

# Machine Learning

In Son Zeng

University of Michigan

*insonz@umich.edu*

February 6, 2019

# Overview

- 1 Concept overview
  - Pattern Recognition
  - Data Mining
- 2 Concept overview
  - PCA and FA

# Pattern Recognition: Introduction

- 1 Pattern Recognition is the recognition of patterns and regularities in data. It is closely related to artificial intelligence, machine learning (ML), data mining and knowledge discovery in databases (KDD).
- 2 Machine Learning is one approach to pattern recognition. Some other methods for pattern recognition involve hand-crafted but not learned rules and heuristics.
- 3 Pattern Recognition is a path to artificial intelligence. Some other approaches to artificial intelligence include symbolic artificial intelligence.
- 4 Definition: The field of pattern recognition is concerned with the automatic discovery of regularities in data through the use of computer algorithms and with the use of these regularities to take actions such as classifying the data into different categories.

# Pattern Recognition: Introduction

- 1 Pattern recognition systems are often trained from labeled "training data".

# Data Mining

- ➊ Data mining is a process to structure the raw data and formulate or recognize the various patterns in the data through the mathematical and computational algorithms.
- ➋ Data mining helps to generate new information and unearth various insights in a data.
- ➌ Data mining has two types: 1) Descriptive - give information about existing data about the structure, 2) Predictive - make forecasts based on the data.
- ➍ Data mining also helps explore trends from the data (one kind of pattern discovery task). Therefore, it requires advanced components of artificial intelligence, pattern distribution and statistics.

# Data Analytics

- 1 Data Analytics is a superset of data mining. It involves extracting, cleaning, transforming, modeling and visualization of data with an intention to uncover meaningful and useful information that can help in deriving conclusion and take decisions.
- 2 Data analytics is the process of ordering and organizing raw data to determine useful insights and decisions.
- 3 Data analytics requires the knowledge of CS, statistics, maths, AI/Machine learning and subject knowledge.
- 4 Data analytics is of several types: exploratory, descriptive, text analytics, predictive analysis, data mining
- 5 Data analytics can work on both structured, semi-structured and unstructured data

# Data Mining vs Data Analytics Responsibilities

- ➊ A Data mining specialist usually builds algorithms to identify meaningful structure in the data.
- ➋ A data analyst usually cannot be a single person. The job profile involves preparation of raw data, its cleansing, transformation and modeling and finally its presentation in the form of chart/non-chart-based visualizations.
- ➌ Data mining is responsible for extracting and discovering meaningful patterns and structure in the data.
- ➍ Data analytics is responsible for developing models, explanations, testing and proposing hypotheses using analytical tools and methods.
- ➎ The output of a data mining task is a data pattern.
- ➏ The output of data analysis is a varied hypothesis or insight on the data.

# Natural Language Processing: Problems

- 1 Case-sensitivity and abbreviation problems are common in electronic health records (EHRs) and patient history records. We need strong NLP system to make corresponding match of abbreviations to the context.
- 2 NLP helps transform the unstructured data to structured database, which in turn increases the interoperability of EHR and patient data. Interoperability, the level of clarity of delivering message through healthcare stakeholders, is crucial for providing effective healthcare.
- 3



# Fixed Effect Model

- 1  $y_i = \alpha + X \cdot \beta$
- 2 Each individual has same regression effects.
- 3 We add a random effect for each FIPS because regression effects cannot represent the housing price disparity by FIPS (geographical difference)
- 4 For each of the FIPS (county), we have an additional intercept called **random intercept**. We already have regression (fixed effect), this random effect gives additional information based on the geographical differences.
- 5 Now we get each random intercept  $\alpha_i \cdot 1(FIPS = i)$ . And we also get the output **standard deviation/variance** across all random effects. If we would like to cover as much as the variability of geographical differences, say 95% or 99.5 %, we need to plus and minus 2 or 3 times the standard deviation.

# PCA Functions

- 1 R: `prcomp()`, `princomp()`: return all parts of PCA
- 2 R: `loadings()`: display variable loading for first  $n$  PCs.
- 3 R: `eigen()`, `svd()`: perform eigen-decomposition, SVD of any matrix
- 4 R: `biplot()`, `ggbiplot()`: make a PCA biplot out of an object created by `prcomp`
- 5 R package `factoextra`: `fviz_fig()` - nice formatted scree plot; `fviz_pca_var()`: variable plot; `fviz_pca_biplot()`: biplot
- 6 We use PCA to find the largest  $p$  eigenvalues and eigenvectors of  $S$  to capture 95% variability, reducing the dimension.

# PCA Functions

- ① Goal: Project a vector in  $q$  dimensions to  $p$  dimensions, where  $p \ll q$ .
- ② Notation:  $y_{q \times 1} \approx u_1 x_1 + \dots + u_p x_p$ , where  $u^T u = I_{p \times p}$ , orthonormal.
- ③ We find  $u_1$  such that the empirical variance of  $x_1, \dots, x_N$  given by  $\frac{1}{N-1} \sum_{i=1}^N (u_1^T y_i - u_1^T \bar{y})^2$  is maximized.
- ④ Algorithm:
  - 1) Maximize  $u_1^T S u_1$ ,  $u_1 \in R^q$ , such that  $u_1^T u_1 = 1$ .
  - 2) Introduce Lagrange Multiplier  $\lambda$  so  $\max u_1^T S u_1 - \lambda(u_1^T u_1 - 1) \rightarrow S u_1 = \lambda u_1$
  - 3) Therefore,  $\lambda$  is an eigenvalue of  $S$ , and  $u_1$  is the corresponding eigenvector of  $S$ .
  - 4) We in short find  $\max \lambda$ , the largest eigenvalue of  $S$ , with  $u_1$  being the corresponding eigenvector.

# FA Functions

- ① R:
- ② Factors are **latent**.
- ③ Goal: Find  $\hat{\Sigma} \approx \hat{\Lambda}\hat{\Lambda}^T + \hat{\Phi}$ , one approach is  $\min \|\hat{\Sigma} - \hat{\Lambda}\hat{\Lambda}^T + \hat{\Phi}\|_F$
- ④ Algorithm: Start with initial estimate of  $\hat{\Phi}$ , we iterate until convergence
  - 1) Apply PCA to  $\hat{\Sigma} - \hat{\Phi}$ , K largeest eigenvalues to obtain  $\hat{\Lambda}$
  - 2) Update  $\hat{\Phi} = \text{diag}(\hat{\Sigma} - \hat{\Lambda}\hat{\Lambda}^T)$
- ⑤ Algorithm 2: MLE method through EM algorithm.

# Important FA Detail

- ① Assume all  $p$  variables are quantitative. Assume  $X$  satisfies  $X_j = \sum_{k=1}^K \lambda_{jk} F_k + u_j$ . Matrix:  $X_{p \times 1} = \Lambda_{p \times K} F_{K \times 1} + U_{p \times 1}$
- ②  $\lambda_{jk}$  is factor loading;  $\Lambda, F, U$  are unobserved.
- ③ Assumptions: 1)  $F$  and  $U$  are independent; 2)  $E(F) = E(U) = 0$ ; 3)  $Cov(F) = I$ ; 4)  $Cov(U) = \Phi_{p \times p}$
- ④ Covariance decomposition: Let  $\Sigma = Var(X) \rightarrow \Sigma = \Lambda \Lambda^T + \Phi$
- ⑤ Diagonal terms of  $\Sigma$  are  $Var(X_j) = \sum_{k=1}^K \lambda_{jk}^2 = \Phi_{jj} = \text{communality} + \text{uniqueness}$
- ⑥ Degree of freedom:  $d = \frac{p(p+1)}{2} - (pK + p - \frac{K(K-1)}{2})$ . If  $d < 0$ , the model cannot be fitted

# Important FA Detail

- ① Factor rotations: For any orthogonal matrix  $T$ ,  
$$\Sigma = \Lambda \Lambda^T + \Phi = (\Lambda T)(\Lambda T)^T + \Phi$$
- ② EM algorithm: 1) The E-step , 2) The M-step updates the variable based on maximized likelihood (local)

# Multi-dimensional Scaling

- ① R: `cmdscale()`;
- ② Double-centering trick:  $G = -\frac{1}{2}(I_n - \frac{1}{n}\mathbf{1}\mathbf{1}^T)D(I_n - \frac{1}{n}\mathbf{1}\mathbf{1}^T)$
- ③ Distance functions:
  - 1) R: `gower.dist()` -
  - 2) R:

# Examples for classification

- ➊ Identifying tumor cells as benign or malignant
- ➋ Classifying credit card transactions as legitimate or fraudulent
- ➌ Classifying secondary structures of protein as alpha-helix, beta-sheet or random coil
- ➍ Categorizing news stories as finance, weather, or sports, etc.
- ➎ Filtering spams emails out of user's mailbox
- ➏ Recognizing handwritten letters and digits as one of A,B,C...,0,1,...,9
- ➐ Recognizing animal pictures: is this a dog? is this a koala?
- ➑ Assign people to either guilty or crime-free.
- ➒ We need to find **decision rule** to make prediction and draw **boundary** to separate the classes as accurate as possible.



# Examples of visualization for classification

- ① R: `require(ggplot2); ggpairs()` - Scatter plot for data
- ② R: `require(GGally); ggparcoord()` - Parallel coordinates plot for data.  
This works well if the number of class is small, but poorly if the number of class is large.
- ③

# Methods to Split Test and Train Dataset

- ① For example, we want 75% of the sample size to be training set.
- ② R Method 1:
  - 1) `smpsize = floor(0.75 * nrow(data)); set.seed(123)`
  - 2) `trainindex = sample(seq_len(nrow(data)), size = smpsize)`
  - 3) `train = data[trainindex, ]; test = data[-trainindex, ]`
- ③ R Method 2:

# Cross Validation (CV)

- 1 Divide the data into  $V$  parts ( $V$ -fold cross-validation).  $V$  can be anything; popular values are 5, 10, 15 and  $n$ .
- 2 Hold out one of the parts to be the **test data**.
- 3 Use all remaining data as **training data** to develop a classification (clustering) rule.
- 4 Apply the classification rule to the held-out part and compute the error rate (**test error**).
- 5 Repeat (using for loop) steps 2-4 for each of the  $V$  parts to ensure that each partition of the data become the test data once
- 6 Average all test errors to obtain the **cross-validation error**.
- 7 The cross-validation error tends to be closer to the **true error rate** than to the apparent error rate???
- 8 The computational cost of CV can become a concern, particularly if there are many **tuning parameters**.

# Brute force approach

1

2

# Generative models for classification

- 1 If we know  $P(X,Y)$ , we are done. In modeling, we need 1) Prior class prob  $Y \sim \pi$  and class conditional distributions  $X|Y = k \sim p_k$

- 2 Parametric model:

$$p_k(x) := p_k(x|\theta_k) \rightarrow \hat{p}(Y = k|X = x) \propto \hat{\pi}_k p_k(x|\hat{\theta}_k)$$

- 3 Suppose  $X|Y \sim p_k = N(\mu_k, \sigma_k)$

- 4 Therefore, the optimal classifier is the maximum of  $k$  classes for  $f(x)$

$$f(x) = \operatorname{argmax}_k P(Y = k) \cdot p_k(x) = \operatorname{argmax}_k \pi_k \cdot N(x|\mu_k, \Sigma_k)$$

Take log we get QDA rule:

$$\operatorname{argmax}_k \delta_k = \operatorname{argmax}_k \left( \log(\pi_k) - \frac{1}{2} \log|\Sigma_k| - \frac{1}{2} (x - \mu_k)^T \Sigma_k^{-1} (x - \mu_k) \right)$$

- 5 For LDA (assume equal  $\Sigma$ ), the function is simplified as

$$\delta_k = \left( \log(\pi_k) - \frac{1}{2} \mu_k^T \Sigma^{-1} \mu_k - x^T \Sigma^{-1} \mu_k \right), \text{ with the quadratic term eliminated.}$$

# Log-Likelihood for Generative Model

- ① For Latent Variable Models, the log-likelihood for  $Y_1, Y_2, \dots, Y_N$  is

$$L(\theta|Y) = \sum_{i=1}^N \log P(Y_i|\theta) = \sum_{i=1}^N \log \int_{x_i} p(x_i, y_i|\theta) dx_i$$

, where  $p(x_i, y_i|\theta) = p(x_i|\theta) \cdot p(y_i|x_i, \theta)$

- ② Example:  $i = 1, 2, \dots, N$  and  $Y_i = (Y_{i1}, \dots, Y_{iq})^T \in \{0, 1\}^q$ , and the model is  $x_i \sim N_p(0, I)$

- ③ For Factor Analysis,  $\log \frac{p(Y_{ij}=1|x_i, \theta)}{p(Y_{ij}=0|x_i, \theta)} = \Lambda_j^T x_i + \mu_j$ , and  $Y_{ij}, j = 1, \dots, q$  are conditionally independent given  $x_i$ . Then,  
 $l(\theta, Y) = \sum_{i=1}^N \log \int_{x_i} p(x_i) p(y_i|x_i, \theta) dx_i$  and have a solution.

- ④ However, in many cases, the EM algorithm does not give analytical form of solution.

# EM algorithm

- ① EM repeats the following E and M steps: in  $(t + 1)^{th}$  step: we have  $\theta^0$  as initial value.
- ② E-step: We compute the conditional expectation of the **complete data likelihood**, which is  $\sum_{i=1}^N \log p(x_i, y_i | \theta)$ , with respect to the posterior distribution, which is  $p(X|Y, \theta^{(t)})$  of  $X$  given  $Y$  and  $\theta^{(t)}$ .
- ③ That is to say, we find

$$E_{X|Y, \theta^{(t)}} l_c(X, Y | \theta) = \sum_{i=1}^N \int_{x_i} \left( \log p(x_i, y_i | \theta) \right) \cdot p(x_i | y_i, \theta^{(t)}) dx_i$$

$$\propto \sum_{i=1}^N \int_{x_i} \log \left( \frac{p(x_i, y_i | \theta) \cdot p(x_i | \theta^{(t)}) \cdot p(y_i | x_i, \theta^{(t)})}{p(y_i | \theta^{(t)})} \right) dx_i$$

- ④ Remark:  $p(x_i | y_i, \theta)$  may not have closed form. We may use Monte Carlo Method to get  $E_{x|Y, \theta^{(t)}}$

# EM algorithm

- 1 Suppose we have the expectation step done; we obtain the conditional expectation of  $\sum_{i=1}^N \log p(x_i, y_i | \theta)$  w.r.t to the posterior distribution. Then we perform M-step to update the latent variable  $\theta^{(t+1)}$ .
- 2 M-step: Update  $\theta^{(t+1)}$  by  $\operatorname{argmax}_{\theta} E_{X|Y, \theta^{(t)}} l(X, Y | \theta)$
- 3 Factor Analysis Example:

$$\text{For } i = 1, \dots, N, \begin{pmatrix} x_i \\ y_i \end{pmatrix} \sim N \left( \begin{pmatrix} 0 \\ \mu \end{pmatrix}, \begin{pmatrix} I, \Lambda^T \\ \Lambda, \Lambda \Lambda^T + E \end{pmatrix} \right)$$

- 4 Therefore,  

$$l_c(\theta | X, Y) = \sum_{i=1}^N \log p(x_i y_i | \theta) = \sum_{i=1}^N \log p(y_i | x_i | \theta) + C = -\frac{N}{2} \log |E| - \frac{1}{2} \sum_{i=1}^N (y_i - \mu)^T \Psi^{-1} (y_i - \mu) + C$$
- 5 Then, take Expectation to  $l_c(\theta | X, Y)$ , making  $x_i | y_i, \theta \sim N(E[x_i | y_i, \theta], \operatorname{Var}(x_i | y_i, \theta))$ , where  

$$E[x_i | y_i, \theta] = \Lambda^T (\Lambda \Lambda^T + \Phi)^{-1} (y_i - \mu) \text{ and}$$

$$\operatorname{Var}[x_i | y_i, \theta] = I - \Lambda^T (\Lambda \Lambda^T + \Phi)^{-1} \Lambda$$



# EM algorithm

- 1 Take Expectation to  $l_c(\theta|X, Y)$ , making  
 $x_i|y_i, \theta \sim N\left(E[x_i|y_i, \theta], \text{Var}(x_i|y_i, \theta)\right)$ , where  
 $E[x_i|y_i, \theta] = \Lambda^T(\Lambda\Lambda^T + \Phi)^{-1}(y_i - \mu) = (I_p + \Lambda^T\Phi\Lambda)^{-1}_{p \times p} \Lambda^T\Phi^{-1}(y_i - \mu)$   
 and  $\text{Var}[x_i|y_i, \theta] = I - \Lambda^T(\Lambda\Lambda^T + \Phi)^{-1} = (I_p + \Lambda^T\Phi^{-1}\Lambda)^{-1}_{p \times p}$
- 2 E-step:  $E_{X|Y, \theta^{(t)}} l_c(\theta|X, Y) = -\frac{N}{2} \log|\Phi| - \frac{N}{2} \text{Tr}\left(\Phi^{-1} E_{X|Y, \theta}[S]\right)$
- 3 Assume that  $\mu = 0$ , then  $S = \frac{1}{N} \sum_{i=1}^N (y_i - \Lambda x_i)(y_i - \Lambda x_i)^T$ , which leads to

$$E_{X|Y, \theta}[S] = \frac{1}{N} \sum_{i=1}^N y_i y_i^T - 2 y_i E[x_i|y_i, \theta^{(t)}]^T \Lambda^T + \Lambda E[x_i x_i^T | y_i, \theta^{(t)}] \Lambda^T$$

# EM algorithm

- ① Finally M-step:

$$\theta = (\Lambda, \Phi) = \left\{ \begin{array}{l} \frac{\partial E_{X|Y, \theta^{(t)}} l_c(\theta|X, Y)}{\partial \Lambda} = 0 \\ \frac{\partial E_{X|Y, \theta^{(t)}} l_c(\theta|X, Y)}{\partial \Phi^{-1}} = 0 \end{array} \right\}$$

- ② (1)  $0 = \sum_{i=1}^N (y_i E[x_i | y_i, \theta^{(t)}] - \Lambda E[x_i x_i^T | y_i, \theta^{(t)}])$ . Therefore,  
 $\hat{\Lambda}^{(t+1)} \rightarrow \sum_{i=1}^N y_i E[x_i^T | y_i, \theta^{(t)}] \left( \sum_{i=1}^N E[x_i x_i^T | y_i, \theta^{(t)}] \right)^{-1}$
- ③ (2) Diagonal matrix.

# Logistic Regression

- 1 Maximize the conditional likelihood using multinomial likelihood with probability  $P(Y = k|X)$
- 2 R: `glm(formula, family = binomial)` - logistic regression
- 3 R: `require(nnet); multinom()` - perform multinomial logistic regression
- 4 R: `update(model, . -var)` - delete variables
- 5 R: `predict(model, type="response")` - predict accuracy in training data; `ifelse()` to separate response class
- 6 R: `predict(model, newdata= test, type="response")` - predict accuracy in test data
- 7 Logistic regression does not work properly if the response classes are fully separated from each other

# Linear Discriminant Analysis

- ① Assumptions: 1) Equal Variance for every class, 2) data from each class follow normal distribution  $N(\mu_k, \Sigma^2)$
- ② Maximize the full joint likelihood based on the joint density  
$$P(x, Y = k) = \pi_k \cdot \phi(x; \mu_k, \Sigma)$$
- ③ R: `require(MASS); qda()`: perform LDA
- ④ R: `predict(lda.model, data=train)` - predict accuracy in training data
- ⑤ R: `predict(lda.model, newdata=test)` - predict accuracy in test data
- ⑥ LDA discriminant direction vector is  $\Sigma^{-1}(\mu_k - \mu_l)$ , the only term including data  $x$ .
- ⑦ Mahalanobis distance:  $d_M(x, \mu_k)^2 = (x - \mu_k)^T \Sigma^{-1}(x - \mu_k)$
- ⑧ LDA is not robust, easily influenced by outliers; Logistic regression is more robust. LDA has more assumptions than Logistic regression.

# Quadratic Discriminant Analysis

- ① Assumptions: 1) Different Variance for every class, 2) data from each class follow normal distribution  $N(\mu_k, \Sigma_k^2)$
- ② R: `require(MASS); lda()`: perform LDA
- ③ R: `predict(lda.model, data=train)` - predict accuracy in training data
- ④ R: `predict(lda.model, newdata=test)` - predict accuracy in test data
- ⑤ R: `require(caret); train(); confusionMatrix()` - Derive the confusion matrix for accuracy

# K nearest neighbor classifier (KNN)

- 1 R: `require(class); knn()`: Perform k-nearest neighbour classification for test set from training set. For each row of the test set, the k nearest (in Euclidean distance) training set vectors are found, and the classification is decided by majority vote, with ties broken at random. If there are ties for the kth nearest vector, all candidates are included in the vote.
- 2 R: `knn.cv()` - Perform k-nearest neighbour cross-validatory classification from training set. This uses leave-one-out cross validation.
- 3 k-NN is an adapted kernel method, a kernel method depending on the data. In the nearest neighbor kernel, K is constant over the nearest k data points, and zero otherwise. As such, k-NN is a local method which relies on the rule not changing too fast in a neighborhood.

# K nearest neighbor classifier (KNN)

- 1 Algorithm: Given  $n$  data pairs  $(x_i, y_i)_{i=1}^n$ , we first 1) determine the  $k$  nearest neighbours of the new point  $x$  among the  $x_i$  by (default) using Euclidean distance. Then, 2) we assign  $x$  to the class corresponding to the majority votes of the  $k$  nearest neighbours. We often choose odd  $k$  to avoid ties.
- 2 Strength: Euclidean distance is convenient and popular for quantitative predictors;  $k$ -NN has no parametric assumption.
- 3 Shortcoming: Curse of dimensionality:  $k$ -NN fails in high dimensions since data points will be very sparsely population in high dimensions. In other words, it is difficult to find enough observations close to a target point  $x$ .
- 4 Indeed, since the volume of  $d$ -dimensional object with same length  $l$  is  $\rho = l^d \rightarrow l = \rho^{1/d}$

# K nearest neighbor classifier (KNN)

- 1 Distance function: Euclidean distance is popular, but weighted distances have been widely used to emphasize some variables over others.
- 2 We may or may not have an explicit classification rule in k-NN.
- 3 k-NN approximates the conditional probability  $P(Y = c|X = x)$  directly. When  $n, k \rightarrow \infty, \frac{k}{n} \rightarrow 0$ , then k-NN estimate is consistent.
- 4 The effective number of parameters is  $\frac{n}{k}$ . So, the degree of freedom of k-NN highly depends on the choice of k (the smaller k, the more complex the model).
- 5 When  $k = 1$ , the training error is 0, but it suffers from overfitting.
- 6 We should choose k to minimize the mis-classification error on a separate data set, or use CV. Ideally, we use validation data to choose the k minimizing the error, and take an independent test set to estimate test error.




# How KNN works (Example)

- ① Step 1 : Calculate Similarity based on distance function
- ② Euclidean Distance:  $d(x, y) = \sqrt{\sum_{i=1}^m (x_i - y_i)^2}$
- ③ Manhattan Distance:  $d(x, y) = \sum_{i=1}^m |x_i - y_i|$


# Support Vector Machine (SVM)

- 1 R: `library(e1071); svm()` - perform support vector machine
- 2 R:
- 3 SVM works for linear data, with hyperplane separating data with different class, with a maximized margin of  $\frac{2}{||\beta||}$ . Graphical interpretation of SVM objective:




# Support Vector Machine (SVM)

- 1 SVM also works for data with clearly segregated vectors. We can add a third dimension or even more. For example, the next example shows  $z = 1 = x^2 + y^2$  is the best separating hyperplane. Again, the margin is still maximized as  $\frac{2}{\|\beta\|}$ .
- 2 By mapping our space to a higher dimension, we can classify nonlinear data using SVM.



3-Dimensional Representation



2-Dimensional Representation

# Support Vector Machine (SVM)

- ① Calculating transformation function is computationally intensive. SVM does not need the actual vectors to work, so **Kernel function** is useful here.
- ② Common types of kernel functions to separate non-linear data:
  - 1) Polynomial kernels
  - 2) Radial basis kernels
  - 3) Linear kernels

# Support Vector Machine (SVM)

## 1 Advantages:

- High dimensionality: SVM is effective tool in high-dimensional space, useful for **document classification** and **sentiment analysis**, where dimensionality can be extremely large
- Memory Efficiency: Since only a subset of the training points are used in the actual decision process of assigning new members, just there points need to be stored in memory when making decisions
- Versatility: The ability to apply new kernels allows substantial flexibility for the decision boundaries, leading to greater classification performance

## 2 Disadvantages:

- Kernel Parameters Selection: SVMs are very sensitive to the choice of the kernel parameters. In situations where the number of features for each object exceeds the number of training data samples, SVMs can perform poorly.
- Non-Probabilistic: Since the classifier works by placing objects

# Support Vector Machine (SVM)

1

# Decision Tree

- 1 R: `library(rpart); rpart()` - build decision tree
- 2 R: `library(rpart.plot); rpart.plot()` - plot decision tree output
- 3 R: `library(party); ctree(); plot()` - build and plot decision tree
- 4 R: `rpart.control(minsplit, minbucket, maxdepth)` - Tune the maximum depth, minimum number of sample a node must have before it can split, and minimum number of sample a leaf node must have
- 5 Gini impurity index:

# Random Forest (RF)

- 1 R: `library(randomForest); randomForest()`
- 2 `importance(); varImpPlot()` - check variable importance
- 3