

3 Superposition

First-order calculi considered so far:

Resolution: for first-order clauses without equality.

(Unfailing) Knuth-Bendix Completion: for unit equations.

Goal:

Combine the ideas of ordered resolution (overlap maximal literals in a clause) and Knuth-Bendix completion (overlap maximal sides of equations) to get a calculus for equational clauses.

3.1 Recapitulation

Atom: either $P(s_1, \dots, s_m)$ with $P \in \Pi$ or $s \approx t$.

Literal: Atom or negated atom.

Clause: (possibly empty) disjunction of literals (all variables implicitly universally quantified).

For refutational theorem proving, it is sufficient to consider sets of clauses: every first-order formula F can be translated into a set of clauses N such that F is unsatisfiable if and only if N is unsatisfiable.

In the non-equational case, unsatisfiability can for instance be checked using the (ordered) resolution calculus.

(Ordered) resolution: inference rules:

	Ground case:	Non-ground case:
<i>Resolution:</i>	$\frac{D' \vee A \quad C' \vee \neg A}{D' \vee C'}$	$\frac{D' \vee A \quad C' \vee \neg A'}{(D' \vee C')\sigma}$ <p style="text-align: center;">where $\sigma = \text{mgu}(A, A')$.</p>
<i>Factoring:</i>	$\frac{C' \vee A \vee A}{C' \vee A}$	$\frac{C' \vee A \vee A'}{(C' \vee A)\sigma}$ <p style="text-align: center;">where $\sigma = \text{mgu}(A, A')$.</p>

Ordering restrictions:

Let \succ be a well-founded and total ordering on ground atoms.

Literal ordering \succ_L : compares literals by comparing lexicographically first the respective atoms using \succ and then their polarities (negative $>$ positive).

Clause ordering \succ_C : compares clauses by comparing their multisets of literals using the multiset extension of \succ_L .

Ordering restrictions (ground case):

Inference are necessary only if the following conditions are satisfied:

- The left premise of a Resolution inference is not larger than or equal to the right premise.
- The literals that are involved in the inferences ($[\neg] A$) are maximal in the respective clauses (strictly maximal for the left premise of Resolution).

Ordering restrictions (non-ground case):

Lift the ground ordering to non-ground literals: A literal L is called [strictly] maximal in a clause C if and only if there exists a ground substitution σ such that for all other literals L' in C : $L\sigma \not\prec L'\sigma$ [$L\sigma \not\prec L'\sigma$].

Resolution is (even with ordering restrictions) refutationally complete:

Dynamic view of refutational completeness:

If N is unsatisfiable ($N \models \perp$) then *fair* derivations from N produce \perp .

Static view of refutational completeness:

If N is *saturated*, then N is unsatisfiable if and only if $\perp \in N$.

Proving refutational completeness for the ground case:

We have to show:

If N is saturated (i. e., if sufficiently many inferences have been computed), and $\perp \notin N$, then N is satisfiable (i. e., has a model).

Constructing a candidate interpretation:

Suppose that N be saturated and $\perp \notin N$. We inspect all clauses in N in ascending order and construct a sequence of Herbrand interpretations (starting with the empty interpretation: all atoms are false).

If a clause C is false in the current interpretation, and has a positive and strictly maximal literal A , then extend the current interpretation such that C becomes true: add A to the current interpretation. (Then C is called *productive*.)

Otherwise, leave the current interpretation unchanged.

The sequence of interpretations has the following properties:

- (1) If an atom is true in some interpretation, then it remains true in all future interpretations.
- (2) If a clause is true at the time where it is inspected, then it remains true in all future interpretations.
- (3) If a clause $C = C' \vee A$ is productive, then C remains true and C' remains false in all future interpretations.

Show by induction: if N is saturated and $\perp \notin N$, then every clause in N is either true at the time where it is inspected or productive.

Note:

For the induction proof, it is not necessary that the conclusion of an inference is contained in N . It is sufficient that it is redundant w. r. t. N .

N is called *saturated up to redundancy* if the conclusion of every inference from clauses in $N \setminus Red(N)$ is contained in $N \cup Red(N)$.

Proving refutational completeness for the non-ground case:

If $C_i\theta$ is a ground instance of the clause C_i for $i \in \{0, \dots, n\}$ and

$$\frac{C_n, \dots, C_1}{C_0}$$

and

$$\frac{C_n\theta, \dots, C_1\theta}{C_0\theta}$$

are inferences, then the latter inference is called a *ground instance* of the former.

For a set N of clauses, let $G_\Sigma(N)$ be the set of all ground instances of clauses in N .

Construct the interpretation from the set $G_\Sigma(N)$ of all ground instances of clauses in N :

- N is saturated and does not contain \perp
- $\Rightarrow G_\Sigma(N)$ is saturated and does not contain \perp
- $\Rightarrow G_\Sigma(N)$ has a Herbrand model I
- $\Rightarrow I$ is a model of N .

It is possible to encode an arbitrary predicate P using a function f_P and a new constant *true*:

$$\begin{aligned} P(t_1, \dots, t_n) &\rightsquigarrow f_P(t_1, \dots, t_n) \approx \text{true} \\ \neg P(t_1, \dots, t_n) &\rightsquigarrow \neg f_P(t_1, \dots, t_n) \approx \text{true} \end{aligned}$$

In equational logic it is therefore sufficient to consider the case that $\Pi = \emptyset$, i. e., equality is the only predicate symbol.

Abbreviation: $s \not\approx t$ instead of $\neg s \approx t$.

3.2 The Superposition Calculus – Informally

Conventions:

From now on: $\Pi = \emptyset$ (equality is the only predicate).

Inference rules are to be read modulo symmetry of the equality symbol.

We will first explain the ideas and motivations behind the superposition calculus and its completeness proof. Precise definitions will be given later.

Ground inference rules:

$$\text{Pos. Superposition: } \frac{D' \vee t \approx t' \quad C' \vee s[t] \approx s'}{D' \vee C' \vee s[t'] \approx s'}$$

$$\text{Neg. Superposition: } \frac{D' \vee t \approx t' \quad C' \vee s[t] \not\approx s'}{D' \vee C' \vee s[t'] \not\approx s'}$$

$$\text{Equality Resolution: } \frac{C' \vee s \not\approx s}{C'}$$

(Note: We will need one further inference rule.)

Ordering wishlist:

Like in resolution, we want to perform only inferences between (strictly) maximal literals.

Like in completion, we want to perform only inferences between (strictly) maximal sides of literals.

Like in resolution, in inferences with two premises, the left premise should not be larger than the right one.

The ordering should be total on ground literals.

Consequences:

The literal ordering must depend primarily on the larger term of an equation.

As in the resolution case, negative literals must be a bit larger than the corresponding positive literals.

Additionally, we need the following property: If $s \succ t \succ u$, then $s \not\approx u$ must be larger than $s \approx t$. In other words, we must compare first the larger term, then the polarity, and finally the smaller term.

The following construction has the required properties:

Let \succ be a *reduction ordering that is total on ground terms*.

To a positive literal $s \approx t$, we assign the multiset $\{s, t\}$, to a negative literal $s \not\approx t$ the multiset $\{s, s, t, t\}$. The *literal ordering* \succ_L compares these multisets using the multiset extension of \succ .

The *clause ordering* \succ_C compares clauses by comparing their multisets of literals using the multiset extension of \succ_L .

Constructing a candidate interpretation:

We want to use roughly the same ideas as in the completeness proof for resolution.

But: a Herbrand interpretation does not work for equality: The equality symbol \approx must be interpreted by equality in the interpretation.

Solution: Productive clauses contribute ground rewrite rules to a TRS R .

The interpretation has the universe $T_\Sigma(\emptyset)/R = T_\Sigma(\emptyset)/\approx_R$; a ground atom $s \approx t$ holds in the interpretation, if and only if $s \approx_R t$ if and only if $s \leftrightarrow_R^* t$.

We will construct R in such a way that it is terminating and confluent. In this case, $s \approx_R t$ if and only if $s \downarrow_R t$.

One problem:

In the completeness proof for the resolution calculus, the following property holds:

If $C = C' \vee A$ with a strictly maximal and positive literal A is false in the current interpretation, then adding A to the current interpretation cannot make any literal of C' true.

This does not hold for superposition:

Let $b \succ c \succ d$. Assume that the current rewrite system (representing the current interpretation) contains the rule $c \rightarrow d$. Now consider the clause $b \approx d \vee b \approx c$.

We need a further inference rule to deal with clauses of this kind, either the “Merging Paramodulation” rule of Bachmair and Ganzinger or the following “Equality Factoring” rule due to Nieuwenhuis:

$$\text{Equality Factoring: } \frac{C' \vee s \approx t' \vee s \approx t}{C' \vee t \not\approx t' \vee s \approx t}$$

Note: This inference rule subsumes the usual factoring rule.

How do the non-ground versions of the inference rules for superposition look like?

Main idea as in the resolution calculus:

Replace identity by unifiability. Apply the mgu to the resulting clause. In the ordering restrictions, replace \succ by $\not\prec$.

However:

As in Knuth-Bendix completion, we do not want to consider overlaps at or below a variable position.

Consequence: there are inferences between ground instances $D\theta$ and $C\theta$ of clauses D and C which are *not* ground instances of inferences between D and C .

Such inferences have to be treated in a special way in the completeness proof.

3.3 The Superposition Calculus – Formally

Until now, we have seen most of the ideas behind the superposition calculus and its completeness proof.

We will now start again from the beginning giving precise definitions and proofs.

Inference rules:

$$\text{Pos. Superposition: } \frac{D' \vee t \approx t' \quad C' \vee s[u] \approx s'}{(D' \vee C' \vee s[t'] \approx s')\sigma}$$

where $\sigma = \text{mgu}(t, u)$ and
 u is not a variable.

$$\text{Neg. Superposition: } \frac{D' \vee t \approx t' \quad C' \vee s[u] \not\approx s'}{(D' \vee C' \vee s[t'] \not\approx s')\sigma}$$

where $\sigma = \text{mgu}(t, u)$ and
 u is not a variable.

$$\text{Equality Resolution: } \frac{C' \vee s \not\approx s'}{C'\sigma}$$

where $\sigma = \text{mgu}(s, s')$.

$$\text{Equality Factoring: } \frac{C' \vee s' \approx t' \vee s \approx t}{(C' \vee t \not\approx t' \vee s \approx t')\sigma}$$

where $\sigma = \text{mgu}(s, s')$.

Theorem 3.1 *All inference rules of the superposition calculus are correct, i. e., for every rule*

$$\frac{C_n, \dots, C_1}{C_0}$$

we have $\{C_1, \dots, C_n\} \models C_0$.

Proof. Exercise. □

Orderings:

Let \succ be a *reduction ordering that is total on ground terms*.

To a positive literal $s \approx t$, we assign the multiset $\{s, t\}$, to a negative literal $s \not\approx t$ the multiset $\{s, s, t, t\}$. The *literal ordering* \succ_L compares these multisets using the multiset extension of \succ .

The *clause ordering* \succ_C compares clauses by comparing their multisets of literals using the multiset extension of \succ_L .

Inferences have to be computed only if the following ordering restrictions are satisfied:

- In superposition inferences, after applying the unifier to both premises, the left premise is not greater than or equal to the right one.
- The last literal in each premise is maximal in the respective premise, i. e., there exists no greater literal (strictly maximal for positive literals in superposition inferences, i. e., there exists no greater or equal literal).
- In these literals, the lhs is not smaller than the rhs (in superposition inferences: neither smaller nor equal).

A ground clause C is called *redundant w. r. t. a set of ground clauses N* , if it follows from clauses in N that are smaller than C .

A clause is *redundant w. r. t. a set of clauses N* , if all its ground instances are redundant w. r. t. $G_\Sigma(N)$.

The set of all clauses that are redundant w. r. t. N is denoted by $Red(N)$.

N is called *saturated up to redundancy*, if the conclusion of every inference from clauses in $N \setminus Red(N)$ is contained in $N \cup Red(N)$.

3.4 Superposition: Refutational Completeness

For a set E of ground equations, $T_\Sigma(\emptyset)/E$ is an E -interpretation (or E -algebra) with universe $\{[t] \mid t \in T_\Sigma(\emptyset)\}$.

One can show (similar to the proof of Birkhoff's Theorem) that for every *ground equation* $s \approx t$ we have $T_\Sigma(\emptyset)/E \models s \approx t$ if and only if $s \leftrightarrow_E^* t$.

In particular, if E is a convergent set of rewrite rules R and $s \approx t$ is a ground equation, then $T_\Sigma(\emptyset)/R \models s \approx t$ if and only if $s \downarrow_R t$. By abuse of terminology, we say that an equation or clause is valid (or true) in R if and only if it is true in $T_\Sigma(\emptyset)/R$.

Construction of candidate interpretations (Bachmair & Ganzinger 1990):

Let N be a set of clauses not containing \perp . Using induction on the clause ordering we define sets of rewrite rules E_C and R_C for all $C \in G_\Sigma(N)$ as follows:

Assume that E_D has already been defined for all $D \in G_\Sigma(N)$ with $D \prec_C C$. Then $R_C = \bigcup_{D \prec_C C} E_D$.

The set E_C contains the rewrite rule $s \rightarrow t$, if

- (a) $C = C' \vee s \approx t$.
- (b) $s \approx t$ is strictly maximal in C .
- (c) $s \succ t$.
- (d) C is false in R_C .
- (e) C' is false in $R_C \cup \{s \rightarrow t\}$.
- (f) s is irreducible w. r. t. R_C .

In this case, C is called *productive*. Otherwise $E_C = \emptyset$.

Finally, $R_\infty = \bigcup_{D \in G_\Sigma(N)} E_D$.

Lemma 3.2 *If $E_C = \{s \rightarrow t\}$ and $E_D = \{u \rightarrow v\}$, then $s \succ u$ if and only if $C \succ_C D$.*

Corollary 3.3 *The rewrite systems R_C and R_∞ are convergent.*

Proof. Obviously, $s \succ t$ for all rules $s \rightarrow t$ in R_C and R_∞ .

Furthermore, it is easy to check that there are no critical pairs between any two rules: Assume that there are rules $u \rightarrow v$ in E_D and $s \rightarrow t$ in E_C such that u is a subterm of s . As \succ is a reduction ordering that is total on ground terms, we get $u \prec s$ and therefore $D \prec_C C$ and $E_D \subseteq R_C$. But then s would be reducible by R_C , contradicting condition (f). \square

Lemma 3.4 *If $D \preceq_C C$ and $E_C = \{s \rightarrow t\}$, then $s \succ u$ for every term u occurring in a negative literal in D and $s \succeq v$ for every term v occurring in a positive literal in D .*

Corollary 3.5 *If $D \in G_\Sigma(N)$ is true in R_D , then D is true in R_∞ and R_C for all $C \succ_C D$.*

Proof. If a positive literal of D is true in R_D , then this is obvious.

Otherwise, some negative literal $s \not\approx t$ of D must be true in R_D , hence $s \not\downarrow_{R_D} t$. As the rules in $R_\infty \setminus R_D$ have left-hand sides that are larger than s and t , they cannot be used in a rewrite proof of $s \downarrow t$, hence $s \not\downarrow_{R_C} t$ and $s \not\downarrow_{R_\infty} t$. \square

Corollary 3.6 *If $D = D' \vee u \approx v$ is productive, then D' is false and D is true in R_∞ and R_C for all $C \succ_C D$.*

Proof. Obviously, D is true in R_∞ and R_C for all $C \succ_C D$.

Since all negative literals of D' are false in R_D , it is clear that they are false in R_∞ and R_C . For the positive literals $u' \approx v'$ of D' , condition (e) ensures that they are false in $R_D \cup \{u \rightarrow v\}$. Since $u' \preceq u$ and $v' \preceq u$ and all rules in $R_\infty \setminus R_D$ have left-hand sides that are larger than u , these rules cannot be used in a rewrite proof of $u' \downarrow v'$, hence $u' \not\downarrow_{R_C} v'$ and $u' \not\downarrow_{R_\infty} v'$. \square

Lemma 3.7 (“Lifting Lemma”) *Let C be a clause and let θ be a substitution such that $C\theta$ is ground. Then every equality resolution or equality factoring inference from $C\theta$ is a ground instance of an inference from C .*

Proof. Exercise. \square

Lemma 3.8 (“Lifting Lemma”) *Let $D = D' \vee u \approx v$ and $C = C' \vee [\neg] s \approx t$ be two clauses (without common variables) and let θ be a substitution such that $D\theta$ and $C\theta$ are ground.*

If there is a superposition inference between $D\theta$ and $C\theta$ where $u\theta$ and some subterm of $s\theta$ are overlapped, and $u\theta$ does not occur in $s\theta$ at or below a variable position of s , then the inference is a ground instance of a superposition inference from D and C .

Proof. Exercise. \square

Theorem 3.9 (“Model Construction”) *Let N be a set of clauses that is saturated up to redundancy and does not contain the empty clause. Then we have for every ground clause $C\theta \in G_\Sigma(N)$:*

- (i) $E_{C\theta} = \emptyset$ if and only if $C\theta$ is true in $R_{C\theta}$.
- (ii) If $C\theta$ is redundant w. r. t. $G_\Sigma(N)$, then it is true in $R_{C\theta}$.
- (iii) $C\theta$ is true in R_∞ and in R_D for every $D \in G_\Sigma(N)$ with $D \succ_C C\theta$.

Proof. We use induction on the clause ordering \succ_c and assume that (i)–(iii) are already satisfied for all clauses in $G_\Sigma(N)$ that are smaller than $C\theta$. Note that the “if” part of (i) is obvious from the construction and that condition (iii) follows immediately from (i) and Corollaries 3.50 and 3.51. So it remains to show (ii) and the “only if” part of (i).

Case 1: $C\theta$ is redundant w. r. t. $G_\Sigma(N)$.

If $C\theta$ is redundant w. r. t. $G_\Sigma(N)$, then it follows from clauses in $G_\Sigma(N)$ that are smaller than $C\theta$. By part (iii) of the induction hypothesis, these clauses are true in $R_{C\theta}$. Hence $C\theta$ is true in $R_{C\theta}$.

Case 2: $x\theta$ is reducible by $R_{C\theta}$.

Suppose there is a variable x occurring in C such that $x\theta$ is reducible by $R_{C\theta}$, say $x\theta \rightarrow_{R_{C\theta}} w$. Let the substitution θ' be defined by $x\theta' = w$ and $y\theta' = y\theta$ for every variable $y \neq x$. The clause $C\theta'$ is smaller than $C\theta$. By part (iii) of the induction hypothesis, it is true in $R_{C\theta}$. By congruence, every literal of $C\theta$ is true in $R_{C\theta}$ if and only if the corresponding literal of $C\theta'$ is true in $R_{C\theta}$; hence $C\theta$ is true in $R_{C\theta}$.

Case 3: $C\theta$ contains a maximal negative literal.

Suppose that $C\theta$ does not fall into Case 1 or 2 and that $C\theta = C'\theta \vee s\theta \not\approx s'\theta$, where $s\theta \not\approx s'\theta$ is maximal in $C\theta$. If $s\theta \approx s'\theta$ is false in $R_{C\theta}$, then $C\theta$ is clearly true in $R_{C\theta}$ and we are done. So assume that $s\theta \approx s'\theta$ is true in $R_{C\theta}$, that is, $s\theta \downarrow_{R_{C\theta}} s'\theta$. Without loss of generality, $s\theta \succeq s'\theta$.

Case 3.1: $s\theta = s'\theta$.

If $s\theta = s'\theta$, then there is an *equality resolution* inference

$$\frac{C'\theta \vee s\theta \not\approx s'\theta}{C'\theta}.$$

As shown in the Lifting Lemma, this is an instance of an *equality resolution* inference

$$\frac{C' \vee s \not\approx s'}{C'\sigma}$$

where $C = C' \vee s \not\approx s'$ is contained in N and $\theta = \rho \circ \sigma$. (Without loss of generality, σ is idempotent, therefore $C'\theta = C'\sigma\rho = C'\sigma\sigma\rho = C'\sigma\theta$, so $C'\theta$ is a ground instance of $C'\sigma$.) Since $C\theta$ is not redundant w. r. t. $G_\Sigma(N)$, C is not redundant w. r. t. N . As N is saturated up to redundancy, the conclusion $C'\sigma$ of the inference from C is contained in $N \cup \text{Red}(N)$. Therefore, $C'\theta$ is either contained in $G_\Sigma(N)$ and smaller than $C\theta$, or it follows from clauses in $G_\Sigma(N)$ that are smaller than itself (and therefore smaller than $C\theta$). By the induction hypothesis, clauses in $G_\Sigma(N)$ that are smaller than $C\theta$ are true in $R_{C\theta}$, thus $C'\theta$ and $C\theta$ are true in $R_{C\theta}$.

Case 3.2: $s\theta \succ s'\theta$.

If $s\theta \downarrow_{R_{C\theta}} s'\theta$ and $s\theta \succ s'\theta$, then $s\theta$ must be reducible by some rule in some $E_{D\theta} \subseteq R_{C\theta}$. (Without loss of generality we assume that C and D are variable disjoint; so we can use the same substitution θ .) Let $D\theta = D'\theta \vee t\theta \approx t'\theta$ with $E_{D\theta} = \{t\theta \rightarrow t'\theta\}$. Since $D\theta$ is productive, $D'\theta$ is false in $R_{C\theta}$. Besides, by part (ii) of the induction hypothesis, $D\theta$ is not redundant w. r. t. $G_\Sigma(N)$, so D is not redundant w. r. t. N . Note that $t\theta$ cannot occur in $s\theta$ at or below a variable position of s , say $x\theta = w[t\theta]$, since otherwise $C\theta$ would be subject to Case 2 above. Consequently, the *left superposition* inference

$$\frac{D'\theta \vee t\theta \approx t'\theta \quad C'\theta \vee s\theta[t\theta] \not\approx s'\theta}{D'\theta \vee C'\theta \vee s\theta[t'\theta] \not\approx s'\theta}$$

is a ground instance of a *left superposition* inference from D and C . By saturation up to redundancy, its conclusion is either contained in $G_\Sigma(N)$ and smaller than $C\theta$, or it follows from clauses in $G_\Sigma(N)$ that are smaller than itself (and therefore smaller than $C\theta$). By the induction hypothesis, these clauses are true in $R_{C\theta}$, thus $D'\theta \vee C'\theta \vee s\theta[t'\theta] \not\approx s'\theta$ is true in $R_{C\theta}$. Since $D'\theta$ and $s\theta[t'\theta] \not\approx s'\theta$ are false in $R_{C\theta}$, both $C'\theta$ and $C\theta$ must be true.

Case 4: $C\theta$ does not contain a maximal negative literal.

Suppose that $C\theta$ does not fall into Cases 1 to 3. Then $C\theta$ can be written as $C'\theta \vee s\theta \approx s'\theta$, where $s\theta \approx s'\theta$ is a maximal literal of $C\theta$. If $E_{C\theta} = \{s\theta \rightarrow s'\theta\}$ or $C'\theta$ is true in $R_{C\theta}$ or $s\theta = s'\theta$, then there is nothing to show, so assume that $E_{C\theta} = \emptyset$ and that $C'\theta$ is false in $R_{C\theta}$. Without loss of generality, $s\theta \succ s'\theta$.

Case 4.1: $s\theta \approx s'\theta$ is maximal in $C\theta$, but not strictly maximal.

If $s\theta \approx s'\theta$ is maximal in $C\theta$, but not strictly maximal, then $C\theta$ can be written as $C''\theta \vee t\theta \approx t'\theta \vee s\theta \approx s'\theta$, where $t\theta = s\theta$ and $t'\theta = s'\theta$. In this case, there is a *equality factoring* inference

$$\frac{C''\theta \vee t\theta \approx t'\theta \vee s\theta \approx s'\theta}{C''\theta \vee t'\theta \not\approx s'\theta \vee t\theta \approx t'\theta}$$

This inference is a ground instance of an inference from C . By saturation, its conclusion is true in $R_{C\theta}$. Trivially, $t'\theta = s'\theta$ implies $t'\theta \downarrow_{R_{C\theta}} s'\theta$, so $t'\theta \not\approx s'\theta$ must be false and $C\theta$ must be true in $R_{C\theta}$.

Case 4.2: $s\theta \approx s'\theta$ is strictly maximal in $C\theta$ and $s\theta$ is reducible.

Suppose that $s\theta \approx s'\theta$ is strictly maximal in $C\theta$ and $s\theta$ is reducible by some rule in $E_{D\theta} \subseteq R_{C\theta}$. Let $D\theta = D'\theta \vee t\theta \approx t'\theta$ and $E_{D\theta} = \{t\theta \rightarrow t'\theta\}$. Since $D\theta$ is productive, $D\theta$ is not redundant and $D'\theta$ is false in $R_{C\theta}$. We can now proceed in essentially the

same way as in Case 3.2: If $t\theta$ occurred in $s\theta$ at or below a variable position of s , say $x\theta = w[t\theta]$, then $C\theta$ would be subject to Case 2 above. Otherwise, the *right superposition* inference

$$\frac{D'\theta \vee t\theta \approx t'\theta \quad C'\theta \vee s\theta[t\theta] \approx s'\theta}{D'\theta \vee C'\theta \vee s\theta[t'\theta] \approx s'\theta}$$

is a ground instance of a *right superposition* inference from D and C . By saturation up to redundancy, its conclusion is true in $R_{C\theta}$. Since $D'\theta$ and $C'\theta$ are false in $R_{C\theta}$, $s\theta[t'\theta] \approx s'\theta$ must be true in $R_{C\theta}$. On the other hand, $t\theta \approx t'\theta$ is true in $R_{C\theta}$, so by congruence, $s\theta[t\theta] \approx s'\theta$ and $C\theta$ are true in $R_{C\theta}$.

Case 4.3: $s\theta \approx s'\theta$ is strictly maximal in $C\theta$ and $s\theta$ is irreducible.

Suppose that $s\theta \approx s'\theta$ is strictly maximal in $C\theta$ and $s\theta$ is irreducible by $R_{C\theta}$. Then there are three possibilities: $C\theta$ can be true in $R_{C\theta}$, or $C'\theta$ can be true in $R_{C\theta} \cup \{s\theta \rightarrow s'\theta\}$, or $E_{C\theta} = \{s\theta \rightarrow s'\theta\}$. In the first and the third case, there is nothing to show. Let us therefore assume that $C\theta$ is false in $R_{C\theta}$ and $C'\theta$ is true in $R_{C\theta} \cup \{s\theta \rightarrow s'\theta\}$. Then $C'\theta = C''\theta \vee t\theta \approx t'\theta$, where the literal $t\theta \approx t'\theta$ is true in $R_{C\theta} \cup \{s\theta \rightarrow s'\theta\}$ and false in $R_{C\theta}$. In other words, $t\theta \downarrow_{R_{C\theta} \cup \{s\theta \rightarrow s'\theta\}} t'\theta$, but not $t\theta \downarrow_{R_{C\theta}} t'\theta$. Consequently, there is a rewrite proof of $t\theta \rightarrow^* u \leftarrow^* t'\theta$ by $R_{C\theta} \cup \{s\theta \rightarrow s'\theta\}$ in which the rule $s\theta \rightarrow s'\theta$ is used at least once. Without loss of generality we assume that $t\theta \succeq t'\theta$. Since $s\theta \approx s'\theta \succ_L t\theta \approx t'\theta$ and $s\theta \succ s'\theta$ we can conclude that $s\theta \succeq t\theta \succ t'\theta$. But then there is only one possibility how the rule $s\theta \rightarrow s'\theta$ can be used in the rewrite proof: We must have $s\theta = t\theta$ and the rewrite proof must have the form $t\theta \rightarrow s'\theta \rightarrow^* u \leftarrow^* t'\theta$, where the first step uses $s\theta \rightarrow s'\theta$ and all other steps use rules from $R_{C\theta}$. Consequently, $s'\theta \approx t'\theta$ is true in $R_{C\theta}$. Now observe that there is an *equality factoring* inference

$$\frac{C''\theta \vee t\theta \approx t'\theta \vee s\theta \approx s'\theta}{C''\theta \vee t'\theta \not\approx s'\theta \vee t\theta \approx t'\theta}$$

whose conclusion is true in $R_{C\theta}$ by saturation. Since the literal $t'\theta \not\approx s'\theta$ must be false in $R_{C\theta}$, the rest of the clause must be true in $R_{C\theta}$, and therefore $C\theta$ must be true in $R_{C\theta}$, contradicting our assumption. This concludes the proof of the theorem. \square

A Σ -interpretation \mathcal{A} is called *term-generated*, if for every $b \in U_{\mathcal{A}}$ there is a ground term $t \in T_{\Sigma}(\emptyset)$ such that $b = \mathcal{A}(\beta)(t)$.

Lemma 3.10 *Let N be a set of (universally quantified) Σ -clauses and let \mathcal{A} be a term-generated Σ -interpretation. Then \mathcal{A} is a model of $G_{\Sigma}(N)$ if and only if it is a model of N .*

Proof. (\Rightarrow): Let $\mathcal{A} \models G_\Sigma(N)$; let $(\forall \vec{x}C) \in N$. Then $\mathcal{A} \models \forall \vec{x}C$ iff $\mathcal{A}(\gamma[x_i \mapsto a_i])(C) = 1$ for all γ and a_i . Choose ground terms t_i such that $\mathcal{A}(\gamma)(t_i) = a_i$; define θ such that $x_i\theta = t_i$, then $\mathcal{A}(\gamma[x_i \mapsto a_i])(C) = \mathcal{A}(\gamma \circ \theta)(C) = \mathcal{A}(\gamma)(C\theta) = 1$ since $C\theta \in G_\Sigma(N)$.

(\Leftarrow): Let \mathcal{A} be a model of N ; let $C \in N$ and $C\theta \in G_\Sigma(N)$. Then $\mathcal{A}(\gamma)(C\theta) = \mathcal{A}(\gamma \circ \theta)(C) = 1$ since $\mathcal{A} \models N$. \square

Theorem 3.11 (Refutational Completeness: Static View) *Let N be a set of clauses that is saturated up to redundancy. Then N has a model if and only if N does not contain the empty clause.*

Proof. If $\perp \in N$, then obviously N does not have a model. If $\perp \notin N$, then the interpretation R_∞ (that is, $T_\Sigma(\emptyset)/R_\infty$) is a model of all ground instances in $G_\Sigma(N)$ according to part (iii) of the model construction theorem. As $T_\Sigma(\emptyset)/R_\infty$ is term-generated, it is a model of N . \square

So far, we have considered only inference rules that add new clauses to the current set of clauses (corresponding to the *Deduce* rule of Knuth-Bendix Completion).

In other words, we have derivations of the form $N_0 \vdash N_1 \vdash N_2 \vdash \dots$, where each N_{i+1} is obtained from N_i by adding the consequence of some inference from clauses in N_i .

Under which circumstances are we allowed to delete (or simplify) a clause during the derivation?

A *run* of the superposition calculus is a sequence $N_0 \vdash N_1 \vdash N_2 \vdash \dots$, such that

- (i) $N_i \models N_{i+1}$, and
- (ii) all clauses in $N_i \setminus N_{i+1}$ are redundant w. r. t. N_{i+1} .

In other words, during a run we may add a new clause if it follows from the old ones, and we may delete a clause, if it is redundant w. r. t. the remaining ones.

For a run, $N_\infty = \bigcup_{i \geq 0} N_i$ and $N_* = \bigcup_{i \geq 0} \bigcap_{j \geq i} N_j$. The set N_* of all *persistent* clauses is called the *limit* of the run.

Lemma 3.12 *If $N \subseteq N'$, then $Red(N) \subseteq Red(N')$.*

Proof. Obvious. \square

Lemma 3.13 *If $N' \subseteq Red(N)$, then $Red(N) \subseteq Red(N \setminus N')$.*

Proof. Follows from the compactness of first-order logic and the well-foundedness of the multiset extension of the clause ordering. \square

Lemma 3.14 *Let $N_0 \vdash N_1 \vdash N_2 \vdash \dots$ be a run. Then $Red(N_i) \subseteq Red(N_\infty)$ and $Red(N_i) \subseteq Red(N_*)$ for every i .*

Proof. Exercise. □

Corollary 3.15 *$N_i \subseteq N_* \cup Red(N_*)$ for every i .*

Proof. If $C \in N_i \setminus N_*$, then there is a $k \geq i$ such that $C \in N_k \setminus N_{k+1}$, so C must be redundant w.r.t. N_{k+1} . Consequently, C is redundant w.r.t. N_* . □

A run is called *fair*, if the conclusion of every inference from clauses in $N_* \setminus Red(N_*)$ is contained in some $N_i \cup Red(N_i)$.

Lemma 3.16 *If a run is fair, then its limit is saturated up to redundancy.*

Proof. If the run is fair, then the conclusion of every inference from non-redundant clauses in N_* is contained in some $N_i \cup Red(N_i)$, and therefore contained in $N_* \cup Red(N_*)$. Hence N_* is saturated up to redundancy. □

Theorem 3.17 (Refutational Completeness: Dynamic View) *Let $N_0 \vdash N_1 \vdash N_2 \vdash \dots$ be a fair run, let N_* be its limit. Then N_0 has a model if and only if $\perp \notin N_*$.*

Proof. (\Leftarrow): By fairness, N_* is saturated up to redundancy. If $\perp \notin N_*$, then it has a term-generated model. Since every clause in N_0 is contained in N_* or redundant w.r.t. N_* , this model is also a model of $G_\Sigma(N_0)$ and therefore a model of N_0 .

(\Rightarrow): Obvious, since $N_0 \models N_*$. □