Artificial Intelligence
09/11/2002

Thinking Rationally: Laws of Thoughts

- Aristotle (~450 BC) attempted to codify "right thinking"
  - What are correct arguments/thought processes?
  - e.g., "Socrates is a man, all men are mortal; therefore Socrates is mortal"

- Several Greek schools developed various forms of logic: **notation plus rules of derivation for thoughts**

- Problems:
  - Uncertainty: not all facts are certain
  - Resource limitations: there is a difference between solving a problem in principle and solving it in practice under various resource limitations such as time, computation, accuracy, etc.

Acting Rationally: The Rational Agent

- Rational behavior: Doing the right thing!
- The right thing: That which is expected to maximize the expected return
- Provides the most general view of AI because it includes:
  - Correct inference ("Laws of thoughts")
  - Uncertainty handling
  - Resource limitation consideration (e.g., reflex vs. deliberation)
  - Cognitive skills (NLP, knowledge representation, ML, etc.)
- Advantages:
  - More general
  - Its goal of rationality is well defined

How to achieve AI?

- How is AI research done?
- AI research has both theoretical and experimental sides
  - The experimental side has both basic and applied aspects
- There are two main lines of research
  - One is biological, based on the idea that since humans are intelligence, AI should study humans and imitate their psychology or physiology
  - The other is phenomenal, based on studying and formalizing common sense facts about the world and the problems that the world presents to the achievements of goals
- The two approaches interact to some extent, and both should eventually succeed. It is a race but not racers seem to be walking (John McCarthy)

What is an (intelligent) agent?

- An over-used, over-loaded, and misused term
- Anything that can be viewed as perceiving its environment through sensors and acting upon that environment thought its effectors to maximize the progress towards its goals
- **PAGE (Percepts, Actions, Goals, Environment)**
- Task-specific & specialized
  - Well-defined goals and environment

Intelligent Agents and AI

- Human mind as network of thousands or millions of agents all working in parallel
  - To produce real AI, we should build computer systems that also contain many agents and systems for arbitrating among the agents competing results
- Distributed decision-making and control
- Challenges
  - Action selection: what next action to choose
  - Conflict resolution

Agent Types

- We can split agent research into two main strands
  - Distributed AI (DAI) or Multi-Agent Systems (MAS): 1980-1990
  - Much broader notion of “agent”
    - Interface, reactive, mobile, information (1990’s – present)

A windshield wiper agent

- How do we design a agent that can wipe the windshields when needed?
  - Goals?
  - Percepts?
  - Sensors?
  - Effectors?
  - Actions?
  - Environment?

A windshield wiper agent (cont'd)

- How do we design a agent that can wipe the windshields when needed?
  - Goals: To keep windshield clean and maintain good visibility
  - Percepts: Raining, dirty
  - Sensors: Camera (moist sensor)
  - Effectors: Wipers (left, right, back)
  - Actions: Off, slow, medium, fast
  - Environment: Inner city, freeway, highway, weather, …

Example: Autonomous Vehicle

- **Beobots**: goal is to build a robot that can operate in unconstrained environments and that can solve a wide variety of tasks (http://ilab.usc.edu/beobots/home.shtml)
  - Lot of CPU power
  - Prototype robotic platform
  - Visual system to find interesting objects in the world
  - Visual system to recognize/identify some of these objects
  - Visual system to know the type of scenery the robot is in
  - Need to:
    - Build an internal representation of the world
    - Understand what the user wants
    - Act upon user requests/solve user problems

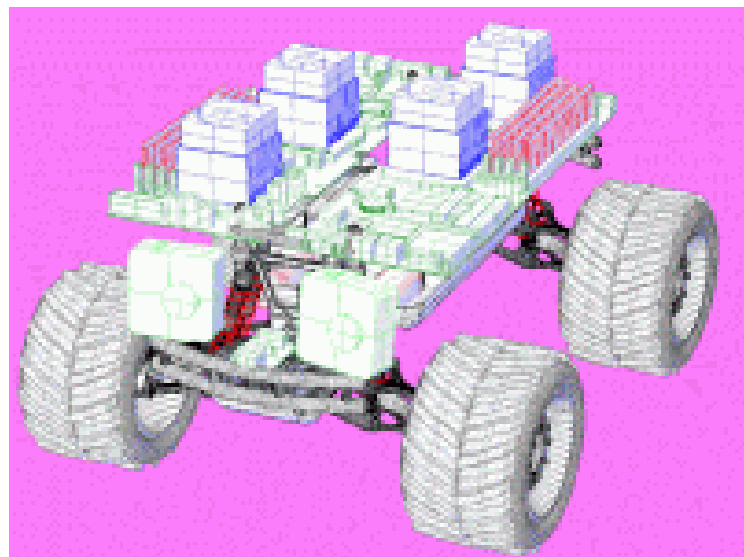

Interacting Agents

	Collision Avoidance Agent (CAA)
Goals	
Percepts	
Sensors	
Effectors	
Actions	
Environment	
	Lane Keeping Agent (LKA)
Goals	
Percepts	
Sensors	
Effectors	
Actions	
Environment	

Interacting Agents (cont'd)

	Collision Avoidance Agent (CAA)
Goals	Avoid running into obstacles
Percepts	Obstacle distance, velocity, trajectory
Sensors	Vision, proximity sensing
Effectors	Steering wheel, accelerator, brakes, horn, headlights
Actions	Steer, speed up, brake, blow horn, signal (headlights)
Environment	Freeway
	Lane Keeping Agent (LKA)
Goals	Stay in current lane
Percepts	Lane center, lane boundaries
Sensors	Vision
Effectors	Steering wheel, accelerator, brakes
Actions	Steer, speed up, brake
Environment	freeway

Conflict resolution by action selection agent

- **Override**: CAA overrides LKA
- **Arbitrate**: if obstacle is close then CAA else LKA
- **Compromise**: choose action that satisfies both agents
- Any combination of the above
- Challenge: Doing the right thing …

The right thing = The rational action

- **Rational Action:** the action that maximizes the expected value of the performance measure given the percept sequence to date

  - Rational: Best ?
  - Rational: Optimal ?
  - Rational: Omniscience ?
  - Rational: Clairvoyant ?
  - Rational: Successful ?

The right thing = The rational action

- **Rational Action:** the action that maximizes the expected value of the performance measure given the percept sequence to date

  - Rational: Best — Yes, to the best of its knowledge
  - Rational: Optimal — Yes, to the best of its abilities (incl. its constraints)

  - ~~Rational: Omniscience ?~~
  - ~~Rational: Clairvoyant ?~~
  - ~~Rational: Successful ?~~

Behavior and performance of IA

- **Perception** (sequence) to **Action Mapping**
  - Ideal mapping: specifies which actions an agent ought to take at any point in time
  - e.g., look-up table

- **Performance measures**
  - A subjective measure to characterize how successful an agent is (e.g., speed, power usage, accuracy, cost, …)

- (degree of) **Autonomy**
  - To what extent is the agent able to make decisions and actions on its own

How is agent different from other software?

- Agents are **autonomous**, that is they act on behalf of the user
- Agents contain some level of **intelligence**, from fixed rules to learning engines that allow them to adapt to changes in the environment
- Agents do not only act **reactively**, but sometimes also **proactively**
- Agents have **social ability**, that is they communicate with the user, the system, and other agents as required
- Agents may also **cooperate** with other agents to carry out more complex tasks than they themselves can handle
- Agents may **migrate** from one system to another to access remote resources or even to meet other agents

Environment Types

- Characteristics
  - Accessible vs. inaccessible
  - Deterministic vs. non-deterministic
  - Episodic vs. non-episodic
  - Hostile vs. friendly
  - Static vs. dynamic
  - Discrete vs. continuous

Environment Types (cont'd)

Environment	Accessible	Deterministic	Episodic	Static	Discrete
Operating System					
Virtual Reality					
Office Environment					
Mars					

Environment Types (cont'd)

Environment	Accessible	Deterministic	Episodic	Static	Discrete
Operating System	**YES**	**YES**	**NO**	**NO**	**YES**
Virtual Reality	**YES**	**YES**	**YES/NO**	**NO**	**YES/NO**
Office Environment	**NO**	**NO**	**NO**	**NO**	**NO**
Mars	**NO**	**Semi**	**NO**	**SEMI**	**NO**

Structure of Intelligent Agents

- Agent = architecture + program

- **Agent program**: the implementation of the agent's perception-action mapping
  - Function

    Skeleton-Agent(Percept) **returns** Action

    Memory <- UpdateMemory(memory, Percept)

    Action <- ChooseBestAction(memory)

    Memory <- UpdateMemory(memory, Action)

    Return Action

- Architecture: a device that can execute the agent program (e.g., general-purpose computer, specialized device, beobot, etc.)

Agent Types

- Reflex agents
  - Works by finding a rule whose condition matches the current situation and then doing the action associated with that rule
- Reflex agents with internal states
  - Maintain some sort of internal state in order to choose an action
- Goal-based agents
  - Agent needs some sort of goal information, which describes situations that are desirable in addition to the current state description
- Utility-based agents
  - Goals alone are not enough to generate a high-quality behavior
  - Utility is a function that maps a state onto a real number, which describes the associated degree of happiness

Mobile Agents

- Program that can migrate from one machine to another
- Execute in a platform-independent execution environment
- Requires agent execution environment (places)
- Mobility not necessary or sufficient condition for agenthood
- Two types
  - One-hop mobile agent (migrate from one to another place)
  - Multi-hop mobile agent (roam the network from place to place)
- Applications
  - Distributed information retrieval
  - Telecommunication network routing

Information Agents

- Manage the explosive growth of information
- Manipulate or collate information from many distributed sources
- Information agents can be mobile or static
- Example
  - Bargain finder
  - FIDO (the shopping doggie)
  - Internet softbot
- Challenge: ontologies for annotating web pages (e.g., SHOE)