Automated Reasoning

6

AI Slides 10e©Lin Zuoquan@PKU 1998-2025

6 Automated Reasoning

- 6.1 Automated theorem proving
- 6.2 Forward and backward chaining
- 6.3 Resolution
- 6.4 Model checking⁺

A brief history of reasoning

Automated reasoning: reasoning completely automatically by computer programs

450B.C.	Stoics	propositional logic
322B.C.	Aristotle	syllogisms (inference rules), quantifiers
1565	Cardano	probability theory (propositional logic + uncertainty)
1847	Boole	propositional logic (again)
1879	Frege	first-order logic
1922	Wittgenstein	proof by truth tables
1930	Gödel	complete algorithm for FOL
1930	Herbrand	complete algorithm for FOL (reduce to propositional)
1931	Gödel	incomplete algorithm for arithmetic
1960	Davis/Putnam	"practical" algorithm for propositional logic
1965	Robinson	"practical" algorithm for FOL—resolution

Automated theorem proving (ATP): proving (mathematical) theorems by computer programs

Proof methods divide into (roughly) two kinds

Application of inference rules

- Legitimate (sound) generation of new sentences from old
- Proof = a sequence of inference rule applications

Can use inference rules as operators in a standard search alg. Inference rules include

- forward chaining, backward chaining, resolution

Model checking

truth table enumeration (always exponential in *n*) improved backtracking, e.g., DPLL algorithm heuristic search in model space (sound but incomplete) e.g., min-conflicts-like hill-climbing algorithms

Proofs

Sound inference: find α such that $KB \vdash \alpha$ Proof process is a <u>search</u>, operators are inference rules

Modus Ponens (MP)

$$\frac{\alpha, \quad \alpha \Rightarrow \beta}{\beta} \qquad \frac{At(lin, pku) \quad At(lin, pku) \Rightarrow Ok(lin)}{Ok(lin)}$$

And-Introduction (AI)

lpha eta	Ok(lin)	AImajor(lin)
$\overline{\alpha \wedge \beta}$	Ok(Lin)	$\land AImajor(in)$

Universal instantiation (UI)

Every instantiation of a universally quantified sentence is entailed by it: $\forall v \ \alpha$

 $\frac{\sqrt{v}}{\text{SUBST}(\{v/g\},\alpha)}$

for any variable \boldsymbol{v} and ground term \boldsymbol{g}

E.g., $\forall x \ King(x) \land Greedy(x) \Rightarrow Evil(x)$ yields $King(john) \land Greedy(john) \Rightarrow Evil(john)$ $King(richard) \land Greedy(richard) \Rightarrow Evil(richard)$ $King(father(john)) \land Greedy(father(john)) \Rightarrow Evil(father(john))$: For any sentence α , variable v, and constant symbol k that does not appear elsewhere in the knowledge base:

 $\frac{\exists v \ \alpha}{\text{Subst}(\{v/k\}, \alpha)}$

provided \boldsymbol{c} is a new constant symbol, called a Skolem constant

Another example: from $\exists x \ d(x^y)/dy = x^y$ we obtain

 $d(e^y)/dy = e^y$

provided e is a new constant symbol

Instantiation

UI can be applied several times to **add** new sentences; the new KB is logically equivalent to the old

El can be applied once to **replace** the existential sentence; the new KB is **not** equivalent to the old, but is satisfiable iff the old KB was satisfiable

Example: a proof

bob is a buffalo	1.	Buffalo(bob)
pat is a pig	2.	Pig(pat)
Buffaloes outrun pigs	3.	$\forall x,y \;\; Buffalo(x) \land Pig(y) \Rightarrow Faster(x,y)$
bob outruns pat		$Buffalo(bob) \land Pig(pat) \Rightarrow Faster(bob, pat)$
UE 3, $\{x/bob, y/pat\}$		

Al 1 & 2	4. $Buffalo(bob) \wedge Pig(pat)$



UE 3, $\{x/bob, y/pat\}$ 5. $Buffalo(bob) \land Pig(pat) \Rightarrow Faster(bob, pat)$

AI Slides 10e©Lin Zuoquan@PKU 1998-2025

MP 6 & 7	6. Faster(bob, pat)

AI Slides 10e©Lin Zuoquan@PKU 1998-2025

Search with inference rules

Operators are inference rules States are sets of sentences Goal test checks state to see if it contains a query sentence



AI, UE, MP are common inference patterns

<u>Problem</u>: branching factor huge, esp. for UE

<u>Idea</u>: find a substitution that makes the rule premise match some known facts

 \Rightarrow a single, more powerful inference rule

Forward and backward chaining

Modus Ponens (for Horn Form): complete for Horn KBs $\frac{\alpha_1, \dots, \alpha_n, \qquad \alpha_1 \wedge \dots \wedge \alpha_n \Rightarrow \beta}{\beta}$

Can be used with forward chaining or backward chaining

These algorithms are very natural and run in **linear** time

Clause form

Clause Form (restricted) KB =conjunction of clauses (CNF) Recall: Clause = disjunction of literals • proposition symbol; or • (conjunction of symbols) \Rightarrow symbol (i.e., conjunction of literals) E.g., $C \land (B \Rightarrow A) \land (C \land D \Rightarrow B)$ i.e., $C \land (\neg B \lor A) \land (\neg C \lor \neg D \lor B)$

Horn clause = a clause in which at most one is positive literal

Definite clause = a clause in which exactly one is positive literal all definite clauses are Horn clauses

Goal clauses = clauses with no positive literals

FC Idea: fire any rule whose premises are satisfied in the KB add its conclusion to the KB, until query is found



















Forward chaining algorithm

```
def PL-FC-Ask(KB, q)
   inputs: KB, the knowledge base, a set of propositional definite clauses
             q, the query, a proposition symbol
   count \leftarrow a table, where count[c] is the number of symbols in c's premise
   inferred \leftarrow a table, where inferred[s] is initially false for all symbols
   queue \leftarrow a queue of symbols, initially symbols known to be true in KB
   while queue is not empty do // not yet processed
        p \leftarrow \text{POP}(queue)
        if p=q then return true
        if inferred[p]=false then
             inferred[p] \leftarrow true
             for each clause c in KB where p is in c.PREMISE do//implication
                  decrement count[c]
                  if count[c] = 0 then add c.CONCLUSION to queue
   return false
```

FC derives every atomic sentence that is entailed by ${\bf Horn}~KB$

1. FC reaches a fixed point where no new atomic sentences are derived

2. Consider the final state as a model m, assigning true/false to symbols

- 3. Every clause in the original KB is true in m **Proof**: Suppose a clause $a_1 \land \ldots \land a_k \Rightarrow b$ is false in mThen $a_1 \land \ldots \land a_k$ is true in m and b is false in mTherefore the algorithm has not reached a fixed point
- 4. Hence m is a model of KB
- 5. If $KB \models q$, q is true in **every** model of KB, including m
- Idea: construct any model of KB by sound inference, check α

BC Idea: work backwards from the query qto prove q by BC check if q is known already, or prove by BC all premises of some rule concluding q

Avoid loops: check if new subgoal is already on the goal stack

Avoid repeated work: check if new subgoal 1) has already been proved true, or

2) has already failed

```
Algorithm: PL-BC-Ask?
(ref. FOL-BC-Ask in later)
```


















Example: backward chaining



Example: backward chaining



Forward vs. backward chaining

FC is data-driven, cf. automatic, unconscious processing e.g., object recognition, routine decisions

May do lots of work that is irrelevant to the goal

BC is goal-driven, appropriate for problem-solving e.g., Where are my keys? How do I get into a PhD program?

Complexity of BC can be **much less** than linear in size of KB

Forward and backward chaining are **complete for Horn KBs** but incomplete for full FOL

E.g., from

 $\begin{array}{l} PhD(x) \Rightarrow HighlyQualified(x) \\ \neg PhD(x) \Rightarrow EarlyEarnings(x) \\ HighlyQualified(x) \Rightarrow Rich(x) \\ EarlyEarnings(x) \Rightarrow Rich(x) \end{array}$

should be able to infer Rich(Me), but FC/BC won't do it Does a complete algorithm exist??

Resolution

- Propositional resolution
- Unification
- First-order resolution

Propositional resolution

Entailment in PL is decidable:

can prove that α if $KB \models \alpha$ or $KB \not\models \alpha$

Resolution is a refutation procedure:

to prove $KB \models \alpha$, show that $KB \land \neg \alpha$ is unsatisfiable

Resolution uses KB, $\neg \alpha$ in CNF

Resolution inference rule combines two clauses to make a new one



C is called a resolvent of input clauses C_1 , C_2

Inference continues until an empty clause { } is derived (contrad.)

Resolution



Given a clause of the form $\ell_1 \lor \cdots \lor \ell_k$ containing some literal ℓ_i , and a clause of the form $m_1 \lor \cdots \lor m_n$ containing some literal m_j , where ℓ_i and m_j are complementary literals, infer the clause consisting of those literals in the first clause other than ℓ_i and those in the second other than m_j , i.e.,

 $\ell_1 \lor \cdots \lor \ell_{i-1} \lor \ell_{i+1} \lor \cdots \lor \ell_k \lor m_1 \lor \cdots \lor m_{j-1} \lor m_{j+1} \lor \cdots \lor m_n$ which is a resolvent of the two input clauses w.r.t. ℓ_i and m_j

A resolution derivation (or proof) of a clause c from a set of clauses S is a sequence of clauses c_1, \dots, c_n , where the last clause, c_n , is c, and where each c_i is either an element of S or a resolvent of two earlier clauses in the derivation

write $S \vdash_i c$ (*i* is resolution, hereafter simply \vdash) if there is a derivation of *c* from *S* write $\{\} \vdash c$, simply $\vdash c$, called *c* is a theorem

Conversion to CNF

 $B_{1,1} \Leftrightarrow (P_{1,2} \lor P_{2,1})$

- 1. Eliminate \Leftrightarrow , replacing $\alpha \Leftrightarrow \beta$ with $(\alpha \Rightarrow \beta) \land (\beta \Rightarrow \alpha)$ $(B_{1,1} \Rightarrow (P_{1,2} \lor P_{2,1})) \land ((P_{1,2} \lor P_{2,1}) \Rightarrow B_{1,1})$
- 2. Eliminate \Rightarrow , replacing $\alpha \Rightarrow \beta$ with $\neg \alpha \lor \beta$

 $(\neg B_{1,1} \lor P_{1,2} \lor P_{2,1}) \land (\neg (P_{1,2} \lor P_{2,1}) \lor B_{1,1})$

- 3. Move \neg inwards using de Morgan's rules and double-negation $(\neg B_{1,1} \lor P_{1,2} \lor P_{2,1}) \land ((\neg P_{1,2} \land \neg P_{2,1}) \lor B_{1,1})$
- 4. Apply distributivity law (\lor over \land) and flatten

 $(\neg B_{1,1} \lor P_{1,2} \lor P_{2,1}) \land (\neg P_{1,2} \lor B_{1,1}) \land (\neg P_{2,1} \lor B_{1,1})$

Resolution algorithm

```
Proof by contradiction, i.e., show KB \wedge \neg \alpha unsatisfiable
```

```
\begin{aligned} & \operatorname{def} \operatorname{PL-Resolution}(KB, \alpha) \\ & \operatorname{inputs:} KB, \text{ the knowledge base, a sentence in propositional logic} \\ & \alpha, \text{ the query, a sentence in propositional logic} \\ & clauses \leftarrow \text{ the set of clauses in the CNF representation of } KB \land \neg \alpha \\ & new \leftarrow \{ \} \\ & \operatorname{while} true \operatorname{do} \\ & \operatorname{for each} C_i, C_j \operatorname{in} clauses \operatorname{do} \\ & resolvents \leftarrow \operatorname{PL-RESOLVE}(C_i, C_j) \\ & \operatorname{if} resolvents \operatorname{contains} \text{ the empty clause then return } true \\ & new \leftarrow new \cup \ resolvents \\ & \operatorname{if} new \subseteq clauses \operatorname{then return} false // \operatorname{unsatisfiable} \\ & clauses \leftarrow clauses \cup new \end{aligned}
```

Example: resolution

 $KB = (B_{1,1} \iff (P_{1,2} \lor P_{2,1})) \land \neg B_{1,1} \ \alpha = \neg P_{1,2}$



Note: need only convert KB to CNF once

- can handle multiple queries with same KB
- \bullet after addition of new fact $\alpha,$ can simply add new clauses α' to KB

Derivation and entailment^{*}

Claim: resolvent is entailed by input clauses

Proof: Suppose $m \models p \lor \alpha$ and $m \models \neg p \lor \beta$ Case 1: $m \models p$ then $m \models \beta$, so $m \models (\alpha \lor \beta)$ Case 2: $m \not\models p$ then $m \models \beta$, so $m \models (\alpha \lor \beta)$ Either way, $m \models (\alpha \lor \beta)$ $\{(p \lor \alpha), (\neg p \lor \beta)\} \models (\alpha \lor \beta)$

Special case: c and $\neg c$ resolve to $\{ \}$ i.e., $\{c, \neg c\}$ is unsatisfiable

Derivation and entailment^{*}

Can extend the previous argument to derivations

```
If KB \vdash c then KB \models c

Proof: by induction on the length of the derivation

Show (by looking at the two cases) that KB \models c_i

But the converse does not hold in general
```

Can have $KB \models c$ without having $KB \vdash c$ E.g., $\neg p \models \neg p \lor \neg q$ but no derivation

Note: resolution is sound but not complete in general

Soundness and completeness of resolution

Theorem: i (resolution) is sound and refutation complete if

$KB \vdash_i \alpha \text{ iff } KB \models \alpha$

A set of clauses is unsatisfiable iff

the resolution closure of those clauses contains the empty clause

– provides method for determining satisfiability: search all derivations for $\{\ \}$

- so provides a method for determining all entailments

 \mathbf{Proof} of soundness

– Consider the complementary literals ℓ_i, m_j , easy to check

Resolution closure RC(S) (of a set of clauses S) denotes the set of all clauses derivable by resolution; RC(S) must be finite 1. Consider the contrapositive: if the closure RC(S) does not contains the empty clause, then S is satisfiable

2. Construct a model for S with suitable truth values for the symbols P_1, \dots, P_k that appear in S:

For i from 1 to k

– If a clause in RC(S) contains $\neg P_i$ and all its other literals are false under the assignment chosen for P_1, \dots, P_{i-1} , then assign false to P_i

– Otherwise, assign true to P_i

3. This assignment to P_1, \dots, P_k is a model of S

Proof by contradiction: at some stage i in the sequence, assigning symbol P_i causes some clause C to become false

We can get the inference immediately if we can find a substitution θ such that King(x) and Greedy(x) match King(john) and Greedy(y)

 $\theta = \{x/john, y/john\} \text{ works}$



We can get the inference immediately if we can find a substitution θ s.t. King(x) and Greedy(x) match King(john) and Greedy(y)

 $\theta = \{x/john, y/john\} \text{ works}$

We can get the inference immediately if we can find a substitution θ such that King(x) and Greedy(x) match King(john) and Greedy(y)

 $\theta = \{x/john, y/john\} \text{ works}$



We can get the inference immediately if we can find a substitution θ such that King(x) and Greedy(x) match King(john) and Greedy(y)

 $\theta = \{x/john, y/john\} \text{ works}$

p	q	$ \theta$
Knows(john, x)	Knows(john, jane)	$\{x/jane\}$
Knows(john, x)	Knows(y, lin)	$\{x/lin, y/john\}$
Knows(john, x)	Knows(y, mother(y))	$\{y/john, x/mother(john)\}$
Knows(john, x)	Knows(x, lin)	

We can get the inference immediately if we can find a substitution θ such that King(x) and Greedy(x) match King(john) and Greedy(y)

 $\theta = \{x/john, y/john\} \text{ works}$

 $\mathrm{Unify}(\alpha,\beta)=\theta \text{ if } \alpha\theta\!=\!\beta\theta$

Standardizing apart eliminates overlap of variables, e.g., Knows(z, lin)

 θ is a most general unifier (MGU, written as UNIFY) of literals l_1 and l_2 iff

- 1. θ unifies l_1 and l_2
- 2. for any other unifier θ' , there is a another substitution θ^* s.t. $\theta' = \theta \theta^*$

where $\theta\theta^*$ requires applying θ^* to terms in θ

E.g., P(g(x), f(x), z), $\neg P(y, f(w), a)$ an MGU is $\theta = \{x/w, y/g(w), z/a\}$

Theorem: Can limit search to most general unifiers only without loss of completeness

There is a better **linear** algorithm

Algorithm of computing MGUs

Given a set of literals $\{l_i\}$ (usually only two literals)

- 1. Start with $\theta := \{\}$.
- 2. If all the $\alpha\theta$ are identical, then done; otherwise, get disagreement set, DS e.g $P(a, f(a, g(z)), P(a, f(a, u), DS = \{u, g(z)\})$
- 3. Find a variable $v \in DS$, and a term $t \in DS$ not containing v; If not, fail.
- 4. $\theta := \theta\{v/t\}$
- 5. Go to 2

There is a better *linear* algorithm

Generalized Modus Ponens (GMP)

$$\begin{array}{ll} \underline{p_1', p_2', \ldots, p_n', (p_1 \wedge p_2 \wedge \ldots \wedge p_n \Rightarrow q)}_{q\theta} & \text{where } p_i'\theta = p_i\theta \text{ for all } i \\ p_1' \text{ is } King(john) & p_1 \text{ is } King(x) \\ p_2' \text{ is } Greedy(y) & p_2 \text{ is } Greedy(x) \\ \theta \text{ is } \{x/john, y/john\} \ q \text{ is } Evil(x) \\ q\theta \text{ is } Evil(john) \end{array}$$

GMP used with KB of definite clauses (exactly one positive literal) All variables assumed universally quantified

Note: Need to replace all variables in its arguments of a rule with new ones that have not been used before

(variable renaming, STANDARDIZE-VARIABLES function).

Hint: Special interesting for rule-based systems

Need to show that

$$p_1', \ldots, p_n', (p_1 \wedge \ldots \wedge p_n \Rightarrow q) \models q\theta$$

provided that $p_i'\theta = p_i\theta$ for all i

Lemma: For any definite clause $p\text{, we have }p\models p\theta$ by UI

- **1.** $(p_1 \land \ldots \land p_n \Rightarrow q) \models (p_1 \land \ldots \land p_n \Rightarrow q)\theta = (p_1 \theta \land \ldots \land p_n \theta \Rightarrow q\theta)$
- **2.** $p_1', \ldots, p_n' \models p_1' \land \ldots \land p_n' \models p_1' \theta \land \ldots \land p_n' \theta$
- 3. From 1 and 2, $q\theta$ follows by ordinary Modus Ponens

Example: a small KB

... it is a crime for an American to sell weapons to hostile nations $American(x) \wedge Weapon(y) \wedge Sells(x,y,z) \wedge Hostile(z) \Rightarrow Criminal(x)$

Nono . . . has some missiles, i.e., $\exists x \ Owns(Nono, x) \land Missile(x)$ $Owns(Nono, M_1)$ and $Missile(M_1)$

... all of its missiles were sold to it by Colonel West

 $\forall x \ Missile(x) \land Owns(Nono, x) \Rightarrow Sells(West, x, Nono)$

Missiles are weapons

 $Missile(x) \Rightarrow Weapon(x)$

An enemy of America counts as "hostile"

 $Enemy(x, America) \Rightarrow Hostile(x)$

West, who is American . . .

American(West)

The country Nono, an enemy of America . . .

Enemy(Nono, America)

Forward and backward chaining

Recall FC and BC in propositional level, and extend to first-order case FC is data-driven BC is goal-oriented the basis for <u>logic programming</u>, e.g., Prolog (More complications help to avoid infinite loops)

Two chainings: find any solution, find all solutions

Forward chaining algorithm $^{\#}$

```
def FOL-FC-Ask(KB, \alpha)
   inputs: KB, a set of first-order definite clauses
               \alpha, the query (an atomic sentence)
   while true do
   new \leftarrow \{\} // The set of new sentences inferred on each iteration
   for each rule in KB do
       (p_1 \land \ldots \land p_n \Rightarrow q) \leftarrow \text{STANDARDIZE-VARIABLES}(rule)
       for each \theta s.t. SUBST(\theta, p_1 \land \ldots \land p_n) = SUBST(\theta, p'_1 \land \ldots \land p'_n)
                  for some p'_1, \ldots, p'_n in KB
             q' \leftarrow \text{SUBST}(\theta, q)
             if q' does not unify with some sentence already in KB or new then
                 add q' to new
                 \theta \leftarrow \text{UNIFY}(q', \alpha)
                 if \theta is not failure then return \theta
         if new={} then return false
         add new to KB
```

American(West)

Missile(M1)

Owns(Nono,M1)

Enemy(*Nono*,*America*)

Hint: can you notice that $FOL\mathchar`EC-Ask$ differs from $PL\mathchar`EC-ENTAIL?$



Forward chaining proof



Properties of forward chaining

Sound and complete for first-order definite clauses (proof similar to propositional proof)

Datalog = first-order definite clauses + **no functions** (e.g., crime KB)

FC terminates for Datalog in poly iterations: at most $p \cdot n^k$ literals

Logica (logic+aggregation, Google 2021) compiles to SQL and run on Google BigQuery

May not terminate in general if α is not entailed

This is unavoidable: entailment with definite clauses is semidecidable

Efficiency of forward chaining

Simple observation: no need to match a rule on iteration k if a premise wasn't added on iteration k-1

 \Rightarrow match each rule whose premise contains a newly added literal

Matching itself can be expensive

Database indexing allows O(1) retrieval of known facts e.g., query Missile(x) retrieves $Missile(M_1)$

Matching conjunctive premises against known facts is NP-hard

Forward chaining is widely used in deductive databases

Hard matching example



 $\begin{array}{ll} \textit{Diff}(wa,nt) \land \textit{Diff}(wa,sa) \land \\ \textit{Diff}(nt,q)\textit{Diff}(nt,sa) \land \\ \textit{Diff}(q,nsw) \land \textit{Diff}(q,sa) \land \\ \textit{Diff}(nsw,v) \land \textit{Diff}(nsw,sa) \land \\ \textit{Diff}(v,sa) \Rightarrow \textit{Colorable}() \\ \textit{Diff}(\textit{Red},\textit{Blue}) \quad \textit{Diff}(\textit{Red},\textit{Green}) \\ \textit{Diff}(\textit{Green},\textit{Red}) \quad \textit{Diff}(\textit{Green},\textit{Blue}) \\ \textit{Diff}(\textit{Blue},\textit{Red}) \quad \textit{Diff}(\textit{Blue},\textit{Green}) \\ \end{array}$

Colorable() is inferred iff the CSP has a solution CSPs include 3SAT as a special case, hence matching is NP-hard

Backward chaining $algorithm^{\#}$

Implement as a generator (multiple substitutions), AND-OR search

```
def FOL-BC-Ask(KB, query)
   return FOL-BC-OR(KB, query, {}) //AND-OR search
def FOL-BC-OR(KB, goal, \theta) // OR because querying goal by any rule
   for each rule in FETCH-RULES-FOR-GOAL(KB, goal) do
      (lhs \Rightarrow rhs) \leftarrow \text{STANDARDIZE-VARIABLES}(rule)
      for each \theta' in FOL-BC-AND(KB, lhs, UNIFY(rhs, goal, \theta)) do
         yield \theta' // return by a generator for multiple substitutions
def FOL-BC-AND(KB, goal, \theta) // AND because lhs is a list of conjuncts
   if \theta = failure then return
   else if LENGTH(qoal) = 0 then yield \theta
   else
      first, rest \leftarrow FIRST(qoal), REST(qoal)
      for each \theta' in FOL-BC-OR(KB, SUBST(\theta, first), \theta) do
         for each \theta'' in FOL-BC-AND(KB, rest), \theta') do
            vield \theta''
```

Example: backward chaining

Criminal(West)

Example: backward chaining












Properties of backward chaining

Depth-first recursive proof search: space is linear in the size of proof

Incomplete due to infinite loops

 \Rightarrow fix by checking the current goal against every goal on the stack

Inefficient due to repeated subgoals (both success and failure) \Rightarrow fix using caching of previous results (extra space!)

Widely used for logic programming

 $\frac{\ell_1 \vee \cdots \vee \ell_k, \quad m_1 \vee \cdots \vee m_n}{(\ell_1 \vee \cdots \vee \ell_{i+1} \vee \cdots \vee \ell_k \vee m_1 \vee \cdots \vee m_{j-1} \vee m_{j+1} \vee \cdots \vee m_n)\theta}$ where UNIFY $(\ell_i, \neg m_j) = \theta$.

E.g.

 $\frac{\neg Rich(x) \lor Unhappy(x)}{Rich(lin)}$

with $\theta = \{x/lin\}$

Apply resolution steps to $CNF(KB \land \neg \alpha)$; complete for FOL

Conjunctive Normal Form

Any FOL KB can be converted to CNF

- 1. Replace $P \Rightarrow Q$ by $\neg P \lor Q$
- 2. Move \neg inwards, e.g., $\neg \forall x P$ becomes $\exists x \neg P$
- 3. Standardize variables apart, e.g., $\forall x\,P \lor \exists x\,Q$ becomes $\forall x\,P \lor \exists y\,Q$
- 4. Move quantifiers left in order, e.g., $\forall x P \lor \exists x Q$ becomes $\forall x \exists y P \lor Q$
- 5. Eliminate \exists by Skolemization (next slide)
- 6. Drop universal quantifiers
- 7. Distribute \wedge over \vee , e.g., $(P \wedge Q) \vee R$ becomes $(P \vee Q) \wedge (P \vee R)$

 $\exists x \operatorname{Rich}(x)$ becomes $\operatorname{Rich}(c)$ where c is a new Skolem constant

More tricky when \exists is inside \forall E.g., "Everyone has a heart" $\forall x \ .Person(x) \Rightarrow \exists y \ .Heart(y) \land Has(x, y)$

Incorrect:

 $\forall x \ .Person(x) \Rightarrow Heart(H1) \land Has(x,H1)$

Correct:

 $\forall x : Person(x) \Rightarrow Heart(H(x)) \land Has(x, H(x))$ where *H* is a new symbol (Skolem function)

Skolem function arguments: all <u>enclosing</u> universally quantified variables

Everyone who loves all animals is loved by someone $\forall x \ .[\forall y \ Animal(y) \Rightarrow Loves(x, y)] \Rightarrow [\exists y \ Loves(y, x)]$

1. Eliminate biconditionals and implications $\forall x \ .[\neg \forall y \ \neg Animal(y) \lor Loves(x, y)] \lor [\exists y \ Loves(y, x)]$

2. Move \neg inwards: $\neg \forall x, p \equiv \exists x \neg p, \neg \exists x, p \equiv \forall x \neg p$

 $\begin{array}{l} \forall x \quad . [\exists y \quad \neg (\neg Animal(y) \lor Loves(x,y))] \lor [\exists y \quad Loves(y,x)] \\ \forall x \quad . [\exists y \quad \neg \neg Animal(y) \land \neg Loves(x,y)] \lor [\exists y \quad Loves(y,x)] \\ \forall x \quad . [\exists y \quad Animal(y) \land \neg Loves(x,y)] \lor [\exists y \quad Loves(y,x)] \end{array}$

Conversion to CNF

3. Standardize variables: each quantifier should use a different one

 $\forall x \ . [\exists y \ Animal(y) \land \neg Loves(x,y)] \lor [\exists z \ Loves(z,x)]$

4. Skolemize: a more general form of existential instantiation Each existential variable is replaced by a Skolem function of the enclosing universally quantified variables

 $\forall x \ .[Animal(F(x)) \land \neg Loves(x, f(x))] \lor Loves(g(x), x)$

5. Drop universal quantifiers

 $[Animal(f(x)) \land \neg Loves(x, f(x))] \lor Loves(g(x), x)$

6. Distribute \wedge over \vee

 $[Animal(f(x)) \lor Loves(g(x), x)] \land [\neg Loves(x, f(x)) \lor Loves(g(x), x)]$

To prove α

- negate it
- convert to CNF
- add to CNF KB
- infer contradiction

E.g., to prove $Rich(me)\text{, add }\neg Rich(me)$ to the CNF KB

 $\begin{array}{l} \neg PhD(x) \lor HighlyQualified(x) \\ PhD(x) \lor EarlyEarnings(x) \\ \neg HighlyQualified(x) \lor Rich(x) \\ \neg EarlyEarnings(x) \lor Rich(x) \end{array}$

Example: resolution derivation



Resolution derivation: definite clauses



Completeness of resolution*

(Refutation) Completeness of resolution: If S is an unsatisfiable set of clauses, then the application of a finite number of resolution steps to S will yield a contradiction

Proof sketch

- If S is unsatisfiable, then there exists a particular set of ground instances of the clauses of S such that this set is also unsatisfiable (Herbrand's theorem)

- The **ground resolution theorem** is hold since propositional resolution is complete for ground sentences

- For any propositional resolution proof using the set of ground sentences, there is a corresponding first-order resolution proof using the first-order sentences from which the ground sentences were obtained (lifting lemma) In full FOL, we have the possibility of deriving $\exists x P(x)$ without being able to derive P(t) for any t

Solution: answer-extraction process

– replace query $\exists x P(x)$ by $\exists x (P(x) \land \neg A(x))$

where A is a new predicate symbol, called the answer predicate

– instead of deriving $\{\,\},$ derive any clause containing just the answer predicate

– can always convert to and from a derivation of $\{ \}$

 $\begin{array}{l} \text{E.g.,} \\ KB = \{Student(john), Student(jane), Happy(john)\} \\ Q = \exists x(Student(x) \land Happy(x)) \\ A(john), \text{ i.e., an answer is } john \end{array}$

Hardness of resolution*

First-order resolution is not guaranteed to terminate

Propositional resolution is (determining if a set of clauses is satisfiable) NP-complete (Cook Theorem)

There are unsatisfiable clauses $\{c_1, c_2, \dots, c_n\}$ s.t. the shortest derivation of $\{\}$ contains on the order of 2^n clauses (Haken, 1985)

Implications

- full theorem-proving may be too difficult
- need to consider other options
- - giving control to user, e.g., procedural representations
- less expressive languages

e.g., Horn clauses (such as Prolog), semantic Web, knowledge graph

Resolution strategies *

Strategies: reduce redundancy

- e.g., mathematical theorem proving, where we care about specific formulas

- automated theorem proving (ATP)

study strategies for automatically proving difficult theorems

- Unit preference
- Set of support
- Input resolution
- Subsumption
- Linear resolution, etc.

Ref. Chang C&Lee R, *Symbolic Logic and Mechanical Theorem Proving*, 2e, 1997 Two efficient algorithms for propositional theorem proving based on model checking

Backtracking

- DPLL (Davis-Putnam-Logemann-Loveland) algorithm: recursive, depth-first enumeration of possible models

Local search

– Similarly, $\rm MIN\text{-}CONFLICTS$ for CSPs, using an evaluation function that counts the number of unsatisfied clauses

DPLL: a complete backtracking algorithm

– improving TT- Entail

- Early termination: a clause is true if *any* literal is true E.g., $(A \lor B) \land (A \lor C)$ is true if A is true, regardless B, C
- Pure symbol heuristic: a pure symbol appears with the same "sign" in all clauses

E.g., $(A \lor \neg B)$, $(\neg B \lor \neg C)$, $(C \lor A)$

A (only positive appears) and B are pure, C is impure A sentence has a model \rightarrow it has a model with the pure symbols assigned so as to make their literals true

Unit clause heuristic: a unit clause with just one literal, with esp. clauses in which all literals but one are already assigned *false* E.g., if B = true, then (¬B ∨ ¬C) simplifies to ¬C
 assigning one unit clause can create another one (unit propagation)

<pre>def DPLL-SATISFIABLE?(s) inputs: s, a sentence in propositional logic</pre>	
$clauses \leftarrow$ the set of clauses in the CNF representation of s $symbols \leftarrow$ a list of the proposition symbols in s return DPLL(<i>clauses</i> , <i>symbols</i> , [])	
def DPLL(clauses, symbols, model)	
if every clause in <i>clauses</i> is true in <i>model</i> then return <i>true</i>	
if some clause in <i>clauses</i> is false in <i>model</i> then return <i>false</i>	
P , $value \leftarrow \text{FIND-PURE-SYMBOL}(symbols, clauses, model)$	
if P is non-null then return DPLL(clauses, symbols – P, model \cup {P =	$value\})$
P , $value \leftarrow \text{FIND-UNIT-CLAUSE}(clauses, model)$	
if P is non-null then return DPLL(<i>clauses, symbols</i> – P, $model \cup \{P =$	$value\})$
$P \leftarrow \text{FIRST}(symbols); rest \leftarrow \text{REST}(symbols)$	
return DPLL(<i>clauses, rest, model</i> \cup { $P = value$ }) or	
$DPLL(clauses, rest, model \cup \{P = value\})$	
	1

Logic programming^{*}

Computation as inference on logical KBs

	Logic programming	Ordinary programming
1.	Identify problem	Identify problem
2.	Assemble information	Assemble information
3.	Tea break	Figure out solution
4.	Encode information in KB	Program solution
5.	Encode problem instance as facts	Encode problem instance as data
6.	Ask queries	Apply program to data
7.	Find false facts	Debug procedural errors

Should be easier to debug Capital(NewYork, US) than x := x + 2

Basis: backward chaining with Horn clauses + bells & whistles Widely used in Europe, Japan (basis of 5th Generation prlinect) Compilation techniques \Rightarrow approaching a billion LIPS

```
Program = set of clauses = head :- literal<sub>1</sub>, ... literal<sub>n</sub>.
```

```
criminal(X) :- american(X), weapon(Y), sells(X,Y,Z), hostile(Z).
```

```
Efficient unification by open coding
Efficient retrieval of matching clauses by direct linking
Depth-first, left-to-right backward chaining
Built-in predicates for arithmetic etc., e.g., X is Y*Z+3
Closed-world assumption ("negation as failure")
    e.g., given alive(X) :- not dead(X).
    alive(joe) succeeds if dead(joe) fails
```

Example: Prolog program*

Depth-first search from a start state X

```
dfs(X) :- goal(X).
dfs(X) :- successor(X,S),dfs(S).
```

No need to loop over S: successor succeeds for each

Appending two lists to produce a third

```
append([],Y,Y).
append([X|L],Y,[X|Z]) :- append(L,Y,Z).
```

Answer set programming (ASP)*

Rule

 $a \leftarrow b_1, \ldots, b_m$, not c_1, \ldots , not c_n - a (head), b_i and c_j (body) are atoms

- true, if all literals to the body are true: a non-negated literal b_i is true if it has a derivation, a negated one, not c_j , is true if the atom c_j does not have one

Programs: finite collections of rules

Prolog: programming language

Need to understand Prolog's evaluation strategy, SLD resolution with unification

- the order of rules in a Prolog program and of subgoals (literals) in rule bodies matters

- Prolog misses true declarativity

ASP: specifications (yet do not allow the programmer to control the search)

– more declarative: it is intuitive, requires less background in logic, and its semantics is robust to changes in the order of literals in rules and rules in programs

- the ground program is fixed and only the data component changes

Automated theorem provers

Stanford Resolution Prover/FOL: one of the most mature subfields of ATP $% \mathcal{A}$

E-prover (E 2.3, github.com/eprover): one of the SOTA FOL /w equality prover

TPTP (Thousands of Problems for Theorem Provers) problem library

CADE ATP System Competition (CASC): a yearly competition of first-order systems

Proof assistant (interactive theorem prover): a software tool to assist with the development of formal proofs by human-machine collaboration

- LEAN, Coq, HOL, Isabelle, etc.

LEAN*

Input: a formal language for expressing math statements (definitions, axioms, conjectures, theorems, and constructions) in a human-readable and machine-verifiable format

Proof assistant: LEAN serves as a proof assistant, allowing users to interactively develop and verify math proofs (correctness and consistency)

Automated reasoning: the resolution-based automated reasoning engine is used to automate the process of proof

Proof checking: The resolution-based proofs are checked for correctness and consistency

Output: Upon successful verification, LEAN provides formalized math theorems and constructions, along with their proofs, in a machine-verifiable format