## Part -1 Artificial Intelligence

## Intelligence

#### intelligence is:

- the ability to reason
- the ability to understand
- the ability to create
- the ability to Learn from experience
- the ability to plan and execute complex tasks

Intelligence involves sensing, thinking, and acting.

SENSING	THINKING	ACTING
Translation of sensory inputs (percepts) into a conceptual representation • Computer Vision • Speech Recognition • Language Understanding	<ul> <li>Manipulation of the conceptual representation</li> <li>Knowledge Representation</li> <li>Problem Solving/Planning</li> <li>Learning (making improvements based on the results of past actions)</li> </ul>	<ul> <li>Translation of intent into (physical) actions (reflexive or deliberative)</li> <li>Robotics</li> <li>Speech and Language Synthesis</li> </ul>

### Consciousness vs Unconsciousness

- Consciousness is the state of being characterized by sensation, emotion, volition, and thought : mind
- Unconsciousness is not marked by conscious thought, sensation, or feeling

Human Intelligence (Conscious) vs Artificial Intelligence (Unconscious)

- Artificially Intelligent machines are unconscious machines or special-purpose devices that support humans in specific, complex tasks.
- As digital machines they are equipped with a completely different operating system (digital vs biological)

### **Artificial Intelligence vs. Human Intelligence**

- Nature: While Human Intelligence looks to adjust to new environments by using a combination of various cognitive processes, Al aims to create machines that can imitate human behavior and perform human-like actions.
- Machines are digital, but the human brain is analogous.
- **Functioning**: Humans use the brain's memory, computing power and ability to think,
- while AI-powered machines depend on data and instructions fed into the system.
- Learning Power: Human Intelligence is all about learning and understanding from different incidents and past experiences.
- But AI falls behind in this area, as AI cannot think.

## **Artificial Intelligence**

- Marvin Minsky
- Artificial Intelligence (AI) is the science of making machines do things that would require intelligence if done by humans
- Barr and Feigenbaum
  - "Artificial Intelligence is the part of computer science concerned with designing intelligence computer systems, that is, systems that exhibit the characteristics we associate with intelligence in human behavior."
- Elaine Rich
  - "Al is the study of how to make computers do things at which, at the moment, people are better"

## Artificial Intelligence:

- There are two main ingredients in the many definition of Artificial Intelligence:
- Thought processes and reasoning (Thinking)
- Behavior and performance (Acting)

### Goals of AI

- The engineering goal is to develop the concepts and practices of building intelligent machines. Emphasis is on system building
- The science goal is to develop concepts, theory, and mechanisms to model intelligence. Emphasis on understanding and automating intelligent behavior

- To Solve ..... Is the goal of AI
- A. Solve Real World Problems
- B. Solve Artificial Problems
- C. Extract causes
- D. explain Various sources of Intelligence
- ..... is about AI
- Playing a computer game
- Program machine with your intelligence
- Making an intelligent machine
- Solving computer programs with your intelligence

## Four Definitions/View of Al

- Acting humanly: The Turing test approach
- Thinking humanly: The cognitive modeling approach
- Thinking rationally: The laws of thought approach
- Acting rationally: The rational agent approach
- 1. Rational agent approach is.....
  - A. Thinking rationally
  - B. Acting rationally
  - C. Acting humanly
  - D. Thinking humanly

### Four Faces of AI Research



#### Acting

Top dimension is concerned with *thought processes and reasoning*, where as bottom dimension addresses the *behavior*.

The definition on the left measures the success in terms of fidelity of *human performance*, whereas definitions on the right measure an *ideal concept of intelligence*, which is called **rationality**.

### Four Views Of Al

- **1. Al means acting humanly**, i.e., acting like a person. The classic example of this is the "Turing test" (details on a later slide).
- **2. AI means thinking humanly**, i.e., thinking like a person. The field of Cognitive Science delves into this topic, trying to model how humans think.
- The difference between "acting humanly" and "thinking humanly" is that acting humanly is only concerned with the actions, the outcome or product of the human's thinking process;
- whereas **thinking humanly** is concerned with modeling human thinking processes.
- **3. AI means thinking rationally,** i.e., modeling thinking as a logical process, where conclusions are drawn based on some type of symbolic logic.
- 4. AI means acting rationally, i.e., performing actions that increase the value of the state of the agent or environment in which the agent is acting.
   Rational behavior: doing the right thing
- The "right thing" is that which is expected to maximize goal given the available information.

### Thinking Rationally Vs Acting Rationally

- Thinking Rationally → the ``laws of thought" approach to AI, the whole emphasis was on correct inferences.
- Making correct inferences is sometimes part of being a rational agent, because one way to act rationally is to reason logically to the conclusion that a given action will achieve one's goals, and then to act on that conclusion.
- Acting Rationally →On the other hand, correct inference is not all of rationality, because there are often situations where there is no provably correct thing to do, yet something must still be done.
- There are also ways of acting rationally that cannot be reasonably said to involve inference. For example, pulling one's hand off of a hot stove is a reflex action that is more successful than a slower action taken after careful deliberation.

- If socrates is a man and man is mortal then socrates is mortal. The example is of:
- 1. Acting Humanly
- 2. Thinking Humanly
- 3. Thinking rationally
- 4. Acting Rationally

Rationality involves ......

- A. Making Inference
- B. Reflex actions
- C. Both A and B
- D. None of Above

### Acting humanly: The Turing Test approach

designed to provide a satisfactory operational definition of intelligence.



 A computer passes the test if a human interrogator, after posing some written questions, cannot tell whether the written responses come from a person or from a computer.

## Capabilities to Pass Turing Test

- A computer/machine would need to possess the following capabilities:
- natural language processing to enable it to communicate successfully in English;
- knowledge representation to store what it knows or hears;
- automated reasoning to use the stored information to answer questions and to draw new conclusions;
- machine learning to adapt to new circumstances and to detect and extrapolate patterns.

## **Total Turing Test**

- Turing's test deliberately avoided direct physical interaction between the interrogator and the computer, because *physical* simulation of a person is unnecessary for intelligence.
- TOTAL TURING TEST includes a video signal so that the interrogator can test the subject's perceptual abilities, as well as the opportunity for the interrogator to pass physical objects "through the hatch."
- To pass the total Turing Test, the computer will need
- computer vision to perceive objects, and
- **robotics** to manipulate objects and move about.

- All of the following are needed to pass turing test except:
- A. Inference mechanism
- B. Natural Language Processing
- C. Identify patterns
- D. None of the above
- Total Turing Test also has .....
- A. the property to manipulate objects
- B. See objects
- C. Both A and B
- D. None Of above

### Foundations Of Al

Philosophy (428 B.C.-present)

- Logic, methods of reasoning
- Mind as physical systems

Mathematics(c.800-present)

- Formal representation and proof
- Algorithms, computation, (un)decidability, (in)tractability
- Probability
- Economics(1776-present)
  - Formal theory of rational decisions
- Neuroscience(1861-present)
  - Memory, Cognitive processes

## History Of AI

- 1943: Warren Mc Culloch and Walter Pitts: a model of artificial boolean neurons to perform computations.
  - First steps toward connectionist computation and learning (Hebbian learning).
  - Marvin Minsky and Dann Edmonds (1951) constructed the first neural network computer
- 1950: Alan Turing's "Computing Machinery and Intelligence"
  - First complete vision of AI.

#### The birth of AI (1956):

- Dartmouth Workshop bringing together top minds on automata theory, neural nets and the study of intelligence.

- Allen Newell and Herbert Simon: The logic theorist (first nonnumeric thinking program used for theorem proving)
- For the next 20 years the field was dominated by these participants.

#### Great expectations (1952-1969):

- Newell and Simon introduced the General Problem Solver.
  - Imitation of human problem-solving
- Arthur Samuel (1952-) investigated game playing (checkers ) with great success.
- John McCarthy(1958-) :
  - Inventor of Lisp (second-oldest high-level language)
  - Logic oriented, Advice Taker (separation between knowledge and reasoning)
- Marvin Minsky (1958 -)
  - Introduction of microworlds that appear to require intelligence to solve: e.g. blocksworld.
  - Anti-logic orientation, society of the mind.

#### Collapse in AI research (1966 - 1973):

- Progress was slower than expected.
  - Unrealistic predictions.
- Some systems lacked scalability.
  - Combinatorial explosion in search.
- Fundamental limitations on techniques and representations.
  - Minsky and Papert (1969) Perceptrons.

#### AI revival through knowledge-based systems (1969-1970):

- General-purpose vs. domain specific
  - E.g. the DENDRAL project (Buchanan et al. 1969) First successful knowledge intensive system.
- Expert systems
  - MYCIN to diagnose blood infections (Feigenbaum et al.)
    - Introduction of uncertainty in reasoning.
- Increase in knowledge representation research.
  - Logic, frames, semantic nets, ...

### **History Of Al**

#### AI becomes an industry (1980 - present):

- R1 at DEC (McDermott, 1982)
- Fifth generation project in Japan (1981)
- American response …

Puts an end to the AI winter.

#### Connectionist revival (1986 - present): (Return of Neural Network):

- Parallel distributed processing (RumelHart and McClelland, 1986); backprop.

#### AI becomes a science (1987 - present):

- In speech recognition: hidden markov models
- In neural networks
- In uncertain reasoning and expert systems: Bayesian network formalism

#### The emergence of intelligent agents (1995 - present):

The whole agent problem:

"How does an agent act/behave embedded in real environments with continuous sensory inputs"

- Who is the father Of AI?
- A. Alan Turing
- B. Andrew NG
- C. John Mcarthy
- D. Charles Babbage

### **AI Applications**

- Autonomous planning and scheduling
- Game playing
- Autonomous Control
- Expert Systems
- Logistics Planning
- Robotics
- Language understanding and problem solving
- Speech Recognition
- Computer Vision

# 1. Which of the following is an example of logistics planning in terms of AI application?

- A. Airline flight scheduling
- B. Trip itineraries
- D. Both A and B
- 2. Which of the following is not an AI application?
- A. Chatbots
- B. Machines differentiating a tomato versus mango by visual perception
- C. A machine with a remote controlled human
- D. A robotic arm in an industry

- Arrange the following in terms of increasing order of PEAS
- Simple reflex agents< Model-based reflex agents < Goal-based agents < Utility-based agent
- 3. Simple reflex agents< Goal-based agents < Model-based reflex agents < Utility-based agent
- Simple reflex agents< Goal-based agents <</li>
   Utility-based agent< Model-based reflex agents</li>

# Artificial Intelligence Intelligent Agents



- Agent: An agent is anything that can be viewed as:
  - perceiving its environment through sensors and
  - acting upon that environment through actuators.

- Agent: An agent is anything that can be viewed as:
  - perceiving its environment through sensors and
  - acting upon that environment through actuators.



 An agent program runs in cycles of: (1)perceive, (2)think, and (3)act.

- Agent: An agent is anything that can be viewed as:
  - perceiving its environment through sensors and
  - acting upon that environment through actuators.



- An agent program runs in cycles of: (1)perceive, (2)think, and (3)act.
- Agent = Architecture + Program

### • Human agent:

- Sensors: eyes, ears, and other organs.
- Actuators: hands, legs, mouth, and other body parts.

### • Robotic agent:

- Sensors: Cameras and infrared range finders.
- Actuators: Various motors.

### • Human agent:

- Sensors: eyes, ears, and other organs.
- Actuators: hands, legs, mouth, and other body parts.

### • Robotic agent:

- Sensors: Cameras and infrared range finders.
- Actuators: Various motors.

### • Agents everywhere!

- Thermostat
- Cell phone
- Vacuum cleaner
- Robot
- Alexa Echo
- Self-driving car
- Human
- etc.

### Vacuum cleaner



- Percepts: location and contents e.g., [A, Dirty]
- Actions: Left, Right, Suck, NoOp
- Agent function: mapping from percepts to actions.

### Vacuum cleaner



- Percepts: location and contents e.g., [A, Dirty]
- Actions: Left, Right, Suck, NoOp
- Agent function: mapping from percepts to actions.

Percept	Action	
[A, clean]	Right	
[A, dirty]	Suck	
[B, clean]	Left	
[B, dirty]	Suck	

### Well-behaved agents

### **Rational Agent:**

"For each possible percept sequence, a rational agent should select an action that is expected to maximize its **performance measure**, given the evidence provided by the percept sequence and whatever built-in knowledge the agent has."

### Rationality

- Rationality is relative to a **performance measure**.
- Judge rationality based on:
  - The performance measure that defines the criterion of success.
  - The agent prior knowledge of the environment.
  - The possible actions that the agent can perform.
  - The agent's percept sequence to date.

### PEAS

- When we define a rational agent, we group these properties under **PEAS**, the problem specification for the task environment.
- The rational agent we want to design for this task environment is the solution.
- PEAS stands for:
  - Performance
  - Environment
  - Actuators
  - Sensors





- Performance:
- Environment:
- Actuators:
- Sensors:





- Performance: Safety, time, legal drive, comfort.
- Environment:
- Actuators:
- Sensors:





- Performance: Safety, time, legal drive, comfort.
- Environment: Roads, other cars, pedestrians, road signs.
- Actuators:
- Sensors:





- Performance: Safety, time, legal drive, comfort.
- Environment: Roads, other cars, pedestrians, road signs.
- Actuators: Steering, accelerator, brake, signal, horn.
- Sensors:





- Performance: Safety, time, legal drive, comfort.
- Environment: Roads, other cars, pedestrians, road signs.
- Actuators: Steering, accelerator, brake, signal, horn.
- Sensors: Camera, sonar, GPS, Speedometer, odometer, accelerometer, engine sensors, keyboard.









- Performance: cleanness, efficiency: distance traveled to clean, battery life, security.
- Environment:
- Actuators:
- Sensors:





- Performance: cleanness, efficiency: distance traveled to clean, battery life, security.
- Environment: room, table, wood floor, carpet, different obstacles.
- Actuators:
- Sensors:





- Performance: cleanness, efficiency: distance traveled to clean, battery life, security.
- Environment: room, table, wood floor, carpet, different obstacles.
- Actuators: wheels, different brushes, vacuum extractor.
- Sensors:





- Performance: cleanness, efficiency: distance traveled to clean, battery life, security.
- Environment: room, table, wood floor, carpet, different obstacles.
- Actuators: wheels, different brushes, vacuum extractor.
- Sensors: camera, dirt detection sensor, cliff sensor, bump sensors, infrared wall sensors.

### **Environment types**

- Fully observable (vs. partially observable): An agent's sensors give it access to the complete state of the environment at each point in time.
- Deterministic (vs. stochastic): The next state of the environment is completely determined by the current state and the action executed by the agent. (If the environment is deterministic except for the actions of other agents, then the environment is strategic)
- Episodic (vs. sequential): The agent's experience is divided into atomic "episodes" (each episode consists of the agent perceiving and then performing a single action), and the choice of action in each episode depends only on the episode itself.

### **Environment types**

- Static (vs. dynamic): The environment is unchanged while an agent is deliberating. (The environment is semi-dynamic if the environment itself does not change with the passage of time but the agent's performance score does.)
- Discrete (vs. continuous): A limited number of distinct, clearly defined percepts and actions. E.g., checkers is an example of a discrete environment, while self-driving car evolves in a continuous one.
- Single agent (vs. multi-agent): An agent operating by itself in an environment.
- Known (vs. Unknown): The designer of the agent may or may not have knowledge about the environment makeup. If the environment is unknown the agent will need to know how it works in order to decide.

### **Environment types**

Environment	Observable	Agents	Deterministic	Static	Discrete
8-puzzle	Fully	Single	Deterministic	Static	Discrete
Chess	Fully	Multi	Deterministic	(Semi)Static	Discrete
Pocker	Partially	Multi	Stochastic	Static	Discrete
Backgammon	Fully	Multi	Stochastic	Static	Discrete
Car	Partially	Multi	Stochastic	Dynamic	Continuous
Roomba	Partially	Single	Stochastic	Dynamic	Continuous

### Agent types

- Four basic types in order of increasing generality:
  - Simple reflex agents
  - Model-based reflex agents
  - Goal-based agents
  - Utility-based agents
- All of which can be generalized into learning agents that can improve their performance and generate better actions.

### Simple reflex agents

- Simple reflex agents select an action based on the current state only ignoring the percept history.
- Simple but limited.
- Can only work if the environment is fully observable, that is the correct action is based on the current percept only.



## Vacuum (reflex) agent



- Let's write the algorithm for the Vacuum cleaner...
- Percepts: location and content (location sensor, dirt sensor).
- Actions: Left, Right, Suck, NoOp

Percept	Action	
[A, clean]	Right	
[A, dirty]	Suck	
[B, clean]	Left	
[B, dirty]	Suck	

## Vacuum (reflex) agent



- Let's write the algorithm for the Vacuum cleaner...
- Percepts: location and content (location sensor, dirt sensor).
- Actions: Left, Right, Suck, NoOp

Percept	Action	
[A, clean]	Right	
[A, dirty]	Suck	
[B, clean]	Left	
[B, dirty]	Suck	

What if the vacuum agent is deprived from its location sensor?

### **Model-based reflex agents**

- Handle partial observability by keeping track of the part of the world it can't see now.
- Internal state depending on the percept history (best guess).
- Model of the world based on (1) how the world evolves independently from the agent, and (2) how the agent actions affects the world.



### **Goal-based agents**

- Knowing the current state of the environment is not enough. The agent needs some goal information.
- Agent program combines the goal information with the environment model to choose the actions that achieve that goal.
- Consider the future with "What will happen if I do A?"
- Flexible as knowledge supporting the decisions is explicitly represented and can be modified.



### **Utility-based agents**

- Sometimes achieving the desired goal is not enough. We may look for quicker, safer, cheaper trip to reach a destination.
- Agent happiness should be taken into consideration. We call it utility.
- A utility function is the agent's performance measure
- Because of the uncertainty in the world, a utility agent choses the action that maximizes the expected utility.



### Learning agents

- Programming agents by hand can be very tedious. "Some more expeditious method seem desirable" Alan Turing, 1950.
- Four conceptual components:
  - Learning element: responsible for making improvements
  - Performance element: responsible for selecting external actions. It is what we considered as agent so far.
  - Critic: How well is the agent is doing w.r.t. a fixed performance standard.
  - Problem generator: allows the agent to explore.



### **Agent's organization**

a) Atomic Representation: Each state of the world is a blackbox that has no internal structure. E.g., finding a driving route, each state is a city. AI algorithms: search, games, Markov decision processes, hidden Markov models, etc.



### **Agent's organization**

b) Factored Representation: Each state has some attribute-value properties. E.g., GPS location, amount of gas in the tank.
AI algorithms: constraint satisfaction, and Bayesian networks.



### **Agent's organization**

c) Structured Representation: Relationships between the objects of a state can be explicitly expressed. AI algorithms: first order logic, knowledge-based learning, natural language understanding.



### **Intelligent agents**

- The concept of intelligent agent is central in AI.
- AI aims to design intelligent agents that are useful, reactive, autonomous and even social and pro-active.
- An agent perceives its environment through percept and acts through actuators.
- A performance measure evaluates the behavior of the agent.
- An agent that acts to maximize its expected performance measure is called a rational agent.
- PEAS: A task environment specification that includes Performance measure, Environment, Actuators and Sensors.

### Agent = Architecture + Program

### **Intelligent agents**

- Four types of agents: Reflex agents, model-based agents, goalbased agents, and utility-based agents.
- Agents can improve their performance through learning.
- This is a high-level present of agent programs.
- States representations: atomic, factored, structured. Increasing expressiveness power.

### **Intelligent agents**

- Four types of agents: Reflex agents, model-based agents, goalbased agents, and utility-based agents.
- Agents can improve their performance through learning.
- This is a high-level present of agent programs.
- States representations: atomic, factored, structured. Increasing expressiveness power.

	Search problems		
	Markov decision processes	Constraint satisfaction problems	
	Adversarial games	Bayesian networks	
Reflex	States	Variables	Logic
"Low-level intellige	nce''	"Hig	gh-level intelligence"

Credit: Courtesy Percy Liang

### Credit

• Artificial Intelligence, A Modern Approach. Stuart Russell and Peter Norvig. Third Edition. Pearson Education.

http://aima.cs.berkeley.edu/