A

Seminar report

On

Artificial Intelligence

Submitted in partial fulfillment of the requirement for the award of degree
Of CSE

SUBMITTED TO:

www.studymafia.org

SUBMITTED BY:

www.studymafia.org

Acknowledgement

I would like to thank respected Mr…….. and Mr. ……..for giving me such a wonderful opportunity to expand my knowledge for my own branch and giving me guidelines to present a seminar report. It helped me a lot to realize of what we study for.

Secondly, I would like to thank my parents who patiently helped me as i went through my work and helped to modify and eliminate some of the irrelevant or un-necessary stuffs.

Thirdly, I would like to thank my friends who helped me to make my work more organized and well-stacked till the end.

Next, I would thank Microsoft for developing such a wonderful tool like MS Word. It helped my work a lot to remain error-free.

Last but clearly not the least, I would thank The Almighty for giving me strength to complete my report on time.

Preface

I have made this report file on the topic **Artificial Intelligence**; I have tried my best to elucidate all the relevant detail to the topic to be included in the report. While in the beginning I have tried to give a general view about this topic.

My efforts and wholehearted co-corporation of each and everyone has ended on a successful note. I express my sincere gratitude to …………..who assisting me throughout the preparation of this topic. I thank him for providing me the reinforcement, confidence and most importantly the track for the topic whenever I needed it.

Table of contents:

Introduction

Thinking machines and artificial beings appear in Greek myths, such as Talos of Crete, the golden robots of Hephaestus and Pygmalion's Galatea. Human likenesses believed to have intelligence were built in every major civilization: animated statues were seen in Egypt and Greece and humanoid automatons were built by Yan Shi, Hero of Alexandria, Al-Jazari and Wolfgang von Kempelen. It was also widely believed that artificial beings had been created by Jābir ibn Hayyān, Judah Loew and Paracelsus. By the 19th and 20th centuries, artificial beings had become a common feature in fiction, as in Mary Shelley's *Frankenstein* or Karel Čapek's *R.U.R. (Rossum's Universal Robots)*. Pamela McCorduck argues that all of these are examples of an ancient urge, as she describes it, "to forge the gods". Stories of these creatures and their fates discuss many of the same hopes, fears and ethical concerns that are presented by artificial intelligence.

Mechanical or "formal" reasoning has been developed by philosophers and mathematicians since antiquity. The study of logic led directly to the invention of the programmable digital electronic computer, based on the work of mathematician Alan Turing and others. Turing's theory of computation suggested that a machine, by shuffling symbols as simple as "0" and "1", could simulate any conceivable act of mathematical deduction. This, along with recent discoveries in neurology, information theory and cybernetics, inspired a small group of researchers to begin to seriously consider the possibility of building an electronic brain.

In the early 1980s, AI research was revived by the commercial success of expert systems, a form of AI program that simulated the knowledge and analytical skills of one or more human experts. By 1985 the market for AI had reached over a billion dollars. At the same time, Japan's fifth generation computer project inspired the U.S and British governments to restore funding for academic research in the field.[36] However, beginning with the collapse of the Lisp Machine market in 1987, AI once again fell into disrepute, and a second, longer lasting AI winter began.

In the 1990s and early 21st century, AI achieved its greatest successes, albeit somewhat behind the scenes. Artificial intelligence is used for logistics, data mining, medical diagnosis and many other areas throughout the technology industry.] The success was due to several factors: the incredible power of computers today (see Moore's law), a greater emphasis on solving specific subproblems, the creation of new ties between AI and other fields working on similar problems, and above all a new commitment by researchers to solid mathematical methods and rigorous scientific standards.

Approaches

There is no established unifying theory or paradigm that guides AI research. Researchers disagree about many issues. A few of the most long standing questions that have remained unanswered are these: should artificial intelligence simulate natural intelligence, by studying psychology or neurology? Or is human biology as irrelevant to AI research as bird biology is to aeronautical engineering? Can intelligent behavior be described using simple, elegant principles (such as logic or optimization)? Or does it necessarily require solving a large number of completely unrelated problems? Can intelligence be reproduced using high-level symbols, similar to words and ideas? Or does it require "sub-symbolic" processing?

Cybernetics and brain simulation

There is no consensus on how closely the brain should be simulated.

In the 1940s and 1950s, a number of researchers explored the connection between neurology, information theory, and cybernetics. Some of them built machines that used electronic networks to exhibit rudimentary intelligence, such as W. Grey Walter's turtles and the Johns Hopkins Beast. Many of these researchers gathered for meetings of the Teleological Society at Princeton University and the Ratio Club in England. By 1960, this approach was largely abandoned, although elements of it would be revived in the 1980s.

Symbolic

When access to digital computers became possible in the middle 1950s, AI research began to explore the possibility that human intelligence could be reduced to symbol manipulation. The research was centered in three institutions: CMU, Stanford and MIT, and each one developed its own style of research. John Haugeland named these approaches to AI "good old fashioned AI" or "GOFAI".

Cognitive simulation

Economist Herbert Simon and Allen Newell studied human problem solving skills and attempted to formalize them, and their work laid the foundations of the field of artificial intelligence, as well as cognitive science, operations research and management science. Their research team used the results of psychological experiments to develop programs that simulated the techniques that people used to solve problems. This tradition, centered at Carnegie Mellon University would eventually culminate in the development of the Soar architecture in the middle 80s.

Logic based

Unlike Newell and Simon, John McCarthy felt that machines did not need to simulate human thought, but should instead try to find the essence of abstract reasoning and problem solving, regardless of whether people used the same algorithms. His laboratory at Stanford (SAIL) focused on using formal logic to solve a wide variety of problems, including knowledge representation, planning and learning. Logic was also focus of the work at the University of Edinburgh and elsewhere in Europe which led to the development of the programming language Prolog and the science of logic programming.

"Anti-logic" or "scruffy"

Researchers at MIT (such as Marvin Minsky and Seymour Papert) found that solving difficult problems in vision and natural language processing required ad-hoc solutions – they argued that there was no simple and general principle (like logic) that would capture all the aspects of intelligent behavior. Roger Schank described their "anti-logic" approaches as "scruffy" (as opposed to the "neat" paradigms at CMU and Stanford). Commonsense knowledge bases (such as Doug Lenat's Cyc) are an example of "scruffy" AI, since they must be built by hand, one complicated concept at a time.

Knowledge based

When computers with large memories became available around 1970, researchers from all three traditions began to build knowledge into AI applications. This "knowledge revolution" led to the development and deployment of expert systems (introduced by Edward Feigenbaum), the first truly successful form of AI software. The knowledge revolution was also driven by the realization that enormous amounts of knowledge would be required by many simple AI applications.

Sub-symbolic

During the 1960s, symbolic approaches had achieved great success at simulating high-level thinking in small demonstration programs. Approaches based on cybernetics or neural networks were abandoned or pushed into the background. By the 1980s, however, progress in symbolic AI seemed to stall and many believed that symbolic systems would never be able to imitate all the processes of human cognition, especially perception, robotics, learning and pattern recognition. A number of researchers began to look into "sub-symbolic" approaches to specific AI problems.

Bottom-up, embodied, situated, behavior-based or nouvelle AI

Researchers from the related field of robotics, such as Rodney Brooks, rejected symbolic AI and focused on the basic engineering problems that would allow robots to move and survive. Their work revived the non-symbolic viewpoint of the early cybernetics researchers of the 50s and reintroduced the use of control theory in AI. This coincided with the development of the embodied mind thesis in the related field of cognitive science: the idea that aspects of the body (such as movement, perception and visualization) are required for higher intelligence.

Computational Intelligence

Interest in neural networks and "connectionism" was revived by David Rumelhart and others in the middle 1980s. These and other sub-symbolic approaches, such as fuzzy systems and evolutionary computation, are now studied collectively by the emerging discipline of computational intelligence.

Statistical

In the 1990s, AI researchers developed sophisticated mathematical tools to solve specific subproblems. These tools are truly scientific, in the sense that their results are both measurable and verifiable, and they have been responsible for many of AI's recent successes. The shared mathematical language has also permitted a high level of collaboration with more established fields (like mathematics, economics or operations research). Stuart Russell and Peter Norvig describe this movement as nothing less than a "revolution" and "the victory of the neats."

Integrating the approaches

Intelligent agent paradigm

An intelligent agent is a system that perceives its environment and takes actions which maximizes its chances of success. The simplest intelligent agents are programs that solve specific problems. The most complicated intelligent agents are rational, thinking humans. The paradigm gives researchers license to study isolated problems and find solutions that are both verifiable and useful, without agreeing on one single approach. An agent that solves a specific problem can use any approach that works — some agents are symbolic and logical, some are sub-symbolic neural networks and others may use new approaches. The paradigm also gives researchers a common language to communicate with other fields—such as decision theory and economics—that also use concepts of abstract agents. The intelligent agent paradigm became widely accepted during the 1990s.

Agent architectures and cognitive architectures

Researchers have designed systems to build intelligent systems out of interacting intelligent agents in a multi-agent system. A system with both symbolic and sub-symbolic components is a hybrid intelligent system, and the study of such systems is artificial intelligence systems integration. A hierarchical control system provides a bridge between sub-symbolic AI at its lowest, reactive levels and traditional symbolic AI at its highest levels, where relaxed time constraints permit planning and world modelling. Rodney Brooks' subsumption architecture was an early proposal for such a hierarchical system.

Tools

In the course of 50 years of research, AI has developed a large number of tools to solve the most difficult problems in computer science. A few of the most general of these methods are discussed below.

Search and optimization

Many problems in AI can be solved in theory by intelligently searching through many possible solutions: Reasoning can be reduced to performing a search. For example, logical proof can be viewed as searching for a path that leads from premises to conclusions, where each step is the application of an inference rule. Planning algorithms search through trees of goals and subgoals, attempting to find a path to a target goal, a process called means-ends analysis. Robotics algorithms for moving limbs and grasping objects use local searches in configuration space. Many learning algorithms use search algorithms based on optimization.

Simple exhaustive searches are rarely sufficient for most real world problems: the search space (the number of places to search) quickly grows to astronomical numbers. The result is a search that is too slow or never completes. The solution, for many problems, is to use "heuristics" or "rules of thumb" that eliminate choices that are unlikely to lead to the goal (called "pruning the search tree"). Heuristics supply the program with a "best guess" for what path the solution lies on.

A very different kind of search came to prominence in the 1990s, based on the mathematical theory of optimization. For many problems, it is possible to begin the search with some form of a guess and then refine the guess incrementally until no more refinements can be made. These algorithms can be visualized as blind hill climbing: we begin the search at a random point on the landscape, and then, by jumps or steps, we keep moving our guess uphill, until we reach the top. Other optimization algorithms are simulated annealing, beam search and random optimization.

Evolutionary computation uses a form of optimization search. For example, they may begin with a population of organisms (the guesses) and then allow them to mutate and recombine, selecting only the fittest to survive each generation (refining the guesses). Forms of evolutionary computation include swarm intelligence algorithms (such as ant colony or particle swarm optimization) and evolutionary algorithms (such as genetic algorithms and genetic programming).

Logic

Logic is used for knowledge representation and problem solving, but it can be applied to other problems as well. For example, the satplan algorithm uses logic for planning and inductive logic programming is a method for learning.

Several different forms of logic are used in AI research. Propositional or sentential logic is the logic of statements which can be true or false. First-order logic also allows the use of quantifiers and predicates, and can express facts about objects, their properties, and their relations with each other. Fuzzy logic, is a version of first-order logic which allows the truth of a statement to be represented as a value between 0 and 1, rather than simply True (1) or False (0). Fuzzy systems can be used for uncertain reasoning and have been widely used in modern industrial and consumer product control systems. Subjective logic models uncertainty in a different and more explicit manner than fuzzy-logic: a given binomial opinion satisfies belief + disbelief + uncertainty = 1 within a Beta distribution. By this method, ignorance can be distinguished from probabilistic statements that an agent makes with high confidence. Default logics, non-monotonic logics and circumscription are forms of logic designed to help with default reasoning and the qualification problem. Several extensions of logic have been designed to handle specific domains of knowledge, such as: description logics; situation calculus, event calculus and fluent calculus (for representing events and time); causal calculus; belief calculus; and modal logics.

Probabilistic methods for uncertain reasoning

Many problems in AI (in reasoning, planning, learning, perception and robotics) require the agent to operate with incomplete or uncertain information. AI researchers have devised a number of powerful tools to solve these problems using methods from probability theory and economics.

Bayesian networks are a very general tool that can be used for a large number of problems: reasoning (using the Bayesian inference algorithm), learning (using the expectation-maximization algorithm), planning (using decision networks) and perception (using dynamic Bayesian networks). Probabilistic algorithms can also be used for filtering, prediction, smoothing and finding explanations for streams of data, helping perception systems to analyze processes that occur over time (e.g., hidden Markov models or Kalman filters).

A key concept from the science of economics is "utility": a measure of how valuable something is to an intelligent agent. Precise mathematical tools have been developed that analyze how an agent can make choices and plan, using decision theory, decision analysis, information value theory. These tools include models such as Markov decision processes, dynamic decision networks, game theory and mechanism design.

Classifiers and statistical learning methods

The simplest AI applications can be divided into two types: classifiers ("if shiny then diamond") and controllers ("if shiny then pick up"). Controllers do however also classify conditions before inferring actions, and therefore classification forms a central part of many AI systems. Classifiers are functions that use pattern matching to determine a closest match. They can be tuned according to examples, making them very attractive for use in AI. These examples are known as observations or patterns. In supervised learning, each pattern belongs to a certain predefined class. A class can be seen as a decision that has to be made. All the observations combined with their class labels are known as a data set. When a new observation is received, that observation is classified based on previous experience.

A classifier can be trained in various ways; there are many statistical and machine learning approaches. The most widely used classifiers are the neural network, kernel methods such as the support vector machine, k-nearest neighbor algorithm, Gaussian mixture model, naive Bayes classifier, and decision tree. The performance of these classifiers have been compared over a wide range of tasks. Classifier performance depends greatly on the characteristics of the data to be classified. There is no single classifier that works best on all given problems; this is also referred to as the "no free lunch" theorem. Determining a suitable classifier for a given problem is still more an art than science.

Neural networks

A neural network is an interconnected group of nodes, akin to the vast network of neurons in the human brain.

The study of artificial neural networks began in the decade before the field AI research was founded, in the work of Walter Pitts and Warren McCullough. Other important early researchers were Frank Rosenblatt, who invented the perceptron and Paul Werbos who developed the backpropagation algorithm.

The main categories of networks are acyclic or feedforward neural networks (where the signal passes in only one direction) and recurrent neural networks (which allow feedback). Among the most popular feedforward networks are perceptrons, multi-layer perceptrons and radial basis networks. Among recurrent networks, the most famous is the Hopfield net, a form of attractor network, which was first described by John Hopfield in 1982. Neural networks can be applied to the problem of intelligent control (for robotics) or learning, using such techniques as Hebbian learning and competitive learning.

Jeff Hawkins argues that research in neural networks has stalled because it has failed to model the essential properties of the neocortex, and has suggested a model (Hierarchical Temporal Memory) that is loosely based on neurological research.

Control theory

Control theory, the grandchild of cybernetics, has many important applications, especially in robotics.

Languages

AI researchers have developed several specialized languages for AI research, including Lisp and Prolog.

Evaluating progress

In 1950, Alan Turing proposed a general procedure to test the intelligence of an agent now known as the Turing test. This procedure allows almost all the major problems of artificial intelligence to be tested. However, it is a very difficult challenge and at present all agents fail.

Artificial intelligence can also be evaluated on specific problems such as small problems in chemistry, hand-writing recognition and game-playing. Such tests have been termed subject matter expert Turing tests. Smaller problems provide more achievable goals and there are an ever-increasing number of positive results.

The broad classes of outcome for an AI test are:

- Optimal: it is not possible to perform better
- Strong super-human: performs better than all humans
- Super-human: performs better than most humans
- Sub-human: performs worse than most humans

For example, performance at draughts is optimal, performance at chess is super-human and nearing strong super-human, and performance at many everyday tasks performed by humans is sub-human.

A quite different approach measures machine intelligence through tests which are developed from *mathematical* definitions of intelligence. Examples of these kinds of tests start in the late nineties devising intelligence tests using notions from Kolmogorov Complexity and data compression. Two major advantages of mathematical definitions are their applicability to nonhuman intelligences and their absence of a requirement for human testers.

Applications of AI

Game playing

You can buy machines that can play master level chess for a few hundred dollars. There is some AI in them, but they play well against people mainly through brute force computation--looking at hundreds of thousands of positions. To beat a world champion by brute force and known reliable heuristics requires being able to look at 200 million positions per second.

Speech recognition

In the 1990s, computer speech recognition reached a practical level for limited purposes. Thus United Airlines has replaced its keyboard tree for flight information by a system using speech recognition of flight numbers and city names. It is quite convenient. On the the other hand, while it is possible to instruct some computers using speech, most users have gone back to the keyboard and the mouse as still more convenient.

Understanding natural language

Just getting a sequence of words into a computer is not enough. Parsing sentences is not enough either. The computer has to be provided with an understanding of the domain the text is about, and this is presently possible only for very limited domains.

Computer vision

The world is composed of three-dimensional objects, but the inputs to the human eye and computers' TV cameras are two dimensional. Some useful programs can work solely in two dimensions, but full computer vision requires partial three-dimensional information that is not just a set of two-dimensional views. At present there are only limited ways of representing three-dimensional information directly, and they are not as good as what humans evidently use.

Expert systems

A ``knowledge engineer'' interviews experts in a certain domain and tries to embody their knowledge in a computer program for carrying out some task. How well this works depends on whether the intellectual mechanisms required for the task are within the present state of AI. When this turned out not to be so, there were many disappointing results. One of the first expert systems was MYCIN in 1974, which diagnosed bacterial infections of the blood and suggested treatments. It did better than medical students or practicing doctors, provided its limitations were observed. Namely, its ontology included bacteria, symptoms, and treatments and did not include patients, doctors, hospitals, death, recovery, and events occurring in time. Its interactions depended on a single patient being considered. Since the experts consulted by the knowledge engineers knew about patients, doctors, death, recovery, etc., it is clear that the knowledge engineers forced what the experts told them into a predetermined framework. In the present state of AI, this has to be true. The usefulness of current expert systems depends on their users having common sense.

Heuristic classification

One of the most feasible kinds of expert system given the present knowledge of AI is to put some information in one of a fixed set of categories using several sources of information. An example is advising whether to accept a proposed credit card purchase. Information is available about the owner of the credit card, his record of payment and also about the item he is buying and about the establishment from which he is buying it (e.g., about whether there have been previous credit card frauds at this establishment).

Advantages for Artificial Intelligence (AI)

- **Jobs** – depending on the level and type of intelligence these machines receive in the future, it will obviously have an effect on the type of work they can do, and how well they can do it (they can become more efficient). As the level of AI increases so will their competency to deal with difficult, complex even dangerous tasks that are currently done by humans, a form of applied artificial intelligence.

- **They don't stop** – as they are machines there is no need for sleep, they don't get ill , there is no need for breaks or Facebook, they are able to go, go, go! There obviously may be the need for them to be charged or refueled, however the point is, they are definitely going to get a lot more work done than we can. Take the Finance industry for example, there are constant stories arising of artificial intelligence in finance and that stock traders are soon to be a thing of the past.

- **No risk of harm** – when we are exploring new undiscovered land or even planets, when a machine gets broken or dies, there is no harm done as they don't feel, they don't have emotions. Where as going on the same type of expeditions a machine does, may simply not be possible or they are exposing themselves to high risk situations.

- **Act as aids –** they can act as 24/7 aids to children with disabilities or the elderly, they could even act as a source for learning and teaching. They could even be part of security alerting you to possible fires that you are in threat of, or fending off crime.

- **Their function is almost limitless** – as the machines will be able to do everything (but just better) essentially their use, pretty much doesn't have any boundaries. They will make fewer mistakes, they are emotionless, they are more efficient, they are basically giving us more free time to do as we please.

Disadvantages for Artificial Intelligence (AI)

- **Over reliance on AI** – as you may have seen in many films such as The Matrix, iRobot or even kids films such as WALL.E, if we rely on machines to do almost everything for us we become very dependent, so much so they have the potential to ruin our lives if something were to go wrong. Although the films are essentially just fiction, it wouldn't be too smart not to have some sort of back up plan to potential issues on our part.

- **Human Feel** – as they are are machines they obviously can't provide you with that 'human touch and quality', the feeling of a togetherness and emotional understanding, that machines will lack the ability to sympathise and empathise with your situations, and may act irrationally as a consequence.

- **Inferior** – as machines will be able to perform almost every task better than us in practically all respects, they will take up many of our jobs, which will then result in masses of people who are then jobless and as a result feel essentially useless. This could then lead us to issues of mental illness and obesity problems etc.

- **Misuse** – there is no doubt that this level of technology in the wrong hands can cause mass destruction, where robot armies could be formed, or they could perhaps malfunction or be corrupted which then we could be facing a similar scene to that of terminator (hey, you never know).

- **Ethically Wrong? –** People say that the gift of intuition and intelligence was God's gift to mankind, and so to replicate that would be then to kind of 'play God'. Therefore not right to even attempt to clone our intelligence.

Conclusion

AI is a common topic in both science fiction and projections about the future of technology and society. The existence of an artificial intelligence that rivals human intelligence raises difficult ethical issues, and the potential power of the technology inspires both hopes and fears.

In fiction, AI has appeared fulfilling many roles, including a servant (R2D2 in *Star Wars*), a law enforcer (K.I.T.T. "Knight Rider"), a comrade (Lt. Commander Data in *Star Trek: The Next Generation*), a conqueror/overlord (*The Matrix*), a dictator (*With Folded Hands*), an assassin (*Terminator*), a sentient race (*Battlestar Galactica*/Transformers), an extension to human abilities (*Ghost in the Shell*) and the savior of the human race (R. Daneel Olivaw in the *Foundation Series*).

Mary Shelley's *Frankenstein* considers a key issue in the ethics of artificial intelligence: if a machine can be created that has intelligence, could it also *feel*? If it can feel, does it have the same rights as a human? The idea also appears in modern science fiction, including the films *I Robot*, *Blade Runner* and *A.I.: Artificial Intelligence*, in which humanoid machines have the ability to feel human emotions. This issue, now known as "robot rights", is currently being considered by, for example, California's Institute for the Future, although many critics believe that the discussion is premature.

The impact of AI on society is a serious area of study for futurists. Academic sources have considered such consequences as a decreased demand for human labor, the enhancement of human ability or experience, and a need for redefinition of human identity and basic values. Andrew Kennedy, in his musing on the evolution of the human personality, considered that artificial intelligences or 'new minds' are likely to have severe personality disorders, and identifies four particular types that are likely to arise: the autistic, the collector, the ecstatic, and the victim. He suggests that they will need humans because of our superior understanding of personality and the role of the unconscious.

Several futurists argue that artificial intelligence will transcend the limits of progress. Ray Kurzweil has used Moore's law (which describes the relentless exponential improvement in digital technology) to calculate that desktop computers will have the same processing power as human brains by the year 2029. He also predicts that by 2045 artificial intelligence will reach a point where it is able to improve *itself* at a rate that far exceeds anything conceivable in the past, a scenario that science fiction writer Vernor Vinge named the "technological singularity".

Robot designer Hans Moravec, cyberneticist Kevin Warwick and inventor Ray Kurzweil have predicted that humans and machines will merge in the future into cyborgs that are more capable and powerful than either. This idea, called transhumanism, which has roots in Aldous Huxley and Robert Ettinger, has been illustrated in fiction as well, for example in the manga *Ghost in the Shell* and the science-fiction series *Dune*.

Edward Fredkin argues that "artificial intelligence is the next stage in evolution," an idea first proposed by Samuel Butler's "Darwin among the Machines" (1863), and expanded upon by George Dyson in his book of the same name in 1998.

REFERENCES

- www.google.com
- www.wikipedia.com
- www.studymafia.org