

## First-Order Logic and Logic Programming

The programming languages based on the Lambda calculus model of computation are convenient for describing functional relationships that hold between the objects of a problem. The focus is on computing functions. For this reason, the programs still preserve a procedural flavor: input data are transformed into output data by means of computations specific to the type of the processed data. Associative languages, such as CLIPS, are more permissive. There are no inputs and no outputs and data are uniformly represented in a symbolic way as facts of a factual base. The program infers all facts that can be derived from the current facts according to a fact-driven set of rules. However, facts have a very simple structure and cannot capture complex relationships between objects. Moreover, the inference procedure does not suit those applications where a specific goal has to be achieved regardless of other goals that could be potentially satisfied according to the current factual base.

Apart from these classes of languages, logic programming languages provide enhanced descriptive power and, moreover, are equipped with goal-oriented inference procedures proved complete and sound. The focus of a logic program is not restricted to the computation of the correct values of some attributes of the problem objects. The main focus is to derive new relationships between objects, starting from a set of given relationships. The theoretical base of such languages is the First-Order Logic (abbreviated FOL, and equivalently known as first-order predicate calculus).

### The syntax and semantics of first-order logic

In a similar way to Lambda calculus, first-order logic can be seen as a language in which sentences are constructed according to strict syntactic rules and have a given interpretation varying with the application domain. In what follows we prefer the term *universe of discourse* (UoD) instead of "solved problem". The reason relies on the descriptive power of predicate calculus. What is described is a generic set of relationships between generic objects that can be subject to different interpretations according to different applications. The UoD can be any problem, domain, and world the sentences of a FOL description may apply to. This is similar to the Markov algorithmic machine, where the algorithms are working with generic alphabets, constants and variables. The basic symbols used to build sentences are:

- Constants: stand for given (generic) objects from the UoD (similar to the constants in the Markov algorithmic machine). By convention, the constants are represented by identifiers starting with a small-case letter<sup>1</sup>; e.g., `pussycat` is a constant that can represent a particular cat or a special beloved person.
- Variables: stand for generic objects from the UoD. A variable is represented as an identifier starting with a capital letter.
- Functional and predicate symbols name relationships between the objects of the UoD. A functional identifier describes a functional relationship. A predicate corresponds to any kind of relation. Recall that an  $n$ -ary relation over a set  $s^n$  is a subset of tuples  $\langle x_1, x_2, \dots, x_n \rangle$ ,  $x_i \in s$  from  $s^n$ . For example, assume that `friend` is a binary relation `friend = {<fred, pussycat>, <bibo, pussycat>, <teo, pussycat>}` and notice that the constants can have different interpretations, as the relation itself.

A functional relation  $f$  is restricted: if a tuple  $\langle x, y \rangle$  is member of  $f$  then no other tuple  $\langle x, z \rangle$  with  $z$  different than  $y$  can be member of  $f$ .

---

<sup>1</sup> The naming conventions follow Prolog

The relation `friend` given above is a functional relation and corresponds to a non-injective function. The distinction between functions and predicates is important. Functions designate generic computations that do not have an expressional implementation, in the sense that a function is not associated with a specific code able to perform the computation. It simply declares what can be the result of a computation for given parameters. A predicate declares an unrestricted relationship that does not necessarily corresponds to a computation.

By convention, predicate and function identifiers start with a lower case letter. Moreover, if the tuple  $\langle x_1, x_2, \dots, x_n \rangle$  is an element of a relation `r` we write `r(x1, x2, ..., xn)`. In the binary case, the reading of `r(x, y)` is “`x` is in the relation `r` with `y`”, e.g. `friend(fred, pussycat)` can be read as *fred is a friend of pussycat*.

- Logical connectives:  $\Rightarrow$  (implication),  $\Leftrightarrow$  (bi-implication or equivalence),  $\wedge$  (conjunction),  $\vee$  (disjunction),  $\neg$  (negation). They have the usual meaning as described by well known truth tables. The precedence of these connectors is (from highest to lowest):  $\neg$ ,  $\wedge$ ,  $\vee$ ,  $\Rightarrow$ ,  $\Leftrightarrow$ .
- Quantifiers:  $\forall$  (the universal quantifier that reads *for all*),  $\exists$  (the existential quantifier that reads *exists*). As far as the precedence is concerned we assume that the quantifiers have lower precedence than the logical connectors.

Using the above symbols, *terms*, *atomic sentences* and *sentences* can be created according to the following syntax rules:

```
term ::= function(term,...,term) | constant | variable
atomic_sentence ::= predicate(term,...,term) | term = term
sentence ::= atomic_sentence | (sentence) | ¬sentence |
            sentence connective sentence |
            quantifier variable,...,variable. sentence
```

Sentences are considered generic. They specify relationships and properties of generic objects that are subject to different interpretations for particular UoDs. Depending on the UoD and the current interpretation these relationships and properties may be true or false.

The equality `term = term` is a way of declaring that the two terms stand for the same object. For example, `pussycat = margot` shows that `pussycat` is an alias name of a generic object identified as `margot` (a constant).

Terms refer objects, atomic sentences stand for relationships between objects, and sentences describe facts. For example, `father(fred)` is a term that under an appropriate interpretation stands for an unnamed object resulting from the computation `father` applied on the object identified as `fred`. The atomic sentence `friend(father(fred), pussycat)` shows that the father of `fred` is in relation `friend` with an object identified as `pussycat`.

Sentences that use quantifiers have enhanced descriptive power. For example, the sentence “everybody is the friend of a different somebody” can be encoded as:

$$\forall Y. \exists X. (\text{friend}(X, Y) \wedge \neg(Y = X))$$

The universal quantifier has the same effect as the binding symbol  $\lambda$  in an expression from the Lambda calculus. It generalizes the domain of a variable in the sense that any object can substitute the bound variable. On the contrary, the variable bound by an existential quantifier stands for a specific object, possibly not unique, which can be substituted for the variable. Variables cannot appear free within sentences. They must be bound either universally or existentially. The variables bound by a quantifier have the quantified sentence as their scope.

The quantifiers have useful properties, similar to the properties of conjunction and disjunction. Such properties are the laws of De Morgan. If  $P$  is a sentence, we have:

$\neg\forall x.P \equiv \exists x.\neg P$	is the transcription of $\neg(A \wedge B) \equiv \neg A \vee \neg B$
$\neg\exists x.P \equiv \forall x.\neg P$	is the transcription of $\neg(A \vee B) \equiv \neg A \wedge \neg B$
$\forall x.P \equiv \neg\exists x.\neg P$	is the transcription of $A \wedge B \equiv \neg(\neg A \vee \neg B)$
$\exists x.P \equiv \neg\forall x.\neg P$	is the transcription of $A \vee B \equiv \neg(\neg A \wedge \neg B)$

## Validity, satisfiability and proof

Assume that, in a description we have the following sentences:

- 1.1)  $\text{pet}(\text{pussycat})$  -- *pussycat is a pet*  
 1.2)  $\forall X, Y. \text{pet}(Y) \wedge \neg(Y = X) \Rightarrow \text{friend}(X, Y)$   
 -- *any pet is the friend of everyone different from that pet*  
 1.3)  $\neg(\text{fred} = \text{pussycat})$

One may ask if the above sentences are true for any universe of discourse or only for some particular universes of discourse. For instance, if the UoD is a neighborhood where *friend* and *pet* have the usual meaning, where pets are cared for and *fred* is a person while *pussycat* is a pet, the sentences are true. In some other UoD, where *fred* and *pussycat* stand for the same pet, sentence (1.4) is not true. We call *interpretation* the mapping between the constants, variables, and the functional and predicate symbols of a given FOL description and the objects and relationships specific to a given UoD. The relationship between the interpretation of sentences and their truth value is highlighted by the following definitions.

- A sentence is *valid* if and only if it is true under all possible interpretations (i.e. it is true in all possible UoDs). In other words, the sentence is true regardless its possible meanings.
- A sentence is *satisfiable* if and only if there is some interpretation, corresponding to some UoD, for which the sentence is true.

In the example above, all three sentences are satisfiable but not valid. A trivially valid sentence such as  $\text{friend}(\text{fred}, \text{pussycat}) \vee \neg \text{friend}(\text{fred}, \text{pussycat})$  is called a *tautology*.

What can we infer intuitively from the sentences above, if we consider them as being true? From sentences (1.1) and (1.3) we derive that, in a particular case,  $y$  can be bound to *pussycat* and  $x$  can be bound to *fred* in sentence (1.2). The sentence (1.2) yields:

(1.4)  $\text{friend}(\text{fred}, \text{pussycat})$

The intuitive derivation of the fact that *pussycat* has *fred* as a friend pinpoints important concepts that relate to the first-order logic (and to any logic system): entailment, inference, and proof. Call  $\text{Descr}$  a set of FOL sentences.

If a sentence  $s$  is true whenever the sentences of the description  $\text{Descr}$  are true we say that  $s$  is *entailed* by  $\text{Descr}$ , and we write  $\text{Descr} \models s$ .

For example, the sentence  $\text{friend}(\text{fred}, \text{pussycat})$  is true provided the sentences (1.1), (1.2) and (1.3) given above are true.

The mechanical derivation of sentences entailed by  $\text{Descr}$  is called *inference*. A mechanic procedure  $\text{inf}$  able to derive sentences entailed by  $\text{Descr}$  is an *inference procedure*. We write  $\text{Descr} \models_{\text{inf}} s$  to show that sentence  $s$  is derived from  $\text{Descr}$  using  $\text{inf}$ .

- If the inference procedure `inf` produces only sentences entailed by `Descr` then it is *sound* or *truth preserving*.
- If `inf` is able to generate all the sentences entailed by `Descr` then it is *complete*.

Given the sentences of a description `Descr`, we can prove or disprove a given sentence `s`. Inference procedures are at the base of the proof procedure. The *proof procedure* uses a given set of *inference procedures* (or inference rules) and a specific *control strategy* for conducting the inference. If the proof procedure is able to prove only sentences entailed by `Descr` it is sound. It is complete if it can prove all sentences entailed by `Descr`.

Gödel completeness theorem shows that in the first-order logic any sentence that is entailed by a set of sentences can be proved true from that set, i.e. there exists a complete proof procedure for the first-order logic. A complete and truth-preserving inference procedure, which can be used in a complete and sound proof procedure, has been found by Robinson in 1965, 35 years after Gödel. It is called *resolution*. Resolution and all other inference procedures for FOL are based on unification. Moreover, mechanical proof is eased if the sentences have a special format called *canonical* or *normal form*. In Prolog, a logic programming language, sentences are further restricted: they are Horn clauses.

## Unification

The process of unification of two formulae has been used in type synthesis. We recall it in the context of FOL. Two atomic sentences  $t_1$  and  $t_2$  unify if there exists a substitution  $s = \{v_1/x_1, v_2/x_2, \dots, v_n/x_n\}$  for the variables  $x_1, x_2, \dots, x_n$  within  $t_1$  and  $t_2$  such that after performing the substitution of  $x_i$  by  $v_i$  within both  $t_1$  and  $t_2$ ,  $t_1$  becomes the same as  $t_2$ . For example, observe that

$$t_1 = \text{friend}(\text{father}(X), Y) \text{ and } t_2 = \text{friend}(Z, \text{father}(Z))$$

unify under the substitution  $s = \{\text{father}(X)/Z, \text{father}(Z)/Y\}$ . This substitution is the most general unifier of  $t_1$  and  $t_2$ . Under the substitution  $s$  the atomic sentences are identical

$$t_1/s \equiv t_2/s \equiv \text{friend}(\text{father}(X), \text{father}(\text{father}(X)))$$

Unification is proved NP-complete. Unification can produce cyclic results. In the unification process of  $\text{friend}(X, \text{father}(X))$  with  $\text{friend}(\text{father}(Y), Y)$  produces the substitution  $s = \{\text{father}(Y)/X, \text{father}(X)/Y\}$  which means that  $x$  corresponds to  $\text{father}(\text{father}(X))$ . The unification algorithm must perform a so called *occurrence check* in order to avoid producing cyclic objects and, eventually, for avoiding the non termination of the unification process itself. The occurrence check searches for occurrences of a variable within the value bound to that variable as result of the unification process. This process is time costly.

## Normal forms

- A set of sentences is in conjunctive normal form (CNF) if each sentence is a disjunction of atoms. An atom is an atomic sentence eventually negated. All the sentences in the set are assumed to be implicitly connected by the conjunction connector (the CNF name derives from this convention).
- A set of sentences is in implicative normal form (INF) if each sentence has the structure  $(P \wedge Q \wedge \dots) \Rightarrow (P' \vee Q' \vee \dots)$ . Therefore, it is an implication with a conjunction of atoms on the left hand side and a disjunction of atoms on the right hand side.

There is a particular format of CNF sentences. This format is restrictive but plays a major role in logic programming.

- A sentence is a Horn sentence (or a Horn clause) if it is a disjunction of atoms with only one positive atom (non-negated atom). The sentence  $\neg a_1 \vee \neg a_2 \dots \vee \neg a_n \vee a$  is a Horn sentence and corresponds to the implication  $(a_1 \wedge a_2 \wedge \dots \wedge a_n) \Rightarrow a$ .

In both CNF and INF normal forms variables appear non-bound. They are considered universally bound implicitly. As an exercise, consider the sentences below and their CNF and INF equivalent.

$$\forall X, Y. (\text{father}(X, Y) \vee \text{mother}(X, Y) \Rightarrow \text{parent}(X, Y))$$

$$\forall X, Y, Z. (\text{parent}(Z, X) \wedge \text{parent}(Z, Y) \Rightarrow \text{sibling}(X, Y))$$

They can be rewritten as the following Horn clauses:

**CNF:**  $\neg \text{father}(X, Y) \vee \text{parent}(X, Y)$   
 $\neg \text{mother}(X, Y) \vee \text{parent}(X, Y)$   
 $\neg \text{parent}(Z, X) \vee \neg \text{parent}(Z, Y) \vee \text{sibling}(X, Y)$

**INF:**  $\text{father}(X, Y) \Rightarrow \text{parent}(X, Y)$   
 $\text{mother}(X, Y) \Rightarrow \text{parent}(X, Y)$   
 $\text{parent}(Z, X) \wedge \text{parent}(Z, Y) \Rightarrow \text{sibling}(X, Y)$

Any FOL sentence can be converted into an equivalent set of sentences in normal form. Here the term *equivalent* means *with the same meaning*. Consider that the sentence converted is part of a set of sentences called  $\text{DESCR}$ . The rules for converting a sentence into a set of sentences in CNF are:

1. Eliminate implications. A sentence  $p \Rightarrow q$  is replaced by  $\neg p \vee q$ .
2. Move negation inwards using De Morgan rules.
  - $\neg(p \vee q)$  is replaced by  $\neg p \wedge \neg q$
  - $\neg(p \wedge q)$  is replaced by  $\neg p \vee \neg q$
  - $\neg \forall x. P$  is replaced by  $\exists x. \neg P$
  - $\neg \exists x. P$  is replaced by  $\forall x. \neg P$
3. Rename variables. For sentences that bind the same variable more than once, the name of the variable is changed to a unique name for each separate binding. For instance,  $\forall x. P(x) \vee \forall x. Q(x) \wedge \exists x. R(x)$  is transformed into  $\forall x. P(x) \vee \forall x'. Q(x') \wedge \exists x''. R(x'')$ .
4. Move quantifiers to the top of the sentence preserving the order in which they occur. Since the variables they bind are distinct, the sentence preserves its meaning. A sentence that has all the quantifiers at the top is in *prenex form*.
5. Remove the existential quantifiers.
  - For an existential quantifier that is not preceded by any universal quantifier simply replace the occurrences of the bound variable by a unique constant (not used in  $\text{DESCR}$ ).
  - For an existential quantifier that is in the scope of several universal quantifiers, replace the occurrences of an existentially bound variable by a Skolem term  $f(x_1, x_2, \dots, x_n)$  where  $x_1, x_2, \dots, x_n$  are the variables bound by the universal quantifiers in the order they appear at the head of the sentence. The function  $f$  (not used elsewhere in  $\text{DESCR}$ ) designates a function which “computes” a unique value for the existentially bound variable depending on the values of  $x_1, x_2, \dots, x_n$ .

For example, in the sentence  $\forall Y. \text{pet}(Y) \Rightarrow \exists X. \text{friend}(X, Y)$  equivalent to  $\forall Y. \neg \text{pet}(Y) \vee \exists X. \text{friend}(X, Y)$  the existential quantifier can be dropped by replacing all the occurrences of the variable  $x$  by the Skolem term  $a\_friend\_of(Y)$ . The sentence becomes

$$\forall Y. \neg \text{pet}(Y) \vee \text{friend}(a\_friend\_of(Y), Y)$$

6. Drop the universal quantifiers and pretend that all variables are universally quantified implicitly.

$$\neg \text{pet}(Y) \vee \text{friend}(a\_friend\_of(Y), Y)$$

7. Distribute  $\wedge$  over  $\vee$ :  $(x \wedge y) \vee z$  is replaced by  $(x \vee z) \wedge (y \vee z)$ .

8. Flatten nested conjunctions and disjunctions.  $(x \wedge y) \wedge z$  is replaced by  $x \wedge y \wedge z$  and  $(x \vee y) \vee z$  is replaced by  $x \vee y \vee z$ . The sentence is in CNF. One more step is necessary for obtaining the INF format.

9. For each disjunction, group the negated atoms  $\neg x_1 \vee \dots \vee \neg x_n$  and the positive atoms  $y_1 \vee \dots \vee y_m$  of the disjunction and replace the disjunction by the implication  $x_1 \wedge \dots \wedge x_n \Rightarrow y_1 \vee \dots \vee y_m$ . For example,  $\text{pet}(Y) \Rightarrow \text{friend}(a\_friend\_of(Y), Y)$ .

## Inference rules

A sound inference rule (or procedure) derives, starting from a given set of sentences, a new sentence entailed by the sentences in the set. The simplest inference rule is *modus ponens*.

$$\frac{P_1 \wedge P_2 \wedge \dots \wedge P_n \quad (q_1 \wedge q_2 \wedge \dots \wedge q_n) \Rightarrow q \quad (\text{unify}(p_i, q_i), i=1, n) = s}{\text{Subst}(s, q)}$$

The rule reads: if the sentences  $p_i$  are true and the antecedents  $q_i$  of the implication are true under the unification substitution  $s$  (that unify  $p_i$  and  $q_i$  for all  $i=1, n$ ) then the sentence  $\text{subst}(s, q)$  is true. Modus ponens is simple but is not complete. The reason is simple: not all sentences can be written as Horn sentences. Therefore, the rule cannot be applied and the sentences entailed by non-Horn sentences cannot be inferred.

A complete inference procedure is *resolution*. The generalized resolution rule for INF is:

$$\frac{P_1 \wedge \dots \wedge P_i \wedge \dots \wedge P_{n1} \Rightarrow r_1 \vee \dots \vee r_{n2} \quad S_1 \wedge \dots \wedge S_{n3} \Rightarrow q_1 \vee q_2 \vee \dots \vee q_k \vee \dots \vee q_{n4} \quad \text{unify}(p_i, q_k) = s}{\text{Subst}(s, P_1 \wedge \dots \wedge P_{i-1} \wedge P_{i+1} \wedge \dots \wedge P_{n1} \wedge S_1 \wedge \dots \wedge S_{n3} \Rightarrow r_1 \vee \dots \vee r_{n2} \vee q_1 \vee \dots \vee q_{k-1} \vee q_{k+1} \vee \dots \vee q_n)}$$

For the particular case of Horn clauses the resolution rule is:

$$\frac{(q_1 \wedge q_2 \wedge \dots \wedge q_{n1}) \Rightarrow q \quad (P_1 \wedge P_2 \wedge \dots \wedge q' \wedge \dots \wedge P_{n2}) \Rightarrow r \quad \text{unify}(q, q') = s}{\text{Subst}(s, P_1 \wedge P_2 \wedge \dots \wedge P_{n2} \wedge q_1 \wedge q_2 \wedge \dots \wedge q_{n1} \Rightarrow r)}$$

This last rule is used in Prolog as the base of the proof procedure. Proof procedures are built using inference rules the application of which is decided by specific control strategies. The

proof problem is: given a set of initial sentences, call it  $\text{Descr}^2$ , and a sentence  $P$ , find out if  $\text{Descr}$  can entail  $P$ . The sentences in  $\text{Descr}$  can be classified into axioms and theorems. An *axiom* specifies a basic fact of the UoD and cannot be derived from the rest of the sentences in  $\text{Descr}$ . A *theorem* is a sentence entailed by the axioms and other theorems of  $\text{Descr}$ . By convention, an axiom  $x$  that corresponds to the implication  $\text{true} \Rightarrow x$ , will be represented simply as  $x$ . It is also useful to observe that  $\neg x$  can be rewritten  $x \Rightarrow \text{false}$ .

## An example of proof by refutation

Consider the following FOL description, which is a modified version of an example from the book of Russell and Norvig “Artificial Intelligence - A Modern Approach”.

```

 $\forall P, W, Q. (\text{person}(P) \wedge \text{weapon}(W) \wedge \text{person}(Q) \wedge \text{hostile}(Q) \wedge \text{sells}(P, W, Q) \Rightarrow \text{criminal}(P))$ 
 $\forall M. (\text{missile}(M) \Rightarrow \text{weapon}(M))$ 
 $\forall W, P. (\text{person}(P) \wedge \text{owns}(P, W) \Rightarrow \exists Q. (\text{person}(Q) \wedge \text{sells}(Q, W, P)))$ 
 $\forall P. (\text{person}(P) \wedge \forall Q. (\text{person}(Q) \wedge \text{hates}(P, Q)) \Rightarrow \text{hostile}(P))$ 

owns(foo, m1)
missile(m1)
person(foo)
 $\forall P. \text{hates}(foo, P)$ 

```

Applying the algorithm for converting FOL sentences to CNF and INF, the description above can be rewritten as the following set of Horn clauses.

```

person(P)  $\wedge$  weapon(W)  $\wedge$  person(Q)  $\wedge$  hostile(Q)  $\wedge$  sells(P, W, Q)  $\Rightarrow$  criminal(P)
missile(M)  $\Rightarrow$  weapon(M)
person(P)  $\wedge$  owns(P, W)  $\Rightarrow$  person(seller(W, P))
person(P)  $\wedge$  owns(P, W)  $\Rightarrow$  sells(seller(W, P), W, P)
person(P)  $\wedge$  person(Q)  $\wedge$  hates(P, Q)  $\Rightarrow$  hostile(P)

owns(foo, m1)
missile(m1)
person(foo)
hates(foo, P)

```

Above,  $\text{seller}(W, P)$  is the Skolem term introduced by the elimination of the existential quantifier. The problem is to check if, according to the description, is there a criminal.

A complete proof procedure using resolution is by *refutation*. We consider that the goal  $g$  to be proved is not true and then try to obtain a contradiction using the sentences of the given description, i.e.  $(\text{Descr} \wedge \neg g \Rightarrow \text{false}) \Leftrightarrow (\text{Descr} \Rightarrow g)$ . Hence, assume there is no criminal, i.e.  $\neg \text{criminal}(x)$  or, equivalently,  $\text{criminal}(x) \Rightarrow \text{false}$  and add this clause to the description above.

The inference rule applied is the resolution for Horn clauses. The order of resolution steps follows a backward chaining path, from the goal to the axioms. During the proof a current substitution called  $s$  is constantly updated. It cumulates the substitutions resulted from all the valid unifying processes performed so far along the current path traversed in the AND/OR proof tree. Recall that  $\forall x$  means the variable  $x$  is bound to  $\forall$  and notice that  $\forall$  can be a variable. All substitutions are performed in situ, although this is not happening in the real proof process (it will be difficult to backtrack). For each use of a Horn clause new variables are created for that clause. The resolvent is written in italic.

---

<sup>2</sup>  $\text{Descr}$  is considered a FOL description of a UoD, which can be generic or particular.

negated goal:  $\text{criminal}(X) \Rightarrow \text{false}$

$\text{criminal}(X) \Rightarrow \text{false}$

$\text{person}(X) \wedge \text{weapon}(Y) \wedge \text{person}(Q) \wedge \text{hostile}(Q) \wedge \text{sells}(P, W, Q) \Rightarrow \text{criminal}(P)$

$S = \{X/P\}$

---

$\text{person}(X) \wedge \text{weapon}(W) \wedge \text{person}(Q) \wedge \text{hostile}(Q) \wedge \text{sells}(X, W, Q) \Rightarrow \text{false}$

$\text{person}(\text{foo})$

$S = \{\text{foo}/Q, X/P\}$

---

$\text{person}(X) \wedge \text{weapon}(W) \wedge \text{hostile}(\text{foo}) \wedge \text{sells}(X, W, \text{foo}) \Rightarrow \text{false}$

$\text{missile}(M) \Rightarrow \text{weapon}(M)$

$S = \{M/W, \text{foo}/Q, X/P\}$

---

$\text{person}(X) \wedge \text{missile}(M) \wedge \text{hostile}(\text{foo}) \wedge \text{sells}(X, M, \text{foo}) \Rightarrow \text{false}$

$\text{missile}(m1)$

$S = \{m1/M, m1/W, \text{foo}/Q, X/P\}$

---

$\text{person}(X) \wedge \text{hostile}(\text{foo}) \wedge \text{sells}(X, m1, \text{foo}) \Rightarrow \text{false}$

$\text{person}(P1) \wedge \text{owns}(P1, W1) \Rightarrow \text{person}(\text{seller}(W1, P1))$

$S = \{\text{seller}(W1, P1)/X, m1/M, m1/W, \text{foo}/Q, \text{seller}(W1, P1)/P\}$

---

$\text{person}(P1) \wedge \text{owns}(P1, W1) \wedge \text{hostile}(\text{foo}) \wedge \text{sells}(\text{seller}(W1, P1), m1, \text{foo}) \Rightarrow \text{false}$

$\text{owns}(\text{foo}, m1)$

$S = \{\text{foo}/P1, m1/W1, \text{seller}(m1, \text{foo})/X, m1/M, m1/W, \text{foo}/Q, \text{seller}(m1, \text{foo})/P\}$

---

$\text{person}(\text{foo}) \wedge \text{hostile}(\text{foo}) \wedge \text{sells}(\text{seller}(m1, \text{foo}), m1, \text{foo}) \Rightarrow \text{false}$

$\text{person}(\text{foo})$

$S = \{\text{foo}/P1, m1/W1, \text{seller}(m1, \text{foo})/X, m1/M, m1/W, \text{foo}/Q, \text{seller}(m1, \text{foo})/P\}$

---

$\text{hostile}(\text{foo}) \wedge \text{sells}(\text{seller}(m1, \text{foo}), m1, \text{foo}) \Rightarrow \text{false}$

$\text{person}(P2) \wedge \text{person}(Q2) \wedge \text{hates}(P2, Q2) \Rightarrow \text{hostile}(P2)$

$S = \{\text{foo}/P2, \text{foo}/P1, m1/W1, \text{seller}(m1, \text{foo})/X, m1/M, m1/W, \text{foo}/Q, \text{seller}(m1, \text{foo})/P\}$

---

$\text{person}(Q2) \wedge \text{hates}(\text{foo}, Q2) \wedge \text{sells}(\text{seller}(m1, \text{foo}), m1, \text{foo}) \Rightarrow \text{false}$

$\text{hates}(\text{foo}, P3)$

$S = \{P3/Q2, \text{foo}/P2, \text{foo}/P1, m1/W1, \text{seller}(m1, \text{foo})/X, m1/M, m1/W, \text{foo}/Q, \text{seller}(m1, \text{foo})/P\}$

---

$\text{person}(P3) \wedge \text{sells}(\text{seller}(m1, \text{foo}), m1, \text{foo}) \Rightarrow \text{false}$

$\text{person}(\text{foo})$

$S = \{\text{foo}/P3, \text{foo}/Q2, \text{foo}/P2, \text{foo}/P1, m1/W1, \text{seller}(m1, \text{foo})/X, m1/M, m1/W, \text{foo}/Q, \text{seller}(m1, \text{foo})/P\}$

---

$\text{sells}(\text{seller}(m1, \text{foo}), m1, \text{foo}) \Rightarrow \text{false}$

$\text{person}(P4) \wedge \text{owns}(P4, W4) \Rightarrow \text{sells}(\text{seller}(W4, P4), W4, P4)$

$S = \{\text{foo}/P4, m1/W4, \text{foo}/P3, \text{foo}/Q2, \text{foo}/P2, \text{foo}/P1, m1/W1, \text{seller}(m1, \text{foo})/X, m1/M, m1/W, \text{foo}/Q, \text{seller}(m1, \text{foo})/P\}$

---

$\text{person}(\text{foo}) \wedge \text{owns}(\text{foo}, m1) \Rightarrow \text{false}$

$\text{person}(\text{foo})$

$S = \{\text{foo}/P4, m1/W4, \text{foo}/P3, \text{foo}/Q2, \text{foo}/P2, \text{foo}/P1, m1/W1, \text{seller}(m1, \text{foo})/X, m1/M, m1/W, \text{foo}/Q, \text{seller}(m1, \text{foo})/P\}$

---

$\text{owns}(\text{foo}, m1) \Rightarrow \text{false}$

$\text{owns}(\text{foo}, m1)$

$S = \{\text{foo}/P4, m1/W4, \text{foo}/P3, \text{foo}/Q2, \text{foo}/P2, \text{foo}/P1, m1/W1, \text{seller}(m1, \text{foo})/X, m1/M, m1/W, \text{foo}/Q, \text{seller}(m1, \text{foo})/P\}$

---

$\text{true} \Rightarrow \text{false}$

contradiction



Therefore, there is a criminal in the story. Who? Look at the value bound to the variable  $x$  of the goal `criminal(x)`. Its value is `seller(m1,foo)`. Although we don't know precisely who that person is, we have been able to prove that such a person exists.

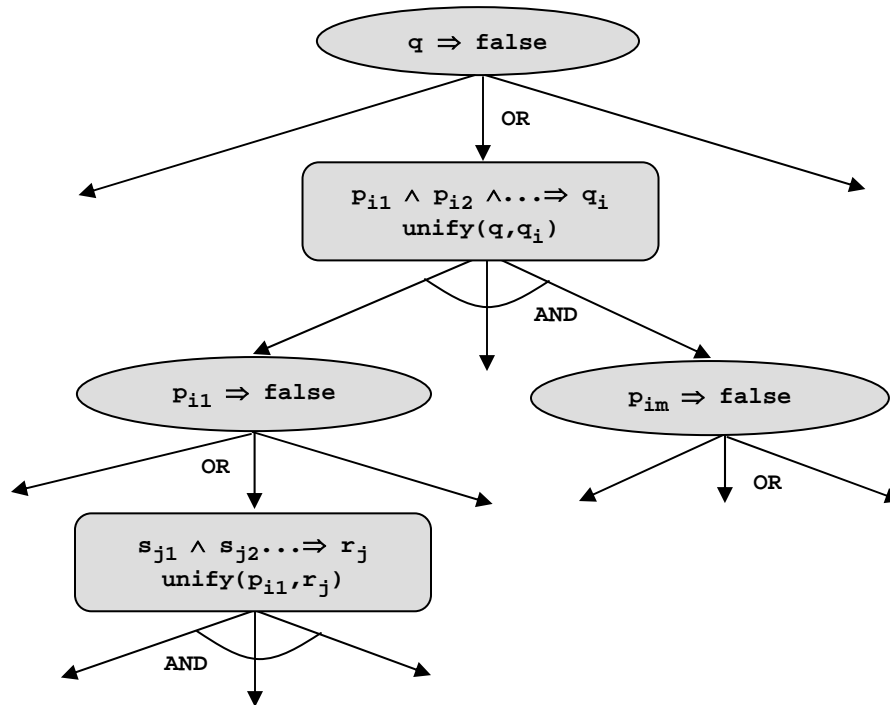


Figure 1. A proof AND-OR tree

In the presence of a control procedure the proof is no longer a linear structure as above. Some paths followed in the process of proof may fail. In this case backtracking occurs in order to try other possible resolvents. The structure of the proof process is an AND/OR tree. An AND node corresponds to a conjunction of resolvents. An OR node corresponds to the choice of the sentence to be used for producing a new resolvent.

To have a complete proof procedure the tree traversal should be breadth-first oriented. The depth-first traversal is making the proof procedure incomplete due to possible infinite paths that are created by circular sentences. It is interesting to notice that in the example above there are such infinite paths created by the circular clause:

$$\text{person}(P) \wedge \text{owns}(P,W) \Rightarrow \text{person}(\text{seller}(W,P))$$

In a depth-first traversal of the proof tree we could enter an infinite loop trying to resolve the goal `person(P)`.

```

person(P) => false
person(P1) ^ owns(P1,W1) => person(seller(W1,P1))
                                                                    S={seller(W1,P1)/P}
-----
person(P1) ^ owns(P1,W1) => false
person(P2) ^ owns(P2,W2) => person(seller(W2,P2))
                                                                    S={seller(W2,P2)/P1,seller(W1,seller(W2,P2))/P}
-----
person(P2) ^ owns(P2,W2) ^ owns(P1,W1) => false
. . .
    
```

## Control strategies

There are two main control strategies: forward chaining and backward chaining. Forward chaining is similar to the `Agenda` control mechanism in CLIPS. It allows an inference procedure to derive all sentences that can be derived starting from `Descr`. It stops when the goal to be proved is eventually derived. Backward chaining focuses the application of the inference rule only on the relevant part of `Descr`, i.e. on these sentences that can effectively contribute to the derivation of the given goal.

The backward-chaining algorithm for Horn clauses resolution is given in Caml below. Except for the unification, it is close the executable kernel of a very simple theorem prover. The `Backward_chaining` function works with sentences that do not contain variables. As an exercise, the `unify` function could be rewritten to work with sentences containing variables.

The function `Backward_chaining` works with a list of rules, the rules from `Descr`. Each rule is represented as a list; the head of the list contains the conclusion of the rule, the tail contains the premises. The result returned is `fail` or `succeed(substitution)`, where `substitution` cumulates all the substitutions resulted from the unification process performed throughout the inference process. In a simple (but inefficient) implementation `substitution` can be an association list containing pairs `(variable,value)`.

The control engine is the local function `solve`. It gets a list of goals, a list of Horn clauses (called here rules), and the current substitution. If the list of goals is exhausted, i.e. all goals are proved, the result is `succeed(final_substitution)`. If the all rules are tried without succeeding to prove the goals, `fail` is signaled. For each goal, `solve` tries all rules the conclusion of which can unify with the goal (recall that the goal stands for a negated sentence while the conclusion is a positive sentence). For each applicable rule, the premises of the rule are appended at the end of the list of goals and `solve` is called recursively with the new goals, all the rules in `Descr`, and the new substitution. In this way, the proof tree is traversed breadth-first.

```

type 'subst Result = fail | succeed of 'subst list;;

let unify(atom1,atom2,substitution) =
  if atom1 = atom2 then succeed(substitution) else fail;;
unify : 'a * 'a * 'b list -> 'b Result = <fun>

let conclusion = hd and premises = tl;;
conclusion : 'a list -> 'a = <fun>
premises : 'a list -> 'a list = <fun>

let Backward_chaining all_rules goal =
  let rec solve =
    fun [] rules substitution -> succeed(substitution)
    | goals [] substitution -> fail
    | (goal::rest_goals as goals) (rule::rest_rules) substitution ->
      match unify(goal,conclusion(rule),substitution)
      with
      fail -> solve goals rest_rules substitution
      | succeed(new_substitution) ->
        match solve (rest_goals@(premises rule)) all_rules
        new_substitution
        with
        fail -> solve goals rest_rules substitution
        | succeed(new_subst) -> succeed(new_subst)
  in
  solve [goal] all_rules [];;
Backward_chaining : 'a list list -> 'a -> 'b Result = <fun>

```

Observe that if `Backward_chaining` would be a coroutine, we could save the status of the proof tree while returning a solution. The traversal of the tree could be resumed from the interruption point producing step by step all the possible proofs of the given goal. This is what happens in Prolog.

The conclusion and the premises of each rule are atomic sentences represented as values of sum types. The representation is close to what a syntax checker would produce.

```

type Term =  constant of string
            |  var of string
            |  funct of string * Term list;;

type AtomicSentence = predicate of string * Term list;;

let Descr = [
  (* friend(bibo,pussycat) ^friend(pussycat,fred) => friend(bibo,fred) *)
  [predicate("friend",[constant "Bibo"; constant "Fred"]);
   predicate("friend",[constant "Bibo"; constant "Pussycat"]);
   predicate("friend",[constant "Pussycat"; constant "Fred"]);

  (* friend(bibo,pussycat) *)
  [predicate("friend",[constant "Bibo"; constant "Pussycat"]);

  (* friend(pussycat,fred) *)
  [predicate("friend",[constant "Pussycat"; constant "Fred"]);
];;

let goal = predicate("friend",[constant "Bibo"; constant "Fred"]);;

Backward_chaining Descr goal;;
- : '_a Result = succeed []

```

The backward-chaining control strategy combined with the refutation-based resolution - the proof procedure - offers an efficient proof engine for Horn clauses. There are logic programming languages (such as Prolog) that use this proof engine. What they add, apart from a convenient syntax, is a set of control mechanisms of the inference engine.

A program in a logic programming language consists of a set of axioms and inference rules that can be used to prove given goals. The axioms, rules, and goals are sentences that state properties and relationships of the objects of a given UoD. The aim of the program is merely to prove rather than compute. This is the reason why logic programming offers less on the side of arithmetic and more on the side of symbolic representation and on the way of combining formulae. A remarkable representative of the logic-programming paradigm is Prolog.