

Code: 23AM3602

III B.Tech - II Semester - Regular Examinations - APRIL 2026**NATURAL LANGUAGE PROCESSING
(CSE - AIML)**

Duration: 3 hours

Max. Marks: 70

- Note: 1. This question paper contains two Parts A and B.
 2. Part-A contains 10 short answer questions. Each Question carries 2 Marks.
 3. Part-B contains 5 essay questions with an internal choice from each unit. Each Question carries 10 marks.
 4. All parts of Question paper must be answered in one place.

BL – Blooms Level

CO – Course Outcome

PART – A

		BL	CO
1.a)	What is tokenization in natural language processing?	L2	CO1
1.b)	Explain Unigrams and trigrams?	L2	CO1
1.c)	What is POS tagging?	L2	CO1
1.d)	What is TF-IDF?	L2	CO1
1.e)	Explain positional encoding in BERT.	L2	CO1
1.f)	State the purpose of Named Entity Recognition.	L2	CO1
1.g)	Explain attention mechanism in transformers.	L2	CO1
1.h)	What is the role of the encoder-decoder architecture in sequence-to-sequence tasks?	L2	CO1
1.i)	Mention one ethical issue associated with large language models.	L2	CO1
1.j)	What is unidirectional context in GPT models?	L2	CO1

PART – B

			BL	CO	Max. Marks
UNIT-I					
2	a)	Explain the phases of NLP with examples using sequential diagram.	L2	CO1	5 M
	b)	Discuss any applications of NLP in day-to-day life.	L2	CO1	5 M
OR					
3	a)	Differentiate between stemming and lemmatization.	L2	CO1	5 M
	b)	Discuss the evolution and history of NLP.	L2	CO1	5 M
UNIT-II					
4	a)	Explain text preprocessing steps in detail.	L2	CO1	5 M
	b)	Apply Bag of Words, TF-IDF, and word embeddings with an example.	L3	CO2	5 M
OR					
5	a)	Explain language models and N-gram models.	L2	CO1	5 M
	b)	Compare and contrast sentence Segmentation and Word segmentation highlighting the differences in their applications.	L3	CO2	5 M
UNIT-III					
6	a)	Elaborate BERT architecture and training process.	L2	CO1	5 M

	b)	Analyze bidirectional context and applications of BERT.	L4	CO4	5 M
OR					
7	a)	Apply BERT in improving language understanding in NLP.	L3	CO3	5 M
	b)	Analyze Next Sentence Prediction (NSP) concept along with Use cases and limitations.	L4	CO4	5 M
UNIT-IV					
8	a)	Compare and contrast encoder and decoder parts of transformers.	L4	CO4	5 M
	b)	Briefly explain all the concepts below a) Learning rate warm-up, b) Adam optimizer, and c) gradient clipping.	L2	CO1	5 M
OR					
9	a)	Explain input embedding in encoder's part of transformer.	L2	CO1	5 M
	b)	Apply encoder-decoder mechanism in transformer.	L3	CO3	5 M
UNIT-V					
10	a)	Outline GPT models architecture and training process.	L2	CO1	5 M
	b)	Compare and contrast GPT model with latest versions of it.	L4	CO4	5 M

OR

11	a)	Analyze the ethical challenges associated with the use of GPT in content creation industries.	L4	CO4	5 M
	b)	Summarize the ethical issues associated with using GPT in sensitive domains like healthcare.	L2	CO1	5 M

III B.Tech

II Semester - Regular Examinations - APRIL 2026**NATURAL LANGUAGE PROCESSING (CSE - AIML)**

Max. Marks: 70

Time: 3 hours

1. This question paper attains two Parts A and B.
2. Part-A contains 10 short answer questions. Each Question carries 2 Marks.
3. Part-B contains 5 essay questions with an internal choice from each unit. Each Question carries 10 marks
4. All parts of Question naper must be answered in one place.

BL-Blooms Level

CO-Course Outcome

1.a) What is tokenization in natural language processing?

Tokenization is the process of breaking text into smaller units called **tokens** (words, subwords, or sentences) for easier processing.

1.b) Explain Unigrams and trigrams?

- **Unigram:** Single word unit (e.g., "I", "love")
- **Trigram:** Sequence of three words (e.g., "I love NLP")

1.c) What is POS tagging?

Part-of-Speech tagging assigns grammatical labels (noun, verb, adjective, etc.) to each word in a sentence.

1.d) What is TF-IDF?

TF-IDF (Term Frequency–Inverse Document Frequency) measures the importance of a word in a document relative to a collection of documents.

1.e) Explain positional encoding in BERT.

Positional encoding adds information about the **position of words** in a sequence so the model understands word order.

1.f) State the purpose of Named Entity Recognition.

NER identifies and classifies entities like **names, locations, organizations, dates** in text.

1.g) Explain attention mechanism in transformers.

Attention allows the model to focus on **important words** in a sentence while processing each word.

1.h) What is the role of the encoder-decoder architecture in sequence-to-sequence tasks?

The encoder converts input into a context representation, and the decoder generates the output sequence from it.

1.i) Mention one ethical issue associated with large language models.

- **Bias in outputs** due to biased training data.
- Privacy
- Hallucination

1.j) What is unidirectional context in GPT models?

GPT uses **left-to-right context**, meaning each word is predicted based only on previous words.

2.a) Explain the phases of NLP with examples [5 Mark]

1. **Lexical Analysis** – Breaking text into tokens [1 Mark]
2. **Syntactic Analysis (Parsing)** – Analyzing grammatical structure [1 Mark]
3. **Semantic Analysis** – Understanding meaning [1 Mark]
4. **Discourse Integration** – Linking sentences [1 Mark]
5. **Pragmatic Analysis** – Understanding intent [1 Mark]

2.b) Applications of NLP in day-to-day life. (Any 3)-[5 Mark]

- **Chatbots & Virtual Assistants** (Alexa, Siri)
- **Machine Translation** (Google Translate)
- **Sentiment Analysis** (reviews, social media)
- **Spam Detection** (emails filtering)
- **Speech Recognition** (voice typing)

3.a) Difference between Stemming and Lemmatization [5 Mark]

1 Mark for each valid point and 2 Marks for example

Feature	Stemming	Lemmatization
Approach	Rule-based cutting	Dictionary-based
Output	May not be real word	Valid word
Accuracy	Less	More
Example	running → runn	running → run

3.b) Discuss the evolution and history of NLP [5 Marks]

- **1950s–60s:** Rule-based systems
- **1980s–90s:** Statistical methods (HMM, n-grams)
- **2000s:** Machine Learning approaches
- **2010s–present:** Deep Learning & Transformers (BERT, GPT)
→ Shift from manual rules to data-driven models

4.a) Explain text preprocessing steps in detail (Any 3)-[5 Mark]

- Tokenization
- Stop word removal
- Stemming/Lemmatization
- Lowercasing

- Removing punctuation
- Normalization

4.b) Apply Bag of Words, TF-IDF, and word embeddings with an example [5 Mark]

- **Bag of Words (BoW):** Counts word frequency
- **TF-IDF:** Assigns importance based on frequency across documents
- **Word Embeddings:** Dense vector representations capturing meaning

5.a) Explain language models and N-gram models [5 Mark]

A **language model** predicts the probability of word sequences.

N-gram model:

- Uses previous (n-1) words to predict next word
- Example:
P(word | previous words)

Types:

- Unigram, Bigram, Trigram

5.b) Compare and contrast sentence segmentation and word segmentation, highlighting differences in their applications [5 Marks]

- **Word Segmentation:** Splitting into words
Example: "I love NLP" → [I, love, NLP]
- **Sentence Segmentation:** Splitting into sentences
Example: "Hi. How are you?" → [Hi, How are you]

6.a) Elaborate BERT architecture and training [5 Marks]

- Based on Transformer encoder
- Uses bidirectional context
- Input: token + segment + position embeddings
- Pre-training tasks:
 - Masked Language Model (MLM)
 - Next Sentence Prediction (NSP)
- Fine-tuned for tasks like QA, classification

6.b) Analyze bidirectional context and applications of BERT [5 Marks]

Answer:

- **BERT** reads text in both left-to-right and right-to-left directions.
- Unlike traditional models, it considers both previous and next words simultaneously.
- This helps in better understanding context and resolving ambiguity.

7.a) Apply BERT in language understanding

- Captures bidirectional context

- Improves:
 - Question answering
 - Sentiment analysis
 - NER
- Provides contextual embeddings

7.b) Analyze Next Sentence Prediction (NSP) with use cases and limitations.[5 Marks]

- Determines if sentence B follows sentence A
- Example:
 - A: "I am hungry."
 - B: "I will eat food." → TRUE

Limitation: Not always useful for all tasks

8.a) Compare and contrast encoder and decoder parts of transformers [5 Marks]

Feature	Encoder	Decoder
Input	Full sentence	Previous outputs
Attention	Self-attention	Masked + encoder-decoder attention
Task	Understanding	Generation

8.b) Explain: (a) Learning rate warm-up (b) Adam optimizer (c) Gradient clipping [5 Marks]

- **Learning rate warmup:** Gradually increases learning rate
- **Adam optimizer:** Adaptive gradient method
- **Gradient clipping:** Prevents exploding gradients

9.a) Explain input embedding in encoder part of transformer [5 Marks]

- Converts words into vectors
- Adds positional encoding
- Passed to attention layers

9.b) Apply encoder-decoder mechanism in transformer [5 Marks]

- Encoder encodes input sequence
- Decoder generates output step-by-step
- Uses attention to focus on relevant parts

10.a) Outline GPT model architecture and training process. [5 Marks]

- Decoder-only Transformer
- Uses unidirectional context
- Trained using next word prediction
- Applications: text generation, chatbots

10.b) Compare GPT with latest versions [5 Marks]

Feature	GPT	Latest Versions
Context	Limited	Larger context
Accuracy	Moderate	High
Training data	Smaller	Massive
Capability	Basic tasks	Advanced reasoning

11.a) Ethical challenges of GPT in content creation industries [5 Marks]

Answer (any 3 points):

- **biased content.**
- Can generate **misinformation or fake content.**
- **Privacy risks** – may expose sensitive data.
- **Lack of accountability** for harmful outputs.
- **Job displacement** in content and media industries.
- **Low transparency** in how decisions are made.

11.b) Ethical issues of GPT in healthcare

[5 Marks]

Answer (any 3 points):

- **Accuracy issues** – wrong medical advice can be harmful.
- **Privacy concerns** – patient data must be protected.
- **Bias** – may lead to unfair treatment suggestions.
- **Lack of accountability** for errors.
- **Over-reliance on AI** reduces human judgment.
- **Need for human oversight** in critical decisions.

II Semester - Regular Examinations - APRIL 2026

NATURAL LANGUAGE PROCESSING (CSE - AIML)

Max. Marks: 70

Time: 3 hours

1. This question paper attains two Parts A and B.
2. Part-A contains 10 short answer questions. Each Question carries 2 Marks.
3. Part-B contains 5 essay questions with an internal choice from each unit. Each Question carries 10 marks
4. All parts of Question naper must be answered in one place.

BL-Blooms Level

CO-Course Outcome

PART-A

1.a) What is tokenization in natural language processing?

Tokenization is the process of breaking text into smaller units called tokens such as words or sentences.

1.b) Explain Unigrams and trigrams?

N-grams are sequences of words used in language modeling. A unigram consists of a single word, while a trigram is a sequence of three words.

1.c) What is POS tagging?

Part-of-Speech tagging assigns grammatical labels like noun, verb, or adjective to each word in a sentence. It helps in understanding sentence structure and improves parsing.

1.d) What is TF-IDF?

TF-IDF measures the importance of a word in a document relative to a collection of documents by combining term frequency and inverse document frequency.

1.e) Explain positional encoding in BERT.

Positional encoding adds information about the order of words in a sequence since transformers process words in parallel.

1.f) State the purpose of Named Entity Recognition.

NER identifies and classifies entities such as names, locations, and organizations in text. It is used for information extraction and knowledge structuring.

1.g) Explain attention mechanism in transformers.

Attention allows the model to focus on **important words** in a sentence while processing each word.

1.h) What is the role of the encoder-decoder architecture in sequence-to-sequence tasks?

The encoder-decoder architecture is used in sequence-to-sequence tasks where the encoder converts input into a context vector and the decoder generates the output sequence.

1.i) Mention one ethical issue associated with large language models.

Large language models raise ethical concerns such as bias in outputs, privacy risks, and generation of incorrect information.

1.j) What is unidirectional context in GPT models?


GPT models use unidirectional context, meaning each word is predicted based only on previous words. It processes text in a left-to-right manner.

PART-B

2.a) Explain the phases of NLP with examples using a sequential diagram [5 Marks]

Answer: [1 mark for each Phase]

1. **Lexical Analysis:**
Breaks text into tokens.
Example: "I love NLP" → [I, love, NLP]
2. **Syntactic Analysis (Parsing):**
Analyzes grammatical structure.
Example: Identifying Subject-Verb-Object.
3. **Semantic Analysis:**
Determines meaning of words.
Example: "Apple" → fruit/company.
4. **Discourse Integration:**
Links sentences together.
Example: Resolving pronouns across sentences.
5. **Pragmatic Analysis:**
Understands user intent.
Example: "Can you open the door?" → request.



2.b) Discuss applications of NLP in day-to-day life.

[5 Marks]

Answer: [Any 3 Application]

- **Chatbots and Virtual Assistants (Alexa, Siri):**
Used to understand user queries and provide intelligent responses through voice or text interaction.
- **Machine Translation (Google Translate):**
Converts text from one language to another while preserving meaning and context.
- **Sentiment Analysis (reviews, social media):**
Determines whether a text expresses positive, negative, or neutral opinions, useful for business insights.
- **Spam Detection (email filtering):**
Automatically identifies and filters unwanted or harmful emails based on text patterns.
- **Speech Recognition (voice typing):**
Converts spoken language into written text, enabling hands-free communication.
- **Text Summarization:**
Generates concise summaries from large documents, saving time in information processing.
- **Search Engines:**
Improves search results by understanding user intent and query context (e.g., Google search suggestions).

3.a) Difference between Stemming and Lemmatization

[5 Marks]

Answer: [3marks for difference + 2 marks for example]

Feature	Stemming	Lemmatization
Approach	Stemming removes suffixes using simple rules without understanding word meaning, which can sometimes produce incorrect forms.	Lemmatization uses vocabulary and linguistic rules to return the correct base form of a word, making it more reliable for real-world applications.
Output	May not be a valid word	Produces meaningful base word
Accuracy	Less accurate	More accurate
Speed	Faster (less computation)	Slower (requires lookup)
Context awareness	Does not consider context	Considers context and part of speech
Example	running → runn	running → run

3.b) Discuss the evolution and history of NLP

[5 Marks]

Answer:

• **1950s–60s: Rule-based Systems**

Early NLP systems relied on manually written linguistic rules and grammar. These systems were rigid, domain-specific, and difficult to scale.

Example: Early machine translation systems.

- **1980s–90s: Statistical Methods**

NLP shifted to probabilistic models like Hidden Markov Models (HMM) and N-grams. These models used large datasets to learn patterns instead of relying on rules. Example: Speech recognition using probabilistic models.

- **2000s: Machine Learning Approaches**

Algorithms like Support Vector Machines (SVM) and Naive Bayes were used for tasks such as text classification and spam detection. Feature engineering played a key role.

- **2010s–Present: Deep Learning & Transformers**

Introduction of neural networks, word embeddings (Word2Vec, GloVe), and transformer models like BERT and GPT. These models capture context and semantics more effectively.

4.a) Explain text preprocessing steps in detail

[5 Marks]

Answer: [Any 3 pre-processing techniques]

- **Tokenization:**

Splits text into smaller units such as words or sentences, making it easier to analyze.

- **Stop-word Removal:**

Removes common words like “the”, “is”, “and” that do not add significant meaning.

- **Stemming / Lemmatization:**

Reduces words to their root or base form to normalize variations.

(e.g., “running” → “run”)

- **Lowercasing:**

Converts all text into lowercase to maintain uniformity and avoid duplication (e.g., “NLP” and “nlp”).

- **Removing Punctuation:**

Eliminates symbols like commas, periods, and special characters that are not useful for analysis.

- **Normalization:**

Standardizes text by handling numbers, abbreviations, and special formats.

4.b) Apply Bag of Words, TF-IDF, and word embeddings with an example

[5 Marks]

Answer:

1. Bag of Words (BoW):

- Represents text as a collection of word frequencies.
- Ignores grammar and word order, focusing only on occurrence.
- Example:
Sentence: “I love NLP” → {I:1, love:1, NLP:1}

2. TF-IDF (Term Frequency–Inverse Document Frequency):

- Measures the importance of a word in a document relative to a collection.
- Words that appear frequently in one document but rarely in others get higher importance.
- Example:
Common word “the” → low score, rare word “NLP” → high score

3. Word Embeddings:

- Represent words as dense, continuous vectors in a low-dimensional space.
- Capture semantic relationships and context between words.

- Example:
“king” – “man” + “woman” \approx “queen”

5.a) Explain language models and N-gram models

[5 Marks]

Answer:

A language model is a probabilistic model that predicts the likelihood of a sequence of words, helping machines understand and generate human language.

- An N-gram model predicts the next word based on the previous $(n-1)$ words.
- It assumes that the probability of a word depends only on a limited context.

Types of N-grams:

- **Unigram:** Considers single words independently
- **Bigram:** Considers one previous word
- **Trigram:** Considers two previous words

5.b) Compare and contrast sentence segmentation and word segmentation, highlighting differences in their applications

[5 Marks]

Answer:

Word Segmentation:

- Word segmentation involves splitting a sentence into individual words or tokens.
- It is the first step in many NLP tasks like text classification and parsing.
- Essential for languages where spaces are not clearly defined (e.g., Chinese).
- Helps in further processes like POS tagging and parsing.
- Example:
“I love NLP” \rightarrow [I, love, NLP]

Sentence Segmentation:

- Sentence segmentation involves dividing a paragraph or text into individual sentences.
- It helps in understanding the structure and meaning of a document.
- Useful for tasks like summarization, translation, and discourse analysis.
- Often uses punctuation marks like periods, question marks, etc.
- Example:
“Hi. How are you?” \rightarrow [Hi, How are you]

6.a) Elaborate BERT architecture and training

[5 Marks]

Answer:

BERT (Bidirectional Encoder Representations from Transformers) is a transformer-based model designed to understand the deep contextual meaning of words by considering both left and right context simultaneously.

- Each input token is represented using:
 - **Token embeddings** (word representation)

- **Segment embeddings** (sentence distinction)
- **Positional embeddings** (word order information)

Attention and feed forward network

- Pre-training tasks:
 - Masked Language Model (MLM)
 - Next Sentence Prediction (NSP)

6.b) Analyze bidirectional context and applications of BERT

[5 Marks]

Answer:[3 marks for bidirectional +2 marks for any 2 application]

BERT (Bidirectional Encoder Representations from Transformers) is a transformer-based model that uses bidirectional context to understand the meaning of words more effectively in a sentence.

- Traditional models process text either **left-to-right** or **right-to-left**, limiting context understanding.
- BERT reads text **in both directions simultaneously**, allowing it to consider both previous and next words.
- This helps in resolving ambiguity and understanding deeper relationships between words.

Applications of BERT:

- **Question Answering:** Extracts precise answers from passages.
- **Sentiment Analysis:** Classifies text as positive, negative, or neutral.
- **Named Entity Recognition (NER):** Identifies entities like names, places, organizations.
- **Text Classification:** Categorizes text into topics or classes.
- **Language Inference:** Determines relationships between sentences.

7.a) Apply BERT in improving language understanding in NLP

[5 Marks]

Answer:

BERT (Bidirectional Encoder Representations from Transformers) improves language understanding in NLP by generating contextual embeddings that capture the meaning of words based on their surrounding context. Unlike traditional models that process text in a single direction, BERT uses bidirectional processing, meaning it considers both previous and next words in a sentence. This helps in resolving ambiguity and understanding the correct meaning of words in different contexts.

BERT is pre-trained using techniques like Masked Language Modeling (MLM), which allows it to learn deep relationships between words. As a result, it performs effectively in various NLP tasks. In **question answering**, BERT can extract precise answers from a given passage. In **sentiment analysis**, it accurately classifies text as positive, negative, or neutral.

In **Named Entity Recognition (NER)**, it identifies entities such as names, locations, and organizations. It is also widely used in **text classification** tasks like spam detection and topic categorization.

By capturing rich contextual and semantic information, BERT significantly improves the accuracy and performance of NLP systems compared to traditional models.

7.b) Analyze Next Sentence Prediction (NSP) with use cases and limitations. [5 Marks]

Answer:

Next Sentence Prediction (NSP) is a pre-training task used in BERT to help the model understand the relationship between pairs of sentences and capture discourse-level context.

- The model is given two sentences, **Sentence A** and **Sentence B**.
- It predicts whether Sentence B logically follows Sentence A in the original text.
- Helps the model understand **sentence relationships**.
- Improves performance in tasks like: Question answering
- Limitations is it only predicts don't generate.

Example:

A: "I am hungry."

B: "I will eat food." → **TRUE**

A: "I am hungry."

B: "The sky is blue." → **FALSE**

8.a) Compare and contrast encoder and decoder parts of transformers [5 Marks]

Answer:

Feature	Encoder	Decoder
description	The encoder processes the entire input sequence at once and creates contextual representations using self-attention.	The decoder generates output step-by-step, using masked attention to prevent access to future tokens and cross-attention to focus on encoder outputs.
Input	Full input sentence	Previously generated output tokens
Attention	Self-attention	Masked self-attention + encoder-decoder attention
Function	Encodes meaning of input	Generates output sequence
Context access	Can see entire input	Can only see past outputs (masked)
Usage	Used in models like BERT	Used in models like GPT

8.b) Explain: (a) Learning rate warm-up (b) Adam optimizer (c) Gradient clipping [5 Marks]

Answer:

Optimization techniques are essential in training deep learning models to ensure stable and efficient convergence.

- **Learning Rate Warmup:**

Gradually increases the learning rate during initial training steps to avoid sudden large updates and instability.

- **Adam Optimizer:**

An adaptive optimization algorithm that adjusts learning rates for each parameter using first and second moment estimates of gradients.

- **Gradient Clipping:**

Limits the magnitude of gradients to prevent exploding gradients, especially in deep networks.

9.a) Explain input embedding in encoder part of transformer

[5 Marks]

Answer:

- **Token Embedding:**

Each word or token is converted into a dense vector representing its meaning in a high-dimensional space.

- **Positional Encoding + Segment Embedding**

Since transformers do not process data sequentially, positional encoding is added to embeddings to provide information about the position of words in a sentence.

- **Combined Representation:**

The final input embedding is a sum of token embeddings and positional encodings.

9.b) Apply encoder-decoder mechanism in transformer

[5 Marks]

Answer:

The encoder-decoder mechanism is a fundamental architecture used in sequence-to-sequence tasks such as machine translation.

- **Encoder:**

Processes the entire input sequence and converts it into a contextual representation (context vector or hidden states).

- **Decoder:**

Generates the output sequence step-by-step using the encoded information.

- **Attention Mechanism:**

Allows the decoder to focus on relevant parts of the input sequence while generating each output token.

- **Sequential Generation:**

The decoder predicts one word at a time based on previous outputs and encoder context.

- **Example:** English → French translation

Input: "I love NLP"

Output: "J'aime le NLP"

10.a) Outline GPT model architecture and training process.

[5 Marks]

Answer:

GPT (Generative Pre-trained Transformer) is a powerful language model designed for text generation using transformer architecture.

- **Decoder-only Transformer:**
GPT uses only the decoder part of the transformer architecture.
- **Unidirectional Context:**
Processes text from left to right, predicting each word based on previous words.
- **Masked Self-Attention:**
Ensures the model does not access future tokens during training.
- **Next Word Prediction:**
The model is trained to predict the next word in a sequence using large-scale text data.

10.b) Compare GPT with latest versions

[5 Marks]

Answer:

Feature	GPT (Early Versions)	Latest Versions
Context	Limited context window	Larger context window (handles long text)
Accuracy	Moderate accuracy	High accuracy with better understanding
Training Data	Smaller datasets	Trained on massive and diverse datasets
Capability	Basic NLP tasks	Advanced reasoning, multi-task learning
Modalities	Text only	Text, images, code (multimodal)

11.a) Analyze the ethical challenges associated with the use of GPT in content creation industries

[5 Marks]

Answer: (any 3 points)

- **Bias and Fairness:** GPT may produce biased outputs, reinforcing harmful stereotypes.
- **Misinformation:** GPT can generate false or misleading information, spreading fake news.
- **Privacy:** Risk of revealing sensitive or private data unintentionally.
- **Accountability:** Unclear responsibility for harmful consequences of model outputs.
- **Job Displacement:** AI automation may lead to job loss in fields like content creation and customer service.
- **Lack of Transparency:** Difficulty in understanding model decisions, affecting trust.

11.b) Summarize the ethical issues associated with using GPT in sensitive domains like healthcare

[5 Marks]

Answer: (any 3 points)

1. **Accuracy and Reliability**
Incorrect medical advice can lead to harmful consequences.
2. **Privacy and Data Security**
Patient data must be protected; misuse can violate confidentiality.
3. **Bias in Medical Decisions**
Biased training data may lead to unequal treatment recommendations.
4. **Lack of Accountability**
Difficult to assign responsibility for incorrect diagnoses or advice.

5. **Over-reliance on AI**

Healthcare professionals may depend too much on AI, reducing critical judgment.

6. **Ethical Responsibility**

Decisions affecting human life require human oversight.