Interpolation Theorems for Nonmonotonic Reasoning Systems

Eyal Amir
Computer Science Division,
University of California, at Berkeley
Berkeley, CA 94720
eyal@cs.berkeley.edu

August 18, 2002

Abstract

Craig's interpolation theorem [Craig, 1957] is an important theorem known for propositional logic and first-order logic. It says that if a logical formula β logically follows from a formula α, then there is a formula γ, including only symbols that appear in both α, β, such that β logically follows from γ and γ logically follows from α. Such theorems are important and useful for understanding those logics in which they hold as well as for speeding up reasoning with theories in those logics. In this paper we present interpolation theorems in this spirit for three nonmonotonic systems: circumscription, default logic and logic programs with the stable models semantics (a.k.a. answer set semantics). These results give us better understanding of those logics, especially in contrast to their nonmonotonic characteristics. They suggest that some *monotonicity* principle holds despite the failure of classic monotonicity for these logics. Also, they sometimes allow us to use methods for the decomposition of reasoning for these systems, possibly increasing their applicability and tractability. Finally, they allow us to build structured representations that use those logics.

Topics: Nonmonotonic Reasoning

1 Introduction

Craig's interpolation theorem [Craig, 1957] is an important theorem known for propositional logic and first-order logic (FOL). It says that if α, β are two logical formulae and $\alpha \vdash \beta$, then there is a formula $\gamma \in \mathcal{L}(\alpha) \cap \mathcal{L}(\beta)$ such that $\alpha \vdash \gamma$ and $\gamma \vdash \beta$ ("$\vdash$" is the classical logical deduction relation; $\mathcal{L}(\alpha)$ is the language of α (the set of formulae built with the nonlogical symbols of α, $L(\alpha)$)). Such interpolation theorems allow us to break inference into pieces associated

with sublanguages of the language of that theory [Amir and McIlraith, 2000; McIlraith and Amir, 2001], for those formal systems in which they hold. In AI, these properties have been used to speed up inference for constraint satisfaction systems (CSPs), propositional logic and FOL (e.g., [Dechter and Pearl, 1988; Darwiche, 1998; McIlraith and Amir, 2001; Dechter and Rish, 1994; Amir and McIlraith, 2000; Dechter, 1999; Darwiche, 1997]) and to build structured representations [Darwiche, 1998; Amir, 2000; Darwiche, 1997].

In this paper we present interpolation theorems for three nonmonotonic systems: *circumscription* [McCarthy, 1980], *default logic* [Reiter, 1980] and *logic programs* with the Answer Set semantics [Gelfond and Lifschitz, 1991; Gelfond and Lifschitz, 1988]. In the nonmonotonic setup there are several interpolation theorems for each system, with different conditions for applicability and different form of interpolation. This stands in contrast to classical logic, where Craig's interpolation theorem always holds. Our theorems allow us to use methods for the decomposition of reasoning (a-la [Amir and McIlraith, 2000; McIlraith and Amir, 2001]) under some circumstances for these systems, possibly increasing their applicability and tractability for structured theories. We list the main theorems that we show in this paper below, omitting some of their conditions for simplicity.

For circumscription we show that, under some conditions, $Circ[\alpha; P; Q] \models \beta$ iff there is some set of formulae $\gamma \subseteq \mathcal{L}(\alpha) \cap \mathcal{L}(\{\beta\} \cup P)$ such that $\alpha \models \gamma$ and $Circ[\gamma; P; Q] \models \beta$. For example, to answer $Circ[BlockW; block; L(BlockW)] \models on(A, B)$, we can compute this formula $\gamma \in \mathcal{L}(\{block, on, A, B\})$ from $BlockW$ *without applying circumscription*, and then solve $Circ[\gamma; block; L(BlockW)] \models on(A, B)$ (where γ may be significantly smaller than $BlockW$).

For default logic, letting $\alpha \mid\!\sim_D \beta$ mean that every extension of $\langle \alpha, D \rangle$ entails β (*cautious entailment*), we show that, under some conditions, if $\alpha \mid\!\sim_D \beta$, then there is a formula $\gamma \in \mathcal{L}(\alpha \cup D) \cap \mathcal{L}(\beta)$ such that $\alpha \mid\!\sim_D \gamma$ and $\gamma \mid\!\sim_D \beta$. For logic programs we show that if P_1, P_2 are two logic programs and $\varphi \in \mathcal{L}(P_2)$ such that $P_1 \cup P_2 \mid\!\sim^b \varphi$, then there is $\gamma \in \mathcal{L}(P_1) \cap \mathcal{L}(P_2)$ such that $P_1 \mid\!\sim^b \gamma$ and $P_2 \cup \gamma \mid\!\sim^b \varphi$ (here $\mid\!\sim^b$ is the *brave* entailment relation for logic programs).

This paper focuses on the form of the interpolation theorems that hold for those nonmonotonic logics. We do not address the possible application of these results to the problem of automated reasoning with those logics. Nonetheless, we mention that direct application of those results is possible along the lines already explored for propositional logic and FOL in [Amir and McIlraith, 2000; McIlraith and Amir, 2001].

No interpolation theorems were shown for nonmonotonic reasoning systems before this paper. Nonetheless, some of our theorems for default logic and logic programs are close to the *splitting theorems* of [Lifschitz and Turner, 1994; Turner, 1996], which have already been used to decompose reasoning for those logics. The main difference between our theorems and those splitting theorems is that the latter change some of the defaults/rules involved to provide the corresponding entailment. Also, they do not talk about an interpolant γ, but rather discuss combining extensions.

Since its debut, the nonmonotonic reasoning line of work has expanded and

several textbooks exist that give a fair view of nonmonotonic reasoning and its uses (e.g., [Brewka, 1991; Antoniou, 1997; Brewka *et al.*, 1997; Gabbay *et al.*, 1993]). The reader is referred to those books for background and further details.

2 Logical Preliminaries

In this paper, we use the notion of *logical theory* for every set of axioms in FOL or propositional logic, regardless of whether the set of axioms is deductively closed or not. For a theory $\mathcal{A}$, we use $L(\mathcal{A})$ to denote the signature of $\mathcal{A}$, i.e., the set of non-logical symbols. $\mathcal{L}(\mathcal{A})$ denotes the language of $\mathcal{A}$, i.e., the set of formulae built with $L(\mathcal{A})$. $Cn(\mathcal{A})$ is the set of logical consequences of $\mathcal{A}$ (i.e., those formulae that are valid consequences of $\mathcal{A}$ in FOL). For a first-order structure, M, in L, we write $U(M)$ for the universe of elements of M. For every symbol, term or formula, s, in $\mathcal{L}$, we write s^M for the interpretation of s in M.

Finally, we note Craig's Interpolation Theorem for FOL (thus, also for propositional logic).

Theorem 2.1 ([Craig, 1957]) *Let α, β be sentences such that $\alpha \vdash \beta$. Then there is a formula γ involving only nonlogical symbols common to both α and β, such that $\alpha \vdash \gamma$ and $\gamma \vdash \beta$.*

3 Circumscription

3.1 McCarthy's Circumscription: Overview

McCarthy's circumscription [McCarthy, 1980; McCarthy, 1986] is a nonmonotonic reasoning system in which inference from a set of axioms, A, is performed by minimizing the extent of some predicate symbols $\overrightarrow{P}$, while allowing some other nonlogical symbols, $\overrightarrow{Z}$ to vary.

Formally, McCarthy's (second-order) circumscription formula

$$Circ[A(P,Z);P;Z] = A(P,Z) \wedge \forall p,z\ (A(p,z) \Rightarrow \neg(p < P)) \qquad (1)$$

says that in the theory A, with parameter relations and function vectors (sequence of symbols) P, Z, P is a minimal element such that $A(P,Z)$ is still consistent, when we are allowed to vary Z in order to allow P to become smaller. Here, $p < P$ means that $\forall \overrightarrow{x}(p(\overrightarrow{x}) \Rightarrow P(x)) \wedge \exists x(P(x) \wedge \neg p(x))$.

Take for example the following simple theory: $T = block(B_1) \wedge block(B_2)$. Then, the circumscription of *block* in T, varying nothing, is

$$Circ[T;block;] = T \wedge \forall p\ [T_{[block/p]} \Rightarrow \neg(p < block)].$$

Roughly, this means that *block* is a minimal predicate satisfying T. Computing circumscription is discussed in length in [Lifschitz, 1993] and others, and we do not expand on it here. Using known techniques we can conclude

$$Circ[T;block;] \equiv \forall x\ (block(x) \Leftrightarrow (x = B_1 \vee x = B_2))$$

This means that there are no other blocks in the world other than those mentioned in the original theory T.

We present the preferential semantics for circumscription that was given by [Lifschitz, 1985; McCarthy, 1986; Etherington, 1986] in the following definition.

Definition 3.1 ([Lifschitz, 1985]) *For any two models M and N of a theory T we write $M \leq_{P,Z} N$ if the models M, N differ only in how they interpret predicates from P and Z and if the extension of every predicate from P in M is a subset of its extension in N. We write $M <_{P,Z} N$ if for at least one predicate in P the extension in M is a strict subset of its extension in N.*

A model M of T is $\leq_{P,Z}$*-minimal* if there is no model N of T such that $N <_{P,Z} M$.

Theorem 3.2 ([Lifschitz, 1985]: Circumscript. Semantics) *Let T be a finite set of sentences. A structure M is a model of $Circ[T; P; Z]$ iff M is a $\leq_{P,Z}$-minimal model of T.*

This theorem allows us to extend the definition of circumscription to infinite sets of sentences. In those cases, $Circ[T; P; Z]$ is defined as the set of sentences that hold in all the $\leq_{P,Z}$-minimal models of T. Theorem 3.2 implies that this extended definition is equivalent to the syntactic characterization of the original definition (equation (1)) if T is a finite set of sentences. In the rest of this paper, we refer to this extended definition of circumscription, if T is an infinite set of FOL sentences (we will note those cases when we encounter them).

Circumscription satisfies Left Logical Equivalence (LLE): $T \equiv T'$ implies that $Circ[T; P; Z] \equiv Circ[T'; P; Z]$. It also satisfies Right Weakening: $Circ[T; P; Z] \models \varphi$ and $\varphi \Rightarrow \psi$ implies that $Circ[T; P; Z] \models \psi$).

3.2 Model Theory

The following reviews some classical results in *model theory* (see [Hodges, 1997] for a reference).

Definition 3.3 *Let M, N be L-structures, for FOL signature L and language $\mathcal{L}$. We say that N is an* elementary extension *of M (or M is an* elementary substructure *of N), written $M \preceq N$, if $U(M) \subseteq U(N)$ and for every $\varphi(\overrightarrow{x}) \in \mathcal{L}$ and vector of elements $\overrightarrow{a}$ of $U(M)$, $M \models \varphi(\overrightarrow{a})$ iff $N \models \varphi(\overrightarrow{a})$.*

$f : M \to N$ is an *elementary embedding* if f is an injective (one-to-one) homomorphism from M to N and for every $\varphi(\overrightarrow{x}) \in \mathcal{L}$ and vector $\overrightarrow{a} = \langle a_1, ..., a_n \rangle$ of elements from $U(M)$ (i.e., $a_1, .., a_n \in U(M)$), $M \models \varphi(\overrightarrow{a})$ iff $N \models \varphi(f(a_1), ..., f(a_n))$.

For FOL signatures $L \subseteq L^+$, and for N an L^+-structure, we say that $N \restriction L$ is the *reduct* of N to L, the L-structure with the same universe of elements as N, and the same interpretation as N for those symbols from L^+ that are in L (there is no interpretation for symbols not in L). For A theory T in a language of L^+, let $Cn^L(T)$ be the set of all classical consequences of T in the language of L.

The following theorem is a model-theoretic property that is analogous to Craig's interpolation theorem (Theorem 2.1).

Theorem 3.4 (See [Hodges, 1997] p.148) *Let L, L^+ be FOL signatures with $L \subseteq L^+$ and T a theory in the language of L^+. Let M be an L-structure. Then, $M \models Cn^L(T)$ if and only if for some model N of T, $M \preceq N \restriction L$ (M is an elementary substructure of the reduct of N to L).*

3.3 Interpolation in Circumscription

In this section we present two interpolation theorems for circumscription. Those theorems hold for both FOL and propositional logic. Roughly speaking, the first (Theorem 3.8) says that if α nonmonotonically entails β (here this means $Circ[\alpha; P; Q] \models \beta$), then there is $\gamma \subseteq \mathcal{L}(\alpha) \cap \mathcal{L}(\beta \cup P)$ such that α classically entails γ ($\alpha \models \gamma$) and γ nonmonotonically entails β ($Circ[\gamma; P; Q] \models \beta$). In the FOL case this γ can be an infinite set of sentences, and we use the extended definition of Circumscription for infinite sets of axioms for this statement.

The second theorem (Theorem 3.12) is similar to the first, with two main differences. First, it requires that $L(\alpha) \subseteq (P \cup Q)$. Second, it guarantees that γ as above (and some other restrictions) exists iff α nonmonotonically entails β. This is in contrast to the first theorem that guarantees only the *if* part. The actual technical details are more fine than those rough statements, so the reader should refer to the actual theorem statements below.

In addition to these two theorems, we present another theorem that addresses the case of reasoning from the union of theories (Theorem 3.11). Before we state and prove those theorems, we prove several useful lemmas.

Our first lemma says that if we are given two theories T_1, T_2, and we know the set of sentences γ that follows from T_2 in the intersection of their languages, then every model of this set of sentences together with T_1 can be extended to a model of $T_1 \cup T_2$.

Lemma 3.5 *Let T_1, T_2 be two theories, with signatures in L_1, L_2, respectively. Let γ be a set of sentences logically equivalent to $Cn^{L_1 \cap L_2}(T_2)$. For every L_1-structure, $\mathcal{M}$, that satisfies $T_1 \cup \gamma$ there is a $(L_1 \cup L_2)$-structure, $\widehat{\mathcal{M}}$, that is a model of $T_1 \cup T_2$ such that $\mathcal{M} \preceq \widehat{\mathcal{M}} \restriction L_1$.*

PROOF Let $\mathcal{M}$ be a L_1-structure that is a model of $T_1 \cup \gamma$. Then $\mathcal{M} \models \gamma$. Noticing that γ is logically equivalent to $Cn^{L_1 \cap L_2}(T_2)$ (by definition of γ), we get that $\gamma \models Cn^{L_1 \cap L_2}(T_2)$. Consequently, $\gamma \models Cn^{L_1}(T_2)$ because $Cn^{L_1 \cap L_2}(T_2) \equiv Cn^{L_1}(Cn^{L_1 \cap L_2}(T_2)) = Cn^{L_1}(T_2)$.

Now we use Theorem 3.4 with $L = L_1$, $L^+ = L_1 \cup L_2$, $M = \mathcal{M}$ and $T = T_1 \cup T_2$. We know that $\mathcal{M} \models T_1 \cup \gamma$. Thus, $M \models T_1 \cup Cn^{L_1}(T_2)$. To use Theorem 3.4 we need to show that $M \models Cn^{L_1}(T_1 \cup T_2)$. We use Craig's interpolation theorem (Theorem 2.1) to show this is indeed the case.

First notice that $Cn^{L_1}(T_1 \cup T_2) \supseteq Cn^{L_1}(T_1 \cup \gamma)$ is true because $T_2 \models \gamma$. We show that $Cn^{L_1}(T_1 \cup T_2) \subseteq Cn^{L_1}(T_1 \cup \gamma)$. Take $\varphi \in Cn^{L_1}(T_1 \cup T_2)$. By

definition, $T_1 \cup T_2 \models \varphi$ and $\varphi \in \mathcal{L}(T_1)$. The deduction theorem for FOL implies that $T_2 \models T_1' \Rightarrow \varphi$, for some finite subset $T_1' \subseteq T_1$. Craig's interpolation theorem for FOL implies that there is $\delta \in \mathcal{L}(T_2) \cap \mathcal{L}(T_1' \Rightarrow \varphi) = \mathcal{L}(T_2) \cap \mathcal{L}(T_1)$ such that $T_2 \models \delta$ and $\delta \models T_1' \Rightarrow \varphi$. Thus, $\delta \in Cn^{L_1 \cap L_2}(T_2) \equiv \gamma$. Consequently, $\gamma \models T_1' \Rightarrow \varphi$. Using the deduction theorem again we get that $T_1' \cup \gamma \models \varphi$, implying that $T_1 \cup \gamma \models \varphi$.

Thus, we showed that $Cn^{L_1}(T_1 \cup T_2) = Cn^{L_1}(T_1 \cup \gamma)$. From $M \models T_1 \cup Cn^{L_1}(T_2)$ and $\gamma = Cn^{L_1}(T_2)$ we get that $M \models Cn^{L_1}(T_1 \cup T_2)$.

Finally, the conditions of Theorem 3.4 for $L = L_1$, $L^+ = L_1 \cup L_2$, $M = \mathcal{M}$ and $T = T_1 \cup T_2$ hold. We conclude that there is a $(L_1 \cup L_2)$-structure, $\widehat{\mathcal{M}}$, that is a model of $T_1 \cup T_2$ such that $\mathcal{M} \preceq \widehat{\mathcal{M}} \restriction L_1$. ■

Our second lemma says that every $<_{P,Q}$-minimal model of T that is also a model of T' is a $<_{P,Q}$-minimal model of $T \cup T'$.

Lemma 3.6 *Let T be a theory and P, Q vectors of nonlogical symbols (P includes only predicate symbols). If $\mathcal{M} \models Circ[T; P; Q]$ and $\mathcal{M} \models T \cup T'$, then $\mathcal{M} \models Circ[T \cup T'; P; Q]$.*

PROOF Let $\mathcal{M}$ be a model of $T \cup T'$ such that $\mathcal{M} \models Circ[T; P; Q]$. If there is $\mathcal{M}' <_{P,Q} \mathcal{M}$ such that $\mathcal{M}' \models T \cup T'$, then $\mathcal{M}' \models T$ and $\mathcal{M} \not\models Circ[T; P; Q]$. Contradiction. Thus, there is no such $\mathcal{M}'$ and $\mathcal{M} \models Circ[T \cup T'; P; Q]$. ■

The following theorem is central to the rest of our results in this section. It says that when we circumscribe P, Q in $T_1 \cup T_2$ we can replace T_2 by its consequences in $\mathcal{L}(T_1)$, for some purposes and under some assumptions.

Theorem 3.7 *For T_1, T_2 theories and P, Q vectors of symbols from $L(T_1) \cup L(T_2)$ such that $P \subseteq L(T_1)$, let γ be a set of sentences logically equivalent to $Cn^{L(T_1) \cap L(T_2)}(T_2)$. Then, for all $\varphi \in \mathcal{L}(T_1)$, if $Circ[T_1 \cup T_2; P; Q] \models \varphi$, then $Circ[T_1 \cup \gamma; P; Q] \models \varphi$.*

PROOF We show that for every model of $Circ[T_1 \cup \gamma; P; Q]$ there is a model of $Circ[T_1 \cup T_2; P; Q]$ whose reduct to $L(T_1)$ is an elementary extension of the reduct of the first model to $L(T_1)$.

Let $\mathcal{M}$ be a $L(T_1 \cup T_2)$-structure that is a model of $Circ[T_1 \cup \gamma; P; Q]$. Then, $\mathcal{M} \models T_1 \cup \gamma$. From Lemma 3.5 we know that there is a $(L_1 \cup L_2)$-structure, $\widehat{\mathcal{M}}$, that is a model of T_2 such that $\mathcal{M} \restriction L(T_1) \preceq \widehat{\mathcal{M}} \restriction L(T_1)$.

Thus, $\widehat{\mathcal{M}}$ is a $\leq_{P,Q}$-minimal model of $T_1 \cup \gamma$. To see this, assume otherwise. Then, there is a model $\mathcal{M}'$ for the signature $L(T_1 \cup T_2)$ such that $\mathcal{M}' <_{P,Q} \widehat{\mathcal{M}}$ and $\mathcal{M}' \models T_1 \cup \gamma$. Take $\mathcal{M}''$ such that the interpretation of all the symbols in $L(T_1)$ is exactly the same as that of $\mathcal{M}'$ and such that the interpretation of all symbols in $L(T_2) \setminus L(T_1)$ is exactly the same as that of $\mathcal{M}$. Then, $\mathcal{M}'' \models T_1 \cup \gamma$ because $T_1 \cup \gamma \subseteq \mathcal{L}(T_1)$. Also, $\mathcal{M}'' <_{P,Q'} \mathcal{M}$, for $Q' = Q \cap L(T_1)$ because $P \subseteq L(T_1)$ and $\mathcal{M}, \widehat{\mathcal{M}}$ agree on the interpretation of symbols in $L(T_1)$ ($\mathcal{M} \restriction L(T_1) \preceq \widehat{\mathcal{M}} \restriction L(T_1)$). Thus, $\mathcal{M}'' <_{P,Q} \mathcal{M}$, since $\mathcal{M}'', \mathcal{M}$ agree on all the interpretation

of all symbols in $L(T_2) \setminus L(T_1)$. This contradicts $\mathcal{M} \models Circ[T_1 \cup \gamma; P; Q]$, so $\widehat{\mathcal{M}}$ is a $\leq_{P,Q}$-minimal model of $T_1 \cup \gamma$.

Thus, $\widehat{\mathcal{M}} \models Circ[T_1 \cup \gamma; P; Q]$, and $\widehat{\mathcal{M}} \models T_1 \cup T_2$. From Lemma 3.6 we get that $\widehat{\mathcal{M}} \models Circ[T_1 \cup T_2; P; Q]$. Now, let $\varphi \in \mathcal{L}(T_1)$ such that $Circ[T_1 \cup T_2; P; Q] \models \varphi$. Then every model of $Circ[T_1 \cup T_2; P; Q]$ satisfies φ. Let $\mathcal{M}$ be a model of $Circ[T_1 \cup \gamma; P; Q]$ in the language $\mathcal{L}(T_1 \cup T_2)$. Then there is $\widehat{\mathcal{M}}$ as above, i.e., $\widehat{\mathcal{M}} \models Circ[T_1 \cup T_2; P; Q]$ and $\mathcal{M} \restriction L(T_1) \preceq \widehat{\mathcal{M}} \restriction L(T_1)$. Thus, $\widehat{\mathcal{M}} \models \varphi$. Since $\mathcal{M} \restriction L(T_1) \preceq \widehat{\mathcal{M}} \restriction L(T_1)$, $\mathcal{M} \models \varphi$. Thus every model of $Circ[T_1 \cup \gamma; P; Q]$ is a model of φ. ■

Theorem 3.8 (Interpolation for Circumscription 1) *Let T be a theory, P, Q vectors of symbols, and φ a formula. If $Circ[T; P; Q] \models \varphi$, then there is $\gamma \subseteq \mathcal{L}(T) \cap \mathcal{L}(\varphi \cup P)$ such that*

$$T \models \gamma \quad \textit{and} \quad Circ[\gamma; P; Q] \models \varphi.$$

Furthermore, this holds for every γ that is logically equivalent to the consequences of T in $L(T) \cap L(\varphi \cup P)$.

PROOF We use Theorem 3.7 to find this γ. For T, φ as in the statement of the theorem we define T_1, T_2 as follows. We choose T_1 such that $\varphi \in \mathcal{L}(T_1)$ and $P \subseteq L(T_1)$: Let $T_1 = \{\varphi \vee \neg\varphi\} \cup \tau_1$ for τ_1 a set of tautologies such that $L(\tau_1) = P$. We choose T_2 such that it includes T and has a rich enough vocabulary so that $P, Q \subseteq L(T_1) \cup L(T_2)$. Let $T_2 = T \cup \tau_2$, for τ_2 a set of tautologies such that $L(\tau_2) = Q \setminus L(T_1)$. Let $L_1 = L(T_1)$, $L_2 = L(T_2)$.

Theorem 3.7 guarantees that if $P \subseteq L_1$ then γ from that theorem satisfies $Circ[T_1 \cup T_2; P; Q] \models \psi \Rightarrow Circ[T_1 \cup \gamma; P; Q] \models \psi$ for every $\psi \in \mathcal{L}(T_1)$. This implies that for every $\psi \in \mathcal{L}(\{\varphi\} \cup \tau_1)$, $Circ[T; P; Q] \models \psi \Rightarrow Circ[\gamma; P; Q] \models \psi$. In particular, $Circ[\gamma; P; Q] \models \varphi$, and this γ satisfies our current theorem. ■

Example 3.9 *For example, if $\alpha = block(A) \wedge block(B) \wedge \forall b\ (clear(b) \Leftrightarrow (block(b) \wedge \forall x \neg on(x, b)))$, $\beta = clear(A)$, $P = on$, $Q = clear$, then one possible interpolant is $\gamma = clear(A) \Leftrightarrow \forall x \neg on(x, A)$ because $\alpha \models \gamma$ and $Circ[\gamma; P; Q] \models \beta$ because*

$$Circ[clear(A) \Leftrightarrow \forall x \neg on(x, A); on; clear] \equiv \forall x, b \neg(x, b) \wedge clear(A).$$

This theorem does not hold if we require $\gamma \subseteq \mathcal{L}(T) \cap \mathcal{L}(\varphi)$ instead of $\gamma \subseteq \mathcal{L}(T) \cap \mathcal{L}(\varphi \cup P)$. For example, take $\varphi = Q$, $T = \{\neg P \Rightarrow Q\}$, where P, Q are propositional symbols. $Circ[T; P; Q] \models \varphi$. However, every logical consequence of T in $L(\varphi)$ is a tautology. Thus, if we use this modified requirement, then γ will be equivalent to $\emptyset$ and $Circ[\gamma; P; Q] \not\models \varphi$.

Theorem 3.10 *Let T_1, T_2 be two theories, P, Q two vectors of symbols from $L(T_1) \cup L(T_2)$ such that $P \subseteq L(T_1)$ and $P \cup Q \supseteq L(T_2)$. Let γ be a set of sentences logically equivalent to $Cn^{L(T_1) \cap L(T_2)}(T_2)$. Then, for all $\varphi \in \mathcal{L}(T_1)$, if $Circ[T_1 \cup \gamma; P; Q] \models \varphi$, then $Circ[T_1 \cup T_2; P; Q] \models \varphi$.*

PROOF We show that every model of $Circ[T_1 \cup T_2; P; Q]$ is also a model of $Circ[T_1 \cup \gamma; P; Q]$. Let $\mathcal{M}$ be a $L(T_1 \cup T_2)$-structure that is a model of $Circ[T_1 \cup T_2; P; Q]$. Then $\mathcal{M} \models T_1 \cup T_2$, implying that also $\mathcal{M} \models T_1 \cup \gamma$.

Assume that there is $\mathcal{M}' <_{P,Q} \mathcal{M}$ such that $\mathcal{M}' \models T_1 \cup \gamma$. From Lemma 3.5, there is $\bar{\mathcal{M}}'$ such that $\bar{\mathcal{M}}' \models T_1 \cup T_2$ and $\mathcal{M}' \restriction L(T_1) \preceq \bar{\mathcal{M}}' \restriction L(T_1)$. Since $\mathcal{M}', \bar{\mathcal{M}}'$ agree on all the symbols of $L(T_1)$, we get that $\bar{\mathcal{M}}' \leq_{P,Q} \mathcal{M}'$ (because $P \cup Q \supseteq L(T_2)$). Finally, we get that $\bar{\mathcal{M}}' \leq_{P,Q} \mathcal{M}' <_{P,Q} \mathcal{M}$, contradicting the assumption of $\mathcal{M}$ being $\leq_{P,Q}$-minimal satisfying $T_1 \cup T_2$. Thus, $\mathcal{M}$ is a model of $Circ[T_1 \cup \gamma; P; Q]$.

Let $\varphi \in \mathcal{L}(T_1)$ such that $Circ[T_1 \cup \gamma; P; Q] \models \varphi$. Then every model of $Circ[T_1 \cup \gamma; P; Q]$ satisfies φ. Let $\mathcal{M}$ be a model of $Circ[T_1 \cup T_2; P; Q]$ in $\mathcal{L}(T_1 \cup T_2)$. Then, $\mathcal{M} \models Circ[T_1 \cup \gamma; P; Q]$ and $\mathcal{M} \models \varphi$. Thus, $Circ[T_1 \cup T_2; P; Q] \models \varphi$. ■

From Theorem 3.7 and Theorem 3.10 we get the following theorem.

Theorem 3.11 (Interpolation Between Theories) *Let T_1, T_2 be two theories, P, Q vectors of symbols in $L(T_1) \cup L(T_2)$ such that $P \subseteq L(T_1)$ and $P \cup Q \supset L(T_2)$. Let γ be a set of sentences logically equivalent to $Cn^{L(T_1) \cap L(T_2)}(T_2)$. Then, for every $\varphi \in \mathcal{L}(T_1)$,*

$$Circ[T_1 \cup \gamma; P; Q] \models \varphi \iff Circ[T_1 \cup T_2; P; Q] \models \varphi$$

Theorem 3.12 (Interpolation for Circumscription 2) *Let T be a theory, P, Q vectors of symbols such that $(P \cup Q) \supseteq L(T)$. Let L_2 be a set of nonlogical symbols. Then, there is $\gamma \in \mathcal{L}(T) \cap \mathcal{L}(L_2 \cup P)$ such that $T \models \gamma$ and for all $\varphi \in \mathcal{L}(L_2)$,*

$$Circ[T; P; Q] \models \varphi \iff Circ[\gamma; P; Q] \models \varphi.$$

Furthermore, this γ can be logically equivalent to the consequences of T in $L(T) \cap (L_2 \cup P)$.

PROOF Let T_1 be a set of tautologies such that $L(T_1) = L_2 \cup P$. Also, let $T_2 = T \cup \tau_2$, for τ_2 a set of tautologies such that $L(\tau_2) = Q \setminus L(T_1)$. Let $L_1 = L(T_1)$, $L_2 = L(T_2)$. Theorem 3.11 guarantees that γ from that theorem satisfies $Circ[\gamma; P; Q] \models \psi \iff Circ[T; P; Q] \models \psi$ for every $\psi \in \mathcal{L}_1 = L_2 \cup P$. ■

Example 3.13 *If α, β and P are taken to be as in Example 3.9, and $Q =$ {clear,block,A,B}, then one interpolant for Theorem 3.12 is $\gamma = clear(A) \Leftrightarrow \forall x \neg on(x, A)$.*

The theorems we presented are for parallel circumscription, where we minimize all the minimized predicates in parallel without priorities. The case of prioritized circumscription is outside the scope of this paper.

4 Default Logic

In this section we present interpolation theorems for FOL default logic and propositional default logic. The results for FOL immediately imply the same results for propositional default logic theories. The results that are unique to propositional logic rely on a the additional condition of a finite language.

4.1 Reiter's Default Logic: Overview

In Reiter's default logic [Reiter, 1980] one has a set of facts W (in either propositional or FOL) and a set of defaults D (in a corresponding language). Defaults in D are of the form $\frac{\alpha:\beta_1,...,\beta_n}{\delta}$ with the intuition that if α is proved, and $\beta_1, ..., \beta_n$ are consistent (throughout the proof), then δ is proved. α is called the *prerequisite*, $pre(d) = \{\alpha\}$; $\beta_1, ..., \beta_n$ are the *justifications*, $just(d) = \{\beta_1, ..., \beta_n\}$ and δ is the *consequent*, $cons(d) = \{\delta\}$. We use similar notation for sets of defaults (e.g., $cons(D) = \bigcup_{d \in D} cons(d)$).

Take, for example, the following default theory $T = \langle W, D \rangle$:

$$D = \{\frac{bird(x) : fly(x)}{fly(x)}\} \quad W = \{bird(Tweety)\} \tag{2}$$

The intuition behind this theory is that birds normally fly and that $Tweety$ is a bird.

Definition 4.1 *An* extension *of $\langle W, D \rangle$ is a set of sentences E that satisfies W, follows the defaults in D, and is minimal. More formally, E is an extension if it is minimal (as a set) such that $\Gamma(E) = E$, where we define $\Gamma(S_0)$ to be S, a minimal set of sentences such that*

1. *$W \subseteq S$; $S = Cn(S)$.*
2. *For all $\frac{\alpha:\beta_1,...,\beta_n}{\delta} \in D$ if $\alpha \in S$ and $\forall i \ \neg\beta_i \notin S_0$, then $\delta \in S$.*

The following theorem provides an equivalent definition that was shown in [Marek and Truszczyński, 1993; Risch and Schwind, 1994] and others. A set of defaults, $\mathbb{D}$ is grounded in a set of formulae W iff for all $d \in \mathbb{D}$, $pre(d) \in Cn_{Mon(\mathbb{D})}(W)$, where $Mon(\mathbb{D}) = \{\frac{pre(d)}{cons(d)} \mid d \in \mathbb{D}\}$.

Theorem 4.2 (Extensions in Terms of Generating Defaults) *A set of formulae E is an extension of a default theory $\langle W, D \rangle$ iff $E = Cn(W \cup \{cons(d) \mid d \in D'\})$ for a minimal set of defaults $D' \subseteq D$ such that*

1. *D' is grounded in W and*
2. *for every $d \in D$, $d \in D'$ if and only if $pre(d) \in Cn(W \cup cons(D'))$ and for every $\psi \in just(d)$, $\neg\psi \notin Cn(W \cup cons(D'))$.*

Every minimal set of defaults $D' \subseteq D$ as mentioned in this theorem is said to be a set of *generating defaults*.

Normal defaults are defaults of the form $\frac{\alpha:\beta}{\beta}$. These defaults are interesting because they are fairly intuitive in nature (if we proved α then β is proved unless previously proved inconsistent). A default theory is *normal*, if all of its defaults are normal.

We define $W \mid\!\sim_D \varphi$ as cautious entailment sanctioned by the defaults in D, i.e., φ follows from every extension of $\langle W, D\rangle$. We define $W \mid\!\sim^b_D \varphi$ as brave entailment sanctioned by the defaults in D, i.e., φ follows from at least one extension of $\langle W, D\rangle$.

4.2 Interpolation in Default Logic

In this section we present several flavors of interpolation theorems, most of which are stated for cautious entailment. Most of those theorems use the following lemma.

Lemma 4.3 *Let $T = \langle W, D\rangle$ be a default theory and L a set of nonlogical symbols (possibly empty). Let γ be a set of sentences such that $\gamma \equiv Cn^{\mathcal{L}(D\cup L)\cap W}(W)$. Then, for every extension E' of $\langle \gamma, D\rangle$ there is an extension E of $\langle W, D\rangle$ such that $Cn(E' \cup W) = Cn(E)$.*

PROOF Let γ be the set of consequences of W in $\mathcal{L}(D \cup L) \cap \mathcal{L}(W)$. Let $\mathbb{E}$ be the set of extensions of $\langle W, D\rangle$ and $\mathbb{E}'$ the set of extensions of $\langle \gamma, D\rangle$. We show that every extension $E' \in \mathbb{E}'$ has an extension $E \in \mathbb{E}$ such that $Cn(E' \cup W) = Cn(E)$.

Take $E' \in \mathbb{E}'$ and define $E_0 = Cn(E' \cup W)$. We assume that $L(E') \subseteq L(\mathbb{D})$ because otherwise we can take a logically equivalent extension whose sentences are in $\mathcal{L}(\mathbb{D})$. We show that E_0 satisfies the conditions for extensions of $\langle W, D\rangle$ (Definition 4.1).

1. $W \subseteq E_0$,
2. For all $\frac{\alpha:\beta_1,\ldots,\beta_n}{\delta} \in D$, if $\alpha \in E_0$ and $\forall i\ \neg\beta_i \notin E_0$, then $\delta \in E_0$.

The first condition holds by definition of E_0. The second condition holds because every default that is consistent with E_0 is also consistent with E' and vice versa. We detail the second condition in what follows.

For the first direction (every default that is consistent with E_0 is also consistent with E'), let $\frac{\alpha:\beta_1,\ldots,\beta_n}{\delta} \in D$ be such that $\alpha \in E_0$. We show that $\alpha \in E'$. By definition, $\alpha \in \mathcal{L}(D)$. $\alpha \in E_0$ implies that $E' \cup W \models \alpha$ because $Cn(E' \cup W) = E_0$. By compactness of FOL, there is a finite set of sentences $\mathcal{E} \subseteq E'$ such that $\mathcal{E} \cup W \models \alpha$. Using the deduction theorem we get that $W \models \mathcal{E} \Rightarrow \alpha$ (taking $\mathcal{E}$ here to be the conjunction of a finite set of sentences that is logically equivalent to $\mathcal{E}$ in $\mathcal{L}(\mathbb{D})$). Using Craig's interpolation theorem, there is $\gamma' \in \mathcal{L}(W) \cap \mathcal{L}(\mathcal{E} \Rightarrow \alpha)$ such that $W \models \gamma'$ and $\gamma' \models \mathcal{E} \Rightarrow \alpha$. However, this means that $\gamma \models \gamma'$, by the way we chose γ. Thus $\gamma \models \mathcal{E} \Rightarrow \alpha$. Since $\mathcal{E} \subseteq E' \subseteq \gamma$ we get that $E' \models \alpha$. Since $E' = Cn(E')$ we get that $\alpha \in E'$.

The case is similar for δ: if $\delta \in E_0$ then $\delta \in E'$ by the same argument as given above for $\alpha \in E_0$ implying $\alpha \in E'$. Finally, if $\forall i\ \neg\beta_i \notin E_0$, then $\forall i\ \neg\beta_i \notin E'$ because $E' \subseteq E_0$. The opposite direction (every default that is consistent with E' is also consistent with E_0) is similar to the first one.

Thus, E_0 satisfies those two conditions. However, it is possible that E_0 is not a minimal such set of formulae. If so, Theorem 4.2 implies that there is a strict subset of the generating defaults of E_0 that generates a different extension. However, we can apply this new set of defaults to generate an extension that is smaller than E', contradicting the fact that E' is an extension of $\langle \gamma, D\rangle$. Thus, for every $E' \in \mathbb{E}'$ there is $E \in \mathbb{E}$ such that $Cn(E' \cup W) = Cn(E)$. ■

The following theorem results from Lemma 4.3. It says that there is an interpolant in the intersection of the languages of W and $D \cup \{\varphi\}$, when φ is the consequence we are trying to prove from $\langle W, D\rangle$.

Theorem 4.4 (Interpolation for Cautious DL 1) *Let $T = \langle W, D\rangle$ be a first-order default theory and φ a first-order formula. If $W \mid\!\sim_D \varphi$, then there are γ_1, γ_2 such that $\gamma_1 \subseteq \mathcal{L}(W) \cap \mathcal{L}(D \cup \{\varphi\})$, $\gamma_2 \subseteq \mathcal{L}(W \cup D) \cap \mathcal{L}(\varphi)$ and all the following hold:*

$$W \models \gamma_1 \qquad \gamma_1 \mid\!\sim_D \gamma_2 \qquad \gamma_2 \models \varphi \qquad W \mid\!\sim_D \gamma_2 \qquad \gamma_1 \mid\!\sim_D \varphi$$

PROOF We are going to show that it is enough to have the consequences of W in $\mathcal{L}(D \cup \{\varphi\})$ to provide the proper set of extensions. We then extend them all with W. Let γ_1 be the set of consequences of W in $\mathcal{L}(D \cup \{\varphi\}) \cap \mathcal{L}(W)$. Let $\mathbb{E}$ be the set of extensions of $\langle W, D\rangle$ and $\mathbb{E}'$ the set of extensions of $\langle \gamma_1, D\rangle$. From Lemma 4.3 we know that every extension $E' \in \mathbb{E}'$ has an extension $E \in \mathbb{E}$ such that $Cn(E' \cup W) = Cn(E)$.

Now we show that, if φ logically follows in all the extensions of $\langle W, D\rangle$ then it must also follow from every extension of $\langle \gamma_1, D\rangle$. Take $E' \in \mathbb{E}'$. We know already that $\mathbb{E}' \cup W \models \varphi$. But $E' \supseteq \gamma_1 \supseteq Cn^{\mathcal{L}(\varphi)\cap\mathcal{L}(W)}(W)$ and $Cn^{\mathcal{L}(\varphi)\cap\mathcal{L}(W)}(W)$ is logically equivalent to $Cn^{\mathcal{L}(\varphi)}(W)$. From Craig's interpolation theorem for FOL we get that $E' \cup W \models \varphi$ implies that there is a sentence $\gamma \in \mathcal{L}(\varphi) \cap \mathcal{L}(W \cup E')$ such that $W \cup E' \models \gamma$ and $\gamma \models \varphi$. Since $\gamma \in \mathcal{L}(\varphi) \cap \mathcal{L}(W \cup E')$ we get that $E' \models \gamma$. Thus, $E' \models \varphi$.

Finally, take $\gamma_2 = \bigcap_{E' \in \mathbb{E}'} Cn^{\mathcal{L}(\varphi)}(E')$, Since $L(E') \subseteq L(W \cup D)$, $\gamma_2 \subseteq \mathcal{L}(W \cup D) \cap \mathcal{L}(\varphi)$. Clearly, $\gamma_2 \models \varphi$ and $\gamma_1 \mid\!\sim_D \gamma_2$. $W \mid\!\sim_D \gamma_2$ follows from an argument that is identical to our previous conclusion for φ.

Thus, these γ_1, γ_2 are those promised by the current theorem: $W \models \gamma_1$, $W \mid\!\sim_D \gamma_2$, $\gamma_2 \models \varphi$, $\gamma_1 \mid\!\sim_D \gamma_2$ and $\gamma_1 \mid\!\sim_D \varphi$. ■

The second interpolation theorem differs from the first in two important ways. First, it says that there is an interpolant (of a different kind from the first) in the intersection of the languages of W and D. This interpolant is chosen disregarding the consequence that we are trying to prove from $\langle W, D\rangle$. Secondly, it holds only for the propositional case (or more generally, when there are only finitely many extensions of $\langle W, D\rangle$).

Theorem 4.5 (Interpolation for Cautious DL 2) *Let $T = \langle W, D\rangle$ be a propositional default theory with a finite signature and φ a propositional formula. If $W \mid\sim_D \varphi$, then there are $\gamma_1, \gamma_2 \in \mathcal{L}(W) \cap \mathcal{L}(D)$, and all the following hold:*

$$W \models \gamma_1 \qquad \gamma_1 \mid\sim_D \gamma_2 \qquad \{\gamma_2\} \cup W \models \varphi \qquad W \mid\sim_D \gamma_2$$

PROOF Let γ_1 be the set of consequences of W in $\mathcal{L}(D) \cap \mathcal{L}(W)$. Let $\mathbb{E}$ be the set of extensions of $\langle W, D\rangle$ and $\mathbb{E}'$ the set of extensions of $\langle \gamma_1, D\rangle$. From Lemma 4.3 we know that every extension $E' \in \mathbb{E}'$ has an extension $E \in \mathbb{E}$ such that $Cn(E' \cup W) = Cn(E)$. Clearly, $W \models \gamma_1$.

Since $W \mid\sim_D \varphi$, then every extension of $\langle W, D\rangle$ entails φ. Thus, for every $E' \in \mathbb{E}'$, $E' \cup W \models \varphi$. From the deduction theorem, $E' \models W \Rightarrow \varphi$. Let $\gamma_{E'}$ be the interpolant for this last implication, as guaranteed by Craig's interpolation theorem. Then, $E' \models \gamma_{E'}$ and $\gamma_{E'} \models W \Rightarrow \varphi$ and $\gamma_{E'} \in \mathcal{L}(E') \cap \mathcal{L}(\varphi \cup W) = (\mathcal{L}(D) \cap \mathcal{L}(W)) \cap \mathcal{L}(\varphi)$.

Take $\gamma_2 = \bigvee_{E' \in \mathbb{E}'} \gamma_{E'}$ (we can take γ_2 this way because there are only finitely many extensions of $\langle \gamma_1, D\rangle$). Then, $\gamma_2 \models W \Rightarrow \varphi$ because $\gamma_{E'} \models W \Rightarrow \varphi$ for every $E \in \mathbb{E}$. Thus, $\gamma_1 \mid\sim_D \gamma_2$ and $\{\gamma_2\} \cup W \models \varphi$.

Finally, to show that $W \mid\sim_D \gamma_2$, let $E \in \mathbb{E}$ and assume that $E \not\models \gamma_2$. Let $E' = E \cap \mathcal{L}(D)$. We show that E' is an extension of $\langle \gamma_1, D\rangle$. Since E is an extension of $\langle W, D\rangle$, every default of D holds in E. Thus, every default of D also holds in E'. This is because every formula $\alpha \in \mathcal{L}(D)$ is in E' iff it is in E, by the way we defined E'. If E' is minimal, then we are done because we found an extension of $\langle \gamma_1, D\rangle$ that does not entail γ_2, contradicting our conclusion from the previous paragraph (thus, it cannot be that $E \not\models \gamma_2$).

If E' is not minimal, then there is E'_0 which is smaller than E' and is an extension. However, from Lemma 4.3 we know that there is an extension E_0 of $\langle W, D\rangle$ such that $Cn(E_0) = Cn(W \cup E'_0)$. Since $E'_0 \subsetneq E'$, we get that $E_0 \subseteq E$. Because of minimality of E (it is an extension of $\langle W, D\rangle$) it must be that $E_0 = E$. Thus, $Cn(W \cup E'_0) = Cn(E)$. However, $E \not\models \gamma_2$, by our assumption. Thus, $W \cup E'_0 \not\models \gamma_2$, contradicting our conclusion from above, that every extension of $\langle \gamma_1, D\rangle$ entails γ_2. Thus, E'_0 cannot be an extension of $\langle \gamma_1, D\rangle$ that is smaller than E', so E' is an extension. This is a contradiction as well because then E' is an extension of $\langle \gamma_1, D\rangle$ that does not entail γ_2. Thus, it cannot be that $E \not\models \gamma_2$. This proves that every extension of $\langle W, D\rangle$ entails γ_2, i.e., that $W \mid\sim_D \gamma_2$. ■

It is important to notice that both interpolation theorems that we have just presented are supplying the reverse direction as well. I.e., if there are interpolants that satisfy the granted conditions, then the formula φ follows from $\langle W, D\rangle$. This is immediate from the conditions $W \mid\sim_D \gamma_2$ and $\gamma_2 \models \varphi$ (or $\gamma_2 \cup W \models \varphi$), which appear in both theorems (it follows that $W \mid\sim_D \varphi$ because default logic satisfies Right Weakening [Kraus *et al.*, 1990; Makinson, 1993]).

Example 4.6 *If $\alpha = wet \wedge (wet \Rightarrow mud)$, $\beta = (rain \Rightarrow clouds) \Rightarrow clouds$, $D = \{\frac{mud:rain}{rain}\}$, then one possible pair of interpolants is $\gamma_1 = mud$, $\gamma_2 = rain$.*

Example 4.7 *If the only objects are A, B, and $\alpha = block(A) \wedge block(B) \wedge \forall b\ (clear(b) \Leftrightarrow (block(b) \wedge \forall x \neg on(x,b)))$, $\beta = clear(A)$ and $D = \{\frac{:\neg on(a,b)}{\neg on(a,b)}\}$, then one possible pair of interpolants is $\gamma_1 = TRUE$, $\gamma_2 = \neg on(B,A) \wedge \neg on(A,A)$. To see that γ_1, γ_2 satisfy Theorem 4.5 notice that $\langle \alpha, D \rangle$ has one extension:*

$$E = \left\{ \begin{array}{c} \neg on(A,B), \neg on(B,A), \neg on(A,A), \neg on(B,B), \\ block(A), block(B), clear(A), clear(B) \end{array} \right\}$$

Corollary 4.8 *Let $\langle W, D \rangle$ be a first-order default theory and φ a formula. If $W \mid\sim_D \varphi$, then there is a set of formulae, $\gamma \subseteq \mathcal{L}(W \cup D) \cap \mathcal{L}(\varphi)$ such that $W \mid\sim_D \gamma$ and $\gamma \mid\sim_D \varphi$.*

PROOF Follows if γ_2 in Theorem 4.4 corresponds to our needed γ. ■

We do not get stronger interpolation theorems for prerequisite-free normal default theories. [Imielinski, 1987] provided a modular translation of normal default theories with no prerequisites into circumscription, but Theorem 3.8 does not lead to better results. In particular, the counter example that we presented after Theorem 3.8 can be massaged to apply here too.

Theorem 4.9 (Interpolation Between Default Extensions) *Let $\langle W_1, D_1 \rangle$ and $\langle W_2, D_2 \rangle$ be first-order default theories such that $L(cons(D_2)) \cap L(pre(D_1) \cup just(D_1) \cup W_1) = \emptyset$. Let φ be a first-order formula such that $\varphi \in \mathcal{L}(W_2 \cup D_2)$. If there is an extension E of $\langle W_1 \cup W_2, D_1 \cup D_2 \rangle$ in which φ holds, then there is a set of formulae $\gamma \subseteq \mathcal{L}(W_1 \cup D_1) \cap \mathcal{L}(W_2 \cup D_2)$ and E_1, E_2 such that*

$$\begin{array}{ll} E_1 \in Ext(\langle W_1, D_1 \rangle) & E_2 \in Ext(\langle W_2 \cup \{\gamma\}, D_2 \rangle) \\ Cn(E_1) \cap \mathcal{L}(W_2 \cup D_2) = \gamma & E_2 \models \varphi \end{array}$$

PROOF Let $D_1' \subseteq D_1$ be the set of generating defaults of E that belong to D_1. Notice that these defaults are grounded in W_1 because there is no information that may have come from applying the rest of the generating defaults in E (we required that $cons(D_2) \cap (pre(D_1) \cup just(D_1) \cup W) = \emptyset$). Let E_1 be the extension of $\langle W_1, D_1 \rangle$ defined using the generating defaults in D_1'.

Let $\gamma \subseteq \mathcal{L}(W_1 \cup D_1) \cap \mathcal{L}(W_2 \cup D_2)$ be equivalent to the set of sentences in that language that follow from E_1. Let $D_2' \subseteq D_2$ be the set of generating defaults of E that belong to D_2. Notice that these defaults are grounded in $W_2 \cup \gamma$ because there is no information that D_1' may contribute that is not already in γ (we required that $cons(D_2) \cap (pre(D_1) \cup just(D_1) \cup W) = \emptyset$). Let E_2 be the extension of $\langle W_2 \cup \gamma, D_2 \rangle$ defined with the generating defaults in D_2'.

Now, $E_1 \cup E_2 \equiv E$, and γ is the set of sentences that follow from E_1 in $\mathcal{L}(E_1) \cap \mathcal{L}(E_2 \cup \varphi)$. $E_2 \models \varphi$ because of Craig's interpolation theorem (Theorem 2.1): $E_1 \cup E_2 \models \varphi$ implies that $E_1 \models E_2 \Rightarrow \varphi$, and Craig's interpolation theorem guarantees the existence of $\gamma' \in \mathcal{L}(E_1) \cap \mathcal{L}(E_2 \cup \varphi)$ such that $E_1 \models \gamma'$ and $\gamma' \models E_2 \Rightarrow \varphi$. Thus, $\gamma' \in \gamma$ and $\gamma \models E_2 \Rightarrow \varphi$. This implies $\gamma \cup E_2 \models \varphi$ which implies that $E_2 \models \varphi$. ■

It is interesting to notice that the reverse direction of this theorem does not hold. For example, if we have two extensions E_1, E_2 as in the theorem

statement, it is possible that E_1 uses a default with justification β, but $W_2 \models \neg\beta$. Strengthening the condition of the theorem, i.e., demanding that $L(W_2 \cup cons(D_2)) \cap L(pre(D_1) \cup just(D_1) \cup W_1) = \emptyset$, is not sufficient either. For example, if D_1 includes two defaults $d_1 = \frac{:}{a \Rightarrow \neg\beta}$, and $d_2 = \frac{:\beta}{\varphi}$, $W_1 = \emptyset$, D_2 includes no defaults and $W_2 = \{a\}$ then there is no extension of $\langle W_1 \cup W_2, D_1 \cup D_2\rangle$ that implies φ, for $\varphi = \{c\}$.

Further strengthening the conditions of the theorem gives the following:

Theorem 4.10 (Reverse Direction of Theorem 4.9 (Some Cases)) *Let $\langle W_1, D_1\rangle, \langle W_2, D_2\rangle$ be propositional default theories such that $L(D_1 \cup W_1) \cap L(W_2 \cup cons(D_2)) = \emptyset$. Let φ be a propositional formula such that $\varphi \in \mathcal{L}(W_2 \cup D_2)$. There is an extension E of $\langle W_1 \cup W_2, D_1 \cup D_2\rangle$ in which φ holds, if there is a formula $\gamma \in \mathcal{L}(W_1 \cup D_1) \cap \mathcal{L}(W_2 \cup D_2)$, an extension E_1 of $\langle W_1, D_1\rangle$ such that $Cn(E_1) \cap \mathcal{L}(W_2 \cup D_2) = \gamma$, and an extension E_2 of $\langle W_2 \cup \{\gamma\}, D_2\rangle$ such that $E_2 \models \varphi$.*

PROOF Let E_1, E_2 be as in the statement of the theorem. Let π_1, π_2 be the sets of defaults applied in E_1, E_2, respectively. Let $E = Cn(E_1 \cup E_2)$, and let $\pi = \pi_1 \cup \pi_2$. We show that E is an extension of $\langle W_1 \cup W_2, D_1 \cup D_2\rangle$ such that $E \models \varphi$ as needed.

First, $E \supseteq W_1 \cup W_2$ by E's definition. For every $d = \frac{\alpha:\beta_1,\ldots,\beta_n}{\delta} \in \pi$, α and δ hold because they hold in one of E_1, E_2 ($d \in \pi = \pi_1 \cup \pi_2$). Assume that $\beta_i \in E$ for some $i \leq n$. Then, $E_1 \cup E_2 \models \neg\beta_i$, implying that $E_1 \models E_2 \Rightarrow \neg\beta_i$ (we treat E_1, E_2 here as finite sets of formulae because they are in propositional logic).

If $d \in \pi_1$, then clearly $\neg\beta_i \in \mathcal{L}(W_1 \cup D_1)$. $E_1 \cup \{\beta_i\} \models \neg E_2$. Using Craig's interpolation theorem we get that γ from the theorem's statement satisfies $\{\gamma\} \cup \{\beta_i\} \models \neg E_2$. Consequently, $E_2 \models \gamma \Rightarrow \neg\beta_i$. However, $E_2 = Cn(cons(\pi_2) \cup W_2 \cup \{\gamma\})$, and $L(W_2 \cup cons(D_2)) \cap L(W_1 \cup D_1) = \emptyset$. This means that $L(cons(\pi_2) \cup W_2) \cap L(\gamma \wedge \beta_i) = \emptyset$, implying that $\gamma \models \neg\beta$, contradicting $\neg\beta \notin E_1$.

If $d \in \pi_2$, then $\neg\beta_i \in \mathcal{L}(W_2 \cup D_2)$. Since $E_1 \models E_2 \Rightarrow \neg\beta_i$, we get from Craig's interpolation theorem that $\gamma \models E_2 \Rightarrow \neg\beta_i$. However, $\gamma \in E_2$ by the definition of E_2. Thus, $E_2 \models \neg\beta$, contradicting the fact that d is a default applied in E_2.

In conclusion, $\neg\beta \notin E$. Thus, all the defaults in π are applied in E. It is also simple to see that no other default is applied in E.

If there is a default $d = \frac{\alpha:\beta_1,\ldots,\beta_n}{\delta} \in D_1$ that should apply in E but is not in π, then its preconditions and justifications hold in E. However, this means that α follows from $E_1 \cup E_2$, and $E_2 = Cn(cons(\pi_2) \cup W_2 \cup \{\gamma\})$. Similar to the argument above we get that $E_1 \models \alpha$. Similarly, we get that if $\beta_i \notin E$ then $E_1 \not\models \neg\beta_i$, implying that d should have applied in E_1, contradicting the fact that E_1 is an extension of $\langle W_1, D_1\rangle$.

If there is a default $d = \frac{\alpha:\beta_1,\ldots,\beta_n}{\delta} \in D_2$ that should apply in E but is not in π, then its preconditions and justifications hold in E. A similar argument to the one above shows that it should have applied in E_2 too, contradicting the fact that E_2 is an extension.

Minimality of E follows from that of E_1, E_2. Thus, E is an extension of $\langle W_1 \cup W_2, D_1 \cup D_2\rangle$ as needed. ■

Corollary 4.11 (Interpolation for Brave DL) *Let $\langle W_1, D_1\rangle, \langle W_2, D_2\rangle$ be propositional default theories such that $L(cons(D_2)) \cap L(pre(D_1) \cup just(D_1) \cup W_1) = \emptyset$. Let φ be a formula such that $\varphi \in \mathcal{L}(W_2 \cup D_2)$. If $W_1 \cup W_2 \mid\sim^b_{D_1 \cup D_2} \varphi$, then there is a formula, $\gamma \in \mathcal{L}(W_1 \cup D_1) \cap \mathcal{L}(W_2 \cup D_2)$, such that $W_1 \mid\sim^b_{D_1} \gamma$ and $W_2 \cup \{\gamma\} \mid\sim^b_{D_2} \varphi$.*

Corollary 4.11 does not hold for the cautious case (where we look at all the extensions and choose φ and γ satisfied by all of them): Let $W_1 = W_2 = \emptyset$, $D_1 = \{\frac{:b}{b}, \frac{:\neg b}{\neg b}\}$ and $D_2 = \{\frac{b:c}{c}, \frac{\neg b:c}{c}\}$. There are two extensions, in both of which c is proved, but $W_1 \mid\sim_{D_1} \gamma$ only for $\gamma \equiv TRUE$.

Better interpolation theorems may hold (e.g., theorems that do not depend on $cons(D_2)$, $pre(D_1)$, etc., and provide $\gamma \in \mathcal{L}(\langle W_1, D_1\rangle) \cap \mathcal{L}(\langle W_2, D_2\rangle))$, if we consider the entailment between two default theories ($\langle W_1, D_1\rangle \mid\sim \langle W_2, D_2\rangle$). These are outside the scope of this paper.

Finally, Corollary 4.11 and Theorem 4.9 are similar to the *splitting theorem* of [Turner, 1996], which is provided for default theories with $W = \emptyset$ (there is a modular translation that converts every default theory to one with $W = \emptyset$). A *splitting set* for a set of defaults D is a subset A of $L(D)$ such that $pre(D), just(D), cons(D) \subseteq \mathcal{L}(A) \cup \mathcal{L}(L(D) \setminus A)$ and $\forall d \in D$ $(cons(d) \notin \mathcal{L}(L(D) \setminus A) \Rightarrow L(d) \subseteq A)$. Let $B = L(D) \setminus A$. The base of D relative to A is $b_A(D) = \{d \in D \mid L(d) \subseteq A)$. For a set of sentences $X \subseteq \mathcal{L}(A)$, we define $e_A(D, X)$ to be

$$\left\{ \frac{\bigwedge(\{a_i\}_{i\leq n} \cap \mathcal{L}(B)) : \{b_i\}_{i\leq m} \cap \mathcal{L}(B)}{c} \;\middle|\; \begin{array}{l} \frac{\bigwedge_{i=1}^{n} a_i : b_1, \ldots, b_m}{c} \in D \setminus b_A(D), \\ \forall i \leq n (a_i \in \mathcal{L}(A) \Rightarrow a_i \in Cn^A(X)), \\ \forall i \leq m (\neg b_i \notin Cn^A(X)) \end{array} \right\}$$

Theorem 4.12 ([Turner, 1996]) *Let A be a splitting set for a default theory D over $\mathcal{L}(U)$. A set E of formulae is a consistent extension of D iff $E = Cn^{L(D)}(X \cup Y)$, for some consistent extension X of $b_a(D)$ over $\mathcal{L}(A)$ and Y a consistent extension of $e_A(D, X)$ over $\mathcal{L}(L(D) \setminus A)$.*

Roughly speaking, this theorem finds an extension X of the *base* ($b_A(D)$) and converts $D \setminus b_A(D)$ using this X into a theory $e_A(D, X)$. Then, an extension Y for $e_A(D, X)$ completes the extension for D if $X \cup Y$ is consistent. In contrast, our theorem does not change $D \setminus b_A(D)$.

5 Logic Programs

In this section we provide interpolation theorems for logic programs with the stable models semantics. We use the fact that logic programs are a special case of default logic, and the results are straightforward. An *extended disjunctive logic program* [Gelfond and Lifschitz, 1988; Gelfond and Lifschitz, 1990; Gelfond and Lifschitz, 1991; Przymusiński, 1991] is a set of *rules*. Each rule, r, is written as an expression of the form

$$L_1 | \ldots | L_l \leftarrow A_1, \ldots, A_n, notB_1, \ldots, notB_m$$

where $L_1, ..., L_l, A_1, ..., A_n, B_1, ..., B_m$ are literals, that is, atomic formulae or their (classic) negations, $L_1, ..., L_l$ are the *head literals*, $head(r)$, $A_1, ..., A_n$ are the *positive subgoals*, $pos(r)$, and $B_1, ..., B_m$ are the *negated subgoals*, $neg(r)$.

A program P is *positive* if none of its rules includes negated subgoals. A set of literals, X, is *closed* under a positive program, P, if, for every rule $r \in P$ such that $pos(r) \subseteq X$, $head(r) \cap X \neq \emptyset$. A set of literals is *logically closed* if it consistent or contains all literals. An *answer set* for a positive program, P is a minimal set of literals that is both closed under P and logically closed.

For an arbitrary logic program, P, and a set of literals, X, we say that X is an *answer set* for a program P if X is an answer set for P^X, where P^X is defined to include a rule r' iff $neg(r') = \emptyset$ and there is $r \in P$ such that $head(r') = head(r)$, $pos(r') = pos(r)$, and $neg(r) \cap X = \emptyset$.

Logic programs with answer-set semantics were shown equivalent to default logic in several ways. Normal rules are rules of the form

$$A \leftarrow B_1, ..., B_m, \ not\ C_1, ..., \ not\ C_n$$

where $A, B_1, ..., B_m, C_1, ..., C_n$ are atoms (i.e., no disjunction or classic negation is allowed). For normal rules, [Gelfond and Lifschitz, 1991] translated every normal rule of the form $A \leftarrow B_1, ..., B_m, \ not\ C_1, ..., \ not\ C_n$, into a default

$$\frac{B_1 \wedge ... \wedge B_m \ : \ \neg C_1, ..., \neg C_n}{A}.$$

Under this mapping, the stable models of a logic program coincide with the extensions of the corresponding default theory (Facts in the logic program are translated to facts in the default theory, while rules are translated to defaults).

[Sakama and Inoue, 1993] showed that disjunctive logic programs (no classic negation) with the stable model semantics can be translated to prerequisite-free default theories as follows:

1. For a rule $A_1|...|A_l \leftarrow B_1, ..., B_m, \ not\ C_1, ..., \ not\ C_n$ in P, we get the default
$$\frac{: \neg C_1, ..., \neg C_n}{B_1 \wedge ... \wedge B_m \Rightarrow A_1 \vee ... \vee A_l}$$
2. For each atom A appearing in P, we get the default $\frac{: \neg A}{\neg A}$

Each stable model of P is the set of atoms in some extension of D_P, and the set of atoms in an extension of D_P is a stable model of P (notice that, in general, an extension of D_P can include sentences that are not atoms and are not subsumed by atoms in that extension). [Sakama and Inoue, 1993] provide a similar translation to extended disjunctive logic programs by first translating those into disjunctive logic programs (a literal $\neg A$ is translated to a new symbol, A'), showing that a similar property holds for this class of programs.

We define $P \mid\sim \varphi$ as cautious entailment sanctioned from the logic program P, i.e., φ follows from stable model of P. We define $P \mid\sim^b \varphi$ as brave entailment sanctioned from the logic program P, i.e., φ follows from at least one stable model of P.

From the last translation above we get the following interpolation theorems.

Theorem 5.1 (Interpolation for Stable Models (Cautious)) *Let P be a logic program and let φ be a formula such that $P \mid\!\sim \varphi$. Then, there is a formula $\gamma \in \mathcal{L}(P) \cap \mathcal{L}(\varphi)$ such that $P \mid\!\sim \gamma$ and $\gamma \models \varphi$.*

PROOF Follows if γ_2 in Theorem 4.4 corresponding to our needed γ. ■

Theorem 5.2 (Interpolation for Stable Models (Brave)) *Let P_1, P_2 be logic programs such that $head(P_2) \cap body(P_1) = \emptyset$. Let $\varphi \in \mathcal{L}(P_2)$ be a formula such that $P_1 \cup P_2 \mid\!\sim^b \varphi$. Then, there is a formula $\gamma \in \mathcal{L}(P_1) \cap \mathcal{L}(P_2)$ such that $P_1 \mid\!\sim^b \gamma$ and $\gamma \cup P_2 \mid\!\sim^b \varphi$.*

PROOF Follows from the reduction to default logic and Corollary 4.11. ■

The last theorem is similar to the *splitting theorem* of [Lifschitz and Turner, 1994]. This theorem finds an answer set X of the *bottom* (P_1) and converts P_2 using this X into a program P_2'. Then, an answer set Y for P_2' completes the answer set for $P_1 \cup P_2$ if $X \cup Y$ is consistent. In contrast, our theorem does not change P_2, but it is somewhat weaker, in that it does only provides a necessary condition for $P_1 \cup P_2 \mid\!\sim^b \varphi$ (this is the typical form of an interpolation theorem).

6 Summary

We presented interpolation theorems that are applicable to the nonmonotonic systems of circumscription, default logic and Answer Set Programming (a.k.a. Stable Models Semantics). These results are somewhat surprising and revealing in that they show particular structure for the nonmonotonic entailments associated with the different systems. They promise to help in reasoning with larger systems that are based on these nonmonotonic systems.

Several questions remain open. First, γ promised by our theorems is not always finite (in the FOL case). This is in contrast to classical FOL, where the interpolant is always of finite length. What conditions guarantee that it is finite in our setup? We conjecture that this will require the partial order involved in the circumscription to be *smooth* [Kraus *et al.*, 1990]. Second, are there better interpolation theorems for the prioritized case of those systems? Also, what is the shape of the interpolation theorems specific for prerequisite-free semi-normal defaults?

Finally, the theorems for default logic and Logic Programming promise that $\alpha \mid\!\sim_D \beta$ implies the existence of γ such that $\alpha \mid\!\sim_D \gamma$ and $\gamma \mid\!\sim_D \beta$. Currently, we get the reverse direction (concluding that $\alpha \mid\!\sim_D \beta$) only if we take γ to be equivalent to a large set of consequences of W. Can we find weaker conditions under which this result holds? Also, can we do better than Theorem 4.10 for different cases?

7 Acknowledgments

I wish to thank Esra Erdem, Vladimir Lifschitz and Leora Morgenstern for reading and commenting on different drafts of this manuscript. This research was supported by DARPA grant N66001-00-C-8018 (RKF program) and National Science Foundation grant ECS-9873474.

References

[Amir and McIlraith, 2000] Eyal Amir and Sheila McIlraith. Paritition-based logical reasoning. In *Principles of Knowledge Representation and Reasoning: Proc. Seventh Int'l Conference (KR '2000)*, pages 389–400. Morgan Kaufmann, 2000.

[Amir, 2000] Eyal Amir. (De)Composition of situation calculus theories. In *Proc. National Conference on Artificial Intelligence (AAAI '00)*, pages 456–463. AAAI Press/MIT Press, 2000.

[Antoniou, 1997] Grigoris Antoniou. *Nonmonotonic Reasoning.* MIT Press, Cambridge, Massachusetts, 1997.

[Brewka *et al.*, 1997] Gerd Brewka, Jürgen Dix, and Kurt Konolige. *Nonmonotonic Reasoning: An Overview*, volume 73 of *CSLI Lecture Notes.* CSLI Publications, 1997.

[Brewka, 1991] Gerhard Brewka. *Nonmonotonic Reasoning: Logical Foundations of Common Sense.* Cambridge University Press, 1991.

[Craig, 1957] William Craig. Linear reasoning. a new form of the herbrand-gentzen theorem. *Journal of Symbolic Logic*, 22:250–268, 1957.

[Darwiche, 1997] Adnan Darwiche. A logical notion of conditional independence. *Artificial Intelligence*, 97(1–2):45–82, 1997.

[Darwiche, 1998] A. Darwiche. Model-based diagnosis using structured system descriptions. *Journal of Artificial Intelligence Research*, 8:165–222, 1998.

[Dechter and Pearl, 1988] Rina Dechter and Judea Pearl. Tree Clustering Schemes for Constraint Processing. In *Proc. National Conference on Artificial Intelligence (AAAI '88)*, 1988.

[Dechter and Rish, 1994] Rina Dechter and Irina Rish. Directional resolution: The davis-putnam procedure, revisited. In *Principles of Knowledge Representation and Reasoning: Proc. Fourth International Conference (KR '94)*, pages 134–145. Morgan Kaufmann, 1994.

[Dechter, 1999] Rina Dechter. Bucket elimination: A unifying framework for reasoning. *Artificial Intelligence*, 113(1–2):41–85, 1999.

[Etherington, 1986] David Etherington. *Reasoning with incomplete Information*. PhD thesis, University of British Columbia, 1986.

[Gabbay *et al.*, 1993] D.M. Gabbay, C.J Hogger, and J.A.Robinson, editors. *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 3: Nonmonotonic Reasoning and Uncertain Reasoning*. Oxford University Press, Great Britain, 1993.

[Gelfond and Lifschitz, 1988] Michael Gelfond and Vladimir Lifschitz. The stable model semantics for logic programming. In *Logic Programming: Proceedings of the Fifth International Conference and Symposium*, pages 1070–1080, 1988.

[Gelfond and Lifschitz, 1990] Michael Gelfond and Vladimir Lifschitz. Logic Program with Classical Negation. In David H. D. Warren and Peter Szeredi, editors, *Proceedings of the 7th Int. Conf. on Logic Programming*, pages 579–597. MIT, June 1990.

[Gelfond and Lifschitz, 1991] Michael Gelfond and Vladimir Lifschitz. Classical Negation in Logic Programs and Disjunctive Databases. *New Generation Computing*, 9, 1991.

[Hodges, 1997] Wilfrid Hodges. *A shorter model theory*. Cambridge University Press, 1997.

[Imielinski, 1987] Tomasz Imielinski. Results on translating defaults to circumscription. *Artificial Intelligence*, 32(1):131–146, April 1987.

[Kraus *et al.*, 1990] Sarit Kraus, Daniel Lehmann, and Menachem Magidor. Nonmonotonic Reasoning, Preferential Models and Cumulative Logics. *Artificial Intelligence*, 44(1):167–207, 1990.

[Lifschitz and Turner, 1994] Vladimir Lifschitz and Hudson Turner. Splitting a logic program. In *Proc. Eleventh Int'l Conf. on Logic Programming*, pages 23–37. MIT Press, 1994.

[Lifschitz, 1985] Vladimir Lifschitz. Computing circumscription. In *Proc. of IJCAI-85*, pages 121–127, 1985.

[Lifschitz, 1993] Vladimir Lifschitz. Circumscription. In D.M. Gabbay, C.J Hogger, and J.A Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 3: Nonmonotonic Reasoning and Uncertain Reasoning*. Oxford University Press, 1993.

[Makinson, 1993] David Makinson. General patterns in nonmonotonic reasoning. In D.M. Gabbay, C.J.Hogger, and J.A.Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming, Volume 3: Nonmonotonic Reasoning and Uncertain Reasoning*. Oxford University Press, 1993.

[Marek and Truszczyński, 1993] Viktor M. Marek and Mirek Truszczyński. *Nonmonotonic Logics; Context-Dependent Reasoning*. Springer Verlag, Berlin-Heidelberg-New York, 1st edition, 1993.

[McCarthy, 1980] John McCarthy. Circumscription—A Form of Non-Monotonic Reasoning. *Artificial Intelligence*, 13:27–39, 1980.

[McCarthy, 1986] John McCarthy. Applications of Circumscription to Formalizing Common Sense Knowledge. *Artificial Intelligence*, 28:89–116, 1986.

[McIlraith and Amir, 2001] Sheila McIlraith and Eyal Amir. Theorem proving with structured theories. In *Proc. Seventeenth International Joint Conference on Artificial Intelligence (IJCAI '01)*, pages 624–631. Morgan Kaufmann, 2001.

[Przymusiński, 1991] Teodor Przymusiński. Stable semantics for disjunctive programs. *New Generation Computing*, 9:401–424, 1991.

[Reiter, 1980] Raymond Reiter. A logic for default reasoning. *Artificial Intelligence*, 13 (1–2):81–132, 1980.

[Risch and Schwind, 1994] V. Risch and Camila Schwind. Tableau-based characterization and theorem proving for default logic. *Journal of Automated Reasoning*, 13:223–242, 1994.

[Sakama and Inoue, 1993] Chiaki Sakama and Katsumi Inoue. Relating disjunctive logic programs to default theories. In *Logic Programming and Nonmonotonic Reasoning (LPNMR'93)*, pages 266–282, 1993.

[Turner, 1996] Hudson Turner. Splitting a default theory. In *Proc. National Conference on Artificial Intelligence (AAAI '96)*, pages 645–651. AAAI Press, 1996.