

**Artificial Intelligence (AI):** The simulation of human intelligence in machines, enabling them to perform tasks that typically require human intelligence, such as visual perception, speech recognition, decision-making, and translation.

**AutoML:** The process of automating the process of applying machine learning to real-world problems.

**Bellman Equation:** A fundamental equation in reinforcement learning that expresses the value of a state in terms of the expected future rewards.

**Bias:** The error due to overly simplistic assumptions in the learning algorithm.

**Boosting:** A method of converting weak learners into strong learners by building a model from the training data, then creating a second model to correct the errors from the first model.

**Bootstrap Aggregating (Bagging):** A type of ensemble learning that combines the results of multiple models to get a generalized result.

**Chatbot:** A software application used to conduct an on-line chat conversation via text or text-to-speech, instead of providing direct contact with a live human agent.

**Classification:** A type of supervised learning where the goal is to categorize input into two or more classes.

**Clustering:** An unsupervised learning technique where you organize a dataset into groups with similar characteristics.

**Computer Vision:** An interdisciplinary field that trains computers to interpret and understand the visual world.

**Convolutional Neural Network (CNN):** A type of neural network, especially powerful for tasks like image recognition.

**Cross-Validation:** A resampling procedure used to evaluate machine learning models on a limited data sample.

**Deep Learning:** A subfield of ML utilizing neural networks with three or more layers, enabling the processing of non-linear features and the performance of complex tasks.

**Dimensionality Reduction:** The process of reducing the number of input variables in a dataset.

**Ensemble Learning:** A method in which multiple models are trained to solve the same problem and combined to get better results.

**Expert Systems:** Computer systems that emulate the decision-making ability of a human expert.

**Exploration vs Exploitation:** The dilemma in reinforcement learning of whether the agent should try new actions (exploration) or choose the known optimal action (exploitation).

**Feature Engineering:** The process of selecting and transforming variables when creating a predictive model.

**Fuzzy Logic:** A mathematical logic that attempts to solve problems by assigning values to an imprecise spectrum of data in order to arrive at the most accurate conclusion possible.

**Generative Adversarial Network (GAN):** A class of ML designed to generate new data that is similar to some input data.

**Grid Search:** A method for hyperparameter tuning that performs exhaustively searches through a manually specified subset of the hyperparameter space.

**Hyperparameter Tuning:** The process of choosing a set of optimal hyperparameters for a learning algorithm.

**Machine Learning (ML):** A subset of AI that involves the use of algorithms and statistical models to enable machines to improve their performance on a task without being explicitly programmed.

**Markov Decision Process (MDP):** A mathematical framework for modeling decision making in situations where outcomes are partly random and partly under the control of a decision maker.

**Monte Carlo Methods:** A statistical technique that allows numerical approximation based on random sampling.

**Natural Language Processing (NLP):** A field of AI focusing on the interaction between computers and humans using natural language.

**Neural Network:** A model inspired by the human brain, consisting of interconnected nodes (neurons), used for pattern recognition and learning.

**Optical Character Recognition (OCR):** The technology used to recognize text inside images, such as scanned documents and photos.

**Overfitting:** A modeling error that occurs when a function is too closely fit to a limited set of data points.

**Policy Gradient Methods:** A type of reinforcement learning methods which optimize the policy function directly.

**Q-Learning:** A model-free reinforcement learning algorithm to learn the value of actions in each state.

**Random Search:** A method for hyperparameter tuning that samples algorithm configurations from a distribution over the hyperparameter space.

**Recurrent Neural Network (RNN):** A type of neural network with loops to allow information persistence, useful for sequence prediction like time series or speech recognition.

**Regression Analysis:** A statistical process for estimating the relationships among variables.

**Reinforcement Learning Environment:** The world (real or simulated) with which the reinforcement learning agent interacts and learns from.

**Reinforcement Learning:** A type of ML where an agent learns to make decisions by taking actions in an environment to achieve maximum cumulative reward.

**Robotics:** The branch of technology that deals with the design, construction, operation, and application of robots.

**Semi-Supervised Learning:** A middle ground between supervised and unsupervised learning, where the model is trained with a combination of labeled and unlabeled data.

**Speech Recognition:** The interdisciplinary subfield of computational linguistics that develops methodologies and technologies that enable the recognition and translation of spoken language into text by computers.

**State-Action-Reward-State-Action (SARSA):** An on-policy reinforcement learning algorithm.

**Supervised Learning:** A type of ML where the model is trained using labeled data.

**Swarm Intelligence:** An AI technique based on the collective behavior of decentralized, self-organized systems.

**Temporal Difference Learning (TD Learning):** A combination of Monte Carlo ideas and dynamic programming (DP) ideas.

**Testing Data:** The dataset used to evaluate the performance of a model.

**Training Data:** The dataset used to train the model.

**Transfer Learning:** The practice of using a pre-trained model on a new, but similar problem.

**Underfitting:** A modeling error that occurs when a model cannot capture the underlying trend of the data.

**Unsupervised Learning:** A type of ML where the model is trained using unlabeled data.

**Validation Data:** The dataset used to tune model parameters and provide an unbiased evaluation of model fit during training.

**Variance:** The error due to too much complexity in the learning algorithm.