Artificial Intelligence

Course Objectives:

- To understand the concept of Artificial Intelligence (AI) in the form of various Intellectual tasks
- To understand Problem Solving using various peculiar search strategies for AI
- To understand multi-agent environment in competitive environment
- To acquaint with the fundamentals of knowledge and reasoning
- To devise plan of action to achieve goals as a critical part of AI
- To develop a mind to solve real world problems unconventionally with optimality

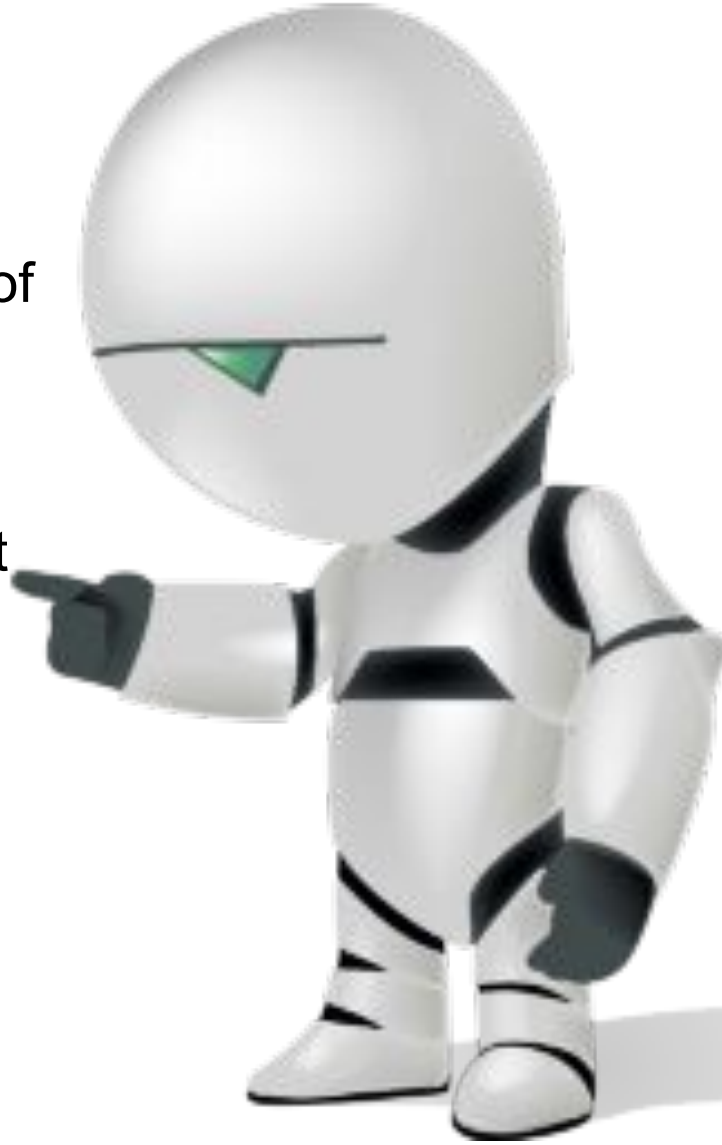

Course Outcomes:

After completion of the course, students should be able to

CO1: Identify and apply suitable Intelligent agents for various AI applications

CO2: Build smart system using different informed search / uninformed search or heuristic approaches

CO3: Identify knowledge associated and represent it by ontological engineering to plan a strategy to solve given problem

CO4: Apply the suitable algorithms to solve AI problems

CO5: Implement ideas underlying modern logical inference systems

CO6: Represent complex problems with expressive yet carefully constrained language of representation

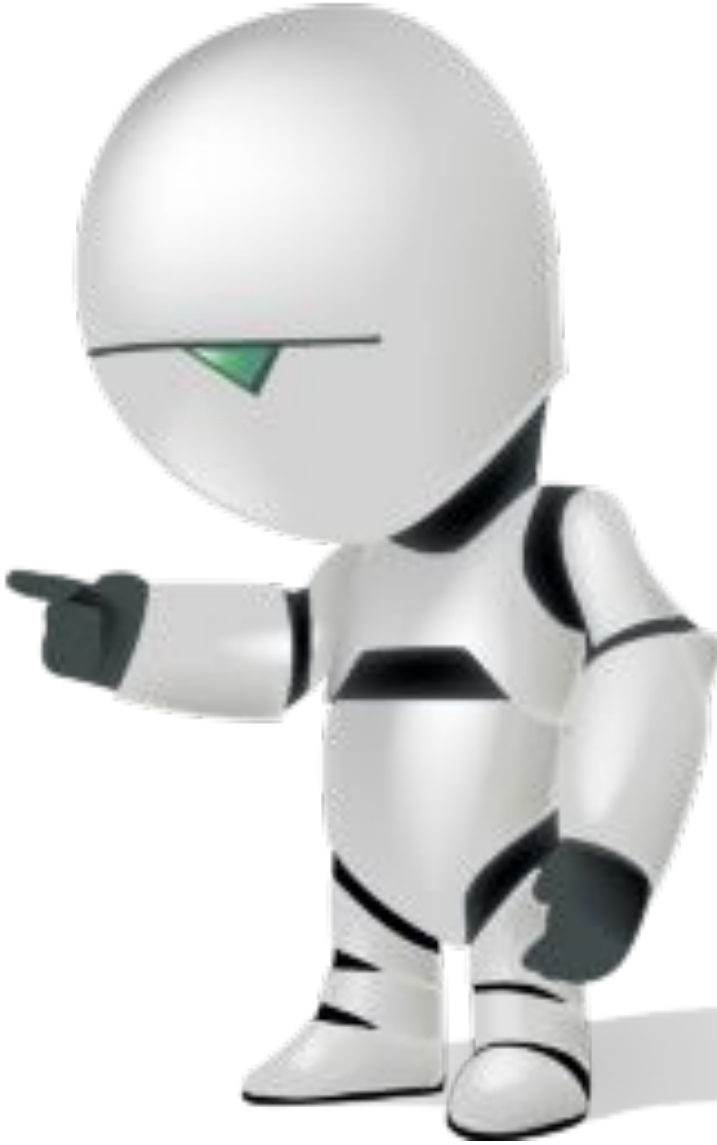

Contents

UNIT 01

Introduction

Unit I	Introduction	07 Hours
Introduction to Artificial Intelligence, Foundations of Artificial Intelligence, History of Artificial Intelligence, State of the Art, Risks and Benefits of AI, Intelligent Agents, Agents and Environments, Good Behavior: Concept of Rationality, Nature of Environments, Structure of Agents.		
#Exemplar/Case Studies	Kroger: How This U.S. Retail Giant Is Using AI And Robots To Prepare For The 4th Industrial Revolution	
***Mapping of Course Outcomes for Unit I**	CO1, CO4	

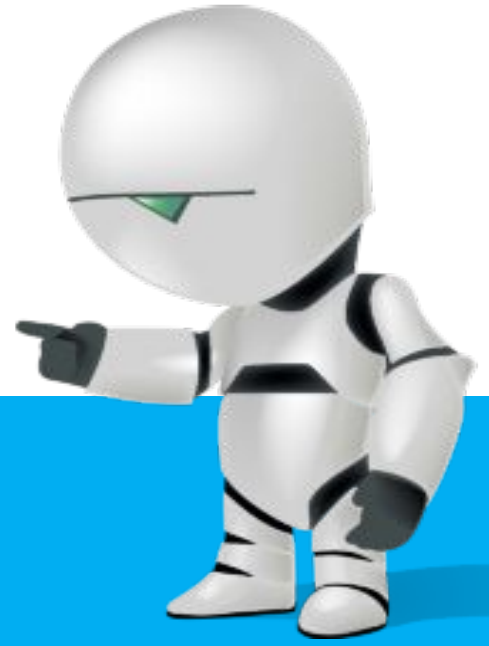

Introduction to Artificial Intelligence

- Artificial intelligence is the ability for a computer to think and learn.
  With AI, computers can perform tasks that are typically done by people, including processing language, problem-solving, and learning.
- Artificial intelligence (AI) is a branch of computer science concerned with building smart machines capable of performing tasks that typically require human intelligence.
- Artificial intelligence allows machines to model, and even improve upon, the capabilities of the human mind.
- "It is a branch of computer science by which we can create intelligent machines which can behave like a human, think like humans, and able to make decisions."
- Example:
  Digital Assistants
  Apple's Siri, Google Now, Amazon's Alexa, and Microsoft's Cortana

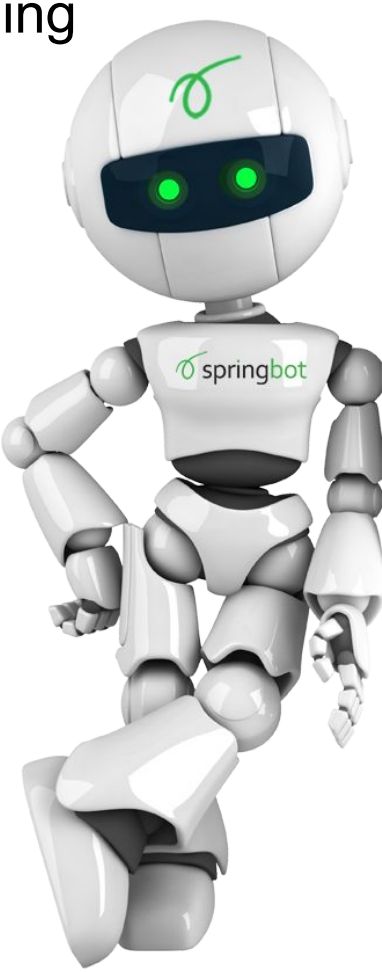

Introduction to Artificial Intelligence

AI Stands two words Artificial and Intelligence:

1. **Artificial**
   Artificial is somewhat that is not real and is simulated. Artificial is not only used in the context of food. Artificial turf is a grass like surface for sports playing fields.
2. **Intelligence**
   Defining and classifying Intelligence is extremely complicated. Intelligence is defined in many different ways like logic, understanding, self awareness, learning, planning, creativity and problem solving.

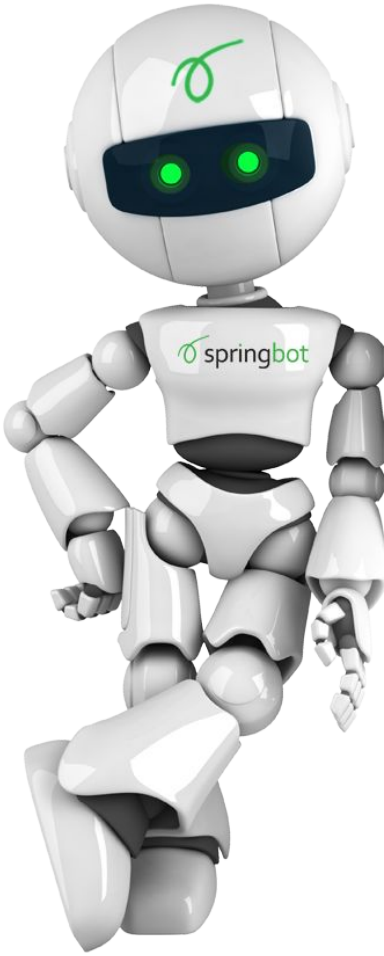

Objectives of AI (Goals of AI)

- AI Solves problems
- AI completes multiple tasks
- AI shapes the future of every company
- AI create synergy between humans and AI
- AI is good at problem-solving.
- AI helps with planning
- AI performs more complex tasks.

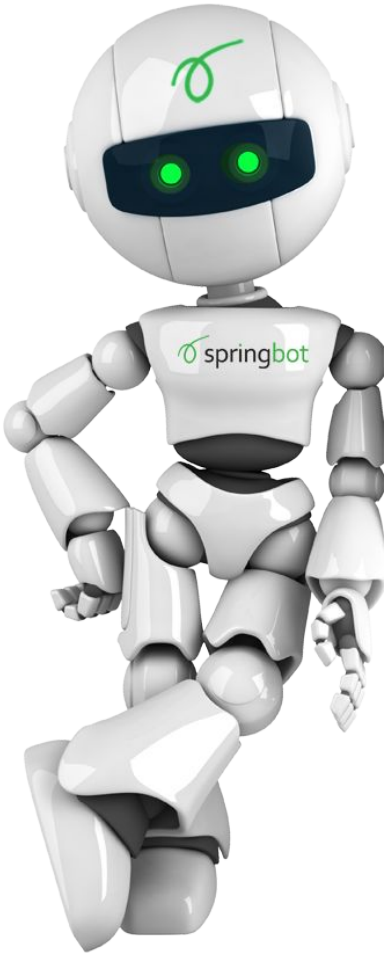

History of AI

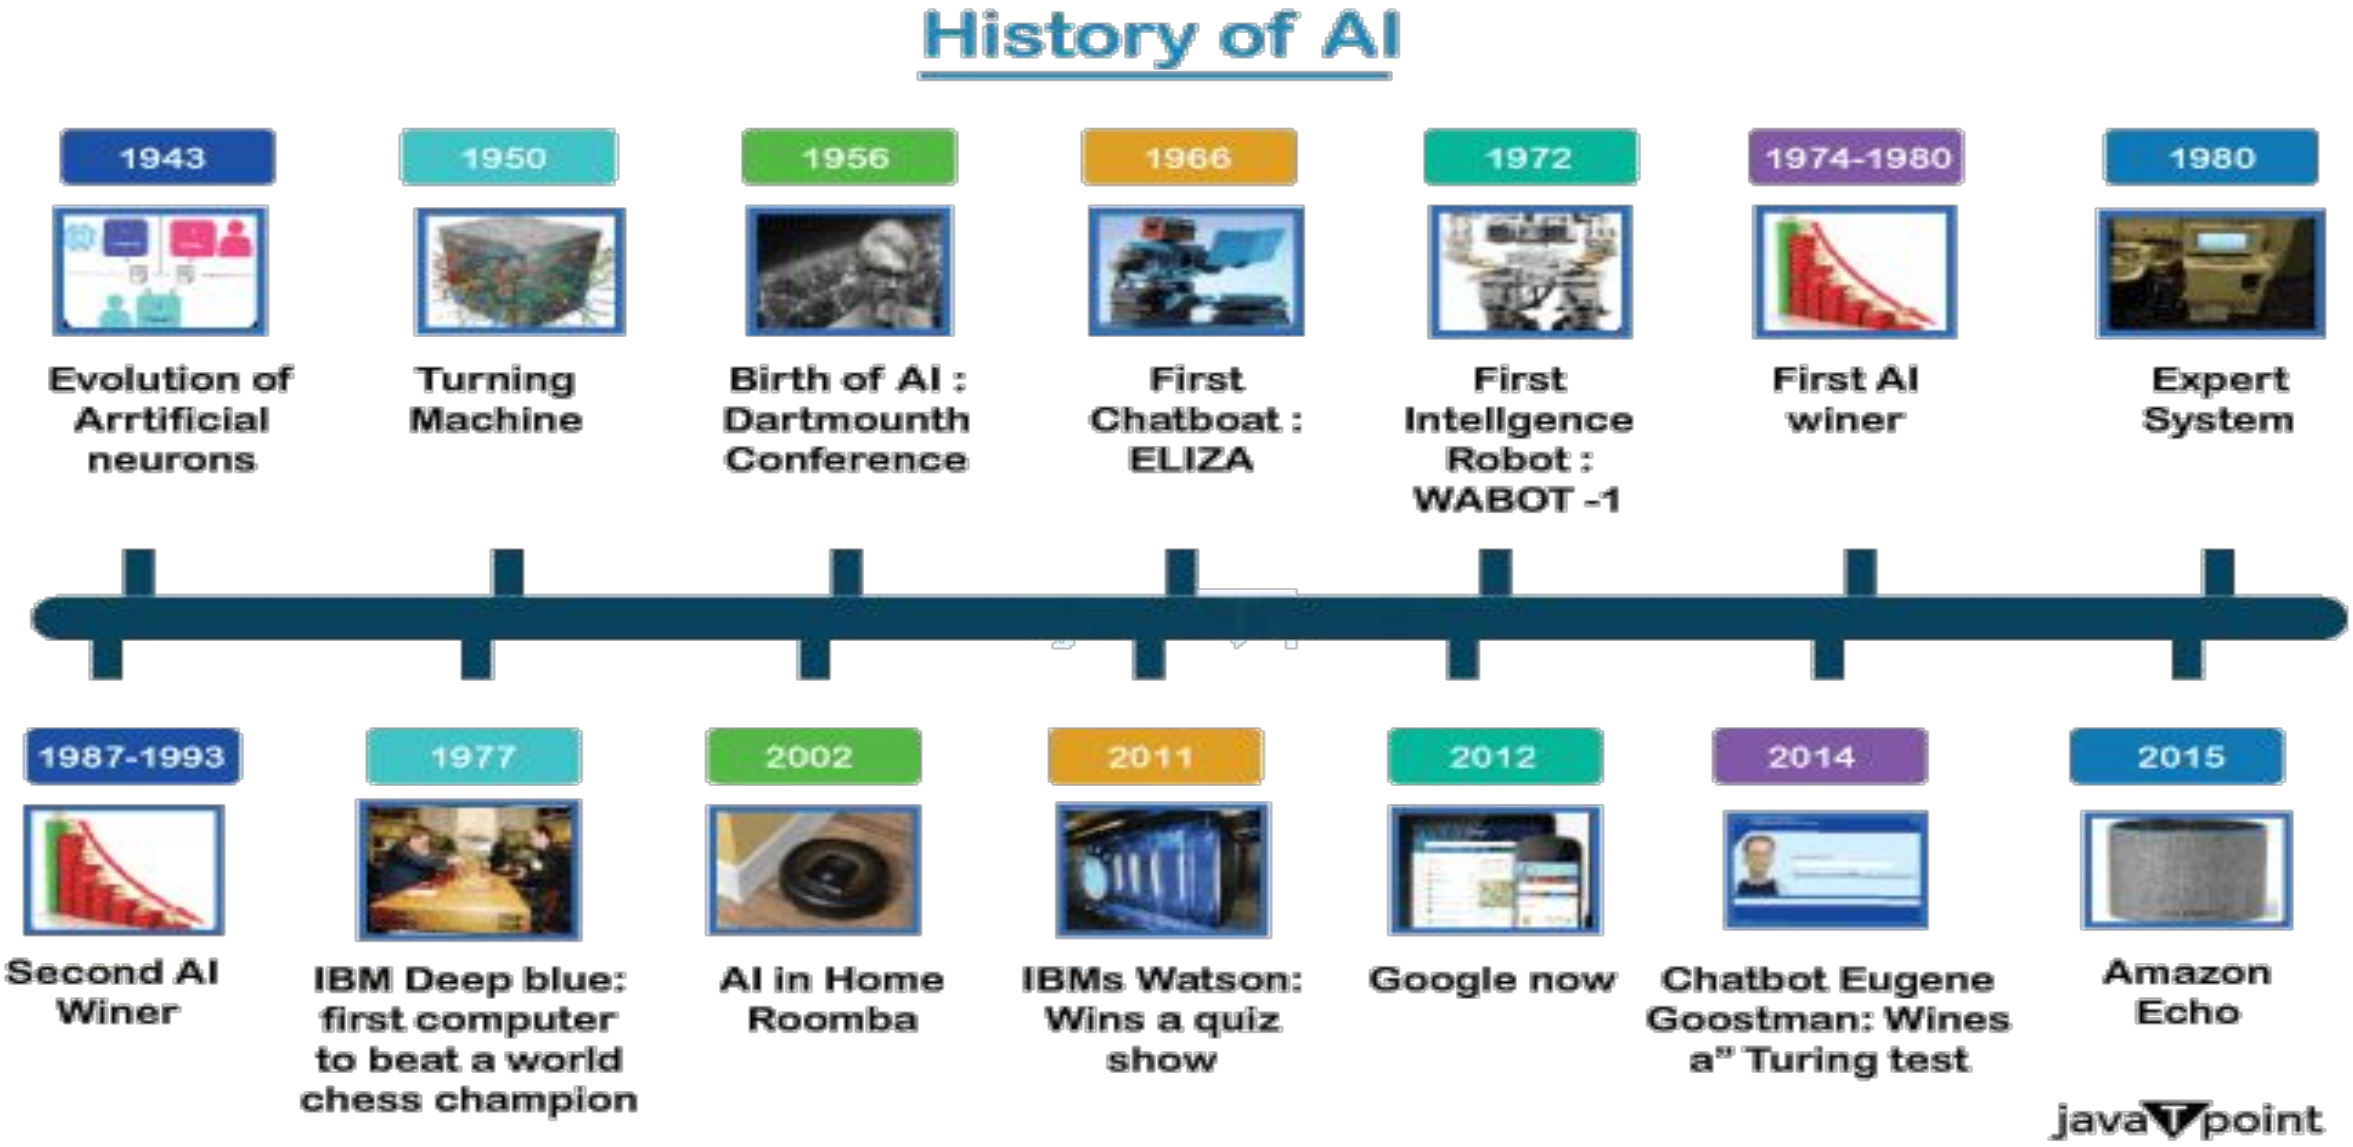

How AI is Transforming Industries?

01

Healthcare

AI-powered diagnostics, personalized treatment plans, drug discovery, remote patient monitoring, predictive analytics for disease prevention.

02

Finance

Algorithmic trading, fraud detection, credit scoring, risk management, customer service automation, personalized financial advice.

03

Transportation

Autonomous vehicles, route optimization, predictive maintenance, traffic management, smart logistics, ride-sharing algorithms.

04

Retail

Personalized recommendations, demand forecasting, inventory optimization, chatbots for customer service, virtual assistants, fraud detection.

The State of the Art

What can AI do today?

- Robotic vehicles
- Speech recognition
- Autonomous planning and scheduling
- Game playing
- Spam fighting
- Logistics planning
- Robotics
- Machine Translation

Types of AI

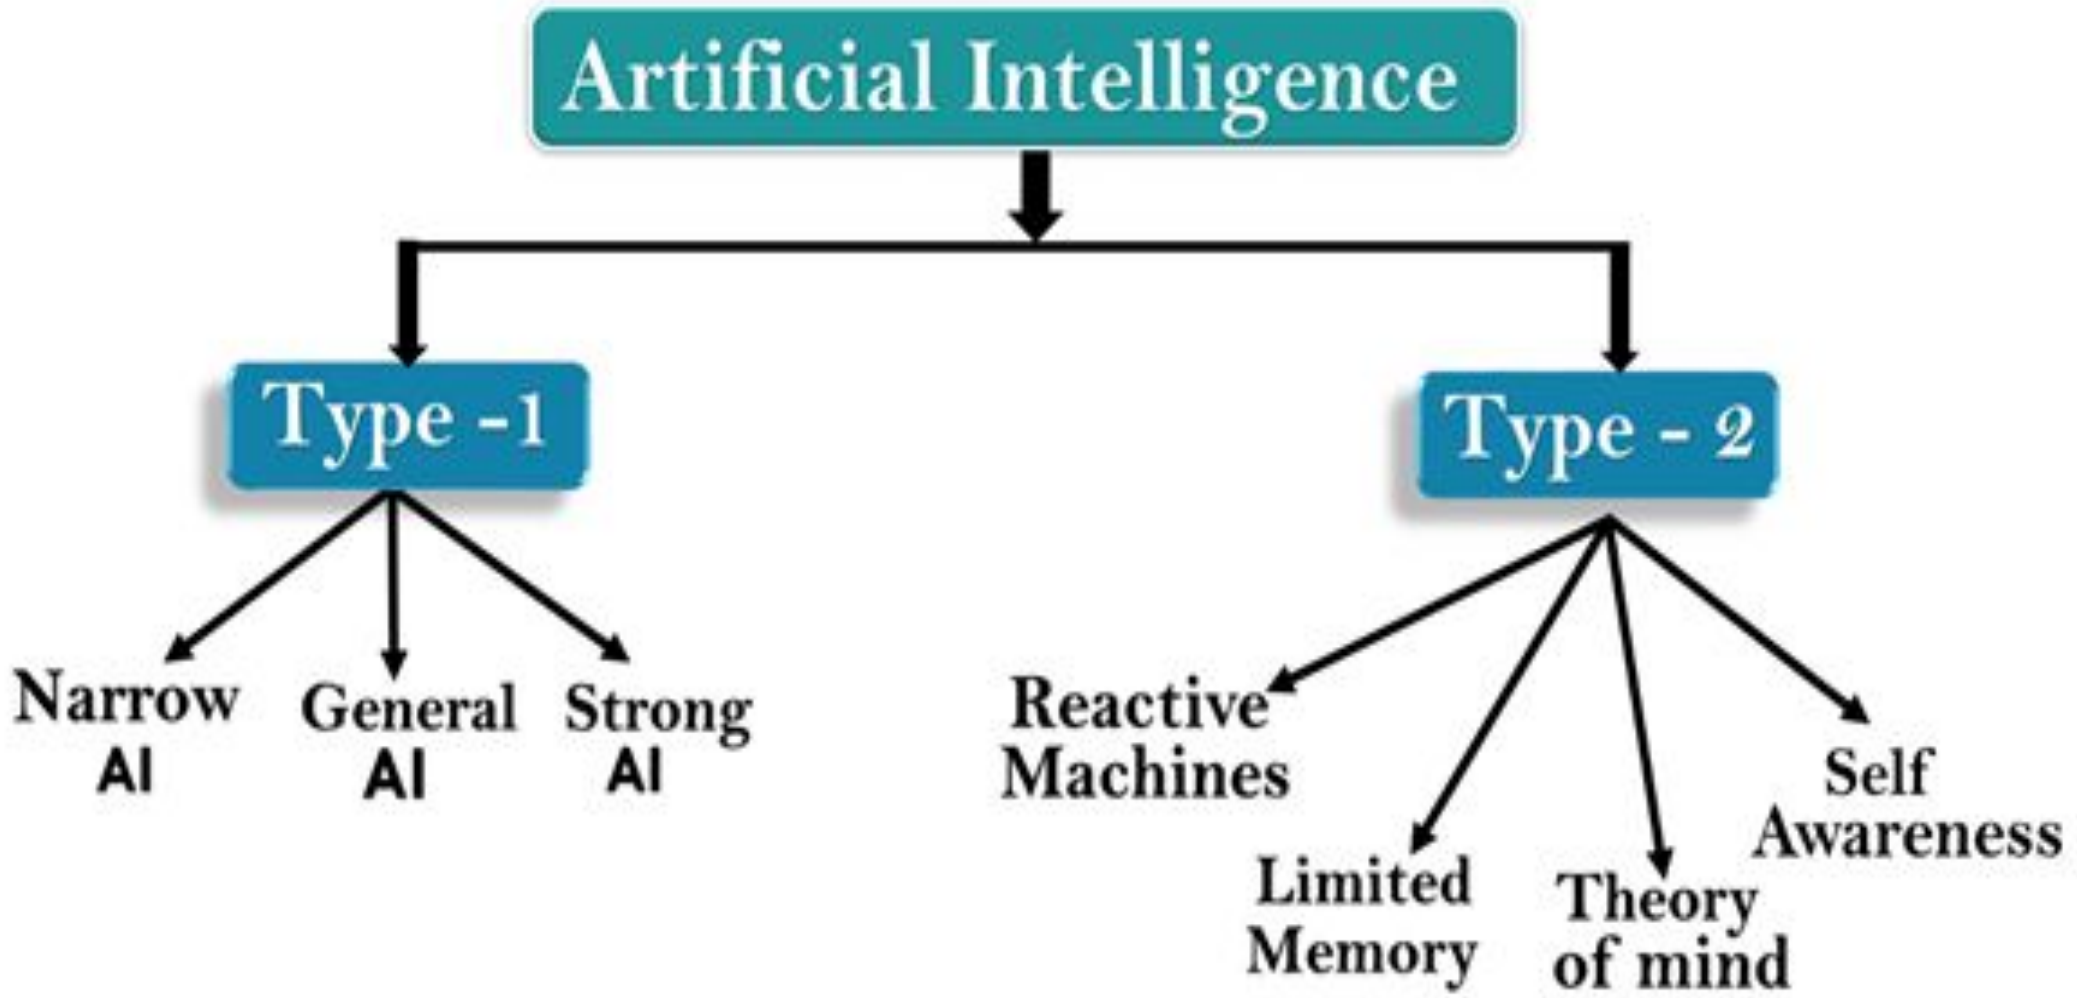
Artificial Intelligence
Type -1
Type - 2
Narrow AI
General AI
Strong AI
Reactive Machines
Limited Memory
Theory of mind
Self Awareness

Types of AI

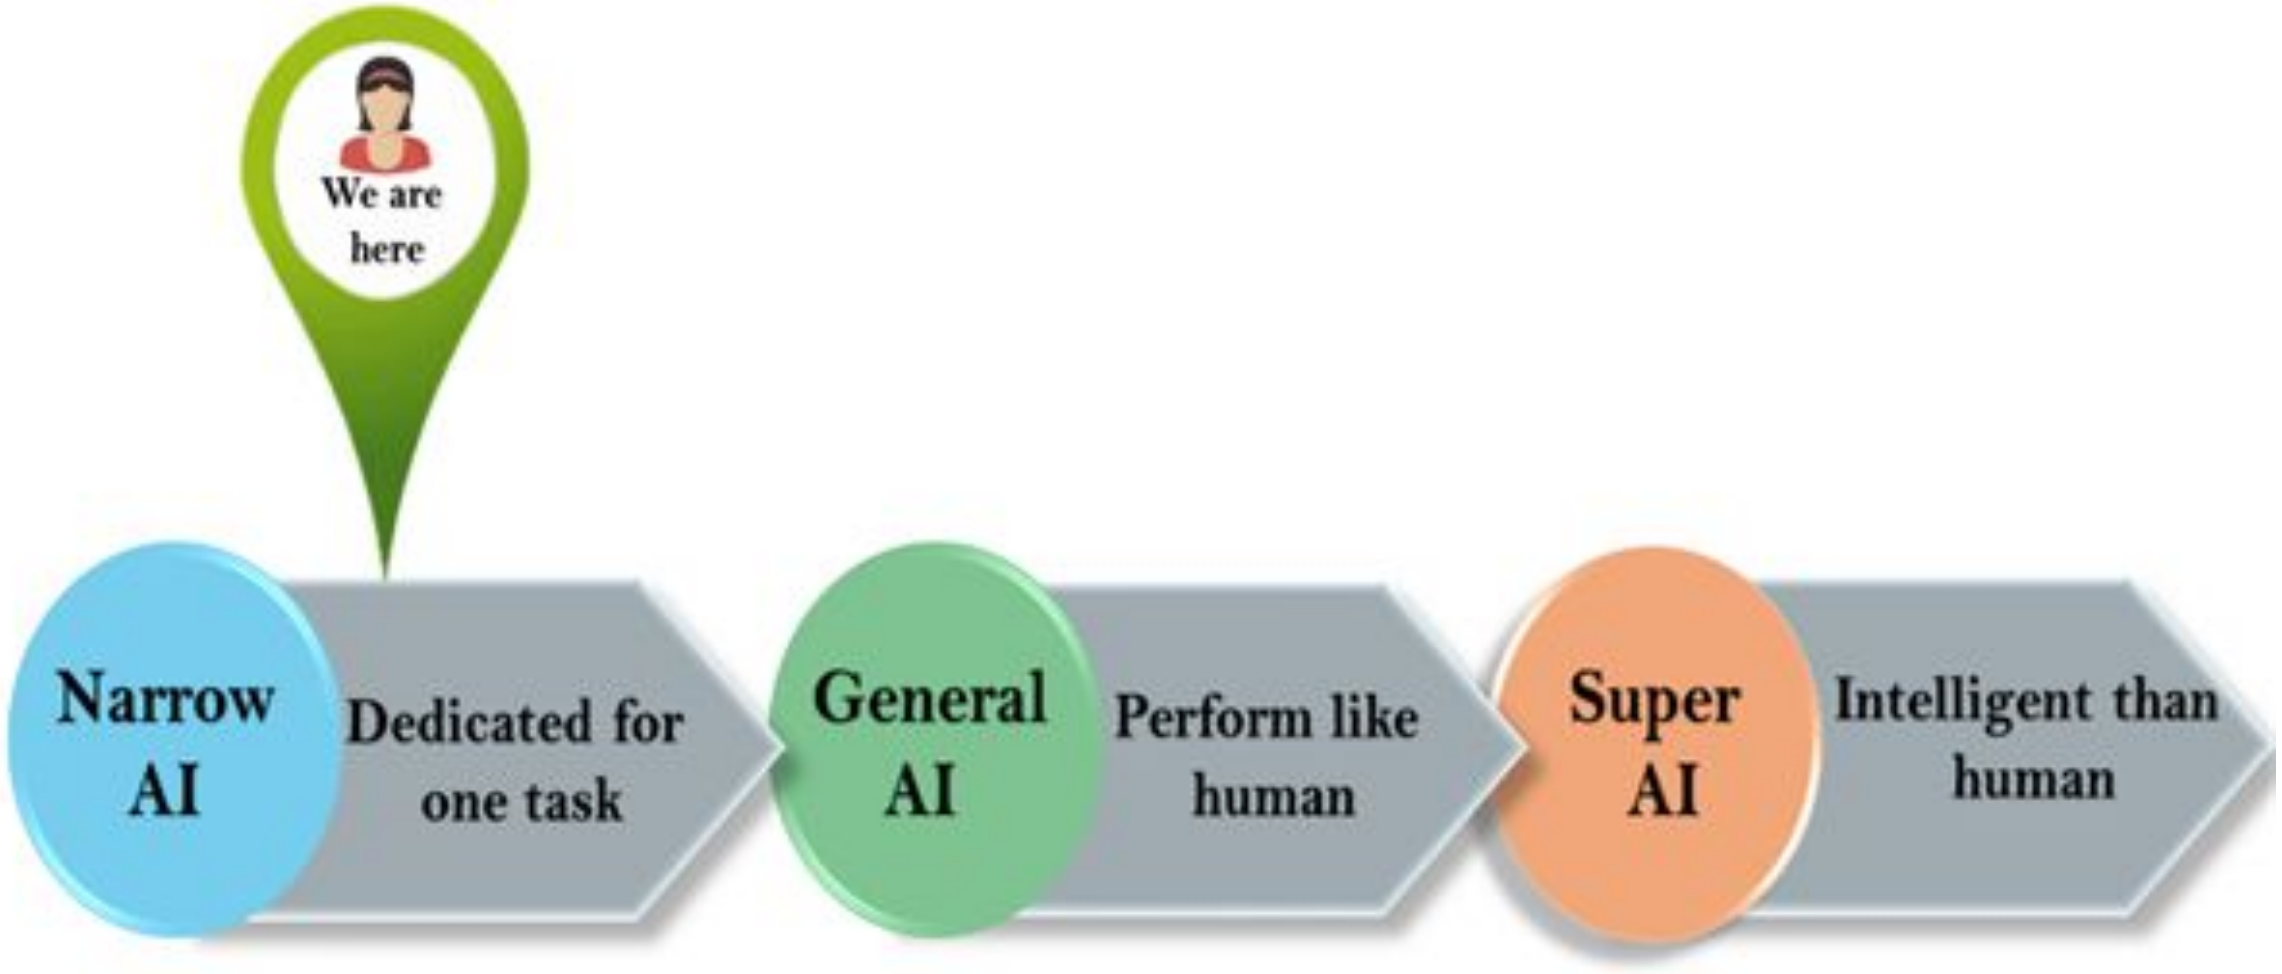
We are here
Narrow AI
Dedicated for one task
General AI
Perform like human
Super AI
Intelligent than human

Applications of AI

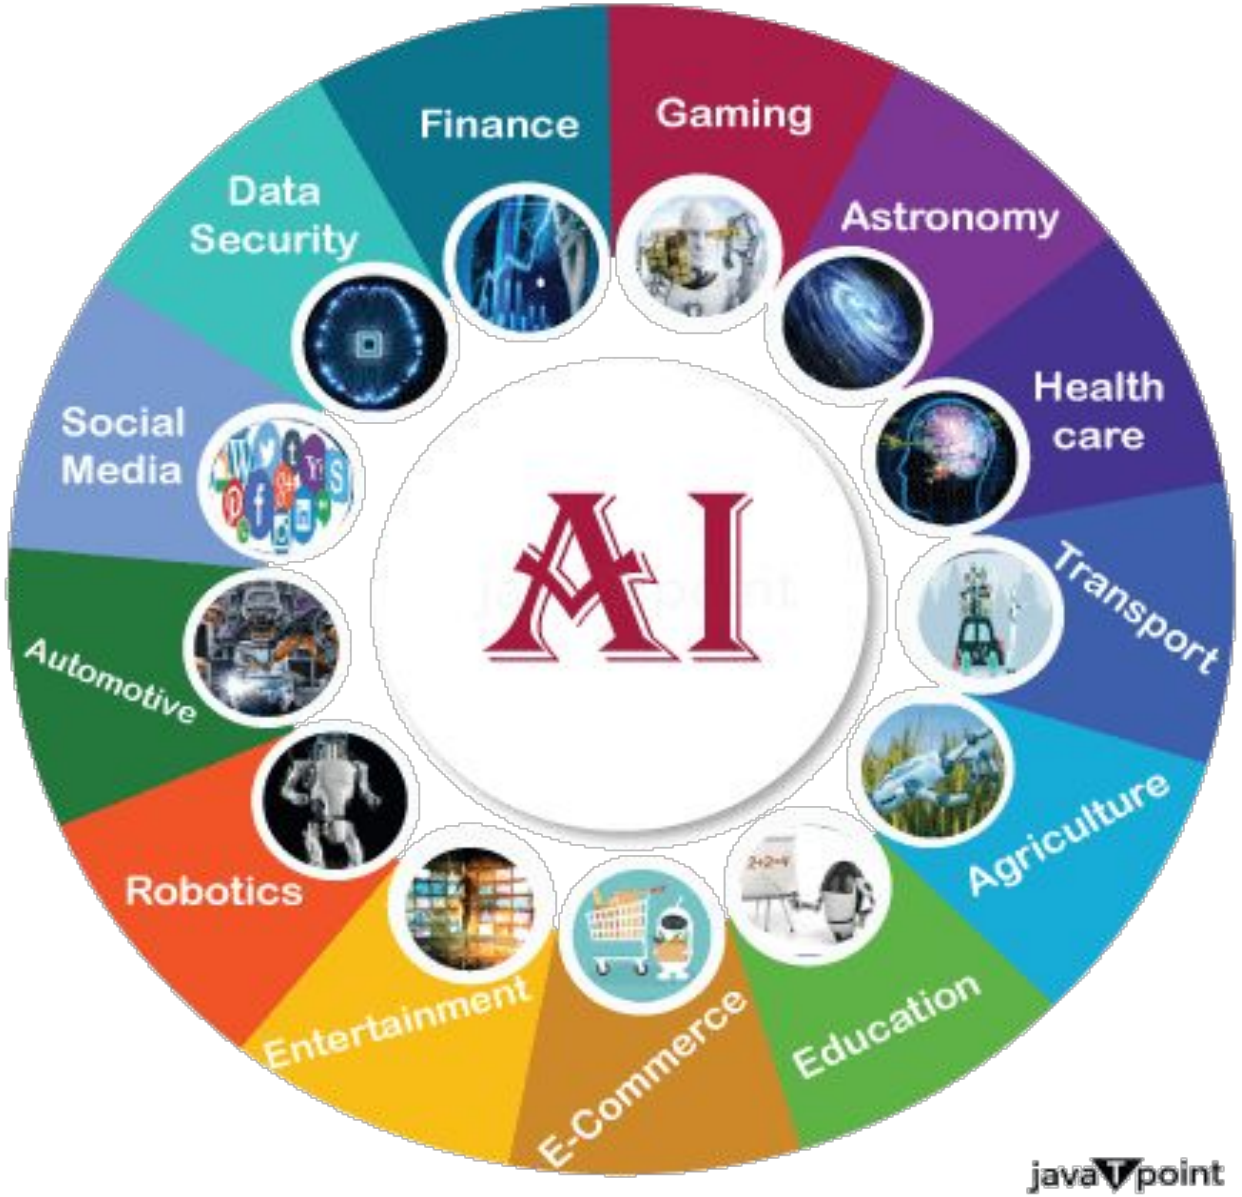

Benefits and Risk of Artificial Intelligence

AI ADVANTAGES

- 24/7 availability
- Scalability
- Improved accuracy and reduced rate of error
- Enhanced safety
- Less mundane and repetitive tasks
- Improved human experiences
- Unbiased decision-making
- Lack of emotion and judgment
- Innovation
- Improved efficiency and productivity
- Democratization of knowledge
- Expanded access to expertise

AI DISADVANTAGES

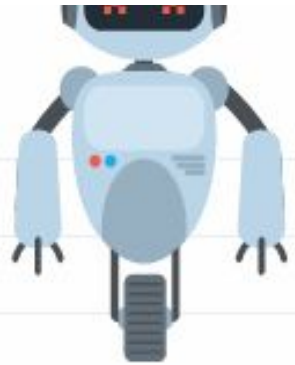

- Lack of creativity
- Absence of empathy
- Skill loss in humans
- Increased laziness in humans, lower productivity
- Job loss and displacement
- Ethical and privacy violations
- High energy requirements, questions of sustainability
- Inaccuracies
- Unforeseen risks stemming from AI flaws
- Risks on a larger scale
- Deepfakes and other fraudulent activities
- Uneven distribution of benefits to society

Intelligent Agents

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

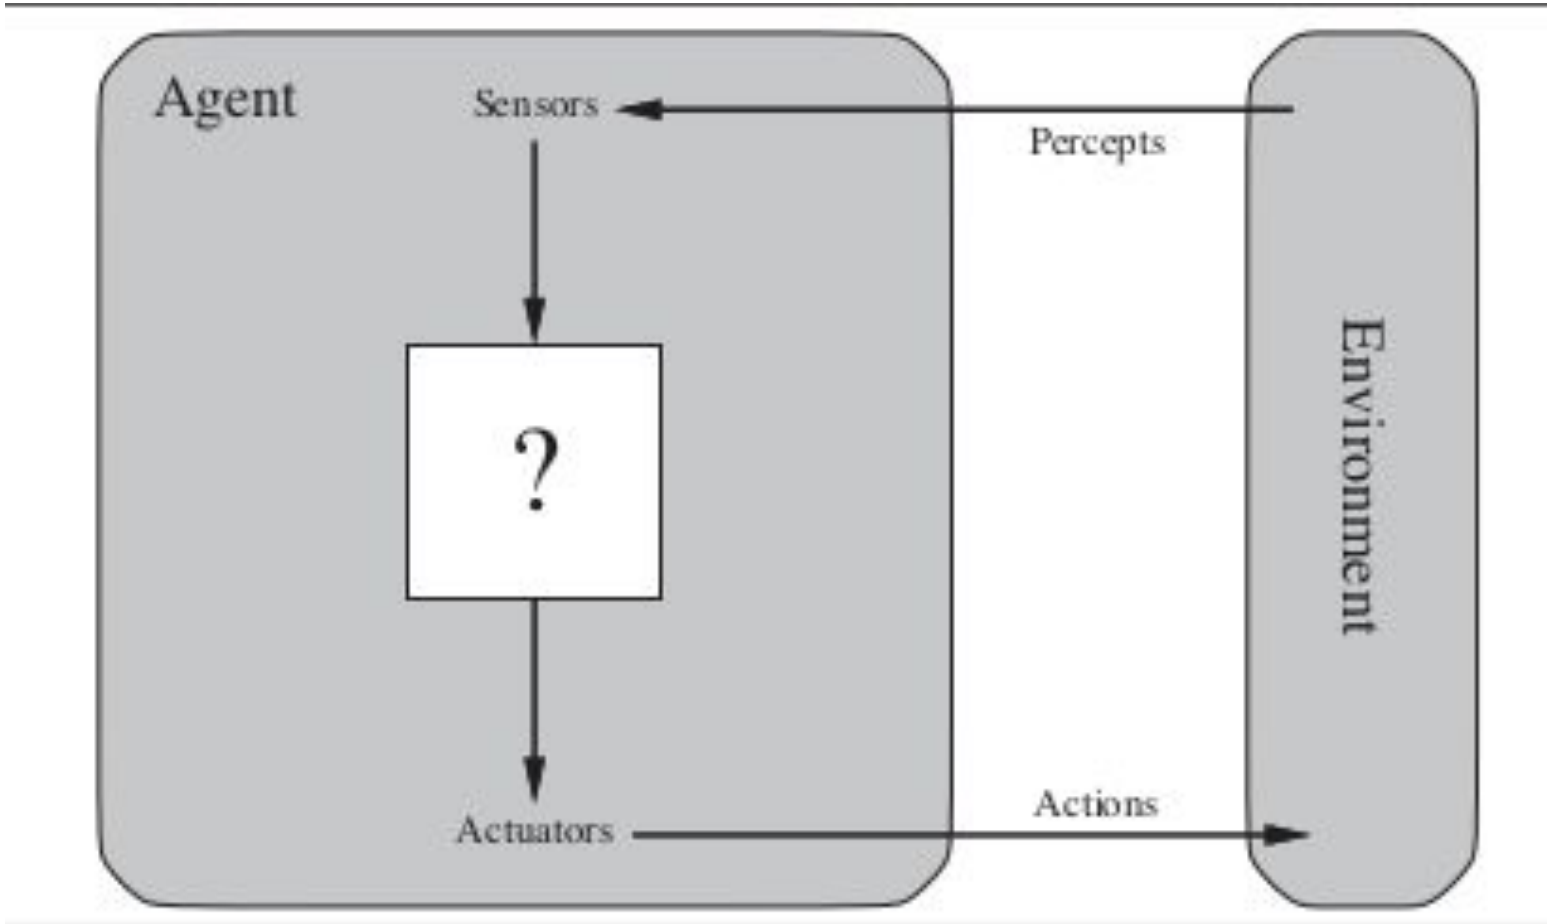

Intelligent Agents

Agent Terminology:

Performance Measure of Agent − It is the criteria, which determines how successful an
agent is.

Behavior of Agent − It is the action that agent performs after any given sequence of percepts.

Percept − It is agent's perceptual inputs at a given instance.

Percept Sequence − It is the history of all that an agent has perceived till date.

Agent Function − It is a map from the precept sequence to an action.

Intelligent Agents

The vacuum-cleaner world

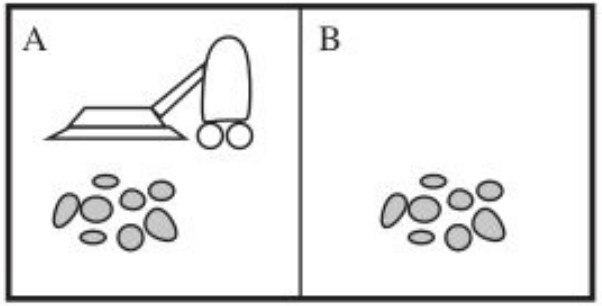

Percept sequence	Action
$[A, Clean]$	*Right*
$[A, Dirty]$	*Suck*
$[B, Clean]$	*Left*
$[B, Dirty]$	*Suck*
$[A, Clean], [A, Clean]$	*Right*
$[A, Clean], [A, Dirty]$	*Suck*
⋮	⋮
$[A, Clean], [A, Clean], [A, Clean]$	*Right*
$[A, Clean], [A, Clean], [A, Dirty]$	*Suck*
⋮	⋮

Good Behavior: The Concept of Rationality

A rational agent is one that does the right thing—conceptually speaking, every entry in the table for the agent function is filled out correctly.

Rationality:
What is 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.

This leads to a 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.

The Nature of Environment

Task environments, which are essentially the "problems" to which rational agents are the "solutions."

Specifying the task environment:

PEAS (Performance, Environment, Actuators, Sensors)

In designing an agent, the first step must always be to specify the task environment as fully as possible.

- **Performance Measure**
  The criteria used by an AI agent to evaluate its actions
- **Environment**
  The external context in which the AI agent operates, including other agents, objects, and conditions
- **Actuators**
  The mechanisms that enable the AI agent to interact with the environment and execute actions
- **Sensors**
  The means by which the AI agent perceives and gathers information about its environment

The Nature of Environment

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.

The Nature of Environment

Examples of agent types and their PEAS descriptions.

Agent Type	Performance Measure	Environment	Actuators	Sensors
Medical diagnosis system	Healthy patient, reduced costs	Patient, hospital, staff	Display of 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 of scene categorization	Color pixel arrays
Part-picking robot	Percentage of parts in correct bins	Conveyor belt with parts; bins	Jointed arm and hand	Camera, joint angle sensors
Refinery controller	Purity, yield, safety	Refinery, operators	Valves, pumps, heaters, displays	Temperature, pressure, chemical sensors
Interactive English tutor	Student's score on test	Set of students, testing agency	Display of exercises, suggestions, corrections	Keyboard entry

The Nature of Environment

An environment in artificial intelligence is the surrounding of the agent. The agent takes input from the environment through sensors and delivers the output to the environment through actuators. There are several types of environments:

- Fully Observable vs Partially Observable
- Deterministic vs Stochastic
- Competitive vs Collaborative
- Single-agent vs Multi-agent
- Static vs Dynamic
- Discrete vs Continuous
- Episodic vs Sequential
- Known vs Unknown

The Nature of Environment

Fully Observable vs Partially Observable

- When an agent sensor is capable to sense or access the complete state of an agent at each point in time, it is said to be a fully observable environment else it is partially observable.
- Maintaining a fully observable environment is easy as there is no need to keep track of the history of the surrounding.
- An environment is called **unobservable** when the agent has no sensors in all environments.
- **Examples:**
  - **Chess –** the board is fully observable, and so are the opponent's moves.
  - **Driving –** the environment is partially observable because what's around the corner is not known.

The Nature of Environment

Deterministic vs Stochastic

- When a uniqueness in the agent's current state completely determines the next state of the agent, the environment is said to be deterministic.
- The stochastic environment is random in nature which is not unique and cannot be completely determined by the agent.
- **Examples:**
  - **Chess –** there would be only a few possible moves for a chess piece at the current state and these moves can be determined.
  - **Self-Driving Cars-** the actions of a self-driving car are not unique, it varies time to time.

The Nature of Environment

Competitive vs Collaborative

- An agent is said to be in a competitive environment when it competes against another agent to optimize the output.
- The game of chess is competitive as the agents compete with each other to win the game which is the output.
- An agent is said to be in a collaborative environment when multiple agents cooperate to produce the desired output.
- When multiple self-driving cars are found on the roads, they cooperate with each other to avoid collisions and reach their destination which is the output desired.

The Nature of Environment

Single-agent vs Multi-agent

- An environment consisting of only one agent is said to be a single-agent environment.
- A person left alone in a maze is an example of the single-agent system.
- An environment involving more than one agent is a multi-agent environment.
- The game of football is multi-agent as it involves 11 players in each team.

The Nature of Environment

Dynamic vs Static

- An environment that keeps constantly changing itself when the agent is up with some action is said to be dynamic.
- A roller coaster ride is dynamic as it is set in motion and the environment keeps changing every instant.
- An idle environment with no change in its state is called a static environment.
- An empty house is static as there’s no change in the surroundings when an agent enters.

The Nature of Environment

Discrete vs Continuous

- If an environment consists of a finite number of actions that can be deliberated in the environment to obtain the output, it is said to be a discrete environment.
- The game of chess is discrete as it has only a finite number of moves. The number of moves might vary with every game, but still, it's finite.
- The environment in which the actions are performed cannot be numbered i.e. is not discrete, is said to be continuous.
- Self-driving cars are an example of continuous environments as their actions are driving, parking, etc. which cannot be numbered.

The Nature of Environment

Episodic vs Sequential

- In **an Episodic task environment**, each of the agent's actions is divided into atomic incidents or episodes. There is no dependency between current and previous incidents. In each incident, an agent receives input from the environment and then performs the corresponding action.
- **Example:** Consider an example of **Pick and Place robot**, which is used to detect defective parts from the conveyor belts. Here, every time robot(agent) will make the decision on the current part i.e. there is no dependency between current and previous decisions.
- In a **Sequential environment**, the previous decisions can affect all future decisions. The next action of the agent depends on what action he has taken previously and what action he is supposed to take in the future.
- **Example:**
    - **Checkers-** Where the previous move can affect all the following moves.

The Nature of Environment

Known vs Unknown

- In a known environment, the output for all probable actions is given.
- Obviously, in case of unknown environment, for an agent to make a decision, it has to gain knowledge about how the environment works.

The Nature of 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

Figure 2.6 Examples of task environments and their characteristics.

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. This program will run on some sort of computing device with physical sensors and actuators—we call this the architecture:

agent = architecture + program

Agent programs:

It takes the current percept as input from the sensors and return an action to the actuators.

The table—an example of which is given for the vacuum world represents explicitly the agent function that the agent program embodies.

To build a rational agent in this way, we as designers must construct a table that contains the appropriate action for every possible percept sequence.

The Structure of Agents

Types of Agents-

- **Simple reflex agents;**
- **Model-based reflex agents;**
- **Goal-based agents; and**
- **Utility-based agents.**
- **Learning agents.**

Each kind of agent program combines particular components in particular ways to generate actions.

The Structure of Agents

Simple reflex agents:

- The Simple reflex agents are the simplest agents. These agents take decisions on the basis of the current percepts and ignore the rest of the percept history.
- These agents only succeed in the fully observable environment.
- The Simple reflex agent does not consider any part of percepts history during their decision and action process.
- The Simple reflex agent works on Condition-action rule, which means it maps the current state to action. Such as a Room Cleaner agent, it works only if there is dirt in the room.
- Problems for the simple reflex agent design approach:
  - They have very limited intelligence
  - They do not have knowledge of non-perceptual parts of the current state
  - Mostly too big to generate and to store.
  - Not adaptive to changes in the environment.

The Structure of Agents

Simple reflex agents:

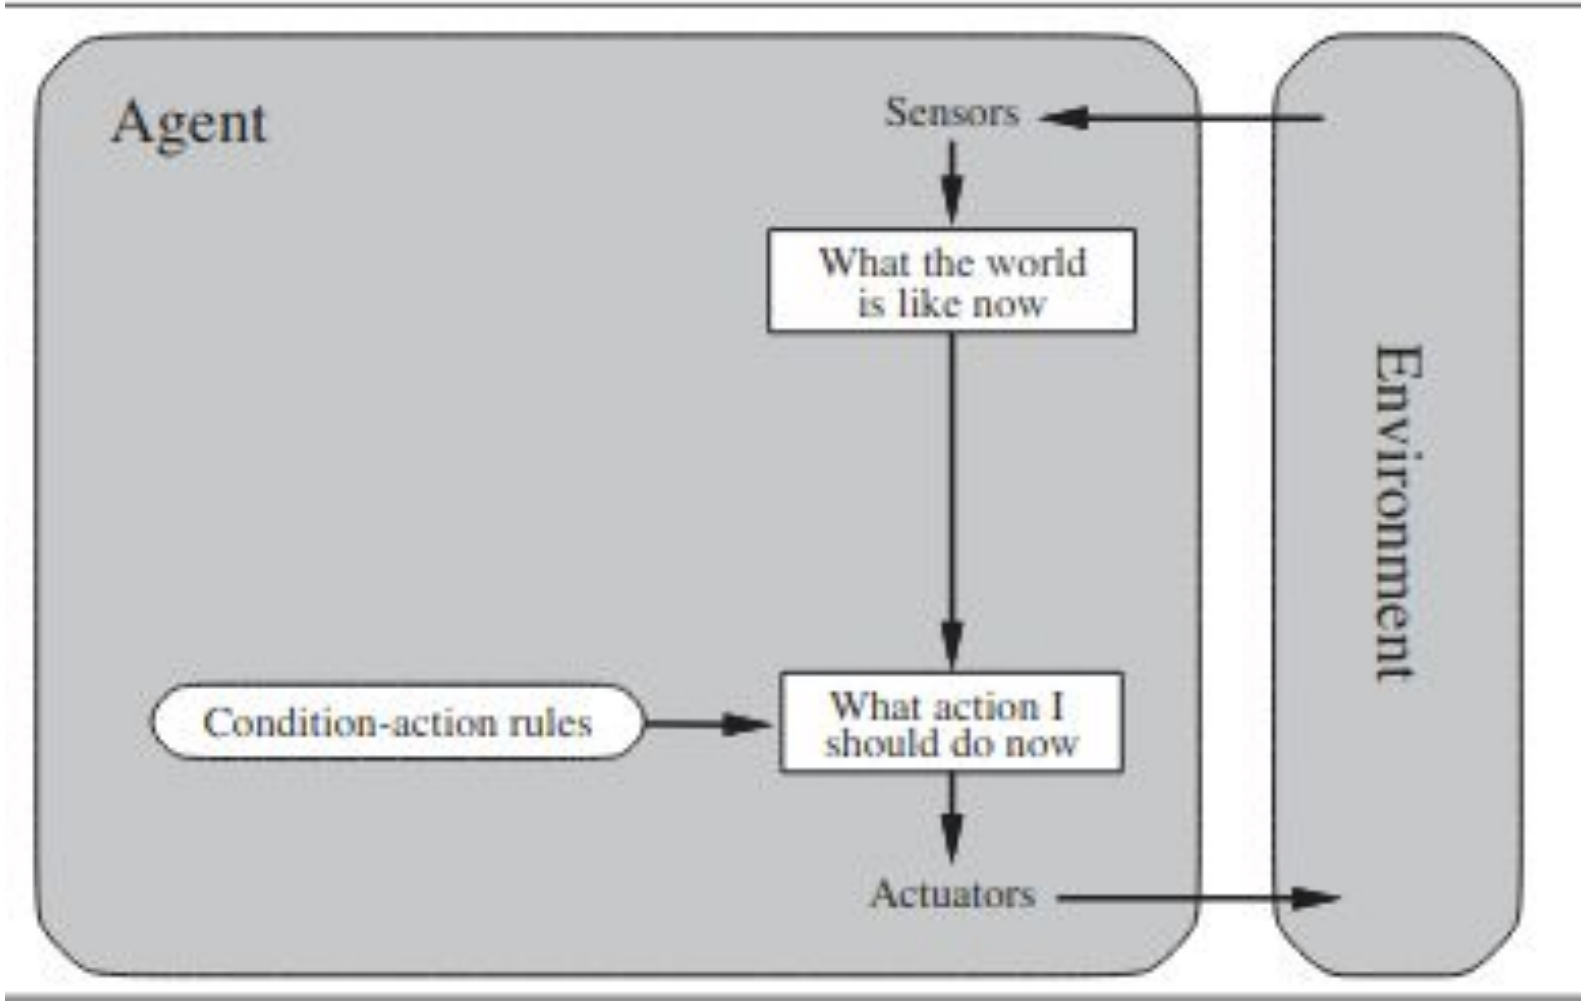

The Structure of Agents

Simple reflex agents:

```
function SIMPLE-REFLEX-AGENT(percept) returns an action
  persistent: rules, a set of condition–action rules

  state ← INTERPRET-INPUT(percept)
  rule ← RULE-MATCH(state, rules)
  action ← rule.ACTION
  return action
```

- Simple reflex agents have the admirable property of being simple, but they turn out to be of limited intelligence.
- Its have the admirable property of being simple, but they turn out to be of limited intelligence. The agent will work only if the correct decision can be made on the basis of only the current percept—that is, only if the environment is fully observable.

The Structure of Agents

Model-based reflex agents:

- The Model-based agent can work in a partially observable environment, and track the situation.
- A model-based agent has two important factors:
  1. **Model:** It is knowledge about "how things happen in the world," so it is called a Model-based agent.
  2. **Internal State:** It is a representation of the current state based on percept history.
- These agents have the model, "which is knowledge of the world" and based on the model they perform actions.
- Updating the agent state requires information about:
  1. How the world evolves
  2. How the agent's action affects the world.

The Structure of Agents

Model-based reflex agents:

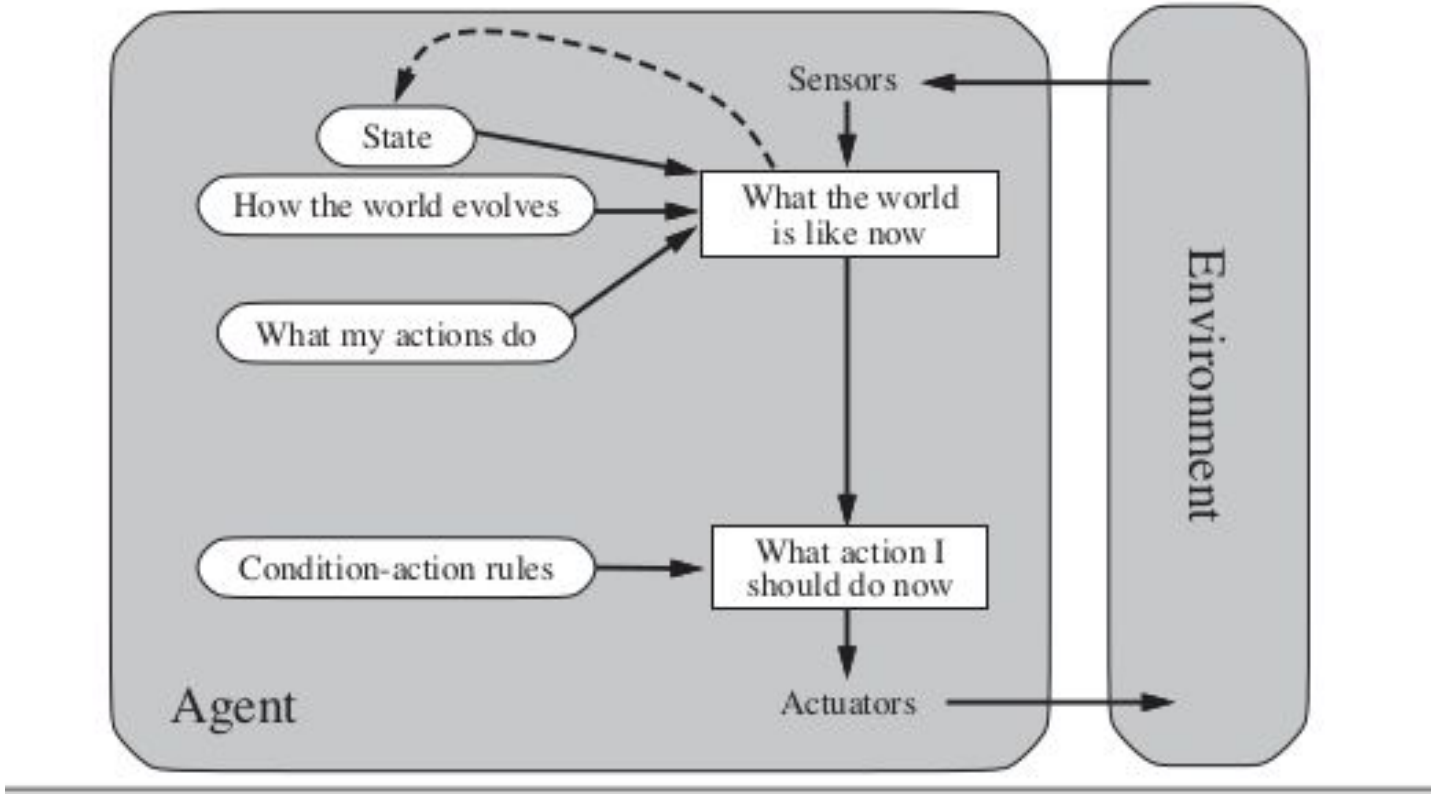

The Structure of Agents

Goal-based agents:

- The knowledge of the current state environment is not always sufficient to decide for an agent to what to do.
- The agent needs to know its goal which describes desirable situations.
- Goal-based agents expand the capabilities of the model-based agent by having the "goal" information.
- They choose an action, so that they can achieve the goal.
- These agents may have to consider a long sequence of possible actions before deciding whether the goal is achieved or not. Such considerations of different scenario are called searching and planning, which makes an agent proactive.

The Structure of Agents

Goal-based agents:

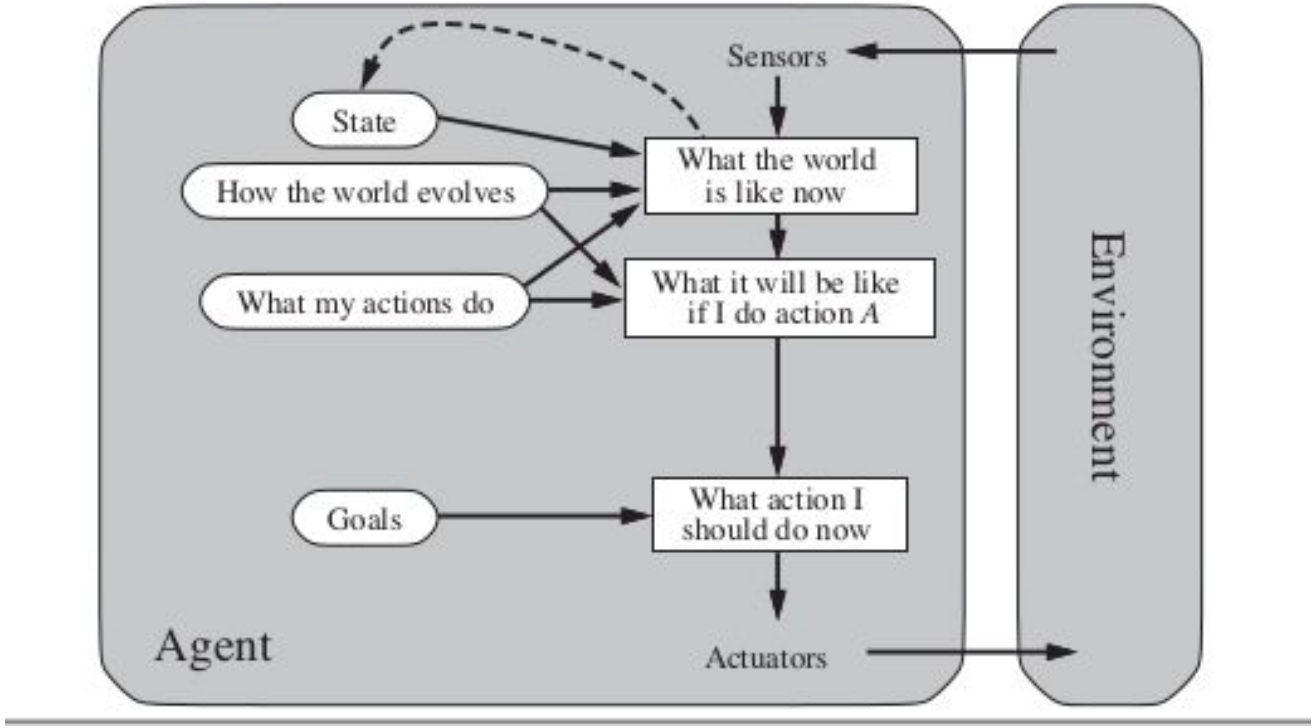

The Structure of Agents

Utility-based agents:

- These agents are similar to the goal-based agent but provide an extra component of utility measurement which makes them different by providing a measure of success at a given state.
- Utility-based agent act based not only goals but also the best way to achieve the goal.
- The Utility-based agent is useful when there are multiple possible alternatives, and an agent has to choose in order to perform the best action.
- The utility function maps each state to a real number to check how efficiently each action achieves the goals.

The Structure of Agents

Utility-based agents:

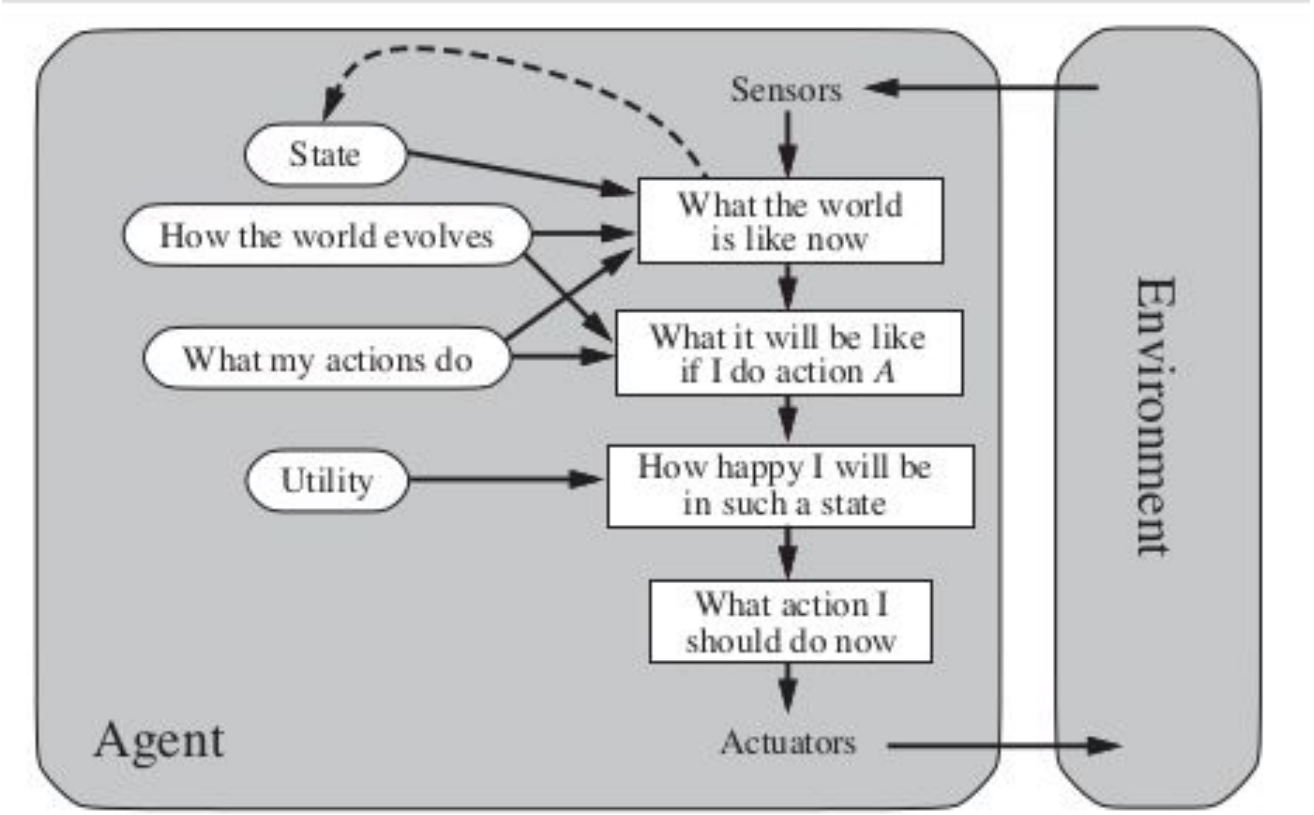

The Structure of Agents

Learning Agents:

- A learning agent in AI is the type of agent which can learn from its past experiences, or it has learning capabilities.
- It starts to act with basic knowledge and then able to act and adapt automatically through learning.
- A learning agent has mainly four conceptual components, which are:
  a. **Learning element:** It is responsible for making improvements by learning from environment
  b. **Critic:** Learning element takes feedback from critic which describes that how well the agent is doing with respect to a fixed performance standard.
  c. **Performance element:** It is responsible for selecting external action
  d. **Problem generator:** This component is responsible for suggesting actions that will lead to new and informative experiences.
- Hence, learning agents are able to learn, analyze performance, and look for new ways to improve the performance.

The Structure of Agents

Learning Agents：

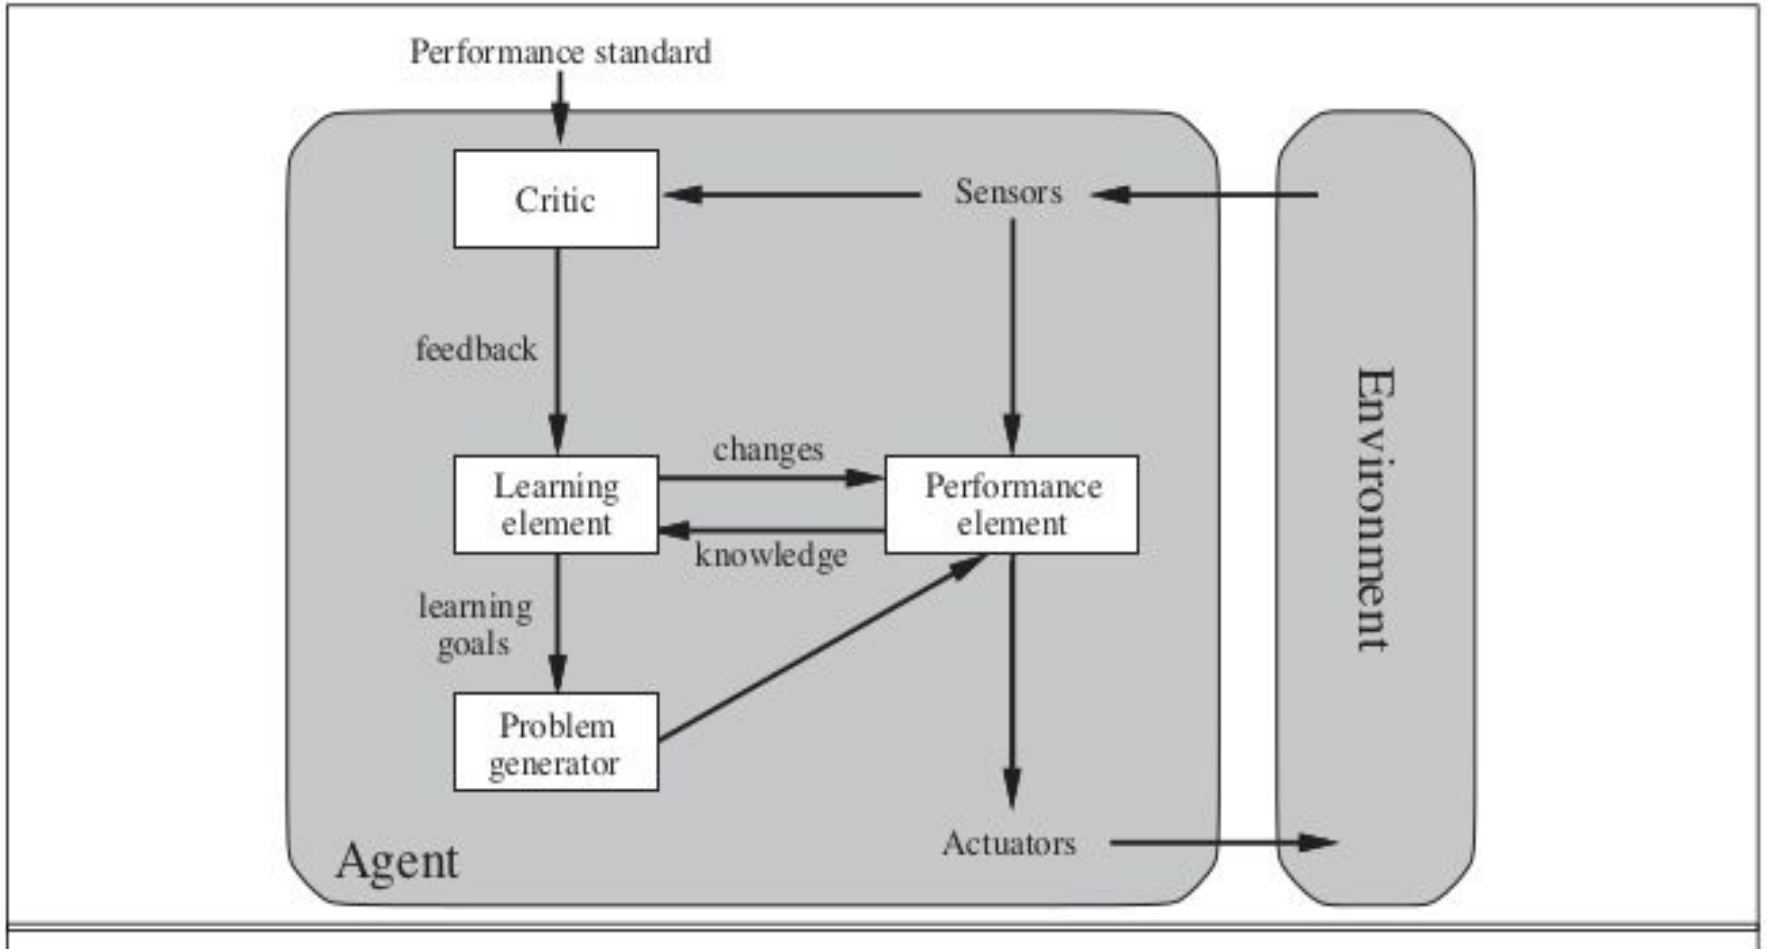

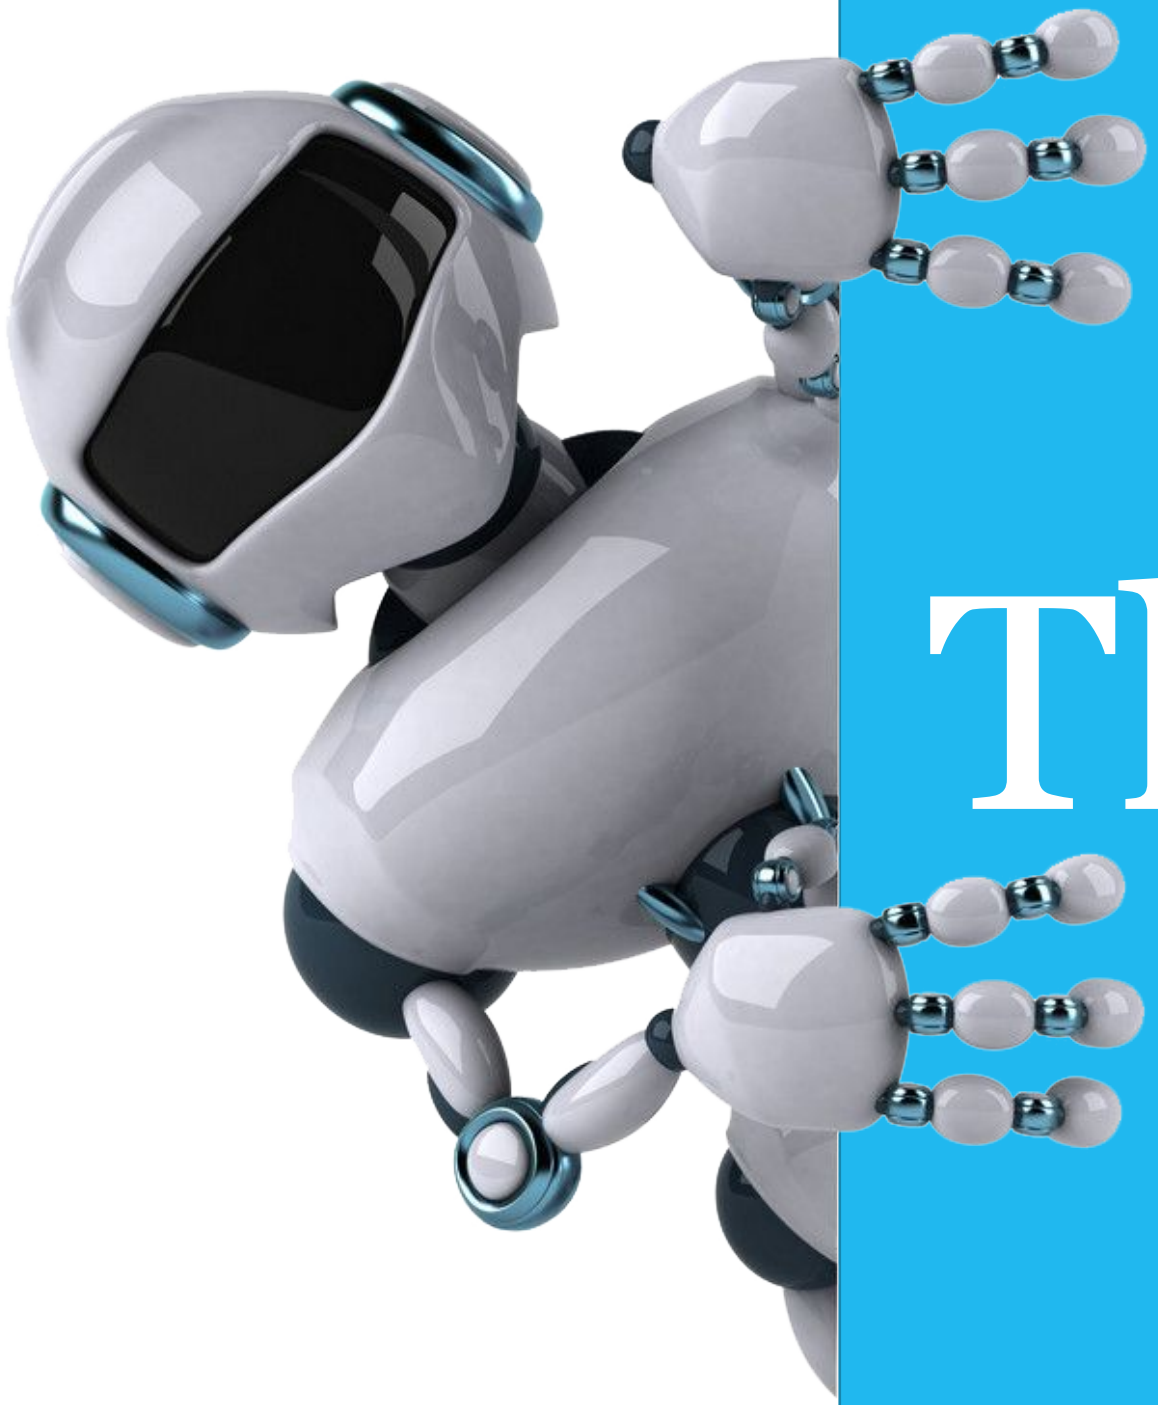

Thank You