Intelligent Agents

- An *agent* is anything that can perceive its environment through *sensors* and act upon that environment through *actuators*.
- For example: a robot agent might have cameras and infrared range finders for sensors and various motors for actuators.
- An *agent's percept sequence* is the complete history of everything the agent has ever perceived.
- Mathematically, we say that an agents behaviour is a *function* that maps any given percept sequence to an action. Such a function is implemented by agent **program**.
- A *rational agent* is that does the right thing. But what does it mean to do right thing?

Rationality depends on four things:

1- The performance measure that defines the criterion of success.
2- The agent's prior knowledge of the environment.
3- The actions that the agent can perform.
4- The agent's percept sequence to data.

So the **rational agent** *should select, for each possible percept sequence, an action that is expected to maximize its performance measure based on its own knowledge.*

PEAS Description

For each agent we can describe it by **PEAS** (Performance, Environment, Actuators, Sensors) description.

Example:

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

Environment Properties:

1- *Fully observable vs. partially observable*: If an agent's sensors give it access to the complete state of the environment at each point in time, then we say that the task environment is fully observable.

2- *Single agent vs. multiagent*: 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.

3- *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.

4- *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. In sequential environments the current decision could affect all future decisions. Chess and taxi driving arc sequential: in both cases, short-term actions can have long-term consequences.

5- *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.

6- *Discrete vs. continuous*: The discrete/continuous distinction applies to the *state* of the environment.

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

Agent Program

The agent program takes just the current percept as input because nothing more is available from the environment; if the agent's actions need to depend on the entire percept sequence, the agent will have to remember the precepts.

Agents can be classified into five classes based on their degree of perceived intelligence and capability

1- Simple reflex agents

These agents select actions on the basis of the *current* percept, ignoring the rest of the percept history. The agent function is based on the *condition-action rule*: if condition then action.

Example: function reflex-vacuum-agent(location, status) returns an action

if status = Dirty then return Suck

else if location = A then return Right

else if location = B then return Left

2- Model-based reflex agents

A model-based agent can handle a partially observable environment. Its current state is stored inside the agent maintaining some kind of structure which describes the part of the world which cannot be seen. This knowledge about "how the world works" is called a *model* of the world, hence the name "model-based agent".

3- Goal-based agents

Goal-based agents further expand on the capabilities of the model-based agents, by using "goal" information. Goal information describes situations that are desirable. This allows the agent a way to choose among multiple possibilities, selecting the one which reaches a goal state. Search and planning are the subfields of artificial intelligence devoted to finding action sequences that achieve the agent's goals.

- In some instances the goal-based agent appears to be less efficient; it is more flexible because the knowledge that supports its decisions is represented explicitly and can be modified.

4- Utility-based agents

Goals alone are not enough to generate high-quality behaviour in most environments. For example, many action sequences will get the taxi to its destination (thereby achieving the goal) but some are quicker, safer, more reliable, or cheaper than others. So we need a utility function to score each possible solution.

- A rational utility-based agent chooses the action that maximizes the expected utility of the action outcomes- that is, the agent expects to derive, on average, given the probabilities and utilities of each outcome.

- A utility-based agent has to model and keep track of its environment, tasks that have involved a great deal of research on perception, representation, reasoning, and learning.

5- Learning agents

Learning has an advantage that it allows the agents to initially operate in unknown environments and to become more competent than its initial knowledge alone might allow.