

PICASSO 2.0.0 User Guide

PICASSO authors

July 2026

Abstract

PICASSO fits sparse regularization paths with lasso, MCP, and SCAD penalties. This guide summarizes the current R interface, solver choices, accuracy modes, diagnostics, prediction, assessment, and cross-validation.

1 Supported models

family	response	native optimization
gaussian	finite numeric	active-set coordinate descent
binomial	exactly two levels; factor levels are retained	active-set Proximal Newton/IRLS
poisson	non-negative integer counts; at least one positive	active-set Proximal Newton/IRLS
sqrtilasso	finite numeric	active-set quadratic-MM updates
multinomial	at least three observed classes	class-coupled active-set Proximal Newton/IRLS

Every family supports `method = "l1"`, `"mcp"`, or `"scad"`. MCP requires `gamma > 1` and SCAD requires `gamma > 2`. Generated paths request up to 100 decreasing lambda values by default, with a nominal endpoint at 0.05 times the null-model maximum. The returned prefix may be shorter after normal path stopping. An explicit path must be finite, non-negative, and strictly decreasing.

2 Basic fitting

```
library(picasso)
set.seed(1)
n <- 120; d <- 30
X <- matrix(rnorm(n * d), n, d)
y <- 0.5 + X[, 1] - 0.7 * X[, 2] + rnorm(n)

fit <- picasso(X, y, family = "gaussian", method = "l1",
              nlambdas = 30)
fit$type.gaussian.requested
fit$type.gaussian
coef(fit, lambda.idx = fit$nlambdas, beta.idx = 1:5)
predict(fit, X[1:5, ], lambda.idx = fit$nlambdas)
```

R uses one-based `lambda.idx` and `beta.idx` values. Prediction by lambda value is also available through `s`; coefficients are linearly interpolated inside the fitted path and clamped at its endpoints.

3 Gaussian backend selection

The public default is `type.gaussian = "auto"`. It uses covariance updates only when the path has at least eight lambdas, the design is not wide, and at most 1024 features are present. Within those guards, covariance updates are selected for at least one of the following: at most 160 features, a requested/generated path endpoint ratio of at least 0.10, a ratio of at least 0.05 with $n \geq 4d$, or $n \geq 16d$. The feature guard caps a fully populated numeric Gram cache at 8 MiB.

Set `type.gaussian = "naive"` or `"covariance"` to pin a backend. Fit objects record both the requested and resolved modes. Cross-validation reuses the full-data decision in every fold.

4 Adaptive LLA and accuracy

For MCP and SCAD binomial, Poisson, square-root-lasso, and multinomial fits, PICASSO solves a sequence of weighted-L1 subproblems using adaptive local linear approximation (LLA). The minimum and default maximum are three total stages: one L1 master and two weighted-L1 updates. A larger `lla.max.stages` allows additional stages; after the minimum three, the solver stops when target stationarity is no larger than `prec`. Gaussian MCP/SCAD uses its direct coordinate-descent implementation instead.

`fast.mode = FALSE` is the default; `prec` then defaults to `1e-7` and may be set to another positive value. `fast.mode = TRUE` uses benchmark-calibrated achieved-accuracy tolerances: `4e-4` for Poisson, `1e-4` for binomial, square-root-lasso, and multinomial, and `1e-7` for Gaussian. These values are not literal copies of another solver's stopping parameter. Custom `prec` values are available only with fast mode disabled.

5 Offsets, standardization, and path stopping

Binomial and Poisson fits accept one offset per observation:

```
exposure <- runif(n, 0.5, 2)
off <- log(exposure)
count <- rpois(n, exp(off + 0.2 + X[, 1]))
pfit <- picasso(X, count, family = "poisson", offset = off)
mu <- predict(pfit, X[1:5, ], lambda.idx = pfit$nlambda,
              newoffset = off[1:5])
```

Offsets are added on the link scale. Response/link prediction, assessment, and binomial class prediction from an offset-fitted model require a matching `newoffset`; support extraction does not. With `standardize = TRUE`, predictors are centered only when an intercept is included; no-intercept models preserve the original zero point.

`dfmax` stops a path when its nonzero count becomes too large. Generated multinomial paths may also omit a saturated tail after at least five fitted lambdas. Explicit multinomial paths

disable only this saturation rule; they may still return a prefix after `dfmax` stopping or a hard failure.

6 Diagnostics and usable termination

Non-Gaussian fit objects include `status`, `status.code`, `failure`, and per-lambda `diagnostics`. Diagnostics report native runtime, objective, weighted-L1 KKT residual, and target stationarity. Scalar families also report completed LLA stages; multinomial fits report outer iterations, inner sweeps, and coordinate updates.

Status codes 0 (completed), 1 (`dfmax_reached`), and 10 (`lla_stationarity_limit`) return usable models. Code 10 means the stage budget ended before stationarity certification; it is not a hard failure. A later hard failure retains only the committed path prefix and records the failed lambda and stage.

7 Assessment and cross-validation

```
metrics <- assess.picasso(fit, X, y)
cv <- cv.picasso(X, y, family = "gaussian", nfolds = 5,
                 nlambda = 30)
cv$lambda.min
coef(cv, s = "lambda.1se", beta.idx = 1:5)
```

Assessment reports family-appropriate deviance, MSE, MAE, or class error. Confusion matrices are available for binomial and multinomial fits. Categorical CV folds are stratified by default and retain the full-data class map. Normal stopping may shorten the initial full-data generated path. That fixed retained path is then passed explicitly to every fold, and a truncated or otherwise unusable fold is reported as an error. If a very deep scalar path triggers such saturation, request fewer lambdas or a less aggressive endpoint.

8 Further information

The package help page `?picasso` is the authoritative argument and return value reference. Reproducible R and Python examples are in the repository's `tutorials` directory. The [original PICASSO paper](#) describes the pathwise active-set framework; newer interfaces documented here were added after that publication.