

Ridge Regression

Cross-validated L2-regularized regression

Dataset: D:\DATA ANALYSIS\H Regression Tests and Models\Ridge Regression\dataset.csv

Complete cases: 649

Outcome: G3

Predictors: G1, G2, studytime, failures, absences, age, Medu, Fedu

Training/testing split: 519/130

Cross-validation folds: 10

Selected lambda: 0.0001

Training R-squared: 0.846445

Testing R-squared: 0.861117

Testing RMSE: 1.163763

Full-data ridge R-squared: 0.850777

Full-data OLS R-squared: 0.850777

Interpretation: ridge regression applies an L2 penalty, shrinking correlated predictor coefficients while retaining all predictors.

Train/test metrics

split	n	r_squared	rmse	mae
Training	519	0.846445	1.27134	0.797773
Testing	130	0.861117	1.16376	0.739991
Full data ridge	649	0.850777	1.24702	0.780019
Full data OLS comparison	649	0.850777	1.24702	0.780019

Ridge coefficients


term	ridge_coefficient_original_scale	ridge_standardized_coefficient	ols_coefficient_original_scale	absolute_ridge_standardized_coefficient	ridge_minus_ols
Intercept	-0.501155		-0.501155		-3.23414e-07
G1	0.143397	0.121758	0.143397	0.121758	4.05177e-07
G2	0.884807	0.797368	0.884807	0.797368	-4.87638e-07
studytime	0.0966318	0.0247923	0.0966318	0.0247923	1.59351e-08
failures	-0.235361	-0.0431854	-0.235361	0.0431854	-1.78222e-07
absences	0.022762	0.0326719	0.022762	0.0326719	-6.09296e-09
age	0.0226856	0.00854713	0.0226855	0.00854713	7.04029e-08
Medu	-0.0449511	-0.0157739	-0.0449512	0.0157739	5.9576e-08
Fedu	0.0220254	0.00749312	0.0220253	0.00749312	2.13359e-08

Selected lambda

selection_rule	lambda_alpha	mean_cv_mse	sd_cv_mse	mean_train_mse	selected_lambda
Minimum cross-validated mean square error	0.0001	1.70275	1.00814	1.61184	1


Outcome distribution

Ridge regression outcome: G3




Cross-validation curve

Selected lambda = 0.0001




Ridge coefficient paths

Increasing lambda shrinks coefficients toward zero




Selected ridge coefficients

Standardized effect magnitudes at the selected lambda




Ridge versus OLS coefficients

Ridge reduces coefficient magnitude when predictors overlap




Observed versus predicted

Held-out testing sample



Residuals versus predicted values

Held-out testing sample



Test residual distribution

Residuals should be centered near zero

