
Professional Certificate in AI for Lean Practitioners

AI Fundamentals

Artificial Intelligence (AI) is a branch of computer science that deals with creating intelligent machines that can think and learn like humans. In the Professional Certificate in AI for Lean Practitioners, you will learn about the key terms and vocabulary used in AI fundamentals. Here's a comprehensive explanation of the important terms you'll encounter in the course.

1. Machine Learning (ML)

Machine learning is a subset of AI that enables machines to learn from data and improve their performance over time. ML algorithms use statistical models to analyze and draw inferences from patterns in data, without explicit programming.

2. Deep Learning (DL)

Deep learning is a subset of ML that uses artificial neural networks with many layers to perform complex tasks such as image and speech recognition. DL algorithms can automatically learn features and representations from raw data, eliminating the need for manual feature engineering.

3. Neural Networks

Neural networks are computational models inspired by the structure and function of the human brain. They consist of interconnected nodes or neurons that process information and learn from data. Neural networks can be used for various tasks, including classification, regression, and prediction.

4. Supervised Learning

Supervised learning is a type of ML where the algorithm is trained on labeled data, i.e., data with known outputs. The algorithm learns to map inputs to outputs by minimizing the difference between its predictions and the true labels. Supervised learning can be used for classification and regression tasks.

5. Unsupervised Learning

Unsupervised learning is a type of ML where the algorithm is trained on unlabeled data, i.e., data without known outputs. The algorithm learns to discover patterns and structures in the data by itself. Unsupervised learning can be used for clustering, dimensionality reduction, and anomaly detection tasks.

6. Reinforcement Learning

Reinforcement learning is a type of ML where the algorithm learns to take actions in an environment to maximize a reward signal. The algorithm interacts with the environment, receives feedback in the form of rewards or penalties, and adjusts its behavior accordingly. Reinforcement learning can be used for control, game playing, and robotics tasks.

7. Overfitting

Overfitting is a common problem in ML where the algorithm learns the training data too well, including its noise and outliers. Overfitting results in poor generalization performance, i.e., the algorithm performs well on the training data but poorly on new, unseen data. Regularization techniques, such as L1 and L2 regularization, can be used to prevent overfitting.

8. Underfitting

Underfitting is a problem in ML where the algorithm fails to learn the underlying patterns and structures in the data. Underfitting results in poor performance on both the training and new data. Increasing the complexity of the algorithm, such as adding more layers or neurons, or using different features, can help prevent underfitting.

9. Bias-Variance Tradeoff

The bias-variance tradeoff is a fundamental concept in ML that refers to the balance between the bias and variance of the algorithm. Bias is the error due to assumptions and simplifications in the algorithm, while variance is the error due to sensitivity to the training data. A high bias algorithm is simplistic and underfits the data, while a high variance algorithm is complex and overfits the data. The goal of ML is to find the sweet spot between bias and variance that results in optimal performance.

10. Activation Function

An activation function is a non-linear function applied to the output of a neural network layer. Activation functions introduce non-linearity into the model, allowing it to learn complex, non-linear relationships between inputs and outputs. Common activation functions include the sigmoid, tanh, and ReLU functions.

11. Loss Function

A loss function is a mathematical function that measures the difference between the predicted and true outputs. The goal of ML is to minimize the loss function, i.e., to find the model parameters that result in the smallest difference between the predictions and the true labels. Common loss functions include the mean squared error (MSE) and cross-entropy loss functions.

12. Optimization Algorithm

An optimization algorithm is a method for finding the model parameters that minimize the loss function. Gradient descent is a common optimization algorithm used in ML. It iteratively adjusts the parameters in the direction of the negative gradient of the loss function, i.e., the direction that reduces the loss.

13. Hyperparameter

A hyperparameter is a parameter that is not learned from the data but is set before training the model. Hyperparameters include the learning rate, the number of layers and neurons, and the type of activation and loss functions. Tuning the hyperparameters can improve the performance of the model.

14. Cross-Validation

Cross-validation is a technique for evaluating the performance of a ML model. It involves dividing the data into k folds, training the model on k-1 folds, and testing it on the remaining fold. This process is repeated k times, with each fold serving as the test set once. The average performance across the k folds is used as the final performance metric.

15. Natural Language Processing (NLP)

NLP is a subfield of AI that deals with the interaction between computers and human language. NLP algorithms can analyze, understand, and generate human language, enabling applications such as language translation, sentiment analysis, and chatbots.

16. Computer Vision

Computer vision is a subfield of AI that deals with the interpretation and analysis of visual data, such as

images and videos. Computer vision algorithms can recognize objects, detect patterns, and track movements, enabling applications such as image and video recognition, object detection, and autonomous vehicles.

17. Explainable AI (XAI)

Explainable AI is a growing field in AI that aims to make ML models more interpretable and transparent. XAI algorithms can provide insights into how the model makes decisions, enabling users to understand and trust the model's outputs.

Challenge:

Now that you have learned about the key terms and vocabulary used in AI fundamentals, try to apply them to a real-world problem. For example, you can use a neural network to classify images of animals or use NLP to analyze the sentiment of social media posts. Make sure to choose a problem that interests you and has practical applications. Good luck!

In summary, AI is a rapidly evolving field with many exciting opportunities and challenges. Understanding the key terms and vocabulary used in AI fundamentals is essential for leveraging the power of AI in practical applications. By mastering the concepts and techniques covered in this explanation, you will be well on your way to becoming a proficient AI practitioner.