

A Survey on Complexity Results for Non-monotonic Logics*

Marco Cadoli, Marco Schaerf

Dipartimento di Informatica e Sistemistica

Università di Roma “La Sapienza”

via Salaria 113, 00198 Roma, Italia

cadoli@assi.ing.uniroma1.it schaerf@assi.ing.uniroma1.it

Abstract

This paper surveys the main results appeared in the literature on the computational complexity of non-monotonic inference tasks. We not only give results about the tractability/intractability of the individual problems but we also analyze sources of complexity and explain intuitively the nature of easy/hard cases. We focus mainly on non-monotonic formalisms, like default logic, autoepistemic logic, circumscription, closed-world reasoning and abduction, whose relations with logic programming are clear and well studied. Complexity as well as recursion-theoretic results are surveyed.

*Work partially supported by the ESPRIT Basic Research Action COMPULOG and the Progetto Finalizzato Informatica of the CNR (Italian Research Council). The first author is supported by a CNR scholarship

1 Introduction

Non-monotonic logics and negation as failure in logic programming have been defined with the goal of providing formal tools for the representation of *default* information. One of the ideas underlying both areas is that the use of default assumptions should lead to a more compact representation of knowledge. As a consequence, non-monotonic knowledge bases should be more space-effective than ordinary ones and this may hopefully have an impact on the performances of theorem provers.

This expectation generated an interesting activity in the development of algorithms for reasoning under non-monotonicity, as well as many studies about the inherent complexity of the inference tasks.

The goal of this paper is to survey the main results appeared in the literature on the computational complexity of non-monotonic inference tasks. We not only give results about the tractability/intractability of the individual problems but we also analyze sources of complexity and explain intuitively the nature of easy/hard cases. Furthermore, we create an extensive list of references to the literature that, in our opinion, will provide an useful tool for researchers interested in designing algorithms for non-monotonic reasoning (NMR in the sequel).

We focus mainly on non-monotonic formalisms, like default logic, autoepistemic logic, circumscription, closed-world reasoning and abduction, whose relations with logic programming are clear and well studied. Other reasoning problems, like reasoning on inheritance networks or belief revision are only briefly mentioned.

We refer the reader to works like [39, 89] for extensive surveys on NMR and to [123] for a survey on the relations between NMR and logic programming. Works on NMR that survey complexity results have been done by Minker [103] and Schlipf [139].

Many results that have appeared in the literature are concerned with decidable fragments of non-monotonic logics. In the spirit of a well-established trend in knowledge representation [20], they aim at the characterization of the expressive power of languages having polynomial time reasoning procedures. Other works deal with fully expressive languages and try to characterize the precise complexity of the inference task. In this paper we are interested in both kinds of results, privileging the analysis of propositional, decidable languages.

We use the jargon of computational complexity and recursion theory, as found in [54] and in [131], respectively. In particular we make use of the notions of polynomial, arithmetical and analytical hierarchies.

Using notions of higher-order complexity is necessary, since a general property of non-monotonic inference is that its computational complexity is higher than the complexity of the underlying monotonic logic. As an example, restricting the expressiveness of the language to Horn clauses allows for polynomial inference as far as classical propositional logic is concerned [40], but the inference task becomes NP-hard when propositional default logic [150] or circumscription [26] are considered.

The fact that non-monotonicity adds complexity to reasoning was clear from the first studies: in his seminal paper, Reiter [129] showed that inference in default logic is not

r. e. The issue of determining precise lower and upper bounds was addressed later and some formalisms have been proved to be complete for precise levels of the arithmetical or analytical hierarchy [5, 25, 34, 135]. Analogous completeness results have been found for the propositional versions of the same formalisms wrt some levels of the polynomial hierarchy [41, 64].

Part of the increase in the complexity of inference can be explained by noticing that the semantic definitions of most non-monotonic formalisms are either based on fixpoint constructions or on conditions requiring some form of *minimality*. This apparently gives a completely orthogonal source of complexity.

Examples of formalisms based on fixpoint semantics are default [129] and autoepistemic logic [104] where deduction is performed wrt *extensions* or *expansions*, which are solutions of fixpoint equations. As we see in the following sections, this fixpoint construction requires an additional non-deterministic choice that cannot be polynomially reduced to a deterministic one unless the polynomial hierarchy collapses.

An example of non-monotonic formalism based on a form of minimality is circumscription [85, 98, 99], in which inference is performed wrt the models of a first-order formula in which the extension of some selected predicates is minimized. As it turns out, the minimality requirement gives rise to a computational overhead that is analogous to that given by the fixpoint constructions.

The paper is organized as follows: after an introductory section on complexity concepts (Section 2), we first survey the formalisms based on fixpoint constructions, starting with default logic, which is presented in Section 3. The structure of reasoning in default logic makes it easy to precisely point out the various sources of complexity present in this kind of non-monotonic formalisms. We then discuss in Section 4 modal non-monotonic logics with particular attention to autoepistemic logic. The introduction of negation in logic programming has produced a proliferation of semantics, some based on fixpoint constructions (e. g. stable model semantics) and other on minimal models (e. g. well-founded semantics). Complexity results for the various semantics are surveyed in Section 5. In Section 6 we discuss non-monotonic formalisms whose semantics is based on some form of minimality, which are the different forms of closed-world reasoning and circumscription. In Section 7 we carry on an analysis of the complexity of the so-called logic-based abduction and, very briefly, mention results for other forms of abduction. While abduction is not strictly a formalism for NMR, nevertheless it is tightly related both to it and to logic programming. In Section 8 we discuss in some detail the various reductions between NMR problems pointing out their importance from the point of view of computational complexity analysis. Finally, in Section 9 we draw some conclusions.

2 Complexity classes

In this section we give a brief overview of complexity concepts that will be used throughout the paper. We refer the reader to [54, 74] for a thorough introduction in the field of complexity.

In this paper we deal most of the times with *decision problems*, i. e. problems that admit a boolean answer. For decision problems the class P is the set of problems that

can be answered by a Turing machine in polynomial time. Often we refer to computations done by non-deterministic Turing machines. The class of decision problems that can be solved by a non-deterministic Turing machine in polynomial time — where it is understood that the answer is **yes** provided *at least one* of the computations done in parallel by the machine ends in an accepting state — is denoted by NP. The class of problems whose answer is always the complement of those in NP, is denoted by co-NP. Also problems in co-NP can be solved by a non-deterministic Turing machine in polynomial time, but it is understood that the answer is **yes** provided *all* the computations done in parallel by the machine end in an accepting state. The class P is obviously contained both in NP and in co-NP.

An example of a problem in NP is testing satisfiability of a propositional formula: a formula T is satisfiable iff *at least one* truth assignment M such that $M \models T$ exists. An example of a problem in co-NP is testing if a propositional formula T entails a propositional formula γ : $T \models \gamma$ iff *for all* truth assignments M it holds that $(M \models T) \implies (M \models \gamma)$. In fact propositional satisfiability (entailment) is an NP-*complete* (co-NP-*complete*) problem, i. e. the “toughest” — wrt many-one polynomial reducibility — problem in the class NP (co-NP). We recall that the best algorithms known for solving NP-complete or co-NP-complete problems require exponential time in the worst case, and that the following relations are conjectured: $P \subset NP$, $P \subset \text{co-NP}$, $NP \neq \text{co-NP}$.

In the following we refer to a particular type of computation called computation with *oracles*. Oracles are intuitively subroutines without cost. Given a class of decision problems C , the class P^C (NP^C) is the class of decision problems that can be solved in polynomial time by a deterministic (non-deterministic) machine that uses an oracle for the problems in C , i. e. a subroutine for any problem in C that can be called several times, spending just one time-unit for each call.

The definition of *polynomial hierarchy* is based on oracle computations. The classes Σ_k^p , Π_k^p and Δ_k^p of the polynomial hierarchy are defined by

$$\Sigma_0^p = \Pi_0^p = \Delta_0^p = P$$

and for all $k \geq 0$,

$$\Sigma_{k+1}^p = NP^{\Sigma_k^p}, \Pi_{k+1}^p = \text{co-}\Sigma_{k+1}^p, \Delta_{k+1}^p = P^{\Sigma_k^p}.$$

Notice that $\Sigma_1^p = NP$, $\Pi_1^p = \text{co-NP}$, $\Delta_1^p = P$. The following relations have been conjectured:

$$\Delta_k^p \subset \Sigma_k^p, \Delta_k^p \subset \Pi_k^p, \Sigma_k^p \subset \Delta_{k+1}^p, \Pi_k^p \subset \Delta_{k+1}^p, \Sigma_k^p \neq \Pi_k^p$$

for all $k \geq 1$.

We say that a problem is at the k -th level of the polynomial hierarchy if it is Δ_{k+1}^p -complete under polynomial Turing reductions, i. e. it is in Δ_{k+1}^p and it is either Σ_k^p -hard or Π_k^p -hard. Propositional satisfiability and entailment are both at the first level.

In the following we focus on problems in Σ_2^p ($=NP^{NP}$) and Π_2^p ($=\text{co-NP}^{NP}$), as they are important in this paper. Σ_2^p contains all problems solvable in non-deterministic polynomial time provided it is possible to use for free a subroutine for a problem in NP, for example propositional satisfiability. The prototypical Σ_2^p -complete problem is testing validity of a quantified boolean formula Q (called $QB F_{2,\exists}$ formula) of this kind:

$$\exists x_1 \cdots \exists x_n \forall y_1 \cdots \forall y_m E$$

where $x_1, \dots, x_n, y_1, \dots, y_m$ are distinct propositional letters and E is a purely propositional formula built on such letters. The formula Q is valid iff there exists a truth value assignment to the propositional letters x_1, \dots, x_n such that for each possible extension of such an assignments to the letters y_1, \dots, y_m , Q is true. The prototypical Π_2^p -complete problem is testing validity of a quantified boolean formula of this kind (called $QBF_{2,\forall}$ formula):

$$\forall x_1 \cdots \forall x_n \exists y_1 \cdots \exists y_m E$$

We refer the reader to [111] for examples of Δ_2^p -complete problems.

In virtue of the conjectures $\Sigma_k^p \subset \Sigma_{k+1}^p$ and $\Pi_k^p \subset \Pi_{k+1}^p$, Σ_2^p -complete and Π_2^p -complete problems are considered more difficult to solve than both NP-complete and co-NP complete problems. A practical difference exists between NP-complete and Σ_2^p -complete problems. Suppose we have a good heuristic for solving an NP-complete problem (as an example propositional satisfiability testing), that normally solves the problem in an acceptable amount of time — even if exponential time is needed in the worst case. It is still not immediate how to use such a good heuristic for solving efficiently a Σ_2^p -complete problem. As a matter of fact, if the conjecture $\text{NP} \subset \Sigma_2^p$ is true, then in the worst case an exponential number of calls to an algorithm for satisfiability testing would be necessary in order to solve any Σ_2^p -complete problem. The same holds for Π_2^p -complete problems. Methods such as GSAT [145] for efficient handling of NP-complete problems are therefore most likely not applicable to problems at the second level of the polynomial hierarchy. On the other hand Δ_2^p -complete (=P^{NP}-complete) problems are “mildly” harder than NP-complete ones, since they can be solved by means of a polynomial number of calls to an algorithm for satisfiability testing.

As we shall see in the sequel, many problems in propositional non-monotonic reasoning are Σ_2^p -complete or Π_2^p -complete, and many algorithms designed for solving such problems use satisfiability testers as subroutines.

The last class of decision problems we introduce are those polynomially solvable by a deterministic Turing machine with no more than $f(n)$ calls to a Σ_k^p oracle. Such class is denoted by $\text{P}^{\Sigma_k^p[f(n)]}$, where $f(n)$ is a polynomial function of the size n of the problem instance. In particular we will mention in the paper the class $\text{P}^{\text{NP}[\log(n)]}$. We notice that $\text{NP} \subseteq \text{P}^{\text{NP}[\log(n)]} \subseteq \Delta_2^p$ and that the containments are conjectured to be strict.

Throughout this paper we assume that all the above mentioned conjectures are true.

If two decision problems A and B are complete for the same class, then there is always a way to solve any instance of A by solving a single instance of B and viceversa. In Section 8 we will discuss in detail applications of this property to NMR problems.

Sometimes we refer to *search* problems, i. e. problems whose answer is more complex than just a boolean value. As an example, *finding* a satisfying truth assignment for a propositional formula is a search problem. An interesting class of search problems is $\text{F}\Delta_2^p$, which is the set of problems solvable in polynomial time by a machine with access to an oracle for an NP problem. If the oracle can be accessed only a logarithmic number of times, then we have the class $\text{FP}^{\text{NP}[\log(n)]}$.

While the polynomial hierarchy is an attempt at characterizing “how polynomially uncomputable” is a decidable function, the arithmetical and analytical hierarchies — which are precursors of the polynomial hierarchy — are characterizations of “how undecidable”

is a function. For the sake of brevity we do not give precise definitions but refer the reader to [131, 70]. We just recall that the elements $\Delta_k^0, \Sigma_k^0, \Pi_k^0$ of the arithmetical hierarchy and the classes $\Delta_k^1, \Sigma_k^1, \Pi_k^1$ of the analytical hierarchy are defined similarly to the elements of the polynomial hierarchy. While separations among the levels of the polynomial hierarchy is only conjectured, separation in both the the arithmetical and the analytical hierarchies has been formally proved. Moreover any element of the analytical hierarchy is harder than any element of the arithmetical hierarchy.

3 Default Logic

Default logic has been defined by Reiter in [129] and it is one of the more extensively studied non-monotonic formalisms. Interesting relations between default logic and logic programming have been shown by Bidoit and Froidevaux in [13]. They used default logic for defining a semantics for negation in logic programming.

In default logic the knowledge about the world is divided into two parts, representing certain knowledge and defeasible rules, respectively. The first part (denoted with W) is a set of closed first-order formulae, while the second one (denoted with D) is a collection of special inference rules called *defaults*. A default is a rule of the form

$$\frac{\alpha(\mathbf{x}) : \beta_1(\mathbf{x}), \dots, \beta_n(\mathbf{x})}{\gamma(\mathbf{x})}$$

where $\alpha(\mathbf{x}), \beta_1(\mathbf{x}), \dots, \beta_n(\mathbf{x}), \gamma(\mathbf{x})$ are wffs whose free variables are among those of $\mathbf{x} = x_1, \dots, x_m$. $\alpha(\mathbf{x})$ is called the *prerequisite* of the default, $\beta_1(\mathbf{x}), \dots, \beta_n(\mathbf{x})$ ($n \geq 0$) are called *justifications* and $\gamma(\mathbf{x})$ is the *consequence*. When $n = 0$, then the propositional constant **true** is implicitly assumed as the justification of the default. A default is *closed* if none of $\alpha, \beta_1, \dots, \beta_n, \gamma$ contains free variables. A default theory $\langle D, W \rangle$ is closed iff all the defaults in D are closed.

There are two special cases of default rules which are frequently used. Default rules of the form $\alpha(\mathbf{x}) : \beta(\mathbf{x})/\beta(\mathbf{x})$ and $\alpha(\mathbf{x}) : \beta(\mathbf{x}) \wedge \gamma(\mathbf{x})/\beta(\mathbf{x})$ are called *normal* and *seminormal*, respectively.

The semantics of a closed default theory $\langle D, W \rangle$ is based on the notion of *extension*, which is a possible state of the world according to the knowledge base. Formally, an extension is a fixed point of the operator Γ defined as follows. $\Gamma(A)$ is the smallest set such that:

1. $W \subseteq \Gamma(A)$;
2. $\Gamma(A) = \{\alpha \mid \Gamma(A) \models \alpha\}$;
3. if $(\alpha : \beta_1, \dots, \beta_n/\gamma) \in D$, $\alpha \in \Gamma(A)$ and $\forall_{j=1}^n \neg\beta_j \notin A$ then $\gamma \in \Gamma(A)$.

A set of formulae E is an extension of $\langle D, W \rangle$ iff $E = \Gamma(E)$. This definition is extended to *open* default theories by assuming that the defaults with free variables implicitly stand for the infinite set of closed defaults obtained by replacing the free variables with terms of the Herbrand Universe of the default theory.

The three computational problems which are most relevant in default logic, and that have been extensively studied in the literature, are: deciding whether a default theory $\langle D, W \rangle$ has an extension, deciding whether a formula α belongs to at least one extension of $\langle D, W \rangle$ (also known as *credulous* default reasoning) and deciding whether α belongs to all the extensions (*skeptical* default reasoning). Notice that these are decision problems.

Each extension E of $\langle D, W \rangle$ is identified by a subset D' of D , called the set of *generating defaults* of E , having the property that E is the deductive closure of

$$W \cup \{\gamma \mid (\alpha : \beta_1, \dots, \beta_n / \gamma) \in D'\}.$$

The set of generating defaults gives a compact representation of an extension of a closed default theory, which by definition is a deductively closed set of formulae, hence infinite. The *search* problem of finding a set of generating defaults has also been studied from the computational point of view.

We would like to explain precisely the sources of complexity of default logic, focusing on the propositional case. A simple algorithm for deciding whether a formula δ follows by credulous default reasoning from $\langle D, W \rangle$ is as follows:

```

for each subset  $D'$  of  $D$ ;
do begin
  (*  $D'$  is potentially a set of generating defaults *)
  if  $D'$  corresponds to an extension of  $\langle D, W \rangle$   (* 1st test *)
  then begin
     $W' := W \cup \{\gamma \mid (\alpha : \beta_1, \dots, \beta_n / \gamma) \in D'\}$ .
    if  $W' \models \delta$   (* 2nd test *)
    then return true
  end;
end;
return false.

```

Since there are $(2^{|D|})$ subsets of D , the body of the **for each** loop could be executed an exponential number of times. As far as the computational cost of the body of the loop is concerned, the first test — according to the definition of extension given previously — requires to perform several satisfiability and entailment checks for propositional formulae. Analogously, the second test requires an entailment check. As a consequence, each execution of the body of the loop requires most likely an exponential time.

If non-deterministic computations are considered, then the situation is different. As an example, it is possible to guess in parallel all the subsets D' of D , and this can be done in polynomial time by a non-deterministic machine. Clearly we want that such a machine to return **true** if *at least one* subset D' satisfies the property in the loop. As far as the body of the loop is concerned: We saw in Section 2 that both satisfiability and entailment tests can be answered by NP machines. In other words, credulous default reasoning can be done by a non-deterministic polynomial Turing machine which guesses in parallel all the possible subsets D' of D and, for each subset, uses the answers given by an *oracle*. The oracle is in charge of performing the satisfiability and entailment checks required by the body of the loop, which are computations doable in non-deterministic

polynomial time as well. It is clear from the definition of extension that for each subset D' only a polynomial number of calls to the oracle are needed.

The above argument can be made more formal and it shows that the problem of credulous default reasoning is in the class $\text{NP}^{\text{NP}} = \Sigma_2^p$ of the polynomial hierarchy. This establishes an upper bound, but it is of obvious practical interest to know whether (at least one of) the non-deterministic polynomial computations that we were referring previously to can be turned into deterministic polynomial ones, by means of smart search techniques.

Gottlob in [64] and, independently, Stillman in [151] give a negative answer to this problem, by showing that credulous reasoning is Σ_2^p -complete, i. e. the hardest problem among those of the class Σ_2^p .

This result can be interpreted by saying that the source of complexity of consistency checking and that of the choice of the generating defaults are independent, and their interaction gives a problem which is complete for the second level of the polynomial hierarchy. This fact formalizes, for the propositional case, the intuition that, if $\text{NP} \neq \Sigma_2^p$, default reasoning is computationally harder than monotonic reasoning, which is complete for the first level of the polynomial hierarchy.

It is interesting to see explicitly Gottlob's proof of Σ_2^p -hardness, in order to understand better the sources of complexity. The proof consists in a polynomial reduction from the problem of deciding the validity of a $QBF_{2,\exists}$ formula.

Let $Q = \exists p_1 \dots \exists p_n \forall q_1 \dots \forall q_m G$ be an arbitrary $QBF_{2,\exists}$ formula. Let $\langle D, \emptyset \rangle$ be a default theory, where

$$D = \{(: p_1/p_1), (: \neg p_1/\neg p_1), \dots, (: p_n/p_n), (: \neg p_n/\neg p_n)\}.$$

All the possible truth assignments to p_1, \dots, p_n are in one-to-one correspondence with the extensions of $\langle D, \emptyset \rangle$. Hence, Q is valid iff there exists at least one extension of $\langle D, \emptyset \rangle$ containing G , i. e. if G follows from $\langle D, \emptyset \rangle$ in a credulous way. This shows that validity of a $QBF_{2,\exists}$ formula can be reduced to credulous default reasoning by means of a polynomial many-one reduction.

We want to spend few words for commenting the above reduction. Extensions can be chosen in non-deterministic polynomial time (symbolized by the sequence of the existential quantifiers), but apparently not in deterministic polynomial time. An intuitive explanation of this fact is that defaults may interact in a combinatorial fashion, thus generating an exponential number of extensions. Moreover, even if a deterministic polynomial choice were possible, the complexity of the inference wrt the chosen extension should still be faced: it can be done in non-deterministic polynomial time (sequence of universal quantifiers), but apparently not in deterministic polynomial time.

The above proof uses normal defaults without prerequisites. As a consequence, restricting to normal default theories does not help reducing the complexity of credulous reasoning. From the algorithmic point of view, normal default theories have been considered more promising, since they support goal-directed algorithms for credulous reasoning (see [129]), while the algorithms proposed for non-normal theories by Etherington in [50] and Zhang and Marek in [158] use exponential amount of space to avoid non-termination.

Skeptical default reasoning in the prerequisite-free normal case has been shown to be Π_2^p -complete with a similar proof in [64] and [151] In the same works the problem of

showing the existence of an extension has been proven to be Σ_2^p -complete for semi-normal default theories (it becomes trivial for normal default theories which are guaranteed to have extensions). Papadimitriou and Sideri prove in [113] similar results for the last problem.

Several researchers studied restrictions of the expressiveness of default theories so that inference could be done in polynomial time. In particular restricted forms of the purely propositional part (the W), like Horn clausal form, 2CNF or 1CNF, have been considered (2CNF formulae are conjunction of clauses of the form $x \vee y$, where x, y are literals; 1CNF formulae are just conjunction of literals). Generally speaking, restricting the expressiveness of the W is not useful from the computational point of view, since the complexity of default reasoning can be completely “hidden” in the default part: the default theories $\langle D, W \rangle$ and $\langle D \cup \{(\ : \ /W)\}, \emptyset \rangle$ are equivalent. Therefore, the expressiveness of the default part must be limited as well, so that the non-deterministic choice of the subsets of D that are to form a set of generating defaults can be turned into a deterministic one. This can be done by imposing conditions on the syntactic form of the defaults in order to control their interaction.

Let us analyze one of such restrictions that have been proposed, the so-called *Horn* defaults. In Horn defaults the prerequisite is a conjunction of positive atoms, the justification and the consequence are the same literal, i. e. the default has the form

$$\alpha : y/y$$

Kautz and Selman in [76] prove that, as long as W is 1CNF and credulous reasoning is concerned, a set of Horn defaults D can be mapped in linear time into a set of Horn clauses H so that the inference of a single literal x from $\langle D, W \rangle$ is equivalent to the inference of x from H , hence doable in linear time in $|\langle D, W \rangle|$. In this case clearly the interaction among the defaults is not arbitrarily complex, and can be easily controlled. As proven by Stillman in [150], the hypothesis that W is 1CNF cannot be fully relaxed, since if W is in Horn clausal form, then credulous reasoning is NP-complete.

In [76] and [150] other examples of restrictions that lead to polynomially tractable cases are shown. Kautz and Selman in [76] focus on default theories in which W is 1CNF, finding results for all three decision problems and the search problem of default reasoning. It is interesting to notice that credulous and skeptical reasoning have different complexity: as an example skeptical reasoning of a single literal is co-NP-complete when W is 1CNF and D is Horn.

Finding a set of defaults generating an extension has still a different complexity, and it is polynomial for some class in which both skeptical and credulous reasoning are intractable. Kautz and Selman for example analyze the case in which W is 1CNF and D has only *disjunction free ordered* defaults, which are strictly more expressive than Horn defaults. Orderedness is a property of defaults defined by Etherington in [51], analogous to the idea of stratification in logic programming. A disjunction-free default has the form $a_1 \wedge \dots \wedge a_l : b_1 \wedge \dots \wedge b_m \wedge c_1 \wedge \dots \wedge c_n / b_1 \wedge \dots \wedge b_m$, where all a_i , b_j and c_k are literals. Finding an extension in this case has been proven to be polynomial in [76]. This result has been improved by Papadimitriou and Sideri in [113], who show that some form of cycles in the defaults can be allowed. In particular orderedness can be relaxed to a property that they call *evenness*, still having polynomiality.

It is clear that finding a set of defaults generating an extension is at least as hard as proving the existence of an extension. The last problem is shown in [76] to be NP-complete when W is 1CNF and D disjunction-free. This result is improved by Dimopoulos and Magirou in [38], in which NP-completeness is shown for disjunction-free and prerequisite-free defaults.

Kautz and Selman give also upper bounds for the complexity of default reasoning, showing that as long as W is 1CNF and D is disjunction-free (the most general class they take into account) credulous reasoning and proving the existence of an extension are in NP, skeptical reasoning is in co-NP and finding an extension is in $F\Delta_2^p$, i. e. it can be done by a polynomial machine which uses an NP oracle. These results confirm the intuitive complexity analysis that we gave before: the complexity of the consistency/entailment checking is always linear in 1CNF-disjunction free default theories, hence the only source of complexity is the interaction of defaults. As a consequence, the decision problems are at the first level of the polynomial hierarchy.

Stillman in [150] tries to enhance the expressiveness of W , by allowing either 2CNF or Horn clauses, with disjunction-free defaults. For the reasons given before, the upper bound of the four problems is not affected by this enhancement. He focuses on prerequisite-free defaults, showing for example that if the defaults are also normal and the W is 2CNF, then credulous reasoning is polynomial, while it is NP-complete if W is in Horn clausal form.

Bidoit and Froidevaux in [13] give results that complement Stillman's ones, by showing that propositional logic programs can be translated into default theories whose W is Horn. If the program is stratified, then the resulting default theory has a unique extension which can be computed in polynomial time. Further polynomial cases can be obtained using the similar translations by Marek and Truszczyński in [94].

Dimopoulos and Magirou in [38] obtain further results, by using a network representation of seminormal disjunction-free default theories. Using a graph-theoretic analysis they show that for theories without odd cycles in the associated graph, finding an extension can be done in polynomial time. This class generalizes the class of disjunction-free ordered theories, thus improving the previous results in [76].

Ben-Eliyahu and Dechter in [9, 10] propose a different technique for performing default reasoning: translating a default theory $\langle D, W \rangle$ into a propositional theory \mathcal{P} so that there is a one-to-one correspondence between the models of \mathcal{P} and the extensions of $\langle D, W \rangle$. As proven by the results of [64] and [151] a transformation of this kind cannot in general be done in polynomial time unless $\Sigma_2^p = \text{NP}$. Actually the size of \mathcal{P} , according to the transformation of [10], is exponential in the size of $\langle D, W \rangle$. Nevertheless, there are subcases in which the size of \mathcal{P} is polynomial wrt the size of the default theory. The most important such case shown in [10] concerns the class of *2-default theories*, in which W is 2CNF, the prerequisite of each default is in 2CNF, each justification is in 2DNF and the conclusion is a clause of size two. Using this reduction, Ben-Eliyahu and Dechter prove that the three decision problems are either NP-complete or co-NP-complete for this class of default theories. Furthermore, each subclass of 2-default theories that translates into a tractable subclass of propositional satisfiability is a tractable subset of default logic. Two such subclasses are shown in [10].

With respect to default theories in which first-order formulae occur, it is worth re-

membering that a default of the form $(: p(x)/p(x))$, where x is a variable, is actually an abbreviation for the infinite set of defaults $(: p(t)/p(t))$, where t is a term of the Herbrand universe of the default theory.

Apart from the results by Reiter in [129] who showed that credulous default reasoning is not r. e. , few results on the complexity of fully first-order default reasoning have appeared. Reiter proved also that, as long as closed normal default theories built on decidable fragments of first-order logic are concerned, decidability of credulous reasoning is guaranteed. When open defaults are considered this is no longer true: Baader and Hollunder showed in [7] that credulous default reasoning is undecidable for default theories built on a decidable fragment of first-order logic and containing a finite number of open defaults. Apt and Blair in [5] performed an analysis on complexity of negation in logic programs, giving some results on default logic by using the existing correspondences between the two formalisms. In particular they show that skeptical and credulous inference of ground atoms is, for all finite n , hard wrt the class Σ_n^0 of the arithmetic hierarchy, hence strictly harder than first-order inference. These results hold for finite first-order W and for infinite D , in the sense specified above. In Section 5 we give more details about the results by Apt and Blair.

Marek, Nerode and Rummel in [91] give recursion-theoretic results on a non-monotonic formalism that generalizes default logic.

We now briefly mention results on variants of default logic. The complexity of model-preference default logic introduced by Kautz and Selman in [142], has been analyzed by the same authors, Papadimitriou [112] and Cadoli [24]. The complexity of another variant, known as default logic with stationary extensions introduced by Przymusinska and Przymusinsky [120], has been analyzed by Gottlob [63].

A computationally appealing alternative to default logic is based on the use of probabilistic semantics. For a general survey on the topic we refer the reader to Pearl's paper [116]. In this setting default rules implicitly define a preference ordering (ranking) among sentences that is used to infer conclusions. This different reading has a stronger causal content than classical default logic. Computationally, this results in a more tractable system. In fact, the problem of deciding whether a formula follows from a conditional knowledge base is in Δ_2^p . This is due to the simpler interaction between defaults which can be accommodated using only a polynomial number of calls to a satisfiability tester. As a consequence, restricting the attention to base languages where satisfiability can be checked in polynomial time (e. g. Horn clauses) delivers polynomial deduction tasks (see the works of Goldszmidt and Pearl [62]).

We would like to conclude this section by summarizing the results previously presented. Propositional default reasoning is in some formal sense strictly harder than classical propositional reasoning. Apparently it has two independent sources of complexity (selection of an appropriate set of generating defaults and consistency/entailment checks) whose interaction leads to problems complete for the second level of the polynomial hierarchy. Many successful attempts at finding polynomially tractable subcases have been proposed, and all of them attack both sources of complexity at the same time.

Algorithms for performing default reasoning have been presented by Reiter [129], Etherington [50], Zhang and Marek [158], Schwind [141] and Poole et al. [117]. Algorithms described in [129, 50, 158] use satisfiability testers as subroutines.

4 Modal Non-Monotonic Logics

In this section we review complexity results presented in the literature concerning non-monotonic versions of modal logics, with special attention to the best known of such logics, autoepistemic logic with all its variants. We consider a propositional language augmented by the modal operator L .

Historically, the first non-monotonic modal logic has been the Nonmonotonic Logic (NM1 in the sequel) introduced by McDermott and Doyle [101, 100]. This logic has been refined by Moore in [104] who defined autoepistemic logic. In autoepistemic logic the semantics of a set of modal formulae A (called *premises*) is given by *stable expansions* which are fixpoints of the following equation:

$$\Delta = \{\alpha \mid A \cup L\Delta \cup \neg L\bar{\Delta} \models \alpha\}$$

where $L\Delta = \{L\beta \mid \beta \in \Delta\}$ and $\neg L\bar{\Delta} = \{\neg L\beta \mid \beta \notin \Delta\}$.

Stable expansions in autoepistemic logic have the same role played by extensions in default logic. Analogously to default logic, the main reasoning tasks are deciding if a set of premises has a stable expansion, if a formula α belongs to at least one stable expansion of a set of premises (credulous autoepistemic reasoning) and if a formula belongs to all stable expansions of a set of premises (skeptical autoepistemic reasoning).

There is a close correlation between the interpretation of $\neg L\alpha$ in autoepistemic logic and the meaning of *not* α in logic programming, where *not* denotes negation as failure. This analogy has been stressed and used by Gelfond in [56] to define the iterative fixed-point semantics and by Gelfond and Lifschitz in [57] to define the stable model semantics for logic programs with negation.

An issue of obvious computational interest concerns the compact representation of stable expansions, which is analogous to the issue of a compact representation of an extension of a default theory. As shown by Konolige in [79] any set of premises A can be transformed into an equivalent one A' which is composed of formulae f_i of the kind:

$$L\alpha_i \wedge \neg L\beta_i^1 \wedge \dots \wedge \neg L\beta_i^n \rightarrow \gamma_i$$

where $\alpha_i, \beta_i^1, \dots, \beta_i^n, \gamma_i$ are formulae not containing the modal operator L . When formulae are in this syntactic form then subsets of A' can play the same role that generating defaults have in default logic. Unfortunately, the process of transforming arbitrary formulae into this form may exponentially increase their size. For this reason Niemelä in [108] and Shvarts in [147] have proposed methods which extract directly from A *kernels*, i. e. sets of subformulae of A which are in one-to-one relation with stable expansions.

The complexity results for the general case are the following. The problem of deciding whether a set of premises A has a stable expansion is Σ_2^p -complete, credulous autoepistemic reasoning is also Σ_2^p -complete, while skeptical autoepistemic reasoning is Π_2^p -complete. The upper bounds have been proven by Niemelä in [108] and the lower bounds by Gottlob in [64].

Notice that these results are exactly the same already obtained for default logic. Since all problems complete for the same class of the polynomial hierarchy can be polynomially reduced one to the other, it is clear that, for example, credulous default reasoning

is polynomially reducible to credulous autoepistemic reasoning. More precisely, the existence of a polynomial mapping \mathcal{M} from the problem of deciding credulous inference of a formula γ from a default theory $\langle D, W \rangle$ into the problem of credulous inference of another formula δ from a set of premises A is guaranteed. The existence of an inverse mapping \mathcal{M}' is guaranteed as well.

This may seem in contradiction with the negative results obtained by Konolige in [79], and Marek and Truszczyński in [95] where they could not find an embedding of default logic in classical autoepistemic logic, but only in some of its variants. This contradiction is only apparent, in fact they were interested in translations satisfying some condition of locality, in which a single default rule is mapped into a modal formula. Transformation like $\mathcal{M}, \mathcal{M}'$ do not necessarily satisfy locality. A more detailed analysis of the various translations proposed for different NMR formalisms is presented in Section 8.

Like in the case of default reasoning, restricting the expressiveness of the language may lead to a decrease of the complexity. For example, Marek and Truszczyński have shown in [97] that deciding whether an atom is member of one expansion of a set of premises of the following form

$$La_1 \wedge \dots \wedge La_n \wedge \neg Lb_1 \wedge \dots \wedge \neg Lb_m \rightarrow c$$

where $a_1, \dots, a_n, b_1, \dots, b_m, c$ are literals, is NP-complete while membership to all expansions is co-NP-complete. In a very recent paper Niemelä and Rintanen [109] show that restricting the attention to stratified autoepistemic theories, a class defined by Marek and Truszczyński in [97] in the style of stratified logic programs, it is possible to greatly reduce the computational complexity of all the decision problems. In particular, they show that for stratified autoepistemic theories the only source of complexity is given by the satisfiability checking in the underlying propositional language, hence, when this check can be done in polynomial time (e. g. Horn clauses) then all the decision problems are polynomially tractable.

Although there are no other direct results on the complexity of autoepistemic inference, other complexity results can be obtained by means of the techniques described in Section 8, using the translations from default logic [65, 87] and logic programs [57, 93] into autoepistemic logic.

In the rest of this section we deal with the several variants of autoepistemic logic that have been proposed in the literature.

A more complex variant is the *moderately grounded* autoepistemic logic introduced by Konolige in [79]. A stable expansion T is moderately grounded if and only if there is no other stable expansion T' such that the part of T' with no occurrence of L is strictly contained in that of T . This minimality requirement for T can be checked by non-deterministically selecting a T' and checking whether the condition holds, but this obviously adds one further level of non-determinism to the computation. Eiter and Gottlob in [46] show that it cannot be eliminated by proving that credulous reasoning wrt moderately grounded expansions is Σ_3^p -complete and skeptical reasoning Π_3^p -complete.

Autoepistemic logic is a refinement of the non-monotonic logic NM1 introduced by McDermott and Doyle [101, 100]. NM1 is based on a fixpoint construction slightly different from the one used in autoepistemic logic. Given the analogies of NM1 with

autoepistemic logic, it is not surprising that the decision problems of NM1 have the same complexity of the corresponding problems of autoepistemic logic, see Niemelä [108] (upper-bounds) and Gottlob [64] (lower-bounds).

Using the fixpoint construction of McDermott and Doyle but changing the underlying modal system Marek and Truszczyński in [95] and Marek, Shvarts and Truszczyński in [92] present a whole new family of non-monotonic modal systems. The complexity analysis of these systems has shown that for most of them all the decision problems lie at the second level of the polynomial hierarchy (Σ_2^P or Π_2^P -complete) even when underlying modal systems of different complexity are used. A result which is worth mentioning is due to Schwarz and Truszczyński in [140] where they show that skeptical reasoning in the non-monotonic modal logic *S4* is Π_2^P -complete. This is somehow surprising since reasoning in the **monotonic** modal system *S4* is PSPACE-complete as proven by Ladner in [82]. Marek, Shvarts and Truszczyński in [92] and Gottlob in [64] give other interesting complexity results.

As far as first-order languages are concerned, we want to point out that in the original paper by Moore [104] only propositional languages are allowed. By the way, Konolige in [79] allows the modal operator to be applied to closed formulae of first-order logic. Under this restriction, Niemelä proves in [108] that whenever the satisfiability test in the underlying first-order language is decidable so are all the reasoning tasks of autoepistemic logic.

5 Negation in Logic Programming

In this section we survey the complexity results for the various semantics for negation in logic programming. This is not intended as an overview of such semantics and their relative merits, a detailed and interesting discussion on this issue can be found in the survey paper by Bidoit [12] and in the paper by Schlipf [138]. We will only marginally refer to the huge body of literature on complexity in rule-based database languages (DATALOG), since this survey focuses on the logic programming semantics more directly related to NMR. For a survey on logic programming and databases we refer the reader to the book by Ceri, Gottlob and Tanca [30].

We survey the perfect model, stable model, default model, well-founded, supported model, positivistic, inflationary semantics and the program completion. The expressive power of the various semantics has been studied and compared extensively. For an analysis of this issue see, for example, the work of Schlipf [137]. All these semantics are extensions of the classical minimal model semantics of positive logic programs and they all agree on the class of stratified and locally stratified logic programs, exception made for the supported and inflationary semantics and the program completion.

Due to lack of space, this survey will necessarily be incomplete. For a more complete survey on complexity and undecidability results on logic programming we recommend the excellent paper of Schlipf [139].

In the field of logic programming the most important inference task is deciding whether a ground atom belongs to the (usually unique) preferred model of a program. This is the reasoning task we will be referring to in the sequel, where not otherwise specified.

First of all let us recall the main results on the complexity of logic programs without negation. In this case, the preferred model is the unique least Herbrand model. For propositional languages reasoning under such a model can be done in linear time using Dowling and Gallier’s algorithm [40]. When dealing with first-order languages, membership of positive and negative literals in the least Herbrand model of a Horn program has been shown r. e.-complete and co-r. e.-complete, respectively, by Smullyan in [148] and by Andreka and Nemeti in [2].

For a first-order language, when negation is allowed there are two different membership problems depending on which sets of models we are interested in. Traditionally, the semantics of logic programs is given in terms of Herbrand models, however, in some works (e. g. Kunen’s paper [81]) it is assumed that models are interpretations over a larger universe containing an infinite number of constants as well as function symbols. As already pointed out by Blair in [16] the second alternative leads, in general, to computationally simpler problems. Where not otherwise indicated, we implicitly deal with Herbrand models.

We now analyze the complexity of reasoning tasks in stratified programs. From the computational point of view, these results apply to all the semantics, again with the exception of supported semantics, inflationary semantics and program completion. In this case it is well known that the unique preferred model of a stratified program can be computed incrementally through the strata by first assigning truth values to the atoms in the lower stratum and then moving on to the next ones [6]. This holds for locally stratified programs as well. This procedure converges to the iterated least model as defined by Apt, Blair and Walker in [6]. For a propositional language, this immediately implies tractability of the computation, since the computations performed at each step are polynomial and the number of strata is at most linear in the size of the program. Notice also that deciding whether a first-order program is stratified can be done in polynomial time [6], while deciding whether a program is locally stratified is undecidable, as proven by Cholak in [33].

The perfect model semantics, introduced by Przymusiński in [121], uses the syntactic form of the program to infer a preferential relation on the models. The semantics is not defined for all programs, but for a strict superset of the class of locally stratified programs.

Apt and Blair in [5] have proven that if P is a stratified program with n strata, then deciding membership in the perfect model of P is Σ_n^0 . Furthermore, for each $n \geq 1$ there is a stratified program P with n strata for which the same problem is Σ_n^0 -complete. This result can be specialized to the case of recursion-free programs, in which the problem is r. e. The more general class of locally stratified programs turns out to be computationally more complex. In fact, membership is Δ_1^1 -complete over ω as shown by Cholak and Blair [32] and Blair, Marek and Schlipf [17]. A further class of programs that has been analyzed is the class of acyclic programs. For this class of programs Apt and Bezem have proven in [4] that deciding membership of ground atoms in the perfect model is a decidable task. The class of acyclic programs is a proper subset of the class of locally stratified ones, and checking acyclicity of a program is a Π_2^0 -complete problem.

Using the translations from default logic and circumscription given by Bidoit and Froidevaux in [13] and Lifschitz in [86], Apt and Blair in [5] show that Σ_n^0 -hardness is a lower bound for ground inference in restricted classes of default and circumscriptive

theories.

For general propositional programs, Eiter and Gottlob prove in [42] that deciding whether a perfect model exists is co-NP-hard while deciding whether a model of a program is perfect is co-NP-complete.

The stable model and default model semantics have been independently introduced by Gelfond and Lifschitz in [57] and Bidoit and Froidevaux in [13]. While the first one is based on autoepistemic logic, the second one is based on default logic. As it turned out, these two semantics are equivalent. Since for stratified programs they also coincide with the perfect model semantics, we only consider general programs. Three problems have been analyzed in the propositional case: deciding the existence of a stable model has been proven NP-complete by Marek and Truszczyński in [96] and independently by Bidoit and Froidevaux in [14]. In [97] Marek and Truszczyński prove that deciding whether an atom is a member of one stable model is NP-complete and membership in all stable models is co-NP-complete.

The supported model semantics has been introduced by Apt, Blair and Walker in [6] to declaratively capture the behavior of PROLOG interpreters. Marek and Truszczyński have shown in [96, 97] that deciding the existence of a supported model is an NP-complete problem, membership of an atom in one supported model is also NP-complete and membership in all supported models is co-NP-complete. A variant of the supported model semantics is the positivistic model semantics introduced by Bidoit and Hull in [15]. Schaerf shows in [134] that deciding whether a formula is true in one positivistic model is a Σ_2^p -complete problem, while deciding whether a formula is true in all the positivistic models is Π_2^p -complete.

Differently from the previous semantics, the well-founded semantics, introduced by Van Gelder, Ross and Schlipf in [154], is based on three-valued or partial models. Partial models do not necessarily assign the values *true* or *false* to all atoms, but may leave someone undefined. Well-founded semantics always provides a unique preferred model (called well-founded model) to any program, furthermore, if the program is stratified the well-founded model is completely defined (not partial). The well-founded model of any propositional program can be computed in polynomial time [154].

When a first-order language is used, the computational differences between stable and well-founded model semantics disappear. In fact the problem of deciding membership in all stable (well-founded) models turns out to be Π_1^1 -complete over ω for both Herbrand models and models over larger universes, as shown by Schlipf in [137, 139] and van Gelder [153]. Two other problems have been considered by Marek, Nerode and Remmel. In [91] they show that for the class of programs admitting a unique stable model the membership problem is still Π_1^1 -complete over ω , while in [90] they prove that deciding whether a program has a stable model is a Σ_1^1 -complete problem over ω . Furthermore, Marek and Subrahmanian show in [93] that deciding whether a model is a stable model of a program P is Π_2^0 -hard and is in Π_3^0 .

As far as the predicate completion defined by Clark [35] is concerned, for propositional languages it has been proven by Kolaitis and Papadimitriou in [78] that deciding whether a ground atom belongs to the completion of a program is a co-NP-complete problem. If quantification and function symbols are introduced, then the membership problem over Herbrand models is Π_1^1 -complete over ω and over models with larger universes is r.e.-complete.

We do not have the space to refer all the complexity results in the DATALOG area, but we still want to mention the works of Kolaitis and Papadimitriou [78] and Papadimitriou and Yannakakis [114]. In their analysis Kolaitis and Papadimitriou take into account DATALOG[¬] programs, that is DATALOG programs with negation and with a finite universe, and give two semantics for it. The first one is given as the fixpoint of a natural consequence operator. The data-complexity, i. e. complexity wrt a fixed extensional database, of finding a fixpoint is NP-complete and finding a least fixpoint is in Δ_2^p (for a more precise characterization see [78]). In order to overcome these computational limitations Kolaitis and Papadimitriou introduce the inflationary semantics, which is based on an iterative construction rather than a fixpoint one. Inflationary semantics, as the well-founded semantics, is defined for all programs and its data-complexity is polynomial.

This analysis has been further developed by Papadimitriou and Yannakakis. The main issue addressed in the paper is to find a semantics whose data complexity is polynomial and which extends the well-founded semantics. The proposed semantics, the tie-breaking semantics is in some provable sense more general than the well-founded semantics and can be computed in polynomial time. In the same paper they also show that proving whether a fixpoint model exists for any extension of the EDB is undecidable.

An extension of logic programming which has been well analyzed from the computational point of view is disjunctive logic programming, where disjunctions of literals are allowed in the head of the rules. For a comprehensive survey of the semantical and computational aspects of disjunctive logic programming we refer to the book by Lobo, Minker and Rajasekar [88]. Due to lack of space we can only briefly mention the major results and give pointers to the literature. Most of the semantics for negation in logic programs have been extended to deal with disjunctive logic programs. In particular, Przymusiński in [124] introduces the disjunctive stable model semantics, which extends the stable model semantics, and the partial disjunctive stable model semantics which extends the well-founded semantics. Eiter and Gottlob in [42] have analyzed the complexity of these and other semantics. In particular they show that deciding whether a literal is true in all the preferred models of a disjunctive logic program under the perfect model semantics and the (partial) disjunctive stable model semantics is a Π_2^p -complete problem. A different extension of the stable model semantics has been given by Gelfond and Lifschitz in [58] through the notion of answer set. While inference has been shown to be a problem complete for the second level of the polynomial hierarchy, see Eiter and Gottlob [44], nevertheless some classes have been shown to be computationally simpler by Ben-Eliyahu and Dechter in [11], where they show a polynomial mapping from a subclass of disjunctive logic programs into a propositional theory. Another semantics, called disjunctive database rule has been defined by Ross and Topor in [132] and it has been shown to be polynomially tractable by Chan [31]. Chan has also extended this rule to the possible worlds semantics which correctly handles negative clauses. However, literal inference under this new semantics is co-NP-complete [31, 42].

6 Circumscription and Closed-World Reasoning

Circumscription has been defined by McCarthy in [98] and refined later by McCarthy himself in [99] and Lifschitz in [85]. We refer to the definition which can be found in [99, 85]. Let $T[P; Z]$ be a first-order formula in which at least the predicates of the disjoint lists of predicates P and Z occur. The *circumscription* $CIRC(T; P; Z)$ of the predicates P in $T[P; Z]$ with varying predicates Z is the following second-order formula

$$T[P; Z] \wedge (\forall P'Z'. T[P'; Z'] \rightarrow \neg(P' < P)) \quad (1)$$

where $P'Z'$ is a list of predicates isomorphic to PZ and $T[P'; Z']$ is T with all the occurrences of predicates of PZ substituted by the corresponding ones in $P'Z'$. The meaning of $P' < P$ is defined in terms of the relation \leq . In particular $P' < P$ is $(P' \leq P) \wedge \neg(P \leq P')$, and $P' \leq P$ stands for the conjunction of the formulae

$$(\forall \vec{x}. p'_i(\vec{x}) \rightarrow p_i(\vec{x}))$$

for each p_i in P . In the above formula p_i, p'_i are corresponding predicates, and \vec{x} is an appropriate list of variables. The predicates occurring in P are called *minimized* while those not occurring in PZ are denoted with Q and are called *fixed*.

The closed-world assumption (CWA) of a first-order formula T is the following formula, defined by Reiter in [128]:

$$CWA(T) = T \cup \{\neg p \mid T \not\models p\} \quad (2)$$

where p is a ground atom. This rule has been refined by several authors: Minker [102] introduced the *generalized* CWA, Rajasekar, Lobo and Minker [126] the *weak generalized* CWA, Yahya and Henschen [156] the *extended generalized* CWA, Gelfond and Przymusinska [59] the *careful* CWA, Gelfond, Przymusinski and Przymusinska [60] the *extended* CWA and the *iterated* CWA. The notion of fixed and varying predicates has been used in the careful, in the extended and in the iterated CWA. Extended CWA is the most general of all the above rules and is defined as follows:

$$ECWA(T; P; Q; Z) = T \cup \{\neg K \mid \exists B. (T \not\models B) \wedge (T \models K \vee B)\} \quad (3)$$

where $\langle P; Q; Z \rangle$ is a partition of the predicate symbols of T in minimized/fixed/varying ones, K is any formula not involving letters from Z and B is a disjunction of ground literals whose predicate symbols come from Q and positive ground literals whose predicate symbols come from P . The formulae K whose negations are added to T in the above formula are called *free for negation*.

Relationships between circumscription, closed-world reasoning and Clark's [35] negation as failure rule have been very well studied in the literature by many researchers, see for example Reiter [130], Sheperdson [146], Lifschitz [84], and in all the works defining closed-world rules. In particular it has been shown that an abstract notion of minimality underlies definitions (1), (2) and (3). In the propositional case, formulae (1) and (3) are equivalent (see [60]).

6.1 Inference

Several computational tasks for circumscription and closed-world reasoning have been addressed from the point of view of the complexity. Inference, i. e. to decide whether a first-order formula logically follows from formulae (1) or (3), is the first one we survey.

Circumscription is a second-order formula, and second-order logic is well-known to be computationally strictly harder than first-order logic (see for example van Benthem and Doets in [152]). In fact circumscription appears to be more expressive than first-order logic as noticed by Lifschitz. In [84] he shows that second-order formulae like transitive closure, which are not collapsible to first-order ones, are expressible as the circumscription of a first order formula.

This fact doesn't *per se* show that circumscription is strictly harder than first-order logic. Nevertheless there have been many attempts at finding classes of first-order formulae whose circumscription is (equivalent to) a first-order formula. From the computational point of view, results of this kind show classes of circumscription in which inference is a r. e. problem.

One of these classes (called of *separable* formulae) has been defined by Lifschitz in [85]. The class of separable formulae is based on a generalization of the idea of predicate completion by Clark [35] and contains all quantifier-free formulae. This result has been improved by Rabinov in [125], who showed a class of collapsible formulae that subsumes separable formulae. In both cases collapsibility is guaranteed only when there are no varying predicates, i. e. $Z = \emptyset$.

Kolaitis and Papadimitriou [77] investigated the fragment of existentially quantified formulae, showing that their circumscription is first-order even if varying predicates are allowed.

All the conditions that imply collapsibility are sufficient and not necessary, and the complexity of deciding whether a given formula is in one of the above collapsible cases is in general not known. Actually, the task of proving collapsibility has been proven to be undecidable by Krishnaprasad [80] for general first-order formulae, and by Kolaitis and Papadimitriou for logic programs [77].

All the above results seem to suggest that circumscriptive inference is harder than first-order inference. Actually, Cadoli, Eiter and Gottlob proved in [25] that inference of a formula γ from a circumscription is as hard as testing validity of any formula in second-order logic.

It is interesting to notice that the last result holds even when γ is propositional, regardless of the presence of varying and/or fixed predicates. The fact that fixed predicates do not contribute to the complexity of circumscription was proved by de Kleer and Konolige in [37]. On the other hand the fact that varying predicates don't affect the complexity may seem surprising, in view of the fact that the algorithms for circumscription that have been proposed (for example Przymusiński's in [122]) are more efficient when varying predicates are not allowed. In the following we will see that the presence of fixed/varying predicates may affect the complexity of inference for some restricted class of formulae.

Other interesting results have been obtained by Schlipf in [135], who focused on countably infinite models of the formula (1). In particular he shows that, when inference

is restricted to such models, then the problem of deciding whether a first-order formula follows from (1) is complete for the class Π_2^1 of the analytical hierarchy over the integers. It is interesting to notice that this result holds even if $Z = \emptyset$.

In [5] Apt and Blair showed a lower bound of Σ_n^0 -hardness for any finite n when ground inference in stratified logic programs is concerned. They take into account a slightly more general version of the formula (1) called *prioritized* circumscription defined by Lifschitz in [85].

As far as the propositional version of (1) is concerned, Eiter and Gottlob showed in [41] that the inference problem (with $Q = \emptyset$, $Z = \emptyset$ or not) is complete for the class Π_2^p of the polynomial hierarchy. In [42] they give a more tight result, showing that Π_2^p -completeness holds even if formulae T with no negative clauses are considered. Comparing this result to those in Section 3, we can say that the minimality requirement imposed by circumscription gives a further level of non-determinism which is analogous to the choice of the right set of defaults. As a consequence a propositional circumscriptive reasoner and a propositional skeptical default reasoner have exactly the same computational power. In Section 8 we see that many researchers attempted at finding analogous relations between the two formalisms.

Papalaskari and Weinstein [115] analyzed the infinitary propositional (sentential) case of circumscription. They showed that when the underlying propositional language is countable and there are neither fixed nor varying predicates, then inference is a problem in Π_2^0 and not in Σ_2^0 . It is worth recalling that inference in sentential logic is r. e.

Cadoli and Lenzerini in [26] studied the complexity of circumscriptive inference for eight classes of (finite) propositional formulae in which monotonic inference is polynomial. In such a case the upper bound for inference is co-NP.

When the inferred formula γ is a clause, then the problem is co-NP-complete for very restricted cases, for example when T is both Horn and 2CNF or when T is 2CNF and only positive literals occur in it. The only polynomial case found is when T is Horn, 2CNF and with no clauses of the form $\neg x \vee \neg y$. Further polynomial cases have been found for restricted versions of the inference problem, for example when T is Horn and $Q = \emptyset$, or when T is 2CNF, γ is a literal and $Z = \emptyset$.

It is interesting to notice that, while for the general propositional case clause inference is not any harder than literal inference, and fixed/varying letters propositional letters do not add complexity to inference, this is not the case for restricted cases. As an example, Cadoli and Lenzerini show several cases in which literal inference is polynomial and clause inference is co-NP-complete. A similar tradeoff exists when either $Q = \text{emptyset}$ or $Z = \emptyset$.

One source of complexity of circumscriptive inference seems to be the number of free for negation formulae in (3), which is potentially exponential. Actually, in all the polynomial cases found in [26], it has been shown that only a polynomial number of free for negation formulae has to be taken into account.

Given the equivalence of circumscription and extended closed world reasoning in the propositional case, the same results hold for the extended CWA as well. For what concerns other forms of closed-world inference, Cadoli and Lenzerini show in the same work that many intractable cases for the extended CWA are tractable for other definitions. As an example, careful closed-world reasoning (defined in [59]) is polynomial for 2CNF

formulae. This can be explained by observing that the definition of careful CWA is (3) with K restricted to be a single positive literal. As a consequence the space of the potential free for negation formulae is linear in the size of the input. In view of this fact, it is interesting to notice that Eiter and Gottlob proved in [41] that careful closed-world inference is Π_2^p -hard, while the best upper bound is $P^{\Sigma_2^p[O(\log n)]}$, which is worse than the upper bound of Π_2^p that holds for extended closed-world inference.

Other authors have studied the complexity of closed-world inference. Apt and Blair have proven in [5] that reasoning under the closed world assumption (2) is Π_1^0 -complete when T is a logic program. Apt in [3] had previously shown that the task is polynomial for propositional Horn formulae. The same bounds hold for the weak generalized CWA defined in [126].

Weak generalized CWA has also been studied by Chan in [31], who proved that literal inference is polynomial for propositional formulae with no negative clauses. This has also been noticed by Cadoli and Lenzerini in [26]. If negative clauses are allowed, then co-NP-hardness holds (see [31]).

As far as the inference under the CWA in an unrestricted propositional language is concerned, Eiter and Gottlob in [41] show a lower bound of NP-hardness and an upper bound of $P^{NP[O(\log n)]}$.

In a more recent paper [42] Eiter and Gottlob analyzed iterated CWA, which is a variant of extended CWA defined in [60], showing Π_2^p -completeness results analogous to those obtained for extended CWA.

Chomicki and Subrahmanian analyzed in [34] the complexity of generalized closed-world inference for logic programs. The generalized CWA (see [102]) is the rule (3) when $Q = Z = \emptyset$ and K is a positive ground literal. The main result of [34] is that inference of a ground atom under the generalized CWA is complete for the class Π_2^0 of the arithmetic hierarchy.

Borgida and Etherington in [18] have shown a very specialized language for representing taxonomies in which some form of generalized closed-world reasoning is polynomial.

The complexity of closed-world reasoning in relational databases has been studied by many authors, like for example Vardi in [155]. The problems addressed are outside the scope of this survey, and we refer to the book by Grahne [67] for a general overview.

Several algorithms for inferencing under circumscription have been proposed: Przymusiński [122], Ginsberg [61], Inoue and Helft [73], Bossu and Siegel [19], Olivetti [110], Bell et al. [8, 107]. Each of them presupposes languages of limited expressiveness. The algorithms in [122, 73, 110] use satisfiability testers as subroutines.

6.2 Satisfiability

Satisfiability is another computational task that has been addressed. It is well-known that both the circumscription of a first-order formula T and its closure (2) could be unsatisfiable even if T is satisfiable. On the other hand the extended closure preserves consistency.

Schlipf proved in [135] that deciding whether the formula (1) has a countably infinite model is Σ_2^1 -complete over the integers.

Eiter and Gottlob in [41] proved a lower bound of NP-hardness and an upper bound of $P^{NP^{O(\log n)}}$ for consistency checking of the CWA of an arbitrary propositional formula.

These results complement some other by Schlipf, who proved in [136] that consistency checking of a slightly different version of the CWA is both NP-hard and co-NP-hard and is in Δ_2^p .

6.3 Model Checking

Kolaitis and Papadimitriou in [77] noticed that the complexity of circumscription seems to arise even in the problem of *model checking*. Model checking is the problem of deciding whether a finite structure satisfies a given formula Ψ . If Ψ is first-order, then the task is polynomial. On the other hand in [77] it is shown that when Ψ is the formula (1) then the task is co-NP-complete. The result holds when T is a universal-existential first-order sentence and $Z \neq \emptyset$. co-NP is actually the upper bound for the model checking of the circumscription of all first-order formulae T .

Cadoli in [23] strengthens this result in several directions, by proving co-NP-hardness of model checking when the formula T to be circumscribed is propositional, satisfiable, and $Q = Z = \emptyset$. co-NP-hardness in the case $Z = \emptyset$ is implicit in the work by Schlipf [136].

In the same work [23] Cadoli shows several polynomial subcases, for example when T is a Horn or 2CNF propositional formula. The complexity analysis has been carried out by examining subcases in which Q and/or Z are the empty set. In some interesting cases the task is polynomial if $Q = \emptyset$ and co-NP-complete otherwise. Some polynomial cases are also shown by Eiter and Gottlob in [41].

It is interesting to notice that propositional circumscription is in co-NP if model checking is polynomial (see [26]).

Polynomiality of the task when T is a first-order logic program has been shown by Kolaitis and Papadimitriou in [77].

6.4 Model Finding

Papadimitriou in [112] addressed the issue of finding a satisfying truth assignment for the formula (1) when T is propositional. Notice that model finding is a search problem.

The upper bound of the analogous task for plain propositional logic is $F\Delta_2^p$, hence a polynomial number of calls to an NP oracle is sufficient. Gottlob and Fermüller prove in [66] that unless $P=NP$ a logarithmic number of calls to an oracle in NP is necessary.

Papadimitriou proves that, if *minimal* satisfying truth assignments are searched, then a polynomial number of calls is still sufficient.

Cadoli in [24] addresses the case in which $Q = Z = \emptyset$, proving that the task is $FP^{NP^{O(\log n)}}$ -hard, even if T is in CNF and a satisfying truth assignment of T is known.

7 Abduction

Abduction is the process of finding explanations for observations in a given theory. This abstract characterization has led to several different definitions of abduction in the

literature, which can be roughly divided into two main areas: logic-based abduction as defined by Selman and Levesque in [144] and set-covering methods as defined by Reggia, Nau and Wang in [127]. Although the two methods look formally different, strong connections have been pointed out by Bylander in [21] and by Friedrich, Gottlob and Nejdl in [53]. In this paper we focus on logic-based abduction, since it has been shown to be tightly related to negation in logic programming (see for example Eshghi and Kowalski [49] and Kakas and Mancarella [75]) and default logic (see [144]). Furthermore, we only consider propositional languages, since very little is known about first-order abduction.

Logic-based abduction has been defined as follows in [144]: Given a set of propositional clauses Σ and a letter q , an *explanation* for q is a set of literals α such that (1) $\Sigma \cup \alpha \models q$, (2) $\Sigma \cup \alpha$ is consistent and (3) α is minimal wrt set inclusion. Notice that a trivial explanation satisfying the above requirements is q itself. All explanations not containing q will be called *non trivial*. A condition sometimes imposed on explanations is that the set α must be a subset of a predefined set of literals A which represent the plausible hypotheses. Explanations satisfying this condition are called *assumption-based* explanations in [144].

There are at least three decision tasks which have been considered in the various works on abduction: deciding whether an explanation exists at all, deciding whether an individual hypothesis $h \in A$ belongs to at least one acceptable explanation (credulous abductive reasoning) and deciding whether an individual hypothesis $h \in A$ belongs to all the acceptable explanations (skeptical abductive reasoning). The search problem of finding one (best) explanation has also been studied.

The most comprehensive work on the complexity of logic-based abduction has been done by Eiter and Gottlob [47]. In this paper, they analyze the complexity of all the above problems under different preference criteria. In particular they show that when no restrictions are imposed on the syntactic form of the theory Σ and no preference criteria is specified, then all the problems are complete for the second level of the polynomial hierarchy. Adding the requirement of minimality wrt set inclusion does not change the complexity, while using other preference criteria such as minimum cardinality and minimum cardinality with priorities the problems are Δ_3^p -complete. Finally, for minimality wrt set containment with priorities and in the general case of an arbitrary (polynomial) ordering of the explanations, the problems are complete for the third level of the polynomial hierarchy.

This complexity can be reduced of exactly one level if we restrict to theories Σ composed of Horn clauses. A further reduction is sometimes obtained if we restrict to definite Horn clauses.

Under the Horn clause restriction, Selman and Levesque have proven in [144] that a non-trivial explanation can be found in polynomial time. They also prove that finding an assumption-based explanation is instead an NP-hard problem, thus showing that the restriction to a subset of the literals may affect the complexity of inference. A similar phenomenon has been noticed for propositional circumscription with and without varying predicates (see Section 6). Bylander in [21] has focused on the more restricted case in which Σ is a conjunction of definite clauses, showing a polynomial method for finding assumption-based explanations. Another polynomial subcase has been found by Eshghi in [48].

As far as Horn clauses are concerned, Selman and Levesque showed that finding an

explanation, with the additional constraint that it has to contain a predefined letter q , is NP-hard.

Friedrich, Gottlob and Nejdil in [53] address problems of skeptical and credulous abductive reasoning in definite Horn clauses, proving that deciding whether a letter belongs to at least one assumption-based explanation is NP-complete, while the set of letters belonging to all the assumption-based explanations can be computed in polynomial time.

The complexity of enumeration problems has also been addressed for abduction. The number of explanations can clearly be exponential wrt the size of Σ , and actually Bylander has proven in [21] that determining the number of assumption-based explanations for definite clauses is complete wrt the class $\#P$ (see [74]). A related result has been obtained by Friedrich, Gottlob and Nejdil in [53].

Logic-based diagnosis has been proposed as a logical reconstruction of de Kleer's Assumption-based Truth Maintenance Systems [36]. Among the computational tasks that have been addressed in this area there is finding explanations, called *nogoods*, of the atom \perp (contradiction). Provan in [118, 119] and Rutenburg [133] address the issue of finding the size of a nogood.

In particular Rutenburg analyzes several definitions of TMS's, considering problems like deciding whether there exists a nogood with size less than a predefined number k . He shows problems whose complexity ranges from polynomial to Σ_2^P -complete. These results confirm the intuition that many computational tasks of non-monotonic reasoning hide two independent sources of complexity, whose interaction leads to problems complete for the second level of the polynomial hierarchy.

Other forms of abduction have been presented in the literature, but they are less related to NMR, and for this reason we will not survey them here. However it is worth mentioning the work of Bylander et al. in [1, 21, 22], who take into account a very general definition of abduction which subsumes that of [144], proving many interesting polynomial as well as NP-hardness results.

8 Polynomial Reductions between NMR Problems

Most non-monotonic formalisms have been presented and motivated independently. Nevertheless there have been several attempts at relating them and at finding suitable translations from one logic into another. The existence of these reductions is not only important for a better semantical understanding of the various logics, but it is also useful from the point of view of the complexity analysis.

We focus on polynomial transformations, but recursive non-polynomial, translations can be used in the same spirit for obtaining new decidability/undecidability results. Actually, every polynomial mapping from one formalism into another can be used for obtaining new complexity results on both formalisms. We would like to illustrate this point by means of a couple of examples, using a well-known translation from circumscriptive into default reasoning proposed by Etherington in [50].

Let W be a propositional formula on the alphabet L and $\langle P; Z \rangle$ a partition of L .

Let D be the following set of default rules:

$$D = \{(: \neg p / \neg p) \mid p \in P\}$$

and Δ the default theory $\langle D, W \rangle$. Etherington [50, Theorem 8.3] shows that inference in $CIRC(W; P; Z)$ and Δ are closely related. Actually, for any formula γ built with the variables of L , $CIRC(W; P; Z) \models \gamma$ if and only if γ follows skeptically from (i. e. in every extension of) Δ . In the jargon of [150], the default rules of D are called *prerequisite-free normal unary* (PFNU).

The above result is clearly a translation that can be computed in polynomial time, and proves that skeptical default reasoning is at least computationally as hard as circumscriptive inference, since the latter can be simulated by the former. The translation can be used along with a result by Cadoli and Lenzerini in [26] for obtaining a lower bound on the complexity of skeptical default reasoning for a restricted language. In [26] it is shown that when W is 2CNF with no occurrence of negative literals and γ is a literal, then determining whether $CIRC(W; P; Z) \models \gamma$ holds is a co-NP-complete problem. Using Etherington's translation, we have the following new complexity result:

Theorem 1 *Let W be a 2CNF formula with no occurrence of negative literals, D a set of PFNU default rules, and l any literal. Then determining if l is true in all the extensions of the default theory $\langle D, W \rangle$ is a co-NP-hard problem.*

Theorem 1 can be interestingly compared to a result by Stillman [150, Theorem 2], stating that if W is a 2CNF formula, D is a set of *prerequisite free normal* default rules (superset of PFNU), and l is a literal, then the problem of determining whether l is true in at least one extension of $\langle D, W \rangle$ (credulous inference) can be solved in polynomial time. This confirms a previous result by Kautz and Selman, discussed in Section 3, which shows that in languages of limited expressiveness credulous and skeptical default reasoning may have different complexity.

We have just seen how a polynomial reduction can be used for obtaining new intractability results on the complexity of non-monotonic inference. More formally, if we have a polynomial transformation from a decision problem A into a decision problem B and we have a lower bound (e. g. co-NP-hardness) on the complexity of A , then we can prove the same lower bound for B by means of a many-one reduction.

In fact polynomial transformations can be used for obtaining polynomial results as well. As an example the result by Etherington can be read also in the following way: if we had a polynomial algorithm for skeptical reasoning with PFNU defaults, then we could use that algorithm along with the transformation for performing circumscriptive inference in polynomial time.

Actually Kautz and Selman in [76] show that, if W is a 1CNF formula, D is a set of *normal unary* default rules (superset of PFNU) and l is a literal, then the problem of determining whether l is true in every extension of $\langle D, W \rangle$ can be solved in polynomial time. Following the previous argument, we have the following result:

Theorem 2 *Let W be a 1CNF formula on the alphabet L , $\langle P; Z \rangle$ a partition of L , and l any literal. The problem of determining if $CIRC(W; P; Z) \models l$ is polynomial.*

Although Theorem 2 is implied by a result by Cadoli and Lenzerini [26, Theorem 7], the above construction gives the flavor of the kind of results that can be obtained.

In [27] Cadoli and Lenzerini show a generalization of the reduction by Etherington which allows for fixed predicates is shown. By using this reduction and the complexity results surveyed in Sections 3 and 6, they obtain new complexity results both on default reasoning and on circumscription

The above analysis shows that polynomial translations between formalisms are very important tools for designers of algorithms for non-monotonic reasoning. This kind of technique has been used for example by Gottlob in [64] for proving a lower bound for credulous reasoning in the non-monotonic modal logic N defined by Marek and Truszczyński in [95]. Gottlob used the polynomial embedding (defined by Marek and Truszczyński) of default logic into the logic N along with the fact that credulous default reasoning is Σ_2^p -hard, for proving that credulous reasoning in N is Σ_2^p -hard as well.

In the rest of this section we survey some of the polynomial reductions between non-monotonic reasoning problems that have been proposed in the literature. It is worth noticing that those reductions have been studied with the primary goal of comparing semantics and expressiveness and were not intended for complexity analysis. In particular, transformations are usually required to fulfill some abstract criteria. *Modularity*, defined by Imielinski in [72] in the context of translations from default logic into circumscription, is one of these criteria. Loosely speaking, a translation is modular if the introduction of new facts does not require recomputation of the whole translation from the beginning.

Grosz [69], Imielinski [72] and Etherington [50, 51] addressed the issue of translating default theories into circumscriptive ones. Transformation for some restricted classes of default theories, such as propositional semi-normal defaults without prerequisites, are shown.

Several authors have shown the relations within autoepistemic and default logic. In particular, Konolige in [79] has shown a translation from default logic into strongly grounded autoepistemic logic, while Marek and Truszczyński in [95] have shown an embedding of default logic in the non-monotonic logic N . In a recent paper [65], Gottlob has shown a non-modular embedding of default logic into standard autoepistemic logic, also proving that no modular translation is possible. All the above translations are defined for the whole class of propositional default theories. If we restrict our attention to prerequisite-free default theories then there is a very simple and modular translation which maps single defaults into single modal formulae. This transformation is shown by Gottlob in [65] but it is already implicit in the work of Lin and Shoham [87]. We notice that several results surveyed in Section 3 are valid for prerequisite-free default theories.

It is proven by Gelfond and Lifschitz in [57] that whenever a program has a single stable model then the corresponding set of autoepistemic premises has a single stable expansion and the two coincide in the non-modal part. As a consequence, all the reasoning tasks for the class of autoepistemic theories which are translations of stratified logic programs can be decided in polynomial time. Other reductions from logic programs to default and autoepistemic logic have been shown by Marek and Truszczyński in [94].

Cadoli, Eiter and Gottlob in [25] showed a polynomial transformation of any infer-

ence problem relative to a circumscription with varying predicates into another inference relative to a circumscription with no varying predicates. Previous attempts at eliminating the varying predicates (see Yuan and Wang in [157]) caused exponential growth of the underlying formula. The method in [25] can be used together with another method defined by de Kleer and Konolige in [37] for eliminating fixed predicates. Notice, however, that both methods change the syntactic form of the underlying formula, e. g. a Horn formula may be changed into a non-Horn one.

9 Conclusions

In this paper we surveyed complexity results for many non-monotonic formalisms. In particular, we focused on those which are more directly related to the field of logic programming. We have seen that all forms of non-monotonicity add a new source of complexity to reasoning. In particular NMR is strictly harder than classical reasoning for first-order languages and it is harder in the propositional case unless the polynomial hierarchy collapses at a sufficiently low level.

Due to the lack of space we have not been able to report on complexity results for other areas of NMR, such as belief revision, non-monotonic inheritance networks and the axiomatic approaches to NMR. Nevertheless, we want to give some pointers to the complexity results presented in the literature.

In the area of belief revision and update important results can be found in the works by Eiter and Gottlob [45, 43], Nebel [105, 106] and Grahne and Mendelzon [68].

The complexity of reasoning in non-monotonic inheritance networks has been analyzed by Horty, Thomason and Touretzky [71], Selman and Levesque [143], Geffner and Verma [55], Cadoli et al. [28] and Stein [149].

Lehmann and Magidor in [83] give some results on the computational complexity of various axiomatic systems for non-monotonic reasoning.

In all the works we have surveyed, tractability is obtained through restricting the expressive power of the language. A rather different approach is based on the use of procedures for *approximated* inference. Some results on the use of approximated inference in NMR can be found in the work by Cadoli and Schaerf [29] and Etherington and Crawford [52].

Acknowledgments

We are indebted to many persons for their help. We want to thank George Gottlob, Thomas Eiter, Tom Bylander, Rina Dechter and Rachel Ben-Eliyahu, for finding out mistakes in a previous version and suggesting improvements. Judea Pearl, Jack Minker, Yannis Dimopoulos, Mirosław Truszczyński, Lynn Andrea Stein, George Gottlob, Thomas Eiter, Rina Dechter and Rachel Ben-Eliyahu sent us their most recent results. Daniel Lehmann, Henry Kautz and Krzysztof Apt provided us with more references. Luigia Carlucci Aiello, Andrea Schaerf, Francesco Donini and Maurizio Lenzerini proofread earlier versions. The anonymous referees found several mistakes, suggested important improvements in the structure of the paper and noticed many missing references. Their suggestions greatly helped improving the quality of the paper. Furthermore,

many thanks are due to all the persons who attended the 4th International Workshop on Nonmonotonic Reasoning, held in Vermont in May 1992, who encouraged us in writing this survey paper.

References

- [1] D. Allemang, M. C. Tanner, T. Bylander, and J. R. Josephson. Computational complexity of hypothesis assembly. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-87)*, pages 1112–1117, 1987.
- [2] H. Andreka and I. Nemeti. The generalized completeness of Horn predicate logic as a programming language. *Acta Cybernetica*, 4:3–10, 1978.
- [3] K. R. Apt. Introduction to logic programming. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B, chapter 10. Elsevier Science Publishers B. V. (North Holland), 1990.
- [4] K. R. Apt and M. Bezem. Acyclic programs. *New Generation Computing*, 9:335–363, 1991.
- [5] K. R. Apt and H. A. Blair. Arithmetic classification of perfect models of stratified programs. *Fundamenta Informaticae*, XIII:1–18, 1990. (with addendum in vol. XIV:339–343, 1991).
- [6] K. R. Apt, H. A. Blair, and A. Walker. Towards a theory of declarative knowledge. In J. Minker, editor, *Foundation of Deductive Databases and Logic Programming*, pages 89–142. Morgan Kaufmann, 1988.
- [7] F. Baader and B. Hollunder. Embedding defaults into terminological knowledge representation formalisms. In *Proceedings of the Third International Conference on the Principles of Knowledge Representation and Reasoning (KR-92)*, pages 306–317, 1992.
- [8] C. Bell, A. Nerode, R. T. Ng, and V. S. Subrahmanian. Implementing deductive databases by linear programming. In *Proceedings of the Eleventh Conference on Principle Of Database Systems (PODS-92)*, pages 283–292, 1992.
- [9] R. Ben-Eliyahu and R. Dechter. Default logic, propositional logic and constraints. In *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI-91)*, pages 379–385, 1991.
- [10] R. Ben-Eliyahu and R. Dechter. Propositional semantics for default logic. Technical Report R-163, Cognitive Systems Laboratory, Computer Science Department, University of California at Los Angeles, July 1992. Also presented at the 4th International Workshop on Nonmonotonic Reasoning, Vermont, May 1992. Notes edited by H. Kautz and D. W. Etherington.

- [11] R. Ben-Eliyahu and R. Dechter. Propositional semantics for disjunctive logic programming. In *Proceedings of the Joint International Conference and Symposium on Logic Programming (JICSLP-92)*, pages 379–385, 1992. Full paper to appear in *Annals of Mathematics and Artificial Intelligence*.
- [12] N. Bidoit. Negation in rule-based database languages: a survey. *Theoretical Computer Science*, 78:3–83, 1991.
- [13] N. Bidoit and C. Froidevaux. General logic databases and programs: default logic semantics and stratification. *Information and Computation*, 91:15–54, 1991.
- [14] N. Bidoit and C. Froidevaux. Negation by default and unstratifiable logic programs. *Theoretical Computer Science*, 78:85–112, 1991.
- [15] N. Bidoit and R. Hull. Minimalism, justification and non-monotonicity in deductive databases. *Journal of Computer and System Sciences*, 38:290–325, 1989.
- [16] H. A. Blair. The recursion-theoretic complexity of the semantics of predicate logic as a programming language. *Information and Control*, 54:25–47, 1982.
- [17] H. A. Blair, W. Marek, and J. S. Schlipf. The expressiveness of locally stratified programs. Technical Report 92-8, Mathematical Sciences Institute, Cornell University, 1992. Submitted.
- [18] A. Borgida and D. W. Etherington. Hierarchical knowledge bases and efficient disjunctive reasoning. In *Proceedings of the First International Conference on the Principles of Knowledge Representation and Reasoning (KR-89)*, pages 33–43, 1989.
- [19] P. Bossu and P. Siegel. Saturation, nonmonotonic reasoning and the closed world assumption. *Artificial Intelligence Journal*, 25:16–63, 1985.
- [20] R. J. Brachman and H. J. Levesque. A fundamental tradeoff in knowledge representation and reasoning. In R. J. Brachman and H. J. Levesque, editors, *Readings in knowledge representation*, pages 41–70. Morgan Kaufmann, 1985.
- [21] T. Bylander. The monotonic abduction problem: A functional characterization on the edge of tractability. In *Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning (KR-91)*, pages 70–77, 1991.
- [22] T. Bylander, D. Allemang, M. C. Tanner, and J. R. Josephson. The computational complexity of abduction. *Artificial Intelligence Journal*, 49:25–60, 1991.
- [23] M. Cadoli. The complexity of model checking for circumscriptive formulae. *Information Processing Letters*, 44:113–118, 1992.
- [24] M. Cadoli. On the complexity of model finding for nonmonotonic propositional logics. In A. Marchetti Spaccamela, P. Mentrasti, and M. Venturini Zilli, editors, *Proceedings of the Fourth Italian Conference on Theoretical Computer Science*, pages 125–139. World Scientific Publishing Co., October 1992.

- [25] M. Cadoli, T. Eiter, and G. Gottlob. An efficient method for eliminating varying predicates from a circumscription. *Artificial Intelligence Journal*, 54:397–410, 1992.
- [26] M. Cadoli and M. Lenzerini. The complexity of closed world reasoning and circumscription. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90)*, pages 550–555. MIT press, July 1990. Extended version as Technical Report RAP. 17.91, Università di Roma "La Sapienza", Dipartimento di Informatica e Sistemistica, 1991. To appear in *Journal of Computer and System Sciences*.
- [27] M. Cadoli and M. Lenzerini. Mapping complexity results from circumscription to default logic and viceversa. Unpublished manuscript, July 1991.
- [28] M. Cadoli, M. Lenzerini, D. Nardi, and F. Pirri. Circumscription and nonmonotonic inheritance. In *Proceedings of the 1st Pacific Rim Conference on Artificial Intelligence (PRICAI-90)*, pages 760–765, November 1990.
- [29] M. Cadoli and M. Schaerf. Approximate inference in default reasoning and circumscription. In B. Neumann, editor, *Proceedings of the Tenth European Conference on Artificial Intelligence (ECAI-92)*, pages 319–323. John Wiley, August 1992. Extended version presented at the 4th International Workshop on Nonmonotonic Reasoning, Vermont, May 1992. Notes edited by Kautz H. and Etherington D. To appear in *Fundamenta Informaticae*.
- [30] S. Ceri, G. Gottlob, and L. Tanca. *Logic programming and databases*. Springer-Verlag, Berlin, 1990.
- [31] E. Chan. A possible worlds semantics for disjunctive databases. *IEEE Transactions on Data and Knowledge Engineering*, 1992. To appear.
- [32] C. Cholak and H. Blair. The complexity of local stratification. Technical Report SU-CIS-91-27, School of Computer and Information Science, Syracuse University, July 1991.
- [33] P. Cholak. Post correspondence problem and Prolog programs. Technical report, University of Wisconsin, Madison, 1988.
- [34] J. Chomicki and V. S. Subrahmanian. Generalized closed world assumption is Π_2^0 -complete. *Information Processing Letters*, 34:289–291, 1990.
- [35] K. L. Clark. Negation as failure. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, pages 293–322. Plenum, 1978.
- [36] J. de Kleer. An assumption-based TMS. *Artificial Intelligence Journal*, 28:127–162, 1986.
- [37] J. de Kleer and K. Konolige. Eliminating the fixed predicates from a circumscription. *Artificial Intelligence Journal*, 39:391–398, 1989.

- [38] Y. Dimopoulos and V. Magirou. A graph theoretic approach to default logic. *Information and Computation*, 1992. To appear.
- [39] F. M. Donini, M. Lenzerini, D. Nardi, F. Pirri, and M. Schaerf. Nonmonotonic reasoning. *Artificial Intelligence Review*, 4:163–210, 1990.
- [40] W. P. Dowling and J. H. Gallier. Linear-time algorithms for testing the satisfiability of propositional horn formulae. *Journal of Logic Programming*, 1:267–284, 1984.
- [41] T. Eiter and G. Gottlob. Propositional circumscription and extended closed world reasoning are Π_2^p -complete. Technical Report CD-TR 91/20, Technische Universität Wien, Vienna Austria, Christian Doppler Labor für Expertensysteme, May 1991. To appear in *Theoretical Computer Science*.
- [42] T. Eiter and G. Gottlob. Complexity aspects of various semantics for disjunctive databases. Technical Report CD-TR 92/46, Technische Universität Wien, Vienna Austria, Christian Doppler Labor für Expertensysteme, November 1992.
- [43] T. Eiter and G. Gottlob. The complexity of nested counterfactuals and iterated knowledge base revision. Technical Report CD-TR 92/44, Technische Universität Wien, Vienna Austria, Christian Doppler Labor für Expertensysteme, October 1992.
- [44] T. Eiter and G. Gottlob. Complexity results for disjunctive logic programming and application to nonmonotonic logics. Technical Report CD-TR 92/41, Technische Universität Wien, Vienna Austria, Christian Doppler Labor für Expertensysteme, October 1992.
- [45] T. Eiter and G. Gottlob. On the complexity of propositional knowledge base revision, updates and counterfactuals. *Artificial Intelligence Journal*, 57:227–270, 1992.
- [46] T. Eiter and G. Gottlob. Reasoning with parsimonious and moderately grounded expansions. *Fundamenta Informaticae*, 17:31–54, 1992.
- [47] T. Eiter and G. Gottlob. The complexity of logic-based abduction. In *Tenth Symposium on Theoretical Aspects of Computer Science (STACS-93)*, 1993. In press, extended version as Technical Report CD-TR 92/35 Technische Universität Wien, Vienna Austria, Christian Doppler Labor für Expertensysteme.
- [48] K. Eshghi. Tractable abduction. Manuscript, 1992.
- [49] K. Eshghi and R. A. Kowalski. Abduction compared with negation by failure. In *Proceedings of the Sixth International Conference on Logic Programming (ICLP-89)*, pages 234–254, 1989.
- [50] D. V. Etherington. *Reasoning with incomplete information*. Morgan Kaufman Publishers, Los Altos, CA, 1987.

- [51] D. V. Etherington. Relating default logic and circumscription. In *Proceedings of the Tenth International Joint Conference on Artificial Intelligence (IJCAI-87)*, pages 489–494, 1987.
- [52] D. V. Etherington and J. M. Crawford. Towards efficient default reasoning. Presented at the 4th International Workshop on Nonmonotonic Reasoning, Vermont, May 1992. Notes edited by Kautz H. and Etherington D., 1992.
- [53] G. Friedrich, G. Gottlob, and W. Nejdl. Hypothesis classification, abductive diagnosis and therapy. In *Proceedings of the International Workshop on Expert Systems in Engineering*, volume 462 of *Lecture Notes In Artificial Intelligence*, pages 69–78. Springer-Verlag, 1990.
- [54] M. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness*. W. H. Freeman, 1979.
- [55] H. Geffner and T. Verma. Inheritance = chaining + defeat. In *Proceedings of the 4th International Symposium on Methodologies for Intelligent Systems (ISMIS-89)*, pages 411–418, New York, 1989. North Holland.
- [56] M. Gelfond. On stratified autoepistemic theories. In *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, pages 207–211. Morgan Kaufmann, 1987.
- [57] M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In *Proceedings of the Fifth Logic Programming Symposium*, pages 1070–1080. MIT Press, 1988.
- [58] M. Gelfond and V. Lifschitz. Classical negation in logic programs and disjunctive databases. *New Generation Computing*, 9:365–385, 1991.
- [59] M. Gelfond and H. Przymusinska. Negation as failure: Careful closure procedure. *Artificial Intelligence Journal*, 30:273–287, 1986.
- [60] M. Gelfond, H. Przymusinska, and T. Przymusinsky. On the relationship between circumscription and negation as failure. *Artificial Intelligence Journal*, 38:49–73, 1989.
- [61] M. L. Ginsberg. A circumscriptive theorem prover. *Artificial Intelligence Journal*, 39:209–230, 1989.
- [62] M. Goldszmidt and J. Pearl. Rank-based systems: A simple approach to belief revision, belief update and reasoning about evidence and actions. In *Proceedings of the Third International Conference on the Principles of Knowledge Representation and Reasoning (KR-92)*, pages 661–672, 1992.
- [63] G. Gottlob. The complexity of propositional default reasoning under the stationary fixed point semantics. Technical Report CD-TR 92/42, Technische Universität Wien, Vienna Austria, Christian Doppler Labor für Expertensysteme, 1992.

- [64] G. Gottlob. Complexity results for nonmonotonic logics. *Journal of Logic and Computation*, 2:397–425, 1992.
- [65] G. Gottlob. On the power of pure beliefs or embedding default logic into standard autoepistemic logic. Technical Report CD-TR 92/34, Technische Universität Wien, Vienna Austria, Christian Doppler Labor für Expertensysteme, 1992. Presented at the ECAI-92 Workshop on Theoretical Foundations of Knowledge Representation and Reasoning, Vienna, August 1992. Notes edited by G. Lakemeyer and B. Nebel.
- [66] G. Gottlob and C. G. Fermüller. Removing redundancy from a clause. *Artificial Intelligence Journal*, 1993. To appear.
- [67] G. Grahne. *The Problem of Incomplete Information in Relational Databases*, volume 554 of *Lecture Notes in Computer Science*. Springer-Verlag, 1991.
- [68] G. Grahne and A. O. Mendelzon. Updates and subjunctive queries. Presented at the Workshop on Nonstandard queries and answers, Toulouse, France, July 1991. Notes edited by R. Demolombe, L. Fariñas del Cerro and T. Imielinski, 1991.
- [69] B. Grosz. Default reasoning as circumscription. Presented at the First International Workshop on Nonmonotonic Reasoning, 1984.
- [70] P. G. Hinman. *Recursion-Theoretic Hierarchies*. Springer-Verlag, 1978.
- [71] J. F. Horty, R. H. Thomason, and D. S. Touretzky. A skeptical theory of inheritance in nonmonotonic semantic networks. *Artificial Intelligence Journal*, 42:311–348, 1990.
- [72] T. Imielinski. Results on translating defaults to circumscription. *Artificial Intelligence Journal*, 32:131–146, 1987.
- [73] K. Inoue and N. Helft. On theorem provers for circumscription. In *Proceedings of the Eight Biennial Conference of the Canadian Society for Computational Studies of Intelligence (CSCSI '90)*. Morgan Kaufmann Publishers, 1990.
- [74] D. S. Johnson. A catalog of complexity classes. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A, chapter 2. Elsevier Science Publishers B. V. (North Holland), 1990.
- [75] A. C. Kakas and P. Mancarella. Negation as stable hypothesis. In *Proceedings of the First International Workshop on Logic Programming and Non Monotonic Reasoning*, pages 275–288. The MIT press, 1991.
- [76] H. A. Kautz and B. Selman. Hard problems for simple default logics. *Artificial Intelligence Journal*, 49:243–279, 1991.
- [77] P. G. Kolaitis and C. H. Papadimitriou. Some computational aspects of circumscription. *Journal of the ACM*, 37:1–14, 1990.
- [78] P. G. Kolaitis and C. H. Papadimitriou. Why not negation by fixpoint? *Journal of Computer and System Sciences*, 43:125–144, 1991.

- [79] K. Konolige. On the relationship between default and autoepistemic logic. *Artificial Intelligence Journal*, 35:343–382, 1988.
- [80] T. Krishnaprasad. On the computability of circumscription. *Information Processing Letters*, 27:237–243, 1988.
- [81] K. Kunen. Negation in logic programming. *Journal of Logic Programming*, 4:289–308, 1987.
- [82] R. E. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM Journal of Computing*, 6:467–480, 1977.
- [83] D. Lehmann and M. Magidor. What does a conditional knowledge base entail? *Artificial Intelligence Journal*, 55:1–60, 1992.
- [84] V. Lifschitz. Closed-world databases and circumscription. *Artificial Intelligence Journal*, 27:229–235, 1985.
- [85] V. Lifschitz. Computing circumscription. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence (IJCAI-85)*, pages 121–127, 1985.
- [86] V. Lifschitz. On the declarative semantics of logic programs with negation. In J. Minker, editor, *Foundation of Deductive Databases and Logic Programming*, pages 177–192. Morgan Kaufmann, 1988.
- [87] F. Lin and Y. Shoham. A logic of knowledge and justified assumptions. *Artificial Intelligence Journal*, 57:271–290, 1992.
- [88] J. Lobo, J. Minker, and A. Rajasekar. *Foundations of Disjunctive Logic Programming*. MIT Press, Cambridge, Massachusetts, 1992.
- [89] W. Lukaszewicz. *Non-Monotonic Reasoning*. Ellis-Horwood, Chichester, West Sussex, England, 1990.
- [90] W. Marek, A. Nerode, and J. Remmel. The stable models of a predicate logic program. In *Joint International Conference and Symposium on Logic Programming*, pages 446–460, 1992.
- [91] W. Marek, A. Nerode, and J. Remmel. A theory of nonmonotonic rule systems II. *Annals of Mathematics and Artificial Intelligence*, 1992. to appear.
- [92] W. Marek, G. F. Shvarts, and M. Truszczyński. Modal nonmonotonic logics: Ranges, characterization, computation. In *Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning (KR-91)*, pages 395–404, 1991.
- [93] W. Marek and V. S. Subrahmanian. The relationship between logic program semantics and non-monotonic reasoning. In *Proceedings of the Sixth International Conference on Logic Programming (ICLP-89)*, pages 600–617, 1989.

- [94] W. Marek and M. Truszczyński. Stable semantics for logic programs and default theories. In E. Lusk and R. Overbeek, editors, *Proc. of the North American Conference on Logic Programming*, pages 243–256. MIT Press, 1989.
- [95] W. Marek and M. Truszczyński. Modal logic for default reasoning. *Annals of Mathematics and Artificial Intelligence*, 1:275–302, 1990.
- [96] W. Marek and M. Truszczyński. Autoepistemic logic. *Journal of the ACM*, 38:588–619, 1991.
- [97] W. Marek and M. Truszczyński. Computing intersection of autoepistemic expansions. In *Proceedings of the First International Workshop on Logic Programming and Non Monotonic Reasoning*, pages 37–50. The MIT press, 1991.
- [98] J. McCarthy. Circumscription - a form of non-monotonic reasoning. *Artificial Intelligence Journal*, 13:27–39, 1980.
- [99] J. McCarthy. Applications of circumscription to formalizing common-sense knowledge. *Artificial Intelligence Journal*, 28:89–116, 1986.
- [100] D. McDermott. Non-monotonic logic II: Non-monotonic modal theories. *Journal of the ACM*, 29:33–57, 1982.
- [101] D. McDermott and J. Doyle. Non-monotonic logic I. *Artificial Intelligence Journal*, 13:41–72, 1980.
- [102] J. Minker. On indefinite databases and the closed world assumption. In *Proceedings of the Sixth Conference on Automated Deduction (CADE-82)*, pages 292–308, 1982.
- [103] J. Minker. An overview of nonmonotonic reasoning and logic programming. Technical Report CS-TR-2736, Department of Computer Science, University of Maryland, August 1991. Invited banquet address, First International Workshop on LP and NMR, 1991.
- [104] R. C. Moore. Semantical considerations on nonmonotonic logic. *Artificial Intelligence Journal*, 25:75–94, 1985.
- [105] B. Nebel. A knowledge-level analysis of belief revision. In *Proceedings of the First International Conference on the Principles of Knowledge Representation and Reasoning (KR-89)*, pages 301–311, 1989.
- [106] B. Nebel. Belief revision and default reasoning: Syntax-based approaches. In *Proceedings of the Second International Conference on the Principles of Knowledge Representation and Reasoning (KR-91)*, pages 417–428, 1991.
- [107] A. Nerode, R. T. Ng, and V. S. Subrahmanian. Computing circumscriptive databases, part i: Theory and algorithms. Technical Report CS-TR-2824, University of Maryland. Department of Computer Science, 1992.
- [108] I. Niemelä. On the decidability and complexity of autoepistemic reasoning. *Fundamenta Informaticae*, 17:117–156, 1992.

- [109] I. Niemelä and J. Rintanen. On the impact of stratification on the complexity of nonmonotonic reasoning. In *Proceedings of the Third International Conference on the Principles of Knowledge Representation and Reasoning (KR-92)*, pages 627–638, 1992.
- [110] N. Olivetti. Tableaux and sequent calculus for minimal entailment. *Journal of Automated Reasoning*, 9:99–139, 1992.
- [111] C. H. Papadimitriou. On the complexity of unique solutions. *Journal of the ACM*, 31:392–400, 1984.
- [112] C. H. Papadimitriou. On selecting a satisfying truth assignment (extended abstract). In *Proceedings of the 32nd Annual Symposium on the Foundations of Computer Science (FOCS-91)*, pages 163–169, 1991.
- [113] C. H. Papadimitriou and M. Sideri. On finding extensions of default theories. In *Proceedings of the International Conference on Database Theory (ICDT-92)*, pages 276–281, 1992.
- [114] C. H. Papadimitriou and M. Yannakakis. Tie-breaking semantics and structural totality. In *Proceedings of the Eleventh Conference on Principle Of Database Systems (PODS-92)*, pages 16–22, 1992.
- [115] M. Papalaskari and S. Weinstein. Minimal consequence in sentential logic. *Journal of Logic Programming*, 9:19–31, 1990.
- [116] J. Pearl. Probabilistic semantics for nonmonotonic reasoning: A survey. In R. Cummins and J. Pollock, editors, *Philosophy and AI: Essays at the Interface*, pages 157–187. Bradford Books/MIT Press, Cambridge, MA, 1991.
- [117] D. L. Poole, R. G. Goebel, and R. Aleliunas. Theorist: A logical reasoning system for defaults and diagnosis. In N. Cercone and G. McCalla, editors, *The Knowledge Frontier: Essays in the Representation of Knowledge*, pages 331–352. Springer, New York, 1987.
- [118] G. M. Provan. A complexity analysis of assumption-based truth maintenance systems. In B. Smith and G. Kelleher, editors, *Reason Maintenance Systems and their Applications*, pages 98–113. Ellis Horwood, New York, 1988.
- [119] G. M. Provan. The computational complexity of multiple-context truth maintenance systems. In *Proceedings of the Ninth European Conference on Artificial Intelligence (ECAI-90)*, pages 522–527, 1990.
- [120] H. Przymusinska and T. Przymusinsky. Stationary default extensions. Presented at the 4th International Workshop on Nonmonotonic Reasoning, Vermont, May 1992. Notes edited by Kautz H. and Etherington D., 1992.
- [121] T. Przymusinski. On the declarative semantics of stratified deductive databases and logic programs. In J. Minker, editor, *Foundation of Deductive Databases and Logic Programming*, pages 193–216. Morgan Kaufmann, 1988.

- [122] T. Przymusiński. An algorithm to compute circumscription. *Artificial Intelligence Journal*, 38:49–73, 1989.
- [123] T. Przymusiński. Non-monotonic formalisms and logic programming. In *Proceedings of the Sixth International Conference on Logic Programming*, pages 655–674. The MIT Press, 1989.
- [124] T. Przymusiński. Stable semantics for disjunctive programs. *New Generation Computing*, 9:401–424, 1991.
- [125] A. Rabinov. A generalization of collapsible cases of circumscription. *Artificial Intelligence Journal*, 38:111–117, 1989.
- [126] A. Rajasekar, J. Lobo, and J. Minker. Weak generalized closed world assumption. *Journal of Automated Reasoning*, 5:293–307, 1989.
- [127] J. A. Reggia, D. S. Nau, and P. Wang. Diagnostic expert systems based on a set covering model. *Int. J. Man-Machine Studies*, 19:437–460, 1983.
- [128] R. Reiter. On closed world data bases. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, pages 119–140. Plenum, 1978.
- [129] R. Reiter. A logic for default reasoning. *Artificial Intelligence Journal*, 13:81–132, 1980.
- [130] R. Reiter. Circumscription implies predicate completion (sometimes). In *Proceedings of the National Conference on Artificial Intelligence (AAAI-82)*, pages 418–420, 1982.
- [131] H. Rogers. *Theory of Recursive Functions and Effective Computability*. McGraw-Hill, 1967.
- [132] K. Ross and R. Topor. Inferring negative information from a disjunctive database. *Journal of Automated Reasoning*, 4:397–424, 1988.
- [133] V. Rutenburg. Computational complexity of truth maintenance systems. In *Eight Symposium on Theoretical Aspects of Computer Science (STACS-91)*, volume 480 of *Lecture Notes in Computer Science*. Springer-Verlag, 1991.
- [134] M. Schaerf. Logic programming and autoepistemic logics: new relations and complexity results. Submitted, 1992.
- [135] J. S. Schlipf. Decidability and definability with circumscription. *Annals of Pure and Applied Logic*, 35:173–191, 1987.
- [136] J. S. Schlipf. When is closed world reasoning tractable? In *Proceedings of the 3rd International Symposium on Methodologies for Intelligent Systems (ISMIS-88)*, pages 485–494, New York, 1988. North Holland.

- [137] J. S. Schlipf. The expressive powers of the logic programming semantics. In *Proceedings of the Ninth Conference on Principle Of Database Systems (PODS-90)*, pages 196–204, 1990. Expanded version available as University of Cincinnati Computer Science Department Tech. Rep. CIS-TR-90-3.
- [138] J. S. Schlipf. Formalizing a logic for logic programming. *Annals of Mathematics and Artificial Intelligence*, 5:279–302, 1992.
- [139] J. S. Schlipf. A survey of complexity and undecidability results in logic programming. In *Proc. of the Workshop on Structural Complexity and Recursion-theoretic methods in Logic Programming*, pages 143–164, Washington, DC, 1992.
- [140] G. Schwarz and M. Truszczyński. Nonmonotonic reasoning is sometimes easier. Technical Report 213-92, Department of Computer Science, University of Kentucky, 1992.
- [141] C. B. Schwind. A tableaux-based theorem prover for a decidable subset of default logic. In *Proceedings of the Tenth Conference on Automated Deduction (CADE-90)*, pages 528–542, 1990.
- [142] B. Selman and H. A. Kautz. Model-preference default theories. *Artificial Intelligence Journal*, 45:287–322, 1990.
- [143] B. Selman and H. J. Levesque. The tractability of path-based inheritance. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence (IJCAI-89)*, pages 1140–1145, 1989.
- [144] B. Selman and H. J. Levesque. Abductive and default reasoning: A computational core. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90)*, pages 343–348, 1990.
- [145] B. Selman, H. J. Levesque, and D. Mitchell. A new method for solving hard satisfiability problems. In *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92)*, pages 440–446, 1992.
- [146] J. C. Sheperdson. Negation as failure: A comparison of Clark’s completed data base and Reiter’s closed-world assumption. *Journal of Logic Programming*, 1:51–79, 1984.
- [147] G. Shvarts. Autoepistemic modal logics. In *Proceedings of the Third Conference on Theoretical Aspects of Reasoning about Knowledge (TARK-90)*, pages 97–109, 1990.
- [148] R. M. Smullyan. *First Order Logic*. Springer, Berlin, 1968.
- [149] L. A. Stein. Resolving ambiguity in nonmonotonic inheritance hierarchies. *Artificial Intelligence Journal*, 55:259–310, 1992.
- [150] J. Stillman. It’s not my default: The complexity of membership problems in restricted propositional default logics. In *Proceedings of the Eighth National Conference on Artificial Intelligence (AAAI-90)*, pages 571–578, 1990.

- [151] J. Stillman. The complexity of propositional default logics. In *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI-92)*, pages 794–799, 1992.
- [152] J. van Benthem and K. Doets. Higher order logic. In D. Gabbay and F. Guenther, editors, *Handbook of Philosophical Logic*, volume 1, chapter 4. Reidel, Dordrecht, Netherlands, 1983.
- [153] A. van Gelder. The alternating fixpoint of logic programs with negation. In *Proceedings of the Eighth Conference on Principle Of Database Systems (PODS-89)*, pages 1–10, 1989.
- [154] A. van Gelder, K. A. Ross, and J. S. Schlipf. The well-founded semantics for general logic programs. *Journal of the ACM*, 38:620–650, 1991.
- [155] M. Y. Vardi. Querying logical databases. *Journal of Computer and System Sciences*, 33:142–160, 1986.
- [156] A. Yahya and L. J. Henschen. Deduction in non-Horn databases. *Journal of Automated Reasoning*, 1:141–160, 1985.
- [157] L. Y. Yuan and Wang. C. H. On reducing parallel circumscription. In *Proceedings of the Seventh National Conference on Artificial Intelligence (AAAI-88)*, pages 460–464, 1988.
- [158] A. Zhang and W. Marek. On the classification and existence of structures in default logic. *Fundamenta Informaticae*, XIII:485–499, 1990.