Introduction to AI and Intelligent Agent

Unit 9.1

Concept of Artificial Intelligence

Meaning of the word: ``intelligence"

(a) The capacity to acquire and apply knowledge.

(b) The faculty of thought and reason.

(c)Superior powers of mind. See Synonyms at mind.

What Behaviors are Intelligent?

Everyday tasks: recognize a friend, recognize who is calling, translate from one language to another, interpret a photograph, talk, cook a dinner

Formal tasks: prove a logic theorem, geometry, calculus, play chess, checkers, or Go

Expert tasks: engineering design, medical designers, financial analysis

Artificial Intelligence

Based on the above, `artificial intelligence' is about the science and engineering necessary to create artifacts that can

- acquire knowledge, i.e., can learn and extract knowledge; and

- reason with knowledge (leading to doing tasks such as planning, explaining, diagnosing, acting rationally, etc.),

Formal Definitions

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

> "AI is the study of how to make computers do things at which, at the moment, people are better"

Acting humanly: Turing Test

- Turing (1950) "Computing machinery and intelligence":
- "Can machines think?" or "Can machines behave intelligently?"
- Operational test for intelligent behavior: the Imitation Game
  - The interrogator can communicate two with sources: one is human and the other is a machine
  - He must decide which is which
  - If he is wrong half the time, then the machine is intelligent

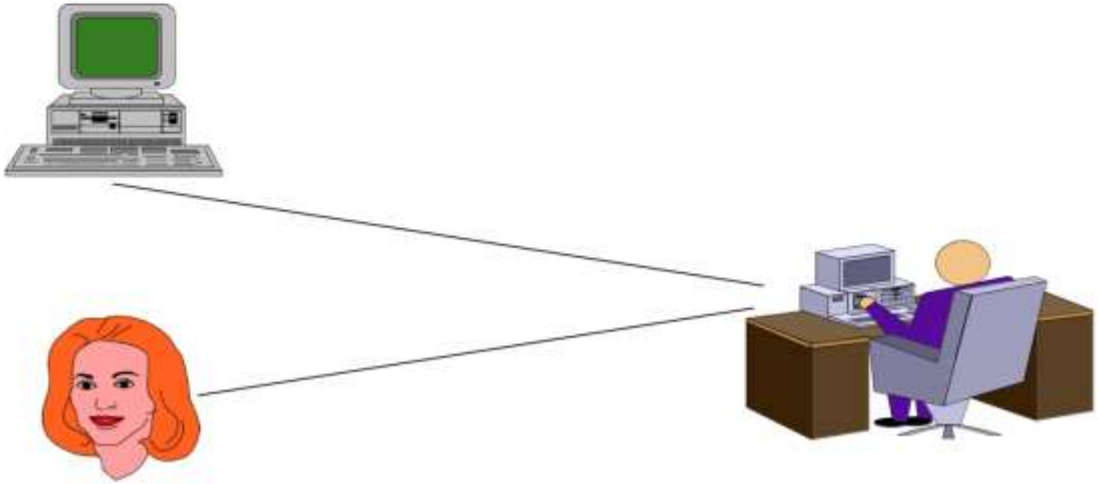

AI Today

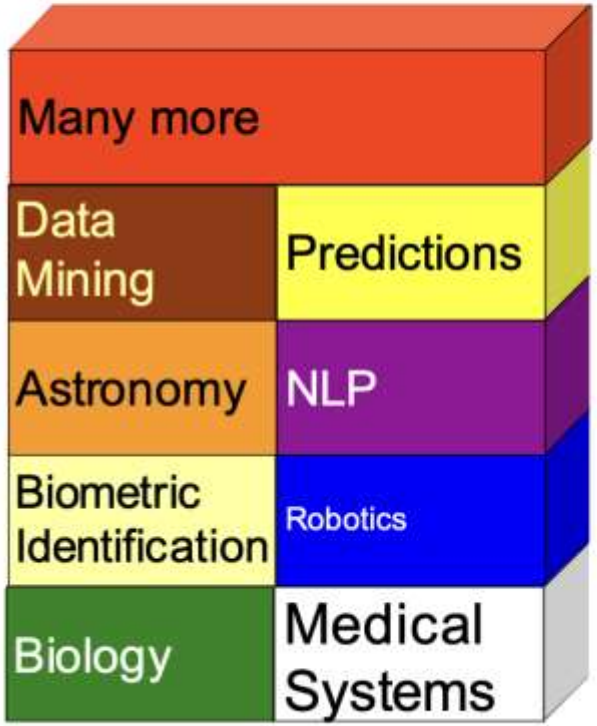
Many more
Data Mining
Predictions
Astronomy
NLP
Biometric Identification
Robotics
Biology
Medical Systems

AI Today

- *Diagnose lymph- node diseases* [Heckerman, 91]
- *Monitor space shuttle missions* [Horvitz, 92]
- *Automatic vehicle control* [Jochem et al., 96]
- *Large- scale scheduling* [Smith et al., 96]
- *Classify astronomical objects* [Goebel et al., 89]
- *Automatic design and configuration systems*
- *First commercial speech understanding systems*
- *Beat world's best players in chess, checkers, and backgammon.*

Prolog

First Prolog program: France, 1970

- Based on theorem proving research
- Major development at University of Edinburgh, 1975-79
- Adopted by the Japanese Fifth Generation Computing Project

Logic programming language:

- Programs composed of facts and rules
- Executes by applying first-order predicate calculus/unification to programs
- Interactive interpreter, compiler
- Tell the computer what is true and what needs to be done, rather than how to do

```
likes(deb, Y) :- horse(Y).
horse(robin).
?- likes(deb, robin).
yes
```

```
likes(deb, horses).
likes(deb, dogs).
?- likes(deb, horses).
yes
?- likes(deb, X).
X=horses
X=dogs
```

LISP

- Proposed by McCarthy, late 1950s; contemporary of COBOL, FORTRAN
- Functional programming language based on lambda calculus/recursive function theory
- Intended as a language for symbolic rather than numeric computation
- Interactive interpreter, compiler
- Uses atoms, lists, functions.

```
(defun hypotenuse (x y)
   (sqrt (+ (square x)
            (square y))))

> (hypotenuse 4 3)
5
```

History of AI

- 1943 McCulloch & Pitts: Boolean circuit model of brain
- 1950 Turing's "Computing Machinery and Intelligence"
- 1956 Dartmouth meeting: "Artificial Intelligence" adopted
- 1950s Early AI programs, including Samuel's checkers program, Newell & Simon's Logic Theorist, Gelernter's Geometry Engine
- 1965 Robinson's complete algorithm for logical reasoning
- 1966—73 AI discovers computational complexity
  Neural network research almost disappears
- 1969—79 Early development of knowledge-based systems
- 1980-- AI becomes an industry
- 1986-- Neural networks return to popularity
- 1987-- AI becomes a science
- 1995-- The emergence of intelligent agents

Artificial Intelligence

- Common sense reasoning
- Reasoning under uncertain conditions
- Learning from experience
- Planning and executing complex tasks
- Understanding and communicating in spoken/written language
- Visual comprehension

Note: "easy for human" != "easy for machine".

Chess is easy for computer, hard for human.

Visual recognition of chess pieces is easy for human, hard for computer.

"Tasks that require knowledge of some domain"

Knowledge Definition

- "The fact or condition of *knowing something* with familiarity gained through experience or association." (Webster's Dictionary, 1988)(Knowing something via seeing, hearing, touching, feeling, and tasting.)

- "The fact or condition of *being aware of* something" . (Ex. Sun is hot, balls are round, sky is blue,…)

Knowledge

Note: *knowledge != data.*

Example: determining voltage from current and resistance

Data	Knowledge
V=12, I=6, R=2 V=28, I=4, R=7 V=9, I=3, R=3 V=12, I=3, R=4 V=15, I=3, R=5	V = I * R

- Knowledge is *more compact* and *faster to manipulate.*
- Knowledge is *more general* (that is, may be applied to situations we have *not* been programmed for).
- Important feature of intelligence: *creating knowledge* from data.

Knowledge

"Heuristics for making good decisions when no algorithm exists for doing so"

That is, for many of the above problem, we must often make *intelligent guesses* due to:

- Lack of *complete knowledge* about how to solve problem.
- Lack of *complete data* about current situation.
- Lack of *time* to completely explore situation.

Example: walking across room.

Do not have complete knowledge of physical laws associated with motion.

Do not have complete knowledge of room (such as what might be behind objects).

Do not have time to map out and compare all possible paths through room.

What is Learning?

Learning is one of those everyday terms which is broadly and vaguely used in the English language

- **Learning is making useful changes in our minds**
- **Learning is constructing or modifying representations of what is being experienced**
- **Learning is the phenomenon of knowledge acquisition in the absence of explicit programming**

Herbert Simon, 1983

Learning denotes changes in the system that are adaptive in the sense that they enable the system to do the same task or tasks drawn from the same population more efficiently and more effectively next time.

What is Learning?

Learning involves 3 factors:

changes

Learning changes the learner: for machine learning the problem is determining the nature of these changes and how to best represent them

generalization

Learning leads to generalization: performance must improve not only on the same task but on similar tasks

improvement

Learning leads to improvements: machine learning must address the possibility that changes may degrade performance and find ways to prevent it.

AI Applications: Consumer Marketing

- Have you ever used any kind of credit/ATM/store card while shopping?
  - if so, you have very likely been "input" to an AI algorithm
- All of this information is recorded digitally
- Companies like Nielsen gather this information weekly and search for patterns
  - general changes in consumer behavior
  - tracking responses to new products
  - identifying customer segments: targeted marketing, e.g., they find out that consumers with sports cars who buy textbooks respond well to offers of new credit cards.
  - Currently a very hot area in marketing

- *How do they do this?*
  - Algorithms ("data mining") search data for patterns
  - based on mathematical theories of learning
  - completely impractical to do manually

AI Applications: Identification Technologies

- *ID cards*
  - e.g., ATM cards
  - can be a nuisance and security risk:

    cards can be lost, stolen, passwords forgotten, etc
- *Biometric Identification*
  - walk up to a locked

    door camera

    fingerprint device

    microphone
  - computer uses your biometric signature for

    identification face, eyes, fingerprints, voice pattern

AI Applications: Predicting the Stock Market

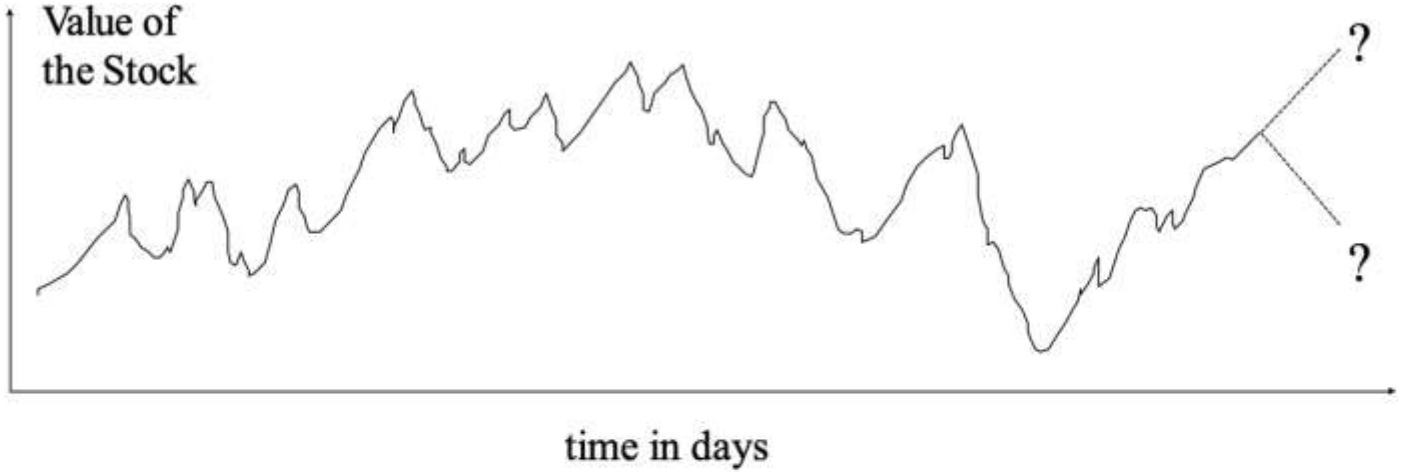

- *The Prediction Problem*
  - given the past, predict the future
  - very difficult problem!
  - we can use learning algorithms to learn a predictive model from historical data

    `prob(increase at day t+1 | values at day t, t-1,t-2....,t-k)`
  - such models are routinely used by banks and financial traders to manage portfolios worth millions of dollars

AI-Applications: Machine Translation

- Language problems in international business
  - e.g., at a meeting of Japanese, Korean, Vietnamese and Swedish investors, no common language
  - or: you are shipping your software manuals to 127 countries
  - solution; hire translators to translate
  - would be much cheaper if a machine could do this!

- How hard is automated translation
  - very difficult!
  - e.g., English to Russian

    "The spirit is willing but the flesh is weak" (English)

    "the vodka is good but the meat is rotten" (Russian)
  - not only must the words be translated, but their meaning also!

- Nonetheless....
  - commercial systems can do alot of the work very well (e.g.,restricted vocabularies in software documentation)
  - algorithms which combine dictionaries, grammar models, etc.
  - see for example babelfish.altavista.com

AI-Applications: Face Detection

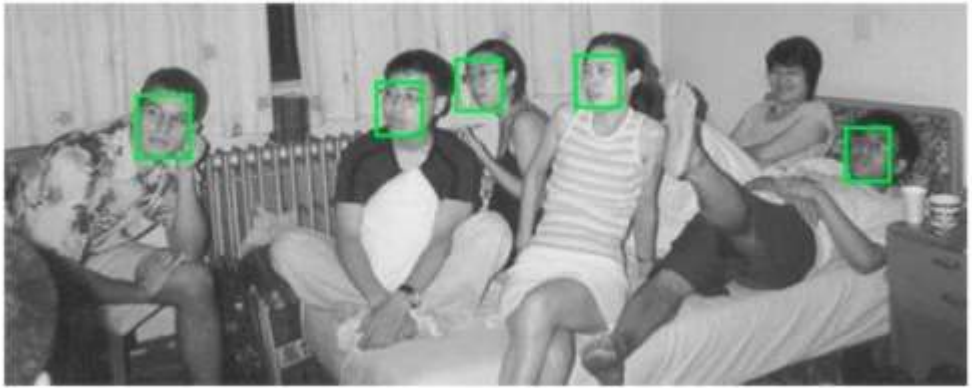

AI-Applications: Product Recommendation

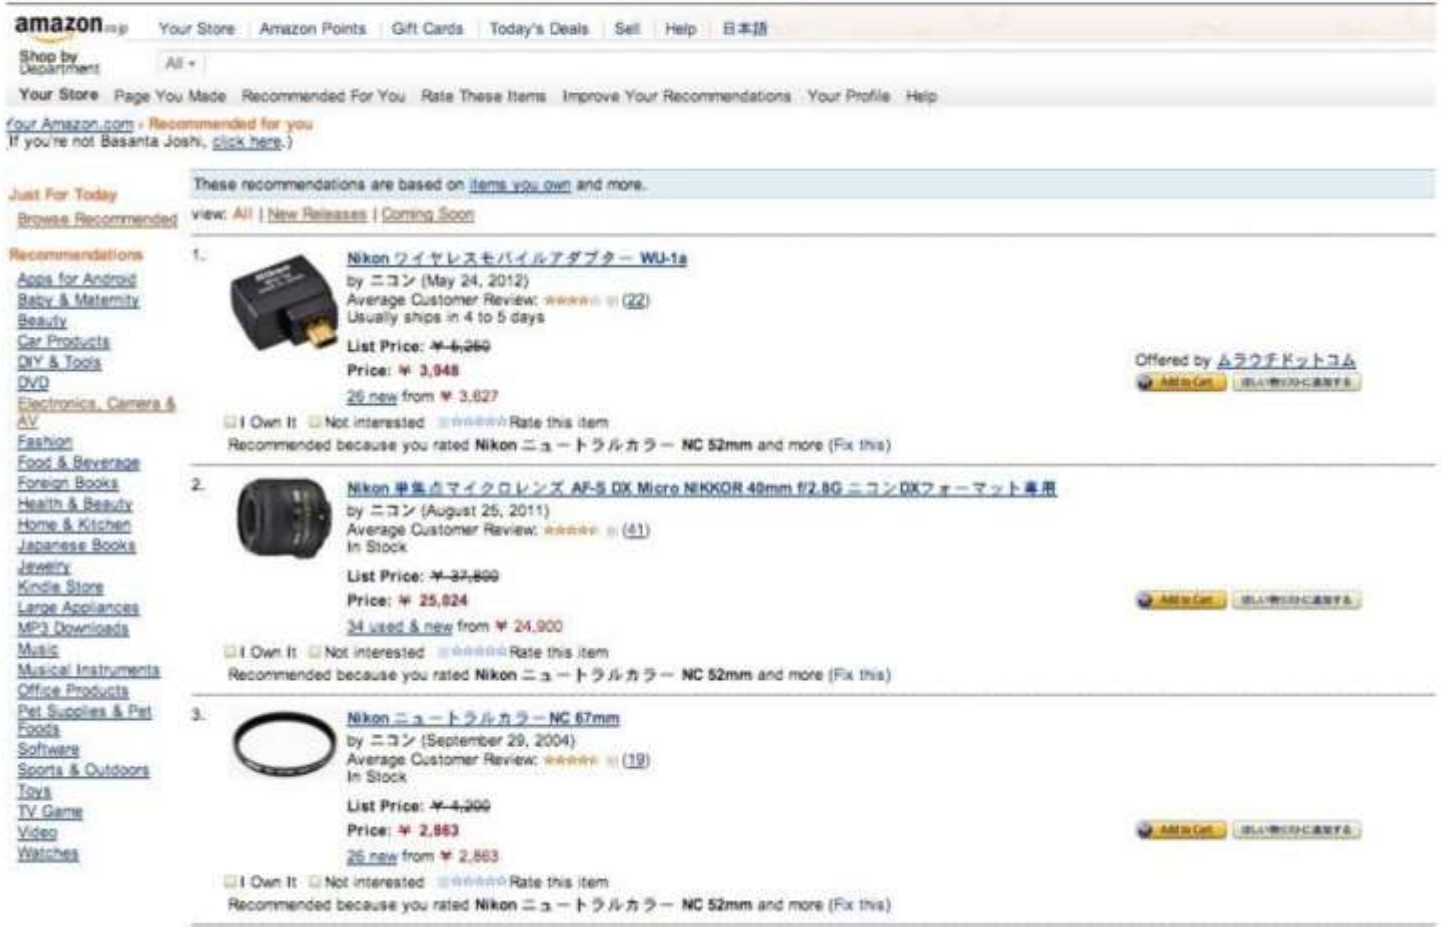

AI-Applications: Image Guided Surgery

Enable surgeons to visualize internal structures through an automated overlay of 3D reconstructions of internal anatomy on top of live video views of a patient

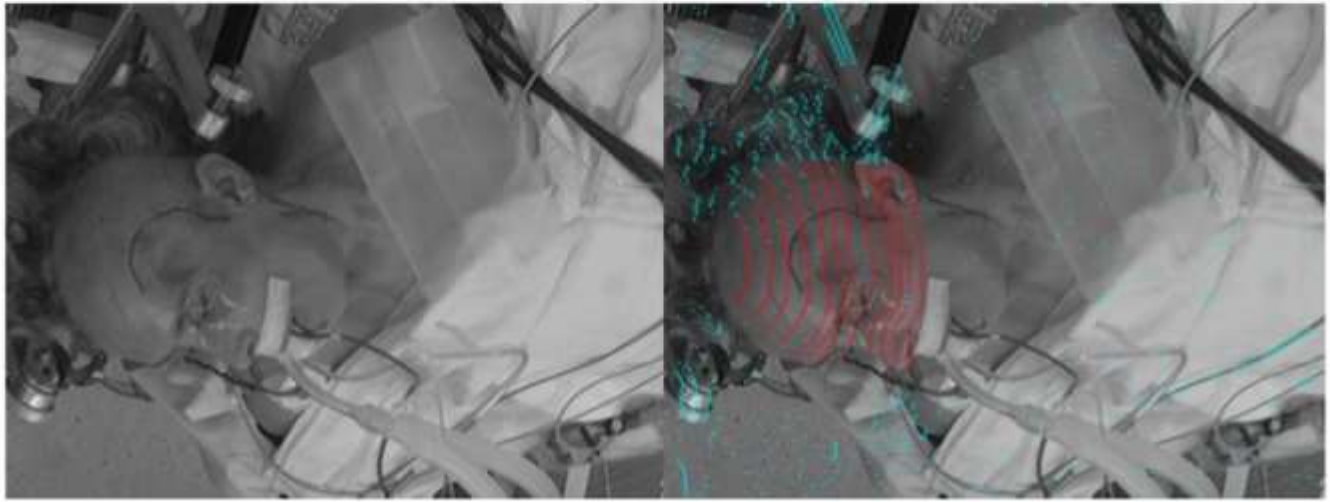

AI-Applications: Speech Recognition

AI-Applications: Exploration of the Universe

Agents

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

Human agent:

eyes, ears, and other organs for sensors;

hands, legs, mouth, and other body parts for actuators

Robotic agent:

cameras and infrared range finders for sensors; various motors for actuators

Agents and environments

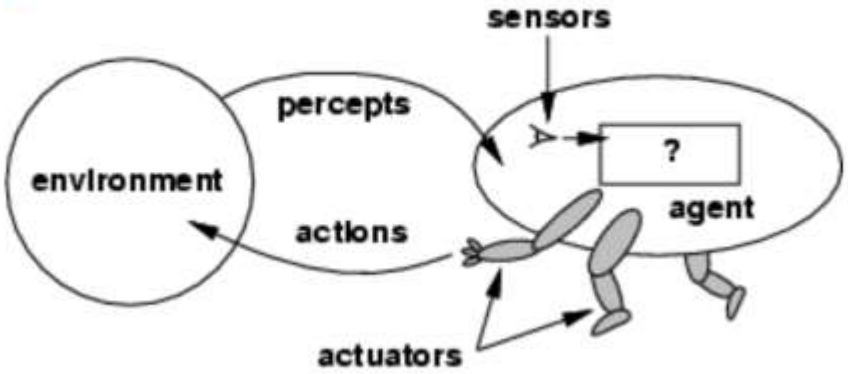

The agent function maps from percept histories to actions:

$$[f: P^* \rightarrow A]$$

The agent program runs on the physical architecture to produce f

agent = architecture + program

Vacuum-cleaner world

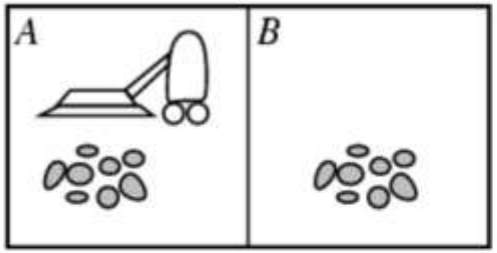

Percepts: location and state of the environment, e.g., [A,Dirty], [B,Clean]

Actions: *Left*, *Right*, *Suck*, *NoOp*

Rational Agents

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

Performance measure: An objective criterion for success of an agent's behavior

E.g., performance measure of a vacuum-cleaner agent could be amount of dirt cleaned up, amount of time taken, amount of electricity consumed, amount of noise generated, etc.

Rational Agents

Rationality is distinct from omniscience (all-knowing with infinite knowledge)

Agents can perform actions in order to modify future percepts so as to obtain useful information (information gathering, exploration)

An agent is autonomous if its behavior is determined by its own percepts & experience (with ability to learn and adapt) without depending solely on build-in knowledge

Task Environment

Before we design an intelligent agent, we must specify its “task environment”:

PEAS:

Performance measure

Environment

Actuators

Sensors

PEAS

Example: Agent = taxi driver

Performance measure: Safe, fast, legal, comfortable trip, maximize profits

Environment: Roads, other traffic, pedestrians,customers

Actuators: Steering wheel, accelerator, brake, signal, horn

Sensors: Cameras, sonar, speedometer, GPS, odometer, engine sensors, keyboard

PEAS

Example: Agent = Medical diagnosis system

Performance measure: Healthy patient, minimize costs, lawsuits

Environment: Patient, hospital, staff

Actuators: Screen display (questions, tests, diagnoses, treatments, referrals)

Sensors: Keyboard (entry of symptoms, findings, patient's answers)

PEAS

Example: Agent = Part-picking robot

Performance measure: Percentage of parts in correct bins

Environment: Conveyor belt with parts, bins

Actuators: Jointed arm and hand

Sensors: Camera, joint angle sensors

AI Classification category

1. Sensing

- Sensing refers to the ability of an agent to perceive its environment through sensors. This could involve collecting data about the surroundings, such as temperature, light levels, or obstacles.
- Example: A robot vacuum uses sensors to detect dirt on the floor and navigate around furniture.

AI Classification category

2. Acting

- Acting involves the agent taking actions based on its perceptions. The actions are determined by the agent's internal decision-making process and can include moving, manipulating objects, or communicating.
- Example: A self-driving car acts by steering, accelerating, or braking in response to road conditions and traffic signals.

AI Classification category

3. Reasoning

Reasoning is the cognitive process that allows an agent to make decisions based on its perceptions and knowledge. This includes evaluating options, predicting outcomes, and planning actions to achieve specific goals.

Example: A chess-playing AI reasons about potential moves by evaluating the positions of pieces on the board and predicting the opponent's responses.

Agent types

Five basic types in order of increasing generality:

Table Driven agent

Simple reflex agents

Model-based reflex agents

Goal-based agents

Utility-based agents

Table Driven Agent.

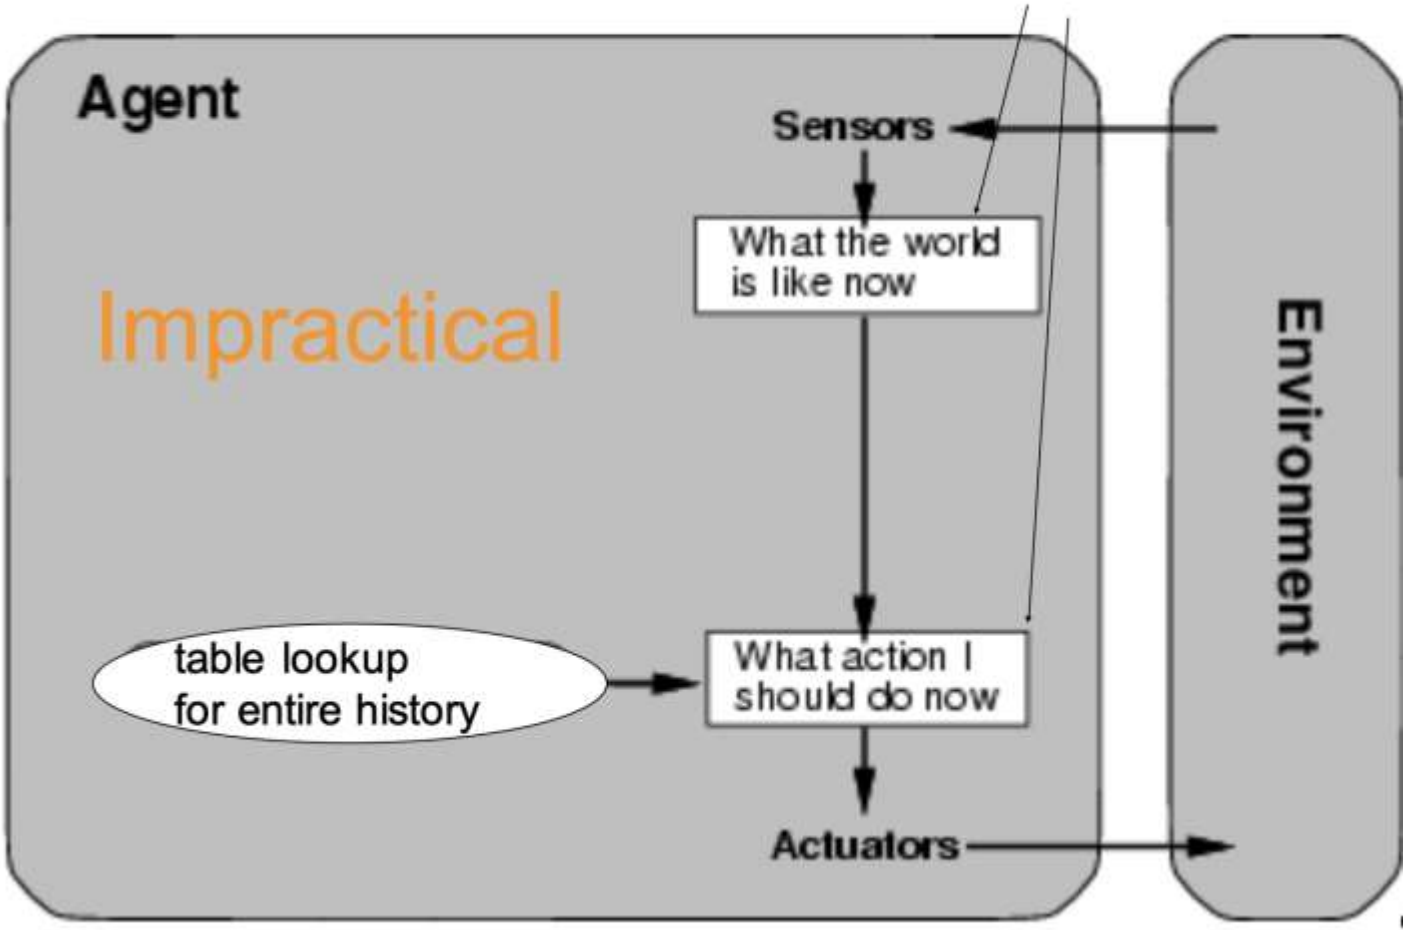

Simple reflex agents

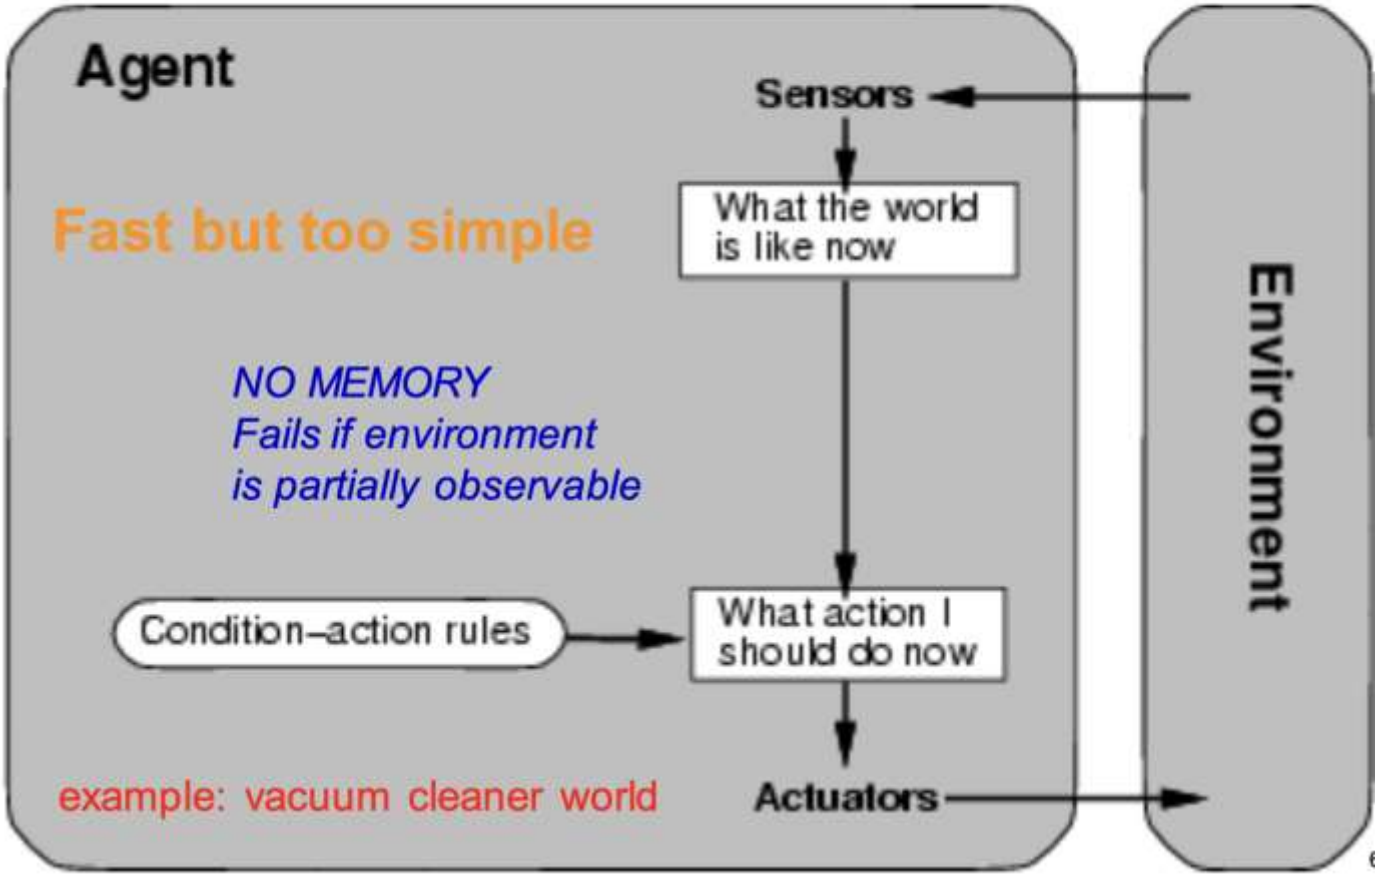

Model-based reflex agents

description of current world state

Model the state of the world by:
modeling how the world chances
how it's actions change the world

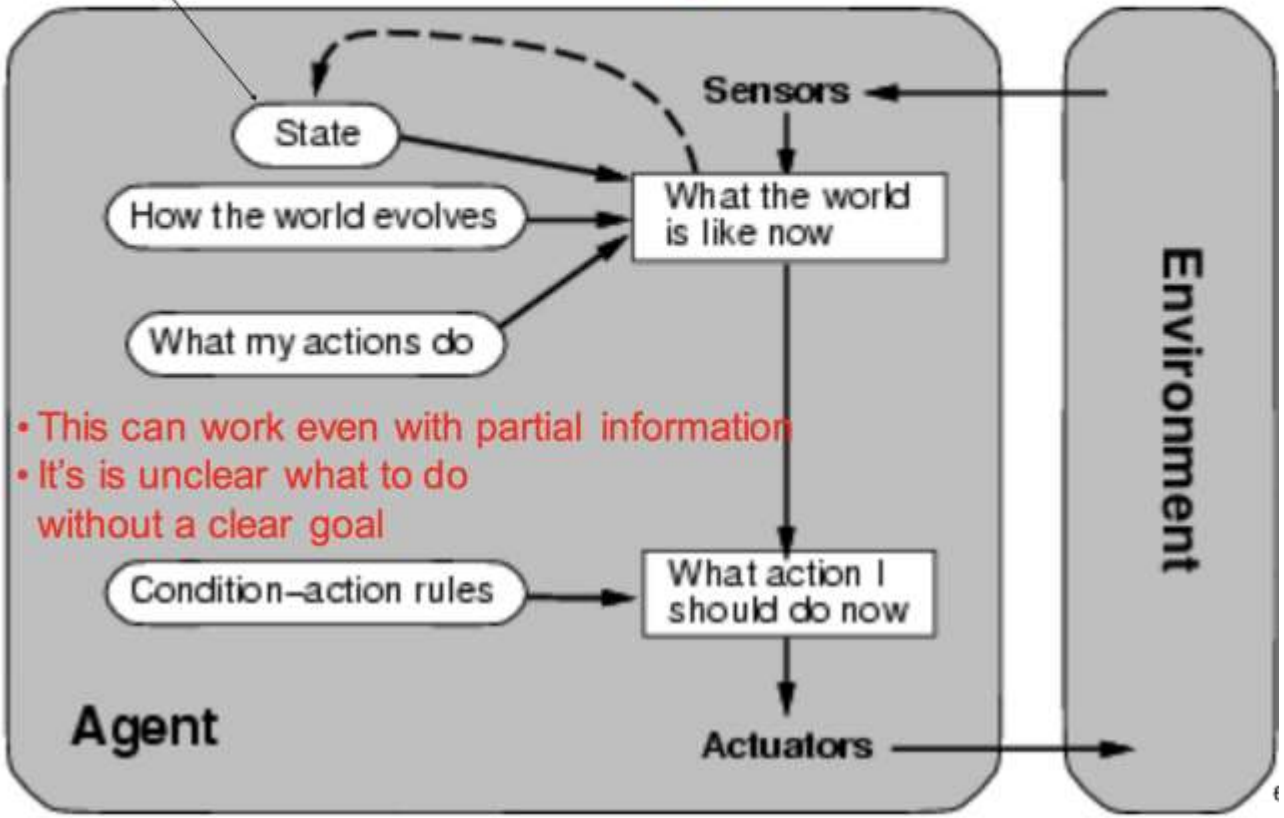

Goal-based agents

Goals provide reason to prefer one action over the other.
We need to predict the future: we need to plan & search

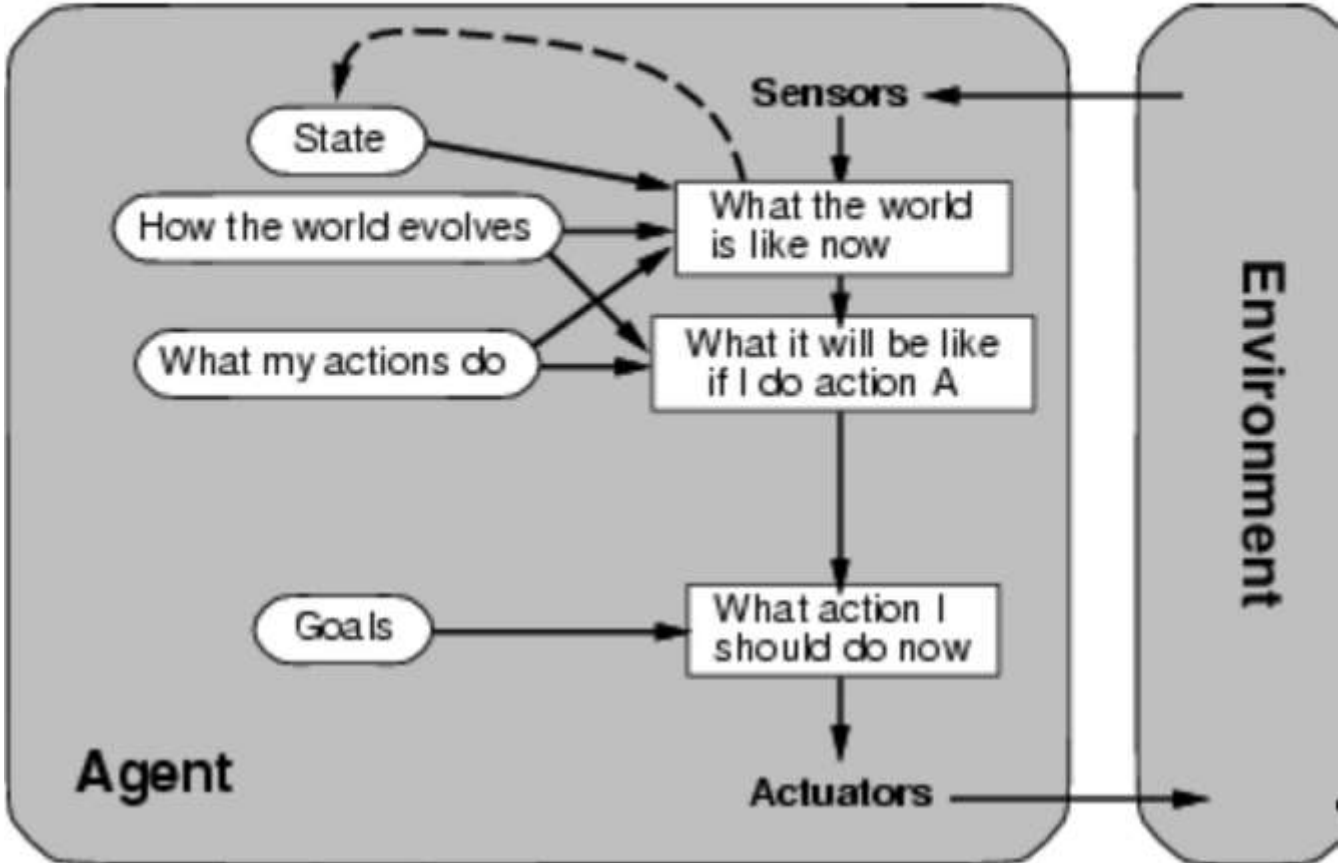

Learning agents

How does an agent improve over time?
By monitoring it's performance and suggesting
better modeling, new action rules, etc.

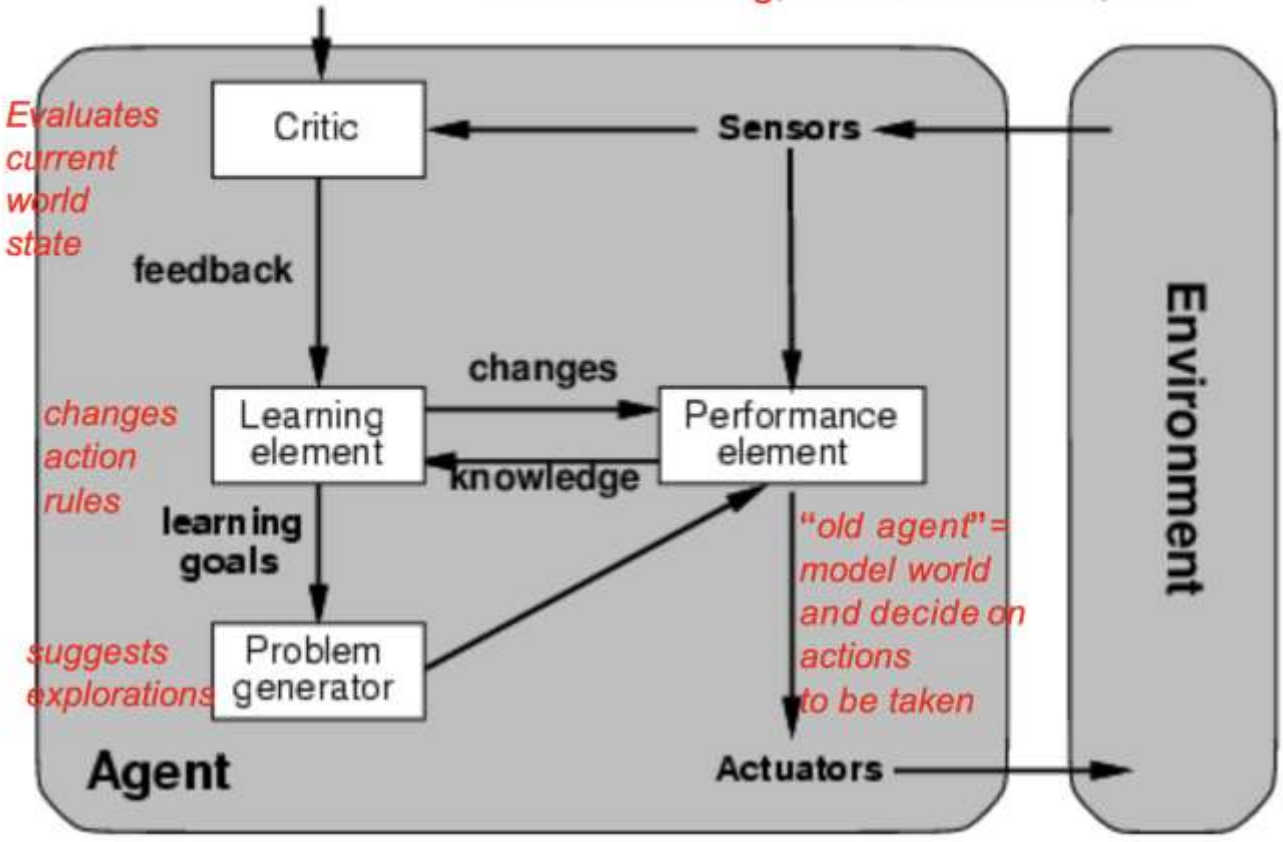

Type of Agent	Description	Example
Simple Reflex Agent	**Operates on a condition-action rule; reacts to current percepts without memory.**	**Spam Email Filter**
Model-Based Agent	**Maintains an internal model of the world; can handle partially observable environments.**	**Autonomous Vehicle**
Goal-Based Agent	**Acts to achieve specific goals; evaluates actions based on their potential to reach goals.**	**Game AI (e.g., Chess AI)**
Utility-Based Agent	**Selects actions based on a utility function to maximize overall satisfaction.**	**Recommendation Systems (e.g., Netflix)**
Learning Agent	**Capable of learning from experiences to improve performance over time.**	**AutoGPT or ChatGPT**

Types of Environment in AI

Fully Observable vs Partially Observable

- When an agent sensor can perceive or access the whole state of an agent at any point in time, the environment is said to be fully observable; otherwise, it is partially observable.
- Maintaining a completely visible environment is simple since there is no need to keep track of the surrounding history.
- When the agent has no sensors in all environments, the environment is said to be unobservable.

Examples:

- Chess – the board and the opponent's movements are both fully observable.
- Driving – the environment is partially observable because what's around the corner is not known.

Deterministic vs Stochastic

- A deterministic environment is one in which an agent's present state and chosen action totally determine the upcoming state of the environment.
- A stochastic environment is unpredictable and cannot be totally predicted by an agent.

Examples:

- Chess – In its current state, a coin has just a few alternative moves, and these moves can be determined.
- Self-Driving Cars– The activities of self-driving cars are not consistent; they change over time.

Dynamic vs Static

- A dynamic environment is one that changes frequently when the agent is doing some action.
- A static environment is one that is idle and does not change its state.

Examples:

- A roller coaster ride is dynamic since it is in motion and the environment changes all the time.
- An empty house is static because nothing changes when an agent arrives.

Single-agent vs Multi-agent

- A single-agent environment is defined as one that has only one agent.
- A multi-agent environment is one in which more than one agent exists.

Examples:

- A person left alone in a maze is an example of the single-agent system.
- Football is a multi-agent game since each team has 11 players.

Fully Observable vs Partially Observable

- When an agent sensor is capable to sense or access the complete state of an agent at each point in time, it is said to be a fully observable environment else it is partially observable.
- Maintaining a fully observable environment is easy as there is no need to keep track of the history of the surrounding.
- An environment is called unobservable when the agent has no sensors in all environments.
- Examples:
    - Chess – the board is fully observable, and so are the opponent’s moves.
    - Driving – the environment is partially observable because what’s around the corner is not known.

MCQ

1. Who is known as the inventor of Artificial Intelligence?

    a) Charles Babbage

    b) John McCarthy

    c) Alan Turing

    d) Andrew Ng

Answer: b)

2. What is meant by Artificial Intelligence?

a) Artificial intelligence is defined as a field aiming to make humans more intelligent.

b) Artificial intelligence is defined as a field aiming to improve security.

c) Artificial intelligence is defined as a field aiming to mine the data.

d) Artificial intelligence is defined as a field aiming to develop intelligent machines.

Answer: d

3. Which is considered the branch of Artificial Intelligence?

a) Cyber Forensics

b) Machine Learning

c) Network Design

d) Full-stack Developer

Answer: b

4. In case a machine is capable of changing its course of action based on the external environment without any external help then the machine is called ____.

a) Intelligent

b) Mobile

c) Both A and B

d) None of the above

Answer: a

5. What is the main aim of Artificial Intelligence?

a) To solve real-world issues

b) To explain different sorts of intelligence

c) To solve artificial problems

d) To obtain information about scientific causes

Answer: b

6 Artificial Intelligence is classified into ____ category processes.

a) 2

b) 5

c) 3

d) 4

Answer: 3

7. ____ is the common language for AI.

a) Lisp

b) Python

c) PHP

d) Java

Answer: b

8. ____ is not an application of AI.

a) Database Management System

b) Digital Assistants

c) Natural language processing

d) Computer Vision

Answer: a

9. ____ is considered a type of artificial intelligence agent.

a) Simple Reflex AI Agent

b) Learning AI agent

c) Goal-Based AI Agent

d) All of the above

Answer: d

10. ____ is a component of AI.

a) Training

b) Designing

c) Learning

d) Puzzling

Answer: c

11. ____ was originally called ‘the Imitation game’.

a) The Turning Test

b) LISP

c) The Halting Problem

d) None of the above

Answer: a

12. There are ____ types of observing environments?

a) 4

b) 3

c) 2

d) 0

Example: c

13. ____ is an example of an Artificial Intelligence Agent.

a) Human

b) Autonomous spacecraft

c) Robot

d) All of the above

Answer: d

14. ____ environment is considered strategic.

a) Partial

b) Stochastic

c) Deterministic

d) Rational

Answer: c

15. LISP was created by?

a) John McCarthy

b) Marvin Minsky

c) Alan Turing

d) Allen Newell and Herbert Simon

Answer: a

16. The performance of an agent can be improved by _____

a) Learning

b) Observing

c) Perceiving

d) None of the mentioned

Answer: a

17. The action of the Simple reflex agent completely depends upon _____

a) Perception history

b) Current perception

c) Learning theory

d) Utility functions

Answer: b

18. Which of the following task/tasks Artificial Intelligence could not do yet?

a) Understand natural language robustly

b) Web mining

c) Construction of plans in real time dynamic systems

d) All of the mentioned

Answer: d

19. What among the following is/are the example of the intelligent agent/agents?

a) Human

b) Robot

c) Autonomous Spacecraft

d) All of the mentioned

Answer: d