



TensorFlow Interview Questions

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

TensorFlow is a popular library that is used to design Machine Learning and Deep Learning models easily. Being open-source and powerful, it is used in a variety of domains. Since Deep Learning and Artificial Intelligence form to be among the top career paths of the 21st century, understanding TensorFlow and acing TensorFlow interviews become very important. These Top TensorFlow interview questions and answers will help you approach the questions in a structured manner and help you understand how best to answer them.

This Top TensorFlow Interview Questions blog is divided into three segments as shown below:

[1. Basic](#)

[2. Intermediate](#)

[3. Advanced](#)

Basic Interview Questions

1. What is TensorFlow?

TensorFlow is the world's most used library for Machine Learning. Developed in 2015 by the Google Brain Team, it ensures to provide an easy-to-use low-level toolkit that can handle complex mathematical operations and learning architectures.

2. What are tensors?

Tensors are similar to arrays in programming languages, but here, they are of higher dimensions. It can be considered as a generalization of matrices that form an n-dimensional array. TensorFlow provides methods that can be used to create tensor functions and compute their derivatives easily. This is what sets tensors apart from the NumPy arrays.

3. What is the meaning of TensorBoard?

TensorBoard is a Graphical User Interface (GUI) that is provided by TensorFlow to help users visualize graphs, plots, and other metrics easily without having to write a lot of code. TensorBoard provides an ample number of advantages in terms of readability, ease of use, and performance metrics.

4. What are some of the advantages of using TensorFlow?

TensorFlow has numerous advantages, and this is why it is the most used framework for Machine Learning in the world. Some of its advantages are given below:

- Platform independency
- Usage of GPU for distributed computing
- Auto-differentiation capability
- Open-source and large community
- Highly customizable based on requirements
- Support for asynchronous computations

5. Are there any limitations to using TensorFlow?

Even though TensorFlow provides numerous advantages, it has one or two caveats in the current versions:

- No support for OpenCL (Open Computing Language) yet
- GPU memory conflicts when used with Theano
- Can be overwhelming for beginners to get started

Next up on these TensorFlow coding interview questions, let us check out about the types of tensors.

6. What are the types of tensors available in TensorFlow?

There are three main types of tensors:

- Constant tensors
- Variable tensors
- Placeholder tensors

7. How can data be loaded into TensorFlow?

There are two ways that you can use to load data into TensorFlow before training [Machine Learning algorithms](#):

- Data load into memory: Here, the data is loaded into the memory as a single array unit. It is the easiest way to load the data.
- TensorFlow data pipeline: It is making use of the built-in APIs to load the data and feed it across to the algorithm.

8. What is the simple working of an algorithm in TensorFlow?

There are five main steps that govern the working of the majority of algorithms in [TensorFlow](#). They are as follows:

1. Data import or data generation, alongside setting up a data pipeline
2. Data input through computational graphs
3. Generation of the loss function to evaluate the output
4. Backpropagation to modify the data
5. Iterating until output criteria are met

9. What are the methods that can be used to handle overfitting in TensorFlow?

There are three methods, which can be used to easily handle the condition of overfitting when using TensorFlow:

- Batch normalization
- Regularization technique
- Dropouts

10. What are the languages that are supported in TensorFlow?

TensorFlow supports a wide variety of languages for programmers to write the code in. The preferred language presently is Python.

However, experimental support is being implemented for other languages, such as Go, Java, and C++. Also, language bindings for Ruby, Scala, and Julia are being developed by the open-source community.

11. What are placeholder tensors?

Placeholder tensors are entities that provide an advantage over a regular variable. It is used to assign data at a later point in time.

Placeholders can be used to build graphs without any prior data being present. This means that they do not require any sort of initialization for usage.

12. What are managers in TensorFlow?

TensorFlow managers are entities that are responsible for handling the following activities for servable objects:

- Loading
- Unloading
- Lookup
- Lifetime management



13. Where is TensorFlow mostly used?

TensorFlow is used in all of the domains that cover Machine Learning and Deep Learning. Being the most essential tool, the following are some of the main use cases of TensorFlow:

- Time series analysis
- Image recognition
- Voice recognition
- Video upscaling
- Test-based applications

14. What are TensorFlow servables?

Servables in TensorFlow are simply the objects that client machines use to perform computations. The size of these objects is flexible. Servables can include a variety of information like any entity from a lookup table to a tuple needed for inference models.

15. How does the Python API work with TensorFlow?

Python is the primary language when it comes to working with TensorFlow. TensorFlow provides ample number of functionalities when used with the API, such as:

- Automatic checkpoints
- Automatic logging

- Simple training distribution
- Queue-runner design methods

16. What are some of the APIs outside of the TensorFlow project?

Following are some of the APIs developed by Machine Learning enthusiasts across the globe:

- TFLearn: A popular Python package
- TensorLayer: For layering architecture support
- Pretty Tensor: Google's project providing a chaining interface
- Sonnet: Provides a modular approach to programming

17. What are TensorFlow loaders?

Loaders are used in TensorFlow to load, unload, and work with servable objects. The loaders are primarily used to add algorithms and data into TensorFlow for working.

The load() function is used to pre-load a model from a saved entity easily.

18. What makes TensorFlow advantageous over other libraries?

Following are some of the benefits of TensorFlow over other libraries:

- Pipelines: data is used to build efficient pipelines for text and image processing.
- Debugging: tfdbg is used to track the state and structure of objects for easy debugging.
- Visualization: TensorBoard provides an elegant user interface for users to visualize graphs.

- Scalability: It can scale Deep Learning applications and their associated infrastructure easily.

19. What are TensorFlow abstractions?

TensorFlow contains certain libraries used for abstraction such as Keras and TF-Slim. They are used to provide high-level access to data and model life cycle for programmers using TensorFlow. This can help them easily maintain clean code and also reduce the length of the code exponentially.

Next up on this top TensorFlow interview questions and answers post, we will take a look at the intermediate set of questions.

Intermediate Interview Questions

20. What are the differences between `tf.variable` and `tf.placeholder` in TensorFlow?

<code>tf.variable</code>	<code>tf.placeholder</code>
Defines values for variables that change with time	Defines inputs that do not change with time
Requires initialization when defined	Does not require initialization during defining

21. What is a graph explorer in TensorFlow?

A graph explorer is used to visualize a graph on TensorBoard. It is also used for the inspection operations of a model in TensorFlow. To easily understand the flow in a graph, it is recommended to use a graph visualizer in TensorBoard.

Next up on these TensorFlow coding interview questions, let us check out variables and their lifetimes.

22. How is variable lifetime tracked in TensorFlow?

The lifetime of a variable is automatically tracked after its initialization, using the `tf.Variable.initializer` operation.

Later, after the usage, the session can be closed and the variable can be destroyed, using the `tf.Session.close` operation.

23. What are the types of dashboards supported by TensorFlow?

TensorFlow supports a variety of dashboards that are used to perform numerous tasks on TensorBoard easily:

- Scalar dashboard
- Image dashboard
- Graph dashboard
- Text dashboard
- Distributer dashboard
- Histogram dashboard

24. Can TensorFlow be deployed onto a container architecture?

Yes, TensorFlow can be easily used with containerization tools like Docker. The containerization tools alongside TensorFlow are mostly used to deploy various models that require text classification using [convolutional neural networks](#).

If you are looking forward to becoming an expert in TensorFlow and AI, make sure to check out [Intellipaat's AI Engineer Course](#).

25. Differentiate between TensorFlow and PyTorch.

TensorFlow	PyTorch
Developed by Google	Developed by Facebook
No support for runtime graph operations	Provides computational graph operations at runtime
Offers TensorBoard for visualization	No visualization tools offered in the bundle
Based on the Theano library	Based on the Torch library

26. Is word embedding supported in TensorFlow?

Yes, word embedding is supported in TensorFlow. It is widely used in the field of [Natural Language Processing](#). When TensorFlow is being used, it is called Word2vec.

Two models are used for word embedding in TensorFlow:

- The Continuous Bag of Words model
- The Skip-Gram model

Next, it is vital that you understand the use of estimators and that is exactly what we will look at in these TensorFlow coding interview questions.

27. What is the use of estimators in TensorFlow?

Estimators in TensorFlow are high-level APIs used to provide a high amount of code reusability when training a model. They can also override the default behavior of any aspect of the model.

There are two ways of the model building using estimators:

- Premade estimator: Used to create a specific model like DNNClassifier
- Base class estimator: Used to control a model using a model_fn function

28. What statistical distribution functions are provided by TensorFlow?

Numerous statistical distribution functions are offered by TensorFlow. They are all located inside the `tf.contrib.distributions` package.

The distributions supported are:

- Beta
- Bernoulli
- Chi2
- Dirichlet
- Gamma
- Uniform

29. Can you use TensorBoard without installing TensorFlow?

If TensorFlow is not installed, users can still make use of TensorBoard (versions above 1.14) in a standalone mode with redacted features.

Following plugins are supported:

- Scalars
- Image
- Audio
- Graph
- Projector
- Histograms
- Mesh



30. What is the meaning of the embedding projector in TensorFlow?

Embedding projector is an entity in TensorFlow that is used to easily visualize high-dimensional data.

It is used to read the data from the model checkpoint file prior to visualization and to view the input data after it has been embedded into a high-dimensional space by the model.

31. What are the differences between CNN and RNN?

Convolutional Neural Network (CNN)	Recurrent Neural Network (RNN)
Used to handle image data	Best suited to handle sequential data

Fixed input and output data types	Flexible input and output data lengths
Ideal for image and video processing	Ideal for speech and text analysis
Efficient and powerful when compared to RNN	Provides less number of feature sets

32. What is the difference between Type 1 and Type 2 errors?

In simple terms, Type 1 errors refer to the occurrence of a false positive outcome, and Type 2 errors denote the occurrence of a false negative value when performing complex computations.

33. When using TensorFlow, is performance always preferred over accuracy?

No, performance is not always preferred over accuracy when you use TensorFlow. This completely depends on the type of requirement and what the model is trying to achieve. The general rule of thumb is to provide equal weightage to model accuracy and performance.

The next set of TensorFlow interview questions will show the importance of using an example along with concepts to explain.

34. Can you give an example to create a tensor using the constant() function in TensorFlow?

Tensors are most commonly created using the constant() function. The values to be input into the tensor are given as arguments as shown below:

import tensorflow as tf

```
1 | t1 = tf.constant([[1, 2, 3], [4, 5, 6], [7, 8, 9]])
2 | t2 = tf.constant(["String One", "String Two", "String Three"])
3 | sess = tf.Session()
4 | print(t1)
5 | print(sess.run(t1))
6 | print("\n")
7 | print(t2)
8 | print(sess.run(t2))
```

Next up on this top TensorFlow interview questions and answers post, we will take a look at the advanced set of questions.

Advanced Interview Questions

35. What are some of the products that are built using TensorFlow?

There are many products that are built completely using TensorFlow. Some of them are as follows:

- Teachable Machine
- Handwriting Recognition
- Giorgio Cam
- NSynth

36. What is the meaning of Deep Speech?

Deep Speech is a speech-to-text engine that is open-source and uses TensorFlow. It is trained based on Machine Learning techniques and uses a simple syntax to process speech from an input to produce textual output on the other end.

The following syntax can be used to view all of the CLI options for Deep Speech:

```
1 | ./deepspeech.py
```

37. What is the use of a histogram dashboard in TensorFlow?

Histogram dashboards are widely used to display complex statistical distributions of a tensor in a simple way. Every histogram chart will have a slice of data that denotes the data that the tensor has at the point of representation.

38. How is audio stored in the audio dashboard?

The audio dashboard serves to primarily help users embed playable widgets stored in files. Tf.summary.audio is used for the storage of these files, and the tagging system is used to embed the latest audio based on the storage policies.

39. What are some of the components needed to deploy a Lite model file?

In TensorFlow, three main components are used to deploy a Lite model:

1. Java API: Used as a wrapper around the C++ API for Android
2. C++ API: Used to load the TensorFlow Lite model and call the interpreter
3. Interpreter: Used to handle kernel loading and the execution of the model

40. What is TensorFlow JS?

TensorFlow JS is a library that gives users the functionality of using browsers to run Machine Learning models. High-level APIs work with JavaScript to support a variety

of entities in the backend, such as WebGL, to use a GPU to render functionality (if available). Models can be imported, re-trained, and executed easily by just using a browser.

41. What are activation functions in TensorFlow?

Activation functions are functions applied to the output side of a neural network that serves to be the input of the next layer. It forms a very important part of [neural networks](#) as it provides nonlinearity that sets apart a neural network from logistic regression.

42. What is the code that is used to check the version of TensorFlow using Python?

There are two commands depending on the Python version:

Python 2:

```
1 | python -c 'import tensorflow as tf; print(tf.__version__)'
```

Python 3:

```
1 | python3 -c 'import tensorflow as tf; print(tf.__version__)'
```

43. What is model quantization in TensorFlow?

The process of handling the complexity that follows when optimizing inferences can be greatly minimized using TensorFlow. Model quantization is primarily used to reduce the representation of weights and also for the storage and computation of the activation function.

Using model quantization provides users with two main advantages:

- Support for a variety of CPU platforms
- SIMD instruction handling capabilities

44. What is the simple syntax that can be used to convert a NumPy array into a tensor?

There are two ways a NumPy array can be converted into a tensor when working with Python. The first one is as follows:

- `train.shuffle_batch()`

And the second way is:

- `convert_to_tensor(tensor1d, dtype = tf.float64)`

The high-level code offers a good amount of readability and ease-of-use and denoted by the above piece of code.

45. How is the weighted standard error computed in TensorFlow?

The weighted standard error is a standard metric that is used to compute the coefficient of determination when working with a linear regression model.

It provides an easy way to evaluate the model and can be used as shown below:

Used along with TFLearn estimators

```
1 | weighted_r2 = WeightedR2()  
2 | regression = regression(net, metric=weighted_r2)
```

46. What are some of the commonly used optimizers when training a model in TensorFlow?

You can use many optimizers based on various factors, such as the learning rate, performance metric, dropout, gradient, and more.

Following are some of the popular optimizers:

- AdaDelta
- AdaGrad
- Adam
- Momentum
- RMSprop
- Stochastic Gradient Descent

47. What is the use of ArrayFlow and FeedDictFlow in TensorFlow?

ArrayFlow is used to convert array entities into tensors and store them automatically in a queue data structure.

- `data_flow.ArrayFlow()`

FeedDictFlow is used to generate a stream of batch data from the input dataset. The working is based on two queues, where one is used to generate batches and the other is used to load the data and apply preprocessing methods to it.

- `data_flow.FeedDictFlow()`

48. What are some of the parameters to consider when implementing the Word2vec algorithm in TensorFlow?

The Word2vec algorithm is used to compute the vector representations of words from an input dataset.

There are six parameters that have to be considered:

- `embedding_size`: Denotes the dimension of the embedding vector

- `max_vocabulary_size`: Denotes the total number of unique words in the vocabulary
- `min_occurrence`: Removes all words that do not appear at least 'n' number of times
- `skip_window`: Denotes words to be considered or not for processing
- `num_skips`: Denotes the number of times you can reuse an input to generate a label
- `num_sampled`: Denotes the number of negative examples to sample from the input

This next one on the top TensorFlow interview questions and answers blog discusses a tricky question, so pay attention!

49. What are some of the important parameters to consider when implementing a random forest algorithm in TensorFlow?

There are six main parameters you should think about and plan when implementing a [random forest algorithm](#) in TensorFlow:

- Number of inputs
- Feature count
- Number of samples per batch
- Total number of training steps
- Number of trees
- Maximum number of nodes

50. What are some of the numerical and categorical loss functions supported when working with TensorFlow?

Following are some of the widely used numerical and categorical loss functions supported when working with TensorFlow:

Numerical loss functions:

- L1 loss
- L2 loss
- Pseudo-Huber loss

Categorical loss functions:

- Hinge loss
- Cross-entropy loss
- Sigmoid-entropy loss
- Weighted cross-entropy loss

Each of the loss functions mentioned above has a specific use based on the input data and the type of modeling involved.

If you are looking forward to becoming an expert in TensorFlow and AI, make sure to check out IntelliPaat's [AI Course](#). With this program, you can become proficient in all of the concepts of Deep Learning and AI and earn a course certificate as well.