



CMR INSTITUTE OF TECHNOLOGY

**Affiliated to VTU, Approved by AICTE,
Accredited by NBA and NAAC with “A++” Grade
ITPL MAIN ROAD, BROOKFIELD, BENGALURU-560037,
KARNATAKA, INDIA**

DEPARTMENT OF ARTIFICIAL INTELLIGENCE & DATA SCIENCE ENGINEERING

**COURSE: INTRODUCTION TO ARTIFICIAL
INTELLIGENCE**

COURSE CODE: 1BAIA103

**VTU SOLUTION
2025-2026**

**BY
PROF. ANUSHREE PAUL
CMRIT, AIDS**

CBCS SCHEME

USN

1BAIA103

First Semester B.E./B.Tech. Degree Examination, Dec.2025/Jan.2026 Introduction to AI and Applications

Time: 3 hrs.

Max. Marks: 100

*Note: 1. Answer any FIVE full questions, choosing ONE full question from each module.
2. M : Marks, L: Bloom's level, C: Course outcomes.*

Module - 1			M	L	C
Q.1	a.	Explain in detail the different types of Artificial Intelligence based on functionality and capabilities.	7	L2	CO1
	b.	Explain different types of knowledge used in AI with suitable examples.	6	L2	CO1
	c.	Compare traditional programming with machine learning.	7	L2	CO1
OR					
Q.2	a.	What is Machine Learning? Relate AI and Machine Learning?	7	L2	CO1
	b.	What is AI Agent? Classify the agents in an AI system.	7	L2	CO1
	c.	Compare weak AI and strong AI.	6	L2	CO1
Module - 2					
Q.3	a.	Differentiate between Zero-shot, One-shot, and Few-shot prompting.	6	L2	CO4
	b.	Explain Self-Consistency Prompting and how does it improve model performance.	6	L2	CO4
	c.	Apply LLM-based writing prompts to produce creative content for digital marketing campaigns.	8	L3	CO4
OR					
Q.4	a.	Build a zero-shot prompt to classify comments as positive or negative. Explain why it qualifies as zero-shot.	6	L3	CO4
	b.	Explain the steps for Effective Prompt Engineering.	6	L2	CO4
	c.	Build a creative prompt that encourages ChatGPT to propose futuristic transportation ideas. Explain how this promotes imaginative thinking.	8	L3	CO4
Module - 3					
Q.5	a.	Is Labelled data supervised or unsupervised machine learning? Extend your answer explaining different types of machine learning.	6	L2	CO2
	b.	Outline K-Means Algorithm with merits and demerits.	8	L2	CO2
	c.	Summarize the features of: (i) Reinforcement Learning (ii) Support Vector Machines (SVM)	6	L2	CO2
OR					

Q.6	a.	Explain the four steps to create Decision Trees with suitable example for each step.	6	L2	CO2
	b.	Define conditional probability and the Bayes rule with examples.	8	L2	CO2
	c.	How does computer vision work with deep learning? Explain the tasks involved in computer vision.	6	L2	CO2
Module – 4					
Q.7	a.	List and explain any four Trusted AI principles.	8	L2	CO5
	b.	Describe the concept of AI as a Service (AIaaS). Outline two advantages and two challenges of using AIaaS in organizations.	6	L2	CO5
	c.	Relate the role of ethics in AI.	6	L2	CO5
OR					
Q.8	a.	What is expert system? Explain three components of expert system.	6	L2	CO5
	b.	Differentiate between the following: (i) AI Programs and Robots (ii) Human-controlled and fully – autonomous bots	6	L2	CO5
	c.	What is Neuromorphic Computing? Explain its architecture, features resembling the human brain, and how it contributes toward achieving Artificial General Intelligence (AGI).	8	L2	CO5
Module – 5					
Q.9	a.	Relate the role of AI in Biology and Environmental Sciences.	6	L2	CO3
	b.	What is No-Code AI. Explain why No-Code AI Must be Used?	6	L2	CO3
	c.	Identify the role of AI in scientific experimentation by examining how it supports different disciplines and breaking down the specific experimental activities such as data collection, pattern identification, simulation, and hypothesis testing that AI enhances.	8	L3	CO3
OR					
Q.10	a.	Explain the role of AI in early disease prevention.	6	L2	CO3
	b.	What is Low Code AI. Compare Traditional tools with Low Code AI.	6	L2	CO3
	c.	Explain how AI contributes to environmental science by breaking down its role in climate modelling, air and water quality monitoring, waste management, and resource conservation. Describe the specific data, techniques, and decision-making processes involved in each area.	8	L3	CO3

Answers

Q1 (a) Explain in detail the different types of Artificial Intelligence based on functionality and capabilities

Artificial Intelligence can be classified based on **functionality** and **capabilities** depending on how intelligently the system behaves and the level of intelligence it possesses.

Based on **functionality**, the first type is **Reactive Machines**. These systems do not have memory and cannot use past experiences to make decisions. They react only to the present input. A classic example is IBM's Deep Blue chess system, which evaluates current board positions without learning from previous games.

The second type is **Limited Memory AI**, which can use past data for a short duration to make decisions. Most real-world AI systems today fall under this category. Examples include self-driving cars that observe traffic patterns and recommender systems used by e-commerce platforms.

The third type is **Theory of Mind AI**, which is still under research. Such systems are expected to understand human emotions, beliefs, intentions, and social interactions. This type of AI would be capable of interacting naturally with humans.

The final type is **Self-aware AI**, which possesses self-consciousness and awareness of its own existence. This form of AI is purely hypothetical and does not exist at present.

Based on **capabilities**, AI is classified into **Weak AI**, **Strong AI**, and **Super AI**. Weak AI is designed to perform a specific task such as voice assistants or image recognition systems. Strong AI refers to systems that can perform any intellectual task a human can do, while Super AI surpasses human intelligence. Strong and Super AI remain theoretical concepts.

Q1 (b) Explain different types of knowledge used in AI with suitable examples

Knowledge representation plays a crucial role in Artificial Intelligence. AI systems use different types of knowledge to reason and make decisions.

Declarative knowledge represents factual information about the world. It includes statements and facts such as "The sun rises in the east." This knowledge is easy to represent and modify.

Procedural knowledge describes how to perform a task. It is represented as a sequence of steps or instructions. For example, the steps involved in solving a mathematical equation represent procedural knowledge.

Heuristic knowledge is experience-based knowledge that helps in problem-solving using rules of thumb. It does not guarantee optimal solutions but provides efficient results. Chess strategies used by AI programs are examples of heuristic knowledge.

Meta-knowledge is knowledge about other knowledge. It helps the system decide which knowledge or method to use in a particular situation.

Structural knowledge represents relationships between different concepts. It is often stored using semantic networks or frames. For example, defining relationships such as “A car has an engine” represents structural knowledge.

Q1 (c) Compare traditional programming with machine learning

Traditional programming and machine learning differ fundamentally in how solutions are developed. In traditional programming, the programmer explicitly writes rules and logic to process the input data and produce the output. The system strictly follows predefined instructions.

In contrast, machine learning systems are trained using data and expected outputs. Instead of writing rules manually, the system learns patterns from the data and builds a model. Machine learning systems can improve performance over time as more data becomes available. Traditional programming is rigid and less adaptive, whereas machine learning systems are flexible and capable of handling complex and uncertain problems.

Q2 (a) What is Machine Learning? Relate AI and Machine Learning

Machine Learning is a branch of Artificial Intelligence that focuses on enabling machines to learn from data and improve their performance without being explicitly programmed. It uses statistical and mathematical techniques to identify patterns in data and make predictions or decisions.

Artificial Intelligence is the broader concept that aims to create intelligent machines capable of performing tasks that normally require human intelligence. Machine Learning is a subset of AI

and serves as one of the primary approaches to achieving artificial intelligence. While AI includes areas such as robotics, natural language processing, and expert systems, machine learning provides the learning capability that allows AI systems to adapt and improve over time.

Q2 (b) What is an AI Agent? Classify the agents in an AI system

An AI agent is an entity that perceives its environment through sensors and acts upon that environment using actuators. The goal of an agent is to perform actions that maximize its performance measure.

A **simple reflex agent** works based on condition–action rules and responds directly to the current percept. A **model-based agent** maintains an internal representation of the environment to handle partially observable situations. A **goal-based agent** selects actions that help achieve specific goals. A **utility-based agent** chooses actions that maximize overall utility. A **learning agent** improves its performance by learning from experience and feedback from the environment.

Q2 (c) Compare Weak AI and Strong AI

Weak AI (Narrow AI)

Weak AI, also known as Narrow AI, refers to artificial intelligence systems that are designed and trained to perform a specific task or a limited range of tasks. These systems operate under predefined algorithms and learned patterns but do not possess consciousness, self-awareness, or genuine understanding. They simulate intelligence rather than actually thinking or reasoning like humans.

Key Characteristics:

- Task-specific functionality
- Operates within limited boundaries
- Cannot generalize knowledge beyond its training
- No emotions, self-awareness, or consciousness
- Relies on data-driven learning models

Examples:

- Virtual assistants like Siri and Alexa
- Chatbots (e.g., customer support bots)

- Recommendation systems (Netflix, Amazon)
- Spam email filters
- Face recognition systems
- Self-driving car perception systems

For example, a recommendation system can suggest movies based on your viewing history, but it cannot understand cinema in a broader intellectual sense like a human can.

Strong AI (Artificial General Intelligence – AGI)

Strong AI refers to a theoretical form of artificial intelligence that would possess the ability to **understand, learn, reason, and apply knowledge across a wide variety of tasks**, similar to human intelligence.

Unlike Weak AI, Strong AI would:

- Exhibit genuine understanding
- Have reasoning and problem-solving abilities
- Transfer learning across domains
- Potentially possess self-awareness or consciousness

Key Characteristics:

- Human-level cognitive abilities
- Ability to generalize knowledge
- Independent decision-making
- Emotional and contextual understanding
- Adaptive learning across different domains

For example, a Strong AI system could:

- Learn mathematics
 - Write poetry
 - Drive a car
 - Conduct scientific research
 - Understand emotions
- all without being specifically programmed for each individual task.
-

Q3 (a) Differentiate between Zero-shot, One-shot, and Few-shot prompting

Three commonly used prompting techniques are:

- Zero-shot prompting
- One-shot prompting
- Few-shot prompting

Zero-Shot Prompting

Zero-shot prompting refers to giving a task to a language model without providing any examples of how the output should look. The model must rely entirely on its pre-trained knowledge and general understanding of language patterns.

In this case, the instruction alone guides the model.

Example:

Prompt:

Translate the following sentence into French: "Good morning."

Here, no example translation is provided. The model uses its prior training to generate the correct translation.

Characteristics:

- No demonstration examples
- Faster and simpler prompting
- Relies heavily on model training
- May produce less precise outputs for complex tasks

One-Shot Prompting

One-shot prompting involves providing a single example along with the instruction to guide the model. The example demonstrates the expected format or reasoning style.

The model uses that one demonstration to understand how to respond.

Example:

Prompt:

Example:
Input: "I love this movie."
Output: Positive

Now classify the sentiment:
Input: "This product is terrible."
Output:

Here, one example helps the model understand the task (sentiment classification).

Characteristics:

- Includes one demonstration
- Improves clarity over zero-shot
- Helps define output structure
- Moderately improves accuracy

Few-Shot Prompting

Few-shot prompting provides multiple examples before asking the model to perform the task. These examples help the model better understand patterns, structure, and expectations.

This technique significantly improves performance for more complex or reasoning-based tasks.

Example:

Prompt:

Example 1:
Input: "I love this movie."
Output: Positive

Example 2:
Input: "The service was disappointing."
Output: Negative

Example 3:
Input: "The experience was fantastic."
Output: Positive

Now classify:
Input: "The product quality is poor."
Output:

The multiple examples clearly define the task pattern.

Characteristics:

- Multiple demonstrations
 - Higher accuracy
 - Better pattern recognition
 - Suitable for complex reasoning tasks
-

Q3 (b) Explain Self-Consistency Prompting and how it improves model performance

Self-Consistency Prompting

Self-consistency prompting is an advanced prompting technique used mainly with large language models (LLMs) for solving complex reasoning problems. Instead of generating only one reasoning path and one final answer, the model generates **multiple independent reasoning chains** for the same problem. The final answer is selected based on the **most consistent or most frequently occurring result** among those multiple outputs.

This method improves reliability and reduces reasoning errors. In traditional prompting (even with chain-of-thought reasoning), the model generates:

- One reasoning path
- One final answer

However, a single reasoning chain may:

- Contain logical mistakes
- Be influenced by randomness in generation
- Lead to incorrect intermediate steps

Self-consistency addresses this limitation by exploring **multiple reasoning possibilities** and choosing the most stable answer.

Key Characteristics

- Generates multiple reasoning chains
- Uses stochastic sampling (introducing variation)
- Applies majority voting for final selection
- Improves robustness and reliability
- Particularly effective in arithmetic and logical reasoning

Advantages

1. Reduces reasoning errors
 2. Improves logical consistency
 3. Enhances mathematical accuracy
 4. Minimizes the impact of random mistakes
-

Q3 (c) Apply LLM-based writing prompts for digital marketing campaigns

Role of Large Language Models (LLMs) in Digital Marketing

Large Language Models (LLMs) are advanced AI systems capable of understanding and generating human-like text. When provided with well-structured prompts, these models can assist marketers in creating high-quality, engaging, and personalized content efficiently.

In digital marketing, content plays a crucial role in attracting, engaging, and converting customers. LLMs help automate and enhance this process.

Social Media Content Creation

LLMs can generate:

- Instagram captions
- Twitter posts
- LinkedIn content
- Facebook ad copy
- Hashtag suggestions

Example Prompt:

Generate an engaging Instagram caption for a new organic skincare product targeting young professionals.

The model can produce creative, audience-specific content within seconds.

Product Descriptions

LLMs can create compelling and SEO-friendly product descriptions by highlighting:

- Key features
- Benefits
- Unique selling points
- Emotional appeal

Example Prompt:

Write a persuasive product description for a wireless noise-canceling headphone targeting students.

This improves:

- Conversion rates
- Customer engagement
- Search engine visibility

Email Marketing Campaigns

LLMs assist in drafting:

- Promotional emails
- Newsletter content
- Follow-up emails
- Abandoned cart reminders
- Personalized offers

Example Prompt:

Draft a promotional email for a festive discount sale with a professional and friendly tone.

Personalized email content increases:

- Open rates
- Click-through rates
- Customer retention

Advertisement Copywriting

LLMs can generate:

- Google Ads headlines
- Call-to-action (CTA) phrases
- Landing page content
- Campaign slogans

They can adjust tone based on:

- Target audience
- Brand voice
- Industry type

Audience Personalization

One of the biggest advantages of LLMs is their ability to tailor content for specific segments:

- Age groups
- Geographic regions
- Interests
- Buying behavior

For example:

Generate a product pitch for fitness enthusiasts aged 20–30.

This personalization improves campaign effectiveness and customer satisfaction.

Benefits of Using LLMs in Digital Marketing

1. **Time Efficiency** – Automates repetitive content creation tasks.
2. **Cost Reduction** – Reduces dependency on large content teams.
3. **Creativity Enhancement** – Generates multiple variations of content quickly.
4. **Scalability** – Enables large-scale campaign generation.
5. **Consistency** – Maintains uniform brand tone across platforms.
6. **A/B Testing Support** – Easily generates multiple versions for testing.

Importance of Well-Structured Prompts

The effectiveness of LLMs depends on how clearly the prompt is written. A good prompt should include:

- Target audience
- Tone (professional, friendly, persuasive, etc.)
- Platform (Instagram, email, website)
- Content goal (engagement, sales, awareness)
- Word limit (if required)

Better prompts produce more relevant and high-quality outputs.

Q4 (a) Build a zero-shot prompt to classify comments as positive or negative and explain

Zero-Shot Prompting for Sentiment Classification

Zero-shot prompting is a technique in which a task is given to a language model **without providing any example outputs**. The model relies entirely on its **pre-trained knowledge and learned language patterns** to perform the task.

In the given example:

“Classify the following comment as Positive or Negative: ‘The service quality was extremely poor.’”

The prompt directly instructs the model to perform **sentiment classification**, but it does not include any demonstration of how positive or negative comments should be labeled.

This prompt qualifies as **zero-shot** because:

- No example input-output pair is given.
- The model is not shown how to classify a sample sentence.
- The instruction alone guides the model’s response.

The model must depend on its prior training, during which it learned patterns such as:

- Words like “poor,” “bad,” “terrible” → Negative sentiment
- Words like “excellent,” “great,” “amazing” → Positive sentiment

Since the phrase “**extremely poor**” indicates dissatisfaction, the model classifies the sentence as **Negative**.

Large Language Models are trained on vast amounts of text data. During training, they learn:

- Word associations
- Emotional tone patterns
- Contextual meaning
- Linguistic sentiment indicators

When given the sentence:

“The service quality was extremely poor.”

The model identifies:

- The adjective “poor”
- The intensifier “extremely”
- The negative tone of the overall sentence

Using this learned knowledge, it predicts the appropriate label: **Negative**.

Advantages of Zero-Shot Prompting in Sentiment Analysis

1. Simple and easy to implement
2. Requires no example demonstrations
3. Fast and efficient for straightforward tasks
4. Works well when the task is clear and commonly understood

Limitations

Although zero-shot prompting is convenient, it may:

- Produce incorrect results for ambiguous sentences
- Struggle with sarcasm or complex emotions
- Be less accurate compared to few-shot prompting

For example:

“The movie was so bad that it was actually good.”

Such sentences may confuse a zero-shot model due to contextual nuance.

Q4 (b) Explain the steps for effective prompt engineering

Effective Prompt Engineering

Prompt engineering is the process of designing and refining inputs (prompts) given to a language model in order to obtain accurate, relevant, and high-quality outputs. Since large language models generate responses based on patterns learned during training, the clarity and structure of the prompt significantly influence the result.

Effective prompt engineering ensures that the model clearly understands:

- What task needs to be performed
- How the output should be structured
- What constraints must be followed
- What context is necessary

Clearly Defining the Task

The first and most important step in prompt engineering is to **explicitly define the task**. Ambiguous prompts may lead to vague or incorrect outputs.

For example:

Weak Prompt:

Explain marketing.

This is too broad.

Effective Prompt:

Explain digital marketing strategies in 5 points suitable for undergraduate students.

The improved prompt clearly defines:

- Topic (digital marketing strategies)
- Scope (5 points)
- Target audience (undergraduate students)

Clear task definition reduces ambiguity and improves output relevance.

Providing Sufficient Context

Context helps the model understand the situation or background of the task. Without context, the model may generate generic responses.

For example:

Without context:

Write an email.

With context:

Write a professional email to a customer informing them about a delay in product delivery due to supply chain issues.

The second prompt gives:

- Purpose
- Audience
- Situation

This leads to more accurate and meaningful output.

Including Precise Instructions and Constraints

Effective prompts should specify:

- Tone (formal, informal, persuasive, academic)
- Length (word limit or number of points)
- Format (bullet points, paragraph, table)
- Style (technical, beginner-friendly)

Example:

Write a 150-word academic explanation of machine learning in paragraph form.

Constraints guide the model and reduce irrelevant or overly lengthy responses.

Specifying Expected Output Format

Clearly defining the output structure ensures consistency and usability.

For example:

Summarize the following text in bullet points.

or

Provide the answer in a table format with three columns: Feature, Advantage, Example.

This is especially important in technical, academic, and professional applications.

Iterative Refinement

Prompt engineering is not a one-time process. It involves:

- Testing the prompt

- Evaluating the response
- Modifying unclear parts
- Adding or removing constraints

This process is known as **iterative refinement**.

For example:

First attempt:

Explain AI.

After refinement:

Explain Artificial Intelligence in 6 points with real-world examples suitable for first-year engineering students.

Each refinement improves clarity and output quality.

Testing and Adjustment

To ensure reliability, prompts should be tested multiple times to:

- Check consistency of responses
- Identify ambiguities
- Improve logical accuracy
- Enhance clarity

Adjusting prompts based on observed weaknesses helps achieve more accurate and stable results.

Benefits of Effective Prompt Engineering

- Improves accuracy
- Reduces irrelevant output
- Saves time
- Enhances creativity
- Produces consistent results
- Increases reliability of AI systems

Q4 (c) Build a creative prompt for futuristic transportation ideas and explain

A creative prompt can be:

“Imagine transportation systems in the year 2050. Propose innovative, eco-friendly transportation methods and explain their benefits.”

This prompt promotes imaginative thinking by encouraging open-ended ideas, future-oriented solutions, and creativity without imposing strict limitations.

Q5 (a) Is labelled data supervised or unsupervised? Explain types of ML

Labelled data is used in **supervised learning**, where each input is associated with a known output. In supervised learning, models are trained using labeled datasets for tasks such as classification and regression.

In **unsupervised learning**, data is unlabeled, and the system identifies hidden patterns, such as clustering. **Reinforcement learning** involves learning through interaction with the environment using rewards and penalties to optimize actions.

Q5 (b) Outline K-Means algorithm with merits and demerits

K-Means is an unsupervised clustering algorithm that groups data into K clusters. The algorithm initializes K centroids, assigns data points to the nearest centroid, recalculates centroids, and repeats the process until convergence.

The advantages of K-Means include simplicity, efficiency, and scalability. However, it is sensitive to the choice of K, affected by initial centroid selection, and performs poorly with non-spherical clusters.

Q5 (c) Summarize the features of Reinforcement Learning and Support Vector Machines

Reinforcement Learning involves an agent interacting with an environment and learning optimal actions based on rewards. It does not require labeled data and follows a trial-and-error approach.

Support Vector Machines are supervised learning algorithms used for classification and regression. They work by finding an optimal hyperplane that maximizes the margin between data points. SVMs are effective in high-dimensional spaces and can handle non-linear data using kernel functions.

Q6 (a) Explain the four steps to create Decision Trees with suitable example

A Decision Tree is a supervised learning algorithm used for classification and regression. It represents decisions in the form of a tree structure consisting of nodes and branches.

The **first step** is **selecting the best attribute** for splitting the data. This is done using measures such as Information Gain, Gini Index, or Gain Ratio. For example, in a weather dataset, "Outlook" may be selected as the root attribute if it provides maximum information gain.

The **second step** is **splitting the dataset** based on the selected attribute. Each branch represents a possible value of the attribute. For instance, "Outlook" may have values such as Sunny, Rainy, and Overcast.

The **third step** involves **creating decision nodes and leaf nodes**. If all instances in a subset belong to the same class, a leaf node is created. Otherwise, the process continues recursively.

The **fourth step** is **stopping or pruning the tree** to avoid overfitting. The process stops when all attributes are used or when further splitting does not improve accuracy.

Q6 (b) Define conditional probability and Bayes' rule with examples

Conditional probability is the probability of an event occurring given that another event has already occurred. It is denoted as $P(A|B)$, where A and B are events. It is mathematically defined as

$$P(A|B) = P(A \cap B) / P(B).$$

For example, the probability that a person has a disease given that they tested positive is a conditional probability.

Bayes' Rule is used to calculate the probability of a hypothesis given observed evidence. It is expressed as

$$P(A|B) = [P(B|A) \times P(A)] / P(B).$$

For example, Bayes' theorem is widely used in medical diagnosis, spam filtering, and recommendation systems to update probabilities based on new evidence.

Q6 (c) How does Computer Vision work with Deep Learning? Explain the tasks involved

Computer Vision enables machines to interpret and understand visual data. With deep learning, convolutional neural networks (CNNs) are used to automatically extract features from images.

The process begins with **image acquisition**, followed by **preprocessing** such as resizing and normalization. CNN layers perform **feature extraction** using convolution and pooling operations. Finally, fully connected layers perform **classification or prediction**.

Common computer vision tasks include **image classification**, **object detection**, **image segmentation**, **face recognition**, and **optical character recognition (OCR)**. Deep learning improves accuracy by learning hierarchical features directly from data.

Q7 (a) List and explain any four Trusted AI principles

Trusted AI ensures that AI systems are ethical, reliable, and safe.

Transparency ensures that AI decisions are explainable and understandable by humans.

Fairness ensures that AI systems do not discriminate against individuals or groups.

Accountability ensures that responsibility can be assigned for AI outcomes.

Privacy and Security ensure that personal data is protected and securely handled.

Q7 (b) Explain AI as a Service (AlaaS). Advantages and challenges

AI as a Service (AlaaS) is a cloud-based model where AI tools and services are provided over the internet.

Advantages include **cost-effectiveness**, as no heavy infrastructure is needed, and **scalability**, allowing organizations to scale AI solutions easily.

Challenges include **data privacy concerns** and **dependency on cloud service providers**, which may limit control over AI systems.

Q7 (c) Relate the role of ethics in AI

Ethics in Artificial Intelligence (AI)

Ethics in AI refers to the moral principles and guidelines that govern the development, deployment, and use of artificial intelligence systems. As AI technologies increasingly influence decision-making in healthcare, finance, education, law enforcement, and social media, ensuring ethical standards becomes essential to protect individuals and society.

Ethical AI ensures that AI systems are designed and used responsibly, safely, and in alignment with human values.

Fairness

Fairness in AI means that systems should make decisions **without discrimination or bias** based on gender, race, religion, age, or other personal attributes.

AI models are trained on data, and if the data contains historical bias, the model may produce biased outcomes.

Example:

- A recruitment AI system should not favor one gender over another.
- A loan approval system should not discriminate against specific communities.

Transparency

Transparency refers to making AI systems understandable and explainable. Users should know:

- How decisions are made
- What data is being used
- Why a particular output was generated

This concept is also known as **Explainable AI (XAI)**.

Example:

If an AI system denies a loan application, it should provide clear reasons rather than giving a black-box decision.

Transparency builds trust between AI systems and users.

Privacy

AI systems often rely on large amounts of personal data. Ethical AI ensures:

- Secure data storage
- Protection against unauthorized access
- Responsible data usage
- Compliance with data protection regulations

Example:

Healthcare AI systems must protect patient records and sensitive medical information.

Strong privacy measures prevent data misuse and identity theft.

Accountability

Accountability means that developers, organizations, and stakeholders must take responsibility for AI system outcomes.

If an AI system causes harm or makes incorrect decisions:

- There should be clear responsibility
- Corrective measures must be implemented
- Legal and ethical consequences should apply

AI systems should not operate without human oversight.

Q8 (a) What is an Expert System? Explain its components

An **Expert System** is a computer-based artificial intelligence system designed to simulate the decision-making ability of a human expert in a specific domain. It uses stored knowledge and logical reasoning to solve complex problems that typically require human expertise.

Expert systems are widely used in fields such as:

- Medical diagnosis
- Financial analysis
- Troubleshooting
- Engineering design
- Legal advisory systems

The primary goal of an expert system is to provide accurate advice or decisions similar to those made by a trained professional.

Knowledge Base

The **knowledge base** is the core component of an expert system. It stores domain-specific knowledge required to solve problems.

It contains:

- Facts (basic information about the domain)
- Rules (IF–THEN rules for decision-making)
- Heuristics (experience-based techniques)

For example, in a medical expert system:

- Fact: Patient has fever and cough
- Rule: IF fever AND cough THEN possible flu

Inference Engine

The **inference engine** is the reasoning mechanism of the expert system. It applies logical rules to the knowledge stored in the knowledge base to derive conclusions.

It performs:

- Rule matching
- Logical reasoning
- Decision-making

User Interface

The **user interface (UI)** allows communication between the user and the expert system.

Through the interface, users can:

- Input data
- Answer system questions
- Receive advice or recommendations

Q8 (b) Differentiate between AI Programs and Robots, and human-controlled and autonomous bots

AI programs are software-based systems that perform intelligent tasks, while robots are physical machines capable of interacting with the environment.

Human-controlled bots require human intervention for decision-making, whereas fully autonomous bots operate independently using AI algorithms and sensors.

Q8 (c) Explain Neuromorphic Computing and its contribution to AGI

Neuromorphic Computing is inspired by the structure and functioning of the human brain. It uses spiking neural networks and specialized hardware to mimic neurons and synapses.

Its architecture supports parallel processing, low power consumption, and real-time learning. By closely resembling human cognitive processes, neuromorphic computing contributes toward achieving Artificial General Intelligence (AGI).

Q9 (a) Role of AI in Biology and Environmental Sciences

Role of AI in Biology

Artificial Intelligence plays a transformative role in biological sciences by analyzing complex biological data and accelerating research.

- **Genomics and DNA Analysis**

AI helps in sequencing, gene identification, and mutation detection. Machine learning models analyze massive genomic datasets quickly and accurately.

- **Drug Discovery**

AI predicts molecular interactions and identifies potential drug candidates, reducing research time and cost.

- **Protein Structure Prediction**

Deep learning models can predict protein folding patterns, which is crucial in understanding diseases and designing medicines.

- **Medical Imaging & Diagnosis**

AI assists in analyzing microscopic images, detecting abnormalities in cells, and diagnosing diseases.

Role of AI in Environmental Sciences

AI supports sustainable environmental management and climate monitoring.

- **Climate Change Prediction**

AI models analyze temperature patterns, greenhouse gas levels, and weather data to predict climate changes.

- **Pollution Monitoring**

AI systems track air and water pollution levels using sensor data and satellite imagery.

- **Wildlife Conservation**

AI-powered image recognition helps track endangered species and prevent poaching.

- **Natural Disaster Prediction**

AI predicts floods, earthquakes, and storms by analyzing historical and real-time environmental data.

Q9 (b) What is No-Code AI? Why must it be used?

No-Code AI refers to AI development platforms that allow users to build, train, and deploy machine learning models without writing programming code.

These platforms use:

- Drag-and-drop interfaces
- Automated workflows
- Pre-built algorithms

No-Code AI Must Be Used

- **Accessibility:** Enables non-programmers (business analysts, marketers, researchers) to use AI.
 - **Faster Development:** Reduces time required to develop AI models.
 - **Cost-Effective:** Decreases dependency on highly skilled developers.
 - **Democratization of AI:** Makes AI available to small businesses and startups.
 - **Rapid Prototyping:** Allows quick experimentation and testing.
-

Q9 (c) Role of AI in scientific experimentation

AI significantly enhances scientific research by automating experimentation and improving analytical accuracy.

- **Data Analysis**

AI processes large experimental datasets quickly and identifies patterns that humans may overlook.

- **Hypothesis Generation**

AI models suggest new research hypotheses based on previous findings.

- **Simulation and Modeling**

AI simulates complex physical, chemical, and biological processes.

- **Automation in Laboratories**

Robotic systems powered by AI conduct experiments with high precision.

- **Error Reduction**

AI minimizes manual errors and improves reproducibility.

Q10 (a) Role of AI in early disease prevention

AI helps in early disease prevention by analyzing medical data, identifying risk factors, and predicting disease outbreaks. It enables early diagnosis and personalized treatment plans. AI plays a crucial role in identifying diseases at early stages, improving survival rates.

- **Predictive Analytics**

AI analyzes patient data to predict disease risk.

- **Early Detection**

AI detects abnormalities in medical images before symptoms become severe.

- **Wearable Devices**

AI monitors heart rate, blood sugar, and other health parameters in real time.

- **Personalized Medicine**

AI suggests preventive measures based on individual health profiles.

- **Outbreak Prediction**

AI analyzes public health data to predict disease outbreaks.

Q10 (b) What is Low-Code AI? Compare with traditional tools

Low-Code AI refers to artificial intelligence development platforms that require only minimal programming knowledge. These platforms provide visual interfaces, drag-and-drop tools, automated workflows, and pre-built machine learning models. While users may write small amounts of code for customization, most of the development process is simplified through graphical tools.

Comparison of Low-code with the traditional tools:

Coding Requirement

Low-Code AI requires minimal coding. Users can build, train, and deploy models using visual interfaces and pre-configured components. In contrast, traditional AI tools require extensive programming knowledge in languages such as Python or R, along with expertise in machine learning libraries.

Development Speed

Low-Code AI platforms enable faster model development because much of the technical complexity is automated. Traditional AI development is comparatively slower since developers must write code from scratch, preprocess data manually, and fine-tune algorithms.

Accessibility

Low-Code AI is more accessible to beginners, business analysts, domain experts, and non-technical professionals. Traditional AI tools are mainly suitable for experienced data scientists and software engineers due to their technical complexity.

Customization and Flexibility

Traditional AI tools offer higher customization and flexibility. Developers can design complex architectures and implement advanced algorithms. Low-Code AI provides moderate customization but may have limitations when handling highly specialized or complex tasks.

Cost

Low-Code AI can reduce development costs by minimizing the need for highly skilled programmers and shortening project timelines. Traditional AI development may be more expensive due to longer development cycles and the requirement for expert professionals.

Q10 (c) Role of AI in Environmental Science

AI supports environmental protection and sustainability through intelligent analysis and prediction.

- **Climate Modeling**

AI predicts temperature changes and global warming patterns.

- **Renewable Energy Optimization**

AI improves efficiency of solar and wind energy systems.

- **Smart Agriculture**

AI optimizes irrigation, crop monitoring, and soil analysis.

- **Waste Management**

AI improves recycling processes using image recognition.

- **Ecosystem Monitoring**

AI tracks deforestation and biodiversity changes using satellite data.
