

SRI International



DEDUCTIVE SYNTHESIS OF THE UNIFICATION ALGORITHM

Technical Note 246

July 1981


By: Zohar Manna
Computer Science Department
Stanford University
and
Applied Mathematics Department
Weizmann Institute

Richard Waldinger
Artificial Intelligence Center
SRI International

This paper will appear in *Automatic Program Construction* (G. Guiho, ed.), NATO Scientific Series, D. Reidel Pub. Co., Dordrecht, Holland, 1981.

The research was supported in part by the National Science Foundation under Grants MCS-78-02591 and MCS-79-09495, in part by the Office of Naval Research under Contracts N00014-75-C-0816 and N00014-76-C-0687, and in part by the Air Force Office of Scientific Research under Contract AFOSR-81-0014.

333 Ravenswood Ave. • Menlo Park, California 94025
(415) 326-6200 • Cable: SRI INTL MPK • TWX: 910-373-1246



ABSTRACT

The *deductive approach* is a formal program construction method in which the derivation of a program from a given specification is regarded as a theorem-proving task. To construct a program whose output satisfies the conditions of the specification, we prove a theorem stating the existence of such an output. The proof is restricted to be sufficiently constructive so that a program computing the desired output can be extracted directly from the proof. The program we obtain is applicative and may consist of several mutually recursive procedures. The proof constitutes a demonstration of the correctness of this program.

To exhibit the full power of the deductive approach, we apply it to a nontrivial example — the synthesis of a *unification algorithm*. Unification is the process of finding a common instance of two expressions. Algorithms to perform unification have been central to many theorem-proving systems and to some programming-language processors.

The task of deriving a unification algorithm automatically is beyond the power of existing program synthesis systems. In this paper we use the deductive approach to derive an algorithm from a simple, high-level specification of the unification task. We will identify some of the capabilities required of a theorem-proving system to perform this derivation automatically.

INTRODUCTION

In an earlier paper (Manna and Waldinger [1980]) we described a deductive approach to program synthesis. In this approach, program synthesis is regarded as a theorem-proving task. Given a high-level specification of the purpose of the program, we prove a theorem that establishes the existence of an output satisfying this specification. The proof is restricted to be sufficiently constructive so that the desired program can be extracted directly. This approach is the direct descendant of the technique applied, *e.g.*, by Green [1969] and by Waldinger and Lee [1969].

In the earlier paper, we applied the technique only to very simple examples. In this paper we consider a somewhat more difficult task: the synthesis of a unification algorithm.

Unification is the process of finding a common instance of two expressions. If such an instance exists, the algorithm is to produce a substitution that will yield that instance when applied to either of the expressions. If no common instance exists, the algorithm is to produce a special symbol indicating this situation. The first unification algorithm appeared in Herbrand's [1930] thesis, but the procedure did not come to widespread attention until it was rediscovered by Prawitz [1960] and employed by Robinson [1965] in his resolution principle for automatic theorem-proving. Since then, the algorithm has been used not only for resolution theorem-proving but also in many nonresolution theorem provers (see Bledsoe [1977]) and some programming-language processors (*e.g.*, PLANNER, see Hewitt [1971] or PROLOG, see Warren *et al.* [1977] or Colmerauer *et al.* [1979]).

Because of its importance in theorem-proving and other applications, some effort has gone into the design of efficient unification algorithms (*e.g.*, Martelli and Montanari [1976]; Paterson and Wegman [1978]) and the extension of the algorithm to more complex logical theories (*e.g.*, higher-order logic, Huet [1975]; associative and commutative theories, Stickel [1975], Livesay *et al.* [1979]).

The unification algorithm was the subject of verification efforts (*e.g.*, Waldinger and Levitt [1974]) and served as an example of automatic debugging (von Henke and Luckham [1974]). An early attempt to synthesize such a program appeared in Manna and Waldinger [1975]. Nevertheless, no complete automatic synthesis, or even verification, of the algorithm has been accomplished by any system.

The derivation presented in this paper depends on the formulation of a theory of expressions and substitutions. Intuitive observations about these objects can then be expressed and proved within the theory. In this paper we have set down, without proof, those results necessary for the derivation. A full presentation of the theory of expressions and substitutions is included in our forthcoming book (Manna and Waldinger [1982]).

The proof on which the derivation is based is presented in full. We include a summary of those aspects of the deductive approach necessary to understand the derivation. Although this proof can be expressed in the deductive-tableau formalism of our [1980] paper, it is given here informally. We do not attempt to describe strategies that would enable the proof to be generated automatically.

However, we consider later what capabilities would be required of a theorem-proving system to discover such a proof.

THE DEDUCTIVE APPROACH

The specification of a program allows us to express the purpose of the desired program without indicating an algorithm that achieves this purpose. In general, we are considering the synthesis of programs whose specifications have the form

$$f(a) \Leftarrow \text{find } z \text{ such that } R(a, z) \\ \text{where } P(a).$$

Here a denotes the input of the desired program and z denotes its output. The *input condition* $P(a)$ expresses the class of legal inputs to which the program is expected to apply. The *output condition* $R(a, z)$ describes the relation the output z is intended to satisfy.

For example, to specify a program to compute the integer square root of a nonnegative integer n , we would write

$$\text{sqrt}(n) \Leftarrow \text{find } z \text{ such that} \\ \text{integer}(z) \text{ and } z^2 \leq n < (z + 1)^2 \\ \text{where } \text{integer}(n) \text{ and } 0 \leq n.$$

A specification of the above form describes an applicative program, one that yields an output but produces no side effects. To derive a program from such a specification, we attempt to prove a theorem of the form

$$(\forall a)(\exists z)[\text{if } P(a) \text{ then } R(a, z)].$$

This theorem states that, for every input a , there exists an output z satisfying the output condition, provided that the input satisfies the input condition. The proof of this theorem must be constructive, in the sense that, in proving the existence of a satisfactory output z , it must tell us how to find such an output. From this proof a program to compute z can be extracted.

WELL-FOUNDED INDUCTION

The formation of repetitive program constructs in the deductive approach depends on the application of the principle of mathematical induction. The induction principle we use is the principle of "well-founded induction," which applies to a wide variety of mathematical structures and results in the formation of a recursive procedure in the program being constructed. Before we can present the induction principle, we must introduce the notion of a "well-founded ordering."

Definition: If \succ is a relation over a set S , we will say that \succ satisfies the *decreasing sequence condition* if there are no infinite decreasing sequences x_1, x_2, x_3, \dots of elements of S ; i.e.,

there are no sequences such that

$$x_1 \succ x_2 \succ x_3 \succ \dots$$

For example, the ordinary “greater-than” relation \succ over the nonnegative integers satisfies the decreasing-sequence condition. The same relation over all the integers does not.

Definition: If \succ is a relation over a set S , we will say that \succ is a *well-founded ordering* and S is a well-founded set under this ordering if

- \succ is transitive and
- \succ satisfies the decreasing-sequence condition.

We will regard $y \prec x$ as synonymous with $x \succ y$.

For example, the following are all well-founded orderings:

- The $>$ relation over the nonnegative integers
- The subset relation over the finite sets
- The subtree relation over the finite trees.

The *principle of well-founded induction* may be expressed as follows:

Let \succ be a well-founded ordering over a set S . Then, to prove that a statement $Q(a)$ is true for all elements a of S , consider an arbitrary element a of S , assume the induction hypothesis

$$(\forall x)[\text{if } x \prec a \text{ then } Q(x)]$$

and prove that the conclusion

$$Q(a)$$

then follows.

When we are applying the well-founded induction principle in a program-synthesis context, we may use the following special form:

Let \succ be a well-founded ordering over a set S . To construct a program f satisfying the specification

$$f(a) \Leftarrow \text{find } z \text{ such that } R(a, z)$$

where $P(a)$,

where the inputs a belong to a well-founded set, consider an arbitrary input a , assume the induction hypothesis

$$(\forall x) \left[\begin{array}{l} \text{if } x \prec a \\ \text{then if } P(x) \\ \text{then } R(x, f(x)) \end{array} \right]$$

and then prove the conclusion

$$(\exists z)R(a, z).$$

In other words, we consider an arbitrary input a and find an output z satisfying the given specification under the following induction hypothesis: the program $f(x)$ we are trying to construct will satisfy the specification for all inputs x that are less than a in the well-founded ordering.

Application of the induction hypothesis during a proof will cause a recursive call $f(x)$ to appear in the program being constructed. The condition $P(x)$ will ensure that the input x of the recursive call will be a legal input; *i.e.*, it will satisfy the given input condition. The condition $x \prec a$ will ensure that the new recursive call cannot result in an infinite computation.

ℓ -EXPRESSIONS

In this section, we define a class of ℓ -expressions that will contain not only the expressions, but also nested lists of expressions formed from a given alphabet.

- *The alphabet*

Suppose that S is an alphabet of symbols, consisting of three disjoint sets:

C : the *constants*

X : the *variables*

F : the *function symbols*.

Together, the constants and variables will be referred to as the *atoms* of S . With each function symbol of F is associated a unique positive integer, called its *arity*, indicating how many arguments the function takes.

- *Generation rules*

The *expressions* of S are constructed by repeated application of the following *generation rules*:

Any constant of C is an expression;

Any variable of X is an expression;

If f is a function symbol (of arity n)
 and ℓ is a list of expressions (of length n),
 then the result of applying f to the expressions in ℓ ,
 denoted by $f \bullet \ell$, is an expression.

Note that if $\ell = [\ell_1, \ell_2, \dots, \ell_n]$ then $f \bullet \ell$ is the expression informally denoted by $f(\ell_1, \ell_2, \dots, \ell_n)$.

The ℓ -expressions of S are constructed by repeated application of the following generation rules:

The empty list $[\]$ is an ℓ -expression;

Any expression is an ℓ -expression;

If s is an ℓ -expression

and m is a list of ℓ -expressions

then the result of inserting s before the first element of m ,
 denoted by $s \circ m$, is an ℓ -expression.

Note that if $m = [m_1, m_2, \dots, m_n]$, then $s \circ m$ is the list of ℓ -expressions informally denoted by $[s, m_1, m_2, \dots, m_n]$.

• *Uniqueness properties*

We assume that each ℓ -expression can only be produced in a unique way from the above rules. This assumption is expressed by the following properties:

$$c \neq x$$

$$c \neq f \bullet \ell$$

$$x \neq f \bullet \ell$$

$$\text{if } f \bullet \ell = f' \bullet \ell'$$

$$\text{then } f = f' \text{ and } \ell = \ell'$$

$$c \neq [\]$$

$$c \neq s \circ m$$

$$x \neq [\]$$

$$x \neq s \circ m$$

$$f \bullet \ell \neq [\]$$

$$f \bullet \ell \neq s \circ m$$

$$[] \neq s \circ m$$

$$\begin{aligned} &\text{if } s \circ m = s' \circ m' \\ &\text{then } s = s' \text{ and } m = m' \end{aligned}$$

for all constants c , variables x , function symbols f and f' (of arity n and n' , respectively), lists of expressions ℓ and ℓ' (of length n and n' , respectively), ℓ -expressions s and s' , and lists of ℓ -expressions m and m' .

• *The occurs-in relation*

We will say that an ℓ -expression s *occurs in* an ℓ -expression s' , denoted by

$$s \preceq s' \text{ or } s' \succeq s,$$

if s is a subexpression of s' ; we will say that s *occurs properly in* s' , denoted by

$$s \prec s' \text{ or } s' \succ s$$

if s occurs in s' but is distinct from s' . We will abbreviate

$$\text{not } s \prec s' \text{ as } s \not\prec s',$$

$$\text{not } s \preceq s' \text{ as } s \not\preceq s',$$

and so forth. Formally these relations are defined by the following properties:

$$s \preceq s' \text{ if and only if } s \prec s' \text{ or } s = s' \quad (\text{partiality})$$

$$s \not\prec a \quad (\text{atom})$$

$$s \not\prec [] \quad (\text{empty})$$

$$s \prec (f \bullet \ell) \text{ if and only if } s \preceq \ell \quad (\text{application})$$

$$s \prec (s' \circ m) \text{ if and only if } s \preceq s' \text{ or } s \preceq m \quad (\text{insertion})$$

for all atoms a , function symbols f (of arity n), lists of expressions ℓ (of length n), ℓ -expressions s and s' , and lists of ℓ -expressions m .

We assume as part of its definition that the occurs-in relation is well-founded. This is a way of expressing formally that all the ℓ -expressions are finite. It follows that the relation is irreflexive, i.e.,

$$s \not\prec s \quad (\text{irreflexivity})$$

for all ℓ -expressions s .

The definition implies the following *component* properties of \prec :

$$\ell \prec f \bullet \ell$$

$$s \prec s \circ m$$

$$m \prec s \circ m$$

for every function symbol f (of arity n), list of expressions ℓ (of length n), ℓ -expression s , and list of ℓ -expressions m .

• *The vars function*

The value of $vars(s)$ is the set of variables that occur in the ℓ -expression s . Formally we define

$$vars(c) = \{\} \quad (\text{constant})$$

$$vars(x) = \{x\} \quad (\text{variable})$$

$$vars([\]) = \{\} \quad (\text{empty})$$

$$vars(f \bullet \ell) = vars(\ell) \quad (\text{application})$$

$$vars(s \circ m) = vars(s) \cup vars(m) \quad (\text{insertion})$$

for all constants c , variables x , function symbols f (of arity n), lists of expressions ℓ (of length n), ℓ -expressions s , and lists of ℓ -expressions m .

Now let us state a proposition relating the occurs-in relation with the $vars$ function.

Proposition (variables): For every variable y and ℓ -expression s ,

$$y \in vars(s) \text{ if and only if } y \preceq s.$$

In other words, the elements of $vars(s)$ are indeed those variables that occur in s . The proof is by well-founded induction over the occurs-in relation itself.

SUBSTITUTIONS

Substitution is the operation of replacing certain variables of an ℓ -expression by other expressions. We begin by giving an informal exposition of substitutions; subsequently we give a formal treatment of the same notion.

Informally we will represent a substitution θ as a set of *replacements*

$$\theta = \{x_1 \leftarrow e_1, x_2 \leftarrow e_2, \dots, x_n \leftarrow e_n\},$$

where x_1, x_2, \dots, x_n are distinct variables in X and e_1, e_2, \dots, e_n are expressions such that $e_i \neq x_i$. Thus, replacements of the form $x \leftarrow x$ are excluded from substitutions, and substitutions of the form $\{x \leftarrow e, x \leftarrow e'\}$ are not allowed.

If $x_i \leftarrow e_i$ is a replacement in the substitution θ , we will refer to x_i as the *variable* and e_i as the *expression* of the replacement. We will denote by

$dom(\theta)$: the set of variables $\{x_1, x_2, \dots, x_n\}$
affected by the substitution θ ,

called the *domain* of θ , and

$range(\theta)$: the set of variables that occur in e_1, e_2, \dots , or e_n ,

called the *range* of θ .

The result $s \leftarrow \theta$ of applying such a substitution θ to an ℓ -expression s is obtained by simultaneously replacing every instance of the variables x_1, x_2, \dots , and x_n by the corresponding expressions e_1, e_2, \dots , and e_n .

For example, if θ is the substitution

$\{x \leftarrow f(y), y \leftarrow g(a, z)\}$,

then

$dom(\theta) = \{x, y\}$,

$range(\theta) = \{y, z\}$;

furthermore, if s is the ℓ -expression

$[x, [g(a, x), y]]$,

then the result $s \leftarrow \theta$ of applying θ to s is

$[f(y), [g(a, f(y)), g(a, z)]]$.

Note that the replacements are performed simultaneously; thus, the variable y in $f(y)$ above was *not* replaced by $g(a, z)$ even though θ contains a replacement $y \leftarrow g(a, z)$.

Let us be more precise. Suppose that S is an alphabet consisting of the constants C , the variables X , and the function symbols F , as before.

• *Generation rules*

The *substitutions* are constructed by repeated application of the following *generation rules*. We define the domain and range sets for each substitution at the same time.

- The empty substitution $\{ \}$ is a substitution,

- $dom(\{\}) = \{\}$,
- $range(\{\}) = \{\}$;

If θ is a substitution,
 x is a variable not in $dom(\theta)$, and
 e is an expression distinct from x ,
then

- the result of adding the replacement $x \leftarrow e$ to the substitution θ , denoted by $(x \leftarrow e) \circ \theta$, is also a substitution,
- $dom((x \leftarrow e) \circ \theta) = \{x\} \cup dom(\theta)$, and
- $range((x \leftarrow e) \circ \theta) = vars(e) \cup range(\theta)$.

Note that if $\theta = \{x_1 \leftarrow e_1, x_2 \leftarrow e_2, \dots, x_n \leftarrow e_n\}$ then $(x \leftarrow e) \circ \theta$ is the substitution informally denoted by $\{x \leftarrow e, x_1 \leftarrow e_1, \dots, x_n \leftarrow e_n\}$. Furthermore,

$$dom((x \leftarrow e) \circ \theta) = \{x, x_1, \dots, x_n\}$$

$$range((x \leftarrow e) \circ \theta) = vars(e) \cup vars(e_1) \cup \dots \cup vars(e_n).$$

We call \circ the *addition* function for substitutions.

Substitutions do not have uniqueness properties; we may regard two substitutions as equal even though they have been constructed in different ways. We will say that two substitutions are equal if they have the same effect when applied to an arbitrary ℓ -expression. But first let us define more precisely what we mean by applying a substitution to an ℓ -expression.

• *The apply function*

If θ is a substitution and s is an ℓ -expression, then the *apply* function $s \triangleleft \theta$ is defined to satisfy the following properties:

$$s \triangleleft \{\} = s \quad (\text{empty substitution})$$

$$c \triangleleft \theta = c \quad (\text{constant})$$

$$x \triangleleft ((x \leftarrow e) \circ \theta) = e \quad (\text{same variable})$$

$$y \triangleleft ((x \leftarrow e) \circ \theta) = y \triangleleft \theta \quad \text{if } x \neq y \quad (\text{distinct variable})$$

$$[] \triangleleft \theta = [] \quad (\text{empty list})$$

$$(f \circ \ell) \triangleleft \theta = f \circ (\ell \triangleleft \theta) \quad (\text{application})$$

$$(s \circ m) \triangleleft \theta = (s \triangleleft \theta) \circ (m \triangleleft \theta) \quad (\text{insertion})$$

for all substitutions θ , constants c , variables x and y , expressions e , function symbols f (of arity n), lists of expressions ℓ (of length n), ℓ -expressions s , and lists of ℓ -expressions m .

Note that, in the *same-variable* and *distinct-variable* properties, we do not require that x and e be distinct or that $x \notin \text{dom}(\theta)$ even though these are conditions on the generation rules for substitutions. This means that $(x \leftarrow e) \circ \theta$ is defined in these cases.

Let us now introduce a simple property relating ℓ -expressions and substitutions.

Proposition (monotonicity): For all ℓ -expressions s and s' , and substitutions θ ,

$$(a) \quad \text{if } s \prec s' \text{ then } s \triangleleft \theta \prec s' \triangleleft \theta,$$

$$(b) \quad \text{if } s \preceq s' \text{ then } s \triangleleft \theta \preceq s' \triangleleft \theta.$$

In other words, the subexpression and proper subexpression relations are maintained after the application of a substitution.

- *Agreement and equality*

We will say that two substitutions θ_1 and θ_2 *agree* on an ℓ -expression s if

$$s \triangleleft \theta_1 = s \triangleleft \theta_2.$$

Now we can define the notion of equality for substitutions. We will say that two substitutions θ_1 and θ_2 are *equal*, denoted by $\theta_1 = \theta_2$, if they agree on all ℓ -expressions; i.e.,

$$\theta_1 = \theta_2$$

if and only if

$$s \triangleleft \theta_1 = s \triangleleft \theta_2 \quad \text{for all } \ell\text{-expressions } s.$$

We assume that the only substitutions are those that have been constructed by a finite number of applications of the generation rules; it follows that any substitution θ is a finite set of replacements and that the sets $\text{dom}(\theta)$ and $\text{range}(\theta)$ are finite sets of variables. This finiteness can be expressed formally by an appropriate induction principle over the substitutions.

- *Characterization of domain and range*

Let us state two propositions that characterize the domain and range of a substitution.

Proposition (domain): For every substitution θ and variable x ,

$$x \in \text{dom}(\theta) \quad \text{if and only if} \quad x \triangleleft \theta \neq x.$$

That is, the domain is the set of all variables affected by the substitution.

Proposition (range): For every substitution θ and variable y ,

$$y \in \text{range}(\theta) \text{ if and only if } \left[\begin{array}{l} \text{there exists a variable } x \text{ such that} \\ x \in \text{dom}(\theta) \text{ and } y \in \text{vars}(x \leftarrow \theta) \end{array} \right].$$

That is, the range is the set of all variables that may be introduced by a substitution.

Let us now introduce the notion of the composition of two substitutions.

COMPOSITION OF SUBSTITUTION

We define the composition $\theta \diamond \theta'$ of two substitutions θ and θ' to be the substitution satisfying the following property:

$$s \leftarrow (\theta \diamond \theta') = (s \leftarrow \theta) \leftarrow \theta'$$

for all ℓ -expressions s . In other words, applying the composition $\theta \diamond \theta'$ to an ℓ -expression is the same as applying θ first, and then applying θ' to the result.

For example, if

$$\theta = \{x \leftarrow f(y)\}$$

$$\theta' = \{y \leftarrow g(a, x), x \leftarrow b\}$$

and

$$s = h(x, y, z),$$

then

$$s \leftarrow \theta = h(f(y), y, z)$$

and

$$s \leftarrow (\theta \diamond \theta') = (s \leftarrow \theta) \leftarrow \theta' = h(f(g(a, x)), g(a, x), z).$$

It follows from the definition that composition has the following properties:

$$\theta \diamond \{\} = \theta \quad (\text{right empty})$$

$$\{\} \diamond \theta = \theta \quad (\text{left empty})$$

$$(\theta_1 \diamond \theta_2) \diamond \theta_3 = \theta_1 \diamond (\theta_2 \diamond \theta_3) \quad (\text{associativity})$$

for any substitutions θ , θ_1 , θ_2 , and θ_3 . Because \diamond is associative, we may write expressions of form $(\theta_1 \diamond \theta_2) \diamond \theta_3$ and $\theta_1 \diamond (\theta_2 \diamond \theta_3)$ as $\theta_1 \diamond \theta_2 \diamond \theta_3$ without fear of ambiguity.

SYNTACTIC CATEGORIES

We will regard constants, variables, expressions, and lists as being of distinct "syntactic categories."

Definition: Let S be an arbitrary alphabet. Then the *syntactic categories* of the ℓ -expressions of S are the following five sets:

- The set C of constants.
- The set X of variables.
- The set consisting of the empty list $[\]$.
- The set of functional expressions of form $f \bullet \ell$,
where f is a function symbol (of arity n)
and ℓ is a list of expressions (of length n).
- The set of nonempty lists of form $t \circ m$,
where t is an ℓ -expression and m is a list of ℓ -expressions.

By the uniqueness properties of ℓ -expressions, every ℓ -expression belongs to precisely one syntactic category. The power of a substitution to change the syntactic category of an ℓ -expression is severely limited by the following observation:

Proposition (syntactic category): For any substitution θ and ℓ -expression s ,

- (a) if s is not a variable
then s and $s \blacktriangleleft \theta$ are in the same syntactic category.
- (b) if e is an expression
then $e \blacktriangleleft \theta$ is an expression.

Note that it is not necessarily true that, if s is a variable, then $s \blacktriangleleft \theta$ is also a variable; it may be a constant or a functional expression. However, the converse is true:

Corollary: For any substitution θ and ℓ -expression s ,

- if $s \blacktriangleleft \theta$ is a variable
then s is also a variable.

THE AGREEMENT PROPOSITION AND ITS CONSEQUENCES

The following proposition has many useful consequences:

Proposition (agreement): For all substitutions θ and θ' and ℓ -expressions s , we have

$$s \triangleleft \theta = s \triangleleft \theta'.$$

if and only if

$$\text{for every } x, \text{ if } x \in \text{vars}(s) \text{ then } x \triangleleft \theta = x \triangleleft \theta'.$$

In other words, two substitutions agree on an ℓ -expression precisely when they agree on all the variables of the ℓ -expression.

An immediate consequence of this proposition indicates what happens when all the variables of an ℓ -expression remain unchanged by a substitution:

Corollary (invariance): For every substitution θ and ℓ -expression s ,

$$s \triangleleft \theta = s$$

if and only if

$$\text{vars}(s) \cap \text{dom}(\theta) = \{\}.$$

In other words, a substitution has no effect on an ℓ -expression precisely when no variable in the domain of the substitution actually occurs in the ℓ -expression. The proof depends on taking θ' to be the empty substitution $\{\}$ in the agreement proposition.

Corollary (replacement invariance): For every ℓ -expression s , expression e , and variable x ,

$$\begin{aligned} &\text{if } x \notin \text{vars}(s) \\ &\text{then } s \triangleleft \{x \leftarrow e\} = s. \end{aligned}$$

In other words, applying a single replacement to an ℓ -expression has no effect if the variable of the replacement does not occur in the ℓ -expression.

Another consequence of the agreement proposition provides a useful characterization of the equality between substitutions. We have defined two substitutions to be equal if they agree on all ℓ -expressions. In fact it suffices to show that they agree on all variables:

Proposition (equality): For all substitutions θ and θ' , we have

$$\theta = \theta'$$

if and only if

for every variable x , $x \triangleleft \theta = x \triangleleft \theta'$.

According to this proposition, to prove equality between substitutions it suffices to show that they agree on all variables. In fact, it suffices to consider only variables in their domains:

Corollary (equality): For all substitutions θ and θ' , we have

$$\left[\begin{array}{l} \text{for every variable } x, \\ \text{if } x \in \text{dom}(\theta) \cup \text{dom}(\theta') \\ \text{then } x \triangleleft \theta = x \triangleleft \theta' \end{array} \right] \text{ if and only if } \theta = \theta'.$$

The following proposition relates the addition and composition functions:

Proposition (addition composition): For every substitution θ , variable x , and expression e ,

$$(x \leftarrow e \triangleleft \theta) \circ \theta = (x \leftarrow e) \diamond \theta.$$

The proof relies on the equality proposition.

• *The subtraction function*

We denote by $\theta - x$ the substitution that has no effect on the variable x , but agrees with θ on all other variables. Formally, we define $\theta - x$ by the following properties:

$$\begin{aligned} x \triangleleft (\theta - x) &= x \\ y \triangleleft (\theta - x) &= y \triangleleft \theta \quad \text{if } x \neq y \end{aligned}$$

for all substitutions θ and variables x and y . It follows (by the domain proposition) that

$$x \notin \text{dom}(\theta - x)$$

for all substitutions θ and variables x and y .

Proposition (subtraction): For any variable x , expression e , and substitution θ , where $x \not\triangleleft e$, i.e., $x \notin \text{vars}(e)$, we have

$$e \triangleleft (\theta - x) := e \triangleleft \theta.$$

The following proposition enables us to break a substitution down into its component replacements:

Proposition (decomposition): For every substitution θ and variable x ,

$$\theta = (x \leftarrow x \triangleleft \theta) \circ (\theta - x).$$

SET-THEORETIC PROPERTIES OF SUBSTITUTIONS

In this section, we give some properties that relate the domain and range of a substitution with the variables of the ℓ -expressions.

Proposition (variable elimination): For any ℓ -expression s , substitution θ , and variable y ,

if $y \in \text{dom}(\theta)$ and $y \notin \text{range}(\theta)$
then $y \notin \text{vars}(s \triangleleft \theta)$.

Thus, if a variable occurs in the domain of a substitution but not in its range, then the substitution will remove that variable from any ℓ -expression in which it occurs.

Proposition (variable introduction): For any ℓ -expression s , substitution θ , and variable y ,

if $y \in \text{vars}(s \triangleleft \theta)$
then $y \in \text{range}(\theta)$ or $y \in \text{vars}(s)$,

i.e.,

$$\text{vars}(s \triangleleft \theta) \subseteq \text{range}(\theta) \cup \text{vars}(s).$$

In other words, if a variable occurs in an ℓ -expression after a substitution, then it was introduced by the substitution or it occurred in the ℓ -expression originally.

UNIFIERS

Suppose that s and s' are ℓ -expressions and θ is a substitution of some alphabet S of constants, variables, and function symbols. We will say that θ is a *unifier* of s and s' if

$$s \triangleleft \theta = s' \triangleleft \theta.$$

Example: If

$$s = g(x, z)$$

and

$$s' = g(y, f(y))$$

then the substitution

$$\theta = \{x \leftarrow y, z \leftarrow f(y)\}$$

is a unifier, because

$$s \triangleleft \theta = s' \triangleleft \theta = g(y, f(y)).$$

Note that θ is not the only unifier of s and s' ; e.g.,

$$\lambda = \{x \leftarrow b, y \leftarrow b, z \leftarrow f(b)\}$$

is a unifier, because

$$s \triangleleft \lambda = s' \triangleleft \lambda = g(b, f(b));$$

also,

$$\rho = \{y \leftarrow x, z \leftarrow f(x)\}$$

is a unifier, because,

$$s \triangleleft \rho = s' \triangleleft \rho = g(x, f(x)). \quad \blacksquare$$

Not every pair of ℓ -expressions has a unifier. For example, there is no unifier for

$$s = g(a, b)$$

and

$$s' = g(x, x).$$

For the result of applying any substitution θ to $g(a, b)$ will be the expression

$$g(a, b)$$

itself, in which the arguments are distinct. On the other hand, the result of applying θ to $g(x, x)$ will be an expression of form

$$g(e, e),$$

in which the arguments are identical. These two results can never be the same.

Two ℓ -expressions will be said to be *unifiable* if they have (at least one) unifier. Thus, $g(x, z)$ and $g(y, f(y))$ are unifiable, but $g(a, b)$ and $g(x, x)$ are not.

The following propositions characterize the unifiers for different categories of ℓ -expressions.

Proposition (application unifier): For any function symbol f (of arity n) and lists ℓ and ℓ' (of length n), we have

$$\lambda \text{ is a unifier of } f \bullet \ell \text{ and } f \bullet \ell'$$

if and only if

λ is a unifier of ℓ and ℓ' .

Proposition (insertion unifier): For any ℓ -expressions t and t' and lists of expressions m and m' ,

λ is a unifier of $t \circ m$ and $t' \circ m'$

if and only if

λ is a unifier of t and t'

and

λ is a unifier of m and m' .

GENERALNESS

We will say that a substitution θ is *more general than* a substitution λ , denoted by $\theta \succeq_{gen} \lambda$ or $\lambda \preceq_{gen} \theta$, if there exists some substitution ψ such that

$$\lambda = \theta \diamond \psi.$$

In this case, we will also say that λ is an *instance* of θ .

Example:

The substitution

$$\theta = \{y \leftarrow b\}$$

is more general than the substitution

$$\lambda = \{x \leftarrow a, y \leftarrow b\},$$

i.e., $\theta \succeq_{gen} \lambda$, because (taking ψ to be $\{x \leftarrow a\}$ in the definition)

$$\lambda = \theta \diamond \{x \leftarrow a\}.$$

In other words, λ is an instance of θ . \square

Note that, according to our terminology, a substitution is always more general than itself, i.e.,

$$\theta \succeq_{gen} \theta \quad (\text{reflexivity})$$

for every substitution θ . Also, the empty substitution is more general than any substitution, i.e.,

$$\{\} \succeq_{gen} \theta \quad (\text{empty})$$

for every substitution θ .

MOST-GENERAL UNIFIERS

We have observed that there may be several distinct unifiers for a given pair of ℓ -expressions. In fact, any instance of a unifier is also a unifier.

Proposition (instance of a unifier): For all substitutions θ and λ and ℓ -expressions s and s' ,

if θ is a unifier of s and s'
 and $\theta \succeq_{gen} \lambda$
 then λ is a unifier of s and s' .

Definition: A substitution θ is a *most-general unifier* of two ℓ -expressions s and s' if

- θ is a unifier of s and s' ; i.e.,
 $s \triangleleft \theta = s' \triangleleft \theta$,

and

- θ is more general than any other unifier of s and s' ; i.e.,
 if $s \triangleleft \lambda = s' \triangleleft \lambda$
 then $\theta \succeq_{gen} \lambda$
 for any substitution λ .

Combining the above proposition and definition, we see that the unifiers of two ℓ -expressions are precisely the instances of a most-general unifier.

Corollary (most-general unifier): For all substitutions θ and λ and ℓ -expressions s and s' ,

θ is a most general unifier of s and s'

if and only if

$$\left[\begin{array}{l} \lambda \text{ is a unifier of } s \text{ and } s' \\ \text{if and only if} \\ \theta \succeq_{gen} \lambda \end{array} \right]$$

Example:

We have seen that the two ℓ -expressions

$$s = g(x, z)$$

and

$$s' = g(y, f(y))$$

have many unifiers; e.g.,

$$\theta = \{x \leftarrow y, z \leftarrow f(y)\},$$

$$\lambda = \{x \leftarrow b, y \leftarrow b, z \leftarrow f(b)\}, \text{ and}$$

$$\rho = \{y \leftarrow x, z \leftarrow f(x)\}.$$

It turns out that θ is a most-general unifier of s and s' . In particular, θ is more general than λ and ρ , because

$$\lambda = \theta \diamond \{y \leftarrow b\}$$

and

$$\rho = \theta \diamond \{y \leftarrow x\}.$$

The unifier λ is not most-general. In particular, λ is not more general than θ ; i.e.,

$$\theta \not\equiv \lambda \diamond \psi$$

for any substitution ψ . For instance,

$$x \leftarrow \theta = y$$

but (by the definition of composition and the *constant* property of the apply function \leftarrow)

$$x \leftarrow (\lambda \diamond \psi) = (x \leftarrow \lambda) \leftarrow \psi = b \leftarrow \psi = b \neq y.$$

Most-general unifiers are not unique; for example, the above substitution ρ is also a most-general unifier of s and s' . In particular, ρ is more general than θ and λ , because

$$\theta = \rho \diamond \{x \leftarrow y\}$$

and

$$\lambda = \rho \diamond \{x \leftarrow b\}.$$

The following proposition concerns the unifiers of a variable and an expression:

Proposition (variable unifier): For any variable x and expression e such that $x \not\equiv e$, we have

$$\{x \leftarrow e\} \text{ is a most-general unifier of } x \text{ and } e.$$

We include the proof of this proposition, because it is not straightforward and because it may be regarded as an integral part of the synthesis of the unification algorithm.

Proof:

It suffices to show (by the most-general unifier corollary) that, for an arbitrary substitution λ ,

$$\begin{aligned} \lambda \text{ is a unifier of } x \text{ and } e \\ \text{if and only if} \\ \{x \leftarrow e\} \succeq_{gen} \lambda, \end{aligned}$$

i.e. (by the definition of the generalness relation \succeq_{gen}),

$$\begin{aligned} x \triangleleft \lambda = e \triangleleft \lambda \\ \text{if and only if} \\ \lambda = \{x \leftarrow e\} \diamond \lambda^* \text{ for some substitution } \lambda^*. \end{aligned}$$

On the one hand, if

$$\lambda = \{x \leftarrow e\} \diamond \lambda^*$$

for some substitution λ^* , then

$$\begin{aligned} x \triangleleft \lambda &= x \triangleleft (\{x \leftarrow e\} \diamond \lambda^*) \\ &= (x \triangleleft \{x \leftarrow e\}) \triangleleft \lambda^* \\ &\quad \text{by the definition of composition} \\ &= e \triangleleft \lambda^* \\ &\quad \text{by the same-variable property of the apply function } \triangleleft, \end{aligned}$$

and

$$\begin{aligned} e \triangleleft \lambda &= e \triangleleft (\{x \leftarrow e\} \triangleleft \lambda^*) \\ &= (e \triangleleft \{x \leftarrow e\}) \triangleleft \lambda^* \\ &\quad \text{by the definition of composition} \\ &= e \triangleleft \lambda^* \\ &\quad \text{by the replacement invariance corollary, because } x \notin vars(e). \end{aligned}$$

In short,

$$x \triangleleft \lambda = e \triangleleft \lambda.$$

On the other hand, suppose

$$x \triangleleft \lambda = e \triangleleft \lambda.$$

Then

$$\begin{aligned} \lambda &= (x \leftarrow x \triangleleft \lambda) \circ (\lambda - x) && \text{by the decomposition proposition} \\ &= (x \leftarrow e \triangleleft \lambda) \circ (\lambda - x) && \text{by our supposition that } x \triangleleft \lambda = e \triangleleft \lambda \\ &= (x \leftarrow e \triangleleft (\lambda - x)) \circ (\lambda - x) && \text{by the subtraction proposition, because } x \not\triangleleft e \\ &= \{x \leftarrow e\} \diamond (\lambda - x) && \text{by the addition composition proposition.} \end{aligned}$$

Therefore, if λ^* is taken to be $\lambda - x$, we have

$$\lambda = \{x \leftarrow e\} \diamond \lambda^*$$

as we had intended.

This concludes the proof. ■

IDEMPOTENT SUBSTITUTIONS

We will say that a substitution θ is *idempotent* if it has the special property that

$$\theta = \theta \diamond \theta.$$

Example:

The substitution

$$\theta = \{x \leftarrow f(y)\}$$

is idempotent, because

$$\begin{aligned} \theta \diamond \theta &= \{x \leftarrow f(y)\} \diamond \{x \leftarrow f(y)\} \\ &= \{x \leftarrow f(y)\} \end{aligned}$$

$$= \theta.$$

On the other hand,

$$\phi = \{x \leftarrow f(x)\}$$

is not idempotent, because

$$\begin{aligned} \phi \diamond \phi &= \{x \leftarrow f(x)\} \diamond \{x \leftarrow f(x)\} \\ &= \{x \leftarrow f(f(x))\} \\ &\neq \phi. \quad \blacksquare \end{aligned}$$

The property of idempotence is characterized by the following proposition.

Proposition (idempotence): A substitution is idempotent if and only if its domain and range are disjoint; i.e.,

$$\theta = \theta \diamond \theta \text{ if and only if } \text{dom}(\theta) \cap \text{range}(\theta) = \{\},$$

for all substitutions θ .

MOST-GENERAL, IDEMPOTENT UNIFIERS

Let us return to the example at the beginning of this section.

Example:

We have seen that the two ℓ -expressions

$$s = g(x, z)$$

and

$$s' = g(y, f(y))$$

have among their most-general unifiers the substitutions

$$\theta = \{x \leftarrow y, z \leftarrow f(y)\}$$

and

$$\rho = \{y \leftarrow x, z \leftarrow f(x)\}.$$

Both of these substitutions happen to be idempotent, *i.e.*,

$$\theta = \theta \diamond \theta \quad \text{and} \quad \rho = \rho \diamond \rho.$$

However, not all most-general unifiers are idempotent.

For instance, the substitution

$$\phi = \{x \leftarrow z, z \leftarrow f(z), y \leftarrow z\}$$

also turns out to be a most-general unifier of s and s' . It is a unifier, because

$$s \leftarrow \phi = g(z, f(z)) = s' \leftarrow \phi.$$

It is more general than the most-general unifier θ , because

$$\theta = \phi \diamond \{z \leftarrow y\}.$$

But ϕ is not idempotent, because

$$z \in \{x, y, z\} = \text{dom}(\phi)$$

and

$$z \in \{z\} = \text{range}(\phi);$$

therefore

$$\text{dom}(\phi) \cap \text{range}(\phi) = \{z\} \neq \{\},$$

and hence (by the idempotence proposition) ϕ is not idempotent. ■

Most-general, idempotent unifiers have some properties we will find useful.

Proposition (most-general, idempotent unifier): If θ is a unifier of two ℓ -expressions s and s' , then

θ is most-general and idempotent

if and only if

for every unifier λ of s and s' , $\lambda = \theta \diamond \lambda$.

Proposition (domain and range): If θ is a most-general, idempotent unifier of two ℓ -expressions s and s' , then

$$(a) \quad \text{dom}(\theta) \subseteq \text{vars}(s) \cup \text{vars}(s')$$

$$(b) \quad \text{range}(\theta) \subseteq \text{vars}(s) \cup \text{vars}(s').$$

In other words, the only variables that may appear in θ are those that occur in s or s' .

THE UNIFICATION ALGORITHM

A *unification algorithm* is a procedure for finding a most-general, idempotent unifier for two ℓ -expressions, if any unifiers exist at all. Otherwise, it produces a special symbol *nil*, which is assumed to be distinct from any substitution.

The specification for the unification algorithm may be expressed as follows:

$unify(s, s') \Leftarrow$ find θ such that

- θ is a most-general, idempotent unifier of s and s' and
- $\theta \neq nil$
- or
- s and s' are not unifiable and
- $\theta = nil$

for all ℓ -expressions s and s' .

According to the deductive approach, therefore, we must prove the existence of an output θ satisfying this specification; *i.e.*, we must prove the following theorem:

$$(\forall s)(\forall s')(\exists \theta) \left[\begin{array}{l} \theta \text{ is a most-general, idempotent unifier of } s \text{ and } s' \text{ and} \\ \theta \neq nil \\ \text{or} \\ s \text{ and } s' \text{ are not unifiable and} \\ \theta = nil \end{array} \right]$$

In other words,

$$(\forall s)(\forall s')(\exists \theta) \left[\begin{array}{l} s \triangleleft \theta = s' \triangleleft \theta \text{ and} \quad (\theta \text{ is a unifier}) \\ (\forall \lambda)[\text{if } s \triangleleft \lambda = s' \triangleleft \lambda \text{ then } \theta \succeq_{gen} \lambda] \text{ and} \quad (\theta \text{ is most-general}) \\ \theta = \theta \diamond \theta \text{ and} \quad (\theta \text{ is idempotent}) \\ \theta \neq nil \\ \text{or} \\ (\forall \lambda)[s \triangleleft \lambda \neq s' \triangleleft \lambda] \text{ and} \quad (s \text{ and } s' \text{ are not unifiable}) \\ \theta = nil \end{array} \right]$$

This theorem will also establish that, if two ℓ -expressions are unifiable, they have a most-general, idempotent unifier.

Before we give the proof of the theorem, let us look ahead at the program we will ultimately extract from the proof. For clarity, we present the program as a set of properties of the *unify* function; actually the deductive approach will produce the corresponding LISP-like applicative program. Of course, these properties will not be available to us during the proof.

- general

For all ℓ -expressions s and s' :

$$\text{unify}(s, s') = \{\} \quad \text{if } s = s' \quad (\text{same})$$

$$\text{unify}(s, s') = \text{nil} \quad (\text{distinct})$$

if s and s' are nonvariables in distinct syntactic categories.

- constant

For all constants c and c' :

$$\text{unify}(c, c') = \text{nil} \quad \text{if } c \neq c' \quad (\text{distinct})$$

- variable

For all variables x , expressions e , and lists of ℓ -expressions m :

$$\text{unify}(x, e) = \{x \leftarrow e\} \quad \text{if } x \not\prec e \quad (\text{not-in, left})$$

$$\text{unify}(e, x) = \{x \leftarrow e\} \quad (\text{not-in, right})$$

if $x \not\prec e$ and e is not a variable

$$\text{unify}(x, e) = \text{unify}(e, x) = \text{nil} \quad \text{if } x \prec e \quad (\text{in})$$

$$\text{unify}(x, m) = \text{unify}(m, x) = \text{nil} \quad (\text{list})$$

- function

For all function symbols f and f' (of arities n and n' , respectively) and all lists of expressions ℓ and ℓ' (of lengths n and n' , respectively):

$$\text{unify}(f \bullet \ell, f' \bullet \ell') = \text{unify}(\ell, \ell') \quad \text{if } f = f' \quad (\text{same})$$

$$\text{unify}(f \bullet \ell, f' \bullet \ell') = \text{nil} \quad \text{if } f \neq f' \quad (\text{distinct})$$

• list

For all ℓ -expressions t and t' and all lists of ℓ -expressions m and m' :

let $\theta_{hd} = \text{unify}(t, t')$	
if $\theta_{hd} = \text{nil}$	
then $\text{unify}(t \circ m, t' \circ m') = \text{nil}$	(no)
else let $\theta_{tl} = \text{unify}(m \triangleleft \theta_{hd}, m' \triangleleft \theta_{hd})$	
if $\theta_{tl} = \text{nil}$	
then $\text{unify}(t \circ m, t' \circ m') = \text{nil}$	(yes-no)
else $\text{unify}(t \circ m, t' \circ m') = \theta_{hd} \diamond \theta_{tl}$	(yes-yes)

In expressing the list properties, we have used the notation

let $x = a$
 $P(x)$

as an abbreviation for

$P(a)$.

The virtue of this notation is that if $P(x)$ has many instances of x and if a is a lengthy expression, we would be required to rewrite a many times in writing $P(a)$. Thus, without this abbreviation, the final equality above would read

$$\text{unify}(t \circ m, t' \circ m') = \text{unify}(t, t') \diamond \text{unify}(m \triangleleft \text{unify}(t, t'), m' \triangleleft \text{unify}(t, t')).$$

In the *list* properties, for easy reference, we have given separate names for the three equalities. The *no* property corresponds to the case that $\theta_{hd} = \text{nil}$, the *yes-no* property to the case that $\theta_{hd} \neq \text{nil}$ and $\theta_{tl} = \text{nil}$, and the *yes-yes* property to the case that $\theta_{hd} \neq \text{nil}$ and $\theta_{tl} \neq \text{nil}$.

Now let us examine a proof of the theorem to see how the above program can be constructed.

The proof is by well-founded induction over an ordering \prec_{un} between pairs $\langle s, s' \rangle$ of ℓ -expressions. Rather than choose this ordering in advance, we will proceed with the proof under the assumption that a satisfactory well-founded ordering can be defined. Afterwards the proof will motivate the definition of an appropriate ordering \prec_{un} .

For two arbitrary ℓ -expressions s and s' , we want to find an output θ that will satisfy the specification for s and s' . In other words, we want to prove the conclusion

$$P(s, s') : (\exists \theta) \left[\begin{array}{l} s \triangleleft \theta = s' \triangleleft \theta \text{ and} \\ (\forall \lambda)[\text{if } s \triangleleft \lambda = s' \triangleleft \lambda \text{ then } \theta \succeq_{gen} \lambda] \text{ and} \\ \theta = \theta \diamond \theta \text{ and} \\ \theta \neq \text{nil} \\ \text{or} \\ (\forall \lambda)[s \triangleleft \lambda \neq s' \triangleleft \lambda] \text{ and} \\ \theta = \text{nil} \end{array} \right]$$

We will be pleased if θ satisfies either of the two disjuncts in this desired conclusion.

As our induction hypothesis, we assume that the program $unify(r, r')$ that we are trying to construct will satisfy its specification for all inputs r and r' such that the pair $\langle r, r' \rangle$ is strictly less than $\langle s, s' \rangle$ in the selected ordering \prec_{un} . In other words, we assume

$$\begin{array}{l} \text{if } \langle r, r' \rangle \prec_{un} \langle s, s' \rangle \\ \text{then } \left[\begin{array}{l} r \triangleleft unify(r, r') = r' \triangleleft unify(r, r') \text{ and} \\ (\forall \lambda)[\text{if } r \triangleleft \lambda = r' \triangleleft \lambda \text{ then } unify(r, r') \succeq_{gen} \lambda] \text{ and} \\ unify(r, r') = unify(r, r') \diamond unify(r, r') \text{ and} \\ unify(r, r') \neq nil \\ \text{or} \\ (\forall \lambda)[r \triangleleft \lambda \neq r' \triangleleft \lambda] \text{ and} \\ unify(r, r') = nil \end{array} \right] \end{array}$$

for all ℓ -expressions r and r' .

The proof distinguishes among several cases, corresponding to the properties of the apply function \triangleleft . At the end of each case we will give the property of the $unify$ function provided by the proof. Together these properties constitute the final program.

GENERAL CASES

Case (same): $s = s'$.

Then any substitution θ is a unifier, *i.e.*,

$$s \triangleleft \theta = s' \triangleleft \theta.$$

• To find a unifier that is most-general, *i.e.*, such that

$$(\forall \lambda)[\text{if } s \triangleleft \lambda = s' \triangleleft \lambda \text{ then } \theta \succeq_{gen} \lambda],$$

we note that (in this case) any substitution λ satisfies the antecedent $s \triangleleft \lambda = s' \triangleleft \lambda$. Therefore, θ must have the property that

$$\theta \succeq_{gen} \lambda$$

for any substitution λ . The *empty* property of the most-general relation \succeq_{gen} , *i.e.*,

$$\{\} \succeq_{gen} \lambda$$

for any substitution λ , suggests that, in this case, to satisfy the first disjunct of the desired conclusion $P(s, s')$, θ can be taken to be the empty substitution $\{\}$.

- The empty substitution $\{\}$ is idempotent, *i.e.*,

$$\{\} = \{\} \diamond \{\},$$

by either of the *empty* properties of composition.

- We ignore the final requirement, that $\{\} \neq nil$, because any substitution is distinct from *nil*.

Therefore, the empty substitution $\{\}$ is a most-general, idempotent unifier of s and s' , and satisfies the first disjunct of the desired conclusion $P(s, s')$ in this case.

$unify(s, s') = \{\} \quad \text{if } s = s'$

Case (distinct): s and s' are nonvariables of distinct syntactic categories.

Then we can show that s and s' are not unifiable, *i.e.*

$$s \triangleleft \lambda \neq s' \triangleleft \lambda,$$

for any substitution λ . To see this, let λ be an arbitrary substitution. Recall that, because s and s' are not variables, we have (by the syntactic-categories proposition) that

$s \triangleleft \lambda$ is in the same syntactic category as s ,

$s' \triangleleft \lambda$ is in the same syntactic category as s' ,

and therefore

$s \triangleleft \lambda$ and $s' \triangleleft \lambda$ are in distinct syntactic categories.

Hence,

$$s \triangleleft \lambda \neq s' \triangleleft \lambda;$$

i.e., s and s' are not unifiable.

It follows that, in this case, we can satisfy the second disjunct of the desired conclusion $P(s, s')$ by taking θ to be *nil*.

$unify(s, s') = nil$ <p style="text-align: center;">if s and s' are nonvariables of distinct syntactic categories</p>

CONSTANT CASE

Case (distinct): s and s' are distinct constants c and c' respectively.

Then we can show that c and c' are not unifiable, *i.e.*,

$$c \triangleleft \lambda \neq c' \triangleleft \lambda$$

for any substitution λ . For, let λ be an arbitrary substitution. Then (by the *constant* property of the apply function),

$$c \triangleleft \lambda = c$$

and

$$c' \triangleleft \lambda = c'.$$

But, since $c \neq c'$,

$$c \triangleleft \lambda \neq c' \triangleleft \lambda;$$

i.e., c and c' are not unifiable.

It follows that, in this case, we can satisfy the second disjunct of the desired conclusion $P(c, c')$ by taking θ to be *nil*.

$unify(c, c') = nil$ if $c \neq c'$

VARIABLE CASES

Case (not-in, left): s is a variable x , s' is an expression e , and $x \not\triangleleft e$.

- We want to find a most-general unifier of x and e . However, by the variable unifier proposition, $\{x \leftarrow e\}$ is a most-general unifier of x and e . This suggests taking

$$\theta \text{ to be } \{x \leftarrow e\}.$$

- We also want $\theta = \{x \leftarrow e\}$ to be idempotent, *i.e.*, that

$$\{x \leftarrow e\} \diamond \{x \leftarrow e\} = \{x \leftarrow e\}.$$

We show the equivalent condition (by the idempotence proposition) that

$$dom(\{x \leftarrow e\}) \cap range(\{x \leftarrow e\}) = \{\}.$$

We have (from the definition of the domain and range),

$$\text{dom}(\{x \leftarrow e\}) = \{x\} \text{ and}$$

$$\text{range}(\{x \leftarrow e\}) = \text{vars}(e).$$

Because $x \not\prec e$, i.e. (by the variables proposition) $x \notin \text{vars}(e)$, we have

$$\text{dom}(\{x \leftarrow e\}) \cap \text{range}(\{x \leftarrow e\}) = \{x\} \cap \text{vars}(e) = \{\},$$

as we wanted to show.

We have succeeded in showing that $\{x \leftarrow e\}$ is a most-general, idempotent unifier in this case.

$\text{unify}(x, e) = \{x \leftarrow e\} \quad \text{if } x \not\prec e$
--

Case (not-in, right): s is an expression e , s' is a variable x , $\text{not } x \preceq e$, and e is not a variable.

As in the previous case, we find that $\{x \leftarrow e\}$ is a most-general, idempotent unifier of e and x . The condition that the first argument e not be variable is not required in the proof; it is included because the possibility that s and s' are both variables is covered by the previous case.

$\text{unify}(e, x) = \{x \leftarrow e\} \quad \text{if } x \not\prec e \text{ and } e \text{ is not a variable}$

Case (in): s is a variable x , s' is an expression e , and $x \prec e$.

In this case, we can show that x and e are not unifiable. To see this, let λ be an arbitrary substitution. Because $x \prec e$, we have (by the monotonicity proposition)

$$x \triangleleft \lambda \prec e \triangleleft \lambda$$

and hence (by the *irreflexivity* of \prec)

$$x \triangleleft \lambda \not\prec e \triangleleft \lambda.$$

Therefore, x and e are not unifiable and we can satisfy the second disjunct of the desired conclusion $P(x, e)$ by taking θ to be *nil*.

The symmetric case, in which s is an expression e , s' is a variable x , and $x \prec e$, is treated similarly.

$\text{unify}(x, e) = \text{unify}(e, x) = \text{nil} \quad \text{if } x \prec e$

Case (list): s is a variable x and s' is a list of ℓ -expressions m .

In this case, we can show that x and m are not unifiable. To see this, let λ be an arbitrary substitution. Then (by the syntactic categories proposition), $x \triangleleft \lambda$ is an expression but $m \triangleleft \lambda$ is a list of ℓ -expressions; hence (by the disjointness of the syntactic categories)

$$x \triangleleft \lambda \neq m \triangleleft \lambda;$$

i.e., x and m are not unifiable. Therefore, we can satisfy the second disjunct of the desired conclusion $P(x, m)$ by taking θ to be *nil*.

The symmetric case, in which s is a list of ℓ -expressions m and s' is a variable x , is treated similarly.

$$\text{unify}(x, m) = \text{unify}(m, x) = \text{nil}$$

FUNCTIONAL CASES

Case (same): s and s' are functional expressions $f \bullet \ell$ and $f' \bullet \ell'$, respectively, where $f = f'$.

Recall that (in this case) we are attempting to prove the conclusion

$$P(f \bullet \ell, f \bullet \ell');$$

i.e., we want to find an output θ such that

$$\begin{aligned} & (f \bullet \ell) \triangleleft \theta = (f \bullet \ell') \triangleleft \theta \text{ and} \\ & (\forall \lambda)[\text{if } (f \bullet \ell) \triangleleft \lambda = (f \bullet \ell') \triangleleft \lambda \text{ then } \theta \succeq_{\text{gen}} \lambda] \text{ and} \\ & \theta = \theta \diamond \theta \text{ and} \\ & \theta \neq \text{nil} \\ & \text{or} \\ & (\forall \lambda)[(f \bullet \ell) \triangleleft \lambda \neq (f \bullet \ell') \triangleleft \lambda] \text{ and} \\ & \theta = \text{nil}. \end{aligned}$$

This reduces (by the application unifier proposition) to finding an output θ such that

$$\begin{aligned} (*) \quad & \ell \triangleleft \theta = \ell' \triangleleft \theta \text{ and} \\ & (\forall \lambda)[\text{if } \ell \triangleleft \lambda = \ell' \triangleleft \lambda \text{ then } \theta \succeq_{\text{gen}} \lambda] \text{ and} \\ & \theta = \theta \diamond \theta \text{ and} \\ & \theta \neq \text{nil} \\ & \text{or} \\ & (\forall \lambda)[\ell \triangleleft \lambda \neq \ell' \triangleleft \lambda] \text{ and} \end{aligned}$$

$\theta = nil.$

Recall that we have assumed as our induction hypothesis that (in this case)

$$P(r, r') : \text{ if } \langle r, r' \rangle \prec_{un} \langle f \bullet \ell, f \bullet \ell' \rangle$$

$$\text{ then } \left[\begin{array}{l} r \triangleleft unify(r, r') = r' \triangleleft unify(r, r') \text{ and} \\ (\forall \lambda)[\text{if } r \triangleleft \lambda = r' \triangleleft \lambda \text{ then } unify(r, r') \succeq_{gen} \lambda] \text{ and} \\ unify(r, r') = unify(r, r') \diamond unify(r, r') \text{ and} \\ unify(r, r') \neq nil \\ \text{ or} \\ (\forall \lambda)[r \triangleleft \lambda \neq r' \triangleleft \lambda] \text{ and} \\ unify(r, r') = nil \end{array} \right]$$

for all ℓ -expressions r and r' .

The required condition (*) and consequent of the induction hypothesis are identical if we take r to be ℓ , r' to be ℓ' , and θ to be $unify(\ell, \ell')$. Therefore, we can satisfy the conclusion if we can establish the appropriate instance of the induction hypothesis's antecedent, *i.e.*,

$$\langle \ell, \ell' \rangle \prec_{un} \langle f \bullet \ell, f \bullet \ell \rangle. \quad (\text{application ordering})$$

The well-founded ordering \prec_{un} will be chosen subsequently. Assuming it will satisfy this condition, we have found that the desired conclusion $P(f \bullet \ell, f \bullet \ell')$ in this case is satisfied if we take θ to be $unify(\ell, \ell')$.

$$\boxed{unify(f \bullet \ell, f' \bullet \ell') = unify(\ell, \ell') \quad \text{if } f = f'}$$

Case (distinct): s and s' are functional expressions $f \bullet \ell$ and $f' \bullet \ell'$, respectively, where $f \neq f'$.

In this case, we can show that $f \bullet \ell$ and $f' \bullet \ell'$ are not unifiable. For any arbitrary substitution λ , we have (by the *application* property of the apply function)

$$(f \bullet \ell) \triangleleft \lambda = f \bullet (\ell \triangleleft \lambda)$$

and

$$(f' \bullet \ell') \triangleleft \lambda = f' \bullet (\ell' \triangleleft \lambda).$$

Because f and f' are distinct, we have (by the *uniqueness* properties of ℓ -expressions),

$$f \bullet (\ell \triangleleft \lambda) \neq f' \bullet (\ell' \triangleleft \lambda),$$

and hence

$$(f \bullet \ell) \triangleleft \lambda \neq (f' \bullet \ell') \triangleleft \lambda.$$

Thus, $f \bullet \ell$ and $f' \bullet \ell'$ are not unifiable.

Therefore, in this case, we can satisfy the second disjunct of the desired conclusion if we take θ to be *nil*.

$$\boxed{\text{unify}(f \bullet \ell, f' \bullet \ell') = \text{nil} \quad \text{if } f \neq f'}$$

LIST CASES

In all of the list cases, s and s' are nonempty lists $t \circ m$ and $t' \circ m'$, respectively, where t and t' are ℓ -expressions and m and m' are lists of ℓ -expressions.

Recall that (in the list cases) we are attempting to prove the conclusion

$$P(t \circ m, t' \circ m');$$

i.e., we want to find an output θ such that

$$\begin{aligned} &(t \circ m) \triangleleft \theta = (t' \circ m') \triangleleft \theta \text{ and} \\ &(\forall \lambda)[\text{if } (t \circ m) \triangleleft \lambda = (t' \circ m') \triangleleft \lambda \text{ then } \theta \succeq_{\text{gen}} \lambda] \text{ and} \\ &\theta = \theta \diamond \theta \text{ and} \\ &\theta \neq \text{nil} \\ &\text{or} \\ &(\forall \lambda)[(t \circ m) \triangleleft \lambda \neq (t' \circ m') \triangleleft \lambda] \text{ and} \\ &\theta = \text{nil}. \end{aligned}$$

By the Insertion Unifier Proposition, this decomposes into finding θ such that

- (*)
- | | |
|--|--------------------|
| (1) $t \triangleleft \theta = t' \triangleleft \theta$ and | (head unifier) |
| (2) $m \triangleleft \theta = m' \triangleleft \theta$ and | (tail unifier) |
| (3) $(\forall \lambda) \left[\begin{array}{l} \text{if } (t \circ m) \triangleleft \lambda = (t' \circ m') \triangleleft \lambda \\ \text{then } \theta \succeq_{\text{gen}} \lambda \end{array} \right]$ and | (most-generalness) |
| (4) $\theta = \theta \diamond \theta$ and | (idempotence) |
| (5) $\theta \neq \text{nil}$ | (nonnil) |
| or | |
| (6) $(\forall \lambda)[t \triangleleft \lambda \neq t' \triangleleft \lambda \text{ or } m \triangleleft \lambda \neq m' \triangleleft \lambda]$ and | (nonunifiability) |
| (7) $\theta = \text{nil}$. | (nil) |

The separate conditions of (*) are numbered for future reference. We attempt to establish conditions (1) to (5) or, alternatively, conditions (6) and (7).

Recall that we have assumed as our induction hypothesis (in this case)

$$P(r, r') : \text{if } \langle r, r' \rangle \prec_{\text{un}} \langle t \circ m, t' \circ m' \rangle$$

$$\text{then } \left[\begin{array}{l} r \triangleleft \text{unify}(r, r') = r' \triangleleft \text{unify}(r, r') \text{ and} \\ (\forall \lambda)[\text{if } r \triangleleft \lambda = r' \triangleleft \lambda \text{ then } \text{unify}(r, r') \succeq_{\text{gen}} \lambda] \text{ and} \\ \text{unify}(r, r') = \text{unify}(r, r') \diamond \text{unify}(r, r') \text{ and} \\ \text{unify}(r, r') \neq \text{nil} \\ \text{or} \\ (\forall \lambda)[r \triangleleft \lambda \neq r' \triangleleft \lambda] \text{ and} \\ \text{unify}(r, r') = \text{nil} \end{array} \right]$$

for all ℓ -expressions r and r' .

Let us compare this induction hypothesis with our required conclusion (*). A natural approach would be to observe that one of the first two conditions, say the *tail unifier* condition (2),

$$m \triangleleft \theta = m' \triangleleft \theta,$$

is identical to the condition that

$$r \triangleleft \text{unify}(r, r') = r' \triangleleft \text{unify}(r, r'),$$

asserted in our induction hypothesis, if we take r to be m , r' to be m' , and θ to be $\text{unify}(m, m')$. However, we still would have to show the *head unifier* condition (1),

$$t \triangleleft \theta = t' \triangleleft \theta,$$

i.e.,

$$t \triangleleft \text{unify}(m, m') = t' \triangleleft \text{unify}(m, m').$$

But this condition is not necessarily true: a unifier of m and m' need not be a unifier of t and t' . Therefore, we would fail to prove the condition.

An attempt to do the same for the *head unifier* condition (1),

$$t \triangleleft \theta = t' \triangleleft \theta,$$

would fail for the same reason; these two required conditions are symmetric.

A less straightforward approach is to observe that one of the first two conditions, say the *tail unifier* condition (2),

$$m \triangleleft \theta = m' \triangleleft \theta,$$

is also equivalent to the same condition of the induction hypothesis

$$r \triangleleft \text{unify}(r, r') = r' \triangleleft \text{unify}(r, r')$$

under a more complex substitution than we considered earlier.

Proof of the tail unifier condition (2):

Recall the definition of the composition of substitutions

$$m^* \triangleleft (\theta_{hd} \diamond \theta_{tl}) = (m^* \triangleleft \theta_{hd}) \triangleleft \theta_{tl}$$

for all ℓ -expressions m^* and substitutions θ_{hd} and θ_{tl} . (We have renamed the variables to reflect the way we intend to use them.)

Applying this equality to the left-hand side of our required condition

$$m \triangleleft \theta = m' \triangleleft \theta$$

(taking m^* to be m and θ to be $\theta_{hd} \diamond \theta_{tl}$), we obtain the condition

$$(m \triangleleft \theta_{hd}) \triangleleft \theta_{tl} = m' \triangleleft (\theta_{hd} \diamond \theta_{tl}).$$

Applying the same equality to the right-hand side (taking m^* to be m') yields the condition

$$(m \triangleleft \theta_{hd}) \triangleleft \theta_{tl} = (m' \triangleleft \theta_{hd}) \triangleleft \theta_{tl}.$$

This condition is identical to the condition in our induction hypothesis

$$\tau \triangleleft \text{unify}(\tau, \tau') = \tau' \triangleleft \text{unify}(\tau, \tau')$$

if we take τ to be $m \triangleleft \theta_{hd}$, τ' to be $m' \triangleleft \theta_{hd}$, and θ_{tl} to be $\text{unify}(m \triangleleft \theta_{hd}, m' \triangleleft \theta_{hd})$.

Let us retain the abbreviation θ_{tl} for the term $\text{unify}(m \triangleleft \theta_{hd}, m' \triangleleft \theta_{hd})$. Thus, the above match has suggested taking

$$\theta \quad \text{to be} \quad \theta_{hd} \diamond \theta_{tl},$$

where

$$\theta_{tl} \quad \text{is} \quad \text{unify}(m \triangleleft \theta_{hd}, m' \triangleleft \theta_{hd})$$

and θ_{hd} is any substitution.

Let us rewrite the induction hypothesis for this case, making the substitutions suggested by the above match:

$$\begin{array}{l} \text{if } (m \triangleleft \theta_{hd}, m' \triangleleft \theta_{hd}) \prec_{un} (t \circ m, t' \circ m') \\ \text{then } \left[\begin{array}{l} (m \triangleleft \theta_{hd}) \triangleleft \theta_{tl} = (m' \triangleleft \theta_{hd}) \triangleleft \theta_{tl} \text{ and} \\ (\forall \lambda)[\text{if } (m \triangleleft \theta_{hd}) \triangleleft \lambda = (m' \triangleleft \theta_{hd}) \triangleleft \lambda \text{ then } \theta_{tl} \succeq_{gen} \lambda] \text{ and} \\ \theta_{tl} = \theta_{tl} \diamond \theta_{tl} \text{ and} \\ \theta_{tl} \neq nil \\ \text{or} \\ (\forall \lambda)[(m \triangleleft \theta_{hd}) \triangleleft \lambda \neq (m' \triangleleft \theta_{hd}) \triangleleft \lambda] \text{ and} \\ \theta_{tl} = nil \end{array} \right] \end{array}$$

We will refer to this as the *tail induction hypothesis*.

To apply this tail induction hypothesis, we must establish

- the antecedent

$$(8) \quad \langle m \triangleleft \theta_{hd}, m' \triangleleft \theta_{hd} \rangle \triangleleft_{un} \langle t \circ m, t' \circ m' \rangle, \quad (\text{tail ordering})$$

to ensure that the consequent of the tail induction hypothesis will be true;

- the condition

$$(9) \quad \theta_{tl} \neq \text{nil}, \quad (\text{tail nonnil})$$

to ensure that the second disjunct of the consequent of the tail induction hypothesis will be false, so that the first disjunct must be true.

Regardless of the choice of θ_{hd} , this will establish the *tail unifier* condition (2) of the required conclusion (*). We must also establish the *head unifier* condition (1), the *most-generalness* condition (3), and the idempotence condition (4), where θ is taken to be $\theta_{hd} \diamond \theta_{tl}$. We know $\theta_{hd} \diamond \theta_{tl}$ will satisfy the *nonnil* condition (5), since any substitution is distinct from *nil*. In cases where we fail to establish one or more of these conditions, we can alternatively attempt to establish the *nonunifiability* condition (6) and the *nil* condition (7) of the required conclusion (*), taking θ to be *nil*.

We consider these conditions one by one, but not in the given order.

Proof of the head unifier condition (1):

We must find a substitution θ_{hd} such that $\theta_{hd} \diamond \theta_{tl}$ is a unifier of t and t' ; i.e., that

$$t \triangleleft (\theta_{hd} \diamond \theta_{tl}) = t' \triangleleft (\theta_{hd} \diamond \theta_{tl}),$$

or, equivalently (by the definition of composition), that

$$(t \triangleleft \theta_{hd}) \triangleleft \theta_{tl} = (t' \triangleleft \theta_{hd}) \triangleleft \theta_{tl}.$$

It thus suffices to find θ_{hd} such that

$$t \triangleleft \theta_{hd} = t' \triangleleft \theta_{hd}.$$

We observe that the above condition is identical to the condition that

$$r \triangleleft \text{unify}(r, r') = r' \triangleleft \text{unify}(r, r'),$$

asserted in our induction hypothesis, if we take r to be t , r' to be t' and θ_{hd} to be $\text{unify}(t, t')$. We retain the abbreviation

$$\theta_{hd} \text{ is } \text{unify}(t, t').$$

Let us rewrite the induction hypothesis for this case, making the substitutions suggested by the above match:

$$\text{if } \langle t, t' \rangle \prec_{un} \langle t \circ m, t' \circ m' \rangle$$

$$\text{then } \left[\begin{array}{l} t \triangleleft \theta_{hd} = t' \triangleleft \theta_{hd} \text{ and} \\ (\forall \lambda)[\text{if } t \triangleleft \lambda = t' \triangleleft \lambda \text{ then } \theta_{hd} \succeq_{gen} \lambda] \text{ and} \\ \theta_{hd} = \theta_{hd} \diamond \theta_{hd} \text{ and} \\ \theta_{hd} \neq nil \\ \text{or} \\ (\forall \lambda)[t \triangleleft \lambda \neq t' \triangleleft \lambda] \text{ and} \\ \theta_{hd} = nil \end{array} \right]$$

We will refer to this as the *head induction hypothesis*.

To apply this head induction hypothesis, we must establish

- the antecedent

$$\langle t, t' \rangle \prec_{un} \langle t \circ m, t' \circ m' \rangle, \quad (\text{head ordering})$$

to ensure that the consequent of the head induction hypothesis will be true.

- the condition

$$\theta_{hd} \neq nil \quad (\text{head nonnil})$$

to ensure that the second disjunct of the consequent of the head induction hypothesis will be false, so that the first disjunct must be true.

As usual, we defer discussion of the *head ordering* condition, that

$$\langle t, t' \rangle \prec_{un} \langle t \circ m, t' \circ m' \rangle,$$

until we have accumulated all such conditions, so that we can define an ordering \prec_{un} to satisfy them all at once.

The *head nonnil* condition $\theta_{hd} \neq nil$ is not necessarily true: t and t' need not be unifiable. Let us now consider the alternate possibility.

Case(no): $\theta_{hd} = nil$.

Then, by our head induction hypothesis, t and t' are not unifiable; *i.e.*,

$$t \triangleleft \lambda \neq t' \triangleleft \lambda$$

for all substitutions λ . Therefore, we can satisfy the *nonunifiability* condition (6) and the *nil* condition (7) of the required conclusion (*), in this case, by taking θ to be *nil*.

$$\text{unify}(t \circ m, t' \circ m') = \text{nil} \quad \text{if } \theta_{hd} = \text{nil}$$

Case: $\theta_{hd} \neq \text{nil}$.

That is, $\text{unify}(t, t') \neq \text{nil}$. In this case, our head induction hypothesis establishes that θ_{hd} is indeed a most-general, idempotent unifier of t and t' , and therefore the *head unifier* condition (1) is satisfied. It remains to show the *tail ordering* condition (8) and the *tail nonnil* condition (9); these conditions ensure that we can apply the tail induction hypothesis to establish the *tail unifier* condition (2) of the required conclusion (*). It also remains to show the original *most-generalness* condition (3) and *idempotence* condition (4). As usual, we assume that we can establish the *tail ordering* condition (8), and so defer its proof.

Proof of the tail nonnil condition (9):

The condition that $\theta_{tl} \neq \text{nil}$, i.e., $\text{unify}(m \triangleleft \theta_{hd}, m' \triangleleft \theta_{hd}) \neq \text{nil}$, is not necessarily true: $m \triangleleft \theta_{hd}$ and $m' \triangleleft \theta_{hd}$ need not be unifiable. Let us consider the alternative possibility.

Subcase (yes-no): $\theta_{tl} = \text{nil}$.

That is, $\text{unify}(m \triangleleft \theta_{hd}, m' \triangleleft \theta_{hd}) = \text{nil}$. Then, by our tail induction hypothesis (where r and r' were taken to be $m \triangleleft \theta_{hd}$ and $m' \triangleleft \theta_{hd}$, respectively), $m \triangleleft \theta_{hd}$ and $m' \triangleleft \theta_{hd}$ are not unifiable, i.e.,

$$(m \triangleleft \theta_{hd}) \triangleleft \lambda \neq (m' \triangleleft \theta_{hd}) \triangleleft \lambda,$$

for all substitutions λ . We establish, in this case, the *nonunifiability* condition (6), that $(\forall \lambda)[t \triangleleft \lambda \neq t' \triangleleft \lambda \text{ or } m \triangleleft \lambda \neq m' \triangleleft \lambda]$. For suppose, to the contrary, that, for some substitution λ ,

$$t \triangleleft \lambda = t' \triangleleft \lambda,$$

i.e., λ is a unifier of t and t' , and

$$m \triangleleft \lambda = m' \triangleleft \lambda,$$

i.e., λ is a unifier of m and m' .

Because λ is a unifier of t and t' and because θ_{hd} is a most-general, idempotent unifier of t and t' , we have (by the most-general, idempotent unifier proposition)

$$\lambda = \theta_{hd} \diamond \lambda.$$

Hence, because λ is a unifier of m and m' , i.e., $m \triangleleft \lambda = m' \triangleleft \lambda$, we have

$$m \triangleleft (\theta_{hd} \diamond \lambda) = m' \triangleleft (\theta_{hd} \diamond \lambda),$$

or equivalently (by the definition of composition)

$$(m \triangleleft \theta_{hd}) \triangleleft \lambda = (m' \triangleleft \theta_{hd}) \triangleleft \lambda.$$

i.e., λ is a unifier of $m \triangleleft \theta_{hd}$ and $m' \triangleleft \theta_{hd}$. But this contradicts our earlier finding, that $m \triangleleft \theta_{hd}$ and $m' \triangleleft \theta_{hd}$ are not unifiable. Hence, the *nonunifiability* condition (6) is established and, in this case, we can establish the required conclusion (*) by taking θ to be *nil*.

$\text{unify}(t \circ m, t' \circ m') = \text{nil} \quad \text{if } \theta_{hd} \neq \text{nil} \text{ and } \theta_{tl} = \text{nil}$
--

Subcase (yes-yes): $\theta_{tl} \neq \text{nil}$.

That is, the *tail nonnil* condition (9), that $\text{unify}(m \triangleleft \theta_{hd}, m' \triangleleft \theta_{hd}) \neq \text{nil}$, is true. Let us again retrace our steps, to see what we have established.

Since we have assumed the *tail ordering* condition (8), we know that the consequent of the tail induction hypothesis is true (where r and r' were taken to be $m \triangleleft \theta_{hd}$ and $m' \triangleleft \theta_{hd}$, respectively.) Because (in this case), the *tail nonnil* condition (9) is true, we know that the second disjunct of the consequent is false, and that therefore the first disjunct must be true. This implies that θ_{tl} is a most-general, idempotent unifier of $m \triangleleft \theta_{hd}$ and $m' \triangleleft \theta_{hd}$, and hence (by applying the definition of composition twice) that $\theta_{hd} \diamond \theta_{tl}$ is a unifier of m and m' . This establishes the *tail unifier* condition (2), which was our reason for applying the tail induction hypothesis in the first place.

Proof of the most-generalness and idempotence conditions (3 and 4):

It still remains to show the *most-generalness* condition (3),

$$(\forall \lambda) [\text{if } (t \circ m) \triangleleft \lambda = (t' \circ m') \triangleleft \lambda \text{ then } (\theta_{hd} \diamond \theta_{tl}) \succeq_{\text{gen}} \lambda]$$

and the *idempotence* condition (4), that

$$\theta_{hd} \diamond \theta_{tl} = (\theta_{hd} \diamond \theta_{tl}) \diamond (\theta_{hd} \diamond \theta_{tl}).$$

For this purpose, it suffices (by the most-general, idempotent unifier proposition) to show the single condition that

$$(\forall \lambda) \left[\begin{array}{l} \text{if } (t \circ m) \triangleleft \lambda = (t' \circ m') \triangleleft \lambda \\ \text{then } \lambda = (\theta_{hd} \diamond \theta_{tl}) \diamond \lambda \end{array} \right]$$

i.e. (by the insertion unifier proposition),

$$(\forall \lambda) \left[\begin{array}{l} \text{if } t \triangleleft \lambda = t' \triangleleft \lambda \text{ and } m \triangleleft \lambda = m' \triangleleft \lambda \\ \text{then } \lambda = (\theta_{hd} \diamond \theta_{tl}) \diamond \lambda \end{array} \right]$$

Suppose that λ is an arbitrary substitution such that

$$t \triangleleft \lambda = t' \triangleleft \lambda$$

and

$$m \triangleleft \lambda = m' \triangleleft \lambda.$$

We would like to show that then

$$\lambda = (\theta_{hd} \diamond \theta_{tl}) \diamond \lambda.$$

Because λ is a unifier of t and t' , and because θ_{hd} is a most-general, idempotent unifier of t and t' , we have (by the most-general, idempotent unifier proposition, again) that

$$\lambda = \theta_{hd} \diamond \lambda.$$

Therefore, because $m \triangleleft \lambda = m' \triangleleft \lambda$, we have

$$m \triangleleft (\theta_{hd} \diamond \lambda) = m' \triangleleft (\theta_{hd} \diamond \lambda),$$

i.e. (by the definition of composition),

$$(m \triangleleft \theta_{hd}) \triangleleft \lambda = (m' \triangleleft \theta_{hd}) \triangleleft \lambda.$$

In other words, λ is a unifier of $m \triangleleft \theta_{hd}$ and $m' \triangleleft \theta_{hd}$.

But then, because θ_{tl} is a most-general, idempotent unifier of $m \triangleleft \theta_{hd}$ and $m' \triangleleft \theta_{hd}$, we have (by the most-general, idempotent unifier proposition, yet again) that

$$\lambda = \theta_{tl} \diamond \lambda.$$

Therefore,

$$\lambda = \theta_{hd} \diamond \lambda = \theta_{hd} \diamond (\theta_{tl} \diamond \lambda) = (\theta_{hd} \diamond \theta_{tl}) \diamond \lambda.$$

In short, we obtain the condition,

$$\lambda = (\theta_{hd} \diamond \theta_{tl}) \diamond \lambda,$$

that we wanted to show.

We conclude that $\theta_{hd} \diamond \theta_{tl}$ satisfies the *most-generalness* condition (3) and the *idempotence* condition (4) of the required conclusion (*), and thus in this case we are justified in taking θ to be $\theta_{hd} \diamond \theta_{tl}$.

$$\text{unify}(t \circ m, t' \circ m') = \theta_{hd} \diamond \theta_{tl} \quad \text{if } \theta_{hd} \neq \text{nil} \text{ and } \theta_{tl} \neq \text{nil}$$

This concludes the final case.

THE ORDERING

We have deferred the choice of an ordering \prec_{un} to satisfy the ordering conditions we have accumulated during the proof. The choice of this ordering is not so well-motivated formally as the other steps of this derivation. The ordering conditions to be satisfied by \prec_{un} are as follows:

the *application ordering* condition

$$\langle \ell, \ell' \rangle \prec_{un} \langle f \bullet \ell, f \bullet \ell' \rangle,$$

the *head ordering* condition

$$\langle t, t' \rangle \prec_{un} \langle t \circ m, t' \circ m' \rangle,$$

and the *tail ordering* condition

$$\langle m \triangleleft \theta_{hd}, m' \triangleleft \theta_{hd} \rangle \prec_{un} \langle t \circ m, t' \circ m' \rangle$$

for all function symbols f , lists of expressions ℓ and ℓ' , ℓ -expressions t and t' , and lists of ℓ -expressions m and m' , where $\theta_{hd} \neq \text{nil}$, i.e., $\text{unify}(t, t') \neq \text{nil}$.

It would be natural to attempt to use as the definition of \prec_{un} the subexpression ordering upon one of the two arguments. However, if we take \prec_{un} to be, say, the ordering \prec_1 upon the first argument, defined by

$$\langle r, r' \rangle \prec_1 \langle s, s' \rangle$$

if and only if

$$r \prec s,$$

it will satisfy the first two of these conditions, and will satisfy the third condition in the case that $m \triangleleft \theta_{hd} = m$. However, this ordering may fail to satisfy the third condition if $m \triangleleft \theta_{hd} \neq m$, because $m \triangleleft \theta_{hd}$ may no longer be a subexpression of $t \circ m$, and may in fact be much larger. For example, if

$$t \text{ is } x$$

t' is $g(a, y, b)$

and

m is $[f(x, x, x)]$,

m' is $[z]$

then

$t \circ m$ is $[x, f(x, x, x)]$

$t' \circ m'$ is $[g(a, y, b), z]$.

In this case

θ_{hd} is $\{x \leftarrow g(a, y, b)\}$

and

$m \triangleleft \theta_{hd}$ is $[f(g(a, y, b), g(a, y, b), g(a, y, b))]$

$m' \triangleleft \theta_{hd}$ is $[z]$.

Thus, $m \triangleleft \theta_{hd}$ is not a subexpression of $t \circ m$; indeed, it is quite a bit larger.

In the case that $m \triangleleft \theta_{hd} \neq m$, however, it can be shown that the variables of $m \triangleleft \theta_{hd}$ and $m' \triangleleft \theta_{hd}$ are a proper subset of the variables of $t \circ m$ and $t' \circ m'$; i.e.,

$$\text{vars}(m \triangleleft \theta_{hd}) \cup \text{vars}(m' \triangleleft \theta_{hd}) \subset \text{vars}(t \circ m) \cup \text{vars}(t' \circ m').$$

In other words, θ_{hd} removes variables from m and m' without introducing any that are not in $t \circ m$ or $t' \circ m'$. Thus, the ordering \prec_{vars} , defined by

$$\langle r, r' \rangle \prec_{\text{vars}} \langle s, s' \rangle$$

if and only if

$$\text{vars}(r) \cup \text{vars}(r') \subset \text{vars}(s) \cup \text{vars}(s'),$$

will satisfy the *tail ordering* condition in this case. Thus, in the above example,

$$\text{vars}(m \triangleleft \theta_{hd}) \cup \text{vars}(m' \triangleleft \theta_{hd}) = \{y, z\}$$

and

$$\text{vars}(t \circ m) \cup \text{vars}(t' \circ m') = \{x, y, z\},$$

and hence

$$\langle m \triangleleft_{hd}, m' \triangleleft_{hd} \rangle \prec_{vars} \langle t \circ m, t' \circ m' \rangle.$$

However, this ordering will fail to satisfy the first two conditions of \prec_{un} , and will also fail to satisfy the third condition in the case that $m \triangleleft_{hd} = m$ and $m' \triangleleft_{hd} = m'$. For example (by the *application* property of *vars*),

$$vars(\ell) = vars(f \bullet \ell)$$

$$vars(\ell') = vars(f \bullet \ell')$$

and hence

$$vars(\ell) \cup vars(\ell') = vars(f \bullet \ell) \cup vars(f \bullet \ell'),$$

i.e.,

$$\text{not } [\langle \ell, \ell' \rangle \prec_{vars} \langle f \bullet \ell, f \bullet \ell' \rangle].$$

In other words, the first condition, $\langle \ell, \ell' \rangle \prec_{un} \langle f \bullet \ell, f \bullet \ell' \rangle$, is never satisfied under the \prec_{vars} ordering.

The successful ordering \prec_{un} is a lexicographic combination of these two orderings \prec_{vars} and \prec_1 , defined by the property

$$\langle r, r' \rangle \prec_{un} \langle s, s' \rangle$$

if and only if

$$vars(r) \cup vars(r') \subset vars(s) \cup vars(s')$$

or

$$vars(r) \cup vars(r') = vars(s) \cup vars(s') \text{ and } r \prec_1 s.$$

- To see that, under this definition, the *application ordering* condition,

$$\langle \ell, \ell' \rangle \prec_{un} \langle f \bullet \ell, f \bullet \ell' \rangle,$$

is satisfied, note that (as we mentioned above)

$$vars(\ell) \cup vars(\ell') = vars(f \bullet \ell) \cup vars(f \bullet \ell')$$

and (by a *component* property of \prec_1)

$$\ell \prec_1 f \bullet \ell.$$

Hence, by the definition of the ordering \prec_{un} , the *application ordering* condition is indeed satisfied.

- To see that, under this definition, the *head ordering* condition,

$$\langle t, t' \rangle \prec_{un} \langle t \circ m, t' \circ m' \rangle,$$

is satisfied, note that (by the *insertion* property of *vars*)

$$\begin{aligned} vars(t) &\subseteq vars(t) \cup vars(m) = vars(t \circ m) \\ vars(t') &\subseteq vars(t') \cup vars(m') = vars(t' \circ m'), \end{aligned}$$

and hence

$$vars(t) \cup vars(t') \subseteq vars(t \circ m) \cup vars(t' \circ m').$$

In case the inclusion is proper, *i.e.*,

$$vars(t) \cup vars(t') \subset vars(t \circ m) \cup vars(t' \circ m'),$$

we have

$$\langle t, t' \rangle \prec_{un} \langle t \circ m, t' \circ m' \rangle$$

immediately. On the other hand, if the two sets are equal, *i.e.*,

$$vars(t) \cup vars(t') = vars(t \circ m) \cup vars(t' \circ m'),$$

we note that (by a *component* property of \prec)

$$t \prec t \circ m.$$

Hence, by the definition of the ordering \prec_{un} , the *head ordering* condition is also satisfied in this case.

- Finally, we must show that the *tail ordering* condition,

$$\langle m \triangleleft \theta_{hd}, m' \triangleleft \theta_{hd} \rangle \prec_{un} \langle t \circ m, t' \circ m' \rangle,$$

is satisfied, where $\theta_{hd} = unify(t, t') \neq nil$. First, we have

$$vars(m \triangleleft \theta_{hd}) \cup vars(m' \triangleleft \theta_{hd})$$

$$\subseteq vars(m) \cup range(\theta_{hd}) \cup vars(m') \cup range(\theta_{hd})$$

by the variable introduction proposition

$$= vars(m) \cup vars(m') \cup range(\theta_{hd})$$

$$\subseteq vars(m) \cup vars(m') \cup vars(t) \cup vars(t')$$

by the domain and range proposition,
because θ_{hd} is a most-general, idempotent unifier for t and t'

$$= \text{vars}(t \circ m) \cup \text{vars}(t' \circ m')$$

by the *insertion* property of *vars*

In short,

$$\text{vars}(m \triangleleft \theta_{hd}) \cup \text{vars}(m' \triangleleft \theta_{hd}) \subseteq \text{vars}(t \circ m) \cup \text{vars}(t' \circ m').$$

By the definition of the ordering \prec_{un} , we must either show that this inclusion is proper, *i.e.*,

$$(*) \quad \text{vars}(m \triangleleft \theta_{hd}) \cup \text{vars}(m' \triangleleft \theta_{hd}) \subset \text{vars}(t \circ m) \cup \text{vars}(t' \circ m'),$$

or show that

$$(**) \quad m \triangleleft \theta_{hd} \prec t \circ m.$$

For this purpose, we distinguish between two subcases.

Subcase: $m \triangleleft \theta_{hd} = m$.

Then (by a *component* property of \prec)

$$m \triangleleft \theta_{hd} = m \prec t \circ m.$$

Subcase: $m \triangleleft \theta_{hd} \neq m$.

In this case, we will show (*), that the inclusion is proper, *i.e.*,

$$\text{vars}(m \triangleleft \theta_{hd}) \cup \text{vars}(m' \triangleleft \theta_{hd}) \subset \text{vars}(t \circ m) \cup \text{vars}(t' \circ m').$$

We have already shown the \subseteq inclusion; therefore, it suffices to show the existence of a variable z such that

$$(\ddagger) \quad z \in \text{vars}(t \circ m) \cup \text{vars}(t' \circ m')$$

but

$$(\dagger\dagger) \quad z \notin \text{vars}(m \triangleleft \theta_{hd}) \cup \text{vars}(m' \triangleleft \theta_{hd}).$$

First, because $m \triangleleft \theta_{hd} \neq m$, we know (by the invariance corollary) that

$$\text{vars}(m) \cap \text{dom}(\theta_{hd}) \neq \{\},$$

i.e., there is a variable z such that

$$z \in \text{vars}(m)$$

and

$$z \in \text{dom}(\theta_{hd}).$$

We know (by the *insertion* property of *vars*) that

$$\text{vars}(m) \subseteq \text{vars}(t) \cup \text{vars}(m) = \text{vars}(t \circ m).$$

Thus, because $z \in \text{vars}(m)$, we have the desired property (†),

$$z \in \text{vars}(t \circ m) \cup \text{vars}(t' \circ m').$$

Next, because θ_{hd} is idempotent, we have (by idempotence proposition)

$$\text{dom}(\theta_{hd}) \cap \text{range}(\theta_{hd}) = \{\},$$

and, therefore, because

$$z \in \text{dom}(\theta_{hd}),$$

we have

$$z \notin \text{range}(\theta_{hd}).$$

It follows (by the variable elimination proposition) that

$$z \notin \text{vars}(m \triangleleft \theta_{hd})$$

$$z \notin \text{vars}(m' \triangleleft \theta_{hd}),$$

and hence

$$z \notin \text{vars}(m \triangleleft \theta_{hd}) \cup \text{vars}(m' \triangleleft \theta_{hd}).$$

This is the desired property (††). We have thus established the proper inclusion

$$\text{vars}(m \triangleleft \theta_{hd}) \cup \text{vars}(m' \triangleleft \theta_{hd}) \subset \text{vars}(t \circ m) \cup \text{vars}(t' \circ m').$$

In both subcases, we can conclude that

$$(m \triangleleft \theta_{hd}, m' \triangleleft \theta_{hd}) \prec_{un} (t \circ m, t' \circ m').$$

Thus, the *tail ordering* condition for \prec_{un} is satisfied.

This concludes the entire derivation proof. \blacksquare

ALTERNATE DERIVATIONS

We have followed only one proof of the specification of the desired theorem. Had we followed other proofs, different programs would have resulted. For example, in the *list* cases of the above derivation, we first matched the *tail unifier* condition $m \triangleleft \theta = m' \triangleleft \theta$ against the induction hypothesis (after applying the definition of composition); had we instead matched the symmetric *head unifier* condition $t \triangleleft \theta = t' \triangleleft \theta$, another proof would have been obtained and the *list* cases of the resulting program would have been as follows:

```

let  $\theta_{tl}^* = unify(m, m')$ 
if  $\theta_{tl}^* = nil$ 
then  $unify(t \circ m, t' \circ m') = nil$ 
else let  $\theta_{hd}^* = unify(t \triangleleft \theta_{tl}^*, t' \triangleleft \theta_{tl}^*)$ 
      if  $\theta_{hd}^* = nil$ 
      then  $unify(t \circ m, t' \circ m') = nil$ 
      else  $unify(t \circ m, t' \circ m') = \theta_{tl}^* \diamond \theta_{hd}^*$ 

```

This program will also satisfy the same specification as the original program but, because it examines the list from right to left rather than left to right, it may produce a different most-general, idempotent unifier.

In general, by exploring different branches of the proof tree, we may obtain families of different unification algorithms analogous to the families of different sort programs obtained from a single specification by Clark and Darlington [1980]. The particular derivation we obtained did not take the efficiency of the final program into account. Other branches of the derivation tree lead to more efficient unification algorithms.

AUTOMATION OF THE PROOF

Our primary objective in examining the above derivation in such detail is to consider the computational prerequisites for discovering the proof automatically. Let us review the proof from this point of view.

The first requirement of a theorem-proving system for program synthesis is that it be able to prove theorems that contain existential quantifiers and require mathematical induction. Existential quantifiers are necessary to transform the specification into a theorem, and induction is necessary to introduce repetitive constructs into the target program. Although resolution theorem provers, say, can prove theorems with existential quantifiers, and several theorem provers (e.g., Boyer and

Moore [1975], Huet and Hullot [1980]) can do proofs by induction, it is rare to see these abilities combined.

The amount of knowledge about ℓ -expressions and substitutions necessary to produce the above proof is formidable. If such knowledge were built into the system, the system would then be tailored to this specific subject domain and would thus lose generality. On the other hand, if the knowledge were introduced into the system as a set of axioms, the system would also need to know how to use the knowledge efficiently.

Much of the derivation proof is fairly mechanical. At each stage, one must decide which property, from a finite collection of legal next steps, to apply. However, certain steps are not straightforward: they are motivated only by their ultimate success. For example, in the list cases, the straightforward use of the induction hypothesis failed, but a surprising application of the definition of composition allowed us to use the induction hypothesis in a more general way, which resulted in the introduction of the composition $\theta_{hd} \diamond \theta_{tl}$ in the final program.

In finding the well-founded ordering $<_{un}$, the use of the sets $vars(s)$ and $vars(s')$ of variables in the ℓ -expressions s and s' was not suggested by the specification, which makes no reference to this notion.

The idempotence condition $\theta = \theta \diamond \theta$ was included in the initial specification. This condition played a vital part in the proof; however, the unification algorithm would be equally useful without this property. Had the idempotence condition not been required initially, either it or an equivalent condition would have had to be invented and added to the specification in the middle of the proof.

Even with the idempotence condition provided, the proof seems somewhat more difficult than current theorem provers can produce. Our hope is that studying hand derivations of this sort will enable us to improve the power of automatic systems.

INTERACTIVE SYNTHESIS

Although the above proof may be beyond the power of current automatic systems, a partially interactive system could be used to produce it with known techniques. This approach requires more human effort, but it would still convey many of the benefits of automatic synthesis:

- The person would provide those steps that require cleverness, but the system would take care of the routine details.
- Whatever mistakes the person might make, the system would not permit him to produce a program that did not meet its specification.
- The program would be accompanied by a full proof of its correctness.
- The derivation could be retained so that, if the program needed modification, the appropriate portions of the program could be updated without endangering its correctness.

- The assumptions on which the correct operation of the program depends would be made explicit.

Of course, for an interactive system to be successful, it would have to communicate in terms the person would be able to understand.

ACKNOWLEDGMENTS

We would like to thank Yoni Malachai, Pierre Wolper, and Frank Yellin, for their careful reading of the manuscript and their many suggestions; and Evelyn Eldridge-Diaz for her TEXing of the paper.

REFERENCES

- Bledsoe, W. W. [Aug. 1977], "Non-resolution theorem proving," *Artificial Intelligence*, Vol. 9, No. 1, pp. 1-35.
- Boyer, R. S. and J. S. Moore [Jan. 1975], "Proving theorems about LISP functions," *JACM*, Vol. 22, No. 1, pp. 129-144.
- Clark, K. and J. Darlington [Feb. 1980], "Algorithm classification through synthesis," *Computer Journal*, Vol. 23, No. 1, pp. 61-65.
- Colmerauer, A., H. Kanoui and M. van Caneghem [1979], "Etude et realisation d'un systeme Prolog," Internal Report, Groupe d'Intelligence Artificielle, U.E.R. de Luminy, Université d' Aix-Marseille II.
- Green, C. C. [May 1969], "Application of theorem proving to problem solving," *Proceedings of the International Joint Conference on Artificial Intelligence*, Washington DC, pp. 219-239.
- Herbrand, J. [1930], "Researches in the theory of demonstration," in *From Frege to Gödel: a source book in mathematical logic, 1879-1931* (J. Van Heijenoort, ed.), Harvard University Press, Cambridge, MA, 1967, pp. 525-581.
- Hewitt, C. [Apr. 1971], "Description and theoretical analysis (using schemata) of PLANNER: A language for proving theorems and manipulating models in a robot," Ph.D. thesis, MIT, Cambridge, MA.
- Huet, G. P. [June 1975], "A unification algorithm for typed λ -calculus," *Theoretical Computer Science*, Vol. 1, No. 1, pp. 27-57.
- Huet, G. P. and J.-M. Hullot [Oct. 1980], "Proofs by induction in equational theories with constructors," *Proceedings of Symposium on Foundations of Computer Science*, Syracuse, NY, pp. 96-107.

- Livesay, M., J. Siekmann, P. Szabó and E. Unvericht [Feb. 1979], "Unification problems for combinations of associativity, commutativity, distributivity and idempotence axioms," *Proceedings of the Fourth Workshop on Automated Deduction*, Austin, TX, pp. 175-184.
- Manna, Z. and R. Waldinger [Summer 1975], "Knowledge and reasoning in program synthesis," *Artificial Intelligence*, Vol. 6, No. 2, pp. 175-208.
- Manna, Z. and R. Waldinger [Jan. 1980], "A deductive approach to program synthesis," *ACM Transactions on Programming Languages and Systems*, Vol. 2, No. 1, pp. 92-121.
- Manna, Z. and R. Waldinger [1982], *Deductive Basis for Computer Programming*, forthcoming.
- Martelli, A. and U. Montanari [July 1976], "Unification in linear time and space: a structured presentation," Internal Report, IEL, Pisa.
- Paterson, M. S. and M. N. Wegman [April 1978], "Linear unification," *Journal of Computer and System Sciences*, Vol. 16, No. 2, pp. 158-167.
- Prawitz, D. [1960], "An improved proof procedure," *Theoria*, Vol. 26, pp. 102-139.
- Robinson, J. A. [Jan. 1965], "A machine-oriented logic based on the resolution principle", *JACM*, Vol. 12, No. 1, pp. 23-41.
- Stickel, M. E. [Sept. 1975], "A complete unification algorithm for associative-commutative functions," *Proceedings of the Fourth International Joint Conference on Artificial Intelligence*, Tbilisi, Georgia, USSR, pp. 71-76.
- Von Henke, F. W. and D. C. Luckham [April 1974], "A methodology for verifying programs," *Proceedings of the International Conference on Reliable Software*, Los Angeles, CA, pp. 156-164.
- Waldinger, R. J. and R. C. T. Lee [May 1969], "PROW: a step toward automatic program writing," *Proceedings of the International Joint Conference on Artificial Intelligence*, Washington, DC, pp. 241-252.
- Waldinger, R. J. and K. N. Levitt [Fall 1974], "Reasoning about programs," *Artificial Intelligence*, Vol. 5, No. 3, pp. 235-316.
- Warren, D., L. M. Pereira and F. Pereira [Aug 1977], "PROLOG - the language and its implementation compared with LISP," *Proceedings of Symposium on Artificial Intelligence and Programming Languages, SIGPLAN Notices (ACM)*, Vol. 12, No. 8, pp. 109-115.

