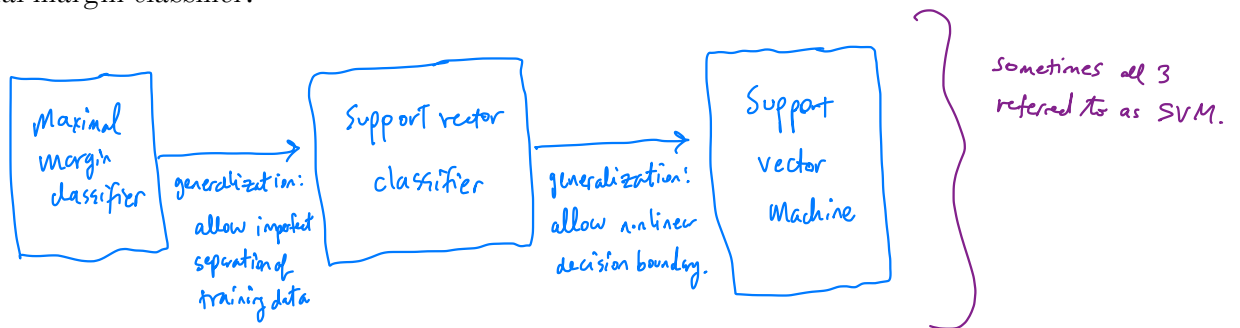


Chapter 9: Support Vector Machines

The *support vector machine* is an approach for classification that was developed in the computer science community in the 1990s and has grown in popularity. *↙ Categorical response.*


- SVMs perform well in a variety of settings.
- Considered one of the best "out of the box" classifiers.

The support vector machine is a generalization of a simple and intuitive classifier called the *maximal margin classifier*.



Support vector machines are intended for binary classification, but there are extensions for more than two classes.

categorical response w/ only 2 classes.




Credit: <https://dilbert.com/strip/2013-02-02>

1 Maximal Margin Classifier

based on a hyperplane separator.

In p -dimensional space, a *hyperplane* is a flat affine ^{→ extension of euclidean space.} subspace of dimension $p - 1$.



e.g. In 2 dimensions, a hyperplane is a flat 1 dimensional subspace. — line.

In 3 dimensions, a hyperplane is a flat 2 dimensional subspace — plane.

⋮

In $p > 3$ dimension, harder to visualize, but still a flat $p-1$ dimensional subspace.

The mathematical definition of a hyperplane is quite simple,

In 2 dimensions, a hyperplane is defined by $\beta_0 + \beta_1 x_1 + \beta_2 x_2 = 0$ ^{parameters}

i.e. any $x = (x_1, x_2)$ for which this equation holds lies on the hyperplane.

Note this is just the equation for a line.

This can be easily extended to the p -dimensional setting.

$\beta_0 + \beta_1 x_1 + \dots + \beta_p x_p = 0$ defines a p -dim hyperplane.

i.e. and $\underline{x} = (x_1, \dots, x_p)$ for which this equation holds lies on the hyperplane.

We can think of a hyperplane as dividing p -dimensional space into two halves.

If $\beta_0 + \beta_1 x_1 + \dots + \beta_p x_p > 0$ then $\underline{x} = (x_1, \dots, x_p)$ lies on one side of the hyperplane

$\beta_0 + \beta_1 x_1 + \dots + \beta_p x_p < 0$ then $\underline{x} = (x_1, \dots, x_p)$ lies on other side of the hyperplane.

You can determine which side of the hyperplane by just determining the sign of

$$\beta_0 + \beta_1 x_1 + \dots + \beta_p x_p.$$

1.1 Classification Using a Separating Hyperplane

Suppose that we have a $n \times p$ data matrix \mathbf{X} that consists of n training observations in p -dimensional space.

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

↑
training observations.

and that these observations fall into two classes.

$$y_1, \dots, y_n \in \{-1, 1\}.$$

where -1 represents one class
 1 represents the other.

We also have a test observation.

$$\underline{x}^* = (x_1^*, \dots, x_p^*)^T \quad p\text{-vector of observed features.}$$


Our Goal: Develop a classifier based on training data that will correctly classify the test observation based on feature measurements.

We already have many approaches:

- random forest, bagging, boosting, trees.
- logistic regression
- LDA, QDA
- KNN

We will see a new approach using a separating hyperplane.

Suppose it is possible to construct a hyperplane that separates the training observations perfectly according to their class labels.



Then a separating hyperplane has the property that

$$\beta_0 + \beta_1 x_{i1} + \dots + \beta_p x_{ip} > 0 \iff y_i = 1 \quad \text{and}$$

$$\beta_0 + \beta_1 x_{i1} + \dots + \beta_p x_{ip} < 0 \iff y_i = -1$$

$$\iff$$

$$y_i (\beta_0 + \beta_1 x_{i1} + \dots + \beta_p x_{ip}) \geq 0 \quad \forall i = 1, \dots, n.$$

If a separating hyperplane exists, we can use it to construct a very natural classifier:

a test observation is assigned a class depending on which side of the hyperplane it is located.

That is, we classify the test observation x^* based on the sign of $f(x^*) = \beta_0 + \beta_1 x_1^* + \dots + \beta_p x_p^*$.

$$\text{If } f(x^*) > 0 \quad \text{assign } x^* \text{ to class } 1$$

$$\text{If } f(x^*) < 0 \quad \text{assign } x^* \text{ to class } -1.$$

We can also use the magnitude of $f(x^*)$.

If $f(x^*)$ is far from zero, this means x^* lies far from the hyperplane.

\Rightarrow we can be confident about our class assignment for x^* .


If $f(x^*)$ is close to zero, this means x^* is located near the hyperplane.

\Rightarrow we are less sure about class assignment.

Note: a classifier based on a separating hyperplane leads to a linear decision boundary.

1.2 Maximal Margin Classifier

If our data can be perfectly separated using a hyperplane, then there will exist an infinite number of such hyperplanes.




a given separating hyperplane can be shifted a tiny bit or rotated without coming into contact w/ any observations.

⇒ which one to use for our classifier?

A natural choice for which hyperplane to use is the maximal margin hyperplane (aka the optimal separating hyperplane), which is the hyperplane that is farthest from the training observations.

- We compute the perpendicular distance from each observation to a given separating hyperplane
- the smallest distance is known as the margin.

The maximal margin hyperplane is the one w/ largest margin, i.e. furthest from all training data points.



$$M_2 < M_1$$

⇒ blue hyperplane larger margin.


⇒ blue is my preferred hyperplane.

We can then classify a test observation based on which side of the maximal margin hyperplane it lies – this is the maximal margin classifier.

- hopefully a large margin on training data will lead to a large margin on test data

⇒ classify test data correctly.

- When p is large, overfitting will occur.



if these point moves the maximal margin hyperplane would move as well.

These are known as support vectors because they are p -dim vectors that "support" the hyperplane.

NOTE: the maximal margin hyperplane only depends on the support vectors!

The rest of the points can move and it doesn't matter.

← a small # of points

We now need to consider the task of constructing the maximal margin hyperplane based on a set of n training observations and associated class labels.

$$x_1, \dots, x_n \in \mathbb{R}^p \quad y_1, \dots, y_n \in \{-1, 1\}.$$

The maximal margin hyperplane is the solution to the optimization problem

① Maximize M ← margin

$\beta_0, \dots, \beta_p, M$

Subject to $\sum_{j=1}^p \beta_j^2 = 1$ ②

$$y_i (\beta_0 + \beta_1 x_{i1} + \dots + \beta_p x_{ip}) \geq M \quad \forall i=1, \dots, n. \quad \textcircled{3}$$

③ means each observation will be on the correct side of the hyperplane ($M \geq 0$) with some cushion (if $M > 0$).

② ensure $y_i (\beta_0 + \beta_1 x_{i1} + \dots + \beta_p x_{ip})$ is perp. distance to hyperplane and ③ means the point x_i is at least M distance away $\Rightarrow M$ is the margin.

① chooses $\beta_0, \dots, \beta_p, M$ to maximize the margin.

\Rightarrow maximal margin hyperplane!

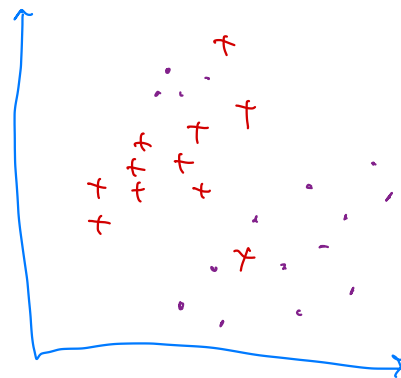
This problem can be solved efficiently, but the details are outside the scope of this course.

\hookrightarrow we'll talk a little bit more later.

What happens when no separating hyperplane exists?

\Rightarrow no maximal margin hyperplane!

We can develop a hyperplane that almost separates the classes - a "soft margin"



We can't draw a hyperplane to separate these perfectly!

2 Support Vector Classifiers

It's not always possible to separate training observations by a hyperplane. In fact, even if we can use a hyperplane to perfectly separate our training observations, it may not be desirable.

We might be willing to consider a classifier based on a hyperplane that does *not perfectly* separate the two classes in the interest of

-
-

The *support vector classifier* does this by finding the largest possible margin between classes, but allowing some points to be on the “wrong” side of the margin, or even on the “wrong” side of the hyperplane.

The support vector classifier classifies a test observation depending on which side of the hyperplane it lies. The hyperplane is chosen to correctly separate **most** of the training observations.

Once we have solved this optimization problem, we classify \boldsymbol{x}^* as before by determining which side of the hyperplane it lies.

ϵ_i

C

The optimization problem has a very interesting property.


Observations that lie directly on the margin or on the wrong side of the margin are called *support vectors*.

The fact that only support vectors affect the classifier is in line with our assertion that C controls the bias-variance tradeoff.

Because the support vector classifier's decision rule is based only on a potentially small subset of the training observations means that it is robust to the behavior of observations far away from the hyperplane.

3 Support Vector Machines

The support vector classifier is a natural approach for classification in the two-class setting...



We've seen ways to handle non-linear classification boundaries before.

In the case of the support vector classifier, we could address the problem of possible non-linear boundaries between classes by enlarging the feature space.

Then our optimization problem would become

The *support vector machine* allows us to enlarge the feature space used by the support classifier in a way that leads to efficient computation.

It turns out that the solution to the support vector classification optimization problem involves only *inner products* of the observations (instead of the observations themselves).


It can be shown that

-

-

-

Now suppose every time the inner product shows up in the SVM representation above, we replaced it with a generalization.



4 SVMs with More than Two Classes

So far we have been limited to the case of binary classification. How can we extend SVMs to the more general case with some arbitrary number of classes?

Suppose we would like to perform classification using SVMs and there are $K > 2$ classes.

One-Versus-One

One-Versus-All