

Logic Programming

Theory Lecture 5: (In)completeness for Definite Clause Predicate Logic

Richard Mayr

School of Informatics

27th 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 (Lecture 4): 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])
```

Example program, query and result (Lecture 4)

Program

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

Goal

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

Result

```
{Q = [r,l], Y = a}
```

Prolog proof search and inference

Given a program

$$F_1, \dots, F_n$$

and goal

$$G_1, \dots, G_m ,$$

a Prolog proof search, if successful, results in a substitution θ , witnessing the implicit existential quantification in the goal.

As in the case of propositional logic (Lectures 1 and 2), it is helpful to understand the proof search procedure as constructing a derivation (proof) in an inference system for definite clause predicate logic.

Inference system (a.k.a. SLD resolution)

As before, we use a single inference rule to derive new *established goals* from goals already established.

$$\frac{(C_1, \dots, C_{l-1}, A_1, \dots, A_k, C_{l+1}, \dots, C_m)\theta \quad A_1 \wedge \dots \wedge A_k \rightarrow B}{C_1, \dots, C_{l-1}, C_l, C_{l+1}, \dots, C_m}$$

where:

- ▶ $1 \leq l \leq m$,
- ▶ $k \geq 0$,
- ▶ $A_1 \wedge \dots \wedge A_k \rightarrow B$ is one of the program clauses F_i , with its variables renamed to variables not appearing in C_1, \dots, C_m .
- ▶ θ is the most general unifier of B and C_l .

A *derivation* (or *proof*) is a sequence of applications of the above rule, in which the topmost rule has the empty goal as its premise.

Technical points

The general format $A_1 \wedge \dots \wedge A_k \rightarrow B$ of clauses includes Prolog *facts* as the special case in which $k = 0$.

In this case, the inference rule specialises to

$$\frac{(C_1, \dots, C_{l-1}, C_{l+1}, \dots, C_m)\theta \quad B}{C_1, \dots, C_{l-1}, C_l, C_{l+1}, \dots, C_m}$$

Technical Lemma 1

If the two goals G_1, \dots, G_m and G'_1, \dots, G'_n are derivable and have no variables in common then the juxtaposed goal

$G_1, \dots, G_m, G'_1, \dots, G'_n$ is also derivable.

Technical Lemma 2

If a substitution instance $(G_1, \dots, G_m)\theta$ of a goal G_1, \dots, G_m is derivable then the goal G_1, \dots, G_m is itself derivable.

Program:

$$\begin{aligned} F_1 &= \text{path}(X, \text{lf}(X), []) \\ F_2 &= \text{path}(X, T, P) \rightarrow \text{path}(X, \text{nd}(T, U), [l | P]) \\ F_3 &= \text{path}(X, T, P) \rightarrow \text{path}(X, \text{nd}(S, T), [r | P]) \end{aligned}$$

Example derivation:

$$\begin{array}{c} \epsilon \quad F_1 \\ \hline \text{path}(Y, \text{lf}(a), []) \quad F_2 \\ \hline \text{path}(Y, \text{nd}(\text{lf}(a), \text{lf}(b)), [l]) \quad F_3 \\ \hline \text{path}(Y, \text{nd}(\text{lf}(c), \text{nd}(\text{lf}(a), \text{lf}(b))), [r, l]) \quad F_1 \\ \hline \text{path}(b, \text{lf}(b), Q''), \text{path}(Y, \text{nd}(\text{lf}(c), \text{nd}(\text{lf}(a), \text{lf}(b))), [r, l | Q'']) \quad F_2 \\ \hline \text{path}(b, \text{nd}(\text{lf}(b), \text{lf}(c)), Q'), \text{path}(Y, \text{nd}(\text{lf}(c), \text{nd}(\text{lf}(a), \text{lf}(b))), [r | Q']) \quad F_3 \\ \hline \text{path}(b, \text{nd}(\text{lf}(a), \text{nd}(\text{lf}(b), \text{lf}(c))), Q), \text{path}(Y, \text{nd}(\text{lf}(c), \text{nd}(\text{lf}(a), \text{lf}(b))), Q) \end{array}$$

Prolog proof search finds derivations

Every successful branch in a Prolog search tree gives rise to a derivation in the inference system.

Essentially, we construct the derivation by turning the branch upside down.

This observation underlies:

Proposition

If Prolog proof search, for a program F_1, \dots, F_n and goal G_1, \dots, G_m succeeds then the goal has a derivation in the inference system.

Soundness of inference system

Theorem

If goal G_1, \dots, G_m is derivable in the inference system for a program F_1, \dots, F_n , then

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

Proof

Similar to proof of soundness in the propositional case (Lecture 2), though the definition of logical consequence in terms of structures, rather than truth-value interpretations, adds (significant but tedious) complications.

Soundness of inference system

Theorem

If goal G_1, \dots, G_m is derivable in the inference system for a program F_1, \dots, F_n , then

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

Proof (Non-examinable)

Similar to proof of soundness in the propositional case (Lecture 2), though the definition of logical consequence in terms of structures, rather than truth-value interpretations, adds (significant but tedious) complications.

Completeness and Incompleteness

Completeness of inference system

For any program F_1, \dots, F_n and goal G_1, \dots, G_m such that

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

there exists a derivation of G_1, \dots, G_m in the inference system ...

Incompleteness of Prolog proof search

... however, Prolog proof search for the goal G_1, \dots, G_m need not succeed (it may fail to terminate).

Completeness and Incompleteness

Completeness of inference system

For any program F_1, \dots, F_n and goal G_1, \dots, G_m such that

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

there exists a derivation of G_1, \dots, G_m in the inference system ...

Incompleteness of Prolog proof search

... however, Prolog proof search for the goal G_1, \dots, G_m need not succeed (it may fail to terminate).

The example used to demonstrate incompleteness of propositional Prolog search in Lecture 2 is again an example of the incompleteness of proof search in the case of predicate logic.

The proof of completeness of the inference system is more involved for predicate logic than for propositional logic. We consider this in some detail, since it involves an important construction.

Outline proof of completeness

Suppose that

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

We shall construct a very special *structure* (in the sense of Lecture 3), the *minimal Herbrand model* of the program F_1, \dots, F_n .

Because the minimal Herbrand model \mathcal{H} is a model, we have

$$\mathcal{H} \models \exists \text{Vars}(G_1, \dots, G_m). G_1 \wedge \dots \wedge G_m$$

Due to the cunning way \mathcal{H} is defined, it will follow that the goal G_1, \dots, G_m is derivable.

Jacques Herbrand (1908–1931)



The Herbrand universe

A term is said to be *ground* if it contains no variables.

$$\begin{aligned} \textit{ground_term} ::= & \textit{constant} \\ & | \textit{fn_symbol}(\textit{ground_term_list}) \end{aligned}$$
$$\begin{aligned} \textit{ground_term_list} ::= & \textit{ground_term} \\ & | \textit{ground_term}, \textit{ground_term_list} \end{aligned}$$

Note that just the variables have been omitted from the grammar for terms from Lecture 3.

The *Herbrand universe* is just the set of all ground terms.

(We need to ensure that there is at least one constant symbol in our vocabulary in order that the Herbrand universe is non-empty.)

The minimal Herbrand model

We define the structure \mathcal{H} as follows.

- ▶ The universe is the Herbrand universe.
- ▶ A constant c is interpreted by $c^{\mathcal{H}} = c$.
- ▶ A function symbol f/k is interpreted by $f^{\mathcal{H}}(u_1, \dots, u_k) = f(u_1, \dots, u_k)$.
- ▶ A predicate symbol p/k is interpreted by

$$p^{\mathcal{H}}(u_1, \dots, u_k) = \mathbf{true} \Leftrightarrow \text{the goal } p(u_1, \dots, u_k) \text{ is derivable}$$

(N.B., we are using u_1, u_2, \dots to range over ground terms.)

Technical observation

Consider any atomic formula $p(t_1, \dots, t_k)$ with $\text{Vars}(t_1, \dots, t_k) = X_1, \dots, X_l$, and let u_1, \dots, u_l be ground terms.

By definition of $p^{\mathcal{H}}$, we have that

$$\mathcal{H} \models_{[X_1=u_1, \dots, X_l=u_l]} p(t_1, \dots, t_k)$$

holds if and only if the ground atomic goal below has a derivation.

$$p(t_1, \dots, t_k)\{X_1 = u_1, \dots, X_l = u_l\}$$

Proof of completeness (started)

Thus far, we have defined the structure \mathcal{H} .

The next step is to show that \mathcal{H} is indeed a model of the program.
That is, we show that

$$\mathcal{H} \models \forall \text{Vars}(F_i). F_i$$

for every F_i in F_1, \dots, F_n .

The detailed argument is given in the Appendix of this lecture.

Now, since

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

and $\mathcal{H} \models \forall \text{Vars}(F_1). F_i$, for every F_i in F_1, \dots, F_n , it follows that:

$$\mathcal{H} \models \exists \text{Vars}(G_1, \dots, G_m). G_1 \wedge \dots \wedge G_m$$

Proof of completeness (completed)

We have:

$$\mathcal{H} \models \exists \text{Vars}(G_1, \dots, G_m). G_1 \wedge \dots \wedge G_m$$

Let Y_1, \dots, Y_l be $\text{Vars}(G_1, \dots, G_m)$. By definition of satisfaction, there exist ground terms u_1, \dots, u_l such that

$$\mathcal{H} \models_{[Y_1:=u_1, \dots, Y_l:=u_l]} G_1 \wedge \dots \wedge G_m$$

So, by the technical observation, each of the atomic formulas $G_i\theta$, where θ is the substitution $\{Y_1 = u_1, \dots, Y_l = u_l\}$, is individually derivable as a ground atomic goal.

By technical lemma 1, it follows that the amalgamated goal

$$G_1\theta, \dots, G_m\theta, \text{ which is the same as } (G_1, \dots, G_m)\theta$$

is derivable. Whence, by technical lemma 2, so is

$$G_1, \dots, G_m$$

Complete proof search

We have a complete inference system for programs consisting of definite clauses in predicate logic, but Prolog's proof search procedure is incomplete.

As in propositional logic, incompleteness can be remedied by changing the search strategy.

For example, a complete procedure is obtained by making both the following two changes.

1. Adopt a breadth-first search strategy.
2. Allow the search to choose *any* goal G_i from the current goal list G_1, \dots, G_m (instead of always choosing G_1)

This provides a complete proof strategy, but *not* a decision procedure. The search does not terminate in cases in which a goal is not provable.

Undecidability

A *decision procedure* for definite clause predicate logic would be an algorithm that, given a goal and a program, outputs **yes** if the goal is a logical consequence of the program, and **no** otherwise. Note that a decision procedure is required to always terminate with one answer or the other.

Unlike in propositional logic, it is *impossible* to provide a decision procedure for definite clause predicate logic.

One can show that the *halting problem* (the question of deciding whether the execution of a *Turing machine* eventually halts) can be encoded as a query in definite clause predicate logic.

In fundamental work, in 1936, Alan Turing showed that the halting problem is an *undecidable problem*

(The undecidability of the halting problem is covered in detail in UG3 “Computability and Intractability”)

Alan Turing (1912–1954)



Conclusions (reprising Lecture 2)

Prolog proof search is an example of good engineering design based on an interplay between theoretical and practical considerations.

- ▶ It has a strong theoretical foundation, in being based on a complete inference system for definite-clause predicate logic.
- ▶ However, its incomplete proof search procedure is a good implementation choice, allowing efficient proof search in the context of predicate logic, and permitting the programmer to tailor programs with regard to efficiency issues.

Main points today

inference system (SLD resolution) for definite clause predicate logic

soundness of inference system

completeness of inference system

incompleteness of Prolog proof search strategy

Herbrand universe

minimal Herbrand model

undecidability of definite clause predicate logic

Appendix: The minimal Herbrand model is a model

Suppose F_i is $A_1 \wedge \dots \wedge A_k \rightarrow B$ with $\text{Vars}(F_i) = x_1, \dots, x_l$.
To verify that $\mathcal{H} \models \forall \text{Vars}(F_i). F_i$, we must show that, for all ground terms u_1, \dots, u_l it holds that

$$\mathcal{H} \models_{[x_1:=u_1, \dots, x_l:=u_l]} A_1 \wedge \dots \wedge A_k \rightarrow B$$

Suppose then that $\mathcal{H} \models_{[x_1:=u_1, \dots, x_l:=u_l]} A_j$ for every $j \in \{1, \dots, k\}$.
By the technical observation, every ground atomic goal

$$A_j\{x_1 = u_1, \dots, x_l = u_l\}$$

is derivable. Hence, applying technical lemma 1, we can derive the premise whence conclusion of the inference rule below

$$\frac{(A_1, \dots, A_k)\{x_1 = u_1, \dots, x_l = u_l\} \quad F_i}{B\{x_1 = u_1, \dots, x_l = u_l\}}$$

Now, by our technical observation, we have as required:

$$\mathcal{H} \models_{[x_1:=u_1, \dots, x_l:=u_l]} B$$