Artificial Intelligence

08. Predicate Logic Reasoning, Part II: Reasoning

And Now: How to Actually *Think* in Terms of Predicates

Jana Koehler Álvaro Torralba

SAARLAND UNIVERSITY
COMPUTER SCIENCE

Summer Term 2019

Thanks to Prof. Hoffmann for slide sources

Agenda

1. Introduction
2. Reduction to Propositional Reasoning
3. Substitutions, and Unification
4. PL1 Resolution
5. On Criminals and Cats: PL1 Resolution Examples
6. Conclusion

Let's *Reason* About Blocks, Baby . . .

Reminder: Our Agenda for This Topic

→ Our treatment of the topic "Predicate Logic Reasoning" consists of Chapters 12 and 13.

- **Chapter 12:** Basic definitions and concepts; normal forms.

  → Sets up the framework and basic operations.

- **This Chapter:** Compilation to propositional reasoning; unification; lifted resolution.

  → Algorithmic principles for reasoning about predicate logic.

Our Agenda for This Chapter

- **Reduction to Propositional Reasoning:** Can we reduce PL1 reasoning to propositional reasoning?

  → Yes we can! (But it's tricky, and involves generating huge grounded encodings . . .)

- **Substitutions, and Unification:** What basic operations are required to avoid grounding everything out?

  → Specifies how to instantiate variables with terms.

- **PL1 Resolution:** How do we reason directly at PL1 level?

  → The foundational procedure for doing so.

- **On Criminals and Cats:** And now, in practice?

  → Gives some examples.

Reasoning About PL1 *Via Propositional Logic?*

What for do we need PL1, then?

- "First-order logic as syntactic sugar for propositional logic."
- Remember all these propositions in the Wumpus world?
- Anyway, it's of course not that easy in general (cf. slide 12).

How?

Reasoning About PL1 Via Propositional Logic

1. Bring into Skolem normal form (SNF).
2. Generate (the finite subsets of) the Herbrand expansion (up next).
3. Use propositional reasoning.

→ Apply DPLL, clause learning, ...

- Herbrand expansion may be very large (infinite, in general).
- Still, this often works well in practice.

Herbrand Expansion

We assume: Skolem normal form. (We don't require φ to be in CNF.)

universal prefix + (quantifier-free) matrix

$$\forall x_1 \forall x_2 \forall x_3 \ldots \forall x_n \; \varphi$$

Notation: For any (finite) set θ^* of PL1 formulas, denote by $CF(\theta^*)$ the set of constant symbols, and of function symbols (arity ≥ 1), occuring in θ^*. If no constant symbol occurs in θ^*, we add a new such symbol c into $CF(\theta^*)$.

Definition (Herbrand Universe). *Let θ^* be a set of PL1 formulas in SNF. Then the Herbrand universe $HU(\theta^*)$ over θ^* is the set of all ground terms that can be formed from $CF(\theta^*)$.*

Example: $\theta^* = \{\forall x[\neg Dog(x) \vee Chases(x, f(x))]\}$

$CF(\theta^*) =$ $\qquad$ $HU(\theta^*) =$

Herbrand Expansion, ctd.

Definition (Herbrand Expansion). *Let θ^* be a set of PL1 formulas in SNF. The Herbrand expansion $HE(\theta^*)$ is defined as:*

$$HE(\theta^*) = \{\varphi\frac{x_1}{t_1}, \ldots, \frac{x_n}{t_n} \mid (\forall x_1 \ldots \forall x_n \varphi) \in \theta^*, t_i \in HU(\theta^*)\}$$

$\rightarrow$ Instantiate each matrix φ with all terms from $HU(\theta^*)$. As $HE(\theta^*)$ contains ground atoms only, it can be interpreted as propositional logic.

Example: $\theta^* = \{\forall x[\neg Dog(x) \vee Chases(x, f(x))]\}$
$\rightarrow HE(\theta^*) = \{[\neg Dog(c) \vee Chases(c, f(c))],$
$[\neg Dog(f(c)) \vee Chases(f(c), f(f(c)))], \ldots\}.$

Theorem (Herbrand). *Let θ^* be a set of PL1 formulas in SNF. Then θ^* is satisfiable iff $HE(\theta^*)$ is satisfiable.* (Proof omitted.)

$\rightarrow$ **Observe:** Without function symbols, the Herbrand expansion is finite, and PL1 reasoning is equivalent to propositional reasoning.

When Herbrand Reasons About Blocks . . .

Herbrand: The Infinite Case

→ **Recall:** Without function symbols, the Herbrand expansion is finite, and PL1 reasoning is equivalent to propositional reasoning.

→ But what if there *are* function symbols?

Theorem (Compactness of Propositional Logic). *Any set θ of propositional logic formulas is unsatisfiable if and only if at least one finite subset of θ is unsatisfiable.* (Proof omitted.)

Method: Enumerate all finite subsets θ_1 of the Herbrand expansion $HE(\theta^*)$, and test propositional satisfiability of θ_1. θ is unsatisfiable if and only if one of the θ_1 is. Only ... which θ_1 will do the job?

→ If the Herbrand expansion is *infinite*, to show unsatisfiability (= to prove that some property does indeed follow from the KB), we must somehow choose a "relevant" finite subset thereof.

→ Direct PL1 reasoning ameliorating this caveat: later in this chapter.

Herbrand, Infinite Case: What If θ is Satisfiable?

Theorem (A). *The set of unsatisfiable PL1 formulas is recursively enumerable.*

Proof. Enumerate all PL1 formulas φ. Incrementally for all of these in parallel, enumerate all finite subsets θ_1 of the Herbrand expansion $HE(\varphi^*)$. Test propositional satisfiability of each θ_1. By compactness of propositional logic, if $HE(\varphi^*)$ is unsatisfiable then one of the θ_1 is.

Theorem (B). *It is undecidable whether a PL1 formula is satisfiable.* (Proof omitted.)

Corollary. *The set of satisfiable PL1 formulas is not recursively enumerable.* (Proof: Else, with Theorem (A), PL1 satisfiability would be decidable, in contradiction to Theorem (B).)

$\rightarrow$ If a PL1 formula is unsatisfiable, then we can confirm this. Otherwise, we might end up in an infinite loop.

Questionnaire

Question!

What is the Herbrand universe $HU(\theta^*)$ of
$\theta^* = \{\forall x[Equals(x, succ(f(x)))], \forall x \neg Equals(1, succ(x))\}$**?**

(A): $\{1\}$.

(B): $\{1, f, succ\}$.

(C): $\{1, succ(1), succ(succ(1)), \ldots\}$.

(D): $\{1, f(1), succ(1), succ(f(1)), f(succ(1)), \ldots\}$.

Questionnaire, ctd.

Question!

Is the Herbrand expansion of $\theta^* = \{\forall x[Equals(x, succ(f(x)))], \forall x \neg Equals(1, succ(x))\}$ **satisfiable?**

(A): Yes.

(B): No.

Towards PL1 Resolution

Clausal normal form:

universal prefix + disjunction of literals

$$\forall x_1 \forall x_2 \forall x_3 \ldots \forall x_n (l_1 \vee \cdots \vee l_n)$$

$\rightarrow$ Written $\{l_1, \ldots, l_n\}$.

$\rightarrow$ The quantifiers are omitted in the notation!

Example: $\{\{Nat(s(x)), \neg Nat(x)\}, \{Nat(1)\}\}$

We want to somehow apply/adapt the resolution rule:

$$\frac{C_1 \dot{\cup} \{l\}, C_2 \dot{\cup} \{\bar{l}\}}{C_1 \cup C_2}$$

Towards PL1 Resolution, ctd.

What about this:

→ $\{\{Nat(s(1)), \neg Nat(1)\}, \{Nat(1)\}\} \models \{Nat(s(1))\}$?

→ And $\{\{Nat(s(1)), \neg Nat(1)\}, \{Nat(1)\}\} \vdash \{Nat(s(1))\}$?

And what about this?

→ $\{\{Nat(s(x)), \neg Nat(x)\}, \{Nat(1)\}\} \models \{Nat(s(1))\}$?

→ But $\{\{Nat(s(x)), \neg Nat(x)\}, \{Nat(1)\}\} \vdash \{Nat(s(1))\}$?

Substitutions

Definition (Substitution). *A substitution* $s = \{\frac{x_1}{t_1}, \ldots, \frac{x_n}{t_n}\}$ *is a function that substitutes variables* x_i *for terms* t_i*, where* $x_i \neq t_i$ *for all* i*. Applying substitution* s *to an expression* φ *yields the expression* φs*, which is* φ *with all occurrences of* x_i *simultaneously replaced by* t_i*.*

$\rightarrow$ Variable instantiation and renaming, as used in the prenex and Skolem transformations as well as in the Herbrand expansion, is a special case of substitution.

Example: For $s = \{\frac{x}{y}, \frac{y}{h(a,b)}\}$, $P(x,y)s = P(y, h(a,b))$.

Remember: $x, y, z, v, w, \ldots$: variables; $a, b, c, d, e, \ldots$: constants.

Substitution Examples

Remember: $x, y, z, v, w, \ldots$: variables; $a, b, c, d, e, \ldots$: constants.

Examples: Can we apply a substitution to $P(x, f(y), b)$ so that it becomes:

1. $P(z, f(w), b)$: Yes: $s = \{\frac{x}{z}, \frac{y}{w}\}$
2. $P(x, f(a), c)$?
3. $P(y, f(h(a, b, w)), b)$?
4. $Q(x, f(y), b)$?
5. $P(x, f(f(y)), b)$?

Composing Substitutions

Definition (Composition). *Given substitutions s_1 and s_2, by $s_1 s_2$ we denote the composed substitution, a single substitution whose outcome is identical to $s_2 \circ s_1$.*

Example: With $s_1 = \{\frac{z}{g(x,y)}, \frac{v}{w}\}$ and $s_2 = \{\frac{x}{a}, \frac{y}{b}, \frac{w}{v}, \frac{z}{d}\}$, we have $P(x, y, z, v)s_1 s_2 =$

How to obtain $s_1 s_2$ given s_1 and s_2?

(i) Apply s_2 to the replacement terms t_i in s_1.

(ii) For any variable x_i replaced by s_2 but not by s_1, apply the respective variable/term pair $\frac{x_i}{t_i}$ of s_2.

(iii) Remove any pairs of variable x and term t where $x = t$.

Example: $\{\frac{z}{g(x,y)}, \frac{v}{w}\}\{\frac{x}{a}, \frac{y}{b}, \frac{w}{v}, \frac{z}{d}\} =$

Properties of Substitutions

For any formula φ and substitutions s_1, s_2, s_3:

$\rightarrow$ $(\varphi s_1)s_2 = \varphi(s_1 s_2)$ by definition (composing functions).

$\rightarrow$ $(s_1 s_2)s_3 = s_1(s_2 s_3)$ by definition (composing functions).

$\rightarrow$ $s_1 s_2 = s_2 s_1$?

(And by the way:) (idempotence)

Proposition. *A substitution $s = \{\frac{x_1}{t_1}, \ldots, \frac{x_n}{t_n}\}$ is idempotent, i.e., $\varphi ss = \varphi s$ for all φ, iff t_i does not contain x_j for $1 \leq i, j \leq n$.*

Proof. "$\Leftarrow$": The second application of s does not do anything because al x_i have been removed. "$\Rightarrow$": if t_i contains x_j then the second application of s replaces x_j with $t_j \neq x_j$.

Example: For $s = \{\frac{x}{y}, \frac{y}{h(a,b)}\}$,
$P(x,y)s = P(y, h(a,b)) \neq P(x,y)ss = P(h(a,b), h(a,b))$.

Unification

Definition (Unifier). *We say that a substitution s is a unifier for a set of expressions $\{E_1, \ldots, E_k\}$ if $E_i s = E_j s$ for all i, j.*

Notation: We'll usually write $\{E_i\}$ for $\{E_1, \ldots, E_k\}$.

Example: $\{P(x, f(y, z), b), P(x, f(b, w), b)\}$

$\rightarrow s = \{\frac{y}{b}, \frac{z}{w}, \frac{x}{h(a,b)}\}$?

$\rightarrow s = \{\frac{y}{b}, \frac{z}{w}\}$?

Definition (mgu). *We say that g is an mgu of $\{E_i\}$ if, for any unifier s of $\{E_i\}$, there exists a substitution s' such that $\{E_i\}s = \{E_i\}gs'$.*

$\rightarrow$ If any unifier exists, then an idempotent mgu exists.

$\rightarrow$ We'll next introduce an algorithm that finds it.

Unification Algorithm: What We Can *Not* Do

Example:

→ Can we unify $\{P(x, f(y), b), P(x, f(f(y)), b)\}$?

Unification Algorithm: Disagreement Set

Our unification algorithm (next slide) makes use of this notion:

Terminology: The disagreement set of a set of expressions $\{E_i\}$ is the set of sub-expressions $\{t_i\}$ of $\{E_i\}$ at the first position in $\{E_i\}$ for which some of the $\{E_i\}$ disagree.

Examples:

$\rightarrow \{P(x, c, f(y)), P(x, z, z)\}$: $\{c, z\}$

$\rightarrow \{P(x, a, f(y)), P(y, a, f(y))\}$: $\{x, y\}$

$\rightarrow \{P(v, f(z), g(w)), P(v, f(z), g(f(z)))\}$?

$\rightarrow \{P(v, f(z), g(w)), P(v, f(z), g(f(z))), P(v, f(z), f(x))\}$?

Unification Algorithm

Theorem. *The following algorithm succeeds if and only if there exists a unifier for $\{E_i\}$. In the positive case, the algorithm returns an idempotent mgu of $\{E_i\}$.* (Proof omitted.)

$k \leftarrow 0$, $T_k = \{E_i\}$, $s_k = \{\}$;
while T_k is not a singleton **do**
  Let D_k be the disagreement set of T_k;

/* if t_k contains x_k then unification is impossible, cf. slide 24 */
  Let $x_k, t_k \in D_k$ be a variable and term s.t. t_k does not contain x_k;
  if such x_k, t_k do not exist **then exit** with output "failure";

  $s_{k+1} \leftarrow s_k\{\frac{x_k}{t_k}\}$; /* t_k does not contain any of $x_1, \ldots, x_k$ */

  $T_{k+1} \leftarrow T_k\{\frac{x_k}{t_k}\}$; /* x_k does not occur in T_{k+1} */

  $k \leftarrow k + 1$;
endwhile
exit with output s_k;

Unification Algorithm: An Example

$$\{P(x, f(y), y), P(z, f(b), b)\}$$

Questionnaire

Question!

Can $\{Knows(John, x), Knows(x, Elizabeth)\}$ **be unified?**

(A): Yes (B): No

Question!

What about $\{Knows(John, x), Knows(y, Elizabeth)\}$**?**

(A): Yes (B): No

PL1 Resolution: Setup

We assume: Clausal normal form, variables standardized apart.

universal prefix + disjunction of literals

$$\forall x_1 \forall x_2 \forall x_3 \ldots \forall x_n (l_1 \vee \cdots \vee l_n)$$

$\rightarrow$ Written $\{l_1, \ldots, l_n\}$.

Example: $\{\{Nat(s(x)), \neg Nat(x)\}, \{Nat(y)\}\}$

Reminder: Terminology and Notation

- A literal l is an atom or the negation thereof; the negation of a literal is denoted $\bar{l}$ (e.g., $\overline{\neg Q} = Q$).
- A clause C is a set (=disjunction) of literals.
- Our input is a set Δ of clauses.
- The empty clause is denoted $\square$.
- A calculus is a set of inference rules.

PL1 Resolution: Setup, ctd.

Derivations: *We say that a clause C can be derived from Δ using calculus $\mathcal{R}$, written $\Delta \vdash_{\mathcal{R}} C$, if (starting from Δ) there is a sequence of applications of rules from $\mathcal{R}$, ending in C.*

$\rightarrow$ In contrast to propositional resolution, we will consider here three different resolution calculi $\mathcal{R}$.

Soundness: *A calculus $\mathcal{R}$ is sound if $\Delta \vdash_{\mathcal{R}} C$ implies $\Delta \models C$.*

Completeness: *A calculus $\mathcal{R}$ is refutation-complete if $\Delta \models \bot$ implies $\Delta \vdash_{\mathcal{R}} \square$, i.e., if Δ is unsatisfiable then we can derive the empty clause.*

- Together: Δ is unsatisfiable iff we can derive the empty clause.
- Propositional resolution is sound & refutation-complete or propositional Δ.

Reminder: Deduction $\approx$ Proof by Contradiction

To decide whether KB $\models \varphi$, decide satisfiability of $\psi := \text{KB} \cup \{\neg\varphi\}$: ψ is unsatisfiable iff KB $\models \varphi$.

Reminder: Propositional Resolution

Definition (Propositional Resolution). *Resolution uses the following inference rule (with exclusive union $\dot\cup$ meaning that the two sets are disjoint):*

$$\frac{C_1 \dot\cup \{l\}, C_2 \dot\cup \{\bar{l}\}}{C_1 \cup C_2}$$

If Δ contains parent clauses of the form $C_1 \dot\cup \{l\}$ and $C_2 \dot\cup \{\bar{l}\}$, the rule allows to add the resolvent clause $C_1 \cup C_2$. l and $\bar{l}$ are called the resolution literals.

Example: $\{P, \neg R\}$ resolves with $\{R, Q\}$ to $\{P, Q\}$.

Lemma. *The resolvent follows from the parent clauses.*

Proof. If $I \models C_1 \dot\cup \{l\}$ and $I \models C_2 \dot\cup \{\bar{l}\}$, then I must make at least one literal in $C_1 \cup C_2$ true.

Binary PL1 Resolution

Definition (Binary PL1 Resolution). *Binary PL1 resolution is the following inference rule:*

$$\frac{C_1 \dot\cup \{P_1\}, C_2 \dot\cup \{\neg P_2\}}{[C_1 \cup C_2]g}$$

If Δ contains parent clauses of the form $C_1 \dot\cup \{P_1\}$ and $C_2 \dot\cup \{\neg P_2\}$, where $\{P_1, P_2\}$ can be unified and g is an mgu thereof, the rule allows to add the resolvent clause $[C_1 \cup C_2]g$. P_1 and $\neg P_2$ are called the resolution literals.

Example: From $\{Nat(s(x)), \neg Nat(x)\}$ and $\{Nat(1)\}$

Lemma (Soundness). *The resolvent follows from the parent clauses.*

Why Do We Need To Standardize Variables Apart?

Example: $\Delta = \{\{Knows(John, x)\}, \{\neg Knows(x, Elizabeth), King(x)\}\}$

$\rightarrow$ We should be able to conclude that?

Unification 1: $\{P_1, P_2\} = \{Knows(John, x), Knows(x, Elizabeth)\}$

$\rightarrow$ Is there a unifier for $\{E_i\}$? No. We would have to substitute two different constants for x. (Cf. slide 28)

Unification 2: $\{P_1, P_2\} = \{Knows(John, x), Knows(y, Elizabeth)\}$

$\rightarrow$ Is there a unifier for $\{E_i\}$? Yes: $\{\frac{x}{Elizabeth}, \frac{y}{John}\}$. (Cf. slide 28)

$\rightarrow$ Standardizing the variables in clauses apart is sometimes necessary to allow unification.

($\rightarrow$ An alternative would be to not use unification, and instead substitute atoms separately to the same outcome; we don't consider this here.)

Questionnaire

Question!

Which are PL1 resolvents of
$\{\neg Chases(x, Garfield), Chases(Lassie, x)\}$ **and** $\{Chases(Bello, y)\}$**?**

(A): $\{Chases(Lassie, Bello)\}$

(B): $\{Chases(Garfield, Bello)\}$

(C): $\square$

(D): $\{Chases(Bello, Garfield)\}$

Binary PL1 Resolution: Examples

Clauses: $\{P(x), Q(f(x))\}$, $\{R(g(x)), \neg Q(f(a))\}$

→ Standardizing variables apart: $\{P(x), Q(f(x))\}$, $\{R(g(y)), \neg Q(f(a))\}$

→ MGU:

→ Resolvent:

Clauses: $\{P(x, g(c)), Q(x, a)\}$, $\{\neg P(y, g(c)), \neg R(b, z)\}$

→ Standardizing variables apart: (Nothing to do.)

→ MGU:

→ Resolvent:

Where Binary PL1 Resolution Fails

Example: $\Delta = \{\{P(x_1, x_2), P(x_2, x_1)\}, \{\neg P(y_1, y_2), \neg P(y_2, y_1)\}\}$

$\rightarrow$ Is Δ satisfiable?

$\rightarrow$ Can we derive $\square$ with binary PL1 resolution?

Notation: Define $\mathcal{R}_{Binary} :=$ {binary PL1 resolution}.

Theorem. *The calculus $\mathcal{R}_{Binary}$ is not refutation-complete.*

Proof. See example above.

However, $\mathcal{R}_{Binary}$ is sound:

Theorem. *The calculus $\mathcal{R}_{Binary}$ is sound.* (Proof: Lemma slide 33)

Solution 1: Full PL1 Resolution

→ Allow to unify *several* resolution literals:

Definition (Full PL1 Resolution). *Full PL1 resolution is the following inference rule:*

$$\frac{C_1 \dot\cup \{P_1^1, \ldots, P_1^n\}, C_2 \dot\cup \{\neg P_2^1, \ldots, \neg P_2^m\}}{[C_1 \cup C_2]g}$$

where $\{P_1^1, \ldots, P_1^n, P_2^1, \ldots, P_2^m\}$ *can be unified and* g *is an mgu thereof.*

Example: $\Delta = \{\{P(x_1, x_2), P(x_2, x_1)\}, \{\neg P(y_1, y_2), \neg P(y_2, y_1)\}\}$

→ Can we derive $\square$ with full PL1 resolution?

Notation: Define $\mathcal{R}_{Full} := \{\text{full PL1 resolution}\}$.

Theorem. *The calculus* $\mathcal{R}_{Full}$ *is sound.*

Proof. It suffices to show that, for each application of the rule, the resolvent follows from the parents. That can be shown with the same argument as for binary PL1 resolution (Lemma slide 33).

Solution 2: Binary PL1 Resolution + Factoring

→ Allow to unify literals *within* a clause:

Definition (Factoring). *Factoring is the following inference rule:*

$$\frac{C_1 \dot\cup \{l_1\} \dot\cup \{l_2\}}{[C_1 \cup \{l_1\}]g}$$

where $\{l_1, l_2\}$ *can be unified and* g *is an mgu thereof.* $[C_1 \cup \{l_1\}]g$ *is called a factor of the parent clause* $C_1 \dot\cup \{l_1\} \dot\cup \{l_2\}$.

Example: $\Delta = \{\{P(x_1, x_2), P(x_2, x_1)\}, \{\neg P(y_1, y_2), \neg P(y_2, y_1)\}\}$

→ How can we apply factoring?

Notation: Define $\mathcal{R}_{FactBin} :=$ {binary PL1 resolution,factoring}.

Theorem. *The calculus* $\mathcal{R}_{FactBin}$ *is sound.*

Proof. Due to the universal quantification, the factor follows from its parent. Done with Lemma slide 33.

What About Completeness? The Lifting Lemma

Lemma (Lifting Lemma). *Let C_1 and C_2 be two clauses with no shared variables, and let C_1^g and C_2^g be ground instances of C_1 and C_2. Say that C^g is a resolvent of C_1^g and C_2^g. Then there exists a clause C such that C^g is a ground instance of C, and:*

(i) *C can be derived from C_1 and C_2 using $\mathcal{R}_{Full}$.*

(ii) *C can be derived from C_1 and C_2 using $\mathcal{R}_{FactBin}$.*

Proof Sketch.

Example: $C_1 = \{P(x_1), P(x_2), R(z)\}$, $C_2 = \{\neg P(y_1), \neg P(y_2), R(z')\}\}$
E.g., $C_1^g = \{P(o), R(o)\}$ and $C_2^g = \{\neg P(o), R(o)\}$. Then

What About Completeness? Proof

Theorem. *The calculi $\mathcal{R}_{Full}$ and $\mathcal{R}_{FactBin}$ are refutation-complete.*

Proof:

Any set θ of PL1 formulas is representable in clausal form Δ.

$\downarrow$

Assume Δ is unsatisfiable.

$\downarrow$ $\longleftarrow$ Herbrand, prop. compactness

Some finite set Δ' of ground instances is unsatisfiable.

$\downarrow$ $\longleftarrow$ Prop. resolution completeness

Propositional resolution can derive $\square$ from Δ'.

$\downarrow$ $\longleftarrow$ Lifting Lemma

Each of $\mathcal{R}_{Full}$ and $\mathcal{R}_{FactBin}$ can derive $\square$ from Δ.

Questionnaire

Question!

Is full PL1 resolution guaranteed to terminate after a finite number of rule applications?

(A): Yes.

(B): No.

When PL1 Resolution Reasons About Blocks . . .

Example "Integers"

Formula: $\forall x[\exists y(Equals(x, succ(y)))]$ ("For every integer x, there is y so that $x = y + 1$")

$\rightarrow$ Is this satisfiable?

Axiomatizing $succ, Equals$:

"1 is not the successor of anybody:" $\forall x \neg Equals(1, succ(x))$

Resolution refutation:

Col. West, *a Criminal?*

From [Russell and Norvig (2010)]:

> *The law says it is a crime for an American to sell weapons to hostile nations. The country Nono, an enemy of America, has some missiles, and all of its missiles were sold to it by Colonel West, who is American.*

$\rightarrow$ Prove that Col. West is a criminal.

Convention: In what follows, for better readability we will sometimes write implications $P \wedge Q \wedge R \rightarrow S$ instead of clauses $\neg P \vee \neg Q \vee \neg R \vee S$.

Col. West, *a Criminal?* Clauses

Col. West, *a Criminal!* PL1 Resolution Proof

Curiosity Killed the Cat?

From [Russell and Norvig (2010)]:

Everyone who loves all animals is loved by someone.
Anyone who kills an animal is loved by noone.
Jack loves all animals.
Cats are animals.
Either Jack or curiosity killed the cat (whose name is "Garfield").

→ Prove that curiosity killed the cat.

Convention: In what follows, for better readability we will sometimes write implications $P \wedge Q \wedge R \rightarrow S$ instead of clauses $\neg P \vee \neg Q \vee \neg R \vee S$.

Curiosity Killed the Cat? Clauses

Curiosity Killed the Cat! PL1 Resolution Proof

Summary

- The Herbrand universe is the set of all ground terms that can be built from the symbols used in a set θ of PL1 formulas. The (propositional-logic) Herbrand expansion instantiates the formulas with these terms, and is satisfiable iff θ is.
- For unsatisfiable θ, we can always find an unsatisfiable finite subset of the Herbrand expansion.
- PL1 resolution reasons directly about PL1 formulas (in clausal normal form) It relies on unification to compare PL1 terms.
- Binary PL1 resolution is like propositional resolution with unification. It is not refutation-complete.
- To obtain a complete PL1 resolution calculus, we can either allow to unify sets of resolution literals (full PL1 resolution), or to unify literals within clauses (factoring).
- The set of satisfiable PL1 formulas is not recursively enumerable. Thus, neither the reduction to propositional logic, nor PL1 resolution, guarantee to terminate in finite time.

Topics We Didn't Cover Here

PL1 is very expressive, but: (some people just can't get enough)

- Second-Order Logic: Quantification over predicates.

$$\forall x, y[Equals(x, y) \leftrightarrow [\forall p(p(x) \leftrightarrow p(y))]]$$

"We define x and y to be "Equal" iff their behavior with respect to all predicates is identical."

- Temporal Logic: Quantification over future behaviors.

$$\mathbf{AG}[\varphi \implies \mathbf{EF}\psi]$$

"For **A**ll futures, we **G**lobally have that, if $s \models \varphi$, then there **E**xists a future from s on which **F**inally we have ψ."

And what else? There's of course also *lots* of algorithmic stuff *within* PL1 that we didn't cover.

→ If you want to know all about this, take the "Automated Reasoning" courses.

Reading

- *Chapter 9: Inference in First-Order Logic*, Section 9.1, 9.2, and 9.5 [Russell and Norvig (2010)].

  Content: What I cover in "Reduction to Propositional Reasoning" is distributed in RN over Section 9.1, which gives a very brief sketch of the idea, and Section 9.5.4 which contains a summary of Herbrand's results.

  Section 9.2.2 contains a much less formal/detailed account of what I cover in "Substitutions, and Unification".

  Section 9.5.2 briefly outlines (in half a page!) what I cover in "Predicate Logic Resolution". Section 9.5.3 pretty much coincides with my "On Criminals and Cats".

  Sections 9.3 and 9.4 describe "forward chaining" and "backward chaining", relevant to Databases and Logic Programming. Nice background reading! The same applies to a few gems here and there, such as a summary of Gödel's incompleteness theorem.

  $\rightarrow$ Overall: As usual, lacks rigor, but covers a great breadth of subjects and provides nice complementary reading. Can't replace the lecture.

References I

Stuart Russell and Peter Norvig. *Artificial Intelligence: A Modern Approach (Third Edition)*. Prentice-Hall, Englewood Cliffs, NJ, 2010.