



Machine Learning Interview Questions

[Click here](#) to view the live version of the page

Top 70 Machine Learning Interview Questions and Answers

Machine Learning and Artificial Intelligence are among the most popular technologies in the world today. This comprehensive blog consists of some of the most frequently asked Machine Learning interview questions that aim to help you revise all the necessary concepts and skills to land your dream job. This blog is specifically designed for you to do a thorough Machine Learning interview preparation before going for the interview.

[Basic Machine Learning Interview Questions for Freshers](#)

[Machine Learning Interview Questions for Intermediate](#)

[Machine Learning Interview Questions for Experienced](#)

[Role-Specific Machine Learning Questions](#)

[FAANG Machine Learning Engineer Questions](#)

[How to Prepare for the Machine Learning Interview](#)

[Machine Learning Salary Based on Experience](#)

[Machine learning trends in 2024](#)

[Job Opportunities in Machine Learning](#)

[Roles and Responsibilities of a Machine Learning Engineer](#)

[Conclusion](#)

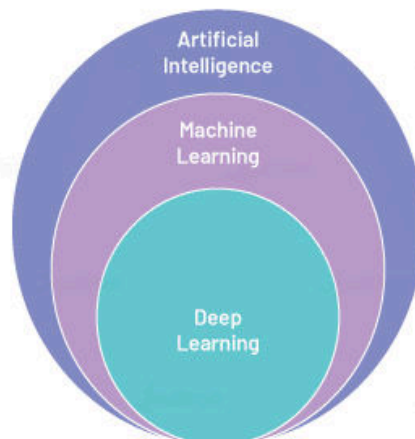
Did You Know?

1. According to [Forbes](#), machine learning is implemented in more than 75% of businesses across multiple business units.
2. According to [Forbes](#), at least 90% of the total data that was generated by the world happened in the last 3 years, and the volume has only doubled in size in the last 2 years.
3. According to Netflix software engineer Aish Fenton, [recommendations account for 80% of the movies watched on Netflix](#).

Basic Machine Learning Interview Questions for Freshers

1. Explain Machine Learning, Artificial Intelligence, and Deep Learning

It is common to get confused between the three in-demand technologies, Machine Learning, Artificial Intelligence, and Deep Learning. These three technologies, though a little different from one another, are interrelated. While Deep Learning is a subset of Machine Learning, Machine Learning is a subset of Artificial Intelligence. Since some terms and techniques may overlap in these technologies, it is easy to get confused among them.



So, let us learn about these technologies in detail:

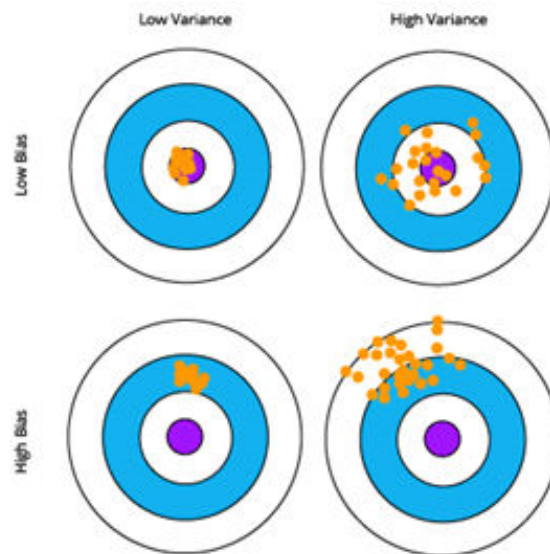
- **Machine Learning:** Machine Learning involves various statistical and Deep Learning techniques that allow machines to use their past experiences and get better at performing specific tasks without having to be monitored.
- **Artificial Intelligence:** Artificial Intelligence uses numerous Machine Learning and Deep Learning techniques that enable computer systems to perform tasks using human-like intelligence with logic and rules. Artificial intelligence is used in every sector hence it is necessary to pursue [Artificial Intelligence Course](#) to make your career in AI.
- **Deep Learning:** Deep Learning comprises several algorithms that enable software to learn from themselves and perform various business tasks including image and speech recognition. Deep Learning is possible when systems expose their multilayered neural networks to large volumes of data for learning.

Willing to master AI & ML skills? Check our [AI and Machine Learning Courses](#) in collaboration with top universities Now!

2. What is Bias and Variance in Machine Learning?

- Bias is the difference between the average prediction of a model and the correct value of the model. If the bias value is high, then the prediction of the model is not accurate. Hence, the bias value should be as low as possible to make the desired predictions.
- Variance is the number that gives the difference of prediction over a training set and the anticipated value of other training sets. High variance may lead to large fluctuation in the output. Therefore, a model's output should have low variance.

The following diagram shows the [bias-variance trade-off](#):



Here, the desired result is the blue circle at the center. If we get off from the blue section, then the prediction goes wrong.

Interested in learning Machine Learning? Enroll in our [Machine Learning Training](#) now!

3. What is Clustering in Machine Learning?

Clustering is a technique used in unsupervised learning that involves grouping data points. The clustering algorithm can be used with a set of data points. This technique will allow you to classify all data points into their particular groups. The data points that are thrown into the same category have similar features and properties, while the data points that belong to different groups have distinct features and properties. Statistical data analysis can be performed by this method. Let us take a look at three of the most popular and useful clustering algorithms.

- K-means clustering: This algorithm is commonly used when there is data with no specific group or category. K-means clustering allows you to find the hidden patterns in the data, which can be used to classify the data into various groups. The variable k is used to represent the number of groups the data is divided into, and the data points are clustered using the similarity of features. Here, the centroids of the clusters are used for labeling new data.

- Mean-shift clustering: The main aim of this algorithm is to update the center-point candidates to be mean and find the center points of all groups. In mean-shift clustering, unlike k-means clustering, the possible number of clusters need not be selected as it can automatically be discovered by the mean shift.
- Density-based spatial clustering of applications with noise (DBSCAN): This clustering algorithm is based on density and has similarities with mean-shift clustering. There is no need to preset the number of clusters, but unlike mean-shift clustering, DBSCAN identifies outliers and treats them like noise. Moreover, it can identify arbitrarily-sized and -shaped clusters without much effort.

4. What is Linear Regression in Machine Learning?

Linear Regression is a supervised Machine Learning algorithm. It is used to find the linear relationship between the dependent and independent variables for predictive analysis.

The equation for Linear Regression:

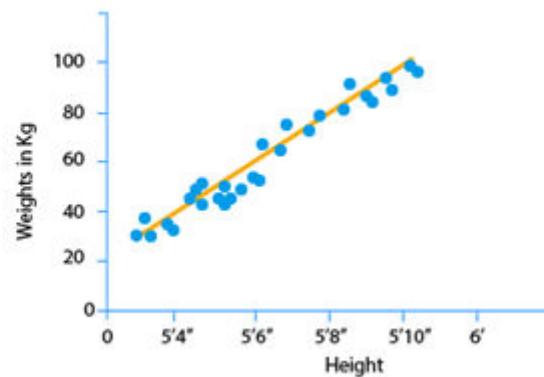
$$Y = A + B.X$$



where:

- X is the input or independent variable
- Y is the output or dependent variable
- a is the intercept, and b is the coefficient of X

Below is the best-fit line that shows the data of weight, Y or the dependent variable, and the



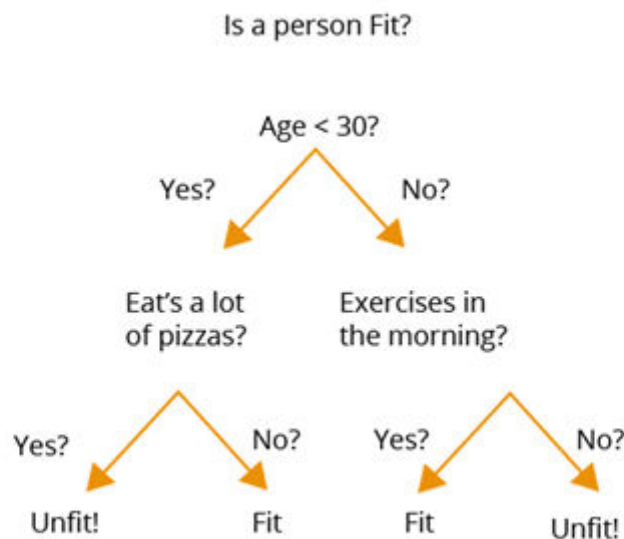
data of height, X or the independent variable, of 21-year-old candidates scattered over the plot. The straight line shows the best linear relationship that would help in predicting the weight of candidates according to their height.

To get this best-fit line, the best values of a and b should be found. By adjusting the values of a and b , the errors in the prediction of Y can be reduced.

This is how linear regression helps in finding the linear relationship and predicting the output.

5. What is a Decision Tree in Machine Learning?

A [decision tree](#) is used to explain the sequence of actions that must be performed to get the desired output. It is a hierarchical diagram that shows the actions.



An algorithm can be created for a decision tree on the basis of the set hierarchy of actions.

In the above decision-tree diagram, a sequence of actions has been made for driving a vehicle with or without a license.

6. What are the types of Machine Learning?

- **Supervised learning:** The algorithms of supervised learning use labeled data to get trained. The models take direct feedback to confirm whether the output that is being predicted is, indeed, correct. Moreover, both the input data and the output data are provided to the model, and the main aim here is to train the model to predict the output upon receiving new data. Supervised learning offers accurate results and can largely be divided into two parts, classification and regression.
- **Unsupervised learning:** The algorithms of unsupervised learning use unlabeled data for training purposes. In unsupervised learning, the models identify hidden data trends and do not take any feedback. The unsupervised learning model is only provided with input data. Unsupervised learning's main aim is to identify hidden patterns to extract information from unknown sets of data. It can also be classified

into two parts, clustering, and associations. Unfortunately, unsupervised learning offers results that are comparatively less accurate.

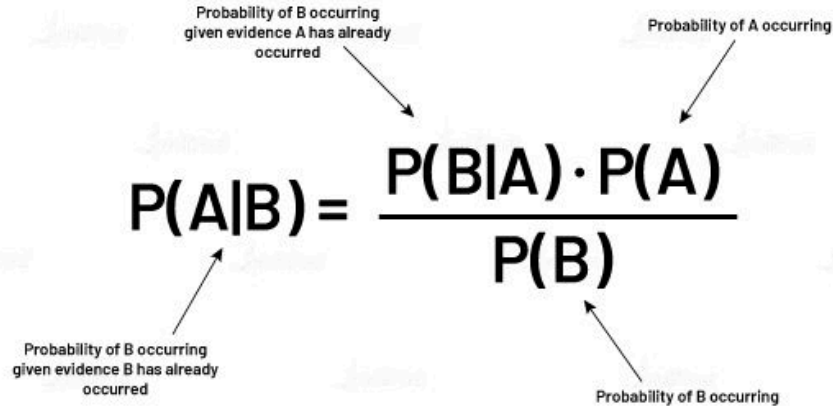
- **Reinforcement learning:** Reinforcement Learning (RL) is a type of machine learning where an agent learns to make decisions by taking actions in an environment to achieve some notion of cumulative reward. It's about trial and error, where the agent discovers through feedback which actions yield the most reward over time. Unlike supervised learning, Reinforcement Learning does not require labeled input/output pairs, and unlike unsupervised learning, the focus is on finding a balance between exploration (of uncharted territory) and exploitation (of current knowledge). The applications of RL range from robotics, where it can help machines learn complex tasks, to web systems, where it can be used to improve the user experience.

Learn more about Machine Learning from this [Machine Learning tutorial!](#)

7. What is Bayes's Theorem in Machine Learning?

Bayes's theorem offers the probability of any given event to occur using prior knowledge. In mathematical terms, it can be defined as the true positive rate of the given sample condition divided by the sum of the true positive rate of the said condition and the false positive rate of the entire population.

Two of the most significant applications of Bayes's theorem in Machine Learning are Bayesian optimization and Bayesian belief networks. This theorem is also the foundation behind the Machine Learning brand that involves the [Naive Bayes](#) classifier.



Probability of B occurring given evidence A has already occurred

Probability of A occurring

$$P(A|B) = \frac{P(B|A) \cdot P(A)}{P(B)}$$

Probability of B occurring given evidence B has already occurred

Probability of B occurring

8. What is PCA in Machine Learning?

Multidimensional data is at play in the real world. [Data visualization](#) and computation become more challenging with the increase in dimensions. In such a scenario, the dimensions of data might have to be reduced to analyze and visualize it easily. This is done by:

- Removing irrelevant dimensions
- Keeping only the most relevant dimensions


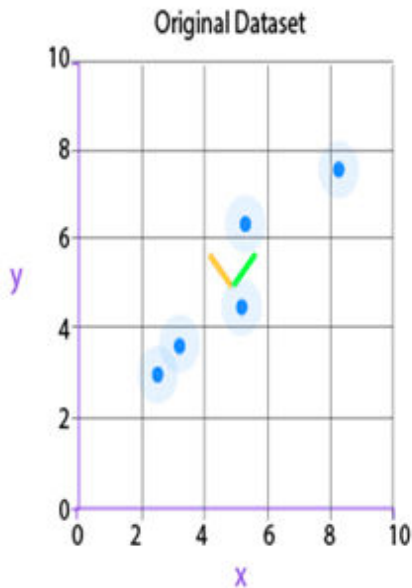
This is where [Principal Component Analysis](#) (PCA) is used.

The goal of PCA is to find a fresh collection of uncorrelated dimensions (orthogonal) and rank them on the basis of variance.

Mechanism of PCA:

- Compute the covariance matrix for data objects
- Compute eigenvectors and eigenvalues in descending order
- Select the initial N eigenvectors to get new dimensions
- Finally, change the initial n -dimensional data objects into N -dimensions

Example: Below are two graphs showing data points or objects and two directions, one is green and the other is yellow. Graph 2 is arrived at by rotating Graph 1 so that the x-axis and y-axis represent the green and yellow direction respectively.



After the rotation of data points, it can be inferred that the green direction, the x-axis, gives the line that best fits the data points.

Here, two-dimensional data is being represented; but in real life, the data would be multidimensional and complex. So, after recognizing the importance of each direction, the area of dimensional analysis can be reduced by cutting off the less-significant directions.

Now, we will go through another important Machine Learning interview question on PCA.

9. What are the types of Machine Learning?

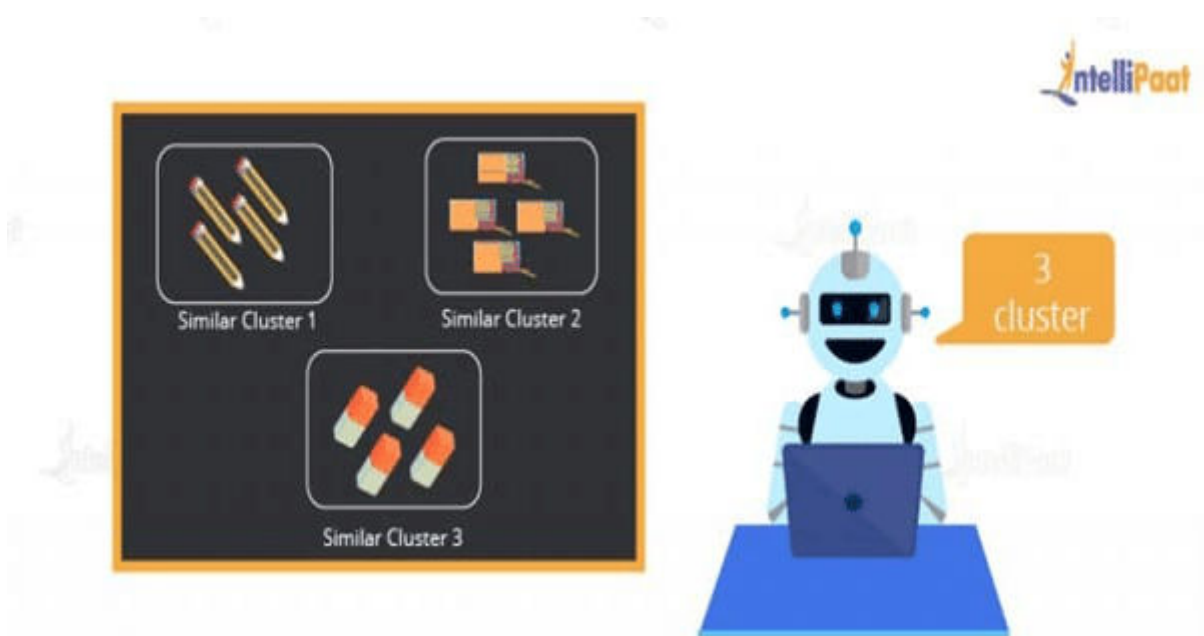

This is one of the most basic interview questions that everyone must know.

So, basically, there are three [types of Machine Learning](#). They are described as follows:


Supervised learning: In this type of Machine Learning, machines learn under the supervision of labeled data. There is a training dataset on which a machine is trained, and it gives the output according to its training.



Unsupervised learning: This type of Machine Learning has unlabeled data unlike supervised learning. Unsupervised learning works on data under absolutely no supervision. Unsupervised learning tries to identify patterns in data and makes clusters of similar entities. After that, when a new input data is fed into the model, it does not identify the entity; rather, it puts the entity in a cluster of similar objects.



Reinforcement learning: Reinforcement learning includes models that learn and traverse to find the best possible move. The algorithms for reinforcement learning are constructed in a way that they try to find the best possible suite of action on the basis of the reward and punishment theory.



For next time....



10. Differentiate between Classification and Regression in Machine Learning

In Machine Learning, there are various types of prediction problems based on supervised and unsupervised learning. They are classification, regression, clustering, and association. Here, we will discuss classification and regression.

Classification: In [classification](#), a Machine Learning model is created that assists in differentiating data into separate categories. The data is labeled and categorized based on the input parameters.

For example, predictions have to be made on the churning out customers for a particular product based on some recorded data. Either the customers will churn out or they will not. So, the labels for this would be “Yes” and “No.”

Regression: It is the process of creating a model for distinguishing data into continuous real values, instead of using classes or discrete values. It can also identify the distribution movement depending on historical data. It is used for predicting the occurrence of an event depending on the degree of association of variables.

For example, the prediction of weather conditions depends on factors such as temperature, air pressure, solar radiation, elevation, and distance from the sea. The relation among these factors assists in predicting the weather condition.

Check out this Executive [M.Tech in Artificial Intelligence & Machine Learning](#) by IIT Jammu to enhance your resume!

11. What is a Confusion Matrix?

[Confusion matrix](#) is used to explain a model’s performance and gives a summary of predictions of the classification problems. It assists in identifying the uncertainty between classes.

Confusion matrix gives the count of correct and incorrect values and error types.
Accuracy of the model:

$$\text{Accuracy} = \frac{\text{TP} + \text{TN}}{\text{TP} + \text{TN} + \text{FP} + \text{FN}}$$



For example, consider the following confusion matrix. It consists of values as true positive, true negative, false positive, and false negative for a classification model. Now, the accuracy of the model can be calculated as follows:

	Predicted: No	Predicted: Yes	
Actual: No	TP = 200	FN = 60	260
Actual: Yes	FP = 10	TN = 50	60
	210	110	



So, in the example:

$$\text{Accuracy} = (200 + 50) / (200 + 50 + 10 + 60) = 0.78$$

This means that the model's accuracy is 0.78, corresponding to its True Positive, True Negative, False Positive, and False Negative values.

12. Explain Logistic Regression

Logistic regression is the proper regression analysis used when the dependent variable is categorical or binary. Like all regression analyses, logistic regression is a technique for **predictive analysis**. Logistic regression is used to explain data and the relationship between one dependent binary variable and one or more independent

variables. Logistic regression is also employed to predict the probability of categorical dependent variables.

Logistic regression can be used in the following scenarios:

- To predict whether a citizen is a Senior Citizen (1) or not (0)
- To check whether a person has a disease (Yes) or not (No)

There are three types of logistic regression:

- Binary logistic regression: In this type of logistic regression, there are only two outcomes possible.

Example: To predict whether it will rain (1) or not (0)

- Multinomial logistic regression: In this type of logistic regression, the output consists of three or more unordered categories.

Example: Predicting whether the prize of the house is high, medium, or low.

- Ordinal logistic regression: In this type of logistic regression, the output consists of three or more ordered categories.

Example: Rating an Android application from one to five stars.

Interested in learning Machine Learning? Enroll in this [Machine Learning Training in Bangalore!](#)

13. Why are Validation and Test Datasets Needed?

Data is split into three different categories while creating a model:

- Training dataset: Training dataset is used for building a model and adjusting its variables. The correctness of the model built on the training dataset cannot be relied on as the model might give incorrect outputs after being fed new inputs.

- Validation dataset: Validation dataset is used to look into a model's response. After this, the hyperparameters on the basis of the estimated benchmark of the validation dataset data are tuned. When a model's response is evaluated by using the validation dataset, the model is indirectly trained with the validation set. This may lead to the overfitting of the model to specific data. So, this model will not be strong enough to give the desired response to real-world data.
- Test dataset: Test dataset is the subset of the actual dataset, which is not yet used to train the model. The model is unaware of this dataset. So, by using the test dataset, the response of the created model can be computed on hidden data. The model's performance is tested on the basis of the test dataset. Note: The model is always exposed to the test dataset after tuning the hyperparameters on top of the validation dataset.

As we know, the evaluation of the model on the basis of the validation dataset would not be enough. Thus, the test dataset is used for computing the efficiency of the model.




Looking for Advanced Courses?
Certification in Big Data Analytics
In partnership with E&ICT, IIT Guwahati
[KNOW MORE](#) 400+ Hrs of Instructor-led Training | Job Assistance


The advertisement features a man with glasses and a red sweater smiling in a classroom setting. The text is overlaid on a white and blue background.

14. Explain the difference between KNN and K-means Clustering

K-nearest neighbors (KNN): It is a supervised Machine Learning algorithm. In KNN, identified or labeled data is given to the model. The model then matches the points based on the distance from the closest points.



K-means clustering: It is an unsupervised Machine Learning algorithm. In K-means clustering, unidentified or unlabeled data is given to the model. The algorithm then creates batches of points based on the average of the distances between distinct points.



15. What is Dimensionality Reduction?

In the real world, **Machine Learning models** are built on top of features and parameters. These features can be multidimensional and large in number.

Sometimes, the features may be irrelevant and it becomes a difficult task to visualize them.

This is where dimensionality reduction is used to cut down irrelevant and redundant features with the help of principal variables. These principal variables conserve the features, and are a subgroup, of the parent variables.

16. What is meant by Parametric and Non-parametric Models?

Parametric models refer to the models having a limited number of parameters. In case of parametric models, only the parameter of a model is needed to be known to make predictions regarding the new data.

Non-parametric models do not have any restrictions on the number of parameters, which makes new data predictions more flexible. In case of non-parametric models, the knowledge of model parameters and the state of the data needs to be known to make predictions.

17. Outlier Values can be Discovered from which Tools?

The various tools that can be used to discover outlier values are scatterplots, boxplots, Z-score, etc.

Machine Learning Interview Questions For Intermediate

18. What is Support Vector Machine (SVM) in Machine Learning?

SVM is a Machine Learning algorithm that is majorly used for classification. It is used on top of the high dimensionality of the characteristic vector.

The following is the code for SVM classifier:

```
1 # Introducing required libraries
2 from sklearn import datasets
3 from sklearn.metrics import confusion_matrix
4 from sklearn.model_selection import train_test_split
5 # Stacking the Iris dataset
6 iris = datasets.load_iris()
7 # A -> features and B -> label
8 A = iris.data
9 B = iris.target
10 # Breaking A and B into train and test data
11 A_train, A_test, B_train, B_test = train_test_split(A, B, random_state = 0)
12 # Training a linear SVM classifier
13 from sklearn.svm import SVC
14 svm_model_linear = SVC(kernel = 'linear', C = 1).fit(A_train, B_train)
15 svm_predictions = svm_model_linear.predict(A_test)
16 # Model accuracy for A_test
17 accuracy = svm_model_linear.score(A_test, B_test)
18 # Creating a confusion matrix
19 cm = confusion_matrix(B_test, svm_predictions)
```

19. What is Cross-validation in Machine Learning?

Cross-validation allows a system to increase the performance of the given Machine Learning algorithm, which is fed a number of sample data from the dataset. This sampling process is done to break the dataset into smaller parts that have the same number of rows, out of which a random part is selected as a test set and the rest of the parts are kept as train sets. Cross-validation consists of the following techniques:

- Holdout method
- K-fold cross-validation
- Stratified k-fold cross-validation
- Leave p-out cross-validation

20. What is Entropy in Machine Learning?

Entropy in Machine Learning measures the randomness in the data that needs to be processed. The more entropy in the given data, the more difficult it becomes to draw any useful conclusion from the data. For example, let us take the flipping of a coin. The result of this act is random as it does not favor heads or tails. Here, the result for any number of tosses cannot be predicted easily as there is no definite relationship between the action of flipping and the possible outcomes.

21. What is Epoch in Machine Learning?

Epoch in Machine Learning is used to indicate the count of passes in a given training dataset where the Machine Learning algorithm has done its job. Generally, when there is a large chunk of data, it is grouped into several batches. All these batches go through the given model, and this process is referred to as iteration. Now, if the batch size comprises the complete training dataset, then the count of iterations is the same as that of epochs.

In case there is more than one batch, $d * e = i * b$ is the formula used, wherein d is the dataset, e is the number of epochs, i is the number of iterations, and b is the batch size.

22. What are Type I and Type II Errors?

Type I Error: Type I Error, false positive, is an error where the outcome of a test shows the nonacceptance of a true condition.

For example, suppose a person gets diagnosed with depression even when they are not suffering from the same, it is a case of false positive.

Type II Error: Type II Error, false negative, is an error where the outcome of a test shows the acceptance of a false condition.

For example, the CT scan of a person shows that they do not have a disease but in fact they do have the disease. Here, the test accepts the false condition that the person does not have the disease. This is a case of false negative.

23. How to handle Missing or Corrupted Data in a Dataset?

In [Python pandas](#), there are two methods to locate lost or corrupted data and discard those values:

- `isNull()`: It can be used for detecting the missing values.
- `dropna()`: It can be used for removing columns or rows with null values.

`fillna()` can be used to fill the void values with placeholder values.

24. When to use mean and when to use median to handle a missing numeric value?

We choose the mean to impute missing values when the data distribution is normal and there are no significant outliers, as the mean is sensitive to both. In contrast, we use the median in cases of skewed distributions or when outliers are present, because the median is more robust to these factors and provides a better central tendency measure under these conditions.

25. Both being Tree-based Algorithms, how is Random Forest different from Gradient Boosting Machine (GBM)?

The main difference between a random forest and GBM is the use of techniques. [Random forest](#) advances predictions using a technique called bagging. On the other hand, GBM advances predictions with the help of a technique called boosting.

- **Bagging:** In bagging, we apply arbitrary sampling and we divide the dataset into N . After that, we build a model by employing a single training algorithm. Following that, we combine the final predictions by polling. Bagging helps to increase the efficiency of a model by decreasing the variance to eschew overfitting.

- Boosting: In boosting, the algorithm tries to review and correct the inadmissible predictions at the initial iteration. After that, the algorithm's sequence of iterations for correction continues until we get the desired prediction. Boosting assists in reducing bias and variance for strengthening the weak learners.

26. Differentiate between Sigmoid and Softmax Functions

Sigmoid and Softmax functions differ based on their usage in Machine Learning task classification. Sigmoid function is used in the case of binary classification, while Softmax function is used in case of multi-classification.

27. In Machine Learning, for how many classes can Logistic Regression be used?

Logistic regression cannot be used for more than two classes. Logistic regression is, by default, a binary classifier. However, in cases where multi-class classification problems need to be solved, the default number of classes can be extended, i.e., multinomial logistic regression.

28. What are the Two Main Types of Filtering in Machine Learning? Explain.

The two types of filtering are:

- Collaborative filtering
- Content-based filtering

Collaborative filtering refers to a recommender system where the interests of the individual user are matched with preferences of multiple users to predict new content.

Content-based filtering is a recommender system where the focus is only on the preferences of the individual user and not on multiple users.

29. What is meant by Ensemble Learning?

Ensemble learning refers to the combination of multiple Machine Learning models to create more powerful models. The primary techniques involved in ensemble learning are bagging and boosting.

30. What are the Various Kernels that are present in SVM?

The various kernels that are present in SVM are:

- Linear
- Polynomial
- Radial Basis
- Sigmoid

Machine Learning Interview Questions for Experienced

31. Suppose you found that your model is suffering from high variance. Which algorithm do you think could handle this situation and why?

Handling High Variance

- For handling issues of high variance, we should use the bagging algorithm.

- The bagging algorithm would split data into subgroups with a replicated sampling of random data.
- Once the algorithm splits the data, we can use random data to create rules using a particular training algorithm.
- After that, we can use polling for combining the predictions of the model.

32. What is Rescaling of Data and how is it done?

In real-world scenarios, the attributes present in data are in a varying pattern. So, rescaling the characteristics to a common scale is beneficial for algorithms to process data efficiently.

We can rescale data using Scikit-learn. The code for rescaling the data using MinMaxScaler is as follows:

```
1 #Rescaling data
2 import pandas
3 import scipy
4 import numpy
5 from sklearn.preprocessing import MinMaxScaler
6 names = ['Abhi', 'Piyush', 'Pranay', 'Sourav', 'Sid', 'Mike', 'pedi', 'Jack', 'Tim']
7 Dataframe = pandas.read_csv(url, names=names)
8 Array = dataframe.values
9 # Splitting the array into input and output
10 X = array[:,0:8]
11 Y = array[:,8]
12 Scaler = MinMaxScaler(feature_range=(0, 1))
13 rescaledX = scaler.fit_transform(X)
14 # Summarizing the modified data
15 numpy.set_printoptions(precision=3)
16 print(rescaledX[0:5,:])
```

Apart from the theoretical concepts, some interviewers also focus on the implementation of Machine Learning topics. The following Interview Questions are related to the implementation of theoretical concepts.

33. What is the difference between Standard scalar and MinMax Scaler?

StandardScaler and Min-Max scaling are two common data preprocessing techniques used in machine learning. The key differences are:

1. StandardScaler (Z-score normalization):
 - Scales data to have a mean of 0 and a standard deviation of 1.
 - Suitable for algorithms assuming normal distribution and is robust to outliers.
2. Min-Max Scaling:
 - Scales data to a specific range, often between 0 and 1.
 - Useful for models sensitive to feature magnitudes, but can be influenced by outliers.

34. What is Binarizing of Data? How to Binarize?

Converting data into binary values on the basis of threshold values is known as binarizing of data. The values that are less than the threshold are set to 0 and the values that are greater than the threshold are set to 1. This process is useful when feature engineering has to be performed. This can also be used for adding unique features. Data can be binarized using Scikit-learn. The code for binarizing data using Binarizer is as follows:

```
1 from sklearn.preprocessing import Binarizer
2 import pandas
3 import numpy
4 names = ['Abhi', 'Piyush', 'Pranay', 'Sourav', 'Sid', 'Mike', 'pedi', 'Jack', 'Tim']
5 dataframe = pandas.read_csv(url, names=names)
6 array = dataframe.values
7 # Splitting the array into input and output
8 X = array[:,0:8]
9 Y = array[:,8]
10 binarizer = Binarizer(threshold=0.0).fit(X)
11 binaryX = binarizer.transform(X)
12 # Summarizing the modified data
13 numpy.set_printoptions(precision=3)
14 print(binaryX[0:5,:])
```

35. How to Standardize Data?

Standardization is the method that is used for rescaling data attributes. The attributes are likely to have a mean value of 0 and a value of the standard deviation of 1. The main objective of standardization is to prompt the mean and standard deviation for the attributes.

Data can be standardized using [Scikit-learn](#). The code for standardizing the data using StandardScaler is as follows:

```
1 # Python code to Standardize data (0 mean, 1 stdev)
2 from sklearn.preprocessing import StandardScaler
3 import pandas
4 import numpy
5 names = ['Abhi', 'Piyush', 'Pranay', 'Sourav', 'Sid', 'Mike', 'pedi', 'Jack', 'Tim']
6 dataframe = pandas.read_csv(url, names=names)
7 array = dataframe.values
8 # Separate the array into input and output components
9 X = array[:,0:8]
10 Y = array[:,8]
11 scaler = StandardScaler().fit(X)
12 rescaledX = scaler.transform(X)
13 # Summarize the transformed data
14 numpy.set_printoptions(precision=3)
15 print(rescaledX[0:5,:])
```

36. We know that one-hot encoding increases the dimensionality of a dataset, but label encoding doesn't. How?

When one-hot encoding is used, there is an increase in the dimensionality of a dataset. The reason for the increase in dimensionality is that every class in categorical variables, forms a different variable.

Example: Suppose there is a variable "Color." It has three sublevels, "Yellow," "Purple," and "Orange." So, one-hot encoding "Color" will create three different variables as Color.Yellow, Color.Purple, and Color.Orange.

In label encoding, the subclasses of a certain variable get the value 0 and 1. So, label encoding is only used for binary variables.

This is why one-hot encoding increases the dimensionality of data and label encoding does not.

Now, if you are interested in doing an end-to-end certification course in Machine Learning, you can check out IntelliPaat's [Machine Learning Course with Python](#).

37. Executing a binary classification tree algorithm is a simple task. But how does tree splitting take place? How does the tree determine which variable to break at the root node and which at its child nodes?

Gini index and Node Entropy assist the binary classification tree to make decisions. Basically, the tree algorithm determines the feasible feature that is used to distribute data into the most genuine child nodes.

According to the Gini index, if we arbitrarily pick a pair of objects from a group, then they should be of identical class and the probability for this event should be 1.

The following are the steps to compute the Gini index:

1. Compute Gini for sub-nodes with the formula: The sum of the square of probability for success and failure ($p^2 + q^2$)
2. Compute Gini for split by weighted Gini rate of every node of the split

Now, Entropy is the degree of indecency that is given by the following:

Where a and b are the probabilities of success and failure of the node

When Entropy = 0, the node is homogenous

When Entropy is high, both groups are present at 50–50 percent in the node.

Finally, to determine the suitability of the node as a root node, the entropy should be very low.

38. Imagine you are given a dataset consisting of variables having more than 30% missing values. Let's say, out of 50 variables, 16 variables have missing

values, which is higher than 30%. How will you deal with them?

To deal with the missing values, we will do the following:

- We will specify a different class for the missing values.
- Now, we will check the distribution of values, and we will hold those missing values that are defining a pattern.
- Then, we will charge these values into yet another class while eliminating others.

39. What is F1-score and How Is It Used?

F-score or F1-score is a measure of overall accuracy of a binary classification model. Before understanding F1-score, it is crucial to understand two more measures of accuracy, i.e., precision and recall.

Precision is defined as the percentage of True Positives to the total number of positive classifications predicted by the model. In other words,

$$\text{Precision} = (\text{No. of True Positives} / \text{No. True Positives} + \text{No. of False Positives})$$

Recall is defined as the percentage of True Positives to the total number of actual positive labeled data passed to the model. In other words,

$$\text{Precision} = (\text{No. of True Positives} / \text{No. True Positives} + \text{No. of False Negatives})$$

Both precision and recall are partial measures of accuracy of a model. F1-score combines precision and recall and provides an overall score to measure a model's accuracy.

$$\text{F1-score} = 2 \times (\text{Precision} \times \text{Recall}) / (\text{Precision} + \text{Recall})$$

This is why, F1-score is the most popular measure of accuracy in any Machine-Learning-based binary classification model.

40. How to Implement the KNN Classification Algorithm?

Iris dataset is used for implementing the KNN classification algorithm.

```
1 # KNN classification algorithm
2 from sklearn.datasets import load_iris
3 from sklearn.neighbors import KNeighborsClassifier
4 import numpy as np
5 from sklearn.model_selection import train_test_split
6 iris_dataset=load_iris()
7 A_train, A_test, B_train, B_test = train_test_split(iris_dataset["data"], iris_dataset["target"], random_state=0)
8 kn = KNeighborsClassifier(n_neighbors=1)
9 kn.fit(A_train, B_train)
10 A_new = np.array([[8, 2.5, 1, 1.2]])
11 prediction = kn.predict(A_new)
12 print("Predicted target value: {}".format(prediction))
13 print("Predicted feature name: {}".format
14 (iris_dataset["target_names"][prediction]))
15 print("Test score: {:.2f}".format(kn.score(A_test, B_test)))
16 Output:
17 Predicted Target Name: [0]
18 Predicted Feature Name: ['Setosa']
19 Test Score: 0.92
```

Come to IntelliPaat's [Machine Learning Community](#) if you have more queries on Machine Learning Interview Questions!

Role-Specific Machine Learning Questions

41. How come logistic regression is labeled as a regression method when it is primarily used for classification tasks?

Logistic regression earns its classification primarily due to its historical connection with linear regression. However, its paramount utility lies in addressing classification tasks, given its remarkable ability to model the probability of an observation belonging to a specific class or category. In practice, it quantifies the likelihood of an event's occurrence, endowing it with great significance in tackling classification challenges, such as discerning spam emails or making medical

diagnoses. Thus, despite its nomenclature as “regression,” its predominant function is in terms of classification, which explains its frequent association with classification algorithms. This explanation is intended to be informative, ensuring originality and search engine optimization.

42. What is Overfitting in Machine Learning and how can it be avoided?

Overfitting happens when a machine has an inadequate dataset and tries to learn from it. So, overfitting is inversely proportional to the amount of data.

For small databases, overfitting can be bypassed by the cross-validation method. In this approach, a dataset is divided into two sections. These two sections will comprise the testing and training dataset. To train a model, the training dataset is used, and for testing the model for new inputs, the testing dataset is used. This is how to avoid overfitting.

43. What is Hypothesis in Machine Learning?

Machine Learning allows the use of available dataset to understand a specific function that maps input to output in the best possible way. This problem is known as function approximation. Here, approximation needs to be used for the unknown target function that maps all plausible observations based on the given problem in the best manner. Hypothesis in Machine learning is a model that helps in approximating the target function and performing the necessary input-to-output mappings. The choice and configuration of algorithms allow defining the space of plausible hypotheses that may be represented by a model.

In the hypothesis, lowercase h (h) is used for a specific hypothesis, while uppercase h (H) is used for the hypothesis space that is being searched. Let us briefly understand these notations:

- Hypothesis (h): A **hypothesis** is a specific model that helps in mapping input to output; the mapping can further be used for evaluation and prediction.
- Hypothesis set (H): Hypothesis set consists of a space of hypotheses that can be used to map inputs to outputs, which can be searched. The general constraints include the choice of problem framing, the model, and the model configuration.

44. How is the suitability of a Machine Learning Algorithm determined for a particular problem?

To identify a **Machine Learning Algorithm** for a particular problem, the following steps should be followed:

Step 1: Problem classification: Classification of the problem depends on the classification of input and output:

- Classifying the input: Classification of the input depends on whether there is data labeled (supervised learning) or unlabeled (unsupervised learning), or whether a model has to be created that interacts with the environment and improves itself (reinforcement learning.)
- Classifying the output: If the output of a model is required as a class, then some classification techniques need to be used.

If the output is a number, then regression techniques must be used; if the output is a different cluster of inputs, then clustering techniques should be used.

Step 2: Checking the algorithms in hand: After classifying the problem, the available algorithms that can be deployed for solving the classified problem should be considered.

Step 3: Implementing the algorithms: If there are multiple algorithms available, then all of them are to be implemented. Finally, the algorithm that gives the best performance is selected.

45. What is the Variance Inflation Factor?

Variance inflation factor (VIF) is the estimate of the volume of multicollinearity in a collection of many regression variables.

$VIF = \text{Variance of the model} / \text{Variance of the model with a single independent variable}$

This ratio has to be calculated for every independent variable. If VIF is high, then it shows the high collinearity of the independent variables.

Courses you may like



Advanced Certification in Data Science & Artificial Intelligence

- Learn from IIT Madras Faculty & Industry Experts
- #1 in NIRF 2020 Ranking

[Enroll Now](#)



PG Certification in Data Science & Machine Learning

- Live Classes from MNIT Faculty & Industry Experts
- #35 in NIRF 2020 Ranking

[Enroll Now](#)



Data Science Master Program

Co-created with **IBM**

47 Projects 200 Hours 12 Courses

[ENROLL NOW](#)

46. When should Classification be used over Regression?

Both classification and regression are associated with prediction. Classification involves the identification of values or entities that lie in a specific group. Regression entails predicting a response value from consecutive sets of outcomes.

Classification is chosen over regression when the output of the model needs to yield the belongingness of data points in a dataset to a particular category.

For example, If you want to predict the price of a house, you should use regression since it is a numerical variable. However, if you are trying to predict whether a house situated in a particular area is going to be high-, medium-, or low-priced, then a classification model should be used.

47. Why is rotation required in PCA? What will happen if the components are not rotated?

Rotation is a significant step in principal component analysis (PCA.) Rotation maximizes the separation within the variance obtained by the components. This makes the interpretation of the components easier.


The motive behind conducting PCA is to choose fewer components that can explain the greatest variance in a dataset. When rotation is performed, the original coordinates of the points get changed. However, there is no change in the relative position of the components.

If the components are not rotated, then there needs to be more extended components to describe the variance.

48. What is ROC Curve and what does it represent?

ROC stands for receiver operating characteristic. [ROC Curve](#) is used to graphically represent the trade-off between true and false-positive rates.

In ROC, the area under the curve (AUC) gives an idea about the accuracy of the model.



The above graph shows a ROC curve. The greater the AUC, the better the performance of the model.

Next, we will be taking a look at Machine Learning interview questions on rescaling, binarizing, and standardizing.

49. What do you understand about the P-value?

P-value is used in decision-making while testing a hypothesis. The null hypothesis is rejected at the minimum significance level of the P-value. A lower P-value indicates that the null hypothesis is to be rejected.

50. What is meant by Correlation and Covariance?

Correlation is a mathematical concept used in statistics and probability theory to measure, estimate, and compare data samples taken from different populations. In simpler terms, correlation helps in establishing a quantitative relationship between two variables.

Covariance is also a mathematical concept; it is a simpler way to arrive at a correlation between two variables. Covariance basically helps in determining what change or affect does one variable has on another.

51. What are the Various Tests for Checking the Normality of a Dataset?

In Machine Learning, checking the normality of a dataset is very important. Hence, certain tests are performed on a dataset to check its normality. Some of them are:

- D'Agostino Skewness Test
- Shapiro-Wilk Test
- Anderson-Darling Test
- Jarque-Bera Test
- Kolmogorov-Smirnov Test

52. Explain False Negative, False Positive, True Negative, and True Positive with a simple example.

True Positive (TP): When the Machine Learning model correctly predicts the condition, it is said to have a True Positive value.

True Negative (TN): When the Machine Learning model correctly predicts the negative condition or class, then it is said to have a True Negative value.

False Positive (FP): When the Machine Learning model incorrectly predicts a negative class or condition, then it is said to have a False Positive value.

False Negative (FN): When the Machine Learning model incorrectly predicts a positive class or condition, then it is said to have a False Negative value.

53. What do you mean by the term Overfitting, and How can you avoid It?

Overfitting is a situation when the model learns too well from the training data set but when set to perform in some unknown data, results in low accuracy.

To avoid this situation, we make use of:

1. Regularization
2. Making a simple model
3. Making use of cross-validation methods

54. What are the 'training set' and 'test sets'? How much data will you allocate for your training, validation, and test sets?

The training set is the dataset on which you will train your machine-learning model. A test set is used to test the model if it can perform on an unknown set of data or not.

Usually, we make a 70:30 split of the existing dataset as a training and test dataset. For example, if we have 100 records, then 70 random records from the dataset will be used to train the model, while 30 random records will be used to test the model.

55. What are the three stages of building a model in machine learning?

The three stages of building the machine learning model are:

1. Development
2. Testing
3. Deployment

56. How will you know which machine learning algorithm to choose for your classification problem?

There are no fixed rules for choosing a machine learning algorithm for a classification problem. However, to reduce the number of algorithms, we can use the following guidelines:

1. For small training datasets, use a model with high bias and low variance.
2. For large training datasets, use a model with high variance and low bias

Lastly, if accuracy is something that you are looking for, then you have to individually test the models.

57. Define precision and recall.

Precision = True Positive / True Positive + False Positive

Recall = True Positive / True Positive + False Negative

58. What do you mean by the term Kernel SVM?

Kernel methods are a class of algorithms that are mostly used for problem statements like pattern analysis. It is used for solving both classification and regression problems. Kernel SVM is just an abbreviated form of Kernel Support Vector Machine. It is one of the most common ones in the Kernel method list.

59. What do you understand by the F1 score?

It is an evaluation matrix for a classification model. It combines both precision and recall.

$$F1 = 2 * (P * R) / (P + R)$$

FAANG Machine Learning Engineer Questions

60. How is Adam Optimizer different from Rmsprop?

Adam (short for Adaptive Moment Estimation) and RMSprop (Root Mean Square Propagation) are optimization algorithms used to train neural networks. The differences between them are:

Adam Optimizer	Rmsprop Optimizer
For every parameter, Adam keeps track of two moving averages: the mean (first moment) and the uncentered variance (second moment).	RMSprop also uses moving averages; it only maintains a running average of squared gradients for each parameter.
Adam blends the ideas of adaptive learning rates with momentum.	RMSprop adapts the learning rate for each parameter based on the magnitude of the recent gradients.
Adam performs bias correction for the moving averages.	RMSprop does not typically use bias correction.

61. What are the different types of activation functions and explain the vanishing gradient problem?

Activation functions are functions used in a neural network to compute the weighted sum of inputs and biases, which decides whether a neuron can be activated or not.

There are multiple types of activation functions present, each with its characteristics. A few are listed below:

- Sigmoid Function:
 - Output values between 0 and 1.
 - Commonly used in the output layer of binary classification models.

- Hyperbolic Tangent Function (tanh):
 - Output values between -1 and 1.
 - Similar to the sigmoid, but it has a wider output range.
- Rectified Linear Unit (ReLU):
 - Outputs the input for positive values, zero otherwise.
 - Simple and computationally efficient, commonly used in hidden layers.
- Leaky ReLU:
 - Similar to ReLU but allows a small, non-zero gradient for negative values (α is a small positive constant).
- Parametric ReLU (PReLU):
 - Similar to leaky ReLU, the negative slope (α) is learned during training.
- Exponential Linear Unit (ELU):
 - Smooth for negative values, allowing for improved learning.

62. Explain the bias-variance tradeoff.

The concept of the bias-variance tradeoff describes the tradeoff between two sources of error, bias and variance. This might also affect the performance of the predictive model you are building.

The model can be visualized as having three different parts :

- High Bias, Low Variance:
 - In this case, understand that the model needs to be more complex to capture true data. It consistently makes the same errors across all training sets.
 - It results in underfitting.
- Low Bias, High Variance:
 - In this case, understand that the model needs to be simpler and fit the data too closely. It performs well on the training set but poorly on new data.
 - It results in overfitting.

- Balanced Bias-Variance:
 - In this situation, you understand that the model has achieved a good balance between bias and variance, capturing the true data without being overly sensitive to noise.

63. What does the “minus” in cross-entropy mean?

The “minus” in cross-entropy is used to define it as a loss function, where the higher the number, the worse the model is, while the lower the number, the better the model is. The goal is to minimize this loss during the training of a model to improve the model's predictive accuracy.

64. What do L1 and L2 regularization mean and when would you use L1 vs. L2? Can you use them both?

L1 Regularization, adds the absolute magnitude of the coefficient as a penalty to the loss function.

L2 Regularization, adds the squared magnitude of the coefficient as the penalty to the loss function.

The choice between L1 and L2 depends on your modeling goal and the data present. L1 is used when you suspect that many features are irrelevant, and you want a simple model with feature selection. L2 is used when all features are relevant, and you want to control the magnitudes of the weights to prevent them from becoming too large.

Yes, you can use both L1 and L2, also known as Elastic Net. It can be a good choice when you want a combination of L1 and L2 regularization, providing a tradeoff between sparsity and weight shrinkage.

65. What is an activation function in machine learning?

In layman's terms, an activation function defines if the neuron should be activated or not. The [activation function](#) helps the neural network define the important data points to be activated while ignoring the irrelevant ones.

According to the definition, "Activation functions are functions used in a neural network to compute the weighted sum of inputs and biases, which decides whether a neuron can be activated or not."

66. What do eigenvalues and eigenvectors mean in PCA?

In Principal Component Analysis (PCA), eigenvalues represent the amount of information that a given principal component can explain.

Eigenvectors represent the weight of each eigenvalue.

How to Prepare for the Machine Learning Interview

You have to go through many rounds of interviews in every company! Following are some of the interview rounds that you will be subjected to:

1. On-call Assessment Round
2. Technical Assessment Round
3. Machine Learning Theory Round
4. Machine Learning System Design Round
5. Case Study Round
6. Behavioral Round

67. How do you prepare for the on-call assessment round at top companies?

The on-call assessment round is the first round of the interview process, where you usually have a short and simple discussion with the hiring manager or HRs about the job role. They ask you about some basic information like your experience, CTC, notice period, etc., and identify if your skills match their requirements. Meanwhile, they will share details about your job, like what projects you will be working on, what tech stacks are, whether it is an on-site job or not, etc. Usually, it doesn't include any technical questions.

68. How do you prepare for the technical assessment round in top companies?

In the technical assessment round, you receive an email from the recruiter with a test link. The test comprises either some hands-on activities or a simple questionnaire. It usually covers a few analytical, problem-solving questions to help them understand your skills in the domain.

To be prepared, you should have a basic understanding of machine learning. Revise all your notes, go through your cheat sheets, and solve some analytical questions.

69. How do you prepare for the machine learning theory round in top companies?

The machine learning theory round is an on-site or virtual round where you interact with your hiring manager and answer questions based on machine learning algorithms. It is more of a round to test your water. The round can last anywhere from 45 minutes to an hour. You can expect straightforward theory questions as well as derivation-based questions.

To prepare, you have to take time and brush up on your algorithmic skills along with the derivation-based questions. In this round, your communication skills and interpersonal skills also play a major part, so make sure you work on that as well.

70. How do you prepare for the machine learning system design round in top companies?

In the machine learning system design round, the interviewer will ask you to design some systems like a recommendation system, a contact ranking system, etc. You just need to design an overview. It's important to understand the requirements and feasibility to answer better, so always ask for clarification.

To prepare, you should have a clear concept of the model you are designing. Make sure to make a script or a schema in which you will answer the question. Most importantly, know the requirements and the feasibility.

71. How do you prepare for the case study round in top companies?

The case study round is mostly the last technical round. There are two phases for a case study round: a discussion on one of the most interesting projects you have mentioned in your resume, or a discussion with the interviewer on a problem statement. The interviewer wants to test your skills, like how you approach a new problem statement, the problems you might face, etc.

To prepare for this round, make sure you have an overall understanding of the projects that are listed in your resume. Prepare all the important points of your projects.

72. How do you prepare for the behavioral round in top companies?

After all the technical rounds, there is a last round before your HR round (in many companies, they merge both), the behavioral round. In this round, you will be having a conversation with the HRs on some situation-based questions. This will help them understand if you fit into their company culture.

73. Few tips to follow throughout the interview process.

A few tips to follow throughout the interview process are :

- Be confident in all the rounds, and work on your communication and interpersonal skills.
- Read the job description carefully.
- Answer questions in a structured manner.
- Be on time for the interview. It will help you build a good impression on the interviewer.

Machine Learning Salary Based on Experience

The average salary for an entry-level machine learning engineer is ₹12,32,000 per year in India and \$1,52,360 per year in the United States. The average additional cash compensation for a machine learning engineer in India is ₹1,32,000, with a range from ₹55,000 – ₹2,50,000 in India and \$26,243, with a range from \$19,682 – \$36,741 in the United States.

Job Role	Experience	Salary Range
Machine Learning Engineer	0 – 2 years	₹08L – ₹13L /yr
Senior Machine Learning Engineer	2 – 4 years	₹14L – ₹17L /yr
Lead Machine Learning Engineer	5 – 7 years	₹14L – ₹37L /yr
Principal Machine Learning Engineer	8+ years	₹30L – ₹47L /yr

Machine Learning Trends in 2024

1. Global Demand: According to [LinkedIn](#), there are currently more than 80000+ open positions for a machine learning engineer in the United States.
2. Projected Growth: As per the [Future of Jobs Report 2023](#), there is a very high demand for machine learning engineers, and It is expected to grow by 40%, or 1 million jobs per year, globally.
3. Regional Trends: According to [LinkedIn](#), there are currently more than 18000+ open positions for a machine learning engineer in India. And the hiring trend is set to increase by 8.3% in 2024.

Job Opportunities in Machine Learning

Multiple job roles in the Industry require machine learning. Here are a few of them:

Job Role	Description
Machine Learning Engineer	They are responsible for designing, building, and developing machine learning models.
Machine Learning Developer	They are responsible for building and implementing machine learning models and algorithms into various applications.

Artificial Intelligence Engineer	They are responsible for building and developing AI models, which can include machine learning models, NLP-based models, and computer vision models.
Computer Vision Engineer	They specifically work in the field of computer vision, which involves interpreting visual data by the computer.
Research Scientist	They research and develop new machine-learning models and algorithms.

Roles and Responsibilities of a Machine Learning Engineer

A machine learning engineer is responsible for creating deep learning and machine learning models, implementing the right machine learning algorithms into practice, and conducting experiments and tests to check their accuracy for the given problem statement.

According to a job description posted by [Siemens](#) on LinkedIn:

Job Role: [AI /ML Engineer](#)

Responsibilities:

- Knowledge of various data mining and machine learning techniques to extract valuable insights from large datasets and communicate the findings to stakeholders.
- Should know how to evaluate the effectiveness of models and algorithms using statistical tests.
- Should understand the basic deployment process using DevOps and any of the public cloud services (AWS, Azure, or GCP).

Technical Skills:

- Strong programming skills in languages such as Python, R, or C++.
- Should be familiar with libraries and concepts such as Tensorflow, Pytorch, Keras, Sklearn, Statmodels, Pandas, Numpy, Scipy, OpenCV, PIL, SkImage, SQL, HQL, or similar.
- Proficient in the use of computer vision tools
- Excellent verbal, written, and presentation skills.

Conclusion

I hope this set of interview questions on Machine Learning will help you prepare for your interviews. Best of luck!

Looking to start your career or even elevate your skills in the field of machine learning? You can enroll in our [comprehensive Machine Learning course](#) or enroll in the [Executive Post Graduate Certification in AI and Machine Learning from IIT Roorkee in collaboration with IBM and Microsoft](#) with IntelliPaat and get certified today.

If you want to deep dive into more machine learning interview questions, feel free to join [IntelliPaat's vibrant Machine Learning Community](#) and get answers to your queries from like-minded enthusiasts.