

On The Complexity of Inference About Probabilistic Relational Models

Manfred Jaeger

*Max-Planck-Institut für Informatik
Im Stadtwald, 66123 Saarbrücken,
Germany*

Abstract

We investigate the complexity of probabilistic inference from knowledge bases that encode probability distributions on finite domain relational structures. Our interest here lies in the complexity in terms of the domain under consideration in a specific application instance. We obtain the result that assuming $\text{NETIME} \neq \text{ETIME}$ this problem is not polynomial for reasonably expressive representation systems. The main consequence of this result is that it is unlikely to find inference techniques with a better worst-case behavior than the commonly employed strategy of constructing standard Bayesian networks over ground atoms (knowledge based model construction).

Key words: Knowledge based model construction, Bayesian networks, first-order logic

1 Introduction

A recent development in probabilistic reasoning in AI is the emergence of various systems for the specification of probability distributions on relational structures, or, in the terminology of Friedman et al. (1999), the construction of *probabilistic relational models* (Ngo & Haddawy 1997, Jaeger 1997, Koller & Pfeffer 1998). These systems have evolved out of earlier frameworks that were developed as specification languages for structurally uniform classes of Bayesian networks (Poole 1993, Breese 1992, Saffiotti & Umkehrer 1994). Given a particular probabilistic query, a specification in such a language would serve as the blueprint for the automatic generation of a Bayesian network in which the probability of the query then is computed. This method has been called *knowledge based model construction* (Wellman, Breese & Goldman 1992).

Initially, only representation languages were considered that are based on some form of probabilistic Horn clauses. Ignoring many particular features present in the representation languages proposed by various authors, these probabilistic Horn clauses are essentially of the form

$$\begin{aligned} \mathbf{p}(u, v) &\stackrel{0.3}{\leftarrow} \mathbf{q}(u, v) \\ \mathbf{p}(u, v) &\stackrel{0.5}{\leftarrow} \mathbf{r}(u), \mathbf{s}(v) \end{aligned} \tag{1}$$

where $\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{s}$ are relation symbols, and u, v, w logical variables. The intuitive meaning of e.g. the first clause is: for all u, v , the conditional probability of $\mathbf{p}(u, v)$ given that $\mathbf{q}(u, v)$ holds is 0.3. Given constants a, b for which we have evidence $\mathbf{q}(a, b)$ the rule allows us to compute a posterior probability of 0.3 for $\mathbf{p}(a, b)$ ¹ (if (1) expresses statistical knowledge, this computation would be an instance of *direct inference*, cf. (Bacchus 1990)). This coincides with the interpretation of similar rules in certain probabilistic logics (Ng & Subrahmanian 1992, Lakshmanan & Sadri 1994). The difference between knowledge based model construction and its outgrowths on the one hand, and probabilistic logics on the other emerges when we consider conditional probabilities that are not fully determined by the rules. The conditional probability $\rho := P(\mathbf{p}(a, b) \mid \mathbf{q}(a, b), \mathbf{r}(a), \mathbf{s}(b))$, for instance, is not defined by either of the rules in (1). Moreover, only the trivial bounds $[0, 1]$ are strictly implied for ρ by instantiations of the rules (1) with a, b, c . In most probabilistic logics, therefore, one will be unable to derive from (1) any nontrivial bounds for ρ .

In purely propositional settings, Bayesian networks have proven to be more useful in practice than propositional probabilistic logics (Nilsson 1986, Frisch & Haddaway 1994) because they define a unique probability distribution on the set of propositional models (i.e. truth assignments), and therefore (at the cost of a greater specification effort) allow us to derive a unique probability value for every query. It is natural to extend this approach to certain forms of first-order probabilistic information, and to develop tools for defining probability distributions on models for first-order logic. In knowledge based model construction this is done by interpreting the probabilistic rules (1) as rules for the construction of standard Bayesian networks over ground atoms. Given a ground query $P(\mathbf{p}(a, b) \mid \mathbf{q}(a, b), \mathbf{r}(a), \mathbf{s}(b)) = ?$ the model construction will yield a Bayesian network containing nodes for the atoms $\mathbf{p}(a, b), \mathbf{q}(a, b), \mathbf{r}(a), \mathbf{s}(b)$ (and possibly a large number of additional nodes),

¹ As long as no constant symbols appear in the rules, the same will be true for constants c, d for which we have the same evidence $\mathbf{q}(c, d)$. Most concrete representation systems provide for constants in the rules, so that the probabilities entailed by the rules are not necessarily invariant under substituting different constants. For the purpose of the present paper we may focus on rules without constants, because our main result is a lower complexity bound, which, obviously, also is applicable to richer systems admitting constant symbols.

and thereby determine a unique value for $P(p(a, b) \mid q(a, b), r(a), s(b))$. As noted above, the intended semantics of the rules (1) alone will not uniquely determine the desired probabilities, so that it is clear that at some point additional information or assumptions – not directly expressed by (1) – must enter the construction process. Essentially, these additional assumptions have to determine how the conditional probabilities in clauses with the same head are to be combined to obtain the conditional probability of the head given the conjunction of the bodies of the various clauses.

In early approaches (Breese 1992) this information was supplied implicitly by certain implementation details of the construction algorithm, and consequently the primary representation language did not possess a declarative semantics independent from the network construction process. Haddawy (1994) and Ngo and Haddawy (1997) have argued that this is unsatisfactory, and have proposed representation systems with additional syntactic constructs that in the knowledge base declare how several applicable clauses are to be combined.

Relational Bayesian networks (Jaeger 1997, Jaeger 1998) can be understood as a representation formalism that goes one step further by compiling sets of clauses (1), and the necessary additional conventions for their combination, into a single functional expression F , so that the knowledge base now consists of exactly one declaration of the form

$$\mathbf{r}(\mathbf{v}) := F(\mathbf{s}_1, \dots, \mathbf{s}_k, \mathbf{v}), \quad (2)$$

for each relation symbol \mathbf{r} (we use boldface letters $\mathbf{v}, \mathbf{a}, \dots$ as an abbreviation for tuples $(v_1, \dots, v_k), (a_1, \dots, a_l), \dots$ of variables or constants). Formal semantics for this set of declarations then can be defined in a straightforward manner. Another related framework that uses a representation language different from probabilistic Horn clauses are the probabilistic frame-based systems of Koller and Pfeffer (1998).

Once one has taken the step to supply the primary representation formalism with descriptive semantics independent from any construction algorithm for standard Bayesian networks, the question arises whether standard Bayesian networks are still needed at all. Their role now has changed from being the subject of our primary representation to being merely a tool of inference: if there were more efficient ways to compute the answer to a probabilistic query than by constructing a Bayesian network over ground atoms, we would be happy to dispense with Bayesian networks altogether. To emphasize this shift of perspective, we refer as *auxiliary network construction* to the process of constructing standard Bayesian networks as an inference technique for representation languages with independent semantics.

It does not seem to be unreasonable to expect more efficient inference techniques than auxiliary network construction to exist, because this approach

amounts to a complete “propositionalization” of first-order information. For logic inference problems from (deterministic) Horn-clauses we know that we can avoid this, and, for example, by unification and resolution deduce from

$$\begin{aligned} p(v) &\leftarrow q(v, w) \\ q(a, u) &\leftarrow \end{aligned}$$

that $p(a)$ holds, without first constructing all the ground atoms $p(c), q(c, c'), \dots$ for all constants c, c', \dots in the language.

It is natural to look for corresponding techniques for probabilistic inference from first-order probabilistic rules like (1) or (2) – techniques that compute probabilities by manipulating more abstract logical expressions than ground atoms. In this paper we show that it is very unlikely that with such algorithms we can obtain inference techniques that are more efficient than auxiliary network construction.

2 Model Representation Systems

It is our aim to derive our complexity results in as general terms as possible, showing their applicability to a great variety of different representation systems. In order to do this, we have to abstract from the concrete syntactical constructs used in various systems, and analyze these systems in terms of their semantic expressiveness.

To achieve this goal we develop in this section the general concept of a *probabilistic model representation system*, which (very loosely) can be seen as specialized counterpart of the general concept of a logic. Just as different logics can be compared, and their complexity be analyzed, by considering the classes of models they can define, we derive results for model representation systems in terms of the class of models they can describe, where models now are probability distributions.

First, we have to describe the structure of the models that are defined by the representation systems we deal with. To motivate the following definition, consider again the case of probabilistic Horn clauses as the representation language. It is clear that, e.g., the semantics of knowledge base (1) will be used to assign probability values to sentences such as $p(a, b) \wedge s(b)$. However, it is not enough to say that the semantics of a knowledge base is given by a

probability distribution over sentences: to see why, consider the two rules

$$\begin{aligned} \mathbf{p}(v) &\stackrel{0.5}{\leftarrow} \mathbf{q}(v, u) \\ \mathbf{q}(v, u) &\stackrel{0.8}{\leftarrow} \end{aligned}$$

where the second clause means that the marginal probability of $\mathbf{q}(v, u)$ is 0.8. Also assume that the semantical conventions adopted (perhaps via some additional declarations in the knowledge base) make the conditional probability of $\mathbf{p}(a)$ increase in the number of valid instantiations for u in $\mathbf{q}(a, u)$. In particular, we would have $P(\mathbf{p}(a) \mid \mathbf{q}(a, b)) < P(\mathbf{p}(a) \mid \mathbf{q}(a, b), \mathbf{q}(a, c))$. But more than that, since each possible instantiation of u in $\mathbf{q}(a, u)$ has a positive probability of 0.8 of being valid, the probability of $\mathbf{p}(a)$ should also increase in the number of possible instantiations, whether or not they appear in the evidence. Thus, the probability of $\mathbf{p}(a)$ as defined by the given rules, can only be determined with respect to a certain (finite) domain D of elements that we can substitute for v and u .

Consequently, the semantics of a knowledge base does not consist of a single probability distribution over sentences, but of one distribution for each (finite) domain D . Formally, a probability distribution on sentences containing relation symbols from a vocabulary S and constants from D is most conveniently represented by a distribution on the set of all structures (or models) that interpret the symbols in S over D . We denote the set of these structures by $\text{Mod}_D(S)$. As the particular names of the elements of D should be irrelevant, we may restrict attention to the case where $D = n = \{0, 1, \dots, n-1\}$ for some $n \in \mathbb{N}$ (we here avail ourselves of the set-theoretic convention to identify the number $n \in \mathbb{N}$ with the set $\{0, \dots, n-1\}$).

Definition 1 *A finite domain probabilistic relational model representation system \mathbf{M} consists of*

- *A syntax that defines for every relational vocabulary $S = \{\mathbf{r}_1, \dots, \mathbf{r}_k\}$ a set $M(S)$ of well-formed model representations.*
- *A semantics that assigns to every $\Phi \in M(S)$ and every $n \in \mathbb{N}$ a probability measure P_n^Φ on $\text{Mod}_n(S)$.*

Note that a probabilistic model representation system differs from a *probabilistic logic* in that it is required that every model representation Φ defines for every n a unique measure on $\text{Mod}_n(S)$, whereas a theory Φ in a probabilistic logic will usually define a (possibly empty) set of such measures. It should also be noted that most existing systems are somewhat more general than described in definition 1 in that they allow for a set R of predefined, deterministic relations on the domain, so that the semantics maps R -structures over n to probability measures over S -expansions of the R -structure.

In the sequel we simply write “model representation system” and “vocabulary” for “finite domain probabilistic relational model representation system” and “relational vocabulary”, respectively.

Next we describe minimal requirements for the expressiveness of model representation systems. Our complexity results will hold for those systems that satisfy these requirements.

The first requirement is very simple: we should be able to represent the uniform distribution on $\text{Mod}_n(S)$. The second requirement is to have the ability to condition the probability that \mathbf{v} belongs to some relation \mathbf{r} on certain logical properties of \mathbf{v} with respect to other relations $\mathbf{s}_1, \mathbf{s}_2, \dots$. In a clause based representation language, for instance, this requirement will demand the availability of rules of the form

$$\mathbf{r}(\mathbf{v}) \stackrel{P}{\leftarrow} \alpha(\mathbf{v}, \mathbf{s}_1, \dots, \mathbf{s}_k), \quad (3)$$

where $\alpha(\mathbf{v}, \mathbf{s}_1, \dots, \mathbf{s}_k)$ is some logical expression in the variables \mathbf{v} and the relation symbols $\mathbf{s}_1, \dots, \mathbf{s}_k$. Our minimal requirement will be that rules of this form are available for α being an equality constraint $v_i = v_j$, a conjunction $\mathbf{s}_1(\mathbf{v}') \wedge \mathbf{s}_2(\mathbf{v}'')$, a negated atom $\neg \mathbf{s}(\mathbf{v})$, and an existentially quantified atom $\exists w \mathbf{s}(\mathbf{v}, w)$. Of these types of rules only the case of α being a conjunction is readily recognized as being provided by existing systems for knowledge based model construction. Existential quantification, on the other hand, might look like a rather strong assumption about a system’s expressiveness. It should be noted, however, that a rule like

$$\mathbf{r}(v) \stackrel{1}{\leftarrow} \mathbf{s}(v, w)$$

together with the common convention that multiple instantiations of the right hand side of a rule are to be combined by *noisy-or*, just amounts to existential quantification.

The following definition formulates the availability of rules like (3) in syntax-independent, general semantic terms. In this definition, and in the remainder of the paper, we need some notation for restrictions of structures to sub-vocabularies, and restrictions of measures to sub-algebras: when $\mathcal{M}' \in \text{Mod}_n(S')$ and $S \subseteq S'$, then $\mathcal{M}' \upharpoonright S$ denotes the S -structure over domain n that has the same interpretations of the symbols in S as \mathcal{M}' . Conversely, a structure $\mathcal{M} \in \text{Mod}_n(S)$ can be identified with the subset $\{\mathcal{M}' \mid \mathcal{M}' \upharpoonright S = \mathcal{M}\} \subseteq \text{Mod}_n(S')$. When P' is a probability measure on $\text{Mod}_n(S')$, then $P' \upharpoonright \text{Mod}_n(S)$ denotes the probability measure P on $\text{Mod}_n(S)$ defined by $P(\mathcal{M}) = P'(\{\mathcal{M}' \mid \mathcal{M}' \upharpoonright S = \mathcal{M}\})$, and $P'(\cdot \mid \mathcal{M})$ denotes the conditional distribution on $\text{Mod}_n(S')$ given $\{\mathcal{M}' \mid \mathcal{M}' \upharpoonright S = \mathcal{M}\}$. We use the notation $\mathbf{v}' \subseteq \mathbf{v}$ to express that all variables in the tuple \mathbf{v}' are variables that also appear in \mathbf{v} .

Definition 2 A model representation system \mathbf{M} allows first-order conditioning if

- For every vocabulary S there exists $\Phi \in M(S)$ such that P_n^Φ is the uniform distribution on $\text{Mod}_n(S)$ for all $n \in \mathbb{N}$.
- For every vocabulary S , every $\Phi \in M(S)$, every k -ary $\mathbf{r} \notin S$, for $\mathbf{v} := (v_1, \dots, v_k)$, and for every expression $\alpha(\mathbf{v})$ of one of the four forms
 - $v_i = v_j$ ($1 \leq i, j \leq k$),
 - $\mathbf{s}_1(\mathbf{v}') \wedge \mathbf{s}_2(\mathbf{v}'')$ ($\mathbf{s}_1, \mathbf{s}_2 \in S; \mathbf{v}', \mathbf{v}'' \subseteq \mathbf{v}$),
 - $\neg \mathbf{s}(\mathbf{v}')$ ($\mathbf{s} \in S, \mathbf{v}' \subseteq \mathbf{v}$),
 - $\exists w \mathbf{s}(\mathbf{v}', w)$ ($\mathbf{s} \in S, \mathbf{v}' \subseteq \mathbf{v}$)
 there exists $\Phi^\alpha \in M(S \cup \{\mathbf{r}\})$, such that $\forall n \in \mathbb{N}$:

$$P_n^{\Phi^\alpha} \upharpoonright \text{Mod}_n(S) = P_n^\Phi, \quad (4)$$

and for all $n \in \mathbb{N}$, all $\mathcal{M} \in \text{Mod}_n(S)$, and all $\mathbf{m} \in n^k$:

$$P_n^{\Phi^\alpha}(\mathbf{r}(\mathbf{m}) \mid \mathcal{M}) = \begin{cases} 1 & \text{if } \mathcal{M} \models \alpha(\mathbf{m}) \\ 0 & \text{if } \mathcal{M} \not\models \alpha(\mathbf{m}) \end{cases} \quad (5)$$

The conditions of definition 2 demand that the probability of $\mathbf{r}(\mathbf{v})$ can be conditioned on very simple logical properties of \mathbf{v} . For relational Bayesian networks it is straightforward to show that, in fact, $\mathbf{r}(\mathbf{v})$ can be conditioned on arbitrary first-order expressible properties of \mathbf{v} (Jaeger 1997). For systems based on probabilistic Horn rules, on the other hand, it is not so obvious that with rules (3) for simple formulas α we can also encode more complicated conditions like

$$\mathbf{r}(\mathbf{v}) \leftarrow \frac{1}{2} \neg \exists w (\mathbf{s}_1(\mathbf{v}', w)) \vee \mathbf{s}_2(\mathbf{v}''). \quad (6)$$

The following lemma, which is instrumental to the proof of theorem 4 in the next section, shows that the elementary requirements of definition 2 are sufficient to guarantee that rules like (6) can be encoded.

Lemma 3 Let \mathbf{M} be a model representation system that allows first-order conditioning. Let Φ be a model representation for a vocabulary S , $\mathbf{r} \notin S$ a k -ary relation symbol, and $\phi(\mathbf{v})$ a first-order S -formula whose free variables are among $\mathbf{v} = (v_1, \dots, v_k)$. Then there exists a model representation Φ^ϕ for a vocabulary $S^\phi \supseteq S \cup \{\mathbf{r}\}$, such that for all n

$$P_n^{\Phi^\phi} \upharpoonright \text{Mod}_n(S) = P_n^\Phi, \quad (7)$$

and for all $n \in \mathbb{N}$, all $\mathcal{M} \in \text{Mod}_n(S)$, and all $\mathbf{m} \in n^k$:

$$P_n^{\Phi^\phi}(\mathbf{r}(\mathbf{m}) \mid \mathcal{M}) = \begin{cases} 1 & \text{if } \mathcal{M} \models \phi(\mathbf{m}) \\ 0 & \text{if } \mathcal{M} \not\models \phi(\mathbf{m}) \end{cases}. \quad (8)$$

In particular, for all $\mathbf{m} \in n^k$:

$$P_n^{\Phi^\phi}(\mathbf{r}(\mathbf{m})) = P_n^\Phi(\phi(\mathbf{m})). \quad (9)$$

PROOF. First note that (7) and (8) directly imply (9). We prove the existence of Φ^ϕ with (7) and (8) by induction on the structure of ϕ .

First, assume that ϕ is of the form $\mathbf{s}(\mathbf{v}')$ for some $\mathbf{s} \in S$, $\mathbf{v}' \subseteq \mathbf{v}$. Then the lemma follows from the case $\alpha(\mathbf{v}) \equiv \mathbf{s}(\mathbf{v}') \wedge \mathbf{s}(\mathbf{v}')$ in definition 2. The case for ϕ of the form $v_i = v_j$ is similar.

Now consider $\phi(\mathbf{v})$ of the form $\psi(\mathbf{v}) \wedge \chi(\mathbf{v})$. According to the induction hypothesis, the lemma holds for ψ and χ . Applying the induction hypothesis first to ψ , let \mathbf{r}_ψ be a new k -ary relation symbol, and let Φ^ψ be a model representation for a vocabulary $S^\psi \supseteq S \cup \{\mathbf{r}_\psi\}$, such that

$$P_n^{\Phi^\psi} \upharpoonright \text{Mod}_n(S) = P_n^\Phi,$$

and for all $\mathcal{M} \in \text{Mod}_n(S)$, $\mathbf{m} \in n^k$:

$$P_n^{\Phi^\psi}(\mathbf{r}_\psi(\mathbf{m}) \mid \mathcal{M}) = \begin{cases} 1 & \text{if } \mathcal{M} \models \psi(\mathbf{m}) \\ 0 & \text{else.} \end{cases}$$

Now we apply the induction hypothesis to χ and the already constructed Φ^ψ . This gives us a model representation Φ^χ for a vocabulary $S^\chi \supseteq S^\psi$ containing another new k -ary relation symbol \mathbf{r}_χ , such that

$$\begin{aligned} P_n^{\Phi^\chi} \upharpoonright \text{Mod}_n(S) &= (P_n^{\Phi^\chi} \upharpoonright \text{Mod}_n(S^\psi)) \upharpoonright \text{Mod}_n(S) \\ &= P_n^{\Phi^\psi} \upharpoonright \text{Mod}_n(S) \\ &= P_n^\Phi, \end{aligned}$$

and for all $\mathcal{M} \in \text{Mod}_n(S^\psi)$, $\mathbf{m} \in n^k$:

$$\begin{aligned} P_n^{\Phi^\chi}(\mathbf{r}_\chi(\mathbf{m}) \mid \mathcal{M}) &= \begin{cases} 1 & \text{if } \mathcal{M} \models \chi(\mathbf{m}) \\ 0 & \text{else.} \end{cases} \\ &= \begin{cases} 1 & \text{if } \mathcal{M} \upharpoonright S \models \chi(\mathbf{m}) \\ 0 & \text{else.} \end{cases} \end{aligned}$$

Given Φ^χ we now can use the case $\alpha \equiv \mathbf{r}_\psi(\mathbf{v}) \wedge \mathbf{r}_\chi(\mathbf{v})$ of definition 2 to find a final model representation Φ^ϕ for the vocabulary $S^\phi = S^\chi \cup \{\mathbf{r}\}$, such that

$$\begin{aligned} P_n^{\Phi^\phi} \upharpoonright \text{Mod}_n(S) &= (P_n^{\Phi^\phi} \upharpoonright \text{Mod}_n(S^\chi)) \upharpoonright \text{Mod}_n(S) \\ &= P_n^{\Phi^\chi} \upharpoonright \text{Mod}_n(S) \\ &= P_n^\Phi, \end{aligned}$$

and for all $\mathcal{M} \in \text{Mod}_n(S^\chi)$, $\mathbf{m} \in n^k$:

$$\begin{aligned} P_n^{\Phi^\phi}(\mathbf{r}(\mathbf{m}) \mid \mathcal{M}) &= \begin{cases} 1 & \text{if } \mathcal{M} \models \mathbf{r}_\psi(\mathbf{m}) \wedge \mathbf{r}_\chi(\mathbf{m}) \\ 0 & \text{else} \end{cases} \\ &= \begin{cases} 1 & \text{if } \mathcal{M} \upharpoonright S \models \psi(\mathbf{m}) \wedge \chi(\mathbf{m}) \\ 0 & \text{else.} \end{cases} \end{aligned}$$

The last identity establishes (8) for $\mathcal{M} \in \text{Mod}_n(S)$.

The case for $\phi(\mathbf{v})$ of the form $\neg\psi(\mathbf{v})$ is dealt with in a similar manner.

Finally, consider $\phi(\mathbf{v})$ of the form $\exists w\psi(\mathbf{v}, w)$. We apply the induction hypothesis to a relation symbol \mathbf{r}_ψ of arity $k+1$ and the formula $\psi(\mathbf{v}, w)$ to obtain a model representation Φ^ψ . We then obtain Φ^ϕ by applying definition 2 for the case $\alpha(\mathbf{v}) \equiv \exists w\mathbf{r}_\psi(\mathbf{v}, w)$ to \mathbf{r} and Φ^ψ . \square

3 Complexity: Deterministic, Exact Inference

Given a model representation system \mathbf{M} we now are interested in the complexity of answering probabilistic queries, i.e. of computing $P_n^\Phi(\phi(\mathbf{m}))$ for a model representation Φ , a domainsize n , and a proposition $\phi(\mathbf{m})$. Obviously, with the assumptions we have made we cannot derive exact bounds for the complexity of this computation, because these would depend on many specific features of the system \mathbf{M} that we have left unspecified. Our aim here, therefore, only is to investigate one specific aspect of the overall computational complexity,

namely its dependence on the domain size n . The dependency on this parameter is of particular interest, because it is with regard to this parameter that we would expect to obtain a gain in efficiency by replacing auxiliary network construction with more sophisticated inference techniques: when Φ and $\phi(\mathbf{m})$ are fixed, then the number of nodes in an auxiliary network constructed to compute $P_n^\Phi(\phi(\mathbf{m}))$ will usually be polynomial in n , and the complexity of inference exponential in n (because, in general, we will also have in the auxiliary network a polynomial growth of the maximal number of parents of single nodes). It is not obvious that this exponential blowup in n is inherent in the problem, and cannot be avoided by other inference techniques. Note, in particular, that the well-known complexity results for inference in Bayesian networks (Cooper 1990) are not applicable here, because we cannot represent a suitable class of Bayesian networks that shows that inference is NP-hard in the network size as the set of auxiliary networks constructed for a set of queries $P_n^\Phi(\phi(\mathbf{m}))$ ($n \in \mathbb{N}$; Φ, ϕ fixed).

Thus, we here will be concerned with the complexity of computing $P_n^\Phi(\phi(\mathbf{m}))$ as a function of n with Φ and $\phi(\mathbf{m})$ being fixed. Moreover, following a common strategy, we will first concentrate on the simpler problem of deciding whether $P_n^\Phi(\phi(\mathbf{m})) > 0$. Formally, our problem then becomes that of deciding predicates of the form

$$\text{NONZERO}(\Phi, \phi(\mathbf{m})) := \{n \in \mathbb{N} \mid P_n^\Phi(\phi(\mathbf{m})) > 0\}$$

defined by model representations Φ (in some representation system \mathbf{M}), and a formula $\phi(\mathbf{m})$ containing constants $m_1, \dots, m_k \in \mathbb{N}$ (use the convention that $P_n^\Phi(\phi(\mathbf{m})) = 0$ when $m_i > n$ for some $i \leq k$, and therefore $\phi(\mathbf{m})$ cannot not be interpreted over the domain n). For arbitrary subsets $A \subseteq \mathbb{N}$ we use A^{un} and A^{bin} to denote the sets of unary and binary encodings, respectively, of the members of A . Since we are interested in the complexity in terms of n of deciding $\text{NONZERO}(\Phi, \phi(\mathbf{m}))$, not in terms of $\log(n)$, we really are talking about the complexity of deciding $\text{NONZERO}(\Phi, \phi(\mathbf{m}))^{\text{un}}$, when complexity is measured in input size.

Adopting the notation of Johnson (1990), we denote by (N)ETIME the class of subsets of $\{0, 1\}^*$ that can be decided in (nondeterministic) time $O(2^{cn})$ for some $c > 0$. Note that this class is distinct from (N)EXPTIME, which is characterized by time bounds of the form $O(2^{n^c})$ ($c > 0$). We can now formulate our main theorem.

Theorem 4 *Let \mathbf{M} be a model construction system that allows first-order conditioning. If $\text{NETIME} \neq \text{ETIME}$, then there exist an S -model representation Φ in \mathbf{M} , and a ground S -atom $\mathbf{r}(\mathbf{m})$ such that $\text{NONZERO}(\Phi, \mathbf{r}(\mathbf{m}))^{\text{un}} \notin \text{P}$.*

The proof of theorem 4 is quite straightforward using lemma 3 and established results due to Jones and Selman (1972) on the connection between the class

NETIME and spectra of first-order sentences. We briefly review the relevant definitions and results here.

The *spectrum* of a first-order sentence ϕ in the vocabulary S is the set $\{n \in \mathbb{N} \mid \exists \mathcal{M} \in \text{Mod}_n(S) : \mathcal{M} \models \phi\}$, i.e. the set of all finite cardinalities for which ϕ has a model. A subset of \mathbb{N} is called a spectrum if it is the spectrum of some first-order sentence ϕ (over an arbitrary vocabulary – but note that without loss of generality we can assume a relational vocabulary). The result of Jones and Selman (1972) that we shall use is: a subset A of \mathbb{N} is a spectrum iff $A^{\text{bin}} \in \text{NETIME}$.

Proof of theorem 4 Let \mathbf{M} be as stated in the theorem, and assume that $A^{\text{bin}} \in \text{NETIME} \setminus \text{ETIME}$. By Jones and Selman’s (1972) theorem there exists a first-order sentence ψ in a relational vocabulary S , such that A is the spectrum of ψ . Let Φ be a model representation for S such that P_n^Φ is the uniform distribution on $\text{Mod}_n(S)$ for all $n \in \mathbb{N}$. By lemma 3 there exists a model representation Φ^ϕ for $S^\phi \supseteq S$, such that S^ϕ contains a unary relation symbol \mathbf{r} , and (7) and (9) hold. Since the right hand side of (9) is nonzero iff n is in the spectrum of ϕ , we obtain

$$\text{NONZERO}(\Phi^\phi, \mathbf{r}(m)) = A.$$

By the assumption $A^{\text{bin}} \notin \text{ETIME}$ it follows that

$$A^{\text{un}} = \text{NONZERO}(\Phi^\phi, \mathbf{r}(m))^{\text{un}} \notin \text{P}.$$

□

Our proof of theorem 4 relies crucially on the requirement of definition 2 that we can condition on equality constraints. If we deleted the case $\alpha \equiv v_i = v_j$ from definition 2 then our arguments would only show that we can encode as sets of the form $\text{NONZERO}(\Phi, \mathbf{r}(\mathbf{m}))$ spectra of first-order sentences without equality. These, however, are simply sets of the form $\mathbb{N} \setminus \{1, 2, \dots, k-1, k\}$, which can be decided in constant time.

Theorem 4 gives us a lower complexity bound for some $\text{NONZERO}(\Phi, \mathbf{r}(\mathbf{m}))$. For most concrete model construction systems proposed so far, on the other hand, we have the upper bound $\text{NONZERO}(\Phi, \mathbf{r}(\mathbf{m})) \in \text{NP}$ for all $\Phi, \mathbf{r}(\mathbf{m})$. This suggests to check whether $\text{NONZERO}(\Phi, \mathbf{r}(\mathbf{m}))$ might be an NP-complete problem for some representation system \mathbf{M} , and suitable $\Phi, \mathbf{r}(\mathbf{m})$. A general result in complexity theory, however, says that this is unlikely to be the case.

Theorem 5 *If $\text{P} \neq \text{NP}$ then $\text{NONZERO}(\Phi, \mathbf{r}(\mathbf{m}))^{\text{un}}$ is not NP-complete for any representation system \mathbf{M} , model representation Φ , and query $\mathbf{r}(\mathbf{m})$.*

PROOF. This follows immediately from results of Mahaney (1986) that so-called *sparse* sets cannot be NP-complete if $P \neq NP$. All sets in unary encoding (also called tally languages) are examples of sparse sets. \square

While it is customary to simplify questions about the complexity of computing a function to a simpler decision problem, it is of course the complexity of computing the value of $P_n^\Phi(\mathbf{r}(\mathbf{m}))$ that we are ultimately interested in. In practice, one will usually not need to compute the precise probability value, but only an approximation with a certain given precision. This leads us to the subject of approximate inference, which we deal with in the following section.

4 Approximate Inference

Following Dagum and Luby (1993), we may distinguish between four principal variants of approximate inference: the approximation may be within a specified *absolute* or *relative error*, and the approximation algorithm may be either *deterministic* or *randomized*.

When lower complexity bounds for the computation of exact probabilities $P(\cdot)$ are derived by a reduction to a decision problem of the form $P(\cdot) > 0$? (as we did here, and as did Cooper (1990) for standard Bayesian network inference), then we cannot gain much by turning from exact inference to approximate inference with a bounded relative error, because an approximation of $P(\cdot)$ with a bounded relative error will still show whether $P(\cdot) > 0$ or $P(\cdot) = 0$. For this reason we here concentrate on computing approximations z for $P_n^\Phi(\mathbf{r}(\mathbf{m}))$ with a bound ϵ on the absolute error, i.e. a number z that satisfies

$$z \in [P_n^\Phi(\mathbf{r}(\mathbf{m})) - \epsilon, P_n^\Phi(\mathbf{r}(\mathbf{m})) + \epsilon].$$

We first turn to deterministic approximations. It turns out that now the complexity of computing $P_n^\Phi(\mathbf{r}(\mathbf{m}))$ can be as well-behaved as one might hope – at least in theory.

Theorem 6 *There exist model construction systems that allow first-order conditioning such that for every model representation Φ , every ground atom $\mathbf{r}(\mathbf{m})$, and every $\epsilon > 0$ the complexity (in n) of computing an approximation of $P_n^\Phi(\mathbf{r}(\mathbf{m}))$ with absolute error at most ϵ is $O(1)$.*

PROOF. In (Jaeger 1998) it is shown that for a certain subclass of relational Bayesian networks the probabilities $P_n^\Phi(\mathbf{r}(\mathbf{m}))$ converge to some limit as $n \rightarrow \infty$ for every network Φ in that subclass, and every query $\mathbf{r}(\mathbf{m})$. The

subclass identified in (Jaeger 1998) is rich enough to allow first-order conditioning.

Now assume that Φ and $\mathbf{r}(\mathbf{m})$ are such that $P_n^\Phi(\mathbf{r}(\mathbf{m})) \rightarrow p \in [0, 1]$ as $n \rightarrow \infty$, and let $\epsilon > 0$ be given. Then there exists $n_0 \in \mathbb{N}$, such that $P_n^\Phi(\mathbf{r}(\mathbf{m})) \in [p - \epsilon, p + \epsilon]$ for all $n \geq n_0$. Thus, we obtain an algorithm for computing an ϵ -approximation of $P_n^\Phi(\mathbf{r}(\mathbf{m}))$ by exact computation of $P_n^\Phi(\mathbf{r}(\mathbf{m}))$ (using any available algorithm) when $n < n_0$, and by simply outputting p when $n \geq n_0$. The time requirement of this procedure is asymptotically constant in n . \square

Clearly this result is of theoretical rather than practical interest, because neither does it tell us how to compute the number n_0 , nor does it provide any bound on the constant characterizing the time requirement. Furthermore, the theorem is not applicable for representation systems in which $P_n^\Phi(\mathbf{r}(\mathbf{m}))$ need not converge.

In practice, randomized approximation algorithms can be particularly well-suited for computing probabilities $P_n^\Phi(\mathbf{r}(\mathbf{m}))$. To see why, consider an algorithm that produces random samples $\mathcal{M}_i \in \text{Mod}_n(S)$ according to the distribution P_n^Φ . As in *logic sampling* for standard Bayesian networks (Henrion 1988) we could use the fraction of structures \mathcal{M}_i with $\mathcal{M}_i \models \mathbf{r}(\mathbf{m})$ in a random sample $\mathcal{M}_1, \dots, \mathcal{M}_n$ as an estimate for $P_n^\Phi(\mathbf{r}(\mathbf{m}))$. This is usually not the best use we can make of the sample $\mathcal{M}_1, \dots, \mathcal{M}_n$, though: when the distribution P_n^Φ is invariant under renaming (as we have always assumed), then we have

$$P_n^\Phi(\mathbf{r}(\mathbf{m})) = E_n^\Phi(\|\mathbf{r}(\mathbf{v})\|), \quad (10)$$

where by $\|\mathbf{r}(\mathbf{v})\|$ we denote the fraction of tuples \mathbf{m} that satisfy $\mathbf{r}(\mathbf{m})$ in a structure $\mathcal{M} \in \text{Mod}_n(S)$, and by E_n^Φ the expected value under the distribution P_n^Φ . Therefore, we also gain an estimate for $P_n^\Phi(\mathbf{r}(\mathbf{m}))$ by averaging over the structures $\mathcal{M}_1, \dots, \mathcal{M}_n$ the values of $\|\mathbf{r}(\mathbf{v})\|$. The variance of the random variable $\|\mathbf{r}(\mathbf{v})\|$ is at most as large as that of the indicator variable for $\mathbf{r}(\mathbf{m})$, and usually decreasing in n . Depending on how fast the variance of $\|\mathbf{r}(\mathbf{v})\|$ decreases the reduction of the size of a random sample needed to estimate $P_n^\Phi(\mathbf{r}(\mathbf{m}))$ with given error and confidence bounds can offset the increased complexity of sampling a single structure \mathcal{M}_i . In the best case we will really obtain a time requirement that is constant in n .

5 Conclusion

The purpose of this paper was twofold: first, we wanted to develop a conceptual framework that permits us to treat in a unified way a number of systems that

have been proposed in the literature for integrating some first-order reasoning capabilities into Bayesian networks. We here have obtained this unified view by characterizing such systems entirely in terms of their semantics, without imposing any restrictions on the specific syntax used.

Our main objective then was to derive within this general setting results on the complexity of probabilistic inference as a function of the size of the specific domain to which the generic knowledge is applied. This is a new complexity problem that is distinctive of the emergent class of probabilistic relational model representation systems. It does not appear in probabilistic logics, because there entailment always is with regard to all models of a knowledge base, not the models over a specific domain. It also is distinct from complexity questions about inference in standard Bayesian networks, because there a change of the domain (by way of a different set of nodes in the network) always is accompanied by a new model representation (i.e. a new network).

Our complexity problem is of particular interest, because polynomial bounds here would have shown that there are more efficient ways for probabilistic inference than auxiliary network construction – an inference technique that a priori looks rather wasteful, because it involves a complete propositionalization of originally first-order knowledge. However, theorem 4 shows that for reasonably expressive representation systems we are unlikely to find inference techniques that have a better worst-case behavior than auxiliary network construction. The proof of the theorem points to reasoning about equality as the main cause for the complexity bounds we obtained. Investigations of weaker systems with potentially lower complexity should therefore be directed towards systems without equality reasoning.

References

- Bacchus, F. (1990), *Representing and Reasoning With Probabilistic Knowledge*, MIT Press.
- Breese, J. S. (1992), ‘Construction of belief and decision networks’, *Computational Intelligence* **8**(4), 624–647.
- Cooper, G. F. (1990), ‘The computational complexity of probabilistic inference using bayesian belief networks’, *Artificial Intelligence* **42**, 393–405.
- Dagum, P. & Luby, M. (1993), ‘Approximating probabilistic inference in bayesian belief networks is np-hard’, *Artificial Intelligence* **60**, 141–153.
- Friedman, N., Getoor, L., Koller, D. & Pfeffer, A. (1999), Learning probabilistic relational models, *in* ‘Proceedings of the 16th International Joint Conference on Artificial Intelligence (IJCAI)’.

- Frisch, A. & Haddaway, P. (1994), ‘Anytime deduction for probabilistic logic’, *Artificial Intelligence* **69**, 93–122.
- Haddawy, P. (1994), Generating bayesian networks from probability logic knowledge bases, *in* ‘Proceedings of the Tenth Conference on Uncertainty in Artificial Intelligence (UAI-94)’.
- Henrion, M. (1988), Propagating uncertainty in bayesian networks by probabilistic logic sampling, *in* J. F. Lemmer & L. N. Kanal, eds, ‘Uncertainty in Artificial Intelligence 2’, pp. 149–163.
- Jaeger, M. (1997), Relational bayesian networks, *in* ‘Proceedings of the Thirteenth Conference on Uncertainty in Artificial Intelligence (UAI-97)’.
- Jaeger, M. (1998), Convergence results for relational bayesian networks, *in* ‘Proceedings of the Thirteenth Annual IEEE Symposium on Logic in Computer Science (LICS-98)’.
- Johnson, D. S. (1990), A catalog of complexity classes, *in* J. van Leeuwen, ed., ‘Handbook of Theoretical Computer Science’, Vol. 1, Elsevier, Amsterdam, pp. 67–161.
- Jones, N. D. & Selman, A. L. (1972), Turing machines and the spectra of first-order formulas with equality, *in* ‘Proceedings of the Fourth ACM Symposium on Theory of Computing’, pp. 157–167.
- Koller, D. & Pfeffer, A. (1998), Probabilistic frame-based systems, *in* ‘Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)’, pp. 580–587.
- Lakshmanan, V. & Sadri, F. (1994), Probabilistic deductive databases, *in* M. Bruynooghe, ed., ‘Proceedings of the International Symposium on Logic Programming (ILPS’94)’, MIT Press.
- Mahaney, S. R. (1986), Sparse sets and reducibilities, *in* R. V. Book, ed., ‘Studies in Complexity Theory’, Pitman, London, pp. 63–118.
- Ng, R. & Subrahmanian, V. (1992), ‘Probabilistic logic programming’, *Information and Computation* **101**, 150–201.
- Ngo, L. & Haddawy, P. (1997), ‘Answering queries from context-sensitive probabilistic knowledge bases’, *Theoretical Computer Science* **171**, 147–177.
- Nilsson, N. (1986), ‘Probabilistic logic’, *Artificial Intelligence* **28**, 71–88.
- Poole, D. (1993), ‘Probabilistic horn abduction and bayesian networks’, *Artificial Intelligence* **64**, 81–129.
- Saffiotti, A. & Umkehrer, E. (1994), ‘Inference-driven construction of valuation systems from first-order clauses’, *IEEE Transactions on Systems, Man, and Cybernetics* **24**(11), 1611–1624.
- Wellman, M. P., Breese, J. S. & Goldman, R. P. (1992), ‘From knowledge bases to decision models’, *The Knowledge Engineering Review* **7**(1), 35–53.