UNIT

1

AN OVERVIEW OF ARTIFICIAL INTELLIGENCE

1.1. DEFINITIONS OF ARTIFICIAL INTELLIGENCE

Artificial Intelligence is the branch of computer science that is concerned with the automation of intelligent behaviour.

Artificial Intelligence is the study and design of intelligent agents, where an intelligent agent is a system that perceives its environment and takes actions that maximize its chance of success.

Artificial Intelligence is the science and engineering of making intelligent machines.

Artificial Intelligence is the study of how to make computers do things which, at the moment, people do better.

Artificial intelligence (AI) refers to the simulation of human intelligence in machines that are programmed to think like humans and mimic their actions.

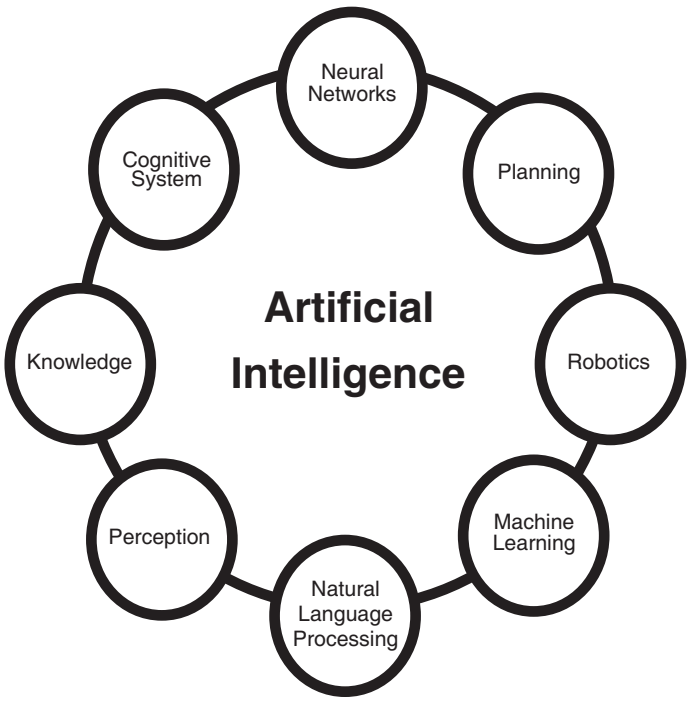

Fig. 1.1.

AI initiatives are often talked about in terms of their belonging to one of four categories:

Reactive AI relies on real-time data to make decisions.

- Limited Memory AI relies on stored data to make decisions.
- Theory of Mind AI can consider subjective elements such as user intent when making decisions.
- Self-Aware AI possesses a human-like consciousness that is capable of independently setting goals and using data to decide the best way to achieve an objective.

WHAT IS INTELLIGENCE?

Intelligence is what we use when we don't know what to do.

Intelligence relates to tasks involving higher mental processes.

The ability to understand, learn and think.

Intelligence has been defined in many ways: the capacity for logic, understanding, self-awareness, learning, emotional knowledge, reasoning, planning, creativity, critical thinking, and problem-solving.

Examples: Creativity, Solving problems, Pattern recognition, Classification, Learning, Induction, Deduction, Building analogies, Optimization, Language processing, Knowledge and many more.

Human intelligence—Human intelligence, mental quality that consists of the abilities to learn from experience, adapt to new situations, understand and handle abstract concepts, and use knowledge to manipulate *one's environment*.

Difference between Artificial Intelligence and Human Intelligence

- **Origin of AI and Human Intelligence**

AI is an innovation created by human intelligence; its early development is credited to Norbert Weiner who theorized on feedback mechanisms while the father of AI is John McCarthy for coining the term and organizing the first conference on research projects regarding machine intelligence. On the other hand, human beings are created with the innate ability to think, reason, recall, etc.

- **Speed of AI and Human Intelligence**

As compared to humans, computers can process more information at a faster rate. For instance, if the human mind can solve a math problem in 5 minutes, AI can solve 10 problems in a minute.

- **Decision Making**

AI is highly objective in decision making as it analyzes based on purely gathered data. However, humans' decisions may be influenced by subjective elements which are not based on figures alone.

- **Accuracy**

AI often produces accurate results as it functions based on a set of programmed rules. As for human intelligence, there is usually a room for "human error" as certain details may be missed at one point or the other.

- **Energy Used**

The human brain uses about 25 watts while modern computers only generally use 2 watts.

- **Adaptation of AI and Human Intelligence**

Human intelligence can be flexible in response to the changes to its environment. This makes people able to learn and master various skills. On the other hand, AI takes much more time to adapt to new changes.

- **Multitasking**

The human intellect supports multitasking as evidenced by diverse and simultaneous roles while AI can only perform fewer tasks at the same time as a system can only learn responsibilities one at a time.

- **Self-Awareness**

AI is still working on its ability regarding self-awareness while humans become naturally aware of themselves and strive to establish their identities as they mature.

- **Social Interaction**

As social beings, humans are much better at social interaction since they can process abstract information, have self-awareness, and are sensitive to others' emotions. On the other hand, AI has not mastered the ability to pick up on pertinent social and emotional cues.

- **General Function**

The general function of human intelligence is innovation as it can create, collaborate, brainstorm, and implement. As for AI, its general function is more on optimization as it efficiently performs tasks according to how it is programmed.

1.1.1. Turing Test

The test was introduced by Turing in his 1950 paper "*Computing Machinery and Intelligence*" while working at the *University of Manchester*. It opens with the words: "I propose to consider the question", **'Can machines think'?**

The Turing test, originally called the **imitation game** by ***Alan Turing*** **in 1950**, is a test of a machine's ability to *exhibit intelligent behaviour* equivalent to, or indistinguishable from, that of a human. Turing proposed that a human evaluator would judge natural language conversations between a human and a machine designed to generate human-like responses. The evaluator would be aware that one of

the two partners in conversation is a machine, and all participants would be separated from one another. The conversation would be limited to a text-only channel such as a computer keyboard and screen so the result would not depend on the machine's ability to render words as speech. If the evaluator cannot reliably tell the machine from the human, the machine is said to have passed the test. The test results do not depend on the machine's ability to give correct *answers to questions*, only how closely its answers resemble those a human would give.

The Turing test involves three participants: a computer, a human interrogator, and a human foil. The interrogator attempts to determine, by asking questions of the other two participants, which is the computer.

The earliest successful AI program was written in 1951 by Christopher Strachey. The first AI program to run in the United States also was a checkers program, written in 1952 by Arthur Samuel for the *prototype* of the *IBM* 701.

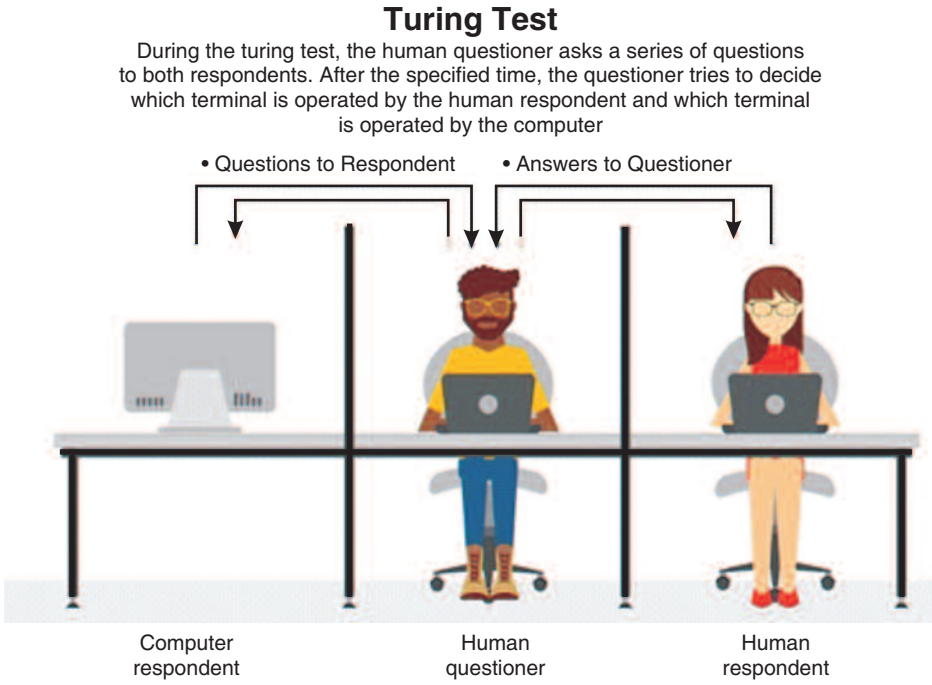

Fig. 1.2. Turing Test

1.1.2. Chinese Room Test

The **Chinese room argument** holds that a digital computer executing a program cannot be shown to have a "*mind*", "*understanding*" or "*consciousness*", regardless of how intelligently or human-like the program may make the computer behave. The argument was first presented by philosopher *John Searle* in his paper, "Minds, Brains, and Programs", published in *Behavioural and Brain Sciences* in 1980.

The Chinese Room conundrum argues that a computer cannot have a mind of its own programmed to mimic the activities of a conscious human being but they can't have an understanding of what they are simulating on their own.

"A human mind has meaningful thoughts, feelings, and mental contents generally. Formal symbols by themselves can never be enough for mental contents, because the symbols, by definition, have no meaning," *said* Searle when questioned about his argument.

Searle's thought experiment goes like this: Suppose a closed room has a non-Chinese speaker with a list of Mandarin characters and an instruction book. This book explains in detail the rules according to which the strings (sequences) of characters may be formed but without giving the meaning of the characters.

Suppose now that we pass to this man through a hole in the wall a sequence of Mandarin characters which he is to complete by following the rules he has learned. We may call the sequence passed to him from the outside a "question" and the completion an "answer."

Now, this non-Chinese speaker masters this sequencing game so much that even a native Chinese person will not be able to spot any difference in the answers given by this man in an enclosed room.

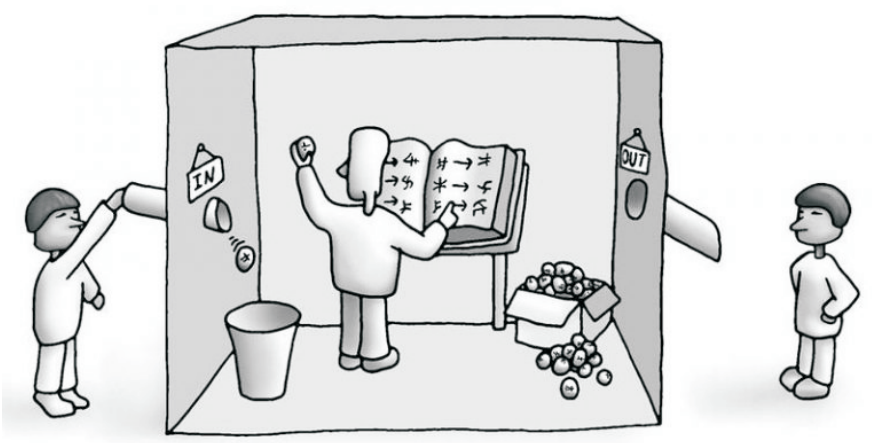

Fig. 1.3. Chinese room Test

But the fact remains that not only is he not Chinese, but he does not even understand Chinese, far less *think* in it.

Now, the argument goes on, a machine, even a Turing machine, is just like this man, in that it does nothing more than follow the rules given in an instruction book (the program). It does not understand the meaning of the questions given to it nor its own answers, and thus cannot be said to be *thinking*.

Making a case for Searle, if we accept that a book has no mind of its own, we cannot then endow a computer with intelligence and remain consistent.

The question Searle wants to answer is this: does the machine *literally* "understand" Chinese? Or is it merely simulating the ability to understand Chinese? Searle calls the first position ***"strong AI"*** and the latter "weak AI".

1.1.3. 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.

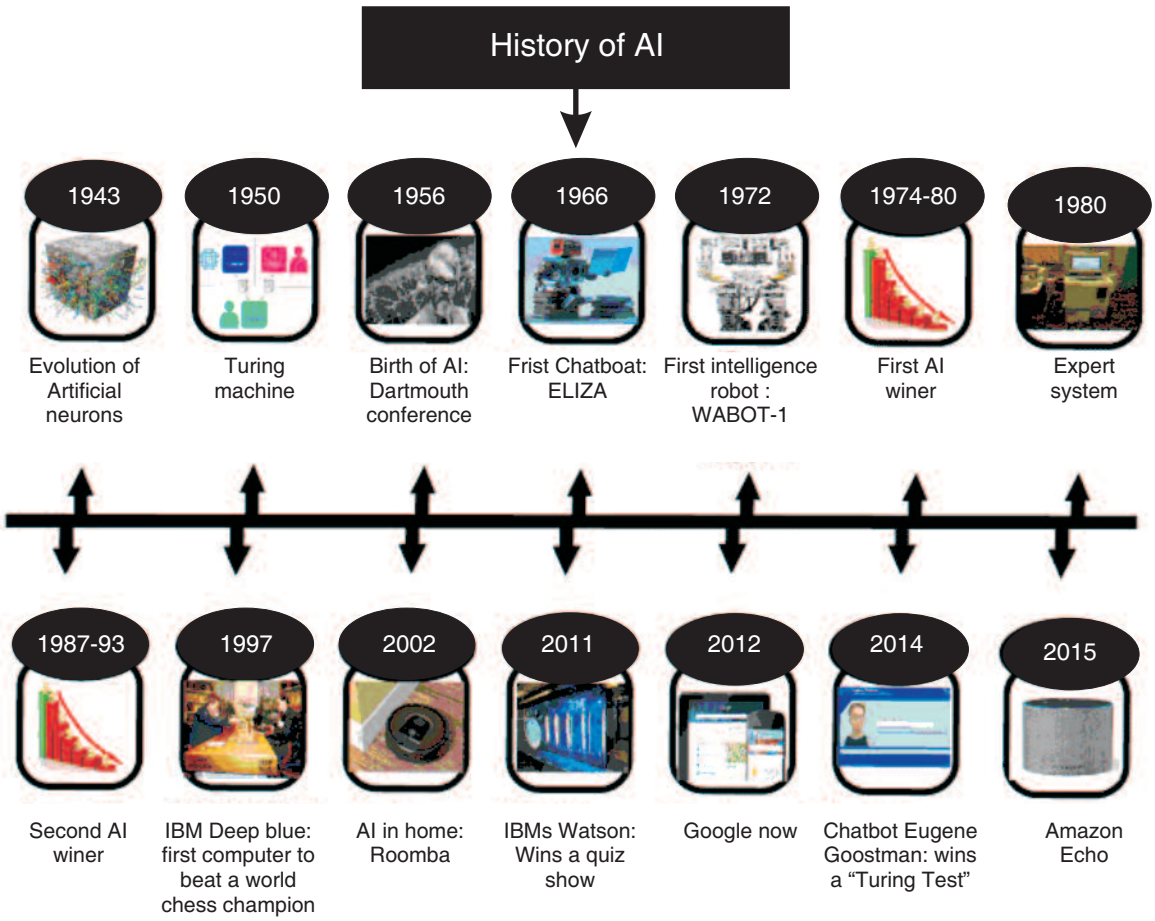

Fig. 1.4. History of AI

Maturation of Artificial Intelligence (1943-1952)

- **Year 1943:** The first work which is now recognized as AI was done by Warren McCulloch and Walter pits 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:** The Alan Turing who was an English mathematician and pioneered Machine learning in 1950. Alan Turing publishes **"Computing Machinery and Intelligence"** in which he proposed a test. The test can check the machine's

ability to exhibit intelligent behaviour equivalent to human intelligence, called a **Turing test.**

The birth of Artificial Intelligence (1952-1956)

- **Year 1955:** An 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

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 Goostman" 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.

Now AI has developed to a remarkable level. The concept of Deep learning, big data, and data science are now trending like a boom. Now-a-days companies like Google, Face book, IBM, and Amazon are working with AI and creating amazing devices. The future of Artificial Intelligence is inspiring and will come with high intelligence.

Goals of AI

The definition of AI gives four possible goals to pursue:

- Systems that think like humans
- Systems that think rationally
- Systems that act like humans
- Systems that act rationally

1.2. VARIOUS APPROACHES TO AI

- Hard or Strong AI
- Soft or Weak AI
- Applied AI
- Cognitive AI

Hard or Strong AI

Strong AI refers to a machine that approaches or supersedes human intelligence. if it can do typically human tasks, If it can apply a wide range of background Knowledge and If it has some degree of self-consciousness. Strong AI aims to build machines whose overall intellectual ability is Indistinguishable from that of a human being.

Soft or Weak AI

Weak AI refers to the use of software to study or accomplish specific problem solving or reasoning tasks that do not encompass the full range of human cognitive abilities. Example: a chess program such as Deep Blue.

Weak AI does not achieve self-awareness; it demonstrates wide range of human level cognitive abilities; it is merely an intelligent, a specific problem-solver.

Applied AI

Applied AI is the branch of artificial intelligence that brings it out of the lab and into the real world, enabling computers and computer-controlled robots to execute real tasks. Applied AI enhances software applications and puts advanced machine learning to use, providing high levels of accuracy and adaptation over time.

Applied AI is contextualizing business models and industry processes, as well as improving the way we interact with everything around us.

Cognitive AI

Computers are used to test theories about how the human mind works—for example theories about how we recognize faces and other objects, or about how we solve abstract problems.

Cognitive science

Aims to develop explore and evaluate theories of how the mind works through the use of computational models.

The important is not what is done but how it is done; means intelligent behaviour is not enough, the program must operate in an intelligent manner.

EXAMPLE

The chess programs are successful, but say little about the ways humans play chess.

Cybernetics

"Cybernetics" comes from a Greek word meaning "the art of steering".

Cybernetics is about having and taking action to achieve that goal. Knowing whether you have reached your goal (or at least are getting closer to it) requires, a concept that comes from cybernetics.

Cybernetics grew from a desire to understand and build systems that can achieve goals, whether complex human goals or just goals like maintaining the temperature of a room under changing conditions.

1.3. WHAT SHOULD ALL ENGINEERS KNOW ABOUT AI?

Artificial intelligence (AI) systems by their nature are software-intensive. To create viable and trusted AI systems, engineers need technologies and standards, similar to those in software engineering.

Introductory Concepts: AI, ML, and Deep Learning

At a very basic level, many practitioners use the terms AI and machine learning as if they are separate entities. They are not. The *DoD AI Strategy* defines AI as

The ability of machines to perform tasks that normally require human intelligence for example, recognizing patterns, learning from experience, drawing conclusions, making predictions, or taking action whether digitally or as the smart software behind autonomous physical systems.

Machine learning (ML), a part of AI, which the SEI defines as:

A system that learns and improves its performance at some task by using data and experience. Further a *recent SEI blog post* defined *deep learning as*

A family of machine learning techniques whose models extract important features by iteratively transforming the data, "going deeper" toward meaningful patterns in the dataset with each transformation. Unlike traditional machine learning methods, in which the creator of the model has to choose and encode features ahead of time, deep learning enables a model to automatically learn features that matter.

AI Engineering Concepts

AI depends on the human element. AI augments, but does not replace, human knowledge and expertise. This basic understanding affects engineers of AI systems in two dimensions: human-machine teaming and the probabilistic nature of AI "answers." Engineers developing AI systems must account for human-machine teaming—the interactions between the system and the people who build and use it. Often, the success of those interactions comes down to trust and transparency: How should AI systems be deployed in environments where people have become accustomed to ignoring automation? How can you address ethics—accounting for algorithms not having a sense of morality? Further, AI will produce probabilistic answers: How does an AI present result to a human as based on a probability distribution, not as a discrete answer? How does the human know when a prediction is bad?

AI depends on labelled and unlabelled data as well as the systems that store and access it. The availability of data and the speed at which today's computers can process it are reasons why AI is exploding today. AI systems are really good at classifying, categorizing, and partitioning massive amounts of data to make the most relevant pieces available for humans to analyze and make decisions. Engineers must

consider the data itself—provenance, security, quality, and aligning test and training data--and the hardware and software systems that support that data. Large amounts of data require a computing environment that has the capacity to handle it. Managing data requires designing storage solutions around physical data constraints and types of queries desired.

One AI, many algorithms. When we talk about AI, ML, and deep learning, we are referring to many different algorithms, many different approaches, not all of which are neural-network based. AI is not a new field, and many of the algorithms in use today were generated in the 1950s, 1960s, or 1970s. For example, *the A* shortest path algorithm* was conceived in the 1950s, and improved on in the 1960s.

The insight is the benefit of AI. Engineers face the reality that it is impossible to test a system in every situation it will ever encounter. An AI system adds capability for the engineering because it can find an answer to never-seen-before situations that is insightful and has a very good probability of being correct. However, it is not necessarily correct, but probabilistic. Thus, gaining increased confidence in AI is hard for engineers who need to focus on creating and validating a system.

An AI system depends on the system under which it runs. When building a system that does not incorporate AI, you can build it in isolation, test it in isolation, and then deploy it and be certain it is going to behave just as it did in the lab. An AI system depends on the conditions under which the AI runs and what the AI system is sensing, and this context adds another level of complexity.

1.4. OTHER EMERGING TECHNOLOGIES IN AI

Even though AI is poised to "completely reframe how businesses operate and consumers interact" many firms are not ready to invest. The two top reasons for the standstill are:

We are unclear what AI can be used for in our business

There is no defined business case for AI investment

One reason behind the standstill is the *AI landscape is complex*. Navigating the AI ecosystem and prioritizing investment is becoming increasingly difficult.

To help us answer these questions and navigate AI, *Forrester recently analyzed the 12* most important AI categories for customer insights.

The 12 AI technologies Forrester evaluated are:

- AI-enhanced analytics solutions
- Conversational service solutions
- Deep learning platforms
- Facial recognition
- Image and video analysis
- Intelligent recommendation solutions

- Intelligent research solutions
- Machine learning platforms
- Natural language generation
- Pre-trained vertical solutions
- Speech analytics
- Text analytics

1. Natural language generation

Machines process and communicate in a different way than the human brain. *Natural language generation* is a trendy technology that converts the structured data into the native language. The machines are programmed with algorithms to convert the data in a desirable format of the user. Natural language is a subset of artificial intelligence which helps the content developers to automate content and deliver in the desired format. The content developers can use the automated content to promote on various social media platforms, and other media platforms to reach the targeted audience. Human intervention will significantly reduce as data will be converted into desired formats. The data can be visualized in the form of charts, graphs etc.

2. Speech recognition

Speech recognition is another important subset of artificial intelligence which converts human speech into a useful and understandable format by computers. Speech recognition is a bridge between human and computer interactions. The technology recognizes and converts human speech in several languages. Siri of iPhone is a classic example of speech recognition.

3. Virtual agents

Virtual agents have become valuable tools for instructional designers. A *virtual agent* is a computer application that interacts with humans. Web and mobile applications provide chatbots as their customer service agents to interact with humans to answer their queries. Google Assistant helps to organize meetings, and Alexia from Amazon helps to make your shopping easy. A virtual assistant also acts like a language assistant, which picks cues from your choice and preference. The IBM Watson understands the typical customer service queries which are asked in several ways. Virtual agents act as software-as-a-service too.

4. Decision management

Modern organizations are implementing decision management systems for data conversion and interpretation into predictive models. Enterprise-level applications implement decision management systems to receive up-to-date information to perform business data analysis to aid in organizational decision making. Decision management helps in making quick decisions, avoidance of risks, and in the automation of the

process. The decision management system is widely implemented in the financial sector, the health care sector, trading, insurance sector, e-commerce, etc.

5. Biometrics

Deep learning another branch of artificial intelligence which functions based on *artificial neural networks*. This technique teaches computers and machines to learn by example just the way humans do. The term "deep" is coined because it has hidden layers in neural networks. Typically, a neural network has 2-3 hidden layers and can have a maximum of 150 hidden layers. Deep learning is effective on huge data to train a model and a graphic processing unit. The algorithms work in a hierarchy to automate predictive analytics. Deep learning has spread its wings in many domains like aerospace and military to detect objects from satellites, helps in improving worker safety by identifying risk incidents when a worker gets close to a machine, helps to detect cancer cells etc.

6. Machine learning

Machine learning enables a machine to automatically learn from data, improve performance from experiences, and predict things without being explicitly programmed Machine Learning is said as a subset of **artificial intelligence** that is mainly concerned with the development of algorithms which allow a computer to learn from the data and past experiences on their own. The term machine learning was first introduced by **Arthur Samuel** in **1959**.

A Machine Learning system learns from historical data, builds the prediction models, and whenever it receives new data, predicts the output for it.

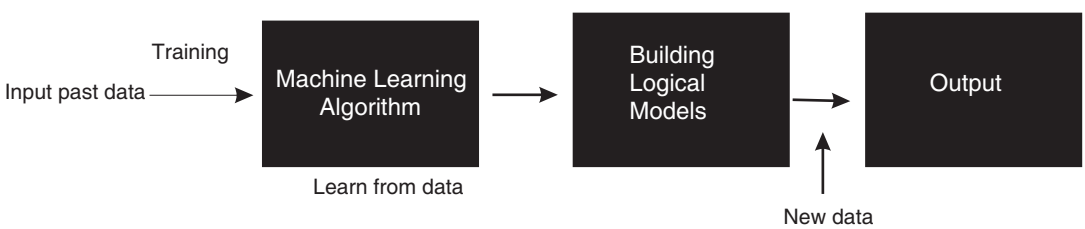

Fig. 1.5.

Machine learning is a division of artificial intelligence which empowers machine to make sense from data sets without being actually programmed. Machine learning technique helps businesses to make informed decisions with data analytics performed using algorithms and statistical models. Enterprises are investing heavily in machine learning to reap the benefits of its application in diverse domains. Healthcare and medical profession need machine learning techniques to analyze patient data for the prediction of diseases and effective treatment. The banking and financial sector needs machine learning for customer data analysis to identify and suggest investment options to customers and for risk and fraud prevention. Retailers utilize machine learning for predicting changing customer preferences, consumer behaviour, by analyzing customer data.

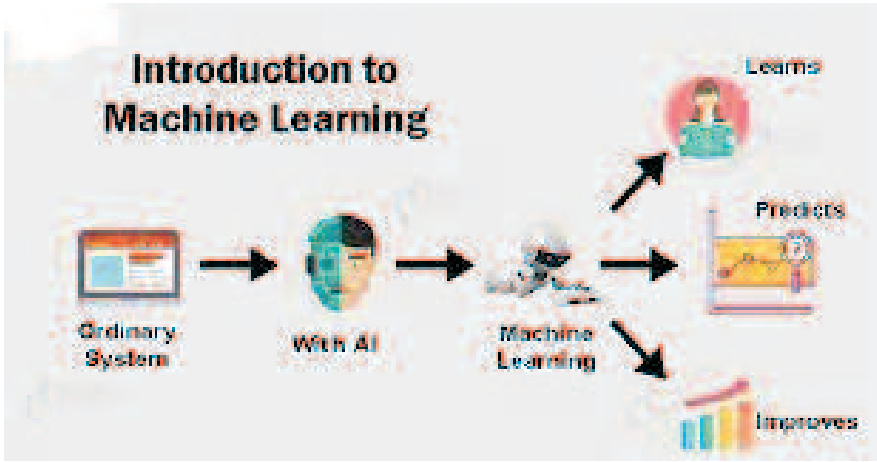

Fig. 1.6.

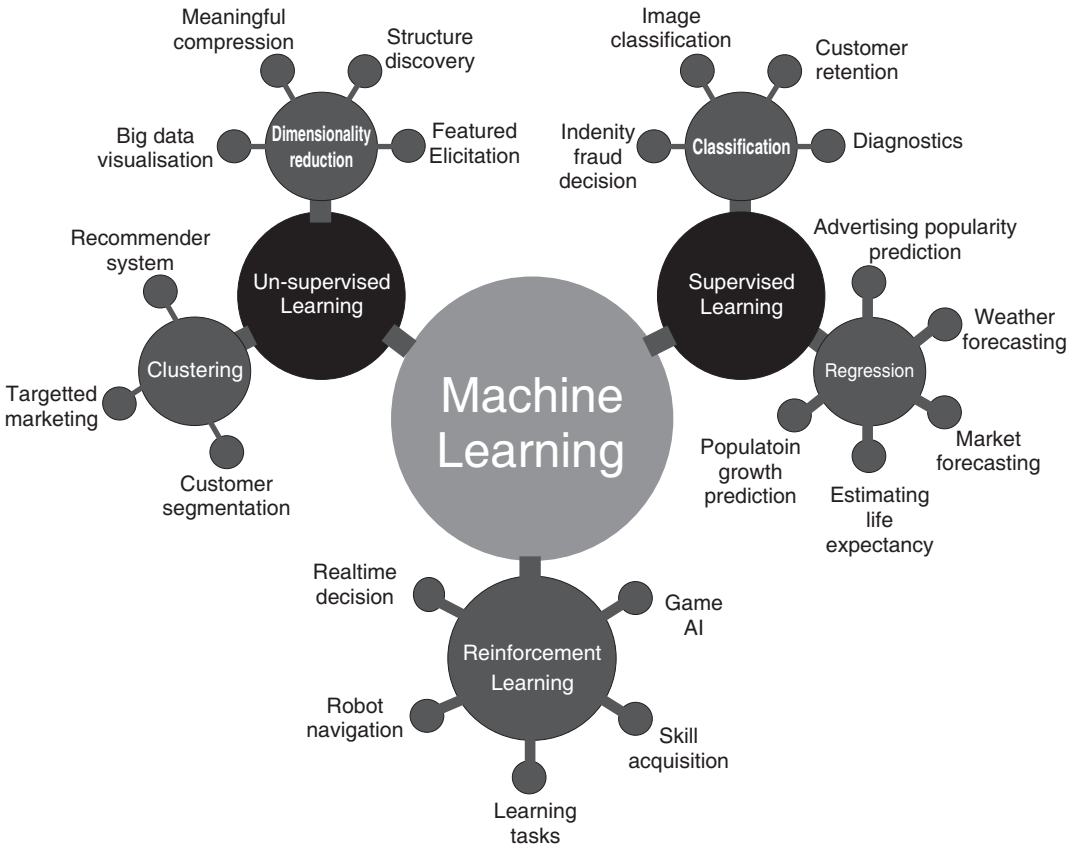

Fig. 1.7.

Machine learning is a growing technology which enables computers to learn automatically from past data. Machine learning uses various algorithms for building mathematical models and making predictions using historical data or information. Currently, it is being used for various tasks such as image recognition, speech

recognition, email filtering, Face book auto-tagging, recommender system, and many more.

History of Machine Learning

Before some years (about 40-50 years), machine learning was science fiction, but today it is the part of our daily life. Machine learning is making our day-to-day life easy from self-driving cars to Amazon virtual assistant **"Alexa"**. However, the idea behind machine learning is so old and has a long history. Below some milestones are given which have occurred in the history of machine learning:

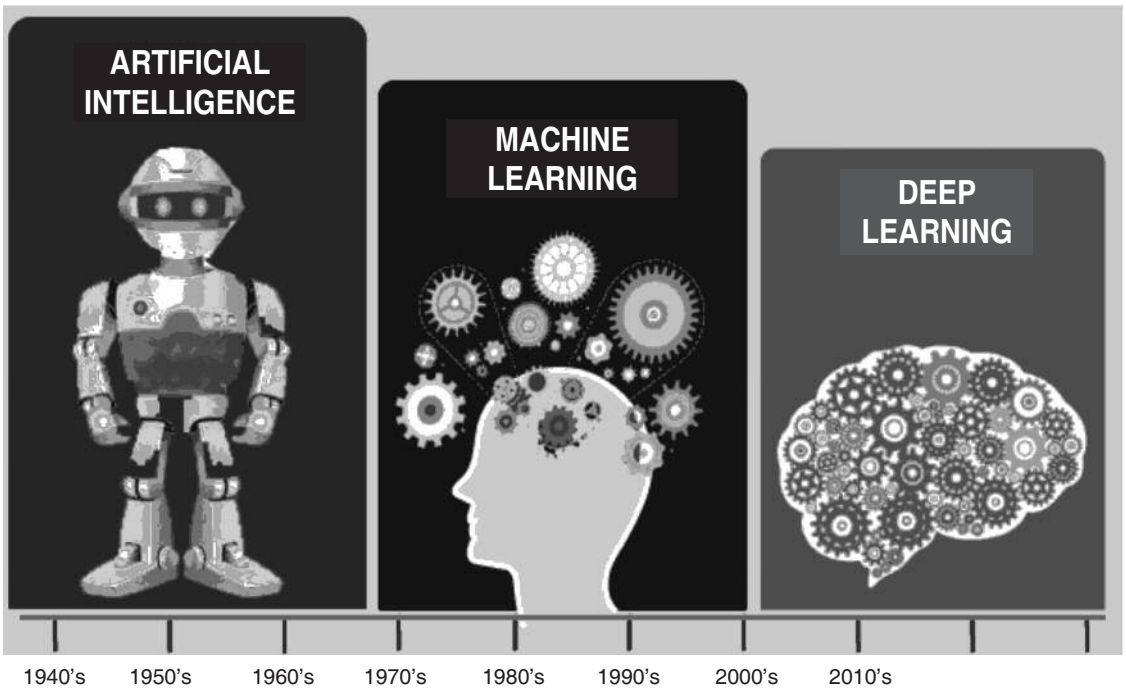

Fig. 1.8.

Now machine learning has got a great advancement in its research, and it is present everywhere around us, such as self-driving cars, Amazon Alexa, Catboats, recommender system, and many more. It includes Supervised, unsupervised, and reinforcement learning with clustering, classification, decision tree, SVM algorithms, etc.

7. Robotic process automation

Robotic process automation is an application of artificial intelligence which configures a robot (software application) to interpret, communicate and analyze data. This discipline of artificial intelligence helps to automate partially or fully manual operations that are repetitive and rule-based.

Robotics is **a branch of engineering that involves the conception, design, manufacture and operation of robots.** The objective of the

robotics field is to create intelligent machines that can assist humans in a variety of ways. Robotics can take on a number of forms.

A robot is the product of the robotics field, where programmable machines are built that can assist humans or mimic human actions. Robots were originally built to handle monotonous tasks (like building cars on an assembly line), but have since expanded well beyond their initial uses to perform tasks like fighting fires, cleaning homes and assisting with incredibly intricate surgeries. Each robot has a differing level of autonomy, ranging from human-controlled bots that carry out tasks that a human has full control over to fully-autonomous bots that perform tasks without any external influences.

8. Peer-to-peer network

The peer-to-peer network helps to connect between different systems and computers for data sharing without the data transmitting via server. Peer-to-peer networks have the ability to solve the most complex problems. This technology is used in crypto currencies. The implementation is cost-effective as individual workstations are connected and servers are not installed.

9. Deep learning

Deep learning another branch of artificial intelligence which functions based on artificial neural networks. This technique teaches computers and machines to learn by example just the way humans do. The term “deep” is coined because it has hidden layers in neural networks. Typically, a neural network has 2-3 hidden layers and can have a maximum of 150 hidden layers. Deep learning is effective on huge data to train a model and a graphic processing unit. The algorithms work in a hierarchy to automate predictive analytics. Deep learning has spread its wings in many domains like aerospace and military to detect objects from satellites, helps in improving worker safety by identifying risk incidents when a worker gets close to a machine, helps to detect cancer cells etc.

Deep learning utilizes both structured and unstructured data for training. Practical examples of deep learning are Virtual assistants, vision for driverless cars, money laundering, face recognition and many more.

Fig. 1.9.

Deep learning methods

Various methods can be used to create strong deep learning models. These techniques include learning rate decay, transfer learning, training from scratch and dropout.

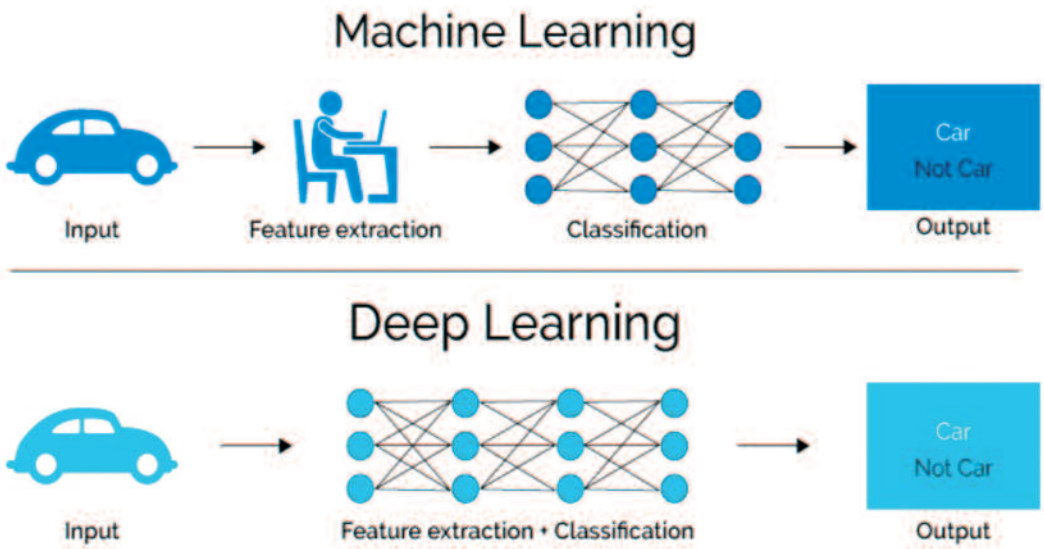

Fig. 1.10.

Learning rate decay. The learning rate is a hyperparameter—a factor that defines the system or set conditions for its operation prior to the learning process—that controls how much change the model experiences in response to the estimated error every time the model weights are altered. Learning rates that are too high may result in unstable training processes or the learning of a suboptimal set of weights. Learning rates that are too small may produce a lengthy training process that has the potential to get stuck.

The learning rate decay method—also called learning rate annealing or adaptive learning rates—is the process of adapting the learning rate to increase performance and reduce training time. The easiest and most common adaptations of learning rate during training include techniques to reduce the learning rate over time.

Transfer learning. This process involves perfecting a previously trained model; it requires an interface to the internals of a preexisting network. First, users feed the existing network new data containing previously unknown classifications. Once adjustments are made to the network, new tasks can be performed with more specific categorizing abilities. This method has the advantage of requiring much less data than others, thus reducing computation time to minutes or hours.

Training from scratch. This method requires a developer to collect a large labelled data set and configure a network architecture that can learn the features and model. This technique is especially useful for new applications, as well as applications with a large number of output categories. However, overall, it is a less common approach, as it requires inordinate amounts of data, causing training to take days or weeks.

Dropout. This method attempts to solve the problem of over fitting in networks with large amounts of parameters by randomly dropping units and their connections from the neural network during training. It has been proven that the dropout method can improve the performance of neural networks on supervised learning tasks in areas such as speech recognition, document classification and computational biology.

10. AL optimized hardware

Artificial intelligence software has a high demand in the business world. As the attention for the software increased, a need for the hardware that supports the software also arises. A conventional chip cannot support artificial intelligence models. A new generation of artificial intelligence chips is being developed for neural networks, deep learning and computer vision. The AL hardware includes CPUs to handle scalable workloads, special purpose built-in silicon for neural network, neuro morphic chips etc. Organizations like Nvidia, Qualcomm. AMD is creating chips which can perform complex AI calculations. Healthcare and automobile may be the industries which will benefit from these chips.

11.Text analytics

The act of converting unstructured text documents into useful, organized data is known as text analytics. Text analysis works by dissecting words and phrases into

their constituent parts and assessing the role and significance of each portion using complicated software rules and machine learning techniques.

Text analytics is the process of extracting meaning out of text. For example, this can be analyzing text written by customers in a customer survey, with the focus on finding common themes and trends.

Techniques of Text Analytics

(*a*) **Clustering :** One of the most important text mining approaches is clustering. Its goal is to find inherent patterns in textual data and arrange them into pertinent subgroups or 'clusters' for further study. The creation of meaningful clusters from unmarked textual material without any prior knowledge is a key problem in the clustering process. Cluster analysis is a common text mining method that aids in data distribution or serves as a pre-processing phase for other text mining techniques that operate on clusters that have been discovered.

(*b*) **Summarization.** The technique of automatically creating a compressed version of a given text that contains useful information for the end-user is known as text summarization. This text mining approach tries to search numerous text sources for abstracts of texts that provide a significant amount of information in a compact style while maintaining the original documents' general sense and intent. Text summarization combines and incorporates numerous text classification approaches such as decision trees, neural networks, logistic regression, and swarm optimization.

(*c*) **Information Extraction.** This is the most widely used text mining method. The process of extracting relevant information from large amounts of textual material is referred to as information exchange. The goal of this text mining approach is to extract entities, characteristics, and their connections from semi-structured or unstructured texts. Whatever data is collected is then saved in a database for easy access and retrieval in the long term. Precision and recall methods are used to examine and enhance the effectiveness and relevance of the outputs.

(*d*) **Categorization.** This is a type of "supervised" learning in which regular language texts are allocated to a pre-set of subjects based on their content using text mining techniques. As a result, categorization, or Natural Language Processing (NLP), is the act of gathering textual data and processing and analyzing them to find the appropriate themes or indexes for each one.

Application of Text Analytics

(*a*) **Fraud Detection.** Text analytics, alongwith text mining techniques, presents a huge potential for domains that collect the entirety of their data in text format. This is an opportunity that insurance and finance businesses are seizing. These businesses are now capable of processing claims quickly as well as preventing and detecting fraud by integrating the results of text analysis with pertinent structured data.

(*b*) **Social Media Analysis**. Many text mining methods have been developed specifically for assessing the performance of social media sites. These aid with the tracking and interpretation of online content created by the news, blogs, emails, and

other sources. Text mining technologies can also quickly evaluate the number of posts, likes, and connections your brand has on social media, helping you to better understand how people are reacting to your company and online content. The research will help you figure out "what's popular and what isn't" for your target market.

(*c*) **Customer Care Service**. In the realm of customer service, text mining applications, notably natural language processing (NLP), are becoming increasingly important. Companies are adopting text analytics tools to improve their entire customer experience by gaining access to textual data from a variety of sources, including surveys, user feedback, and user conversations, among others. Text analysis seeks to minimize the company's reaction time and assist in quickly and efficiently resolving client complaints.

(*d*) **Knowledge Management.** Managing a large volume of text data has become a difficulty in several areas, such as healthcare. It would certainly fly to the moon if you began designing systems and kept all of the papers relevant to healthcare on a single upwardly scalable rack. The amount of data collected per hour is enormous. All of this information must be kept in such a way that it may be accessed whenever it is needed. It's possible that an epidemic could break out, and hospitals will need to work together to analyze all of their data in order to locate the source or the first affected individual.

(*e*) **Risk Management**. The absence of appropriate or inadequate risk analysis is one of the leading reasons for company failure. Implementing and adopting risk management tools based on text mining technologies, like SAS Text Miner, may assist organizations in staying current with industry trends and enhancing their ability to mitigate possible hazards. Because text mining tools and technologies can aggregate relevant data from hundreds of text data sources and build linkages between the retrieved insights, they enable companies to access the appropriate information at the right time, therefore improving the risk management framework.

1.5. AI AND ETHICAL CONCERN

AI is not an end in itself, but rather a promising means to increase human flourishing, thereby enhancing individual and societal well-being and the common good, as well as bringing progress and innovation.

AI presents three major areas of ethical concern for society: privacy and surveillance, bias and discrimination, and perhaps the deepest, most difficult philosophical question of the era, the role of human judgement, said Sandel, who teaches a course in the moral, social, and political implications of new technologies.

When we speak of ethical issues of AI, there tends to be an implicit assumption that we are speaking of morally bad things. And, of course, most of the AI debate revolves around such morally problematic outcomes that need to be addressed. However, it is worth highlighting that AI promises numerous benefits. As noted earlier, many AI policy documents focus on the economic benefits of AI that are expected to arise from higher levels of efficiency and productivity. These are ethical values insofar as they

promise higher levels of wealth and wellbeing that will allow people to live better lives and can thus be conducive to or even necessary for human flourishing. It is worth pointing out that this implies certain levels of distribution of wealth and certain assumptions about the role of society and the state in redistributing wealth in ethically acceptable manners which should be made explicit.

The ethics of artificial intelligence is the branch of the *ethics of technology* specific to *artificially intelligent* systems. It is sometimes divided into a concern with the moral behaviour of *humans* as they design, make, use and treat artificially intelligent systems, and a concern with the behaviour of machines, in *machine ethics*. It also includes the issue of a possible singularity due to super intelligent AI .

Ethics of technology is a sub-field of *ethics* addressing the ethical questions specific to the *Technology Age*, the transitional shift in society where personal computers and subsequent devices have been introduced to provide users an easy and quick way to transfer information.

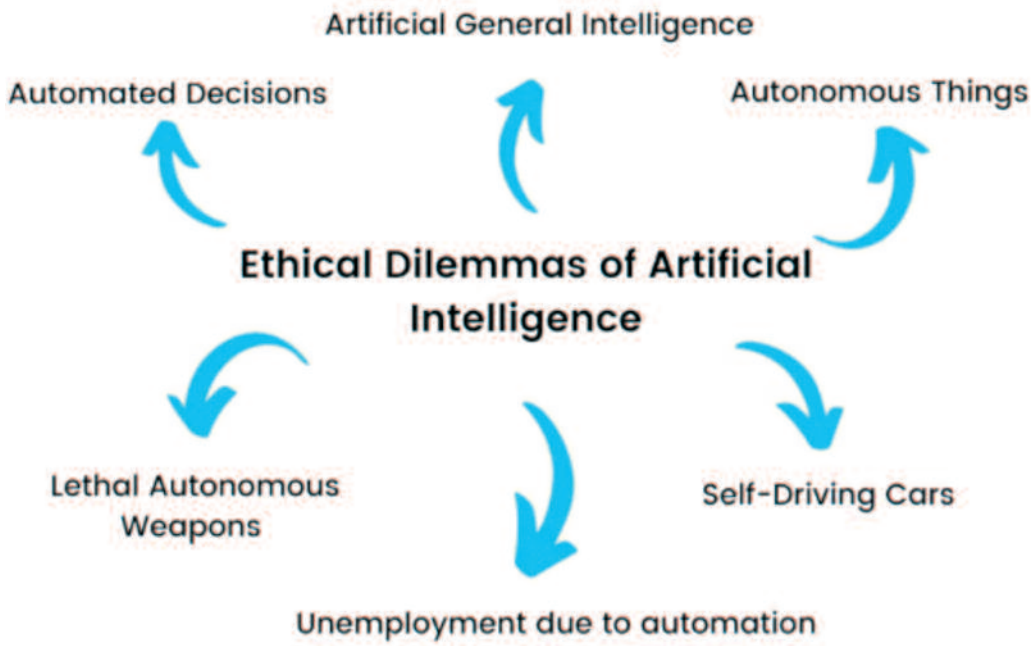

Fig. 1.11.

1. Unemployment. What happens after the end of jobs?

The hierarchy of labour is concerned primarily with automation. As we've invented ways to automate jobs, we could create room for people to assume more complex roles, moving from the physical work that dominated the pre-industrial globe to the cognitive labour that characterizes strategic and administrative work in our globalized society.

2. Inequality. How do we distribute the wealth created by machines?

Our economic system is based on compensation for contribution to the economy, often assessed using an hourly wage. The majority of companies are still dependent on

hourly work when it comes to products and services. But by using artificial intelligence, a company can drastically cut down on relying on the human workforce, and this means that revenues will go to fewer people. Consequently, individuals who have ownership in AI-driven companies will make all the money.

3. Humanity. How do machines affect our behaviour and interaction?

Artificially intelligent bots are becoming better and better at modelling human conversation and relationships. In 2015, a boot named *Eugene Goostman won the Turing Challenge* for the first time. In this challenge, human raters used text input to chat with an unknown entity, and then guessed whether they had been chatting with a human or a machine. Eugene Goostman fooled more than half of the human raters into thinking they had been talking to a human being.

4. Artificial stupidity. How can we guard against mistakes?

Intelligence comes from learning, whether you're human or machine. Systems usually have a training phase in which they "learn" to detect the right patterns and act according to their input. Once a system is fully trained, it can then go into test phase, where it is hit with more examples and we see how it performs.

Obviously, the training phase cannot cover all possible examples that a system may deal with in the real world. These systems *can be fooled* in ways that humans wouldn't be. For example, random dot patterns can lead a machine to "see" things that aren't there. If we rely on AI to bring us into a new world of labour, security and efficiency, we need to ensure that the machine performs as planned, and that people can't overpower it to use it for their own ends.

5. Racist robots. How do we eliminate AI bias?

Though artificial intelligence is capable of a speed and capacity of processing that's far beyond that of humans, it cannot always be trusted to be fair and neutral. Google and its parent company Alphabet are one of the leaders when it comes to artificial intelligence, as seen in Google's Photos service, where AI is used to identify people, objects and scenes. But it can go wrong, such as when a *camera missed the mark* on racial sensitivity, or when a *software used to predict future criminals* showed bias against black people.

We shouldn't forget that AI systems are created by humans, who can be biased and judgemental. Once again, if used right, or if used by those who strive for social progress, artificial intelligence can become a catalyst for positive change.

6. Security. How do we keep AI safe from adversaries?

The more powerful a technology becomes, the more can it be used for nefarious reasons as well as good. This applies not only to robots produced to replace human

soldiers, or autonomous weapons, but to AI systems that can cause damage if used maliciously. Because these fights won't be fought on the battleground only, cyber security will become even more important. After all, we're dealing with a system that is faster and more capable than us by orders of magnitude.

7. Evil genies. How do we protect against unintended consequences?

It's not just adversaries we have to worry about. What if artificial intelligence itself turned against us? This doesn't mean by turning "evil" in the way a human might, or the way AI disasters are depicted in Hollywood movies. Rather, we can imagine an advanced AI system as a "genie in a bottle" that can fulfill wishes, but with terrible unforeseen consequences.

8. Singularity. How do we stay in control of a complex intelligent system?

The reason humans are on top of the food chain is not down to sharp teeth or strong muscles. Human dominance is almost entirely due to our ingenuity and intelligence. We can get the better of bigger, faster, stronger animals because we can create and use tools to control them: both physical tools such as cages and weapons, and cognitive tools like training and conditioning.

9. Robot rights. How do we define the humane treatment of AI?

While neuroscientists are still working on unlocking the secrets of conscious experience, we understand more about the basic mechanisms of reward and aversion. We share these mechanisms with even simple animals. In a way, we are building similar mechanisms of reward and aversion in systems of artificial intelligence. For example, reinforcement learning is similar to training a dog: improved performance is reinforced with a virtual reward.

MULTIPLE CHOICE QUESTIONS

1. What is the full form of "AI"?

(*a*) Artificially Intelligent (*b*) Artificial Intelligence
(*c*) Artificially Intelligence (*d*) Advanced Intelligence

2. Who is the inventor of Artificial Intelligence?

(*a*) Geoffrey Hinton (*b*) Andrew Ng
(*c*) John McCarthy (*d*) Jürgen Schmidhuber

3. What is the goal of Artificial Intelligence?

(*a*) To solve artificial problems
(*b*) To extract scientific causes
(*c*) To explain various sorts of intelligence
(*d*) To solve real-world problems

4. In how many categories process of Artificial Intelligence is categorized?

(*a*) categorized into 5 categories
(*b*) processes are categorized based on the input provided
(*c*) categorized into 3 categories
(*d*) process is not categorized

5. Based on which of the following parameter Artificial Intelligence is categorized?

(*a*) Based on functionally only
(*b*) Based on capabilities only
(*c*) Based on capabilities and functionally
(*d*) It is not categorized

6. What is the function of an Artificial Intelligence "Agent"?

(*a*) Mapping of goal sequence to an action
(*b*) Work without the direct interference of the people
(*c*) Mapping of precept sequence to an action
(*d*) Mapping of environment sequence to an action

7. Which of the following is not a type of Artificial Intelligence agent?

(*a*) Learning AI agent (*b*) Goal-based AI agent
(*c*) Simple reflex AI agent (*d*) Unity-based AI agent

8. Which of the following is not the commonly used programming language for Artificial Intelligence?

(*a*) Perl (*b*) Java
(*c*) PROLOG (*d*) LISP

9. Which of the following machine requires input from the humans but can interpret the outputs themselves?

(*a*) Actuators (*b*) Sensor
(*c*) Agents (*d*) AI system

10. The total number of proposition symbols in AI are _____

(*a*) 3 proposition symbols (*b*) 1 proposition symbols
(*c*) 2 proposition symbols (*d*) No proposition symbols

ANSWERS

1. (*b*)	2. (*c*)	3. (*c*)	4. (*c*)	5. (*c*)
6. (*c*)	7. (*d*)	8. (*a*)	9. (*d*)	10. (*c*)

LONG TYPE QUESTIONS

1. What is Intelligence?
2. Describe the four categories under which AI is classified with examples.
3. Define Artificial Intelligence.
4. List the fields that form the basis for AI.
5. What are various approaches to AI.
6. What is emerging technologies? Give some examples.
7. What is the importance of ethical issue in AI?
8. Write the history of AI.
9. What are applications of AI?
10. What should all engineers know about AI?