Artificial Intelligence

Chandan Gajbhiye
Department of Mechanical Engineering
G.H. Raisoni College of Engineering, Nagpur, India

Abstract— In Computer Science, artificial intelligence (AI), sometimes called machine intelligence, is intelligence demonstrated by Machine, in contrast to the natural intelligence displayed by humans. Leading AI textbooks define the field as the study of intelligent agent any device that perceives its environment and takes actions that maximize its chance of successfully achieving its goals. Artificial intelligence (AI), deep learning, machine learning and neural networks represent incredibly exciting and powerful machine learning-based techniques used to solve many real-world problems. For a primer on machine learning, you may want to read this five-part series that I wrote. While human-like deductive reasoning, inference, and decision-making by a computer are still a long time away, there have been remarkable gains in the application of AI techniques and associated algorithms.

Keywords: Artificial Intelligence, Artificial neural networks (ANNs)

I. Introduction

Computer science defines AI research as the study of intelligent agent any device that perceives its environment and takes actions that maximize its chance of successfully achieving its goals. A more elaborate definition characterizes AI as a system s ability to correctly interpret external data, to learn from such data, and to use those learning's to achieve specific goals and tasks through flexible adaptation.

Artificial intelligence (AI) is an area of computer science that emphasizes the creation of intelligent machines that work and reacts like humans. Some of the activities computers with artificial intelligence are designed for include Speech recognition, Learning, Planning, Problem solving. In this topic we shall discus the following subjects; deep learning, Machine learning, Computer Programming, Medical field. Deep Learning has enabled many practical applications of Machine Learning and by extension the overall field of AI. Deep Learning breaks down tasks in ways that makes all kinds of machine assists seem possible, even likely. Driverless cars, better preventive healthcare, even better movie recommendations, are all here today or on the horizon. AIis the present and the future. With Deep Learning's help, AI may even get to that science fiction state we've so long imagined.

Machine learning at its most basic is the practice of using algorithms to parse data, learn from it, and then decide or prediction about something in the world. So rather than hand-coding software routines with a specific set of Instructions to accomplish a task, the machine are "trained" using large amounts of data and algorithms that give it the ability to learn how to perform the task.

II. AI

Artificial intelligence (AI), the ability of a digital computer or computer- controlled robot to perform tasks commonly associated with intelligent beings. The term is frequently applied to the project of developing systems endowed with the intellectual processes characteristic of humans, such as the ability to reason, discover meaning, generalize, or learn from past experience. Since the development of the digital computer in the 1940s, it has been demonstrated that computers can be programmed to carry out very complex tasks as, for example, discovering proofs for mathematical theorems or playing chess with great proficiency. Still, despite continuing advances in computer processing speed and memory capacity, there are as yet no programs that can match human flexibility over wider domains or in tasks requiring much everyday knowledge. On the other hand, some programs have attained the performance levels of human experts and professionals in performing certain specific tasks, so that artificial intelligence in this limited sense is found in applications as diverse as medical diagnosis, computer search engine, and voice or handwriting recognition.

A. *What is intelligence?*

All but the simplest human behavior is ascribed to intelligence, while even the most complicated insect's behavior is never taken as an indication of intelligence. What is the difference? Consider the behavior of the digger wasp, Sphinx ichneumons. When the female wasp returns to her burrow with food, she first deposits it on the threshold, checks for intruders inside her burrow, and only then, if the coast is clear, carries her food inside. The real nature of the wasp's instinctual behavior is revealed if the food is moved a few inches away from the entrance to her burrow while she is inside: on emerging, she will repeat the whole procedure as often as the food is displaced. Intelligence conspicuously absent in the case of Sphinx must include the ability to adapt to new circumstances. Psychologist generally does not characterize human Intellectual by just one trait but by the combination of many diverse abilities. Research in AI has focused chiefly on the following components of intelligence: learning, reasoning, problem solving, perception, and using language.

Artificial intelligence is a branch of computer science that aims to create intelligent machines. It has become an essential part of the technology industry. Research associated with artificial intelligence is highly technical and specialized. The core problems of artificial intelligence include programming computers for certain traits such as: Knowledge, Reasoning, and Problem solving, Perception, Learning, Planning, and Ability to manipulate and move objects. Knowledge engineering is a core part of AI research.

Machines can often act and react like humans only if they have abundant information relating to the world. Artificial intelligence must have access to objects, categories, properties and relations between all of them to implement knowledge engineering. Initiating common sense, reasoning and problem-solving power in machines is a difficult and tedious approach. Machine learning is another core part of AI. Learning without any kind of supervision requires an ability to identify patterns in streams of inputs, whereas learning with adequate supervision involves classification and numerical regressions.

Classification determines the category an object belongs to and regression deals with obtaining a set of numerical input or output examples, thereby discovering functions enabling the generation of suitable outputs from respective inputs. Mathematical analysis of machine learning algorithms and their performance is a well-defined branch of theoretical computer science often referred to as computational learning theory. Machine perception deals with the capability to use sensory inputs to deduce the different aspects of the world, while computer vision is the power to analyses visual inputs with a few sub- problems such as facial, object and gesture recognition.

Robotics is also a major field related to AI. Robots require intelligence to handle tasks such as object manipulation and navigation, along with sub- problems of localization, motion planning and mapping.

Deep learning while flashy is just a term to describe certain types of neural networks and related algorithms that consume often very raw input data. They process this data through many layers of nonlinear transformations of the input data in order to calculate a target output. Unsupervised feature extraction is also an area where deep learning excels. Feature extraction is when an algorithm can automatically derive or construct meaningful features of the data to be used for further learning, generalization, and understanding.

The burden is traditionally on the data scientist or programmer to carry out the feature extraction process in most other machine learning approaches, along with feature selection and engineering. Feature extraction usually involves some amount dimensionality reduction as well, which is reducing the amount of input features and data required to generate meaningful results. This has many benefits, which include simplification, computational and memory power reduction, and so on. Programmers would train a neural network to detect an object or phoneme by blitzing the network with digitized versions of images containing those objects or sound waves containing those phonemes. If the network didn't accurately recognize a pattern, an algorithm would adjust the weights. The eventual goal of this training was to get the network to consistently recognize the patterns in speech or sets of images that we humans know as, say, the phoneme "d" or the image of a dog. This is much the same way a child learns what a dog is by noticing the details of head shape, behavior, and the like in furry, barking animals that other people call dogs.

Machine learning came directly from minds of the early AI crowd and the algorithmic approaches over the years included decision tree learning, inductive logic programming. Clustering, reinforcement learning, and Bayesian networks among others. As we know, none achieved the goal of General AI, and even Narrow AI was mostly out of reach with early machine learning approaches. As it turned out, one of the very best application areas for machine learning for many years was computer vision, though it still required a great deal of hand coding to get the job done. People would go in and write hand- coded classifiers like edge detection filters so the program could identify where an object started and stopped; shape detection to determine if it had eight sides a classifier to recognize the letters S- T-O-P From all those hand coded classifiers they would develop algorithms to make sense of the image and learn to determine whether it was a stop sign. Good, but not mind bindingly great. Especially on a foggy day when the sign isn't perfectly visible, or a tree obscures part of it. There's a reason computer vision and image detection didn't come close to rivaling humans until very recently it was too brittle and too prone to error. Time and the right learning algorithms made all the difference.

III. Solving Problems by Searching

Early researchers developed algorithms that imitated step-by-step reasoning that humans use when they solve puzzles or make logical deductions. By the late 1980s and 1990s, AI research had developed methods for dealing with uncertain or incomplete information, employing concepts from probability and economics.

These algorithms proved to be insufficient for solving large reasoning problems, because they experienced a combinatorial explosion they became exponentially slower as the problems grew larger. In fact, even humans rarely use the step by step deduction that early AI research was able to model. They solve most of their problems using fast, intuitive judgments.

Simple reflex agents directly maps states to actions. Therefore they cannot operate well in environments where the mapping is too large to store or takes too much to learn Goal- based agents can succeed by considering future actions and desirability of their outcomes Problem solving agent is a goal- based agent that decides what to do by finding sequences of actions that lead to desirable states.

A. Problem Solving Agents

Intelligent agents are supposed to maximize their performance measure This can be simplified if the agent can adopt a goal and aim at satisfying it Goals help organize behavior by limiting the objectives that the agent is trying to achieve Goal formulation, based on the current situation and theagent's performance measure, is the first step in problem solving Goal is a set of states. The agent's task is to find out which sequence of actions will get it to a goal state Problem formulation is the process of deciding what sorts of actions and states to consider, given a goal.

An agent with several immediate options of unknown value can decide what to do by first examining different possible sequences of actions that lead to states of known value, and then choosing the best sequence Looking for such a sequence is called search A search algorithm takes a problem as input and returns a solution in the form of action sequence. One a solution is found the actions it recommends can be carried out execution phase Formulate, search and execute design for the agent. After formulating a goal and a problem to solve the agent calls a search procedure to solve it. It then uses the solution to guide its actions doing whatever the solution recommends as the next thing to do (typically the first action in the sequence). Then removing that step from the sequence. Once the solution has been executed, the agent will formulate a new goal.

B. Environment Assumptions

Static, formulating and solving the problem is done without paying attention to any changes that might be occurring in the environment. Initial state is known, and the environment is observable. Discrete enumerate alternative courses of actions. Deterministic, solutions to problems are single sequences of actions, so they cannot handle any unexpected events and solutions are executed without paying attention to the precepts.

IV. Knowledge Representation and Reasoning

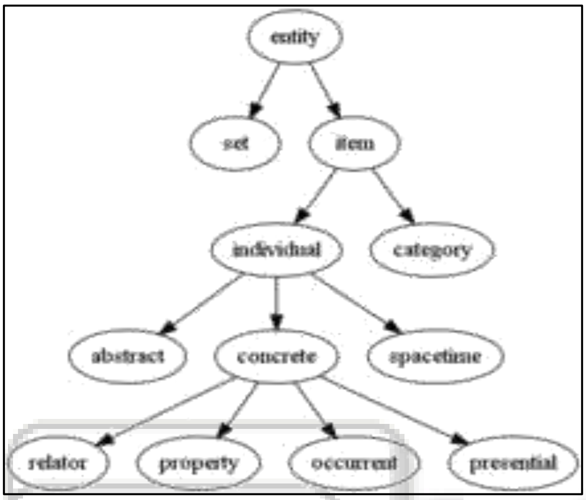

Fig. 1: Ontology represents knowledge as a set of concepts within a domain and the relationships between those concepts.

Knowledge Representation and Knowledge engineering are central to classical AI research. Some expert systems attempt to gather together explicit knowledge possessed by experts in some narrow domain. In addition, some projects attempt to gather the commonsense knowledge known to the average person into a database containing extensive knowledge about the world. Among the things a comprehensive commonsense knowledge base would contain are: objects, properties, categories and relations between objects; situations, events, states and time causes and effects knowledge about knowledge and many other, less well researched domains. A representation of what exists is ontology the set of objects, relations, concepts, and properties formally described so that software agents can interpret them. The semantics of these are captured as description logic concepts, roles, and individuals, and typically implemented as classes, properties and individuals in the Web ontology Language the most general ontologies are called upper Ontologies, which attempt to provide a foundation for all other knowledge by acting as mediators between domain ontologies that cover specific knowledge about a particular knowledge domain. Such formal knowledge representations can be used in content based indexing and retrieval, scene interpretation clinical decision support, knowledge discovery and other areas.

KR Hypothesis Any mechanically embodied intelligent process will be comprised of structural ingredients that

- We as external observers naturally take to represent a propositional account of the knowledge that the overall process exhibits
- Independent of such external semantic attribution, play a formal but causal and essential role in engendering the behavior that manifests that knowledge
- Two issues: existence of structures that We can interpret Determine how the system behaves Benefits of Explicit Representation
- We can add new tasks and easily make them depend on previous knowledge Enumerating objects vs. painting objects
- Extend the existing behavior by adding new beliefs Assert that canaries are yellow
- Debug faulty behavior by locating the erroneous beliefs By changing the color of sky we change any routine that uses that information
- Explain and Justify the behavior of the system The program did X because Y Benefits of Reasoning
- Given Patient X allergic to medication M Anyone allergic to medication M is also
- allergic to medical allergic to medication M
- Reasoning helps us derive Patient X is allergic to medication M KR&R and AI
- KR&R started as a field in the context of AI research Need explicitly represented knowledge to achieve intelligent behavior
- Expert systems, language understanding
- Many of the AI problems today heavily rely on statistical representation and reasoning
- Speech understanding, vision, machine learning, natural language processing. For example, the recent Watson system relies on statistical methods but also uses some symbolic representation and reasoning
- Some AI problems require symbolic representation and reasoning
  Explanation, story generation Planning, diagnosis Abstraction, reformulation, approximation
- KR&R today has many applications outside AI Biomedicine, Engineering, Business and commerce, Databases, Software engineering, Education

V. Reasoning under Uncertainty

A. Uncertainty

1) *Back to planning:*

Let action A denote leaving for the airport t minutes before the flight

For a given value of t, will A get me there on time?

2) *Problems:*

- Partial observability
- Noisy sensors
- Uncertainty in action outcomes
- Immense complexity of modeling and predicting traffic

B. How to Deal with Uncertainty

1) *Implicit methods:*

- Ignore uncertainty as much as possible
- Build procedures that are robust to uncertainty
- This is the approach in the planning methods studied so far

2) *Explicit methods*

- Build a model of the world that describes the uncertainty

- Reason about the effect of actions given the model

C. Methods for Handling Uncertainty

- Default logic: make assumptions unless contradicted by evidence. –
  E.g. Assume my car doesn't have a flat tire.
  What assumptions are reasonable? What about contradictions?
- Rules with fudge factor: E.g. Sprinkler →0.99 Wet Grass", "Wet Grass →0.7 Rain But Problems with combination
- Probability: E.g. Given what I know, A (25) succeed with probability 0.2
- Fuzzylogic: E.g. Wet Grassistrueto degree 0.2 But: Handles degree of truth, NOT uncertainty.

VI. Machine learning

Machine learning is an application of artificial intelligence (AI) that provides systems the ability to automatically learn and improve from experience without being explicitly programmed. Machine learning focuses on the development of computer programs that can access data and use it learn for themselves.

The process of learning begins with observations or data, such as examples, direct experience, or instruction, in order to look for patterns in data and make better decisions in the future based on the examples that we provide. The primary aim is to allow the computers learn automatically without human intervention or assistance and adjust actions accordingly.

A. Some Machine Learning Methods

Machine learning algorithms are often categorized as supervised or unsupervised. Supervised machine learning algorithms can apply what has been learned in the past to new data using labeled examples to predict future events. Starting from the analysis of a known training dataset, the learning algorithm produces an inferred function to make predictions about the output values. The system can provide targets for any new input after enough training. The learning algorithm can also compare its output with the correct, intended output and find errors in order to modify the model accordingly.

In contrast, unsupervised machine learning algorithms are used when the information used to train is neither classified nor labeled. Unsupervised learning studies how systems can infer a function to describe a hidden structure from unlabeled data. The system doesn't figure out the right output, but it explores the data and can draw inferences from datasets to describe hidden structures from unlabeled data.

Semi-supervised machine learning algorithms fall somewhere in between supervised and unsupervised learning, since they use both labeled and unlabeled data for training – typically a small amount of labeled data and a large amount of unlabeled data. The systems that use this method can considerably improve learning accuracy. Usually, semi-supervised learning is chosen

When the acquired labeled data requires skilled and relevant resources in order to train it learn from it. Otherwise, acquiring unlabeled data generally doesn't require additional resources.

Reinforcement machine learning algorithms is a learning method that interacts with its environment by producing actions and discovers errors or rewards. Trial and error search and delayed reward are the most relevant characteristics of reinforcement learning. This method allows machines and software agents to automatically determine the ideal behavior within a specific context in order to feedback is required for the agent to learn which action is best; this is known as the reinforcement signal.

Machine learning enables analysis of massive quantities of data. While it generally delivers faster, more accurate results in order to identify profitable opportunities or dangerous risks, it may also require additional time and resources to maximize its performance. Simple reward trains it properly. Combining machine learning with AI and cognitive technologies can make it even more effective in processing large volumes of information.

VII. Conclusion

AI is an extremely powerful and exciting field. It's only going to become more important and ubiquitous moving forward and will certainly continue to have very significant impacts on modern society.

Artificial neural networks (ANNs) and the more complex deep learning technique are some of the most capable AI tools for solving very complex problems and will continue to be developed and leveraged in the future.