



مجموعة جامعة
الجميع الذكية
للتعليم القابضة



Master of Science in Artificial intelligent

Lecture 2

AI511

Foundations of Artificial Intelligence

Prepared by
Prof. Lamiaa Fattouh Ibrahim

www.midocean.ae

Outline

Intelligent “Rational” Agents?

- AI as the Study & Design of Intelligent Agents
 - Intelligent Agents in the World
 - Example: Vacuum _ Agent
 - Type of Agents
- Specifying the task environment
 - Environment types

Acting Rationally
Intelligent
Agents



Main Resources for this Section

- This lecture covers mainly the following chapter(s):
 - **Chapter 2 (Intelligent Agents)** from Stuart J. Russell and Peter Norvig, "Artificial Intelligence: A Modern Approach," Third Edition (2010), by Pearson Education Inc.

What is Artificial Intelligence?


What is Artificial Intelligence?

Recap ..

Four Main Approaches that have been followed, each by different people with different methods.

the Cognitive Modelling approach

**Thinking
Humanly**



the Turing Test approach

**Acting
Humanly**



**Thinking
Rationally**



the Laws-of-Thought approach

**Acting
Rationally**



the Rational Agent approach

Recap: What is Artificial Intelligence?

Systems that **act** like humans

The study of how to make computers do things at which, at the moment, people are better" (Rich and Knight, 1991)

"The art of creating machines that perform functions that require intelligence when performed by people." (Kurzweil, 1990)

Systems that **think** like humans

The automation of activities that we associate with human thinking, such as decision making, problem solving, learning" (Bellman, 1978)

"The exciting new effort to make computers think ... machines with minds, in the full and literal sense." (Haugeland, 1985)

Systems that **think** rationally

The study of mental faculties through the use of computational models" (Charniack and McDermott, 1985).


"The study of the computations that make it possible to perceive, reason, and act." (Winston, 1992)

Systems that **act** rationally

"AI .. is concerned with intelligent behavior in artifacts (Nilsson, 1998)

"Computational Intelligence is the study of the design of intelligent agents." (Poole et al., 1998)

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




Recap: AI as the Study & Design of Intelligent Agents

(Poole and Mackworth, 1999)

- An *intelligent agent* is such that:
 - Its *actions* are *appropriate* for its goals and circumstances.
 - It is *flexible* to changing environments and goals.
 - It *learns* from experience.
 - It makes *appropriate choices* given *perceptual limitations* and *limited resources* (bounded rationality or bounded optimality).

Thus, a **rational agent** acts to optimally achieve its goals (does the right thing). The right thing: that which is expected to maximize goal achievement, given the available information



- It's all about how the control policy function of this agent takes the decision


Agent	Sensors (inputs)	Actuators
Human Agent	eyes, ears, and other organs	hands, legs,
robotic agent	cameras and infrared range finders	various motors
software agent	keystrokes, file contents, and network packets	displaying on the screen, writing files, and sending network packets.


THE STRUCTURE OF AGENTS


The job of AI is to design an **agent program** that implements the **agent function**— the mapping from 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**:


Structure of Agent = architecture + agent program


Application of AI















Intelligent Agents in the World..

Knowledge Representation
Machine Learning

.. Capacities/ Fields?




Example: Vacuum – Agent

Percepts:

Location and status,
e.g., [A, Dirty]

Actions:

Left, Right, Suck, NoOp



```
function 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
```

Rational Agents

- 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 the agent's built-in knowledge*
- **Performance measure (utility function):**
An objective criterion for success of an agent's behavior.

Specifying the task environment

PEAS: Performance measure, Environment, Actuators, Sensors

P: a function the agent is maximizing (or minimizing);

Assumed given ...

In practice, needs to be computed somewhere.

E: a formal representation for *world states*;

For concreteness, a tuple $(var_1 = val_1, var_2 = val_2, \dots, var_n = val_n)$.

A: actions that change the state according to a *transition model*;

Given a state and action, what is the successor state (or distribution over successor states)?

S: observations that allow the agent to infer the world state;

Often come in very different form than the state itself ...

E.g., in tracking, observations may be pixels and state variables 3D coordinates.

PEAS Example: Autonomous taxi

- Performance measure

Safe, fast, legal, comfortable trip, maximize profits

- Environment

Roads, other traffic, pedestrians, customers

- Actuators

Steering wheel, accelerator, brake, signal, horn

- Sensors

Cameras, LIDAR, speedometer, GPS, odometer, engine sensors, keyboard

Another PEAS example: Spam filter

- **Performance measure**

Minimizing false positives, false negatives

- **Environment**

A user's email account, email server

- **Actuators**

Mark as spam, delete, etc.

- **Sensors**

Incoming messages, other information about user's account


Agent types

FIVE basic types in order of increasing generality:

- Simple reflex agents
- Model-based reflex agents
- Goal-based agents
- Utility-based agents
- Learning agents

Simple reflex agents


- Selects action based on current perception of the environment.
- Simple but has a limited intelligence.




Reflex agents with state

Model-based reflex agents

Keeps track of perception history in the form of internal state (reflecting unobserved aspects of the current state of the World).




Goal-based agents




Problem solving by searching


A puzzle search tree



Utility-based agents



Learning agents



Goal-based Agents versus Cost-based Agents

Goal-based agents: **the actions depend on the goal;** *E.g., a mobile robot which should move from room 112 to room 179 in a building takes actions different from those of a robot that should move to room 105.*

Goal-based Agents versus Cost-based Agents

Cost-based agents: the goal is to minimize the cost of erroneous decisions in the long term; E.g., a spam filter is an agent that puts incoming emails into wanted or unwanted (spam) categories & deletes any unwanted emails. Its goal as a goal-based agent is to put all emails in the right category. In the course of this not-so-simple task, the agent can occasionally make mistakes. Because its goal is to classify all emails correctly, it will attempt to make as few errors as possible. However, that is not always what the user has in mind. Let us compare the following two agents. Out of 1,000 emails, Agent 1 makes only 12 errors. Agent 2 on the other hand makes 38 errors with the same 1,000 emails. Is it therefore worse than Agent 1? The errors of both agents are shown in more detail in the following confusion matrix:

Agent 1:

		correct class	
		wanted	spam
spam filter decides	wanted	189	1
	spam	11	799

Agent 2:

		correct class	
		wanted	spam
spam filter decides	wanted	200	38
	spam	0	762

Agent 1 in fact makes fewer errors than Agent 2, but those few errors are severe because the user loses 11 potentially important emails. Because there are in this case two types of errors of differing severity, each error should be weighted with the appropriate cost factor.

Environment Types

Fully Observable	vs.	Partially Observable
Deterministic	vs. (vs. Strategic)	Stochastic
Episodic	vs.	Sequential
Static	vs. (vs. Semi-Dynamic)	Dynamic
Discrete	vs.	Continuous
Single-Agent	vs.	Multi-Agent
Known	vs.	Unknown

Fully Observable vs. Partially Observable

Do the agent's sensors give it access to the complete state of the environment?

For any given world state, are the values of all the variables



Deterministic vs. Stochastic

- Is the next state of the environment completely determined by the current state and the agent's action?
- Is the transition model deterministic (unique successor state given current state and action) or stochastic (distribution over successor states given current state and action)?
- **Strategic:** the environment is deterministic except for the actions of other agents.



VS




Episodic vs. Sequential

Is the agent's experience divided into unconnected single decisions/actions, or is it a coherent sequence of observations and actions in which the world evolves according to the transition model?




VS.



Static vs. Dynamic

Is the world changing while the agent is thinking?

Semidynamic: the environment does not change with the passage of time, but the agent's performance score does.




VS.




Discrete vs. Continuous

Does the environment provide a fixed number of distinct percepts, actions, and environment states?

Are the values of the state variables discrete or continuous? Time can also evolve in a discrete or continuous fashion



VS.




Single-Agent vs. Multi-Agent

Is an agent operating by itself in the environment?



vs.



Known vs. Unknown

Are the rules of the environment (*transition model and rewards associated with states*) known to the agent?

Strictly speaking, not a property of the environment, but of the agent's state of knowledge.




VS.



- **Deterministic:** Deterministic environment is one where your agent's actions uniquely determine the outcome.
- **Stochastic:** In stochastic environment, there is certain amount of randomness.

- **Discrete v.s. continues:** A discrete environment is one where you have finitely many action choices, and finitely many things you can sense. For example, in chess there's finitely many board positions, and finitely many things you can do.

Examples of the different environments



	Word jumble solver	Chess with a clock	Scrabble	Autonomous driving
Observable	Fully	Fully	Partially	Partially
Deterministic	Deterministic	Strategic	Stochastic	Stochastic
Episodic	Episodic	Sequential	Sequential	Sequential
Static	Static	Semi-dynamic	Static	Dynamic
Discrete	Discrete	Discrete	Discrete	Continuous
Single agent	Single	Multi	Multi	Multi

Summary

- AI as the Study & Design of Intelligent Agents
 - Intelligent Agents in the World
 - Example: Vacuum _ Agent
 - Type of Agents
- Specifying the task environment
 - Environment types



Thank you