Nonmonotonic logics and their algebraic foundations

Mirosław Truszczyński

Department of Computer Science
University of Kentucky
Lexington, KY 40506-0046, USA

Abstract. The goal of this note is to provide a background and references for the invited lecture presented at Computer Science Logic 2006. We briefly discuss motivations that led to the emergence of nonmonotonic logics and introduce two major nonmonotonic formalisms, default and autoepistemic logics. We then point out to algebraic principles behind the two logics and present an abstract algebraic theory that unifies them and provides an effective framework to study properties of nonmonotonic reasoning. We conclude with comments on other major research directions in nonmonotonic logics.

1 Why nonmonotonic logics

In the late 1970s, research on languages for knowledge representation, and considerations of basic patterns of commonsense reasoning brought attention to rules of inference that admit *exceptions* and are used only under the assumption of normality of the world in which one functions or to put it differently, when things are as expected.

For instance, a knowledge base concerning a university should support an inference that, given no information that might indicate otherwise, if Dr. Jones is a professor at that university, then Dr. Jones teaches. Such conclusion might be sanctioned by an inference rule stating that *normally* university professors teach. In commonsense reasoning rules with exceptions are ubiquitous. Planning our day and knowing we are to have lunch with a friend, we might use the following rule: *normally*, lunches end by 1:00pm. If nothing we know indicates that the situation we are in is not normal, we use this rule and conclude that our lunch will be over by 1:00pm.

The problem with such rules is that they do not lend themselves in any direct way to formalizations in terms of first-order logic, unless *all* exceptions are known and explicitly represented — an unrealistic expectation in practice. The reason is that standard logical inference is *monotone*: whenever a sentence α is a consequence of a set T of sentences then α is also a consequence of any set of sentences T' such that $T \subseteq T'$. On the other hand, it is clear that reasoning with normality rules when complete information is unavailable, is not monotone. In our lunch scenario, we may conclude that the lunch will be over by 1:00pm. However, if we learn that our friend will be delayed, the normality assumption is no longer valid our earlier inference is unsupported; we have to withdraw it.

Such reasoning, where additional information may invalidate conclusions, is called *nonmonotonic*. As we briefly noted above, it is common. It has been a focus of extensive studies by the knowledge representation community since the early eighties of the last century. This research developed along two major directions.

The first direction is concerned with the design of nonmonotonic logics — formalisms with direct ways to model rules with exceptions and with ways to use them. Arguably, two most studied nonmonotonic formalisms are default logic [1] and autoepistemic logic [2,3]. These two logics are the focus of this note. Our main goal in this paper is to introduce default and autoepistemic logics, identify algebraic principles that underlie them, and show that both logics can be viewed through a single abstract unifying framework of operators on complete lattices.

The second direction focused on studies of nonmonotone inference relations either in terms of classes of models or abstract postulates, the two perspectives being quite closely intertwined. Circumscription [4] and, more generally, preference logics [5] fall in this general research direction, as do studies of abstract properties of nonmonotonic inference relations [6,7,8,9]. Although outside our focus, for the sake of completeness, we will provide a few comments on preference logics and nonmonotonic inference relations in the last section of the paper.

2 Default Logic — an introduction

In his ground-breaking paper [1] Ray Reiter wrote: *Imagine a first order formalization of what we know about any reasonably complex world. Since we cannot know everything [...] — there will be gaps in our knowledge — this first order theory will be incomplete. [...] The role of a default is to help fill in some of the gaps in the knowledge base [...]. Defaults therefore function somewhat like meta-rules: they are instructions about how to create an* extension *of this incomplete theory. Those formulas sanctioned by the defaults and which extend the theory can be viewed as beliefs about the world. Now in general there are many different ways of extending an incomplete theory, which suggests that the default rules may be nondeterministic. Different applications of the defaults yield different extensions and hence different sets of beliefs about the world.*

According to Reiter defaults are meta-rules of the form "in the absence of any information to the contrary, assume ..." (hence, they admit exceptions), and default reasoning consists of applying them. Reiter's far-reaching contribution is that he provided a formal method to do so.

We will now present basic notions of default logic. We consider the language $\mathcal{L}(At)$ (or simply, $\mathcal{L}$) of propositional logic determined by a set At of propositional variables. A *default* is an expression

$$d = \frac{\alpha\colon \beta_1, \ldots, \beta_k}{\gamma}, \tag{1}$$

where α, β_i, $1 \leq i \leq k$, and γ are formulas from $\mathcal{L}$. We say that α is the *prerequisite*, β_i, $1 \leq i \leq k$, are *justifications*, and γ is the *consequent* of default d. If α is a tautology, we omit it from the notation. For a default d, we write $p(d)$, $c(d)$ and $j(d)$ for its prerequisite, consequent, and the set of justifications, respectively.

An informal reading of a default (1) is: *conclude γ if α holds and if all justifications β_i are possible*. In other words, to apply a default and assert its consequent, we must derive the prerequisite and establish that all justifications are possible. We will soon formalize this intuition. For now, we note that we can encode the rule arising in the university example by the following default:

$$\frac{prof_J : teaches_J}{teaches_J}$$

saying that if $prof_J$ holds and it is possible that $teaches_J$ holds (no information contradicts $teaches_J$), then $teaches_J$ does hold.

A *default theory* is a pair (D, W), where D is a set of defaults and W is a theory in the language $\mathcal{L}$. The role of W is to represent our knowledge (which is, in general, incomplete) while the role of defaults in D is to serve as "meta-rules" we might use to fill in gaps in what we know.

Let $\Delta = (D, W)$ be a default theory and let S be a propositional theory closed under consequence. If we start with S as our beliefs, Δ could be used to revise them. The revised belief set should contain W. Further, it should be closed under propositional consequence (to be a belief set) and under those defaults whose justifications are not contradicted by the current belief set S (are possible with respect to S). This revision process can be formalized by an operator Γ_Δ such that for a any set S of formulas (not necessarily closed under propositional consequence), $\Gamma_\Delta(S)$ is defined as the inclusion-least set U of propositional formulas satisfying the following conditions:

1. U is closed under propositional provability
2. $W \subseteq U$
3. for every default $d \in D$, if $p(d) \in U$ and for every $\beta \in j(d)$, $S \not\vdash \neg\beta$, then $c(d) \in U$.

Fixpoints of the operator Γ_Δ represent belief sets (by (1) they are indeed closed under propositional consequence) that are in a way *stable* with respect to Δ — they cannot be revised away. Reiter [1] proposed them as belief sets associated with Δ and called them *extensions*.

Definition 1. *Let Δ be a default theory. A propositional theory S is an* extension *of Δ if $S = \Gamma_\Delta(S)$.*

Let us look again at the university scenario, which we expand slightly. We know that Dr. Jones is a professor. We also know that if Dr. Jones is chair of the department then Dr. Jones does not teach. Finally we have the default rule saying that normally Dr. Jones teaches. This knowledge can be captured by a default theory (D, W), where

$$W = \{prof_J, chair_J \supset \neg teaches_J\}$$

and

$$D = \left\{ \frac{prof_J : teaches_J}{teaches_J} \right\}.$$

One can check that this default theory has only one extension and it contains $teaches_J$. However, if we append W by additional information that Dr. Jones is chair of the department ($chair_J$), then the resulting default theory has also one extension but it does not contain $teaches_J$, anymore (it contains $\neg teaches_J$). Thus, default theories with the semantics of extension can model nonmonotonic inferences.

Much of the theory of default logic is concerned with properties of extensions. A detailed studies of extensions can be found in [10,11].

3 Autoepistemic logic

Autoepistemic logic is a logic in a *modal* propositional language $\mathcal{L}_K(At)$ (or simply, $\mathcal{L}_K$), where At is the a of propositional variables and K stands for the modal operator. It was proposed to formalize how a rational agent with perfect introspection might construct belief sets [2,3].

The first modal nonmonotonic logic was introduced by McDermott and Doyle [12]. They proposed to use modal-free formulas to represent facts about an application domain, and "proper" modal formulas to encode nonmonotonic reasoning patterns. An informal reading of a modal formula $K\alpha$ is "α is believed" or "α is known." It suggests that a formula $\neg K\neg\alpha \supset \beta$ could be read "if $\neg\alpha$ is not believed (or, to put it differently, if α is possible) then β. Given this intuition, McDermott and Doyle [12] proposed to use the formula $\neg K\neg\alpha \supset \beta$ to represent a reasoning pattern "*in the absence of information contradicting* α*, infer* β" and gave a method to reason with such formulas supporting nonmonotonic inferences.

The logic of McDermott and Doyle was found to have counterintuitive properties [13,2,3]. Moore proposed autoepistemic logic [2,3] as a way to address this problem. As in the case of default logic, the goal was to describe a mechanism to assign to a theory belief sets that can be justified on its basis. Unlike in default logic, a specific objective for autoepistemic logic was to formalize belief sets a rational agent reasoning with perfect introspection might form.

Given a theory $T \subseteq \mathcal{L}_K$, Moore [3] defined an *expansion* of T to be a theory $E \subseteq \mathcal{L}_K$ such that

$$E = Cn(T \cup \{K\alpha \mid \alpha \in E\} \cup \{\neg K\alpha \mid \alpha \notin E\})$$

(Cn stands for the operator of propositional consequence which treats formulas $K\alpha$ as propositional variables). Moore justified this fixpoint equation by arguing that expansions should consist precisely of formulas that can be inferred from T and from formulas obtained by positive and negative introspection on the agent's beliefs.

Moore's expansions of T indeed have properties that make them adequate for modeling belief sets a rational agent reasoning with perfect introspection may built out of a theory T. In particular, expansions satisfy postulates put forth by Stalnaker [14] for belief sets in a modal language:

B1: $Cn(E) \subseteq E$ (rationality postulate)
B2: if $\alpha \in E$, then $K\alpha \in E$ (closure under positive introspection)
B3: if $\alpha \notin E$, then $\neg K\alpha \in E$ (closure under negative introspection).

Although motivated differently, autoepistemic logic can capture similar reasoning patterns as default logic does. For instance, the university example can be described in the modal language by a single theory

$$T = \{prof_J, chair_J \supset \neg teaches_J, Kprof_J \wedge \neg K\neg teaches_J \supset teaches_J\}.$$

This theory has exactly one expansion and it contains $teaches_J$. When extended with $chair_J$, the new theory also has just one expansion but it contains $\neg teaches_J$.

Examples like this one raised the question of the relationship between default and autoepistemic logics. Konolige suggested to encode a default

$$d = \frac{\alpha\colon \beta_1, \ldots, \beta_k}{\gamma}$$

with a modal formula

$$k(d) = K\alpha \wedge \neg K\neg\beta_1 \wedge \ldots \wedge \neg K\neg\beta_k \supset \gamma$$

and to represent a default theory $\Delta = (D, W)$ by a modal theory

$$k(\Delta) = W \cup \{k(d)\colon d \in D\}.$$

The translation seemed intuitive enough. In particular, it worked in the university example in the sense that extension of the default logic representation correspond to expansions of the modal logic representation obtained by translating the default logic one. However, it turned not to align extensions with expansions in general (a default theory $(\{\frac{p\colon q}{p}\}, \emptyset)$ has one extension but its modal counterpart has two expansions).

4 Default and autoepistemic logics — algebraically

Explaining the relationship between the two logics became a major research challenge. We will present here a recent algebraic account of this relationship [15]. As the first step, we will describe expansions and extensions within the framework of operators on the lattice of possible-world structures.

A *possible-world structure* is a set (possibly empty) of truth assignments to atoms in At. Possible-world structures can be ordered by the *reverse set inclusion*: for $Q, Q' \in \mathcal{W}$, $Q \sqsubseteq Q'$ if $Q' \subseteq Q$. The ordering $\sqsubseteq$ can be thought of as an ordering of increasing knowledge. As we move from one possible-world structure to another, greater with respect to $\sqsubseteq$, some interpretations are excluded and our knowledge of the world improves. We denote the set of all possible-world structures with $\mathcal{W}$. One can check that $\langle \mathcal{W}, \sqsubseteq \rangle$ is a complete lattice.

A possible-world structure Q and an interpretation I, determine the truth function $\mathcal{H}_{Q,I}$ inductively as follows:

1. $\mathcal{H}_{Q,I}(p) = I(p)$, if p is an atom.
2. $\mathcal{H}_{Q,I}(\varphi_1 \wedge \varphi_2) = \mathbf{t}$ if $\mathcal{H}_{Q,I}(\varphi_1) = \mathbf{t}$ and $\mathcal{H}_{Q,I}(\varphi_2) = \mathbf{t}$. Otherwise, $\mathcal{H}_{Q,I}(\varphi_1 \wedge \varphi_2) = \mathbf{f}$.
3. $\mathcal{H}_{Q,I}(\varphi_1 \vee \varphi_2) = \mathbf{t}$ if $\mathcal{H}_{Q,I}(\varphi_1) = \mathbf{t}$ or $\mathcal{H}_{Q,I}(\varphi_2) = \mathbf{t}$. Otherwise, $\mathcal{H}_{Q,I}(\varphi_1 \vee \varphi_2) = \mathbf{f}$.
4. $\mathcal{H}_{Q,I}(\neg\varphi) = \mathbf{t}$ if $\mathcal{H}_{Q,I}(\varphi) = \mathbf{f}$. Otherwise, $\mathcal{H}_{Q,I}(\varphi) = \mathbf{f}$.
5. $\mathcal{H}_{Q,I}(K\varphi) = \mathbf{t}$, if for every interpretation $J \in Q$, $\mathcal{H}_{Q,J}(\varphi) = \mathbf{t}$. Otherwise, $\mathcal{H}_{Q,I}(K\varphi) = \mathbf{f}$.

It is clear that for every formula $\varphi \in \mathcal{L}_K$, the truth value $\mathcal{H}_{Q,I}(K\varphi)$ does not depend on I. Thus, and we will denote it by $\mathcal{H}_Q(K\varphi)$, dropping I from the notation.

The *modal theory* of a possible-world structure Q, denoted by $Th_K(Q)$, is the set of all modal formulas that are believed in Q. Formally,

$$Th_K(Q) = \{\varphi\colon \mathcal{H}_Q(K\varphi) = \mathbf{t}\}.$$

The *(modal-free) theory* of Q, denoted $Th(Q)$, is defined by

$$Th(Q) = Th_K(Q) \cap \mathcal{L}.$$

(As an aside, we note here a close relation between possible-world structures and Kripke models with universal accessibility relations.)

Default and autoepistemic logics can both be defined in terms of fixpoints of operators on the lattice $\langle \mathcal{W}, \sqsubseteq\rangle$. A characterization of expansions in terms of fixpoints of an operator on $\mathcal{W}$ has been known since Moore [2]. Given a theory $T \subseteq \mathcal{L}_K$ and a possible-world structure Q, Moore defined a possible-world structure $D_T(Q)$ as follows:

$$D_T(Q) = \{I : \mathcal{H}_{Q,I}(\varphi) = \mathbf{t},\ \text{for every } \varphi \in T\}.$$

The intuition behind this definition is as follows (perhaps not coincidentally, as in the case of default logic, we again refer to belief-set revision intuitions). The possible-world structure $D_T(Q)$ is a revision of a possible-world structure Q. This revision consists of the worlds that are acceptable given the constraints on agent's beliefs captured by T. That is, the revision consists precisely of these worlds that make all formulas in T true (in the context of Q — the current belief state). Fixpoints of the operator D_T represent "stable" belief sets — they cannot be revised any further and so take a special role in the space of belief sets. It turns out [3] that they correspond to expansions!

Theorem 1. *Let $T \subseteq \mathcal{L}_K$. A theory $E \subseteq \mathcal{L}_K$ is an* expansion *of T if and only if $E = Th_K(Q)$, for some possible-world structure Q such that $Q = D_T(Q)$.*

A default theory defines a similar operator. With the Konolige's interpretation of defaults in mind, we first define a truth function on the set of all propositional formulas and defaults. Namely, for a propositional formula φ, we define

$$\mathcal{H}^{dl}_{Q,I}(\varphi) = I(\varphi),$$

and for a default $d = \frac{\alpha\colon \beta_1,\ldots,\beta_k}{\gamma}$, we set

$$\mathcal{H}^{dl}_{Q,I}(d) = \mathbf{t}$$

if at least one of the following conditions holds:

1. there is $J \in Q$ such that $J(\alpha) = \mathbf{f}$.
2. there is i, $1 \le i \le k$, such that for every $J \in Q$, $J(\beta_i) = \mathbf{f}$.
3. $I(\gamma) = \mathbf{t}$

(we set $\mathcal{H}^{dl}_{Q,I}(d) = \mathbf{f}$, otherwise).

Given a default theory $\Delta = (D, W)$, for a possible-world structure Q, we define a possible-world structure $D_\Delta(Q)$ as follows:

$$D_\Delta(Q) = \{I : \mathcal{H}_{Q,I}(\varphi) = \mathbf{t},\ \text{for every } \varphi \in W \cup D\}.$$

Do fixpoints of D_Δ correspond to extensions? The answer is no. Fixpoints of D_Δ correspond to *weak extensions* [16], another class of belief sets one can associate with default theories.

To characterize extensions a different operator is needed. The following definition is due (essentially) to Guerreiro and Casanova [17]. Let $\Delta = (D, W)$ be a default theory and let Q be a possible-world structure. We define $\Gamma'_\Delta(Q)$ to be the least possible-world structure Q' (with respect to $\sqsubseteq$) satisfying the conditions:

1. $W \subseteq Th(Q')$
2. for every default $d \in D$, if $p(d) \in Th(Q')$ and for every $\beta \in j(d)$, $\neg\beta \notin Th(Q)$, then $c(d) \in Th(Q')$.

One can show that $\Gamma'_\Delta(Q)$ is well defined. Moreover, for every possible-world structure Q,

$$Th(\Gamma'_\Delta(Q)) = \Gamma_\Delta(Th(Q))$$

Consequently, we have the following result connecting fixpoints of $\Gamma'_\Delta(Q)$ and extensions of Δ [17].

Theorem 2. *Let Δ be a default theory. A theory $S \subseteq \mathcal{L}$ is an extension of Δ if and only if $S = Th(Q)$ for some possible-world structure Q such that $Q = \Gamma'_\Delta(Q)$.*

Several questions arise. Is there a connection between the operators D_Δ and Γ'_Δ? Is there a counterpart to the operator Γ'_Δ in autoepistemic logic? Can these operators, their fixpoints and their interrelations be considered in a more abstract setting? What are abstract algebraic principles behind autoepistemic and default logics? We provide some answers in the next section.

5 Approximation theory

Possible-world structures form a complete lattice. As we argued, default and autoepistemic theories determine "revision" operators on this lattice. These operators formalize a view of a theory (default or modal) as a device for *revising* possible-world structures. Possible-world structures that are stable under the revision or, more formally, which are fixpoints of the revision operator give a semantics to the theory (of course, with respect to the revision operator used).

Operators on a complete lattice of propositional truth assignments and their fixpoints were used in a similar way to study the semantics of logic programs with negation. Fitting [18,19,20] characterized all major 2-, 3- and 4-valued semantics of logic programs, specifically, supported-model semantics [21], stable-model semantics [22], Kripke-Kleene semantics [18,23] and well-founded semantics [24], in terms of fixpoints of the van Emden-Kowalski operator [25,26] and its generalizations and variants.

These results suggested the existence of more general and abstract principles underlying these characterizations. [27,28] identified them and proposed a comprehensive unifying *abstract* framework of *approximating* operators as an algebraic foundation for nonmonotonic reasoning. We will now outline the theory of approximating operators and use it to relate default and autoepistemic logics. For details, we refer to [27,28].

Let $\langle L, \leq\rangle$ be a poset. An element $x \in L$ is a *pre-fixpoint* of an operator $O\colon L \to L$ if $O(x) \leq x$; x is a *fixpoint* of O if $O(x) = x$. We denote a least fixpoint of O (if it exists) by $\mathit{lfp}(O)$.

An operator $O : L \to L$ is *monotone* if for every $x, y \in L$ such that $x \leq y$, $O(x) \leq O(y)$. Monotone operators play a key role in the algebraic approach to nonmonotonic reasoning. Tarski and Knaster's theorem asserts that monotone operators on complete lattices (from now on L will always stand for a complete lattice) have least fixpoints [29].

Theorem 3. *Let L be a complete lattice and let O be a monotone operator on L. Then O has a least fixpoint and a least pre-fixpoint, and these two elements of L coincide. That is, we have $\mathit{lfp}(O) = \bigwedge\{x \in L\colon O(x) \leq x\}$.*

The *product bilattice* [30] of a complete lattice L is the set $L^2 = L \times L$ with the following two orderings $\leq_p$ and $\leq$:

1. $(x, y) \leq_p (x', y')$ if $x \leq x'$ and $y' \leq y$
2. $(x, y) \leq (x', y')$ if $x \leq x'$ and $y \leq y'$.

Both orderings are complete lattice orderings in L^2. For the theory of approximating operators, the ordering $\leq_p$ is of primary importance.

If $(x, y) \in L^2$ and $x \leq z \leq y$, then $(x, y) \in L^2$ *approximates* z. The "higher" a pair (x, y) in L^2 with respect to $\leq_p$, the more *precise* estimate it provides to elements it approximates. Therefore, we call this ordering the *precision* ordering. Most precise approximations are provided by pairs $(x, y) \in L^2$ for which $x = y$. We call such pairs *exact*.

For a pair $(x, y) \in L^2$, we define its *projections* as:

$$(x, y)_1 = x \qquad \text{and} \qquad (x, y)_2 = y.$$

Similarly, for an operator $A : L^2 \to L^2$, if $A(x, y) = (x', y')$, we define

$$A(x, y)_1 = x' \qquad \text{and} \qquad A(x, y)_2 = y'.$$

Definition 2. *An operator $A\colon L^2 \to L^2$ is* symmetric *if for every $(x, y) \in L^2$, $A(x, y)_1 = A(y, x)_2$; A is* approximating *if A is symmetric and $\leq_p$-monotone.*

Every approximating operator A on L^2 maps exact pairs to exact pairs. Indeed, $A(x, x) = (A(x, x)_1, A(x, x)_2)$ and, by the symmetry of A, $A(x, x)_1 = A(x, x)_2$.

Definition 3. *If A is an approximating operator and O is an operator on L such that for every $x \in L$ $A(x, x) = (O(x), O(x))$, then A is an* approximating operator for O.

Let $A\colon L^2 \to L^2$ be an approximating operator. Then for every $y \in L$, the operator $A(\cdot, y)_1$ (on the lattice L) is $\leq$-monotone. Thus, by Theorem 3, it has a least fixpoint. This observation brings us to the following definition.

Definition 4. *Let $A\colon L^2 \to L^2$ be an approximating operator. The* stable operator *for A, $\mathcal{C}_A$, is defined by*

$$\mathcal{C}_A(x, y) = (C_A(y), C_A(x)),$$

where $C_A(y) = \mathit{lfp}(A(\cdot, y)_1)$ (or equivalently, as A is symmetric, $C_A(y) = \mathit{lfp}(A(y, \cdot)_2)$).

The following result states two key properties of stable operators.

Theorem 4. *Let $A\colon L^2 \to L^2$ be an approximating operator. Then*

1. *$\mathcal{C}_A$ is $\leq_p$-monotone, and*
2. *if $\mathcal{C}_A(x,y) = (x,y)$, then $A(x,y) = (x,y)$.*

Operators A and $\mathcal{C}_A$ are $\leq_p$-monotone. By Theorem 3, they have least fixpoints. We call them the *Kripke-Kleene* and the *well-founded* fixpoints of A, respectively (the latter term is justified by Theorem 4).

Let A be an approximating operator for an operator O. An *A-stable fixpoint* of O is any element x such that (x,x) is a fixpoint of $\mathcal{C}_A$. By Theorem 4, if (x,x) is a fixpoint of $\mathcal{C}_A$ then it is a fixpoint of A and so, x is a fixpoint of O. Thus, our terminology is justified. The following result gathers some basic properties of fixpoints of approximating operators.

Theorem 5. *Let O be an operator on a complete lattice L and A its approximating operator. Then,*

1. *fixpoints of the operator $\mathcal{C}_A$ are minimal fixpoints of A (with respect to the ordering $\leq$ of L^2); in particular, A-stable fixpoints of O are minimal fixpoints of O*
2. *the Kripke-Kleene fixpoint of A approximates all fixpoints of O*
3. *the well-founded fixpoint of A approximates all A-stable fixpoints of O*

How does it all relate to default and autoepistemic logic? In both logics operators D_Δ and D_T have natural generalizations, $\mathcal{D}_\Delta$ and $\mathcal{D}_T$, respectively, defined on the lattice $\mathcal{W}^2$ — the product lattice of the lattice $\mathcal{W}$ of possible-world structures [15]. One can show that $\mathcal{D}_\Delta$ and $\mathcal{D}_T$ are approximating operators for the operators D_Δ and D_T. Fixpoints of operators $\mathcal{D}_\Delta$ and $\mathcal{D}_T$ and their stable counterparts define several classes of belief sets one can associate with default and autoepistemic theories.

Exact fixpoints of the operators $\mathcal{D}_\Delta$ and $\mathcal{D}_T$ (or, more precisely, the corresponding fixpoints of operators D_Δ and D_T) define the semantics of *expansions* (in the case of autoepistemic logic, proposed originally by Moore; in the case of default logic, expansions were known as weak extensions [16]). The stable fixpoints of the operators D_Δ and D_T define the semantics of *extensions* (in the case of default logic, proposed originally by Reiter, in the case of autoepistemic logic the concept was not identified until algebraic considerations in [15] revealed it). Finally, the Kripke-Kleene and the well-founded fixpoints provide three-valued belief sets that approximate expansions and extensions (except for [31], these concepts received essentially no attention in the literature, despite their useful computational properties [15]).

Moreover, these semantics are aligned when we cross from default to autoepistemic logic by means of the Konolige's translation. One can check that the operators $\mathcal{D}_\Delta$ and $\mathcal{D}_{k(\Delta)}$ coincide. The Konolige's translation preserves expansions, extensions, the Kripke-Kleene and the well-founded semantics. However, clearly, it does not align default extensions with autoepistemic expansions. Different principles underlie these two concepts. Expansions are fixpoints of the basic revision operator $\mathcal{D}_\Delta$ or $\mathcal{D}_T$, while extensions are fixpoints of the stable operators for $\mathcal{D}_\Delta$ or $\mathcal{D}_T$, respectively.

Properties of fixpoints of approximating operators we stated in Theorems 4 and 5 specialize to properties of expansions and extensions of default and autoepistemic theories. One can prove several other properties of approximating operators that imply known or new results for default and autoepistemic logics. In particular, one can generalize the notion of stratification of a default (autoepistemic) theory to the case of operators and obtain results on the existence and properties of extensions and expansions of stratified theories as corollaries of more general results on fixpoints of stratified operators [32,33].

Similarly, one can extend to the case of operators concepts of strong and uniform equivalence of nonmonotonic theories and prove characterization results purely in the algebraic setting [34].

6 Additional comments

In this note, we focused on nonmonotonic logics which use fixpoint conditions to define belief sets and we discussed abstract algebraic principles behind these logics. We will now briefly mention some other research directions in nonmonotonic reasoning.

Default extensions are in some sense minimal (cf. Theorem 5(1)) and minimality was identified early as one of the fundamental principles in nonmonotonic reasoning. McCarthy [4] used it to define *circumscription*, a nonmonotonic logic in the language of first-order logic in which entailment is defined with respect to minimal models only. Circumscription was extensively studied [35,36]. Computational aspects were studied in [37,38,39]; connections to fixpoint-based logics were discussed in [40,41,42].

Preferential models [5,8] generalize circumscription and provide a method to define nonmonotonic inference relations. Inference relations determined by preferential models were shown in [8] to be precisely inference relations satisfying properties of *Left Logical Equivalence*, *Right Weakening*, *Reflexivity*, *And*, *Or* and *Cautious Monotony*. Inference relations determined by *ranked* preferential models were shown in [9] to be precisely those preferential inference relations that satisfy *Rational Monotony*.

Default conditionals that capture statements "if α then normally β" were studied in [43,9,44]. [9,44] introduce the notion of rational closure of sets of conditionals as as a method of inference ([44] uses the term *system Z*).

Default extensions and autoepistemic expansions also define nonmonotonic inference relations. For instance, given a set of defaults D, we might say that a formula β can be inferred from a formula α given D if β is in every extension of the default theory $(D, \{\alpha\})$. A precise relationship (if any) between this and similar inference relations based on the concept of extension or expansion to preferential or rational inference relations is not know at this time. Discovering it is a major research problem.

We conclude this paper by pointing to several research monographs on nonmonotonic reasoning [45,46,10,47,11,48,49,50].

Acknowledgments

We acknowledge the support of NSF grant IIS-0325063 and KSEF grant 1036-RDE-008.

References

1. Reiter, R.: A logic for default reasoning. Artificial Intelligence **13** (1980) 81–132
2. Moore, R.: Possible-world semantics for autoepistemic logic. In: Proceedings of the Workshop on Non-Monotonic Reasoning. (1984) 344–354 Reprinted in: M. Ginsberg, ed., *Readings on Nonmonotonic Reasoning*, pages 137–142, Morgan Kaufmann, 1990.
3. Moore, R.: Semantical considerations on nonmonotonic logic. Artificial Intelligence **25** (1985) 75–94
4. McCarthy, J.: Circumscription — a form of non-monotonic reasoning. Artificial Intelligence **13** (1980) 27–39
5. Shoham, Y.: A semantical approach to nonmonotic logics. In: Proceedings of the Symposium on Logic in Computer Science, LICS-87, IEEE Computer Society (1987) 275–279
6. Gabbay, D.M.: Theoretical foundations for non-monotonic reasoning in expert systems. In: Proceedings of the NATO Advanced Study Institute on Logics and Models of Concurrent Systems, Springer (1989) 439–457
7. Makinson, D.: General theory of cumulative inference. In Reinfrank, M., de Kleer, J., Ginsberg, M., Sandewall, E., eds.: Nonmonotonic reasoning (Grassau, 1988). Volume 346 of Lecture Notes in Computer Science., Berlin-New York, Springer (1989) 1–18
8. Kraus, S., Lehmann, D., Magidor, M.: Nonmonotonic reasoning, preferential models and cumulative logics. Artificial Intelligence Journal **44** (1990) 167–207
9. Lehmann, D., Magidor, M.: What does a conditional knowledge base entail? Artificial Intelligence Journal **55** (1992) 1–60
10. Marek, W., Truszczyński, M.: Nonmonotonic Logic; Context-Dependent Reasoning. Springer, Berlin (1993)
11. Antoniou, G.: Nonmonotonic Reasoning. MIT Press (1997)
12. McDermott, D., Doyle, J.: Nonmonotonic logic I. Artificial Intelligence **13** (1980) 41–72
13. McDermott, D.: Nonmonotonic logic II: nonmonotonic modal theories. Journal of the ACM **29** (1982) 33–57
14. Stalnaker, R.: A note on nonmonotonic modal logic. Unpublished manuscript (1980)
15. Denecker, M., Marek, V., Truszczyński, M.: Uniform semantic treatment of default and autoepistemic logics. Artificial Intelligence Journal **143** (2003) 79–122
16. Marek, W., Truszczyński, M.: Relating autoepistemic and default logics. In: Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning (Toronto, ON, 1989), San Mateo, CA, Morgan Kaufmann (1989) 276–288
17. Guerreiro, R., Casanova, M.: An alternative semantics for default logic. Preprint. The 3rd International Workshop on Nonmonotonic Reasoning, South Lake Tahoe (1990)
18. Fitting, M.C.: A Kripke-Kleene semantics for logic programs. Journal of Logic Programming **2** (1985) 295–312
19. Fitting, M.C.: Bilattices and the semantics of logic programming. Journal of Logic Programming **11** (1991) 91–116
20. Fitting, M.C.: Fixpoint semantics for logic programming – a survey. Theoretical Computer Science **278** (2002) 25–51
21. Clark, K.: Negation as failure. In Gallaire, H., Minker, J., eds.: Logic and data bases. Plenum Press, New York-London (1978) 293–322
22. Gelfond, M., Lifschitz, V.: The stable semantics for logic programs. In: Proceedings of the 5th International Conference on Logic Programming, MIT Press (1988) 1070–1080
23. Kunen, K.: Negation in logic programming. Journal of Logic Programming **4** (1987) 289–308
24. Van Gelder, A., Ross, K., Schlipf, J.: The well-founded semantics for general logic programs. Journal of the ACM **38** (1991) 620–650

25. van Emden, M., Kowalski, R.: The semantics of predicate logic as a programming language. Journal of the ACM **23** (1976) 733–742
26. Apt, K., van Emden, M.: Contributions to the theory of logic programming. Journal of the ACM **29** (1982) 841–862
27. Denecker, M., Marek, V., Truszczyński, M.: Approximations, stable operators, well-founded fixpoints and applications in nonmonotonic reasoning. In Minker, J., ed.: Logic-Based Artificial Intelligence. Kluwer Academic Publishers (2000) 127–144
28. Denecker, M., Marek, V., Truszczyński, M.: Ultimate approximation and its application in nonmonotonic knowledge representation systems. Information and Computation **192** (2004) 84–121
29. Tarski, A.: Lattice-theoretic fixpoint theorem and its applications. Pacific Journal of Mathematics **5** (1955) 285–309
30. Ginsberg, M.: Multivalued logics: a uniform approach to reasoning in artificial intelligence. Computational Intelligence **4** (1988) 265–316
31. Baral, C., Subrahmanian, V.: Dualities between alternative semantics for logic programming and nonmonotonic reasoning (extended abstract). In Nerode, A., Marek, W., Subrahmanian, V., eds.: Logic programming and non-monotonic reasoning (Washington, DC, 1991), Cambridge, MA, MIT Press (1991) 69–86
32. Vennekens, J., Gilis, D., Denecker, M.: Splitting an operator: an algebraic modularity result and its applications to logic programming. In Lifschitz, V., Demoen, B., eds.: Logic programming, Proceedings of the 20th International Conference on Logic Programming, ICLP-04. (2004) 195–209
33. Vennekens, J., Denecker, M.: An algebraic account of modularity in id-logic. In Baral, C., Leone, N., eds.: Proceedings of 8th International Conference on Logic Programming and Nonmonotonic Reasoning, LPNMR-05. LNAI 3662 (2005) 291–303
34. Truszczyński, M.: Strong and uniform equivalence of nonmonotonic theories — an algebraic approach. In Doherty, P., Mylopoulos, J., Welty, C.A., eds.: Proceedings of the 10th International Conference on Principles of Knowledge Representation and Reasoning, KR 2006, AAAI Press (2006) 389–399
35. Lifschitz, V.: Pointwise circumscription. In M.Ginsberg, ed.: Readings in Nonmonotonic Reasoning, Los Altos, CA, Morgan Kaufmann (1987) 179–193
36. Lifschitz, V.: Circumscriptive theories: a logic-based framework for knowledge representation. Journal of Philosophical Logic **17** (1988) 391–441
37. Lifschitz, V.: Computing circumscription. In: Proceedings of IJCAI-85, Los Altos, CA, Morgan Kaufmann (1985) 121–127
38. Ginsberg, M.: A circumscriptive theorem prover. In: Nonmonotonic reasoning (Grassau, 1988). Volume 346 of Lecture Notes in Computer Science., Berlin, Springer (1989) 100–114
39. Przymusiński, T.C.: An algorithm to compute circumscription. Artificial Intelligence **38** (1989) 49–73
40. Gelfond, M., Przymusinska, H.: On the relationship between circumscription and autoepistemic logic. In: Proceedings of the 5th International Symposium on Methodologies for Intelligent Systems. (1986)
41. Bidoit, N., Froidevaux, C.: Minimalism subsumes default logic and circumscription. In: Proceedings of IEEE Symposium on Logic in Computer Science, LICS-87, IEEE Press (1987) 89–97
42. Imieliński, T.: Results on translating defaults to circumscription. Artificial Intelligence **32** (1987) 131–146
43. Lehmann, D.: What does a conditional knowledge base entail? In: Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning, KR-89, Morgan Kaufmann (1989) 212–222

44. Pearl, J.: System Z: A natural ordering of defaults with tractable applications to nonmonotonic reasoning. In: Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge, TARK-90, Morgan Kaufmann (1990) 121–135
45. Besnard, P.: An Introduction to Default Logic. Springer, Berlin (1989)
46. Brewka, G.: Nonmonotonic Reasoning: Logical Foundations of Commonsense. Volume 12 of Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge, UK (1991)
47. Brewka, G., Dix, J., Konolige, K.: Nonmonotonic Reasoning, An Overview. CSLI Publications (1997)
48. Bochman, A.: A Logical Theory of Nonmonotonic Inference and Belief Change. Springer, Berlin (2001)
49. Bochman, A.: Explanatory Nonmonotonic Reasoning. Volume 4 of Advances in Logic. World Scientific (2005)
50. Makinson, D.: Bridges from Classical to Nonmonotonic Logic. Volume 5 of Texts in Computing. King's College Publications (2005)