

Package: RidgeLRGLH (via r-universe)

October 28, 2026

Type Package

Title Ridge-Regularized Largest Root Tests for General Linear Hypotheses

Version 0.1.0

Description Tests general linear hypotheses in high-dimensional multivariate regression using a ridge-regularized largest root statistic. Estimates the population spectral distribution by linear programming and computes Tracy-Widom centering and scaling constants. Supports fixed, Bayesian, and minimax regularization choices. Implements the methods of Li (2025) [doi:10.48550/arXiv.2504.15510](https://doi.org/10.48550/arXiv.2504.15510).

License MIT + file LICENSE

Encoding UTF-8

Depends R (>= 4.1.0)

Imports lpSolve, RMTstat, stats

NeedsCompilation no

Author Haoran Li [aut, cre]

Maintainer Haoran Li <hz10152@auburn.edu>

Repository <https://cran.r-universe.dev>

Date/Publication 2026-10-28 10:30:27 UTC

RemoteUrl <https://github.com/cran/RidgeLRGLH>

RemoteRef HEAD

RemoteSha 7fb41bccdce548353e144b8cd37b6e46cb5f4a7b

Contents

estimate_spectrum	2
ridge_calibration	3
ridge_glht	4
ridge_root	5

Index	7
--------------	----------

estimate_spectrum	<i>Estimate a Population Covariance Spectrum by Linear Programming</i>
-------------------	--

Description

Fits nonnegative spectral masses using the complex-valued relative discrepancy constraints in Algorithm 1 of Li (2025).

Usage

```
estimate_spectrum(eigenvalues, gamma2, lambda = NULL,
  size = 500L, grid_size = 500L, sigma_grid = NULL,
  z_grid = NULL, order = 2L, truncation = 2)
```

Arguments

eigenvalues	All eigenvalues of the normalized residual covariance matrix, including zeros, in any order.
gamma2	Dimension divided by residual degrees of freedom.
lambda	Positive ridge used for fitting. Defaults to half the mean eigenvalue. This fitting parameter need not equal the eventual test ridge.
size	Number of complex transform grid points, at least 3.
grid_size	Number of population spectrum grid points, at least 3.
sigma_grid	Optional nonnegative population spectrum grid, in original covariance units.
z_grid	Optional complex points in the open upper half-plane, in original covariance units. Overrides the default transform grid.
order	1 for first-order constraints; 2 (default) also uses derivative constraints.
truncation	Positive exponent: discard weights at most $10^{-truncation}/K$, where K is the spectrum grid size.

Details

The default grid is obtained by inverting the empirical companion Stieltjes transform. Its real parts run from the transform at 1.05 times the largest sample eigenvalue to the transform at minus the fitting ridge. The imaginary transform component is 0.01 divided by the largest eigenvalue. The default population grid follows Supplement S.2. A bounded Newton iteration with step halving performs complex inversion; failed inversion or optimization raises an error. The returned loss is measured before truncation. Finite grids are approximations; inspect sensitivity to grid size, fitting ridge and constraint order.

Value

A list containing positive weights, their sigma locations, pre-truncation loss, solver status (zero on success), fitting lambda, z_grid, and constraint order.

References

Li, H. (2025). Ridge-Regularized Largest Root Test For High-Dimensional General Linear Hypotheses. [doi:10.48550/arXiv.2504.15510](https://doi.org/10.48550/arXiv.2504.15510). Accepted by the Annals of Statistics.

Examples

```
set.seed(12)
Z <- matrix(rnorm(20 * 60), 20)
ev <- eigen(tcrossprod(Z) / 60, symmetric = TRUE, only.values = TRUE)$values
spectrum <- estimate_spectrum(ev, 20 / 60, size = 40, grid_size = 60)
sum(spectrum$weights)
```

ridge_calibration	<i>Compute Tracy-Widom Centering and Scaling Constants</i>
-------------------	--

Description

Computes the spectral edge and calibration constants for a specified population covariance spectrum.

Usage

```
ridge_calibration(sigma, gamma1, gamma2, lambda,
  weights = rep(1 / length(sigma), length(sigma)))
```

Arguments

sigma	Nonnegative population covariance eigenvalues or spectral mass locations, not all zero.
gamma1	Dimension divided by numerator degrees of freedom.
gamma2	Dimension divided by denominator degrees of freedom.
lambda	Positive regularization parameter, in covariance units.
weights	Nonnegative spectral masses, normalized internally to sum to one.

Details

Uses the scalar parameterization $x(h) = h + [1 + \gamma_2 H_1(h)]^{-1}$ from Section 3. It solves for the regular branch and $\beta^2 s'(\beta) = 1/\gamma_1$ by bracketed root finding. Derivatives follow by implicit differentiation. This is equivalent to the manuscript's ODE characterization, and avoids evaluating a dense ODE trajectory near the edge. Root tolerances are applied after normalization of the population spectrum. The Tracy-Widom approximation assumes the high-dimensional regime and spectral regularity conditions in the article; it is not an exact finite-sample calibration.

Value

A named numeric vector: center, scale, beta, rho, s, s_prime, and s_double_prime, with derivatives evaluated at beta. The normalized statistic is $p^{2/3}(l_{max} - center)/scale$.

References

Li, H. (2025). Ridge-Regularized Largest Root Test For High-Dimensional General Linear Hypotheses. doi:10.48550/arXiv.2504.15510. Accepted by the Annals of Statistics.

Examples

```
ridge_calibration(sigma = 1, gamma1 = 2, gamma2 = 0.5, lambda = 1)
```

ridge_glht

Test a General Linear Hypothesis in Multivariate Regression

Description

Tests the hypothesis $BC = 0$ in the model $Y = BX + E$ using a ridge-regularized largest root statistic.

Usage

```
ridge_glht(Y, X, C, lambda = NULL,
  selection = c("fixed", "bayes_identity", "bayes_covariance", "minimax"),
  lambda_grid = NULL, spectrum = NULL, control = list())
```

Arguments

Y	Real p by n response matrix: variables in rows and observations in columns.
X	Real m by n fixed design matrix with full row rank, $m < n$. Include an intercept explicitly if needed.
C	Real m by q full column rank contrast matrix specifying $BC = 0$.
lambda	Positive fixed ridge, defaulting to the mean eigenvalue of W_2 . Supply only for fixed selection.
selection	Fixed ridge, Bayes identity prior, Bayes covariance prior, or minimax over non-negative mixtures of these priors with equal total signal strength.
lambda_grid	Positive candidate ridges for adaptive selection. Defaults to ten equally spaced values from $\min(0.1p/n_2, 1) \operatorname{tr}(W_2)/p$ to $5 \operatorname{tr}(W_2)/p$. Ties select the smallest candidate.
spectrum	Optional list with population mass locations sigma and weights. If omitted, estimates the spectrum from W_2 . Known spectra may be supplied for oracle experiments.
control	Named list of options passed to estimate_spectrum, excluding eigenvalues and gamma2. The lambda option here controls spectral fitting only.

Details

QR factorizations construct the tested and residual subspaces, without forming observation-by-observation projection matrices. Numerator and denominator degrees of freedom are q and n - m. Nonzero right-hand sides and rank-deficient designs or contrasts are not supported. This is an asymptotic test under the moment, covariance, dimensional and regular-edge conditions in Li (2025). Very small tested ranks or dimensions may give poor finite-sample approximation.

Value

An object of class `htest`, with statistic, selected ridge in parameter, asymptotic upper-tail `p.value`, and standard test metadata. Additional fields include `largest.root`, `center`, `scale`, `n1`, `n2`, `dimension`, `selection`, `tuning` (candidate ridges, constants and scores), `spectrum`, and selected calibration. Fixed selection has an undefined (NA) tuning score.

References

Li, H. (2025). Ridge-Regularized Largest Root Test For High-Dimensional General Linear Hypotheses. [doi:10.48550/arXiv.2504.15510](https://doi.org/10.48550/arXiv.2504.15510). Accepted by the Annals of Statistics.

See Also

[ridge_root](#)

Examples

```
set.seed(8)
Y <- matrix(rnorm(20 * 80), 20)
X <- rbind(1, matrix(rnorm(10 * 80), 10))
C <- diag(11)[, -1, drop = FALSE]
ridge_glht(Y, X, C, selection = "minimax",
           control = list(size = 40, grid_size = 60))
```

ridge_root	<i>Ridge-Regularized Largest Root Test for Two Covariance Matrices</i>
------------	--

Description

Tests using the largest eigenvalue of W_1 times the inverse of W_2 plus a positive ridge, with asymptotic real Tracy-Widom calibration.

Usage

```
ridge_root(W1, W2, n1, n2, lambda = NULL,
           selection = c("fixed", "bayes_identity", "bayes_covariance", "minimax"),
           lambda_grid = NULL, spectrum = NULL, control = list())
```

Arguments

W1, W2	Real symmetric positive semidefinite matrices of equal dimensions, normalized by <code>n1</code> and <code>n2</code> respectively. <code>W2</code> must have positive trace. Under the null they must satisfy the common-covariance double Wishart structure and assumptions of the article.
n1, n2	Positive numerator and denominator degrees of freedom.
lambda	Positive fixed ridge, defaulting to the mean eigenvalue of <code>W2</code> . Supply only for fixed selection.

selection	Fixed ridge, Bayes identity prior, Bayes covariance prior, or minimax over non-negative mixtures of these priors with equal total signal strength.
lambda_grid	Positive candidate ridges for adaptive selection. Defaults to ten equally spaced values from $\min(0.1p/n_2, 1) \operatorname{tr}(W_2)/p$ to $5 \operatorname{tr}(W_2)/p$. Ties select the smallest candidate.
spectrum	Optional list with population mass locations sigma and weights. If omitted, estimates the spectrum from W2. Known spectra may be supplied for oracle experiments.
control	Named list of options passed to estimate_spectrum, excluding eigenvalues and gamma2. The lambda option here controls spectral fitting only.

Details

Bayes criteria are the identity or covariance signal estimates divided by the scaling constant. The minimax criterion is the smaller of the identity signal estimate and the trace-normalized covariance signal estimate, divided by that same scaling constant. This solves the minimization over non-negative mixtures exactly at an endpoint. All adaptive choices use a finite candidate set. Their asymptotic validity requires the separation and consistency conditions in Section 4 of the article; choosing a ridge from the observed data does not provide a finite-sample error guarantee. The returned p-value is for the selected test and does not adjust for reporting multiple tests or searching multiple hypotheses. Inputs must be real; complex-valued Wishart models are not supported.

Value

An object of class `htest`, with statistic, selected ridge in parameter, asymptotic upper-tail p.value, and standard test metadata. Additional fields include `largest.root`, `center`, `scale`, `n1`, `n2`, `dimension`, `selection`, `tuning` (candidate ridges, constants and scores), `spectrum`, and `selected calibration`. Fixed selection has an undefined (NA) tuning score.

References

Li, H. (2025). Ridge-Regularized Largest Root Test For High-Dimensional General Linear Hypotheses. [doi:10.48550/arXiv.2504.15510](https://doi.org/10.48550/arXiv.2504.15510). Accepted by the Annals of Statistics.

See Also

[ridge_glht](#), [estimate_spectrum](#)

Examples

```
set.seed(17)
Z1 <- matrix(rnorm(20 * 30), 20)
Z2 <- matrix(rnorm(20 * 60), 20)
ridge_root(tcrossprod(Z1) / 30, tcrossprod(Z2) / 60,
           n1 = 30, n2 = 60, control = list(size = 40, grid_size = 60))
```

Index

`estimate_spectrum`, [2](#), [6](#)

`ridge_calibration`, [3](#)

`ridge_glht`, [4](#), [6](#)

`ridge_root`, [5](#), [5](#)