

Fundamentals of AI and Machine Learning

Activation Function: A function used in artificial neural networks to introduce non-linearity into the output of a node or neuron. Common activation functions include the sigmoid, tanh, and ReLU functions.

Artificial Intelligence (AI): The simulation of human intelligence in machines that are programmed to think and learn. AI can be categorized into two main types: Narrow AI, which is designed to perform a narrow task (e.g., facial recognition), and General AI, which can perform any intellectual task that a human being can do.

Artificial Neural Network (ANN): A type of machine learning algorithm inspired by the structure and function of the human brain. ANNs are composed of interconnected nodes or neurons that process information in parallel.

Backpropagation: A training algorithm used in artificial neural networks to adjust the weights of the connections between nodes. Backpropagation uses the chain rule to calculate the gradient of the loss function with respect to the weights, and then updates the weights in the opposite direction of the gradient.

Bias: A parameter in machine learning algorithms that determines the prior probability of a given class or outcome. Bias can lead to systematic errors in the predictions made by the model, and can be reduced through techniques such as cross-validation and regularization.

Convolutional Neural Network (CNN): A type of deep learning algorithm that is commonly used for image recognition tasks. CNNs use convolutional layers to extract features from images, and pooling layers to reduce the dimensionality of the data.

Deep Learning: A subset of machine learning that uses artificial neural networks with multiple layers to learn complex patterns in data. Deep learning algorithms can automatically extract features from raw data, and have been used to achieve state-of-the-art performance in many domains, including computer vision, natural language processing, and speech recognition.

Feature Engineering: The process of selecting and transforming raw data into features that can be used by machine learning algorithms. Feature engineering can include techniques such as normalization, scaling, and dimensionality reduction.

Feature Selection: The process of selecting a subset of relevant features from a larger set of available features. Feature selection can improve the performance of machine learning algorithms by reducing overfitting and improving interpretability.

Gradient Descent: A optimization algorithm used to minimize the loss function in machine learning. Gradient descent updates the parameters of the model in the direction of the negative gradient of the loss

function with respect to the parameters.

Hyperparameter: A parameter that is set before training a machine learning model, and is not learned from the data. Examples of hyperparameters include the learning rate, regularization strength, and number of hidden layers in a neural network.

Label: The target variable or outcome that a machine learning model is trying to predict.

Loss Function: A function that measures the difference between the predicted output of a machine learning model and the true output. The loss function is used to optimize the parameters of the model during training.

Machine Learning (ML): A subset of artificial intelligence that uses statistical models and algorithms to learn patterns in data. ML algorithms can be categorized into three main types: supervised learning, unsupervised learning, and reinforcement learning.

Natural Language Processing (NLP): A field of AI that focuses on the interaction between computers and human language. NLP tasks include language translation, sentiment analysis, and text summarization.

Overfitting: A phenomenon in machine learning where a model learns the training data too well, and performs poorly on new, unseen data. Overfitting can be reduced through techniques such as regularization, cross-validation, and early stopping.

Principal Component Analysis (PCA): A dimensionality reduction technique that projects high-dimensional data onto a lower-dimensional space while preserving as much of the variance in the data as possible. PCA is often used for feature extraction and data visualization.

Reinforcement Learning (RL): A type of machine learning where an agent learns to take actions in an environment to maximize a reward signal. RL algorithms are used in applications such as robotics, game playing, and autonomous vehicles.

Regularization: A technique used in machine learning to prevent overfitting by adding a penalty term to the loss function. Regularization encourages the model to have smaller weights, which can improve generalization performance.

Regression: A supervised learning task where the goal is to predict a continuous output variable. Common regression algorithms include linear regression, polynomial regression, and support vector regression.

Supervised Learning: A type of machine learning where the model is trained on labeled data, and the goal is to predict the label for new, unseen data. Supervised learning algorithms can be categorized into two main types: regression and classification.

Support Vector Machine (SVM): A supervised learning algorithm used for classification and regression tasks. SVMs find the hyperplane that maximally separates the data points of different classes, and use kernel functions to transform the data into higher-dimensional spaces.

TensorFlow: An open-source software library for machine learning and deep learning. TensorFlow provides a flexible platform for defining and training machine learning models, and can be used for a wide range of applications, including image recognition, natural language processing, and reinforcement learning.

Unsupervised Learning: A type of machine learning where the model is trained on unlabeled data, and the goal is to discover patterns or structure in the data. Unsupervised learning algorithms can be categorized into two main types: clustering and dimensionality reduction.

Validation Set: A subset of the training data that is used to evaluate the performance of a machine learning model during training. The validation set is used to tune hyperparameters and prevent overfitting.

Feature Extraction: The process of transforming raw data into a set of features that can be used by a machine learning algorithm. Feature extraction can include techniques such as principal component analysis, independent component analysis, and non-negative matrix factorization.

Cross-Validation: A technique used to evaluate the performance of a machine learning model by splitting the data into multiple folds, and training and testing the model on different subsets of the data. Cross-validation can reduce the variance of the performance estimate and prevent overfitting.

Underfitting: A phenomenon in machine learning where a model fails to learn the underlying patterns in the data, and performs poorly on both the training and test data. Underfitting can be addressed by increasing the complexity of the model, adding features, or using a different algorithm.

Ensemble Learning: A machine learning technique that combines the predictions of multiple models to improve the overall performance. Ensemble learning can reduce the variance and bias of the predictions, and can be used for both regression and classification tasks.

Precision: A performance metric used in classification tasks to measure the proportion of true positive predictions out of all positive predictions. Precision is calculated as the number of true positives divided by the sum of true positives and false positives.

Recall: A performance metric used in classification tasks to measure the proportion of true positive predictions out of all actual positive instances. Recall is calculated as the number of true positives divided by the sum of true positives and false negatives.

F1 Score: A performance metric used in classification tasks to balance the trade-off between precision and recall. The F1 score is the harmonic mean of precision and recall, and is calculated as twice the product of precision and recall divided by the sum of precision and recall.

Area Under the ROC Curve (AUC-ROC): A performance metric used in binary classification tasks to measure the ability of the model to distinguish between positive and negative instances. The AUC-ROC is the area under the receiver operating characteristic (ROC) curve, which plots the true positive rate against the false positive rate.

Grid Search: A hyperparameter tuning technique used in machine learning to search for the best hyperparameter values in a grid-like space. Grid search can be computationally expensive, but can ensure

that the best hyperparameters are found.

Random Search: A hyperparameter tuning technique used in machine learning to search for the best hyperparameter values by randomly sampling from a given range. Random search can be less computationally expensive than grid search, and can still find good hyperparameters.

Early Stopping: A technique used in machine learning to prevent overfitting by stopping the training process before the model starts to memorize