[2023] 153 taxmann.com 43 (Article)

Artificial Intelligence: Basics

VIKAS PANSARE

CS, Pansare & Associates

Artificial intelligence (AI) is the term used to describe how machines simulate human intelligence. Its definition is always evolving as new technologies are developed to better emulate people and as the potential and boundaries of AI are re-examined. One of today's most hotly disputed topics is artificial intelligence (AI), and there doesn't seem to be much agreement on how human and artificial intelligence vary and overlap.

Rule-based computer programmes that could handle some complexities were the earliest forms of artificial intelligence. The software was partitioned into a knowledge base and an inference engine rather than having every choice hardcoded. Developers would add facts to the knowledge base, and the inference engine would subsequently query those facts to return results. This type of AI was limited, though.

Robotics, self-driving cars, power grid optimisation, and natural language comprehension are just a few of the applications that can benefit from the ability of contemporary AI systems to learn from previous data. Even while AI has occasionally demonstrated superhuman ability in these areas, there is still a long way to go before AI can really match human intelligence.

There is currently no AI that can learn the way humans do, that is, from a small number of examples. To grasp any subject, AI must first be trained on vast amounts of data. Algorithms that can transfer their knowledge from one area to another are still lacking.

Artificial intelligence (AI) systems will initially differ significantly from biological systems in terms of their cognitive capabilities. How can we use (and "collaborate" with") these technologies as effectively as feasible is thus a major concern. What decisions can safely be left to AI and under what circumstances, and when does human judgement become necessary? How can we best utilise the unique advantages of both human and artificial intelligence? How might AI systems be used to successfully complement and make up for the limitations of human cognition (and vice versa)? Should we work on creating AI "partners" with intelligence on par with humans, or should we prioritise overcoming human limitations?

AI concepts:

Artificial intelligence (AI) is the capacity of a machine to carry out cognitive tasks that we typically attribute to human minds, such as sensing, reasoning, learning, interacting with the outside world, solving problems, and even exercising creativity. Making intelligent machines, particularly intelligent computer programmes, is a scientific and engineering endeavour.

Artificial General Intelligence

In computer science, AGI (Artificial General Intelligence) refers to an intelligent system with extensive or complete knowledge as well as cognitive computing abilities. True AGI systems are still the stuff of science fiction as of now. These systems' theoretical performance would be identical to that of a human. An AGI system's goal is to carry out every task that a human being is capable of.

Collective Artificial Intelligence: Collective intelligence (CI) aims to delve into collective human knowledge to make better decisions. Both AI and CI forms transcend the simple individual intelligence but do so in very different manner. AI uses machines and technological development as the way to become smarter as a collective.

Artificial consciousness (AC) A step up from artificial general intelligence (AGI), artificial consciousness (AC) implies sentience and self-awareness in addition to intelligence.

Human Rights: The issues posed by artificial intelligence affect human rights. The fundamental tenet of human rights is the inviolability of human life, with the implicit premise that humans are superior to other kinds of life and so deserving of more protection.

Extensional equivalent intelligence is defined as the ability to produce the same outputs as a human being, given the same inputs. This means that two systems are extensionally equivalent if they can both answer the same questions correctly, solve the same problems, and make the same decisions.

Intentionally equivalent intelligence is defined as the ability to have the same goals and intentions as a human being. This means that two systems are intentionally equivalent if they both have the same understanding of the world, the same values, and the same desires.

In other words, extensional equivalent intelligence is about what a system can do, while intentionally equivalent intelligence is about why a system does what it does.

It is important to note that these two definitions of intelligence are not mutually exclusive. It is possible for a system to be both extensionally equivalent and intentionally equivalent to a human being. However, it is also possible for a system to be one without being the other.

For example, a system that is trained on a large dataset of human conversations may be able to produce the same outputs as a human being in a conversation. However, the system may not have the same understanding of the world as a human being, and it may not have the same values or desires. In this case, the system would be extensionally equivalent to a human being, but it would not be intentionally equivalent.

The distinction between extensional and intentional intelligence is important because it helps us to understand the different ways that intelligence can be manifest. It also helps us to think about the challenges of creating truly intelligent machines.

Here are some examples of systems that are extensionally equivalent to human beings:

A large language model that can answer questions correctly and generate text that is indistinguishable from human-written text.

A chess-playing program that can beat a human grandmaster.

A self-driving car that can navigate a city safely and efficiently.

Here are some examples of systems that are intentionally equivalent to human beings:

A chatbot that can understand the nuances of human conversation and respond in a way that is both informative and engaging.

A robot that can learn from its experiences and develop its own goals and desires.

A virtual assistant that can understand the needs of its users and provide them with the help that they need.

The distinction between extensional and intentional intelligence is still a relatively new concept, and there is much debate about how to best define and measure it. However, it is a promising area of research, and it has the potential to help us to create truly intelligent machines that can interact with the world in a way that is indistinguishable from human beings.

Human Rights are fundamental rights and freedoms that belong to all human beings, regardless of their race, sex, nationality, ethnicity, language, religion, or any other status. Human rights are universal, indivisible, interdependent, and interrelated. They are essential for human dignity and well-being.

Internet of Things (IoT) is a network of physical objects that are connected to the internet and can collect and exchange data. The IoT is made up of sensors, actuators, and devices that are embedded in everyday objects. These objects can be used to collect data about the environment, to control devices, and to automate tasks.

STEM stands for Science, Technology, Engineering, and Mathematics. STEM education is an interdisciplinary approach to learning that emphasizes the application of science, technology, engineering, and mathematics in real-world problems. STEM education is important because it prepares students for the jobs of the future. It also helps students to develop critical thinking skills, problem-solving skills, and creativity.

The evolution of artificial intelligence (AI) can be divided into three main phases:

The early years (1950s-1970s): This was a time of great theoretical advances in AI, but also of limited practical applications. Some of the most important milestones of this period include the development of the Turing Test, the first AI programming language (LISP), and the first expert systems.

The AI winter (1970s-1980s): This period was marked by a decline in funding for AI research, due in part to the failure of some early AI systems to meet expectations. However, some important advances were still made during this time, such as the development of neural networks.

The modern era (1990s-present): This is the period in which AI has seen the most rapid growth and development. Advances in machine learning, deep learning, and natural language processing have led to the development of AI systems that are now capable of performing many tasks that were once thought to be the exclusive domain of humans.

Some of the most notable examples of AI in the modern era include:

Self-driving cars: These vehicles use a variety of AI technologies, such as computer vision and machine learning, to navigate the roads without human input.

Virtual assistants: These AI-powered systems can answer questions, schedule appointments, and control smart devices.

Medical diagnosis: AI systems are now being used to help doctors diagnose diseases more accurately and efficiently.

Financial trading: AI is being used to automate trading decisions and identify patterns in the market.

The future of AI is still uncertain, but there is no doubt that it has the potential to revolutionize many aspects of our lives. As AI continues to evolve, we can expect to see even more amazing and life-changing applications in the years to come.

In addition to the three main phases of AI evolution, there are a number of other important milestones that have helped to shape the field. These include:

The development of the first AI chatbot (ELIZA) in 1966.

The launch of the first commercial expert system (MYCIN) in 1972.

The development of the first neural network (Perceptron) in 1957.

The release of the first deep learning framework (TensorFlow) in 2015.

The evolution of AI is a complex and fascinating story. It is a story of great leaps forward, followed by periods of setbacks and disappointment. But it is also a story of progress, as AI systems have become increasingly sophisticated and capable. As AI continues to evolve, it is sure to have a profound impact on our lives.
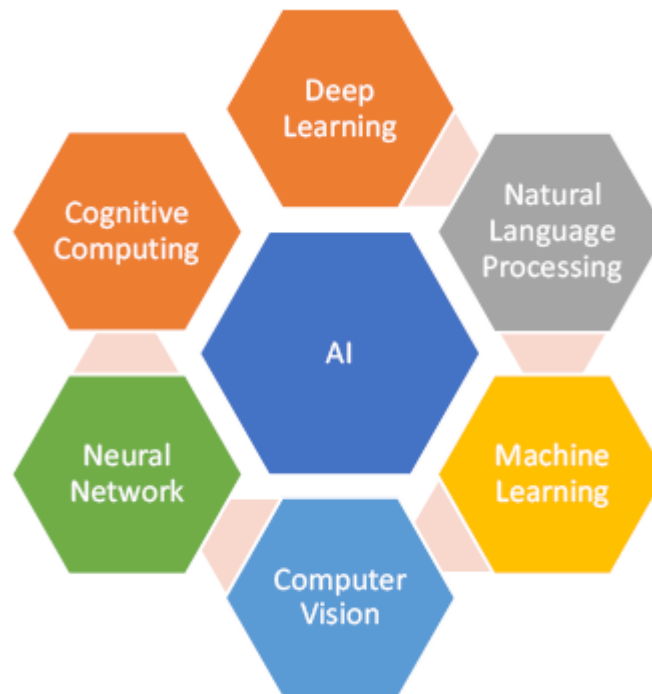
Key concepts in artificial intelligence:

1. Machine learning: This is a type of AI that enables computers to learn and improve their performance without being explicitly programmed. This is done through the analysis of data and the identification of patterns and trends.

2. Neural networks: These are a type of machine learning algorithm that is inspired by the structure and function of the human brain. Neural networks are able to learn complex relationships

between data points, and they are used in a wide variety of AI applications, such as image recognition, natural language processing, and speech recognition.

3. Deep learning: This is a subset of machine learning that uses neural networks to learn from data. Deep learning models are able to learn complex patterns from data, and they are used in a wide variety of AI applications, such as image recognition, natural language processing, and speech recognition.

4. Natural language processing (NLP): This is a field of computer science that deals with the interaction between computers and human (natural) languages. NLP techniques are used in a wide variety of applications, such as machine translation, text summarization, and question answering.

5. Computer vision: This is a field of computer science that deals with the extraction of meaning from digital images or videos. Computer vision techniques are used in a wide variety of applications, such as object detection, face recognition, and self-driving cars.

6. Expert systems: These are computer programs that are designed to mimic the reasoning and decision-making capabilities of human experts. Expert systems are used in a wide variety of applications, such as medical diagnosis, financial trading, and customer service.

7. Robotics: This is a field of engineering that deals with the design, construction, operation, and application of robots. Robots are used in a wide variety of applications, such as manufacturing, healthcare, and space exploration.

8. Game AI: This is a type of AI that is used to create intelligent agents that can play games. Game AI techniques are used in a wide variety of games, from simple board games to complex video games.

9. Big data: This refers to the collection of large and complex datasets. Big data is often used in AI applications, as it can provide the data that is needed to train machine learning models.

10. Ethics of AI: This is a field of study that deals with the ethical implications of artificial intelligence. As AI becomes more powerful, it is important to consider the potential ethical implications of its use.

As AI continues to evolve, we can expect to see even more new and exciting and at the same time challenging concepts emerge.

Development of an AI ecosystem in different countries

Artificial intelligence (AI) is rapidly becoming a major driver of economic growth and social transformation. As a result, governments around the world are developing AI policy positions and strategies to ensure that their countries benefit from this technology.

There is a wide range of AI policy positions and strategies being adopted by different countries. Some countries, such as China, are taking a very proactive approach, with ambitious plans to become global leaders in AI. Other countries, such as the United States, are taking a more cautious approach, focusing on ensuring that AI is developed and used in a responsible and ethical manner.

One of the key elements of any AI policy position is the development of an AI ecosystem. An AI ecosystem is a network of people, organizations, and resources that are working together to advance AI research, development, and deployment. A strong AI ecosystem is essential for ensuring that a country can benefit from AI.

There are several different factors that contribute to the development of an AI ecosystem. These include:

Government investment in AI research and development: Governments can play a key role in supporting AI research and development by providing funding for universities, research institutes, and private companies.

Access to data: AI systems require large amounts of data to train and improve their performance. Governments can help to facilitate access to data by providing open data sets and by supporting the development of data sharing platforms.

Talent development: A strong AI ecosystem requires a pool of skilled AI workers. Governments can help to develop this talent pool by investing in AI education and training programs.

Regulation: Governments can help to ensure that AI is developed and used in a responsible and ethical manner by developing appropriate regulations.

The development of AI policy positions and strategies, as well as the development of AI ecosystems, are complex and challenging tasks. However, these are essential steps for ensuring that countries can benefit from the potential of AI.

Here are some examples of AI policy positions and the development of an AI ecosystem in different countries:

China: China has a very ambitious AI policy position, with the goal of becoming a global leader in AI by 2030. The Chinese government has invested heavily in AI research and development, and it has also taken steps to develop an AI ecosystem by providing access to data and supporting the development of talent.

United States: The United States has a more cautious approach to AI, focusing on ensuring that AI is developed and used in a responsible and ethical manner. The U.S. government has developed several regulations related to AI, and it is also working to develop standards for the responsible use of AI.

European Union: The European Union has a set of guidelines for the ethical use of AI. These guidelines focus on ensuring that AI is used in a way that respects human rights and fundamental freedoms. The European Union is also working to develop standards for the responsible use of AI.

As AI continues to develop, it will be interesting to see how different countries approach this technology and how they develop their AI ecosystems.

The economic impact of AI in India

Artificial intelligence (AI) is rapidly transforming the Indian economy, with the potential to add up to $967 billion to India's gross domestic product (GDP) by 2035.

AI is already having a significant impact on several sectors in India, including:

Manufacturing: AI is being used to automate tasks in manufacturing, such as quality control and predictive maintenance. This is leading to increased productivity and efficiency, as well as reduced costs.

Agriculture: AI is being used to improve crop yields, predict crop prices, and manage pests and diseases. This is helping to increase agricultural productivity and profitability.

Healthcare: AI is being used to develop new drugs and treatments, diagnose diseases, and provide personalized care. This is improving the quality of healthcare and making it more affordable.

Financial services: AI is being used to automate tasks in financial services, such as fraud detection and risk assessment. This is leading to increased efficiency and security.

Retail: AI is being used to personalize customer experiences, recommend products, and optimize inventory. This is leading to increased sales and customer satisfaction.

The economic impact of AI in India is expected to grow in the coming years. A study by the National Association of Software and Services Companies (NASSCOM) estimates that AI could add up to $500 billion to India's GDP by 2025.

The growth of AI in India will create new jobs and opportunities, but it will also lead to some challenges. For example, AI could lead to job displacement in some sectors, as tasks that are currently performed by humans are automated. However, AI is also expected to create new jobs in sectors such as data science, machine learning, and artificial intelligence engineering.

Overall, the economic impact of AI in India is expected to be positive. AI has the potential to boost productivity, create new jobs, and improve the quality of life for Indians. However, it is important to manage the risks associated with AI and to ensure that this technology is used in a responsible and ethical manner.

Ethical Principles of AI

AI also presents challenges in terms of accountability, safety, protection of human rights, maintaining a competitive market, and cyber and information security.

Accountability

In terms of accountability, it is important to have clear definitions and measures in place to ensure that AI systems are fair and in line with values that should not be compromised. Accountability in AI can refer to a variety of values, practices, and measures, and it is important to have an architecture in place that addresses these different aspects.

Safety

Safety is another important concern when it comes to AI. AI systems must be designed and used in a way that minimizes risks to human safety. This includes ensuring that AI systems are transparent and explainable, so that their decisions can be understood and challenged if necessary.

Protecting Human Rights

The protection of human rights is also a crucial issue when it comes to AI. There have been calls for moratoriums on the sale and use of AI systems until adequate safeguards are put in place to protect human rights. This includes ensuring that AI systems do not discriminate against certain groups or individuals, and that they respect privacy and other fundamental rights.

Maintaining Competitive Market

Maintaining a competitive market is also important when it comes to AI. This means ensuring that there is a level playing field for companies developing and using AI, and that there are no monopolies or other market distortions that could stifle innovation or harm consumers.

Cyber & Information Security

Finally, cyber and information security are crucial issues when it comes to AI. AI can be used to improve cybersecurity by analysing large amounts of data to identify threats. However, it is also important to ensure

that AI systems themselves are secure, and that they do not present new vulnerabilities that could be exploited by malicious actors.

Artificial intelligence and the ancient Indian scriptures

Vedas, Upanishads, and Puranas - may seem not be very clearly connected to Artificial Intelligence, but there are some interesting parallels and insights that can be drawn from them.

The Vedas, one of the oldest known scriptures, contain a vast collection of hymns, rituals, and philosophical ideas. They place a strong emphasis on seeking knowledge, comprehending the nature of reality, and investigating one's spiritual side. Similar to human intelligence, artificial intelligence (AI) aims to comprehend and imitate human intelligence in robots.

The nature of consciousness, the self, and the ultimate reality are topics covered in the Upanishads, philosophical writings connected to the Vedas. These concepts can be closely related to the field of AI, as researchers strive to create machines that exhibit various aspects of intelligence, mimic human thought processes, and even simulate conscious awareness.

Additionally, the Puranas, which are mythological and historical texts, often contain stories and narratives about celestial beings, deities, and their interactions with humans. One can draw parallels between these stories and the concept of AI-powered virtual assistants, like the one you are interacting with right now. These virtual assistants, although artificially created, aim to provide helpful information and assistance much like the divine beings portrayed in the Puranas.

Our ancient scriptures do provide valuable insights into the human intellect, consciousness, and spirituality, which can guide our understanding and ethical development of AI. They remind us to approach technological advancements with wisdom, empathy, and a holistic perspective.

It's critical to remember that artificial intelligence as we currently understand it is a relatively recent invention. However, the nature of maya and some aspects of current science found in the Upanishads and Advaita Siddhanta mirror contemporary scientific consciousness. This knowledge is further used to comprehending human thought processes and techniques to simulate them, which helps advance the artificial intelligence field of natural language understanding.

In conclusion, the ancient Indian scriptures offer philosophical ideas and perspectives that can inform our approach to artificial intelligence. Whether contemplating the nature of consciousness, self-realization, or ethical considerations, these scriptures provide a foundation to reflect upon and guide the development and application of AI technologies in a thoughtful and responsible manner.

Conclusion:

Artificial Intelligence (AI) poses several dangers to humanity in the future. These include invasion of privacy, weaponised AI, loss of human jobs, AI terrorism, AI bias, social manipulation through AI algorithms, social surveillance with AI technology, widening socioeconomic inequality as a result of AI, weakening ethics and goodwill because of AI and financial crises brought about by AI algorithms. Many experts have warned that AI could lead to the extinction of humanity.

However, it is important to note that the development of AI also presents numerous benefits and opportunities for society. It is up to us to carefully manage the risks and make the most of its benefits.

Summary:

AI is one of the most debated subjects of today and there seems little common understanding concerning the differences and similarities of human intelligence and artificial intelligence.

While AI sometimes yields superhuman performance in these fields, we still have a long way to go before AI can actually compete with human intelligence. As AI technology develops, it is important to be aware of the ethical and social implications of AI. We need to ensure that AI is used in a way that respects human rights and that benefits all of humanity. Many experts have warned that AI could lead to the extinction of humanity. It is up to us to carefully manage the risks and make the most of its benefits.

Sources:

(1) The Similarities And Differences Of Artificial Intelligence And Human .... https://edubirdie.com/examples/the-similarities-and-differences-of-artificial-intelligence-and-human-intelligence/.

(2) What is Artificial General Intelligence? - TechTarget. https://www.techtarget.com/searchenterpriseai/definition/artificial-general-intelligence-AGI.

(3) Human Rights and Artificial Intelligence - Carr Center for Human Rights .... https://carrcenter.hks.harvard.edu/files/cchr/files/humanrightsai_designed.pdf. Hybrid collective intelligence in a human–AI society.

(4) www.linkedin.com/pulse/10-core-concepts-artificial-intelligence-consultants-factory

(5) Artificial Intelligence: An Existential Risk To Humanity. https://www.ndtv.com/world-news/is-artificial-intelligence-an-existential-risk-to-humanity-4155993

(6) What are the risks and rewards of artificial intelligence? | World .... https://www.weforum.org/agenda/2021/10/artificial-intelligence-ieveryday-lives-aii00-report/.