



College of Engineering & Technology
Computer Techniques Engineering Department
Artificial Intelligence – Stage 3
Introduction To AI



Introduction

1. What is Artificial Intelligence (AI)?

Human call themselves as wise because of their intelligence. For thousands of years, human have tried to understand how they think, perceive, understand, predict and manipulate a large and complicated world. The field of artificial intelligence (AI), is an attempts not just to understand but also to build intelligent entities.

AI is one of the relatively new fields in science and engineering. Work started soon after World War II, and the name itself was coined in 1956. AI currently encompasses a huge variety of subfields, ranging from general to specific, such as plying chess, proving mathematical theorems, writing poetry, driving a car on crowded street, diagnosis disease. AI is relevant to any intellectual task; it is truly a universal field.

Figure 1 explains eight definitions of AI. These definitions are laid out along two dimensions. The definitions on top are concerned with though processes and reasoning, whereas the ones on the bottom address behavior. The definitions on the left measure success in terms of fidelity to human performance. Whereas the ones on the right measure against an ideal performance measure, called rationality. A system is rational if it does the “right thing” given what it knows.



College of Engineering & Technology
Computer Techniques Engineering Department
Artificial Intelligence – Stage 3
Introduction To AI



<p>Thinking Humanly</p> <p>“The exciting new effort to make computers think . . . <i>machines with minds</i>, in the full and literal sense.” (Haugeland, 1985)</p> <p>“[The automation of] activities that we associate with human thinking, activities such as decision-making, problem solving, learning . . .” (Bellman, 1978)</p>	<p>Thinking Rationally</p> <p>“The study of mental faculties through the use of computational models.” (Charniak and McDermott, 1985)</p> <p>“The study of the computations that make it possible to perceive, reason, and act.” (Winston, 1992)</p>
<p>Acting Humanly</p> <p>“The art of creating machines that perform functions that require intelligence when performed by people.” (Kurzweil, 1990)</p> <p>“The study of how to make computers do things at which, at the moment, people are better.” (Rich and Knight, 1991)</p>	<p>Acting Rationally</p> <p>“Computational Intelligence is the study of the design of intelligent agents.” (Poole <i>et al.</i>, 1998)</p> <p>“AI . . . is concerned with intelligent behavior in artifacts.” (Nilsson, 1998)</p>

Figure 1. Some definitions of artificial intelligence

Historically, all four approaches to AI have been followed, each by different people with different methods. A human-centered approach must be in part an empirical science, involving observation and hypotheses about human behavior. A rationalist approach involves a combination of mathematics and engineering. The various group have helped each other. Let us look at the four approaches in more detail.

1. Acting humanly: The Turing Test approach

The Turing Test, proposed by Alan Turing (1950), was designed to provide a satisfactory operational definition of intelligence. A computer passes the test if a human interrogator, after posing some written questions, cannot tell whether the written responses come from a person or from a computer. For a computer to pass a rigorously applied test provides plenty to work on. The computer would need to possess the following capabilities:



College of Engineering & Technology
Computer Techniques Engineering Department
Artificial Intelligence – Stage 3
Introduction To AI



- Natural language processing to enable it to communicate successfully in English.
- Knowledge representation to store what it knows or hears.
- Automated reasoning to use the stored information to answer questions and draw new conclusions.
- Machine learning to adapt to new circumstances and to detect and extrapolate patterns.

Turing's test deliberately avoided direct physical interaction between the interrogator and the computer, because physical simulation of a person is unnecessary for intelligence. However, the so-called total Turing Test includes a video signal so that the interrogator can test the subject's perceptual abilities, as well as the opportunity for the interrogator to pass physical objects. To pass the total Turing Test, the computer will need.

- Computer vision to perceive objects
- Robotics to manipulate objects and move around.

These six disciplines compose most of AI. Yet AI researchers have devoted little effort to passing the Turing Test, believing that it is more important to study the underlying principles of intelligence than to duplicate an exemplar. For example, the quest for "artificial flight" succeeded when the researchers stopped imitating birds and started using wind tunnels and learn about aerodynamics.

2. Thinking humanly: The cognitive modeling approach

If we are going to say that a given program thinks like a human, we must have some way of determining how humans think. We need to get inside the actual workings of human minds. There are three ways to this: through introspection – trying to catch our own thoughts as they go by; through psychological experiments – observing a person in action; and through brain imaging – observing the brain in action. Once we have a sufficiently precise theory of the mind, it becomes possible



College of Engineering & Technology
Computer Techniques Engineering Department
Artificial Intelligence – Stage 3
Introduction To AI



to express the theory as a computer program. If the programs input-output behavior matches corresponding human behavior, that is evidence that some of the program's mechanisms could also be operating in humans. For example, the developer of "General Problem Solver GPS", were not content merely to have their program solve problems correctly. They were more concerned with comparing the trace of its reasoning steps to traces of human subjects solving the same problems. The interdisciplinary field of cognitive science brings together computer models from AI and experimental techniques for psychology to construct precise and testable theories of the human mind.

3. Thinking rationally: The "laws of thought" approach

The Greek philosopher Aristotle was one of the first to attempt to codify "right thinking", that is, irrefutable reasoning process. His syllogisms provided patterns for argument structures that always yielded correct conclusions when give correct premises. For example, "Socrates is a man, all men are mortal; therefore, Socrates is mortal. These laws of thought were supposed to govern the operation of the mind, their study initiated the field called logic.

Logicians in the 19th century developed a precise notation for statements about all kinds of objects in the world and the relations among them. By 1965, programs existed that could, in principle, solve any solvable problem described in logical notation. The so-called logicians tradition within artificial intelligence hopes to build on such program to create intelligent systems.

There are two main obstacles to this approach. First, it is not easy to take informal knowledge and state it in the formal terms required by logical notation, particularly when the knowledge is less than 100% certain. Second, there is a big difference between solving a problem "in principle" and solving it in practice. Even problems with just a few hundred facts can exhaust the computational resources of any computer unless it has some guidance as to which reasoning steps to try first.

4. Acting rationally: The rational agent approach



College of Engineering & Technology
Computer Techniques Engineering Department
Artificial Intelligence – Stage 3
Introduction To AI



An agent is just something that acts (agent comes from the Latin agere, to do). Of course, all computer programs do something, but computer agents are expected to do more: operate autonomously, perceive their environment, persist over a prolonged time period, adapt to change, and create and pursue goals. A rational agent is one that acts so as to achieve the best outcome or, when there is uncertainty, the best expected outcome.

In AI, the emphasis was on correct inferences. Making correct inferences is sometimes part of being a rational agent, because one way to act rationally is to reason logically to conclusion that a given action will achieve one's goals and then to act on that conclusion.

2. The foundations of Artificial Intelligence

In this section, we provide a brief history of the disciplines that contributed ideas, viewpoints, and techniques to AI. We organized these disciplines a series of questions. We certainly would not wish to give the impression that these questions are the only ones working toward AI.

A. Philosophy:

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

B. Mathematics

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

C. Economics

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

D. Neuroscience



College of Engineering & Technology
Computer Techniques Engineering Department
Artificial Intelligence – Stage 3
Introduction To AI



- How do brains process information?
- E. Psychology
- How do humans and animals think and act?
- F. Computer engineering
- How can we build an efficient computer?
- G. Control theory and cybernetics
- How can artifacts operate under their own control?
- H. Linguistics
- How dose language related to thought?
3. The History of Artificial Intelligence
- With the background material behind us, we are ready to cover the development of AI itself.
- A. The gestation of artificial intelligence (1943-1955)
- The firs work that is now generally recognized as AI was done by Warren McCulloch and Walter Pitts (1943). They drew on three sources: knowledge of the basic physiology and function of neurons in the brain; a formal analysis of propositional logic due to Russell and Whitehead; and Turing’s theory of computation. The proposed a model of artificial neurons in which each neuron is characterized as being “on” or “off”, with a switch to “on” occurring in response to stimulation by a sufficient number of neighboring neurons. They showed, for example, that any computable function could be computed by some network of connected neurons, and all the logical connectives (and, or, not, etc.) could be implemented by simple net structures. McCulloch and Pitts also suggested that suitable defined networks could learn.
- Donald Hebb (1949) demonstrated a simple updating rule for modifying the connection strengths between neurons. his rule, now called Hebbian learning, remains an influential model to this day.



College of Engineering & Technology
Computer Techniques Engineering Department
Artificial Intelligence – Stage 3
Introduction To AI



Two undergraduate students at Harvard, Marvin Minsky and Dean Edmonds, build the first neural network computer in 1950. The SNARC, as it was called, used 3000 vacuum tubes and a surplus automatic pilot mechanism from a B-24 bomber to simulate a network of 40 neurons.

B. The birth of artificial intelligence (1956)

US researchers interested in automata theory, neural nets, and study of intelligence. They organized a two-month workshop at Dartmouth college in Hanover in the summer of 1956.

C. Early enthusiasm, great expectation (1952-1969)

The early of AI were full of success, in a limited way. Given the primitive computers and programming tools of the time and the fact that only a few years earlier computers were seen as things that could do arithmetic and no more, it was astonishing whenever a computer did anything remotely clever. An example of an early AI projects is the General Problem Solver (GPS). The GPS was probably the first program to embody the “thinking humanly” approach. A second example is the famous physical symbol system hypothesis, which states that “a physical symbol system has the necessary and sufficient means for general intelligent action” what they meant is that any system (human and machine) exhibiting intelligence must operate by manipulating data structures composed of symbols. Another example is the IBM geometry theorem prover, which was able to prove theorems that many students of mathematics would find quite difficult. And many other examples.

D. A dose of reality (1966-1973)

AI researchers were not shy about making prediction of their coming successes. They made a concrete prediction: that within 10 years a computer would be chess champion, and a significant mathematical theorem would be proved by machine. These predictions came true within 40 years rather than 10 years. The early AI system faced some difficult problems: the first kind of difficulty arose because earliest programs knew nothing of their subject matter, they succeeded



College of Engineering & Technology
Computer Techniques Engineering Department
Artificial Intelligence – Stage 3
Introduction To AI



by means of simple syntactic manipulation. In other word, the absence of background knowledge was the first difficulty. The second kind of difficulty is the need for faster hardware and larger memories.

E. Knowledge-based systems: the key to power? (1969-1979)

The picture of problem solving that had arisen during the first decade of AI research was of a general-purpose search mechanism trying to string together elementary reasoning steps to find complete solutions. Such approaches have been called weak methods because, although general, they do not scale up to large or difficult problem instances. The alternative to weak methods is to use more powerful, domain-specific knowledge that allows larger reasoning steps and can more easily handle typically occurring cases in narrow areas of expertise (expert system).

F. AI becomes an industry (1980 - present)

The first successful commercial AI (expert system), R1, began operation at the Digital Equipment Corporation. The program helped configure orders for new computer systems; by 1986, it was saving the company and estimated \$40 million a year. By 1988, Digital Equipment Corporation group had 40 expert system deployed, with more on the way. Saving and estimate of \$10 million a year. Nearly every major U.S. corporation had its own AI group and was either using or investigating expert systems.

G. The return of neural networks (1986-present)

In the mid of 1980s at least four different group reinvented the back-propagation learning algorithm first found in 1969. The Algorithm was applied to many learning problems in computer science and psychology, and the widespread dissemination of the results in the collection of Parallel Distributed Processing caused great excitement.

H. AI adopts the scientific method (1987-present)

Recent years have seen a revolution in both the content and methodology of work in artificial intelligence. It is now more common to build on existing theories



College of Engineering & Technology
Computer Techniques Engineering Department
Artificial Intelligence – Stage 3
Introduction To AI



than to propose brand-new one, to base claims on rigorous theorems or hard experimental evidence rather than on intuition, and to show relevance to real-world applications rather than toy examples.

I. The emergence of intelligent agents (1995-present)

Perhaps encouraged by the progress in solving sub-problems of AI, researchers have also started to look at the “Whole agent” problem. One of the most important environments for intelligent agents is the Internet. AI systems have become so common in Web-based applications that the “-bot” suffix has entered everyday language. Moreover, AI technologies underlie many internet tools, such as search engines, recommender systems and web site aggregators.

J. The availability of very large data sets (2001-present)

Throughout the 60-year history of computer science, the emphasis has been on the algorithm as the main subject of study. But some recent work in AI suggests that for many problems, it makes more sense to worry about the data and be less picky about the what algorithm to apply. This is true because of the increasing availability of very large data sources. For example, trillions of words of English and billions of images from the web and so on.

4. The state of the art

What can AI do today? A concise answer is difficult because there are so many activities in so many subfields. Here we sample a few applications:

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