

KNOWLEDGE REPRESENTATION WITH LOGIC PROGRAMS

Gerhard Brewka and Jürgen Dix

ABSTRACT: In this tutorial-overview, which resulted from a lecture course given by the authors at the *European Summer School in Logic, Language and Information 1997* in Aix-en-Provence (<http://www.lpl.univ-aix.fr/~ess11i97/>), we show how knowledge representation (KR) can be done with the help of *generalized* logic programs. We start by introducing the core of PROLOG, which is based on *definite* logic programs. Although this class is very restricted (and will be enriched by various additional features in the rest of the paper), it has a very nice property for KR-tasks: there exist efficient *query-answering procedures*—a *top-down* approach and a *bottom-up* evaluation. In addition we can not only handle ground queries but also queries with variables and compute *answer-substitutions*.

It turns out that more advanced KR-tasks can not be properly handled with definite programs. Therefore we extend this basic class of programs by additional features like *negation-as-finite-failure*, *default-negation*, *explicit negation*, *preferences*, and *disjunction*. The need for these extensions is motivated by suitable examples and the corresponding semantics are discussed in detail.

Clearly, the more expressive the respective class of programs under a certain semantics is, the less efficient are potential query-answering methods. This point will be illustrated and discussed for every extension. By well-known recursion-theoretic results, it is obvious that there do not exist complete query-answering procedures for the general case where variables and function symbols are allowed. Nevertheless we consider it an important topic of further research to extract *feasible* classes of programs where answer-substitutions can be computed.

1 INTRODUCTION

One of the major reasons for the success story (if one is really willing to call it a success story) of human beings on this planet is our ability to invent tools that help us improve our—otherwise often quite limited—capabilities. The invention of machines that are able to do interesting things, like transporting people from one place to the other (even through the air), sending moving pictures and sounds around the globe, bringing our email to the right person, and the like, is one of the cornerstones of our culture and determines to a great degree our everyday life.

Among the most challenging tools one can think of are machines that are able to handle knowledge adequately. Wouldn't it be great if, instead of the stupid device which brings coffee from the kitchen to your office every day at 9 am, and which needs complete reengineering whenever your coffee preferences change, you could (for the same price, admitted) get a smart robot who simply can be told that you want your coffee black this morning, and that you need an extra Aspirin since it was your colleague's birthday yesterday? To react in the right way to your needs such a robot would have to know a lot, for instance that Aspirin should come with a glass of water, or that people in certain situations need their coffee extra strong.

Building smart machines of this kind is at the heart of Artificial Intelligence (AI). Since such machines will need tremendous amounts of knowledge to work

properly, even in very limited environments, the investigation of techniques for representing knowledge and reasoning is highly important.

In the early days of AI it was still believed that modeling general purpose problem solving capabilities, as in Newell and Simon's famous GPS (General Problem Solver) program, would be sufficient to generate intelligent behaviour. This hypothesis, however, turned out to be overly optimistic. At the end of the sixties people realized that an approach using available knowledge about narrow domains was much more fruitful. This led to the expert systems boom which produced many useful application systems, expert system building tools, and expert system companies. Many of the systems are still in use and save companies millions of dollars per year¹.

Nevertheless, the simple knowledge representation and reasoning methods underlying the early expert systems soon turned out to be insufficient. Most of the systems were built based on simple rule languages, often enhanced with ad hoc approaches to model uncertainty. It became apparent that more advanced methods to handle incompleteness, defeasible reasoning, uncertainty, causality and the like were needed.

This insight led to a tremendous increase of research on the foundations of knowledge representation and reasoning. Theoretical research in this area has blossomed in recent years. Many advances have been made and important results were obtained. The technical quality of this work is often impressive.

On the other hand, most of these advanced techniques have had surprisingly little influence on practical applications so far. To a certain degree this is understandable since theoretical foundations had to be laid first and pioneering work was needed. However, if we do not want research in knowledge representation to remain a theoreticians' game more emphasis on computability and applicability seems to be needed. We strongly believe that the kind of research presented in this tutorial, that is research aiming at interesting combinations of ideas from logic programming and nonmonotonic reasoning, provides an important step into this direction.

1.1 Some History

Historically, logic programs have been considered in the logic programming community for more than 20 years. It began with [Colmerauer *et al.*, 1973; Kowalski, 1974; van Emden and Kowalski, 1976] and led to the definition and implementation of *PROLOG*, a by now theoretically well-understood programming language (at least the declarative part consisting of Horn-clauses: *pure PROLOG*). Extensions of *PROLOG* allowing negative literals have been also considered in this area: they rely on the idea of *negation-as-finite-failure*, we call them *logic-programming-semantics* (or shortly *LP-semantics*).

¹We refer the interested reader to [Russel and Norvig, 1995] which gives a very detailed and nice exposition of what has been done in AI since its very beginning until today.

In parallel, starting at about 1980, *Nonmonotonic Reasoning* entered into computer science and began to constitute a new field of active research. It was originally initiated because *Knowledge Representation* and *Common-Sense Reasoning* using classical logic came to its limits. Formalisms like classical logic are inherently monotonic and they seem to be too weak and therefore inadequate for such reasoning problems.

In recent years, independently of the research in logic programming, people interested in knowledge representation and nonmonotonic reasoning also tried to define declarative semantics for programs containing *default* or *explicit* negation and even *disjunctions*. They defined various semantics by appealing to (different) intuitions they had about programs.

This second line of research started in 1986 with the *Workshop on the Foundations of Deductive Databases and Logic Programming* organized by Jack Minker: the revised papers of the proceedings were published in [Minker, 1988]. The *stratified* (or the similar *perfect*) semantics presented there can be seen as a splitting-point: it is still of interest for the logic programming community (see [Cavedon and Lloyd, 1989]) but its underlying intuitions were inspired by nonmonotonic reasoning and therefore much more suitable for knowledge representation tasks. Semantics of this kind leave the philosophy underlying classical logic programming in that their primary aim is not to model *negation-as-finite-failure*, but to construct new, more powerful semantics suitable for applications in knowledge representation. Let us call such semantics *NMR-semantics*.

Nowadays, due to the work of Apt, Blair and Walker, Fitting, Lifschitz, Przymusiński and others, very close relationships between these two independent research lines became evident. Methods from logic programming, e.g. least fixpoints of certain operators, can be used successfully to define NMR-semantics.

The NMR-semantics also shed new light on the understanding of the classical nonmonotonic logics such as *Default Logic*, *Autoepistemic Logic* and the various versions of *Circumscription*. In addition, the investigation of possible semantics for logic programs seems to be useful because

1. parts of nonmonotonic systems (which are usually defined for full predicate logic, or even contain additional (modal)-operators) may be “implemented” with the help of such programs,
2. nonmonotonicity in these logics may be described with an appropriate treatment of negation in logic programs.

1.2 Non-Monotonic Formalisms in KR

As already mentioned above, research in nonmonotonic reasoning has begun at the end of the seventies. One of the major motivations came from reasoning about actions and events. John McCarthy and Patrick Hayes had proposed their situation calculus as a means of representing changing environments in logic. The basic idea is to use an extra situation argument for each fact which describes the situation in

which the fact holds. Situations, basically, are the results of performing sequences of actions. It soon turned out that the problem was not so much to represent what changes but *to represent what does not change* when an event occurs. This is the so-called *frame problem*. The idea was to handle the frame problem by using a default rule of the form

If a property P holds in situation S then P typically also holds in the situation obtained by performing action A in S .

Given such a rule it is only necessary to explicitly describe the changes induced by a particular action. All non-changes, for instance that the *real colour* of the kitchen wall does not change when the light is turned on, are handled implicitly. Although it turned out that a straightforward formulation of this rule in some of the most popular nonmonotonic formalisms may lead to unintended results the frame problem was certainly the challenge motivating many people to join the field.

In the meantime a large number of different nonmonotonic logics have been proposed. We can distinguish four major types of such logics:

1. Logics using nonstandard inference rules with an additional consistency check to represent default rules. Reiter's default logic (see Appendix A.3) and its variants are of this type.
2. Nonmonotonic modal logics using a modal operator to represent consistency or (dis-) belief. These logics are nonmonotonic since conclusions may depend on disbelief. The most prominent example is Moore's autoepistemic logic.
3. Circumscription (see Appendix A.4) and its variants. These approaches are based on a preference relation on models. A formula is a consequence iff it is true in all most preferred models of the premises. Syntactically, a second order formula is used to eliminate all non-preferred models.
4. Conditional approaches which use a non truth-functional connective \rightsquigarrow to represent defaults. A particularly interesting way of using such conditionals was proposed by Kraus, Lehmann and Magidor. They consider p as a default consequence of q iff the conditional $q \rightsquigarrow p$ is in the closure of a given conditional knowledge base under a collection of rules. Each of the rules directly corresponds to a desirable property of a nonmonotonic inference relation.

The various logics are intended to handle different intuitions about nonmonotonic reasoning in a most general way. On the other hand, the generality leads to problems, at least from the point of view of implementations and applications. In the first order case the approaches are not even semi-decidable since an implicit consistency check is needed. In the propositional case we still have tremendous complexity problems. For instance, the complexity of determining whether a formula is contained in all extensions of a propositional default theory is on the second

level of the polynomial hierarchy. As mentioned earlier we believe that logic programming techniques can help to overcome these difficulties.

Originally, nonmonotonic reasoning was intended to provide us with a *fast* but *unsound* approximation of classical reasoning in the presence of incomplete knowledge. Therefore one might ask whether the higher complexity of NMR-formalisms (compared to classical logic) is not a real drawback of this aim? The answer is that NMR-systems allow us to formulate a problem in a very *compact* way as a theory T . It turns out that for some problems any equivalent formulation in classical logic (if possible at all) as a theory T' is much larger: the size of T' is exponential in the size of T ! We refer to [Gogic *et al.*, 1995] and [Cadoli *et al.*, 1996; Cadoli *et al.*, 1997; Cadoli *et al.*, 1995] where such problems are investigated.

1.3 How this Paper is organized

In this overview paper we show how Knowledge Representation can be done with the help of *generalized* logic programs. We start by introducing the core of PROLOG, which is based on *definite* logic programs. Although this class is very restricted (and will be enriched by various additional features in the rest of the paper), it has a very nice property for KR-tasks: there exist efficient *query-answering procedures*— a *top-down* approach and a *bottom-up* evaluation. In addition we can not only handle ground queries but also queries with variables and compute *answer-substitutions*.

It turns out that more advanced KR-tasks can not be properly handled with definite programs. Therefore we extend this basic class of programs by additional features like *negation-as-finite-failure*, *default-negation*, *explicit negation*, *preferences*, and *disjunction*. The need for these extensions is motivated by suitable examples and the corresponding semantics are also discussed.

Clearly, the more expressive the respective class of programs under a certain semantics is, the less efficient are potential query-answering methods. This point will be illustrated and discussed for every extension. By well-known recursion-theoretic results, it is obvious that there do not exist complete query-answering procedures for the general case where variables and function symbols are allowed. Nevertheless we consider it an important topic of further research to extract *feasible* classes of programs where answer-substitutions can be computed.

2 DEFINITE LOGIC PROGRAMS

In this section we consider the most restricted class of programs: *definite* logic programs, programs without any negation at all. All the extensions of this basic class we will introduce later contain at least some kind of negation (and perhaps additional features). But here we also allow the occurrence of free variables as well as function symbols.

In Section 2.1 we introduce as a representative for the *top-down* approach the SLD-resolution. Section 2.2 presents the main competing approach of SLD: *bottom-up evaluation*. This approach is used in the database community and it is efficient when additional assumptions are made (*finiteness-assumptions, no function symbols*). In Section 2.3 we consider the influence and appropriateness of Herbrand models and their underlying intuition. Finally in Section 2.4 we present and discuss two important examples in KR: *reasoning in inheritance hierarchies* and *reasoning about actions*. Both examples clearly motivate the need of extending definite programs by a kind of *default-negation* “not”.

First some notation used throughout this paper. A language \mathcal{L} consists of a set of relation symbols and a set of function symbols (each symbol has an associated arity). Nullary functions are called constants. Terms and atoms are built from \mathcal{L} in the usual way starting with variables, applying function symbols and relation-symbols.

Instead of considering arbitrary \mathcal{L} -formulae, our main object of interest is a program:

Definition 2.1 (Definite Logic Program).

A *definite* logic program consists of a finite number of *rules* of the form

$$A \leftarrow B_1, \dots, B_m,$$

where A, B_1, \dots, B_m are positive atoms (containing possibly free variables). We call A the *head* of the rule and B_1, \dots, B_m its *body*. The comma represents conjunction \wedge .

We can think of a program as formalizing our knowledge about the world and how the world behaves. Of course, we also want to derive new information, i.e. we want to ask queries:

Definition 2.2 (Query).

Given a definite program we usually have a definite query in mind that we want to be solved. A definite query Q is a conjunction of positive atoms $C_1 \wedge \dots \wedge C_l$ which we denote by

$$?- C_1, \dots, C_l.$$

These C_i may also contain variables. Asking a query Q to a program P means asking for all possible substitutions Θ of the variables in Q such that $Q\Theta$ follows from P . Often, Θ is also called an answer to Q . Note that $Q\Theta$ may still contain free variables.

Note that if a program P is given, we usually assume that it also determines the underlying language \mathcal{L} , denoted by \mathcal{L}_P , which is generated by exactly the symbols occurring in P . The set of all these atoms is called the *Herbrand base* and denoted by $B_{\mathcal{L}_P}$ or simply B_P . The corresponding set of all ground terms is the *Herbrand universe*. Another important notion that we are not explaining in detail here is that of *unification*. Given two atoms A and B with free variables we can ask if we can

compute two substitutions Θ_1, Θ_2 for the variables such that

$$A\Theta_1 \text{ is identical to } B\Theta_2,$$

or if we can decide that this is not possible at all. In fact, if the two atoms are *unifiable* we can indeed compute a *most general unifier*, called mgU (see [Lloyd, 1987]). The mgU Θ is a substitution defined on the set of variables occurring in both A and B such that $A\Theta$ is identical to $B\Theta$.

This will be important in our framework because if an atom appears as a subgoal in a query, we may want to determine if there are rules in the program whose heads unify with this atom.

How are our programs related to classical predicate logic? Of course, we can map a program-rule into classical logic by interpreting “ \leftarrow ” as material implication “ \supset ” and universally quantifying. This means we view such a rule as the following universally quantified formula

$$B_1 \wedge \dots \wedge B_m \supset A.$$

However, as we will see later, there is a great difference: a logic program-rule takes some orientation with it. This makes it possible to formulate the following principle as an underlying intuition of all semantics of logic programs:

Principle 1 (Orientation).

If a ground atom A does not unify with some head of a program rule of P , then this atom is considered to be false. In this case we say that “not A ” is derivable from P to distinguish it from classical $\neg A$.

The orientation principle is nothing but a weak form of *negation-by-failure*. Given an intermediate goal not A , we first try to prove A . But if A does not unify with any head, A fails and this is the reason to derive not A .

2.1 Top-Down

SLD-Resolution² is a special form of Robinson’s general resolution rule. While Robinson’s rule is complete for full first order logic, SLD is complete for definite logic programs (see Theorem 2.1 on page 9). We do not give a complete definition of SLD-resolution (see [Lloyd, 1987]) but rather prefer to illustrate its behaviour on the following example.

Example 2.1 (SLD-Resolution).

Let the program P_{SLD} consist of the following three clauses

- (1) $p(x, z) \leftarrow q(x, y), p(y, z)$
- (2) $p(x, x)$
- (3) $q(a, b)$

²SL-resolution for Definite clauses. SL-resolution stands for Linear resolution with Selection function.

Definite programs have the nice feature that the intersection of all Herbrand-models exists and is again a Herbrand model of P . It is denoted by M_P and called the least Herbrand-model of P . Note that our original aim was to find substitutions Θ such that $Q\Theta$ is derivable from the program P . This task as well as M_P is closely related to SLD:

Theorem 2.1 (Soundness and Completeness of SLD).

The following properties are equivalent:

- $P \models \forall Q\Theta$, i.e. $\forall Q\Theta$ is true in all models of P ,
- $M_P \models \forall Q\Theta$,
- SLD computes an answer τ that subsumes³ Θ wrt Q .

Note that not any correct answer is computed, only the most general one is (which of course subsumes all the correct ones).

The main feature of SLD-resolution is its *goal-orientedness*. SLD automatically ensures (because it starts with the Query) that we consider only those rules that are relevant for the query to be answered. Rules that are not at all related are simply not considered in the course of the proof.

2.2 Bottom-Up

We mentioned in the last section the least Herbrand model M_P . The bottom-up approach can be described as computing this least Herbrand model from below. We start first with rules with empty bodies (in our example these are all instantiations of rules (2) and (3)). We get as facts all atoms that are in the heads of rules with empty bodies (namely $p(a, a)$, $p(b, b)$, $q(a, b)$ in Example 2.1 on page 7). In the next round we use the facts that we computed before and try to let the rules “fire”, i.e. when their bodies are true, we add their heads to the atoms we already have (this gives us $p(a, b)$).

To be more precise we introduce the immediate consequence operator T_P which associates to any Herbrand model another Herbrand model.

Example 2.2 (T_P).

Given a definite program P let $T_P : 2^{B_P} \mapsto 2^{B_P}; \mathcal{I} \mapsto T_P(\mathcal{I})$

$$T_P(\mathcal{I}) := \{A \in B_P : \begin{array}{l} \text{there is an instantiation of a rule in } P \\ \text{s.t. } A \text{ is the head of this rule and all} \\ \text{body-atoms are contained in } \mathcal{I} \end{array}\}$$

It turns out that T_P is monotone and continuous so that (by a general theorem of Knaster-Tarski) the least fixpoint is obtained after ω steps. Moreover we have

³i.e. $\exists \sigma : Q\tau\sigma = Q\Theta$.

Theorem 2.2 (T_P and M_P).

$$M_P = T_P \uparrow^\omega = \text{lfp}(T_P).$$

This approach is especially important in database applications, where the underlying language does not contain function symbols (DATALOG) — this ensures the Herbrand universe to be finite. Under this condition the iteration stops after finitely many steps. In addition, rules of the form

$$p \leftarrow p$$

do not make any problems. They simply can not be applied or do not produce anything new. Note that in the top-down approach, such rules give rise to infinite branches! Later, elimination of such rules will turn out to be an interesting property. We therefore formulate it as a principle:

Principle 2 (Elimination of Tautologies).

Suppose a program P has a rule which contains the same atom in its body as well as in its head (i.e. the head consists of exactly this atom). Then we can eliminate this rule without changing the semantics.

Unfortunately, such a bottom-up approach has two serious shortcomings. First, the goal-orientedness of SLD-resolution is lost: we are always computing the whole M_P , even those facts that have nothing to do with the query. The reason is that in computing T_P we do not take into account the query we are really interested in. Second, in any step facts that are already computed before are recomputed again. It would be more efficient if only new facts were computed. Both problems can be (partially) solved by appropriate refinements of the naive approach:

- *Semi-naive* bottom-up evaluation ([Bry, 1990; Ullman, 1989a]),
- *Magic Sets* techniques ([Beeri and Ramakrishnan, 1991; Ullman, 1989b]).

2.3 Herbrand-Models and the underlying language

Usually when we represent some knowledge in first order logic or even in logic programs, it is understood that the underlying language is given exactly by the symbols that occur in the formal theory. Suppose we have represented some knowledge about the world as a theory T in a language \mathcal{L} . Classical predicate logic formalizes the notion of a formula ϕ *entailed* by the theory T . This means that ϕ is true in all models of T (we denote this set by $\text{MOD}(T)$). Why are we considering all models? Doesn't it make sense to look only at Herbrand models, i.e. to models generated by the underlying language? After all we are not interested in models that contain elements which are not representable as terms in our language. These requirements are usually called *unique names assumption* and *domain closure assumption*:

Definition 2.3 (UNA and DCA).

Let a language \mathcal{L} be given. We understand by the *unique names assumption* the

restriction to those models \mathcal{I} , where syntactically different ground \mathcal{L} -terms t_1, t_2 are interpreted as nonidentical elements: $t_1^{\mathcal{I}}$ is not identical to $t_2^{\mathcal{I}}$.

By the *domain closure assumption* we mean the restriction to those models \mathcal{I} where for any element a in \mathcal{I} there is a \mathcal{L} -term t that represents this element: $a = t^{\mathcal{I}}$.

As an example, in Theorem 2.1 on page 9 of Section 2.1 we referred to M_P , the least Herbrand model of P . The reason that the first equivalence in this theorem holds is given by the fact that for universal theories T and existential formulae ϕ the following holds

$$\text{MOD}(T) \models \phi \text{ iff } \text{Herb}_{\mathcal{L}}\text{-MOD}(T) \models \phi.$$

In our particular case, where T is a definite program P , we can even replace $\text{Herb}_{\mathcal{L}}\text{-MOD}(T)$ in the above equation by the single model M_P .

This last result does not hold in general. But what happens if we nevertheless are interested in only the Herbrand-models of a theory T (and therefore automatically⁴ assume UNA and DCA)? At first sight one can argue that such an approach is much simpler: in contrast to *all* models we only need to take care about the very specific Herbrand models. But it turns out that determining the truth of a formula in all Herbrand models is a much more complex task (namely Π_1^1 -complete) than to determine if it is true in all models. This latter task is also undecidable in general, but it is recursively enumerable, i.e. Π_1^0 -complete. The fact that this task is recursively enumerable is the content of the famous completeness theorem of Gödel, where “truth of a formula in all models” is shown to be equivalent to deriving this formula in a particular axiomatization of the predicate calculus of first order. We refer to the appendix (Section A.1 and Section A.2) where the necessary notions are introduced.

But we have still a problem with Theorem 2.1 on page 9 in our restricted setting:

Example 2.3 (Universal Query Problem).

Consider the program $P := p(a)$, the query $Q := p(x)$ and the empty substitution $\Theta := \epsilon$. We have

- $M_P \models \forall x p(x)$,
- but SLD only computes the answer x/a .

Przymusiński called this *the universal query problem*.

There are essentially two solutions to avoid this behaviour: to use a language which is *rich enough* (i.e. contains sufficiently many terms, not only those occurring in the program P itself) or to consider arbitrary models, not only Herbrand models. Both approaches have been followed in the literature but they are beyond the scope of this paper.

⁴The only difference between Herbrand models and models satisfying UNA and DCA is that the interpretation of terms is uniquely determined in Herbrand models. It is required that a term “ $f(t_1, \dots, t_n)$ ” is interpreted in a Herbrand model \mathcal{I} as “ $f^{\mathcal{I}}(t_1^{\mathcal{I}}, \dots, t_n^{\mathcal{I}})$ ”.

2.4 Why going beyond Definite Programs?

So far we have a nice query-answering procedure, SLD-Resolution, which is goal-oriented as well as sound and complete with respect to general derivability. But note that up to now we are not able to derive any *negative* information. Not even our queries allow this. From a very pragmatic viewpoint, we can consider “not A ” to be derivable if A is not. Of course, this is not sound with respect to classical logic but it is with respect to M_P .

In KR we do not only want to formulate negative queries, we also want to express *default-statements* of the form

Normally, unless something abnormal holds, then ψ implies ϕ .

Such statements were the main motivation for nonmonotonic logics, like Default Logic or Circumscription (see Section A.3 and Section A.4 of the appendix). How can we formulate such a statement as a logic program? The most natural way is to use negation “not”

$$\phi \leftarrow \psi, \text{ not } ab$$

where ab stands for *abnormality*. Obviously, this forces us to extend definite programs by *negative* atoms, we call them default atoms.

A typical example for such statements occurs in Inheritance Reasoning. We take the following example from [Baral and Gelfond, 1994]:

Example 2.4 (Inheritance Hierarchies).

Suppose we know that birds typically fly and penguins are non-flying birds. We also know that Tweety is a bird. Now an agent is hired to build a cage for Tweety. Should the agent put a roof on the cage? After all it could be still the case that Tweety is a penguin and therefore can not fly, in which case we would not like to pay for the unnecessary roof. But under normal conditions, it should be obvious that one should conclude that Tweety is flying.

A natural axiomatization is given as follows:

$$\begin{array}{lll}
 P_{Inheritance} : & flies(x) & \leftarrow bird(x), \quad \text{not } ab(r_1, x) \\
 & bird(x) & \leftarrow penguin(x) \\
 & ab(r_1, x) & \leftarrow penguin(x) \\
 & make_top(x) & \leftarrow flies(x)
 \end{array}$$

together with some particular facts, like e.g. $bird(Tweety)$ and $penguin(Sam)$. The first rule formalizes our default knowledge, while the third formalizes that the default rule should not be applied in abnormal or exceptional cases. In our example, it expresses the famous *specificity principle* which says that more specific knowledge should override more general one ([Touretzky, 1986; Touretzky *et al.*, 1988; Horty *et al.*, 1990]).

For the query “ $make_top(Tweety)$ ” we expect the answer “yes” while for “ $make_top(Sam)$ ” we expect the answer “no”.

Another important KR task is to formalize knowledge for reasoning about action. We again consider a particular important instance of such a task, namely *temporal projection*. The overall framework consists in describing the initial state of the world as well as the effects of all actions that can be performed. What we want to derive is how the world looks like after a sequence of actions has been performed.

Example 2.5 (Temporal Projection: Yale-Shooting Problem).

We distinguish between three sorts⁵ of variables:

- situation variables: S, S', \dots ,
- fluent variables: F, F', \dots ,
- action variables: A, A', \dots .

The initial situation is denoted by the constant s_0 , and the two-ary function symbol $res(A, S)$ denotes the situation that is reached when in situation S the action A has been performed. The relation symbol $holds(F, S)$ formalizes that the fluent F is true in situation S .

For the YSP there are three actions (*wait*, *load* and *shoot*) and two fluents (*alive* and *loaded*). Initially a turkey called *Fred* is alive. We then load a gun, wait and shoot. The effect should be that Fred is dead after this sequence of actions. The common-sense argument from which this should follow is the

Law of Inertia: Things normally tend to stay the same.

Using our intuition from the last example, a natural formalization is given as follows:

$$\begin{array}{ll}
 P_{YSP} : \quad holds(F, res(A, S)) & \leftarrow \quad holds(F, S), \quad not \ ab(r_1, A, F, S) \\
 \quad \quad holds(loaded, res(load, S)) & \leftarrow \\
 \quad \quad ab(r_1, shoot, alive, S) & \leftarrow \quad holds(loaded, S) \\
 \quad \quad holds(alive, s_0) & \leftarrow
 \end{array}$$

Such a straightforward formalization leads in most versions of classical non-monotonic logic to the unexpected result, that Fred is not necessarily dead. But obviously we expect to derive $holds(alive, res(load, s_0))$ and

$$not \ holds(alive, res(shoot, res(wait, res(load, s_0))))$$

Up to now we only have stated some very “natural” axiomatizations of given knowledge. We have motivated that something like default-negation “not” should be added to definite programs in order to do so and we have explicitly stated the answers to particular queries. What is still missing are solutions to the following very important problems

⁵To be formally correct we have to use many-sorted logic. But since all this could also be coded in predicate logic by using additional relation symbols, we do not emphasize this fact. We also understand that instantiations are done in such a way that the sorts are respected.

- *How should an appropriate query answering mechanism handling default-negation “not” look like?*
- *What is the formal semantics that such a procedural mechanism should be checked against?*

Such a semantics is certainly not classical predicate logic because of the default character of “not”—not is not classical \neg . Both problems will be considered in detail in Section 3.

2.5 What is a Semantics?

In the last subsections we have introduced two principles (*Orientation* and *Elimination of Tautologies*) and used the term *semantics of a program* in a loose, imprecise way. We end this section with a precise notion of what we understand by a semantics.

As a first attempt, we can view a semantics as a mapping that associates to any program a set of positive atoms and a set of default atoms. In the case of SLD-Resolution the positive atoms are the ground instances of all derivable atoms. But sometimes we also want to derive default atoms (like in our two examples above). Our *Orientation-Principle* formalizes a minimal requirement for deriving such default-atoms.

Of course, we also want that a semantics SEM should *respect* the rules of P , i.e. whenever SEM makes the body of a rule true, then SEM should also make the head of the rule true. But it can (and will) happen that a semantics SEM does not always decide *all* atoms. Some atoms A are not derivable nor are their default-counterparts $\text{not } A$. This means that a semantics SEM can view the body of a rule as being *undefined*.

This already happens in classical logic. Take the theory

$$T := \{(A \wedge B) \supset C, \neg A \supset B\}.$$

What are the atoms and negated atoms derivable from T , i.e. true in all models of T ? No positive atom nor any negated atom is derivable! The classical semantics therefore makes the truthvalue of $A \wedge B$ undefined in a sense.

Suppose a semantics SEM treats the body of a program rule as undefined. What should we conclude about the head of this rule? We will only require that this head is not treated as false by SEM—it could be true or undefined as well. This means that we require a semantics to be compatible with the program *viewed as a 3-valued theory*—the three values being “*true*”, “*false*” and “*undefined*”. For the understanding it is not necessary to go deeper into 3-valued logic. We simply note that we interpret “ \leftarrow ” as the Kleene-connective which is *true* for “*undefined* \leftarrow *undefined*” and *false* for “*false* \leftarrow *undefined*”.

Our discussion shows that we can view a semantics SEM as a 3-valued model of a program. In classical logic, there is a different viewpoint. For a given theory

T we consider there the set of all classical models $\text{MOD}(T)$ as the semantics. The intersection of all these models is of course a 3-valued model of T , but $\text{MOD}(T)$ contains more information. In order to formalize the notion of semantics as general as possible we define

Definition 2.4 (SEM).

A semantics SEM is a mapping from the class of all programs into the powerset of the set of all 3-valued structures. SEM assigns to every program P a set of 3-valued models of P :

$$\text{SEM}(P) \subseteq \text{MOD}_{3\text{-val}}^{\mathcal{L}_P}(P).$$

This definition covers both the classical viewpoint (classical models are 2-valued and therefore special 3-valued models) as well as our first attempt in the beginning of this subsection. Later on, in most cases we will be really interested only in Herbrand models.

Formally, we can associate to any semantics SEM in the sense of Definition 2.4 two entailment relations

sceptical: $\text{SEM}^{\text{scept}}(P)$ is the set of all atoms or default atoms that are true in *all* models of $\text{SEM}(P)$.

credulous: $\text{SEM}^{\text{cred}}(P)$ is the set of all atoms or default atoms that are true in *at least one* model of $\text{SEM}(P)$.

In this tutorial we only consider the sceptical viewpoint. Also, to facilitate notation, we will not formally distinguish between SEM and $\text{SEM}^{\text{scept}}$. In cases where by definition SEM can only contain a single model (like in the case of well-founded semantics) we will omit the outer brackets and write

$$\text{SEM}(P) = M$$

instead of $\text{SEM}(P) = \{M\}$. We will also slightly abuse notation and write $l \in \text{SEM}(P)$ as an abbreviation for $l \in M$ for all $M \in \text{SEM}(P)$.

3 ADDING DEFAULT-NEGATION

In the last section we have illustrated that logic programs with negation are very suitable for KR—they allow a natural and straightforward formalization of default-statements. The problem still remained to define an appropriate semantics for this class and, if possible, to find efficient query-answering methods. Both points are addressed in this section.

We can distinguish between two quite different approaches:

LP-Approach: This is the approach taken mainly in the logic programming community. There one tried to stick as close as possible to SLD-resolution and

treat negation as “finite-failure”. This resulted in an extension of SLD, called SLDNF-resolution, a procedural mechanism for query answering. For a nice overview, we refer to [Apt and Bol, 1994].

NML-Approach: This is the approach suggested by non-monotonic reasoning people. Here the main question is “*What is the right semantics?*” I.e. we are looking first for a semantics that correctly fits to our intuitions and treats the various KR-Tasks in the right (or appropriate) way. It should allow us to jump to conclusions even when only little information is available. Here it is of secondary interest how such a semantics can be implemented with a procedural calculus. Interesting overviews are [Minker, 1993; Minker, 1996] and [Dix, 1995c; Dix *et al.*, 2001a].

The LP-approach is dealt with in Section 3.1. It is still very near to classical predicate logic—default negation is interpreted as *finite-failure*. To get a stronger semantics, we interpret “not” as *failure* in Section 3.2. The main difference is that the principle *Elimination of Tautologies* holds. We then introduce a principle GPPE which is related to partial evaluation. In KR one can see this principle as allowing for definitional extensions—names or abbreviations can be introduced without changing the semantics.

All these principles do not yet determine a unique semantics—there is still room for different semantics and a lot of them have been defined in the last years. We do not want to present the whole zoo of semantics nor to discuss their merits or shortcomings. We refer the reader to the overview articles [Apt and Bol, 1994] and [Dix, 1995c] and the references given therein. We focus on the two main competing approaches that still have survived. These are the wellfounded semantics WFS (Section 3.3) and the stable semantics STABLE (Section 3.4). Finally, in Section 3.5 we discuss complexity and expressibility results for the semantics presented so far.

3.1 Negation-as-Finite-Failure

The idea of negation treated as *finite-failure* can be best illustrated by still considering definite programs, but queries containing default-atoms. How should we handle such default-atoms by modifying our SLD-resolution? Let us try this:

- If we reach a default-atom “not *A*” as a subgoal of our original query, we keep the current SLD-tree in mind and start a new SLD-tree by trying to solve “*A*”.
- If this succeeds, then we falsified “not *A*”, the current branch is failing and we have to backtrack and consider a different subquery.
- But it can also happen that the SLD-tree for “*A*” is *finite with only failing branches*. Then we say that *A* *finitely fails*, we turn back to our original

SLD-tree, consider the subgoal “not A ” as successfully solved and go on with the next subgoal in the current list.

It is important to note that an SLD-tree for a positive atom can fail without *being finite*. The SLD-tree for the program consisting of the single rule $p \leftarrow p$ with respect to the query p is infinite but failing (it consists of one single infinite branch). In Figure 1 on page 8 the leftmost branch is also failing but infinite.

Although this idea of *finite-failure* is very procedural in nature, there is a nice model theoretical counterpart—Clark’s completion $comp(P)$ ([Clark, 1978]). The idea of Clark was that a program P consists not only of the implications, but also of the information that *these are the only ones*. Roughly speaking, he argues that one should interpret the “ \leftarrow ”-arrows in rules as equivalences “ \equiv ” in classical logic. We do not give the exact definitions here, as they are very complex; in the non-propositional case, a symbol for equality, together with axioms describing it⁶, has to be introduced. However, for the propositional case, $comp(P)$ is obtained from P by just

1. collecting all given clauses with the same head into one new “clause” with this respective head and a disjunctive body (containing all bodies of the old clauses), and
2. replacing the implication-symbols “ \leftarrow ” by “ \equiv ”.

Definition 3.1 (Clark’s Completion $comp(P)$).

Clark’s semantics for a program P is given by the set of all classical models of the theory $comp(P)$.

We can now see the classical theory $comp(P)$ as the information contained in the program P . $comp(P)$ is like a sort of closed world assumption applied to P . We are now able to derive negative information from P by deriving it from $comp(P)$. In fact, the following soundness and completeness result for definite programs P and definite queries $Q = \bigwedge_i A_i$ (consisting of only *positive* atoms) holds:

Theorem 3.1 (COMP and Fair FF-Trees).

The following conditions are equivalent:

- $comp(P) \models \forall \neg Q$
- Every fair SLD-tree for P with respect to Q is finitely failed.

Note that in the last theorem we did not use default negation but classical negation \neg because we just mapped all formulae into classical logic. We need the fairness assumption to ensure that the selection of atoms is reasonably well-behaving: we want that every atom or default-atom occurring in the list of preliminary goals will eventually be selected.

⁶CET: Clark’s Equational Theory. $CET(\mathcal{L}_P)$ axiomatizes the equality theory of all Herbrand(\mathcal{L}_P)-models. See [Mancarella *et al.*, 1988; Shepherdson, 1992] for the problem of equality and the underlying language.

But even this result is still very weak—after all we want to handle not only negative queries but programs containing default-atoms. From now on we consider programs with default-atoms in the body. We usually denote them by

$$A \leftarrow B^+ \wedge \text{not } B^-,$$

where B^+ contains all the positive body atoms and $\text{not } B^-$ all default atoms “not C ”.

Our two motivating examples in Section 2.4 contain such default atoms. This gives rise to an extension of SLD, called SLDNF, which treats negation as *Finite-Failure*

$$\text{SLDNF} = \text{SLD} + \text{not } L \text{ succeeds, if } L \text{ finitely fails.}$$

The precise definitions of SLDNF-*resolution*, *tree*, etc. are very complex: we refer to [Lloyd, 1987; Apt, 1990]. Apt and Bol gave interesting improved versions of these notions: see [Apt and Bol, 1994, Section 3.2]. In order to get an intuitive idea, it is sufficient to describe the following underlying principle:

Principle 3 (A “Naive” SLDNF-Resolution).

If in the construction of an SLDNF-tree a default-atom $\text{not } L_{ij}$ is selected in the list $\mathcal{L}_i = \{L_{i1}, L_{i2}, \dots\}$, then we try to prove L_{ij} .

If this *fails finitely* (it fails because the generated subtree is finite and failing), then we take $\text{not } L_{ij}$ as proved and we go on to prove $L_{i(j+1)}$.

If L_{ij} succeeds, then $\text{not } L_{ij}$ fails and we have to backtrack to the list \mathcal{L}_{i-1} of preliminary subgoals (the next rule is applied: “backtracking”).

Does SLDNF-resolution properly handle Examples 2.4 on page 12 and 2.5 on page 13? It does indeed:

Inheritance: The query $\text{make_top}(\text{Tweety})$ generates an SLD-tree with one main branch, the nodes of which are:

flies(*Tweety*),
bird(*Tweety*), $\text{not } ab(r_1, \text{Tweety})$,
 $\text{not } ab(r_1, \text{Tweety})$,
Success.

The third node has a sibling-node *penguin*(*Tweety*), $\text{not } ab(r_1, \text{Tweety})$ which immediately fails because *Tweety* does not unify with *Sam*. The *Success*-node is obtained from $\text{not } ab(r_1, \text{Tweety})$ because the corresponding SLD-tree for $ab(r_1, \text{Tweety})$ fails finitely (this tree consists only of $ab(r_1, \text{Tweety})$ and *penguin*(*Tweety*)).

YSP: The crucial query is

$$?- \text{not holds}(\text{alive}, \text{res}(\text{shoot}, \text{res}(\text{wait}, \text{res}(\text{load}, s_0))))).$$

So we consider $?\text{-holds}(\text{alive}, \text{res}(\text{shoot}, \text{res}(\text{wait}, \text{res}(\text{load}, s_0))))$. Again the SLD-tree for this query consists mainly of one branch: the nodes are




Figure 2. The Floundering-Problem

obtained from the query by applying successively the first program rule (*law of inertia*). By evaluation of the holds-predicate, we eventually arrive at the fact $holds(alive, s_0)$ and the “not *ab*” predicates remain to be solved. For any of these predicates we again have to consider separate SLD-trees. But for $ab(r_1, shoot, alive, res(wait, res(load, s_0)))$ it is easy to see that the associated tree already finitely fails (because it generates the subgoal “not $ab(r_1, wait, loaded, res(load, s_0))$ ” the corresponding SLD-tree of which immediately finitely fails) and therefore, since no backtracking is possible, the tree for

$$?- holds(alive, res(shoot, res(wait, res(load, s_0))))$$

finitely fails and our original query succeeds: *Fred is dead*.

Up to now it seems that SLDNF-resolution solves all our problems. It handles our examples correctly, and is defined by a procedural calculus strongly related to SLD. There are two main problems with SLDNF:

- SLDNF can not handle free variables in negative subgoals,
- SLDNF is still too weak for knowledge representation.

The latter problem is the most important one. By looking at a particular example, we will motivate in Section 3.2 the need for a stronger semantics. This will lead us in the remaining sections to the wellfounded and the stable semantics.

For the rest of this section we consider the first problem, known as the *floundering problem*. This problem will also occur later in implementations of the wellfounded or the stable semantics. We consider the program $P_{flounder}$ consisting of the three facts

$$p(c, c), q(b), r(f(c)).$$

Our query is $? - p(x, c), \text{not } q(x), r(f(x))$, that is, we are interested in instantiations of x such that the query follows from the program. The situation is illustrated in Figure 2 on the page before. Let us suppose that we always select the first atom or default-atom: it is underlined in the sequel. The SLDNF-tree of this trivial example is linear and has three nodes: the first node is the query itself

$$? - \underline{p(x, c)}, \text{not } q(x), r(f(x)),$$

the second node is $? - \text{not } \underline{q(c)}, r(f(c))$. Now, we enter the negation-as-failure mode and ask $? - q(c)$. This query immediately fails (the generated tree exists, is finite and fails) so that we give back the answer “yes, the default atom $\text{not } q(c)$ succeeds and can be skipped from the list”. The last node is $? - \underline{r(f(c))}$ which immediately succeeds.

Note that in the last step, the test for $? - q(c)$ has to be finished before the tree can be extended. If we get no answer, the SLDNF-tree simply does not exist: this can not happen with SLD-trees.

So far everything was fine. But what happens if we select the second atom in the first step

$$? - p(x, c), \underline{\text{not } q(x)}, r(f(x))?$$

Example 3.1 (Floundering).

We again consider the program $P_{flounder}$ consisting of the three facts

$$p(c, c), q(b), r(f(c)).$$

Our query is $? - p(x, c), \text{not } q(x), r(f(x))$, and in the first step we will select the second default-atom, i.e. one with a free variable. Thus we enter the negation-as-failure mode with the query $? - \text{not } q(x)$. In this case, x may be instantiated to b so that we have to give back the answer “no, the default-atom $\text{not } q(x)$ fails” and the whole query will fail. This is because SLDNF treats the subgoal as “ $\forall x \text{not } q(x)$ ” instead of “ $\exists x \text{not } q(x)$ ” which is intended. There exist approaches to overcome this shortcoming by treating negation as *constructive negation*: see [Chan, 1988; Chan and Wallace, 1989; Drabent, 1994].

In the classical SLDNF-resolution *negation-as-finite-failure is only a test, no bindings are produced*. On the one hand this may be considered a shortcoming, on the other hand, it makes the SLDNF procedure more tractable. Note that the problem to decide if a given program flounders is undecidable [Börger, 1987]. See also [Shepherdson, 1991] for more unsolvable problems related to SLDNF.

SLDNF is a procedural mechanism. It would be nice to have a modeltheoretical counterpart. In Theorem 3.1 on page 17 we already related a restricted form of finite failure to Clark’s completion. We will see later that $\text{comp}(P)$ is inconsistent even in cases where we would not expect it. Therefore Fitting [Fitting, 1985] introduced a three-valued formulation $\text{comp}_3(P)$ of the original completion. Kunen ([Kunen, 1987]) then proved in the propositional case *SLDNF is sound and complete with respect to $\text{comp}_3(P)$* .

In the predicate logic case, SLDNF is not complete but it is always correct [Shepherdson, 1988, Theorem 39]) with respect to $comp_3(P)$: given a query Q ,

- if SLDNF succeeds with answer Θ , then $comp_3(P) \models_3 \forall Q \Theta$, and
- if SLDNF fails, then $comp_3(P) \models_3 \neg \exists Q$.

This correctness result is also the reason for the incompleteness of SLDNF with respect to two-valued $comp(P)$. It states that any formula derivable by SLDNF is a three-valued consequence of $comp_3(P)$. But, since there are two-valued consequences of a theory that are not three-valued ones (three-valued logic is weaker than two-valued logic), SLDNF can not be complete. Extensions of the above completeness result to certain subclasses of predicate logic programs require severe restrictions on the syntactic form of P . To define these syntactic restrictions, we need the notion of the dependency-graph:

Definition 3.2 (Dependency-Graph \mathcal{G}_P).

For a logic program P with negation, the *dependency graph* \mathcal{G}_P is a finite directed graph whose vertices are the predicate symbols from P . There is a *positive* (respectively *negative*) edge from R to R' iff there is a clause in P with R in its head and R' occurring positively (respectively negative) in its body.

We also say

- R *depends on* R' if there is a path in \mathcal{G}_P from R to R' (by definition, R depends on itself),
- R *depends positively* (resp. *negatively*) on R' if there is a path in \mathcal{G}_P from R to R' containing only *positive* edges (resp. *at least one negative* edge). (by definition R depends positively on itself),
- R *depends evenly* (resp. *oddly*) on R' if there is a path in \mathcal{G}_P from R to R' containing an *even* (resp. *odd*) number of negative edges (by definition R depends evenly on itself).

The following properties of a program P turn out to be very important:

- stratified*: no predicate depends negatively on itself⁷,
- strict*: there are no dependencies that are both even and odd,
- call-consistent*: no predicate depends oddly on itself⁸,
- allowedness*: every variable occurring in a clause must occur in at least one positive atom of the body of that clause.

Strictness and allowedness turn out to be the most important restrictions that imply completeness results for SLDNF:

⁷or: there are no cycles containing at least one *negative* edge.

⁸or: there are no odd cycles.

Prog. P	Semantics	Completeness
allowed + hierarchical allowed + stratified allowed	$comp(P)$ $comp(P)$ $comp_3(P)$	yes, no recursion at all yes, if $P \cup \{\leftarrow A\}$ strict yes, w.r.t. \vdash_3
allowed + call-consistent	$comp(P)$	yes, if $P \cup \{\leftarrow A\}$ strict: $comp(P) \vdash \forall A$ iff $comp_3(P) \vdash_3 \forall A$

Table 1. Completeness for SLDNF

While *strictness* excludes situations of the form $p(x) \leftarrow q(x), p(c) \leftarrow \neg q(f(c'))$, *allowedness* excludes constructs of the form $equal(x, x) \leftarrow$ and also solves the floundering-problem.

Strictness implies that $comp_3(P)$ and $comp(P)$ are equivalent [Kunen, 1989]

$$comp_3(P) \models_3 \forall Q \Theta \text{ iff } comp(P) \models \forall Q \Theta.$$

Table 1 gives an overview of the different completeness results. Note that the query A is always considered to be allowed.

Much work was done in LP (see [Decker and Cavedon, 1990; Barbuti and Martelli, 1986; Stärk, 1994]) to find other syntactically characterizable classes, for which SLDNF is also *complete*.

3.2 Negation-as-Failure

Let us first illustrate that SLDNF answers quite easily our requirements of a semantics SEM (stated explicitly in Definition 2.4 on page 15). We can formulate these requirements as two program-transformations (they will be used later for computing a semantics). We call them *reductions* for obvious reasons.

Principle 4 (Reduction).

Suppose we are given a program P with possibly default-atoms in its body. If a ground atom A does not unify with any head of the rules of P , then we can delete in every rule any occurrence of “not A ” without changing the semantics.

Dually, if there is an instance of a rule of the form “ $B \leftarrow$ ” then we can delete all rules that contain “not B ” in their bodies.

It is obvious that SLDNF “implements” these two reductions automatically. The weakness of SLDNF for knowledge representation is in a sense inherited from SLD. When we consider rules of the form “ $p \leftarrow p$ ”, then SLD resolution gets

into an infinite loop and no answer to the query $?- p$ can be obtained. This has often the effect that when we enter into negation-as-failure mode, the SLD-tree to be constructed is not finite, although it is not successful and therefore should be considered as failed.

Let us discuss this point with a more serious example.

Example 3.2 (The Transitive Closure).

Assume we are given a graph consisting of nodes and edges between some of them. We want to know which nodes are reachable from a given one. A natural formalization of the property “reachable” would be

$$reachable(x) \leftarrow edge(x, y), reachable(y).$$

What happens if we are given the following facts

$$edge(a, b), edge(b, a), edge(c, d)$$

and $reachable(c)$? Of course, we expect that neither a nor b are reachable because there is no path from c to either a or b .

But SLDNF-Resolution does not derive “not $reachable(a)$ ”!

How does this result relate to Theorem 3.1 on page 17? Note that our query has exactly the form as required there. Clark’s completion of our program rule is

$$reachable(x) \equiv (x \doteq c \vee \exists y (reachable(y) \wedge edge(y, x)))$$

from which, together with our facts about the edge-relation, $\neg reachable(a)$ is indeed not derivable. This is due to the wellknown fact that transitive closure is not expressible in first order predicate logic.

Note also that our Principle 2 on page 10 does not help, because it simply does not apply. It turns out that we can augment our two principles by a third one, that constitutes together with them a very nice calculus handling the above example in the right way. This principle is related to *partial evaluation*, hence its name GPPE⁹. Let us motivate this principle with the last example. The query “not $reachable(a)$ ” leads to “ $reachable(a) \leftarrow edge(a, b), reachable(b)$ ” and “ $reachable(b)$ ” leads to “ $reachable(b) \leftarrow edge(b, a), reachable(a)$ ”. Both rules can be seen as definitions for $reachable(a)$ and $reachable(b)$ respectively. So it should be possible to replace in these rules the body atoms of $reachable$ by their definitions. Thus we obtain the two rules

$$\begin{aligned} reachable(a) &\leftarrow edge(a, b), edge(b, a), reachable(a) \\ reachable(b) &\leftarrow edge(b, a), edge(a, b), reachable(b) \end{aligned}$$

that can both be eliminated by applying Principle 2 on page 10. So we end up with a program that does neither contain $reachable(a)$ nor $reachable(b)$ in one

⁹Generalized Principle of Partial Evaluation

of the heads. Therefore, according to Principle 1 on page 7 both atoms should be considered false. The precise formulation of this principle is as follows:

Principle 5 (GPPE, [Brass and Dix, 1994; Sakama and Seki, 1994]).

We say that a semantics SEM satisfies GPPE, if the following transformation does not change the semantics. *Replace a rule $A \leftarrow B^+ \wedge \text{not } B^-$ where B^+ contains a distinguished atom B by the rules*

$$A \leftarrow (B^+ \setminus \{B\}) \cup B_i^+ \wedge \text{not } (B^- \cup B_i^-) \quad (i = 1, \dots, n)$$

where $B \leftarrow B_i^+ \wedge \text{not } B_i^-$ ($i = 1, \dots, n$) are all rules with head B .

Note that any semantics SEM satisfying GPPE and Elimination of Tautologies can be seen as extending SLD by doing some *Loop-checking*. We will call such semantics *NMR-semantics* in order to distinguish them from the classical *LP-semantics* which are based on SLDNF or variants of Clark's completion $\text{comp}(P)$:

- *NMR-Semantics* = *SLDNF* + *Loop-check*.

The following, somewhat artificial example illustrates this point.

Example 3.3 (COMP vs. NMR).

$$\begin{array}{ll} P_{NMR} : & p \leftarrow p \\ & q \leftarrow \text{not } p \end{array} \qquad \begin{array}{ll} P'_{NMR} : & p \leftarrow p \\ & q \leftarrow \text{not } p \\ & r \leftarrow \text{not } r \end{array}$$

$$\begin{array}{ll} \text{comp}(P_{NMR}) : & p \equiv p \\ & q \equiv \neg p \end{array} \qquad \begin{array}{ll} \text{comp}(P'_{NMR}) : & p \equiv p \\ & q \equiv \neg p \\ & r \equiv \neg r \end{array}$$

$$\begin{array}{ll} ?-q: & \text{No (COMP).} \\ & \text{Yes (NMR).} \end{array} \qquad \begin{array}{ll} ?-p: & \text{Yes (COMP).} \\ & \text{No (NMR).} \end{array}$$

For both programs, the answers of the completion-semantics do not match our NMR-intuition! In the case of P_{NMR} we expect q to be derivable, since we expect $\text{not } p$ to be derivable: the only possibility to derive p is the rule $p \leftarrow p$ which, obviously, will never succeed. But $q \notin Th(\{q \equiv \neg p\}) = \text{comp}(P_{NMR})!$ In the case of P'_{NMR} we expect p not to be derivable, for the same reason: the only possibility to derive p is the rule $p \leftarrow p$. But $p \in Fml = Th(\{r \equiv \neg r\}) = \text{comp}(P'_{NMR})!$

Note that the answers of the completion-semantics agree with the mechanism of SLDNF: $p \leftarrow p$ represents a loop. The completion of P' is inconsistent: this led Fitting to consider the three-valued version of $\text{comp}(P)$ mentioned at the end of Section 3.1. This approach avoids the inconsistency (the query $?-p$ is not answered “yes”) but it still does not answer “no” as we would like to have.

The last principle in this section is related to *subsumption*: we can get rid of non-minimal rules by simply deleting them.

Principle 6 (Subsumption).

In a program P we can delete a rule $A \leftarrow B^+ \wedge \text{not } B^-$ whenever there is another rule $A \leftarrow B'^+ \wedge \text{not } B'^-$ with

$$B'^+ \subseteq B^+ \text{ and } B'^- \subseteq B^-.$$

As a simple example, the rule $A \leftarrow B, C, \text{not } D, \text{not } E$ is subsumed by the 3 rules $A \leftarrow C, \text{not } D, \text{not } E$ or $A \leftarrow B, C, \text{not } E$ and by $A \leftarrow C, \text{not } E$.

3.3 The Wellfounded Semantics: WFS

The wellfounded semantics, originally introduced in [Van Gelder *et al.*, 1988], is the *weakest* semantics satisfying our 4 principles (see [Brass and Dix, 1999; Brass and Dix, 1998; Dix, 1995b]). We call a semantics

$$\text{SEM}_1 \text{ weaker than } \text{SEM}_2, \text{ written } \text{SEM}_1 \leq_k \text{SEM}_2,$$

if for all programs P and all atoms or default-atoms l the following holds:

$$\text{SEM}_1(P) \models l \text{ implies } \text{SEM}_2(P) \models l.$$

I.e. all atoms derivable from SEM_1 with respect to P are also derivable from SEM_2 . The notion \leq_k refers to the knowledge ordering in three-valued logic. This is a nice theorem and gives rise to the following definition:

Theorem 3.2 (WFS, [Brass and Dix, 1999]).

There exists the weakest semantics satisfying our four principles Elimination of Tautologies, Reduction, Subsumption and GPPE. This semantics is called wellfounded semantics WFS.

It can also be shown, that for propositional programs, our transformations can be applied to compute this semantics.

Theorem 3.3 (Confluent Calculus for WFS, [Brass and Dix, 1998]).

The calculus consisting of these four transformations is confluent, i.e. whenever we arrive at an irreducible program, it is uniquely determined. The order of the transformations does not matter.

For finite propositional programs, it is also terminating: any program P is therefore associated a unique normalform $\text{res}(P)$. The wellfounded semantics of P can be read off from $\text{res}(P)$ as follows

$$\text{WFS}(P) = \{A : A \leftarrow \in \text{res}(P)\} \cup \{\text{not } A : A \text{ is in no head of } \text{res}(P)\}$$

We note that the size of the residual program is in general *exponential* in the size of the original program. Recently it was shown in [Brass *et al.*, 2001b] how a

small modification of the residual program, which still satisfies the nice characterization of computing WFS as given in Theorem 3.3 on the page before, results in a *polynomial* computation.

Therefore the wellfounded semantics associates to every program P with negation a set consisting of atoms and default-atoms. This set is a 3-valued model of P . It can happen, of course, that this set is empty. But it is always consistent, i.e. it does not contain an atom A and its negation $\text{not } A$. Moreover, it extends SLDNF: whenever SLDNF derives an atom or default-atom and does not flounder, then WFS derives it as well. Therefore the two examples of Section 2.4 are handled in the right way. But also for Example 3.2 on page 23 we get the desired answers.

Let us discuss whether *every sequence of program transformations terminates*, i.e. if our calculus is *strongly terminating*. Already the simple program consisting of just the loop “ $A \leftarrow A$ ” shows a problem. Applying GPPE leads to the same program, so GPPE can be applied infinitely often without leading to the residual program. Of course, in this particular case an application of the *Elimination of Tautologies* immediately leads to the residual program. But still another problem can occur¹⁰:

$$P_{loop} : \begin{array}{l} A \leftarrow B \\ B \leftarrow A \\ C \leftarrow A, \neg C \end{array}$$

If we apply $GPPE_A$ to the third clause, this clause is replaced by $C \leftarrow B, \neg C$. We can now apply $GPPE_B$ again to this clause and get the original program. So we have an oscillation

$$GPPE_B \circ GPPE_A(P_{loop}) = P_{loop}$$

To summarize, not every sequence of transformations terminates. But a simple additional property will ensure this.

Definition 3.3 (Fair Sequences).

We call a sequence of program transformations *fair*, if in the corresponding sequence of programs

1. every positive body-atom is eventually removed (either by removing the whole clause using a suitable transformation or by an application of GPPE), and
2. every tautology clause is eventually removed (either by applying *Elimination of Tautologies* or another suitable transformation).

Theorem 3.4 (Strong Termination for Fair Sequences).

Our calculus of transformations is strongly terminating for fair sequences of transformations. Such fair sequences therefore always lead to the residual program.

¹⁰brought to our attention by Frieder Stolzenburg

As we said above, loop-checking is in general undecidable. Therefore WFS is in the most general case where variables and function-symbols are allowed, undecidable. Only for finite propositional programs it is decidable. In fact, it is of quadratic complexity (see Section 3.5).

Let us end this section with another example, which contains negation.

Example 3.4 (Van Gelder's Example).

Assume we are describing a two-players game like checkers. The two players alternately move a stone on a board. The moving player wins when his opponent has no more move to make. We can formalize that by

- $wins(x) \leftarrow move_from_to(x, y), \text{ not } wins(y)$

meaning that

- the situation x is won (for the moving player A), if he can lead over¹¹ to a situation y that can never be won for B.

Assume we also have the facts

$$move_from_to(a, b), move_from_to(b, a) \text{ and } move_from_to(b, c).$$

Our query to this program P_{game} is $?- wins(b)$. Here we have no problems with floundering, but using SLDNF we get an infinite sequence of oscillating SLD-trees (none of which finitely fails).

WFS, however, derives the right results

$$WFS(P_{game}) = \{ \text{not } wins(c), wins(b), \text{ not } wins(a) \}$$

which matches completely with our intuitions.

3.4 The Stable Semantics: STABLE

We defined WFS as the weakest semantics satisfying our four principles. This already indicates that there are even stronger semantics. One of the main competing approaches is the stable semantics STABLE. The stable semantics associates to any program P a set of 2-valued models, like classical predicate logic. STABLE satisfies the following property, in addition to those that have been already introduced:

Principle 7 (Elimination of Contradictions).

Suppose a program P has a rule which contains the same atom A and $\text{not } A$ in its body. Then we can eliminate this rule without changing the semantics.

This principle can be used, in conjunction with the others to define the stable semantics

¹¹With the help of a regular move, given by the relation $move_from_to/2$.

Theorem 3.5 (STABLE, [Brass and Dix, 1997]).

There exists the weakest semantics satisfying our five principles Elimination of Tautologies, Reduction, Subsumption, GPPE and Elimination of Contradictions.

If a semantics SEM satisfies *Elimination of Contradictions* it is based on 2-valued models ([Brass and Dix, 1997]). The underlying idea of STABLE is that any atom in an intended model should have a definite reason to be true or false. This idea was made explicit in [Bidoit and Froidevaux, 1991a; Bidoit and Froidevaux, 1991b] and, independently, in [Gelfond and Lifschitz, 1988]. We use the latter terminology and introduce the Gelfond-Lifschitz transformation: for a program P and a model $N \subseteq B_P$ we define

$$P^N := \{rule^N : rule \in P\}$$

where $rule := A \leftarrow B_1, \dots, B_n, \text{not } C_1, \dots, \text{not } C_m$ is transformed as follows

$$(rule)^N := \begin{cases} A \leftarrow B_1, \dots, B_n, & \text{if } \forall j : C_j \notin N, \\ \mathbf{t}, & \text{otherwise.} \end{cases}$$

Note that P^N is always a *definite* program. We can therefore compute its least Herbrand model M_{P^N} and check whether it coincides with the model N with which we started:

Definition 3.4 (STABLE).

N is called a *stable model*¹² of P iff $M_{P^N} = N$.

What is the relationship between STABLE and WFS? We have seen that they are based on rather identical principles.

- Stable models N extend WFS: $l \in \text{WFS}(P)$ implies $N \models l$.
- If $\text{WFS}(P)$ is two-valued, then $\text{WFS}(P)$ is the unique stable model.

But there are also differences. We refer to Example 3.4 on the page before and consider the program P consisting of the clause

$$\text{wins}(x) \leftarrow \text{move_from_to}(x, y), \text{not } \text{wins}(y)$$

together with the following facts: $\text{move_from_to}(a, b)$, $\text{move_from_to}(b, a)$, as well as $\text{move_from_to}(b, c)$, and $\text{move_from_to}(c, d)$. In this particular case we have two stable models: $\{\text{wins}(a), \text{wins}(c)\}$ and $\{\text{wins}(b), \text{wins}(c)\}$ and therefore

$$\text{WFS}(P) = \{\text{wins}(c), \text{not } \text{wins}(d)\} = \bigcap_{\mathcal{N} \text{ a stable model of } P} \mathcal{N}.$$

This means that the 3-valued wellfounded model is exactly the set of all atoms or default-atoms true in all stable models. But this is not always the case, as the program of $P_{\text{splitting}}$ shows:

¹²Note that we only consider Herbrand models.

Example 3.5 (Reasoning by cases).

$$\begin{array}{l}
 P_{\text{splitting}} : \quad a \leftarrow \text{not } b \\
 \quad \quad \quad b \leftarrow \text{not } a \\
 \quad \quad \quad p \leftarrow a \\
 \quad \quad \quad p \leftarrow b
 \end{array}$$

Although neither a , nor b can be derived in any semantics based on two-valued models (as STABLE for example), the disjunction $a \vee b$, thus also p , is true. In this way the example is handled by the completion semantics, too. $WFS(P)$, however, is empty; if the WFS cannot decide between a or $\text{not } a$, then a is undefined.

The main differences between STABLE and WFS are

- STABLE is not always consistent,
- STABLE does not allow for a goal-oriented implementation.

The inconsistency comes from odd, negative cycles

$$STABLE(p \leftarrow \text{not } p) = \emptyset.$$

The idea to consider 2-valued models for a semantics necessarily implies its inconsistency ([Brass and Dix, 1997]). Note that $WFS(p \leftarrow \text{not } p) = \{\emptyset\}$ which is quite different! Sufficient criteria for the existence of stable models are contained in [Dung, 1992; Fages, 1993].

That STABLE does not allow for a Top-Down evaluation is a more serious drawback and has nothing to do with inconsistency. This behaviour led Dix to define the notion of *Relevance* and *Modularity* (see Section 7.1 and [Dix, 1992a; Dix, 1992b; Dix, 1995b]). Bry reinvented *Modularity* (he termed it *compositionality*) and argued that a semantics should satisfy it.

Example 3.6 (STABLE is not Goal-Oriented).

$$\begin{array}{l}
 P_{\text{rel}(a)} : \quad a \leftarrow \text{not } b \quad P : \quad a \leftarrow \text{not } b \\
 \quad \quad \quad b \leftarrow \text{not } a \quad \quad \quad b \leftarrow \text{not } a \\
 \quad \quad \quad \quad \quad \quad \quad \quad \quad p \leftarrow \text{not } p \\
 \quad \quad \quad \quad \quad \quad \quad \quad \quad p \leftarrow a
 \end{array}$$

$P_{\text{rel}(a)}$ is the subprogram of P that consists of all rules that are relevant to answer the query $?- a$. It has two stable models $\{a\}$ and $\{b\}$ — a is not true in all of them. But the program P has the unique stable model $\{p, a\}$, so a is true in all stable models of P .

The last example shows that the truthvalue of an atom a also depends on atoms that are totally unrelated with a ! This is considered a drawback of STABLE by many people. Note that a straightforward modification of STABLE is not possible ([Dix and Müller, 1994b; Dix and Müller, 1994c]).

We end this section with another description of WFS and STABLE that will be useful in later sections. It was introduced in [Baral and Subrahmanian, 1991; Baral and Subrahmanian, 1992]:

Definition 3.5 (Antimonotone Operator γ_P).

For a program P and a set $N \subset B_P$ we define an operator γ_P mapping Herbrand-structures to Herbrand structures:

$$\gamma_P(N) := M_{PN}.$$

It is easy to see that γ_P is antimonotone. Therefore its twofold application γ^2 is monotone ([Tarski, 1955]).

Obviously, the stable models of a program P are exactly the fixpoints of γ_P . This is just a reformulation of Definition 3.4 on page 28. WFS is related to γ as follows

Theorem 3.6 (WFS and γ^2).

A positive atom A is in WFS(P) iff $A \in lfp(\gamma_P^2)$. A default-atom $\text{not } A$ is in WFS(P) iff $A \notin gfp(\gamma_P^2)$:

$$WFS(P) = lfp(\gamma_P^2) \cup \{ \text{not } A : A \notin gfp(\gamma_P^2) \}.$$

Atom or default-atoms that do occur in neither of the two sets are undefined.

3.5 Complexity and Expressibility

In this section we collect some complexity results for the semantics considered so far. The reason why NMR-semantics are in the general case (free variables and function symbols) undecidable is strongly related to loop-checking. Let us consider the program

$$P(x) \leftarrow P(f(x))$$

or, equivalently, the infinite propositional program

$$p_0 \leftarrow p_1, p_1 \leftarrow p_2, \dots, p_i \leftarrow p_{i+1}, \dots$$

Any NMR-semantics should derive “not $P(t)$ ” (resp. “not p_i ”) for all terms t , but a procedure to detect such infinite loops is impossible in general. Our principles *GPPE* and *Elimination of Tautologies* can detect *finite* loops.

From a model theoretic point of view it is easy to *define* a semantics that derives “not $P(t)$ ”: we could just take all minimal Herbrand models as the intended semantics. Of course, this does not change the general undecidability.

For the exact terminology, definitions and results presented in this section we refer the interested reader to the following interesting overviews [Schlipf, 1990; Schlipf, 1992; Cadoli and Schaerf, 1993]. Further results are contained in [Eiter *et al.*, 1993; Sacca, 1993; Chomicki and Subrahmanian, 1990; Eiter and Gottlob, 1993b].

While Table 2 on the next page treats the complexity Table 3 on page 33 treats the expressibility problem. Some general explanations are appropriate.

	Complexity	
	<i>1. ord. prog.</i> <i>(with functions)</i>	<i>prop. prog.</i> <i>(no variables)</i>
M_P <i>(P is Horn)</i>	A: Σ_1^0 -compl. not A: Π_1^0 -compl.	linear in $ P $
M_P^{supp} <i>(P is stratified)</i>	arithm.-compl. <i>(M_P^{supp} is Σ_n^0)</i>	linear in $ P $
COMP	Π_1^1 -compl. over \mathbb{N}	co-NP -compl.
COMP₃	Π_1^1 -compl. over \mathbb{N}	linear in $ P $
STABLE	Π_1^1 -compl. over \mathbb{N}	co-NP -compl.
REG-SEM	Π_1^1 -compl. over \mathbb{N}	co-NP -compl.
WFS	Π_1^1 -compl. over \mathbb{N}	linear in $\#At \times P $
WFS'	Π_1^1 -compl. over \mathbb{N}	co-NP -compl.
WFS⁺	Π_1^1 -compl. over \mathbb{N}	co-NP -compl.

Table 2. Complexity of Non-Disjunctive Semantics

Table 2

We consider the complexity of deciding if a *given ground atom or default-atom is contained in the respective semantics* (i.e. if it is true in all *intended models*).

For the first column, we consider arbitrary first-order programs with function symbols. We therefore get undecidability results of varying strength. Since we restrict to Herbrand models, we can assume (by standard recursive encoding techniques, like *Gödel-numberings*) that all models have universes which are subsets of the natural numbers \mathbb{N} . The completeness results mean that for every set of the respective complexity class there is a program that defines this set under the respective sceptical semantics. Unless indicated otherwise, there is no difference between deciding ground *atoms* or ground *negated atoms*.

For the second column, we consider *propositional* programs. Hence we get decidable problems of various degrees. We denote by $|P|$ the total length of the program and by $\#At$ the number of distinct proposition letters in P . See also [Ben-Eliyahu and Dechter, 1992; Schlipf, 1992; Imielinski, 1991; Marek *et al.*, 1992; Witteveen, 1991b] for more results on the complexity of propositional programs.

Table 3

Here we consider the *expressibility* (or expressive power) of first order programs *without* function symbols. The idea is to distinguish between EDB-relations (relations that do not appear in the head of a program) and IDB-relations (which are contained in some heads). For a given program P we can view any instance \mathcal{D} of the (finite) EDB-relations as an *input* argument and then compute the (finite) IDB-relations (the *output*) under the respective sceptical semantics. So we are asking

What are the relations expressible with logic programs under certain semantics?

Roughly speaking, a relation R over finite EDB's \mathcal{D} (i.e. for every finite \mathcal{D} is associated a relation $R^{\mathcal{D}}$ on \mathcal{D}) is *expressible* if there is a program P containing an IDB-symbol r s. t. for every relational database \mathcal{D} and tuple \underline{t} corresponding to r :

$$r(\underline{t}) \in SEM(P + \mathcal{D}) \text{ if and only if } R(\underline{t}) \text{ holds in } \mathcal{D}.$$

This is the classical notion of *expressibility* ([Schlipf, 1990; Eiter *et al.*, 1993]).

We are in particular interested to express all relations of some complexity class (note that the complexity is always with respect to the finite relational database as input, the program is fixed). It is well-known that the relations *inductively definable* over \mathcal{D} , we denote them by $\mathbf{IND}(\mathcal{D})$ (or simply \mathbf{IND} to avoid the explicit occurrence of the EDB), is a strict subclass of the relations that are *polynomial* over \mathcal{D} (see [Barwise, 1975; Moschovakis, 1974; Gurevich, 1988]).

It is worth noting that in the general predicate logic case, all semantics are highly undecidable. The entries for comp and comp_3 are to be understood as restricted to Herbrand models.

In the propositional case, WFS is of quadratic complexity (a folklore result—for a proof see [Witteveen, 1991a]), while STABLE is **co-NP**-complete. The low complexity of WFS can be traced back to Dowling and Gallier's result whereby satisfiability of Horn clauses can be tested in linear time ([Dowling and Gallier, 1984]). In Dowling and Gallier's approach it is actually a minimal model of a Horn theory that is computed in linear time. Since minimal models of Horn theories are equivalent to closures of rules without negation the result is directly applicable to well-founded semantics for logic programs with default-atoms.

As far as expressibility is concerned, STABLE is more expressive: all **co-NP**-relations can be expressed, while WFS can only describe all inductively definable relations. As an example, STABLE can express the satisfiability problem. WFS is not able to do this (unless the polynomial hierarchy collapses).

4 ADDING EXPLICIT NEGATION

So far we have considered programs with one special type of negation, namely default negation. Default negation is particularly useful in domains where complete

	Expressibility <i>1. ord. prog.</i> <i>(no functions)</i>
M_P <i>(P is Horn)</i>	$\subset \text{IND}$ (thus $\subset \text{P}$)
COMP	= co-NP
COMP₃	= IND (thus $\subset \text{P}$)
STABLE	= co-NP
REG-SEM	= Π_2^P
WFS	= IND (thus $\subset \text{P}$)
WFS'	= IND (thus $\subset \text{P}$)
WFS⁺	= IND (thus $\subset \text{P}$)

Table 3. Expressibility of Non-Disjunctive Semantics

positive information can be obtained. For instance, if one wants to represent flight connections from Budapest to the US it is very convenient to represent all existing flights and to let default negation handle the derivation of negative information. There are domains, however, where the lack of positive information cannot be assumed to support (or support with enough strength) that this information is false. In such domains it becomes important to distinguish between cases where a query does not succeed and cases where the negated query succeeds. The following example was used by McCarthy to illustrate the issue. Assume one wants to represent the rule: cross the railroad tracks if no train is approaching. The straightforward representation of this rule with default negation would be

$$crosstracks \leftarrow \text{not } train$$

It seems obvious that in many practical settings the use of such a rule would not lead to intended behaviour, in fact it might even have disastrous consequences. What seems to be needed here is the possibility of using a different negation symbol representing a stronger form of negation. This new negation—we will call it explicit negation—should be true only if the corresponding negated literal can actually be derived. We will use the classical negation symbol \neg to represent explicit negation. The track crossing rule will be represented as

$$crosstracks \leftarrow \neg train$$

The idea is that this latter rule will only be applicable if $\neg train$ has been proved, contrary to the first rule which is applicable whenever $train$ is not provable.

In the next subsection we will shortly discuss that explicit negation is not (or should not be) *classical* negation and how it should interfere with default negation. In the two following subsections we will generalize the semantics STABLE and WFS, respectively, to programs with explicit negation.

4.1 Explicit vs. Classical and Strong Negation

First we define the language we are using more precisely.

Definition 4.1 (Extended Logic Program).

An extended logic program consists of rules of the form

$$c \leftarrow a_1, \dots, a_n, \text{not } b_1, \dots, \text{not } b_m$$

where the a_i, b_j and c are literals, i.e., either propositional atoms or such atoms preceded by the classical negation sign. The symbol “not” denotes negation by failure (default negation), “ \neg ” denotes explicit negation.

We have already motivated the need of a second kind of negation “ \neg ” different from “not”. What should the semantics of “ \neg ” be? Should it be just like in classical logic? Note that classical negation satisfies the law of excluded middle

$$A \vee \neg A.$$

The following example taken from [Alferes *et al.*, 1996] shows that classical negation is sometimes inappropriate for KR-tasks.

Example 4.1 (Behaviour of Classical Negation).

Suppose an employer has several candidates that apply for a job. Some of them are clearly qualified while others are not. But there may also be some candidates whose qualifications are not clear and who should therefore be interviewed in order to find out about their qualifications. If we express the situation by

$$\text{hire}(X) \leftarrow \text{qualified}(X) \text{ and } \text{reject}(X) \leftarrow \neg \text{qualified}(X)$$

then, interpreting “ \neg ” as classical negation, we are forced to derive that every candidate must either be hired or rejected! There is no room for those that should be interviewed. Also, applying the law of excluded middle has a highly non-constructive flavor.

Let us now consider again the example $\text{crosstracks} \leftarrow \neg \text{train}$ from the beginning of this section. Suppose that we replace $\neg \text{train}$ by free_track . We obtain

$$\text{crosstracks} \leftarrow \text{free_track}.$$

From this program, “not *crosstracks*” will be derivable for any semantics. Therefore we should make sure that “not *crosstracks*” is also derivable from

$$\text{crosstracks} \leftarrow \neg \text{train}$$

After all, the second program is obtained from the first one by a simple syntactic operation. This means we have to make sure that default negation “not” treats positive and negated atoms symmetrically.

Such a negation, we will call it *explicit* will be introduced in the next two subsections. Note that Gelfond/Lifschitz called the negation they introduced in their stable semantics *classical*, although it is not classical in the sense that we just discussed. Sometimes explicit negation is also called *strong* negation and denotes still a variant of our explicit negation. In [Alferes *et al.*, 1996] the authors introduce both a strong and explicit negation and discuss their relation with classical and default negation at length.

4.2 STABLE for Extended Logic Programs

The extension of STABLE to extended logic programs is based on the notion of answer sets which generalize the original notion of stable models in a rather straightforward manner. Let us first introduce some useful notation. We say a rule $r = c \leftarrow a_1, \dots, a_n, \text{not } b_1, \dots, \text{not } b_m \in P$ is defeated by a literal l iff $l = b_i$ for some $i \in \{1, \dots, m\}$. We say r is defeated by a set of literals X if X contains at least one literal that defeats r . Furthermore, we call the rule obtained by deleting default negated antecedents from r the monotonic counterpart of r and denote it with $Mon(r)$. We also apply Mon to sets of rules with the obvious meaning.

Definition 4.2 (X -reduct).

Let P be an extended logic program, X a set of literals. The X -reduct of P is the set of rules

$$P^X := \{Mon(r) : r \text{ not defeated by } X\}.$$

Note that the only difference between the definition of P^N in Section 3.4 and this definition is that the new one handles literals and not only than atoms.

The definition of stable models (Definition 3.4) was based on the minimal model of the reduct. Since the reduct now may contain explicit negation we need a different notion here, namely a notion of consequence:

Definition 4.3 (Consequences of Rules).

Let R be a set of rules without default negation. $Cn(R)$ denotes the smallest set of literals that is

1. closed under R , and
2. logically closed, i.e., either consistent or equal to the set of all literals.

Definition 4.4 (γ_P).

Let P be a logic program, X a set of literals. Define an operator γ_P as follows:

$$\gamma_P(X) = Cn(P^X)$$

X is an answer set of P iff $X = \gamma_P(X)$.

A literal l is a consequence of a program P under the new semantics, denoted $l \in STABLE(P)$, iff l is contained in all answer sets of P .

It is not difficult to see that for programs without explicit negation stable models and answer sets coincide. Here is an example involving both types of negation. The example describes the strategy of a certain college for awarding scholarships to its students. It is taken from [Baral and Gelfond, 1994]:

$$\begin{aligned}
 P_{el} : \quad (1) \quad & eligible(x) \quad \leftarrow \quad highGPA(x) \\
 (2) \quad & eligible(x) \quad \leftarrow \quad minority(x), fairGPA(x) \\
 (3) \quad & \neg eligible(x) \quad \leftarrow \quad \neg fairGPA(x), \neg highGPA(x) \\
 (4) \quad & interview(x) \quad \leftarrow \quad not\ eligible(x), not\ \neg eligible(x)
 \end{aligned}$$

Assume in addition to the rules above the following facts about Anne are given:

$$fairGPA(Anne), \neg highGPA(Anne)$$

We obtain exactly one answer set, namely

$$\{fairGPA(Anne), \neg highGPA(Anne), interview(Anne)\}$$

Anne will thus be interviewed before a decision about her eligibility is made. If we use the above rules together with the facts

$$minority(Mike), fairGPA(Mark)$$

then the program entails $eligible(Mike)$.

The following results are taken from [Lifschitz, 1996]:

Lemma 4.1 (Program Types).

Let P be an extended logic program. P satisfies exactly one of the following conditions:

- P has no answer sets,
- the only answer set for P is \perp lit,
- P has an answer set, and all its answer sets are consistent.

A program is consistent if the set of its consequences is consistent, and inconsistent otherwise. The former corresponds to the first two cases listed in the proposition, the latter to the third case.

We say that a set X of literals is supported by P if, for each literal $l \in X$, there exists a rule $l \leftarrow a_1, \dots, a_n, not\ b_1, \dots, not\ b_m$ in P such that

1. $\{a_1, \dots, a_n\} \subseteq X$, and
2. $\{b_1, \dots, b_m\} \cap X = \emptyset$.

Lemma 4.2 (Properties of answer sets).

Let P be an extended logic program. The following properties hold:

- Any consistent answer set for P is supported by P .
- If X and Y are answer sets of P and $X \subseteq Y$ then $X = Y$.
- Each element of a consistent answer set of P is a head literal¹³ of P .

From the last property it follows immediately that every consequence of P is a head literal of P whenever P is consistent. We would finally like to mention the following theorem:

Theorem 4.3 (Head Consistency).

If the set of head literals of an extended program P is consistent then every answer set of P is consistent.

Note that a program satisfying the conditions of the last theorem can still be inconsistent since it may have no answer set at all.

It should be noted that extended logic programs under answer set semantics can be reduced to general logic programs as follows: for any predicate p occurring in a program P we introduce a new predicate symbol p' of the same arity representing the explicit negation of p . We then replace each occurrence of $\neg p$ in the program with p' , thus obtaining the general logic program P' . It can be proved that a consistent set of literals S is an answer set of P iff the set S' is a stable model of P' , where S' is obtained from S by replacing $\neg p$ with p' .

4.3 WFS for Extended Logic Programs

We now show how the second major semantics for general logic programs, WFS, can be extended to logic programs with explicit negation. For our purposes the characterization of WFS given in Theorem 3.6 on page 30 will be useful. WFS is based on a particular three-valued model. To simplify our presentation in this section we will restrict ourselves to the literals which are true in this three-valued model. The literals which are false will be left implicit. They can be added in a canonical way as follows: let T , the set of true literals, be defined as the least fixpoint of a monotone operator composed of two antimonotone operators $op_1 op_2$. Then the literals which are false in the three-valued model are exactly those which are not contained in $op_2(T)$. Given this canonical extension to the full three-valued model we can safely leave the false literals implicit from now on.

We will first present a formulation which can be found in various papers, e.g. [Baral and Gelfond, 1994; Lifschitz, 1996]. We then slightly modify this formulation to obtain stronger results. We finally discuss a further modification by Pereira and Alferes.

¹³A head literal of a program P is the head of a rule of P (see also Principle 1 on page 7 and Definition 6.3 on page 51).

Like answer set semantics well-founded semantics for extended logic programs can be based on the operator γ_P . However, the operator is used in a totally different way. Since γ_P is anti-monotone the operator $\Gamma_P = (\gamma_P)^2$ is monotone. According to the famous Knaster-Tarski theorem [Tarski, 1955] every monotone operator has a least fixpoint. We can thus define

Definition 4.5 (WFS for extended programs).

Let P be an extended logic program. The set of well-founded conclusions of P , denoted $WFS(P)$, is the least fixpoint of Γ_P .

The fixpoint can be approached from below by iterating Γ_P on the empty set. In case P is finite this iteration is guaranteed to actually reach the fixpoint.

The intuition behind this use of the operator is as follows: whenever γ_P is applied to a set of literals X known to be true it produces the set of all literals that are still potentially derivable. Applying it to such a set of potentially derivable literals it produces a set of literals known to be true, often larger than the original set X . Starting with the empty set and iterating until the fixpoint is reached thus produces a set of true literals.

We first want to illustrate this using an example without explicit negation:

$$P : \begin{array}{l} (1) \quad b \leftarrow \text{not } a \\ (2) \quad c \leftarrow \text{not } b \\ (3) \quad e \leftarrow \text{not } d \\ (4) \quad d \leftarrow \text{not } e \end{array}$$

In the beginning we know nothing about derivable literals, i.e., we start with empty set X . The X -reduct of the program is

$$\begin{array}{l} (1) \quad b \\ (2) \quad c \\ (3) \quad e \\ (4) \quad d \end{array}$$

The set of consequences of this program, or in other words, the literals still considered to be potentially derivable, is thus $\{b, c, d, e\}$. If we now reduce the program with this set we obtain

$$(1) \quad b$$

that is, the first iteration of the two-fold application of γ_P tells us that b is provable.

If we now use $X = \{b\}$ to continue the iteration we obtain the reduced program

$$\begin{array}{l} (1) \quad b \\ (3) \quad e \\ (4) \quad d \end{array}$$

that is $\{b, d, e\}$ is the current set of potential conclusions. Using this set to reduce the program gives us again

$$(1) \quad b$$

We thus have reached the least fixpoint of γ_P^2 and b is the single literal provable under WFS.

It can be shown that every well-founded conclusion is a conclusion under the answer set semantics. Well-founded semantics can thus be viewed as an approximation of answer set semantics.

Unfortunately it turns out that for many programs the set of well-founded conclusions as defined in Definition 4.5 is extremely small and provides a very poor approximation of answer set semantics. Consider the following program P_0 which has also been discussed by Baral and Gelfond [Baral and Gelfond, 1994]:

$$P_0 : \begin{array}{lll} (1) & b & \leftarrow \text{not } \neg b \\ (2) & a & \leftarrow \text{not } \neg a \\ (3) & \neg a & \leftarrow \text{not } a \end{array}$$

The set of well-founded conclusions is empty since $\gamma_{P_0}(\emptyset)$ equals *Lit*, the set of all literals, and the *Lit*-reduct of P_0 contains no rule at all. This is surprising since, intuitively, the conflict between (2) and (3) has nothing to do with $\neg b$ and b .

This problem arises whenever the following conditions hold:

1. a complementary pair of literals is provable from the monotonic counterparts of the rules of a program P , and
2. there is at least one proof for each of the complementary literals whose rules are not defeated by $Cn(P')$, where P' consists of the “strict” rules in P , i.e., those without negation as failure.

In this case well-founded semantics concludes l iff $l \in Cn(P')$. It should be obvious that such a situation is not just a rare limiting case. To the contrary, it can be expected that many commonsense knowledge bases will give rise to such undesired behaviour. Let us consider again our Example 2.4 on page 12 from Section 2.

$$\begin{array}{ll} (1) & fly(x) \leftarrow \text{not } \neg fly(x), bird(x) \\ (2) & \neg fly(x) \leftarrow \text{not } fly(x), penguin(x) \end{array}$$

Assume further that the knowledge base contains the information that Tweety is a penguin bird. Now if neither $fly(Tweety)$ nor $\neg fly(Tweety)$ follows from strict rules in the knowledge base we are in the same situation as with P_0 : well-founded semantics does not draw any “defeasible” conclusion, i.e. a conclusion derived from a rule with default negation in the body, at all.

We want to show that a minor reformulation of the fixpoint operator can overcome this weakness and leads to better results. Consider the following operator

$$\gamma_P^*(X) = Cl(P^X)$$

where $Cl(R)$ denotes the minimal set of literals closed under the (classical) rules R . $Cl(R)$ is thus like $Cn(R)$ without the requirement of logical closedness. Now

define

$$\Gamma_P^*(X) = \gamma_P(\gamma_P^*(X))$$

Again we iterate on the empty set to obtain the well-founded conclusions of a program P which we will denote $WFS^*(P)$.

Consider the effects of this modification on our example P_0 :

$$\gamma_{P_0}^*(\emptyset) = \{a, \neg a, b\}.$$

Rule (1) is contained in the $\{a, \neg a, b\}$ -reduct of P_0 and thus $\Gamma_{P_0}^*(\emptyset) = \{b\}$. Since b is also the only literal contained in all answer sets of P_0 our approximation actually coincides with answer set semantics in this case.

In the Tweety example both $fly(Tweety)$ and $\neg fly(Tweety)$ are provable from the \emptyset -reduct of the knowledge base. However, this has no influence on whether a rule not containing the default negation of one of these two literals in the body is used to produce $\gamma_P^*(\emptyset)$ or not. The effect of the conflicting information about Tweety's flying ability is thus kept local and does not have the disastrous consequences it has in the original formulation of well-founded semantics.

It is not difficult to see that the new monotone operator is equivalent to the original one whenever P does not contain negation as failure. In this case the X -reduct of P , for arbitrary X , is equivalent to P and for this reason it does not make any difference whether to use γ_P or γ_P^* as the operator to be applied first in the definition of Γ_P . The same is obviously true for programs without classical negation: for such programs Cn can never produce complementary pairs of literals and for this reason the logical closedness condition is obsolete.

In the general case the new operator produces more conclusions than the original one:

Lemma 4.4. Let P be an extended logic program. For an arbitrary set of literals X we have

$$\Gamma_P(X) \subseteq \Gamma_P^*(X).$$

It can also be shown that the new operator produces no unwanted results, i.e., that our new semantics can still be viewed as an approximation of answer set semantics.

Lemma 4.5. Let P be an extended logic program. WFS^* is correct wrt. $STABLE$, i.e., $l \in WFS^*(P)$ implies $l \in STABLE(P)$.

An alternative, somewhat stronger approach, was developed by Pereira and Alferes [Pereira and Alferes, 1992; Alferes and Pereira, 1995; Alferes and Pereira, 1996], the semantics WFSX. This semantics implements the intuition that a literal with default negation should be derivable from the corresponding explicitly negated literal. The authors call this the coherence principle. To satisfy the principle they use the seminormal version of a program P , denoted $S(P)$, which is obtained from P by replacing each rule

$$c \leftarrow a_1, \dots, a_n, \text{not } b_1, \dots, \text{not } b_m$$

by the rule

$$c \leftarrow a_1, \dots, a_n, \text{ not } b_1, \dots, \text{ not } b_m, \text{ not } \neg c$$

where $\neg c$ is the complement of c , i.e. $\neg c$ if c is an atom and a if $c = \neg a$. Based on this notion Pereira and Alferes consider the following monotone operator:

$$\Omega_P(X) = \gamma_P^* \gamma_{S(P)}^*(X)$$

The use of the seminormal version of the program in the first application of γ^* guarantees that a literal l is not considered a potential conclusion whenever the complementary literal is already known to be true. In the general case $S(P)^X$ contains fewer rules than P^X . Therefore, fewer literals are considered as potential conclusions and thus more conclusions are obtained in each iteration of the monotone operator. Here is an example [Baral and Gelfond, 1994]:

$$P_{WFSX} : \begin{array}{l} (1) \quad a \leftarrow \text{ not } b \\ (2) \quad b \leftarrow \text{ not } a \\ (3) \quad \neg a \leftarrow \end{array}$$

The original version of WFS does not conclude b . In WFSX the set $X = \{\neg a\}$ is obtained after the first iteration of the monotone operator. Since rule (1) is not contained in the X -reduct of the seminormal version of the program the monotonic counterpart of (2) produces b after the second iteration.

5 ADDING PREFERENCES

In this section we describe how preferences among rules can be taken into account in logic programs with two types of negation. The basic idea is the following: in case of a conflict between rules preferences are used to break ties wherever possible. For programs under answer set semantics this means that some of the answer sets are preferred, the others disregarded. The conclusions of a prioritized logic program are then defined as the intersection of all preferred answer sets. In the general case this leads to a larger set of conclusions.

For well-founded semantics we will take preferences into account by modifying the monotone operator whose least fixpoint defines the well-founded conclusions. Again more well-founded conclusions will be generated than without preferences.

We first give some motivation in Section 5.1. We then describe how preferred answer sets can be defined on the basis of preferences among rules. Our presentation of this subsection is based on [Brewka and Eiter, 1999]. Section 5.3 adds preferences to well-founded semantics and is based on [Brewka, 1996b] as well as [Brewka, 2001]. Subsection 5.4 illustrates the expressiveness of our approach using a decision making example.

Note that we use $r_1 < r_2$ to express that r_1 is preferred over r_2 , that is the smaller rules are the better ones in this section.

5.1 Motivation

Preferences among defaults play a crucial role in nonmonotonic reasoning. One source of preferences that has been studied intensively is specificity [Poole, 1985; Touretzky, 1986; Touretzky *et al.*, 1991]—we already discussed it in Example 2.4 on page 12. In case of a conflict between defaults we tend to prefer the more specific one since this default provides more reliable information. E.g., if we know that students are adults, adults are normally employed, students are normally not employed, we want to conclude “Peter is not employed” from the information that Peter is a student, thus preferring the student default over the conflicting adult default.

Specificity is an important source of preferences, but not the only one, and at least in some applications not necessarily the most important one. In the legal domain it may, for instance, be the case that a more general rule is preferred since it represents federal law as opposed to state law [Prakken, 1993]. In these cases preferences may be based on some basic principles regulating how conflicts among rules are to be resolved.

Also in other application domains, like model based diagnosis or configuration, preferences play a fundamental role. Model based diagnosis uses logical descriptions of the normal behaviour of components of a device together with a logical description of the actually observed behaviour. One tries to assume normal behaviour for as many components as possible. A diagnosis corresponds to a set of components for which these normalcy assumptions lead to inconsistency. Very often a large number of possible diagnoses is obtained. In real life some components are less reliable than others. To eliminate less plausible diagnoses one can give the normalcy assumptions for reliable components of higher priority.

In configuration tasks it is often impossible to achieve all of the design goals. Often one can distinguish more important goals from less important ones. To construct the best possible configurations goals then have to be represented as defaults with different preferences according to their desirability.

Preferences also turn out to be relevant for reasoning about action. For instance, Kakas, Miller and Toni developed an approach where preferences between different types of rules (inertia rules, effect rules etc.) are used to model commonsense reasoning involving actions [Kakas *et al.*, 2001].

Prioritized versions for most of the existing nonmonotonic formalisms have been proposed, e.g., prioritized circumscription [Lifschitz, 1994], hierarchic autoepistemic logic [Konolige, 1988], prioritized default logic [Marek and Truszczyński, 1993; Brewka, 1994; Baader and Hollunder, 1995], prioritized theory revision [Benferhat *et al.*, 1993; Nebel, 1998], or prioritized abduction [Eiter and Gottlob, 1995]. Somewhat surprisingly, at least for some time preferences have received less attention in logic programming. This may be explained by the fact that for a long period, logic programming was mainly conceived as a logical paradigm for declarative programming, and to a less extent as a tool for knowledge representation and reasoning. However, it has become evident that logic programming can

serve as a powerful framework for knowledge representation, cf. [Gelfond, 1994; Baral and Gelfond, 1994]. If logic programming wants to successfully stand this challenge, it must provide the features which have been recognized as indispensable in the context of knowledge representation. One such feature is the possibility to handle specificity and priority of knowledge.

5.2 Preferred Answer Sets

In this section we will briefly describe the approach developed in [Brewka and Eiter, 1999]. For a more general treatment, more detail and more motivation consult the original paper.

Let us first define what we mean by a prioritized logic program. For simplicity we consider only propositional programs in this and the following section, i.e. we assume that programs are ground and finite:

Definition 5.1 (Prioritized Program).

A (propositional) prioritized logic program is a pair

$$\mathcal{R} = (R, <)$$

where R is a finite set of ground rules and $<$ is a strict partial order on P .

The treatment of partially ordered prioritized programs is reduced to that of totally ordered programs as follows:

Definition 5.2 (Preferred Answer Set).

Let $\mathcal{R} = (R, <)$ be a prioritized logic program. A is a preferred answer set of \mathcal{R} iff A is a preferred answer set of $\mathcal{R}' = (R, <')$, where $<'$ is an arbitrary total order on R extending $<$.

What remains to be defined then are preferred answer sets for totally ordered programs. The approach presented here was originally introduced using a new reduct which can be viewed as dual to the Gelfond-Lifschitz reduct. Rather than eliminating negation based on a given set of literals, as the Gelfond-Lifschitz reduct, the new reduct eliminates prerequisites. The idea is that the treatment of preferences can thus be reduced to prerequisite-free programs which can be handled much easier.

The exact definitions are somewhat involved. For the purposes of this overview it is sufficient to give a characterization of preferred answer sets which is based on Proposition 5.1 in [Brewka and Eiter, 1999]. We use the following notation: $pre(r)$ denotes the set of antecedents of the rule r are which are not default negated (called prerequisites of r), $head(r)$ denotes the head of r , and a literal l defeats a rule r if *not* l appears in the body of r . r is said to be a generating rule of answer set A iff $pre(r) \in A$ and r is not defeated by any literal in A .

PROPOSITION 1 (Preferred Answer Sets and Generating Rules).

Let $\mathcal{R} = (R, <)$ be a totally prioritized logic program and let A be an answer set of R . A is a preferred answer set of \mathcal{R} if and only if for each rule $r \in R$ such that

$pre(r) \subseteq A$ and $head(r) \notin A$ there is a generating rule $r' \in R$ such that $r' < r$ and $head(r')$ defeats r .

Here is a simple example illustrating this characterization of preferred answer sets:

- (1) $b \leftarrow \text{not } a$
- (2) $a \leftarrow \text{not } b$

Assume $(1) < (2)$. The program has two answer sets, $A_1 = \{b\}$ and $A_2 = \{a\}$, the former generated by rule (1), the latter by (2). Rule (2) was not applied in A_1 although its prerequisites are in A_1 . This is perfectly ok since rule (1) which is a generating rule of higher priority defeats (2). A_1 is thus a preferred answer set.

Now consider A_2 . Rule (1) was not applied, yet there is no generating rule of A_2 defeating (1) which is of higher priority than (1). For this reason, A_2 is not a preferred answer set.

In [Brewka and Eiter, 1999] it was proven that the approach presented here, contrary to several competing approaches, satisfies two natural principles of preference handling in rule based systems. The paper also contains a discussion of many alternative approaches found in the literature. A more recent comparison of some of these approaches can be found in [Schaub and Wang, 2001].

5.3 Prioritized WFS

We now show how preferences can adequately be added to well-founded semantics. We will use the variant of WFS for extended logic programs which is based on the operator γ^* as described in Section 4.3.

For this section a minor reformulation turns out to be convenient. Instead of using the monotonic counterparts of undefeated rules we will work with the original rules and extend the definitions of the consequence operator Cn and the closure operator Cl accordingly. We simply require that default negated preconditions be neglected, i.e., for an arbitrary set of rules P possibly containing default negation we define $Cn(P) = Cn(Mon(P))$ and $Cl(P) = Cl(Mon(P))$. We can now equivalently characterize γ_P and γ_P^* by the equations

$$\gamma_P(X) = Cn(P_X)$$

$$\gamma_P^*(X) = Cl(P_X)$$

where P_X denotes the set of rules not defeated by X .

As mentioned earlier, the intuition behind well-founded semantics can be described as follows: given a set of literals S already known to be derivable, $\gamma^*(S)$ produces a set of potential conclusions which still might defeat rules in P . The conclusions of rules not defeated by any of the potential defeaters are clearly derivable. Starting with the empty set, we thus generate larger and larger sets S until a fixpoint is reached. The following terminology - somewhat influenced by argumentation theory - reflects this intuition:

Definition 5.3 (Defeated Rules).

Let P be an extended logic program.

- A literal l is an S -potential defeater iff l is in the closure of the rules in P not defeated by S .
- A rule r is S -undefeatable iff r is not defeated by any S -potential defeater.
- A literal l is S -derivable iff l is a consequence of the S -undefeatable rules in P .

It is obvious that l is S -derivable iff $l \in \gamma(\gamma^*(S))$. The least fixpoint of $\gamma\gamma^*$, $WFS(P)$, can thus equivalently be characterized as the least fixpoint of the operator which, given a set S , produces the set of S -derivable literals. It turns out that this reformulation is most adequate for introducing preferences.

To take preferences into account we first introduce a notion of dominance. Intuitively, a rule r dominates a rule r' in the context of a set of literals S if r has higher priority and if the application of r in context S actually defeats r' . As pointed out in [Brewka, 1996a] the second condition is necessary to guarantee that prioritized well-founded semantics is an extension of well-founded semantics. Here is the formal definition

Definition 5.4 (S -dominates r').

Let r and r' be rules, S a set of literals. We say r S -dominates r' iff

1. $r < r'$, and
2. $\text{Cl}(\{r\} \cup \{s : s \text{ is } S\text{-undefeatable}\})$ defeats r' .

For the case of prioritized programs Definition 5.3 becomes

Definition 5.5 (Potential Defeater).

Let $(P, <)$ be a prioritized logic program.

- A literal l is an S -potential r -defeater iff l is in the closure of rules in P which are (1) not defeated by S and (2) not S -dominated by r .
- A rule r is S -safe iff r is not defeated by any S -potential r -defeater.
- A literal l is S -derivable iff l is a consequence of the set of S -safe rules in P .

The definition for prioritized logic programs differs from the one for non-prioritized programs in two respects. Firstly, there is not a single set of potential defeaters for all rules but each rule r has its own set of potential defeaters. Secondly, the rules which are used to derive potential defeaters must satisfy an additional condition: to potentially defeat r a rule must not be dominated by r in context S . Since fewer rules can be used to derive potential defeaters for a rule the safe rules are a superset

of the undefeatable rules. We thus obtain more derivable literals. For the special case where $<$ is empty the two definitions of S -derivable clearly coincide.

The set of S -derivable literals grows monotonically with S . We thus can start as usual with the empty set of literals and iterate the computation of S -derivable formulae until a fixpoint is reached.

Here is a small example illustrating the definition.

- (1) $c \leftarrow \text{not } \neg c, a$
- (2) $\neg c \leftarrow \text{not } c$
- (3) a

Let (1) $<$ (2). Clearly, rule (3) is \emptyset -safe since there is no way of defeating a rule without default negation. But also (1) is \emptyset -safe since the closure of (1) together with the \emptyset -undefeatable rule (3) defeats (2) and thus (1) \emptyset -dominates (2). Therefore, the set of \emptyset -derivable literals is $\{c, a\}$. This set is already the least fixpoint. Note that c is not a well-founded conclusion of the program without priorities.

5.4 An Application: Qualitative Decision Making

In this section we want to discuss a somewhat more realistic example and show how prioritized logic programs can be used for qualitative decision making. Although we considered finite propositional programs only in the last two subsections we will, for simplicity, use rules with variables in this subsection. Since we do not need functions the Herbrand universe is finite and we clearly do not go beyond finite propositional programs. Note that $r_1 < r_2$ for rules with variables means that all ground instantiations of r_1 are preferred over all ground instantiations of r_2 . We also write $R_1 < R_2$ for sets of rules R_1 and R_2 to express that each element of R_1 is preferred over each element of R_2 .

Assume you want to buy a car. You have collected the following information about different types of cars:

$$\begin{aligned} & \text{expensive}(\text{Chevrolet}), \text{safe}(\text{Chevrolet}), \text{safe}(\text{Volvo}), \\ & \text{nice}(\text{Porsche}), \text{fast}(\text{Porsche}) \end{aligned}$$

Your decision which car to buy is based on different criteria. We can use rules corresponding to normal defaults [Reiter, 1980] in order to represent the properties which you consider relevant. Let's assume you like fast and nice cars. On the other hand, your budget does not allow you to purchase a very expensive car. Moreover, you have to take your wife's wishes into account, and she insists on a car which is known to be safe.

- (1) $\neg \text{buy}(x) \leftarrow \text{not } \text{buy}(x), \text{expensive}(x)$
- (2) $\text{buy}(x) \leftarrow \text{not } \neg \text{buy}(x), \text{safe}(x)$
- (3) $\text{buy}(x) \leftarrow \text{not } \neg \text{buy}(x), \text{nice}(x)$

$$(4) \quad buy(x) \leftarrow \text{not } \neg buy(x), fast(x)$$

Since buying more than one car is out of the question these different decision criteria may obviously lead to conflict, and a preference ordering is necessary. Since there is not much you can do about your restricted budget, rule (1) gets highest preference. Moreover, since your wife is very concerned about safety you better give (2) higher priority than (3) and (4). Since there is tremendous traffic on highways in your area, (3) is more important to you than (4). This means we have

$$(1) < (2) < (3) < (4)$$

We still have to represent that you cannot afford more than one car. A straightforward idea would be to use the rule

$$(0) \quad \neg buy(y) \leftarrow buy(x), x \neq y$$

with highest priority. Unfortunately, this does not work. The reason is that the high priority of the instances of this rule would allow us to defeat instances of (2), (3) and (4) even if this is not intended. In our example we would obtain two preferred answer sets, the intended one containing $buy(Volvo)$, but also an unintended one (at least from your wife's point of view) containing $buy(Porsche)$. In the latter, the instance of (2) with $x = Volvo$ would be defeated by the instance of (0) with $x = Porsche$ and $y = Volvo$.

To represent our problem adequately, we have to make sure that the consequences of a certain decision, namely that certain cars are not purchased, do not have higher priority than the decision itself, that is, the decision to buy a specific car. Instead of adding the single rule (0) we therefore represent our criteria for buying a car as pairs of rules. To each rule of the form

$$(r) \quad buy(x) \leftarrow \text{not } \neg buy(x), c_1(x), \dots, c_n(x)$$

we add a second one of the form

$$(r') \quad \neg buy(y) \leftarrow c_1(x), \dots, c_n(x), buy(x), x \neq y$$

with the same priority as (r). In our example we have to add

$$(2') \quad \neg buy(y) \leftarrow safe(x), buy(x), x \neq y$$

$$(3') \quad \neg buy(y) \leftarrow nice(x), buy(x), x \neq y$$

$$(4') \quad \neg buy(y) \leftarrow fast(x), buy(x), x \neq y$$

and use the following preferences

$$(1) < \{(2), (2')\} < \{(3), (3')\} < \{(4), (4')\}$$

Given this information we obtain a single preferred answer set containing $buy(Volvo)$.

This leaves you somewhat dissatisfied since you really like the Porsche. You might try to convince your wife that in case a car is both nice and fast you would have heard about any safety problems. That is, you would like to add the following pair of rules:

- (1.5) $buy(x) \leftarrow \text{not } \neg buy(x), \text{not } \neg safe(x), nice(x), fast(x)$
 (1.5') $\neg buy(y) \leftarrow \text{not } \neg safe(x), nice(x), fast(x), buy(x), x \neq y$

If your wife accepts this rule with preferences $\{(1.5), (1.5')\} < \{(2), (2')\}$ you are happy since now the single preferred answer set contains $buy(Porsche)$.

6 ADDING DISJUNCTION

In this section we will extend our programs to *disjunctive* statements. In Knowledge Representation it often occurs that we know $A \vee B \vee C$ without being sure which of these propositions hold. In fact, such a disjunction leaves it open: there might be states in the world where A holds or B or C or any combination thereof. Nevertheless, we can have information that A implies D and B implies D and C implies D from which we would like to derive that D holds for sure. We will see in Section 6.5 that even with disjunctive programs without negation we can already express relations which belong to the second level of the polynomial hierarchy.

Concerning the right semantics for such programs, we are in the same situation as in Section 3—for positive programs there is general agreement while for disjunctive programs with default-negation there exist several competing approaches.

We present in Section 6.1 the generalized closed world assumption introduced by Minker. In Section 6.2 we show that our definition of WFS from Section 3.3 immediately carries over to the disjunctive case. The original definition of STABLE (Definition 3.4 on page 28) also carries over—we present it in Section 6.3. We mention some other attempts to define disjunctive semantics in Section 6.4. Finally we discuss complexity and expressibility in Section 6.5.

6.1 GCWA

GCWA was defined by Minker ([Minker, 1982]) and can be seen as a refined version of the CWA introduced by Reiter ([Reiter, 1978]):

Definition 6.1 (CWA).

$$\text{CWA}(\text{DB}) = \text{DB} \cup \{\neg P(t) : \text{DB} \not\models P(t)\},$$

where $P(t)$ is a ground predicate instance.

That is, if a ground term cannot be inferred from the database, its negation is added to the closure. A weakness of CWA is that already for very simple theories, like $A \vee B$ it is inconsistent. Since neither $\neg A$ nor $\neg B$ is derivable, we have to add them both which makes the whole set inconsistent.

GCWA is defined for positive disjunctive programs consisting of rules of the form

$$A_1 \vee \dots \vee A_n \leftarrow B_1, \dots, B_m$$

by declaring all the minimal models to be the intended ones:

Definition 6.2 (GCWA).

The generalized closed world assumption GCWA of P is the semantics given by the set of all minimal Herbrand models of P :

$$\text{GCWA}(P) := \text{Min-MOD}(P)$$

Originally, Minker denoted by $\text{GCWA}(P)$ a set of negated atoms with the property that $P \cup \text{GCWA}(P) \models \text{not } A$ if and only if $\text{MinMOD}(P) \models \text{not } A$ but we prefer here to denote by GCWA a semantics in the sense of Definition 2.4 on page 15.

GCWA is very important because it plays the same role for positive disjunctive programs as the least Herbrand model M_P does for definite programs. In addition it turns out that some semantics SEM defined for arbitrary disjunctive programs (i. e. with default-negation) can be characterized, sometimes even implemented, by reducing them to positive programs and then applying recursively GCWA. Thus an appropriate procedure iterating GCWA can “implement” such semantics SEM.

Note also that as far as we consider deriving *positive* disjunctions, we stay entirely within classical logic—a positive disjunction is true in GCWA if and only if it follows from the program considered as a classical theory. Therefore this task can be accomplished by methods and techniques developed in theorem proving in the last 30 years. In fact this was one of the main starting points of the DisLoP-project in Koblenz (see Section 7.2).

Of course, GCWA is nothing else than Circumscription (see Section A.4) for a special class of theories. Methods developed for CIRC can be used to compute GCWA. For recent approaches that work in *polynomial space* see [Niemelä, 1996a; Niemelä, 1996b].

In Sections 2 and 3 we have introduced the general notion of a semantics and various principles. Do they carry over to the disjunctive case? Fortunately, the answer is yes. In addition, GCWA not only satisfies all these properties, it is also uniquely characterized by them as the next theorem shows (we will introduce these properties in the next section).

Theorem 6.1 (Characterization of GCWA, [Brass and Dix, 1997]).

Let SEM be a semantics satisfying *GPPE* and *Elimination of Tautologies*.

a) Then: $\text{SEM}(P) \subseteq \text{Min-MOD}_{2\text{-val}}(P)$ for positive disj. programs P .

I.e. any such semantics is already based on 2-valued minimal models. In particular, GCWA is the weakest semantics with these properties.

b) If SEM is non-trivial and satisfies in addition¹⁴ *Isomorphy* and *Relevance*, then it coincides with GCWA on positive disjunctive programs.

We end this section with the discussion of a well-known example that can not be handled adequately by Circumscription:

¹⁴See Section 7.1 for the precise definitions of Relevance and Isomorphy.

Example 6.1 (Poole's Broken Arm).

Usually, a person's left arm is useable. But if the left arm is broken, it is an exception. The same statement holds for the right arm. Suppose that we saw Fred yesterday with a broken arm but we do not remember if it was the left or the right one. We also know that Fred can make out a cheque if he has at least one useable arm (he is ambidextrous) but that he is completely disabled if both arms are broken. Here is the natural formalization:

$left_use(x)$	\leftarrow	$not\ ab(left, x)$
$ab(left, x)$	\leftarrow	$left_brok(x)$
$right_use(x)$	\leftarrow	$not\ ab(right, x)$
$ab(right, x)$	\leftarrow	$right_brok(x)$
$left_brok(Fred) \vee right_brok(Fred)$	\leftarrow	
$make_cheque(x)$	\leftarrow	$left_use(x)$
$make_cheque(x)$	\leftarrow	$right_use(x)$
$disabled(x)$	\leftarrow	$left_brok(x), right_brok(x)$

Of course, we expect that Fred is able to make out a cheque even without knowing which arm he is actually using. Also we derive that he is not (completely) disabled.

For general Circumscription, the problem is to rule out the unintended model where both arms are broken and Fred is disabled. As we will see later, both D-WFS and DSTABLE derive that Fred is not disabled but only DSTABLE is strong enough to also conclude that Fred can make out a cheque.

6.2 D-WFS

Before we can state the definition of D-WFS we have to extend our principles to disjunctive programs with default-negation. We abbreviate general rules

$$A_1 \vee \dots \vee A_k \leftarrow B_1, \dots, B_m, not\ C_1, \dots, not\ C_n,$$

by

$$\mathcal{A} \leftarrow \mathcal{B}^+, not\ \mathcal{B}^-$$

where $\mathcal{A} := \{A_1, \dots, A_k\}$, $\mathcal{B}^+ := \{B_1, \dots, B_m\}$, $\mathcal{B}^- := \{C_1, \dots, C_n\}$. We also generalize our notion of a semantics slightly:

Definition 6.3 (Operator \models , Semantics \mathcal{S}_{\models}).

By a semantic operator \models we mean a binary relation between logic programs and pure disjunctions which satisfies the following three arguably obvious conditions:

1. *Right Weakening:* If $P \models \psi$ and $\psi \subseteq \psi'$ ¹⁵, then $P \models \psi'$.
2. *Necessarily True:* If $\mathcal{A} \leftarrow true \in P$ for a disjunction \mathcal{A} , then $P \models \mathcal{A}$.

¹⁵I. e. ψ is a subdisjunction of ψ' .

3. *Necessarily False*: If $A \notin \text{Head_atoms}(P)$ ¹⁶ for some \mathcal{L} -ground atom A , then $P \not\sim \text{not } A$.

Given such an operator \sim and a logic program P , by the semantics $\mathcal{S}_{\sim}(P)$ of P determined by \sim we mean the set of all pure disjunctions derivable by \sim from P , i.e., $\mathcal{S}_{\sim}(P) := \{\psi \mid P \sim \psi\}$.

In order to give a unified treatment in the sequel, we introduce the following notion:

Definition 6.4 (Invariance of \sim under a Transformation).

Suppose that a program transformation $\text{Trans} : P \mapsto \text{Trans}(P)$ mapping logic programs into logic programs is given. We say that the operator \sim is invariant under Trans (or that Trans is a \sim -equivalence transformation) iff

$$P \sim \psi \iff \text{Trans}(P) \sim \psi$$

for any pure disjunction ψ and any program P .

All our principles introduced below can now be naturally extended.

Definition 6.5 (Elimination of Tautologies, Non-Minimal Rules).

Semantics \mathcal{S}_{\sim} satisfies **a)** the *Elimination of Tautologies*, resp. **b)** the *Elimination of Non-Minimal Rules* iff \sim is invariant under the following transformations:

- a)** Delete a rule $A \leftarrow B^+ \wedge \text{not } B^-$ with $A \cap B^+ \neq \emptyset$.
- b)** Delete a rule $A \leftarrow B^+ \wedge \text{not } B^-$ if there is another rule $A' \leftarrow B^{+'} \wedge \text{not } B^{-'}$ with $A' \subseteq A$, $B^{+'} \subseteq B^+$, and $B^{-'} \subseteq B^-$.

Our partial evaluation principle has now to take into account disjunctive heads. The following definition was introduced independently by Sakama/Seki and Brass/Dix ([Brass and Dix, 1994; Brass and Dix, 1997; Sakama and Seki, 1994]):

Definition 6.6 (GPPE).

Semantics \mathcal{S}_{\sim} satisfies GPPE iff it is invariant under the following transformation: *Replace a rule $A \leftarrow B^+ \wedge \text{not } B^-$ where B^+ contains a distinguished atom B by the rules*

$$A \cup (A_i \setminus \{B\}) \leftarrow (B^+ \setminus \{B\}) \cup B_i^+ \wedge \text{not } (B^- \cup B_i^-) \quad (i = 1, \dots, n)$$

where $A_i \leftarrow B_i^+ \wedge \text{not } B_i^-$ ($i = 1, \dots, n$) are all the rules with $B \in A_i$.

Note that we are free to select a specific positive occurrence of an atom B and then perform the transformation. The new rules are obtained by replacing B by the bodies of all rules r with head literal B and adding the remaining head atoms of r to the head of the new rule.

¹⁶We denote by $\text{Head_atoms}(P)$ the set of all (instantiations of) atoms occurring in some rule-head of P .

Here is the analogue of Principle 4 on page 22:

Definition 6.7 (Positive and Negative Reduction).

Semantics \mathcal{S}_\perp satisfies **a) Positive**, resp. **b) Negative Reduction** iff \sim is invariant under the following transformations:

- a)** Replace $\mathcal{A} \leftarrow \mathcal{B}^+ \wedge \text{not } \mathcal{B}^-$ by $\mathcal{A} \leftarrow \mathcal{B}^+ \wedge \text{not } (\mathcal{B}^- \cap \text{Head_atoms}(P))$.
- b)** Delete $\mathcal{A} \leftarrow \mathcal{B}^+ \wedge \text{not } \mathcal{B}^-$ if there is a rule $\mathcal{A}' \leftarrow \text{true}$ with $\mathcal{A}' \subseteq \mathcal{B}^-$.

Now the definition of a disjunctive counterpart of WFS is straightforward:

Definition 6.8 (D-WFS).

There exists the weakest semantics satisfying positive and negative Reduction, GPPE, Elimination of Tautologies and non-minimal Rules. We call this semantics D-WFS.

As it was the case for WFS, our calculus of transformations is also confluent ([Brass and Dix, 1998]).

Theorem 6.2 (Confluent Calculus for D-WFS).

The calculus consisting of our four transformations is confluent and terminating for propositional programs. I.e. we always arrive at an irreducible program, which is uniquely determined. The order of the transformations does not matter.

Therefore any program P is associated a unique normalform $\text{res}(P)$. The disjunctive wellfounded semantics of P can be read off from $\text{res}(P)$ as follows

$$\psi \in \text{D-WFS}(P) \iff \begin{array}{l} \text{there is } \mathcal{A} \subseteq \psi \text{ with } \mathcal{A} \leftarrow \text{true} \in \text{res}(P) \text{ or} \\ \text{there is } \text{not } A \in \psi \text{ and } A \notin \text{Head_atoms}(\text{res}(P)). \end{array}$$

Note that the original definition of WFS, or any of its equivalent characterizations, does not carry over to disjunctive programs in a natural way.

Let us see how Example 6.1 on page 51 is handled by D-WFS. Applying GPPE and Reduction gives us the following residual program (we consider just the *Fred*-instantiations):

$$\begin{array}{ll} \text{left_use}(F) & \leftarrow \text{not } ab(\text{left}, F) \\ ab(\text{left}, F) \vee \text{right_brok}(F) & \leftarrow \\ \text{right_use}(F) & \leftarrow \text{not } ab(\text{right}, F) \\ ab(\text{right}, F) \vee \text{left_brok}(F) & \leftarrow \\ \text{left_brok}(F) \vee \text{right_brok}(F) & \leftarrow \\ \text{make_cheque}(F) & \leftarrow \text{not } ab(\text{left}, F) \\ \text{make_cheque}(F) & \leftarrow \text{not } ab(\text{right}, F) \end{array}$$

Therefore we derive $\text{not } \text{disabled}(F)$, because it does not appear in any head of the residual program. All the remaining atoms are undefined.

Two properties of D-WFS are worth noticing

- For positive disjunctive programs, D-WFS coincides with GCWA.
- For non-disjunctive programs with negation, D-WFS coincides with WFS.

6.3 DSTABLE

Unlike the wellfounded semantics, the original definition of stable models carries over to disjunctive programs quite easily:

Definition 6.9 (DSTABLE).

N is called a *stable model*¹⁷ of P iff $N \in \text{Min-Mod}(P^N)$.

In the last definition P^N is the positive disjunctive program obtained from P by applying the Gelfond/Lifschitz transformation (as introduced before Definition 3.4 on page 28—its generalization to disjunctive programs is obvious).

Analogously to D-WFS the following two properties of DSTABLE hold:

- For positive disjunctive programs, DSTABLE coincides with GCWA.
- For non-disjunctive programs with negation, DSTABLE coincides with STABLE.

What about our transformations introduced to define D-WFS? Do they hold for DSTABLE? Yes, they are indeed true. The most difficult proof is the one for GPPE. It was proved in [Brass and Dix, 1999; Sakama and Seki, 1994] independently that stable models are preserved under GPPE. Moreover, Brass/Dix proved in [Brass and Dix, 1997] that STABLE can be almost uniquely determined by GPPE:

Theorem 6.3 (Characterization of DSTABLE).

Let SEM be a semantics satisfying GPPE, *Elimination of Tautologies*, and *Elimination of Contradictions*. Then: $\text{SEM}(P) \subseteq \text{STABLE}(P)$.

Moreover, DSTABLE is the weakest semantics satisfying these properties.

DSTABLE is stronger than D-WFS as can be seen from Example 6.1 on page 51. There we have exactly two stable models

1. $\text{left_use}(F)$, $\text{not } ab(\text{left}, F)$, $ab(\text{right}, F)$, $\text{not } \text{right_use}(F)$,
 $\text{left_brok}(F)$, $\text{not } \text{right_brok}(F)$, $\text{make_cheque}(F)$, $\text{not } \text{disabled}(F)$,
2. $\text{right_use}(F)$, $\text{not } ab(\text{right}, F)$, $ab(\text{left}, F)$, $\text{not } \text{left_use}(F)$,
 $\text{right_brok}(F)$, $\text{not } \text{left_brok}(F)$, $\text{make_cheque}(F)$, $\text{not } \text{disabled}(F)$.

In all of them, Fred is not disabled and can make out a cheque.

Of course, DSTABLE inherits the shortcomings of STABLE such as inconsistency and no goal-orientedness.

¹⁷Note that we only consider Herbrand models.

6.4 Other Semantics

In this section we just want to mention some other disjunctive semantics proposed in the last years. First, there are semantics differing from GCWA in that they interpret “ \vee ” inclusively, rather than exclusively (like GCWA does).

The corresponding semantics is called WGCWA (see [Rajasekar *et al.*, 1989]) and is equivalent to the disjunctive database rule DDR considered in [Ross and Topor, 1988]. WGCWA has been considered as a more tractable (weaker) variant of GCWA (from the procedural point of view developed in [Rajasekar *et al.*, 1989]).

Example 6.2 (Inclusive versus Exclusive).

$$P_{incl/excl} : \begin{array}{l} a \vee b \\ c \quad \leftarrow \quad a, b \end{array}$$

Under an *exclusive* interpretation, $\text{not } c$ should be derivable. Indeed, we have $\text{GCWA}(P_{incl/excl}) = \{\text{not } c\}$.

Under an *inclusive* interpretation however, $\text{not } c$ should not be derivable. This is the case for WGCWA: $M_{P_{incl/excl}}^* = \{a, b, c\}$. The set of *positive* derivable literals is in both cases the same! If we replace the first clause with a or b , then $\text{not } c$ is derivable.

There are extensions of WGCWA to disjunctive programs with negation: [Sakama and Inoue, 1993; Ross, 1989; Dix, 1992b; Dix and Müller, 1994a].

There is also the book [Lobo *et al.*, 1992]—the first in-depth-study of disjunctive semantics with negation. However, we feel that these semantics have a drawback in that they are based on rather technical, complicated and not-easy-to-understand fixpoint definitions. These definitions leave a lot of room for modifications. But small modifications usually have a tremendous impact on the outcoming semantics. In addition these semantics do not allow for a proper treatment of definitional extensions (see Example 7.1 on page 60).

Let us discuss one more time the *inclusive/exclusive* meaning of \vee . Chiaki Sakama noted that inclusive vs. exclusive is not always fully determined by the underlying semantics. For example in the program $p \vee q, p \leftarrow q, q \leftarrow p$, GCWA derives both p and q so it is better to say that GCWA *tends to interpret* \vee *exclusively unless specified otherwise*. But weaker semantics cannot specify exclusive \vee . Therefore Sakama proposed in [Sakama, 1989] to consider programs with two different kinds of semantics: one for inclusive and one for exclusive behaviour. He defined a corresponding semantics, called *possible model semantics* PMS (see also [Sakama and Inoue, 1993; Sakama and Inoue, 1994]). Chan introduced this same idea under the name *possible world semantics* in [Chan, 1993]. PMS has the nice feature that it lies on the first level of the polynomial hierarchy (see [Eiter and Gottlob, 1993a]).

Other approaches are due to Przymusiński: *stationary-semantics* STN and

	Complexity	
	<i>1. ord. prog.</i> (with functions)	<i>prop. prog.</i> (no variables)
GCWA (<i>P is positive</i>)	A: Σ_1^0 -compl. not A: Π_2^0 -compl.	A: co-NP -compl. not A: Π_2^P -compl.
WGCWA (<i>P is positive</i>)	A: Σ_1^0 -compl. not A: Π_1^0 -compl.	A: co-NP -compl. not A: linear in $ P $
PERFECT (<i>P is stratified</i>)	arithm. -compl.	Π_2^P -compl.
PMS	not yet studied	co-NP -compl
WPERFECT (<i>P is stratified</i>)	arithm. -compl.	Π_2^P -compl.
D-WFS	Π_1^1 -compl. over N	Π_2^P -compl.
DSTABLE	Π_1^1 -compl. over N	Π_2^P -compl.

Table 4. Complexity of Disjunctive Semantics

static-semantics STATIC (see [Przymusiński, 1995; Brass *et al.*, 1999; Brass *et al.*, 2001a]). STATIC is an improvement of his former stationary semantics that is very close to D-WFS: in fact it coincides with D-WFS if it is restricted to a common sublanguage ([Brass *et al.*, 2001a]). This approach also allows us to consider a larger class of programs, namely those that contain $\text{not}(A_1 \wedge \dots \wedge A_n)$ in their bodies. Such programs are more expressible and therefore turn out to be even better suited for representation tasks.

Another approach differing from GCWA and WGCWA is considered in [Dix *et al.*, 1994; Dix *et al.*, 1996a; Bonatti, 1993].

6.5 Complexity and Expressibility

From the complexity point of view GCWA lies between CWA (which is Π_1^0 -complete, see [Apt and Blair, 1990] and general Circumscription (Σ_1^1 -complete, see [Cadoli *et al.*, 1992]): GCWA is Π_2^0 -complete. For propositional programs we have to distinguish between deriving an *atom* or a *literal*. The first problem is **co-NP**-complete while the second is even Π_2^P -complete (see [Imielinski, 1991]).

For deriving negated literals $\text{not } A$, WGCWA is Π_1^0 -complete (like CWA) and therefore “better” than GCWA (Π_2^0 -complete). In the propositional case, WGCWA is polynomial while GCWA is Π_2^P -complete (both for the derivation of literals of the form $\text{not } A$).

	Expressibility <i>1. ord. prog.</i> <i>(no functions)</i>
GCWA <i>(P is positive)</i>	$\subset \Pi_2^P$
WGCWA <i>(P is positive)</i>	$\subset \Pi_2^P$
PERFECT <i>(P is stratified)</i>	$= \Pi_2^P$
WPERFECT <i>(P is stratified)</i>	$= \Pi_2^P$
D-WFS	$= \Pi_2^P$
DSTABLE	$= \Pi_2^P$

Table 5. Expressibility of Disjunctive Semantics

7 WHAT DO WE WANT AND WHAT IS IMPLEMENTED?

In this part we first consider the question *Is there an optimal semantics?* (Section 7.1) and give in Section 7.2 an overview of all the existing implementations we are aware of. We also describe theoretical approaches that have not yet been implemented.

7.1 What is the Best Semantics?

Most probably there is no definite answer to the question in the title. Different knowledge representation tasks may ask for different semantics. Some might be better suited in special domains than others. What are reasonable properties that semantics should be checked against?

While many people defined in the last years new semantics by considering only few examples and appealing to their own personal intuitions they had about how these few examples should be handled, Dix tried to adjust and investigate abstract properties known in general nonmonotonic reasoning to semantics of logic programs ([Dix, 1991; Dix, 1992b; Dix, 1995a; Dix, 1995b]). He showed for example that WFS is cumulative and rational and that a semantics defined independently by Schlipf and Dix is the weakest extension of WFS satisfying *Cut* and *Supraclassicality*. Figure 3 on the following page illustrates the properties and the relationship between many semantics (note that \leq_k refers to the knowledge ordering discussed

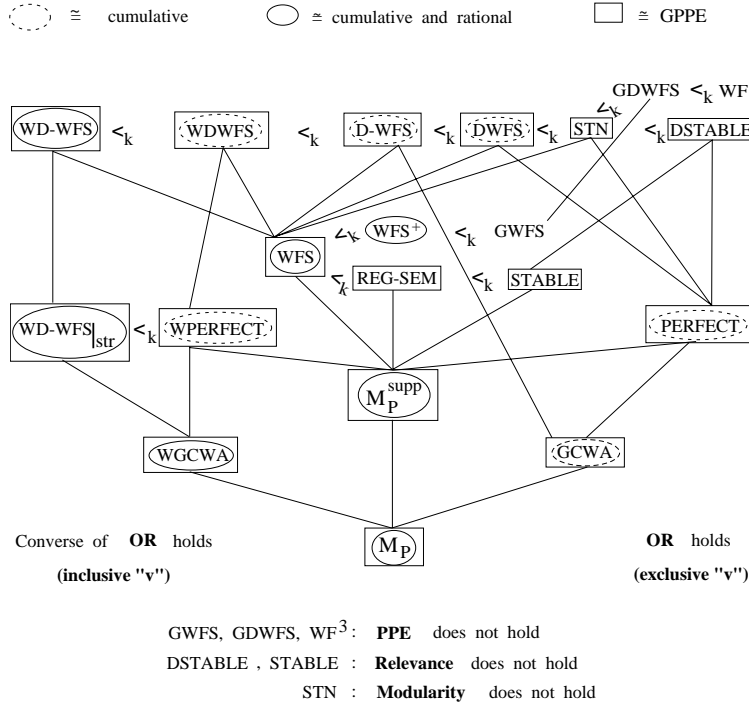


Figure 3. Semantics for Disjunctive Programs

in the beginning of Subsection 3.3: a semantics is *strictly weaker*, i.e. \leq_k , than another semantics if it derives strictly less literals). The distinction between *inclusive* and *exclusive* semantics is perhaps a bit misleading. We refer to the discussion in Subsection 6.4.

In Figure 4 on the next page normal programs are considered.

Besides such properties (which he calls *strong*) he defined also *weak* properties—these are conditions that any reasonable semantics should satisfy ([Dix, 1992a; Dix, 1995b]). The principles we have introduced in Sections 2, 3 belong to this sort. Let us take a closer look into some weak properties already mentioned (but not yet defined). We start with a property that is satisfied for any semantics we know:

Definition 7.1 (Isomorphy).

A semantics SEM satisfies Isomorphy, iff

$$SEM(\mathcal{I}(P)) = \mathcal{I}(SEM(P))$$

for all programs P and isomorphisms \mathcal{I} on the Herbrand base B_P .

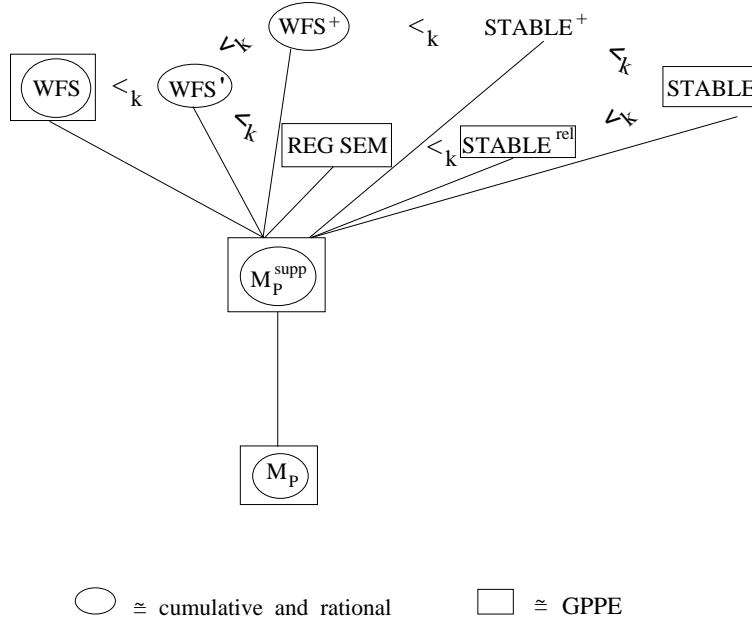


Figure 4. Semantics for Normal Programs

Isomorphism formalizes the intuition that a renaming of the program should have no influence on the semantics, as long as we also apply this same renaming to the semantics.

The next property gives a formal definition of the notion *Goal-Orientedness*. To state this conditions, we need the notion of the Dependency-Graph (Definition 3.2 on page 21) and the two definitions

- $dependencies_of(X) := \{A : X \text{ depends on } A\}$, and
- $rel_rul(P, X)$ is the set of *relevant rules* of P with respect to X , i.e. the set of rules that contain an $A \in dependencies_of(X)$ in their head.

Given any semantics SEM and a program P , it is perfectly reasonable that the truthvalue of a literal L , with respect to SEM(P), only depends on the subprogram formed from the *relevant rules* of P with respect to L .¹⁸ This idea is formalized by:

Definition 7.2 (Relevance).

The principle of Relevance states: $L \in SEM(P)$ iff $L \in SEM(rel_rul(P, L))$.

¹⁸Let $dependencies_of(not X) := dependencies_of(X)$, and $rel_rul(P, not X) := rel_rul(P, X)$.

Note that the set of relevant rules of a program P with respect to a literal L contains all rules, that could ever contribute to L 's derivation (or to its nonderivability). In general, L depends on a large set of atoms: $dependencies_of(L) := \{A : L \text{ depends on } A\}$. But rules that do not contain these atoms in their heads, will never contribute to their derivation or non-derivation. Therefore, these rules should not affect the meaning of L in P . STABLE does not satisfy this principle. This is due to the nonexistence of stable models by adding a clause " $c \leftarrow \text{not } c$ " to a program.

We have already introduced GPPE above. It is an extension of the following property for non-disjunctive programs:

Definition 7.3 (PPE).

Let P be an instantiated program and let the atom c occur positively in P . Let $c \leftarrow rhs_1, \dots, c \leftarrow rhs_n$ be all the rules of P with c in their heads.

Any program clause of the form " $head \leftarrow c, body$ " can be replaced by the rules

$$\begin{array}{l} head \leftarrow rhs_1, body \\ \vdots \\ head \leftarrow rhs_n, body \end{array}$$

Note that the rules $c \leftarrow rhs_1, \dots, c \leftarrow rhs_n$ are not removed (in contrast to the weak version of PPE). We call the program obtained in this way P' .

The principle of partial evaluation is: $SEM(P') = SEM(P)$.

GPPE is obtained from PPE by weakening the assumption that c only occurs positively. We note that most semantics defined by Minker and his group do not satisfy this condition:

Example 7.1 (Extension-by-Definition, [Dix, 1991]).

We consider the following two programs:

$$\begin{array}{ll} P_{GWFS} : & p \leftarrow \text{not } b \\ & a \leftarrow \text{not } b \\ & b \leftarrow c \\ & c \leftarrow p, \text{not } a \end{array} \quad \begin{array}{ll} P_{GWFS_c} : & p \leftarrow \text{not } b \\ & a \leftarrow \text{not } b \\ & b \leftarrow p, \text{not } a \\ & c \leftarrow p, \text{not } a \end{array}$$

$GWFS(P_{GWFS})$ entails $\text{not } c$, because $\text{Min-MOD}(P_{GWFS}) = \{ \{p, a\}, \{b\} \}$ and thus also (by simple negation-as-failure reasoning) $\text{not } b, p$ and a . Also $\text{Min-MOD}(P_{GWFS_c}) = \{ \{p, a\}, \{b\} \}$ but negation-as-failure can not be applied like before. Therefore $GWFS(P_{GWFS_c})$ does not entail $\text{not } b$, nor p nor a .

P_{GWFS_c} partial evaluates P_{GWFS} : the last but one clause was transformed into another one by expanding the definition of c . Obviously, a semantics should assign the same meaning to these programs: unfortunately GWFS does not!

The next principle, *Modularity*, has some similarities with PPE. It enables us to compute a semantics by modularizing it into certain "subprograms" (formed of

the relevant rules). The semantics of these modules can be computed first and the semantics of the whole program can be determined by reducing this program with literals that were already determined.

Definition 7.4 (Modularity).

Let $P = P_1 \cup P_2$ and for every $A \in B_{P_2}$: $rel_rul(P, A) \subseteq P_2$.

The principle of Modularity is: $SEM(P) = SEM(P_1^{SEM(P_2)} \cup P_2)$.

To illustrate this property, we compare the program

$$\begin{array}{l}
 P_1 : \quad b \quad \leftarrow \quad y \\
 \quad \quad y \vee x \\
 \quad \quad z \vee y \\
 \quad \quad m \quad \leftarrow \quad x, z, b \\
 \quad \quad y \quad \leftarrow \quad \text{not } m
 \end{array}$$

with the union of the following two programs

$$\begin{array}{l}
 P'_1 : \quad b \quad \leftarrow \quad y \\
 \quad \quad y \vee x \quad \quad a \\
 \quad \quad z \vee y \quad \quad a \\
 \quad \quad m \quad \leftarrow \quad x, z, b \\
 \quad \quad y \quad \leftarrow \quad \text{not } m \\
 \\
 P_2 : \quad a \quad \quad \leftarrow \quad e, \text{ not } g \\
 \quad \quad a \quad \quad \leftarrow \quad f, \text{ not } g \\
 \quad \quad a \quad \quad \leftarrow \quad f, \text{ not } e \\
 \quad \quad a \quad \quad \leftarrow \quad g, \text{ not } e \\
 \quad \quad e \vee f \vee g
 \end{array}$$

P_2 is a stratified program and STN derives a . Concerning P_1 , different intuitions seem possible. One can argue, that $\text{not } m$ should be derivable, since the only way to derive m is by using the fourth clause, which means deriving b , which means deriving y which excludes deriving x or z . This is the way, P_1 is handled by the first version of STN . The second (final) version STN does not derive $\text{not } m$. But if we apply STN to $P'_1 \cup P_2$, then $\text{not } m$ is derivable. This shows that *weak Modularity* is not satisfied: we consider this to be a serious shortcoming.

Typical results of Dix are

- WFS is the weakest semantics satisfying some of these weak properties,
- WFS can be uniquely characterized if some strong properties are added.

We conclude with Table 6 on the next page: an overview of the properties of some semantics mentioned above.

The bad properties of the PMS (failure of Relevance) stem from the fact that it was originally based on stable models. But the underlying idea of PMS is to transform disjunctive programs into non-disjunctive ones and then applying a semantics for non-disjunctive programs. By choosing semantics different from STABLE, PMS inherits other properties (see [Sakama and Inoue, 1994]).

7.2 Query-Answering Systems and Implementations

In this section we give a rough overview of what semantics have been implemented so far and where they are available. As already explained in Sections 3.5, 6.5,

Properties of Logic-Programming Semantics							
Semantics	Reference	Domain	Taut.	GPPE	Red.	NMin.	Rel.
comp	Cla78	Nondis.	—	•	•	•	—
GCWA	Min82	Pos.	•	•	•	•	•
WGCWA	RosTop88	Pos.	—	•	•	—	•
DSTABLE	GeLif91	Dis.	•	•	•	•	—
WFS	vGeld. etal88	Nondis.	•	•	•	•	•
$ST\mathcal{N}$	Prz91	Dis.	•	•	•	•	•
STATIC	Prz95	Dis.	•	•	•	•	•
D-WFS	BraDix95	Dis.	•	•	•	•	•
DWFS	Dix92	Dis.	•	•	•	•	•
Str. WFS	Ros92	Dis.	—	—	•	—	•
WD-WFS	BraDix95	Dis.	—	•	•	—	•
WDWFS	Dix92	Dis.	—	•	•	—	•
PMS	SakIno94	Dis.	—	—	•	—	—

Table 6. Semantics and Their Equivalence-Transformations

our NMR-semantics are undecidable in general. Nevertheless we think it is very important to have running systems that

1. can handle programs with free variables, and
2. are Goal-Oriented.

To ensure *completeness* (or *termination*) we need then additional requirements like *allowedness* (to prevent floundering, see Section 3.1) and no function symbols.

Although these restrictions ensure the Herbrand-universe to be finite (and thus we are really considering a propositional theory) we think that such a system has great advantages over a system that can just handle ground programs. For a language \mathcal{L} , the fully instantiated program can be quite large and difficult to handle effectively. The goal-orientedness (or *Relevance* as introduced in Section 7.1) is also important—after all this was one reason of the success of SLD-Resolution. As noted above, such a goal-oriented approach is not possible for STABLE.

LP-Semantics

Various commercial PROLOG-systems perform variants of SLDNF-Resolution. Chan's constructive negation has also been implemented as part of the master-theses [Ludäscher, 1991; Vorbeck, 1991].

Currently, a library of implemented logic programming systems and interesting test-cases for such systems is collected as a project of the artificial intelligence group at Koblenz. We refer to <http://www.uni-koblenz.de/ag-ki/LP/>.

Non-Disjunctive NMR-Semantics

There are many theoretical papers that deal with the problem of implementation ([Bol and Degerstedt, 1993; Kemp *et al.*, 1991; Degerstedt and Nilsson, 1995; Fernández *et al.*, 1993]) but only few running systems. The problem of handling and representing ground programs given a non-ground one has also been addressed [Kagan *et al.*, 1994; Kagan *et al.*, 1995; Eiter *et al.*, 1997a].

In [Bell *et al.*, 1993; Bell *et al.*, 1994] the authors showed how the problem of computing stable models can be transformed to an Integer-Linear Programming Problem. This has been extended in [Dix and Müller, 1993] to disjunctive programs.

Inoue *et al.* show in [Inoue *et al.*, 1992] how to compute stable models by transforming programs into propositional theories and then using a model-generation theorem prover.

In Berne, Switzerland, a group around G. Jäger has built a non-monotonic reasoning system which incorporates various monotonic and non-monotonic logics. We refer to <http://lwbwww.unibe.ch:8080/LWBinfo.html>.

Extended logic programs under the well-founded semantics are considered by Pereira and his colleagues: [Pereira *et al.*, 1993; Alferes and Pereira, 1996]. The REVISE system, which deals with contradiction removal for paraconsistent programs in this semantics, can be found in <http://www.uni-koblenz.de/ag-ki/LP/> too.

In [Niemelä and Simons, 1996], an implementation of STABLE with a special eye on complexity is described. The resulting system, *smodels*, is publicly available (see <http://www.tcs.hut.fi/Software/smodels/>) and seems to outperform most other approaches to implementing STABLE. More references can be found in [Dix *et al.*, 2001a]. Many problems in model checking, planning, diagnosis and various other areas can be translated into logic programs in such a way, that stable models of these programs correspond exactly to solutions of the original problems.

The most advanced system, XSB, has been implemented by David Warren and his group in Stony Brook based on OLDT-algorithm of [Tamaki and Sato, 1986]. They first developed a meta-interpreter (SLG, see [Chen and Warren, 1996]) in PROLOG and then directly modified the WAM for a direct implementation of WFS (XSB). They use tabling-methods and a mixture of Top-Down and bottom-up evaluation to detect loops. The XSB system is complete and terminating for non-floundering DATALOG. It also works for general programs but termination is not guaranteed. This system is described in [Chen and Warren, 1993; Chen *et al.*, 1995; Chen and Warren, 1995; Swift, 1999], and is available by anonymous ftp from <ftp.cs.sunysb.edu/pub/XSB/>. We also refer to [Apt *et al.*, 1999].

Disjunctive NMR-Semantics

There are theoretical descriptions of implementations that have not yet been implemented: [Fernández and Minker, 1995; Minker and Ruiz, 1995; Costantini and Lanzarone, 1995]. Also Sakama and Seki describe an approach for first-order disjunctive programs ([Sakama and Seki, 1997]).

Here are some implemented systems. Inoue et. al. show in [Inoue *et al.*, 1992] how to compute stable models for extended disjunctive programs in a bottom-up-fashion using a theorem prover.

The approach of Bell et. al. ([Nerode *et al.*, 1991]) was used by Dix/Müller to implement versions of the stationary semantics of Przymusinski ([Przymusinski, 1991]): [Müller and Dix, 1993; Dix and Müller, 1992].

Brass/Dix have implemented both D-WFS and DSTABLE for allowed DATA-LOG programs ([Brass and Dix, 1995]). An implementation of static semantics is described in [Brass *et al.*, 1999].

Seipel has implemented in his DisLog-system various (modified versions of) semantics of Minker and his group (<http://sunwww.informatik.uni-tuebingen.de:8080/dislog/dislog.tar.Z>).

The DisLoP project (1995–2000, see [Aravindan *et al.*, 1997]) aimed at extending certain theorem proving concepts, such as restart model elimination and hyper tableaux calculi, for disjunctive logic programming. This system can be extended to handle non-monotonic semantics such as D-WFS, STATIC etc. In particular, an implementation of D-WFS for general disjunctive programs which works *in polynomial space* is available ([Brass *et al.*, 2001a]). An extension to first-order programs is proposed in [Dix and Stolzenburg, 1998]. Information on the DisLoP project and related publications can be obtained from the WWW page <<http://www.uni-koblenz.de/~dix/DLP/>>.

The most advanced system however is dlv ([Eiter *et al.*, 1997b; Eiter *et al.*, 1998]). It constitutes a knowledge representation system, based on disjunctive logic programming, which offers front-ends to several advanced KR formalisms (developed since the end of 1996). Major emphasis has been put on advanced knowledge modelling features. The kernel language, which extends disjunctive logic programming by true negation and integrity constraints, allows for representing complex knowledge based problems in a highly declarative fashion [Eiter *et al.*, 1998]. The system runs in *polynomial space* and *single exponential time*, and is able to efficiently recognize and process syntactical subclasses of disjunctive logic programs which have lower computational complexity than the general case (like, e.g., programs with head-cycle free disjunction or stratified negation).

An important outcome of the Dagstuhl Seminar 9627 ([Dix *et al.*, 1996b]) was to construct a web page to collect and disseminate information on various logic programming systems that concentrate on non-monotonic aspects (different kinds of negation, disjunction, abduction etc.). This web page is actively maintained at the URL <<http://www.uni-koblenz.de/ag-ki/LP/>>. In addition the *Logic Programming and Nonmonotonic Reasoning*-conference 1997 ([Dix *et*

al., 1997a]) contained for the first time a special track on implementations and working systems. Various LP-systems were also demonstrated at NMR 2000 <<http://www.cs.engr.uky.edu/nmr2000>>.

OUTLOOK

Currently, a new logic programming paradigm seems to emerge: *answer set programming*. The idea is to take the credulous viewpoint under the stable semantics seriously: the set of all stable models of a program really is the main object described by this program. Viewed in this way, even propositional programs can encode second order objects. Thus when a problem is translated into a logic program under the stable semantics, the program itself can be seen as a *compact encoding* of the set of solutions of the original problem. We refer to [Dix, 1998; Apt *et al.*, 1999; Niemelä, 1999].

Recently, methods of logic programming have also been successfully applied in the area of multi agent systems [Eiter *et al.*, 2000; Eiter *et al.*, 1999; Dix *et al.*, 2000b; Dix *et al.*, 2000a; Dix *et al.*, 2001b; Subrahmanian *et al.*, 2000; Dix, 2001].

ACKNOWLEDGEMENTS

We thank Chandrabose Aravindan, Norbert E. Fuchs, Ulrich Furbach, Ilkka Niemelä and Chiaki Sakama who provided us with many useful remarks. Katrin Erk and Dorothea Schäfer proofread parts of the paper—special thanks to them.

BIBLIOGRAPHY

- [Alferes and Pereira, 1995] Jose Julio Alferes and Luiz Moniz Pereira. An argumentation theoretic semantics based on non-refutable falsity. In J. Dix, L. Pereira, and T. Przymusiński, editors, *Non-monotonic Extensions of Logic Programming*, LNAI 927, pages 3–22. Springer, Berlin, 1995.
- [Alferes and Pereira, 1996] Jose Julio Alferes and Luiz Moniz Pereira, editors. *Reasoning with Logic Programming*, LNAI 1111, Berlin, 1996. Springer.
- [Alferes *et al.*, 1996] Jose Julio Alferes, Luiz Moniz Pereira, and Teodor Przymusiński. Strong and Explicit Negation in Non-Monotonic Reasoning and Logic Programming. In J.J Alferes, L.M. Pereira, and E. Orłowska, editors, *Logics in Artificial Intelligence (JELIA '96)*, LNCS 1126, pages 143–163. Springer, 1996.
- [Apt and Blair, 1990] Krzysztof R. Apt and Howard A. Blair. Arithmetic Classification of perfect Models of stratified Programs. *Fundamenta Informaticae*, XIII:1–17, 1990. Addendum in Vol. XIV, pages 339–344, 1991.
- [Apt and Bol, 1994] Krzysztof R. Apt and Roland N. Bol. Logic Programming and Negation: A Survey. *Journal of Logic Programming*, 19-20:9–71, 1994.
- [Apt *et al.*, 1999] K. R. Apt, V. Marek, M. Truszczyński, and D. S. Warren, editors. *The Logic Programming Paradigm: Current Trends and Future Directions*, Berlin, 1999. Springer.
- [Apt, 1990] Krzysztof R. Apt. Logic programming. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science, Vol. B*, chapter 10, pages 493–574. Elsevier Science Publishers, 1990.
- [Aravindan *et al.*, 1997] Chandrabose Aravindan, Jürgen Dix, and Ilkka Niemelä. Dislop: A research project on disjunctive logic programming. *AI Communications*, 10(3/4):151–165, 1997.

- [Baader and Hollunder, 1995] Franz Baader and Bernd Hollunder. Priorities on defaults with prerequisite and their application in treating specificity in terminological default logic. *Journal of Automated Reasoning*, 15:41–68, 1995.
- [Balcázar et al., 1988] J.L. Balcázar, I. Díaz, and J. Gabarró. *Structural Complexity I*. Springer-Verlag, Berlin, 1988.
- [Baral and Gelfond, 1994] Chitta Baral and Michael Gelfond. Logic Programming and Knowledge Representation. *Journal of Logic Programming*, 19-20:73–148, 1994.
- [Baral and Subrahmanian, 1991] Chitta Baral and V.S. Subrahmanian. Dualities between Alternative Semantics for Logic Programming and Non-monotonic Reasoning. In Anil Nerode, Wiktor Marek, and V. S. Subrahmanian, editors, *Logic Programming and Non-Monotonic Reasoning, Proceedings of the first International Workshop*, pages 69–86, Cambridge, Mass., July 1991. Washington D.C, MIT Press.
- [Baral and Subrahmanian, 1992] Chitta Baral and V.S. Subrahmanian. Stable and Extension Class Theory for Logic Programs and Default Logics. *Journal of Automated Reasoning*, 8, No. 3:345–366, 1992.
- [Barbuti and Martelli, 1986] R. Barbuti and M. Martelli. Negation as Failure. Completeness of the Query Evaluation Process for Horn Clause Programs with Recursive Definition. *Journal of Automated Reasoning*, 2:155–170, 1986.
- [Barwise, 1975] Jon Barwise. *Admissible Sets and Structures*. Springer, 1975.
- [Beeri and Ramakrishnan, 1991] Catril Beeri and Raghu Ramakrishnan. On the power of magic. *The Journal of Logic Programming*, 10:255–299, 1991.
- [Bell et al., 1993] Colin Bell, Anil Nerode, Raymond T. Ng, and V. S. Subrahmanian. Implementing Stable Semantics by Linear Programming. In Luis Moniz Pereira and Anil Nerode, editors, *Logic Programming and Non-Monotonic Reasoning, Proceedings of the Second International Workshop*, pages 23–42, Cambridge, Mass., July 1993. Lisbon, MIT Press.
- [Bell et al., 1994] Colin Bell, Anil Nerode, Raymond T. Ng, and V. S. Subrahmanian. Mixed Integer Programming Methods for Computing Non-Monotonic Deductive Databases. *Journal of the ACM*, 41(6):1178–1215, November 1994.
- [Ben-Eliyahu and Dechter, 1992] Rachel Ben-Eliyahu and Rina Dechter. Propositional Semantics for Disjunctive Logic Programs. In K. R. Apt, editor, *LOGIC PROGRAMMING: Proceedings of the 1992 Joint International Conference and Symposium*, pages 813–, Cambridge, Mass., November 1992. MIT Press.
- [Benferhat et al., 1993] Salem Benferhat, Claudette Cayrol, Didier Dubois, Jerome Lang, and Henri Prade. Inconsistency management and prioritized syntax-based entailment. In *Proc. 13th International Joint Conference on Artificial Intelligence*, pages 640–645. Morgan Kaufman, 1993.
- [Bidoit and Froidevaux, 1991a] Nicole Bidoit and Christine Froidevaux. General logical Databases and Programs: Default Logic Semantics and Stratification. *Information and Computation*, 91:15–54, 1991.
- [Bidoit and Froidevaux, 1991b] Nicole Bidoit and Christine Froidevaux. Negation by Default and unstratifiable logic Programs. *Theoretical Computer Science*, 78:85–112, 1991.
- [Bol and Degerstedt, 1993] Roland N. Bol and L. Degerstedt. Tabulated resolution for well-founded semantics. In *Proc. Int. Logic Programming Symposium '93*, Cambridge, Mass., 1993. MIT Press.
- [Bonatti, 1993] Piero Bonatti. Shift-based semantics: general results and applications. Technical Report CD-TR-93-59, Technical University of Vienna, Inst. für Informationssysteme, 1993.
- [Börger, 1987] Egon Börger. Unsolvable Decision Problems For Prolog Programs. In Egon Börger, editor, *Computation Theory and Logic*, LNCS 270, pages 37–47, Berlin, 1987. Springer.
- [Brass and Dix, 1994] Stefan Brass and Jürgen Dix. A disjunctive semantics based on unfolding and bottom-up evaluation. In Bernd Wolfinger, editor, *Innovationen bei Rechen- und Kommunikationssystemen*, (IFIP '94-Congress, Workshop FG2: Disjunctive Logic Programming and Disjunctive Databases), pages 83–91, Berlin, 1994. Springer.
- [Brass and Dix, 1995] Stefan Brass and Jürgen Dix. A General Approach to Bottom-Up Computation of Disjunctive Semantics. In J. Dix, L. Pereira, and T. Przymusiński, editors, *Nonmonotonic Extensions of Logic Programming*, LNAI 927, pages 127–155. Springer, Berlin, 1995.
- [Brass and Dix, 1997] Stefan Brass and Jürgen Dix. Characterizations of the Disjunctive Stable Semantics by Partial Evaluation. *Journal of Logic Programming*, 32(3):207–228, 1997.
- [Brass and Dix, 1998] Stefan Brass and Jürgen Dix. Characterizations of the Disjunctive Well-founded Semantics: Confluent Calculi and Iterated GCWA. *Journal of Automated Reasoning*, 20(1):143–165, 1998.

- [Brass and Dix, 1999] Stefan Brass and Jürgen Dix. Semantics of (Disjunctive) Logic Programs Based on Partial Evaluation. *Journal of Logic Programming*, 38(3):167–213, 1999.
- [Brass et al., 1999] Stefan Brass, Jürgen Dix, and Teodor Przymusiński. Computation of the semantics of autoepistemic belief theories. *Artificial Intelligence*, 112(1-2):104–123, 1999. (Partial results appeared in: Super Logic Programs. *Proceedings of the 5-th International Conference on Principles of Knowledge Representation and Reasoning*. Boston, MA, pages 529–541, 1996. Morgan Kaufman.)
- [Brass et al., 2001a] Stefan Brass, Jürgen Dix, Ilkka Niemelä, and Teodor Przymusiński. On the Equivalence of the Static and Disjunctive Well-Founded Semantics and its Computation. *Theoretical Computer Science*, 258(1-2):523–553, 2001.
- [Brass et al., 2001b] Stefan Brass, Jürgen Dix, Burkhardt Freitag, and Ulrich Zukowski. Transformation-based bottom-up computation of the well-founded model. *Theory and Practice of Logic Programming*, to appear, 2001.
- [Brewka and Eiter, 1999] Gerhard Brewka and Thomas Eiter. Preferred answer sets. *Artificial Intelligence Journal*, 109:297–356, 1999.
- [Brewka et al., 1997] Gerd Brewka, Jürgen Dix, and Kurt Konolige. *Nonmonotonic Reasoning: An Overview*. CSLI Lecture Notes 73. CSLI Publications, Stanford, CA, 1997.
- [Brewka, 1994] G. Brewka. Adding priorities and specificity to default logic. In *Logics in Artificial Intelligence, Proc. JELIA-94*, York. Springer, 1994.
- [Brewka, 1996a] Gerd Brewka. *Principles of Knowledge Representation*. CSLI. Dikran, 1996.
- [Brewka, 1996b] Gerhard Brewka. Well-founded semantics for extended logic programs with dynamic preferences. *Journal of Artificial Intelligence Research*, 4:19–36, 1996.
- [Brewka, 2001] Gerhard Brewka. On the relation between defeasible logic and well-founded semantics. In *Proc. LPNMR 2001*. Springer, 2001.
- [Bry, 1990] Francois Bry. Query evaluation in recursive databases: bottom-up and top-down reconciled. *Data & Knowledge Engineering*, 5:289–312, 1990.
- [Cadoli and Schaerf, 1993] Marco Cadoli and Marco Schaerf. A Survey of Complexity Results for Non-Monotonic Logics. *Journal of Logic Programming*, 17:127–160, 1993.
- [Cadoli et al., 1992] Marco Cadoli, Thomas Eiter, and Georg Gottlob. An efficient method for eliminating varying predicates from a circumscription. *Artificial Intelligence Journal*, 54:397–410, 1992.
- [Cadoli et al., 1995] M. Cadoli, F. M. Donini, P. Liberatore, and M. Schaerf. The size of a revised knowledge base. In *PODS '95*, pages 151–162, 1995.
- [Cadoli et al., 1996] Marco Cadoli, Francesco M. Donini, and Marco Schaerf. Is intractability of non-monotonic reasoning a real drawback? *Artificial Intelligence Journal*, 88:215–251, 1996.
- [Cadoli et al., 1997] Marco Cadoli, Francesco M. Donini, Marco Schaerf, and Riccardo Silvestri. On compact representations of propositional circumscription. *Theoretical Computer Science*, 182:183–202, 1997.
- [Cavedon and Lloyd, 1989] L. Cavedon and J.W. Lloyd. A Completeness Theorem for SLDNF-Resolution. *Journal of Logic Programming*, 7:177–191, 1989.
- [Chan and Wallace, 1989] David Chan and Mark Wallace. An Experiment with programming using pure Negation. Technical Report TR, ECRC, July 1989.
- [Chan, 1988] David Chan. Constructive negation based on the completed database. In *Proc. 1988 Conf. and Symp. on Logic Programming*, pages 111–125, September 1988.
- [Chan, 1993] Edward P.F. Chan. A Possible World Semantics for Disjunctive Databases. *IEEE Trans. on Knowledge and Data Engineering*, 5(2):282–292, 1993.
- [Chen and Warren, 1993] Weidong Chen and David S. Warren. A Goal Oriented Approach to Computing The Well-founded Semantics. *Journal of Logic Programming*, 17:279–300, 1993.
- [Chen and Warren, 1995] Weidong Chen and David S. Warren. Computing of Stable Models and its Integration with Logical Query Processing. *IEEE Transactions on Knowledge and Data Engineering*, 17:279–300, 1995.
- [Chen and Warren, 1996] Weidong Chen and David S. Warren. Tabled Evaluation with Delaying for General Logic Programs. *Journal of the ACM*, 43(1):20–74, January 1996.
- [Chen et al., 1995] Weidong Chen, Terrance Swift, and David S. Warren. Efficient Top-Down Computation of Queries under the Well-Founded Semantics. *Journal of Logic Programming*, 24(3):219–245, 1995.
- [Chomicki and Subrahmanian, 1990] Jan Chomicki and V.S. Subrahmanian. Generalized Closed World Assumption is Π_2^0 -Complete. *Information Processing Letters*, 34:289–291, 1990.

- [Clark, 1978] Keith L. Clark. Negation as Failure. In H. Gallaire and J. Minker, editors, *Logic and Data-Bases*, pages 293–322. Plenum, New York, 1978.
- [Colmerauer *et al.*, 1973] A. Colmerauer, H. Kanoui, R. Pasero, and P. Roussel. Un système de communication homme-machine en français. Technical report, Groupe de Intelligence Artificielle Université de Aix-Marseille II, 1973.
- [Costantini and Lanzarone, 1995] Stefania Costantini and Gaetano A. Lanzarone. Static Semantics as Program Transformation and Well-founded Computation. In J. Dix, L. Pereira, and T. Przymusiński, editors, *Nonmonotonic Extensions of Logic Programming*, LNAI 927, pages 156–180. Springer, Berlin, 1995.
- [Decker and Cavedon, 1990] Hendrik Decker and Lawrence Cavedon. Generalizing syntactic properties which ensure that SLDNF-Resolution is complete and flounder-free. Technical report, ECRC Munich, January 1990.
- [Degerstedt and Nilsson, 1995] Lars Degerstedt and Ulf Nilsson. Magic Computation of Well-founded Semantics. In J. Dix, L. Pereira, and T. Przymusiński, editors, *Nonmonotonic Extensions of Logic Programming*, LNAI 927, pages 181–204. Springer, Berlin, 1995.
- [Dix and Müller, 1992] Jürgen Dix and Martin Müller. Abstract Properties and Computational Complexity of Semantics for Disjunctive Logic Programs. In *Proc. of the Workshop W1, Structural Complexity and Recursion-theoretic Methods in Logic Programming, following the JICSLP '92*, pages 15–28. H. Blair and W. Marek and A. Nerode and J. Remmel, November 1992. also available as Technical Report 13/93, University of Koblenz, Department of Computer Science.
- [Dix and Müller, 1993] Jürgen Dix and Martin Müller. Implementing Semantics for Disjunctive Logic Programs Using Fringes and Abstract Properties. In Luis Moniz Pereira and Anil Nerode, editors, *Logic Programming and Non-Monotonic Reasoning, Proceedings of the Second International Workshop*, pages 43–59, Cambridge, Mass., July 1993. Lisbon, MIT Press.
- [Dix and Müller, 1994a] Jürgen Dix and Martin Müller. An Axiomatic Framework for Representing and Characterizing Semantics of Disjunctive Logic Programs. In Pascal Van Hentenryck, editor, *Proceedings of the 11th Int. Conf. on Logic Programming, S. Margherita Ligure*, pages 303–322. MIT, June 1994.
- [Dix and Müller, 1994b] Jürgen Dix and Martin Müller. Partial Evaluation and Relevance for Approximations of the Stable Semantics. In Z.W. Ras and M. Zemankova, editors, *Proceedings of the 8th Int. Symp. on Methodologies for Intelligent Systems, Charlotte, NC, 1994*, LNAI 869, pages 511–520, Berlin, 1994. Springer.
- [Dix and Müller, 1994c] Jürgen Dix and Martin Müller. The Stable Semantics and its Variants: A Comparison of Recent Approaches. In L. Dreschler-Fischer and B. Nebel, editors, *Proceedings of the 18th German Annual Conference on Artificial Intelligence (KI '94), Saarbrücken, Germany*, LNAI 861, pages 82–93, Berlin, 1994. Springer.
- [Dix and Stolzenburg, 1998] Jürgen Dix and Frieder Stolzenburg. A Framework to incorporate Nonmonotonic Reasoning into Constraint Logic Programming. *Journal of Logic Programming*, 37(1,2,3):47–76, 1998. Special Issue on *Constraint Logic Programming*, Guest Editors: Kim Marriott and Peter Stuckey.
- [Dix *et al.*, 1994] Jürgen Dix, Georg Gottlob, and Viktor Marek. Causal Models for Disjunctive Logic Programs. In Pascal Van Hentenryck, editor, *Proceedings of the 11th Int. Conf. on Logic Programming, S. Margherita Ligure*, pages 290–302. MIT, June 1994.
- [Dix *et al.*, 1995] J. Dix, L. Pereira, and T. Przymusiński, editors. *Non-Monotonic Extensions of Logic Programming*, LNAI 927, Berlin, 1995. Springer.
- [Dix *et al.*, 1996a] Jürgen Dix, Georg Gottlob, and Viktor Marek. Reducing disjunctive to non-disjunctive semantics by shift-operations. *Fundamenta Informaticae*, XXVIII(1/2):87–100, 1996.
- [Dix *et al.*, 1996b] Jürgen Dix, Donald Loveland, Jack Minker, and David. S. Warren. Disjunctive Logic Programming and databases: Nonmonotonic Aspects. Technical Report Dagstuhl Seminar Report 150, IBFI GmbH, Schloß Dagstuhl, 1996.
- [Dix *et al.*, 1997a] J. Dix, U. Furbach, and A. Nerode, editors. *Logic Programming and Nonmonotonic Reasoning*, LNAI 1265, Berlin, 1997. Springer.
- [Dix *et al.*, 1997b] J. Dix, L. Pereira, and T. Przymusiński, editors. *Non-Monotonic Extensions of Logic Programming*, LNAI 1216, Berlin, 1997. Springer.
- [Dix *et al.*, 2000a] Jürgen Dix, Mirco Nanni, and V. S. Subrahmanian. Probabilistic agent reasoning. *Transactions of Computational Logic*, 1(2):201–245, 2000.
- [Dix *et al.*, 2000b] Jürgen Dix, V. S. Subrahmanian, and George Pick. Meta Agent Programs. *Journal of Logic Programming*, 46(1-2):1–60, 2000.

- [Dix *et al.*, 2001a] Jürgen Dix, Ulrich Furbach, and Ilkka Niemelä. Nonmonotonic Reasoning: Towards Efficient Calculi and Implementations. In Andrei Voronkov and Alan Robinson, editors, *Handbook of Automated Reasoning*. Elsevier-Science-Press, 2001.
- [Dix *et al.*, 2001b] Jürgen Dix, Sarit Kraus, and VS Subrahmanian. Temporal agent reasoning. *Artificial Intelligence*, 127(1):87–135, 2001.
- [Dix, 1991] Jürgen Dix. Classifying Semantics of Logic Programs. In Anil Nerode, Wiktor Marek, and V. S. Subrahmanian, editors, *Logic Programming and Non-Monotonic Reasoning, Proceedings of the first International Workshop*, pages 166–180, Cambridge, Mass., July 1991. Washington D.C., MIT Press.
- [Dix, 1992a] Jürgen Dix. A Framework for Representing and Characterizing Semantics of Logic Programs. In B. Nebel, C. Rich, and W. Swartout, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Third International Conference (KR '92)*, pages 591–602. San Mateo, CA, Morgan Kaufmann, 1992.
- [Dix, 1992b] Jürgen Dix. Classifying Semantics of Disjunctive Logic Programs. In K. R. Apt, editor, *LOGIC PROGRAMMING: Proceedings of the 1992 Joint International Conference and Symposium*, pages 798–812, Cambridge, Mass., November 1992. MIT Press.
- [Dix, 1995a] Jürgen Dix. A Classification-Theory of Semantics of Normal Logic Programs: I. Strong Properties. *Fundamenta Informaticae*, XXII(3):227–255, 1995.
- [Dix, 1995b] Jürgen Dix. A Classification-Theory of Semantics of Normal Logic Programs: II. Weak Properties. *Fundamenta Informaticae*, XXII(3):257–288, 1995.
- [Dix, 1995c] Jürgen Dix. Semantics of Logic Programs: Their Intuitions and Formal Properties. An Overview. In Andre Fuhrmann and Hans Rott, editors, *Logic, Action and Information – Essays on Logic in Philosophy and Artificial Intelligence*, pages 241–327. DeGruyter, 1995.
- [Dix, 1998] Jürgen Dix. The Logic Programming Paradigm. *AI Communications*, Vol. 11, No. 3:39–43, 1998. Short version in Newsletter of ALP, Vol. 11(3), 1998, pages 10–14.
- [Dix, 2001] Jürgen Dix. A Computational Logic Approach to Heterogenous Agent Systems. In Th. Eiter, M. Truszczyński, and W. Faber, editors, *Logic Programming and Non-Monotonic Reasoning, Proceedings of the Sixth International Conference*, LNCS, Berlin, September 2001. Springer.
- [Dowling and Gallier, 1984] W.F. Dowling and J.H. Gallier. Linear Time Algorithms for Testing the Satisfiability of Propositional Horn Formulae. *Journal of Logic Programming*, 1:267–284, 1984.
- [Drabent, 1994] Włodzimierz Drabent. What is failure? A constructive approach to negation. *Acta Informatica*, 32(1):27–29, 1994.
- [Dung, 1992] P. M. Dung. On the relations between stable and wellfounded semantics of logic programs. *Theoretical Computer Science*, 105:7–25, 1992.
- [Eiter and Gottlob, 1993a] T. Eiter and G. Gottlob. Complexity Aspects of Various Semantics for Disjunctive Databases. In *Proceedings of the Twelfth ACM SIGACT SIGMOD-SIGART Symposium on Principles of Database Systems (PODS-93)*, pages 158–167, June 1993.
- [Eiter and Gottlob, 1993b] Thomas Eiter and Georg Gottlob. Propositional Circumscription and Extended Closed World Reasoning are Π_2^P -complete. *Theoretical Computer Science*, 144(2):231–245, Addendum: vol. 118, p. 315, 1993, 1993.
- [Eiter and Gottlob, 1995] Thomas Eiter and Georg Gottlob. The complexity of logic-based abduction. *Journal of the ACM*, 42(1):3–42, 1995.
- [Eiter *et al.*, 1993] Thomas Eiter, Georg Gottlob, and Heikki Mannila. Expressive Power and Complexity of Disjunctive DATALOG. In *Proceedings of Workshop on Logic Programming with Incomplete Information, Vancouver Oct. 1993, following ILPS' 93*, pages 59–79, 1993.
- [Eiter *et al.*, 1997a] T. Eiter, J. Lu, and V. S. Subrahmanian. Computing Non-Ground Representations of Stable Models. In J. Dix, U. Furbach, and A. Nerode, editors, *Logic Programming and Non-Monotonic Reasoning, Proceedings of the Fourth International Conference*, LNAI 1265, pages 198–217, Berlin, July 1997. Springer.
- [Eiter *et al.*, 1997b] Thomas Eiter, Nicola Leone, Cristinel Mateis, Gerald Pfeifer, and Francesco Scarcello. A Deductive System for Nonmonotonic Reasoning. In Jürgen Dix, Ulrich Furbach, and Anil Nerode, editors, *Proceedings of the 4th International Conference on Logic Programming and Nonmonotonic Reasoning (LPNMR '97)*, number 1265 in Lecture Notes in AI (LNAI), Berlin, 1997. Springer.
- [Eiter *et al.*, 1998] Thomas Eiter, Nicola Leone, Cristinel Mateis, Gerald Pfeifer, and Francesco Scarcello. The KR System dlv: Progress Report, Comparisons and Benchmarks. In *Proceedings Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, pages 406–417, 1998.

- [Eiter *et al.*, 1999] Thomas Eiter, V. S. Subrahmanian, and Georg Pick. Heterogeneous Active Agents, I: Semantics. *Artificial Intelligence*, 108(1-2):179–255, 1999.
- [Eiter *et al.*, 2000] T. Eiter, V.S. Subrahmanian, and T. J. Rogers. Heterogeneous Active Agents, III: Polynomially Implementable Agents. *Artificial Intelligence*, 117(1):107–167, 2000.
- [Fages, 1993] F. Fages. Consistency of Clark’s completion and existence of stable models. *Methods of Logic in Computer Science*, 2, 1993.
- [Fernández and Minker, 1995] J. A. Fernández and J. Minker. Bottom-Up Computation of Perfect Models for Disjunctive Theories. *Journal of Logic Programming*, 25(1):33–51, 1995.
- [Fernández *et al.*, 1993] J. A. Fernández, J. Lobo, J. Minker, and V.S. Subrahmanian. Disjunctive LP + Integrity Constraints = Stable Model Semantics. *Annals of Mathematics and Artificial Intelligence*, 8(3-4), 1993.
- [Fitting, 1985] Melvin C. Fitting. A Kripke-Kleene Semantics of logic Programs. *Journal of Logic Programming*, 4:295–312, 1985.
- [Garey and Johnson, 1979] M.R. Garey and D.S. Johnson. *Computers and Intractability*. W.H. Freeman and Company, San Francisco, 1979.
- [Gelfond and Lifschitz, 1988] Michael Gelfond and Vladimir Lifschitz. The Stable Model Semantics for Logic Programming. In R. Kowalski and K. Bowen, editors, *5th Conference on Logic Programming*, pages 1070–1080. MIT Press, 1988.
- [Gelfond, 1994] Michael Gelfond. Logic programming and reasoning with incomplete information. *Annals of Mathematics and Artificial Intelligence*, 12:89–116, 1994.
- [Gogic *et al.*, 1995] Goran Gogic, Christos Papadimitriou, Bart Selman, and Henry Kautz. The Comparative Linguistics of Knowledge Representation. In *Proceedings of the 14th International Joint Conference on Artificial Intelligence*, pages 862–869, Montreal, Canada, August 1995. Morgan Kaufmann Publishers.
- [Gurevich, 1988] Y. Gurevich. Logic and the Challenge of Computer Science. In E. Börger, editor, *Trends in Theoretical Computer Science*, chapter 1. Computer Science Press, 1988.
- [Horty *et al.*, 1990] Jeff Horty, Richmond Thomason, and D. S. Touretzky. A skeptical Theory of Inheritance in Nonmonotonic Semantic Networks. *Artificial Intelligence*, 42:311–348, 1990.
- [Imielinski, 1991] T. Imielinski. Incomplete Deductive Databases. *Annals of Mathematics and Artificial Intelligence*, 3:259–294, 1991.
- [Inoue *et al.*, 1992] Katsumi Inoue, M. Koshimura, and R. Hasegawa. Embedding negation-as-failure into a model generation theorem prover. In Deepak Kapur, editor, *Automated Deduction — CADE-11*, number 607 in LNAI, Berlin, 1992. Springer.
- [Johnson, 1990] D.S. Johnson. A catalog of complexity classes. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume A. Algorithms and Complexity, pages 67–161. Elsevier, 1990.
- [Kagan *et al.*, 1994] Vadim Kagan, Anil Nerode, and V. S. Subrahmanian. Computing Definite Logic Programs by Partial Instantiation. *Annals of Pure and Applied Logic*, 67:161–182, 1994.
- [Kagan *et al.*, 1995] Vadim Kagan, Anil Nerode, and V. S. Subrahmanian. Computing Minimal Models by Partial Instantiation. *Theoretical Computer Science*, 155:157–177, 1995.
- [Kakas *et al.*, 2001] Toni Kakas, Robert Miller, and Francesca Toni. E-res: Reasoning about actions, events and observations. In *Proc. LPNMR 2001*. Springer, 2001.
- [Kemp *et al.*, 1991] David B. Kemp, Peter J. Stuckey, and Divesh Srivastava. Magic Sets and Bottom-Up Evaluation of Well-Founded Models. In Vijay Saraswat and Kazunori Ueda, editors, *Proceedings of the 1991 Int. Symposium on Logic Programming*, pages 337–351. MIT, June 1991.
- [Konolige, 1988] Kurt Konolige. Hierarchic autoepistemic theories for nonmonotonic reasoning. In *Proceedings AAAI ’88*. Morgan Kaufmann, 1988.
- [Kowalski, 1974] R.A. Kowalski. Predicate logic as a programming language. In *Proceedings IFIP’74*, pages 569–574. North Holland Publishing Company, 1974.
- [Kunen, 1987] Kenneth Kunen. Negation in Logic Programming. *Journal of Logic Programming*, 4:289–308, 1987.
- [Kunen, 1989] Kenneth Kunen. Signed Data Dependencies. *Journal of Logic Programming*, 7:231–245, 1989.
- [Lifschitz, 1985] Vladimir Lifschitz. Computing Circumscription. In *Proceedings of the International Joint Conference on Artificial Intelligence, Los Angeles, California*, pages 121–127, 1985.
- [Lifschitz, 1994] V. Lifschitz. *Circumscription*, pages 297–353. Clarendon, Oxford, 1994.
- [Lifschitz, 1996] V. Lifschitz. Foundations of declarative logic programming. In G. Brewka, editor, *Principles of Knowledge Representation*, chapter 3, pages 69–128. CSLI, 1996.

- [Lloyd, 1987] John W. Lloyd. *Foundations of Logic Programming*. Springer, Berlin, 1987. 2nd edition.
- [Lobo *et al.*, 1992] Jorge Lobo, Jack Minker, and Arcot Rajasekar. *Foundations of Disjunctive Logic Programming*. MIT-Press, 1992.
- [Ludäscher, 1991] Bertram Ludäscher. CNF-Prolog: A Meta-Interpreter for Chan's Constructive Negation, Implementation. Technical report, Master Thesis, Karlsruhe University (in German), 1991.
- [Mancarella *et al.*, 1988] Paolo Mancarella, Simone Martini, and Dino Pedreschi. Complete logic Programs with domain-closure Axiom. *Journal of Logic Programming*, 5:263–276, 1988.
- [Marek and Truszczyński, 1993] Wiktor Marek and Mirek Truszczyński. *Nonmonotonic Logics – Context-Dependent Reasoning*. Springer, 1993.
- [Marek *et al.*, 1992] Wiktor Marek, Arcot Rajasekar, and Mirek Truszczyński. Complexity of Computing with Extended Propositional Logic Programs. In *Proc. of the Workshop W1, Structural Complexity and Recursion-theoretic Methods in Logic Programming, following the IJCSLP '92*, pages 93–102. H. Blair and W. Marek and A. Nerode and J. Remmel, November 1992.
- [McCarthy, 1979] John McCarthy. First order theories of individual concepts and propositions. In B. Meltzer and D. Michie, editors, *Machine Intelligence 9*, pages 120–147. Edinburgh University Press, Edinburgh, 1979.
- [McCarthy, 1980] John McCarthy. Circumscription: A Form of Nonmonotonic Reasoning. *Artificial Intelligence*, 13:27–39, 1980.
- [McCarthy, 1986] John McCarthy. Applications of Circumscription to formalizing Common-Sense Reasoning. *Artificial Intelligence*, 28:89–116, 1986.
- [Minker and Ruiz, 1995] Jack Minker and Carolina Ruiz. Computing stable and partial stable models of extended disjunctive logic programs. In J. Dix, L. Pereira, and T. Przymusiński, editors, *Nonmonotonic Extensions of Logic Programming*, LNAI 927, pages 205–229. Springer, Berlin, 1995.
- [Minker, 1982] Jack Minker. On indefinite databases and the closed world assumption. In *Proceedings of the 6th Conference on Automated Deduction, New York*, pages 292–308, Berlin, 1982. Springer.
- [Minker, 1988] Jack Minker. *Foundations of Deductive Databases*. Morgan Kaufmann, 95 First Street, Los Altos, CA 94022, 1st edition, 1988.
- [Minker, 1993] Jack Minker. An Overview of Nonmonotonic Reasoning and Logic Programming. *Journal of Logic Programming, Special Issue*, 17(2/3/4):95–126, 1993.
- [Minker, 1996] Jack Minker. Logic and databases: A 20 year retrospective. In Dino Pedreschi and Carlo Zaniolo, editors, *Proceedings of the International Workshop on Logic in Databases (LID)*, LNCS 1154, pages 3–58. Springer, Berlin, 1996.
- [Moschovakis, 1974] Y. N. Moschovakis. *Elementary Induction on Abstract Structures*. North-Holland, 1974.
- [Müller and Dix, 1993] Martin Müller and Jürgen Dix. Implementing Semantics for Disjunctive Logic Programs Using Fringes and Abstract Properties. In Luis Moniz Pereira and Anil Nerode, editors, *Logic Programming and Non-Monotonic Reasoning, Proceedings of the Second International Workshop*, pages 43–59, Cambridge, Mass., July 1993. Lisbon, MIT Press.
- [Nebel, 1998] Bernhard Nebel. How hard is it to revise a belief base? In Didier Dubois and Henri Prade, editors, *Handbook on Defeasible Reasoning and Uncertainty Management Systems*, pages 77–145. Kluwer Academic, 1998.
- [Nerode *et al.*, 1991] Anil Nerode, Raymond T. Ng, and V.S. Subrahmanian. Computing Circumscriptive Deductive Databases. CS-TR 91-66, Computer Science Dept., Univ. Maryland, University of Maryland, College Park, Maryland, 20742, USA, December 1991.
- [Niemelä and Simons, 1996] Ilkka Niemelä and Patrik Simons. Efficient Implementation of the Well-founded and Stable Model Semantics. In M. Maher, editor, *Proceedings of the Joint International Conference and Symposium on Logic Programming*, pages 289–303, Bonn, Germany, September 1996. The MIT Press.
- [Niemelä, 1996a] Ilkka Niemelä. Implementing circumscription using a tableau method. In W. Wahlster, editor, *Proceedings of the European Conference on Artificial Intelligence*, pages 80–84, Budapest, Hungary, August 1996. John Wiley.
- [Niemelä, 1996b] Ilkka Niemelä. A tableau calculus for minimal model reasoning. In P. Miglioli, U. Moscato, D. Mundici, and M. Ornaghi, editors, *Proceedings of the Fifth Workshop on Theorem Proving with Analytic Tableaux and Related Methods*, pages 278–294, Terrasini, Italy, May 1996. LNAI 1071, Springer-Verlag.

- [Niemelä, 1999] Ilkka Niemelä. Logic Programs with Stable Model Semantics as a Constraint Programming Paradigm. *Annals of Mathematics and Artificial Intelligence, Special Issue on Logics in Artificial Intelligence*, edited by J. Dix and J. Lobo, 25(3–4):241–273, 1999.
- [Odifreddi, 1989] P. Odifreddi. *Classical Recursion Theory*. North-Holland, 1989.
- [Papadimitriou, 1994] C.H. Papadimitriou. *Computational Complexity*. Addison-Wesley, 1994.
- [Pereira and Alferes, 1992] L.M. Pereira and J.J. Alferes. Well founded semantics for logic programs with explicit negation. In Bernd Neumann, editor, *Proc. of 10th European Conf. on Artificial Intelligence ECAI 92*, pages 102–106. John Wiley & Sons, 1992.
- [Pereira et al., 1993] L. M. Pereira, J. N. Aparício, and J. J. Alferes. Non-Monotonic Reasoning with Logic Programming. *Journal of Logic Programming*, 17:227–264, 1993.
- [Poole, 1985] D. Poole. On the comparison of theories: Preferring the most specific explanation. In *Proc. IJCAI-85, Los Angeles*, 1985.
- [Prakken, 1993] H. Prakken. *Logical Tools for Modelling Legal Argument*. PhD thesis, VU Amsterdam, 1993.
- [Przymusiński, 1991] Teodor Przymusiński. Stationary Semantics for Normal and Disjunctive Logic Programs. In C. Delobel, M. Kifer, and Y. Masunaga, editors, *DOOD '91, Proceedings of the 2nd International Conference*, Berlin, December 1991. Muenchen, Springer. LNCS 566.
- [Przymusiński, 1995] Teodor Przymusiński. Static Semantics For Normal and Disjunctive Logic Programs. *Annals of Mathematics and Artificial Intelligence*, 14:323–357, 1995.
- [Rajasekar et al., 1989] Arcot Rajasekar, Jorge Lobo, and Jack Minker. Weak Generalized Closed World Assumption. *Journal of Automated Reasoning*, 5:293–307, 1989.
- [Reiter, 1978] R. Reiter. On Closed-World Databases. In H. Gallaire and J. Minker, editors, *Logic and Data Bases*, pages 55–76. Plenum Press, New York, 1978.
- [Reiter, 1980] Raymond Reiter. A Logic for Default-Reasoning. *Artificial Intelligence*, 13:81–132, 1980.
- [Ross and Topor, 1988] Kenneth A. Ross and Rodney A. Topor. Inferring negative Information from disjunctive Databases. *Journal of Automated Reasoning*, 4:397–424, 1988.
- [Ross, 1989] Kenneth A. Ross. The well-founded semantics for disjunctive logic programs. In *Proceedings of the first International Conference on Deductive and Object Oriented Databases, Kyoto, Japan*, pages 1–22, 1989.
- [Russel and Norvig, 1995] Stuart Russel and Peter Norvig. *Artificial Intelligence — A Modern Approach*. Prentice Hall, New Jersey 07458, 1995.
- [Sacca, 1993] Domenico Sacca. The Expressive Power of Stable Models For DATALOG Queries with Negation. In *Proceedings of Workshop on Logic Programming with Incomplete Information, Vancouver Oct. 1993, following ILPS' 93*, pages 150–162, 1993.
- [Sakama and Inoue, 1993] Chiaki Sakama and Katsumi Inoue. Negation in Disjunctive Logic Programs. In D. Warren and Peter Szeredi, editors, *Proceedings of the 10th Int. Conf. on Logic Programming, Budapest*, pages 703–719, Cambridge, Mass., July 1993. MIT Press.
- [Sakama and Inoue, 1994] Ch. Sakama and K. Inoue. An Alternative Approach to the Semantics of Disjunctive Logic Programs and Deductive Databases. *Journal of Automated Reasoning*, 13:145–172, 1994.
- [Sakama and Seki, 1994] Chiaki Sakama and Hirohisa Seki. Partial Deduction of Disjunctive Logic Programs: A Declarative Approach. In *Logic Program Synthesis and Transformation – Meta Programming in Logic*, LNCS 883, pages 170–182, Berlin, 1994. Springer.
- [Sakama and Seki, 1997] Chiaki Sakama and Hirohisa Seki. Partial Deduction in Disjunctive Logic Programming. *Journal of Logic Programming*, 32(3):229–245, 1997.
- [Sakama, 1989] Chiaki Sakama. Possible Model Semantics for Disjunctive Databases. In Won Kim, Jean-Marie Nicolas, and Shojiro Nishio, editors, *Deductive and Object-Oriented Databases, Proceedings of the First International Conference (DOOD89)*, pages 369–383, Kyoto, Japan, 1989. North-Holland Publ.Co.
- [Schaub and Wang, 2001] T. Schaub and K. Wang. A comparative study of logic programs with preference: Preliminary report. In A. Proveti and S. Cao, editors, *Proceedings of AAAI Spring Symposium on Answer Set Programming*, pages 151–157. AAAI Press, 2001.
- [Schlipf, 1990] John S. Schlipf. The Expressive Powers of the Logic Programming Semantics. In *Proceedings of the Ninth ACM Symposium on Principles of Databases*, pages 196–204, 1990.

- [Schlipf, 1992] John S. Schlipf. A Survey of Complexity and Undecidability Results in Logic Programming. In H. Blair, W. Marek, A. Nerode, and J. Remmel, editors, *Proceedings of the Workshop on Complexity and Recursion-theoretic Methods in Logic Programming, following the JICSLP'92*, informal, 1992.
- [Shepherdson, 1988] John C. Shepherdson. Negation in Logic Programming. In Jack Minker, editor, *Foundations of Deductive Databases*, chapter 1, pages 19–88. Morgan Kaufmann, 1988.
- [Shepherdson, 1991] John C. Shepherdson. Unsolvability Problems for SLDNF-Resolution. *Journal of Logic Programming*, 10:19–22, 1991.
- [Shepherdson, 1992] John C. Shepherdson. SLDNF-Resolution with Equality. *Journal of Automated Reasoning*, 8, No. 2:297–306, 1992.
- [Stärk, 1994] Robert F. Stärk. Input/output dependencies of normal logic programs. *Journal of Logic and Computation*, 4(3):249–262, 1994.
- [Subrahmanian *et al.*, 2000] V.S. Subrahmanian, Piero Bonatti, Jürgen Dix, Thomas Eiter, Sarit Kraus, Fatma Özcan, and Robert Ross. *Heterogenous Active Agents*. MIT-Press, 2000.
- [Swift, 1999] Terry Swift. Using Tabling for nonmonotonic Programming. *Annals of Mathematics and Artificial Intelligence, Special Issue on Logics in Artificial Intelligence, edited by J. Dix and J. Lobo*, 25(3–4):201–240, 1999.
- [Tamaki and Sato, 1986] H. Tamaki and T. Sato. OLD Resolution with Tabulation. In *Proceedings of the Third International Conference on Logic Programming, London*, LNAI, pages 84–98, Berlin, June 1986. Springer.
- [Tarski, 1955] A. Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5:285–309, 1955.
- [Touretzky *et al.*, 1988] David S. Touretzky, Jeff Harty, and Richmond Thomason. A Clash of Intuitions: The current State of Nonmonotonic Multiple IHS. In *Proceedings IJCAI*, 1988.
- [Touretzky *et al.*, 1991] D. S. Touretzky, R. H. Thomason, and J. F. Harty. A skeptic's menagerie: Conflictors, preemptors, reinstaters, and zombies in nonmonotonic inheritance. In *Proc. 12th IJCAI, Sydney*, 1991.
- [Touretzky, 1986] D. S. Touretzky. *The Mathematics of Inheritance*. Research Notes in Artificial Intelligence. Pitman, London, 1986.
- [Ullman, 1989a] J. D. Ullman. *Principles of Database and Knowledge Base Systems*. Computer Science Press, 1989.
- [Ullman, 1989b] Jeffrey D. Ullman. Bottom-up Beats Top-down for Datalog. In *Proc. of the Eight ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, Philadelphia, Pennsylvania*, pages 140–149. ACM Press, March 1989.
- [van Emden and Kowalski, 1976] M.H. van Emden and R.A. Kowalski. The semantics of predicate logic as a programming language. *JACM*, 23:733–742, 1976.
- [Van Gelder *et al.*, 1988] Allen Van Gelder, Kenneth A. Ross, and J. S. Schlipf. Unfounded Sets and well-founded Semantics for general logic Programs. In *Proceedings 7th Symposium on Principles of Database Systems*, pages 221–230, 1988.
- [Vorbeck, 1991] Martin Vorbeck. CNF-Prolog: A Meta-Interpreter for Chan's Constructive Negation, Theory. Technical report, Master Thesis, Karlsruhe University (in german), 1991.
- [Witteveen, 1991a] Cees Witteveen. Partial Semantics for Truth Maintenance. In J. van Eijck, editor, *Logics in AI*, LNAI 478, Berlin, 1991. Springer.
- [Witteveen, 1991b] Cees Witteveen. Skeptical Reason Maintenance is Tractable. In J. Allen, R. Fikes, and B. Sandewall, editors, *Proceedings of the second Conference on Principles of Knowledge Representation and Reasoning, Cambridge, Massachusetts*, pages 570–581. Morgan Kaufmann, 1991.

A APPENDIX

A.1 Predicate Logic

We assume the reader is familiar with the basic notions of predicate logic such as *models*, *formulae*, *satisfiability* \models and *derivability* \vdash . There exist several calculi for first-order predicate logic like Hilbert-style, Resolution-style, Gentzen-style or

natural deduction-style calculi. One of the main theorems states the completeness of such calculi with respect to the semantics given by models:

Theorem A.1 (Completeness).

A formula φ follows semantically from a theory T (is true in all models of T) iff φ is derivable from T by means of a particular calculus.

$$T \vdash \varphi \text{ iff } T \models \varphi$$

This theorem tells us that we can enumerate all the theorems of a theory, but it does not provide us with a decision-method to do so. In fact, as we will explain now, such a method does not exist.

Before turning to *undecidability*, let us emphasize that in the whole paper we are dealing with predicate logic *without equality* “ \doteq ”. But we can try to simulate “ \doteq ” as follows. We introduce a binary relation-symbol eq and require that it satisfies the following axioms with respect to an underlying language \mathcal{L} :

$$\forall x eq(x, x)$$

for all function-symbols f of suitable arity:

$$\forall x_1 \dots x_n, y_1 \dots y_n \\ (eq(x_1, y_1) \dots eq(x_n, y_n)) \rightarrow eq(f(x_1, \dots, x_n), f(y_1, \dots, y_n))$$

for all predicate-symbols P of suitable arity:

$$\forall x_1 \dots x_n y_1 \dots y_n \\ (eq(x_1, y_1) \dots eq(x_n, y_n)) \rightarrow (P(x_1, \dots, x_n) \rightarrow P(y_1, \dots, y_n)).$$

This set, is denoted by $EQ_{\mathcal{L}}$. It can be shown that transitivity and symmetry of eq follow from these axioms. Let us consider the language of *Arithmetic* \mathcal{L}_{Ar} which contains: 0 (a constant), s (a unary function-symbol), eq (a two-ary relation-symbol) and \oplus, \otimes (ternary relation-symbols).

We have in mind to axiomatize the theory of natural numbers. Before we do so we introduce the following abbreviation. The formula $\exists!z \phi(z)$ stands for

$$\exists z (\phi(z) \wedge \forall y \phi(y) \rightarrow eq(y, z)).$$

Definition A.1 (Arithmetic Ar_{fin}).

Ar_{fin} is the finite set consisting of $EQ_{\mathcal{L}_{Ar}}$ and the following axioms:

$$\begin{aligned} \forall x \forall y \exists!z & \quad \oplus(x, y, z) \\ \forall x & \quad \oplus(x, 0, x) \\ \forall x \forall y \forall z & \quad \oplus(x, y, z) \rightarrow \oplus(x, s(x), s(z)) \end{aligned}$$

$$\begin{aligned} \forall x \forall y \exists!z & \quad \otimes(x, y, z) \\ \forall x & \quad \otimes(x, 0, 0) \\ \forall x \forall y \forall z \forall z' & \quad \otimes(x, y, z) \rightarrow (\otimes(x, s(y), z') \wedge \oplus(z, x, z')) \end{aligned}$$

The set of natural numbers $\mathcal{N} := (\mathbb{N}, 0^{\mathcal{N}}, s^{\mathcal{N}}, \oplus^{\mathcal{N}}, \otimes^{\mathcal{N}}, eq^{\mathcal{N}})$ is a model of Ar_{fin} . Here $0^{\mathcal{N}}$ is the “true” 0, $s^{\mathcal{N}}$ is the successor-function, $\oplus^{\mathcal{N}}$ is addition and $\otimes^{\mathcal{N}}$ is multiplication (viewed as relations), $eq^{\mathcal{N}}$ is identity. We note the following facts:

1. The set $\{\phi : Ar_{fin} \models \phi\}$ is recursively enumerable but not recursive.
2. The set $\{\phi : \mathcal{N} \models \phi\}$ is not even recursively enumerable.

We even have

Theorem A.2 (Gödel).

No set of formulae containing Ar_{fin} and having \mathcal{N} as model, is recursive.

Every recursively enumerable set of formulae Φ that contains Ar_{fin} and has \mathcal{N} as a model, is incomplete, i.e. there is ψ with: $\mathcal{N} \models \psi$ but $\Phi \not\models \psi$. Therefore no complete axiomatization of \mathcal{N} is possible.

Note that, although \mathcal{N} formally is not a Herbrand model, it is isomorphic to such a model. In fact, the axioms immediately imply that there is, up to isomorphy, only one single Herbrand-model of Ar_{fin} with respect to \mathcal{L}_{Ar} . Therefore to determine if a formula is true in all Herbrand-models of Ar_{fin} is just as complicated as the theory of \mathcal{N} itself. \mathcal{N} contains, for example, famous statements (or there negation) from number theory like Goldbach-conjecture or Fermat’s last theorem.

A.2 Complexity Theory

We assume some familiarity with the classes P(problems solvable in deterministic polynomial time) and NP(problems solvable in nondeterministic polynomial time). The class **co-NP** is the complement of NP, i.e. a problem is in **co-NP** if its complement is in NP. From these sets we can build larger classes by considering problems solvable in deterministic (resp. nondeterministic) time where we allow to ask queries to an NP-oracle: i.e. whenever we come up with a subproblem that lies in NP, we just ask an oracle which immediately gives us the answer (we count this as just one step). This gives rise to the polynomial hierarchy:

Definition A.2 (Polynomial Hierarchy).

For a complexity class C we denote by P^C (resp. NP^C) the class of problems solvable in deterministic polynomial (resp. nondeterministic polynomial) time using C -oracles. Let $\Sigma_0 := \Pi_0 := P$ and

$$\begin{aligned}\Sigma_{k+1} &:= NP^{\Sigma_k} \\ \Pi_{k+1} &:= \text{co-NP}^{\Sigma_k} \\ \Delta_{k+1} &:= P^{\Sigma_k}\end{aligned}$$

Thus Σ_1 is NP with queries to a P-oracle, i.e. $\Sigma_1 = NP$. Similarly we have $\Pi_1 = \text{co-NP}$ and $\Delta_1 = P$. A problem is in $\Delta_2 = P^{NP}$ if it can be solved in deterministic polynomial time with subcalls to an NP-oracle. Although the index is 2, Δ_2 is considered to belong to the first level of the polynomial hierarchy.

The second level of this hierarchy consists of Σ_2 , Π_2 and Δ_3 . Here $\Sigma_2 := \text{NP}^{\text{NP}}$: nondeterministic polynomial time with queries to an NP-oracle. $\Pi_2 := \text{co-NP}^{\text{NP}}$ and $\Delta_3 := \text{P}^{[\text{NP}^{\text{NP}}]}$.

It is immediate that

$$\Sigma_k \cup \Pi_k \subseteq \Delta_{k+1} \subseteq \Sigma_{k+1} \cap \Pi_{k+1}$$

but it has not yet been proved that the inclusions are proper. That is, it is not known if the hierarchy collapses at some point or not.

The polynomial hierarchy classifies a subclass of all decidable problems, namely those that are NP-hard. A problem is called NP-hard if any other problem in NP can be polynomially reduced to it. Of particular interest are those problems in a class Π_k or Σ_k that are *the hardest ones*: they are called *complete*. This means that all problems in the respective class can be polynomially reduced to such a complete problem and the problem itself belongs to this class. As an example, to determine if a formula is valid is **co-NP**-complete. Thus, satisfiability of a propositional formula is NP-complete.

An analogue hierarchy exists (in fact it was the prototype of the polynomial hierarchy) for undecidable problems. The notation is analog to the one just introduced. Therefore one often adds a superscript P to the Π_k and Σ_k which stands for *polynomial* (but not for an oracle) to denote the polynomial hierarchy.

To introduce the arithmetical hierarchy we consider the model \mathcal{N} of the natural numbers and \mathcal{L}_{Ar} -formulae. We call such formulae for short *arithmetical*. We classify arithmetical formulae according to their quantifier-alternations:

Definition A.3 (Arithmetical Hierarchy).

We call an arithmetical formula Σ_k^0 (resp. Π_k^0) if it is of the form $\exists \forall \dots \phi$ (resp. $\forall \exists \dots \phi$) where ϕ is quantifier-free and there are at most $k - 1$ alternations of quantifier-blocks.

We call a set M of natural numbers Σ_k^0 -definable, if M is definable by a Σ_k^0 -formula. This means that there is a Σ_k^0 -formula $\phi(x)$ with one free variable x such that

$$\mathcal{N} \models \phi(i) \text{ iff } i \in M.$$

Note that the Σ_0^0 -definable sets coincide with the Π_0^0 -definable ones: they are exactly the recursive sets. The recursive enumerable sets are the Σ_1^0 -definable ones, the Π_1^0 -definable sets are their complements. The set corresponding to the famous *Halting Problem*, i.e. the set of all Gödel numbers of those Turing-machines that stop on their own Gödel number, is Σ_1^0 , so this problem is located very low in the hierarchy.

The higher a problem lies in the hierarchy, the more undecidable it is. For example a problem located at the second level, say Σ_2^0 , can be thought of as *being recursively enumerable using an oracle which solves Σ_1^0 -problems* (like the halting problem).

Analogously to the polynomial hierarchy we have the notions of Σ_k^0 -complete and Π_k^0 -complete. As an example, the halting problem is Σ_1^0 -complete.

In contrast to the polynomial hierarchy, the arithmetical hierarchy is strict. We denote by Δ_k^0 the intersection of Σ_{k+1}^0 and Π_{k+1}^0 . We have

$$\Sigma_k^0 \cup \Pi_k^0 \subset \Delta_k^0 = \Sigma_{k+1}^0 \cap \Pi_{k+1}^0.$$

Are there more undecidable problems, not yet captured by our hierarchy? Yes, take for example the theory of \mathcal{N} considered in Section A.1. Obviously, the general problem to determine if an arbitrary formula is true or not in \mathcal{N} can not be captured at a certain level, because the class of formulae in question can have unlimited alternations of quantifiers. The careful reader may have asked himself what the superscript 0 means in Σ_k^0 ? It just means that we consider just first-order formulae and we do not allow our arithmetical formulae to contain second-order quantifiers.

This remark gives rise to the analytical hierarchy, denoted by Σ_k^1, Π_k^1 , where we consider *second-order* arithmetical formulae. We only count the alternations of the quantifiers over sets. So any Σ_k^0 -formula is in Σ_0^1 .

Note that for the arithmetical hierarchy the identity $\Sigma_0^0 = \Sigma_1^0 \cap \Pi_1^0$ holds. The analogue for the analytical hierarchy does not hold. A counterexample is given by the theory of the natural numbers \mathcal{N} : the set of true sentences in arithmetic is in $\Sigma_1^1 \cap \Pi_1^1$ but not in Σ_0^1 . This set is also called *hyperarithmetical* for obvious reasons.

For a more detailed treatment of the topics in this section we refer the reader to the standard literature: [Balcázar *et al.*, 1988; Garey and Johnson, 1979; Johnson, 1990] and [Papadimitriou, 1994; Odifreddi, 1989] for undecidability.

A.3 Default Logic

Reiter's default logic [Reiter, 1980] is one of the most prominent nonmonotonic logics. Default logic assumes knowledge to be represented in terms of a default theory. A default theory is a pair (D, W) . W is a set of first order formulae representing the facts which are known to be true with certainty. D is a set of defaults the form

$$\frac{A : B_1, \dots, B_n}{C}$$

where A, B_i and C are classical formulae. We will also frequently use the alternative, less space consuming notation $A:B_1, \dots, B_n/C$ for this default. The default has the intuitive reading: if A is provable and, for all i ($1 \leq i \leq n$), $\neg B_i$ is not provable, then derive C . A is called the *prerequisite*, B_i a *consistency condition* or *justification*, and C the *consequent* of the default. For a default d we use $pre(d)$, $just(d)$, and $cons(d)$ to denote the prerequisite, the set of justifications, and the consequent of d , respectively. Open defaults, i.e., defaults with free variables, are usually interpreted as schemata representing all of their closed instances.¹⁹

¹⁹Reiter treats open defaults somewhat differently and uses a more complicated method to define extensions for them.

Default theories induce so-called extensions which represent acceptable belief sets a reasoner may adopt based on the available information. A formula p is called a skeptical consequence of (D, W) iff p is contained in all extensions of (D, W) . p is called a credulous consequence of (D, W) iff p is contained in at least one extension of (D, W) .

We will first present a definition of extensions which is slightly different from (but equivalent to) Reiter's original definition. We have found that this definition is somewhat easier to digest. The original definition will be presented later.

Intuitively, E is an extension of (D, W) iff E is a deductively closed (in the sense of classical logic) superset of W satisfying the following two properties

1. all defaults that are "applicable" with respect to E have been applied,
2. every formula in E has a "derivation" from W and applicable defaults.

To make the two requirements more precise we introduce the following notion:

Definition A.4 (Default Proof).

Let (D, W) be a default theory, S a set of formulae, and p a formula. A (D, W) -default proof for p is a finite sequence $P = (d_1, \dots, d_n)$ of defaults in D such that:

1. $W \cup \{cons(d_1), \dots, cons(d_{i-1})\} \vdash pre(d_i)$, for $i \in \{1, \dots, n\}$,
2. $W \cup \{cons(d_1), \dots, cons(d_n)\} \vdash p$.

P is valid in S iff S does not contain the negation of a justification of a default in P .

As usual \vdash denotes classical provability. We now can state the definition of extensions formally:

Definition A.5 (Extension 1).

Let (D, W) be a default theory. E is an extension of (D, W) iff E is a deductively closed superset of W satisfying the conditions

1. if $A:B_1, \dots, B_n/C \in D$, $A \in E$ and for all i ($1 \leq i \leq n$) $\neg B_i \notin E$, then C in E , and
2. $p \in E$ implies there is a (D, W) -default proof for p valid in E .

Reiter's equivalent original definition is more compact. It defines extensions as fixpoints of a certain operator.

Definition A.6 (Extension 2).

Let (D, W) be a default theory, S a set of formulae. Let $\Gamma(S)$ be the smallest set such that:

1. $W \subseteq \Gamma(S)$,

$$2. Th(\Gamma(S)) = \Gamma(S),$$

$$3. \text{ if } A:B_1, \dots, B_n/C \in D, A \in \Gamma(S), \neg B_i \notin S \ (1 \leq i \leq n), \text{ then } C \in \Gamma(S).$$

E is an extension of (D, W) iff $E = \Gamma(E)$, that is, if E is a fixpoint of Γ .

We finally give a third, quasi-inductive characterization of extensions, also due to Reiter. This version is often used in proofs about default logic and makes the way in which formulae have to be grounded in the premises more explicit. Let E be a set of formulae and define, for a given default theory (D, W) , a sequence of sets of formulae as follows:

$$E_0 = W, \text{ and for } i \geq 0 \\ E_{i+1} = Th(E_i) \cup \{C \mid A:B_1, \dots, B_n/C \in D, A \in E_i, \neg B_i \notin E\}.$$

It can be shown that E is an extension of (D, W) iff $E = \bigcup_{i=0}^{\infty} E_i$. The appearance of E in the definition of E_{i+1} is what renders this alternative definition of extensions non-constructive.

Default theories may have an arbitrary number of extensions (including zero). Extensions are always consistent if W is and if there are no degenerate defaults without consistency conditions. If W is inconsistent then the single extension of (D, W) is the set of all formulae. Extensions are maximal in the following sense: if E is an extension then there is no extension E' such that $E' \subset E$.

A.4 Circumscription

Circumscription is a method of computing the closure of a theory by restricting its models to those that have minimal extensions of some of the predicates and functions. Since its first formulation by McCarthy [McCarthy, 1980], it has taken on several different forms, including domain circumscription [McCarthy, 1979] (minimizing the elements in the universe of models), and the most popular and useful version, parallel predicate circumscription [McCarthy, 1980; McCarthy, 1986; Lifschitz, 1985] which we present here.

Although circumscription was originally presented as a schema for adding more formulae to a theory (just as Clark's completion does), here we describe it in terms of restricting the models of the theory. This view leads to the generalization of circumscription by model preference theories, and is more useful analytically in relating circumscription to other nonmonotonic formalisms. More detailed references to circumscription can be found in Lifschitz' excellent survey article [Lifschitz, 1994].

Choose a language \mathcal{L} , and let P be the set of predicate symbols that we are interested in minimizing, and Z another set of predicate symbols whose interpretation we allow to vary across compared models. For example, if we wish to minimize the number of cannibals, we would let $P = \{C\}$, and Z be all other predicate symbols (the importance of Z will be indicated later). Suppose A is a theory containing the statements $C(p_1)$, $C(p_2)$, and $C(p_3)$, but no other assertions




Figure 5. Two Models of a Theory with the Same Valuation.

using C . Then every model of A will have at least the individuals referred to by p_1, p_2 , and p_3 with property C . Now consider two models with the same valuation function from terms to individuals, as in Figure 5. In model M_1 , the extension of the predicate C includes just the three individuals e_1, e_2 , and e_3 . In model M_2 there is a fourth individual, e_4 , who is a cannibal. Circumscription would prefer M_1 to M_2 , since the extension of C in M_1 is a proper subset of its extension in M_2 . Under appropriate assumptions (that these terms refer to different individuals), circumscription would yield the result $\neg C(p_4)$, which is not present in the original theory.

Let $A(P, Z)$ be a first-order sentence containing the symbols P and Z . Circumscription prefers models of $A(P, Z)$ that are minimal in the predicates P , assuming that these models have the same interpretation for all symbols not in P or Z . A may contain predicates other than P and Z ; these are called the *fixed symbols*.

To state this more formally, let M_1 and M_2 be two models of $A(P, Z)$. $|M|$ is the universe of model M , and $M[[K]]$ is the interpretation of the symbol K in M . Then

Definition A.7 (Minimal Models).

$$M_1 \leq^{P;Z} M_2 \text{ iff } \begin{cases} 1. & |M_1| = |M_2|. \\ 2. & M_1[[K]] = M_2[[K]] \text{ for all } K \text{ not in } P, Z. \\ 3. & M_1[[P_i]] \subseteq M_2[[P_i]] \text{ for all } P_i \in P. \end{cases}$$

$\leq^{P;Z}$ is a preorder relation (reflexive and transitive) on models, but not necessarily a partial order, since it is not antireflexive. We define the strict order $M_1 <^{(P;Z)} M_2$ as $M_1 \leq^{P;Z} M_2$ and not $M_2 \leq^{P;Z} M_1$. The preferred models of $A(P, Z)$ are those that are minimal according to the strict ordering.