MACHINE LEARNING PROGRAMMING

1. What is the difference between supervised and unsupervised learning?

Supervised learning is a type of machine learning in which the model is trained on labeled data, where the output variable is known. The model learns to predict the output variable based on the input data and the corresponding labels. The goal of supervised learning is to learn a mapping between the input data and the output variable so that the model can accurately predict the output for new, unseen data.

Unsupervised learning, on the other hand, is a type of machine learning in which the model is trained on unlabeled data, where the output variable is unknown. The model learns to identify patterns and structure in the data without any prior knowledge of the output variable. The goal of unsupervised learning is to discover underlying patterns and relationships in the data that can be used for further analysis or downstream tasks.

In summary, the main difference between supervised and unsupervised learning is that supervised learning requires labeled data, while unsupervised learning does not. Supervised learning is used for prediction tasks, while unsupervised learning is used for data exploration and pattern discovery.

2.Can you explain the bias-variance tradeoff in machine learning?

The bias-variance tradeoff is a fundamental concept in machine learning that describes the relationship between a model's bias and variance and its ability to generalize to new data.



Bias refers to the error that is introduced by approximating a real-world problem with a simplified model. High bias means that the model is too simple, and it is not able to capture the underlying patterns in the data. This can result in underfitting, where the model is not able to fit the training data well or perform well on new data.

Variance, on the other hand, refers to the error that is introduced by the model's sensitivity to fluctuations in the training data. High variance means that the model is too complex, and it is overfitting the training data, meaning that it is fitting the training data too well and is not able to generalize well to new data.

In machine learning, the goal is to find the right balance between bias and variance to create a model that can generalize well to new, unseen data. This is known as the bias-variance tradeoff. A model with high bias and low variance may underfit the data, while a model with low bias and high variance may overfit the data. The ideal model has low bias and low variance, meaning it fits the training data well and generalizes well to new data.

In summary, the bias-variance tradeoff is a balancing act between creating a model that is complex enough to capture the underlying patterns in the data but not so complex that it overfits the data and is unable to generalize well to new, unseen data.

Regularization is a technique in machine learning that is used to prevent overfitting of models. Overfitting occurs when a model fits the training data too well and captures the noise in the data, rather than the underlying patterns. This can lead to poor performance when the model is used to make predictions on new data.

Regularization works by adding a penalty term to the loss function that the model is trying to minimize during training. The penalty term encourages the model to have smaller weights or coefficients, which results in a simpler model that is less likely to overfit the data.

There are different types of regularization techniques used in machine learning, including L1 regularization (also known as Lasso), L2 regularization (also known as Ridge), and Elastic Net regularization. L1 regularization adds a penalty term to the loss function that is proportional to the absolute value of the weights, while L2 regularization adds a penalty term that is proportional to the square of the weights. Elastic Net regularization is a combination of L1 and L2 regularization.

Regularization is an important technique in machine learning because it helps to improve the generalization performance of models by reducing overfitting. However, it is also important to choose the right regularization technique and regularization strength for the specific problem at hand, as using too much regularization can result in underfitting and poor performance on both training and test data.

4. What is cross-validation in machine learning?

Cross-validation is a technique in machine learning that is used to evaluate the performance of a model and to tune its hyperparameters. The goal of cross-validation is to estimate the performance of a model on new, unseen data by using the available data for both training and testing.

Cross-validation works by dividing the available data into multiple sets or folds, typically using a k-fold approach. In k-fold cross-validation, the data is split into k equal-sized subsets or folds. The model is trained on k-1 folds and tested on the remaining fold. This process is repeated k times, with each fold serving as the test set once. The performance of the model is then averaged over the k runs to obtain an estimate of its generalization performance.

Cross-validation is useful for selecting the best hyperparameters for a model, as it provides an estimate of how well the model will perform on new, unseen data. By evaluating the performance of the model on multiple subsets of the data, cross-validation helps to reduce the variance of the estimate of the model's performance and provides a more reliable estimate of its generalization performance.

In summary, cross-validation is a technique in machine learning that is used to estimate the generalization performance of a model by using the available data for both training and testing. It is a useful technique for selecting the best hyperparameters for a model and for evaluating the performance of different models.

5. What is a confusion matrix in machine learning?

A confusion matrix is a table used in machine learning to evaluate the performance of a classification model. The matrix provides a summary of the model's predictions on the test data, showing how many true positives, true negatives, false positives, and false negatives were produced by the model.

The confusion matrix is typically a square matrix of size n x n, where n is the number of classes in the classification problem. The rows of the matrix correspond to the true class labels, while the columns correspond to the predicted class labels. Each entry in the matrix represents the number of instances that belong to a particular true class and were predicted to belong to a particular predicted class.

The four entries in the confusion matrix are defined as follows:

- True positives (TP): the number of instances that belong to a particular true class and were correctly predicted to belong to that class.
- True negatives (TN): the number of instances that do not belong to a particular true class and were correctly predicted not to belong to that class.
- False positives (FP): the number of instances that do not belong to a particular true class but were incorrectly predicted to belong to that class.
- False negatives (FN): the number of instances that belong to a particular true class but were incorrectly predicted not to belong to that class.

The confusion matrix is useful for evaluating the performance of a classification model, as it provides a more detailed view of the model's performance than a single accuracy score. From the confusion matrix, various metrics can be derived, such as precision, recall, and F1-score, which can provide more insights into the model's performance for specific classes.

In summary, a confusion matrix is a table used in machine learning to evaluate the performance of a classification model by showing the number of true positives, true negatives, false positives, and false negatives produced by the model.

6.Can you explain the difference between precision and recall?

Precision is a measure of how many of the predicted positive cases are actually positive. Recall is a measure of how many of the actual positive cases are correctly predicted as positive. Precision is often used when the cost of a false positive is high, while recall is used when the cost of a false negative is high.

7. What is gradient descent in machine learning?

Gradient descent is an optimization algorithm used to minimize the error of a machine learning model. It works by iteratively adjusting the model's parameters in the direction of the negative gradient of the objective function.

8. What is overfitting in machine learning?

Overfitting is a common problem in machine learning where the model is too complex and fits the training data too closely. This can lead to poor generalization performance on new data

9. What is the curse of dimensionality in machine learning?

The curse of dimensionality refers to the difficulty of learning in high-dimensional spaces. As the number of features or dimensions increases, the amount of data required to cover the space increases exponentially, making it more difficult to learn a model that generalizes well to new data.

10.What is deep learning and how does it differ from traditional machine learning?

Deep learning is a type of machine learning that uses neural networks with multiple layers to learn hierarchical representations of the data. Deep learning differs from traditional machine learning in that it is able to automatically learn features from the data, rather than relying on hand-engineered features.

11.Can you explain the difference between a neural network and a decision tree?

A neural network is a type of machine learning model that consists of interconnected nodes or neurons that process and transmit information. A decision tree, on the other hand, is a type of machine learning model that represents decisions and their possible consequences as a tree-like structure.

12. What is transfer learning in deep learning?

Transfer learning is a technique in deep learning where a pre-trained neural network is used as a starting point for a new model. This can significantly reduce the amount of training data required for the new model and improve its performance.

13. What is backpropagation in neural networks?

Backpropagation is an algorithm used to train neural networks by adjusting the weights of the connections between neurons in the network. It works by propagating the error back through the network and using it to update the weights in a way that reduces the error.

14. What is batch normalization in deep learning?

Batch normalization is a technique used to improve the training of deep neural networks by normalizing the input data to each layer. This can help to reduce the internal covariate shift and stabilize the training process

15. What is a generative model in machine learning?

A generative model is a type of machine learning model that learns to generate new data that is similar to the training data. This can be used for tasks such as data augmentation, anomaly detection, and image synthesis.

16. What is reinforcement learning?

Reinforcement learning is a type of machine learning where an agent learns to make decisions based on feedback from its environment. The agent receives rewards or punishments for its actions and learns to optimize its behavior to maximize its rewards.

17.Can you explain the difference between overfitting and underfitting?

Overfitting occurs when a machine learning model is too complex and fits the training data too closely, resulting in poor generalization performance on new data. Underfitting occurs when the model is too simple and fails to capture the underlying patterns in the data, resulting in poor performance on both the training and test data.

18. What is the curse of big data in machine learning?

The curse of big data refers to the challenges of working with large datasets, such as increased computational requirements, difficulties in data storage and management, and the risk of overfitting due to the large number of features.

19.Can you explain the difference between a support vector machine (SVM) and a logistic regression model?

A support vector machine (SVM) is a type of machine learning model that is used for classification and regression tasks. It works by finding a hyperplane that maximally separates the data into different classes. A logistic regression model is a type of linear model that is used for classification tasks. It works by modeling the probability of each class using a logistic function.

20. What is cross-validation in machine learning?

Cross-validation is a technique used to evaluate the performance of a machine learning model by dividing the data into multiple subsets and training the model on different subsets while evaluating its performance on the remaining subset.

21. What is the bias-variance tradeoff in machine learning?

The bias-variance tradeoff is a fundamental concept in machine learning that refers to the tradeoff between a model's ability to fit the training data (low bias) and its ability to generalize to new data (low variance).

22. What is the difference between a regression model and a classification model?

A regression model is used to predict continuous values, while a classification model is used to predict discrete values or classes.

23. What is unsupervised learning in machine learning?

Unsupervised learning is a type of machine learning where the model learns to identify patterns in the data without any labeled examples or guidance from a human expert.

24. What is the difference between a parametric model and a non-parametric model?

A parametric model makes assumptions about the underlying distribution of the data and has a fixed number of parameters, while a non-parametric model makes no assumptions about the underlying distribution and can have an unlimited number of parameters.

25. What is regularization in machine learning?

Regularization is a technique used to prevent overfitting in machine learning by adding a penalty term to the loss function that encourages the model to have simpler weights.

26. What is the difference between deep learning and machine learning?

Deep learning is a subset of machine learning that involves training neural networks with multiple layers to learn hierarchical representations of the data.

27. What is the difference between supervised learning and unsupervised learning?

Supervised learning involves learning from labeled examples, while unsupervised learning involves learning from unlabeled data.

28.What is the difference between a feedforward neural network and a recurrent neural network?

A feedforward neural network processes data in a single direction, while a recurrent neural network has loops that allow it to process sequences of data.

29. What is the difference between a convolutional neural network and a recurrent neural network?

A convolutional neural network is designed for image and signal processing tasks, while a recurrent neural network is designed for sequence modeling tasks.

30. What is the difference between deep reinforcement learning and supervised learning?

Deep reinforcement learning involves learning from feedback in the form of rewards or punishments, while supervised learning involves learning from labeled examples.

31. What is the difference between a decision tree and a random forest?

A decision tree is a single tree-like structure that represents decisions and their consequences, while a random forest is an ensemble of decision trees that combine their predictions.

32. What is transfer learning and how does it work?

Transfer learning is a technique that involves using a pre-trained model as a starting point for a new model. The pre-trained model is fine-tuned on a new task, which can reduce the amount of training data required and improve the performance of the new model.

33. What is a confusion matrix in machine learning?

A confusion matrix is a table that summarizes the performance of a classification model by showing the number of true positives, true negatives, false positives, and false negatives.

34. What is gradient descent in machine learning?

Gradient descent is an optimization algorithm used to minimize the loss function of a machine learning model by iteratively adjusting the parameters in the direction of the negative gradient.

35. What is a hyperparameter in machine learning?

A hyperparameter is a parameter of a machine learning model that is set before training and determines the behavior of the model, such as the learning rate, regularization strength, or number of hidden layers.

36.What is the difference between overfitting and underfitting in machine learning?

Overfitting occurs when a model is too complex and fits the training data too well, leading to poor generalization to new data. Underfitting occurs when a model is too simple and cannot capture the patterns in the data, leading to poor performance on both the training and test data.

37. What is data augmentation in machine learning?

Data augmentation is a technique used to increase the amount and diversity of training data by applying transformations to the existing data, such as flipping, rotating, scaling, or adding noise.

38. What is the curse of dimensionality in machine learning?

The curse of dimensionality refers to the fact that as the number of features or dimensions of the data increases, the amount of data required to maintain the same level of accuracy increases exponentially.

39. What is ensemble learning in machine learning?

Ensemble learning is a technique that involves combining multiple models to improve the accuracy and robustness of predictions. Common ensemble methods include bagging, boosting, and stacking.

40.What is the difference between overfitting and underfitting in machine learning?

Overfitting occurs when a model is too complex and fits the training data too well, leading to poor generalization to new data. Underfitting occurs when a model is too simple and cannot capture the patterns in the data, leading to poor performance on both the training and test data.

41. What is data augmentation in machine learning?

Data augmentation is a technique used to increase the amount and diversity of training data by applying transformations to the existing data, such as flipping, rotating, scaling, or adding noise.

42. What is the curse of dimensionality in machine learning?

The curse of dimensionality refers to the fact that as the number of features or dimensions of the data increases, the amount of data required to maintain the same level of accuracy increases exponentially.

43. What is ensemble learning in machine learning?

Ensemble learning is a technique that involves combining multiple models to improve the accuracy and robustness of predictions. Common ensemble methods include bagging, boosting, and stacking.

44. What is the difference between precision and recall in machine learning?

Precision is the proportion of true positive predictions out of all positive predictions, while recall is the proportion of true positive predictions out of all actual positive cases.

45. What is the difference between online learning and batch learning in machine learning?

Online learning involves updating the model with each new data point, while batch learning involves updating the model on batches of data.

46. What is the difference between L1 and L2 regularization?

L1 regularization adds a penalty term proportional to the absolute value of the weights, while L2 regularization adds a penalty term proportional to the square of the weights.

47. What is the difference between an autoencoder and a generative adversarial network (GAN)?

An autoencoder is a type of neural network that learns to reconstruct the input data, while a GAN is a type of generative model that learns to generate new data that is similar to the input data.

48. What is transfer learning and why is it useful in machine learning?

Transfer learning is the process of leveraging knowledge learned from one task or domain to improve performance on a different but related task or domain. It is useful in machine learning because it allows models to learn more efficiently from smaller datasets and can improve performance on tasks with limited labeled data.

49.What is the difference between a support vector machine (SVM) and a neural network?

An SVM is a type of linear classifier that maximizes the margin between classes, while a neural network is a nonlinear classifier that learns complex representations of the data.

50. What is the difference between a clustering algorithm and a classification algorithm?

A clustering algorithm groups similar data points together based on their similarity, while a classification algorithm assigns data points to predefined categories or classes.

51. What is the difference between a feedforward neural network and a recurrent neural network?

A feedforward neural network processes data in a single direction, while a recurrent neural network has loops that allow it to process sequences of data.

52. What is the difference between a decision tree and a random forest?

A decision tree is a single tree-like structure that represents decisions and their consequences, while a random forest is an ensemble of decision trees that combine their predictions.

53. What is the difference between batch normalization and layer normalization in deep learning?

Batch normalization normalizes the input to each layer across the entire batch, while layer normalization normalizes the input to each layer across the feature dimension.

54. What is the difference between supervised and unsupervised learning in machine learning?

Supervised learning involves training a model on labeled data, where the correct output is known for each input. Unsupervised learning involves training a model on unlabeled data, where the goal is to discover patterns or structure in the data.

55. What is the difference between a convolutional neural network (CNN) and a recurrent neural network (RNN)?

A CNN is typically used for image and video processing, while an RNN is typically used for natural language processing and sequence modeling.

56.What is the difference between a local and a global optimum in optimization problems?

A local optimum is a solution that is the best in a particular neighborhood, while a global optimum is the best solution overall.

57. What is the difference between k-means and hierarchical clustering?

K-means is a centroid-based clustering algorithm that partitions data into k clusters based on the distance between each data point and its nearest cluster center. Hierarchical clustering builds a hierarchy of clusters by merging or splitting clusters based on their similarity.

58. What is the difference between a hyperparameter and a parameter in machine learning?

A hyperparameter is a configuration setting for the model that is set before training and cannot be learned from the data. A parameter is a learned value of the model that is updated during training.

59. What is the difference between a neural network and a deep neural network?

A neural network is a single layer or a few layers of interconnected nodes, while a deep neural network has many layers and can learn complex representations of the data.

60. What is the difference between reinforcement learning and supervised learning?

Reinforcement learning involves learning to take actions to maximize a reward signal, while supervised learning involves learning to predict an output given an input.

61. What is the difference between a decision tree and a naive Bayes classifier?

A decision tree is a flowchart-like structure that represents decisions and their consequences, while a naive Bayes classifier is a probabilistic model that uses Bayes' theorem to make predictions.

62. What is the difference between a softmax and a sigmoid function in neural networks?

A softmax function is used for multi-class classification problems and outputs probabilities that sum to one, while a sigmoid function is used for binary classification problems and outputs probabilities between 0 and 1.

63. What is the difference between batch gradient descent and stochastic gradient descent?

Batch gradient descent updates the model parameters based on the average gradient of the loss function over the entire training set, while stochastic gradient descent updates the model parameters based on the gradient of the loss function for each individual data point in the training set.

64. What is overfitting in machine learning and how can it be prevented?

Overfitting occurs when a model is too complex and learns the noise in the training data, rather than the underlying patterns. It can be prevented by using techniques such as regularization, early stopping, and increasing the size of the training set.

65. What is the difference between a decision boundary and a hyperplane in machine learning?

A decision boundary is a boundary that separates the classes in a classification problem, while a hyperplane is a higher-dimensional version of a decision boundary that separates the classes in a linearly separable problem.

66.What is the difference between a deep belief network and a convolutional neural network?

A deep belief network is an unsupervised learning model that learns to represent the input data in a hierarchical manner, while a convolutional neural network is a supervised learning model that learns to extract features from image data.

67. What is the difference between a one-vs-one and a one-vs-all approach in multi-class classification?

A one-vs-one approach trains a binary classifier for each pair of classes and combines their predictions to make a multi-class prediction, while a one-vs-all approach trains a binary classifier for each class and uses their predictions to make a multi-class prediction.

68. What is the difference between a neural network and a decision tree?

A neural network is a series of connected nodes that learn from data, while a decision tree is a tree-like structure that models decisions and their consequences.

69. What is the difference between feature selection and feature extraction in machine learning?

Feature selection involves selecting a subset of the available features based on their importance, while feature extraction involves transforming the input features into a new space using a mathematical function.

70. What is the difference between a kernel method and a non-kernel method in machine learning?

A kernel method uses a kernel function to transform the input data into a higher-dimensional space, while a non-kernel method does not use a kernel function and operates in the original feature space.

71. What is the difference between a Gaussian mixture model and a k-means clustering algorithm?

A Gaussian mixture model is a probabilistic model that represents data as a mixture of Gaussian distributions, while a k-means clustering algorithm is a non-probabilistic clustering algorithm that partitions data into k clusters based on their distance from the cluster centers.

72. What is the difference between a validation set and a test set in machine learning?

A validation set is used to tune the hyperparameters of the model, while a test set is used to evaluate the performance of the final model on unseen data.

73. What is the difference between a decision tree and a random forest?

A decision tree is a single tree-like structure that represents decisions and their consequences, while a random forest is an ensemble of decision trees that combine their predictions.

74. What is gradient descent in machine learning and how does it work?

Gradient descent is an optimization algorithm used to minimize the cost function in a machine learning model. It works by iteratively adjusting the model parameters in the direction of the steepest descent of the cost function.

75. What is the difference between supervised and unsupervised learning in machine learning?

Supervised learning is a type of machine learning where the model is trained on labeled data, while unsupervised learning is a type of machine learning where the model is trained on unlabeled data and tries to find patterns or structure in the data.

76.What is the difference between L1 and L2 regularization in machine learning?

L1 regularization penalizes the model for high absolute values of the model parameters, while L2 regularization penalizes the model for high squared values of the model parameters.

77. What is the difference between precision and recall in machine learning?

Precision is the ratio of true positives to the total number of positive predictions, while recall is the ratio of true positives to the total number of actual positives.

78. What is the difference between cross-validation and train-test split in machine learning?

Cross-validation involves splitting the data into multiple train-test splits and averaging the results to evaluate the model performance, while train-test split involves splitting the data into a training set and a test set to evaluate the model performance.

79. What is the difference between a generative model and a discriminative model in machine learning?

A generative model learns the joint probability distribution of the input features and the output labels, while a discriminative model learns the conditional probability distribution of the output labels given the input features.

80. What is the difference between batch gradient descent and stochastic gradient descent in machine learning?

Batch gradient descent updates the model parameters using the gradient of the cost function calculated over the entire training set, while stochastic gradient descent updates the model parameters using the gradient of the cost function calculated over a single data point or a small batch of data points.

81. What is the difference between an autoencoder and a variational autoencoder in machine learning?

An autoencoder is an unsupervised learning model that learns to compress and reconstruct the input data, while a variational autoencoder is a generative model that learns to model the probability distribution of the input data in a latent space.

82. What is the difference between a kernel density estimator and a histogram in machine learning?

A kernel density estimator is a non-parametric model that estimates the probability density function of the input data by smoothing it with a kernel function, while a histogram is a discrete representation of the input data by dividing it into bins.

84. What is the difference between a linear regression model and a logistic regression model in machine learning?

A linear regression model is used for continuous output variables, while a logistic regression model is used for binary output variables.

85. What is the difference between a parametric and a non-parametric model in machine learning?

A parametric model assumes a specific functional form for the relationship between the input data and the output labels, while a non-parametric model does not make any assumptions about the functional form and learns it directly from the data.

86. What is the difference between a classification problem and a regression problem in machine learning?

A classification problem is a type of machine learning problem where the output variable is categorical, while a regression problem is a type of machine learning problem where the output variable is continuous.

87. What is the difference between a feature and a label in machine learning?

A feature is an input variable used to predict the output label in a machine learning model, while a label is the output variable that the model is trained to predict.

88.What is the difference between a linear and a non-linear model in machine learning?

A linear model assumes a linear relationship between the input data and the output labels, while a non-linear model assumes a non-linear relationship between the input data and the output labels.

89. What is the difference between a batch normalization and a layer normalization in machine learning?

Batch normalization normalizes the input data across the batch dimension, while layer normalization normalizes the input data across the feature dimension.

90. What is the difference between a parametric and a non-parametric test in machine learning?

A parametric test assumes a specific distribution for the input data, while a non-parametric test does not make any assumptions about the distribution and uses methods such as bootstrapping or permutation testing to make statistical inferences.

91. What is cross-validation in machine learning, and why is it important?

Cross-validation is a technique for evaluating the performance of a machine learning model by splitting the data into multiple subsets, training the model on some subsets, and evaluating it on the remaining subset. It is important because it provides a more accurate estimate of the model's performance on new data and helps prevent overfitting.

92. What is gradient descent in machine learning, and how does it work?

Gradient descent is an optimization algorithm used to minimize the error or loss function in a machine learning model by iteratively adjusting the model's parameters in the direction of the steepest descent of the gradient.

93.What is the difference between a supervised and unsupervised learning in machine learning?

In supervised learning, the model is trained on labeled data, where the output variable is known, while in unsupervised learning, the model is trained on unlabeled data, where the output variable is unknown.

94. What is the difference between precision and recall in machine learning?

Precision is the ratio of true positives to the total predicted positives, while recall is the ratio of true positives to the total actual positives. Precision measures how accurate the model's positive predictions are, while recall measures how well the model identifies all positive instances.

95. What is the difference between a linear and a logistic regression in machine learning?

A linear regression is used for predicting continuous output variables, while logistic regression is used for predicting binary or categorical output variables.

96. What is regularization in machine learning, and why is it important?

Regularization is a technique used to prevent overfitting in machine learning models by adding a penalty term to the loss function that penalizes large parameter values. It is important because it helps improve the model's generalization performance on new data.

97. What is the difference between a local and a global minimum in machine learning?

A local minimum is the point in the optimization landscape where the loss function is minimized locally, while a global minimum is the point where the loss function is minimized globally. A local minimum may not be the best solution as it may not be the global minimum.

98. What is a neural network in machine learning, and how does it work?

A neural network is a type of machine learning model that consists of multiple layers of interconnected neurons that learn the complex relationships between the input data and the output labels by adjusting the weights and biases of the neurons during training.

99. What is the difference between a kernel and a distance metric in machine learning?

A kernel is a function that transforms the input data into a higher-dimensional feature space, where the data can be more easily separated. A distance metric measures the similarity or dissimilarity between two data points in the input space.

100. What is ensemble learning in machine learning, and how does it work?

Ensemble learning is a technique for combining multiple machine learning models to improve the predictive performance. It works by training several models on different subsets of the data or using different algorithms and then combining their predictions through voting or averaging.

101. Explain Artificial Intelligence and give its applications.

Artificial Intelligence (AI) is a field of Computer Science focuses on creating systems that can perform tasks that would typically require human intelligence, such as recognizing speech, understanding natural language, making decisions, and learning. We use AI to build various applications, including image and speech recognition, natural language processing (NLP), robotics, and machine learning models like neural networks.

102. How are machine learning and AI related?

Machine learning and Artificial Intelligence (AI) are closely related but distinct fields within the broader domain of computer science. Al includes not only machine learning but also other approaches, like rule-based systems, expert systems, and knowledge-based systems, which do not necessarily involve learning from data. Many state-of-the-art AI systems are built upon machine learning techniques, as these approaches have proven to be highly effective in tackling complex, data-driven problems.

103. What is Deep Learning based on?

Deep learning is a subfield of machine learning that focuses on the development of artificial neural networks with multiple layers, also known as deep neural networks. These networks are particularly effective in modeling complex, hierarchical patterns and representations in data. Deep learning is inspired by the structure and function of the human brain, specifically the biological neural networks that make up the brain.

Learn more about AI vs ML vs Deep Learning here.

104. How many layers are in a Neural Network?

Neural networks are one of many types of ML algorithms that are used to model complex patterns in data. They are composed of three layers — input layer, hidden layer, and output layer.

105.Explain TensorFlow.

TensorFlow is an open-source platform developed by Google designed primarily for high-performance numerical computation. It offers a collection of workflows that can be used to develop and train models to make machine learning robust and efficient. TensorFlow is customizable, and thus, helps developers create experiential learning architectures and work on the same to produce desired results.

106. What are the pros of cognitive computing?

Cognitive computing is a type of AI that mimics human thought processes. We use this form of computing to solve problems that are complex for traditional computer systems. Some major benefits of cognitive computing are:

- It is the combination of technology that helps to understand human interaction and provide answers.
- Cognitive computing systems acquire knowledge from the data.
- These computing systems also enhance operational efficiency for enterprises.

107. What's the difference between NLP and NLU?

Natural Language Processing (NLP) and Natural Language Understanding (NLU) are two closely related subfields within the broader domain of Artificial Intelligence (AI), focused on the interaction between computers and human languages. Although they are often used interchangeably, they emphasize different aspects of language processing.

NLP deals with the development of algorithms and techniques that enable computers to process, analyze, and generate human language. NLP covers a wide range of tasks, including text analysis, sentiment analysis, machine translation, summarization, part-of-speech tagging, named-entity recognition, and more. The goal of NLP is to enable computers to effectively handle text and speech data, extract useful information, and generate human-like language outputs.

While, NLU is a subset of NLP that focuses specifically on the comprehension and interpretation of meaning from human language inputs. NLU aims to disambiguate the nuances, context, and intent in human language, helping machines grasp not just the structure but also the underlying meaning, sentiment, and purpose. NLU tasks may include sentiment analysis, question-answering, intent recognition, and semantic parsing.

108. Give some examples of weak and strong Al.

Some examples of weak AI include rule-based systems and decision trees. Basically, those systems that require an input come under weak AI. On the other hand, a strong AI includes neural networks and deep learning, as these systems and functions can teach themselves to solve problems.

109. What is the need of data mining?

Data mining is the process of discovering patterns, trends, and useful information from large datasets using various algorithms, statistical methods, and machine learning techniques. It has gained

significant importance due to the growth of data generation and storage capabilities. The need for data mining arises from several aspects, including decision-making.

110. Name some sectors where data mining is applicable.

There are many sectors where data mining is applicable, including:

- **Healthcare** t is used to predict patient outcomes, detection of fraud and abuse, measure the effectiveness of certain treatments, and develop patient and doctor relationships.
- **Finance** The finance and banking industry depends on high-quality, reliable data. It can be used to predict stock prices, predict loan payments and determine credit ratings.
- **Retail** It is used to predict consumer behavior, noticing buying patterns to improve customer service and satisfaction.

111. What are the components of NLP?

There are three main components to NLP:

- Language understanding This defines the ability to interpret the meaning of a piece of text
- Language generation This is helpful in producing text that is grammatically correct and conveys the intended meaning.
- Language processing This helps in performing operations on a piece of text, such as tokenization, lemmatization, and part-of-speech tagging.

112.What is the full form of LSTM?

LSTM stands for Long Short-Term Memory, and it is a type of recurrent neural network (RNN) architecture that is widely used in artificial intelligence and natural language processing. LSTM networks have been successfully used in a wide range of applications, including speech recognition, language translation, and video analysis, among others.

113. What is Artificial Narrow Intelligence (ANI)?

Artificial Narrow Intelligence (ANI), also known as Weak AI, refers to AI systems that are designed and trained to perform a specific task or a narrow range of tasks. These systems are highly specialized and can perform their designated task with a high degree of accuracy and efficiency. This type of technology is also known as Weak AI.

114.What is a data cube?

A data cube is a multidimensional (3D) representation of data that can be used to support various types of analysis and modeling. Data cubes are often used in machine learning and data mining applications to help identify patterns, trends, and correlations in complex datasets.

115. What is the difference between model accuracy and model performance?

Model accuracy refers to how often a model correctly predicts the outcome of a specific task on a given dataset. Model performance, on the other hand, is a broader term that encompasses various aspects of a model's performance, including its accuracy, precision, recall, F1 score, AUC-ROC, etc. Depending on the problem you're solving, one metric may be more important than the other.

116.What are different components of GAN?

Generative Adversarial Network (GAN) are a class of deep learning models that consist of two primary components working together in a competitive setting. GANs are used to generate new, synthetic data that closely resemble a given real-world dataset. The two main components of a GAN are:

- Generator: The generator is a neural network that takes random noise as input and
 generates synthetic data samples. The aim of the generator is to produce realistic data that
 mimic the distribution of the real-world data. As the training progresses, the generator
 becomes better at generating data that closely resemble the original dataset, without
 actually replicating any specific instances.
- Discriminator: The discriminator is another neural network that is responsible for distinguishing between real data samples (from the original dataset) and synthetic data samples (generated by the generator). Its objective is to correctly classify the input as real or synthesized.

117. What are common data structures used in deep learning?

- 1. Deep learning models involve handling various types of data, which require specific data structures to store and manipulate the data efficiently. Some of the most common data structures used in deep learning are:
- 2. Tensors: Tensors are multi-dimensional arrays and are the fundamental data structure used in deep learning frameworks like TensorFlow and PyTorch. They are used to represent a wide variety of data, including scalars, vectors, matrices, or higher-dimensional arrays.
- 3. Matrices: Matrices are two-dimensional arrays and are a special case of tensors. They are widely used in linear algebra operations that are common in deep learning, such as matrix multiplication, transpose, and inversion.
- 4. Vectors: Vectors are one-dimensional arrays and can also be regarded as a special case of tensors. They are used to represent individual data points, model parameters, or intermediate results during calculations.
- 5. Arrays: Arrays are fixed-size, homogeneous data structures that can store elements in a contiguous memory location. Arrays can be one-dimensional (similar to vectors) or multi-dimensional (similar to matrices or tensors).

118. What is the role of the hidden layer in a neural network?

The hidden layer in a neural network is responsible for mapping the input to the output. The hidden layer's function is to extract and learn features from the input data that are relevant for the given task. These features are then used by the output layer to make predictions or classifications.

In other words, the hidden layer acts as a "black box" that transforms the input data into a form that is more useful for the output layer.

119. Mention some advantages of neural networks.

- Some advantages of neural networks include:
- Neural networks need less formal statistical training.
- Neural networks can detect non-linear relationships between variables and can identify all types of interactions between predictor variables.
- Neural networks can handle large amounts of data and extract meaningful insights from it.
 This makes them useful in a variety of applications, such as image recognition, speech recognition, and natural language processing.
- Neural networks are able to filter out noise and extract meaningful features from data. This
 makes them useful in applications where the data may be noisy or contain irrelevant
 information.
- Neural networks can adapt to changes in the input data and adjust their parameters accordingly. This makes them useful in applications where the input data is dynamic or changes over time.

120. What is the difference between stemming and lemmatization?

The main difference between stemming and lemmatization is that stemming is a rule-based process, while lemmatization is a more sophisticated, dictionary-based approach.

Image 17-05-23 at 8.26 PM_11zon.webp

121. What are the different types of text summarization?

There are two main types of text summarization:

Extraction-based: It does not take new phrases and words; instead, it uses the already existing phrases and words and presents only that. Extraction-based summarization ranks all the sentences according to the relevance and understanding of the text and presents you with the most important sentences.

Abstraction-based: It creates phrases and words, puts them together, and makes a meaningful word or sentence. Along with that, abstraction-based summarization adds the most important facts found in the text. It tries to find out the meaning of the whole text and presents the meaning to you.

122. What is the meaning of corpus in NLP?

Corpus in NLP refers to a large collection of texts. A corpus can be used for various tasks such as building dictionaries, developing statistical models, or simply for reading comprehension.

123.Explain binarizing of data.

Binarizing of data is the process of converting data features of any entity into vectors of binary numbers to make classifier algorithms more productive. The binarizing technique is used for the recognition of shapes, objects, and characters. Using this, it is easy to distinguish the object of interest from the background in which it is found.

124. What is perception and its types?

- 1. Perception is the process of interpreting sensory information, and there are three main types of perception: visual, auditory, and tactile.
- 2. Vision: It is used in the form of face recognition, medical imaging analysis, 3D scene modeling, video recognition, human pose tracking, and many more
- 3. Auditory: Machine Auditory has a wide range of applications, such as speech synthesis, voice recognition, and music recording. These solutions are integrated into voice assistants and smartphones.
- 4. Tactile: With this, machines are able to acquire intelligent reflexes and better interact with the environment.

125. Give some pros and cons of decision trees.

Decision trees have some advantages, such as being easy to understand and interpret, but they also have some disadvantages, such as being prone to overfitting.

126. Explain marginalization process. ean coe

The marginalization process is used to eliminate certain variables from a set of data, in order to make the data more manageable. In probability theory, marginalization involves integrating over a subset of variables in a joint distribution to obtain the distribution of the remaining variables. The process essentially involves "summing out" the variables that are not of interest, leaving only the variables that are desired.

127. What is the function of an artificial neural network?

An artificial neural network is a ML algorithm that is used to simulate the workings of the human brain. ANNs consist of interconnected nodes (also known as neurons) that process and transmit information in a way that mimics the behavior of biological neurons.

The primary function of an artificial neural network is to learn from input data, such as images, text, or numerical values, and then make predictions or classifications based on that data. ANNs can be used for a wide range of tasks, such as image recognition, natural language processing, and predictive analytics.

128. Explain cognitive computing and its types?

Cognitive computing is a subfield of AI that focuses on creating systems that can mimic human cognition and perform tasks that require human-like intelligence. The primary goal of cognitive

computing is to enable computers to interact more naturally with humans, understand complex data, reason, learn from experience, and make decisions autonomously.

There is no strict categorization of cognitive computing types; however, the key capabilities and technologies associated with cognitive computing can be grouped as follows:

- **NLP:** NLP techniques enable cognitive computing systems to understand, process, and generate human language in textual or spoken form.
- **Machine Learning:** Machine learning is essential for cognitive computing, as it allows systems to learn from data, adapt, and improve their performance over time.
- Computer Vision: Computer vision deals with the interpretation and understanding of visual information, such as images and videos. In cognitive computing, it is used to extract useful information from visual data, recognize objects, understand scenes, and analyze emotions or expressions.

129. Explain the function of deep learning frameworks.

Deep learning frameworks are software libraries and tools designed to simplify the development, training, and deployment of deep learning models. They provide a range of functionalities that support the implementation of complex neural networks and the execution of mathematical operations required for their training and inference processes. Some popular deep learning frameworks are TensorFlow, Keras, and PyTorch.

130. How are speech recognition and video recognition different?

Speech recognition and video recognition are two distinct areas within AI and involve processing and understanding different types of data. While they share some commonalities in terms of using machine learning and pattern recognition techniques, they differ in the data, algorithms, and objectives associated with each domain.

Speech Recognition focuses on the automatic conversion of spoken language into textual form. This process involves understanding and transcribing the spoken words, phrases, and sentences from an audio signal.

Video Recognition deals with the analysis and understanding of visual information in the form of videos. This process primarily involves extracting meaningful information from a series of image frames, such as detecting objects, recognizing actions, identifying scenes, and tracking moving objects.

131. What is the pooling layer on CNN?

A pooling layer is a type of layer used in a convolutional neural network (CNN). Pooling layers downsample the input feature maps by summary pooled areas. This reduces the dimensionality of the feature map and makes the CNN more robust to small changes in the input.

132. What is the purpose of Boltzmann machine?

Boltzmann machines are a type of energy-based model which learn a probability distribution by simulating a system of diverging and converging nodes. These nodes act like neurons in a neural network, and can be used to build deep learning models.

133. What do you mean by regular grammar?

Regular grammar is a type of grammar that specifies a set of rules for how strings can be formed from a given alphabet. These rules can be used to generate new strings or to check if a given string is valid.

134. How do you obtain data for NLP projects?

There are many ways to obtain data for NLP projects. Some common sources of data include texts, transcripts, social media posts, and reviews. You can also use web scraping and other methods to collect data from the internet.

135. Explain regular expression in layman's terms.

Regular expressions are a type of syntax used to match patterns in strings. They can be used to find, replace, or extract text. In layman's terms, regular expressions are a way to describe patterns in data. They are commonly used in programming, text editing, and data processing tasks to manipulate and extract text in a more efficient and precise way.

136. How is NLTK different from spaCy?

Both NLTK and spaCy are popular NLP libraries in Python, but they have some key differences:

NLTK is a general-purpose NLP library that provides a wide range of tools and algorithms for basic NLP tasks such as tokenization, stemming, and part-of-speech tagging. NLTK also has tools for text classification, sentiment analysis, and machine translation. In contrast, spaCy focuses more on advanced NLP tasks such as named entity recognition, dependency parsing, and semantic similarity.

spaCy is generally considered to be faster and more efficient than NLTK due to its optimized Cython-based implementation. spaCy is designed to process large volumes of text quickly and efficiently, making it well-suited for production environments.

137. Name some best tools useful in NLP.

There are several powerful tools and libraries available for Natural Language Processing (NLP) tasks, which cater to various needs like text processing, tokenization, sentiment analysis, machine translation, among others. Some of the best NLP tools and libraries include:

• **NLTK:** NLTK is a popular Python library for working with human language data. It provides easy-to-use interfaces to over 50 corpora and lexical resources, along with text processing libraries for classification, tokenization, stemming, tagging, parsing, and more.

- spaCy: spaCy is a modern, high-performance, and industry-ready NLP library for Python. It
 offers state-of-the-art algorithms for fast and accurate text processing, and includes features
 like part-of-speech tagging, named entity recognition, dependency parsing, and word
 vectors.
- **Gensim:** Gensim is a Python library designed for topic modeling and document similarity analysis. It specializes in unsupervised semantic modeling and is particularly useful for tasks like topic extraction, document comparison, and information retrieval.
- OpenNLP: OpenNLP is an open-source Java-based NLP library that provides various components such as tokenizer, sentence segmenter, part-of-speech tagger, parser, and named entity recognizer. It is widely used for creating natural language processing applications.

138. Are chatbots derived from NLP?

Yes, chatbots are derived from NLP. NLP is used to process and understand human language so that chatbots can respond in a way that is natural for humans.

139. What is embedding and what are some techniques to accomplish embedding?

Embedding is a technique to represent data in a vector space so that similar data points are close together. Some techniques to accomplish embedding are word2vec and GloVe.

- Word2vec: It is used to find similar words which have similar dimensions and, consequently, help bring context. It helps in establishing the association of a word with another similar meaning word through the created vectors.
- **GloVe:** It is used for word representation. GloVe is developed for generating word embeddings by aggregating global word-word co-occurrence matrices from a corpus. The result shows the linear structure of the word in vector space.

INTERMEDIATE ARTIFICIAL INTELLIGENCE INTERVIEW QUESTIONS AND ANSWERS

140. Why do we need activation functions in neural networks?

Activation functions play a vital role in neural networks, serving as a non-linear transformation applied to the output of a neuron or node. They determine the output of a neuron based on the weighted sum of its inputs, introducing non-linearity into the network. The inclusion of activation functions allows neural networks to model complex, non-linear relationships in the data.

141.Explain gradient descent.

Gradient descent is a popular optimization algorithm that is used to find the minimum of a function iteratively. It's widely used in machine learning and deep learning for training models by minimizing the error or loss function, which measures the difference between the predicted and actual values.

142. What is the purpose of data normalization?

Data normalization is a pre-processing technique used in machine learning and statistics to standardize and scale the features or variables in a dataset. The purpose of data normalization is to bring different features or variables to a common scale, which allows for more accurate comparisons and better performance of learning algorithms.

The main purposes of data normalization are:

- 1. **Improving model performance:** Some machine learning algorithms, like gradient-based optimization methods or distance-based classifiers, are sensitive to the feature scale.
- 2. **Ensuring fair comparison:** Normalization brings all features to a comparable range, mitigating the effect of different magnitudes or units of measurement, and ensuring that each feature contributes equally to the model's predictions.
- 3. **Faster convergence:** Gradient-based optimization algorithms can converge faster when data are normalized, as the search space becomes more uniformly scaled and the gradients have a more consistent magnitude.
- Reducing numerical issues: Normalizing data can help prevent numerical issues like over- or underflow that may arise when dealing with very large or very small numbers during calculations.

143. Name some activation functions.

Some common activation functions include sigmoid, tanh, and ReLU.

- 1. **Sigmoid:** Maps the input to a value between 0 and 1, allowing for smooth gradient updates. However, it suffers from the vanishing gradient problem and is not zero-centered.
- 2. **Tanh:** Maps the input to a value between -1 and 1, providing a zero-centered output. Like the sigmoid function, it can also suffer from the vanishing gradient problem.
- 3. **ReLU (Rectified Linear Unit):** Outputs 0 for negative input values and retains the input for positive values. It helps alleviate the vanishing gradient problem and has faster computation time, but the output is not zero-centered and can suffer from the dying ReLU issue.

144. Briefly explain data augmentation.

Data augmentation is a technique used to increase the amount of data available for training a machine learning model. This is especially important for deep learning models, which require large amounts of data to train.

145. What is the Swish function?

The Swish function is an activation function. It is a smooth, non-linear, and differentiable function that has been shown to outperform some of the traditional activation functions, like ReLU, in certain deep learning tasks.

146. Explain forward propagation and backpropagation.

Forward propagation is the process of computing the output of a neural network given an input. Forward propagation involves passing an input through the network, layer by layer, until the output is produced. Each layer applies a transformation to the output of the previous layer using a set of weights and biases. The activation function is applied to the transformed output, producing the final output of the layer.

On the other hand, backpropagation is the process of computing the gradient of the loss function with respect to the weights of the network. It is used to update the weights and biases of the network during the training process. It involves calculating the gradient of the loss function with respect to each weight and bias in the network. The gradient is then used to update the weights and biases using an optimization algorithm such as gradient descent.

147. What is classification and its benefits?

Classification is a type of supervised learning task in machine learning and statistics, where the objective is to assign input data points to one of several predefined categories or labels. In a classification problem, the model is trained on a dataset with known labels and learns to predict the category to which a new, unseen data point belongs. Examples of classification tasks include spam email detection, image recognition, and medical diagnosis.

Some benefits of classification include:

- **1. Decision-making**: Classification models can help organizations make informed decisions based on patterns and relationships found in the data.
- 2. Pattern recognition: Classification algorithms are capable of identifying and learning complex patterns in data, enabling them to predict the category of new inputs accurately.
- **3. Anomaly detection**: Classification models can be used to detect unusual or anomalous data points that don't fit the learned patterns.
- **4. Personalization and recommendation**: Classification models can be used to tailor content and recommendations to individual users, enhancing user experiences and increasing engagement.

148. What is a convolutional neural network?

Convolutional neural networks are a type of neural network that is well-suited for image classification tasks. In classification, the model learns to classify input data into one or more predefined classes or categories based on the features of the data. There are various benefits of classification, and it has numerous practical applications in different fields,

such as:

- **Object Recognition:** It is used in image and speech recognition to identify objects, faces, or voices.
- **Sentiment Analysis:** It helps understand the polarity of textual data, which can be used to gauge customer feedback, opinions, and emotions.

Email Spam Filtering: It can be used to classify emails into a spam or non-spam categories to improve email communication.

149. Explain autoencoders and its types.

Autoencoders are a type of neural network that is used for dimensionality reduction. The different types of autoencoders include Denoising, Sparse, Undercomplete, etc.

- Denoising Autoencoder: It is used to achieve good representation, meaning it can be
 obtained robustly from a corrupted input, which will be useful for recovering the
 corresponding clean input.
- **Sparse Autoencoder:** This has a sparsity penalty, a value close to zero but not exactly zero. It is applied on the hidden layer in addition to the reconstruction error, which prevents overfitting.
- **Undercomplete Autoencoder:** This does not need any regularization because they maximize the probability of data rather than copying the input to the output.

150. State fuzzy approximation theorem.

Fuzzy approximation theorem states that a function can be approximated as closely as desired using a combination of fuzzy sets. The theorem states that any continuous function can be represented as a weighted sum of linear functions, where the weights are fuzzy sets that capture the input variables' uncertainty.

151. What are the main components of LSTM?

LSTM stands for Long Short-Term Memory. It is a neural network architecture that is used for modeling time series data. LSTM has three main components:

- **The forget gate:** This gate decides how much information from the previous state is to be retained in the current state.
- **The input gate:** This gate decides how much new information from the current input is to be added to the current state.
- The output gate: This gate decides what information from the current state is to be output.

152. Give some benefits of transfer learning.

Transfer learning is a machine learning technique where you use knowledge from one domain and apply it to another domain. This is usually done to accelerate the learning process or to improve performance.

There are several benefits of transfer learning:

- Learn from smaller datasets: If you have a small dataset, you can use transfer learning to learn from a larger dataset in the same domain. This will help you to build better models.
- Learn from different domains: You can use transfer learning to learn from different domains. For example, if you want to build a computer vision model, you can use knowledge from the medical domain.

- **Better performance:** Transfer learning can help you to improve the performance of your models and apply it on other domains to build better models.
- **Pre-trained models:** If you use a pre-trained model, you can save time and resources. This is because you don't have to train the model from scratch.
- Use of fine-tune models: You can fine-tune models using transfer learning. Also, you can adapt the model to your specific needs.

153. Explain the importance of cost/loss function.

The cost/loss function is an important part of machine learning that maps a set of input parameters to a real number that represents the cost or loss. The cost/loss function is used for optimization problems. The goal of optimization is to find the set of input parameters that minimize the cost/loss function.

154. Define the following terms - Epoch, Batch, and Iteration?

Epoch, batch, and iteration are all important terms in machine learning. Epoch refers to the number of times the training dataset is used to train the model; Batch refers to the number of training samples used in one iteration; Iteration is the number of times the training algorithm is run on the training dataset.

156. Explain dropouts.

Dropout is a method used to prevent the overfitting of a neural network. It refers to dropping out some neural network units. The process is similar to that of natural reproduction, where distinct genes combine to produce offspring while the other genes are dropped out instead of strengthening their co-adaptation.

157. Explain vanishing gradient

As more layers are added and the distance from the final layer increases, backpropagation is not as helpful in sending information to the lower layers. As a result, the information is sent back, and the gradients start disappearing and becoming small in relation to network weights. These disappearing gradients are known as vanishing gradients.

158. Explain the function of batch Gradient Descent.

Batch gradient descent is an optimization algorithm that calculates the gradient of the cost function with respect to the weights of the model for each training batch. The weights are updated in the direction that decreases the cost function.

159. What is an Ensemble learning method?

Ensemble learning is a method of combining multiple models to improve predictive accuracy. These methods usually cost more to train but can provide better accuracy than a single model.

160. What are some drawbacks of machine learning?

One of the biggest drawbacks of Machine learning is that it can be biased if the data used to train the algorithm is not representative of the real world. For example, if an algorithm is trained using data that is mostly from one gender or one race, it may be biased against other genders or races.

Here are some other disadvantages of Machine Learning:

- Possibility of high Error
- Algorithm selection
- Data acquisition
- Time and space
- High production costs
- Lacking the skills to innovate

161. Explain Sentimental analysis in NLP?

Sentiment analysis is the process of analyzing text to determine the emotional tone of the text in NLP. This can be helpful in customer service to understand how customers are feeling, or in social media to understand the general public sentiment about a topic.

162. What is BFS and DFS algorithm?

Breadth-First Search (BFS) and Depth-First Search (DFS) are two algorithms used for graph traversal. BFS algorithm starts from the root node (or any other selected node) and visits all the nodes at the same level before moving to the next level

On the other hand, DFS algorithm starts from the root node (or any other selected node) and explores as far as possible along each branch before backtracking.

163. Explain the difference between supervised and unsupervised learning.

Supervised learning involves training a model with labeled data, where both input features and output labels are provided. The model learns the relationship between inputs and outputs to make predictions for unseen data. Common supervised learning tasks include classification and regression.

Unsupervised learning, on the other hand, uses unlabeled data where only input features are provided. The model seeks to discover hidden structures or patterns in the data, such as clusters or data representations. Common unsupervised learning tasks include clustering, dimensionality reduction, and anomaly detection.

164. What is the text extraction process?

Text extraction is the process of extracting text from images or other sources. This can be done with OCR (optical character recognition) or by converting the text to a format that can be read by a text-to-speech system.

165. What are some disadvantages of linear models?

- 1. Here are some disadvantages of using linear models -
- 2. They can be biased if the data used to train the model is not representative of the real world.
- 3. Linear models can also be overfit if the data used to train the model is too small.
- 4. Linear models assume a linear relationship between the input features and the output variable, which may not hold in reality. This can lead to poor predictions and decreased model performance.

167. Mention methods for reducing dimensionality.

Artificial intelligence interview questions like this can be easy and difficult at the same time as you may know the answers but not on the tip of your tongue. Hence, a quick refresher can help a lot. Reducing dimensionality refers to the reduction of the number of random variables. This can be achieved by different techniques including principal component analysis, low variance filter, missing values ratio, high correlation filter, random forest, and others.

168. Explain cost function.

This is a popular AI interview question. A cost function is a scalar function that helps to identify how wrong an AI model is with regard to its ability to determine the relationship between X and Y. In other words, it tells us the neural network's error factor.

The neural network works better when the cost function is lower. For instance, it takes the output predicted by the neural network and the actual output and then computes how incorrect the model was in its prediction.

So, the cost function will give a lower number if the predictions don't differ too much from the actual values and vice-versa

169. Mention hyper-parameters of ANN.

The hyper-parameters of ANN are as follows:

Learning rate: It refers to the speed with which the network gets familiar with its parametersMomentum: This parameter enables coming out of the local minima and smoothening jumps during gradient descent

The number of epochs: This parameter refers to the number of times the whole training dataset is fed to the network during training. One must increase the number of epochs until a decrease in validation accuracy is noticed, even if there is an increase in training accuracy, which is called overfitting.

- **Number of hidden layers:** This parameter specifies the number of layers between the input and output layers.
- **Number of neurons in each hidden layer:** This parameter specifies the number of neurons in each hidden layer.
- Activation functions: Activation functions are responsible for determining a neuron's output based on the weighted sum of its inputs. Widely used activation functions include Sigmoid, ReLU, Tanh, and others.

170. Explain intermediate tensors. Do sessions have a lifetime?

Intermediate tensors are temporary data structures in a computational graph that store intermediate results when executing a series of operations in Artificial Intelligence, particularly in deep learning frameworks. These tensors represent the values produced during the forward pass of a neural network while processing input data before reaching the final output.

Yes, sessions have a lifetime, which starts when the session is created and ends when the session is closed or the script is terminated. In TensorFlow 1.x, sessions were used to execute and manage operations in a computational graph. A session allowed the allocation of memory for tensor values and held necessary resources to execute the operations. In TensorFlow 2.x, sessions and computational graphs have been replaced with a more dynamic and eager execution approach, allowing for simpler and more Pythonic code.

171. Explain Exploding variables

Exploding variables are a phenomenon in which the magnitude of a variable grows rapidly over time, often leading to numerical instability and overflow errors. This can happen when a variable is repeatedly multiplied or divided by a value that is greater than 1 or less than -1. As a result, the variable's value grows exponentially or collapses to zero, causing computational problems.

172. Is it possible to build a deep learning model only using linear regression?

Linear regression is a basic tool in statistical learning, but it cannot be used to build a deep learning model. Deep learning models require non-linear functions to learn complex patterns in data.

173. What is the function of Hyperparameters?

Hyperparameters are parameters that are not learned by the model. They are set by the user and used to control the model's behavior.

174. What is Artificial Super Intelligence (ASI)?

An Artificial Super Intelligence system is not one that has been achieved yet. Also known as Super AI, it is a hypothetical system that can surpass human intelligence and execute any task better than a human. The concept of ASI suggests that such an AI can exceed all human intelligence. It can even

take complex decisions in harsh conditions and think just like a human would, or even better, develop emotional, sensible relationships.

175. What is overfitting, and how can it be prevented in an AI model?

Overfitting occurs when a model learns the training data too well, including capturing noise and random fluctuations. This often results in a model that performs poorly on unseen or validation data. Techniques to prevent overfitting include:

Regularization (L1 or L2)

Early stopping

Cross-validation

Using more training data

Reducing model complexity

176. What is the role of pipeline for Information extraction (IE) in NLP?

Pipelines are used in information extraction to sequentially apply a series of processing steps to input data. This allows for efficient data processing and helps avoid errors.

177. What is the difference between full listing hypothesis and minimum redundancy hypothesis?

Full listing hypothesis states that all possible values of a variable should be listed in the data dictionary. Minimum redundancy hypothesis states that all values of a variable should be listed in the data dictionary, but that only the most important values should be listed multiple times. Looking for remote developer job at US companies?

178. Mention the steps of the gradient descent algorithm.

The gradient descent algorithm helps in optimization and in finding coefficients of parameters that help minimize the cost function. The steps that help achieve this are as follows:

- **Step 1**: Give weights (x,y) random values and then compute the error, also called Sum of Squares Error (SSE).
- **Step 2:** Compute the gradient or the change in SSE when you change the value of the weights (x,y) by a small amount. This step helps us identify the direction in which we must move x and y to minimize SSE.
- Step 3: Adjust the weights with the gradients for achieving optimal values for the minimal SSE.
- **Step 4:** Change the weights for predicting and calculating the new error. Step 5: Repeat steps 2 and 3 till the time making more adjustments stops producing significant error reduction.

These types of artificial intelligence interview questions help hiring managers properly guage a candidate's expertise in this domain. Hence, you must thoroughly understand such questions and enlist all steps properly to move ahead.

179. How to handle an imbalance dataset?

There are a number of ways to handle an imbalanced dataset, such as using different algorithms, weighting the classes, or oversampling the minority class.

- Algorithm selection: Some algorithms are better suited to handle imbalanced data than
 others. For example, decision trees and random forests tend to work well on imbalanced
 data, while algorithms like logistic regression or support vector machines may struggle.
- **Class weighting:** By assigning higher weights to the minority class, you can make the algorithm give more importance to it during training. This can help prevent the algorithm from always predicting the majority class.
- Oversampling: You can create synthetic samples of the minority class by randomly duplicating existing samples or generating new samples based on the existing ones. This can balance the class distribution and help the algorithm learn more about the minority class.

180. How do you solve the vanishing gradient problem in RNN?

The vanishing gradient problem is a difficulty encountered when training artificial neural networks using gradient-based learning methods. This problem is resolved by replacing the activation function of the network. You can use the Long Short-Term Memory (LSTM) network to solve the problem.

It has three gates called input, forgets, and output gates. Here forget gates constantly observe what information needs to be dropped going through the network. In this way, we have short and long-term memory. So, we can transfer the information through the network and retrieve it even at the last stage to identify the context of prediction.