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

Knowledge Representation

- People use mental models or representations to make sense of their world
  - Inflation
    - One may blame inflation on supply chain disruptions due to Covid
    - Another may blame poor decisions by the Federal Reserve

Knowledge Representation

- People use mental models or representations to make sense of their world
  - David Letterman on Soup Nazi
    - One gets the joke because she is an avid watcher of Seinfeld reruns
    - Another does not get the joke because he is oblivious to pop culture

Knowledge representation

- Even the simplest of agents, to behave intelligently, need
  - Knowledge of their environments
- Thermostat needs to know
  - To turn the heat on if the room is too cold
    - Turning the heat on would make the room warmer
  - To turn the AC on if the room is too hot
    - Turning the AC on would make the room colder
  - How occupancy impacts the decisions

Knowledge representation

- What is knowledge?
- How can knowledge be encoded in a form that can be used by a machine?
- How does the form in which the knowledge is encode impact, independently of content, the behavior of agents?
- What kinds of knowledge are there?
- How can an agent effectively use its knowledge to act in its environment?
- How can an agent communicate what it knows to other agents and humans?
- How can an agent acquire knowledge?

Example: Finding the shortest route from A to F

- How can you find the shortest route?

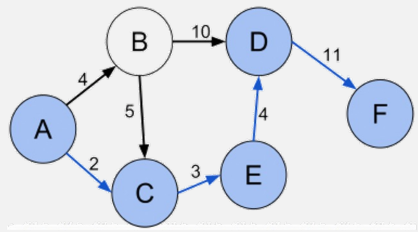

Example: Finding the shortest route from A to F

- How can you find the shortest route?
- Nodes represent locations
- Directed edges denote roads
- Numbers on edges denote distances
- Total distance along a path is the sum of the distances along edges

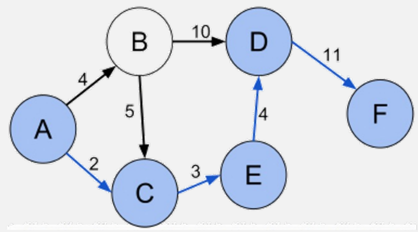

Example: Finding the shortest route from A to F

- Total distance along a path is the sum of the distances along edges
- What is the shortest path from A to F?

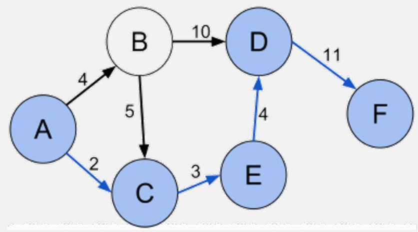

Example: Finding the shortest route from A to F

- What is the shortest path from A to F?
- A → C → E → D → F
- Why?

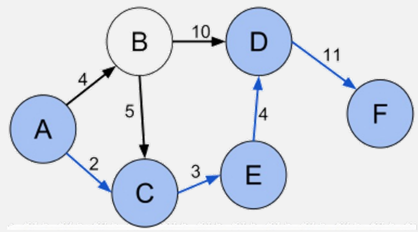

Example: Finding the shortest route from A to F

- Is there a systematic algorithm for solving shortest path problems?

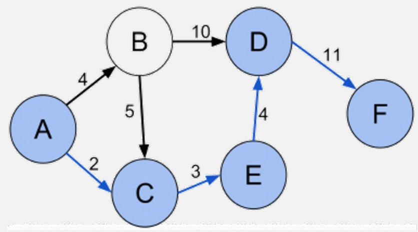

Example: Finding the shortest route from A to F

- Is there a systematic algorithm for solving shortest path problems?
- Yes!

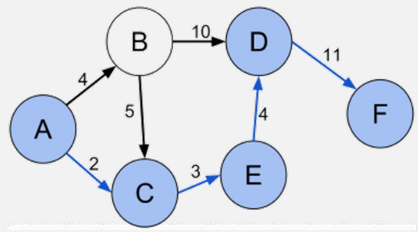

Example: Finding the shortest route from A to F

1. Generate 1-hop paths from the start location A to its neighbors and order them by their cost. AC (2), AB (4).
2. If the first path on the list ends in the destination, you have the solution.
3. If not, extend the paths on our list by one more step, and replace the list by resulting 2-hop paths, ACE (5), ABC (9
4. Go back to step 2 and repeat.

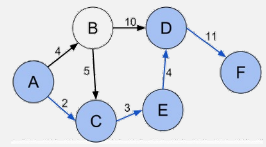

Example: Finding the shortest route from A to F

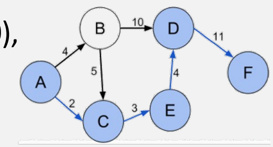

- Eventually we end up with ACEDF (20), ABDF (25), ABCEDF (27).
- Now the first path on the list takes us from A to F.
- Because the list is sorted in increasing order of cost, the cost of ACEDF is cheaper than all other partial or complete paths on the list.

- Note that this algorithm is kind of obvious once you have the graph representation

Role of representation in humans and machines

- How did we solve the shortest path problem?
- We used a representation that
  - Made the relevant aspects explicit
  - Discarded the irrelevant aspects – population of the cities, economy of the cities, weather, what the driver had for breakfast…
  - If we can use such a representation to solve problems, so can a machine
- AI is about creating and manipulating representations of the world to solve problems, to exhibit intelligent behavior

Form versus content of representation

- Suppose you want to add two numbers between 0 and 10,000
  - Solution 1: Use a lookup table that stores sums of all possible pairs of numbers between 0 and 10,000
    - Adding the numbers 5 and 10 will take a single table lookup
  - Solution 2: Add numbers using a counter that starts at 0 and is successively incremented by the two numbers.
    - Adding the numbers 5 and 10 will take 15 increments of the counter starting at 0
- Both solutions are functionally equivalent
  - For any valid inputs (pairs of numbers between 0 and 10,000), they produce the same results
- But they differ greatly in their efficiency!
  - Solution 1 is faster, but requires a large memory
  - Solution 2 is slower, but requires little memory

Representation is a surrogate for the world

- A knowledge representation is fundamentally a surrogate for
  - Individuals
  - Objects
  - Properties
  - Relationships
  - Processes
  - Actions
- in the world
- Representation enables an agent to substitute thinking for acting in the world

Substituting thinking for acting in the world

- Suppose an agent that knows that
  - the coffee that was just poured into a cup that it is holding is hot
  - if one were to drink hot coffee, it will burn one's tongue
- The agent can can infer the coffee that that was just poured into the cup it is holding, if it were to drink it, will burn its tongue
- Note that the agent figured out the effect of its action
  - By reasoning with what it knows
  - And not actually performing the action – drinking hot coffee and getting its tongue burned

Why substitute thinking for interacting with the world?

- Acting in the world can be costly
  - Imagine assembling furniture you bought from Ikea
  - You better have a plan for assembling the furniture before you start
  - Mistakes can be costly
- Acting in the world can be dangerous
  - Imagine the scenario with hot coffee
  - Imagine a self-driving taxi approaching an intersection

Representation as a surrogate for the world

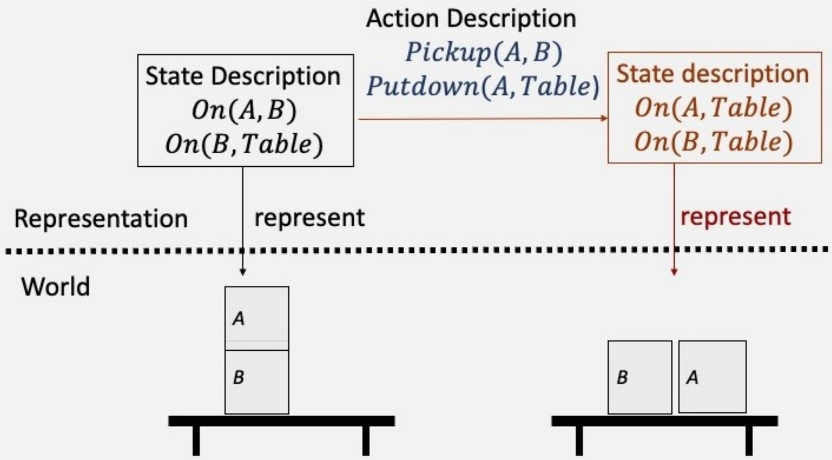

How are surrogates useful?

- Representation offers descriptions of the world states and actions
- The descriptions can be manipulated by an algorithm
- The results of running the algorithm give us the consequences of actions in the world
- For a representation to be useful, there must be a coupling or correspondence between descriptions and their physical counterparts

What makes a representation meaningful?

- The correspondence between descriptions and objects, actions, individuals .. in the world makes the representation meaningful to the agent
- What is an ``apple''?
- ``apple'' has no intrinsic meaning
- The use of ``apple'' to refer to apples is by social convention
- The meaning of ``apple'' is learned from experience
  - What apple looks like
  - What it smells like
  - What it feels like
  - (in the case of humans) what it tastes like

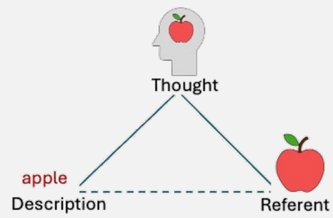

Discuss

- Can a robot assign to `apple’ the same meaning as you?
  - Why or why not?
- Do you assign the same meaning to `apple’ as your human friends?
  - Why or why not?

Human-designed representations in AI systems

- In AI systems with representations built by human designers
  - The correspondence between the surrogate or description and the aspects of the world they refer to is established by the designers
  - The meaning or semantics of the representation are provided by the human designers of the AI system.
  - In AI systems with human-supplied semantics, there is no necessary correspondence between the string of letters ``apple'' and what it refers to, except through the mind of the designer.
- Such machines may not ``understand'' anything the way a human might
- All such machines can do is exhibit competence without comprehension

Self-acquired representations in AI systems

- In robots with representations that are learned from interaction with the world
  - The correspondence between the surrogate or description and the aspects of the world they refer to arises from coupling of the representations with sensory percepts and actions
  - The meaning or semantics of the representation are acquired by the robot
  - The descriptions used by the robot are meaningful to it just as descriptions used by us are meaningful to us
- Such robots ``understand'' the world in ways a human might, although the machine's understanding might differ from ours because of differences in descriptions and their semantics
- Such robots can exhibit competence with comprehension

All representations are wrong, but some are useful

- What a useful thing a pocket-map is! I remarked.
- Map making. That's another thing we've learned from your Nation, said Mein Herr. But we've carried it much further than you. What do you consider the largest map that would be really useful?''
- About six inches to the mile.
- Only six inches! exclaimed Mein Herr. We very soon got to six yards to the mile. Then we tried a hundred yards to the mile. And then came the grandest idea of all ! We actually made a map of the country, on the scale of a mile to the mile!''
- Have you used it much? I enquired.
- It has never been spread out, yet," said Mein Herr: the farmers objected: they said it would cover the whole country, and shut out the sunlight! So we now use the country itself, as its own map, and I assure you it does nearly as well.

Sylvie and Bruno Concluded, Lewis Carroll, 1893

All representations are wrong, but some are useful

- The most accurate representation of the world is the world itself
- Representations are cartoons – they keep some details, throw out other details
- More detail is not necessarily better
- Because all representations are cartoons,
  - all broad-based reasoning about the natural world,
  - whether by humans or machines,
  - because of its reliance on a representation that is in some aspects wrong,
  - will at some point yield conclusions that are necessarily incorrect
- We can assess the usefulness of representations only relative to the tasks we want to perform
- All models are wrong, but some are useful – George Box

Sticks and Squares Problem

- We have 17 sticks arranged to form 6 squares
- Your task is to remove exactly 5 sticks so that we are left with exactly 3 squares and no orphan sticks

Sticks and Squares Problem

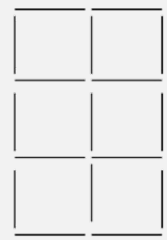

- We have 17 sticks arranged to form 6 squares
- Your task is to remove exactly 5 sticks so that we are left with exactly 3 squares and no orphan sticks
- Level of abstraction – sticks
  - $\binom{17}{5} = \frac{17!}{(17-5)!5!} = 6188$ configurations
- Level of abstraction – squares
  - $\binom{6}{3} = \frac{6!}{(6-3)!3!} = 20$ configurations

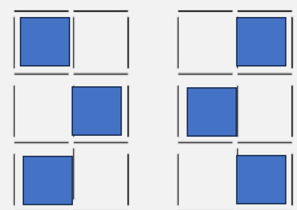

Feature spaces

- A feature space is constructed by assigning a problem-related attribute or feature or measurement to a distinct dimension of a multidimensional space.
- Figure shows a two-dimensional feature space with one axis representing weight and the other color (redness).
- Points that are close together in this representation represent objects that have similar color-weight measurements.
- Feature spaces are used by machine learning systems

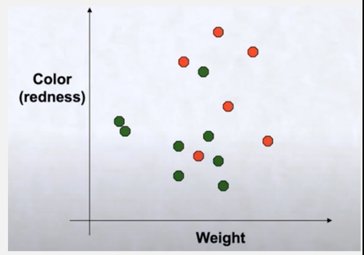

Relational networks

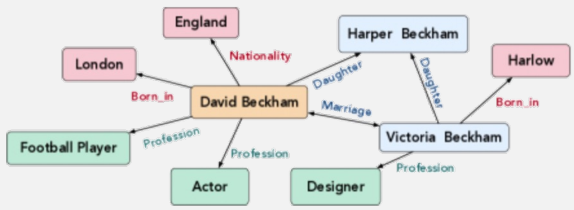

- Relational networks are graphs where the nodes represent individuals or other entities and links denote relationships between entities.
- Such a representation can be used to answer queries about individuals and their relationships to other individuals or entities
  - Who is David Beckham married to?
  - Who is the daughter of David Beckham and Victoria Beckham?
  - Who is both an actor and a Football player? etc.

Game trees

- In a game tree, each node, representing a state, is connected to one or more successor states.
- The goal is to traverse the tree from an initial state to a desired goal state.
- Figure shows the game tree for the familiar 2-person game of tic-tac-toe.
- Two players (X and O) take turns, and at each turn the player can place his or her mark (X or O) in an empty square.
- The player who gets three X's or three O's along a row, column, or diagonal wins.

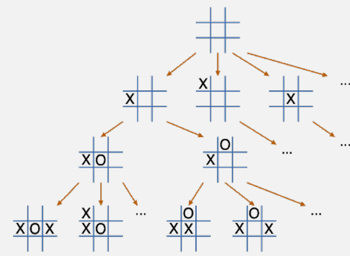

State transition graphs

- This representation is a generalization of finite state automaton. The nodes denote states and labeled links denote conditions that the input must satisfy for the corresponding state transition to occur.
- The state diagram in Figure shows a ``even parity machine''

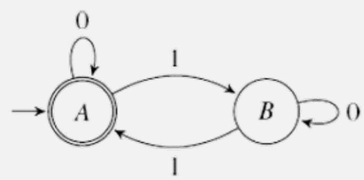

Logic

- The propositional and predicate logic, and their extensions, e.g., logics of knowledge and belief formalize of the process of inferring assertions about the world that are true from existing facts about the world.
- Given All men are mortal, that is, Socrates is a man, we can deduce using predicate logic that Socrates is mortal

Procedural representation

- A procedural representation describes knowledge about the world in terms of a procedure that instructs an agent to perform a specific task.

Procedure Boil Water
(1) Obtain pot, and put water in it
(2) Put pot over the stove, and light the stove
(3) Turn off the stove when steam rises

Frame representation

- A frame provides a means of representing knowledge about the objects, relationships, and events that are prototypical in a setting.
- The aspects of a given scenario are stored as entries in the slots of the frame.
- A partial frame representation of a hotel room is shown in Figure

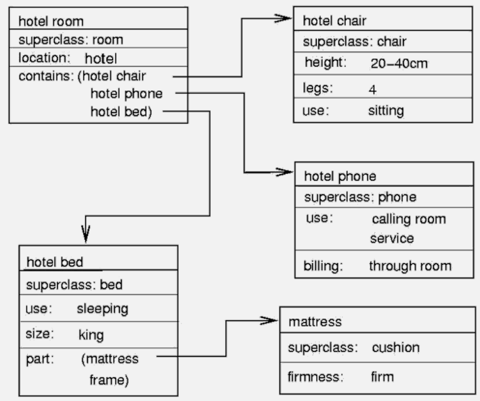

Production(rule) systems

- Production systems use rules of the form IF condition THEN consequence.
- Such rules form the basis of the so-called expert systems
- Applications range from configuring computer systems to medical diagnosis.

Rules of the Road

- IF the stoplight is red AND there is no "No turn on Red" sign AND you have stopped, THEN a right turn is ok.
- IF the stoplight is green THEN a right turn is ok.
- IF right turn is ok THEN turn right.

Isomorphic representations

- Isomorphic representations enforce a direct structural relation between the representation and some of the properties of the world being represented.
- Road maps offer an example of a widely used isomorphic representation.
- For any current location on the map, we might ask,
    - What is the nearest major town, and how far away is it?
    - What is the closest highway intersection where at least three roads meet?