Circumscribing Equality*

Peter K. Rathmann
Computer Science Dept.
Stanford University
Stanford, CA 94305

Marianne Winslett
Computer Science Dept.
University of Illinois
Urbana, IL 61801

Abstract

One important facet of common-sense reasoning is the ability to draw default conclusions about the state of the world. Such an ability enables one to assume, for example, that a given bird flies in the absence of information to the contrary. One drawback of the circumscriptive approach to common-sense reasoning has been its inability to produce default conclusions about equality. For example, generally one cannot tentatively conclude that President(USA) ≠ Fido using circumscription. In this paper we give a second-order axiom and model theory for circumscribing equality, and prove that they are equivalent.

1 Introduction

Circumscription [Bossu 85, Etherington 85, 88, Lifschitz 87, McCarthy 80, 86, Perlis 87, Shoham 87] seeks to solve the problem of common-sense reasoning by "preferring" certain models of a theory *T* to others. More precisely, circumscription picks out those models of a theory that are minimal with respect to some partial order on models [Shoham 87]. Thus, a circumscriptive partial order encodes our intuitive sense of which of the logically plausible alternative models of *T* are the most "normal" and reasonable.

Absent evidence to the contrary, common sense suggests that different names denote different things. For example, this paper is not the same thing as your car, your cat, or your phone number. In a circumscriptive setting, this common-sense principle translates into a preference for those models of the world in which as many terms as possible are unequal, i.e., those in which equality is minimized. Minimizing equality, though, requires the inclusion of preferences between models having different universes - and in the usual types of circumscription two models must have the same universe in order for them to be comparable. Thus, it is not surprising that the simple forms of circumscribing a theory *T* will not produce any facts about equality of terms that were not already deducible from *T* [Etherington 88], This problem with equality has been a drawback of circumscription, especially since other common-sense reasoning approaches such as default logic and autoepistemic logic deal satisfactorily with equality. Sections 2 and 3 of this paper propose a new kind of circumscription that remedies this deficiency. Section 4 shows how to strengthen equality circumscription so that it also circumscribes the domain.

*This work was supported by DARPA under grant N39-84-C-211 (KBMS Project, Gio Wiederhold, principal investigator).

In the past, many authors have sidestepped the equality problem by adopting a Herbrand universe assumption. However, this assumption mandates that different terms *must* denote different objects. For example, in a Herbrand universe, President(USA) cannot be Bush. We desire a mechanism that will allow us to conclude that President(USA) = Bush if this is entailed by T.

McCarthy [86] proposed two methods of dealing with the equality problem. The first was to state "unique name" axioms, formulas such as gutter ≠ cat, explicitly in *T.* In general, however, not all facts of this nature will be known in advance (who *is* the president?). In addition, there may be an infinite number of such facts. The second method proposed by McCarthy was to introduce multiple types of equality, with one variety of equality not enjoying the full principle of substitutivity. This approach is rather unwieldy. Rather than working around the problem, it would be preferable simply to circumscribe equality.

The inspiration for our approach to circumscribing equality comes from the category-theoretic concept of *initial model semantics* [Goguen 85]. Under this approach, the preference ordering on models is determined not by a simple model/submodel relationship, but rather by the existence (or lack thereof) of homomorphisms between one model and another. This results in semantics where all predicates, including the equality predicate, are minimized. We believe that homomorphisms will, in many cases, prove more useful than the model/submodel relationship in determining a partial order on models.

This paper will not adopt a category-theoretic presentation because that formalism is more complex than is needed in this paper. Instead, we directly present the idea of homomorphisms between models in Section 2.

We do not intend equality circumscription as a replacement for the traditional forms of circumscription, or as the ultimate method of non-monotonic reasoning. Rather, we see it as yet another tool in the catalog of non-monotonic reasoning techniques, one which in some cases will prove more appropriate than other tools. A major challenge for artificial intelligence research in this area will be the development of systematic means of choosing *policies* for application of these tools. Simply encouraging the use of equality circumscription does not give much more insight into the solution of a problem than encouraging a would-be author use an editor or suggesting that an AI system architect use an ATMS.

2 A Model Theory for Equality Circumscription

The model theory for equality circumscription is based on homomorphisms between models. We first define the concepts of signature and model, and then define homomorphisms between models. A signature is a description of the terms available to create sentences in a logical theory. It lists the function symbols, predicate symbols, and their arities. Constants are not considered separately, but instead are treated as 0-ary functions. Because this paper only covers finite theories, signatures will also be finite.

A model over a signature consists of a universe *U* and extensions for all the function and predicate symbols. The extension of a function symbol is an actual function (of proper arity) on *U,* while the extension of a predicate symbol is a relation (of proper arity) on *U*.

Let *M'* and *M* be models of a theory T, and let U^r and *U* be the universes of *M'* and *M,* respectively. Then a homomorphism from *M'* to *M* is a function *h* : *V —> U* from the universe of *M'* to the universe of M, satisfying the following conditions:

1. The mapping *h* is a homomorphism on functions. By this we mean that if / is an n-ary function symbol from T's signature, and it has extensions f and f in *M'* and *M* respectively, we require that

$$\forall a_1 \ldots a_n \in U' \quad h(f'(a_1,\ldots,a_n)) = f(h(a_1),\ldots,h(a_n))$$

2. The truth of predicates is preserved. If *P* is an n-ary predicate symbol from T's signature with extensions P' and P in *M'* and *M*, then

$$\forall a_1 \ldots a_n \in U' \quad P'(a_1,\ldots,a_n) \rightarrow P(h(a_1),\ldots,h(a_n)).$$

It will sometimes be convenient to treat n-tuples such as $a_1,\ldots,a_n$ as vectors a, so that we can write *h(a)* as shorthand for $n(a_1),\ldots,h(a_n)$, and so on.

If there is a homomorphism from model *M'* to M, we write *M' —► M.*

The existence of a homomorphism *M' —* M* can represent a number of different relationships. It may be that the true atoms of *M'* are a proper subset of their counterparts in *M*. This is the model/submodel relationship used in ordinary circumscription, and it means that M^1 is preferable to M, from the viewpoint of circumscription. Alternatively, the morphism could map two or more separate elements of Af"s universe onto a single element in *M*. Since we want to minimize equality, intuitively *M'* should be preferable to M, because *M* equates atoms that M^l does not. Thus we might be tempted to create a partial order, based on the idea that M^f is preferred to *M* iff M^1 —> M. Unfortunately, this does not work, since it is possible to have both *M' —> M* and *M —> M'.* Instead, we consider such models with two-way morphisms to be incomparable, and arrive at the following definition:

Definitions. A model *M'* is *preferred to M* iff M^f —► *M* but not *M —* M'.* A model *M* is *preferred* if no other model is preferred to *M.*

(Section 4 describes a different, stronger semantics that also provides a form of domain circumscription.)

Theories typically have multiple preferred models. For example, the theory *P(a)* V *P(b)* will not have a single preferred model. It says explicitly that we don't know what model represents the true state of the world.

Example 1 *Fido and the President.*

Let *T* be the theory containing the two formulas Dog(Fido) and Happy(President(USA)). Ordinary circumscription of this theory would not allow us to conclude that President(USA) ≠ Fido, or that Fido is not happy. The model theory for equality circumscription, on the other hand, produces both these conclusions.

Our signature contains two 0-ary function symbols (Fido, USA), one 1-ary function symbol (President), and two 1-ary predicate symbols (Dog, Happy). A model over this signature must include a universe and interpretations for the constants, functions, and predicates. Two particular models of the theory are shown in figure 1.

In figure 1, *h* is a homomorphism A —► B, a mapping from A to B that preserves functions and the truth of predicates. The dashed arcs represent the mappings of *h* from model A to model B. For example, the constant Fido denotes the universe object $Fido_A$ in model A, but denotes $Fido_B$ in model B. Consequently, *h* maps $Fido_A$ onto $Fido_B$. Note that in model B, President(USA)=Fido, and Fido is happy. There can be no homomorphism from model B to model A, since such a morphism could only map $Fido_B$ onto one object, and there is no object which is both happy and a dog in model A. Hence, model A is preferred to model B, assuming an appropriate completion of the remainder of the information in *A.*

Neither A nor B gives a complete description of a model of the theory: there are many other syntactically correct terms in the language, such as President(President(USA)), and any model must have extensions for them. Equality circumscription will prefer models in which all of these expressions denote unique elements in the universe, i.e., in which the unique name axioms hold. For a model *M* to be preferred, the predicates Dog and Happy must be false on all ground terms, except for terms Fido and President(USA), respectively.

Example 2 *Inheritance.*

Suppose we wish to encode an inheritance hierarchy, in particular the common-sense principle that birds generally fly. This can be expressed [McCarthy 86] by a formula which says that all normal birds fly (Vx. *b(x)* A -iafr(x) —► *f(x)).* Exceptions are allowed - ostriches do not fly (Vx. o(x) —► -»/(x))). If we know that Tweety is

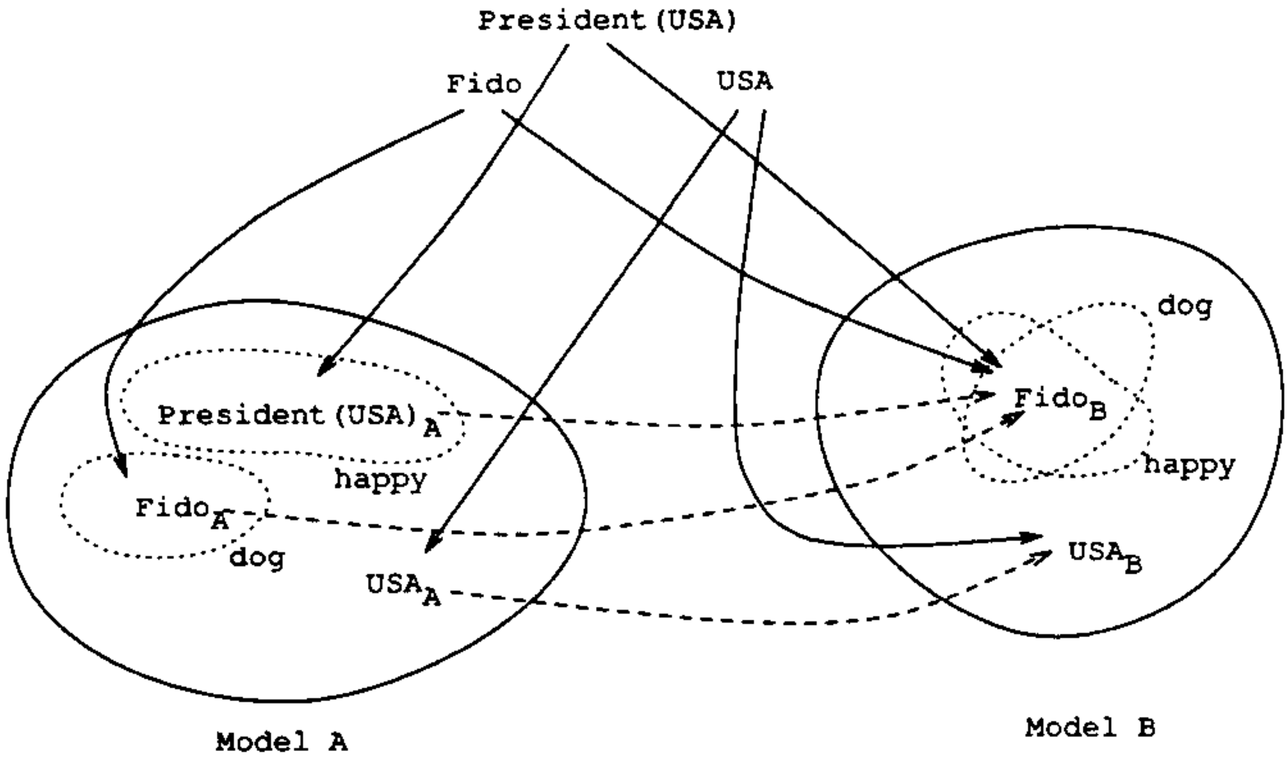

Figure 1: Example of a model homomorphism.

a bird (b(T)), both the appropriate ordinary and equality circumscription[1] of these formulas will conclude that Tweety flies. The difference between the two approaches appears when we add another sentence to the theory. If we know of a particular ostrich, Blutto (o(B)), ordinary circumscription will no longer conclude that Tweety flies, because Tweety and Blutto might designate the same individual. Equality circumscription, which prefers models in which Tweety and Blutto are distinct, will still conclude that Tweety flies.

We now show that the model theory for equality circumscription is in fact based on a partial order.[2]

Proposition 1 *The preference relation on the models of a theory is a partial order (transitive, asymmetric, and irreflexive).*

(Proofs of the theorems in this paper appear in [Rathmann 88].)

Theorem 1 shows that when testing to see whether a model is preferred, it suffices to consider models whose universes are generated by the constants and functions of the underlying signature:

Theorem 1 *Let A be a non-preferred model of a universal theory T. Then there exists a preferred model B of T such ihat B is preferred to A and B's universe is exactly the interpretation of the ground terms of the signature of T.*

[1]This requires an extension to equality circumscription to allow some predicates to vary and hold others fixed.

[2]Authors differ on whether a partial order should be reflexive or irreflexive. In this presentation we have chosen irreflexive, so that the order behaves like $<$, rather than $\leq$.

Theorem 1 has an important corollary:[3]

Corollary 1 *A universal theory T has preferred models iff T is consistent.*

Theorem 2 shows that the unique name axioms follow from the equality circumscription of T, whenever possible. Since we minimize equality at the same priority as all other predicates, the unique name axioms do not follow when that would increase the extension of other predicates, and Theorem 2 reflects this. For example, if T is $(a \neq b) \rightarrow (P(a) \vee P(6))$, then the unique name axioms will not hold in all preferred models of T, as when those axioms are true, P must have a non-empty extension.

Theorem 2 *If the equality predicate does not appear positively (i.e., governed by an even number of negation signs) in the conjunctive normal form of a theory T, then $a \neq b$ is true in all preferred models of T, for all ground terms a and b.*

Of course, when the equality predicate does appear positively in T, then "as many as possible" of the unique name axioms will still hold in the preferred models of *T*. We do not yet know how to state this property formally.

Equality circumscription also avoids some of the overly strong conclusions that plague traditional circumscription. For example, if *T* is the theory $\{P(f(a))\}$, under ordinary circumscription $\forall x (P(x) \leftrightarrow x = f(a))$ is true in all preferred models of *T*. In contrast, under equality circumscription, we reach new conclusions about the truth

[3]In new work, we have shown that this corollary (though not Theorem 1) holds for all first-order theories, so that every consistent theory has a preferred model under equality circumscription.

or falsity of ground formulas, but not in general about non-ground formulas. Under equality circumscription, in all preferred models of T the formula schema $\neg P(t)$ holds, for *t any* ground term except /(a). The formula Vx(P(x) <-> x = /(a)) does *not* follow, however. One can think of this property as saying that we draw conclusions about known and named individuals, rather than about all the individuals in the universe.

Thus equality circumscription in its basic form admits the possible existence of additional unknown individuals who are similar to known individuals. For example, if we know that one penguin exists (Penguin(Opus)), then equality circumscription permits models containing any number of unnamed penguins. In contrast, ordinary circumscription would conclude that Opus was the only penguin extant. If in addition we know that Opus is a bachelor (Bachelor(Opus)), equality circumscription will permit any number of unnamed bachelor and/or penguin elements to appear in models. If we also know that Bill is a cat (Cat(Bill)), there may also be unnamed cats in the world under equality circumscription; however, bachelor cats will *not* be permitted, because under equality circumscription, any unnamed element must enjoy properties that are a subset of those of some particular named element. The existence and nature of these unnamed individuals follow from the homomorphisms used to pick out preferred models; the model having a minimal universe (just Opus and Bill) is not preferred to a model containing "extra" elements because the extra elements "fold onto" the minimal universe: one can construct a homomorphism from the larger to the smaller universe, mapping all cats to Bill and all bachelors and penguins to Opus. One may think of equality circumscription as allowing the existence of unnamed individuals whose properties are analogous, in a homomorphic sense, to those of known individuals.

Example 3 *Lotteries.*

Let T contain formulas stating that the universe contains 100,000 distinct anonymous individuals. (Do this by stating that there exist x_1 through X_{100000} such that each $x_i \neq X_j$ for distinct i, *j,* and for all y, *y* is equal to one of the X_iS.) Include one constant "Harry" in the language, designating our friend Harry. Add the fact that one of these people won the lottery (3x *winner(x)),* and consider the question, "Did Harry win the lottery?" Ordinary circumscription admits models where Harry did win; common sense supposedly requires the answer that Harry did not. Equality circumscription says that Harry did not win, but that some other, unnamed individual did win.[4]

[4] If individuals arc uniquely identifiable by name, then equality circumscription will answer "maybe" to the question; if they are identifiable by some other means, as for example by an ordering on social security number, then if that ordering predicate is considered *relevant* to the question, the answer will again be "maybe." If the predicate is considered irrelevant, as in the case of social security numbers, then it should be allowed to vary when circumscribing, and the answer will be "no." See the remarks in Section 1 on choice of circumscriptive policies.

As Theorem 1 only guarantees consistency in the case of universal theories, in the remainder of this paper we restrict our attention to universal theories.[5]

3 An Axiom for Equality Circumscription

The model theory for equality circumscription, as presented in the previous section, used a single-sorted first-order logic. Like ordinary circumscription, the axiom for equality circumscription uses second-order logic. In addition, for equality circumscription we must introduce a *second sort.* To show the need for this sort, let us first present a straw-man axiom using only a single sort.

By analogy with ordinary circumscription, as a first cut at a second-order axiom for equality circumscription, one might devise the following formula:

$$T(S) \wedge \forall S' \neg[(S' < S) \wedge T(S')],$$

where $S^f < S$ is shorthand for a formula specifying that there is a homomorphism from the "model" defined by S' to that defined by S, but not in the reverse direction. (Readers who are not familiar with the notation used in circumscription axioms need not be alarmed; all will be explained below.) This would not capture the semantics of equality circumscription, however, as a non-preferred model of T will satisfy this axiom *if its universe is smaller than that of any model which is preferred to it.*[6]

The purpose of the second sort (hereafter called the *test sort)* is to allow comparison of a model M with models M' having universes larger than that of M. This is important, since when minimizing equality, it is quite possible that such an M' may have fewer equality atoms true than in M, and hence be preferred over M.

The second sort is used for (intuitively speaking) testing to see whether a particular model is preferred, by constructing a second model (in the test sort) to compare to the original one. The preferred models of the original theory T will be models of the second-order theory Circ(T), *restricted to the original sort.* In what follows, it will be convenient to assume that theory T contains no free variables, and that the signature of T contains at least one constant.

We assume that the elements of the test sort are disjoint from those of the original sort. In light of Theorem 1, it suffices, when testing to see if a model is preferred, to consider only other models with universes as large as the set of (possibly skolemized) ground terms of the language. We therefore assume that the test sort is of that size.[7]

[5] As mentioned in footnote 3, new results show this restriction to be unnecessary.

[6] Two universes are *isomorphic* if there exists a one-to-one, onto mapping between them. Universe U_1 is *larger* than U_2 (and U_2 is smaller than U_1) if U_2 is isomorphic to a proper subset of U_1 and is not isomorphic to U_1.

[7] Both of these conditions on the test sort can be expressed as statements in first order logic, and enforced by adding these statements to the axiom for equality circumscription. We will not add these conditions explicitly in this paper.

Let *F* be the set of all function symbols in the signature of T, and let *P* be the corresponding set of predicate symbols. Recall that constants are actually 0-ary functions, so the constants of *T* appear in *F*. For every function symbol / in *F,* let /' be a function variable that matches / in arity and type. Extending this notation, let 5 be a tuple of all the function symbols and predicates in *F* and P, except equality, and let S' be a tuple of function and predicate variables matching those of *S* in arity and type. The members of *S* have arguments and results only in the original sort, and the members of S' are similarly confined to the test sort. We will change our notation for theories slightly to include the functions and predicates as a parameter of the theory. Thus, *T(S)* denotes *T* with its usual functions and predicates, and *T(S')* denotes the result of replacing in *T* all the function and predicate symbols of S with the appropriate variables of S'. The quantifiers of *T(S)* range only over the original sort, and the quantifiers of T(S') range only over the test sort.

The second order axiom for circumscribing equality looks very similar to that for ordinary circumscription. The circumscription of T, denoted Circ(T), is given by

$$T(S) \wedge \forall S' \forall U \neg[(S' < S) \wedge T_U(S')]. \tag{1}$$

U is a unary predicate variable taking arguments of the second sort. Intuitively, the predicate variables *U* and S' pick out a universe (called the *test* universe) and a set of predicate valuations, respectively, over the test sort. In essence *U* and S" construct a new "test model" of *T* over the test sort, to see whether the test model is preferred to the model given by the original sort and the predicate valuations of *S* (the *original* model). The conjunct T_U(5'), defined precisely below, guarantees that the test model actually satisfies the formulas of T. The biggest difference from ordinary circumscription comes in the construct *S'* < S, which is shorthand for

$$\exists h \ \mathbf{hom}(h) \wedge \neg\exists g \ \mathbf{reverseHom}(g). \tag{2}$$

The unary function variable *h* takes an argument of the test sort and produces a result in the original sort; the reverse is true for *g.* Formula (2) ensures that *h* is a homomorphism from the test model to the original model, and that there is no homomorphism in the reverse direction. As described in Section 2, this is the condition for the test model to be preferable to the original model. More precisely, hom(h) is shorthand for

$$\forall \overline{x} \left[U(\overline{x}) \rightarrow \left(\bigwedge_{f_i \in F} (h(f_i'(\overline{x})) = f_i(h(\overline{x}))) \wedge \bigwedge_{P_i \in P} (P_i'(\overline{x}) \rightarrow P_i(h(\overline{x}))) \right) \right]. \tag{3}$$

The quantifier of (3) ranges over the test sort. Formula (3) is the analog, in second-order logic, of the definition of homomorphism given in Section 2. (To see this, recall that equality is one of the predicates in *P.)* The definition of reverseHom(g) is symmetric to that for hom(h). In this case the universal quantifier ranges over the original sort.

The definitions of hom(h) and reverseHom(h) do some sneaky things with the arity of their arguments. First, $\overline{x}$ is an n-tuple of arguments, $x_1, \cdots, x_n$, where n is the largest arity of the members of S, but at least 1. $U(\overline{x})$ is true iff $U(x_1) \wedge \cdots \wedge U(x_n)$ holds; similarly, $f(\overline{x})$ is the n-tuple $f(x_1), \ldots, f(x_n)$. If a function or predicate s has arity i, where $i < n$, then the last $n - i$ arguments of $\overline{x}$ are to be ignored in computing $s(\overline{x})$.

Finally, $T_U(S')$ is formed from $T(S')$ by (1) ensuring that the test universe is nonempty, (2) restricting the range of quantifiers to the test universe, and (3) ensuring that the results of functions are in the test universe, whenever their arguments are. Step (1) is accomplished by adding $\exists x U(x)$ to $T_U(S')$. Step (2) is accomplished by first putting the formulas of $T(S')$ into prenex form, and then successively replacing all occurrences of subformulas of the form $\forall x \alpha$ by $\forall x(U(x) \rightarrow \alpha)$, and subformulas of the form $\exists x \alpha$ by $\exists x(U(x) \wedge \alpha)$. For example, if T is $\exists x$ sane$(x) \rightarrow \forall y$[President$(y) \neq$ Fido], then $T_U(S')$ contains the formula $\exists x[U(x) \wedge \forall y[U(y) \rightarrow (\text{sane}'(x) \rightarrow (\text{President}'(y) \neq \text{Fido}'))]]$. To accomplish step (3), one must add to $T_U(S')$ the formula

$$\forall \overline{x}[U(\overline{x}) \rightarrow U(f'(\overline{x}))],$$

for each function f in F, where $\overline{x}$ and U are as described in the previous paragraph. For example, the President function will require the addition of $\forall x[U(x) \rightarrow U(\text{President}'(x))]$ to $T_U(S')$. For a 0-ary function f such as Fido, the appropriate formula to add is simply $U(f')$, as in $U(\text{Fido}')$.

Theorem 3 shows that the model-theoretic and axiomatic formulations of equality circumscription are equivalent.

Theorem 3 *The preferred models of a universal theory T are the original models of Circ(T).*

4 Equality Circumscription with Domain Closure

In many applications of circumscription, we expect a slightly stronger form of equality circumscription to be desirable: equality circumscription with domain closure. Domain closure corresponds to the common-sense reasoning assumption that only those things that must exist do exist. Equality circumscription with domain closure (ECDC) differs from plain equality circumscription in that preferred models must have universes that are an interpretation of the ground terms of the signature of T (a *minimal* universe) under ECDC. In other words, preferred models must have universes that are minimal in the sense that every element in them is designated by some ground term. Informally, we are eliminating all models containing unnamed elements. In applications where one knows a priori that all the relevant individuals are known and named, i.e., are designated by ground terms, the use of ECDC rather than ordinary equality circumscription is indicated. The partial order on models that corresponds to ECDC is obtained by adding a requirement that all preferred models have minimal universes:

Definition. A model M' is preferred to model M under ECDC iff

1. M' — M and not M — M'; or
2. M' is the restriction of M to a smaller and minimal universe.

Proposition 2 shows that the new preference relation is a partial order on models.

Proposition 2 The preference relation for ECDC is a partial order.

Preferred models under ECDC are also preferred models under equality circumscription:

Proposition 3 If M is a preferred model of T under ECDC, then M is also a preferred model of T under equality circumscription.

More precisely, for a universal theory the preferred models under ECDC are obtained by restricting the preferred models of equality circumscription:

Proposition 4 Given a universal theory T, its preferred models under ECDC are the restriction to minimal universes of its preferred models under equality circumscription.

Note that Proposition 4 guarantees that a universal theory will always have preferred models under ECDC.

Example 4 Fido and the president, continued.

Continuing the example begun earlier, under ECDC a preferred model must have in addition a minimal universe. Therefore T now has only one preferred model, up to isomorphism. In this model, Fido and USA map to separate elements of the universe; and the President function is injective, i.e., no two different applications of it give the same result. In addition, there is exactly one dog, the interpretation of Fido; and there is one happy thing, the interpretation of the President of the USA.

What is the relationship between ECDC and ordinary global circumscription? All statements, both quantified and ground, that are true under ordinary circumscription will also follow under ECDC, as Theorem 4 shows.

Theorem 4 If a model is preferred under ECDC, then it is preferred under ordinary global circumscription.

A second-order formulation of ECDC is presented in [Rathmann 88]. The formulation is very similar to that for plain equality circumscription, and has the same number of second-order quantifier alternations. The equivalence of the second-order and model-theoretic formulations is also proven there.

5 Extensions

We have also devised varieties of equality circumscription that allow one to circumscribe only certain predicates, to allow selected predicates to vary, and to circumscribe with priorities. These variants change the definition of ordinary equality circumscription in the same manner as these variants change the definition of global circumscription.

6 Conclusions

In this paper we have shown the advantages of the use of a preference relation between models that is based on homomorphisms between models, rather than the model/submodel relationship usual in circumscription. In particular, this new preference relation allows one to prefer models of a theory in which different terms denote different objects, to the maximum extent possible. This principle of common-sense reasoning cannot be implemented using previously proposed forms of circumscription.

7 Acknowledgments

We would like to thank Moshe Vardi, Joseph Goguen, Phokion Kolaitis, Vladimir Lifschitz, Matt Ginsberg, and Ken Forbus for their helpful advice. Joseph Goguen invented the concept of "initial model" and provided guidance on the role of category theory in semantics. Phokion Kolaitis helped us to understand our model theory when we found ourselves confused. Moshe Vardi, Vladimir Lifschitz and Ken Forbus generously volunteered to read our manuscripts, and gave valuable comments. Discussions with Matt Ginsberg inspired example 3.

8 References

[Bossu 85] G. Bossu and P. Siegel, "Saturation, Nonmonotonic Reasoning and the Closed-World Assumption," Artificial Intelligence, 25, 1985, pp. 13-63.

[Etherington 85] D. W. Etherington, R. E. Mercer, and R. Reiter, "On the Adequacy of Predicate Circumscription for Closed-World Reasoning," Computational Intelligence, 1, 1985, pp. 11-15.

[Etherington 88] D. W. Etherington, Reasoning with Incomplete Information, Morgan Kaufmann, 1988.

[Goguen 85] J. A. Goguen and R. M. Burstall, "Institutions: Abstract Model Theory for Computer Science," Technical Report CSLI-85-30, Center for the Study of Language and Information, Stanford University, August, 1985.

[Lifschitz 87] V. Lifschitz, "Pointwise Circumscription," in Readings in Nonmonotonic Reasoning, M. L. Ginsberg, ed., Morgan Kaufmann, 1987.

[McCarthy 80] J. McCarthy, "Circumscription—A Form of Non-Monotonic Reasoning," Artificial Intelligence, 13, 1980, pp. 27-39 and 171-172.

[McCarthy 86] "Applications of Circumscription to Formalizing Common Sense Knowledge," Artificial Intelligence, 28, 1986.

[Perlis 87] D. Perlis, "Circumscribing with Sets," Artificial Intelligence, 31, February 1987.

[Rathmann 88] P. K. Rathmann and M. Winslett, "Circumscribing Equality," TR UIUCDCS-R-88-1473, Univ. of Illinois, Dec. 1988.

[Shoham 87] Y. Shoham, "A Semantical Approach to Nonmonotonic Logics," Proceedings of the 10th IJCAI, Milan, 1987.