

INTRODUCTION TO ARTIFICIAL INTELLIGENCE

**Ms. Roopa
Assistant Professor
Computer Science Department)
JSSCACs-Mysore**



INTRODUCTION

AI is one of the newest disciplines, formally initiated in 1956 when the name was coined. However, the study of intelligence is one of the oldest disciplines being approximately 2000 years old. The advent of computers made it possible for the first time for people to test models they proposed for learning, reasoning, perceiving, etc.

Artificial Intelligence is composed of two words **Artificial** and **Intelligence**, where **Artificial** defines "man-made," and **intelligence** defines "thinking power", hence AI means "a man-made thinking power."

WHAT IS ARTIFICIAL INTELLIGENCE?

Artificial intelligence is **the simulation of human intelligence processes by machines, especially computer systems**. Specific applications of AI include expert systems, natural language processing,

speech recognition and machine vision.

RS

Thinking Humanly

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

“[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . .” (Bellman, 1978)

Acting Humanly

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

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

Thinking Rationally

“The study of mental faculties through the use of computational models.”
(Charniak and McDermott, 1985)

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

Acting Rationally

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

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

ACTING HUMANLY

- The first proposal for success in building a program and acts humanly was the Turing Test. To be considered intelligent a program must be able to act sufficiently like a human to fool an interrogator.
- A human interrogates the program and another human via a terminal simultaneously. If after a reasonable period, the interrogator cannot tell which is which, the program passes.

To pass this test requires:

- ✓ Natural language processing
- ✓ Knowledge representation
- ✓ Automated reasoning
- ✓ Machine learning

- This test avoids physical contact and concentrates on "higher level" mental faculties. A total Turing test would require the program to also do:
 - ✓ Computer vision

✓ Robotics

RRS

THINKING HUMANLY

- This requires "getting inside" of the human mind to see how it works and then comparing our computer programs to this. This is what cognitive science attempts to do.
- Another way to do this is to observe a human problem solving and argue that one's programs go about problem solving in a similar way.

EXAMPLE:

- GPS (General Problem Solver) was an early computer program that attempted to model human thinking.
- The developers were not so much interested in whether or not GPS solved problems correctly.
- They were more interested in showing that it solved problems like people, going through the same steps

and taking around the same amount of time to perform those steps.

RRSS

THINKING RATIONALLY

Aristotle was one of the first to attempt to codify "thinking". His **sylogisms** provided patterns of argument structure that always gave correct conclusions, given correct premises.

EXAMPLE: All computers use energy. Using energy always generates heat. Therefore, all computers generate heat.

This initiated the field of **logic**. Formal logic was developed in the late nineteenth century. This was the first step toward enabling computer programs to reason logically.

By **1965**, programs existed that could, given enough time and memory, take a description of the problem in logical notation and find the solution, if one existed. The **logicist** tradition in AI hopes to build on such programs to create intelligence.

There are two main obstacles to this approach: First, it is difficult to make informal knowledge precise enough to use the

logician approach particularly when there is uncertainty in the knowledge.

Second, there is a big difference between being able to solve a problem in principle and doing so in practice.

RRSS

ACTING RATIONALLY / THE RATIONAL AGENT APPROACH

Acting rationally means acting so as to achieve one's goals, given one's beliefs. An agent is just something that perceives and acts.

In the logical approach to AI, the emphasis is on correct inferences. This is often part of being a rational agent because one way to act rationally is to reason logically and then act on one's conclusions. But this is not all of rationality because agents often find themselves in situations where there is no provably correct thing to do, yet they must do something. There are also ways to act rationally that do not seem to involve inference, e.g., reflex actions.

The study of AI as rational agent design has two advantages:

1. It is more general than the logical approach because correct inference is only a useful mechanism for achieving rationality, not a necessary one.
2. It is more amenable to scientific development than approaches based on human behaviour or human thought because a standard of rationality can be defined independent of humans.

Achieving perfect rationality in complex environments is not possible because the computational demands are too high. However, we will study perfect rationality as a starting place.

RRSS

FOUNDATIONS OF AI

Like any history, this one is forced to concentrate on a small number of people, events, and ideas and to ignore others that also were important. We organize the history around a series of questions. We certainly would not wish to give the impression that these questions are the only ones the disciplines address or that the disciplines have all been working toward AI as their ultimate fruition.

1. PHILOSOPHY

- ✓ Can formal rules be used to draw valid conclusions?
- ✓ How does the mind arise from a physical brain?
- ✓ Where does knowledge come from?
- ✓ How does knowledge lead to action?

Aristotle (384–322 B.C.), was the first to formulate a precise set of laws governing the rational part of the mind. He developed an informal system of syllogisms for proper reasoning, which in principle allowed one to generate conclusions mechanically, given initial premises.

Much later, Ramon Lull (d. 1315) had the idea that useful reasoning could actually be carried out by a mechanical artifact. Thomas Hobbes (1588–1679) proposed that reasoning was like numerical computation that “we add and subtract in our silent thoughts.”

2 MATHEMATICS

- ✓ What are the formal rules to draw valid conclusions?
- ✓ What can be computed?
- ✓ How do we reason with uncertain information?

Philosophers staked out most of the important ideas of AI, but to move to a formal science requires a level of mathematical formalism in three main areas: computation, logic and probability.

Mathematicians have proved that there exists an algorithm to prove any true statement in first-order logic.

However, if one adds the principle of induction required to capture the semantics of the natural numbers, then this is no longer the case. Specifically, the **incompleteness theorem** showed that in any language expressive

enough to describe the properties of the natural numbers, there are true statements that are undecidable: their truth cannot be established by any algorithm.

RS

3 ECONOMICS

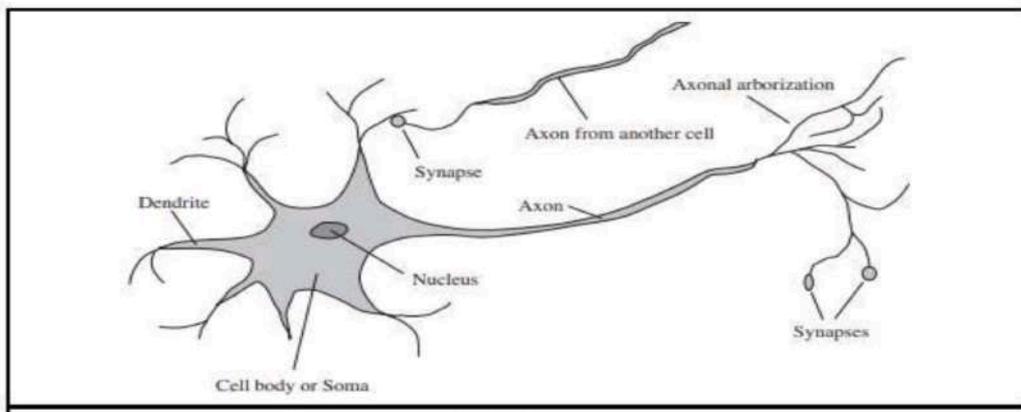
- ✓ How should we make decisions so as to maximize payoff?
- ✓ How should we do this when others may not go along?
- ✓ How should we do this when the payoff may be far in the future?

4 NEUROSCIENCE

- ✓ How do brains process information?
- Neuroscience is the study of the nervous system, particularly the brain. Although the exact way in which the brain enables thought is one of the great mysteries of science, the fact that it does enable thought has been appreciated for thousands of years because of the evidence that strong blows to the head can lead to mental incapacitation.
- It has also long been known that human brains are somehow different; in about 335 B.C. Aristotle wrote, “Of all the animals, man has the largest brain in proportion to his size.”⁵ Still, it was not until the middle of the

18th century that the brain was widely recognized as the seat of consciousness. Before then, candidate locations included the heart and the spleen.

RRSS



The parts of a nerve cell or neuron. Each neuron consists of a cell body, or soma, that contains a cell nucleus. Branching out from the cell body are a number of fibers called dendrites and a single long fiber called the axon. The axon stretches out for a long distance, much longer than the scale in this diagram indicates.

	Supercomputer	Personal Computer	Human Brain
Computational units	10^4 CPUs, 10^{12} transistors	4 CPUs, 10^9 transistors	10^{11} neurons
Storage units	10^{14} bits RAM 10^{15} bits disk	10^{11} bits RAM 10^{13} bits disk	10^{11} neurons 10^{14} synapses
Cycle time	10^{-9} sec	10^{-9} sec	10^{-3} sec

5 PSYCHOLOGY

- ✓ How do humans and animals think and act?

The principle characteristic of **cognitive psychology** is that the brain processes and processes information. The claim is that beliefs, goals, and reasoning steps can be useful components of a theory of human behaviour. The knowledge-based agent has three key steps:

1. Stimulus is translated into an internal representation
2. The representation is manipulated by cognitive processes to derive new internal representations

These are translated into actions

6 COMPUTER ENGINEERING

- ✓ How can we build an efficient computer?

For artificial intelligence to succeed, we need two things: intelligence and an artifact. The computer has

been the artifact of choice. The modern digital electronic computer was invented independently and almost simultaneously by scientists in three countries embattled in World War II.

RS

7 CONTROL THEORY AND CYBERNETICS

- ✓ How can artifacts operate under their own control?

Modern control theory, especially the branch known as stochastic optimal control, has as its goal the design of systems that maximize an objective function over time. This roughly OBJECTIVE FUNCTION matches our view of AI: designing systems that behave optimally. Why, then, are AI and control theory two different fields, despite the close connections among their founders? The answer lies in the close coupling between the mathematical techniques that were familiar to the participants and the corresponding sets of problems that were encompassed in each world view.

8. LINGUISTICS

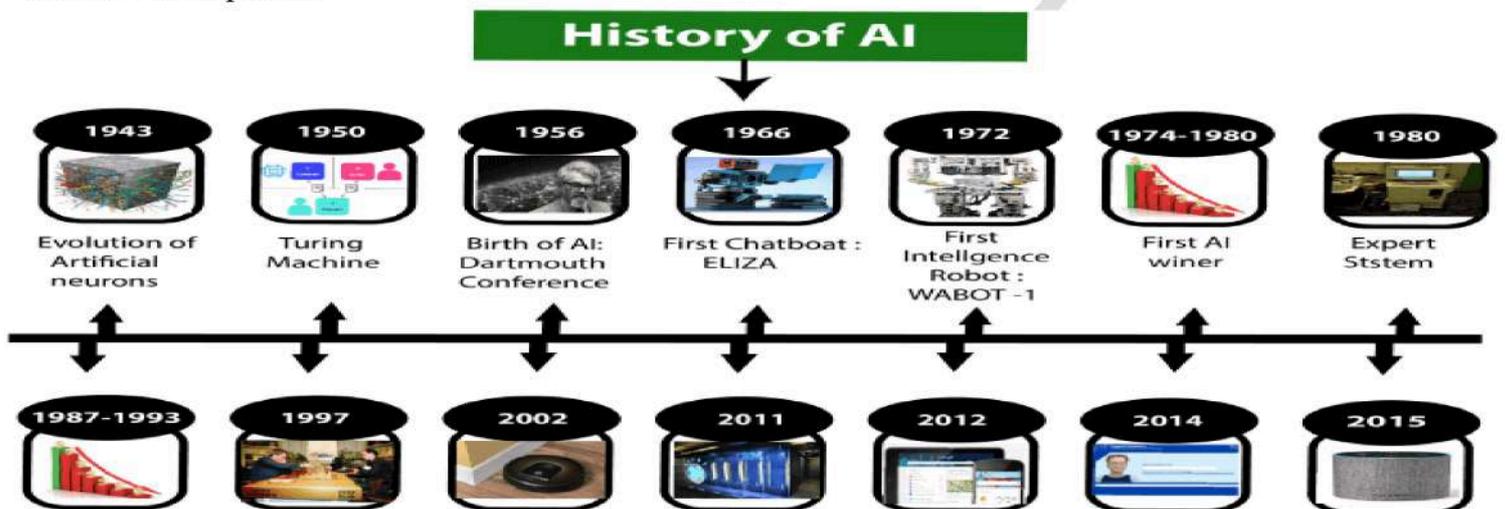
- ✓ Having a theory of how humans successfully process natural language is an AI-complete problem - if we could solve this problem then we would have created a model of intelligence.

- ✓ Much of the early work in knowledge representation was done in support of programs that attempted natural language understanding.

RRSS

HISTORY OF ARTIFICIAL INTELLIGENCE

Artificial Intelligence is not a new word and not a new technology for researchers. This technology is much older than you would imagine. Even there are the myths of Mechanical men in Ancient Greek and Egyptian Myths. Following are some milestones in the history of AI which defines the journey from the AI generation to till date development.



Maturation of Artificial Intelligence (1943-1952)

- **Year 1943:** The first work which is now recognized as AI was done by Warren McCulloch and Walter Pitts in 1943. They proposed a model of artificial neurons.
- **Year 1949:** Donald Hebb demonstrated an updating rule for modifying the connection strength between neurons. His rule is now called Hebbian learning.
- **Year 1950:** Alan Turing, who was an English mathematician and pioneered Machine learning in 1950. Alan Turing "Computing Machinery and Intelligence" in which he proposed a test. The test publishes can check the machine's ability to exhibit intelligent

behavior equivalent to human intelligence, called a Turing test.

RS

The birth of Artificial Intelligence (1952-1956)

- **Year 1955:** Allen Newell and Herbert A. Simon created the "first artificial intelligence program" which was named as "Logic Theorist". This program had proved 38 of 52 Mathematics theorems, and find new and more elegant proofs for some theorems.
- **Year 1956:** The word "Artificial Intelligence" first adopted by American Computer scientist John McCarthy at the Dartmouth Conference. For the first time, AI coined as an academic field.
- At that time high-level computer languages such as FORTRAN, LISP, or COBOL were invented. And the enthusiasm for AI was very high at that time.

The golden years-Early enthusiasm (1956-1974)

- **Year 1966:** The researchers emphasized developing algorithms which can solve mathematical problems. Joseph Weizenbaum created the first chatbot in 1966, which was named as ELIZA.
- **Year 1972:** The first intelligent humanoid robot was built in Japan which was named as WABOT-1.

The first AI winter (1974-1980)

- The duration between years 1974 to 1980 was the first AI winter duration. AI winter refers to the time period where computer scientist dealt with a severe shortage of funding from

government for AI researches.

- During AI winters, an interest of publicity on artificial intelligence was decreased.

A boom of AI (1980-1987)

- **Year 1980:** After AI winter duration, AI came back with "Expert System". Expert systems were programmed that emulate the decision-making ability of a human expert.
- **In the Year 1980,** the first national conference of the American Association of Artificial **Intelligence was held at Stanford University.**

The second AI winter (1987-1993)

- The duration between the years 1987 to 1993 was the second AI Winter duration.
- Again Investors and government stopped in funding for AI research as due to high cost but not

efficient result. The expert system such as XCON was very cost effective.

RS

The emergence of intelligent agents (1993-2011)

- Year 1997: In the year 1997, IBM Deep Blue beats world chess champion, Gary Kasparov, and became the first computer to beat a world chess champion.
- Year 2002: for the first time, AI entered the home in the form of Roomba, a vacuum cleaner.
- Year 2006: AI came in the Business world till the year 2006. Companies like Facebook, Twitter, and Netflix also started using AI.

Deep learning, big data and artificial general intelligence (2011-present)

- **Year 2011:** In the year 2011, IBM's Watson won jeopardy, a quiz show, where it had to solve the complex questions as well as riddles. Watson had proved that it could understand natural language and can solve tricky questions quickly.
- **Year 2012:** Google has launched an Android app feature "Google now", which was able to provide information to the user as a prediction.
- **Year 2014:** In the year 2014, Chatbot "Eugene Goost man" won a competition in the infamous "Turing test."
- **Year 2018:** The "Project Debater" from IBM debated on complex topics with two master debaters and also performed extremely well.
- Google has demonstrated an AI program "Duplex" which was a virtual assistant and which had

taken hairdresser appointment on call, and lady on other side didn't notice that she was talking with the machine.

RRSS

Need for Artificial Intelligence



1. COMPETITIVE EDGE

The organizations which mean to have a serious edge over their adversaries are banking upon AI advancements to acquire this.

Take the case of the **Autopilot highlight** offered by **Tesla** in its vehicles. Tesla is utilizing Deep Learning Algorithms to accomplish Autonomous driving. This was before, when there was only one element out of many, yet now it is characterizing the brand.

2. ACCESSIBILITY

The establishment **speed, availability, and sheer scale** have enabled bolder computations to deal with progressively exciting issues. Not solely is the gear faster, expanded by specific assortments of processors (e.g., GPUs), it is moreover available looking like cloud organizations.

What used to run in explicit labs with access to super PCs would now pass on to the cloud at a lower cost. This has democratized access to the significant hardware stages to run AI, enabling duplication of new organizations.

3. FEAR OF MISSING OUT (FOMO)

No typo, you read that right! Not simply us, organizations additionally feel the **dread of passing up** a major opportunity. To stay competitive and not get tossed out of the market, they need to adjust appropriately. This is done by putting resources into advances that would upset their enterprises. Take the case of the financial part, where practically all the banks have put vigorously in chatbots with the goal that they won't pass up the following rush of interruption.

4. COST-EFFECTIVENESS

As with all other technologies, with time, AI is becoming more and more affordable. This has made it **feasible** for a lot of organizations that couldn't bear the cost of them in the past to use

these advances.

Organizations do not have that barrier of cost to implement AI.

RRSS

5. FUTURE PROOF

One thing that we all need to comprehend is that future in AI is very safe .

Organizations can and ought to **guarantee** themselves to be future confirmation by actualizing AI advancements. On the off chance that this is where the world is going, why not to head in that equivalent course and be versatile to that change.

WHY ARTIFICIAL INTELLIGENCE?

Before Learning about Artificial Intelligence, we should know that what is the importance of AI and why should we learn it.

- With the help of AI, you can create such software or devices which can solve real-world problems very easily and with accuracy such as health issues, marketing, traffic issues, etc.
- With the help of AI, you can create your personal virtual Assistant, such as Cortana, Google Assistant, Siri, etc.
- With the help of AI, you can build such Robots which can work in an environment where survival of humans can be at risk.

AI opens a path for other new technologies, new devices, and new Opportunities.

GOALS OF ARTIFICIAL INTELLIGENCE

- Replicate human intelligence
- Solve Knowledge-intensive tasks
- An intelligent connection of perception and action
- Building a machine which can perform tasks that requires human intelligence such as:
 - ✓ Proving a theorem
 - ✓ Playing chess
 - ✓ Plan some surgical operation

- ✓ Driving a car in traffic
- Creating some system which can exhibit intelligent behaviour, learn new things by itself, demonstrate, explain, and can advise to its user.

ADVANTAGES OF ARTIFICIAL INTELLIGENCE

- **High Accuracy with less errors:** AI machines or systems are prone to less errors and high accuracy as it takes decisions as per pre-experience or information.
- **High-Speed:** AI systems can be of very high-speed and fast-decision making, because of that AI systems can beat a chess champion in the Chess game.
- **High reliability:** AI machines are highly reliable and can perform the same action multiple times with high accuracy.
- **Useful for risky areas:** AI machines can be helpful in situations such as defusing a bomb, exploring the ocean floor, where to employ a human can be risky.
- **Digital Assistant:** AI can be very useful to provide digital assistant to the users such as AI technology is currently used by various E-commerce websites to show the products as per customer requirement.
- **Useful as a public utility:** AI can be very useful for public utilities such as a self-driving car which can make

our journey safer and hassle-free, facial recognition for security purpose, Natural language processing to communicate with the human in human-language, etc.

RS

DISADVANTAGES OF ARTIFICIAL INTELLIGENCE

Every technology has some disadvantages, and the same goes for Artificial intelligence. Being so advantageous technology still, it has some disadvantages which we need to keep in our mind while creating an AI system.

Following are the disadvantages of AI:

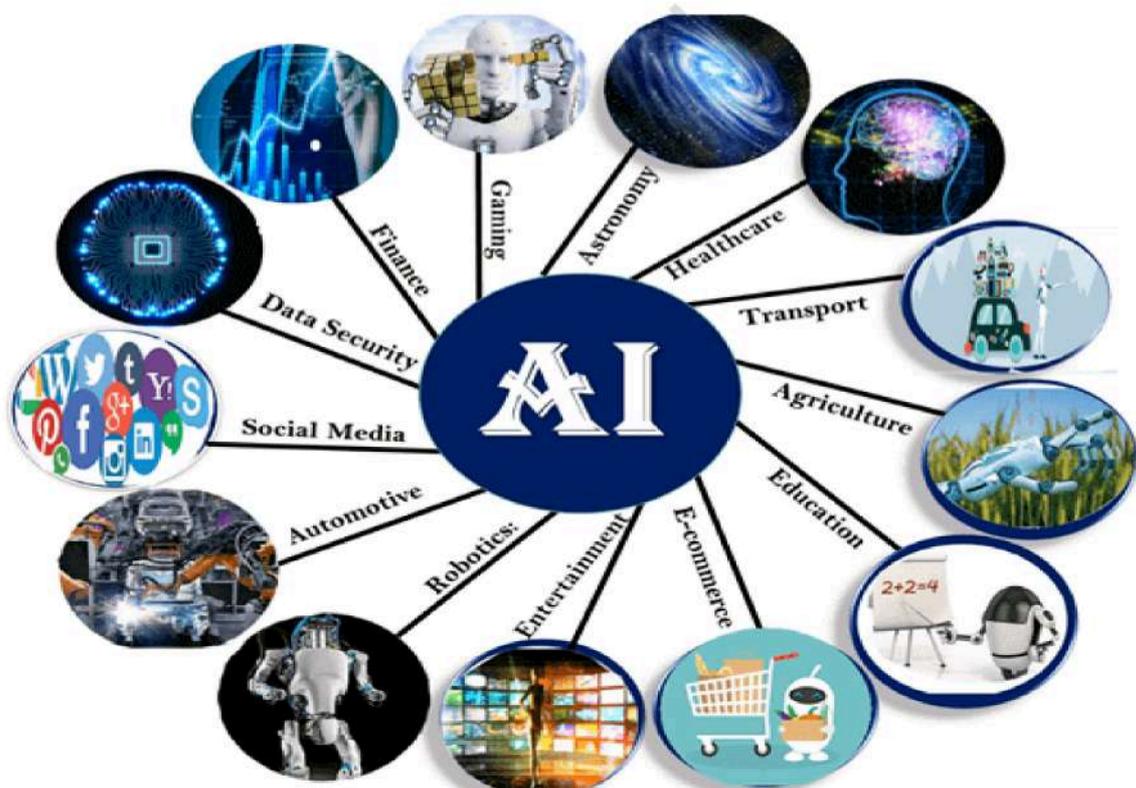
- **High Cost:** The hardware and software requirement of AI is very costly as it requires lots of maintenance to meet current world requirements.
- **Can't think out of the box:** Even we are making smarter machines with AI, but still they cannot work out of the box, as the robot will only do that work for which they are trained, or programmed.
- **No feelings and emotions:** AI machines can be an outstanding performer, but still it does not have the feeling so it cannot make any kind of emotional attachment with human, and may sometime be harmful for users if the proper care is not taken.
- **Increase dependency on machines:** With the increment of technology, people are getting more dependent on

devices and hence they are losing their mental capabilities.

- **No Original Creativity:** As humans are so creative and can imagine some new ideas but still AI machines cannot beat this power of human intelligence and cannot be creative and imaginative.

APPLICATION OF AI





1. AI in Astronomy

Artificial Intelligence can be very useful to solve complex universe problems. AI technology can be helpful for understanding the universe such as how it works, origin, etc.

2. AI in Healthcare

In the last, five to ten years, AI becoming more advantageous for the healthcare industry and going to have a significant impact on this industry.

Healthcare Industries are applying AI to make a better and faster diagnosis than humans. AI can help doctors with diagnoses and can inform when patients are worsening so that medical help can reach to the patient before hospitalization.

3. AI in Gaming

AI can be used for gaming purpose. The AI machines can play strategic games like chess, where the machine needs to think of a large number of possible places.

4. AI in Finance

AI and finance industries are the best matches for each other. The finance industry is implementing automation, chatbot, adaptive intelligence, algorithm trading, and machine learning into financial processes.

5. AI in Data Security

The security of data is crucial for every company and cyber-attacks are growing very rapidly in the digital world. AI can be used to make your data more safe and secure. Some examples such as AEG bot, AI2 Platform, are used to determine software bug and cyber-attacks in a better way.

6. AI in Social Media

Social Media sites such as Facebook, Twitter, and Snapchat contain billions of user profiles, which

need to be stored and managed in a very efficient way. AI can organize and manage massive amounts of data. AI can analyze lots of data to identify the latest trends, hashtag, and requirement of different users.

7. AI in Travel & Transport

AI is becoming highly demanding for travel industries. AI is capable of doing various travel related works such as from making travel arrangement to suggesting the hotels, flights, and best routes to the customers. Travel industries are using AI-powered chatbots which can make human-like interaction with customers for better and fast response.

8. AI in Automotive Industry

Some Automotive industries are using AI to provide virtual assistant to their user for better performance. Such as Tesla has introduced TeslaBot, an intelligent virtual assistant.

Various Industries are currently working for developing self-driven cars which can make your journey more safe and secure.

9. AI in Robotics:

Artificial Intelligence has a remarkable role in Robotics. Usually, general robots are programmed such that they can perform some repetitive task, but with the help of AI, we can create intelligent robots which can perform tasks with their own experiences without pre-programmed.

Humanoid Robots are best examples for AI in robotics, recently the intelligent Humanoid robot named as Erica and Sophia has been developed which can talk and behave like humans.

10. AI in Entertainment

We are currently using some AI based applications in our daily life with some entertainment services such as Netflix or Amazon. With the help of ML/AI algorithms, these services show the recommendations for programs or shows.

11. AI in Agriculture

Agriculture is an area which requires various resources, labor, money, and time for best result. Now a day's agriculture is becoming digital, and AI is emerging in this field. Agriculture is applying AI as agriculture robotics, solid and crop monitoring, predictive analysis. AI in agriculture can be very helpful for farmers

12. AI in E-commerce

AI is providing a competitive edge to the e-commerce industry, and it is becoming more demanding in the e-commerce business. AI is helping shoppers to discover associated products with recommended size, color, or even brand.

13. AI in education:

AI can automate grading so that the tutor can have more time to teach. AI chatbot can communicate with students as a teaching assistant.

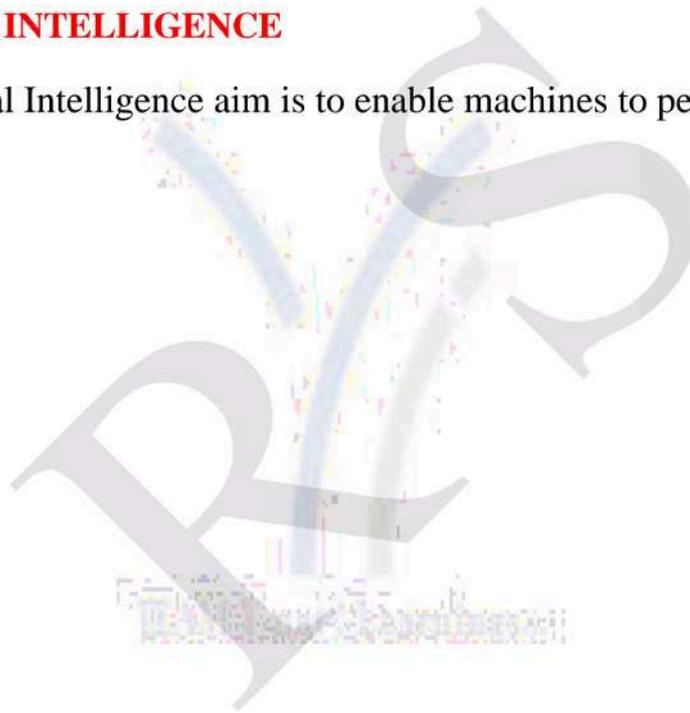
RS

Types of Artificial Intelligence

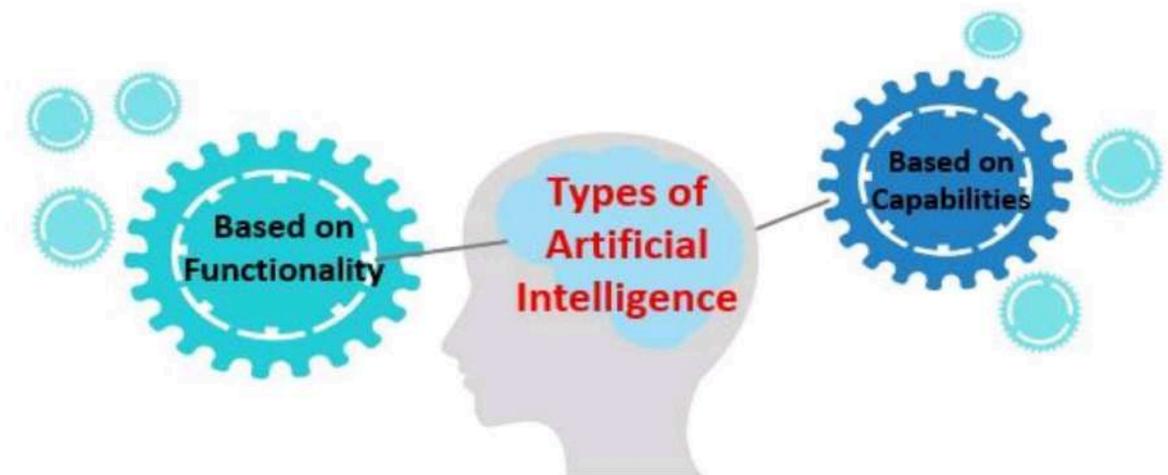


TYPES OF ARTIFICIAL INTELLIGENCE

The main aim of Artificial Intelligence is to enable machines to perform a human-like function



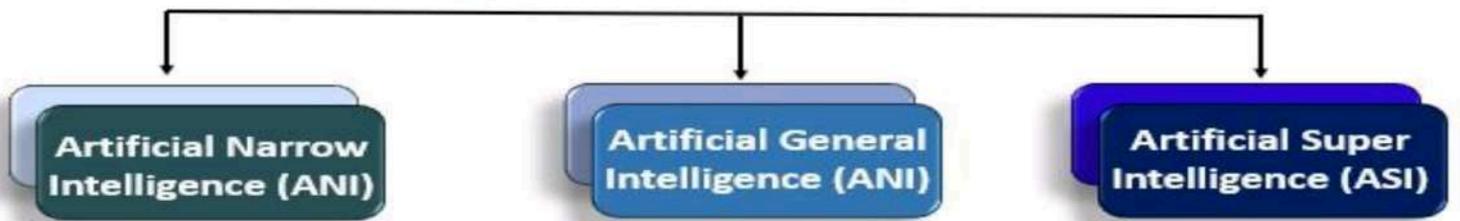
Artificial Intelligence can be divided in various types, there are mainly two types of main categorization which are **based on capabilities** and **based on functionality** of AI.



Type 1 - Based on Functionality



Type 2 - Based on Capabilities



K

BASED ON FUNCTIONALITY

1. REACTIVE MACHINES

- Purely reactive machines are the most **basic types** of Artificial Intelligence.
- Such AI systems **do not store memories** or **past experiences** for **future actions**.
- These machines **only focus** on **current scenarios** and **react on it** as per possible **best action**.
- **IBM's Deep Blue system** is an example of reactive machines.
- **Google's AlphaGo** is also an example of reactive machines.

2. LIMITED MEMORY

- Limited memory machines can store **past experiences** or **some data** for a **short period** of time.
- These machines can use stored data for a **limited time period only**.

- **Self-driving cars are one of the best examples** of Limited Memory systems. These cars can store recent speed of nearby cars, the distance of other cars, speed limit, and other information to navigate the road.



3. THEORY OF MIND

- Theory of Mind AI should understand the **human emotions**, people, beliefs, and be able to **interact socially like humans**.
- This type of AI machines are still not developed, but researchers are making lots of efforts and improvement for developing such AI machines.

4. SELF-AWARENESS

- Self-awareness AI is the **future** of **Artificial Intelligence**. These machines will be **super intelligent**, and will have their own **consciousness, sentiments, and self-awareness**.

- These machines will be **smarter than human mind**.
- Self-Awareness **AI does not exist in reality still** and it is a **hypothetical concept**.

BASED ON CAPABILITIES

1. WEAK AI OR NARROW AI:

- Narrow AI is a type of AI which is able to perform a **dedicated task** with intelligence. **The most common and currently available AI is Narrow AI in the world of Artificial Intelligence.**
- Narrow AI cannot **perform beyond its** field or **limitations**, as it is only **trained for one specific task**. Hence it is also termed as weak AI. Narrow AI can fail in unpredictable ways if it goes beyond its limits.
- **Apple Siri** is a good example of Narrow AI, but it operates with a limited pre-defined range of functions.
- **IBM's Watson supercomputer** also comes under Narrow AI, as it uses an Expert system approach combined with Machine learning and natural language processing.

- **Some Examples of Narrow AI are playing chess, purchasing suggestions on e-commerce site, self-driving cars, speech recognition, and image recognition.**

2. GENERAL AI:

- General AI is a type of intelligence which could perform any intellectual task with efficiency like a human.
- The idea behind the general AI to make such a system which could be **smarter** and **think like a human by its own**.
- Currently, there is **no such system exist which could come under general AI** and can perform any task **as perfect as a human**.
- The worldwide researchers are now focused on developing machines with General AI.

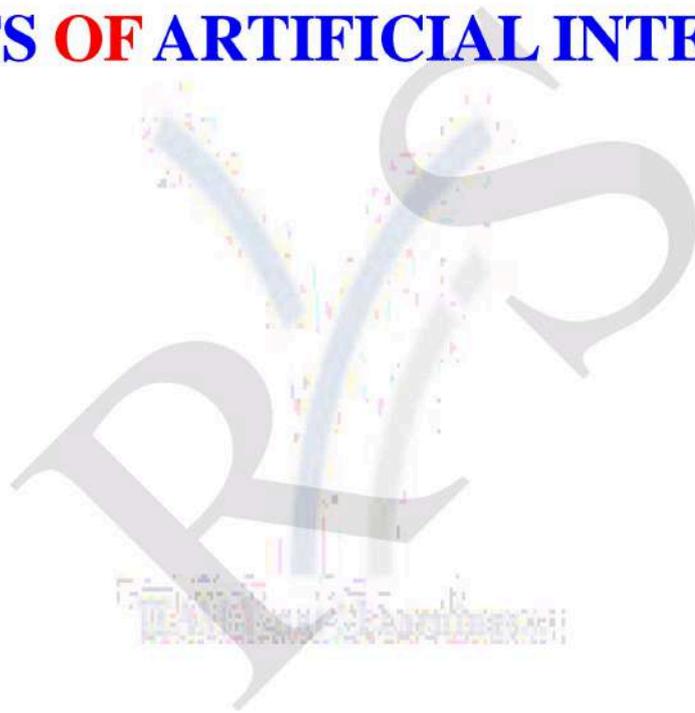
- As systems with general AI are still under research, and it will take lots of efforts and time to develop such systems.

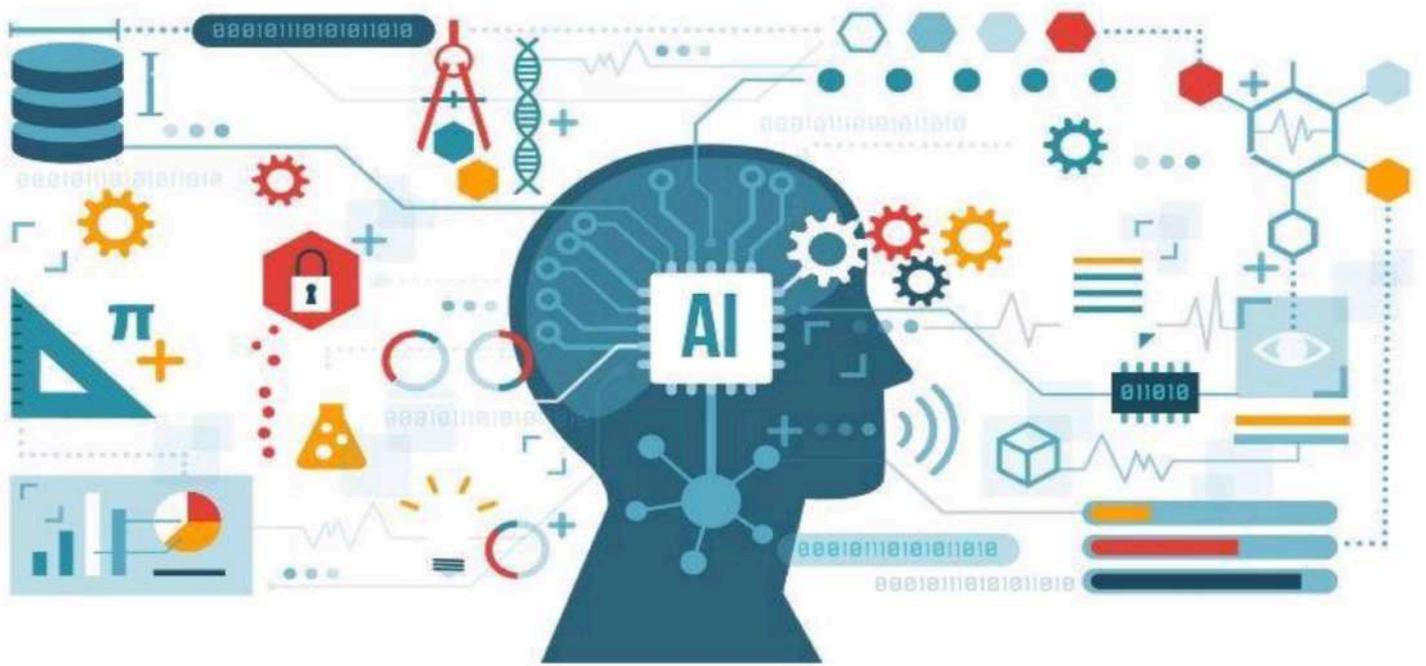
3. SUPER AI:

- Super AI is a level of Intelligence of Systems at which machines could surpass human intelligence, and can perform any task better than human with cognitive properties. It is an outcome of general AI.
- Some key characteristics of strong AI include capability include the ability to think, to reason, solve the puzzle, make judgments, plan, learn, and communicate by its own.
- Super AI is still a hypothetical concept of Artificial Intelligence. Development of such systems in real is still world changing task.



BENEFITS OF ARTIFICIAL INTELLIGENCE





BENEFITS OF ARTIFICIAL INTELLIGENCE

- 1. INCREASED EFFICIENCY**
- 2. REDUCING HUMAN RISK**
- 3. AID IN DECISION MAKING**
- 4. AVAILABILITY**

1. INCREASED EFFICIENCY

Machines are highly efficient in contrast to the human workforce. This is because they can complete monotonous tasks effectively without taking a break and irrespective of the working conditions.

Take the example of Power BI. Earlier, time was spent in **sifting, isolating, and utilizing the information** in **MS Excel**, **Power BI** utilizes AI to do the entirety of this (and significantly more)

autonomously.

This expands the effectiveness of the framework and spares significant time for the investigator to utilize it somewhere else.

RRSS

2. REDUCING HUMAN RISK

A lot of places wherein it is **hazardous** for a person to be truly present.

Man-made brainpower can be useful in such cases since it would empower the machines to have the option to settle on choices continuously. A case of this is a remote ocean boring at the sea bed where there is a high danger of human life.

3. AID IN DECISION MAKING

This is the domain where AI has successfully paid dividends. AI has the capacity to analyze a great deal of information in a brief time frame. This is helping people improve decision making. Take the case of an investor who needs to conclude whether to give out a credit or not.

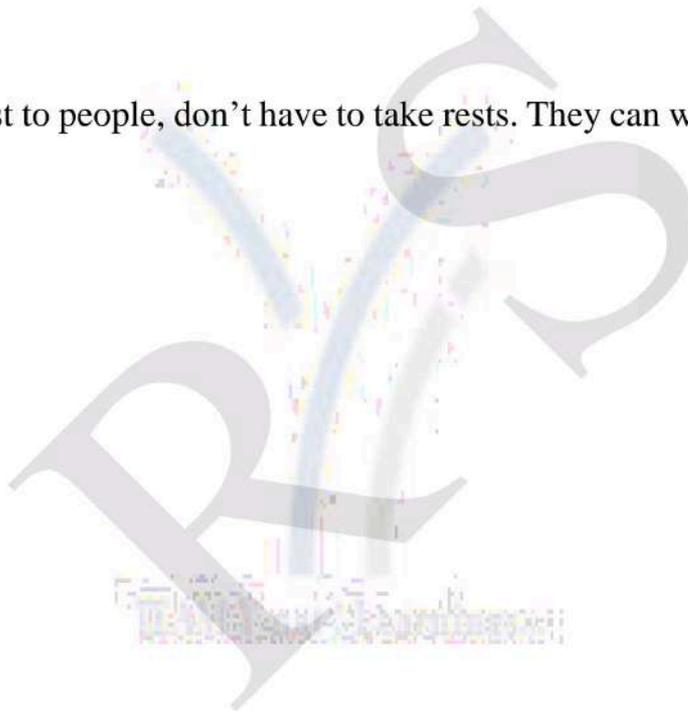
This can be highly effective in verticals like **portfolio management**. This would help in a more astute

and quicker dynamic for the investor.

RRSS

4. AVAILABILITY

Machines, in contrast to people, don't have to take rests. They can work nonstop.

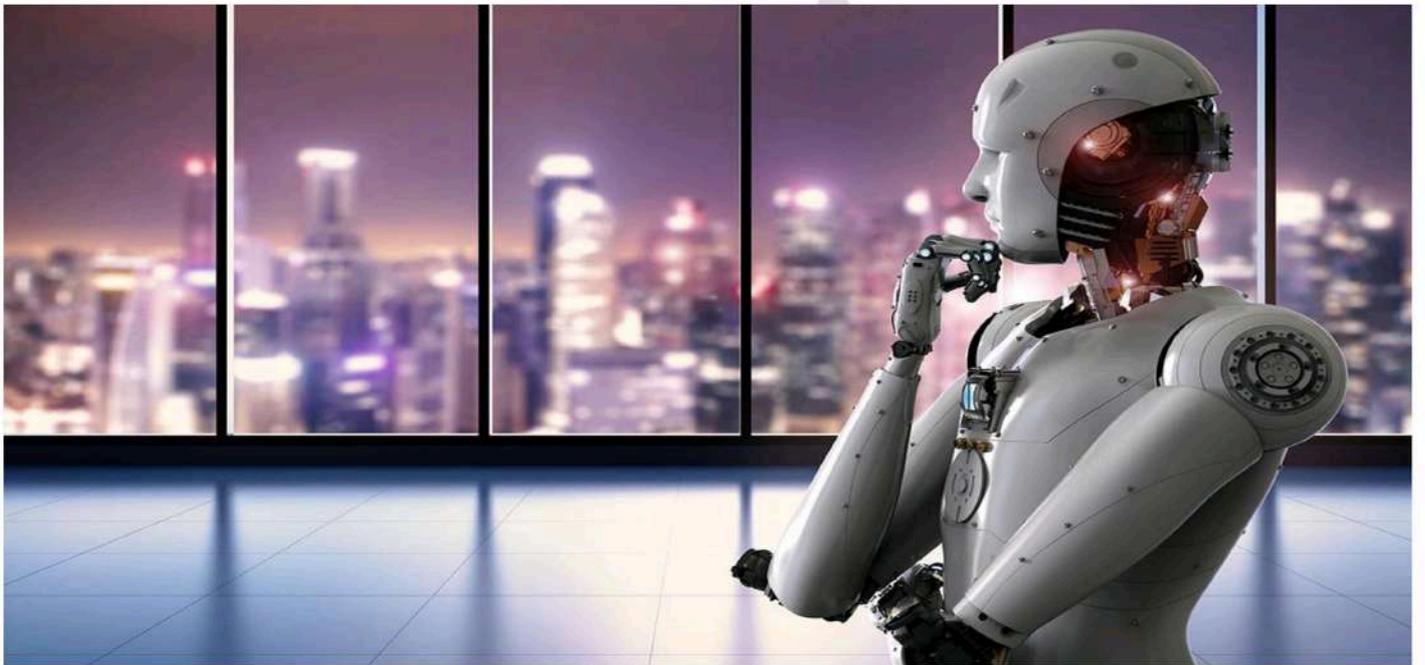


Now, we can depend on machines for keeping up required manufacturing units running with their judgment which would prompt **24×7** creation units and complete mechanization.

-  **Improve efficiency**
-  **Save costs**
-  **Avoid problems**
-  **Anticipate future scenarios**
-  **Reduce human error**
-  **Automate processes**

THE STATE OF ART OR WHAT CAN AI DO TODAY?





➤ **ROBOTIC VEHICLES:** A driverless robotic car named STANLEY sped through the rough terrain of the Mojave Desert at 22 mph, finishing the 132-mile course first to win the 2005 DARPA Grand Challenge.

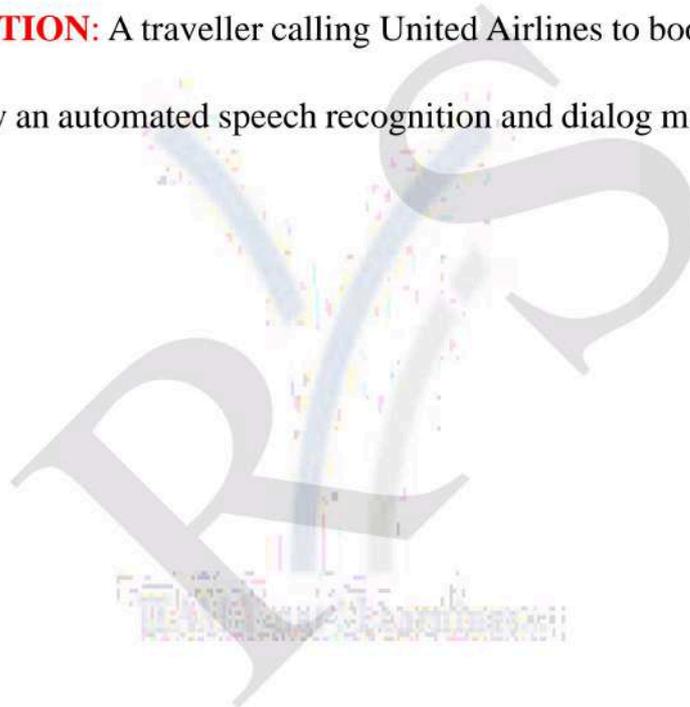
STANLEY is a Volkswagen Touareg outfitted with cameras, radar, and laser rangefinders to sense the environment and onboard software to command the steering, braking, and acceleration (Thrun, 2006).

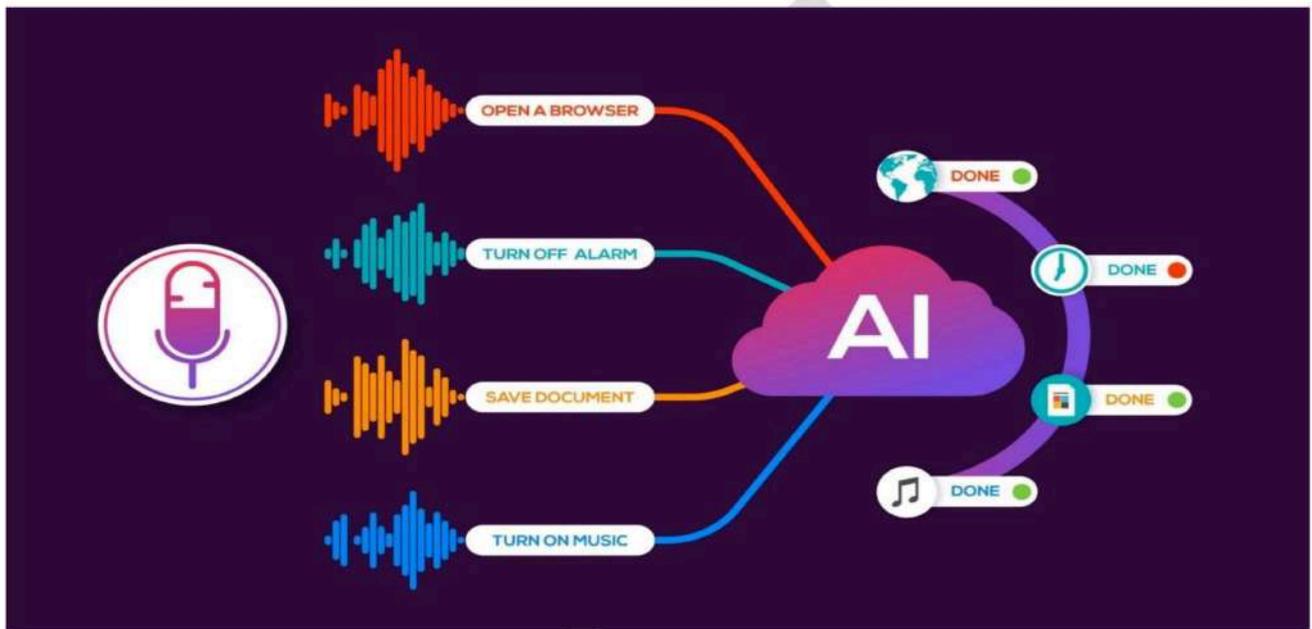
The following year CMU's BOSS won the Urban Challenge, safely driving in traffic through the streets of a closed Air Force base, obeying traffic rules and avoiding pedestrians and other vehicles.



R

➤ **SPEECH RECOGNITION:** A traveller calling United Airlines to book a flight can have the entire conversation guided by an automated speech recognition and dialog management system.





➤ AUTONOMOUS PLANNING AND

SCHEDULING: A hundred million miles from Earth, NASA's Remote Agent program became the first on-board autonomous planning program to control the scheduling of operations for a spacecraft (Jonsson *et al.*, 2000).

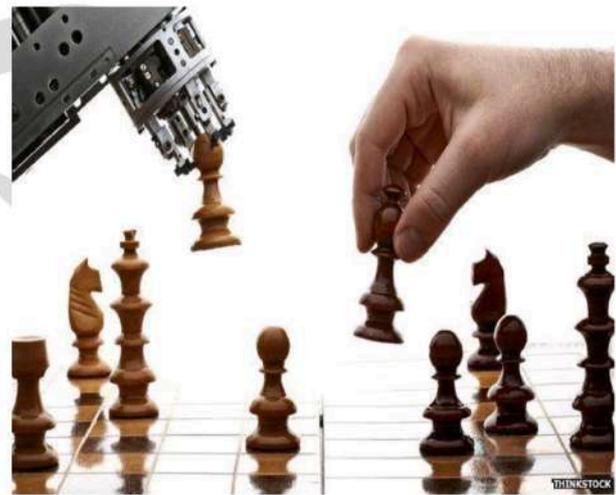
REMOTE AGENT generated plans from high-level goals specified from the ground and monitored the execution of those plans—detecting, diagnosing, and recovering from problems as they occurred.

Successor program MAPGEN (Al-Chang *et al.*, 2004) plans the daily operations for NASA's



Mars Exploration Rovers, and MEXAR2 (Cesta *et al.*, 2007) did mission planning—both logistics and science planning—for the European Space Agency’s Mars Express mission in 2008.

➤ **GAME PLAYING:** IBM's DEEP BLUE became the first computer program to defeat the world champion in a chess match when it bested Garry Kasparov by a score of 3.5 to 2.5 in an exhibition match (Goodman and Keene, 1997). Kasparov said that he felt a "new kind of intelligence" across the board from him. Newsweek magazine described the match as "The brain's last stand." The value of IBM's stock increased by \$18 billion. Human champions studied Kasparov's loss and were able to draw a few matches in subsequent years, but the most recent human-computer matches have been



won convincingly by the computer.

RRSS

- **SPAM FIGHTING:** Each day, learning algorithms classify over a billion messages as spam, saving the recipient from having to waste time deleting what, for many users, could comprise 80% or 90% of all messages, if not classified away by algorithms. Because the spammers are continually updating their tactics, it is difficult for a static programmed approach to keep up, and learning algorithms work best (Sahami *et al.*, 1998; Goodman and Heckerman, 2004).
- **LOGISTICS PLANNING:** During the Persian Gulf crisis of 1991, U.S. forces deployed a Dynamic Analysis and Replanning Tool, DART (Cross and Walker, 1994), to do automated logistics planning and scheduling for transportation. This involved up to 50,000 vehicles, cargo, and people at a time, and had to account for starting points, destinations, routes, and conflict resolution among all parameters. The AI planning techniques generated in hours a plan that would have taken weeks with

older methods. The Defence Advanced Research Project Agency (DARPA) stated that this single application more than paid back DARPA's 30-year investment in AI.

- **ROBOTICS:** The iRobot Corporation has sold over two million Roomba robotic vacuum cleaners for home use. The company also deploys the more rugged PackBot to Iraq and Afghanistan, where it is used to handle hazardous materials, clear explosives, and identify the location of snipers.
- **MACHINE TRANSLATION:** A computer program automatically translates from Arabic to English, allowing an English speaker to see the headline “Ardogan Confirms That Turkey Would Not Accept Any Pressure, Urging Them to Recognize Cyprus.” The program uses a statistical model built from examples of Arabic-to-English translations and from examples of English text totalling two trillion words (Brants *et al.*, 2007). None of the computer scientists on the team speak

Arabic, but they do understand statistics and machine learning algorithms.

RRSS

AGENTS
AND
ENVIRONMENTS

AGENTS AND ENVIRONMENTS

- **Agents** in Artificial Intelligence are the **associated concepts** that the **AI technologies** work upon.
- The **AI software or AI-enabled devices** with **sensors** generally captures the information **from** the **environment setup** and **process** the **data for further actions**.
- There are mainly two ways the agents interact with the environment, such as **perception** and **action**.
- The **person** is only **passive** for capturing the **information** without **changing** the **actual environment**, whereas **action** is the **active** form of **interaction** by **changing** the **actual environment**.

- AI technologies such as **virtual assistance catboats**, AI-enabled devices to work based on the **previous persecution data processing and learning for the actions.**

WHAT IS AN AGENT?

An Agent is anything that takes actions according to the information that it gains from the environment.

- **HUMAN-AGENT:** A human agent has eyes, ears, and other organs which work for sensors and hand, legs, vocal tract work for actuators.
- **ROBOTIC AGENT:** A robotic agent can have cameras, infrared range finder, NLP for sensors and various motors for actuators.
- **SOFTWARE AGENT:** Software agent can have keystrokes, file contents as sensory input and act on those inputs and display output on the screen.

EXAMPLE FOR AGENT AND ENVIRONMENTS?



Agents

Environments

Robot → Room

Chatbot → Chatting

Vehicle → Road

Program → Data & Rules

Machine → Working Field

HOW DOES THE AGENT INTERACT WITH THE ENVIRONMENT?

The agents interact with the environment in two ways:

1. PERCEPTION

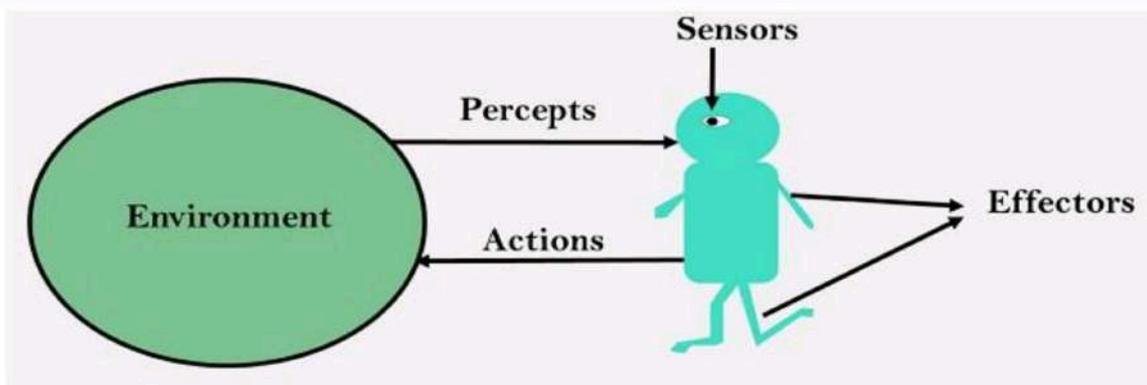
Perception is a **passive interaction**, where the **agent gains information** about the **environment without changing the environment**. The **sensors** of the **robot help** it to **gain information** about the **surroundings** without **affecting** the **surrounding**. Hence, gaining information through sensors is called perception.

2. ACTION

Action is an **active interaction** where the **environment is changed**. When the **robot moves** an **obstacle** using its **arm**, it is called an **action** as the **environment is changed**. The arm of the robot

is called an **“Effector”** as it performs the action.

RS

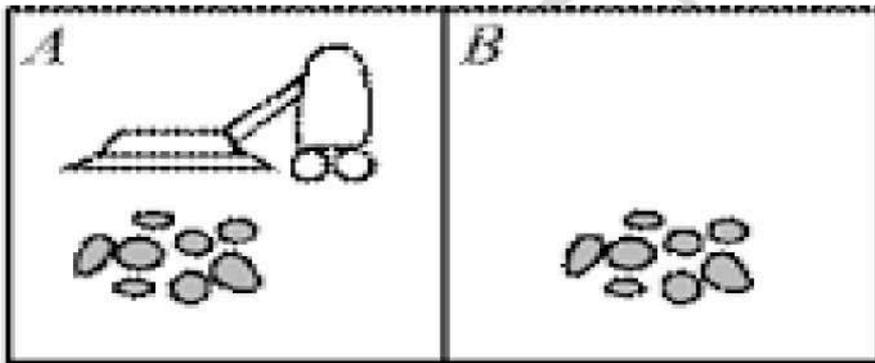


- **SENSOR:** Sensor is a device which detects the change in the environment and sends the information to other electronic devices. An agent observes its environment through sensors.
- **ACTUATORS:** Actuators are the component of machines that converts energy into motion. The actuators are only responsible for moving and controlling a system. An actuator can be an electric

motor, gears, rails, etc.

- **EFFECTORS:** Effectors are the devices which affect the environment. Effectors can be legs, wheels, arms, fingers, wings, fins, and display screen.

CONSIDER A VACUUM CLEANER WORLD



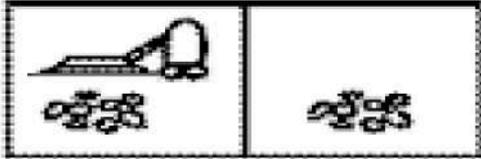
Imagine that our **intelligent agent** is a robot vacuum cleaner.

Let's suppose that the world has just two rooms. The robot can be in either room and there can

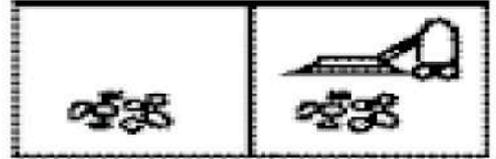
be dirt in zero, one, or two rooms.

RRSS

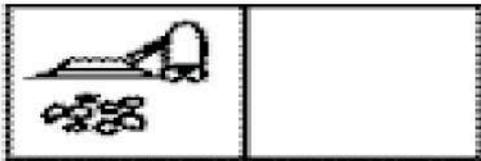
1



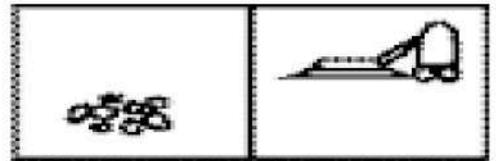
2



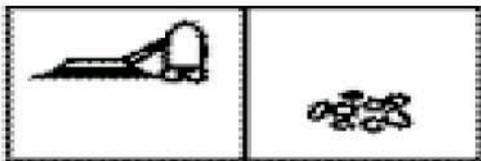
3



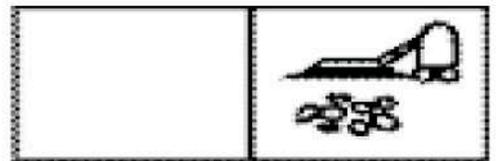
4



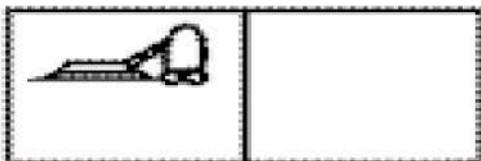
5



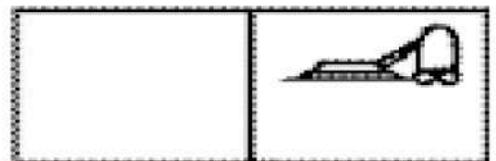
6



7

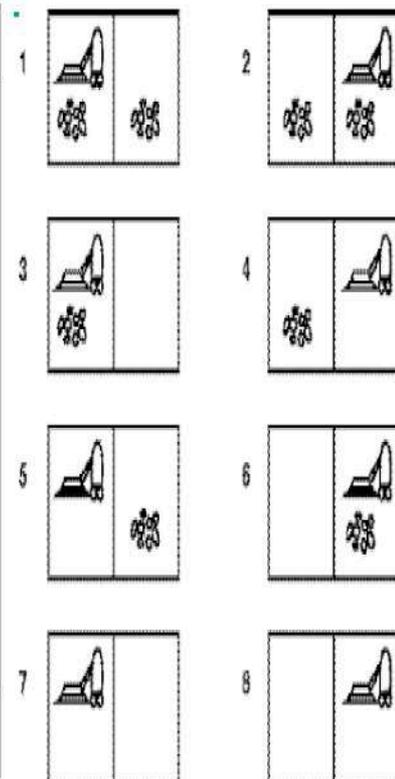


8



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
⋮	⋮

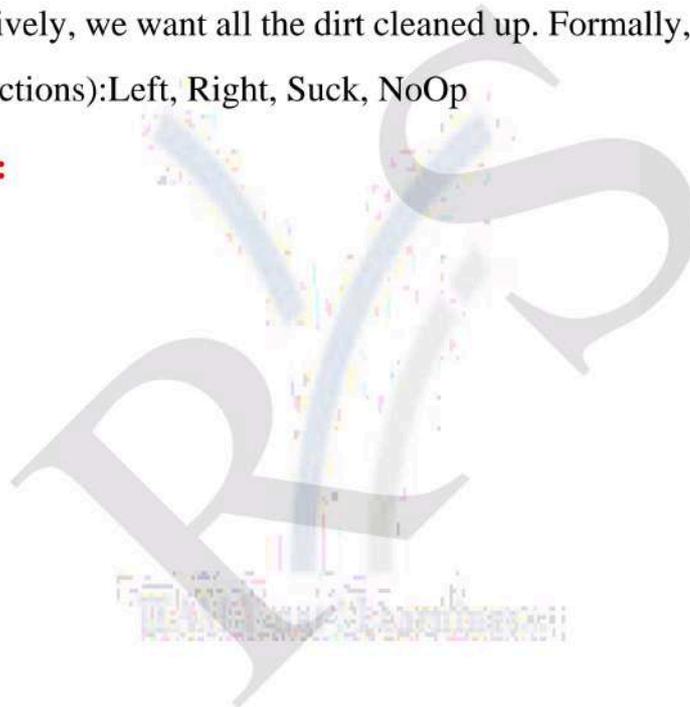
Partial tabulation of a simple agent function for the vacuum-cleaner world

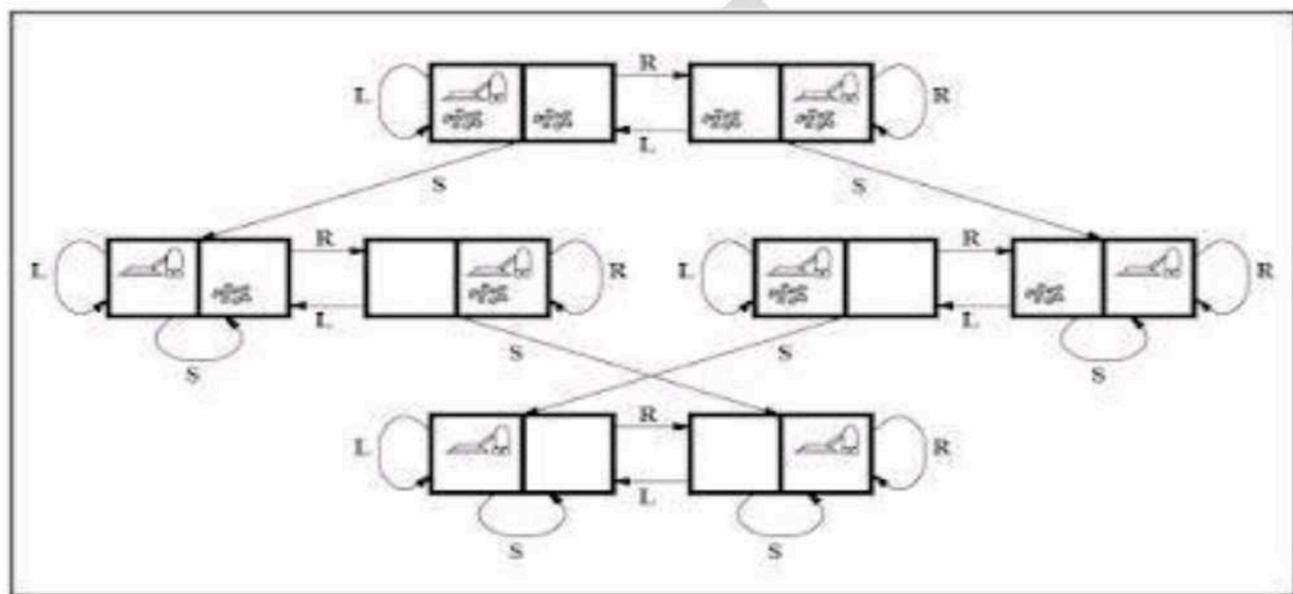


Goal formulation: intuitively, we want all the dirt cleaned up. Formally, the goal is {State 7, state 8}.

Problem formulation (Actions):Left, Right, Suck, NoOp

STATE SPACE GRAPH:





```
function REFLEX-VACUUM-AGENT( [location,status] ) returns  
action  
  
  if status = Dirty then return Suck  
  else if location = A then return Right  
  else if location = B then return Left
```

MEASURING PERFORMANCE

With any intelligent agent, we want it to find a (good) solution and not spend forever doing it.

The interesting quantities are, therefore,

THE SEARCH COST--how long the agent takes to come up with the solution to the problem,
and

THE PATH COST--how expensive the actions in the solution are.

The total cost of the solution is the sum of the above two quantities.

HOW SHOULD THE AGENTS ACT IN ARTIFICIAL INTELLIGENCE?

- A rational agent does the right thing. The right action is the one that causes the agent to be the most successful.
- An omniscient agent knows what impact the action will have and can act accordingly, but it is not possible in reality.
- The degree of success which is defined by the performance measure
- The percept sequence which is the entire sequence of perceptions by the agent until the present moment

Agent Type	Percepts	Action	Goals	Environment
-------------------	-----------------	---------------	--------------	--------------------

Car Driver	Speedometer, GPS, Microphone, Cameras	Steering control, accelerate, brake, talk to passenger	Safe, legal, comfortable journey	Road, Traffic, Pedestrian etc.
-------------------	---------------------------------------------	--------------------------------------------------------------	-------------------------------------	-----------------------------------

GOOD BEHAVIOUR: THE CONCEPT OF RATIONALITY

INTELLIGENT AGENTS:

An **intelligent agent** is an autonomous entity which act upon an environment using sensors and actuators for achieving goals. An intelligent agent may learn from the environment to achieve their goals. A thermostat is an example of an intelligent agent.

Following are the main four rules for an AI agent:

- > **Rule 1:** An AI agent must have the ability to perceive the environment.
- > **Rule 2:** The observation must be used to make decisions.

- **Rule 3:** Decision should result in an action.
- **Rule 4:** The action taken by an AI agent must be a rational action.

RATIONAL AGENT:

- A rational agent is an agent which has **clear preference**, models uncertainty, and acts in a way to maximize its **performance measure with all possible actions**.
- A rational agent is said to perform the **right things**. AI is about creating **rational agents** to use for **game theory** and **decision theory** for **various real- world scenarios**.
- For an AI agent, the **rational action is most important** because in AI reinforcement learning algorithm, **for each best possible action**, agent **gets the positive reward** and for each wrong action, an **agent gets a negative reward**.

NOTE: Rational agents in AI are very similar to intelligent agents.

RATIONALITY:

The rationality of an agent is measured by its performance measure. Rationality can be judged on the basis of following points:

- Performance measure which defines the success criterion.
- Agent prior knowledge of its environment.
- Best possible actions that an agent can perform.
- The actions that the agent can perform.
- The agent's percept sequence to date.

NOTE: Rationality differs from Omniscience because an Omniscient agent knows the actual outcome of its action and act accordingly, which is not possible in reality.

MAPPING OF PERCEPT SEQUENCES TO ACTIONS

When it is known that the action of agent depends completely on the perceptual history – the percept sequence, then the agent can be described by using a mapping. Mapping is a list that maps the percept sequence to the action. When we specify which action an agent should take corresponding to the given percept sequence, we specify the design for an ideal agent.

AUTONOMY

The behaviour of an agent depends on its own experience as well as the built-in knowledge of the agent instilled by the agent designer. A system is autonomous if it takes actions according to its experience. So for the initial phase, as it does not have any experience, it is good to provide built-in knowledge. The agent learns then through evolution. A truly autonomous intelligent agent should be

able to operate successfully in a wide variety of environments if given sufficient time to adapt.

RSS

TASK ENVIRONMENTS

To design a rational agent we need to specify a task environment

- A problem specification for which the agent is a solution

PEAS Representation

PEAS is a type of model on which an AI agent works upon. When we define an AI agent or rational agent, then we can group its properties under PEAS representation model. It is made up of four words:

- **P:** Performance measure
- **E:** Environment
- **A:** Actuators



➤ **S: Sensors**

Here performance measure is the objective for the success of an agent's behaviour.

RRSS

PEAS: SPECIFYING AN AUTOMATED TAXI DRIVER

Performance measure: ?

Environment: ?

Actuators: ?

Sensors: ?

Performance measure:

- safe, fast, legal, comfortable, maximize profits

Environment:

- roads, other traffic, pedestrians, customers

Actuators:

- steering, accelerator, brake, signal, horn



Sensors:

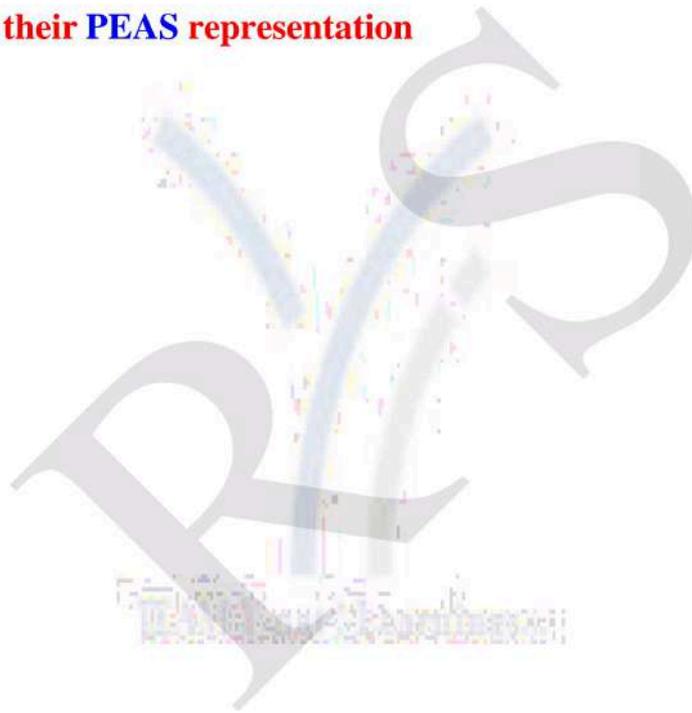
- cameras, sonar, speedometer, GPS

RRSS

PEAS: MEDICAL DIAGNOSIS SYSTEM

- **PERFORMANCE MEASURE:** Healthy patient, minimize costs, lawsuits
- **ENVIRONMENT:** Patient, hospital, staff
- **ACTUATORS:** Screen display (form including: questions, tests, diagnoses, treatments, referrals)
- **SENSORS:** Keyboard (entry of symptoms, findings, patient's answers)

Example of Agents with their PEAS representation



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

Examples of agent types and their PEAS descriptions.



THE
NATURE
OF
ENVIRONMENTS

THE NATURE OF ENVIRONMENTS

- 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**.
- An environment is **everything** in the world **which surrounds the agent**, but it is not a part of an **agent itself**. An environment can be described as a situation in which an agent is present.
- The environment is where agent lives, operate and provide the agent with something to sense and act upon it. **An environment is mostly said to be non-feministic.**

FEATURES OF ENVIRONMENT

As per Russell and Norvig, an environment can have various features from the point of view of an agent:

1. Fully observable vs Partially Observable
2. Static vs Dynamic
3. Discrete vs Continuous
4. Deterministic vs Stochastic
5. Single-agent vs Multi-agent
6. Episodic vs sequential

7. Known vs Unknown

8. Accessible vs Inaccessible

RRSS

1. FULLY OBSERVABLE VS PARTIALLY OBSERVABLE:

- ✓ If an agent sensor can sense or access the complete state of an environment at each point of time then it is a **fully observable** environment, else it is **partially observable**.
- ✓ A fully observable environment is easy as there is no need to maintain the internal state to keep track history of the world.
- ✓ An agent with no sensors in all environments then such an environment is called as **unobservable**.

2. DETERMINISTIC VS STOCHASTIC:

- ✓ If an agent's current state and selected action can completely determine the next state of the environment, then such environment is called a deterministic environment.

- ✓ A stochastic environment is random in nature and cannot be determined completely by an agent.
- ✓ In a deterministic, fully observable environment, agent does not need to worry about uncertainty.

RRSS

3. EPISODIC VS SEQUENTIAL:

- ✓ In an episodic environment, there is a series of one-shot actions, and only the current percept is required for the action.
- ✓ However, in Sequential environment, an agent requires memory of past actions to determine the next best actions.

4. SINGLE-AGENT VS MULTI-AGENT

- ✓ If only one agent is involved in an environment, and operating by itself then such an environment is called single agent environment.
- ✓ However, if multiple agents are operating in an environment, then such an environment is called a multi-agent environment.

- ✓ The agent design problems in the multi-agent environment are different from single agent environment.

RRS

5. STATIC VS DYNAMIC:

- ✓ If the environment can change itself while an agent is deliberating then such environment is called a dynamic environment else it is called a static environment.
- ✓ Static environments are easy to deal because an agent does not need to continue looking at the world while deciding for an action.
- ✓ However for dynamic environment, agents need to keep looking at the world at each action.
- ✓ Taxi driving is an example of a dynamic environment whereas Crossword puzzles are an example of a static environment.

6. DISCRETE VS CONTINUOUS:

- ✓ If in an environment there are a finite number of percepts and actions that can be performed within it, then such an environment is called a discrete environment else it is called continuous environment.
- ✓ A chess game comes under discrete environment as there is a finite number of moves that can be performed.
- ✓ A self-driving car is an example of a continuous environment.

7. ACCESSIBLE VS INACCESSIBLE

- ✓ If an agent can obtain complete and accurate information about the state's environment, then such an environment is called an Accessible environment else it is called inaccessible.

- ✓ An empty room whose state can be defined by its temperature is an example of an accessible environment.

RRSS

8. 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.
- ✓ It is quite possible that a known environment to be partially observable and an Unknown environment to be fully observable.

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

Examples of task environments and their characteristics.



THE
STRUCTURE
OF
AGENTS

THE STRUCTURE OF AGENTS

The task of AI is to design an agent program which implements the agent function. The structure of an intelligent agent is a combination of architecture and agent program. It can be viewed as:

$$\text{Agent} = \text{Architecture} + \text{Agent program}$$

Architecture: Architecture is machinery that an AI agent executes on.

Agent Function: Agent function is used to map a percept to an action.

$$f: P^* \rightarrow A$$

Agent program: Agent program is an implementation of agent function. An agent program executes on the physical architecture to produce function f .

TYPES OF AI AGENTS

Agents can be grouped into five classes based on their degree of perceived intelligence and capability. All these agents can improve their performance and generate better action over the time.

These are given below:

- Simple Reflex Agent
- Model-based reflex agent
- Goal-based agents
- Utility-based agent

- Learning agent

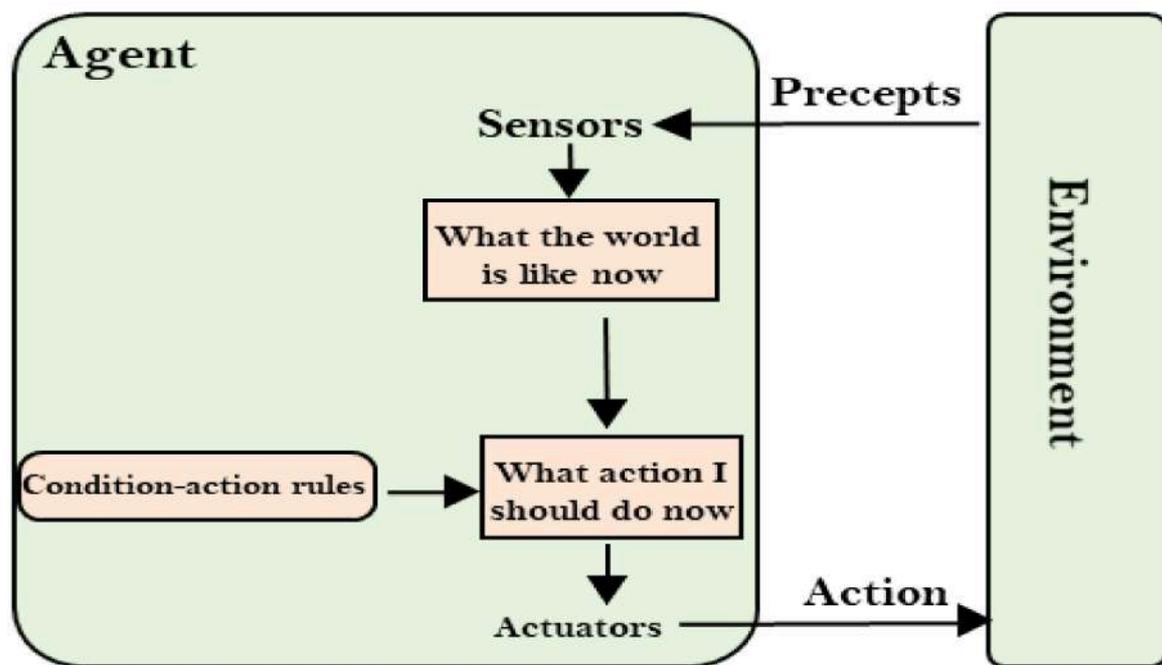
RRS

1. SIMPLE REFLEX AGENT:

- 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.

RRSS



Schematic diagram of a simple reflex agent.

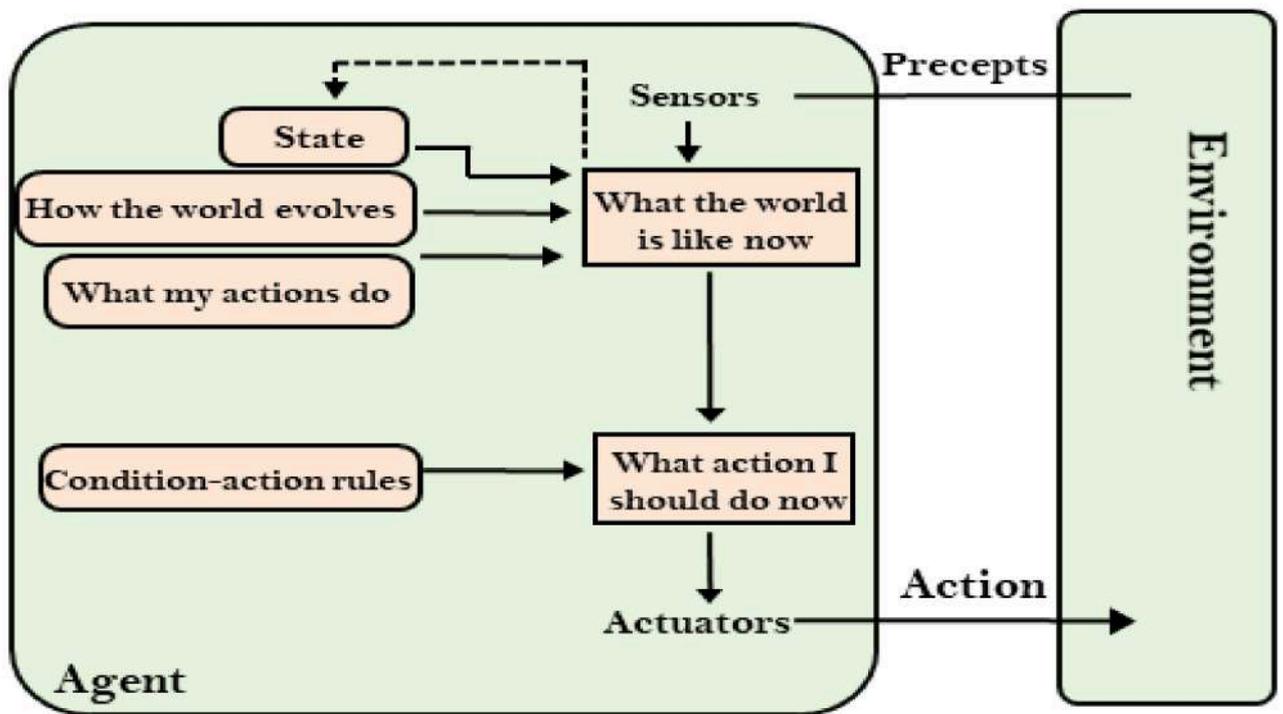
RRS

2. MODEL-BASED REFLEX AGENT

- The Model-based agent can work in a partially observable environment, and track the situation.
- A model-based agent has two important factors:
 - ✓ **Model:** It is knowledge about "how things happen in the world," so it is called a Model-based agent.
 - ✓ **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:

- ✓ How the world evolves
- ✓ How the agent's action affects the world.

RS



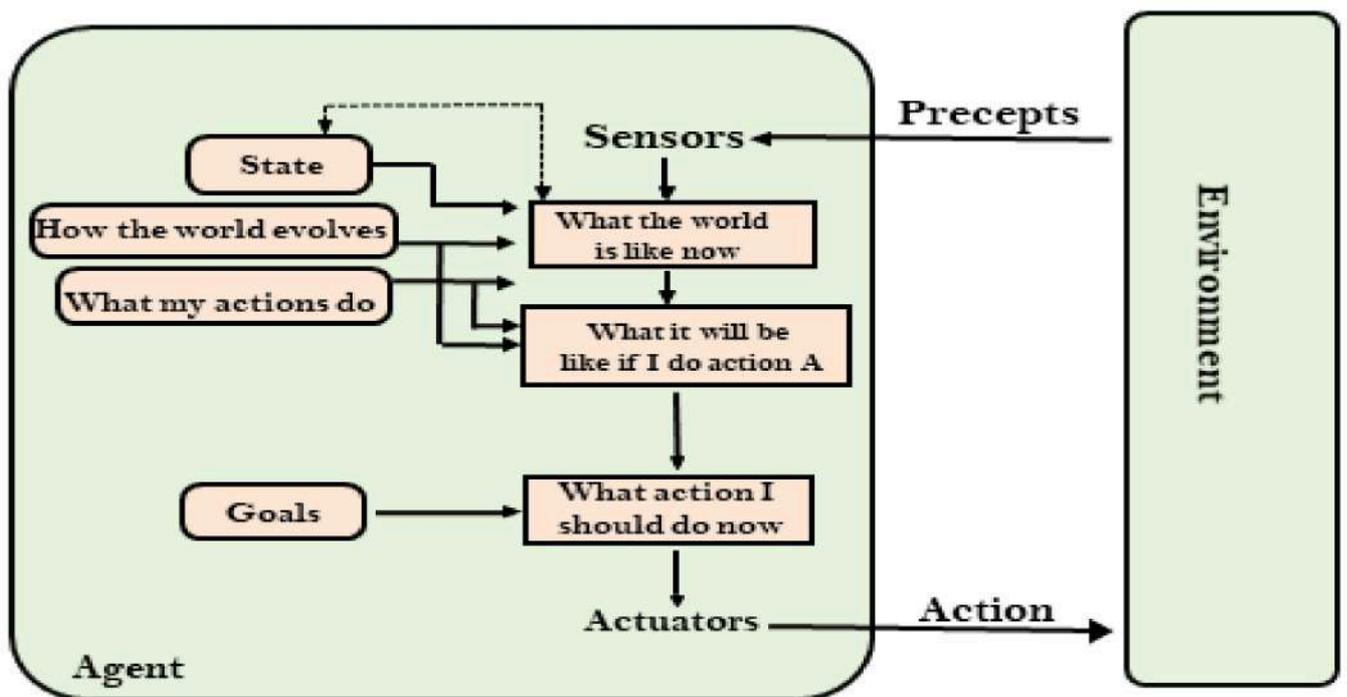
A model-based reflex agent.

RRS

3. GOAL-BASED AGENTS

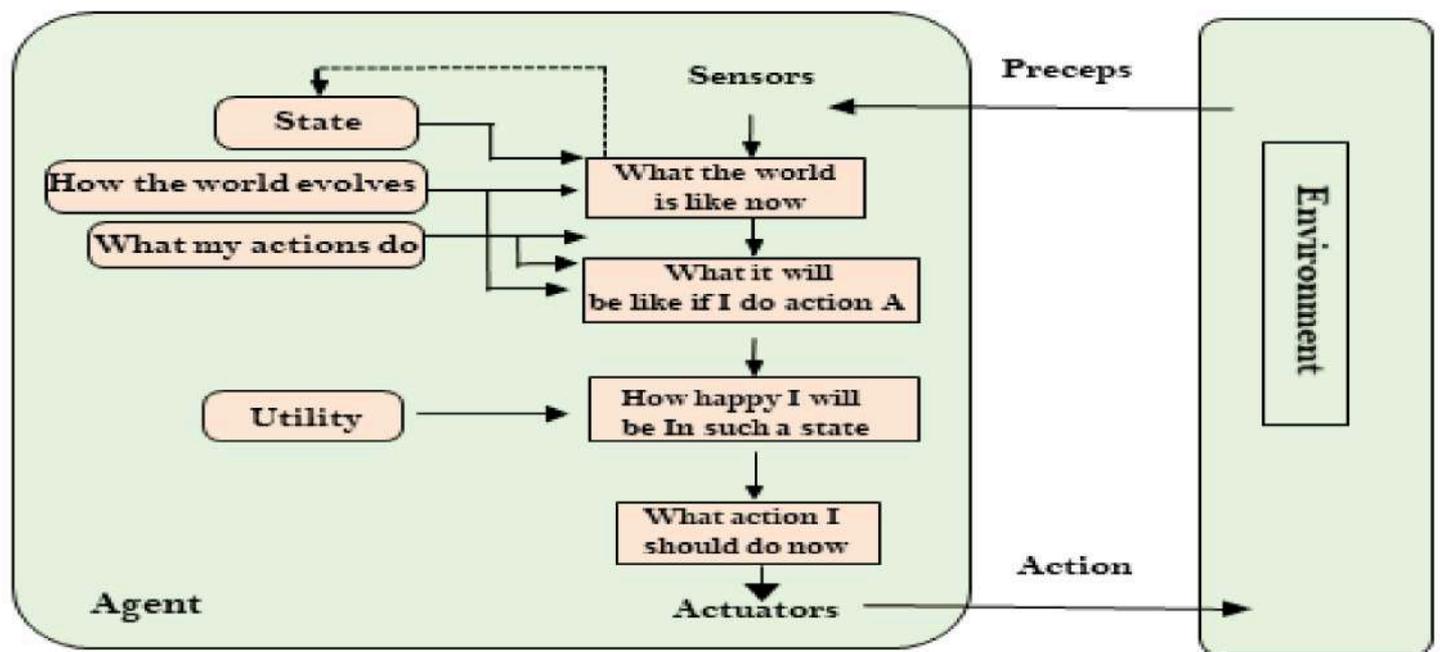
- The knowledge of the current state environment is not always sufficient to decide for an agent to what to do.
- The agent needs to know its goal which describes desirable situations.
- Goal-based agents expand the capabilities of the model-based agent by having the "goal" information.
- They choose 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.



4. 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.

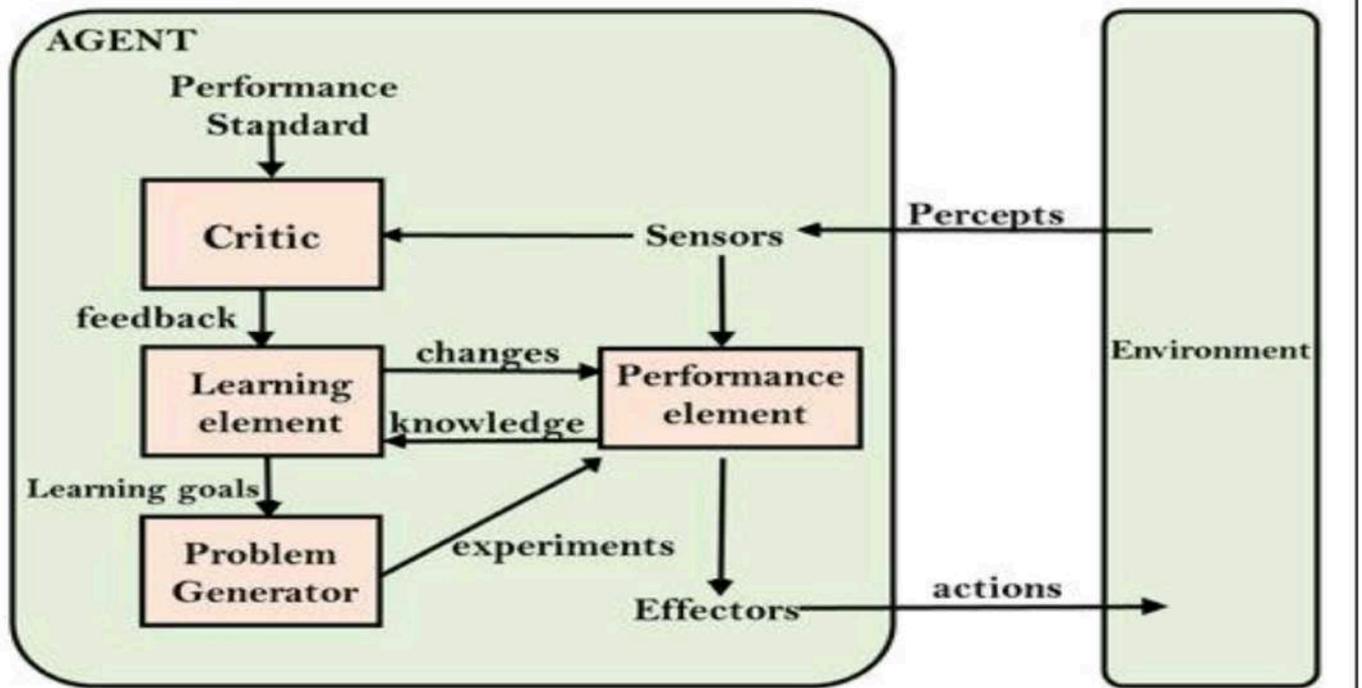


5. LEARNING AGENTS

- 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:
 - ✓ **Learning element:** It is responsible for making improvements by learning from environment
 - ✓ **Critic:** Learning element takes feedback from critic which describes that how well the agent is doing with respect to a fixed performance standard.
 - ✓ **Performance element:** It is responsible for selecting external action
 - ✓ **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.

RRSS



A general learning agent.

RRS

