

Artificial Intelligence

Joint class 5314 / 4320

Agents

1 What is an agent?

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

Agents perceive the environment through sensors, process the perceived information, and acts accordingly on the environment.

	Sensors	Actuators
Human	eyes, hands, legs, nose, etc.	hands, legs, voice, etc.
Robotic agent	cameras, infrared range finders, etc.	various actuators
Software agent	keystrokes, file contents, network packets	displaying on the screen, writing files, sending network packets, etc.

Perception → Action

Action is taken according to the sequence of perceptual inputs. The **agent function** is what maps any given percept sequence to an action.

2 Behavior of agents

As mentioned in the first lecture, a **rational agent** is one that does things right. But what does this mean to do things right? Maybe the one that will cause the agent to be most successful.

“Most successful” → we need a way to measure success.

2.1 Performance measures

= criterion for success of an agent’s behavior

There is not one measure suitable for all agents... In general such a measure, the objective (to be maximized, or a criterion to be satisfied), is defined ad hoc by the designer, but is very important to provide some control.

2.2 Rationality

Rationality 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 sequence of perceptions to date

This leads to a possible 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.

Let us note that, depending on the performance measure, the same agent may be interpreted as rational or not. *e.g. vacuum-cleaner example of the textbook.*

2.3 Omniscience, learning, and autonomy

Being **omniscient**, for an agent, means knowing all possible outcomes of its agents. It does not mean being perfect. *e.g., crossing the Champs Élysées example in the textbook.* Being rational does not imply to be omniscient either.

Learning is necessary when little is known about the environment.

Being **autonomous** means that the agent does not rely only on the a priori knowledge given by its designer. Being autonomous is important to be able to correct possibly erroneous prior knowledge.

3 Environments of agents

Before to go ahead with building agents, it is important to spend some time on defining their environment to make sure the agents fits it the best. In particular, we are going to consider what are called **task environments**, which are basically the problems to which agents are the solutions.

3.1 Specifying the task environment

PEAS (Performance measure, Environment, Actuators, Sensors)

Agent Type	Performance measure	Environment	Actuators	Sensors
Taxi driver	Safe, fast, legal, comfortable, trip, maximize profit	Roads, other traffic, pedestrians, customers	Steering, accelerator, brake, signal, horn	Cameras, sonar, speedometer, GPS, odometer, accelerometer, engine sensors, keyboards
Medical diagnosis system	Healthy patient, minimize costs, lawsuits	Patient, hospital, staff	Display 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 categorization of scene	Color pixel arrays
Interactive English tutor	Maximize student's score on test	Set of students, testing agency	Display exercises, suggestions, corrections	Keyboard entry

3.2 Properties of task environments

Even if the number of possible task environments is huge, we consider several properties and compare for each two opposite situations.

- **Fully observable vs. partially observable.**

If sensors provide all information that is relevant for taking actions, then we say that the task environment is fully observable. This is convenient because the agent does not need to keep track of the state of the world in memory, for lack of information.

An environment is partially observable when for instance the sensors are not 100 % accurate. Let us note that this is most of the time the case, but that in most of the cases however, if only the sensors are inaccurate, the environment is supposed to be fully observable.

- **Deterministic vs. stochastic.**

This is a property of the agent function. In partially observable environments, agent functions may not always be deterministic, but rather stochastic, because this is therefore difficult to keep track of all the unobserved aspects of the world.

- *e.g., taxi driving is stochastic for this is difficult to predict the behavior of traffic*
- *e.g., in a multi-player game, agent function is necessarily stochastic for this is impossible to always predict the behavior of the players*

- **Episodic vs. sequential.**

In an episodic task environment, the agent's experience can be divided into atomic episodes. Each episode consists of the agent perceiving and then performing a single action. The main property is that the next episodes does not depend on the actions taken in the previous ones.

e.g., classification tasks in general: agents responsible for detecting defective parts on an assembly line base their decision on the current part, regardless of previous decisions

In sequential environments, the current decision could affect all future decisions.
chess and taxi driving

Of course, episodic environments are far easier to manage than sequential ones, because the agent does not need to think ahead.

- **Static vs. dynamic.**

The environment is dynamic when it is likely to change when the agent is determining which action to take, making this action possibly obsolete by the time it is decided. If the environment is not likely to change, but the performance measure is, then we say that the environment is **semi-dynamic**.

e.g., taxi driving = dynamic

chess = semi-dynamic

crossword puzzle = static

- **Discrete vs. continuous.**

Discrete if there is a finite number of possible states. Continuous otherwise.

e.g., taxi driving is continuous: speed, location, etc.

- **Single agent vs. multiagent.** The definition is straightforward.

e.g., crossword = single-agent

game playing = multiagent

Let us add that multiagent systems can be cooperative (*driving: everybody try to avoid collisions*) or competitive (*game playing*). In addition, in multiagent systems, communication between agents has also to be taken into account. Information does not come any longer from sensors only.

4 The structure of agents

4.1 Agent programs

4.2 Simple reflex agents

4.3 Model-based reflex agents

4.4 Goal-based agents

4.5 Utility-based agents

4.6 Learning agents

5 Assignments

1. Read the slides about an introduction to AI, available on my website <http://www.cs.utep.edu/mceberio>
2. Read Chapter 1 of your textbook.
3. Exercise 1.7 of your textbook (p. 31): choose four of the points and answer them.

Contents

1	What is an agent?	1
2	Behavior of agents	1
3	Environments of agents	2
4	The structure of agents	5
5	Assignments	5
