iJRASET
International Journal For Research in
Applied Science and Engineering Technology

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY

Volume: 14 **Issue:** IV **Month of publication:** April 2026

DOI: https://doi.org/10.22214/ijraset.2026.80079

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

Reasoning in Artificial Intelligence: Foundations, Challenges, and Research Directions

Hardik Kumar[1], Aanya Arora[2], Tanisha Tayal[3], Divesh Aherwar[4], Yogita Thareja[5]
[1, 2, 3, 4]Research Scholar, [5]Assistant Professor, Vivekananda Institute of Professional Studies - Technical Campus

Abstract: A lot of work has been carried out in order to develop an artificially intelligent system that thinks similarly to a human being. Some of the abilities that make up such a thinking process include reasoning and decision making. The reasoning process can be broadly classified into two categories, which include deductive reasoning and inductive reasoning. Deductive reasoning entails deriving specific conclusions based on theories and rules while inductive reasoning entails deriving general conclusions based on specific facts. This research paper focuses on defining deductive and inductive reasoning and comparing the differences between the two. The applicability of both deductive and inductive reasoning in artificially intelligent systems is considered in this paper together with a comparative study. Additionally, the advantages and disadvantages of both deductive and inductive reasoning are provided along with their importance in the domain of machine learning. Apart from the aforementioned information, some of the challenges associated with the development of artificial intelligence systems are discussed in this paper. Among other things, algorithmic biases and explainability influence the decision-making capabilities of intelligent machines.
Keywords—Artificial Intelligence, Deductive Reasoning, Inductive Reasoning, Machine Learning, Deep Learning, Neuro-Symbolic AI, Explainability, Causal Inference, Knowledge Representation, Algorithmic Bias

I. INTRODUCTION

However, looking back on the matter from the perspective of hindsight, it is possible to overlook the scale of the initial vision for what was intended by the concept of artificial intelligence. Although it may seem quite insignificant now, the notion that in 1956 a group of scientists gathered together at Dartmouth College and proposed the idea that everything about human cognitive ability could somehow be written down well enough to allow a computer to imitate its abilities was a highly bold claim that virtually no one took seriously, let alone considered controversially.

And the other factor that has remained consistent through all these years of progress is the significance of the problem that has to be solved, namely reasoning. It takes no little effort to program a computer in a way that would allow it to identify some patterns in the stream of visual information, translate from one language into another, recommend movies, and even rank Web pages according to the results of a certain search query. There have been many successes in each of these fields of activity. Reasoning is an entirely different story. And it is the problem under discussion in this article. This paper is arranged in such a manner that there will be three important tasks carried out during it. First, in subsection 2.1, AI will be explained in terms of its definition, background, and current state. Second, in subsection 2.2, reasoning will be explained as well as its importance, kinds of reasoning, and efficiency of existing reasoning methods. Third, the goals of this study will be presented in subsection 2.3. It is suggested that readers who are familiar with the background of AI skip this first subsection; others who do not know much about AI can read it before anything else.

A. Background of Artificial Intelligence

According to the conventional history of AI, found in most academic texts on the topic, AI began in 1956. This dating of the origin of AI is a little bit deceptive because the problems studied by AI — can thinking be reduced to a machine, can we devise a formal logic of reason, what is actually meant by intelligence — have existed long before the creation of the first computer. Leibniz envisaged a formal calculus for resolving disputes by means of mechanical calculation. Boolean algebra was created for formalizing reasoning. Frege, Russell and Whitehead even tried to build mathematics based on a few logical principles. Of course, they met insuperable barriers — as the famous theorems of Gödel have proven.

In the 1956 Dartmouth Conference, McCarthy, Minsky, Shannon et al shared the common objective of finding out if machines could be programmed to use languages, make abstractions and concepts, solve problems that only humans could solve, and learn to do things better.

In a few years, Newell and Simon created the Logic Theorist, which proved theorems in Principia Mathematica by Whitehead and Russell, and the General Problem Solver, which tried to implement a general approach to problem solving. This was no mean feat at the time, and this gave rise, predictably, to an optimistic belief in rapid solutions to the rest of the problems.

This was not to be the case, however. It became increasingly clear by the mid-1970s that any form of AI involving symbolic reasoning, despite its aesthetic appeal, had absolutely nothing going for it where the practicalities of the real world were concerned. Issues regarding combinatorial explosions within search space, the difficulty of representing commonsense knowledge, and the inadequacies of rule-based programming left little doubt that there was much that such software would be incapable of doing. Most notable of these was the report by Sir James Lighthill submitted to the British government in 1973.

There was a sort of resurgence during the 1980s, but not quite in the same sense as before. Expert systems, which implemented the knowledge of experts through production rules, were actually used for practical purposes in areas such as medical diagnosis (MYCIN), chemical analysis (DENDRAL), and information technology configuration (XCON). They did work, up to a point, and businesses spent considerable amounts of money on their development and implementation. However, there was a catch — AI could indeed serve a practical purpose, but only within highly defined boundaries, and expanding the boundaries involved an awful lot of knowledge engineering. If anything, outside of the defined boundaries came up, the system would have no idea what to do.

Another contraction occurred in the late 1980s and the early 1990s, due to high cost and inefficiency of expertise-based methods and problems with the approach. Nevertheless, beyond the pessimism another trend began. Neuroscientists like Rumelhart, Hinton, and LeCun continued to work on the improvement of backpropagation in multi-layered neural networks. Approaches to statistical analysis of natural languages also became popular. All the premises for the domination of machine learning were put in place.

This shift can be seen quite clearly during the '90s and 2000s, with the increasing popularity of statistical and probabilistic methods — SVM, BN, HMM. It became clear that the top results in numerous AI-related tasks do not necessarily require explicitly programmed reasoning algorithms — it only requires careful construction of the learning algorithm to learn from the data. From a philosophical point of view, this was a big leap forward compared to the symbolic approach. Instead of programming the knowledge base itself, scientists started to work with algorithms learning from huge amounts of data. This yielded excellent results in areas like speech recognition, handwriting recognition, and document classification.

And then came 2012 and AlexNet. The triumph of a deep convolutional neural network on the ImageNet image recognition benchmark heralded the dawn of the age of deep learning. The subsequent decade saw a series of breakthroughs happen at a dizzying pace: neural networks that outperformed humans on vision, games like AlphaGo and AlphaZero, protein structure prediction through AlphaFold, and most recently, language modelling. In 2017, the introduction of the transformer model by Vaswani et al. created a whole new era for NLP applications, and in just a few years, massive pre-trained language models on large amounts of text data such as GPT-2, GPT-3, BERT, and T5 exhibited abilities beyond those seen previously.

What happens in the current context is that the size of the models employed is unprecedented. GPT-4, Gemini, Claude, and others like them have been trained on hundreds of billions of words and billions of parameters, with capabilities such as programming, legal reasoning, scientific explanations, and complicated conversations that would have been impossible for any previous AI model a decade ago. However, as will be observed below, some of the oldest issues regarding AI are still relevant. The point is that it appears like these systems can sometimes use logic; whether they do or not remains to be seen.

B. Importance of Reasoning in AI

We need to think about what exactly is meant by reasoning, as reasoning comprises many different concepts. The main idea of reasoning involves an agent deriving new facts based on existing knowledge, where these facts cannot just be remembered but have to be inferred in one way or another. There are different types of inference procedures, each significant both cognitively and computationally.

The most common form of reasoning, which is taught in courses in formal logic, is called deduction. Given certain premises, we make conclusions that logically follow from those premises. Thus, if mammals are warm-blooded animals, and if dolphins belong to mammals, then we can safely conclude that dolphins must be warm-blooded as well — period. The significance of such a kind of reasoning is that it is truth-preserving: if your premises are true, the conclusions drawn from them are also true; otherwise, there is nothing to prove. This is very handy when doing verification, proving theorems automatically, or building constraint-based reasoning engines. However, a downside of deductive inference is obvious: you can never derive anything outside premises.

Induction is the contrary to deduction where an observation leads to a general law. In this case, we observe that the sun always rises each and every day starting from morning till evening for as long as we remember, and therefore deduce that tomorrow morning the sun will rise. This is the essence of machine learning: given some instances, figure out the pattern.

The technique is highly useful since it is this approach that has largely propelled research on AI in recent decades. There is however one downside to this method: inferences by induction are defeasible in nature; that is, they are valid only in normal cases.

Abductive reasoning might be the most frequent form of deduction and also the most difficult one to formulate in a precise way. It involves inferring from an observation to the most reasonable explanation of its occurrence. If a doctor looks at some symptoms and comes up with a diagnosis, he uses abductive reasoning. The same applies to detectives or scientists trying to account for something unexpected. The difficulty with such type of deduction lies in determining how "plausible" can be formulated, because plausibility usually relies on other factors like background knowledge and probability distributions.

Analogy is the ability to see similarities between two different scenarios and then transfer information across from one scenario to another based on those similarities. The theory of evolution via natural selection created by Darwin relied heavily on an analogy between evolution through natural selection and artificial selection performed by farmers for hundreds of years. The formulation of Kepler's laws of motion was aided by analogies between light and physical forces. Analogical reasoning is central to case-based reasoning AI programs, transfer learning, and the generalization abilities that sometimes emerge in large language models, which is among their most fascinating attributes.

For example, causal inference warrants particular focus since it is not only highly significant but also extremely underdeveloped in modern artificial intelligence applications. The difference between correlation and causality goes far beyond simple mathematical distinctions and becomes relevant in almost all real-world scenarios. When a model is trained using observational data, it will notice that patients of hospitals are generally less healthy than average citizens, without knowing that this correlation tells absolutely nothing about the cause-and-effect relationship between hospitals and illnesses. The theory of causality developed by Judea Pearl incorporates those concepts using formal tools such as graphical models and do-calculus. It allows one to distinguish among things we know based on observational data, things we know about interventions, and counterfactual statements. The vast majority of contemporary machine learning algorithms belong to the first category only.

The development of large language models has further complicated these debates. While on certain tests, GPT-4 and its peers exhibit capabilities that would have been unimaginable only a handful of years ago — such as solving complex mathematics problems, answering queries requiring the synthesis of information from multiple sources, making logical arguments — they simultaneously make simple mistakes that anyone who thinks carefully would be able to avoid — contradicting themselves over the course of an extended conversation, fabricating believable falsehoods, failing to understand problems that are identical in logic to others that they correctly solve but are posed in different ways. How should we interpret this? To be frank, we do not know.

But it is important not just in terms of what machines can or cannot do from an academic standpoint. For instance, when developing a system to aid in the diagnosis of patients, it becomes critical to think carefully about how the system infers conclusions based on evidence, and what it is capable of knowing about its uncertainty. When dealing with law, an AI-driven solution to help in legal case analysis must make the distinction between the data available and the information derived from the data. In the finance sector, an AI-powered solution for assessing financial risk without the ability to discern causal relationships may come up with fatally flawed conclusions if there are even slight variations in the distribution from the training set.

There is even a more important element related to the concept of trust. Human beings tend to be less resistant to an outcome that they do not agree with if they know why the conclusion has been reached. Where the machine is capable of defending its decisions as opposed to simply stating them, the machine will become verifiable, improvable, and extendable. In other words, reasoning is not just a case of giving the correct response, but offering an approach that can be verified and refined.

C. Objectives of the Research Paper

This research is focused on achieving three specific objectives that are related to each other in that one depends on the other. Yet, each objective raises its own question through its own methodological approach.

1) Objective 1: Characterising What Current AI Systems Can and Cannot Do: The initial challenge is descriptive: to construct a detailed and theoretically informed description of the reasoning abilities displayed by existing AI models. It turns out to be more complicated than one might expect. There is an abundance of benchmarks available — commonsense reasoning benchmarks, math problems, logic puzzles, science question answering benchmarks — but many of these numbers are hard to make sense of. An AI model that performs excellently on a multiple-choice reasoning benchmark may simply have found some statistical trick to exploit properties of the answer set. An AI model that apparently fails to solve an easy logical puzzle may succeed when presented with a slight variation of the question. We want to move beyond overall performance and consider in detail the patterns of success and failure.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

2) Objective 2: Diagnosing the Structural Sources of Limitation: Objective two arises directly from objective one: if we have a better understanding of what goes wrong with the process of reasoning, we will naturally wonder why. The reasons for some failures are architectural — reflecting constraints on information representation and processing imposed by the structure of our current neural networks. Some limitations are inherent to the way our networks are trained, namely being optimized around fluency rather than structural coherence. Yet others arise due to intrinsic limitations of learning, based on observation alone, and in the absence of explicit biases toward causality and structural relations. The point is that each type of limitation requires a very different approach; thus, it is necessary to identify the source of each of them.

3) Objective 3: Sketching a Principled Research Agenda: Thirdly, this is a future-oriented goal — now that we have identified our strengths and weaknesses regarding our current capabilities, our next step is to propose a research agenda that considers both sides of this picture. This does not involve the idea of offering an instant solution to the problem, since experience from the history of artificial intelligence research suggests that there was never a full-fledged solution in its early days — but rather involves an examination of various avenues worth investigating and possible links between them. Neuro-symbolic integration, chain-of-thought prompting, process supervision, retrieval-augmented generation, causal modelling, and formal verification are all approaches that, on their own levels, relate to improving reasoning in artificial intelligence systems. Ethics of enhancing reasoning skills of artificial intelligence systems should also be discussed here.

These three goals, one would hope, are mutually reinforcing in a productive manner. Simply describing the world without analysing its condition may result in the collection of symptoms without an explanation of their cause. The same may be true for diagnosing problems without suggesting an agenda for change. It is also possible for the agenda to lack credibility without an adequate evaluation of the existing situation. We strive to combine all three simultaneously, presenting both a candid description of today's condition, along with an understanding of why that condition exists.

II. THE IDEA OF REASONING

A. The Essence of Reasoning

At the most fundamental level, reasoning is the act of progressing from recognizing an entity to formulating a judgment concerning it.

The matter of reasoning alone has perplexed many philosophers for more than two thousand years, and their debates are still ongoing. Nevertheless, a number of key differences have proven resilient to time and continue to endure despite this fact. One such difference was established by the ancient thinkers and particularly Aristotle, who delineated between reasoning that proceeded from the universal to the particular and reasoning that went from the particular to the universal. It was in modern times that this difference evolved into the current distinction of deductive reasoning as true-preserving but not ampliative, and of inductive reasoning as ampliative but not true-preserving.

Deductive reasoning and inductive reasoning can be found in our day-to-day activities. Whenever we take actions based on our conclusions from the information given by an expert, then we apply deductive reasoning. However, whenever we base our conclusions on personal experience, then we apply inductive reasoning.

B. Reasoning Environment in AI

In research related to AI, numerous types of reasoning have been formulated, and each emphasizes one aspect of the way intelligent agents draw inferences based on their knowledge base:

- Deduction: It is the use of universal truths to derive absolute certainties concerning particulars.
- Induction: It is the development of general principles from specific facts and is thus associated with uncertain results.
- Abduction: This is a strategy used in determining how one should account for a phenomenon. Abduction is a technique used by detectives and medical practitioners.
- Analogy: It utilizes information acquired in one context to make sense of the structural aspects of another issue.
- Commonsense: Common sense entails the application of one's intuitive understanding of the natural world.

This research is mainly concerned with deductive and inductive reasoning since they are the two types of reasoning that have been thoroughly studied and utilized in artificial intelligence. The contrast between the two is instrumental in understanding a lot about artificial intelligence.

III. DEDUCTIVE LOGIC

A. The Nature of Deductive Logic

Deductive logic is an ideal form of logic in the process of making arguments. With proper use of deduction in your logic, the truth value of your conclusion is not a probability — it is guaranteed to be true.

Here is an example of an excellent case for discussion. A syllogism can be presented at the initial level of studying logic:

- All men are mortal.
- Socrates is a man.
- Hence, Socrates is mortal.

It is neither a guesswork nor approximation but an inference from premises that are always true due to their intrinsic nature.

In other words, deduction is done in logical systems, which give rules on how to infer conclusions. In the logical systems, there are inference rules that have names like modus ponens and modus tollens, and these rules specify what forms of arguments are regarded as valid arguments.

B. The Features That Make It Unique

There are a number of characteristics of the deductive approach which distinguish it from other approaches and affect the way it functions within applications of artificial intelligence.

1) *Virtue of Certainty:* It has already been shown that deduction allows one to get an absolutely reliable result as a consequence of applying a valid reasoning process having true premises. Hence, deduction may be called a tool with great value if there is a necessity to make the audit process more reliable.
2) *Virtue of Monotonicity:* The principle of monotonicity implies that no additional premises that are introduced into the reasoning framework should make the obtained results unreliable.
3) *Limitations Regarding Lack of Ampliativeness:* Since deductions are built using information contained in the premises, finding some other information which has not been included in the initial premises is impossible.
4) *Virtue of Formal Precision:* The formal approach is inherent to deduction.

C. Deduction in Artificial Intelligence

Deduction is one of the types of reasoning that serve as the foundations of artificial intelligence. In the early development stages of artificial intelligence, there were efforts geared towards developing deduction and automated theorem proving to show the machine's logical capabilities.

The Logic Theorist created by Allen Newell and Herbert Simon in 1956 was able to prove 38 of the first 52 theorems in the work of Principia Mathematica of Whitehead & Russell, demonstrating that computers were capable of reasoning in a logical fashion. Newell and Simon went on to develop the General Problem Solver.

Over the following decades, the use of deduction in AI took many shapes, which include the following:

In the case of Logic Programming, as used in Prolog programming language, computer programs could be created where facts and rules were expressed logically, and then a query would be asked to receive answers automatically.

As opposed to this, the present-day automated theorem provers such as Coq, Isabelle, and Z3 are widely used in order to prove the correctness of both software and hardware implementations mathematically.

Yet another application included ontological reasoning which allowed making deductions from knowledge bases, and hence new drugs could be discovered based on comparison with the existing ones.

D. Rule-Based Expert Systems

The most straightforward and effective use of deductive inference in AI is evident in rule-based expert systems. Despite the apparent simplicity of this notion, a lot remains to be accomplished when it comes to the development of an appropriate system of "if-then" rules based on one's expertise.

Rule-based systems have two ways in which the engine works. In the scenario where the engine uses forward chaining, it starts from facts towards conclusions, from causes to their effects. In backward chaining, the engine begins with the conclusion and moves backward in order to fulfil the conditions for making that conclusion, for example, the suitability of the drug.

However, their efficiency has been highly improved through the use of the RETE algorithm, which was developed by Charles Forgy in 1979. The RETE algorithm allows rule-based systems to efficiently compare thousands of rules against thousands of facts without any form of repetitive calculations. Today's rule-based engines are direct offspring of RETE.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

E. Famous Deductive Systems

MYCIN is one of the examples of a deductive system. This system was invented by some of the scientists from Stanford University in the seventies decade. MYCIN is one of the most famous AI systems ever designed. Approximately, there were around 600 rules in MYCIN. MYCIN demonstrated performance on par with human experts during tests — an impressive feat considering that the program knew nothing other than the rules embedded within its architecture.

MYCIN was also capable of explaining its reasoning process; it could guide a doctor through the reasoning process behind its recommendation using individual rules.

DENDRAL, created prior to MYCIN, employed deductive rules based on chemical principles to deduce the composition of an organic compound using mass spectrography results. DENDRAL was among the early AI systems that proved its capabilities in solving complex scientific questions.

In addition to their importance in solving problems, such forms of artificial intelligence are regarded as the predecessors of machines that must be developed to facilitate artificial intelligence.

IV. INDUCTIVE REASONING

A. Definition of Inductive Reasoning

Inductive reasoning differs from deductive reasoning in that the former does not rely on general statements when making specific conclusions. In this case, the purpose is to deduce generalizations based on specifics. It is more probabilistic than deductive reasoning.

This is how the Scottish philosopher David Hume describes this problem very briefly: no matter how many times we observe the sunrise in the east, our observation cannot be proof of the occurrence of the same event the next day; however, at the same time, there is no doubt in this matter. It is on this kind of contradiction that induction is based.

Put simply, induction is the way in which we draw conclusions based on our experiences. A child who has suffered from burns due to the hot stove knows that he must always be careful when he comes close to any stove. Similarly, a doctor who finds out that his patients can be cured of the disease using some particular medicine will prescribe the same medicine to other patients with the same disease. The stockbroker who sees a pattern emerging in his stock charts creates a theory about it.

B. Speciality of It

The speciality of induction which makes it stand apart from all other methods of reasoning is the fact that the inferential process is ampliative in nature and therefore it is possible to draw a conclusion which goes beyond the premises.

On the flip side, defeasibility is an intrinsic feature of induction. Inductive arguments can be disproven, defeated, or refuted by the emergence of new evidence. Every inductive argument is vulnerable to counterexample at all times.

Data sensitivity is another important feature of induction. The validity of the generalization will hinge on the validity of the information gathered using the technique of observation. In case the information collected does not have any validity whatsoever, then the generalization too will not be valid.

Lastly, there is scalability. Induction is highly advantageous because it does not require manually codifying knowledge — unlike most other forms of inference, it scales very well to extremely large datasets.

C. Inductive Reasoning and Machine Learning

Inductive reasoning can be considered a more organized approach to learning. The computer learns through repeated exposure to many examples, and from there, patterns emerge, which the computer uses to hypothesize about future cases that have yet to be encountered.

Theoretical underpinning of this method is strong enough. The PAC learning framework was introduced by Leslie Valiant in 1984. This theory states that what are the precise conditions needed for the working of an inductive learning algorithm, that is, how many examples are needed in order to come up with a hypothesis which is probably correct. Another theory presented by Vapnik and Chervonenkis provides us with a different angle to view the problem.

Inductive reasoning in Bayesian machine learning is regarded as probabilistic reasoning. Beliefs about theories are updated based on evidence mathematically. It is an approach that represents scientific reasoning.

D. Various Aspects of Machine Learning

Inductive artificial intelligence can be formulated in numerous manners in modern times.

Supervised learning: The learning process involves learning a function that maps the input vector into the output vector based on the training data set.
Unsupervised learning: The challenge to be addressed is that of discovering a pattern underlying an unlabelled data set by means of clustering, compressing the data set into the low-dimensional space containing important data only, and even creating new data points.
Reinforcement learning: The intelligent agent learns to behave in the environment based on the rewards received for completing particular actions.
Deep learning: The major success associated with deep learning can be attributed to the fact that the deep neural network was used to learn from high-dimensional and unstructured data inputs.

E. Transformative Systems

AlexNet (2012) was more than just an excellent participant in a machine vision competition. Using a large convolutional neural network trained with one million labeled images, Krizhevsky et al. showed that inductive learning with large-scale datasets would prevail over all hand-crafted computer vision approaches developed over the past half-century. Checkmate.
Google subsidiary DeepMind made several contributions to the progress of science. The inductive learning model invented by the researchers helped solve the long-standing problem of predicting the 3D structure of a protein from its amino acid sequence — almost as accurate as the experiment.
Models like GPT-4, Gemini, and Claude for inductive learning for language modelling show that when inductive learning is applied at a grand scale, our models are capable of not only thinking but also generating, translating, and solving puzzles across numerous areas. What is even more surprising is that these models learn grammar, knowledge, and logic rules without being explicitly taught how to do so.

V. DEDUCTIVE VS. INDUCTIVE: AN HONEST COMPARISON

A situation where one of the approaches would always prove better than the other approach is highly desirable. But such a boring situation does not arise in life.
The two modes operate in reverse directions. From the general to the specific is the route that deduction follows, whereby a general rule is used to deduce something concerning a concrete situation. The other route taken by induction involves going from the specific to the general. These contrasting directions inform all other aspects of their operation.
When it comes to the issue of certainty, the difference is quite clear. For instance, when a valid deduction has been made using true premises, the outcome is certain. This means that nothing else can result apart from what is already in place. When talking about induction, however, the conclusion made is probabilistic rather than being certain. It is the best explanation of the data at hand.
Knowledge sources for both methods are equally clear-cut. In deductive AI systems, knowledge sources come in the form of carefully codified knowledge — rules, facts, and relations, which had to be written down by human experts. For inductive AI systems, however, sources of knowledge come from data itself, finding regularities among large amounts of examples without human expertise involved to state them explicitly. As a result, while being easier to create because of reduced need for human input, inductive knowledge depends completely on quality and representativeness of training data.
Scale is definitely an area where inductive approaches have the upper hand in practice. Deductive knowledge struggles when applied to domains which become richer and more complex because there is simply too much to account for. In contrast, inductive knowledge tends to perform better the harder and larger problem becomes.
It is almost the opposite case regarding the ability to explain decisions. While the workings of deductive systems are clear-cut and easy to follow because each decision derives from specific rules applied to certain premises, inductive systems, especially contemporary neural networks, tend to operate through mechanisms that their creators cannot entirely elucidate.
Adaptability presents yet another difference. As mentioned above, inductive systems adjust automatically by updating their generalizations whenever they receive fresh data. In other words, they learn as they operate. On the contrary, deductive systems are unable to gain new experience because, apart from being updated manually by humans, their database remains unchanged.
Each method has its own set of weaknesses. The weakness of deductive approaches arises due to the inadequacy of information for an unforeseen event. Inductive approaches can suffer from overfitting and biases as well as make wrong correlations of the relationship if the latter no longer holds.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

Neither approach can be considered superior. Although deductive approaches provide clarity, dependability, and audibility, they are brittle and challenging to construct. Conversely, inductive approaches are sturdy, scalable, and flexible, but they lack clarity and certainty, in addition to being susceptible to garbage-in, garbage-out data issues.

The best AI systems currently available actually do not have a preference for one approach over the other — they embrace both. Self-driving cars use inductive deep learning to understand the environment around them and deductive formal rule-based systems to implement any safety restrictions. Medical AI applications employ inductive methods for data analysis and deduction of clinical guidelines to ensure that suggestions do not go outside the safe limit.

However, one overlooked aspect of such a comparison is how each approach deals with novelty. Deduction is, by nature, very conservative since it relies solely on information that has already been formalized. In an unprecedented situation, there is simply no room for deduction to improvise in any manner; all the system can do is state that the situation does not fit into its framework. On the other hand, induction systems rely, by definition, on generalizing to novel scenarios, something which they are meant to be capable of doing.

However, this capacity also comes with a downside — if an induction system is presented with an unexpected scenario, its generalization could prove catastrophically inaccurate. What is worse is that the result could actually appear perfectly reasonable, and the user would have no way to recognize that the system has reached its own limitations. It should be evident that this aspect becomes extremely significant when dealing with critical industries such as aviation, medicine, and robotics.

The issue of trust — technical and social — distinguishes the two approaches in practice as well as theory. The deductive approach builds its legitimacy on transparency; the stakeholder who comprehends the logic behind the algorithmic structure may review the results of its application and dispute any step in the process of inference. It is well-suited for those contexts in which accountability cannot be avoided: courtrooms, regulatory agencies, clinical governance panels. The inductive approach, by contrast, garners trust based on track record and performance evaluations.

In other words, a successful track record in large-scale samples, in which the algorithm demonstrated high levels of accuracy and precision, is evidence of its trustworthiness. While this is clearly an important quality, it is retrospective and aggregate in nature. While the system's predictions are accurate for the majority of cases, the fact that the algorithm failed to accurately predict outcomes in one particular instance is of limited concern.

VI. APPLICATIONS IN THE REAL WORLD

A. Expert Systems in Practice

However, up to the middle of the 1970s and through much of the 1980s, the expert systems were predominant forms of applied AI and, indeed, there are many problems that expert systems intended to solve and which still remain unresolved up to the present day.

For example, expert systems were applied in the sphere of medicine where they helped doctors conduct differential diagnosis, identify drug interactions, and develop treatment alternatives. One of such systems was the DXplain system of the Massachusetts General Hospital where there is simply an enormous amount of information regarding medicines and which cannot be memorized by any doctor.

Finance: Rule-based systems include regulation and risk assessment criteria together with credit regulations. Whenever a certain bank provides credits or identifies deals as suspicious, the likelihood that a rule-based system will be involved is rather high.

Law: In legal science, contract analysis, and regulation, rule-based systems involve rules together with machine learning procedures that analyse a huge number of court cases.

B. Data Analysis and Prediction

The influence that applications of inductive AI exert on data analysis is groundbreaking and extends across nearly every sphere of scientific study and commercial enterprise.

Within the realm of medicine, artificial intelligence systems employing neural networks and making use of medical image datasets have shown themselves capable of detecting particular cancers, diabetic eye disease, and neurological disorders at levels comparable to or even better than the diagnostic proficiency of radiologists. This is especially significant when one considers that, in screening operations, the rate of growth is always limited by the quantity of specialists available.

Within climate science, predictive models generated via induction from observations gathered over several decades, alongside simulations, have assisted in developing more accurate weather forecasts and gaining insight into the workings of climate change.

Within the financial services industry, superfast trading software that operates faster than any human could possibly execute trades is largely founded upon induction, wherein patterns can be recognized that cannot be by humans.

C. Decision-Making Systems

Increasingly, these reasoning systems become involved in decisions that have a significant impact on the lives of people.

Self-driving cars: Perception of the extremely complicated environment and decision-making about what to do within certain limits in this environment in a time frame that allows real-time processing are important. Inductive learning for object and situation recognition, as well as deductive verification of constraints (in order to ensure that suggested actions meet the safety requirements) is needed.

Criminal justice: Systems of many countries used inductively trained tools for estimating criminal risk for decision-making regarding pretrial detention, parole, and sentencing. Ethical concerns regarding use of such systems are considered in the next section.

Medical recommendation systems suggest treatment plans on the basis of models for disease progression learned from experience and formalization of clinical guidelines.

VII. LIMITATIONS OF DEDUCTIVE AND INDUCTIVE REASONING IN ARTIFICIAL INTELLIGENCE

A. Introduction

The reasoning mechanisms act as the core component in AI systems because they enable them to be able to conduct reasoning tasks similar to human beings. There exist many reasoning mechanisms, although the two common reasoning mechanisms are deductive reasoning mechanism and inductive reasoning mechanisms. Although it is true that deductive and inductive reasoning mechanisms play a critical role in intelligent systems, they pose some challenges in their implementation.

B. Limitations of Deductive Reasoning in AI

Deductive reasoning plays an incredibly important role in AI and especially in those AI's based on logical reasoning. Nevertheless, even though it is highly effective, there are numerous limitations for such a type of reasoning that negatively impact its efficiency. The limitations are enumerated below.

1) *Dependence on Predefined Knowledge:* In this method, the AI process operates under the assumption of completeness and accuracy of the rules set out in advance. Thus, the design of an AI system using this technique will be heavily dependent on the availability of defined knowledge bases. The problem arises when inaccurate knowledge bases are used.
2) *Lack of Flexibility:* That the system is very stringent makes perfect sense. This is due to the fact that the system has absolutely no capacity for adaptation whatsoever. It is also unable to discover truths that lie outside of itself.
3) *Knowledge Engineering Bottleneck:* Among the limitations associated with the use of deduction based AI systems in decision making is the issue known as the "knowledge engineering bottleneck." The reason for this limitation is the challenge of constructing an efficient knowledge base, which comprises facts and rules that have to be correctly developed by experts.
4) *Inability to Deal With Ambiguity:* The deductive approach is based on a two-valued logic system (either true or false), making it less effective at dealing with ambiguous or incomplete information.
5) *Inefficiency in Handling Large Amounts of Data:* Another drawback that can be highlighted in connection with deductive methodology used in artificial intelligence is inefficiency in handling large amounts of data or complex rules. Deductive methodologies make use of pre-existing rules and logical reasoning processes in order to arrive at a conclusion. More the volume of data, more complex the rules become.
6) *Learning Capability Is Restricted:* An important limitation associated with deductive reasoning as applied in artificial intelligence is that its learning capability is highly restricted. In this regard, deductive systems in AI only work by following predefined rules and logic. Consequently, there is no learning capability involved in the process since these systems cannot learn anything new.

C. Limitations of Inductive Reasoning in AI

Inductive reasoning is a common practice employed in today's systems of artificial intelligence. In fact, such reasoning helps the systems acquire knowledge from the data by finding certain patterns. Nevertheless, although flexible and versatile, inductive reasoning possesses various crucial limitations that should be taken into account when employing it.

1) *Importance of Data Quality and Amount:* Inductive reasoning within the scope of artificial intelligence heavily relies on data; thus, there is a high dependency on data quality and amount when it comes to developing reliable and efficient AI systems. The collection of well-representative, high-quality, and diverse data sets plays an important role in creating a proper model. Nevertheless, in practice, there might be cases when data is imperfect and even misleading.

2) *Uncertainty:* Inductive reasoning does not have the same guarantee as deductive reasoning, where one can be sure of obtaining logical certainty from their premises. Inductive reasoning results in probable answers. This is because one draws conclusions based on what is observed, which may not always be right.
3) *Overfitting and Underfitting:* The major limitation of the inductive reasoning approach can be said to be overfitting and underfitting. Overfitting is a scenario where the machine learning process is too perfect in training based on the training data, including all unnecessary aspects that make it unable to generalize using new data. Underfitting occurs in simple models that cannot train using the training data.
4) *Bias and Fairness Concerns:* Inductive approaches to AI are more likely to be biased since they derive insights from past information. If the information that has been provided during the process of training contains any form of social, cultural, or institutional bias, then the model created will likely contain the same level of bias.
5) *Explanation Problem:* Numerous models created from the inductive reasoning approach are often referred to as "black boxes," which implies that even though these models achieve remarkable precision levels, they lack excellent explanatory powers. It is difficult for users and creators, among other stakeholders, to understand the output generated by these models.
6) *High Computational Demands:* The high computational demands of inductive reasoning in artificial intelligence is another notable drawback, especially when applied to sophisticated machine learning techniques like deep neural networks. Such machine learning algorithms are computationally intensive, consuming vast amounts of time and resources during the training process.

D. Shared Limitations of Both Approaches

While deduction and induction may be regarded as two necessary methods in artificial intelligence, some weaknesses are observed regarding each of them. These weaknesses can be considered similar to those of both approaches.

1) *Inability to Completely Imitate Human Thinking:* A major limitation that exists within deductive as well as inductive reasoning techniques is the inability to imitate completely human thought processes. Human thought processes rely not only on logic or pattern recognition based on available data but also include intuitive and creative thinking processes.
2) *Ethical and Societal Issues:* The use of both forms of logic within AI poses some serious ethical and societal issues. Both the deductive and inductive forms of logic can lead to biased and discriminatory outcomes. The reason behind this in the case of deductive logic is the limitations and imperfections of the rules involved while in the case of inductive logic, it is the biases present in the training data.
3) *Problems of Generalization:* The second limitation that applies to both methods concerns the generalization of the acquired knowledge. The deductive approach has limited capabilities due to its dependence on pre-defined rules and inability to extend beyond its existing repository of knowledge. On the other hand, even though the inductive approach has gained knowledge, it still fails to generalize due to the data used.
4) *Absence of Contextual Understanding:* Neither of the two kinds of reasoning used by AI models possesses context sensitivity. Even if the model is able to analyze and understand information, it cannot understand the meaning behind the analysis of the information. Therefore, it is possible that even when the model produces the right answer to the information inputted into it or to the rules defined for it, the outcome might be completely illogical to the greater picture.

VIII. ETHICAL CONSIDERATIONS IN AI-BASED REASONING SYSTEMS

The incorporation of reasoning algorithms in artificial intelligence (AI) programs has improved decision-making capabilities, data analysis processes, and performance tasks. Nevertheless, with the growing involvement of AI in crucial sectors such as health care, finance, policing, employment recruitment, and self-driving cars, ethical considerations have emerged as one of the most important aspects for exploration. Ethical dilemmas associated with AI-based reasoning programs can be attributed to both deductive reasoning systems and inductive reasoning frameworks.

A. Bias in Machine Learning Systems

The issue of bias is one of the main ethical issues faced by inductive reasoning systems. Inductive reasoning relies heavily on learning patterns from historical data, and in case there are any biases present in the society or institution or culture, then the same bias will be replicated into the outputs of the system.

International Journal for Research in Applied Science & Engineering Technology (IJRASET)
ISSN: 2321-9653; IC Value: 45.98; SJ Impact Factor: 7.538
Volume 14 Issue IV Apr 2026- Available at www.ijraset.com

It could result in discrimination against certain categories of people. For instance, a biased recruitment model would consider some particular demographic categories higher than others because of biased data related to hiring practices in the past, irrespective of whether the model had an explicit programming bias.

In the case of deductive reasoning systems, the biases could stem from humans designing the rules and basing their judgment on subjective information about the domain.

Therefore, both types of reasoning could generate unethical outcomes. Some of the methods for handling bias include:

- Data preprocessing and balancing
- Biased algorithms in machine learning
- Systems audit of AI
- Include diverse data

B. Transparency and Explainability

Transparency implies the capability of AI algorithms to explain their actions in a coherent way. In addition, transparency is a primary necessity for gaining the trust of AI users.

Deductive inference systems demonstrate higher levels of transparency since they are based on explicit rule bases. Every decision made by such a system is justified by the specific rule or inference rule used in the decision-making process.

On the contrary, induction-based inference systems usually operate as black boxes; even if they show good performance, it is difficult to interpret their decision-making logic. Thus, explainability issues in the case of inductive models become critical when it comes to accountability because of possible adverse consequences of errors made by AI systems.

One solution proposed in AI research is Explainable Artificial Intelligence (XAI). XAI implies the development of techniques for explaining the decision-making process of an algorithm without decreasing its efficiency. These techniques include:

- Analysis of feature importance
- Model agnostic methods of explanations (LIME, SHAP)
- Visualization of network activation
- Rule extraction techniques

C. Privacy and Data Security Problems

Inductive algorithms are based on extensive data where most contain information related to the private lives of individuals. The problem becomes more complicated considering that such data is not necessarily obtained through proper consent or security protocols.

Privacy problems include the following:

- Illegal access to private data
- Data breach
- Unapproved surveillance and tracking of individuals
- Securing secondary purposes of data utilization

Moreover, the use of private information to create machine learning models may lead to an accidental disclosure of such information.

Methods employed to handle this include the following:

- Hiding data with anonymization and encryption
- Following proper privacy policies (like GDPR)
- Collection of data with user consent
- Machine learning training within secure environments

D. AI-Based Decision-Making Process — Accountability and Responsibility

Accountability could be regarded as yet another issue in ethics that deserves to be discussed. In this case, it might be quite complicated to hold a person accountable since there may be some consequences due to the application of AI-based decisions in making conclusions. Following the deductive approach, the developer of the algorithm should bear responsibility under certain criteria.

In turn, taking the inductive approach, the person making the decision should be responsible for his/her work together with the data scientist and algorithm developer.

E. Process of Decision Making Using Artificial Intelligence — Accountability and Responsibility

Accountability is another ethical problem which needs to be mentioned in this respect. In this context, it will be rather difficult to make someone responsible because there could be some negative consequences resulting from the use of decisions made on the basis of artificial intelligence. From the standpoint of the deductive approach, it will be the developer of the algorithm who will be responsible according to specific criteria.

From the perspective of the inductive approach, both the decision maker and data scientist will be responsible.

F. Regulatory and Ethical Governance Issues

There have been many advancements made in artificial intelligence technology, but not in the field of regulation. It is only now that governments around the world are trying to come up with ethical regulations for artificial intelligence applications.

The following are some issues concerning ethical governance:
- The formulation of universal ethics
- The implementation of ethical regulation within various countries
- The control of autonomous systems
- The abuse of artificial intelligence technology

IX. THE ROAD AHEAD

A. Bringing Deduction and Induction Together

The most promising area of research in terms of advancing AI reasoning is to find a way to reconcile deductive and inductive systems by creating AI that will be able to learn and reason simultaneously. Neuro-symbolic AI systems are aimed at solving just that problem.

A number of hybrid models have demonstrated considerable success, such as AlphaProof — an AI program created by DeepMind which merges a language model based on mathematical text with a mathematical proof assistant, giving the model an ability to discover proofs of mathematical problems, which would be impossible otherwise. Other researchers are working on building models with embedded logic into a neural network, or vice versa — training a neural network according to logic-driven criteria and using less data than before.

The ultimate goal of this research is to create an AI program capable of reading a medical research paper, making generalizations based on information contained in the paper, and then applying these generalizations in accordance with the rules established in clinical practice in the form of deductive reasoning.

B. Learning to Reason About Causation

One of the most fundamental problems with existing inductive AI is that while it can learn correlation, it cannot learn causation. For example, an AI system trained on medical records might learn that patients taking a certain drug have a high recovery rate, without knowing whether the drug itself is causing recovery or some other factor is influencing both the use of the drug and the patient's likelihood of recovery.

It becomes critical in such cases because when the system is used for decision-making, it will be accurate when everything is stable but fail dramatically when conditions change.

The theory of causal inference by Judea Pearl gives formal mathematical techniques to represent and reason about causes. Connecting causal inference to modern machine learning is one of the most exciting directions in AI today.

C. Commonsense and World Knowledge

Despite their remarkable prowess, it would be necessary to mention that there is one very crucial shortcoming associated with such models — their inability to conduct common sense reasoning, whereby a simple child of five can surpass them.

However, it should be understood that it is not only superficial since it is a genuine flaw in the manner through which present-day machines learn about the world. Fixing such problems would require completely new approaches and ways of storing both learned and structural data about the world, as well as different methods of training.

D. Safety and Formal Verification

As the influence of AI increases into critical areas such as medicine, construction, transportation, and national security, there will be a need for more formal assurances about how these systems perform. Merely being probabilistically accurate is not enough; a failure may result in casualties.

This is leading to a trend towards using formal verification techniques from computer science to ensure that a system behaves as expected in all possible situations. While formally verifying systems with learned components is a difficult task, it is becoming ever more necessary, and this is one of the areas where deduction will remain indispensable in AI systems.

E. Governance and Institutional Responsibility

Just as critical to the design and functioning of tomorrow's AI reasoning systems as technical considerations is the issue of governance. The reliability of hybrid neuro-symbolic approaches, causal reasoning engines, and formal verifiers depends just as much on institutional contexts as it does on their technical capabilities. It is possible to build a reasoning engine with perfectly sound reasoning processes in the lab that can nonetheless wreak havoc if it is put into operation without adequate human oversight, without a means for disputing the conclusions it makes, or without transparency about the training process that produced it.

Thus, the pathway forward for the future of AI reasoning is both technological and political in nature. The very best technological advances will not, in themselves, solve the problem of responsibility when an AI comes to the wrong conclusion in a manner that has serious consequences. This issue must be addressed by society at large and poses just as urgent a problem.

F. Reasoning under Deep Uncertainty

Another frontier that has only just begun to be investigated is the ability of AI systems to engage in reasoning under deep uncertainty—uncertainty that is not the well-formed, probability-oriented uncertainty of a well-trained classifier, but the more fundamental uncertainty of a situation where even the ground rules of the problem at hand are either unknown or unclear. This sort of uncertainty confronts human experts on a regular basis: the physician faced with a novel disease, the jurist called upon to interpret a statute that has never been applied to technology that did not exist when the statute was drafted, the engineer who must design for a future climate with no historical precedents. Neither of these two types of system is well-suited to dealing with such a problem. The deductive system requires that we know the pertinent rules of the matter beforehand, while the inductive one requires us to have appropriate training data.

The development of AI reasoning systems capable of doing this kind of reasoning—ones that are capable not only of recognizing their own uncertainty but also of reasoning productively in spite of that uncertainty—is likely to be the most significant challenge in the decade ahead.

X. CONCLUSIONS

The goal of this paper was relatively clear from the outset: namely, to understand the true position of artificial intelligence in relation to reasoning neither its position according to any kind of hype, nor its position according to any kind of detractor, but its actual standing. By covering everything from the history of artificial intelligence to practical application, limitations, and moral implications, it becomes possible to come to a conclusion regarding the state of things. The following summary will not try to answer all the questions that remain, simply because there is much that remains to be explored in the field; however, it will summarize some of the findings of the present paper.

A. What This Paper Has Argued

The primary thesis of the whole writing revolves around the concept that reasoning is an integral part of artificial intelligence. It is really obvious that machines that can see what is around them look at information and talk to people like we are having a conversation are very impressive. There is still a lot of work to be done. Making machines like this is one part of a big job that needs to be finished and it might even be one of the easier parts compared to the other things that need to be done with these machines. Reasoning, which involves the interpretation of the observed events, conclusion-making, assessing the validity of the arguments, and rejection of incorrect assumptions based on available evidence, is the other half, and the main topic of this paper.

As evident from the brief history provided in section two above, the issue has existed even from the time the field of artificial intelligence was first explored. The initial researchers at the Dartmouth University were not only interested in developing AI algorithms to perform certain tasks; their ultimate goal was to develop thinking machines. Early strategies in this regard consisted of symbolic logic and formal methods of reasoning together with expertise knowledge bases.

Although such efforts brought good results in closed systems, these became extremely fragile in open systems. The later inclusion of statistical and induction-based strategies since the 1980s to date, culminating in the emergence of deep learning since 2012, has greatly improved AI performance but also created types of fragility.

The foregoing two sections reviewed two prominent approaches in the field of logical reasoning, namely deduction and induction. As was mentioned before, deduction, which utilizes classical logic, has been applied to the development of rule-based expert systems such as MYCIN and DENDRAL. There are a number of strong points about this approach. First, the approach guarantees certainty in its results. Second, the output generated by a deductive approach is easily interpretable since all steps taken within the process follow certain rules. The downside of the approach, however, is quite obvious: deduction does not allow self-learning, fails to deal with uncertain situations, and requires much effort due to its dependency on knowledge bases. In contrast, the second approach, namely induction, which has become the basis of machine learning, starts from specific observations and generates a general conclusion. The main advantages of induction include scalability, potentiality for discovering unseen patterns by humans, and self-learning ability. However, unlike deduction, it produces probabilities only.

The comparison in Section 6 showed us quite clearly that pitting these two paradigms against each other is an entirely flawed approach. The most advanced AI systems currently in operation, such as self-driving cars, decision-support applications in medicine, and risk models for banks, all rely on both inductively trained neural networks for perception and modelling of the world, as well as on deductive constraints on recommended actions to make sure these actions will not lead to anything outside of safe boundaries. It is not a question of which one is superior; it is a question of how they can work together.

From sections 7 through 9, this discussion was tied to practical matters. In medicine, economics, jurisprudence, and law enforcement, there are already decisions being made by artificial intelligence (AI) or reasoned about in such a way that influences the outcomes. The deficiencies noted in section 8 — how brittle the deductive reasoning becomes when faced with unprecedented cases; how prone the induction model is to overfitting, to exacerbating historical prejudice, and to failing in unforeseeable ways once the context shifts — are more than just minor obstacles. These are sources of injury when these models are applied without proper precautions in place. And the ethical challenges identified in section 9 — bias, opacity, accountability, privacy, and the new regulatory lacunae — are not matters of algorithmic refinement. There must be structural fixes: better protocols, more robust oversight systems, and an attitude change in the research community regarding the consideration of these issues upfront rather than after the fact.

B. The Three Objectives Revisited

Objective 1 was asked for a characterization of the capabilities and limitations of current AI systems as far as reasoning is concerned. This survey suggests that the answer is quite mixed. Current AI is very impressive when it comes to recognizing large-scale patterns, using language fluently, and carrying out complex chains of steps in well-bounded domains, but less so when it comes to maintaining consistency of reasoning, inferring cause and effect relations, making commonsense judgements, and even identifying gaps in one's knowledge. The results on benchmarks that dominate the field's sense of its own accomplishments tend to obscure these distinctions instead of illuminating them. A model that succeeds on a hard math problem on one set of benchmarks could flounder on the same exact problem stated slightly differently. Fluency and plausibility alone can make a model produce compelling but vacuous reasoning, without doing anything other than statistical pattern matching on the training data. We still lack effective methods for discerning reasoning from clever imitation.

Objective 2 The purpose here was to identify the underlying structural reasons behind such limitations. Three specific sources were identified based on such an investigation. The first one relates to the architectural nature of current neural networks. The reason being that these networks are designed to achieve accurate predictions rather than logical reasoning or structured inference. The second source relates to the training paradigm employed in training these networks. Since neural networks learn skills useful in predicting the next token in a language model or maximising a reward function in an artificial simulation, they are well suited for such tasks. They are however not built in a way that allows them to carry out reasoning processes necessary to solve difficult tasks in the real world. Finally, there are concepts whose acquisition cannot be done using purely observational learning regardless of the amount of observational data used. These concepts include causal relations that require experimentation or structural assumptions in order to be understood.

Objective 3 Their recommendation was for a research agenda. The topics suggested in chapter 10 numbered six. They included: the use of neuro-symbolic methods to integrate deductive and inductive methods; development of AI models that use causal reasoning and not just correlation in their reasoning; solving the problem of common sense through new ways of modelling worlds; applying formal verification to AI applications that operate in critical environments;

designing governance frameworks that move in tandem with advances in AI applications; and lastly, developing AI systems that reason effectively with deep uncertainties in their reasoning processes due to lack of knowledge of rules. This is not simply building on what we already know. Most of these tasks would entail new methods altogether. However, these are the correct issues to address and we have made good strides towards achieving each of them.

C. The Deeper Point

The simple truth behind all the technical talk in this paper is easy to forget. Reasoning is important because reasoning links information with understanding and understanding with trust. An extremely well-informed system that cannot reason about all that information in a coherent, clear, and honest way is no intelligent system at all; rather, it is merely a very elaborate database. And applying such an unreasoning machine in high-risk situations without truly comprehending the machine's failings is a dangerous ethical venture as much as a technical one.

However, the science of AI has made immense strides in the easy aspects of intelligence over the last ten years. The difficult aspects of intelligence — reasoning, comprehension, and judgment — continue to resist. This should not be cause for despair. All significant achievements in the science of AI in the past seventy years have occurred when many experts were convinced that the next achievement would be impossible. However, it is certainly cause for honesty. It requires acknowledging the truth of what current systems can and cannot do, resisting the urge to overstate the capabilities of benchmark performance, and undertaking the painstaking task of developing trustable systems.

This is not going to be a straight forward progression from specialized and limited forms of reasoning in artificial intelligence to more complex and universal forms. It is going to take much work on the part of scientists in many disciplines such as computer science, cognitive science, philosophy, linguistics, and mathematics. Much patience and humility will be required on the part of the scientific community and will not be easily attained when conducting research in advanced fields such as artificial intelligence.

D. Final Remarks

If one were to ask Turing's famous question about machine thinking, one would be asking more than just a technical question. One would be questioning the very nature of mind and thinking itself, as well as what constitutes the building of a thinking thing. Despite seventy-five years since his seminal paper in which this famous question was posed, we still have yet to provide a definitive answer to the question. In fact, we have done even less than that, for the task proved to be much more difficult than once thought.

In sum, this paper has made the case that reasoning — real, principled, open, and sound reasoning — is not only the outstanding challenge in artificial intelligence but also the solution to creating machines that can be trusted. Certainly, there are significant obstacles standing in our way. However, they are obstacles that can be cleared through thorough investigation, sincere contemplation, and responsible collaboration. This means that the case made in this paper does not argue for a particular device or structure but instead for an approach to solving such problems that recognizes the difficulty of the reasoning involved and takes into consideration both the strengths and weaknesses of the approaches we use and our hope for better ways in the future.

These machines are not yet thinking. However, figuring out why not, and how it could be done, is a challenge that should be given the same time as the effort currently devoted to the problem.

REFERENCES

[1] St. Anne's College of Engineering and Technology — Introduction to AI and Production Systems.
[2] Taylor & Francis Online — Deductive qualitative research: an integrative approach to designing studies.
[3] P. N. Johnson-Laird, "Deductive Reasoning," Department of Psychology, Princeton University, Princeton, NJ 08544.
[4] P. K. Ray, A Textbook of Deductive Logic for Students. [Publisher details as per source].
[5] J. St. B. T. Evans, The Psychology of Deductive Reasoning. [Publisher details as per source].
[6] L. K. Soiferman, "Compare and Contrast Inductive and Deductive Research Approaches," University of Manitoba, Apr. 2010.
[7] Simon Fraser University — Inductive Reasoning. [Course/publication details as per source].
[8] A. Feeney and E. Heit, Eds., Inductive Reasoning: Experimental, Developmental, and Computational Approaches. Cambridge: Cambridge University Press.
[9] "Inductive Reasoning," Living Reference Work Entry, First Online: 24 May 2017, pp. 1–8.
[10] Journal of Mixed Methods Research, "Inductive/Deductive Hybrid Thematic Analysis in Mixed Methods Research," vol. 17(3), pp. 308–326, 2023.
[11] O. Stock and M. Schaerf, Eds., Reasoning, Action and Interaction in AI Theories and Systems. [Publisher details as per source].
[12] R. Brachman and H. Levesque, Knowledge Representation and Reasoning. San Francisco, CA: Morgan Kaufmann, 2004.
[13] S. Raschka, "Components of a Coding Agent," Ahead of AI. [URL/publication details as per source].

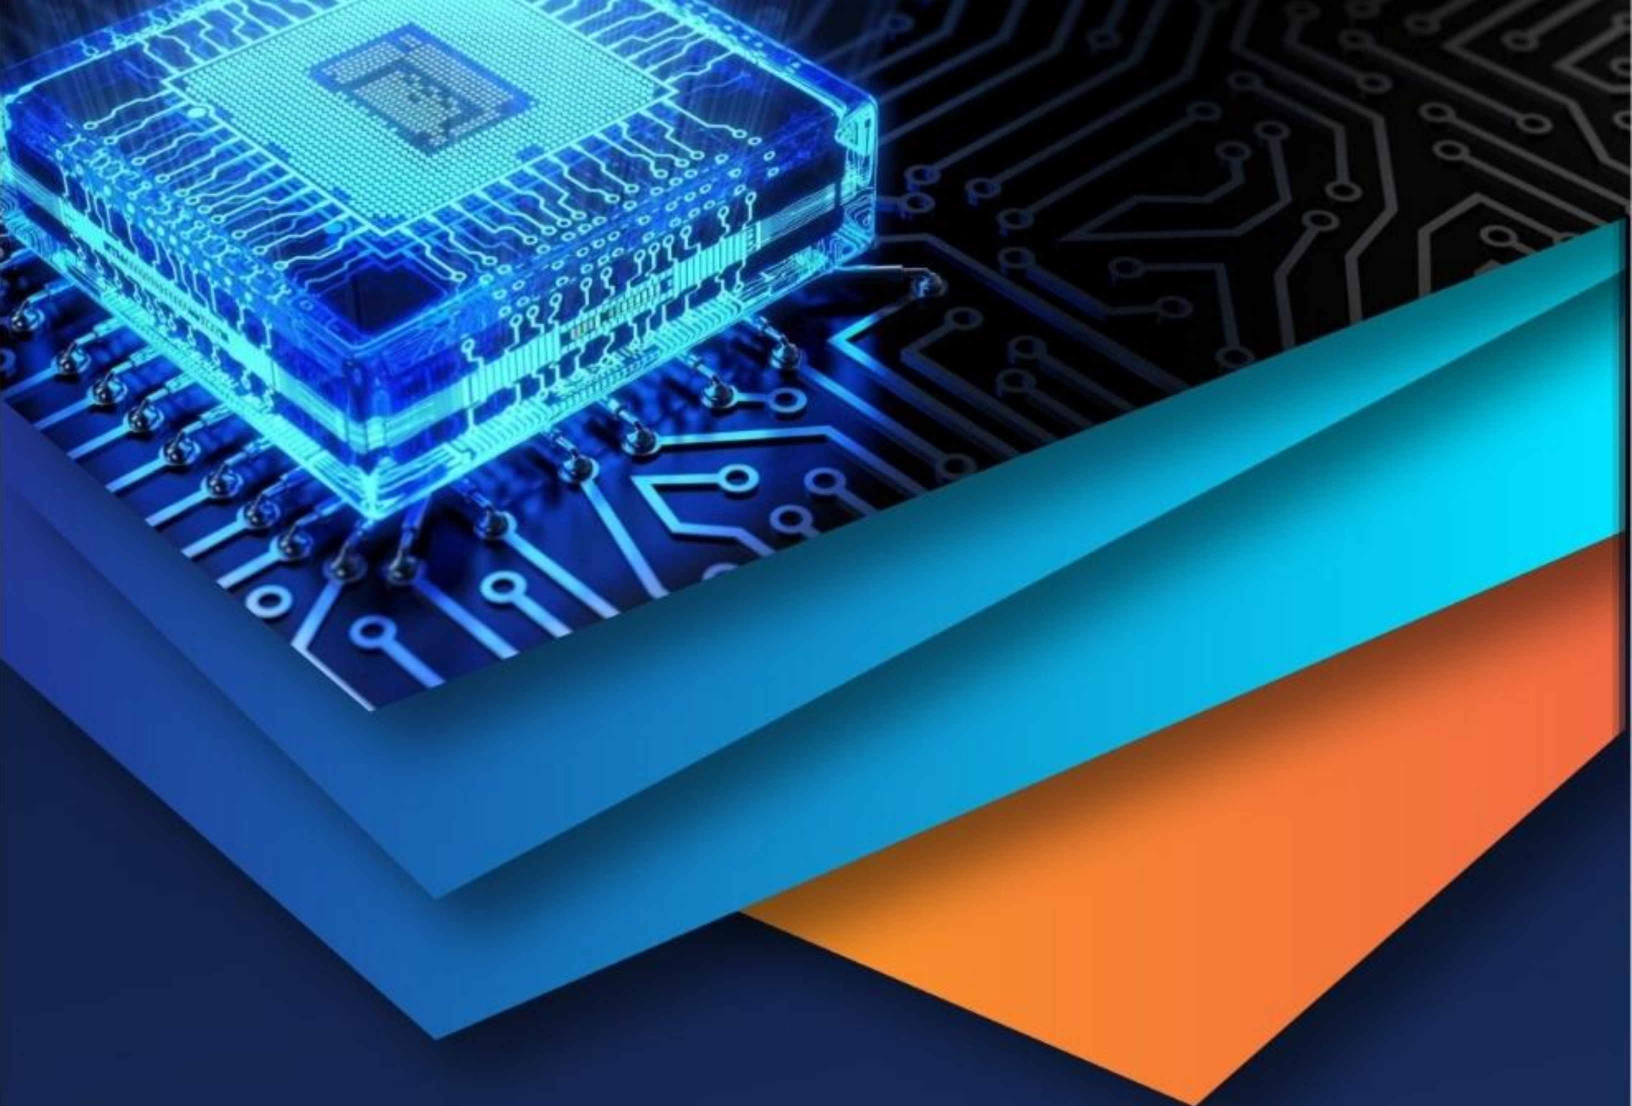

10.22214/IJRASET

45.98

IMPACT FACTOR: 7.129

IMPACT FACTOR: 7.429

INTERNATIONAL JOURNAL FOR RESEARCH

IN APPLIED SCIENCE & ENGINEERING TECHNOLOGY