
LOGIC PROGRAM SYNTHESIS

YVES DEVILLE AND KUNG-KIU LAU

- ▷ This paper presents an overview and a survey of logic program synthesis. Logic program synthesis is interpreted here in a broad way; it is concerned with the following question: given a specification, how do we get a logic program satisfying the specification? Logic programming provides a uniquely nice and uniform framework for program synthesis since the specification, the synthesis process and the resulting program can all be expressed in logic.

Three main approaches to logic program synthesis by formal methods are described: constructive synthesis, deductive synthesis and inductive synthesis. Related issues such as correctness and verification as well as synthesis by informal methods are briefly presented.

Our presentation is made coherent by employing a unified framework of terminology and notation, and by using the same running example for all the approaches covered. This paper thus intends to provide an assessment of existing work and a framework for future research in logic program synthesis. ◁

1. INTRODUCTION

Program synthesis refers to the elaboration of a program in some systematic manner, starting from a specification, that is a statement describing what the program should do. A specification may have many forms. We can distinguish formal specifications from informal ones. For the latter, the synthesis process cannot be totally formalised and can only be partially automated. We thus distinguish between *synthesis by formal methods* and *synthesis by informal methods*.

Address correspondence to Unité d'Informatique, Université Catholique de Louvain, Place Ste Barbe, 2, B-1348 Louvain-la-Neuve, Belgium.

Address correspondence to Department of Computer Science, University of Manchester, Oxford Road, Manchester M13 9PL, United Kingdom.

THE JOURNAL OF LOGIC PROGRAMMING

©Elsevier Science Publishing Co., Inc., 1993
655 Avenue of the Americas, New York, NY 10010

0743-1066/93/\$3.50

Program synthesis in general has been an active area of research outside logic programming. See [BGK84, Bal85, Gol86, Bie92], for example, for a presentation of the major achievements. In the early days of logic programming, logic program synthesis was one of the first areas of active research, mainly focusing on manual program derivation. By the middle eighties, however, this work had dwindled considerably. More recently, it has become an active area of research once again, this time focusing mainly on automated or semi-automated synthesis. In this paper, we give a brief survey of the work to date in logic program synthesis.

Logic program synthesis is concerned with the following:

Given a (non-executable) specification, how to get an (executable) logic program satisfying the specification?

where the notions of “specification” and “executability” are interpreted broadly. For tackling this problem, logic programming provides a uniquely uniform framework because the specification, the synthesis process and the resulting program can all be expressed in logic.

There are three main approaches to logic program synthesis by formal methods:

- In the constructive approach,¹ a conjecture based on the specification is constructively proved, and from this proof the specified program is extracted. We call this approach *constructive synthesis*.
- A more direct approach is to deduce clauses for the specified program directly from the specification. We shall call this approach *deductive synthesis*.
- Another approach can induce a program from a partial specification of the program. The program is a generalisation of the partial specification. We shall call this approach *inductive synthesis*.

This survey covers all the above approaches. Related issues such as correctness and verification, as well as synthesis by informal methods will briefly be presented. This paper does not cover some other synthesis approaches, such as knowledge-based synthesis [Smi90], synthesis by inspection (see [BF82, BGK84]), because very little work has been done using these approaches in logic programming.

Thus this survey covers only the *program synthesis* part of the so-called *program development* area. Program development is a much larger area since it also includes other aspects such as program analysis (e.g. abstract interpretation, termination) and program transformation (e.g. fold/unfold, partial evaluation), *inter alia* (see related papers in this Special Issue for other aspects of Program Development). Indeed for the whole process of program development, we should consider the realistic scenario depicted by the spiral: (informal specification \rightarrow incomplete formal specification \rightarrow unsatisfactory program \rightarrow better specification \rightarrow more satisfactory program, and so on). This scenario shows the development process as a life-cycle, within which program synthesis enables a transition from specification to program. We will not deal with this entire life-cycle in this survey. Rather, we concentrate on synthesis methods. However, the reader should bear in mind the larger context of program development.

The distinction between deductive program synthesis and program transformation is rather subjective and often depends on the context of the research. A

¹Also known as the *proofs-as-programs* approach, after [BC85].

possible difference could be that synthesis starts from some non-executable specification, which usually means a non-recursive description of the problem, whilst transformation usually starts from an already executable description. Program transformation will be covered by a separate survey in this volume. We shall not cover other related logic programming topics such as programming environments, inductive logic programming, and semantics.

Our aim for this survey is to provide a broad, extensive (and hopefully complete) survey in the field of logic program synthesis. We shall give a coherent presentation using a unified framework of terminology and notation, and the same running example for all the approaches covered. By so doing, we shall attempt to expose the similarities and dissimilarities between the different approaches, and to evaluate their respective strengths and weaknesses.

Given the space limit, it is not possible to describe all the theoretical background necessary for a full treatment of the examples, or the theoretical difficulties and unsolved theoretical problems for each method. Nevertheless, a basic knowledge of logic programming will be sufficient to get a precise idea of the synthesis methods presented. However, the reader will need to read the original papers for a deeper understanding.

We hope this paper provides an assessment of existing work and a framework for future research in logic program synthesis.

The paper is organised as follows. In Section 2 we define and explain the basic concepts. In Sections 3, 4, 5, and 6, we present an account of constructive synthesis, deductive synthesis, inductive synthesis and synthesis by informal methods respectively. In each of these sections, we describe the approach, show the running example for this approach, and give a brief overview of existing methods using the approach. Finally, in Section 7, we conclude with an assessment of existing work and proffer our views on the way ahead.

2. BASIC CONCEPTS

In this section, we first define and describe basic concepts such as specifications and programs, and then we introduce commonly used notions of program correctness, in the framework of logic programming. Finally, we give an overview of correctness criteria and verification methods that have been proposed.

2.1. Specifications and Programs

A specification can be formal or informal; a program can be either a pure logic program or a program in an existing logic programming language such as Prolog. An informal specification describes what will be called the intended relation. This is the relation the programmer/specifier has in mind when synthesising a program.

Definition 1. The *intended relation* for a predicate r (of arity n), denoted by $\mathcal{I}(r)$, is a set (of n -tuples) of ground terms.

Example 2. A pair $\langle l_1, l_2 \rangle$ belongs to $\mathcal{I}(\textit{included})$ iff l_1, l_2 are ground lists, and all the elements of l_1 belong to l_2 .

The intended relation for a predicate r can also be seen as a specific (Herbrand) interpretation for the predicate r .

In logic programming, formal specifications are usually expressed in some logic. Such a specification will be called a logic specification.

Definition 3. A *logic specification* of a predicate r , denoted $Spec(r)$ is a set of logical formulas involving r .

In the above definition, the form of the logic formulas defining a logic specification is deliberately vague, to allow specifications by examples and incomplete specifications. The logic formulas of a logic specification $Spec(r)$ necessarily involve predicate r , but can also contain other predicates defined within the specification, or defined elsewhere (primitives or other specifications).

In Example 4, the predicate *member* is assumed to be a primitive such that $member(H, L)$ holds iff H is a member of list L . In practice, *member* can be specified in another specification, or together with the specification of *included*. The second logic specification in Example 4 is a “specification by examples”.

Example 4.

$$Spec_1(included) = \{ \textit{included}(L_1, L_2) \iff \forall X(\textit{member}(X, L_1) \Rightarrow \textit{member}(X, L_2)) \}$$

$$Spec_2(included) = \{ \textit{included}([], [2, 1]), \\ \textit{included}([1, 2], [1, 3, 2, 1]), \\ \neg \textit{included}([2, 1], []) \}$$

A (pure) logic program can be seen as a particular case of a logic specification, where the language is restricted to definite Horn clauses (also called program clauses). This is called a program because it also has a procedural or operational semantics (SLD-resolution). Sometimes, normal program clauses are used [Llo87], allowing negations in the bodies of the clauses.

Definition 5. A *logic program* for a predicate r , denoted by $Prog(r)$, is a set of program clauses.

A logic program for a predicate r will normally contain clauses with the predicate r in the head. In Example 6 it is assumed that the predicate *member* and *remove_all* are primitives (with $remove_all(H, L, NL)$ holding iff NL is the list L without all the occurrences of H). This restriction amounts to assuming that the subproblems involved have been, or will be correctly implemented. They can thus be seen as primitives for $Prog(r)$. This simplification for programs (and for logic specifications) is made for ease of understanding, and can be overcome by simultaneously considering $Prog(r)$ (or $Spec(r)$) and its subproblems [Dev90]. In the context of synthesis, such predicates are often included in a *background theory*.

Example 6.

$$Prog_1(included) = \{ \textit{included}([], L_2) \leftarrow \\ \textit{included}([H|T], L_2) \leftarrow \textit{member}(H, L_2), \\ \textit{included}(T, L_2) \}$$

$$\begin{aligned}
Prog_2(included) = \{ & \text{included}([], []) \leftarrow \\
& \text{included}(L_1, [H|T]) \leftarrow \text{remove_all}(H, L_1, NL_1), \\
& \qquad \qquad \qquad \text{remove_all}(H, T, NT_2), \\
& \qquad \qquad \qquad \text{included}(NL_1, NT_2) \} \\
Prog_3(included) = \{ & \text{included}(L_1, L_2) \leftarrow \neg q(L_1, L_2) \\
& q(L_1, L_2) \leftarrow \text{member}(X, L_1), \\
& \qquad \qquad \qquad \neg \text{member}(X, L_2) \}
\end{aligned}$$

Finally, an (executable) Prolog program is a program written in the Prolog language. The difference between Prolog programs and (pure) logic programs is that most Prolog programs contain control information and/or side-effect procedure calls that destroy the correspondence between the declarative semantics and the procedural one.

Most of the existing synthesis methods focus on the declarative aspects of logic programming, and thus usually produce logic programs rather than (executable) Prolog programs. This is not a weakness of the methods, but rather a deliberate separation of concerns. The transition from a correct logic program to an executable correct Prolog program usually amounts to the introduction of suitable control information [Dev90]. We will not cover this issue in this survey.

2.2. Correctness

Correctness (and verification) is complementary to program synthesis. We have to verify that the synthesised program fulfills its specification, or more generally, that the synthesis method is correct (or sound), that is always producing correct programs.

Correctness criteria relate the intended relation, the logic specification and the logic program, one to another. A logic specification and a logic program denote some relation, according to some semantics. We here introduce correctness criteria which are parametric with respect to the chosen underlying semantics.

Definition 7. Let $Spec(r)$ be a logic specification. The *meaning* of $Spec(r)$ is the set of ground terms²

$$\mathcal{S}(r) = \{ t \mid Spec(r) \models_s r(t) \}$$

where \models_s denotes the chosen semantics for logic specifications.

Definition 8. Let $Prog(r)$ be a logic program. The *meaning* of $Prog(r)$ is the set of ground terms

$$\mathcal{P}(r) = \{ t \mid Prog(r) \models_p r(t) \}$$

where \models_p denotes the chosen semantics for logic programs.

The semantics defines how a logic specification (program) should be interpreted. The meaning of a logic specification (program) defines how the predicate r should be interpreted, according to the specification (program) and the semantics.

²We assume an underlying Herbrand universe built out of the specification language.

The chosen semantics for logic specifications and logic programs also involve interpreting (primitive) predicates not defined in the specification or in the program – for instance a possible approach is to add suitable formulas defining these predicates. An example of possible semantics for a logic specification could be its classical logical consequences. For a logic program, the semantics could be defined by its least Herbrand model, or by the set of models of the completed programs $Comp(P)$. We refer the reader to the survey on semantics for a detailed presentation of semantic issues.

2.2.1. Logic Programs versus Logic Specifications We can define three correctness criteria for logic programs with respect to logic specifications.

Definition 9. A logic program $Prog(r)$ is *partially correct* wrt a logic specification $Spec(r)$ iff $\mathcal{P}(r) \subseteq \mathcal{S}(r)$.

Partial correctness requires that the meaning of the program is included in the meaning of the specification. In other words, every answer computed by the program belongs to the specification.

Definition 10. A logic program $Prog(r)$ is *complete* wrt a logic specification $Spec(r)$ iff $\mathcal{P}(r) \supseteq \mathcal{S}(r)$.

Completeness is the converse of partial correctness. It requires that the meaning of the specification is included in the meaning of the program. It ensures that every specified answer is computed by the program.

Definition 11. A logic program $Prog(r)$ is *totally correct* wrt a logic specification $Spec(r)$ iff $\mathcal{P}(r) = \mathcal{S}(r)$.

Total correctness is thus the combination of partial correctness and completeness. Note that total correctness does not necessarily imply, from a procedural point of view, that all SLD-derivations for a given query are finite. By the completeness of SLD-resolution, each correct answer corresponds with at least one finite SLD-derivation.

2.2.2. Logic Specifications versus Intended Relations The usual correctness criterion for logic specifications with respect to intended relations is the following:

Definition 12. A logic specification $Spec(r)$ is *totally correct* wrt an intended relation $\mathcal{I}(r)$ iff $\mathcal{S}(r) = \mathcal{I}(r)$.

Total correctness here can be decomposed into partial correctness (also called *consistency* here) ($\mathcal{S}(r) \subseteq \mathcal{I}(r)$) and completeness ($\mathcal{S}(r) \supseteq \mathcal{I}(r)$). In program synthesis one usually assumes the consistency of the specification, but not its completeness. Some inductive synthesis methods explicitly consider incomplete specifications (e.g. specification by examples). For instance, $Spec_2(included)$ is consistent with $\mathcal{I}(included)$, but not complete.

2.2.3. Logic Programs versus Intended Relations Similarly, the usual correctness criterion for logic programs with respect to intended relations is the following:

Definition 13. A logic program $Prog(r)$ is *totally correct* wrt an intended relation $\mathcal{I}(r)$ iff $\mathcal{P}(r) = \mathcal{I}(r)$.

Total correctness here can also be decomposed into partial correctness ($\mathcal{P}(r) \subseteq \mathcal{I}(r)$) and completeness ($\mathcal{P}(r) \supseteq \mathcal{I}(r)$). When the intended relation is informal, such a criterion cannot be established either formally or automatically. However, since the intended relation is a mathematical one, the correctness criterion can be established by mathematically rigorous, although not formal, methods. Such correctness criteria can thus be used to guide the synthesis process, especially when incomplete logic specifications are considered.

Similarly, one could also define equivalence criteria between logic specifications, or between logic programs. Equivalence could then be decomposed into specialisation and generalisation.

When negations are allowed in the logic programs and in the queries, the correctness criteria (and the meaning of programs) are usually strengthened. For instance, besides having $\mathcal{P}(r) = \mathcal{S}(r)$, one may also require some equivalence between the negative atoms in the meaning of the program, and the negative atoms in the meaning of the specification.

Whereas verification of a program proves its correctness, *testing* a program can show its incorrectness. Testing a program can detect differences (if any) between the executable program and the specification. Tools for debugging can provide some help to locate where the problems are. Declarative debugging aims at detecting inconsistencies between the logic program (i.e. its declarative meaning) and the specification. See the survey on programming environments for a development of this subject.

2.3. Overview of Correctness Literature

Most of the literature on correctness could be described in terms of the above framework. Depending on the specification language, the form of the logic programs (with or without negations), and the chosen semantics, various approaches can be taken. For definite programs, the underlying semantics is usually the least Herbrand model (or something equivalent). We will not give a precise account of each piece of existing work. The reader is referred to the cited references for a detailed presentation.

A first formulation of correctness criteria appeared in [CT77]. This has been systematised and extended in [Cla79] (see also [Cla81, CS77]). The programs considered here are definite programs. Partial correctness and completeness include the idea of preconditions on the input. A termination criterion is also proposed. Various verification methods are described: consequence verification, computational induction, and structural induction.

Hogger's work [Hog78, Hog81, Hog84] is also based on definite programs. He proposes two families of criteria: the first one relating the declarative semantics of the logic programs to the logic specification, and the second one relating the procedural semantics of the logic program to the logic specification. He also describes

verification methods based on transformation and derivation rules.

A similar approach is proposed in [BC89, BCD90] which discusses verification methods and program equivalence. Program equivalence is also treated in [Mah88, Lev91a, Lev91b]. Notice that Lever considers programs with negations. In [GLM92], equivalence captures some observational behaviour (i.e. computation) of programs. Generalisation of programs, called extensions, are proposed in [PS90] in a framework for program development.

In [Dev90], correctness criteria relate intended relations to logic programs with possible negations. A first set of criteria relates the intended relation to the declarative semantics of the (completion of the) logic program. A second set of criteria relates the intended relation to the sequence of computed answer substitutions of the Prolog programs (derived from the logic program). These criteria are used within a methodology for constructing programs, rather than for verifying programs. The methodology also addresses the problem of types in correctness criteria. Type problems in correctness criteria and in verification are also handled in [Nai87, Nai93].

In [CD88, DF88, DF89], correctness criteria and verification methods based on proof trees are proposed for definite programs. Programs with negations are also treated in [FD92a].

Other verification methods are proposed in [Kan86, KF86, KS86, KK88]. However, the specification language is here restricted.

In [Bal78], a verification technique based on a Hoare-like inference rule is proposed for verifying the partial correctness and the completeness of a definite program. A similar approach is described in [CM91]. Another verification method based on annotations and procedural semantics, is presented in [DM87].

Correctness criteria have also been proposed in the framework of Inductive Logic Programming. (ILP). In this framework, the logic specification is a set of positive examples (denoted by $Spec^+(r)$), and possibly a set of negative examples (denoted by $Spec^-(r)$). Correctness criteria includes coverage of the positive examples by the logic program, what is also called consistency or completeness ($\mathcal{P}(r) \supseteq Spec^+(r)$), and consistency of the logic program with respect to the negative examples ($\overline{\mathcal{P}(r)} \subseteq Spec^-(r)$), where $\overline{\mathcal{P}(r)}$ is the complement of $\mathcal{P}(r)$. See [MR93] for further references.

For defining the meaning of logic programs with negations, semantics other than the usual completed program has been proposed, but not especially in the context of correctness: the perfect model semantics [Prz88], the stable model semantics [GL88], the iterated least fixpoint model [ABW88], the well-founded semantics [VGRS88], and others.

3. CONSTRUCTIVE SYNTHESIS

Constructive synthesis is an approach that originated in the functional programming paradigm,³ and is also known as the *proofs-as-programs* approach [BC85]. It has been the basis of various existing program synthesis systems [MW80, *Cet al.* 86, Hay86]. Over the past few years, constructive synthesis has also become an active area of research within logic programming [BSW90, Fri90].

³We use this term in a wide sense, to encompass any related work such as Kleene's work on realisable predicates and Martin-Löf's type theory for instance.

In functional programming, constructive synthesis is based on the Curry-Howard isomorphism in constructive type theory [How80], which states that there is a one-to-one relationship between a constructive proof of an existence theorem and a program (i.e. a function) that computes witnesses of the existentially quantified variables of the theorem. That is, from a (constructive) proof of a formula of the form

$$\forall i. \exists o. r(i, o) \quad (1)$$

one can extract a program such that for all inputs i , it computes an output o that satisfies the specified relation r . Thus the constructive synthesis process consists of two steps:

1. construct the formula (1) and prove it in a constructive logic;
2. extract from the proof a program for computing r .

It is worth emphasising that the type theory is usually a higher-order (typed) logic, and the extracted program is a function.

3.1. Description of Approach

The original constructive synthesis approach based on the functional programming paradigm may be adapted for logic program synthesis, as has been done by Bundy and Wiggins [BSW90, WBKH92] for instance. A different constructive synthesis formulation for logic program synthesis, used by Fribourg [Fri90], is based more directly on logic programming. However, to acknowledge its origin in functional programming, we shall present constructive logic program synthesis along the lines of the work of Bundy and Wiggins.

The key idea of Bundy and Wiggins' adaptation is that a predicate (in a typed logic) can be regarded as a truth-valued function, i.e. a predicate $p(X_1 : t_1, \dots, X_n : t_n)$ can be regarded as a function of type $t_1 \times \dots \times t_n \rightarrow \text{bool}$. This enables them to use the constructive synthesis approach to synthesise predicates as functions. To extract a (first-order) typed logic program from such a synthesis proof, they use a proof system based on a first-order (typed) logic with a set of rules specially devised⁴ for constructing logic program fragments from the proof rules (see [Wig92b] for their definition and proof of their correctness). It is worth emphasising that Bundy and Wiggins extract (first-order) typed, *pure* logic programs.

3.1.1. The Starting Point: In the functional paradigm, constructive synthesis of a program to compute a relation r starts from a theorem of the form

$$\vdash \forall X_1 : t_1, \dots, X_n : t_n. \exists Y_1 : t'_1, \dots, Y_m : t'_m. r(X_1, \dots, X_n, Y_1, \dots, Y_m) \quad (2)$$

where X_1, \dots, X_n are *input variables* of types t_1, \dots, t_n and Y_1, \dots, Y_m are *output variables* of types t'_1, \dots, t'_m respectively. This specification thus defines a relation between the input and output variables of the program.

This theorem is usually referred to as the *specification theorem*. However, it is more accurate to call it the *synthesis conjecture*, since it is the start of the *synthesis* process and it has yet to be proved.

⁴After the Curry-Howard isomorphism.

To adapt this synthesis conjecture for (typed) logic program synthesis, we encounter two problems due to the differences between functional and logic programs:

- Unlike a functional program, a logic program can be used in more than one way, i.e. in different input-output modes.
- Even for a chosen input-output mode, a logic program may produce many outputs or none at all.

A solution to these problems is to consider only the *all-ground* mode [BSW90]. The relation r can then be seen as a boolean-valued function. Thus the synthesis conjecture (2) becomes:

$$\vdash \forall X_1 : t_1, \dots, X_n : t_n . \exists B : boole . r(X_1, \dots, X_n) \hookrightarrow B \quad (3)$$

where we have no output variables (as in (2)) but only input variables X_1, \dots, X_n (with types t_1, \dots, t_n respectively); $boole = \{true, false\}$.

The meaning of the operator \hookrightarrow is defined by

$$\vdash Formula \hookrightarrow B \text{ iff } \left\{ \begin{array}{l} \vdash Formula \iff (B =_{boole} true) \\ \vdash \neg Formula \iff (B =_{boole} false) \end{array} \right\}$$

Thus the function to be synthesised from (3) will be a logic program that is a decision procedure for the predicate $r(X_1, \dots, X_n)$.

Note that in this *all-ground mode* approach, after a logic program has been synthesised, it is necessary to verify that the program can be used in some specific mode. This is usual and it amounts to separating the logic part from the procedural one. Such a verification of modes can be performed by using existing tools such as abstract interpretation.

3.1.2. The End Result: Thus a successful proof of the synthesis conjecture means that a logic program exists which can answer a *goal* of the form $r(X_1, \dots, X_n)$. The (pre-defined) construction rules corresponding to the proof rules used in the steps of the proof allow us to extract such a logic program. The end result of constructive logic program synthesis is usually a typed, first-order, pure logic program.

3.1.3. The Synthesis Process: Starting from a synthesis conjecture, the first step of constructive synthesis is to produce a proof of this conjecture, for instance in a typed (first-order) constructive logic. Obviously this is not a simple task in itself. It is usually carried out on a mechanised proof system (which embodies the typed constructive logic), and requires the use of very sophisticated (mechanised) development tools and proof assistants. For instance, Bundy and Wiggins [BSW90, WBKH92, Wig92b] use a proof planner to guide and automate parts of their proofs carried out in a proof development system.

From the proof of the synthesis conjecture, a program can be extracted. One possible approach consists of regarding the proof itself as a program. This interpretive approach is possible if we have an operational semantics for such proofs. Usually though, a logic program is mechanically extracted from the proof. This is possible because each proof rule used in the proof system has an associated construction rule which has been pre-defined and proved to be correct, allowing each step of the proof to generate the corresponding logic clause(s).

Programs extracted from the constructive proofs are totally correct, assuming that the proof system is sound, and that the construction rules associated with the proof system are correct.

3.2. Example

For a logic program for $included(L_1, L_2)$, the synthesis conjecture is

$$\vdash \forall L_1 : lists, L_2 : lists . \exists B : boole . included(L_1, L_2) \leftrightarrow B$$

or, using the logic specification $Spec_1(included)$ in Example 4 (Section 2.1),

$$\vdash \forall L_1 : lists, L_2 : lists . \exists B : boole . \forall X . (member(X, L_1) \Rightarrow member(X, L_2)) \leftrightarrow B.$$

Applying induction on L_1 gives two sub-conjectures:

- (base case)

$$\vdash \forall L_2 : lists . \exists B : boole . \forall X . (member(X, []) \Rightarrow member(X, L_2)) \leftrightarrow B \quad (4)$$

- (step case)

$$\begin{aligned} T &: lists \\ \forall L_2 : lists . \exists B : boole . \forall X . (member(X, T) \Rightarrow member(X, L_2)) &\leftrightarrow B \\ \vdash \forall L_2 : lists . \exists B : boole . \forall X . (member(X, [H|T]) \Rightarrow member(X, L_2)) &\leftrightarrow B \end{aligned} \quad (5)$$

The program fragment that can be extracted at this point of the proof is

$$\begin{aligned} included(L_1, L_2) &\leftarrow L_1 = [], \dots \\ included(L_1, L_2) &\leftarrow L_1 = [H|T], \dots \end{aligned} \quad (6)$$

Since the base case is always true, the first ‘...’ can be replaced by *true*.

By using the definition of *member*, the step case gives rise to two more conjectures:

$$\vdash \exists B : boole . \forall X . (X = H \Rightarrow member(X, L_2)) \leftrightarrow B \quad (7)$$

$$\vdash \exists B : boole . \forall X . (member(X, T) \Rightarrow member(X, L_2)) \leftrightarrow B \quad (8)$$

(7) can be proved by a further application of induction and (8) can be proved using the induction hypothesis

$$\forall L_2 : lists . \exists B : boole . \forall X . (member(X, T) \Rightarrow member(X, L_2)) \leftrightarrow B$$

in the step case.

The second ‘...’ in (6) can now be replaced by

$$member(H, L_2) \wedge included(T, L_2)$$

and so the complete program extracted from the proof is

$$\begin{aligned} included(L_1, L_2) &\leftarrow L_1 = [] \\ included(L_1, L_2) &\leftarrow L_1 = [H|T], member(H, L_2), included(T, L_2) \end{aligned}$$

The synthesised logic program is thus $Prog_3(included)$ from Example 6. It is worth noting that a proof by induction on L_2 of the synthesis conjecture would lead to $Prog_2(included)$ [LWD93].

3.3. Overview of Methods

Our description of constructive logic program synthesis is based on the work of Bundy and Wiggins. They use a proof development system called *Whelk*, which is based on a Gentzen sequent calculus and a first-order typed constructive logic. *Whelk* has been implemented in a proof development environment called *Mollusc*. The precise notation and the details of their proof system and proof planning techniques can be found in [BSW90, BvHHS90, BSH90, BvHSI90, BSvH⁺91, BvHHS91, WBKH92, Wig92b]. It is worth noting that they synthesise programs either in Prolog, or the new logic programming language Gödel [HL91]. An analysis of modes for the synthesised programs is made in [Wig92a].

In contrast, Fribourg [Fri90, Fri93] uses a constructive approach based directly on logic programming. His method starts from a set P of logic procedures for pre-defined predicates and a goal G of the form⁵

$$\forall X . \exists Y . q(X, Y) \Leftarrow r(X)$$

where $q(X, Y)$ and $r(X)$ are conjunctions of atoms defined in P . The variables in X and Y are regarded as *input* and *output* variables respectively. A proof of this goal will define a procedure (for a new predicate) for computing Y from X . Fribourg performs the proof using the *extended execution system* of Kanamori and Seki [KS86] on the pre-defined predicates. The extended execution system is the standard Prolog interpreter with an extended form of SLD-resolution and a restricted form of structural induction. Each inference rule applied during the proof yields a corresponding logic procedure, thus enabling a logic program to be extracted for the new predicate on completion of the proof.

Unusually, Fribourg's method synthesises programs that are guaranteed not only to be (partially) correct with respect to the specification, but also to terminate.⁶ Moreover, tail-recursive programs can also be synthesised, though termination is not guaranteed for such programs. To help automate proofs, his method makes use of simplification lemmas [Fri91].

In [Fle93], constructive synthesis techniques are used to deductively add atoms to a logic program so that some correctness criteria with respect to a set of given logic properties are satisfied.

4. DEDUCTIVE SYNTHESIS

Deductive synthesis starts from a specification, and derives or deduces a logic program according to some pre-defined deduction rules. If the specification is a set of logic sentences, then the synthesis process consists of deducing program clauses directly from the specification. In this case, we can exploit fully the uniquely uniform framework provided by logic programming for program synthesis. Verification of partial correctness of the synthesised program reduces to showing that the deduction rules used are sound with respect to the underlying specification semantics. The resulting logic program will then be a logical consequence of the logic specification. It is therefore hardly surprising that almost all existing deductive synthesis

⁵Fribourg calls such a goal an *implicative goal*.

⁶He considers a notion of *existential termination*.

methods fall into this category. In general, however, synthesis may have to be done using more sophisticated deduction strategies (possibly involving theorem proving) which will ensure that the logic program synthesised will be correct with respect to the specification.

4.1. Description of Approach

To give a general description of (first-order) deductive synthesis of logic programs, we follow a formalisation along the lines of [LO92, LO93].

4.1.1. The Starting Point: The starting point for deductive synthesis is a pair $\langle \mathcal{M}, Q \rangle$ where

1. \mathcal{M} is a set of axioms (in some first-order language), containing the logic specification $Spec(r)$. This specification is a predicate defined by means of a *definition axiom*, i.e. an *iff*-formula (in the language of \mathcal{M}) whose head is the defined predicate. Other auxiliary predicates can also be defined in this way.
2. Q is an instance (or a set of instances)⁷ of the specified relation r for which we want a program. Q thus represents a *query*.

\mathcal{M} provides a general mathematical framework in which we can specify a large class of programs (or problems). For example, \mathcal{M} may contain a theory of *lists* (complete with induction schemas for instance).

The meaning of the pair $\langle \mathcal{M}, Q \rangle$ is the set of atoms that are instances of Q (and hence of r) and that logically follow from \mathcal{M} , according to the underlying specification semantics.

4.1.2. The End Result: To deduce a logic program only for the query specified by Q , a small subset of the initial axiomatisation \mathcal{M} suffices in general. The synthesis process tries to derive (step by step) such a subset in the form of a set of definite or normal clauses (i.e. a logic program P), in such a way that *SLD* or *SLDNF* (instead of full first-order logic) can be used on these clauses to compute the answers to Q in an efficient way. In other words, the synthesis process derives (using some methods) the program P in such a way that the set of atoms that are instances of Q and that logically follow from \mathcal{M} (under the specification semantics) is equivalent to the set of atoms which logically follow from the completion of P , $Comp(P)$ (under the program semantics). The synthesised program P is thus totally correct with respect to its logic specification $Spec(r)$, for the query Q considered.

4.1.3. The Synthesis Process: Starting from the pair $\langle \mathcal{M}, Q \rangle$, a typical deductive synthesis method performs a synthesis process that can be formalised as a sequence of the form

$$\begin{aligned} \langle \mathcal{M} \cup D_0 \cup Comp(P_0), Q \rangle &\Rightarrow \langle \mathcal{M} \cup D_1 \cup Comp(P_1), Q \rangle \Rightarrow \dots \\ &\Rightarrow \langle \mathcal{M} \cup D_n \cup Comp(P_n), Q \rangle \end{aligned}$$

⁷Not necessarily ground instances.

where $D_0 \subseteq \dots \subseteq D_n$ are sets of definition axioms for defining (new) predicates, $D_0 = \{ \}$; $P_0 \subseteq \dots \subseteq P_n$ are logic programs, $P_0 = \{ \}$; such that

$$\mathcal{M} \cup D_k \cup \text{Comp}(P_k) \models_s \text{Comp}(P_{k+1}), \quad \text{for } 0 \leq k < n, \quad (9)$$

where \models_s denotes the underlying specification semantics.

Condition (9) ensures that every program P_k of the sequence is *partially correct* with respect to $\text{Spec}(r)$.

Each step of the synthesis process thus consists of adding either a definition axiom or a program clause that has been derived. The logic programs P_1, \dots, P_n are thus derived incrementally, clause by clause, so that

$$P_1 \subseteq \dots \subseteq P_n.$$

Any chosen method for deriving the clauses will guarantee partial correctness as long as it satisfies (9).

Total correctness, however, is a much more complex issue. In the formalisation in [LO92, LO93], for example, an alternative to the standard completion of P is considered. This yields a criterion for determining when the synthesised program is totally correct (and hence when to stop the synthesis process).

Finally, it is worth pointing out that, in general, this description of deductive synthesis according to [LO92, LO93] does not apply to partial deduction (partially evaluating a logic program), although at first sight it may appear to do so. The distinction, formalised in [LO94], is that partial deduction derives $\langle \text{Comp}(P_2), Q \rangle$ from $\langle \text{Comp}(P_1), Q \rangle$, where P_1 and P_2 are logic programs and Q is the chosen goal, whereas deductive synthesis derives $\langle \text{Comp}(P), Q \rangle$ from $\langle \mathcal{M}, Q \rangle$ where $\text{Comp}(P)$ (or the completion of any program in general) is only a small subsystem of \mathcal{M} (or a *specification framework* [LO94] in general).

Similarly, deductive synthesis can be distinguished from program transformation (based on unfold or fold rules, for instance).

4.2. Example

Now we show an example of a typical deductive synthesis process, where first-order logic and SLD provide the specification and program semantics respectively.

Suppose we have a theory of lists, S_{list} , and we want to synthesise a program for the query $included(L_1, L_2)$ from the logic specification $\text{Spec}_1(included)$ in Example 4 (Section 2.1). Then the starting point will be

$$\langle S_{list} \cup D_{member} \cup D_{included}, included(L_1, L_2) \rangle$$

where D_{member} is the following definition axiom for *member*

$$member(X, L) \iff L = [H|T] \wedge X = H \vee member(X, T) \quad (10)$$

and $D_{included}$ is of course just $\text{Spec}_1(included)$.

The condition for partial correctness (9) here is

$$S_{list} \cup D_{member} \cup D_{included} \cup D_k \cup \text{Comp}(P_k) \models_s \text{Comp}(P_{k+1}), \quad \text{for } 0 \leq k < n, \quad (11)$$

where \models denotes first-order consequence.

To synthesise a clause for *included*, we deduce it directly from $Spec_1(\textit{included})$

$$\textit{included}(L_1, L_2) \iff \forall X. (\textit{member}(X, L_1) \Rightarrow \textit{member}(X, L_2))$$

as follows. Since we have

$$\begin{aligned} \textit{included}([H|T], L_2) &\iff \forall X. (\textit{member}(X, [H|T]) \Rightarrow \textit{member}(X, L_2)) \\ &\iff \forall X. (\neg \textit{member}(X, [H|T]) \vee \textit{member}(X, L_2)) \end{aligned}$$

we can deduce

$$\textit{included}([H|T], L_2) \iff \forall X. ((\neg X = H \wedge \neg \textit{member}(X, T)) \vee \textit{member}(X, L_2))$$

from $D_{\textit{member}}$ (10).

Applying the distributivity law to the right-hand-side, we get

$$\begin{aligned} \textit{included}([H|T], L_2) &\iff \forall X. ((\neg X = H \vee \textit{member}(X, L_2)) \wedge \\ &\quad (\neg \textit{member}(X, T) \vee \textit{member}(X, L_2))) \\ &\iff \forall X. ((X = H \Rightarrow \textit{member}(X, L_2)) \wedge \\ &\quad (\textit{member}(X, T) \Rightarrow \textit{member}(X, L_2))) \end{aligned}$$

which gives

$$\textit{included}([H|T], L_2) \iff \textit{member}(H, L_2) \wedge \textit{included}(T, L_2) \quad (12)$$

Thus we have deduced (12) from the axioms in

$$S_{\textit{list}} \cup D_{\textit{member}} \cup D_{\textit{included}}$$

and so we can use its *if*-part, namely the clause

$$\textit{included}([H|T], L_2) \leftarrow \textit{member}(H, L_2), \textit{included}(T, L_2)$$

and put it in P_1 . That is we have performed the synthesis step

$$\begin{aligned} &\langle S_{\textit{list}} \cup D_{\textit{member}} \cup D_{\textit{included}}, \textit{included}(L_1, L_2) \rangle \\ \Rightarrow &\langle S_{\textit{list}} \cup D_{\textit{member}} \cup D_{\textit{included}} \cup \textit{Comp}(P_1), \textit{included}(L_1, L_2) \rangle \end{aligned}$$

in such a way that (9) is satisfied, that is

$$S_{\textit{list}} \cup D_{\textit{member}} \cup D_{\textit{included}} \models \textit{Comp}(P_1) \quad (13)$$

thus ensuring the partial correctness of

$$P_1 = \{ \textit{included}([H|T], L_2) \leftarrow \textit{member}(H, L_2), \textit{included}(T, L_2) \} .$$

Similarly, we could also derive the clauses

$$\begin{aligned} \textit{included}([], []) &\leftarrow \\ \textit{included}([H], [H]) &\leftarrow \\ \textit{included}([], L_2) &\leftarrow \end{aligned}$$

from

$$\langle S_{\textit{list}} \cup D_{\textit{member}} \cup D_{\textit{included}} \cup \textit{Comp}(P_1), \textit{included}(L_1, L_2) \rangle$$

and add them to P_1 to get the final program. The partial correctness of this program is ensured by condition (9).

4.3. Overview of Methods

As mentioned earlier, most existing methods use (first-order) logic sentences for \mathcal{M} , and deduce logic clauses by correct inference rules directly from these logic sentences. That is, first-order logic and SLD respectively provide the underlying specification and program semantics.

Hansson and Tärnlund [HT79, Han80], and Clark [CT77, Cla79, Cla81] axiomatise the relations that they wish to compute, as well as the data structures involved, as definition axioms in the form of *iff*-formulas in predicate logic. That is their \mathcal{M} is a set of *iff*-formulas, and their Q is a single atomic query defined (or definable) in terms of atoms already defined in \mathcal{M} . They then deduce logic programs from the axiomatisation either by logical deduction (natural deduction in the case of [HT79, Han80]) or by symbolic execution (or re-writing) of Q together with other rules for simplifying formulas such as equivalence substitutions (as in the case of [CT77, Cla79, Cla81]). These methods guarantee partial correctness, but they require proofs of total correctness.

Hogger [Hog78, Hog81, Hog84] also starts with a set of *iff*-formulas for \mathcal{M} and a single *iff* formula for $Spec(r)$. He then treats the *if*-part without the head as a goal which is to be solved by a resolution-like mechanism using a ‘logic procedure’ consisting of the *only-if* part of $Spec(r)$, as well as other relevant ‘clauses’ from \mathcal{M} which are necessary for solving this goal. He calls this *goal-directed derivation*. Alternatively, he also carries out the derivation using non-resolution inference rules. Hogger’s method also guarantees partial correctness, but it requires proofs of total correctness.

Similar methods to the above have been proposed that are based on standard techniques for logic program transformation, for example, by Kanamori and Horiuchi [KH87].

In contrast, the method of Lau and Prestwich [LP90, LP91] is designed to be mechanisable. Here $Spec(r)$ is also an *iff*-formula and \mathcal{M} is a set of *iff*-formulas. However, the deduction of the clause(s) for solving Q is automatically decomposed into sub-deductions which when completed are automatically composed into their parent deductions. This automation is possible because the user has to specify the recursion pattern in the required procedure. This method also guarantees partial correctness. Lau and Ornaghi [LO94] have proposed a method for synthesising totally correct programs using their formalisation of deduction synthesis in terms of *SLDNF* in [LO92, LO93].

For specifications expressed by restricted classes of first-order logic formulas, it is possible to synthesise totally correct programs automatically. Such methods have been proposed by Dayantis [Day87], Sato and Tamaki [ST89], who have implemented a compiler for translating a class of first-order formulas directly into logic programs, and Kawamura [Kaw91, Kaw92].

Finally, a couple of methods which may not at first sight seem to fall into this category. Starting from the work of Bundy and Wiggins (see previous section), Kraan [KBB93a, KBB93b] developed a method for program synthesis that is based on proof planning. In planning the proof that a (not yet synthesised) program meets its given specification, the program’s body is represented by a meta-variable. The proof plan is completed by instantiating this meta-variable to logical formulas deduced from the specification.

In the LOPS synthesis system [BH84], the specification is really an ‘input-output’

synthesis conjecture (as in constructive synthesis), and the specified program is derived by (various strategies for) re-writing formulas as well as using domain knowledge to generate relevant theorems. However, we may regard the domain knowledge as \mathcal{M} , and Neugebauer [Neu92, Neu93] has shown LOPS can be re-cast as a deductive synthesis method for logic programs (as well as programs in other target languages, even C!).

At this point, it is worth noting that Kreitz [Kre90, Kre93a] has studied program synthesis at a meta-level, and has shown that the constructive and deductive approaches are fully equivalent.

5. INDUCTIVE SYNTHESIS

Inductive synthesis refers to the process of formulating general rules from incomplete information, such as examples. Inductive synthesis of programs is performed by means of inductive inference, and is part of machine learning, a branch of AI. Inductive inference is related to the concept of generalisation (deductive synthesis is related to specialisation) and has received much attention in functional programming during the 1970s. It has been an active area of research in logic programming since the early 1980s.

We shall first give a more precise description of inductive synthesis and show that inductive synthesis of recursive logic programs has a specific niche within Inductive Logic Programming (ILP). As an example of inductive synthesis, we shall then briefly present the Model Inference System [Sha81, Sha83], before overviewing other existing approaches.

This section will not cover the entire Inductive Logic Programming area. Focus will be put on methods aiming at solving “programming problems” (rather than at concept learning), that is problems where some recursion has to be synthesised. See [MR93, Mug92] for a complete survey and references on ILP.

5.1. Description of approach

In the specific framework of (recursive) program synthesis from examples, it will also be assumed here that the specifier/programmer “knows” (even if only informally) the intended relation $\mathcal{I}(r)$. He is thus able to decide whether a given example belongs to the intended relation or not.

In an inductive synthesis of logic programs, the logic specification $Spec(r)$ is usually a set of positive examples (denoted by $Spec^+(r)$), and possibly a set of negative examples (denoted by $Spec^-(r)$). Examples are ground atoms. In some methods, the specification can be constructed incrementally during the synthesis process.

The assumption that the specifier “knows” the intended relation is formalised by assuming the consistency of the logic specification with respect to the intended relation. More precisely:

$$\begin{aligned} Spec^+(r) &\subseteq \mathcal{I}(r) \\ Spec^-(r) &\subseteq \overline{\mathcal{I}(r)} \end{aligned}$$

where $\overline{\mathcal{I}(r)}$ denotes the complement of $\mathcal{I}(r)$. It is clear that a specification by examples is usually intrinsically incomplete (i.e. $Spec^+(r) \neq \mathcal{I}(r)$).

The objective of inductive synthesis is to infer a logic program $Prog(r)$ that covers at least all the examples: $Prog(r)$ must be consistent with respect to $Spec^+(r)$ (i.e. $\mathcal{P}(r) \supseteq Spec^+(r)$), and with respect to $Spec^-(r)$ (i.e. $\overline{\mathcal{P}(r)} \subseteq Spec^-(r)$). Given the incompleteness of the specification, the synthesised program must also cover other unspecified examples. Partial correctness with respect to the logic specification is thus irrelevant here. The objective is to get a program that is totally correct with respect to the intended relation, although such an objective cannot always be achieved in a fully automatic way. Inductive synthesis thus aims at inferring some “natural” extension of the given examples.

Within the methods for inductive program synthesis, one can distinguish between the trace-based approach from the model-based approach. In the trace-based approach, example traces are first generated. A trace is a sequence of instructions executed by an unknown program on some given input data. Then the traces are generalised into a program. This program may be obtained by folding, matching, and generalising the traces. Generalisation is required since traces are related to some specific inputs; folding is required in order to form loops and recursion. In the model-based approach, synthesis aims at constructing a finite axiomatisation of a model of the examples. It thus makes an intensional representation of a relation (i.e. a program) from the given (incomplete) extensional representation (i.e. the examples).

The model-based approach to inductive synthesis of logic program is better known as Inductive Logic Programming (ILP). ILP is at the intersection of empirical (inductive) learning and logic programming [Mug91]. By empirical learning, we mean the elaboration of a concept description from incomplete definitions. However, we concentrate here on a specific class of logic programs, namely the recursive ones. In this specific case, we assume that a human specifier knows the intended relation. This underlines the algorithmic focus of inductive synthesis compared to the more general scope of ILP (which also covers concept learning).

5.2. Model Inference System

One of the first systems for synthesising logic programs from examples is Shapiro’s Model Inference System (MIS) [Sha81, Sha83]. It can also be seen as a special case of program debugging [Sha83], where the initial program is empty. MIS is model-based. It is also incremental in the sense that examples are introduced one by one. For each new example, the program induced from the previous examples is updated to correctly handle this new example. A key feature of MIS is the clause generator, which has the capacity of “enumerating” possible program clauses according to some subsumption relation computed by a refinement operator. Such an enumeration is actually performed by searching the refinement graph induced by the chosen refinement operator.

The general strategy behind a synthesis with MIS is the following. The initial program is empty. For each new example, if it is a positive example that is not covered by the program, a new clause covering this example is added to the program. If the new example is a negative example that is covered by the program, then the covering clause is removed. If the resulting program is inconsistent with respect to the previous examples, the program is modified, using the above strategy. The resulting new program is then proposed to the user. The generated programs are always consistent with respect to all the introduced examples.

Example 14. Let us sketch a possible dialogue between the specifier and MIS to synthesise the *included* relation.

(Type) <i>included(list, list)</i>	
(Mode) <i>included(+, +), determinate</i>	
(Possibly used predicates) <i>member(-, -), included(-, -)</i>	
(E1) <i>included([], [1, 2])</i>	(P1) <i>included(L₁, L₂) ←</i>
(E2) <i>¬included([1], [2])</i>	(P2) <i>included([], L₂) ←</i>
(E3) <i>included([1], [2, 1])</i>	(P3) <i>included([], L₂) ←</i> <i>included([H T₁], L₂) ← member(H, L₂)</i>
(E4) <i>¬included([1, 2], [1, 3])</i>	(P4) <i>included([], L₂) ←</i> <i>included([H T₁], L₂) ← member(H, L₂),</i> <i>included(T₁, L₂)</i>

The specifier must first declare the predicate to be synthesised, its type, mode, as well as the possible predicates used by the program. The declaration of the possibly used predicates is necessary for the system to limit the size of the refinement graph. This will only contain clauses involving the *included* or *member* predicates. After the presentation of example (E1), the synthesised program is (P1), the most general clause for *included*. Example (E2) forces the system to review this choice, and to take something less general. With example (E3), the program must be generalised to cover the new example. A new clause is chosen by the clause generator. It is as general as possible, while yielding a program consistent with the previous examples. The presentation of example (E4) forces the system to reconsider the second program clause, and the clause generator produces a less general one.

5.3. Overview of methods

Fundamental notions for inductive synthesis are subsumption and generalisation, as developed by [Plo70, Plo71, Rey70]. Plotkin's idea of least general generalisation has been the basis of most model-based approaches to the induction of logic programs. Generalisation can be used in two different ways, bottom-up or top-down. In a bottom-up approach, the example *included([], [1, 2])* would yield the bottom element (least general) among the generalisations of the examples, that is the program clause *included([], [1, 2]) ←*. In a top-down approach, such as in MIS, the clauses are enumerated from the most general to the most specific. The example *included([], [1, 2])* would yield the top elements (most general) among the generalisation of the example, that is the program clause *included(L₁, L₂) ←*.

Top-down approaches as well as extensions and improvements of MIS have received much attention in the ILP framework. A complete account of this work can be found in [MR93].

Among the possible improvements of the MIS method, we mention the definition of more sophisticated refinement operators for the clause generator, the introduction of background knowledge, and predicate invention for the used predicates.

The combination of MIS and program schemata allows a further organisation of

the search space [Tin90, KS91]. This approach is especially adaptable to our specific case where recursive programs have to be synthesised, since recursive programs can often be classified according to their design strategy (see Section 6.2).

The trace-based approach to program synthesis has received much attention in the context of functional programming (see the survey [Smi84]). In the logic programming context, the trace-based approach has been reformulated in [Hag90] by means of higher-order unification in a type theory with recursion. There, logic program synthesis from examples is actually also based on the constructive paradigm. A constructive proof for a concrete example of the theory is first generated, then the proof is generalised into an inductive proof from which a program can be extracted. In the context of program transformation, Compiling Control techniques (e.g. [BSK89]), is also related to the trace-based approach.

Specifications by examples can also be extended by allowing examples and properties (i.e. logic formulas). In [FD92b, FD93b, FD93a, Fle93], logic program synthesis is performed by instantiating a divide-and-conquer program schema. The specification is composed of examples and properties. The whole synthesis process combines inductive, deductive, and constructive synthesis. Different synthesis methods are used for instantiating the different place-holders of the program schema. One of the proposed methods, called Most Specific Generalisation, aims at inductively inferring a logic program from examples, but within a restrictive setting. This method can successfully be applied to synthesise parts of a divide-and-conquer schema.

6. SYNTHESIS BY INFORMAL METHODS

Some of the methods that have been studied in the literature are informal in the sense that they start from an informal description of the intended relation. The primary objective of such methods is not necessarily the full automation of the synthesis process, but rather the elaboration of practical methods for the construction of logic programs. Usually, parts of such construction processes are, or can be, automated, hence providing a computer-aided environment for the development of logic programs (see also the survey on Programming Environments).

Broadly speaking, there are two main informal approaches. The first one constructs a logic program by structural induction, starting from the intended relation. It is informal in the sense that the resulting logic program cannot be *formally* proven correct with respect to the intended relation. We will not consider here the direct construction of Prolog programs where the construction process is based on the operational semantics of Prolog. The second approach starts with a program schema and “instantiates” it to obtain a logic program. In logic programming, methods based on program schemata basically fall into the category of informal methods because logic specifications are usually absent.

6.1. Program Construction by Structural Induction

Structural induction [Bur69, Bur74] is a major technique for the construction and the proof of correctness of programs. Basically, structural induction is a proof-by-induction method, where the induction is on the structural form of some terms. The construction of a program by structural induction is a construction where the

reasoning is based on the structure of some input parameter. Such a construction implicitly contains a correctness proof by structural induction. Although initially introduced in the context of functional programming, it is also well-adapted for logic program construction. Structural induction is also used in constructive synthesis.

The construction of a logic program by structural induction can be seen as a framework allowing a precise presentation of the “natural” (manual) construction of a logic program, but based purely on declarative semantics.

To simplify notation, let us assume that we are dealing with a binary relation $r(X, Y)$. Given an intended relation $\mathcal{I}(r)$, the constructed logic program will have the following form:

$$\begin{array}{l} r(X, Y) \leftarrow C_1 \wedge F_1 \\ \vdots \\ r(X, Y) \leftarrow C_n \wedge F_n \end{array}$$

where typically, each $C_i \wedge F_i$ will deal with one of the various cases of the induction parameter, with C_i determining a particular case of the induction parameter and the corresponding F_i verifying that the intended relation holds in this case. In practice, each C_i will often be a literal and each F_i a conjunction of literals (otherwise a straightforward transformation can lead directly to the form of a logic program).

The construction process consists of the following:

1. Choice of an induction parameter (X or Y).
2. Choice of a well-founded relation⁸ over the type of the induction parameter.
3. Construction of the structural forms C_i of the induction parameter.
4. Construction of the structural cases.

In the construction process, the predicate r as well as the other predicates involved are interpreted according to their intended relations. The construction process is thus performed within some intended Herbrand interpretation \mathcal{H} .

The structural forms of the induction parameter must cover all the possible cases. More formally, the formula⁹

$$\forall X, Y . \exists(C_1 \vee \dots \vee C_n)$$

must be true in the intended interpretation \mathcal{H} . The formula F_i should satisfy the condition that

$$\forall X, Y . \exists(C_i) \Rightarrow (r(X, Y) \Leftrightarrow \exists(C_i \wedge F_i))$$

is true in the intended interpretation \mathcal{H} .

Such an F_i formula can be obtained by reduction to simpler subproblems (because of the particular form of the induction parameter) and/or by a recursive use of $r(s, t)$. It is however crucial to show that s (or t) is smaller than the induction parameter according to the chosen well-founded relation. The construction of the F_i is certainly one of the creative tasks.

⁸A relation $<$ is *well-founded* over a set E iff there is no infinite decreasing sequence $x_1 > x_2 > \dots > x_i > \dots$ of elements of E .

⁹The subformula $\exists(F)$ in the formula $\forall X, Y \dots \exists(F) \dots$ denotes the existential closure of F , except for the variables X and Y .

One can show that under the hypothesis that the construction process has been correctly applied, the (completion of the) resulting logic program is totally correct with respect to the intended relation [Dev90]. We also have that the intended interpretation \mathcal{H} is a model of $Comp(P)$, and that the interpretation of r is the same (i.e. the intended relation $\mathcal{I}(r)$) in all the Herbrand models of $Comp(P)$. The choice of an induction parameter, a well-founded relation, and the structural forms are important since different choices can lead to different, though correct, logic programs.

The role of a well-founded relation is crucial to the correctness of the resulting program. Without a well-founded relation, (incorrect) programs of the form $r(X, Y) \leftarrow r(X, Y)$ could be constructed. From a procedural point of view, the well-founded relation also ensures the termination properties of the program when the induction parameter is ground in the query and in the recursive calls.

Example 15. Let us consider the intended relation $\mathcal{I}(included)$ specified in Example 2. In the intended interpretation \mathcal{H} , *included* is interpreted as $\mathcal{I}(included)$, *member*(H, L) is true iff H is a member of the list L , and $X = Y$ is true iff Y and X are syntactically identical. The construction proceeds as follows:

1. *Choice of an induction parameter:*
We choose L_1 (arbitrarily).
2. *Choice of a well-founded relation (over lists):*
Given two lists l_1, l_2 , we define $l_1 < l_2$ iff l_1 is the tail of l_2 .
3. *Construction of the structural forms of L_1 :*
The possible structural forms of L_1 are L_1 empty and L_1 non empty. Hence the two cases:
 - $L_1 = []$
 - $L_1 = [H|T]$

These two forms covers all the possible forms because the following formula is true in the intended interpretation \mathcal{H} :

$$\forall L_1 . \exists H, T. (L_1 = [] \vee L_1 = [H|T])$$

4. *Construction of the structural cases:*
For each structural form, we have to find a necessary and sufficient condition to have *included*(L_1, L_2) true in the intended relation.
 - For $L_1 = []$, the intended relation holds whatever the list L_2 is. The structural form is thus simply *true*. One can easily verify that the following formula is true in the intended interpretation \mathcal{H} :

$$\forall L_1, L_2 . L_1 = [] \Rightarrow (included(L_1, L_2) \Leftrightarrow true)$$

- For $L_1 = [H|T]$, a necessary and sufficient condition to have all the elements of L_1 belonging to L_2 is to have H belonging to L_2 and, all the elements of T belonging to L_2 . That is,

$$member(H, L_2) \wedge included(T, L_2)$$

Notice that T is smaller than L_1 according to the well-founded relation.

One could also easily verify that the following formula is true in the intended interpretation \mathcal{H} :

$$\forall L_1, L_2 . \quad (\exists H, T. L_2 = [H|T]) \Rightarrow \\ (\text{included}(L_1, L_2) \Leftrightarrow L_2 = [H|T] \wedge \text{member}(H, L_2) \\ \wedge \text{included}(T, L_2))$$

The resulting program is then

$$\begin{aligned} \text{included}(L_1, L_2) &\leftarrow L_1 = [] \\ \text{included}(L_1, L_2) &\leftarrow L_1 = [H|T], \quad \text{member}(H, L_2), \\ &\quad \text{included}(T, L_2) \end{aligned}$$

This program can be easily transformed into $Prog_1(\text{included})$ given in Example 6. The alternative choice for the induction parameter, that is L_2 , would lead to $Prog_2(\text{included})$ in Example 6.

6.2. Schema-guided Program Construction

Programs can be classified according to their design strategies (divide-and-conquer, generate-and test, and so on). Informally, a program schema is a program template representing a whole family of particular programs all based on the same design strategy. These programs can be obtained by instantiating the place-holders in the template to particular parameters or predicates. It is therefore interesting to guide the construction of a program by a schema capturing the essence of the chosen strategy.

Example 16 presents a (simplified) version of the divide-and-conquer schema, where the “divide” (i.e. induction) is performed on the second parameter.

Example 16.

$$\begin{aligned} r(X, Y) &\leftarrow \text{Minimal}(Y), \quad \text{Solve}(X, Y) \\ r(X, Y) &\leftarrow \text{NonMinimal}(Y), \quad \text{Decompose}(Y, \text{First}Y, \text{Rest}Y), \\ &\quad r(\text{Rest}X, \text{Rest}Y), \\ &\quad \text{Process}(\text{First}Y, \text{First}X), \\ &\quad \text{Compose}(\text{First}X, \text{Rest}X, X) \end{aligned}$$

Various methods can be used (knowledge-based, schema composition, deductive/-constructive/inductive synthesis, uses of algebraic properties of the specification, etc.), and they can be combined for instantiating the different parts of the schema. For example, a possible instantiation for the *included* problem could be the following:

$$\begin{aligned} \text{included}(L_1, L_2) &\leftarrow L_2 = [], \quad L_1 = [] \\ \text{included}(L_1, L_2) &\leftarrow L_2 = [H_2|T_2], \quad \text{remove_all}(H_2, L_2, NL_2), \\ &\quad \text{included}(NL_1, NL_2), \\ &\quad H_2 = H_1, \\ &\quad \text{insert}(H_1, NL_1, L_1) \end{aligned}$$

where $remove_all(H, L, NL)$ holds iff NL is the list L without all the occurrences of H , and $insert(H, NL, L)$ holds iff L is the list NL where k occurrences (for some $k > 0$) of H have been added.

Given that NL_1 has no occurrence of H_1 , the atoms $insert$ can be replaced by $remove_all$, yielding $Prog_2(included)$ in Example 6.

6.3. Overview of Methods

Structural induction in logic program construction has already been seen in [CT77]. The construction of an axiomatic definition of a relation is performed by case analysis on the structural form of a parameter.

In Prolog textbooks, the usual guidelines for program construction are mainly based on a very procedural approach, and mostly disconnected with structural induction [CM84, Bra86, GKPVC86]. It should be noted that in [SS86] there is a clear distinction between the concept of a logic program and a Prolog program.

The above presentation of program construction by structural induction is based on [Dev90] where methods are proposed for the systematic development of logic programs. These methods cover the whole process, starting from the intended relation, constructing a logic program, and deriving an executable Prolog program.

In functional programming, program schemata are used in deductive synthesis, such as in the KIDS system [Smi88, Smi90], or in program transformation [HL78]. A formalization of a strategy deriving global search algorithms from specifications is described in [Kre93a]. Details can also be found in [Kre93b]. In the context of logic programming, schemata were mostly used for assisting the manual construction of logic programs.

In [GH89, GH90], a hierarchy of logic program schemata is proposed. These are set in a second-order logic framework, and reflect a divide-and-conquer design strategy. Divide-and-conquer schemata are also proposed in [DB89, Dev90] which incorporate generalisation techniques. The schemata are integrated in an environment for logic program development [HC92]. Various divide-and-conquer logic program schemata are carefully detailed in [Fle93]. These are used to guide inductive logic program synthesis.

Logic program schemata proposed in [O’K90], cover different classes of problems, and different design strategies.

Stepwise enhancement is proposed in [LS88, LS90, Lak89b, Lak89a, KS90] as a structured and procedural approach to Prolog program development. Program schemata, called “skeletons”, isolate the basic control flow structures. Skeletons can be extended by means of “techniques” which can be applied to include extra computations in the skeletons. Different extensions can also be combined. [BP90] discusses a similar system based on what the authors call Clichés.

7. EVALUATION AND PERSPECTIVES

As we pointed out earlier, in this survey we only intend to give a short introduction to the various synthesis methods. It would be folly to pretend that these methods on their own can tackle all the remaining problems or unsolved theoretical difficulties in logic program synthesis, let alone program synthesis in general! The synthesis of a program from a specification cannot be reduced to the choice of a method and the

application of well-defined rules to synthesise a correct program. With this caveat, we now conclude with a brief summary, assessment, and discussion of existing work and potential future trends.

It is generally recognised that to achieve the goals of program synthesis, the best formalisms to use are declarative ones, such as functional and logic programming. The functional programming community has been very actively pursuing this objective, mainly doing constructive synthesis based on constructive type theory, and inductive synthesis from examples. In contrast, logic programmers have mainly concentrated on deductive synthesis. Each of these approaches has its own strengths and weaknesses.

In constructive synthesis, though program extraction can be mechanised, producing the proofs remains a nontrivial task and needs human interaction. In deductive synthesis, program extraction is unnecessary (in logic programming at any rate), and each deduction step can be automated, but the overall deduction strategy also needs human guidance.

Constructive and deductive synthesis are usually applied starting with a complete logic specification to start with. A problem with these approaches [Bie93] is that writing a logic specification for a program is sometimes very much like writing the program itself. A precise syntax has to be devised to completely codify the desired behaviour, and one might prefer to write the program directly in this syntax rather than using automated synthesis systems. On the other hand, inductive synthesis from examples works very well. However, it can create programs automatically only if they are small (two or three or four lines of code), and the cost in execution time is exponential!

Synthesis by informal methods stresses what are the crucial creative steps within the design of a program. It also enables us to abstract programming concepts such as program schemas. As the starting specification is informal, these methods cannot be totally automated, but can yield tools supporting interactive program synthesis.

Logic programming provides a nice uniform framework for program synthesis. On the one hand, the specification, the synthesis, and the resulting program can all be expressed in logic. On the other hand, logic specifications can describe complete specifications as well as incomplete ones such as examples or properties of the relation to compute. The logic programming paradigm thus offers a chance to present both kinds of information within the same language, and treat them uniformly in a synthesis process.

Although presented separately in this paper, the different methods can be combined in various ways. Constructive and deductive synthesis do not have to start with complete specifications. It is reasonable to believe that the key to a general synthesis method lies in a combination of the strengths of the different synthesis approaches. By studying these different approaches in the framework of logic programming, we hope we have taken a first step in the right direction.

Finally, in order to suggest or predict the future trends or directions of program synthesis, it is useful to return to the general context of program synthesis. If we view computer programming as a process of constructing executable code from (fragmentary) information, then program synthesis shares with automatic programming the same objective of using a machine to do computer programming. However, to paraphrase [RW88], it would be in vain to hope that thanks to automatic syn-

thesis there will be no more programming. It is impossible to have user-oriented, general-purpose, and fully automatic programming systems. At least one of these three desirable qualities has to be sacrificed. The required input of such automatic systems needs to be carefully crafted, debugged and maintained. Thus some “programming” task will still have to be done. To quote [RW88],

“Automatic programming systems of the future will be more like vacuum cleaners than like self-cleaning ovens”.

Realistically, then, program synthesis aims at abstracting the programming process, letting the programmer concentrate on the really creative tasks involved. In this perspective, the synthesis system thus becomes a partner rather than an independent agent and we have IA (Intelligence Amplification) rather than AI (Artificial Intelligence) [Bie93]. Automatic programming will begin to have an impact on realistic programming by offering users tools for interactive synthesis and not by delivering some ultimate solution.

Acknowledgements

We would like to thank the referees for their valuable, detailed comments, and for their constructive criticisms and suggestions, which have greatly improved this paper. We also thank Pierre Flener and Geraint Wiggins for reading the draft and for their helpful remarks.

REFERENCES

- [ABW88.] K.R. Apt, H. Blair, and A. Walker. Towards a theory of declarative knowledge. In J. Minker, editor, *Foundations of Deductive Databases*. Morgan Kaufmann, 1988.
- [Bal78.] K. Balogh. On an interactive program verifier for Prolog programs. In *Colloquia Mathematica Societatis Janos Bolyai 26, Mathematical Logic in Computer Science, Salgotarjan, Hungary*, pages 111–142, 1978.
- [Bal85.] Robert Balzer. A 15 year perspective on automatic programming. *IEEE Transactions on Software Engineering*, 11(11):1257–1268, November 1985.
- [BC85.] Joseph L. Bates and Robert L. Constable. Proofs as programs. *ACM Transactions on Programming Languages and Systems*, 7(1):113–136, January 1985.
- [BC89.] A. Bossi and N. Cocco. Verifying correctness of logic programs. In *Proc. TAPSOFT '89, LNCS 352*, pages 96–110. Springer-Verlag, 1989.
- [BCD90.] A. Bossi, N. Cocco, and S. Dulli. A method for specializing logic programs. *ACM Transactions on Programming Languages and Systems*, 12(2):253–302, 1990.
- [BF82.] A. Barr and E.A. Feigenbaum, editors. *The Handbook of Artificial Intelligence*. Morgan Kaufmann, 1982.
- [BGK84.] Alan W. Biermann, Gerard Guiho, and Yves Kodratoff, editors. *Automatic Program Construction Techniques*. Macmillan, 1984.
- [BH84.] W. Bibel and K. M. Hörnig. LOPS - a system based on a strategical approach to program synthesis. In A. W. Biermann, G. Guiho, and Y. Kodratoff, editors, *Automatic program construction techniques*, chapter 3, pages 69–89. MacMillan, 1984.

-
- [Bie92.] A. W. Biermann. Automatic programming. In S. C. Shapiro, editor, *Encyclopedia of Artificial Intelligence*, pages 59–83. John Wiley, second, extended edition, 1992.
- [Bie93.] A.W. Biermann. Personal Communication, 1993.
- [BP90.] Dave Barker-Plummer. Cliché programming in Prolog. In M. Bruynooghe, editor, *Proceedings of Meta'90*, pages 47–256, 1990.
- [Bra86.] I. Bratko. *PROLOG Programming for Artificial Intelligence*. International Computer Science. Addison-Wesley, 1986.
- [BSH90.] A. Bundy, A. Smaill, and J. Hesketh. Turning eureka steps into calculations in automatic program synthesis. In S.L.H. Clarke, editor, *Proceedings of UK IT 90*, pages 221–226, 1990.
- [BSK89.] M. Bruynooghe, D. De Schreye, and B. Krekels. Compiling control. *Journal of Logic Programming*, 6(1-2):135–162, 1989.
- [BSvH⁺91.] A. Bundy, A. Stevens, F. van Harmelen, A. Ireland, and A. Smaill. Rippling: A heuristic for guiding inductive proofs. Research Paper 567, DAI, 1991. Submitted to Artificial Intelligence.
- [BSW90.] A. Bundy, A. Smaill, and G. Wiggins. The synthesis of logic programs from inductive proofs. In J.W. Lloyd, editor, *Proc. Esprit Symposium on Computational Logic*, pages 135–149. Springer-Verlag, 1990.
- [Bur69.] R.M. Burstall. Proving properties of programs by structural induction. *The Computer Journal*, 72:41–48, 1969.
- [Bur74.] R.M. Burstall. Program proving as hand simulation with a little induction. In *IFIP 74*, pages 308–312. North Holland, 1974.
- [BvHHS90.] A. Bundy, F. van Harmelen, C. Horn, and A. Smaill. The Oyster-Clam system. In M.E. Stickel, editor, *Proc. 10th International Conference on Automated Deduction*, pages 647–648. Springer-Verlag, 1990. Lecture Notes in Artificial Intelligence No. 449.
- [BvHHS91.] A. Bundy, F. van Harmelen, J. Hesketh, and A. Smaill. Experiments with proofs plans for induction. *Journal of Automated Reasoning*, 7:303–324, 1991.
- [BvHSI90.] A. Bundy, F. van Harmelen, A. Smaill, and A. Ireland. Extensions to the rippling-out tactic for guiding inductive proofs. In M.E. Stickel, editor, *Proc. 10th International Conference on Automated Deduction*, pages 132–146. Springer-Verlag, 1990. Lecture Notes in Artificial Intelligence No. 449.
- [CD88.] B. Courcelle and P. Deransart. Proofs of partial correctness for attribute grammars with applications to recursive procedures and logic programming. *Information and Computation*, 78(1):1–55, 1988.
- [C^{et al.} 86.] R.L. Constable and *et al.*. *Implementing Mathematics with the NuPRL Proof Development System*. Prentice Hall, 1986.
- [Cla79.] K.L. Clark. Predicate logic as a computational formalism. Technical Report 79/59, Imperial College of Science and Technology, Univ. of London, 1979.
- [Cla81.] K.L. Clark. The synthesis and verification of logic programs. Technical Report DOC 81/36, Imperial College, September 1981. (Revised version of a document which first appeared in 1977).
- [CM84.] W.F. Clocksin and C.S. Mellish. *Programming in Prolog*. Springer Verlag, New York, second edition, 1984.
- [CM91.] L. Colussi and E. Marchiori. Proving correctness of logic programs using axiomatic semantics. In K. Furukawa, editor, *Proceedings of ICLP'91*, pages 629–642. MIT Press, 1991.
- [CS77.] K.L. Clark and S. Sickel. Predicate logic: A calculus for deriving programs. In *Proc. IJCAI-77*, pages 419–420, 1977.

- [CT77.] K.L. Clark and S.-Å. Tärnlund. A first order theory of data and programs. In *Proc. IFIP 77*, pages 939–944. North-Holland, 1977.
- [Day87.] G. Dayantis. Logic program derivation for a class of first order logic relations. In *Proc. IJCAI-87*, pages 9–14, 1987.
- [DB89.] Y. Deville and J. Burnay. Generalization and program schemata: A step towards computer-aided construction of logic programs. In *Proc. North American Conf. on Logic Programming 1989*, pages 409–425. MIT Press, 1989.
- [Dev90.] Y. Deville. *Logic Programming: Systematic Program Development*. International Series in Logic Programming. Addison-Wesley, 1990.
- [DF88.] P. Deransart and G. Ferrand. Logic programming: Methodology and teaching. In K. Fuchi and L. Kott, editors, *Proc. French Japan Symposium*, pages 133–147, August 1988.
- [DF89.] P. Deransart and G. Ferrand. A methodological view of logic programming with negation. Technical Report RR 1011, INRIA, April 1989.
- [DM87.] W. Drabent and J. Maluszynsky. Inductive assertion method for logic programs. In *Proc. TAPSOFT '87, 2, LNCS 250*, pages 167–181. Springer-Verlag, 1987.
- [FD92a.] G. Ferrand and P. Deransart. Proof methods of partial correctness and weak completeness for normal logic programs. In K. Apt, editor, *Proceedings of JICSLP'92*, pages 161–174. MIT Press, 1992.
- [FD92b.] P. Flener and Y. Deville. Towards stepwise, schema-guided synthesis of logic programs. In T.P. Clement and K.K. Lau, editors, *Logic Program Synthesis and Transformation*, Workshops in Computing, pages 46–64. Springer-Verlag, 1992.
- [FD93a.] P. Flener and Y. Deville. Synthesis of composition and discriminate operators for divide-and-conquer logic programs. In J.-M. Jacquet, editor, *Constructing Logic Programs*. Wiley & Sons, 1993. (forthcoming).
- [FD93b.] Pierre Flener and Yves Deville. Logic program synthesis from incomplete specifications. *Journal of Symbolic Computation: Special Issue on Automatic Programming*, 1993. (W. Bibel and A. W. Biermann Editors), Accepted for publication.
- [Fle93.] P. Flener. *Logic Algorithm Synthesis from Examples and properties*. PhD thesis, Unité d'Informatique, Université Catholique de Louvain, Belgium, 1993.
- [Fri90.] L. Fribourg. Extracting logic programs from proofs that use extended Prolog execution and induction. In *Proc. 7th Int. Conf. on Logic Programming*, pages 685–699. MIT Press, 1990.
- [Fri91.] L. Fribourg. Automatic generation of simplification lemmas for inductive proofs. In V. Saraswat and K. Ueda, editors, *Proc. 1991 Int. Logic Programming Symposium*, pages 103–116. MIT Press, 1991.
- [Fri93.] L. Fribourg. Extracting logic programs that use extended Prolog execution and induction. In J.-M. Jacquet, editor, *Constructing Logic Programs*. Wiley & Sons, 1993. (forthcoming).
- [GH89.] T. S. Gegg-Harrison. Basic Prolog schemata. Technical Report CS-1989-20, Dept. of Computer Science, Duke University, 1989.
- [GH90.] T. S. Gegg-Harrison. A scheme-based approach to teaching recursive Prolog programming. Technical Report CS-1990-4, Dept. of Computer Science, Duke University, 1990.
- [GKPV86.] F. Giannesini, H. Kanoui, R. Passero, and M. Van Caneghem. *Prolog*. Addison-Wesley, 1986.
- [GL88.] M. Gelfond and V. Lifschitz. The stable model semantics for logic programming. In B Kowalski and K. Bowen, editors, *Proc. 1988 Int. Logic*

- Programming Symposium*, pages 1070–1080. MIT Press, 1988.
- [GLM92.] M. Gabrielli, G. Levi, and M.C. Meo. Observational equivalences for logic programs. In K. Apt, editor, *Proceedings of JICSLP'92*, pages 131–145. MIT Press, 1992.
- [Gol86.] Allen T. Goldberg. Knowledge-based programming: A survey of program design and construction techniques. *IEEE Trans. on Software Engineering*, 12(7):752–768, July 1986.
- [Hag90.] Masami Hagiya. Programming by example and proving by example using higher-order unification. In M. E. Stickel, editor, *Proceedings of CADE'90*, pages 588–602. Springer Verlag LNCS 449, 1990.
- [Han80.] Å. Hansson. *A Formal Development of Programs*. PhD thesis, Dept. of Information Processing and Computer Science and Computer Science, The Royal Institute of Technology and The University of Stockholm, 1980.
- [Hay86.] S. Hayashi. A system extracting programs from proofs. In *3rd Working Conference on the formal description of programming concepts*, Ebberup, Denmark, 1986.
- [HC92.] Jean Henrard and Baudouin Le Charlier. Folon: An environment for declarative construction of logic programs. In M. Bruynooghe and M. Wirsing, editors, *Proceedings of PLILP'92*, pages 217–231. Springer Verlag LNCS 631, 1992.
- [HL78.] G. Huet and B. Lang. Proving and applying program transformations expressed with second-order patterns. *Acta Informatica*, 11:31–55, 1978.
- [HL91.] P. Hill and J. Lloyd. The Gödel report. Technical Report TR-91-02, Department of Computer Science, University of Bristol, March 1991.
- [Hog78.] C.J. Hogger. Program synthesis in predicate logic. In *Proc. AISB/GI Conf. on Artificial Intelligence, University of Hamburg*, pages 22–27, 1978.
- [Hog81.] C.J. Hogger. Derivation of logic programs. *J. ACM*, 28(2):372–392, April 1981.
- [Hog84.] C.J. Hogger. *Introduction to Logic Programming*. Academic Press, 1984.
- [How80.] W.A. Howard. The formulae-as-type notion of construction. In J.P. Seldin and J.R. Hindley, editors, *To H.B. Curry; Essays on combinatory logic, lambda calculus and formalism*, pages 479–490. Academic Press, 1980.
- [HT79.] Å. Hansson and S.-Å. Tärnlund. A natural programming calculus. In *Proc. IJCAI-79*, pages 348–355, 1979.
- [Kan86.] T. Kanamori. Soundness and completeness of extended execution for proving properties of Prolog programs. Technical Report TR-175, ICOT, May 1986.
- [Kaw91.] T. Kawamura. Derivation of efficient logic programs by synthesizing new predicates. In V. Saraswat and K. Ueda, editors, *Proc. 1991 Int. Logic Programming Symposium*, pages 611–612. MIT Press, 1991.
- [Kaw92.] T. Kawamura. Logic program synthesis from first order logic specifications. In *Proc. Fifth Generation Computer Systems 92*, pages 463–472, Tokyo, 1992. Ohmsha.
- [KBB93a.] I. Kraan, D. Basin, and A. Bundy. Logic program synthesis via proof planning. In K.K. Lau and T. Clement, editors, *Logic Program Synthesis and Transformation, Workshops in Computing*, pages 1–14. Springer-Verlag, 1993.
- [KBB93b.] I. Kraan, D. Basin, and A. Bundy. Middle-out reasoning for logic program synthesis. In D.S. Warren, editor, *Proc. 10th Int. Conf. on Logic Programming*, pages 441–455. MIT Press, 1993.
- [KF86.] T. Kanamori and H. Fujita. Formulation of induction formulas in verification of Prolog programs. In J.H. Siekmann, editor, *Proc. *th International Conference on Automated Deduction*, LNCS 225, pages 281–299. Springer-

- Verlag, 1986.
- [KH87.] T. Kanamori and K. Horiuchi. Construction of logic programs based on generalized unfold/fold rules. In *Proc. 4th Int. Conf. on Logic Programming, Melbourne*, pages 744–768, 1987.
 - [KK88.] T. Kanamori and T. Kawamura. Preservation of stronger equivalence in unfold/fold logic program transformation (ii). Technical Report TR-403, ICOT, June 1988.
 - [Kre90.] C. Kreitz. Towards a formal theory of program construction. *Revue d'Intelligence Artificielle*, 4(3):53–79, 1990.
 - [Kre93a.] C. Kreitz. Formal mathematics for verifiably correct program synthesis, 1993. Forschungsbericht, Technische Hochschule Darmstadt, Germany.
 - [Kre93b.] C. Kreitz. Meta-synthesis — Deriving programs that develop programs, 1993.
 - [KS86.] T. Kanamori and H. Seki. Verification of Prolog programs using an extension of execution. In E. Shapiro, editor, *Proc. 3rd Int. Conf. on Logic Programming*, pages 475–489. Springer-Verlag, 1986. LNCS 255.
 - [KS90.] M. Kirschenbaum and L. Sterling. Prolog programming using skeletons and techniques. Technical Report TR-90-109, Center for Automation and Intelligent Systems Research, Case Western Reserve University, 1990.
 - [KS91.] M. Kirschenbaum and L. Sterling. Refinement strategies for inductive learning of simple Prolog programs. In *Proceedings of IJCAI'91*, pages 757–761, 1991.
 - [Lak89a.] A. Lakhota. *A Workshop for Developing Logic Programs by Stepwise Enhancement*. PhD thesis, Dept. of Computer Engineering and Science, Case Western Reserve University, 1989.
 - [Lak89b.] Arun Lakhota. Incorporating ‘programming techniques’ into Prolog programs. In E. L. Lusk and R. A. Overbeek, editors, *Proceedings of NA-CLP'89*, pages 426–440. MIT Press., 1989.
 - [Lev91a.] J.M. Lever. Combining induction with resolution in logic programming. Technical Report PhD Thesis, Dpt. of Computing, Imperial College, London, 1991.
 - [Lev91b.] J.M. Lever. Proving program properties by means of SLS-resolution. In K. Furukawa, editor, *Proceedings of ICLP'91*, pages 614–628. MIT Press, 1991.
 - [Llo87.] J. W. Lloyd. *Foundations of Logic Programming*. Springer-Verlag, New York, second edition, 1987.
 - [LO92.] K.K. Lau and M. Ornaghi. Towards a formal framework for deductive synthesis of logic programs. Technical Report UMCS-92-11-2, Dept. of Computer Science, University of Manchester, November 1992.
 - [LO93.] K.K. Lau and M. Ornaghi. An incompleteness result for deductive synthesis of logic programs. In D.S. Warren, editor, *Proc. 10th Int. Conf. on Logic Programming*, pages 456–477. MIT Press, 1993.
 - [LO94.] K.K. Lau and M. Ornaghi. A formal view of specification, deductive synthesis and transformation of logic programs. In Y. Deville, editor, *Logic Program Synthesis and Transformation, Workshops in Computing*. Springer-Verlag, 1994. Proceedings of Third International Workshop on Logic Program Synthesis and Transformation.
 - [LP90.] K.K. Lau and S.D. Prestwich. Top-down synthesis of recursive logic procedures from first-order logic specifications. In D.H.D. Warren and P. Szeredi, editors, *Proc. 7th Int. Conf. on Logic Programming*, pages 667–684. MIT Press, 1990.
 - [LP91.] K.K. Lau and S.D. Prestwich. Synthesis of a family of recursive sorting procedures. In V. Saraswat and K. Ueda, editors, *Proc. 1991 Int. Logic*

- Programming Symposium*, pages 641–658. MIT Press, 1991.
- [LS88.] A. Lakhotia and L. Sterling. Composing recursive logic programs with clausal join. *New Generation Computing*, 6:211–225, 1988.
- [LS90.] A. Lakhotia and L. Sterling. Stepwise enhancement: A variant of incremental programming. In *Proc. Conf. on Software Engineering*, 1990.
- [LWD93.] V. Lombart, G. Wiggins, and Y. Deville. Guiding synthesis proofs. In *LOPSTR'93*, 1993. Extended Abstract.
- [Mah88.] M.J. Maher. Equivalences of logic programs. In J. Minker, editor, *Foundations of Deductive Databases and Logic Programming*, pages 627–658. Morgan Kaufmann, 1988.
- [MR93.] S. Muggleton and L. De Raedt. Inductive logic programming: a survey. *Journal of Logic Programming, Special Issue on "Ten Years of Logic Programming"*, 1993. (Submitted).
- [Mug91.] S. Muggleton. Inductive logic programming. *New Generation Computing*, 8(4):295–317, 1991.
- [Mug92.] S. Muggleton, editor. *Inductive Logic Programming*. Volume APIC-38, Academic Press, 1992.
- [MW80.] Z. Manna and R. Waldinger. A deductive approach to program synthesis. *ACM Trans. on Programming Languages and Systems*, 2(1):90–121, January 1980.
- [Nai87.] L. Naish. Specification = Program + Types. In *Proc. FST & TCS*. Springer-Verlag, LNCS, 1987.
- [Nai93.] L. Naish. Verification of logic programs and imperative programs. In J.-M. Jacquet, editor, *Constructing Logic Programs*. Wiley & Sons, 1993. (forthcoming).
- [Neu92.] G. Neugebauer. The IO-graph method: Algorithm design and implementation. 1992. Submitted to the *Journal of Symbolic Computation*.
- [Neu93.] G. Neugebauer. The LOPS approach: A transformation point of view. In K.K. Lau and T.P. Clement, editors, *Logic Program Synthesis and Transformation*, Workshops in Computing, pages 80–81. Springer-Verlag, 1993.
- [O'K90.] R.A. O'Keefe. *The Craft of Prolog*. Logic Programming Series. MIT Press, 1990.
- [Plo70.] Gordon D. Plotkin. A note on inductive generalization. In B. Meltzer and D. Michie, editors, *Machine Intelligence 5*, pages 153–163. Edinburgh University Press, Edinburgh (UK), 1970.
- [Plo71.] Gordon D. Plotkin. A further note on inductive generalization. In B. Meltzer and D. Michie, editors, *Machine Intelligence 6*, pages 101–124. Edinburgh University Press, Edinburgh (UK), 1971.
- [Prz88.] T. Przymusiński. Perfect model semantics. In B Kowalski and K. Bowen, editors, *Proc. 1988 Int. Logic Programming Symposium*, pages 1081–1096. MIT Press, 1988.
- [PS90.] A.J. Power and L. Sterling. A notion of map between logic programs. In *Proc. 7th Int. Conf. on Logic Programming*, pages 390–404. MIT Press, 1990.
- [Rey70.] J.C. Reynolds. Transformational systems and the algebraic structure of atomic formulas. In B. Meltzer and D. Michie, editors, *Machine Intelligence 5*, pages 135–151. Edinburgh University Press, Edinburgh (UK), 1970.
- [RW88.] Charles Rich and Richard C. Waters. Automatic programming: Myths and prospects. *IEEE Computer*, 21(8):40–51, August 1988.
- [Sha81.] Ehud Y. Shapiro. An algorithm that infers theories from facts. In *Proceedings of IJCAI'81*, pages 446–451, 1981.
- [Sha83.] Ehud Y. Shapiro. *Algorithmic Program Debugging*. MIT Press, 1983.

-
- [Smi84.] D.R. Smith. The synthesis of LISP programs from examples. In A. W. Biermann, G. Guiho, and Y. Kodratoff, editors, *Automatic program construction techniques*, chapter 15, pages 307–324. MacMillan, 1984.
- [Smi88.] D.R. Smith. The structure and design of global search algorithms. Technical Report KES.U.87.12, Kestrel Institute, Palo Alto (CA, USA), 1988.
- [Smi90.] D.R. Smith. KIDS: A semiautomatic program development system. *IEE Trans. on Soft. Engineering*, 16(9):1024–1043, 1990.
- [SS86.] L. Sterling and E. Shapiro. *The Art of Prolog*. MIT Press, 1986.
- [ST89.] T. Sato and H. Tamaki. First order compiler: A deterministic logic program synthesis algorithm. *J. Symbolic Computation*, 8:605–627, 1989.
- [Tin90.] Nancy L. Tinkham. *Induction of Schemata for Program Synthesis*. PhD thesis, Duke University, Durham (NC, USA), 1990.
- [VGRS88.] A. Van Gelder, K. Ross, and J.S. Schlipf. Unfounded sets and well-founded semantics for general logic programs. In *Proc. 7th Symposium on Principles of Database Systems*, pages 221–230, 1988.
- [WBKH92.] G. Wiggins, A. Bundy, I. Kraan, and J. Hesketh. Synthesis and transformation of logic programs from constructive, inductive proof. In T.P. Clement and K.K. Lau, editors, *Logic Program Synthesis and Transformation*, Workshops in Computing, pages 27–45. Springer-Verlag, 1992.
- [Wig92a.] G.A. Wiggins. Negation and control in automatically generated logic programs. In A. Pettorossi, editor, *Proceedings of META-92*, 1992.
- [Wig92b.] Geraint Wiggins. Synthesis and transformation of logic programs in the whelk proof development system. In K. Apt, editor, *Proceedings of JIC-SLP'92*, pages 351–365. MIT Press, 1992.