ARTIFICIAL INTELLIGENCE

The Very Idea

Vasant G. Honavar
Dorothy Foehr Huck and J. Lloyd Huck Chair in Biomedical Data Sciences and Artificial Intelligence
Professor of Data Sciences, Informatics, Computer Science, Bioinformatics & Genomics and Neuroscience
Director, Artificial Intelligence Research Laboratory
Director, Center for Artificial Intelligence Foundations and Scientific Applications
Associate Director, Institute for Computational and Data Sciences
Pennsylvania State University

vhonavar@psu.edu
http://faculty.ist.psu.edu/vhonavar
http://ailab.ist.psu.edu

Reasoning with Knowledge

"Knowledge is power" – Francis Bacon

- The approach to automating problem-solving using search
  - is general, and hence broadly applicable
  - But lacks knowledge – an essential ingredient of intelligence
- Knowledge-based systems or so-called expert systems
  - encode human knowledge in specific domains, e.g., medicine, chemistry, etc., to solve problems in those domains.
  - Do not seek to exhibit general intelligence, but rather, to solve problems that require expertise in settings where such expertise is scarce.

Reasoning with Knowledge

- Consider a company that manufactures automobiles.
- As part of its customer support efforts, the company has set up call center to help customers experiencing problems with their automobiles.
- For example, the call center may receive a call from a customer complaining that his car does not start.
- The call center worker must help the customer by diagnosing the cause of the problem.

Reasoning with Knowledge

- Suppose the call is answered by an expert mechanic who knows
  - automobile engines will not start if either the ignition or the battery or both are faulty
- To pinpoint the problem, the mechanic might use her knowledge
  - if the battery is faulty, the headlights won't work either.
- To keep things simple, we are ignoring other possible causes like faulty spark-plugs, carburetor, fuel injector, and so on
- The operator might ask the caller to try turning on the headlights.
  - If she is told that the headlights work just fine, she can diagnose the problem to be faulty ignition.
  - If on the other hand, she is told that the lights don't work, she may suggest that the caller might need to get the battery replaced.

Reasoning with Knowledge

- Human expertise is scarce – as are knowledgeable mechanics
- Human experts are prone to stress, fatigue, etc. that impacts their performance
- The automobile company might be interested in an expert system that plays the role of the mechanic on call
- Can AI help?
- How?

How can we design an expert system for auto diagnosis?

- Suppose we limit ourselves to malfunctions of only the Battery, Bulbs (headlights), Wiper Motor, and Ignition
- Suppose the observable symptoms are limited to: Headlights work/don't work; Engine starts/doesn't start; Wipers work/don't work
- In this simplified setting, if the engine does not start, we can conclude that either the battery or the ignition are to blame.
- Real-world settings can involve hundreds of possible observable symptoms involving a large number of components
- While this increases the complexity of the problem, it does not necessarily impact the overall approach to the problem

How can we design an expert system for auto diagnosis?

- Develop a set of rules that capture the mechanic's knowledge
- Knowledge engineers pick the brains of experts to elucidate the knowledge needed

Rule1 : IF Headlights don't work THEN faulty Bulbs or/and Battery
Rule 2: IF faulty Bulbs or/and Battery AND Engine does not start THEN faulty Battery
Rule 3: IF faulty Bulbs or/and Battery AND Engine starts THEN faulty Bulbs

How does an expert system work?

- A knowledge base encodes relevant knowledge, e.g., rules of the form IF antecedents then consequences
- An inference engine draws conclusions by applying the rules in the knowledge base to the assertions supplied by the user.
- A working memory is used as a scratch pad to store user-supplied assertions and any assertions derived during the inference process.

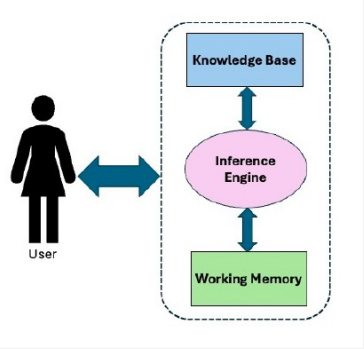

How does an expert system work?

- During a consultation with an expert system the user enters the facts relevant for the problem to be solved into the working memory.
- The system matches the user-supplied facts with the
- Knowledge contained in the knowledge base to infer additional facts which are entered into the working memory
- The process continues until the system either completes
- the task presented to it or runs into a dead end.

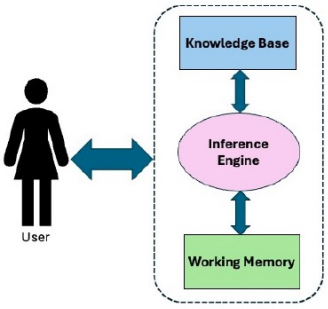

How does an expert system work?

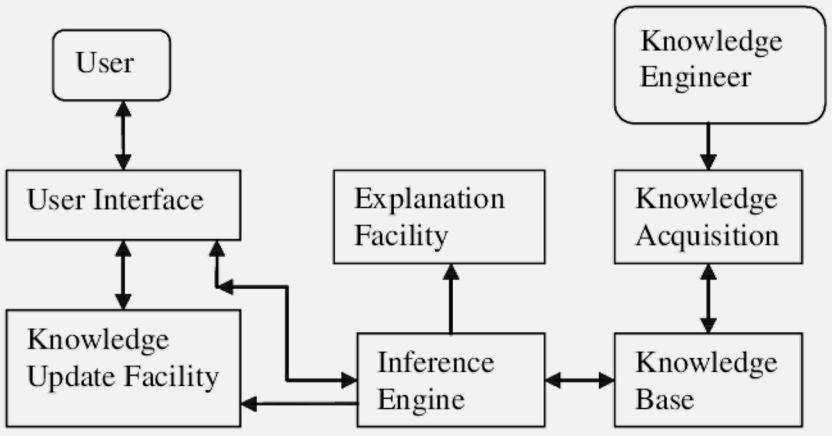

Expert auto mechanic

Rule1 : IF Headlights don't work THEN faulty Bulbs or/and Battery

Rule 2: IF faulty Bulbs or/and Battery AND Engine does not start THEN faulty Battery

Rule 3: IF faulty Bulbs or/and Battery AND Engine starts THEN faulty Bulbs

Engine starts?

Engine starts

Headlights don't work

Headlights don't work
Faulty Bulbs and/or Battery

Rule 1

Headlights don't work
Faulty Bulbs and/or Battery
Engine starts

Rule3

Headlights don't work
Faulty Bulbs and/or Battery
Engine starts
Faulty Bulbs

Rule-based systems

- Rule based systems use IF-THEN rules
  - IF – antecedents
  - THEN – consequent

Forward chaining

- If all the antecedents match the working memory, then the rule is triggered
- A triggered rule is allowed to fire (based on some prioritization)
- When a rule fires, the consequents are added to the working memory
- The process repeats

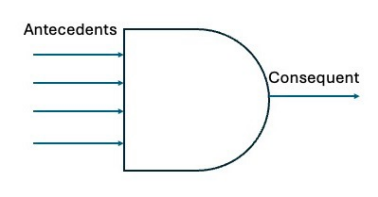

Rule-based systems

- Deductive systems
  - Rule firing simply adds new assertions to the working memory
- Reactive systems
  - Rule firing results in an action
  - Action may result in some working memory assertions being deleted and new ones added

Expert system for managing a fleet of trucks

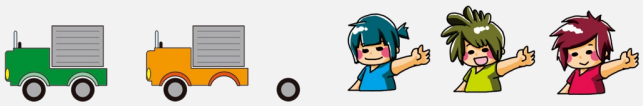

Rules

- Fix a truck if broken and parts are available
- Assign driver to truck when truck is without a driver and a driver is available

Expert system for managing a fleet of trucks

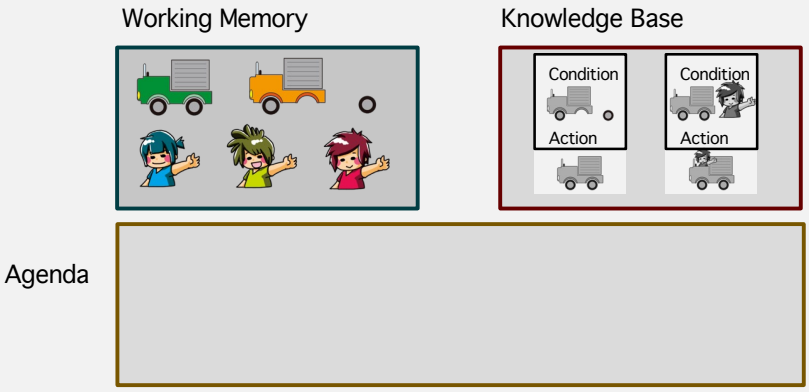

Expert system for managing a fleet of trucks

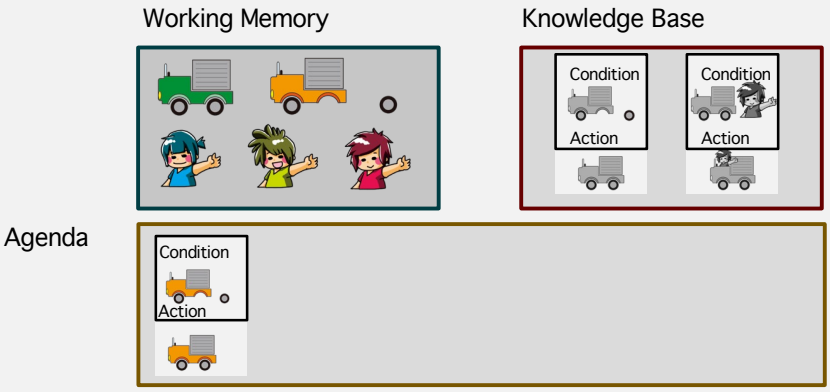

Expert system for managing a fleet of trucks

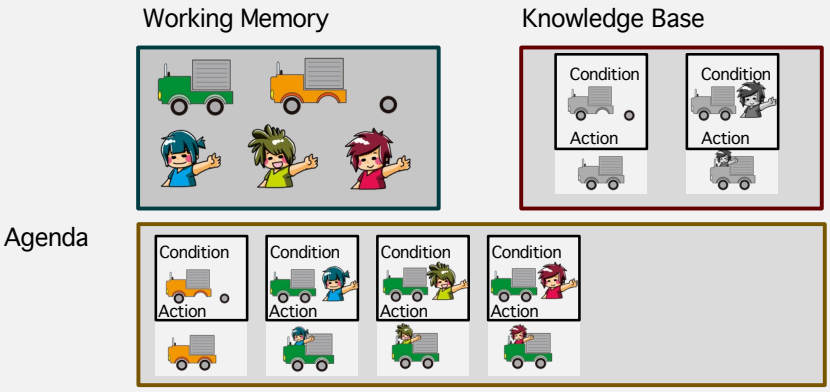

Expert system for managing a fleet of trucks

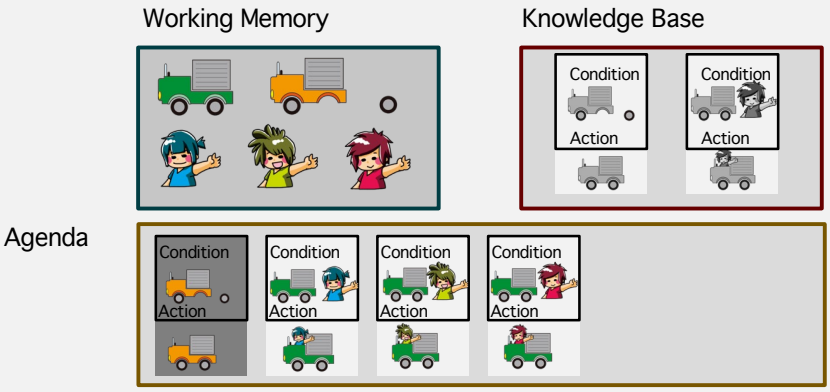

Expert system for managing a fleet of trucks

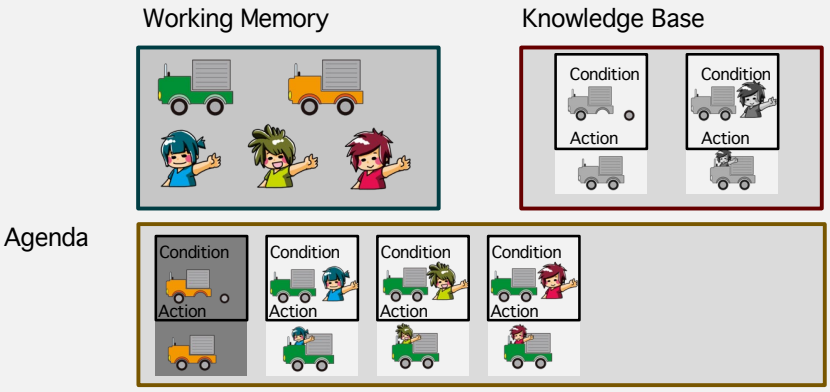

Expert system for managing a fleet of trucks

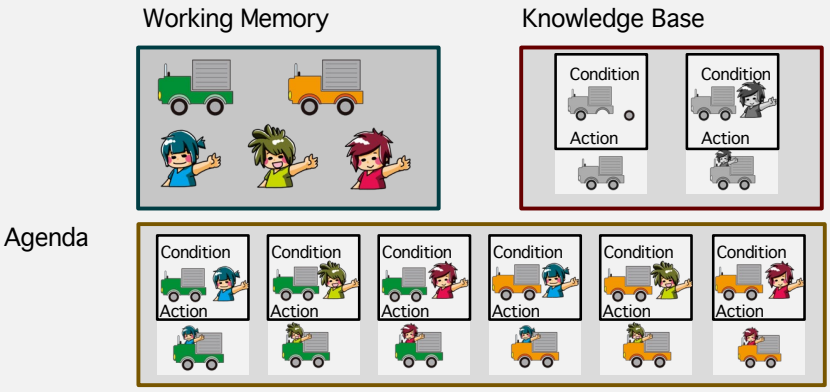

Expert system for managing a fleet of trucks

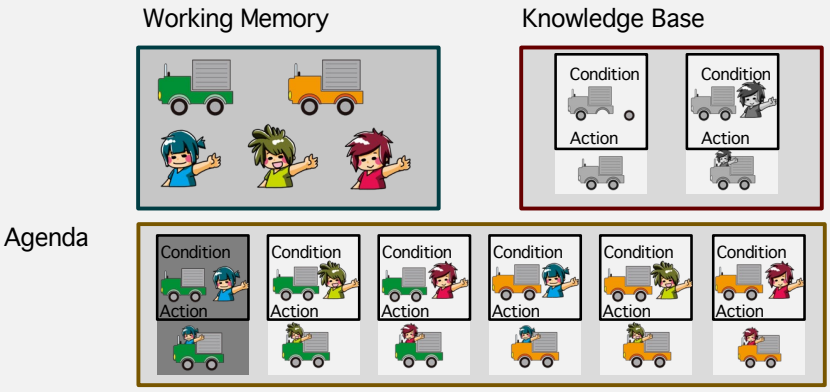

Expert system for managing a fleet of trucks

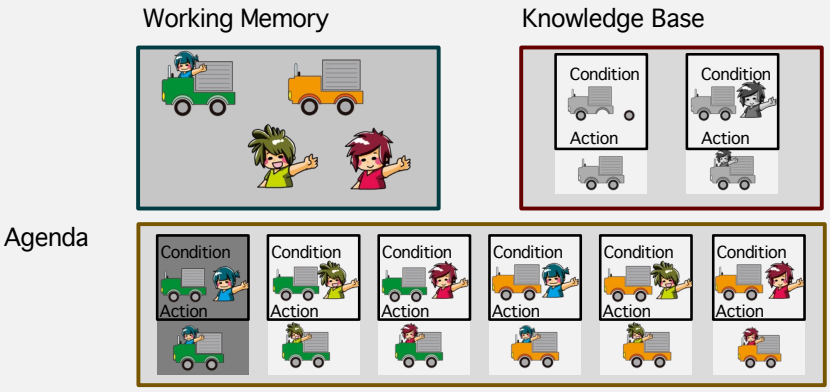

Expert system for managing a fleet of trucks

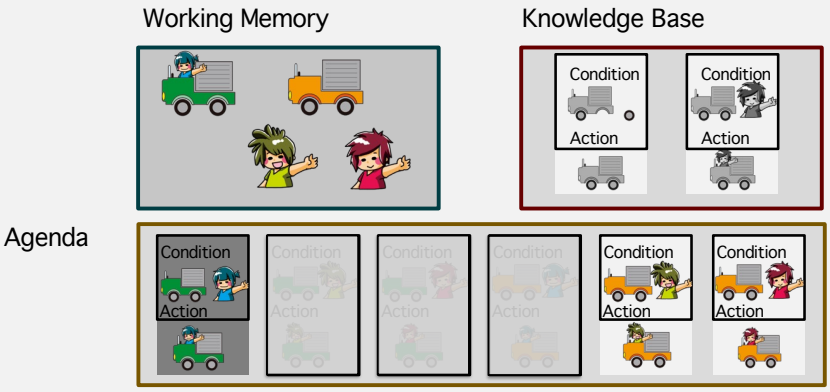

Expert system for managing a fleet of trucks

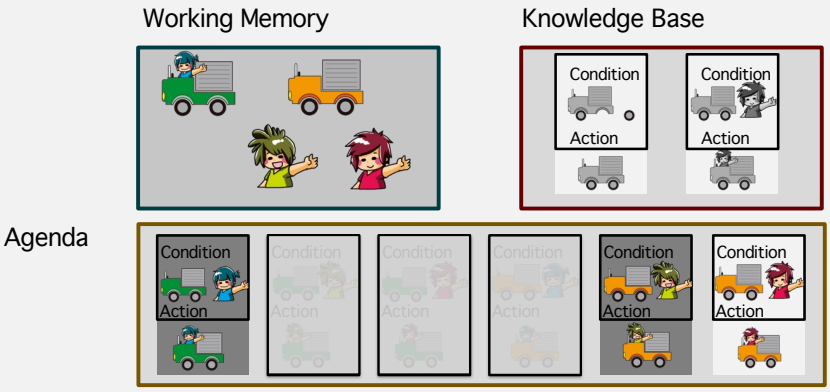

Expert system for managing a fleet of trucks

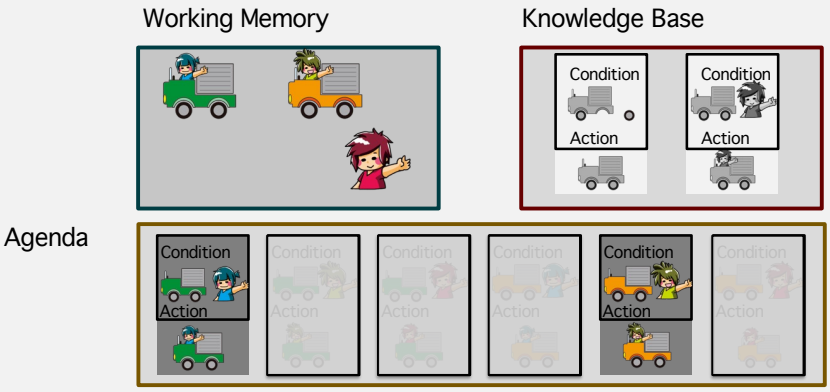

Expert system for managing a fleet of trucks

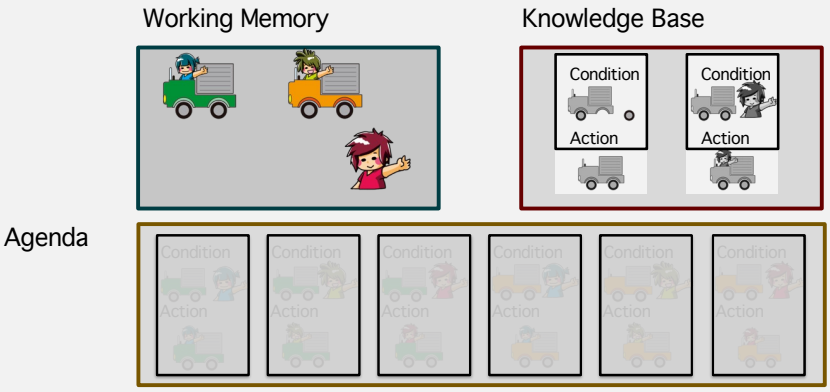

Rob, the robot Zoo-Guide: Rules for recognizing zoo animals

- Z_1: IF ?x has hair THEN ?x is a mammal
  - ?x is a variable that can be instantiated through a match with working memory
  - Suppose working memory contains Maya has hair
  - Then the antecedent of rule Z_1 match the working memory
  - Upon firing of rule Z_1 Maya is a mammal is added to the working memory
- When multiple rules match, they get on an Agenda with some mechanism for prioritization among the different matching rules

Rob, the robot Zoo-Guide: Rules for recognizing zoo animals

- Z_2: IF ? x gives milk THEN ? x is a mammal
- Z_3 IF ? x has feathers THEN ? x is a bird
- Z_4 : IF ? x flies AND ? x lays eggs THEN ? x is a bird
- Z_5: IF ? x is a mammal AND ? x eats meat THEN ? x is a carnivore
- Z_6 : IF ? x is a mammal AND ? x has pointed teeth AND ? x has claws AND ? x has forward-pointing eyes THEN ? x is a carnivore
- Z_7 : IF ? x is a mammal AND ? x has hoofs THEN ? x is an ungulate
- Z_8 : IF ? x is a mammal AND ? x chews cud THEN ? x is an ungulate
- Z_9 : IF ? x is a carnivore AND ? x has tan color AND ? x has dark spots THEN ? x is a cheetah
- Z_{10} : IF ? x is a carnivore AND ? x has tan color AND ? x has black stripes THEN ? x is a tiger
- Z_{11} : IF ? x is an ungulate AND ? x has long legs AND ? x has long neck AND ? x has tan color AND ? x has dark spots THEN ? x is a giraffe
- ..

Rob, the robot Zoo-Guide: Forward chaining

- Maya has hair.
- Maya chews cud.
- Maya has long legs.
- Maya has a long neck.
- Maya has tan color.
- Maya has dark spots.
- Because Maya has hair, rule Z_1 fires, and asserts that Maya is a mammal.
- Because Maya is a mammal and chews cud, rule Z_8 fires and asserts that Maya is an ungulate.
- At this point, all the antecedents for rule Z_{11} match the working memory, allowing Z_{11} to fire and inform Rob that Maya is a giraffe.

Rob, the robot Zoo-Guide: Forward Chaining

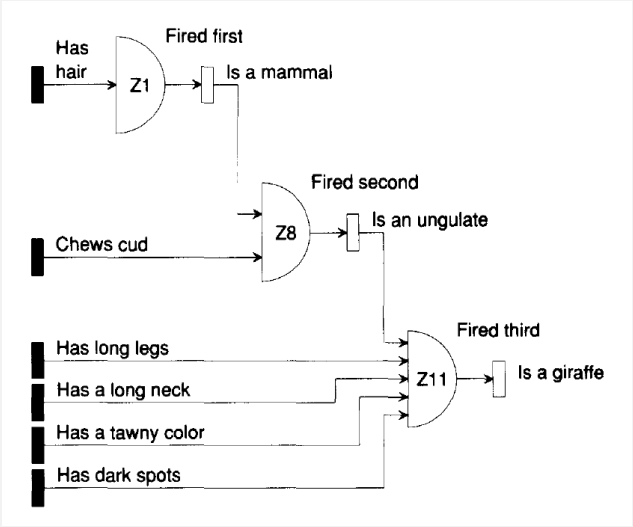

Forward chaining proceeds from facts to conclusions

Rob, the robot Zoo-Guide: Forward Chaining

- Until no rule produces a new assertion or the animal is identified,
  - For each rule,
    - Try to support each of the rule's antecedents by matching it to known facts.
    - If all the rule's antecedents are satisfied, assert the consequent unless there is an identical assertion already.
    - Repeat for all alternative matches.

Rob, the robot Zoo-Guide: Backward Chaining

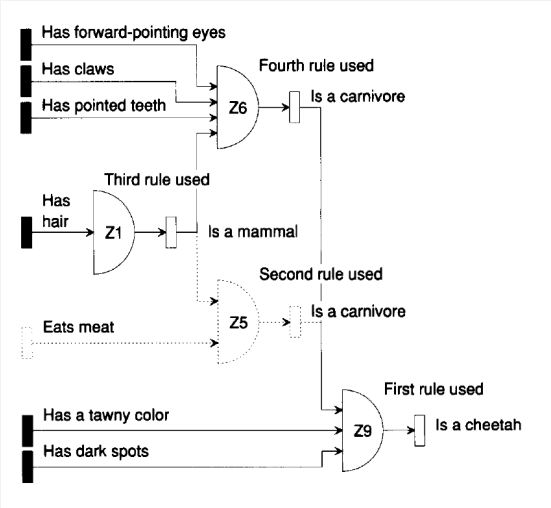

Backward chaining proceeds from hypothesis to facts

Rob, the robot Zoo-Guide: Backward Chaining

- For each hypothesis,
  - For each rule whose consequent matches the current hypothesis,
    - Try to support each of the rule's antecedents by matching it to assertions in working memory or by backward chaining through another rule, creating new hypotheses. Be sure to check all matching and instantiation alternatives.
    - If all the rule's antecedents are supported, announce success and conclude that the hypothesis is true.

To Forward chain or backward chain?

- Whenever the rules are such that a typical set of assertions can lead to many conclusions, the rule system has a high degree of fan-out which argues for backward chaining.
- On the other hand, whenever the rules are such that a typical hypothesis can lead to many questions, rule system has a high degree of fan-in which argues for forward chaining.

To Forward chain or backward chain?

- If the assertions that you have or may establish can lead to a large number of conclusions, but the number of ways to reach the particular conclusion in which you are interested is small, then there is more fan-out than fan-in, which argues for backward chaining
- If the number of ways to reach the particular conclusion in which you are interested is large, but the number of conclusions that you are likely to reach using the assertions is small, then there is more fan-in than fan-out, which argues for forward chaining.

To Forward chain or backward chain?

- In many situations, neither fan-out nor fan-in dominates
- If you have not yet gathered any facts, and you are interested in only whether one of many possible conclusions is true, use backward chaining.
  - Example: You do not care about the identity of an animal. All you care about is whether it is a carnivore.
- If you already have in hand all the facts you are ever going to get, and you want to know everything you can conclude from those facts, use forward chaining.
  - Example: You have just seen an animal run by. You want to know what you can deduce about it.

Explaining inferences: Forward chaining

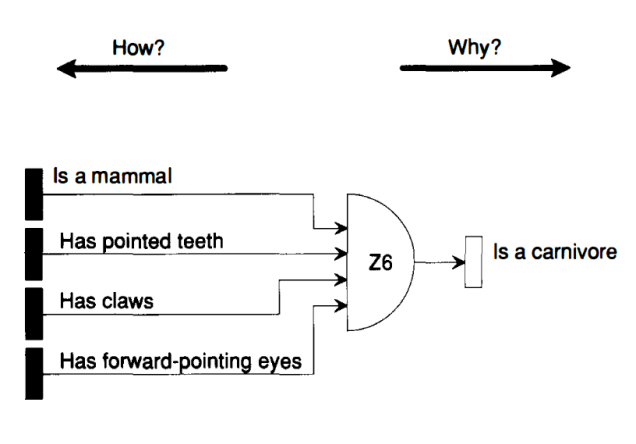

Explaining inferences

- If the question is a how question, report the assertions connected to the IF side of the rule that established the assertion referenced in the question.
- If the question is a why question, report the assertions connected to the THEN sides of all rules that used the assertion referenced in the question

Explaining inferences: Backward chaining

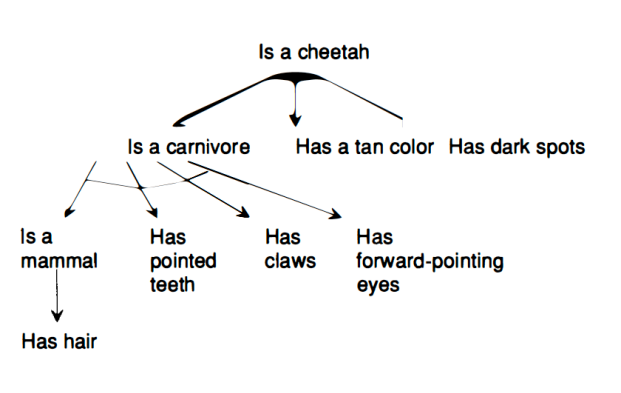

Making reasoning efficient: RETE network

- The naïve approach to rule matching repeats work because IF conditions can be shared by multiple rules
- This can be avoided by compiling the rules into a network so that once a particular IF condition is matched, the match is shared with all of the rules that depend on it
- A rule is triggered when all of its IF conditions are matched

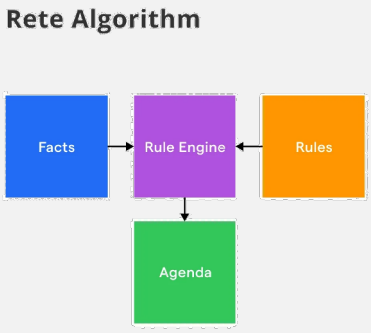

Making reasoning efficient: RETE network

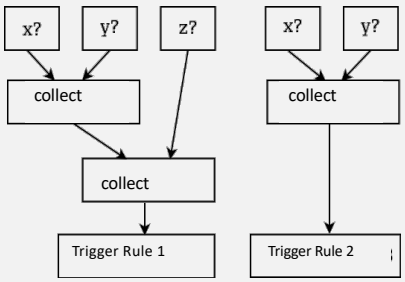

The naïve approach to rule matching repeats work because IF conditions can be shared by multiple rules

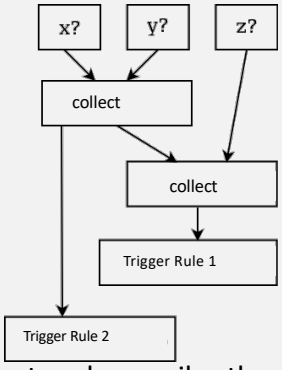

Rete network compiles the rules so that once a particular IF condition is matched, the match is shared with all of the rules that depend on it

Expert systems in practice: R1/XCON

- Expert systems were among the first commercial successes of AI
  - R1 – an expert system designed by John McDermott of CMU in collaboration with Digital Equipment Corporation (DEC) later named XCON (Expert configurer)
  - Configuring computer system orders
  - Started with about 750 rules in 1982 and grew to 2500 rules by 1984
  - Had processed processed 80,000 orders, and achieved 95–98 percent accuracy by 1986
  - Saved DEC between 10 and 40 million dollars per year

Expert systems in practice: MYCIN

- MYCIN – an expert system for clinical consultation about diagnosis and treatment of bacterial infections designed by Ted Shortliffe of Stanford
- MYCIN informs itself about patients by asking questions symptoms, history, and laboratory-test results
- Question asked by MYCIN is determined by MYCIN's current hypothesis and the answers it has received so far
- MYCIN could also explain its reasoning in simplified English.
- MYCIN matched the performance of human specialists in blood infections and outperformed general practitioners.
- MYCIN never saw deployment in actual clinical setting due to legal and regulatory concerns

Expert systems in practice

- According to the results of a survey by Durkin, as of 1996, more than 2500 expert systems were deployed.
- Durkin speculated that this figure represents only about a fifth of systems that might have been developed
- Anecdotally, rule-based expert systems are widely used in industrial applications

Expert systems in practice

Knowledge Engineering Bottleneck

- The experts providing the domain knowledge must be able to express their knowledge in the form of simple rules.
- The knowledge engineer working with the experts should be able to grasp the expert's knowledge well enough to encode it in the knowledge base
- Knowledge engineering bottleneck is responsible for the exorbitant cost and effort involved in developing expert systems for real-world applications.
- Machine learning offers a way out -- by giving expert systems the ability to acquire the knowledge needed by observing experts or interacting with the world

Engineering challenge

- The interaction between the knowledge base, the inference mechanism, the conflict resolution strategies and the order in which facts enter the working memory can be complex