

# Elastic Net Regression

Test name: Elastic Net Regression

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

Rows in original dataset: 649

Columns in original dataset: 33

Rows used after cleaning: 649

Outcome variable: G3

Numeric predictors: G1, G2, studytime, failures, absences, age, Medu, Fedu

Categorical predictors: school, sex, address, famsize, Pstatus

Encoded predictor count: 13

Train rows: 486

Test rows: 163

Cross-validation folds: 5

Selected l1\_ratio / alpha mix: 1.0000

Selected alpha/lambda penalty: 0.084017

Retained non-zero coefficients: 5 of 13

Test RMSE: 1.2675

Test MAE: 0.7774

Test R-squared: 0.8491

Method note: Elastic Net combines L1 lasso-style selection and L2 ridge-style shrinkage.

Interpretation note: coefficients are standardized because predictors are scaled before fitting.

Decision guide: retained coefficients identify variables selected by the penalty; larger absolute coefficients have stronger predictive contribution.

Reporting guide: state outcome, predictors, selected l1\_ratio, selected lambda/alpha penalty, cross-validation folds, retained predictors, RMSE, MAE and R-squared.

Charts generated: 7

# Elastic Net model settings

| test_name   | outcome | numeric_pre | categorical | rows_origin | columns_or | rows_used | train_rows | test_rows | cv_folds | selected_l1 | selected_al | nonzero_co | encoded_pre |
|-------------|---------|-------------|-------------|-------------|------------|-----------|------------|-----------|----------|-------------|-------------|------------|-------------|
| Elastic Net | G3      | G1, G2, stu | school, sex | 649.0000    | 33.0000    | 649.0000  | 486.0000   | 163.0000  | 5.0000   | 1.0000      | 0.0840      | 5.0000     | 13.0000     |

# Elastic Net coefficients


| term        | standardized_coefficient | abs_standardized_coefficient | coefficient_sign | selection_status |
|-------------|--------------------------|------------------------------|------------------|------------------|
| (Intercept) | 11.7798                  | 11.7798                      | Intercept        | Model intercept  |
| G2          | 2.5841                   | 2.5841                       | Positive         | Retained         |
| G1          | 0.2783                   | 0.2783                       | Positive         | Retained         |
| failures    | -0.0546                  | 0.0546                       | Negative         | Retained         |
| school_MS   | -0.0487                  | 0.0487                       | Negative         | Retained         |
| sex_M       | -0.0212                  | 0.0212                       | Negative         | Retained         |
| absences    | 0.0000                   | 0.0000                       | Zero             | Shrunk to zero   |
| studytime   | 0.0000                   | 0.0000                       | Zero             | Shrunk to zero   |
| Medu        | -0.0000                  | 0.0000                       | Zero             | Shrunk to zero   |
| age         | 0.0000                   | 0.0000                       | Zero             | Shrunk to zero   |
| Fedu        | 0.0000                   | 0.0000                       | Zero             | Shrunk to zero   |
| address_U   | 0.0000                   | 0.0000                       | Zero             | Shrunk to zero   |
| famsize_LE3 | 0.0000                   | 0.0000                       | Zero             | Shrunk to zero   |
| Pstatus_T   | -0.0000                  | 0.0000                       | Zero             | Shrunk to zero   |

# Model performance metrics


| model                  | split | n        | n_predictors_after_end | rmse   | mae    | r_squared |
|------------------------|-------|----------|------------------------|--------|--------|-----------|
| Elastic Net Regression | Train | 486.0000 | 13.0000                | 1.2518 | 0.7775 | 0.8476    |
| Elastic Net Regression | Test  | 163.0000 | 13.0000                | 1.2675 | 0.7774 | 0.8491    |
| Baseline Linear Regres | Train | 486.0000 | 13.0000                | 1.2370 | 0.7846 | 0.8512    |
| Baseline Linear Regres | Test  | 163.0000 | 13.0000                | 1.2509 | 0.7593 | 0.8530    |

# Elastic Net outcome distribution

The model predicts the selected continuous outcome using regularized regression.




**Observed vs predicted test values**  
Points near the diagonal show better out-of-sample prediction accuracy.




# Elastic Net residuals vs predicted values

Random scatter around zero supports better regression fit; patterns suggest remaining structure.




# Elastic Net retained coefficient sizes

Positive and negative coefficients show the direction of retained predictor effects after shrinkage.




# Cross-validation curve for selected l1\_ratio

The selected penalty minimizes average validation error across folds.



# Elastic Net train/test prediction error

Similar train and test errors suggest better generalization.



# Elastic Net regularization path

Coefficient paths show how predictors enter or shrink as penalty strength changes.

