
- *closed-world assumption:* $\forall x((x = john) \supset Smoker(x)) \land \forall x((x \neq john) \supset \neg Smoker(x))$ says that John is the only one in the universe who smokes.

In general, such a first-order logical language is undecidable. However, we show that when restricted to universally quantified clauses, as is usual in the literature on template models, a finite version of the ground knowledge base suffices for reasoning purposes.

When probabilities are further accorded to infinitely many atoms, the natural question is in which sense are conditional probabilities correct? Prior work has focused on various semantics with infinitely many random variables, based on topological orderings and locality properties. In this work, we contribute a definition of conditional probabilities that is coherent wrt first-order logical entailment: this has the benefit that no syntactical constraints on the encoding are necessary. We show its application in a number of examples, including a parametrized version of Pearl's (1988) Alarm Bayesian network discussed in (Russell 2015).

References

Belle, V. 2017. Open-universe weighted model counting. In *AAAI*.

Chavira, M., and Darwiche, A. 2008. On probabilistic inference by weighted model counting. *Artif. Intell.* 172(6-7):772–799.

Pearl, J. 1988. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann.

Russell, S. J. 2015. Unifying logic and probability. *Commun. ACM* 58(7):88–97.

Srivastava, S.; Russell, S. J.; Ruan, P.; and Cheng, X. 2014. First-order open-universe pomdps. In *UAI*, 742–751.