

On Specificity in Default Logic

Jussi Rintanen
Department of Computer Science
Helsinki University of Technology
Otakaan 1, 02150 ESPOO
Finland

Abstract

The applicability of lexicographic comparison in nonmonotonic reasoning with specificity is investigated. A priority mechanism based on lexicographic comparison is defined for Reiter's default logic. The following principle - earlier used by Geffner and Pearl in conditional entailment - is used as the basis for specificity-based priorities for normal default theories: for each rule there is a context where it may not be defeated by any other rule. A method for computing priorities according to the principle is given. Connection to earlier work is discussed.

1 Introduction

Two kinds of weaknesses have been identified in nonmonotonic logics: insensitivity to specificity [Poole, 1985] and sensitivity to irrelevant premises [Geffner and Pearl, 1992]. Only the first kind of weakness is present in systems that are defined using some form of maximal consistency [Reiter, 1980, Moore, 1985, Poole, 1988]. In the well-known example, Tweety is a penguin, penguins are birds, penguins usually cannot fly, birds usually can. Reiter's default logic does not say anything about the flying ability of Tweety. The rule concerning the flying ability of penguins is more specific than the rule about birds, and the rule about penguins should be used. Reiter and Cnscuolo [1981] suggest that the application of less specific defaults can be explicitly blocked by using semi-normal defaults. Others introduce *priorities* [Lifschitz, 1985, Konohge, 1988, Brewka, 1989, Baader and Hollunder, 1993, Brewka, 1994, Rintanen, 1994]. Priorities are given as orderings of predicates, formulae, or default rules: in conflicting situations preference is given to items with a higher priority.

The second kind of weakness shows up in logics based on probability [Adams, 1975, Pearl, 1988] and in logics that are defined directly by stating the properties a consequence relation should satisfy, for example the preferential logics of Kraus et al. [1990]. These logics give the correct answers in the standard examples about specificity, for instance concerning Tweety's flying ability. Their weakness is apparent when a set of premises is extended with an irrelevant premise: most of the useful consequences are cancelled. For example, if Tweety

is a black penguin, Tweety's inability to fly cannot be inferred.

In this paper, we investigate the applicability of lexicographic priorities as a mechanism for resolving default conflicts according to specificity. Our approach is conservative: we do not propose any original notion of specificity. To date, the most successful area in defeasible reasoning - from the point of view of handling both specificity and irrelevant premises - is inheritance reasoning [Hortj, 1994]. Therefore it seems desirable that solutions to these problems in a more general setting would subsume some form of inheritance reasoning. Different inheritance theories propose different specificity notions. A weak but reasonable principle that is found for example in the preferential logics of Kraus et al. and that is also respected by inheritance theories is the following: given a default that allows the conclusion of Q from a, the premise a and no other premises, conclude Q . This principle is used in the definition of conditional entailment [Geffner and Pearl, 1992] and we shall use it in deriving priorities from specificity. We strengthen Reiter's default logic with a priority mechanism that uses lexicographic comparison. Lexicographic comparison has been used in earlier work on explicit priorities, for example in [Lifschitz, 1985, Brewka, 1989, Geffner and Pearl, 1992, Ryan, 1992].

The outline of the paper is as follows. In Section 2 we present the specificity notion, and in Section 3 we extend default logic with a priority mechanism. In Section 4 we show how nonmonotonic reasoning that respects the specificity notion can be achieved in the extended default logic and how it subsumes an inheritance theory. In Section 5 we discuss how the specificity principle is too weak for default theories in general, claim that priority mechanisms that lexicographically order extensions or models are in general too weak for achieving inheritance-like reasoning, and discuss the possibility of stronger priorities in more general but still restricted cases.

2 Specificity in default logic

Many formalizations of defeasible reasoning that address specificity [Kraus et al., 1990, Lehmann and Magidor, 1992, Adams, 1975, Pearl, 1988, 1990, Goldszmidt et al., 1990] employ or fulfill a principle not fulfilled by default logic. For default theories (D, BUC) with normal de-

faults D the principle can be expressed as follows

$$(D, B \cup \{\alpha\}) \models_x \beta \text{ for all defaults } \frac{\alpha}{\beta} \in D \quad (1)$$

Above D is the set of default rules of a default theory, B is a set of sentences that forms the background¹, and \models_x is a consequence relation. The consequence relation usually discussed in the context of default logic is *cautious reasoning* $\Delta \models_c \phi$, which means that the formula ϕ is contained in all extensions of the default theory Δ . This consequence relation does not fulfill Condition 1.

Example 2.1 Consider the following set of defaults D

$$\frac{P}{B} \quad \frac{B}{F} \quad \frac{P}{\neg F}$$

The default theory $\Delta = (D, \{P\})$ has two extensions. Both contain P and B , one contains $\neg F$ and the other contains F . Therefore $\Delta \not\models_c \neg F$. ■

In ϵ -semantics [Pearl, 1988] a counterpart of Condition 1 is directly justified by the definition of conditional probability. Outside probabilistic systems, this condition has been used for example by Poole [1991]. A default can be viewed as expressing a rule that applies unconditionally in the context specified by the antecedent of the rule. In Example 2.1 above, if P is read as "x is a penguin", B as "x is a bird", and F as "x is able to fly". If B is all the information available, conclude that x flies. If P is all the information available, conclude that x does not fly.

Consequence relations for default logic that treat specificity properly should fulfill at least Condition 1. As shown above, \models_c is not such a relation. A natural way to strengthen it is to consider only a subset of the extensions, those in which conflicts are resolved without violating specificity, and take their intersection. Such a subset can be selected for example by lexicographic comparison.

3 Priorities in default logic

Lexicographic priorities in the context of expressive non-monotonic logics have been earlier investigated in [Rintanen, 1994]. In that paper a definition of preferred stable expansions for autoepistemic logic was given. A similar definition can be devised for default logic. The significance of defaults is represented by a strict partial order, that is, a transitive and asymmetric relation \mathcal{P} . If defaults δ and δ' are related as $\delta \mathcal{P} \delta'$, then the application of δ is more desirable than the application of δ' . The default δ has higher priority.

Below, \mathcal{L} denotes the language of the propositional logic, and \mathcal{D} the set of all defaults $\frac{\alpha}{\beta}$ where $\{\alpha, \beta\} \subseteq \mathcal{L}$.

¹This - possibly controversial - division between background B and evidence C (necessary and contingent facts) is made in all the afore-mentioned formalisms that address specificity. The background may affect the priorities, the evidence may not.




Figure 1 An extended penguin triangle

Definition 3.1 (Application) A default $\frac{\alpha}{\beta}$ is applied in $E \subseteq \mathcal{L}$, if $\alpha \wedge \beta \in E$. This is written as $app(\frac{\alpha}{\beta}, E)$.

Definition 3.2 (Preferred extensions) Let $\Delta = (D, W)$ be a default theory, and \mathcal{P} a strict partial order on D . Let E be an extension of Δ . Then E is a \mathcal{P} -preferred extension if there is a strict total order $\mathcal{T} \supseteq \mathcal{P}$ on D such that for all extensions E' of Δ , for all $\delta \in D$,

$$app(\delta, E' \setminus E) \text{ implies that for some } \delta' \in D, \delta' \mathcal{T} \delta \text{ and } app(\delta', E' \setminus E)$$


Definition 3.3 The fact that formula ϕ is contained in all \mathcal{P} -preferred extensions of Δ is written as $\Delta \models_p \phi$.

Example 3.1 For the default theory (D, \emptyset) depicted in Figure 1 the following priorities are needed

$$\mathcal{P} = \{ \{ \frac{P}{\neg R}, \frac{Q}{R} \}, \{ \frac{P}{\neg R}, \frac{R}{S} \} \}$$

The default theory $(D, \{P\})$ has two extensions that respectively apply $\{ \frac{P}{\neg R}, \frac{P}{Q} \}$ and $\{ \frac{P}{Q}, \frac{Q}{R}, \frac{R}{S} \}$. Only the first one is \mathcal{P} -preferred. The two extensions and all strict total orders $\mathcal{T} \supseteq \mathcal{P}$ on the defaults are depicted below on the left (the default $\frac{P}{Q}$ is omitted because it is applied in both extensions). The most significant defaults are the lowest. Application is depicted as \bullet and non-application as \circ .

If $(\frac{P}{\neg R}, \frac{R}{S})$ were omitted in \mathcal{P} , the strict total order $\frac{R}{S}, \frac{P}{\neg R}, \frac{P}{Q}, \frac{Q}{R}$ would extend \mathcal{P} , and consequently there would be a \mathcal{P} -preferred extension in which $\frac{P}{R}$ were not applied. The diagram on the right depicts the unwanted total order.



4 Specificity-based priorities

Pre-computed explicit priorities are useful if the computation can be performed once and the priorities can be used in answering a wide variety queries concerning the knowledge base. Priorities for a default theory $\langle D, BuC \rangle$ depend only on the defaults D and the background B . As long as changes are made only to the contingent information C , no recomputation of priorities is needed.

Next we show how Condition 1 can be fulfilled in default logic with the lexicographic priority mechanism defined in the preceding section.

Definition 4.1 Let $\langle D, B \rangle$ be a default theory and f a function from $\mathcal{D} \times 2^{\mathcal{D}} \times 2^{\mathcal{C}}$ to $2^{\mathcal{D}}$. Define $\mathcal{P}_{\Delta, f}$ as follows:

$$\begin{aligned} D_0 &= \left\{ \frac{\alpha}{\beta} \in D \mid \langle D, B \cup \{\alpha\} \models \beta \right\} \\ D_{i+1} &= \left\{ \delta \in D \setminus \bigcup_{0 \leq k \leq i} D_k \mid f(\delta, D, B) \subseteq \bigcup_{0 \leq k \leq i} D_k \right\} \\ \mathcal{P}_0 &= \emptyset \\ \mathcal{P}_{i+1} &= \{ \langle \delta, \delta' \rangle \mid \delta \in D_{i+1}, \delta' \in f(\delta, D, B) \} \\ \mathcal{P}_{\Delta, f} &= \text{the transitive closure of } \bigcup_{i \geq 0} \mathcal{P}_i \end{aligned}$$

If $D_0 \cup \dots \cup D_n = D$ for no $n \geq 0$, then we leave $\mathcal{P}_{\Delta, f}$ undefined.

Given a function f , the above definition imposes an ordering on defaults. For a default δ , $f(\delta, D, B)$ can be thought of as the set of defaults that may become applied if a default conflict is resolved so that Condition 1 or some stronger specificity notion is violated. To avoid this, δ has to be given a higher priority than any of the defaults in $f(\delta, D, B)$. Violation of Condition 1 is sometimes unavoidable, for example when $D = \left\{ \frac{A \wedge B}{A}, \frac{A \wedge \neg B}{\neg B} \right\}$.

Next we give one such function f . The notation $\Delta \models \phi$ means that ϕ is in at least one extension of Δ .

Theorem 4.2 Let $\Delta = \langle D, B \rangle$. Define $f_1(\frac{\alpha}{\beta}, D, B)$ as

$$\begin{aligned} &\left\{ \delta \in D \mid E \text{ and } E' \text{ are extensions of } \langle D, B \cup \{\alpha\} \rangle, \right. \\ &\quad \text{app}(\delta, E' \setminus E), \text{app}(\frac{\alpha}{\beta}, E \setminus E'), \\ &\quad \text{for no extension } E'' \text{ of } \langle D, B \cup \{\alpha, \beta\} \rangle \\ &\quad \left. \left\{ \delta \in D \mid \text{app}(\delta, E' \setminus E'') \text{ or } \text{app}(\delta, E'' \setminus E') \right\} \subset \right. \\ &\quad \left. \left\{ \delta \in D \mid \text{app}(\delta, E \setminus E') \text{ or } \text{app}(\delta, E' \setminus E) \right\} \right\} \end{aligned}$$

Then $\models_{\mathcal{P}_{\Delta, f_1}}$ fulfills Condition 1 for $\langle D, B \rangle$.

Proof. The relation $\mathcal{P} = \mathcal{P}_{\Delta, f_1}$ is transitive and asymmetric by construction. Let $\frac{\alpha}{\beta}$ be any default in D_0 . Clearly $\langle D, B \cup \{\alpha\} \rangle \models \beta$. Let $\delta = \frac{\alpha}{\beta}$ be any default in D_i for some $i \in \{1, \dots, n\}$. Let $\mathcal{T} \supseteq \mathcal{P}$ be any strict total order on D . Let E' be any extension of Δ that does not apply $\frac{\alpha}{\beta}$ (hence $\neg\beta \in E'$). By construction, there is an extension E of $\langle D, B \cup \{\alpha, \beta\} \rangle$ such that for all defaults δ' such that $\text{app}(\delta', E' \setminus E)$, $\delta \mathcal{P} \delta'$. Hence $\delta \mathcal{T} \delta'$, and E' cannot be \mathcal{P} -preferred. Therefore all \mathcal{P} -preferred extensions contain β . \square

Gelfond and Przymusinska [1990] present theories of inheritance that are based on translating inheritance networks into sets of formulae in autoepistemic logic. In Section 4b of their paper, they give two formalizations of inheritance. It turns out that the one that explicitly orders extensions of an inheritance network is closely related to our definition of specificity-based priorities and our priority mechanism². Our result below provides a conceptual simplification of Gelfond's and Przymusinska's work. Reification of inheritance networks is avoided as links can be directly represented as default rules, and the ordering on extensions is justified by Condition 1. We restrict to inheritance networks with only defeasible links between classes. Strict links require stronger priorities or non-normal defaults.

Theorem 4.3 Assume that G is an inheritance network with only one object x , and that there are no strict links between classes. The translation of G into default logic is $\Delta = \langle D, B \cup C \rangle$, where

$$\begin{aligned} D &= \left\{ \frac{P}{Q} \mid \text{there is link } P \rightarrow Q \text{ in } G \right\} \\ &\cup \left\{ \frac{P}{\neg Q} \mid \text{there is link } P \nrightarrow Q \text{ in } G \right\} \\ B &= \emptyset \\ C &= \left\{ P \mid \text{there is strict link } x \Rightarrow P \text{ in } G \right\} \\ &\cup \left\{ \neg P \mid \text{there is a strict link } x \nRightarrow P \text{ in } G \right\} \end{aligned}$$

Let $\mathcal{P}_{\Delta, f_1}$ be the priorities for Δ as given in Definition 4.1. Then for any propositional variable P , P belongs to all $\mathcal{P}_{\Delta, f_1}$ -preferred extensions of Δ if and only if $H(x, P)$ belongs to all belief sets of $Th(G)$.

Condition 1 imposes only a lower bound on the consequence relation. Reasoning with lexicographic priorities is monotonic with respect to the priorities: the more priorities there are, the more conclusions can be obtained. Consider the Nixon Diamond. No priorities on the defaults involved violate Condition 1. Some priorities prefer "quakers are pacifists" to "republicans are not pacifists". There is no reason for this. It would seem that demanding that priorities are minimal would solve this problem. This is not the case.

$$\left\{ \frac{P}{Q}, \frac{P \wedge A \vee B}{A \vee B}, \frac{Q \wedge \neg A}{\neg A}, \frac{Q \wedge \neg B}{\neg B} \right\}$$

If the second default is preferred to one or both of the last two, Condition 1 is satisfied. However, minimal such priorities give arbitrarily preference to one of the last two defaults. In conditional entailment [Geffner and Pearl, 1992], this problem is avoided by considering all (minimal) priorities. There is no obvious way to formalize this requirement about the fairness of priorities. Our priorities are not minimal and do not seem to be unfair because defaults are treated symmetrically in cases like the above.

² On page 412 in [Gelfond and Przymusinska, 1990], the relation *better* is defined. The results in the paper assume that this relation is asymmetric although it is not. Our results use a corrected definition.




Figure 2 Two independent penguin triangles

5 Beyond inheritance networks

Example 3.1 demonstrates that it is not sufficient to prioritize defaults the conclusions of which are contradictory. To guarantee that a conflict is solved properly, it is necessary to give a lower priority also to those defaults that potentially become applied if a conflict is resolved in a wrong way. In inheritance networks, these defaults can be easily identified, as shown in Section 4. This is not the case with default theories in general. Next we give a variant of Example 3.1.

Example 5.1 Let $\Delta = \langle D, \emptyset \rangle$ where

$$D = \left\{ \frac{P \neg R}{\neg R}, \frac{P \ Q}{Q}, \frac{Q \ R}{R}, \frac{R \wedge A \ S}{S} \right\}$$

Priorities according to Definition 4.1 are $\mathcal{P}_{\Delta, I_1} = \left\{ \left(\frac{P \neg R}{\neg R}, \frac{Q \ R}{R} \right) \right\}$, but $\langle D, \{P \ A\} \rangle \not\models_{\mathcal{P}_{\Delta, I_1}} \neg R$. The conclusions with $\mathcal{P}_{\Delta, I_1}$ are too weak. ■

This is a problem not recognized in earlier work on lexicographic priority mechanisms. To overcome it, it does not suffice to consider Condition 1 only. A more general definition of specificity has to be devised. Such a definition should cover also cases where the contingent sentences are not just a prerequisite of one default. As hinted in Example 5.1, the contingent facts may in an intricate way affect the set of defaults that potentially become applied. As priorities are to be independent of the contingent facts, this creates new challenges for the definition of specificity based priorities, as illustrated in the following example.

Example 5.2 Consider the defaults D in Figure 2. We want $\langle D, \{A, C \rightarrow P\} \rangle \models_P \neg C$. We also want $\langle D, \{P, R \rightarrow A\} \rangle \models_P \neg R$. To fulfill these requirements respectively the priorities $\frac{A \rightarrow C \ P \ \neg R}{\neg C}$ and $\frac{P \ \neg R \ P \ A \rightarrow C}{\neg R}$ are needed. Hence no \mathcal{P} fulfills these requirements. ■

Similar examples can be constructed for other systems of prioritized nonmonotonic reasoning that order extensions or classical models by lexicographic or other form of comparison, for example [Geffner and Pearl, 1992, Brewka, 1989, Lifschitz, 1985, Ryan, 1992, Gardenfors and Makinoji, 1994]. The presence of disjunctive formulae in the sets of contingent facts is an obstacle to the possibility of inheritance-like reasoning in systems of default reasoning that use fixed priorities with lexicographic comparison.

A natural line of research is to consider cases where the contingent facts may not contain the kind of problematic implicational formulae used in Example 5.2. A suitable restriction is to require that the consistent facts are propositions, variables or their negations. With this restriction satisfactory defeasible reasoning with specificity can hopefully be achieved for a wide class of default theories.

6 Alternative definitions of prioritized reasoning

The priority mechanism introduced in Definition 3.2 is one of several alternatives. That definition classifies defaults to those that are applied in an extension and to those that are not. Another reasonable classification is to defaults that are not defeated and to defaults that are defeated.

Definition 6.1 (Defeat) A default $\frac{\alpha \beta}{\beta}$ is defeated in $E \subseteq \mathcal{L}$, if $\alpha \wedge \neg \beta \in E$.

For prerequisite-free defaults application and non-defeat coincide. The notions differ for a default $\frac{\alpha \beta}{\beta}$ and an extension E when the prerequisite α is not in E . Then $\frac{\alpha \beta}{\beta}$ is neither defeated nor applied in E . Another choice - orthogonal to the choice between application and non-defeat - is the way partiality in the priority partial orders is interpreted.

Definition 6.2 (Preferred₂ extensions) Let $\Delta = \langle D, W \rangle$ be a default theory, and \mathcal{P} a strict partial order on D . Let E be an extension of Δ . Then E is a \mathcal{P} -preferred₂ extension if for all extensions E' of Δ , there is a strict total order $T \supseteq \mathcal{P}$ on D such that for all $\delta \in D$,

$$\text{app}(\delta, E' \setminus E) \text{ implies that for some } \delta' \in D, \\ \delta' T \delta \text{ and } \text{app}(\delta', E \setminus E')$$

Each "P-preferred extension is P-preferred₂ but not vice versa. The generalization of lexicographic comparison to partially ordered positions as given in the definition of preferred₂ has been earlier used by Brewka [1989]. The generalization used in preferred₂ has been earlier used for example by Geffner and Pearl [1992] and Ryan [1992]. The last two do not explicitly reduce lexicographic comparison with partially ordered positions to lexicographic comparison with totally ordered positions. Brewka performs lexicographic comparison implicitly in a procedure for computing preferred sets of formulae.

There are no obvious differences in the expressivity of the priority mechanisms suggested by the above definitions and the one given in Section 3. A theorem corresponding to Theorem 4.3 for a priority mechanism that uses defeat instead of application requires stronger priorities than those obtained merely on the basis of Condition 1. The reason for this is that in inheritance networks, the number of defaults that are potentially defeated in the intended extensions but not in the unintended extensions, may be increased by extending the set of contingent facts with propositional atoms.

7 Related work

Pearl's system Z is an improvement over p-entailment [Adams, 1975] and preferential logics of Kraus, Lehmann and Magidor [1990] because it handles irrelevant premises satisfactorily. System Z imposes an ordering on defaults (the Z-ordering), and models of the defaults are ordered according to the highest defeated default. This introduces a new kind of weakness: if a member of some class is in some respect exceptional, it is allowed to be exceptional in other respects as well. Conditional entailment of Geffner and Pearl [1992] remedies this problem: comparison by highest defeated default is replaced by lexicographic comparison that considers significant also violations to defaults below the highest defeated one. Instead of the Z-ordering, all strict partial orders that - together with lexicographic comparison - satisfy a counterpart of Condition 1, are used. Within our framework, conditional entailment can be defined as follows:

Definition 7.1 Let D_c be a set of defaults $\alpha \Rightarrow \beta$ and $B \subseteq \mathcal{L}$ a set of formulae. Let $D = \{ \frac{\top \alpha \rightarrow \beta}{\alpha \rightarrow \beta} \mid \alpha \Rightarrow \beta \in D_c \}$. Then $((D_c, B), C)$ conditionally entails γ , if γ is in all \mathcal{P} -preferred₂ extensions of $(D, B \cup C)$ for all \mathcal{P} such that $(D, B \cup \{\alpha\}) \models_{\mathcal{P}} \beta$ for all $\alpha \Rightarrow \beta \in D_c$.

Delgrande and Schaub [1994] and Brewka [1994] present methods for computing priorities for Reiter's default logic. Both identify minimal sets of conflicting defaults, Delgrande and Schaub directly using the Z-ordering of Pearl, and Brewka using a closely related idea that takes into account the fact that default rules are unidirectional. Delgrande and Schaub do not use priorities explicitly, but translate sets of defaults and the priority information to sets of semi-normal default rules. Brewka [1994] uses the priorities within his prioritized default logic that is based on a variant of Reiter's semiconstructive definition of extension. Similar priority mechanisms have been independently investigated elsewhere [Marek and Truszczynski, 1993, Baader and Hollunder, 1993]. Neither Delgrande and Schaub nor Brewka verify their systems against any correctness criteria or otherwise explicate the specificity notion they use. Both systems seem to fulfill Condition 1.

Default logic with specificity as defined in this paper differs in some aspects from most of the inheritance theories. The priorities are fixed, that is, they are the same for all extensions and all sets C of contingent information. Theories of inheritance have priorities implicitly hidden in the definition of preemption. There are examples where fixed priorities necessarily seem to produce too strong conclusions [Horty, 1994].

8 Conclusions

We have extended Reiter's default logic with a lexicographic priority mechanism and presented a method for computing priorities according to specificity. The method produces priorities that are sufficiently strong for inheritance networks. Relaxing two restrictions made in inheritance networks - only propositional atoms as prerequisites and factual knowledge that is not disjunctive - reveals the priorities insufficient. It seems that

non-atomic prerequisites can be handled with a lexicographic priority mechanism for a large class default theories, but disjunctive contingent facts cannot.

Our research is closely related to Geffner's and Pearl's work on conditional entailment [Geffner and Pearl, 1992]. Conditional entailment does not support inheritance reasoning. Delgrande [1994] extends conditional entailment towards inheritance reasoning. Our results are likely to have implications on research on that topic.

The computation of priorities from specificity is rather complex and tightly intertwined with default reasoning itself. In argument-based systems of defeasible reasoning [Simari and Loui, 1992, Pollock, 1994] these two things are combined: there is no division between the reasoning and priority components in these systems. This raises questions concerning the usefulness of such division, for example from the computational point of view.

Acknowledgements

This research was funded by the Academy of Finland. The author gratefully acknowledges additional financial support from the following foundations: Suomen Kulttuurirahasto, Emil Aaltosen saatio.

References

- [Adams, 1975] Ernest W. Adams. *The Logic of Conditionals: An Application of Probability to Deductive Logic*, volume 86 of *Synthese Library*. D. Reidel Publishing Company, Dordrecht, Holland, 1975.
- [Baader and Hollunder, 1993] Franz Baader and Bernhard Hollunder. How to prefer more specific defaults in terminological default logic. In Ruzena Bajcsy, editor, *Proceedings of the 13th International Joint Conference on Artificial Intelligence*, volume 1, pages 669-674, Chambéry, France, 1993.
- [Brewka, 1989] Gerhard Brewka. Preferred subtheories: an extended logical framework for default reasoning. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, pages 1043-1048, Detroit, 1989. Morgan Kaufmann Publishers.
- [Brewka, 1994] Gerhard Brewka. Adding priorities and specificity to default logic. In *Logic in Artificial Intelligence: Proceedings of the 4th European Workshop JELIA'94*, number 838 in *Lecture Notes in Artificial Intelligence*, pages 247-260, York, England, 1994. Springer-Verlag.
- [Delgrande and Schaub, 1994] James P. Delgrande and Torsten H. Schaub. A general approach to specificity in default reasoning. In *Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning*, pages 146-157, Bonn, May 1994.
- [Delgrande, 1994] James P. Delgrande. A preference-based approach to default reasoning. Preliminary report. In *Proceedings of the 12th National Conference on Artificial Intelligence*, pages 902-908, 1994.
- [Gardenfors and Makinson, 1994] Peter Gardenfors and David Makinson. Nonmonotonic inference based on expectations. *Artificial Intelligence*, 65: 197-245, 1994.

- [Geffner and Pearl, 1992] Hector Geffner and Judea Pearl Conditional entailment bridging two approaches to default reasoning *Artificial Intelligence* 53 209-244, 1992
- [Gelfond and Przymusinska, 1990] Michael Gelfond and Habna Przymueineka Formalization of inheritance reasoning in autoepistemic logic *Fundamenta Informaticae*, 13(4) 403-443, 1990
- [Goldszmidt et al, 1990] Moises Goldszmidt, Paul Morris, and Judea Pearl A maximum entropy approach to nonmonotonic reasoning In *Proceedings of the 9th National Conference on Artificial Intelligence*, pages 646-652, Boston, 1990
- [Horty, 1994] John F Horty Some direct theories of nonmonotonic inheritance In D M Gabbay, C J Hogger, and J A Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming Nonmonotonic Reasoning and Uncertain Reasoning*, volume 3, pages 111-187 Clarendon Press Oxford, 1994
- (Konohge, 1988] Kurt Konohge Hierarchic autoepistemic theories for nonmonotonic reasoning Preliminary report In M Remfrank, J de Kleer, M L Ginsberg, and E Sandewall, editors, *Proceedings of the 2nd Workshop on Non-Monotonic Reasoning*, number 346 in *Lecture Notes in Artificial Intelligence*, pages 42-59, Grassau Germany, June 1988 Springer-Verlag
- [Kraus et al, 1990] Sarit Kraus, Daniel Lehmann, and Menachem Magidor Nonmonotonic reasoning preferential models and cumulative logics *Artificial Intelligence*, 44 167-207, 1990
- [Lehmann and Magidor, 1992] Daniel Lehmann and Menachem Magidor What does a conditional knowledge base entail? *Artificial Intelligence*, 55 1-60, 1992
- [Lifschitz, 1985] Vladimir Lifschitz Computing circumscription In A Joshi, editor, *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pages 121-127, Los Angeles, August 1985 Morgan Kaufmann Publishers
- [Marek and Truszczyrski, 1993] V W Marek and M Truszc} riski *Nonmonotonic Logic Context-Dependent Reasoning* Springer-Verlag, 1993
- [Moore, 1985] R C Moore Semantical considerations on nonmonotonic logic *Artificial Intelligence* 25 75-94, 1985
- [Pearl, 1988] Judea Pearl *Probabilistic Reasoning in Intelligent Systems Networks of Plausible Inference* Morgan Kaufmann Publishers, San Mateo, CA, 1988
- [Pearl, 1990] Judea Pearl System Z a natural ordering of defaults with tractable applications to nonmonotonic logic In *Proceedings of the 3rd Conference on Theoretical Aspect* of Reasoning about Knowledge*, pages 121-135 Pacific Grove, CA, March 1990 Morgan Kaufmana Publishers
- [Pollock, 1994] John L Pollock Justification and defeat *Artificial Intelligence*, 67 377-407 1994
- [Poole, 1985] David L Poole On the comparison of theories preferring the most specific explanation In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, pages 144-147, Los Angeles, CA, 1985
- [Poole, 1988] David Poole A logical framework for default reasoning *Artificial Intelligence*, 36 27-47, 1988
- [Poole, 1991] David Poole The effect of knowledge on belief Conditioning, specificity and the lottery paradox in default reasoning *Artificial Intelligence*, 49(1-3)281-307, 1991
- [Reiter and Cnscuolo, 1981] Raymond Reiter and Giovanni Cnscuolo On interacting defaults In *Proceedings of the 7th International Joint Conference on Artificial Intelligence*, pages 270-276, 1981
- [Reiter, 1980] R Reiter A logic for default reasoning *Artificial Intelligence*, 13 81-132, 1980
- [Rintanen, 1994] Jussi Rintanen Prioritized autoepistemic logic In *Logic in Artificial Intelligence Proceedings of the 4th European Workshop JELIA '94* number 838 in *Lecture Notes in Artificial Intelligence*, pages 232-246, York, England, 1994 Springer-Verlag
- [Rvan, 1992] Mark Rvan Representing defaults as sentences with reduced priority In B Nebel, C Rich, and W Swartout, editors, *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning*, pages 649-660, Cambridge, MA October 1992 Morgan Kaufmann Pubhshers
- [Simari and Loui, 1992] Guillermo R Simari and Ronald P Loui A mathematical treatment of defeasible reasoning and its implementation *Artificial Intelligence*, 53 125-157, 1992