

50+ Deep Learning Interview Questions

- Deep Learning Coding Interview Questions
- Deep Learning NLP Interview Questions
- Deep Learning Computer Vision Interview Questions
- More Sample Deep Learning Interview Questions
- FAQs on Deep Learning Interview Questions

Deep Learning Coding Interview Questions

Following are some deep learning coding interview questions and answers to get you started.

Q1. What are the different types of activation functions?

This is an important Deep Learning coding interview question. You must know the following types of activation functions:

1. **Sigmoid function:** It is a nonlinear function in an ANN that is mostly used in feedforward neural networks. It's a differentiable real function with positive derivatives and a certain degree of smoothness. It is written as: $\sigma(x) = \frac{1}{1 + e^{-x}}$.
2. **Hyperbolic tangent function (Tanh):** It is a smoother, zero-centered function (range of -1 to +1). The output is represented by: $\tanh(x) = \frac{e^x - e^{-x}}{e^x + e^{-x}}$. The primary advantage of this function is that it gives a zero-centered output that helps in backpropagation.
3. **Softmax function:** It is used to generate probability distribution from a vector of real numbers. This function returns the output between 0 and 1, with the sum of probabilities equals to 1. This is written as: $\text{softmax}(x_i) = \frac{e^{x_i}}{\sum_j e^{x_j}}$. It is used in multi-class models, returning probabilities of each class, with the target having the highest probability.
4. **Softsign function:** It is commonly used in regression computation issues and text-to-speech applications. It's a quadratic polynomial, written as: $\text{softsign}(x) = \frac{x}{1 + |x|}$.
5. **Rectified linear unit of function:** It outperforms other AFs in generalization and performance. The function is roughly linear and preserves the features of linear models, making gradient-descent approaches easier to optimize. It is written as: $\text{ReLU}(x) = \max(0, x)$.
6. **Exponential linear unit of function:** The major advantage of this function is that it can solve the vanishing gradient problem by employing identity for positive values and boosting the model's learning properties. It is represented by: $\text{ELU}(x) = \max(0, x) + \min(0, e^x - 1)$.

Q2. How does recurrent neural network backpropagation vary from artificial neural network backpropagation?

Each node in a recurrent neural network has an additional loop. This makes it different from artificial neural network propagation. This loop incorporates a temporal component into the network. The main advantage of recurrent neural networks is that they allow for sequential data information. This is usually impossible with a generic artificial neural network.

Q3. Can a deep learning model be solely built on linear regression?

If you are well-versed in Deep Learning, you can answer these types of Deep Learning interview questions with ease.

A deep learning model may be solely built on linear regression. However, the problem should be represented by a linear equation, which does not boost the machine learning model's predictive capacity due to the addition of nodes. Hence, building a deep learning model solely on linear regression creates no spectacular results.

Q4. What is a computational graph in Deep Learning?

This is one of the important topics asked in Deep Learning interview questions.

A computational graph is a series of operations performed to take inputs and arrange them as nodes in a graph. It is a way of implementing mathematical calculations into a graph. This way, it will help in parallel processing and provide high performance in terms of computational capability.

Q5. What are the types of autoencoders, and where are they used?

This is a commonly asked in Deep Learning interview question. You must have a sound understanding of what autoencoders are to answer this.

Autoencoders are used worldwide. Some of the popular usages of autoencoders are:

1. Adding color to black-white images
2. Removing noise from images
3. Dimensionality reduction
4. Feature removal and variation

You must know there are four types of autoencoders. They are:

1. Deep autoencoders
2. Convolutional autoencoders
3. Sparse autoencoders
4. Contractive autoencoders

Recommended Reading: [Google Machine Learning Engineer Interview Prep](#)

Deep Learning NLP Interview Questions

To excel at ML or data science interviews, you must have profound knowledge of natural language processing (NLP) in Deep Learning. Following are a few questions that you must practice to nail your Deep Learning interview at FAANG and tier-1 tech companies.

Q1. What do you understand about text normalization in NLP?

These types of Deep Learning interview questions test your fundamental knowledge of the subject.

When developing NLP tools to work with exceptional data, it's beneficial to attain a canonical representation of textual content. This is known as textual normalization. Textual normalization captures different kinds of variations into one representation.

Q2. Do you know what feature engineering is?

When you employ machine learning methods to complete your modeling, you need to input pre-processed text into an NLP algorithm. This set of strategies used for this process is known as feature engineering or feature extraction. The main purpose of feature extraction is to convert the text's qualities into a numeric vector that NLP algorithms can understand. This stage is known as text representation.

Q3. Explain TF-IDF in NLP.

TF-IDF is known as Term-Frequency-Inverse Document Frequency. It helps you get the importance of a particular word relative to other words in the corpus. It converts words into vectors and adds semantic information, resulting in weighted unusual words. These words can be utilized in various NLP applications. Moreover, it's a common scoring metric in information retrieval and summarization.

Q4. What do you understand about POS tagging?

A part-of-speech (POS) tagger reads the text in a language and assigns specific parts to each word, such as noun, verb, adverb, and others. POS taggers employ an algorithm to label terms in text bodies. These labels create various complex categories with tags like "noun plural" or other complicated labels.

Q5. What is the difference between NLP and NLU?

This is one of the most asked Deep Learning interview questions. The differences between NLP and NLU are:

Natural Language Understanding (NLU)	Natural Language Programming (NLP)
Aids in solving AI's complex problems.	A system that manages end-to-end conversations between computers and people simultaneously.
Allows machines to interpret unstructured input by transforming them into structured text.	Humans and machines are involved in NLP.
Concentrates on extracting meaning and context.	Focuses on interpreting language in its most literal sense.
Helps machines deduce the meaning behind the language content.	Can parse text-based on grammar, typography, structure, and point of view.

Recommended Reading: [Amazon Machine Learning Engineer Interview Prep](#)

Deep Learning Computer Vision Interview Questions

If you are applying for a role of a Computer Vision Engineer in any top company, you must practice the following Deep Learning computer vision interview questions to uplevel your preparation:

Q1. What are the features detected by the initial layers of a neural network used for computer vision? How is it different from what is detected by the later neural network layers?

Neural network's earlier layers detect simple features of an image (for example, edges or corners). As you go deeper, the features become increasingly complex, detecting patterns and shapes in the neural network. The later layers can detect intricate patterns, such as complete objects.

Q2. How will you address the edge pixels issue during convolutional operation?

You can use padding to address the issue of filter or kernel extracting information from the edge pixels less compared to the central pixel. Padding is the addition of one or more rows or columns of pixels along the boundary of the image.

It forms the new pixels of the picture. Therefore, it results in insufficient extraction of information from the original edge pixels. It also prevents the shrinking of an image due to the convolution operations.

Q3. You are given a 5x5 image with a 3x3 filter and a padding $p = 1$. What will be the resultant image's size if a convolutional stride of $s = 2$ is used?

You should know that for an $n \times n$ image with an $f \times f$ filter, padding p , and stride length s , resultant image's size after convolution has the shape $(n + 2p - fs + 1) \times (n + 2p - fs + 1)$. Therefore, per the data provided, the resulting size of the image will be $((5 + 2 * 1 - 3) / 2 + 1) \times ((5 + 2 * 1 - 3) / 2 + 1) = 3 \times 3$.

Q4. What will be the resultant image size for an RGB image of 10x10x3 convolved with a 3x3 filter?

The convolution operation is not possible for such dimensions of an RGB image. The third dimension (number of channels) should be the same to achieve convolution. However, if the 10x10x3 image is convolved in a 3x3x3 filter, the dimensions of the resultant image will be 4x4.

Q5. How many parameters need to be learned in pooling layers?

The pooling layer contains hyperparameters describing the filter size and the stride length. These parameters are set and work as a fixed computation. Hence, no parameters are to be learned in the pooling layers.

Recommended Reading: [Amazon Machine Learning Engineer Interview Process](#)

More Sample Deep Learning Interview Questions

Here are a few sample Deep Learning interview questions that you must prepare to enhance your preparation for your next tech interview:

1. What is an ensemble method in NLP?
2. State the steps to build a text classification system.
3. How is parsing done in NLP?
4. Differentiate between deep learning and machine learning.
5. What is a bag of words (BOW)?
6. What is Latent Semantic Indexing (LSI) in NLP?
7. What are some metrics on which NLP models are evaluated?
8. Explain the pipeline for information extraction.
9. What do you understand about autoencoders?
10. Explain the meaning of masked language modeling.
11. Explain pragmatic analysis in NLP.
12. What is the meaning of N-gram in NLP?
13. What do you mean by perplexity in NLP?

14. Explain why the inputs in computer vision problems can get huge. Provide a solution to overcome this challenge.
15. What should the padding be for a 10x10 image used with a 5x5 filter to get an image of the same size as the original image?
16. What method can be used to evaluate an object localization model? How does it work?
17. How will you use IoU for resolving the issue of multiple detections of the same object?
18. Give us an example of a scenario that would require the use of anchor boxes.
19. How is the Siamese Network beneficial in addressing the one-shot learning problem?
20. What purpose does grayscaling serve?
21. Explain translational equivariance.
22. Explain the object detection algorithm YOLO.
23. What do you know about dropouts?
24. Explain exploding and vanishing gradients.
25. Differentiate between bias and variance in the context of deep learning models. How can you achieve a balance between the two?
26. According to you, which one is more powerful — a two-layer neural network without any activation function or a two-layer decision tree?
27. While building a neural network architecture, how will you decide how many neurons and hidden layers should the neural network have?
28. What is an activation function? What is the use of an activation function?
29. What deep learning algorithm works best for face detection?
30. What is Stochastic Gradient Descent and how is it different from Batch Gradient Descent?
31. Explain how you would fix the constant validation accuracy in a Convolutional Neural Network (CNN)?
32. What are the differences between a shallow network and a deep network.
33. What is a tensor in deep learning?
34. What are the advantages of transfer learning?
35. Difference between multi-class and multi-label classification problems.
36. What are the different techniques to achieve data normalization?
37. What are Forward and Back Propagation in the context of deep learning?
38. List the different types of deep neural networks.