Outline

Unification and Lifting

Substitution

- A substitution θ is a mapping from variables to terms such that the set $\{x \mid \theta(x) \neq x\}$ is finite.
- This set is called the domain of θ.
- Notation: $\{x_1 \mapsto t_1, \ldots, x_n \mapsto t_n\}$, where $x_1, \ldots, x_n$ are pairwise different variables, denotes the substitution θ such that

$$\theta(x) = \begin{cases} t_i & \text{if } x = x_i; \\ x & \text{if } x \notin \{x_1, \ldots, x_n\}. \end{cases}$$

- Application of this substitution to an expression E: simultaneous replacement
of x_i by t_i.
- Application of a substitution θ to E is denoted by $E\theta$.
- Since substitutions are functions, we can define their composition (written $\sigma\tau$ instead of $\tau \circ \sigma$). Note that we have $E(\sigma\tau) = (E\sigma)\tau$.

Exercise

Exercise: Suppose we have two substitutions

$$\{x_1 \mapsto s_1, \ldots, x_m \mapsto s_m\} \text{ and}$$
$$\{y_1 \mapsto t_1, \ldots, y_n \mapsto t_n\}.$$

How can we write their composition using the same notation?

Instances, Ground

An instance of an expression (that is term, atom, literal, or clause) E is obtained by applying a substitution to E. Examples:

- some instances of the term $f(x, a, g(x))$ are:
  $f(x, a, g(x))$,
  $f(y, a, g(y))$,
  $f(a, a, g(a))$,
  $f(g(b), a, g(g(b)))$;
- but the term $f(b, a, g(c))$ is not an instance of this term.

Ground instance: instance with no variables.

Herbrand's Theorem

For a set of clauses S denote by S^* the set of ground instances of clauses in S.

Theorem Let S be a set of clauses. The following conditions are equivalent.

1. S is unsatisfiable;
2. S^* is unsatisfiable;

Herbrand's Theorem

For a set of clauses S denote by S^* the set of ground instances of clauses in S.

Theorem Let S be a set of clauses. The following conditions are equivalent.

1. S is unsatisfiable;
2. S^* is unsatisfiable;

By compactness the last condition is equivalent to

3. there exists a finite unsatisfiable set of ground instances of clauses in S.

Herbrand's Theorem

For a set of clauses S denote by S^* the set of ground instances of clauses in S.

Theorem Let S be a set of clauses. The following conditions are equivalent.

1. S is unsatisfiable;
2. S^* is unsatisfiable;

By compactness the last condition is equivalent to

3. there exists a finite unsatisfiable set of ground instances of clauses in S.

The theorem reduces the problem of checking unsatisfiability of sets of arbitrary clauses to checking unsatisfiability of sets of ground clauses ...
The only problem is that S^* can be infinite even if S is finite.

Note on Herbrand's Theorem, Compactness and Completeness

The proofs of completeness of resolution and superposition with redundancy elimination does not use any of these theorems.

Note on Herbrand's Theorem, Compactness and Completeness

The proofs of completeness of resolution and superposition with redundancy elimination does not use any of these theorems.

Interestingly, they all can be derived as simple corollaries of this proof of completeness!

Lifting

Lifting is a technique for proving completeness theorems in the following way:

1. Prove completeness of the system for a set of ground clauses;
2. Lift the proof to the non-ground case.

Lifting, Example

Consider two (non-ground) clauses $p(x, a) \vee q_1(x)$ and $\neg p(y, z) \vee q_2(y, z)$. If the signature contains function symbols, then both clauses have infinite sets of instances:

$$\begin{array}{rcl} \{p(r, a) \vee q_1(r) & \mid & r \text{ is ground}\} \\ \{\neg p(s, t) \vee q_2(s, t) & \mid & s, t \text{ are ground}\} \end{array}$$

We can resolve such instances if and only if $r = s$ and $t = a$. Then we can apply the following inference

$$\frac{p(s, a) \vee q_1(s) \quad \neg p(s, a) \vee q_2(s, a)}{q_1(s) \vee q_2(s, a)} \text{ (BR)}$$

But there is an infinite number of such inferences.

Lifting, Idea

The idea is to represent an infinite number of ground inferences of the form

$$\frac{p(s,a) \vee q_1(s) \quad \neg p(s,a) \vee q_2(s,a)}{q_1(s) \vee q_2(s,a)} \text{ (BR)}$$

by a single non-ground inference

$$\frac{p(x,a) \vee q_1(x) \quad \neg p(y,z) \vee q_2(y,z)}{q_1(y) \vee q_2(y,a)} \text{ (BR)}$$

Is this always possible?

Yes!

$$\frac{p(x,a) \vee q_1(x) \quad \neg p(y,z) \vee q_2(y,z)}{q_1(y) \vee q_2(y,a)} \text{ (BR)}$$

Note that the substitution $\{x \mapsto y, z \mapsto a\}$ is a solution of the "equation" $p(x,a) = p(y,z)$.

What should we lift?

- Ordering $\succ$;
- Selection function σ;
- Calculus $\mathbb{S}\mathrm{up}_{\succ,\sigma}$.

Most importantly, for the lifting to work we should be able to solve equations $s = t$ between terms and between atoms. This can be done using most general unifiers.

Unifier

Unifier of expressions s_1 and s_2: a substitution θ such that $s_1\theta = s_2\theta$. In other words, a unifier is a solution to an "equation" $s_1 = s_2$. In a similar way we can define solutions to systems of equations $s_1 = s'_1, \dots, s_n = s'_n$. We call such solutions simultaneous unifiers of $s_1, \dots, s_n$ and $s'_1, \dots, s'_n$.

(Most General) Unifiers

A solution θ to a set of equations E is said to be a most general solution if for every other solution σ there exists a substitution τ such that $\theta\tau = \sigma$. In a similar way can define a most general unifier.

(Most General) Unifiers

A solution θ to a set of equations E is said to be a most general solution if for every other solution σ there exists a substitution τ such that $\theta\tau = \sigma$. In a similar way can define a most general unifier.

Consider terms $f(x_1, g(x_1), x_2)$ and $f(y_1, y_2, y_2)$.
(Some of) their unifiers are
$\theta_1 = \{y_1 \mapsto x_1, y_2 \mapsto g(x_1), x_2 \mapsto g(x_1)\}$ and
$\theta_2 = \{y_1 \mapsto a, y_2 \mapsto g(a), x_2 \mapsto g(a), x_1 \mapsto a\}$:

$f(x_1, g(x_1), x_2)\theta_1 = f(x_1, g(x_1), g(x_1))$;
$f(y_1, y_2, y_2)\theta_1 = f(x_1, g(x_1), g(x_1))$;
$f(x_1, g(x_1), x_2)\theta_2 = f(a, g(a), g(a))$;
$f(y_1, y_2, y_2)\theta_2 = f(a, g(a), g(a))$.

But only θ_1 is most general.

Unification

Let E be a set of equations. An isolated equation in E is any equation $x = t$ in it such that x has exactly one occurrence in E.

input:
  A finite set of equations E
output:
  A solution to E or failure.
begin
  while there exists a non-isolated equation $(s = t) \in E$
  do
    case (s, t) **of**
      $(t, t) \Rightarrow$ Remove this equation from E
      $(x, t) \Rightarrow$
        if x occurs in t
        then halt with failure
        else replace x by t in all other equations of E
      $(t, x) \Rightarrow$ replace this equation by $x = t$
        and do the same as in the case (x, t)
      $(c, d) \Rightarrow$ halt with failure
      $(c, f(t_1, \ldots, t_n)) \Rightarrow$ halt with failure
      $(f(t_1, \ldots, t_n), c) \Rightarrow$ halt with failure
      $(f(s_1, \ldots, s_m), g(t_1, \ldots, t_n)) \Rightarrow$ halt with failure
      $(f(s_1, \ldots, s_n), f(t_1, \ldots, t_n)) \Rightarrow$ replace this equation by the set
        $s_1 = t_1, \ldots, s_n = t_n$
      end
    od
    Now E has the form $\{x_1 = r_1, \ldots, x_l = r_l\}$ and every equation in it is isolated
    return the substitution $\{x_1 \mapsto r_1, \ldots, x_l \mapsto r_l\}$
  end

Examples

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

Properties

Theorem Suppose we run the unification algorithm on $s = t$. Then

- If s and t are unifiable, then the algorithms terminates and outputs a most general unifier of s and t.
- If s and t are not unifiable, then the algorithms terminates with failure.

Notation (slightly ambiguous):

- $mgu(s, t)$ for a most general unifier;
- $mgs(E)$ for a most general solution.

Exercise

Consider a trivial system of equations $\{\}$ or $\{a = a\}$.
What is the set of solutions to it?
What is the set of most general solutions to it?

Properties

Theorem Let C be a clause and E a set of equations. Then

$$\{D \in C^* \mid \exists\theta(C\theta = D \text{ and } \theta \text{ is a solution to } E)\} = (C\,mgs(E))^*.$$

In other words, to find a set of ground instances of a clause C that also satisfy an equation E, take the most general solution σ of E and use ground instances of $C\sigma$.

Non-Ground Superposition Rule

Superposition:

$$\frac{\underline{l = r} \vee C \quad \underline{s[l'] = t} \vee D}{(s[r] = t \vee C \vee D)\theta}\ \text{(Sup)}, \qquad \frac{\underline{l = r} \vee C \quad \underline{s[l'] \neq t} \vee D}{(s[r] \neq t \vee C \vee D)\theta}\ \text{(Sup)},$$

where

1. θ is an mgu of l and l';
2. l' is not a variable;
3. $r\theta \not\succeq l\theta$;
4. $t\theta \not\succeq s[l']\theta$.
5. …

Non-Ground Superposition Rule

Superposition:

$$\frac{\underline{l = r} \vee C \quad \underline{s[l'] = t} \vee D}{(s[r] = t \vee C \vee D)\theta}\ (\text{Sup}), \qquad \frac{\underline{l = r} \vee C \quad \underline{s[l'] \neq t} \vee D}{(s[r] \neq t \vee C \vee D)\theta}\ (\text{Sup}),$$

where

1. θ is an mgu of l and l';
2. l' is not a variable;
3. $r\theta \not\succeq l\theta$;
4. $t\theta \not\succeq s[l']\theta$.
5. ...

Observations:

- ordering is partial, hence conditions like $r\theta \not\succeq l\theta$;
- these conditions must be checked a posteriori, that is, after the rule has been applied.

Note, however, that $l \succ r$ implies $l\theta \succ r\theta$, so checking orderings a priory helps.

More rules

Equality Resolution:

$$\frac{\underline{s \neq s'} \lor C}{C\theta} \text{ (ER)},$$

where θ is an mgu of s and s'.

Equality Factoring:

$$\frac{\underline{l = r} \lor l' = r' \lor C}{(l = r \lor r \neq r' \lor C)\theta} \text{ (EF)},$$

where θ is an mgu of l and l', $r\theta \not\succeq l\theta$, $r'\theta \not\succeq l\theta$, and $r'\theta \not\succeq r\theta$.