

Short Questions

1. What is supervised learning?
2. Define linear regression.
3. Explain the least squares method.
4. What is multiple regression?
5. How does multiple outputs regression work?
6. Describe subset selection in regression analysis.
7. Explain ridge regression.
8. Define Lasso regression.
9. What is Linear Discriminant Analysis (LDA)?
10. Describe logistic regression.
11. What is the Perceptron learning algorithm?
12. How does regularization affect regression models?
13. Compare Lasso and ridge regression.
14. When would you use LDA over logistic regression?
15. How does multiple regression handle collinearity?
16. What are the assumptions of linear regression?
17. How do you interpret regression coefficients?
18. What is the purpose of subset selection?
19. How does the Perceptron algorithm update weights?
20. Discuss the importance of feature scaling in linear models.
21. Explain the concept of multicollinearity.
22. How can Lasso regression be used for feature selection?
23. What is the difference between linear and logistic regression?
24. Describe the process of validating a linear regression model.
25. What criteria can be used for subset selection?
26. What is the difference between simple and multiple linear regression?
27. How do you assess the fit of a linear regression model?
28. What is multicollinearity, and why is it a problem?
29. How does ridge regression address multicollinearity?
30. What is the significance of the alpha parameter in Lasso and ridge regression?

31. How can you interpret logistic regression coefficients?
32. In what scenarios is logistic regression preferred over linear regression?
33. Describe the steps in performing Linear Discriminant Analysis.
34. What are the limitations of the Perceptron learning algorithm?
35. How do you choose between Lasso and ridge regression for a given dataset?
36. What is the bias-variance tradeoff?
37. Explain model complexity in supervised learning.
38. Define the concept of overfitting.
39. How does cross-validation work?
40. What is bootstrapping in model assessment?
41. Describe the Bayesian Information Criterion (BIC).
42. How do you estimate in-sample prediction error?
43. What is the effective number of parameters?
44. Explain optimism in the training error rate.
45. How is the bias-variance tradeoff related to model performance?
46. Discuss the role of cross-validation in model selection.
47. What are the advantages of bootstrapping methods?
48. How does BIC compare to AIC in model selection?
49. What is conditional or expected test error?
50. How can model complexity influence bias and variance?
51. Describe the process of selecting the optimal model complexity.
52. Explain the concept of regularization in reducing overfitting.
53. How does bias affect machine learning models?
54. What strategies can be used to reduce variance in predictions?
55. How do you interpret the results of cross-validation?
56. What is the purpose of the training error rate?
57. How can the effectiveness of a model's in-sample prediction error be evaluated?
58. Describe the concept of the effective number of parameters in a model.
59. How does the Bayesian approach influence model selection?
60. What are the benefits and drawbacks of using cross-validation?
61. Explain how bootstrap methods are applied in model assessment.

62. What is meant by conditional or expected test error in model evaluation?
63. How does one balance bias and variance in a machine learning model?
64. Why might one prefer Bayesian Information Criterion (BIC) over other model selection criteria?
65. What are the implications of high variance in model predictions?
66. How do Generalized Additive Models (GAM) differ from linear models?
67. What are the key considerations in building regression trees?
68. Describe the methodology behind classification trees.
69. How does the AdaBoost algorithm function?
70. What is meant by exponential loss in the context of boosting?
71. How do regression trees handle continuous and categorical data?
72. What are the advantages of using boosting methods over single models?
73. How can overfitting be controlled in boosted models?
74. In what ways can trees be pruned to improve model performance?
75. What are the considerations in choosing the depth of a decision tree?
76. What are Generalized Additive Models (GAM)?
77. Describe regression trees.
78. Explain classification trees.
79. What is boosting in machine learning?
80. Define AdaBoost.
81. How does boosting improve model accuracy?
82. What is the exponential loss function in boosting?
83. Discuss the advantages of using additive models.
84. How do trees handle categorical variables?
85. Compare boosting to bagging.
86. How does AdaBoost select weak learners?
87. What role does the loss function play in boosting?
88. Describe the process of building a regression tree.
89. How can additive models be applied to classification problems?
90. What is tree pruning and why is it important?
91. Explain the concept of feature importance in tree models.
92. How do boosting algorithms reduce bias and variance?

93. Describe the splitting criteria in decision trees.
94. How does GAM handle non-linear relationships?
95. What are the limitations of tree-based models?
96. Discuss the importance of tree depth in model performance.
97. How do ensemble methods improve prediction accuracy?
98. What are the key parameters in boosting algorithms?
99. Explain how decision trees can be used for regression.
100. How do you interpret the results from a boosted model?
101. What are the key differences between supervised and unsupervised learning?
102. How do you handle missing data in regression analysis?
103. Discuss the importance of data preprocessing in machine learning.
104. What metrics can be used to evaluate a regression model?
105. What metrics can be used to evaluate a classification model?
106. How can machine learning models be deployed in production?
107. Explain the concept of ensemble learning.
108. How do you select the right algorithm for a machine learning problem?
109. What are the challenges in applying linear methods to real-world data?
110. Discuss the role of hyperparameter tuning in model performance.
111. How can model interpretability be improved?
112. What are the ethical considerations in machine learning?
113. How do you ensure the robustness of machine learning models?
114. What are the common pitfalls in model assessment and selection?
115. How does feature engineering affect model performance?
116. How is model complexity related to the training and testing error?
117. What strategies are effective for handling imbalanced datasets in classification?
118. How do ensemble methods leverage multiple models to improve accuracy?
119. Describe the trade-offs between interpretability and accuracy in model selection.
120. How does feature selection impact the performance of machine learning models?
121. What role does data scaling play in the performance of linear models?
122. How can cross-validation be used to select hyperparameters?
123. Discuss the importance of domain knowledge in feature engineering.

124. What are the common causes of overfitting, and how can it be prevented?
125. How do you validate the assumptions of a linear regression model?