Package 'MR.RGM'

November 7, 2026

Title Bayesian Multivariable Bidirectional Mendelian Randomization

Version 0.1.3

Description Fits multivariable bidirectional Mendelian randomization networks with Bayesian directed cyclic graphical models (reciprocal graphical models; Ni, Ji and Mueller (2018) <doi:10.1214/17-BA1087>), an earlier version of which is described in Sarkar and Ni (2025) <doi:10.1093/bioinformatics/btaf130>. The causal effects among the responses, which may form directed cycles, are estimated with spike-and-slab or threshold priors, and posterior probabilities of network motifs are provided. With a known instrument-response map the instrument effects obey the exclusion restrictions; with an unrestricted map they are selected by a spike-and-slab prior, which allows invalid instruments (horizontal pleiotropy). The input can be individual-level data, the second-moment matrices of one sample, or marginal instrument-response regressions, with optional measured covariates. The error covariance can be diagonal, inverse-Wishart, or a stochastic search structure learning prior (Wang (2015) <doi:10.1214/14-BA916>) whose inclusion indicators describe unmeasured confounding among the responses. Inference is by Markov chain Monte Carlo with exact Gibbs updates of the instrument effects and adaptive random-walk Metropolis updates of the causal effects (Roberts and Rosenthal (2009) <doi:10.1198/jcgs.2009.06134>); the fitting function reports diagnostics of weak instruments, weak identification and convergence.

License GPL (>= 3)

Encoding UTF-8

URL https://github.com/bitansa/MR.RGM

BugReports https://github.com/bitansa/MR.RGM/issues

Imports Rcpp, stats, utils, igraph, GIGrvg

Suggests testthat (>= 3.1.5)

LinkingTo Rcpp, RcppArmadillo, GIGrvg

Config/testthat/edition 3

Config/roxygen2/version 8.1.0

NeedsCompilation yes

Author Bitan Sarkar [aut, cre],
Yang Ni [aut]

Maintainer Bitan Sarkar <bitansarkar010899@gmail.com>

Repository CRAN

Date/Publication 2026-11-07 22:20:29 UTC

Contents

NetworkMotif 2
RGM 3

Index **11**

`NetworkMotif` *Posterior probability of a network motif*

Description

`NetworkMotif` quantifies the uncertainty of a given network structure among the response variables: it returns the proportion of posterior draws in which every edge of the motif is present, that is, the posterior probability that all edges of the motif are present simultaneously. The input draws are the `GammaPst` output of `RGM`.

Usage

```
NetworkMotif(Gamma, GammaPst)
```

Arguments

`Gamma`	A p * p binary matrix (numeric 0/1 or logical) with a zero diagonal, specifying the motif: `Gamma[i, j] = 1` for the directed edge from response j to response i (the same orientation as `A` and `GammaEst` in `RGM`). Entries equal to 0 are ignored, so the motif can be any subgraph (a single edge, a cascade, a feed-forward loop, a feedback loop, ...); only the presence of the specified edges is required, not the absence of others.
`GammaPst`	A p * p * n_pst array of posterior draws of the causal edge indicators, the `GammaPst` component returned by `RGM`.

Value

A number between 0 and 1: the proportion of posterior draws that contain all edges of `Gamma`.

References

Ni, Y., Ji, Y., & Müller, P. (2018). Reciprocal graphical models for integrative gene regulatory network analysis. *Bayesian Analysis*, **13(4)**, 1095-1110. doi:10.1214/17BA1087.

Examples

```
# Simulate a cascade Y1 -> Y2 -> Y3 with two own instruments per response
set.seed(9154)
n = 2000; p = 3; k = 6
A = matrix(0, p, p); A[2, 1] = 0.3; A[3, 2] = -0.3     # A[i, j]: effect of Y_j on Y_i
D = matrix(0, p, k)
for (i in 1:p) D[i, (2 * i - 1):(2 * i)] = 1
X = matrix(rnorm(n * k), n, k)
Y = (X %*% t(D) + matrix(rnorm(n * p), n, p)) %*% t(solve(diag(p) - A))
X = scale(X); Y = scale(Y)                              # the model has no intercept

# Fit MR.RGM (a short chain to keep the example fast; the defaults are the recommended length)
fit = RGM(X = X, Y = Y, D = D, nIter = 2000, nBurnin = 1000, Thin = 2)

# Posterior probability of the cascade Y1 -> Y2 -> Y3 (both edges present)
Gamma = matrix(0, p, p); Gamma[2, 1] = 1; Gamma[3, 2] = 1
NetworkMotif(Gamma = Gamma, GammaPst = fit$GammaPst)

# Posterior probability of the feedback loop Y1 -> Y2 -> Y3 -> Y1
Gamma[1, 3] = 1
NetworkMotif(Gamma = Gamma, GammaPst = fit$GammaPst)
```

RGM *Fitting Multivariate Bidirectional Mendelian Randomization Networks Using Bayesian Directed Cyclic Graphical Models*

Description

RGM fits the Bayesian directed cyclic graphical model

$$Y = AY + BX + CU + E, \qquad E \sim N_p(0, \Sigma^*),$$

where Y are p response variables (traits), X are k instrumental variables (for example genetic variants), U are l optional measured covariates, A (p x p, zero diagonal) contains the direct causal effects among the responses (A[i, j] is the effect of Y_j on Y_i; directed cycles are allowed), B (p x k) the instrument effects, C (p x l) the covariate effects and Σ^* the error covariance, whose off-diagonal entries represent unmeasured confounding between responses.

Two versions of the method differ only in how the instruments are mapped to the responses (in the accompanying methods paper they are called IV-RGM and IV-RGM+, since the model applies to instrumental variables in general):

MR.RGM (known instrument map) D is a binary p x k matrix with D[i, j] = 1 if instrument j is allowed to affect response i (for example a cis-eQTL for its own gene). Entries of B with D = 0 are fixed at zero (the exclusion restriction) and the allowed entries receive a slab prior N(0, eta) without any selection: every allowed instrument stays in the model. This is the default when D contains zeros (selectB = FALSE).

MR.RGM+ (instrument selection) `D` is a matrix of ones: every instrument may affect every response, and each entry of `B` receives a spike-and-slab prior whose inclusion indicator is sampled, so that the instrument–response map is learned from the data (horizontal pleiotropy). This is the default when `D` has no zeros (`selectB = TRUE`).

The argument `selectB` overrides these defaults (see its description).

Inference is by MCMC on the sufficient statistics of the Gaussian model, so the function accepts individual-level data (`X`, `Y`, optionally `U`), the second-moment matrices (`Syy`, `Syx`, `Sxx`, optionally `Syu`, `Sxu`, `Suu`) computed from one sample, or the matrices `Sxx`, `Beta` and `SigmaHat` of marginal instrument–response regressions; responses alone (`Y` or `Syy`) fit the model without instruments. The marginal-regression input does not contain the covariances between the responses: they are reconstructed under independent errors (no unmeasured confounding), so this input requires `SigmaStarModel = "diagonal"`, no covariates and, for every response, an instrument that affects only that response, and it is exact only for uncorrelated instruments. Use individual-level data or the second-moment matrices to model confounding. The model has no intercept and the stochastic-search prior on Σ^* assumes responses on roughly unit scale: center and scale the columns of `Y`, `X` and `U` (or compute the summary matrices from centered and scaled data) before fitting.

The output contains posterior means of all parameters, posterior inclusion probabilities of the causal edges (`GammaEst`), of the instrument effects (`PhiEst`, MR.RGM+ only) and of the confounding links (`ZEst`, SSSL only), the graph estimated by thresholding the inclusion probabilities at 0.5, the posterior draws of the causal edge indicators (`GammaPst`, used by `NetworkMotif`), the log-likelihood trace, acceptance rates and convergence/identification diagnostics. Warnings are issued for signs of weak identification or poor convergence (see `diagnostics` in the Value section and the Note).

MCMC algorithm. Every iteration updates, in turn, the instrument effects `B`, the covariate effects `C`, the error covariance Σ^* (with its confounding indicators) and the causal effects `A`, each with its prior hyperparameters. Given the other parameters, the likelihood is quadratic in `B` and its (spike-and-slab or slab) prior is normal given the inclusion indicators, so each row of `B` is drawn exactly from its multivariate normal full conditional (Gibbs; `gibbsB = TRUE`). The entries of `A` enter the likelihood also through $\log|\det(I - A)|$ and are updated one at a time by random-walk Metropolis–Hastings with adaptive proposals: each entry has its own proposal variance, started at 2.38^2 times an approximation of its posterior variance and tuned during the burn-in towards an acceptance rate of 0.44 (adaptive Metropolis-within-Gibbs of Roberts and Rosenthal, 2009, with larger adaptation steps early in the burn-in); the variances are frozen after the burn-in, so the kept draws come from a fixed Markov kernel. With more than 100 allowed instruments for a response, its row of `B` is drawn entry by entry (each entry exactly from its scalar full conditional) instead of jointly. The same adaptive scheme updates `B` when it cannot be drawn by Gibbs (Threshold prior with selection, or `gibbsB = FALSE`). The chain starts at `A = 0`, at the least-squares effects of each response on its allowed instruments (known map; `B = 0` with selection) and at the corresponding residual covariance.

Chain length and convergence. One call runs one chain. The defaults (50,000 iterations, 10,000 burn-in, thinning 10) are the settings validated in the simulation study of the accompanying paper for networks of up to about 20 responses; larger networks or weak instruments may need longer chains. For a reported analysis, run several chains (for example four) with different seeds (`set.seed()` before each call), check their agreement (for example the rank-normalized split-$\hat{R}$ of `LLPst` and of the edge indicators in `GammaPst`, computed with the package **posterior**, and the posterior inclusion probabilities of the chains), and pool the chains only if they agree. The function also reports diagnostics and warnings of its own (see `diagnostics` in the value).

Usage

```
RGM(
  X = NULL,
  Y = NULL,
  U = NULL,
  Syy = NULL,
  Syx = NULL,
  Sxx = NULL,
  Syu = NULL,
  Sxu = NULL,
  Suu = NULL,
  Beta = NULL,
  SigmaHat = NULL,
  D,
  n,
  nIter = 50000,
  nBurnin = 10000,
  Thin = 10,
  prior = c("Spike and Slab", "Threshold"),
  SigmaStarModel = c("SSSL", "IW", "diagonal"),
  aRho = 3,
  bRho = 1,
  nu1 = 0.001,
  aPsi = 0.5,
  bPsi = 0.5,
  nu2 = 1e-04,
  aSigma = 0.01,
  bSigma = 0.01,
  PropVarA = NULL,
  PropVarB = NULL,
  gibbsB = TRUE,
  selectB = NULL,
  tauC = 2,
  omega1 = 5,
  omega2 = 0.005,
  varpi = 0.5,
  lambda = 1
)
```

Arguments

X	A numeric matrix of dimension n * k of instrumental variables (rows = observations). Default NULL.
Y	A numeric matrix of dimension n * p of response variables (rows = observations). Default NULL.
U	A numeric matrix of dimension n * l of measured covariates (rows = observations). Default NULL.

`Syy` A p * p matrix, `t(Y) %*% Y / n`, computed from centered (and preferably scaled) responses. Default NULL.

`Syx` A p * k matrix, `t(Y) %*% X / n`. Default NULL.

`Sxx` A k * k matrix, `t(X) %*% X / n`. Default NULL.

`Syu` A p * l matrix, `t(Y) %*% U / n`. Default NULL.

`Sxu` A k * l matrix, `t(X) %*% U / n`. Required together with `Syu` and `Suu` when instruments and covariates are both given as summary matrices. Default NULL.

`Suu` An l * l matrix, `t(U) %*% U / n`. Default NULL.

`Beta` A p * k matrix of regression coefficients of each response on each instrument separately (marginal regressions on centered data). Used with `Sxx` and `SigmaHat`; this input treats the marginal effects as joint effects, which is exact for uncorrelated instruments, and reconstructs the covariances between the responses under independent errors. It requires `SigmaStarModel = "diagonal"`, no covariates, and a map `D` in which every response has an instrument that affects only that response (so it is not available for MR.RGM+ with `D` all ones). Default NULL.

`SigmaHat` A p * k matrix of the residual variances of the same marginal regressions: the residual sum of squares divided by n (not by the residual degrees of freedom). Default NULL.

`D` A binary p * k matrix (numeric 0/1 or logical): `D[i, j] = 1` if instrument j is allowed to affect response i. A matrix of ones gives MR.RGM+; a matrix with zeros gives MR.RGM (see Description). Required whenever instruments are given; not used without instruments.

`n` A positive integer, the number of observations. Required when summary-level matrices are the input; with individual-level data it is `nrow(Y)`.

`nIter` A positive integer, the number of MCMC iterations. Default 50,000 (the setting of the simulation study of the accompanying paper).

`nBurnin` A non-negative integer smaller than `nIter`, the number of burn-in iterations. Default 10,000. The proposal variances of the Metropolis updates are tuned during the burn-in (in batches of 50 iterations); a burn-in of a few thousand iterations or more is recommended. With the spike-and-slab prior and a known instrument map (`D` with zeros, `selectB = FALSE`), the inclusion indicators of `A` are held at 1 during the first half of the burn-in, so that every causal effect first moves under its slab prior before selection starts.

`Thin` A positive integer, the thinning interval of the kept draws (draws are kept at iterations `nBurnin + Thin, nBurnin + 2 Thin, ...`). It must not exceed `nIter - nBurnin`. Default 10 (4,000 kept draws with the default chain length).

`prior` The prior on the graph structure: `"Spike and Slab"` (default) or `"Threshold"`.

`SigmaStarModel` The model for the error covariance Σ^*: `"SSSL"` (default; stochastic search structure learning prior with correlated errors and inclusion indicators for the covariances), `"IW"` (inverse-Wishart prior $IW(p+1, I_p)$, correlated errors) or `"diagonal"` (independent errors with inverse-gamma priors $IG(aSigma, bSigma)$ on the variances).

`aRho, bRho` Positive shape parameters of the Beta prior on the inclusion probability of a causal edge (spike-and-slab prior on `A`). Defaults 3 and 1 (prior inclusion probability 0.75).

nu1 — A positive scalar, the spike variance of the spike-and-slab prior on A relative to the slab variance (spike-and-slab prior only). Default 0.001.

aPsi, bPsi — Positive shape parameters of the Beta prior on the inclusion probability of an instrument effect (spike-and-slab prior with selection of the instrument effects, e.g. MR.RGM+, only). Defaults 0.5 and 0.5.

nu2 — A positive scalar, the spike variance of the spike-and-slab prior on B relative to the slab variance (spike-and-slab prior with selection, e.g. MR.RGM+, only). Default 0.0001.

aSigma, bSigma — Positive shape and rate parameters of the inverse-gamma prior on the error variances when SigmaStarModel = "diagonal". Defaults 0.01 and 0.01.

PropVarA, PropVarB — Starting variances of the adaptive random-walk Metropolis–Hastings proposals for the entries of A and (when B is not drawn by Gibbs) B; the proposals are tuned during the burn-in. NULL (default): an automatic start, $2.38^2/(n\Omega_{ii}S_{jj}+1)$ for entry (i, j), where Ω is the inverse of the starting error covariance and S is Syy (for A) or Sxx (for B); it scales with the sample size. A positive number is used as the starting variance of every entry.

gibbsB — Logical. TRUE (default): the instrument effects B are drawn exactly from their normal full conditional (Gibbs), row by row; with the Threshold prior this applies when the instrument effects are not selected (selectB = FALSE). FALSE: adaptive random-walk Metropolis–Hastings for every entry of B.

selectB — Logical or NULL. Whether the instrument effects in B are selected by a spike-and-slab (or threshold) prior. NULL (default) means all(D == 1): selection for MR.RGM+ (D all ones) and no selection for MR.RGM (D with zeros). TRUE with a D that contains zeros selects among the allowed entries only; FALSE with D all ones gives every entry a slab prior without selection. The value used is returned in settings.

tauC — A positive scalar, the column variance of the matrix-normal prior C ~ MN(0, SigmaStar, tauC I_l) on the covariate effects. Default 2.

omega1, omega2 — Positive slab and spike variances of the SSSL prior on the off-diagonal entries of Σ^* (SigmaStarModel = "SSSL"). omega1 must exceed omega2. Defaults 5 and 0.005; they assume responses on unit scale.

varpi — The prior inclusion probability of a confounding link under the SSSL prior, in (0, 1). Default 0.5.

lambda — A positive scalar; under the SSSL prior the diagonal entries of Σ^* have an exponential prior with rate lambda / 2 (the parametrization of Wang, 2015). Default 1.

Value

A list with the following components (components that do not apply to the chosen prior or model are omitted):

AEst — p * p matrix, posterior mean of A (AEst[i, j] is the estimated direct effect of Y_j on Y_i).

BEst — p * k matrix, posterior mean of B.

`CEst` p * l matrix, posterior mean of `C` (with covariates).

`zAEst` p * p binary matrix, `GammaEst > 0.5` (median probability model).

`zBEst` p * k binary matrix, `PhiEst > 0.5`; equal to `D` without selection.

`Graph` An `igraph` object of the estimated causal network among the responses (edges of `zAEst`, labelled by `AEst`).

`A0Est`, `B0Est` Posterior means of the unthresholded `A` and `B` (Threshold prior).

`GammaEst` p * p matrix of posterior inclusion probabilities of the causal edges (`GammaEst[i, j]` for the edge `Y_j -> Y_i`).

`TauEst` p * p matrix, posterior mean of the slab variances of `A`.

`RhoEst` p * p matrix, posterior mean of the edge inclusion probabilities `rho` (spike-and-slab prior).

`PhiEst` p * k matrix of posterior inclusion probabilities of the instrument effects when they are selected (MR.RGM+); equal to `D` otherwise.

`EtaEst` p * k matrix, posterior mean of the slab variances of `B`.

`PsiEst` p * k matrix, posterior mean of the inclusion probabilities `psi` of the instrument effects (spike-and-slab prior with selection only).

`tAEst`, `tBEst` Posterior means of the thresholds (Threshold prior; `tBEst` with selection only).

`SigmaEst` The estimated error covariance: a p * p matrix for `"SSSL"` and `"IW"`, a vector of length p of error variances for `"diagonal"`.

`ZEst` p * p matrix of posterior inclusion probabilities of the confounding links (`"SSSL"` only; the diagonal is 1).

`AccptA`, `AccptB` Acceptance rates (in percent, over the iterations after burn-in) of the Metropolis–Hastings updates of `A` and `B`; `AccptB` is omitted when `B` is drawn by Gibbs.

`AccpttA`, `AccpttB` Acceptance rates (in percent, after burn-in) of the threshold updates (Threshold prior; `AccpttB` only with selection of the instrument effects).

`PropVarA`, `PropVarB` The tuned (frozen) proposal variances of the entries of `A` (p * p, zero diagonal) and, when `B` is updated by Metropolis–Hastings, of `B` (p * k, zero where `D = 0`).

`LLPst` Numeric vector of the log-likelihood at the kept draws.

`GammaPst` p * p * nPst array of the kept draws of the causal edge indicators (input of `NetworkMotif`).

`diagnostics` A list of convergence and identification diagnostics, each with a warning when it signals a problem: `instrumentF` (only without instrument selection, `selectB = FALSE`: first-stage F statistic of each response on its allowed instruments, computed from the second-moment matrices without adjusting for the covariates, `NA` where it cannot be computed; warning below 10), `spectralRadiusA` (spectral radius of `AEst`; warning at 1 or above, where the estimated feedback loops do not converge and the equilibrium interpretation fails), `halvesMaxDiff` (largest difference of the causal inclusion probabilities between the two halves of the kept draws; with at least 200 kept draws, warning when, for at least one edge, the difference exceeds 0.2 and four times that edge's batch-means Monte Carlo standard error; returned only with at least 200 kept draws; a heuristic, which can

also fire for edges that mix slowly), `halvesMCSE` (the Monte Carlo standard error of that largest difference), `acceptanceA` (= `AccptA`; warning below 10 percent) and `errorVarianceRatio` (ratio of each estimated error variance to the second moment of its response; warning above 5).

`settings` A list of the settings used: `input` (`"Y"`, `"XY"`, `"S"` or `"Beta"`), `prior`, `SigmaStarModel`, `instruments`, `covariates`, `selectB`, `gibbsB` (whether `B` was drawn by Gibbs), `n`, `p`, `k` (number of instruments), `l` (number of covariates), `nIter`, `nBurnin`, `Thin`, the hyperparameters, the starting proposal variances (`"auto"` for the automatic start) and the package version. The Threshold prior also uses a fixed upper bound of 0.1 on the thresholds and a truncated-normal proposal with standard deviation 0.1 for them.

Note

Weak identification. With `D = 1` (MR.RGM+), the instrument effects are identified only through sparsity and an instrument-diversity condition, which require instruments that explain a non-negligible share of the variance of several responses. With weak, response-specific instruments (for example cis-eQTLs) the posterior can stay close to the prior; in that case `RGM` warns (error variances far above the second moments of the responses) and the instrument map `D` should encode which instrument affects which response (MR.RGM). Compare the posterior inclusion probabilities with their prior values (0.75 for causal edges under the default `aRho = 3, bRho = 1`; 0.5 for confounding links) and the causal inclusion probabilities computed from the two halves of `GammaPst`.

Identification of MR.RGM. With a known map, the causal effects are identified when, for every response, the instruments excluded from its equation satisfy the classical rank condition; a sufficient condition is that every response has at least one instrument that affects only that response. `RGM` warns when this sufficient condition fails and no selection is performed.

References

Roberts, G. O., & Rosenthal, J. S. (2009). Examples of adaptive MCMC. *Journal of Computational and Graphical Statistics*, **18(2)**, 349-367. doi:10.1198/jcgs.2009.06134.

Wang, H. (2015). Scaling it up: stochastic search structure learning in graphical models. *Bayesian Analysis*, **10(2)**, 351-377. doi:10.1214/14BA916.

Ni, Y., Ji, Y., & Müller, P. (2018). Reciprocal graphical models for integrative gene regulatory network analysis. *Bayesian Analysis*, **13(4)**, 1095-1110. doi:10.1214/17BA1087.

Sarkar, B. & Ni, Y. (2025). MR.RGM: an R package for fitting Bayesian multivariate bidirectional Mendelian randomization networks. *Bioinformatics*, **41(4)**, btaf130. doi:10.1093/bioinformatics/btaf130.

Examples

```
# ------------------------------------------------------------------------------------------
# Simulated data: Y1 -> Y2 -> Y3 and Y3 -> Y1 (a feedback loop), two own instruments per
# response, one covariate, and an unmeasured confounder shared by Y1 and Y3
set.seed(9154)
n = 2000; p = 3; k = 6; l = 1
A = matrix(0, p, p); A[2, 1] = 0.3; A[3, 2] = -0.3; A[1, 3] = 0.2  # A[i, j]: effect of Y_j on Y_i
D = matrix(0, p, k)                                                 # known instrument map
```

```
for (i in 1:p) D[i, (2 * i - 1):(2 * i)] = 1
X = matrix(rnorm(n * k), n, k); U = matrix(rnorm(n * l), n, l); W = rnorm(n)
E = cbind(W, 0, W) + matrix(rnorm(n * p), n, p)                  # confounding of Y1 and Y3
Y = (X %*% t(D) + U %*% matrix(0.5, l, p) + E) %*% t(solve(diag(p) - A))
colnames(Y) = paste0("Y", 1:p); colnames(X) = paste0("SNP", 1:k); colnames(U) = "age"
X = scale(X); Y = scale(Y); U = scale(U)                       # the model has no intercept

# Example 1: MR.RGM (known map D) from individual-level data, spike-and-slab prior and
# SSSL error covariance (the defaults); a short chain to keep the example fast, whereas the
# default chain length is the recommended one
fit = RGM(X = X, Y = Y, U = U, D = D, nIter = 2000, nBurnin = 1000, Thin = 2)
round(fit$GammaEst, 2)   # posterior inclusion probabilities of the causal edges
round(fit$AEst, 2)       # posterior means of the causal effects
round(fit$ZEst, 2)       # posterior inclusion probabilities of the confounding links
fit$diagnostics          # weak-instrument, stability and convergence diagnostics
plot(fit$Graph, main = "Estimated causal network")
plot(fit$LLPst, type = "l", xlab = "Kept draw", ylab = "Log-likelihood")

# Example 2: the same model from the second-moment matrices of the same sample
fit_s = RGM(Syy = crossprod(Y) / n, Syx = crossprod(Y, X) / n, Sxx = crossprod(X) / n,
            Syu = crossprod(Y, U) / n, Sxu = crossprod(X, U) / n, Suu = crossprod(U) / n,
            D = D, n = n, nIter = 2000, nBurnin = 1000, Thin = 2)
round(fit_s$GammaEst, 2)

# Example 3: MR.RGM+ (D = 1): every instrument may affect every response and the instrument
# effects are selected by a spike-and-slab prior (horizontal pleiotropy)
fit_plus = RGM(X = X, Y = Y, U = U, D = matrix(1, p, k), nIter = 2000, nBurnin = 1000, Thin = 2)
round(fit_plus$PhiEst, 2)   # posterior inclusion probabilities of the instrument effects
fit_plus$settings$selectB   # TRUE

# Example 4: threshold prior on A and a diagonal error covariance
fit_t = RGM(X = X, Y = Y, D = D, prior = "Threshold", SigmaStarModel = "diagonal",
            nIter = 2000, nBurnin = 1000, Thin = 2)
fit_t$zAEst

# Example 5: marginal-regression input (Sxx, Beta, SigmaHat), with uncorrelated instruments,
# no covariates and independent errors (SigmaStarModel = "diagonal" is required)
Y0 = scale((X %*% t(D) + matrix(rnorm(n * p), n, p)) %*% t(solve(diag(p) - A)))
Beta = SigmaHat = matrix(0, p, k)
for (i in 1:p) for (j in 1:k) {
  b = sum(X[, j] * Y0[, i]) / sum(X[, j]^2)                    # centered data, no intercept
  Beta[i, j] = b; SigmaHat[i, j] = sum((Y0[, i] - b * X[, j])^2) / n  # residual sum of squares / n
}
fit_b = RGM(Sxx = crossprod(X) / n, Beta = Beta, SigmaHat = SigmaHat, D = D, n = n,
            SigmaStarModel = "diagonal", nIter = 2000, nBurnin = 1000, Thin = 2)
round(fit_b$GammaEst, 2)
```

Index

NetworkMotif, 2, *4*, *8*

RGM, *2*, 3