

Intelligent Agents, & AI Related Disciplines

Dr. Mohamed Mostafa

AI 511 Foundations of Artificial Intelligence

Lecture 2



2.1 Intelligent “Rational” Agents?

- Recap: AI as the Study & Design of Intelligent Agents
- Intelligent Agents in the World
- An Example: Vacuum_Agent
- Specifying the Task Environment [PEAS]
- Goal-based vs. Cost-based Agents

2.3 AI: Related Disciplines

- Learning Agents
- AI vs. Machine Learning vs. Deep Learning
- Machine Learning?
- Datamining?
- AI vs. Data Science
- Why Data Science?
- Why Big Data?

Resources for this
lecture


- This lecture covers the following chapters:
 - **Chapter 2 (Intelligent Agents)** from Stuart J. Russell and Peter Norvig, "Artificial Intelligence: A Modern Approach," Third Edition (2010), by Pearson Education Inc.
- .. AND ..
 - Introduction to Artificial Intelligence Course,. Dr. Amr S. Ghoneim

What is Artificial Intelligence?

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**



**Acting
Rationally**



the Laws-of-Thought approach

*the Rational Agent
approach*

What is Artificial Intelligence?

Recap ..

Systems that act like humans	Systems that think rationally
“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)	“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 think like humans	Systems that act rationally
“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)	“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)

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**).
- This definition drops the constraint of *cognitive plausibility*;
 - Same as building flying machines by understanding general principles of flying (aerodynamic) vs. by reproducing how birds fly.

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.

- In AI, artificial agents that have a physical presence in the world are usually known as **Robots**.
- Robotics is the field primarily concerned with the implementation of the physical aspects of a robot (*i.e. perception of the physical environment, actions on the environment*).
- Another class of artificial agents include **interface agents**, for either stand alone or Web-based applications (*e.g. intelligent desktop assistants, recommender systems, intelligent tutoring systems*).
- Interface agents don't have to worry about interaction with the physical environment, but share all other fundamental components of intelligent behavior with robots.
- We will focus on these agents in this course (*i.e. software agents not hardware agents "autonomous robots"*).


Pac-Man .. *as an* .. Intelligent Agent

Recap ..



Intelligent Agents

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



YESTERDAY DR. WILL CASTER WAS ONLY HUMAN.

Dr. Will Caster

JOHNNY DEPP REBECCA HALL PAUL BETTANY MATT MOSS JULIAN MARICA MORGAN FREEMAN

TRANSCENDENCE

CASTING BY JAMES NEWTON HOWARD COSTUME DESIGNER JAMES NEWTON HOWARD EXECUTIVE PRODUCERS JAMES NEWTON HOWARD PRODUCED BY JAMES NEWTON HOWARD WRITTEN BY JAMES NEWTON HOWARD DIRECTED BY JAMES NEWTON HOWARD

www.transcendencemovie.com

SCHWARZENEGGER

In the Year of Darkness, 2029, the future of this planet hangs on the ultimate prize: The survival of mankind. The future is by changing the Past.


The great computer, something that can do anything, has been something unstoppable. They created THE TERMINATOR.

THE TERMINATOR.

Presented by Pacific Western Productions, a division of Warner Bros. Entertainment, "The Terminator" Richard Matheson, Linda Hamilton, and Paul Winfield by James Cameron. Director of Photography: Michael Ballhaus. Screenplay by James Cameron. Produced by James Cameron. Produced by Gale Anne Hand. Directed by James Cameron. Released by P. Buick Film Distributors. (TM) 1984 Warner Bros.

Agents?

BB-8 “the Star Wars franchise”

A collage featuring two robots. On the left is BB-8, a small, white and orange spherical droid from the Star Wars franchise. On the right is WALL-E, a boxy, yellow and white robot from the Pixar movie WALL-E. WALL-E is holding a small, multi-colored cube in his hand. The background is a dark, textured surface.

Arthur, an android bartender on the *Avalon*, *Passengers*



Sonny in I, Robot (2004)

Sonny in I, Robot (2004)

Intelligent Rational Agents?



TARS and CASE
Interstellar



Vision
(Marvel
Comics)




**L3 (Star Wars –
Han Solo)**



Ava,
in *Ex Machina*


Intelligent Agents in the World *Learning?* .. *Feedback?*



Intelligent Agents in the World ..


.. Capacities / Fields?

Knowledge Representation
Machine Learning



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 behaviour.

Performance measure (utility function)

$$\text{Expected Utility(action)} = \sum_{\text{outcomes}} [\text{Utility(outcome)} * P(\text{outcome}|\text{action})]"$$

- represents a formula used in **decision theory** to calculate the expected utility of taking a specific action.

Action:

Consider an individual who is deciding whether to . invest in a particular stock

Outcomes:

There are various potential outcomes associated with this investment, such as the **stock price increasing, staying stable, or decreasing.**

Utility(outcome):

The utility function assigns a numerical value to each outcome, representing the individual's subjective satisfaction or preference. For instance, gaining a profit might have a positive utility, while a loss might have a negative utility.

P(outcome | action):

This represents the probability of each outcome occurring given that the individual takes the specific action (investing in the stock). The probabilities are subjective estimates or based on historical data.



Performance measure (utility function)

$$\text{Expected Utility(action)} = \sum_{\text{outcomes}} [\text{Utility(outcome)} * P(\text{outcome}|\text{action})]$$

Let's say the individual assigns the following utilities and probabilities:

- If the stock price increases (Outcome A):
 - $\text{Utility(A)} = +10$
 - $P(A|\text{action}) = 0.6$
- If the stock price stays the same (Outcome B):
 - $\text{Utility(B)} = +2$
 - $P(B|\text{action}) = 0.3$
- If the stock price decreases (Outcome C):
 - $\text{Utility(C)} = -8$
 - $P(C|\text{action}) = 0.1$

Using the formula: $\text{Expected Utility(action)} = (10 \times 0.6) + (2 \times 0.3) + (-8 \times 0.1)$

$\text{Expected Utility(action)} = 6 + 0.6 - 0.8 = 5.8$ **Expected Utility(action) = 6 + 0.6 - 0.8 = 5.8 +ve utility**



Specifying the Task Environment [PEAS]

- **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 1: 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

PEAS Example 2: 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

Goal-based Agents versus Cost-based Agents

- **Goal-based agents:** **the actions depend on the goal;** *E.g.,*

Autonomous Delivery Drone

- **Goal:** Deliver packages to specified locations.
- **Actions:** Fly to destination, avoid obstacles, optimize routes.
- **Utility:** Successfully delivering packages contributes to goal achievement.



Goal-based Agents versus Cost-based Agents

- **Cost-based agents:** **the goal is minimize costs associated with actions.**
- Costs can include time, resources, energy, or any other measurable factor ; *E.g.,*

Ride-Sharing Service

- **Goal:** Efficiently transport passengers from pick-up to drop-off.
- **Actions:** Choose optimal routes, minimize travel time, conserve fuel.
- **Cost:** Time, fuel consumption, vehicle wear and tear.



Goal-based Agents versus Cost-based Agents

- **Comparison:**
- **Goal-Based Agents:**
 - Emphasize achieving specific objectives or outcomes.
 - Assess actions based on their contribution to goal attainment.
 - Appropriate for scenarios where the focus is on accomplishing predefined objectives.
- **Cost-Based Agents:**
 - Emphasize minimizing costs or resource usage.
 - Assess actions based on the efficiency of resource utilization.
 - Suitable for scenarios where resource conservation and efficiency are paramount.

2.1 Intelligent “Rational” Agents?

- Recap: AI as the Study & Design of Intelligent Agents
- Intelligent Agents in the World
- An Example: Vacuum_Agent
- Specifying the Task Environment [PEAS]
- Goal-based vs. Cost-based Agents

2.2 Specifying the Task Environment

- Environment Types

2.3 AI: Related Disciplines

- Learning Agents
- AI vs. Machine Learning vs. Deep Learning
- Machine Learning?
- Datamining?
- AI vs. Data Science
- Why Data Science?
- Why Big Data?

Environment Types


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

Fully Observable vs. Partially Observable

- In a fully observable environment, the agent can access complete and accurate information about the current state at any given time.
- The agent has full knowledge of the environment's current conditions, and there is no uncertainty in its perception.
- In a partially observable environment, the agent does not have complete access to the current state of the environment.
- The agent's perception is limited, and it may need to infer or estimate the current state based on partial or incomplete information.



VS.



Deterministic vs. Stochastic (vs. Strategic)

- - In a deterministic environment, the next state of the environment is entirely determined by the current state and the actions taken by the agent. There is no randomness or uncertainty in the outcomes.
- In a stochastic environment, there is inherent randomness or uncertainty in the outcomes, even if the agent takes the same actions in the same state.
- - The next state is not entirely predictable, and probabilities are associated with different outcomes.



VS.



Deterministic vs. Stochastic (vs. Strategic)

- **Strategic:** the environment is deterministic except for the actions of other agents.
- The outcomes are influenced not only by the agent's actions but also by the actions of other entities or agents in the environment.
- Agents need to consider the strategies of other entities when making decisions.



- [illegible]

VS.



Driving a car in traffic

Static vs. Dynamic (vs. Semi-dynamic)

Static Environment: Elements remain constant unless explicitly modified (e.g., chessboard).

Dynamic Environment: Elements change autonomously over time (e.g., autonomous driving in a city).

Semi-Dynamic Environment: Combination of static and dynamic elements (e.g., online marketplace with product listings and user interactions).

Discrete vs. Continuous

- In a discrete environment, the set of possible states and actions is finite and countable. The agent perceives and interacts with distinct, separate states.
 - **Example:** Consider a chess game. The chessboard has a finite number of squares, and each piece can occupy a specific square. The possible moves and positions are well-defined, making the environment discrete.
- In a continuous environment, the set of possible states and actions is infinite and uncountable. The agent deals with a continuous range of possible values.
 - **Example:** An autonomous vehicle navigating through city streets operates in a continuous environment.

Discrete vs. Continuous

- **Discrete Environment Examples:**

- **Grid World Navigation:**

- **Puzzle Solving:**

- **Explanation:** Solving puzzles with a finite number of possible configurations. Each move or action leads to a new discrete state.
- **Example:** Solving a Rubik's Cube, where each turn of a face results in a new, distinct configuration.

- **Continuous Environment Examples:**

- **Robot Arm Control:**

- **Explanation:** Controlling the movement of a robot arm in a continuous space. Positions and angles can vary smoothly.
- **Example:** An industrial robot performing precise movements to assemble components on a production line.

Single-Agent vs. Multi-Agent

- Is an agent operating by itself in the environment?
- Is the environment of an autonomous taxi driver a competitive multiagent environment or a cooperative multiagent environment?



VS.



Single-Agent vs. Multi-Agent

- **Single-Agent Environment:**


- **Example:** Consider a robotic vacuum cleaner in a room. The vacuum cleaner is the single agent, and its goal is to clean the room efficiently. It navigates the space, avoids obstacles, and determines its cleaning strategy without considering interactions with other agents.

- **Multi-Agent Environment:**

- **Example:** In a traffic control system, there are multiple agents, such as individual vehicles, pedestrians, and traffic lights. Each vehicle is an agent with the goal of reaching its destination. The traffic lights act as agents coordinating the flow of traffic. The interactions among these agents impact the overall traffic flow.

Known vs. Unknown


- Known and unknown are not actually a feature of an environment, but it is an agent's state of knowledge to perform an action.
- In a known environment, the results for all actions are known to the agent. While in unknown environment, agent needs to learn how it works in order to perform an action.



Known vs. Unknown

- **Known Environment Example:** Chess-playing AI in a closed system. The rules of chess are well-defined, and the AI has complete knowledge of the current board state. It can explore all possible moves and outcomes, making it a known environment for the agent.
- **Unknown Environment Example:** Exploratory robot navigating an unfamiliar terrain. The robot may encounter obstacles, varying terrains, and unknown features. It must use its sensors and learning mechanisms to adapt to the unknown aspects of the environment.

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

Exercises

- **Identify a specific task that an AI agent could perform, such as playing a game, managing a smart home, or recommending movies. For each component of PEAS (Performance measure, Environment, Actuators, Sensors), describe how they would apply to the chosen task.**
- **Classify different environments as known, unknown, static, dynamic, or semi-dynamic. For example, consider a weather forecasting system, an autonomous drone, and an online shopping platform. Explain the characteristics of their environments and how these characteristics impact the design of intelligent agents.**
- **Choose a real-world application, such as an autonomous vehicle or a healthcare monitoring system. Define the PEAS components for the chosen application. Discuss the challenges and considerations for designing intelligent agents to operate in these complex and dynamic environments.**

Additional Resources

- Intelligent Agents Intro

https://www.youtube.com/watch?v=5dg_pou-ZHg

- Wooldridge, M. (1999). Intelligent agents. Multiagent systems: A modern approach to distributed artificial intelligence, 1, 27-73.

<https://citeseerx.ist.psu.edu/document?repid=rep1&type=pdf&doi=e6396112e1b6c7b7ecf05d84f1e46f66ed33cd06>



Thank You