

JAYPEE UNIVERSITY OF INFORMATION TECHNOLOGY, WAKNAGHAT

TEST -2EXAMINATION- 2025

B.Tech-V Semester (CSE/IT)

COURSE CODE (CREDITS): 25B1WCI511 (2)

MAX. MARKS: 25

COURSE NAME: PROMPT ENGINEERING

COURSE INSTRUCTORS: VANI SHARMA

MAX. TIME: 1 Hour 30 Min

Note: (a) All questions are compulsory.

(b) The candidate is allowed to make Suitable numeric assumptions wherever required for solving problems.

(c) The use of calculator is allowed.

Q.No	Question	CO	Marks												
Q1	Consider a document containing 100 words wherein the word cat appears 3 times. Now, assume we have 10 million documents, and the word cat appears in one thousand of these. Now compute (a) TF, (b) IDF, and (c) TF-IDF.	2	[3]												
Q2	Explain in detail the concepts of self-attention mechanisms in Transformer models. Discuss thoroughly how query, key, and value vectors interact to compute attention probabilities.	2	[6]												
Q3	a) Discuss the limitations of self-attention that lead to the inclusion of additional components in the Transformer architecture. b) How does masked decoding prevent information leakage during language model training?	2	[3+2]												
Q4	What are the main components of a GRU, and how does it solve the problems of traditional RNNs?	2	[3]												
Q5	A language model at one decoding timestep assigns the following probabilities to 5 candidate tokens (already softmaxed): <table border="1"><thead><tr><th>Token</th><th>Probability</th></tr></thead><tbody><tr><td>A</td><td>0.40</td></tr><tr><td>B</td><td>0.25</td></tr><tr><td>C</td><td>0.15</td></tr><tr><td>D</td><td>0.12</td></tr><tr><td>E</td><td>0.08</td></tr></tbody></table>	Token	Probability	A	0.40	B	0.25	C	0.15	D	0.12	E	0.08	2	[3]
Token	Probability														
A	0.40														
B	0.25														
C	0.15														
D	0.12														
E	0.08														

	(a) For $\text{top-p} = 0.8$, determine the nucleus set (the smallest set of tokens whose cumulative probability $\geq p$). (b) Renormalize the probabilities over the nucleus and show the resulting sampling distribution.		
Q6.	a) What is the main purpose of instruction tuning in large language models? b) Suppose you are building a document-based QA system using prompt chaining. • Outline the sequence of prompts you would design (e.g., retrieval $\rightarrow$ summarization $\rightarrow$ answer generation). • Explain how chaining improves transparency and control over outputs.	3	[2+3]