

Logic, inference and expert systems

How an agent can know things, prove new things, and explain itself: from truth tables to MYCIN.

01 COLD OPEN 0-4 min

Minesweeper is a logic exam

Minesweeper is model checking

Click a hidden cell. The tool enumerates every mine layout consistent with the numbers.



Pick any hidden cell next to a number. Predict first: safe, mine, or can't tell?

FRONTIER CELLS

9

MODELS CHECKED

$2^9 = 512$

CONSISTENT

2

PROVABLE NOW

6



POP CULTURE

Every Windows PC shipped with a logic tutor for decades. Good players never guess: they **prove** a cell is safe from the numbers. Sherlock calls it deduction; we call it **entailment**.

Each number is a sentence about its hidden neighbours. A cell is safe if it is safe in **every** world that agrees with all the numbers.

By the end of today you'll know exactly what the tool on the left is doing, and why it doesn't scale.

02 CONCEPT 4-8 min

A knowledge-based agent: TELL, ASK, repeat

A **knowledge base (KB)** is a set of sentences in a formal language. The agent **TELL**s it what it perceives and **ASK**s it what to do; an **inference engine** works out answers that were never

stored explicitly.

```
function KB-AGENT(percept):           ; t = time step
  TELL(KB, MAKE-PERCEPT-SENTENCE(percept, t))
  action = ASK(KB, MAKE-ACTION-QUERY(t))
  TELL(KB, MAKE-ACTION-SENTENCE(action, t))
  t = t + 1
  return action
```

FACT
One specific truth. Fever(Rahim)

RULE
A general truth that makes new facts. $\text{Fever}(x) \wedge \text{LowPlt}(x) \Rightarrow \text{Suspect}(x)$

REPRESENTATION	GOOD AT
Propositional logic	Small, fixed worlds (a Minesweeper corner)
First-order logic	Objects and relations: "every patient..."
IF-THEN rules	Expert know-how, easy to explain
Semantic nets, frames	Categories, inheritance, defaults
Probabilities	Uncertainty (Week 8: Bayes)

 **PRO TIP**

Declarative: say *what* is true and let the engine decide how to use it. **Procedural:** hard-code *how*. A KB can be updated without rewriting the program.

03 **CONCEPT** 8-12 min

Propositional logic in five symbols

Atoms (P , Q , *Dengue*) are true or false. Connectives build sentences, tightest first: $\neg$, $\wedge$, $\vee$, $\Rightarrow$, $\Leftrightarrow$. A **model** is one assignment of T/F to every atom: one possible world.

P	Q	$\neg P$	$P \wedge Q$	$P \vee Q$	$P \Rightarrow Q$	$P \Leftrightarrow Q$
F	F	T	F	F	T	T
F	T	T	F	T	T	F
T	F	F	F	T	F	F
T	T	F	T	T	T	T

 **GOTCHA**

$\Rightarrow$ **is not causation**. "If Messi is from Bogura, then $2 + 2 = 5$ " is *true*: a false premise makes $P \Rightarrow Q$ true (the highlighted rows). $P \Rightarrow Q$ only promises: never P true with Q false.

$$P \Rightarrow Q \equiv \neg P \vee Q$$

Entailment, α follows from the KB: $KB \models \alpha \iff M(KB) \subseteq M(\alpha)$, where $M(\cdot)$ is the set of models that make a sentence true.

Precedence: $\neg A \wedge B \vee C \Rightarrow D$ means $((\neg A \wedge B) \vee C) \Rightarrow D$. When in doubt, bracket.

04 **LIVE LAB** 12-18 min

Entailment: check every possible world

Entailment by model checking $KB \models \alpha$ when α is true in *every* row where the KB is true.

R: Red is the impostor · B: Blue is the impostor · V: someone saw Blue vent

KNOWLEDGE BASE

KB1 $R \vee B$

KB2 $B \Rightarrow V$

KB3 $\neg V$

ASK $\alpha =$

R	B	V	KB1	KB2	KB3	KB	α
F	F	F	.	.	.	.	.
F	F	T	.	.	.	.	.
F	T	F	.	.	.	.	.
F	T	T	.	.	.	.	.
T	F	F	.	.	.	.	.
T	F	T	.	.	.	.	.
T	T	F	.	.	.	.	.
T	T	T	.	.	.	.	.

3 symbols $\rightarrow 2^3 = 8$ models · KB true in ? so far

05 **WORKED EXAMPLE** 18-22 min

Worked example: who is the impostor?

GIVEN

An Among Us meeting. (1) "It's Red or Blue." (2) "If Blue is the impostor, someone saw Blue vent." (3) "Nobody saw Blue vent." Symbols: R, B, V.

FIND

Prove by truth table that the KB entails R. Then prove it again with inference rules.

R	F	F	F	F	T	T	T	T
B	F	F	T	T	F	F	T	T

V	F	T	F	T	F	T	F	T
KB	F	F	F	F	T	F	F	F

Method: truth-table entailment, Russell & Norvig "Logical Agents"

SOLUTION

1. Three symbols $\rightarrow 2^3 = 8$ models. $\neg V$ kills the 4 rows with $V = T$.

remaining: $V = F$

Write rows in binary order (000 $\rightarrow$ 111) so none is missed. Exams give marks for the full table.

2. $B \Rightarrow V$ with $V = F$ kills every $B = T$ row; $R \vee B$ with $B = F$ then forces $R = T$.

$$M(KB) = \{ R = T, B = F, V = F \}$$

3. R is true in every (here: the only) model of the KB:

$$KB \models R \quad \text{and also} \quad KB \models \neg B$$

4. By inference instead: modus tollens on $B \Rightarrow V$ and $\neg V$ gives $\neg B$; unit resolution on $R \vee B$ and $\neg B$ gives R.

Two proof steps instead of 2^n rows. That trade-off is the next slide.

ANSWER

$KB \models R$ (and $KB \models \neg B$). Vote Red.

06 CONCEPT 22-26 min

Valid, satisfiable, and proofs you can trust

VALID true in **every** model (a tautology). $L \vee \neg L$

SATISFIABLE true in **some** model. $D \Rightarrow L$

UNSATISFIABLE true in **no** model. $D \wedge \neg D$

Two bridges that power every prover:

$$\begin{aligned} KB \models \alpha &\iff (KB \Rightarrow \alpha) \text{ is valid} \\ &\iff KB \wedge \neg \alpha \text{ is unsatisfiable} \end{aligned}$$

The last one is proof by contradiction: the idea behind resolution.

MODUS PONENS

$$\frac{\alpha \Rightarrow \beta, \alpha}{\beta}$$

AND-ELIMINATION

$$\frac{\alpha \wedge \beta}{\alpha}$$

UNIT RESOLUTION

RESOLUTION

$$\frac{\alpha \vee \beta, \quad \neg\beta}{\alpha}$$

$$\frac{\alpha \vee \beta, \quad \neg\beta \vee \gamma}{\alpha \vee \gamma}$$

Sound: derives only entailed sentences (no lies). **Complete:** can derive *every* entailed sentence (no blind spots).

07 YOUR TURN 26-29 min

Your turn • 1: entailed or not?

1 EXAM CQ

Let **D** = "the patient has dengue", **L** = "platelets are low".

- (a) Does $\{D \Rightarrow L, \neg L\}$ entail $\neg D$?
- (b) Does $\{D \Rightarrow L, L\}$ entail D ?
- (c) Classify as valid, satisfiable or unsatisfiable: $L \vee \neg L$, $D \wedge \neg D$, $D \Rightarrow L$

SOLUTION

1. (a) Rows over (D, L): only $D = F, L = F$ satisfies both $D \Rightarrow L$ and $\neg L$. There $\neg D$ is true.

$$\{D \Rightarrow L, \neg L\} \models \neg D \quad (\text{modus tollens})$$

2. (b) The KB $\{D \Rightarrow L, L\}$ is true in two rows: ($D = T, L = T$) and ($D = F, L = T$). In the second, D is false.

$$\{D \Rightarrow L, L\} \not\models D$$

| Affirming the consequent. Clinically: malaria, viral fevers and some drugs also drop platelets.

3. (c) $L \vee \neg L$: valid. $D \wedge \neg D$: unsatisfiable. $D \Rightarrow L$: satisfiable but not valid (false at $D = T, L = F$).

ANSWER

(a) Entailed. (b) Not entailed; counter-example $D = F, L = T$. (c) Valid / unsatisfiable / satisfiable.

08 CONCEPT 29-32 min

Horn clauses: logic's fast lane

A **definite clause** is a disjunction with *exactly one* positive literal. It reads as a rule:

$$\neg A \vee \neg B \vee C \equiv A \wedge B \Rightarrow C$$

A **fact** is a definite clause with an empty body (just C). A **Horn clause** has *at most* one positive literal, so it also allows goal clauses like $\neg A \vee \neg B$ (i.e. $A \wedge B \Rightarrow \text{False}$).



PRO TIP

With Horn KBs, modus ponens alone is complete, and entailment takes **linear time** in the KB size. That is why Prolog, CLIPS and most expert systems live here.


	FORWARD CHAINING	BACKWARD CHAINING
Starts from	Known facts	The query (goal)
Style	Data-driven	Goal-driven
Does	Fire every rule whose premises hold	Find rules that conclude the goal; prove their premises
Risk	Derives lots of irrelevant facts	Loops; repeated sub-goals
Used in	ICU alarms, monitoring, CLIPS	Diagnosis, MYCIN, Prolog

09 **LIVE LAB** 32-39 min

Forward vs backward chaining, live

Chaining engine Rules are AND-joins (circles). Click a bottom node to toggle it as a known fact.

Goal **ADMIT** ▼



STEP

0 / 8

RULES FIRED

0

AGENDA (QUEUE) →

Fever

Eye pain

Plt < 100k

Abdo pain

Agenda starts with the known facts: Fever, EyePain, LowPlt, AbdoPain.

Worked example: forward chaining with an agenda

GIVEN

Rules: R1: $P \Rightarrow Q$, R2: $L \wedge M \Rightarrow P$, R3: $B \wedge L \Rightarrow M$, R4: $A \wedge P \Rightarrow L$, R5: $A \wedge B \Rightarrow L$. Facts: A, B.

FIND

Is Q entailed? Trace the agenda and the premise counts of every rule.



PRO TIP

Draw it as an AND-OR graph first (the sim's textbook rule base). Arcs across the edges into a rule mean AND; two rules into the same node mean OR.

Method: PL-FC-ENTAILS, Russell & Norvig "Logical Agents"

SOLUTION

1. Counts (unproved premises): R1 = 1, R2 = 2, R3 = 2, R4 = 2, R5 = 2. Agenda = [A, B].

2. Pop **A**: R4 $\rightarrow$ 1, R5 $\rightarrow$ 1. Pop **B**: R3 $\rightarrow$ 1, R5 $\rightarrow$ 0, so R5 fires: add **L**.

$$A \wedge B \Rightarrow L$$

3. Pop **L**: R2 $\rightarrow$ 1, R3 $\rightarrow$ 0, so R3 fires: add **M**.

$$B \wedge L \Rightarrow M$$

4. Pop **M**: R2 $\rightarrow$ 0, so R2 fires: add **P**.

$$L \wedge M \Rightarrow P$$

5. Pop **P**: R1 $\rightarrow$ 0 fires: add **Q**. R4 also reaches 0, but L is already known, so nothing new.

$$P \Rightarrow Q$$

6. Pop **Q**: it is the query. Stop (or run on to the fixed point: nothing else is added).

Each symbol is processed at most once and each rule count only goes down, so forward chaining is linear in the size of the KB.

ANSWER

Q is entailed. Derived in order L, M, P, Q; rules fired R5, R3, R2, R1 (R4 reaches 0 but adds nothing new).

Your turn · 2: Feluda chains backward

2 EXAM CQ

Feluda's rule book:

R1: Motive $\wedge$ Opportunity $\Rightarrow$ Suspect

R2: AtScene $\Rightarrow$ Opportunity

R3: NoAlibi $\Rightarrow$ Opportunity

R4: Inherits $\Rightarrow$ Motive

R5: Suspect $\wedge$ Fingerprints $\Rightarrow$ Arrest

Facts: **Inherits**, **NoAlibi**, **Fingerprints**. Backward-chain the query **Arrest**. Show the goal stack, where it backtracks, and count rules tried and fired.

SOLUTION

1. Goal **Arrest**: try R5, needs Suspect $\wedge$ Fingerprints. Stack: [Arrest].
2. **Suspect**: try R1, needs Motive $\wedge$ Opportunity. **Motive**: R4 needs Inherits: a fact $\checkmark$. Motive $\checkmark$.
3. **Opportunity**: try R2, needs AtScene: not a fact, no rule $\rightarrow \times$. Backtrack, try R3: NoAlibi is a fact $\checkmark$. Opportunity $\checkmark$.
 | Deepest stack: [Arrest, Suspect, Opportunity, AtScene], depth 4.
4. Suspect $\checkmark$ (R1). **Fingerprints**: a fact $\checkmark$. Arrest $\checkmark$ (R5).
5. Rules tried: R5, R1, R4, R2, R3 = 5. Rules fired: R4, R3, R1, R5 = 4.

ANSWER

Arrest is proved. R2 is the dead end (AtScene); backtracking to R3 saves the case. 5 rules tried, 4 fired, max stack depth 4.

First-order logic: objects, relations, “for all”

Propositional logic needs a new symbol for every patient: Dengue_{Rahim}, Dengue_{Karim}... FOL talks about **objects** and their **relations**.

Constants	Rahim, Dhaka, AedesAegypti
Predicates	Patient(x), Lives(x , Dhaka)
Functions	Mother(x), Platelets(x) (return an object)
Quantifiers	$\forall x$ for all, $\exists x$ there exists

✓ $\forall x \text{Dengue}(x) \Rightarrow \text{Fever}(x)$

$\forall$ goes with $\Rightarrow$. "For everything: *if* it has dengue, then fever."

✓ $\exists x \text{Patient}(x) \wedge \text{Malaria}(x)$

$\exists$ goes with $\wedge$. "Something is a patient *and* has malaria."

≠ $\forall x \exists y \text{Treats}(y, x)$ vs $\exists y \forall x \text{Treats}(y, x)$

Order matters: "everyone has *a* doctor" vs "*one* doctor treats everyone".

$$\neg \forall x P(x) \equiv \exists x \neg P(x)$$

13 THINK DEEPER 50–52 min

Critical thinking • the smart-student trap

3

TRAP

A student translates "All students are smart" and "Some student is smart" as:

(a) $\forall x \text{Student}(x) \wedge \text{Smart}(x)$

(b) $\exists x \text{Student}(x) \Rightarrow \text{Smart}(x)$

What does each sentence *actually* say? Fix both.

SOLUTION

1. ✗ (a) says *every object* in the domain is a student and is smart: the chair, Dhaka, Rahim's cat. One non-student makes it false.

$$\forall x \text{Student}(x) \wedge \text{Smart}(x)$$

2. ✓ Intended: every student is smart.

$$\forall x \text{Student}(x) \Rightarrow \text{Smart}(x)$$

3. ✗ (b) is true as soon as *anything* is not a student (the chair makes the implication vacuously true). It says almost nothing.

$$\exists x \text{Student}(x) \Rightarrow \text{Smart}(x)$$

4. ✓ Intended: some student is smart.

$$\exists x \text{Student}(x) \wedge \text{Smart}(x)$$

ANSWER

(a) claims everything is a smart student; (b) is almost always trivially true. $\forall$ pairs with $\Rightarrow$, $\exists$ pairs with $\wedge$.

Your turn • 3: English to FOL

4

EXAM CQ

Translate into first-order logic:

1. Every dengue patient has a fever.
2. Some patient in Dhaka has malaria.
3. No mosquito that carries the dengue virus is harmless.
4. Every patient is treated by some doctor.

SOLUTION

1. $\forall x \text{ Dengue}(x) \Rightarrow \text{Fever}(x)$

2. $\exists x \text{ Patient}(x) \wedge \text{In}(x, \text{Dhaka}) \wedge \text{Malaria}(x)$

3. $\forall x \text{ Mosquito}(x) \wedge \text{Carries}(x, \text{DengueVirus}) \Rightarrow \neg \text{Harmless}(x)$

Equivalent: $\neg \exists x \text{ Mosquito}(x) \wedge \text{Carries}(x, \text{DengueVirus}) \wedge \text{Harmless}(x)$. Either earns full marks.

4. $\forall x \text{ Patient}(x) \Rightarrow \exists y \text{ Doctor}(y) \wedge \text{Treats}(y, x)$

Swapping to $\exists y \forall x$ would claim one doctor treats every patient in Bangladesh.

ANSWER

$\forall$ with $\Rightarrow$ (1, 3, 4 outer), $\exists$ with $\wedge$ (2, 4 inner). Quantifier order in 4 decides who treats whom.

Unification: making a rule fit a fact

Generalized modus ponens lifts modus ponens to FOL: find a substitution θ that makes the rule's premises match the facts.


RULE	$\forall x \text{ Mosquito}(x) \wedge \text{Infected}(x) \Rightarrow \text{Risky}(x)$
FACTS	$\text{Mosquito}(M_7), \text{Infected}(M_7)$
UNIFY	$\theta = \{x/M_7\}$
CONCLUDE	$\text{SUBST}(\theta, \text{Risky}(x)) = \text{Risky}(M_7)$

UNIFY returns the **most general unifier** (MGU): the substitution that commits to as little as possible.

P	Q	UNIFY(P, Q)
Treats(Rahman, x)	Treats(Rahman, Karim)	{x/Karim}
Treats(Rahman, x)	Treats(y, Mother(y))	{y/Rahman, x/Mother(Rahman)}
Treats(Rahman, x)	Treats(x, Karim)	fail: x can't be both. Rename apart (x → z) and it works: {z/Rahman, x/Karim}
Treats(y, y)	Treats(Rahman, Karim)	fail

16 **CONCEPT** 58-61 min

Anatomy of an expert system



Inference engine. The reasoning part: matches rules against facts and decides which to fire next (forward or backward chaining, conflict resolution, certainty-factor arithmetic). Domain-independent.

An **expert system** captures a specialist's know-how as rules so non-specialists get expert-level advice, *with reasons*.

Domain expert: knows the field. **Knowledge engineer:** turns that into rules and tests them.

Click a block to see what it does.

17 **CONCEPT** 61-63 min

Inside the shell: CLIPS, semantic nets, frames


```
(deftemplate patient ; the shape of a fact
  (slot name)
  (slot fever (default no))
  (slot platelets (type NUMBER)))

(defrule dengue-suspect ; IF ... THEN
  (patient (name ?n) (fever yes)
    (platelets ?p < (< ?p 100000)))
  =>
  (assert (suspect ?n dengue))
  (printout t ?n ": suspect dengue, order NS1" crlf))

(assert (patient (name Rahim) (fever yes) (platelets 85000)))
```

```
(run)
; output → Rahim: suspect dengue, order NS1
```

Facts live in working memory; rules pattern-match them; $\Rightarrow$ separates IF from THEN. CLIPS chains **forward**.



Frame: Dengue

is-a	Viral fever
vector	Aedes aegypti
incubation	4–10 days
platelets (default)	low

18 **LIVE LAB** 63–68 min

MYCIN in miniature: certainty factors

RULE FIRING

$$CF = CF_{\text{rule}} \times \min(CF_{\text{premises}})$$

TWO POSITIVE CFS, SAME CONCLUSION

$$CF = CF_1 + CF_2 (1 - CF_1)$$

THRESHOLD

A premise counts only if its CF > 0.2

Mini-MYCIN: fever consult

Answer yes / no / unknown with a confidence. Ask it WHY and HOW.

⚠ Teaching toy with made-up rule strengths. **Not medical advice**; real triage follows national guidelines and a clinician.

QUESTION 1

Fever of 38.5 °C or more?

How sure is the answer? **1.0**

CONCLUSIONS (CERTAINTY FACTOR)

Dengue		0.00
Malaria		0.00
Typhoid		0.00

REASONING TRACE

Nothing yet. Answer the first question.

19 WORKED EXAMPLE 68-71 min

Worked example: combining certainty

GIVEN

Evidence CFs: fever 0.9, eye pain 0.7, platelets < 100k 0.8, rash 0.5, monsoon 1.0. Rules: D1: fever $\wedge$ eye pain $\Rightarrow$ dengue (0.4); D2: fever $\wedge$ low platelets $\Rightarrow$ dengue (0.6); D3: rash $\wedge$ monsoon $\Rightarrow$ dengue (0.3).

FIND

The certainty factor for dengue after all three rules fire.

Method: MYCIN / EMYCIN certainty factors, Patterson's chapter on uncertainty in expert systems

SOLUTION

1. Each rule: rule CF $\times$ the *minimum* premise CF (a chain is as strong as its weakest link).

$$CF_{D1} = 0.4 \times \min(0.9, 0.7) = 0.28, \quad CF_{D2} = 0.6 \times \min(0.9, 0.8) = 0.48$$

2. Two positive CFs for the same hypothesis combine:

$$CF = 0.28 + 0.48(1 - 0.28) = 0.6256$$

3. Third rule:

$$CF_{D_3} = 0.3 \times \min(0.5, 1.0) = 0.15$$

4. Combine again:

$$CF = 0.6256 + 0.15(1 - 0.6256) = 0.6818$$

Each rule closes only part of the gap to 1, so CF never exceeds 1. Order does not matter.

ANSWER

CF(dengue) $\approx$ 0.68: "suggestive evidence", not proof. Order of combination doesn't change it.

29 THINK DEEPER 71-74 min

Critical thinking • brittle experts, fluent LLMs

THINK DEEPER

Expert systems boomed in the 1980s and then faded. Why did they become *brittle*?

Writing the rules was the slow, expensive step, and the result only worked inside a narrow box. Ask a dengue rule base about a pregnant patient on steroids and it answers confidently anyway.

A STRONG ANSWER MENTIONS

- Knowledge-acquisition bottleneck: tacit know-how is hard to state as rules.
- No common sense outside the narrow domain, and no idea when it is out of its depth.
- Thousands of interacting rules are hard to maintain; CFs are ad hoc.

THINK DEEPER

Why might an LLM *plus* explicit rules or probabilities be safer in medicine than an LLM alone?

An LLM alone can sound certain, be wrong, and not show its reasoning. An explicit Bayesian or rule layer brings back MYCIN's audit trail. Your lecturer's **mLabLLM** (AAAI 2025) combines Bayesian reasoning with an LLM for dengue, malaria and chikungunya diagnosis along these lines.

A STRONG ANSWER MENTIONS

- The explicit layer is auditable and checkable against guidelines and local data.
- The LLM reads messy symptom descriptions; the reasoning layer constrains the conclusion.
- WHY/HOW-style explanations become possible again.

From MYCIN to dengue season in Dhaka

MYCIN (Stanford, 1970s) recommended antibiotics for blood infections using a few hundred backward-chaining rules with certainty factors. In blinded evaluations its advice was rated about as good as infectious-disease specialists'. It was never used on real patients: typing in every answer took too long, and nobody knew who would be liable.

Today its descendants are clinical decision support inside hospital record systems: drug-interaction alerts, sepsis warnings, dosing checks. Their big problem is *alert fatigue*: too many rules firing, so clinicians click "ignore".



BME CORNER

Bangladesh's 2023 dengue season was its worst on record, with over 300,000 hospital admissions. Triage rules like today's (fever + warning signs $\Rightarrow$ admit) are what a clinic app or a hospital's intake form encodes, so a junior doctor at 3 am applies the same logic as the national guideline.



FROM OUR LAB

mHealth Lab angle: explicit rules and probabilities give you something you can check against local data. That is the design idea behind combining Bayesian reasoning with an LLM in mLabLLM.

Pop quiz

1. $KB \models \alpha$ means...

- A. α is true in at least one model of the KB
- B. α is true in every model in which the KB is true ✓**
- C. The KB and α have the same truth table
- D. α can be derived by modus ponens only

Entailment: $M(KB) \subseteq M(\alpha)$. One model is satisfiability, not entailment.

2. A sentence that is true in *no* model is...

- A. Valid
- B. Satisfiable
- C. Unsatisfiable ✓**
- D. A Horn clause

Unsatisfiable (a contradiction), e.g. $D \wedge \neg D$. Resolution proves $KB \models \alpha$ by showing $KB \wedge \neg \alpha$ is unsatisfiable.

3. Which is the correct FOL for "Some student in this class likes AI"?

A. $\exists x \text{ InClass}(x) \Rightarrow \text{Likes}(x, \text{AI})$

B. $\exists x \text{ InClass}(x) \wedge \text{Likes}(x, \text{AI})$ ✓

C. $\forall x \text{ InClass}(x) \wedge \text{Likes}(x, \text{AI})$

D. $\forall x \text{ InClass}(x) \Rightarrow \text{Likes}(x, \text{AI})$

$\exists$ goes with $\wedge$. The $\Rightarrow$ version is true as soon as anything (a chair) is not in the class.

4. MYCIN mostly reasoned...

A. Forward, from every fact to every conclusion

B. Backward, from a hypothesis (which organism?) to the questions it needed ✓

C. With neural networks

D. By full truth tables

Goal-driven backward chaining. That is why it could answer WHY: every question belonged to a rule it was trying to use.

5. Two rules support dengue with CF 0.5 and 0.4. The combined CF is...

A. 0.9

B. 0.45

C. 0.7 ✓

D. 0.2

$0.5 + 0.4 \times (1 - 0.5) = 0.5 + 0.2 = 0.7$. Never add CFs directly: they would exceed 1.

23 RECAP 79-80 min

Three things to remember

01

$\text{KB} \models \alpha$: α holds in **every** model of the KB. Check models (2^n) or prove with sound rules; Horn KBs let forward/backward chaining do it in linear time.

02

FOL adds objects, predicates, functions and quantifiers. $\forall$ with $\Rightarrow$, $\exists$ with $\wedge$. Unification lets one rule serve every patient.

03

Expert system = knowledge base + inference engine + explanation. MYCIN's CFs: x min for AND, $\text{CF}_1 + \text{CF}_2(1 - \text{CF}_1)$ to combine.

Evidence: fever 1.0, cyclic chills 0.6, recent Bandarban trip 0.9. Rules M1: fever $\wedge$ chills $\Rightarrow$ malaria (0.4); M2: chills $\wedge$ hill trip $\Rightarrow$ malaria (0.7). Find CF(malaria), and say which piece of evidence to firm up first.

SOLUTION

1. M1: fever $\wedge$ chills $\Rightarrow$ malaria (0.4):

$$0.4 \times \min(1.0, 0.6) = 0.24$$

2. M2: chills $\wedge$ hill trip $\Rightarrow$ malaria (0.7):

$$0.7 \times \min(0.6, 0.9) = 0.42$$

3. Combine (either order):

$$0.24 + 0.42(1 - 0.24) = 0.42 + 0.24(1 - 0.42) = 0.5592$$

ANSWER

CF(malaria) $\approx$ 0.56. Chills (0.6) is the weakest link in both rules, so a firmer chills history raises it the most.

S. M. Sakeef Sani

Adjunct Lecturer, CSE · Northern University Bangladesh

smsakeefsani.hbitcorp.com