Default Reasoning with Specificity

Phan Minh Dung* Tran Cao Son†

* Department of Computer Science and Information Managemen,
School of Advanced Technology, Asian Institute of Technology
PO Box 2754, Bangkok 10501, Thailand
dung@cs.ait.ac.th

† Department of Computer Science,
University of Texas at El Paso
El Paso, Texas 79968, USA
tson@cs.utep.edu

Abstract. We present a new approach to reasoning with specificity which subsumes inheritance reasoning. The new approach differs from other approaches in the literature in the way priority between defaults is handled. Here, it is context sensitive rather than context independent as in other approaches. We show that any context independent handling of priorities between defaults as advocated in the literature until now is not sufficient to capture general defeasible inheritance reasoning. We propose a simple and novel argumentation semantics for reasoning with specificity taking the context-dependency of the priorities between defaults into account. Since the proposed argumentation semantics is a form of stable semantics of nonmonotonic reasoning, it inherits a common problem of the later where it is not always defined for every default theory. We identify the class of stratified default theories for which the argumentation semantics is always defined and show that acyclic defeasible inheritance networks are stratified. We also prove that the argumentation semantics satisfies the basic properties of a nonmonotonic consequence relation such as deduction, reduction, conditioning, and cumulativity for stratified default theories.

1 Introduction

Reasoning with specificity constitutes an inseparable part of default reasoning as specificity is an important source for conflict resolution in human's commonsense reasoning. In fact, the famous example of whether penguins fly because they are birds [23] in default reasoning is an example of reasoning with specificity. Reasoning with specificity also constitutes a difficult problem which has been studied extensively in the literature [1, 5, 7, 14, 16, 29, 31, 38].

Formally a default theory T could be defined as a pair (E, K) where E is a set of evidence or facts representing what we call the concrete context of T, and $K = (D, B)$ constitutes the domain knowledge consisting of a set of default rules D and a first order theory B representing the background knowledge. In the literature [1, 5, 7, 14, 16, 29] the principle of reasoning with specificity is "enforced" by first determining a set of priority orders between defaults in D using the information given by the domain knowledge K. Based on these priorities between defaults and following some sensible and intuitive criteria, the semantics of T is then defined either model-theoretically by selecting a subset of the set of all models of $E \cup B$ as the set of preferred models of T or proof-theoretically by selecting certain extensions as preferred extensions. The problem of these approaches is that their obtained semantics is rather weak. They do not capture general defeasible inheritance reasoning. There are many intuitive examples of reasoning with specificity (see below) that can not be handled in these approaches. The reason is that the priorities between defaults are defined independent of the context.

Priority orders are partial orders [1] between defaults in D. Let PO_K be the set of all these priority orders. For each priority order $\alpha \in PO_K$, where $(d, d') \in \alpha$ means that d is of lower priority than d', a priority order $<_\alpha$ between the sets of defaults in D is defined where $S <_\alpha S'$ means that S is preferred to S'. There are many ways to define $<_\alpha$ [1, 5, 7, 14, 16, 29, 31]. But whatever the definition of $<_\alpha$ is, $<_\alpha$ has to satisfy the following property.

Let d, d' be two defaults in D such that $(d, d') \in \alpha$. Then $\{d'\} <_\alpha \{d\}$.

$<_\alpha$ can be extended into an partial order between models of $B \cup E$ as follows:

$M <_\alpha M'$ iff $D_M <_\alpha D_{M'}$

where D_M is the set of all defaults in D which are satisfiable in M whereas a default $p \rightarrow q$ is said to be satisfiable in M iff the material implication $p \Rightarrow q$ is satisfiable in M.

A model M of $B \cup E$ is defined as a preferred model of T if there exists a partial order α in PO_K such that M is minimal with respect to $<_\alpha$. We then say that a formula β is defeasibly derived from T if β holds in each preferred model of T.

In a previous paper [11], we formally proved that any preferential semantics based on $<_\alpha$ can not account in full for defeasible inheritance reasoning. We include below this proof for the self-containment of the paper.

Example 1. Let us consider the following defeasible inheritance network[2]

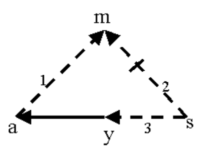

[1] Partial orders are transitive, irreflexive and antisymmetric relations

[2] *Throughout the paper, solid lines and dotted lines represent strict rules and default rules respectively.*

where the links $s \not\to m$, $a \to m$, and $s \to y$ represent the normative sentences "normally, students are not married", "normally, adults are married", and "normally, students are young adults", respectively, and, the strict link $y \Rightarrow a$ represents the subclass relation "young adults are adults."

This defeasible inheritance network represents the domain knowledge (B, D) with $B = \{y \Rightarrow a\}$, and $D = \{d_1 : \ a \to m, \ d_2 : \ s \to \neg m, \ d_3 : \ s \to y\}$.

Consider now the marital status of a young adult who is also a student. This problem is represented by the default theory $T = (E, B, D)$ with $E = \{s, y, a\}$. The desirable semantics here is represented by the model $M = \{s, y, a, \neg m\}$. To deliver this semantics, all priority-based approaches in the literature *[1, 5, 7, 16, 29]* assigns default 1 a lower priority than default 2.

Let us consider now the marital status of another student who is an adult but not a young one. Let $T' = (E', B, D)$ with $E' = \{s, \neg y, a\}$. Now, since y does not hold, default 2 can not be considered more specific than default 1. Hence, it is intuitive to expect that neither m nor $\neg m$ should be concluded in this case. This is also the result sanctioned by all semantics of defeasible inheritance networks [17, 18, 37, 41, 39, 40]. In any priority-based system employing the same priorities between defaults with respect to E' as with respect to E, we have $M = \{\neg m, s, \neg y, a\} <_\alpha M' = \{m, s, \neg y, a\}$ since $D_M = \{2\} <_\alpha D'_M = \{1\}$ (due to $(1, 2) \in \alpha$). That means priority-based approaches in the literature conclude $\neg m$ given (E', K) which is not the intuitive result we expect.

To produce a correct semantics, 1 should have lower priority than 2 only in the context $\{s, y, a\}$ where the considered student is young (i.e. default 3 can be applied). In other words, the priority order under the context $\{s, \neg y, a\}$ is different than the priority order under the context $\{s, y, a\}$. In general, the example shows that *specificity cannot be treated independently from the context* in which it is defined. $\square$

Argumentation has been recognized lately as an important and natural approach to nonmonotonic reasoning [6, 9, 14, 30, 32, 33, 38, 45]. It has been showed in [9] that many major nonmonotonic logics [22, 23, 25, 28, 34] represent in fact different forms of a simple system of argumentation reasoning. Based on the results in [9], a simple logic-based argumentation system has been developed in [6] which captures well-known nonmonotonic logics like autoepistemic logics, Reiter's default logics and logic programming as special cases. In [14] argumentation has been employed to give a proof procedure for conditional logics. In [38], an argumentation system for reasoning with specificity has been developed. Like the proposals based on context-independent priorities, this system is rather weak. It can not deal with many intuitive examples and also fails to capture inheritance reasoning. It does not satisfy many basic properties of defeasible reasoning like the cumulativity. But despite these shortcomings, works like [38] suggests that argumentation offers a natural and intuitive framework for dealing with specificity. As we will show in this paper, argumentation indeed provides a simple and intuitive framework for reasoning with specificity.

In this paper, we extend the approach to reasoning with specificity [11] to allow *general default theories*. In the process, we simplify the notion of *more* *specific.* We propose a simple and novel argumentation semantics for reasoning with specificity taking the context-dependency of the priorities between defaults into account. We then identify a large class of stratified default theories for which the argumentation semantics is always defined and show that acyclic and consistent inheritance networks are stratified. We prove that the argumentation semantics satisfies the basic properties of a nonmonotonic consequence relation such as deduction, reduction, conditioning, and cumulativity for stratified default theories.

2 A General Framework

We assume a first order language $\mathcal{L}$ that is finite but large enough to contain all constants, function and predicate symbols of interest. The set of ground literals of $\mathcal{L}$ is denoted by $lit(\mathcal{L})$. Literals of $\mathcal{L}$ will be called hereafter simply as literals (or $\mathcal{L}$-literals) for short. Following the literature, a default theory is defined as follows:

Definition 1. *A default theory T is a triple (E, B, D) where*
(i) E is a set of ground literals representing the evidences of the theory;
(ii) B is a set of ground clauses;
(iii) D is a set of defaults of the form $l_1 \wedge \ldots \wedge l_n \rightarrow l_0$ where l_i's are ground literals; and
(iv) $E \cup B$ is a consistent first order theory.

Notice that in the above definition, we use $\rightarrow$ to denote a default implication. The material implication is represented by the $\Rightarrow$ symbol. Intuitively, $a \rightarrow b$ means that "typically, if a holds then b holds" while $a \Rightarrow b$ means that "whenever a holds then b holds." It is worth noting that default theories considered in [11] do not contain ground clauses, i.e., $B = \emptyset$.

For a default $d \equiv l_1 \wedge \ldots \wedge l_n \rightarrow l_0$, we denote $l_1 \wedge \ldots \wedge l_n$ and l_0 by $bd(d)$ and $hd(d)$ respectively.

In the following, we often use clauses and defaults with variables as a shorthand for the sets of their ground instantiations.

Example 2. Consider the famous penguin and bird example:

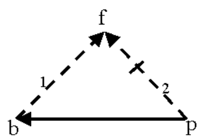

We have that $B = \{p \Rightarrow b\}$ (penguins are birds) and D consisting of two defaults $p \rightarrow \neg f$ (normally, penguins do not fly) and $b \rightarrow f$ (normally, birds fly)[3].

The question is to determine whether penguins fly. This problem is epresented by the default theory $T = (E, B, D)$ where $E = \{p\}$. $\square$

We next define the notion of *defeasible derivation*.

Definition 2. *Let $T = (E, B, D)$ be a default theory and l be a ground literal.*
• *A sequence of defaults $d_1, \dots, d_n$ $(n \geq 0)$ is said to be a defeasible derivation of l if following conditions are satisfied:*

1. $n = 0$ and $E \cup B \vdash l$ where the relation $\vdash$ represents the first-order consequence relation, or

2. (a) $E \cup B \vdash bd(d_1)$; and

(b) $E \cup B \cup \{hd(d_1), \dots, hd(d_i)\} \vdash bd(d_{i+1})$; and

(c) $E \cup B \cup \{hd(d_1), \dots, hd(d_n)\} \vdash l$.

• *We say l is a* possible consequence of E with respect to B and a set of defaults *$K \subseteq D$, denoted by $E \cup B \vdash_K l$, if there exists a defeasible derivation $d_1, \dots, d_n$ of l such that for all $1 \leq i \leq n$, $d_i \in K$.* $\square$

For a set of literals L we write $E \cup B \vdash_K L$ iff $\forall l \in L: \ E \cup B \vdash_K l$.

We write $E \cup B \vdash_K \perp$[4] iff there is an atom a such that both $E \cup B \vdash_K a$ and $E \cup B \vdash_K \neg a$ hold.

For the default theory from example 1, it is easy to check that $E \cup B \vdash_{\{s \rightarrow \neg m\}} \neg m$ and $E \cup B \vdash_{\{s \rightarrow y, \ a \rightarrow m\}} m$. Hence $E \cup B \vdash_D \perp$.

A set of defaults K is said to be *consistent in* T if $E \cup B \not\vdash_K \perp$. K is *inconsistent* if it is not consistent [5].

The "More Specific" Relation

We now define the notion of "more specific" between defaults generalizing the specificity principle of Touretzky in inheritance reasoning. Consider for example the network from the example 1, it is clear that being a student is normally a specific case of being a young adult. Since being a young adult is always a specific case of being an adult, it follows that being a student is a specific case of being an adult if the respective individual is a young adult. This stipulates us to say that the default $s \rightarrow \neg m$ (students are normally not married) is more specific than the default $a \rightarrow m$ (adults are normally married) provided that the default $s \rightarrow y$ (students are normally young adults) can be applied. Similarly, in example 2, since being a penguin is always a specific case of being a bird, we can conclude the default $p \rightarrow \neg f$ (penguins don't fly) is always more specific than $b \rightarrow f$ (birds fly).

---

[3] On the other hand, we have to change $p \Rightarrow b$ to $p \rightarrow b$ if we were to use the notion of default theories in [11].

[4] Throughout the paper, we use $\top$ and $\perp$ to denote $True$ and $False$ respectively.

[5] If there is no possibility for misunderstanding, we often simply say consistent instead of consistent in T

Definition 3. *Let d_1, d_2 be two defaults in D. We say that d_1 is* more specific *than d_2 with respect to a set of defaults $K \subseteq D$, denoted by $d_1 \prec_K d_2$, if*
(i) $B \cup \{hd(d_1), hd(d_2)\}$ is inconsistent;
(ii) $bd(d_1) \cup B \vdash_K bd(d_2)$; and
(iii) $bd(d_1) \cup B \not\vdash_K \bot$.

In the above definition (i) guarantees that a priority is defined between two defaults only if they are conflicted, (ii) ensures that being $bd(d_1)$ is a special case of being $bd(d_2)$ provided that the defaults in K can be applied, and (iii) guarantees that K is a sensible set of defaults. *We could say that this is a generalization of Touretzky's specificity principle to general default theories.* In [11], the more specific relation is defined based on the minimal conflict set notion, which in turn is defined based on the defeasible derivation notion. As it can be seen, the above definition is much simpler than that was proposed in [11]. Besides, it allows us to deal with default theories with nonempty background knowledge.

If $K = \emptyset$ we say that d_1 is *strictly more specific* than d_2 and write $d_1 < d_2$ instead of $d_1 \prec_\emptyset d_2$

Example 3. In example 1, $d_2 \prec_{\{d_3\}} d_1$ holds, i.e. d_2 is more specific than d_1 if d_3 is applied. In the context $E = \{s, y, a\}$, d_3 can be applied, and hence d_2 is more specific than d_1 in the context E. But in the context $E' = \{s, \neg y, a\}$, d_3 can not be applied, and hence, d_2 is not more specific than d_1 in E'.

In example 2, it is obvious that $d_2 < d_1$, i.e. d_2 is always more specific than d_1. □

Stable Semantics of Default Reasoning with Specificity

The semantics of a default theory is defined by determining which defaults can be applied to draw new conclusions from the evidences. For example, the semantics of the network in example 1 is defined by determining that the defaults which could be applied are 2 and 3.

In the following, we will see that an argumentation-theoretic notion of attack between a set of defaults K and a default d lies at the heart of the semantics of reasoning with specificity.

Suppose that $K \subseteq D$ is a set of defaults we can apply. Further let d be a default such that $E \cup B \vdash_K \neg hd(d)$. It is obvious that d should not be applied together with K. In this case, we say that *K attacks d by conflict.*

For illustration of attack by conflict, consider the default theory T in example 1. Let $K = \{d_3, d_2\}$. Since $E \cup B \vdash_K \neg m$, K attacks d_1 by conflict. Similarly, $K' = \{d_3, d_1\}$ attacks d_2 by conflict because $E \cup B \vdash_{K'} m$.

The other case where d should not be applied together with K is where it is less specific than some default with respect to K. Formally, this means that if there exists $d' \in D$ such that $d' \prec_K d$ and $E \cup B \vdash_K bd(d')$ then d should not be applied together with the defaults in K. In this case we say that *K attacks d by specificity.*

For illustration of attack by specificity, consider again the default theory T in example 1. Let $K = \{d_3\}$. Because $d_2 \prec_{\{d_3\}} d_1$ and $E \cup B \vdash_{\{d_3\}} bd(d_2)$, K attacks d_1 by specificity.

The following definition summarize what we have just discussed:

Definition 4. *Let $T = (E, B, D)$ be a default theory. A set of defaults K is said to attack a default d in T* [6] *if following conditions are satisfied:*

(i) (Attack by Conflict) $E \cup B \vdash_K \neg hd(d)$; or

(ii) (Attack by Specificity) There exists $d' \in D$ such that $d' \prec_K d$ and $E \cup B \vdash_K bd(d')$.

Note that there is a distinct difference between attack by conflict and inconsistency. It is possible that though K is consistent and $K \cup \{d\}$ is inconsistent but K does not attack d by conflict. It is also possible that K attacks some default d by conflict though $K \cup \{d\}$ is consistent. The Nixon diamond example illustrates these points.

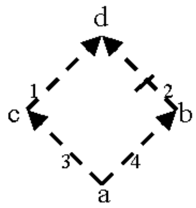

Let $E = \{a\}$, $B = \emptyset$, and $D = \{d_1 : c \rightarrow d,\ d_2 : b \rightarrow \neg d,\ d_3 : a \rightarrow c,\ d_4 : a \rightarrow b\}$. Though $K = \{d_1, d_2, d_4\}$ is consistent and $K \cup \{d_3\}$ is inconsistent, K does not attack d_3 by conflict. Further, though $K' = \{d_2, d_4\}$ attacks d_1 by conflict, $K = K' \cup \{d_1\}$ is consistent.

It is obvious that if K attacks d then every superset of K attacks d. K is said to attack some set $H \subseteq D$ if K attacks some default in H. K is said to *attack itself* if K attacks K.

Now we can give a precise definition of what constitutes the semantics of a default theory with specificity.

Definition 5. *Let T = (E,B,D) be a default theory. A set of defaults S is called an* extension *of T if S does not attack itself and attacks every default not belonging to it.*

Definition 6. *Let T = (E,B,D) be a default theory. Let l be a ground literal. We say T* entails *l, denoted by $T \mid\!\sim l$, if for every extension S of T, $E \cup B \vdash_S l$*

[6] *if there is no possibility for misunderstanding then T often is not mentioned*

Because the defeasible consequence relation $\vdash_K$ subsumes the first order consequence relation (definition 2), it is obvious that an inconsistent set of defaults attacks every default. Therefore it is clear that an extension is always consistent.

Example 4. Consider the theory in example 2. We have that $d_2 < d_1$, i.e., d_2 is strictly more specific than d_1. This can be used to prove that $\{d_2\}$ is the unique extension of T. Therefore $T \mathrel{|\!\sim} \neg f$. □

Example 5.
1. Consider the theory T in example 1. Let $H = \{d_3, d_2\}$. Because $\{s, y, a\} \cup B \vdash_H \neg m$, H attacks d_1 by conflict. Furthermore, since $\{s, y, a\} \cup B \not\vdash_H m$ and $\{s, y, a\} \cup B \not\vdash_H \neg y$, H does not attack itself by conflict. Because there is no default which is more specific than d_2 or d_3 with respect to H, H does not attack itself by specificity. Hence H does not attack itself and attacks every default not belonging to it. Therefore H is an extension of T. Let $K = \{d_1, d_3\}$. Because $d_2 \prec_K d_1$ and $\{s, y, a\} \cup B \vdash_K bd(d_2)$, K attacks d_1 by specificity. Hence K is not an extension of T. It should be obvious now that H is the only extension of T. Hence, $T \mathrel{|\!\sim} \neg m$.
2. Consider the theory T' in example 1. Let $H = \{d_2\}$ and $K = \{d_1\}$. Since $\{s, \neg y, a\} \vdash_H \neg m$ and $\{s, \neg y, a\} \vdash_K m$, and $\{s, \neg y, a\} \vdash_\emptyset \neg y$, H attacks d_1, d_3 by conflict while K attacks d_2, d_3 by conflict. Due to the fact that d_3 can not be applied, there are no defaults d, d' such that $d \prec_H d'$ or $d \prec_K d'$. Hence both H and K do not attack themselves. Thus, both H and K are extensions of T', and so, $T' \not\mathrel{|\!\sim} \neg m$ and $T' \not\mathrel{|\!\sim} m$.

□

The definition 5 of an extension of a default theory corresponds to the stable semantics of argumentation which has been first introduced in [9] and later further studied in [6]. There are also a number of other semantics for argumentation which could be applied to reasoning with specificity. But in this paper we will limit ourselves to the stable semantics.

Existence of Extensions

A well-known problem of stable semantics in nonmonotonic reasoning is that it is not always defined for every nonmonotonic logics. As our semantics is a form of stable semantics of argumentation, it is expected that the same problem will be encountered in our framework. The following example originated from [7] confirms our expectation.

Example 6 ([7]). Consider $T = (E, \emptyset, D)$ with $E = \{a, b, c\}$ and D consists of the following defaults

$$\begin{aligned}
&d_1 : a \wedge q \rightarrow \neg p \\
&d_2 : a \rightarrow p \\
&d_3 : b \wedge r \rightarrow \neg q \\
&d_4 : b \rightarrow q \\
&d_5 : c \wedge p \rightarrow \neg r \\
&d_6 : c \rightarrow r
\end{aligned}$$

It is easy to see that for each $K \subseteq D$, there is no $d \in D$ such that $d \prec_K d_1$ or $d \prec_K d_3$ or $d \prec_K d_5$.

We will prove that T does not have an extension.

Assume the contrary that T has an extension S. We want to prove that $d_1 \notin S$. Assume the contrary that $d_1 \in S$. Since $E \vdash_{\{d_2\}} p$ and S does not attack itself, we conclude that $d_2 \notin S$. This implies that S attacks d_2. There are two cases:

1. S attacks d_2 by conflict. This means that $E \vdash_S \neg p$, which implies that $E \vdash_S q$.
2. S attacks d_2 by specificity. Since the only default in D, that is more specific than d_2, is d_1, S attacks d_2 by specificity implies that $E \vdash_S bd(d_1)$. Thus $E \vdash_S q$.

It follows from the above two cases that $E \vdash_S q$. Therefore S contains d_4. Now, consider the two defaults d_5 and d_6. Since $d_2 \notin S$, $E \not\vdash_S bd(d_5)$. Therefore S does not attack d_6 by specificity. Further $E \not\vdash_S bd(d_5)$ implies that $E \not\vdash_S \neg r$. So, S does not attack d_6 by conflict either. Again, because S is an extension, we have that $d_6 \in S$. However, $E \vdash_{\{d_6\}} bd(d_3)$, which implies that S attacks d_4 by specificity, i.e., S attacks itself. This contradicts the assumption that S is an extension of T. Thus the assumption that $d_1 \in S$ leads to a contradiction. Therefore $d_1 \notin S$.

Similarly, we can prove that $d_3 \notin S$ and $d_5 \notin S$. Since S is a stable extension of T, S attacks d_1. This implies that S must attack d_1 by conflict because there is no default in D which is more specific than d_1. Thus $d_2 \in S$. Similar arguments lead to $d_4 \in S$ and $d_6 \in S$, i.e., $S = \{d_2, d_4, d_6\}$. However, S attacks d_2 by specificity because $d_1 < d_2$ and $E \cup B \vdash_S bd(d_1)$. This means that S attacks itself which contradicts the assumption that S is a stable extension of T. Thus the assumption that there exists an extension leads to a contradiction. Therefore, we can conclude that there exists no extension of T. $\square$

In the next section we will introduce the class of stratified default theories for which extensions always exist.

3 Stratified Default Theories

The definition of stratified default theories is based on the notion of a *rank function* which is a mapping from the set of ground literals $lit(\mathcal{L}) \cup \{\top, \bot\}$ to the set of nonnegative integers.

Definition 7. *A default theory $T = (E, B, D)$ over $\mathcal{L}$ is* stratified *if there exists a rank function of T, denoted by rank, satisfying the following conditions:*

(i) $rank(\top) = rank(\bot) = 0$;

(ii) for each ground atom l, $rank(l) = rank(\neg l)$;

(iii) for all literals l and l' occurring in a clause in B, $rank(l) = rank(l')$;

(iv) for each default $l_1, \ldots, l_m \to l$ in D, $rank(l_i) < rank(l)$, $i = 1, \ldots, m$;

It is not difficult to see that all the default theories in examples 1 and 2 are stratified. We prove that

Theorem 1. *Every stratified default theory has at least one extension.* □

3.1 General Properties of $\mid\!\sim$

There is a large body of works in the literature [2, 14, 19] on what properties characterize a defeasible consequence relation like $\mid\!\sim$. In general, it is agreed that such relation should extend the monotonic logical consequence relation. Further, since the intuition of a default rule d is that $bd(d)$ normally implies $hd(d)$, we expect that in the context $E = \{bd(d)\}$, $T \mid\!\sim hd(d)$ holds. Another important property of defeasible consequence relations is related to the adding of proved conclusions to a theory. Intuitively, this means that if $T \mid\!\sim a$ then we expect T and $T + a$ [7] to have the same set of conclusions. Formally, the discussed key properties are given below:

- Deduction: $T \mid\!\sim l$ if $E \cup B \vdash l$;
- Conditioning: If $E = \{bd(d)\}$ for $d \in D$, then $T \mid\!\sim hd(d)$;
- Reduction: If $T \mid\!\sim a$ and $T + a \mid\!\sim b$ then $T \mid\!\sim b$;
- Cumulativity: If $T \mid\!\sim a$ and $T \mid\!\sim b$ then $T + a \mid\!\sim b$;

In the next two theorems, we show that $\mid\!\sim$ satisfies deduction and reduction:

Theorem 2 (Deduction). *Let T=(E,B,D) be an arbitrary default theory. Then, for every $l \in lit(\mathcal{L})$, $E \cup B \vdash l$ implies $T \mid\!\sim l$.* □

Theorem 3 (Reduction). *Given T=(E,B,D) be an arbitrary default theory and $a, b \in lit(\mathcal{L})$ such that $T \mid\!\sim a$ and $T + a \mid\!\sim b$. Then, $T \mid\!\sim b$.* □

In general $\mid\!\sim$ does not satisfy cumulativity as the following example shows.

Example 7. Consider the default theory $T = (E, B, D)$

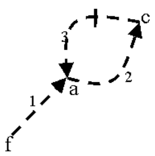

where $E = \{f\}$, $B = \emptyset$, $D = \{d_1{:}f \to a,\ d_2{:}a \to c,\ d_3{:}c \to \neg a\}$ Because the only instance of the more-specific-relation is $d_1 \prec_{\{d_1,d_2\}} d_3$, T has a unique extension $\{d_1, d_2\}$. Hence, $T \mid\!\sim a$ and $T \mid\!\sim c$.

Now consider $T + c$. $T + c$ has two extensions: $\{d_1, d_2\}$ and $\{d_2, d_3\}$. Thus, $T + c \not\mid\!\sim a$. This implies that $\mid\!\sim$ is not cumulative.

[7] $T + a$ denotes the default theory $(E \cup \{a\}, B, D)$.

The next theorem proves that stratification is sufficient for cumulativity.

Theorem 4 (Cumulativity). *Let $T = (E, B, D)$ be a stratified default theory and a, b be literals such that $T \mid\sim a$, and $T \mid\sim b$. Then $T + a \mid\sim b$.* $\square$

Because stratification does not rule out the coexistence of defaults like $a \rightarrow \neg c, \quad a \rightarrow c$, conditioning does not hold for stratified theories as the next example shows.

Example 8. Let $T = (\{a\}, \emptyset, \{d_1 : \ a \rightarrow \neg c, \quad d_2 : \ a \rightarrow c\})$. It is obvious that T is stratified. Because $d_1 < d_2$ and $d_2 < d_1$, both d_1, d_2 are attacked by specificity by the empty set of defaults. Thus the only extension of T is the empty set. Hence, $T \not\mid\sim \neg c$, and $T \not\mid\sim c$. That means that conditioning is not satisfied.

The coexistence of defaults like $a \rightarrow \neg c, \quad a \rightarrow c$ means that a is normally c and normally not c at the same time which is obviously not sensible. Hence it should not be a surprise that conditioning is not satisfied in such cases. The conditioning property would hold for a default d if in the context of $bd(d)$, d is the most specific default. The following definition formalizes this intuition. For simplicity, we often write $d \prec d'$ if $d \prec_K d'$ for some K. Let $\prec^*$ be the transitive closure of $\prec$.

Definition 8. *A default theory T = (E,B,D) is said to be* conditioning-sensible *if for every default d following conditions are satisfied:*

(i) $d \not\prec^* d$;

(ii) For every set $K \subseteq D$ such that $bd(d) \cup B \vdash_{K \cup \{d\}} \bot$ and $bd(d) \cup B \not\vdash_K \bot$, there exist $d' \in K$ such that $d \prec_K d'$

Theorem 5. *Let $T = (E, B, D)$ be a conditioning-sensible default theory, d be a default in D, and $E = bd(d)$. Then $T \mid\sim hd(d)$.* $\square$

It is interesting to note that conditioning-sensibility and stratification are two independent concepts. Default theories like the one in example 7 are conditioning-sensible but not stratified while default theories like that in example 8 are stratified but not conditioning-sensible. Further while example 8 shows that stratification does not imply conditioning, example 7 shows that conditioning-sensibility does not imply cumulativity.

We prove that our approach captures inheritance reasoning by transforming each acyclic and consistent inheritance network Γ into a default theory T_Γ and show that the conclusions sanctioned by the credulous semantics of Γ are also the conclusions of $\mid\sim$ with respect to T_Γ[8].

4 Discussion and Conclusion

Reiter and Criscuolo [35] are among the first to discuss the importance of specificity (or default interaction, in their terminology) in default reasoning. They discussed various situations, in which the interaction between defaults of a normal

[8] Due to the lack of space, we omit the transformation here. It can be found in the full paper.

default theory can be compiled into the original theory to create a new default theory whose semantics yields the intuitive results. It has been recognized relatively early that priorities between defaults can help in dealing with specificity. In prioritized circumscription, first defined by McCarthy [24], a priority order between predicates is added into each circumscription theory. Lifschitz [27] later proved that prioritized circumscription is a special case of parallel circumscription. A similar approach has been taken by Konolige [20] in using autoepistemic logic to reason with specificity. He defined hierarchical autoepistemic theories in which a preference order between sub-theories and a syntactical condition on the sub-theories ensure that higher priority conclusions will be concluded. Brewka [4] - in defining prioritized default logic - also adds a preference order between defaults into a Reiter's default theory and modifies the semantics of default logic in such a way that guarantees that default of higher priority is preferred. Baader and Hollunder [5] develops prioritized default logic to handle specificity in terminological systems. All of the approaches in [4, 5, 24, 27, 20] assume that priorities between defaults are given by the users.

Computing specificity is another important issue in approaches to reasoning with specificity. Work from Poole [31] is an early attempt to extract the preference between defaults from the theory. Poole defines a notion of more specific between pairs consisting of a conclusion and an argument supporting this conclusion. Moinard [26] pointed out that Poole's definition yields unnecessary priority, for example, it can arise even in consistent default theories. Simari and Loui [38] noted that Poole's definition does not take into consideration the interaction between arguments. To overcome this problem they combined Poole's approach and Pollock's theory [30] to define an approach that unifies various approaches to argument-based defeasible reasoning. We have discussed the shortcoming of Simari and Loui's system in the introduction.

Touretzky's specificity principle [42] in inheritance reasoning is a major step in reasoning with specificity. Although this principle is generally accepted, different intuitions on "*what does more specific mean?*" leads to numerous approaches to reasoning with specificity. More interestingly, some seem to contradict the others. Detailed discussions about this problem in inheritance reasoning can be found in Touretzky et al. [43, 44]. Moinard [26] showed that Touretzky's approach does not work well for general default theories. He proposed several principles for determining a preference relation based on specificity in default logic but does not discuss how this preference would change the semantics of a default theory. Furthermore, like Poole he does not take into consideration the interaction between arguments either.

Conditional entailment of Geffner and Pearl [14] bridges the extensional and conditional approaches to default reasoning and is the first approach to reasoning with specificity which satisfies the basic properties of a nonmonotonic consequence relation. Because the priority order between assumptions in [14] is context-independent, conditional entailment, however, is too weak (as also noted by Geffner and Pearl) to capture inheritance reasoning. Pearl also discussed how a preference relation between defaults can be established. In System Z [29], Pearl uses consistency check to determine the order of a default. The lower the order of a default is, the higher is its priority. As in Poole's approach, sometimes System Z introduces unwanted priorities.

The idea of compiling specificities into a general nonmonotonic framework is also used in [7, 11] and in this work. Delgrande and Schaub [7] compiled the preference order between defaults (defined using a order similar to a Z-order of [29]) into the original theory and create a Reiter's default theory whose semantics defines the semantics of the original theory. The compilation of the preference order, however, does not take the context into consideration. As a consequence their approach cannot capture inheritance reasoning. In our approach, the compilation of the more specific relation into the original theory is done in such a way that the context will affect the decision process determining which default can be applied.

Our approach to specificity in this paper is a continuation of our own work in [11]. It could be viewed as a kind of a hybrid between the above approaches. For an intuitive semantical foundation of reasoning with specificity, we develop a general framework, but for implementation, we translate our framework into Reiter's default logics. Wang, You and Yang [46] has applied our idea to give a semantics for possibly cyclic inheritance networks.

Even though our work is not directly related to the recent works on prioritized default theories [8, 15, 36] or adding priority into extended logic programming [3], we believe that there is a mutual benefit between the research done in these works and ours. For example, the more specific relation defined here can be used to specify the priorities between defaults in [8] or the preference relation in [15]. Thus, these two approaches can be extended to realize two different modes of reasoning: one with explicit priority ordering and the other with implicit priority ordering. On the other hand, programs such as that in [46] can be extended to compute the more specific relation and hence allows a fully automatic translation from a default theory $T = (E, B, D)$ into its corresponding Reiter's default theory, R_T. The result of [46] also shows that this can be done in polynomial time for defeasible inheritance networks.

Our work also shows that inheritance networks can be modularly translated into equivalent general nonmonotonic formalism such as Reiter's default theory. We want to note that there are other works on formulating inheritance networks using general nonmonotonic formalisms such as [10, 12, 13, 46] or the works listed in [17]. To the best of our knowledge, our work is the first general approach to default reasoning with specificity which is capable of capturing inheritance reasoning in full.

References

1. Brewka, G.: Adding Priorities and Specificity to Default Logic, Proc. JELIA 94, LNAI 838, Springer Verlag, 247–260.
2. Brewka, G.: Cumulative default logic: In defense of nonmonotonic inference rules, AIJ'50, 183–205.

3. Brewka, G. and Eiter, T.: Preferred Answer Sets for Extended Logic Programs, AIJ'109, (1999) 297–356.
4. Brewka, G.: Reasoning about Priorities in Default Logic, Proc. AAAI'94, Seatle, 1994.
5. Baader, F. and Hollunder, B.: How to prefer more specific defaults in terminological default logic, Proc. IJCAI'93.
6. Bondarenko A., Dung P.M., Kowalski R.A., Toni F.: An abstract, argumentation-theoretic approach to default reasoning, AIJ'93, (1997) 63–101.
7. Delgrande , J.P. and Schaub, T.H.: A General Approach to Specificity in Default Reasoning, Proc. KR'94, (1994) 47–158.
8. Delgrande , J.P. and Schaub, T.H.: Compiling reasoning with and about preferences into default logic, Proc. IJCAI'97, 1997.
9. Dung, P.M.: On the acceptability of arguments and its fundamental role in non-monotonic reasoning and logic programming and N-person game, AIJ'77, 2 (1995) 321–357.
10. Dung, P.M. and Son, T.C.: Nonmonotonic Inheritance , Argumentation, and Logic Programming, Proc. LPNMR'95, (1995) 316–329.
11. Dung, P.M. and Son, T.C.: An Argumentation-theoretic Approach to Default Reasoning with Specificity, Proc. KR'96, (1996) 407–418.
12. Etherington, D.W. and Reiter, R.: On Inheritance Hierarchies With Exceptions, Proc. AAAI'84, (1984) 104–108.
13. Gelfond, M., Przymusińska, H.: Formalization of Inheritance Reasoning in Autoepistemic Logic, Fundamental Informaticae XIII, 4(1990), 403–445
14. Geffner, H., and Pearl J.: Conditional entailment: bridging two approaches to default reasoning, AIJ'53, (1992) 209–244.
15. Gelfond, M. and Son T.C.: Prioritized Default Theory, LNAI'1471, Proceeding of the Workshop LPKR'97, (1997) 164–223.
16. Geerts, P. and Viemeir, D.: A nonmonotonic reasoning formalism using implicit specificity information, Proc. LPNRM'94, (1994) 380–396.
17. Horty, J.F.: Some direct Theories of Non-monotonic Inheritance in Handbook of Logic and Artificial Intelligence and Logic Programming, D.Gabbay and C. Hogger, Oxford Uni., 1991.
18. Horty, J.F., Thomason, R.H., and Touretzky, D.S.: A skeptical theory of inheritance in non-monotonic semantic networks, AIJ'42, (1987) 311–348.
19. Kraus, S., Lehmann, D., and Magidor, M.: Nonmonotonic Reasoning, Preferential Models and Cumulative Logics, AIJ'44, (1990) 167–207
20. Konolige, K.: Hierarchic Autoepistemic Theories for Nonmonotonic Reasoning, Proc. AAAI'88, (1988) 439–443.
21. Marek, V.M. and Truszczynski, M.: Nonmonotonic Logic, Springer Verlag, 1993.
22. Moore, R.C.: Semantical Considerations on Nonmonotonic Logics, AIJ'25, 75–94.
23. McCarthy, J.: Circumscription - a form of nonmonotonic reasoning , AIJ'13, (1980) 27–39.
24. McCarthy, J.: Applications of Circumscription to Formalizing Commonsense Knowledge, AI, 1984.
25. McDermott, D. and Doyle, J.: Nonmonotonic Logic I, AIJ'13, (1980) 41–72.
26. Moinard, Y.: Preference by Specificity in Default Logic, 1990.
27. Lifschitz, V.: Computing Circumscription, Proc. IJCAI'85, (1985) 121–127.
28. Lloyd, J.W.: Foundations of Logic Programming, Springer Verlag, 1987.
29. Pearl, J.: System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning: A survey; Proc. KR-89, (1989) 505–516.

30. Pollock, J.L.: Defeasible reasoning. Cognitive Science **17** (1987) 481–518.
31. Poole, D.: On the comparison of theories: Preferring the most specific explanation, Proc. IJCAI'85, (1985) 144–147.
32. Prakken, H., and Sartor, G.: A semantics for argument-based systems with weak and strong negation and with explicit priorities. Draft, July 1995. An extended abstract can be found in Proceedings of the 5th International Conf. on AI and Laws, ACM Press, 1995, 1-9.
33. Prakken, H., and Vreeswijk, G.A.W.: Logics for Defeasible Argumentation, To appear in Handbook of Philosophical Logic, D.Gabbay, Oxford University.
34. Reiter R.: A Logic for Default Reasoning, AIJ'13 1-2 (1980), 81-132.
35. Reiter R. and Criscuolo G.: On interacting defaults, in Readings in Nonmonotonic Reasoning, Edited by M. L. Ginsberg, Morgan Kaufmann Publishers, Inc., Los Altos, California (1987) 94–100
36. Rintanen, J.: Lexicographic priorities in default logics, AIJ'106, 2 (1998) 221–265.
37. Sandewall, E.: Nonmonotonic Inference Rules for Multiple Inheritance with Exceptions, Proceedings of the IEEE, Vol. 74, No. 10, October (1986) 1345–1353
38. Simari, G.R. and Loui, R.P.: A mathematical treatment of defeasible reasoning and its implementation, Elsevier, AIJ'53, (1992) 125–257.
39. Simonet, G. and Ducournau R.: On Stein's paper: resolving ambiguity in nonmonotonic inheritance hierarchies, AIJ'71, 1 (1994) 183–193.
40. Simonet, G.: On Sandewall's paper: Nonmonotonic inference rules for multiple inheritance with exceptions, AIJ'86, 2 (1996) 359–374.
41. Stein, L.A.: Resolving ambiguity in non-monotonic inheritance hierarchies, AIJ'55, 259–310.
42. Touretzky, D. S.: The mathematics of Inheritance Systems, Morgan Kaufmann Pub., Inc., Los Altos, CA, 1986.
43. Touretzky, D. S., Horty, J.F., and Thomason, R.H.: A clash of Intuition: The current state of Non-monotonic Multiple Inheritance Systems, Proc. IJCAI'87, (1987) 476–482
44. Touretzky, D. S., Horty, J.F., and Thomason, R.H.: A skeptic's Menagerie: Conflictors, Preemptors, Reinstaters, and zombies in nonmonotonic Inheritance, Proc. IJCAI'91, (1991) 478–483
45. Vreeswijk, G.A.W.: The feasibility of defeat in defeasible reasoning, Proc. of KR-91, Cambridge, MA, 1991
46. Wang, X., You, J.H., and Yuan, L.Y.: A default interpretation of defeasible network, Proc. IJCAI'97, 1997.