32002: Artificial intelligence (Introduction , roots, goals, and sub-areas) **Spring 2015**

Lecture 1: January 1, 2015

Lecturer: K.R. Chowdhary *: Professor of CS (VF)*

1.1 Introduction

AI may be defined as a branch of Computer Science, that is concerned with *automation of Intelligent behavior*. A very 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." W. F. Dearborn
- "Ability to adapt oneself adequately to relatively new situations in life." R. Pinter
- "A person possesses intelligence insofar as he has learned, or can learn, to adjust himself to his environment." - S. S. Colvin
- "The ability of an organism to solve new problems." - W. V. Bingham
- "A global concept that involves an individual's ability to act purposefully, think rationally, and deal effectively with the environment." - D. Wechsler
- "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." - L. S. Gottfredson and 52 expert signatories

AI's principles includes - data structures used and knowledge representation, the algorithms needed to apply the knowledge and language, and programming techniques used in their implementation. In part, the appeal of AI is that it offers unique and powerful tools for exploring exactly these questions. AI offers the medium, and the test-bed for theories of intelligence. Such theories may be stated in the language of computer programs, and can be tested and verified using the execution of these programs on computer.

1.2 The Turing Test

In 1950, in a article "Computing Machinery and Intelligence," *Alan M. Turing* proposed an *empirical test* for machine intelligence, now called *Turing Test*(see figure 1.1). It measures the performance of an intelligent machine against human, 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 machine, and communicates to these two solely by textual device like terminal.

The interrogator is asked to distinguish machine from human, solely on the basis of their answers and questions over the device. If interrogator is not able to distinguish machine from human, then, Turing argues that machine can be assumed to be intelligent. Interrogator may ask highly computation oriented questions to identify machine, and other questions related to poetry etc., to identify the human.

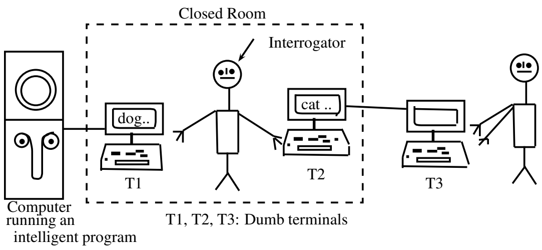

Figure 1.1: Imitation Game: Turing test.

The game (with the player machine omitted) is frequently used in practice under the name of viva-voce to discover whether some one really understands something or has 'learnt it parrot fashion'.

Some of the arguments *for* and *against* the above test can be as follows:

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, as a terminal is used as 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*.

1.3 Goals of AI

AI is the area of computer science, concerned with designing intelligent computer systems, i.e., systems that exhibit characteristics we associate with intelligence in human behavior, like, understanding language, learning, reasoning, solving problems, and so on.

The ultimate goal of this field for some researchers is to emulate human cognition, to others creation of intelligence without concern for human characteristics, and still to others to create useful artifacts without concern for abstract notions of intelligence.

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

The *engineering goal* is to solve real world problems using AI techniques such as knowledge representation, learning, rule systems, search, and so on.

1.4 Roots of AI

Artificial Intelligence has identifiable roots in a number of older disciplines, particularly:

- Philosophy
- Logic/Mathematics
- Computation
- Psychology/Cognitive Science
- Biology/Neuroscience
- Evolution

There is inevitably much overlap, e.g. between philosophy and logic, or between mathematics and computation. By looking at each of these in turn, we can gain a better understanding of their role in AI, and how these underlying disciplines have developed to play that role.

1.5 Artificial Consciousness

Since the beginnings of computer technology, researchers have speculated about the possibility of building smart machines that could compete with human intelligence. Given the current pace of advances in artificial intelligence and neural computing, such an evolution seems to be a more concrete possibility. Many people now believe that artificial consciousness is possible and that, in the future, it will emerge in complex computing machines.

However, a discussion of artificial consciousness gives rise to several philosophical issues:

- Can computers think or do they just calculate?
- Is consciousness a human prerogative?
- Does consciousness depend on the material that comprises the human brain, and
- Can computer hardware replicate consciousness?

1.6 Common Techniques used in AI

Even apparently radically different AI systems (such as rule based expert systems and neural networks) have many common techniques. Four important ones are:

Representation: Knowledge needs to be represented somehow perhaps as a series of *if-then rules*, as a *frame based* system, as a *semantic network*, or in the connection weights of an *artificial neural network*.

Learning: Automatically building up knowledge from the environment such as acquiring the rules for a rule based expert system, or determining the appropriate connection weights in an artificial neural network.

Rules: These could be explicitly built into an expert system by a knowledge engineer, or implicit in the connection weights learn by a neural network.

Search: This can take many forms perhaps searching for a sequence of states that leads quickly to a problem solution, or searching for a good set of connection weights for a neural network by minimizing a fitness function.

1.7 Sub-fields of AI

If we mean AI to be a realization of real human intelligence in the machine, its current state may be considered primitive. In this sense, the name artificial "intelligence" can be misleading. However, when AI is looked at as advanced computing, it can be seen as much more. In the past few years, the repertory of AI techniques has evolved and expanded, and applications have been made in everyday commercial and industrial domains. AI applications today span the realm of manufacturing, consumer products, finance, management and medicine. Implementation of the correct AI technique in an application is often a must to stay competitive. Truly profitable AI techniques are even kept secret.

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

- *Neural Networks* brain modeling, time series prediction, classification
- *Evolutionary Computation* genetic algorithms, genetic programming
- *Vision* object recognition, image understanding
- *Robotics* intelligent control, autonomous exploration
- *Expert Systems* decision support systems, teaching systems
- *Speech Processing* speech recognition and production
- *Natural Language Processing* machine translation
- *Planning* scheduling, game playing
- *Machine Learning* decision tree learning, version space learning

Most of these have both engineering and scientific aspects. Many of these we are going to discuss in this text. following is brief Introduction to each area.

1.8 Considerations for Knowledge Representation

The following are the aspects of knowledge representation:

- What is Knowledge?
- What is a Knowledge Representation?
- Requirements of a Knowledge Representation
- Practical Aspects of Representations
- Knowledge Representation using Natural Language
- Databases as a Knowledge Representation
- Frame Based Systems and Semantic Networks
- First Order Logic as a Representation
- Rule Based Systems
- Which Knowledge Representation is Best?

References

[GFL09] D. George F. Ludger, "Artificial Intelligence - Structures and Strategies for Complex Problem Solving," *5th Edition, Pearson Education, India*, 2009, Chapter 1.