Business Intelligent

Unit –I

Chap-II

Agent

- An agent is anything that can be ENVIRONMENT viewed as perceiving its environment through sensors and SENSOR acting upon that environment through actuators.
- This simple idea is illustrated in Figure

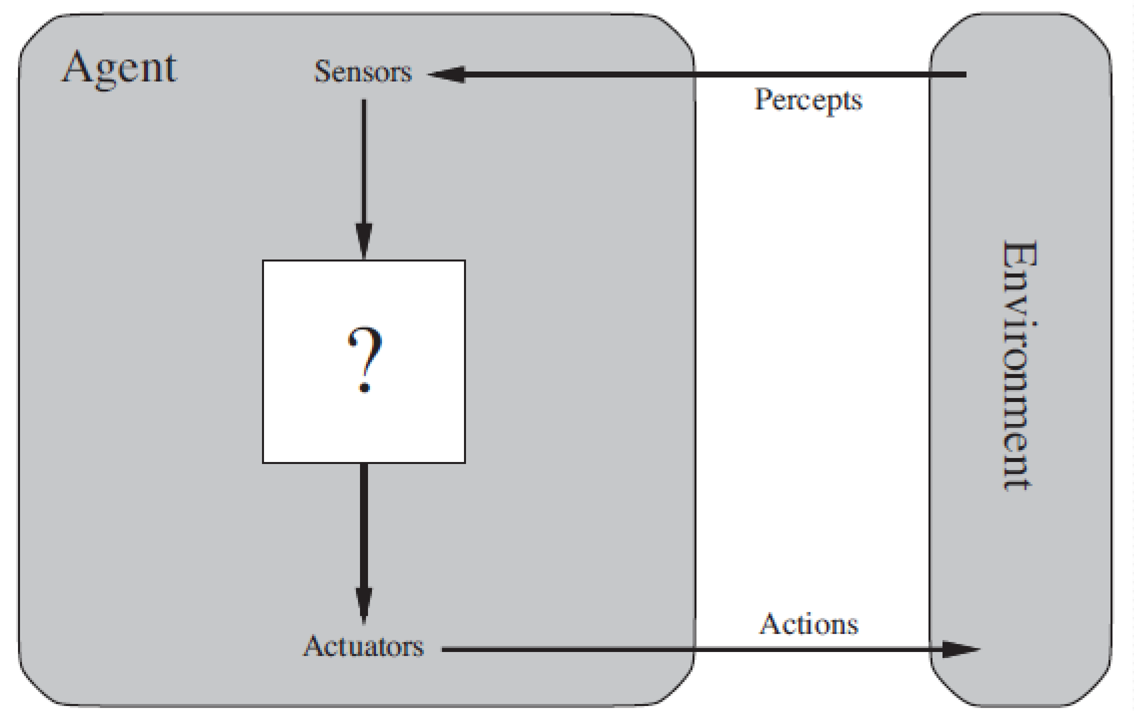

Types of Intelligent Agent

- A human agent has eyes, ears, and other organs for sensors and hands, legs, vocal tract, and so on for actuators.
- A robotic agent might have cameras and infrared range finders for sensors and various motors for actuators. A software agent receives keystrokes, file contents, and

  network packets as sensory inputs and acts on the environment by displaying on the screen, writing files, and sending network packets.

PERCEPT SEQUENCE

- The term percept to refer to the agent’s perceptual inputs at any given instant.
- An PERCEPT SEQUENCE agent’s percept sequence is the complete history of everything the agent has ever perceived.

AGENT PROGRAM

- We can imagine *tabulating the agent function that describes any given agent; for most* agents, this would be a very large table—infinite, in fact, unless we place a bound on the length of percept sequences we want to consider.
- The table is, of course, an *external characterization* of the agent. *Internally, the agent function for an artificial agent will be implemented by an* AGENT PROGRAM.

RATIONAL AGENT

- A rational agent is one RATIONAL AGENT that does the right thing—conceptually speaking, every entry in the table for the agent function is filled out correctly.
- As a general rule, it is better to design performance measures according to what one actually wants in the environment, rather than according to how one thinks the agent should behave.
- 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

- What is rational at any given time depends on four things:
- The performance measure that defines the criterion of success.
- The agent's prior knowledge of the environment.
- The actions that the agent can perform.
- The agent's percept sequence to date.

Demonstration of agent program : Vacuum Cleaner

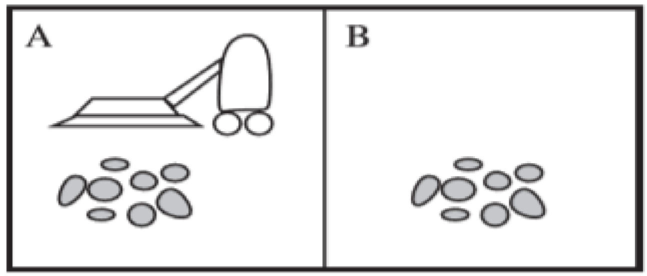

Figure A vacuum-cleaner world with just two locations.

Percept sequence	Action
[*A*, *Clean*]	*Right*
[*A*, *Dirty*]	*Suck*
[*B*, *Clean*]	*Left*
[*B*, *Dirty*]	*Suck*
[*A*, *Clean*], [*A*, *Clean*]	*Right*
[*A*, *Clean*], [*A*, *Dirty*]	*Suck*
⋮	⋮
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Clean*]	*Right*
[*A*, *Clean*], [*A*, *Clean*], [*A*, *Dirty*]	*Suck*
⋮	⋮

Figure Partial tabulation of a simple agent function for the vacuum-cleaner world shown

Explanation:

Vacuum Cleaner world has just two locations:

squares A and B.

The vacuum agent perceives which square it is in and whether there is dirt in the square.

It can choose to move left, move right, suck up the dirt, or do nothing.

One very simple agent function is the following: if the current square is dirty, then suck; otherwise, move to the other square.

NATURE OF ENVIRONMENT

- For Business Intelligent, we must think about task environments, which are essentially the “problems” to which rational agents are the “solutions.”

Specifying the task environment

- For Specifying the task environment, we call this the PEAS (Performance, Environment, Actuators, Sensors) description.
- Table below summarizes the PEAS description for the taxi's task environment.

Agent Type	Performance Measure	Environment	Actuators	Sensors
Taxi driver	Safe, fast, legal, comfortable trip, maximize profits	Roads, other traffic, pedestrians, customers	Steering, accelerator, brake, signal, horn, display	Cameras, sonar, speedometer, GPS, odometer, accelerometer, engine sensors, keyboard

Figure PEAS description of the task environment for an automated taxi.

Agent Type	Performance Measure	Environment	Actuators	Sensors
Medical diagnosis system	Healthy patient, reduced costs	Patient, hospital, staff	Display of questions, tests, diagnoses, treatments, referrals	Keyboard entry of symptoms, findings, patient's answers
Satellite image analysis system	Correct image categorization	Downlink from orbiting satellite	Display of scene categorization	Color pixel arrays
Part-picking robot	Percentage of parts in correct bins	Conveyor belt with parts; bins	Jointed arm and hand	Camera, joint angle sensors
Refinery controller	Purity, yield, safety	Refinery, operators	Valves, pumps, heaters, displays	Temperature, pressure, chemical sensors
Interactive English tutor	Student's score on test	Set of students, testing agency	Display of exercises, suggestions, corrections	Keyboard entry

Figure Examples of agent types and their PEAS descriptions.

Fully observable vs. partially observable:

- If an agent's sensors give it access to complete state of the environment at each point in time, then we say that the task environment is fully observable.
- A task environment is effectively fully observable if the sensors detect all aspects that are relevant to the choice of action; relevance, in turn, depends on the performance measure.
- Fully observable environments are convenient because the agent need not maintain any internal state to keep track of the world.

Fully observable vs. partially observable:

- An environment might be partially observable because of noisy and inaccurate sensors or because parts of the state are simply missing from the sensor data—for example, a vacuum agent with only a local dirt sensor cannot tell whether there is dirt in other squares, and an automated taxi cannot see what other drivers are thinking.
- If the agent has no sensors at all then the environment is unobservable

Single agent vs. Multi agent:

- The distinction between single-agent and multi agent environments may seem simple enough. For example, an agent solving a crossword puzzle by itself is clearly in a single-agent environment, whereas an agent playing chess is in a two agent environment.
- There are, however, some subtle issues, First, we have described how an entity may be viewed as an agent, but we have not explained which entities must be viewed as agents.

COMPETITIVE vs. COOPERATIVE MULTI AGENT

- Chess is a competitive multi agent environment. In the taxi-driving environment, on the other hand, avoiding collisions maximizes the performance measure of all agents, so it is a partially cooperative multi agent environment.

DETERMINISTIC vs. STOCHASTIC

- If the next state of the environment is completely determined by the current state and the action executed by the agent, then we say the environment is deterministic; otherwise, it is stochastic.
- In principle, an agent need not worry about uncertainty in a fully observable, deterministic environment.
- If the environment is partially observable, however, then it could appear to be stochastic.
- Most real situations are so complex that it is impossible to keep track of all the unobserved aspects; for practical purposes, they must be treated as stochastic.
- Eg. Taxi driving is clearly stochastic.

DETERMINISTIC vs. STOCHASTIC

- The vacuum world as we described it is deterministic, but variations can include stochastic elements such as randomly appearing dirt and an unreliable suction mechanism.
- We say an environment is **uncertain if it is not fully observable or not deterministic.**

Note:

- "Stochastic" generally implies that uncertainty about outcomes is quantified in terms of probabilities; a **nondeterministic environment is one in which actions are** characterized by their *possible outcomes, but no probabilities are attached to them.*
- *Nondeterministic* environment descriptions are usually associated with performance measures that require the agent to succeed for *all possible outcomes of its actions.*

Episodic vs. sequential

- In an episodic task environment, the agent’s experience is divided into atomic episodes.
- In each episode the agent receives a percept and then performs a single action.
- Crucially, the next episode does not depend on the actions taken in previous episodes.
- Many classification tasks are episodic.
- For example, an agent that has to spot defective parts on an assembly line bases each decision on the current part, regardless of previous decisions; moreover, the current decision doesn’t affect whether the next part is defective.

Episodic vs. sequential

- In sequential environments, on the other hand, the current decision could affect all future decisions.
- Chess and taxi driving are sequential: in both cases, short-term actions can have long-term consequences.
- **Note:**
- Episodic environments are much simpler than sequential environments because the agent does not need to think ahead.

Static vs. dynamic

- If the environment can change while an agent is deliberating, then we say the environment is dynamic for that agent; otherwise, it is static.
- Static environments are easy to deal with because the agent need not keep looking at the world while it is deciding on an action, nor need it worry about the passage of time.
- Dynamic environments, on the other hand, are continuously asking the agent what it wants to do; if it hasn't decided yet, that counts as deciding to do nothing.

SEMIDYNAMIC

- If the environment itself does not change with the passage of time but the agent’s performance score does, then we say the environment is **semidynamic. Taxi driving is clearly dynamic: the other cars and the taxi itself keep moving** while the driving algorithm dithers about what to do next. Chess, when played with a clock, is semidynamic.
- Crossword puzzles are static.

Discrete vs. continuous

- The discrete/continuous distinction applies to the *state of the* environment, to the way *time is handled, and to the precepts and actions of the agent.*
- For example, the chess environment has a finite number of distinct states (excluding the clock).
- Chess also has a discrete set of precepts and actions. Taxi driving is a continuous-state and continuous-time problem: the speed and location of the taxi and of the other vehicles sweep through a range of continuous values and do so smoothly over time.

Discrete vs. continuous

- Taxi-driving actions are also continuous (steering angles, etc.). Input from digital cameras is discrete, strictly speaking, but is typically treated as representing continuously varying intensities and locations.

Known vs. unknown

- Strictly speaking, this distinction refers not to the environment itself but to the agent's (or designer's) state of knowledge about the "laws of physics" of the environment.
- In a known environment, the outcomes (or outcome probabilities if the environment is stochastic) for all actions are given.
- It is quite possible for a *known environment* to be *partially observable—for example, in solitaire card games*

Final Note:

- As one might expect, the hardest case is *partially observable, multiagent, stochastic, sequential, dynamic, continuous, and unknown. Taxi driving is hard in all these senses, except* that for the most part the driver's environment is known.
- Driving a rented car in a new country with unfamiliar geography and traffic laws is a lot more exciting.

Task Environment	Observable	Agents	Deterministic	Episodic	Static	Discrete
Crossword puzzle	Fully	Single	Deterministic	Sequential	Static	Discrete
Chess with a clock	Fully	Multi	Deterministic	Sequential	Semi	Discrete
Poker	Partially	Multi	Stochastic	Sequential	Static	Discrete
Backgammon	Fully	Multi	Stochastic	Sequential	Static	Discrete
Taxi driving	Partially	Multi	Stochastic	Sequential	Dynamic	Continuous
Medical diagnosis	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Image analysis	Fully	Single	Deterministic	Episodic	Semi	Continuous
Part-picking robot	Partially	Single	Stochastic	Episodic	Dynamic	Continuous
Refinery controller	Partially	Single	Stochastic	Sequential	Dynamic	Continuous
Interactive English tutor	Partially	Multi	Stochastic	Sequential	Dynamic	Discrete

Figure Examples of task environments and their characteristics.