

A feature selection algorithm can be seen as the combination of a search technique for proposing new feature subsets, along with an evaluation measure which scores the different feature subsets. The simplest algorithm is to test each possible subset of features finding the one which minimises the error rate. This is an exhaustive search of the space, and is computationally intractable for all but the smallest of feature sets. The choice of evaluation metric heavily influences the algorithm, and it is these evaluation metrics which distinguish between the three main categories of feature selection algorithms: wrappers, filters and embedded methods.^[1]

task - PCA

Wrapper methods use a predictive model to score feature subsets. Each new subset is used to train a model, which is tested on a hold-out set. Counting the number of mistakes made on that hold-out set (the error rate of the model) gives the score for that subset. As wrapper methods train a new model for each subset, they are very computationally intensive, but usually provide the best performing feature set for that particular type of model.

Filter methods use a proxy measure instead of the error rate to score a feature subset. This measure is chosen to be fast to compute, whilst still capturing the usefulness of the feature set. Common measures include the mutual information,^[1] the pointwise mutual information,^[2] Pearson product-moment correlation coefficient, inter/intra class distance or the scores of significance tests for each class/feature combinations.^{[2][3]} Filters are usually less computationally intensive than wrappers, but they produce a feature set which is not tuned to a specific type of predictive model. Many filters provide a feature ranking rather than an explicit best feature subset, and the cut off point in the ranking is chosen via cross-validation. Filter methods have also been used as a preprocessing step for wrapper methods, allowing a wrapper to be used on larger problems.

Embedded methods are a catch-all group of techniques which perform feature selection as part of the model construction process. The exemplar of this approach is the LASSO method for constructing a linear model, which penalises the regression coefficients, shrinking many of them to zero. Any features which have non-zero regression coefficients are 'selected' by the LASSO algorithm. One other popular approach is the Recursive Feature Elimination algorithm, commonly used with Support Vector Machines to repeatedly construct a model and remove features with low weights. These approaches tend to be between filters and wrappers in terms of computational complexity.

In statistics, the most popular form of feature selection is stepwise regression. It is a greedy algorithm that adds the best feature (or deletes the worst feature) at each round. The main control issue is deciding when to stop the algorithm. In machine learning, this is typically done by cross-validation. In statistics, some criteria are optimized. This leads to the inherent problem of nesting. More robust methods have been explored, such as branch and bound and piecewise linear network.

Subset selection

Subset selection evaluates a subset of features as a group for suitability. Subset selection algorithms can be broken up into Wrappers, Filters and Embedded. Wrappers use a search algorithm to search through the space of possible features and evaluate each subset by running a model on the subset. Wrappers can be computationally expensive and have a risk of over fitting to the model. Filters are similar to Wrappers in the search approach, but instead of evaluating against a model, a simpler filter is evaluated. Embedded techniques are embedded in and specific to a model.

Many popular search approaches use greedy hill climbing, which iteratively evaluates a candidate subset of features, then modifies the subset and evaluates if the new subset is an improvement over the old. Evaluation of the subsets requires a scoring metric that grades a subset of features. Exhaustive search is generally impractical,

Model Sparsity and Feature Selection

1 The "Bet on Sparsity" Principle

book: Elements, 16.2.2

book: Statistics for high dimensional data, introduction

$$Y_i = \mu + \sum_{j=1}^p \beta_j X_i^j + \epsilon_i$$

$$h_{\vec{w}}(\vec{x}) = w^0 + w^1 x^1 + w^2 x^2 + \dots + w^D x^D$$

w.l.o.g. let $x^0 = 1$

Roughly speaking, for High-dimensional statistical inference to achieve reasonable accuracy or asymptotic consistency, we need

$$h_{\vec{w}}(\vec{x}) = \vec{w}^T \vec{x} = \sum_{d=0}^D w^d \cdot x^d$$

$$\log(p)(\text{sparsity}(\beta)) \ll n$$

2 Forward Selection

Forward selection starts with no feature(variable) in the model, and adds features to the model one at a time. At each step, the feature that can contribute most to the model is added. The procedure is repeated until one new feature can improve the model significantly(defined by some statistical test threshold).

For a complete survey of feature selection methods, see [3].

3 Regularized Linear Models

3.1 Regularized Linear Regression

Consider the linear regression model

$$Y = \beta_0 + x^T \beta$$

Suppose we have N data points, and p features. Each feature is standardized to have mean 0 and variance 1. Regularized Linear Regression solves the following problem:

$$\min_{(\beta_0, \beta) \in \mathcal{R}^{p+1}} \left[\frac{1}{2N} \sum_{i=1}^N (y_i - \beta_0 - x_i^T \beta)^2 + \lambda P(\beta) \right],$$

$$\text{SSE: } J(\vec{w}) = \sum_i (h_{\vec{w}}(\vec{x}_i) - y_i)^2$$

$$\text{Now, } J(\vec{w}) = \sum_i (h_{\vec{w}}(\vec{x}_i) - y_i)^2 + \lambda P(\vec{w})$$

where $\sum_{i=1}^N (y_i - \beta_0 - x_i^T \beta)^2$ is the square loss term, $P(\beta)$ is a penalty term, and λ controls the strength of the penalty.

There are three kinds of commonly used penalties in linear regression:

- $P(\beta) = \frac{1}{2} \|\beta\|_2^2 \leq \frac{1}{2} \sum_{j=1}^p \beta_j^2$ is called the RIDGE regression penalty.
- $P(\beta) = \|\beta\|_1 = \sum_{j=1}^p |\beta_j|$ is called the LASSO penalty.
- $P_\alpha(\beta) = (1 - \alpha) \frac{1}{2} \|\beta\|_2^2 + \alpha \|\beta\|_1 = \sum_{j=1}^p [\frac{1}{2}(1 - \alpha)\beta_j^2 + \alpha|\beta_j|]$ is called the elastic-net penalty. The elastic-net penalty is a compromise between the ridge-regression penalty and the lasso penalty.

$$P(\vec{w}) = \sum_d (w_d)^2$$

$$P(\vec{w}) = \sum_d |w_d|$$

$$P_\alpha(\vec{w}) = (1 - \alpha) \cdot \sum_d (w_d)^2 + \alpha \cdot \sum_d |w_d|$$

Ridge Regression shrinks the size the regression coefficients. In linear regression, if there are two correlated features, there coefficients can be poorly determined and have high variance. One of them can have a very large positive coefficient, and the other correlated feature can have a very large negative coefficient. They cancel each other. By adding the ridge penalty, the problem is alleviated, as it shrinks the coefficients towards 0. In the extreme case of k identical features, they each get small identical coefficients. So ridge penalty encourages features to borrow strength from each other. From a Bayesian point of view, the ridge regression estimation assumes that β_j has a Gaussian distribution with 0 mean as its prior distribution. And the solution to ridge regression is the mean (or mode) of the posterior distribution.

Lasso behaves differently than Ridge. If there are several correlated features, Lasso tends to pick one and ignore the rest. That is, some features will have coefficients exactly 0. So Lasso can be used to perform continuous feature selection. From a Bayesian point of view, the Lasso penalty corresponds to a Laplace prior.

To illustrate the behaviors of Ridge and Lasso, we write them as constrained optimization problems.

Ridge regression can be equivalently formulated as

$$\hat{\beta}^{ridge} = \arg \min_{\beta} \sum_{i=1}^N (y_i - \beta_0 - \sum_{j=1}^p x_{ij} \beta_j)^2$$

subject to $\sum_{j=1}^p \beta_j^2 \leq t$

$$\arg \min_{\vec{w}} \sum_i (h_w(\vec{x}_i) - y_i)^2$$

s.t. $\sum_d (w_d)^2 \leq t$

There is a one-to-one correspondence between λ and t .

Lasso regression can be equivalently formulated as

$$\hat{\beta}^{lasso} = \arg \min_{\beta} \sum_{i=1}^N (y_i - \beta_0 - \sum_{j=1}^p x_{ij} \beta_j)^2$$

subject to $\sum_{j=1}^p |\beta_j| \leq t$

$$\arg \min_{\vec{w}} \sum_i (h_w(\vec{x}_i) - y_i)^2$$

s.t. $\sum_d |w_d| \leq t$

Similarly, there is a one-to-one correspondence between λ and t .




FIGURE 3.11. Estimation picture for the lasso (left) and ridge regression (right). Shown are contours of the error and constraint functions. The solid blue areas are the constraint regions $|\beta_1| + |\beta_2| \leq t$ and $\beta_1^2 + \beta_2^2 \leq t^2$, respectively, while the red ellipses are the contours of the least squares error function.

Figure 1: Source: Figure 3.11 of [4]

Figure 1 shows the difference between lasso and ridge regression estimations when there are only two features. The square loss has elliptical contours. Ridge regression has a disk constraint region, while lasso has a diamond constraint region. In both constrained optimization problems, the optimal solution is the first point where the elliptical contours hit the constraint region. In Lasso regression, if the solution occurs at a corner, then it has one parameter β_j equal to zero. When p is large, there are many corners so that many parameters are likely to become zero.

3.2 Regularized Logistic Regression

The regularized Logistic Regression has the form

$$\max_{\beta_0, \beta} \left\{ \frac{1}{N} \sum_{i=1}^N [y_i \log(P(y=1|x_i)) + (1-y_i) \log(P(y=0|x_i))] - \lambda P_\alpha(\beta) \right\}$$

If we use RIDGE penalty, we get $J(\theta)$ as following:

$$\max_{\beta_0, \beta} \left\{ \frac{1}{N} \sum_{i=1}^N [y_i \log(P(y=1|x_i)) + (1-y_i) \log(P(y=0|x_i))] - \frac{\lambda}{2N} \sum_{j=1}^p \beta_j^2 \right\}$$

batch G.D.

$$\begin{aligned}
 \max \text{ log likelihood } & - \\
 J(\vec{w}) &= \sum_i (y_i \ln(h_{\vec{w}}(\tilde{x}_i)) + (1-y_i) \ln(1-h_{\vec{w}}(\tilde{x}_i))) - \lambda P(\vec{w}) \\
 &= \dots - \lambda \frac{1}{2} \sum_d (w_d)^2 \\
 \frac{\partial J(\vec{w})}{\partial w^d} &= \sum_i (y_i - h_{\vec{w}}(\tilde{x}_i)) \cdot x_i^d - \lambda w^d \\
 w^d &= w^d + \alpha \left[\sum_i (y_i - h_{\vec{w}}(\tilde{x}_i)) x_i^d - \lambda w^d \right] \\
 &= w^d - \alpha \left[\sum_i (h_{\vec{w}}(\tilde{x}_i) - y_i) x_i^d + \lambda w^d \right]
 \end{aligned}$$

maximize in pos. dir., so