Chapter 1:




Observations


| Distance | Time |
| :---: | :---: |
| 4.9 m | 1 s |
| 19.6 m | 2 s |
| 44.1 m | 3 s |
| 78.5 m | 4 s |

$$
g=9.8 \mathrm{~m} / \mathrm{s}^{2}
$$



| features |  |  |  |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
|  |  |  |  |  |  |
| $\bigcirc$ |  |  |  |  | 1 |
| year | model | price | mileage | color | transmission |
| 2011 | SEL | 21992 | 7413 | Yellow | AUTO |
| 2011 | SEL | 20995 | 10926 | Gray | AUTO |
| 2011 | SEL | 19995 | 7351 | Silver | AUTO |
| 2011 | SEL | 17809 | 11613 | Gray | AUTO |
| 2012 | SE | 17500 | 8367 | White | MANUAL |
| 2010 | SEL | 17495 | 25125 | Silver | AUTO |
| 2011 | SEL | 17000 | 27393 | Blue | AUTO |
| 2010 | SEL | 16995 | 21026 | Silver | AUTO |
| 2011 | SES | 16995 | 32655 | Silver | AUTO |

## Chapter 2:



Histogram of Used Car Prices


Histogram of Used Car Mileage



Right Skew


No Skew


Left Skew


## Uniform Distribution


$\operatorname{Var}(\mathrm{X})=\sigma^{2}=\frac{1}{n} \sum_{i=1}^{n}\left(x_{i}-\mu\right)^{2}$

## $\operatorname{StdDev}(\mathrm{X})=\sigma=\sqrt{\frac{1}{n} \sum_{i=1}^{n}\left(x_{i}-\mu\right)^{2}}$




Unimodal Distribution


Bimodal Distribution

Scatterplot of Price vs. Mileage


## Cell Contents



Total Observations in Table: 150


## Chapter 3:



how sweet the food tastes

how sweet the food tastes
$\operatorname{dist}(p, q)=\sqrt{\left(p_{1}-q_{1}\right)^{2}+\left(p_{2}-q_{2}\right)^{2}+\ldots+\left(p_{n}-q_{n}\right)^{2}}$
$\operatorname{dist}($ tomato, green bean $)=\sqrt{(6-3)^{2}+(4-7)^{2}}=4.2$

> Larger k
> Smaller k
> $\mathrm{X}_{\text {new }}=\frac{\mathrm{X}-\min (\mathrm{X})}{\max (\mathrm{X})-\min (\mathrm{X})}$
> $\mathrm{X}_{\text {new }}=\frac{\mathrm{X}-\mu}{\sigma}=\frac{\mathrm{X}-\operatorname{Mean}(\mathrm{X})}{\mathrm{StdDev}(\mathrm{X})}$
> male $= \begin{cases}1 & \text { if } x=\text { male } \\ 0 & \text { otherwise }\end{cases}$

## hot $= \begin{cases}1 & \text { if } x=\text { hot } \\ 0 & \text { otherwise }\end{cases}$

## medium $=\left\{\begin{array}{l}1 \\ 0\end{array}\right.$ if $\mathrm{x}=$ medium otherwise



## kNN classification syntax

using the knn () function in the class package

## Building the classifier and making predictions:

$\mathrm{p}<-\mathrm{knn}($ train, test, class, k)

- train is a data frame containing numeric training data
- test is a data frame containing numeric test data
- class is a factor vector with the class for each row in the training data
- $\mathbf{k}$ is an integer indicating the number of nearest neighbors

The function returns a factor vector of predicted classes for each row in the test data frame.

## Example:

```
wbcd_pred <- knn(train = wbcd_train, test = wbcd_test,
    c1 = wbcd_train_1abels, k = 3)
```



Chapter 4:

## all email



ham<br>(80\%)








|  | Viagra |  | Total |  | Viagra |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Frequency | Yes | No |  | Likelihood | Yes | No | Total |
| spam | 4 | 16 | 20 | spam | 4 / 20 | 16 / 20 | 20 |
| ham | 1 | 79 | 80 | ham | $1 / 80$ | $79 / 80$ | 80 |
| Total | 5 | 95 | 100 | Total | $5 / 100$ | $95 / 100$ | 100 |


|  | Viagra ( $\mathrm{W}_{1}$ ) |  | Money ( $\mathrm{W}_{2}$ ) |  | Groceries ( $\mathrm{W}_{3}$ ) |  | Unsubscribe ( $\mathrm{W}_{4}$ ) |  |  |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| Likelihood | Yes | No | Yes | No | Yes | No | Yes | No | Total |
| spam | 4 / 20 | 16 / 20 | 10 / 20 | $10 / 20$ | 0/20 | 20/20 | 12 / 20 | $8 / 20$ | 20 |
| ham | $1 / 80$ | 79 / 80 | 14 / 80 | 66 / 80 | $8 / 80$ | 71/80 | $23 / 80$ | $57 / 80$ | 80 |
| Total | 5/100 | 95/100 | 24/100 | 76/100 | 8/100 | 91/100 | 35/100 | $65 / 100$ | 100 |

$P\left(\right.$ spam $\left.\mid W_{1} \cap \neg W_{2} \cap \neg W_{3} \cap W_{4}\right)=\frac{P\left(W_{1} \cap \neg W_{2} \cap \neg W_{3} \cap W_{4} \mid \text { spam }\right) P(\text { spam })}{P\left(W_{1} \cap \neg W_{2} \cap \neg W_{3} \cap W_{4}\right)}$
$P\left(\right.$ spam $\left.\mid W_{1} \cap \neg W_{2} \cap \neg W_{3} \cap W_{4}\right) \propto P\left(W_{1} \mid\right.$ spam $) P\left(\neg W_{2} \mid\right.$ spam $) P\left(\neg W_{3} \mid\right.$ spam $) P\left(W_{4} \mid\right.$ spam $) P($ spam $)$
$P\left(\right.$ ham $\left.\mid W_{1} \cap \neg W_{2} \cap \neg W_{3} \cap W_{4}\right) \propto P\left(W_{1} \mid\right.$ ham $) P\left(\neg W_{2} \mid\right.$ ham $) P\left(\neg W_{3} \mid\right.$ ham $) P\left(W_{4} \mid\right.$ ham $) P($ ham $)$

$$
P\left(C_{L} \mid F_{1}, \ldots, F_{n}\right)=\frac{1}{Z} p\left(C_{L}\right) \prod_{i=1}^{n} p\left(F_{i} \mid C_{L}\right)
$$



| message \# | balloon | balls | bam | bambling | band |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 0 | 0 | 0 | 0 | 0 |
| 2 | 0 | 0 | 0 | 0 | 0 |
| 3 | 0 | 0 | 0 | 0 | 0 |
| 4 | 0 | 0 | 0 | 0 | 0 |
| 5 | 0 | 0 | 0 | 0 | 0 |

 finish win pIs ask stop tell mobil even always care work $\underset{\rightarrow}{\boldsymbol{x}}$ lo one need hope really ${ }_{6}^{\circ}$ place hey still $\underset{+}{(1)}$ like KnOW miss min year take + went

 soon prize $\underset{=}{\leq}$ JUSt free for say sort help meet $\xlongequal{\leftrightarrows}$ claim wish hello end
money great $\stackrel{\oplus}{\otimes}$ gOOd tut ont $\begin{gathered}\text { said } \\ \text { the late } \\ \text { later find gonna }\end{gathered}$ nice hap dun lo


## Naive Bayes classification syntax

## using the naiveBayes () function in the e1071 package

## Building the classifier:

```
m <- naiveBayes(train, class, laplace = 0)
```

- train is a data frame or matrix containing training data
- c1 ass is a factor vector with the class for each row in the training data
- 1 ap 1 ace is a number to control the Laplace estimator (by default, 0 )

The function will return a naive Bayes model object that can be used to make predictions.

## Making predictions:

```
p <- predict(m, test, type = "class")
```

- $m$ is a model trained by the naiveBayes () function
- test is a data frame or matrix containing test data with the same features as the training data used to build the classifier
- type is either "class" or "raw" and specifies whether the predictions should be the most likely class value or the raw predicted probabilities

The function will return a vector of predicted class values or raw predicted probabilities depending upon the value of the type parameter.

## Example:

```
sms_classifier <- naiveBayes(sms_train, sms_type)
sms_predictions <- predict(sms_classifier, sms_test)
```

Total Observations in Table: 1390


Total Observations in Table: 1390


## Chapter 5:







$\operatorname{InfoGain}(F)=\operatorname{Entropy}\left(S_{1}\right)-\operatorname{Entropy}\left(S_{2}\right)$
$\operatorname{Entropy}(S)=\sum_{i=1}^{n} w_{i} \operatorname{Entropy}\left(P_{i}\right)$

## C5.0 decision tree syntax

using the C5.0 () function in the C50 package

## Building the classifier:

$m<-C 5.0(t r a i n, ~ c l a s s, ~ t r i a l s=1, ~ c o s t s=N U L L)$

- train is a data frame containing training data
- class is a factor vector with the class for each row in the training data
- trials is an optional number to control the number of boosting iterations (set to 1 by default)
- costs is an optional matrix specifying costs associated with various types of errors

The function will return a C5.0 model object that can be used to make predictions.

## Making predictions:

$p<-$ predict(m, test, type = "class")

- m is a model trained by the C5.0() function
- test is a data frame containing test data with the same features as the training data used to build the classifier.
- type is either "c1ass" or "prob" and specifies whether the predictions should be the most probable class value or the raw predicted probabilities

The function will return a vector of predicted class values or raw predicted probabilities depending upon the value of the type parameter.

## Example:

credit_mode1 <- C5.O(credit_train, 1oan_default)
credit_prediction <- predict(credit_mode1, credit_test)

```
C5.0 [Release 2.07 GPL Edition]
```

Class specified by attribute `outcome'
Read 900 cases (17 attributes) from undefined.data Decision tree:
checking_balance in \{> 200 DM, unknown\}: no (412/50) checking_balance in \{< 0 DM,1 - 200 DM\}:
:...credit_history in \{perfect, very good\}: yes (59/18) credit_history in \{critical,good, poor\}:
:...months_loan_duration <= 22:
:...credit_history = critical: no (72/14)
: credit_history = poor:
: :...dependents > 1: no (5)
: : dependents <= 1:
: : :...years_at_residence <= 3: yes (4/1)
: : years_at_residence > 3: no (5/1)





| Animal | Travels By | Has Fur | Mammal |
| :---: | :---: | :---: | :---: |
| Bats | Air | Yes | Yes |
| Bears | Land | Yes | Yes |
| Birds | Air | No | No |
| Cats | Land | Yes | Yes |
| Dogs | Land | Yes | Yes |
| Eels | Sea | No | No |
| Elephants | Land | No | Yes |
| Fish | Sea | No | No |
| Frogs | Land | No | No |
| Insects | Air | No | No |
| Pigs | Land | No | Yes |
| Rabbits | Land | Yes | Yes |
| Rats | Land | Yes | Yes |
| Rhinos | Land | No | Yes |
| Sharks | Sea | No | No |

Full Dataset

| Travels By | Predicted | Mammal |
| :---: | :---: | :---: |
| Air | No | Yes |
| Air | No | No |
| Air | No | No |
| Land | Yes | Yes |
| Land | Yes | Yes |
| Land | Yes | Yes |
| Land | Yes | Yes |
| Land | Yes | No |
| Land | Yes | Yes |
| Land | Yes | Yes |
| Land | Yes | Yes |
| Land | Yes | Yes |
| Sea | No | No |
| Sea | No | No |
| Sea | No | No |

Rule for "Travels By"
Error Rate = 2 / 15

| Has Fur | Predicted | Mammal |
| :---: | :---: | :---: |
| No | No | No |
| No | No | No |
| No | No | Yes |
| No | No | No |
| No | No | No |
| No | No | No |
| No | No | Yes |
| No | No | Yes |
| No | No | No |
| Yes | Yes | Yes |
| Yes | Yes | Yes |
| Yes | Yes | Yes |
| Yes | Yes | Yes |
| Yes | Yes | Yes |
| Yes | Yes | Yes |

Rule for "Has Fur"
Error Rate = 3 / 15



## 1 R classification rule syntax

## using the OneR () function in the RWeka package

## Building the classifier:

```
m <- OneR(class ~ predictors, data = mydata)
```

- c1ass is the column in the mydata data frame to be predicted
- predictors is an R formula specifying the features in the mydata data frame to use for prediction
- data is the data frame in which class and predictors can be found

The function will return a 1R model object that can be used to make predictions.

## Making predictions:

p <- predict(m, test)

- $m$ is a model trained by the OneR() function
- test is a data frame containing test data with the same features as the training data used to build the classifier.

The function will return a vector of predicted class values.

## Example:

$$
\begin{array}{r}
\text { mushroom_classifier <- OneR(type } \left.\sim \begin{array}{r}
\sim \text { odor + cap_color, } \\
\text { data }
\end{array}=\text { mushroom_train }\right) \\
\text { mushroom_prediction }<- \text { predict(mushroom_classifier }, \\
\text { mushroom_test })
\end{array}
$$

## RIPPER classification rule syntax

using the JRip() function in the RWeka package

## Building the classifier:

```
m <- JRip(class ~ predictors, data = mydata)
```

- c1ass is the column in the mydata data frame to be predicted
- predictors is an R formula specifying the features in the mydata data frame to use for prediction
- data is the data frame in which class and predictors can be found

The function will return a RIPPER model object that can be used to make predictions.

## Making predictions:

p <- predict(m, test)

- $m$ is a model trained by the JRip() function
- test is a data frame containing test data with the same features as the training data used to build the classifier.

The function will return a vector of predicted class values.

## Example:

$$
\begin{array}{r}
\text { mushroom_classifier }<- \text { JRip(type } \sim \text { odor }+ \text { cap_color, } \\
\text { data }=\text { mushroom_train) } \\
\text { mushroom_prediction }<- \text { predict(mushroom_classifier, } \\
\text { mushroom_test) }
\end{array}
$$



## Chapter 6:









$$
\begin{aligned}
& \sum\left(y_{i}-\hat{y}_{i}\right)^{2}=\sum e_{i}^{2} \\
& a=\bar{y}-b \bar{x}
\end{aligned}
$$

$b=\frac{\sum\left(x_{i}-\bar{x}\right)\left(y_{i}-\bar{y}\right)}{\sum\left(x_{i}-\bar{x}\right)^{2}}$
$\operatorname{Var}(x)=\underline{\sum\left(x_{i}-\bar{x}\right)^{2}}$
$n$
$\operatorname{Cov}(x, y)=\frac{\sum\left(x_{i}-\bar{x}\right)\left(y_{i}-\bar{y}\right)}{n}$


$$
\begin{aligned}
& \rho_{x, y}=\operatorname{Corr}(x, y)=\frac{\operatorname{Cov}(x, y)}{\sigma_{x} \sigma_{y}} \\
& y=\alpha+\beta_{1} x_{1}+\beta_{2} x_{2}+\ldots+\beta_{i} x_{i}+\varepsilon \\
& y=\beta_{0}+\beta_{1} x_{1}+\beta_{2} x_{2}+\ldots+\beta_{i} x_{i}+\varepsilon \\
& y=\beta_{0} x_{0}+\beta_{1} x_{1}+\beta_{2} x_{2}+\ldots+\beta_{i} x_{i}+\varepsilon
\end{aligned}
$$

regression coefficients
$\begin{array}{llll}\boldsymbol{\beta}_{0} & \boldsymbol{\beta}_{1} & \boldsymbol{\beta}_{2} & \boldsymbol{\beta}_{3}\end{array}$


## $\mathbf{Y}=\boldsymbol{\beta} \mathbf{X}+\varepsilon$ <br> $\hat{\boldsymbol{\beta}}=\left(\mathbf{X}^{\mathbf{T}} \mathbf{X}\right)^{-1} \mathbf{X}^{\mathbf{T}} \mathbf{Y}$

Histogram of insurance\$expenses



| $\infty$ | $\circ \circ \circ \circ$ | $000 \infty \circ$ | $\circ$ |  |  |
| :--- | :--- | :--- | :--- | :--- | :--- | :--- |
| $\circ$ | $\circ \infty \circ \circ \infty \infty \infty$ | $\infty \infty \infty$ | $\infty$ | $\circ$ |  |

$000000000000000000000000000000 \infty 00000$


0000000000000000000000000000000000




expenses


## Multiple regression modeling syntax

using the 7 m () function in the stats package

## Building the model:

$\mathrm{m}<-1 \mathrm{~m}(\mathrm{dv} \sim \mathrm{iv}$, data = mydata)

- $d v$ is the dependent variable in the mydata data frame to be modeled
- $\quad i v$ is an $R$ formula specifying the independent variables in the mydata data frame to use in the model
- data specifies the data frame in which the $d v$ and $i v$ variables can be found

The function will return a regression model object that can be used to make predictions. Interactions between independent variables can be specified using the * operator.

## Making predictions:

$\mathrm{p}<-\mathrm{predict}(\mathrm{m}$, test)

- $m$ is a model trained by the 1 m () function
- test is a data frame containing test data with the same features as the training data used to build the model.

The function will return a vector of predicted values.

## Example:

```
ins_mode1 <- 1m(charges ~ age + sex + smoker,
    data = insurance)
ins_pred <- predict(ins_model, insurance_test)
```

Call:
$\operatorname{lm}$ (formula $=$ expenses $\sim$., data $=$ insurance)
Residuals:

| Min | $1 Q$ | Median | $3 Q$ | Max |
| ---: | ---: | ---: | ---: | ---: |
| -11302.7 | -2850.9 | -979.6 | 1383.9 | 29981.7 |

Coefficients:
Estimate Std. Error $t$ value $\operatorname{Pr}(>|\mathrm{t}|)$
(Intercept) -11941.6 987.8-12.089 < 2e-16
age
sexmale 256.8
11.9 21.586 < 2e-16

2
bmi
children
smokeryes
-131.3
$332.9-0.3950 .693255$
regionnorthwest
339.3
$28.611 .864<2 \mathrm{e}-16$
$475.7 \quad 137.8 \quad 3.452 \quad 0.000574$
23847.5
413.157 .723 < 2e-16 ***
regionsoutheast -1035.6 478.7 -2.163 0.030685 *
regionsouthwest -959.3 477.9-2.007 0.044921 *
Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 '.’ 0.1 ' ' 1

Residual standard error: 6062 on 1329 degrees of freedom Multiple R-squared: 0.7509, Adjusted R-squared: 0.7494 F-statistic: 500.9 on 8 and 1329 DF, p-value: < 2.2e-16

$$
y=\alpha+\beta_{1} x
$$



Call:
$\operatorname{lm}$ (formula $=$ expenses $\sim$ age + age $2+$ children $+b m i+s e x+b m i 30 *$ smoker + region, data = insurance)

Residuals:

| Min | $1 Q$ | Median | 3Q | Max |
| ---: | ---: | ---: | ---: | ---: |
| -17297.1 | -1656.0 | -1262.7 | -727.8 | 24161.6 |

Coefficients:

|  | Estimate | Std. Error | value | $\operatorname{Pr}(>\|t\|)$ |  |
| :---: | :---: | :---: | :---: | :---: | :---: |
| (Intercept) | 139.0053 | 1363.1359 | 0.102 | 0.918792 |  |
| age | -32.6181 | 59.8250 | -0.545 | 0.585690 |  |
| age2 | 3.7307 | 0.7463 | 4.999 | 6.54e-07 |  |
| children | 678.6017 | 105.8855 | 6.409 | 2.03e-10 |  |
| bmi | 119.7715 | 34.2796 | 3.494 | 0.000492 |  |
| sexmale | -496.7690 | 244.3713 | -2.033 | 0.042267 |  |
| bmi30 | -997.9355 | 422.9607 | -2.359 | 0.018449 |  |
| smokeryes | 13404.5952 | 439.9591 | 30.468 | < 2e-16 |  |
| regionnorthwest | -279.1661 | 349.2826 | -0.799 | 0.424285 |  |
| regionsoutheast | -828.0345 | 351.6484 | -2.355 | 0.018682 |  |
| regionsouthwest | -1222.1619 | 350.5314 | -3.487 | 0.000505 |  |
| bmi30:smokeryes | 19810.1534 | 604.6769 | 32.762 | < 2e-16 |  |

Signif. codes: 0 ‘***’ 0.001 ‘**’ 0.01 ‘*’ 0.05 '.’ 0.1 ' ’ 1
Residual standard error: 4445 on 1326 degrees of freedom
Multiple R-squared: 0.8664, Adjusted R-squared: 0.8653
F-statistic: 781.7 on 11 and 1326 DF, p-value: < $2.2 \mathrm{e}-16$


| original data | 1 | 1 | 1 | 2 | 2 | 3 | 4 | 5 | 5 | 6 | 6 | 7 | 7 | 7 | 7 |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| split on feature A | 1 | 1 | 1 | 2 | 2 | 3 | 4 | 5 | 5 | 6 | 6 | 7 | 7 | 7 | 7 |
| split on feature B | 1 | 1 | 1 | 2 | 2 | 3 | 4 | 5 | 5 | 6 | 6 | 7 | 7 | 7 | 7 |

Histogram of wine\$quality


## Regression trees syntax

using the rpart () function in the rpart package

## Building the model:

```
m <- rpart(dv ~ iv, data = mydata)
```

- $d v$ is the dependent variable in the mydata data frame to be modeled
- $\quad i v$ is an $R$ formula specifying the independent variables in the mydata data frame to use in the model
- data specifies the data frame in which the $d v$ and $i v$ variables can be found

The function will return a regression tree model object that can be used to make predictions.

## Making predictions:

```
p <- predict(m, test, type = "vector")
```

- $m$ is a model trained by the rpart () function
- test is a data frame containing test data with the same features as the training data used to build the model
- type specifies the type of prediction to return, either "vector" (for predicted numeric values), "class" for predicted classes, or "prob" (for predicted class probabilities)

The function will return a vector of predictions depending on the type parameter.

## Example:

```
wine_mode1 <- rpart(quality ~ alcohol + sulfates,
    data = wine_train)
wine_predictions <- predict(wine_mode1, wine_test)
```



## $\mathrm{MAE}=\frac{1}{n}$ $\sum_{i=1}^{n}\left|e_{i}\right|$

## Model trees syntax

using the M5 P () function in the RWeka package

## Building the model:

$m<-$ M5P(dv ~ iv, data = mydata)

- $d v$ is the dependent variable in the mydata data frame to be modeled
- $\quad i v$ is an $R$ formula specifying the independent variables in the mydata data frame to use in the model
- data specifies the data frame in which the $d v$ and $i v$ variables can be found

The function will return a model tree object that can be used to make predictions.

## Making predictions:

$\mathrm{p}<-\mathrm{predict}(m$, test)

- $m$ is a model trained by the M5P() function
- test is a data frame containing test data with the same features as the training data used to build the model

The function will return a vector of predicted numeric values.

## Example:

```
wine_mode1 <- M5P(quality ~ alcohol + sulfates,
    data = wine_train)
wine_predictions <- predict(wine_mode1, wine_test)
```


## Chapter 7:



## $y(x)=f$ <br> 

1.00-

$$
f(x)=\left\{\begin{array}{l}
0 \text { if } x<0 \\
1 \text { if } x \geq 0
\end{array}\right.
$$

Output Signal





Hyperbolic Tangent





Input Nodes

## Neural network syntax

using the neuralnet () function in the neuralnet package

## Building the model:

$m$ <- neuralnet (target ~ predictors, data = mydata, hidden = 1 )

- target is the outcome in the mydata data frame to be modeled
- predictors is an R formula specifying the features in the mydata data frame to use for prediction
- data specifies the data frame in which the target and predictors variables can be found
- $\quad h i d d e n$ specifies the number of neurons in the hidden layer (by default, 1)

The function will return a neural network object that can be used to make predictions.

## Making predictions:

$\mathrm{p}<-\operatorname{compute}(m$, test)

- $m$ is a model trained by the neuralnet () function
- test is a data frame containing test data with the same features as the training data used to build the classifier

The function will return a list with two components: \$ neurons, which stores the neurons for each layer in the network, and \$net. result, which stores the model's predicted values.

## Example:

```
concrete_mode1 <- neuralnet(strength ~ cement + slag
    + ash, data = concrete)
mode1_results <- compute(concrete_mode1,
    concrete_data)
strength_predictions <- mode1_results$net.resu1t
```



Error: 5.077438 Steps: 4882


Error: 1.626684 Steps: 86849

Two Dimensions


Three Dimensions



$\vec{w} \cdot \vec{x}+b=0$
$\vec{w} \cdot \vec{x}+b \geq+1$
$\vec{w} \cdot \vec{x}+b \leq-1$


$\min \frac{1}{2}\|\vec{w}\|^{2}+C \sum_{i=1}^{n} \xi_{i}$
s.t. $y_{i}\left(\vec{w} \cdot \vec{x}_{i}-b\right) \geq 1-\xi_{i}, \forall \vec{x}_{i}, \xi_{i} \geq 0$

$$
\begin{aligned}
& \mathrm{K}\left(\overrightarrow{x_{i}}, \overrightarrow{x_{j}}\right)=\phi\left(\overrightarrow{x_{i}}\right) \cdot \phi\left(\overrightarrow{x_{j}}\right) \\
& K\left(\overrightarrow{x_{i}}, \overrightarrow{x_{j}}\right)=\overrightarrow{x_{i}} \cdot \overrightarrow{x_{j}} \\
& \mathrm{~K}\left(\overrightarrow{x_{i}}, \overrightarrow{x_{j}}\right)=\left(\overrightarrow{x_{i}} \cdot \overrightarrow{x_{j}}+1\right)^{d} \\
& \mathrm{~K}\left(\overrightarrow{x_{i}}, \overrightarrow{x_{j}}\right)=\tanh \left(\kappa \overrightarrow{x_{i}} \cdot \overrightarrow{x_{j}}-\delta\right) \\
& K\left(\vec{x}_{i}, \vec{x}_{j}\right)=e^{2 \sigma^{2}}
\end{aligned}
$$



## Support vector machine syntax

using the ksvm () function in the kern 1 ab package

## Building the model:

m <- ksvm(target ~ predictors, data = mydata,

$$
\text { kerne1 = "rbfdot", } c=1)
$$

- target is the outcome in the mydata data frame to be modeled
- predictors is an R formula specifying the features in the mydata data frame to use for prediction
- data specifies the data frame in which the target and predictors variables can be found
- kerne1 specifies a nonlinear mapping such as "rbfdot" (radial basis), "polydot" (polynomial), "tanhdot" (hyperbolic tangent sigmoid), or "vanilladot" (linear)
- C is a number that specifies the cost of violating the constraints, i.e., how big of a penalty there is for the "soft margin." Larger values will result in narrower margins
The function will return a SVM object that can be used to make predictions.


## Making predictions:

$\mathrm{p}<-\mathrm{predict}(m$, test, type = "response")

- $m$ is a model trained by the $k s v m()$ function
- test is a data frame containing test data with the same features as the training data used to build the classifier
- type specifies whether the predictions should be "response" (the predicted class) or "probabilities" (the predicted probability, one column per class level).

The function will return a vector (or matrix) of predicted classes (or probabilities) depending on the value of the type parameter.

## Example:

```
1etter_classifier <- ksvm(1etter ~ ., data =
    letters_train, kerne1 = "vanilladot")
letter_prediction <- predict(letter_classifier,
    1etters_test)
```


## Multiple Output Nodes



## Multiple Hidden Layers




## true minimum

## Chapter 8:

$\{$ bread, peanut butter, jelly $\}$
$\{$ peanut butter, jelly $\} \rightarrow\{$ bread $\}$

## $\operatorname{support}(X)=\frac{\operatorname{count}(X)}{N}$

confidence $(X \rightarrow Y)=\frac{\operatorname{support}(X, Y)}{\operatorname{support}(X)}$

|  | V1 | V2 | V3 | V4 |
| :---: | :--- | :--- | :--- | :--- |
| $\mathbf{1}$ | citrus fruit | semi-finished bread | margarine | ready soups |
| $\mathbf{2}$ | tropical fruit | yogurt | coffee |  |
| $\mathbf{3}$ | whole milk |  |  |  |
| $\mathbf{4}$ | pip fruit | yogurt | cream cheese | meat spreads |
| $\mathbf{5}$ | other vegetables | whole milk | condensed milk | long life bakery product |







## Association rule syntax

using the apriori() function in the arules package

## Finding association rules:

```
    myrules <- apriori(data = mydata, parameter =
```

        list(support \(=0.1\), confidence \(=0.8\), minlen \(=1\) ))
    - data is a sparse item matrix holding transactional data
- support specifies the minimum required rule support
- confidence specifies the minimum required rule confidence
- min1en specifies the minimum required rule items

The function will return a rules object storing all rules that meet the minimum criteria.

## Examining association rules:

inspect(myrules)

- myrules is a set of association rules from the apriori() function

This will output the association rules to the screen. Vector operators can be used on myrules to choose a specific rule or rules to view.

## Example:

```
groceryrules <- apriori(groceries, parameter =
    1ist(support = 0.01, confidence = 0.25, min1en = 2))
inspect(groceryrules[1:3])
```


## $\operatorname{lift}(X \rightarrow Y)=$ confidence $(X \rightarrow Y)$ support $(Y)$




Chapter 9:



Math and Statistics Publications


$$
\operatorname{dist}(x, y)=\sqrt{\sum_{i=1}\left(x_{i}-y_{i}\right)^{2}}
$$

## $n$







Within-Group Heterogeneity


## Clustering syntax

## using the kmeans () function in the stats package

## Finding clusters:

myclusters <- kmeans(mydata, k)

- mydata is a matrix or data frame with the examples to be clustered
- $\mathbf{k}$ specifies the desired number of clusters

The function will return a cluster object that stores information about the clusters.

## Examining clusters:

- myc1usters\$c1uster is a vector of cluster assignments from the kmeans () function
- myclusters\$centers is a matrix indicating the mean values for each feature and cluster combination
- myclusters\$size lists the number of examples assigned to each cluster


## Example:

```
teen_clusters <- kmeans(teens, 5)
teens$cluster_id <- teen_clusters$cluster
```

|  | basketball | football | soccer | softball | volleyball | swimming |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| 1 | 0.16001227 | 0.2364174 | 0.10385512 | 0.07232021 | 0.18897158 | 0.23970234 |
| 2 | -0.09195886 | 0.0652625 - | -0.09932124 - | -0.01739428 - | -0.06219308 | 0.03339844 |
| 3 | 0.52755083 | 0.4873480 | 0.29778605 | 0.37178877 | 0.37986175 | 0.29628671 |
| 4 | 0.34081039 | 0.3593965 | 0.12722250 | 0.16384661 | 0.11032200 | 0.26943332 |
| 5 | -0.16695523 | -0.1641499 - | -0.09033520 - | -0.11367669 - | -0.11682181 -0 | 0.10595448 |
|  | cheerleading | baseball | 1 tennis | s sports | s cute | x |
| 1 | 0.3931445 | 0.02993479 | 90.13532387 | $7 \quad 0.10257837$ | 70.37884271 | 0.020042068 |
| 2 | -0.1101103 | -0.11487510 | 00.04062204 | -0.09899231 | $1-0.03265037$ | -0.042486141 |
| 3 | 0.3303485 | 0.35231971 | 10.14057808 | 80.32967130 | 0.54442929 | 0.002913623 |
| 4 | 0.1856664 | 0.27527088 | 80.10980958 | 80.79711920 | 0.47866008 | 2.028471066 |
|  | -0.1136077 | -0.10918483 | $3-0.05097057$ | $7-0.13135334$ | -0.18878627 | -0.097928345 |


|  | sexy | hot | kissed | dance | band | marching | music |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 1 | 0.11740551 | 0.41389104 | 0.06787768 | 0.22780899 | -0.10257102 | -0.10942590 | 0.1378306 |
| 2 | -0.04329091 | -0.03812345 | -0.04554933 | 0.04573186 | 4.06726666 | 5.25757242 | 0.4981238 |
| 3 | 0.24040196 | 0.38551819 | -0.03356121 | 0.45662534 | -0.02120728 | -0.10880541 | 0.2844999 |
| 4 | 0.51266080 | 0.31708549 | 2.97973077 | 0.45535061 | 0.38053621 | -0.02014608 | 1.1367885 |
| 5 | -0.09501817 | -0.13810894 | -0.13535855 | -0.15932739 | -0.12167214 | -0.11098063 | -0.1532006 |


| Cluster 1 $(N=3,376)$ | Cluster 2 $(N=601)$ | Cluster 3 $(N=1,036)$ | Cluster 4 $(N=3,279)$ | Cluster 5 $(N=21,708)$ |
| :---: | :---: | :---: | :---: | :---: |
| swimming cheerleading cute sexy hot dance dress hair mall hollister abercrombie shopping clothes | band marching music rock | sports <br> sex <br> sexy <br> hot <br> kissed <br> dance <br> music <br> band <br> die <br> death <br> drunk <br> drugs | basketball football soccer softball volleyball baseball sports god church Jesus bible | ??? |
| Princesses | Brains | Criminals | Athletes | Basket Cases |

## Chapter 10:



Three Classes
Predicted Class


| Predicted to be Spam <br> yes |  |
| :---: | :---: |
| yctually <br> Spam <br> yes |  |

$$
\begin{aligned}
& \text { accuracy }=\frac{\mathrm{TP}+\mathrm{TN}}{\mathrm{TP}+\mathrm{TN}+\mathrm{FP}+\mathrm{FN}} \\
& \text { error rate }=\frac{\mathrm{FP}+\mathrm{FN}}{\mathrm{TP}+\mathrm{TN}+\mathrm{FP}+\mathrm{FN}}=1-\text { accuracy }
\end{aligned}
$$

Cell Contents


Total Observations in Table: 1390


Confusion Matrix and Statistics

| Reference |  |  |
| :---: | ---: | ---: |
| Prediction | ham | spam |
| ham | 1203 | 31 |
| spam | 4 | 152 |

Accuracy : 0.9748
95\% CI : (0.9652, 0.9824)
No Information Rate : 0.8683 P-Value [Acc > NIR] : < 2.2e-16
$\begin{aligned} & \text { Kappa }: 0.8825 \\ & \text { Mcnemar's Test P-Value }: 1.109 e-05\end{aligned}$

Sensitivity : 0.8306
Specificity : 0.9967
Pos Pred Value : 0.9744
Neg Pred Value : 0.9749
Prevalence : 0.1317
Detection Rate : 0.1094
Detection Prevalence : 0.1122
Balanced Accuracy : 0.9136
'Positive' Class : spam
$\kappa=\frac{\operatorname{Pr}(a)-\operatorname{Pr}(e)}{1-\operatorname{Pr}(e)}$


## TP

## sensitivity $=$ <br>  <br>  <br> specificity $=$ <br> 

$$
\begin{aligned}
& \text { precision }=\frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FP}} \\
& \text { recall }=\frac{\mathrm{TP}}{\mathrm{TP}+\mathrm{FN}} \\
&
\end{aligned}
$$



error $=0.632 \times$ error $_{\text {test }}+0.368 \times$ error $_{\text {train }}$

## Chapter 11:



1000 samples
16 predictor
2 classes: 'no', 'yes'
2 No pre-processing
2 Resampling: Bootstrapped (25 reps)
Summary of sample sizes: 1000, 1000, 1000, 1000, 1000, 1000, ...

3 Resampling results across tuning parameters:

| model | winnow | trials | Accuracy | Kappa | Accuracy SD | Kappa SD |
| :--- | :---: | ---: | :--- | :--- | :--- | :--- |
| rules | FALSE | 1 | 0.6847204 | 0.2578421 | 0.02558775 | 0.05622302 |
| rules | FALSE | 10 | 0.7112829 | 0.3094601 | 0.02087257 | 0.04585890 |
| rules | FALSE | 20 | 0.7221976 | 0.3260145 | 0.01977334 | 0.04512083 |
| rules | TRUE | 1 | 0.6888432 | 0.2549192 | 0.02683844 | 0.05695277 |
| rules | TRUE | 10 | 0.7113716 | 0.3038075 | 0.01947701 | 0.04484956 |
| rules | TRUE | 20 | 0.7233222 | 0.3266866 | 0.01843672 | 0.03714053 |
| tree | FALSE | 1 | 0.6769653 | 0.2285102 | 0.03027647 | 0.07001131 |
| tree | FALSE | 10 | 0.7222552 | 0.2880662 | 0.02061900 | 0.05601918 |
| tree | FALSE | 20 | 0.7297858 | 0.3067404 | 0.02007556 | 0.05616826 |
| tree | TRUE | 1 | 0.6771020 | 0.2219533 | 0.02703456 | 0.05955907 |
| tree | TRUE | 10 | 0.7173312 | 0.2777136 | 0.01700633 | 0.04358591 |
| tree | TRUE | 20 | 0.7285714 | 0.3058474 | 0.01497973 | 0.04145128 |

4
Accuracy was used to select the optimal model using the largest value.
The final values used for the model were trials $=20$, model $=$ tree and winnow $=$ FALSE.

1000 samples
16 predictor
2 classes: 'no', 'yes'
No pre-processing
Resampling: Cross-Validated (10 fold)
Summary of sample sizes: 900, 900, 900, 900, 900, 900, ...
Resampling results across tuning parameters:

| trials | Accuracy | Kappa | Accuracy SD | Kappa SD |
| ---: | :--- | :--- | :--- | :--- |
| 1 | 0.724 | 0.3124461 | 0.02547330 | 0.05897140 |
| 5 | 0.713 | 0.2921760 | 0.02110819 | 0.06018851 |
| 10 | 0.719 | 0.2947271 | 0.03107339 | 0.06719720 |
| 15 | 0.721 | 0.3009258 | 0.01969207 | 0.05105480 |
| 20 | 0.717 | 0.2929875 | 0.02790858 | 0.07912362 |
| 25 | 0.728 | 0.3150336 | 0.03224903 | 0.09367152 |
| 30 | 0.729 | 0.3104144 | 0.02766867 | 0.08069045 |
| 35 | 0.741 | 0.3389908 | 0.03142893 | 0.09352673 |

Tuning parameter 'model' was held constant at a value of tree Tuning parameter 'winnow' was held constant at a value of FALSE Kappa was used to select the optimal model using the one SE rule. The final values used for the model were trials $=1$, model $=$ tree and winnow = FALSE.



## Random forest syntax

## using the randomForest() function in the randomForest package

## Building the classifier:

m <- randomForest(train, class, ntree $=500$, mtry $=\operatorname{sqrt}(p)$ )

- train is a data frame containing training data
- c1 ass is a factor vector with the class for each row in the training data
- ntree is an integer specifying the number of trees to grow
- mtry is an optional integer specifying the number of features to randomly select at each split (uses sqrt(p) by default, where p is the number of features in the data)

The function will return a random forest object that can be used to make predictions.

## Making predictions:

```
p <- predict(m, test, type = "response")
```

- $m$ is a model trained by the randomForest() function
- test is a data frame containing test data with the same features as the training data used to build the classifier
- type is either "response", "prob", or "votes" and is used to indicate whether the predictions vector should contain the predicted class, the predicted probabilities, or a matrix of vote counts, respectively.

The function will return predictions according to the value of the type parameter.

## Example:

credit_mode1 <- randomForest(credit_train, loan_default)
credit_prediction <- predict(credit_mode1, credit_test)

Chapter 12:



Source: local data frame [1,000 x 17]

| checking_balance | months_loan_duration | credit_history | purpose | amount |
| :---: | :---: | :---: | :---: | :---: |
| < 0 DM | 6 | critical | furniture/appliances | 1169 |
| 1-200 DM | 48 | good | furniture/appliances | 5951 |
| unknown | 12 | critical | education | 2096 |
| < 0 DM | 42 | good | furniture/appliances | 7882 |
| $<0 \mathrm{DM}$ | 24 | poor | car | 4870 |
| unknown | 36 | good | education | 9055 |
| unknown | 24 | good | furniture/appliances | 2835 |
| 1-200 DM | 36 | good | car | 6948 |
| unknown | 12 | good | furniture/appliances | 3059 |
| 1-200 DM | 30 | critical | car | 5234 |
| ... | $\ldots$ | . | ... |  |

Variables not shown: savings_balance (fctr), employment_duration (fctr), percent_of_income (int), years_at_residence (int), age (int), other_credit (fctr), housing (fctr), existing_loans_count (int), job (fctr), dependents (int), phone (fctr), default (fctr)

Serial computing:


Parallel computing:



CPU with 16 cores


GPU with 1000+ cores

