
An Explanation Based Generalization Approach to Inductive Synthesis of Functional Programs

Emanuel Kitzelmann
Ute Schmid

EMANUEL.KITZELMANN@WIAI.UNI-BAMBERG.DE
UTE.SCHMID@WIAI.UNI-BAMBERG.DE

Department of Information Systems and Applied Computer Science, Otto-Friedrich-University, Bamberg

Abstract

We describe an approach to the inductive synthesis of recursive equations from input/output-examples which is based on the classical two-step approach to induction of functional Lisp programs of Summers (1977). In a first step, I/O-examples are rewritten to traces which explain the outputs given the respective inputs based on a datatype theory. This traces can be integrated into one conditional expression which represents a non-recursive program. In a second step, this initial program term is generalized into recursive equations by searching for syntactical regularities in the term. Our approach extends the classical work in several aspects. The most important extensions are that we are able to induce a *set* of recursive equations in one synthesizing step, the equations may contain more than one recursive call, and additionally needed parameters are automatically introduced.

1. Introduction

Automatic induction of recursive programs from input/output-examples (I/O-examples) is an active area of research since the sixties and of interest for AI research as well as for software engineering (Lowry & McCarthy, 1991; Flener & Partridge, 2001). In the seventies and eighties, there were several approaches to the synthesis of Lisp programs from examples or traces (see Biermann et al., 1984 for an overview). The most influential approach was developed by Summers (1977), who put inductive synthesis on a firm theoretical foundation.

Summers' early approach is an explanation based generalization (EBG) approach, thus it widely relies on algorithmic processes and only partially on search: In a first step, traces—steps of computations executed from a program to yield an output from a particular input—and predicates for distinguishing the inputs are calculated for each I/O-pair. Construction of traces, which are *terms* in the

classical functional approaches, relies on knowledge of the inductive datatype of the inputs and outputs. That is, traces *explain* the outputs based on a theory of the used datatype given the respective inputs. The classical approaches for synthesizing Lisp-programs used the general Lisp datatype *S-expression*. By integrating traces and predicates into a conditional expression a non-recursive program explaining all I/O-examples is constructed as result of the first synthesis step. In a second step, regularities are searched for between the traces and predicates respectively. Found regularities are then inductively generalized and expressed in form of the resulting recursive program.

The programs synthesized by Summers' system contain exactly one recursive function, possibly along with one constant term calling the recursive function. Furthermore, all synthesizable functions make use of a small fixed set of Lisp-primitives, particularly of exactly one predicate function, *atom*, which tests whether its argument is an atom, e.g., the empty list. The latter implies two things: First, that Summers' system is restricted to induce programs for *structural* problems on S-expressions. That means, that execution of induced programs depends only on the structure of the input S-expression, but never on the semantics of the atoms contained in it. E.g., *reversing* a list is a structural problem, yet not *sorting* a list. The second implication is, that calculation of the traces is a deterministic and algorithmic process, i.e. does not rely on search and heuristics.

Due to only limited progress regarding the class of programs which could be inferred by functional synthesis, interest decreased in the mid-eighties. There was a renewed interest of inductive program synthesis in the field of inductive logic programming (ILP) (Flener & Yilmaz, 1999; Muggleton & De Raedt, 1994), in genetic programming and other forms of evolutionary computation (Olsson, 1995) which rely heavily on search.

We here present an EBG approach which is based on the methodologies proposed by Summers (1977). We regard the functional two-step approach as worthwhile for the following reasons: First, algebraic datatypes provide guid-

ance in expressing the outputs in terms of the inputs as first synthesis step. Second, it enables a separate and thereby specialized handling of a knowledge dependent part and a purely syntactic driven part of program synthesis. Third, both using algebraic datatypes and separating a knowledge-dependent from a syntactic driven part enables a more accurate utilization of search than in ILP or evolutionary programming. Fourth, the two-step approach using algebraic datatypes provides a systematic way to introduce auxiliary recursive equations if necessary.

Our approach extends Summers in several important aspects, such that we overcome fundamental restrictions of the classical approaches to induction of Lisp programs: First, we are able to induce a *set* of recursive equations in one synthesizing step, second, the equations may contain more than one recursive call, and third, additionally needed parameters are automatically introduced. Furthermore, our generalization step is domain-independent, in particular independent from a certain programming language. It takes as input a first-order term over an arbitrary signature and generalizes it to a recursive program scheme, that is, a set of recursive equations over that signature. Hence it can be used as learning component in all domains which can represent their objects as recursive program schemes and provide a system for solving the first synthesis step. E.g., we use the generalization algorithm for learning recursive control rules for AI planning problems (cp. Schmid & Wysotzki, 2000; Wysotzki & Schmid, 2001).

2. Central Concepts and an Example

The three central objects dealt with by our system are (1) sets of *I/O-examples* specifying the algorithm to be induced, (2) *initial (program) terms* explaining the I/O-examples, and (3) *recursive program schemes (RPSs)* representing the induced algorithms. Their functional role in our two-step synthesis approach is shown in Fig. 1.

An example for I/O-examples is given in Tab. 1. The examples specify the *lasts* function which takes a list of lists as input and yields a list of the last elements of the lists as output. In the first synthesis step, an initial term is constructed from these examples. An initial term is a term respecting an arbitrary first-order signature extended by the special constant symbol Ω , meaning the *undefined* value and directing generalization in the second synthesis step. Suitably interpreted, an initial term evaluates to the specified output when its variable is instantiated with a particular input of the example set and to *undefined* for all other inputs.

Tab. 2 gives an example of an initial term. It shows the result of applying the first synthesis step to the I/O-examples for the *lasts* function as shown in Tab. 1. *if* means the 3-ary non-strict function which returns the value of its second pa-

Table 1. I/O-examples for *lasts*

$[]$	\mapsto	$[],$
$[[a]]$	\mapsto	$[a],$
$[[a, b]]$	\mapsto	$[b],$
$[[a, b, c]]$	\mapsto	$[c],$
$[[a, b, c, d]]$	\mapsto	$[d],$
$[[a], [b]]$	\mapsto	$[a, b],$
$[[a], [b, c]]$	\mapsto	$[a, c],$
$[[a, b], [c], [d]]$	\mapsto	$[b, c, d],$
$[[a, b], [c, d], [e, f]]$	\mapsto	$[b, d, f],$
$[[a], [b], [c], [d]]$	\mapsto	$[a, b, c, d]$

rameter if its first parameter evaluates to *true* and otherwise returns the value of its third parameter; *empty* is a predicate which tests, whether its argument is the empty list; *hd* and *tl* yield the first element and the rest of a list respectively; *cons* constructs a list from one element and a list; and $[]$ denotes the empty list.

Table 2. Initial term for *lasts*

```

if(empty(x), [],
  cons(
    hd(
      if(empty(tl(hd(x))), hd(x),
        if(empty(tl(tl(hd(x)))), tl(hd(x)),
          if(empty(tl(tl(tl(hd(x))))), tl(tl(hd(x)),
             $\Omega$ ))),
        if(empty(tl(x)), [],
          cons(
            hd(
              if(empty(tl(hd(tl(x)))), hd(tl(x)),
                 $\Omega$ )),
              if(empty(tl(tl(x))), [],
                cons(
                  hd(
                    if(empty(tl(hd(tl(tl(x))))), hd(tl(tl(x)),
                       $\Omega$ ))),
                    if(empty(tl(tl(tl(x))))), [],
                       $\Omega$ ))))))

```

Calculation of initial terms relies on knowledge of the datatypes of the example inputs and outputs. For our exemplary *lasts* program inputs and outputs are lists. Lists are uniquely constructed by means of the empty list $[]$ and the constructor *cons*. Furthermore they are uniquely decomposed by the functions *hd* and *tl*. That allows to calculate a unique term which expresses an example output in terms of the input. For example, consider the third I/O-example from Tab. 1: If x denotes the input $[[a, b]]$, then the term $cons(hd(tl(hd(x))), [])$ expresses the specified output $[b]$ in terms of the input. Such traces are constructed for each I/O-pair. The overall concept for integrating the resulting traces into one initial term is to go through all traces in parallel position by position. If the same function symbol is contained at the current position in all traces, then it is in-

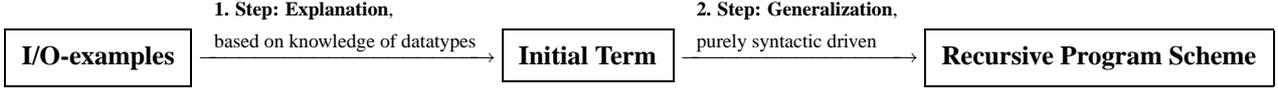


Figure 1. Two synthesis steps

roduced to the initial term at this position. If at least two traces differ at the current position, then it is introduced an *if*-expression. Therefore a predicate function is calculated to discriminate the inputs according to the different traces. Construction of the initial term proceeds from the discriminated inputs and traces for the second and third branch of the *if*-tree respectively. We describe the calculation of initial terms from I/O-examples, i.e. the first synthesis step, in Sec. 4.

In the second synthesis step, initial ground terms are generalized to a recursive program scheme. Initial terms are considered as (*incomplete*) *unfoldings* of an RPS which is to be induced by generalization. An RPS is a set of recursive equations whose left-hand-sides consist of the names of the equations followed by their parameter lists and whose right-hand-sides consist of terms over the signature from the initial terms, the set of the equation names, and the parameters of the equations. One equation is distinguished to be the main one. An example is given in Tab. 3. This RPS, suitably interpreted, computes the *lasts* function as described above and specified by the examples in Tab. 1. It

 Table 3. Recursive Program Scheme for *lasts*

```

lasts(x) = if(empty(x), [],
             cons(hd(last(hd(x))), lasts(tl(x))))
last(x) = if(empty(tl(x)), x, last(tl(x)))
    
```

results from applying the second synthesis step to the initial term shown in Tab. 2. Note that it is a *generalization* from the initial term in that it not merely computes the *lasts* function for the example inputs but for input-lists of arbitrary length containing lists of arbitrary length.

The second synthesis step does not depend on domain knowledge. The meaning of the function symbols is irrelevant, because the generalization is completely driven by detecting syntactical regularities in the initial terms. To understand the link between initial terms and RPSs induced from them, we consider the process of incrementally unfolding an RPS. *Unfolding* of an RPS is a (non-deterministic and possibly infinite) rewriting process which starts with the instantiated head of the main equation of an RPS and which repeatedly rewrites a term by substituting any instantiated head of an equation in the term with either the equally instantiated body or with the special sym-

bol Ω . Unfolding stops, when all heads of recursive equations in the term are rewritten to Ω , i.e., the term contains no rewritable head any more. Consider the *last* equation from the RPS shown in Tab. 3 and the initial instantiation $\{x \mapsto [a, b, c]\}$. We start with the instantiated head $last([a, b, c])$ and rewrite it to the term:

$$\text{if}(\text{empty}(\text{tl}([a, b, c])), [a, b, c], \text{last}(\text{tl}([a, b, c])))$$

This term contains the head of the *last* equation instantiated with $\{x \mapsto \text{tl}([a, b, c])\}$. When we rewrite this head again with the equally instantiated body we obtain:

$$\begin{aligned} &\text{if}(\text{empty}(\text{tl}([a, b, c])), [a, b, c], \\ &\quad \text{if}(\text{empty}(\text{tl}(\text{tl}([a, b, c])), \text{tl}([a, b, c]), \\ &\quad \quad \text{last}(\text{tl}(\text{tl}([a, b, c]))) \end{aligned}$$

This term now contains the head of the equation instantiated with $\{x \mapsto \text{tl}(\text{tl}([a, b, c]))\}$. We rewrite it once again with the instantiated body and then replace the head in the resulting term with Ω and obtain:

$$\begin{aligned} &\text{if}(\text{empty}(\text{tl}([a, b, c])), [a, b, c], \\ &\quad \text{if}(\text{empty}(\text{tl}(\text{tl}([a, b, c])), \text{tl}([a, b, c]), \\ &\quad \quad \text{if}(\text{empty}(\text{tl}(\text{tl}(\text{tl}([a, b, c])), \text{tl}(\text{tl}([a, b, c])), \Omega))) \end{aligned}$$

The resulting *finite* term of a *finite* unfolding process is also called *unfolding*. Unfoldings of RPSs contain regularities if the heads of the recursive equations are more than once rewritten with its bodies before they are rewritten with Ω s. The second synthesis step is based on detecting such regularities in the initial terms.

We describe the generalization of initial terms to RPSs in the following section. The reason why we first describe the second synthesis step and only afterwards the first synthesis step is, that the latter is governed by the goal of constructing a term which can be generalized in the second step. Therefore, for understanding the first step, it is necessary to know the connection between initial terms and RPSs as established in the second step.

3. Generalizing an Initial Term to an RPS

Since our generalization algorithm exploits the relation between an RPS and its unfoldings, in the following we will first introduce the basic terminology for terms, substitutions, and term rewriting as for example presented in Dershowitz and Jouanaud (1990). Then we will present definitions for RPSs and the relation between RPSs and their

unfoldings. The set of all possible RPSs constitutes the hypothesis language for our induction algorithm. Some restrictions on this general hypothesis language are introduced and finally, the components of the generalization algorithm are described.

3.1. Preliminaries

We denote the set of natural numbers starting with 0 by \mathbb{N} and the natural numbers greater 0 by \mathbb{N}_+ . A signature Σ is a set of (function) symbols with $\alpha : \Sigma \rightarrow \mathbb{N}$ giving the arity of a symbol. We write T_Σ for the set of ground terms, i.e. terms without variables, over Σ and $T_\Sigma(X)$ for the set of terms over Σ and a set of variables X . We write $T_{\Sigma,\Omega}$ for the set of ground terms—called partial ground terms—constructed over $\Sigma \cup \{\Omega\}$, where Ω is a special constant symbol denoting the *undefined* value. Furthermore, we write $T_{\Sigma,\Omega}(X)$ for the set of partial terms constructed over $\Sigma \cup \{\Omega\}$ and variables X . With $T_{\Sigma,\Omega}^\infty(X)$ we denote the set of infinite partial terms over Σ and variables X . Over the sets $T_{\Sigma,\Omega}$, $T_{\Sigma,\Omega}(X)$ and $T_{\Sigma,\Omega}^\infty(X)$ a complete partial order (CPO) \leq is defined by: a) $\Omega \leq t$ for all $t \in T_{\Sigma,\Omega}, T_{\Sigma,\Omega}(X), T_{\Sigma,\Omega}^\infty(X)$ and b) $f(t_1, \dots, t_n) \leq f(t'_1, \dots, t'_n)$ iff $t_i \leq t'_i$ for all $i \in [1; n]$.

Terms can uniquely be expressed as labeled trees: If a term is a constant symbol or a variable, then the corresponding tree consists of only one node labeled by the constant symbol or variable. If a term has the form $f(t_1, \dots, t_n)$, then the root node of the corresponding tree is labeled with f and contains from left to right the subtrees corresponding to t_1, \dots, t_n . We use the terms *tree* and *term* as synonyms. A *position* of a term/tree is a sequence of positive natural numbers, i.e. an element from \mathbb{N}_+^* . The set of positions of a term t , denoted $\mathbf{pos}(t)$, contains the empty sequence ε and the position iu , if the term has the form $t = f(t_1, \dots, t_n)$ and u is a position from $\mathbf{pos}(t_i), i \in [1; n]$. Each position of a term uniquely denotes one subterm. We write $t|_u$ for denoting that subterm which is determined as follows: (a) $t|_\varepsilon = t$, (b) if $t = f(t_1, \dots, t_n)$ and u is a position in t_i , then $t|_u = t_i|_u, i \in [1; n]$. We say that position u is smaller than position u' , $u \leq u'$, if u is a prefix of u' . If u is a position of term t and $u' \leq u$, then u' is a position of t . For a term t and a position u , $\mathbf{node}(t, u)$ denotes the fixed symbol $f \in \Sigma$, if $t|_u = f(t_1, \dots, t_n)$ or $t|_u = f$ respectively. The set of all positions at which a fixed symbol f appears in a term is denoted by $\mathbf{pos}(t, f)$. The replacement of a subterm $t|_u$ by a term s in a term t at position u is written as $t[u \leftarrow s]$. Let U denote a set of positions in a term t . Then $t[U \leftarrow s]$ denotes the replacement of all subterms $t|_u$ with $u \in U$ by s in t .

A *substitution* σ is a mapping from variables to terms. Substitutions are naturally continued to mappings from terms to terms by $\sigma(f(t_1, \dots, t_n)) = f(\sigma(t_1), \dots, \sigma(t_n))$. Substitutions are written in postfix notation, i.e. we write $t\sigma$ instead of $\sigma(t)$. Substitutions $\beta : X \rightarrow T_\Sigma$ from variables to

ground terms are called (variable) *instantiations*. A term p is called *pattern* of a term t , iff $t = p\sigma$ for a substitution σ . A pattern p of a term t is called *trivial*, iff p is a variable and *non-trivial* otherwise. We write $t \leq_s p$ iff p is a pattern of t and $t <_s p$ iff additionally holds, that p and t can not be unified by variable renaming only.

A *term rewriting system (TRS)* over Σ and X is a set of pairs of terms $\mathcal{R} \subseteq T_\Sigma(X) \times T_\Sigma(X)$. The elements (l, r) of \mathcal{R} are called rewrite rules and are written $l \rightarrow r$. A term t' can be derived in one rewrite step from a term t using \mathcal{R} ($t \rightarrow_{\mathcal{R}} t'$), if there exists a position u in t , a rule $l \rightarrow r \in \mathcal{R}$, and a substitution $\sigma : X \rightarrow T_\Sigma(X)$, such that (a) $t|_u = l\sigma$ and (b) $t' = t[u \leftarrow r\sigma]$. \mathcal{R} implies a rewrite relation $\rightarrow_{\mathcal{R}} \subseteq T_\Sigma(X) \times T_\Sigma(X)$ with $(t, t') \in \rightarrow_{\mathcal{R}}$ if $t \rightarrow_{\mathcal{R}} t'$.

3.2. Recursive Program Schemes

Definition 1 (Recursive Program Scheme). Given a signature Σ , a set of *function variables* $\Phi = \{G_1, \dots, G_n\}$ for a natural number $n > 0$ with $\Sigma \cap \Phi = \emptyset$ and arity $\alpha(G_i) > 0$ for all $i \in [1; n]$, a natural number $m \in [1; n]$, and a set of equations

$$\mathcal{G} = \begin{cases} G_1(x_1, \dots, x_{\alpha(G_1)}) = t_1, \\ \vdots \\ G_n(x_1, \dots, x_{\alpha(G_n)}) = t_n \end{cases}$$

where the t_i are terms with respect to the signature $\Sigma \cup \Phi$ and the variables $x_1, \dots, x_{\alpha(G_i)}$, $\mathcal{S} = (\mathcal{G}, m)$ is an RPS. $G_m(x_1, \dots, x_{\alpha(G_m)}) = t_m$ is called the main equation of \mathcal{S} .

The function variables in Φ are called *names* of the equations, the left-hand-sides are called *heads*, the right-hand-sides *bodies* of the equations. For the *lasts* RPS shown in Tab. 3 holds: $\Sigma = \{if, empty, cons, hd, tl, []\}$, $\Phi = \{G_1, G_2\}$ with $G_1 = lasts$ and $G_2 = last$, and $m = 1$. \mathcal{G} is the set of the two equations.

We can identify a TRS with an RPS $\mathcal{S} = (\mathcal{G}, m)$:

Definition 2 (TRS implied by an RPS). Let be $\mathcal{S} = (\mathcal{G}, m)$ an RPS over Σ , Φ and X , and Ω the bottom symbol in $T_{\Sigma,\Omega}(X)$. The equations in \mathcal{G} constitute rules $\mathcal{R}_{\mathcal{S}} = \{G_i(x_1, \dots, x_{\alpha(G_i)}) \rightarrow t_i \mid i \in [1; n]\}$ of a term rewriting system. The system additionally contains rules $\mathcal{R}_{\Omega} = \{G_i(x_1, \dots, x_{\alpha(G_i)}) \rightarrow \Omega \mid i \in [1; n], G_i \text{ is recursive}\}$.

The standard interpretation of an RPS, called *free interpretation*, is defined as the supremum in $T_{\Sigma,\Omega}^\infty(X)$ of the set of all terms in $T_{\Sigma,\Omega}(X)$ which can be derived by the implied TRS from the head of the main equation. Two RPSs are called *equivalent*, iff they have the same free interpretation, i.e. if they compute the same function for every interpretation of the symbols in Σ . Terms in $T_{\Sigma,\Omega}$ which can be derived by the *instantiated* head of the main equation

regarding some instantiation $\beta : X \rightarrow T_{\Sigma}$ are called *unfoldings* of an RPS relative to β . Note, that terms derived from RPSs are partial and do not contain function variables, i.e. all heads of the equations are eventually rewritten by Ω s.

The goal of the generalization step is to find an RPS which *explains* a set of initial terms, i.e. to find an RPS such that the initial terms are unfoldings of that RPS. We denote initial terms by \bar{t} and a set of initial terms by \mathcal{I} . We liberalize \mathcal{I} such that it may include *incomplete* unfoldings. Incomplete unfoldings are unfoldings, where some subtrees containing Ω s are replaced by Ω s.

We need to define four further concepts, namely *recursion positions* which are positions in the equation bodies where recursive calls appear, *substitution terms* which are the argument terms in recursive calls, *unfolding positions* which are positions in unfoldings at which the heads of the equations are rewritten with their bodies, and finally *parameter instantiations in unfoldings* which are subterms of unfoldings resulting from the initial parameter instantiation and the substitution terms:

Definition 3 (Recursion Positions and Substitution Terms). Let $G(x_1, \dots, x_{\alpha(G)}) = t$ with parameters $X = \{x_1, \dots, x_{\alpha(G)}\}$ be a recursive equation. The set of *recursion positions* of G is given by $R = \mathbf{pos}(t, G)$. Each recursive call of G at position $r \in R$ in t implies substitutions $\sigma_r : X \rightarrow T_{\Sigma}(X) : x_j \mapsto t|_{r_j}$ for all $j \in [1; \alpha(G)]$ for the parameters in X . We call the terms $t|_{r_j}$ *substitution terms* of G .

For equation *lasts* of the *lasts* RPS (Tab. 3) holds $R = \{32\}$ and $x\sigma_{32} = tl(x)$. For equation *last* holds $R = \{3\}$ and $x\sigma_3 = tl(x)$.

Now consider an unfolding process of a recursive equation and the positions at which rewrite steps are applied in the intermediate terms. The first rewriting is applied at root-position ε , since we start with the instantiated head of the equation which is completely rewritten with the instantiated body. In the instantiated body, rewrites occur at recursion positions R . Assume that on recursion position $r \in R$ the instance of the head is rewritten with an instance of the body. Then, relative to the resulting *subtree* at position r , rewrites occur again at recursion positions, e.g. at position $r' \in R$. Relative to the entire term these latter rewrites occur therefore at compositions of position r and recursion positions, e.g. at position rr' and so on. We call the infinite set of positions at which rewrites can occur in the intermediate terms within an unfolding of a recursive equation *unfolding positions*. They are determined by the recursion positions as follows:

Definition 4 (Unfolding Positions). Let be R the recursion positions of a recursive equation G . The set of *unfolding positions* U of G is defined as the smallest set of positions

which contains the position ε and, if $u \in U$ and $r \in R$, the position ur .

The unfolding positions of equation *lasts* of the *lasts* RPS are $\{32, 3232, 323232, \dots\}$.

Now we look at the variable instantiations occurring during unfolding a recursive equation. Recall the unfolding process of the *last* equation (see Tab. 3) described at the end of Sec. 2. The initial instantiation was $\beta_{\varepsilon} = \beta = \{x \mapsto [a, b, c]\}$, thus in the body of the equation (replaced for the instantiated head as result of the first rewrite step), its variable is instantiated with this initial instantiation. Due to the substitution term $tl(x)$, the variable of the head in this body is instantiated with $\beta_3 = \sigma_3 \beta_{\varepsilon} = \{x \mapsto tl([a, b, c])\}$, i.e. the variable in the body replaced for this instantiated head is instantiated with $\sigma_3 \beta_{\varepsilon}$. A further rewriting step implies the instantiation $\beta_{33} = \sigma_3 \sigma_3 \beta_{\varepsilon} = \sigma_3 \beta_3 = \{x \mapsto tl(tl([a, b, c]))\}$ and so on. We index the instantiations occurring during unfolding with the unfolding positions at which the particular instantiated heads were placed. They are determined by the substitutions implied by recursive calls and an initial instantiation as follows:

Definition 5 (Instantiations in Unfoldings). Let be $G(x_1, \dots, x_{\alpha(G)}) = t$ a recursive equation with parameters $X = \{x_1, \dots, x_{\alpha(G)}\}$, R and U the recursion positions and unfolding positions of G resp., σ_r the substitutions implied by the recursive call of G at position $r \in R$, and $\beta : X \rightarrow T_{\Sigma}$ an initial instantiation. Then a family of instantiations indexed over U is defined as $\beta_{\varepsilon} = \beta$ and $\beta_{ur} = \sigma_r \beta_u$ for $u \in U, r \in R$.

3.3. Restrictions and the Generalization Problem

An RPS which can be induced from initial terms is restricted in the following way: First, it contains no mutual recursive equations, second, there are no calls of recursive equations within calls of recursive equations (no nested recursive calls). The first restriction is not a semantic restriction, since each mutual recursive program can be transformed to an equivalent (regarding a particular algebra) non-mutual recursive program. Yet it is a *syntactical* restriction, since unfoldings of mutual RPSs can not be generalized using our approach. A restriction similar to the second one was stated by Rao (2004). He names TRSs complying with such a restriction *flat* TRSs.

Inferred RPSs conform to the following syntactical characteristics: First, all equations, potentially except of the main equation, are recursive. The main equation may be recursive as well, but, as only equation, it is not required to be recursive. Second, inferred RPSs are minimal, in that (i) each equation is directly or indirectly (by means of other equations) called from the main equation, and (ii) no parameter of any equation can be omitted without changing the free

interpretation. RPSs complying with the stated restrictions and characteristics are called *minimal, non-mutual, flat recursive program schemes*.

There might be several RPSs which explain an initial term \bar{t} , but have different free interpretations. For example, Ω is an unfolding of *every* RPS with a recursive main equation. Therefore, an important question is which RPS will be induced. Summers (1977) required that recurrence relations hold at least over three succeeding traces and predicates to justify a generalization. A similar requirement would be that induced RPSs explain the initial terms *recurrently*, meaning that \mathcal{S} contains at least one term \bar{t} which can be derived from an unfolding process, in which each recursive equation had to be rewritten at least three times with its body. We use a slightly different requirement: One characteristic of minimal RPSs is, that if at least one substitution term is replaced by another, then the resulting RPS has a different free interpretation. We call this characteristic *substitution uniqueness*. Thus, it is sensible to require that induced RPSs are substitution unique regarding the initial terms, i.e. that if some substitution term is changed, then the resulting RPS *no longer* explains the initial terms. It holds, that a minimal RPS explains a set of initial trees recurrently, if it explains it substitution uniquely.

Thus the problem of generalizing a set of initial terms \mathcal{S} to an RPS is to find an RPS which explains \mathcal{S} and which is substitution unique regarding \mathcal{S} .

3.4. Solving the Generalization Problem

We will not state the generalization algorithm in detail in this section but we will describe the underlying concepts and the algorithm in a more informal manner. For this section and its subsections we use the term *body* of an equation for terms which are strictly speaking *incomplete* bodies: They contain only the name of the equation instead of complete recursive calls including substitution terms at recursion positions. For example, we refer to the term $if(empty(x), [], cons(hd(last(hd(x))), lasts))$ as the body for equation *lasts* of the *lasts* RPS (see Tab. 3). The reason is, that we infer the complete body in two steps: First the term which we name body in this context, second the substitution terms for the recursive calls.

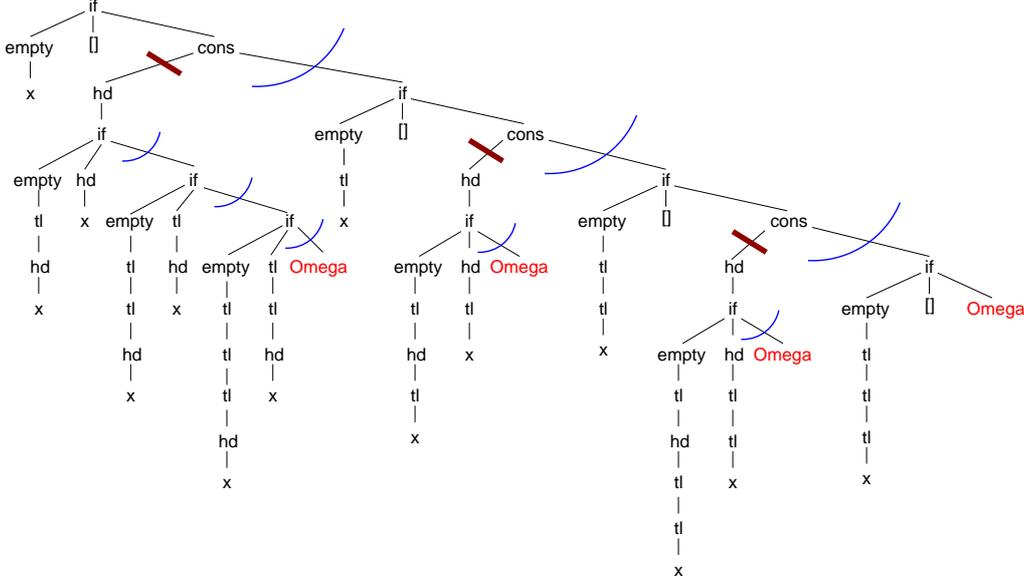
Generalization of a set of initial terms to an RPS is done in three successive steps, namely *segmentation of the terms*, *construction of equation bodies* and *calculation of substitution terms*. These three generalization steps are organized in a divide-and-conquer algorithm, where backtracking can occur to the divide-phase. *Segmentation* constitutes the divide-phase which proceeds top-down through the initial terms. Within this phase recursion positions (see Def. 3) and positions indicating further recursive equations are searched for each induced equation. The latter set of

positions is called *subscheme positions* (see Def. 6 below). Found recursion positions imply unfolding positions (see Def. 4). As result of the divide-phase the initial terms are divided into several parts by the subscheme positions, such that—roughly speaking—each particular part is assumed to be an unfolding of *one* recursive equation. Furthermore, the particular parts are segmented by the unfolding positions, such that—roughly speaking—each segment is assumed to be the result of *one* unfolding step of the respective recursive equation.

Consider the initial tree in Fig. 2, it represents the initial term for *lasts*, shown in Tab. 2. The curved lines on the path to the rightmost Ω divide the tree into three segments which correspond to unfolding steps of the main equation, i.e. equation *lasts*. The short broad lines denote three subtrees which are—except of their root *hd*—unfoldings of the *last* equation. The curved lines within these subtrees divide each subtree into segments, such that each segment correspond to one unfolding step of the *last* equation.

When the initial trees are segmented, calculation of equation bodies and of substitution terms follows within the conquer-phase. These two steps proceed bottom-up through the divided initial trees and reduce the trees during this process. The effect is, that bodies and substitution terms for each equation are calculated from trees which are unfoldings of *only* the currently induced equation and hence, each segment in these trees is an instantiation of the body of the currently induced equation. E.g., for the *lasts* tree shown in Fig. 2, a body and substitution terms are first calculated from the three subtrees, i.e. for the *last* equation. Since there are no further recursive equations called by the *last* equation—i.e. the segments of the three subtrees contain themselves no subtrees which are unfoldings of further equations—each segment is an instantiation of the body of the *last* equation. When this equation is completely inferred, the three subtrees are replaced by suitable instantiations of the head of the inferred *last* equation. The resulting reduced tree is an unfolding of merely *one* recursive equation, the *lasts* equation. The three segments in this reduced tree—indicated by the curved lines on the path to the rightmost Ω —are instantiations of the body of the searched for *lasts* equation. From this reduced tree, body and substitution terms for the *lasts* equation are induced and the RPS is completely induced.

Segmentations are searched for, whereas calculation of bodies and substitution terms are algorithmic. Construction of bodies always succeeds, whereas calculation of substitution terms—such that the inferred RPS explains the initial terms—may fail. Thus, an inferred RPS can be seen as the result of a search through a hypothesis space where the hypotheses are segmentations (divide-phase), and a *constructive* goal test, including construction of bodies and calcu-


 Figure 2. Initial Tree for *lasts*

lation of substitution terms (conquer-phase), which tests, whether the completely inferred RPS explains the initial terms (and is substitution unique regarding them). In the following we describe each step in more detail:

3.4.1. SEGMENTATION

When induction of an RPS from a set of initial trees \mathcal{S} starts, the hypothesis is, that there exists an RPS with a recursive main equation which explains \mathcal{S} . First, recursion and subscheme positions for the hypothetical main equation G_m are searched for.

Definition 6 (Subscheme Positions). Subscheme positions are all smallest positions in the body of a recursive equation G which denote subterms, in which calls of further recursive equations from the RPS appear, but no recursive call of equation G .

E.g., the only subscheme position of equation *lasts* of the *lasts* RPS (Tab. 3) is $u = 31$. A priori, only particular positions from the initial trees come into question as recursion and subscheme positions, namely those which belong to a path leading from the root to an Ω . The reason is, that eventually *each* head of a recursive equation at any unfolding position in an intermediate term while unfolding this equation is rewritten with an Ω :

Lemma 1 (Recursion and Subscheme Positions imply Ω s). Let $\bar{t} \in T_{\Sigma, \Omega}$ be an (incomplete) unfolding of an RPS $\mathcal{S} = (\mathcal{G}, m)$ with a recursive main equation G_m . Let R, U and S be the sets of recursion, unfolding and subscheme positions of G_m respectively. Then for all $u \in U \cap \mathbf{pos}(\bar{t})$ holds:

1. $\mathbf{pos}(\bar{t}|_u, \Omega) \neq \emptyset$
2. $\forall s \in S : \text{if } us \in \mathbf{pos}(\bar{t}) \text{ then } \mathbf{pos}(\bar{t}|_{us}, \Omega) \neq \emptyset$

It is not very difficult to see that this lemma holds. For a lack of space we do not give the proof here. It can be found in (Kitzelmann, 2003) where Lem. 1 and Lem. 2 are proven as one lemma. Knowing Lem. 1, before search starts, the initial trees can be reduced to their *skeletons* which are terms resulting from replacing subtrees without Ω s with variables.

Definition 7 (Skeleton). The skeleton of a term $t \in T_{\Sigma, \Omega}(X)$, written $\mathbf{skeleton}(t)$ is the minimal pattern of t for which holds $\mathbf{pos}(t, \Omega) = \mathbf{pos}(\mathbf{skeleton}(t), \Omega)$.

For example, consider the subtree indicated by the leftmost short broad line of the tree in Fig. 2. Omitting the root *hd*, it is an unfolding of the *last* equation of the *lasts* RPS shown in Tab. 3. Its skeleton is the *substantially* reduced term $\text{if}(x_1, x_2, \text{if}(x_3, x_4, \text{if}(x_5, x_6, \Omega)))$. Search for recursion and subscheme positions is done on the skeletons of the original initial trees. Thereby the hypothesis space is substantially narrowed without restricting the hypothesis language, since only those hypotheses are ruled out which are a priori known to fail the goal test.

Ω s are not only implied by recursion and subscheme positions, but also imply Ω s recursion and subscheme positions since Ω s in unfoldings result *only* from rewriting an instantiated head of a recursive equation in a term with an Ω :

Lemma 2 (Ω s imply recursion and subscheme positions). Let $\bar{t} \in T_{\Sigma, \Omega}$ be an (incomplete) unfolding of an RPS $\mathcal{S} = (\mathcal{G}, m)$ with a recursive main equation G_m . Let R, U

and S be the sets of recursion, unfolding and subscheme positions of G_m respectively. Then for all $v \in \mathbf{pos}(\bar{t}, \Omega)$ hold

- It exists an $u \in U \cap \mathbf{pos}(\bar{t}), r \in R$ with $u \leq v < ur$ or
- it exists an $u \in U \cap \mathbf{pos}(\bar{t}), s \in S$ with $us \leq v$.

Proof: in (Kitzelmann, 2003).

From the definition of subscheme positions and the previous lemma follows, that subscheme positions are determined, if a set of recursion positions has been fixed. Lem. 1 restricts the set of positions which come into question as recursion and subscheme positions. Lem. 2 together with characteristics from subscheme positions suggests to organize the search as a search for recursion positions with a depending parallel calculation of subscheme positions. When hypothetical recursion, unfolding, and subscheme positions are determined they are checked regarding the labels in the initial trees on paths leading to Ω s. The nodes between one unfolding position and its successors in unfoldings result from the same body (with different instantiations). Since variable instantiations only occur in subtrees at positions *not* belonging to paths leading to Ω s, for each unfolding position the nodes between it and its successors are necessarily equal:

Lemma 3 (Valid Segmentation). Let $\bar{t} \in T_{\Sigma, \Omega}$ be an unfolding of an RPS $\mathcal{S} = (\mathcal{G}, m)$ with a recursive main equation G_m . Then it exists a term $\check{t}_G \in T_{\Sigma, \Omega}(X)$ with $\mathbf{pos}(\check{t}_G, \Omega) = R \cup S$ such that for all $u \in U \cap \mathbf{pos}(\bar{t})$ hold: $\check{t}_G \leq_{\Omega} \bar{t}|_u$ where \leq_{Ω} is defined as (a) $\Omega \leq_{\Omega} t$ if $\mathbf{pos}(t, \Omega) \neq \emptyset$, (b) $x \leq_{\Omega} t$ if $x \in X$ and $\mathbf{pos}(t, \Omega) = \emptyset$, and (c) $f(t_1, \dots, t_n) \leq_{\Omega} f(t'_1, \dots, t'_n)$ if $t_i \leq_{\Omega} t'_i$ for all $i \in [1; n]$.

Proof: in (Kitzelmann, 2003).

This lemma has to be slightly extended, if one allows for initial trees which are incomplete unfoldings. Lem. 3 states the requirements to assumed recursion and subscheme positions which can be assured at segmentation time. They are necessary for an RPS which explains the initial terms, yet not sufficient to assure, that an RPS complying with them exists which explains the initial trees. That is, later a backtrack can occur to search for other sets of recursion and subscheme positions. If found recursion and subscheme positions R and S comply with the stated requirements, we call the pair (R, S) a *valid segmentation*.

In our implemented system the search for recursion positions is organized as a greedy search through the space of sets of positions in the skeletons of the initial trees. When a valid segmentation has been found, compositions of unfolding and subscheme positions denote subtrees in the initial trees assumed to be unfoldings of further recursive equations. Segmentation proceeds recursively on each set

of (sub)trees denoted by compositions of unfolding positions and one subscheme position $s \in S$. We denote such a set of initial (sub)trees \mathcal{S}_s .

3.4.2. CONSTRUCTION OF EQUATION BODIES

Construction of each equation body starts with a set of initial trees \mathcal{S} for which at segmentation time a valid segmentation (R, S) has been found, and an already inferred RPS for each subscheme position $s \in S$ which explains the subtrees \mathcal{S}_s . These subtrees of the trees in \mathcal{S} are replaced by the suitably instantiated heads or respectively bodies of the main equations of the already inferred RPSs. For example, consider the initial tree for *lasts* shown in Fig. 2. When calculation of a body for the main equation *lasts* starts from this tree, an RPS containing only the *last* equation which explains all three subtrees indicated by the short broad lines has already been inferred. The initial tree is reduced by replacing these three subtrees by suitable instantiations of the head of the *last* equation. We denote the set of reduced initial trees also with \mathcal{S} and its elements also with \bar{t} . By reducing the initial trees based on already inferred recursive equations, the problem of inducing a set of recursive equations is reduced to the problem of inducing merely *one* recursive equation (where the recursion positions are already known from segmentation).

An equation body is induced from the segments of an initial tree which is assumed to be an unfolding of one recursive equation.

Definition 8 (Segments). Let be \bar{t} an initial tree, R a set of (hypothetical) recursion positions and U the corresponding set of unfolding positions. The set of *complete segments* of \bar{t} is defined as: $\{\bar{t}|_u[R \leftarrow G] \mid u \in U \cap \mathbf{pos}(\bar{t}), R \subset \bar{t}|_u\}$

For example, consider the subtree indicated by the leftmost short broad line of the initial tree in Fig. 2 without its root *hd*. It is an unfolding of the *last* equation as stated in Tab. 3. When the only recursion position 3 has been found it can be splitted into three segments, indicated by the curved lines:

1. $\text{if}(\text{empty}(\text{tl}(\text{hd}(x))), \text{hd}(x), G)$
2. $\text{if}(\text{empty}(\text{tl}(\text{tl}(\text{hd}(x))))), \text{tl}(\text{hd}(x)), G)$
3. $\text{if}(\text{empty}(\text{tl}(\text{tl}(\text{tl}(\text{hd}(x))))), \text{tl}(\text{tl}(\text{hd}(x))), G)$

Expressed according to segments, the fact of a repetitive pattern between unfolding positions (see Lem. 3) becomes the fact, that the sequences of nodes between the root and each G are equal for each segment. Each segment is an instantiation of the body of the currently induced equation. In general, the body of an equation contains other nodes among those between its root and the recursive calls. These further nodes are also equal in each segment. Differences in segments of unfoldings of a recursive equation can *only*

result from different instantiations of the variables of the body. Thus, for inducing the body of an equation from segments, we assume each position in the segments which is equally labeled in *all* segments as belonging to the body of the assumed equation, but each position which is variably labeled in at least two segments as belonging to the instantiation of a variable. This assumption can be seen as an inductive bias since it might occur, that also positions which are equal over all segments belong to a variable instantiation. Nevertheless it holds, that if an RPS exists which explains a set of initial trees, then it also exists an RPS which explains the initial trees and is constructed based on the stated assumption. Based on the stated assumption, the body of the equation to be induced is determined by the segments and defined as follows:

Definition 9 (Valid Body). Given a set of reduced initial trees, the most specific maximal pattern of all segments of all the trees is called *valid body* and denoted \hat{t}_G .

The maximal pattern of a set of terms can be calculated by first order anti-unification (Plotkin, 1969).

Calculating a valid body regarding the three segments enumerated above results in the term $if(empty(tl(x)),x,G)$. The *different* subterms of the segments are assumed to be instantiations of the parameters in the calculated valid body. Since each segment corresponds to one unique unfolding position, instantiations of parameters in unfoldings as defined in Def. 5 are now given. E.g., from the three segments enumerated above we obtain:

1. $\beta_\varepsilon(x) = hd(x)$
2. $\beta_3(x) = tl(hd(x))$
3. $\beta_{33}(x) = tl(tl(hd(x)))$

3.4.3. INDUCING SUBSTITUTION TERMS

Induction of substitution terms for a recursive equation starts on a set of reduced initial trees which are assumed to be unfoldings of one recursive equation, an already inferred (incomplete) equation body which contains only a G at recursion positions, and variable instantiations in unfoldings according to Def. 5. The goal is to complete each occurrence of G to a recursive call including substitution terms for the parameters of the recursive equation.

The following lemma follows from Def. 5 and states characteristics of parameter instantiations in unfoldings more detailed. It characterizes the instantiations in unfoldings against the substitution terms of a recursive equation considering each single position in them.

Lemma 4 (Instantiations in Unfoldings). Let be $G(x_1, \dots, x_{\alpha(G)}) = t$ a recursive equation with parameters $X = \{x_1, \dots, x_{\alpha(G)}\}$, recursion positions R and unfolding

positions U , $\beta : X \rightarrow T_\Sigma$ an instantiation, σ_r substitution terms for each $r \in R$ and β_u instantiations as defined in Def. 5 for each $u \in U$. Then for all $i, j \in [1; \alpha(G)]$ and positions v hold:

1. If $(x_i \sigma_r)|_v = x_j$ then for all $u \in U$ hold $(x_i \beta_{ur})|_v = x_j \beta_u$.
2. If $(x_i \sigma_r)|_v = f((x_i \sigma_r)|_{v_1}, \dots, (x_i \sigma_r)|_{v_m}), f \in \Sigma, \alpha(f) = n$ then for all $u \in U$ hold $\mathbf{node}(x_i \beta_{ur}, v) = f$.

We can read the implications stated in the lemma in the inverted direction and thus we get almost immediately an algorithm to calculate the substitution terms of the searched for equation from the known instantiations in unfoldings.

One interesting case is the following: Suppose a recursive equation, in which at least one of its parameters *only* occurs within a recursive call in its body, for example the equation $G(x, y, z) = if(zerop(x), y, +(x, G(prev(x), z, succ(y))))$ in which this is the case for parameter z ¹. For such a variable no instantiations in unfoldings are given when induction of substitution terms starts. Also such variables are not contained in the (incomplete) valid equation body. Our generalizer introduces them each time, when none of the both implications of Lem. 4 hold. Then it is assumed, that the currently induced substitution term contains such a “hidden” variable at the current position. Based on this assumption the instantiations in unfoldings of the hidden variable can be calculated and the inference of substitution terms for it proceeds as described for the other parameters.

When substitution terms have been found, it has to be checked, whether they are substitution unique with regard to the reduced initial terms. This can be done for each substitution term that was found separately.

3.4.4. INDUCING AN RPS

We have to consider two further points: The first point is that segmentation presupposes the initial trees to be explainable by an RPS with a *recursive* main equation. Yet in Sec. 3.3 we characterized the inferable RPSs as liberal in this point, i.e. that also RPSs with a non-recursive main equation are inferable. In such a case, the initial trees contain a constant (not repetitive) part at the root such that no recursion positions can be found for these trees (as for example the three subtrees indicated by the short broad lines in Fig. 2 which contain the constant root hd). In this case, the root node of the trees is assumed to belong to the body of a non-recursive main equation and induction of RPSs recursively proceeds at each subtree of the root nodes.

The second point is that RPSs explaining the subtrees which are assumed to be unfoldings of further recursive equations at segmentation time are already inferred. Based

¹A practical example is the Tower of Hanoi problem

on these already inferred RPSs, the initial trees are reduced and then a body and substitution terms are induced. Calculation of a body always succeeds, whereas calculation of substitution terms may fail. One important question is, whether success of calculating substitution terms depends on the already inferred RPSs. Is the (problematic) case possible, that two different sets of RPSs both explain the subtrees which will be replaced, such that calculation of substitution terms from the reduced initial trees succeeds presupposed one particular set of RPSs, but fails for the other set? Fortunately we could prove, that this is *not* possible. If a set of RPSs explaining the subtrees exists such that substitution terms can be calculated, then substitution terms can be calculated presupposed *any* set of RPSs explaining the subtrees. Thus, inducing RPSs for the subtrees can be dealt with as an independent problem as it is done in our divide-and-conquer approach.

4. Generating an Initial Term

Our theory and prototypical implementation for the first synthesis step uses the datatype *List*, defined as follows: The empty list \square is an (α -)list and if a is in element of type α and l is an α -list, then $\text{cons}(a, l)$ is an α -list. Lists may contain lists, i.e. α may be of type *List* α' .

4.1. Characterization of the Approach

The constructed initial terms are composed from the list constructor functions \square, cons , the functions for decomposing lists hd, tl , the predicate *empty* testing for the empty list, one variable x , the 3ary (non-strict) conditional function *if* as control structure, and the bottom constant Ω meaning *undefined*. Similar to Summers (1977), the set of functions used in our term construction approach implies the restriction of induced programs to solve *structural* list programs. An extension to Summers is that we allow the example inputs to be *partially* ordered instead of only totally ordered. This is related to the extension of inducing *sets* of recursive equations as described in Sec. 3 instead of only one recursive equation.

We say that an initial term *explains* I/O-examples, if it evaluates to the specified output when applied to the respective input or to *undefined*. The goal of the first synthesis step is to construct an initial term which explains a set of I/O-examples and which can be explained by an RPS.

4.2. Basic Concepts

Definition 10 (Subexpressions). The set of subexpressions of a list l is defined to be the smallest set which includes l itself and, if l has the form $\text{cons}(a, l')$, all subexpressions of a and of l' . If a is an atom, then a itself is its only subexpression.

Since hd and tl —which are defined by $hd(\text{cons}(a, l)) = a$ and $tl(\text{cons}(a, l)) = l$ —decompose lists uniquely, each subexpression can be associated with the unique term which computes the subexpression from the original list. E.g., consider the following I/O-pair which is the third one from Tab. 1: $[[a, b]] \mapsto [b]$. The set of all subexpressions of the input list $[[a, b]]$ together with their associated terms is: $\{x = [[a, b]], hd(x) = [a, b], tl(x) = \square, hd(hd(x)) = a, tl(hd(x)) = [b], hd(tl(hd(x))) = b, tl(tl(hd(x))) = \square\}$.

Since lists are uniquely constructed by the constructor functions \square and cons , traces which compute the specified output can uniquely be constructed from the terms for the subexpressions of the respective input:

Definition 11 (Construction of Traces). Let $i \mapsto o$ be an I/O-pair (i is a list). If o is a subexpression of i , then the trace is defined to be the term associated with o . Otherwise o has the form $\text{cons}(a, l)$. Let t and t' be the traces for the I/O-pairs $i \mapsto a$ and $i \mapsto l$ respectively. The trace for $i \mapsto o$ is defined to be the term $\text{cons}(t, t')$.

E.g., the trace for computing the output $[b]$ from its input $[[a, b]]$ is the term $\text{cons}(hd(tl(hd(x))), \square)$.

Similar to Summers, we discriminate the inputs with respect to their structure, more precisely wrt a partial order over them implied by their structural complexity. As stated above, we allow for arbitrarily nested lists as inputs. A partial order over such lists is given by: $\square \leq l$ for all lists l and $\text{cons}(a, l) \leq \text{cons}(a', l')$, iff $l \leq l'$ and, if a and a' are again lists, $a \leq a'$.

Consider any unfolding of an RPS. Generally it holds, that greater positions on a path leading to an Ω result from more rewritings of a head of a recursive equation with its body compared to some smaller position. In other words, the computation represented by a node at a greater position is one on a deeper recursion level than a computation represented by a smaller position. Since we use only the complexity of an input list as criterion whether the recursion stops or whether another call appears with the input decomposed in some way, deeper recursions result from more complex inputs in the induced programs.

4.3. Solving the Term Construction Problem

The overall concept of constructing the initial tree is to introduce the nodes from the traces position by position to the initial tree *as long as the traces are equal* and to introduce an *if*-expression as soon as at least two (sub)traces differ. The predicate in the *if*-expression divides the inputs into two sets. The “then”-subtree is recursively constructed from the input/trace-pairs whose inputs evaluate to *true* with the predicate and the “else”-subtree is recursively constructed from the other input/trace-pairs. Eventually only one single input/trace-pair remains when an *if*-

expression is introduced. In this case an Ω indicating a recursive call on this path is introduced as leaf at the current position in the initial term and (this subtree of) the initial tree is finished. The reason for introducing an Ω in this case is, that we assume, that *if* the input/trace-set would contain a pair with a more complex input, than the respective trace would at some position differ from the remaining trace and thus it would imply an *if*-expression, i.e. a recursive call at some deeper position. Since we do not know the position at which this difference would occur, we can not use this single trace, but have to indicate a recursive call on this path by an Ω . Thus, for principal reasons, the constructed initial terms are undefined for the most complex inputs of the example set.

We now consider the both cases that all roots of the traces are equal and that they differ respectively more detailed.

4.3.1. EQUAL ROOTS

Suppose all generated traces have the same root symbol. In this case, this symbol constitutes the root of the initial tree. Subsequently the sub(initial)trees are calculated through a recursive call to the algorithm. Suppose the initial tree has to explain the I/O-examples $\{[a] \mapsto a, [a, b] \mapsto b, [a, b, c] \mapsto c\}$. Calculating the traces and replacing them for the outputs yields the input/trace-set $\{[a] \mapsto hd(x), [a, b] \mapsto hd(tl(x)), [a, b, c] \mapsto hd(tl(tl(x)))\}$. All three traces have the same root *hd*, thus we construct the root of the initial tree with this symbol. The algorithm for constructing the initial subterm of the constructed root *hd* now starts recursively on the set of input/trace-pairs where the traces are the subterms of the roots *hd* from the three original traces, i.e. on the set $\{[a] \mapsto x, [a, b] \mapsto tl(x), [a, b, c] \mapsto tl(tl(x))\}$.

The traces from these new input/trace-set have different roots, that is, an *if*-expression is introduced as subtree of the constructed initial tree.

4.3.2. INTRODUCING CONTROL STRUCTURE

Suppose the traces (at least two of them) have different roots, as for example the traces of the second input/trace-set in the previous subsection. That means that the initial term has to apply different computations to the inputs corresponding to the different traces. This is done by introducing the conditional function *if*, i.e. the root of the initial term becomes the function symbol *if* and contains from left to right three subtrees: First, a predicate term with the predicate *empty* as root to distinguish between the inputs which have to be computed differently wrt their complexity; second, a tree explaining all I/O-pairs whose inputs are evaluated to *true* from the predicate term; third, a tree explaining the remaining I/O-examples. It is presupposed, that all traces corresponding to inputs evaluating to *true* with the predicate are equal. These equal subtraces be-

come the second subtree of the *if*-expression, i.e. they are evaluated, if an input evaluates to *true* with the predicate. That means that never an Ω occurs in a “then”-subtree of a constructed initial tree, i.e. that recursive calls in the induced RPSs may only occur in the “else”-subtrees. For the “else”-subtree the algorithm is recursively processed on all remaining input/trace-pairs.

5. Experimental Results

We have implemented prototypes (without any thoughts about efficiency) for both described steps, construction of the initial tree and generalization to an RPS. The implementations are in Common-Lisp. In Tab. 4 we have listed experimental results for a few sample problems. The first column lists the names for the induced functions, the second column lists the number of given I/O-pairs, the third column lists the number of induced equations, and the fourth column lists the times consumed by the first step, the second step, and the total time respectively. The experiments were performed on a Pentium 4 with Linux and the program runs are interpreted with the *clisp* interpreter.

Table 4. Some inferred functions

function	#expl	#eqs	times in sec
<i>last</i>	4	2	.003 / .001 / .004
<i>init</i>	4	1	.004 / .002 / .006
<i>evenpos</i>	7	2	.01 / .004 / .014
<i>switch</i>	6	1	.012 / .004 / .016
<i>unpack</i>	4	1	.003 / .002 / .005
<i>lasts</i>	10	2	.032 / .032 / .064
<i>mult-lasts</i>	11	3	.04 / .49 / .53
<i>reverse</i>	6	4	.031 / .036 / .067

All induced programs compute the intended function. The number of given examples is in each case the minimal one. When given one example less, the system does *not* produce an *unintended* program, but produces *no* program. Indeed, an initial term is produced in such a case which is correct on the example set, but no RPS is generalized, because it exists no RPS which explains the initial term *and is substitution unique wrt it* (see Sec. 3.3).

last computes the last element of a list. The main equation is not recursive and only applies a *hd* to the result of the induced recursive equation which computes a one element list containing the last element of the input list. *init* returns the input list without the last element. *evenpos* computes a list containing each second element of the input list. The main equation is not recursive and only deals with the empty input list as special case. *switch* returns a list, in which each two successive elements of the input list are on switched positions, e.g., $switch([a, b, c, d, e]) = [b, a, d, c, e]$. *unpack* produces an output list, in which each element

from the input list is encapsulated in a one element list, e.g., $unpack([a,b,c]) = [[a],[b],[c]]$. $unpack$ is the classical example in (Summers, 1977). $lasts$ is the program described in Sec. 2. The given I/O-examples are those from Tab. 1. $mult-lasts$ takes a list of lists as input just like $lasts$. It returns a list of the same structure as the input list where each inner list contains repeatedly the last element of the corresponding inner list from the input. For example, $mult-lasts([[a,b],[c,d,e],[f]]) = [[b,b],[e,e,e],[f]]$. All three induced equations are recursive. The third equation computes a one element list containing the last element of an input list. The second equation utilizes the third equation and returns a list of the same structure as a given input list where the elements of the input list are replaced by the last element. The first equation utilizes the second equation to compute the inner lists. Finally $reverse$ reverses a list. The induced program has an unusual form, nevertheless it is correct.

6. Conclusion and Further Research

We presented an EBG approach to inducing sets of recursive equations representing functional programs from I/O-examples. The underlying methodologies are inspired by classical approaches to induction of functional Lisp-programs, particularly by the approach of Summers (1977). The presented approach goes in three main aspects beyond Summers' approach: Sets of recursive equations can be induced at once instead of only one recursive equation, each equation may contain more than one recursive call, and additionally needed parameters are introduced systematically. We have implemented prototypes for both steps. The generalizer works domain-independent and all problems which comply to our general program scheme (Def. 1) with the restrictions described in Sec. 3.3 can be solved, whereas construction of initial terms as described in Sec. 4 relies on knowledge of datatypes.

We are investigating several extensions for the first synthesis step: First, we try to integrate knowledge about further datatypes such as trees and natural numbers. For example, we believe, that if we introduce $zero$ and $succ$, denoting the natural number 0 and the successor function resp. as constructors for natural numbers, $prev$ for "decomposing" natural numbers and the predicate $zerop$ as bottom test on natural numbers, then it should be possible to induce a program returning the length of a list for example. Another extension will be to allow for more than one input parameter in the I/O-examples, such that $append$ becomes inducible for example. A third extension should be the ability to utilize user-defined or in a previous step induced functions within an induction step.

Until now our approach suffers from the restriction to structural problems due to the principal approach to calculate

traces deterministically without search in the first synthesis step. We work on overcoming this restriction, i.e. on extending the first synthesis step to the ability of dealing with problems which are not (only) structural, list sorting for example. A strong extension to the second step would be the ability to deal with nested recursive calls, yet this would imply a much more complex structural analysis on the initial terms.

References

- Biermann, A. W., Guiho, G., & Kodratoff, Y. (Eds.). (1984). *Automatic program construction techniques*. Collier Macmillan.
- Dershowitz, N., & Jouanaud, J.-P. (1990). Rewrite systems. In J. Leeuwen (Ed.), *Handbook of theoretical computer science*, vol. B. Elsevier.
- Flener, P., & Partridge, D. (2001). Inductive programming. *Autom. Softw. Eng.*, 8, 131–137.
- Flener, P., & Yilmaz, S. (1999). Inductive synthesis of recursive logic programs: Achievements and prospects. *Journal of Logic Programming*, 41, 141–195.
- Kitzelmann, E. (2003). Inductive functional program synthesis – a term-construction and folding approach. Master's thesis, Dept. of Computer Science, TU Berlin. <http://www.cogsys.wiai.uni-bamberg.de/kitzelmann/documents/thesis.ps>.
- Lowry, M. L., & McCarthy, R. D. (1991). *Automatic software design*. Cambridge, Mass.: MIT Press.
- Muggleton, S., & De Raedt, L. (1994). Inductive logic programming: Theory and methods. *Journal of Logic Programming, Special Issue on 10 Years of Logic Programming*, 19-20, 629–679.
- Olsson, R. (1995). Inductive functional programming using incremental program transformation. *Artificial Intelligence*, 74, 55–8.
- Plotkin, G. D. (1969). A note on inductive generalization. In *Machine intelligence*, vol. 5, 153–163. Edinburgh University Press.
- Rao, M. R. K. K. (2004). Inductive inference of term rewriting systems from positive data. *ALT* (pp. 69–82).
- Schmid, U., & Wysotzki, F. (2000). Applying inductive program synthesis to macro learning. *Proc. 5th International Conference on Artificial Intelligence Planning and Scheduling (AIPS 2000)* (pp. 371–378). AAAI Press.
- Summers, P. D. (1977). A methodology for LISP program construction from examples. *Journal ACM*, 24, 162–175.
- Wysotzki, F., & Schmid, U. (2001). *Synthesis of recursive programs from finite examples by detection of macro-functions* (Technical Report 01-2). Dept. of Computer Science, TU Berlin, Germany.