

Learning Ridge Regression with R: A Step-by-Step Guide

Authored by
Mohammed loot

November 6, 2025

RECOMMENDED CITATION

Mohammed loot (2025). *Learning Ridge Regression with R: A Step-by-Step Guide*.
PSYCHOLOGICAL STATISTICS. Retrieved from
<https://statistics.arabpsychology.com/?p=11800>

[Ridge regression](#) is an indispensable **regularization** technique in statistical modeling, specifically designed to address stability issues when fitting linear models that suffer from [multicollinearity](#). Multicollinearity arises when predictor variables within the model are highly correlated with one another. This high correlation can lead to highly inflated variance in the standard coefficient estimates, making them unstable and difficult to interpret. By applying a penalty, Ridge regression provides a robust mechanism to stabilize these estimates and produce a more reliable predictive model.

To fully appreciate the mechanism of Ridge regression, it is essential to compare it against the traditional statistical benchmark: Ordinary [Least Squares regression](#) (OLS). OLS operates by finding the set of coefficient estimates (β) that minimize the Residual Sum of Squares (RSS). This minimization is the sole objective of OLS, mathematically represented as:

$$\text{RSS} = \sum (y_i - \hat{y}_i)^2$$

where the terms are defined as:

Σ : Represents the mathematical operation of summation.

y_i : The actual observed response value for the i th observation.

\hat{y}_i : The predicted response value derived from the multiple linear regression model.

In sharp contrast, Ridge regression modifies this objective function by introducing a penalty term, known as the L2 penalty. This compels the model not only to achieve a good fit (minimizing RSS) but simultaneously to keep the magnitude of the coefficient estimates small. This dual optimization process--fitting the data while constraining the coefficients--is the core principle that stabilizes estimates, especially when facing correlated predictors, thus mitigating the effects of multicollinearity.

The objective function that Ridge regression seeks to minimize integrates both the goodness-of-fit (RSS) and the penalty constraint:

$$\text{RSS} + \lambda \sum \beta_j^2$$

Here, the index j spans from 1 to p predictor variables, and λ (lambda) is a crucial non-negative tuning parameter ($\lambda \geq 0$). The term $\lambda \sum \beta_j^2$ is the L2 [shrinkage penalty](#). The fundamental challenge in implementing Ridge regression is the selection of the optimal λ , which is typically chosen through cross-validation to minimize the estimated test [Mean Squared Error](#) (MSE). This comprehensive tutorial provides a step-by-step guide to executing and interpreting a Ridge regression model using the **R** programming environment and the specialized **glmnet** package.

Step 1: Prepare and Structure Data for Analysis in R

Our practical implementation begins by preparing the dataset and ensuring it adheres to the specific formatting requirements of the **glmnet** package in R. For this demonstration, we will leverage the well-established built-in R dataset, **mtcars**. Our modeling goal is to predict the vehicle's horsepower (**hp**), which will serve as our response variable. We will utilize a subset of four specific variables as our predictors, selected to illustrate the modeling process effectively:

mpg (Miles per gallon)

wt (Weight)

drat (Rear axle ratio)

qsec (1/4 mile time)

The specialized **glmnet** package, designed for fitting penalized regression models, requires strict input data structures. Specifically, the dependent variable (response) must be defined as a simple **vector**, while the independent variables (predictors) must be grouped and structured as a **data.matrix**. Failing to adhere to these structural requirements will prevent the model from fitting correctly or efficiently.

The following R code snippet demonstrates the correct procedure for defining and structuring both our response variable (y) and our predictor matrix (x) extracted from the **mtcars** dataset, making them compatible with the **glmnet** framework:

```
#define response variable
```

```
y <- mtcars$hp
```

```
#define matrix of predictor variables
```

```
x <- data.matrix(mtcars)
```

Step 2: Initialize and Fit the Initial Ridge Regression Model

To initiate the Ridge regression modeling process using the **glmnet()** function, we must explicitly set the **alpha** parameter to 0. This parameter dictates the type of penalized regression being performed: setting **alpha** to 1 yields [Lasso Regression](#), while any intermediate value ($0 < \alpha < 1$) results in an Elastic Net model. For pure Ridge regression, **alpha = 0** is mandatory.

A critical theoretical requirement for applying Ridge regression is that all predictor variables must be **standardized**. Standardization ensures that every predictor contributes equally to the L2 penalty term. Without standardization, variables with larger inherent scales or variances would disproportionately dominate the shrinkage process, biasing the results. Standardization transforms the data such that each predictor variable has a mean of 0 and a standard deviation of 1.

Fortunately for the practitioner, the **glmnet()** function simplifies this process by automatically handling data standardization internally before fitting the model. If your data has already been pre-scaled outside of the package, you can instruct the function to bypass this automatic scaling by specifying the argument **standardize=False** within the function call. We now fit the initial model, allowing **glmnet** to automatically calculate coefficient paths across a comprehensive range of potential lambda values:

library(glmnet)

```
#fit ridge regression model (alpha=0 for Ridge)
```

```
model <- glmnet(x, y, alpha = 0)
```

```
#view summary of the resulting model object
```

```
summary(model)
```

```
Length Class Mode
```

```
a0 100 -none- numeric
```

```
beta 400 dgCMatrix S4
```

```
df 100 -none- numeric
```

```
dim 2 -none- numeric
```

```
lambda 100 -none- numeric
```

```
dev.ratio 100 -none- numeric
```

```
nulldev 1 -none- numeric
```

```
npasses 1 -none- numeric
```

```
jerr 1 -none- numeric
```

```
offset 1 -none- logical
```

```
call 4 -none- call
```

```
nobs 1 -none- numeric
```

Step 3: Optimize Lambda (λ) Using K-Fold Cross-Validation

The ultimate predictive power and stability of the fitted Ridge regression model are entirely dependent on selecting the correct value for the hyperparameter, λ . Our methodological objective is to pinpoint the specific lambda value that minimizes the model's generalization error, which is the error observed on unseen data, typically quantified by the test [Mean Squared Error](#) (MSE).

To achieve a robust estimate of this test error and avoid overfitting to the training data, we employ [k-fold cross-validation](#) (KCV). KCV is a rigorous resampling technique that divides the training data into k equal subsets (folds). The model is trained on $k-1$ folds and validated on the remaining fold, repeating the process k times. This ensures that the estimated error for each lambda value is

reliable and representative of the model's true performance.

The [glmnet](#) package conveniently provides the dedicated `cv.glmnet()` function to automate this entire optimization procedure. By default, this function conducts 10-fold cross-validation ($k=10$) across the full spectrum of potential lambda values calculated in the previous step, efficiently identifying the optimal shrinkage magnitude.

#perform k-fold cross-validation to find optimal lambda value

```
cv_model <- cv.glmnet(x, y, alpha = 0)
```

```
#extract the optimal lambda value that minimizes test MSE
```

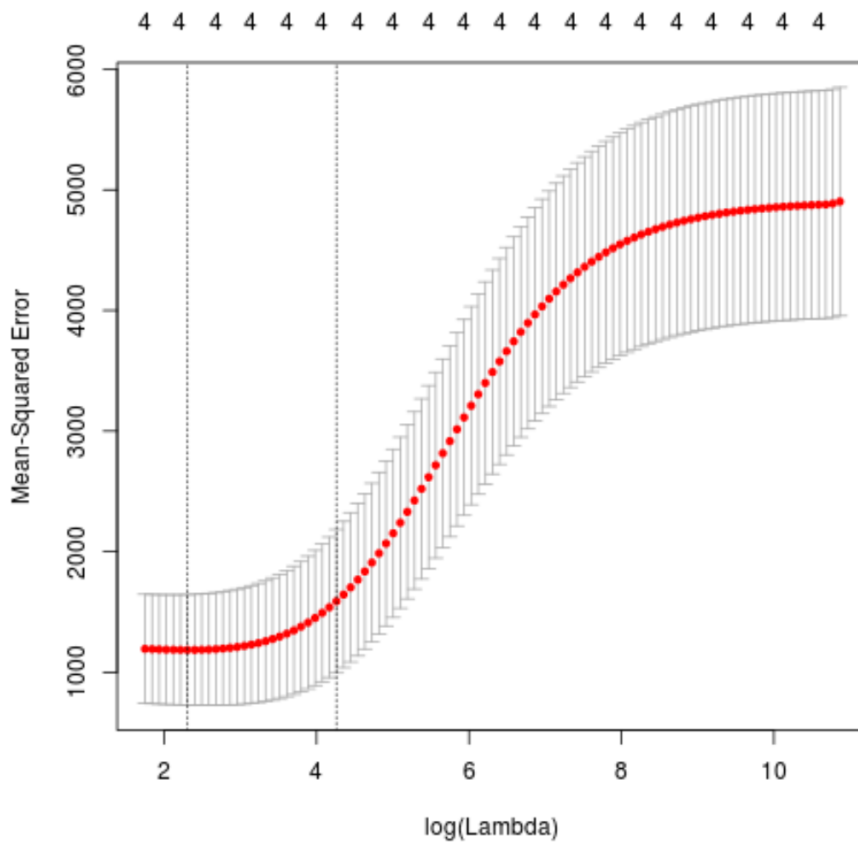
```
best_lambda <- cv_model$lambda.min
```

```
best_lambda
```

```
10.04567
```

```
#produce plot of test MSE by lambda value
```

```
plot(cv_model)
```



As evidenced by the output and the visualization, the cross-validation process successfully

identified the ideal shrinkage parameter. The optimal lambda value, which corresponds to the point of minimum test MSE, was determined to be **10.04567**.

Step 4: Analyze the Final Shrinkage Coefficients and Trace Plot

With the optimal lambda value precisely determined, the next step is to refit the Ridge regression model using this specific hyperparameter. This refitting procedure yields the final, stable coefficient estimates that will be used for interpretation and prediction. Analyzing these coefficients allows us to fully understand the effects of the L2 shrinkage applied to our predictor variables.

The following R code applies the determined **best_lambda** to the model to generate the final set of standardized coefficient estimates for all our predictors:

```
#find coefficients of best model
```

```
best_model <- glmnet(x, y, alpha = 0, lambda = best_lambda)
```

```
coef(best_model)
```

```
5 x 1 sparse Matrix of class "dgCMatrix"
```

```
s0
```

```
(Intercept) 475.242646
```

```
mpg -3.299732
```

```
wt 19.431238
```

```
drat -1.222429
```

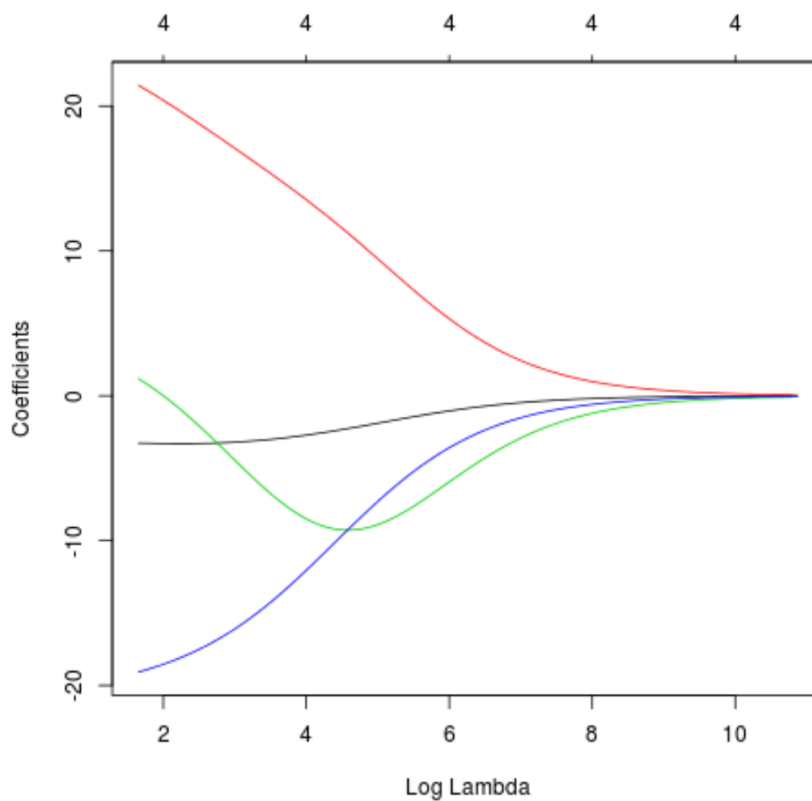
```
qsec -17.949721
```

These derived coefficients represent the weights assigned to the standardized predictor variables in the final Ridge regression model. A key characteristic of **Ridge regression** is visible here: while the coefficients have been significantly shrunk (moved closer to zero) due to the L2 penalty, none of them have been set **exactly to zero**. This retention of all predictors is a defining difference between Ridge regression and Lasso regression.

To visualize the continuous effect of the lambda tuning parameter across its entire range, we generate a Ridge trace plot. This plot is essential for illustrating how the magnitude of the coefficient estimates dynamically changes (shrinks) as the penalization factor (lambda) increases:

```
#produce Ridge trace plot
```

```
plot(model, xvar = "lambda")
```



The Ridge trace plot confirms the expected regularization behavior. When lambda is very small (far left on the log-scaled x-axis), the coefficients closely resemble the traditional OLS estimates. As the value of lambda increases (moving right), the penalty term gains influence, causing the coefficients to be progressively penalized and smoothly compressed toward zero, thereby stabilizing the model estimates.

Step 5: Evaluate Model Performance Using R-squared

The final stage of the modeling process involves quantifying the overall goodness-of-fit achieved by the optimal Ridge regression model. We assess performance by calculating the [R-squared](#) (Coefficient of Determination) value on our training data. R-squared provides a clear, interpretable metric: it measures the proportion of the total variance in the response variable (horsepower) that is successfully explained or predicted by the set of independent predictor variables in the model.

To calculate R-squared, we must first use the fitted model and the optimal lambda value to generate predicted response values (**y_predicted**). Subsequently, we calculate the essential components required for the R-squared formula: the Total Sum of Squares (SST), which measures the variance around the mean, and the Sum of Squared Errors (SSE), which measures the unexplained residual variance.

```
#use fitted best model to make predictions
```

```
y_predicted <- predict(model, s = best_lambda, newx = x)
```

```
#find SST (Total Sum of Squares) and SSE (Sum of Squared Errors)
```

```
sst <- sum((y - mean(y))^2)
```

```
sse <- sum((y_predicted - y)^2)
```

```
#calculate R-Squared
```

```
rsq <- 1 - sse/sst
```

```
rsq
```

```
0.7999513
```

The calculation yields an R-squared value of approximately **0.7999513**. This result indicates that the final, optimized Ridge regression model is capable of explaining roughly **79.99%** of the observed variation in vehicle horsepower within the training dataset. This high level of explained variance suggests that the model provides a strong and stable fit, successfully managing the inherent complexities and potential multicollinearity present in the predictor data through the application of the L2 penalty.

For those interested in reviewing the complete, executable R script utilized throughout this practical implementation, the code is readily available for review [here](#).