

MADANAPALLE INSTITUTE OF TECHNOLOGY & SCIENCE, MADANAPALLE

(UGC-AUTONOMOUS INSTITUTION)

B. Tech III Year II Semester (R23) Regular End Semester Examinations, May- 2026**BIG DATA ANALYTICS**

(CSE- AI & ML)

Time: 3Hrs**Max Marks: 70**

Attempt all the questions. All parts of the question must be answered in one place only.

All parts of Q.no 1 are compulsory. In Q.no 2 to 6 answer either A or B only

S.No.	Question	Marks	CO	BL
1.	i) Define Big Data.	1	1	1
	ii) Define structured data.	1	1	1
	iii) Name the two main components of HDFS.	1	2	1
	iv) List any two features of Hadoop.	1	2	1
	v) Define mapper.	1	3	1
	vi) List any two MapReduce applications.	1	3	1
	vii) Name the stages in Spark execution.	1	4	1
	viii) Define fault tolerance in Spark.	1	4	1
	ix) List types of NoSQL databases.	1	5	1
	x) Define capped collection.	1	5	1
2(A)	(i) List and explain the different types of Big Data with examples.	6	1	3
	(ii) Analyze the challenges in handling Big Data.	6	1	4
OR				
2(B)	Analyze the characteristics of Big Data (5Vs) and critically examine how each characteristic influences the design and performance of Big Data processing systems with suitable examples.	12	1	4
3(A)	(i) Demonstrate the role of YARN by explaining its architecture and illustrating how it manages resources and schedules tasks in a Hadoop cluster.	6	2	3
	(ii) Demonstrate the architecture of HDFS with a neat diagram.	6	2	3
OR				
3(B)	Demonstrate the process of setting up a Hadoop cluster by explaining configuration steps, hardware requirements, and deployment considerations.	12	2	3
4(A)	(i) Apply the MapReduce model to a simple example.	6	3	3
	(ii) Explain the mapper and reducer functions in data processing.	6	3	3
OR				
4(B)	Apply the MapReduce programming model to solve a word count problem by clearly explaining mapper, reducer, and intermediate processing steps.	12	3	3
5(A)	(i) Discuss the Resilient Distributed Datasets (RDDs).	6	4	2
	(ii) Compare and contrast Spark vs Hadoop MapReduce.	6	4	3
OR				
5(B)	Analyze Apache Spark architecture by explaining its components and examining how it improves performance compared to traditional MapReduce	12	4	4
6(A)	(i) Explain the MongoDB data types and structure.	6	5	2

(ii)	Analyze CRUD operations in MongoDB.	6	5	3
------	-------------------------------------	---	---	---

OR

6(B)	Evaluate NoSQL solutions for real-world applications by examining their suitability for use cases such as credit fraud detection or weather forecasting.	12	5	4
------	--	----	---	---

*****END*****

MADANAPALLE INSTITUTE OF TECHNOLOGY & SCIENCE, MADANAPALLE

(UGC-AUTONOMOUS INSTITUTION)

B. Tech III Year II Semester (R23) Regular End Semester Examinations, May – 2026**CLOUD COMPUTING FOR AI**

(CSE- AI&ML)

Time: 3Hrs**Max Marks: 70**

Attempt all the questions. All parts of the question must be answered in one place only.
All parts of Q. no. 1 are compulsory. In Q. nos. 2 to 6, answer either A or B only.

S.No.	Question	Marks	CO	BL
1.	i) What is a hybrid cloud, and what are its advantages?	1	1	1
	ii) Define NLP and its applications.	1	1	1
	iii) Mention the purpose of Azure Blob storage.	1	2	1
	iv) State the purpose of using Spark.	1	2	1
	v) State the purpose of the Azure ML pipeline.	1	3	1
	vi) Define ML pipeline.	1	3	1
	vii) What is a container and Docker?	1	4	1
	viii) List the types of autoscaling.	1	4	1
	ix) Define bias in AI.	1	5	1
	x) Name any one of the security issues in cloud computing.	1	5	1
2(A)	Examine IaaS, PaaS, and SaaS for AI workloads, highlighting trade-offs in scalability and cost.	12	1	3
OR				
2(B)	(i) Explain the following cloud characteristics with an example. (a) On-demand self-service (b) Elastic Computing (c) Resource pooling (d) Measured service	6	1	2
	(ii) Compare and contrast the on-premise and cloud computing environments.	6	1	4
3(A)	Compare and contrast data lakes and data warehouses with a neat diagram.	12	2	4
OR				
3(B)	Illustrate the role of virtualization in improving cloud efficiency. Explain any three types of virtualization with neat diagrams.	12	2	4
4(A)	Analyze GPU vs TPU for ML training, comparing performance and cost implications.	12	3	4
OR				
4(B)	Explain the differences between AutoML and custom machine learning workflows in cloud platforms. Discuss their advantages, limitations, and suitable use cases.	12	3	2
5(A)	Elaborate on the following key Characteristics of Cloud-Native Applications: i. Microservice Architecture ii. Containerization iii. DevOps and CI/CD	12	4	2

OR

5(B)	Describe the stages of a CI/CD pipeline for AI models and explain how it supports automated model development, testing, deployment, and monitoring in cloud environments, using a neat diagram.	12	4	2
6(A)	Analyze how cloud-based AI solutions are transforming the healthcare sector. Discuss the architecture, benefits, and challenges involved in deploying AI healthcare applications on cloud platforms.	12	5	4
OR				
6(B)	Evaluate the importance of cost management and billing strategies for AI services in cloud environments and discuss methods to optimize resource utilization and reduce operational costs.	12	5	4

*****END*****

MADANAPALLE INSTITUTE OF TECHNOLOGY & SCIENCE, MADANAPALLE
(UGC-AUTONOMOUS INSTITUTION)**B. Tech III Year II Semester (R23) Regular End Semester Examinations, May – 2026****DEEP LEARNING**

(CSE-AI&ML)

Time: 3Hrs**Max Marks: 70**

Attempt all the questions. All parts of the question must be answered in one place only.
All parts of Q.no 1 are compulsory. In Q.no 2 to 6 answer either A or B only

S.No.	Question	Marks	CO	BL
1.	i) Differentiate between global minima and local minima in an optimization landscape.	1	1	2
	ii) If a model has 0 training error but high-test error, which phenomenon is it likely exhibiting?	1	1	3
	iii) What is the range of the output of a Logistic Regression classifier?	1	2	1
	iv) Explain how maximum likelihood estimation works in finding optimal parameters.	1	2	2
	v) Explain the role of the "Summation" and "Activation" steps in a neuron.	1	3	2
	vi) List any four common activation functions used in Neural Networks.	1	3	1
	vii) Define a "Kernel" in the context of CNNs.	1	4	1
	viii) Given an input image of size 32×32, filter size 3×3, stride 1, and padding 0, compute the output feature map size.	1	4	3
	ix) Define the vanishing gradient problem in RNNs.	1	5	1
	x) Compare the computational complexity of RNNs vs. LSTMs.	1	5	4
2(A)	(i) Explain the significance of eigenvalues and eigenvectors in machine learning. Where are they applied?	6	1	2
	(ii) Describe the fundamental difference between Batch Gradient Descent and Stochastic Gradient Descent.	6	1	2
OR				
2(B)	Explain the complete taxonomy of errors in Deep learning.	12	1	2
3(A)	(i) Explain the basic concepts of regression problems. What are the key characteristics and evaluation metrics?	6	2	2
	(ii) Assess the effectiveness of logistic regression for non-linear classification problems.	6	2	3
OR				
3(B)	Explain in detail maximum likelihood estimation. How is it used in logistic regression?	12	2	2
4(A)	(i) Explain the single-layer perceptron model. What types of problems can it solve?	6	3	2
	(ii) Discuss why non-linear activation functions are necessary for deep networks.	6	3	2
OR				
4(B)	(i) Describe all major components of a deep neural network: neurons, layers, weights, biases, activation functions, forward propagation, and loss functions, with a neat architecture	12	3	2
5(A)	(i) Define stride and padding and explain their role in convolution operations.	6	4	2
	(ii) Explain the difference between L1 and L2 regularization effects.	6	4	2

OR

5(B)	Briefly explain famous CNN architectures (ResNet, GoogLeNet, MobileNet).	12	4	3
6(A)	(i) Explain backpropagation through time (BPTT) algorithm.	6	5	2
	(ii) Critically analyze the causes of vanishing and exploding gradients. Examine solutions and their effectiveness.	6	5	4
OR				
6(B)	Sketch the architecture and explain the working of LSTM and its applications.	12	5	3

*****END*****

MADANAPALLE INSTITUTE OF TECHNOLOGY & SCIENCE, MADANAPALLE

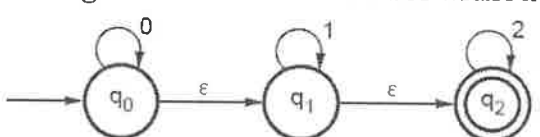
(UGC-AUTONOMOUS INSTITUTION)

B. Tech III Year II Semester (R23) Regular End Semester Examinations, May – 2026**AUTOMATA THEORY AND COMPILER DESIGN**

(CSE – AI & ML)

Time: 3Hrs**Max Marks: 70**

Attempt all the questions. All parts of the question must be answered in one place only.
All parts of Q.no 1 are compulsory. In Q.no 2 to 6 answer either A or B only

S.No.	Question	Marks	CO	BL
1.	i) Distinguish between DFA and NFA with example	1	1	2
	ii) When do you say a given grammar G is ambiguous?	1	1	1
	iii) Is the grammar $E \rightarrow E + E \mid id$ is ambiguous? Justify.	1	2	2
	iv) What is parse Tree?	1	2	2
	v) Differentiate between Turing machine and pushdown automata.	1	3	2
	vi) Justify whether PDA more powerful than a DFA.	1	3	2
	vii) How does top-down parsing handle recursion in a grammar?	1	4	1
	viii) State the role of lexical analyser.	1	4	1
	ix) Translate the arithmetic expression $a := b * - (c + d)$ into syntax tree.	1	5	2
	x) How does constant folding help in optimizing code?	1	5	1
2(A)	(i) Define a Finite Automaton (FA) and explain its components with an example.	6	1	2
	(ii) Differentiate between Deterministic Finite Automata (DFA) and Non-Deterministic Finite Automata(NFA).	6	1	2
OR				
2(B)	Convert the given NFA with ϵ to NFA without ϵ .	12	1	5
				
3(A)	(i) Describe the procedure to convert a DFA into an equivalent Regular Expression.	6	2	2
	(ii) Explain the conversion of a Regular Expression into DFA using subset construction method.	6	2	2
OR				
3(B)	Prove that the language $L = \{a^n b^n c^n \mid n \geq 0\}$ is not context-free using the pumping lemma for context-free languages.	12	2	4
4(A)	(i) Explain Deterministic Pushdown Automaton (DPDA) features and limitations.	6	3	2
	(ii) Write short notes on Equivalence of PDA's with example.	6	3	1
OR				
4(B)	Construct an empty store Push Down Automata (PDA) for the below mentioned language: $L(G) = \{w \mid w \in (a,b) \text{ and } w \text{ is of the form } a^n b^n \text{ and } n \geq 1\}$. Also mention the state transition of this PDA while parsing the string $w = 'aaabbb'$.	12	3	4
5(A)	(i) Describe bottom-up parsing with suitable example	6	4	2
	(ii) Differentiate between compiler and interpreter with example	6	4	2

OR

5(B)	Construct the LR(0) parsing table for the following grammar. $S \rightarrow cA \mid ccB$ $A \rightarrow cA \mid a$ $B \rightarrow ccB \mid b$	12	4	5
-------------	--	----	---	---

6(A)	(i) Compare Three Address Code and syntax trees in intermediate representation	6	5	2
	(ii) Explain how DAG representation helps in optimizing basic blocks	6	5	4

OR

6(B)	Construct the syntax tree and DAG for the following expression $x = (a+b) * (a+b+c) * (a+b+c+d)$	12	5	5
-------------	---	----	---	---

*****END*****

MADANAPALLE INSTITUTE OF TECHNOLOGY & SCIENCE, MADANAPALLE
(UGC-AUTONOMOUS INSTITUTION)**B. Tech III Year II Semester (R23) Regular End Semester Examinations, May – 2026**
REINFORCEMENT LEARNING
(CSE – AI&ML)

Time: 3Hrs

Max Marks: 70

Attempt all the questions. All parts of the question must be answered in one place only.
All parts of Q.no 1 are compulsory. In Q.no 2 to 6 answer either A or B only

S.No.	Question	Marks	CO	BL
1.	i) Define environment in RL.	1	1	1
	ii) What are optimistic initial values?	1	1	2
	iii) Define optimal policy.	1	2	1
	iv) What is dynamic programming?	1	2	1
	v) Define policy evaluation in Monte Carlo methods.	1	3	2
	vi) What is Q-learning?	1	3	1
	vii) Define Sarsa(λ).	1	4	2
	viii) What are replacing traces?	1	4	1
	ix) Define approximation error.	1	5	2
	x) What is a policy gradient?	1	5	1
2(A)	Explain the agent-environment interface in reinforcement learning with a neat diagram. Describe the four main elements of an RL system (policy, reward signal, value function, model) with suitable examples.	12	1	2
OR				
2(B)	Explain the history and development of Reinforcement Learning, including key milestones.	12	1	2
3(A)	(i) Write short notes on approximation methods in MDPs.	6	2	2
	(ii) Compare value iteration vs policy iteration.	6	2	3
OR				
3(B)	Explain the Markov Decision Process in detail, including its components, assumptions, and applications.	12	2	2
4(A)	(i) Differentiate between on-policy and off-policy learning.	6	3	3
	(ii) Discuss advantages and limitations of Monte Carlo methods.	6	3	2
OR				
4(B)	Compare SARSA, Q-learning, and R-learning with suitable examples.	12	3	3
5(A)	(i) Discuss the importance of λ in controlling bias and variance.	6	4	2
	(ii) Explain Q(λ) and how it differs from standard Q-learning.	6	4	2
OR				
5(B)	Derive and explain n-step TD prediction and its advantages over 1-step TD.	12	4	3
6(A)	(i) Discuss average reward setting in reinforcement learning.	6	5	2
	(ii) Compare tabular methods with function approximation methods.	6	5	3
OR				
6(B)	Explain Actor-Critic methods in detail with architecture and working.	12	5	2

*****END*****