

KNOWLEDGE REPRESENTATION,
REASONING AND DECLARATIVE
PROBLEM SOLVING

CHITTA BARAL

Arizona State University



CAMBRIDGE
UNIVERSITY PRESS

PUBLISHED BY THE PRESS SYNDICATE OF THE UNIVERSITY OF CAMBRIDGE
The Pitt Building, Trumpington Street, Cambridge, United Kingdom

CAMBRIDGE UNIVERSITY PRESS
The Edinburgh Building, Cambridge CB2 2RU, UK
40 West 20th Street, New York, NY 10011-4211, USA
477 Williamstown Road, Port Melbourne, VIC 3207, Australia
Ruiz de Alarcón 13, 28014 Madrid, Spain
Dock House, The Waterfront, Cape Town 8001, South Africa

<http://www.cambridge.org>

© Cambridge University Press 2003

This book is in copyright. Subject to statutory exception
and to the provisions of relevant collective licensing agreements,
no reproduction of any part may take place without
the written permission of Cambridge University Press.

First published 2003

Printed in the United Kingdom at the University Press, Cambridge

Typeface Times 11/14 pt *System* L^AT_EX 2_ε [TB]

A catalogue record for this book is available from the British Library

Library of Congress Cataloguing in Publication data

Baral, Chitta.

Knowledge representation, reasoning, and declarative problem solving / Chitta Baral.

p. cm.

Includes bibliographical references and index.

ISBN 0 521 81802 8

1. Expert systems (Computer science) 2. Artificial intelligence. 3. Knowledge
representation (Information theory). I. Title.

QA76.76.E95 B265 2002

006.3'3-dc21 2002025622

ISBN 0 521 81802 8 hardback

Contents

	<i>page</i>
<i>Preface</i>	<i>ix</i>
1 Declarative programming in AnsProlog*: introduction and preliminaries	1
1.1 Motivation: Why AnsProlog*?	3
1.2 Answer set frameworks and programs	8
1.3 Semantics of AnsProlog* programs	16
1.4 Database queries and AnsProlog* functions	40
1.5 Notes and references	44
2 Simple modules for declarative programming with answer sets	46
2.1 Declarative problem solving modules	47
2.2 Knowledge representation and reasoning modules	73
2.3 Notes and references	81
3 Principles and properties of declarative programming with answer sets	83
3.1 Basic notions and basic properties	84
3.2 Some AnsProlog* sub-classes and their basic properties	93
3.3 Restricted monotonicity and signed AnsProlog* programs	108
3.4 Analyzing AnsProlog* programs using ‘splitting’	113
3.5 Language independence and language tolerance	120
3.6 Interpolating an AnsProlog program	126
3.7 Building and refining programs from components: functional specifications and realization theorems	137
3.8 Filter-abducible AnsProlog ^{¬,or} programs	144
3.9 Equivalence of programs and semantics preserving transformations	154
3.10 Notes and references	168
4 Declarative problem solving and reasoning in AnsProlog*	170
4.1 Three well-known problem solving tasks	170
4.2 Constraint satisfaction problems (CSPs)	183

4.3	Dynamic constraint satisfaction problems (DCSPs)	186
4.4	Combinatorial graph problems	188
4.5	Prioritized defaults and inheritance hierarchies	192
4.6	Notes and references	197
5	Reasoning about actions and planning in AnsProlog*	199
5.1	Reasoning in the action description language \mathcal{A}	199
5.2	Reasoning about actions and plan verification in richer domains	229
5.3	Answer set planning examples in extensions of \mathcal{A} and STRIPS	244
5.4	Approximate planning when initial state is incomplete	261
5.5	Planning with procedural constraints	262
5.6	Explaining observations through action occurrences and application to diagnosis	269
5.7	Case study: Planning and plan correctness in a space shuttle reaction control system	274
5.8	Notes and references	277
6	Complexity, expressiveness, and other properties of AnsProlog* programs	278
6.1	Complexity and expressiveness	278
6.2	Complexity of AnsDatalog* sub-classes	288
6.3	Expressiveness of AnsDatalog* sub-classes	301
6.4	Complexity and expressiveness of AnsProlog* sub-classes	304
6.5	Compact representation and compilability of AnsProlog	311
6.6	Relationship with other knowledge representation formalisms	313
6.7	Notes and references	341
7	Answer set computing algorithms	345
7.1	Branch and bound with WFS: wfs-bb	346
7.2	The assume-and-reduce algorithm of SLG	358
7.3	The smodels algorithm	363
7.4	The dlv algorithm	372
7.5	Notes and references	379
8	Query answering and answer set computing systems	382
8.1	Smodels	382
8.2	The dlv system	403
8.3	Applications of answer set computing systems	412
8.4	Pure PROLOG	440
8.5	Notes and references	456
9	Further extensions of and alternatives to AnsProlog*	458
9.1	AnsProlog ^{not, or, \neg, \perp} : allowing not in the head	458
9.2	AnsProlog ^{{not, or, \neg, \perp}*} : allowing nested expressions	461
9.3	AnsProlog ^{\neg, or, K, M} : allowing knowledge and belief operators	466

9.4	Abductive reasoning with AnsProlog: AnsProlog^{abd}	470
9.5	Domain closure and the universal query problem	471
9.6	AnsProlog_{set} : adding set constructs to AnsProlog	475
9.7	$\text{AnsProlog}\text{-}\langle\bar{\text{ }}\bar{\text{ }}\bar{\text{ }}$ programs: $\text{AnsProlog}\bar{\text{ }}$ programs with ordering	477
9.8	Well-founded semantics of programs with AnsProlog syntax	482
9.9	Well-founded semantics of programs with $\text{AnsProlog}\bar{\text{ }}$ syntax	490
9.10	Notes and references	492
Appendix A: Ordinals, lattices, and fixpoint theory		494
Appendix B: Turing machines		496
<i>Bibliography</i>		498
<i>Index of notation</i>		519
<i>Index of terms</i>		522

Chapter 1

Declarative programming in AnsProlog*: introduction and preliminaries

Among other characteristics, an intelligent entity – whether an intelligent autonomous agent, or an intelligent assistant – must have the ability to go beyond just following direct instructions while in pursuit of a goal. This is necessary to be able to behave intelligently when the assumptions surrounding the direct instructions are not valid, or there are no direct instructions at all. For example even a seemingly direct instruction of ‘bring me coffee’ to an assistant requires the assistant to figure out what to do if the coffee pot is out of water, or if the coffee machine is broken. The assistant will definitely be referred to as lacking intelligence if he or she were to report to the boss that there is no water in the coffee pot and ask the boss what to do next. On the other hand, an assistant will be considered intelligent if he or she can take a high level request of ‘make travel arrangements for my trip to International AI conference 20XX’ and figure out the lecture times of the boss; take into account airline, hotel and car rental preferences; take into account the budget limitations, etc.; overcome hurdles such as the preferred flight being sold out; and make satisfactory arrangements. This example illustrates *one benchmark of intelligence – the level of request an entity can handle*. At one end of the spectrum the request is a detailed algorithm that spells out *how* to satisfy the request, which no matter how detailed it is may not be sufficient in cases where the assumptions inherent in the algorithm are violated. At the other end of the spectrum the request spells out *what* needs to be done, and the entity has the knowledge – again in the *what* form rather than the *how* form – and the knowledge processing ability to figure out the exact steps (that will satisfy the request) and execute them, and when it does not have the necessary knowledge it either knows where to obtain the necessary knowledge, or is able to gracefully get around its ignorance through its ability to reason in the presence of incomplete knowledge.

The languages for spelling out *how* are often referred to as *procedural* while the languages for spelling out *what* are referred to as *declarative*. Thus our initial thesis that intelligent entities must be able to comprehend and process descriptions of *what*

leads to the necessity of inventing suitable declarative languages and developing support structures around those languages to facilitate their use. We consider the development of such languages to be fundamental to knowledge based intelligence, perhaps similar to the role of the language of calculus in mathematics and physics. *This book is about such a declarative language – the language of **AnsProlog***.* We now give a brief history behind the quest for a suitable declarative language for knowledge representation, reasoning, and declarative problem solving.

Classical logic which has been used as a specification language for procedural programming languages was an obvious initial choice to represent declarative knowledge. But it was quickly realized that classical logic embodies the monotonicity property according to which the conclusion entailed by a body of knowledge stubbornly remains valid no matter what additional knowledge is added. This disallowed human like reasoning where conclusions are made with the available (often incomplete) knowledge and may be withdrawn in the presence of additional knowledge. This led to the development of the field of *nonmonotonic logic*, and several nonmonotonic logics such as circumscription, default logic, auto-epistemic logic, and nonmonotonic modal logics were proposed. The AI journal special issue of 1980 (volume 13, numbers 1 and 2) contained initial articles on some of these logics. In the last twenty years there have been several studies on these languages on issues such as representation of small common-sense reasoning examples, alternative semantics of these languages, and the relationship between the languages. But the dearth of efficient implementations, use in large applications – say of more than ten pages, and studies on building block support structures has for the time being diminished their applicability. Perhaps the above is due to some fundamental deficiency, such as: all of these languages which build on top of the classical logic syntax and allow nesting are quite complex, and all except default logic lack structure, thus making it harder to use them, analyze them, and develop interpreters for them.

An alternative nonmonotonic language paradigm with a different origin whose initial focus was to consider a subset of classical logic (rather than extending it) is the programming language PROLOG and the class of languages clubbed together as ‘logic programming’. PROLOG and logic programming grew out of work on automated theorem proving and Robinson’s resolution rule. One important landmark in this was the realization by Kowalski and Colmerauer that logic can be used as a programming language, and the term PROLOG was developed as an acronym from PROgramming in LOGic. A subset of first-order logic referred to as Horn clauses that allowed faster and simpler inferencing through resolution was chosen as the starting point. The notion of closed world assumption (CWA) in databases was then imported to PROLOG and logic programming and the negation as failure operator **not** was used to refer to negative information. The evolution of PROLOG was guided by concerns that it be made a full fledged programming language with

efficient implementations, often at the cost of sacrificing the declarativeness of logic. Nevertheless, research also continued on logic programming languages with declarative semantics. In the late 1980s and early 1990s the focus was on finding the right semantics for agreed syntactic sub-classes. One of the two most popular semantics proposed during that time is the *answer set semantics*, also referred to as the *stable model semantics*.

This book is about the language of logic programming with respect to the answer set semantics. We refer to this language as AnsProlog*, as a short form of ‘**Programming in logic with Answer sets**’¹. In the following section we give an overview of how AnsProlog* is different from PROLOG and also the other non-monotonic languages, and present the case for AnsProlog* to be the most suitable declarative language for knowledge representation, reasoning, and declarative problem solving.

1.1 Motivation: Why AnsProlog*?

In this section², for the purpose of giving a quick overview without getting into a lot of terminology, we consider an AnsProlog* program to be a collection of rules of the form:

$$L_0 \text{ or } \dots \text{ or } L_k \leftarrow L_{k+1}, \dots, L_m, \mathbf{not} L_{m+1}, \dots, \mathbf{not} L_n.$$

where each of the L_i s is a literal in the sense of classical logic. Intuitively, the above rule means that if L_{k+1}, \dots, L_m are to be true and if L_{m+1}, \dots, L_n can be safely assumed to be false then at least one of L_0, \dots, L_k must be true.

This simple language has a lot going for it to be the leading language for knowledge representation, reasoning, and declarative problem solving. To start with, the nonclassical symbols \leftarrow , and **not** in AnsProlog* give it a structure and allow us to easily define syntactic sub-classes and study their properties. It so happens that these various sub-classes have a range of complexity and expressiveness thus allowing us to choose the appropriate sub-classes for particular applications. Moreover, there exists a more tractable approximate characterization which can be used – at the possible cost of completeness – when time is a concern. Unlike the other nonmonotonic logics, AnsProlog* now has efficient implementations which have been used to program large applications. In addition, the expressiveness studies show AnsProlog* to be as expressive as some of these logics, while syntactically it seems less intimidating as it does not allow arbitrary formulas. Finally, the most important reason to study and use AnsProlog* is that there is now a large body (much larger than for any other knowledge representation language) of support structure around AnsProlog*

¹ In the recent literature it has also been referred to as A-Prolog [BGN00, Gel01].

² In Section 1.2 we introduce more specific terminologies and use those in the rest of the book.

that includes the above mentioned implementations and theoretical building block results that allow systematic construction of AnsProlog* programs, and assimilation of new information. We now expand on these points in greater detail.

1.1.1 AnsProlog* vs PROLOG

Although, PROLOG grew out of programming with Horn clauses – a subset of first-order logic, several nondeclarative features were included in PROLOG to make it programmer friendly. We propose AnsProlog* as a declarative alternative to PROLOG. Besides the fact that AnsProlog* allows disjunction in the head of rules, the following are the main differences between AnsProlog* and Prolog.

- The ordering of literals in the body of a rule matters in PROLOG as it processes them from left to right. Similarly, the positioning of a rule in the program matters in PROLOG as it processes them from start to end. The ordering of rules and positioning of literals in the body of a rule do not matter in AnsProlog*. From the perspective of AnsProlog*, a program is a *set* of AnsProlog* rules, and in each AnsProlog* rule, the body is a *set* of literals and literals preceded by **not**.
- Query processing in PROLOG is top-down from query to facts. In AnsProlog* query-processing methodology is not part of the semantics. Most sound and complete interpreters with respect to AnsProlog* do bottom-up query processing from facts to conclusions or queries.
- Because of the top-down query processing, and start to end, and left to right processing of rules and literals in the body of a rule respectively, a PROLOG program may get into an infinite loop for even simple programs without negation as failure.
- The *cut* operator in PROLOG is extra-logical, although there have been some recent attempts at characterizing it. This operator is not part of AnsProlog*.
- There are certain problems, such as floundering and getting stuck in a loop, in the way PROLOG deals with negation as failure. In general, PROLOG has trouble with programs that have recursions through the negation as failure operator. AnsProlog* does not have these problems, and as its name indicates it uses the *answer set* semantics to characterize negation as failure.

In this book, besides viewing AnsProlog* as a declarative alternative to PROLOG, we also view PROLOG systems as top-down query answering systems that are correct with respect to a sub-class of AnsProlog* under certain conditions. In Section 8.4 we present these conditions and give examples that satisfy these conditions.

1.1.2 AnsProlog* vs Logic programming

AnsProlog* is a particular kind of logic programming. In AnsProlog* we fix the semantics to *answer set semantics*, and only focus on that. On the other hand logic programming refers to a broader agenda where different semantics are considered

as alternatives. We now compare AnsProlog (a sub-class of AnsProlog* with only one atom in the head, and without classical negation in the body) with the alternative semantics of programs with AnsProlog syntax.

Since the early days of logic programming there have been several proposals for semantics of programs with AnsProlog syntax. We discuss some of the popular ones in greater detail in Chapter 9. Among them, the most popular ones are the *stable model semantics* and the *well-founded semantics*. The stable models are same as the answer sets of AnsProlog programs, the main focus of this book. The well-founded semantics differs from the stable model semantics in that:

- Well-founded models are three-valued, while stable models are two valued.
- Each AnsProlog program has a unique well-founded model, while some AnsProlog programs have multiple stable models and some do not have any.

For example, the program $\{p \leftarrow \mathbf{not} p.\}$ has no stable models while it has the unique well-founded model where p is assigned the truth value *unknown*.

The program $\{b \leftarrow \mathbf{not} a., a \leftarrow \mathbf{not} b., p \leftarrow a., p \leftarrow b.\}$ has two stable models $\{p, a\}$ and $\{p, b\}$ while its unique well-founded model assigns the truth value *unknown* to $p, a,$ and b .

- Computing the well-founded model or entailment with respect to it is more tractable than computing the entailment with respect to stable models. On the other hand the latter increases the expressive power of the language.

As will be clear from many of the applications that will be discussed in Chapters 4 and 5, the nondeterminism that can be expressed through multiple stable models plays an important role. In particular, they are important for enumerating choices that are used in planning and also in formalizing aggregation. On the other hand, the absence of stable models of certain programs, which was initially thought of as a drawback of the stable model semantics, is useful in formulating integrity constraints whose violation forces elimination of models.

1.1.3 AnsProlog* vs Default logic

The sub-class AnsProlog can be considered as a particular subclass of default logic that leads to a more efficient implementation. Recall that a default logic is a pair (W, D) , where W is a first-order theory and D is a collection of defaults of the type $\frac{\alpha:\beta_1,\dots,\beta_n}{\gamma}$, where $\alpha, \beta,$ and γ are well-founded formulas. AnsProlog can be considered as a special case of a default theory where $W = \emptyset$, γ is an atom, α is a conjunction of atoms, and β_i s are literals. Moreover, it has been shown that AnsProlog* and default logic have the same expressiveness. In summary, AnsProlog* is syntactically simpler than default logic and yet has the same expressiveness, thus making it more usable.

1.1.4 AnsProlog* vs Circumscription and classical logic

The connective ‘ \leftarrow ’ and the negation as failure operator ‘**not**’ in AnsProlog* add structure to an AnsProlog* program. The AnsProlog* rule $a \leftarrow b$. is different from the classical logic formula $b \supset a$, and the connective ‘ \leftarrow ’ divides the rule of an AnsProlog* program into two parts: the head and the body.

This structure allows us to define several syntactic and semi-syntactic notions such as: *splitting*, *stratification*, *signing*, etc. Using these notions we can define several subclasses of AnsProlog* programs, and study their properties such as: *consistency*, *coherence*, *complexity*, *expressiveness*, *filter-abducibility*, and *compi-lability to classical logic*.

The sub-classes and their specific properties have led to several building block results and realization theorems that help in developing large AnsProlog* programs in a systematic manner. For example, suppose we have a set of rules with the predicates p_1, \dots, p_n in them. Now if we add additional rules to the program such that p_1, \dots, p_n only appear in the body of the new rules, then if the overall program is consistent the addition of the new rules does not change the meaning of the original predicates p_1, \dots, p_n . Additional realization theorems deal with issues such as: When can closed world assumption (CWA) about certain predicates be explicitly stated without changing the meaning of the modified program? How to modify an AnsProlog* program which assumes CWA so that it reasons appropriately when CWA is removed for certain predicates and we have incomplete information about these predicates?

The non-classical operator \leftarrow encodes a form of directionality that makes it easier to encode causality, which can not be expressed in classical logic in a straight-forward way. AnsProlog* is more expressive than propositional and first-order logic and can express transitive closure and aggregation that are not expressible in them.

1.1.5 AnsProlog* as a knowledge representation language

There has been extensive study about the suitability of AnsProlog* as a knowledge representation language. Some of the properties that have been studied are:

- When an AnsProlog* program exhibits *restricted monotonicity*. That is, it behaves monotonically with respect to addition of literals about certain predicates. This is important when developing an AnsProlog* program where we do not want future information to change the meaning of a definition.
- When is an AnsProlog* program *language independent*? When is it *language tolerant*? When is it *sort-ignorable*; i.e., when can sorts be ignored?
- When can new knowledge be added through filtering?

In addition it has been shown that AnsProlog* provides *compact representation* in certain knowledge representation problems; i.e., an equivalent representation in a tractable language would lead to an exponential blow-up in space. Similarly, it has been shown that certain representations in AnsProlog* can not be *modularly* translated into propositional logic. On the other hand problems such as constraint satisfaction problems, dynamic constraint satisfaction problems, etc. can be modularly represented in AnsProlog*. In a similar manner to its relationship with default logic, subclasses of other nonmonotonic formalisms such as auto-epistemic logic have also been shown to be equivalent to AnsProlog*.

Finally, the popular sub-class AnsProlog has a sound approximate characterization, called the well-founded semantics, which has nice properties and which is computationally more tractable.

1.1.6 AnsProlog* implementations: Both a specification and a programming language

Since AnsProlog* is fully declarative, representation (or programming) in AnsProlog* can be considered both as a specification and a program. Thus AnsProlog* representations eliminate the ubiquitous gap between specification and programming.

There are now some efficient implementations of AnsProlog* sub-classes, and many applications are built on top of these implementations. Although there are also some implementations of other nonmonotonic logics such as default logic (DeReS at the University of Kentucky) and circumscription (at the Linköping University), these implementations are very slow and very few applications have been developed based on them.

1.1.7 Applications of AnsProlog*

The following is a list of applications of AnsProlog* to database query languages, knowledge representation, reasoning, and planning.

- AnsProlog* has a greater ability than Datalog in expressing database query features. In particular, AnsProlog* can be used to give a declarative characterization of the standard *aggregate operators*, and recently it has been used to define new aggregate operators, and even data mining operators. It can also be used for querying in the presence of different kinds of incomplete information, including *null values*.
- AnsProlog* has been used in planning and allows easy expression of different kinds of (procedural, temporal, and hierarchical) domain control knowledge, ramification and qualification constraints, conditional effects, and other advanced constructs, and can be used for approximate planning in the presence of incompleteness. Unlike propositional logic, AnsProlog* can be used for conformant planning, and there are attempts to use AnsProlog* for planning with sensing and diagnostic reasoning. It has also been used for

assimilating observation of an agent and planning from the current situation by an agent in a dynamic world.

- AnsProlog* has been used in product configuration, representing constraint satisfaction problems (CSPs) and dynamic constraint satisfaction problems (DCSPs).
- AnsProlog* has been used for scheduling, supply chain planning, and in solving combinatorial auctions.
- AnsProlog* has been used in formalizing deadlock and reachability in Petri nets, in characterizing monitors, and in cryptography.
- AnsProlog* has been used in verification of contingency plans for shuttles, and also has been used in verifying correctness of circuits in the presence of delays.
- AnsProlog* has been used in benchmark knowledge representation problems such as reasoning about actions, plan verification, and the frame problem therein, in reasoning with inheritance hierarchies, and in reasoning with prioritized defaults. It has been used to formulate normative statements, exceptions, weak exceptions, and limited reasoning about what is known and what is not.
- AnsProlog* is most appropriate for reasoning with incomplete information. It allows various degrees of trade-off between computing efficiency and completeness when reasoning with incomplete information.

1.2 Answer set frameworks and programs

In this section we define the syntax of an AnsProlog* program (and its extensions and sub-classes), and the various notations that will be used in defining the syntax and semantics of these programs and in their analysis in the rest of the book.

An *answer set framework*³ consists of two alphabets (an axiom alphabet and a query alphabet), two languages (an axiom language, and a query language) defined over the two alphabets, a set of axioms, and an entailment relation between sets of axioms and queries. The query alphabet will be closely associated with the axiom alphabet and the query language will be fairly simple and will be discussed later in Section 1.3.5. We will now focus on the axiom language.

Definition 1 The axiom alphabet (or simply the *alphabet*) of an answer set framework consists of seven classes of symbols:

- (1) variables,
- (2) object constants (also referred to as constants),
- (3) function symbols,
- (4) predicate symbols,
- (5) connectives,
- (6) punctuation symbols, and
- (7) the special symbol \perp ;

³ In contrast logical theories usually have a single alphabet, a single language, and have inference rules to derive theorems from a given set of axioms. The theorems and axioms are both in the same language.

where the connectives and punctuation symbols are fixed to the set $\{\neg, or, \leftarrow, \mathbf{not}, ', '\}$ and $\{ '(', ')', '., '\}$ respectively; while the other classes vary from alphabet to alphabet. \square

We now present an example to illustrate the role of the above classes of symbols. Consider a world of blocks in a table. In this world, we may have object constants such as *block1*, *block2*, ... corresponding to the particular blocks and the object constant *table* referring to the table. We may have predicates *on_table*, and *on* that can be used to describe the various properties that hold in a particular instance of the world. For example, *on_table(block1)* means that *block1* is on the table. Similarly, *on(block2, block3)* may mean that *block2* is on top of *block3*. An example of a function symbol could be *on_top*, where *on_top(block3)* will refer to the block (if any) that is on top of *block3*.

Unlike the earlier prevalent view of considering logic programs as a subset of first order logic we consider answer set theories to be different from first-order theories, particularly with some different connectives. Hence, to make a clear distinction between the connectives in a first-order theory and the connectives in the axiom alphabet of an answer set framework, we use different symbols than normally used in first-order theories: *or* instead of \vee , and *'* instead of \wedge .

We use some informal notational conventions. In general, variables are arbitrary strings of English letters and numbers that start with an upper-case letter, while constants, predicate symbols and function symbols are strings that start with a lower-case letter. Sometimes – when dealing with abstractions – we use the additional convention of using letters *p*, *q*, ... for predicate symbols, *X*, *Y*, *Z*, ... for variables, *f*, *g*, *h*, ... for function symbols, and *a*, *b*, *c*, ... for constants.

Definition 2 A *term* is inductively defined as follows:

- (1) A variable is a term.
- (2) A constant is a term.
- (3) If *f* is an *n*-ary function symbol and t_1, \dots, t_n are terms then $f(t_1, \dots, t_n)$ is a term. \square

Definition 3 A term is said to be *ground*, if no variable occurs in it. \square

Definition 4 Herbrand Universe and Herbrand Base

- The Herbrand Universe of a language \mathcal{L} , denoted by $HU_{\mathcal{L}}$, is the set of all ground terms which can be formed with the functions and constants in \mathcal{L} .
- An *atom* is of the form $p(t_1, \dots, t_n)$, where *p* is a predicate symbol and each t_i is a term. If each of the t_i s is ground then the atom is said to be ground.
- The Herbrand Base of a language \mathcal{L} , denoted by $HB_{\mathcal{L}}$, is the set of all ground atoms that can be formed with predicates from \mathcal{L} and terms from $HU_{\mathcal{L}}$.

- A *literal* is either an atom or an atom preceded by the symbol \neg . The former is referred to as a positive literal, while the latter is referred to as a negative literal.

A literal is referred to as ground if the atom in it is ground.

- A *naf-literal* is either an atom or an atom preceded by the symbol **not**.

The former is referred to as a positive naf-literal, while the latter is referred to as a negative naf-literal.

- A *gen-literal* is either a literal or a literal preceded by the symbol **not**. □

Example 1 Consider an alphabet with variables X and Y , object constants a, b , function symbol f of arity 1, and predicate symbols p of arity 1. Let \mathcal{L}_1 be the language defined by this alphabet.

Then $f(X)$ and $f(f(Y))$ are examples of terms, while $f(a)$ is an example of a ground term. Both $p(f(X))$ and $p(Y)$ are examples of atoms, while $p(a)$ and $p(f(a))$ are examples of ground atoms.

The Herbrand Universe of \mathcal{L}_1 is the set $\{a, b, f(a), f(b), f(f(a)), f(f(b)), f(f(f(a))), f(f(f(b))), \dots\}$.

The Herbrand Base of \mathcal{L}_1 is the set $\{p(a), p(b), p(f(a)), p(f(b)), p(f(f(a))), p(f(f(b))), p(f(f(f(a))))\}$. □

Definition 5 A rule is of the form:

$$L_0 \text{ or } \dots \text{ or } L_k \leftarrow L_{k+1}, \dots, L_m, \mathbf{not} L_{m+1}, \dots, \mathbf{not} L_n. \quad (1.2.1)$$

where L_i s are literals or when $k = 0$, L_0 may be the symbol \perp , and $k \geq 0$, $m \geq k$, and $n \geq m$.

A rule is said to be ground if all the literals of the rule are ground.

The parts on the left and on the right of ' \leftarrow ' are called the *head* (or *conclusion*) and the *body* (or *premise*) of the rule, respectively.

A rule with an empty body and a single disjunct in the head (i.e., $k = 0$) is called a *fact*, and then if L_0 is a ground literal we refer to it as a ground fact.

A fact can be simply written without the \leftarrow as:

$$L_0. \quad (1.2.2)$$

□

When $k = 0$, and $L_0 = \perp$, we refer to the rule as a *constraint*.

The \perp s in the heads of constraints are often eliminated and simply written as rules with empty head, as in

$$\leftarrow L_1, \dots, L_m, \mathbf{not} L_{m+1}, \dots, \mathbf{not} L_n. \quad (1.2.3)$$

Definition 6 Let r be a rule in a language \mathcal{L} . The grounding of r in \mathcal{L} , denoted by $ground(r, \mathcal{L})$, is the set of all rules obtained from r by all possible substitutions of elements of $HU_{\mathcal{L}}$ for the variables in r . \square

Example 2 Consider the rule $p(f(X)) \leftarrow p(X)$. and the language \mathcal{L}_1 from Example 1. Then $ground(r, \mathcal{L}_1)$ will consist of the following rules:

$$p(f(a)) \leftarrow p(a).$$

$$p(f(b)) \leftarrow p(b).$$

$$p(f(f(a))) \leftarrow p(f(a)).$$

$$p(f(f(b))) \leftarrow p(f(b)).$$

\vdots

\square

Definition 7 The *answer set language* given by an alphabet consists of the set of all ground rules constructed from the symbols of the alphabet. \square

It is easy to see that the language given by an alphabet is uniquely determined by its constants O , function symbols F , and predicate symbols P . This triple $\sigma = (O, F, P)$ is referred to as the *signature* of the answer set framework and often we describe a language by just giving its signature.

1.2.1 AnsProlog* programs

An AnsProlog* program is a finite set of rules of the form (1.2.1), and is used to succinctly express a set of axioms of an answer set framework. ‘AnsProlog’ is a short form for *Answer set programming in logic*, and the ‘*’ denotes that we do not place any restrictions on the rules.

With each AnsProlog* program Π , when its language is not otherwise specified, we associate the language $\mathcal{L}(\Pi)$ that is defined by the predicates, functions, and constants occurring in Π . If no constant occurs in Π , we add some constants to $\mathcal{L}(\Pi)$ for technical reasons. Unless stated otherwise, we use the simplified notation HU_{Π} and HB_{Π} instead of $HU_{\mathcal{L}(\Pi)}$ and $HB_{\mathcal{L}(\Pi)}$, respectively. When the context is clear we may just use HU and HB , without the subscripts.

Example 3 Consider the following AnsProlog* program Π :

$$p(a).$$

$$p(b).$$

$$p(c).$$

$$p(f(X)) \leftarrow p(X).$$

Then $\mathcal{L}(\Pi)$ is the language defined by the predicate p , function f , and constants a , b , and c .

HU_{Π} is the set $\{a, b, c, f(a), f(b), f(c), f(f(a)), f(f(b)), f(f(c)), f(f(f(a))), f(f(f(b))), f(f(f(c))), \dots\}$.

HB_{Π} is the set $\{p(a), p(b), p(c), p(f(a)), p(f(b)), p(f(c)), p(f(f(a))), p(f(f(b))), p(f(f(c))), p(f(f(f(a))))\}$.

□

Throughout this book we consider several distinct sub-classes of AnsProlog* programs. The important ones are:

- AnsProlog program: A set of rules where L_i s are atoms and $k = 0$. This is the most popular sub-class, and to make it easier to write and refer, it does not have a superscript.

Such programs are syntactically⁴ referred to as *general logic programs* and *normal logic programs* in the literature. The program in Example 3 is an AnsProlog program.

Example 4 Following is an example of another AnsProlog program from which we can conclude that *tweety* flies while *skippy* is abnormal and does not fly.

$fly(X) \leftarrow bird(X), \mathbf{not} ab(X).$

$ab(X) \leftarrow penguin(X).$

$bird(X) \leftarrow penguin(X).$

$bird(tweety) \leftarrow.$

$penguin(skippy) \leftarrow.$

□

- AnsProlog^{-not} program: A set of rules where L_i s are atoms, $k = 0$, and $m = n$.

Such programs are referred to as *definite programs* and *Horn logic programs* in the literature.

Example 5 The following is an example of an AnsProlog^{-not} program from which we can make conclusions about the *ancestor* relationship between the constants a , b , c , d , and e , for the particular parent relationship specified in the program:

$anc(X, Y) \leftarrow par(X, Y).$

$anc(X, Y) \leftarrow par(X, Z), anc(Z, Y).$

$par(a, b) \leftarrow.$

$par(b, c) \leftarrow.$

$par(d, e) \leftarrow.$

The first two rules of the above program can be used to define the ancestor relationship over an arbitrary set of *parent* atoms. This is an example of ‘transitive closure’ and in general it cannot be specified using first-order logic. □

⁴ AnsProlog programs also denote a particular semantics, while several different semantics may be associated with general logic programs.

- AnsProlog[¬] program: A set of rules where $k = 0$.

Such programs are syntactically referred to as *extended logic programs* in the literature.

Example 6 The following is an example of an AnsProlog[¬] program from which we can conclude that *tweety* flies while *rocky* does not.

```
fly(X) ← bird(X), not ¬fly(X).
¬fly(X) ← penguin(X).
bird(tweety) ←.
bird(rocky) ←.
penguin(rocky) ←.
```

□

- AnsProlog^{OR} program: A set of rules where L_i s are atoms.

Such programs are syntactically referred to as *normal disjunctive logic programs* in the literature. Sub-classes of them where $m = n$ are syntactically referred to as *disjunctive logic programs*.

Example 7 The following is an example of an AnsProlog^{OR} program from which we can conclude that *slinky* is either a bird or a reptile but not both.

```
bird(X) or reptile(X) ← lays_egg(X).
lays_egg(slinky) ←.
```

□

- In each of the above classes if we allow constraints (i.e., rules with \perp in the head) then we have AnsProlog[⊥], AnsProlog^{¬**not**,⊥}, AnsProlog^{¬,⊥}, and AnsProlog^{OR,⊥} programs respectively.
- AnsProlog^{¬, OR,⊥} program: It is the same as an AnsProlog^{*} program.
- AnsDatalog program: An AnsProlog program, with the restriction that the underlying language does not have function symbols. The programs in Example 4 and Example 5 are also AnsDatalog programs, while the program in Example 3 is not an AnsDatalog program.
- AnsDatalog^X program, $X \in \{ \text{‘-not’}, \text{‘*’}, \text{‘¬’}, \text{‘or’}, \text{‘¬, or’}, \text{‘-not, ⊥’}, \text{‘¬, ⊥’}, \text{‘or, ⊥’}, \text{‘¬, or, ⊥’} \}$: An AnsProlog^X program, with the restriction that the underlying language does not have function symbols. The programs in Example 5, Example 6, and Example 7 are examples of AnsDatalog^{¬**not**}, AnsDatalog[¬], and AnsDatalog^{OR} programs, respectively.
- Propositional Y program, where Y is one of the above classes: A program from the class Y with the added restriction that all the predicates are of arity 0, i.e., all atoms are propositional ones. An example of a propositional AnsProlog program is the program $\{a \leftarrow \text{not } b, b \leftarrow \text{not } a.\}$.
- AnsProlog^{*}(n) program: An AnsProlog^{*}(n) program is an AnsProlog^{*} program that has at most n literals in the body of its rules. We can make similar restrictions for other sub-classes of AnsProlog^{*}.

The following table relates our terminologies to the various terminologies used in the literature.

AnsProlog terminology	Earlier terminologies
answer sets (of AnsProlog programs)	stable models
AnsProlog ^{-not}	definite programs, Horn logic programs
AnsProlog	general logic programs, normal logic programs (with stable model semantics)
AnsProlog ⁻	extended logic programs (with answer set semantics)
AnsProlog ^{or}	normal disjunctive logic programs (with stable model semantics)
AnsProlog ^{-not,or}	disjunctive logic programs
AnsDatalog ^{-not}	Datalog
AnsDatalog	Datalog ^{not} (with stable model semantics)
AnsDatalog ^{or}	Datalog ^{not,or} (with stable model semantics)

1.2.2 AnsProlog* notations

In this section we will present an almost comprehensive list of additional notations that will be used in the rest of this book. The reader should not become concerned with the large range of notations, and does not need to grasp them all at the same time. Only a small subset of them will be used together in a section or a chapter. The reason we give them here together instead of distributing them over the various chapters is that some of the notations are very similar and may create confusion if presented separately. By presenting them together we can contrast them easily.

- Given a rule r of the form (1.2.1):

$$head(r) = \{L_0, \dots, L_k\},$$

$$body(r) = \{L_{k+1}, \dots, L_m, \mathbf{not} L_{m+1}, \dots, \mathbf{not} L_n\},$$

$$pos(r) = body^+(r) = \{L_{k+1}, \dots, L_m\},$$

$$neg(r) = body^-(r) = \{L_{m+1}, \dots, L_n\},$$

$$lit(r) = head(r) \cup pos(r) \cup neg(r), \text{ and}$$

r is said to be *active* with respect to a pair $\langle X, Y \rangle$ of sets of literals, if $pos(r) \subseteq X$, and $neg(r) \cap Y = \emptyset$.

Given a set of literals S , $\mathbf{not} S$ denotes the set of naf-literals $\{\mathbf{not} l : l \in S\}$. Using this notation we write $\mathcal{A} \leftarrow \mathcal{B}^+$, $\mathbf{not} \mathcal{B}^-$ to denote the rule r where \mathcal{A} is $head(r)$, \mathcal{B}^+ is $pos(r)$, and \mathcal{B}^- is $neg(r)$.

- For any program Π , $Head(\Pi) = \bigcup_{r \in \Pi} head(r)$.
- Various notations for sets of atoms.

For a predicate p , $atoms(p)$ will denote the subset of HB_Π formed with predicate p .

For a set of predicates A , $atoms(A)$ will denote the subset of HB_Π formed with the predicates in A .

For a list of predicates p_1, \dots, p_n , $atoms(p_1, \dots, p_n)$ denotes the set of atoms formed with predicates p_1, \dots, p_n .

For a signature σ , $atoms(\sigma)$ denotes the set of atoms over σ .

Given a set of naf-literals S , $atoms(S)$ denotes the set $\{a : a \text{ is an atom, and } a \in S\} \cup \{a : a \text{ is an atom, and } \mathbf{not} \ a \in S\}$.

- Various notations for sets of literals.

For a program Π , $lit(\Pi) = \bigcup_{r \in \Pi} lit(r)$.

For a predicate p , $Lit(p)$ denotes the collection of ground literals formed by the predicate p .

For a language \mathcal{L} , $Lit(\mathcal{L})$ denotes the set of all literals in \mathcal{L} .

For a program Π , Lit_Π denotes the set of all literals in its associated language; and when the context is clear we may just use Lit .

For a list of predicates p_1, \dots, p_n , $lit(p_1, \dots, p_n)$ denotes the set of literals formed with predicates p_1, \dots, p_n .

For a signature σ , $lit(\sigma)$ denotes the set of literals over σ .

- For any logic program Π , we define

$$ground(\Pi, \mathcal{L}) = \bigcup_{r \in \Pi} ground(r, \mathcal{L})$$

and write $ground(\Pi)$ for $ground(\Pi, \mathcal{L}(\Pi))$.

Example 8 Consider the program Π from Example 3. The program $ground(\Pi)$ consists of the following rules:

$p(a) \leftarrow.$

$p(b) \leftarrow.$

$p(c) \leftarrow.$

$p(f(a)) \leftarrow p(a).$

$p(f(b)) \leftarrow p(b).$

$p(f(c)) \leftarrow p(c).$

$p(f(f(a))) \leftarrow p(f(a)).$

$p(f(f(b))) \leftarrow p(f(b)).$

$p(f(f(c))) \leftarrow p(f(c)).$

\vdots

□

- Signature $\sigma_1 = \{O_1, F_1, P_1\}$ is said to be a sub-signature of signature $\sigma_2 = \{O_2, F_2, P_2\}$ if $O_1 \subseteq O_2$, $F_1 \subseteq F_2$ and $P_1 \subseteq P_2$.

$\sigma_1 + \sigma_2$ denotes the signature $\{O_1 \cup O_2, F_1 \cup F_2, P_1 \cup P_2\}$.

The sets of all ground terms over signature σ are denoted by $terms(\sigma)$.

Consistent sets of ground literals over signature σ are called *states* of σ and denoted by $states(\sigma)$.

- For any literal l , the symbol \bar{l} denotes the literal opposite in sign to l . That is for an atom a , if $l = \neg a$ then $\bar{l} = a$, and if $l = a$ then $\bar{l} = \neg a$. Moreover, we say l and \bar{l} are *complementary* or *contrary* literals.

Similarly for a literal l , $not(l)$ denotes the gen-literal **not** l , while $not(\mathbf{not} l)$ denotes l .

- For a set of literals S , \bar{S} denotes the set $HB \setminus S$.
- For a set of literals S , $\neg S$ denotes the set $\{\bar{l} : l \in S\}$.
- For a set of literals S , $Cn(S) = Lit$ if S has complementary literals; otherwise $Cn(S) = S$.
- Two sets of literals S_1 and S_2 are said to disagree if $S_1 \cap \neg S_2 \neq \emptyset$. Otherwise we say that they agree.
- Given a set of literals L and an AnsProlog* program Π , $\Pi \cup L$ means the AnsProlog* program $\Pi \cup \{l \leftarrow . : l \in L\}$.
- Let Π be an AnsProlog program, A be a set of naf-literals, and B be a set of atoms. B is said to *agree* with A , if $\{a : a \text{ is an atom, and } a \in A\} \subseteq B$ and $\{a : a \text{ is an atom, and } \mathbf{not} a \in A\} \cap B = \emptyset$.
- A set S of literals is said to be *complete* with respect to a set of literals P if for any atom in P either the atom or its negation is in S . When $P = S$ or P is clear from the context, we may just say S is complete.
- A set X of literals is said to be *saturated* if every literal in X has its complement in X .
- A set X of literals is said to be *supported* by an AnsProlog ^{\neg, \perp} program Π , if for every literal L in X there is a rule in Π with L in its head and $L_1, \dots, L_m, \mathbf{not} L_{m+1}, \mathbf{not} L_n$ as its body such that $\{L_1, \dots, L_m\} \subseteq X$ and $\{L_{m+1}, \dots, L_n\} \cap X = \emptyset$.
- A rule is said to be *range restricted* (or *allowed*) if every variable occurring in a rule of the form 1.2.1 occurs in one of the literals L_{k+1}, \dots, L_m . In the presence of built-in comparative predicates such as *equal*, *greater than*, etc., the variables must occur in a nonbuilt-in literal among L_{k+1}, \dots, L_m . A program Π is range restricted (or allowed) if every rule in Π is range restricted.

The programs in Examples 3–7 are all range restricted. The program consisting of the following rules is not range restricted, as its first rule has the variable X in the head which does not appear in its body at all.

$$p(X) \leftarrow q.$$

$$r(a) \leftarrow.$$

The program consisting of the following rules is also not range restricted, as its first rule has the variable Y , which appears in **not** $r(X, Y)$ in the body, but does not appear in a positive naf-literal in the body.

$$p(X) \leftarrow q(X), \mathbf{not} r(X, Y).$$

$$r(a, b) \leftarrow.$$

$$q(c) \leftarrow.$$

1.3 Semantics of AnsProlog* programs

In this section we define the semantics of AnsProlog* programs. For that we first define the notion of answer sets for the various sub-classes and then define query

languages appropriate for the various sub-classes and define the entailment between programs and queries. While defining the answer sets we start with the most specific sub-class and gradually consider the more general sub-classes.

The answer sets of an AnsProlog* program Π , are defined in terms of the answer sets of the ground program $ground(\Pi)$. Hence, in the rest of the section we can assume that we are only dealing with ground programs.

1.3.1 Answer sets of AnsProlog^{-not} and AnsProlog^{-not,⊥} programs

AnsProlog^{-not} programs form the simplest class of declarative logic programs, and its semantics can be defined in several ways. We present two of them here, and refer to [Llo84, Llo87, LMR92] for other characterizations. In particular, we present a model theoretic characterization and a fixpoint characterization.

Model theoretic characterization

A *Herbrand interpretation* of an AnsProlog[⊥] program Π is any subset $I \subseteq HB_{\Pi}$ of its Herbrand base. Answer sets are defined as particular Herbrand interpretations that satisfy certain properties with respect to the program and are ‘minimal’. We say an interpretation I is *minimal* among the set $\{I_1, \dots, I_n\}$ if there does not exist a j , $1 \leq j \leq n$ such that I_j is a strict subset of I . We say an interpretation I is *least* among the set $\{I_1, \dots, I_n\}$ if for all j , $1 \leq j \leq n$ $I \subseteq I_j$.

A *Herbrand interpretation* S of Π is said to *satisfy* the AnsProlog[⊥] rule

$$L_0 \leftarrow L_1, \dots, L_m, \mathbf{not} L_{m+1}, \dots, \mathbf{not} L_n.$$

if (i) $L_0 \neq \perp$: $\{L_1, \dots, L_m\} \subseteq S$ and $\{L_{m+1}, \dots, L_n\} \cap S = \emptyset$ implies that $L_0 \in S$.

(ii) $L_0 = \perp$: $\{L_1, \dots, L_m\} \not\subseteq S$ or $\{L_{m+1}, \dots, L_n\} \cap S \neq \emptyset$.

A *Herbrand model* A of an AnsProlog[⊥] program Π is a Herbrand interpretation of Π such that it satisfies all rules in Π . We also refer to this as A is closed under Π .

Definition 8 An answer set of an AnsProlog^{-not,⊥} program Π is a Herbrand model of Π which is minimal among the Herbrand models of Π . \square

Example 9 Consider the following AnsProlog^{-not} program:

$$p \leftarrow a.$$

$$q \leftarrow b.$$

$$a \leftarrow .$$

The set $\{a, b, p, q\}$ is a model⁵ of this program as it satisfies all rules of this program. The sets $\{a, p, q\}$ and $\{a, p\}$ are also models of this program. But the set $\{a, b, p\}$ is not a model of this program as it does not satisfy the second rule.

Since $\{a, p\}$ is a model of this program, the sets $\{a, b, p, q\}$ and $\{a, p, q\}$ which are strict supersets of $\{a, p\}$ are not minimal models of this program. None of the sets $\{a\}$, $\{p\}$, and $\{\}$ are models of this program as each of them does not satisfy at least one of the rules of the program. Thus since all the strict subsets of $\{a, p\}$ are not models of this program, $\{a, p\}$ is a minimal model and answer set of the program. \square

Example 10 The program Π in Example 5 has an answer set S_1 given by the set $\{par(a, b), par(b, c), par(d, e), anc(a, b), anc(b, c), anc(a, c), anc(d, e)\}$, which is also its unique minimal Herbrand model. It is easy to see that S_1 satisfies all rules of $ground(\Pi)$. Hence, it is a model of Π . We now have to show that it is a minimal model. To show that S_1 is a minimal model, we will show that none of the strict subsets of S_1 are models of $ground(\Pi)$. Suppose we were to remove one of the *par* atoms of Π . In that case it will no longer be a model of $ground(\Pi)$. Now suppose we were to remove $anc(a, b)$ from S_1 . The resulting interpretation is not a model of $ground(\Pi)$ as it does not satisfy one of the ground instances of the first rule of Π . The same goes for $anc(b, c)$ and $anc(d, e)$. Hence, we cannot remove one of those three and still have a model. Now, if we remove $anc(a, c)$, it will no longer be a model as it will not satisfy one of the ground instances of the second rule of Π . Hence, S_1 is a minimal model and answer set of $ground(\Pi)$ and therefore of Π .

The set $S_2 = \{par(a, b), par(b, c), par(d, e), anc(a, b), anc(b, c), anc(a, c), anc(d, e), par(d, c), anc(d, c)\}$ is a Herbrand model of $ground(\Pi)$, as it satisfies all rules in $ground(\Pi)$. S_2 is not minimal among the models of Π , as S_1 , a model of Π , is a strict subset of S_2 . Hence, S_2 is also not an answer set of Π .

The set $\{par(a, b), par(b, c), par(d, e), anc(a, b), anc(b, c), anc(a, c), anc(d, e), par(d, c)\}$ is not a Herbrand model of $ground(\pi)$ as it does not satisfy the rule:

$$anc(d, c) \leftarrow par(d, c).$$

which is one of the ground instances of the first rule of Π . \square

⁵ In the rest of this section whenever it is clear from the context we may simply say ‘model’ instead of ‘Herbrand model.’

The notion of model although useful, is a relic from the semantics of first-order logic. So alternatively, answer sets can be defined without using the notion of a model in the following way:

Definition 9 An *answer set* of an $\text{AnsProlog}^{-\text{not}, \perp}$ program Π is a minimal subset (with respect to subset ordering) S of HB that is closed under $\text{ground}(\Pi)$. \square

Proposition 1 $\text{AnsProlog}^{-\text{not}}$ programs have unique answer sets. \square

The above is not true in general for $\text{AnsProlog}^{-\text{not}, \perp}$ programs. For example, the program $\{p \leftarrow \cdot, \perp \leftarrow p.\}$ does not have an answer set. We will denote the answer set of an $\text{AnsProlog}^{-\text{not}, \perp}$ program Π , if it exists, by $\mathcal{M}_0(\Pi)$. Otherwise, $\mathcal{M}_0(\Pi)$ is undefined. For an $\text{AnsProlog}^{-\text{not}}$ program Π we will denote its unique minimal Herbrand model by $MM(\Pi)$.

Proposition 2 The intersection of the Herbrand models of an $\text{AnsProlog}^{-\text{not}}$ program is its unique minimal Herbrand model. \square

Exercise 1 Consider the program consisting of the following rules.

$p \leftarrow p.$

What are the models of this program? What are its answer sets? \square

Iterated fixpoint characterization

From a computational viewpoint, a more useful characterization is an iterated fixpoint characterization. To give such a characterization let us assume Π to be a possibly infinite set of ground $\text{AnsProlog}^{-\text{not}}$ rules. Let 2^{HB_Π} denote the set of all Herbrand interpretations of Π . We define an operator $T_\Pi^0 : 2^{HB_\Pi} \rightarrow 2^{HB_\Pi}$ as follows:

$$T_\Pi^0(I) = \{L_0 \in HB_\Pi \mid \Pi \text{ contains a rule } L_0 \leftarrow L_1, \dots, L_m. \text{ such that } \{L_1, \dots, L_m\} \subseteq I \text{ holds}\}. \quad (1.3.4)$$

The above operator is referred to as the *immediate consequence operator*. Intuitively, $T_\Pi^0(I)$ is the set of atoms that can be derived from a single application of Π given the atoms in I .

We will now argue that T_Π^0 is monotone, i.e., $I \subseteq I' \Rightarrow T_\Pi^0(I) \subseteq T_\Pi^0(I')$. Suppose X is an arbitrary element of $T_\Pi^0(I)$. Then there must be a rule $X \leftarrow L_1, \dots, L_m$.

in Π such that $\{L_1, \dots, L_m\} \subseteq I$. Since $I \subseteq I'$, we have that $\{L_1, \dots, L_m\} \subseteq I'$. Hence, X must be in $T_\Pi^0(I')$. Therefore, $T_\Pi^0(I) \subseteq T_\Pi^0(I')$.

Now, let us assign the empty set to $T_\Pi^0 \uparrow 0$. Let us also define $T_\Pi^0 \uparrow (i + 1)$ to be $T_\Pi^0(T_\Pi^0 \uparrow i)$. Clearly, $T_\Pi^0 \uparrow 0 \subseteq T_\Pi^0 \uparrow 1$; and by monotonicity of T_Π^0 and transitivity of \subseteq , we have $T_\Pi^0 \uparrow i \subseteq T_\Pi^0 \uparrow (i + 1)$. In the case of a finite Herbrand base it can be easily seen that repeated application of T_Π^0 starting from the empty set will take us to a fixpoint of T_Π^0 . We will now argue that this fixpoint – let us refer to it as a – that is reached is the least fixpoint of T_Π^0 . Suppose this is not the case. Then there must be a different fixpoint b . Since b is the least fixpoint and a is only a fixpoint, $b \subseteq a$. Since $\emptyset \subseteq b$, by using the monotonicity property of T_Π^0 and by repeatedly applying T_Π^0 to both sides we will obtain $a \subseteq b$. Thus $a = b$, contradicting our assumption that b is different from a . Hence, a must be the least fixpoint of T_Π^0 .

In the case of an infinite Herbrand base, the case is similar and we refer to Appendix A. We can summarize the result from Appendix A as being that the operator T_Π^0 satisfies a property called *continuity*, and the ordering \subseteq over the elements in 2^{HB_Π} is a *complete lattice*, both of which guarantee that iterative application of T_Π^0 starting from the empty set will take us to the least fixpoint of T_Π^0 . More formally, $\text{lfp}(T_\Pi^0) = T_\Pi^0 \uparrow \omega =$ least upper bound of the set $\{T_\Pi^0 \uparrow \beta : \beta < \omega\}$, where ω is the first limit ordinal.

An AnsProlog^{-not} program Π can now be characterized by its least fixpoint. Recall that we assumed Π to be a possibly infinite set of ground rules. When this is not the case, and Π is non-ground, we characterize Π by the least fixpoint of the program $\text{ground}(\Pi)$. It can be shown that $\text{lfp}(T_\Pi^0)$ is also the unique minimal Herbrand model of Π .

Proposition 3 For any AnsProlog^{-not} program Π , $\text{lfp}(T_\Pi^0) =$ the unique minimal Herbrand model of $\Pi =$ the answer set of Π . \square

We now give two examples showing how the answer set of AnsProlog^{-not} programs can be computed by the iterated fixpoint method.

Example 11 Consider the following program Π from Example 9.

$p \leftarrow a.$

$q \leftarrow b.$

$a \leftarrow .$

By definition, $T_\Pi^0 \uparrow 0 = \emptyset$.

$T_\Pi^0 \uparrow 1 = T_\Pi^0(T_\Pi^0 \uparrow 0) = \{a\}.$

$T_\Pi^0 \uparrow 2 = T_\Pi^0(T_\Pi^0 \uparrow 1) = \{a, p\}.$

$T_\Pi^0 \uparrow 3 = T_\Pi^0(T_\Pi^0 \uparrow 2) = \{a, p\} = T_\Pi^0 \uparrow 2.$

Hence $\text{lfp}(T_\Pi^0) = \{a, p\}$, and therefore $\{a, p\}$ is the answer set of Π . \square