Chapter 0
Introducing AI

Artificial Intelligence is intended to understand, and build, intelligent entities. AI has produced many interesting and important results. E.g., on May 11, 1997, an IBM program, *DEEP BLUE,* beat the reigning world class champion. Although hard to predict in the details, it is clear that computers coupled with human-level intelligence will have a huge impact on our lives.

Another example is that in the Third DARPA Grand Challenge took place on November 3, 2007, an autonomous vehicle built by a CMU Racing Team completed the 55 mile course in an Urban setting, Victorville, California, in about four hours and won a USD 2M prize.

What to do with AI

With AI, we want to study how it is possible for a brain to *perceive, understand, predict,* and then *manipulate* the world, which is much larger and complicated than itself.

A much more difficult question is that, if we understand such a process, how could we build something with these features?

Both the newest and the oldest

AI is one of the newest discipline. It was formally initiated in the mid 1950's. Although it has accomplished quite a bit, there is still a long way to go.

On the other hand, the study of intelligence is also one of the oldest. For 2,000 years, philosophers have tried to understand how seeing, learning, remembering, and reasoning, could, as well as should, be done.

The emergence of computers in the early 1950's made speculations into a real experimental and theoretical science.

Computer and AI

With its huge memory and stunning speed, besides providing a vehicle for creating artificial intelligence entities, the computer also provides a *tool* for testing theories of intelligence, many of which could not stand the test, just like many of the physics theories.

AI currently consists of many sub-fields. From such general areas as logic reasoning, to some very specific ones, such as playing chess.

Many scientists of other areas often move into AI, where they find the theories and tools that systemize what they have been doing for a long time; while AI workers may choose to apply their methods to any area of human intelligence.

What is AI?

There are various definitions of artificial intelligence. Some of them emphasize more on the *thought process* and *reasoning;* while others emphasize on the behavior.

Another angle is that if it tries to characterize human performance; or some ideal performance, i.e., if *it* always does the right thing, or, *rationally.*

Thus, there could be four approaches, which have all been followed. A human-centered approach must be an experimental one, involving assumptions, and experimental confirmations; while a rationalist approach involves a combination of mathematics and engineering.

Act humanly: Turing Test

Turing proposed a test to provide an operational definition of intelligence.

He defines intelligent behavior as the ability to achieve human-level performance in all cognitive tasks, sufficient to fool an interrogator. More specifically, a human interrogator asks questions to, and get answers, from two sources, one is controlled by a computer, the other by a human being, without knowing which is which.

If the interrogator cannot distinguish the sources based on the answers, then we may say the compute exhibits sufficiently human intelligence to pass the *Turing test,*

What does it take to pass?

The Loebner Prize, that started in 1990, is an annual competition that awards prizes to the Chatter*bot* considered by the judges to be the most human-like among the competitors. The format of the competition is that of a standard Turing test.

This year's test will be held on September 6 at Brighton, UK with the top prize of $3,000.

To pass such a test, a computer would have to possess such abilities as *natural language processing*, *knowledge representation*, *automated reasoning*, and *machine learning*.

Think humanly: Cognitive modeling

Following this approach, we have to set up a theory on human being's thinking first by studying our own thinking process, or via psychological experiments.

Then, we have to express the theory as a computer program. If the program's I/O and timing behavior match with the real thing, we will have some evidence that some of the program's mechanisms may also be operating in humans.

The filed of *cognitive science* brings together computer models of AI and experimental techniques of Psychology to construct precise and testable models of the workings of human mind.

The GPS (General Problem Solver) program (1957, Simon and Newell) is to compare its reasoning steps with that of human subjects solving the same problems.

Think rationally: The laws of thought

Aristotle (384-322 BC) was one of the first to attempt to formalize the thinking process. His famous *syllogisms* provides pattern for argument structures that always gave correct conclusions given correct premises. E.g., "Socrates is a man; all men are mortal; therefore Socrates is mortal."

By 1965, programs came to life which, taking a problem represented in logic notation, is guaranteed to find a solution, if there is one, given enough space and time. The trick is that, if no solution exists, we might never know it (CS3780).

But it is not easy to covert informal knowledge into formal notation (How does a baby recognize her mother's face in just a few days?) Also, there is a big gap between being able to do it in principle and solve it in practice.

Act rationally: The rational agent

It is to act to achieve one's goals, given one's beliefs. An *agent* is just something that perceives and acts. In this approach, AI is regarded as the study and construction of rational agents.

Although making correct inferences is sometimes part of being a rational agent, it is not always the case, since there might not be any provably correct things to do; or it is simply an act of gut feeling, e.g., pulling one's hand off a hot stove.

All the cognitive skills listed under Turing Test are needed to allow rational actions.

This approach has two advantages. It is more general than the "laws of thought" approach since correct inference is only a useful mechanism, but not necessary.

It is also more amenable to scientific development than human centered approaches, because the concept of rationality is clearly defined and completely general.

The Grand Challenge might be regarded as a good test of this theory.

Philosophical foundation

At around 450 B.C., Aristotle, trying to formulate the laws that governs the rational part of the mind, developed an informal system of syllogism for proper reasoning, which in principle allowed one to mechanically generate conclusion, given initial premises. He also noted intuitive reason.

With at least part of the mind understood, the next step is to *consider the mind as a physical system.* Descartes (1596–1650) told the difference between *mind* and *matter,* as well as the associated problems.

One problem with a purely physical conception of the mind is that it leaves little room for free will. Although Descartes strongly believed the power of reasoning, he also believed that part of the mind is exempt from the physical laws. Thus, the dualism.

On the other hand, materialist held that all the world operates according to certain physical laws. In between there could be an intermediate position, in which one accepts that the mind has a material basis, but denies that it can be explained by a reduction to ordinary physical processes.

By and large, philosophy laid a tradition in which the mind was conceived of as physical device operating mainly by reasoning with the knowledge it has.

The next problem is that *where does knowledge come from?* Eventually, Rusell (1872–1970) suggested that all knowledge can be characterized by logical theories connected to observation sentences corresponding to sensory inputs. Confirmation theorists also tried to understand the process of acquiring knowledge from experience.

The last element of philosopher's perception of the mind is the connection between knowledge and action. Particularly, *how can an action be justified?*

A typical technique in this regard is the means-ends analysis as put forward in GPS program by Newell and Simon around 1970's. The following argument is fairly common:

> I want to get an 'A' in this course. The difference between what I have (nothing) and what I want is one of a number, at least 92.5.
>
> To reach that goal, I have to do well in the following areas: homework, projects, midterm, and final. Thus the ends of getting of an A justifies my action of working hard in reading the text books, staying in programming lab until 2 a.m., bugging the professor in his office the whole time, etc..

Mathematical foundation

Mathematicians provided a solid, and formal, foundation for AI through three main areas: computation, logic and probability.

The notion of expressing a computation as a formal algorithm goes back to the Arabian world in the 9th century.

Logic goes back at least to Aristotle, but it was Frege (1848–1925) that coined the basic notation of the first-order logic that is used today as the most basic knowledge representation system.

Tarski (1902–1983) introduced a way to relate the objects in logic to those in the real world.

Logic is the basis of the Prolog language we will use in this course.

How far can it go?

But, how far could we push the idea of logic and computation in the real world?

Hilbert (1862–1943) asked if there is an algorithm for deciding the truth of any logic proposition involving the natural numbers, as one of the 23 problems that he presented in 1900 for the new, now old, century.

Gödel (1906–1978) showed that there is a procedure to prove a true statement in the first-order logic; but, much more deeply, he proved the *incompleteness result* in 1931, namely, in any language that can express the properties of natural numbers, there exist true statements whose truth cannot be established by any algorithm.

Now what?

Turing (1912–1954) then tried to characterize the class of functions that can be computed, via his Turing Machine. Later, the Church-Turing Thesis states that a function is computable iff it can be computed by a Turing Machine. It turns out that some functions are not computable, e.g., the halting problem.

Besides the notion of *non-decidability* and *UN-computability*, the notion of *intractability* has an even bigger impact on AI. A class of problems is called intractable if the time required to solve *instances* of those problems in that class grows exponentially in terms of the size of those instances.

For example, the problem of Hanoi Tower is intractable.

Computability

Another important concept is *reduction.* It is a general transformation from one class of problems to another, such that solutions to the first class can be found by transforming to the problems of the second class, and then solving the latter problems.

For example, the problem of cutting the lawn can be reduced to the problem of buying a mower.

A problem is called *NP-complete* if it can be solved by a non-deterministic machine in polynomial time, and all such problems can be reduced to it.

It is commonly accepted that such problems are *most likely* intractable. The *Satisfiability* problem, i.e., how long does it take to show if a Boolean expression can be satisfied, is among the first problem to be proved NP-complete.

To recognize an intractable problem, one can reduce an *NP-complete* problem to it. If this is doable, then the underlying problem is also NP-complete, thus *likely* intractable as well.

We will discuss more about such basic concepts in CS3780.

Probability

The third great mathematical component of AI is the *probability theory.* Initially studied in terms of possible outcomes of gambling events, by Cardano (1501–1576), many mathematicians advanced the theory and introduced new statistical methods.

Now, probability has become an invaluable part of mathematics, particularly useful in dealing with *uncertainty* and *incompleteness.*

Bayes (1702–1761) proposed a rule for updating subjective probabilities in terms of new evidence. Bayes' rule forms the basis of the modern approach to uncertain reasoning in AI systems.

We will look into the AI application of probability in Expert System.

Psychological foundation

The view that the brain possesses and processes information constitutes the essence of *cognitive psychology.* It can be traced back at least to the works of William James (1842–1910).

Around 1943, Crack added the missing mental step between stimulus and response. He specified three key components of a knowledge-based agent:

1) the stimulus must be translated into an internal representation, which is

2) manipulated by cognitive processes to derive new internal representation, and

3) these are then retranslated into action.

A contemporary view

Nowadays, it is generally accepted that “a cognitive theory should be like a computer program.”

That is to say cognition consists of well-defined transformation processes operating at the level of information carried by the input signals.

This view constitutes the psychological foundations of Expert System, one of the most commercially successful areas of AI.

Computer engineering

For AI to succeed, we need two things: intelligence and an artifact. The computer has been unanimously accepted as the artifact with the best chance of demonstrating intelligence.

Starting in the mid-40's, each generation of computer hardware has brought an increase in speed and capacity, and decrease in price.

Regularly doubling performance every two years, there is no immediate end in sight for this rate of increase. Massively parallel machines promises to add several more zeros to the overall performance.

The other side

AI also owes a great deal to the software side of computer science, which has supplied the operating systems, programming languages, and tools needed to write up modern programs.

Many ideas initiated in AI have found their ways in the mainstream computer science, including time sharing, interactive interpreters, the linked list data structure, etc..

Past, present and future

Starting around mid-40's, and officially born in Dartmouth in 1956, AI has gone through many stages: the great expectation between 52 and 69; dose of reality in its middle stage; emphasis on knowledge based system between 69 and 79; .... AI Has been growing up.

Nowadays, it is more common to build on existing theories than to propose new ones, to base claims on rigorous theorems or hard experimental experience rather than on intuition, and to show relevance to real-world applications rather than toy examples.

Some of the more eye-catching applications include speech recognition (voice activated GPS), expert systems, planning, etc..