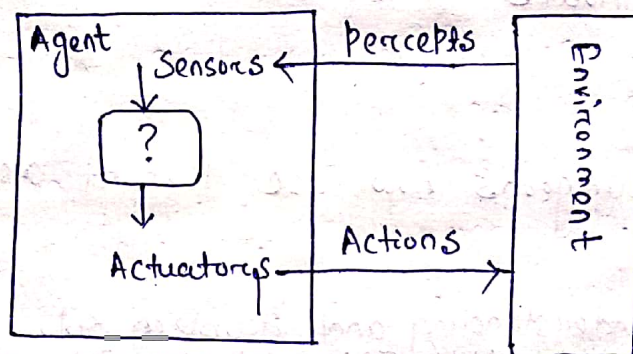


Agents

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

- (1). A human agent has eyes, ears and other organs for sensors and hands, legs, mouth and other body parts for actuators.
- (2). A robotic agent might have cameras and infrared range finders for sensors and various motors for actuators.
- (3). 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.




We use the term percept to refer to the agent's perceptual inputs at any given instant. An agent's percept sequence is the complete history of everything the agent has ever perceived.

In general, an agent's choice of action at any given instant can depend on entire percept sequence observed till date. If we can specify agent's choice of action for every possible percept sequence then we have said more or less everything there is to say about the agent.

Mathematically speaking, we say that an agent's behaviour is described by the agent function that maps any given percept sequence to an action.

Illustration

let us take an example of vacuum cleaner as shown in the below figure.



<u>Percept sequence</u>	<u>Action</u>
[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
⋮	⋮

The above tabulation depicts the partial tabulation of a simple agent function for the vacuum cleaner world.

Rational Agent

A rational agent is one that does the right thing - conceptually speaking, every entry in the table for the agent function is filled out correctly.

Obviously, doing the right thing is better than doing the wrong thing, but what does it mean to do the right thing?

As a first approximation, we will say that the right action is the one that will cause the agent to be most successful. Therefore, we need some way to measure success.

Performance measures

A performance measure embodies the criterion for success of an agent's behaviour. When an agent is plunged down in an environment, it generates a sequence of actions according to the percepts it receives. This sequence of actions causes the environment to go through a sequence of states. If the sequence is desirable, the agent has performed well.

Obviously, there is not one fixed measure suitable for all agents. 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.

Rationality

Rationality depends on four things.

- (1). The performance measure that defines the criterion of success.

- (11)
- (3). Agent's Prior knowledge about the environment.
 - (3). The actions that the agent can perform.
 - (4). The agent's percept sequence till date.

This leads to a definition of 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.

Omniscience, learning and autonomy

An omniscient agent knows the actual outcome of its action and can act accordingly. But omniscience is impossible in reality.

Our definition of rationality does not require omniscience. Because the rational choice depends only on the percept sequence till date.

A rational agent not only gather information but also learn as much as possible from what it perceives. The agent's initial configuration could reflect some prior knowledge of the environment but as the agent gains experience this may be modified and augmented. A rational agent should be autonomous — it should learn what it can to compensate for partial or incorrect prior knowledge. We say that agent lacks autonomy if it relies on the prior knowledge of its designer, rather than on its own percepts.

Specifying the task environment

A task environment specification includes the performance measure, the external environment, the actuators and the sensors (PEAS).

In designing an agent, the first step must always be to specify the task environment as fully as possible. For example, PEAS description of the task environment for an automated taxi is as follows.

1. Agent Type — Taxi Driver.
2. Performance measure — safe, fast, legal, comfortable trip, maximize profit.
3. Environment — Roads, other traffic, pedestrians, customers.
4. Actuators — steering, accelerator, Break, signal, Horn, Display, ~~sensors~~.
5. Sensors — Cameras, Sonar, speedometer, GPS, odometer, accelerometers, engine sensors, keyboard.

Properties of task environments

The range of task environments that might arise in AI is obviously vast. However, we can identify a fairly small number of dimensions along which task environments can be categorized.

1. Fully observable vs Partially observable

If an agent's sensors give it access to the complete state of the environment at each point of time, then we say that the task environment is fully 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.

2. Deterministic vs Stochastic

If the next state of the environment is completely determined by the current state and action executed by the agent, then we say the environment is deterministic; otherwise, it is stochastic.

3. Episodic vs Sequential

In an episodic task environment, the agent's experience is divided into atomic episodes. Each episode consists of agent perceiving and then performing a single action. Here, the choice of action in each episode depends only on the episode itself. Many classification tasks are episodic.

In sequential environment, the current decision could affect all future decisions. For example playing chess, taxi driving etc.

Episodic environment are ~~more~~ much simpler than sequential environment because the agent does not need to think ahead.

4. 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.

5. Discrete vs Continuous

The discrete/continuous distinction can be applied to the state of the environment, to the way time is handled, and

to the percepts and actions of the agent.

6. Single agent vs multi agent

In a single agent environment, only one agent interact with the environment whereas in a multiagent environment more than one agent act upon the environment.

Task Environment	Observable	Deterministic	Episodic	static	Discrete	Agent
Crossword puzzle chess with a clock	Fully fully	Deterministic strategic	sequential sequential	static Semi	Discrete Discrete	Single Multi
Poker Backgammon	Partially Fully	strategic stochastic	sequential sequential	static static	Discrete Discrete	Multi Multi
Taxi driving Medical diagnosis	Partially partially	stochastic stochastic	Sequential sequential	Dynamic Dynamic	Continuous continuous	Multi single
Image analysis Part-picking robot	Fully Partially	Deterministic Stochastic	Episodic episodic	semi Dynamic	Continuous Continuous	Single Single
Reflexory Controllers Interactive English Tutors	Partially partially	Stochastic Stochastic	sequential sequential	Dynamic Dynamic	Continuous Discrete	Single Multi

The structure of agents

The job of AI is to design the agent program that implements the agent function mapping percepts to actions. We assume this program will run on some sort of computing device with physical sensors and actuators - we call this the architecture.

$$\text{Agent} = \text{Architecture} + \text{Program}$$

obviously, the program we choose has to be one that is appropriate for the architecture. If the program is going to recommend actions like walk, the architecture had better have legs.

Agent Programs

An agent program takes the current percept as the input from the sensors and returns an action to the actuators. The agent program takes just the current percept as the input because nothing more is available from the environment. If the agent's actions depend on the entire percept sequence, the agent will have to remember the percepts.

The table driven agent program

Function TABLE-DRIVEN-AGENT (Percept) returns an action

static: percepts, a sequence, initially empty
table, a table of actions, indexed by percept
sequences, initially fully specified

append Percept to the end of percepts

action ← LOOK UP (percepts, table)

return action

Types of agents Program

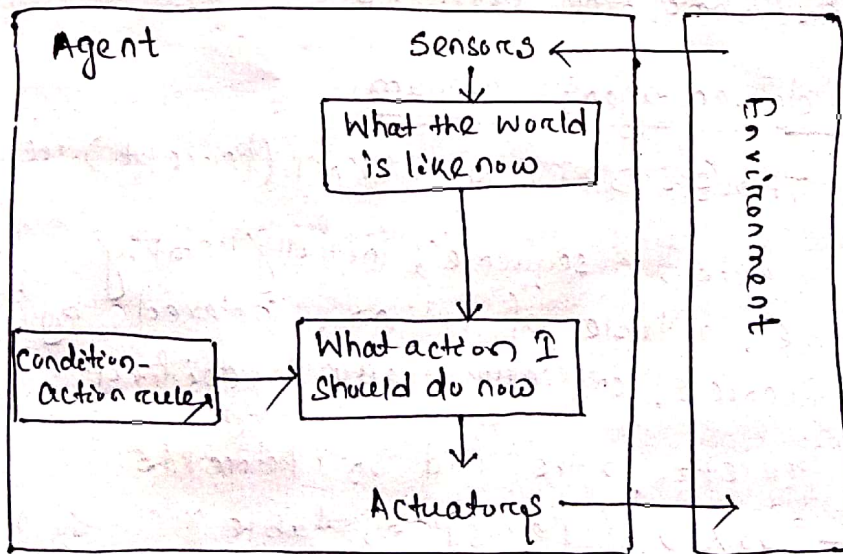
Four basic kinds of agent program that embody the principles ~~of~~ underlying almost all intelligent system are

- (i) Simple reflex agents
- (ii) Model based reflex agents
- (iii) goal based agents
- (iv) utility based agents.

(i) Simple reflex agents

This is a simplest kind of agent. These agents select actions on the basis of current percept, ignoring the rest of the perceived history.

For example, the vacuum cleaner agent whose agent function is a simple reflex agent because its decision is based on only on the current location and on whether that contains dirt or not.



A simple reflex agent program

function SIMPLE-REFLEX-AGENT(percept) returns an action
static: rules, a set of condition-action rules

state ← INTERPRET-INPUT(percept)

rule ← RULE-MATCH(state, rules)

action ← RULE-ACTION [rule]

return action

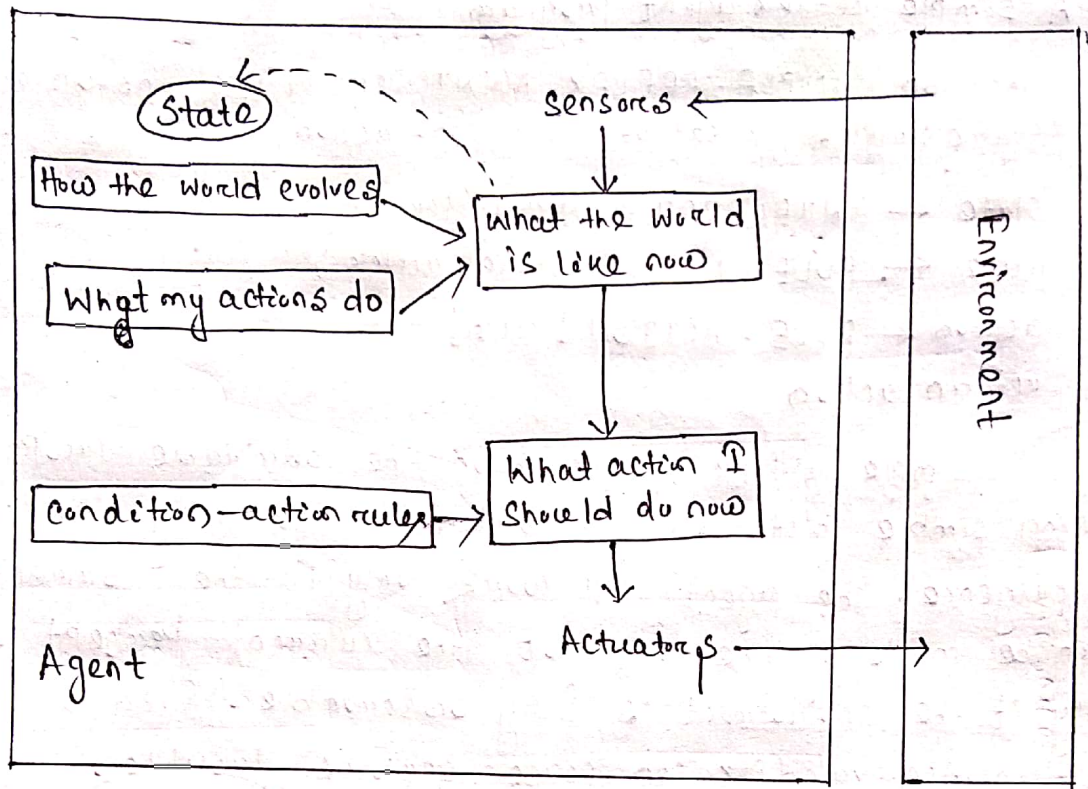
Simple reflex agents have the admirable property of being simple but they turn out to be of very limited intelligence. The agent will work only if the correct decision can be made on the basis of the current percept only - i.e., only if the environment is fully observable. Even a little bit of unobservability can cause serious trouble.

(ii) Model based reflex agents

The most effective way to handle partial observability is for the agent to keep track of the part of the world it can not see now. i.e., the agent should maintain some sort of internal state that depends on the percept history and thereby reflects at least some of the unobserved aspects of the current state.

Updating this internal state information as time goes by requires two kinds of knowledge to be encoded in the agent program.

1. First we need some information about how the world evolves independently of the agent? for example overtaking a car generally will be closer behind than it was a moment ago.
2. We need some information about how the agent's own action affect the world?



Model based reflex agent program

function REFLEX-AGENT-WITH-STATE(percept) returns an action
static: state, a description of the current world state
rules, a set of condition-action rules
action, the most recent action, initially none

state ← UPDATE-STATE(state, action, percept)

rule ← RULE-MATCH(state, rules)

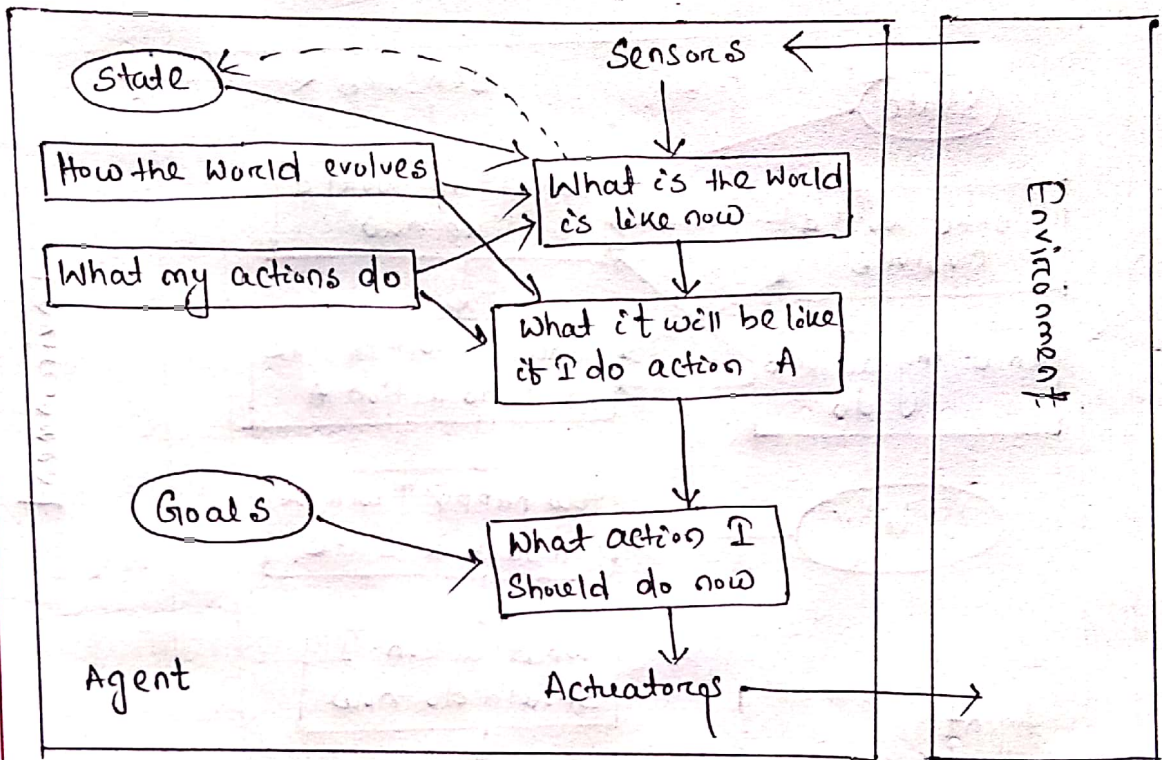
action ← RULE-ACTION[rule]

return action

(iii) Goal-based agents

In goal-based agents, the agent needs some sort of goal information that describes situations that are desirable along with the current state description.

The agent program can combine the information about the results of possible actions in order to choose actions that achieve the goal. Sometimes goal-based action selection is straight forward, when goal satisfaction results immediately from a single action.



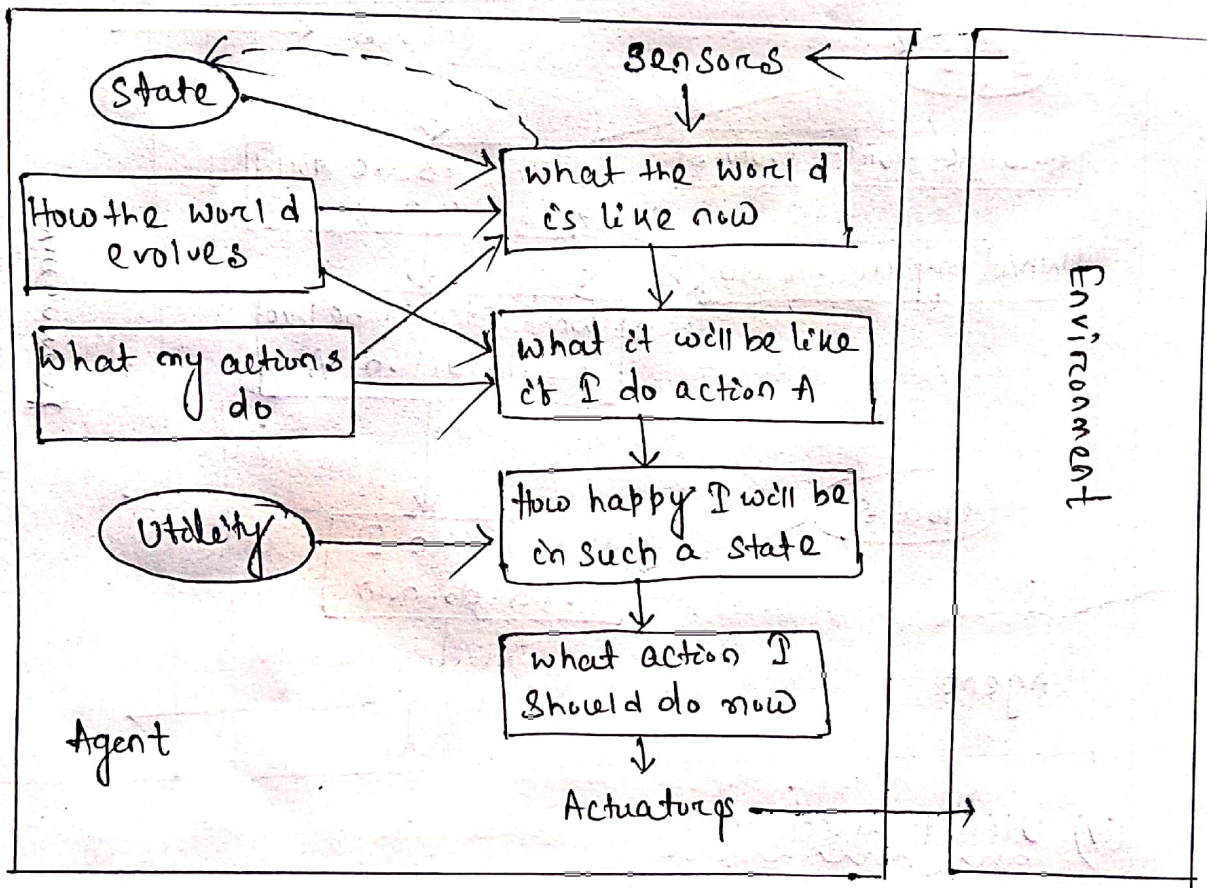
(iv) Utility based agent

Goals alone are not really enough to generate high quality behaviours in most environments. A utility function maps a state (or a sequence of states) on to a real number, which describes the associated degree of happiness.

A complete specification of the utility function allows rational decisions in two kinds of cases where goals are inadequate.

- 1. when there are conflicting goals, only some of which can be achieved (speed or safety), the utility functions specify the appropriate tradeoffs.

2. when there are several goals that the agent can aim for, none of which can be achieved with certainty, utility provides a way in which the likelihood of success can be ~~waited~~ weighed up against the importance of the goals.



Learning agents

A learning agent can be divided into four conceptual components.

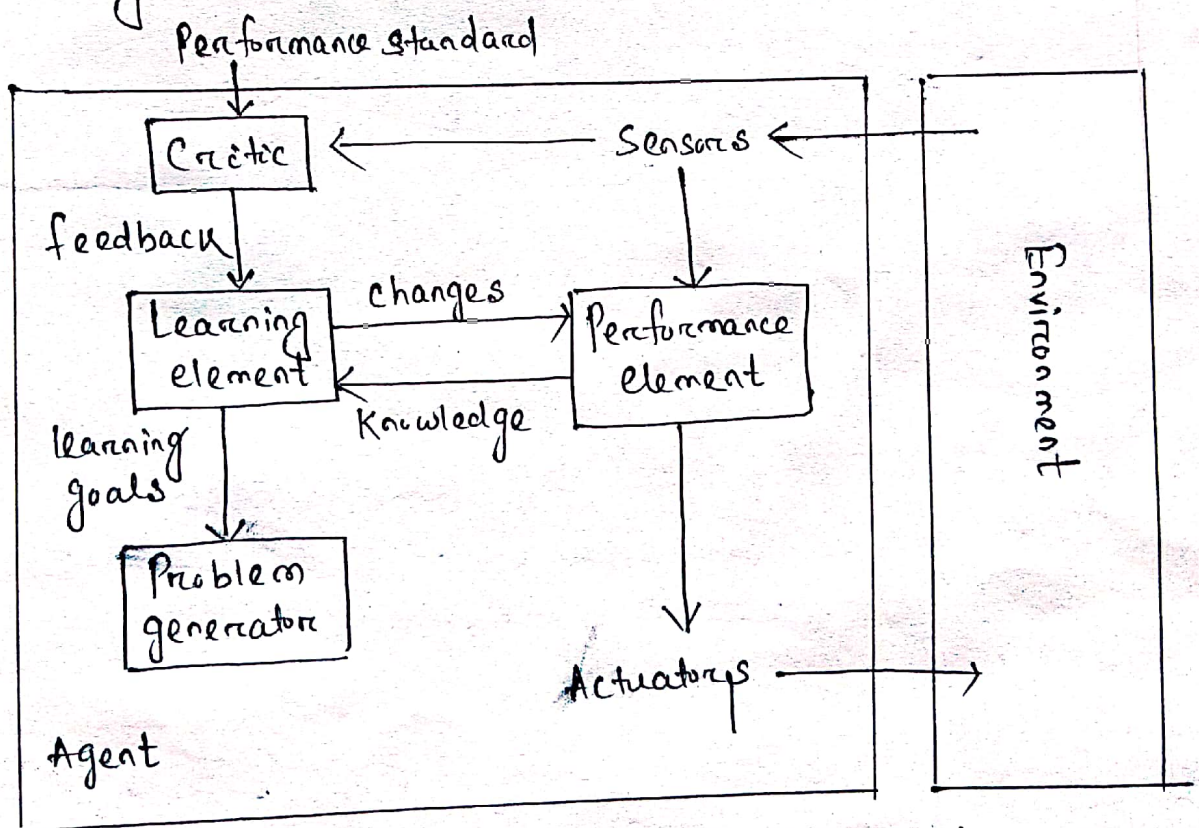
1. Critic
2. Learning element
3. Performance element
4. Problem generator.

A learning element is responsible for making improvements. Performance element is responsible for selecting external actions.

The learning element uses feedback from the critic on how the agent is doing and determines how the performance element should be modified to do better in future.

The problem generator component is responsible for suggesting actions that will lead to new and informative experiences.

All agents can improve their performance through learning.



[General model of learning agents]