

A Non-Deterministic Semantics for Tractable Inference

James M. Crawford

i2 Technologies
909 E. Las Colinas Blvd.
12th Floor
Irving, TX 75038

David W. Etherington

Computational Intelligence Research Laboratory
University of Oregon
1269 University of Oregon
Eugene, OR 97403-1269

Abstract

Unit resolution is arguably the most useful known algorithm for tractable reasoning in propositional logic. Intuitively, if one knows a , b , and $a \wedge b \supset c$, then c should be an obvious implication. However, devising a tractable semantics that allows unit resolution has proven to be an elusive goal. We propose a 3-valued semantics for a tractable fragment of propositional logic that is inherently non-deterministic: the denotation of a formula is not uniquely determined by the denotation of the variables it contains. We show that this semantics yields a tractable, sound and complete, decision procedure. We generalize this semantics to a family of semantics, tied to Dalal's notion of intricacy, of increasing deductive power and computational complexity.

Introduction

Despite the recent advances in propositional reasoning power [Selman, Kautz, & Cohen 1993; Crawford & Auton 1996; Bayardo & Miranker 1996; Bayardo & Schrag 1997], knowledge-bases of the size required for common-sense reasoning are still far beyond the reach of complete reasoners. This means that, for the foreseeable future, common-sense reasoners will necessarily be based on tractable, classically-incomplete inference algorithms.

Many such algorithms have been proposed, and these algorithms are generally reasonably easy to understand procedurally. However, understanding inference procedurally is no substitute for understanding what sort of models of the world implicitly underlie the inference. This kind of deeper understanding is the domain of semantics. A sound and complete semantics for an inference algorithm tells us what it means for the world to be consistent with a theory, and thus gives a clearer picture of what the inference algorithm is, and is not, assuming about the world.

Early attempts at devising a semantics for tractable limited deduction [Patel-Schneider 1990; Frisch 1987] were oriented toward limiting potentially unbounded "chaining," i.e., inferences such as deducing b from a and $(a \supset b)$, or deducing $(a \supset c)$ from $(a \supset b)$ and

$(b \supset c)$. It was assumed that such inferences were the source of intractability. While this restriction did, in fact, produce polynomial-time reasoners, the resulting preclusion of unit propagation was a significant restriction. Unit propagation is a powerful, yet computationally cheap, reasoning technique [Schrag & Crawford 1996], and its loss severely limited the utility of these systems.

As a case in point, in a recent paper we attempted to develop a tractable reasoner by restricting inference to a limited context and pruning deductions that branched beyond the bounds of the context [Etherington & Crawford 1996]. Paradoxically, our experimental results showed that such context limitation actually degrades performance in certain cases, making it harder to detect inconsistency. Analysis of the data shows that pruning inference can block unit propagation that would pass outside the context, and this propagation can help quickly detect contradictions in theories with significant numbers of unit clauses. The loss of unit propagation can overwhelm the benefits gained by restricting inference. Such observations have led us to investigate mechanisms for avoiding this problem.

Crawford and Kuipers [1989; 1991] approached tractable deduction by limiting chaining to defined "access paths" in the knowledge base. Intuitively, when attempting to deduce a , the focus of attention is restricted to only rules for deriving a . If such a rule says that a follows from b then attention is extended to b , and then to c if b follows from c , and so on. The essential idea here is the same as unit resolution: follow "linear" chains of inference but do not perform reasoning by cases. Reasoning by cases (i.e., considering the effects of choosing each of several alternatives and then concluding whatever holds in all cases) was then permitted by a dialectical process of making assumptions that then enabled further access-limited reasoning.

Recently, Dalal has furthered this approach, developing a family of tractable deduction systems that do not sacrifice unit propagation [Dalal 1996a]. The essential insight that underlies his work, Crawford and Kuipers', and ours is that it is reasoning by cases, not chaining or disjunction per se, that makes propositional reasoning intractable. However, Dalal's systems have been limited to theories in clausal form. We present a semantics for tractable propositional deduction that corresponds

in many respects to Dalal’s systems, but extends naturally to non-clausal formulae and to limited, tractable, first-order reasoners.

Model Theory

As have previous authors, we begin by basing our semantics on three truth values: **t**, **f**, and **u**. Intuitively we can imagine **u** as corresponding to 0.5 in an integer programming encoding of a propositional formula. In such an encoding **t** maps to 1, **f** maps to 0, $x \vee y$ and $x \wedge y$ map to $x + y$ and $x * y$, respectively, over the 3-element domain. This gives the initial truth table shown in Table 1.

x	y	$\neg y$	$x \vee y$	$x \wedge y$
u	t	f	t	u
u	f	t	u	f
u	u	u	t	f

Table 1: A simple 3-valued truth table¹

This is essentially the semantics used by Dalal [1996b]² but extended to include conjunction (Dalal deals only with clausal theories, first converting non-clausal theories to clausal form).

Unfortunately there is a problem with this semantics: determining satisfiability is NP-complete for non-clausal theories. To see this, observe that we can take any theory and conjoin it with:

$$(x \wedge x') \vee (\neg x \wedge \neg x')$$

for each variable x in the theory, where x' is a unique new variable. These additional clauses force all the variables to take the values **t** or **f**, and thus force the semantics to behave classically. Dalal avoided this problem by dealing only with clausal theories, but we must face it if we are to achieve tractability while allowing general propositional formulae.

There is an additional, more subtle, problem with the initial semantics given above. Every non-unit clause evaluates to **t** when its constituent literals are valued **u**, which means that uncertainty about the constituents does not carry over into uncertainty about their disjunction, which seems peculiar. For example, consider the theory consisting of the clauses $(x \vee y)$, $(x \vee \neg y)$, $(\neg x \vee y)$, and $(\neg x \vee \neg y)$. If we focus on all models in which $x = y = \mathbf{u}$ then all four of these clauses evaluate to **t**. This is, in a sense, more non-classical than we want to be. Furthermore, if x is changed from **u** to **t** then the truth value of $\neg x \vee y$ goes from **t** to **f**. This is even more unfortunate since, ideally, increased certainty about variable values should translate into increased certainty about clause values. To put this another way, as variables values are refined from non-classical to classical values we would like to see the set

¹The purely classical cases are omitted for succinctness.

²For details see the “related work” section.

of models also refine toward the classical models. We formalize this notion of compositional refinement below.

We can address both semantic and tractability issues by adopting a nondeterministic semantics. In order to allow unit propagation without allowing arbitrary reasoning by cases, it is necessary to develop a mechanism that can allow, for example, $(a \vee b)$ and $(\neg b \vee c)$ to evaluate to **t**, without forcing $(a \vee c)$ to do so too. Clearly, the truth values for a and c must be weaker than **t**, and b cannot be classically valued. However, setting all three atoms to **u** and applying the above semantics won’t work either, since all three clauses would evaluate to **t**. Allowing nondeterminism in the semantics for \vee and \wedge provides a way around this dilemma. Intuitively, this nondeterminism reflects the fact that it is computationally too expensive, in general, to work out all the ramifications of uncertainty about particular facts.

x	y	$\neg y$	$x \vee y$	$x \wedge y$
u	t	f	t	u
u	f	t	u	f
u	u	u	t u	f u

Table 2: A nondeterministic truth table³

The modified truth table in Table 2 constitutes such a nondeterministic semantics for propositional logic. Multiple values in a cell indicate a non-deterministic choice in interpretation construction. Thus if $x = y = \mathbf{u}$ then $x \vee y$ will take on the value **u** in some interpretations and the value **t** in others. This solves the tractability problem above because there are models that set $x = x' = \mathbf{u}$, choose $x \wedge x' = \mathbf{u}$ and $\neg x \wedge \neg x' = \mathbf{u}$, and then choose to evaluate the disjunction of the two as **t**. This approach also solves the semantic problem because if $x = y = \mathbf{u}$ then there will be models in which $x \vee y$ is **t** and others in which it is **u**. Thus when x and y are refined to **t**, we can regard the classical model produced as a refinement of the non-classical model that set $\neg x \vee \neg y$ to **u**. In the formal development below we precisely define this notion of compositional refinement and show that it holds in general for the semantics defined by Table 2.

The following definitions make the details of the above discussion precise. Except as explicitly noted, we will use lowercase letters to represent propositional variables, and Greek letters to represent propositional formulae. Theories will be defined in terms of the primitive connectives, \wedge , \vee , and \neg , although we will occasionally use $(\alpha \supset \beta)$ as shorthand for $(\neg \alpha \vee \beta)$. We will restrict our attention to propositional theories in negation-normal form (all negations pushed in as far as possible), and use T and T' to represent propositional

³We assume the obvious generalization of this truth table for n-ary disjunctions and conjunctions: $x_1 \vee \dots \vee x_m$ is valued **t** if some x_i is true, false if all x_i are false, **u** if exactly one x_i is **u**, and **t u** otherwise. Conjunctions are handled similarly.

theories.

Definition 1 A basic assignment for a propositional theory, T , is an assignment of \mathbf{t} , \mathbf{f} , or \mathbf{u} to the propositional variables of T .

Definition 2 A 3-interpretation, M , for a theory, T , consists of a basic assignment for T together with an assignment to the subformulae of T that respects Table 2. We write $M(\alpha)$ for the value assigned by M to α .

Definition 3 M is a 3-model for a formula, α , (written $M \models \alpha$) if M is a 3-interpretation for α and $M(\alpha) = \mathbf{t}$. M is a 3-model for a theory, T , (written $M \models T$) if M is a 3-interpretation for T and $\forall \alpha \in T. M \models \alpha$.

To give the reader a feel for how the semantics is applied, as well as some of its properties, consider the following examples, where a , b , and c are literals. To invalidate a reasoning schema, we must show a model where the premises are all \mathbf{t} , but the conclusion is \mathbf{u} or \mathbf{f} (i.e., *not* \mathbf{t}). If such a model does not exist, the schema is valid.

Example 1 (Unit propagation) Unit propagation allows us to conclude b from a and $a \supset b$. Consider a basic assignment where a is \mathbf{t} . To make $a \supset b$ (i.e. $\neg a \vee b$) \mathbf{t} , b must be \mathbf{t} —i.e., the conclusion must be true. Hence the schema is valid.

Unit propagation is a simple form of modus ponens. We now show that the general form of modus ponens is not valid under our semantics.

Example 2 (Failure of modus ponens) We first construct a model where $a \vee b$ and $(a \vee b) \supset (b \vee c)$ are both \mathbf{t} . Start with the basic assignment $a = b = c = \mathbf{u}$, and for the first formula choose $\mathbf{u} \vee \mathbf{u}$ to be \mathbf{t} . The second formula is a notational shorthand for $(\neg a \wedge \neg b) \vee (b \vee c)$. We can choose $(\neg \mathbf{u} \wedge \neg \mathbf{u}) \vee (\mathbf{u} \vee \mathbf{u})$ to be \mathbf{t} (choose $\mathbf{u} \mathbf{f} \vee \mathbf{u} \mathbf{t}$ as $\mathbf{u} \vee \mathbf{u}$, which reduces to $\mathbf{u} \mathbf{t}$ and choose this to be \mathbf{t}). However, $(b \vee c)$ evaluates to $\mathbf{u} \vee \mathbf{u}$, which can be \mathbf{u} .

Finally, we show that chaining fails in general.

Example 3 (Failure of chaining) Given premises $(a \vee b)$ and $(\neg b \vee c)$, the basic assignment $a = b = c = \mathbf{u}$ can make both true (let $\mathbf{u} \vee \mathbf{u}$ be \mathbf{t} in both cases). However, in the case of $(a \vee c)$, $\mathbf{u} \vee \mathbf{u}$ can also be \mathbf{u} .

The third truth value, \mathbf{u} , represents ambiguities in models that could potentially be refined away, thus reducing uncertainty. This refinement process only restricts the possibilities available, it never creates new alternatives. This is important, since it means that a model containing \mathbf{u} 's should never lead us to dismiss as impossible the “true” model of the world. We can formally capture this idea with the notion of compositional refinement. The essential idea here is that if we refine the basic assignment (by taking some variable from \mathbf{u} to \mathbf{t} or \mathbf{f}) we must be able to find, for each model consistent with the new assignment, a “parent” that was consistent with the old assignment:

Definition 4 An assignment A' is a refinement of an assignment A iff A' differs from A only in that some variables assigned to \mathbf{u} by A are assigned to \mathbf{t} or \mathbf{f} by A' . Refinement relations between interpretations are defined similarly.

Theorem 5 (Compositional Refinement)

Assume that A is the basic assignment for some 3-model of a theory T . For any refinement A' of A , if M' is a 3-model of T with basic assignment A' , then there is some model M of T with basic assignment A such that M' is a refinement of M .

One important corollary of this result is that if V is the set of propositional variables valued classically (\mathbf{t} or \mathbf{f}) by A , then any formula valued \mathbf{t} in all models of T with basic assignment A will also be valued \mathbf{t} in all classical models of T and V .

Proof Theory

Two of the rules of inference used most commonly in tractable reasoners are unit resolution and unit subsumption. Unit resolution allows x to be deduced from $(x \vee y \vee z)$, $\neg y$, and $\neg z$. Unit subsumption allows $(x \vee y \vee z)$ to be deduced from x . In terms of proof rules, these can be written (for δ , δ_i literals, α , β formulae) as:

$$R1 : \text{unit subsumption} \quad \frac{T \models \delta}{T \models \delta \vee \alpha}$$

$$R2 : \text{unit resolution} \quad \begin{array}{l} T \models \delta_1 \vee \dots \vee \delta_m \\ T \models \neg \delta_1 \\ \vdots \\ T \models \neg \delta_{m-1} \\ \hline T \models \delta_m \end{array}$$

To this we add the obvious rule that a theory's constituent terms can be deduced from the theory, and a rule for contradictory theories:

$$R3 : \text{containment} \quad \frac{\alpha \in T}{T \models \alpha}$$

$$R4 : \text{contradiction} \quad \frac{T \models \delta}{T \models \neg \delta} \\ \hline T \models \alpha$$

(where by $\alpha \in T$ we mean that T is a conjunction of formulae, one of which is α). We also need the following rules for handling conjunctions in the non-clausal case:

$$R5 : \wedge\text{-elimination} \quad \frac{T \models \alpha \wedge \beta}{T \models \alpha}$$

$$R6 : \wedge\text{-introduction} \quad \frac{T \models \alpha}{T \models \beta} \\ \hline T \models \alpha \wedge \beta$$

This set of rules is slightly weaker than the tractable inference theory BCP defined by Dalal [1996b], as it must be to preserve tractability in the non-clausal case. It is, however, precisely equivalent to the semantics defined above:

Theorem 6 $T_1 \Vdash T_2$ iff $T_1 \models T_2$.

In a real sense the core of this proof theory is unit resolution. This observation is captured in the following proposition.

Proposition 7 A propositional theory, T , has a 3-model iff unit propagation over T does not produce both a literal and its negation.

To see why this must hold, note that we can construct an interpretation, I , for a theory, T , using the following algorithm:

Unit propagate T to completion. Assign \mathbf{t} to each variable appearing positively in a unit clause (or a top-level unit conjunct), and \mathbf{f} to each variable appearing negatively in a unit clause in the resulting theory. Assign \mathbf{u} to the remaining variables. Beginning with the most deeply nested connectives, using the truth table of Table 2, choose the greatest allowable truth value ($\mathbf{t} > \mathbf{u} > \mathbf{f}$) to assign.

If no variable is assigned both \mathbf{t} and \mathbf{f} , then I is a 3-model for T .

Discussion

Using non-deterministic truth tables does introduce some interesting behavior. The primary difference from classical semantics hinges on an issue we call “pegging.” Consider the formula:

$$(x \vee y) \wedge (x \vee y \vee z)$$

The non-deterministic nature of the semantics above allows the first clause to take on the truth value \mathbf{t} and the second (subsumed) clause to take on the value \mathbf{u} . In a sense an interpretation has no memory of the decisions that are made on the value of formulae. As a result, the value assigned to $x \vee y$ in one part of the theory is not pegged to the value assigned to subsuming clause in another part of the theory. We have explored a variety of “pegged” versions of this semantics, and found them all to be incoherent or intractable (or both). While this is not a proof that no suitable pegged version could be found, pegging *seems* to provide a back-door entry for intractability, and we have decided to content ourselves with an unpegged semantics. This is, however, an area that deserves further study.

It is also interesting to note that some of the truisms of classical model theory do not hold for this semantics. In particular, there are theories, T , for which $M(T) \cap M(\neg T) \neq \emptyset$. For example, $(a \vee c) \wedge (b \vee d)$ shares a model where $a = b = c = d = \mathbf{u}$ with $(\neg a \wedge \neg c) \vee (\neg b \vee \neg d)$. Similarly, there are theories, T , for which $M(T) \cup M(\neg T) \neq M(\emptyset)$ (no variable appearing in a unit clause in T can be valued \mathbf{u}).

Perhaps most disturbing, however, is that certain transformations that preserve classical equivalence do not preserve equivalence under our new semantics. For example, conversion to clausal form may change the set of interpretations. Thus, $(a \wedge b) \vee (a \wedge c)$ can be valued \mathbf{t} , \mathbf{f} , or \mathbf{u} when $a = b = c = \mathbf{u}$, but $(a \wedge (b \vee c))$ can be only valued \mathbf{f} or \mathbf{u} .

A Family of Tractable Semantics

The semantics described so far allows for sound but classically incomplete tractable inference from a set of premises. This notion of inference essentially amounts to avoiding reasoning by cases. In some circumstances, however, such reasoning is essential. We thus seek a way of reintroducing a limited form of reasoning by cases without at the same time losing tractability. One obvious way to do this is to limit the “depth” of reasoning by cases. Our goal is to allow assumptions to be made that allow “lemmas” to be derived that further the inference process, but to restrict the complexity of these lemmas to avoid a combinatorial explosion.

Essentially, we allow k assumptions to be made. If a set of assumptions allows us to derive a contradiction, we deduce the negation of that set of assumptions. An interesting question is then, “What level of nesting is required to reach a particular conclusion, α ?” This depth corresponds to how difficult it is to deduce α , and is essentially analogous to Dalal’s notion of *intricacy* [Dalal 1996a].

Using this notion of nested assumptions, we define a hierarchical family of semantics corresponding to computationally more expensive reasoning systems.

Definition 8 $M \models_0 T$ iff $M \models T$.

For $k > 0$, $M \models_k T$ iff $M \models T$ and, for all clauses, γ , such that $|\gamma| \leq k$ and $M(\gamma) = \mathbf{u}$:

$$\begin{aligned} &\exists M'. M' \models_{(k-|\gamma|)} (T \wedge \gamma), \text{ and} \\ &\exists M'. M' \models_{(k-|\gamma|)} (T \wedge \neg\gamma). \end{aligned}$$

If $M \models_k T$, we say M is a k -model for T and T is k -satisfiable.

Basically, a k -model is forced to be “honest” about the possible truth values of clauses with no more than k literals, in the sense that they cannot be labeled \mathbf{u} without explicitly allowing the possibility that the actual value may be either of \mathbf{f} and \mathbf{t} .

We define a notion of entailment in the obvious way.

Definition 9

$T \models_k \alpha$ iff $M(\alpha) = \mathbf{t}$ for all k -models, M , of T .

To generate a sound and complete proof theory we augment the proof theory given earlier with the following rules (for clauses γ and formulae α and β):

$$R7 : \text{base case} \quad \frac{T \Vdash \alpha}{T \Vdash_k \alpha}$$

$$R8 : \text{contradiction} \quad \frac{T, \neg\gamma \Vdash \mathbf{f}}{\frac{|\gamma| \leq k}{T \Vdash_k \gamma}}$$

$$R9 : \text{weak modus ponens} \quad \begin{array}{l} T \Vdash_k \alpha \\ T \wedge \alpha \Vdash_k \beta \\ |\alpha| \leq k \\ \hline T \Vdash_k \beta \end{array}$$

Theorem 10 (Soundness+Completeness)

If $T \Vdash_k \gamma$ then $T \models_k \gamma$. If T is k -satisfiable and $T \models_k \gamma$ then $T \Vdash_k \gamma$.

The complexity of determining whether $T \Vdash_k \gamma$ is polynomial in the sizes of T and γ , but it increases exponentially with k . We conjecture that when k equals the number of variables in T , the \Vdash relation is classically complete, but we have not yet proven this.

Related Work

Cadoli and Schaerf [1992] present a semantics in which chaining is restricted to a set S of literals in the theory. In the clausal case, entailment can be determined in PTIME for fixed S . However, their approach differs from ours in that unit resolution is allowed only within S , not throughout the theory. Based on our earlier experience, described in the introduction, however, we expect this restriction to be a significant limitation.

The basic idea of limiting the depth of assumption nesting goes back at least to Access-Limited Logic (ALL) [Crawford & Kuipers 1989] and Socratic sequent systems [McAllester, Givan, & Fatima 1989]. The results reported here grew out of our attempts to semantically characterize ALL.

The first attempt to define a semantics for such limited inference is Dalal’s recent work [Dalal 1996b]. Dalal’s semantics assigns real numbers in the range $[0,1]$ to all variables in the theory, and assigns $1 - x$ to $\neg x$. The value of any clause is the sum of the values in the clause. E.g., if $x = 0.4$ and $y = 0.8$ then $(x \vee y) = 1.2$. A theory is satisfied by a model if the model makes every clause in the theory ≥ 1 .

The precise relationship between our two semantics is given by the following results.

Let \models_D mean “models under Dalal’s semantics”. We can define a mapping g from his models to ours by:

$$\begin{array}{ll} g(M)(x) = \mathbf{t} & \text{if } M(x) = 1 \\ g(M)(x) = \mathbf{f} & \text{if } M(x) = 0 \\ g(M)(x) = \mathbf{u} & \text{otherwise} \end{array}$$

then for any disjunction:

$$\begin{array}{ll} g(M)(x_1 \vee \dots \vee x_k) = \mathbf{t} & \text{if } M(x_1 \vee \dots \vee x_k) \geq 1 \\ g(M)(x_1 \vee \dots \vee x_k) = \mathbf{f} & \text{if } M(x_1 \vee \dots \vee x_k) = 0 \\ g(M)(x_1 \vee \dots \vee x_k) = \mathbf{u} & \text{otherwise} \end{array}$$

Note that we can do this only because our semantics is non-deterministic (i.e., sometimes taking $\mathbf{u} \vee \mathbf{u}$ to \mathbf{u} and sometimes taking it to \mathbf{t}). Since Dalal handles only clausal theories directly⁴ this completely defines the mapping g .

Lemma 11 $M \models_D T$ iff $g(M) \models T$.

⁴Non-clausal theories are converted to clausal form.

To map the other way we take any model M under our semantics and map:

$$\begin{array}{ll} h(M)(x) = 1 & \text{if } M(x) = \mathbf{t} \\ h(M)(x) = 0 & \text{if } M(x) = \mathbf{f} \\ h(M)(x) = 0.5 & \text{if } M(x) = \mathbf{u} \end{array}$$

Since Dalal’s semantics is compositional this completely defines the mapping h . The correspondence in this case is more limited since we must restrict our attention to models that are deterministic.

Lemma 12 If M always maps $\mathbf{u} \vee \mathbf{u}$ to \mathbf{t} then $M \models T$ iff $h(M) \models_D T$.

Dalal defines $T_1 \models_D T_2$ iff all models of T_1 do not assign T_2 the value \mathbf{f} . This is different from our (we believe more intuitive) requirement that T_2 takes the value \mathbf{t} in all models of T_1 . Still, one can show the following relationship:

Theorem 13 If $T_1 \models T_2$ then $T_1 \models_D T_2$.

The converse is not true. Take $T_1 = (x \vee y)$ and $T_2 = (x \vee y \vee z)$. $T_1 \models_D T_2$ (in fact, $T_1 \models T_2$). However, our semantics allows models that map $(x \vee y)$ to \mathbf{t} and $(x \vee y \vee z)$ to \mathbf{u} (this is, in fact, a necessary consequence of our lack of pegging).

Our extension to \Vdash_k is also quite similar to Dalal’s. Our proof rules $R7$ and $R9$ are identical to his. The differences in the base semantics necessitated our addition of rule $R8$ (for Dalal our rule $R8$ is a special case of rule $R9$).

Conclusions and Future Work

We have presented a nondeterministic 3-valued semantics for propositional logic, and have shown that it is equivalent to a tractable proof theory based on unit resolution. The semantics has the virtue that it supports compositional refinement: increasing certainty about the truth values of the atomic formulae leads to increased certainty about those of general formulae. It also has the unusual property that the truth value of a formula is not uniquely determined by the truth value of the terms in the formula.

We then generalized our semantics to a family of semantics for a range of tractable inference mechanisms similar to that proposed by Dalal. For each member of the family we described a set of sound and complete proof rules.

We have been exploring extensions of our approach to the first-order case, where it shows promise in providing a tractable notion of context-limited first-order reasoning. The extension is non-trivial because in classical first-order logic, unit-resolution alone is undecidable. Intuitively this is because one can write axioms saying things like, “Every man has a father who is a man,” which force infinite chains of “linear” inferences. The reason these chains do not seem to bother common-sense reasoning seems to be that we focus our attention on a limited set of objects that we are reasoning about and let this focus “grey out” around the edges (e.g.

we will think about John and perhaps be aware of his father, but we will only mentally instantiate John's father's father if we have some reason to believe he is significant).

Semantically we capture this notion of a focus of attention by introducing an explicit model-theoretic context object at the first-order level; all universal quantifications are scoped to just objects in the context and witnesses for existential quantifiers are required to be in the context or to have an unknown (**u**) relationship to the context. One nice property of this extension is that the truth value **u** plays a similar role in propositional reasoning and in first-order reasoning—in both cases it marks the points at which reasoning was cut off because of computational limits. Variables bound to **u** are primary targets for reasoning by cases—the source of intractability in propositional logic. Domain objects with an unknown relationship to the context are targets for extending the context since they represent existential witnesses that are not fully instantiated.

The details of the first-order case will be described in a future paper.

Acknowledgments

We would like to thank Mukesh Dalal and members of CIRL for productive discussions, as well as the anonymous referees for their helpful comments.

This work was sponsored in part by the National Science Foundation (NSF) under grant number IRI-94-12205, and by the Defense Advanced Research Projects Agency (DARPA) and the Air Force Research Laboratory (AFRL) under agreement numbers F30602-95-1-0023 and F30602-97-1-0294. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. The views, findings, and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the NSF, DARPA, AFRL, or the U.S. Government.

References

Bayardo, R. J., and Miranker, D. P. 1996. A complexity analysis of space-bounded learning for the constraint satisfaction problem. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-96)*, 298–304.

Bayardo, R. J., and Schrag, R. C. 1997. Using CSP look-back techniques to solve real-world sat instances. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI-97)*, 203–208.

Cadoli, M., and Schaerf, M. 1992. Tractable reasoning via approximation. In *Proceedings of the AAAI Workshop on Tractable Reasoning*, 12–15.

Crawford, J. M., and Auton, L. D. 1996. Experimental results on the crossover point in random 3SAT. *Artificial Intelligence* 81:13–59.

Crawford, J., and Kuipers, B. 1989. Towards a theory of access-limited logic for knowledge representation. In *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning (KR'89)*, 67–78.

Crawford, J., and Kuipers, B. 1991. ALGERNON – A tractable system for knowledge-representation. In *AAAI Spring Symposium on Implemented Knowledge Representation and Reasoning Systems*.

Dalal, M. 1996a. Anytime families of tractable propositional reasoners. In *Proceedings of the Fourth International Symposium on Artificial Intelligence and Mathematics (AI/MATH-96)*, 42–45.

Dalal, M. 1996b. Semantics of an anytime family of reasoners. In *Proceedings of the European Conference on Artificial Intelligence (ECAI-96)*, 360–364.

Etherington, D., and Crawford, J. 1996. Toward efficient default reasoning. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-96)*, 627–632.

Frisch, A. 1987. Inference without chaining. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-87)*, 515–519.

McAllester, D.; Givan, R.; and Fatima, T. 1989. Taxonomic syntax for first order inference. In *Proceedings of the First International Conference on Principles of Knowledge Representation and Reasoning (KR'89)*, 289–300.

Patel-Schneider, P. F. 1990. A decidable first-order logic for knowledge representation. *Journal of Automated Reasoning* 6:361–388.

Schrag, R., and Crawford, J. 1996. Implicates and prime implicates in random 3SAT. *Artificial Intelligence* 81:199–222.

Selman, B.; Kautz, H.; and Cohen, B. 1993. Local search strategies for satisfiability testing. In *Proceedings 1993 DIMACS Workshop on Maximum Clique, Graph Coloring, and Satisfiability*, 521–523.