Study Unit 2: Intelligent Agents

Introduction

AI is about practical reasoning: reasoning in order to do something. A coupling of perception, reasoning, and acting comprises an **agent**. An agent acts in an **environment**. An agent's environment may well include other agents.

- **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.

Box 2.1: Definition of an Agent
An agent is anything that can be viewed as perceiving its environment through sensors and acting upon that environment through actuators.

Agents and Environment

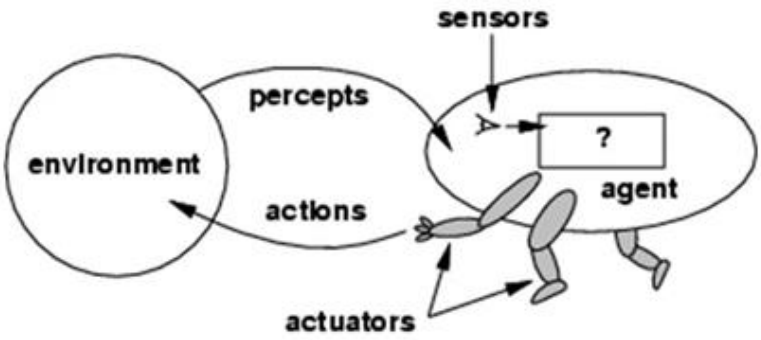

Figure 1: An Agent interact with environments through sensors and actuators

- An agent's percept sequence is the complete history of everything the agent has ever perceived.
- *An agent's choice of action at any given instant can depend on the entire percept sequence observed to date, but not on anything it hasn't perceived.*
- Mathematically speaking, an agent's behavior is described by the **agent function** that maps any given percept sequence to an action.

At any time, what an agent does depends on its

- **Prior knowledge** about the agent and the environment;
- **History** of interaction with the environment, which is composed of
  - **Observations** of the current environment
  - **Past experiences** of previous actions and observations, or other data, from which it can learn;
- **Goals** that it must try to achieve or preferences over states of the world;
- **Abilities**, which are the primitive actions it is capable of carrying out.
-

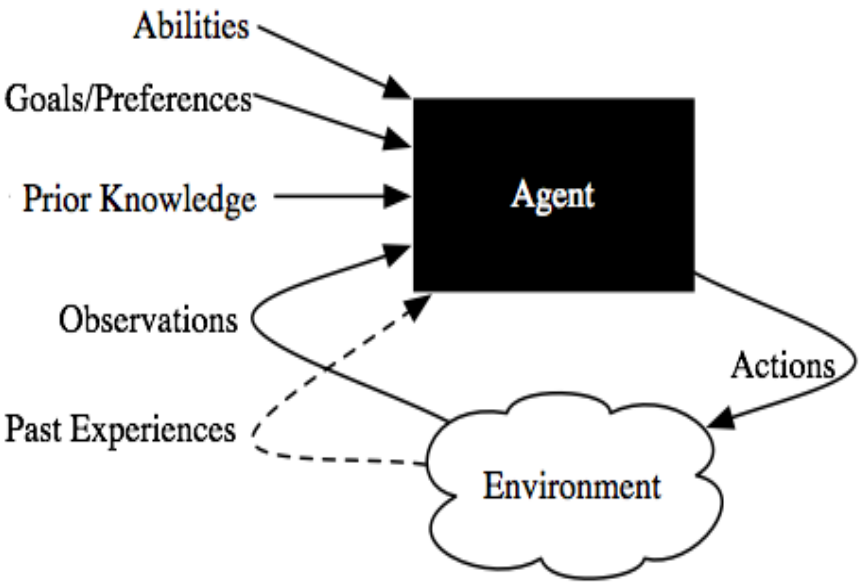

Figure 2: shows the inputs and outputs of an agent

The Concept of Rationality

- An agent should strive to "do the right thing", based on what it can perceive and the actions it can perform. The right action is the one that will cause the agent to be most successful

- Therefore, 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

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

- A ***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.

Definition of a 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.

PEAS:

- **PEAS:** Performance Measure, Environment, Actuators, Sensors. In the development of an agent
- we must first specify the setting (PEAS) for intelligent agent design
- Consider a task of designing an automated taxi driver
    - Performance Measure
    - Environment
    - Actuators
    - Sensors

Example: Consider a task of designing an automated taxi driver

- **Performance Measure:** Safe, fast, legal, comfortable trip, maximize profits
- **Environment:** Roads, Other Traffic, Pedestrians, Customers
- **Actuators:** Steering, Accelerator, Brake, Signal, Horn, Display
- **Sensors:** Cameras, Sonar, Speedometer, GPS, odometer, Accelerometer, Engine sensors, keyboard

PEAS Description an Automated Taxi

- **The performance measure to which we would like our automated driver** to aspire? Desirable qualities include getting to the correct destination; minimizing fuel consumption and wear and tear; minimizing the trip time or cost; minimizing violations of traffic laws and disturbances to other drivers; maximizing safety and passenger

comfort; maximizing profits. Obviously, some of these goals conflict, so tradeoffs will be required.

- **The environment:** Any taxi driver must deal with a variety of roads. The roads contain other traffic, pedestrians, stray animals, road works, police cars, puddles, and pot-holes. The taxi must also interact with potential and actual passengers.

- **The actuators** include those available to a human driver: control over the engine through the accelerator and control over steering and braking, display screen or voice synthesizer to talk back to the passengers, and perhaps some way to communicate with other vehicles, politely or otherwise.

- **The basic sensors**

    - one or more controllable video cameras so that it can see the road; it might augment these with infrared or sonar sensors to detect distances to other cars and obstacles.

    - To avoid speeding tickets, the taxi should have a speedometer, it should have an accelerometer.

    - To determine the mechanical state of the vehicle, it will need the usual array of engine, fuel, and electrical system sensors. Like many human drivers, it might want a global positioning system (GPS) so that it doesn't get lost.

    - Finally, it will need a keyboard or microphone for the passenger to request a destination.

Properties of task environments

- The range of task environments that might arise in AI is obviously vast. We can, however, identify a fairly small number of dimensions along which task environments can be categorized. These dimensions determine, to a large extent, the appropriate agent design and the applicability of each of the principal families of techniques for agent implementation.
    - Fully observable vs. partially observable
    - Single agent vs multi agent
    - Deterministic vs stochastic
    - Episodic vs sequential
    - Static vs dynamic
    - Discrete vs continuous

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.
- A task environment is effectively fully observable if the sensors detect all aspects that are *relevant* to the choice of action;
- 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.
- E.g, an automated taxi cannot see what other divers are thinking If the agent has no sensors at all then the environment is **unobservable.**

Single agent vs multi agent

- 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.
- Does an agent A (the taxi driver for example) have to treat an object *B* (another vehicle) as an agent. or can it be treated merely as an object behaving according to the laws of physics, analogous to waves at the beach or leaves blowing in the wind? The key distinction is whether B's behavior is best described as maximizing a performance measure whose value depends on agent A's behavior.

- For example, in chess, the opponent entity *B* is trying to maximize its performance measure, which, by the rules of chess, minimizes agent *A's* performance measure. Thus, chess is a **competitive** 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 env't. If the env't 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.
- Taxi driving is clearly stochastic, because one can never predict the behavior of traffic exactly; moreover, one's tires blow out and one's engine seizes up without warning

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.
- The next episode does not depend on the actions taken in previous episodes.
- In sequential environments, 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.
- 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; 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.
- Dynamic environments, are continuously asking the agent what it wants to do; if it hasn't decided yet that counts as deciding to do nothing.

- Taxi driving is dynamic: other cars and the taxi keep moving while the driving algorithm dithers about what to do next.
- Chess, when played with a clock, is semi-dynamic. Crossword puzzles are static.

Discrete vs. Continuous

- The discrete or continuous distinction applies to the state of the environment, to the way time is handled, and to the percepts and actions of the agent.
- For example, the chess environment has a finite number of distinct states, Chess also has a discrete set of percepts 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. Taxi-driving actions are also continuous (steering angles, etc.).

Types of Agents

- There are four basic kinds of agent programs that embody the principles underlying almost all intelligent systems. Each kind of agent program combines particular components in particular ways to generate actions:
    - Simple reflex agents;
    - Model-based reflex agents;
    - Goal-based agents; and
    - Utility-based agents.

Read and briefly explain the four types of agents listed above. (Give a brief explanation and an example) to be submitted next week. (maximum 2 pages)