

Logic Programming

Theory Lecture 4: Proof Search for Definite Clause Predicate Logic

Richard Mayr

School of Informatics

23th October 2014

Recap (Lecture 3): Definite clause predicate logic

A *definite clause* is a formula of one of the two shapes below

B (a Prolog *fact* B .)

$A_1 \wedge \dots \wedge A_k \rightarrow B$ (a Prolog *rule* $B :- A_1, \dots, A_k$.)

where A_1, \dots, A_k, B are all *atomic formulas*.

A *logic program* is a list F_1, \dots, F_n of definite clauses


A *goal* is a list G_1, \dots, G_m of atomic formulas.

The job of the system is to ascertain whether the logical consequence below holds.

$$\forall \text{Vars}(F_1). F_1, \dots, \forall \text{Vars}(F_n). F_n \models \exists \text{Vars}(G_1, \dots, G_m). G_1 \wedge \dots \wedge G_m$$

Example: Leaf-labelled binary trees

E.g.



```
nd(lf(a),nd(lf(b),nd(lf(c),lf(a))))
```

The path predicate :

```
path(a, nd(lf(a),nd(lf(b),nd(lf(c),lf(a))))), [l])
path(b, nd(lf(a),nd(lf(b),nd(lf(c),lf(a))))), [r,l])
path(c, nd(lf(a),nd(lf(b),nd(lf(c),lf(a))))), [r,r,l])
path(a, nd(lf(a),nd(lf(b),nd(lf(c),lf(a))))), [r,r,r])
```

Logic program for path

Prolog notation

```
path(X,lf(X), []).  
path(X,nd(T,_), [l|P]) :- path(X,T,P).  
path(X,nd(_,T), [r|P]) :- path(X,T,P).
```

Logical notation (with explicit universal quantification)

```
∀X. path(X,lf(X), [])  
∀X,T,U,P. path(X,T,P) → path(X,nd(T,U), [l|P])  
∀X,S,T,P. path(X,T,P) → path(X,nd(S,T), [r|P])
```

Example goal

Prolog notation

`path(b,nd(lf(a),nd(lf(b),lf(c))),Q), path(Y,nd(lf(c),nd(lf(a),lf(b))),Q)`

Logical notation (with explicit existential quantification)

$\exists Q, Y. \text{path}(b, \text{nd}(\text{lf}(a), \text{nd}(\text{lf}(b), \text{lf}(c))), Q) \wedge \text{path}(Y, \text{nd}(\text{lf}(c), \text{nd}(\text{lf}(a), \text{lf}(b))), Q)$

The next slide depicts (relevant parts of) the search tree for this goal.

$\text{path}(b, \text{nd}(\text{lf}(a), \text{nd}(\text{lf}(b), \text{lf}(c))), Q), \text{path}(Y, \text{nd}(\text{lf}(c), \text{nd}(\text{lf}(a), \text{lf}(b))), Q)$

$\{Q = [1|Q']\}$

$\text{path}(b, \text{lf}(a), Q')$,

$\text{path}(Y, \text{nd}(\text{lf}(c), \text{nd}(\text{lf}(a), \text{lf}(b))), [1|Q'])$

fail

$\{Q = [r|Q']\}$

$\text{path}(b, \text{nd}(\text{lf}(b), \text{lf}(c)), Q')$,

$\text{path}(Y, \text{nd}(\text{lf}(c), \text{nd}(\text{lf}(a), \text{lf}(b))), [r|Q'])$

$\{Q' = [1|Q'']\}$

$\text{path}(b, \text{lf}(b), Q'')$,

$\text{path}(Y, \text{nd}(\text{lf}(c), \text{nd}(\text{lf}(a), \text{lf}(b))), [r, 1|Q''])$

$\{Q' = [r|Q'']\}$

\vdots

$\{Q'' = []\}$

$\text{path}(Y, \text{nd}(\text{lf}(c), \text{nd}(\text{lf}(a), \text{lf}(b))), [r, 1])$

$\text{path}(Y, \text{nd}(\text{lf}(a), \text{lf}(b)), [1])$

$\text{path}(Y, \text{lf}(a), [])$

$\{Y = a\}$

yes

Result of example query

The query

$\exists Q, Y. \text{path}(b, \text{nd}(\text{lf}(a), \text{nd}(\text{lf}(b), \text{lf}(c))), Q) \wedge \text{path}(Y, \text{nd}(\text{lf}(c), \text{nd}(\text{lf}(a), \text{lf}(b))), Q)$

succeeds, returning the result of the combined substitution

$$\{ Q = [r|Q'] \} \{ Q' = [l|Q''] \} \{ Q'' = [] \} \{ Y = a \}$$

which is the substitution

$$\{ Q = [r, l], Y = a \}$$

The resulting substitution provides *witnessing terms* for the existentially quantified variables. That is, the following formula is a logical consequence of the program.

$\text{path}(b, \text{nd}(\text{lf}(a), \text{nd}(\text{lf}(b), \text{lf}(c))), [r, l]) \wedge \text{path}(a, \text{nd}(\text{lf}(c), \text{nd}(\text{lf}(a), \text{lf}(b))), [r, l])$

Substitutions

A substitution is an assignment of terms to variables:

$$\{ X_1 = t_1, \dots, X_n = t_n \}$$

where the variables X_1, \dots, X_n are all distinct.

Other notation often used for substitutions is:

$$\begin{aligned} & \{ X_1 \setminus t_1, \dots, X_n \setminus t_n \} \\ & \{ X_1 / t_1, \dots, X_n / t_n \} \\ & \{ t_1 / X_1, \dots, t_n / X_n \} \end{aligned}$$

Note that the last two are mutually inconsistent.

Applying substitutions

A substitution $\{X_1 = t_1, \dots, X_n = t_n\}$ is *applied* to a term t or a formula F by replacing all the (free) occurrences of the variables X_1, \dots, X_n with the terms t_1, \dots, t_n respectively.

We write:

$$\begin{aligned} & t \{ X_1 = t_1, \dots, X_n = t_n \} \\ & F \{ X_1 = t_1, \dots, X_n = t_n \} \end{aligned}$$

for the resulting substituted term and formula.

Technical restriction on substitutions

In logic programming, the application of the substitution $\{X_1 = t_1, \dots, X_n = t_n\}$ can be restricted only to terms t and formulas F for which

$$(Vars(t) \cup \{X_1, \dots, X_n\}) \cap Vars(t_1, \dots, t_n) = \emptyset$$

$$(Vars(F) \cup \{X_1, \dots, X_n\}) \cap Vars(t_1, \dots, t_n) = \emptyset$$

Informally, the restriction says that variables in the new terms being substituted in must be “fresh”.

The restriction is technically motivated. We shall comment further on the effects of this restriction in purple. Such comments are for the benefit of the mathematically inquisitive.

Example

The result of the substitution

$$f(X, g(Y, h(X))) \{ X = h(Z), Y = f(W, W) \}$$

is the term

$$f(h(Z), g(f(W, W), h(h(Z))))$$

Note how the substitutions for X and Y are performed simultaneously.

Example (continued)

Substitutions are performed in sequence in a composite substitution:

$$f(X, g(W, h(X))) \{ X = h(Y) \} \{ Y = f(Z, Z) \}$$

(W is not allowed to be Y or Z here). The result of this is

$$f(h(f(Z, Z)), g(W, h(h(f(Z, Z))))))$$

The composite substitution $\{ X = h(Y) \} \{ Y = f(Z, Z) \}$ has the same effect (on terms not containing Y, Z) as the substitution:

$$\{ X = h(f(Z, Z)) \}$$

We write:

$$\{ X = h(Y) \} \{ Y = f(Z, Z) \} = \{ X = h(f(Z, Z)) \}$$

Composition of substitutions

In general, given substitutions θ_1 and θ_2 , one can (when there is no clash of variables) define a *composite* substitution

$$\theta_1 \theta_2$$

that satisfies, for every term t (for which both sides are legitimate),

$$t(\theta_1 \theta_2) = (t\theta_1)\theta_2$$

Let $\{\}$ be the *identity* substitution, which has no effect. That is $t\{\} = t$.

Composition is an example of a (partial) *monoid*, i.e.,

$$\begin{aligned}\{\} \theta &= \theta = \theta \{\} \\ \theta_1 (\theta_2 \theta_3) &= (\theta_1 \theta_2) \theta_3\end{aligned}$$

Accordingly, we can write $\theta_1 \theta_2 \theta_3$ without ambiguity.

Generality of substitutions

We say that a substitution θ_2 is *more general than* a substitution θ_1 , notation

$$\theta_1 \preceq \theta_2$$

if there exists a substitution θ such that

$$\theta_1 = \theta_2 \theta$$

Example

$$\{X = f(a), Y = g(a, b), W = f(b)\} \preceq \{X = f(Z), Y = g(Z, V)\}$$

because

$$\{X = f(a), Y = g(a, b), W = f(b)\} = \{X = f(Z), Y = g(Z, V)\} \{Z = a, V = b, W = f(b)\}$$

Properties of generality

The relation \preceq is a *preorder*, that is:

$$\begin{aligned}\theta &\preceq \theta \\ \theta_1 \preceq \theta_2 \wedge \theta_2 \preceq \theta_3 &\rightarrow \theta_1 \preceq \theta_3\end{aligned}$$

We say that substitutions θ_1 and θ_2 are *equivalent* if

$$\theta_1 \preceq \theta_2 \wedge \theta_2 \preceq \theta_1$$

θ_1 and θ_2 are equivalent if and only if there exists a *renaming*, that is, a substitution

$$\{ X_1 = Y_1, \dots, X_n = Y_n \}$$

where the variables Y_1, \dots, Y_n are all distinct, such that

$$\theta_1 = \theta_2 \{ X_1 = Y_1, \dots, X_n = Y_n \}$$

Most general unifiers

A *unifier* for two terms t_1, t_2 is a substitution θ such that

$$t_1 \theta = t_2 \theta$$

We say that t_1, t_2 are *unifiable* if there exists a unifier.

Theorem.

If t_1, t_2 are unifiable then they have a *most general unifier*. That is, there exists a unifier θ satisfying

for every unifier θ' , it holds that $\theta' \preceq \theta$.

Proof. Find θ using the algorithm in Programming Lecture 2. \square

Note that most general unifiers are unique up to equivalence.

Similarly, most general unifiers exist for unifiable formulas A_1, A_2 .

Example search tree revisited

For the goal

$$\text{path}(\text{b}, \text{nd}(\text{lf}(\text{b}), \text{lf}(\text{c})), \text{Q}'), \text{path}(\text{Y}, \text{nd}(\text{lf}(\text{c}), \text{nd}(\text{lf}(\text{a}), \text{lf}(\text{b}))), [\text{r}|\text{Q}']$$

We unify the head atomic formula $\text{path}(\text{b}, \text{nd}(\text{lf}(\text{b}), \text{lf}(\text{c})), \text{Q}')$ with the right-hand-side of the rule

$$\text{path}(\text{X}, \text{T}, \text{P}) \rightarrow \text{path}(\text{X}, \text{nd}(\text{T}, \text{U}), [\text{l}|\text{P}])$$

Giving the most general unifier

$$\theta = \{ \text{X} = \text{b}, \text{T} = \text{lf}(\text{b}), \text{U} = \text{lf}(\text{c}), \text{Q}' = [\text{l}|\text{Q}''], \text{P} = \text{Q}'' \} \quad (*)$$

We replace the goal with

$$\begin{aligned} & \text{path}(\text{X}, \text{T}, \text{P}) \theta, \text{path}(\text{Y}, \text{nd}(\text{lf}(\text{c}), \text{nd}(\text{lf}(\text{a}), \text{lf}(\text{b}))), [\text{r}|\text{Q}'] \theta \\ & = \text{path}(\text{b}, \text{lf}(\text{b}), \text{Q}'), \text{path}(\text{Y}, \text{nd}(\text{lf}(\text{c}), \text{nd}(\text{lf}(\text{a}), \text{lf}(\text{b}))), [\text{r}, \text{l}|\text{Q}'] \end{aligned}$$

and we record θ restricted to goal variables $\{ \text{Q}' = [\text{l}|\text{Q}''] \}$

Building the search tree in general

At each node in the search tree, the *current goal* is a list G_1, \dots, G_m of atomic formulas.

1. If G_1 unifies with some axiom B (a Prolog fact) in the program then generate a new goal

$$G_2\theta, \dots, G_m\theta ,$$

where θ is the most general unifier of G_1 and B .

Record the effect of the substitution θ on variables occurring in G_1, \dots, G_m .

Building the search tree in general (continued)

2. If G_1 unifies with the right-hand side B of an axiom (a Prolog rule) of the form

$$A_1 \wedge \dots \wedge A_k \rightarrow B$$

in the program, then generate a new goal of the form:

$$A_1\theta, \dots, A_k\theta, G_2\theta, \dots, G_m\theta$$

where θ is the most general unifier of G_1 and B .

Again, record the effect of the substitution θ on variables occurring in G_1, \dots, G_m .

The goal of the proof search procedure is to find a branch ending in a leaf with the *empty* goal list. Whence Prolog returns the composite substitution along this branch, which provides a substitution witnessing the original query.

A subtle point

In the above, G_1 is always unified with an atomic formula B , where B is either itself one of the program axioms F_i , or it is the right-hand side of a rule F_i . In either case, the variables in F_i are implicitly universally quantified. Similarly, the variables in G_1, \dots, G_m are implicitly existentially quantified.

Because of this, it does not matter what the variables are called. The meaning of the clause is unaffected by a different choice of variable names.

However, for the unification in the proof search to work correctly, it is essential that the variables names in F_i and G_1, \dots, G_m are chosen so that they do not overlap.

This motivates the choice of variable names in the most general unifier (*) on the Example Search Tree Revisited slide.

Prolog proof search

Prolog proof search is *depth first*:

- ▶ The search always moves to an *unvisited* (i.e., not previously visited) *child* (i.e., immediately below) node of the current node, whenever such a node exists.
- ▶ If there is no such node, the search backtracks to the most recently visited node from which such an unvisited child node is available.

Prolog proof search follows *program order*.

- ▶ Child nodes are visited in the order that the axioms (Prolog rules) that determine the child node appear in the program (i.e., from left to right in the trees as we are drawing them).

The strategy here is identical to that for propositional Prolog

Main points today

substitutions

unifiers and most general unifiers

search tree for definite clause predicate logic

Prolog proof search strategy

Appendix 1: Restriction on application of substitutions

We would like to have that

$$\{X = a, Y = a\} \preceq \{X = Z, Y = Z\}$$

via the expected composition of substitutions

$$\{X = a, Y = a\} = \{X = Z, Y = Z\} \{Z = a\}$$

But if we were to allow the application of $\{X = Z, Y = Z\}$ to terms containing Z we would get, for example,

$$\begin{aligned} f(X, Y, Z) \{X = Z, Y = Z\} \{Z = a\} &= f(Z, Z, Z) \{Z = a\} \\ &= f(a, a, a) \end{aligned}$$

$$f(X, Y, Z) \{X = a, Y = a\} = f(a, a, Z)$$

Appendix 2: Non-overlapping variables requirement

Consider the simple logic program

$$p(a, X)$$

and the simple query $p(X, b)$.

The query should succeed with substitution $\{ X = a \}$

As written, the goal $p(X, b)$ and the axiom $p(a, X)$ do not unify, because there is no consistent assignment for the variable X .

It is thus necessary to rename variables so that the variables in the program do not overlap with variables in the goal.

The unification of $p(X, b)$ and the axiom $p(a, X')$ does succeed with most general unifier $\{ X = a, X' = b \}$, and so the correct substitution $\{ X = a \}$ is returned.