

SRM VALLIAMMAI ENGINEERING COLLEGE

SRM Nagar, Kattankulathur – 603 203

DEPARTMENT OF ARTIFICIAL INTELLIGENCE & DATA SCIENCE & DEPARTMENT OF INFORMATION TECHNOLOGY

QUESTION BANK



V SEMESTER AD3562-MACHINE LEARNING Regulation – 2023

Academic Year 2025-2026 (Odd Semester)

Prepared by

Dr. S. Narayanan, Professor & Head, IT

Dr. S. Sandhya, Assistant Professor (Sr.G), IT

Ms. R. Lakshmi, Assistant Professor (Sel.G), AI-DS



SRM VALLIAMMAI ENGINEERING COLLEGE

SRM Nagar, Kattankulathur – 603 203.

DEPARTMENT OF INFORMATION TECHNOLOGY

QUESTION BANK



SUBJECT : AD3562-MACHINE LEARNING

SEM / YEAR: V / III

UNIT I – INTRODUCTION TO MACHINE LEARNING

Review of Linear Algebra for machine learning; Introduction and motivation for machine learning; Examples of machine learning applications, Vapnik-Chervonenkis (VC) dimension, Probably Approximately Correct (PAC) learning, Hypothesis spaces, Inductive bias, Generalization, Bias variance trade-off.

PART-A

Q.No	Questions	BT Level	Competence
1	List the role of eigenvalues and eigenvectors in machine learning.	BTL1	Remembering
2	Define the dot product and state its significance in machine learning.	BTL1	Remembering
3	Identify the rank of a matrix and state its importance.	BTL1	Remembering
4	Name how matrix inversion is used in linear regression.	BTL1	Remembering
5	Define the difference between scalar, vector, and matrix.	BTL1	Remembering
6	Summarize supervised learning with an example.	BTL2	Understanding
7	Classify the key differences between AI, machine learning, and deep learning.	BTL2	Understanding
8	State why machine learning is preferred over rule-based systems.	BTL1	Remembering
9	Outline two real-world applications of machine learning.	BTL2	Understanding
10	Differentiate between training and testing data.	BTL2	Understanding
11	Outline two machine learning applications in healthcare.	BTL2	Understanding
12	Show how machine learning is used in recommendation systems.	BTL2	Understanding
13	List the type of learning used in spam detection and state why it is suitable.	BTL1	Remembering
14	Give an example of a machine learning application in finance.	BTL2	Understanding
15	Illustrate one use case of machine learning in autonomous vehicles.	BTL2	Understanding
16	Define VC dimension in the context of hypothesis spaces.	BTL2	Understanding
17	Outline what a high VC dimension indicates about a model.	BTL2	Understanding
18	Interpret how VC dimension relates to model complexity.	BTL2	Understanding
19	List what PAC stands for in machine learning.	BTL2	Understanding
20	Name the term "probably approximately correct" in one sentence.	BTL1	Remembering
21	Define how PAC learning formalizes the learning process.	BTL1	Remembering
22	Identify a hypothesis space in machine learning.	BTL1	Remembering

23	State inductive bias with an example.	BTL1	Remembering
24	List why inductive bias is necessary for generalization.	BTL1	Remembering
PART-B			
1	Explain the importance of Eigen values and Eigen vectors in machine learning with an example application such as Principal Component Analysis (PCA)	(16)	BTL3 Applying
2	Discuss how linear algebra concepts such as matrix multiplication, inversion, and transposition are used in linear regression.	(16)	BTL3 Applying
3	Determine how vector spaces and subspaces are relevant in understanding the feature space of machine learning models.	(16)	BTL5 Evaluating
4	What is machine learning? Discuss its motivation, and differentiate it from traditional programming with suitable examples.	(16)	BTL5 Evaluating
5	Describe the different types of machine learning (supervised, unsupervised, reinforcement learning) and provide one real-world example of each.	(16)	BTL3 Applying
6	What are the major challenges in machine learning today, and how do they influence model development?	(16)	BTL3 Applying
7	Discuss at least three real-world applications of machine learning and explain how machine learning is used in each.	(16)	BTL3 Applying
8	Explain how machine learning is applied in healthcare, focusing on predictive diagnostics or personalized medicine.	(16)	BTL5 Evaluating
9	Illustrate the use of machine learning in natural language processing (NLP) with examples.	(16)	BTL3 Applying
10	Define VC dimension and explain its significance in understanding the capacity of a learning algorithm.	(16)	BTL4 Analyzing
11	Derive and interpret the VC dimension for simple hypothesis spaces such as intervals on the real line or linear classifiers in 2D.	(16)	BTL6 Creating
12	How does the VC dimension affect overfitting and underfitting in a learning model?	(16)	BTL3 Applying
13	(i) Explain the PAC learning framework. Discuss its assumptions and how it quantifies learnability. (ii) Derive the PAC learning bound and explain the significance of sample complexity in terms of error and confidence.	(8) (8)	BTL3 Applying
14	How does PAC learning help in understanding the feasibility of learning a concept?	(16)	BTL3 Applying
15	How can you apply the concepts of hypothesis spaces, inductive bias, and generalization to design a machine learning model?	(16)	BTL3 Applying
16	What is a hypothesis space? How does its size and complexity impact model generalization?	(16)	BTL3 Applying
17	Define inductive bias and explain its role in generalization. Use an example to show how different biases lead to different learned hypotheses.	(16)	BTL3 Applying

UNIT II- SUPERVISED LEARNING

Linear Regression Models: Least squares, single & multiple variables, Bayesian linear regression, gradient descent, Linear Classification Models: Discriminant function – Perceptron algorithm, Probabilistic discriminative model - Logistic regression, Probabilistic generative model – Naive Bayes, Maximum margin classifier – Support vector machine, Decision Tree, Random Forests.

PART-A			
Q.No	Questions	BT Level	Competence
1	Identify the objective of least squares in linear regression.	BTL1	Remembering
2	Summarize the assumption behind linear regression.	BTL2	Understanding
3	Define the cost function used in simple linear regression.	BTL2	Understanding
4	State what is meant by multicollinearity in multiple linear regression.	BTL1	Remembering
5	Mention one advantage and one limitation of Bayesian linear regression.	BTL2	Understanding
6	How does a decision tree classify data?	BTL1	Remembering
7	Compare the difference between batch and stochastic gradient descent.	BTL2	Understanding
8	Identify the role of the learning rate in gradient descent.	BTL1	Remembering
9	Illustrate how gradient descent minimizes the cost function in linear regression.	BTL2	Understanding
10	Define a discriminant function with an example.	BTL1	Remembering
11	Outline the basic idea behind the perceptron algorithm.	BTL2	Understanding
12	Name the type of data the perceptron algorithm can classify correctly.	BTL1	Remembering
13	Summarize one limitation of the perceptron learning algorithm.	BTL2	Understanding
14	State the decision boundary for logistic regression.	BTL2	Understanding
15	Compare why logistic regression is preferred over linear regression for classification.	BTL2	Understanding
16	Identify the main objective of the Maximum Margin Classifier.	BTL1	Remembering
17	Define the sigmoid function and its role in logistic regression.	BTL1	Remembering
18	Summarize the assumption made by the Naive Bayes classifier about features.	BTL2	Understanding
19	What is a random forest?.	BTL1	Remembering
20	Identify one use case of Naive Bayes in real-world applications.	BTL1	Remembering
21	Illustrate how Naive Bayes handles categorical data.	BTL2	Understanding
22	State what is meant by the margin and support vectors in SVM.	BTL1	Remembering
23	Define Gini index, and how is it used in decision trees?	BTL1	Remembering
24	Identify why SVM is called a maximum margin classifier.	BTL1	Remembering

	PART-B			
1	Discover the least squares solution for linear regression with a single variable. Explain the assumptions made and interpret the regression coefficients.	(16)	BTL3	Applying
2	Examine the least squares method to multiple linear regression. Derive the normal equation and discuss the geometric interpretation of the solution.	(16)	BTL4	Analyzing
3	Identify the concept of overfitting in linear regression. How can regularization (Ridge or Lasso) help in preventing it?	(16)	BTL3	Applying
4	Explain Bayesian linear regression in detail. Derive the posterior distribution and compare it with traditional linear regression.	(16)	BTL5	Evaluating
5	Explain the working of gradient descent in linear regression. Discuss batch, stochastic, and mini-batch variants, along with their pros and cons.	(16)	BTL5	Evaluating
6	Construct gradient descent for linear regression with one variable and show convergence using a sample dataset	(16)	BTL6	Creating
7	Evaluate the perceptron learning algorithm. Derive the update rule, explain its convergence for linearly separable data, and discuss its limitations.	(16)	BTL5	Evaluating
8	Identify the discriminant function used in linear classification. How does it help in decision boundary formation?	(16)	BTL3	Applying
9	Discover the logistic regression model from the sigmoid function. Explain how maximum likelihood estimation is used to find model parameters.	(8) (8)	BTL4	Analyzing
10	Compare linear regression and logistic regression. Discuss their differences in terms of output, loss function, and use cases.	(16)	BTL5	Evaluating
11	Analyze the Naive Bayes classification algorithm in detail. Derive the decision rule and discuss the effect of the independence assumption.	(8) (8)	BTL4	Analyzing
12	Apply Naive Bayes to a binary classification problem. Use a small dataset and show the full step-by-step computation.	(8) (8)	BTL3	Applying
13	Compose the concept of maximum margin classification. Derive the hard margin SVM optimization problem and interpret the role of support vectors.	(8) (8)	BTL6	Creating
14	Explain the kernel trick in SVM. Provide examples of common kernels and explain how they help in non-linear classification.	(16)	BTL5	Evaluating
15	Model how a decision tree is constructed? Discuss entropy, information gain, and the splitting criteria used in ID3 or CART.	(16)	BTL3	Applying
16	Estimate overfitting in decision trees. How does pruning help? Explain pre-pruning and post-pruning strategies.	(16)	BTL6	Creating
17	Analyze the working of Random Forests. Explain how ensemble learning and bagging help improve accuracy and reduce overfitting in classification problems.	(16)	BTL4	Analyzing

UNIT III- ENSEMBLE TECHNIQUES AND UNSUPERVISED LEARNING				
Combining multiple learners: Model combination schemes, Voting, Ensemble Learning - bagging, boosting, stacking, Unsupervised learning: K-means, Instance Based Learning: KNN, Gaussian mixture models and Expectation maximization.				
	PART-A			
Q.No	Questions		BT Level	Competence
1	What is ensemble learning?		BTL1	Remembering
2	List one benefit of using ensemble methods.		BTL1	Remembering
3	Compare hard voting and soft voting.		BTL2	Understanding
4	State the definition of bagging in the context of ensemble learning.		BTL1	Remembering
5	Summarize the concept of bootstrapping in bagging.		BTL2	Understanding
6	Mention one algorithm that uses bagging.		BTL2	Understanding
7	What is boosting, and how does it improve model performance?		BTL1	Remembering
8	Differentiate boosting from bagging.		BTL2	Understanding
9	What are two popular boosting algorithms?		BTL1	Remembering
10	Outline the concept of stacking in ensemble learning.		BTL2	Understanding
11	What is the role of the meta-learner in stacking?		BTL2	Understanding
12	Why does ensemble learning often improve model performance?		BTL2	Understanding
13	What is the purpose of model combination?		BTL1	Remembering
14	List two common strategies for model combination.		BTL1	Remembering
15	Summarize the role of weighting in model averaging.		BTL2	Understanding
16	Define majority voting in ensemble methods.		BTL2	Understanding
17	Why does model diversity affect ensemble performance?		BTL2	Understanding
18	What is unsupervised learning, and provide one example.		BTL2	Understanding
19	What is the main objective of clustering?		BTL2	Understanding
20	What is the importance of instance based learning in machine learning and give an example algorithm.		BTL2	Understanding
21	State the objective function of K-means clustering.		BTL1	Remembering
22	What is the role of centroids in K-means?		BTL1	Remembering
23	Mention one limitation of K-means clustering.		BTL2	Understanding
24	What is the Expectation-Maximization (EM) algorithm and how is it used in GMMs?		BTL1	Remembering
	PART-B			
1	Identify the concept of ensemble learning. Discuss its advantages with examples of where it outperforms individual models.	(16)	BTL3	Applying

2	Analyze bagging technique in ensemble learning. Explain its advantages and limitations.	(16)	BTL4	Analyzing
3	What is boosting? Explain the bootstrapping process, and how it helps in reducing variance using an example like Random Forests.	(16)	BTL3	Applying
4	Compare and contrast bagging and boosting. Discuss their objectives, mechanisms, and impact on bias and variance.	(16)	BTL4	Analyzing
5	Assess the stacking ensemble technique. Explain the role of base learners and meta-learners with an appropriate example.	(16)	BTL4	Analyzing
6	What are voting-based ensemble methods? Differentiate between hard voting and soft voting with examples.	(16)	BTL5	Evaluating
7	Explain model combination schemes. Describe model averaging, majority voting, and weighted combination strategies with examples.	(16)	BTL5	Evaluating
8	Apply the importance of diversity in ensemble methods. How does it contribute to better generalization?	(16)	BTL3	Applying
9	Define unsupervised learning. Explain how it differs from supervised learning and provide examples of tasks and algorithms.	(16)	BTL5	Evaluating
10	Describe the K-means clustering algorithm. Explain the steps with mathematical formulation and give a real-world example.	(16)	BTL3	Applying
11	Determine the limitations of K-means clustering? How can they be addressed (e.g., by using K-medoids or DBSCAN)?	(16)	BTL5	Evaluating
12	Explain the working of the K-Nearest Neighbors (KNN) algorithm. How is it used for both classification and regression?	(8) (8)	BTL5	Evaluating
13	Compare KNN and K-means. Explain their differences in learning type, working mechanism, and use cases.	(16)	BTL4	Analyzing
14	Compile the working of Gaussian Mixture Model (GMM). How does it differ from K-means in terms of cluster assignment?	(16)	BTL6	Creating
15	Make use of the Expectation-Maximization (EM) algorithm. Derive its steps and show how it is used in training GMMs.	(16)	BTL3	Applying
16	Compose the application of EM in unsupervised learning. Provide an example of clustering using EM and GMM.	(16)	BTL6	Creating
17	Analyze elbow method and silhouette analysis for determining the optimal number of clusters in K means clustering.	(16)	BTL4	Analyzing

UNIT IV- NEURAL NETWORKS

Multilayer perceptron, activation functions, network training – gradient descent optimization – stochastic gradient descent, error backpropagation, from shallow networks to deep networks –Unit saturation (aka the vanishing gradient problem) – ReLU, hyperparameter tuning, batch normalization, regularization, dropout.

PART-A			
Q.No	Questions	BT Level	Competence
1	Define a multilayer perceptron (MLP).	BTL1	Remembering
2	List how MLP differs from a single-layer perceptron.	BTL1	Remembering
3	What is the role of hidden layers in an MLP?	BTL1	Remembering
4	Interpret the purpose of an activation function in neural networks.	BTL2	Understanding
5	Write the mathematical form of the sigmoid function.	BTL2	Understanding
6	State one limitation of the sigmoid activation function.	BTL2	Understanding
7	Compare the advantage of using ReLU over sigmoid or tanh.	BTL2	Understanding
8	Define the ReLU activation function.	BTL2	Understanding
9	What is the objective of gradient descent in training neural networks?	BTL2	Understanding
10	Compare batch and stochastic gradient descent.	BTL2	Understanding
11	State one advantage and one disadvantage of stochastic gradient descent.	BTL1	Remembering
12	Demonstrate what the learning rate controls in gradient descent.	BTL2	Understanding
13	Summarize the concept of backpropagation.	BTL2	Understanding
14	Why is the chain rule important in backpropagation?	BTL2	Understanding
15	What kind of loss function is typically used in classification tasks?	BTL1	Remembering
16	Summarize the vanishing gradient problem.	BTL2	Understanding
17	Outline how ReLU helps mitigate the vanishing gradient problem.	BTL2	Understanding
18	Classify the two common regularization techniques.	BTL2	Understanding
19	Give two examples of hyperparameters in neural network training.	BTL2	Understanding
20	Outline the role of cross-validation in hyperparameter tuning.	BTL2	Understanding
21	What is batch normalization, and why is it used?	BTL2	Understanding
22	Mention one benefit of using batch normalization during training.	BTL1	Remembering
23	State the purpose of regularization in neural networks.	BTL1	Remembering
24	Illustrate what dropout does in a neural network during training.	BTL1	Remembering

	PART-B			
1	Elaborate the architecture of a multilayer perceptron (MLP). Describe how inputs are transformed through layers using weights and activation functions.	(16)	BTL6	Creating
2	Compare single-layer and multilayer perceptrons. Explain why multilayer perceptrons are more powerful in modeling complex data.	(16)	BTL5	Evaluating
3	Describe different activation functions such as sigmoid, tanh, and ReLU. Compare their properties and impact on training deep networks.	(16)	BTL4	Analyzing
4	Explain the limitations of sigmoid and tanh functions. How does the ReLU activation function help overcome these limitations?	(16)	BTL5	Evaluating
5	Formulate the gradient descent algorithm and its role in training neural networks. Include mathematical derivation for weight updates.	(16)	BTL6	Creating
6	Explain stochastic gradient descent (SGD). Compare it with batch and mini-batch gradient descent in terms of convergence and computation.	(16)	BTL5	Evaluating
7	Analyze the importance of the learning rate in gradient descent optimization. What happens when it is too small or too large?	(16)	BTL4	Analyzing
8	Explain the backpropagation algorithm in detail. Derive the equations used to compute the gradients of the loss function.	(16)	BTL5	Evaluating
9	Illustrate backpropagation with an example for a simple neural network with one hidden layer.	(8) (8)	BTL3	Applying
10	What is the vanishing gradient problem? Identify why it occurs in deep networks and its impact on learning.	(8) (8)	BTL3	Applying
11	Choose the strategies to address the vanishing gradient problem. Include the role of ReLU and proper weight initialization.	(16)	BTL3	Applying
12	What are hyperparameters in neural networks? Discuss strategies for tuning them and the effect on model performance.	(16)	BTL6	Creating
13	Compare grid search and random search for hyperparameter tuning. Mention their pros and cons.	(16)	BTL4	Analyzing
14	Explain batch normalization. Describe how it works and its benefits during training of deep neural networks.	(16)	BTL5	Evaluating
15	Discuss how batch normalization affects learning rate, convergence speed, and regularization.	(16)	BTL5	Evaluating
16	Discover the concept of regularization in neural networks. Describe L1 and L2 regularization and how they prevent overfitting.	(16)	BTL4	Analyzing
17	Make use of dropout as a regularization technique. How does it work, and why is it effective in deep learning?	(16)	BTL3	Applying

UNIT V- : DESIGN AND ANALYSIS OF MACHINE LEARNING EXPERIMENTS

Guidelines for machine learning experiments, Cross Validation (CV) and resampling – K-fold CV, bootstrapping, measuring classifier performance, assessing a single classification algorithm and comparing two classification algorithms – t test, McNemar’s test, K-fold CV paired t test.

PART-A			
Q.No	Questions	BT Level	Competence
1	What is the purpose of conducting machine learning experiments?	BTL1	Remembering
2	Mention two good practices when designing a machine learning experiment.	BTL2	Understanding
3	Classify and name two common resampling techniques used in cross-validation.	BTL2	Understanding
4	Why is it important to separate training and testing datasets?	BTL2	Understanding
5	Define model evaluation in the context of machine learning.	BTL2	Understanding
6	What is cross-validation?	BTL2	Understanding
7	Summarize the main idea behind K-fold cross-validation.	BTL2	Understanding
8	Differentiate stratified K-fold cross-validation from regular K-fold.	BTL2	Understanding
9	Define leave-one-out cross-validation (LOOCV).	BTL1	Remembering
10	State the purpose of using cross-validation.	BTL2	Understanding
11	Illustrate resampling in machine learning experiments.	BTL2	Understanding
12	What is bootstrapping in the context of model evaluation?	BTL2	Understanding
13	List how bootstrapping differs from K-fold cross-validation.	BTL1	Remembering
14	Summarize the key advantage of bootstrapping over traditional train-test splits.	BTL2	Understanding
15	What is a confusion matrix?	BTL2	Understanding
16	Define precision and recall.	BTL1	Remembering
17	State what the F1-score represents.	BTL1	Remembering
18	What does ROC stand for in machine learning?	BTL2	Understanding
19	Interpret in what scenario is McNemar’s test most appropriate.	BTL2	Understanding
20	State the purpose of assessing a single classification algorithm’s performance.	BTL2	Understanding
21	List why repeated evaluation is important for a classification algorithm.	BTL1	Remembering
22	What is the purpose of using a t-test in comparing classifiers?	BTL1	Remembering
23	Interpret the use of McNemar’s test in machine learning.	BTL2	Understanding
24	Summarize how K-fold CV paired t-test differs from a standard t-test in classifier comparison.	BTL1	Remembering

	PART-B			
1	Explain the importance of designing proper machine learning experiments. Discuss common pitfalls and best practices.	(16)	BTL3	Applying
2	Describe K-fold cross-validation in detail. Explain how it helps in estimating the generalization error of a model.	(16)	BTL3	Applying
3	What is bootstrapping? Explain the bootstrapping procedure for model evaluation and compare it with cross-validation.	(16)	BTL3	Applying
4	Explain the difference between hold-out validation and cross-validation. Discuss the advantages and disadvantages of each.	(16)	BTL3	Applying
5	Define the confusion matrix and explain its components. How is it used to measure classifier performance?	(16)	BTL4	Analyzing
6	Discuss various metrics used to measure classifier performance, such as accuracy, precision, recall, F1-score, and ROC-AUC and explain their significance.	(16)	BTL3	Applying
7	Examine the ROC curve and AUC metric. How do they help in evaluating classifiers?	(16)	BTL4	Analyzing
8	Discuss the need for statistical testing when comparing two classification algorithms.	(16)	BTL4	Analyzing
9	Explain the paired t-test used for comparing two classifiers. Describe its assumptions, procedure, and interpretation.	(16)	BTL5	Evaluating
10	What is McNemar's test? Explain how it is used to compare two classification models.	(16)	BTL4	Analyzing
11	Describe the K-fold cross-validation paired t-test. How does it differ from the standard paired t-test?	(16)	BTL5	Evaluating
12	Estimate how to assess the performance of a single classification algorithm using cross-validation.	(16)	BTL6	Creating
13	Formulate the importance of resampling techniques in machine learning experiments. Give examples where they are useful.	(16)	BTL6	Creating
14	Compare and contrast different resampling methods used in machine learning evaluation.	(16)	BTL5	Evaluating
15	What is overfitting in machine learning experiments? How can cross-validation help detect overfitting?	(16)	BTL5	Evaluating
16	Elaborate the concept of bias-variance trade-off in the context of model evaluation.	(16)	BTL6	Creating
17	Develop an experiment to compare two classification algorithms on a given dataset. Outline the steps, evaluation metrics, and statistical tests you would use.	(16)	BTL4	Analyzing

Faculty In-Charge

Verified by

HOD