Roll No: | | | | | | | | | | | |

PRESIDENCY UNIVERSITY

GAIN MORE KNOWLEDGE
REACH GREATER HEIGHTS

Presidency University Act, 2013 of the Karnataka Act No. 41 of 2013 | Established under Section 2(f) of UGC Act, 1956
Approved by AICTE, New Delhi | Approved By BCI
Bengaluru

Mid Term Examinations - October 2026

Date: 08/10/2026 **Time:** 02:00 PM - 03:30 PM

Course Code: CSA4710 **Course Name:** Generative AI

Semester: Third Semester **Max. Marks:** 50 **Weightage:** 50%

CO - Levels	CO1	CO2	CO3
Marks	90	68	22

PART A: Answer Following Questions. 10 M

Q.N.	Questions	M	CO	BT
1	Define Generative AI.	2	CO1	BT1
2	Outline the notable challenges for training Generative AI models.	2	CO1	BT2
3	What is a Large Language Model (LLM)?	2	CO1	BT1
4	Summarize the typical workflow of an LLM-based application in brief.	2	CO3	BT2
5	Compare RNN and LSTM in brief.	2	CO2	BT2

PART B : Answer Any 1 Following Questions. 10 M

Q.N.	Questions	M	CO	BT
6	Consider the sample input “A sunset over the mountains” given to different types of generative AI models used for different data modalities — for example, a text-based model generating a descriptive caption, an image-based model generating a picture, and an audio-based model generating a matching background sound. Explain how each of these different types of generative models processes this input and produces its output, highlighting relevant factors such as efficiency, speed, and number of parameters.	10	CO1	BT2
7	Compare these four LLMs: GPT, PaLM 2, LLaMA and Claude in terms of their developer organization and key characteristics.	10	CO1	BT2

PART C: Answer Any 1 Following Questions. 10 M

Q.N.	Questions	M	CO	BT
8	Identify the components of a prompt namely, task, system instructions, few shot examples and context information with suitable examples for each.	10	CO1	BT3
9	Consider a simple vocabulary with 5 words: [“I”, “like”, “to”, “read”, “books”]. Each word is represented by a vector in a 3-dimensional space with the word embeddings as follows: **“I” -> [0.1,0.2,0.3]** **“like” -> [0.4,0.5,0.6]** **“to” -> [0.2,0.1,0.3]** **“read” -> [0.6,0.5,0.4]** **“books” -> [0.3,0.4,0.7]** Identify the sentence embedding for the sentence.	10	CO1	BT3

PART D: Answer Any 1 Following Questions. 10 M

PTO (Please Turn Over)

Q.N.	Questions	M	CO	BT
10	Following are the embedding matrix of an input sequence and the weight matrices. Solve for the attention weighted feature matrix Z of the input tokens as discussed in the transformer architecture. (Assume one attention head) **Embeddings : X =** $\begin{bmatrix} 2 & 3 & 8 \\ 3 & 2 & 1 \end{bmatrix}$ $W^Q =$ $\begin{bmatrix} 2 & 3 \\ 7 & 8 \\ 9 & 0 \end{bmatrix}$ $W^k =$ $\begin{bmatrix} 9 & 10 \\ 1 & 2 \\ 4 & 6 \end{bmatrix}$ $W^v =$ $\begin{bmatrix} 8 & 7 \\ 3 & 2 \\ 9 & 5 \end{bmatrix}$	10	CO2	BT3
11	Experiment the two phases of GPT namely Pre-training and fine-tuning for text classification with neat illustrations.	10	CO2	BT3

PART E: Answer Any 1 Following Questions. 10 M

Q.N.	Questions	M	CO	BT
12	Identify the various components of RAG in Langchain with neat illustrations.	10	CO3	BT3
13	You are given the following input text for the decoder: “The weather today is”. Use the following simplified vocabulary to predict the next word in the sequence: **Vocabulary: [“sunny”, “rainy”, “cloudy”, “snowy”, “windy”]** **Logits=[2.5,1.0,3.5,0.5,1.2]** Your Task is to find: 1. **Identify the next word predicted by the decoder.** Apply softmax to find the probabilities for each word in the vocabulary.	10	CO3	BT3