Chapter 4
Nonmonotonic Reasoning

Alexander Bochman

Abstract Nonmonotonic reasoning is a theory of the rational use of assumptions. We describe the relations between NMR and Logic, and two main paradigms of NMR, preferential and explanatory one.

4.1 Nonmonotonic Reasoning Versus Logic

Nonmonotonic reasoning (NMR) is an essential part of the logical approach to Artificial Intelligence. Its birth is due to the research methodology suggested in McCarthy [16] whose objective was a logical formalization of *common sense reasoning* for dealing with AI problems. NMR itself was born, however, as a result of dissatisfaction with traditional logical methods. Reasoning necessary for an intelligent behavior and decision making has appeared to be impossible to represent as deductive inferences in a logical system. The essence of the problem was formulated in Minsky [21] that questioned the suitability of representing commonsense knowledge in a form of a deductive system. Minsky also pointed to monotonicity of logical systems as a source of the problem:

> *Monotonicity*: ... In any logistic system, all the axioms are necessarily "permissive" - they all help to permit new inferences to be drawn. Each added axiom means more theorems, none can disappear. There simply is no direct way to add information to tell such the system about kinds of conclusions that should not be drawn!

Long before the first nonmonotonic formalisms, there have been problems and applications in AI that required some forms of nonmonotonic reasoning. Initial solutions to these problems worked, and this was an incentive for trying to provide them with a more systematic logical basis [15].

A. Bochman (✉)
Computer Science Department, Holon Institute of Technology, Holon, Israel
e-mail: bochmana@hit.ac.il

S. O. Hansson, V. F. Hendricks (eds.), *Introduction to Formal Philosophy*, Springer Undergraduate Texts in Philosophy, https://doi.org/10.1007/978-3-319-77434-3_4

NMR is intimately related to traditional philosophical problems of natural kinds and ceteris paribus laws. These notions resist precise logical definition, but involve description of normal cases. Reasoning with such concepts is inherently *defeasible*, so it fails to 'preserve truth' under all circumstances, which has always been considered a standard for logical reasoning.

Natural kinds have reappeared in AI as a practical problem of building taxonomic hierarchies for large knowledge bases that are allowed to have exceptions. The theory of reasoning in such taxonomies has been called *nonmonotonic inheritance* (see [10]). The guiding principle in resolving conflicts in such hierarchies was a *specificity principle*: more specific information should override more generic information in cases of conflict. Thus, a knowledge base may contain both *Birds fly* and *Penguins don't fly*, but then, given that Tweety is a penguin, we univocally infer that it does not fly, since *Birds fly* is a less specific claim. Though nonmonotonic inheritance relied more on graph-based representations than on traditional logical tools, it has managed to provide a plausible analysis of reasoning in this restricted context.

Nonmonotonicity of a different kind occurs in databases, logic programming and planning algorithms. A common assumption in such systems is that positive assertions that are not explicitly stated or derivable should be considered false. Thus, a database of students enrolled in a particular course implicitly presupposes that students that do not appear in the list are not enrolled in the course. Databases embody such negative information by appealing to the *closed word assumption*, which states that if a positive fact is not derivable from the database, its negation is assumed to hold. A similar principle is employed in programming languages for AI such as Prolog and Planner. Thus, in Prolog, the goal **not** *G* succeeds if the attempt to find a proof of G fails. Prolog's negation **not** is a nonmonotonic operator: if G is not provable from some axioms, it needn't remain nonprovable from an enlarged axiom set. This negation-as-failure has been used to implement important forms of commonsense reasoning, which eventually has led to developing modern declarative logic programming as a general representation formalism for AI (see [1]).

But first and foremost, nonmonotonicity has appeared in reasoning about actions. The main problem here was the *frame problem*: how efficiently determine which things remain the same in a changing world (e.g., a red block remains red after we have put it on top of another block). The frame problem arises in the context of predictive reasoning that is essential for planning and formalizing intelligent behavior, though neglected in traditional logic. Prediction involves the inference of later states from earlier ones. Changes in this setting do not merely occur, but occur for a reason. Furthermore, we usually assume that most things will be unchanged by the performance of an action. It is this *inertia assumption* that connects reasoning about action and change with NMR. What complicates the problem, however, is a *ramification problem*, the necessity of taking into account derived effects (ramifications) of actions. Suppose we have a suitcase with two locks, and it is opened if both locks are open. Then the action of opening one lock produces an indirect effect of opening the suitcase if the other lock is open. Such derived effects override the inertia assumption. The ramification problem has raised general questions about the role of causation in dynamic reasoning, and has led, eventually, to the so-called causal approach to the frame problem (see [8]).

Last but not least, there was the *qualification problem*, the problem of specifying conditions for a given action to have its intended effect. If I turn the ignition key in my car, I expect the car to start. However, many conditions have to be true for this: the battery must be alive, there must be gas in the tank, there is no potato in the tailpipe, etc. – an open-ended list of qualifications. Still, we normally *assume* that turning the key will start the car. This is obviously a special instance of a general philosophical problem of ceteris paribus laws, laws or generalizations that are valid under 'normal' circumstances which are usually impossible to specify exactly. It has become, however, an urgent practical problem for the representation of action and change in AI.

The above problems and their first solutions provided the starting point and basic objectives for the first nonmonotonic theories. These origins explain, in particular, an eventual discrepancy that has developed between NMR and commonsense reasoning. Though the latter has often appeared to be a promising way of solving AI problems, the study of 'artificial reasoning' need not be committed to it. Still, in trying to cope with principal commonsense reasoning tasks, the suggested formalisms have succeeded in capturing important features of the latter and thereby have broken new territory for logical reasoning. Today, nonmonotonic reasoning is not yet another application of logic, but a relatively independent field of logical research that has a great potential in informing, in turn, general logical theory and many areas of philosophical inquiry.

4.2 What Is Nonmonotonic Reasoning?

In everyday reasoning, we usually have incomplete information about a given situation, and we use a lot of assumptions about how things normally are in order to carry out further reasoning. Without such assumptions, it would be impossible to accomplish the simplest human reasoning tasks. Speaking generally, human reasoning is not reducible to collecting facts and deriving their consequences, but involves also making assumptions (and wholesale theories) about the world and acting in accordance with them. In this sense, commonsense reasoning is a simplified form of a general scientific methodology.

NMR is a theory of the rational use of assumptions. Now, assumptions are just beliefs, so they are abandoned when we learn new facts that contradict them. However, NMR assigns a special status to assumptions; it makes them *default* assumptions. Default assumptions are seen as always acceptable unless they conflict with current evidence. This presumptive reading has a semantic counterpart in the notion of *normality*; defaults are considered as holding for normal circumstances, and the nonmonotonic reasoning always assumes that the world is as normal as is compatible with known facts. This kind of belief commitment is a novel contribution of NMR to a general theory of reasoning.

This form of reasoning is distinct from deductive inference already because the latter is monotonic: if C is provable from a set a, it will be provable from a larger set $a \cup \{A\}$. Assumption-based reasoning is not monotonic, however, because adding new facts may invalidate some of the assumptions.

The default *Birds fly* is not a statement that is true or not of the world; some birds fly, some do not. Rather, it is an assumption used in building our theory of the world. NMR does not make any claims about the objective status of the assumptions it uses, so it does not depend on the objective confirmation of the latter. What it cares about, however, is the internal coherence of the choice of assumptions in particular situations. Of course, if we make an entirely inappropriate claim a default assumption, it will either be useless (inapplicable in most situations) or, worse, it may produce wrong conclusions. This makes nonmonotonic reasoning a risky business. Still, in most cases assumptions we make are useful and give desired results, and hence they are worth the risk of making an error. But what is even more important, more often than not we simply have no 'safe' replacement for such a reasoning strategy. That is why it is worth to teach robots and computers to reason in this way.

4.3 Two Problems of Default Assumptions

The primary problem of NMR is how we can make and consistently use default assumptions. Three initial nonmonotonic formalisms, namely circumscription [17], default logic [23] and modal nonmonotonic logic [20] have provided rigorous answers to this problem. The formalisms used three different languages – the classical language in circumscription,[1] a set of inference rules in default logic, and a modal language in modal nonmonotonic logic. Still, a common idea was to represent commonsense conditionals as ordinary conditionals with additional assumptions that could readily be accepted in the absence of contrary information. The differences between the three theories amounted, however, to different mechanisms of making default assumptions. In fact, default logic and modal nonmonotonic logics embodied the same nonmonotonic mechanism. However, the differences between both of them and circumscription were more profound. In order to articulate them, we should consider yet another important problem of default assumptions.

In order to preserve consistency of the resulting solutions, default assumptions should not be used when they contradict known facts and other defaults. Clearly, if a default plainly contradicts the facts, it should be 'canceled'. But if a number of defaults are jointly inconsistent with the facts, although each of them taken alone is consistent with them, then we have a *selection problem*: which of the defaults should be retained, and which abandoned in each particular case? An apparent solution is to choose all maximal consistent subsets of defaults; this solution was implicitly

[1] Circumscription amounts to using only *minimal* models satisfying a first-order description.

used in the circumscription approach of [17]. Unfortunately, it has turned out to be inadequate as a general solution to the selection problem. The main reason is that commonsense defaults are not born equal, and in most cases there is an additional structure of dependence and priority among the defaults themselves. As a result, not all consistent combinations of defaults turn out to be adequate as options for choice. We mentioned, for instance, that the choice of defeasible rules in nonmonotonic inheritance is constrained by the specificity principle: the two rules *Birds fly* and *Penguins don't fly* are jointly incompatible with a fact that Tweety is a penguin, but we univocally drop only the first rule in this situation, since it is a less specific claim than *Penguins don't fly*. Speaking generally, commonsense defaults involve much more structure than just a set of assumptions. That is why a solution to the primary problem of NMR, how to make default assumptions, does not necessary provide a solution to the selection problem. The latter requires a deeper understanding of the use of assumptions in commonsense reasoning.

A general way of handling the selection problem in the framework of circumscription, called *prioritized* circumscription, has been suggested by Lifschitz and endorsed in McCarthy [18]. The solution amounted to imposing priorities among minimized predicates. In fact, it was one of the origins of a general *preferential approach* to NMR (see below).

Default and modal nonmonotonic logics suggested a different, *explanatory* approach to the selection problem. In fact, this approach has 'borrowed' a much larger piece of commonsense methodology than circumscription. In both scientific and commonsense discourse, a particular law may fail to explain the actual outcome due to interference with other mechanisms and laws that contribute to the combined result. In other words, violations of laws are always *explainable* (at least in principle) by other laws that are active. It is this justificational aspect of reasoning that has been formalized in the notion of extension in default logic and corresponding models of modal nonmonotonic logic. An extension is a model generated by a set of defaults that is not only consistent, but also, and most importantly, explains away, or refutes, all other defaults that are left out. The latter requirement constitutes a very strong constraint on the coherence of potential choices, which goes far beyond plain consistency. Using this requirement, an explanatory theory can be 'tuned' to intended combinations of defaults by supplying the underlying logic with appropriate refutation rules for default assumptions. In a hindsight, this might be seen as one of the reasons why these formalisms have been relatively slow in realizing the complexity of the selection problem. In fact, the problem has 'survived' initial attempts of formalization, and has reappeared in a most dangerous form as a *Yale Shooting Anomaly* in Hanks and McDermott [9], where it was demonstrated that apparently plausible representations of defaults in default logic and other formalisms still do not provide an intended choice of assumptions for the solution of the frame problem. Nevertheless, despite initial, radically anti-logicist, reactions (cf. [19]), subsequent studies have shown that the Yale Shooting problem can be resolved, after all, in the framework of these formalisms.

4.4 Logic in Nonmonotonic Reasoning

The first nonmonotonic systems have re-shaped the initial contrast between NMR and logic. Namely, it has been shown that a nonmonotonic formalism can be defined by supplying some logical formalism with a *nonmonotonic semantics*, which forms a distinguished subset of the corresponding *logical* semantics determined by the logical formalism itself. Thus, for circumscription, the underlying logical formalism is just the classical logic (and its semantics), while the nonmonotonic semantics is given be the set of minimal models.

Unfortunately, this latter description has also brought to life a problematic 'shortcut' notion of *nonmonotonic logic* as a formalism determined directly by syntax and associated nonmonotonic semantics. On this view, a nonmonotonic logic has become just yet another logic determined by an unusual (nonmonotonic) semantics. However, this view has actually hindered in a number of ways an adequate understanding of nonmonotonic reasoning.

In ordinary logical systems, the semantics determines the set of logical consequences of a given theory, but also, and most importantly, it provides an interpretation for the syntax itself. Namely, it provides propositions and rules of a formalism with *meaning*, and its theories with *informational content*. By its very design, however, the nonmonotonic semantics is defined as a certain subset of logically possible models, and consequently it does not determine, in turn, the meaning of the propositions and rules of the syntax. Two radically different theories may (accidentally) have the same nonmonotonic semantics. Furthermore, such a difference cannot be viewed as apparent, since it may well be that by adding further rules or facts to both these theories, we obtain new theories that already have different nonmonotonic models (see [3] for further discussion).

The above situation is remarkably similar to the distinction between meaning (intension) and extension of logical concepts, a distinction that is fundamental for modern logic. Nonmonotonic semantics provides, in a sense, the extensional content of a theory in a particular context of its use. In order to determine the meaning, or informational content, of a theory, we have to consider all potential contexts of its use, and hence 'retreat' to the underlying logic. This distinction suggests the following more adequate understanding of nonmonotonic reasoning:

$$Nonmonotonic\ Reasoning = Logic\ + Nonmonotonic\ Semantics$$

Logic and its associated logical semantics are responsible for providing the meaning of the rules of the formalism, while the nonmonotonic semantics provides us with nonmonotonic consequences of a theory in particular situations.

In addition to a better understanding of the structure of nonmonotonic formalisms, the above two-layered structure has important benefits in comparing different formalisms. In particular, it allows us to see many of them as instantiations of the same nonmonotonic mechanisms in different underlying logics.

4.5 Preferential Nonmonotonic Reasoning

In solving the selection problem of default assumptions, preferential approach follows the slogan *"Choice presupposes preference"*, which makes it an instance of a general methodology that is at least as old as decision theory and the theory of social choice. According to this approach, the choice of assumptions should be made by establishing preference relations among them.

Generalizing prioritized circumscription, [25] defined a model preference logic based on an arbitrary preference ordering on interpretations.

Definition An interpretation i is a *preferred model* of A if it satisfies A and there is no better interpretation $j > i$ satisfying A. A *preferentially entails* B (written $A \mid\!\sim B$) if all preferred models of A satisfy B.

Shoham's approach was very appealing, and apparently suggested a unifying perspective on NMR. Kraus et al. [11] provided an axiomatization of such inference relations. This has established logical foundations for a research program that attracted many researchers both in AI and in logic. A detailed description of the preferential approach can be found in [15].

A representation of preferential entailment more suitable for real NMR can be based on the following model, where belief states correspond to admissible combinations of default assumptions (see [2]):

Definition An *epistemic state* is a triple $(\mathcal{S}, l, \prec)$, where $\mathcal{S}$ is a set of *belief states*, $\prec$ a preference relation on $\mathcal{S}$, while l is a labeling function assigning a deductively closed *belief set* to every belief state from $\mathcal{S}$.

Epistemic states can determine what to believe in particular situations. Changes in facts do not automatically lead to changes in epistemic states: the actual assumptions made in particular situations are obtained by choosing preferred belief states that are consistent with the facts.

A *preferentially entails* B in an epistemic state if $A \supset B$ holds in all preferred belief states consistent with A. Though apparently different from the original definition of Shoham, it is actually equivalent to the latter.

It is tempting to conclude from the above that preferential approach has assimilated nonmonotonic reasoning to plain deductive reasoning in a certain 'nonmonotonic' logic. This conclusion would be premature, however.

Preferential entailment is called nonmonotonic for the obvious reason that its rules do not admit Strengthening: $A \mid\!\sim B$ does not imply $A \wedge C \mid\!\sim B$. However, it is a monotonic, logical system in the more important sense that addition of new rules preserves previous derivations. Furthermore, the above semantics determines the *meaning* of conditionals, and hence preferential entailment describes precisely their *logic*. This inevitably implies, however, that it cannot capture the associated nonmonotonic reasoning with such defaults.

Preferential inference is severely sub-classical and does not allow us, for example, to infer *Red birds fly* from *Birds fly*. Clearly, there are good reasons for not accepting such a derivation as a *logical* rule; otherwise *Birds fly* would imply also

Penguins fly. Still, we could accept *Red birds fly* as a reasonable *default* conclusion from *Birds fly* in the absence of contrary information. By doing this, we would follow the general strategy of NMR of making reasonable assumptions on the basis of available information. This kind of reasoning will be defeasible, or *globally* nonmonotonic, since addition of new rules can block some of the conclusions made earlier. We can follow the idea of NMR also on the semantic side, namely by choosing 'most normal' epistemic states that satisfy a given set of conditionals. By doing this, we will accept rules that would not be derivable by preferential inference alone.

Summing up, the logic of preferential entailment should be extended to a nonmonotonic formalism by defining the associated nonmonotonic semantics. In fact, the literature is abundant with attempts to define such a theory.

Lehmann and Magidor [12] described a semantic construction, called *rational closure*, that allows us to make default conclusions from a set of conditionals.[2] This was a starting point in the quest for an adequate theory of defeasible entailment. A large number of modifications have been suggested, but a consensus has not been achieved. A general approach to this problem can be found in Geffner [6]. Finally, nonmonotonic inheritance (see [10]) can be viewed as a syntactic approach to defeasible entailment. Though it deals with conditionals restricted to literals, it has achieved a remarkable correspondence between what is derived and what is expected intuitively.

Most systems of defeasible entailment assume that classical implications corresponding to conditionals should serve as defaults in the associated nonmonotonic reasoning. Already this choice allows us to derive *Red birds fly* from *Birds fly* in the absence of conflicting information about redness. It is still insufficient, however, for capturing some further reasoning patterns. Suppressing details[3], what needs to be added here is a principled way of constructing a preference order on default sets. Recall, however, that establishing preferences among defaults is the main tool used by the preferential approach for resolving the selection problem of NMR. Accordingly the problem of defeasible entailment boils down again to the general selection problem for defaults. Unfortunately, this problem has turned out to be far from being trivial, or even univocal. Geffner's conditional entailment and nonmonotonic inheritance still remain the most plausible solutions suggested in the literature on preferential reasoning.

The preferential approach to NMR has suggested a powerful research program that significantly advanced our understanding of nonmonotonic reasoning and even of commonsense reasoning in general. Its most important achievement consists in formalizing a plausible logic of default conditionals that could serve as a logical basis for a full, nonmonotonic theory of defeasible reasoning. Unfortunately, it has not succeeded in achieving this latter goal.

---

[2] An equivalent construction, called system Z, has been suggested in Pearl [22].

[3] See [2].

4.6 Explanatory Nonmonotonic Reasoning

The explanatory approach encompasses almost all nonmonotonic formalisms that are actively investigated in AI today, including logic programming, argumentation and causal reasoning. Explanation can be seen as its basic ingredient. Propositions may not only hold in a model, but some of them are explainable (or caused) by other facts and rules. Furthermore, explanatory NMR is based on principles of *Explanation Closure* or *Causal Completeness* (see [24]), according to which any fact holding in a model should be explained.

By the above description, abduction and causation are integral parts of explanatory NMR. In some domains, explanatory reasoning adopts simplifying assumptions that exempt certain facts from the burden of explanation. Thus, the *Closed World Assumption* stipulates that negative assertions do not require explanation. In fact, minimization of models employed in McCarthy's circumscription can be seen as a by-product of this stipulation.

Simple default theories. Recall that a Tarski consequence relation is a set of rules $a \vdash A$ (where A is a conclusion, and a a set of premises) that satisfies the usual postulates. Its associated provability operator is $\mathrm{Cn}(u) = \{A \mid u \vdash A\}$. A consequence relation is *supraclassical* if it subsumes classical entailment.

For a set Δ of rules, let Cn_Δ denote the provability operator of the least supraclassical consequence relation containing Δ. Then $A \in \mathrm{Cn}_\Delta(u)$ precisely when A is derivable from u using the rules from Δ and classical entailment.

Now, a simple way of defining a nonmonotonic theory consists in combining a logical theory, given by a set of (Tarski) rules, and a set of default assumptions:

Definition A *simple default theory* is a pair $(\Delta, \mathcal{A})$, where Δ is a set of rules, and $\mathcal{A}$ a distinguished set of propositions called *defaults*.

Reasoning in this setting amounts to deriving plausible conclusions using rules and defaults. Explanatory reasoning requires here that a reasonable set of defaults explains why the rest of the defaults should be rejected.

Definition

- A set $\mathcal{A}_0$ of defaults is *stable* if and only if it is consistent and refutes any other default: $(\neg A) \in \mathrm{Cn}_\Delta(\mathcal{A}_0)$, for any $A \in \mathcal{A} \backslash \mathcal{A}_0$.
- A set s of propositions is an *extension* of a simple default theory iff $s = \mathrm{Cn}_\Delta(\mathcal{A}_0)$, for some stable set of defaults $\mathcal{A}_0$. Extensions determine the *nonmonotonic semantics* of a default theory.

Simple default theories provide a transparent description of explanatory NMR. Despite its simplicity, however, this formalism is equivalent to Reiter's default logic (see [4]). It is also closely related to the general argumentation (or assumption-based) framework of [5].

Generalizing the logic. For actual reasoning tasks of AI, we have to generalize the logical basis from Tarski rules to disjunctive rules $a \vdash b$, where b is a set of propositions. Informally, such a rule says that if all a's hold, then at least one of b's should hold. The theory of disjunctive inference is actually a well-developed part of general logical theory. A set of such rules forms a *Scott consequence relation* if and only if it satisfies the following postulates:

(Reflexivity) $A \vdash A$.
(Monotonicity) If $a \vdash b$ and $a \subseteq a'$, $b \subseteq b'$, then $a' \vdash b'$;
(Cut) If $a \vdash b, A$ and $a, A \vdash b$, then $a \vdash b$.

Let $\overline{u}$ denote the complement of a set u of propositions. Then u is a *theory* of a Scott consequence relation if $u \nvdash \overline{u}$.[4] A Scott consequence relation in a classical language is *supraclassical*, if it satisfies:

Supraclassicality If $a \vDash A$, then $a \vdash A$.
Falsity $\mathbf{f} \vdash$.

The Falsity postulate excludes, in effect, classically inconsistent models.

Simple default theories can be naturally extended to disjunctive rules. The resulting formalism will be equivalent to a disjunctive generalization of default logic [7], and even to powerful formalisms suggested in Lin and Shoham [14] and Lifschitz [13] as unified formalisms for nonmonotonic reasoning and logic programming.

Biconsequence Relations. For a detailed analysis of explanatory NMR, we can employ reasoning with respect to a *pair* of contexts. On the interpretation suitable for NMR, one of these contexts is the main (objective) one, while the other context provides assumptions that justify inferences in the main context.

A *bisequent* is an inference rule of the form $a : b \Vdash c : d$, where a, b, c, d are sets of propositions. On the explanatory interpretation, it says 'If a's hold then one of c's holds *provided* no b is assumed, and all d's are assumed'.

A *biconsequence relation* is a set of bisequents satisfying the rules:

Monotonicity $\dfrac{a : b \Vdash c : d}{a' : b' \Vdash c' : d'}$, if $a \subseteq a', b \subseteq b', c \subseteq c', d \subseteq d'$;

Reflexivity $A : \Vdash A :$ and $: A \Vdash : A$;

Cut $\dfrac{a : b \Vdash A, c : d \quad A, a : b \Vdash c : d}{a : b \Vdash c : d} \qquad \dfrac{a : b \Vdash c : A, d \quad a : A, b \Vdash c : d}{a : b \Vdash c : d}$.

A biconsequence relation can be seen as a product of two Scott consequence relations. A pair (u, v) of sets of propositions is a *bitheory* of a biconsequence relation if $u : \overline{v} \nVdash \overline{u} : v$. A set u is a *theory* if (u, u) is a bitheory. A bitheory (u, v) is *positively minimal*, if there is no bitheory (u', v) such that $u' \subset u$. Finally, a biconsequence relation is *supraclassical* if both its component contexts respect the classical entailment.

Nonmonotonic semantics of a biconsequence relation is a set of theories that are explanatory closed in the sense that all their propositions are explained (i.e., derived) when the theory itself is taken as the assumption context.

[4] Or, equivalently, if $a \vdash b$ and $a \subseteq u$, then $u \cap b \neq \emptyset$.

Definition A set u is an *extension* of a biconsequence relation, if (u, u) is a positively minimal bitheory. A *default nonmonotonic semantics* of a biconsequence relation is the set of its extensions.

A direct correspondence between default logic and biconsequence relations can be established by representing Reiter's default rules $a : b/A$ as bisequents $a{:}\neg b \Vdash A{:}$. Then the above nonmonotonic semantics will correspond precisely to the semantics of extensions in default logic. Moreover, many other nonmonotonic formalisms, such as logic programming, modal and autoepistemic logics, and the causal calculus can be expressed in this framework by varying the underlying logic (see [3] for details).

References and Recommended Readings

1. Baral, C. (2003). *Knowledge representation, reasoning and declarative problem solving*. Cambridge/New York: Cambridge University Press.
2. Bochman, A. (2001). *A logical theory of nonomonotonic inference and belief change*. New York: Springer.
3. Bochman, A. (2005). *Explanatory nonmonotonic reasoning*. Hackensack: World Scientific.
4. Bochman, A. (2008). Default logic generalized and simplified. *Annals of Mathematics and Artificial Intelligence, 53*, 21–49.
5. Bondarenko, A., Dung, P. M., Kowalski, R. A., & Toni, F. (1997). An abstract, argumentation-theoretic framework for default reasoning. *Artificial Intelligence, 93*, 63–101.
6. Geffner, H. (1992). *Default reasoning. Causal and conditional theories*. Cambridge: MIT Press.
7. Gelfond, M., Lifschitz, V., Przymusińska, H., & Truszczyński, M. (1991). Disjunctive defaults. In *Proceedings of Second International Conference on Principles of Knowledge Representation and Reasoning, KR'91*, Cambridge, MA (pp. 230–237).
8. Giunchiglia, E., Lee, J., Lifschitz, V., McCain, N., & Turner, H. (2004). Nonmonotonic causal theories. *Artificial Intelligence, 153*, 49–104.
9. Hanks, S., & McDermott, D. (1987). Non-monotonic logics and temporal projection. *Artificial Intelligence, 33*, 379–412.
10. Horty, J. F. (1994). Some direct theories of nonmonotonic inheritance. In D. M. Gabbay, C. J. Hogger, & J. A. Robinson (Eds.), *Handbook of logic in artificial intelligence and logic programming 3: Nonmonotonic reasoning and uncertain reasoning*. Oxford: Oxford University Press.
11. Kraus, S., Lehmann, D., & Magidor, M. (1990). Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence, 44*, 167–207.
12. Lehmann, D., & Magidor, M. (1992). What does a conditional knowledge base entail? *Artificial Intelligence, 55*, 1–60.
13. Lifschitz, V. (1994). Minimal belief and negation as failure. *Artificial Intelligence, 70*, 53–72.
14. Lin, F., & Shoham, Y. (1992). A logic of knowledge and justified assumptions. *Artificial Intelligence, 57*, 271–289.
15. Makinson, D. (2005). *Bridges from classical to nonmonotonic logic*. London: King's College Publications.
16. McCarthy, J. (1959). Programs with common sense. In *Proceedings of the Teddington Conference on the Mechanization of Thought Processes* (pp. 75–91). London: Her Majesty's Stationary Office.

17. McCarthy, J. (1980). Circumscription – a form of non-monotonic reasoning. *Artificial Intelligence, 13*, 27–39.
18. McCarthy, J. (1986). Applications of circumscription to formalizing common sense knowledge. *Artificial Intelligence, 13*, 27–39.
19. McDermott, D. (1987). Critique of pure reason. *Computational Intelligence, 3*(3), 149–160.
20. McDermott, D., & Doyle, J. (1980). Nonmonotonic logic. *Artificial Intelligence, 13*, 41–72.
21. Minsky, M. (1974). A framework for representing knowledge (Technical report 306). Artificial Intelligence Laboratory, MIT.
22. Pearl, J. (1990). System Z: A natural ordering of defaults with tractable applications to default reasoning. In *Proceedings of the Third Conference on Theoretical Aspects of Reasoning About Knowledge (TARK'90)* (pp. 121–135), San Mateo: Morgan Kaufmann.
23. Reiter, R. (1980). A logic for default reasoning. *Artificial Intelligence, 13*, 81–132.
24. Reiter, R. (2001). *Knowledge in action: Logical foundations for specifying and implementing dynamic systems*. Cambridge/London: MIT Press.
25. Shoham, Y. (1988). *Reasoning about change*. Cambridge, MA: Cambridge University Press.