Automated Theorem Proving

Prof. Dr. Jasmin Blanchette, Lydia Kondylidou,
Yiming Xu, PhD, and Tanguy Bozec
based on exercises by Dr. Uwe Waldmann

Winter Term 2024/25

Exercises 5: Resolution

Exercise 5.1: Let $\Sigma = (\Omega, \Pi)$ be a first-order signature with $\Omega = \{b/0,\ f/1\}$ and $\Pi = \{P/1\}$. Determine for each of the following statements whether they are true or false:

(1) There is a Σ-model $\mathcal{A}$ of $P(b) \land \neg P(f(b))$ such that $U_{\mathcal{A}} = \{7, 8, 9\}$.

(2) There is a Σ-model $\mathcal{A}$ of $P(b) \land \neg P(f(f(b)))$ such that $f_{\mathcal{A}}(a) = a$ for every $a \in U_{\mathcal{A}}$.

(3) $P(b) \land \neg P(f(b))$ has a Herbrand model.

(4) $P(b) \land \forall x\, \neg P(x)$ has a Herbrand model.

(5) $\forall x\, P(f(x))$ has a Herbrand model with a two-element universe.

(6) $\forall x\, P(x)$ has exactly one Herbrand model.

(7) $\forall x\, P(f(x))$ entails $\forall x\, P(f(f(x)))$.

Proposed solution. (1) True. Define $b_{\mathcal{A}} = 7$, $f_{\mathcal{A}}(a) = 8$ for $a \in \{7, 8, 9\}$ and $P_{\mathcal{A}} = \{7\}$.

(2) False. If there existed a model $\mathcal{A}$ such that $f_{\mathcal{A}}(a) = a$ for every $a \in U_{\mathcal{A}}$, then we would have $b_{\mathcal{A}} = f_{\mathcal{A}}(b_{\mathcal{A}}) = f_{\mathcal{A}}(f_{\mathcal{A}}(b_{\mathcal{A}}))$, but $b_{\mathcal{A}} \in P_{\mathcal{A}}$ and $f_{\mathcal{A}}(f_{\mathcal{A}}(b_{\mathcal{A}})) \notin P_{\mathcal{A}}$.

(3) True. Define $P_{\mathcal{A}} = \{b\}$.

(4) False. The formula is contradictory; it has no model and in particular no Herbrand model.

(5) False. Every Herbrand interpretation (and therefore every Herbrand model) over the signature Σ has the infinite universe $\mathrm{T}_{\Sigma} = \{b, f(b), f(f(b)), \dots\}$.

(6) True. The Herbrand interpretation in which $P_{\mathcal{A}} = \mathrm{T}_{\Sigma}$ is the only Herbrand model.

(7) True. If $f_{\mathcal{A}}(a) \in P_{\mathcal{A}}$ for every $a \in U_{\mathcal{A}}$, since $f_{\mathcal{A}}(a) \in U_{\mathcal{A}}$ we have $f_{\mathcal{A}}(f_{\mathcal{A}}(a)) \in P_{\mathcal{A}}$ for every $a \in U_{\mathcal{A}}$.

Exercise 5.2: Let $\Sigma = (\Omega, \Pi)$ be a first-order signature with $\Omega = \{b/0,\ f/1\}$ and $\Pi = \{P/1\}$. Let F be the Σ-formula

$$\neg P(b) \wedge P(f(f(b))) \wedge \forall x\, (\neg P(x) \vee P(f(x))).$$

Determine for each of the following statements whether they are true or false:

(1) There is a Σ-model $\mathcal{A}$ of F such that $U_{\mathcal{A}} = \{7, 8, 9\}$.

(2) There is a Σ-model $\mathcal{A}$ of F such that $f_{\mathcal{A}}(a) = a$ for every $a \in U_{\mathcal{A}}$.

(3) F has exactly two Σ-models.

(4) Every Σ-model of F is a model of $\exists x\, P(x)$.

(5) Every Σ-model of F is a model of $\forall x\, P(f(f(x)))$.

(6) There are infinitely many Herbrand interpretations over Σ.

(7) There is a Herbrand model of F over Σ whose universe has exactly two elements.

(8) There is a Herbrand model of F over Σ with an infinite universe.

(9) F has exactly two Herbrand models over Σ.

Proposed solution. (1) True. E.g., $U_{\mathcal{A}} = \{7, 8, 9\}$, $b_{\mathcal{A}} = 7$, $f_{\mathcal{A}}(7) = 8$, $f_{\mathcal{A}}(8) = 8$, $f_{\mathcal{A}}(9) = 9$, $P_{\mathcal{A}} = \{8\}$.

(2) False. If $f_{\mathcal{A}}(a) = a$ for every $a \in U_{\mathcal{A}}$, then $f_{\mathcal{A}}(f_{\mathcal{A}}(b_{\mathcal{A}})) = f_{\mathcal{A}}(b_{\mathcal{A}}) = b_{\mathcal{A}}$, but $f_{\mathcal{A}}(f_{\mathcal{A}}(b_{\mathcal{A}})) \in P_{\mathcal{A}}$ and $b_{\mathcal{A}} \notin P_{\mathcal{A}}$.

(3) False. F has infinitely many models.

(4) True. In every model of F, $P(x)$ holds for the assignment that maps x to $f_{\mathcal{A}}(f_{\mathcal{A}}(b_{\mathcal{A}}))$.

(5) False. E.g., in the model given for (1), $P(f(f(x)))$ does not hold for the assignment that maps x to 9.

(6) True. The universe of a Herbrand interpretation over Σ is the set of ground Σ-terms, i.e., $\mathrm{T}_{\Sigma}(\emptyset) = \{b, f(b), f(f(b)), f(f(f(b))), \ldots\}$. Since the universe is infinite, there are infinitely many ways to interpret P.

(7) False. For every Herbrand model of F over Σ, the universe is infinite, see (6).

(8) True. In fact, *every* Herbrand model over Σ has an infinite universe, see (6).

(9) True. In every Herbrand model for F, $P(b)$ must be false and $P(f^n(b))$ must be true for every $n \geq 2$. Since $P(f(b))$ can be either true or false, there are two Herbrand models for F.

Exercise 5.3: Let $\Sigma = (\Omega, \Pi)$ be a first-order signature with $\Omega = \{f/1,\ b/0,\ c/0\}$ and $\Pi = \{P/1\}$. Are the following statements correct?

(1) The formula $\forall x\, P(x)$ has infinitely many Σ-models.

(2) Every model of $\forall x\, P(x)$ is a model of $\forall x\, P(f(x))$.

(3) The formula $\neg P(b) \wedge \forall x\, P(x)$ has a Σ-model with an infinite universe.

(4) The formula $\neg P(b) \wedge \forall x\, P(f(x))$ has a Σ-model with a two-element universe.

(5) Every Σ-model of $P(b) \wedge P(c) \wedge \forall x\, P(f(x))$ is a model of $\forall x\, P(x)$.

(6) Every Herbrand model over Σ of $P(b) \wedge P(c) \wedge \forall x\, P(f(x))$ has an infinite universe.

(7) The formula $P(b) \vee P(c)$ has exactly three Herbrand models over Σ.

(8) The formula $\forall x\, P(f(x))$ has exactly four Herbrand models over Σ.

Proposed solution. (1) True. In particular, it has models with arbitrarily large universes.

(2) True. $\forall x\, P(x) \models \forall x\, P(f(x))$.

(3) False. The formula is unsatisfiable, so it has no models at all.

(4) True. Take $U_{\mathcal{A}} = \{1, 2\}$, $b_{\mathcal{A}} = 1$, $c_{\mathcal{A}} = 1$, $f_{\mathcal{A}} : x \mapsto 2$, $P_{\mathcal{A}} = \{2\}$.

(5) False. Take $U_{\mathcal{A}} = \{1, 2\}$, $b_{\mathcal{A}} = 1$, $c_{\mathcal{A}} = 1$, $f_{\mathcal{A}} : x \mapsto 1$, $P_{\mathcal{A}} = \{1\}$.

(6) True. In fact all Herbrand interpretations over Σ have the same infinite universe $\{b, c, f(b), f(c), f(f(b)), f(f(c)), \dots\}$.

(7) False. $P(b) \vee P(c)$ has infinitely many Herbrand models over Σ, which differ in the interpretation of P on ground terms different from b and c.

(8) True. The interpretation of P on all ground terms with f at the root is fixed, but P can be either true or false for b and either true or false for c; this leaves four combinations.

Exercise 5.4: Let $\Sigma = (\Omega, \Pi)$ be a first-order signature with $\Omega = \{b/0,\ f/1\}$ and $\Pi = \{P/1\}$. Let F be the Σ-formula

$$\neg P(b) \wedge P(f(f(b))) \wedge \forall x\, (P(x) \vee P(f(x))).$$

Determine for each of the following statements whether they are true or false:

(1) If $\mathcal{A}$ is a Σ-model of F, then $P_{\mathcal{A}} \neq \emptyset$ and $P_{\mathcal{A}} \neq U_{\mathcal{A}}$.

(2) There is a Σ-model $\mathcal{A}$ of F such that $U_{\mathcal{A}} = \{7, 8, 9\}$.

(3) There is a Σ-model $\mathcal{A}$ of F such that $f_{\mathcal{A}}(a) = f_{\mathcal{A}}(a')$ for all $a, a' \in U_{\mathcal{A}}$.

(4) F has exactly four Σ-models.

(5) There are infinitely many Herbrand interpretations over Σ.

(6) There is an Herbrand model of F over Σ with a finite universe.

(7) There is an Herbrand model $\mathcal{A}$ of F over Σ and an assignment β such that $\mathcal{A}(\beta)(f(b)) = \mathcal{A}(\beta)(f(f(b)))$.

Proposed solution. (1): True. $P_{\mathcal{A}}$ cannot equal $U_{\mathcal{A}}$, since $b_{\mathcal{A}} \notin P_{\mathcal{A}}$; $P_{\mathcal{A}}$ cannot be empty, since $f_{\mathcal{A}}(f_{\mathcal{A}}(b_{\mathcal{A}})) \in P_{\mathcal{A}}$.

(2) True. Let $U_{\mathcal{A}} = \{7, 8, 9\}$, let $b_{\mathcal{A}} = 7$, let $f_{\mathcal{A}}$ map every element of $U_{\mathcal{A}}$ to 8, and let $P_{\mathcal{A}} = \{8\}$.

(3) True. See (2).

(4) False. F has infinitely many Σ-models; in particular it has Σ-models with any universe with at least 2 elements.

(5) True. Since $\mathrm{T}_{\Sigma}(\emptyset)$ is infinite, there are infinitely many different possibilities to choose a subset $P_{\mathcal{A}} \subseteq \mathrm{T}_{\Sigma}(\emptyset)$.

(6) False. All Herbrand models of F over Σ have the same universe $\mathrm{T}_{\Sigma}(\emptyset)$ (which is infinite).

(7) False. If $\mathcal{A}$ is an Herbrand model over Σ, then $\mathcal{A}(\beta)(t) = t$ for every ground term $t \in \mathrm{T}_{\Sigma}(\emptyset)$, so $\mathcal{A}(\beta)(f(b))$ and $\mathcal{A}(\beta)(f(f(b)))$ are different elements of the universe.

Exercise 5.5: Determine for each of the following statements whether it is true or false:

(1) If $\Sigma = (\{b/0,\ c/0\},\ \{P/1\})$, then $P(b) \vee \neg P(c)$ has exactly three Herbrand models over Σ.

(2) If $\Sigma = (\{f/1,\ c/0\},\ \{P/1\})$, then $P(c) \vee P(f(c))$ has an Herbrand model over Σ whose universe has exactly four elements.

(3) If $\Sigma = (\{f/1,\ c/0\},\ \{P/1\})$, then $\neg P(c) \wedge \forall x\, P(f(x))$ has a model whose universe has exactly five elements

(4) If $\Sigma = (\{b/0,\ c/0,\ d/0\},\ \{P/1\})$, then $P(b) \vee \neg P(b)$ and $P(c) \vee \neg P(d)$ are equisatisfiable.

(5) If $\Sigma = (\{f/1,\ c/0\},\ \{P/1\})$, N is a set of universally quantified Σ-clauses, and every clause in N has at least one positive literal, then N has an Herbrand model.

(6) If $\Sigma = (\{f/1,\ c/0\},\ \{P/1\})$, N is a set of universally quantified Σ-clauses, and $N \models \neg P(x) \vee P(f(x))$, then N has a model.

(7) If $\Sigma = (\{f/1,\ c/0\},\ \{P/1\})$, then $\forall x\, P(f(x)) \models \forall y\, P(c) \vee P(f(f(y)))$.

Proposed solution. (1) True. There are exactly four Herbrand interpretations over Σ, namely $\emptyset$, $\{P(b)\}$, $\{P(c)\}$, and $\{P(b), P(c)\}$, and three of them (the first, the second, and the fourth) are models of $P(b) \vee \neg P(c)$.

(2) False. The universe of every Herbrand model is the set of ground terms. Since Σ contains a unary function symbol, there are infinitely many ground terms.

(3) True. Take $\mathcal{A}$ with $U_{\mathcal{A}} = \{1, 2, 3, 4, 5\}$, $c_{\mathcal{A}} = 1$, $f_{\mathcal{A}} : n \mapsto 2$, and $P_{\mathcal{A}} = \{2\}$.

(4) True. Both formulas are satisfiable, therefore the are equisatisfiable.

(5) True. Take an Herbrand interpretation in which all atoms are true; then every clause that has at least one positive literal is true in that interpretation.

(6) False. Take $N = \{\bot\}$.

(7) True. By Lemma 3.3.8, every model of $\forall x\, P(f(x))$ is also a model of $\forall y\, P(f(f(y)))$ and thus a model of $\forall y\, P(c) \vee P(f(f(y)))$.

Exercise 5.6: Let N be the set consisting of the following ground clauses:

$$
\begin{aligned}
&P \vee Q && (1)\\
&P \vee \neg Q && (2)\\
&\neg P \vee Q && (3)\\
&\neg P \vee \neg Q && (4)
\end{aligned}
$$

(a) Show that $N \vdash_{Res} \bot$, that is, derive $\bot$ from N using the "Resolution" and "Factorization" rules.

(b) Why is it impossible to derive the empty clause from N without using "Factorization"?

Proposed solution. (a) From (1) and (2), using "Resolution" we obtain

$$P \vee P \qquad (5)$$

From (5), using "Factorization"' we obtain

$$P \qquad (6)$$

From (3) and (4), using "Resolution" we obtain

$$\neg P \vee \neg P \qquad (7)$$

From (6) and (7), using "Resolution" we obtain

$$\neg P \qquad (8)$$

From (6) and (8), using "Resolution" we obtain the empty clause.

(b) Given two clauses with $m \geq 1$ and $n \geq 1$ literals, respectively, the result of "Resolution" always has $m + n - 2$ literals. Since the problem consists exclusively of two-literal clauses, the result of all inferences from N also consist of $2 + 2 - 2 = 2$ literals, which in turn can only yield two-literal clauses, and so on. The empty clause, which has zero literals, can never be generated.

Exercise 5.7 (∗)**:** Find a finite set N of ground clauses such that no clause in N is a tautology and such that $Res^*(N)$ is infinite.

Proposed solution. We take N to be the set consisting of the following clauses:

$$P \quad (1)$$
$$\neg P \vee Q \vee Q \quad (2)$$
$$\neg Q \vee P \vee P \quad (3)$$

Using the "Resolution" rule, from (1) and (2) we obtain

$$Q \vee Q \quad (4)$$

Using the "Resolution" rule, from (4) and (3) we obtain

$$P \vee P \vee Q \quad (5)$$

Using the "Resolution" rule, from (5) and (2) we obtain

$$Q \vee Q \vee P \vee Q \quad (6)$$

Using the "Resolution" rule, from (6) and (3) we obtain

$$P \vee P \vee Q \vee P \vee Q \quad (7)$$

And so on. At each step, we derive a clause with one more literal than we started with. Thus $Res^*(N)$ is infinite.

Exercise 5.8: Let $\Sigma = (\Omega, \Pi)$ with $\Omega = \{b/0,\ c/0\}$ and $\Pi = \{P/1,\ Q/0,\ R/0\}$. Use the ground resolution calculus *Res* to check whether the following clause set is satisfiable:

$$\begin{aligned} \neg P(b) \vee Q & \quad (1) \\ \neg P(b) \vee R & \quad (2) \\ \neg P(c) \vee Q & \quad (3) \\ \neg Q \vee \neg R & \quad (4) \\ Q \vee R & \quad (5) \\ P(b) & \quad (6) \\ \neg P(c) & \quad (7) \end{aligned}$$

Proposed solution. From (6) and (1) we obtain via "Resolution" Q (8), from (6) and (2) we obtain via "Resolution" R (9), from (8) and (4) we obtain via "Resolution" $\neg R$ (10), and from (9) and (10) we obtain via "Resolution" $\bot$. Since resolution is sound, the clause set is unsatisfiable.

Exercise 5.9: Use the ground resolution calculus to show that

$$\big\{(P \leftrightarrow (Q \wedge R)),\ (P \leftrightarrow Q)\big\} \models Q \rightarrow R$$

Hint: You will need some preprocessing.

Proposed solution.

1. Convert to CNF:
   - $P \leftrightarrow (Q \wedge R)$:

$$\begin{aligned} & P \leftrightarrow (Q \wedge R) \\ \Rightarrow_{CNF}\ & (P \rightarrow (Q \wedge R)) \wedge ((Q \wedge R) \rightarrow P) \\ \Rightarrow^{+}_{CNF}\ & (\neg P \vee (Q \wedge R)) \wedge (\neg(Q \wedge R) \vee P) \\ \Rightarrow^{+}_{CNF}\ & (\neg P \vee Q) \wedge (\neg P \vee R) \wedge (\neg Q \vee \neg R \vee P) \end{aligned}$$

   - $P \leftrightarrow Q$:

$$\begin{aligned} & P \leftrightarrow Q \\ \Rightarrow_{CNF}\ & (P \rightarrow Q) \wedge (Q \rightarrow P) \\ \Rightarrow^{+}_{CNF}\ & (\neg P \vee Q) \wedge (\neg Q \vee P) \end{aligned}$$

- Negation of $Q \rightarrow R$:

$$\begin{aligned} & \neg(Q \rightarrow R) \\ \Rightarrow^+_{CNF} \; & (Q \wedge \neg R) \end{aligned}$$

2. Combine all the clauses: $\{\neg P \vee Q, \neg P \vee R, \neg Q \vee \neg R \vee P, \neg P \vee Q, \neg Q \vee P, Q, \neg R\}$.

3. Use the resolution calculus to derive a contradiction:

1	$\neg P \vee Q$	(given)
2	$\neg P \vee R$	(given)
3	$\neg Q \vee \neg R \vee P$	(given)
4	$\neg P \vee Q$	(given)
5	$\neg Q \vee P$	(given)
6	Q	(given)
7	$\neg R$	(given)
8	P	(Res. 6 into 5)
9	R	(Res. 8 into 2)
10	$\bot$	(Res. 9 into 7)

Exercise 5.10: Prove or refute: $Res(N)$ is satisfiable if and only if N is satisfiable.

Proposed solution. The statement does not hold. For example, if $N = \{P \vee Q,\ \neg P,\ \neg Q\}$, then $Res(N) = \{Q,\ P\}$; if $N = \{\neg P \vee \neg P,\ P\}$, then $Res(N) = \{\neg P\}$; and if $N = \{\bot\}$, then $Res(N) = \emptyset$. In all three cases, the set N is unsatisfiable, but $Res(N)$ is satisfiable.

Exercise 5.11: Prove or refute: All clauses in $Res^*(N)$ are tautologies if and only if all clauses in N are tautologies.

Proposed solution. The statement holds.

Ground resolution is sound. This means that for every algebra $\mathcal{A}$, whenever the premises of an inference hold in $\mathcal{A}$, then the conclusion holds in $\mathcal{A}$ as well. In particular, if the premises are tautological (i.e., hold in every algebra $\mathcal{A}$), then the conclusion holds in every algebra $\mathcal{A}$, so it is also tautological. Thus, if all clauses in a set M are tautologies, then all clauses in $Res(M)$ are tautologies. By induction over n we can now show that, if all clauses in N are tautologies, then all clauses in $Res^n(N)$ are tautologies. So, all clauses in $Res^*(N) = \bigcup_{n \geq 0} Res^n(N)$ are tautologies.

The reverse direction follows immediately from the fact that $N \subseteq Res^*(N)$.

Exercise 5.12 ($*$)**:** Prove the following statement: If N is a set of propositional formulas and C is a propositional formula such that $N \models C$, then there exists a finite subset $M \subseteq N$ such that $M \models C$.

Proposed solution. Let $C = L_1 \vee \cdots \vee L_n$. If $N \models C$, then by definition of $\models$ we have $N \cup \{\overline{L_1}, \ldots, \overline{L_n}\} \models \bot$, where $\overline{L_i}$ denotes the complementary literal of L_i. By refutational completeness of ground resolution, we have $\bot \in \mathit{Res}^*(N \cup \{\overline{L_1}, \ldots, \overline{L_n}\})$. This means that there exists a finite derivation tree with $\bot$ at the root and clauses from $N \cup \{\overline{L_1}, \ldots, \overline{L_n}\}$ on its leaves. Take M to be the finite subset of clauses from N that appear on the leaves. The existence of the derivation tree means that $\bot \in \mathit{Res}^*(M \cup \{\overline{L_1}, \ldots, \overline{L_n}\})$. By soundness of ground resolution, we have $M \cup \{\overline{L_1}, \ldots, \overline{L_n}\} \models \bot$. Equivalently, $M \models C$.