On The Adequacy of Predicate Circumscription For Closed-World Reasoning

David W. Etherington[1]
Robert E. Mercer[2]
Raymond Reiter[2,3]
Department of Computer Science
University of British Columbia

Technical Report 84-5
September 3, 1984

Abstract

We focus on McCarthy's method of predicate circumscription in order to establish various results about its consistency, and about its ability to conjecture new information. A basic result is that predicate circumscription cannot account for the standard kinds of default reasoning. Another is that predicate circumscription yields no new information about the equality predicate. This has important consequences for the unique names and domain closure assumptions.

[1] This research was supported in part by an I.W.Killam Memorial Predoctoral Fellowship and an NSERC Postgraduate Scholarship.

[2] This research was supported in part by the Natural Sciences and Engineering Research Council of Canada, grant A7462, and the U.S. National Science Foundation, grant MCS-8203954.

[3] Fellow of the Canadian Institute for Advanced Research.

Table of Contents

1. Introduction

There are many settings in artificial intelligence and in the theory of databases where we must assume that the available information consists of all and only the relevant facts. A variety of different intuitions appear to underly this so-called *Closed-World Assumption.* These differing intuitions have lead to different formalizations of this important notion:

1) Negation as failure to derive [Reiter 1978a,1978b] – If a ground atomic formula cannot be proved using the given information as premises, then assume the formula's negation.

2) Predicate Completion [Clark 78, Kowalski 78] – When the given information about a predicate consists of a set of sufficient conditions on that predicate, assume that these conditions are also necessary.

3) Domain Circumscription [McCarthy 1977, 1980] – Assume that the individuals required by the given information are all there are.

4) Predicate Circumscription [McCarthy 1980] – Assume those facts which are true in all models of the given theory having minimal extensions for certain predicates.

The relationships among these different formalisms are not completely understood, but some results are known. Domain circumscription can be reduced to predicate circumscription [McCarthy 1980]. Under certain circumstances, predicate completion is a special case of predicate circumscription [Reiter 1982]. Shepherdson [1984] provides a careful analysis of the relationships between predicate completion and negation as failure to derive.

In this paper we focus on predicate circumscription, as presented in [McCarthy 1980]. Our objective is to establish various results concerning the consistency of this formalism, and to describe some limitations of its ability to conjecture new information. One such limitation is that predicate circumscription cannot account for the standard kinds of default reasoning. Another limitation relates to equality; predicate circumscription yields no new information about the equality predicate for a large class of first-order theories. This has important consequences for the so-called "unique names assumption" and "domain closure assumption".

Recently, McCarthy has proposed a generalization of predicate circumscription which minimizes arbitrary first-order expressions rather than simple predicates. Certain predicates are allowed to act as variables during this minimization

[McCarthy 1984]. Some of the limitations of predicate circumscription which we describe in this paper do not apply to this generalized form of circumscription, although equality appears to remain problematic.

2. Predicate Circumscription: Formal Preliminaries

The semantic intuition underlying predicate circumscription is that closed-world reasoning about one or more predicates of a theory corresponds to truth in all models of the theory which are minimal in those predicates. Specifically, let $T(P_1,...,P_n)$ be a first-order sentence, some (but not necessarily all) of whose predicates are $\mathbf{P} = \{P_1,...,P_n\}$. A model M of the sentence T is a **P**-*submodel* of a model M' of T iff the extension of each P_i in M is a subset of its extension in M', and M and M' are otherwise identical. M is a **P**-*minimal model of* T iff every **P**-submodel of M is identical to M.

For finite theories, $T(P_1,...,P_n)$, McCarthy [1980] proposes realizing predicate circumscription syntactically by adding the following axiom schema to T:

$$\left[T(\Phi_1,...,\Phi_n) \wedge \bigwedge_{i=1}^{n} [\forall \bar{x}.\ (\Phi_i \bar{x} \supset P_i \bar{x})]\right] \supset \bigwedge_{i=1}^{n} [\forall \bar{x}.\ (P_i \bar{x} \supset \Phi_i \bar{x})].$$

Here $\Phi_1,...,\Phi_n$ are predicate variables with the same arities as $P_1,...,P_n$, respectively, and $T(\Phi_1,...,\Phi_n)$ is the result of replacing every occurrence of $P_1,...,P_n$ in T by $\Phi_1,...,\Phi_n$, respectively. The above schema is called *the (joint) circumscription schema of* $P_1,...,P_n$ *in* T. Let $CLOSURE_{\mathbf{P}}(T)$ – *the closure of* T *with respect to* $\mathbf{P} = \{P_1,...,P_n\}$ – denote the theory consisting of T together with the above axiom schema. McCarthy formally identifies reasoning about T under the closed-world assumption with respect to the predicates **P** with first-order deductions from the theory $CLOSURE_{\mathbf{P}}(T)$.

McCarthy [1980] shows that any instance of the schema resulting from circumscribing a single predicate P in a sentence $T(P)$ is true in all $\{P\}$-minimal models of T. This generalizes directly to the joint circumscription of multiple predicates; we omit the proof of this. We use this generalization extensively in the proofs of the results of this paper. Because predicate circumscription is applicable only to finitely axiomatizable theories, we will restrict our attention to such theories.

3. On the Consistency of Predicate Circumscription

The minimal model semantics of predicate circumscription suggests that certain consistent first-order theories lacking minimal models may have inconsistent closures. Indeed, this can happen, as we now show.

Example 3.1 – An inconsistent circumscription

The following consistent theory has no $\{P\}$-minimal models:

$$\left\{\begin{array}{l} \exists x.\ Px \wedge \forall y.\ [Py \supset x \neq succ(y)] \\ \forall x.\ Px \supset Psucc(x) \\ \forall xy.\ succ(x) = succ(y) \supset x = y \end{array}\right\}$$

Circumscribing P in this theory, and letting Φx be $[Px \wedge \exists y. x = succ(y) \wedge Py]$ yields $\forall x.\ Px \supset \exists y.\ [Py \wedge x = succ(y)]$ which contradicts the first axiom. ∎

In view of this example, it is natural to seek classes of first-order theories for which predicate circumscription does not introduce inconsistencies. The well-founded theories form such a class. We say that a first-order theory is *well-founded* iff each of its models has a **P**-minimal submodel for every finite set of predicates **P**. Any consistent well-founded theory obviously has at least one **P**-minimal model. Since every instance of the circumscription schema of **P** in a theory T is true in all **P**-minimal models of T, we have:

Theorem 3.1

If T is a consistent well-founded theory, then $CLOSURE_{\mathbf{P}}(T)$ is consistent for any set **P** of predicates of T, i.e. predicate circumscription preserves consistency for well-founded theories. ∎

Which theories are well-founded? We know of no complete syntactic characterization, but a partial answer comes from a generalization of a result on universal theories due to Bossu and Siegel [1984]. A first-order theory is *universal* iff the prenex normal form of each of its formulae contains no existential quantifiers.

Theorem 3.2

Universal theories are well-founded. ∎

In view of Theorem 3.1, we know that predicate circumscription preserves consistency for universal theories:

Corollary 3.3

If T is a consistent universal theory, then $CLOSURE_{\mathbf{P}}(T)$ is consistent for any set $\mathbf{P}$ of predicates of T. ∎

Notice that the class of universal theories includes the Horn theories, which have attracted considerable attention from the PROLOG, AI, and Database communities.

4. Well-Founded Theories and Predicate Circumscription

In this section we describe some limitations of predicate circumscription with respect to well-founded theories. The first such result is that predicate circumscription yields no new positive ground instances of any of the predicates being circumscribed.

Theorem 4.1

Suppose that T is a well-founded theory, $P \in \mathbf{P}$ is an n-ary predicate, and $\vec{\alpha}$ is an n-tuple of ground terms. Then

$$CLOSURE_{\mathbf{P}}(T) \vdash P\vec{\alpha} \iff T \vdash P\vec{\alpha}.$$ ∎

On reflection, this is not too surprising, since circumscription is intended to minimize the extensions of those predicates being circumscribed. New positive instances of such predicates should not arise from this minimization.

A more interesting result is that no new ground instances, positive or negative, of uncircumscribed predicates can be derived by predicate circumscription.

Theorem 4.2

Suppose that T is a well-founded theory, $P \notin \mathbf{P}$ is an n-ary predicate, and $\vec{\alpha}$ is an n-tuple of ground terms. Then

(i) $CLOSURE_{\mathbf{P}}(T) \vdash P\vec{\alpha} \iff T \vdash P\vec{\alpha}$, and

(ii) $CLOSURE_{\mathbf{P}}(T) \vdash \neg P\vec{\alpha} \iff T \vdash \neg P\vec{\alpha}$. ∎

In summary, Theorems 4.1 and 4.2 tell us that the only new ground literals that can be conjectured by predicate circumscription of well-founded theories are negative instances of one of the predicates being circumscribed. An unfortunate consequence of this result is that the usual kinds of default reasoning cannot be realized by predicate circumscription. To see why, consider the standard "flying birds" example. The relevant facts may be represented in various ways, two of which follow:

1) In this representation, all of the exceptions to flight are listed explicitly in the axiom sanctioning the conclusion that birds can fly.

   $\forall x.\ Bird(x) \wedge \neg Penguin(x) \wedge \neg Ostrich(x) \wedge \neg Dead(x) \wedge \cdots \supset Can\text{-}Fly(x)$

   In addition, there are various IS-A axioms, as well as mutual exclusion axioms:

   $\forall x.\ Canary(x) \supset Bird(x)$
   $\forall x.\ Penguin(x) \supset Bird(x)$
   etc.
   $\forall x.\ \neg(Canary(x) \wedge Penguin(x))$
   $\forall x.\ \neg(Penguin(x) \wedge Ostrich(x))$
   etc.

2) In this representation, due to McCarthy [1984], a new predicate, ab, standing for "abnormal", is introduced. One then states that "normal" birds can fly:

   $\forall x.\ Bird(x) \wedge \neg ab(x) \supset Can\text{-}Fly(x)$

   The abnormal birds are listed:

   $\forall x.\ Penguin(x) \supset ab(x)$
   $\forall x.\ Ostrich(x) \supset ab(x)$
   etc.

   Finally, one includes the IS-A and mutual exclusion axioms as in (1) above.

Both representations (1) and (2) are universal, and hence well-founded, theories. Therefore, if $Bird(Tweety)$ is given, Theorems 4.1 and 4.2(i) tell us that the default assumption $Can\text{-}Fly(Tweety)$ cannot be conjectured by predicate circumscription.

Careful readers of [McCarthy 1980] might find Theorems 4.1 and 4.2 inconsistent with the results in Section 7 of that paper. In the blocks-world example presented there to illustrate predicate circumscription, the ground instance $on(A,C,result(move(A,C),s_0))$ can be derived by circumscribing a different predicate, $\lambda z.prevents(z,move(A,C),s_0)$. This appears to violate our Theorem 4.2(i). This discrepancy stems from the fact that in formulating the circumscription schema for this example, McCarthy uses specializations of some of the original axioms, and omits one of the axioms. Thus, only part of the theory enters into the circumscription for his example, whereas our Theorems 4.1 and 4.2 suppose that the entire theory is used in proposing a circumscription schema.

A generalization of predicate circumscription has been recently formulated by McCarthy [1984]. This generalization provides for the minimization of arbitrary first-order expressions rather than simple predicates. It also provides for the treatment of designated predicates as variables of the minimization. In this version of circumscription the limitations of our Theorem 4.2 no longer apply. Thus, as some of McCarthy's examples show, it is possible to circumscribe a predicate P, treating another predicate Q as variable, and derive new positive and negative ground instances of Q. In particular, McCarthy's new formalism appears adequate for the treatment of default reasoning, as his "flying birds" example shows.

5. Equality

We now consider some limitations of predicate circumscription with respect to the treatment of equality. These limitations will be seen to have consequences for two special cases of closed-world reasoning, namely deriving the "unique names assumption" and the "domain closure assumption".

5.1. The Unique Names Assumption

When told that Tom, Dick and Harry are friends, one naturally assumes that 'Tom', 'Dick' and 'Harry' denote distinct individuals: Tom $\neq$ Dick, Tom $\neq$ Harry, Dick $\neq$ Harry. For a more general example, consider a setting in which one is told that Tom's telephone number is the same as Sue's, and that Bill's number is 555-1234, which is different from Mary's number. Thus, we have:

$$\text{tel-no(Tom)} = \text{tel-no(Sue)}$$
$$\text{tel-no(Bill)} = 555\text{-}1234$$
$$\text{tel-no(Mary)} \neq 555\text{-}1234$$

One would naturally assume from this information that tel-no(Tom) $\neq$ 555-1234, and that tel-no(Tom) $\neq$ tel-no(Mary).

In general, the unique names assumption is invoked whenever one can assume that all of the relevant information about the equality of individuals has been specified. All pairs of individuals not specified as identical are assumed to be different. This assumption arises in a number of settings, for example in the theory of databases [Reiter 1980], and in connection with the semantics of negation in PROLOG [Clark 1978]. Virtually every AI reasoning system, with the exception of those based on theorem-provers, implicitly makes this assumption. Because of Clark's results, we know that this is also the case for PROLOG based AI systems.

Unique names axioms are also important for closed-world reasoning using predicate circumscription. For example, if all we know is that *Opus* is a *Penguin*, we can circumscriptively conjecture $\forall x.Penguin(x) \equiv x = Opus$. We cannot use this to deduce $\neg Penguin(Tweety)$, however, unless we know $Opus \neq Tweety$.

How then can we formalize reasoning under the unique names assumption? The natural first attempt is to circumscribe the equality predicate in the theory under consideration, but this will not work; nothing new is derivable by circumscribing the equality predicate.

Theorem 5.1

Let T be a first-order theory containing axioms which define the equality predicate, =. Then $T \vdash CLOSURE_{\{=\}}(T)$. ∎

In view of this result, one might attempt to capture the unique names assumption by jointly circumscribing several predicates of the theory, not just the equality predicate. We do not know whether there are any theories for which this might work, but it cannot succeed for well-founded theories. No new ground equalities or inequalities can be derived by circumscribing a well-founded theory, regardless of the predicates circumscribed.

Theorem 5.2

Suppose that T is a well-founded theory containing axioms which define the equality predicate, α and β are ground terms, and **P** is a set of some of the predicates of T. Then

(i) $CLOSURE_{\mathbf{P}}(T) \vdash \alpha = \beta \Leftrightarrow T \vdash \alpha = \beta$, and

(ii) $CLOSURE_{\mathbf{P}}(T) \vdash \alpha \neq \beta \Leftrightarrow T \vdash \alpha \neq \beta$. ∎

Returning to the "Penguin" example above, we see that predicate circumscription cannot conjecture $\neg Penguin(Tweety)$ unless it is known that $Opus \neq Tweety$; otherwise we could derive $Opus \neq Tweety$ from

CLOSURE ({*Penguin* (*Opus*)}), contradicting Theorem 5.2.

In a recent paper, McCarthy [1984] proposes a circumscriptive approach to the unique names assumption by introducing two equality predicates. One of these is the standard equality predicate, but restricted to arguments which are names of objects. The other equality predicate, e(x,y), means that the names x and y denote the same object. e is axiomatized as an equivalence relation which does not, however, satisfy the full principle of substitution, in contrast to "normal" equality. This failure of full substitutivity for the predicate e prevents our Theorems 5.1 and 5.2 from applying to e. Benjamin Grosof (personal communication) has independently proposed a similar approach to the unique names assumption. He has also observed that our Theorem 5.1 applies to McCarthy's [1984] more general notion of circumscription.

5.2. The Domain Closure Assumption

The domain closure assumption is the assumption that, in a given first-order theory T, the universe of discourse is restricted to the smallest set which contains those individuals mentioned in T, and which is closed under the application of those functions mentioned in T. Domain circumscription [McCarthy 1977, 1980] is a proposed formalization of this assumption. McCarthy [1980] shows that domain circumscription can be reduced to predicate circumscription.

The simplest setting in which the domain closure assumption can arise is for a theory with a finite Herbrand Universe $\{c_1,...,c_n\}$. In this case we might want to conjecture the *domain closure axiom* for this theory: $\forall x . x=c_1 \vee ... \vee x=c_n$. Such an axiom is important for the theory of first-order databases [Reiter 1980]. No such axiom can arise from predicate circumscription for well-founded theories.

Theorem 5.3

Suppose that T is a well-founded theory; $t_1,...,t_n$ are ground terms; and **P** is a set of some of the predicate symbols of T. Then

$$CLOSURE_{\mathbf{P}}(T) \vdash \forall x . x=t_1 \vee ... \vee x=t_n \iff T \vdash \forall x . x=t_1 \vee ... \vee x=t_n . \qquad \blacksquare$$

6. Discussion

One obvious problem with using circumscription in a given setting is knowing just what to circumscribe. Some of our results provide clues in this direction. Theorem 4.2 tells us that if we wish to use predicate circumscription to conjecture $\neg P(\vec{\alpha})$ in some well-founded theory then we must include P among the predicates being circumscribed. Theorems 4.1 and 4.2 tell us that predicate circumscription will not do at all if we wish to conjecture $P(\vec{\alpha})$, as is the case for most forms of default reasoning, so that we must appeal to some other mechanism, such as McCarthy's more general form of circumscription.

A natural question is the extent to which our results translate to McCarthy's generalized circumscription. Grosof has observed that an appropriate version of our Theorem 5.1 continues to hold. Most of the examples of [McCarthy 1984] invoke the following pattern: Jointly circumscribe predicates $P_1,...,P_m$ treating predicates $Q_1,...,Q_n$ as variables. Our results apply only to the case with no variables. The natural question, therefore, is what role do the variables play in circumscriptively conjecturing ground instances, $P(\vec{\alpha})$ and $\neg P(\vec{\alpha})$, of some predicate P. Is there some way to determine which predicates should be taken as variable? Similarly, one might seek classes of theories for which such circumscriptions preserve consistency. Notice, however, that many of our proofs appeal to the model theory of predicate circumscription. It would seem, therefore, that a prior problem is to determine the appropriate model theory for predicate circumscription with variables.

Acknowledgments

We would like to thank Paul Gilmore, Akira Kanda, and Philippe Besnard for their insights and suggestions. Thanks also to UBC's Laboratory for Computational Vision for access to their coveted typesetting facilities.

References

Bossu, G. and Siegel, P. (1984), "The Saturation to the Rescue of the Non-Monotonicity", to appear in *Artificial Intelligence*.

Clark, K.L. (1978), "Negation as Failure", in [Gallaire and Minker 78], 293-322.

Gallaire, H. and Minker, J. (1978), (eds), *Logic and Data Bases*, Plenum Press, 1978.

Kowalski, R. (1978), "Logic For Data Description", in [Gallaire and Minker 78], 77-103.

McCarthy, J. (1977), "Epistemological Problems of Artificial Intelligence", *Proc. IJCAI-5*, 1977, 1038-1044.

McCarthy, J. (1980), "Circumscription – a Form of Non-Monotonic Reasoning", *Artificial Intelligence 13*, 1980, 27-39.

McCarthy, J. (1984), "Applications of Circumscription to Formalizing Commonsense Knowledge", Technical Report, Dept. of Computer Science, Stanford University, 1984.

Reiter, R. (1978a), "On Closed-World Data Bases", in [Gallaire and Minker 78], 55-76.

Reiter, R. (1978b), "On Reasoning by Default", *Proc. Second Symposium on Theoretical Issues in Natural Language Processing,* Urbana, Illinois, 1978, 210-218.

Reiter, R. (1980), "Equality and Domain Closure in First-Order Data Bases", *JACM 27(2)*, 1980, 235-249.

Reiter, R. (1982), "Circumscription Implies Predicate Completion (Sometimes)", *Proc. AAAI-82*, 1982, 418-420.

Shepherdson, J.C. (1984), "Negation as Failure: A Comparison of Clark's Completed Data Base and Reiter's Closed-World Assumption", Report PM-84-01, School of Mathematics, University of Bristol, 1984.

Appendix
Proofs of Theorems

Theorem 3.2

Universal theories are well-founded.

Proof.

The proof is identical to that of Proposition 4 in [Bossu and Seigel 1984]. The definition of submodel used there is less restrictive than that used here, but this does not alter the form of the proof. ∎

Theorem 4.1.

If T is a well-founded theory, $\vec{\alpha}$ is an n-tuple of ground terms, and $P \in \mathbf{P}$, is an n-ary predicate, then

$$CLOSURE_{\mathbf{P}}(T) \vdash P\vec{\alpha} \Leftrightarrow T \vdash P\vec{\alpha}.$$

Proof.

The if-half of the theorem is immediate. We prove the contrapositive of the only-if-half. Assume that $CLOSURE_{\mathbf{P}}(T) \vdash P\vec{\alpha}$ and $T \not\vdash P\vec{\alpha}$. Then T has a model, M, in which $P\vec{\alpha}$ is false. Since T is well-founded, there is a **P**-minimal submodel, M', of M. Furthermore, since the circumscription is true in all **P**-minimal submodels, $P\vec{\alpha}$ is true in M'. But then M' is not a **P**-submodel of M, and this contradicts the fact that M' is a **P**-minimal submodel of M. Therefore $CLOSURE_{\mathbf{P}}(T) \not\vdash P\vec{\alpha}$. ∎

Theorem 4.2

If T is a well-founded theory, $\vec{\alpha}$ is an n-tuple of ground terms, and $P \notin \mathbf{P}$ is an n-ary predicate, then

(i) $CLOSURE_{\mathbf{P}}(T) \vdash P\vec{\alpha} \Leftrightarrow T \vdash P\vec{\alpha}$, and

(ii) $CLOSURE_{\mathbf{P}}(T) \vdash \neg P\vec{\alpha} \Leftrightarrow T \vdash \neg P\vec{\alpha}$.

Proof.

(i) The if-half is immediate. We prove the contrapositive of the only-if-half. Assume $T \not\vdash P\vec{\alpha}$. Then there is a model, M, for T in which $P\vec{\alpha}$ is false. Since T is well-founded, there is a **P**-minimal submodel, M', of M. By the definition of submodel, the interpretation of P remains the same in M and M', since $P \notin \mathbf{P}$. Hence $P\vec{\alpha}$ is false in M'. Since the circumscription is true in all minimal models, $CLOSURE_{\mathbf{P}}(T) \not\vdash P\vec{\alpha}$. The proof for (ii) is similar. ∎

In the proofs of Theorems 5.1 and 5.2 we use the following notational conventions:

1. *SCHEMA* $(T,\mathbf{P})$ is the circumscription schema resulting from circumscribing the predicates of **P** in T.
2. $CLOSURE_{\{\}}(T) = T$. (The closure of T with respect to the empty set of predicates is defined to be T itself.)
3. $\bigwedge_{i=1}^{0} Q_i$, the empty conjunction, stands for any tautology.

Theorem 5.1

If T is an arbitrary theory containing axioms which define the equality predicate, $=$, then $T \vdash CLOSURE_{\{=\}}(T)$.

Proof.

Consider the schema resulting from circumscribing $=$ in T:

$$SCHEMA\,(T,\{=\}) = [T(\Phi) \wedge \forall xy.\ \Phi xy \supset x=y] \supset \forall xy.x=y \supset \Phi xy.$$

We show that $T \vdash SCHEMA\,(T,\{=\})$, whence $T \vdash CLOSURE_{\{=\}}(T)$. It is sufficient to show for any instance Ψ of the predicate parameter Φ that

$$T\,,T(\Psi) \vdash \forall xy.x=y \supset \Psi xy. \tag{5.1}$$

Now $\forall x.\ \Psi xx$ is one of the conjuncts of $T(\Psi)$, since T contains the axiom $\forall x.x=x$. Moreover, by basic properties of equality,

$$T \vdash \forall xy.\ \Psi xx \wedge x=y \supset \Psi xy.$$

The result (5.1) now follows. ∎

Theorem 5.2

If T is a well-founded theory containing axioms which define the equality predicate, $=$; α and β are ground terms; and $\mathbf{P}$ is a set of some of the predicates of T; then

(i) $CLOSURE_{\mathbf{P}}(T) \vdash \alpha = \beta \Leftrightarrow T \vdash \alpha = \beta$, and

(ii) $CLOSURE_{\mathbf{P}}(T) \vdash \alpha \neq \beta \Leftrightarrow T \vdash \alpha \neq \beta$.

Proof.

(i) This is a corollary of Theorems 4.1 and 4.2(i).

(ii) The if-half is immediate. To prove the only-if-half, notice that if $\mathbf{P}$ does not include the equality predicate then the result follows directly from Theorem 4.2(ii). So assume $\mathbf{P}$ does include the equality predicate, say $\mathbf{P} = \{=, P_1,\ldots, P_n\}$. Writing $\vec{\Phi}$ for $\Phi_1,\ldots,\Phi_n$; $A(\vec{\Phi})$ for $\bigwedge_{i=1}^{n}(\forall \vec{x}.\Phi_i \vec{x} \supset P_i \vec{x})$;

and $B(\vec{\Phi})$ for $\bigwedge_{i=1}^{n}(\forall \vec{x}.P_i\vec{x} \supset \Phi_i\vec{x})$; we have:

$$SCHEMA(T,\{P_1,\ldots, P_n\}) = T(\vec{\Phi}) \wedge A(\vec{\Phi}) \supset B(\vec{\Phi})$$

$$SCHEMA(T,\{=, P_1,\ldots, P_n\}) = T(\Psi, \vec{\Phi}) \wedge A(\vec{\Phi}) \wedge (\forall xy.\, \Psi xy \supset x=y)$$
$$\supset B(\vec{\Phi}) \wedge (\forall xy.x=y \supset \Psi xy)$$

To establish the result we seek, it is sufficient to show that

$$CLOSURE_{\{P_1,\ldots, P_n\}}(T) \vdash \text{[any instance of } SCHEMA(T, \{=, P_1,\ldots, P_n\})] \tag{5.2}$$

since, by hypothesis,

$$CLOSURE_{\{=, P_1,\ldots, P_n\}}(T) \vdash \alpha \neq \beta.$$

Therefore, if (5.2) holds, we would have

$$CLOSURE_{\{P_1,\ldots, P_n\}}(T) \vdash \alpha \neq \beta.$$

By Theorem 4.2(ii), $T \vdash \alpha \neq \beta$, as required, since $\{P_1,\ldots, P_n\}$ does not include the equality predicate. To show (5.2), we first establish the following result:
For fixed predicates, Ψ' and $\vec{\Phi}'$,

$$T,\ T(\Psi', \vec{\Phi}'),\ \forall xy.\, \Psi' xy \supset x=y \vdash \forall xy.\, \Psi' xy \equiv x=y \tag{5.3}$$

Clearly, to show this, it is sufficient to show

$$T,\ T(\Psi', \vec{\Phi}') \vdash \forall xy.x=y \supset \Psi' xy$$

whose proof is essentially the same as the proof of (5.1) in Theorem 5.1. Now (5.3) says that under the stated assumptions, the predicates Ψ' and $=$ are indistinguishable. Hence

$$T,\ T(\Psi', \vec{\Phi}'),\ \forall xy.\, \Psi' xy \supset x=y \vdash T(\Psi', \vec{\Phi}') \equiv T(=, \vec{\Phi}') \tag{5.4}$$

Noting that $T(=, \vec{\Phi}')$ is what we are denoting by $T(\vec{\Phi}')$ we have, using (5.3) and (5.4):

$$T,\ T(\Psi', \vec{\Phi}'),\ \forall xy.\, \Psi' xy \supset x=y,\ T(\vec{\Phi}') \wedge A(\vec{\Phi}') \supset B(\vec{\Phi}'),\ A(\vec{\Phi}')$$
$$\vdash B(\vec{\Phi}') \wedge (\forall xy.x=y \supset \Psi' xy)$$

which is equivalent to

$$T,\ T(\vec{\Phi}') \wedge A(\vec{\Phi}') \supset B(\vec{\Phi}')$$
$$\vdash T(\Psi', \vec{\Phi}') \wedge A(\vec{\Phi}') \wedge (\forall xy.\ \Psi' xy \supset x=y)$$
$$\supset B(\vec{\Phi}') \wedge (\forall xy. x=y \supset \Psi' xy).$$

Thus we have

T, [instance of *SCHEMA* $(T,\{P_1,\ldots, P_n\})$]

$\vdash$ [instance of *SCHEMA* $(T,\{=, P_1,\ldots, P_n\})$]

from which (5.2) follows. ∎

Theorem 5.3

If T is a well-founded theory; $\alpha_1,\ldots,\alpha_n$ are ground terms; and **P** is a set of some of the predicate symbols of T; then

$$CLOSURE_{\mathbf{P}}(T) \vdash \forall x. x=\alpha_1 \vee \ldots \vee x=\alpha_n \Leftrightarrow T \vdash \forall x. x=\alpha_1 \vee \ldots \vee x=\alpha_n .$$

Proof.

The if-half is immediate. For the only-if-half, assume that $T \not\vdash \forall x. x=\alpha_1 \vee \ldots \vee x=\alpha_n$. Then T has a model which falsifies $\forall x. x=\alpha_1 \vee \ldots \vee x=\alpha_n$. Since T is well-founded, this model has a **P**-minimal submodel. But $\forall x. x=\alpha_1 \vee \ldots \vee x=\alpha_n$ is false in this submodel, because the extension of the equality predicate in this submodel must be a subset of its extension in the original model. Since the circumscription is true in all minimal models, $CLOSURE_{\mathbf{P}}(T) \not\vdash \forall x. x=\alpha_1 \vee \ldots \vee x=\alpha_n$. ∎