Math 260A — Mathematical Logic — Scribe Notes
UCSD — Winter Quarter 2012
Instructor: Sam Buss

Notes by: Andy Parrish
Friday, March 16, 2012

1 Robinson resolution refutation

Let Γ be a set of clauses of first order literals — the terms have the form $P(t_1, \ldots, t_k)$ or $\neg P(t_1, \ldots, t_k)$ for terms $t_1, \ldots, t_k$ and k-ary function P. Without loss of generality, we will assume the clauses in Γ use distinct variables (though terms within a clause cannot be assumed distinct).

Throughout these notes, we will assume the language L contains at least one constant symbol.

Definition A ground resolution refutation of Γ is a sequence of clauses $C_1, C_2, \ldots, C_k = \emptyset$ where each C_i is either a ground instance of a clause in Γ, or is inferred by a resolution inference from two previous clauses C_j and C_ℓ.

Definition A Robinson resolution refutation of Γ is a sequence of clauses $C_1, C_2, \ldots, C_k = \emptyset$ where each C_i is either a relabeling[1] of a clause in Γ, or is obtained by a Robinson resolution inference from two previous clauses C_j and C_ℓ.

To define a Robinson resolution inference, take two sets of clauses A and B, and nonempty subsets $A' \subset A$, $B' \subset B$, where A' contains only positive clauses, and B' has only negative clauses. Let

$$F = \{\varphi \mid \varphi \in A'\} \cup \{\varphi \mid \neg\varphi \in B'\}.$$

Choose an mgu σ unifying F, so that $\varphi\sigma = P(t_1, \ldots, t_k)$ for every $\varphi \in F$. If such a σ exists, then we make this resolution inference:

$$\frac{A\sigma \qquad B\sigma}{C}\text{ Resolution}$$

where $C = (A \setminus A')\sigma \cup (B \setminus B')\sigma$.

[1] Traditionally a Robinson resolution does not allow for relabeling the variables in a clause. We allow it here as it does not add any power, but removes some technical concerns from the upcoming proof.

For C determined from A and B by such an inference, we have

$$\frac{A \qquad B}{C} \text{ Robinson resolution}$$

The selection of A' and B' is called *factoring.*

Note At first glance, this inference rule may seem needlessly complex. Why not do resolution on individual terms of the clause? There's a good reason: such inferences are not complete. Here is a simple example where things go wrong.

$\Gamma = \big\{\{P(x), P(y)\}, \{\neg P(u), \neg P(v)\}\big\}$. Γ corresponds to the sentence

$$(\forall x \forall y P(x) \vee P(y)) \wedge (\forall x \forall y \neg P(u) \vee \neg P(v)).$$

This is clearly unsatisfiable,

However, the only inference possible from these clauses (up to variable names) is to resolve $P(x)$ against $\neg P(u)$ (after appropriate unification), which leaves us with the resolvent $\{P(y), \neg P(v)\}$, which corresponds to the sentence

$$\forall y \forall v P(x) \vee \neg P(y),$$

which is a tautology, and thus not any help.

1.1 Relation to ground resolution refutation

Theorem If Γ has a ground resolution refutation, then Γ has a Robinson resolution refutation.

Note This theorem saves us from choosing terms for the ground instances, instead requiring a good factoring strategy.

Proof Let $C_1, \ldots, C_k = \emptyset$ be a ground resolution refutation. Without loss of generality, we assume that $C_i \neq C_j$

We will find a Robinson resolution refutation $D_1, \ldots, D_k$ on distinct variables, and substitutions $\sigma_1, \ldots, \sigma_k$ such that $C_i = D_i \sigma_i$. In particular, $D_k = \emptyset$.

We will show that, if the above property holds for the initial sequence $C_1, \ldots, C_{i-1}$, then it also holds for $C_1, \ldots, C_i$.

Case 1 C_i is a ground instance of $C \in \Gamma$. Let D_i be an instance of C with new variables (not yet seen), so C_i is a substitution instance of D_i. Pick such a substitution σ_i.

Case 2 C_i is the resolvent of $C_j = D_j\sigma_j$ and $C_\ell = D_\ell\sigma_\ell$, with respect to $P(\mathbf{t})$. Select $D'_j = \{\varphi \in D_j \mid \varphi\sigma = P(\mathbf{t})\}$, and $D'_\ell = \{\varphi \in D_\ell \mid \varphi\sigma = \neg P(\mathbf{t})\}$. Let

$$F = \{\varphi \mid \varphi \in D'_j\} \cup \{\varphi \mid \neg\varphi \in D'_\ell\}.$$

Since the D's are chosen to have distinct variables, the domains of σ_j, σ_ℓ are disjoint.

By construction, $\sigma_j \cup \sigma_\ell$ unifies F, so F must have an mgu — call it τ — so that $\exists \pi, \tau\pi = \sigma_j \cup \sigma_\ell$. Choose such a τ which sends all variables in C_j and C_ℓ to a new set of unused variables.

Let D_i = Robinson resolvent = $(D_j \setminus D'_j)\tau \cup (D_\ell \setminus D'_\ell)\tau$.

Claim $C_i = D_i\pi$.

Proof

$$\begin{aligned}
\psi \in C_i \quad &\iff \quad \psi \in (C_j \setminus \{P(\mathbf{t})\}) \cup (C_\ell \setminus \{\neg P(\mathbf{t})\}) \\
&\iff \quad \text{Either } \exists \psi' \in D_j \setminus D'_j, \psi = \psi'\sigma_j, \text{ or } \exists \psi' \in D_\ell \setminus D'_\ell, \psi = \psi'\sigma_\ell \\
&\iff \quad \exists \psi' \in (D_j \setminus D'_j) \cup (D_\ell \setminus D'_\ell), \psi = \psi'(\sigma_j \cup \sigma_\ell) = \psi'\tau\pi \\
&\iff \quad \exists \psi' \in D_i, \psi = \psi'\pi
\end{aligned}$$

which was the goal. Take $\sigma_i = \pi$, so $C_i = D_i\sigma_i$, completing the proof.