• Sensitive to Irrelevant Features: If the dataset has irrelevant or highly variable features, they can distort distance calculations and lead to poor performance.

Applications of k-NN

- **Recommender Systems**: k-NN can be used to recommend items based on similarity with other users or items.
- Image and Text Recognition: k-NN is often used in pattern recognition tasks, such as handwriting or facial recognition.
- **Medical Diagnosis**: Helps in classifying patients based on symptoms or medical imaging features.

Naïve Bayes

Naïve Bayes is a probabilistic machine learning algorithm based on **Bayes' Theorem**. It is mainly used for **classification tasks** and is especially popular in text classification, spam detection, and sentiment analysis. The "naïve" part of Naïve Bayes comes from the assumption that all features are independent of each other, which is rarely true in real-world data but simplifies calculations and often yields good performance.

How Naïve Bayes Works

Naïve Bayes calculates the probability of each class given the input features and classifies the input by choosing the class with the highest probability. The formula relies on Bayes' Theorem:

$$P(C|X) = \frac{P(X|C) \times P(C)}{P(X)}$$

where:

- P(C| X) is the posterior probability (probability of class C given feature set X),
- P(X| C) is the likelihood (probability of observing features X given class C),
- P(C) is the **prior probability** of class C,
- P(X) is the **evidence** (probability of observing feature set X across all classes).

To classify a new instance, Naïve Bayes calculates P(C| X) for each class and assigns the class with the highest probability.

Simplification (Independence Assumption)

Naïve Bayes assumes that all features in $X=\{x1,x2,...,xn\}$ are conditionally independent, which simplifies the probability calculation:

$$P(X|C) = P(x_1|C) \times P(x_2|C) \times \cdots \times P(x_n|C)$$

This allows for efficient computation even with many features.

Types of Naïve Bayes Classifiers

1. Gaussian Naïve Bayes:

- Assumes that the continuous features follow a Gaussian (normal) distribution.
- o Often used when dealing with continuous data.

2. Multinomial Naïve Bayes:

- o Commonly used in text classification.
- Assumes that feature occurrences are counts (e.g., word counts in documents).

3. Bernoulli Naïve Bayes:

- O Used for binary features.
- Commonly applied in tasks where features represent binary indicators (e.g., presence or absence of a word in text).

Example of Naïve Bayes Classification

Imagine a spam detection system where we have two classes: **Spam** and **Not Spam**. Given a new email, we want to classify it based on the presence of certain words (features).

1. Calculate the Prior Probability:

P(Spam): Probability of an email being spam based on past data.
P(Not Spam): Probability of an email not being spam.

2. Calculate the Likelihood:

 For each word in the new email, compute P(word| Spam) and P(word| Not Spam), representing the probability of the word appearing in each class.

3. Apply Bayes' Theorem:

 Calculate P(Spam|Email) and P(Not Spam|Email) using the independence assumption.

4. Classify:

 Choose the class with the highest probability as the predicted class for the email.

Advantages of Naïve Bayes

- **Simple and Fast**: Naïve Bayes is computationally efficient and easy to implement, making it suitable for large datasets.
- **Performs Well with Small Data**: With a limited amount of data, Naïve Bayes can perform surprisingly well, as it doesn't require complex parameter tuning.
- Handles Multiple Classes: It works effectively with multiple classes, unlike some algorithms that are designed for binary classification.

Disadvantages of Naïve Bayes

- Independence Assumption: The assumption that features are independent is often unrealistic and can limit model performance.
- Zero-Frequency Issue: If a feature value wasn't present in the training data for a class, the likelihood for that class becomes zero, which can cause issues. This is typically handled with **smoothing techniques** (e.g., Laplace smoothing).
- Limited in Real-World Data: For complex data with correlated features, Naïve Bayes may not perform as well as more sophisticated models.

Applications of Naïve Bayes

- **Spam Detection**: Naïve Bayes is used to classify emails as spam or not spam based on word frequencies and patterns.
- **Sentiment Analysis**: Useful in classifying text data (e.g., positive or negative sentiments in reviews).
- **Document Classification**: Widely used in categorizing documents or news articles by topic.
- **Medical Diagnosis**: Can assist in diagnosing diseases based on symptoms due to its probabilistic nature.

Decision Trees

Decision Trees are a supervised learning algorithm used for **classification** and **regression** tasks. The model works by splitting the dataset into smaller subsets based on feature values, resulting in a tree structure where each branch represents a

decision rule, and each leaf represents an outcome or prediction. Decision Trees are widely used because they are simple to understand, interpret, and visualize.

How Decision Trees Work

A Decision Tree algorithm splits data into subsets based on the values of input features. The splits are chosen to maximize the separation of classes or reduce prediction error in regression.

- 1. **Start with the Root Node**: The algorithm begins with the root node, representing the entire dataset. It then looks for the best feature to split the data, which minimizes classification or prediction error.
- 2. Choose the Best Split:
 - For classification tasks, the best split is chosen to maximize **purity**, meaning each subset of data should contain mostly one class.
 - For regression tasks, splits are chosen to minimize the **variance** of target values in each subset.
- 3. **Repeat Splitting**: This process continues for each resulting subset, recursively splitting nodes until the stopping criteria are met (e.g., maximum tree depth, minimum node size, or purity level).
- 4. Leaf Nodes: Once splitting stops, each final subset (leaf node) represents a decision or prediction. In classification, the leaf node represents the predicted class; in regression, it represents the average of the target values in the subset.

Types of Decision Trees

- 1. Classification Trees:
 - Used for categorical outcomes.
 - Each leaf represents a class label, and each split represents a rule based on feature values.
 - For example, a tree might classify an email as **spam** or **not spam**.

2. Regression Trees:

- o Used for continuous outcomes.
- Each leaf represents the average value of the target variable for that subset of data.
- For example, a tree might predict house prices based on features like size and location.

Example of a Decision Tree for Classification

Imagine a decision tree for classifying whether a customer will buy a product based on features like **Age** and **Income Level**:

- 1. **Root Node**: Start with the entire dataset and choose a feature that best splits it, such as **Income Level** (e.g., High, Medium, Low).
- 2. First Split: Split on Income Level.
 - For High Income, classify directly as Buy.
 - For **Medium Income**, check further conditions.
- Second Split: For Medium Income, split based on Age (e.g., Age < 30, Age >= 30).
 - Continue until all conditions are met, creating final leaf nodes where each customer is classified as either **Buy** or **Not Buy**.

Advantages of Decision Trees

- **Easy to Understand and Interpret**: Decision Trees closely mirror human decision-making processes, making them easy to understand and explain.
- **Nonlinear Relationships**: Decision Trees can capture complex, nonlinear relationships in data.
- **No Need for Feature Scaling**: Unlike some other algorithms, Decision Trees don't require normalization or scaling of input features.
- Handles Both Numerical and Categorical Data: Decision Trees work with both types of data, making them versatile.

Disadvantages of Decision Trees

- **Prone to Overfitting**: Decision Trees can grow very deep, leading to overfitting, especially if the tree is not pruned or limited in depth.
- **Instability**: Small changes in data can lead to very different tree structures, making them sensitive to data variations.
- Less Accurate Compared to Ensemble Methods: Standalone Decision Trees may not be as accurate as ensemble methods like Random Forest or Gradient Boosting.

Applications of Decision Trees

- **Customer Segmentation**: Grouping customers by purchasing behavior or demographic data.
- Medical Diagnosis: Assisting in diagnosing conditions based on symptoms or test results.
- **Loan Approval**: Determining whether to approve a loan based on applicant features.

• **Fraud Detection**: Identifying fraudulent transactions by analyzing transaction characteristics.

Decision Tree Variants

- **Pruning**: Pruning reduces the complexity of a tree by removing branches that have little importance, reducing overfitting.
- Ensemble Methods: Algorithms like Random Forests and Gradient Boosted Trees combine multiple Decision Trees to improve accuracy and reduce overfitting.

Linear Regression

Linear Regression is a simple yet powerful algorithm in machine learning and statistics used to model the relationship between a **dependent variable (target)** and one or more **independent variables (features)**. The main objective of linear regression is to find the best-fit line that represents the relationship between the variables, allowing us to make predictions for the target variable.

Types of Linear Regression

1. Simple Linear Regression:

• Models the relationship between a single independent variable and the dependent variable.

• The best-fit line is defined by the equation:

y = mx + b

where y is the predicted output, m is the slope of the line (the coefficient of x), and b is the intercept (the value of y when x=0).

2. Multiple Linear Regression:

- Models the relationship between two or more independent variables and the dependent variable.
- o The equation for multiple linear regression becomes:

 $y = b_0 + b_1 x_1 + b_2 x_2 + \dots + b_n x_n$

intercept, and b1,b2,...,bn are the coefficients for each independent variable.

Objective of Linear Regression