
Professional Certificate in AI for Graphic Designers

Machine Learning Fundamentals

Machine Learning Fundamentals:

Machine learning is a subset of artificial intelligence that focuses on developing algorithms and models that allow computers to learn from and make predictions or decisions based on data. In this course, we will cover key terms and vocabulary essential for understanding the fundamentals of machine learning.

Data:

Data is the foundation of machine learning. It refers to the information that is used to train a machine learning model. Data can be structured (e.g., tabular data in a spreadsheet) or unstructured (e.g., text, images, or audio). High-quality and relevant data is crucial for building accurate machine learning models.

Features:

Features are individual measurable properties or characteristics of the data. In a dataset, each column represents a feature. For example, in a dataset of housing prices, features could include the number of bedrooms, square footage, and location. Features are used as input to machine learning algorithms to make predictions or classifications.

Label:

In supervised learning, the label is the output variable that we are trying to predict. It is also known as the target variable. For example, in a dataset of images of cats and dogs, the label would be the class (cat or dog) associated with each image.

Model:

A model in machine learning is a mathematical representation of a system that is trained on data to make predictions or decisions. Models can range from simple linear regression models to complex deep learning neural networks.

Algorithm:

An algorithm is a set of rules or instructions that a machine learning model follows to learn from data and make predictions. There are many different types of machine learning algorithms, each suited to different types of problems and data.

Training:

Training is the process of feeding data into a machine learning model to allow it to learn the patterns and relationships in the data. During training, the model adjusts its parameters to minimize the difference

between its predictions and the actual labels in the training data.

Testing:

Testing is the process of evaluating a trained machine learning model on a separate dataset called the testing or validation set. This allows us to assess how well the model generalizes to new, unseen data.

Accuracy:

Accuracy is a common metric used to evaluate the performance of a machine learning model. It measures the proportion of correct predictions the model makes on the test data. High accuracy indicates that the model is making accurate predictions.

Overfitting:

Overfitting occurs when a machine learning model performs well on the training data but poorly on new, unseen data. This is often a result of the model memorizing noise or irrelevant patterns in the training data. Regularization techniques can help prevent overfitting.

Underfitting:

Underfitting occurs when a machine learning model is too simple to capture the underlying patterns in the data. An underfit model performs poorly on both the training and test data. Increasing the complexity of the model or adding more features can help mitigate underfitting.

Supervised Learning:

Supervised learning is a type of machine learning where the model is trained on labeled data. The model learns to map input features to output labels by minimizing a loss function. Common supervised learning algorithms include linear regression, logistic regression, and support vector machines.

Unsupervised Learning:

Unsupervised learning is a type of machine learning where the model is trained on unlabeled data. The goal is to find patterns or structure in the data without explicit guidance. Clustering algorithms and dimensionality reduction techniques are examples of unsupervised learning.

Reinforcement Learning:

Reinforcement learning is a type of machine learning where an agent learns to make decisions by interacting with an environment and receiving rewards or penalties. The agent learns through trial and error to maximize its cumulative reward over time.

Deep Learning:

Deep learning is a subset of machine learning that focuses on neural networks with multiple layers (deep neural networks). Deep learning models are capable of learning complex patterns in data and are commonly

used for tasks such as image recognition, natural language processing, and speech recognition.

Neural Network:

A neural network is a computational model inspired by the structure and function of the human brain. It consists of interconnected nodes (neurons) organized into layers. Neural networks are used in deep learning to learn hierarchical representations of data.

Convolutional Neural Network (CNN):

A convolutional neural network is a type of neural network commonly used for image recognition tasks. CNNs are designed to automatically learn spatial hierarchies of features from images by applying convolutional filters and pooling operations.

Recurrent Neural Network (RNN):

A recurrent neural network is a type of neural network designed to handle sequential data, such as time series or natural language. RNNs have feedback connections that allow them to maintain a memory of past inputs, making them suitable for tasks like language modeling and speech recognition.

Loss Function:

A loss function is a measure of how well a machine learning model is performing on the training data. It quantifies the difference between the predicted output and the actual label. During training, the model aims to minimize the loss function to improve its predictive accuracy.

Gradient Descent:

Gradient descent is an optimization algorithm used to update the parameters of a machine learning model during training. It works by iteratively moving in the direction of steepest descent of the loss function to find the optimal set of parameters that minimize the loss.

Hyperparameters:

Hyperparameters are parameters of a machine learning model that are set before training and remain constant during training. Examples of hyperparameters include the learning rate, batch size, and number of hidden layers. Tuning hyperparameters is essential for optimizing model performance.

Feature Engineering:

Feature engineering is the process of selecting, transforming, and creating new features from the raw data to improve the performance of a machine learning model. This can involve techniques such as one-hot encoding, scaling, and dimensionality reduction.

Cross-Validation:

Cross-validation is a technique used to assess the generalization performance of a machine learning model.

It involves splitting the data into multiple subsets (folds), training the model on some folds, and evaluating it on the remaining fold. This helps to reduce the risk of overfitting.

Bias-Variance Tradeoff:

The bias-variance tradeoff is a fundamental concept in machine learning that describes the balance between the bias (error due to simplifying assumptions) and variance (sensitivity to fluctuations in the training data) of a model. A high-bias model is likely to underfit, while a high-variance model is likely to overfit.

Ensemble Learning:

Ensemble learning is a technique that combines multiple machine learning models to improve predictive performance. Common ensemble methods include bagging (e.g., random forests), boosting (e.g., AdaBoost), and stacking. Ensemble learning can help reduce overfitting and improve generalization.

Regularization:

Regularization is a technique used to prevent overfitting in machine learning models. It involves adding a penalty term to the loss function that discourages the model from learning complex patterns that may not generalize well. L1 and L2 regularization are common regularization techniques.

Transfer Learning:

Transfer learning is a machine learning technique where a pre-trained model is used as a starting point for a new task. By leveraging knowledge learned from a related task, transfer learning can help improve model performance with limited data or computational resources.

Unstructured Data:

Unstructured data refers to data that does not have a predefined format or structure, making it difficult to analyze using traditional methods. Examples of unstructured data include text, images, videos, and sensor data. Machine learning techniques such as natural language processing and computer vision are used to extract insights from unstructured data.

Supervised Learning:

Supervised learning is a type of machine learning where the model is trained on labeled data. The model learns to map input features to output labels by minimizing a loss function. Common supervised learning algorithms include linear regression, logistic regression, and support vector machines.

Unsupervised Learning:

Unsupervised learning is a type of machine learning where the model is trained on unlabeled data. The goal is to find patterns or structure in the data without explicit guidance. Clustering algorithms and dimensionality reduction techniques are examples of unsupervised learning.

Reinforcement Learning:

Reinforcement learning is a type of machine learning where an agent learns to make decisions by interacting with an environment and receiving rewards or penalties. The agent learns through trial and error to maximize its cumulative reward over time.

Deep Learning:

Deep learning is a subset of machine learning that focuses on neural networks with multiple layers (deep neural networks). Deep learning models are capable of learning complex patterns in data and are commonly used for tasks such as image recognition, natural language processing, and speech recognition.

Neural Network:

A neural network is a computational model inspired by the structure and function of the human brain. It consists of interconnected nodes (neurons) organized into layers. Neural networks are used in deep learning to learn hierarchical representations of data.

Convolutional Neural Network (CNN):

A convolutional neural network is a type of neural network commonly used for image recognition tasks. CNNs are designed to automatically learn spatial hierarchies of features from images by applying convolutional filters and pooling operations.

Recurrent Neural Network (RNN):

A recurrent neural network is a type of neural network designed to handle sequential data, such as time series or natural language. RNNs have feedback connections that allow them to maintain a memory of past inputs, making them suitable for tasks like language modeling and speech recognition.

Loss Function:

A loss function is a measure of how well a machine learning model is performing on the training data. It quantifies the difference between the predicted output and the actual label. During training, the model aims to minimize the loss function to improve its predictive accuracy.

Gradient Descent:

Gradient descent is an optimization algorithm used to update the parameters of a machine learning model during training. It works by iteratively moving in the direction of steepest descent of the loss function to find the optimal set of parameters that minimize the loss.

Hyperparameters:

Hyperparameters are parameters of a machine learning model that are set before training and remain constant during training. Examples of hyperparameters include the learning rate, batch size, and number of hidden layers. Tuning hyperparameters is essential for optimizing model performance.

Feature Engineering:

Feature engineering is the process of selecting, transforming, and creating new features from the raw data to improve the performance of a machine learning model. This can involve techniques such as one-hot encoding, scaling, and dimensionality reduction.

Cross-Validation:

Cross-validation is a technique used to assess the generalization performance of a machine learning model. It involves splitting the data into multiple subsets (folds), training the model on some folds, and evaluating it on the remaining fold. This helps to reduce the risk of overfitting.

Bias-Variance Tradeoff:

The bias-variance tradeoff is a fundamental concept in machine learning that describes the balance between the bias (error due to simplifying assumptions) and variance (sensitivity to fluctuations in the training data) of a model. A high-bias model is likely to underfit, while a high-variance model is likely to overfit.

Ensemble Learning:

Ensemble learning is a technique that combines multiple machine learning models to improve predictive performance. Common ensemble methods include bagging (e.g., random forests), boosting (e.g., AdaBoost), and stacking. Ensemble learning can help reduce overfitting and improve generalization.

Regularization:

Regularization is a technique used to prevent overfitting in machine learning models. It involves adding a penalty term to the loss function that discourages the model from learning complex patterns that may not generalize well. L1 and L2 regularization are common regularization techniques.

Transfer Learning:

Transfer learning is a machine learning technique where a pre-trained model is used as a starting point for a new task. By leveraging knowledge learned from a related task, transfer learning can help improve model performance with limited data or computational resources.

Unstructured Data:

Unstructured data refers to data that does not have a predefined format or structure, making it difficult to analyze using traditional methods. Examples of unstructured data include text, images, videos, and sensor data. Machine learning techniques such as natural language processing and computer vision are used to extract insights from unstructured data.

Supervised Learning:

Supervised learning is a type of machine learning where the model is trained on labeled data. The model

learns to map input features to output labels by minimizing a loss function. Common supervised learning algorithms include linear regression, logistic regression, and support vector machines.

Unsupervised Learning:

Unsupervised learning is a type of machine learning where the model is trained on unlabeled data. The goal is to find patterns or structure in the data without explicit guidance. Clustering algorithms and dimensionality reduction techniques are examples of unsupervised learning.

Reinforcement Learning:

Reinforcement learning is a type of machine learning where an agent learns to make decisions by interacting with an environment and receiving rewards or penalties. The agent learns through trial and error to maximize its cumulative reward over time.

Deep Learning:

Deep learning is a subset of machine learning that focuses on neural networks with multiple layers (deep neural networks). Deep learning models are capable of learning complex patterns in data and are commonly used for tasks such as image recognition, natural language processing, and speech recognition.

Neural Network:

A neural network is a computational model inspired by the structure and function of the human brain. It consists of interconnected nodes (neurons) organized into layers. Neural networks are used in deep learning to learn hierarchical representations of data.

Convolutional Neural Network (CNN):

A convolutional neural network is a type of neural network commonly used for image recognition tasks. CNNs are designed to automatically learn spatial hierarchies of features from images by applying convolutional filters and pooling operations.

Recurrent Neural Network (RNN):

A recurrent neural network is a type of neural network designed to handle sequential data, such as time series or natural language. RNNs have feedback connections that allow them to maintain a memory of past inputs, making them suitable for tasks like language modeling and speech recognition.

Loss Function:

A loss function is a measure of how well a machine learning model is performing on the training data. It quantifies the difference between the predicted output and the actual label. During training, the model aims to minimize the loss function to improve its predictive accuracy.

Gradient Descent:

Gradient descent is an optimization algorithm used to update the parameters of a machine learning model during training. It works by iteratively moving in the direction of steepest descent of the loss function to find the optimal set of parameters that minimize the loss.

Hyperparameters:

Hyperparameters are parameters of a machine learning model that are set before training and remain constant during training. Examples of hyperparameters include the learning rate, batch size, and number of hidden layers. Tuning hyperparameters is essential for optimizing model performance.

Feature Engineering:

Feature engineering is the process of selecting, transforming, and creating new features from the raw data to improve the performance of a machine learning model. This can involve techniques such as one-hot encoding, scaling, and dimensionality reduction.

Cross-Validation:

Cross-validation is a technique used to assess the generalization performance of a machine learning model. It involves splitting the data into multiple subsets (folds), training the model on some folds, and evaluating it on the remaining fold. This helps to reduce the risk of overfitting.

Bias-Variance Tradeoff:

The bias-variance tradeoff is a fundamental concept in machine learning that describes the balance between the bias (error due to simplifying assumptions) and variance (sensitivity to fluctuations in the training data) of a model. A high-bias model is likely to underfit, while a high-variance model is likely to overfit.

Ensemble Learning:

Ensemble learning is a technique that combines multiple machine learning models to improve predictive performance. Common ensemble methods include bagging (e.g., random forests), boosting (e.g., AdaBoost), and stacking. Ensemble learning can help reduce overfitting and improve generalization.

Regularization:

Regularization is a technique used to prevent overfitting in machine learning models. It involves adding a penalty term to the loss function that discourages the model from learning complex patterns that may not generalize well. L1 and L2 regularization are common regularization techniques.

Transfer Learning:

Transfer learning is a machine learning technique where a pre-trained model is used as a starting point for a new task. By leveraging knowledge learned from a related task, transfer learning can help improve model performance with limited data or computational resources.

Unstructured Data:

Unstructured data refers to data that does not have a predefined format or structure, making it difficult to analyze using traditional methods. Examples of unstructured data include text, images, videos, and sensor data. Machine learning techniques such as natural language processing and computer vision are used to extract insights from unstructured data.