

Inference Strategy: Backward Chaining

Idea:

- Check whether a particular fact q is true.

Backward Chaining:

Given a fact q to be “proven”,

1. See if q is already in the KB. If so, return TRUE.
2. Find all implications, I , whose conclusion “matches” q .
3. Recursively establish the premises of all i in I via backward chaining.

✚ Avoids inferring unrelated facts.

Example: Backward Chaining

Knowledge-base describing when the car should brake?

```
( PersonInFrontOfCar  $\Rightarrow$  Brake )  
 $\wedge$  ((( YellowLight  $\wedge$  Policeman )  $\wedge$  (  $\neg$ Slippery ))  $\Rightarrow$  Brake )  
 $\wedge$  ( Policecar  $\Rightarrow$  Policeman )  
 $\wedge$  ( Snow  $\Rightarrow$  Slippery )  
 $\wedge$  ( Slippery  $\Rightarrow$   $\neg$ Dry )  
 $\wedge$  ( RedLight  $\Rightarrow$  Brake )  
 $\wedge$  ( Winter  $\Rightarrow$  Snow )
```

Observation from sensors:

YellowLight \wedge \neg RedLight \wedge \neg Snow \wedge Dry \wedge Policecar \wedge \neg PersonInFrontOfCar

Should the agent brake (i.e. can “brake” be inferred)?

- Goal: Brake
 - Modus Ponens (brake): PersonInFrontOfCar
 - Failure: PersonInFrontOfCar Backtracking
- Goal: Brake
 - Modus Ponens (brake): YellowLight \wedge Policeman \wedge \neg Slippery
 - Known (YellowLight): Policeman \wedge \neg Slippery
 - Modus Ponens (Policeman): Policecar \wedge \neg Slippery
 - Known (Policecar): \neg Slippery
 - Modus Tollens (\neg Slippery): Dry
 - Known (Dry)

First-Order Logic

• Idea:

- Don’t treat propositions as “atomic” entities.

• First-Order Logic:

- Objects: cs472, fred, oh155, emptylist ...
- Relations/Predicates: is_Man(fred), Located(cs472, oh155) ...
 - Note: Relations typically correspond to verbs
- Functions: Pair(search,Pair(learning,Pair(kbsystems, emptylist)))
- Connectives: \wedge , \vee , \neg , \Rightarrow , \Leftrightarrow
- Quantifiers:
 - Universal: $\forall x: (\text{is_Man}(x) \Rightarrow \text{is_Mortal}(x))$
 - Existential: $\exists y: (\text{is_Father}(y, \text{fred}))$

Example:

Representing Facts in First-Order Logic

1. Lucy* is a professor
2. All professors are people.
3. Fuchs is the dean.
4. Deans are professors.
5. All professors consider the dean a friend or don’t know him.
6. Everyone is a friend of someone.
7. People only criticize people that are not their friends.
8. Lucy criticized Fuchs.

* Name changed for privacy reasons.

Example: Proof

Knowledge base:

- is-prof(lucy)
- $\forall x (\text{is-prof}(x) \rightarrow \text{is-person}(x))$
- is-dean(fuchs)
- $\forall x (\text{is-dean}(x) \rightarrow \text{is-prof}(x))$
- $\forall x (\forall y (\text{is-prof}(x) \wedge \text{is-dean}(y) \rightarrow \text{is-friend-of}(y,x) \vee \neg \text{knows}(x,y)))$
- $\forall x (\exists y (\text{is-friend-of}(y,x)))$
- $\forall x (\forall y (\text{is-person}(x) \wedge \text{is-person}(y) \wedge \text{criticize}(x,y) \rightarrow \neg \text{is-friend-of}(y,x)))$
- criticize(lucy,fuchs)

Question: Is Fuchs no friend of Lucy?

$\neg \text{is-friend-of}(\text{fuchs}, \text{lucy})$

Knowledge Engineering

1. Identify the task.
2. Assemble the relevant knowledge.
3. Decide on a vocabulary of predicates, functions, and constants.
4. Encode general knowledge about the domain.
5. Encode a description of the specific problem instance.
6. Pose queries to the inference procedure and get answers.
7. Debug the knowledge base.

Inference Procedures: Theoretical Results

- **There exist complete and sound proof procedures for propositional and FOL.**
 - Propositional logic
 - Use the definition of entailment directly. Proof procedure is exponential in n , the number of symbols.
 - In practice, can be much faster...
 - Polynomial-time inference procedure exists when KB is expressed as **Horn clauses**: $P_1 \wedge P_2 \wedge \dots \wedge P_n \Rightarrow Q$ where the P_i and Q are non-negated atoms.
 - First-Order logic
 - Godel's completeness theorem showed that a proof procedure exists...
 - But none was demonstrated until Robinson's 1965 *resolution algorithm*.
 - Entailment in first-order logic is *semidecidable*.

Resolution Rule of Inference

General Rule:

Assume: $E \vee E_{12} \vee \dots \vee E_{1k}$
 and $\neg E \vee E_{22} \vee \dots \vee E_{2l}$

 Then: $E_{12} \vee \dots \vee E_{1k} \vee E_{22} \vee \dots \vee E_{2l}$

Note: E_{ij} can be negated.

Example:

Assume: $E_1 \vee E_2$ playing tennis or raining
 and $\neg E_2 \vee E_3$ not raining or working

 Then: $E_1 \vee E_3$ playing tennis or working

Algorithm: Resolution Proof

- **Negate the theorem to be proved, and add the result to the knowledge base.**
- **Bring knowledge base into conjunctive normal form (CNF)**
 - CNF: conjunctions of disjunctions
 - Each disjunction is called a clause.
- **Until there is no resolvable pair of clauses,**
 - Find resolvable clauses and resolve them.
 - Add the results of resolution to the knowledge base.
 - If NIL (empty clause) is produced, stop and report that the (original) theorem is true.
- **Report that the (original) theorem is false.**

Resolution Example: Propositional Logic

- **To prove: $\neg P$**
- **Transform Knowledge Base into CNF**

	Regular	CNF
Sentence 1:	$P \rightarrow Q$	$\neg P \vee Q$
Sentence 2:	$Q \rightarrow R$	$\neg Q \vee R$
Sentence 3:	$\neg R$	$\neg R$

Proof

1. $\neg P \vee Q$	Sentence 1
2. $\neg Q \vee R$	Sentence 2
3. $\neg R$	Sentence 3
4. P	Assume opposite
5. Q	Resolve 4 and 1
6. R	Resolve 5 and 2
7. nil	Resolve 6 with 3

Resolution Example: FOL

Example: Prove bird (tweety)

Axioms: Regular

1: $\forall x: feathers(x) \rightarrow bird(x)$
 2: $feathers(tweety)$
 3: $\neg bird(tweety)$

CNF

$\neg feathers(x) \vee bird(x)$
 $feathers(tweety)$
 $\neg bird(tweety)$
 $\neg feathers(tweety)$

4: Resolution Proof

1. Resolve 3 and 1, specializing (i.e. "unifying") tweety for x. Add $\neg feathers(tweety)$
2. Resolve 4 and 2. Add NIL.

Resolution Theorem Proving

Properties of Resolution Theorem Proving:

- sound (for propositional and FOL)
- (refutation) complete (for propositional and FOL)

Procedure may seem cumbersome but note that can be easily automated. Just "smash" clauses until empty clause or no more new clauses.

Unification

Unify procedure: Unify(P,Q) takes two atomic (i.e. single predicates) sentences P and Q and returns a substitution that makes P and Q identical.

Rules for substitutions:

- Can replace a variable by a constant.
- Can replace a variable by a variable.
- Can replace a variable by a function expression, as long as the function expression does not contain the variable.

Unifier: a substitution that makes two clauses resolvable.

$$v_1/C; v_2/v_3; v_4/f(\dots)$$

Unification - Purpose

Given:

$$\neg \text{Knows}(\text{John}, x) \vee \text{Hates}(\text{John}, x) \\ \text{Knows}(\text{John}, \text{Jim})$$

Derive:

$$\text{Hates}(\text{John}, \text{Jim})$$

Unification:

$$\text{unify}(\text{Knows}(\text{John}, x), \text{Knows}(\text{John}, \text{Jim})) = \{x/\text{Jim}\}$$

Need unifier $\{x/\text{Jim}\}$ for resolution to work.

Add to knowledge base:

$$\neg \text{Knows}(\text{John}, \text{Jim}) \vee \text{Hates}(\text{John}, \text{Jim})$$

Unification (example)

Who does John hate?

$$\exists x: \text{Hates}(\text{John}, x)$$

Knowledge base (in clause form):

1. $\neg \text{Knows}(\text{John}, v) \vee \text{Hates}(\text{John}, v)$
 2. $\text{Knows}(\text{John}, \text{Jim})$
 3. $\text{Knows}(y, \text{Leo})$
 4. $\text{Knows}(z, \text{Mother}(z))$
 5. $\neg \text{Hates}(\text{John}, x)$ (since $\neg \exists x: \text{Hates}(\text{John}, x) \quad \forall x: \neg \text{Hates}(\text{John}, x)$)
- Resolution with 5 and 1:
 $\text{unify}(\text{Hates}(\text{John}, x), \text{Hates}(\text{John}, v)) = \{x/v\}$
6. $\neg \text{Knows}(\text{John}, v)$
- Resolution with 6 and 2:
 $\text{unify}(\text{Knows}(\text{John}, v), \text{Knows}(\text{John}, \text{Jim})) = \{v/\text{Jim}\}$
 or resolution with 6 and 3:
 $\text{unify}(\text{Knows}(\text{John}, v), \text{Knows}(y, \text{Leo})) = \{y/\text{John}, v/\text{Leo}\}$
 or Resolution with 6 and 4:
 $\text{unify}(\text{Knows}(\text{John}, v), \text{Knows}(z, \text{Mother}(z))) = \{z/\text{John}, v/\text{Mother}(z)\}$

Answers:

1. $\text{Hates}(\text{John}, x)$ with $\{x/v, v/\text{Jim}\}$ (i.e. John hates Jim)
2. $\text{Hates}(\text{John}, x)$ with $\{x/v, y/\text{John}, v/\text{Leo}\}$ (i.e. John hates Leo)
3. $\text{Hates}(\text{John}, x)$ with $\{x/v, v/\text{Mother}(z), z/\text{John}\}$ (i.e. John hates his mother)

Most General Unifier

In cases where there is more than one substitution choose the one that makes the least commitment (most general) about the bindings.

$$\text{UNIFY} \quad (\text{Knows}(\text{John}, x), \text{Knows}(y, z)) \\ = \{y/\text{John}, x/z\} \\ \text{not } \{y/\text{John}, x/z, z/\text{Freda}\} \\ \text{not } \{y/\text{John}, x/\text{John}, z/\text{John}\} \\ \dots$$

See R&N for general unification algorithm. $O(n^2)$ with Refutation

Converting More Complicated Sentences to CNF

Sentence:

$$\forall x: \text{brick}(x) \rightarrow ((\exists y: \text{on}(x, y) \wedge \neg \text{pyramid}(y)) \\ \wedge (\neg \exists y: \text{on}(x, y) \wedge \text{on}(y, x)) \\ \wedge (\forall y: \neg \text{brick}(y) \rightarrow \neg \text{equal}(x, y)))$$

CNF:

$$\neg \text{brick}(x) \vee \text{on}(x, \text{support}(x)) \\ \neg \text{brick}(w) \vee \neg \text{pyramid}(\text{support}(w)) \\ \neg \text{brick}(u) \vee \neg \text{on}(u, y) \vee \neg \text{on}(y, u) \\ \neg \text{brick}(v) \vee \text{brick}(z) \vee \neg \text{equal}(v, z)$$

1. Eliminate Implications

Substitute $\neg E_1 \vee E_2$ for $E_1 \rightarrow E_2$

$$\forall x: \text{brick}(x) \rightarrow ((\exists y: \text{on}(x, y) \wedge \neg \text{pyramid}(y)) \\ \wedge (\neg \exists y: \text{on}(x, y) \wedge \text{on}(y, x)) \\ \wedge (\forall y: \neg \text{brick}(y) \rightarrow \neg \text{equal}(x, y)))$$

$$\forall x: \neg \text{brick}(x) \vee ((\exists y: \text{on}(x, y) \wedge \neg \text{pyramid}(y)) \\ \wedge (\neg \exists y: \text{on}(x, y) \wedge \text{on}(y, x)) \\ \wedge (\forall y: \neg (\neg \text{brick}(y)) \vee \neg \text{equal}(x, y)))$$

2. Move negations down to the atomic formulas

Equivalence Transformations:

$$\neg(E_1 \wedge E_2) \iff (\neg E_1) \vee (\neg E_2)$$

$$\neg(E_1 \vee E_2) \iff (\neg E_1) \wedge (\neg E_2)$$

$$\neg(\neg E_1) \iff E_1$$

$$\neg\forall x : E_1(x) \iff \exists x : \neg E_1(x)$$

$$\neg\exists x : E_1(x) \iff \forall x : \neg E_1(x)$$

Result:

$$\forall x : \neg brick(x) \vee$$

$$((\exists y : on(x, y) \wedge \neg pyramid(y))$$

$$\wedge (\neg\exists y : on(x, y) \wedge on(y, x))$$

$$\wedge (\forall y : \neg(\neg brick(y)) \vee \neg equal(x, y)))$$

3. Eliminate Existential Quantifiers: Skolemization

Harder cases:

$$\forall x : \exists y : father(y, x) \text{ becomes } \forall x : father(S1(x), x)$$

There is one argument for each universally quantified variable whose scope contains the Skolem function.

Easy case:

$$\exists x : President(x) \text{ becomes } President(S2)$$

$$\forall x : \neg brick(x) \vee ((\exists y : on(x, y) \wedge \neg pyramid(y)) \wedge \dots$$

4. Rename variables as necessary

We want no two variables of the same name.

$$\begin{aligned} \forall x : \neg brick(x) \vee & (on(x, S1(x)) \wedge \neg pyramid(S1(x))) \\ & \wedge (\forall y : (\neg on(x, y) \vee \neg on(y, x))) \\ & \wedge (\forall y : (brick(y) \vee \neg equal(x, y))) \end{aligned}$$

$$\begin{aligned} \forall x : \neg brick(x) \vee & (on(x, S1(x)) \wedge \neg pyramid(S1(x))) \\ & \wedge (\forall y : (\neg on(x, y) \vee \neg on(y, x))) \\ & \wedge (\forall z : (brick(z) \vee \neg equal(x, z))) \end{aligned}$$

5. Move the universal quantifiers to the left

This works because each quantifier uses a unique variable name.

$$\begin{aligned} \forall x : \neg brick(x) \vee & (on(x, S1(x)) \wedge \neg pyramid(S1(x))) \\ & \wedge (\forall y : (\neg on(x, y) \vee \neg on(y, x))) \\ & \wedge (\forall z : (brick(z) \vee \neg equal(x, z))) \end{aligned}$$

$$\begin{aligned} \forall x \forall y \forall z : \neg brick(x) \vee & (on(x, S1(x)) \wedge \neg pyramid(S1(x))) \\ & \wedge (\neg on(x, y) \vee \neg on(y, x)) \\ & \wedge (brick(z) \vee \neg equal(x, z)) \end{aligned}$$

6. Move disjunctions down to the literals

$$E_1 \vee (E_2 \wedge E_3) \iff (E_1 \vee E_2) \wedge (E_1 \vee E_3)$$

$$\begin{aligned} \forall x \forall y \forall z : & (\neg brick(x) \vee (on(x, S1(x)) \wedge \neg pyramid(S1(x)))) \\ & \wedge (\neg brick(x) \vee \neg on(x, y) \vee \neg on(y, x)) \\ & \wedge (\neg brick(x) \vee brick(z) \vee \neg equal(x, z)) \end{aligned}$$

$$\begin{aligned} \forall x \forall y \forall z : & (\neg brick(x) \vee on(x, S1(x))) \\ & \wedge (\neg brick(x) \vee \neg pyramid(S1(x))) \\ & \wedge (\neg brick(x) \vee \neg on(x, y) \vee \neg on(y, x)) \\ & \wedge (\neg brick(x) \vee brick(z) \vee \neg equal(x, z)) \end{aligned}$$

7. Eliminate the conjunctions

$$\begin{aligned} \forall x \forall y \forall z : & (\neg brick(x) \vee on(x, S1(x))) \\ & \wedge (\neg brick(x) \vee \neg pyramid(S1(x))) \\ & \wedge (\neg brick(x) \vee \neg on(x, y) \vee \neg on(y, x)) \\ & \wedge (\neg brick(x) \vee brick(z) \vee \neg equal(x, z)) \end{aligned}$$

$$\begin{aligned} \forall x : \neg brick(x) \vee & on(x, S1(x)) \\ \forall x : \neg brick(x) \vee & \neg pyramid(S1(x)) \\ \forall x \forall y : \neg brick(x) \vee & \neg on(x, y) \vee \neg on(y, x) \\ \forall x \forall z : \neg brick(x) \vee & brick(z) \vee \neg equal(x, z) \end{aligned}$$

8. Rename all variables, as necessary, so no two have the same name

$\forall x : \neg brick(x) \vee on(x, S1(x))$
 $\forall x : \neg brick(x) \vee \neg pyramid(S1(x))$
 $\forall x \forall y : \neg brick(x) \vee \neg on(x, y) \vee \neg on(y, x)$
 $\forall x \forall z : \neg brick(x) \vee brick(z) \vee \neg equal(x, z)$

$\forall x : \neg brick(x) \vee on(x, S1(x))$
 $\forall w : \neg brick(w) \vee \neg pyramid(S1(w))$
 $\forall u \forall y : \neg brick(u) \vee \neg on(u, y) \vee \neg on(y, u)$
 $\forall v \forall z : \neg brick(v) \vee brick(z) \vee \neg equal(v, z)$

9. Eliminate the universal quantifiers

$\neg brick(x) \vee on(x, S1(x))$
 $\neg brick(w) \vee \neg pyramid(S1(w))$
 $\neg brick(u) \vee \neg on(u, y) \vee \neg on(y, u)$
 $\neg brick(v) \vee brick(z) \vee \neg equal(v, z)$

Algorithm: Putting Axioms into Clausal Form

1. Eliminate the implications.
2. Move the negations down to the atomic formulas.
3. Eliminate the existential quantifiers.
4. Rename the variables, if necessary.
5. Move the universal quantifiers to the left.
6. Move the disjunctions down to the literals.
7. Eliminate the conjunctions.
8. Rename the variables, if necessary.
9. Eliminate the universal quantifiers.

Resolution Proofs as Search

- **Search Problem**
 - States: Content of knowledge base in CNF
 - Initial state: Knowledge base with negated theorem to prove
 - Successor function: Resolution inference rule with unify
 - Goal test: Does knowledge base contain the empty clause 'nil'?
- **Search Algorithm**
 - Depth first search (used in PROLOG)
 - Note: Possibly infinite state space
 - Example:
 - IsPerson(Fred)
 - IsPerson(y) IsPerson(mother(y))
 - Goal: $\exists x: \text{IsPerson}(x)$
 - Answers: {x/Fred} and {x/mother(Fred)} and {x/mother(mother(Fred))} and ...

Example

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?

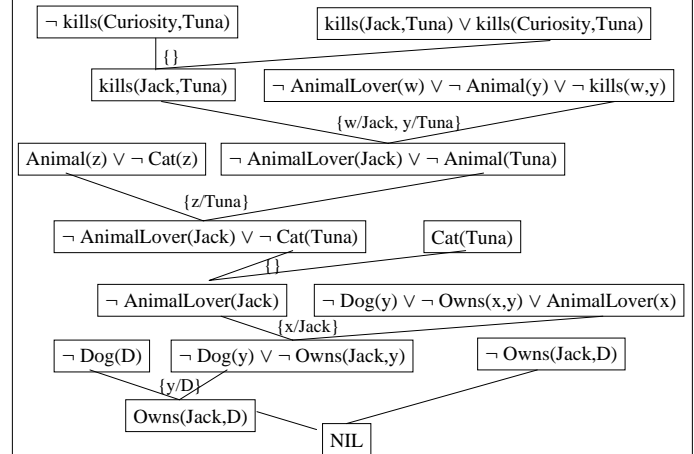
Original Sentences (Plus Background Knowledge)

1. $\exists x : Dog(x) \wedge Owns(Jack, x)$
2. $\forall x; (\exists y \text{ } Dog(y) \wedge Owns(x, y)) \rightarrow AnimalLover(x)$
3. $\forall x; AnimalLover(x) \rightarrow (\forall y \text{ } Animal(y) \rightarrow \neg Kills(x, y))$
4. $Kills(Jack, Tuna) \vee Kills(Curiosity, Tuna)$
5. $Cat(Tuna)$
6. $\forall x : Cat(x) \rightarrow Animal(x)$

Conjunctive Normal Form

$Dog(D)$ (D is a placeholder for the dogs unknown name (i.e. Skolem symbol/function). Think of D like "JohnDoe")
 $Owns(Jack, D)$
 $\neg Dog(y) \vee \neg Owns(x, y) \vee AnimalLover(x)$
 $\neg AnimalLover(w) \vee \neg Animal(y) \vee \neg Kills(w, y)$
 $Kills(Jack, Tuna) \vee Kills(Curiosity, Tuna)$
 $Cat(Tuna)$
 $\neg Cat(z) \vee Animal(z)$
 $\neg Kills(Curiosity, Tuna)$

Proof by Resolution



Completeness

Resolution with unification applied to clausal form, is a complete inference procedure.

In practice, still significant search problem!
Many different search strategies: resolution strategies.

Strategies for Selecting Clauses

unit-preference strategy: Give preference to resolutions involving the clauses with the smallest number of literals.

set-of-support strategy: Try to resolve with the negated theorem or a clause generated by resolution from that clause.

subsumption: Eliminates all sentences that are subsumed (i.e., more specific than) an existing sentence in the KB.

May still require exponential time.

Proofs can be Lengthy

A relatively straightforward KB can quickly overwhelm general resolution methods.

Resolution strategies reduce the problem somewhat, but not completely.

As a consequence, many practical Knowledge Representation formalisms in AI use a restricted form and specialized inference.

- Logic programming (Prolog)
- Production systems
- Frame systems and semantic networks
- Description logics

Successes in Rule-Based Reasoning

Expert systems

- DENDRAL (Buchanan *et al.*, 1969)
- MYCIN (Feigenbaum, Buchanan, Shortliffe)
- PROSPECTOR (Duda *et al.*, 1979)
- R1 (McDermott, 1982)

Successes in Rule-Based Reasoning

- **DENDRAL (Buchanan *et al.*, 1969)**
 - Infers molecular structure from the information provided by a mass spectrometer
 - Generate-and-test method
 - if there are peaks at x_1 and x_2 s.t.
 - $x_1 + x_2 = M + 28$
 - $x_1 - 28$ is a high peak
 - $x_2 - 28$ is a high peak
 - At least one of x_1 and x_2 is high
 - then there is a ketone subgroup

Successes in Rule-Based Reasoning

- **MYCIN (Feigenbaum, Buchanan, Shortliffe)**
 - Diagnosis of blood infections
 - 450 rules; performs as well as experts
 - Incorporated **certainty factors**
- If: (1) the stain of the organism is gram-positive, and
(2) the morphology of the organism is coccus, and
(3) the growth conformation of the organism is clumps,
then there is suggestive evidence (0.7) that the identity of the organism is staphylococcus.

Successes in Rule-Based Reasoning

- **PROSPECTOR (Duda *et al.*, 1979)**
 - Correctly recommended exploratory drilling at geological site
 - Rule-based system founded on probability theory
 - **R1 (McDermott, 1982)**
 - Designs configurations of computer components
 - About 10,000 rules
 - Uses meta-rules to change context
- If: current context is ?x
then: deactivate ?x context
and activate ?y context

Cognitive Modeling with Rule-Based Systems

SOAR is a general architecture for building intelligent systems.

- Long term memory consists of rules
- Working memory describes current state
- All problem solving, including deciding what rule to execute, is state space search
- Successful rule sequences are *chunked* into new rules
- Control strategy embodied in terms of meta-rules

Properties of Knowledge-Based Systems

Advantages

1. Expressibility*
2. Simplicity of inference procedures*
3. Modifiability*
4. Explainability
5. Machine readability
6. Parallelism*

Disadvantages

1. Difficulties in expressibility
2. Undesirable interactions among rules
3. Non-transparent behavior
4. Difficult debugging
5. Slow
6. Where does the knowledge base come from???