Adding Number Restrictions to a Four-Valued Terminological Logic

Peter F. Patel-Schneider
Schlumberger Palo Alto Research
3340 Hillview Avenue
Palo Alto, California 94304

Abstract

An intuitive four-valued semantics can be used to develop expressively powerful terminological logics which have tractable subsumption. If a four-valued identity is also used, number restrictions can be added to the logic while retaining tractability. The subsumptions supported by the logic are a type of "structural" subsumption, where each structural component of one concept must have an analogue in the other concept. Structural subsumption captures an important set of subsumptions, similar to the subsumptions computed in KL-ONE and NIKL. This shows that the trade-off between expressive power and computational tractability which plagues terminological logics based on standard, two-valued semantics can be defeated while still retaining a useful and semantically supported set of subsumptions.

1 Introduction

Terminological logics formalize the notion of frames—a notion present in many current knowledge representation systems—as structured types, often called *concepts*. These logics include a set of syntactic constructs that form concepts, and other, related, notions such as *roles*. Terminological logics are based on formal model-theoretic semantics which provide firm definitions for the syntactic constructs of the logic.

The allowable concepts vary between different terminological logics but generally concepts are the conjunction of a set of more general concepts and a set of restrictions on the attributes of instances of the concept. Such concepts can be loosely rendered as noun phrases such as

> a **student** and a **female** whose **major** is a **computer science major**, and who has at least **three enrolled courses**, each of which is a **graduate course** whose **department** is an **engineering department**.

Terminological logics are part of KL-ONE [Brachman and Schmolze, 1985], NIKL [Moser, 1983], KRYPTON [Brachman *et al.*, 1983, Brachman *et al.*, 1985], and KANDOR [Patel-Schneider, 1984], as well as several other knowledge representation systems.

The most important operation in terminological logics is determining if one concept *subsumes*—is more general than—another. A formal definition of subsumption is part of the semantics of terminological logics. Informally, one concept subsumes another if any object that satisfies the requirements of the second concept (i.e., is an instance of the second concept) must also satisfy the requirements of the first. For example, the concept

> **person** with at least **two children**

subsumes the concept

> **person** with at least **three children** who are **lawyers**

in standard terminological logics. This is so because, in the standard semantics for terminological logics, all instances of the second concept must also be instances of the first.

However, as shown by Levesque and Brachman [1987], computing subsumption is intractable in expressively powerful terminological logics based on standard semantics. This intractability is a severe problem, since terminological reasoners would be useful in many knowledge-based systems, and it is not desirable to have components of knowledge-based systems that may take an extremely long time to compute common operations. To achieve tractable subsumption, at least in the worst case, the logic must be expressively weak—too weak to be usable in knowledge-based systems.

The trade-off between expressive power and computational tractability can be defused by using a weak semantics for terminological logics—one that supports fewer subsumption relationships—resulting in tractable subsumption for expressively powerful logics. This solution retains a firm semantic foundation for the system, as opposed to the more usual method of achieving tractability by providing a sound but not complete reasoner (as in KL-ONE and NIKL).

A tractable terminological logic using a weak semantics based on the four truth values of tautological entailment [Belnap, 1977, Levesque, 1984] has been developed [Patel-Schneider, 1986]. The logic is more expressive than the terminological logic that Levesque and Brachman discovered to be computationally intractable in the standard semantics. However, it is still expressively weak, as it does not include number restrictions—a construct appearing in many semantic networks and frame-based knowledge representation systems.

A number restriction is a way of restricting the number of values that a role has. For example, "at least **three children** who are **lawyers**" is a number restriction. As number restrictions are useful in many domains, a terminological logic without number restrictions is lacking an extremely useful construct.

Number restrictions introduce a new source of complexity to terminological logics. When they are added, new semantic ideas—such as identity and cardinality—have to be considered, thus complicating the semantics and the analysis of the logic.

The computational problem with number restrictions is that the standard two-valued notion of identity sanctions subsumptions that are hard to compute, even in an otherwise four-valued semantics. This paper presents a four-valued notion of identity that solves this computational problem, resulting in a terminological logic incorporating number restrictions that has tractable subsumption—subsumption similar to the "structural" subsumption of KL-ONE and NIKL.

Of course, there is a price to be paid for using a four-valued identity. A four-valued identity is yet another change from the standard semantics, resulting in less correspondence between the semantics of the terminological logic and the standard semantics. However, the change is not too unappealing, and incorporating number restrictions while still retaining tractable subsumption and a similarity to subsumption in KL-ONE and NIKL is useful.

2 Syntax and Intuitive Meaning

The terminological logic developed here has two major syntactic types—*concepts* and *roles*—corresponding to the frames and slots of most frame-based knowledge representation systems. Concepts represent collections of related individuals and roles describe relations between these individuals. The intuitive meaning of the various constructs in the language are derived from the intuitive meanings of constructs in frame-based knowledge representation systems.

The logic mostly is an extension of the logic in [Patel-Schneider, 1986]. It is closely related to the terminological logics of KL-ONE, NIKL, KRYPTON, and KANDOR.

Concepts can be formed in the following ways:

```
<concept> ::=   <atomic concept> |
                (and <concept>+) |
                (all <role> <concept>) |
                (atleast <minimum> <role>) |
                (atmost <maximum> <role>)
<minimum> ::=   <positive integer>
<maximum> ::=   <non-negative integer>
```

The construct (**and** <concept>$^+$) is a conjunction construct. Informally, an individual belongs to (**and** C_1 C_2) if it belongs to both C_1 and C_2. The construct (**all** <role> <concept>) is a role restriction construct. Informally, an individual α belongs to (**all** R C) if, for every individual β, either α is not related to β by R, or β belongs to C.

The constructs (**atleast** <minimum> <role>) and (**atmost** <maximum> <role>) are number restriction constructs. Informally, an individual belongs to (**atleast** n R) if it is related to at least n distinct individuals by R. If n is 1, this reduces to a role filler existence construct. An individual belongs to (**atmost** n R) if it is related to at most n distinct individuals by R.

Roles can be formed in the following ways:

```
<role> ::=      <atomic role> |
                (and <role>+) |
                (restrict <role> <concept>)
```

The construct (**and** <role>$^+$) is a conjunction construct, similar to (**and** <concept>$^+$). The construct (**restrict** <role> <concept>) is a restriction construct. Two individuals are related by (**restrict** R C) if they are related by R and the second is also an instance of C.

The addition of number restrictions is the major change between this language and the language of [Patel-Schneider, 1986]. This addition brings the expressive power of the language nearly up to par with the terminological languages of KL-ONE and NIKL.

3 Formal Semantics

The formal semantics of the logic is an extension of the semantics of [Patel-Schneider, 1986]—supporting number restrictions via a notion of cardinality based on a four-valued identity. The basic ideas underlying the semantics are similar to the ideas underlying other denotational semantics. The semantics is based on semantic structures or possible worlds, each of which contains a set of individuals and a mapping from syntactic constructs—concepts and roles—into their meaning in the semantic structure. The truth values of this semantics are $\{t\}$ or *true*, $\{f\}$ or *false*, $\{\}$ or *unknown*, and $\{t, f\}$ or *contradictory*.[1] Thus the set of truth values form the powerset of $\{t, f\}$, written $2^{\{t,f\}}$.

A *semantic structure* is a triple, $\langle D, V, I\rangle$, where D is a set of individuals, V is a function that takes concepts and roles into their *extension*, and I is an identity relationship over D. The extension of a concept is a mapping from D to $2^{\{t,f\}}$. The extension of a concept is thus a four-valued characteristic function—not a two-valued characteristic function. Similarly, the extension of a role is a mapping from $D \times D$ to $2^{\{t,f\}}$. The identity relationship is also a mapping from $D \times D$ to $2^{\{t,f\}}$, which must satisfy

1. $I(d, d) = \{t\}$,
2. $I(d, e) = I(d, e)$, and
3. if $t \in I(d, e)$ and $t \in I(e, f)$ then $t \in I(d, f)$,

for all $d, e, f \in D$. These restrictions make the identity relationship into a four-valued version of an equivalence relation.

Although this semantics is not too far distant from a standard two-valued extensional semantics, there are some differences that need explanation. One way of motivating these differences is to treat treat the extension of a concept, and also of a role, as two extensions, the *positive extension* and the *negative extension*. The positive extension of the concept C is the set of individuals that belong to the concept—defined as $\{d \in D : t \in V[\mathsf{C}](d)\}$. The negative extension of the concept C is the set of individuals that definitely do not belong to the concept—defined

[1] A slightly different set of truth values that could be used is the set $\{\{t\}, \{f\}, \{t, f\}\}$ used by Frisch [1985]. This set of truth values gives a slightly stronger logic, which may be tractable here, at the expense of removing a useful symmetry. Note that the set of truth values $\{\{\}, \{t\}, \{f\}\}$, used by some of the popular three-valued logics, is usually as intractable as two-valued logics.

as $\{d \in D : f \in V[\mathsf{C}](d)\}$. Unlike the case in two-valued semantics, these two sets need not be complements of each other—there may be individuals that are members of neither of these sets, and also individuals that are members of both of these sets.

Individuals that are members of neither set are not known to belong to the concept and are also not known not to belong to the concept. This is a perfectly reasonable state for a system that is not a perfect reasoner or does not have complete information. Individuals which are in both the positive and negative extension of a concept can be thought of as inconsistent with respect to that concept in that there is evidence to indicate that they are both in the extension of the concept and also (conflicting) evidence to indicate that they are not in the extension of the concept. (Such individuals need not be contradictory with respect to other concepts). This is a slightly harder state to rationalize but can be considered a possibility in the light of inconsistent information.

The difference between this semantics and the one in [Patel-Schneider, 1986] is the presence of the four-valued identity relationship. This relationship is easier to understand if viewed in a manner similar to the positive and negative extension viewing of the extension function. Under this view, if $t \in I(d,e)$ then d and e are known to be identical, and if $f \in I(d,e)$ then d and e are known not to be identical, *i.e.*, known to be distinct. As above, it is possible that two individuals are neither known to be identical nor known not to be identical, and it is also possible that two individuals are both known to be identical and known not to be identical. The reflexive, symmetric, and transitive nature of the identity relationship do, however, serve to make it similar to the standard two-valued notion of equality, and thus makes the change reasonably palatable.

A notion of cardinality can be derived from this four-valued identity. A set does not have a unique cardinality, but instead has a minimum cardinality, based on which of its members are known to be distinct, and a maximum cardinality, based on how many members it has which are not known to be identical. The minimum cardinality of a set, X, is defined to be the size of its largest subset for which all elements are known to be distinct,

$$\mathrm{minc}(X) = \max\left\{|Y| : \begin{array}{c} Y \subseteq X \wedge \\ f \in I(d,e), \forall d,e \in Y, d \neq e \end{array}\right\}.$$

Similarly, the maximum cardinality of a set, X, is defined to be the size of its largest subset for which no two elements are known to be identical,

$$\mathrm{maxc}(X) = \max\left\{|Y| : \begin{array}{c} Y \subseteq X \wedge \\ t \notin I(d,e), \forall d,e \in Y, d \neq e \end{array}\right\}.$$

It is possible for the maximum and minimum cardinality of a set to be different. For example, if no identity relationships, positive or negative, are known, then the maximum cardinality of a set is its standard cardinality and its minimum cardinality is 1. It is also possible for the maximum cardinality of a set to be less than the minimum cardinality of a set. For example, if the identity relationship is total—everything is both identical and distinct from everything else—then the maximum cardinality of a set is 1 and its minimum cardinality is its standard cardinality.

The four-valued identity and the derived notions of minimum and maximum cardinality form another departure from the standard semantics. The main problem with this departure is not the change from a two-valued identity to a four-valued one, which is in keeping with the basic four-valued nature of the semantics, but the associated divorcing of identity from equality in the domain. The four-valued identity weakens the connection between elements of the domain and objects in the world, suggesting instead an interpretation where elements of the domain are more akin to descriptions. Going from individuals to descriptions is not a fatal problem, but requires some rethinking of how well the semantics corresponds to its desired role.

The semantics can perhaps best be viewed as a semantics of belief, where the elements of the domain are descriptions in some agent's belief space. In this view of the semantics, if $f \in I(d,e)$, then d and e are believed to be descriptions of distinct objects. Similarly, if $t \in V[\mathsf{C}](e)$, then d is believed to be a description of an object that belongs to the extension of C. Of course, this view does not change the underlying four-valued nature of the semantics, so it is possible to have incomplete and inconsistent beliefs about identity.

The extensions of non-atomic concepts and roles are specified in terms of conditions that they have to meet:

$t \in V[(\mathsf{and}\ \mathsf{C}_1 \ldots \mathsf{C}_n)](d)$ iff for each i, $t \in V[\mathsf{C}_i](d)$
$f \in V[(\mathsf{and}\ \mathsf{C}_1 \ldots \mathsf{C}_n)](d)$ iff for some i, $f \in V[\mathsf{C}_i](d)$
$t \in V[(\mathsf{all}\ \mathsf{R}\ \mathsf{C})](d)$ iff $\forall e\, f \in V[\mathsf{R}](d,e)$ or $t \in V[\mathsf{C}](e)$
$f \in V[(\mathsf{all}\ \mathsf{R}\ \mathsf{C})](d)$ iff $\exists e\, t \in V[\mathsf{R}](d,e)$ and $f \in V[\mathsf{C}](e)$
$t \in V[(\mathsf{atleast}\ \mathsf{m}\ \mathsf{R})](d)$ iff $\mathrm{minc}\{e : t \in V[\mathsf{R}](d,e)\} \geq m$
$f \in V[(\mathsf{atleast}\ \mathsf{m}\ \mathsf{R})](d)$ iff $\mathrm{maxc}\{e : f \notin V[\mathsf{R}](d,e)\} < m$
$t \in V[(\mathsf{atmost}\ \mathsf{m}\ \mathsf{R})](d)$ iff $\mathrm{maxc}\{e : f \notin V[\mathsf{R}](d,e)\} \leq m$
$f \in V[(\mathsf{atmost}\ \mathsf{m}\ \mathsf{R})](d)$ iff $\mathrm{minc}\{e : t \in V[\mathsf{R}](d,e)\} > m$
$t \in V[(\mathsf{and}\ \mathsf{R}_1 \ldots \mathsf{R}_n)](d,e)$ iff for each i, $t \in V[\mathsf{R}_i](d,e)$
$f \in V[(\mathsf{and}\ \mathsf{R}_1 \ldots \mathsf{R}_n)](d,e)$ iff for some i, $f \in V[\mathsf{R}_i](d,e)$
$t \in V[(\mathsf{restrict}\ \mathsf{R}\ \mathsf{C})](d,e)$ iff $t \in V[\mathsf{R}](d,e)$ and $t \in V[\mathsf{C}](e)$
$f \in V[(\mathsf{restrict}\ \mathsf{R}\ \mathsf{C})](d,e)$ iff $f \in V[\mathsf{R}](d,e)$ or $f \in V[\mathsf{C}](e)$

These conditions are designed so that the formal semantics corresponds closely to the previously-discussed informal meaning of concepts and roles.

For example, the positive extension of (and C_1 C_2) must be the intersection of the positive extension of C_1 and C_2 and its negative extension must be the union of their negative extensions. In this way the intuitive notion of conjunction is made formal. Similarly, the conditions above require that if an element of the domain is in the positive extension of (atleast m R) then it must be related to at least m domain elements, known to be pairwise distinct, by the positive extension of R. Also, if an element of the domain, d, is in the positive extension of (atmost m R) then any set of domain elements, no two of which are known to be identical, that are not known to be related to d by the negative extension of R must have cardinality at most m. In this way the intuitive semantics of number restrictions are captured in a four-valued framework.

The final part of the semantics is the definition of subsumption:

Definition 1 *One concept or role is subsumed by another, written* C $\Rightarrow$ C′, *if the positive extension of the first is always a subset of the positive extension of the second and* *the negative extension of the second is always a subset of the negative extension of the first.*

This definition again corresponds closely to the informal notion of one concept being more general than another.

4 Discussion

The semantics defined here has a close relationship to standard, two-valued semantics for terminological logics as defined by Levesque and Brachman [1987].

Define a *model* as a semantic structure where

1. for every concept C, the positive and negative extensions of C are disjoint and together exhaust the set of individuals of the model,
2. the positive and negative extensions of roles are also disjoint and exhaustive, and
3. the identity relationship is equality.

In such semantic structures the above semantics, including the definition of subsumption, reduces to a standard two-valued semantics for terminological logics. Because of this inclusion relationship, all reasoning in this logic is sound with respect to standard terminological logics.

The conditions for concepts and roles, and also the definitions of cardinality and subsumption, are just a reinterpretation, in a four-valued setting, of the standard two-valued conditions and definitions. There is nothing added besides what is needed to get from two truth values to four truth values. Thus the semantics is closely related to intuitions about the meanings of concepts and roles.

The changes in the semantics—going from two to four truth values and a four-valued identity—are reasonable for systems with limited reasoning power. Such systems do not have total information, thus the presence of truth-value gaps, and also cannot resolve inconsistencies, thus allowing for inconsistent situations. The four truth values of the logic have also been previously used to develop limited reasoners in other areas [Levesque, 1984; Patel-Schneider, to appear].

The set of subsumptions supported by this logic forms an interesting and useful set. Since subsumption is sound with respect to standard terminological logics, if one concept subsumes another in this logic then it will also do so in a standard, two-valued terminological logic. Soundness of subsumption is an important requirement if the semantics is to capture some of the intuitive ideas behind terminological logics.

The sort of subsumption relationships that are valid in this logic are the simple ones, such as

```
(and person (atleast 2 child))
```

subsuming

```
(and person (atleast 3 (restrict child lawyer))),
```

and

```
(and person (atmost 4 (restrict child doctor)))
```

subsuming

```
(and person female (atmost 3 child)).
```

As these examples show, the valid subsumption relationships are not trivial, and include at least some interesting subsumption relationships.

Subsumption relationships involving *modus ponens* are not valid here. For example,

```
(and person
     (all friend doctor)
     (all (restrict friend doctor) (atleast 1 speciality)))
```

is not subsumed by

```
(and person
     (all friend (atleast 1 speciality))),
```

because in four-valued semantic structures it is possible that some friend might both be a doctor and not be a doctor, as well as not specializing. Because the friend is a doctor, (all friend doctor) is not falsified; because the friend is not a doctor, (all (restrict friend doctor) (atleast 1 speciality)) is not falsified; however, because the friend does not specialize, (all friend (atleast 1 speciality)) is falsified, and thus the subsumption relationship does not hold. Also,

```
(and person
     (atleast 2 friend)
     (all friend doctor))
```

is not subsumed by

```
(and person
     (atleast 2 (restrict friend doctor))),
```

because some individual might both be a friend and not be a friend.

Subsumption relationships that require reasoning from the law of the excluded middle for identity are also not valid here. For example[2], in a two-valued terminological logic

```
(and (atleast 1 (restrict child lawyer))
     (atleast 1 (restrict child doctor)))
```

would be subsumed by

```
(or (atleast 2 child)
    (atleast 1 (restrict child (and lawyer doctor)))),
```

because either the child that is a lawyer is different from the child that is a doctor, in which case there are two children, or they are identical, in which case there is one child which is both a doctor or a lawyer. In the four-valued logic this is not a valid subsumption because it is possible to be uncertain about whether the doctor and the lawyer are identical. These subsumptions are hard to compute, which forms one of the reasons for the switch to a four-valued identity.

The subsumption relationships that are valid form a sort of "structural" subsumption[3]—where each structural component of one concept or role must have an analogue in the other—similar to the subsumption relationships computed by KL-ONE and NIKL. This close correspondence indicates that the subsumption relationships of this logic form a useful set, and, moreover, provides a way of semantically justifying the incomplete subsumption algorithm for KL-ONE and NIKL.

---

[2]This example cannot be expressed in the logic described here because it includes a disjunction operator. However, a more complicated example which embeds this one *can* be expressed in the logic.

[3]As will be shown in the next section.

5 Computing Subsumption

Subsumption in this logic *is* weaker than subsumption in logics using the standard semantics, however this does not imply that subsumption is easy to compute here. Even the fact that subsumption is easy in the logic of [Patel-Schneider, 1986] is no assurance that it will be easy here. The addition of number restrictions is a major extension and, as Levesque and Brachman have shown [1987], even small changes in the expressive power of a formal system can result in large changes in the computational tractability of its operations.

Fortunately, subsumption is tractable in this logic. The subsumption algorithm for the full form of the logic is too long to fit in this paper, so an indirect argument has to be used to show its tractability. This is done by converting concepts and roles to a canonical form, giving a subsumption algorithm for concepts and roles in this canonical form, and then showing how this algorithm can be converted into a tractable subsumption algorithm for concepts and roles in arbitrary form.

Concepts and roles in canonical form take the following form:

```
<concept> ::=   (and <primary>*)
<primary> ::=   <atomic concept> |
                ¬ <atomic concept> |
                (atleast <minimum> <role>) |
                ¬ (atleast <minimum> <role>)
<role> ::=      (restrict (and <atomic role>+)
                          <concept>)
<minimum> ::=   <positive integer>
```

This canonical form introduces a new operator, ¬, which is a classical negation operator defined as $t \in V[\neg \mathsf{C}](d)$ iff $f \in V[\mathsf{C}](d)$ and $f \in V[\neg \mathsf{C}](d)$ iff $t \in V[\mathsf{C}](d)$.

The conversion can be done by using the following equivalences:

1. commutativity and associativity for conjunctions of concepts and roles
2. $\mathsf{C} \equiv (\mathsf{and}\ \mathsf{C})$
3. $(\mathsf{all\ R\ (and\ C_1\ C_2)}) \equiv (\mathsf{and\ (all\ R\ C_1)\ (all\ R\ C_2)})$
4. $(\mathsf{all\ R\ E}) \equiv \neg(\mathsf{atleast\ 1\ (restrict\ R\ \neg E)})$
5. $(\mathsf{atmost\ n\ R}) \equiv \neg(\mathsf{atleast\ n{+}1\ R})$
6. $\mathsf{R} \equiv (\mathsf{restrict\ R\ (and\ )})$
7. $(\mathsf{and\ (restrict\ R_1\ C)\ R_2}) \equiv (\mathsf{restrict\ (and\ R_1\ R_2)\ C})$
8. $(\mathsf{restrict\ (restrict\ R\ C_1)\ C_2}) \equiv (\mathsf{restrict\ R\ (and\ C_1\ C_2)})$

A canonical form concept will often be viewed as a set of primaries. Similarly, the two parts of a canonical form role will often be viewed as a set of atomic roles and a set of primaries.

The conversion to canonical form does not change the extension of concepts or roles:

Theorem 1 *Let* C' *be the canonical form of the concept or role* C. *Then for any semantic structure,* $V[\mathsf{C}'] = V[\mathsf{C}]$.
Proof: By simple structural induction on C.[4]

[4]Proofs of the theorems of this paper can be found in or are very similar to proofs in [Patel-Schneider, 1988].

Once concepts and roles are in canonical form then the following characterization of subsumption is both sound and complete.

Theorem 2 *Let* C *and* C' *be canonical form concepts. Then* $\mathsf{C} \Rightarrow \mathsf{C}'$ *iff for all top-level conjuncts,* D', *in* C', *there exists a top-level conjunct,* D, *in* C *such that*

1. *if* D' *is an atomic concept or the negation of an atomic concept, then* $\mathsf{D} = \mathsf{D}'$,
2. *if* D' *is of the form* $(\mathsf{atleast\ m\ R}')$, *then* D *is of the form* $(\mathsf{atleast\ n\ R})$, *with* $n \geq m$ *and* $\mathsf{R} \Rightarrow \mathsf{R}'$, *and*
3. *if* D' *is of the form* $\neg(\mathsf{atleast\ m\ R}')$, *then* D *is of the form* $\neg(\mathsf{atleast\ n\ R})$, *with* $m \geq n$ *and* $\mathsf{R}' \Rightarrow \mathsf{R}$.

Let $\mathsf{R} = (\mathsf{restrict\ S\ C})$ *and* $\mathsf{R}' = (\mathsf{restrict\ S'\ C'})$ *be canonical form roles. Then* $\mathsf{R} \Rightarrow \mathsf{R}'$ *iff* $\mathsf{S}' \subseteq \mathsf{S}$ *and* $\mathsf{C} \Rightarrow \mathsf{C}'$.

This characterization confirms that subsumption in this logic is weak. Only "structural" subsumptions are valid, and inference rules that chain together separate pieces of a concept or role are not valid, except for those involving the conversion to canonical form. Thus subsumption is this logic is very close to the subsumption relationships computed by KL-ONE and NIKL.

Given this characterization of subsumption, it is simple to derive a subsumption algorithm that runs in time proportional to the product of the sizes of its arguments.

Theorem 3 *Subsumption for canonical form concepts and roles can be performed in time proportional to the product of the sizes of the two concepts or roles involved.*

The process of converting concepts and roles to canonical form can exponentially increase their size, and thus the tractability of subsumption on arbitrary form concepts and roles has not yet been demonstrated. Two modifications are needed to produce a tractable algorithm for subsumption. First, the conversion of concepts and roles to canonical form must be done by means of structure sharing. If this is done the "size" of the canonical form of a concept or role—not the length of its printed form but the size of the data structure—will be proportional to the size of the original concept or role, and the canonicalization can be done in linear time. Second, the subsumption algorithm has to be changed so as not to redo computations. This can be done by storing previously performed subsumption tests at the appropriate places in the canonical form of the concept or role, and querying these results when applicable. The obvious method of storing and querying the cached subsumption tests results in a subsumption algorithm that runs in time proportional to the product of the "size"s of its arguments. Thus the entire subsumption process can be done in time proportional to the product of the sizes of the two concepts or roles, resulting in

Theorem 4 *Subsumption for arbitrary concepts and roles can be performed in time proportional to the product of the sizes of the two concepts or roles involved.*

6 Summary

The extension of the four-valued semantics for terminological logics to encompass number restrictions shows that four-valued semantics can be of use in expressive terminological logics. The logic used here contains most of the constructs of the languages of KL-ONE and NIKL, and contains some useful constructs that do not occur in them. Several of the constructs in KL-ONE and NIKL, such as role-value maps and structural descriptions, that are not in this logic can be easily formulated in it, and, moreover, do not seem to provide any computational difficulties for subsumption.[5] One extension that has been investigated is the ability to specify fillers (or values) for roles [Patel-Schneider, 1988]. This extension retains the desirable computational properties of subsumption.

The four-valued semantics used here is a reasonable semantics, especially when considering systems with limited reasoning capabilities. Subsumption in this semantics is easy to compute, at least for the language given here. The valid subsumption relationships form an interesting set—one that includes the easy subsumptions and leaves out the less obvious ones. This set corresponds closely to the set of subsumption relationships computed in KL-ONE and NIKL, lending a degree of credence to that set.

This extension is not without problems, however. It shares the problems of the semantics in [Patel-Schneider, 1986]—being not as intuitive as the standard two-valued semantics and sanctioning a very limited set of subsumption relationships. These seem to be unavoidable problems if a uniform, simple semantics with a fast subsumption algorithm is required. The extension also weakens the relationship between the elements of the domain set and objects in the world.

This extension shows that there are even more tradeoffs in the relationship between expressive power, deductive power, and computational tractability in terminological logics. It justifies a limited set of subsumption relationships for an expressively powerful terminological logic that is easy to compute and, moreover, captures an interesting subset of the standard subsumption relationships. This is not a total solution, because no total solutions are possible (unless $P = NP$), but it does demonstrate that it is possible to to alleviate the computational problems of expressively powerful terminological logics by weakening deduction in a principled manner.

Acknowledgments

Hector Levesque and Ron Brachman, through their investigation of the complexity of computing subsumption in standard semantics for frame-based description languages, provided the impetus for this research. Anonymous reviewers pointed out a problem with a previous version of this paper.

References

[Belnap, 1977] Nuel D. Belnap, Jr. A useful four-valued logic. In G. Epstein and J. M. Dunn, editors, *Modern Uses of Multiple-Valued Logic*, pages 8–37. Reidel, 1977.

[Brachman and Schmolze, 1985] Ronald J. Brachman and James G. Schmolze. An overview of the KL-ONE knowledge representation system. *Cognitive Science*, 9(2):171–216, April–June 1985.

[5]Of course, statements like this are notoriously prone to error. Adding extra expressive power to the logic via adding more constructs is an area for further research.

[Brachman *et al.*, 1983] Ronald J. Brachman, Richard E. Fikes, and Hector J. Levesque. KRYPTON: A functional approach to knowledge representation. *IEEE Computer*, 16(10):67–73, October 1983.

[Brachman *et al.*, 1985] Ronald J. Brachman, Victoria Pigman Gilbert, and Hector J. Levesque. An essential hybrid reasoning system: Knowledge and symbol level accounts of KRYPTON. In IJCAI-85 [1985], pages 532–539.

[Frisch, 1985] Alan M. Frisch. Using model theory to specify AI programs. In IJCAI-85 [1985], pages 148–154.

[IJCAI-85, 1985] International Joint Committee on Artificial Intelligence. *Proceedings IJCAI-85*, August 1985.

[Levesque and Brachman, 1987] Hector J. Levesque and Ronald J. Brachman. Expressiveness and tractability in knowledge representation and reasoning. *Computational Intelligence*, 3(2):78–93, May 1987.

[Levesque, 1984] Hector J. Levesque. A logic of implicit and explicit belief. In *Proceedings AAAI-84*, pages 198–202, Austin, Texas, August 1984. American Association for Artificial Intelligence. A revised version available as FLAIR Technical Report Number 32, Fairchild Laboratory for Artificial Intelligence Research, August 1984.

[Moser, 1983] M. G. Moser. An overview of NIKL, the new implementation of KL-ONE. Technical Report 5421, BBN Laboratories, 1983. Part of a collection entitled "Research in Knowledge Representation for Natural Language Understanding—Annual Report, 1 September 1982–31 August 1983".

[Patel-Schneider, 1984] Peter F. Patel-Schneider. Small can be beautiful in knowledge representation. In *Proceedings IEEE Workshop on Principles of Knowledge-Based Systems*, pages 11–16. IEEE Computer Society, December 1984. A revised and extended version is available as AI Technical Report Number 37, Schlumberger Palo Alto Research, October 1984.

[Patel-Schneider, 1986] Peter F. Patel-Schneider. A four-valued semantics for frame-based description languages. In *Proceedings AAAI-86*, pages 344–348, Philadelphia, Pennslyvania, August 1986. American Association for Artificial Intelligence.

[Patel-Schneider, 1988] Peter F. Patel-Schneider. A four-valued semantics for terminological logics. Technical Report 72, Schlumberger Palo Alto Research, February 1988.

[Patel-Schneider, to appear] Peter F. Patel-Schneider. A decidable first-order logic for knowledge representation. *Journal of Automated Reasoning*, to appear.