

TMA4268 Statistical Learning V2018

Module 9: SUPPORT VECTOR MACHINES

Mette Langaas and Thea Roksvåg, Department of Mathematical
Sciences, NTNU

week 12, version 18.03.2018

Before we start

Learning material

- ▶ James et al (2013): An Introduction to Statistical Learning. Chapter 9.
- ▶ Classnotes 19.03.2018


Some of the figures in this presentation are taken from (or are inspired by) “An Introduction to Statistical Learning, with applications in R” (Springer, 2013) with permission from the authors: G. James, D. Witten, T. Hastie and R. Tibshirani.

Topics

- ▶ Motivation
- ▶ Maximal margin classifier
- ▶ Support vector classifier
- ▶ Support vector machines
- ▶ Extensions
- ▶ Comparisons
- ▶ Summing up
- ▶ Recommended exercises
- ▶ Packages to install
- ▶ References

Motivation

Suppose that you are interested in the distribution of two tree types: redwood and pines. You have three different study areas in which these trees grow. Your study areas are visualized in the three figures below with orange points indicating the position of a redwood tree and green points indicating the position of a pine tree in a forest.




Assume you want to build one continuous fence to separate the two tree types in each of the three study areas. **Where should you build the fence?**

For Forest 1 the choice seems easy: The orange and green points are clearly separated and the fence can be built anywhere inside the band that separates them. However, we can draw infinitely many straight lines that all separate the two tree types, and we should take into account that the trees reproduce and that we want future pines and future redwoods to grow up on the correct side of the fence.


The problem gets more complicated for Forest 2. A linear fence still seems like a good idea, but in this case the two tree types can not be perfectly separated. You have to allow some of the trees to be on the wrong side of the fence. The complexity of the problem is further increased for Forest 3. It is not possible to separate the two tree types by a straight line without getting a large number of misclassifications. Here, a circular fence around the pine trees seems like a reasonable choice.

Forest 1 illustrates the problem of finding an **optimal separating hyperplane** for a dataset. In this module you are going to learn a method for finding optimal hyperplanes called **Maximal Margin hyperplanes**.




You are also going to learn how you can find an optimal separating hyperplane when your data cannot be perfectly separated by a straight line, as in Forest 2. This leads to a classifier called a **Support Vector Classifier** or a **Soft Margin Classifier**.

Forest 2



The Support vector classifier can be generalised to an approach that produces non-linear decision boundaries. This is called the **Support Vector Machine** (SVM) and is useful when the data is distributed as illustrated in Forest 3.

Forest 3



The reason why we want to separate the data by finding an optimal hyperplane or a separating curve, is that we want to use it for classification. In our example: Is it likely that a random seed found at location $(0.8, 0.4)$ becomes a redwood or a pine tree given the observed data? We classify a new observation based on which side of the decision boundary it falls into.

The methods presented in this module are intended for binary classification, i.e. **classification with two classes**, but extensions to more than two classes are briefly mentioned.


Dataset

The three forests will be used as illustrative examples throughout the module. In this dataset we have two covariates,

*the x_1 - and the x_2 -coordinate of each tree in the forest, and the response is the tree type, either *pine* or *redwood*.*

The class *pine* is stored as 1 in our dataset, while *redwood* is stored as -1. (Yes, we have used 0 and 1 before.)

Our goal is to make a classifier for random seeds that we find on the ground for each of the three forests. The locations of the seeds are shown in the figure below (black circles), and we want to classify the seeds as either pine seeds or redwood seeds given our observations. Thus, the point patterns above can be thought of as the training set, while the point pattern generated by the black circles can be thought of as the test set.



Maximal Margin Classifier

Hyperplane

A **hyperplane** in p -dimensions is defined as

$$\beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_p X_p = \beta_0 + \mathbf{x}^T \boldsymbol{\beta} = 0.$$

and is a $p - 1$ dimensional subspace.

- ▶ If a point $X = (X_1, X_2, \dots, X_p)^T$ satisfies the above equation, it lies on the hyperplane.
- ▶ If $\beta_0 = 0$ the hyperplane goes through the origin (origo).
- ▶ The vector β_1, \dots, β_p (not including β_0) is called the normal vector and points in the direction orthogonal to the hyperplane.

If a point $\mathbf{X} = (X_1, X_2, \dots, X_p)^T$ satisfies


- ▶ $\beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_p X_p = \beta_0 + \mathbf{x}^T \boldsymbol{\beta} > 0$ it lies on one side of the hyperplane, while if it satisfies
- ▶ $\beta_0 + \beta_1 X_1 + \beta_2 X_2 + \dots + \beta_p X_p = \beta_0 + \mathbf{x}^T \boldsymbol{\beta} < 0$ it lies on the opposite side of the hyperplane.
- ▶ For normalized β s ($\sum_{j=1}^p \beta_j^2 = 1$) the value of $\beta_0 + \mathbf{x}^T \boldsymbol{\beta}$ gives the distance from the hyperplane.

Assumptions

Assume that we have n training observations with p predictors

$$\mathbf{x}_1 = \begin{pmatrix} x_{11} \\ \vdots \\ x_{1p} \end{pmatrix}, \dots, \mathbf{x}_n = \begin{pmatrix} x_{n1} \\ \vdots \\ x_{np} \end{pmatrix}$$

and that the responses \mathbf{y} fall into two classes $y_1, \dots, y_n \in \{-1, 1\}$.
Further, assume that it is possible to separate the training observations perfectly according to their class.



The three lines displayed in the figure are three possible separating hyperplanes for this dataset which contains two predictors x_1 and x_2 ($p = 2$). The hyperplanes have the property that

$$\beta_0 + \beta_1 x_{i1} + \beta_2 x_{i2} + \dots + \beta_p x_{ip} = \beta_0 + \mathbf{x}_i^T \boldsymbol{\beta} > 0$$

if $y_i = 1$ (green points) and

$$\beta_0 + \beta_1 x_{i1} + \beta_2 x_{i2} + \dots + \beta_p x_{ip} = \beta_0 + \mathbf{x}_i^T \boldsymbol{\beta} < 0$$

if $y_i = -1$ (orange points).

This means that for all observations (all are correctly classified)

$$y_i(\beta_0 + \mathbf{x}_i^T \boldsymbol{\beta}) > 0$$

The hyperplane leads to a natural classifier:

We can assign a class to a new observation depending on which side of the hyperplane it is located. We denote the new observation $\mathbf{x}^* = (x_1^*, \dots, x_n^*)$ and classify it as $y^* = 1$ if

$$f(\mathbf{x}^*) = \beta_0 + \beta_1 x_1^* + \beta_2 x_2^* + \dots + \beta_p x_p^* > 0$$

and as $y^* = -1$ otherwise.

The next question is which hyperplane we should choose.

In the above figure we plotted three possible hyperplanes, but there exist infinitely many possible separating hyperplanes for this dataset.

A natural choice is the **maximal margin hyperplane**. This is the separating hyperplane that is farthest from the training observations.

(From a statistical point of view we might be afraid that we are overfitting the data now.)

Optimization problem

- ▶ The maximal margin hyperplane is found by computing the perpendicular distance from each training observation to a given separating hyperplane.
- ▶ The smallest such distance is the minimal distance from the observations to the hyperplane, also known as the margin. (See illustration below.)
- ▶ We want to maximize this margin.




Figure 1: ISLR Figure 9.3

The process of finding the maximal margin hyperplane for a dataset with p covariates and n training observations can be formulated through the following optimization problem:

$$\text{maximize}_{\beta_0, \beta_1, \dots, \beta_p} \quad M$$

$$\text{subject to } \sum_{j=1}^p \beta_j^2 = 1,$$

$$y_i(\beta_0 + \beta_1 x_{i1} + \beta_2 x_{i2} + \dots + \beta_p x_{ip}) \geq M \quad \forall i = 1, \dots, n$$

where M is the width of the margin.

Observe: $y_i(\beta_0 + \mathbf{x}^T \boldsymbol{\beta})$ is the (signed) distance from the i th point to the hyperplane defined by the β s. We want to find the hyperplane, where each observation is at least M units away - on the correct side, where M is as big as possible.

Above three of the observations are equidistant from the hyperplane. These are called **support vectors**. If one of the support vectors changes its position, the whole hyperplane will move. This is a property of the maximal margin hyperplane: It only depends on the support vectors, and not on the other observations.

It can be shown, see for example Efron and Hastie (2016) Section 19.1 and Friedman, Hastie, and Tibshirani (2001) Section 4.5, that the optimization problem can be reformulated using Lagrange multipliers (primal and dual problem) into a quadratic convex optimization problem that can be solved efficiently.

However, we do of course have to solve the optimization problem to identify the support vectors and the unknown parameters for the separating hyperplane.


Since we in TMA4268 Statistical learning do not require a course in optimization - we do not go into details here.

Questions

- ▶ Explain briefly the idea behind the maximal margin classifier.
- ▶ Is there any tuning parameters that need to be chosen?
- ▶ What if our problem is not separable by a hyperplane?

Support Vector Classifiers

For some data sets a separating hyperplane does not exist, it is *non-separable*.



Also, in some situation we what to allow for some misclassifications to make the class boundaries more robust to future observations - that is, we have noisy data or outliers are present.

In the special case where we have more predictors than observations it is possible to find a separating hyperplane, but the might not be the “best” hyperplane for us.

We now relax the maximal margin classifier to allow for a *soft-margin classifier*.

Optimization problem

$$\text{maximize}_{\beta_0, \beta_1, \dots, \beta_p, \epsilon_1, \dots, \epsilon_n} M \text{ subject to } \sum_{j=1}^p \beta_j^2 = 1$$

$$y_i(\beta_0 + \beta_1 x_{i1} + \beta_2 x_{i2} + \dots + \beta_p x_{ip}) \geq M(1 - \epsilon_i) \quad \forall i = 1, \dots, n.$$

$$\epsilon_i \geq 0, \quad \sum_{i=1}^n \epsilon_i \leq C.$$

- ▶ M is the width of the margin.
- ▶ $\epsilon_1, \dots, \epsilon_n$ are *slack variables*.
 - ▶ If $\epsilon_i = 0$ it means that observation i is on the correct side of the margin,
 - ▶ if $\epsilon_i > 0$ observation i is on the wrong side of the margin, and
 - ▶ if $\epsilon_i > 1$ observation i is on the wrong side of the hyperplane.
- ▶ C is a *tuning (regularization) parameter* (chosen by cross-validation) giving the *budget for slacks*. It restricts the number of the training observations that can be on the wrong

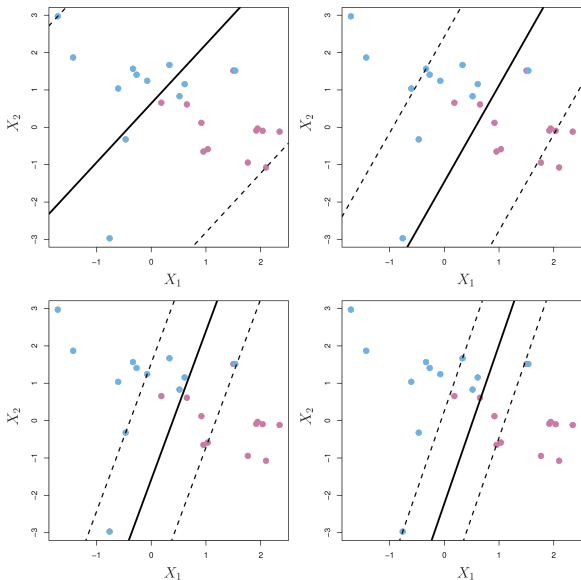


Figure 2: ISLR Figure 9.7: Top left=large C , smaller top right, bottom left and bottom right. As C decreases the tolerance for observations being on the wrong side of the margin decreases and the margin narrows.

The hyperplane has the property that it **only** depends on the observations that **either lie on the margin or on the wrong side of the margin**.

These observations are called our **support vectors**. The observations on the correct side of the margin do not affect the support vectors. The length of distance for the support vectors to the class boundary is proportional to the slacks.

Classification rule: We classify a test observation \mathbf{x}^* based on the sign of $f(\mathbf{x}^*) = \beta_0 + \beta_1 x_1^* + \dots + \beta_p x_p^*$ as before:

- ▶ If $f(\mathbf{x}^*) < 0$ then $y^* = -1$.
- ▶ If $f(\mathbf{x}^*) > 0$ then $y^* = 1$.

More on solving the optimization problem: Friedman, Hastie, and Tibshirani (2001) Section 12.2.1 (primal and dual Lagrange problem, quadratic convex problem).

Questions

- ▶ Should the variables be standardized before used with this method?
- ▶ The support vector classifier only depends on the observations that violate the margin. How does C affect the width of the margin?
- ▶ Discuss how the tuning parameter C affects the bias-variance trade-off of the method.

Example

We will now find a support vector classifier for the second training dataset (forest2) and use this to classify the observations in the second test set (seeds2).

- ▶ There are 100 observations of trees: 45 pines ($y_i = 1$) and 55 redwood trees ($y_i = -1$).
- ▶ In the test set there are 20 seeds: 10 pine seeds and 10 redwood seeds.

The function `svm` in the package `e1071` is used to find the maximal margin hyperplane. The response needs to be coded as a factor variable, and the data set has to be stored as a dataframe.

```
library("e1071")
forest2=read.table(file="https://www.math.ntnu.no/emner/TMA4268/2018v/d
seeds2=read.table(file="https://www.math.ntnu.no/emner/TMA4268/2018v/d
train2=data.frame(x=forest2[,1:2], y=as.factor(forest2[,3]))
test2=data.frame(x=seeds2[,1:2], y=as.factor(seeds2[,3]))
```

The `svm` function uses a slightly different formulation from what we wrote above.


We had in our presentation a budget for errors C , but in `svm` we instead have an argument `cost` that allows us to specify the cost of violating the margin.

- ▶ When `cost` is set to a low value, the margin will be wider than if set to a large value.

We first try with `cost=1`. We set `kernel='linear'` as we are interested in a linear decision boundary. `scale=TRUE` scales the predictors to have mean 0 and standard deviation 1. We choose not to scale.

```
svmfit_linear1=svm(y ~ ., data=train2, kernel='linear', cost=1, scale=F)  
plot(svmfit_linear1,train2,col=c("lightcoral","lightgreen"))
```

SVM classification plot



```
summary(svmfit_linear1)
```


Observations

- ▶ Remark that the x_1 is plotted on the vertical axis, and the the implementation of the plotting function is made in a way that the linear boundary looks jagged.
- ▶ The crosses in the plot indicate the support vectors. With $cost = 1$, we have 56 support vectors, 28 in each class.
- ▶ All other observations are shown as circles.

Next, we set `cost = 100`:

```
svmfit_linear2=svm(y ~ ., data=train2, kernel='linear', cost=100, scale=1)
plot(svmfit_linear2,train2,col=c("lightcoral","lightgreen"))
```

SVM classification plot



With $cost = 100$ we have 31 support vectors, i.e the width of the margin is decreased.

How do we find an optimal $cost$ parameter? By using the `tune()` function we can perform ten-fold cross-validation and find the cost-parameter that gives the lowest cross-validation error:

```
set.seed(1)
CV_linear=tune(svm,y~.,data=train2,kernel="linear",ranges=list(cost=c(0
summary(CV_linear)
```

```
##
## Parameter tuning of 'svm':
##
## - sampling method: 10-fold cross validation
##
## - best parameters:
##   cost
##   0.1
##
## - best performance: 0.15
##
## - Detailed performance results:
##   cost error dispersion
```

According to the `tune()` function we should set the cost parameter to 0.1. The function also stores the best model obtained and we can access it as follows:


```
bestmod_linear=CV_linear$best.model
```

Next, we want to predict the class label of the seeds in the test set. We use the `predict` function and make a confusion table:


```
ypred_linear=predict(bestmod_linear,test2)
table(predict=ypred_linear,truth=test2[,3])
```

```
##           truth
## predict -1  1
##        -1  9  2
##         1  1  8
```

True class



Predicted class




In this case three of the test observations are misclassified: These three observations are marked with a black circle in the plot, and we observe that they lie on the border between the green and the orange points which is reasonable: The test observations located on the border between green and orange are hardest to predict.

Missing: the `svm` function is not (directly) outputting the equation for the class boundary, and not the value for the width of the margin. Want to see how to find this? Go to the recommended exercises.

Support Vector Machines

For some datasets a non-linear decision boundary between the classes is more suitable than a linear decision boundary. In such cases you can use a **Support Vector Machine** (SVM). This is an extension of the support vector classifier.



Expanding the feature space

We saw in Module 7 that in regression we could fit non-linear curves by using a polynomial basis - adding polynomials of different order as covariates. This was a linear regression in the transformed variables, but non-linear in the original variables. Maybe we may add many such extra features and find a nice linear boundary in that high-dimensional space?

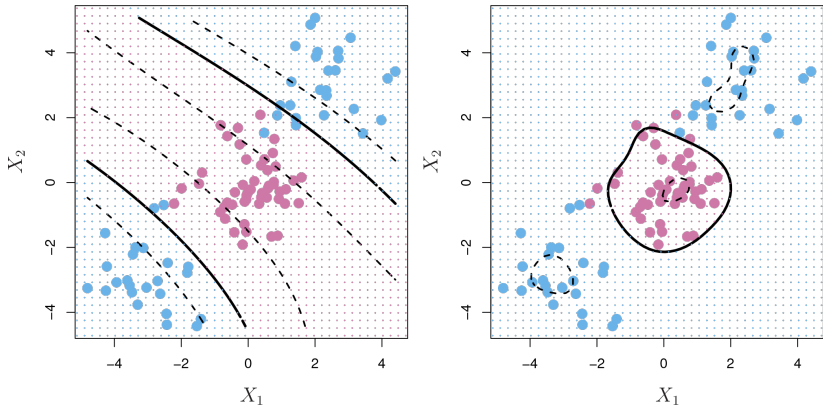


Figure 3: ISLR Figure 9.9

Left: expanding feature space to include cubic polynomials (9 parameters to estimate), and also observe the margins. (Right: radial basis function kernel - wait a bit.)

Next: replace polynomials with *kernels* for elegance and computational issues.

Inner products

We have not focused on how to solve the optimisation problem of finding the support vector classifier hyperplane, because this is outside the scope of this course.

However, it can be *shown* that the solution to the support vector classifier problem can be expressed as

$$f(\mathbf{x}) = \beta_0 + \sum_{i=1}^n \alpha_i \langle \mathbf{x}, \mathbf{x}_i \rangle$$

where α_i is some parameter and $i = 1, \dots, n$. The term $\langle \mathbf{x}_i, \mathbf{x}_{i'} \rangle$ denotes the inner product between two observations and is defined as:

$$\langle \mathbf{x}_i, \mathbf{x}_{i'} \rangle = \sum_{j=1}^p x_{ij} x_{i'j}.$$

This means that we need to estimate n parameters instead of p (and for our expanded feature space then p might be larger than n). (For the interested reader: See Eq. 19.22 and 19.23 of Efron and Tibshirani (2016).)

Further, it then turns out that to estimate the parameters $\beta_0, \alpha_1, \dots, \alpha_n$ this can be based on the $\binom{n}{2}$ inner products $\langle x_i, x_j' \rangle$ between all pair of training observations.

Also, $\alpha_i = 0$ for the observations i that are *not* the support vectors. Thus, we only need the inner product between the training observations and the observations corresponding to support vectors, and

$$f(x) = \beta_0 + \sum_{i \in \mathcal{S}} \alpha_i \langle x, x_i \rangle,$$

where \mathcal{S} contains the indices of the support points. So, we have sparsity in the observations.

Remark: we could alternatively say that $\alpha_i = 0$ define the support vectors.

Q: Find the support vectors

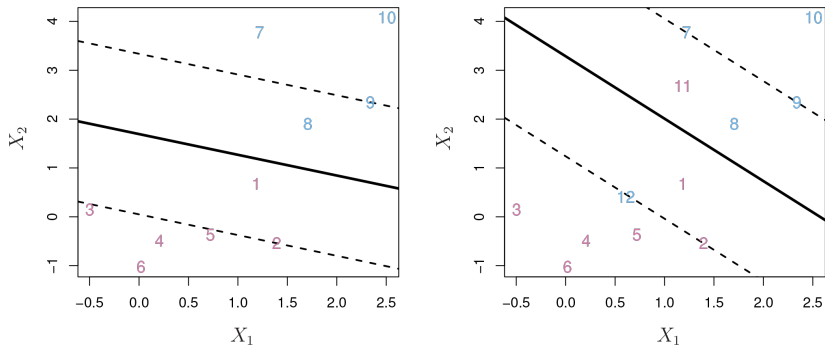


Figure 4: ISLR Figure 9.6

Observe: we need all observations to find the support vectors.

Kernels

The next step is now to *replace the inner product* $\langle x, x_i \rangle$ with a function $K(x_i, x_i')$ referred to as the **kernel**:

$$f(x) = \beta_0 + \sum_{i \in S} \alpha_i K(x, x_i).$$

For the linear case (which is what we have considered so far), the kernel is simply the inner product $K(x_i, x_i') = \sum_{j=1}^p x_{ij} x_{i'j}$.

The two arguments to the kernel are two p -vectors.

If we want a more flexible decision boundary we could instead use a **polynomial kernel**. This polynomial kernel of degree $d > 1$ is given by:

$$K(x_i, x_i') = \left(1 + \sum_{i=1}^p x_{ij} x_{i'j}\right)^d.$$

(This kernel is not so much used in practice, but is popular for proofs.)

Using these kernels our solution for the class boundary can be written

$$\hat{f}(\mathbf{x}) = \beta_0 + \sum_{i \in \mathcal{S}} \hat{\alpha}_i K(\mathbf{x}, \mathbf{x}_i)$$

The nice thing here is that we only need to calculate the kernels, not the basis functions (what we in Modul 7 did as extra columns of the design matrix).

A very popular choice is the radial kernel,

$$K(x_i, x'_i) = \exp(-\gamma \sum_{j=1}^p (x_{ij} - x'_{ij})^2),$$

where γ is a positive constant (a tuning parameter).

Observe the connection to a multivariate normal density. If γ is small the decision boundaries are smoother than for larger γ .

It turns out that this computes the inner product in a very high (infinite) dimensional feature space. But, this does not give overfitting because some of the dimensions are “squashed down”.

The radial kernel is convenient if we want a circular decision boundary, and γ can be chosen by cross-validation.

Study Figures 19.5 and 19.6 from Efron and Hastie (2016) to see how the radial kernel can make smooth functions.

Computer Age Statistical Inference

Kernels and our optimization

We now merge our optimization problem (from our support vector classifier) with our kernel representation $f(\mathbf{x})$ to get the Support Vector Machine (SVM).

$$\begin{aligned} & \text{maximize}_{\beta_0, \alpha_1, \dots, \alpha_n, \epsilon_1, \dots, \epsilon_n} \quad M \\ & y_i(f(\mathbf{x}_i)) \geq M(1 - \epsilon_i) \quad \forall i = 1, \dots, n. \\ & \epsilon_i \geq 0, \quad \sum_{i=1}^n \epsilon_i \leq C. \end{aligned}$$

where

$$f(\mathbf{x}) = \beta_0 + \sum_{i \in \mathcal{S}} \alpha_i K(\mathbf{x}, \mathbf{x}_i)$$

Regularization parameter example

Heart data - predict heart disease from $p = 13$ predictors.

Training errors as ROC and AUC.

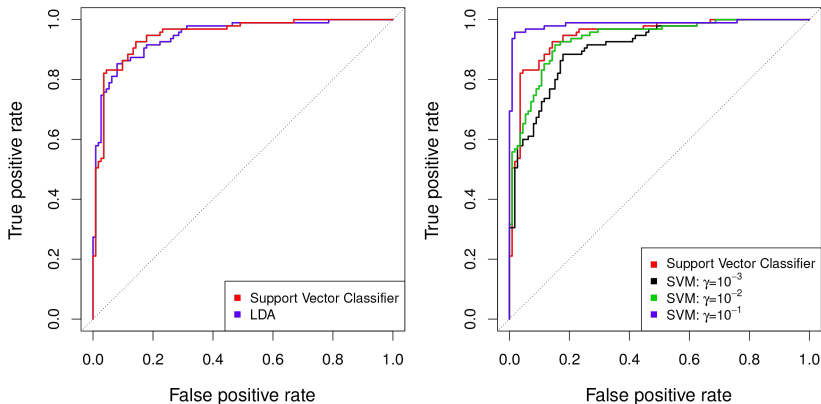


Figure 5: ISLR Figure 9.10

Heart data - test error.

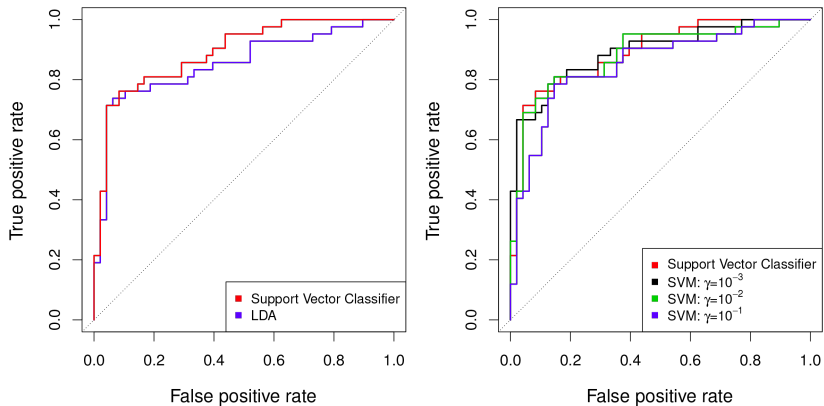


Figure 6: ISLR Figure 9.11

Example: forest 3


To illustrate the SVM we use the third training dataset (forest3) and the third test set (seeds3). We use the `svmfunction` as before. However, we now set `kernel='radial'` as we want a non-linear decision boundary:

```
library("e1071")
forest3=read.table(file="forest3.txt");
seeds3=read.table(file="seeds3.txt")

train3=data.frame(x=forest3[,1:2], y=as.factor(forest3[,3]))
test3=data.frame(x=seeds3[,1:2], y=as.factor(seeds3[,3]))

svmfit_kernel1=svm(y ~ ., data=train3, kernel='radial', gamma=1, cost=1)
plot(svmfit_kernel1,train3,col=c("lightcoral","lightgreen"))
```

SVM classification plot




We could also try with a polynomial kernel with degree 4 as follows:

```
library("e1071")
forest3=read.table(file="forest3.txt");
seeds3=read.table(file="seeds3.txt")

train3=data.frame(x=forest3[,1:2], y=as.factor(forest3[,3]))
test3=data.frame(x=seeds3[,1:2], y=as.factor(seeds3[,3]))

svmfit_kernel2=svm(y ~ ., data=train3, kernel='polynomial')
plot(svmfit_kernel2,train3,col=c("lightcoral","lightgreen"))
```

SVM classification plot



For this dataset a radial kernel is a natural choice: A circular decision boundary seems like a good idea. Thus, we proceed with `kernel='radial'`, and use the `tune()` function to find the optimal tuning parameter C :

```
set.seed(1)
CV_kernel=tune(svm,y~.,data=train3,kernel="radial",gamma=1)
summary(CV_kernel)
```

```
##
## Parameter tuning of 'svm':
##
## - sampling method: 10-fold cross validation
##
## - best parameters:
##   cost
##     10
##
## - best performance: 0.1232143
##
```

The optimal C is 10. Next, we predict the class label of the seeds in the test set with a model with $C=10$, make a confusion table and plot the results:

```
bestmod_kernel=CV_kernel$best.model
ypred_kernel=predict(bestmod_kernel,test3)

par(mfrow=c(1,3)); par(pty="s")
plot(NA,xlab="x1",ylab="x2",xlim=c(0,1),ylim=c(0,1));title
points(seeds3[seeds3[,3]==-1,1:2],pch=19,col="lightcoral",c
points(seeds3[seeds3[,3]==1,1:2],pch=19,col="darkseagreen"
points(seeds3[which(ypred_kernel!=seeds3[,3]),1:2],pch=21)

plot(NA,xlab="x1",ylab="x2",xlim=c(0,1),ylim=c(0,1));title
points(seeds3[ypred_kernel==-1,1:2],pch=19,col="lightcoral"
points(seeds3[ypred_kernel==1,1:2],pch=19,col="darkseagreen"
points(seeds3[which(ypred_kernel!=seeds3[,3]),1:2],pch=21)

table(predict=ypred_kernel,truth=test3[,3])
```

Extensions

More than two classes

What if we have k classes?

- ▶ OVA: one-versus-all. Fit k different two-class SVMs $f_k(x)$ where one class is compared to all other classes. Classify a test observation to the class where $f_k(x^*)$ is largest.
- ▶ OVO: For multiclass-classification with k levels, $k > 2$, libsvm uses the 'one-against-one'-approach, in which $k(k-1)/2$ binary classifiers are trained; the appropriate class is found by a voting scheme (the class that wins the most pairwise competitions are chosen).

Comparisons

Focus is comparing the support vector classifier and logistic regression

It is possible to write the optimization problem for the support vector classifier as a “loss”+“penalty”:

$$\text{minimize}_{\beta} \left\{ \sum_{i=1}^n \max(0, 1 - y_i f(\mathbf{x}_i)) + \lambda \sum_{j=1}^p \beta_j^2 \right\}$$

- ▶ the loss is called *hinge loss* - observe the max and 0 to explain why only support vectors contribute
- ▶ the penalty is a ridge penalty
- ▶ large λ gives β s small and more violations=high bias, but low variance
- ▶ small λ gives β s large and less violations=low bias, but high variance




Figure 7: ISLR Figure 9.12: hinge loss - loss 0 for observations on the correct side of the margin

Hinge loss:

$$\max(0, 1 - y_i f(\mathbf{x}_i))$$

For comparison a logistic regression (with ridge penalty) would be (binomial deviance with -1,1 coding of y)

$$\log(1 + \exp(-y_i f(\mathbf{x}_i)))$$

It can be shown that in logistic regression all observations contribute weighted by $p_i(1 - p_i)$ (where p_i is probability for class 1), that fade smoothly with distance to the decision boundary

It is possible to extend the logistic regression to include non-linear terms, and ridge penalty,

When to use SVM?

- ▶ If classes are nearly separable SVM will perform better than logistic regression. (LDA will also perform better than logistic regression.)
- ▶ and if not, then a ridge penalty version of logistic regression are very similar to SVM, and logistic regression will also give you probabilities for each class.
- ▶ If class boundaries are non-linear then SVM is more popular, but kernel versions of logistic regression is possible, but more computationally expensive.

Summing up

- ▶ We use methods from computer science, not probability models
- but looks for a separating hyperplane in (an extended) feature space in the classification setting.
- ▶ SVM is a widely successful and a “must have tool”
- ▶ Interpretation of SVM: all features are included and maybe not so easy to interpret.
- ▶ Not so easy to get class probabilities from SVM (what is done is actually to fit a logistic regression after fitting SVM).

Recommended exercises

1. Understanding the algorithms:

- ▶ Exercise 1, 2 and 3 in the book.

2. Data analysis

- ▶ Go back and read in the forest1 data (is located in the same place as forest2) and run the svm with a very high value for cost. The forest1 is a separable problem.
- ▶ Linear version of SVM: Making nicer plots for SVM from Lab video. Go through the code and see what is happening (and see the video if you want more explanation).

```
# code taken from video by Trevor Hastie linked above  
library(e1071)  
# fake data  
set.seed(10111)  
x=matrix(rnorm(40),20,2)
```

R packages

These packages needs to be install before knitting this R Markdown file.

```
install.packages("e1071")
```

References

- ▶ Videos on YouTube by the authors of ISL, Chapter 9, and corresponding slides
- ▶ Solutions to exercises in the book, chapter 9

Efron, Bradley, and Trevor Hastie. 2016. *Computer Age Statistical Inference - Algorithms, Evidence, and Data Science*. Cambridge University Press.

Friedman, Jerome, Trevor Hastie, and Robert Tibshirani. 2001. *The Elements of Statistical Learning*. Vol. 1. Springer series in statistics New York.

James, Gareth, Daniela Witten, Trevor Hastie, and Robert Tibshirani. 2013. *An Introduction to Statistical Learning*. Vol. 112. Springer.

Karatzoglou, A., D. Meyer, and K. Hornik. 2006. "Support Vector Machines in R." *Journal of Statistical Software* 15 (9).