Chapter 1
Introducing Artificial Intelligence

Abstract The field of Artificial Intelligence (AI) has interfaces with almost every discipline of engineering, and beyond, and all those can be benefited by the use of AI. This chapter presents the introduction to AI, its roots, sub-domains, Turing test to judge if the given program is intelligent, what are the goals of AI for Engineers and Scientists, what are the basic requirements for AI, symbol system, what are the requirements for knowledge representation, and concludes with a chapter summary, and an exhaustive list of exercises. In addition raises many questions to ponder over, like, consciousness, and whether it is possible to create artificial consciousness.

Keywords Artificial intelligence · Turing test · Imitation game · Philosophy · Logic · Computation · Cognitive science · Evolution · Artificial consciousness · Speech processing · Expert systems · Physical symbol system hypothesis · Formal system · Knowledge representation

1.1 Introduction

The Artificial Intelligence (AI) is a branch of Computer Science, which is mainly concerned with *automation of Intelligent behavior*. This behavior we may consider from all domains—the human, animal world, and vegetation. A compact definition of Intelligence is:

$$Intelligence = Perceive + Analyze + React.$$

The following are often quoted definitions, all expressing this notion of intelligence but with different emphasis in each case:

- "The capacity to learn or to profit by experience."
- "Ability to adapt oneself adequately to relatively new situations in life."
- "A person possesses intelligence insofar as he has learned, or can learn, to adjust himself to his environment."
- "The ability of an organism to solve new problems."

- "A global concept that involves an individual's ability to act purposefully, think rationally, and deal effectively with the environment."
- "Intelligence is a very general mental capability that, among other things, involves the ability to reason, plan, solve problems, think abstractly, comprehend complex ideas, learn quickly and learn from experience."

The foundation materials of AI comprises—data structures, knowledge representation techniques, algorithms to apply the knowledge and language, and the programming techniques to implement all these.

To get an idea of *Intelligence* it requires answering these and many similar questions:

- Is intelligence due to a *single faculty* or it is a name for a collection of distinct unrelated faculties?
- Is it a priori existence or it can be learned? What does exactly happen when we learn some thing, that is, in terms of information and storage structures.
- What is truly the process of *creativity* and *intuition* in human? These again, in terms of knowledge and its structures.
- Does the intelligence require an internal mechanism, or it can be concluded from the behavior observed?
- What is the mechanism for representing the knowledge in the living cells?
- Are the machines self aware like humans? What are the basic requirements for creating the facility of self-awareness in machines?
- Is it that computer intelligence can be defined only when we know the intelligence in reference to human beings?
- Would it ever be possible to achieve the intelligence in computers? Or, is it true that achievement of intelligence is possible only when or is it that an intelligent entity requires the richness of *sensation* and *experience* which might be found only in a biological existence?

Partly, the aim of Artificial Intelligence (AI) is to find the answer to these questions, through the tools provided by AI. These tools are because, the AI offers the medium, as well as the test-bed for theories of intelligence, which can be expressed in the form of computer programs, and can be tested as well as verified by running these programs on computers.

Unlike Physics and Chemistry, AI is still a premature field, hence, its structure, objectives, and procedures are less clearly defined, and not clear like those in physics and chemistry. The AI has been more concerned about expanding limits of computers, apart from defining itself.

Learning Outcomes of this Chapter[1]:

1. Defining AI. [Familiarity]
2. Describe Turing test thought experiment. [Familiarity]

---

[1] https://www.acm.org/binaries/content/assets/education/cs2013_web_final.pdf (pp. 121–129).

3. Differentiate between the concepts of optimal reasoning/behavior and humanlike reasoning/behavior. [Familiarity]
4. Sub-fields of AI. [Familiarity]
5. Determine the characteristics of a given problem that an intelligent system (i.e., AI-based system) must solve. [Assessment]
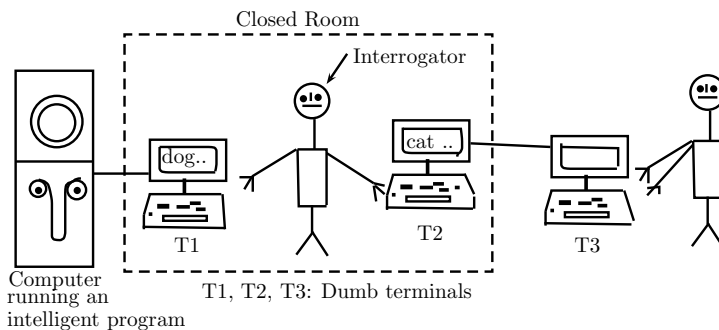
1.2 The Turing Test

In 1950, in an article "Computing Machinery and Intelligence," *Alan M. Turing* proposed an *empirical test* for machine intelligence, now called *Turing Test* (see Fig. 1.1). It is designed to measure the performance of an intelligent machine against humans, for its intelligent behavior. Turing called it *imitation game*, where machine and human counter-part are put in different rooms, separate from a third person, called *interrogator*. The interrogator is not able to see or speak directly to any of the other two, and does not know which entity is a machine, and communicates to these two solely by textual devices like a dumb terminal [11, 12].

The interrogator is supposed to distinguish the machine from the human solely based on the answers received for the questions asked over the interface device, which is a keyboard (or teletype). Even after having asked the number of questions, if the interrogator is not able to distinguish the machine from the human, then as per the argument of Alan Turing, the machine can be considered intelligent. Interrogator may ask highly computation oriented questions to identify the machine, and other questions related to general awareness, poetry, etc., to identify the human [6].

The game (with the "player machine" omitted) is often in practice under the name of viva-voce to discover whether someone really understands something or has "learned it parrot fashion".

Many researchers argue that the Turing test is not sufficient to establish the presence of intelligence. Some of the arguments *for* and *against* the above test can be as follows:

Fig. 1.1 Turing test (imitation game)

1. It takes human being as a reference for intelligent behavior, rather than debating over the true nature of intelligence: *against*.
2. The unmeasurable things are not considered, e.g., whether a computer uses *internal structures*, or for example, whether the machine is *conscious* of its actions, which are currently not answerable: *against*.
3. Eliminates any bias to human oriented interaction mechanisms, since a computer terminal is used as a communication device: *for*.
4. Biased towards only symbolic problem solving: *against*.
5. Perceptual skills or dexterity cannot be checked: *against*.
6. Unnecessarily constrains the machine intelligence to human intelligence: *against*.

Though, the number of counts of *against* are far more than *for*, there is yet no known test which is considered better than the Turing Test.

In the part success of the Turing Test, a powerful computer has deceived humans in the thinking process, where this machine modeled intelligence of a young boy, to become the first machine to pass the Turing test, conducted in June 2014. In this experiment, five machines were tested at the Royal Society in central London to check if these machines could fool the people into thinking. The machines behaved like humans, and the conversation was in the form of text. A computer program, by the name "Eugene Goostman" was developed to simulate a young boy, which came out to convince the one-third of the judges that it was human [5].

1.3 Goals of AI

AI is the area of computer science aiming for the design of intelligent computer systems, i.e., systems that have characteristics of intelligence, like, we observe in human behavior, for example, to understand language(s), and have abilities of learning, reasoning, and problem solving [9].

For many researchers, the goal of AI is to emulate human cognition, while to some researchers, it is the creation of intelligence without considering any human characteristics. To many other researchers, AI is aimed to create useful artifacts for the comforts and needs of human, without any criteria of an abstract notion of intelligence.

The above variation in aims is not necessarily a wrong, as each approach uncovers new ideas and provides a base for pursuing the research in AI. However, there is a convincing argument that due to the absence of a proper definition of AI, it is difficult to establish as what can and what cannot be done through AI.

One of the goals for studying AI is to create intelligence in machines as a general property—not necessarily based on any attribute of humans. When this is the goal, it also includes the objective of creation of artifacts of human comforts and needs, which can be the driving force of technological development. However, this goal also requires a notion of intelligence, to start with.

Further, the problem gets compounded, since artifacts manufacturer may say their product is better in terms of saving labor and money, while cognitive scientists may say that their system correctly predicted human behavior. Apart from this dispute, without proper theoretical base of AI, it is not a wise idea to build complex system—on which one has confidence, can be analyzed for performance and error analysis can be carried out.

The definition of AI be such that for any system, it covers input, output, and their relationship based on the structure of the system. There is a need of such definition to be as general as possible so that it can be uniformly applicable. In the absence of such a definition, one takes the AI as that exists in Chess playing, the one in automated vehicle driving, and the one existing in medical expert system for diagnosis—all these approaches for the definition of AI are varying from case to case.

The *scientific goal* of AI is to determine theories about knowledge representation, learning, rule-based systems, and search that explain various sorts of intelligence.

The *engineering goal* of AI is to acquire ability in the machine so that it can solve the real-life problems. The basic techniques used by AI for this purpose are knowledge representation, machine learning, rule systems, and state space search.

In the past, computer scientists and engineers used to be more concerned with the engineering goals, but the psychologists, philosophers, and cognitive scientists were more keen on the scientific goals. In spite of these opposite concerns, there are common techniques that the two approaches can feed to each other. Hence, we will proceed with both the goals in mind.

1.4   Roots of AI

The field of AI does not live in isolation, and has significant roots in number of older disciplines, particularly,

Philosophy,
Logic/Mathematics,
Computing,
Psychology/Cognitive Science,
Biology/Neuroscience, and
Evolution.

In the above domains, there is significant overlap, for example, between philosophy and logic, and between mathematics and computation. By looking at each of these in turn, we get a better understanding about their role in AI, and how these fields have developed to play that role in AI [6, 10].

1.4.1   Philosophy

The evidence of philosophy goes as back as Socrates times (∼400 BC) where he asks for an algorithm to distinguish between *piety* (a reverence of supreme being) from *non-piety*. In around 300 BC, *Aristotle* formulated various types of *deductive* reasoning approaches, to mechanically generate conclusions using initial premises. One approach of deductive reasoning he used was *modus ponens*, now also a standard technique of inference in predicate and propositional logic. It is stated as

> If "A is True" → "B is True", and "A holds True" then conclude that "B holds True".

As an example of inference, "If it's raining then you get wet, and it's raining, then you should have got wet".

The philosopher *Rene Descartes* (1596–1650) introduced the concept of mind-body dualism, which says that part of the mind is exempted from following the physical laws. The conclusion of which he has drawn as *free will*. In the present time also, when the AI, machine learning, and data science are dominantly used, but it is argued that machine cannot supersede human in intelligence, as they do not have free will, and they need to be assigned the goal by a human.

Gottfried Wilhelm Leibnitz, a German philosopher and mathematician (1646–1716), who supported the materialist nature of mind, said that mind operates by ordinary physical processes. In the present context, this means that mental processes can be performed by the machines.

1.4.2   Logic and Mathematics

The logic, has history of development from the time of Greece philosophers—Plato and Aristotle (∼300–400 BC), however, there has been more recent development at a rapid pace. These were due to the following

- *Earl Stanhope's* Logic (1777), using which Earl demonstrated a machine capable of solving the problem using the inference rule, called *syllogisms*, and the numerical problems of logical nature, and elementary questions based on the theory of probability.
- *George Boole* (1815–1864) introduced the language-based formal logic for a drawing logical inference in 1847, later became popular by name *Boolean Algebra*.
- *Gottlob Frege* (1848–1925) introduced the *first-order logic* that today forms the most common knowledge representation system, called as FOPL (first-order predicate logic).
- *Kurt* Gödel (1906–1978), in 1931, demonstrated that there are limits of logic. Through his *incompleteness theorem* he showed that in any formal logic, which is powerful enough to describe the properties of natural numbers, there exist true statements, but their truth cannot be proved using any algorithm.

- *Roger Penrose* in 1995 tried to prove that human mind has non-computable capabilities.

1.4.3 Computation

In the nineteenth and twentieth-century, many scientists defined the formalism of what is *computation*, basic theory of it, and that there are things that are not computable irrespective of whatever are the computing resources and time provided.

In the year 1869, *William Jevon* constructed a *Logic Machine* capable of handling Boolean Algebra and Venn Diagrams, and could solve logical problems faster than human beings.

Alan M. Turing (1912–1954) tried to characterize exactly what are the functions that can be computed. He used, what is now called as Turing Machine. Unfortunately, it is difficult to give the notion of computation as a formal definition, however, the *Church-Turing thesis*, due to Alonzo Church and Turing, states that a Turing machine is capable of computing any computable function, which is now, accepted as a sufficient definition of computability. Turing also showed that there are some functions which no Turing machine can compute (e.g., *Halting Problem*)—these are non-computable functions.

John von Neumann (1903–1957) gave, now what is called as, von Neumann architecture—a description of a logical model of computation and computer, without any physical realization of a computer.

In 1960s, two important concepts emerged—*Intractability* (the solution time of a problem grows at least exponentially) and *Reduction* of complex problems into simpler problems.

1.4.4 Psychology and Cognitive Science

Cognitive Psychology or Cognitive Science is the study about the functioning of the mind, human behavior, and the processing of information about the human brain. An important consequence of human intelligence is human languages. The early work on knowledge representation in AI was about human language, and was produced through research in linguistics.

It is humans' quest to understand as to how our and other animals' brains lead to intelligent behavior, with the aim to ultimately build AI systems. On the other hand, it is also aimed to explore the properties of artificial systems, like, computer models/simulations to test our hypotheses concerning human systems.

Many people working in sub-fields of AI are in the process of building models of how the human system operates, and use artificial systems for solving real-world problems.

1.4.5   Biology and Neuroscience

The field of neuroscience says that human brains, which provide intelligence, are made up of tens of billions of neurons, and each neuron is connected to hundreds or thousands of other neurons. A neuron is an elementary processing unit, performing a function called *firing*, depending on the total amount of activity feeding into it. When a large number of neurons are connected together, it gives rise to a very powerful computational device that can compute, as well as learn how to compute.

The concept of human brains, having the capability to compute, as well as learn, is used to build artificial neurons in the form of electronics circuits, and connect them as circuits (called ANN—artificial neural networks) in large quantities, to build powerful AI systems. In addition, the ANN are used to model various human abilities.

The major difference between the functions of neurons and the process of human reasoning is that neurons work at *sub-symbolic level*, whereas much of conscious human reasoning appears to operate at a symbolic level, for example, we do most of the reasoning in the form of thoughts, which are manipulations of sentences.

The collection of neurons in the form of programs called Artificial Neural Networks (ANN), perform well in executing simple tasks, and provide good models of many human abilities. However, there are many tasks of AI that they are not so good at ANN, and other approaches are more promising in those areas, compared to ANN. For example, for natural language processing (NLP) and reasoning, the symbolic logic, called predicate logic is better suited.

1.4.6   Evolution

Unlike the machines, the humans (intelligence) has a very long history of evolution, of millions of years, compared to less than hundred years for electronic machines and computers. The first exhaustive document of human evolution, the evolution by *natural selection* is due to Charles Darwin (1809–1882). The idea is that fitter individuals will naturally tend to live longer and produce more children (may not be truly valid in the modern world). Hence, after many generations, a population will automatically emerge with good innate properties [3].

Due to this evolution, the structure of the human brain, and even the knowledge, are to a sufficient extent built-in at the time of the birth. This is an advantage over ANNs, which have no pre-stored knowledge, hence they need to acquire the entire knowledge by learning only. However, the present-day computers are powerful enough that even the evolution can be simulated using them, and can evolve the AI systems. It has now become possible to evolve the neural networks to some extent so that they are efficient at learning. But, may still be challenging to recreate the long history of the evolution of humans in the ANNs.

A closely related field to ANNs is *genetic programming*, which is concerned with writing the programs that evolve over time, and do not need to modify them as it is done in usual programs when the system requirement changes [4].

1.5 Artificial Consciousness

Right from the time, automated machines like computers came into existence, it was the quest of researchers to build machines that can compete in intelligence with humans. Looking at the current progress rate of smart machines, like smartphones, it is believed that in the not very far future, it may be possible to build machines which may be, if not more, but have comparable intelligence to that of humans. Using such machines, it may be possible to produce human like consciousness in machines—called as artificial consciousness.

On the contrary, even a far of realization of artificial consciousness gives rise to several philosophical nature of questions:

- Can the computers be made to think or they will just calculate?
- Is consciousness a human prerogative only, or it can be created in machines also?
- Is the consciousness due to the material comprised in the human brain, or it can be created in silicon also (the computer hardware)?

To provide the answers to these questions is difficult as of now, mainly because it requires combining the knowledge from the fields of computer science, neurophysiology, and philosophy.

On the other hand, the very talk of artificial consciousness—a possible product of the human imagination, express human desires, and fears about future technologies—may influence the course of progress.

At a social level, the science fiction stories simulate future scenarios that can help prepare us for crucial transitions by predicting the consequences of such technological advances [1].

1.6 Techniques Used in AI

The AI systems have a lot of variations, for example, the rule-based systems are based on symbolic representations, and work on inferences. There are other extremes, the ANN-based system work on the interface with other neurons, and connection weights. In spite of all these, there are four common features among all of them.

Representation

All AI systems have an important feature of knowledge representation. The rule-based systems, frame-based systems, and semantic networks make use of a sequence of *if-then rules*, while the artificial neural networks make use of connections along with connection weights.

Learning

All AI systems have capability of learning, using which they automatically build up the knowledge from the environment, e.g., acquiring the rules for a rule-based expert system, or determining the appropriate connection weights in an artificial neural network.

Rules

The rules of an AI-based system can be implicit or explicit. When explicit, the rules are created by a knowledge engineer, say, for an expert system, and when implicit, they can be, for example, in the form of connection weights in a neural network.

Search

The search can be in many forms, for example, searching the sequence of states that lead to solution faster, or searching for an optimum set of connection weights in an ANN by minimizing the fitness function.

1.7   Sub-fields of AI

Considering AI as replication of human intelligence may be misleading, and primitive. The later is because the true process of human intelligence and its sources are still under debate. However, if AI is taken to mean the advanced computing, it means more justified. In the past two decades, particularly after the year 2k, the AI applications have evolved and expanded, in the commercial, industrial, medicines and drug decide, medical science, consumer products, manufacturing processes, and even in management, to list only a few of its total domain. The use of AI techniques in every organization has become necessary to maintain competitiveness in the market. Many organizations keep secret of the true AI techniques they use.

AI now consists of many sub-fields, using a variety of techniques, such as the following:

Speech Processing: To understand speech, speech generation, machine dialog, machine user-interface.
Natural Language Processing: Information retrieval, Machine translation, Question/Answering, summarization.
Planning: Scheduling, game playing.
Engineering and Expert Systems: Troubleshooting medical diagnosis, Decision support systems, teaching systems.
*Fuzzy Systems*: For fuzzy controls.
Models of Brain and Evolutionary: Genetic algorithms, genetic programming, Brain modeling, time series prediction, classification.
Machine Vision and Robotics: Object recognition, image understanding, Intelligent control, autonomous exploration.
Machine Learning: Decision tree learning, version space learning.

Most of these have both engineering and scientific aspects. Many of these are going to be discussed in this text. Following is brief Introduction to some of these areas.
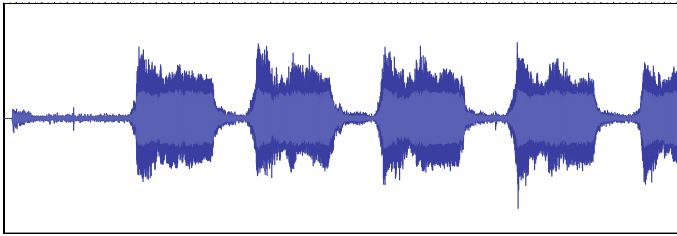
Fig. 1.2 Sound waves for the text "Hello" repeated five times by a five year child

1.7.1 Speech Processing

The speech processing and understanding of human speech has number of applications, some of them we come across quite often, like, speech recognition for dictation systems, speech production for automated announcements, voice-activated control, human–computer interface (HCI), and voice-activated transactions are a few examples.

One of the primary goals is, how do we get from sound waves to text streams and vice-versa? The Fig. 1.2 is an example, showing the sound wave pattern for the text "Hello" repeated five times.

To be precise, how should we go about segmenting the stream into words? How can we distinguish between "Recognize speech" and "Wreck a nice beach"?

1.7.2 Natural Language Processing

Consider the machine understanding and translation of simple sentences given below.

Ram saw the boy in the park with a telescope.
Ram saw the boy in the park with a dog.

In the parse-tree in Fig. 1.3a, the sentence structure is "Ram saw, the boy in the park, with a Telescope." Whereas in Fig. 1.3b, it is "Ram saw, the boy in the park with a dog." In first, it shows the association of verb *saw* and *telescope*, i.e., someone is seeing using telescope. In Fig. 1.3b the association of *boy*, and *dog* is shown, and all are in the park. The further deeper contexts help in resolving this ambiguity.

Though the sentences appear simple, finding out the meaning of each using the machine is difficult, as the parse-tree of each need to be analyzed for the meaning associated in each, in addition, the context knowledge is important.

Fig. 1.3 Parse-trees with different semantics
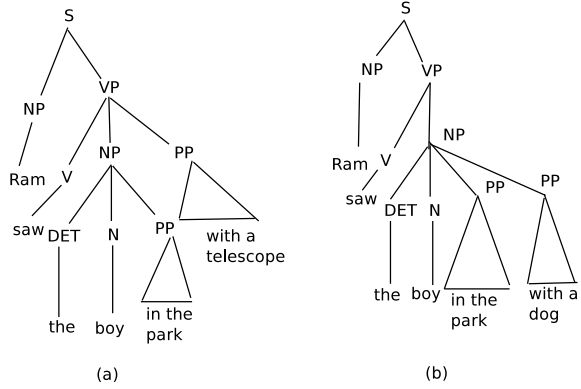
Some of the common applications of Natural Language Processing (NLP) are [2]:

1. Word Processing and Desktop Publishing
2. Spell check and Correction
3. Information Retrieval
4. Information Extraction
5. Information Categorization
6. Question-answering
7. Information summarization
8. Machine Translation.

1.7.3 Planning

Any planning, or specifically the robotic planning, is concerned with choosing (computing) the correct sequence of actions to complete a specific task. This requires a convenient and efficient representation of the problem domain. The plan steps are called states defined in formal language, like, in predicate language, or in the form of rules, depending on what type of planning has been used. A plan may be taken as a sequence of operations that transform the initial state into the goal state ultimately, the later is the solution. The best planning seeks to explore the best path of states for reaching the goal state. Hence, the best path or an optimum path seeks exploration or searching, to find out the best possible path efficiently, in terms of time and space required for running the planning algorithm.

1.7.4 Engineering and Expert Systems

These are primarily based on symbolic processing—which is the main stream of AI. The fundamental problems such as, how to represent knowledge, is the con-

tent of the following chapters. Various representation schemes are classified into two categories: *Symbolic* representation and *Graphical* representation. The first is *propositional* and *predicate logic* based, while the other is graph-based, like—*semantic networks, frames, ontologies* and *conceptual dependencies*.

1.7.5 Fuzzy Systems

The primary aim of fuzzy or *soft-computing* is to exploit the tolerance for *imprecision* and *uncertainty* to achieve tractability, robustness, and low cost applications. Fuzzy systems can be integrated with other techniques such as neural networks, and *probabilistic reasoning*. In fuzzy systems, the membership are partial (fuzzy/unclear). Examples of fuzzy sets are: old, cloudy, high speed, rainy, etc. This is in contrast to the classical-based systems, where boundaries are *crisp*, like, member of computer science class, an Indian national, chair (member of the set of chair), where all three are representing the full membership of each category of the given set.

1.7.6 Models of Brain and Evolution

The models of the human brain and evolution correspond to two major approaches to AI. The AI as a model of the brain corresponds to *Symbolic AI*, and has remained the major field of AI. This has the property of high level of mathematical abstractions, and considered as the macroscopic view of AI. The human psychology operates at a symbolic level, as well as the AI programming languages, and early engineered systems fall in this type of AI.

The other approach of AI is like human evolution, which is based on low level biological and genetic models of living beings. The neural-based computing (artificial neural networks), and genetic algorithms derived from these concepts from life. These biological models of AI may not necessarily be resembling verbatim to their living being counterparts, but the AI techniques that are based on GA (genetic algorithms) evolve the solution like the populations of human or like other forms of the life evolves [7].

Neural networks, similar to expert systems, are modeled on the human brain and learn by themselves from patterns. Thus, learning can then be applied to classification, prediction, or control applications. The GAs, are computer models based on genetic and evolution. Their basic idea is that genetic programs work towards finding better and better solutions to problems, just as species evolve to better adapt to their environments. The GAs comprises three basic processes: *reproduction* of solution based on their fitness, *crossover* of genes, and *mutation* for random change of genes. A broader definition GA is *evolutionary computing*, which includes not only GAs but classifier systems, genetic programming in which each solution is a computer program, and a part of *artificial life*.

1.8 Perception, Understanding, and Action

These fields are concerned with vision, speech processing, and robotics. The basic theme is applications that make machine sense (e.g., to see, hear, or touch), then understand and think, and finally take action.

For example, the basic objective of machine vision may be to make the machine "understand" the input consisting of reflected brightness values. Once this understanding is achieved, the results can be used for interpretation of patterns, inspecting parts, action of robots, and so forth. Developing an understanding presents the same difficulty in all areas of AI, including knowledge based systems—people understand what they see by integrating an optical image with complex background knowledge. Such background knowledge has been built over years of experiencing perceptions. Creating this type of information processing in the machine is a challenging task; however, some interesting applications have already appeared as the evidence to support future progress.

The speech processing uses two major technologies. One of these areas focus on input or speech recognition where acoustic input, like optical input in machine vision, is difficult task to automate. People understand what they hear with complex background knowledge. Speech recognition technology includes signal detection, pattern recognition, and possibly semantics—a feature closely related to Natural Language understanding. The other technology concerns to the creation of output or text-to-speech (tts) synthesis. Speech synthesis is easier than recognition, and its commercialization has been well established.

The field of Robotics integrate many techniques of sensing, and, is one of the AI areas in which industrial applications have the longest and widest successful records. The abilities of these robots are relatively limited. For example, only in a limited task such as welding seams and installing windshields, etc [7].

1.9 Physical Symbol System Hypothesis

The symbols are the basic requirements of intelligent activity, e.g., by human the symbols are basic number systems, alphabet of our languages, sign language, etc. Similar is the case with entire computer science, the languages, commands, computations, all have symbols as the base. When the information is processed by computers, on the completion of the task, we measure the progress, as well as the quality of results and efficiency of computations, only based on its symbols' contents in the end results [8].

1.9.1   Formal System

Basic requirement of achieving AI is *formal system*, which is based on *physical symbol system hypothesis*. The term was coined by *Allen Newell and Herbert Simon*, which states that a "physical symbol system" is a necessary requirement for AI to function. As per this, the physical patterns, called symbols, are combined to produce structures (i.e., expressions), and the processes act on these to manipulate to produce new expressions [8].

This symbol system hypothesis claims that human intelligence is due to this symbol system, which comprises all the alphabets, numerals, and other punctuation symbols. Thus, the symbol system is a "necessary" requirement for achieving the intelligence. Based on this argument, one can say that, if the machines are provided with symbol system with symbol manipulation capabilities, it is "sufficient" for achieving the intelligence in the machines.

As per the Physical Symbol System Hypothesis (PSSH), the capabilities of symbol manipulation are the essence for human's, as well as machines' intelligence. Hence, it is a necessary and sufficient tool for achieving intelligence in both the machines and humans. There is also experimental evidence, that, in various problem solving, like, in mathematical puzzles, planning of activities, and execution, the symbol system is the key requirement. By the term "necessary" here means, that the system possessing general intelligence, on analysis, will prove to be based on a physical symbol system. And, the term "sufficient" means that the physical symbol system can be organized to be exhibiting the general intelligence. When the problem-solving process of humans were simulated step by step on computers by the researchers, it was found to be simply the process of symbols' manipulations.

Of course, various researchers have criticized this hypothesis strongly, but still, it forms the central part of AI research. The critics argue that the symbol systems work only for high level processes like chess, games, and puzzles, but not suitable for low level systems like vision and speech recognition. This distinction is based on the fact that high-level symbols directly correspond to objects, like ⟨*cat*⟩, ⟨*house*⟩, ⟨*hill*⟩, etc, but not to the low-level symbols that are present in the machinelike neural networks (or ANN).

1.9.2   Symbols and Physical Symbol Systems

If we look at the entire knowledge of computer science, it is the symbols, which have been used to explain this knowledge at the most fundamental level. The explanation is nothing but the scientific proposition of nature, which is empirically derived over a long period of time, through a graduate development. Hence, the symbols are at the root of artificial intelligence, and are also the primary topic of artificial intelligence.

For all the information processed by computers in the service of finding the end goals, the intelligence of the system (computers) is their ability to reach goals in

the face of difficulties and complexity of the solution, as well as the complexity introduced by the environment. The fundamental requirements to achieve artificial intelligence is to store and manipulate the symbols, however, there is no uniformity and specific requirement of storage structures, as the structures vary from method to method used for implementation of AI, which are mostly the variants of network-based representation and predicate-based representations.

The "physical systems" used have two important characteristics: 1. Operation of systems is governed by the laws of physics, when they are realized as engineered systems, and made of engineered components; and 2. The "symbols" are not limited only to the symbols used by human beings.

1.9.3 Formal Logic

The "physical symbol system" hypothesis has its root to Russel's *formalizing logic*, which states that one need to capture the basic conceptual notion of mathematics in logic and put that notion to proof and deduction as sound base. This notion, with the effort ultimately grew in the form of mathematical logic—the propositional logic, predicate logic, and their variants [13].

1.9.4 The Stored Program Concept

The second generation of computers brought the concept of stored program concept in the mid-forties, after the *Eniac* computer. The arrival of these computers was considered as a milestone in terms of conceptual progress, as well as practical availability of systems. In such systems, the programs are treated as data, which, in turn, was processed by other programs, like a compiler program processing another program as data to generate object code. Interestingly, this capability was already verified and existed in Turing machine, which came as early as 1936. The Turing machine is a model of computing given by Alan M. Turing, where, in a universal Turing machine, an algorithm (another Turing machine) and data are on the very same tape. This idea was realized practically when machines were built with enough memory to make it practicable to store actual programs in some internal place, along with the data on which the program will act, as well as the data which will be produced as a result of the execution of the programs.

1.10 Considerations for Knowledge Representation

As far as AI is concerned, the following are the aspects of knowledge representation:

- What is the meaning of Knowledge?
- How the Knowledge can be represented in the machine?
- What are the requirements of representation of knowledge, e.g., structures, methods, size, etc.
- How the practical and theoretical aspects differ for knowledge representation?
- Can it be? Or, if yes, how to represent the knowledge using Natural Language?
- Can we call the databases as a form of knowledge representation?
- What are the semantic networks, and what are the frames? How the knowledge can be represented using these approaches?
- How the knowledge can be represented using the First-Order Predicate Logic (FOPL)?
- What is a Rule-Based Systems?
- What is an expert system?
- Out of the many techniques, which is the best technique for knowledge representation?

1.10.1 Defining the Knowledge

As per the Webster English language dictionary, the following are the meanings of knowledge:

1. The act or state of knowing; clear perception of fact, truth, or duty; certain apprehension; familiar cognizance; cognition. [1913 Webster]
   Knowledge, which is the highest degree of the speculative faculties, consists in the perception of the truth of affirmative or negative propositions—Locke. [1913 Webster]
2. That which is or may be known; the object of an act of knowing; a cognition—Chiefly used in the plural. [1913 Webster]
3. That which is gained and preserved by knowing; instruction; acquaintance; enlightenment; learning; scholarship; erudition. [1913 Webster]

1.10.2 Objective of Knowledge Representation

The objective of knowledge representation is to express the knowledge in computer so that the AI programs can use it to perform reasoning and inferences using this in an efficient way. The knowledge is represented using certain representation language,

for example, a predicate like language. The language has two important components in it.

Syntax

The system of a language defines the methods using which we or the machine can distinguish the correct structures from incorrect, i.e., it makes possible to identify the structurally valid sentences.

Semantics

The semantics of a language defines the world, or facts in the world of the concerned domain. And, hence defines the meaning of the sentence in reference, to the world.

1.10.3   Requirements of a Knowledge Representation

A good knowledge representation system for any particular domain should possess the following properties.

Adequacy of representation

The representation system should be able to represent all kind of knowledge needed in the concerned AI-based system.

Adequacy of Inference

The representation should be such that all that can be inferable by manipulating the given knowledge structures should be inferred by the system, when needed.

Inference Efficiency

The knowledge structures in the representation are so organized that the attention of the system, in the form of deductions, navigates in such a direction that it can reach the goal quickly.

Efficient acquisition

It should be able to acquire the new information automatically and efficiently, as and when needed, and also to update the knowledge regularly. In addition, there should be a provision that knowledge engineer can update the information in the system.

1.10.4   Practical Aspects of Representations

We are aware of good and bad knowledge representation, when we consider the knowledge representation in English or any other natural language. These are due to factors, like, syntax, semantics, partial versus full knowledge on any subject, depth and breadth of knowledge, etc.

There are many theoretical requirements for good knowledge representations, which can be met by dealing with a number of practical aspects, as follows:

- The representations should be *complete*, so that everything needs to be represented, can easily be represented.
- The representations should be simple and clear, so that one can easily understand what is being communicated by the representation,
- The important objects and their relations should be explicit and accessible, so that it becomes easy to see what is going on, and how the components of knowledge interact with each other.
- The irrelevant detail of the knowledge should be suppressed in the representation, so that they do not introduce complications. However, when needed, these are still available.
- The representation should be concise, so that information can be stored, retrieved and manipulated rapidly.
- The representation should be such that the overall system is fast.
- They must be *computable* and implementable with standard computing procedures.

To realize the above, a lot depends on algorithms used, representation structures, hardware, as well as dissemination of knowledge before it is represented.

1.10.5 Components of a Representation

To carry out the analysis of any representation system, it is useful to break the entire representation into smaller (smallest) components, which are in the most fundamental form. Accordingly, the components of an AI representation are divided into the four fundamental components:

Lexical components

The lexical components of knowledge representation are the symbols and words of the vocabulary used for representation.

Syntactic/Structural components

It describes how the symbols can be arranged systematically to create meaningful sentences. These structures are the grammar of the language used for representation.

Semantic components

They help is associating real-world meaning to objects and entities.

Procedural components

These procedures are used for creating and modifying the representations, and also for answering the questions using these procedures.

1.11   Knowledge Representation Using Natural Language

We humans are intelligent beings, who make use of the knowledge represented in the form of natural language (like, English, Hindi, Chinese, etc.), we update that knowledge (i.e., acquisition), and do the reasoning and inferences using this representation. Of course, there are many other types of knowledge representation and inferencing with the humans, which are not symbolic-based, like those acquired through smell, touch, hearing, and taste (through tongue). Hence, why not use the natural language for knowledge representation for machines also? The following are the trade-offs for representation using natural language.

Advantages

There are very strong advantages in favor of using natural language for knowledge representation.

- The natural language is strong at expressiveness, using which we can represent almost everything (real-world situations, pictures, symbols, ideas, emotions) and can carry out the reasoning using that.
- It is the most abundantly used source for knowledge representation by humans, for example, can we list the name of textbooks not written in natural language? It is hard to reply!

Disadvantages

In spite of strong points in favor of natural language-based representation, there are serious difficulties in realizing such representation for machines, due to the following reasons:

- The syntax and semantics of the natural language are very complex, which are not so easily understood. Hence, it becomes challenging and risky if solely depended on machines.
- The uniformity in representation is lacking—the sentences carrying the identical meaning can be represented in many different syntax (structures).
- There is a lot of ambiguity in the natural language, a sentence/word may have many different meanings, and the meanings are context-dependent. Hence, it is overly risky to try these for machines, unless the machines are having intelligence at par with the human.

1.12   Summary

Intelligence is defined as:

$$Intelligence = Perceive + Analyze + React$$

AI has its inter-related goals for scientific, as well as engineering areas. Its roots are in several historical disciplines, which include, philosophy, logic, computation, psychology, cognitive science, neuroscience, biology, and evolution.

The major sub-fields of AI now include: neural networks, machine learning, evolutionary computation, speech recognition, text-to-speech translation, fuzzy logic, genetic algorithms, vision systems and robotics, expert systems, natural language processing, and planning. Many of these domains have dependency and are inter-related, for example, neural network is one of the techniques for machine learning. The common techniques used across these sub-fields are: knowledge representation, search, and information manipulations.

Human brain and evolution are also the areas of AI modeling.

The study of logic and computers have demonstrated that intelligence lies in the *physical symbol system* (PSS)—a collection of patterns and processes. The PSS needs the capability to manipulate the patterns, i.e., it should be able to create the patterns, modify the patterns, and should be able to destroy the patterns. The patterns have important properties, that they can designate objects, processes, and other patterns. When the patterns designate processes, the later can be interpreted, i.e., to perform the process steps. The two significant classes of symbol systems we are familiar with those that are used by human beings, and those by computers. The later uses binary strings or patterns.

The PSSH (physical symbol system hypothesis) says that to achieve the intelligence, it is sufficient to have three things,

  i. a representation system, using which anything can be represented,
 ii. a manipulation system, using which the symbols can be manipulated, and
iii. search using which the solution can be searched.

For the above, it is in fact, not important as whether the medium of storage is the human brain (neurons) or the electronic memory of computer systems.

Various approaches for knowledge representations (KR) are:

  i. Natural languages versus databases.
 ii. Frame versus semantic network-based representation.
iii. Propositional and predicate logic-based representation.
 iv. Rule-based representation.

Knowledge representation helps to know the object or the concerned concept. Various characteristics of KR are:

  i. KR has syntax and semantics.
 ii. Requirements for knowledge representation are: adequacy of representation and inferencing, and efficiency of inference and acquisition.
iii. Its practical aspects are: complete, computable, and suppression of irrelevant data.
 iv. Components for KR are: lexical, structural, semantic, and procedural.

Exercises

1. Try to analyze your own learning behavior, and list the goals of learning, in the order of their difficulty level of learning.
2. List the living beings—human, dog, cow, camel, elephant, cat, birds, insects, in the order of their intelligence levels. Also, justify your argument.
3. Suggest some method to combine large number of human beings, in a group, and formulate a method/algorithm to perform the pipeline or any fast computing work.
4. Write an essay, describing the various anticipated theories as to how the human processes the information?
5. Consider a task requiring knowledge like baking a cake. Out of your imagination, suggest what are the knowledge requirements to complete this task.
6. Out of your reasoning, explain the distinction between knowledge and belief.
7. What is the basic difference between neural network level processing and processing carried out for human reasoning?
8. What are the major advantages of humans over modern computers?
9. List the examples where PSSH is not sufficient or not the basis of achieving Intelligence. Justify your claims.
10. Write an essay, describing how the things (objects/activities) are memorized by

    a. Plants.
    b. Birds.
    c. Sea animals.
    d. Land animals.
    e. Human.

    What can be the size of memories in each of the above cases?
11. Discuss the potential uses of AI in the following applications:

    a. Word processing systems.
    b. Smartphones.
    c. Web-based auction sites.
    d. Scanner machines.
    e. Facebook.
    f. Twitter.
    g. Linkedin.
    h. Amazon and Flipkart.

12. How the artificial neural networks (ANN) and genetic algorithms differ from each other in respect of leaning to be used for problem solution? (Note: You need to explore the AI resources to answer this question).

13. It is an accepted scientific base that physical characteristics of life are genetically transferred. Do you believe that information and knowledge are also genetically transferred? Justify for yes/no?

14. Are the beliefs of rebirth and reincarnation are also the goals of AI research? How and how not?

References

1. Buttazzo G (2001) Artificial consciousness: utopia or real possibility? IEEE Comput 34(7):24–30
2. Church KW, Lisa FR (1995) Commercial applications of natural language processing. Commun ACM 38(11)
3. Forrest S (1993) Genetic algorithms: principles of natural selection applied to computation. Science 261(13):872–878
4. Goldberg DE (1994) Genetic and evolutionary algorithms coming of age. Commun ACM 37(3):113–119
5. http://www.theguardian.com/technology/2014/jun/08/super-computer-simulates-13-year-old-boy-passes-turing-test. Accessed 19 Dec 2017
6. Ludger GF (2009) Artificial intelligence - structures and strategies for complex problem solving, 5th edn. Pearson, New Delhi
7. Munakata T (1994) Commercial and industrial applications of AI. Commun ACM 37(03)
8. Newell A, Herbert AS (1976) Computer science as an empirical inquiry: symbols and search. Commun ACM 19(3):113–126
9. Russell SJ (1997) Rationality and intelligence. Artif Intell 94:57–77
10. Russell SJ, Norvig P (2005) Artificial intelligence, a modern approach, 2nd edn. Pearson, New Delhi
11. Stanford Encyclopedia of Philosophy. http://plato.stanford.edu/entries/turing-test/. Accessed 19 Dec 2017
12. Turing AM (1950) Computing machinery and intelligence. MIND - Q Rev Psychol Philos Lix(236):433
13. Whitehead AN, Russell B (1910) Principia mathematica, vol 1 (Part I. Mathematical logic). Merchant Books. ISBN 978-1-60386-182-3