INCREMENTAL RECOMPILATION OF KNOWLEDGE
(Extended Abstract)

Goran Gogic[1], Christos H. Papadimitriou[1], and Martha Sideri[2]

ABSTRACT: *Approximating a general formula from above and below by Horn formulas (its* Horn envelope *and* Horn core, *respectively) was proposed in [SK] as a form of "knowledge compilation," supporting rapid approximate reasoning; on the negative side, this scheme is static in that it supports no updates, and has certain complexity drawbacks pointed out in [KPS]. On the other hand, the many frameworks and schemes proposed in the literature for theory update and revision are plagued by serious complexity-theoretic impediments, even in the Horn case, as was pointed out in [EG2] and the present paper. More fundamentally, these schemes are not* inductive, *in that they lose in a single update any positive properties of the represented sets of formulas (small size, Horn, etc.).* In this paper we propose a new scheme, incremental recompilation, combining Horn approximation and model-based updates; this scheme is inductive and very efficient, *free of the problems facing its constituents. A set of formulas is represented by an upper and lower Horn approximation. To update, we replace the upper Horn formula by the Horn envelope of its minimum-change update, and similarly the lower one by the Horn core of its update; the key fact is that Horn envelopes and cores are easy to compute when the underlying formula is the result of a minimum-change update of a Horn formula by a clause. We conjecture that efficient algorithms are possible for more complex updates.*

1. INTRODUCTION

Starting with the ideas of Levesque, in recent years there has been increasing interest in computational models for *rapid approximate reasoning*, starting from a "vivid" representation of knowledge [Le]. One important proposal in this regard has been the *knowledge compilation* idea of [SK], whereby a propositional formula is represented by its optimal upper (relaxed) and lower (strict) approximations by Horn formulas —the corresponding Horn formulas are called in the present paper the *Horn envelope* and the *Horn core* of the original formula. The key idea of course is that, since these approximate theories are Horn, one can use them for rapid (linear-time) approximate reasoning.

Despite the computational advantages and attractiveness of this idea, some obstacles to its implementation have been pointed out. First, although the Horn envelope of a formula is unique, the Horn core is not; that is, there may be exponentially many most relaxed Horn formulas implying the given one. As was proved in [KPS], selecting the one with the largest set of models, or one that is approximately optimal in this respect (within *any* bounded ratio), is NP-complete. Another disadvantage is that the Horn envelope may have to be exponentially larger, as a Boolean formula, than the given formula. What is more alarming is that, even if the Horn envelope is small, it may take exponential time to produce. Even if we are given the set of models of the original formula, there is no known *output-polynomial* algorithm for producing all clauses of the Horn envelope. An algorithm is output-polynomial if it runs in time that is polynomial *in both the size of its input and its output;* this novel and little-studied concept of tractability (and, unfortunately, related concepts of *intractability*), have proved very relevant to various aspects of AI. In fact, it was shown in [KPS] that generating the Horn envelope from the models of a formula is what we call in the present paper *TRANSVERSAL-hard* —suggesting that it is unlikely to have an output-polynomial algorithm. These negative complexity results for knowledge compilation (admittedly, quite mild when compared with the serious obstacles to other approaches to knowledge representation and common-sense reasoning, see for example [EG1, EG2]) are summarized in Theorem 1.

Our knowledge about the world changes dynamically —and the world itself changes as well. The knowledge compilation idea has no provisions for incorporating such belief revisions or updates. There are, of course, in the literature many formalisms for updating knowledge bases and databases with in-

[1] University of California San Diego, La Jolla, California 92093-0114. e-mail: goran/christos@cs.ucsd.edu. Research supported by an NSF grant.

[2] Athens Univ. of Economics and Business, Athens, Greece. e-mail: sideri@aueb.ariadne-t.gr.

complete information [Da, Sa, Bo, We, Gi, FUV, Wil, Fo]; see [Wi2] and [EG2] for two systematic surveys. As was established in [EG2], all these systems are plagued with tremendous complexity obstacles —even making the next inference, which is known as the *counterfactual problem*, is complete at some high level of the polynomial hierarchy for all of them. We point out in this paper (Theorem 2) some serious problems associated with computing the updated formula in the two formula-based frameworks *even in the Horn case*. The only ray of hope from [EG2] —namely that when the formula is Horn, the update is small, and the approach is any one of the model-based ones, then counterfactuals are easy— is tarnished by our observation that, in all these cases, the updated formula is not Horn (this is part (iii) of Theorem 2); hence, such an update scheme would fail to be *inductive*, retaining its positive computational properties in the face of an update.

To summarize, knowledge compilation of arbitrary formulas is not easy to do. And all known approaches to the update problem encounter serious complexity obstacles, or result in loss of the Horn property. What hope is there then for a system that supports both rapid approximate reasoning *and* updates?

Quite surprisingly, combining these two ideas, both shackled as they are by complexity-theoretic obstacles, seems to remove the obstacles from both, thus solving the combined problem, at least in some interesting and heretofore intractable cases. In particular we propose the following scheme: Suppose that formula Γ is represented by its Horn envelope $\overline{\Gamma}$ and its Horn core $\underline{\Gamma}$ (to start the process, we incur a one-time computational cost for computing these bounds; alternatively, we may insist that we start with a Horn formula). Suppose now that we update our formula by ϕ, a "simple enough" formula (how "simple" it has to be for our scheme to be efficient is an important issue which we have only partially explored; we know how to handle a single Horn clause, as well as several other special cases). *We represent the updated formula by the two formulas* $\overline{\overline{\Gamma} + \phi}$ *and* $\underline{\underline{\Gamma} + \phi}$, where '+' stands for an appropriate model-based update formalism; that is, by the Horn envelope of the updated upper bound and the Horn core of the updated lower bound. These are our new $\overline{\Gamma}$ and $\underline{\Gamma}$. In other words, we apply the update to the two approximations, and approximate the two results, each in the safe direction. And so on, starting from the new approximations. The key technical point is that, although updating Horn formulas, even by Horn clauses, does not preserve the Horn property, and finding Horn envelopes and cores is hard in general, *it is easy when the formula to be approximated is the result of the update of a Horn formula by a Horn clause.* To our knowledge, our proposal, with all its restrictions, is the first computationally feasible approach to belief revision and updates.

Our proposal exhibits a desirable and intuitively expected "minimum-change" behavior, best demonstrated in the case in which a Horn formula Γ is updated by a Horn clause, say $\phi = (x \& y \rightarrow z)$. Suppose that Γ can be written as $x \& y \& \neg z \& \Gamma'$, where Γ' does not involve x, y, or z —otherwise $\Gamma + \phi = \Gamma \& \phi$. Then the upper and lower approximations are these: $\overline{\Gamma + \phi}$ is $(x \& y \leftrightarrow z) \& \Gamma'$, while $\underline{\Gamma + \phi}$ is $x \& (y \leftrightarrow z) \& \Gamma'$ (or $y \& (x \leftrightarrow z) \& \Gamma'$, recall that cores are not unique). Notice the "circumscriptive" nature of the updates (resulting from the minimum-change update formalisms that we are using).

There is an interesting and satisfying *methodological* aspect of our work. The main motivation and justification for applying the concepts and techniques of Complexity Theory to any application area is that, this way, research is supposed to be redirected by negative complexity results to the study of the right problems, to the adoption of the right approaches. For AI, there is an added argument why such results are relevant: In attacking computationally a problem in an application area, one cannot in principle exclude the possibility that this problem might be *totally insusceptible* to computational solution, that there may be *no right approach* to be discovered by a sequence of trials and negative complexity results. In contrast, a tacit ideological assumption in AI research is that the right approach *must* exist *—because intelligence exists.* Although the recent literature is teeming with complexity-theoretic criticism of approaches to various aspects of AI, the present work is an unusually clear example of a new approach that was arrived at by a tight complexity-theoretic argument excluding almost everything else.

2. NEGATIVE RESULTS

Let Γ be a *propositional* formula. Define [KS] its *Horn envelope* $\overline{\Gamma}$ to be the strictest Horn formula implied by Γ, and its *Horn core* to be the weakest Horn formula implying Γ. Naturally, one could not hope that the Horn envelope and core can be efficiently computed for all Boolean formulas. The reason is simple: Γ is unsatisfiable iff both $\overline{\Gamma}$ and $\underline{\Gamma}$ coincide with the **false** formula. But what if Γ is given in some more convenient form, say in terms of its set of models $\mu(\Gamma)$ (that is, in "full disjunctive form")? A first problem is that $\overline{\Gamma}$ may have exponentially many clauses —there is little that can be done in this case, we need them all to best approximate our formula. But can we hope to output these clauses, however many they may be, in time polynomial both in the size of input —$\mu(\Gamma)$— and of the output —$\overline{\Gamma}$? There are systematic ways that output all clauses of $\overline{\Gamma}$, but unfortunately in all known algorithms there may be exponential delay between the production of two consecutive clauses. There is no known *output-polynomial* algorithm for this problem.

There are many instances of such enumeration problems in the literature, for which no output-polynomial algorithm is known (despite the fact that, in contrast to NP-complete problems, it is trivial to output the first solution). The most famous one is to compute *all transversals of a hypergraph* [EG3]. As was pointed out in [EG3], many enumeration problems arising in AI, databases, distributed computation, and other areas of Computer Science, turn out to be what we call in this paper *TRANSVERSAL-hard*, in the sense that, if they are solvable in output polynomial time, then the transversal problem is likewise solvable.

Theorem 1 [KPS]: Enumerating all clauses of the Horn envelope of a given set M of models is TRANSVERSAL-hard. As for the Horn core, (i) it is not unique, that is, there may be exponentially many inequivalent most relaxed Horn formulas not satisfied by any model outside M; (ii) selecting the Horn core with the maxim*um* number of models (i.e., the one that best approximates M) is NP-complete; furthermore (iii) even approximating the maximum within any constant ratio is NP-complete. □

The computational problems related to updates and belief revisions are in fact much harder. Let Γ be a set of Boolean formulas, and let ϕ be another formula; ϕ will usually be assumed to be of size bounded by a small constant k. We want to compute a new set of formulas $\Gamma+\phi$ —intuitively, the result of updating Γ by ϕ. There are many formalisms in the literature for updating and revising knowledge bases. First, if $\Gamma\&\phi$ *is satisfiable*, then all (with the single exception of [Wi1]) approaches define $\Gamma+\phi$ to be precisely $\Gamma\&\phi$ (we often blur the distinction between a set of formulas and their conjunction). So, suppose that $\Gamma\&\phi$ is unsatisfiable.

1. In the approach introduced by Fagin, Ullman, and Vardi [FUV], and later elaborated on by Ginsberg [Gi], we take $\Gamma+\phi$ to be not a single set of formulas, but the set of all maximal subsets of Γ that are consistent with ϕ, with ϕ added to each. A variant takes the "cross product" of all these formulas.

2. In a more conservative approach, we take $\Gamma+\phi$ to be ϕ plus the *intersection* of all these maximal sets —this is the "when-in-doubt-throw-it-out," or WIDTIO, approach.

3–7. The remaining approaches define $\Gamma+\phi$ implicitly, by its set of models $\mu(\Gamma+\phi)$, given in terms of the set of models of Γ, $\mu(\Gamma)$, and that of ϕ, $\mu(\phi)$ —notice that, since $\Gamma\&\phi$ is unsatisfiable, these two sets are disjoint. All five approaches take $\mu(\Gamma+\phi)$ to be *the projection* of $\mu(\Gamma)$ on $\mu(\phi)$, the subset of $\mu(\phi)$ that is closest to $\mu(\Gamma)$ —and they differ in their notions of a "projection" and "closeness." In Satoh's [Sa] and Dalal's [Da] models, the projection is the subset of $\mu(\phi)$ that achieves minimal distance from *any* model in $\mu(\Gamma)$ (in Dalal's it is minimum Hamming distance, in Satoh's minimal set-theoretic difference). In Borgida's [Bo] and Forbus's [Fo] models, the projection is the subset of $\mu(\phi)$ that achieves minimal distance from *some* model in $\mu(\Gamma)$ (in Forbus it is minimum Hamming distance, in Borgida's minimal set-theoretic difference). Finally, Winslett's [Wi1] approach is a variant of Borgida's, in which the "projection" is preferred over the intersection even if $\Gamma\&\phi$ is satisfiable.

In [EG2], Eiter and Gottlob embark on a systematic study of the complexity issues involved in the various formalisms for updates and revisions. They show that telling whether $\Gamma+\phi\models\psi$ in any of these approaches (this is known as the *counterfactual problem*) is complete for levels in the polynomial hierarchy beyond NP —that is to say, hopelessly complex. When Γ and ϕ are Horn, and ϕ is of bounded size, [EG2] show their only positive result (for adverse complexity results, even in extremely simple cases, in approaches 1 and 2, see Theorem 2 parts (i) and (ii) below): The problem is polynomial in the approaches 3–7. This seems at first sight very promising, since we are interested in updating Horn approximations by bounded formulas. The problem is that *the updated formulas cease being Horn* (part (iii)).

Theorem 2: Computing $\Gamma+\phi$, where Γ is a set of

Horn formulas and ϕ is a single Horn clause with at most three literals:

(i) Is TRANSVERSAL-hard in the Fagin-Ullman-Vardi-Ginsberg [FUV, Gi] approach.

(ii) Is $\mathrm{FP}^{\mathrm{NP}[\log n]}$-complete in the WIDTIO approach (that is, as hard as any problem that requires for its solution the interactive use of an NP oracle $\log n$ times).

(iii) May result in formulas that are not Horn in the model-based approaches. □

3. INCREMENTAL RECOMPILATION

We now describe our scheme for representing *propositional* knowledge in a manner that supports rapid approximate reasoning and minimum-change updates. At time i we represent our knowledge base with two Horn formulas $\underline{\Gamma}_i$ and $\overline{\Gamma}_i$, such that $\underline{\Gamma}_i \models \overline{\Gamma}_i$. We start the process by computing the Horn envelope and core of the initial formula Γ_0, incurring a start-up computational cost —alternatively, we may insist that we always start with a Horn formula. Notice that we are slightly abusing notation, in that $\underline{\Gamma}_i$ and $\overline{\Gamma}_i$ may not necessarily be the Horn envelope and core of some formula Γ_i; they are simply convenient upper (weak) and lower (strict) bounds of the knowledge base being represented.

When the formula is updated by the formula ϕ_i, the new upper and lower bounds are as follows:

$$\overline{\Gamma}_{i+1} := \overline{\overline{\Gamma}_i + \phi_i},$$

$$\underline{\Gamma}_{i+1} := \underline{\underline{\Gamma}_i + \phi_i}.$$

Here '+' denotes any one of the update formalisms discussed (we address towards the end of this section the issue of selecting the appropriate formalism). That is, the new upper bound is the Horn envelope of the updated upper bound, and the new lower bound is the Horn core of the updated lower bound.

Obviously, implementing this knowledge representation proposal relies on computing the Horn envelopes and cores of updated Horn formulas. We therefore now turn to this computational problem.

Updating Horn Formulas

To understand the basic idea, suppose that we want to update a Horn formula Γ by a *clause* $\phi = (\neg x \vee \neg y)$. If $\Gamma \& \phi$ is satisfiable, then the updated formula is precisely this conjunction. So, suppose that $\Gamma \& \phi$ is unsatisfiable; that is, $\Gamma = x \& y \& \Gamma'$ for some Horn formula Γ' not involving x and y. Consider now any model of Γ; it is of the form $m = 11m'$, where 11 is the truth values of x and y, and m' is the remaining part. The models of ϕ that are closest to it (both in minimum Hamming distance and in minimal set difference, as dictated by all five approaches) are the two models $01m'$ and $10m'$. Taking the union over all models of Γ, as the fomalisms by Borgida and Forbus suggest, we conclude that $\Gamma + \phi$, the updated formula, is $(x \neq y) \& \Gamma'$. The Horn envelope of this is easy: It is just $(\neg x \vee \neg y) \& \Gamma'$, while the Horn core is either $x \& \neg y \& \Gamma'$ or $y \& \neg x \& \Gamma'$ —we can choose either one.

As we mentioned in the introduction, in case of a Horn implication, such as $\phi = (x \& y \rightarrow z)$ with Γ of the form $x \& y \& \neg z \& \Gamma'$, the upper and lower approximations are these: $\overline{\Gamma + \phi}$ is $(x \& y \leftrightarrow z) \& \Gamma'$, while $\underline{\Gamma + \phi}$ is $x \& (y \leftrightarrow z) \& \Gamma'$ or $y \& (x \leftrightarrow z) \& \Gamma'$. The generalization to arbitrary Horn formulas is obvious.

Suppose next that $\phi = (\bigwedge_{i=1}^m x_i \rightarrow \bigvee_{j=1}^n y_i)$ is a general clause update, with $n > 1$, and $\Gamma = \bigwedge_{i=1}^m x_i \& \bigwedge_{j=1}^n (\neg y_i) \& \Gamma'$ —again, this is the interesting case. A similar calculation shows that the Horn envelope of $\Gamma + \phi$ is precisely $\bigwedge_{i \neq k} (\neg y_i \vee \neg y_k) \& \bigwedge_{i,j} (y_i \rightarrow x_j) \& \Gamma'$, whereas the Horn core is $(\bigwedge_{i=1}^m x_i \leftrightarrow y_j) \& \Gamma'$, for our choice of j.

Theorem 3: The Horn envelope and core of the update of a Horn formula Γ by a Horn clause ϕ, in any one of the model-based update formalisms, can be computed in time $O(|\Gamma| + |\phi|^2)$. □

There is however a serious problem with our scheme when the updates are non-Horn: As can be seen from the calculation that preceded Theorem 4, the Horn envelope of the updated formula *fails to logically imply the update* —contrary to the intuitive meaning of an "update" or "belief revision," and in violation of the accepted axioms which such formalisms are supposed to satisfy (see, for example, [EG2]). It can be shown that this does not happen not only when ϕ is a single Horn clause, but also whenever it is *any Horn formula.*

Suppose now that ϕ has several clauses. In fact, suppose that ϕ is the conjunction of several *negative* clauses, with no positive literals in them, and that Γ is of the form $x_1 \& \ldots \& x_k \& \Gamma'$, where $x_1 \ldots x_k$ are the variables appearing in ϕ. Consider a model $11 \ldots 1m'$ of Γ; what is the closest in Hamming distance model of ϕ? The answer is *the model that has zeros in those variables among $x_1 \ldots x_k$ which correspond to a minimum hitting set of the clauses* (considered as sets of variables). Therefore, telling whether the Horn envelope of the updated formula (in the Forbus model) implies x_i is equivalent to asking whether i is not involved in any minimum-size hitting set —an coNP-complete problem!

Theorem 4: Computing the Horn envelope of the update of a Horn formula by the conjunction of negative clauses in the Forbus or Dalal formalisms is NP-hard *if the update is allowed to be arbitrarily long.* □

Notice however that, if we restore our assumption that the update is bounded, Theorem 4 is no threat. Also, in the other three formalisms, updates such as these turn out to be easy.

We conjecture that the Horn envelope and core of a Horn formula updated by any bounded formula can be computed in polynomial time in all five model-based update formalisms. In our view, this is an important and challenging technical problem suggested by this work. We know the conjecture is true in several special cases —for example, the one whose unbounded variant was shown NP-complete in Theorem 4— and we have some partial results and ideas that might work for the general case.

Choosing the Right Update Formalism

Of the five model-based update formalisms, which one should we adopt as the update vehicle in our representation scheme? Besides computational efficiency, there is another important desideratum: The property that $\underline{\Gamma}_i \models \overline{\Gamma}_i$, that is, that the "upper and lower bound" indeed imply one another in the desirable direction, must be retained inductively. We can show:

Theorem 5: If $\underline{\Gamma}_i \models \overline{\Gamma}_i$, and the update formalism of Winslett is adopted, then $\underline{\Gamma}_{i+1} \models \overline{\Gamma}_{i+1}$. □

There are examples, to be included in the full paper, which show that the remaining four model-based formalisms may lead to situations in which the conclusion of Theorem 5 is violated.

Characteristic Models and 2SAT Approximations

As it turns out, much of this work can be extended in two directions: To the case in which the Horn formula is represented by its *characteristic models* [KKS2], and to the one in which we approximate from above and below not by a Horn formula but by a 2SAT formula (at most two literals per clause). Details will be presented in the full paper.

The Quality of Approximation

Our approach responds to updates by producing approximations of the knowledge base which become, with new updates, more and more loose. Naturally, its practical applicability rests with the quality of these approximations, and their usefulness in reasoning. This important aspect of our proposal should be evaluated experimentally; we also plan to apply it to situations in AI in which reasoning in a dynamically updated world is well-known to be challenging, such as reasoning about action.

Acknowledgment: We are indebted to Bart Selman for many helpful comments on a preliminary version of the manuscript.

REFERENCES

[Bo] A. Borgida "Language Features for Flexible Handling of Exceptions in Information Systems," *ACM Trans. on Database Systems*, 1993.

[Da] M. Dalal "Investigations into a Theory of Knowledge Baase Revision: Preliminary Report," *Proc. AAAI 88*, 475-479, 1988.

[DP] R. Dechter and J. Pearl "Structure identification in relational data," *Artificial Intelligence*, 58:237-270, 1992.

[EG1] T. Eiter, G. Gottlob "Propositional Cirsumscription and Extended Closed World Reasoning are Π_2^P-complete," *Theoretical Computer Science,* 1993.

[EG2] T. Eiter, G. Gottlob "On the Complexity of Propositional Knowledge Base Revision, Updates, and Counterfactuals," *Artificial Intell., 57* pp. 227–270, 1992.

[EG3] T. Eiter, G. Gottlob "Identifying the minimal transversals of a hypergraph and related problems," to appear in *SIAM J. of Computing.*

[Fo] K. D. Forbus "Introducing Actions in Qualitative Simulation," *Proc. IJCAI 89*, 1273–1278, 1989.

[FUV] R. Fagin, J. D. Ullman, M. Vardi "On the Semantics of Updates in Databases," *Proc. PODS 83*, 352–365, 1983.

[Gi] M. L. Ginsberg "Counterfactuals," *Artificial Intelligence*, 30:35-79, 1986.

[KKS1] H. A. Kautz, M. J. Kearns, B. Selman "Horn approximations of empirical data," to appear in *Artificial Intelligence*, 1994.

[KKS2] H. A. Kautz, M. J. Kearns, B. Selman "Reasoning with Characteristic Models," AAAI 1993.

[KPS] D. Kavvadias, C. H. Papadimitriou, M. Sideri "On Horn Envelopes and Hypergraph Transversals," *Proc. International Symposium on Algorithms and Complexity, Hong-Kong 1993*, Springer-Verlag, 1993.

[Le] H. Levesque, "Making believers out of computers," *Artificial Intelligence*, 30:81-108, 1986.

[Pa] C. H.Papadimitriou *Computational Complexity*, Addison Wesley, 1993.

[Sa] K. Satoh, "Nonmonotonic Reasoning by Minimal Belief Revision," *Proc. of the International Conference on Fifth Generation Computer Systems*, 455-462, 1988.

[SK] B. Selman, H. A. Kautz "Knowledge compilation using Horn approximation," *Proc. AAAI 1991*, 904-909, 1991.

[We] A. Weber, "Updating Proposotional Formulas," *Proc. of the First Conference on Expert Database Systems*, 10:563-603, 1985.

[Wi1] M. Winslett "Reasoning about Action Using a Possible Models Approach," *Proc. AAAI 88*, 88-93, 1988.

[Wi2] M. Winslett *Updating Logical Databases,* Cambridge University Press, 1990.