

Logical Inference

Chapter 9

Some material adopted from notes
by Andreas Geyer-Schulz,
Chuck Dyer, and Mary Getoor

Overview

- Model checking
- Inference in first order logic
 - Inference rules and generalized modes ponens
 - Forward chaining
 - Backward chaining
 - Resolution
 - Clausal form
 - Unification
 - Resolution as search

Model checking

- Given KB, does sentence S hold?
- Basically generate and test:
 - Generate all the possible models
 - Consider the models M in which KB is TRUE
 - If $\forall M S$, then S is **provably true**
 - If $\forall M \neg S$, then S is **provably false**
 - Otherwise ($\exists M1 S \wedge \exists M2 \neg S$): S is **satisfiable** but neither provably true or provably false

Efficient model checking

- Davis-Putnam algorithm (DPLL) is a Generate-and-test model checking with:
 - Early termination (short-circuiting of disjunction and conjunction)
 - Pure symbol heuristic: Any symbol that only appears negated or unnegated must be FALSE/TRUE respectively. (Can “conditionalize” based on instantiations already produced)
 - Unit clause heuristic: Any symbol that appears in a clause by itself can immediately be set to TRUE or FALSE
- WALKSAT: Local search for satisfiability: Pick a symbol to flip (toggle TRUE/FALSE), either using min-conflicts *or* choosing randomly
- ...or you can use *any* local or global search algorithm!

Reminder: Inference rules for FOL

- Inference rules for propositional logic apply to FOL as well
 - Modus Ponens, And-Introduction, And-Elimination, ...
- New (sound) inference rules for use with quantifiers:
 - Universal elimination
 - Existential introduction
 - Existential elimination
 - Generalized Modus Ponens (GMP)

5

Automating FOL inference with Generalized Modus Ponens

6

Automated inference for FOL

- Automated inference using FOL is harder than PL
 - Variables can potentially take on an *infinite* number of possible values from their domains
 - Hence there are potentially an *infinite* number of ways to apply the Universal Elimination rule of inference
- *Godel's Completeness Theorem* says that FOL entailment is only *semidecidable*
 - If a sentence is **true** given a set of axioms, there is a procedure that will determine this
 - If the sentence is **false**, then there is no guarantee that a procedure will ever determine this—i.e., it **may never halt**

7

Generalized Modus Ponens (GMP)

- Apply *modus ponens* reasoning to generalized rules
- Combines And-Introduction, Universal-Elimination, and Modus Ponens
 - From $P(c)$ and $Q(c)$ and $(\forall x)(P(x) \wedge Q(x)) \rightarrow R(x)$ derive $R(c)$
- General case: **Given**
 - **atomic sentences** P_1, P_2, \dots, P_N
 - **implication sentence** $(Q_1 \wedge Q_2 \wedge \dots \wedge Q_N) \rightarrow R$
 - Q_1, \dots, Q_N and R are atomic sentences
 - **substitution** $\text{subst}(\theta, P_i) = \text{subst}(\theta, Q_i)$ for $i=1, \dots, N$
 - **Derive new sentence: subst(θ, R)**
- Substitutions
 - $\text{subst}(\theta, \alpha)$ denotes the result of applying a set of substitutions defined by θ to the sentence α
 - A substitution list $\theta = \{v_1/t_1, v_2/t_2, \dots, v_n/t_n\}$ means to replace all occurrences of variable symbol v_i by term t_i
 - Substitutions are made in left-to-right order in the list
 - $\text{subst}(\{x/\text{IceCream}, y/\text{Ziggy}\}, \text{eats}(y,x)) = \text{eats}(\text{Ziggy}, \text{IceCream})$

8

Horn clauses

- A Horn clause is a sentence of the form:
 $(\forall x) P_1(x) \wedge P_2(x) \wedge \dots \wedge P_n(x) \rightarrow Q(x)$
 where
 - ≥ 0 P_i s and 0 or 1 Q
 - the P_i s and Q are positive (i.e., non-negated) literals
- Equivalently: $P_1(x) \vee P_2(x) \dots \vee P_n(x)$ where the P_i are all atomic and *at most one* of them is positive
- Prolog is based on Horn clauses
- Horn clauses represent a *subset* of the set of sentences representable in FOL

9

Horn clauses II

- Special cases
 - *Typical rule*: $P_1 \wedge P_2 \wedge \dots \wedge P_n \rightarrow Q$
 - *Constraint*: $P_1 \wedge P_2 \wedge \dots \wedge P_n \rightarrow \text{false}$
 - *A fact*: $\text{true} \rightarrow Q$
- These are not Horn clauses:
 - $\neg p(a) \vee q(a)$
 - $(P \wedge Q) \rightarrow (R \vee S)$

10

Forward chaining

- Proofs start with the given axioms/premises in KB, deriving new sentences using GMP until the goal/query sentence is derived
- This defines a **forward-chaining** inference procedure because it moves “forward” from the KB to the goal [eventually]
- Inference using GMP is **sound** and **complete** for KBs containing **only Horn clauses**

11

Forward chaining algorithm

```

procedure FORWARD-CHAIN(KB, p)
  if there is a sentence in KB that is a renaming of p then return
  Add p to KB
  for each  $p_1 \wedge \dots \wedge p_n \Rightarrow q$  in KB such that for some  $\theta$ ,  $\text{UNIFY}(p_1, p_i) = \theta$  succeeds do
    FIND-AND-INFER(KB,  $\theta$ ,  $p_1, \dots, p_n, q, \theta$ )
  end

procedure FIND-AND-INFER(KB, premises, conclusion,  $\theta$ )
  if premises = {} then
    FORWARD-CHAIN(KB, SUBST( $\theta$ , conclusion))
  else for each  $p$  in KB such that  $\text{UNIFY}(p, \text{FIRST}(\text{premises})) = \theta_1$  do
    FIND-AND-INFER(KB, REST(premises), conclusion, COMPOSE( $\theta, \theta_1$ ))
  end
  
```

12

Forward chaining example

- KB:
 - allergies(X) \rightarrow sneeze(X)
 - cat(Y) \wedge allergicToCats(X) \rightarrow allergies(X)
 - cat(felix)
 - allergicToCats(mary)
- Goal:
 - sneeze(mary)

13

Backward chaining

- **Backward-chaining** deduction using GMP is also **complete** for KBs containing **only Horn clauses**
- Proofs start with the goal query, find rules with that conclusion, and then prove each of the antecedents in the implication
- Keep going until you reach premises
- Avoid loops: check if new subgoal is already on the goal stack
- Avoid repeated work: check if new subgoal
 - Has already been proved true
 - Has already failed

14

Backward chaining algorithm

```

function BACK-CHAIN(KB, q) returns a set of substitutions
    BACK-CHAIN-LIST(KB, [q], [])

function BACK-CHAIN-LIST(KB, qlist,  $\theta$ ) returns a set of substitutions
    inputs: KB - a knowledge base
           qlist - a list of conjuncts forming a query ( $\theta$  already applied)
            $\theta$  - the current substitution
    static: answers, a set of substitutions, initially empty
    if qlist is empty then return  $\theta$ 
    q  $\leftarrow$  FIRST(qlist)
    for each  $r$  in KB such that  $\theta \wedge$  UNIFY(q,  $r$ ) succeeds do
        Add CONJUNCTS( $\theta, r$ ) to answers
    end
    for each conjunct  $p_1 \wedge \dots \wedge p_n$  in  $r$  such that  $\theta \wedge$  UNIFY(q,  $r$ ) succeeds do
        answers  $\leftarrow$  BACK-CHAIN-LIST(KB, SUBST( $\theta, [p_1, \dots, p_n]$ ), COMPOSE( $\theta, r$ ))
    end
    return the union of BACK-CHAIN-LIST(KB, RES, [qlist,  $\theta$ ]) for each  $\theta \in$  answers
    
```

15

Backward chaining example

- KB:
 - allergies(X) \rightarrow sneeze(X)
 - cat(Y) \wedge allergicToCats(X) \rightarrow allergies(X)
 - cat(felix)
 - allergicToCats(mary)
- Goal:
 - sneeze(mary)

16

Forward vs. backward chaining

- FC is data driven
 - Automatic, unconscious processing
 - E.g., object recognition, routine decisions
 - May do lots of work that is irrelevant to the goal
 - Efficient when you want to compute all conclusions
- BC is goal driven, better for problem solving
 - Where are my keys? How do I get to my next class?
 - Complexity of BC can be much less than linear in the size of the KB
 - Efficient when you want one or a few decisions

17

Mixed strategy

- Many practical reasoning systems do both forward and backward chaining
- The way you encode a rule determines how it is used, as in
 - % this is a forward chaining rule
 - spouse(X,Y) => spouse(Y,X).
 - % this is a backward chaining rule
 - wife(X,Y) <= spouse(X,Y), female(X).
- Given a model of the rules you have and the kind of reason you need to do, it's possible to decide which to encode as FC and which as BC rules.

18

Completeness of GMP

- GMP (using forward or backward chaining) is complete for KBs that contain only Horn clauses
- It is **not complete** for simple KBs that contain **non-Horn clauses**
- The following entail that S(A) is true:
 1. $(\forall x) P(x) \rightarrow Q(x)$
 2. $(\forall x) \neg P(x) \rightarrow R(x)$
 3. $(\forall x) Q(x) \rightarrow S(x)$
 4. $(\forall x) R(x) \rightarrow S(x)$
- If we want to conclude S(A), with GMP we cannot, since the second one is not a Horn clause
- It is equivalent to $P(x) \vee R(x)$

19

How about in Prolog?

- Let's try encoding this in Prolog
 1. $q(X) :- p(X).$
 2. $r(X) :- \text{neg}(p(X)).$
 3. $s(X) :- q(X).$
 4. $s(X) :- r(X).$
- We should not use `\+` or `not` (in SWI) for negation since it means “negation as failure”
- Prolog explores possible proofs independently
- It can't take a larger view and realize that one branch must be true, since $p(x) \vee \sim p(x)$ is always true

20

Automating FOL Inference with Resolution

21

Resolution

- Resolution is a **sound** and **complete** inference procedure for FOL
- Reminder: Resolution rule for propositional logic:
 - $P_1 \vee P_2 \vee \dots \vee P_n$
 - $\neg P_1 \vee Q_2 \vee \dots \vee Q_m$
 - Resolvent: $P_2 \vee \dots \vee P_n \vee Q_2 \vee \dots \vee Q_m$

22

Resolution covers many cases

- Modes Ponens
 - from P and $P \rightarrow Q$ derive Q
 - from P and $\neg P \vee Q$ derive Q
- Chaining
 - from $P \rightarrow Q$ and $Q \rightarrow R$ derive $P \rightarrow R$
 - from $(\neg P \vee Q)$ and $(\neg Q \vee R)$ derive $\neg P \vee R$
- Contradiction detection
 - from P and $\neg P$ derive false
 - from P and $\neg P$ derive the empty clause (=false)


23

Resolution in first-order logic

- Given sentences in *conjunctive normal form*:
 - $P_1 \vee \dots \vee P_n$ and $Q_1 \vee \dots \vee Q_m$
 - P_i and Q_i are literals, i.e., positive or negated predicate symbol with its terms
- if P_j and $\neg Q_k$ **unify** with substitution list θ , then derive the resolvent sentence:
 $\text{subst}(\theta, P_1 \vee \dots \vee P_{j-1} \vee P_{j+1} \dots P_n \vee Q_1 \vee \dots \vee Q_{k-1} \vee Q_{k+1} \vee \dots \vee Q_m)$
- Example
 - from clause $P(x, f(a)) \vee P(x, f(y)) \vee Q(y)$
 - and clause $\neg P(z, f(a)) \vee \neg Q(z)$
 - derive resolvent $P(z, f(y)) \vee Q(y) \vee \neg Q(z)$
 - using $\theta = \{x/z\}$

24

A resolution proof tree



25

Resolution refutation

- Given a consistent set of axioms KB and goal sentence Q, show that $KB \models Q$
- Proof by contradiction:** Add $\neg Q$ to KB and try to prove false.
 - i.e., $(KB \vdash Q) \leftrightarrow (KB \wedge \neg Q \vdash \text{False})$
- Resolution is **refutation complete**: it can establish that a given sentence Q is entailed by KB, but can't (in general) be used to generate all logical consequences of a set of sentences
- Also, it cannot be used to prove that Q is **not entailed** by KB.
- Resolution **won't always give an answer** since entailment is only semidecidable
 - And you can't just run two proofs in parallel, one trying to prove Q and the other trying to prove $\neg Q$, since KB might not entail either one

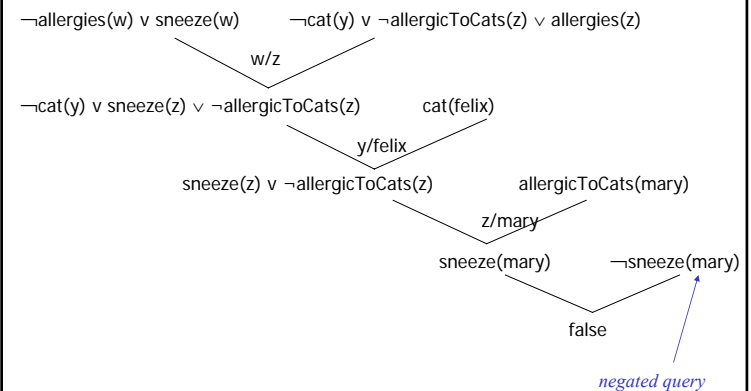
26

Resolution example

- KB:
 - $\text{allergies}(X) \rightarrow \text{sneeze}(X)$
 - $\text{cat}(Y) \wedge \text{allergicToCats}(X) \rightarrow \text{allergies}(X)$
 - $\text{cat}(\text{felix})$
 - $\text{allergicToCats}(\text{mary})$
- Goal:
 - $\text{sneeze}(\text{mary})$

27

Refutation resolution proof tree



28

questions to be answered

- How to convert FOL sentences to conjunctive normal form (a.k.a. CNF, clause form): **normalization and skolemization**
- How to unify two argument lists, i.e., how to find their most general unifier (**mgu**) q : **unification**
- How to determine which two clauses in KB should be resolved next (among all resolvable pairs of clauses) : **resolution (search) strategy**

29

Converting to CNF

30

Converting sentences to CNF

1. Eliminate all \leftrightarrow connectives
 $(P \leftrightarrow Q) \Rightarrow ((P \rightarrow Q) \wedge (Q \rightarrow P))$
2. Eliminate all \rightarrow connectives
 $(P \rightarrow Q) \Rightarrow (\neg P \vee Q)$
3. Reduce the scope of each negation symbol to a single predicate
 $\neg\neg P \Rightarrow P$
 $\neg(P \vee Q) \Rightarrow \neg P \wedge \neg Q$
 $\neg(P \wedge Q) \Rightarrow \neg P \vee \neg Q$
 $\neg(\forall x)P \Rightarrow (\exists x)\neg P$
 $\neg(\exists x)P \Rightarrow (\forall x)\neg P$
4. Standardize variables: rename all variables so that each quantifier has its own unique variable name

31

Converting sentences to clausal form Skolem constants and functions

5. Eliminate existential quantification by introducing Skolem constants/functions
 $(\exists x)P(x) \Rightarrow P(C)$
C is a Skolem constant (a brand-new constant symbol that is not used in any other sentence)
 $(\forall x)(\exists y)P(x,y) \Rightarrow (\forall x)P(x, f(x))$
since \exists is within the scope of a universally quantified variable, use a **Skolem function f** to construct a new value that **depends on** the universally quantified variable
f must be a brand-new function name not occurring in any other sentence in the KB.
E.g., $(\forall x)(\exists y)\text{loves}(x,y) \Rightarrow (\forall x)\text{loves}(x,f(x))$
In this case, $f(x)$ specifies the person that x loves
a better name might be **oneWhoIsLovedBy(x)**

32

Converting sentences to clausal form

6. Remove universal quantifiers by (1) moving them all to the left end; (2) making the scope of each the entire sentence; and (3) dropping the “prefix” part

Ex: $(\forall x)P(x) \Rightarrow P(x)$

7. Put into conjunctive normal form (conjunction of disjunctions) using distributive and associative laws

$(P \wedge Q) \vee R \Rightarrow (P \vee R) \wedge (Q \vee R)$

$(P \vee Q) \vee R \Rightarrow (P \vee Q \vee R)$

8. Split conjuncts into separate clauses
9. Standardize variables so each clause contains only variable names that do not occur in any other clause

33

An example

$(\forall x)(P(x) \rightarrow ((\forall y)(P(y) \rightarrow P(f(x,y))) \wedge \neg(\forall y)(Q(x,y) \rightarrow P(y))))$

2. Eliminate \rightarrow

$(\forall x)(\neg P(x) \vee ((\forall y)(\neg P(y) \vee P(f(x,y))) \wedge \neg(\forall y)(\neg Q(x,y) \vee P(y))))$

3. Reduce scope of negation

$(\forall x)(\neg P(x) \vee ((\forall y)(\neg P(y) \vee P(f(x,y))) \wedge (\exists y)(Q(x,y) \wedge \neg P(y))))$

4. Standardize variables

$(\forall x)(\neg P(x) \vee ((\forall y)(\neg P(y) \vee P(f(x,y))) \wedge (\exists z)(Q(x,z) \wedge \neg P(z))))$

5. Eliminate existential quantification

$(\forall x)(\neg P(x) \vee ((\forall y)(\neg P(y) \vee P(f(x,y))) \wedge (Q(x,g(x)) \wedge \neg P(g(x))))$

6. Drop universal quantification symbols

$(\neg P(x) \vee ((\neg P(y) \vee P(f(x,y))) \wedge (Q(x,g(x)) \wedge \neg P(g(x))))$

34

Example

7. Convert to conjunction of disjunctions

$(\neg P(x) \vee \neg P(y) \vee P(f(x,y))) \wedge (\neg P(x) \vee Q(x,g(x))) \wedge$
 $(\neg P(x) \vee \neg P(g(x)))$

8. Create separate clauses

$\neg P(x) \vee \neg P(y) \vee P(f(x,y))$

$\neg P(x) \vee Q(x,g(x))$

$\neg P(x) \vee \neg P(g(x))$

9. Standardize variables

$\neg P(x) \vee \neg P(y) \vee P(f(x,y))$

$\neg P(z) \vee Q(z,g(z))$

$\neg P(w) \vee \neg P(g(w))$

35

Unification

36

Unification

- Unification is a “**pattern-matching**” procedure
 - Takes two atomic sentences, called literals, as input
 - Returns “Failure” if they do not match and a substitution list, θ , if they do
- That is, $unify(p, q) = \theta$ means $subst(\theta, p) = subst(\theta, q)$ for two atomic sentences, p and q
- θ is called the **most general unifier** (mgu)
- All variables in the given two literals are implicitly universally quantified
- To make literals match, replace (universally quantified) variables by terms

37

Unification algorithm

```
procedure unify(p, q,  $\theta$ )
  Scan p and q left-to-right and find the first corresponding
  terms where p and q “disagree” (i.e., p and q not equal)
  If there is no disagreement, return  $\theta$  (success!)
  Let r and s be the terms in p and q, respectively,
  where disagreement first occurs
  If variable(r) then {
    Let  $\theta = \text{union}(\theta, \{r/s\})$ 
    Return unify(subst( $\theta$ , p), subst( $\theta$ , q),  $\theta$ )
  } else if variable(s) then {
    Let  $\theta = \text{union}(\theta, \{s/r\})$ 
    Return unify(subst( $\theta$ , p), subst( $\theta$ , q),  $\theta$ )
  } else return “Failure”
end
```

38

Unification: Remarks

- *Unify* is a linear-time algorithm that returns the most general unifier (mgu), i.e., the shortest-length substitution list that makes the two literals match.
- In general, there is not a **unique** minimum-length substitution list, but unify returns one of minimum length
- A variable can never be replaced by a term containing that variable
Example: $x/f(x)$ is illegal.
- This “occurs check” should be done in the above pseudo-code before making the recursive calls

39

Unification examples

- Example:
 - $\text{parents}(x, \text{father}(x), \text{mother}(\text{Bill}))$
 - $\text{parents}(\text{Bill}, \text{father}(\text{Bill}), y)$
 - $\{x/\text{Bill}, y/\text{mother}(\text{Bill})\}$
- Example:
 - $\text{parents}(x, \text{father}(x), \text{mother}(\text{Bill}))$
 - $\text{parents}(\text{Bill}, \text{father}(y), z)$
 - $\{x/\text{Bill}, y/\text{Bill}, z/\text{mother}(\text{Bill})\}$
- Example:
 - $\text{parents}(x, \text{father}(x), \text{mother}(\text{Jane}))$
 - $\text{parents}(\text{Bill}, \text{father}(y), \text{mother}(y))$
 - Failure

40

Resolution example

41

Practice example

Did Curiosity kill the cat

- Jack owns a dog. Every dog owner is an animal lover. No animal lover kills an animal. Either Jack or Curiosity killed the cat, who is named Tuna. Did Curiosity kill the cat?
- These can be represented as follows:
 - A. $(\exists x) \text{Dog}(x) \wedge \text{Owns}(\text{Jack}, x)$
 - B. $(\forall x) ((\exists y) \text{Dog}(y) \wedge \text{Owns}(x, y)) \rightarrow \text{AnimalLover}(x)$
 - C. $(\forall x) \text{AnimalLover}(x) \rightarrow ((\forall y) \text{Animal}(y) \rightarrow \neg \text{Kills}(x, y))$
 - D. $\text{Kills}(\text{Jack}, \text{Tuna}) \vee \text{Kills}(\text{Curiosity}, \text{Tuna})$
 - E. $\text{Cat}(\text{Tuna})$
 - F. $(\forall x) \text{Cat}(x) \rightarrow \text{Animal}(x)$
 - G. $\text{Kills}(\text{Curiosity}, \text{Tuna})$ ← GOAL

42

• Convert to clause form

A1. $(\text{Dog}(D))$ ← *D is a skolem constant*

A2. $(\text{Owns}(\text{Jack}, D))$

B. $(\neg \text{Dog}(y), \neg \text{Owns}(x, y), \text{AnimalLover}(x))$

C. $(\neg \text{AnimalLover}(a), \neg \text{Animal}(b), \neg \text{Kills}(a, b))$

D. $(\text{Kills}(\text{Jack}, \text{Tuna}), \text{Kills}(\text{Curiosity}, \text{Tuna}))$

E. $\text{Cat}(\text{Tuna})$

F. $(\neg \text{Cat}(z), \text{Animal}(z))$

• Add the negation of query:

$\neg G: \neg \text{Kills}(\text{Curiosity}, \text{Tuna})$


43

The resolution refutation proof

- | | |
|---|---|
| R1: $\neg G, D, \{\}$ | $(\text{Kills}(\text{Jack}, \text{Tuna}))$ |
| R2: R1, C, $\{a/\text{Jack}, b/\text{Tuna}\}$ | $(\sim \text{AnimalLover}(\text{Jack}), \sim \text{Animal}(\text{Tuna}))$ |
| R3: R2, B, $\{x/\text{Jack}\}$ | $(\sim \text{Dog}(y), \sim \text{Owns}(\text{Jack}, y), \sim \text{Animal}(\text{Tuna}))$ |
| R4: R3, A1, $\{y/D\}$ | $(\sim \text{Owns}(\text{Jack}, D), \sim \text{Animal}(\text{Tuna}))$ |
| R5: R4, A2, $\{\}$ | $(\sim \text{Animal}(\text{Tuna}))$ |
| R6: R5, F, $\{z/\text{Tuna}\}$ | $(\sim \text{Cat}(\text{Tuna}))$ |
| R7: R6, E, $\{\}$ | FALSE |

44

• **The proof tree**



45

Resolution search strategies

46

Resolution TP as search

- Resolution can be thought of as the **bottom-up construction of a search tree**, where the leaves are the clauses produced by KB and the negation of the goal
- When a pair of clauses generates a new resolvent clause, add a new node to the tree with arcs directed from the resolvent to the two parent clauses
- **Resolution succeeds** when a node containing the **False** clause is produced, becoming the **root node** of the tree
- A strategy is **complete** if its use guarantees that the empty clause (i.e., false) can be derived whenever it is entailed

47

Strategies

- There are a number of general (domain-independent) strategies that are useful in controlling a resolution theorem prover
- We'll briefly look at the following:
 - Breadth-first
 - Length heuristics
 - Set of support
 - Input resolution
 - Subsumption
 - Ordered resolution

48

Example

1. \neg Battery-OK \vee \neg Bulbs-OK \vee Headlights-Work
 2. \neg Battery-OK \vee \neg Starter-OK \vee Empty-Gas-Tank \vee Engine-Starts
 3. \neg Engine-Starts \vee Flat-Tire \vee Car-OK
 4. Headlights-Work
 5. Battery-OK
 6. Starter-OK
 7. \neg Empty-Gas-Tank
 8. \neg Car-OK
 9. \neg Flat-Tire
- ← negated goal

49

Breadth-first search

- Level 0 clauses are the original axioms and the negation of the goal
- Level k clauses are the resolvents computed from two clauses, one of which must be from level k-1 and the other from an earlier level
- Compute all possible level 1 clauses, then all possible level 2 clauses, etc.
- Complete, but very inefficient

50

BFS example

1. \neg Battery-OK \vee \neg Bulbs-OK \vee Headlights-Work
2. \neg Battery-OK \vee \neg Starter-OK \vee Empty-Gas-Tank \vee Engine-Starts
3. \neg Engine-Starts \vee Flat-Tire \vee Car-OK
4. Headlights-Work
5. Battery-OK
6. Starter-OK
7. \neg Empty-Gas-Tank
8. \neg Car-OK
9. \neg Flat-Tire
- 1,4 10. \neg Battery-OK \vee \neg Bulbs-OK
- 1,5 11. \neg Bulbs-OK \vee Headlights-Work
- 2,3 12. \neg Battery-OK \vee \neg Starter-OK \vee Empty-Gas-Tank \vee Flat-Tire \vee Car-OK
- 2,5 13. \neg Starter-OK \vee Empty-Gas-Tank \vee Engine-Starts
- 2,6 14. \neg Battery-OK \vee Empty-Gas-Tank \vee Engine-Starts
- 2,7 15. \neg Battery-OK \neg Starter-OK \vee Engine-Starts
16. ... [and we're still only at Level 1!]

51

Length heuristics

- **Shortest-clause heuristic:**
Generate a clause with the fewest literals first
- **Unit resolution:**
Prefer resolution steps in which at least one parent clause is a “unit clause,” i.e., a clause containing a single literal
 - Not complete in general, but complete for Horn clause KBs

52

Unit resolution example

1. \neg Battery-OK \vee \neg Bulbs-OK \vee Headlights-Work
2. \neg Battery-OK \vee \neg Starter-OK \vee Empty-Gas-Tank \vee Engine-Starts
3. \neg Engine-Starts \vee Flat-Tire \vee Car-OK
4. Headlights-Work
5. Battery-OK
6. Starter-OK
7. \neg Empty-Gas-Tank
8. \neg Car-OK
9. \neg Flat-Tire
- 1,5 10. \neg Bulbs-OK \vee Headlights-Work
- 2,5 11. \neg Starter-OK \vee Empty-Gas-Tank \vee Engine-Starts
- 2,6 12. \neg Battery-OK \vee Empty-Gas-Tank \vee Engine-Starts
- 2,7 13. \neg Battery-OK \neg Starter-OK \vee Engine-Starts
- 3,8 14. \neg Engine-Starts \vee Flat-Tire
- 3,9 15. \neg Engine-Starts \neg Car-OK
16. ... [this doesn't seem to be headed anywhere either!]

53

Set of support

- At least one parent clause must be the negation of the goal *or* a “descendant” of such a goal clause (i.e., derived from a goal clause)
- (*When there's a choice, take the most recent descendant*)
- Complete (assuming all possible set-of-support clauses are derived)
- Gives a goal-directed character to the search

54

Set of support example

1. \neg Battery-OK \vee \neg Bulbs-OK \vee Headlights-Work
2. \neg Battery-OK \vee \neg Starter-OK \vee Empty-Gas-Tank \vee Engine-Starts
3. \neg Engine-Starts \vee Flat-Tire \vee Car-OK
4. Headlights-Work
5. Battery-OK
6. Starter-OK
7. \neg Empty-Gas-Tank
8. \neg Car-OK
9. \neg Flat-Tire
- 9,3 10. \neg Engine-Starts \vee Car-OK
- 10,2 11. \neg Battery-OK \vee \neg Starter-OK \vee Empty-Gas-Tank \vee Car-OK
- 10,8 12. \neg Engine-Starts
- 11,5 13. \neg Starter-OK \vee Empty-Gas-Tank \vee Car-OK
- 11,6 14. \neg Battery-OK \vee Empty-Gas-Tank \vee Car-OK
- 11,7 15. \neg Battery-OK \vee \neg Starter-OK \vee Car-OK
16. ... [a bit more focused, but we still seem to be wandering]

55

Unit resolution + set of support example

1. \neg Battery-OK \vee \neg Bulbs-OK \vee Headlights-Work
 2. \neg Battery-OK \vee \neg Starter-OK \vee Empty-Gas-Tank \vee Engine-Starts
 3. \neg Engine-Starts \vee Flat-Tire \vee Car-OK
 4. Headlights-Work
 5. Battery-OK
 6. Starter-OK
 7. \neg Empty-Gas-Tank
 8. \neg Car-OK
 9. \neg Flat-Tire
 - 9,3 10. \neg Engine-Starts \vee Car-OK
 - 10,8 11. \neg Engine-Starts
 - 12,2 12. \neg Battery-OK \vee \neg Starter-OK \vee Empty-Gas-Tank
 - 12,5 13. \neg Starter-OK \vee Empty-Gas-Tank
 - 13,6 14. Empty-Gas-Tank
 - 14,7 15. FALSE
- [Hooray! Now that's more like it!]

56

Simplification heuristics

- **Subsumption:**
Eliminate all sentences that are subsumed by (more specific than) an existing sentence to keep the KB small
 - If $P(x)$ is already in the KB, adding $P(A)$ makes no sense – $P(x)$ is a superset of $P(A)$
 - Likewise adding $P(A) \vee Q(B)$ would add nothing to the KB
- **Tautology:**
Remove any clause containing two complementary literals (tautology)
- **Pure symbol:**
If a symbol always appears with the same “sign,” remove all the clauses that contain it

57

Example (Pure Symbol)

1. ~~Battery-OK \vee Bulbs-OK \vee Headlights-Work~~
2. \neg Battery-OK \vee \neg Starter-OK \vee Empty-Gas-Tank \vee Engine-Starts
3. \neg Engine-Starts \vee Flat-Tire \vee Car-OK
4. ~~Headlights-Work~~
5. Battery-OK
6. Starter-OK
7. \neg Empty-Gas-Tank
8. \neg Car-OK
9. ~~\neg Flat-Tire~~

58

Input resolution

- At least one parent must be one of the input sentences (i.e., either a sentence in the original KB or the negation of the goal)
- Not complete in general, but complete for Horn clause KBs
- **Linear resolution**
 - Extension of input resolution
 - One of the parent sentences must be an input sentence *or* an ancestor of the other sentence
 - Complete

59

Ordered resolution

- Search for resolvable sentences in order (left to right)
- This is how Prolog operates
- Resolve the first element in the sentence first
- This forces the user to define what is important in generating the “code”
- The way the sentences are written controls the resolution

60

Prolog

- A logic programming language based on Horn clauses
 - Resolution refutation
 - Control strategy: goal-directed and depth-first
 - always start from the goal clause
 - always use the new resolvent as one of the parent clauses for resolution
 - backtracking when the current thread fails
 - complete for Horn clause KB
 - Support answer extraction (can request single or all answers)
 - Orders the clauses and literals within a clause to resolve non-determinism
 - $Q(a)$ may match both $Q(x) \Leftarrow P(x)$ and $Q(y) \Leftarrow R(y)$
 - A (sub)goal clause may contain more than one literals, i.e., $\Leftarrow P1(a), P2(a)$
 - Use “closed world” assumption (negation as failure)
 - If it fails to derive $P(a)$, then assume $\sim P(a)$

61

Summary

- Logical agents apply inference to a knowledge base to derive new information and make decisions
- Basic concepts of logic:
 - Syntax: formal structure of sentences
 - Semantics: truth of sentences wrt models
 - Entailment: necessary truth of one sentence given another
 - Inference: deriving sentences from other sentences
 - Soundness: derivations produce only entailed sentences
 - Completeness: derivations can produce all entailed sentences
- FC and BC are linear time, complete for Horn clauses
- Resolution is a sound and complete inference method for propositional and first-order logic

62