Syntactic Conditional Closures for Defeasible Reasoning

James P Delgrande
School of Computing Science,
Simon Fraser University,
Burnaby, B C ,
Canada V5A 1S6
email jim@cs sfu ca

Abstract

An approach to nonmonotonic inference based on a closure operation on a conditional knowledge base is presented The central idea is that, given a theory of default conditionals, an extension to the theory le defined that satisfies certain intuitive restrictions Two notions for forming an extension are given, corresponding to the incorporation of irrelevant properties in conditionals and of transitivity among conditionals, in this approach these notions coincide Several equivalent definitions for an extension are developed general nonconstructive definitions, and a general "pseudo-iterative" definition Reasoning with irrelevant properties is correctly handled, as is specificity, reasoning within exceptional circumstances, and inheritance reasoning Tina approach is intented to ultimately serve as the proof-theoretic analogue to an extant semantic development based on preference orderings among possible worlds

1 Introduction

A continuing problem in Artificial Intelligence is dealing with general, generic sentences that admit exceptions For example, suppose we are given that birds fly birds have wings, penguins are birds, and penguins don't fly We can write this in a propositional gloss[1] as

$$B \Rightarrow F,\ B \Rightarrow W,\ P \Rightarrow B,\ P \Rightarrow \neg F \tag{1}$$

In nonmonotonic reasoning with such statements there are several principles that one would want to hold According to the principle of *specificity,* a more specific default should apply over a less specific default Thus given that *P* is true, we would want to conclude *->F by* default, since being a penguin is a more specific notion than that of being a bird Second, one should obtain transitivity of default conditionals, or *inheritance* of properties by default So, given that *P* is true, one would also want to conclude by default that *W* was true, and so penguins have wings by virtue of being birds Third, *irrelevant* properties should be properly handled and so, all other things being equal, we would want to conclude that a green bird flies Fourth, one should be able to reason in the presence of *exceptions,* and so given that *B, ~>F* was true one would still conclude that W was true

It has proven difficult to specify an approach that achieves all of these properties Earlier systems (such as Circumscription and Default Logic) provide accounts of general mechanisms for nonmonotonic inference that generally handle inheritance and relevance well, but do not deal with specificity Rather it is up to the user to hand-code specificity information For example in the naive representation of the above example in Default Logic, if *P* is true we obtain a set of default conclusions in which *->F* is true and another bet in which *F* is true

Rpcently much attention has been paid to *conditional* systems of default inferencing The starting point here is not so much to provide a general mechanism for nonmonotonic inference, as it is to provide a *theory* of some phenomenon, based on intuitions concerning, for example, possibility exceptionalness, or nonmonotonic consequence operators The initial systems were, on the whole, quite weak, subsequent work has focussed on means to extend the system's basic inferences As discussed in the next section, while these systems generally handle specificity well, none satisfactorily handles all of inheritance, relevance, and reasoning in exceptional circumstances

This paper explores approaches to syntactically closing off a conditional knowledge base The idea is to formalise directly notions corresponding to the incorporation of irrelevant properties and, separately, to the inheritance of properties Essentially then "proof-theoretic" approaches to default reasoning are investigated, based on intuitions regarding desirable properties for a nonmonotonic reasoning system I begin with a specific conditional logic capable of expressing relations among defaults So, given our initial example, we can derive the fact that birds are normally not penguins *(\ e B = ->P)* This logic however does not address the problems of irrelevant properties or inheritance Thus we cannot conclude that a green bird flies, nor that a penguin has wings by "transitivity" We address this by defining an *extension* of a theory, or a plausible "extended" default theory In the first place, an extension includes condi-

[1] That IS to say we will eventually have to deal with first-order issues but do not do so here For this paper *B => F* can be given the reading "if x-is-a-bird then normally x-flies"

tionals incorporating irrelevant properties In the second place, an extension includes conditionals obtained from (default) transitivities among conditionals, and so accounts for inheritance

The appropriateness of the approach is argued in a number of ways First, the approach formalises plausible, familiar intuitions concerning properties of default inference Second, various canonical and non-canonical examples are shown to be appropriately handled here Third, it proves to be the case that incorporating irrelevant properties corresponds exactly to incorporating transitivity of conditionals Hence the same (nonmonotonic) phenomenon underlies these seemingly distinct notions The approach then is of independent interest, nonetheless, it is intended ultimately to serve as a proof-theoretic analogue to an extant semantic approach, based on preference orderings among possible worlds The next section discusses related work In Section 3 the logic of defaults on which the present approach is based is presented Section 4 gives the formal details of the approach This is followed by a discussion and a concluding section Proofs of theorems and further details are found in [Delgrande, 1995]

2 Background

Related Work Many of the earlier systems of default reasoning deal with mechanisms for effecting nonmonotonic reasoning Autoepistemic Logic [Moore, 1985], Circumscription [McCarthy, 1980], Default Logic [Reiter, 1980], and, for our purposes. Theorist [Poole, 1988] are examples of such approaches In this group, issues of specificity are not addressed within the system, although other properties such as inheritance, are ade quately handled Various modifications have been proposed to handle specificity in these systems but these modifications are build on top of the system Without a formal theory, it is not clear if such modifications are appropriate or in any sense complete Since defaults per se are not part of the formal system one also cannot reason *about* defaults Thus, one could not conclude from our initial example that birds are normally not penguins

Recently, much attention has been paid to conditional systems of default inferencing Such systems address specific forms of nonmonotonic inference, or deal with specific defeasible conditionals based, for example, on notions of preference or exceptionalness There has been a remarkable convergence or agreement on what constitutes a core set of inferences that ought to be common to all nonmonotonic systems Systems such as ent-entailment [Pearl, 1988] (or 0-entailment or p-entailment [Adams, 1975]), possibihstic logic [Dubois and Prade 1994], preferential entailment [Kraus *et al*, 1990], and *CTA* [Boutiher, 1992], among others, essentially allow the same inferences, and may be taken as specifying a *conservative core* [Pearl, 1989] that ought to be common to all nonmonotonic inference systems These approaches deal satisfactorily with specificity However, not unexpectedly, they are much too weak In particular, relevance and inheritance of properties are not handled Hence, even though a bird may be assumed to fly by default, a green bird cannot be assumed to fly by default (since there may be models of a theory where a green bird does not fly)

Equally surprising, there has also been a strong convergence on a means of strengthening these systems Approaches including System Z and 1-entailment [Pearl, 1990], *CO'* [Boutilier. 1992], possibilistic entailment [Benferhat *et al,* 1992], and rational closure [kraus *et al*, 1990] all assume, in a semantic sense, that a world is as unexceptional as possible Thus essentially, since there is no reason to suppose that greenness has any bearing on flight, one assumes that greenness has no effect on flight While this assumption seems reasonable enough, its realisation in these systems is not unproblematic, as described below In brief these approaches fail to allow full inheritance of properties, as well they allow unwanted specificity relations In the next subsection, System Z is described as a represent ative of these approaches These approaches have been extended in various wavs, including [Goldszmidt and Pearl, 199] Goldszmidt *et al*, 1990, Benferhat *et al*, 1993] Howevtr these extensions allow the unwanted specificities found in the original approaches)

Of other work, [Delgrande, 1988] gives an iterative strengthening of defaults using meta-theoretic assumptions However, since it provides an iterative procedure for strengthening it is difficult to formally characterise the set of default inferences This approach is strictly subsumed by the first approach described here [Geffner and Pearl, 1992] presents another strengthening of the above-mentioned systems called *conditional entailment* There are two difficulties with this approach first, it is quite complex in its formulation and, second, it does not sanction full inheritance of default properties

System Z This subsection describes System **Z** as a representative of the set of systems mentioned above for forming the closure of a conditional knowledge base As well as being simple to describe, System **Z** has the advantage that its description does not rely on prior knowledge of any system corresponding to the "conservative core" In System **Z**, a set of rules R representing defaults is partitioned into an list of mutually exclusive sets of rules R_0 , R_n One begins with a set $R = \{r \mid \alpha_r \rightarrow \beta_r\}$ where each α_r and β_r are propositional formulas A set $R' \subseteq R$ *tolerates* a rule r if $\{\alpha_r \wedge \beta_r\} \cup R'$ is satisfiable From this, an ordering on the rules in R is defined

1 Find all rules tolerated by R and call this set R_0

2 Next, find all rules tolerated by $R - R_0$, and call this set R_1

3 Continue in this fashion until all rules have been accounted for

We obtain a partition $(R_0, \quad R_n)$ of R, where $R_i = \{r \mid r$ is tolerated by $R - R_0 - \quad - R_{i-1}\}$ The *rank* of rule r, written $\mathbf{Z}(r)$, is given by $\mathbf{Z}(r) = i$ iff $r \in R_i$ For our example, we obtain the ordering

$$R_0 = \{B \rightarrow F \; B \rightarrow W\}, \; R_1 = \{P \rightarrow B, P \rightarrow \neg F\}$$

An interpretation M is given a **Z**-rank, $\mathbf{Z}(M)$, according to the highest ranked rule it falsifies $\mathbf{Z}(M) = min\{ n \mid M \models \alpha_r \supset \beta_r, \; \mathbf{Z}(r) \geq n\}$ The rank of a formula φ is defined as the lowest $\mathbf{Z}$-rank of all interpretations satisfying φ $\mathbf{Z}(\varphi) = min\{\mathbf{Z}(M) \mid M \models \varphi\}$ A form of default entailment *1-entailment*, is defined by

$$\varphi \vdash_1 \phi \text{ iff } \mathbf{Z}(\varphi \wedge \phi) < \mathbf{Z}(\varphi \wedge \neg\phi)$$

This gives a form of default inference that has some very nice properties Irrelevant facts are handled well and for example we have $B \wedge G \vdash_1 F$, so green birds fly There are two weaknesses with this approach First, one cannot inherit properties across exceptional subclasses So one cannot conclude that penguins have wings (since $\neg W$ will be true at some least-ranked P world) Second, undesirable specificities may be obtained consider where we add to our example the default that calm animals (C) have low blood pressure (L) Intuitively, $C \rightarrow L$ is irrelevant to the other defaults yet we obtain the default conclusion that calm animals aren't penguins, since $\mathbf{Z}(C \wedge \neg P) < \mathbf{Z}(C \wedge P)$ The first difficulty is addressed in [Benferhat *et al*, 1993], the second is not Moreover, this seems to be a problem endemic to all such approaches

3 A Logic for Default Properties

The approach is founded on a specific *conditional logic* This logic corresponds to an extension of the aforementioned "conservative core" for default inferences [2] See [Delgrande, 1987, Boutilier, 1992] for details The fundamental idea is straightforward possible worlds are arranged according to a notion of "exceptionalness", a default $\alpha \Rightarrow \beta$ is true just when there is a world in which $\alpha \wedge \beta$ is true and $\alpha \supset \beta$ is true at all worlds that are not more exceptional Thus, "birds fly", $B \Rightarrow F$, is true *if, in the least B-worlds,* $B \supset F$ *is true* Intuitively, we disregard exceptional circumstances such as being a penguin, having a broken wing, etc

More formally, our language is that of propositional logic augmented with a binary operator $\Rightarrow$ We reserve $\supset$ for material implication For readability (especially in Section 4) conjunction is sometimes represented by juxtaposition, hence $A \wedge B$ and AB stand for the same formula For simplicity there are no nested occurrences of the $\Rightarrow$ operator Sentences are interpreted in terms of a *model* $M = \langle W, E, P\rangle$ where

1 W is a set (of worlds),
2 E is a binary *accessibility* relation on worlds, with properties

   Reflexive Eww for every $w \in W$

   Transitive If Ew_1w_2 and Ew_2w_3 then Ew_1w_3

   Forward Connected If Ew_1w_2 and Ew_1w_3 then Ew_2w_3 or Ew_3w_2

3 P is a mapping of atomic sentences and worlds onto $\{true, false\}$

Truth at a world w in model M ($\models^M_w$) is as for propositional logic, except that

[2]There are two reasons for adopting a slightly stronger system First this system includes (in the author's opinion) other essential relations governing generic statements (most importantly that $\alpha \Rightarrow \gamma \supset ((\alpha \wedge \beta \Rightarrow \gamma) \vee (\alpha \wedge \neg\beta \Rightarrow \gamma))$ is valid) Second it has a simpler semantics However the results presented here apply equally to these weaker systems

$\models^M_w \alpha \Rightarrow \beta$ iff there is a w_1 such that Ew, w_1 and $\models^M_{w_1} \alpha \wedge \beta$ and for every w_2 where Ew_1, w_2 we have $\models^M_{w_2} \alpha \supset \beta$

In addition, $\Box\alpha$ ("necessarily α") is defined as $\neg\alpha \Rightarrow \alpha$, and $Th(T)$ is the set of formulas logically entailed by default theory T

The accessibility relation between worlds is defined so that from a particular world w one "sees" a sequence of successively "less exceptional" sets of worlds For $\alpha \Rightarrow \beta$ to be true we require that there be a world in which α is true, and so we exclude the case of a vacuously-true default [3] This has no effect on the expressibility of the language (and indeed the following exposition can be developed allowing vacuously-true conditionals), but it does simplify things somewhat

This logic supplies us with a weak, but semantically justified, system of default inferencing Given a set of defaults and necessitations T, β follows by default from α just when $\alpha \Rightarrow \beta$ is true in all models of T Thus

Definition 3 1
β follows from α by default in T iff $T \models \alpha \Rightarrow \beta$

Hence (as previously discussed) we can conclude that a penguin does not fly, while a bird does, and if something flies then it is not a penguin However we cannot conclude that green birds fly, since there are models in which green birds do not fly Nor can we conclude that penguins, by virtue of being birds, have wings In both these cases it seems that, plausibly, we would want to make the given inference, from green bird to fly, and from penguin to winged In the next section we describe two approaches for extending a default theory, wherein relevance and inheritance are properly handled

4 The Approach

In what follows, I will identify a default theory T with a set of default conditionals (corresponding to assertions such as "birds fly") together with a set of assertions taken to be necessarily true (corresponding to assertions such as "penguins are necessarily birds") A default theory then contains information that is taken to be true across all domains of interest Contingent world knowledge (such as "individual x flies") is not part of a default theory, but rather represents case-specific information

The logic of Section 3 supplies us with a means of determining the relative *specificity* of default conditionals β is *strictly more specific* than α, written $\alpha \prec \beta$, can be defined as follows

Definition 4 1 $T \models \alpha \prec \beta$ *iff* $T \models \alpha \vee \beta \Rightarrow \neg\beta$

The right side of the definition says that for every model of T, at some $\alpha \vee \beta$ world, w, $\neg\beta$ is true, and $\alpha \vee \beta \supset \neg\beta$ is true at all equally- or less-exceptional worlds Thus α is true at w, and there are no equivalently-exceptional or less-exceptional worlds in which β is true We have that $\prec$ is irreflexive, asymmetric, and transitive, also, $\prec$ and $\Rightarrow$ are interdefinable [Lewis, 1973]

$\alpha \preceq \beta$ is defined as $\neg(\beta \prec \alpha)$ Using this, the following notion of specificity is defined

[3]Thus this conditional corresponds to the "would" conditional of [Lewis, 1973]

Definition 4 2
β is more specific than α in theory T, written $\alpha \ll_T \beta$ iff $T \models \alpha \preceq \beta$ and it is not the case that $T \models \beta \preceq \alpha$

The general idea of the approach is straightforward Beginning with a default theory T, we "appropriately' extent this theory, so that we obtain "reasonable" conclusions in the extension Default reasoning is given as in Definition 3 1, but with respect to extensions Thus, since birds fly and we have no reason to believe that a green bird does not fly, we would include $B \wedge Gr \Rightarrow F$ in an extension However there is a good reason to believe that a bird that is a green penguin does not fly, namely that penguins do not fly, and so we would not include $B \wedge P \wedge Gr \Rightarrow F$ The issue then is to distinguish "reasonable" from "unreasonable" additions to a theory

In this regard, there are two common notions for augmenting a default theory These are strengthening the antecedent of a conditional, which I will call *STR*, and transitivity among conditionals, or *TRANS* Informally, these notions correspond to the incorporation of irrelevant properties in a conditional and inheritance of properties, respectively In what follows I first develop a definition for STR Following this, a definition for TRANS is given and the two notions are shown to be equivalent

For STR, there are three properties that can be expected to hold for an extension E of default theory T

1 It should contain the original theory
2 It should be logically closed
3 If $\alpha \Rightarrow \beta \notin T$ then $\alpha \Rightarrow \beta \in E$ iff
   (a) There is a reason to accept the conditional
   (b) There is no equally strong (or stronger) reason to not accept the conditional

Thus in the case of green penguins there 16 a reason to suppose that such an animal flies (since birds fly) and a reason to suppose that it does not (since penguins don't fly) Since the notion of penguinhood is more specific than that of being a bird, we would conclude that a green penguin does not fly by default

More formally, we have the following definition

Definition 4 3 (STR)
An extension E from strengthening a *default theory T is a minimal set of formulas such thai*

1 $T \subseteq E$,
2 $Th(E) = E$, *and*
3 *If $\alpha' \Rightarrow \beta \in E$ then $\alpha \wedge \alpha' \Rightarrow \beta \in E$ if $\not\models \alpha \wedge \alpha' \supset \bot$ and whenever $\models \alpha \wedge \alpha' \supset \alpha''$ and $\alpha'' \Rightarrow \neg\beta \in E$ then $\alpha'' \ll_E \alpha'$*

For default transitivity, we can again identify three properties that can be reasonably expected to hold for any extension E In addition to Properties 1 and 2

3 If $\alpha \Rightarrow \gamma \notin T$ but $\alpha \Rightarrow \beta, \beta \Rightarrow \gamma \in E$ then $\alpha \Rightarrow \gamma \in E$ if
   (a) There is a reason to accept $\alpha \Rightarrow \gamma \in E$
   (b) There is no equally strong or stronger reason to not accept $\alpha \Rightarrow \gamma \in E$

Hence, appropriately realised, these properties would permit the inference that a penguin has wings, based on the fact that penguins are birds and birds have wings We have the definition

Definition 4 4 (TRANS)
An extension E from transitivity *of a default theory T is a minimal set of formulas such that*

1 $T \subseteq E$,
2 $Th(E) = E$,
3 *If $\alpha \Rightarrow \alpha', \alpha' \Rightarrow \beta \in E$ then $\alpha \Rightarrow \beta \in E$ if whenever $\alpha \Rightarrow \alpha'' \in E$ and $\alpha'' \Rightarrow \neg\beta \in E$ then $\alpha'' \ll_E \alpha'$*

In this case, for the third part, $\alpha \Rightarrow \beta \in E$ if

1 there is a reason to accept the conditional via transitivity we have conditionals $\alpha' \Rightarrow \alpha \in E$ and $\alpha \Rightarrow \beta \in E$, and
2 if there are also formulas $\alpha' \Rightarrow \alpha''$ and $\alpha'' \Rightarrow \neg\beta$ in E, then α'' is less specific than α' That is, we would conclude α' and α'' from α However, since $\alpha'' \ll_F \alpha'$ we would conclude β and not $\neg\beta$

Informally, this definition allows us to assert that "birds fly based on the defaults "birds normally have wings" and "winged things normally fly" Hence birds inherit the property of flight from that of winged things

Interestingly, these two notions for constructing an extension are equivalent

Theorem 4 1 *E is an extension of T from strengthening iff E is an extension of T from transitivity*

In light of Theorem 4 1, we need not distinguish extensions obtained from strengthening versus extensions obtained from transitivity Consequently I will simply refer to *extensions* of a default theory, and will just make use of Definition 4 3 [4]

The next theorem provides a quasi-iterative specification for an extension While this may be of limited use in constructing an extension (since the extension appears in the specification), it does provide a means of *verifying* whether a set of formulas is an extension of a default theory

Theorem 4 2 *Let T and E be sets of conditionals Define*

$$E_0 = T$$

and for every $i \geq 0$

$$E_{i+1} = Th(E_i) \cup \{\alpha \Rightarrow \beta \mid \not\models \alpha \supset \bot \text{ and (i) there is } \alpha' \text{ such that } \models \alpha \supset \alpha' \text{ where } \alpha' \Rightarrow \beta \in E_i, \text{ and (ii) if } \models \alpha \supset \alpha'' \text{ and } \alpha'' \Rightarrow \neg\beta \in E \text{ then } \alpha'' \ll_E \alpha' \}$$

Then E is an extension of T iff $E = \bigcup_{i=0}^{\infty} E_i$

[4] It is worth noting that for any sufficiently strong system of defeasible inferencing these notions coincide, in that any system that "reasonably" implements STR will also obtain TRANS and vice versa See [Delgrande, 1995] for details

Default inference is defined as follows

Definition 4 5
β follows from α by default in theory T, $\alpha \mid\!\sim_T \beta$,[5] iff for every extension E of T, we have $E \models \alpha \Rightarrow \beta$

The remainder of this section presents examples and discusses properties of the approach The first example is a variant on an example presented at the outset of the paper

Example 4 1
$T = \{B \Rightarrow W, W \Rightarrow F, P \Rightarrow B, P \Rightarrow \neg F\}$

So birds have wings, winged things fly, penguins are nonflying birds The extension E of T contains the following sentences

$$B \wedge Gr \Rightarrow F, \quad B \wedge Gr \wedge P \Rightarrow \neg F, \quad P \Rightarrow W$$

In the notation of Definition 4 5 this could be written as

$$B \wedge Gr \mid\!\sim_T F, \quad B \wedge Gr \wedge P \mid\!\sim_T \neg F, \quad P \mid\!\sim_T W$$

In the first formula, birds fly by transitivity, and greenness is irrelevant, so green birds fly Green is also irrelevant with respect to a penguin s being able to fly, and so green penguins don't fly While penguins are exceptional birds, in that they do not fly nonetheless penguins have wings by default

Note that there are two sources of 'nonmonotonicity" The first is in a sense illusory, deriving from the properties of the conditional logic itself For the above example we have $B \Rightarrow F, B \wedge P \Rightarrow \neg F \in Th(T)$ and so $B \mid\!\sim_T F$, and $B \wedge P \mid\!\sim_T \neg F$ That is despite appearances, we are here dealing with an underlying monotonic system However there is a second source, deriving from the definition of an extension Thus if T is as above and $T' = \{B \Rightarrow F\}$, then we have $T' \subseteq T$, and $B \wedge P \mid\!\sim_{T'} F$ but $B \wedge P \mid\!\sim_T \neg F$

Example 4 2 $T = \{Q \Rightarrow P, R \Rightarrow \neg P\}$

This is a familiar example involving conflicting defaults There is one extension E of T, we have, among other conclusions $Q \wedge \neg R \Rightarrow P, Q \Rightarrow \neg R \in E$

Example 4 3
$T = \{A_1 \Rightarrow B_1, A_2 \Rightarrow B_2, A_3 \Rightarrow B_3, \Box\neg(B_1 B_2 B_3)\}$

This is an extension of the previous example there are three defaults which cannot be simultaneously applied Again there is only one extension The formula $A_1 A_2 \Rightarrow B_1$ is in the extension, as is $A_1 A_2 \Rightarrow B_1 B_2$ Clearly we would not want $A_1 A_2 A_3 \Rightarrow B_1$ to be in the extension, since by symmetry we would also have the same antecedent support B_2 and B_3 But these statements together are only trivially satisfiable Note that in order to block the presence of $A_1 A_2 A_3 \Rightarrow B_1 \in E$ based on $A_1 \Rightarrow B_1 \in T$ we use the fact that $A_2 A_3 \Rightarrow B_2 B_3 \in E$, whence $A_2 A_3 \Rightarrow \neg B_1 \in E$ However we do have that the following formula is in the extension

$$A_1 A_2 A_3 \Rightarrow (B_1 B_2 \neg B_3) \vee (B_1 \neg B_2 B_3) \vee (\neg B_1 B_2 B_3)$$

[5] By rights α should be a *set* of formulas For simplicity I will sometimes abuse notation and write a formula to the left of $\mid\!\sim_T$

That is, given $A_1 A_2 A_3$ we have that two of B_1, B_2, B_3 hold, but we don't know which Furthermore we have that

$$A_1 A_2 A_3 \neg B_3 \Rightarrow B_1 B_2 \text{ and } A_1 A_2 \neg B_2 A_3 \neg B_3 \Rightarrow B_1$$

are in the extension Thus, in the presence of falsified conditionals, we obtain the desired conclusions Finally if we were to add the information that $A_2 \prec A_1$ and $A_3 \prec A_1$ to T, we would obtain $A_1 A_2 A_3 \Rightarrow B_1$ in the extension (in fact we also obtain $A_1 A_2 A_3 \Rightarrow B_1(B_2 \vee B_3)$)

There are default theories where there is more than one extension Perhaps the simplest such theory is the following

Example 4 4
$T = \{B_1 \Rightarrow C_1, B_2 \Rightarrow C_2, A \Rightarrow (B_1 \neg C_1) \vee (B_2 \neg C_2)\}$

The third default asserts that one of the first two does not hold, however it does not specify which does not hold We obtain one extension which contains $A \wedge B_1 \Rightarrow C_1$ (and $A \wedge B_2 \Rightarrow \neg C_2$) and a second which contains $A \wedge B_2 \Rightarrow C_2$ (and $A \wedge B_1 \Rightarrow \neg C_1$) Both extensions contain $A \Rightarrow \neg(B_1 \supset C_1) \equiv (B_2 \supset C_2)$

Extensions have quite reasonable properties The following theorems are fundamental

Theorem 4 3 *If T is a default theory then T has an extension*

Theorem 4 4 *For extension E of T, E is inconsistent iff T is inconsistent*

Theorem 4 5 *If E_1, E_2 are extensions of T and $E_1 \neq E_2$ then $E_1 \cup E_2 \models \bot$*

Thus extensions exist, and are inconsistent only if the underlying theory is inconsistent Further, if there is more than one extension of a theory then the extensions are mutually inconsistent

A default theory may be ambiguous or incomplete with respect to specificity in the following sense

Definition 4 6 *T is* ambiguous with respect to specificity *if there are γ_1, γ_2, α such that $T \models \gamma_1 \vee \gamma_2 \prec \alpha$ but $T \not\models \gamma_1 \prec \alpha$ and $T \not\models \gamma_2 \prec \alpha$*

We obtain

Theorem 4 6
No extension is ambiguous with respect to specificity

Thus specificity ambiguities are resolved in one fashion or another in forming an extension In Example 4 4 we have $T \models B_1 \vee B_2 \prec A$ but $T \not\models B_1 \prec A$ and $T \not\models B_2 \prec A$ In one extension we obtain $B_1 \prec A$ and in another $B_2 \prec A$ In Example 4 2, on the other hand, $T \models Q \vee R \prec Q \wedge R$ but $T \not\models Q \prec Q \wedge R$ and $T \not\models R \prec Q \wedge R$ In the resulting (single) extension we obtain both $Q \prec Q \wedge R$ and $R \prec Q \wedge R$

Finally, for a conditional belonging to an extension, either a given property is irrelevant with respect to the conditional or its negation is

Theorem 4 7
If $\alpha \Rightarrow \beta \in E$ then $\alpha \wedge \gamma \Rightarrow \beta \in E$ or $\alpha \wedge \neg\gamma \Rightarrow \beta \in E$

Thus, given that $B \Rightarrow F$ is in an extension, then either $B \wedge P \Rightarrow F$ is in the extension or $B \wedge \neg P \Rightarrow F$ is Also presumably we would obtain that both $B \wedge Gr \Rightarrow F$ and $B \wedge \neg Gr \Rightarrow F$ are in the extension

5 Discussion

Comparison with Related Work The technical development of the last section is somewhat reminiscent of that of Default Logic [Reiter, 1980], and it is worthwhile outlining the similarities and differences In Default Logic (DL) a default theory is given by a pair *(D,W)* where *D* is a set of default rules (expressed as domain-specific "rules of inference") and *W* is a set of classical (first-order or propositional) formulas, giving information about the domain at hand An extension is a maximal acceptable set of beliefs, obtained by "applying" as many rules from *D* (given *W)* as possible An extension then is a superset of *W*

In our case the object is to extend the set of default conditionals, in DL this would correspond to augmenting *D* An extension in the present approach then is intended to be "applicable' to *any* domain of application (1 e to any "W")

In both approaches three properties are given that can be expected to hold for any extension

1 The thing extended *(W* in DL, *T* here) should he contained in an extension

2 The extension should be deductively closed

3 As many "rules" as possible should be applied or added

In both cases a nonconstructive definition of an exten won is given the present approach does not require a fixed-point definition as DL does (roughly) due to the comparatively restricted form of the conditional defaults

With respect to more recent and more closely-related work, the present approach can be compared to ap proaches exemplified by System Z [Pearl, 1990][6] As mentioned in Section 2, there arc two weaknesses with these systems one cannot inherit properties across exceptional subclasses, and undesirable specificities, are sometimes obtained The first difficulty has been rectified in, for example [Benferhat *et al*, 1993] However since this system subsumes System Z the second problem remains The difficulty here (and the difference with the approach at hand) is that System Z assumes that things are as "unexceptional' (l e ranked as low) as consistently possible Hence for our birds and penguins example, any irrelevant conditional (e g "students like cheap but good restaurants'') is ranked at level 0, and consequently lower than any conditional involving penguins This, as described previously, leads to undesirable specificities In the approach at hand, unrelated conditionals remain unrelated and so there are no direct inferences between penguins and students

As mentioned, the approach of strengthening presented here strictly subsumes one of the approaches given in [Delgrande, 1988], where an iterative means of augmenting a theory is given

6'Recall from Section 2 that systems such as CO" [Boutiher 1992] possibilistic entailment [Benferhat el *al*, 1992], or rational closure [Lehmann and Magidor, 1992] are essentially equivalent with respect to derivability

Further work There are two major areas for further work The first concerns a semantics for the approach, while the second deals with implementation issues

The assertion was made at the outset that this approach is of independent interest Nonetheless, it has been formulated with a specific semantic development in mind [Delgrande, 1994] In this latter development, a conditional in a theory *T prefers* a world in which the classical counterpart of the conditional is true, over a world in which it is false That is, $B \Rightarrow F$ prefers a world in which $B \supset F$ is true over one in which $B \wedge \neg F$ is true Following on this intuition, a set of conditionals *T* determines an ordering (or ordenngs) on a set of possible worlds In this approach *(3* follows by default from ct just if B is Lrue in the most preferred worlds in which ft is true It is conjectured that these approaches are equivalent they were formulated with this equivalence in mind but it remains to be shown that they are equivalent

With respect to implementation concerns, although the definitions for an extension involve the addition of an infinite number of defaults, it is clear that an implementation would not need to contend with the full set of defaults in an extension [7] Rather, in determining whether alpha |—7- B given a theorem prover for the logic (see [Lamarre 1992] for example), the issue is to find a set of supported conditionals which when added to *T* allows or => *0* to be derived Second, we might expect that when the form of sentences in a theory is appropriately restricted (say, to something looking like Horn clauses) we would obtain good complexity results

6 Conclusion

This paper has presented an approach for extending a default theory so as to obtain a system for default inferencing via the incorporation of irrelevant properties or inheritance of properties The approach takes as a starting point an extant (monotonic) logic of defaults From this, given a default theory, definitions for an ex tension were given corresponding to intuitions regarding irrelevant properties and inheritance of properties In addition to the logic of defaults described in Section 3, the results of this paper apply also to the "conservative core" for default inferences Arguably the notion of an extension satisfactorily formalises intuitions concern ing augmenting a default theory That the approach captures a reasonable notion of "extension' is also supported by the fact that the concepts of irrelevance and inheritance prove to be interdefinable The approach also handles standard and non-standard examples of default reasoning, and avoids the difficulties encountered by other approaches Finally, it is conjectured that an extant semantic approach is equivalent to the present work

The approach is intended to be applicable to defeasible conditionals, or to generic sentences which allow possible exceptions Consequently it will not be, nor

[T]Any more so than in classical logic, in determining whether T l- o, one would first, form the deductive closure of

IS it intended to be, applicable in all areas For example, in this approach defaults are implicitly "applied" wherever possible Consequently, given a chain of defaults $T = \{A_1 \Rightarrow A_2, A_2 \Rightarrow A_3 \quad , A_{n-1} \Rightarrow A_n\}$ lacking contradictory information we would obtain that $A_1 \mid\!\sim_T A_n$ This may be fine for default reasoning, but it leads to unintuitive results for temporal reasoning, as has been noted for example for *chronological ignorance* Hence the present approach would appear to produce results too strong for such reasoning

There are two principal areas for further work First a provably equivalent semantics, presumably based on preference orderings among worlds, is obviously desirable Second, computational concerns have not been addressed However the immediate goal of this paper has been to characterisation proof-theoretic approaches to default inference Given this, computational issues can subsequently be addressed

Acknowledgements

This research was funded by the Natural Science and Engineering Research Council of Canada grant A0884 The author was a visitor at IRISA, Rennes, France while part of work was being carried out The author wishes to thank Philippe Besnard and Repco for their hospitality, and Torsten Schaub for comments on the initial draft

References

[Adams, 1975] EW Adams *The Logic of Conditionals* D Reidel Publishing Co , Dordrecht, Holland, 1975

[Benferhat *et al,* 1992] Sal em Benferhat Didier Dubois, and Henri Prade Representing default rules in possibilisitc logic In *Third International Conference on Principles of Knowledge Representation and Reasoning,* pages 673-684, Cambridge Massachusetts, October 1992

[Benferhat *et al*, 1993] Salem Benferhat, Claudette Cayrol, Didier Dubois, Jerome Lang, and Henri Prade Inconsistency management and prioritized syntax-based entailment In *Proc IJCAI 93,* pages 640-645, Chambery, Fr , 1993

[Boutiher, 1992] Craig Boutiher *Conditional Logics for Default Reasoning and Belief Revision* PhD thesis, Department of Computer Science, University of Toronto, 1992

[Delgrande, 1987] J P Delgrande A first order conditional logic for prototypical properties *Artificial Intelligence,* 33(1) 105-130,1987

[Delgrande, 1988] J P Delgrande An approach to default reasoning based on a first-order conditional logic Revised report *Artificial Intelligence,* 36(1) *63-90,* 1988

[Delgrande, 1994] J P Delgrande A preference-based approach to default reasoning Preliminary report In *American Association for Artificial Intelligence Conference,* Seattle, WA, July 1994

[Delgrande, 1995] J P Delgrande A preference-based approach to default reasoning In preparation, 1995

[Dubois and Prade, 1994] D Dubois and H Prade Possibilistic Jogje In Dov M Gabbay, C J Hogger, and J A Robinson, editors, *Nonmonotonic Reasontng and Uncertain Reasoning,* volume 3 of *Handbook of Logic in Artifical Intelligence and Logic Programming,* pages 439-513 Oxford Clarendon Press, 1994

[Geffner and Pearl, 1992] Hector Geffner and Judea Pearl Conditional entailment Bridging two approaches to default reasoning *Artificial Intelligence,* 53(2-3) 209-244, 1992

[Goldszmidt and Pearl, 1991] Moises Goldszmidt and Judea Pearl System-Z+ A formalism for reasoning with variable-strength defaults In Proc *AAAI-9J,* pages 399-404, Anaheim, CA, 1991

[GoldBzrmdt *et al*, 1990] Moises Goldszmidt, Paul Morns and Judea Pearl A maximum entropy approach to nonmonotonic reasoning In Proc *AAAI-90,* Boston, MA, 1990

[kraus *et al* , 1990] S Kraus, D Lehmann, and M Magidor Nonmonotonic reasoning, preferential models and cumulative logics *Artificial Intelligence,* 44(1-2) 167-207, 1990

[Lamarre, 1992] Philippe Lamarre A promenade from monotomcity to non-monotomcity following a theorem prover In *Third International Conference on Principles of Knowledge Representation and Reason ing* pages 572-580, Cambridge, Massachusetts, October 1992

[Lehmann and Magidor, 1992] D Lehmann and M Magidor What does a conditional knowledge base entail"[7] *Artificial Intelligence,* 55(1) 1-60,1992

[Lewis, 1973] D Lewis *Counterj'actuals* Harvard University Press, 1973

[McCarthy, 1980] J McCarthy Circumscription - a form of non-monotomc reasoning *Artificial Intelligence,* 13 27-39,1980

[Moore, 1985] R C *Moore* Semantical considerations on nonmonotonic logic *Artificial Intelligence,* 25 75-94, 1985

[Pearl, 1988] J Pearl *Probabilistic Reasontng* in *Intelligent Systems Networks of Plausible Inference* Morgan Kaufman, San Mateo, CA, 1988

[Pearl, *1989]* J Pearl Probabilistic semantics for nonmonotonic reasoning A survey In Proc *KR-89,* pages 505-516, Toronto, May 1989 Morgan Kaufman

[Pearl, 1990] J Pearl System Z A natural ordering of defaults with tractable applications to nonmonotonic reasoning In *Proc of the Third Conference on Theorettcal Aspects of Reasoning About Knowledge,* pages 121-135, Pacific Grove, Ca , 1990

[Poole, 1988] D L Poole A logical framework for default reasoning *Artificial Intelligence,* 36(1)27-48, 1988

[Reiter, 1980] R Reiter A logic for default reasoning *Artificial Intelligence,* 13 81-132, 1980