A Decision Method for Nonmonotonic Reasoning Based on Autoepistemic Reasoning

Ilkka Niemelä
Department of Computer Science
Helsinki University of Technology
Otakaari 1, FIN–02150 Espoo, Finland
Ilkka.Niemela@hut.fi

Abstract

A novel decision method for autoepistemic reasoning is developed and proved correct. The method is applicable in a general setting, i.e. for an autoepistemic logic based on a given classical logic. It provides a decision procedure for a tightly grounded form of autoepistemic reasoning based on L-hierarchic expansions as well as for autoepistemic reasoning based on Moore style expansions and N-expansions. Prominent formalizations of nonmonotonic reasoning, such as default logic and circumscription, can be embedded into autoepistemic logic based on L-hierarchic expansions using simple local translations. Hence, the method can serve as a unified reasoning tool for a wide range of forms of nonmonotonic reasoning. The method is conceptually simple and the inherent sources of complexity and targets for optimization are clearly identifiable. As an example of exploiting optimization possibilities a new decision method for Reiter's default logic is developed where ideas from autoepistemic reasoning are used to efficiently prune the search space of applying default rules when constructing extensions of a default theory.

1 INTRODUCTION

We study the problem of automating nonmonotonic reasoning. Several forms of nonmonotonic reasoning have been proposed and in recent years research has been focused on relationships between different formalizations. Despite of different starting points and underlying intuitions, leading approaches to nonmonotonic reasoning have turned out to be closely related. Our aim is to exploit the close connections 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 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 decision problems in a given formalization to decision problems in the basic system. In this paper we consider the possibility of using autoepistemic logic, originally introduced by Moore (1985), as the basic system and we develop a novel decision method for autoepistemic reasoning which can serve as a basis for a unified reasoning tool for nonmonotonic reasoning.

Besides autoepistemic logic, there are many other interesting alternatives to be used as a basic system including logic programming, default logic (Reiter 1980), McDermott and Doyle style nonmonotonic modal logics (Marek, Shvarts and Truszczyński 1991), and systems based on two modalities (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 captured using simple local translations. When using a basic language with limited expressivity such as that of logic programs more complicated translations are needed. Autoepistemic logic has a uniform syntax unlike default logic. This enables a unified treatment of default rules, queries and integrity constraints. 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 where only straightforward principles of positive and negative introspection are added but no modal axioms 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, Elkan 1990). However, some of the embeddings are rather complicated because the original autoepistemic logic allows weakly grounded conclusions which have to be eliminated in the embeddings using extra conditions (Konolige 1988, Marek and Truszczyński 1989) or complicated translations (Gottlob 1993). In this paper we adopt a tightly grounded form of autoepistemic logic based on L-hierarchic expansions (Niemelä 1991) as the basic system. In L-hierarchic expansions the troublesome weakly grounded conclusions are excluded but without introducing any dependence on the syntactic representation of the premises. Furthermore, simple local translations can be used to reduce decision problems in other forms of nonmonotonic reasoning to decision problems in autoepistemic logic based on L-hierarchic expansions (Niemelä 1992b).

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 specific underlying proof method and some (Marek and Truszczyński 1991a, Junker and Konolige 1990, Marek and Truszczyński 1991b) are based on the idea of mapping a decision problem in autoepistemic logic into another problem. 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. At this stage we do not want to commit to a particular underlying 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 1990, Marek and Truszczyński 1991b). This aspect is further discussed in Section 4.

2 AUTOEPISTEMIC LOGIC

Autoepistemic logic is a modal logic with an operator L which is interpreted as "is believed". Autoepistemic logic models the beliefs of an ideally rational fully introspective agent. The main object of interest in autoepistemic logic is the set of beliefs of the agent given a set of sentences as the initial 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 entails 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 $\neg L\chi$ is).

We consider a general setting where a classical calculus 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). 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: if $\phi \in \mathcal{L}_{\mathrm{ae}}$ is a sentence, then $L\phi \in \mathcal{L}_{\mathrm{ae}}$. 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 $\forall x L P(x)$ are not allowed). To simplify the treatment we consider only closed formulae. So from here on we let $\mathcal{L}$ and $\mathcal{L}_{\mathrm{ae}}$ contain only closed formulae. In the general setting an ideally rational agent believes the logical consequences given by a consequence relation $\models_{\mathrm{ae}}$ which is a simple extension of the consequence relation $\models_{\mathrm{cl}}$ of CL where the $L\phi$ formulae are treated like atomic formulae in the propositional calculus.

To formalize a tight notion of autoepistemic reasoning Niemelä (1991) has proposed an enumeration-based method 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 which 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 agent having initial assumptions Σ after introspecting sentences in the order ε.

Definition 2.1 *Let $\Sigma \subseteq \mathcal{L}_{\mathrm{ae}}$. Let $\varepsilon = \psi_1, \psi_2, \ldots$ be an enumeration of sentences in $\mathcal{L}_{\mathrm{ae}}$. Let $\mathbf{B}^{\varepsilon}_0(\Sigma) = \emptyset$ and define $\mathbf{B}^{\varepsilon}_{i+1}(\Sigma)$ for $i = 0, 1, \ldots$ as follows:*

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

Finally let

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

To guarantee that the introspection principles of autoepistemic reasoning are respected the enumeration ε is required to be acceptable with respect to the set of premises.

Definition 2.2 *An enumeration ε is* Σ-acceptable *if there is no i and no formula ϕ such that $\neg L\phi \in \mathbf{B}^{\varepsilon}_i(\Sigma)$ but $\Sigma \cup \mathbf{B}^{\varepsilon}_i(\Sigma) \models_{\mathrm{ae}} \phi$.*

Example 1 Consider the propositional case, a set of premises $\Sigma = \{\neg Lp \rightarrow q\}$ and an enumeration $\varepsilon = q, p, \ldots$ Then $\mathbf{B}^{\varepsilon}_1(\Sigma) = \{\neg Lq\}$ as $\Sigma \not\models_{\mathrm{ae}} q$ and $\mathbf{B}^{\varepsilon}_2(\Sigma) = \{\neg Lq, \neg Lp\}$ as $\Sigma \cup \mathbf{B}^{\varepsilon}_1(\Sigma) \not\models_{\mathrm{ae}} p$. This enumeration is not Σ-acceptable because $\neg Lq \in \mathbf{B}^{\varepsilon}_2(\Sigma)$ but $\Sigma \cup \mathbf{B}^{\varepsilon}_2(\Sigma) \models_{\mathrm{ae}} q$. ■

We assume that the underlying consequence relation $\models_{\mathrm{ae}}$ is compact, i.e., if $\Sigma \models_{\mathrm{ae}} \phi$, there exists a finite subset Σ' of Σ such that $\Sigma' \models_{\mathrm{ae}} \phi$. For compact logics the acceptability condition implies that the introspection principles are not violated and the resulting set of beliefs is a Moore style expansion (Moore 1985).

Theorem 2.3 (Niemelä (1992a)) *For every Σ-acceptable enumeration ε of $\mathcal{L}_{\text{ae}}$, $\mathbf{SE}^{\varepsilon}(\Sigma)$ is a Moore style stable expansion of Σ, i.e.*

$$\mathbf{SE}^{\varepsilon}(\Sigma) = \{\phi \mid \Sigma \cup L\mathbf{SE}^{\varepsilon}(\Sigma) \cup \neg L\overline{\mathbf{SE}^{\varepsilon}(\Sigma)} \models_{\text{ae}} \phi\}$$

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

Enumeration-based expansions, i.e. sets of beliefs induced by acceptable enumerations, are a proper subclass of Moore style expansions. In the propositional case they coincide with the **N**-expansions (Marek et al. 1991) which are the same as iterative expansions (Marek and Truszczyński 1989) (see (Niemelä 1991)). The weakly grounded beliefs in **N**-expansions can be eliminated by requiring enumerations to be *L-hierarchic*. The resulting L-hierarchic expansions provide an interesting tightly grounded form of autoepistemic reasoning.

Definition 2.4 *Let ε be an enumeration $\psi_1, \psi_2, \ldots$ of sentences in $\mathcal{L}_{\text{ae}}$. The enumeration ε is* L-hierarchic *iff 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 iff $\Delta = \mathbf{SE}^{\varepsilon}(\Sigma)$ for some Σ-acceptable L-hierarchic enumeration ε of $\mathcal{L}_{\text{ae}}$.*

Example 2 Consider the propositional case, a set of premises

$$\Sigma = \{\neg L\neg Lp \rightarrow p\} \tag{1}$$

where p is an atomic formula and an enumeration $\varepsilon = \neg Lp, p, \ldots$ Now ε leads to an enumeration-based expansion containing an ungrounded belief p (see also (Niemelä 1991)). However, ε is not L-hierarchic. ■

An expansion of premises is a possible set of conclusions derivable from the premises. But as there are premises with multiple expansions, the concepts of cautious and brave reasoning are introduced. Given a class of expansions, a formula is derivable from a set of premises in the *cautious* sense iff it is in every expansion of the premises and in the *brave* sense iff it is at least in one of the expansions. We denote the derivability relations with respect to L-hierarchic, enumeration-based and Moore style expansions by $\mid\!\sim_{\text{c(LE)}}$, $\mid\!\sim_{\text{b(LE)}}$, $\mid\!\sim_{\text{c(E)}}$, $\mid\!\sim_{\text{b(E)}}$, $\mid\!\sim_{\text{c}}$, $\mid\!\sim_{\text{b}}$, respectively.

Expansions are infinite sets of sentences and a finitary characterization is needed to enable a computational treatment of autoepistemic reasoning. We use a characterization based on full sets (Niemelä 1990, Niemelä 1992a). The notion of full sets provides a simple and compact representation of expansions which is straightforwardly applicable in the general setting considered here. An overview of other approaches to characterizing expansions can be found in (Niemelä 1992a).

First we introduce some notation. For an $L\chi$ formula, q-atom$(L\chi)$ = q-atom$(\neg L\chi)$ = $L\chi$ and for a set Λ, q-atom$(\Lambda) = \{\text{q-atom}(\phi) \mid \phi \in \Lambda\}$. The set of all subformulae of the form $L\chi$ of a formula ϕ is denoted by $\mathrm{Sf}^{\mathrm{L}}(\phi)$ and $\mathrm{Sf}^{\mathrm{qL}}(\phi)$ is the set of all subformulae of the form $L\chi$ of ϕ which are not in the scope of another L operator in ϕ. We define a simple consequence relation $\models_{\mathrm{L}}$ which is given recursively on top of the underlying consequence relation $\models_{\text{ae}}$.

Definition 2.5 *Given a set of sentences Σ and a sentence ϕ,*

$$\Sigma \models_{\mathrm{L}} \phi \textit{ iff } \Sigma \cup \mathrm{SB}_{\Sigma}(\phi) \models_{\text{ae}} \phi$$

$$\textit{where} \quad \mathrm{SB}_{\Sigma}(\phi) \;=\; \begin{array}{l}\{L\chi \in \mathrm{Sf}^{\mathrm{qL}}(\phi) \mid \Sigma \models_{\mathrm{L}} \chi\} \cup \\ \{\neg L\chi \in \neg\mathrm{Sf}^{\mathrm{qL}}(\phi) \mid \Sigma \not\models_{\mathrm{L}} \chi\}.\end{array}$$

A finitary characterization for Moore style expansions can be provided in terms of sets of sentences constructed from the $L\chi$ and $\neg L\chi$ 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. In fact there is a one-to-one correspondence between full sets and expansions.

Definition 2.6 *For a set of sentences Σ, a set $\Lambda \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma) \cup \neg\mathrm{Sf}^{\mathrm{L}}(\Sigma)$ is Σ-full iff the following two conditions hold for every $L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)$:*

1. *$\Sigma \cup \Lambda \models_{\text{ae}} \chi$ iff $L\chi \in \Lambda$.*
2. *$\Sigma \cup \Lambda \not\models_{\text{ae}} \chi$ iff $\neg L\chi \in \Lambda$.*

Theorem 2.7 (Niemelä (1992a)) *Let Σ be a set of sentences of $\mathcal{L}_{\text{ae}}$. Then a function SE_{Σ} defined as*

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

gives a bijective mapping from the set of Σ-full sets to the set of Moore style stable expansions of Σ and $\mathrm{SE}_{\Sigma}(\Lambda)$ is the unique stable expansion Δ of Σ such that $\Lambda \subseteq L\Delta \cup \neg L\overline{\Delta}$.

Example 3 The premise (1) has two Σ-full sets $\Lambda_1 = \{L\neg Lp, \neg Lp\}$ and $\Lambda_2 = \{\neg L\neg Lp, Lp\}$. For example, Λ_1 is Σ-full as $\Sigma \cup \Lambda_1 \models_{\text{ae}} \neg Lp$ and $\Sigma \cup \Lambda_1 \not\models_{\text{ae}} p$. This implies that Σ has exactly two Moore style expansions $\mathrm{SE}_{\Sigma}(\Lambda_1)$ and $\mathrm{SE}_{\Sigma}(\Lambda_2)$. As $\Sigma \cup \Lambda_1 \not\models_{\mathrm{L}} p$, $p \notin \mathrm{SE}_{\Sigma}(\Lambda_1)$. However, $\neg LLp \in \mathrm{SE}_{\Sigma}(\Lambda_1)$ as $\Sigma \cup \Lambda_1 \models_{\mathrm{L}} \neg LLp$. ■

A finitary characterization for enumeration-based expansions is obtained by observing that full sets corresponding to enumeration-based (L-hierarchic) expansions are exactly those generated by acceptable (L-hierarchic) enumerations. First we note that for a finite enumeration the test for acceptability can be stated in the following simple form.

Proposition 2.8 (Niemelä (1992a)) *Let $\varepsilon_n = \psi_1, \ldots, \psi_n$ be an enumeration of a finite set.*

Then ε_n is Σ-acceptable iff ε_n satisfies the condition: for all $\neg L\phi \in \mathbf{B}_n^{\varepsilon_n}(\Sigma), \Sigma \cup \mathbf{B}_n^{\varepsilon_n}(\Sigma) \not\models_{\mathrm{ae}} \phi$.

Theorem 2.9 (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 $\{\phi \mid L\phi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ is Σ-acceptable, then $\mathbf{B}_n^{\varepsilon_n}(\Sigma)$ is a Σ-full set and there exists a Σ-acceptable enumeration ε of $\mathcal{L}_{\mathrm{ae}}$ such that $\mathbf{SE}^{\varepsilon}(\Sigma)$ is the unique stable expansion Δ of Σ for which $\mathbf{B}_n^{\varepsilon_n}(\Sigma) \subseteq L\Delta \cup \neg L\overline{\Delta}$ (i.e. $\mathbf{SE}^{\varepsilon}(\Sigma) = \mathrm{SE}_{\Sigma}(\mathbf{B}_n^{\varepsilon_n}(\Sigma))$). Furthermore, if ε_n is L-hierarchic, then ε is also.*

Theorem 2.10 (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 a Σ-acceptable (L-hierarchic) enumeration ε_L of $\{\phi \mid L\phi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ such that $\mathbf{B}^{\varepsilon_L}(\Sigma)$ is the unique Σ-full set Λ for which $\Lambda \subseteq L\mathbf{SE}^{\varepsilon}(\Sigma) \cup \neg L\overline{\mathbf{SE}^{\varepsilon}(\Sigma)}$ holds and thus $\mathbf{SE}^{\varepsilon}(\Sigma) = \mathrm{SE}_{\Sigma}(\mathbf{B}^{\varepsilon_L}(\Sigma))$.*

Example 4 The Moore style expansions in Example 3 are both characterizable using enumerations of $\{\phi \mid L\phi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\} = \{p, \neg Lp\}$: $\mathrm{SE}_{\Sigma}(\Lambda_1)$ by the enumeration $\varepsilon_2 = p, \neg Lp$ and $\mathrm{SE}_{\Sigma}(\Lambda_2)$ by the enumeration $\varepsilon_2' = \neg Lp, p$, i.e. $\Lambda_1 = \mathbf{B}_2^{\varepsilon_2}(\Sigma)$ and $\Lambda_2 = \mathbf{B}_2^{\varepsilon_2'}(\Sigma)$. This implies that Σ has exactly two enumeration-based expansions but only one L-hierarchic expansion, i.e. $\mathrm{SE}_{\Sigma}(\mathbf{B}_2^{\varepsilon_2}(\Sigma))$. ■

3 OTHER FORMS OF NONMONOTONIC REASONING

Autoepistemic logic based on L-hierarchic expansions provides an interesting unified basis for a wide range of forms of nonmonotonic reasoning. Leading formalizations such as default logic, circumscription, logic programs, justification-based truth maintenance systems and forms of abduction can be embedded into this variant of autoepistemic logic using simple local translations and without extra conditions or restriction (Niemelä 1992b). This implies that decision problems in these approaches can be reduced to decision problems in autoepistemic logic based on L-hierarchic expansions without introducing additional computational overhead.

Default logic (Reiter 1980) can be embedded into autoepistemic logic based on L-hierarchic expansions using the following translation.

$$\mathrm{AE}_{\mathrm{DL}}\left(\frac{a : b_1, \ldots, b_n}{c}\right) = (La \wedge L\neg L\neg b_1 \wedge \cdots \wedge L\neg L\neg b_n) \rightarrow c \qquad (2)$$

$$\mathrm{AE}_{\mathrm{DL}}(D, W) = W \cup \{\mathrm{AE}_{\mathrm{DL}}(d) \mid d \in D\} \qquad (3)$$

Theorem 3.1 (Niemelä (1992b)) *Let (D, W) be a default theory where D is finite. Then a set $E \subseteq \mathcal{L}$ is an extension of (D, W) iff $E = \Delta \cap \mathcal{L}$ for an L-hierarchic expansion Δ of $\mathrm{AE}_{\mathrm{DL}}(D, W)$.*

Default logic is closely related to the answer set semantics (Gelfond and Lifschitz 1990) of extended logic programs, i.e. logic programs where classical negation ($\neg$) is allowed even in the scope of the negation-as-failure operator (*not*). Consequently, it is not surprising that answer set semantics can be embedded into autoepistemic logic using the following mapping.

$$\mathrm{LP}_{\mathrm{AE}}(l_0 \leftarrow l_1, \ldots, l_m, not\, l_{m+1}, \ldots, not\, l_n) = Ll_1 \wedge \ldots \wedge Ll_m \wedge L\neg Ll_{m+1} \wedge \ldots \wedge L\neg Ll_n \rightarrow l_0 \qquad (4)$$

Theorem 3.2 (Niemelä (1992b)) *Let P be a finite extended logic program. S is an answer set of P iff S is the set of literals in an L-hierarchic expansion of $\mathrm{LP}_{\mathrm{AE}}(P)$.*

When considering logic programs where classical negation is not allowed, i.e. general logic programs, the answer set semantics coincides with the stable model semantics (Gelfond and Lifschitz 1990).

Corollary 3.3 *Let P be a finite general logic program. S is a stable model of P iff S is the set of atomic sentences in an L-hierarchic expansion of $\mathrm{LP}_{\mathrm{AE}}(P)$.*

Elkan (1990) has shown that given a set of justifications J the grounded model of J computed by a truth maintenance system is just a stable model of J when J is seen as a propositional logic program.

Corollary 3.4 *Let J be a finite set of justifications. M is a grounded model of J iff M is the set of atomic sentences in an L-hierarchic expansion of $\mathrm{LP}_{\mathrm{AE}}(J)$.*

4 AUTOMATING NONMONOTONIC REASONING

In the general setting autoepistemic reasoning is decidable if the underlying consequence relation $\models_{\mathrm{ae}}$ is decidable (Niemelä 1992a). Moreover, decision procedures can be obtained from the finitary characterization based on full sets in a very straightforward fashion. For brave reasoning an expansion containing a given formula must be found and for cautious reasoning it must be verified that a given formula is in all the expansions. This can be achieved by constructing the characterizing full set for each of the expansions and performing the membership test on the basis of the full set as implied by Theorem 2.7 using the consequence relation $\models_{\mathrm{L}}$ (see (Niemelä 1992a) for details). However, a decision procedure which is based directly on the finitary characterization is not very practical because to construct the full sets using a direct approach it is necessary to iterate over every subset of the set of $L\chi$ subformulae of the premises. If the cardinality of a set is n, it has 2^n subsets. Moreover, each iteration might be computationally costly, i.e. involve a number of calls to a theorem prover for the underlying consequence relation $\models_{\mathrm{ae}}$.

Decision problems in propositional autoepistemic reasoning are complete problems with respect to the second level of the polynomial time hierarchy (Niemelä 1992a) which implies that any general decision method for autoepistemic reasoning most likely has an exponential worst-case *time* complexity. 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

$$\Sigma = \{p_0, Lp_0 \rightarrow p_1, Lp_1 \rightarrow p_2, \ldots, Lp_{n-1} \rightarrow p_n\}. \quad (5)$$

A straightforward 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 $\models_{\rm ae}$ it is possible to find the full set without iterating over the subsets: for any Σ-full set Λ, $\Sigma \cup \Lambda \models_{\rm ae} p_0$ holds and thus $Lp_0 \in \Lambda$ must hold which in turn implies $Lp_1 \in \Lambda$ and so on.

We develop first a decision method for autoepistemic reasoning based on L-hierarchic expansions. The decision method builds Σ-full sets for L-hierarchic expansions of Σ. The Σ-full sets can be found by examining the enumerations of the set $\{\phi \mid L\phi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$. However, the enumerations are not important but the sets of beliefs induced by the enumerations. This observation cuts the combinatorial explosion considerably as typically there is a large number of enumerations inducing the same Σ-full set. The full sets are developed together as far as possible so that the same computation is not repeated by exploiting the monotonicity of the underlying consequence relation $\models_{\rm ae}$. The novel decision method for autoepistemic reasoning needs only a theorem prover for the underlying monotonic consequence relation $\models_{\rm ae}$ as a subroutine.

The approach differs from techniques where nonmonotonic reasoning is reduced to another problem such as a truth maintenance problem (Junker and Konolige 1990), a theorem proving problem (Marek and Truszczyński 1991b), or a constraint satisfaction problem (Ben-Eliyahu and Dechter 1991). These reductions provide valuable insights to the relationships between the corresponding problems and enable new techniques to be used for solving nonmonotonic reasoning problems. However, while there are interesting subclasses of nonmonotonic reasoning problems where the problem size increases only polynomially as a result of the reduction mapping (see, e.g. (Marek and Truszczyński 1991b, Ben-Eliyahu and Dechter 1991)), in the general case the reductions can lead to an exponential increase in the problem size. Hence the worst case *space* complexity becomes exponential. Moreover, the reductions are computationally complex.

As an example we consider the approach of Junker and Konolige (1990). 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 rather complex. This is because the set of justifications corresponding to premises Σ includes all monotonic justifications $\phi \leftarrow \omega_1, \ldots, \omega_n$ where $L\phi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)$ 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 $(\Sigma \cap \mathcal{L}) \cup \{\omega_1, \ldots, \omega_n\} \models_{\rm cl} \phi$ holds. For an $L\phi$ subformula of the premises, ϕ can possess a very large number of minimal proofs. 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 $\Sigma \cup H' \models_{\rm cl} \phi$, is closely related to logic-based abduction (Eiter and Gottlob 1992). 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 1992). As propositional brave autoepistemic reasoning is also Σ_2^p-complete (Niemelä 1992a), the search problem of finding a 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.

4.1 A NOVEL DECISION METHOD FOR AUTOEPISTEMIC REASONING

Figure 1 presents the function $\mathbf{der}_{\mathrm{LE}}$ which is the key 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 iff there exists an L-hierarchic expansion Δ of Σ containing $B \cup F$ such that $\mathbf{test}(\Sigma, \Lambda, \phi)$ 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 correspond to L-hierarchic expansions until a full set Λ is found for which $\mathbf{test}(\Sigma, \Lambda, \phi)$ returns true.

In the correctness proof of $\mathbf{der}_{\mathrm{LE}}$ we use the following

FUNCTION $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \phi)$
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 $\mathbf{test}(\Sigma, B, \phi)$
  ELSE
    return false
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**
  IF $\mathbf{der}_{\mathrm{LE}}(\Sigma, B \cup \{\neg L\chi\}, F, \phi)$ returns false **THEN**
    return $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F \cup \{L\chi\}, \phi)$
  ELSE
    return true
ELSE
  return false.

Figure 1: A Skeleton for the Decision Procedures for Autoepistemic Reasoning Based on L-hierarchic Expansions

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. 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 2.10 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. 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 $\{\phi \mid L\phi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ from Σ, i.e. $B' = \mathbf{B}^{\varepsilon_m}_m(\Sigma)$. 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 B *covers* all $L\chi'$ subformulae of χ, i.e. $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq$ q-atom(B), and $\Sigma \cup B \models_{\mathrm{ae}} \chi$ holds. 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 $\Sigma \cup B \models_{\mathrm{ae}} \phi$, then $L\phi$ belongs to the full set provided it exists. It is straightforward to verify that $\mathbf{pos}_{\mathrm{LE}}$ has the following properties.

Proposition 4.1 *Let* $\langle B', F'\rangle = \mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$.

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

REPEAT
  $B' := B$;
  FOR every $L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma) -$ q-atom(B') **DO**
    IF $\Sigma \cup B \models_{\mathrm{ae}} \chi$ **THEN**
      IF $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq$ q-atom(B) **THEN**
        $B := B \cup \{L\chi\}$
      ELSE
        $F := F \cup \{L\chi\}$
UNTIL $B = B'$;
return $\langle B, F\rangle$.

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

(I1) $B \subseteq B'$ *and* $F \subseteq F'$ *and if* q-atom$(B) \cup F \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$*, then* q-atom$(B') \cup F' \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$.

(I2) If there is some fully L-hierarchic enumeration generating B *from* Σ*, then there is a fully L-hierarchic enumeration generating* B' *from* Σ.

(I3) If $B \cup F \subseteq \mathbf{B}^{\varepsilon_n}_n(\Sigma)$ *for some* Σ*-acceptable L-hierarchic enumeration* ε_n *of* $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$*, then* $B' \cup F' \subseteq \mathbf{B}^{\varepsilon_n}_n(\Sigma)$.

(I4) For all $L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)$*, if* $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq$ q-atom(B') *and* $\Sigma \cup B' \models_{\mathrm{ae}} \chi$*, then* $L\chi \in$ q-atom(B').

Example 5 Consider the premises Σ in (5). Then $\mathbf{pos}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset)$ returns $\langle\{Lp_0, Lp_1, \ldots, Lp_{n-1}\}, \emptyset\rangle$. ■

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

If there is some $L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)$ not covered by $B \cup F$, $L\chi$ 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 $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq$ q-atom(B) 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 $\mathbf{pos}_{\mathrm{LE}}$.

There are two possibilities for such a formula $L\chi$: either $\neg L\chi$ is in the full set or $L\chi$ is in the full set. The two alternatives are handled by backtracking. First B is extended by $\neg L\chi$ and if $\mathbf{der}_{\mathrm{LE}}(\Sigma, B \cup \{\neg L\chi\}, F, \phi)$

returns false, $L\chi$ is taken as a frozen belief and $\mathbf{der}_{\mathrm{LE}}$ returns what $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F \cup \{L\chi\}, \phi)$ returns.

Example 6 (i) Consider the propositional case and the execution of $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \phi)$ where the premises Σ are given in (5). 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 $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \phi)$ is that of

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

(ii) Consider the premises $\Sigma = \{\neg Lp \rightarrow q, \neg Lq \rightarrow p\}$ and the execution of $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \phi)$. Here then $\mathbf{pos}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset)$ returns $\langle\emptyset, \emptyset\rangle$. As Lp and Lq are not covered by $\emptyset$ and p and q have no $L\chi$ subformulae, we can choose either of them. Let us choose Lp. Hence $\mathbf{der}_{\mathrm{LE}}(\Sigma, \{\neg Lp\}, \emptyset, \phi)$ is executed next. Now $\mathbf{pos}_{\mathrm{LE}}(\Sigma, \{\neg Lp\}, \emptyset)$ returns $\langle\{\neg Lp, Lq\}, \emptyset\rangle$ and

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

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

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

is executed. If this returns true, then $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \phi)$ returns true. Otherwise $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \{Lp, Lq\}, \phi)$ is executed and because $\mathbf{pos}_{\mathrm{LE}}(\Sigma, \emptyset, \{Lp, Lq\})$ returns $\langle\emptyset, \{Lp, Lq\}\rangle$, $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \phi)$ 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 $\mathbf{pos}_{\mathrm{LE}}$ which is used in $\mathbf{der}_{\mathrm{LE}}$ satisfies the four conditions in Proposition 4.1. 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 4.2 *Let Σ be a finite set of sentences. If B and F are sets of sentences such that* $\text{q-atom}(B) \cup F \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$*, there is a fully L-hierarchic enumeration generating B from Σ and $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \phi)$ 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 $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \phi)$ 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 $\text{q-atom}(B) \cup F \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$, let there exist a fully L-hierarchic enumeration generating B from Σ and let $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \phi)$ return true.

The function $\mathbf{der}_{\mathrm{LE}}$ starts by introspection. Using Proposition 2.8 and Proposition 4.1 (I1, I2) the following basic result can be established which shows that if $\mathrm{Sf}^{\mathrm{L}}(\Sigma) \subseteq \text{q-atom}(B) \cup F$ 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 $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \phi)$ returns true.*

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 4.1 (I1). By (**B**) there exists an acceptable enumeration for which **test** succeeds.

Let the cardinality of $\mathrm{SL}(\Sigma, B, F)$ be i. Let $\langle B', F'\rangle = \mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$. 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 $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \phi)$ returns true, there is $L\chi \in \mathrm{SL}(\Sigma, B', F')$ such that $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq \text{q-atom}(B')$ and either $\mathbf{der}_{\mathrm{LE}}(\Sigma, B' \cup \{\neg L\chi\}, F', \phi)$ returns true or $\mathbf{der}_{\mathrm{LE}}(\Sigma, B', F' \cup \{L\chi\}, \phi)$ returns true.

Assume that $\mathbf{der}_{\mathrm{LE}}(\Sigma, B', F' \cup \{L\chi\}, \phi)$ returns true. The cardinality of $\mathrm{SL}(\Sigma, B', F' \cup \{L\chi\})$ is less than i and $\text{q-atom}(B') \cup F' \cup \{L\chi\} \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$. As B is generated by a fully L-hierarchic enumeration, by Proposition 4.1 (I2) 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 $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \phi)$ returns true. By Proposition 4.1 (I1) $B \subseteq B'$ and $F \subseteq F'$ which implies $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$.

On the other hand, consider now the case where $\mathbf{der}_{\mathrm{LE}}(\Sigma, B' \cup \{\neg L\chi\}, F', \phi)$ returns true. The cardinality of $\mathrm{SL}(\Sigma, B' \cup \{\neg L\chi\}, F')$ less than i and $\text{q-atom}(B' \cup \{\neg L\chi\}) \cup F' \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$. 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 $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq \text{q-atom}(B') = \text{q-atom}(\mathbf{B}_{m'}^{\varepsilon_{m'}}(\Sigma))$ and as $\varepsilon_{m'}$ is fully L-hierarchic, $\varepsilon_{m'+1}$ is fully L-hierarchic. As $L\chi \notin \text{q-atom}(B')$, by Proposition 4.1 (I4) $\Sigma \cup \mathbf{B}_{m'}^{\varepsilon_{m'+1}}(\Sigma) \not\models_{\mathrm{ae}} \chi$. 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 $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \phi)$ returns true. By Proposition 4.1 (I1) $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$. □

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 4.3 *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 $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \phi)$ returns true. Then for every set B of sentences and F of $L\chi$ sentences such that $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$, $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \phi)$ returns true.*

Proof. We prove the lemma 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 $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \phi)$ 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 4.1 (I3, I4) and Proposition 2.8 it is straightforward to establish a result showing that if $\mathrm{SL}(\Sigma, B, F) = \emptyset$ after introspection, the function $\mathbf{der}_{\mathrm{LE}}$ returns true.

(**B**) $\mathrm{SL}(\Sigma, B', F') = \emptyset$ *implies that* $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \phi)$ *returns true where* $\langle B', F' \rangle = \mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$.

Let $\mathrm{SL}(\Sigma, B, F) = \emptyset$. Then by Proposition 4.1 (I1) $\mathrm{SL}(\Sigma, B', F') = \emptyset$ where $\langle B', F' \rangle = \mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$. Then by (**B**) $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \phi)$ 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 by (**B**) $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \phi)$ returns true.

Assume that $\mathrm{SL}(\Sigma, B', F') \neq \emptyset$. First we show that there exists a formula $L\chi \in \mathrm{SL}(\Sigma, B', F')$ such that $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq \text{q-atom}(B')$. 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 $\mathrm{Sf}^{\mathrm{L}}(\psi_i) \subseteq \text{q-atom}(B')$. Let $L\psi_j \in \mathrm{Sf}^{\mathrm{L}}(\psi_i)$. As ε_n is L-hierarchic, $j < i$. Therefore $L\psi_j \in \text{q-atom}(B') \cup F'$ by the minimality requirement for i. It sufficient to show that $L\psi_j \in F'$ implies $L\psi_j \in \text{q-atom}(B')$ because this implies $L\psi_j \in \text{q-atom}(B')$ and thus $\mathrm{Sf}^{\mathrm{L}}(\psi_i) \subseteq \text{q-atom}(B')$. Let $L\psi_j \in F'$. By Proposition 4.1 (I3) $L\psi_j \in \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ and $\Sigma \cup \mathbf{B}_{j-1}^{\varepsilon_n}(\Sigma) \models_{\mathrm{ae}} \psi_j$ holds. Using Proposition 4.1 (I3, I4) it can proved that

$$\mathbf{B}_{j-1}^{\varepsilon_n}(\Sigma) \subseteq B' \tag{6}$$

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.

As the enumeration ε_n is L-hierarchic, then by (6) $\mathrm{Sf}^{\mathrm{L}}(\psi_j) \subseteq \text{q-atom}(\mathbf{B}_{j-1}^{\varepsilon_n}(\Sigma)) \subseteq \text{q-atom}(B')$ and $\Sigma \cup B' \models_{\mathrm{ae}} \psi_j$ holds. By Proposition 4.1 (I4) $L\psi_j \in \text{q-atom}(B')$. Hence $\mathrm{Sf}^{\mathrm{L}}(\psi_i) \subseteq \text{q-atom}(B')$. So there is $L\chi (= L\psi_i) \in \mathrm{SL}(\Sigma, B', F')$ such that $\mathrm{Sf}^{\mathrm{L}}(\chi) \subseteq \text{q-atom}(B')$.

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 $\mathbf{der}_{\mathrm{LE}}(\Sigma, B' \cup \{\neg L\chi\}, F', \phi)$ returns true which implies that $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \phi)$ 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 $\mathbf{der}_{\mathrm{LE}}(\Sigma, B', F' \cup \{L\chi\}, \phi)$ returns true which implies that $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \phi)$ returns true. □

As a direct consequence of Lemmata 4.2 and 4.3 we have a soundness and completeness theorem for $\mathbf{der}_{\mathrm{LE}}$.

Theorem 4.4 *Let Σ be a finite set of sentences. Let B and F be sets for which $\text{q-atom}(B) \cup F \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$ and let there be a fully L-hierarchic enumeration generating B from Σ. Then $\mathbf{der}_{\mathrm{LE}}(\Sigma, B, F, \phi)$ returns true iff 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 $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \phi)$ returns true.*

The theorem implies that if the underlying consequence relation $\models_{\mathrm{ae}}$ is decidable, decision procedures for $\mid\!\sim_{\mathrm{c(LE)}}$, $\mid\!\sim_{\mathrm{b(LE)}}$, 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 iff $\Sigma \cup \Lambda \models_{\mathrm{L}} \phi$.

Theorem 4.5 *Let $\Sigma \subseteq \mathcal{L}_{\mathrm{ae}}$ be a finite set of sentences and $\phi \in \mathcal{L}_{\mathrm{ae}}$ a sentence.*

1. If $\mathbf{test}(\Sigma, \Lambda, \phi)$ returns true, then $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \phi)$ returns true iff Σ has an L-hierarchic expansion.

2. If $\mathbf{test}(\Sigma, \Lambda, \phi)$ returns true iff $\Sigma \cup \Lambda \models_{\mathrm{L}} \phi$, then $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \phi)$ returns true iff $\Sigma \mid\!\sim_{\mathrm{b(LE)}} \phi$.

3. If $\mathbf{test}(\Sigma, \Lambda, \phi)$ returns true iff $\Sigma \cup \Lambda \not\models_{\mathrm{L}} \phi$, then $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \phi)$ returns true iff $\Sigma \not\mid\!\sim_{\mathrm{c(LE)}} \phi$.

Proof. Clearly, $\text{q-atom}(\emptyset) \cup \emptyset \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$ and $\emptyset$ is generated from Σ by the empty enumeration which is fully L-hierarchic, of course. Then by Theorem 4.4 $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \phi)$ returns true iff there exists an L-hierarchic Σ-acceptable enumeration ε_n of $\{\chi \mid L\chi \in \mathrm{Sf}^{\mathrm{L}}(\Sigma)\}$ such that $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \phi)$ returns true.

(1) Let $\mathbf{test}(\Sigma, \Lambda, \phi)$ return true for every input. Then by Theorems 2.9 and 2.10 $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \phi)$ returns true iff Σ has an L-hierarchic expansion.

(2) Let $\mathbf{test}(\Sigma, \Lambda, \phi)$ return true iff $\Sigma \cup \Lambda \models_{\mathrm{L}} \phi$. ($\Rightarrow$) Let $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \phi)$ 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 $\Sigma \cup \mathbf{B}_n^{\varepsilon_n}(\Sigma) \models_{\mathrm{L}} \phi$. By Theorem 2.9 there is an L-hierarchic expansion Δ of Σ such that

$$\begin{aligned} \Delta &= \mathrm{SE}_{\Sigma}(\mathbf{B}_n^{\varepsilon_n}(\Sigma)) \\ &= \{\phi \in \mathcal{L}_{\mathrm{ae}} \mid \Sigma \cup \mathbf{B}_n^{\varepsilon_n}(\Sigma) \models_{\mathrm{L}} \phi\}. \end{aligned} \tag{7}$$

Thus $\phi \in \Delta$ and $\Sigma \mid\!\sim_{\mathrm{b(LE)}} \phi$. ($\Leftarrow$) Let $\Sigma \mid\!\sim_{\mathrm{b(LE)}} \phi$ hold. Then there exists an L-hierarchic expansion Δ of Σ containing ϕ. By Theorem 2.10 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 (7). Thus $\Sigma \cup \mathbf{B}_n^{\varepsilon_n}(\Sigma) \models_{\mathrm{L}} \phi$ and $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \phi)$ returns true. Hence $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \phi)$ returns true.

(3) Let $\mathbf{test}(\Sigma, \Lambda, \phi)$ return true iff $\Sigma \cup \Lambda \not\models_{\mathrm{L}} \phi$. Similarly, by Theorems 2.9 and 2.10 it can be shown

that $\mathbf{der}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset, \phi)$ returns true iff 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) imply 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 required monotonic reasoning is given by the underlying consequence relation $\models_{\mathrm{ae}}$ and in the decision method this task can be handled by a separate theorem prover. The conflict resolution task emerges when deciding which formula is included in the full set from the available $L\chi$ subformulae of the premises. In the decision method the conflict resolution task is solved by a simple backtracking scheme. If there is a conflict for some $\neg L\chi'$ included in the full to be constructed, the method backtracks to the last choice point where some $\neg L\chi$ was added and retracts this choice. 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). Notice that the decision method can be implemented to run in polynomial space provided that the theorem prover of $\models_{\mathrm{ae}}$ runs in polynomial space.

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 $\mathbf{pos}_{\mathrm{LE}}$. Now the frozen beliefs can be used in positive introspection because the groundedness requirement for Moore style expansions is 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. Properties of the underlying logic as well as possible regularities of the sets of premises can be utilized to optimize the computationally costly subtasks in 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 or restricted form of the premises to achieve highest possible performance. An interesting topic of further research is to study possibilities of exploiting the regular pattern in which the theorem prover for $\models_{\mathrm{ae}}$ is used and make use of 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.

4.2 A DECISION METHOD FOR DEFAULT LOGIC

As an example of developing optimization techniques we derive a new complete decision method for Reiter's default logic. By Theorem 3.1 default logic can be seen as a special case of autoepistemic logic based on L-hierarchic expansions. The simple logical character of the $L\chi$ sentences can be exploited when deciding $\models_{\mathrm{ae}}$ and the set of premises can be reduced considerably when the premises are restricted to the form resulting from the translation (2).

Proposition 4.6 (Niemelä(1992a)) *Let the underlying logic be a first-order calculus. Let $\Sigma \subseteq \mathcal{L}_{\mathrm{ae}}$ be a set of sentences of the form $(Lb_1 \wedge \cdots \wedge Lb_n) \rightarrow c$ where $c \in \mathcal{L}$ and $n \geq 0$ and Λ is a consistent set of $L\psi$ sentences and their negations. Then for all $\phi \in \mathcal{L}$,*

$$\Sigma \cup \Lambda \models_{\mathrm{ae}} \phi \text{ iff } \mathrm{Hds}(\Sigma, \Lambda) \models_{\mathrm{cl}} \phi$$

where $\mathrm{Hds}(\Sigma, \Lambda) = \{c \mid (Lb_1 \wedge \cdots \wedge Lb_n) \rightarrow c \in \Sigma$ *and* $Lb_i \in \Lambda$ *for all* $i = 1, \ldots, n\}$.

An interesting notion of safe disbeliefs can be developed for the subclass of autoepistemic sentences resulting from the translation (2). The idea is to strengthen the premises Σ by "assuming" $\neg L\phi$ for every $L\phi \notin B \cup F$ in addition to those $\neg L\phi$ already in B. It turns out that a disbelief $\neg L\psi$ is safe if ψ does not follow from the strengthened premises Σ^* and B using positive introspection. The function $\mathbf{safeneg}_{\mathrm{LE}}$ in Figure 3 computes these safe disbeliefs

It is easy to exploit safe disbeliefs in the decision method. Instead of propagating positive introspection by using $\mathbf{pos}_{\mathrm{LE}}$, both positive and safe negative introspection are propagated. This means that the function $\mathbf{intro}_{\mathrm{LE}}$ in Figure 3 is used instead of $\mathbf{pos}_{\mathrm{LE}}$ in the beginning of $\mathbf{der}_{\mathrm{LE}}$. Notice that $\mathbf{intro}_{\mathrm{LE}}$ also computes safe disbeliefs which are based on the fact that $\neg L\neg L\psi$ belongs to every consistent L-hierarchic expansion containing $L\psi$. We call the resulting function $\mathbf{der}_{\mathrm{DL}}$. The function $\mathbf{der}_{\mathrm{DL}}$ is sound and complete when given premises of the form resulting from the translation (2).

Example 7 Consider the propositional case, the premises $\Sigma = \{L\neg Lp \rightarrow q, L\neg Lq \rightarrow p\}$ and the ex-

FUNCTION $\mathbf{safeneg}_{\mathrm{LE}}(\Sigma, B, F)$
INPUT: Finite sets of sentences Σ, B, and F.

$\Sigma^* := \{La \to c \mid La \wedge L\neg Lb_1 \wedge \cdots \wedge L\neg Lb_n \to c \in \Sigma,$
for all $i = 1, \ldots, n, \neg Lb_i \in B$ or $Lb_i \notin B \cup F\}$;
$\langle B^*, F^* \rangle := \mathbf{pos}_{\mathrm{LE}}(\Sigma^*, B, F)$;
return $\{\neg L\psi \in \neg(\mathrm{Sf}^{\mathrm{L}}(\Sigma) - \text{q-atom}(B)) \mid \psi \in \mathcal{L},$
$\Sigma^* \cup B^* \not\models_{\mathrm{ae}} \psi\}$.

FUNCTION $\mathbf{intro}_{\mathrm{LE}}(\Sigma, B, F)$
INPUT: Finite sets of sentences Σ, B, and F.

REPEAT
$B' := B$;
$\langle B, F \rangle := \mathbf{pos}_{\mathrm{LE}}(\Sigma, B, F)$;
$B := B \cup \mathbf{safeneg}_{\mathrm{LE}}(\Sigma, B, F)$;
IF $\Sigma \cup B$ is consistent and
there exist $\neg L\chi \in B$ **THEN**
$B := B \cup \{\neg L \neg L\psi \in \neg(\mathrm{Sf}^{\mathrm{L}}(\Sigma) - \text{q-atom}(B)) \mid$
$L\psi \in B\}$
UNTIL $B = B'$;
return $\langle B, F \rangle$.

Figure 3: Functions for Safe Negative Introspection and Propagation of Introspection

ecution of $\mathbf{der}_{\mathrm{DL}}(\Sigma, \emptyset, \emptyset, \phi)$. Compare this to Example 6 where safe negative introspection is not used. First $\mathbf{intro}_{\mathrm{LE}}(\Sigma, \emptyset, \emptyset)$ returns $\langle \emptyset, \emptyset \rangle$. Here we can extend the full set by $(\neg)Lp$ or $(\neg)Lq$. We choose Lp. Hence first $\mathbf{der}_{\mathrm{DL}}(\Sigma, \{\neg Lp\}, \emptyset, \phi)$ is executed. Because $\langle \{\neg Lp, L\neg Lp, Lq, \neg L\neg Lq\}, \emptyset \rangle$ is returned by $\mathbf{intro}_{\mathrm{LE}}(\Sigma, \{\neg Lp\}, \emptyset)$,

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

is executed. If this test returns true, $\mathbf{der}_{\mathrm{DL}}(\Sigma, \emptyset, \emptyset, \phi)$ returns true. Otherwise $\mathbf{der}_{\mathrm{DL}}(\Sigma, \emptyset, \{Lp\}, \phi)$ is executed. As $\langle \{Lp, \neg L\neg Lp, \neg Lq, L\neg Lq\}, \emptyset \rangle$ is returned by $\mathbf{intro}_{\mathrm{LE}}(\Sigma, \emptyset, \{Lp\})$, $\mathbf{der}_{\mathrm{DL}}(\Sigma, \emptyset, \emptyset, \phi)$ returns what $\mathbf{test}(\Sigma, \{Lp, \neg L\neg Lp, \neg Lq, L\neg Lq\}, \phi)$ returns. ■

Theorem 4.7 *Let Σ be a finite set of sentences of the form $(La \wedge L\neg L\neg b_1 \wedge \cdots \wedge L\neg L\neg b_n) \to c$. Let B and F be sets of sentences such that $\text{q-atom}(B) \cup F \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$ and let there be a fully L-hierarchic enumeration generating B from Σ. Then $\mathbf{der}_{\mathrm{DL}}(\Sigma, B, F, \phi)$ returns true iff 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 $\mathbf{test}(\Sigma, \mathbf{B}_n^{\varepsilon_n}(\Sigma), \phi)$ returns true.*

Proof. This is Theorem 4.4 except that the form of the premises is restricted and the function $\mathbf{intro}_{\mathrm{LE}}$ is employed instead of just $\mathbf{pos}_{\mathrm{LE}}$ in $\mathbf{der}_{\mathrm{DL}}$. To prove Theorem 4.4 it was sufficient that $\mathbf{pos}_{\mathrm{LE}}$ satisfied the four conditions in Proposition 4.1. Hence to prove Theorem 4.7 it is sufficient to show that $\mathbf{intro}_{\mathrm{LE}}$ satisfies the same four conditions (I1–I4):

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

(I1) *$B \subseteq B'$ and $F \subseteq F'$ and if $\text{q-atom}(B) \cup F \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$, then $\text{q-atom}(B') \cup F' \subseteq \mathrm{Sf}^{\mathrm{L}}(\Sigma)$.*

(I2) *If there is some fully L-hierarchic enumeration generating B from Σ, then there is a fully L-hierarchic enumeration generating B' from Σ.*

(I3) *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 $B' \cup F' \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$.*

(I4) *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 $L\chi \in \text{q-atom}(B')$.*

We prove that each condition is satisfied by showing for every step in the body of the repeat loop in $\mathbf{intro}_{\mathrm{LE}}$ that if the condition holds before the step, then it holds after the step. By Proposition 4.1 the conditions hold for $\mathbf{pos}_{\mathrm{LE}}$ and it remains to show that the conditions are satisfied by the two other steps where B and F are modified in $\mathbf{intro}_{\mathrm{LE}}$.

(I1) This clearly holds for $\mathbf{intro}_{\mathrm{LE}}$.

(I2) Let B be generated by a fully L-hierarchic enumeration from Σ.

($\mathbf{safeneg}_{\mathrm{LE}}$) Let $B_{\mathrm{sn}} = \mathbf{safeneg}_{\mathrm{LE}}(\Sigma, B, F)$. We establish that there exists a fully L-hierarchic enumeration generating $B \cup B_{\mathrm{sn}}$ from Σ. We show that for all $\neg L\psi \in B_{\mathrm{sn}}$,

$$\Sigma \cup B \cup B_{\mathrm{sn}} \not\models_{\mathrm{ae}} \psi. \tag{8}$$

Now $B \cup B_{\mathrm{sn}}$ is generated by an enumeration $\varepsilon_{m'}$ which is obtained by extending ε_m with formulae $\{\psi \mid \neg L\psi \in B_{\mathrm{sn}}\}$ in some order. This is because it can be shown that for every i, $m < i \leq m'$, $\mathbf{B}_{i-1}^{\varepsilon_{m'}}(\Sigma) \subseteq B \cup B_{\mathrm{sn}}$ which by (8) implies $\Sigma \cup \mathbf{B}_{i-1}^{\varepsilon_{m'}}(\Sigma) \not\models_{\mathrm{ae}} \psi_i$ and thus $\neg L\psi_i \in \mathbf{B}_{m'}^{\varepsilon_{m'}}(\Sigma)$.

So it remains to show that (8) holds. We use the following two results in the proof. The results are straightforward to establish.

NegBels: *Let B be generated by some fully L-hierarchic enumeration from Σ. If $L\neg Lb \in B$ and $\Sigma \cup B$ is consistent, then $\neg Lb \in B$.*

Incons: *Let B be generated by some fully L-hierarchic enumeration from Σ. If $\Sigma \cup B$ is inconsistent, then $\Sigma^* \cup \{L\phi \in B \mid \phi \in \mathcal{L}\}$ is inconsistent.*

As for all $\neg L\psi \in B_{\mathrm{sn}}$, $\Sigma^* \cup B^* \not\models_{\mathrm{ae}} \psi$, (8) can be proved by showing that for all $\neg L\psi \in B_{\mathrm{sn}}$,

$$\Sigma^* \cup B^* \not\models_{\mathrm{ae}} \psi \text{ implies } \Sigma \cup B \cup B_{\mathrm{sn}} \not\models_{\mathrm{ae}} \psi. \tag{9}$$

If $\Sigma \cup B$ is inconsistent, then by (**Incons**) $\Sigma^* \cup \{L\phi \in B \mid \phi \in \mathcal{L}\}$ is inconsistent. By Proposition 4.1 (I1) $B \subseteq B^*$. Then also $\Sigma^* \cup B^*$ is inconsistent. This implies $B_{\mathrm{sn}} = \emptyset$.

Let $\Sigma \cup B$ be consistent. Let $\Sigma \cup B \cup B_{\rm sn} \models_{\rm ae} \psi$ hold. Then by Proposition 4.6 $\mathrm{Hds}(\Sigma, B \cup B_{\rm sn}) \models_{\rm cl} \psi$ holds. By Proposition 4.1 (I1) $B \subseteq B^*$. Using this and (**NegBels**) we can show

$$\mathrm{Hds}(\Sigma, B \cup B_{\rm sn}) \subseteq \mathrm{Hds}(\Sigma^*, B^*). \qquad (10)$$

By (10) $\mathrm{Hds}(\Sigma^*, B^*) \models_{\rm cl} \psi$ holds and by Proposition 4.6 $\Sigma^* \cup B^* \models_{\rm ae} \psi$ holds. Hence (8) holds.

($\neg L \neg L \psi$) Let $\Sigma \cup B$ be consistent and there exist $\neg L\chi \in B$. Let $\neg L \neg L \psi \in B_{\rm sn'} = \{\neg L \neg L \psi \in \neg(\mathrm{Sf}^{\rm L}(\Sigma) - \text{q-atom}(B)) \mid L\psi \in B\}$ but assume $\Sigma \cup B \cup B_{\rm sn'} \models_{\rm ae} \neg L\psi$. As $L\psi \in B$, $\Sigma \cup B \cup B_{\rm sn'}$ is inconsistent. As $\mathrm{Hds}(\Sigma, B \cup B_{\rm sn'}) = \mathrm{Hds}(\Sigma, B)$, by Proposition 4.6 $\Sigma \cup B$ is inconsistent, a contradiction. Hence for every $\neg L \neg L\psi \in B_{\rm sn'}$, $\Sigma \cup B \cup B_{\rm sn'} \not\models_{\rm ae} \neg L\psi$ which implies that there is a fully L-hierarchic enumeration generating $B \cup B_{\rm sn'}$ from Σ. Thus we have shown (I2).

(I3) Let there be a Σ-acceptable L-hierarchic enumeration $\varepsilon_n = \psi_1, \ldots, \psi_n$ of $\{\chi \mid L\chi \in \mathrm{Sf}^{\rm L}(\Sigma)\}$ such that $B \cup F \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ holds.

(**safeneg**$_{\rm LE}$) Let $B_{\rm sn} = \mathbf{safeneg}_{\rm LE}(\Sigma, B, F)$ and we prove

$$B_{\rm sn} \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma). \qquad (11)$$

Using Proposition 4.6, (**NegBels**), (**Incons**), and Proposition 4.1 (I4) for B^* we can show by induction on i that for all $i = 1, \ldots, n$,

$$\psi_i \in \mathcal{L} \text{ and } L\psi_i \in \mathbf{B}_i^{\varepsilon_n}(\Sigma) \text{ implies}$$
$$\Sigma^* \cup B^* \models_{\rm ae} \psi_i. \qquad (12)$$

Using (12) it is straightforward to establish (11). Let $\neg L\psi \in B_{\rm sn}$. Then $\Sigma^* \cup B^* \not\models_{\rm ae} \psi$ holds, $L\psi \in \mathrm{Sf}^{\rm L}(\Sigma)$ and $\psi \in \mathcal{L}$. Hence ψ is some ψ_i in the enumeration ε_n. By (12) $L\psi_i \notin \mathbf{B}_i^{\varepsilon_n}(\Sigma)$ which implies that $\neg L\psi_i \in \mathbf{B}_n^{\varepsilon_n}(\Sigma)$. Hence (11) holds.

($\neg L \neg L \psi$) Let there exist $\neg L\chi \in B$. Let $\neg L \neg L \psi \in B_{\rm sn'} = \{\neg L \neg L \psi \in \neg(\mathrm{Sf}^{\rm L}(\Sigma) - \text{q-atom}(B)) \mid L\psi \in B\}$. As $\neg L\chi \in B \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$, $\Sigma \cup \mathbf{B}_n^{\varepsilon_n}(\Sigma)$ is consistent. As $L\psi \in B \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$, by (**NegBels**) $\neg L \neg L\psi \in \mathbf{B}_n^{\varepsilon_n}(\Sigma)$. Hence $B_{\rm sn'} \subseteq \mathbf{B}_n^{\varepsilon_n}(\Sigma)$. Thus we have established (I3).

(I4) By Proposition 4.1 (I4) holds for B_1 where $\langle B_1, F_1 \rangle = \mathbf{pos}_{\rm LE}(\Sigma, B, F)$ and this implies that (I4) holds for B'.

□

The resulting decision method does not use the obvious idea of building an extension of a default theory by trying to find a suitable order of applying the default rules but it constructs the set of applicable rules by cautiously building a set representing the valid premises of the rules, i.e. the set B. Autoepistemic logic offers a convenient language for this as it easy to represent the fact that a prerequisite a of a rule holds (La) and that it holds that a justification b is consistent ($L\neg L\neg b$). Every time a choice is made an approximation of the intersection of the possible extensions is constructed using safe disbeliefs and propagation of positive introspection in a way which is closely related to stationary default extensions (Przymusinska and Przymusinski 1992). This can reduce the number of choice points in backtracking considerably. However, here we have integrated the approximation method efficiently in the backtracking search: not only the partially constructed set of valid premises (B) is used but also the information from the backtracking search, i.e. the frozen beliefs F, are employed when constructing the next approximation.

5 CONCLUSIONS

We have developed 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 has been to devise a conceptually clear method where sources of complexity and targets for further optimization are clearly identifiable. Two orthogonal sources can be identified: classical reasoning and conflict resolution. The classical reasoning task can be 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 simple 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 are obtained by minor adjustments. The method can be implemented to run in polynomial space in the propositional case. Leading forms of nonmonotonic reasoning, such as default logic and circumscription, can be embedded into autoepistemic logic based on L-hierarchic expansions using simple local translations. Thus the method can be used as a unified decision method for a large range of forms of nonmonotonic reasoning. As an example of the optimization possibilities we develop a new complete decision method for default logic where the search space of applying default rules is pruned efficiently by using ideas from autoepistemic reasoning. The method provides an interesting approach to implementing truth maintenance systems as well as stable model semantics and answer set semantics of logic programs.

References

Ben-Eliyahu, R. and Dechter, R. (1991). Default logic, propositional logic and constraints, *Proceedings of the 9th National Conference on Artificial Intelligence*, The MIT Press, pp. 370–385.

Eiter, T. and Gottlob, G. (1992). The complexity of logic-based abduction, *Technical Report CD-TR 92/35*, Christian Doppler Labor für Expertensysteme, Institut für Informationssysteme, Technische Universität Wien, Vienna, Austria.

Elkan, C. (1990). A rational reconstruction of nonmonotonic truth maintenance systems, *Artificial Intelligence* **43**: 219–234.

Gelfond, M. and Lifschitz, V. (1990). Logic programs with classical negation., *Proceedings of the 7th International Conference on Logic Programming*, The MIT Press, Jerusalem, Israel, pp. 579–597.

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, *Proceedings of the 13th International Joint Conference on Artificial Intelligence*, Morgan Kaufmann Publishers, Chambéry, France, pp. 570–575.

Junker, U. and Konolige, K. (1990). Computing the extensions of autoepistemic and default logics with a truth maintenance system, *Proceedings of the 8th National Conference on Artificial Intelligence*, The MIT Press, Boston, MA, USA, pp. 278–283.

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, *Proceedings of the 11th International Joint Conference on Artificial Intelligence*, Morgan Kaufmann Publishers, Detroit, Michigan, USA, pp. 1213–1218.

Lifschitz, V. (1991). Nonmonotonic databases and epistemic queries, *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, Morgan Kaufmann Publishers, Sydney, Australia, pp. 381–386.

Marek, W. and Truszczyński, M. (1989). Relating autoepistemic and default logics, *Proceedings of the 1st International Conference on Principles of Knowledge Representation and Reasoning*, Toronto, Canada, pp. 276–288.

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

Marek, W. and Truszczyński, M. (1991b). Computing intersection of autoepistemic expansions, *Proceedings of the 1st International Workshop on Logic Programming and Non-monotonic Reasoning*, The MIT Press, Washington, D.C., USA, pp. 37–50.

Marek, W., Shvarts, G. and Truszczyński, M. (1991). Modal nonmonotonic logics: Ranges, characterization, computation, *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning*, Morgan Kaufmann Publishers, Cambridge, MA, USA, pp. 395–404.

Moore, R. (1985). Semantical considerations on nonmonotonic logic, *Artificial Intelligence* **25**: 75–94.

Moore, R. (1988). Autoepistemic logic, *in* P. Smets, E. Mamdani and D. Dubois (eds), *Non-Standard Logics for Automated Reasoning*, Academic Press, London, pp. 105–136.

Niemelä, I. (1988). Decision procedure for autoepistemic logic, *Proceedings of the 9th International Conference on Automated Deduction*, Springer-Verlag, Argonne, USA, pp. 675–684.

Niemelä, I. (1990). Towards automatic autoepistemic reasoning, *Proceedings of the European Workshop on Logics in Artificial Intelligence—JELIA'90*, Springer-Verlag, Amsterdam, The Netherlands, pp. 428–443.

Niemelä, I. (1991). Constructive tightly grounded autoepistemic reasoning, *Proceedings of the 12th International Joint Conference on Artificial Intelligence*, Morgan Kaufmann Publishers, Sydney, Australia, pp. 399–404.

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, *Proceedings of the 10th European Conference on Artificial Intelligence*, John Wiley, Vienna, Austria, pp. 334–338.

Przymusinska, H. and Przymusinski, T. (1992). Stationary default extensions, *Working Notes of the 4th International Workshop on on Nonmonotonic Reasoning*, Plymouth, Vermont, USA, pp. 179–193.

Reiter, R. (1980). A logic for default reasoning, *Artificial Intelligence* **13**: 81–132.

Stärk, R. (1990). On the existence of fixpoints in Moore's autoepistemic logic and the nonmonotonic logic of McDermott and Doyle, *Proceedings of the 4th Workshop on Computer Science Logic*, Springer-Verlag, Heidelberg, Germany, pp. 354–365.