

INTRODUCTION TO AI AND APPLICATIONS

SUBJECT CODE: 1BAIA103/203

MODULE-4

Introduction to AI and Applications 1BAIA103/203**Module-4**

Trends in AI: AI and Ethical Concerns, AI as a Service (AIaaS), Recent trends in AI, Expert System, Internet of Things, Artificial Intelligence of Things (AIoT).

Textbook 1: Chapter 8 (8.1, 8.2, 8.4), Chapter 9 (9.1- 9.3)

Current Trends in Artificial Intelligence**4.1 AI and Ethical Concerns**

We know that ethics define the discipline that deals with moral obligations and duties of humans. Correspondingly, ethics of AI deals with ethics specific to robots and other artificially intelligent beings. These ethics can be divided into two groups.

1. **Roboethics:** Roboethics consider the moral behaviour of humans as they design, construct, use and treat artificially intelligent beings.
2. **Machine ethics:** These ethics manage the moral behaviour of artificial moral agents (AMAs) that play an important role in delivering artificial general intelligence (AGIs) within existing legal and social frameworks.




FIGURE 8.1 Trusted AI principles

The bigger concerns are as follows:

1. If AI generates human-like output, can it also make human like decisions?
2. If AI makes human-like decisions, then are these decisions human-like?
3. If AI takes decisions of sanctioning or rejecting a bank loan, is that decision justified?
4. If AI decides whether a student should be enrolled in a college or not, is AI taking a fair decision?
5. If AI makes a human-like decision, is it human-like trustworthy?

AI is basically data + mathematical model + training based on data + predictions. What if the data provided for the training is unintentionally wrong/biased?

Issues like these are endless. Hence, getting a good understanding of them is critical to the success of this technology. Refer Fig. 8.1 to learn about some guiding principles of a trusted

AI Tool. As a citizen of the AI society, we must be aware about how AI works and the framework of AI ethics.

4.1.1 Ethical Use of Artificial Intelligence

AI systems no doubt offer a wide range of functionality for businesses, but their use also raises ethical questions. An ideal AI system must not be biased. This is especially true when AI uses inherently unexplainable deep learning and generative adversarial network (GAN) algorithms.




FIGURE 8.2 Features of Ethical AI System

The ability to explain is a potential stumbling block to using AI systems in industries that operate under strict regulatory compliance requirements. For example, financial institutions must explain their credit-issuing decisions. The main problem with such systems is that it will be difficult to explain how the decision was taken to refuse a credit as such decisions are made by teasing out subtle correlations between thousands of variables. In such a scenario, the AI program that is not able to explain decisions is known as black box AI.

A responsible AI system is explainable, monitorable, reproducible, secure, human-centered, unbiased and justifiable (refer Fig. 8.2).

Formulating laws to regulate AI is not at all easy as AI uses multiple technologies. Any strict law can hamper its progress and development. AI technology is changing so rapidly that new technology breakthroughs sometimes make existing laws obsolete. For example, existing laws regulating the privacy of conversations and recorded conversations do not cover the challenge posed by voice assistants like Amazon's Alexa and Apple's Siri. These digital assistants collect data not to distribute them but to improve its machine learning algorithms. This is a golden chance for criminals having malicious intent.

4.1.2 Is AI Dangerous? Will Robots Take Over the World?

Rapid growth and increase in the capabilities of AI systems have made us wonder about the risks involved in using AI. Some of us think that AI will quickly take over tasks performed by humans. But most researchers believe that super-intelligent AI is unlikely to exhibit human emotions, so there is no point thinking that AI will become malevolent. AI can become a risk only in two conditions.

The AI system is specifically programmed to do something devastating. For example, autonomous weapons are AI systems that are programmed to kill. If acquired by scrupulous people having ill intentions, such weapons could inadvertently lead to an AI war, and mass casualties, potentially even the end of humanity.

Autonomous weapons are designed to be extremely difficult to 'turn off', and once they become operational, humans could themselves rapidly lose control over them. This risk is also present with narrow AI when given autonomy.

Just imagine what would happen if an autonomous drone with facial recognition as well as a 3D-printed rifle, pistol or other gun became available. Or if a self-driving car, connected to the Internet, is hacked to get into some serious accident.

Even in hospitals, more and more equipment is now connected to the Internet. If any or all of these are hacked, then can you imagine what the hacker can do with the patient's body?

AI could be programmed to do something beneficial, but the method used to achieve its goal can be highly destructive.

This is especially true when the AI programmer asks the machine to complete a task, but does not clearly outline the goals. For example, you can instruct a self-driven car to program to take you to a particular destination as soon as possible. However, the instruction 'as soon as possible' fails to address safety, road rules, etc. The smart car may successfully complete its task, but after creating a great havoc. So, there must be some provision to continuously monitor and control the machine, its goals and process being performed to achieve that goal.

Of 9 million low-skilled services and BPO roles, 30 per cent or around 3 million will be lost by 2022, principally driven by the impact of robot process automation or RPA. As of now, there is no liability for actions on machines. There is no clarity on what legal aspects bind machines when they become increasingly smart. The following questions may arise:

1. Do we judge AI systems the same way as we judge a human?
2. Who is responsible if the AI system becomes self-learning and autonomous to a greater extent?
3. Is there any error margin for AI machines, even if it has fatal consequences?

Apart from these two serious concerns, AI also poses some additional threats/risks that calls for attention. These threats are discussed below.



Credit: PaO_STUDIO / Shutterstock

FIGURE 8.3 Effect of AI on jobs

1.The immediate risk posed by job automation: As AI robots become smarter and more dexterous, they will quickly replace humans. For example, you must have heard of restaurants having robots as waiters. Just imagine, where the less educated people will go if machines start taking their jobs (as seen in Fig. 8.3)? It is often rightly feared that the deployment of AI systems for job automation will eventually replace certain types of jobs, especially those that are predictable and repetitive. According to a 2019 Brookings Institution study, 36 million people work in jobs that they may soon lose owing to automation as at least 70 percent of their tasks (varying from retail sales, market analysis to hospitality and warehouse labour) will be done using AI. In fact, a newer Brookings report even states that white collar jobs may be worst affected. As per McKinsey & Company report (2018), the African American workforce will be hardest hit by automation.

2.Biased algorithms: We know that computers work on GIGO concept which says Garbage-In-Garbage-Out. This is a very serious limitation in AI applications. If we feed our algorithms data sets that contain biased data, then the output generated from such systems will only produce biased results. This biased result, if applied to solve a real-world problem, may lead to an even bigger problem. A dynamic advertising billboard in Utrecht was switched off because the spy software installed on these billboards aroused public outrage.

3.Too little privacy: When using IoT or AI systems, a lot of data is generated. It is said that approximately 2.5 quintillion bytes (or 2.5 million terabytes) of data is added each day. It is interesting to note that 90% of the digital data has been created in the last two years. A lot of data is required for the proper functioning of the smart systems. As a result, much data about us is collected thereby eroding our privacy. Once the systems collect our data, we have no means to know what data about us is used by whom and for what purpose. These days, cameras can easily be fitted with facial recognition software to capture details about us (including our gender, age, ethnicity, gesture and state of mind). Do you know that in China, some police officers wear glasses with facial recognition technology having a database with facial pictures of thousands of ‘suspects’ (judged on the basis of certain behaviour)?




FIGURE 8.4 AI Tools can be used for Face Swapping and Face Manipulation

4. Everything becomes unreliable: These days, fake news and filter bubbles are not uncommon. Smart systems can create faces, compose texts, produce tweets, manipulate images, clone voices and engage in smart advertising.

An AI system can be used to depict day into night or create highly realistic faces of people who have never existed. For example, open-source software Deepfake can easily stick pictures of faces on moving video footage. Due to such fake videos, celebrities are the worst affected. People with malicious intentions easily create pornographic videos starring them. Even normal citizens are blackmailed with morphed images and/or videos. This is often known as Faceswap video blackmailing (refer Fig. 8.4).

Even in politics, fake or manipulated videos are produced at a high speed to influence the opinions of the masses.

Did you know that the company Cambridge Analytica, was found to manage access to data from 87 million Facebook profiles of Americans to campaign for President Trump so that he could come to power?

4.1.3 Ethics in AI

The current world is quickly getting transformed by artificial intelligence and those who use this technology are living a richer and easy-going life. However, this is possible only when we talk about ethical AI, not just AI. The term, ethical AI, is composed of two words: ethical and AI. We are now already familiar with the term AI. Ethical means one that is according to ethics. And ethics relates to moral principles that govern the behaviour and actions of a group or an individual. So, ethical AI is the theory behind developing computer systems that can perform tasks that require human intelligence. It focuses on how right, how fair and how just is the AI's output, outcome and impact.

Example: We know that AI can produce biased results if it is trained with biased data or if it does not do anything to avoid bias. So, any AI system that works on biased data is actually not ethical (or unethical). For example, Microsoft released a chatbot called Tay on Twitter in 2016, to learn by engaging people in dialogue through tweets or direct messages. However, trolls made Tay learn all negative words that spread hatred for women, a particular race and all sorts of biased data within hours. Tay was also taught to give repulsive and toxic responses. This aroused the anger of Twitter users and finally Microsoft had to silence Tay forever in even less than 24 hours.

An ethical AI system designed to solve a particular problem, uses unbiased data and is trained using the right learning model for the problem and is monitored to evaluate if the results produced by it are right and fair.

4.1.4 AI and Bias

Think about the following statements.

1. Most images that show up when you do an image search for 'doctor' are white men.
2. Images and videos of a 'doctor' show usually shows a male character and that of a 'nurse' is a woman.

3. All virtual assistants like Alexa, Siri, Google assistant, etc. work with female voices.
4. Computer vision systems frequently give an error while recognizing people of different colours.

Can you think of more such cases where we witness biasness towards a particular gender, colour, caste, religion, country, etc? Technically, AI bias occurs when an algorithm produces results that are systematically prejudiced towards a particular group of people and thus produces skewed output. AI applications may have built-in biases because ultimately, they are created by humans who have conscious or unconscious preferences that may go undiscovered until the algorithms are used at a massive scale. Let us understand three basic sources of bias in AI systems.

Data

Any AI system is just as good as the data we put into them. Naturally, adding biased or skewed data can never give fair results. Machines do not have the intelligence to differentiate between right or wrong training data.

For example, many-a-time, you must have noticed that Google or any other voice assistant replies in an American accent and at times gives wrong results when we give a command in Indian accent. This happens because the AI system is trained on recorded speech from white, upper-middle class Americans, making it difficult for the technology to understand commands from people not belonging to this category.

Algorithm

Algorithms can further amplify the biases injected by skewed data. For example, if you Google Image Search 'Teacher', then the image classifier algorithm trained on the images available in public domain shows more women teaching as opposed to male teachers. AI algorithms must be designed to maximize accuracy. But in this case, the AI algorithm may decide that all people teaching are women, despite the fact the training data has some images of men in the classroom. So, this is a case of gender bias in the AI system that becomes a part of AI algorithm due to biasness in data.

People

People who are developing the AI system, that is, engineers, scientists, developers, etc. are the next big source of bias. Humans design AI systems to give the most accurate results with the available data. So, data feeders that may be humans may again hit the success unknowingly. **Therefore, it is rightly said that ethics and bias are not the problem of the machine but that of the humans behind the machine.**

4.1.5 Towards Ethical and Trustworthy AI

Today, several companies worldwide are using algorithm-based 'emotional AI' hiring platforms to augment and lower the financial burden of their recruitment processes. While these systems may hire in a fair and unbiased manner, there are cases reporting women applicants being disproportionately rejected based on years of biased data in a male-dominated sector. Let us now see a few ways in which the world is trying to deal with this situation.

Regulating a More Ethical AI

In April 2021, the European Commission launched its first ever legal framework on AI to guarantee the safety and fundamental rights of people and businesses, while strengthening AI uptake, investment and innovation across the EU.

A new risk-based approach will not only set strict requirements for AI systems but also immediately ban AI systems which are considered to ‘be a threat to the safety, livelihood and rights of people’ including ‘systems that manipulate human behaviour, circumvent users’ free will and allow social scoring by governments’.

Later, in June 2021, Australia launched a similar AI ethics framework to guide businesses, governments and other organizations to responsibly design, develop and use AI.

Company and Organizational Engagement

Besides necessary regulation and standards, these days, organizations are building trust in AI technologies through cultural and educational programmes, risk assessments and third-party audit programmes. This would help them to move towards a prevention, detection and response framework (like anti-corruption, prevention of tax evasion frameworks).

Rights and Activist Groups

AI systems are yet to go a long way. In this journey, the relationship between people and technology will be redefined continuously. To ensure that our societies move forward on the path of ethics, it is critical that AI systems are human-centric. Civil rights and activist groups need to challenge the discourse and amplify the voices of those who are most adversely affected by new technologies. We must all ask the question—Is the technology necessary or is there any alternative? Only if the benefits outweigh the harm, we must check from whom the AI systems work.

Ensuring Data Privacy

Companies are now taking measures that are important to keep data safe and protect it from potential misuse. Privacy protection is now becoming an integral part of any ethical AI system. Data ethics and privacy in such systems is ensured by the following:

1. Encrypting the data.
2. Ensuring that the AI algorithm cannot learn from outside its dataset.
3. Using secure computation so that even the people who developed or work on those AI systems are not able to see or access the data.
4. Restrict reverse engineering the AI model so that no user can access data used for model training.

Diversify Your Team

Since bias is an important issue and results from an AI tool will be used by people from different racial, gender and economic identities, we need to make sure that there should be no bias. For this, build diverse teams to reduce the potential risk of bias falling through the cracks.

A diverse team has data scientists, business leaders, professionals with different educational backgrounds and experiences, such as lawyers, accountants, sociologists and ethicists. This would help the company to get their views on any sort of bias so that it could be mitigated in a timely manner. The team should continually analyse data and algorithms ensure fairness. There are many tools like Bias Analyzer which automate this process and also show the costs and benefits associated with a variety of possible mitigation actions.

4.1.6 Why is Ethical AI Important?

Since AI is used in areas like medicine, law enforcement, recruiting, data privacy, military defense or self-driving vehicles, it makes it mandatory that the AI system must produce accurate, transparent and understandable results that are in synchronization with the ethical standards and norms of our society. To understand the need of ethical AI systems, we must first understand that a biased or incorrect output from AI system that assist in law enforcement, job recruitment, defense work, self-driving vehicles can result in eroding people's privacy (by misusing data in unintended ways), and taking decisions that are impossible for people to understand and to access liability for damages when harm is caused.

Apart from being unbiased and accurate (as shown in Fig. 8.5), the credibility of AI systems also depends on how we choose to develop and use it. We can just hope that AI will be used for purposes that benefit mankind. But there is nothing that can stop a person from creating systems that can be disastrous for humanity.

In October 2019, researchers found that an algorithm used in US hospitals to predict which patients may require extra medical care heavily favoured white patients over blacks. This bias crept in because of the fact that white patients paid more bills and availed extra medical facilities for the same medical problem as compared to black.

Amazon's hiring algorithm also suffered from bias related issues. The algorithm was found to be biased against women. This was probably due to the fact that the system saw a greater number of male applicants than women.



FIGURE 8.5 An ethical system must be unbiased and accurate

Nonetheless, efforts are being made by numerous agencies, committees, coalitions and expert groups to ensure that we are on the right track to make AI systems do more to heal than harm. For example, the IEEE Global Initiative on Ethics of Autonomous and Intelligent Systems that

put out Ethically Aligned Design, a publication from 2016 outlines the guiding standards for developing and administering ethical AI solutions. Therefore, ethical AI aims to ensure that AI is doing good for the planet.

Going further on this, the European Commission High-Level Expert Group on AI identified seven guiding principles for creating Trustworthy AI. These guidelines are as follows:

1. AI systems should support human autonomy and decision making.
2. AI systems should be technically robust and safe to use. They must have a fallback plan that can be used when something goes wrong.
3. AI systems should ensure data privacy by protecting user's data, and providing adequate mechanisms to maintain the quality and integrity of the data.
4. AI systems should not work on biased data. The system should be clear, explicable and transparent about its underlying data and the models.
5. AI system should be trained with diverse, non discriminatory and fair data and models to avoid unfair bias.
6. AI systems should be designed to benefit everyone, now and in future. They must focus on societal and environmental wellbeing to be sustainable and environmentally friendly.
7. An AI system should have the accountability to ensure accurate, unbiased outcomes with utmost responsibility.
8. An AI system must give due importance to data privacy and data security. A system that provides design confidentiality, transparency and security into their AI programs must ensure that data is collected, used, managed and stored safely and responsibly.
9. An AI system must be accountable for its decisions. There must be someone in the organization who will be held responsible for any problematic situation.



Have you ever wondered how much information Google has about you? In an Android phone, Google has your entire Contact List. Google knows the names and numbers of people you talk to. Google knows your location. It knows where you go every day and on holidays. Google knows what email you write to whom. Google has your photographs (Google Photos), knows details about your friends (Google Hangouts). All this is done in lieu of free services it provides. But imagine what big risk it poses to our privacy.

4.1.7 Impact of AI on Jobs

Today, AI is on the tip of everyone's tongue and one of the most in-demand areas of expertise for job seekers. The World Economic Forum has predicted that AI and ML will displace 75 million jobs but would also create 133 million new ones in a few months. Some of our favourite sci-fi movies have shown that the advent of this technology has created a fear that

AI will one day make human beings obsolete in the workforce. As the technology is advancing, many tasks that were once executed by humans have now been automated




FIGURE 8.6 Emerging Areas in IT

A two-year study from McKinsey Global Institute suggests that by 2030, intelligent agents and robots could replace as much as 30% of the world's current human labour. That is, automation will displace between 400 and 800 million jobs and require approximately 375 million people to switch job categories entirely. However, job losses due to AI is a fear far from reality. It will just introduce a paradigm shift, similar to the one which occurred after the Industrial Revolution. When the Industrial Revolution took place, many feared that people will lose jobs to machines. But now we see that revolution has generated a lot more opportunities and jobs. Similarly, with AI, many professions will become obsolete and disappear, but other new emerging occupations in areas listed in Fig. 8.6 will become popular. Learning AI skills will help you gain a significant career advantage.

Jobs which will grow with the help of AI include the following.

Creative jobs will get refined and advanced by use of AI. These jobs are usually performed by professionals like artists, doctors and scientists. AI will not only enhance efficiency but also make these jobs less complex for humans. AI will probably not make human workers obsolete, at least not for a long time.

Tech jobs including programmers, data scientists, data analysts, big data engineers and all those who work on the creation and maintenance of AI systems will always be in high demand. However, a few tech jobs which are in demand today may become less common, while others may become more vital.

A few job titles which can be expected to appear by 2030 include Chief Bias Officer, Data detective, Man–Machine teaming manager, AI business development manager, AI assisted medical professional, AI tutor, to name a few.

In many sectors of business, AI has grown by 270% over the last four years. Rather than promoting the obsolescence of humans, AI will continue to drive massive innovation that will fuel many existing industries and could have the potential to create many new sectors for growth, ultimately leading to the creation of more jobs.

We must understand that although AI has done a lot to replicate the efficacy of human intelligence in executing certain tasks, AI programs suffer from major limitations. They can only solve one problem at a time. These systems become rigid, and are unable to think outside their prescribed programming. In contrast, humans have generalized intelligence which helps them solve any kind of problem, think abstractly to make critical judgement across all sectors. Moreover, AI can learn only with massive amounts of relevant data that is available at the right time. Many times, this data is protected under issues of privacy and security. In addition, data may have bias. And it is quite obvious that biased data will never give accurate results.

AI is becoming the standard in all businesses, not just in the world of tech.

Machines may also suffer because of the availability of limited computation and processing power. For example, the cost of electricity required to power one supercharged language model AI was estimated around \$4.6 million.

So, AI has the potential to ultimately create more jobs, not less. The question is not ‘humans or computers’ but ‘humans and computers’ involved in complex systems that advance industry and prosperity. Nobody would want to be behind the curve when it comes to AI. In fact, 90% of leading businesses have already invested in AI technologies. More than half of businesses that have implemented some manner of AI-driven technology report experiencing greater productivity.

AI and machine learning top the skills required by companies these days and jobs in this area are expected to increase by 71% in the next five years. AI is likely to have a strong impact on certain sectors in particular. These are discussed below

1. Medical: The medical industry has huge data which can be utilized to create predictive models related to healthcare and by physicians for diagnosis. And at a higher level, AI and automation will also help to eliminate disease and world poverty.

2. Automotive: Autonomous vehicles, autonomous navigation and manufacturing within the automotive sector have all been possible with AI technology.

3. Cybersecurity: During the pandemic, cyber frauds rose by 600% as hackers capitalized on people working from home, on less secure technological systems and Wi-Fi networks. AI and machine learning provide powerful tools to identify and predict threats. AI is also playing a vital role in the field of finance as it can process large amounts of data to predict and catch instances of fraud.

4. E-commerce: On e-commerce websites (like Amazon), AI is used to design chatbots, recommend products, use image based targeted advertising, and assist in warehouse and inventory automation.

5. HR applications: An HR department gets thousands of applications every day. AI tools can be used to reject 75% of resumes that do not fit the current context. Prior to AI, HR managers had to devote considerable time to filter applications to select suitable candidates. Data from LinkedIn shows that recruiters spend up to 23 hours looking over resumes for one successful hire.

6. Legal profession: Support services with document handling classification, discovery, summarization, comparison, knowledge extraction, and management use AI for efficiency and accuracy. Thus, the AI revolution will lead to a new era of prosperity, creativeness and well-

being. Humans will no longer be asked to perform routine, limited-value, full-time jobs. Rather, jobs will be flexible and offer selective premium services. There will be more jobs related to programming, robotics, engineering, etc. And, blue-collar and white-collar jobs will be eliminated.

4.2 AI as a Service (AIaaS)

In artificial intelligence as a service (AIAAS), companies use off the-shelf AI tools to implement and scale AI techniques at a very low price as compared to building a complete in-house AI system. When a software is provided as a service across the network, it usually uses cloud computing. Similarly, the idea of providing AI services using cloud computing has allowed even small business organizations to use cost effective solutions to improve their performance in a big way.

4.2.1 Factors Triggering Growth of AIaaS

In the last couple of years, AIaaS is growing fast as a business opportunity and many startups have joined to lead the AI revolution. Some major advancements in IT in the last few years have fostered the adoption of AIaaS. Some examples are as follows:

- 1.Availability of cloud platform with a variety of affordable options for enterprise data management.
- 2.Data storage technologies have improved and have become cheaper and reliable.
- 3.Streaming devices and IoT technologies generate massive amounts of data which when analysed can render vital information required to gain competitive advantage.
4. Availability of semi or fully automated data management, analytics and BI products changed the way business was analysed.

Moreover, some major factors that have motivated business to use AIaaS are listed below.

- 1.Expensive infrastructure
- 2.Lack of trained programmers
3. Very high charges charged by programmers for implementing AI in a particular organization
- 4.Many companies lacked sufficient data to analyse

4.2.2 The Growth of AIaaS

With cloud services becoming incredibly accessible, AI has seen a big boom. Now, AI can be made available to a greater number of companies and even those companies can collect and store any amount of data. Previously, companies were hesitant to build their own clouds to develop, test and utilize their own AI systems. But now with AIaaS, every other company can take advantage of data insights without making huge investment in talent and resources. Moreover, AIaaS has been a lucrative option as business can now easily accomplish the following goals.

- 1.Focus on its core area rather than getting involved in becoming an expert in machine learning.
- 2.Reduce operational and infrastructural cost to a big extent. A complex AI system may require many parallel machines and speedy GPUs. Not many companies can afford such huge

investment especially in its initial years. But with AIaaS, every small company can also harness the power of machine learning at significantly lower costs.

3. Minimize the risk of investment.
4. The time required to deploy a solution is short.
5. Pre-built algorithms make complex data management accessible to even small business organizations.
6. Availability of on-demand services empower businesses with a competitive edge.
7. Increase the benefits obtained by analysing data trends.
8. Develop a cost effective, flexible and transparent solution. AIaaS provides software to the organization that enables them to access machine learning capabilities using templates, pre-built models and drag-and-drop tools to assist developers in building a more customized machine learning framework. Companies no longer have to maintain a team with AI skilled professionals, train them, upgrade their knowledge and spend on areas that only partially support decision-making.
9. Improve its strategic flexibility. Companies can only pay for what they actually use. Though machine learning requires a lot of computing power to run, small companies may actually be using only in short amounts that too in spurts. So, in such a scenario, AIaaS is a huge saving.




FIGURE 8.7

10. Scalability is another big advantage of AIaaS. Now companies can delve into AI technology with smaller projects. If they find it beneficial, then probably, they can experiment with more services. At any point in time with the change in demands of the ongoing project, companies can scale up or down their demand for services. Figure 8.7 summarizes the advantages and shortcomings of providing AI as a Service.

4.2.3 Challenges of AIaaS

1.Reduced security: To avail services over cloud, companies have no option but to share their data with third-party vendors. Lot of data security techniques need to be incorporated to ensure that data is safe during storage and transmission and illegal access is strictly avoided. **2.**

Reliance: When working with third parties, the organization becomes dependent on them to satisfy their information needs. This results in delay in getting responses.

3.Reduced transparency: In AIaaS, organizations are paying to get service, but not the access. Third party vendors take the input and give the output. The entire process works like a black box. The organization has no idea about the inner workings (like algorithms used). This may result in confusion or miscommunication regarding the stability of data or the output.

4.Data governance: Every industry and organization has its own data policy. These policies may restrict data storage on a third-party cloud, thereby preventing a company from taking full advantage of AIaaS.

4.2.4 Vendors of AIaaS

There are multiple AI provider platforms offering various machine learning and AI services. Organizations can match their needs with the services and compare the cost at which those services can be accessed to select the offering that is best suited to their current requirements. Cloud AI service providers allow users to use GPU-based processing for intensive workloads along with their staffing and maintenance staff as well as catering to the hardware changes for different tasks. Some major vendors of AIaaS are Amazon Web Services (AWS), Microsoft Azure and Google Cloud Platform (GCP). Each of these vendors offer different types of bots, APIs and machine learning frameworks to cater to the organization's demands. Apart from these, countless start-ups are focusing on AIaaS. It is not uncommon for larger companies to purchase smaller companies to add the developed services to their portfolios. An organization decides on an AI service based on its goals, business size, available budget, technical capabilities of its in house teams and the amount of data that needs to be processed. MonkeyLearn is an AI platform that simplifies text analysis through intuitive, no-code tools. Users can either use pre trained model like survey analyzer to classify customer feedback by topic or build customized machine learning models to detect sentiment, keywords, and topics in the data. Once developed, these models can be integrated with apps through point-and-click integrations or via the API. IBM Watson hosts a suite of AI tools and offers many pre-built applications, like Watson Assistant and Watson Natural Language Understand to build virtual assistants and to perform advanced text analysis tasks respectively. Similarly, IBM Watson Studio can be used to build, train and deploy machine learning models across any cloud. With these offerings even users with no expertise on machine learning or data science can make the most of their data. Microsoft Azure is Microsoft's public cloud computing platform that provides a range of AI and machine learning solutions for developers. For example, Azure Cognitive Services can be used to add computer vision or text extraction capabilities to apps using APIs. Similarly, Azure Bot Service can be used to incorporate any type of bot in the application. Google Cloud ML is Google's AI platform that basically facilitates data scientists and developers working with big data to create and deploy machine learning projects. AutoML is being extensively used to train custom machine learning models for text analysis, image classification, translation and more. Data in the datasets can be visualized, 'what-if tools' can be used and metrics can be analysed to assess performance of the model. AIPaaS or AI Platform-as-a-Service is an end-to-end solution like a cloud platform,

using which business organizations can utilize required services on a pay-per-use or pay-per-service basis. With AIPaaS, users can use third-party APIs that provide a complete, intelligent data management platform. The Software-as-a-Service (SaaS) market is valued at \$133 billion and the advanced technology platforms-as-a-service market globally is forecast to reach about \$11 billion by 2023 and surpass '\$88,500 million by the end of 2025.'

4.3 Recent Trends in AI

In this section, we will read about some latest trends in AI. Extensive research is being done in the terms discussed here.

These days, a lot of research is being done to develop collaborative systems in which humans and machines (with AI) can become effective partners by overcoming each other's limitations and augmenting human abilities. There are many examples of machine-human collaboration as follows.

1. The computer game Foldit uses machine-human collaboration fold simulated proteins. The folding task is performed to understand how real-world proteins involved in the causes of human disease are formed. In the game, both the player as well as AI performs operations. The AI is put to use where it excels and humans' intuition and imagination is applied where they exceed machines.
2. Two amateur chess players and an AI system worked collaboratively and won against a field of supercomputers and grandmasters.
3. In 2014, Knowledge Ventures, a Japanese venture capital company, elected an AI system to its board of directors, to realize the benefits of machine-human collaboration at the highest levels of business.
4. In the military, UAV drones are being used in which a machine takes the front lines with human provides virtual support.
5. The Talos suit in the US military augments a front-line human with a power exoskeleton. The exoskeleton's integrated heating cooling, vital monitoring and heads-up display (HUD) collaborate with humans much similar to that seen in the Ironman movies.

While many people believe that humans and machines can be good partners, others, however, fear that relying too much on AI will result in a massive job loss and making common man people dumb and lazy.

However, the good news is that research involving 1,500 companies says that many companies that have used AI to automate processes and displaced employees will see only short-term productivity gains. Better performance is expected only when humans and machines work together to actively enhance each other's complementary strengths- leadership, teamwork, creativity, social skills of humans and speed, scalability and quantitative capabilities of the machines.

AI can boost our analytic and decision-making abilities and heighten creativity.

Humans Assisting Machines

In a human-machine collaborative system, humans need to perform three important roles. First, they must train machines to perform a particular task; explain the outcome of that

task (especially when results are counterintuitive or controversial) and sustain the responsible use of machines so that machines do not harm mankind. For example, the European Union's new General Data Protection Regulation (GDPR) empowers consumers with the right to get an explanation for any algorithm-based decision (like reason for acceptance/rejection of a loan). This feature will not be a big step in delivering responsible AI but will also create about 75,000 new jobs to administer the GDPR requirements.

Last but not the least, companies also need employees who can continuously monitor to ensure that AI systems are functioning properly, safely and responsibly. For example, a self-driving car must not hit objects and any living being. Similarly, tech companies including Apple use AI to collect personal details about users to improve the user experience, but collecting too much data can compromise privacy, anger customers and even go against the law. Therefore, employees make sure that AI uses data only for analysis and the privacy of individual users is not violated.

Points to remember: Humans must also train AI systems to interact with humans in the best possible way. For example, all virtual assistants were given extensive training to develop to react as a confident, caring and helpful personality without being bossy. Such training requires inputs from a team of experts (even including a poet, a novelist and a playwright). These days, AI assistants are also being trained to display sympathetic feelings. For example, the start-up Koko, an offshoot of the MIT Media Lab, is developing an empathetic assistant. For example, if a user had a bad day, then Koko would feel sorry for it, ask for more information and offer advice to help the person see his issues in a different light.

4.3.2 Machines Assisting Humans

Smart machines are helping humans enhance their abilities in three ways—amplify, interact and embody. AI systems boost human creativity, analytic and decision-making abilities by providing the right information at the right time. They can also interact with company's employees and customers in more effective ways. Even at homes, digital assistants are assisting humans to perform their instructions.

SEB, a Swedish bank, uses Aida, a virtual assistant to interact with millions of customers and assist them in opening a bank account, making cross-border payments and interpreting their voice tone (satisfied or frustrated) to provide better service later. In 30% of the cases when the system is not able to resolve an issue, it directs the call to a human customer-service representative and then monitors that interaction to learn how to resolve similar problems in the future.

Case study Question: Collect information about how Autodesk's Dreamcatcher AI enhances the imagination of even exceptional designers.

Embodying

AI systems are embodied in machines (like robots) to augment the capabilities of a human worker. Using their sophisticated sensors, motors and actuators, these machines can recognize people/objects and work safely in factories, warehouses and laboratories. Such robots also known as cobots handle repetitive actions that require heavy lifting or some

dangerous tasks. In such a collaborative environment, humans are free to do other complementary tasks that require dexterity.

For example, Hyundai has extended the concept of cobot with exoskeletons.

Exoskeletons are wearable robotic devices, which adapt to an industrial worker's location in real time to enable him to perform his job with superhuman endurance and strength. Similarly, Mercedes-Benz uses cobot arms to customize cars as per user's expectations according to the real-time choices consumers make at the dealer's showroom.

4.3.3 Algorithmic Game Theory and Computational Social Choice

These days, (especially enhanced due to pandemic) with increased Internet connectivity and speed, more people prefer to play games like chess, checker, poker and solitaire. These games are played using a clear set of defined rules. So, there was a need to embed algorithms in machines that allowed machines to play ethically according to pre-defined rules. This led to the onset of game theory. In a multi-agent situation (when more than one person is involved in solving a logical problem), game theory is used to choose from a set of options knowing that our choice would affect the choices of the opponent and their decision would affect our choices.

Von Newman is credited for the invention of game theory. We can understand the concept of game theory by studying the five categories.

1. Cooperative vs Non-cooperative Games: In cooperative games, participants can establish alliances to increase their chances of winning the game. Correspondingly, in a non cooperative game, participants cannot form any alliance.

2. Symmetric vs Asymmetric Games: In a symmetric game, each participant has the same goal(s). However, their strategy to win may be different. In contrast to this, in an asymmetric game, the participants have different (and at times conflicting) goals.

3. Perfect vs Imperfect Information Games: In perfect information games like chess, all the players can see what moves the other player(s) are making. But, in an imperfect information game (like a card game), other players' moves are hidden.

4. Simultaneous vs Sequential Games: In simultaneous games, multiple players can take actions concurrently. In a sequential game like in a board game, each player is aware of the other players' previous actions.

5. Zero-sum vs Non-zero Sum Games: In zero sum games, a player gains something that causes a loss to other players but in a non-zero sum game, multiple players get the advantages when another player gains.

When using AI techniques in game theory, two concepts are very important—Nash equilibrium and inverse game theory.

Nash Equilibrium

The Nash equilibrium is a condition in which all the players playing the game accept that there is no better solution to the game than the current situation. No player would gain anything after changing their current strategy. So, any new decision made by the players will not result in anything better.

Example: A typical example of Nash equilibrium is the Prisoner's Dilemma. If two criminals get arrested and are kept in confinement without being allowed to communicate with each other, then look at Fig. 8.21 to understand the possibilities.

- If either of the two prisoners confess that the other committed a crime, the first will be set free and the other will be sent to prison for 10 years.
- If none of them confess, then they will be sent to prison only for 1 year.
- If both confess, then both will be sent to prison for 5 years. FIGURE 8.21 Nash equilibrium

		PLAYER B	
		COOPERATE	DEFECT
PLAYER A	COOPERATE	A: 1 year jail B: 1 year jail	A: 10 years jail B: 0 years jail
	DEFECT	A: 0 years jail B: 10 years jail	A: 5 years jail B: 5 years jail

FIGURE 8.21 Nash equilibrium

In this case, the Nash equilibrium is reached when both criminals betray each other.

Inverse Game Theory

Game theory is used to understand the dynamics of a game to optimize the possible outcome of its players. In contrast to this, inverse game theory designs a game based on players' strategies and aims. It is key to designing AI agents' environments.

Practical Examples A generative adversarial network (GAN) is a machine learning (ML) model in which two neural networks compete with each other to become more accurate in their predictions.

Nash equilibrium can be attained more easily in symmetric than asymmetric games. But, in the real world, asymmetric games are more common. GANs contain two models—a generative model and a discriminative model (Fig. 8.22).




FIGURE 8.22 GAN Architecture [3]

Generative models accept some features as input to examine their distributions and understand how they have been produced. Discriminative models take the input features to predict the class to which our sample belongs. For example, in a GAN, the generative model uses input features to create new samples that resemble the main features of the original samples. The original samples and the newly generated samples are then passed to the discriminative model to identify which samples are genuine and which are fake. A common application of GANs is to generate images (as given in Fig. 8.23) and then distinguish between real and fake ones.



Credit: doglikehorse / Shutterstock

FIGURE 8.23 Generate Images using GAN Models

Even in game playing, GANs are used extensively. The players (the two models) challenge each other. While one player creates fake samples to confuse the other, the other tries to identify the right samples. The game is played iteratively until Nash equilibrium is reached. In each iteration, the learning parameters are updated to minimize loss.

4.3.4 Multi-Agents Reinforcement Learning (MARL)

Reinforcement learning (RL) is used to make an agent (model) learn through the interaction with its environment. RL was initially developed based on Markov decision processes. This means that the agent placed in a stochastic stationary environment tries to learn a policy through a reward/punishment mechanism. The agent will finally converge to a satisfactory policy.

However, this does not hold true when multiple agents are placed in the same environment. Therefore, instead of depending only on learning through interaction between the agent and the environment, now it is also dependent on the interaction between agents.

Modelling systems with multiple agents is not a trivial task as increase in the number of agents exponentially increases the number of possible ways in which they can interact with each other.

Therefore, modelling multi-agents reinforcement learning models can be done using mean field scenarios (MFS) to reduce its complexity by making the assumption a priori that all agents have similar reward functions.

Example: To improve the traffic flow in a city, AI-powered self driving cars can perfectly interact with the external environment. But things can get complicated if these cars think as a group. In such a case, a car might get in conflict with another when both of them find it most convenient to follow a certain route.

The above situation can be easily modelled using game theory. Our cars would be different players and the Nash equilibrium would be the point between the collaboration between the different cars.

4.3.5 Neuromorphic Computing

Neuromorphic computing is a technique in which a computer (both hardware and software) is modelled based on the architecture of the human brain and its nervous system. It is a science that draws concepts from multiple disciplines including computer science, biology, mathematics, electronic engineering and physics.

The main aim of neuromorphic computing is to create artificial neural systems inspired by biological structures. Such a structure would enable the machine to behave like a human brain in the following ways.

First, to learn, retain information and even make logical decisions as the human brain makes. Second, to acquire new information and work like a human brain. Thus, a computer needs to be transformed into a cognition machine.

Did you know that our brain uses roughly 20 watts of power on average, which is about half that of a standard laptop?

How Does Neuromorphic Computing Work?

Traditional neural network and machine learning algorithms are best suited for existing algorithms. They either provide fast computation or focus on using low power and there is always a trade-off between the what to achieve. Neuromorphic systems perform fast computation while consuming less power consumption. Additionally, they also resemble human brain as they have the following features:

1. They are massively parallel, to handle multiple tasks at once.
2. They are event-driven, as they respond to events dynamically based on changing environmental conditions.
3. They are economical in terms of power used.
4. They are flexible as they are highly adaptable and are able to generalize.
5. They are strong and fault-tolerant, since they produce results even when some component(s) have failed.

Neuromorphic computing achieves this brain-like function and efficiency by building artificial neural systems that implement 'neurons' (nodes that process information) and 'synapses' (the connections between those nodes) to transfer electrical signals using analog circuitry. This structure controls the amount of electricity being passed from one node to another to mimic the varying degrees of strength that usually occurs in a human brain. Such a system of neurons and

synapses that transmit these electric pulses is known as a spiking neural network (SNN). This facility was not available in traditional neural networks (that uses digital signals).

Neuromorphic systems deviate from the Von Neumann architecture and introduce a new chip architecture that collocates memory and processing together on each individual neuron. Hence, there is no need to have separate memory unit (MU), central processing unit (CPU) and data paths. Having separate parts requires information to be repeatedly moved back and forth between different components to complete a given task, creating a bottleneck for time and energy efficiency. This is known as the von Neumann bottleneck.

By bringing memory on the neuromorphic chip, information can be processed much more efficiently, thereby making the chip more powerful and efficient.

Neuromorphic Computing and Artificial General Intelligence (AGI)

We have seen that the term ‘artificial general intelligence’ (AGI) exhibits intelligence equal to that of humans has not yet been achieved and may not even be reached with the traditional computers. But with neuromorphic computing, researchers are feeling optimistic about it. Example: The Human Brain Project, featuring the neuromorphic supercomputer SpiNNaker produces a functioning simulation of the human brain and is one of the hottest research projects in AGI.

A machine is said to achieve AGI when the machine can do the following: Reason and make judgments in an uncertain situation Plan Learn Communicate using natural language Represent knowledge Integrate the above skills to fulfil a goal An AGI may even have the ability to imagine, gain from experience and be self-aware. To confirm whether a system has developed AGI or not, it can be made to pass the Turing Test or the Robot College Student Test, in which a machine enrolls in classes and obtains a degree like a human would. **Debatable Issue**

Even if a machine develops human intelligence, then a question arises whether it should be handled ethically and legally. While some researchers argue that it should be treated as a nonhuman animal in the eyes of the law, others say that still a lot has to be done to inculcate consciousness in machines.

Comparing the Human Brain with Neuromorphic Computing

Brains need very less energy than most supercomputers. While a human brain uses about 20 watts, the Fugaku supercomputer needs 28 megawatts (about 0.00007% of Fugaku’s) power supply.

Supercomputers need elaborate cooling systems but our brain sits at 37°C in between bones. Supercomputers can make millions and billions of complex calculations in second(s) but the human brain is very adaptable. A single brain can write poetry, identify a familiar face out of a crowd, drive a car, learn a new language, instantly make decisions, respond to events and do much more.

While traditional computers designed on von Neumann architecture are largely serial, brains use massively parallel computing. Brains are also more fault-tolerant than computers. Therefore, harnessing techniques used by a human brain will be the key to making more powerful neuromorphic computers in the future.

So How Can You Make a Computer That Works Like the Human Brain?

To understand how neuromorphic computers work, let us first recapitulate how the brain works.

When we feel hot or cold, neurons (a type of nerve cell) carry messages to the brain. To transfer messages, several neurons release chemicals across a gap called a synapse. This happens until the message reaches the brain. Once the brain identifies the sensation, it sends an appropriate message to the point which felt the sensation using neurons.

An action potential is triggered either through sending multiple inputs at once (spatial), or by building multiple inputs over time (temporal). The brain transfers information quickly and efficiently through the huge interconnectivity of synapses where one synapse might be connected to 10,000 other synapses.

Neuromorphic computers work like human brains using spiking neural networks (SNN). While transistors used in conventional computing can be either on (1) or off (0), SNN can produce more than one output and transfers information in a temporal and spatial way as the human brain does.

Neuromorphic systems can be either digital or analogue. Synapses are controlled either by software or memristors. Memristors can store a range of values instead of storing just one and zero. Like the human brain, they can vary the strength of a connection between two synapses to allow the brain-based systems to learn.

Uses of Neuromorphic Systems

As of now, to perform heavy computational tasks, devices like smartphones hand off processing to a cloud-based system. The cloud processes the query and sends the output to the device. This scenario will change with neuromorphic systems. In such a system, instead of passing queries and responses back and forth, computations would be done within the device itself.

The current generation of AI is rules-based. It is extensively trained on datasets to generate a particular outcome. However, the human brain does not work like this. Researchers are trying hard to make the next generation of artificial intelligence deal with more brain-like problems. For example, constraint satisfaction, in which a system has to find the optimum solution to a problem with a lot of restrictions. Neuromorphic systems also perform well with noisy and uncertain data.

Examples of Neuromorphic Computer Systems Used Today 64 Intel's neuromorphic chips named Loihi are used to make an 8 million synapse system called Pohoiki Beach, consisting of 8 million neurons (may soon reach 100 million neurons in the near future). They are being used to create artificial skin and also develop powered prosthetic limbs. IBM's neuromorphic system, TrueNorth, launched in 2014, consists of 64 million neurons and 16 billion synapses. IBM in a partnership with the US Air Force Research Laboratory is trying to create a 'neuromorphic supercomputer' known as Blue Raven for creating smarter, lighter, less energy-demanding drones. The EU-funded Human Brain Project (HBP), a 10-year project that started in 2013, has led to two major neuromorphic initiatives, SpiNNaker and BrainScaleS. In 2018, the largest neuromorphic supercomputer of that time, a million-core SpiNNaker system went

live. It is hoped that this model would soon scale up to one million neurons. BrainScaleS also has a similar goal and its architecture is now on its second generation, BrainScaleS-2.

Challenges to Using Neuromorphic Systems

Shifting from Von Neumann to neuromorphic computing is not that easy and straightforward. For example, when analysing visual input, traditional systems treat them as a series of individual frames, but a neuromorphic processor would encode this information as changes in a visual field over time. New programming languages need to be developed to utilize the change in the underlying architecture. New kinds of memory, storage and sensor technology need to be created to take full advantage of neuromorphic devices. This may even change the entire process of creating and integrating hardware and software components.

Where AI Is Heading Today?

4.4 Expert System

An expert system is a computer program that solves complex problems in a particular domain, at the level of extra-ordinary human intelligence and expertise. That is, it is designed to provide decisions just like a human expert. Expert systems are the computer applications developed to solve complex problems.

Expert systems are designed for a specific domain, such as medicine, science, etc. To perform its task, the expert system analyses user queries to extract knowledge from its knowledge base using the reasoning and inference rules.

The first expert system was developed in 1970. This system makes decisions for complex problems using both facts and heuristics like a human expert. The performance of an expert system depends on the knowledge stored in its knowledge base. More the amount of knowledge stored in the KB, better is the performance of the expert system. **The suggestion of spelling errors while typing in the Google search box is a perfect example of expert systems used these days by all users.**

An expert system is not used to replace the human experts; instead, it is used to assist the human in making a complex decision. The working of an expert system can be understood through the block diagram given in Fig. 9.1. The main components of an expert system include user interface, knowledge base and inference engine.




FIGURE 9.1 Working of an expert system

4.4.1 Popular Examples of the Expert System

1.DENDRAL: It is a chemical analysis expert system. DENDRAL was used in organic chemistry to detect unknown organic molecules with the help of their mass spectra and knowledge base of chemistry. The expert system studied a substance's spectrographic data to predict its molecular structure.

2.MYCIN: It is one of the earliest backward chaining expert systems that was designed to diagnose blood clotting diseases and find the bacteria causing infections like bacteraemia and meningitis. MYCIN could also recommend antibiotics and drugs based on the person's physical conditions.

3.PXDES: It is an expert system that determines the type and level of lung cancer by analysing picture(s) from the upper body.

4.CaDeT: CaDeT is an expert system that works like a diagnostic support system to detect cancer at early stages.

5.R1/XCON: The expert system could select specific software to generate a customized computer system based on user's requirements. **6. DXplain:** It is a clinical support system that could suggest the presence of one or more diseases based on the findings of the doctor.

4.4.2 Characteristics of an Expert System

Expert systems are among the first truly successful forms of artificial intelligence (AI) software. Let us read about their distinguishing characteristics (refer Fig. 9.2).




FIGURE 9.2

1.High performance: An expert system provides high performance for solving any type of complex problem of a specific domain with high efficiency and accuracy. The success of an expert system depends on its knowledge base. So, to get accurate knowledge continuously, the knowledge base is updated regularly.

2.Understandable: The results of an expert system are easily understood by users. The system takes input in human language and provides the output in the same language.

3.Reliable: An expert system generates efficient and accurate output. This makes it a reliable system.

4. Highly responsive: An expert system generates results for any complex query within a very short period of time, making it highly responsive. The total time by an expert system should be less than that taken by a human expert.

5. No memory limitations: An expert system can store and exploit huge amounts of data.

6.Expertise in a domain: There are many human experts in a particular domain and each of them specialize in a different skill and has different experiences. When we put the knowledge gained from these human experts into the expert system, the results are provided by utilizing all the facts and knowledge

7.Not affected by emotions: An expert system is not affected by human emotions such as fatigue, anger, depression, anxiety, etc. This helps them to give a consistent performance every time they are used.

8. Not biased: To respond to any query, the expert systems check and consider all the available facts.

9. Reduce cost: Human experts charge a heavy amount for examining, diagnosing and treating a patient. But with an expert system, multiple experts can be consulted at much cheaper rates.

10. Non-perishable: Human experts are perishable as they have a limited life span but an expert system is permanent. Once developed, it can be used without any limitations of expiry.

11. Intelligent: These systems use knowledge base and inference engine to solve complex problems. This is done by deducing new facts from existing facts that are represented as if-then rules.

4.4.3 Components of an Expert System

An expert system mainly consists of three components—user interface, inference engine and knowledge base (as shown in Fig. 9.3)




FIGURE 9.3 Components of an Expert System

User Interface

The user interacts with an expert system through the user interface. It is through this interface that the user enters queries in a readable format. The interface then passes it to the inference engine. After getting a response from the inference engine, it displays the output to the user. The user interface is designed to be very simple and intuitive so that even a non **expert user can communicate with the expert system to find a solution to a complex problem.**

The user interface also explains how the expert system derived a particular conclusion. This explanation may be text displayed on screen, verbal narrations in natural language or a listing of rules on the screen. Through this, the user interface makes it easy to trace the credibility of the conclusions. An effective user interface must have the following features:

1. Help users to accomplish their goals in the shortest possible way.
2. Be easy to customize as per user's requirements.
3. Make efficient use of user input.

Inference Engine (Rules of Engine)

The inference engine is the main processing unit and hence known as brain of the expert system. The engine derives conclusion or deduces new information by applying inference rules to the knowledge base. There are two types of inference engine:

Deterministic Inference Engine

This engine uses facts and rules to draw conclusions. These conclusions are thus assumed to be true

Probabilistic Inference Engine

This inference engine contains uncertainty in conclusions as it is based on the probability. Efficient procedures and rules used by the inference engine helps to formulate correct, flawless solutions. In case there is a conflict when multiple rules are applicable to a particular case, the inference engine resolves the problem by selecting the best rule that is applicable in the current context. Inference engines can also include explanation and debugging abilities. The explanation module, if incorporated, can be used by the expert system to give the user an explanation about how it reached a particular conclusion.

The inference engine generally uses two strategies for acquiring knowledge from the knowledge base and deriving solutions. These are as follows:

1. Forward Chaining: As discussed earlier, forward chaining starts from the known facts and rules (refer Fig. 9.4). After applying the inference rules, conclusions are drawn and then added to the known facts. Forward chaining strategy facilitates an expert system to answer the question, ‘What can happen next?’ For example, forward chaining can be used to predict stock prices after changes in interest rates.




FIGURE 9.4 Forward Chaining

2. Backward Chaining: Backward reasoning, on the other hand, starts from the goal and works backward to prove the known facts (as shown in Fig. 9.5). This strategy enables an expert system to answer the question, ‘Why this happened?’ If an event has already taken place, then the inference engine can be used to find out which conditions occurred in the past for this result. Thus, backward chaining strategy is used for discovering cause or reason. For example, diagnosis of blood cancer in humans.




FIGURE 9.5 Backward Chaining

Knowledge Base

Before going to knowledge base, let us first understand what knowledge is. From junior classes, we are taught that data is a collection of facts. Data, when processed, gives us information. Data, information and past experience combined together is known as knowledge.

We have read earlier that the knowledge base stores knowledge acquired from different experts of a particular domain. More the knowledge in the knowledge base, more precise is the result from the expert system.

The knowledge base contains information and rules of a particular domain or subject. The two main components of a knowledge base are factual knowledge and heuristic knowledge.

Factual knowledge is based on facts and accepted by knowledge engineers in that domain.

Heuristic knowledge is based on practice, probability, evaluation and experiences.

Knowledge in a knowledge base is usually represented using the If-else rules and is acquired by extracting, organizing and structuring the domain knowledge under the guidance of human experts. This knowledge is domain-specific and of high quality



FIGURE 9.6 Creation of Knowledge Base

CASE STUDY: Working of MYCIN as an expert system.

Step 1: Knowledge base is created with expert knowledge about bacterial infection—their causes, symptoms, treatment, etc. (refer Fig. 9.6).

Step 2: The knowledge base is regularly updated. In order to test it, the doctor gives the expert system a new problem of identifying the presence of the bacteria by inputting the details of a patient, including the symptoms, current condition, and medical history.

Step 3: The system prompts the patient to provide general information like gender, age, etc.

Step 4: The system applies if-then rules using the inference engine and the facts stored within the knowledge base.

Step 5: The output is displayed to the patient through the user interface.

4.4.4 Participants in the Development of Expert System

In Fig. 9.7, we see that there are three primary participants in an expert system:




FIGURE 9.7 Participants in an expert system

Expert

The performance of any expert system depends on the knowledge fed in its knowledge base. This knowledge is provided by human experts who are specialized in that specific domain.

Knowledge Engineer

The term 'knowledge engineering' refers to the process of building an expert system and its practitioners are called knowledge engineers. The main task of a knowledge engineer is to ensure that the computer possesses all the knowledge required to solve a problem. For this, he/she gathers the knowledge from the domain experts and then represent that knowledge in the format accepted by the knowledge base.

End User

The end user who uses the expert system may or may not be well versed in that domain. He/she just uses the expert system to find solution or advice for his queries, which are usually complex

4.4.5 Capabilities of the Expert System

An expert system can do the following things:

- 1. Advise:** The system can give its advice to users for any query —simple or complex— that fits its domain.
- 2. Provide decision-making capabilities:** An expert system can be designed to make complex decisions in any field, be it finance, medical science or any other.
- 3. Demonstrate a device:** It can be used to demonstrate features, specifications, usage of any new product.
- 4. Problem-solving:** It can be used to solve complex problems.
- 5. Explaining a problem:** An expert system can be used to provide a detailed description of an input problem.
- 6. Interpreting the input:** It can interpret the input given by the user.
- 7. Predicting results:** An expert system can predict results using past data.
- 8. Instructing and assisting human in decision-making**
- 9. Diagnosing:** Expert systems, especially in the field of medical science, are used to diagnose any disease.
- 10. Justifying the conclusion:** Expert systems not only find solutions to complex real-world problems by drawing conclusions from known facts but also justify those conclusions.
- 11. Suggesting:** The best part is that an expert system also suggests alternative options to a problem. **However, expert systems cannot do the following:**

1. Substitute human decision-makers
2. Possess human capabilities
3. Produce accurate output for inadequate knowledge base
4. Refining knowledge base
5. Use emotions to make decisions

4.4.6 Advantages of Expert Systems

1. Results obtained by an expert system are highly reproducible.
2. They can be used in risky places where it is not safe for humans to work.

3. They give accurate results, especially when the knowledge base is updated regularly.
4. They give consistent performance at a fast speed as these systems are not affected by emotions, tension or fatigue.
5. They are easily available due to mass production of software.
6. They are affordable as the production cost is reasonable.
7. They work at great speed and reduce human effort.
8. They are capable of explaining how the solution was obtained.
9. They improve the quality of decisions.
10. They gather scarce expertise knowledge in good amount and use it efficiently.

4.4.7 Limitations of Expert Systems

1. Results obtained from an expert system may get wrong if the knowledge base contains wrong/obsolete information.
2. Each problem is different; therefore, the solution from a human expert may be strikingly different from that of an expert system.
3. They lack creativity for different scenarios, especially in extraordinary situations.
4. They have high development and maintenance costs.
5. Knowledge acquisition is a challenging task.
6. Success depends on knowledge acquired from human experts.
7. These systems cannot learn from themselves and hence require manual updates.
8. They cannot produce correct results from less knowledge.
9. They need excessive training.
10. They do not have any emotions.
11. They do not possess any common sense.
12. They are developed only for a particular domain.
13. They are not capable of making decisions in extraordinary situations.
14. They work on Garbage-in Garbage-out (GIGO) principle. This means that if there is an error in the knowledge base, the expert system will give incorrect results.
15. They have high maintenance costs.

4.4.8 Applications of Expert Systems

Expert systems are one of the prominent research domains of AI. Some common applications of expert system are as follows:

1. Designing and manufacturing devices such as VLSI systems, camera lenses and automobiles.
2. Publishing relevant knowledge content to users. An advisor and a tax advisor are two popularly used expert systems in this domain.
3. Finance industries use expert systems to detect any type of possible fraud, suspicious activity, stock market trading and advise bankers on whether they should approve the loans or not.
4. Expert systems are extensively used for medical diagnosis and treatment. In fact, it was the first area where these systems were used.

5. Expert systems are used for planning and scheduling tasks for achieving the project's goal. Companies also use expert systems for scheduling airlines, cargo customer order, computer resources and various manufacturing tasks.
6. Tracking the progress of a software development project.
7. Monitoring applications compare data continuously with observed systems or with prescribed behaviour to detect leakage in long petroleum pipelines using expert systems.
8. Process control systems use expert systems to control a physical process based on continuous monitoring.
9. The automobile and electronics industries use expert systems to identify faults in vehicles and computers manufactured.
10. Diagnosing complex electronic, electromechanical and diesel-electric locomotive systems.
11. Forecasting crop damage or in general predicting results of any experiment/activity.
12. Identifying the structure of a chemical compound.
13. Assessing geologic structure from dip meter logs and space structure through satellite and robot.
14. Evaluating civil cases, product liability, employee performance, etc. using logs.
15. Detecting virus attacks in a computer system.
16. Repair and maintenance projects.
17. Optimizing warehouses.
18. Configuring manufactured objects.

4.4.9 Expert System Technology

Expert systems technologies include the following.

1. Expert System Development Environment: The development environment of an expert system includes hardware and tools like

1. Workstations, minicomputers, mainframes.
2. High level symbolic programming languages (like LISP Programming (LISP) and PROgrammation en LOGique (PROLOG)).
3. Large databases.

2. Tools: The cost of developing an expert system gets reduced to a large extent with the help of tools that help developers by providing,

1. Powerful editors
2. Debugging capabilities with multi-windows
3. Rapid prototyping
4. Having in-built definitions of model, knowledge representation and inference design.

3. Shells: A shell is an expert system without a knowledge base. It facilitates developers by providing modules for knowledge acquisition, inference engine, user interface and explanation facility. Two popularly used shells are JESS and Vidwan. While Java Expert System Shell (JESS) provides Java API for creating an expert system, Vidwan, enables knowledge encoding in the form of IF-THEN rules.

4.4.10 Development of Expert Systems

The process of developing an expert system is iterative. The steps to develop such a system are described below.

Identify Problem Domain

The first step to develop any system is to identify the problem that it intends to solve. Ensure that this problem is solvable by an expert system. If the problem can be solved, identify the human experts in the task domain and analyse the cost effectiveness of the system. Knowledge engineer(s) and domain experts work in coherence to define the problem and list the characteristics of the system.

Design the System

To design the expert system, first identify the technology (hardware, tools and shells) that will be used. Understand how well the system will be integrated with the other systems and databases. Prepare a framework for representing concepts in domain knowledge.

Develop the Prototype

In this step, the knowledge engineer acquires domain knowledge from the human experts and represent it in the form of If-THEN-ELSE rules. The knowledge expert also determines how heuristic knowledge could be integrated in the reasoning process and what type of explanation would be useful.

Test and Refine the Prototype

The knowledge engineer tests the prototype to identify any deficiencies in performance or results by using sample cases. Testing is also done along with end-users to unleash additional errors or discrepancies. The results of testing are then used to refine the system.

Develop and Complete the System

Once the expert system is developed and tested, it is then tested along with other elements of its environment, including end users, databases and other information systems to ensure that it interacts in the prescribed manner and smoothly integrates with them to render the desired functionality. At this step, the expert system is clearly documented and users are trained for its usage. An expert system is also used for information management and help desks management.

Maintain the System

Maintenance is an ongoing process. In this phase, the knowledge base is regularly reviewed and updated. New requirements from the users or blending with new interfaces of other information systems as those systems evolve are steps that need to be done time to time

TABLE 9.1 Traditional System vs Expert System

Conventional System	Expert System
Knowledge and processing are combined in one unit.	Knowledge database and the processing mechanism are two separate components.
The program does not make errors (unless there is an error in input data or programming logic).	The expert system may give less accurate results.
The system is operational only when it is fully developed.	The expert system is optimized on an ongoing basis and can be used even with a small number of rules.
Program is executed step-by-step according to a fixed algorithm.	Execution is done logically and heuristically.
The problem expertise is encoded in both program as well as data structures	Problem-related expertise is encoded only in data structures.
It needs full information.	It can be functional even with insufficient information.
It uses data more efficiently.	It utilizes knowledge efficiently.
It gives solution without explanation.	It gives solution and explains how it was obtained.
It does not represent knowledge symbolically for computations.	It symbolically represents knowledge for computations.

TABLE 9.2 Human expert vs Expert System

Human Expert	Artificial Expertise
Perishable	Permanent
Difficult to transfer	Transferable
Difficult to document	Easy to document
Unpredictable	Consistent
Expensive	Cost-effective system

4.5 Internet of Things

Although the concept was not named until 1999, the Internet of Things (IoT) has been in development for decades. IoT means a system of interrelated computing devices, machines, objects, animals or people that have unique identifiers for identification and also the ability to transfer data over a network without requiring human-to-human or human-to-computer interaction (as shown in Fig. 9.8).



Credit: 363868 / Shutterstock

FIGURE 9.8 Interconnected Things

If the thing, in the IoT, is a person, then he can have heart monitor implant; if it is an animal, it can have a biochip transponder, if it is an automobile, it can have sensors to generate alert alarms (e.g., when the pressure in a tyre is too low) and if it is any other natural or man-made object, then it is assigned an IP address and provided with the ability to transfer data over a network.

The huge availability of address space in IPv6 has led to the development of the IoT. With more IP addresses available, more devices can be connected to exchange information with each other.

Today, IoT is being widely used for precision agriculture, building management, healthcare, energy and transportation. The first Internet appliance was a Coke machine at Carnegie Mellon University in the early 1980s when the programmers succeeded in connecting the machine over the Internet to check the status of the machine and determine the availability of a cold drink. This check helped them to save their time as they did not have to go to the machine to buy Coke when it was not available in the machine.

Later, other devices like cell phones, coffee makers, washing machines, headphones, lamps, wearable devices and almost anything you can think of could be connected. It is interesting to know that in 2022, the market for the IoT is expected to grow 18% to 14.4 billion active connections and by 2025, there will be approximately 27 billion connected IoT devices. So, we can conclude that, IoT is a giant network of connected ‘things’ (including people) which will have relationships between people-people, people-things and things-things. In simple terminology, IoT is the concept of connecting any devices over the Internet.

4.5.1 Examples of Applications of IoT

Whenever we read about a new technology, the first question that comes to our mind is how is this technology going to help us? While with IoT, anything that can be connected, will be connected, but the question is why should we want these devices to be connected? There are many examples to justify the answers to these questions.

Case 1: Imagine you are going for a meeting and your car tells you the best route to take. If there is a traffic jam, then your car sends a text message to the other party, notifying them that you will be late.

Case 2: Imagine that the moment your alarm rings to wake you, a message is sent to the coffee machine to start brewing coffee for you.

Case 3: The smart watch that your wear in office tells you when and where you were most active and productive.

Case 4: You have a self-driving car with complex sensors to detect objects in their path. In 2019, IoT devices generated an estimated 18.3 zettabytes of data, which is expected to grow to 73.1 ZB by 2025.

Case 5: A smart football that can track how far and fast it is thrown and record those statistics via an app for future training purposes.

Case 6: You have a smart refrigerator that can text message you that milk packets are over in the fridge and you need to buy them before you come home. Or, your fridge checking the expiry date of an ice-cream and notifying you that you should not eat it now.

Case 7: Imagine that you have an AC installed in your home which is connected with your smartphone. On a very hot day, you can instruct your AC to start and set the room temperature to 18 degrees before you finally reach the house.

In an article, Ashton wrote that if computers knew everything about things (using data gathered by them) without any help from us, then we would be able to track and count everything. It would also reduce waste, loss and cost. We would know exactly when things need to be replaced, repaired or recalled, and whether they were fresh or past their best.



Credit: monicaodo. Shutterstock

FIGURE 9.9 A Smart City

In this way, IoT can help organizations save a lot of money through improved process efficiency, asset utilization and enhanced productivity. With improved tracking of objects using sensors and connectivity, companies can better analyse them and make smart decisions. For example, if you own a car manufacturing company, then you can know which accessories are the most popular by using sensors to detect which areas in the showroom are the most popular, and where customers linger longest. You can even use the available sales data to identify which components are selling fastest and then automatically align sales data with supply, so that popular items never go out of stock.

The information collected by IoT devices can be used to detect patterns, make recommendations, and detect possible problems before they occur. We have heard that the

Indian government is working hard to develop smart cities. Do you know that smart cities use IoT for efficient utilization of resources? Figure 9.9 gives an overview of such a city. IoT has the potential to transform entire cities by solving real problems citizens which are faced by the people every day. With the proper connections and data, the IoT can solve traffic congestion issues and reduce noise, crime and pollution.

4.5.2 IoT Products

Big companies like Honeywell, Hitachi, GE, Cisco, AT&T, Apple, Google, IBM, Microsoft, Skyworks, Iridium Communications, Red Hat, Zebra Technologies, InterDigital are already playing in the market to realize the benefits of IoT. Given below are some examples.

Amazon Echo for Smart Home works through its voice assistant, Alexa. Users can talk to Alexa and give order to perform a variety of functions. For example, users can tell Alexa to play music, provide a weather report, get sports scores, order an Uber and do much more. **Fitbit One – Wearables** tracks your steps, floors climbed, calories burned, and quality of sleep. The device wirelessly connects with computers and smartphones to transmit your fitness data in understandable charts to monitor your progress.

Barcelona - Smart Cities: The Barcelona city in Spain is one of the foremost smart cities in the world. It has implemented several IoT initiatives that have helped to enhance smart parking and the environment.

Amazon Web Services, Microsoft Azure, IBM's Watson, Cisco IoT Cloud Connect, Salesforce IoT Cloud, Oracle Integrated Cloud and GE Predix are popular IoT platforms. **AT&T - Connected Car:** AT&T added 1.3 million cars to its network in the second quarter of 2016. With this, the count of total number of connected cars rose to 9.5 million. Thus, we see that IoT allows for virtually endless opportunities and challenges thereby making it a hot topic for research. According to a report, nearly \$6 trillion will be spent on IoT solutions over the next five years.




FIGURE 9.10 Benefits of IoT

4.5.3 Challenges

Security is the biggest issue that we often come across while talking about IoT. With billions of devices connected together, concerns of privacy and data sharing always haunt our minds. In fact, the protection of sensitive data was ranked as the top concern among enterprises according to the 2016 Vormetric Data Threat Report. Hackers try to penetrate connected cars, critical infrastructure and even people's homes. Therefore, the main focus of companies is to ensure security of all the data generated by these devices. Most IoT devices do not encrypt communications if the data is transferred over a local Wi-Fi network. If the Wi-Fi network is unsecured, then it opens the gates for security threats wide open. IoT devices used on patients as wearable devices, if left unattended (with technical errors), can be life-threatening for patient. Another issue with IoT is that massive amounts of data is being generated by these devices. So, companies need to figure out how they would store, monitor, analyse and deduce results from this vast amount of data that is continuously being generated.




FIGURE 9.11 Key challenges in IOT project development

The graph in Fig. 9.11 shows top challenges in the world of IoT. Besides security, other key concerns are as follows.

- 1. Scalability:** When billions of Internet-enabled devices are connected, large volumes of data are generated that need to be processed. To process this huge data, devices need big data analytics and cloud storage for interpretation of useful data. Moreover, these systems need to be scalable to accommodate ever expanding data.
- 2. Interoperability:** Technological standards in most areas are still fragmented and need to be converged to establish a common framework and standard for the IoT devices.
- 3. Lack of standardization:** Since the standardization process is still lacking, interoperability of IoT with legacy devices is critical which, in turn, prevent us to move towards the vision of truly connected interoperable smart objects.
- 4. Lack of government support:** Government and regulatory bodies like the FDA should set up regulations by forming a standard committee for safety and security of devices and people.
- 5. IoT Malware and ransomware:** Ransomware can lock out users from various devices and platforms and still use a user's valuable data and information. For example, a hacker can hijack a computer camera and take pictures. He may also lock a device and then demand ransom to unlock it and return the data.
- 6. Connectivity:** Poor Internet connectivity is a challenge where IoT sensors are required to monitor process data and supply information.

7. Lack of effective and informed government regulations: IoT is a fast-developing area of technology and many legislators are either not aware or do not fully understand the workings of IoT, so they are reluctant to create or enforce regulations.

8. No privacy policies: There are no clear regulations for IoT devices on which information is private or public. For example, an IoT toy could collect information from a child without taking consent from the parents. The toy can then share that data with the manufacturer who could then share it with others. IoT devices should be updated regularly. Updates should be small. Also, IoT devices should be backward compatible with older devices.

9. Bandwidth strain: As the number of IoT devices connected in a geographical area grows, more text/audio/video data will flow through those devices, thereby requiring higher bandwidth and placing strain on the server farms that process all that information.

10. IoT professional skills gap: IoT suffers from a professional skills gap. The main reason behind this is that IoT is growing rapidly and many people are yet to develop a skillset working with this technology. As of now, a few IoT experts are working in the industry, training new people even more difficult. In a survey, it was found that 38% of respondents reported that they lack in-house expertise in IoT network management, while 27% were facing a dearth of in-house experience to deploy IoT.

CASE STUDY EXAMPLE -Several high-profile IoT security breaches have been reported till date. One example is the 2015 Jeep Cherokee hack, when two hackers remotely took control of a car using vulnerabilities in the entertainment system to access its dashboard functions. The hacker initiated a series of unexpected disturbances and finally disabled the brakes, causing the driver to swerve into a ditch. Wireless IoT SIM cards should automatically update the system, removing the possibility of users not updating and leaving holes in system security.

CASE STUDY EXAMPLE -Inter-operability is a serious issue. Once it was seen that there was a problem when an insulin pump communicated only with a particular operating system. It makes the functionality unavailable for patients who may have a computer with a different operating system installed in it. There are issues around the device's wireless standard. For example, if it uses Bluetooth, then the device could inadvertently connect to the wrong nearby device.

CASE STUDY EXAMPLE- A manufacturing plant where a large number of precise, high resolution pictures are sent to 3-D printers could slow down the servers or even cause them to crash if the demands are too great. Similarly, in a hospital, connected video cameras transfer of thousands of high-resolution radiology images and hundreds of telehealth video calls going on in a single day could significantly overload a server farm. However, despite all odds, IoT is being used by several industries including manufacturing, defence, transportation, banks, retail, oil & gas mining, health care, connected home, banks to name a few.

4.5.4 Sensors

A sensor is a device that detects and responds to some type of input that it receives from the physical environment. This input could be heat, light, motion, pressure, moisture or any other environmental phenomena. As the output, sensors usually generate a signal that is

converted to human-readable form and then displayed at the sensor location or transmitted electronically over a network for reading or further processing.

For example, an oxygen sensor in a car's emission control system detects the gasoline/oxygen ratio. If the mixture is not optimal, the balance is readjusted. Another example is motion sensors in home security lights, automatic doors and bathroom fixtures that send out microwaves, ultrasonic waves or light waves and detect when the flow of energy is interrupted by something entering its path.

Vision and imaging sensors detect the presence of objects or colours within their fields of view. They display a visual image of whatever it detects.

Temperature sensors detect thermal parameters and provide signals to the inputs of control and display devices. They are used to measure the thermal characteristics of gases, liquids and solids in many industrial processes.

Radiation sensors sense the presence of alpha, beta or gamma particles and provide signals to counters and display devices. They are usually used for surveys and sample counting.

Proximity sensors detect the presence of nearby objects through non-contacting means within a range of up to several millimetres. They are used in manufacturing operations to detect the presence of parts and machine components.

Pressure sensors detect forces per unit area in gases or liquids and provide signals to the inputs of control and display devices.

Position sensors sense the positions of valves, doors, throttles, etc. and supply signals to the inputs of control or display devices.

Photoelectric sensors sense objects passing within their field of detection. They can even detect colour, cleanliness and location, if required. Photo sensors are commonly used in manufacturing and material handling automation. For example, they are used for counting, robotic picking and automatic doors and gates.

Particle sensors sense dust and other airborne particulates and supply signals to the inputs of control or display devices. Particle sensors are common in bin and baghouse monitoring.

Motion sensors sense the movement or stoppage of parts, people, etc. and supply signals to the inputs of control or display devices. They are usually used for detecting the stalling of conveyors or the seizing of bearings. Metal detectors sense the presence of metal in a variety of situations ranging from packages to people.

Level sensors are used to determine the height of gases, liquids or solids in tanks or bins. The information is then passed as signals to the inputs of control or display devices.

Leak sensors are used to identify or monitor the unwanted discharge of liquids or gases. Some leak detectors are used to measure the effectiveness of the seals in vacuum packages. **Humidity sensors** measure the amount of water in the air and convert these measurements into signals that can be used as inputs to control or display devices.

Gas and chemical sensors sense the presence and properties of various gases or chemicals and relay signals to the inputs of controllers or visual displays.

Force sensors measure various parameters related to forces such as weight, torque, load, etc. and provide signals to the inputs of control or display devices.

Flow sensors sense the movement of gases, liquids or solids and provide signals to the inputs of control or display devices. They are used extensively in the processing industries.

Flaw sensors detect inconsistencies on surfaces or in underlying materials such as welds. They are used in a variety of manufacturing processes.

Flame detectors sense the presence and quality of fire and provide signals to the inputs of control devices. They are used in many combustion control applications like burners. **Electrical sensors** sense current, voltage, etc. and provide signals to the inputs of control devices or visual displays.

Contact sensors detect physical touch or contact between the sensor and the object being observed or monitored. They are often used in alarm systems to monitor doors, windows and other access points. For example, when a door or window is opened/closed, a magnetic switch provides an indication to the alarm control unit so that the status of that entry point is known. They are also used as proximity sensors in robotics applications and automated machinery.

Non-contact sensors do not require a physical touch between the sensor and the object being monitored in order to function. Radar guns used by law enforcement to monitor the speed of vehicles is an example of a non-contact sensor.

Speed sensors detect the speed of an object or a vehicle.

Ultrasonic sensors detect noise and measure the distance between two objects. For this, high-frequency sound waves generated by active ultrasonic sensors are received back by the ultrasonic sensor for evaluating the echo. Time delay between transmitting and receiving the echo is used to calculate the distance to an object. Passive ultrasonic sensors are used for detecting ultrasonic noise present under specific conditions.

4.6 Artificial Intelligence of Things (AIoT)

AIoT, or the combination of artificial intelligence (AI) technologies and the Internet of Things (IoT) infrastructure aims to make IoT operations more efficient, improve interactions between humans and machines and facilitates better data management and analytics. We now clearly know that while AI simulates human intelligence processes by machines, IoT, on the other hand, is a system of interrelated computing devices with unique identifiers and the ability to transfer data over a network without requiring human-to-human or human-to-computer interaction. A thing or a computing device in IoT can be a person's heart monitor implant, a self-driving car with built-in sensors, a door mat or any other object that can be assigned an Internet Protocol address and transfer data over a network. AIoT enhances the capability of IoT by adding the power of machine learning algorithms to improve decision-making processes.

4.6.1 How Does AIoT Work?

In AIoT devices, AI is embedded into programs and infrastructure components, like chips, which are all connected using IoT networks. Efficient exchange of data amongst all hardware, software and platform components without human intervention is ensured using pre-written APIs.





FIGURE 9.12 Autonomous Vehicle

IoT devices generate as well as collect data, that is then analysed using AI techniques to provide better insights into data. This improves efficiency and productivity. Figure 9.12 shows an autonomous vehicle in which AI is embedded into chipsets, which are all connected using IoT networks. AIoT data when processed using Edge AI (discussed later in the chapter) minimizes the bandwidth needed to move data while avoiding possible delays to data analysis. This further adds to the efficiency of IoT.

4.6.2 Where Does AI Unlock IoT?

IoT is all about implanting sensors into machines. These sensors collect streams of data through Internet connectivity. All IoT related services follow five basic steps to perform the assigned task—Create, Communicate, Aggregate, Analyse and Act. Since the Act depends on Analyse, the efficiency of any IoT system depends on the Analysis step. AI technology when used as the analysis stage, improves both efficiency and productivity manifold.



AIoT helps in achieving the following agile solutions:

1. Manage, analyse and obtain meaningful insights from data (refer Fig. 9.13)
2. Provide fast and accurate data analysis
3. Maintain a balance between requirements for localized and centralized intelligence
4. Protect data privacy and confidentiality

5. Ensure security against cyber attack

4.6.3 Applications and Examples of AIoT

AIoT is extensively used in the following domains.

1. Smart cities: Smart technology such as sensors, lights and meters collect data that is processed using AI to improve operational efficiency, drive economic growth and enhance quality of life for residents.

2. Smart retail: Smart cameras are now being used at retail outlets by managers to identify faces of customers who have scanned their items at the self-checkout counter before leaving the store. Moreover, cameras and sensors are installed at multiple places to observe customers' movement and predict when they will reach the checkout line. This helps in maintaining dynamic staffing levels to enhance the productivity of the cashiers and reduce the customer's checkout time. Major retailers can use AIoT solutions to grow sales through customer insights. Data such as mobile-based user behaviour and proximity detection offer valuable insights to deliver personalized marketing campaigns to customers while they shop, increasing traffic in brick-and-mortar locations.

3. Smart home: In smart appliances installed at homes, data collected from human interaction with devices is analysed to understand user habits to provide customized support.

4. Manufacturing: These days, manufacturers use smart chips to detect when a particular equipment is not functioning properly, or when it is the right time to replace.

5. Hiring: AIoT technology when used to integrate data from social media and HR-related platforms can help to identify right talent without any bias.

6. Autonomous vehicles: These vehicles collect real-time data about nearby vehicles, pedestrians and other objects using multiple video cameras and sensors. Data is also analysed to monitor driving conditions and take appropriate decision instantly.

7. Robots: Autonomous robots have multiple sensors installed to gather data about the environment. AI is then used to make decisions that help robot to make smart and optimized moves.

8. Healthcare: Medical devices and wearables collect and monitor real-time health data, such as heart rate. Moreover, healthcare facilities produce high volumes of data including patient information, imaging and test results. This information is valuable and necessary to take good care of the patient care. So, this data needs to be accessed and processed quickly to suggest decision regarding patient's diagnostic and treatment. IoT combined with AI not only improves diagnostic accuracy, enable telemedicine and remote patient care but also reduces the administrative burden of tracking patient health in the facility.

9. Smart Thermostat Solution: The smartphone integration with smart thermostat solution can be used by users to check and manage the temperature from anywhere. Users can set the temperature based on their work schedule and temperature preferences. For example, Nest's smart thermostat solution uses AI-powered IoT.

10. Drone Traffic Monitoring: AIoT can be used by drones to monitor real-time traffic and make adjustments to the traffic flow to reduce congestion. While drones can collect and transmit traffic data, AI can be used to analyse that data and make decisions about how traffic

can be regulated- make diversions, adjustments to speed limits and timing of traffic lights without human intervention. For example, the ET City Brain, a product of Alibaba Cloud, uses AIoT to optimize the use of urban resources. It can detect accidents, illegal parking and can alter traffic lights to help ambulances reach patients/ hospitals faster.

11. Office Buildings: AIoT is used in smart office buildings. In such buildings, usually, a network of smart environmental sensors is installed that detect number of people present in the office and adjust temperatures and lighting accordingly to improve energy efficiency. These smart buildings are often accessed through facial recognition technology using a combination of connected cameras and AI to compare images taken in real-time against a database to determine who should be granted access. this technique is also used by employers to maintain attendance of employees for mandatory meetings.

12. Fleet Management: AIoT is used in fleet management to monitor a fleet's vehicles, reduce fuel costs, track vehicle maintenance and identify unsafe driver behaviour. Through IoT devices such as GPS and other sensors and an artificial intelligence system, companies are able to manage their fleet better thanks to AIoT.

4.6.4 Benefits and Challenges of AIoT

The benefits of AIoT include the following:

1. With AI, operational efficiency of IoT is increased. AIoT devices analyse data to reveal patterns and insights and automatically adjusts system operations.

2. On the fly decision-making: With AIoT and edge computing, vital decisions can be taken instantly. Therefore, data generated by the devices are analysed to identify points of failure/concerns to make appropriate adjustments as and when required.

3. Reduced workload: Since data is analysed by AI and not humans, lot of time and money spent for this time is saved.

4. Scalability: The number of devices connected to an IoT system can be increased anytime as per requirements to optimize existing processes or introduce new features. IoT devices range from mobile devices and high-end computers to low-end sensors. Low-end sensors offer most of the data. An AI-powered IoT ecosystem analyses and summarizes the data from one device before transferring it to other devices. This not only reduces large volumes of data to a handy level but also enables connecting a large number of IoT devices.

5. Boosting Operational Efficiency: AIoT instantly analyses constant streams of data to predict the operational conditions and identify the parameters to be modified to obtain better results. Thus, intelligent IoT can deduce which processes are redundant and time-consuming, and which tasks can be fine tuned to enhance efficiency. For example, Google uses AIoT to reduce its data centre cooling costs.

6. Better Risks Management: Pairing AI with IoT facilitates businesses to understand as well as predict a wide variety of risks and automate processes to provide a prompt response. This helps them to handle financial loss, ensure employee safety and prevent potential cyber threats. For example, Fujitsu ensures safety of the employees by using AI for analysing data coming from connected wearable devices.

7. Triggering New and Enhanced Products and Services: With NLP, users can easily communicate with devices. With this human-machine interaction, IoT and AI can together be used to create new products and services. Even existing products and services can be enhanced by quickly processing and analyzing real-time data. For example, Rolls Royce uses AI technologies in the implementation of IoT-enabled airplane engine maintenance to identify patterns and gain useful operational insights.

8. Eliminates Costly Unplanned Downtime: In industrial manufacturing and in some fields like offshore oil and gas exploration, equipment breakdown can result in costly unplanned downtime. Predictive maintenance with AI enabled IoT predicts equipment failure well in advance and schedules orderly maintenance procedures. This reduces the side effects of downtime. For example, Deloitte, integrates AI and IoT for the following:

1. Reducing 20%–50% of their time spent in maintenance planning
2. Increasing 10%–20% of equipment availability and uptime
3. Reducing 5%–10% amount spent as maintenance costs

All this portrays a rosy picture but do not forget that there are always two faces of the same coin. AIoT may fail. For example, autonomous delivery robots that fail might cause delays in product delivery; smart retail stores may lead to accidentally stealing a product if it fails to read a customer's face, a road accident can be caused if an autonomous vehicle fails to read its surroundings, like a red light on a traffic signal.

4.6.5 Future of AIoT


Integration of AI and IoT (refer Fig. 9.14) creates a much smarter system that makes accurate judgments without human intervention. AIoT will get a tremendous boost with 5G networks that are designed to enable faster transfer of large data files through higher bandwidth and lower latency. AIoT can revolutionize supply chains and delivery models and solve existing operational problems very efficiently. Several businesses have already adopted AI and IoT as part of their processes and products. In fact, AI and IoT are the top technologies in which companies are investing money to increase their operational efficiency and provide a competitive advantage. The graph in Fig. 9.15 depicts the details.



Credit: elenabsi / Shutterstock



Credit: marigranula / 123RF

FIGURE 9.14 An AIoT Machine**FIGURE 9.15** Visualization of emerging trends

It is surprising to note that every day, around one billion gigabytes of data is being generated by IoT devices. By 2025, it is estimated that 42 billion IoT-connected devices will be present globally and with growth in the number of devices and networks, the data generated is set to rise exponentially. With more data, there will be more challenges than opportunities. With AI, IoT networks and devices can learn from past decisions, predict future activity, and continuously improve their decision-making capabilities and thus performance. AI allows the devices to ‘think for themselves’ by interpreting data without the delays and congestion that occur from data transfers.

