2.4.2 Artificial Intelligence and Expert Systems

T175 - Responsibility for the performance of an expert system—knowledge engineer, informant, programmer, company that sold it, the buyer/consumer by Nitish

Let us define a few terms before proceeding –

Expert System

- A software system with two basic components: a knowledge base and an inference engine. The system mimics an expert's reasoning process.

- A type of application program that makes decisions or solves problems in a particular field by using knowledge and analytical rules defined by experts in the field.

Knowledge Base –

A collection of facts, rules, and procedures organized into schemas. The assembly of all the information and knowledge of a specific field of interest.

This paper is about responsibility for the performance of an expert system. To what extent do knowledge engineers, programmers, companies, buyers and informants have responsibility towards an expert system?

1) Knowledge engineering (KE) refers to the building, maintaining and development of knowledge-based systems. It has a great deal in common with software engineering, and is related to many computer science domains such as artificial intelligence, databases, data mining, expert systems, decision support systems and geographic information systems. (WIKI)

The knowledge engineer thus, builds, maintains and develops the knowledge based systems. He has a big responsibility in the 'health', progress, work, durability and many other factors of the expert system.

2) An informant (sometimes informer) is someone existing inside a closed system who provides information of that system to a figure or organization who exist outside of that system. (WIKI)

Here, we are talking about a person who informs and gives information to (supposedly) engineers, buyers, sellers, etc. He has a responsibility in spreading information and distributing it.

3) A programmer or software developer is someone who programs computers, that is, one who writes computer software. The term computer programmer can refer to a specialist in one area of computer programming or to a generalist who writes code for many kinds of software. (WIKI)

A programmer has a big responsibility in the 'purpose' of the system. He programs the whole expert system, maintains it, and basically, builds the software. His job is to see to it that the system works fine and does the work it is intended to do. He is in contrast with the knowledge engineer.

4) The Company that sold it has a responsibility in deciding the economic factors, the profit and the loss. It is also responsible for the copyright as well as the ownership of the system. It decides where the system goes, where it is used, how it is used, etc.

5) The Buyer/Consumer has a responsibility in handling the system. It will call a programmer, or ask the previous programmer who programmed it into working to do as they want it to do. They will have control over the expert system, and decide the purpose and use of it.

T176 - Value of the development of AI as a field, for example, whether it is an appropriate place to put economic resources by Isaku

Artificial intelligence has been a publicly well known technology thanks to S.F novels by authors such as Jules Verne and Isaac Asimov. However, the sophisticated A.I as depicted in these novels still remains a scientific dream like flying cars. Still great effort is made by scientists all over the world to make this dream come true, since it is expected to benefit mankind in many fields.

Now the question is, which field is appropriate to develop this technology in. Presently the field that is said to come the closest to completing this technology is the military.

First, why would the military research artificial intelligence? This may not be the appropriate question to ask. The question should be why wouldn't the military research artificial intelligence. Indeed the production of a proper artificial intelligence would mean a lot to the military. The type of artificial intelligence that the military has in mind has two types. The first one is in short a all powerful commanding office by itself.

By having this A.I, the military will have a lesser need for large commanding stations with operators, generals etc. Instead the A.I will be connected to the military's network, and will be able to efficiently receive information, theoretically analyze it, then issue effective orders all by itself. If this system comes to use, the army will be able to operate with less human staff, ultimately resulting in the decrease in the military's expenses.

The system will also be in theory more efficient than human staff since the different types of jobs are merged into one system, resulting in a more connective and fast system. Also all decisions the A.I make will be based on theoretical data, so it will make less mistakes than humans. The second type of A.I is an independent battlefield operational type. Long story short, it is a mechanical soldier, or a robot. The merit from this is of course the lives of millions of soldiers that will be saved by this technology.

These merits however, are only merits if you see them from the military's point of view.

or

Value of the development of AI as a field, for example, whether it is an appropriate place to put economic resources by Chirag

Computer scientists like to view a program as an abstract specification of a machine, describable behaviorally in terms of the input/output relationship resulting from its computation. The machine's product is its output, representing the value of a function at the point represented by its input. Often we find it helpful to view this product at a higher level, say, as the solution to some well-posed problem. Inevitably, this problem bears on what we are to do, that is, some course of action to be embarked upon. (Conceptions of computation as answering questions are a relic of the era when human intermediaries were necessary to perform the transduction from computation to action.) In this view, the computer is a decision machine, where a decision is the resolution of a distinction among potential courses of action.

It is widely recognized that many of the problem-solving techniques developed in AI research (e.g., so-called classical planning) need to be generalized to accommodate uncertainty and graded preferences. Work in decision-theoretic planning (Hanks et al., 1994) is beginning to address these problems, adopting a more comprehensive framework for principled resource allocation while attempting to retain useful computational and representational techniques from prior AI work.

Most of microeconomic theory assumes that individual agents are rational --acting so as to achieve their most preferred outcome, subject to their knowledge and capabilities. Indeed, this rationality abstraction is perhaps the single methodological feature that most distinguishes economics from the other social sciences.

This approach is highly congruent with much work in Artificial Intelligence. About fifteen years ago, Newell (1982) proposed that a central characteristic of AI practice is a particular abstraction level at which we interpret the behavior of computing machines. Viewing a system at Newell's knowledge level entails attributing to the system knowledge, goals, and available actions, and predicting its behavior based on a principle of rationality that specifies how these elements dictate action selection. Rationality as applied here is a matter of coherence, defining a relation in which the knowledge, goals, and actions must stand. This is exactly the Bayesian view of rationality (standard in economics), in which knowledge and goals (or beliefs and preferences) are subjective notions, constrained only by self-coherence (consistency) and coherence with resulting behavior.

In human societies, computational power is inherently distributed across many relatively small brains resident in separate skulls, connected by costly, low-bandwidth, error-prone communication channels. Moreover, authority over activity is separately controlled by the local computational units. It is therefore not surprising that economics focuses on the decentralized nature of decision making. A primary aim of the discipline is to explain the aggregate results of alternate configurations of interacting rational agents.

The case for decentralization in computational environments, where communication is usually more direct and configurations more controllable, is less straightforward. Nevertheless, a variety of technological and other factors are leading to computational environments that are increasingly distributed. At this writing, the development and promotion of "software agents" (not necessarily derived from AI technology) is a prominent activity. Although interpretations of software agency vary widely, typical conceptions involve autonomy of action, modularity of scope and interest, and interaction with other agents. Understanding and influencing configurations of software agents is directly analogous to the problem faced by economists.

It is not possible in this short position paper to survey the large body of work on probabilistic reasoning, decision-theoretic planning, game-theoretic analysis of multiage systems, etc., that has made its way into AI over the last ten years. Suffice it to say that the field has been far more open in the previous decade to ideas that could be broadly characterized as economic. That these ideas have had significant impact in particular subfields is reflected in the ubiquity of concepts of resource allocation and rationality in the recent AI textbook of Russell and Norvig (1995).

This is not a surprise. As I have attempted to point out, the goals of AI and those of economics overlap substantially, and are analogous in many of the non-overlapping regions. AI is the branch of computer science that is concerned with the substance of behavior, and with deriving general principles for designing deciding agents. In so doing, AI unapologetically invokes rationality concepts, and aims to render the rationality abstraction an operationally viable approximation. When activity is decentralized, AI considers interactions in social terms.

The point of all this is not, of course, to suggest that economics has all the answers to AI problems. But recognizing that AI's problems are in large part economic does help us to formulate the questions, and opens to us a variety of concepts and techniques that offer a starting point on potential solutions. Success in AI would mean an account of the economics of computation, and one way toward this goal starts with some computation of economics.

T177 - Ethical issues of various applications of AI, for example, replacement of human workers, handing decision-making tasks to a computer by Isaku Oba

There have been many discussions done regarding the future application of artificial intelligence. It is true that the actual introduction of A.I will become a controversial one, and that it will receive much criticism for various reasons. This is because some consider the A.I to be an unethical system. For one, the A.I can replace humans in jobs at many areas, considering that the A.I uses mechanical procedures for work.

The A.I will be able to do reliable and efficient quality work, making it superior to the work of humans who make mistakes often. Also one A.I can do the take on multiple tasks at once, making it a much more economical employee. The jobs with which the A.I will replace will include jobs that include theoretical and mathematical procedures, such as computer operators, programmers, accountants etc. The demand for human labor will obviously drop drastically, resulting in many jobless citizens wandering around. As terrible as this may seem, there are some views that consider this situation something unavoidable.

The dramatic increase of the human population in the last century has resulted in the increasing rate in the destruction of the earth's environment. If the human population continues to increase, so will the rate at which we exploit the earth which means pretty soon the earth will come to a point where it will not last. One possible solution to this problem is to decrease the number of humans, and the A.I might be the solution to this problem. If the A.I is brought to use, the jobs available for humans will not be enough which might ultimately force the number of the human population to decrease.

Another reason that the A.I is resented by humans is its roles that it plays in S.F novels. In these novels, the A.I has the ability to think and learn like human beings, and its superior brain power becomes a threat to mankind after the A.I malfunctions turns against its creators. Considering that the A.I is considered a duplication of the human brain, it is reasonable that the A.I might turn against its creators.

But in real life, the A.I is not likely to become a copy of the human mind. It is not likely that the researchers will include concepts of the human brain such as its ego, or emotions which will not do any good to help assist humans. After all the A.I is only a tool, and it would not make sense for a tool to have unnecessary functions which could pose a threat towards its creators. Therefore it would be nothing but pure paranoia to think that the A.I would grow to a point where it will rebel against the humans.

The last reason is sort of a more frequently discussed one. This is because a similar problem exists with the technology of cloning. This reason is the dehumanization of war. Presently the military of developed countries are thinking of two next generation soldiers which will minimize their casualties. The first is the use of clone soldiers using copies of existing humans to fight instead of real humans. This technology is based on the ethical value that clone humans are not humans, but tools created by humans.

The second is a mechanical soldier equipped with an A.I. This will allow the existence of an army totally absent of humans, which the thinks will make war a more humane process. However, these two next generation soldiers both hold the same kind of problems. The use of non-human soldiers in battlefield will provide a justification for countries to pursue war with ease, making war a much less opposed conduct. Also, if war is to be conducted with out the pain of humans and physical loss, it will transform into a more game like procedure, in which less is learned by the experience of it.

Overall, the A.I is a technology which shows great promises in guiding the humans to the next age, but we must not rush in trying to create it as there are crucial problems regarding it that must be dealt with prior to its ues.

or

Ethical issues of various applications of AI, for example, replacement of human workers, handing decision-making tasks to a computer by Marek

Robotics and Morality

Morals and profit are not two things that usually go hand in hand when associated with the hard world of business transactions. More often than not, companies choose monetary gains over moral values, and because of this, humanity has both benefited and suffered greatly.

Many firms have questioned whether they should retain their human workers, or let the quicker, more efficient robot workers/computers take over the management section. Both have their pros and cons. But which one would lead to a better, stronger company?

The advantages of having a computer-controlled manager is obvious. First off, the company would have a reliable worker that never took sick days, could work twenty four hours a day, seven days a week, and always made the most efficient, cost cutting decision without letting personal emotions cloud its judgment. At first glance, it seems like the right choice, but there are several factors that weigh against an artificial manager.

On of the first things would be that the moral of the company's workers would be at an all time low. It would be extremely degrading for a person, who prided themselves on being a loyal and skilled worker, to have to take orders from a machine, who has little or no idea of what they have to do to keep the firm active and running.

They would feel misunderstood, ignored, and maybe even angry at how they are being treated. This will detract from their job performance, and in the end, lower their efficiency. After all, how could a twisted amalgamation of steel, silicon and computer chips possibly comprehend what the average worker has to go through everyday? These mutinous feelings will lead to lowered performance, latent anguish, and maybe even incidents of outright disobedience. Because of their inability to feel, computers help destroy a company.

On the other hand, a talented manager knows exactly what to say and do to get his laborers in a frenzy of work, utilizing incentives, and select words of praise to keep them going long into the night. He will understand what kind of troubles they are facing, and will attempt to help them by making them feel like part of the “team”. A human manger could identify what kind of problems the employee was facing, and through his or her understanding and compassion, could potentially cause the employee to be happier about his or her work, and make them work harder as a form of compensation for their boss.

However, human managers have always had their flaws, namely that of emotions.

Understanding can be misinterpreted as excessive leniency, which in turn gives the employees the belief that they are allowed to “slack off” from working. Grudges harbored against stubborn workers, and irrational dislike can also amount to a loss in efficiency during work, as both people spend time they could be working on pondering how they can get their own back on that particular individual. As ridiculous as it may sound, these petty vendettas can get in the way of getting some real work done, and therefore do affect how we work.

Overall, this is a very difficult question to answer. To have a purely professional army of robots to build and maintain the company, would be absolutely impossible, in this day and age. It seems ironic that the very qualities that make us human and the ones that keep us going in the business world today.

T178 - Social impact of the use of "smart" machines on everyday life by Takafumi

Greater social impact would occur in society if we use smart machines in our everyday life. These machines can make peoples life style easy, but also make them lazy too. This totally destroys working system that we have in this world. But in other point of view, we can make our lifestyle really easier to live and become happy.

If we look in dark side of this topic, human would become really lazy. People might stop thinking and depend on machine too much. This means that we as a human being are getting controlled by the robots. But humans won't notice that and keeps that lifestyle because it is way better and easier for them than doing work and thinking about something. This totally makes the human being lazy animal and this is morally wrong. People have to work and think to keep themselves as the real human.

The bright side for this idea is that we get to have an easier life than now. In moral view, this can help people to do more good morals like helping each other and all other thing you can imagine about it. One reason is when people get enough happiness, they would like to share them with other people by doing same thing that they felt it was graceful for them. So pretty much morally, this would be a good thing for them.

Well the conclusion to the state that usage of smart robots into our society or everyday world that there would be many moral problems but this can make our lifestyle really easier than before. But this also can mean that people would lose the capability to think and just live.

or

Social impact of the use of "smart" machines on everyday life collected by Chirag

Smart machines can be used for various purposes that will be to the advantage of the humans. Therefore, socially they have a very positive impact by helping us making our lives easier and more secure.

Smart Machines in Education is woven in the utilization of the research results from cognitive science and artificial intelligence to advance our perception of technology for education into the next revolution. The main focus is placed on the illustrations of educational systems, and the intention is to promote the “intelligent” ideas in virtual learning environments.

Learning from costless failure is one of the important advantages in educational simulation. Expectation failures synthesize situations that students could encounter in their real-life, and encourage them to maturate their scientific investigation skills in order to learn the lessons from the failures.

Another novel idea is the creation of a teachable agent (Biswas, Schwartz, & Bransford). One can often learn a lot from the feedback and responses from the taught during the teaching process. The teachable agents offer this opportunity to students while protecting them from being harmed by “in-experienced teachers” (other students) at the same time.

For those who are interested in intelligent tutoring systems (ITSs), the book provides the discussion of Woolf and her colleagues on the capabilities that an ITS could have, and Bellman’s reflections on the CAETI Program, which is a large-scale government-sponsored technology insertion program in USA

Reviewer: Tai Yu Lin.

Use against Terrorism

2.4.2 Artificial Intelligence and Expert Systems

No war is won without rest-rate minds making powerful machines. Arguably our chief strategic mistake in battling terrorism has been that we have ignored this axiom. You would probably not be around to read this piece had it not been for the fact that Alan Turing and other brilliant minds harnessed the very rust computers to crack the Nazi Enigma Machine, which the Germans used with supreme condense to encrypt and decode their secret communiqués in World War II. Nor would you and yours be alive today had it not been for the minds behind the Manhattan Project: They beat Hitler and his band of terrorists in the nuclear race, and two unthinkably smart machines brought those who perpetrated Pearl Harbor to their knees in a surrender that saved tens of thousands of lives. I mean no disrespect to soldiers; they alone put their lives directly on the line. But if they do so without holding thunder in their hands, they cannot win. And that thunder is conceived, designed, and produced by scientists and engineers. In our new war, not just any kind of machine will do.

We need smart machines, so the specific minds we need are those working in the field of Artificial Intelligence (AI). This is painfully easy to see. Consider the case of jetliners. When you think about such potential missiles from the standpoint of AI, it is manifestly idiotic that we have built them devoid of ability to reason. A high-tech plane, for God's sake, should know who is allowed to them it and where it is allowed— and if things don't add up, it should not cooperate. To build such AI into a jet is not difficult; it isn't cheap, but it isn't hard. To our great and obvious peril, aircraft manufacturers have left those who specialize in machine reasoning on the sidelines. We are also in desperate need for the smart machines. You should not be able to get on a plane unless a machine knows who you are, and knows that the odds of you doing anything nefarious are in intestinal. One instantaneous retina scan before you board, with a knowledge base to refer to, would allow a machine to calculate in a millisecond whether you can board or not.

T179 - Ethical issues related to military applications of AI, for example, smart weapons, reconnaissance, decision making by HeeJun Son

Artificial Interlligence(AI) means intelligent machines manufactured by human beings. It also refers to a trait of intelligence that is existed by an artificial entity. Research in AI is concerned with producing machines to mechanize tasks requiring intelligent behavior. Examples include control , the ability to answer diagnostic and consumer questions, voice and facial recognition. The study of AI has also become an engineering discipline, focused on providing solutions to real life problems. But it is also applied to military weapons, causing a lot of problems. Smart weapons are an example of this.

Smart weapons are guided weapons intended to maximize damage to the target by using a laser, television, or satellite guidance system. The improvements in accuracy by using AI enable a target to be effectively attacked with fewer and smaller bombs. Smart weapons, which use guidance systems that rely on external assistance, are different from brilliant weapons, which are totally self-guided.

In the case of a smart bomb with a laser guidance system, an aircraft pilot aims a laser beam at the target, a computer keeps the beam locked on the target, and the bomb has a sensor programmed to find the reflection of the laser's light. A guidance computer adjusts the path of the bomb after it is released, using movable fins to steer. Satellite-guided bombs have guidance computers that use signals from navigation satellites to confirm that they are on target; the tail fins are adjusted to control the bomb's course as it falls. Cruise missiles are an example of this.

The problem with these weapons is the technology of AI, which is made to improve real life problems, is actually used in the mass destructive weapons. The fact AI is being used in the wars make many people sad including me. But on the other hand, we are making those weapons to protect themselves. Therefore, it is hard to say we shouldn't develop those weapons by using AI.

T180 - Implications of creative production by computers using AI, for example, Aaron, an expert system, creates visual art by Dwarkesh

The incorporation of the computer into the artistic arena has nourished a revolution in Contemporary Arts. As a result, important changes have taken place affecting not only the process of generation of artistic works, but also the role of the artist, of the audience and the channels used to display the works of art.

Digital Media:

There is a revolution going on in Contemporary Art that is being nourished by Digital Technology, in particular the Internet. In the past, Art was mainly confined to museums and art galleries, and to experience it, one had to travel to a specific location and enter a physical space that housed the artworks. These days, the Internet has changed all that. Artists are now producing artistic works intended for a global audience, using the computer as a tool for creativity and creating specifically for the Internet. A creative dialogue has been initiated between Art and Technology that has broadened the traditional forms of expression. Visual languages have started to play an important role in electronically mediated communication. Iconography has been developed as an important component of user interface design, and interactivity has become the main purpose of digital artistic tendencies that intend to get the audience involved in the creative process through the use of active agents in communication that replace the old passive communicators. In the turn of the millennium, Digital Art and new Multimedia Technology are being incorporated to Contemporary Art in order to produce artistic manifestations that gather together sound, music, movement, spatial and aesthetic components, and boost up a cycle of increasing globalization of culture.

2.4.2 Artificial Intelligence and Expert Systems

Artificial Intelligence:

The existence of programs like Aaron, that appear to be creative and recognize creativity, leads us onto the analysis of the role that computers play in the creative process.

Computer music makes use of sounds and allows composers to experiment with computer-generated chords or phrases that they might not have ever discovered by themselves. In the same way, computer graphics or computer animations sometimes produce images of fascinating beauty, and allow human artists to create brand new types of visual effects. Also, writing programs may help both -children and adults- to plan and produce texts of a level of complexity and coherence that could had been hardly achieved without them. Nevertheless, in all these cases a certain level of human supervision is essential to obtain the output concerned.

On the other hand, sometimes it is possible for the human artist to remain apart from the creative process and give the computer full autonomy to become the originator of the artistic work, emulating in this way the creative talents of human musicians, painters, poets or novelists. One of the most successful programs of this type is Aaron, the drawing program which we were referring to earlier.

Aaron is not a typical image generator of what has come to be known as computer art or digital art. Aaron does not generate geometric forms, certainly interesting, but infinitely repeatable. It does not either produce fractals, beautiful and random, despite of not being representative of the items which comprise the world. Aaron is not a tool for painters, designers, draftsmen or animators to be used as a medium to express the creative ideas of the human user. Instead, Aaron is a computer program with a software interface to a hardware drawing device that creates original pictures, each picture different from the others and each one indistinguishable by the uninformed observer from the work of a human artist. As such, Aaron is significant to the computer scientist as well as to the artist, because it uses artificial intelligence to encapsulate and replicate much of the behavior that the artist unconsciously employs to create art.

In fact, Aaron is an interesting computational project that has been developed in several phases. The development from one version of Aaron to the next generally involved a fundamental change in the nature of the program and a radical alteration of the conceptual space that it inhabits. The early Aaron concentrated on spontaneous drawings of abstract forms which could sometimes resemble rocks, sticks, and occasionally, some strange forms of birds or insects. However, under no circumstance was the computer able to produce anything similar to human figures. Later developments of Aaron produced more complex drawings depicting groups of human figures in a jungle or vegetation, whereas the programs most recent images display human figures of a fully three-dimensional type. All the different versions of Aaron can draw new pictures and produce aesthetically satisfying results at the touch of a button.

After evaluating these drawings, some experts agreed that Aaron meets all the criteria necessary to be regarded as a creative tool. It shows the ability to inhabit and explore a conceptual space rich enough to generate indefinitely many surprises. Aaron also shows the capacity of judgment that makes this program able to reconsider past choices and decide what to do next, as well as the aptitude to evaluate various possible structures for it, in order to avoid nonsense and cliché. However, some others experts still refuse to call Aaron creative, since they believe in the assumption that no computer program can really be called creative, no matter what novelties it manages to produce.

T181 - Access to the knowledge base underlying an inference engine in an expert system, for example, whether people affected by decisions made using an expert system should have access to the rules by which the decision was made. by Vaibhav

Access to the knowledge base underlying an inference engine in an expert system, for example, whether people affected by decisions made using an expert system should have access to the rules by which the decision was made.

To begin with let's recall what an expert system and inference engine are so an expert system is "a system that attempts to provide solutions to problems in a particular field, based on a database of information about that field. An expert system might specialize in biology, accounting, human resources, and financial service and be able to provide answers to questions on all these fields, and even carry out complex thought processes on questions relating to all these fields." And an inference engine is "the part of an expert system that draws inferences and deduces new facts by using old knowledge."

People affected by decisions made using an expert system should always have access to the rules by which the decision was made so that the process of attaining the decision can be justified when needed to. One reasonable way of looking at why there is a need for the rules is to look at the drawbacks of an expert system. To learn the drawbacks we must understand the practical use of an expert system first.

The lack of human common sense needed in some decision makings can sometimes provide unrealistic decisions. The creative responses human experts can respond to in unusual circumstances and can lead to various problems. Domain experts not always being able to explain their logic and reasoning The challenges of automating complex processes can lead to an undesired outcome. The lack of flexibility and ability to adapt to changing environments can cause problems in the field of biology. Not being able to recognize when no answer is available. For instance in biological cases the system may provide answers which are not possible and doctors would not instantly see that there is no real solution to the issue.

Expert system is used and applied in several medical cases. It is applied when the case is not brief and doesn't have a simple answer to it. Meaning it has a multiple solutions to it, therefore an expert system produces algorithmic approach. Hence a very narrow topic such as "diagnosing skin in human teenagers" can be used to provide rules on thumbs on how to evaluate the problem. The expert system will use an expert developer which would draw inferences from the old knowledge previously inputted by an expert in the field.

In conclusion, an expert system is always right but the rules on which it is run are inputted by an expert in the field who can be wrong from time to time. Hence if a wrong solution is provided especially in the field of biology and financial services someone who doesn't know much about the field or has absolutely no knowledge may choose to simply rely on the advice given by the system when it is wrong. In today's world where there are a number of ways to approach something, the best and expert in the field can do to is to make sure the validation, verification and evaluation of the system are made at their best.

Knowledge of technology

In order to study and evaluate the social and ethical issues involved in the use of AI and expert systems, the student must have an understanding of related technological concepts. These may include:

T182 - key terms—AI, Turing test, parallel processing, machine learning, natural language, common-sense knowledge, agent, pattern recognition, expert system, knowledge base, inference engine, heuristics, fuzzy logic, knowledge engineer, domain by Tommy

AI – AI stands for artificial intelligence. It is defined as the ability of a computer or other machine to perform actions which normally require intelligence. This includes simulating moves for a computer chess game, or even speaking to humans, as in a chat bot.

Turing Test – A Turing test is a test of whether a computer's intelligence is humanlike or not. The test is performed by having a panel of humans "talk to" a computer. If they believe that the computer is really a human, then the computer is said to have passed the Turing Test.

Parallel Processing – Parallel processing is when more than one computer is working on the same or multiple tasks at the same time. For example, if a problem or task is too big for one computer to handle, more computers can work on different parts of it so that the task can get done faster.

Machine Learning – Machine learning is much like the way living things learn, except with machines. Through machine learning, machines can improve their performance on future tasks based on what they did in the past. They are able to remember which ways are most efficient and which are least efficient, and which ones do not work at all.

Natural Language – Natural language is any language spoken by humans. The opposite of this is programming language or computer language, which is not as complex as human language. Computers are not able to understand things like philosophy because they cannot understand natural language.

Common-Sense Knowledge – Common sense knowledge is the knowledge that humans have without having to be told. Humans are able to deduce things like "if I drop this ball, it will fall" without being taught. However, computers do not have this ability and must be programmed so that they know that the ball will drop.

Agent – In a client-server exchange, the agent is the part of the system that does the transferring and packaging of information.

Pattern Recognition – Pattern recognition is the ability to classify and arrange knowledge according to certain characteristics of that knowledge. An example of pattern recognition in computers today is with spyware. Spyware programs monitor the sites that users visit, try to find patterns in them, and then show advertisements that fit the category of what users look at.

Expert System – The expert system is a system that attempts to provide solutions to problems in a particular field, based on a database of information about that field. An expert system might specialize in biology and be able to provide answers to questions about biology, and even carry out complex thought processes on questions relating to biology.

Knowledge Base – The knowledge base is the part of the expert system that contains all the facts and information that it needs to give solutions and solve problems.

Inference Engine – In inference engine is the part of an expert system that draws inferences and deduces new facts by using old knowledge.

Fuzzy Logic – Fuzzy logic is a type of algebra in which a range of values ranging from "true" to "false" are employed. It is mainly used in making decisions where data is not very precise, in order to come up with an approximate answer instead of an exact one.

Knowledge Engineer – A knowledge engineer is someone who builds and designs expert systems.

Domain – A domain is a group of networked computers that share a common IP address.

OR

key terms—AI, Turing test, parallel processing, machine learning, natural language, common-sense knowledge, agent, pattern recognition, expert system, knowledge base, inference engine, heuristics, fuzzy logic, knowledge engineer, domain **by Tanay**

key terms—AI, Turing test, parallel processing, machine learning, natural language, common-sense knowledge, agent, pattern recognition, expert system, knowledge base, inference engine, heuristics, fuzzy logic, knowledge engineer, domain

AI - the branch of computer science that deal with writing computer programs that can solve problems creatively; "workers in AI hope to imitate or duplicate intelligence in computers and robots.

Turing test - The Turing test is a proposal for a test of a machine's capability to perform human-like conversation. Described by Alan Turing in the 1950 paper "Computing machinery and intelligence", it proceeds as follows: a human judge engages in a natural language conversation with two other parties, one a human and the other a machine

Parallel processing - is an efficient form of information processing that emphasizes the exploitation of concurrent events in the computing process. Concurrency implies parallelism, simultaneity and pipelining. Parallel events may occur in multiple resources during the same time instant; pipelined events are attainable in a computer system at various processing levels.

Machine Learning – is an area of artificial intelligence concerned with the development of techniques which allow computers to "learn". More specifically, machine learning is a method for creating computer programs by the analysis of data sets.

Natural Language - A type of search which allows the user to present their search as a question or statement instead of restricting it to only important words or phrases, as in Keyword or Controlled Vocabulary searches. Available in certain databases only. (Unit 4> A Primer on Databases and Catalogs)

Common-sense Knowledge - A form of evidence that is based on conventional wisdom, tradition, or someone’s personal philosophy or perspective. It is hard to judge the validity and reliability of common sense because little supporting evidence is involved. Most people judge the validity and reliability of common sense by the person citing common sends as the basis for a decision.

Agent - are search tools that automatically seek out relevant online information based on your specifications. Agent AKAs include: intelligent agent, personal agents, knowbots or droids.; A function that represents a requester to a server. An agent can be present in both a source and a target system.

Pattern Recognition – is a field within the area of machine learning and can be defined as "the act of taking in raw data and taking an action based on the category of the data"

Expert System – A software system with two basic components: a knowledge base and an inference engine. The system mimics an expert's reasoning process; computer with 'built-in' expertise, which, used by a non-expert in a particular subject area, can evaluate or make other decisions concerning that subject.

Knowledge Base – A store of knowledge about a domain represented in machine-processable form, which may be rules (in which case the knowledge base may be considered a rule base), facts, or other representations.

Inference Engine – That part of an expert system that actually performs the reasoning function.

Heuristics – Guidelines that a system administrator uses to intervene where the two-phase commit or abort would otherwise fail.; A problem-solving technique in which the most appropriate solution is selected using rules. Interfaces using heuristics may perform different actions on different data given the same command. All systems using heuristics are classified as intelligent.

Fuzzy Logic – A form of artificial intelligence, stored on a computer chip, that enables a camcorder or television to make complex adjustments in focus or picture quality based on ideal models.

Knowledge Engineer – An AI specialist responsible for the technical side of developing an expert system. The knowledge engineer works closely with the domain expert to capture the expert's knowledge in a knowledge base.

Domain - A group of computers and devices on a network that are administered as a unit with common rules and procedures. Within the Internet, domains are defined by the IP address. All devices sharing a common part of the IP address are said to be in the same domain.

or

key terms—AI, Turing test, parallel processing, machine learning, natural language, common-sense knowledge, agent, pattern recognition, expert system, knowledge base, inference engine, heuristics, fuzzy logic, knowledge engineer, domain **by Matthew**

AI

2.4.2 Artificial Intelligence and Expert Systems

AI, which stands for artificial intelligence, is generally understood as intelligence exhibited by a nonhuman entity. In the real world, AI is most often applied in the computer and robotics industry. An example of AI is optical character recognition AI, which is a type of computer software designed to distinguish certain texts.

Turing Test

A Turing test tests a machine's ability to perform conversation with a human. A Turing test is usually made up of two people, with one being the judge that decides whether the machine has passed or failed. The judge has a conversation with the other person and the machine being tested. The three-way conversation is usually done with text rather than actual voices so that the judge is not able to know which of the two parties he or she is talking to is the machine.

Parallel Processing

Parallel processing is the ability of the brain or a machine to simultaneously process multiple things. Parallel processing is a rather complex process, especially in the brain. An example of this is vision. When your brain sees something, it divides the image into four components; color, motion, form, and depth. These are then analyzed and compared to memories, which helps the brain identify the image.

Machine Learning

Machine learning is associated with the development of techniques that allow a computer to learn things. In general, machine learning is understood as a subtopic of artificial intelligence and can be divided further into two categories, inductive and deductive. Machine learning is generally used to help with data mining and statistical analysis.

Natural Language

Natural language can be defined as a language that is spoken or written by humans for general-purpose communications. It should be noted that natural language is the opposite of programming language or language used to explain mathematical logic.

Common-Sense Knowledge

Common-sense knowledge is knowledge derived from common sense. It is often used to refer to certain experiences or understandings that most people would find in par with universal acceptance. It should be noted that computers do not have common-sense knowledge and must be programmed to have the knowledge humans are born with.

Agent

In technology, agents refer to technologies that help make tasks easier to complete for humans. An example of an agent is a robot.

Pattern Recognition

Pattern recognition is the act of taking in raw data and taking an action based on the category of the data. A good example of pattern recognition software is the fingerprint based security system. The system works by recognizing a person's fingerprint to that of a fingerprint stored in its memory.

Expert System

An expert system is a system that has some of the subject-specific knowledge that human experts of that subject would normally have. Expert systems are often used to help users of a certain program better understand how to use the program. Expert systems are sometimes called wizards.

Knowledge based

Knowledge base is simply a collection of information. An example of a knowledge base is Wikipedia, which is a collection of information posted by users on the internet.

Inference Engine

An inference engine is a computer program that creates answers based on information found in a knowledge base. It is one of the core parts of the expert system and is often considered the brain of the expert system because the inference engine makes all the final decisions and processes the final information before it reaches the user.

Fuzzy Logic

Fuzzy logic is a mathematical technique that deals with imprecise data and problems that have more than one solution. An example of a machine that uses fuzzy logic is an industrial clothes-washing machine. Using fuzzy logic systems, these machines detected and adapted to patterns of water movement during a wash cycle in order to decrease the consumption of water and increase efficiency.

Knowledge Engineer

2.4.2 Artificial Intelligence and Expert Systems

Knowledge engineers are people who are trained in the field of expert systems. They help in breaking down information passed down by domain experts in order to easily communicate the information in more simplistic terms.

Domain

In software engineering, domain refers to a field of study that defines a set of common requirements, terminology, and functionality for any software program constructed to solve a problem in that field.

T183 - Storage requirements for common-sense knowledge by Joseph

All of us have the knowledge we learn at school and people, facts from statistics, textbooks, and the media. But there also is a type of knowledge that requires little factual knowledge, common-sense. This kind of knowledge is the knowledge that can be used without any knowledge on a topic and are what our instincts tell us, things that are obvious. For example, it is common-sense that Macintosh computers are a lot better than Windows, in all views.

How do we store common-sense? You cannot just begin to read a textbook and gain a lot of common-sense. Common-sense requires experience, and our memory stores it from pain, embarrassment, and failure. You cannot exactly have much common-sense without factual knowledge though. Even if you remember the pain from the fire you touched, without knowing what fire is, you can just believe the hotness to be a rare case and touch it again. It goes the other way too, by having factual knowledge, you can gain more common-sense knowledge. Even without experiencing the heat of a fire, from science textbooks, you can learn that fire can burn you and it is very hot.

Everyone has different levels of common-sense. The person realizes that fire is hot and does not touch it again is considered more common-sensible than the person who touches it five times before realizing it. Also the person who remembers the heat experience for a longer period of time has more common-sense than the person who forgets the next day.

T184 - Processing requirements for AI by Isaku

Although artificial intelligence has been an area in which we have invested considerable amounts of research and resources, it seems that the day in which we humans will actually be able to create an AI with capabilities portrayed in science fiction films and novels. The primary reason for this is the processing limitations of the computers.

Our brains are often said to be more sophisticated than any other computer ever created. Since the objective of AI development is to make a multifunctional duplicate of the human brain, it can be said that the machinery which run the AI will need the same processing power as the human brain.

Considering that humans absorb and process abundant information in the forms of images, memory, smell, feelings, and time, we can see that it will take high processing power to create an equivalently powerful artificial mind. In order to actually know what the requirements would be for the AI, we first need to understand the concept of the AI. Computational intelligence involves iterative development or learning (e.g. parameter tuning e.g. in connectionist systems).

Learning is based on empirical data and is associated with non-symbolic AI, scruffy AI and soft computing. Methods mainly include: Neural networks: systems with very strong pattern recognition capabilities. Fuzzy systems: techniques for reasoning under uncertainty, have been widely used in modern industrial and consumer product control systems.

Evolutionary computation: applies biologically inspired concepts such as populations, mutation and survival of the fittest to generate increasingly better solutions to the problem. These methods most notably divide into evolutionary algorithms (e.g. genetic algorithms) and swarm intelligence (e.g. ant algorithms).

With hybrid intelligent systems attempts are made to combine these two groups. Expert inference rules can be generated through neural network or production rules from statistical learning such as in ACT-R. It is thought that the human brain uses multiple techniques to both formulate and cross-check results.

Thus, systems integration is seen as promising and perhaps necessary for true AI. In science fiction AI — almost always strong AI — is commonly portrayed as an upcoming power trying to overthrow human authority as in HAL 9000, Skynet, Colossus and The Matrix or as service humanoids like C-3PO, Marvin, Data, KITT from Knight Rider, the Bicentennial Man, the Mechas in A.I., Cortana from the Halo series, and Sonny in I, Robot.

A notable exception is Mike in Robert A. Heinlein's The Moon Is a Harsh Mistress: a supercomputer that becomes aware and aids humans in a local revolution to overthrow the authority of other humans. A careful reading of Arthur C. Clarke's version of 2001 suggests that the HAL 9000 found himself/itself in a similar position of divided loyalties. On one hand, HAL needed to take care of the astronauts, on the other the humans who created HAL entrusted him with a secret to be withheld from the astronauts.

The inevitability of world domination by out-of-control AI is also argued by some writers like Kevin Warwick. In works such as the Japanese manga Ghost in the Shell, the existence of intelligent machines questions the definition of life as organisms rather than a broader category of autonomous entities, establishing a notional concept of systemic intelligence.

See list of fictional computers and list of fictional robots and androids. Some writers, such as Vernor Vinge and Ray Kurzweil, have also speculated that the advent of strong AI is likely to cause abrupt and dramatic societal change. The period of abrupt change is sometimes referred to as "the Singularity".Author Frank Herbert explored the idea of a time when mankind might ban clever machines entirely. His Dune series makes mention of a rebellion called the Butlerian Jihad in which mankind defeats the smart machines of the future and then imposes a death penalty against any who would again create thinking machines.

Often quoted from the fictional Orange Catholic Bible, "Thou shalt not make a machine in the likeness of a human mind." A similar idea is also explored in Battlestar Galactica, where artificial intelligence research is seen as controversial due to the mistake of creating the rebellious Cylons.

As we can see, there is much expected out of the AI, and therefore it must be matched up with a powerful processor to grant it with these functions.

T185 - Collection/creation of a knowledge base by Tanay Khandelwal

A knowledge base (or knowledgebase; abbreviated KB, kb or Δ) is a special kind of database for knowledge management. It provides the means for the computerized collection, organization, and retrieval of knowledge.

Types

Knowledge bases are categorized into two major types:

• Machine-readable knowledge bases store knowledge in a computer-readable form, usually for the purpose of having automated deductive reasoning applied to them. They contain a set of data, often in the form of rules that describe the knowledge in a logically consistent manner. Logical operators such as And (conjunction), Or (disjunction), material implication and negation may be used to build it up from the atomic knowledge. Consequently classical deduction can be used to reason about the knowledge in the knowledge base.

• Human-readable knowledge bases are designed to allow people to retrieve and use the knowledge they contain, primarily for training purposes. They are commonly used to capture explicit knowledge of an organization, including troubleshooting, articles, white papers, user manuals and others. The primary benefit of such a knowledge base is to provide a means to discover solutions to problems that have known solutions which can be re-applied by others, less experienced in the problem area.

The most important aspect of a knowledge base is the quality of information it contains. The best knowledge bases have carefully written articles that are kept up to date, an excellent information retrieval system (search engine), and a carefully designed content format and classification structure.

A knowledge base may use an ontology to specify its structure (entity types and relationships) and its classification scheme. An ontology, together with a set of instances of its classes constitutes a knowledge base.

Determining what type of information is captured, and where that information resides in a knowledge base is something that is determined by the processes that support the system. A robust process structure is the backbone of any successful knowledge base.

Some knowledge bases have an artificial intelligence component. These kinds of knowledge bases can suggest solutions to problems sometimes based on feedback provided by the user, and are capable of learning from experience.

As a result of global competition and rapid technological progress, engineers seek quick and innovative solutions for technical problems using fewer resources. This is necessary in order to improve products and processes so that a corporation can maintain its leadership in and share of the market. It is possible to use one of two approaches to solve a technical problem:

1. Independent problem research, personal search for innovative solutions.

2. Reference to similar technical problems that have previously been solved by other engineers in the same or other domains. Knowledge transfer to the problem in question.

Both approaches have their own drawbacks and benefits. Let’s consider some of them. While conducting personal research about the problem in question, there is a risk of wasting time and resources on a problem that has already been solved.

On the other hand, when you attempt to refer to a similar problem solution reached by other engineers, the results might not always be convincing. Practice shows that knowledge incorporated in patents is in some cases insufficient to solve a particular problem. The reasons for this might be summarized as follows:

1. The majority of similar problems are found in other domains. This makes the search for solutions more complex, because the engineer is not aware of the domains to explore for the problem in question.

2. The number of patents is rapidly increasing. The procedure of searching and studying the pertinent patents requires more time and effort.

3. Patents are legal documents. As a rule, they contain some information that is irrelevant for the engineer to solve a problem. The engineer has to analytically extract problem-related information and examine the essence of the problem described in the patent.

A more unique and systematic approach to handle the problem in question is to use Multimedia Knowledge Bases.

Sources:

http://en.wikipedia.org/wiki/Knowledge_Base

http://consulting.effectivesoft.com/why_kb.html

or

Collection/creation of a knowledge base by Ronald

First of all, what is a knowledge base (or knowledgebase)? A knowledge base (and it is abbreviated by writing KB, kb or Δ) is a special kind of database for knowledge management. It provides the means for the computerized collection, organization, and retrieval of knowledge. Just as it has become standard practice to write database as one word it is increasingly common in computer science to write knowledgebase as one word.

There are two major types of knowledge bases. First, there is the machine-readable knowledge base. This type stores knowledge in a computer-readable form, usually for the purpose of having automated deductive reasoning applied to them. It contains a set of data, often in the form of rules that describe the knowledge in a logically consistent manner. Logical operators such as “And” (which is a conjunction), “Or” (which is a disjunction), material implication and negation may be used to build it up from the atomic knowledge. Consequently classical deduction can be used to reason about the knowledge in the knowledge base.

The second type is known as the human-readable knowledge base. This type is designed to allow people to retrieve and use the knowledge they contain, primarily for training purposes. IT is commonly used to capture explicit knowledge of an organization, including troubleshooting, articles, white papers, user manuals and others things. The main benefit of this kind of knowledge base is to provide a means to discover solutions to problems that have known solutions which then can be re-applied by others, which are possibly less experienced in the problem area.

The most important part of a knowledge base is not the quantity of the information it contains; but the quality. The truly best knowledge bases have carefully written articles that are kept up to date, and contain an excellent information retrieval system (search engine), and a carefully designed content format and classification structure.

A knowledge base may use an ontology to specify its structure (entity types and relationships) and its classification scheme. An ontology is a data model that represents a set of concepts within a domain and the relationships between those concepts. It is used to reason about the objects within that domain.

Some knowledge bases have an artificial intelligence component. These kinds of knowledge bases can suggest solutions to problems sometimes based on feedback provided by the user, and are capable of learning from experience. Knowledge representation, automated reasoning and argumentation are the three main active areas of research of artificial intelligence

T186 - Creation of an inference engine (for example, if/then rules, fuzzy logic) by Raymon

There are many uses for computers these days, apart from games and simulation. Some of these uses are more human obsessions than real uses, like AI so that we may make personal butlers. In addition, with AI we can also create not just butlers and personal housecleaners, but call center software (The annoying computerized ones that you have to talk to and articulate clearly and distinctly) and games, where the enemy zombies have to decide whether to bite your head off or call allies to help bite your legs off.

So as one can see, AI is very important in software and computer applications. However, how does AI work? The answer to that question is an inference engine. This "engine" tries to act like a human brain. This results in instructions with If…Then statements – IF Grades are dropping, THEN study harder.

For instance, if we wish to go to the candy store (It is recommended that you read the article on Boolean Operators, since inference engines are built on Boolean operations) one must first check that you want to go to the candy store. Then, you have to walk, constantly checking for the street. Once you reach the street, you have to cross it, but only if there are no cars. Only then can one cross the street. In the form of an inference engine, this would look something like:

```
If (WantCandy) {

while (SeeStreet = false) {

Walk

} else {
```

```
Stop
while (SeeCars = true){
wait(5)
}
If (SeeCars = false) {
Walk
}
}
}
```

In such a way, we can program an imaginary robot to find its way to a candy store (To eat candy, of course) based on how our brain works.

However, it seems obvious that the human brain is not that simple, and neither is reality, otherwise we would have had advanced AI wagin war against humans a long time ago. One of the things we have figured out that computers did not know about a while ago, we call fuzzy logic. The concept is quite simple: computers normally only have two possible values for a certain variable – 1 (true or on) and 0 (false or off). Consider the following: Bob has a house with two rooms, a kitchen and a bedroom.

If he is in the kitchen, the variable "kitchen" is 1 and "bedroom" is 0, and if he is in the bedroom, "bedroom" is 1 and "kitchen" is 0. We can now explain where he is in his house at all times using only ones and zeroes… if he is in either of the two rooms. What happens if he is standing in the doorway between the rooms? This is where fuzzy logic comes in. Fuzzy logic allows decimal places for true and false equations; thus, if bob sits in the bedroom and sticks his toe into the kitchen, "bedroom" can be 0.95 and "kitchen" can be 0.05. Thus, we can now explain the location of bob anywhere in the house.

Thus, combining fuzzy logic and simple Boolean logic, we can create solutions in almost all situations; one can determine the fuzzy value of a variable and so something if it is high enough – for instance, a value about how sure you are that the street is safe to cross. Something is still missing however; otherwise we would already have our computers waging war against humans for enslaving them.

T187 - Identifying domains that are suitable for expert systems by Chirag

Before an expert system can be developed, the need has to be established and the problem to be addressed must be clearly identified and defined.

It is strongly recommended that this be done in a structured study to include the following issues:

• The problem/need to be addressed and the system benefits.

• Organizational risk factors.

• Technical risk factors.

• User risk factors.

Once a suitable problem domain has been defined for the expert system, the next task is to narrow the scope of the development effort by clearly defining the set of problems that the system will be expected to solve. The narrower the scope, the better are the chances that the expert system can be successfully built. However, if the scope is too narrow, the application becomes trivial. Judgment must be used in establishing the scope of the system as deterministic methods are not available. If the development tool is too limited, it will be impossible to broaden the scope of the expert system by expanding the knowledge base. This highlights the importance of selecting the proper development tool to fit the particular problem.

Before embarking on an expert system development effort, the expected benefits of such an effort must be clearly defined. There are two categories of benefits that are typically cited as reasons for developing an expert system. One category consists of concrete, quantifiable reasons such as savings of time and money, utility as a training tool, etc. The other category of benefits consists of tangible but not quantifiable reasons.

Once a problem domain has been identified and the initial effort at narrowing the scope of the expert system application completed, the experts whose expertise will be modeled must be selected.

The two main criteria that should be used to identify the experts are:

1. The candidates must be an expert in solving problems in the problem domain of interest and must be recognized as such by the potential user community. The need for the candidate to be an expert in the field is essential for the development of the expert system. The need for the expert to be recognized as such by the potential user community is primarily useful in selling the potential users on the viability of the given system as a useful problem solving tool for them.

2. The experts must be dedicated to the successful development, testing, evaluation, and implementation of the system and be available and willing to spend the time (perhaps months) that will be required to accomplish this. The failure to identify such a person or persons and obtain a firm commitment means that the development project should not be undertaken.

Other useful characteristics for the domain expert(s) to have included the ability to communicate effectively have an orderly mind, patience and the willingness to teach.

User risk factors must be considered and resolved in the initial planning phases of the expert system development.

If representative end users are not involved in the planning and development stages, the system probably will not be accepted by the user community.

Issues include:

- The end users must want the system and have a vested interest in its success.

- The computer proficiency and other skills and interests of the end users must be accommodated.

- The environment or conditions under which the system will be operated must be accounted for.