

Package: APML0 (via r-universe)

June 20, 2026

Type Package

Title Augmented and Penalized Minimization Method L0

Version 0.11

Date 2026-06-13

Description Fit linear, logistic and Cox models regularized with L0, lasso (L1), elastic-net (L1 and L2), or net (L1 and Laplacian) penalty, and their adaptive forms, such as adaptive lasso / elastic-net and net adjusting for signs of linked coefficients. It solves the L0 penalty problem by simultaneously selecting regularization parameters and performing hard-thresholding or selecting the number of non-zeros. This augmented and penalized minimization method provides an approximation solution to the L0 penalty problem, but runs as fast as L1 regularization. The package uses a one-step coordinate descent algorithm and runs extremely fast by taking into account the sparsity structure of coefficients. It can handle very high dimensional data and has superior selection performance.

License GPL (>= 2)

Encoding UTF-8

URL <https://github.com/LeeSprite/APML0>

BugReports <https://github.com/LeeSprite/APML0/issues>

Imports Rcpp (>= 0.12.12)

LinkingTo Rcpp, RcppEigen

Depends Matrix (>= 1.2-10)

Config/roxygen2/version 8.0.0

Suggests spelling

Language en-US

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

Date/Publication 2026-06-13 15:27:30 UTC

RemoteUrl <https://github.com/leesprite/apml0>

RemoteRef HEAD

RemoteSha e1495a087a6a61d4f141b2581ceb03e7f456a122

Contents

APML0-package	2
APML0	3
print.APML0	7

Index **9**

APML0-package	<i>Augmented and Penalized Minimization Method L0</i>
---------------	---

Description

Fit linear, logistic and Cox models regularized with L0, lasso (L1), elastic-net (L1 and L2), or net (L1 and Laplacian) penalty, and their adaptive forms. The package solves the L0 penalty problem by simultaneously selecting regularization parameters and performing hard-thresholding or selecting the number of non-zeros. This augmented and penalized minimization method provides an approximation solution to the L0 penalty problem but runs as fast as L1 regularization. A one-step coordinate descent algorithm exploits sparsity and handles very high dimensional data efficiently.

Details

Package: APML0
 Type: Package
 Version: 0.11
 Date: 2026-06-13
 License: GPL (>= 2)

Main function: [APML0](#)
 Print method: [print.APML0](#)

Author(s)

Xiang Li, Shanghong Xie, Donglin Zeng and Yuanjia Wang
 Maintainer: Xiang Li <spiritcoke@gmail.com>

References

- Li, X., Xie, S., Zeng, D., Wang, Y. (2018). *Efficient l0-norm feature selection based on augmented and penalized minimization*. *Statistics in Medicine*, 37(3), 473–486. doi:10.1002/sim.7526
- Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J. (2011). *Distributed optimization and statistical learning via the alternating direction method of multipliers*. *Foundations and Trends in Machine Learning*, 3(1), 1–122.
- Friedman, J., Hastie, T., Tibshirani, R. (2010). *Regularization paths for generalized linear models via coordinate descent*. *Journal of Statistical Software*, 33(1). doi:10.18637/jss.v033.i01

Examples

```
set.seed(1213)
N=100; p=30; p1=5
x=matrix(rnorm(N*p), N, p)
beta=rnorm(p1)
xb=x[, 1:p1] %*% beta
y=rnorm(N, xb)
fiti=APML0(x, y, penalty="Lasso", nlambda=10)
fiti2=APML0(x, y, penalty="Lasso", nlambda=10, nolds=10)
```

 APML0

Fit a Model with Various Regularization Forms

Description

Fit linear, logistic and Cox models regularized with L0, lasso (L1), elastic-net (L1 and L2), or net (L1 and Laplacian) penalty, and their adaptive forms, such as adaptive lasso / elastic-net and net adjusting for signs of linked coefficients.

It solves the L0 penalty problem by simultaneously selecting regularization parameters and performing hard-thresholding (or selecting the number of non-zeros). This augmented and penalized minimization method provides an approximation solution to the L0 penalty problem and runs as fast as L1 regularization.

The function uses a one-step coordinate descent algorithm and runs extremely fast by taking into account the sparsity structure of coefficients. It can handle very high dimensional data.

Usage

```
APML0(x, y, weights=NULL, family=c("gaussian", "binomial", "cox"),
penalty=c("Lasso", "Enet", "Net"),
Omega=NULL, alpha=1.0, lambda=NULL, nlambda=50, rlambda=NULL,
wbeta=rep(1, ncol(x)), sgn=rep(1, ncol(x)), KK=log(nrow(x)),
nolds=1, foldid=NULL, cvll=TRUE, icutB=TRUE, ncutB=10,
ifast=TRUE, isd=FALSE, iysd=FALSE, ifastr=TRUE, keep.beta=FALSE,
thresh=1e-6, maxit=1e+5, threshC=1e-5, maxitC=1e+2, threshP=1e-5)
```

Arguments

x	input matrix. Each row is an observation vector.
y	response variable. For family = "gaussian", y is a continuous vector. For family = "binomial", y is a binary vector with 0 and 1. For family = "cox", y is a two-column matrix with columns named time and status. status is binary, with 1 indicating event and 0 indicating right censoring.
weights	case weights for each observation. Default is 1 for every observation. For family = "cox", observations with weight zero are excluded from the risk set.
family	type of outcome. One of "gaussian", "binomial", or "cox".
penalty	penalty type. One of "Lasso", "Enet" (elastic net), or "Net". For "Net", Omega must be supplied; otherwise "Enet" is used. For penalty = "Net", the penalty is $\lambda\{\alpha\ \beta\ _1 + (1 - \alpha)/2\beta^T L\beta\},$ where L is the Laplacian matrix computed from Omega.
Omega	adjacency matrix with zero diagonal and non-negative off-diagonal, used with penalty = "Net".
alpha	mixing parameter between L1 and Laplacian for "Net", or between L1 and L2 for "Enet". Default alpha = 1.0 (lasso).
lambda	user-supplied decreasing sequence of regularization values. If NULL, a sequence is generated from nlambda and rlambda.
nlambda	number of lambda values. Default is 50.
rlambda	fraction of lambda.max used as the smallest lambda. Default is 0.0001 when $n \geq p$ and 0.01 otherwise.
wbeta	penalty weights for the L1 term (adaptive L1): $\sum_j w_j \beta_j $. Zero entries remove the penalty on the corresponding coefficient. The same weights are applied to the L0 step. Default is 1 for all coefficients.
sgn	sign adjustment for the Laplacian penalty (adaptive Laplacian). A vector of 1 or -1. Default is 1 for all coefficients.
KK	multiplier in the BIC-type criterion used when no cross-validation is performed (nfolds = 1 and foldid = NULL). The criterion is $nzero \times KK - 2 \times \log\text{Lik}$. Default KK = log(nrow(x)) gives BIC; KK = 2 gives AIC.
nfolds	number of cross-validation folds. With nfolds = 1 and foldid = NULL (default), BIC-type selection is used. Minimum value for cross-validation is 3. Specifying foldid overrides nfolds.
foldid	optional vector of fold assignments (integers 1 to nfolds) for each observation.
cvll	logical flag for using the log-likelihood as the cross-validation criterion. Default TRUE. For family = "gaussian", set FALSE to use mean squared prediction error. Not used for family = "cox" (partial likelihood is always used).
icutB	logical flag for L0 selection via hard-thresholding (TRUE, default) rather than by selecting the number of non-zeros (FALSE). Applies to all three families.
ncutB	number of threshold grid points for icutB = TRUE. Default is 10. Larger values may improve selection but increase runtime.

<code>isd</code>	logical flag for returning standardized coefficients. x is always standardized internally; <code>isd = FALSE</code> (default) returns β on the original scale.
<code>iysd</code>	logical flag for standardizing y before fitting, for <code>family = "gaussian"</code> only. Returned coefficients are always on the original y scale. Default is <code>FALSE</code> .
<code>keep.beta</code>	logical flag for returning estimates at all <code>lambda</code> values. When <code>FALSE</code> (default and cross-validation only), only the estimate at the optimal <code>lambda</code> is returned.
<code>ifast</code>	logical flag for the fast computation path. When <code>TRUE</code> (default), BIC/CV selection reuses the regularized fit without re-fitting. Set <code>FALSE</code> to re-fit via maximum likelihood on the selected support.
<code>ifastr</code>	logical flag for <code>family = "cox"</code> only. When <code>TRUE</code> (default), an efficient risk-set update is used; the algorithm may stop before all <code>nlambda</code> values are evaluated. Affects only efficiency, not the estimates.
<code>thresh</code>	convergence threshold for coordinate descent. Default <code>1e-6</code> .
<code>maxit</code>	maximum coordinate descent iterations. Default <code>1e5</code> .
<code>threshC</code>	convergence threshold for the hard-thresholding step. Default <code>1e-5</code> .
<code>maxitC</code>	maximum iterations for the hard-thresholding step. Default <code>100</code> .
<code>threshP</code>	probability truncation bound for <code>family = "binomial"</code> . Default <code>1e-5</code> .

Details

A one-step coordinate descent algorithm is applied for each `lambda`. For `family = "cox"`, `ifastr = TRUE` uses an efficient risk-set update and may stop before all `nlambda` values are evaluated; use `ifastr = FALSE` to force evaluation of all values.

x is always standardized internally and estimates are returned on the original scale. For `family = "gaussian"`, y is centered (no intercept is returned).

L0 variable selection is always performed in addition to the regularized fit. Without cross-validation, the support is selected by a BIC-type criterion $n_{zero} \times KK - 2 \times \log\text{Lik}$ (see `KK`). With cross-validation, selection uses hard-thresholding (`icutB = TRUE`) or the number of non-zeros (`icutB = FALSE`).

Value

An object of S3 class `"APML0"` containing:

<code>Beta0</code>	coefficients after L0 selection. For <code>family = "binomial"</code> , the first element is the intercept.
<code>Beta</code>	sparse matrix of regularized coefficients (class <code>"dgCMatrix"</code>) along the <code>lambda</code> path. For <code>family = "binomial"</code> , the first row is the intercept.
<code>fit</code>	data frame with columns <code>lambda</code> and <code>nzero</code> . Without cross-validation: also criterion (BIC-type score). With cross-validation: also <code>cvm</code> , <code>cvse</code> , and <code>index</code> . For <code>family = "gaussian"</code> : also <code>rsq</code> .
<code>fit0</code>	data frame for the L0-selected solution: <code>lambda</code> , <code>cvm</code> (or criterion), <code>nzero</code> .
<code>lambda.min</code>	<code>lambda</code> giving minimum <code>cvm</code> (cross-validation only).
<code>lambda.opt</code>	<code>lambda</code> for the L0 solution (cross-validation only).

penalty	penalty type used.
adaptive	logical flag(s) for adaptive weighting.
flag	convergence flag. 0 means converged.

Warning

The function may return NULL if no valid lambda values are found.

Author(s)

Xiang Li, Shanghong Xie, Donglin Zeng and Yuanjia Wang
 Maintainer: Xiang Li <spiritcoke@gmail.com>

References

- Li, X., Xie, S., Zeng, D., Wang, Y. (2018). *Efficient l0-norm feature selection based on augmented and penalized minimization*. *Statistics in Medicine*, 37(3), 473–486. doi:10.1002/sim.7526
- Boyd, S., Parikh, N., Chu, E., Peleato, B., Eckstein, J. (2011). *Distributed optimization and statistical learning via the alternating direction method of multipliers*. *Foundations and Trends in Machine Learning*, 3(1), 1–122.
- Friedman, J., Hastie, T., Tibshirani, R. (2010). *Regularization paths for generalized linear models via coordinate descent*. *Journal of Statistical Software*, 33(1). doi:10.18637/jss.v033.i01

See Also

[print.APML0](#)

Examples

```
### Linear model ###
set.seed(1213)
N=100; p=30; p1=5
x=matrix(rnorm(N*p), N, p)
beta=rnorm(p1)
xb=x[, 1:p1] %%% beta
y=rnorm(N, xb)

fiti=APML0(x, y, penalty="Lasso", nlambda=10)
fiti2=APML0(x, y, penalty="Lasso", nlambda=10, nfolds=10)

### Logistic model ###
set.seed(1213)
N=100; p=30; p1=5
x=matrix(rnorm(N*p), N, p)
beta=rnorm(p1)
xb=x[, 1:p1] %%% beta
y=rbinom(n=N, size=1, prob=1.0/(1.0+exp(-xb)))

fiti=APML0(x, y, family="binomial", penalty="Lasso", nlambda=10)
```

```

fiti2=APML0(x, y, family="binomial", penalty="Lasso", nlambda=10, nfolds=10)

### Cox model ###
set.seed(1213)
N=100; p=30; p1=5
x=matrix(rnorm(N*p), N, p)
beta=rnorm(p1)
xb=x[, 1:p1] %*% beta
ty=rexp(N, exp(xb))
td=rexp(N, 0.05)
tcens=ifelse(td < ty, 1, 0)
y=cbind(time=ty, status=1-tcens)

fiti=APML0(x, y, family="cox", penalty="Lasso", nlambda=10)
fiti2=APML0(x, y, family="cox", penalty="Lasso", nlambda=10, nfolds=10)

```

```
print.APML0
```

Print an APML0 Object

Description

Print a summary of results along the regularization path.

Usage

```
## S3 method for class 'APML0'
print(x, digits = 4, ...)
```

Arguments

<code>x</code>	a fitted APML0 object.
<code>digits</code>	number of significant digits in the output. Default is 4.
<code>...</code>	additional arguments (currently unused).

Details

Prints the model family and penalty type, the fit table along the lambda path, and, when cross-validation with L0 selection was performed, the fit0 table for the selected solution.

Value

Called for its side effect (printing). Returns `x` invisibly.

Author(s)

Xiang Li, Shanghong Xie, Donglin Zeng and Yuanjia Wang
 Maintainer: Xiang Li <spiritcoke@gmail.com>

See Also[APML0](#)**Examples**

```
set.seed(1213)
N=100; p=30; p1=5
x=matrix(rnorm(N*p), N, p)
beta=rnorm(p1)
xb=x[, 1:p1] %*% beta
y=rnorm(N, xb)
fiti=APML0(x, y, penalty="Lasso", nlambda=10, nfolds=10)
print(fiti)
```

Index

- * **Hard-thresholding**
 - APML0, 3
 - * **L0**
 - APML0, 3
 - APML0-package, 2
 - * **Package**
 - APML0-package, 2
 - * **Regularization**
 - APML0, 3
 - APML0-package, 2
 - * **print**
 - print.APML0, 7
- APML0, 2, 3, 8
- APML0-package, 2
- print.APML0, 2, 6, 7