

1943–1956 The Beginnings of AI

- 1943** McCulloch and Pitts show networks of neurons can compute and learn any function
- 1950** Shannon (1950) and Turing (1953) wrote chess programs
- 1951** Minsky and Edmonds build the first neural network computer (SNARC)
- 1956** Dartmouth Conference – Newell and Simon brought a reasoning program The Logic Theorist which proved theorems.


Slide CS472 – Knowledge-Based Systems 1

1952–1969 The Good Years

- 1952** Samuel's checkers player
- 1958** McCarthy designed **LISP**, helped invent time-sharing, and created Advice Taker (a domain independent reasoning system)
- 1960's** Microworlds – solving limited problems: SAINT (1963), ANALOGY (1968), STUDENT (1967), blocksworld invented.
- 1962** Perceptron Convergence Theorem is proved.


Slide CS472 – Knowledge-Based Systems 2

Example ANALOGY Problem



Slide CS472 – Knowledge-Based Systems 3

Blocksworld



Slide CS472 – Knowledge-Based Systems 4

History of AI 1966–1974 A Dose of Reality

- Herb Simon (1957)
- Machine translation
- Knowledge-poor programs
- Intractable problems, lack of computing power (Lighthill Report, 1973)
- Limitations in knowledge representation (Minsky and Papert, 1969)

Slide CS472 – Knowledge-Based Systems 5

Knowledge Representation

- Human intelligence relies on a lot of background knowledge (the more you know, the easier many tasks become / “**knowledge is power**”)
- E.g. SEND + MORE = MONEY puzzle.
- Natural language understanding
 - Time flies like an arrow.
 - Fruit flies like bananas.
 - The spirit is willing but the flesh is weak. (English)
 - The vodka is good but the meat is rotten. (Russian)
- Or: Plan a trip to L.A.

Slide CS472 – Knowledge-Based Systems 6

- Q. How did we encode (domain) knowledge so far?
For search problems?

For learning problems?

Fine for limited amounts of knowledge / well-defined domains.
Otherwise: **knowledge-based systems approach**.

Slide CS472 – Knowledge-Based Systems 7

Knowledge-Based Systems / Agents

Key components:

- **knowledge base**: a set of *sentences* expressed in some knowledge representation language
- **inference / reasoning mechanisms** to query what is known and to derive new information or make decisions

Natural candidate: logical language (propositional / first-order) combined with a logical inference mechanism

How close to human thought?

In any case, appears reasonable strategy for machines.

Slide CS472 – Knowledge-Based Systems 8

Logic as a Knowledge Representation

Components:

syntax

semantics (link to the world)

logical reasoning

entailment: $\alpha \models \beta$ iff, in every model in which α is true, β is also true.

inference algorithm

$KB \vdash \alpha$, i.e., α is derived from KB

To make it work: **soundness** and **completeness**.

Slide CS472 – Knowledge-Based Systems 9

Soundness and Completeness

An inference algorithm that derives only entailed sentences is called **sound** or **truth-preserving**.

$KB \vdash \alpha$ implies $KB \models \alpha$.

An inference algorithm is **complete** if it can derive any sentence that is entailed.


$KB \models \alpha$ implies $KB \vdash \alpha$.

Why so important?

Allow computer to ignore semantics and “just push symbols”!

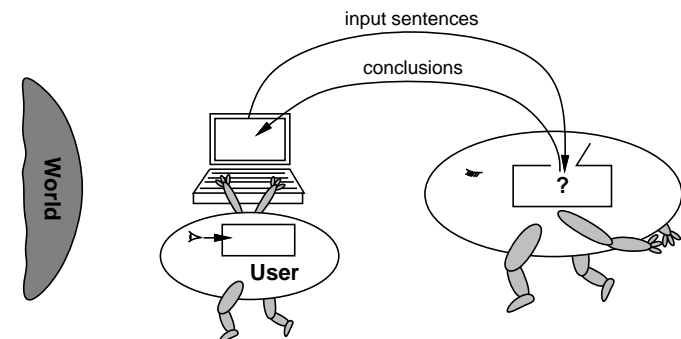
Slide CS472 – Knowledge-Based Systems 10

Connecting Sentences to the World



Slide CS472 – Knowledge-Based Systems 11

Tenuous Link to Real World



All computer has are sentences (hopefully about the world).
Sensors can provide some grounding.

Slide CS472 – Knowledge-Based Systems 12

KR Language: Propositional Logic

Syntax: build sentences from atomic propositions, using
connectives $\wedge, \vee, \neg, \Rightarrow, \Leftrightarrow$.
(and / or / not / implies / equivalence (biconditional))

E.g.: $((\neg P) \vee (Q \wedge R)) \Rightarrow S$

Slide CS472 – Knowledge-Based Systems 13

Semantics

P	Q	$\neg P$	$P \wedge Q$	$P \vee Q$	$P \Rightarrow Q$	$P \Leftrightarrow Q$
<i>False</i>	<i>False</i>	<i>True</i>	<i>False</i>	<i>False</i>	<i>True</i>	<i>True</i>
<i>False</i>	<i>True</i>	<i>True</i>	<i>False</i>	<i>True</i>	<i>True</i>	<i>False</i>
<i>True</i>	<i>False</i>	<i>False</i>	<i>False</i>	<i>True</i>	<i>False</i>	<i>False</i>
<i>True</i>	<i>True</i>	<i>False</i>	<i>True</i>	<i>True</i>	<i>True</i>	<i>True</i>

Note: \Rightarrow somewhat counterintuitive.

What's the truth value of "5 is even implies Sam is smart"?

Slide CS472 – Knowledge-Based Systems 14

First-Order Logic as a Knowledge Representation

propositions: "It is raining" becomes RAINING

operators: \vee, \wedge, \neg or $\sim, =, \Rightarrow, \Leftrightarrow$

predicates: $Man(SOCRATES)$ for "Socrates is a man."

$On(A, B)$. Can be functions: *on - top - of*(A).

quantifiers:

All men are mortal.

$$\forall x : Man(x) \Rightarrow Mortal(x)$$

Some man is mortal.

$$\exists x : Man(x) \Rightarrow Mortal(x)$$

Slide CS472 – Knowledge-Based Systems 15

Reasoning Methods: Rules of Inference

1. Modus Ponens:

Assume: $P \Rightarrow Q$ If raining, then soggy courts.

and P It is raining.

Then: Q Soggy Courts.

2. Modus Tollens:

Assume: $P \Rightarrow Q$ If raining, then soggy courts.

and $\neg Q$ No soggy courts.

Then: $\neg P$ It is not raining.

Slide CS472 – Knowledge-Based Systems 16

Representing Facts

1. Pingali is a CS professor.
2. All CS professors are ENG professors.
3. Fuchs is the dean.
4. All ENG professors are a friend of the dean or don't know him.
5. Everyone is a friend of someone.
6. People only criticize deans they are not friends of.
7. Pingali criticized Fuchs.

Slide CS472 – Knowledge-Based Systems 17

Representing Subset Hierarchies

Member:

$CSPROF(Pingali)$

or

$member(Pingali, CSPROF)$

subset:

$\forall x : CSPROF(x) \Rightarrow ENGPROF(x)$

or

$isa(CSPROF, ENGPROF)$

Slide CS472 – Knowledge-Based Systems 18

Is Pingali a friend of Fuchs?

$\neg \text{friend-of}(\text{Pingali}, \text{Fuchs})$

Slide CS472 – Knowledge-Based Systems 19

Forward Chaining

Given a fact p to be added to the KB,

1. Find all implications I that have p as a premise
2. For each i in I , if the other premises in i are already known to hold
 - (a) Add the consequent in i to the KB

Continue until no more facts can be inferred.

Slide CS472 – Knowledge-Based Systems 20

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. Establish the premises of all i in I via backward chaining.

Slide CS472 – Knowledge-Based Systems 21

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.

Slide CS472 – Knowledge-Based Systems 22

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 nonnegated atoms.

Slide CS472 – Knowledge-Based Systems 23

– First-Order

- * 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*.

Slide CS472 – Knowledge-Based Systems 24

Resolution Rule of Inference

Assume:	$E_1 \vee E_2$	playing tennis or raining
and	$\neg E_2 \vee E_3$	not raining or working
<hr/>		
Then:	$E_1 \vee E_3$	playing tennis or working

Slide CS472 – Knowledge-Based Systems 25

General Resolution Rule

If $(L_1 \vee L_2 \vee \dots L_k)$ is true,
and $(\neg L_k \vee L_{k+1} \vee \dots L_m)$ true,
then we can conclude that
 $(L_1 \vee L_2 \vee \dots L_{k-1} \vee L_{k+1} \vee \dots \vee L_m)$ is true.

Slide CS472 – Knowledge-Based Systems 26

Algorithm: Resolution Proof

- Negate the theorem to be proved, and add the result to the list of axioms.
- Put the list of axioms into conjunctive normal form.
- Until there is no resolvable pair of clauses,
 - Find resolvable clauses and resolve them.
 - Add the results of resolution to the list of clauses.
 - If NIL (empty clause) is produced, stop and report that the (original) theorem is true.
- Report that the (original) theorem is false.

Slide CS472 – Knowledge-Based Systems 27

Resolution Example

Example: Prove $\neg P$

Axioms:

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

Slide CS472 – Knowledge-Based Systems 28

Resolution Example (cont.)

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

Slide CS472 – Knowledge-Based Systems 29

Resolution Example: FOL

Example: Prove $\text{bird}(\text{tweety})$

Axioms:

Regular

CNF

- 1: $\forall x : \text{feathers}(x) \rightarrow \text{bird}(x)$
- 2: $\text{feathers}(\text{tweety})$
- 3:

Slide CS472 – Knowledge-Based Systems 30

Resolution Theorem Proving

- sound (for propositional and FOL)
- 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.

Slide CS472 – Knowledge-Based Systems 31

Resolution

- I **put KB in CNF (clausal) form**
all variables universally quantified
main trick: “skolemization” to remove existentials
idea: invent names for unknown objects known to exist
- II **use unification** to match atomic sentences
- III **apply resolution rule** to the clausal set combined with negated goal. Attempt to generate empty clause.

Slide CS472 – Knowledge-Based Systems 32

Converting more complicated axioms to CNF

Axiom:

$$\begin{aligned}\forall x : brick(x) \rightarrow & ((\exists y : on(x, y) \wedge \neg pyramid(y)) \\ & \wedge (\neg \exists y : on(x, y) \wedge on(y, x)) \\ & \wedge (\forall y : \neg brick(y) \rightarrow \neg equal(x, y))) \\ \neg brick(x) \vee & on(x, support(x)) \\ \neg brick(w) \vee & \neg pyramid(support(w)) \\ \neg brick(u) \vee & \neg on(u, y) \vee \neg on(y, u) \\ \neg brick(v) \vee & brick(z) \vee \neg equal(v, z)\end{aligned}$$

Slide CS472 – Knowledge-Based Systems 33

1. Eliminate Implications

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

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

Slide CS472 – Knowledge-Based Systems 34

2. Move negations down to the atomic formulas

$$\begin{aligned}\neg(E_1 \wedge E_2) & \iff (\neg E_1) \vee (\neg E_2) \\ \neg(E_1 \vee E_2) & \iff (\neg E_1) \wedge (\neg E_2) \\ \neg(\neg E_1) & \iff E_1 \\ \neg \forall x : E_1(x) & \iff \exists x : \neg E_1(x) \\ \neg \exists x : E_1(x) & \iff \forall x : \neg E_1(x)\end{aligned}$$

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

Slide CS472 – Knowledge-Based Systems 35

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:

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

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

Slide CS472 – Knowledge-Based Systems 36

4. Rename variables as necessary

We want no two variables of the same name.

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

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

Slide CS472 – Knowledge-Based Systems 37

5. Move the universal quantifiers to the left

This works because each quantifier uses a unique variable name.

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

Slide CS472 – Knowledge-Based Systems 38

6. Move disjunctions down to the literals

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

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

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

Slide CS472 – Knowledge-Based Systems 39

7. Eliminate the conjunctions

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

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

Slide CS472 – Knowledge-Based Systems 40

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

$\forall x : \neg brick(x) \vee on(x, S1(x))$

$\forall x : \neg brick(x) \vee \neg pyramid(S1(x))$

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

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

$\forall x : \neg brick(x) \vee on(x, S1(x))$

$\forall w : \neg brick(w) \vee \neg pyramid(S1(w))$

$\forall u \forall y : \neg brick(u) \vee \neg on(u, y) \vee \neg on(y, u)$

$\forall v \forall z : \neg brick(v) \vee brick(z) \vee \neg equal(v, z)$

Slide CS472 – Knowledge-Based Systems 41

9. Eliminate the universal quantifiers

$\neg brick(x) \vee on(x, S1(x))$

$\neg brick(w) \vee \neg pyramid(S1(w))$

$\neg brick(u) \vee \neg on(u, y) \vee \neg on(y, u)$

$\neg brick(v) \vee brick(z) \vee \neg equal(v, z)$

Slide CS472 – Knowledge-Based Systems 42

Algorithm: Putting Axioms into Clausal Form

- Eliminate the implications.
- Move the negations down to the atomic formulas.
- Eliminate the existential quantifiers.
- Rename the variables, if necessary.
- Move the universal quantifiers to the left.
- Move the disjunctions down to the literals.
- Eliminate the conjunctions.
- Rename the variables, if necessary.
- Eliminate the universal quantifiers.

Slide CS472 – Knowledge-Based Systems 43

Unification

UNIFY (P,Q) takes two atomic sentences P and Q and returns a substitution that makes P and Q **look the same**.

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.

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

Slide CS472 – Knowledge-Based Systems 44

Unification — Purpose

Given:

$Knows(John, x) \rightarrow Hates(John, x)$

$Knows(John, Jim)$

Derive:

$Hates(John, Jim)$

Need **unifier** $\{x/Jim\}$ for resolution to work.
(simplest case)

Slide CS472 – Knowledge-Based Systems 45

$\neg Knows(John, x) \vee Hates(John, x)$

$Knows(John, Jim)$

How do we resolve? First, match them.

Solution:

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

Gives

$\neg Knows(John, Jim) \vee Hates(John, Jim)$ and

$Knows(John, Jim)$

Conclude by resolution

$Hates(John, Jim)$

Slide CS472 – Knowledge-Based Systems 46

Unification (example)

one rule:

$Knows(John, x) \rightarrow Hates(John, x)$

facts:

$Knows(John, Jim)$

$Knows(y, Leo)$

$Knows(z, Mother(z))$

$Knows(x, Jane)$

Who does John hate?

Slide CS472 – Knowledge-Based Systems 47

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

$UNIFY(Knows(John, x), Knows(y, Leo)) = \{x/Leo, y/John\}$

$UNIFY(Knows(John, x), Knows(z, Mother(z))) =$
 $\{z/John, x/Mother(John)\}$

$UNIFY(Knows(John, x), Knows(x, Jane)) = fail$

Slide CS472 – Knowledge-Based Systems 48

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.

$\text{UNIFY}(\text{Knows}(\text{John}, x), \text{Knows}(y, z))$
 $= \{y/\text{John}, x/z\}$
 or $\{y/\text{John}, x/z, z/\text{Freda}\}$
 or $\{y/\text{John}, x/\text{John}, z/\text{John}\}$
 or

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

Slide CS472 – Knowledge-Based Systems 49

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?

Slide CS472 – Knowledge-Based Systems 50

Original Sentences (Plus Background Knowledge)

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

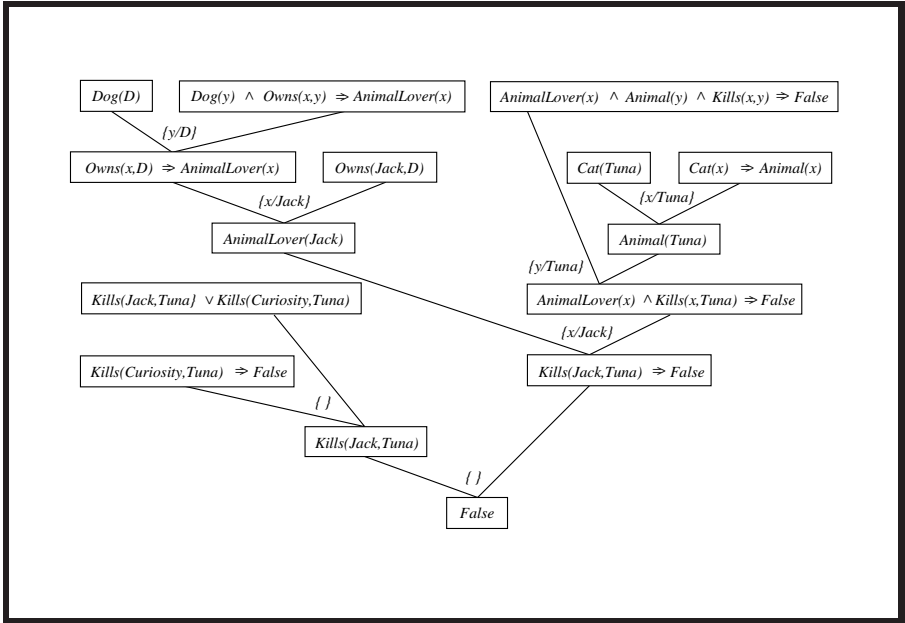
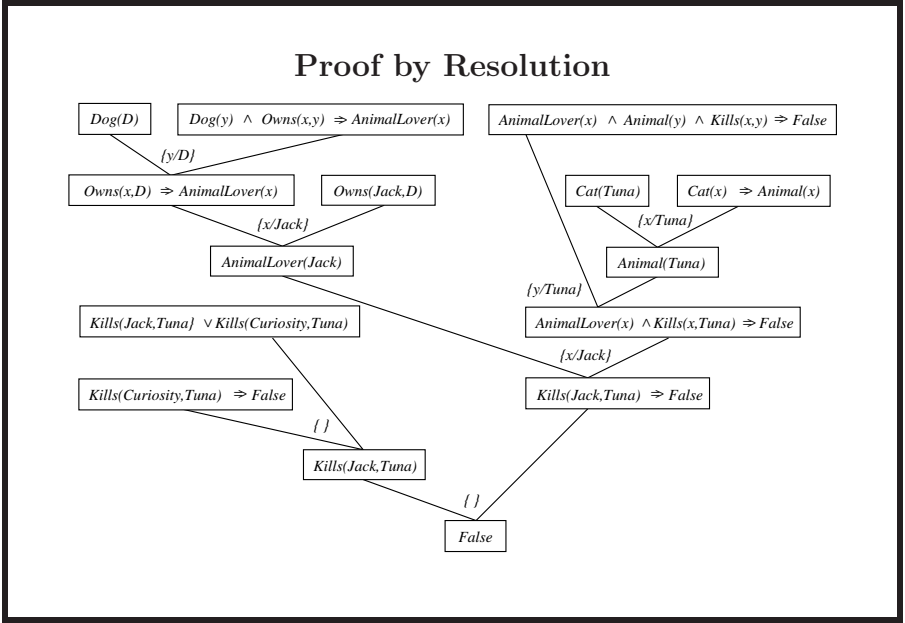
Slide CS472 – Knowledge-Based Systems 51

Conjunctive Normal Form

1. $\text{Dog}(D)$ (D is the function that finds Jack's dog)
2. $\text{Owns}(\text{Jack}, D)$
3. $\neg \text{Dog}(S(x)) \vee \neg \text{Owns}(x, S(x)) \vee \text{AnimalLover}(x)$
4. $\neg \text{AnimalLover}(w) \vee \neg \text{Animal}(y) \vee \neg \text{Kills}(w, y)$
5. $\text{Kills}(\text{Jack}, \text{Tuna}) \vee \text{Kills}(\text{Curiosity}, \text{Tuna})$
6. $\text{Cat}(\text{Tuna})$
7. $\neg \text{Cat}(z) \vee \text{Animal}(z)$

Slide CS472 – Knowledge-Based Systems 52

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.

Slide CS472 – Knowledge-Based Systems 57

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**.

Slide CS472 – Knowledge-Based Systems 58

Practical Knowledge-Based Systems

- **Theorem provers / logic programming**
- **Production systems**
forward chaining / if-then-rules / **expert** systems
- **Frame systems and semantic networks**
- **Description logics**

Slide CS472 – Knowledge-Based Systems 59

Theorem provers / logic programming

Theorem provers: generally based on resolution
many different strategies to improve efficiency

Logic programming: program statements directly in
restricted FOL.

Execution: search for proof of goal/query
using backward chaining with depth first-search.

In certain cases too inefficient.

Slide CS472 – Knowledge-Based Systems 60

Production systems

- rich history in AI
- “**expert system**” boom in 70’s / 80’s

Basic idea:

capture knowledge of human expert in a
large set of “if-then” rules

(really, logical implication \Rightarrow)

“**production rules**”

Slide CS472 – Knowledge-Based Systems 61

Components of Rule-Based Systems

working memory: set of positive literals with no variables

rule memory: set of inference rules

$$p_1 \wedge p_2 \dots \rightarrow a_1 \wedge a_2 \dots$$

where the p_i are literals, and the a_i are actions to take
when the p_i are all satisfied

rule interpreter: inference engine

Slide CS472 – Knowledge-Based Systems 62

Sample Knowledge Base

Working Memory

(in robot room1)
(armempty robot)
(in table room1)
(on cup table)
(object table)
(object cup)
(room room1)
(room room2)

Slide CS472 – Knowledge-Based Systems 63

1. $(in\ robot\ ?x) \wedge (room\ ?y) \rightarrow$
 $(walk\ ?x\ ?y) \wedge (add\ (in\ robot\ ?y)) \wedge (delete\ (in\ robot\ ?x))$
2. $(in\ robot\ ?x) \wedge (in\ ?y\ ?x) \wedge (object\ ?y) \wedge (not\ (at\ robot\ ?y))$
 $\rightarrow (walk\ ?x\ ?y) \wedge (add\ (at\ robot\ ?y))$
3. $(in\ robot\ ?x) \wedge (at\ robot\ ?y) \wedge (clear\ ?y) \wedge (armempty\ robot)$
 $\wedge (room\ ?z) \rightarrow (push\ ?y\ ?z) \wedge (add\ (in\ robot\ ?z)) \wedge (add\ (in\ ?y\ ?z)) \wedge (delete\ (in\ robot\ ?x)) \wedge (delete\ (in\ ?y\ ?x))$
4. $(at\ robot\ ?x) \wedge (armempty\ robot) \wedge (on\ ?y\ ?x) \rightarrow$
 $(pickup\ ?y) \wedge (add\ (holding\ robot\ ?y)) \wedge (add\ (clear\ ?x)) \wedge$
 $(delete\ (armempty\ robot)) \wedge (delete\ (on\ ?y\ ?x))$
5. $(holding\ robot\ ?x) \rightarrow (putdown\ ?x) \wedge (add\ (armempty\ robot)) \wedge (delete\ (holding\ robot\ ?x)) \wedge (add\ (on\ ?x\ floor))$

Slide CS472 – Knowledge-Based Systems 64

Reasoning with Rules

Forward Reasoning

1. Until no rule can fire or goal state is achieved,
 - (a) Find all rules whose *left* sides match assertions in working memory.
 - (b) Pick some to execute; modify working memory by applying the *right* sides of the rules.

Slide CS472 – Knowledge-Based Systems 65

Three Parts to Forward-Chaining Rule Interpreter

Match: identifying which rules are applicable at any given point in the reasoning

Conflict Resolution: selecting which of many rules should be applied at any given point in the reasoning

Execute: execute the right-hand side of the rule

Slide CS472 – Knowledge-Based Systems 66

Example: forward chaining

Goal: (in table room2)

1.

Rule 1	x = room1, y = room1, y = room2
Rule 2	x = room1, y = table

Choose randomly; assume rule 2.

Walk to table.

Add: (at robot table)

2.

Rule 1	same bindings
Rule 2?	
Rule 4	x = table, y = cup

Slide CS472 – Knowledge-Based Systems 67

Choose randomly; assume rule 4.

Pick-up cup.

Add: (holding robot cup)
(clear table)

Delete: (armempty robot)

3.

Rule 1	same bindings
Rule 5	x = cup

Choose randomly; assume rule 5.

Putdown cup.

Add: (armempty robot)
Add: (on cup floor)

Slide CS472 – Knowledge-Based Systems 68

Delete: (holding robot cup)

4. Rule 1 same bindings
Rule 3 $x = \text{room1}$, $y = \text{table}$, $z = \text{room1}$ or room2

Choose randomly; assume rule 3.

Push table (to) room2.

Add: (in robot room2)
(in table room2)

Delete: (in robot room1)
(in table room1)

Slide CS472 – Knowledge-Based Systems 69

Matching for Forward-Chaining

- requires smart indexing of the rules
- requires unification

Problem: For practical systems, applying unification in a straightforward manner will be very inefficient.

Rete algorithm: network-based data structure facilitates unification.

Slide CS472 – Knowledge-Based Systems 70

Successes in Rule-Based Reasoning

Expert systems

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

Slide CS472 – Knowledge-Based Systems 71

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

Slide CS472 – Knowledge-Based Systems 72

- 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.

Slide CS472 – Knowledge-Based Systems 73

- PROSPECTOR (Duda *et al.*, 1979)

- correctly recommended exploratory drilling at a 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

Slide CS472 – Knowledge-Based Systems 74

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.

Slide CS472 – Knowledge-Based Systems 75

Example Syntax for Control Rule

Meta-rule

Under conditions *A* and *B*,
Rules that do {*not*} mention *X*

{*at all*,
in their LHS,
in their RHS }

will

{*definitely be useless*,
probably be useless, ...
probably be especially useful,
definitely be especially useful }

Slide CS472 – Knowledge-Based Systems 76

Advantages of Knowledge-Based Systems

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

Slide CS472 – Knowledge-Based Systems 77

Disadvantages of Knowledge-Based Systems

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???

Slide CS472 – Knowledge-Based Systems 78