

Chapter 7

Technologies to Manage Knowledge: Artificial Intelligence

Section 7.1 - Chapter Objectives

- Introduce artificial intelligence as a facilitating technology for knowledge management
- Introduce knowledge as an important facet of intelligent behavior
- Introduce the early state space search techniques
- Introduce expertise in the context of knowledge
- Introduce knowledge-based systems as a modern evolution of the early state space search techniques

Section 7.2 - Objectives

- Define Artificial Intelligence (AI) as the science that “... *encompasses computational techniques for performing tasks that apparently require intelligence when performed by humans.*”
 - ◆ Turing Test
- Provide a short historical summary of the most significant events and systems. This places artificial intelligence in the context of other significant advances in information technology.

Knowledge vs Intelligence

- People go to college to acquire knowledge and reasoning ability (*not* intelligence)
- People are hired for their ability to solve problems intelligently
- AI systems aim to mimic intelligent problem solving
- Problem-solving requires
 - ◆ Knowledge
 - ◆ Inference (ability to manipulate, acquire, and manage knowledge effectively and efficiently – for recognition, reasoning, learning, etc.)
- Key questions:
 - ◆ How can inference best be implemented in AI models, so as to manipulate, acquire, and manage knowledge most effectively and efficiently?
 - ◆ How can knowledge best be represented in AI models, so as to facilitate effective and efficient inference?
 - ◆ Is an inference model itself a kind of knowledge?

AI history = CS history

- Digital computers, information theory, and computer science were originally developed (by pioneers like Turing, Von Neumann, Shannon, Chomsky, and so on) primarily for the purpose of mimicking human processing of language:
 - ◆ cryptography
 - ◆ machine translation
 - ◆ formal language theory
 - ◆ ...
- Game playing was the other early area of AI modeling:
 - ◆ checkers
 - ◆ chess
 - ◆ backgammon
 - ◆ ...

AI history = CS history

- Machine translation (more generally, natural language processing) has proved more challenging than game playing
 - ◆ Modern game playing models can beat even world champions at checkers, chess, backgammon, ...
 - ◆ But a 3-year old child can process language far better than any AI model
- The two areas have different emphases
 - ◆ Game playing more crucially depends on *inference* mechanisms (and typically use relatively simple knowledge representations)
 - ◆ Language processing more crucially depends on how to represent *knowledge*
- Both are needed!
- So the parallel research efforts in these two areas since the 1940s have, in effect, taken turns driving progress in AI

Figure 7.1

Artificial Intelligence

Knowledge-based Systems

- Rule-based systems
 - Classification
 - Diagnosis
 - Design
 - Decision support
 - Planning
 - Scheduling
- Case-based Reasoning
 - Diagnostics
 - Design
 - Decision support
 - Classification
- Constraint-based reasoning
 - Planning
 - Scheduling
- Model-based reasoning
 - Monitoring
 - Diagnostics
 - Design

Natural Language Processing

- Machine translation
- NL understanding
- NL generation
- Speech understanding
- Speech synthesis


Computer Vision

- Image processing
- Image understanding

Machine Learning

- Inductive learning
- Case-based learning
- Connectionist learning
- Learning from analogy
- Explanation-based learning.
- Data mining
- Others.


Figure 7.1 (cont.)



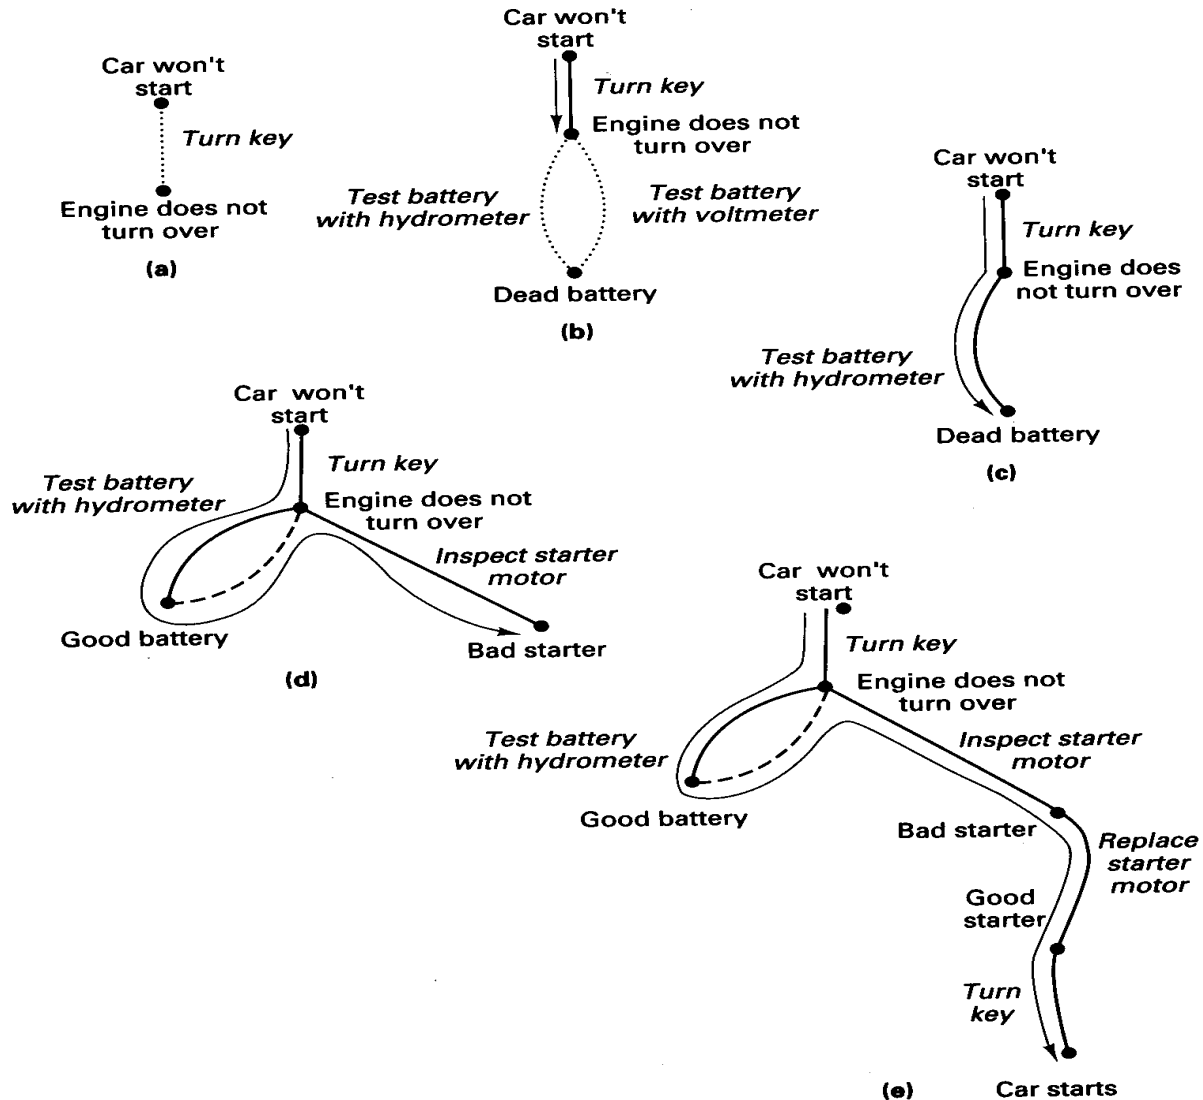
Section 7.3 - Objectives

- Introduce the early approaches to artificial intelligence - the solution space or state space search
- Explain the nature of the knowledge found in state space searches as being general
- Explain the advent of the heuristic function as a way to expedite the state space search
- Present two vignettes as examples
- Conclude that the general knowledge employed in state space searches was not sufficient to solve the difficult problems

Representation of a portion of a solution space for automobile diagnosis



Searching a portion of the solution space



The problem with “dumb search”

- Although the concept of searching a solution space or state space is an early approach to AI, it remains the foundation way to view processing in all sorts of more sophisticated AI models
- “Dumb search” within solution spaces uses only general domain-independent knowledge – no domain-specific knowledge
- Too inefficient to be practical – how can we improve search to solve more realistic difficult problems?

Section 7.4 - Objectives

- Briefly introduce modern knowledge-based systems
- Introduce modern knowledge-based systems in the context of the state space search methods to understand their advantages and disadvantages
- Uses several vignettes to describe the difference between the different approaches
- Provides a transition to the more detailed contents of Chapter 8

Heuristic search

- Consider trying to get from an Origin Point (say, HKUST) to a Destination Point (say, the Peak Galleria)
- Possible search strategies:
 - ◆ Random search
 - ◆ Systematic blind search
 - ◆ Heuristic search (or directed search)
 - Let search be directed by a heuristic function
 - A heuristic is a “rule of thumb” that tries to guess which search directions are more promising
 - E.g.: at any choice point where you can see the Destination Point, choose the road that is aimed most closely at that direction
 - Heuristics may not always produce the optimum choice!
 - Here, the heuristics capture only very general strategies – general domain-independent knowledge

Heuristic search in Knowledge-Based Systems

- Consider this:
 - ◆ You notice a backfiring noise in your car
 - ◆ You remember a similar situation several months earlier
 - The mechanic at that time told you the problem was a loose vacuum hose connection in the engine (and showed you the hose)
 - ◆ You stop the car, discover the hose is loose, reconnect it
 - ◆ You restart the car, notice the problem is gone
- You:
 - ◆ Did not methodically analyze the operation of the engine
 - ◆ Do not know how an internal combustion engine operates
 - ◆ Did not have a solution space in mind, did not search a solution space
- You used heuristic knowledge that:
 - ◆ Was based on prior experience
 - ◆ Is domain-specific
 - ◆ Captured associational expertise
- Good domain-specific heuristic rules are drawn from experience.

Heuristic search

- Heuristics are useful when:
 - ◆ The number of possibilities to be examined is too large,
 - ◆ The exact evaluation function applied to each possible answer to determine correctness is too complex, or
 - ◆ The exact evaluation function is unknown and must be approximated

Heuristic search

- Consider this:
 - ◆ You need the contractor to give you an estimate to build a house before you leave town tonight
 - ◆ The contractor typically prices by doing a detailed cost estimate:
 - Call a building supply warehouse for material prices
 - Evaluate quotations from subcontractors on cost of labor
 - Determine appropriate contractors' fees
 - Allow a reasonable contingency figure
 - ...
 - ◆ This takes too much time to finish by tonight!
 - ◆ The contractor instead takes a heuristic approach:
 - Find another home of similar size that he recently built
 - Normalize for the exact size
 - Adjust for large differences
 - Number of bathrooms
 - Luxury kitchen fittings
 - ◆ Takes only 30 minutes

Knowledge-Based Systems for KM

- Knowledge-based systems are an excellent platform for capturing, sharing, and applying knowledge (of certain kinds).
- Knowledge-based systems were designed primarily for the purpose of being able to apply knowledge automatically.
- In many KM contexts, we just want a tool to support knowledge capture, discovery, and/or sharing – we may not need (or desire) automatic application.
 - ◆ Even then, the ability of knowledge-based systems to apply knowledge can be useful for knowledge sharing:
 - Demonstrate the system on example problems
 - Have the system indicate what heuristic knowledge was used to solve the problems
 - ◆ An excellent way to help humans internalize the knowledge.

Section 7.5 - Objectives

- Provide a historical view of knowledge-based systems juxtaposed to the historical discussion of AI done earlier in this chapter
- Present the basic concepts of a modern knowledge-based system and how MYCIN pioneered that approach
- Presents a list of legacy knowledge-based systems that pioneered advances in the field

Knowledge Based Systems – more precise definitions

- A knowledge-based system is “a computerized system that uses domain knowledge to arrive at a solution to a problem within that domain. This solution is essentially the same as one concluded by a person knowledgeable about the domain, when confronted with the same problem.”
- But this definition isn't strict enough - too many conventional systems could be described this way!
- Key differences from conventional software:
 - ◆ The use of highly specific domain knowledge.
 - ◆ The heuristic nature of the knowledge employed, instead of exact.
 - ◆ The separation of the knowledge from how it is used.

The use of highly specific domain knowledge

- Initially pioneered in the late 1960s to early 1970s in the DENDRAL and Meta-DENDRAL systems [Lindsay 1980]
- DENDRAL infers the molecular structure of unknown compounds from mass spectral and nuclear magnetic resonance data
- Meta-DENDRAL assists in the determination of the dependence of mass spectrometric fragmentation on substructural features

The separation of the knowledge from how it is used

- Note there are two kinds of knowledge involved in automatically making inferences:
 - ◆ The general knowledge of how to infer something (anything): general inference methods
 - ◆ The specific knowledge about that something
- Emphasizes that a declarative (rather than procedural) representation should be used for the specific knowledge.
- Allows general inference engines to be developed
 - ◆ Can be re-used in any domain
 - ◆ Simply encode domain-specific knowledge using the declarative representation language, and then run the general inference engine
- E.g., the CLIPS rule-based system developed at NASA in the early 1980s

Knowledge-Based Systems – MYCIN

- These three key differences from conventional software were first merged in Stanford's MYCIN system in the early 1970s [Shortliffe 1976; Buchanan 1984]:
 - ◆ The use of highly specific domain knowledge.
 - ◆ The heuristic nature of the knowledge employed, instead of exact.
 - ◆ The separation of the knowledge from how it is used.
- MYCIN diagnoses and specifies treatments for blood disorders through a Q&A session with a physician:
 - ◆ asks questions about the signs and symptoms of the patient
 - ◆ requests certain laboratory tests as appropriate
 - ◆ recommends a drug treatment after the set of possible infections has been sufficiently narrowed

Knowledge-Based Systems – some other examples

- Medicine: CASNET, INTERNIST, PUFF, TEIRESIAS
- Understanding natural language: HEARSAY
- Geology: PROSPECTOR
 - ◆ SRI [Duda 1978]
 - ◆ Elicited, preserved, and applied geologic formation knowledge to assist in mineral exploration
- Manufacturing: XCON
 - ◆ DEC & CMU [McDermott 1982]
 - ◆ Elicited, preserved, and applied the knowledge of human configurators of computer systems to automate and duplicate their functions
 - ◆ One of the earliest commercially successful systems
- Manufacturing: COOKER
 - ◆ Texas Instruments for Campbell Soup [AIInteractions 1985]
 - ◆ Assists in the maintenance of soup-making equipment
 - ◆ Captured and preserved the knowledge of a highly experienced employee about to retire
- Credit: AUTHORIZER'S ASSISTANT
 - ◆ Inference Corp & American Express [Leonard-Barton 1988]
 - ◆ Elicited, preserved, and applied human knowledge in handling applications for AmEx credit cards.
 - ◆ Takes information from multiple databases and issues approval/denial of large purchase requests from merchants

Section 7.6 - Objectives

- Distinguish among the various types of knowledge
- Establish a distinction between knowledge and expertise

What kinds of knowledge are in Knowledge-Based Systems?

- Consider the earlier car engine example
- General support simple knowledge (known to all mechanics):
 - ◆ Can test battery (in two different ways), inspect cables, inspect starter switch, inspect starter motor
- Specific simple knowledge (known to only the mechanic working on the particular car):
 - ◆ The outcome of each test or inspection for that particular car
- Specific tactical or strategic complex knowledge (known only to expert mechanics):
 - ◆ What order should the inspections most efficiently/effectively be done in, in all possible situations for all cars

What kinds of knowledge are in Knowledge-Based Systems?

- Consider the game of chess
- Knowledge about the rules of the game
 - ◆ general, simple, support domain knowledge
 - ◆ which moves *can* be made
- Knowledge of how best to move the pieces to defeat the opponent
 - ◆ specific, complex, tactical and strategic
 - ◆ may be based on experience
 - ◆ which moves *should* be made

Knowledge Based Systems vs Expert Systems

- There exist knowledgeable individuals who are not considered experts
 - ◆ Knowledge can be specific without being at expert level
 - ◆ Having the knowledge is not the same thing as being able to apply it effectively
- Expert systems are knowledge-based systems where:
 - ◆ the specific knowledge is at the level equivalent to a human expert, and
 - ◆ the inference engine is able to make use of the knowledge as efficiently and effectively as a human expert
- All expert systems are knowledge-based systems, but not vice versa

What kinds of expertise are in Expert Systems?

- Knowledge-based systems are excellent at representing many forms of associational expertise.
 - ◆ Heuristic ability/knowledge acquired mostly through eliciting from humans via a knowledge engineering process (discussed later in the course)
 - ◆ Typically (but not necessarily) represented as rules in rule-based systems
- Knowledge-based systems are poor at representing most forms of motor skills expertise.
 - ◆ Motor skills tend to involve tacit, physical (not cognitive) abilities
 - ◆ Progress is being made (consider the RoboCup competition)
- Knowledge-based systems can represent limited forms of theoretical expertise.
 - ◆ But more model-based reasoning systems

Section 7.7 - Objectives

- Introduce the advantages of knowledge-based systems
- Introduce the disadvantages of knowledge-based systems

Advantages of Knowledge-Based Systems

- Wide distribution of scarce expertise
 - ◆ especially for direction
- Ease of modification
- Consistency of answers
- 24/7 accessibility
- Preservation of expertise
- Solution of problems involving incomplete data
- Explanation of solution
 - ◆ Especially for sharing/internalization

Disadvantages of Knowledge-Based Systems

- Answers may not always be correct
- Limits not always recognized
- Lack of common sense
- Restricted scalability
 - ◆ systems often become too complex
 - ◆ Typically infeasible to maintain correct rule interaction when the number of rules grows into the thousands

Section 7.8 - Objectives

- Introduce briefly other types of AI reasoning as an alternative to rule-based reasoning:
 - ◆ Model-based reasoning
 - ◆ Constraint-based reasoning
 - ◆ Diagrammatic Reasoning
 - ◆ Fuzzy logic
 - ◆ Evolutionary algorithms

Table 7.1

TABLE 7.1 FUZZY SETS TALL, STATUESQUE, SHORT, AND GIANT

Tall		Statuesque		Short		NBA Players	
5'0"	0.00	5'0"	0.00	5'0"	1.00	5'0"	0.00
5'4"	0.08	5'4"	0.08	5'4"	0.92	5'4"	0.04
5'8"	0.32	5'8"	0.32	5'8"	0.68	5'8"	0.08
6'0"	0.50	6'0"	0.50	6'0"	0.50	6'0"	0.18
6'4"	0.82	6'4"	0.82	6'4"	0.18	6'4"	0.32
6'8"	0.98	6'8"	0.98	6'8"	0.02	6'8"	0.50
7'0"	1.00	7'0"	1.00	7'0"	0.00	7'0"	0.75

Section 7.9 - Objectives

- Summarize the chapter
- Provide Key terms
- Provide Review Questions
- Provide Review Exercises

Conclusions

- The student should be familiar with:
 - ◆ The concept of expertise in the context of knowledge
 - ◆ The state space search methods comprising early AI work
 - ◆ The difference between these and the modern knowledge-based systems
 - ◆ How knowledge-based systems can be used to manage knowledge.
 - ◆ The difference between forward and backward reasoning, and when one or the other should be used.

Individual Assignment

(Due in class Jul 24)

1. How would you draw the solution space for the game of tic-tac-toe? Do it for the first three moves.
2. Consider a jigsaw puzzle. Identify a set of heuristics (say, 3 or 4) that can be used in assembling the puzzle.

Chapter 7

Technologies to Manage Knowledge: Artificial Intelligence