

	Decision tree	Naïve Bayesian Classifier	Logistic Regression	KNN classifier
Input Parameter	The attributes of a tuple	The posterior probabilities of Hypothesis H based on additional information	The measure of the total contribution of all independent variables used in the model.	A new unknown tuple for which a class has to be assigned.
Principle	<p>The attributes of a tuple are tested against the decision tree and a path is traced from the root to a leaf node which holds the prediction for that tuple</p>	<p>Given a tuple X, the classifier will predict that X belongs to the class having the highest posterior probability conditioned on X.</p>	<p>Given x representing the exposure to some set of risk factors, LR predicts the probability of occurrence of an event by fitting data to a logistic curve, $f(x)$, which represents the probability of a particular outcome</p>	<p>Nearest-neighbor classifiers compare a given test tuple with training tuples that are similar and described by n attributes and are stored in n-dimensional space</p> <p>➔ Find the k-nearest tuples from the training set to the unknown tuple</p>
ALG. FORMULA	<p>At start, all the training tuples are at the root. Then, tuples are partitioned recursively based on selected attributes. The test attributes are selected on the basis of a heuristic or statistical measure (e.g., information gain). We stop when all samples for a given node belong to the same class or there are no remaining attributes for further partitioning ➔ majority voting.</p>	<p>At start, compute $P(C)$ The prior probability of C_i. Each class can be computed based on the training tuples. Then compute each independent probability for attribute x_i in reference with Class C and multiply $P(x_i C)$. Example: The probability of class C to be yes is $P(C) = 9/14$. The amount of attribute x_i being test with respect to C is $P(x_i C_i) = \frac{x_i=test}{C_i=yes} = \frac{2}{9}$</p> <p>Finally, compute $P(X C_i)P(C_i)$ for each class ➔ The naïve Bayesian Classifier predicts $C=yes$ for tuple X</p>	<p>Since x is defined as $x = \beta_0 + \beta_1 X_1 + \dots + \beta_k X_k$ the LR is</p> $\frac{e^{\beta_0 + \beta_1 X_1 + \dots + \beta_k X_k}}{1 + e^{\beta_0 + \beta_1 X_1 + \dots + \beta_k X_k}}$ <p>Estimate the parameters using the Maximum Likelihood Function and then by computing the partial derivatives of the log likelihood, equate each partial derivative to zero, and solve the resulting nonlinear equations</p>	<p>The ALG assigns for the unknown tuple the most common class among its k-nearest neighbor. When $k=1$ the unknown tuple is assigned the class of the training tuple that is closest to it. To measure the distance we can use the Euclidean distance $\sqrt{\sum_{i=1}^n (x_{1i} - x_{2i})^2}$</p> <p>Choose K: If $k=1$ the classification will be 1:1 sensitive to other data.</p> <p>If $k=n$ we'll suffer high noise.</p> <p>Go through various K's and choose one giving lowest misclassification error!</p>