

The Logic of Logic Programming

Marc Denecker and David S. Warren

October 2022

Abstract

Our position is that logic programming is not programming in the Horn clause sublogic of classical logic, but programming in a *logic of (inductive) definitions*. Thus, the similarity between prototypical Prolog programs (e.g., `member`, `append`, ...) and how inductive definitions are expressed in mathematical text, is not coincidental but essential. We argue here that this provides a natural solution to the main lingering semantic questions of Logic Programming and its extensions.

1 Introduction

There is much ado about the “declarative semantics” of logic programs. But for a pure positive logic program, virtually every logic programmer around the world will accept that the Least Herbrand Model (LHM) [7] represents the state of affairs of the universe and the predicates as determined by the program¹. E.g., take the `member` program:

```
member(X, [X|_]).
member(X, [_|_]) :- member(X, _).
```

Assuming the vocabulary consists of predicate `member`/2, list functor `|`/2 and the constant symbols `[], 0, 1, 2, ...`, then the domain of the LHM is the Herbrand universe, the set of terms built from these symbols, and the LHM contains *member*(*t*, *l*) iff *t* is a member of list *l* (*l* not necessarily ending with `[]`).

Historically, e.g., in [7], logic programs are explained as Horn theories, i.e., sets of material implications. But, Horn theories are satisfied in an extremely broad class of (Herbrand and first-order) structures, most of which do not match at all with the LHM. For instance, the Herbrand structure in which all `member` atoms are true, satisfies the Horn theory while it is *full of errors*, e.g., `member(0, [1, 2, 3])`. It is proven in [7] that the Horn theory logically entails all atomic facts true in the LHM. However, the Horn theory entails *none* of the intended negative literals (e.g., `¬member(0, [1, 2, 3])`) and only few of the composite formulas true in the LHM and expected to hold in a theory of list-membership in the given universe. Arguably, the Horn theory is inadequate

¹Negation in programs is discussed in Section 5.

as an “axiomatisation” of the LHM and the way programmers interpret their program. This inadequacy emerges not only for the `member` program, but for virtually all logic programs that logic programmers write.

Logic programming is two languages in one: a declarative and a procedural one. Our goal is to investigate its declarative logic. We argue that such a logic should at least play the following roles. (1) The logic should explain what the states of affairs are that programmers associate with their program. Here in this case, it is the LHM and only the LHM. In this logic, a logic program should be an “axiomatisation” of the LHM, in the sense that mathematical logicians see this. (2) The declarative logic should be related to a range of natural language expressions that closely correspond to the formal expressions, both in syntax and semantics. (3) (and related) The declarative logic should give non-ambiguous insight in the meaning of the main connectives of logic programs (“:-” “,” and “not”). (4) The declarative logic should explain not only the meaning of finalized logic programs but also the meaning of programs in development, components of programs. It should explain how to interpret the program while the programmer is writing it, and not only when it is finished. E.g., when applying logic programming to query a family database *DB* for siblings, the programmer may write:

```
sibling(X,Y) :- child_of(X,P), child_of(Y,P), X \== Y.
```

The programmer writes this query without knowledge of *DB* and must have a precise understanding of this query on its own, independent of the *DB*. The declarative logic should explain this. Notice that the LHM semantics itself, for all its virtues, does *not* provide formal semantics for such stand-alone program components.

The goal of this paper is to present such a declarative logic and argue that it satisfies the above criteria. It is based on existing ideas [2, 3, 4]. A logic program is seen as a combination of definitions of all its predicates, along with an implicit “axiom” that constrains the set of function symbols to be the constructors of the universe. E.g., the two rules of the `member` program are read as an inductive definition of membership in mathematical text.

Definition *We define list membership by induction:*

- (base case) *x is a member of list [x|t];*
- (inductive case) *x is a member of [h|t] if x is a member of t.*

This mathematical (non-formal) definition defines the membership as the least relation that satisfies these two rules. Equivalently, it is the limit of the *induction process*: the process that starts with the empty set, and proceeds by iterated application of these rules until a fixpoint is reached. In this view, the similarity between prototypical logic programs and (non-formal) inductive definitions in mathematical text is not coincidental but essential. As a result, the second condition of the previous paragraph will be satisfied.

The paper is structured as follows. Section 2 introduces the logic underlying LP and discusses the meaning of the rule operator. Sections 3 and 4 use it to analyze full logic programs and their components. Section 5 considers negation.

2 The declarative logic of Logic Programming

As explained in the Introduction, we now design the core declarative logic $\mathcal{L}_{\mathcal{D}}$ (syntax and formal semantics) to formalize a logic program as a combination of an *axiom of the universe* and a *definition* of all its predicates.

Most mathematical and philosophical logicians will agree that, to understand an expression, one needs to know its *truth conditions*: the states of affairs in which it is *true*, and the states of affairs in which it is *false*. A logic semantics that specifies this is a *truth conditional semantics*. The common way to formalize this is through a satisfaction relation $M \models \psi$ (or a truth function ψ^M). Here, ψ is an expression, theory or program, and M a structure that is an abstraction of a state of affairs. First-order logic's (FO) satisfaction relation \models_{FO} specifies a truth conditional semantics. A semantics of this type abstracts all computational and operational aspects and formalizes the answer to the essential question of declarative meaning: *when is an expression true, when is it false?* In our case, the expressions ψ will be logic programs Π and their components: an *axiom of the universe* and component definitions. E.g., a structure M interpreting the relation `member` by a relation containing $(0, [1, 2, 3])$ *will not satisfy* the `member` definition and is not a *model* of the definition. Hence, this definition is *false* in this structure.

The semantics of FO is based on *first-order structures*, while in Logic Programming, often only *Herbrand structures* are used. In this paper we will use first-order structures. In particular, we cannot formalize the meaning of an axiom of the universe without considering also structures (abstractions of states of affairs) that *do not* satisfy it.

Definition 2.1. *A (first-order) structure M for a (first-order) vocabulary Σ consists of a non-empty universe \mathcal{U}^M , and appropriate values σ^M in \mathcal{U}^M for all symbols $\sigma \in \Sigma$ (an element of, or a relation or function of appropriate arity on \mathcal{U}^M).*

Definition 2.2. *The Herbrand universe $HU(\Sigma)$ for Σ is the set of terms over Σ . M is a Herbrand structure of Σ if \mathcal{U}^M is $HU(\Sigma)$ and constants c and functors f/n have Herbrand values, i.e., $c^M = c$ and $f^M(t_1, \dots, t_n) = f(t_1, \dots, t_n)$.*

This is the standard notion of Herbrand structure except when Σ contains no constant symbols. In that case, $HU(\Sigma)$ is empty and no Herbrand structures exist. Instead, in logic programming semantics, one will often add an arbitrary constant to Σ so that Herbrand structures always exist.

First-order structures do have one complication compared to Herbrand structures: while two different Herbrand structures represent necessarily different states of affairs, all *isomorphic* first-order structures M, N (notation $M \cong N$) represent the *same* state of affairs. The following is a natural property for any logic \mathcal{L} with satisfaction relation $\models_{\mathcal{L}}$: for any pair of isomorphic structures M, N , and for any expression (theory, program) ψ , $M \models \psi$ iff $N \models \psi$. Stated differently, M is a model of ψ iff N is a model of ψ . We will show that our semantics satisfies this constraint (Definition 2.4 and Theorem 2.1).

Definition 2.3. For a given vocabulary Σ , and two (first-order) structures M, N interpreting at least Σ , we say that M and N are isomorphic relative to Σ (notation $M \cong_{\Sigma} N$) if there exists a 1-1 mapping $b : \mathcal{U}^M \rightarrow \mathcal{U}^N$ such that $b(\sigma^M) = \sigma^N$ for each $\sigma \in \Sigma$.² Two structures M, N are isomorphic (notation $M \cong N$) if they interpret the same set of symbols Σ and $M \cong_{\Sigma} N$.

Below, we define the syntax and the satisfaction relation $\models_{\mathcal{D}}$ of the logic $\mathcal{L}_{\mathcal{D}}$ with a truth conditional semantics for logic programs Π and their components. The outcome should be that $M \models_{\mathcal{D}} \Pi$ iff M is (isomorphic to) the LHM of Π . The logic $\mathcal{L}_{\mathcal{D}}$ will turn out to be a sublogic of the logic FO(ID) [3] containing definitions with a.o., negation in the body and integrating it with FO.

The Herbrand Axiom Logic Programming uses constant and function symbols in a radically more restrictive way than in FO. In LP, they are treated as constructors of the universe of the program; in FO, the universe is arbitrary, and constants and functors may denote arbitrary objects and functions in it. Semantically, this leads LP to use Herbrand structures, while FO admits the much broader class of first order structures.

Why does Prolog consider only Herbrand structures? To the programmer, the universe of the program consists of all data structures. Compound terms (e.g., the term $[1,2,3]$) are used as data structures, i.e., containers of data, from which later data can be retrieved through unification. But this works only when functors are interpreted as *constructors*. For example, consider the first-order structure M' for the **member** Horn theory with universe $\{a\}$. All constants are interpreted by a , every n-ary function symbol is interpreted as the map of n-tuple (a, \dots, a) to a , every n-ary predicate is interpreted as $\{(a, \dots, a)\}$. Structure M' provably satisfies the Horn theory. The data structures represented by terms *all* collapse into the same object a ; all information stored in them has vanished and is completely lost. This is why constructors are important in logic programming. It is the same in functional programming. Also there, the universe of a program is built as the collection of terms formed by a set of constructors. Of course, functional programs have many defined non-constructor functions as well. In LP, the only such functions are the interpreted ones, e.g., $+$, \times .

The *Herbrand Axiom* for a set \mathcal{CF} of constant and function symbols is syntactically denoted as $\mathcal{H}(\mathcal{CF})$. It expresses that the universe is $HU(\mathcal{CF})$ and that symbols in \mathcal{CF} are its *constructors*.

Definition 2.4. A structure M satisfies $\mathcal{H}(\mathcal{CF})$, notation $M \models_{\mathcal{D}} \mathcal{H}(\mathcal{CF})$, if \mathcal{U}^M is $HU(\mathcal{CF})$ and all symbols in \mathcal{CF} have Herbrand values. Or, if M is isomorphic to such a structure.

Recall that if \mathcal{CF} contains no constants, then $HU(\mathcal{CF})$ is empty and $\mathcal{H}(\mathcal{CF})$ is inconsistent.

The Herbrand Axiom can be expressed as a combination of the unique name axiom and the domain closure axiom for \mathcal{CF} . The domain closure axiom cannot be expressed in FO but requires second-order logic or inductive definitions. More discussion is beyond the scope of this article.

²Here b is extended to functions and relations on \mathcal{U}^M in the standard way.

Not every structure M satisfying $\mathcal{H}(\mathcal{CF})$ is a Herbrand structure. If M interprets a function symbol f/n not in \mathcal{CF} , M is not a Herbrand structure and f/n is not interpreted as a constructor. E.g., take again $\mathcal{CF} = \{\square, |/2, 0, 1, 2, \dots\}$. Take its unique Herbrand structure and expand it for the numerical product functor $\times/2$ to the structure M_\times by interpreting \times by the function $\times^{M_\times} : HU(\mathcal{CF})^2 \mapsto HU(\mathcal{CF})$ that maps pairs of numbers n, m to $n \times m$, the product of n and m , and maps all other pairs to \square . Although M_\times is not a Herbrand structure, it satisfies $\mathcal{H}(\mathcal{CF})$. Such structures are needed for the semantics of Constraint LP, e.g., CLP(R); they are also used in examples below.

The simplest definition logic Definitions, non-formally, are a common and precise form of human knowledge. E.g., a non-inductive definition of sibling:

Definition 1. *A person x is a sibling of y if³ x and y are different and they share a parent.*

An inductive definition is that of the reachability relation of a graph.

Definition 2. *We define the reachability relation R of graph G by induction:*

- *if $(x, y) \in G$ then $(x, y) \in R$;*
- *if $(x, y) \in R$ and $(y, z) \in G$ then $(x, z) \in R$.*

(Non-formal) definitions *define* concepts *in terms of* other concepts. We call the latter the *parameters* of the definition. E.g., *sibling* is defined in terms of *parent*, and the reachability relation R in terms of G . A definition does not constrain the parameter concepts but derives, for each possible assignment of values for the parameters, a unique value for the defined concept.

Basically a (non-formal, inductive) definition specifies how the value (or extension) of the defined set is obtained from the value of the parameters. It can be explained in two equivalent ways: non-constructively, the defined set is the least set that satisfies the rules *interpreted as material implications*; constructively, the defined set is the result of the *induction process*: starting from the empty set, the rules are iteratively applied in a bottom up fashion until a fixpoint is reached. The constructive and non-constructive methods are well-known to be equivalent.

Importantly, while the above explanations are typically used for inductive definitions, both work for non-inductive definitions as well. E.g., it may be overkill to use an induction process to construct the value of the sibling relation from a given parent relation, but it does construct the correct relation.

We now introduce the formal syntax for expressing definitions.

Definition 2.5. *A formal definition D is a non-empty set of definitional rules of the form:*

$$A \leftarrow B_1, \dots, B_n.$$

³In non-formal definitions, “if” is often used where, logically speaking, “iff” is intended.

where A, B_1, \dots, B_n are standard atomic formulas. We allow B_i also to be \mathbf{t}, \mathbf{f} and equality and disequality atoms $s = t, s \neq t$. All variables are implicitly quantified universally in front of the rule. A predicate symbol is defined by D if it occurs in the head of a rule; a parameter symbol of D is any other non-variable symbol that occurs in D . The set of defined predicates is denoted $Def(D)$, the set of parameters as $Param(D)$.

This rule-based definition construct is syntactically similar to Aczel's (propositional) definition logic in [1] and Martin-Löf (predicate) definition logic in [6]. However, following intuitionistic tradition, Martin-Löf proposed only proof theory, not formal model semantics.

Definition 2.6. *A formal definition D is inductive if its dependency graph⁴ contains a cycle.*

E.g., the formal representation of the sibling definition:

$$D_s = \{ \textit{sibling}(x, y) \leftarrow \textit{child_of}(x, z), \textit{child_of}(y, z), x \neq y \}$$

E.g., the formal representation of the non-formal definition of reachability:

$$D_R = \left\{ \begin{array}{l} R(x, y) \leftarrow G(x, y) \\ R(x, z) \leftarrow R(x, y), G(y, z) \end{array} \right\}$$

E.g., an inductive definition of the product of a list of numbers:

$$D_L = \left\{ \begin{array}{l} \textit{Listproduct}([], 1) \leftarrow \\ \textit{Listproduct}([h|t], h \times p) \leftarrow \textit{Listproduct}(t, p) \end{array} \right\}$$

A formal definition can define multiple predicates by simultaneous induction. For a non-formal example, take the inductive definition of even and odd numbers: 0 is even; if n is even then $n+1$ is odd; if n is odd, $n+1$ is even.

From now on, sets D of rules can be seen as Horn theories and as definitions, and they can be evaluated in two satisfaction relations \models_{FO} and $\models_{\mathcal{D}}$.

Below, we define the satisfaction relation $\models_{\mathcal{D}}$ for definitions D , basically by copy paste of the above non-constructive explanation of non-formal definitions.

Definition 2.7. *Let D be a definition and M a first-order structure interpreting all symbols in D . We define that M satisfies D (notation $M \models_{\mathcal{D}} D$) if $M \models_{FO} D$ and for all structures N with the same universe and values of parameters in $Param(D)$ as M , if $N \models_{FO} D$ then it holds that $p^M \subseteq p^N$ for all $p \in Def(D)$.*

A definition does not in itself incorporate an assumption about the nature of constants and functors as constructors. This is why $\models_{\mathcal{D}}$ needs to be defined in terms of first order structures. Thanks to this, it specifies the semantics of the definition D_l of *Listproduct* which uses the non-constructor function \times .

⁴The graph consisting of pairs (P, Q) of predicate symbols such that Q appears in the head and P in the body of a rule in D .

Definition 2.7 makes use of the non-constructive characterisation of definitions. Alternatively, the value p^M of defined predicates can be characterized constructively, stating that M is to be the limit of a (potentially transfinite) induction process $M_0, M_1, \dots, M_n, M_{n+1}, \dots$ that starts at the structure M_0 identical to M except that defined predicates are interpreted by the empty set. From then on, M_{n+1} is obtained from M_n by applying one or more or all applicable rules in a bottom up way (and taking the union values for limit ordinals). In the constructive interpretation, rules specify the atomic operations of the induction process. Martin-Löf called them *productions*. The induction process in case of programs with and without negation was formally described in [4].

Example 2.1. *The definition D_R of reachability (given above) contains only predicate symbols and has no Herbrand structures but D_R has infinitely many first-order models. In each model M , the value R^M is the reachability relation (a.k.a. the transitive closure) of G^M . Consider M :*

$$\mathcal{U}^M = \{a, b, c\}, G^M = \{(a, b), (b, a), (c, c)\}, R^M = \{(a, a), (b, b), (c, c), (a, b), (b, a)\}$$

To verify that $M \models_{\mathcal{D}} D_R$, we can verify that R^M minimally satisfies the Horn theory D_R in M . Alternatively, we can build the induction process of D_R in the context of M and verify that it constructs R^M . The following sequence of elements of R can be derived by iterated rule application:

$$\langle (a, b), (b, a), (a, a), (b, b), (c, c) \rangle$$

It is well-known that reachability cannot be expressed in FO.^{5 6}

Example 2.2. *A small part of the (infinite) induction process of the definition D_L of Listproduct in the context of the structure M_{\times} introduced before is:*

$$\langle ([], 1), ([2], 2), ([3, 2], 6), ([5, 3, 2], 30), \dots \rangle$$

Two important theorems follow (no proof provided). Let D be a definition over vocabulary Σ .

Theorem 2.1. *If $M \cong_{\Sigma} N$ then $M \models_{\mathcal{D}} D$ iff $N \models_{\mathcal{D}} D$.*

The following theorem specifies what should be a property of every logic of definitions: given the universe and values of the parameter undefined symbols, D uniquely determines the values of its defined symbols.

Theorem 2.2. *Every structure interpreting $\text{Param}(D)$ has a unique expansion for the defined symbols that satisfies D .*

⁵Here is a folk proof that “ R is the reachability relation of G ” is not expressible in FO. Assume it was expressible in FO, by the FO theory Ψ . Choose new constants A, B , consider the FO theory $\Psi \cup \{R(A, B)\} \cup \{\psi_n | n \in \mathbb{N}_0\}$ where $\psi_n = \neg(\exists x_1, \dots, x_{n-1} : G(A, x_1) \wedge \dots \wedge G(x_i, x_{i+1}) \wedge \dots \wedge G(x_{n-1}, B))$ expresses that there is no path of length n from A to B . This theory is unsatisfiable since it states that B is reachable from A but there are no finite paths from A to B . Therefore by the compactness theorem it has a finite subset Ω that is unsatisfiable, which is impossible since clearly, its superset $\Psi \cup \{R(A, B)\} \cup \Omega$ is satisfiable. Indeed, Ω “forbids” only a finite number of lengths of paths from A to B . QED

⁶Clark completion of a rule set sometimes agrees with its semantics in $\mathcal{L}_{\mathcal{D}}$ but not for many inductive rule sets, e.g., D_R . For example, adding $(a, c), (b, c)$ to R^M yields a model of the Clark completion in which R^M is not the reachability relation.

3 Explaining full logic programs

The position of this paper on the declarative reading of logic programs boils down to the following.

A logic program Π is a theory $\{\mathcal{H}(\mathcal{CF}), D\}$ of the logic $\mathcal{L}_{\mathcal{D}}$, consisting of a definition D that defines every predicate in Π and the Herbrand Axiom $\mathcal{H}(\mathcal{CF})$ for a set \mathcal{CF} of constants and functors.

It is well-known that each program Π has a unique least Herbrand interpretation $LHM^{\mathcal{CF}}(\Pi)$, where \mathcal{CF} contains (at least) all constant and function symbols in Π [7]. Let Σ be the set of all symbols in Π and \mathcal{CF} .

Theorem 3.1. *A structure M satisfies D and $\mathcal{H}(\mathcal{CF})$ iff $M \cong_{\Sigma} LHM^{\mathcal{CF}}(\Pi)$.⁷*

Thus, any Σ -structure M isomorphic with the LHM is also a model of Π , moreover any extension of such M with *arbitrary values* for any set of additional symbols, is also a model of Π . These are the only models of Π . The satisfaction relation $\models_{\mathcal{D}}$ implements the principle of isomorphism. Importantly, it also implements the natural principle that Π contains no information about symbols not occurring in Π .

Example 3.1. *Taking $\mathcal{CF} = \{|\!/2\}$ for the `member` program results in an inconsistent $\mathcal{H}(\mathcal{CF})$. But, for any extended \mathcal{CF} with at least one constant, this program determines the correct membership relation within the intended universe. All predicates other than `member` are unconstrained by it.*

Example 3.2. *A homework problem commonly given early in introductory Prolog courses is to define family relationships, such as `sibling`, `grandparent`, or `ancestor`, using just a binary `child_of` relation, and to test it using the student's own family. To define `sibling`, a student may submit:*

```
sibling(X,Y) :- child_of(X,P), child_of(Y,P), X \== Y.
child_of(tessa,david).
child_of(jonah,david).
```

The LHM is the unique state of affairs that the student had in mind. The universe is as described by $\mathcal{H}(\{tessa, jonah, david\})$. The definition can be interpreted non-constructively through minimal satisfaction, or constructively through the induction process which constructs the intended relations in at most 4 steps. The rules can be interpreted as productions in the induction process. The facts of `child_of` behave, not as a conjunction of true facts, but as an exhaustive enumeration involving that same set of facts. The rule for `sibling` behaves, not as a weak material implication, but as a necessary and sufficient condition to be a sibling.

⁷This theorem reassures us that the LHM semantics is correct as a model semantics for $\mathcal{L}_{\mathcal{D}}$. Importantly, the original paper [7] that introduced LHM did not mention definitions and presented the LHM as the *denotation* of the Horn theory, the set of entailed atomic formulas.

4 Explaining components of logic programs

The previous section defines a program Π as consisting of two modules only: D and $\mathcal{H}(\mathcal{CF})$. A program can be large and complex, in which case it may really only be understood by its programmer(s) in a piecemeal way. Therefore, a large program must be able to be split into a collection of natural, meaningful components which the programmer understands, and develops more or less independently of each other. The definitional view shines bright light on this.

Taking the constructive view of definitions, rules are to be viewed as productions in the induction process. We define predicates by describing how their values must be constructed from the values of parameters. The basic operations are bottom-up execution of the rules. As such, single rules are not truth functional expressions. Satisfaction of a rule in a structure is not defined. Of course a rule entails a material implication, but this captures only a small bit of how to understand it.

This leads us to the following question: what are the least components of logic programs for which it makes sense to ask the question: *when is it true, when is it false?* It is not the rule, that much is clear.

The “formulas” of the logic, the basic truth functional components of a program are its (sub)definitions, i.e., rule subsets of D that define one or more predicates. Take the family program of Example 3.2. To the programmer, the program clearly consists of three components: the (implicit) Herbrand Axiom, and the definitions of `child_of` and of `sibling`. The first definition defines `child_of` by exhaustive enumeration and contains no knowledge of `sibling`. The second defines `sibling` in terms of the parameter `child_of` and contains no knowledge of the value of this parameter.

Not every partition D_1, \dots, D_n of D yields a sensible modularization in subdefinitions of it. A subdefinition D_i that defines one or more predicates, should contain all rules of D involved in the induction process of these predicates. Thus, formally, a component definition D_i of D should contain all rules of D with the same predicates in the head. For the same reason, if two or more predicates are defined by simultaneous induction, their rules should not be spread out over multiple subdefinitions but be concentrated in one definition; otherwise, the induction process cannot be computed. Thus, there should not be cycles in the dependency relation over multiple subdefinitions. In summary, the definitional view on logic programming suggests that a program Π can be naturally split in modules $D_1, \dots, D_n, \mathcal{H}(\mathcal{CF})$, such that each predicate is defined in exactly one module D_i and there are no cycles in the dependency relation involving predicates defined in different modules. The following theorem proves the correctness of this hypothesis.

Theorem 4.1. *For every structure M interpreting all symbols of Π and \mathcal{CF} , the following statements are equivalent:*

1. $M \cong_{\Sigma_D} LHM^{\mathcal{CF}}(\Pi)$
2. M satisfies D and $\mathcal{H}(\mathcal{CF})$;

3. M satisfies D_1, \dots, D_n and $\mathcal{H}(\mathcal{CF})$.

A proof of this theorem can be found in [3]. The theorem says that Π is logically equivalent to $\mathcal{H}(\mathcal{CF})$ and the conjunction of the definitions D_1, \dots, D_n . This indeed shows that a programmer can develop such modules D_i and reason with them independently of each other.

5 Negation

The nature of negation (as failure) is probably the most troubling question in the history of Logic Programming.⁸ For 50 years now, the general conviction is that the negation `not` in Prolog cannot be classical negation \neg . Where does that idea come from? It comes from the fact that Horn theories do not entail the falsity of atoms. E.g., consider the query

```
?- member(0, [1,2,3])
no
```

The answer “no” expresses that the member Horn logic theory does not entail the truth of `member(0, [1,2,3])` but neither does it entail its falsity. As a consequence, Prolog would be unsound if in the following query, the symbol `not` is interpreted as classical negation \neg :

```
?- not member(0, [1,2,3])
yes
```

This is the first and main reason why `not` is believed to be non-classical negation.

But the definitional view sheds a completely different light on the issue. Definitions augmented with the Herbrand Axiom do indeed semantically entail the falsity of many defined atomic facts. In particular, since the LHM is the unique model (modulo isomorphism), any defined atom A that is false in the LHM is false in *every* model. Therefore, this theory semantically entails the truth of the classically negated fact $\neg A$. In particular, `member(0, [1,2,3])` is false in *every* structure that satisfies the member program (as defined in this paper), and hence, \neg `member(0, [1,2,3])` is semantically entailed by the program.

And so, in the following procedure `compress/2` which removes duplicates in a list, `not` can and should be interpreted as classical negation \neg :

```
compress([], []).
compress([X|T], [X|T1]) :- compress(T, T1), not member(X, T1).
compress([X|T], T1) :- compress(T, T1), member(X, T1).
```

⁸The LP community seems much less worried about the nature of the rule operator, while for more than hundred years, the logic science knows there are considerable troubles with material implication.

The resulting program (defining `compress` and `member`) is a stratified program.

Stratification is also a natural principle of great importance in mathematics and science: it is the principle that once a concept is well-defined, it can be used to define new concepts in terms of it. It is a key building principle of science. Mathematical logicians have studied infinite, even transfinite stratified stacks of definitions, called *iterated inductive definitions* [5, 6]. The well-known principle of definition by structural induction is an instance of this [4]. In [2, 3, 4], it was argued that the well-founded semantics implements a semantic stratification principle, and that logic programs under this semantics can be viewed as a finite description of this type of definition.

To conclude, the definitional view on logic programs sheds a different light on the nature of language constructs: negation as failure is indeed classical negation! It is the rule operator that is non-classical: much stronger than a material implication, it is a production operator in the induction process.

We conclude by noting that we have realized the four goals put forward in the Introduction.

References

- [1] Peter Aczel. An introduction to inductive definitions. In J. Barwise, editor, *Handbook of Mathematical Logic*, pages 739–782. North-Holland Publishing Company, 1977.
- [2] Marc Denecker, Maurice Bruynooghe, and Victor Marek. Logic programming revisited: Logic programs as inductive definitions. *ACM Trans. Comput. Log.*, 2(4):623–654, 2001.
- [3] Marc Denecker and Eugenia Ternovska. A logic of nonmonotone inductive definitions. *ACM Trans. Comput. Log.*, 9(2):14:1–14:52, April 2008.
- [4] Marc Denecker and Joost Vennekens. The well-founded semantics is the principle of inductive definition, revisited. In Chitta Baral, Giuseppe De Giacomo, and Thomas Eiter, editors, *KR*, pages 1–10. AAAI Press, 2014.
- [5] Solomon Feferman. Formal theories for transfinite iterations of generalised inductive definitions and some subsystems of analysis. In A. Kino, J. Myhill, and R.E. Vesley, editors, *Intuitionism and Proof theory*, pages 303–326. North Holland, 1970.
- [6] Per Martin-Löf. Hauptsatz for the intuitionistic theory of iterated inductive definitions. In J.E. Fenstad, editor, *Second Scandinavian Logic Symposium*, pages 179–216, 1971.
- [7] Maarten H. van Emden and Robert A. Kowalski. The semantics of predicate logic as a programming language. *J. ACM*, 23(4):733–742, 1976.