

Knowledge Compilation and Theory Approximation

BART SELMAN AND HENRY KAUTZ

AT&T Bell Laboratories, Murray Hill, New Jersey

Abstract. Computational efficiency is a central concern in the design of knowledge representation systems. In order to obtain efficient systems, it has been suggested that one should limit the form of the statements in the knowledge base or use an incomplete inference mechanism. The former approach is often too restrictive for practical applications, whereas the latter leads to uncertainty about exactly what can and cannot be inferred from the knowledge base. We present a third alternative, in which knowledge given in a general representation language is translated (compiled) into a tractable form—allowing for efficient subsequent query answering.

We show how propositional logical theories can be compiled into Horn theories that approximate the original information. The approximations bound the original theory from below and above in terms of logical strength. The procedures are extended to other tractable languages (for example, binary clauses) and to the first-order case. Finally, we demonstrate the generality of our approach by compiling concept descriptions in a general frame-based language into a tractable form.

Categories and Subject Descriptors: F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic; I.2.3 [Artificial Intelligence]: Deduction and Theorem Proving; I.2.4 [Artificial Intelligence]: Knowledge Representation Formalisms and Methods

General Terms: Algorithms, Experimentation, Theory

Additional Key Words and Phrases: Efficient reasoning methods, Horn clauses, knowledge-base optimization, knowledge compilation, query evaluation, theory approximation

1. Introduction

The study of the computational properties of knowledge representation systems has revealed a direct trade-off between tractability and expressiveness [Levesque and Brachman 1985]. In general, in order to obtain a computationally efficient representation system one either restricts the expressive power of the knowledge representation language or one uses an incomplete inference mechanism. In the first approach, the representation language is often too limited for practical applications [Doyle and Patil 1991]. The second approach involves either resource-bounded reasoning or the introduction of a nontraditional semantics. In resource-bounded reasoning, inference is limited

Authors' address: AI Principles Research Department, AT&T Bell Laboratories, 600 Mountain Avenue, Murray Hill, NJ 0797-0636; e-mail: {selman, kautz}@research.att.com.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery (ACM), Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 1996 ACM 0004-5411/96/0300-0193 \$03.50

by bounding the number of inference steps performed by the inference procedure. It therefore becomes difficult to characterize exactly what can and cannot be inferred, that is, the approach lacks a “real” semantics (one that does not simply mimic the proof theory). Moreover, no information is provided if a proof cannot be found within the time bound. (But see Horvitz et al. [1989] for an example of *probabilistic* inference, where confidence in the results increases with the amount of computation.) Accounts of limited inference based on nontraditional semantics [Levesque 1984; Frisch 1985; Patel-Schneider 1986] often provide only a very weak kind of inference. For example, in the four-valued semantics approach of Levesque, given the statements p and $p \supset q$, one cannot infer q .

This paper presents a third alternative for obtaining efficient representation systems, which neither limits the expressive power of the representation language, nor gives up completeness of the inference procedure. In this new approach, the user enters statements in a general, unrestricted representation language, which the system compiles into a restricted language that allows for efficient inference. Since an exact translation into a tractable form is often impossible, the system searches for the best approximation of the original information. We describe how the approximation can be used to speed up inference. We refer to this approach as *knowledge compilation*.

We illustrate our method by studying the approximation of propositional logic theories by Horn theories. Applications of propositional inference arise in many areas of artificial intelligence and computer science in general. Examples include qualitative reasoning about physical systems [Reiter and de Kleer 1987], circuit diagnosis [Larrabee 1992], expert systems [Dean 1985], and learning theory [Valiant 1984]. In addition, many reasoning tasks can be directly mapped into propositional inference problems, such as finite-domain constraint satisfaction problems (CSP) [Dechter 1992] and certain classes of planning problems [Kautz and Selman 1992b].

Finally, note that first-order domain theories in artificial intelligence often involve only finite domains. Such theories are therefore essentially propositional. For example, Reiter and Mackworth [1989] discuss how their first-order domain theory for visual interpretation can be restated in propositional form. By considering the propositional formulation, they show how methods from relational databases can be used to answer queries much more efficiently than by using a general first-order theorem prover on the original domain theory. Another example is our propositional formulation of planning, cited above.

General propositional reasoning is the prototypical co-NP-complete task. However, inference based on Horn approximations is very fast (linear time). Following the formal definition of Horn approximation, we present algorithms for generating such approximations. The algorithms have the property that they can be interrupted at any time to produce some useful intermediate result. We also present an empirical evaluation of our approach, and demonstrate substantial computational savings on computationally challenging propositional theories. The paper concludes with a discussion of extensions to first-order theories and various generalizations of Horn approximations, as well as to description logics. A preliminary discussion of these ideas appeared in Selman and Kautz [1991] and Kautz and Selman [1991].

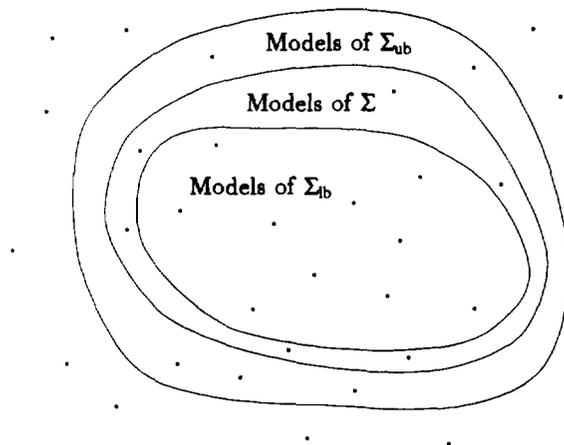


FIG. 1. Two sets of models that bound the original set of models of the original theory Σ from below and above. Models are represented by dots. Each of the bounds can be represented by a set of Horn clauses.

2. Propositional Horn Approximations

In this section, we introduce the idea of knowledge compilation using a concrete example. We show how a propositional theory can be compiled into a tractable form consisting of a set of Horn clauses.

2.1. DEFINITIONS. We assume a standard propositional language, and use a , b , c , p , q , and r to denote propositional letters and x , y , and z to denote literals (a *literal* is either a propositional letter, called a positive literal, or its negation, called a negative literal). A *clause* is a disjunction of literals, and can be represented by the set of literals it contains. A clause is *Horn* if and only if it contains at most one positive literal; a set of such clauses is called a *Horn theory*. Formulas are given in conjunctive normal form (CNF, a conjunction of disjuncts), so they can be represented by a set of clauses. CNF notation and clausal notation are used interchangeably. For example, we may write $(p \vee \neg q) \wedge r$ instead of $\{\{p, \neg q\}, \{r\}\}$, and vice versa.

In general, determining whether a given CNF formula (the *query*) follows from a set of formulas in a knowledge base is presumed to be intractable [Cook 1971]. However, when the knowledge base contains only Horn clauses the problem can be solved in time linear in the length of the knowledge base times the number of clauses in the query [Dowling and Gallier 1984].

So, a useful kind of *knowledge compilation* would be the following: Given a set of arbitrary clauses, compute a logically equivalent set of Horn clauses, and base subsequent inference on that set. Unfortunately, there does not always exist a logically equivalent Horn theory. For example, no Horn theory is equivalent to the theory $p \vee q$. We therefore propose to *approximate* the original set of clauses by a set of Horn clauses. The basic idea is to bound the set of models (satisfying truth assignments) of the original theory from below and from above by Horn theories. Figure 1 illustrates the basic idea.

In the following definition, $\mathcal{M}(\Sigma)$ denotes the set of satisfying truth assignments of the theory Σ .

Definition 2.1.1. Horn lower-bound and Horn upper-bound. Let Σ be a set of clauses. The sets Σ_{lb} and Σ_{ub} of Horn clauses are respectively a Horn lower-bound and a Horn upper-bound of Σ iff

$$\mathcal{M}(\Sigma_{lb}) \subseteq \mathcal{M}(\Sigma) \subseteq \mathcal{M}(\Sigma_{ub})$$

or, equivalently,

$$\Sigma_{lb} \models \Sigma \models \Sigma_{ub}.$$

Note that the bounds are defined in terms of models: the lower-bounds have fewer models than the original theory, and the upper-bound has more models. The reader is cautioned not to associate “lower” with “logically weaker.” In fact, because the lower-bound has fewer models, it is logically stronger than (i.e., implies) the original theory. Similarly, because the upper-bound has more models, it is logically weaker than (i.e., is implied by) the original theory.

Instead of simply using any pair of bounds to characterize the initial theory, we wish to use the best possible ones: a *greatest* Horn lower-bound and a *least* Horn upper-bound.

Definition 2.1.2. Greatest Horn lower-bound (GLB). Let Σ be a set of clauses. The set Σ_{glb} of Horn clauses is a greatest Horn lower-bound of Σ iff $\mathcal{M}(\Sigma_{glb}) \subseteq \mathcal{M}(\Sigma)$ and there is no set Σ' of Horn clauses such that $\mathcal{M}(\Sigma_{glb}) \subset \mathcal{M}(\Sigma') \subseteq \mathcal{M}(\Sigma)$.

Definition 2.1.3. Least Horn upper-bound (LUB). Let Σ be a set of clauses. The set Σ_{lub} of Horn clauses is a least Horn upper-bound of Σ iff $\mathcal{M}(\Sigma) \subseteq \mathcal{M}(\Sigma_{lub})$ and there is no set Σ' of Horn clauses such that $\mathcal{M}(\Sigma) \subseteq \mathcal{M}(\Sigma') \subset \mathcal{M}(\Sigma_{lub})$.

We call these bounds *Horn approximations* of the original theory Σ .¹ The definition of a Horn upper-bound implies that the conjunction of two such bounds is another, possibly smaller upper-bound. It follows that a given theory has a unique LUB (up to logical equivalence). On the other hand, a theory can have many different GLBs.

Example. Consider the non-Horn theory $\Sigma = (\neg a \vee c) \wedge (\neg b \vee c) \wedge (a \vee b)$. The Horn theory $a \wedge b \wedge c$ is an example of a Horn lower-bound; both $a \wedge c$ and $b \wedge c$ are GLBs; $(\neg a \vee c) \wedge (\neg b \vee c)$ is an example of a Horn upper-bound; and c is the LUB. The reader can verify these bounds by noting that

$$(a \wedge b \wedge c) \models (a \wedge c) \models \Sigma \models c \models ((\neg a \vee c) \wedge (\neg b \vee c)).$$

Moreover, there is no Horn theory Σ' logically distinct from $a \wedge c$ such that $(a \wedge c) \models \Sigma' \models \Sigma$. Similar properties hold of the other GLB and of the LUB.

2.2. USING APPROXIMATIONS FOR FAST INFERENCE. Before we discuss how to compute Horn approximations, let us consider how these approximations can be used to improve the efficiency of a knowledge representation system. Suppose a knowledge base (KB) contains the set of clauses Σ , and we want to determine whether the formula α is implied by the KB. We assume that α is in CNF, because one can determine in linear time if a propositional CNF formula follows from a Horn theory. (Note that the query need not be Horn.) The system can proceed as shown in Figure 2. First, it tries to obtain an answer quickly by using

¹ We will sometimes use the term Horn approximation to refer to *any* Horn upper or lower bound, not necessarily the LUB or the GLB.

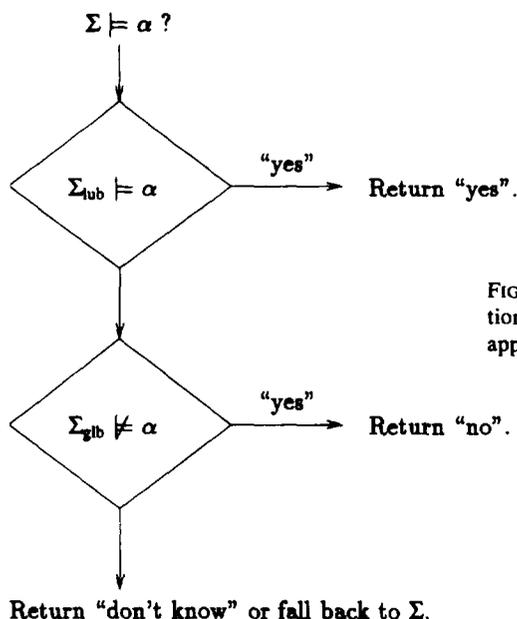


FIG. 2. Fast querying using theory approximation. The original theory is Σ ; Σ_{lub} and Σ_{glb} are its approximations; and α is the query.

the Horn approximations. If $\Sigma_{lub} \models \alpha$, then it returns “yes,” or if $\Sigma_{glb} \not\models \alpha$, then it returns “no.” So far, the procedure takes only time linear in the length of the approximations.² In case no answer is obtained, the system could simply return “don’t know,” or it could decide to spend more time and use a general inference procedure to determine the answer directly from the original theory. The general inference procedure could still use the approximations to prune its search space. Thus, the system can answer certain queries in linear time, resulting in an improvement in its overall response time. Exactly how many queries can be handled directly by the Horn approximations depends on how well the bounds characterize the original theory. (We will return to this issue in Section 2.4 below.) Note that we give up neither soundness nor completeness, because we can always fall back to the original theory.

When the system opts for returning “don’t know” instead of performing general inference with the original theory, we have an example of trading accuracy for speed. This approach may be quite useful in real-time systems that have some fall back mechanism for dealing with the “don’t know” answers. For example, if it is too expensive to answer a query by inference, a robot may instead try to determine the answer by sensing.

This paper concentrates on compiling a given knowledge base before any queries are made. Greiner and Schuurmans [1992] extend our framework by providing compilation algorithms that interleave compilation and query answering. Another interesting direction for future research would be to allow for updates to the knowledge base between queries.

2.3. COMPUTING HORN APPROXIMATIONS. We now turn to the problem of generating Horn approximations. There does not exist a polynomial time proce-

² We assume that the lengths of the Horn approximations are roughly the same as that of the original theory. We will return to this issue later on.

ture for generating such approximations (provided $P \neq NP$). This is consequence of the following theorem:

THEOREM 2.3.1. *Let Σ be a set of clauses. The GLB of Σ is consistent iff Σ is consistent. Similarly, the LUB of Σ is consistent iff Σ is consistent.*

PROOF. If $\mathcal{M}(\Sigma) = \emptyset$, then an inconsistent Horn theory such as $p \wedge \neg p$ is both a GLB and LUB of Σ . Let Σ be consistent and M be a satisfying truth assignment of Σ . By definition, M must be a satisfying assignment of any Horn upper-bound. Moreover, the theory with M as the only satisfying truth assignment is a better (larger) Horn lower-bound than an unsatisfiable (inconsistent) Horn theory. So, both the LUB and the GLB of Σ are consistent since each has at least one satisfying truth assignment. \square

If the length of the Horn approximations is bounded by some polynomial function of the length of Σ , then the task of finding them is NP-hard, because checking the consistency of a general set of clauses is NP-complete, whereas checking the consistency of Horn clauses takes only linear-time. On the other hand, if certain approximations are of exponential length, then it certainly takes exponential time to generate them. Thus, in either case, the problem is intractable. Of course, if the approximations were polynomial time computable, then they could not be very good approximations (e.g., inconsistency could go undetected), and at query time they could save at most a polynomial amount of time on an exponentially hard problem.

Computing the Horn approximations should be viewed as a *compilation* process. The computational cost can be amortized over the total set of subsequent queries to the KB. In some cases, however, the approximations may be needed for query answering before the compilation process finishes. So, instead of waiting for the best Horn bounds, it would be desirable to employ procedures that could output lower- and upper-bounds as intermediate results, generating better and better bounds over time. That is, the approximation algorithms should be “anytime” procedures [Dean and Boddy 1988]. The algorithms presented in this paper have this property.

We discuss a method for generating the GLB first. The following notion is central to our approach:

Definition 2.3.2. Horn-strengthening. A Horn clause C_H is a Horn-strengthening of a clause C iff $C_H \subseteq C$ and there is no Horn clause C'_H such that $C_H \subset C'_H \subseteq C$. A Horn-strengthening of a set of clauses $\{C_1, \dots, C_n\}$ is any set of clauses $\{C'_1, \dots, C'_n\}$ where C'_i is a Horn-strengthening of C_i .

Thus, a Horn-strengthening of a non-Horn clause is formed by removing all but one positive literal.

Example. Consider the clause $C = \{p, q, \neg r\}$. The clauses $\{p, \neg r\}$ and $\{q, \neg r\}$ are the Horn-strengthenings of C .

The following two lemmas state some useful properties of Horn-strengthenings. The first lemma shows that a Horn theory entails a clause only if it entails some Horn-strengthening of the clause.

LEMMA 2.3.3. *Let Σ_H be a Horn theory and C a clause that is not a tautology. If $\Sigma_H \models C$, then there is a clause C_H that is a Horn-strengthening of C such that $\Sigma_H \models C_H$.*

Generate_GLB

Input: a set of clauses $\Sigma = \{C_1, C_2, \dots, C_n\}$.

Output: a greatest Horn lower-bound of Σ .

begin

$L :=$ the lexicographically first Horn-strengthening of Σ

loop

$L' :=$ lexicographically next Horn-strengthening of Σ

if none exists then exit loop

if $L \models L'$ then $L := L'$

end loop

remove subsumed clauses from L

return L

end

FIG. 3. Algorithm for generating a greatest Horn lower-bound.

C_H . More generally, if a Horn theory entails a set of clauses, then it entails a Horn-strengthening of the set.

PROOF. By the subsumption theorem [Lee 1967], there is a clause C' that follows from Σ_H by resolution such that C' subsumes C . Because the resolvent of Horn clauses is Horn, C' is Horn. Therefore, there is some C_H that is a Horn-strengthening of C such that $C' \subseteq C_H \subseteq C$, and $\Sigma_H \models C_H$. The generalization to sets of clauses follows immediately. \square

The next lemma shows that every GLB of a theory is equivalent to some Horn-strengthening of the theory.

LEMMA 2.3.4. Let Σ_{glb} be a GLB of a theory $\Sigma = \{C_1, \dots, C_n\}$. Then there is a set of clauses C'_1, \dots, C'_n such that $\Sigma_{glb} \equiv \{C'_1, \dots, C'_n\}$, where C'_i is a Horn-strengthening of C_i .

PROOF. By the definition of a GLB, we have $\Sigma_{glb} \models \Sigma$. Therefore, $\Sigma_{glb} \models C_i$ for $1 \leq i \leq n$. Because Σ_{glb} is Horn, by Lemma 2.3.3, there exists $\{C'_1, \dots, C'_n\}$, where each C'_i is a Horn-strengthening of C_i , such that $\Sigma_{glb} \models C'_i$ for $1 \leq i \leq n$. Thus $\Sigma_{glb} \models \{C'_1, \dots, C'_n\} \models \Sigma$. Moreover, because Σ_{glb} is a greatest lower-bound, it must be the case that $\Sigma_{glb} \equiv \{C'_1, \dots, C'_n\}$. \square

Moreover, it is not difficult to see that every Horn-strengthening of Σ is a Horn lower-bound of Σ —though not necessarily the greatest lower-bound. This leads us to the **Generate_GLB** algorithm given in Figure 3. The algorithm systematically searches through the various possible Horn-strengthenings of the clauses of the original theory, looking for a most general one. Where $\Sigma = \{C_1, C_2, \dots, C_n\}$ and C'_j is the j th Horn strengthening of clause C_i , the Horn strengthenings of Σ are generated in the lexicographic order $\{C_1^1, C_2^1, \dots, C_n^1\}, \{C_1^2, C_2^1, \dots, C_n^1\}, \{C_1^1, C_2^2, \dots, C_n^1\}, \{C_1^2, C_2^2, \dots, C_n^1\}$, etc.

THEOREM 2.3.5. Given a set of clauses Σ , **Generate_GLB** (Figure 3) computes a greatest Horn lower-bound of Σ of length less than or equal to that of Σ .

Example. Consider the theory $\Sigma = (\neg a \vee b \vee c) \wedge (a \vee b)$. The algorithm first tries the Horn-strengthening $L = ((\neg a \vee b) \wedge a) \equiv (a \wedge b)$, and then $L' = ((\neg a \vee b) \wedge b) \equiv b$. Since $L \models L'$, L is set to L' , and the algorithm proceeds. Since the other two Horn strengthenings do not further improve the bound, the algorithm returns b as an answer (the redundant clause $(\neg a \vee b)$ is removed by the last step of the algorithm).

PROOF OF THEOREM 2.3.5. Ignoring the last step that removes subsumed clauses from L , it is obvious that **Generate_GLB** returns a logically-weakest Horn-strengthening of Σ . So to prove the algorithm correct, we need to show that any logically-weakest Horn-strengthening of Σ is a GLB of Σ . First note that $L \models \Sigma$ and L is Horn so L is a LB. Suppose it were not a greatest LB. There is a set of Horn clauses H such that

$$\begin{aligned} L &\models H \\ H &\not\models L \\ H &\models \Sigma. \end{aligned}$$

By Lemma 2.3.3, there must be some Horn strengthening Σ' of Σ such that

$$H \models \Sigma'.$$

But this means that

$$L \models \Sigma'.$$

Now there are two cases: (1) L and Σ' are logically equivalent. Then this means that L and H are logically equivalent, which is a contradiction. (2) L and Σ' are logically distinct. But then L is not a logically weakest Horn strengthening of Σ , which is also a contradiction. This proves the theorem. \square

Generate_GLB is indeed an anytime algorithm: L represents some lower-bound whenever the algorithm is interrupted. Note also that the total running time of the algorithm is exponential only in the length of the non-Horn part of the original theory, because the only strengthening of a Horn clause is the clause itself.

The querying algorithm shown in Figure 2 uses only a single GLB of the theory. One could extend the algorithm to use several lower bounds (although not necessarily all of them, because certain theories can have exponentially many distinct GLBs). There is a simple and interesting relationship between a theory and the set of all of its GLBs. This follows from the fact that each model of a theory corresponds to a lower-bound of the theory, namely the conjunction of all the literals assigned true by the model. We thus obtain the following theorem:

THEOREM 2.3.6. *Let Σ be a set of clauses. Then Σ is logically equivalent to the disjunction of all the greatest Horn lower-bounds of Σ .*

We now turn our attention to the generation of the LUB. We will use the notion of a *prime implicate* of a theory, which is a strongest clause implied by the theory. The following theorem reveals our basic strategy.

```

Generate_LUB
Input: a set of clauses  $\Sigma = \Sigma_H \cup \Sigma_N$ , where
           $\Sigma_H$  is a set of Horn clauses, and  $\Sigma_N$  is a
          set of non-Horn clauses.
Output: a least Horn upper-bound of  $\Sigma$ .
begin
  loop
    try to choose clause  $C_0 \in \Sigma_N$  and  $C_1 \in \Sigma_H \cup \Sigma_N$ ,
    such that  $C_2 = \text{Resolve}(C_0, C_1)$ 
    is not subsumed by any clause in  $\Sigma_H \cup \Sigma_N$ 
    if no such choice is possible then exit loop
    delete from  $\Sigma_H$  and  $\Sigma_N$  any clauses
    subsumed by  $C_2$ 
    if  $C_2$  is Horn then
       $\Sigma_H := \Sigma_H \cup \{C_2\}$ 
    else
       $\Sigma_N := \Sigma_N \cup \{C_2\}$ 
    end if
  end loop
  return  $\Sigma_H$ 
end

```

FIG. 4. Algorithm for generating the least Horn upper-bound.

THEOREM 2.3.7. *Let Σ be a set of clauses. The LUB of Σ is logically equivalent to the set of all Horn prime implicates of Σ .³*

PROOF. The set of Horn prime implicates is implied by Σ , and thus is a Horn upper-bound. Furthermore, it must be the LUB, because at least one of its clauses subsumes (and therefore implies) any clause in any Horn upper-bound. \square

So, in principle, we could use resolution to generate the prime implicates and simply collect the Horn ones in order to generate the least Horn upper-bound. However, such a method could prove very expensive since even if the original theory contains only Horn clauses, there can be exponentially many Horn prime implicates (for an example, see Selman [1991]). Clearly, such prime implicates add nothing to the best approximation of the original Horn theory, since the least Horn upper-bound is already given by the theory itself.

Fortunately, we can improve upon the procedure of generating all prime implicates by only resolving two clauses if at least one is non-Horn. The intuition behind this improvement is that any resolution proof tree that ends in a Horn clause can be rewritten so that all the resolutions between pairs of Horn clauses appear at the bottom of the tree. The compilation phase essentially generates the top part of the tree, and the all-Horn part is generated at query-answering time. The full inductive proof that this method is complete was provided by del Val [1995]. As with **Generate_GLB**, the algorithm is anytime: Σ_H improves over time.

THEOREM 2.3.8. *Given a set of clauses Σ , **Generate_LUB** algorithm (Figure 4) computes the least Horn upper-bound of Σ .*

³ It follows directly that the LUB is complete with respect to Horn queries. That is, for any Horn clause α , we have $\Sigma \models \alpha$ if and only if $\Sigma_{\text{lub}} \models \alpha$.

Example. Consider the theory $(\neg a \vee b) \wedge (\neg b \vee c) \wedge (a \vee b)$. The **LUB algorithm** resolves the first and the third clause, obtaining the clause b . Σ_H becomes $(\neg b \vee c) \wedge b$, upon which the loop is exited and Σ_H is returned.

2.4. COMPUTATIONAL SAVINGS. In Section 2.2, we described how the bounds can be used to answer queries, and claimed that doing so could lead to computational savings. One might wonder whether it could be the case that queries that can be answered by the bounds are also easy to answer using the original theory. A central aspect of our approach is that this is *not* so. An obvious counterexample is any inconsistent theory. Compilation yields an inconsistent upper bound. Any query against this bound would quickly return “yes” (see Figure 2). However, evaluating a query against the original theory would in general involve proving that the theory was inconsistent, which is co-NP-complete.

Of course, most interesting knowledge bases will be consistent. Let us therefore consider a consistent theory that is equivalent to a Horn theory, but is not in Horn form. Clearly, all queries can be answered efficiently against the bounds. However, it is *not* the case that a theorem prover could also answer queries efficiently against the original theory. This can be shown using a result by Valiant and Vazirani [1986]. They show that even if a propositional theory has a single model (and is thus trivially equivalent to a Horn theory), finding the model is still intractable (unless $NP = RP$, which is unlikely). Therefore, there cannot exist a theorem prover that efficiently handles this special case, because such a prover could be used to find the unique model of the non-Horn theory (by repeatedly testing whether each literal followed from the theory).

The fact that the bounds obtained by our method can be used to answer queries efficiently that would be intractable to test against the original theory is a key feature that distinguishes our work from most related work on reasoning with approximations (related work is discussed in Section 4 below). The most natural application of our framework is therefore in domains where fast response to queries is of the highest importance. Direct practical evidence for this can be found in the work by de Kleer [1990] on qualitative physics. De Kleer uses an Assumption-Based Truth Maintenance System (ATMS), a kind of propositional theorem prover, to reason about logical theories of qualitative physics. The main form of inference of the ATMS is so-called Boolean Constraint Propagation (BCP). BCP is sound and complete for Horn clauses, but incomplete for more general theories. In order to make BCP more useful, the system preprocesses the theory using a limited form of resolution. For example, given the clauses $x \vee \neg y \vee z$ and $\neg x \vee \neg y \vee z$, one can add their resolvent $\neg y \vee z$. This additional clause will allow BCP to derive additional information. Since BCP in effect ignores non-Horn clauses, the process can be viewed as computing a form of a Horn upper-bound approximation to the original theory. De Kleer’s approach is thus a concrete example of the practical usefulness of a limited form of knowledge compilation.⁴

⁴ De Kleer also discusses the possibility of replacing the entire theory by its set of prime implicates. Though this may be useful for abductive reasoning [Reiter and de Kleer 1987; Selman and Levesque 1990; 1996], it is not necessary for deductive reasoning because Horn clauses themselves already allow for efficient inference. Also note that generating all the prime implicates may require exponential space.

TABLE I. SIZE OF BOUNDS AND PERCENTAGE OF RANDOM QUERIES ANSWERED BY UNIT BOUNDS FOR HARD RANDOM 3CNF THEORIES.

vars	clauses	size unit		percent queries answered		
		LUB	GLB	unit	binary	ternary
75	322	53	71	100%	85%	88%
100	430	57	93	100%	76%	79%
150	645	62	139	100%	66%	66%
200	860	132	188	100%	83%	85%

Note: The size of the unit LUB and unit GLB were determined empirically. The table gives the median values over 10 random theories of each size. The percentage of random queries answered by the bounds was then computed probabilistically, as described in the text.

In order to gain some further insight into the practical usefulness of knowledge compilation, we now consider the compilation of hard, randomly generated propositional theories. Mitchell et al. [1992] show that most randomly generated theories are easy to reason with. Such theories tend to be either very over-constrained or very under-constrained; in either case, experiments show that answering queries is easy using the standard Davis–Putnam procedure [Davis and Putnam, 1960].⁵ However, Mitchell et al. also described how to randomly generate computationally challenging theories. The key is to generate formulas with a particular ratio of clauses to variables. For random 3CNF formulas, the ratio is about 4.3 [Kirkpatrick and Selman 1994; Selman et al. 1996; Selman and Kirkpatrick 1996]. We consider hard random 3CNF theories containing between 75 and 200 variables. In order to simplify the following analysis, we computed bounds that consisted of conjunctions of unit clauses. Note that unit clauses are a restricted case of Horn clauses. Therefore, these bounds are not as tight as the full Horn bounds. We will show that even these bounds are useful for answering a high percentage of queries. Because the full Horn bounds are tighter, they would answer an even higher percentage of queries. However, by considering the unit clause bounds we are able to provide a simple exact analysis.

We began by generating a set of 40 random 3CNF theories, with 10 each based on 75, 100, 150, and 200 variables. Then, we computed the unit LUB and a unit GLB of each. Table I gives the average size, in literals, of the LUB and GLB for each size theory.

We computed the unit upper bound using a version of the **Generate_LUB** algorithm that was restricted to output only unit clauses. **Generate_GLB** was similarly adapted for unit bounds; in particular, the first strengthening of the theory was taken to be a satisfying model of the theory found using a fast randomized satisfiability procedure [Selman et al. 1992; Selman and Kautz 1993].

Computation time for the unit LUBs ranged from 5 minutes for the 75 variable theories, to one hour for the 200 variable theories, on a 100-MHz SGI Challenge. Computation of the unit GLBs ranged from 1 minute to 5 minutes each.

⁵ If the theory is over-constrained, it is generally unsatisfiable, so that all queries trivially follow. If it is under-constrained and the CNF query contains short disjunctions, then the query almost certainly does not follow. Finally, if the theory is under-constrained and the CNF query contains only long disjunctions, then the query almost certainly does follow, which can be easily shown by adding the negation of each clause in the query to the theory and using the Davis–Putnam procedure with unit propagation to show inconsistency.

The percentage of queries that could be answered by these bounds is computed using some basic probability theory. We assume that we are dealing with single-clause queries drawn from a uniform fixed-clause length distribution. The simplest case is the unit clause queries. All unit clause queries can be answered using only the unit LUB, because this bound is complete for such queries. Thus, this column is 100% for every size theory.

Next, let us consider the more interesting case of binary queries. Let $x \vee y$ be a random binary clause, where x and y are distinct and not complements. We wish to compute the probability that the bounds answer the query, given that the unit LUB is of size l and the unit GLB is of size m , and there are N variables in the theory. That is, we wish to compute

$$\Pr((\Sigma_{\text{lub}} \vdash x \vee y) \quad \text{or} \quad (\Sigma_{\text{glb}} \not\vdash x \vee y)),$$

which equals

$$\Pr(\Sigma_{\text{lub}} \vdash x \vee y) + \Pr(\Sigma_{\text{glb}} \not\vdash x \vee y)$$

because the two possibilities are disjoint. A disjunction is entailed by a set of literals if and only if one of the disjuncts is so entailed. Thus,

$$\Pr(\Sigma_{\text{lub}} \vdash x \vee y) = \Pr((\Sigma_{\text{lub}} \vdash x) \text{ or } (\Sigma_{\text{lub}} \vdash y)).$$

Note that we were only able to make this step because we are considering the unit bounds, instead of the full Horn bounds. This quantity is equal to

$$\Pr(\Sigma_{\text{lub}} \vdash x) + (\Sigma_{\text{lub}} \vdash y) - (\Sigma_{\text{lub}} \vdash x \wedge y).$$

The first and second terms are equal to the odds of picking a random literal that is in the LUB, namely $l/(2N)$. The third term is equal to the number of ways of choosing two distinct literals from the LUB, divided by the number of ways of choosing two distinct, noncomplementary literals, namely $l(l-1)/((2N)2(N-1))$. Thus,

$$\Pr(\Sigma_{\text{lub}} \vdash x \vee y) = \frac{l}{2N} + \frac{l}{2N} - \frac{l(l-1)}{4N(N-1)}.$$

Next, consider the probability that a binary query is answered by the GLB. We have

$$\Pr(\Sigma_{\text{glb}} \not\vdash x \vee y) = 1 - \Pr(\Sigma_{\text{glb}} \vdash x \vee y).$$

The final term above can be calculated using the same argument as for the LUB to yield

$$\Pr(\Sigma_{\text{glb}} \not\vdash x \vee y) = 1 - \left(\frac{m}{2N} + \frac{m}{2N} - \frac{m(m-1)}{4N(N-1)} \right).$$

Combining the probability that the LUB answers the query with the probability that the GLB answers the query results in the expression

TABLE II. TIME IN SECONDS TO ANSWER 1000 RANDOM QUERIES USING THE BOUNDS TOGETHER WITH THE PROGRAM TABLEAU (A VERSION OF THE DAVIS-PUTNAM ALGORITHM) VERSUS USING TABLEAU ALONE. (ALL EXPERIMENTS WERE ON A 100-Mz SGI CHALLENGE WORKSTATION.)

vars	clauses	bounds and tableau		tableau only	
		binary	ternary	binary	ternary
75	322	51	48	258	248
100	430	54	45	368	341
150	645	61	59	1286	1084
200	860	55	51	12962	8632

$$1 - \frac{4N(m-1) - 3(m-1) + l^2 - m^2}{4N(N-1)}$$

The value of this expression was used to complete the “binary” column of Table I.

The probability that the bounds answer a random *ternary* query can be similarly derived, and was used to complete the final column of the table. As we have noted, a similar analysis could be performed for Horn bounds, but would be much more complicated, because the probability that a Horn theory entails a disjunction is not the same as the probability that it entails one of the disjuncts.

As we can see from Table I, the percentage of queries that can be handled by even such simple unit clause bounds is quite high. Note that the queries handled by the bounds can be answered in linear time. However, the Davis–Putnam procedure scales exponentially on the queries considered in the table (this follows from the experiments in Mitchell et al. [1992]). Thus, this suggests that knowledge compilation on such randomly-generated theories should have a clear payoff.

We verified the computational savings suggested by the preceding analysis by implementing the fast querying algorithm shown in Figure 2, and testing 1000 random binary and 1000 random ternary queries against each of the 40 test theories.

In case both bounds failed to answer a query, it was tested against the original theory using an efficient implementation of the Davis–Putnam procedure called tableau.⁶ Table II lists the average time to run each batch of a 1000 queries, using the bounds together with tableau versus using tableau alone. Thus, in both cases *all* queries were answered. We see that knowledge compilation reduced the overall time by over two orders of magnitude on the largest theories. This eliminates the remote possibility that the bounds are only answering the “easy” queries. Earlier we invoked complexity theory to argue that in *general* the bounds are not limited to easy queries; these experiments verify that the bounds answer hard queries against a computationally interesting distribution of random theories.

⁶ The Davis–Putnam procedure is currently the fastest known complete procedure for propositional satisfiability testing and theorem-proving on the class of formulas considered here [Buro and Büning 1992; Dubois et al. 1995]. Tableau [Crawford and Auton 1993] is one of the fastest implementations of the algorithm.

As an aside, we observe that even when we take into account the time required to compile the theories, we obtain an overall time savings. For example, on the 200 variable theories, computing the bounds takes about an hour and five minutes; thus, the total time to compute the bounds *and* answer 1000 binary queries is 3,955 seconds, versus 12,962 seconds not using the bounds. (Note that difference in overall time will increase even further when we would consider, for example, 10,000 queries.) Thus, in this case, we have gone beyond our stated goal of shifting computational effort from on-line to off-line, and have actually reduced the total amount of work required.

These positive results for random theories are quite surprising, since one would expect that their apparent lack of structure would make them hard to approximate by simpler theories. A further understanding of the costs and benefits of the knowledge compilation approach will come through the application of the techniques to various real-world problems [Kautz and Selman 1994]. Note that, in practice, one may be satisfied with upper and lower bounds that are not necessarily optimal, but are cheaper to compute. Greiner and Schurmans [1992] gives compilation algorithms that begin answering queries with nonoptimal bounds, and gradually improve them at run-time depending on the particular query distribution.

2.5. SIZE OF THE APPROXIMATIONS. The size of the approximations is clearly an important issue. Ideally, one would like the approximations to be of roughly the same size or less than the original theory. If the approximations are much larger, then one may lose the potential computational advantage. For example, if a bound is exponentially larger than the original theory, evaluating a query in linear time using the bound would be as expensive as using the original theory.

From the **Generate_GLB** algorithm it follows immediately that the size of the generated bound is less than or equal to that of the original theory.⁷ The system can safely use this bound; even if the approximation does not provide an answer for a particular query, at most a linear amount of time (in terms of the length of the original theory) is being “wasted.” The LUB is less well behaved. First, we will show that there are theories that have exponentially long upper-bounds. However, such bounds can often be shrunk down to the size of the original theory by introducing new letters in the vocabulary. Such new letters can be viewed as useful abbreviations for repeated structure in the theory. Finally, we will prove that, unfortunately, one cannot *always* shrink down the LUB by introducing new letters. In those cases, the system has to limit the length of the approximation by using a weaker bound. (We discuss some kinds of weaker bounds in Section 3; see also Greiner and Schuurmans [1992] for an example of this.)

⁷ Recently, Cadoli [1993], building on our original paper on knowledge compilation [Selman and Kautz, 1991], has shown that computing the GLB is in the complexity class P^{NP} . (Problems in P^{NP} can be solved in polynomial time by a deterministic Turing machine with access to an NP oracle. See Garey and Johnson [1979] or Johnson [1990].) This means that computing a GLB is only slightly harder than answering a single query, since the best-known algorithms for both problems are singly exponential (D. S. Johnson, personal communications). Note that is indeed a positive result, because one might have feared that, in the worst case, a doubly exponential amount of work was needed to compute a bound that can potentially handle an exponential number of queries efficiently.

2.5.1. *Explosion of the LUB.* The following theory has an LUB with exponentially many clauses. The clauses in the theory can be interpreted as rules for deciding if someone is a cognitive scientist. The clauses are numbered for ease of reference.

$$(\text{CompSci} \wedge \text{Phil} \wedge \text{Psych}) \supset \text{CogSci} \quad (1)$$

$$\text{ReadsMcCarthy} \supset (\text{CompSci} \vee \text{CogSci}) \quad (2)$$

$$\text{ReadsDennett} \supset (\text{Phil} \vee \text{CogSci}) \quad (3)$$

$$\text{ReadsKosslyn} \supset (\text{Psych} \vee \text{CogSci}) \quad (4)$$

Clause (1) states a sufficient condition for being a cognitive scientist: being a computer scientist, *and* a philosopher, *and* a psychologist. The remaining clauses let one deduce a person's profession from his or her reading habits. Clause (2) states that if a person reads papers written by McCarthy, then the person is either a computer scientist or a cognitive scientist (or possibly both). Similarly, a reader of Dennett is either a philosopher or cognitive scientist or both, and a reader of Kosslyn is either a psychologist or cognitive scientist or both.

Reasoning with this theory can be quite complicated. For example, by reasoning by cases, one can prove that a computer scientist who reads Dennett and Kosslyn is also cognitive scientist. In general, for such non-Horn form theories, finding a proof will take time exponential in the length of the theory (provided $P \neq NP$).

Clause (1) can be resolved with subsets of clauses (2–4) to yield many different Horn clauses, such as

$$(\text{ReadsMcCarthy} \wedge \text{Phil} \wedge \text{ReadsKosslyn}) \supset \text{CogSci}$$

$$(\text{CompSci} \wedge \text{ReadsDennett} \wedge \text{Psych}) \supset \text{CogSci}.$$

In fact, the LUB of this theory is equivalent to the set of 2^3 Horn clauses:

$$\left\{ \begin{array}{l} (p \wedge q \wedge r) \supset \text{CogSci} \mid \\ p \in \{\text{CompSci}, \text{ReadsMcCarthy}\} \\ q \in \{\text{Phil}, \text{ReadsDennett}\} \\ r \in \{\text{Psych}, \text{ReadsKosslyn}\} \end{array} \right\}. \quad (5)$$

Furthermore, we can prove that there is *no* smaller set of clauses equivalent to (5). Note that this is a much stronger condition than simply saying that there are no redundant clauses in (5); we are asserting that there is no way to represent the same information in less space in clausal form. This example can be generalized by increasing the number of propositions that appear in the left-hand side of clause (1), and adding a clause of the general form of (2) for each new proposition. The size of the LUB then grows exponentially. In general, we have the following theorem.

THEOREM 2.5.1.1. *There exist clausal theories Σ of size n such that the smallest clausal representation of their LUB is of size $O(2^n)$.*

The proof appears in the appendix. Therefore, although we can tell if any clause follows from the LUB in time linear in the size of the LUB, the explosion in size of the LUB in this example wipes out our savings. Of course, there are also many commonsense theories for which such exponential blowup does *not* occur.

2.5.2. Shrinking the LUB. The syntactic form of a theory can often be modified without changing its basic meaning. For example, any theory can be represented by a set of clauses each containing no more than three literals (3-CNF form) by introducing new propositional letters. The old and new theories are not equivalent, since the new one uses an expanded vocabulary, but they are essentially the same: They both entail or both do not entail any formula that does not contain any of the new letters.

Thus, one might wonder if a large LUB could be represented by a small set of Horn clauses that have basically the same meaning, if not actual logical equivalence. As with the case of 3-CNF formulas, the technique we use depends on the introduction of new propositional letters. Rather than modify the definition of the LUB, we will add these new letters to the source theory itself. If we take the meaning of a letter to be a concept, we will see that the method reduces to the definition of new concepts that generalize concepts in the original theory.

For example, let us modify the theory given by clauses (1–4) by introducing three new concepts, “computer science buff,” “philosophy buff,” and “psychology buff.” The first generalizes the concepts of a computer scientist and of a reader of papers by McCarthy. Similarly, the second generalizes philosopher and reader of Dennett, and the third generalizes psychologist and reader of Kosslyn. Each concept definition requires three clauses: one to assert that the more general concept is divided among its subconcepts, and two to assert that the subconcepts are part of the concept. The added clauses are

$$\text{CompSciBuff} \supset (\text{CompSci} \vee \text{ReadsMcCarthy}) \quad (6)$$

$$\text{CompSci} \supset \text{CompSciBuff} \quad (7)$$

$$\text{ReadsMcCarthy} \supset \text{CompSciBuff} \quad (8)$$

$$\text{PhilBuff} \supset (\text{Phil} \vee \text{ReadsDennett}) \quad (9)$$

$$\text{Phil} \supset \text{PhilBuff} \quad (10)$$

$$\text{ReadsDennett} \supset \text{PhilBuff} \quad (11)$$

$$\text{PsychBuff} \supset (\text{Psych} \vee \text{ReadsKosslyn}) \quad (12)$$

$$\text{Psych} \supset \text{PsychBuff} \quad (13)$$

$$\text{ReadsKosslyn} \supset \text{PsychBuff}. \quad (14)$$

The LUB of the augmented theory containing (1–4) and the clauses above can be represented by just the Horn clauses from the new concept definitions (7, 8, 10, 11, 13, 14) together with the single clause

$$(\text{CompSciBuff} \wedge \text{PhilBuff} \wedge \text{PsychBuff}) \supset \text{CogSci}. \quad (15)$$

The total size of the LUB (counting all occurrences of literals) is half the size of the LUB of the unmodified theory. As we generalize the example by adding new propositions, the size of the LUB grows as $O(n)$, as opposed to $O(2^n)$. Furthermore, the added clauses only increase the size of the source theory by a linear amount.

Returning to our example concerning a computer scientist who reads Dennett and Kosslyn, we can now infer quickly from (7), (11), and (14) that the person is a computer science buff, philosophy buff, and psychology buff, and therefore by (15) a cognitive scientist. Note that this inference can be computed in time linear in the size of the new LUB and therefore linear in the size of the *original* theory (1–4). Thus, by teaching the system new concept definitions, the size of the new source theory grows only slightly, and the LUB shrinks to approximately the size of the source theory itself.

In Kautz and Selman [1992a], we discuss a strategy for *automatically* deriving (“learning”) new concept letters that shrink the original bound. We call this *theory compaction*. The flavor of our approach is very similar to that of the use of additional letters in extended resolution [Tseitin 1968]. See Muggleton and Buntine [1988] for a related approach to learning new generalizations, based on inverting resolution proofs.

2.5.3. Does a Compact Representation of the LUB Always Exist? So far, we have shown that a naive representation of a theory’s LUB can sometimes require an exponential amount of space, and that in some of those cases a clever representation using new propositional letters requires only a polynomial amount of space. The question then becomes how general these results are. Is it *always* possible to produce a compact, tractable representation of the LUB? Unfortunately, the following theorem shows that this is unlikely.

THEOREM 2.5.3.1. *Unless $NP \subseteq$ nonuniform P, it is not the case that the Horn least upper-bound Σ_{lub} of a propositional clausal theory Σ can always be represented by a data structure that allows queries of the form*

$$\Sigma_{lub} \models \alpha$$

to be answered in time polynomial in $(|\Sigma| + |\alpha|)$, where α is a single Horn clause. Note that this is so despite the fact that we allow an arbitrary amount of work to be performed in computing the data structure used to represent Σ_{lub} . The proof of the theorem appears in the appendix.

The notion of “nonuniform P” comes from work in circuit complexity [Boppana and Sipser 1990]. A problem is in nonuniform P (also called P/poly) iff for every integer n there exists a circuit of complexity (size) polynomial in n that solves instances of size n . The adjective “nonuniform” refers to the fact that different circuits may be required for different values of n . Any problem that can be solved by an algorithm that runs in time $O(f(n))$ has circuit complexity $O(f(n)\log n)$. We use this fact implicitly in the proof of the theorem, where we talk about polynomial time algorithms rather than polynomial size circuits.

The class nonuniform P is, however, considerably larger than P. (For example, non uniform P contains noncomputable functions, such as the function that returns “1” on inputs of length n iff Turing Machine number n halts on all inputs. For any n the circuit is simply fixed to return 1 or 0.) Although it is possible that

$P \neq NP$ and yet $NP \subseteq$ nonuniform P , this is considered unlikely. One consequence would be that the polynomial-time hierarchy would collapse to Σ_2^P [Karp and Lipton 1982]. As shown in the appendix, the theorem can be strengthened to say that the claim that there always exists an efficient form of the LUB for answering Horn clausal queries is *equivalent* to the claim that $NP \subseteq$ nonuniform P . Therefore, a proof that such efficient representations do or do not always exist would be a major result in the complexity of finite functions.

An immediate corollary of the theorem is that unless the polynomial hierarchy collapses in this manner, compaction by defining new propositions is an incomplete method for shrinking the LUB.

COROLLARY 2.5.3.2. *Unless $NP \subseteq$ nonuniform P , it is not the case that there is always a compaction (using any number of new variables) of the Horn least upper-bound of a theory Σ that is of size polynomial in $|\Sigma|$.*

This follows from the theorem because one can determine if a Horn clause follows from a compaction (which is itself Horn) in time linear in the length of the compaction plus the length of the query.

To summarize, there are theories for which the Horn LUB cannot be represented in a short, tractable form. In such cases, knowledge compilation can still be of use, but one will have to rely on a shorter and therefore weaker upper-bound. One way to do this would be to simply stop generating the bounds at some point; this is straightforward, given the anytime nature of our compilation algorithms. Another alternative would be to generate so-called k -Horn approximations, which are Horn theories with at most k literals per clause [Dechter and Pearl 1992]. Given the limitation on the clause lengths, such theories are of most polynomial size in the number of variables. (We discuss k -Horn and other alternatives to Horn approximations in Section 3 below.)

3. General Framework

Central to our knowledge compilation proposal is the definition of approximations in terms of subsets and supersets of the set of models of the original theory. This was illustrated in Figure 1. Such a model-theoretic definition leaves much freedom in terms of choice of representation language for the original theory, the approximations, and the queries. Our only additional requirement is that the queries can be evaluated efficiently with respect to the approximations. Note again that we view the computation of the approximations as an off-line compilation process, so that its cost can be amortized over the set of queries. By allowing substantial computational resources to be spend on the compilation the approximations will become of real practical significance, as was shown in Section 2.4.

Formally, a *knowledge compilation system* is a tuple $\langle \mathcal{L}, \models, \mathcal{L}_S, \mathcal{L}_T, \mathcal{L}_Q, f_L, f_U \rangle$ containing the following components:

- \mathcal{L} is a formal language. We identify a language with the set of all its sentences.
- \models is a consequence relation over sets of sentences in \mathcal{L} . In most of the examples we will study, \models has its usual meaning of logical entailment, but the framework allows \models to represent other relationships, such as subsumption.

- \mathcal{L}_S is the “source” sublanguage of \mathcal{L} , used to express a general knowledge base.
- \mathcal{L}_T is the “target” sublanguage of \mathcal{L} , used to express approximations to the general knowledge base. It should be easier in some sense (analytically or empirically) to determine if a query is a consequence of a set of sentences in \mathcal{L}_T than of a set sentences in \mathcal{L}_S .
- \mathcal{L}_O is the “query” sublanguage of \mathcal{L} .
- f_L is a function that models some (particular) anytime lower-bound algorithm, mapping a source theory and time-bound n to the best lower-bound that can be computed in n steps. Formally, $f_L: \mathcal{L}_S \times \mathbb{N} \rightarrow \mathcal{L}_T$. By definition, $f_L(\Sigma, \infty)$ is a GLB of Σ .
- f_U is a similar upper-bound function, where $f_U(\Sigma, \infty)$ is the LUB of Σ .

Suppose a knowledge compilation system is presented a query $\alpha \in \mathcal{L}_O$ after performing i compilation steps on the source theory Σ . As described earlier, if $f_L(\Sigma, i) \not\models \alpha$ then the system answers “no”, if $f_U(\Sigma, i) \models \alpha$ then the system answers “yes”, and otherwise it answers “unknown.” A definite answer of “yes” or “no” always agrees with the answer to the question “Does $\Sigma \models \alpha$?”, even if the lower (upper) bound is not the greatest (least) bound.

The case of knowledge compilation using propositional Horn clauses fits into this framework as follows: \mathcal{L} , \mathcal{L}_S , and \mathcal{L}_O are general, propositional, clausal languages, \mathcal{L}_T is restricted to Horn clauses, and \models is propositional entailment. In this case the query language is strictly more expressive than \mathcal{L}_T . (Note that to answer a general query consisting of a set of clauses, one can determine efficiently whether each of the clauses is implied by the Horn approximation, by using the linear time algorithm for testing the satisfiability of Horn clauses.) The functions f_L and f_U model the **Generate_GLB** and **Generate_LUB** algorithms presented above.

3.1. OTHER INSTANCES OF KNOWLEDGE COMPILATION

3.1.1. *Restricted Propositional Clausal Forms.* A broad class of propositional knowledge compilation systems can be realized by generalizing the algorithms for the Horn case. The idea of a Horn-strengthening is generalized as follows:

Definition 3.1.1.1. θ -strengthening. Let θ be a particular class of propositional clauses. A clause in that class is called a θ -clause, and a set of such clauses is called a θ -theory. A clause C_θ is a θ -strengthening of a clause C iff $C_\theta \subseteq C$ and there is no θ -clause C'_θ such that $C_\theta \subset C'_\theta \subseteq C$.

A knowledge compilation system can be created by letting the source language \mathcal{L}_S be the set of general clauses, and the target language \mathcal{L}_T be such a restricted class of clauses θ . Now further suppose that θ is closed under resolution (i.e., the resolvent of any two θ -clauses is a θ -clause), and that every clause has at least one θ -strengthening. Then we see that the proof of Lemma 2.3.3 generalizes from Horn clauses to θ -clauses: that is, if a θ -theory entails a clause C , then it must also entail a θ -strengthening of C . Furthermore, every greatest lower-bound in terms of θ -clauses of a theory is equivalent to a (logically weakest) θ -strengthening of the theory. Therefore, one can find a GLB by searching the space of θ -strengthenings.

THEOREM 3.1.1.2. *Let θ be a class of propositional clauses such that (i) the resolvent of any two clauses in θ is also in θ , and (ii) every clause is subsumed by some clause in θ . Given a set of clauses Σ , **Generate_GLB** (Fig. 3) computes a greatest θ -lower-bound Σ of length less than or equal to that of Σ (where “ θ ” is substituted for “Horn” in the algorithm).*

Now consider the problem of generating the least upper-bound of a clausal theory. Recall that the “naive” algorithm for finding the Horn LUB of a theory is simply to compute all resolvents of the theory, and then collect the Horn prime implicates from that set. This is correct because of the resolution completeness theorem and the fact that propositional Horn clauses are closed under *subsumption*, that is, a Horn clause can only be subsumed by another Horn clause. Therefore, a condition we place on the general target language θ is that it be closed under subsumption. (An example of a class that violates this condition is the set of definite Horn clauses. These are clauses that contain exactly one positive literal. The resolvent of any two definite Horn clauses is a definite Horn clause. However, a definite Horn clause (for example, $\neg p \vee q$) can be subsumed by a non-definite Horn clause (for example, $\neg p$)). We summarize these observations about generating the θ -LUB in the following theorem:⁸

THEOREM 3.1.1.3. *Let θ be a class of propositional clauses closed under subsumption and such that if $C \in \theta$ and C' subsumes C , then $C' \in \theta$. Given a set of clauses Σ , **Generate_LUB** (Fig. 4) computes the least θ -upper-bound of Σ (where “ θ ” is substituted for “Horn” in the algorithm).*

Several useful clausal target languages meet the conditions of Theorems 3.1.1.2 and 3.1.1.3. These include:

- Horn and Reverse-Horn (clauses containing at most one *negative* literal).
- Binary*. Clauses containing two or fewer literals. Satisfiability and entailment can be determined in linear time for this class [Aspvall et al. 1979]. (Note that unit clauses and the empty clause are included in this class.)
- Unit clauses*. That is, bounds that consist of conjunctions of literals. This class is a specialization of both Horn and binary.
- Clauses not containing a given set of propositional letters*. While such a class may not have better worst-case complexity than the unrestricted language, it may be empirically desirable to “compile away” certain irrelevant propositions.⁹ For example, let $\Sigma = \{p \vee q, \neg q \vee r\}$, and θ be clauses not containing q . Then the θ -GLB of Σ is $\{p \wedge r\}$, and θ -LUB is $\{p \vee r\}$.

There are still other interesting tractable clausal target languages that require more significant modifications to our compilation algorithms. Important ones include:

- k-Horn*. For any fixed k , Horn clauses containing k or fewer literals. Note that reasoning with k -Horn theories takes only linear time, and furthermore, the

⁸ This theorem is due to del Val [1995], and strengthens a proposition from Kautz and Selman [1991] that appeared in earlier drafts of this paper.

⁹ Subramanian and Genesereth [1987] present a formal system for inferring that certain propositions are irrelevant to the computation of a given class of queries. Given this sense of what is irrelevant, knowledge compilation can then be used as way to simplify the theory.

size of the k -Horn approximations is always polynomial in the number of variables. One way to compute a k -Horn GLB is to simply generate the set of all Horn clauses of length k as an initial (inconsistent) lower bound, and then minimizing the set by striking out clauses until it no longer entails the source theory. There is a similar brute-force algorithm to compute the k -Horn LUB: start with the set of all k -length clauses entailed by the source theory, and then repeatedly shrink it by removing clauses that are entailed by the conjunction of the remaining clauses.

—*Renamable Horn*. The class of theories that become Horn by uniformly replacing a subset of the variables by their negation. This class is of particular interest because (linear-time) unit propagation (BCP) is complete for it [Henschen and Wos 1974; Lewis 1978].

In addition to the clausal target languages we have concentrated on in this paper, one could consider tractable *nonclausal* target languages. For example, McAllester and Givan [1992] discuss a logical language based on the structure of natural language, and identify a fragment in which inference can be performed in polynomial time. It would be interesting to see if ordinary first-order theories could be (approximately) compiled into this tractable fragment, using our model-theoretic notion of approximation.

3.1.2. *First-Order Theories*. We will now briefly consider the application of knowledge compilation in first-order languages.

In the propositional case, GLBs and LUBs for any pair of a source and target language always exist, since only a finite number of logically distinct theories can be constructed. In the general first-order case, however, such best approximations may not be well defined. For a simple example, let \mathcal{L}_S be first-order logic, with a single predicate p and the infinite set of constant terms $\{c_1, c_2, \dots\}$, and \mathcal{L}_T consist of ground (nonquantified) clauses. The theory $\{\exists x.p(x)\}$ has an infinite series of better and better ground-clause lower-bounds:

$$\{p(c_1)\} \models \{p(c_1) \vee p(c_2)\} \models \{p(c_1) \vee p(c_2) \vee p(c_3)\} \models \dots \models \{\exists x.p(x)\}.$$

That is, there is no one *greatest* lower-bound.

The notion of a GLB is well defined, however, for the first-order generalizations of the *clausal* languages we have considered. (In fact, clausal representations are the most widely used first-order languages.) A *first-order literal* is constructed in the usual way from an optionally negated predicate symbol applied to a sequence of terms constructed from function symbols, constants, and variables. A *first-order clause* is a sentence in prenex form containing only universal quantifiers, whose matrix is a disjunction of literals.

A first-order Horn clause is a first-order clause whose matrix contains at most one positive literal; a Horn-strengthening of a first-order clause corresponds to a Horn-strengthening of its matrix. Note that even in the first-order case, a clause has only a *finite* number of Horn-strengthenings.

The reader may now verify that Lemmas 2.3.3 and 2.3.4 hold without change in the first-order case. Let us supply **Generate_GLB** in Figure 3 with an oracle for the test

if $L \models L'$ **then** $L := L'$

Then it clear that given a set of first-order clauses, the algorithm will search through a finite set of Horn-strengthenings and return one that is logically weakest; by the argument in the proof of Theorem 2.3.5, this must be a GLB. This shows that first-order Horn GLBs exist and are small.

There is also no difficulty in generalizing the notion of a GLB in other restricted clausal target languages to the first-order case. The θ -strengthening of a first-order clause is again generated by eliminating literals from the clause's matrix. Then, Theorem 3.1.1.2 goes through as before.

It is less straightforward to apply the notion of a LUB to first-order theories. We can construct finite first-order clausal theories that have no finite Horn LUB (even using no function symbols, only variables). Thus, not all first-order Horn LUBs are finitely representable. There are still other aspects of compiling first-order theories to be dealt with. In particular, we are investigating decidable (and possibly tractable) approximations to theories in undecidable languages.¹⁰

3.1.3. *Description Logics.* In this section, we consider description logics, a family of frame-based knowledge representation languages as studied by Levesque and Brachman [1985]. (See also Donini et al. [1991].)

Levesque and Brachman consider a language \mathcal{FL} in which one can describe structured concepts in terms of other concepts, either complex or primitive. For example, if we wished to describe the concept of a person whose male friends are all doctors with some specialty, we could use the \mathcal{FL} expression

```
(AND person
  (ALL (RESTR friend male)
    (AND doctor
      (SOME specialty))))),
```

which contains all the constructs (the capitalized terms) used in the language. Levesque and Brachman consider the complexity of determining whether one concept *subsumes* another. For example, the concept of “a person whose male friends are all doctors”, written

```
(AND person
  (ALL (RESTR friend male)
    doctor)),
```

subsumes the one given above. Now, their central technical result is that determining subsumption in \mathcal{FL} is intractable, but that removing the RESTR construct leads to polynomial time computable subsumption. The restricted language is called \mathcal{FL}^- .

¹⁰ Interesting connections may be drawn with some recent research on the analysis of Prolog programs (for use in, for example, optimization and program specification). For example, Heintze and Jaffar [1990] describes how to construct a *recursive* (i.e., decidable) approximation to a potentially non-recursive logic program. Their method computes a lower-bound of the logic program viewed as a logical theory (but not in general the greatest lower-bound).

So for efficient subsumption, one can use the language \mathcal{FL}^- . But this language may not be sufficiently expressive for practical applications. Knowledge compilation provides again an alternative. In this case, the idea is to take a concept description in the language \mathcal{FL} and to approximate it using two concept descriptions in \mathcal{FL}^- : a best lower-bound, that is, the most general more specific concept in \mathcal{FL}^- , and a best upper-bound, that is, the most specific more general (subsuming) concept in \mathcal{FL}^- .

As an example consider the first concept given above. It is not difficult to see that the concepts person and

```
(AND person
  (ALL friend
    (AND doctor
      (SOME specialty))))
```

in \mathcal{FL}^- are examples of, respectively, an upper-bound and a lower-bound in \mathcal{FL}^- . A knowledge representation system can store such bounds with the original concept description, and use them to try to determine quickly whether the newly given concept subsumes it or is subsumed by it.

Note that in this application we compute a pair of bounds for each concept, instead of a single pair of bounds for an entire knowledge base. The bounds can be used both in subsumption queries to the KB and in updating the KB by adding new concept definitions. In both cases, the system can use the bounds in determining the appropriate place of the new concept in the hierarchy.

4. Related Work

An upper-bound approximation is logically weaker than the original theory, and can thus be viewed as a *generalization* or *abstraction* of the theory. A lower-bound approximation is a *specialization* (logical strengthening) of the original theory. Borgida and Etherington [1989] propose using background knowledge that captures the hierarchical relationship among predicates in order to replace disjunctions by more general concepts. Suppose the background knowledge is

```
doctor(Jill)  $\supset$  professional(Jill)
lawyer(Jill)  $\supset$  professional(Jill)
```

and the KB is

```
doctor(Jill)  $\vee$  lawyer(Jill).
```

They then generate a new KB that contains only `professional(Jill)`. But this is simply the least Horn upper-bound of the original KB together with the background knowledge. Note that the idea of an LUB is more general than Borgida and Etherington's approach, since it can be applied to arbitrary propositional Horn theories—not just concept hierarchies with positive disjunctions. (Related to the work of Borgida and Etherington is that of Imielinski [1987], who

proposes a notion of abstraction based on grouping individuals into equivalence classes.)

A Horn lower-bound corresponds to a more specific theory than the original one. Its use generalizes the use of a counterexample to prune the search space of inference procedures. The best-known example of the use of counterexamples in artificial intelligence can be found in the early work by Gelernter [1959/1963] on proving theorems in geometry. Gelernter used a single model M (given by a diagram) of the original theory Σ to answer certain queries negatively, based on the fact that, if $M \models \alpha$, then $\Sigma \models \alpha$, for a query α . The Horn lower-bound is used in a similar manner, but it will generally involve a *set* of models, and is thus a better characterization of the original theory. In particular, one may avoid some of the “accidental truths” that often hold in a single model or diagram.

Horn lower-bounds are also a generalization of Levesque’s notion of vivid representations [Levesque 1986; Etherington et al. 1989]. To allow for fast inference, a vivid representation contains only complete information (no “reasoning by cases”). Levesque proposes to replace a knowledge base containing incomplete information by a complete, vivid representation of the information. Such a representation can be generated using abstractions or by filling in details. In the propositional case, the proposal comes down to replacing a propositional theory by one of its models (which is presumably “prototypical” of the models of the original theory). Queries are evaluated directly with respect to the model. Our Horn lower-bound can be viewed as a relaxation of the completeness constraint: Instead of picking a single model, we use a collection of models represented by a set of Horn clauses. We maintain the computational advantage because inference is still efficient using the Horn bound.

In work in artificial intelligence on planning and theorem-proving with abstraction [Amarel 1968; Plaisted 1981], one maps a theory to a smaller, simpler theory, generates proofs in the smaller theory, and then uses the proofs to guide generation of proofs in the original theory. Various mechanisms have been suggested for creating the abstractions, such as eliminating operator preconditions [Sacerdoti 1974], finding symmetries in the search space [Ellman 1993], and employing explanation-based learning on previously-generated plans [Knoblock et al. 1991a]. These abstractions are similar to our notion of an upper-bound. (Some of the approaches differ in that they can introduce abstract solution that do not correspond to any concrete solution. Knoblock et al. [1991b] discuss various kinds of abstractions that do preserve such solution and consistency properties.)

A crucial difference between the work on planning and our knowledge compilation framework is that the former concentrates on solving single problem instances, rather than shifting computational effort from on-line to off-line. All but the most restricted forms of planning are NP-hard [Gupta and Nau 1991; Bylander 1991; Erol et al. 1992; Selman 1994]. Therefore, the best known planning algorithms run in exponential time. Work on planning with abstraction attempts to lower the exponent. By contrast, the aim in the knowledge compilation framework is to preprocess (compile) a theory in such a manner that queries (problem instances) can be solved in polynomial time.

The motivation behind our knowledge compilation approach is similar to that behind the work on *universal plans* [Schoppers 1987], a form of reactive planning [Agre and Chapman 1987]. A universal plan is a program (or logical circuit) that

efficiently determines the first step of a plan leading from any initial state toward a predefined goal state. (After this step is taken, the universal plan determines what to do in the new state, and so on, until the goal is reached.) Universal plans can thus be viewed as an attempt to efficiently represent a potentially exponential set of plans, just as knowledge compilation aims at creating a representation that can efficiently answer a potentially exponential set of queries. In universal plans, one generally tries to capture exactly the original domain theory. It would be interesting to explore the question of whether Horn approximations of the domain knowledge could be used to create “approximate” universal plans (e.g., ones that may not specify for *every* state what to do next).

Cadoli and Schaerf [1992] generalize Levesque’s work on implicit and explicit belief [Levesque 1984]. They approximate the inference process by allowing a sequence of more and more powerful inference relations. The more time the system has to evaluate a query, the more powerful an inference relation it can use. Allowing for an unrestricted amount of time leads to standard deductive reasoning. An important difference with our approach is that there is no notion of directly processing the original theory to obtain an approximation to it. This means that in practice only a limited amount of work can be done per query, in order to obtain an answer. Moreover, none of the work done for one query is saved for use in answering the next query. In our framework, the cost of run-time querying is shifted to off-line compilation. The cost of the compilation process is amortized over a potentially very large set of queries, and the result of the compilation can be used in processing each query.

Dalal and Etherington [1992] provide a general framework in which many different forms of approximations can be captured. Our notion of an upper bound corresponds to their notion of a “direct weakening” of a theory. However, their notion of the strengthening of a theory is quite different from ours; it roughly corresponds to the disjunction of all of the lower-bounds of a theory. They introduce approximations to queries as well as theories; in our framework there is no need to approximate queries—as we saw earlier, it is possible to test arbitrary CNF queries against Horn theories in linear time.

Dalal and Etherington do not develop the notion of shifting computation effort from query-answering time to an off-line compilation phase, nor do they present algorithms for compiling specific languages. The emphasis in their work is more on developing a very general framework for describing mappings between different vocabularies, rather than dealing with specific issues of computational complexity.

Dechter and Pearl [1992] investigate the use of Horn approximations to obtain more efficient representations of empirical data. They consider the case where the original theory Σ is given directly by its set of models (satisfying truth assignments). They introduce a weaker notion of Horn approximations, called k -Horn approximations, in which each Horn clause contains at most k literals. See Kautz et al. [1994], for a discussion of this work.

Greiner and Schuurmans [1992] have adapted our compilation algorithms to ones that search for bounds that are optimal with respect to a given query distribution. Generation of the bounds is interleaved with query answering, and the queries themselves are used to optimize the bounds. They show that this can be done efficiently, in that each iteration takes polynomial time.

Finally, Gogic et al. [1994] show how Horn approximations can be used to obtain a tractable form of theory revision. Also, a generalized form of the LUB is used in the “learning to reason” approach of Khardon and Roth [1994].

5. Conclusions

We introduced the notion of knowledge compilation. The basic idea is to compile knowledge from an intractable into a tractable form. Since an exact translation is often not possible, we introduced approximate translations, consisting of two bounds that delimit the original theory.

Knowledge compilation provides an alternative to approaches that force the user to state all knowledge in some restricted (tractable) language. A representation system incorporating a knowledge compilation procedure will allow the user to enter information in a general, unrestricted language, which then the system compiles into a tractable form.

To illustrate our approach, we showed how knowledge represented in a propositional theory can be approximated using two Horn theories, called Horn approximations: a greatest Horn lower-bound and a least Horn upper-bound. Answering a query based on the original knowledge base is intractable, but by using the Horn approximations certain queries can be answered in time linear in the length of the approximations. We gave algorithms for generating such Horn approximations. The algorithms operate incrementally, generating better and better approximations over time. The incremental nature of the approximation algorithms is a key feature of our approach, since in practical applications it would be unacceptable to have to wait until the system has computed the best bounds before answering any queries.

In summary, the main features of our knowledge compilation approach are:

- A guaranteed fast response for queries that can be answered directly using the approximations.
- An incremental, off-line compilation process that provides continuous improvement of the overall response time of the system.

We also presented an empirical evaluation of our approach, and demonstrated substantial computational savings on hard propositional theories.

We showed how the procedures for compiling propositional theories into Horn theories can be generalized to apply to other tractable classes of clauses. Those classes were characterized using various closure conditions. The classes containing reverse-Horn clauses, clauses with two or fewer literals, or clauses not containing a certain set of “irrelevant letters” are examples of classes that satisfy the closure conditions. Finally, we discussed the compilation of concept descriptions given in a terminological representation language. This example showed that our knowledge compilation approach appears suited for dealing with a variety of knowledge representation languages—not only traditional logics.

Appendix A. Proofs

PROOF OF THEOREM 2.5.1.1. Consider a theory Σ of the following form, for arbitrary n :

$$\neg p_1 \vee \neg p_2 \vee \dots \vee \neg p_n \vee s \quad (16)$$

$$\neg q_1 \vee p_1 \vee s \quad (17)$$

$$\neg q_2 \vee p_2 \vee s \quad (18)$$

$$\vdots \quad (19)$$

$$\neg q_n \vee p_n \vee s \quad (20)$$

We see that Σ contains $4n + 1$ literals; that is, it is of size $O(n)$. The LUB Σ_{lub} of Σ is the following set of Horn clauses, of total length 2^n :

$$\left\{ \begin{array}{l} \neg x_1 \vee \neg x_2 \vee \dots \vee \neg x_n \vee s \\ x_i \in \{p_i, q_i\} \text{ for } 1 \leq i \leq n \end{array} \right\}.$$

First, we prove that the set Σ_{lub} is irredundant (no subset of Σ_{lub} implies all of Σ_{lub}). Then, we prove that Σ_{lub} is of minimum size. (Note that being of minimum size is a stronger condition than being irredundant.)

Proof that Σ_{lub} is irredundant: Suppose there is a clause α in Σ_{lub} such that $\Sigma_{\text{lub}} - \{\alpha\} \models \alpha$. Since no two clauses in $\Sigma_{\text{lub}} - \{\alpha\}$ resolve, by completeness of resolution there must be an α' in $\Sigma_{\text{lub}} - \{\alpha\}$ such that α' subsumes α . But this is impossible, since all clauses in Σ_{lub} are of the same length.

We can now prove that there is no smaller set of clauses Σ'_{lub} which is equivalent to Σ_{lub} : Suppose there were a Σ'_{lub} such that $\Sigma_{\text{lub}} \equiv \Sigma'_{\text{lub}}$ and $|\Sigma'_{\text{lub}}| < |\Sigma_{\text{lub}}|$. Then for all α in Σ'_{lub} ,

$$\Sigma_{\text{lub}} \models \alpha$$

and since no clauses in Σ_{lub} resolve, this means that there exists an α' in Σ_{lub} such that α' subsumes α .

That is, every clause in Σ'_{lub} is subsumed by some clause in Σ_{lub} . Suppose each clause in Σ_{lub} subsumed a different clause in Σ'_{lub} ; then $|\Sigma'_{\text{lub}}| \geq |\Sigma_{\text{lub}}|$, a contradiction. Therefore, there is a proper subset Σ''_{lub} of Σ_{lub} such that each clause in Σ'_{lub} is subsumed by some clause in Σ''_{lub} .

Then $\Sigma''_{\text{lub}} \models \Sigma'_{\text{lub}}$, and therefore $\Sigma''_{\text{lub}} \models \Sigma_{\text{lub}}$. But this is impossible, because we saw that Σ_{lub} is irredundant. Therefore, there is no smaller set of clauses equivalent to Σ_{lub} , which is shorter than Σ_{lub} . \square

PROOF OF THEOREM 2.5.3.1. Suppose such a representation of the LUB always existed. We then show that 3-SAT over n variables can be determined in $O(n^3)$ time. Because the choice of n is arbitrary, and 3-SAT is an NP-complete problem, this would mean that $\text{NP} \subseteq \text{nonuniform P}$.

Let the variables be a set of main variables $\{p_1, \dots, p_n\}$ together with a set of auxiliary variables

$$\{c_{x,y,z} \mid x, y, z \in \text{LITS}\},$$

where LITS is the set of literals constructed from the main variables (the variables or their negations). Let the source theory Σ be the conjunction of clauses:

$$\Sigma = \bigwedge_{x,y,z \in \text{LITS}} \{x \vee y \vee z \vee \neg c_{x,y,z}\}.$$

Note that Σ is of length $O(n^3)$. The idea behind the construction is that any 3-SAT clause over n variables can be constructed by selecting a subset of the clauses from Σ and eliminating the auxiliary variables.

Now suppose we have been able to compute a representation of the Horn LUB of Σ that has the property that one can test Horn queries against it in polynomial time. We noted before that a Horn formula follows from the LUB if and only if it follows from the source theory.

Suppose Δ is an arbitrary 3-CNF formula over n variables that we wish to test for satisfiability. We construct a Horn clause α containing only auxiliary variables by including a negative auxiliary variable c that corresponds to each clause δ in Δ . For example, if Δ is

$$(p_1 \vee \neg p_3 \vee p_5) \wedge (\neg p_1 \vee p_2 \vee \neg p_4),$$

then the corresponding Horn clause is

$$\neg c_{p_1, \neg p_3, p_5} \vee \neg c_{\neg p_1, p_2, \neg p_4}.$$

Now we claim that this Horn clause is implied by the LUB if and only if Δ is not satisfiable.

(\rightarrow) Suppose the query is implied by the LUB. It follows that

$$\Sigma \models \neg c \vee \neg c' \vee \neg c'' \vee \dots,$$

where the $\{c, c', c'', \dots\}$ are the auxiliary variables in the query that correspond to clauses $\{\delta, \delta', \delta'', \dots\}$ in Δ . Equivalently,

$$\Sigma \cup \{\neg(\neg c \vee \neg c' \vee \neg c'' \vee \dots)\} \text{ is unsatisfiable.}$$

That is,

$$\Sigma \cup \{c, c', c'', \dots\} \text{ is unsatisfiable.}$$

Note that any clause in Σ containing an auxiliary variable other than $\{c, c', c'', \dots\}$ can be eliminated, since that clause is satisfied in any model in which its auxiliary variable is assigned false, and no other instance of that variable appears in the formula. Thus, it must be the case that

$$\{\delta \vee \neg c, \delta' \vee \neg c', \delta'' \vee \neg c'', \dots\} \cup \{c, c', c'', \dots\}$$

is unsatisfiable. Because the auxiliary variables each appear exactly once negatively and once positively above, they can be resolved away. Therefore

$$\{\delta, \delta', \delta'', \dots\} \equiv \Delta \text{ is unsatisfiable.}$$

(\leftarrow) Note that each step in the previous section can be reversed, to go from the assumption that Δ is unsatisfiable to the conclusion that $\Sigma \models \alpha$. And, since α is Horn, it follows that the LUB implies α .

We assumed that the test whether the LUB implies α could be performed in time polynomial in the length of the source theory plus the length of the query.

Since both the source theory and query are polynomial in the length of Δ , it follows that the satisfiability of Δ can be determined in time polynomial in the length of Δ . Since the choice of n was arbitrary, and 3-SAT is an NP-complete problem, this means that $\text{NP} \subseteq \text{nonuniform P}$. \square

PROOF OF STRENGTHENED THEOREM 2.5.3.1. We can strengthen the theorem to an equivalence by showing that $\text{NP} \subseteq \text{nonuniform P}$ implies that small and tractable representations of the LUB always exist. Suppose we are given a source theory Σ of length m containing n variables. Assuming $\text{NP} \subseteq \text{nonuniform P}$, there exists a circuit that determines satisfiability of formulas of length $m + n$ that has complexity polynomial in $m + n$. We use this circuit to construct program to test queries of the form $\Sigma_{\text{lub}} \models \alpha$ as follows: given α , first check that is not a tautology, and eliminate any duplicated literals. The resulting query is of length $\leq n$. Then pad out the query to exactly length n by duplicating any of its literals. Then, the negation of the query together with Σ is a formula of exactly length $m + n$, so we can use the circuit to determine if the formula is unsatisfiable, or equivalently, that α follows from Σ . Since α is Horn, then the latter condition is equivalent to saying $\Sigma_{\text{lub}} \models \alpha$. Since the circuit is of size polynomial in $m + n$ it must execute in time polynomial in $m + n$ —that is, in time polynomial in $(|\Sigma| + |\alpha|)$. \square

ACKNOWLEDGMENTS. We thank Daniel Bobrow for getting us to think more about the issue of how to make practical use of restricted, tractable representation languages and Ray Reiter for pointing us to the work on prime implicates. We also thank Hector Levesque, Yoav Shoham, and two anonymous referees for many useful comments.

REFERENCES

- AGRE, P. E., AND CHAPMAN, D. 1987. Pengi: An implementation of a theory of activity. In *Proceedings of the 6th National Conference on Artificial Intelligence (AIII-87)* (Seattle, Wash.). Morgan-Kaufmann, San Mateo, Calif., pp. 268–272.
- AMAREL, S. 1968. On representations of problems of reasoning about actions. In *Machine Intelligence*, vol. 3, D. Michie, ed. Edinburgh University Press, Edinburgh, Scotland, pp. 131–171.
- ASPVALL, B., PLASS, M. F., AND TARJAN, R. E. 1979. A linear-time algorithm for testing the truth of certain quantified Boolean formulae. *Inf. Proc. Lett.* 8, 121–127.
- BOPANA, R. B., AND SIPSER, M. 1990. The complexity of finite functions. In *Handbook of Theoretical Computer Science, Volume A: Algorithms and Complexity*, J. van Leeuwen, ed. Elsevier, Amsterdam, The Netherlands (and MIT Press, Cambridge, Mass.), pp. 757–804.
- BORGIDA, A., AND ETHERINGTON, D. W. 1989. Hierarchical knowledge bases and efficient disjunctive reasoning. In *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*. (Toronto, Ont., Canada). Morgan-Kaufmann, San Mateo, Calif., pp. 33–43.
- BURO, M., AND BÜNING, H. K. 1992. Report on a SAT competition. Technical Memorandum 110. Mathematik/Informatik Univ. Paderborn, Paderborn, Germany.
- BYLANDER, T. 1991. Complexity results for planning. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence (IJCAI-91)* (Sidney, Australia). Morgan-Kaufmann, San Mateo, Calif., pp. 274–279.
- CADOLI, M. 1993. Semantical and computational aspects of horn approximations. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI-93)* (Chambery, France). Morgan-Kaufmann, San Mateo, Calif., pp. 39–44.
- CADOLI, M., AND SCHAERF, M. 1992. Approximation in concept description languages. In *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR-92)* (Cambridge, Mass.). Morgan-Kaufmann, San Mateo, Calif., pp. 330–341.

- COOK, S. A. 1971. The complexity of theorem-proving procedures. In *Proceedings of the 3rd Annual ACM Symposium on the Theory of Computing* (Shaker Heights, Ohio, May 3–5). ACM, New York, pp. 151–158.
- CRAWFORD, J., AND AUTON, L. 1993. Experimental results on the crossover point in satisfiability problems. In *Proceedings of the 11th National Conference of Artificial Intelligence (AIII-93)* (Washington, D.C.). MIT Press, Cambridge, Mass., pp. 21–27.
- DALAL, M., AND ETHERINGTON, D. W. 1992. Tractable approximate deduction using limited vocabularies. In *Proceedings of the 9th Canadian Conference on Artificial Intelligence (AI '92)* (Vancouver, B.C., Canada). Morgan-Kaufmann, San Mateo, Calif., pp. 206–212.
- DAVIS, M., AND PUTNAM, H. 1960. A computing procedure for quantification theory. *J. ACM* 7, 3 (July), 201–215.
- DE KLEER, J. 1990. Exploiting locality in the TMS. In *Proceedings of the 8th National Conference on Artificial Intelligence (AAAI-91)* (Boston, Mass.). MIT Press, Cambridge, Mass., pp. 264–271.
- DEAN, T. 1985. *Artificial Intelligence, Theory and Practice*. Benjamin/Cummings, Redwood City, Calif.
- DEAN, T., AND BODDY, M. 1988. An analysis of time-independent planning. In *Proceedings of the 7th National Conference on Artificial Intelligence (AAAI-88)* (St. Paul, Minn.). Morgan-Kaufmann, San Mateo, Calif., pp. 49–54.
- DECHTER, R. 1992. Constraint networks. In *Encyclopedia of Artificial Intelligence*. John Wiley & Sons, New York, pp. 276–285.
- DECHTER, R., AND PEARL, J. 1992. Structure identification in relational data. *Artif. Int.* 58, 1–3, 237–270.
- DEL VAL, A. 1995. An analysis of approximate knowledge compilation. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence (IJCAI-95)* (Montreal, Que., Canada). Morgan-Kaufmann, San Mateo, Calif., pp. 830–836.
- DONINI, F. M., LENZERINI, M., NARDI, D., AND NUTT, W. 1991. The complexity of concept languages. In *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR-91)* (Cambridge, Mass.). Morgan-Kaufmann, San Mateo, Calif., pp. 151–162.
- DOWLING, W. F., AND GALLIER, J. H. 1984. Linear time algorithms for testing the satisfiability of propositional Horn formula. *J. Log. Prog.* 3, 267–284.
- DOYLE, J., AND PATIL, R. 1991. Two theses of knowledge representation: Language restrictions, taxonomic classification, and the utility of representation services. *Artif. Int.* 48, 3, 261–298.
- DUBOIS, O., ANDRE, P., BOUFGHAD, Y., AND CARLIER, J. 1995. SAT versus unSAT. In *Cliques, Coloring, and Satisfiability: Second DIMACS Implementation Challenge*, D. S. Johnson and M. A. Trick, eds. DIMACS Series in Discrete Mathematics and Theoretical Computer Science. AMS Press, Providence, R.I.
- ELLMAN, T. 1993. Abstraction via approximate symmetry. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI-93)*, vol. 2. (Chambery, France). Morgan-Kaufmann, Mateo, Calif., pp. 916–921.
- EROL, K., NAU, D. S., AND SUBRAHMANIAN, V. S. 1992. On the complexity of domain-independent planning. In *Proceedings of the 10th National Conference on Artificial Intelligence (AAAI-92)* (San Jose, Calif.). MIT Press, Cambridge, Mass., pp. 381–386.
- ETHERINGTON, D. W., BORGIDA, A., BRACHMAN, R. J., AND KAUTZ, H. 1989. Vivid knowledge and tractable reasoning: Preliminary report. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI-89)* (Detroit, Mich.). Morgan-Kaufmann, San Mateo, Calif., pp. 1146–1152.
- FRISCH, A. M. 1985. Using model theory to specify AI programs. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence (IJCAI-85)* (Los Angeles, Calif.). Morgan-Kaufmann, San Mateo, Calif., pp. 148–154.
- GAREY, M. R., AND JOHNSON, D. S. 1979. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, San Francisco, Calif.
- GELERNTER, H. 1959/1963. Realization of a geometry theorem-proving machine. In *Proceedings of the International Conference on Information Processing* (Paris, France). UNESCO House, 1959. pp. 273–282. (Reprinted in *Computers and Thought*. E. Feigenbaum and J. Feldman, eds. McGraw-Hill, New York, 1963, 134–152.)

- GOGIC, G., PAPADIMITRIOU, C., AND SIDERI, M. 1994. Incremental recompilation of knowledge. In *Proceedings of the 12th National Conference on Artificial Intelligence (AAAI-94)* (Seattle, Wash.). MIT Press, Boston, Mass., pp. 922–927.
- GREINER, R., AND SCHUURMANS, D. 1992. Learning useful horn approximations. In *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning (KR-92)* (Boston, Mass.). Morgan-Kaufmann, San Mateo, Calif., pp. 383–392.
- GUPTA, N., AND NAU, D. S. 1991. Complexity results for blocks-world planning. In *Proceedings of the 9th National Conference on Artificial Intelligence (AAAI-91)* (Anaheim, Calif.). MIT Press, Cambridge, Mass., pp. 629–635.
- HEINTZE, N., AND JAFFAR, J. 1990. A finite presentation theorem for approximating logic programs. In *Proceedings of POPL-90*. MIT Press, Cambridge, Mass., pp. 197–201.
- HENSCHEN L., AND WOS, L. 1974. Unit refutations and Horn sets. *J. ACM* 21, 4 (Oct.), 590–605.
- HORVITZ, E. J., COOPER, G. F., AND HECKERMAN, D. E. 1989. Reflection and action under scarce resources: Theoretical principles and empirical study. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence (IJCAI-89)* (Detroit, Mich., May). Morgan-Kaufmann, San Mateo, Calif., pp. 1121–1126.
- IMIELINSKI, T. 1987. Domain abstraction and limited reasoning. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence (IJCAI-87)*, vol. 2. Morgan-Kaufmann, San Mateo, Calif., pp. 997–1002.
- JOHNSON, D. S. 1990. A catalog of complexity classes. In *Handbook of Theoretical Computer Science*, vol. A: Algorithms and Complexity, chap. 2. J. Van Leeuwen, ed. Elsevier Science Publishers B. V., Amsterdam, The Netherlands (and MIT Press, Cambridge, Mass.), pp. 67–161.
- KARP, R. M., AND LIPTON, R. 1982. Turing machines that take advice. *Enseign. Math.*, 28, 191–209.
- KAUTZ, H. A., KEARNS, M. J., AND SELMAN, B. 1994. Horn approximations of empirical data. *Artif. Int.* 74, 129–145.
- KAUTZ, H., AND SELMAN, B. 1991. A general framework for knowledge compilation. In *Proceedings of the International Workshop on Processing Declarative Knowledge (PDK-91)* (Kaiserslautern, Germany, July). Lecture Notes in Artificial Intelligence vol. 567. H. Boley and M. M. Richter, eds. Springer-Verlag, Berlin, Germany, pp. 287–300.
- KAUTZ, H., AND SELMAN, B. 1992a. Forming concepts for fast inference. In *Proceedings of the 10th National Conference on Artificial Intelligence (AAAI-92)* (San Jose, Calif.). MIT Press, Cambridge, Mass., pp. 786–793.
- KAUTZ, H., AND SELMAN, B. 1992b. Planning as satisfiability. In *Proceedings of the 10th European Conference on Artificial Intelligence (ECAI-92)* (Vienna, Austria). Wiley, Chichester, UK, pp. 359–364.
- KAUTZ, H., AND SELMAN, B. 1994. An empirical evaluation of knowledge compilation. In *Proceedings of the 12th National Conference on Artificial Intelligence (AAAI-94)* (Seattle, Wash.). MIT Press, Boston, Mass., pp. 155–161.
- KHARDON, R., AND ROTH, D. 1994. Learning to reason. In *Proceedings of the 12th National Conference on Artificial Intelligence (AAAI-94)* (Seattle, Wash.). MIT Press, Boston, Mass. pp. 682–688.
- KIRKPATRICK, S. AND SELMAN, B. 1994. Critical behavior in the satisfiability of random Boolean expressions. *Science*, 264 (May), 1297–1301.
- KNOBLOCK, C. A., MINTON, S., AND ETZIONI, O. 1991. Integrating abstraction and explanation-based learning in prodigy. In *Proceedings of the 9th National Conference on Artificial Intelligence (AAAI-91)* (Anaheim, Calif.). MIT Press, Cambridge, Mass., pp. 541–546.
- KNOBLOCK, C. A., TENENBERG, J. D., AND YANG, Q. 1991. Characterizing abstraction hierarchies for planning. In *Proceedings of the 9th National Conference on Artificial Intelligence (AAAI-91)* (Anaheim, Calif.). MIT Press, Cambridge, Mass., pp. 692–697.
- LARRABEE, T. 1992. Test pattern generation using Boolean satisfiability. *IEEE Trans. Comput. Aid. Des.* 11, 4–15.
- LEE, R. C. T. 1967. A completeness theorem and a computer program for finding theorems derivable from given axioms. Ph.D. dissertation. Univ. California at Berkeley, Berkeley, Calif.
- LEVESQUE, H. J. 1984. A logic of implicit and explicit belief. In *Proceedings of the 4th National Conference on Artificial Intelligence (AAAI-84)*. (Austin, Tex.). Morgan-Kaufmann, San Mateo, Calif., pp. 198–202.
- LEVESQUE, H. 1986. Making believers out of computers. *Artif. Int.* 30, 1, 81–108.
- LEVESQUE, H. J., AND BRACHMAN, R. J. 1985. A fundamental tradeoff in knowledge representation and reasoning (revised version). In *Readings in Knowledge Representation*, R. J. Brachman and H. J. Levesque, eds. Morgan-Kaufmann, San Mateo, Calif., pp. 41–70.

- LEWIS, H. R. 1978. Renaming a set of clauses as a Horn set. *J. ACM*, 25, 1 (Jan.), 134–135.
- MCALLESTER, D. A., AND GIVAN, R. 1992. Natural language syntax and first-order inference. *Artif. Int.* 56, 1–20.
- MITCHELL, D., SELMAN, B., AND LEVESQUE, H. 1992. Hard and easy distribution of SAT problems. In *Proceedings of the 10th National Conference on Artificial Intelligence (AAAI-92)* (San Jose, Calif.). MIT Press, Cambridge, Mass., pp. 459–465.
- MUGGLETON, S., AND BUNTINE, W. 1988. Machine invention of first-order predicates by inverting resolution. In *Proceedings of the 5th International Conference on Machine Learning*. Morgan-Kaufmann, San Mateo, Calif., pp. 339–344.
- PATEL-SCHNEIDER, P. F. 1986. A four-valued semantics for frame-based description languages. In *Proceedings of the 5th National Conference on Artificial Intelligence (AAAI-86)* (Philadelphia, Pa.). Morgan-Kaufmann, San Mateo, Calif., pp. 344–348.
- PLAISTED, D. 1981. Theorem proving with abstraction. *Artif. Int.* 16, 47–65.
- REITER, R., AND DE KLEER, J. 1987. Foundations of assumption based truth maintenance systems: Preliminary report. In *Proceedings of the 6th National Conference on Artificial Intelligence (AAAI-87)* (Austin, Tex.). Morgan-Kaufmann, San Mateo, Calif., pp. 183–187.
- REITER, R., AND MACKWORTH, A. 1989. A logical framework for depiction and image interpretation. *Artif. Int.* 41, 2, 125–155.
- SACERDOTI, E. D. 1974. Planning in a hierarchy of abstraction spaces. *Artif. Int.* 5, 2, 115–135.
- SCHOPPERS, M. J. 1987. Universal plans for reactive robots in unpredictable environments. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence (IJCAI-87)* (Milan, Italy). Morgan-Kaufmann, San Mateo, Calif., pp. 1039–1046.
- SELMAN, B. 1991. Tractable default reasoning. Ph.D. thesis. Dept. Computer Science, Univ. Toronto, Toronto, Ontario, Canada.
- SELMAN, B. 1994. Near-Optimal Plans, Tractability, and Reactivity. In *Proceedings of the 4th International Conference on Knowledge Representation and Reasoning (KR-94)* (Bonn, Germany). Morgan-Kaufmann, San Mateo, Calif., pp. 521–529.
- SELMAN, B., AND KAUTZ, H. 1991. Knowledge compilation using Horn approximations. In *Proceedings of the 9th National Conference on Artificial Intelligence (AAAI-91)* (Anaheim, Calif.). MIT Press, Cambridge, Mass., pp. 904–909.
- SELMAN, B., AND KAUTZ, H. 1993. Domain-independent extensions to GSAT: Solving large structured satisfiability problems. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence (IJCAI-93)* (Chambéry, France). Morgan-Kaufmann, San Mateo, Calif., pp. 290–295.
- SELMAN, B., AND KIRKPATRICK, S. 1996. Critical behavior in the computational cost of satisfiability testing. *Artif. Int.*, 81, in press.
- SELMAN, B., AND LEVESQUE, H. J. 1990. Abductive and default reasoning: a computational Core. In *Proceedings of the 8th National Conference on Artificial Intelligence (AAAI-90)* (Cambridge, Mass.). MIT Press, Cambridge, Mass., pp. 343–348.
- SELMAN, B., AND LEVESQUE, H. J. 1996. Support set selection for abductive and default reasoning. *Artif. Int.*, in press.
- SELMAN, B., LEVESQUE, H. J., AND MITCHELL, D. 1992. A new method for solving hard satisfiability problems. In *Proceedings of the 10th National Conference on Artificial Intelligence (AAAI-92)* (San Jose, Calif.). MIT Press, Cambridge, Mass. pp. 440–446.
- SELMAN, B., MITCHELL, D., AND LEVESQUE, H. J. 1996. Generating Hard Satisfiability Problems. *Artif. Int.*, 81, in press.
- SUBRAMANIAN, D., AND GENESERETH, M. R. 1987. The relevance of irrelevance. In *Proceedings of the 10th International Joint Conference on Artificial Intelligence (IJCAI-87)* (Milan, Italy). Morgan-Kaufmann, San Mateo, Calif., pp. 416–422.
- TSEITIN, G. S. 1968. On the complexity of derivation in propositional calculus. *Studies in Constructive Mathematics and Mathematical Logic, Part II*. A. O. Slisenko, ed. pp. 115–125.
- VALIANT, L. G. 1983. A theory of the learnable. *Commun. ACM* 27, 11 (Nov.), 1134–1142.
- VALIANT, L. G., AND VAZIRANI, V. V. 1986. NP is as easy as detecting unique solutions. *Theoret. Comput. Sci.* 47, 85–93.

RECEIVED FEBRUARY 1993; REVISED OCTOBER 1994; ACCEPTED AUGUST 1995