freebooks.pk

Computer Science Class 9 Notes

Chapter 10

Emerging Technologies in Computer Science

Punjab Board (RNC 2023) | English Medium

Freebooks.pk

Chapter 10: Emerging Technologies in Computer Science

Artificial Intelligence (AI) is a rapidly growing field that is transforming various aspects of our lives, from healthcare to gaming, reshaping industries and redefining how we live, work, and interact with our environment. This chapter introduces AI's definition and historical evolution, its wide-ranging applications, its key subfields (machine learning, deep learning, natural language processing, computer vision, and robotics), and the distinction between explainable (whitebox) and unexplainable (blackbox) AI algorithms.

The chapter also introduces the Internet of Things (IoT) — a network of physical objects equipped with sensors and software that exchange data over the internet — covering its core components, applications across healthcare and transportation, and important security and privacy considerations. It closes by examining the broader implications of AI and IoT, including risks like data privacy and algorithmic bias, policy and regulatory frameworks, and their societal impact on daily life, work environments, and society at large.

Learning Objectives

- Define Artificial Intelligence (AI) and understand its historical context and evolution
- Identify various applications of AI across domains such as healthcare, education, and gaming
- Explain the subfields of AI: machine learning, deep learning, natural language processing, computer vision, and robotics
- Distinguish between explainable (whitebox) and unexplainable (blackbox) AI algorithms
- Define the Internet of Things (IoT) and describe the components of IoT systems
- Explore applications of IoT in domains such as smart homes, healthcare, and transportation
- Discuss security and privacy considerations in IoT deployments

- Analyze the risks, challenges, and societal impact associated with AI and IoT

Key Concepts

10.1 Introduction to Artificial Intelligence

Artificial Intelligence denotes the simulation of human thinking ability in computer systems, enabling them to think and learn in a manner like humans. AI is being applied to solve complex problems and improve daily experiences — for example, AI-driven systems monitor crop health and predict yields using data from sensors and drones to optimize farming practices.

The term AI was first coined by John McCarthy in 1956 during the Dartmouth Conference, regarded as the origin of AI as a discipline. Key milestones include: the 1950s-1960s (early problem-solving and symbolic methods, including the first AI program, the Logic Theorist, created in 1955 by Allen Newell and Herbert A. Simon); the 1970s-1980s (expert systems mimicking human decision-making); the 1990s (the rise of machine learning); the 2000s onward (advances in deep learning, NLP, and robotics); 2011 (voice assistants for voice command and recognition); and 2023-present (ChatGPT, an AI model designed to understand human-like text).

10.2 Applications of AI Across Domains

AI has numerous applications across different fields. In Healthcare, AI diagnoses diseases, personalizes treatment plans, and predicts patient outcomes. In Education, AI-powered tools provide personalized learning and automate administrative tasks. In Gaming, AI enhances game design and creates realistic characters. In Transportation and Automobiles, self-driving cars and driver-assistance systems improve safety and efficiency.

In Finance, AI enables personalized investment recommendations, fraud detection, and algorithmic trading. In Social Media, AI powers personalized content recommendations, sentiment analysis, and targeted advertising. In Agriculture, AI enables precision farming through predictive analytics and computer-vision-based pest detection. In E-Commerce, AI powers product recommendations, intelligent chatbots, and fraud detection systems.

10.3 Subfields of AI

AI encompasses several subfields, each focusing on different aspects of intelligence. Machine Learning is a type of AI where computers learn from experience and improve over time without being explicitly programmed — like teaching a computer by showing it many examples until it figures out the pattern itself. Deep Learning is a special kind of machine learning that uses complex structures called neural networks, inspired by how the human brain works, helping computers learn from large amounts of data and recognize patterns even better.

Natural Language Processing (NLP) helps computers understand and communicate in human language — for example, when you ask Siri or Alexa a question, they use NLP to understand you and respond helpfully. Computer Vision enables computers to see and understand the visual world, interpreting images and videos. Robotics is the science of building and programming robots — machines that can perform tasks for us, like cleaning floors or building cars, with some robots even able to think and make decisions.

10.4 AI Algorithms: Explainable vs. Unexplainable

AI algorithms can be broadly categorized into two types based on their interpretability. Explainable (Whitebox) Algorithms are those where the decision-making process is transparent and understandable, allowing users to see how decisions are made. Examples include Decision Trees (a flowchart-like tool that helps computers make decisions by following a series of questions), Linear Regression (finding the relationship between two features, such as study time and grades, by fitting a straight line to data points), and Rule-Based Systems (a set of 'if-then' rules written by humans, e.g., 'if the character is about to hit an obstacle, then jump').

Unexplainable (Blackbox) Algorithms are those where the decision-making process is not easily interpretable, often involving complex computations that make it difficult to understand how a particular decision was reached — examples include neural networks and deep learning models. Explainable AI algorithms are especially important in fields like healthcare and finance, where understanding the decision-making process is essential for trust and accountability. Google's AlphaGo, a reinforcement learning model, famously defeated a world champion at the game of Go, a feat once considered nearly impossible given the game's complexity.

10.5 Introduction to the Internet of Things (IoT)

The Internet of Things (IoT) is a network of physical objects, or 'things,' equipped with sensors, software, and other technologies to exchange data with other devices and systems over the internet, enabling new innovative services and more intelligent, efficient operations. IoT is significant because it allows the seamless integration of the physical and digital worlds, enabling devices to collect and share data that can be analyzed to improve efficiency and create new opportunities in fields like healthcare, agriculture, and smart homes. The term 'Internet of Things' was coined by Kevin Ashton in 1999 during his work at Procter & Gamble.

An IoT system typically consists of five components. Sensors detect and measure physical properties like temperature, humidity, light, and motion. Actuators convert energy into motion, acting on data to generate an output. Devices are everyday objects (smartwatches, refrigerators, cars) connected to the internet that use sensor data to perform tasks. Networks are the wired or wireless communication pathways connecting sensors and devices to the internet. Data Analysis involves processing the collected data (on the device, in the cloud, or on a central server) to gain insights and make decisions.

10.6 IoT Applications

IoT is transforming many aspects of life across various domains. In Healthcare, IoT is revolutionizing patient monitoring and care — devices can track vital signs, remind patients to take medication, and alert healthcare providers in emergencies. A practical example is a Smart Home System, where internet-connected appliances like temperature control, lighting, and surveillance cameras work together, and can even help save energy by automatically turning off lights or adjusting heating/cooling when no one is home.

In Transportation, IoT is enhancing systems through connected vehicles, smart traffic lights, and real-time tracking systems, making transportation more efficient and safer. IoT can similarly transform schools through smart classrooms, smart libraries, and smart buses. In 2020, there were over 20 billion IoT devices in use worldwide, highlighting the technology's rapid growth and importance.

10.7 Security and Privacy in IoT Deployments

While IoT offers many benefits, it also raises important security and privacy concerns — as more devices connect to the internet, the risk of cyber-attacks increases, making it essential to secure IoT systems to protect personal data and privacy. Key security measures include: Strong Passwords (using strong, unique passwords for all IoT devices to prevent unauthorized access), Regular Updates (keeping software and firmware up to date to protect against known vulnerabilities), and Encryption (ensuring data transmitted between devices is encrypted to prevent interception by hackers).

Always use devices from reputable manufacturers and keep health or personal data secure using strong passwords and regular device updates. Ensuring IoT devices are connected to a secure network is essential to protect data from unauthorized access.

10.8 Implications and Future of AI and IoT

Emerging technologies like AI and IoT present important risks that must be carefully considered. Data Privacy is a major concern, as AI and IoT devices collect vast amounts of data, putting personal and sensitive information at risk of misuse or unauthorized access. Algorithmic Bias occurs when AI systems are trained on datasets that contain biases, causing the AI models to inadvertently perpetuate or amplify those biases, leading to unfair outcomes in areas like hiring, law enforcement, and lending — addressing this requires analyzing training data and implementing bias-mitigation techniques.

To mitigate these risks, comprehensive policy and regulatory frameworks are essential, focusing on Data Protection Laws (e.g., Europe's General Data Protection Regulation, GDPR), Ethical Guidelines (e.g., IEEE's guidelines for ethical AI), Bias Mitigation Standards, and Security Standards for IoT devices. AI and IoT are already reshaping daily life (smart homes, wearable health devices), work environments (automating repetitive tasks, optimizing industrial production), and society at large (smart cities managing resources efficiently, addressing climate change and healthcare accessibility).

Important Definitions

What is Artificial Intelligence (AI)?

The simulation of human thinking ability in computer systems, enabling them to think and learn in a manner similar to humans.

What is Machine Learning?

A type of artificial intelligence where computers learn from experience and improve over time without being explicitly programmed, by learning patterns from many examples.

What is Deep Learning?

A special kind of machine learning that uses complex structures called neural networks, inspired by the human brain, to learn from large amounts of data and recognize patterns.

What is Natural Language Processing (NLP)?

A technology that helps computers understand and communicate in human language, such as reading, writing, and chatting with users.

What is the Internet of Things (IoT)?

A network of physical objects, or 'things,' equipped with sensors, software, and other technologies to exchange data with other devices and systems over the internet.

What is a whitebox (explainable) AI algorithm?

An AI algorithm where the decision-making process is transparent and understandable, allowing users to see and understand how decisions are made, such as decision trees.

What is a blackbox (unexplainable) AI algorithm?

An AI algorithm where the decision-making process is not easily interpretable due to complex computations, such as neural networks and deep learning models.

What is algorithmic bias?

Unfair or skewed outcomes produced by an AI system because it was trained on biased data, causing it to perpetuate or amplify existing societal biases.

Key Facts and Relations

Topic	Key Fact / Relation
AI term coined by	John McCarthy, 1956, at the Dartmouth Conference
First AI program	Logic Theorist (1955) by Allen Newell and Herbert A. Simon
5 subfields of AI	Machine Learning, Deep Learning, NLP, Computer Vision, Robotics
2 AI algorithm types	Explainable (Whitebox) and Unexplainable (Blackbox)
IoT term coined by	Kevin Ashton, 1999, at Procter & Gamble
5 IoT system components	Sensors, Actuators, Devices, Networks, Data Analysis
3 IoT security measures	Strong Passwords, Regular Updates, Encryption
IoT devices worldwide (2020)	Over 20 billion

Diagrams

Key Milestones in the History of AI: A timeline diagram showing the major milestones in AI's development, from the 1950s Dartmouth Conference through to ChatGPT in 2023

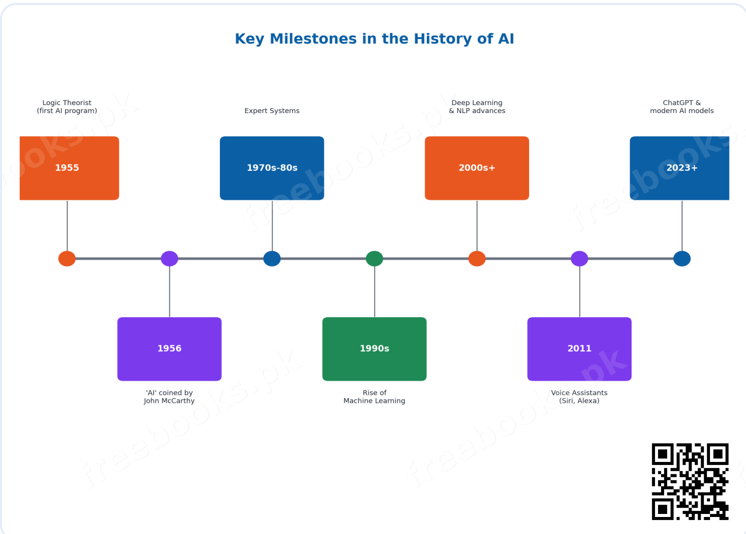

Key Milestones in the History of AI

The Five Subfields of AI: A hub-and-spoke diagram showing Machine Learning, Deep Learning, NLP, Computer Vision, and Robotics as the five key subfields of Artificial Intelligence

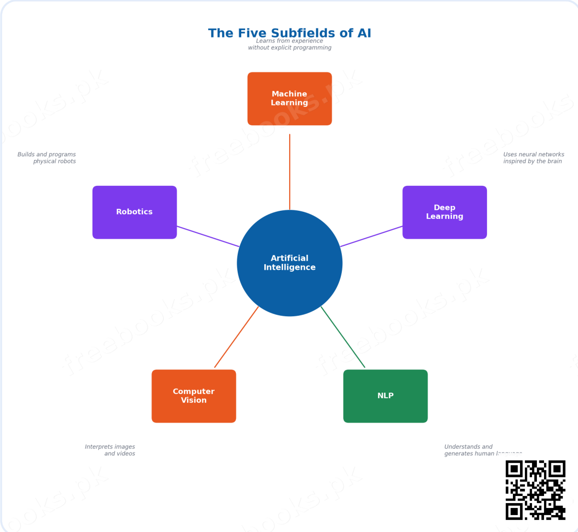

The Five Subfields of AI

Components of an IoT System: A diagram showing the five components of an IoT system — Sensors, Actuators, Devices, Networks, and Data Analysis — and how they connect

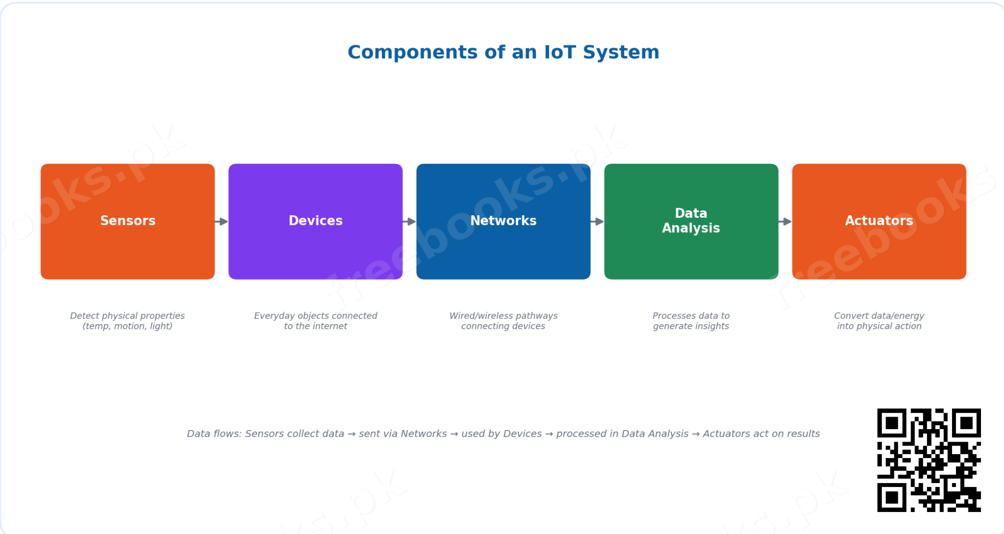

Components of an IoT System

Short Questions & Answers

Define Artificial Intelligence (AI).

Artificial Intelligence is the simulation of human thinking ability in computer systems, enabling them to think and learn in a manner like humans, in order to solve complex problems.

What is the historical context and evolution of AI?

AI was first coined as a term by John McCarthy in 1956 at the Dartmouth Conference, evolving from early symbolic problem-solving methods in the 1950s-60s, through expert systems in the 1970s-80s, machine learning in the 1990s, and deep learning/NLP advances from the 2000s onward, up to modern tools like ChatGPT.

Provide two examples of AI applications in healthcare.

AI is used in healthcare for diagnosing diseases and for personalizing treatment plans based on a patient's individual data; it is also used for predicting patient outcomes and enabling remote patient monitoring.

Define the Internet of Things (IoT).

The IoT is a network of physical objects, or 'things,' equipped with sensors, software, and other technologies that enable them to exchange data with other devices and systems over the internet.

Describe the significance of IoT in connecting devices and systems.

IoT is significant because it allows the seamless integration of the physical and digital worlds, enabling devices to collect and share data that can be analyzed to improve efficiency, provide better services, and create new opportunities in fields like healthcare and smart homes.

What are the potential risks associated with AI and IoT?

Key risks include data privacy concerns (since AI and IoT devices collect vast amounts of personal data) and algorithmic bias (where AI systems trained on biased data can produce unfair outcomes in areas like hiring or lending).

Explain the concept of algorithmic bias.

Algorithmic bias occurs when an AI system is trained on datasets that contain existing biases, causing the resulting AI models to inadvertently perpetuate or even amplify those biases, leading to unfair outcomes in real-world applications.

What are the three main components of IoT security measures?

The three main IoT security measures are using strong, unique passwords for all devices, keeping software and firmware regularly updated, and ensuring data transmitted between devices is encrypted.

Long Questions & Answers

Discuss the various applications of AI in the field of education, providing specific examples and explaining how AI can enhance the educational experience.

Artificial Intelligence has genuinely begun to transform the field of education in a number of distinct but closely related meaningful ways, fundamentally enhancing the overall educational experience for students, teachers, and educational institutions alike through several particular specific distinct applications. One especially important application of AI within education is the specific provision of personalized learning experiences, whereby AI-powered educational tools and platforms are specifically able to properly analyse each individual particular student's own unique specific learning pace, their own particular individual specific strengths, and their own particular individual specific areas that may still require further improvement, and then subsequently adapt the given educational content accordingly, in a way that a single standardized, one-size-fits-all traditional classroom approach could simply never realistically be able to properly achieve; for example, a given adaptive learning software platform might specifically notice that one particular individual student is genuinely struggling with a specific particular mathematical concept, such as fractions, and would then automatically provide that specific particular student with additional targeted practice problems and further supplementary explanatory materials that are specifically tailored to their own particular individual specific learning needs, while simultaneously allowing other students within that same class who have already properly mastered that particular given concept to instead move on ahead to more genuinely advanced material. A second especially important application of AI within education specifically involves the automation of various different everyday routine administrative tasks, thereby freeing up considerably more valuable time for classroom teachers to instead focus more directly on actual teaching itself, rather than being burdened with excessive paperwork; for example, AI-powered grading and marking systems are now increasingly able to automatically and efficiently grade multiple-choice tests and, in certain more advanced particular cases, even properly grade written short-answer responses as well, thereby significantly reducing the sheer amount of time that individual teachers would otherwise need to spend manually marking each individual student's coursework by hand. A third especially important application of AI within education specifically involves providing much deeper, more genuinely valuable insights into overall student performance, thereby helping both individual teachers and entire wider school administrations alike to properly identify at-risk students considerably earlier on; for example, an AI-powered analytics system might specifically notice a clear, gradually declining pattern in one particular individual student's ongoing quiz scores over a period of several consecutive weeks, and would then automatically alert that particular student's classroom teacher, thereby allowing that teacher to properly intervene and offer additional targeted academic support well before that particular student's own overall academic performance has the chance to decline any further. Taken together as a whole, these various different particular specific applications of AI within the field of education collectively demonstrate genuinely significant, real potential to make learning itself considerably more properly personalized to each individual student's needs, considerably more genuinely efficient overall, and considerably more effective for every single individual student, ultimately regardless of each particular student's own unique individual specific learning needs or particular individual pace.

Differentiate between explainable (whitebox) and unexplainable (blackbox) AI models, and discuss why this distinction matters in practice.

Explainable (whitebox) AI models and unexplainable (blackbox) AI models represent two genuinely fundamentally different broad general categories of AI algorithm, and the key underlying distinction between these two particular categories specifically concerns the overall degree to which a given human user is actually able to properly understand precisely how that particular AI model actually arrived at any one of its own particular given individual specific decisions or predictions. Explainable, or whitebox, AI models are those particular models whose own internal decision-making process remains genuinely transparent and properly understandable to a human observer; a clear, concrete example of this particular kind of model is a simple decision tree, which specifically works by following a clear, properly structured sequence of individual logical questions, extremely similar in basic overall structure to a standard flowchart, ultimately arriving at some particular given final decision or specific outcome through a series of individual logical steps that a human user could, at least in principle, quite easily and directly trace back through by hand, one individual step at a time. Unexplainable, or blackbox, AI models, by clear contrast, are instead those particular models whose own internal decision-making process is simply not easily interpretable or understandable by a human user, largely on account of the sheer overall underlying complexity of the many individual internal computations and internal interactions that are actually taking place within that particular model; a clear, concrete example of this particular kind of model is a deep neural network, which specifically works by passing a given individual piece of input data through many separate individual interconnected internal layers of individual artificial 'neurons,' with each one of these individual internal layers separately applying its own particular individual specific mathematical transformation to that given data, such that the final given overall specific outcome genuinely emerges from an extremely large number of individual interacting internal calculations that are simply far too complex for any human user to properly and fully trace back through by hand. This particular distinction between explainable and unexplainable AI models genuinely matters a great deal in real-world practical terms, especially within certain particular high-stakes specific fields such as healthcare and finance, where properly understanding precisely why a given AI system actually reached one particular specific given decision, such as a particular medical diagnosis or a particular loan application decision, is often just as practically important as the actual final given decision itself; if a given doctor, for example, cannot properly understand precisely why a given deep learning blackbox AI model has specifically recommended one particular course of medical treatment over another, that doctor may then reasonably find it very difficult to fully trust that particular given AI recommendation, or to properly explain that particular recommendation clearly to their own individual patient. As a direct result of this, whitebox AI models are very often specifically preferred within these kinds of particular high-stakes specific application contexts, even in various different specific cases where a given corresponding blackbox model might potentially otherwise be capable of achieving somewhat higher raw overall predictive accuracy, precisely because the ability to properly explain a given decision is very often considered to be just as practically important as the raw overall accuracy of that particular decision itself.

Describe the components of an IoT system, and explain how these components work together to enable IoT applications, using a smart home as a worked example.

An IoT system is generally composed of five separate distinct core components that must all properly work together seamlessly in close combination with one another in order to successfully enable any given particular IoT application to actually function correctly and effectively as intended. The first of these five core components is Sensors, which specifically function as devices that detect and properly measure various different particular physical properties, such as temperature, humidity, ambient light levels, or physical motion, and which are specifically responsible for collecting the necessary raw underlying data directly from the surrounding physical environment. The second core component is Actuators, which specifically function as devices that convert incoming electrical energy into some form of physical motion or other specific physical action, and which specifically act upon the previously collected data in order to actually generate some kind of tangible real-world physical output or response. The third core component is Devices themselves, which specifically refer to the actual everyday physical objects, such as smartwatches, refrigerators, or motor vehicles, that are properly connected to the internet and which specifically make direct practical use of the data collected by the sensors in order to properly perform their own particular given specific tasks. The fourth core component is Networks, which specifically function as the underlying communication pathways, whether wired or wireless in nature, that properly connect the sensors and devices together to the wider internet, thereby allowing them to properly share their collected data with one another. The fifth and final core component is Data Analysis, which specifically involves the process of properly processing and carefully analysing all of the collected data, whether directly on the device itself, remotely in the cloud, or on some form of centralized server, in order to properly generate genuinely useful insights and to support well-informed decision-making. In order to properly illustrate precisely how these five separate individual components actually work together in genuine practice, consider the specific concrete practical example of a smart home security system: first, a given motion sensor installed near the home's front door (the sensor component) specifically detects physical movement in that particular area; this detected sensor data is then transmitted over the home's own wireless internet network (the network component) to a smartphone application (the connected device component); the underlying system's own data analysis component then properly determines whether this particular detected movement genuinely represents an actual potential security threat, based on factors such as the particular specific time of day or the home's own current occupancy status; and finally, if this particular movement is indeed properly determined to represent a genuine potential security threat, the system will then trigger a connected actuator, such as an automatic outdoor security light or an audible alarm siren, to properly respond accordingly. This particular smart home example clearly demonstrates precisely how all five of these individual core IoT components must necessarily work together in close, properly coordinated combination with one another, from the initial raw physical sensing stage all the way through to the final given physical action stage, in order to successfully create one single genuinely coherent, properly functioning, and practically useful overall IoT application.

Multiple Choice Questions (MCQs)

Which of the following is not a subfield of AI? (A) Machine Learning (B) Natural Language Processing (C) Computer Vision (D) Web Hosting

Correct answer: (D) Web Hosting. Web Hosting is not a subfield of AI. Machine Learning, NLP, and Computer Vision (along with Deep Learning and Robotics) are all recognized AI subfields.

Which of these AI algorithms is considered an 'explainable' model? (A) Neural Networks (B) Decision Trees (C) Deep Learning Models (D) Convolutional Neural Networks

Correct answer: (B) Decision Trees. Decision Trees are explainable (whitebox) models because their decision-making process follows a transparent, traceable sequence of questions.

Which of these is a security concern in IoT deployments? (A) Device vulnerability (B) Data privacy (C) Lack of standardization (D) All of the above

Correct answer: (D) All of the above. IoT deployments face multiple security concerns including device vulnerability, data privacy risks, and a lack of standardization across devices.

Which of the following is an application of AI in healthcare? (A) Personalized drug development (B) Automated diagnosis (C) Remote patient monitoring (D) All of the above

Correct answer: (D) All of the above. AI is applied across all of these healthcare areas: personalized drug development, automated diagnosis, and remote patient monitoring.

What is the key difference between explainable (whitebox) and unexplainable (blackbox) AI models? (A) The complexity of the model (B) The ability to understand the decision-making process (C) The performance of the model (D) The training data used

Correct answer: (B) The ability to understand the decision-making process. The key difference is the ability to understand the decision-making process — whitebox models are transparent, while blackbox models are not easily interpretable.

Which of the following is an application of IoT in the transportation domain? (A) Smart traffic management (B) Vehicle-to-Vehicle (V2V) communication (C) Predictive maintenance of vehicles (D) All of the above

Correct answer: (D) All of the above. IoT is applied across all of these transportation areas: smart traffic management, V2V communication, and predictive vehicle maintenance.

What is the key concern associated with algorithmic bias in AI-powered decision-making? (A) Lack of transparency (B) Perpetuation of existing societal biases (C) Reduced accuracy of the model (D) All of the above

Correct answer: (B) Perpetuation of existing societal biases. The key concern with algorithmic bias is that it can perpetuate or even amplify existing societal biases, leading to unfair outcomes.

Which of the following is an ethical principle for the responsible development of AI and IoT technologies? (A) Transparency and accountability (B) Respect for privacy and data rights (C) Fairness and non-discrimination (D) All of the above

Correct answer: (D) All of the above. Responsible AI/IoT development requires all of these principles: transparency and accountability, respect for privacy, and fairness/non-discrimination.

Who is credited with coining the term 'Artificial Intelligence'? (A) Alan Turing (B) John McCarthy (C) Kevin Ashton (D) Allen Newell

Correct answer: (B) John McCarthy. John McCarthy first coined the term 'Artificial Intelligence' in 1956 during the Dartmouth Conference.

Who coined the term 'Internet of Things'? (A) John McCarthy (B) Herbert Simon (C) Kevin Ashton (D) Alan Turing

Correct answer: (C) Kevin Ashton. Kevin Ashton coined the term 'Internet of Things' in 1999 while working at Procter & Gamble.

Quick Revision Summary

- AI term coined by John McCarthy (1956, Dartmouth Conference); first AI program was Logic Theorist (1955)
- 5 subfields of AI: Machine Learning, Deep Learning, NLP, Computer Vision, Robotics
- 2 AI algorithm types: Explainable/Whitebox (decision trees, linear regression, rule-based) vs. Unexplainable/Blackbox (neural networks, deep learning)
- IoT term coined by Kevin Ashton (1999); over 20 billion IoT devices worldwide by 2020
- 5 IoT components: Sensors, Actuators, Devices, Networks, Data Analysis
- 3 IoT security measures: Strong passwords, Regular updates, Encryption
- Key AI/IoT risks: Data Privacy and Algorithmic Bias
- Policy frameworks: Data Protection Laws (e.g., GDPR), Ethical Guidelines (e.g., IEEE), Bias Mitigation Standards, Security Standards

Exam Tips

- Remember AI's 5 subfields with 'MDNCR': Machine learning, Deep learning, NLP, Computer vision, Robotics
- Whitebox = transparent/explainable (decision trees); Blackbox = opaque/complex (neural networks) — a common exam distinction
- IoT's 5 components in order of data flow: Sensor → Network → Device → Data Analysis → Actuator (output)
- Remember the 3 IoT security measures: strong passwords, regular updates, encryption
- Algorithmic bias comes from biased training data — always link the cause (biased data) to the effect (unfair outcomes) in your answer
- Know your key dates: AI = 1956 (McCarthy); IoT = 1999 (Ashton) — frequently tested