Statistical learning and Visualization:

Supervised learning - classification (1/2)

Erik-Jan van Kesteren

Department of Methodology and Statistics

Universiteit Utrecht

Applied Data Science

1. Introduction
2. KNN
3. Discriminative
4. Generative
5. Break
6. Trees!
7. Evaluating classifiers

About me

- Assistant professor of data science @ UU M&S
- Team lead for the Social Data Science Team (ODISSEI national consortium)
- Statistics, programming, high-dimensional data, geospatial analysis, supercomputing, structural equation modeling, (Bayesian) probabilistic programming, optimization, synthetic data, and more
- I will teach two classification weeks in this course
- I coordinate the `INFOMDA2` course (sign up everyone)!

Topics this week

- Classification
- KNN
- Logistic regression
- Linear discriminant analysis
- Generative vs discriminative
- Trees
- Confusion matrix

Classification

The thing you're trying to predict is *discrete*:

- *Titanic*: Survival/Nonsurvival
- Banking data: Default on/payment of debt
- GPS/Accelerometer data: Work/Home/Friend/Parking/Other
- Imagenet: gazelle/tank/pirate/sea lion/tandem bicycle/...
- Etc.

KNN

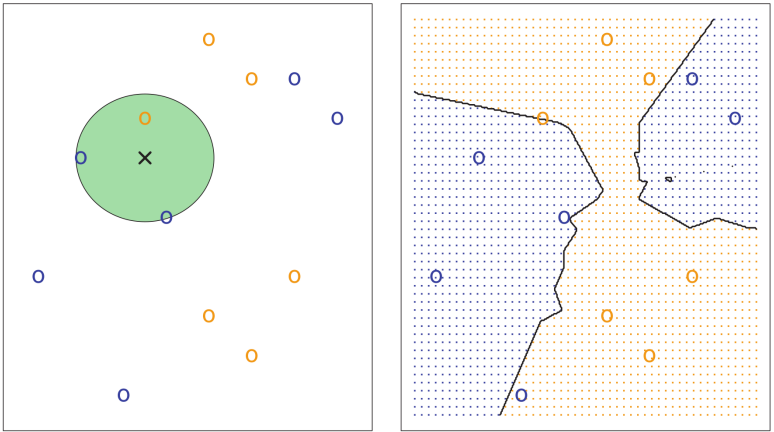

FIGURE 2.14. *The KNN approach, using $K = 3$, is illustrated in a s*

Discriminative classifier

Directly model $p(Y = k|X)$ as a function of X.

$$p(Y = k|X) = f(X)$$

Logistic regression

$$p(Y = 1|X) = \mathit{logit}^{-1}(\beta_0 + \beta_1 X) = \frac{e^{\beta_0+\beta_1 X}}{1 + e^{\beta_0+\beta_1 X}}$$

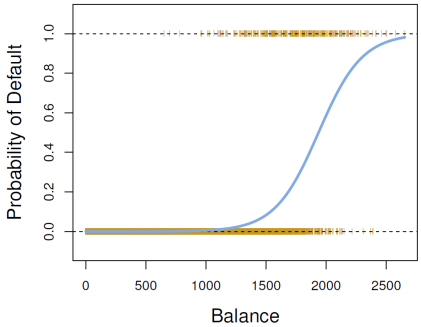

$\beta_0 = -10.65$, $\beta_1 = 0.0055$

Logistic regression

Turning this function around:

$$log\left(\frac{p(Y=1|X)}{1-p(Y=1|X)}\right) = \beta_0 + \beta_1 X$$

Get comfortable with odds, log-odds, the logit, and the inverse logit!

Logistic regression

$$log\left(\frac{p(Y=1|X)}{1-p(Y=1|X)}\right) = \beta_0 + \beta_1 X$$

If $\beta_0 = 0; \beta_1 = 2$: Interpretation for log-odds?
When X increases by 1, the log-odds of $Y = 1$ increase by 2.

Logistic regression

$$\frac{p(Y=1|X)}{1-p(Y=1|X)} = e^{\beta_0+\beta_1 X}$$

If $\beta_0 = 0; \beta_1 = 2$: Interpretation in odds?
When X increases by 1, the odds of $Y = 1$ multiply by $e^2 = 7.39$

Logistic regression

$$p(Y = 1|X) = \frac{e^{\beta_0+\beta_1 X}}{1 + e^{\beta_0+\beta_1 X}}$$

If $\beta_0 = 0; \beta_1 = 2$: Interpretation in probabilities?

- When X increases from 0 to 1, $Pr(Y = 1)$ increases from $logit^{-1}(0 + 2 \cdot 0) = 0.5$ to $logit^{-1}(0 + 2 \cdot 1) \approx 0.88$
- When X increases from 1 to 2, $Pr(Y = 1)$ increases from $logit^{-1}(0 + 2 \cdot 1) \approx 0.88$ to $logit^{-1}(0 + 2 \cdot 2) \approx 0.98$

Tip: use predicted probabilities (`predict(model, type = "response")` function in R)

Generative classifier

Use Bayes' rule to get to $p(Y = k|X)$.

$$p(Y = k|X) = \frac{\pi_k \cdot p(X|Y = k)}{\sum_{k=1}^{K} \pi_k \cdot p(X|Y = k)}$$

Linear discriminant analysis

- π_k is the proportion of observations in class k
- $p(X|Y = x)$ is a normal distribution with mean μ_k and common variance σ^2

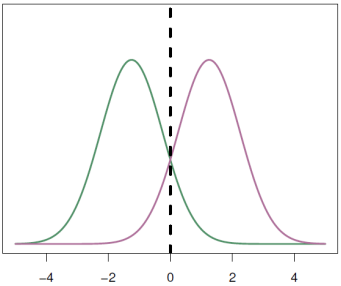

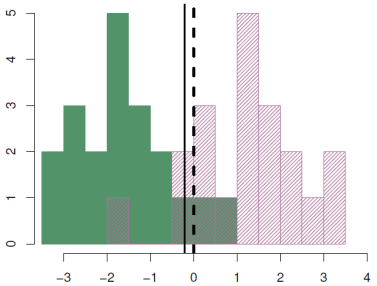

Linear discriminant analysis

Advantages over logistic regression:

- Easy to extend to K > 2 classes
- Really easy to estimate (analytic solution for μ_k and σ^2). You can program it yourself!
- You can generate new *X* from the model (generative model).

Disadvantages:

- Assumption that *X* is normally distributed within each class *k* (categorical predictors???)
- Assumption that the variance of each normal distribution is the same!

Linear discriminant analysis

Discriminative classifiers

- Directly model $p(Y = k|X)$, for example using the logit link function.

Generative classifiers

- Estimate $p(X|Y = k)$ and π_k
- Use Bayes' rule to turn this into $p(Y = k|X)$:

$$p(Y = k|X) = \frac{\pi_k \cdot p(X|Y = k)}{\sum_{k=1}^{K} \pi_k \cdot p(X|Y = k)}$$

Break

Trees!

Using decision trees for prediction

Decision tree: should I buy a car?

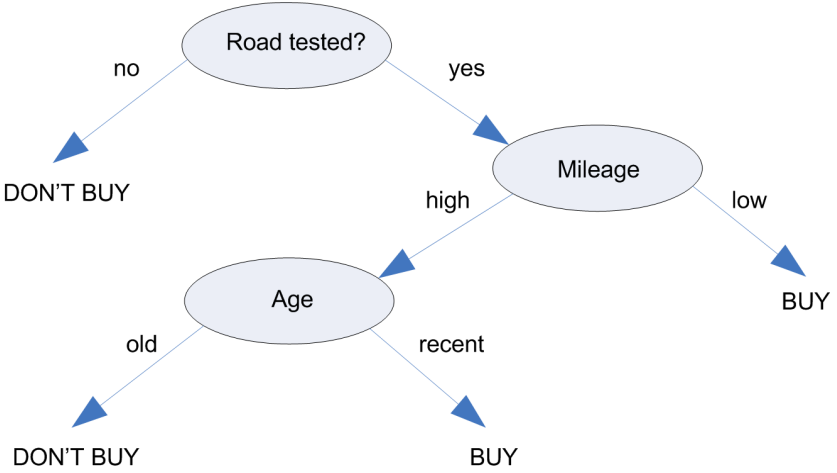

Prediction tree: wood you survive the *Titanic*?

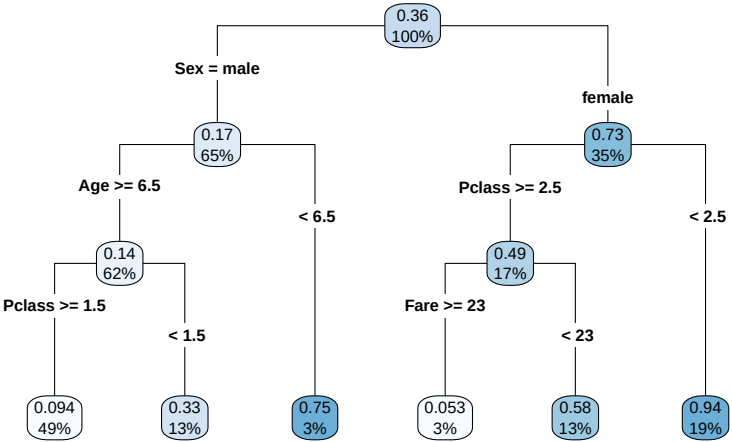

Growing decision trees from data

Recursive partitioning

1. Find the split that makes observations as similar as possible on the outcome within that split;
2. Within each resulting group, do (1).

Recursive partitioning

1. Find the split that makes observations as similar as possible on the outcome within that split;
2. Within each resulting group, do (1).

- Criteria for “as similar as possible”: Purity, Reduction in MSE, ...
- Early stopping: add after (2):
  - “unless there are fewer than n_{min} observations in the group” (typically 10);
  - “unless the total complexity of the model becomes more than *cp*” (typically 0.05);

Simple example

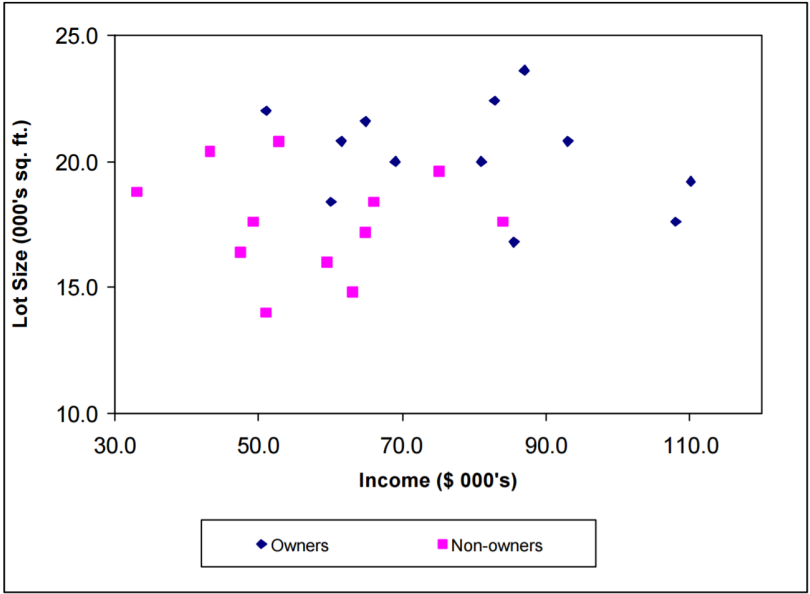
25.0
20.0
15.0
10.0
Lot Size (000's sq. ft.)
30.0
50.0
70.0
90.0
110.0
Income ($ 000's)
Owners
Non-owners

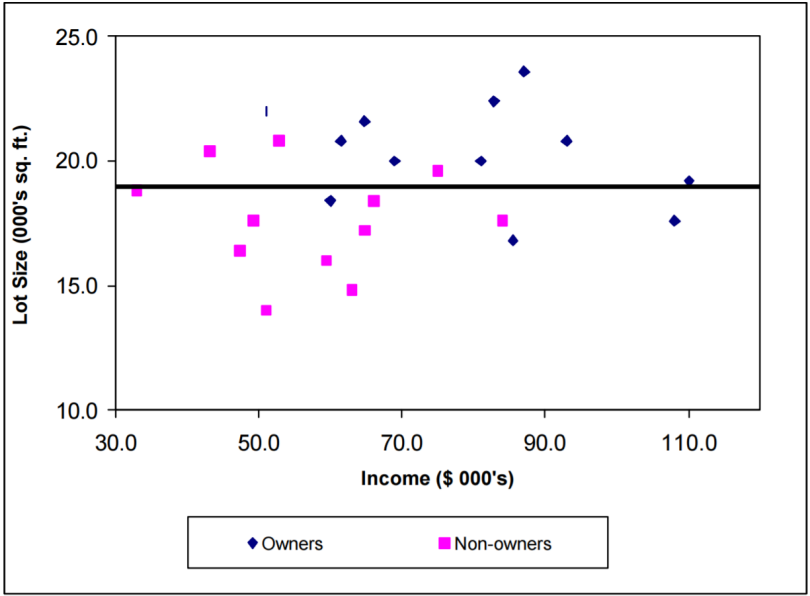
25.0
20.0
15.0
10.0
Lot Size (000's sq. ft.)
30.0
50.0
70.0
90.0
110.0
Income ($ 000's)
Owners
Non-owners

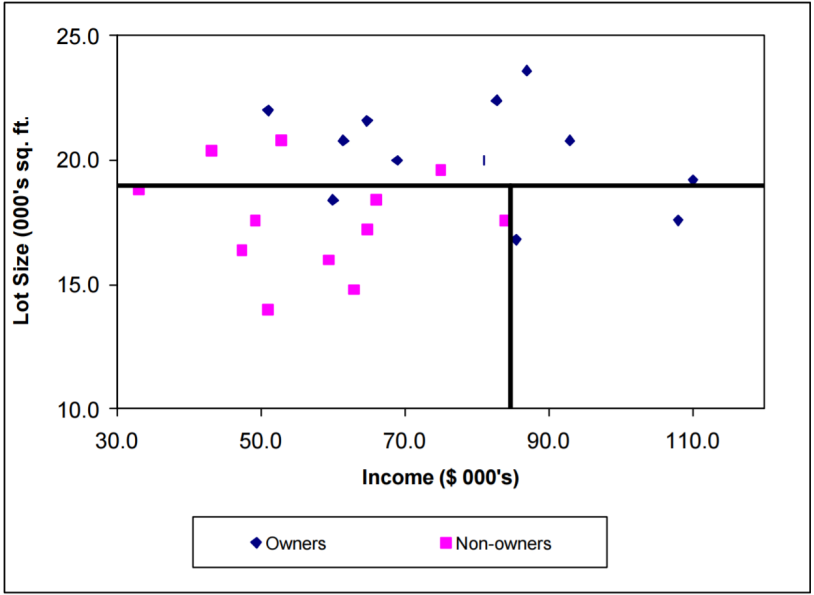
25.0
20.0
15.0
10.0
Lot Size (000's sq. ft.
30.0
50.0
70.0
90.0
110.0
Income ($ 000's)
Owners
Non-owners

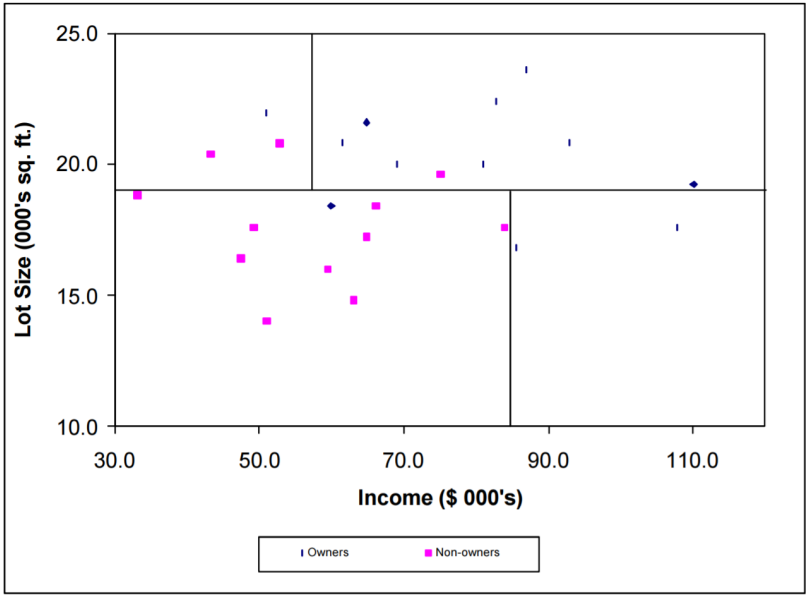
25.0
20.0
15.0
10.0
Lot Size (000's sq. ft.)
30.0
50.0
70.0
90.0
110.0
Income ($ 000's)
Owners
Non-owners

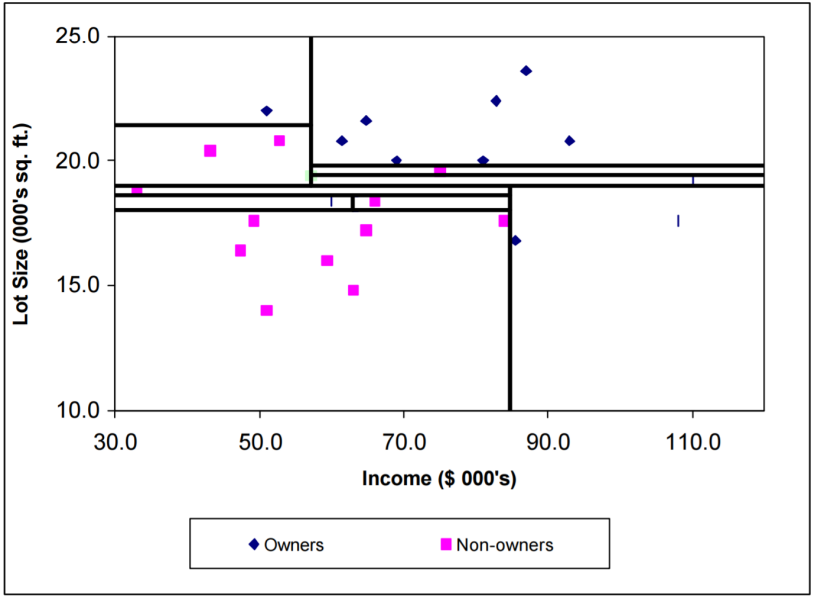
25.0
20.0
15.0
10.0
Lot Size (000's sq. ft.)
30.0
50.0
70.0
90.0
110.0
Income ($ 000's)
Owners
Non-owners

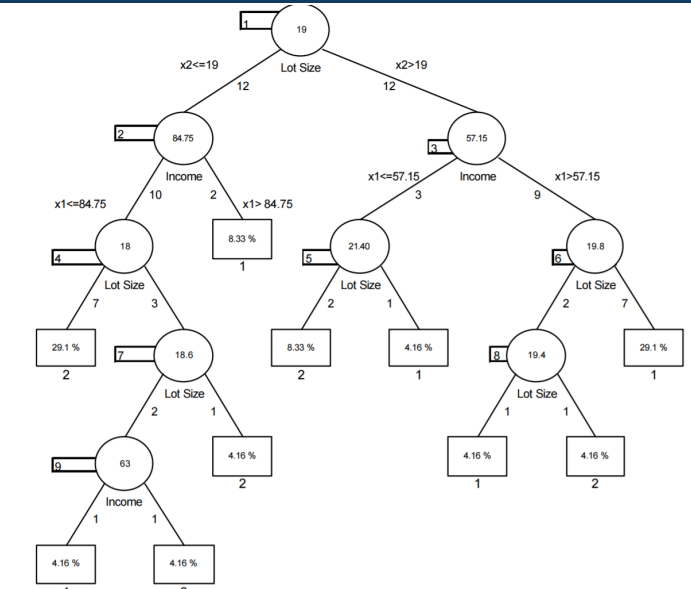
1
19
Lot Size
x2<=19
12
x2>19
12
2
84.75
Income
10
2
x1<=84.75
x1> 84.75
3
57.15
Income
x1<=57.15
3
x1>57.15
9
4
18
Lot Size
7
3
8.33 %
1
5
21.40
Lot Size
2
1
6
19.8
Lot Size
2
7
29.1 %
2
7
18.6
Lot Size
2
1
8.33 %
2
4.16 %
1
8
19.4
Lot Size
1
1
29.1 %
1
4.16 %
2
9
63
Income
1
1
4.16 %
1
4.16 %
2
4.16 %
1
4.16 %
2

More interesting example

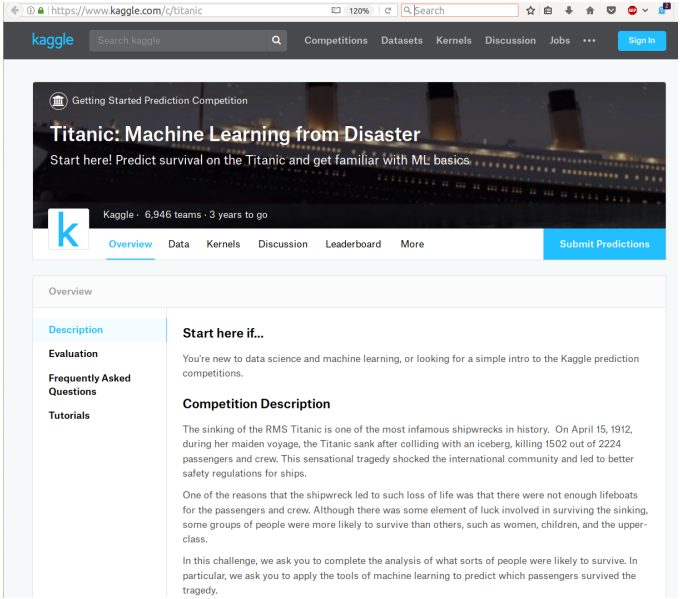
https://www.kaggle.com/c/titanic
120%
Search
kaggle
Search kaggle
Competitions
Datasets
Kernels
Discussion
Jobs
Sign In
Getting Started Prediction Competition
Titanic: Machine Learning from Disaster
Start here! Predict survival on the Titanic and get familiar with ML basics
Kaggle · 6,946 teams · 3 years to go
Overview
Data
Kernels
Discussion
Leaderboard
More
Submit Predictions
Overview
Description
Evaluation
Frequently Asked Questions
Tutorials
Start here if...
You're new to data science and machine learning, or looking for a simple intro to the Kaggle prediction competitions.
Competition Description
The sinking of the RMS Titanic is one of the most infamous shipwrecks in history. On April 15, 1912, during her maiden voyage, the Titanic sank after colliding with an iceberg, killing 1502 out of 2224 passengers and crew. This sensational tragedy shocked the international community and led to better safety regulations for ships.
One of the reasons that the shipwreck led to such loss of life was that there were not enough lifeboats for the passengers and crew. Although there was some element of luck involved in surviving the sinking, some groups of people were more likely to survive than others, such as women, children, and the upper-class.
In this challenge, we ask you to complete the analysis of what sorts of people were likely to survive. In particular, we ask you to apply the tools of machine learning to predict which passengers survived the tragedy.

Data Dictionary

Variable	Definition	Key
survival	Survival	0 = No, 1 = Yes
pclass	Ticket class	1 = 1st, 2 = 2nd, 3 = 3rd
sex	Sex	
Age	Age in years	
sibsp	# of siblings / spouses aboard the Titanic	
parch	# of parents / children aboard the Titanic	
ticket	Ticket number	
fare	Passenger fare	
cabin	Cabin number	
embarked	Port of Embarkation	C = Cherbourg, Q = Queenstown, S = Southampton

Getting the Titanic data from Kaggle

```
library(tidyverse)

# Import the Titanic data from Kaggle
train_url <-
"http://s3.amazonaws.com/assets.datacamp.com/course/Kaggle/train.csv"
titanic_df <- read_csv(train_url)

# Make sure the results are reproducible
set.seed(1027)

# Split the data into 70% train and 30% validation data
N <- nrow(titanic_df)
idx_train <- sample(1:N, size = round(N * 0.7))
train_df <- titanic_df[idx_train, ] # Training data
val_df <- titanic_df[-idx_train, ] # Validation data
```

```
> train_df
# A tibble: 624 x 12
   PassengerId Survived Pclass Name                    Sex     Age SibSp Parch Ticket   Fare Cabin Embarked
         <dbl>    <dbl>  <dbl> <chr>                   <chr> <dbl> <dbl> <dbl> <chr>   <dbl> <chr> <chr>
 1         531        1      2 Quick, Miss. Phyllis M~ fema~     2     1     1 26360   26     NA    S
 2         843        1      1 Serepeca, Miss. Augusta fema~    30     0     0 113798  31     NA    C
 3         414        0      2 Cunningham, Mr. Alfred~ male     NA     0     0 239853   0     NA    S
 4         690        1      1 Madill, Miss. Georgett~ fema~    15     0     1 24160  211.   B5    S
 5         538        1      1 LeRoy, Miss. Bertha     fema~    30     0     0 PC 17~ 106.   NA    C
 6          27        0      3 Emir, Mr. Farred Chehab male     NA     0     0 2631     7.22 NA    C
 7         290        1      3 Connolly, Miss. Kate    fema~    22     0     0 370373   7.75 NA    Q
 8         853        0      3 Boulos, Miss. Nourelain fema~     9     1     1 2678    15.2  NA    C
 9          92        0      3 Andreasson, Mr. Paul E~ male     20     0     0 347466   7.85 NA    S
10         513        1      1 McGough, Mr. James Rob~ male     36     0     0 PC 17~  26.3  E25   S
# ... with 614 more rows
```

Fitting a classification tree in R

```
library(rpart)
library(rpart.plot)

titanic_tree <-
  rpart(
    Survived ~ Pclass + Sex + Age + SibSp + Parch + Fare + Embarked,
    data = train_df,
    control = list(cp = 0.02)
  )

rpart.plot(titanic_tree)
```

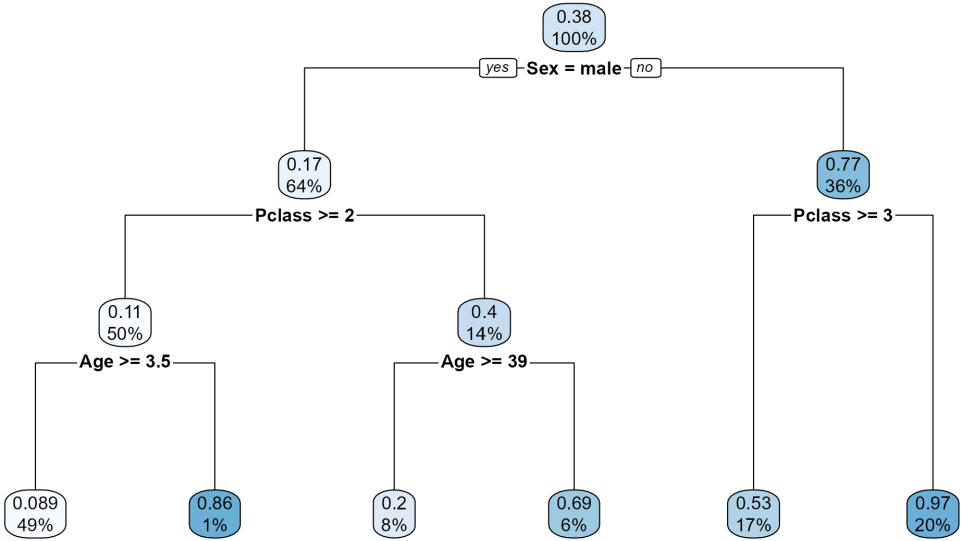
0.38
100%
yes
Sex = male
no
0.17
64%
0.77
36%
Pclass >= 2
Pclass >= 3
0.11
50%
0.4
14%
Age >= 3.5
Age >= 39
0.089
49%
0.86
1%
0.2
8%
0.69
6%
0.53
17%
0.97
20%

Evaluating classifiers

THE INTERNATIONAL JOURNAL OF ROBOTICS RESEARCH / January 2007

Table 5. Place Confusion Matrix

Truth	Inferred labels					FN
	Work	Home	Friend	Parking	Other	
Work	5	0	0	0	0	0
Home	0	4	0	0	0	0
Friend	0	0	3	0	2	0
Parking	0	0	0	8	0	2
Other	0	0	0	0	28	1
FP	0	0	1	1	2	-

More on this next week.
Wednesday: Q&A session for practical.

Have a nice day!