

Package ‘svmpath’

October 21, 2009

Title svmpath: the SVM Path algorithm

Date 2009-10-21

Version 0.93

Author Trevor Hastie

Description Computes the entire regularization path for the two-class svm classifier with essentially the same cost as a single SVM fit.

Maintainer Trevor Hastie <hastie@stanford.edu>

License GPL-2

URL <http://www.jmlr.org/papers/volume5/hastie04a/hastie04a.pdf>

Repository CRAN

Date/Publication 2009-10-21 14:15:53

R topics documented:

balanced.overlap	2
predict.svmpath	2
print.svmpath	4
radial.kernel	5
summary.svmpath	6
svmpath	7

Index	9
--------------	----------

balanced.overlap *simple examples for svmpath*

Description

Datasets for illustrating the svmpath function, that can be plotted while its running

Usage

```
data(svmpath)
```

Format

In each case a list with a component `x` (t column matrix) and a component `y` (vector of +1/-1 values). "Balanced" refers to whether the number of +1s is the same as the -1s. "Overlap" indicates whether the classes are linearly separable. `mixture.data` is a balanced dataset with 100 observations in each class. The others are smaller with between 10-12 obs total.

References

The paper <http://www-stat.stanford.edu/~hastie/Papers/svmpath.pdf>, as well as the talk <http://www-stat.stanford.edu/~hastie/TALKS/svmpathtalk.pdf>.

Examples

```
data(svmpath)
attach(balanced.overlap)
svmpath(x, y, trace=TRUE, plot=TRUE)
detach(2)
```

predict.svmpath *Make predictions from a "svmpath" object*

Description

Provide a value for `lambda`, and produce the fitted lagrange alpha values. Provide values for `x`, and get fitted function values or class labels.

Usage

```
## S3 method for class 'svmpath':
predict(object, newx, lambda, type = c("function", "class",
"alpha", "margin"), ...)
```

Arguments

object	fitted svmpath object
newx	values of x at which prediction are wanted. This is a matrix with observations per row
lambda	the value of the regularization parameter. Note that lambda is equivalent to $1/C$ for the usual parametrization of a SVM
type	type of prediction, with default "function". For type="alpha" or type="margin" the newx argument is not required
...	Generic compatibility

Details

This implementation of the SVM uses a parameterization that is slightly different but equivalent to the usual (Vapnik) SVM. Here $\lambda = 1/C$. The Lagrange multipliers are related via $\alpha_i^* = \alpha_i/\lambda$, where α_i^* is the usual multiplier, and α_i our multiplier. Note that if alpha=0, that observation is right of the elbow; alpha=1, left of the elbow; $0 < \text{alpha} < 1$ on the elbow. The latter two cases are all support points.

Value

In each case, the desired prediction.

Author(s)

Trevor Hastie

References

The paper <http://www-stat.stanford.edu/~hastie/Papers/svmpath.pdf>, as well as the talk <http://www-stat.stanford.edu/~hastie/TALKS/svmpathtalk.pdf>.

See Also

coef.svmpath, svmpath

Examples

```
data(svmpath)
attach(balanced.overlap)
fit <- svmpath(x,y,trace=TRUE,plot=TRUE)
predict(fit, lambda=1,type="alpha")
predict(fit, x, lambda=.9)
detach(2)
```

print.svmpath *Print a summary of the SVM path*

Description

print a summary of the fitted svmpath object

Usage

```
## S3 method for class 'svmpath':  
print(x, digits, maxsteps, ...)
```

Arguments

x	object to be printed
digits	number of significant digits (default 6)
maxsteps	the number of steps to print; default all
...	additional arguments to the generic print function

Value

For each step taken by the algorithm, one or more lines are printed. The step is described in terms of the observation number involved, a coded version of what happened, such as "L->E" meaning "from the Left set" to the "Elbow". Initially all the sets are empty. It gives the margin (sum of the xi), the size of the elbow, and the training error.

Author(s)

Trevor Hastie

References

The paper <http://www-stat.stanford.edu/~hastie/Papers/svmpath.pdf>, as well as the talk <http://www-stat.stanford.edu/~hastie/TALKS/svmpathtalk.pdf>.

See Also

coef.svmpath, svmpath, predict.svmpath

Examples

```
data(svmpath)  
attach(balanced.overlap)  
fit <- svmpath(x, y, trace=TRUE, plot=TRUE)  
print(fit)  
detach(2)
```

radial.kernel *compute the kernel matrix for svm*

Description

compute the kernel matrix for svm

Usage

```
radial.kernel(x, y=x, param.kernel = 1/p)
poly.kernel(x, y=x, param.kernel = 1)
```

Arguments

`x` an $n \times p$ matrix of features
`y` an $m \times p$ matrix of features (if omitted, it defaults to `x`)
`param.kernel` the parameter(s) for the kernel. For this radial kernel, the parameter is known in the fields as "gamma". For the polynomial kernel, it is the "degree"

Details

For the radial kernel, this computes the function $\exp(-\gamma\|x - y\|^2)$ for each pair of rows `x,y` from the input matrices. Here `g` is `param.kernel`. For the polynomial kernel, it computes $(xy^T + 1)^d$, where `d` is `param.kernel`.

Value

An $n \times m$ matrix.

Author(s)

Trevor Hastie

References

The paper <http://www-stat.stanford.edu/~hastie/Papers/svmpath.pdf>, as well as the talk <http://www-stat.stanford.edu/~hastie/TALKS/svmpathtalk.pdf>.

See Also

svmpath

Examples

```
data(svmpath)
attach(balanced.overlap)
fit<-svmpath(x,y,kernel=radial.kernel)
detach(2)
```

summary.svmpath *produce a summary of an svmpath object*

Description

printing an svmpath object can produce a lot of lines. The summary methods gives a more concise description by picking out a subset of the steps

Usage

```
## S3 method for class 'svmpath':  
summary(object, nsteps = 5, digits = 6, ...)
```

Arguments

object	the svmpath object
nsteps	usually omitted, but can be changed to get longer summaries
digits	number of significant digits
...	additional arguments to the generic summary function

Details

Uses the `pretty` function to extract the approximately the desired number of steps. Always includes the first and last step.

Value

returns a dataframe with the steps, value of lambda, training error, size of elbow, number of support points, and the sum of the overlaps

Author(s)

Trevor Hastie

References

The paper <http://www-stat.stanford.edu/~hastie/Papers/svmpath.pdf>, as well as the talk <http://www-stat.stanford.edu/~hastie/TALKS/svmpathtalk.pdf>.

See Also

`coef.svmpath`, `svmpath`, `predict.svmpath`, `print.svmpath`

Examples

```

data(svmpath)
attach(balanced.overlap)
fit <- svmpath(x,y,trace=TRUE,plot=TRUE)
summary(fit)
detach(2)

```

svmpath

*Fit the entire regularization path for a 2-class SVM***Description**

The SVM has a regularization or cost parameter C , which controls the amount by which points overlap their soft margins. Typically either a default large value for C is chosen (allowing minimal overlap), or else a few values are compared using a validation set. This algorithm computes the entire regularization path (i.e. for all possible values of C for which the solution changes), with a cost a small (~ 3) multiple of the cost of fitting a single model.

Usage

```
svmpath(x, y, K, kernel.function = poly.kernel, param.kernel = 1, trace, plot.it, 1
```

Arguments

<code>x</code>	the data matrix ($n \times p$) with n rows (observations) on p variables (columns)
<code>y</code>	The "-1,+1" valued response variable.
<code>K</code>	a $n \times n$ kernel matrix, with default value <code>K= kernel.function(x, x)</code>
<code>kernel.function</code>	This is a user-defined function. Provided are <code>poly.kernel</code> (the default, with parameter set to default to a linear kernel) and <code>radial.kernel</code>
<code>param.kernel</code>	parameter(s) of the kernels
<code>trace</code>	if TRUE, a progress report is printed as the algorithm runs; default is FALSE
<code>plot.it</code>	a flag indicating whether a plot should be produced (default FALSE; only usable with $p=2$)
<code>linear.plot</code>	if TRUE, the plotting routine exploits the knowledge that the solution is linear; otherwise a contour algorithm is used. The default is <code>missing(kernel)</code> (i.e. TRUE if a default linear kernel is used)
<code>eps</code>	a small machine number which is used to identify minimal step sizes
<code>Nmoves</code>	the maximum number of moves
<code>digits</code>	the number of digits in the printout
<code>lambda.min</code>	The smallest value of $\lambda = 1/C$; default is <code>lambda=10e-4</code> , or $C=10000$
<code>...</code>	additional arguments to some of the functions called by <code>svmpath</code>

Details

The algorithm used in `svmpath()` is described in detail in "The Entire Regularization Path for the Support Vector Machine" by Hastie, Rosset, Tibshirani and Zhu (2004). It exploits the fact that the "hinge" loss-function is piecewise linear, and the penalty term is quadratic. This means that in the dual space, the lagrange multipliers will be piecewise linear (c.f. `lars`).

Value

a "svmpath" object is returned, for which there are `print`, `summary`, `coef` and `predict` methods.

Warning

Currently the algorithm can get into machine errors if `epsilon` is too small, or if `lambda.min` is too small. Increasing either from their defaults should make the problems go away, by terminating the algorithm slightly early.

Note

This implementation of the algorithm does not use updating to solve the "elbow" linear equations. This is possible, since the elbow changes by a small number of points at a time. Future version of the software will do this. The author has encountered numerical problems with early attempts at this.

Author(s)

Trevor Hastie

References

The paper <http://www-stat.stanford.edu/~hastie/Papers/svmpath.pdf>, as well as the talk <http://www-stat.stanford.edu/~hastie/TALKS/svmpathtalk.pdf>.

See Also

`print`, `coef`, `summary`, `predict`, and `FilmPath`

Examples

```
data(svmpath)
attach(unbalanced.separated)
svmpath(x, y, trace=TRUE, plot=TRUE)
detach(2)
## Not run: svmpath(x, y, kernel=radial.kernel, param.kernel=.8)
```

Index

*Topic **datasets**

`balanced.overlap`, 1

*Topic **regression**

`predict.svmpath`, 2

`print.svmpath`, 3

`radial.kernel`, 4

`summary.svmpath`, 5

`svmpath`, 6

`balanced.overlap`, 1

`balanced.separated`

(balanced.overlap), 1

`mixture.data(balanced.overlap)`, 1

`poly.kernel(radial.kernel)`, 4

`predict.svmpath`, 2

`print.svmpath`, 3

`radial.kernel`, 4

`summary.svmpath`, 5

`svmpath`, 6

`unbalanced.separated`

(balanced.overlap), 1