9

Autoepistemic Logic as a Basis for Automating Nonmonotonic Reasoning

ILKKA NIEMELÄ

ABSTRACT. Autoepistemic logic is employed as a basis for automating nonmonotonic reasoning. In particular, a recent variant of autoepistemic logic based on L-hierarchic expansions is put forward as an interesting unified approach. Basic results needed for automating autoepistemic reasoning are reviewed in a general setting where autoepistemic logic is defined on top of a given classical logic. A simple finitary characterization for autoepistemic reasoning is presented and recent decidability and complexity are surveyed. Building on the finitary characterization a novel general decision method for autoepistemic reasoning is developed. Then it is shown how default logic, circumscription, logic programs with classical negation, and truth maintenance systems can be embedded into autoepistemic logic based on L-hierarchic expansions using simple local translations. Hence, the novel decision method provides a unified approach to automating a wide range of formalizations of nonmonotonic reasoning.

1 Introduction

We study the problem of automating nonmonotonic reasoning, i.e., developing an automated theorem prover for nonmonotonic reasoning. Several forms of nonmonotonic reasoning have been proposed and in recent years research has been focused on relationships between different formalizations. Despite different starting points and motivations, leading approaches to nonmonotonic reasoning have turned out to be closely related. Our aim is to exploit the close relationships so that a single theorem prover could be used as a unified reasoning tool for a wide range of forms of nonmonotonic reasoning. An interesting approach is to choose some formalization of

Partiality, Modality, and Nonmonotonicity
Patrick Doherty ed.

nonmonotonic reasoning as a basic system and develop a theorem prover for it. Decision procedures for other forms of nonmonotonic reasoning can then be realized by reducing a decision problem in a given formalization to a decision problem in the basic system. In this work we consider the possibility of using a recently proposed variant of autoepistemic logic based on L-hierarchic expansions as the basic system and we develop a novel decision method for autoepistemic reasoning which can serve as a basis for a unified approach to automating nonmonotonic reasoning.

Besides autoepistemic logic there are many other interesting alternatives to be used as the basic system including logic programming, default logic (Reiter 1980), McDermott and Doyle-style nonmonotonic modal logics (Marek et al. 1991), as well as systems based on two modal operators (Lin and Shoham 1990, Lifschitz 1991). Autoepistemic logic provides a rather expressive basic language. This implies that other general approaches to nonmonotonic reasoning such as default logic can be embedded using simple local translations. When using a basic language with limited expressivity like that of logic programs more complicated translations are needed. There is a close correspondence between logic programs and autoepistemic and default logic (Gelfond and Lifschitz 1988, Gelfond and Lifschitz 1990). However, current implementations of logic programming systems do not provide direct implementations of fragments of autoepistemic logic and default logic as the systems are based on an operational semantics where difficult issues such as loop detection are not sufficiently addressed. Autoepistemic logic has a uniform syntax unlike default logic that has two categories of objects, first-order sentences and default rules. In such an approach, e.g., default rules about default rules are not straightforwardly expressible. Epistemic expressions saying that something is believed or is not believed turn out to be useful especially for queries and integrity constraints (Reiter 1990, Niemelä 1992b) but they are not available in default logic.

As opposed to other approaches to formalizing nonmonotonic reasoning based on modal logic, autoepistemic logic is rather uncomplicated. It is a direct extension of the underlying classical logic. The modal operator in autoepistemic logic is governed by the straightforward principles of positive and negative introspection and no modal axioms for expressing the interaction between the modal operator and the connectives of the underlying logic are needed. Moreover, recent results show that autoepistemic logic covers a range of prominent forms of nonmonotonic reasoning (Konolige 1988, Marek and Truszczyński 1989, Konolige 1989, Marek et al. 1991, Gelfond and Lifschitz 1988, Elkan 1990, Gelfond and Przymusinska 1990, Kakas and Mancarella 1990). However, some of the embeddings are rather complicated because the original autoepistemic logic allows quite weakly grounded conclusions that have to be eliminated in the embeddings by employing extra groundedness conditions (Konolige 1988, Marek and Truszczyński 1989) or complicated translations (Gottlob 1993). Using special groundedness conditions it is possible to capture a very tightly grounded form of autoepistemic reasoning that corresponds to default logic. However, the non-logical character of the extra conditions leads to dependence on the syntactic representation of the premises (Konolige 1988).

One of the leading attempts to find a satisfying nonmonotonic modal logic is based on a scheme introduced by McDermott and Doyle (1980, 1982) for obtaining a nonmonotonic modal logic from a given monotonic modal logic. This approach has been further investigated, e.g, by Schwarz (1990, 1992) and Marek et al. (1991). Schwarz (1990) shows that also Moore's original autoepistemic logic can be seen as a McDermott and Doyle-style nonmonotonic modal logic and Marek et al. (1991) prove that there is a range of nonmonotonic modal logic corresponding to Moore's autoepistemic logic. There are many nonmonotonic modal logics corresponding to default logic under suitable embeddings and some consistency restrictions (Marek and Truszczyński 1990, Truszczyński 1991a, Truszczyński 1991b) (see (Niemelä 1992b) for further discussion). Moore's autoepistemic logic as well as the scheme of McDermott and Doyle employ a fixed point equation for defining the possible sets of correct conclusions and McDermott and Doyle-style nonmonotonic logics have similar groundedness problems as Moore's autoepistemic logic. It does not seem possible to eliminate weakly grounded conclusions in McDermott and Doyle-style logics even if the fixed point equation is strengthened and nonmonotonic ground logics (Tiomkin and Kaminski 1990, Tiomkin and Kaminski 1991, Truszczyński 1991c, Konolige 1992) are considered. A more detailed discussion of the groundedness problems in McDermott and Doyle-style logics can be found in (Niemelä 1991).

These observations led the author to propose a variant of autoepistemic logic where correct conclusions are defined, instead of a fixed point equation, more constructively using enumerations of sentences (Niemelä 1991). The resulting sets of conclusions, L-hierarchic expansions, are very tightly grounded in the premises. Autoepistemic logic based on L-hierarchic expansions seems to offer a promising unified framework for several important forms of nonmonotonic reasoning. Default logic, circumscription, logic programs, truth maintenance systems, and abductive reasoning have been shown to be special cases of autoepistemic logic based on L-hierarchic expansions under simple local embeddings (Niemelä 1992b). On the other hand, L-hierarchic expansions do not suffer from the dependency on the syntactic representation of the premises which is typical of other strongly grounded forms of autoepistemic reasoning (Konolige 1988, Marek and Truszczyński 1989). In this work we employ autoepistemic logic based on

L-hierarchic expansions as the basis for a unified approach to automating nonmonotonic reasoning.

Several decision methods have been proposed for autoepistemic reasoning. Some of these (Moore 1988, Stärk 1990, Marek et al. 1991) are straightforward applications of a particular finitary characterization of autoepistemic reasoning, some (Niemelä 1988) use extensively a particular proof method and some of them (Marek and Truszczyński 1991a, Junker and Konolige 1990a, Marek and Truszczyński 1991b) are based on the idea of mapping a decision problem in autoepistemic logic into another problem. Except for a couple of approaches (Junker and Konolige 1990a, Marek et al. 1991) the methods treat only original Moore-style autoepistemic reasoning. In this work we develop a decision method for autoepistemic logic in a general setting, i.e., for an autoepistemic logic based on a given classical logic. The novel method directly addresses tightly grounded autoepistemic reasoning. At this stage we do not want to commit to a particular proof method and we aim at a conceptually clear method where the sources of complexity and targets for optimization are clearly identifiable. Furthermore, we try to avoid exponential worst-case space requirements which seem to be inherent in the more advanced methods (Niemelä 1988, Marek and Truszczyński 1991a, Junker and Konolige 1990a, Marek and Truszczyński 1991b). This aspect is further discussed in Section 5.

The rest of the chapter is organized as follows. We start by introducing Moore-style autoepistemic logic in the general setting considered in this work. A finitary characterization for Moore-style autoepistemic reasoning is presented. The characterization forms the basis of the computational treatment of autoepistemic reasoning. The enumeration-based method for defining autoepistemic reasoning is explained and L-hierarchic expansions are introduced. Fundamental results on decidability and computational complexity of autoepistemic reasoning are surveyed. Then a novel decision method for autoepistemic reasoning is developed. It is shown how default logic, circumscription, logic programs, and truth maintenance can be translated into autoepistemic logic so that the decision method for autoepistemic reasoning can be used as an approach to automating these forms of nonmonotonic reasoning as well. Finally we discuss approaches to nonmonotonic reasoning that are based on partial semantics and, in particularly, their relation to autoepistemic reasoning.

2 Moore-Style Autoepistemic Logic

In this section we introduce Moore-style autoepistemic logic in a general setting. Moore's (1985) original formulation starts from classical propositional logic and builds autoepistemic logic on top of it. In this work we study a general setting where autoepistemic logic is constructed on top of a given classical logic. In this section we first present a method for defining Moore-style autoepistemic logic on top of a given classical logic (Niemelä 1992a). Then we introduce a finitary characterization of autoepistemic reasoning based on the notion of full sets (Niemelä 1992a). This characterization forms the basis of the computational treatment of autoepistemic reasoning developed in Section 5.

Autoepistemic logic is a modal logic with an operator L that is interpreted as "is believed". Robert Moore (1985) developed autoepistemic logic as a reconstruction of McDermott and Doyle's (1980) nonmonotonic logic to avoid some peculiarities of this logic. Autoepistemic logic models the beliefs of an ideally rational fully introspective agent. The main objects of interest in autoepistemic logic are the possible sets of beliefs of the agent given a set of sentences as the initial assumptions or premises of the agent. The agent's ideal rationality implies that the agent believes exactly every logical consequence of the initial assumptions and her/his beliefs. Full introspection means that the agent is capable of using both *positive introspection* (if χ is a belief, so is $L\chi$) and *negative introspection* (if χ is not a belief, then $\neg L\chi$ [1] is). The following example illustrates the concept of autoepistemic reasoning.

Example 2.1 Consider an autoepistemic sentence

$$\neg Lp \rightarrow \neg p.$$

The intuitive reading of this sentence is that if the agent does not believe p, then $\neg p$ holds (a kind of a closed world assumption). Given the sentence as the initial assumption of the agent what is the set of beliefs of the agent? Does, for example, $\neg p$ belong to the set of beliefs? In fact, $\neg p$ is one of the beliefs of the agent. This can be shown informally as follows. The sentence p is not one of the beliefs of the agent because it is not a logical consequence of the initial assumption and the set of beliefs. Thus by negative introspection $\neg Lp$ is one of the beliefs of the agent and $\neg p$ is a logical consequence of this belief and the initial assumption. If p is added to the initial premises, $\neg p$ no longer belongs to the set of beliefs of the agent. ■

The informal description of the set of beliefs of the agent given above is circular: the set of beliefs is defined using the set of beliefs. This leads Moore (1985) to use a fixed point definition. The sets of beliefs are called *stable expansions* of the initial assumptions and they are defined as fixed points of an operator given with the aid of the logical consequence relation used by the agent. Stable expansions provide a notion of possible sets of

[1] We use the following notation for logical connectives: $\neg$ (negation), $\wedge$ (conjunction), $\vee$ (disjunction), $\rightarrow$ (implication), $\leftrightarrow$ (equivalence), $\forall$ (universal quantifier), and $\exists$ (existential quantifier).

autoepistemic conclusions from a set of premises as the possible sets of beliefs of the agent given the premises as the initial assumptions of the agent. The previous example shows that autoepistemic reasoning is nonmonotonic in the sense that the set of conclusions does not grow monotonically when the set of premises is extended.

In this work we consider a general setting where a classical logic CL is given and an autoepistemic logic $\mathrm{CL}_{\mathrm{ae}}$ is built on top of that. This kind of a general setting has already been studied (Niemelä 1992a). It turns out that there is a clear boundary between the underlying classical logic CL and the autoepistemic layer so that a Moore-style autoepistemic logic can be developed straightforwardly following the ideas of Moore's (1985) propositional autoepistemic logic. The stable expansions in the general case are defined by a similar fixed point equation as in the propositional case. The only difference is that the underlying consequence relation used by the agent is changed to reflect the change in the underlying classical logic.

Let $\mathcal{L}$ be the language of CL. We extend $\mathcal{L}$ by adding a monadic operator L not appearing in $\mathcal{L}$ and obtain an autoepistemic language $\mathcal{L}_{\mathrm{ae}}$ which is the language of $\mathrm{CL}_{\mathrm{ae}}$. The language $\mathcal{L}_{\mathrm{ae}}$ is defined recursively as $\mathcal{L}$ but with an extra formation rule:

$$\text{if } \varphi \in \mathcal{L}_{\mathrm{ae}} \text{ is a closed formula, then } L\varphi \in \mathcal{L}_{\mathrm{ae}}.$$

Thus we consider autoepistemic logics where quantification into a modal context is not allowed even if $\mathcal{L}$ contains quantifiers (e.g., formulae of the form $\forall x LP(x)$ are not allowed). To simplify the treatment we consider only sentences, i.e., closed formulae. So from here on we let $\mathcal{L}$ and $\mathcal{L}_{\mathrm{ae}}$ contain only closed formulae.

In the general setting the role of the underlying logic CL is to induce a logical consequence relation $\models_{\mathrm{ae}}$ that determines the strength of the rationality of the agent: an ideally rational agent believes the logical consequences given by $\models_{\mathrm{ae}}$. This consequence relation is a simple extension of the consequence relation $\models_{\mathrm{cl}}$ of CL where the $L\varphi$ formulae are treated like atomic formulae in the propositional calculus. It can be formalized straightforwardly in the following way using a decoupling approach which has been employed by Konolige (1988) in the case of autoepistemic logic based on a first-order logic (see also (Niemelä 1992a)).

We assume that the semantics of the underlying logic CL is defined in the usual way. A set of CL-structures $\mathcal{S}$ is given. Then a relation $\models_{\mathrm{cl}}$ between the structures and sentences in $\mathcal{L}$ is defined. Given a CL-structure I and a sentence $\varphi \in \mathcal{L}$, $I \models_{\mathrm{cl}} \varphi$ means that φ is true in the structure I. For a set of sentences Σ, $I \models_{\mathrm{cl}} \Sigma$ holds if and only if $I \models_{\mathrm{cl}} \varphi$ holds for every $\varphi \in \Sigma$. Then the consequence relation $\models_{\mathrm{cl}}$ of CL is formalized in the usual way: for a set of sentences Σ and a sentence φ, $\Sigma \models_{\mathrm{cl}} \varphi$ holds if and only if for every CL-structure I, $I \models_{\text{cl}} \Sigma$ implies $I \models_{\text{cl}} \varphi$. We also assume that CL contains the classical negation ($\neg$): $\neg\varphi$ is true if and only if φ is not true. This is because in autoepistemic logic CL_{ae} we must be able to talk about not believing.

The decoupling approach is used to extend this semantics to the richer language of CL_{ae} by dividing the structures in which the sentences of $\mathcal{L}_{\text{ae}}$ are interpreted into two parts. One part is a regular CL-structure that is used to interpret the operators of CL. The other is a set of sentences containing the beliefs of the agent. Thus the set of AE-structures is the set of pairs $\langle I, T\rangle$ where $I \in \mathcal{S}$ and $T \subseteq \mathcal{L}_{\text{ae}}$. The relation $\models_{\text{ae}}$ between AE-structures and sentences in $\mathcal{L}_{\text{ae}}$ is defined inductively on the formation of sentences as $\models_{\text{cl}}$ in the case of CL with the addition of the valuation rule for the new $L\varphi$ sentences:

$$\langle I, T\rangle \models_{\text{ae}} L\varphi \text{ if and only if } \varphi \in T.$$

$\langle I, T\rangle \models_{\text{ae}} \varphi$ means that a sentence φ is true in an AE-structure $\langle I, T\rangle$. So $L\varphi$ is true in $\langle I, T\rangle$ if and only if φ is one of the beliefs of the agent given by the set T. Using AE-structures we can define the logical consequence relation $\models_{\text{ae}}$ used by the agent.

Definition 2.2 For a set of sentences Σ and a sentence φ, $\Sigma \models_{\text{ae}} \varphi$ holds if and only if for all CL-structures I and for all sets of sentences $T \subseteq \mathcal{L}_{\text{ae}}$, $\langle I, T\rangle \models_{\text{ae}} \Sigma$ implies $\langle I, T\rangle \models_{\text{ae}} \varphi$.

This relation is very similar to the logical consequence relation $\models_{\text{cl}}$ of the underlying logic. It is monotonic, treats $L\varphi$ formulae like atomic formulae in the propositional calculus and behaves otherwise like $\models_{\text{cl}}$.

Example 2.3 Let the underlying logic CL be some propositional modal logic of knowledge. For example, let CL be **S4** (although the choice of a particular logic plays no role in the development). Hence the language of CL is that of the propositional logic extended with a monadic operator K and the semantics is given using Kripke structures (Chellas 1980, Hughes and Cresswell 1984). For a modal logic with one modality, the Kripke structures can be take as tuples $\langle W, R, P, w\rangle$ where W is a set of worlds, R is a binary relation on W, P is a function assigning a truth-value to every atomic proposition for each world in W, and w is a world. We build an autoepistemic logic CL_{ae} on top of CL. First, the language of CL is extended by a new monadic operator L. Hence, e.g., formulae of the form $K\neg LK\varphi \rightarrow \psi$ belong to the language $\mathcal{L}_{\text{ae}}$ of CL_{ae}. Second, we define the logical consequence relation $\models_{\text{ae}}$ using the decoupling approach. For this we specify when a formula φ in $\mathcal{L}_{\text{ae}}$ is true in an AE-structure $\langle I, T\rangle$ where I is a Kripke structure and $T \subseteq \mathcal{L}_{\text{ae}}$. This is done exactly as in CL when defining whether a formula is true in a world w ($\langle W, R, P, w\rangle \models_{\text{cl}} \varphi$) but with adding the new valuation rule for $L\varphi$ formulae.

Let $I = \langle W, R, P, w \rangle$ and $T \subseteq \mathcal{L}_{\text{ae}}$.

1. $\langle I, T \rangle \models_{\text{ae}} \varphi$ if and only if $\langle W, R, P, w \rangle \models_{\text{cl}} \varphi$ (i.e. if and only if φ is true in the world w) when φ is atomic.
2. $\langle I, T \rangle \models_{\text{ae}} \neg\varphi$ if and only if $\langle I, T \rangle \not\models_{\text{ae}} \varphi$.
3. $\langle I, T \rangle \models_{\text{ae}} \varphi \rightarrow \psi$ if and only if $\langle I, T \rangle \not\models_{\text{ae}} \varphi$ or $\langle I, T \rangle \models_{\text{ae}} \psi$.
4. $\langle I, T \rangle \models_{\text{ae}} K\varphi$ if and only if for all w' for which wRw', $\langle \langle W, R, P, w' \rangle, T \rangle \models_{\text{ae}} \varphi$.
5. $\langle I, T \rangle \models_{\text{ae}} L\varphi$ if and only if $\varphi \in T$.

As CL is **S4**, for example, $\{Kp\} \models_{\text{ae}} p$ but $\{Lp\} \not\models_{\text{ae}} p$. ■

Following Moore we define *stable expansions* of a set of premises employing the following fixed point equation. If the underlying logic CL is the classical propositional calculus, then this definition coincides with that of the original Moore-style stable expansions.

Definition 2.4 For a set of sentences Σ, a set of sentences Δ is a *stable expansion* of Σ if and only if

(1) $$\Delta = \{\varphi \in \mathcal{L}_{\text{ae}} \mid \Sigma \cup L\Delta \cup \neg L\overline{\Delta} \models_{\text{ae}} \varphi\}$$

where $L\Delta = \{L\varphi \mid \varphi \in \Delta\}$, $\neg\Delta = \{\neg\varphi \mid \varphi \in \Delta\}$, and $\overline{\Delta} = \mathcal{L}_{\text{ae}} - \Delta$. Thus $\neg L\overline{\Delta} = \{\neg L\varphi \mid \varphi \in \mathcal{L}_{\text{ae}} - \Delta\}$.

A stable expansion is the key concept in autoepistemic reasoning. But stable expansions are hard to handle and reason about. Firstly, stable expansions are infinite: if χ is in a stable expansion, then so are $L\chi, LL\chi, \ldots$ Secondly, stable expansions are defined non-constructively using a fixed point equation. For example, it is very hard to argue formally that a given set of premises has exactly one stable expansion. A finitary characterization of stable expansions would clarify their structure, provide a tool to reason about expansions and enable their computational treatment.

In this work we use a compact finitary characterization of Moore-style expansions based on the idea that an expansion is determined by the $L\varphi$ subformulae that appear in the set of premises. This characterization, presented in (Niemelä 1992a), builds on the notion of full sets and it is a refinement of an earlier approach developed by the author (Niemelä 1989, Niemelä 1990a, Niemelä 1990b). Similar characterizations have been independently proposed by several researchers (Moore 1988, Shvarts 1990, Stärk 1990). An overview of these approaches and others to characterizing stable expansions can be found in (Niemelä 1992a). The notion of full sets provides a simple and compact representation of expansions that is straightforwardly applicable in the general case considered in this work. The general case is treated in full detail in (Niemelä 1992a) where also proofs of correctness can be found. Here we present only the main definitions and results relevant to automating autoepistemic reasoning.

We introduce some notation. For an $L\chi$ formula, $\text{q-atom}(L\chi) = \text{q-atom}(\neg L\chi) = L\chi$ and for a set Λ, $\text{q-atom}(\Lambda) = \{\text{q-atom}(\varphi) \mid \varphi \in \Lambda\}$. The set of all subformulae of the form $L\chi$ of a formula φ is denoted by $\mathrm{Sf}^{\mathrm{L}}(\varphi)$ and $\mathrm{Sf}^{\mathrm{qL}}(\varphi)$ is the set of all subformulae of the form $L\chi$ of φ that are not in the scope of another L operator in φ. We call the formulae in $\mathrm{Sf}^{\mathrm{qL}}(\varphi)$ *major* subformulae of φ. For a set of formulae Σ, $\mathrm{Sf}^{\mathrm{L}}(\Sigma) = \bigcup_{\varphi\in\Sigma} \mathrm{Sf}^{\mathrm{L}}(\varphi)$ and similarly for $\mathrm{Sf}^{\mathrm{qL}}(\Sigma)$.

Example 2.5 Let the underlying logic be the classical propositional calculus and let $\Sigma = \{Lp, q \wedge L\neg L(Lp \wedge \neg q)\}$. Then

$$\begin{aligned}
&\mathrm{Sf}^{\mathrm{L}}(\Sigma) = \{Lp, L\neg L(Lp \wedge \neg q), L(Lp \wedge \neg q)\}\\
&\mathrm{Sf}^{\mathrm{qL}}(\Sigma) = \{Lp, L\neg L(Lp \wedge \neg q)\}\\
&\text{q-atom}(\neg L(Lp \wedge \neg q))) = L(Lp \wedge \neg q)
\end{aligned}$$

■

The finitary characterization employs a consequence relation $\models_{\mathrm{L}}$ that is defined recursively on top of the underlying consequence relation $\models_{\mathrm{ae}}$. The idea here is to incorporate the introspection principles of autoepistemic reasoning into the consequence relation so that $L\varphi$ sentences are no longer treated like atomic sentences.

Definition 2.6 Given a set of sentences Σ and a sentence φ, the relation $\Sigma \models_{\mathrm{L}} \varphi$ is defined by recursion on the complexity of φ as follows:

$$\Sigma \models_{\mathrm{L}} \varphi \text{ if and only if } \Sigma \cup \mathrm{SB}_{\Sigma}(\varphi) \models_{\mathrm{ae}} \varphi$$

where $\mathrm{SB}_{\Sigma}(\varphi) = \{L\chi \in \mathrm{Sf}^{\mathrm{qL}}(\varphi) \mid \Sigma \models_{\mathrm{L}} \chi\} \cup \{\neg L\chi \in \neg\mathrm{Sf}^{\mathrm{qL}}(\varphi) \mid \Sigma \not\models_{\mathrm{L}} \chi\}$.

The following example shows that $\models_{\mathrm{L}}$ is a nonmonotonic consequence relation.

Example 2.7 Let the underlying logic be the classical propositional calculus and let p be an atomic sentence. Then $\emptyset \models_{\mathrm{L}} \neg Lp$ holds. This can be verified as follows. First note that for a formula φ with no $L\psi$ subformulae, $\mathrm{SB}_{\Sigma}(\varphi) = \emptyset$ and $\Sigma \models_{\mathrm{L}} \varphi$ if and only if $\Sigma \models_{\mathrm{ae}} \varphi$. As $\mathrm{Sf}^{\mathrm{qL}}(p) = \emptyset$, $\emptyset \not\models_{\mathrm{L}} p$ holds because $\emptyset \not\models_{\mathrm{ae}} p$ holds. This implies that $\mathrm{SB}_{\emptyset}(\neg Lp) = \{\neg Lp\}$. Thus $\emptyset \models_{\mathrm{L}} \neg Lp$ holds.

On the other hand, $\{p\} \models_{\mathrm{L}} p$ holds because $\{p\} \models_{\mathrm{ae}} p$ holds. This implies that $\mathrm{SB}_{\{p\}}(\neg Lp) = \{Lp\}$. Thus $\{p\} \models_{\mathrm{L}} \neg Lp$ does not hold. ■

The finitary characterization of stable expansions relies on sets of sentences constructed from the $L\chi$ and $\neg L\chi$ major subformulae of the premises satisfying a special fullness condition. A full set serves as the kernel of a stable expansion; it uniquely characterizes the stable expansion.

Definition 2.8 For a set of sentences Σ, a set $\Lambda \subseteq \mathrm{Sf}^{\mathrm{qL}}(\Sigma) \cup \neg\mathrm{Sf}^{\mathrm{qL}}(\Sigma)$ is Σ-full if and only if the following two conditions hold for every $L\chi \in \mathrm{Sf}^{\mathrm{qL}}(\Sigma)$:

1. $\Sigma \cup \Lambda \models_{\mathrm{L}} \chi$ if and only if $L\chi \in \Lambda$.
2. $\Sigma \cup \Lambda \not\models_{\mathrm{L}} \chi$ if and only if $\neg L\chi \in \Lambda$.

It turns out that for a stable expansion Δ of Σ, $\Lambda = \{\chi \in L\Delta \cup \neg L\overline{\Delta} \mid \text{q-atom}(\chi) \in \mathrm{Sf}^{\mathrm{qL}}(\Sigma)\}$ is a Σ-full set. Conversely, given a Σ-full set Λ the set $\{\varphi \in \mathcal{L}_{\mathrm{ae}} \mid \Sigma \cup \Lambda \models_{\mathrm{L}} \varphi\}$ is a stable expansion of Σ. In fact, there is a one-to-one correspondence between full sets and stable expansions.

Theorem 2.9 (Niemelä 1992a) *Let Σ be a set of sentences of $\mathcal{L}_{\mathrm{ae}}$. Then a function SE_Σ defined as*

$$\mathrm{SE}_\Sigma(\Lambda) = \{\varphi \in \mathcal{L}_{\mathrm{ae}} \mid \Sigma \cup \Lambda \models_{\mathrm{L}} \varphi\}$$

gives a bijective mapping from the set of Σ-full sets to the set of stable expansions of Σ.

So $\mathrm{SE}_\Sigma(\Lambda) = \{\varphi \in \mathcal{L}_{\mathrm{ae}} \mid \Sigma \cup \Lambda \models_{\mathrm{L}} \varphi\}$ is the unique stable expansion Δ of Σ such that $\Lambda \subseteq L\Delta \cup \neg L\overline{\Delta}$ and a sentence φ belongs to the expansion induced by Λ if and only if $\Sigma \cup \Lambda \models_{\mathrm{L}} \varphi$. Full sets provide a technique for verifying the number of stable expansions of a set of premises. A set of premises has exactly as many stable expansions as it has full sets.

Example 2.10 Let the underlying logic CL be the propositional calculus. Let

$$\Sigma = \{\neg L\neg Lp \to p\},$$

where p is atomic. Then $\mathrm{Sf}^{\mathrm{qL}}(\Sigma) = \{L\neg Lp\}$ and there are only two possible candidates for the Σ-full sets: $\Lambda_1 = \{\neg L\neg Lp\}$ and $\Lambda_2 = \{L\neg Lp\}$. Both are Σ-full. For example, Λ_1 is full as $\Sigma \cup \Lambda_1 \not\models_{\mathrm{L}} \neg Lp$. So Σ has exactly two expansions $\mathrm{SE}_\Sigma(\{\neg L\neg Lp\})$ and $\mathrm{SE}_\Sigma(\{L\neg Lp\})$. ■

We can offer a (partial) answer to the question about a criterion for the sets of premises having exactly one stable expansion: a set of sentences has exactly one stable expansion if and only if it has exactly one full set. As a straightforward consequence, it follows that sets of premises containing only sentences without the L operator have exactly one stable expansion. This is because a set of sentences $\Sigma \subseteq \mathcal{L}$ has exactly one full set $\Lambda = \emptyset$. Note that in the propositional case $\mathrm{SE}_\Sigma(\emptyset) = \{\varphi \in \mathcal{L}_{\mathrm{ae}} \mid \Sigma \models_{\mathrm{L}} \varphi\}$ gives a compact characterization of Marek's (1989) operator E that produces the unique stable expansion of an L free set of premises. For a finite set of premises Σ, the set $\mathrm{Sf}^{\mathrm{qL}}(\Sigma)$ is finite and thus there are only a finite number of candidates for the full sets and thus there can be only finitely many stable expansions. In fact, a finite set of sentences Σ has at most 2^n stable expansions where n is the cardinality of $\mathrm{Sf}^{\mathrm{qL}}(\Sigma)$.

A full set can be captured in terms of the underlying consequence relation $\models_{\mathrm{ae}}$ but then all $L\varphi$ subformulae of the premises have be considered.

Proposition 2.11 (Niemelä 1992a) *Let $\Sigma \subseteq \mathcal{L}_{\mathrm{ae}}$ be a set of sentences. If $\Upsilon \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma) \cup \neg\mathrm{Sf}^{\mathrm{L}}(\Sigma)$ and for every $L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)$,*

1. $\Sigma \cup \Upsilon \models_{\text{ae}} \chi$ *if and only if* $L\chi \in \Upsilon$ *and*
2. $\Sigma \cup \Upsilon \not\models_{\text{ae}} \chi$ *if and only if* $\neg L\chi \in \Upsilon$,

then $\Lambda = \{\varphi \in \Upsilon \mid \text{q-atom}(\varphi) \in \text{Sf}^{\text{qL}}(\Sigma)\}$ *is a* Σ*-full set and for all* $\chi \in \mathcal{L}_{\text{ae}}$, $\Sigma \cup \Lambda \models_{\text{L}} \chi$ *if and only if* $\Sigma \cup \Upsilon \models_{\text{L}} \chi$.

3 Enumeration-Based Autoepistemic Reasoning

In order to formalize a tight notion of autoepistemic reasoning an enumeration-based method was proposed (Niemelä 1991) where autoepistemic reasoning is defined as a sequence of introspection steps of an ideally rational agent. The idea is to build a set of beliefs from premises Σ by applying introspection to sentences in the order given by an enumeration ε. A set $\mathbf{B}^{\varepsilon}(\Sigma)$ is constructed that contains all the results of introspection. The set $\mathbf{B}^{\varepsilon}(\Sigma)$ together with Σ induces the set of beliefs $\mathbf{SE}^{\varepsilon}(\Sigma)$ of an introspective agent who has introspected sentences in the order ε with Σ as the set of initial assumptions. Considering only a certain subclass of enumerations, *L-hierarchic* enumerations, a particularly interesting tightly grounded class of expansions is captured. In this section we first present the enumeration-based method in the general setting and then provide a finitary characterization of enumeration-based autoepistemic reasoning following the treatment in (Niemelä 1992a) where also the proofs of the theorems presented in this section can be found.

Definition 3.1 Let $\Sigma \subseteq \mathcal{L}_{\text{ae}}$. Let $\varepsilon = \psi_1, \psi_2, \ldots$ be an enumeration of sentences in $\mathcal{L}_{\text{ae}}$. Let $\mathbf{B}^{\varepsilon}_0(\Sigma) = \emptyset$. Define $\mathbf{B}^{\varepsilon}_{i+1}(\Sigma)$ for $i = 0, 1, \ldots$ as follows:

$$\mathbf{B}^{\varepsilon}_{i+1}(\Sigma) = \begin{cases} \mathbf{B}^{\varepsilon}_i(\Sigma) \cup \{L\psi_{i+1}\} \text{ if } \Sigma \cup \mathbf{B}^{\varepsilon}_i(\Sigma) \models_{\text{ae}} \psi_{i+1} \\ \mathbf{B}^{\varepsilon}_i(\Sigma) \cup \{\neg L\psi_{i+1}\} \text{ otherwise} \end{cases}$$

Finally let

$$\begin{aligned} \mathbf{B}^{\varepsilon}(\Sigma) &= \bigcup_{i=0}^{\infty} \mathbf{B}^{\varepsilon}_i(\Sigma) \\ \mathbf{SE}^{\varepsilon}(\Sigma) &= \{\varphi \in \mathcal{L}_{\text{ae}} \mid \Sigma \cup \mathbf{B}^{\varepsilon}(\Sigma) \models_{\text{ae}} \varphi\}. \end{aligned}$$

To guarantee that the introspection principles of autoepistemic reasoning are respected, i.e., that the set of beliefs $\mathbf{SE}^{\varepsilon}(\Sigma)$ induced by an introspection order ε satisfies the fixed point equation (1) for a stable expansion, we must require the enumeration ε to be acceptable with respect to the set of premises. This is because $\mathbf{SE}^{\varepsilon}(\Sigma)$ is not necessarily a stable expansion of Σ for every enumeration ε as shown in the following example.

Example 3.2 Let the underlying logic be the propositional calculus. Consider a set of premises

$$\Sigma = \{\neg L\neg p \rightarrow p\}$$

and take an enumeration $\varepsilon = p, \neg p, \ldots$ Then $\mathbf{B}_1^\varepsilon(\Sigma) = \{\neg Lp\}$ as $\Sigma \not\models_{\mathrm{ae}} p$ and $\mathbf{B}_2^\varepsilon(\Sigma) = \{\neg Lp, \neg L\neg p\}$ as $\Sigma \cup \mathbf{B}_1^\varepsilon(\Sigma) \not\models_{\mathrm{ae}} \neg p$. Thus $p \in \mathbf{SE}^\varepsilon(\Sigma)$ but $Lp \notin \mathbf{SE}^\varepsilon(\Sigma)$. So $\mathbf{SE}^\varepsilon(\Sigma)$ does not satisfy the fixed point equation (1) for a stable expansion. ■

Definition 3.3 An enumeration ε is Σ-*acceptable* if there is no i and no formula φ such that $\neg L\varphi \in \mathbf{B}_i^\varepsilon(\Sigma)$ but $\Sigma \cup \mathbf{B}_i^\varepsilon(\Sigma) \models_{\mathrm{ae}} \varphi$.

We assume that the underlying consequence relation $\models_{\mathrm{ae}}$ is compact, i.e., if $\Sigma \models_{\mathrm{ae}} \varphi$, there exists a finite subset Σ' of Σ such that $\Sigma' \models_{\mathrm{ae}} \varphi$. For compact logics the acceptability condition implies that the autoepistemic introspection principles are not violated.

Theorem 3.4 (Niemelä 1992a) *For all Σ-acceptable enumerations ε of $\mathcal{L}_{\mathrm{ae}}$, $\mathbf{SE}^\varepsilon(\Sigma)$ is a stable expansion of Σ.*

Enumeration-based expansions, i.e., sets of beliefs induced by acceptable enumerations, are a proper subclass of Moore-style expansions and in the propositional case they coincide with the $\mathbf{N}$-expansions (Marek et al. 1991) which are the same as iterative expansions (Marek and Truszczyński 1989) (see (Niemelä 1991)). Weakly grounded beliefs in $\mathbf{N}$-expansions (iterative expansions) can be eliminated by requiring that all subformulae of a given sentence φ that can effect the result of the introspection of φ must be examined before φ. This new groundedness requirement can be incorporated easily to the enumeration-based method by considering only introspection sequences that are *L-hierarchic*.

Definition 3.5 Let ε be an enumeration $\psi_1, \psi_2, \ldots$ of sentences in $\mathcal{L}_{\mathrm{ae}}$. The enumeration ε is *L-hierarchic* if and only if for all ψ_i, ψ_j in ε holds that if $L\psi_i$ is a subformula of ψ_j, then $i < j$.

A set of sentences Δ is an *L-hierarchic expansion* of Σ if and only if $\Delta = \mathbf{SE}^\varepsilon(\Sigma)$ for some Σ-acceptable L-hierarchic enumeration ε of $\mathcal{L}_{\mathrm{ae}}$.

Example 3.6 The enumeration $\varepsilon = L(\neg Lp \rightarrow q), q, p$ is not L-hierarchic since Lp is a subformula of $L(\neg Lp \rightarrow q)$ but p appears after $L(\neg Lp \rightarrow q)$ in ε. On the other hand, the enumeration $\varepsilon' = p, L(\neg Lp \rightarrow q), q$ is L-hierarchic. ■

L-hierarchic expansions provide an interesting tightly grounded subclass of Moore-style expansions. In the propositional case L-hierarchic expansions are always minimal with respect to their L free parts whereas $\mathbf{N}$-expansions and other McDermott and Doyle-style expansions can be non-minimal (Niemelä 1991). Furthermore, L-hierarchic expansions are not dependent on the syntactic representation of the premises in the sense that classically equivalent sets of premises have exactly the same set of expansions. This independence of the syntactic representation does not hold for other more strongly grounded classes of expansions introduced by Konolige (1988) and Marek and Truszczyński (1989) that satisfy the minimality requirement.

Enumeration-based and L-hierarchic expansions can be given a finitary characterization using the notion of full sets. It turns out that full sets corresponding to enumeration-based (L-hierarchic) expansions are exactly those which have an acceptable (L-hierarchic) inducing enumeration. Note that for a finite enumeration the test for acceptability can be stated in the following simple form.

Proposition 3.7 (Niemelä 1992a) *Let $\varepsilon_n = \psi_1, \ldots, \psi_n$ be an enumeration of a finite set. Then ε_n is Σ-acceptable if and only if for all $\neg L\varphi \in \mathbf{B}_n^{\varepsilon_n}(\Sigma), \Sigma \cup \mathbf{B}_n^{\varepsilon_n}(\Sigma) \not\models_{\mathrm{ae}} \varphi$.*

Theorem 3.8 (Niemelä 1992a) *Let $\Sigma \subseteq \mathcal{L}_{\mathrm{ae}}$ be a set of sentences such that $\mathrm{Sf}^{\mathrm{L}}(\Sigma)$ is finite. If an enumeration $\varepsilon_n = \psi_1, \ldots, \psi_n$ of the finite set $\{\varphi \mid L\varphi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ is Σ-acceptable, then $\Lambda = \{\varphi \in \mathbf{B}_n^{\varepsilon_n}(\Sigma) \mid \text{q-atom}(\varphi) \in \mathrm{Sf}^{\mathrm{qL}}(\Sigma)\}$ is a Σ-full set and there exists an Σ-acceptable enumeration ε of $\mathcal{L}_{\mathrm{ae}}$ which is an extension of ε_n such that $\mathbf{SE}^{\varepsilon}(\Sigma)$ is the unique stable expansion Δ of Σ such that $\Lambda \subseteq L\Delta \cup \neg L\overline{\Delta}$, i.e., $\mathbf{SE}^{\varepsilon}(\Sigma) = \mathrm{SE}_{\Sigma}(\Lambda) = \mathrm{SE}_{\Sigma}(\mathbf{B}_n^{\varepsilon_n}(\Sigma))$. Furthermore, if ε_n is L-hierarchic, then ε is, too.*

Theorem 3.9 (Niemelä 1992a) *Let $\Sigma \subseteq \mathcal{L}_{\mathrm{ae}}$ be a set of sentences. For every Σ-acceptable (L-hierarchic) enumeration ε of $\mathcal{L}_{\mathrm{ae}}$, there exists an Σ-acceptable (L-hierarchic) enumeration $\varepsilon_L = \psi_1, \psi_2, \ldots$ of the set $\{\varphi \mid L\varphi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ such that $\Lambda = \{\varphi \in \mathbf{B}^{\varepsilon_L}(\Sigma) \mid \text{q-atom}(\varphi) \in \mathrm{Sf}^{\mathrm{qL}}(\Sigma)\}$ is the unique Σ-full set for which $\Lambda \subseteq L\mathbf{SE}^{\varepsilon}(\Sigma) \cup \neg L\overline{\mathbf{SE}^{\varepsilon}(\Sigma)}$ holds and thus $\mathbf{SE}^{\varepsilon}(\Sigma) = \mathrm{SE}_{\Sigma}(\Lambda) = \mathrm{SE}_{\Sigma}(\mathbf{B}_n^{\varepsilon_L}(\Sigma))$.*

Example 3.10 Let the underlying logic be the propositional calculus. Consider the set of premises in Example 2.10

$$\Sigma = \{\neg L \neg Lp \to p\}$$

where p is atomic. The set Σ has two stable expansions which both are enumeration-based expansions and thus **N**-expansions as they are both characterizable using acceptable enumerations of $\{\varphi \mid L\varphi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\} = \{p, \neg Lp\}$:

- $\mathrm{SE}_{\Sigma}(\Lambda_1)$ by the enumeration $\varepsilon = \neg Lp, p$ and
- $\mathrm{SE}_{\Sigma}(\Lambda_2)$ by the enumeration $\varepsilon' = p, \neg Lp$.

Hence, $\Lambda_1 = \{\varphi \in \mathbf{B}_2^{\varepsilon}(\Sigma) \mid \text{q-atom}(\varphi) \in \mathrm{Sf}^{\mathrm{qL}}(\Sigma)\}$ where $\mathbf{B}_2^{\varepsilon}(\Sigma) = \{Lp, \neg L \neg Lp\}$ and $\Lambda_2 = \{\varphi \in \mathbf{B}_2^{\varepsilon'}(\Sigma) \mid \text{q-atom}(\varphi) \in \mathrm{Sf}^{\mathrm{qL}}(\Sigma)\}$ where $\mathbf{B}_2^{\varepsilon'}(\Sigma) = \{\neg Lp, L\neg Lp\}$. The non-minimal expansion $\mathrm{SE}_{\Sigma}(\Lambda_1)$ containing p is induced by the enumeration ε that is not L-hierarchic. This implies that Σ has only one L-hierarchic expansion, $\mathrm{SE}_{\Sigma}(\mathbf{B}_2^{\varepsilon'}(\Sigma)) = \mathrm{SE}_{\Sigma}(\Lambda_2)$. ■

4 Decidability and Complexity of Autoepistemic Reasoning

In this section we survey recent decidability and complexity results on autoepistemic reasoning. These results provide a framework where implementations of autoepistemic reasoning can be developed and assessed. We start by examining the notion of autoepistemic reasoning and provide a precise definition for it.

An expansion of a set of premises is an obvious candidate as the set of conclusions derivable from the premises. But as there are sets of premises with multiple expansions, correct autoepistemic conclusions are not uniquely defined by expansions. This leads to the concepts of cautious and brave reasoning. Given a class of expansions, a formula is derivable from a set of premises in the *cautious* sense if and only if it is in every expansion of the premises and in the *brave* sense if and only if it is at least in one of the expansions. We denote the derivability relations with respect to L-hierarchic expansions, enumeration-based expansions and Moore-style expansions by $\mid\!\sim_{\mathrm{c(LE)}}$, $\mid\!\sim_{\mathrm{b(LE)}}$, $\mid\!\sim_{\mathrm{c(E)}}$, $\mid\!\sim_{\mathrm{b(E)}}$, $\mid\!\sim_{\mathrm{c}}$, $\mid\!\sim_{\mathrm{b}}$, respectively. The fundamental decision problems in autoepistemic reasoning are cautious and brave derivability as well as the problem of deciding whether a given set of premises has an expansion. For Moore-style expansions, we denote the decision problem related to the cautious derivability relation $\mid\!\sim_{\mathrm{c}}$ by **DERc** and to the brave derivability relation $\mid\!\sim_{\mathrm{b}}$ by **DERb**. By **SE** we denote the decision problem whether a finite set of premises has a Moore-style stable expansion. Similarly, the problems **DERc**(E), **DERb**(E), **SE**(E) and **DERc**(LE), **DERb**(LE), **SE**(LE) are defined based on enumeration-based and L-hierarchic expansions, respectively.

The decidability and complexity of autoepistemic reasoning has been studied intensively in recent years. The decidability of propositional autoepistemic logic based on Moore-style expansions has been established independently by a number of researchers (Marek and Truszczyński 1988, Marek and Truszczyński 1991a, Moore 1988, Niemelä 1988). The author has shown that first-order autoepistemic logic where quantifying-in is not allowed is decidable if and only if the underlying first-order consequence relation is decidable (Niemelä 1990b). Moore (1988) as well as Junker and Konolige (1990a) report a similar result. Marek and Truszczyński (1988, 1991a) show that the problem of existence of a Moore-style expansion is **NP**-complete for a restricted subclass of propositional autoepistemic logic. In a later paper Marek and Truszczyński (1991b) establish further complexity results on restricted classes of autoepistemic logic. The author has shown that decision problems related to Moore-style expansions are on the second level of the polynomial time hierarchy (Niemelä 1989, Niemelä 1990b). Gottlob (1992) extends this result and establishes the corresponding completeness results. He also shows similar completeness results with respect to the second level of the polynomial time hierarchy for decision problems related to **N**-expansions (Marek and Truszczyński 1990) which coincide with iterative expansions (Marek and Truszczyński 1989). Reiter's default logic is closely related to several variants of autoepistemic logic. Gottlob (1992) and Stillman (1992) have independently established completeness results with respect to the second level of the polynomial time hierarchy for reasoning in Reiter's default logic. These results complement the findings of Kautz and Selman (1991) and Stillman (1990) which indicate that even very restricted subclasses of default reasoning are intractable. Eiter and Gottlob (1992b) establish complexity results for autoepistemic reasoning based on Konolige's moderately grounded expansions.

Using the finitary characterization of expansions it is possible to prove fundamental results on decidability and computational complexity of autoepistemic reasoning also in the general setting.

Theorem 4.1 (Niemelä 1992a) *If the consequence relation $\models_{\mathrm{ae}}$ is decidable, then $\models_{\mathrm{L}}$ is decidable as well as each of the problems* **DERc**, **DERb**, **SE**, **DERc**(E), **DERb**(E), **SE**(E), **DERc**(LE), **DERb**(LE), *and* **SE**(LE).

If the underlying classical logic CL is a first-order logic, then satisfiability in CL is polynomial time reducible to the decision problems in autoepistemic logic (Niemelä 1990b). This means that autoepistemic logic based on a full first-order calculus is not even semi-decidable, i.e., if the underlying logic CL is a full first-order calculus, then **DERc**, **DERb**, **SE**, **DERc**(E), **DERb**(E), **SE**(E), **DERc**(LE), **DERb**(LE), and **SE**(LE) are not recursively enumerable. If the underlying logic CL is the classical propositional calculus, then the result implies that **DERc**, **DERb**, **SE**, **DERc**(E), **DERb**(E), **SE**(E), **DERc**(LE), **DERb**(LE), and **SE**(LE) are all **NP**-hard.

The decision problems in autoepistemic reasoning can be related to the polynomial time hierarchy (Johnson 1990) as follows.

Theorem 4.2 (Niemelä 1992a) *If the underlying logic* CL *is the propositional calculus, then*

1. **SE**, **SE**(E), *and* **SE**(LE) *are all* Σ_2^p*-complete,*
2. **DERb**, **DERb**(E), *and* **DERb**(LE) *are all* Σ_2^p*-complete,*
3. **DERc**, **DERc**(E), *and* **DERc**(LE) *are all* Π_2^p*-complete.*

5 A Novel Decision Method

In this section we develop a novel decision method for autoepistemic reasoning in the general setting. We first present the method for autoepistemic logic based on L-hierarchic expansions and then indicate how to modify the approach to cope with enumeration-based expansions and Moore-style expansions. We end the section by discussing various optimization possibilities in the new decision method. We start the section by critically examining previous approaches to automating autoepistemic reasoning.

5.1 Previous Approaches

As discussed in the previous section, in the general setting autoepistemic reasoning is decidable if the consequence relation $\models_{\text{ae}}$ determined by the underlying classical logic is decidable. Moreover, decision procedures can be obtained from the finitary characterization based on full sets in a very straightforward fashion (see (Niemelä 1992a) for details). However, a decision procedure that is based directly on the finitary characterization is not very practical because in order to find an appropriate full set in a direct approach it is necessary to iterate over every subset of the set of $L\chi$ major subformulae of the premises. If the cardinality of a set is n, it has 2^n subsets. Even for $n = 10$ this means 1024 iterations each of which might be computationally costly, i.e., involve a number of calls to a theorem prover for the underlying consequence relation $\models_{\text{ae}}$. Theorem 4.2 indicates that any general decision method for autoepistemic reasoning has an exponential worst-case *time* complexity even in the propositional case. However, a practically oriented decision method should be able to take advantage of possible regularities in a given set of premises. For example, consider a set of premises

(2) $$\Sigma = \{p_0, Lp_0 \rightarrow p_1, Lp_1 \rightarrow p_2, \ldots, Lp_{n-1} \rightarrow p_n\}.$$

A direct approach would examine the 2^n subsets of $\{Lp_0, \ldots, Lp_{n-1}\}$ to find the unique Σ-full set $\Lambda = \{Lp_0, \ldots, Lp_{n-1}\}$. However, by carefully exploiting the monotonicity of the underlying consequence relation $\models_{\text{ae}}$ it is possible to find the full set without iterating over the subsets: for any Σ-full set Λ, $\Sigma \cup \Lambda \models_{\text{ae}} p_0$ holds and thus $Lp_0 \in \Lambda$ must hold which in turn implies $Lp_1 \in \Lambda$ and so on. The novel decision method aims to take full advantage of the monotonicity of $\models_{\text{ae}}$.

Some of the leading approaches to automating autoepistemic reasoning are based on mapping a nonmonotonic reasoning problem to another problem such as a truth maintenance problem (Junker and Konolige 1990a), a theorem proving problem (Marek and Truszczyński 1991b), or a constraint satisfaction problem (Ben-Eliyahu and Dechter 1991). These reductions provide valuable insights into the relationships between the corresponding problems and enable new techniques to be used for solving nonmonotonic reasoning problems. There are subclasses of nonmonotonic reasoning problems where the problem size increases only polynomially as a result of the reduction mapping (see, for example, (Junker and Konolige 1990b, Marek and Truszczyński 1991b, Ben-Eliyahu and Dechter 1991)). In the general case, however, the reductions can lead to an exponential increase in the problem size. Hence, in these reduction approaches the worst case *space* complexity becomes exponential. Moreover, the reductions are computationally complex even for restricted classes of nonmonotonic reasoning problems. The use of a normal form and the need to encode classical proofs can be seen as factors leading to the increase in the problem size and the high (time) complexity of the reductions. Note that because decision problems in nonmonotonic reasoning are complete problems with respect to the second level of the polynomial time hierarchy in the propositional case, there cannot exist a polynomial time reduction from propositional nonmonotonic reasoning to an **NP**-complete (or **co-NP**-complete) problem unless the polynomial time hierarchy collapses.

As an example we examine the approach of Junker and Konolige (1990a, 1990b) where nonmonotonic reasoning is reduced to a truth maintenance problem. In this approach the premises are required to be in a normal form with no nested L operators. In the worst case, if a set of premises is transformed into the normal form, its size might increase exponentially. Then the premises are mapped to a set of justifications in a truth maintenance system. The size of the resulting set of justifications can be exponential with respect to the size of the premises and the mapping is computationally complex. This is because the set of justifications corresponding to premises Σ includes all monotonic justifications

$$\varphi \leftarrow \omega_1, \ldots, \omega_n$$

where $L\varphi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)$ and $\{\omega_1, \ldots, \omega_n\}$ is a minimal *proof* of φ from $H = \{\omega \mid \neg L\alpha \vee L\beta_1 \vee \ldots \vee L\beta_m \vee \omega \in \Sigma - \mathcal{L}\}$, i.e., a minimal subset of H such that $(\Sigma \cap \mathcal{L}) \cup \{\omega_1, \ldots, \omega_n\} \models_{\mathrm{cl}} \varphi$ holds. For an $L\varphi$ subformula of the premises, φ can possess an exponential number of minimal proofs. In addition, it seems that it is computationally quite hard to find every such minimal proof. The problem of finding a proof of φ from a hypothesis set H given a background theory Σ, i.e., finding a set of $H' \subseteq H$ such that $\Sigma \cup H' \models_{\mathrm{cl}} \varphi$, is closely related to logic-based abduction (Poole 1988, Selman and Levesque 1990, Eiter and Gottlob 1992a) where usually the proof H' is require to be consistent with the background theory Σ. If the background theory is a set of clauses, then deciding whether there is a consistent proof of an atom from a set of atoms is Σ_2^p-complete in the propositional case (Eiter and Gottlob 1992a). As propositional brave autoepistemic reasoning is also Σ_2^p-complete (Niemelä 1992a), the search problem of finding some consistent proof of an atom is at least as hard as brave autoepistemic reasoning.

In Junker and Konolige's approach an expansion corresponds to a suitable model of the justifications. For example, strongly grounded expansions correspond to grounded models but deciding whether there is a grounded model is **NP**-complete (Elkan 1990). The exponential worst case space complexity and the high time complexity of the reductions raise questions about the usability of the reduction approaches as general decision methods for nonmonotonic reasoning.

5.2 The New Method

The novel decision method to be developed in this section employs the full set based finitary characterization of expansions. However, it differs considerably from the straightforward method for examining every candidate for full sets. In the novel approach the idea is to develop the full sets together as long as possible so that the same computation is not repeated. The Σ-full sets for L-hierarchic expansions can be found by examining the enumerations of the set $\{\varphi \mid L\varphi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$. However, the enumerations are not important but the sets of beliefs generated by the enumerations. This observation cuts the combinatorial explosion considerably as typically there is a large number of enumerations generating the same Σ-full set[2]. Moreover, the novel method takes full advantage of the monotonicity of the underlying consequence relation $\models_{\mathrm{ae}}$.

The novel decision method is quite simple and uses only a theorem prover for the underlying monotonic consequence relation $\models_{\mathrm{ae}}$ as a subroutine. It offers an interesting framework for developing further optimization techniques. The novel decision method does not suffer from the problem of exponential space requirements. The space needed by the theorem prover for the underlying monotonic consequence relation $\models_{\mathrm{ae}}$ dominates the space requirements: if the theorem prover works in polynomial space, then the whole method works in polynomial space. Hence, for example, propositional autoepistemic reasoning can be implemented in polynomial space employing the new method.

Figure 1 presents the function $\mathbf{der}_{\mathrm{LE}}$ which is the main part of the novel decision method. The function $\mathbf{der}_{\mathrm{LE}}$ is a skeleton for the decision procedures for brave and cautious reasoning as well as for checking the existence of expansions. For that purpose it contains an unspecified function **test**. By changing this function the various decision procedures are obtained. The function $\mathbf{der}_{\mathrm{LE}}$ takes as input a set of premises Σ, sets B and F which give the common part of the Σ-full sets to be considered and a sentence φ which is just passed as an argument to the function **test**.

The purpose of $\mathbf{der}_{\mathrm{LE}}$ is to return true if and only if there exists an L-hierarchic expansion Δ of Σ containing $B \cup F$ such that $\mathbf{test}(\Sigma, \Lambda, \varphi)$ returns true where Λ is the Σ-full set corresponding Δ, i.e., $\Delta = \mathrm{SE}_{\Sigma}(\Lambda)$. This is achieved by constructing Σ-full sets which contain $B \cup F$ and cor-

[2]In this section we use the term full somewhat loosely. Often a set $\Upsilon \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma) \cup \neg\mathrm{Sf}^{\mathrm{L}}(\Sigma)$ satisfying the two conditions in Proposition 2.11 is called a Σ-full set instead of the actual Σ-full set $\{\varphi \in \Upsilon \mid \text{q-atom}(\varphi) \in \mathrm{Sf}^{\mathrm{qL}}(\Sigma)\}$.

function $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$
input: Finite sets of sentences Σ, B, and F and a sentence φ.

$\langle B, F\rangle := \mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$;
if for some $\neg L\chi \in B$, $\Sigma \cup B \cup F \models_{\mathrm{ae}} \chi$ **then**
  return false
else if $\mathrm{Sf}^{\mathrm{L}}(\Sigma) \subseteq$ q-atom$(B) \cup F$ **then**
  if $F \subseteq B$ **then**
    return $\mathbf{test}(\Sigma, B, \varphi)$
  else
    return false
  endif
else if for some $L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma) - ($q-atom$(B) \cup F)$,
    $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq$ q-atom(B) **then**
  take some $L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma) - ($q-atom$(B) \cup F)$ with $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq$ q-atom(B);
  if $\mathbf{der}_{\mathrm{LE}}(\Sigma, B \cup \{\neg L\chi\}, F, \varphi)$ returns false **then**
    return $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F \cup \{L\chi\}, \varphi)$
  else
    return true
  endif
else
  return false
endif.

FIGURE 1 A Skeleton for the Decision Procedures for Autoepistemic Reasoning Based on L-hierarchic Expansions respond to L-hierarchic expansions until a full set Λ is found for which $\mathbf{test}(\Sigma, \Lambda, \varphi)$ returns true.

In the correctness proof of $\mathbf{der}_{\mathrm{LE}}$ we use the following two notions. An enumeration ε is *fully L-hierarchic* if for every formula φ in ε for each subformula $L\chi$ of φ, χ appears before φ in ε. For example, the enumeration p, LLp is not fully L-hierarchic but p, Lp, LLp is. The set B is *generated* by some enumeration ε if $B = \mathbf{B}^{\varepsilon}(\Sigma)$.

In the decision method the idea is to expand the set B so that finally it contains every formula in a full set. By Theorem 3.9 for every L-hierarchic expansion, the corresponding full set is generated by some fully L-hierarchic enumeration from Σ. This implies that when deciding how to extend B the number of possibilities can be reduced by considering only those extensions which are generated by some fully L-hierarchic enumeration from Σ. Hence in the decision method B is extended only if the resulting set B' is generated by some fully L-hierarchic enumeration ε_m of a subset of $\{\varphi \mid L\varphi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ from Σ, i.e., $B' = \mathbf{B}_m^{\varepsilon_m}(\Sigma)$. The set F contains $L\chi$ sentences which should be included in the full set to be constructed but are "frozen": a sentence $L\chi \in F$ is added to B only after $\Sigma \cup B \models_{\mathrm{ae}} \chi$ holds and B *covers* all $L\chi'$ subformulae of χ, i.e., $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq \text{q-atom}(B)$. This ensures that B remains generatable from Σ by a fully L-hierarchic enumeration.

In $\mathbf{der}_{\mathrm{LE}}$ the results of positive introspection are first propagated as far as possible using the function $\mathbf{pos}_{\mathrm{LE}}$ presented in Figure 2. The function exploits the monotonicity of the underlying consequence relation $\models_{\mathrm{ae}}$ and expands the sets B and F, i.e., if B is included in a full set and $\Sigma \cup B \models_{\mathrm{ae}} \chi$, then $L\chi$ belongs to the full set. It is straightforward to verify that $\mathbf{pos}_{\mathrm{LE}}$ has the following properties.

Proposition 5.1 *Let* $\langle B', F' \rangle = \mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$.

(1) $B \subseteq B'$ *and* $F \subseteq F'$ *and if* $\text{q-atom}(B) \cup F \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$, *then* $\text{q-atom}(B') \cup F' \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$.

(2) If there is some fully L-hierarchic enumeration generating B *from* Σ, *then there is a fully L-hierarchic enumeration generating* B' *from* Σ.

(3) If $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ *for some* Σ*-acceptable L-hierarchic enumeration* ε_n *of* $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$, *then* $B' \cup F' \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$.

(4) For all $L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)$, *if* $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq \text{q-atom}(B')$ *and* $\Sigma \cup B' \models_{\mathrm{ae}} \chi$, *then* $L\chi \in \text{q-atom}(B')$.

Example 5.2 Consider the set of premises Σ in (2) and the execution of $\mathbf{pos}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset)$. The set B is gradually expanded by adding consecutively $Lp_0, Lp_1, \ldots$ and finally $\langle \{Lp_0, Lp_1, \ldots, Lp_{n-1}\}, \emptyset \rangle$ is returned. ■

After positive introspection if there is a *conflict*, i.e., if some $\neg L\chi$ has been included in the full set to be constructed ($\neg L\chi \in B$) but χ becomes a consequence of the premises Σ and the full set under construction ($\Sigma \cup$

function $\mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$
input: Finite sets of sentences Σ, B, and F.

repeat
  $B_o := B$;
  for every $L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma) - \text{q-atom}(B_o)$ **do**
    if $\Sigma \cup B \models_{\mathrm{ae}} \chi$ **then**
      if $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq \text{q-atom}(B)$ **then**
        $B := B \cup \{L\chi\}$
      else
        $F := F \cup \{L\chi\}$
      endif
    endif
  endfor
until $B = B_o$;
return $\langle B, F\rangle$.

FIGURE 2 The Function $\mathbf{pos}_{\mathrm{LE}}$ for Positive Introspection Based on L-hierarchic Enumerations

$B \cup F \models_{\mathrm{ae}} \chi$), there is no full set containing $B \cup F$ and false is returned. If there are no conflicts, $B \cup F$ covers the $L\chi$ subformulae of Σ and all the frozen beliefs F have been included in B, B is a full set corresponding to an L-hierarchic expansion and $\mathbf{der}_{\mathrm{LE}}$ returns what $\mathbf{test}(\Sigma, B, \varphi)$ returns.

If there is some $L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)$ not covered by $B \cup F$, there are two possibilities: either $L\chi$ or $\neg L\chi$ should be included in the full set. To ensure that B extended by $\neg L\chi$ is generated by some fully L-hierarchic enumeration, only $L\chi$ formulae for which $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq \text{q-atom}(B)$ holds are considered. Then we know that if B is generated by a fully L-hierarchic enumeration ε, then $B \cup \{\neg L\chi\}$ is generated by a fully L-hierarchic enumeration, i.e., ε extended by χ. This is because $\Sigma \cup B \not\models_{\mathrm{ae}} \chi$ holds as otherwise $L\chi$ would have been added to B in the function $\mathbf{pos}_{\mathrm{LE}}$ by Proposition 5.1 (4).

The two possibilities for such a formula $L\chi$ are handled by backtracking. First B is extended by $\neg L\chi$ and if $\mathbf{der}_{\mathrm{LE}}(\Sigma, B \cup \{\neg L\chi\}, F, \varphi)$ returns false, $L\chi$ is taken as a frozen belief and $\mathbf{der}_{\mathrm{LE}}$ returns what $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F \cup \{L\chi\}, \varphi)$ returns.

Example 5.3 (i) Consider the execution of $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ where Σ is given in (2) and φ is an arbitrary sentence. As $\mathbf{pos}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset)$ returns $\langle\{Lp_0, Lp_1, \ldots, Lp_{n-1}\}, \emptyset\rangle$, backtracking is not needed at all and the return value of $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ is that of

$$\mathbf{test}(\Sigma, \{Lp_0, Lp_1, \ldots, Lp_{n-1}\}, \varphi).$$

(ii) Consider the premises $\Sigma = \{\neg Lp \to q, \neg Lq \to p\}$ and the execution of $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$. Here $\mathbf{pos}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset)$ returns $\langle \emptyset, \emptyset \rangle$. As $\mathrm{Sf}^{\mathrm{L}}(\Sigma) = \{Lp, Lq\}$ is not covered by $\emptyset$, a new $L\chi$ formula must be chosen. Since p and q have no $L\chi$ subformulae, we can select either of them. Let us choose Lp. Hence $\mathbf{der}_{\mathrm{LE}}(\Sigma, \{\neg Lp\}, \emptyset, \varphi)$ is executed next. Now $\mathbf{pos}_{\mathrm{LE}}(\Sigma, \{\neg Lp\}, \emptyset)$ returns $\langle \{\neg Lp, Lq\}, \emptyset \rangle$ and

$$\mathbf{test}(\Sigma, \{\neg Lp, Lq\}, \varphi)$$

is executed. If this test returns true, $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ returns true. Otherwise $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \{Lp\}, \varphi)$ is executed. Because $\mathbf{pos}_{\mathrm{LE}}(\Sigma, \emptyset, \{Lp\})$ returns $\langle \emptyset, \{Lp\} \rangle$, then $\mathbf{der}_{\mathrm{LE}}(\Sigma, \{\neg Lq\}, \{Lp\}, \varphi)$ is executed. Since now $\mathbf{pos}_{\mathrm{LE}}(\Sigma, \{\neg Lq\}, \{Lp\})$ returns $\langle \{\neg Lq, Lp\}, \{Lp\} \rangle$,

$$\mathbf{test}(\Sigma, \{\neg Lq, Lp\}, \varphi)$$

is executed. If this returns true, then $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ returns true. Otherwise $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \{Lp, Lq\}, \varphi)$ is executed and as $\mathbf{pos}_{\mathrm{LE}}(\Sigma, \emptyset, \{Lp, Lq\})$ returns $\langle \emptyset, \{Lp, Lq\} \rangle$, $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ returns false. ■

We aim to show that the function $\mathbf{der}_{\mathrm{LE}}$ is sound and complete. For the proof it is sufficient that the function $\mathbf{pos}_{\mathrm{LE}}$ that is used in $\mathbf{der}_{\mathrm{LE}}$ satisfies the four conditions in Proposition 5.1. The next lemma shows that the function $\mathbf{der}_{\mathrm{LE}}$ is *sound*, i.e., if it returns true, there is a full set for which **test** succeeds.

Lemma 5.4 *Let Σ be a finite set of sentences. If B and F are sets of sentences such that* q-atom$(B) \cup F \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$*, there is a fully L-hierarchic enumeration generating B from Σ and* $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ *returns true, then there exists an L-hierarchic Σ-acceptable enumeration ε_n of the set $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ such that $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ and* $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \varphi)$ *returns true.*

Proof. We prove the lemma by induction on the cardinality of the set

$$\mathrm{SL}(\Sigma, B, F) = \mathrm{Sf}^{\mathrm{L}}(\Sigma) - (\text{q-atom}(B) \cup F).$$

Let q-atom$(B) \cup F \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$, let there exist a fully L-hierarchic enumeration generating B from Σ and let $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ return true.

The function $\mathbf{der}_{\mathrm{LE}}$ starts by introspection. Using Proposition 3.7 and Proposition 5.1 (1, 2) the following basic result can be established which shows that if $\mathrm{Sf}^{\mathrm{L}}(\Sigma) \subseteq$ q-atom$(B) \cup F$ after introspection, an acceptable enumeration exists.

(B) *If $\mathrm{SL}(\Sigma, B', F') = \emptyset$ holds where $\langle B', F' \rangle = \mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$, then there exists an L-hierarchic Σ-acceptable enumeration ε_n of the set $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ such that $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ and* $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \varphi)$ *returns true.*

This can be shown in the following way. Let $\mathrm{SL}(\Sigma, B', F') = \emptyset$. Then $\mathrm{Sf}^{\mathrm{L}}(\Sigma) \subseteq$ q-atom$(B') \cup F'$. As $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ returns true,

$F' \subseteq B'$ and for all $\neg L\chi \in B'$, $\Sigma \cup B' \not\models_{\mathrm{ae}} \chi$ and, furthermore, $\mathbf{test}(\Sigma, B', \varphi)$ returns true. As B is generated by a fully L-hierarchic enumeration, by Proposition 5.1 (2) there is a fully L-hierarchic enumeration ε_n generating B', i.e., $B' = \mathbf{B}_n^{\varepsilon_n}(\Sigma)$.

We show first that ε_n is a Σ-acceptable enumeration of $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$. By Proposition 5.1 (1) q-atom$(B') \cup F' \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$. As $F' \subseteq B'$, q-atom$(B') \cup F' =$ q-atom(B'). Thus q-atom$(B') = \mathrm{Sf}^{\mathrm{L}}(\Sigma)$. This implies that ε_n is an enumeration of $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$. As for all $\neg L\chi \in \mathbf{B}_n^{\varepsilon_n}(\Sigma)$, $\Sigma \cup \mathbf{B}_n^{\varepsilon_n}(\Sigma) \not\models_{\mathrm{ae}} \chi$, by Proposition 3.7 ε_n is Σ-acceptable.

It remains to be shown that $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ and $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \varphi)$ returns true. If $\varphi \in B \cup F$, then by Proposition 5.1 (1) $\varphi \in B' \cup F' = B' = \mathbf{B}_n^{\varepsilon_n}(\Sigma)$. Thus $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$. Finally, $B' = \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ implies that $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \varphi)$ returns true. So (**B**) holds.

Let $\mathrm{SL}(\Sigma, B, F) = \emptyset$ hold. Then also $\mathrm{SL}(\Sigma, B', F') = \emptyset$ holds where $\langle B', F' \rangle = \mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$ as $B \subseteq B'$ and $F \subseteq F'$ by Proposition 5.1 (1). By (**B**) there exists an acceptable enumeration for which **test** succeeds.

Assume that for all B, F such that the cardinality of $\mathrm{SL}(\Sigma, B, F)$ is less than i if q-atom$(B) \cup F \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$, there exists a fully L-hierarchic enumeration generating B from Σ and $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ returns true, then there exists an L-hierarchic Σ-acceptable enumeration ε_n of $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ such that $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ and $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \varphi)$ returns true.

Let the cardinality of $\mathrm{SL}(\Sigma, B, F)$ be i and $\langle B', F' \rangle = \mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$. If $\mathrm{SL}(\Sigma, B', F') = \emptyset$, then (**B**) implies that there exists an acceptable enumeration for which **test** succeeds.

Assume then that $\mathrm{SL}(\Sigma, B', F') \neq \emptyset$ holds. Because $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ returns true, there is $L\chi \in \mathrm{SL}(\Sigma, B', F')$ such that $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq$ q-atom(B') and either $\mathbf{der}_{\mathrm{LE}}(\Sigma, B' \cup \{\neg L\chi\}, F', \varphi)$ returns true or $\mathbf{der}_{\mathrm{LE}}(\Sigma, B', F' \cup \{L\chi\}, \varphi)$ returns true.

Assume that $\mathbf{der}_{\mathrm{LE}}(\Sigma, B', F' \cup \{L\chi\}, \varphi)$ returns true. Note that the cardinality of $\mathrm{SL}(\Sigma, B', F' \cup \{L\chi\})$ is less than i and q-atom$(B') \cup F' \cup \{L\chi\} \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$. As B is generated by a fully L-hierarchic enumeration, by Proposition 5.1 (2) there is a fully L-hierarchic enumeration $\varepsilon_{m'}$ generating B'. Hence by the inductive hypothesis there exists an L-hierarchic Σ-acceptable enumeration ε_n of $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ such that $B' \cup F' \cup \{L\chi\} \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ and $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \varphi)$ returns true. By Proposition 5.1 (1) $B \subseteq B'$ and $F \subseteq F'$ which implies $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$.

On the other hand, consider the case where $\mathbf{der}_{\mathrm{LE}}(\Sigma, B' \cup \{\neg L\chi\}, F', \varphi)$ returns true. The cardinality of $\mathrm{SL}(\Sigma, B' \cup \{\neg L\chi\}, F')$ less than i and q-atom$(B' \cup \{\neg L\chi\}) \cup F' \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$. Let $\varepsilon_{m'+1}$ be $\varepsilon_{m'}$ extended by χ. Hence $B' = \mathbf{B}_{m'}^{\varepsilon_{m'}}(\Sigma) = \mathbf{B}_{m'}^{\varepsilon_{m'+1}}(\Sigma)$. As $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq$ q-atom$(B') =$ q-atom$(\mathbf{B}_{m'}^{\varepsilon_{m'}}(\Sigma))$ and as $\varepsilon_{m'}$ is fully L-hierarchic, $\varepsilon_{m'+1}$ is fully L-hierarchic. As $L\chi \notin$

q-atom(B'), by Proposition 5.1 (4) $\Sigma \cup \mathbf{B}_{m'}^{\varepsilon_{m'+1}}(\Sigma) \not\models_{\mathrm{ae}} \chi$. Hence $B' \cup \{\neg L\chi\} = \mathbf{B}_{m'+1}^{\varepsilon_{m'+1}}(\Sigma)$. Thus $B' \cup \{\neg L\chi\}$ is generated by the fully L-hierarchic enumeration $\varepsilon_{m'+1}$. By the inductive hypothesis there exists an L-hierarchic Σ-acceptable enumeration ε_n of $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ such that $B' \cup \{\neg L\chi\} \cup F' \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ and $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \varphi)$ returns true. By Proposition 5.1 (1) $B \subseteq B'$ and $F \subseteq F'$ which implies $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$. □

The next lemma shows that the function $\mathbf{der}_{\mathrm{LE}}$ is *complete*, i.e., if there is a full set for which the function **test** succeeds, then $\mathbf{der}_{\mathrm{LE}}$ returns true.

Lemma 5.5 *Let Σ be a finite set of sentences. Let ε_n be an L-hierarchic Σ-acceptable enumeration of $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ such that $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \varphi)$ returns true. Then for every set B of sentences and F of $L\chi$ sentences such that $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$, $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ returns true.*

Proof. The proof is by induction on the cardinality of $\mathrm{SL}(\Sigma, B, F) = \mathrm{Sf}^{\mathrm{L}}(\Sigma) - (\text{q-atom}(B) \cup F)$. Let there exist an L-hierarchic Σ-acceptable enumeration ε_n of $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ such that $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \varphi)$ returns true. Let B be a set of sentences and let F be a set of $L\chi$ sentences such that $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$.

Using Proposition 5.1 (3, 4) and Proposition 3.7 we can establish a result showing that if $\mathrm{SL}(\Sigma, B, F) = \emptyset$ after introspection, the function $\mathbf{der}_{\mathrm{LE}}$ returns true.

(**B**) *Let $\langle B', F'\rangle = \mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$ hold. If $\mathrm{SL}(\Sigma, B', F') = \emptyset$, then $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ returns true.*

This can be shown in the following way. By Proposition 5.1 (3) $B' \cup F' \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ holds. Let $\mathrm{SL}(\Sigma, B', F') = \emptyset$. Then it is straightforward to establish that

$$B' \cup F' = \mathbf{B}_n^{\varepsilon_n}(\Sigma) \tag{3}$$

holds. Using this we show that

$$F' \subseteq B' \tag{4}$$

holds. Assume that there is some $L\chi \in F' - B'$. By (3) χ is some ψ_i in ε_n. Let ψ_i be a sentence with the smallest index (in ε_n) such that $L\psi_i \in F'$ but $L\psi_i \notin B'$. By (3) $L\psi_i \in \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ and thus $\neg L\psi_i \notin \mathbf{B}_n^{\varepsilon_n}(\Sigma)$. Then by (3) $\neg L\psi_i \notin B'$. Thus $L\psi_i \in \mathrm{Sf}^{\mathrm{L}}(\Sigma) - \text{q-atom}(B')$. As $L\psi_i \in \mathbf{B}_n^{\varepsilon_n}(\Sigma)$, $\Sigma \cup \mathbf{B}_{i-1}^{\varepsilon_n}(\Sigma) \models_{\mathrm{ae}} \psi_i$ holds. We show next that $\mathbf{B}_{i-1}^{\varepsilon_n}(\Sigma) \subseteq B'$ holds. If $L\psi_j \in \mathbf{B}_{i-1}^{\varepsilon_n}(\Sigma)$, then $j < i$ and by (3) $L\psi_j \in B' \cup F'$. As $j < i$, $L\psi_j \in B'$. If $\neg L\psi_j \in \mathbf{B}_{i-1}^{\varepsilon_n}(\Sigma)$, then by (3) $\neg L\psi_j \in B' \cup F'$ and thus $\neg L\psi_j \in B'$. So $\mathbf{B}_{i-1}^{\varepsilon_n}(\Sigma) \subseteq B'$ and thus $\Sigma \cup B' \models_{\mathrm{ae}} \psi_i$ holds by the monotonicity of $\models_{\mathrm{ae}}$. But as ε_n is L-hierarchic, $\mathrm{Sf}^{\mathrm{L}}(\psi_i) \subseteq \text{q-atom}(\mathbf{B}_{i-1}^{\varepsilon_n}(\Sigma)) \subseteq \text{q-atom}(B')$.

Because $L\psi_i \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)$, Proposition 5.1 (4) implies $L\psi_i \in$ q-atom(B'), a contradiction. Thus (4) holds.

By (3) and (4) $B' = \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ holds. By Proposition 3.7 as ε_n is Σ-acceptable, if $\neg L\chi \in B'$, then $\Sigma \cup B' \cup F' \not\models_{\mathrm{ae}} \chi$. Hence false is not returned and by (4) $\mathbf{test}(\Sigma, B', \varphi)$ is executed. It returns true as $B' = \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ holds. Thus $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ returns true and (**B**) holds.

Let $\mathrm{SL}(\Sigma, B, F) = \emptyset$. Then by Proposition 5.1 (1) $\mathrm{SL}(\Sigma, B', F') = \emptyset$ where $\langle B', F' \rangle = \mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$. Then by (**B**) $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ returns true.

Assume that for B and F such that the cardinality of $\mathrm{SL}(\Sigma, B, F)$ is less than i, if $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$, then $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ returns true.

Let the cardinality of $\mathrm{SL}(\Sigma, B, F)$ be i and let $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$. Let $\langle B', F' \rangle = \mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$. If $\mathrm{SL}(\Sigma, B', F') = \emptyset$, then $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ returns true by (**B**) .

Assume that $\mathrm{SL}(\Sigma, B', F') \neq \emptyset$. First we show that there exists a sentence $L\chi \in \mathrm{SL}(\Sigma, B', F')$ such that $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq$ q-atom(B'). Consider a sentence ψ_i which is the sentence with the smallest index (in ε_n) such that $L\psi_i \in \mathrm{SL}(\Sigma, B', F')$. We show that $\mathrm{Sf}^{\mathrm{L}}(\psi_i) \subseteq$ q-atom(B'). Let $L\psi_j \in \mathrm{Sf}^{\mathrm{L}}(\psi_i)$. As ε_n is L-hierarchic, $j < i$. Therefore $L\psi_j \in$ q-atom$(B') \cup F'$ by the minimality requirement for i. It sufficient to show that $L\psi_j \in F'$ implies $L\psi_j \in$ q-atom(B') because this implies $L\psi_j \in$ q-atom(B') and thus $\mathrm{Sf}^{\mathrm{L}}(\psi_i) \subseteq$ q-atom(B'). Let $L\psi_j \in F'$. By Proposition 5.1 (3) $L\psi_j \in \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ and $\Sigma \cup \mathbf{B}_{j-1}^{\varepsilon_n}(\Sigma) \models_{\mathrm{ae}} \psi_j$ holds. Using Proposition 5.1 (3, 4) it can proved that

$$\mathbf{B}_{j-1}^{\varepsilon_n}(\Sigma) \subseteq B' \tag{5}$$

holds by showing by induction on l that for all $l = 0, 1, \ldots, j-i$, $\mathbf{B}_l^{\varepsilon_n}(\Sigma) \subseteq B'$ holds. For $l = 0$ it holds trivially. Let $\mathbf{B}_l^{\varepsilon_n}(\Sigma) = \mathbf{B}_{l-1}^{\varepsilon_n}(\Sigma) \cup \{L\psi_l\}$. Then $\Sigma \cup \mathbf{B}_{l-1}^{\varepsilon_n}(\Sigma) \models_{\mathrm{ae}} \psi_l$ holds. By the inductive hypothesis on l $\Sigma \cup B' \models_{\mathrm{ae}} \psi_l$ holds. But as ε_n is L-hierarchic, $\mathrm{Sf}^{\mathrm{L}}(\psi_l) \subseteq$ q-atom$(\mathbf{B}_{l-1}^{\varepsilon_n}(\Sigma)) \subseteq$ q-atom(B'). So Proposition 5.1 (4) implies $L\psi_l \in$ q-atom(B'). But $\neg L\psi_l \notin B'$ as $L\psi_l \in \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ and $B' \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ by Proposition 5.1 (3). Thus $L\psi_l \in B'$.

Let $\mathbf{B}_l^{\varepsilon_n}(\Sigma) = \mathbf{B}_{l-1}^{\varepsilon_n}(\Sigma) \cup \{\neg L\psi_l\}$. Then $L\psi_l \notin B' \cup F' \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$. As $l < j < i$, by assumption that i is the smallest index (in ε_n) such that $L\psi_i \in \mathrm{SL}(\Sigma, B', F')$, $L\psi_l \in$ q-atom$(B') \cup F'$. Then $L\psi_l \in$ q-atom(B') and $\neg L\psi_l \in B'$. Thus (5) holds.

Since the enumeration ε_n is L-hierarchic, then by (5) $\mathrm{Sf}^{\mathrm{L}}(\psi_j) \subseteq$ q-atom$(\mathbf{B}_{j-1}^{\varepsilon_n}(\Sigma)) \subseteq$ q-atom(B') and $\Sigma \cup B' \models_{\mathrm{ae}} \psi_j$ holds. By Proposition 5.1 (4) $L\psi_j \in$ q-atom(B'). Hence $\mathrm{Sf}^{\mathrm{L}}(\psi_i) \subseteq$ q-atom(B'). So there is $L\chi (= L\psi_i) \in \mathrm{SL}(\Sigma, B', F')$ such that $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq$ q-atom(B').

If $\neg L\chi \in \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ holds, then $B' \cup \{\neg L\chi\} \cup F' \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ holds as well.

As the cardinality of $\mathrm{SL}(\Sigma, B' \cup \{\neg L\chi\}, F')$ is less than i, by the inductive hypothesis $\mathbf{der}_{\mathrm{LE}}(\Sigma, B' \cup \{\neg L\chi\}, F', \varphi)$ returns true which implies that $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ returns true. If $\neg L\chi \notin \mathbf{B}_n^{\varepsilon_n}(\Sigma)$, then $L\chi \in \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ and $B' \cup F' \cup \{L\chi\} \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$. As the cardinality of $\mathrm{SL}(\Sigma, B', F' \cup \{L\chi\})$ is less than i, by the inductive hypothesis $\mathbf{der}_{\mathrm{LE}}(\Sigma, B', F' \cup \{L\chi\}, \varphi)$ returns true and hence $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ returns true. $\square$

As a direct consequence of Lemmata 5.4 and 5.5 we have a soundness and completeness theorem for $\mathbf{der}_{\mathrm{LE}}$.

Theorem 5.6 *Let Σ be a finite set of sentences. Let B and F be sets for which* q-atom$(B) \cup F \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$ *and let there be a fully L-hierarchic enumeration generating B from Σ. Then $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \varphi)$ returns true if and only if there exists an L-hierarchic Σ-acceptable enumeration ε_n of the set $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ such that $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ and $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \varphi)$ returns true.*

The theorem implies that if the underlying consequence relation $\models_{\mathrm{ae}}$ is decidable, decision procedures for $\mathrel{|\!\sim}_{\mathrm{c(LE)}}$, $\mathrel{|\!\sim}_{\mathrm{b(LE)}}$, and the L-hierarchic expansion existence problem can be obtained by employing an appropriate function **test** and taking the empty set as the initial common part of the full sets. These **test** functions are straightforward to devise because membership in an expansion is captured by the consequence relation $\models_{\mathrm{L}}$, i.e., given a Σ-full set Λ, φ is in the corresponding unique expansion if and only if $\Sigma \cup \Lambda \models_{\mathrm{L}} \varphi$ (Theorem 2.9).

Theorem 5.7 *Let $\Sigma \subseteq \mathcal{L}_{\mathrm{ae}}$ be a finite set of sentences and $\varphi \in \mathcal{L}_{\mathrm{ae}}$ a sentence.*

1. *Let $\mathbf{test}(\Sigma, \Lambda, \varphi)$ return true for any Σ, Λ, φ. Then $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ returns true if and only if Σ has an L-hierarchic expansion.*
2. *Let $\mathbf{test}(\Sigma, \Lambda, \varphi)$ return true if and only if $\Sigma \cup \Lambda \models_{\mathrm{L}} \varphi$. Then $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ returns true if and only if $\Sigma \mathrel{|\!\sim}_{\mathrm{b(LE)}} \varphi$.*
3. *Let $\mathbf{test}(\Sigma, \Lambda, \varphi)$ return true if and only if $\Sigma \cup \Lambda \not\models_{\mathrm{L}} \varphi$. Then $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ returns true if and only if $\Sigma \mathrel{|\!\not\sim}_{\mathrm{c(LE)}} \varphi$.*

Proof. Clearly, q-atom$(\emptyset) \cup \emptyset \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$ and $\emptyset$ is generated from Σ by the empty enumeration which is fully L-hierarchic, of course. Then by Theorem 5.6 $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ returns true if and only if there exists an L-hierarchic Σ-acceptable enumeration ε_n of $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ such that $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \varphi)$ returns true.

(1) Assume that $\mathbf{test}(\Sigma, \Lambda, \varphi)$ returns true for every input. ($\Rightarrow$) If $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ returns true, then there is an L-hierarchic Σ-acceptable enumeration ε_n of $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$. By Theorem 3.8 there is an L-hierarchic expansion of Σ. ($\Leftarrow$) Let there exist an L-hierarchic expansion of Σ. Then by Theorem 3.9 there exists an L-hierarchic Σ-acceptable enumeration ε_n of $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$. So $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ returns true.

(2) Let $\mathbf{test}(\Sigma, \Lambda, \varphi)$ return true if and only if $\Sigma \cup \Lambda \models_{\mathrm{L}} \varphi$. ($\Rightarrow$) Let $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ return true. Then there exists an L-hierarchic Σ-acceptable enumeration ε_n of $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ and $\Sigma \cup \mathbf{B}_n^{\varepsilon_n}(\Sigma) \models_{\mathrm{L}} \varphi$. By Theorem 3.8 there is an L-hierarchic expansion Δ of Σ such that

$$(6) \qquad \Delta \;=\; \mathrm{SE}_{\Sigma}(\mathbf{B}_n^{\varepsilon_n}(\Sigma)) = \{\varphi \in \mathcal{L}_{\mathrm{ae}} \mid \Sigma \cup \mathbf{B}_n^{\varepsilon_n}(\Sigma) \models_{\mathrm{L}} \varphi\}.$$

Thus $\varphi \in \Delta$ and $\Sigma \mid\!\sim_{\mathrm{b(LE)}} \varphi$. ($\Leftarrow$) Let $\Sigma \mid\!\sim_{\mathrm{b(LE)}} \varphi$ hold. Then there exists an L-hierarchic expansion Δ of Σ containing φ. By Theorem 3.9 there exists an L-hierarchic Σ-acceptable enumeration ε_n of $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ such that Δ is given by (6). Thus $\Sigma \cup \mathbf{B}_n^{\varepsilon_n}(\Sigma) \models_{\mathrm{L}} \varphi$ and $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \varphi)$ returns true. Hence $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ returns true.

(3) Let $\mathbf{test}(\Sigma, \Lambda, \varphi)$ return true if and only if $\Sigma \cup \Lambda \not\models_{\mathrm{L}} \varphi$. As for (2) by Theorems 3.8 and 3.9 it can be shown that $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \varphi)$ returns true if and only if there is an L-hierarchic expansion of Σ not containing φ. $\square$

The completeness results with respect to the second level of the polynomial time hierarchy (Gottlob 1992, Niemelä 1992a, Stillman 1992) indicate that there are two independent sources of complexity in nonmonotonic reasoning:

- required monotonic reasoning and
- conflict resolution.

In the decision method the two sources are clearly identifiable. In the autoepistemic framework the required monotonic reasoning is given by the underlying consequence relation $\models_{\mathrm{ae}}$. The conflict resolution task emerges when deciding which $\neg L\chi$ is included in the full set from the available $L\chi$ subformulae of the premises. Such choices can be in conflict with each other: if we include one such $\neg L\chi$ in the full set to be constructed, that might make some χ' a consequence of the premises Σ and the full set and force us to add $L\chi'$, instead of $\neg L\chi'$, to the full set to be constructed. In the decision method the conflict resolution task is solved by a backtracking scheme. If $\neg L\chi'$ has been included in the full set to be constructed but χ' becomes a consequence of the premises Σ and the full set, the method backtracks to the last choice point where some $\neg L\chi$ was added to the full set and retracts this choice. As for every $L\chi$ subformula of Σ either $\neg L\chi$ or $L\chi$ belongs to the full set, the other possibility is that $L\chi$ belongs to the full set. The other alternative is explored by putting $L\chi$ as a "frozen" belief to be included in the full set later. The same backtracking mechanism can be exploited when searching for an expansion containing a given sentence (brave derivability) or an expansion not containing a given sentence (cautious derivability). When a full set has been found, the function **test** is invoked to determine whether the expansion induced by the full set satisfies the required condition, e.g., contains a given sentence. If this is not the case, i.e., **test** returns false, the method backtracks to the last choice point and tries the other alternative. In the decision method a conflict test is performed immediately after positive introspection has been completed. The aim is to cut the backtracking as early as possible when a conflict emerges.

The decision procedures for $\mid\!\sim_{\mathrm{c(E)}}$, $\mid\!\sim_{\mathrm{b(E)}}$, and the enumeration-based expansion existence problem can be obtained easily by modifying the functions $\mathbf{pos}_{\mathrm{LE}}$ and $\mathbf{der}_{\mathrm{LE}}$ slightly. The only change is that there is no need to require that $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq \text{q-atom}(B)$ for the new $(\neg)L\chi$ formulae to be added to B. The decision procedures for $\mid\!\sim_{\mathrm{c}}$, $\mid\!\sim_{\mathrm{b}}$, and the Moore-style expansion existence problem can be devised by further modifying the function $\mathbf{pos}_{\mathrm{LE}}$. The change involves allowing the frozen sentences to be used in positive introspection because the groundedness requirements for Moore-style expansions are not as strict as for enumeration-based expansions.

The decision method described above is completely general. It does not rely on any additional assumptions about the underlying logic which determines the consequence relation $\models_{\mathrm{ae}}$. For example, no normal form transformation of the premises is required. The method needs a theorem prover for $\models_{\mathrm{ae}}$ as a subroutine but no extra requirements are posed on the theorem prover.

Properties of the underlying logic as well as possible regularities of the sets of premises can be utilized to optimize the computationally costly subtasks of the decision method: required monotonic reasoning and conflict resolution. The theorem prover for $\models_{\mathrm{ae}}$ can exploit any special properties of the underlying logic to achieve highest possible performance. For example, the order of solving the subgoals during theorem proving is not restricted. If the set of premises is of some special form for which a special purpose theorem prover is available, this theorem prover can be applied immediately. It should be noted that despite formulae containing L operators the theorem prover for $\models_{\mathrm{ae}}$ is very similar to that for the underlying logic CL. This is because the $L\chi$ sentences behave like atomic sentences in the propositional calculus as far as the consequence relation $\models_{\mathrm{ae}}$ goes. For example, if the underlying logic is a first-order predicate calculus, we can use as a theorem prover for $\models_{\mathrm{ae}}$, a theorem prover for the predicate calculus which treats $L\chi$ sentence as new nilary predicates. The simple logical character of the $L\chi$ sentences can be exploited when deciding $\models_{\mathrm{ae}}$ and the set of premises can be considerably reduced. For example, it can be shown that in the first-order case if Σ is a set of sentences of the form $(Lb_1 \wedge \cdots \wedge Lb_n) \rightarrow c$ where $c \in \mathcal{L}$ and Λ is a consistent set of $L\psi$ formulae and their negations, then for all $\varphi \in \mathcal{L}$,

$$\Sigma \cup \Lambda \models_{\mathrm{ae}} \varphi \text{ if and only if } \mathrm{Hds}(\Sigma, \Lambda) \models_{\mathrm{cl}} \varphi$$

where $\mathrm{Hds}(\Sigma, \Lambda) = \{c \mid (Lb_1 \wedge \cdots \wedge Lb_n) \rightarrow c \in \Sigma \text{ and } Lb_i \in \Lambda \text{ for all } i = 1, \ldots, n\}$. This result turns out to be applicable when using the decision method as a decision procedure for default logic by employing the translation (7) discussed in Section 6. The theorem prover for $\models_{\mathrm{ae}}$ is used with a regular pattern. An interesting topic of further research is to study possibilities of taking advantage of the regularity and exploiting previous computations when deciding $\models_{\mathrm{ae}}$.

Optimization of conflict resolution is also a challenging area of further research. An interesting approach is to investigate the notion of *safe disbeliefs*. The idea is to identify disbeliefs $\neg L\chi$ such that χ does not follow from Σ and any possible extension of $B \cup \{\neg L\chi\} \cup F$. This reduces backtracking as then B can be extended safely by $\neg L\chi$ since for $\neg L\chi$ a conflict is not possible.

The two independent sources of complexity nicely explain the high computational complexity of some restricted subclasses of default logic and autoepistemic logic (Kautz and Selman 1991, Stillman 1990, Stillman 1992, Marek and Truszczyński 1991a, Marek and Truszczyński 1991b). To obtain a subclass of nonmonotonic reasoning with favorable computational properties it is not enough to restrict the form of formulae to make the classical reasoning task easy but also the conflict resolution task has to be addressed. We have studied the notion of stratification originally introduced in the context of logic programs and shown that when considering only stratified sets of autoepistemic formulae (Marek and Truszczyński 1991a) the conflict resolution task can be solved efficiently and a general iterative decision method can be devised (Niemelä and Rintanen 1992). The iterative decision method can be used also for stratified default logic, logic programs, truth maintenance systems, and McDermott and Doyle-style nonmonotonic modal logics. For example, a linear time decision procedure for stratified propositional logic programs and truth maintenance systems is obtained (Niemelä and Rintanen 1992).

6 Applications

We show how the decision method for autoepistemic reasoning can be used as a decision method for other forms of nonmonotonic reasoning by providing simple local mappings from several other forms of nonmonotonic reasoning into autoepistemic reasoning.

McDermott and Doyle-Style Nonmonotonic Modal Logics

Autoepistemic logic is closely related to McDermott and Doyle-style nonmonotonic modal logics where for a monotonic modal logic S, the possible sets of nonmonotonic conclusions are defined in terms of S-expansions. The following results imply that the methods and results obtained for autoepistemic logic are directly applicable to nonmonotonic modal logics for a wide range of modal logics and vice versa.

Theorem 6.1 (Marek et al. 1991) *Let* $\Sigma \subseteq \mathcal{L}_{\mathrm{ae}}$.

(1) For each logic S such that $\mathbf{5} \subseteq S \subseteq \mathbf{KD45}$, Δ is a consistent stable expansion of Σ if and only if Δ is a consistent S-expansion of Σ.

(2) Let $\Sigma \subseteq \mathcal{L}_{\mathrm{ae}}$ be stratified. Then for each logic S such that $\mathbf{N} \subseteq S$ and $S \subseteq \mathbf{KD45}$ or $S \subseteq \mathbf{S4}$, the unique stable expansion of Σ is the unique S-expansion of Σ.

Theorem 6.2 (Niemelä 1991) *Let Σ be a finite set of propositional autoepistemic formulae. Then Δ is an $\mathbf{N}$-expansion (iterative expansion) of Σ if and only if Δ is an enumeration-based expansion of Σ.*

Default Logic

Default logic introduced by Reiter (1980) is one of the leading formalizations of nonmonotonic reasoning. Truszczyński (1991a) has proposed an appealing translation (7) of default rules to autoepistemic sentences and shown that the consistent expansions of the translated default theory in many McDermott and Doyle-style nonmonotonic modal logics correspond to the extensions of the default theory. When using L-hierarchic expansions there is no need to consider only consistent expansions.

$$\text{(7)} \quad \mathrm{AE_{DL}}\left(\frac{a : b_1, \ldots, b_n}{c}\right) = (La \wedge L\neg L\neg b_1 \wedge \cdots \wedge L\neg L\neg b_n) \rightarrow c$$

$$\text{(8)} \quad \mathrm{AE_{DL}}(D, W) = W \cup \{\mathrm{AE_{DL}}(d) \mid d \in D\}$$

Theorem 6.3 (Niemelä 1992b) *Let (D, W) be a default theory where D is finite. Then $E \subseteq \mathcal{L}$ is an extension of (D, W) if and only if $E = \Delta \cap \mathcal{L}$ for an L-hierarchic expansion Δ of $\mathrm{AE_{DL}}(D, W)$.*

Circumscription

In autoepistemic logic it is possible to capture circumscriptive inference, i.e., the formulae true in every minimal model of a set of premises. We consider the propositional case and a rather general notion of minimal models. Let P and F be sets of atoms. A model I (a set of atoms) of Σ is $\langle P, F\rangle$-minimal if there is no model I' of Σ such that $I' \cap F = I \cap F$ but $I' \cap P \subset I \cap P$. This notion of minimal models corresponds to parallel circumscription (Lifschitz 1985) where the atoms in P are minimized and the atoms not in F or P are allowed to vary. For a set of formulae P, we define $\mathrm{MIN}(P) = \{\neg L\varphi \rightarrow \overline{\varphi} \mid \varphi \in P\}$ and $\mathrm{FIX}(P) = \mathrm{MIN}(P) \cup \mathrm{MIN}(\neg P)$ where $\overline{\varphi}$ denotes the complement of a formula φ.

Theorem 6.4 (Niemelä 1992b) *Let P and F be finite sets of atomic sentences and $\Sigma \subseteq \mathcal{L}$. For every sentence $\varphi \in \mathcal{L}$ holds that φ is true in every $\langle P, F\rangle$-minimal model of Σ if and only if φ belongs to every L-hierarchic expansion of $\Sigma \cup \mathrm{MIN}(P) \cup \mathrm{FIX}(F)$.*

Konolige (1989) shows how to achieve minimalization in the first-order case with non-finite domains using a variant of autoepistemic logic which allows quantification into the scope of L operators. He uses the same scheme as above to minimize and fix predicates.

Logic Programs and Deductive Databases

Let us consider *extended logic programs* which are sets of rules of the form

$$l_0 \leftarrow l_1, \ldots, l_m, \mathrm{not}\, l_{m+1}, \ldots, \mathrm{not}\, l_n \tag{9}$$

where not is the negation-as-failure operator, $n \geq m \geq 0$ and each l_i is a literal, i.e., an atomic formula or a negation of an atomic formula. Extended logic programs are a generalization of *general logic programs* where classical negation is allowed. Gelfond and Lifschitz (1990) have proposed a semantics for extended logic programs which is a generalization of the stable model semantics of general logic programs (Gelfond and Lifschitz 1988). The semantics is given in terms of answer sets which are sets of literals giving answers to queries.

We consider the propositional case where a rule with variables can be seen as a shorthand for the set of its ground instances. Answer sets are closely related to extensions in default logic and can be captured using L-hierarchic enumerations by translating rules in an extended logic program in the following way.

$$\begin{aligned} &\mathrm{LP}_{\mathrm{AE}}(l_0 \leftarrow l_1, \ldots, l_m, \mathrm{not}\, l_{m+1}, \ldots, \mathrm{not}\, l_n) = \\ &\qquad Ll_1 \wedge \cdots \wedge Ll_m \wedge L\neg Ll_{m+1} \wedge \cdots \wedge L\neg Ll_n \rightarrow l_0 \end{aligned} \tag{10}$$

Theorem 6.5 (Niemelä 1992b) *Let P be a finite extended propositional logic program. Then S is an answer set of P if and only if S is the set of literals in an L-hierarchic expansion of* $\mathrm{LP}_{\mathrm{AE}}(P)$.

As the answer sets of a logic program without classical negation coincide with the stable models of the program, the theorem above implies a corresponding result for stable models.

Corollary 6.6 *Let P be a finite general propositional logic program. Then S is a stable model of P if and only if S is the set of atoms in an L-hierarchic expansion of* $\mathrm{LP}_{\mathrm{AE}}(P)$.

Autoepistemic logic provides a framework where also answers to queries can be captured. Here the belief operator turns out to be indispensable. For example, given a stable model the answer to a query is determined by the truth value of the query in the model. In autoepistemic logic truth in a model can be captured in terms of the corresponding expansion by interpreting truth as belief.

Proposition 6.7 (Niemelä 1992b) *Let Δ be a consistent stable expansion and let I be the atoms in Δ. Then for every propositional sentence $\varphi \in \mathcal{L}$, φ is true in I if and only if $\mathrm{B}(\varphi) \in \Delta$ where $\mathrm{B}(\varphi)$ denotes the formula which is obtained by replacing every atomic subformulae χ of φ by $L\chi$.*

Truth Maintenance

Elkan (1990) gives an elegant logical characterization of nonmonotonic justification-based truth maintenance systems (TMS). A TMS is a program that manages a set of nodes and a set of justifications. The task of a TMS is to choose a subset of nodes (a grounded model) as currently believed. Elkan shows that the justifications can be seen as a propositional general logic program and the subset of nodes chosen by a TMS is a stable model of the logic program.

Corollary 6.8 *Let J be a finite set of justifications. Then M is a grounded model of J if and only if M is the set of atoms in an L-hierarchic expansion of* $\mathrm{LP}_{\mathrm{AE}}(J)$.

7 Approaches Based on Partial Semantics

Recent studies indicate that nonmonotonic reasoning as formalized in autoepistemic logic, default logic, circumscription, stable model semantics of logic programs, truth maintenance systems, and abduction is computationally quite complex (Gottlob 1992, Niemelä 1992a, Stillman 1992, Eiter and Gottlob 1993, Elkan 1990, Selman and Levesque 1990). Generally the relevant decision problems are in the propositional case complete problems with respect to the second level of the polynomial time hierarchy. This means that such nonmonotonic reasoning is strictly harder than classical propositional reasoning unless the polynomial time hierarchy collapses which is considered unlikely by the experts (Johnson 1990). As discussed in Section 5 nonmonotonic reasoning appears to have two orthogonal sources of complexity: classical reasoning and conflict resolution. Both of the sources can be addressed using an approach based on partial semantics.

It is possible to restrict the complexity of the required classical reasoning by moving from a two-valued underlying logic to a multi-valued logic (Levesque 1988). Here the gain in efficiency has to be compared with the loss in deductive power. Also the complexity of the conflict resolution task can be addressed by using partial semantics. One of the leading examples is the well-founded semantics (Van Gelder et al. 1991) of logic programs. Here the advantages over the corresponding (two-valued) stable model semantics (Gelfond and Lifschitz 1988) are that there is a unique well-founded (partial) model for each program and the model can be constructed iteratively (in the propositional case in polynomial time). On the other hand, some deductive power is lost. For example, the unique well-founded model of the program $\{p \leftarrow \mathrm{not}\, q; q \leftarrow \mathrm{not}\, p; r \leftarrow \mathrm{not}\, p; r \leftarrow \mathrm{not}\, q\}$ is empty while r belongs to the intersection of the two (two-valued) stable models of the program.

Autoepistemic logics based on partial semantics have also been proposed (Przymusinski 1991b, Bonatti 1992). These approaches are closely related to the well-founded semantics with similar advantages and disadvantages. Each set of premises has a unique set of conclusions, i.e., a unique expansion, and the expansion can be constructed iteratively. But some deductive power is lost as the expansion does not contain all the conclusions given by the intersection of the expansions in the two-valued case. An interesting result is that when using an autoepistemic framework it is not necessary to adopt partial semantics to reconstruct the well-founded semantics. A classical two-valued underlying logic can be used and the partiality appears on the level of the beliefs of the agent: besides formulae that are believed and disbelieved there are formulae whose belief status is undetermined. Examples of this kind of a reconstruction are Pryzymusinski's (1991a) and Kalinski's (1990, 1994) variants of autoepistemic logic as well as a similar variant of default logic (Przymusinska and Przymusinski 1992). It is possible to devise stronger forms of partial autoepistemic reasoning where some of the conclusions missed by the well-founded semantics are captured without loosing an iterative construction for a unique set of conclusions for each set of premises (Janhunen and Niemelä 1993). An interesting topic of further research is to study the trade-off between deductive power and complexity of reasoning which seems to underlie these kinds of results.

8 Summary

In this work we consider the problem of automating nonmonotonic reasoning. Our aim is to exploit the recently discovered close connections between various formalizations of nonmonotonic reasoning and develop a unified approach to automating nonmonotonic reasoning in the following way. We choose some formalization of nonmonotonic reasoning as a basic system and develop a decision method for it. For other forms of nonmonotonic reasoning, decision methods can then be developed by mapping a reasoning problem in a given formalization to a reasoning problem in the basic system. In this work we consider the possibility of using a recently proposed variant of autoepistemic logic based on L-hierarchic expansions as the basic system. This variant provides a tightly grounded form of autoepistemic logic with a wide coverage: major forms of nonmonotonic reasoning can be embedded into it using simple local translations.

We study the automation of autoepistemic reasoning in a general setting where autoepistemic logic is defined on top of a given classical logic. We present a simple finitary characterization of autoepistemic reasoning providing the basis for the computational treatment. We survey the relevant complexity and decidability results which offer a framework for developing and assessing implementations of autoepistemic reasoning.

The main contribution is a novel decision method for autoepistemic reasoning which is applicable in a general setting, i.e., for an autoepistemic logic defined on top of a given classical logic. The aim is to devise a conceptually clear method where the sources of complexity and targets for further optimization are clearly identifiable. There are two independent sources of complexity: classical reasoning and conflict resolution. The classical reasoning task is handled using a separate theorem prover. The decision method imposes no additional requirements leaving abundant room for optimizations of the theorem prover. The conflict resolution task is solved by employing a backtracking scheme. The method exploits efficiently the monotonicity of the underlying consequence relation to reduce choice points in backtracking. The method provides a direct decision procedure for tightly grounded autoepistemic reasoning based on L-hierarchic expansions and no additional groundedness tests are needed. Decision methods for more weakly grounded forms of autoepistemic reasoning based on Moore-style expansions and **N**-expansions can be obtained with minor adjustments. The method can be implemented to run in polynomial space in the propositional case. Prominent forms of nonmonotonic reasoning, such as default logic and circumscription, can be embedded in autoepistemic logic based on L-hierarchic expansions using simple local translations. Thus the method can be used as a unified approach to automating a large range of forms of nonmonotonic reasoning.

There are several interesting topics of further research. One interesting area is to study the consequences of considering only a certain subset of autoepistemic logic, e.g., that corresponding to default logic under the translation (7). As indicated in Section 5 some advantages concerning classical reasoning are readily available and other gains seem possible, as well. Another topic is the development of further optimization techniques both for classical reasoning and conflict resolution. Here partial semantics can play a role. For example for conflict resolution, the well-founded semantics (Van Gelder et al. 1991) and the related stationary extensions (Przymusinska and Przymusinski 1992) offer an interesting approach to implementing the notion of safe disbeliefs introduced in the end of Section 5. These two topics are addressed in subsequent papers (Niemelä 1994, Niemelä 1995a, Niemelä 1995b). Another area of further research where partial semantics offers interesting options is the development of approximations of nonmonotonic reasoning as defined in, e.g., autoepistemic logic and default logic.

References

Ben-Eliyahu, R., and R. Dechter. 1991. Default Logic, Propositional Logic and Constraints. In *Proceedings of the 9th National Conference on Artificial Intelligence*, 370–385. The MIT Press, July.

Bonatti, P.A. 1992. Autoepistemic Logics as a Unifying Framework for the Semantics of Logic Programs. In *Proceedings of the 1992 Joint International*

Conference and Symposium on Logic Programming. Washington, D.C., USA, November.

Chellas, B.F. 1980. *Modal Logic: an Introduction*. Cambridge: Cambridge University Press.

Eiter, T., and G. Gottlob. 1992a. The Complexity of Logic-Based Abduction. Technical Report CD-TR 92/35. Vienna, Austria: Christian Doppler Labor für Expertensysteme, Institut für Informationssysteme, Technische Universität Wien, July.

Eiter, T., and G. Gottlob. 1992b. Reasoning with Parsimonious and Moderately Grounded Expansions. *Fundamenta Informaticae* 17:31–53.

Eiter, T., and G. Gottlob. 1993. Propositional Circumscription and Extended Closed World Reasoning are Π_2^p-complete. *Theoretical Computer Science* 114:231–245.

Elkan, C. 1990. A Rational Reconstruction of Nonmonotonic Truth Maintenance Systems. *Artificial Intelligence* 43:219–234.

Gelfond, M., and V. Lifschitz. 1988. The Stable Model Semantics for Logic Programming. In *Proceedings of the 5th International Conference on Logic Programming*, 1070–1080. Seattle, USA, August. The MIT Press.

Gelfond, M., and V. Lifschitz. 1990. Logic Programs with Classical Negation. In *Proceedings of the 7th International Conference on Logic Programming*, 579–597. Jerusalem, Israel, June. The MIT Press.

Gelfond, M., and H. Przymusinska. 1990. Formalization of Inheritance Reasoning in Autoepistemic Logic. *Fundamenta Informaticae* 13:430–443.

Gottlob, G. 1992. Complexity Results for Nonmonotonic Logics. *Journal of Logic and Computation* 2(3):397–425.

Gottlob, G. 1993. The Power of Beliefs or Translating Default Logic into Standard Autoepistemic Logic. In *Proceedings of the 13th International Joint Conference on Artificial Intelligence*, 570–575. Chambéry, France, August. Morgan Kaufmann Publishers.

Hughes, G.E., and M.J. Cresswell. 1984. *A Companion to Modal Logic*. London: Methuen and Co.

Janhunen, T., and I. Niemelä. 1993. A Scheme for Weakened Negative Introspection in Autoepistemic Reasoning. In *Proceedings of the 3rd Kurt Gödel Colloquium on Computational Logic and Proof Theory*, 211–222. Brno, Czech Republic, August. Springer-Verlag.

Johnson, D.S. 1990. A Catalog of Complexity Classes. In *Handbook of Theoretical Computer Science*, ed. J. van Leeuwen. 67–161. Elsevier.

Junker, U., and K. Konolige. 1990a. Computing the Extensions of Autoepistemic and Default Logics with a Truth Maintenance System. In *Proceedings of the 8th National Conference on Artificial Intelligence*, 278–283. Boston, MA, USA, July. The MIT Press.

Junker, U., and K. Konolige. 1990b. Computing the Extensions of Autoepistemic and Default Logics with a Truth Maintenance System. TASSO-Report 16. St. Augustin, Germany: Gesellschaft für Mathematik und Datenverarbeitung, December.

Kakas, A.C., and P. Mancarella. 1990. Generalized Stable Models: a Semantics for Abduction. In *Proceedings of the 9th European Conference on Artificial Intelligence*, 385–391. Stockholm, Sweden, August. Pitman Publishing.

Kalinski, J. 1990. Autoepistemic Expansions with Incomplete Belief Introspection. In *Proceedings of the 1st International Workshop on Nonmonotonic and Inductive Logic*, 223–232. Karlsruhe, Germany, December. Springer-Verlag. Lecture Notes in Artificial Intelligence 543.

Kalinski, J. 1994. Weak Autoepistemic Reasoning and Well-Founded Semantics. In *Foundations of Knowledge Representation and Reasoning*, ed. G. Lakemeyer and B. Nebel. 179–199. Berlin: Springer-Verlag.

Kautz, H., and B. Selman. 1991. Hard Problems for Simple Default Logics. *Artificial Intelligence* 49:243–279.

Konolige, K. 1988. On the Relation Between Default and Autoepistemic Logic. *Artificial Intelligence* 35:343–382.

Konolige, K. 1989. On the Relation Between Autoepistemic Logic and Circumscription. In *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, 1213–1218. Detroit, Michigan, USA, August. Morgan Kaufmann Publishers.

Konolige, K. 1992. Ideal Introspective Belief. In *Proceedings of the 10th National Conference on Artificial Intelligence*, 635–641. San Jose, CA, USA, July. The MIT Press.

Levesque, H.J. 1988. Logic and Complexity of Reasoning. *Journal of Philosophical Logic* 17:355–389.

Lifschitz, V. 1985. Computing Circumscription. In *Proceedings of the 9th International Joint Conference on Artificial Intelligence*, 121–127. Los Angeles, California, USA, August. Morgan Kaufmann Publishers.

Lifschitz, V. 1991. Nonmonotonic Databases and Epistemic Queries. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, 381–386. Sydney, Australia, August. Morgan Kaufmann Publishers.

Lin, F., and Y. Shoham. 1990. Epistemic Semantics for Fixed-Point Non-Monotonic Logics. In *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge*, 111–120. Pacific Grove, USA, March. Morgan Kaufmann Publishers.

Marek, W. 1989. Stable Theories in Autoepistemic Logic. *Fundamenta Informaticae* 12:243–254.

Marek, W., G.F. Shvarts, and M. Truszczyński. 1991. Modal Nonmonotonic Logics: Ranges, Characterization, Computation. In *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, 395–404. Cambridge, MA, USA, April. Morgan Kaufmann Publishers.

Marek, W., and M. Truszczyński. 1988. Autoepistemic Logic. Technical Report 115-88. Lexington, Kentucky, USA: Department of Computer Science, University of Kentucky.

Marek, W., and M. Truszczyński. 1989. Relating Autoepistemic and Default Logics. In *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, 276–288. Toronto, Canada, May.

Marek, W., and M. Truszczyński. 1990. Modal Logic for Default Reasoning. *Annals of Mathematics and Artificial Intelligence* 1:275–302.

Marek, W., and M. Truszczyński. 1991a. Autoepistemic Logic. *Journal of the ACM* 38:588–619.

Marek, W., and M. Truszczyński. 1991b. Computing Intersection of Autoepistemic Expansions. In *Proceedings of the 1st International Workshop on Logic Programming and Non-monotonic Reasoning*, 37–50. Washington, D.C., USA, July. The MIT Press.

McDermott, D. 1982. Non-Monotonic Logic II. *Journal of the ACM* 29:33–57.

McDermott, D., and J. Doyle. 1980. Non-Monotonic Logic I. *Artificial Intelligence* 13:41–72.

Moore, R.C. 1985. Semantical Considerations on Nonmonotonic Logic. *Artificial Intelligence* 25:75–94.

Moore, R.C. 1988. Autoepistemic Logic. In *Non-Standard Logics for Automated Reasoning*, ed. P. Smets, E.H. Mamdani, and D. Dubois. 105–136. London: Academic Press.

Niemelä, I. 1988. Decision Procedure for Autoepistemic Logic. In *Proceedings of the 9th International Conference on Automated Deduction*, 675–684. Argonne, USA, May. Springer-Verlag.

Niemelä, I. 1989. Decision Problem in Autoepistemic Logic. In *Proceedings of the Workshop on Alternatives of Logic Programming*, 174–188. Kuopio, Finland, May. Kuopion yliopiston painatuskeskus.

Niemelä, I. 1990a. A New Characterization of Autoepistemic Reasoning. In *Proceedings of the Finnish Artificial Intelligence Symposium*, 64–73. Oulu, Finland, June. Finnish Artificial Intelligence Society.

Niemelä, I. 1990b. Towards Automatic Autoepistemic Reasoning. In *Proceedings of the European Workshop on Logics in Artificial Intelligence—JELIA'90*, 428–443. Amsterdam, The Netherlands, September. Springer-Verlag.

Niemelä, I. 1991. Constructive Tightly Grounded Autoepistemic Reasoning. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, 399–404. Sydney, Australia, August. Morgan Kaufmann Publishers.

Niemelä, I. 1992a. On the Decidability and Complexity of Autoepistemic Reasoning. *Fundamenta Informaticae* 17(1,2):117–155.

Niemelä, I. 1992b. A Unifying Framework for Nonmonotonic Reasoning. In *Proceedings of the 10th European Conference on Artificial Intelligence*, 334–338. Vienna, Austria, August. John Wiley.

Niemelä, I. 1994. A Decision Method for Nonmonotonic Reasoning Based on Autoepistemic Reasoning. In *Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning*, 473–484. Bonn, Germany, May. Morgan Kaufmann Publishers.

Niemelä, I. 1995a. A Decision Method for Nonmonotonic Reasoning Based on Autoepistemic Reasoning. *Journal of Automated Reasoning* 14:3–42.

Niemelä, I. 1995b. Towards Efficient Default Reasoning. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*. Montreal, Canada, August. Morgan Kaufmann Publishers. To be published.

Niemelä, I., and J. Rintanen. 1992. On the Impact of Stratification on the Complexity of Nonmonotonic Reasoning. In *Proceedings of the 3rd International Conference on Principles of Knowledge Representation and Reasoning*, 627–638. Cambridge, MA, USA, October. Morgan Kaufmann Publishers.

Poole, D. 1988. A Logical Framework for Default Reasoning. *Artificial Intelligence* 36:27–47.

Przymusinska, H., and T.C. Przymusinski. 1992. Stationary Default Extensions. In *Working Notes of the 4th International Workshop on on Nonmonotonic Reasoning*, 179–193. Plymouth, Vermont, USA, May.

Przymusinski, T.C. 1991a. Autoepistemic Logics of Closed Beliefs and Logic Programming. In *Proceedings of the 1st International Workshop on Logic Programming and Non-monotonic Reasoning*, 3–20. Washington, D.C., USA, July. The MIT Press.

Przymusinski, T.C. 1991b. Three-Valued Nonmonotonic Formalisms and Semantics of Logic Programming. *Artificial Intelligence* 49:309–343.

Reiter, R. 1980. A Logic for Default Reasoning. *Artificial Intelligence* 13:81–132.

Reiter, R. 1990. On Asking What a Database Knows. In *Proceedings of the Symposium on Computational Logic*, 95–113. Brussels, Belgium, November. Springer-Verlag.

Schwarz, G. 1992. Minimal Model Semantics for Nonmonotonic Modal Logics. In *Proceedings of the Seventh Annual IEEE Symposium on Logic in Computer Science*, 34–43. Santa Cruz, California, USA, June. IEEE Computer Society Press.

Selman, B., and H.J. Levesque. 1990. Abductive and Default Reasoning: A Computational Core. In *Proceedings of the 8th National Conference on Artificial Intelligence*, 343–348. Boston, MA, USA, July. The MIT Press.

Shvarts, G. 1990. Autoepistemic Modal Logics. In *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge*, 97–109. Pacific Grove, USA, March. Morgan Kaufmann Publishers.

Stärk, R.F. 1990. On the Existence of Fixpoints in Moore's Autoepistemic Logic and the Non-Monotonic Logic of McDermott and Doyle. In *Proceedings of the 4th Workshop on Computer Science Logic*, 354–365. Heidelberg, Germany, October. Springer-Verlag.

Stillman, J. 1990. It's not my Default: the Complexity of Membership Problems in Restricted Propositional Default Logics. In *Proceedings of the 8th National Conference on Artificial Intelligence*, 571–578. Boston, Massachusetts, USA, July. The MIT Press.

Stillman, J. 1992. The Complexity of Propositional Default Logics. In *Proceedings of the 10th National Conference on Artificial Intelligence*, 794–800. San Jose, California, USA, July. The MIT Press.

Tiomkin, M., and M. Kaminski. 1990. Nonmonotonic Default Modal Logics. In *Proceedings of the 3rd Conference on Theoretical Aspects of Reasoning about Knowledge*, 73–83. Pacific Grove, California, March. Morgan Kaufmann Publishers.

Tiomkin, M., and M. Kaminski. 1991. Nonmonotonic Default Modal Logics. *Journal of the ACM* 38(4):963–984.

Truszczyński, M. 1991a. Embedding Default Logic into Modal Nonmonotonic Logics. In *Proceedings of the 1st International Workshop on Logic Programming and Non-monotonic Reasoning*, 151–165. Washington, D.C., USA, July. The MIT Press.

Truszczyński, M. 1991b. Modal Interpretations of Default Logic. In *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, 393–398. Sydney, Australia, August. Morgan Kaufmann Publishers.

Truszczyński, M. 1991c. Modal Nonmonotonic Logic with Restricted Application of the Negation as Failure to Prove Rule. *Fundamenta Informaticae* 14:355–366.

Van Gelder, A., K.A. Ross, and J.S. Schlipf. 1991. The Well-Founded Semantics for General Logic Programs. *Journal of the ACM* 38(3):620–650.