

# datacamp Top Machine Learning Algorithms

	ALGORITHM	DESCRIPTION	APPLICATIONS	ADVANTAGES	DISADVANTAGES	
Supervised Learning	Linear Models	<b>Linear Regression</b>	A simple algorithm that models a linear relationship between inputs and a continuous numerical output variable	<b>USE CASES</b> <ol style="list-style-type: none"> <li>1. Stock price prediction</li> <li>2. Predicting housing prices</li> <li>3. Predicting customer lifetime value</li> </ol>	<ol style="list-style-type: none"> <li>1. Explainable method</li> <li>2. Interpretable results by its output coefficients</li> <li>3. Faster to train than other machine learning models</li> </ol>	<ol style="list-style-type: none"> <li>1. Assumes linearity between inputs and output</li> <li>2. Sensitive to outliers</li> <li>3. Can underfit with small, high-dimensional data</li> </ol>
		<b>Logistic Regression</b>	A simple algorithm that models a linear relationship between inputs and a categorical output (1 or 0)	<b>USE CASES</b> <ol style="list-style-type: none"> <li>1. Credit risk score prediction</li> <li>2. Customer churn prediction</li> </ol>	<ol style="list-style-type: none"> <li>1. Interpretable and explainable</li> <li>2. Less prone to overfitting when using regularization</li> <li>3. Applicable for multi-class predictions</li> </ol>	<ol style="list-style-type: none"> <li>1. Assumes linearity between inputs and outputs</li> <li>2. Can overfit with small, high-dimensional data</li> </ol>
		<b>Ridge Regression</b>	Part of the regression family — it penalizes features that have low predictive outcomes by shrinking their coefficients closer to zero. Can be used for classification or regression	<b>USE CASES</b> <ol style="list-style-type: none"> <li>1. Predictive maintenance for automobiles</li> <li>2. Sales revenue prediction</li> </ol>	<ol style="list-style-type: none"> <li>1. Less prone to overfitting</li> <li>2. Best suited where data suffer from multicollinearity</li> <li>3. Explainable &amp; interpretable</li> </ol>	<ol style="list-style-type: none"> <li>1. All the predictors are kept in the final model</li> <li>2. Doesn't perform feature selection</li> </ol>
		<b>Lasso Regression</b>	Part of the regression family — it penalizes features that have low predictive outcomes by shrinking their coefficients to zero. Can be used for classification or regression	<b>USE CASES</b> <ol style="list-style-type: none"> <li>1. Predicting housing prices</li> <li>2. Predicting clinical outcomes based on health data</li> </ol>	<ol style="list-style-type: none"> <li>1. Less prone to overfitting</li> <li>2. Can handle high-dimensional data</li> <li>3. No need for feature selection</li> </ol>	<ol style="list-style-type: none"> <li>1. Can lead to poor interpretability as it can keep highly correlated variables</li> </ol>
	Tree-Based Models	<b>Decision Tree</b>	Decision Tree models make decision rules on the features to produce predictions. It can be used for classification or regression	<b>USE CASES</b> <ol style="list-style-type: none"> <li>1. Customer churn prediction</li> <li>2. Credit score modeling</li> <li>3. Disease prediction</li> </ol>	<ol style="list-style-type: none"> <li>1. Explainable and interpretable</li> <li>2. Can handle missing values</li> </ol>	<ol style="list-style-type: none"> <li>1. Prone to overfitting</li> <li>2. Sensitive to outliers</li> </ol>
		<b>Random Forests</b>	An ensemble learning method that combines the output of multiple decision trees	<b>USE CASES</b> <ol style="list-style-type: none"> <li>1. Credit score modeling</li> <li>2. Predicting housing prices</li> </ol>	<ol style="list-style-type: none"> <li>1. Reduces overfitting</li> <li>2. Higher accuracy compared to other models</li> </ol>	<ol style="list-style-type: none"> <li>1. Training complexity can be high</li> <li>2. Not very interpretable</li> </ol>
		<b>Gradient Boosting Regression</b>	Gradient Boosting Regression employs boosting to make predictive models from an ensemble of weak predictive learners	<b>USE CASES</b> <ol style="list-style-type: none"> <li>1. Predicting car emissions</li> <li>2. Predicting ride hailing fare amount</li> </ol>	<ol style="list-style-type: none"> <li>1. Better accuracy compared to other regression models</li> <li>2. It can handle multicollinearity</li> <li>3. It can handle non-linear relationships</li> </ol>	<ol style="list-style-type: none"> <li>1. Sensitive to outliers and can therefore cause overfitting</li> <li>2. Computationally expensive and has high complexity</li> </ol>
		<b>XGBoost</b>	Gradient Boosting algorithm that is efficient & flexible. Can be used for both classification and regression tasks	<b>USE CASES</b> <ol style="list-style-type: none"> <li>1. Churn prediction</li> <li>2. Claims processing in insurance</li> </ol>	<ol style="list-style-type: none"> <li>1. Provides accurate results</li> <li>2. Captures non linear relationships</li> </ol>	<ol style="list-style-type: none"> <li>1. Hyperparameter tuning can be complex</li> <li>2. Does not perform well on sparse datasets</li> </ol>
		<b>LightGBM Regressor</b>	A gradient boosting framework that is designed to be more efficient than other implementations	<b>USE CASES</b> <ol style="list-style-type: none"> <li>1. Predicting flight time for airlines</li> <li>2. Predicting cholesterol levels based on health data</li> </ol>	<ol style="list-style-type: none"> <li>1. Can handle large amounts of data</li> <li>2. Computational efficient &amp; fast training speed</li> <li>3. Low memory usage</li> </ol>	<ol style="list-style-type: none"> <li>1. Can overfit due to leaf-wise splitting and high sensitivity</li> <li>2. Hyperparameter tuning can be complex</li> </ol>
	Unsupervised Learning	Clustering	<b>K-Means</b>	K-Means is the most widely used clustering approach—it determines K clusters based on euclidean distances	<b>USE CASES</b> <ol style="list-style-type: none"> <li>1. Customer segmentation</li> <li>2. Recommendation systems</li> </ol>	<ol style="list-style-type: none"> <li>1. Scales to large datasets</li> <li>2. Simple to implement and interpret</li> <li>3. Results in tight clusters</li> </ol>
<b>Hierarchical Clustering</b>			A "bottom-up" approach where each data point is treated as its own cluster—and then the closest two clusters are merged together iteratively	<b>USE CASES</b> <ol style="list-style-type: none"> <li>1. Fraud detection</li> <li>2. Document clustering based on similarity</li> </ol>	<ol style="list-style-type: none"> <li>1. There is no need to specify the number of clusters</li> <li>2. The resulting dendrogram is informative</li> </ol>	<ol style="list-style-type: none"> <li>1. Doesn't always result in the best clustering</li> <li>2. Not suitable for large datasets due to high complexity</li> </ol>
<b>Gaussian Mixture Models</b>			A probabilistic model for modeling normally distributed clusters within a dataset	<b>USE CASES</b> <ol style="list-style-type: none"> <li>1. Customer segmentation</li> <li>2. Recommendation systems</li> </ol>	<ol style="list-style-type: none"> <li>1. Computes a probability for an observation belonging to a cluster</li> <li>2. Can identify overlapping clusters</li> <li>3. More accurate results compared to K-means</li> </ol>	<ol style="list-style-type: none"> <li>1. Requires complex tuning</li> <li>2. Requires setting the number of expected mixture components or clusters</li> </ol>
<b>Association</b>		<b>Apriori algorithm</b>	Rule based approach that identifies the most frequent itemset in a given dataset where prior knowledge of frequent itemset properties is used	<b>USE CASES</b> <ol style="list-style-type: none"> <li>1. Product placements</li> <li>2. Recommendation engines</li> <li>3. Promotion optimization</li> </ol>	<ol style="list-style-type: none"> <li>1. Results are intuitive and Interpretable</li> <li>2. Exhaustive approach as it finds all rules based on the confidence and support</li> </ol>	<ol style="list-style-type: none"> <li>1. Generates many uninteresting itemsets</li> <li>2. Computationally and memory intensive.</li> <li>3. Results in many overlapping item sets</li> </ol>